The Pentium® II Xeon™ Processor Server Platform System Management Guide

June 1998
Intel Corporation

Order Number: 243835-001

Table Of Content

List Of Figures

List Of Tables

Reference Documents

Additional information on the topics discussed here may be obtained through:

- *Pentium® II Xeon™ Processor at 400MHz*, Order Number 243770-001.
- *Wired for Management*, Intel Corporation. **http://www.intel.com/managedpc**
- *Intelligent Platform Management Interface Specification* , 1997, Intel Corporation. **http://developer.intel.com/design/servers/ipmi/index.htm**
- *Web-Based Enterprise Management (WBEM).* **http://www.microsoft.com**, **http://www.freerange.com**
- *Desktop Management Task Force*, *DMI* Rev.2.0. **http://www.dmtf.org**
- *The I^2C Bus and How to Use It*, January 1992, Signetics/ Philips Semiconductor.
- *System Management Bus Specification* , Intel Corporation. **http://developer.intel.com/ial/powermgm/specs.htm**

1. System Management Overview

This paper establishes a common background in system management by providing a simplified view of the management goals along with software and hardware elements involved in implementation. This discussion is followed by outlining the Pentium® II Xeon™ processor's manageability features and its benefits in various aspects of manageability.

The proliferation of networked client and server platforms has proven effective for business. At the same time, increasing administrative costs and ongoing expenses incurred in deploying and managing systems has become a major burden on IT departments. In a large network environment, IT departments are challenged to accurately maintain and track network components, implement mass upgrade plans, eliminate unreliable equipment, prevent failures, and provide support, service and repair. To properly address these issues, IT managers need access to simple tools for gathering detailed data from departmental networked servers and client workstations. Such databases should be detailed enough to enable IT managers to locate installed hardware and software on the network, detect unreliable hardware, detect stressed operating environments, and ensure intrusion-free networks.

The computing industry's first step towards deployment of such tools was the introduction and promotion of industry-standard management software initiatives. Such standardization allows independent deployment of management application software that articulates a consistent instrumentation methodology across platforms. The availability of such tools has addressed some IT managers' dilemmas. However, the proprietary nature of the hardware architecture, and the fact that software is tightly coupled to the hardware, has proven costly and inflexible for scalability, portability and extensibility. The effectiveness of such tools relies on the availability of accurate information that benefits instrumentation. This addresses both asset and configuration management, as well as the monitoring hardware implementation, which targets system Reliability, Availability, Serviceability, Usability, and Manageability (RASUM).

The industry's commitment to the creation of open specifications has resulted in proposals to industry open system management hardware implementation. Hardware common platform architectures eliminates the drawbacks of discrete and proprietary implementations, which often translates to higher cost, incoherent implementation methodologies, and erroneous data collection and interpretation. Adoption of such industry open architecture can de-couple software from hardware, allowing independent enhancement. The benefit of industry open-based implementation is collection of consistent and accurate data, even when gathered across heterogeneous networks.

The microprocessors at the heart of server platforms, which integrate the highest transistor counts and function at the highest system frequencies ever, are the critical hardware element for proper and continuous system operation. Most hardware management implementations are designed to monitor and check the system's operating environment, referencing the processor's specification. This real-time monitoring is critical to ensure reliable system operation and detection of failure symptoms.

1.0 Introduction

The Pentium® II Xeon™ processor introduces manageability features that enable enhancements in instrumentation, improving asset management and configuration management, as well as making available detailed information about the processor's characteristics. Such features are key to enabling a fine-tuned hardware monitoring logic implementation. A typical hardware management implementation focuses on verifying that the platform operating environment is within the microprocessor's characterized specification. In this process, the processor's core temperature, voltage, ambient temperature, cache interface and internally generated (IERR) errors are monitored. The real-time monitoring of processor's working environment is critical to ensuring reliable system operation and detection of a failing system's symptoms.

In server platforms with Pentium II Xeon processors, such data is pre-programmed into the cartridge and is accessible through the system management bus (SMBus). The SMBus seamlessly connects the processor's system management components to the baseboard in an industry open format. The Pentium II Xeon processor integrates many of the discrete components required to implement robust processor management, such as the thermal sensor, in a manageable platform, enhancing data accuracy at lower cost. The following management features are incorporated into the Pentium II Xeon processor cartridge:

- SMBus
- Processor Information Read Only Memory (PIROM)
- Scratch EEPROM
- Thermal sensor devices

The information contained within the PIROM makes more data available about the cartridge's characteristics and configuration than preceding Intel processors. System management software can apply this information for generating databases that can then be used for superior instrumentation, inventory control, tracking, and networked server platform configuration. IT instrumentation can be enhanced by displaying detailed and accurate data about each processor cartridge in each server platform, including:

- S-spec Number
- Processor Core Type
- Processor Core Family
- Processor Core Model
- Processor Core Stepping
- Maximum Core Frequency
- Core Voltage requirement
- Core Voltage Tolerance, High
- Core Voltage Tolerance, Low
- L2 Cache Size
- Number of SRAM Components
- L2 Cache Voltage requirement
- L2 Cache Voltage Tolerance, Low
- Cache/Tag Stepping ID
- Cartridge mechanical Revision
- Substrate Revision Software ID
- Processor Part Number
- Processor BOM ID
- 64-bit unique identification number
- Thermal Reference Byte
- Processor Core Feature Flags

This information can also be used to automatically configure processor voltage, temperature, and frequency monitoring to match the specified operating ranges. Products that take advantage of this capability can eliminate manual configuration of these ranges as part of system integration or field upgrades.

2.0 Software Elements of Server Platform Management

Instrumentation is a basic requirement of a well-managed server platform. Instrumentation defines manageable data about hardware and software components on a server platform, and the methods used to make those data available to management applications. Standardized access to management data is provided through software interfaces defined by industry standards such as DMI or SNMP (detailed information about these standards can be obtained by visiting the web sites outlined in the references section). Appendix B outlines a simplified sample of DMI software stack as a quick reference summary. An instrumented server platform with the Pentium® II Xeon™ processor can:

- Display a complete inventory of processor cartridges by part number
- Display a complete inventory of processors BOM ID
- Display a complete inventory of cartridges and substrate revisions
- Display a complete configuration listing of processor cartridges installed in each server platform, by speed and cache sizes
- Display complete server performance data, including processor up-time (run time duration), by maintaining a up-time counter in the Scratch EEPROM
- Provide detailed alerts of pending failures by tracking the number of internal or external errors detected by the processor's Machine check Architecture (MCA)
- Automatically set processor and L2 cache voltages and processor core thermal monitoring envelope
- Display processor's maximum thermal specification
- Display processor's configuration data to allow for proper matching
- Display the dynamic operating status by monitoring processor core thermal information
- Display and correlate installed processor core and cache voltage requirements with the voltage modules installed
- Detect a security breach by tracking each processor by its unique electronic ID
- Display the installed Micro-code ID, storing the data in each processor's Scratch EEPROM.

The instrumentation can then map out manageable data about the platform hardware and software components. This covers all the manageable elements in a server platform, from a component's temperature to silicon stepping, the network configuration topology, and versions of the installed software. In an advanced instrumentation environment, it's feasible to engineer proactive communication on critical events and remote enabling or disabling of a component's functions. Component-level instrumentation consists of maintaining attributes with up-to-the minute values so that adjustments to component's operational characteristics, based on these values, can be managed.

In Pentium II Xeon processor-based server platforms, system management data is available through the SMBus, a subset of the industry-standard I^2C serial bus. The SMBus is powered separately from the processor, allowing this information to be available to a remote management application before system power-up boot. Some parameters, which represent processor's operating environment limitations, are preprogrammed in the PIROM. This data can be used to accurately monitor and manage the processor's operating environment. A scratch EEPROM is also provided on the cartridge substrate, which may be used by OEMs in a proprietary fashion. Possible usage for the scratch EEPROM includes:

- Service information
- Inventory management (e.g., asset tag)
- Processor up-time duration
- Processor Micro-code Update revision
- Security

3.0 Hardware Elements of Server Platform Management

In a server platform, the system management goal of reducing total cost of ownership (TCO) is measured by metrics associated with RASUM:

- **Reliability**: Continuous self-checking and data correction, extensive redundancy for higher reliability
- **Availability**: Up-time, automatic data correction, failure tolerance, error-alerting mechanism, boot resiliency
- **Serviceability**: Fault detection and rapid repair
- **Usability**: Flexibility in customization
- **Manageability**: Availability of detailed information about the platform components and characteristics. Address asset management, inventory tracking, and configuration management

The hardware implementation of server management focuses on providing fault prediction, detection, and resilience. In the case of a component failure, server management's goal is to provide rapid servicing of the problem to minimize server downtime. Rapid server restoration requires remote notification, including the Field Replaceable Unit (FRU) that failed, and, if feasible, deployment of remote mechanisms for reconfiguration and recovery.

The hardware monitoring mechanism's efficiency is critical to ensure reliable system operation. The monitoring logic's efficiency is improved once the true characteristics of the processor are compared to the measured parameters in the system. Prior to the Pentium® II Xeon™ processor, a processor's operating environment was compared to the specifications outlined in the processor's data sheet. Since each processor may have unique characteristics, implementing wide-tolerance monitoring that covers all variations can produce a sub-optimal result.

The Pentium II Xeon processor enhances the accuracy of such monitoring by making data about its characteristics and operating limitations available through the PIROM. This methodology maximizes the system up-time by efficiently matching the processor's operating environment to the specification. The PIROM also enables creation of systems that eliminate the manual configuration of monitoring logic, such as system clock speed and cache sizes, to the management software.

3.1 Common Platform Architecture

The first server management specification aimed at common platform architecture of a managed platform is the "Intelligent Platform Management Interface," or IPMI. Implementing an IPMI architecture is just one example of many implementations that can be used to exploit the unique management features of the Pentium II Xeon processor.

The IPMI specification is comprised of three sub-specifications, which define the interface to the platform hardware (IPMI), the internal intelligent platform management bus (IPMB), and the external bus for connecting additional IPMI-enabled systems. The bus operates autonomously; the critical sensors and events are monitored and logged even if the processor is not operating and system management software is not available. IPMI defines a common interface and message-based protocol for accessing platform management hardware. This reduces TCO by improving server platform management functionality and compatibility while de-coupling the hardware and software. This allows hardware advances without impacting server management software.

The Pentium II Xeon processor interfaces the PRIOM, Scratch EEPROM and thermal sensor to a baseboard BMC through the SMBus interface. The SMBus, which is a two-wire bus interface and a subset of the industry-standard I^2C serial bus, is compatible with the IPMI private management bus interface protocol. This allows a seamless, low-cost connection of the management components. The major building blocks of an IPMI manageable platforms are:

1. **Baseboard Management Controller (BMC):** A micro-controller connected to the platform management components in the system through a private bus, typically an I^2C or SMBus, or other direct connection. The controller operates autonomously, providing automatic monitoring and recovery functions independent of the system processors, system software, or operating system. This controller is typically mapped as an I/O device in the system interfacing with the central processing unit through the ISA bus. The baseboard management controller firmware is either download-able or permanently resident in the BMC.

2. **System Interface Ports:** Typically an I/O-mapped interface to the ISA bus. This provides a standardized set of registers that provides communication between instrumentation software and the platform management hardware.

3. **System Management Software (SMS):** Software executing on the system that interprets management application requests to retrieve management information and set management parameters. In addition to providing the interface to the platform management hardware, system management software typically provides access to management information from BIOS, add-in cards, and the operating system, as well.

4. **Private Management Bus (PMB):** Means by which the system management software communicates with the manageable devices; typically an I^2C or SMBus. Access to the manageable devices is accomplished by establishing command communication with the BMC through the system interface ports. The BMC interprets requests from system management software and performs the requested operation on the targeted Private Management Bus device.

5. **Sensors:** Voltage, temperature and fan speed sensors are typical data-retrieving devices used by software management to properly control the operating environment, and can provide failure prevention and fault detection. In IPMI, access to sensors is abstracted behind a message-based interface that isolates system management software from the hardware. For example, to get a sensor reading, system management software sends a ‘Get Sensor Reading’ command to the management controller, rather than performing a low-level access directly to the monitoring hardware.

Figure 1. outlines an example block diagram of hardware interconnects of a managed server platform

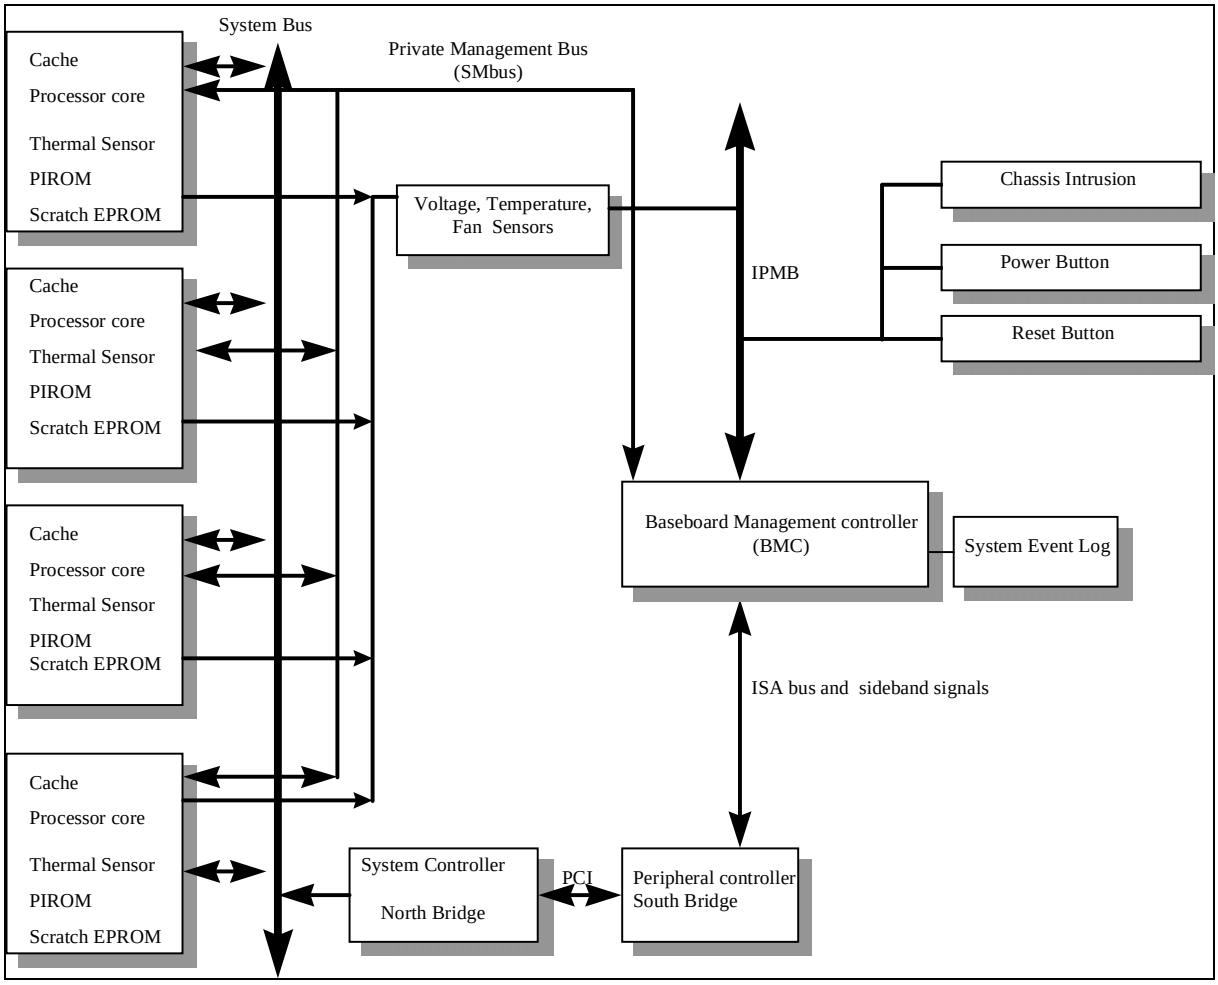

Figure 1. Example Block Diagram of Hardware Interconnects of a Managed Server Platform

4.0 Pentium® II Xeon™ Processor Manageability Features

The Pentium® II Xeon™ processor maintains all the management features offered in previous generations of processors [system management mode (SMM), functional redundancy checking (FRC), low power states and clock control (halt, stop grant and sleep states)], and advances processor manageability by integrating the following management features:

- System Management Bus
- Processor Information ROM
- Scratch EEPROM
- Thermal Sensor Devices

The Pentium II Xeon processor, with initial core frequencies of 400 MHz, has incorporated a system management architectures that lends itself to improve server platform RASUM. Server RASUM is directly impacted by the accuracy of the parameters chosen to monitoring and manage the processors' operating environment. Server platforms based on the Pentium II Xeon processor can enable the creation of management software and firmware that can automatically configure the monitoring hardware based on the parameters obtained from each processor. These parameters are programmed into the PIROM during manufacturing test, closely matching each processor's characterized limitations. This ensures a highly efficient and reliable operating environment.

An IT manager can ensure reliable, improved up-time and decrease time-to-repair by deploying a server platform management that monitors the system and provides alerts once the system parameters begin to drift out of range. Automatic recovery actions can provide information about the cause of the failures and clearly identify the failure unit. Timely and accurate component replacement eliminates the opportunity for the end user to misconfigure the system. This information can be obtained by closely monitoring errors generated internally or on the system bus, monitoring abnormal variations in the core or ambient temperature, and monitoring abnormal variations in the voltage regulator or modules supplying voltages to the processor core and the cache modules.

Integration of management components on to the Pentium II Xeon processor cartridge enhances data accuracy at lower cost. The Pentium II Xeon processor-based server platforms can maintain the operating environment within factory specifications, characterized for the individual processor. This includes:

- Processor core temperature
- Processor voltage requirement
- Processor voltage tolerance
- Processor core frequency
- Cache size
- Cache voltage requirement
- Cache voltage tolerance

The Pentium II Xeon processor enables system management software to utilize the PIROM information for generating databases, which enhances inventory control and tracking and networked server platform configuration. This data includes:

- S-spec/QDF number
- Processor core family
- Processor core model
- Processor core stepping
- Cache/tag stepping ID
- Cartridge revision
- Substrate revision software ID
- Processor part number

This information can also be used to automatically configure processor voltage, temperature, and frequency monitoring to match the specified operating ranges. Products that take advantage of this capability can eliminate manual configuration of these ranges as part of system integration or field upgrades.

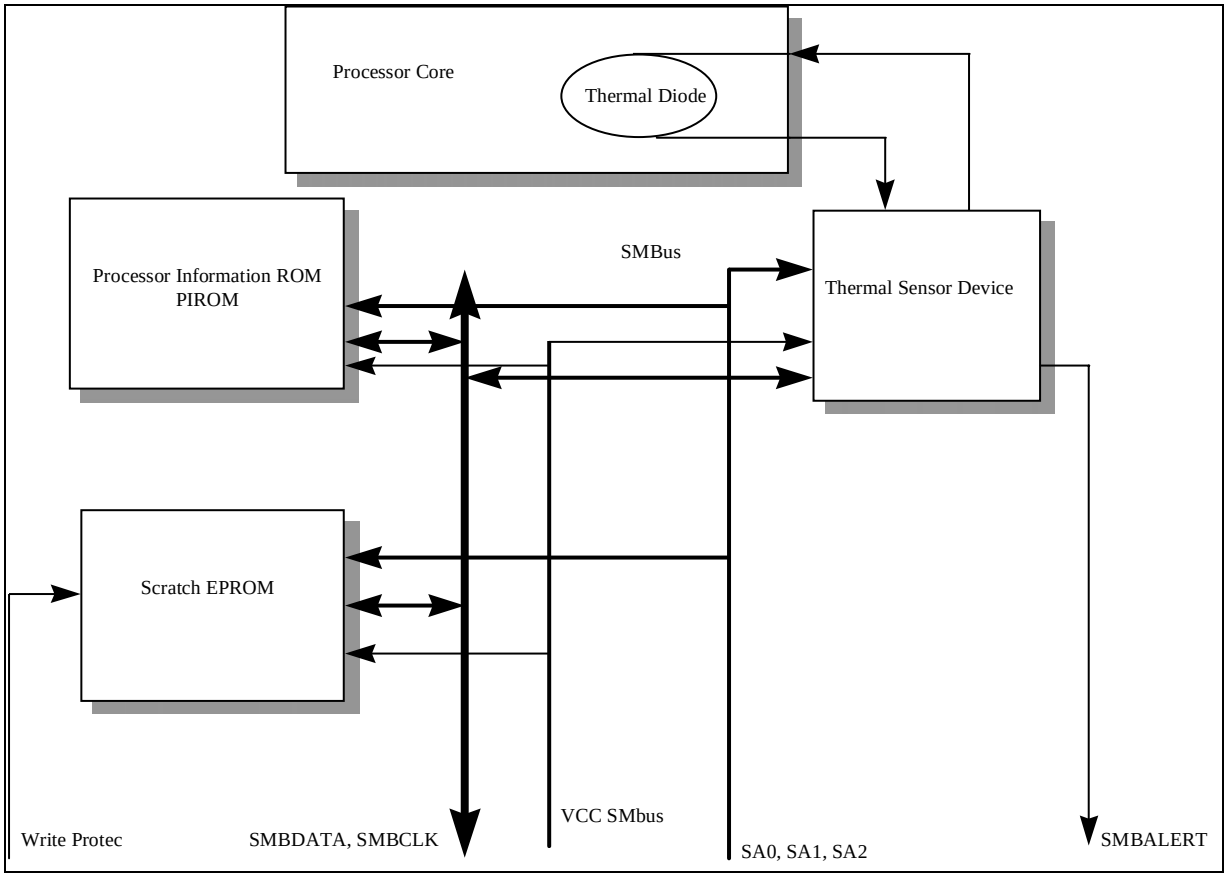

Figure 2. Block Diagram of The Pentium® II Xeon™ Processor Management Components

4.1 System Management Mode (SMM)

SMM is a feature preserved from previous generations of Intel processors. Server platform designers may choose to use SMM in various system management applications. One example might be to use the SMM error recovery mechanism. In the event of detected internal errors, address parity errors or system bus multi-bit data errors, an SMM routine may be deployed to identify the source of the error, log the error, and disable malfunctioning hardware pieces. Another use of SMM is to flush out the processor data cache content in an attempt to preserve the data of an ailing processor.

4.2 Functional Redundancy Checking (FRC)

Functional Redundancy Checking (FRC) allows two Pentium II Xeon processors to be configured as a pair, with one processor acting as the master and the other as a checker. The pair operates as a single processor, increasing system fault detection and data integrity. In the case of a mismatch between the processors, the checker processor asserts FRCERR. This triggers the master processor to issue a Machine Check Exception (MCE), which results in a software fault detection and prevention mechanism to log and possible recover from the error. When this condition occurs, system management may choose to alert the system administer via system management instrumentation tools.

4.3.0 Low Power States and Clock Control

The Pentium II Xeon processor allows three power-down states:

- HALT
- Stop-Grant
- Sleep

System management hardware may take advantage of these features to implement temperature control logic for the purpose of monitoring a processor's core (via a processor thermal sensor on the cartridge), cartridge, and ambient temperatures. Once system management is alerted of an abnormal temperature rise, the management solution may choose to start the processor power-down sequence and/or activate a backup cooling system to prevent damage to the cartridge. In HALT and Stop-Grant states, the processor continues to monitor system bus activities and respond to bus snoops. The system management recovery mechanism may choose to flush the content of the cache in the cartridge while the rest of the unit attempts to stabilize the cartridge's thermal environment. In this way, processor up-time is improved, enhancing the chances of complete data recovery. The multi-stage power-down sequencing adopted by the Pentium II Xeon processor is outlined in the following sections.

4.3.1 Auto-Halt Power-Down State

The first power-down state is entered by executing the HALT instruction. In this state, while the processor is halted, it continues to snoop the internal cache, allowing normal system operation. Depending on the stability of the system temperature, system management may choose to bring the processor back to a normal operating environment by issuing a System Management Interrupt (SMI), or it may choose to enter a lower level power state by going through the stop-clock/stop-clock acknowledge protocol, and entering the stop-grant state.

4.3.2 Stop-Grant State

The Stop-Grant state is entered by asserting the Stopclk# input and receiving a Stop-Clock Acknowledge from the processor. This state is entered from either the normal operating mode or from the Auto HALT power-down state. This is the second-level power-down state where the processor continues to snoop its internal cache and allow system operation. In the case of continuous temperature instability, the system management software may choose to enter the Sleep state. This preventive measure protects the processor from permanent damage.

4.3.3 Sleep State

This is the lowest-power state in which the processor maintains its context. The Sleep State can only be entered from the Stop-Grant state. Since in this state the processor no longer snoops its internal cache to ensure proper system operation, system management software must first flush the processor's cache prior to entering this state. Once system temperature is stabilized, the system management recovery mechanism can bring the processor back to its normal operating state, following the same state transition protocol in reverse.

4.4 System Management Bus (SMBus)

The Pentium II Xeon processor incorporates an SMBus interface, which allows access to management components such as the PIROM, Scratch EEPROM and thermal sensor residing on the processor substrate. This addition is a major enhancement to the Pentium II Xeon processor cartridge, which provides mission-critical data to system management software, enabling implementation of advanced instrumentation and superior RASUM. The SMBus enables thermal monitoring, PIROM access, and Scratch EEPROM access, all via a common two-wire interface.

4.4.1 Processor Information ROM (PIROM)

The PIROM resides on the substrate of the Pentium II Xeon processor, interfacing to the baseboard via the SMBus. The PIROM contains detailed information about the processor cartridge; parameters such as the Thermal Reference Byte is determined during manufacturing burning test. The Thermal Reference Byte closely matches each processor's thermal characteristics, which allows the implementation of efficient monitoring logic, further prolonging system up-time. In the following sections, the contents of the PIROM are discussed and their applications in the management environment are explored.

The addition of the PIROM, to the Pentium II Xeon processor cartridge facilitates the implementation of superior instrumentation and accurate fault detection and prevention mechanisms at the component and cartridge levels. The PIROM content can be used to:

- Accurately configure the platform frequency, thermal monitoring, and voltage settings.
- Implement voltage monitoring logic at higher degrees of accuracy and lower cost.
- Provide software-readable component type and revision information that can be used for inventory and asset management.
- Provide software-readable component type and revision information that can be used to verify that supported processor versions are being used in the system.
- Provide software-readable component type and revision information that can be used to verify that appropriately matched processors are being used in a multi-processor system.

Provide component authentication which can be a preventive measure against manufacturing and installation flaw.

4.4.1.1 PIROM Processor Data

The PIROM processor information provides key information identifying cartridge specifications and characteristics that benefit system management in the areas of asset management, configuration management, server management, and inventory management. Data contained in this segment of the PIROM include:

- Engineering sample processor versus production processor
- S-spec or Qualification Detail Form (QDF) number
- Checksum

These data identify the processor core as either an engineering sample, with the Qualification Detail Form (QDF) number along with specification revision or a production silicon. Management tools may use the PIROM processor data to closely match the installed processors in a server platform.
The PIROM processor data can be used for inventory or configuration management, outlining detail information, such as the S-spec, of the Pentium II Xeon processor cartridges installed. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

4.4.1.2 PIROM Core Data

The PIROM core data provides key information identifying core specifications and characteristics that benefit system management in the areas of asset management, configuration management, performance management, and inventory management. The PIROM core contains:

- Processor core type
- Processor core family
- Processor core model
- Processor core stepping
- Maximum core frequency
- Core voltage I.D.
- Core voltage tolerance, high and low
- Checksum

The PIROM core data outlines the processor type, family, model and stepping, this is similar to the data obtained by executing the CPU-ID instruction. Such data can be helpful in authenticating the processor core installed on the cartridge.

The PIROM data identifying the maximum core frequency can be used to accurately configure a system's clock circuitry. In server platforms where switches are used to configure the clock circuitry, accessing the PIROM allows automatic configuration and eliminates human intervention. The PIROM core data enables additional means of cross checking installed processor's clock rating, assisting the configuration management in ensuring consistent usage of the installed processors.

The PIROM data on core Voltage ID (VID), in millivolts, enables the dynamic monitoring of the processor core voltage. Including the processor core voltage requirement enables the monitoring logic to function at a higher degree of accuracy in comparison to the VID mechanisms adopted in past. The benefits of closely monitoring the core voltage is implementation of a failure detection mechanism, targeting the Voltage Regulating Module (VRM). This is a preventive approach to ensuring that the cartridge continuously operates within the characterized environment, improving system reliability and availability. The cost of voltage monitoring logic can improve by delegating this task to the management software, versus the previous implementations that connects the VID signals to the detection logic.

The system administrator may choose to use the stepping information to identify the cartridges ready for upgrade or replaced, either because of obsolescence or design improvements. In a multiprocessor environment, this information can be used to check for consistent (matched) stepping among the installed

processors. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

4.4.1.3 PIROM L2 Cache Data

The PIROM Level 2 (L2) cache data provides key information about the cartridge L2 cache specification and its characteristics. This information benefits system management in the areas of asset management, configuration management, performance management, and inventory management. The data in this segment of the PIROM includes:

- Level 2 cache size
- Number of synchronous RAM components
- Level 2 cache voltage I.D.
- Level 2 cache voltage tolerance, high and low
- Cache/tag stepping I.D.
- Checksum

The PIROM L2 cache data size is helpful in authenticating cache sizing algorithms, implemented for sizing and testing L2 cache during the Power-On Self-Test (POST). Such authentication can provide an early indication of any cache malfunctioning. The PIROM L2 cache data is an additional means of cross-checking the installed processors' cache size and stepping information.

The PIROM data on cache VID (resolution in millivolts), enables the dynamic monitoring of the cache voltage. Including the cache voltage requirement enables the monitoring logic to function at a higher degree of accuracy, compared to VID mechanisms adopted in past. The benefit of closely monitoring the cache voltage is implementation of a failure detection mechanism, targeting the VRM. This is a preventive approach to ensuring that the cartridge continuously operates within the characterized environment, improving the system reliability and availability.

By returning this information as a voltage, the system software does not need to compute the operating voltage and tolerance (range) values by doing a table lookup on the VID values. This isolates the software in the case of changes to the VID specification. The voltage monitoring logic and the implementation cost can improve by delegating this task to the management software, versus previous implementations that connected the VID signals to the detection logic. This information eliminates the need for bringing the VID lines into the platform management subsystem, which can save a significant amount of wiring space on a multiprocessor system board.

Lastly, the system administrator may choose to use the Cache/TAG stepping ID to identify the cartridges necessary for upgrade due to obsolescence or design improvements. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

4.4.1.4 PIROM Cartridge Data

The PIROM cartridge cache data provides key information identifying cartridge specification and characteristics that benefits system management in the areas of asset management, inventory management and configuration management. This includes:

- Cartridge revision
- Substrate revision
- Checksum

The PIROM cartridge data provides the means of cross-checking the cartridge substrates installed in a server platform ensuring uniform installation. This information may also be used for instrumentation and platform configuration management, as well as inventory control and system management. System administrators may choose to use this data in identifying the cartridges required for upgrade due to

obsolescence or design improvements. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

4.4.1.5 PIROM Part Numbers Data

The PIROM part number data provides key information identifying cartridge specification and characteristics that benefit system management, asset management, security management and configuration management. This data includes:

- Processor part number
- Processor BOM I.D.
- Processor electronic signature
- Checksum

The PIROM part number information provides a means of cross-checking the product's part number and BOM ID. This data provides detailed visibility on the bills of material used for each cartridge, enabling IT manager to mix certain cartridges and eliminate others.

The electronic signature is an unprecedented feature added to the Pentium II Xeon processor cartridge that enables a wide variety of management-related implementations. The electronic signature is unique data for each the Pentium II Xeon processor. The signature is provided as an aid to processor inventory management and asset tracking. The signature is programmed into the PIROM during manufacturing test. Value-added security management features may make use of this data, along with the scratch EEPROM on the cartridge, to track each processor to a specific baseboard and prevent unauthorized removal of the cartridge from a platform. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

4.4.1.6 PIROM Thermal Reference Data

The PIROM thermal reference data provides key information that identifies cartridge specification and characteristics that benefits system management in the areas of reliability and configuration management. The is data includes:

- Processor thermal reference byte
- Checksum

The Thermal Reference Byte identifies the processor's maximum allowable operating temperature, with a high degree of accuracy, +/- 1°c. This value is determined for each Pentium II Xeon processor during manufacturing test. This data is obtained by reading the thermal sensor while the processor is operating in a high-power test environment, with the thermal plate raised to its maximum specified temperature. The system management software may use this data to configure the thermal sensor device on the cartridge (details in 4.4.3). Availability of thermal reference data ensures that each processor is operating within its factory-tested, worst-case temperature thus protecting the cartridge from over temperature hazards. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

4.4.1.7 PIROM Features Data

The PIROM feature data provides key information identifying cartridge features enabled in the PIROM. The system management software uses these bytes to ensure validity of data in each segments of the PIROM.

- Processor core features flag
- Cartridge feature flags:
  - Electronic signature present
  - Thermal sensor present
  - Thermal reference byte present
  - Scratch EEPROM present
  - L2 cache VID present
- Checksum

System management software must use these data to validate the PIROM content for each of the segments. The checksum byte validates accuracy and reliability of the data programmed in this segment of the PIROM.

	Byte	*# of Bits*	*Function*	*Notes*
Header	00h	8	Data Format Revision	Two 4-bit hex digits
		16	EEPROM Size	Size in bytes
		8	Processor Data Address	Pointer to a byte (00h, if not present)
		8	Processor Core Data Address	Pointer to a byte (00h, if not present)
		8	L2 Cache Data Address	Pointer to a byte (00h, if not present)
		8	SEC Cartridge Data Address	Pointer to a byte (00h, if not present)
		8	Part Number Data Address	Pointer to a byte (00h, if not present)
		8	Thermal Reference Data Address	Pointer to a byte (00h, if not present)
		8	Feature Data Address	Pointer to a byte (00h, if not present)
		8	Other Data Address	Pointer to a byte (00h, if not present)
		16	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
Processor	0Eh	48	S-spec/QDF Number	Six 8-bit ASCII characters
		2	Sample/Production	00b = Sample only
		6	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
Core	16h	2	Processor Core Type	From CPUID
		4	Processor Core Family	From CPUID
		4	Processor Core Model	From CPUID
		4	Processor Core Stepping	From CPUID
		42	Reserved	Reserved for future use
		16	Maximum Core Frequency	16-bit binary number in MHz
		16	Core Voltage ID	Voltage in mV
		8	Core Voltage Tolerance, High	Edge finger tolerance in mV, +
		8	Core Voltage Tolerance, Low	Edge finger tolerance in mV, -
		8	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
L2 Cache	25h	32	Reserved	Reserved for future use
		16	L2 Cache Size	16-bit binary number in Kbytes
		4	Number of SRAM Components	One 4-bit hex digit
		4	Reserved	Reserved for future use
		16	L2 Cache Voltage ID	Voltage in mV
		8	L2 Cache Voltage Tolerance, High	Edge finger tolerance in mV, +
		8	L2 Cache Voltage Tolerance, Low	Edge finger tolerance in mV, -
		4	Cache/Tag Stepping ID	One 4-bit hex digit
		4	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
Cartridge	32h	32	Cartridge Revision	Four 8-bit ASCII characters
		2	Substrate Revision Software ID	2-bit revision number
		6	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
Part Numbers	38h	56	Processor Part Number	Seven 8-bit ASCII characters
		80	Processor BOM ID	Ten 8-bit ASCII characters
		64	Processor Electronic Signature	64-bit unique identification number
		240	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
Thermal Ref.	70h	8	Thermal Reference Byte	Reference byte for Thermal Sensor
		16	Reserved	Reserved for future use
		8	Checksum	1 byte checksum
Features	74h	32	Processor Core Features Flags	From CPUID
		32	Cartridge Feature Flags	[5] = Electronic Signature Present
				[4] = Thermal Sensor Present
				[3] = Thermal Reference Byte Present
				[2] = Scratch EEPROM Present
				[1] = Core VID Present
				[0] = L2 Cache VID Present
		4	Number of Devices in TAP Chain	One 4-bit hex digit
		4	Reserved	Reserved for future use
		8	Checksum	1 byte checksum

Table 1, Processor Information ROM Data Position and Format

4.4.2 Scratch EEPROM

The presence of scratch EEPROM (128 bytes of EEPROM) on the Pentium II Xeon processor allows system management software to store data on the processor cartridge. The Scratch EEPROM interfaces to the system management controller through the SMBus. The write-protect pin on the cartridge-edge fingers may be used by the controller on the baseboard to write-protect data in the scratch EEPROM. The EEPROM can be used to store a unique asset tag for the processor. For example, system management software may choose to store an encoded data that allows it to implement a tracking mechanism for the processors on a specific baseboard. The following flow diagram outlines this implementation, which is an attempt to prevent the unauthorized removal of the cartridge from the baseboards:

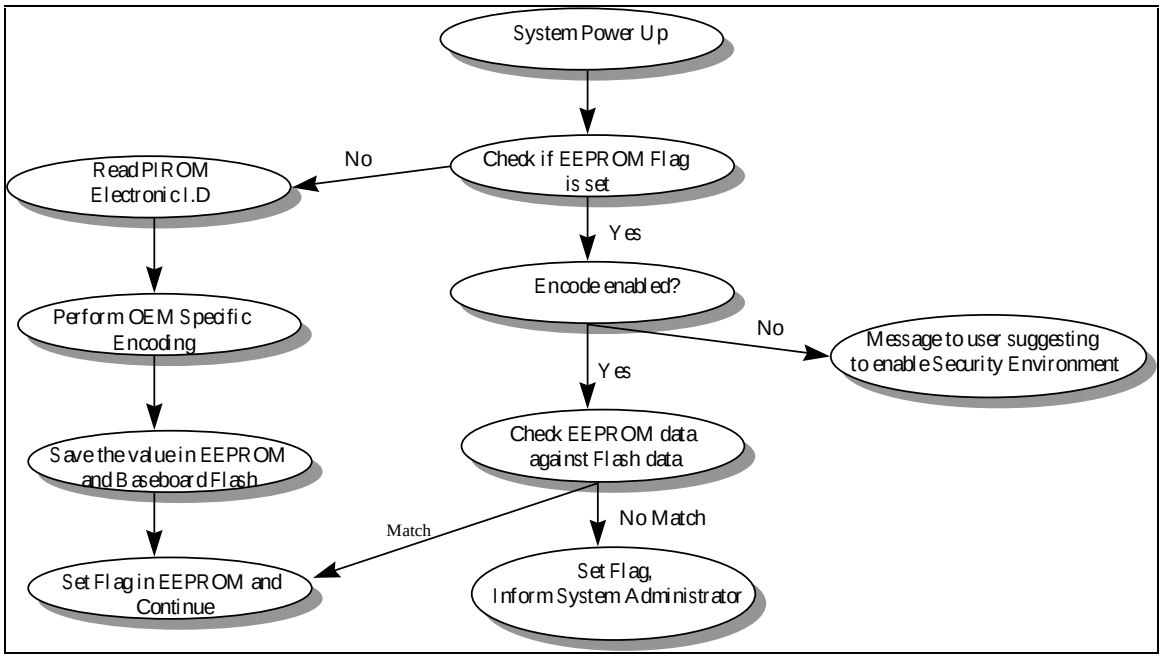

Other uses for the scratch EEPROM include:

- Keeping a log file on the last diagnostic ran on the cartridge
- Maintaining a record of processor micro-code updates
- Keeping track of processor errors, i.e., errors detected on the Level 2 cache or inside the processor
- Maintaining a log file of each processor's up-time, which can be updated by management software running on the BMC monitoring power cycles.

Such error-logging mechanism can be used to detect symptoms of ailing processors so IT managers can eliminate unreliable equipment. For example, if a processor is replaced as part of a system repair, but there was a question as to whether the processor was indeed faulty (no trouble found), it could be 'tagged' as having been part of a previous replacement operation. In a faulty system, the collaborating processor can then be pin-pointed by checking the history log-file. This type of record-keeping allows IT managers to decrease time-to-repair, routinely perform system diagnostics, or replace units as required.

4.4.3 Thermal Sensor Device

The Pentium II Xeon processor includes a thermal sensor on the substrate of the cartridge. This device interfaces with the system management controller through the SMBus. This approach benefits processor thermal management by providing improved degrees of accuracy and correlation with

processor temperature. Such implementation eliminates discrete components from the baseboard and achieves a consistent performance at lower cost.

The thermal sensor is connected to a thermal diode on the processor core. The thermal sensor provides the earliest indication, asserts SMBALERT# of abnormal thermal variation and early indications of a possible violation of a processor's thermal operating environment. The system designer may choose to use the SMBALERT# signal assertion to invoke the backup cooling mechanism and initiate the cool-down sequence via the low-power states (Halt, Stop-Grant, Sleep). This effort may increase system up-time to provide time for data recovery and graceful processor shutdown. Figure 2. illustrates the block diagram of the thermal diode and the thermal sensor on the processor. It should be noted that the thermal diode is integrated in the processor core, the thermal sensor does not report the thermal status of the L2 cache components. To prevent a scenario where the processor core temperature is a lower temperature than the L2 cache, cartridge specifications must be used as the overriding specification in all cases. The thermal sensor values, at no times, guarantees to represent cartridge temperature.

5.0 Pentium® II Xeon™ Processor Manageability Benefits

The Pentium® II Xeon™ processor was designed with the manageability needs of real-world IT managers in mind. These typically include:

1. An easy way to do comprehensive asset management. The lack of simple tools to gather inventory data creates a major roadblock for most IT departments. How can they upgrade hundreds or thousands of desktops or servers if they do not know what is out there?

2. An easier way to do IT support. They need to do quick and accurate repairs, implement error-tracking and prevention mechanisms, and install user- friendly hardware to aid inexperienced operators in resolving issues quickly.

3. Eliminating unreliable equipment.

4. Capture performance data for doing capacity planning. This is needed to eliminate unreliable equipment, plan expansions based on platform performance, and do system bandwidth performance monitoring.

5. Capture dynamic system scanning data to log the last minutes of a platform's operation prior to failure.

5.1 The Pentium® II Xeon™ Processor Fits Common Platform Architecture

Addition of the SMBus to the Pentium II Xeon processor cartridge, allows seamless integration with an IPMI architecture, through BMC via a private management bus. The communication is a standard I^2C serial bus protocol and is readily supported using other management hardware implementations. Using the Pentium II Xeon processor, together with an IPMI architecture, materializes benefits of the common platform architecture-based system management implementation. This allows:

- Flexible access to platform management information
- Industry open implementation for easier server-platform instrumentation
- De-coupling server management software from hardware, which enables the hardware to change without impacting the software
- Reduced development time for new products and enabling remote management

5.2 The Pentium® II Xeon™ Processor Benefits Asset Management

Server platforms with the Pentium II Xeon processor provide detailed information about the cartridge, which helps IT manager perform asset tracking by:

- Cartridge part-number
- BOM ID the cartridge is built on
- Electronic signature of the cartridge
- Cartridge substrate revision
- Cartridge revision
- Cartridge cache size
- Processor core type
- Cartridge QDF number

This enables IT managers to build an extensive database on their existing inventory of Pentium II Xeon processor-based server platforms, re-arrange high-end server platform to high performance demanded users and detecting the nodes which need to be upgraded or serviced.

5.3 The Pentium® II Xeon™ Processor Benefits Configuration Management

Processor information allows the creation of a configuration manager that can inspect the installed Pentium II Xeon processors in a given server platform by reading back information about the cartridge:

- Core type and stepping
- Maximum core frequency
- Level 2 cache size
- Cartridge voltage requirements

The configuration manager can use this information for checking the installed processors are of the same kind, and configured to operate at the maximum clock frequency. This results in best system resource optimization. Additional mechanisms may be deployed to alert the system administrator of inefficient system usage and possible processor over-clocking conditions.

5.4 The Pentium® II Xeon™ Processor Benefits Inventory Management

Server platforms with the Pentium II Xeon processor provide detailed information about the cartridge installed in the platform, which helps IT manager in determining:

- Cartridge part-number
- Cartridge BOM ID
- Cartridge Electronic signature
- Cartridge substrate revision
- Cartridge revision
- Cartridge cache size
- Processor core type
- Cartridge QDF number

Using this information, IT managers can build an extensive database of Pentium II Xeon processor-based servers, optimizing performance level by matching demanding users with the most high-powered servers. IT managers may also use such a database for detecting and eliminating obsolete components.

5.5 The Pentium® II Xeon™ Processor Benefits Performance Management

Server platforms with the Pentium II Xeon processor provide detail information about the cartridge installed in the platform, which enables IT managers to determine:

- Maximum clock frequency
- Processor core type
- L2 cache size

IT managers may choose to use this data to ensure that processors of the same kind and performance are being used on the same node to prevent under-utilization of any processor cartridge.

5.6 The Pentium® II Xeon™ Processor Benefits Security Management

The Pentium II Xeon processor cartridge with the electronic signature and Scratch EEPROM allows the creation of system management software that ties the cartridge to a specific baseboard. This prevents

unauthorized replacement of the cartridge from the nodes in restricted operating environments, or as a tool for aiding warranty tracking and repair.

5.7 The Pentium® II Xeon™ Processor Benefits Server Management

The Pentium II Xeon processor cartridge provides detailed information about its operating characteristics, enabling the management hardware to detect and prevent failure conditions and allowing safe fail-over, impacting server RASUM.

As mentioned, the system management enhancement in the Pentium II Xeon processor addresses key concepts in the instrumentation and RASUM areas. The system management enhancements to the Pentium II Xeon processor dramatically extend and enhance existing processor management capabilities, enabling new and robust methodology for monitoring, controlling, managing and illustrating the processor's operating environments, features and characteristics.

With strong instrumentation capabilities built into the server platforms with the Pentium II Xeon processor, IT departments are able to accurately maintain and track server network components, implementing upgrade planning and eliminating unreliable equipment. Server platforms with the Pentium II Xeon processor can use numerous failure prevention mechanisms to deliver an improved and reliable system operation.

The Pentium II Xeon processor meets the advance requirements of a well-managed server platform by providing detailed information about itself, enhancing instrumentation to levels never achieved by the preceding processors. The Pentium II Xeon processor is designed with manageability in mind.

APPENDIX A

A.1 Management Initiatives Background

The first management initiative primarily focused on cross-platform system management was first introduced by the Desktop Management Task Force(DMTF) a consortium of more than 100 vendors established in 1992 committed to make computing platforms easier to use, configure and manage. The DMTF developed the Desktop Management Interface (DMI) focusing on defining and standardizing software aspects of system management implementation.

Earlier protocols, such as the Simple Network Management Protocol (SNMP), is complemented by the DMI. The DMI-to-SNMP mapping software is available that allows the DMI-based instrumentation to be cleanly integrated into the SNMP-based management environments.

Since then additional system management initiatives, i.e., Wired for Management (WfM), Web-Based Enterprise Management (WBEM), Zero Administration Windows Initiative(ZAW) were formed with a common focus of utilizing system management as the primary means of reducing the Total Cost of Ownership (TCO).

The first server management specification aimed at creating an industry open hardware implementation of a managed platform is the “Intelligent Platform Management interface” or IPMI. The IPMI specification defines a common interface and a message-based protocol for accessing platform management hardware. This helps reduce TCO by improving server platform management functionality and compatibility, while de-coupling the hardware implementation from the software. IPMI, an industry open based implementation, allows hardware advancements be implemented without impacting server management software.

A.2.0 Manageability and Management Areas

Manageability is defined as the use of technologies and products to enhance server RASUM (Reliability, Availability, Serviceability, Usability and Manageability), thus increasing up-time and reducing the cost of deployment, ownership and administration. Manageability rests on having a range of products, systems and system components working together to provide information and interfaces that system management software, management applications and operating systems can use to improve RASUM. Manageability includes asset management, configuration management, inventory management, network management, performance management, security management, server management and system management.

A.2.1 Asset Management

Asset management is the process of maximizing the use of assets to produce revenue while minimizing overall costs. Manageable Server platforms contribute to asset management by capturing inventory and tracking information, enabling organizations to analyze key cost variables and better discern the return on investment (ROI) of their technology purchases. The data gathered by manageable systems can assist asset management issues such as inventory consolidation and rationalizing license issues, leasing considerations, analyzing training costs, analyzing software upgrades for volume purchasing plans, evaluating the cost-efficiency of outsourcing and improving warranty usage.

A.2.2 Configuration Management

Configuration management deals with tasks such as optimizing the user’s configuration for a given task, and discovering what hardware and components are in the Server platform to ensure compatibility among them. This includes features that prevent or detect system mis-configuration.

A.2.3 Inventory Management

Inventory management addresses the need to identify system and software components in a system. Instrumented platforms can show a complete inventory of all components and subsystem's information that can be highly useful for diagnosing problems remotely, maximizing assets and optimizing configurations and performance levels. An important use of inventory information is to providing Field Replaceable Unit (FRU) data that assists in the rapid replacement of modules during servicing. This in turn improves system up-time by decreasing time-to-repair.

A.2.4 Network Management

Network management is one of the three major components of managing a computing environment. Network management includes the performance, configuration, security, failure analysis and repair of the infrastructure components such as switches, routers, bridges and gateways in a LAN,WAN or Internet/Intranet. Network management is a related, but different management area than "System Management" that focuses on the management of the computer system hardware and applications rather than the network.

A.2.5 Performance Management

Performance management analyzes usage pattern and attempts to optimize system responsiveness and network throughput. This enables IT to remotely monitor CPU utilization and pinpoint the user, application or process that is generating the excessive demand.

A.2.6 Security Management

Security management addresses the need to protect systems and data from unauthorized access. Examples include encryption, authentication, and firewalls.

A.2.7 Server Management

Server management aims to optimize the RASUM and performance of network servers and includes monitoring such factors as disk capacity and user accounts. Server hardware management includes reliability and failure tolerance features that are not commonly found in desktop systems. This includes features such as multiprocessor monitoring and recovery features, redundant cooling (fan) and power supply management, redundant driver array (RAID) management, multiple system board management and field replaceable unit identification, and multiple chassis management.

A.2.8 System Management

System management refers to controlling, configuring, installing and monitoring the applications, servers and clients in a distributed computing environment.

APPENDIX B

B.1 Example DMI Software Stack

The DMI architecture defines the following software stack between the application software and managed components. The manageable feature is specified in a standard ASCII text file format called Management Information Format (MIF). Additional information can be obtained from the Desktop Management Task Force via their web site at http://www.dmtf.org. The DMI software stack illustration is following:

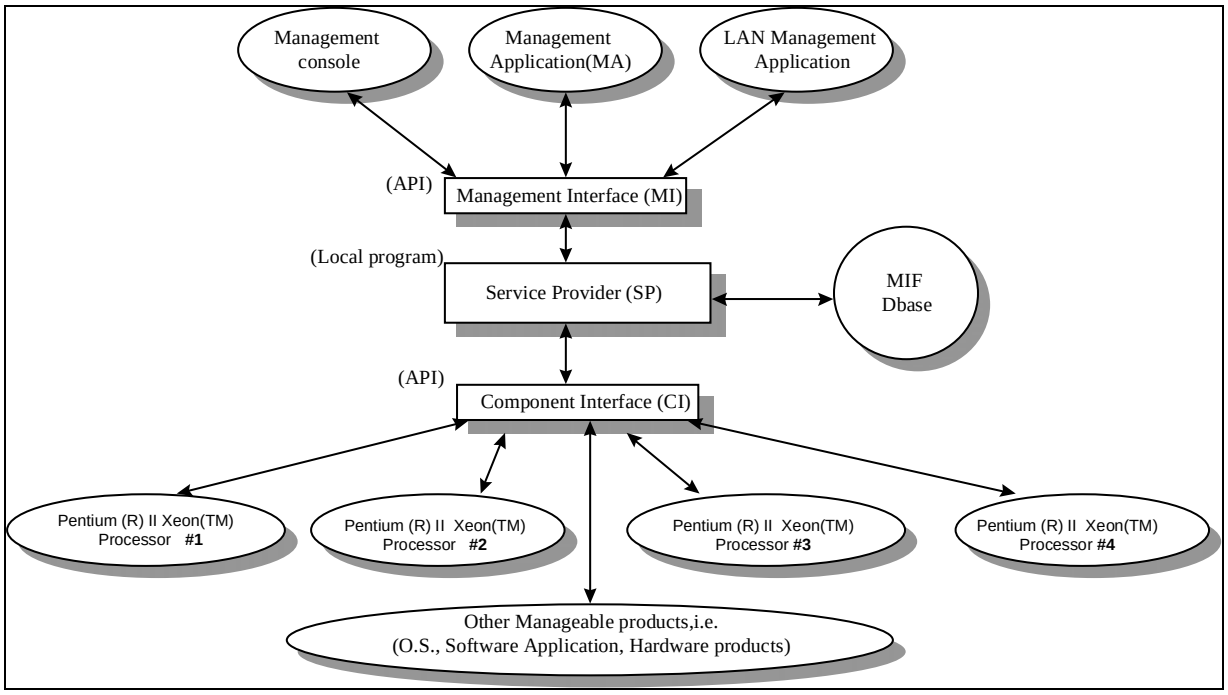

Figure 3, Flow diagram of DMI software stack

1- Management Application (MA): Program that initiates management requests. The MA uses the DMI Service provider (SP) to access data from Component Instrumentation via the CI Manager. DMI also defines a Management Information Format(MIF) where operating systems and management applications make use of for instrumentation purposes. The GUI application is OEM specific with the following common usage:

- Graphical display of the system's manageable components
- Issue detailed alert pending failures
- Watch for indicators of potential security breach
- Provide configuration and dynamic operational status to determine the system health

2- Management Interface(MI): An API that provides the interface between the Service Layer and management applications and allows these applications to access, manage and control the platform components. The MI offers a consistent interface for any management application to the various mechanisms used to obtain information from products and components within a platform.

3- The Service Provider(SP): Program that resides in the platform and is responsible for all DMI activities. This layer collects management information from products (whether system hardware, peripherals or software, and store that information in the DMI's database and passes it to management applications as requested.

4- Component Interface (CI): An API that handles communication between manageable elements and the DMI's Service Layer. The CI gives all hardware or software components a common method for describing their management attributes or features.

UNITED STATES, Intel Corporation
2200 Mission College Blvd., P.O. Box 58119, Santa Clara, CA 95052-8119
Tel: +1 408 765-8080

JAPAN, Intel Japan K.K.
5-6 Tokodai, Tsukuba-shi, Ibaraki-ken 300-26
Tel: + 81-29847-8522

FRANCE, Intel Corporation S.A.R.L.
1, Quai de Grenelle, 75015 Paris
Tel: +33 1-45717171

UNITED KINGDOM, Intel Corporation (U.K.) Ltd.
Pipers Way, Swindon, Wiltshire, England SN3 1RJ
Tel: +44 1-793-641440

GERMANY, Intel GmbH
Dornacher Strasse 1
85622 Feldkirchen/ Muenchen
Tel: +49 89/99143-0

HONG KONG, Intel Semiconductor Ltd.
32/F Two Pacific Place, 88 Queensway, Central
Tel: +852 2844-4555

CANADA, Intel Semiconductor of Canada, Ltd.
190 Attwell Drive, Suite 500
Rexdale, Ontario M9W 6H8
Tel: +416 675-2438