Intel® Xeon® Processor E5-2400 Product Family

Datasheet - Volume One

May 2012

Reference Number: 327248-001

Table of Contents

Figures

Tables

Revision History

Revision Number	Description	Revision Date
001	Initial Release	May 2012

§

1 Overview

1.1 Introduction

The Intel® Xeon® Processor E5-2400 Product Family Datasheet - Volume One provides DC specifications, signal integrity, differential signaling specifications, land and signal definitions, and an overview of additional processor feature interfaces.

The Intel® Xeon® processor E5-2400 product family is the next generation of 64-bit, multi-core enterprise processors built on 32-nanometer process technology. Throughout this document, the Intel® Xeon® processor E5-2400 product family may be referred to as simply the processor. Based on the low-power/high performance 2nd Generation Intel(r) Core(TM) Processor Family microarchitecture, the processor is designed for a two chip platform consisting of a processor and a Platform Controller Hub (PCH) enabling higher performance, easier validation, and improved x-y footprint. The Intel® Xeon® processor E5-2400 product family is designed for server-class platforms including those for embedded and storage applications with up to two processor sockets. The processor family features one Intel® QuickPath Interconnect (Intel® QPI) point-to-point link capable of up to 8.0 GT/s, up to 24 lanes of PCI Express* 3.0 links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI Express* 2.0 interface with a peak transfer rate of 5.0 GT/s. The processor supports up to 46 bits of physical address space and 48-bit of virtual address space.

Included in this family of processors is an integrated memory controller (IMC) and integrated I/O (IIO) (such as PCI Express* and DMI2) on a single silicon die. This single die solution is known as a monolithic processor.

1.1.1 Processor Feature Details

Figure 1-1. Intel® Xeon® Processor E5-2400 Product Family on the 2 Socket Platform

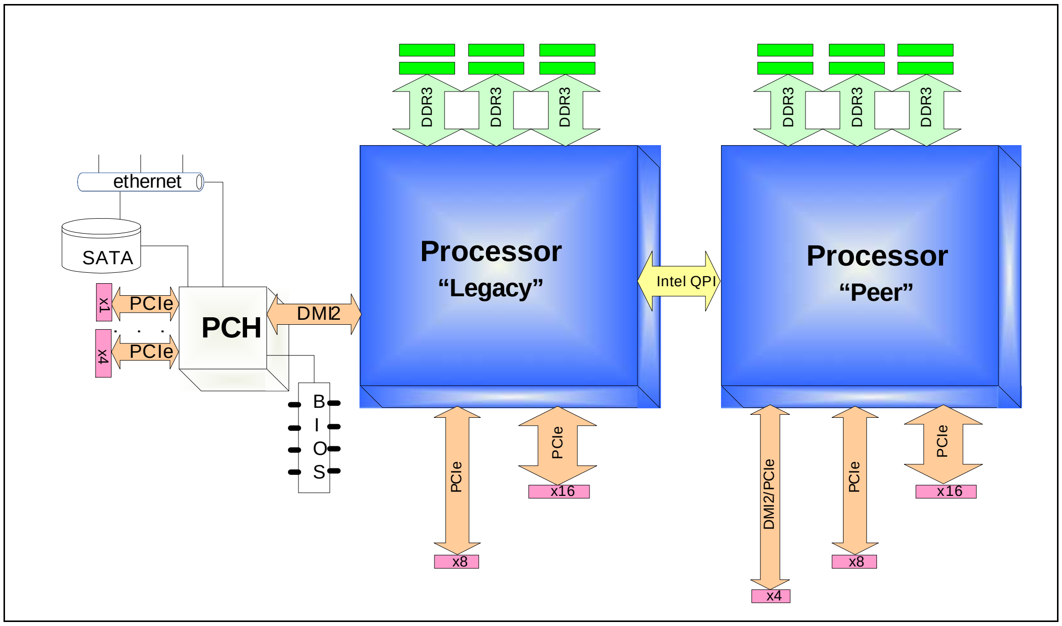

- Up to 8 execution cores

- Each core supports two threads (Intel® Hyper-Threading Technology), up to 16 threads per socket
- 46-bit physical addressing and 48-bit virtual addressing
- 1 GB large page support for server applications
- A 32-KB instruction and 32-KB data first-level cache (L1) for each core
- A 256-KB shared instruction/data mid-level (L2) cache for each core
- Up to 20 MB last level cache (LLC): up to 2.5 MB per core instruction/data last level cache (LLC), shared among all cores

1.1.2 Supported Technologies

- Intel® Virtualization Technology (Intel® VT)
- Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)
- Intel Virtualization Technology Processor Extensions
- Intel® Trusted Execution Technology (Intel® TXT)
- Intel® Advanced Encryption Standard Instructions (Intel® AES-NI)
- Intel 64 Architecture
- Intel® Streaming SIMD Extensions 4.1 (Intel SSE4.1)
- Intel Streaming SIMD Extensions 4.2 (Intel SSE4.2)
- Intel Advanced Vector Extensions (Intel AVX)
- Intel® Hyper-Threading Technology (Intel® HT Technology)
- Execute Disable Bit
- Intel® Turbo Boost Technology
- Intel® Intelligent Power Technology
- Enhanced Intel SpeedStep® Technology
- Intel® Dynamic Power Technology (Intel® DPT) (Memory Power Management)

1.2 Interfaces

1.2.1 System Memory Support

- Intel® Xeon® processor E5-2400 product family product family supports three DDR3 channels
- Unbuffered DDR3 and registered DDR3 DIMMs
- LR DIMM (Load Reduced DIMM) for buffered memory solutions demanding higher capacity memory subsystems
- Independent channel mode or lockstep mode
- Data burst length of eight cycles for all memory organization modes
- Memory DDR3 data transfer rates of 800, 1066, 1333, and 1600 MT/s
- 64-bit wide channels plus 8-bits of ECC support for each channel
- DDR3 standard I/O Voltage of 1.5 V and DDR3 Low Voltage of 1.35 V
- 1-Gb, 2-Gb and 4-Gb DDR3 DRAM technologies supported for these devices:

— UDIMMs x8, x16
— RDIMMs x4, x8
— LRDIMM x4, x8 (2-Gb and 4-Gb only)

- Up to 8 ranks supported per memory channel, 1, 2 or 4 ranks per DIMM
- Open with adaptive idle page close timer or closed page policy
- Per channel memory test and initialization engine can initialize DRAM to all logical zeros with valid ECC (with or without data scrambler) or a predefined test pattern
- Minimum memory configuration: independent channel support with 1 DIMM populated
- Integrated dual SMBus master controllers
- Command launch modes of 1n/2n
- RAS Support (including and not limited to):
  — Rank Level Sparing and Device Tagging
  — Demand and Patrol Scrubbing
  — DRAM Single Device Data Correction (SDDC) for any single x4 or x8 DRAM device failure. Independent channel mode supports x4 SDDC. x8 SDDC requires lockstep mode
  — Lockstep mode where channels 2 & 3 are operated in lockstep mode
  — Data scrambling with address to ease detection of write errors to an incorrect address.
  — Error reporting via Machine Check Architecture
  — Read Retry during CRC error handling checks by iMC
  — Channel mirroring within a socket
- Improved Thermal Throttling with dynamic Closed Loop Thermal Throttling (CLTT)
- Memory thermal monitoring support for DIMM temperature via two memory signals, MEM_HOT_C{1/23}_N

1.2.2 PCI Express*

- The PCI Express* port(s) are fully-compliant to the *PCI Express* Base Specification, Revision 3.0 (PCIe* 3.0)*
- Support for PCI Express* 3.0 (8.0 GT/s), 2.0 (5.0 GT/s), and 1.0 (2.5 GT/s)
- Up to 24 lanes of PCI Express* interconnect for general purpose PCI Express* devices at PCIe* 3.0 speeds that are configurable for up to 6 x4 independent ports
- 4 lanes of PCI Express* at PCIe* 2.0 speeds when not using DMI2 port (Port 0), also can be downgraded to x2 or x1
- Negotiating down to narrower widths is supported, see Figure 1-2:
  — x16 port (Port 3) may negotiate down to x8, x4, x2, or x1.
  — x8 port (Port 1) may negotiate down to x4, x2, or x1.
  — x4 port (Port 0) may negotiate down to x2, or x1.
  — When negotiating down to narrower widths, there are caveats as to how lane reversal is supported.
- Non-Transparent Bridge (NTB) is supported by PCIe* Port3a/IOU1. For more details on NTB mode operation refer to *PCI Express Base Specification - Revision 3.0*:
  — x4 or x8 widths and at PCIe* 1.0, 2.0, 3.0 speeds

  - Two usage models; NTB attached to a Root Port or NTB attached to another NTB
  - Supports three 64-bit BARs
  - Supports posted writes and non-posted memory read transactions across the NTB
  - Supports INTx, MSI and MSI-X mechanisms for interrupts on both side of NTB in upstream direction only
- Address Translation Services (ATS) 1.0 support
- Hierarchical PCI-compliant configuration mechanism for downstream devices.
- Traditional PCI style traffic (asynchronous snooped, PCI ordering).
- PCI Express* extended configuration space. The first 256 bytes of configuration space aliases directly to the PCI compatibility configuration space. The remaining portion of the fixed 4-KB block of memory-mapped space above that (starting at 100h) is known as extended configuration space.
- PCI Express* Enhanced Access Mechanism. Accessing the device configuration space in a flat memory mapped fashion.
- Automatic discovery, negotiation, and training of link out of reset.
- Supports receiving and decoding 64 bits of address from PCI Express*.
  - Memory transactions received from PCI Express* that go above the top of physical address space (when Intel VT-d is enabled, the check would be against the translated HPA (Host Physical Address) address) are reported as errors by the processor.
  - Outbound access to PCI Express* will always have address bits 63 to 46 cleared.
- Re-issues Configuration cycles that have been previously completed with the Configuration Retry status.
- Power Management Event (PME) functions.
- Message Signaled Interrupt (MSI and MSI-X) messages
- Degraded Mode support and Lane Reversal support
- Static lane numbering reversal and polarity inversion support

Figure 1-2. PCI Express* Lane Partitioning and Direct Media Interface Gen 2 (DMI2)

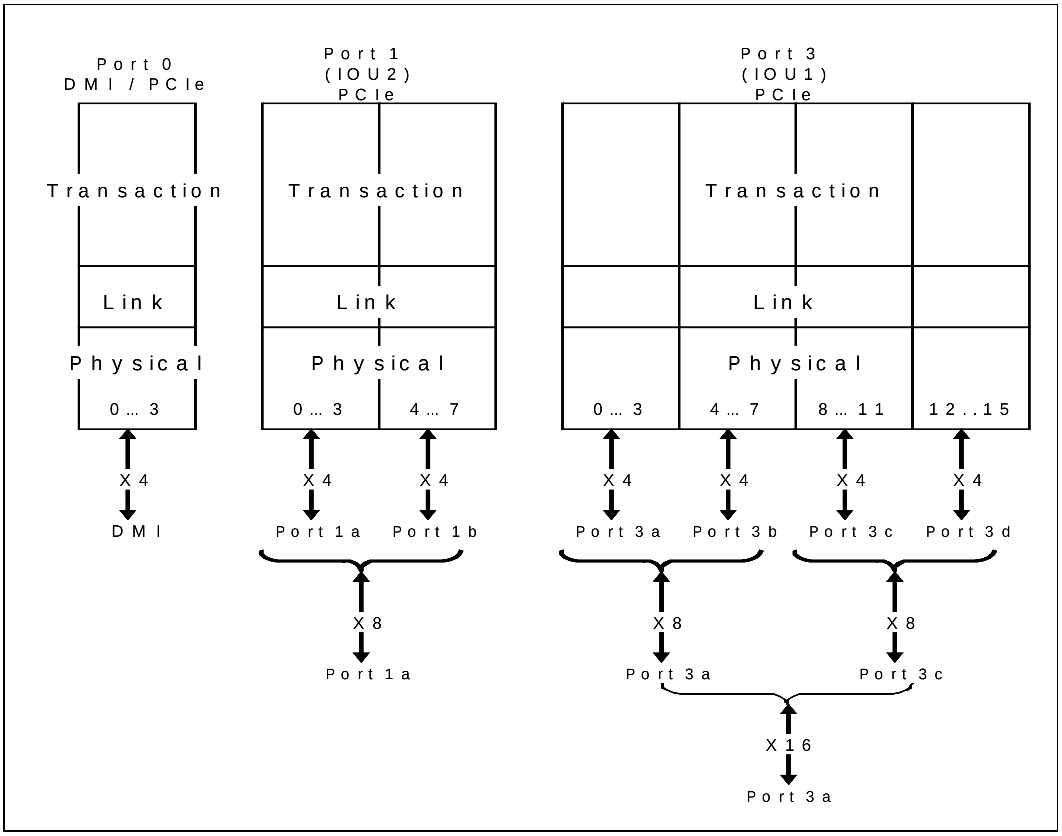

1.2.3 Direct Media Interface Gen 2 (DMI2)

- Serves as the chip-to-chip interface to the Intel® C600 Series Chipset
- The DMI2 port supports x4 link width and only operates in a x4 mode when in DMI2
- Operates at PCI Express* 1.0 or 2.0 speeds
- Transparent to software
- Processor and peer-to-peer writes and reads with 64-bit address support
- APIC and Message Signaled Interrupt (MSI) support. Will send Intel-defined “End of Interrupt” broadcast message when initiated by the processor.
- System Management Interrupt (SMI), SCI, and SERR error indication
- Static lane numbering reversal support
- Supports DMI2 virtual channels VC0, VC1, VCm, and VCp

1.2.4 Intel® QuickPath Interconnect (Intel® QPI)

- Compliant with Intel QuickPath Interconnect v1.1 standard packet formats
- Implements a full width Intel QuickPath Interconnect port
- Full width port includes 20 data lanes and 1 clock lane
- 64 byte cache-lines

- Home snoop based coherency
- 3-bit Node ID
- 46-bit physical addressing support
- No Intel QuickPath Interconnect bifurcation support
- Differential signaling
- Forwarded clocking
- Up to 8.0 GT/s data rate (up to 16 GB/s direction peak bandwidth per port)
  — All ports run at same operational frequency
  — Reference Clock is 100 MHz
  — Slow boot speed initialization at 50 MT/s
- Common reference clocking (same clock generator for both sender and receiver)
- Intel® Interconnect Built-In-Self-Test (Intel® IBIST) for high-speed testability
- Polarity and Lane reversal (Rx side only)

1.2.5 Platform Environment Control Interface (PECI)

The PECI is a one-wire interface that provides a communication channel between a PECI client (the processor) and a PECI master (the PCH).

- Supports operation at up to 2 Mbps data transfers
- Link layer improvements to support additional services and higher efficiency over PECI 2.0 generation
- Services include CPU thermal and estimated power information, control functions for power limiting, P-state and T-state control, and access for Machine Check Architecture registers and PCI configuration space (both within the processor package and downstream devices)
- PECI address determined by SOCKET_ID configuration
- Single domain (Domain 0) is supported

1.3 Power Management Support

1.3.1 Processor Package and Core States

- ACPI C-states as implemented by the following processor C-states:
  — Package: PC0, PC1/PC1E, PC2, PC3, PC6 (Package C7 is not supported)
  — Core: CC0, CC1, CC1E, CC3, CC6, CC7
- Enhanced Intel SpeedStep® Technology

1.3.2 System States Support

- S0, S1, S3, S4, S5

1.3.3 Memory Controller

- Multiple CKE power down modes
- Multiple self-refresh modes

- Memory thermal monitoring via MEM_HOT_C1_N and MEM_HOT_C23_N Signals

1.3.4 PCI Express

- L1 ASPM power management capability

1.3.5 Intel QuickPath Interconnect

- L0p and L1 power management capabilities

1.4 Thermal Management Support

- Digital Thermal Sensor with multiple on-die temperature zones
- Adaptive Thermal Monitor
- THERMTRIP_N and PROCHOT_N signal support
- On-Demand mode clock modulation
- Open and Closed Loop Thermal Throttling (OLTT/CLTT) support for system memory in addition to Hybrid OLTT/CLTT mode
- Fan speed control with DTS
- Two integrated SMBus masters for accessing thermal data from DIMMs
- New Memory Thermal Throttling features via MEM_HOT_C{1/23}_N signals
- Running Average Power Limit (RAPL), Processor and DRAM Thermal and Power Optimization Capabilities

1.5 Package Summary

The Processor socket type is noted as Socket B2. It is a 45 mm x 42.5 mm FCLGA10 package (LGA1356).

1.6 Terminology

Term	Description
ASPM	Active State Power Management
BMC	Baseboard Management Controllers
Cbo	Cache and Core Box. It is a term used for internal logic providing ring interface to LLC and Core.
DDR3	Third generation Double Data Rate SDRAM memory technology that is the successor to DDR2 SDRAM
DMA	Direct Memory Access
DMI	Direct Media Interface
DMI2	Direct Media Interface Gen 2
DTS	Digital Thermal Sensor
ECC	Error Correction Code
Enhanced Intel SpeedStep® Technology	Allows the operating system to reduce power consumption when performance is not needed.

Term	Description
Execute Disable Bit	The Execute Disable bit allows memory to be marked as executable or non-executable, when combined with a supporting operating system. If code attempts to run in non-executable memory the processor raises an error to the operating system. This feature can prevent some classes of viruses or worms that exploit buffer overrun vulnerabilities and can thus help improve the overall security of the system. See the *Intel® 64 and IA-32 Architectures Software Developer's Manuals* for more detailed information.
Flit	Flow Control Unit. The Intel QPI Link layer’s unit of transfer; 1 Flit = 80-bits.
Functional Operation	Refers to the normal operating conditions in which all processor specifications, including DC, AC, system bus, signal quality, mechanical, and thermal, are satisfied.
IMC	The Integrated Memory Controller. A Memory Controller that is integrated in the processor die.
IIO	The Integrated I/O Controller. An I/O controller that is integrated in the processor die.
Intel® ME	Intel® Management Engine (Intel® ME)
Intel® QuickData Technology	Intel QuickData Technology is a platform solution designed to maximize the throughput of server data traffic across a broader range of configurations and server environments to achieve faster, scalable, and more reliable I/O.
Intel® QuickPath Interconnect (Intel® QPI)	A cache-coherent, link-based Interconnect specification for Intel processors, chipsets, and I/O bridge components.
Intel® 64 Technology	64-bit memory extensions to the IA-32 architecture. Further details on Intel 64 architecture and programming model can be found at http://developer.intel.com/technology/intel64/.
Intel® Turbo Boost Technology	Intel® Turbo Boost Technology is a way to automatically run the processor core faster than the marked frequency if the part is operating under power, temperature, and current specifications limits of the Thermal Design Power (TDP). This results in increased performance of both single and multi-threaded applications.
Intel® TXT	Intel® Trusted Execution Technology
Intel® Virtualization Technology (Intel® VT)	Processor virtualization which when used in conjunction with Virtual Machine Monitor software enables multiple, robust independent software environments inside a single platform.
Intel® VT-d	Intel® Virtualization Technology (Intel® VT) for Directed I/O. Intel VT-d is a hardware assist, under system software (Virtual Machine Manager or OS) control, for enabling I/O device virtualization. Intel VT-d also brings robust security by providing protection from errant DMAs by using DMA remapping, a key feature of Intel VT-d.
Intel® Xeon® processor E5-2400 product family	Intel’s 32-nm processor design for use in Intel® Xeon® processor E5-2400 product family-based platforms.
Integrated Heat Spreader (IHS)	A component of the processor package used to enhance the thermal performance of the package. Component thermal solutions interface with the processor at the IHS surface.
Jitter	Any timing variation of a transition edge or edges from the defined Unit Interval (UI).
IOV	I/O Virtualization
LGA1356 Socket	The 1356-land FCLGA package mates with the system board through this surface mount, 1356-contact socket.
LLC	Last Level Cache
LRDIMM	Load Reduced Dual In-line Memory Module
NCTF	Non-Critical to Function: NCTF locations are typically redundant ground or non-critical reserved, so the loss of the solder joint continuity at end of life conditions will not affect the overall product functionality.
NEBS	Network Equipment Building System. NEBS is the most common set of environmental design guidelines applied to telecommunications equipment in the United States.

Term	Description
PCH	Platform Controller Hub (Intel® C600 Series Chipset). The next generation chipset with centralized platform capabilities including the main I/O interfaces along with display connectivity, audio features, power management, manageability, security and storage features.
PCU	Power Control Unit
PCI Express* 3.0	The third generation PCI Express* specification that operates at twice the speed of PCI Express* 2.0 (8 Gb/s); however, PCI Express* 3.0 is completely backward compatible with PCI Express* 1.0 and 2.0.
PCI Express* 3	PCI Express* Generation 3.0
PCI Express* 2	PCI Express* Generation 2.0
PCI Express*	PCI Express* Generation 2.0/3.0
PECI	Platform Environment Control Interface
Phit	Physical Unit. An Intel® QPI terminology defining units of transfer at the physical layer. 1 Phit is equal to 20 bits in 'full width mode' and 10 bits in 'half width mode'
Processor	The 64-bit, single-core or multi-core component (package)
Processor Core	The term "processor core" refers to silicon die itself which can contain multiple execution cores. Each execution core has an instruction cache, data cache, and 256-KB L2 cache. All execution cores share the L3 cache. All DC and signal integrity specifications are measured at the processor die (pads), unless otherwise noted.
RDIMM	Registered Dual In-line Memory Module
Rank	A unit of DRAM corresponding four to eight devices in parallel, ignoring ECC. These devices are usually, but not always, mounted on a single side of a DDR3 DIMM.
Scalable-2S	Intel® Xeon® processor E5-2400 product family-based Platform targeted for scalable designs using third party Node Controller chip. In these designs, Node Controller is used to scale the design beyond one/two/four sockets.
SCI	System Control Interrupt. Used in ACPI protocol.
SSE	Intel® Streaming SIMD Extensions (Intel® SSE)
SKU	A processor Stock Keeping Unit (SKU) to be installed in either server or workstation platforms. Electrical, power and thermal specifications for these SKU's are based on specific use condition assumptions. Server processors may be further categorized as Efficient Performance server, workstation and HPC SKUs. For further details on use condition assumptions, please refer to the latest Product Release Qualification (PRQ) Report available via your Customer Quality Engineer (CQE) contact.
SMBus	System Management Bus. A two-wire interface through which simple system and power management related devices can communicate with the rest of the system. It is based on the principals of the operation of the I2C* two-wire serial bus from Philips Semiconductor.
Storage Conditions	A non-operational state. The processor may be installed in a platform, in a tray, or loose. Processors may be sealed in packaging or exposed to free air. Under these conditions, processor landings should not be connected to any supply voltages, have any I/Os biased or receive any clocks. Upon exposure to "free air" (i.e., unsealed packaging or a device removed from packaging material) the processor must be handled in accordance with moisture sensitivity labeling (MSL) as indicated on the packaging material.
TAC	Thermal Averaging Constant
TDP	Thermal Design Power
TSOD	Thermal Sensor on DIMM
UDIMM	Unbuffered Dual In-line Module
Uncore	The portion of the processor comprising the shared cache, IMC, HA, PCU, UBox, and Intel QPI link interface.

Term	Description
Unit Interval	Signaling convention that is binary and unidirectional. In this binary signaling, one bit is sent for every edge of the forwarded clock, whether it be a rising edge or a falling edge. If a number of edges are collected at instances t_1, t_2, t_n,...., t_k then the UI at instance "n" is defined as: $UI_n = t_n - t_{n-1}$
V_{CC}	Processor core power supply
V_{SS}	Processor ground
V_{CCD}	Variable power supply for the processor system memory interface.
x1	Refers to a Link or Port with one Physical Lane
x4	Refers to a Link or Port with four Physical Lanes
x8	Refers to a Link or Port with eight Physical Lanes
x16	Refers to a Link or Port with sixteen Physical Lanes

1.7 Related Documents

Refer to the following documents for additional information.

Table 1-1. Referenced Documents

Document	Location
Intel® Xeon® Processor E5 Product Family Datasheet Volume Two	http://www.intel.com
Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide	http://www.intel.com
Intel® Xeon® Processor E5-2400 Product Family– Boundary Scan Description Language (BSDL) File	http://www.intel.com
Intel® C600 Series Chipset Data Sheet	http://www.intel.com
Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3	http://www.intel.com
Advanced Configuration and Power Interface Specification 3.0	http://www.acpi.info
PCI Local Bus Specification 3.0	http://www.pcisig.com/specifications
PCI Express Base Specification - Revision 2.1 and 1.1 *PCI Express Base Specification - Revision 3.0*	http://www.pcisig.com
System Management Bus (SMBus) Specification	http://smbus.org/
DDR3 SDRAM Specification	http://www.jedec.org
Low (JESD22-A119) and High (JESD-A103) Temperature Storage Life Specifications	http://www.jedec.org
Intel 64 and IA-32 Architectures Software Developer's Manuals • Volume 1: Basic Architecture • Volume 2A: Instruction Set Reference, A-M • Volume 2B: Instruction Set Reference, N-Z • Volume 3A: System Programming Guide • Volume 3B: System Programming Guide *Intel® 64 and IA-32 Architectures Optimization Reference Manual*	http://www.intel.com/products/processor/manuals/index.htm
Intel® Virtualization Technology Specification for Directed I/O Architecture Specification	http://download.intel.com/technology/computing/vptech/Intel(r)_VT_for_Direct_IO.pdf
Intel® Trusted Execution Technology Software Development Guide	http://www.intel.com/technology/security/

1.8 State of Data

The data contained within this document is the most accurate information available by the publication date of this document.

§

2 Interfaces

This chapter describes the interfaces supported by the processor.

2.1 System Memory Interface

2.1.1 System Memory Technology Support

The Integrated Memory Controller (IMC) supports DDR3 protocols with three independent 64-bit memory channels with 8 bits of ECC for each channel (total of 72-bits) and supports 2 DIMMs per channel. The type of memory supported by the processor is dependent on the target platform:

- Intel® Xeon® processor E5-2400 product family-based platforms support:
  - — ECC registered DIMMs: with a maximum of two DIMMs per channel allowing up to eight device ranks per channel.
  - — ECC and non-ECC unbuffered DIMMs: with a maximum of two DIMMs per channel thus allowing up to four device ranks per channel. Support for mixed non-ECC with ECC un-buffered DIMM configurations.

2.1.2 System Memory Timing Support

The IMC supports the following DDR3 Speed Bin, CAS Write Latency (CWL), and command signal mode timings on the main memory interface:

- tCL = CAS Latency
- tRCD = Activate Command to READ or WRITE Command delay
- tRP = PRECHARGE Command Period
- CWL = CAS Write Latency
- Command Signal modes = 1n indicates a new command may be issued every clock and 2n indicates a new command may be issued every 2 clocks. Command launch mode programming depends on the transfer rate and memory configuration.

2.2 PCI Express* Interface

This section describes the PCI Express* 3.0 interface capabilities of the processor. See the *PCI Express* Base Specification* for details of PCI Express* 3.0.

2.2.1 PCI Express* Architecture

Compatibility with the PCI addressing model is maintained to ensure that all existing applications and drivers operate unchanged. The PCI Express* configuration uses standard mechanisms as defined in the PCI Plug-and-Play specification.

The PCI Express* architecture is specified in three layers: Transaction Layer, Data Link Layer, and Physical Layer. The partitioning in the component is not necessarily along these same boundaries. Refer to Figure 2-1 for the PCI Express* Layering Diagram.

Figure 2-1. PCI Express* Layering Diagram

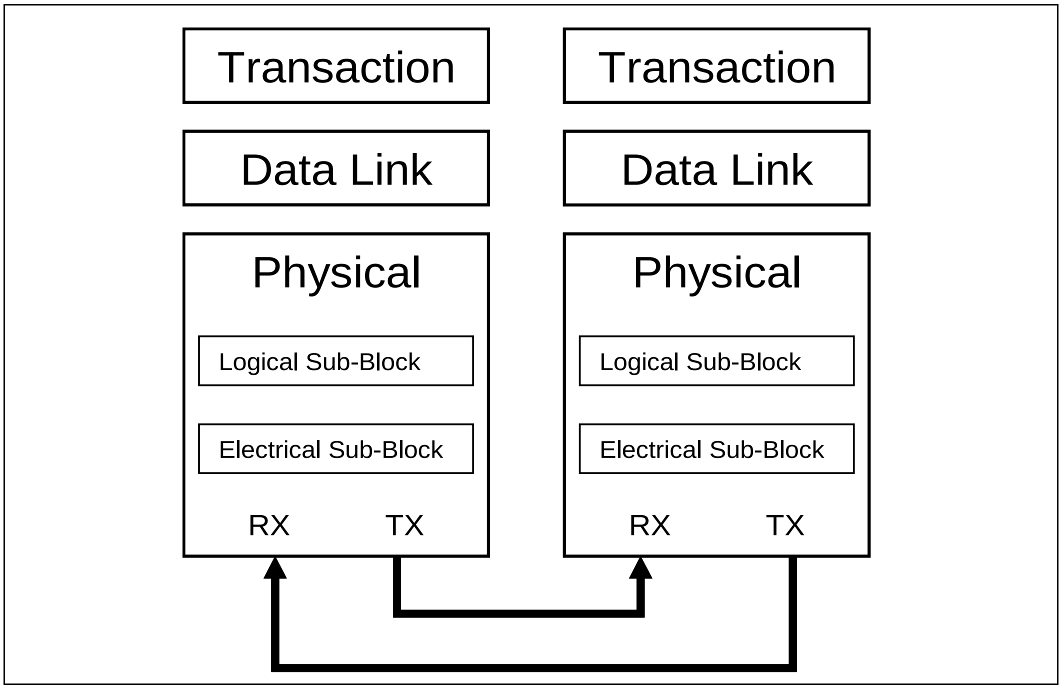

PCI Express* uses packets to communicate information between components. Packets are formed in the Transaction and Data Link Layers to carry the information from the transmitting component to the receiving component. As the transmitted packets flow through the other layers, they are extended with additional information necessary to handle packets at those layers. At the receiving side, the reverse process occurs and packets get transformed from their Physical Layer representation to the Data Link Layer representation and finally (for Transaction Layer Packets) to the form that can be processed by the Transaction Layer of the receiving device.

Figure 2-2. Packet Flow through the Layers

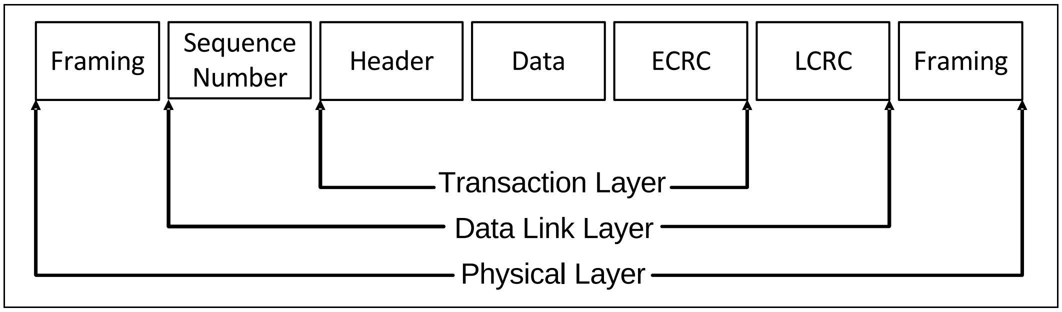

2.2.1.1 Transaction Layer

The upper layer of the PCI Express* architecture is the Transaction Layer. The Transaction Layer's primary responsibility is the assembly and disassembly of Transaction Layer Packets (TLPs). TLPs are used to communicate transactions, such as read and write, as well as certain types of events. The Transaction Layer also manages flow control of TLPs.

2.2.1.2 Data Link Layer

The middle layer in the PCI Express* stack, the Data Link Layer, serves as an intermediate stage between the Transaction Layer and the Physical Layer. Responsibilities of Data Link Layer include link management, error detection, and error correction.

The transmission side of the Data Link Layer accepts TLPs assembled by the Transaction Layer, calculates and applies data protection code and TLP sequence number, and submits them to Physical Layer for transmission across the Link. The receiving Data Link Layer is responsible for checking the integrity of received TLPs and for submitting them to the Transaction Layer for further processing. On detection of TLP error(s), this layer is responsible for requesting retransmission of TLPs until information is correctly received, or the Link is determined to have failed. The Data Link Layer also generates and consumes packets which are used for Link management functions.

2.2.1.3 Physical Layer

The Physical Layer includes all circuitry for interface operation, including driver and input buffers, parallel-to-serial and serial-to-parallel conversion, PLL(s), and impedance matching circuitry. It also includes logical functions related to interface initialization and maintenance. The Physical Layer exchanges data with the Data Link Layer in an implementation-specific format, and is responsible for converting this to an appropriate serialized format and transmitting it across the PCI Express* Link at a frequency and width compatible with the remote device.

2.2.2 PCI Express* Configuration Mechanism

The PCI Express* link is mapped through a PCI-to-PCI bridge structure.

PCI Express* extends the configuration space to 4096 bytes per-device/function, as compared to 256 bytes allowed by the *Conventional PCI Specification*. PCI Express* configuration space is divided into a PCI-compatible region (which consists of the first

256 bytes of a logical device's configuration space) and an extended PCI Express* region (which consists of the remaining configuration space). The PCI-compatible region can be accessed using either the mechanisms defined in the PCI specification or using the enhanced PCI Express* configuration access mechanism described in the PCI Express* Enhanced Configuration Mechanism section.

The PCI Express* Host Bridge is required to translate the memory-mapped PCI Express* configuration space accesses from the host processor to PCI Express* configuration cycles. To maintain compatibility with PCI configuration addressing mechanisms, it is recommended that system software access the enhanced configuration space using 32-bit operations (32-bit aligned) only.

See the *PCI Express* Base Specification* for details of both the PCI-compatible and PCI Express* Enhanced configuration mechanisms and transaction rules.

2.3 DMI2/PCI Express* Interface

Direct Media Interface 2 (DMI2) connects the processor to the Platform Controller Hub (PCH). DMI2 is similar to a four-lane PCI Express* supporting a speed of 5 GT/s per lane. This interface can be configured at power-on to serve as a x4 PCI Express* link based on the setting of the SOCKET_ID[1:0] and FRMAGENT signal for processors not connected to a PCH.

Note: Only DMI2 x4 configuration is supported.

2.3.1 DMI2 Error Flow

DMI2 can only generate SERR in response to errors, never SCI, SMI, MSI, PCI INT, or GPE. Any DMI2 related SERR activity is associated with Device 0.

2.3.2 Processor/PCH Compatibility Assumptions

The processor is compatible with the PCH and is not compatible with any previous MCH or ICH products.

2.3.3 DMI2 Link Down

The DMI2 link going down is a fatal, unrecoverable error. If the DMI2 data link goes to data link down, after the link was up, then the DMI2 link hangs the system by not allowing the link to retrain to prevent data corruption. This is controlled by the PCH.

Downstream transactions that had been successfully transmitted across the link prior to the link going down may be processed as normal. No completions from downstream, non-posted transactions are returned upstream over the DMI2 link after a link down event.

2.4 Intel QuickPath Interconnect

The Intel QuickPath Interconnect is a high speed, packetized, point-to-point interconnect used in the 2nd Generation Intel® Core™ Processor Family. The narrow high-speed links stitch together processors in distributed shared memory and integrated I/O platform architecture. It offers much higher bandwidth with low latency. The Intel QuickPath Interconnect has an efficient architecture allowing more interconnect performance to be achieved in real systems. It has a snoop protocol optimized for low latency and high scalability, as well as packet and lane structures enabling quick completions of transactions. Reliability, availability, and serviceability features (RAS) are built into the architecture.

The physical connectivity of each interconnect link is made up of twenty differential signal pairs plus a differential forwarded clock. Each port supports a link pair consisting of two uni-directional links to complete the connection between two components. This supports traffic in both directions simultaneously. To facilitate flexibility and longevity, the interconnect is defined as having five layers: Physical, Link, Routing, Transport, and Protocol.

- **The Physical layer** consists of the actual wires carrying the signals, as well as circuitry and logic to support ancillary features required in the transmission and receipt of the 1s and 0s. The unit of transfer at the Physical layer is 20-bits, which is called a Phit (for Physical unit).
- **The Link layer** is responsible for reliable transmission and flow control. The Link layer's unit of transfer is 80-bits, which is called a Flit (for Flow control unit).
- **The Routing layer** provides the framework for directing packets through the fabric.
- **The Transport layer** is an architecturally defined layer (not implemented in the initial products) providing advanced routing capability for reliable end-to-end transmission.
- **The Protocol layer** is the high-level set of rules for exchanging packets of data between devices. A packet is comprised of an integral number of Flits.

The Intel QuickPath Interconnect includes a cache coherency protocol to keep the distributed memory and caching structures coherent during system operation. It supports both low-latency source snooping and a scalable home snoop behavior. The coherency protocol provides for direct cache-to-cache transfers for optimal latency.

2.5 Platform Environment Control Interface (PECI)

The Platform Environment Control Interface (PECI) uses a single wire for self-clocking and data transfer. The bus requires no additional control lines. The physical layer is a self-clocked one-wire bus that begins each bit with a driven, rising edge from an idle level near zero volts. The duration of the signal driven high depends on whether the bit value is a logic '0' or logic '1'. PECI also includes variable data transfer rate established with every message. In this way, it is highly flexible even though underlying logic is simple.

The interface design was optimized for interfacing to Intel processor and chipset components in both single processor and multiple processor environments. The single wire interface provides low board routing overhead for the multiple load connections in the congested routing area near the processor and chipset components. Bus speed, error checking, and low protocol overhead provides adequate link bandwidth and reliability to transfer critical device operating conditions and configuration information.

The PECI bus offers:

- A wide speed range from 2 Kbps to 2 Mbps
- CRC check byte used to efficiently and atomically confirm accurate data delivery
- Synchronization at the beginning of every message minimizes device timing accuracy requirements

Note: The PECI commands described in this document apply primarily to the Intel® Xeon® processor E5-2400 product family. The processors utilizes the capabilities described in this document to indicate support for three memory channels. Refer to Table 2-1 for the list of PECI commands supported by the processors.

Table 2-1. Summary of Processor-specific PECI Commands

Command	Supported on the Processor
Ping()	Yes
GetDIB()	Yes
GetTemp()	Yes
RdPkgConfig()	Yes
WrPkgConfig()	Yes
RdIAMSR()	Yes
WrIAMSR()	No
RdPCIConfig()	Yes
WrPCIConfig()	No
RdPCIConfigLocal()	Yes
WrPCIConfigLocal()	Yes

2.5.1 PECI Client Capabilities

The processor PECI client is designed to support the following sideband functions:

- Processor and DRAM thermal management
- Platform manageability functions including thermal, power, and error monitoring
  - The platform 'power' management includes monitoring and control for both the processor and DRAM subsystem to assist with data center power limiting.
- Processor interface tuning and diagnostics capabilities (Intel® Interconnect BIST).

2.5.1.1 Thermal Management

Processor fan speed control is managed by comparing Digital Thermal Sensor (DTS) thermal readings acquired via PECI against the processor-specific fan speed control reference point, or $T_{CONTROL}$. Both $T_{CONTROL}$ and DTS thermal readings are accessible via the processor PECI client. These variables are referenced to a common temperature, the TCC activation point, and are both defined as negative offsets from that reference.

PECI-based access to the processor package configuration space provides a means for Baseboard Management Controllers (BMCs) or other platform management devices to actively manage the processor and memory power and thermal features. Details on the list of available power and thermal optimization services can be found in Section 2.5.2.6.

2.5.1.2 Platform Manageability

PECI allows read access to certain error registers in the processor MSR space and status monitoring registers in the PCI configuration space within the processor and downstream devices. Details are covered in subsequent sections.

PECI permits writes to certain Memory Controller RAS-related registers in the processor PCI configuration space. Details are covered in Section 2.5.2.10.

2.5.1.3 Processor Interface Tuning and Diagnostics

The processor Intel® Interconnect Built In Self Test (Intel® IBIST) allows for in-field diagnostic capabilities in the Intel® QPI and memory controller interfaces. PECI provides a port to execute these diagnostics via its PCI Configuration read and write capabilities in the BMC INIT mode. Refer to Section 2.5.3.7 for more details.

2.5.2 Client Command Suite

PECI command requires at least one frame check sequence (FCS) byte to ensure reliable data exchange between originator and client. The PECI message protocol defines two FCS bytes that are returned by the client to the message originator. The first FCS byte covers the client address byte, the Read and Write Length bytes, and all bytes in the write data block. The second FCS byte covers the read response data returned by the PECI client. The FCS byte is the result of a cyclic redundancy check (CRC) of each data block.

2.5.2.1 Ping()

Ping() is a required message for all PECI devices. This message is used to enumerate devices or determine if a device has been removed, been powered-off, etc. A Ping() sent to a device address always returns a non-zero Write FCS if the device at the targeted address is able to respond.

2.5.2.1.1 Command Format

The Ping() format is as follows:

Write Length: 0x00

Read Length: 0x00

Figure 2-3. Ping()

Byte #	0	1	2	3
Byte Definition	Client Address	Write Length 0x00	Read Length 0x00	FCS

An example Ping() command to PECI device address 0x30 is shown below.

Figure 2-4. Ping() Example

Byte #	0	1	2	3
Byte Definition	0x30	0x00	0x00	0xe1

2.5.2.2 GetDIB()

The processor PECI client implementation of GetDIB() includes an 8-byte response and provides information regarding client revision number and the number of supported domains. All processor PECI clients support the GetDIB() command.

2.5.2.2.1 Command Format

The GetDIB() format is as follows:

Write Length: 0x01

Read Length: 0x08

Command: 0xf7

Figure 2-5. GetDIB()

Byte #	0	1	2	3	4
Byte Definition	Client Address	Write Length 0x01	Read Length 0x08	Cmd Code 0xf7	FCS
	5	6	7	8	9
	Device Info	Revision Number	Reserved	Reserved	Reserved
	10	11	12	13	
	Reserved	Reserved	Reserved	FCS	

2.5.2.2.2 Device Info

The Device Info byte gives details regarding the PECI client configuration. At a minimum, all clients supporting GetDIB will return the number of domains inside the package via this field. With any client, at least one domain (Domain 0) must exist. Therefore, the Number of Domains reported is defined as the number of domains in addition to Domain 0. For example, if bit 2 of the Device Info byte returns a ‘1’, that would indicate that the PECI client supports two domains.

Figure 2-6. Device Info Field Definition

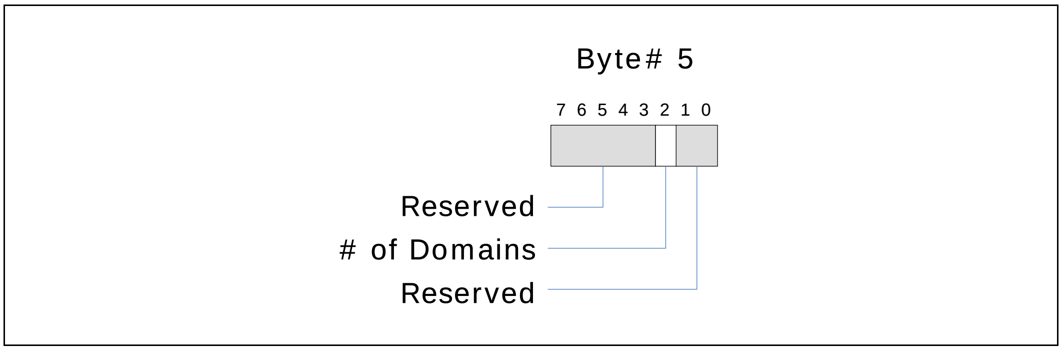

2.5.2.2.3 Revision Number

All clients that support the GetDIB command also support Revision Number reporting. The revision number may be used by a host or originator to manage different command suites or response codes from the client. Revision Number is always reported in the second byte of the GetDIB() response. The ‘Major Revision’ number in Figure 2-7 always maps to the revision number of the PECI specification that the PECI client processor is designed to. The ‘Minor Revision’ number value depends on the exact command suite supported by the PECI client as defined in Table 2-2.

Figure 2-7. Revision Number Definition

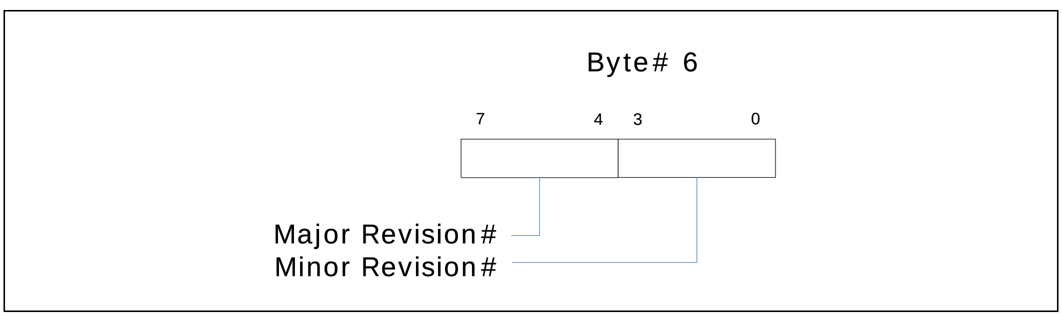

Table 2-2. Minor Revision Number Meaning

Minor Revision	Supported Command Suite
0	Ping(), GetDIB(), GetTemp()
1	Ping(), GetDIB(), GetTemp(), WrPkgConfig(), RdPkgConfig()
2	Ping(), GetDIB(), GetTemp(), WrPkgConfig(), RdPkgConfig(), RdIAMSR()
3	Ping(), GetDIB(), GetTemp(), WrPkgConfig(), RdPkgConfig(), RdIAMSR(), RdPCIConfigLocal(), WrPCIConfigLocal()
4	Ping(), GetDIB(), GetTemp(), WrPkgConfig(), RdPkgConfig(), RdIAMSR(), RdPCIConfigLocal(), WrPCIConfigLocal(), RdPCIConfig()

Table 2-2. Minor Revision Number Meaning

Minor Revision	Supported Command Suite
5	Ping(), GetDIB(), GetTemp(), WrPkgConfig(), RdPkgConfig(), RdIAMSR(), RdPCIConfigLocal(), WrPCIConfigLocal(), RdPCIConfig(), WrPCIConfig()
6	Ping(), GetDIB(), GetTemp(), WrPkgConfig(), RdPkgConfig(), RdIAMSR(), RdPCIConfigLocal(), WrPCIConfigLocal(), RdPCIConfig(), WrPCIConfig(), WrIAMSR()

For the processor PECI client the Revision Number will return ‘0011 0100b’.

2.5.2.3 GetTemp()

The GetTemp() command is used to retrieve the maximum die temperature from a target PECI address. The temperature is used by the external thermal management system to regulate the temperature on the die. The data is returned as a negative value representing the number of degrees centigrade below the maximum processor junction temperature (T_{jmax}). The maximum PECI temperature value of zero corresponds to the processor T_{jmax}. This also represents the default temperature at which the processor Thermal Control Circuit activates. The actual value that the thermal management system uses as a control set point ($T_{CONTROL}$) is also defined as a negative number below T_{jmax}. $T_{CONTROL}$ may be extracted from the processor by issuing a PECI RdPkgConfig() command as described in Section 2.5.2.4 or using a RDMSR instruction. $T_{CONTROL}$ application to fan speed control management is defined in the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

Please refer to Section 2.5.7 for details regarding PECI temperature data formatting.

2.5.2.3.1 Command Format

The GetTemp() format is as follows:

Write Length: 0x01

Read Length: 0x02

Command: 0x01

Description: Returns the highest die temperature for addressed processor PECI client.

Figure 2-8. GetTemp()

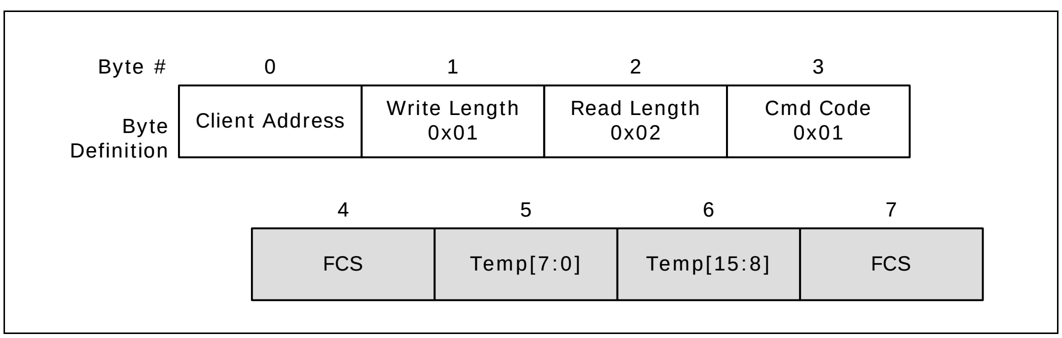

Example bus transaction for a thermal sensor device located at address 0x30 returning a value of negative 10 counts is show in Figure 2-9.

Figure 2-9. GetTemp() Example

Byte #	0	1	2	3
Byte Definition	0x30	0x01	0x02	0x01

Byte #	4	5	6	7
Byte Definition	0xef	0x80	0xfd	0x4b

2.5.2.3.2 Supported Responses

The typical client response is a passing FCS and valid thermal data. Under some conditions, the client's response will indicate a failure. GetTemp() response definitions are listed in Table 2-3. Refer to Section 2.5.7.4 for more details on sensor errors.

Table 2-3. GetTemp() Response Definition

Response	Meaning
General Sensor Error (GSE)[1]	Thermal scan did not complete in time. Retry is appropriate.
Bad Write FCS	Electrical error
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
0x0000[1]	Processor is running at its maximum temperature or is currently being reset.
All other data	Valid temperature reading, reported as a negative offset from the processor T_{jmax}.

Notes:
1. This response will be reflected in Bytes 5 & 6 in Figure 2-9.

2.5.2.4 RdPkgConfig()

The RdPkgConfig() command provides read access to the package configuration space (PCS) within the processor, including various power and thermal management functions. Typical PCS read services supported by the processor may include access to temperature data, energy status, run time information, DIMM temperatures and so on. Refer to Section 2.5.2.6 for more details on processor-specific services supported through this command.

2.5.2.4.1 Command Format

The RdPkgConfig() format is as follows:

Write Length: 0x05

Read Length: 0x05 (dword)

Command: 0xa1

Description: Returns the data maintained in the processor package configuration space for the PCS entry as specified by the 'index' and 'parameter' fields. The 'index' field contains the encoding for the requested service and is used in conjunction with the 'parameter' field to specify the exact data being requested. The Read Length dictates the desired data return size. This command supports only dword responses on the processor PECI clients. All command responses are prepended with a completion code that contains additional pass/fail status information. Refer to Section 2.5.5.2 for details regarding completion codes.

Figure 2-10. RdPkgConfig()

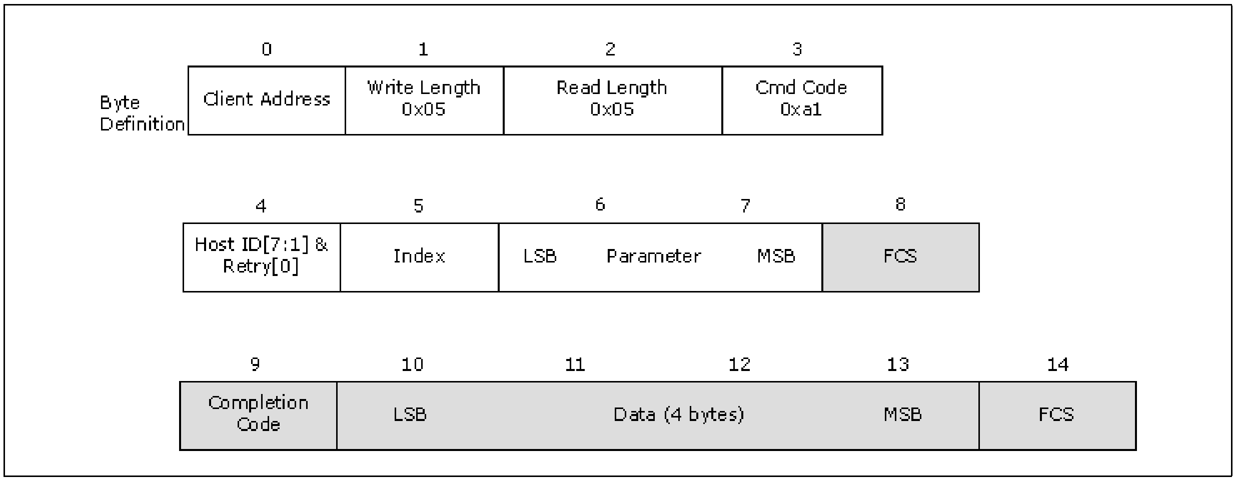

Note: The 2-byte parameter field and 4-byte read data field defined in Figure 2-10 are sent in standard PECI ordering with LSB first and MSB last.

2.5.2.4.2 Supported Responses

The typical client response is a passing FCS, a passing Completion Code and valid data. Under some conditions, the client's response will indicate a failure.

Table 2-4. RdPkgConfig() Response Definition

Response	Meaning
Bad Write FCS	Electrical error
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
CC: 0x40	Command passed, data is valid.
CC: 0x80	Response timeout. The processor is not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor is not able to allocate resources for servicing this command at this time. Retry is appropriate.
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.

2.5.2.5 WrPkgConfig()

The WrPkgConfig() command provides write access to the package configuration space (PCS) within the processor, including various power and thermal management functions. Typical PCS write services supported by the processor may include power limiting, thermal averaging constant programming and so on. Refer to Section 2.5.2.6 for more details on processor-specific services supported through this command.

2.5.2.5.1 Command Format

The WrPkgConfig() format is as follows:

Write Length: 0x0a(dword)

Read Length: 0x01

Command: 0xa5

AW FCS Support: Yes

Description: Writes data to the processor PCS entry as specified by the ‘index’ and ‘parameter’ fields. This command supports only dword data writes on the processor PECI clients. All command responses include a completion code that provides additional pass/fail status information. Refer to Section 2.5.5.2 for details regarding completion codes.

The Assured Write FCS (AW FCS) support provides the processor client a high degree of confidence that the data it received from the host is correct. This is especially critical where the consumption of bad data might result in improper or non-recoverable operation.

Figure 2-11. WrPkgConfig()

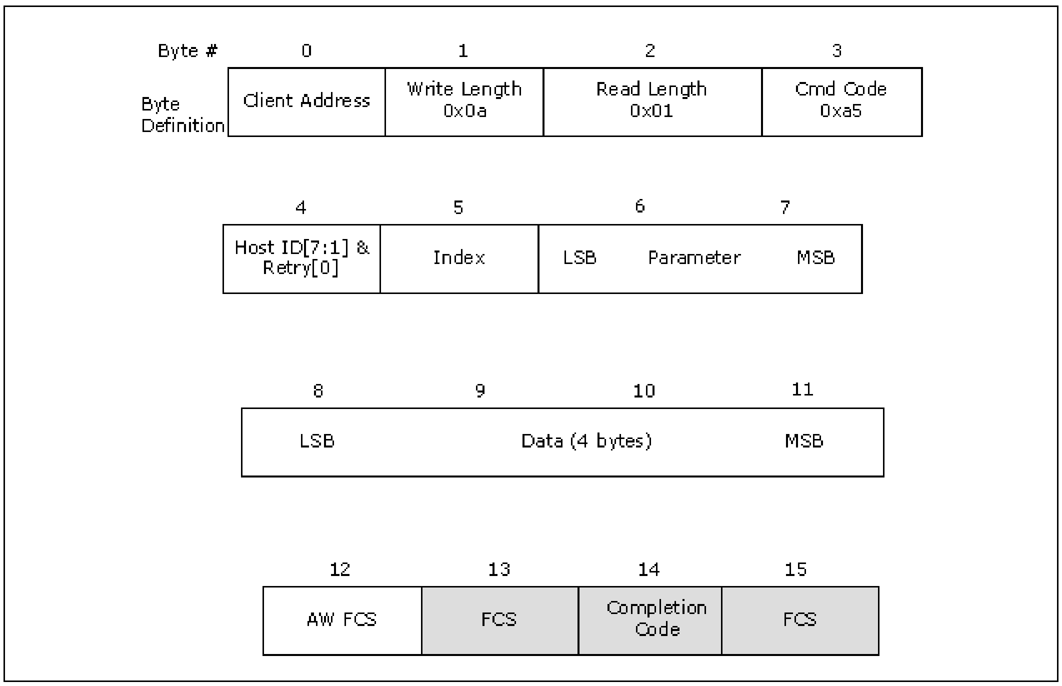

Note: The 2-byte parameter field and 4-byte write data field defined in Figure 2-11 are sent in standard PECI ordering with LSB first and MSB last.

2.5.2.5.2 Supported Responses

The typical client response is a passing FCS, a passing Completion Code and valid data. Under some conditions, the client’s response will indicate a failure.

Table 2-5. WrPkgConfig() Response Definition (Sheet 1 of 2)

Response	Meaning
Bad Write FCS	Electrical error or AW FCS failure
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
CC: 0x40	Command passed, data is valid.

Table 2-5. WrPkgConfig() Response Definition (Sheet 2 of 2)

Response	Meaning
CC: 0x80	Response timeout. The processor was not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor is not able to allocate resources for servicing this command at this time. Retry is appropriate.
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.

2.5.2.6 Package Configuration Capabilities

Table 2-6 combines both read and write services. Any service listed as a “read” would use the RdPkgConfig() command and a service listed as a “write” would use the WrPkgConfig() command. PECI requests for memory temperature or other data generated outside the processor package do not trigger special polling cycles on the processor memory or SMBus interfaces to procure the required information.

2.5.2.6.1 DRAM Thermal and Power Optimization Capabilities

DRAM thermal and power optimization (also known as RAPL or “Running Average Power Limit”) services provide a way for platform thermal management solutions to program and access DRAM power, energy and temperature parameters. Memory temperature information is typically used to regulate fan speeds, tune refresh rates and throttle the memory subsystem as appropriate. Memory temperature data may be derived from a variety of sources including on-die or on-board DIMM sensors, DRAM activity information or a combination of the two. Though memory temperature data is a byte long, range of actual temperature values are determined by the DIMM specifications and operating range.

Note: DRAM related PECI services described in this section apply only to the memory connected to the specific processor PECI client in question and not the overall platform memory in general. For estimating DRAM thermal information in closed loop throttling mode, a dedicated SMBus is required between the CPU and the DIMMs. The processor PCU requires access to the VR12 voltage regulator for reading average output current information through the SVID bus for initial DRAM RAPL related power tuning.

Table 2-6 provides a summary of the DRAM power and thermal optimization capabilities that can be accessed over PECI on the processor. **The Index values referenced in** Table 2-6 **are in decimal format.**

Table 2-6 also provides information on alternate inband mechanisms to access similar or equivalent information through register reads and writes where applicable. The user should consult the *Intel® 64 and IA-32 Architectures Software Developer’s Manual (SDM) Volumes 1, 2, and 3* or *Intel® Xeon® Processor E5 Product Family Datasheet Volume Two* for details on MSR and CSR register contents.

Table 2-6. RdPkgConfig() & WrPkgConfig() DRAM Thermal and Power Optimization Services Summary (Sheet 1 of 2)

Service	Index Value (decimal)	Parameter Value (word)	RdPkgConfig() Data (dword)	WrPkgConfig() Data (dword)	Description	Alternate Inband MSR or CSR Access
DRAM Rank Temperature Write	18	Channel Index & DIMM Index	N/A	Absolute temperature in Degrees Celsius for ranks 0, 1, 2 & 3	Write temperature for each rank within a single DIMM.	N/A
DIMM Temperature Read	14	Channel Index	Absolute temperature in Degrees Celsius for DIMMs 0, 1, & 2	N/A	Read temperature of each DIMM within a channel.	CSR: DIMMTEMPSTAT_[0:2]
DIMM Ambient Temperature Write / Read	19	0x0000	N/A	Absolute temperature in Degrees C to be used as ambient temperature reference	Write ambient temperature reference for activity-based rank temperature estimation.	N/A
DIMM Ambient Temperature Write / Read	19	0x0000	Absolute temperature in Degrees C to be used as ambient temperature reference	N/A	Read ambient temperature reference for activity-based rank temperature estimation.	N/A
DRAM Channel Temperature Read	22	0x0000	Maximum of all rank temperatures for each channel in Degrees Celsius	N/A	Read the maximum DRAM channel temperature.	N/A
Accumulated DRAM Energy Read	04	Channel Index 0x00FF - All Channels	DRAM energy consumed by the DIMMs	N/A	Read the DRAM energy consumed by all the DIMMs in all the channels or all the DIMMs within a specified channel.	MSR 619h: DRAM_ENERGY_STATUS CSR: DRAM_ENERGY_STATUS CSR: DRAM_ENERGY_STATUS_CH[0:3] [1]
DRAM Power Info Read	35	0x0000	Typical and minimum DRAM power settings	N/A	Read DRAM power settings info to be used by power limiting entity.	MSR 61Ch: DRAM_POWER_INFO CSR: DRAM_POWER_INFO
DRAM Power Info Read	36	0x0000	Maximum DRAM power settings & maximum time window	N/A	Read DRAM power settings info to be used by power limiting entity	MSR 61Ch: DRAM_POWER_INFO CSR: DRAM_POWER_INFO

Table 2-6. RdPkgConfig() & WrPkgConfig() DRAM Thermal and Power Optimization Services Summary (Sheet 2 of 2)

Service	Index Value (decimal)	Parameter Value (word)	RdPkgConfig() Data (dword)	WrPkgConfig() Data (dword)	Description	Alternate Inband MSR or CSR Access
DRAM Power Limit Data Write / Read	34	0x0000	N/A	DRAM Plane Power Limit Data	Write DRAM Power Limit Data	MSR 618h: DRAM_POWER_LIMIT CSR: DRAM_PLANE_POWER_LIMIT
DRAM Power Limit Data Write / Read	34	0x0000	DRAM Plane Power Limit Data	N/A	Read DRAM Power Limit Data	MSR 618h: DRAM_POWER_LIMIT CSR: DRAM_PLANE_POWER_LIMIT
DRAM Power Limit Performance Status Read	38	0x0000	Accumulated DRAM throttle time	N/A	Read sum of all time durations for which each DIMM has been throttled	CSR: DRAM_RAPL_PERF_STATUS

Notes:
1. Time, energy and power units should be assumed, where applicable, to be based on values returned by a read of the PACKAGE_POWER_SKU_UNIT MSR or through the Package Power SKU Unit PCS read service.
2. For the Intel® Xeon® processor E5-2400 product family, accumulated DRAM energy status would be reflected in the DRAM_ENERGY_STATUS_CH[1:3] CSR.

2.5.2.6.2 DRAM Thermal Estimation Configuration Data Read/Write

This feature is relevant only when activity-based DRAM temperature estimation methods are being utilized and would apply to all the DIMMs on all the memory channels. The write allows the PECI host to configure the ‘β’ and ‘θ’ variables in Figure 2-12 for DRAM channel temperature filtering as per the equation below:

$$T_N = \beta * T_{N-1} + \theta * \Delta Energy$$

T_N and T_{N-1} are the current and previous DRAM temperature estimates respectively in degrees Celsius, ‘β’ is the DRAM temperature decay factor, ‘ΔEnergy’ is the energy difference between the current and previous memory transactions as determined by the processor power control unit and ‘θ’ is the DRAM energy-to-temperature translation coefficient. The default value of ‘β’ is 0x3FF. ‘θ’ is defined by the equation:

$$\theta = (1 - \beta) * (\text{Thermal Resistance}) * (\text{Scaling Factor})$$

The ‘Thermal Resistance’ serves as a multiplier for translation of DRAM energy changes to corresponding temperature changes and may be derived from actual platform characterization data. The ‘Scaling Factor’ is used to convert memory transaction information to energy units in Joules and can be derived from system/memory configuration information. Refer to the *Intel® 64 and IA-32 Architectures Software Developer’s Manual (SDM) Volumes 1, 2, and 3* for methods to program and access ‘Scaling Factor’ information.

Figure 2-12. DRAM Thermal Estimation Configuration Data

31 – 20	19 – 10	9 – 0
RESERVED	THETA VARIABLE	BETA VARIABLE

Memory Thermal Estimation Configuration Data

2.5.2.6.3 DRAM Rank Temperature Write

This feature allows the PECI host to program into the processor, the temperature for all the ranks within a DIMM up to a maximum of four ranks as shown in Figure 2-13. The DIMM index and Channel index are specified through the parameter field as shown in Table 2-7. This write is relevant in platforms that do not have on-die or on-board DIMM thermal sensors to provide memory temperature information or if the processor does not have direct access to the DIMM thermal sensors. This temperature information is used by the processor in conjunction with the activity-based DRAM temperature estimations.

Table 2-7. Channel & DIMM Index Decoding

Index Encoding	Physical Channel#	Physical DIMM#
000	Reserved	0
001	1	1
010	2	Reserved
011	3	Reserved

Figure 2-13. DRAM Rank Temperature Write Data

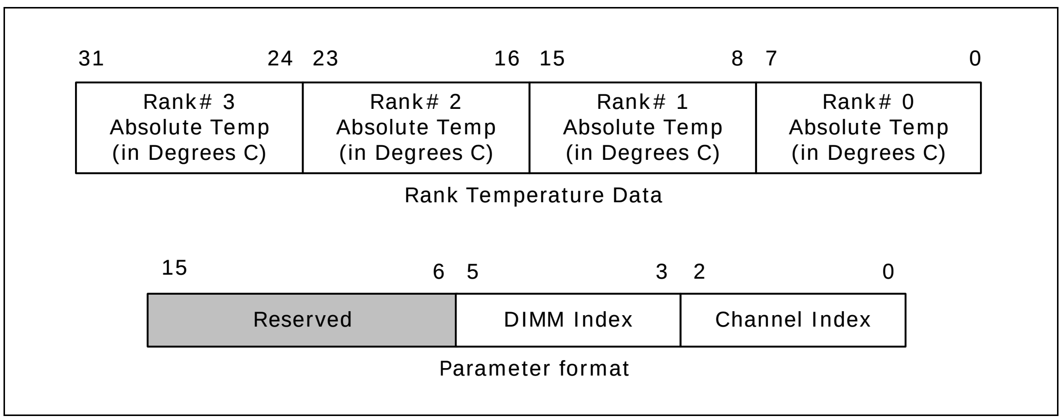

2.5.2.6.4 DIMM Temperature Read

This feature allows the PECI host to read the temperature of all the DIMMs within a channel up to a maximum of three DIMMs. This read is not limited to platforms using a particular memory temperature source or temperature estimation method. For platforms using DRAM thermal estimation, the PCU will provide the estimated temperatures. Otherwise, the data represents the latest DIMM temperature provided by the TSOD or on-board DIMM sensor and requires that CLTT (closed loop throttling mode) be enabled and OLTT (open loop throttling mode) be disabled. Refer to Table 2-7 for channel index encodings.

Figure 2-14. Intel® Xeon® Processor E5-2400 Product Families DIMM Temperature Read / Write

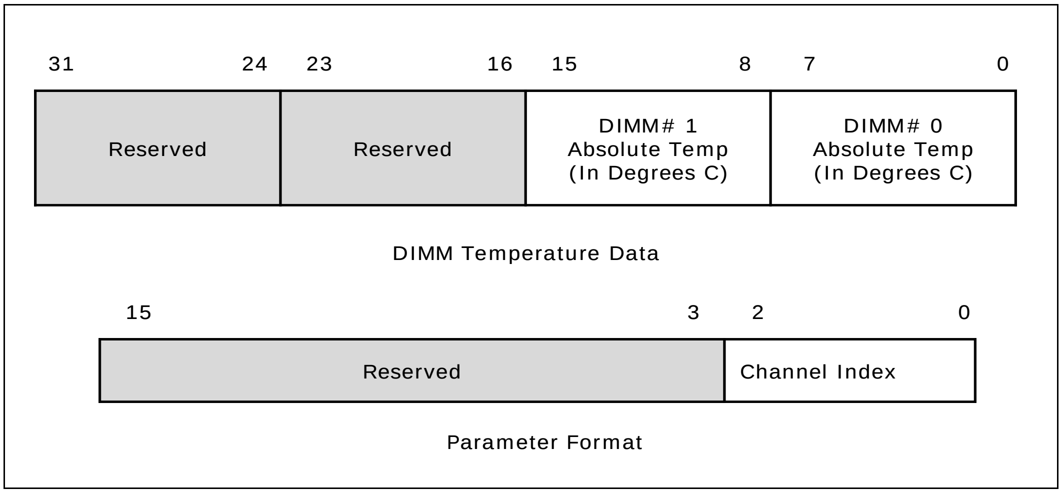

2.5.2.6.5 DIMM Ambient Temperature Write / Read

This feature allows the PECI host to provide an ambient temperature reference to be used by the processor for activity-based DRAM temperature estimation. This write is used only when no DIMM temperature information is available from on-board or on-die DIMM thermal sensors. It is also possible for the PECI host controller to read back the DIMM ambient reference temperature.

Since the ambient temperature may vary over time within a system, it is recommended that systems monitoring and updating the ambient temperature at a fast rate use the ‘maximum’ temperature value while those updating the ambient temperature at a slow rate use an ‘average’ value. The ambient temperature assumes a single value for all memory channel/DIMM locations and does not account for possible temperature variations based on DIMM location.

Figure 2-15. Ambient Temperature Reference Data

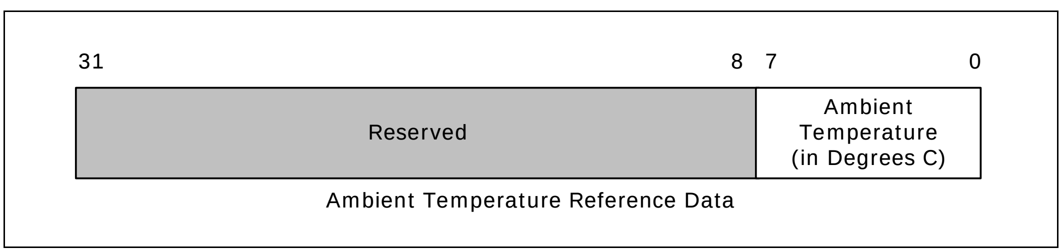

2.5.2.6.6 DRAM Channel Temperature Read

This feature enables a PECI host read of the maximum temperature of each channel. This would include all the DIMMs within the channel and all the ranks within each of the DIMMs. Channels that are not populated will return the ‘ambient temperature’ on systems using activity-based temperature estimations or alternatively return a ‘zero’ for systems using sensor-based temperatures.

Figure 2-16. Intel® Xeon® Processor E5-2400 Product Families DRAM Channel Temperature

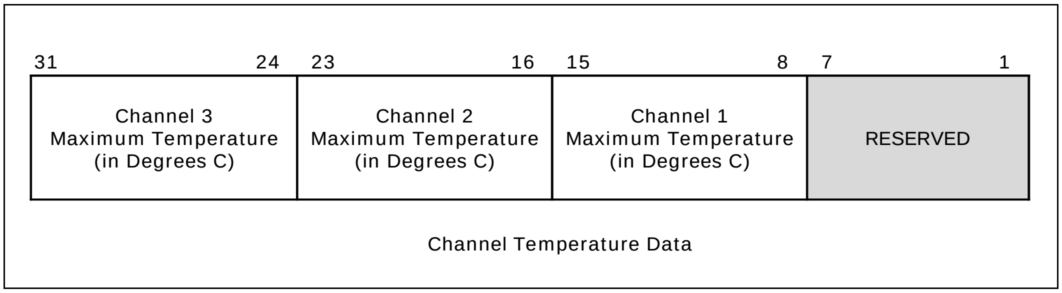

2.5.2.6.7 Accumulated DRAM Energy Read

This feature allows the PECI host to read the DRAM energy consumed by all the DIMMs within all the channels or all the DIMMs within just a specified channel. The parameter field is used to specify the channel index. Units used are defined as per the Package Power SKU Unit read described in Section 2.5.2.6.11. This information is tracked by a 32-bit counter that wraps around. The channel index in Figure 2-17 is specified as per the index encoding described in Table 2-7. A channel index of 0x00FF is used to specify the “all channels” case. While Intel requires reading the accumulated energy data at least once every 16 seconds to ensure functional correctness, a more realistic polling rate recommendation is once every 100 mS for better accuracy. This feature assumes a 200W memory capacity. In general, as the power capability decreases, so will the minimum polling rate requirement.

When determining energy changes by subtracting energy values between successive reads, Intel advocates using the 2’s complement method to account for counter wrap-arounds. Alternatively, adding all ‘F’s (‘0xFFFFFFFF’) to a negative result from the subtraction will accomplish the same goal.

Figure 2-17. Accumulated DRAM Energy Data

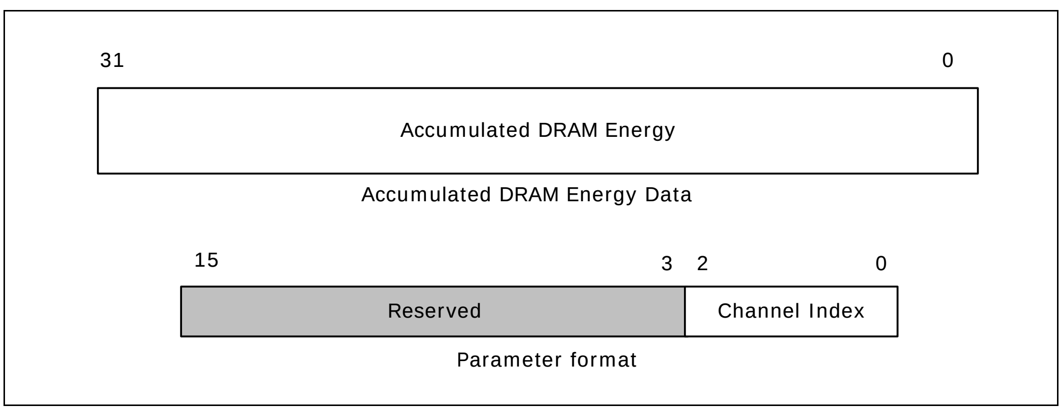

2.5.2.6.8 DRAM Power Info Read

This read returns the minimum, typical and maximum DRAM power settings and the maximum time window over which the power can be sustained for the entire DRAM domain and is inclusive of all the DIMMs within all the memory channels. Any power values specified by the power limiting entity that is outside of the range specified through these settings cannot be guaranteed. Since this data is 64 bits wide, PECI facilitates access to this register by allowing two requests to read the lower 32 bits and upper 32 bits separately as shown in Table 2-6. Power and time units for this read are defined as per the Package Power SKU Unit settings described in Section 2.5.2.6.11.

The minimum DRAM power in Figure 2-18 corresponds to a minimum bandwidth setting of the memory interface. It does 'not' correspond to a processor IDLE or memory self-refresh state. The 'time window' in Figure 2-18 is representative of the rate at which the power control unit (PCU) samples the DRAM energy consumption information and reactively takes the necessary measures to meet the imposed power limits. Programming too small a time window may not give the PCU enough time to sample energy information and enforce the limit while too large a time window runs the risk of the PCU not being able to monitor and take timely action on energy excursions. While the DRAM power setting in Figure 2-18 provides a maximum value for the 'time window' (typically a few seconds), the minimum value may be assumed to be ~100 mS.

The PCU programs the DRAM power settings described in Figure 2-18 when DRAM characterization has been completed by the memory reference code (MRC) during boot as indicated by the setting of the RST_CPL bit of the BIOS_RESET_CPL register. The DRAM power settings will be programmed during boot independent of the 'DRAM Power Limit Enable' bit setting. Please refer to the *Intel® Xeon® Processor E5 Product Family Datasheet Volume Two* for information on memory energy estimation methods and energy tuning options used by BIOS and other utilities for determining the range specified in the DRAM power settings. In general, any tuning of the power settings is done by polling the voltage regulators supplying the DIMMs.

Figure 2-18. DRAM Power Info Read Data

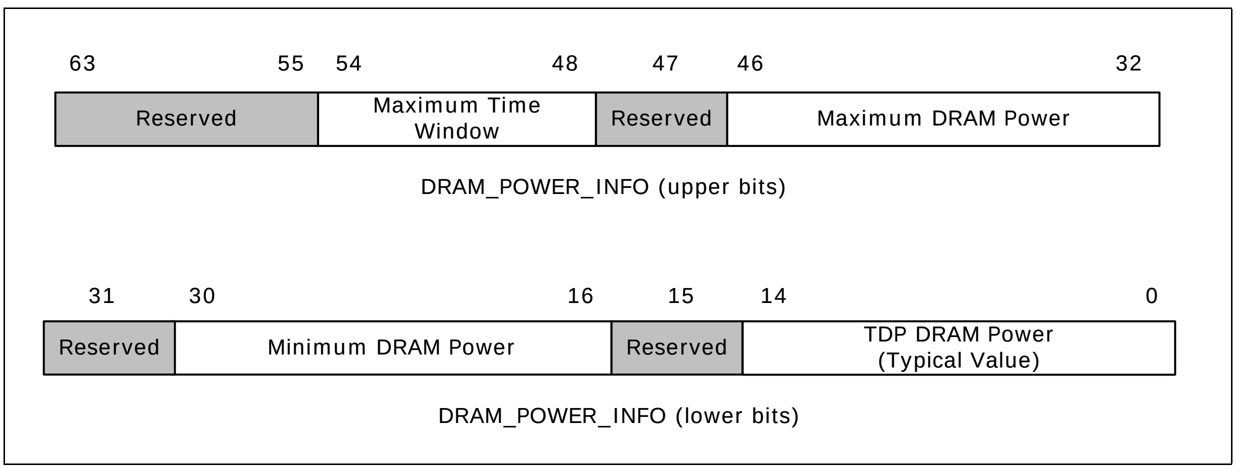

2.5.2.6.9 DRAM Power Limit Data Write / Read

This feature allows the PECI host to program the power limit over a specified time or control window for the entire DRAM domain covering all the DIMMs within all the memory channels. Actual values are chosen based on DRAM power consumption characteristics. The units for the DRAM Power Limit and Control Time Window are determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.11. The DRAM Power Limit Enable bit in Figure 2-19 should be set to activate this feature. Exact DRAM power limit values are largely determined by platform memory configuration. As such, this feature is disabled by default and there are no defaults associated with the DRAM power limit values. The PECI host may be used to enable and initialize the power limit fields for the purposes of DRAM power budgeting. Alternatively, this can also be accomplished through inband writes to the appropriate registers. Both power limit enabling and initialization of power limit values can be done in the same command cycle. All RAPL parameter values including the power limit value, control time window, and enable bit will have to be specified correctly even if the intent is to change just one parameter value when programming over PECI.

The following conversion formula should be used for encoding or programming the 'Control Time Window' in bits [23:17].

Control Time Window (in seconds) = ([1 + 0.25 * 'x'] * $2^{\text{'y'}}$) * 'z' where

'x' = integer value of bits[23:22]

'y' = integer value of bits[21:17]

'z' = Package Power SKU Time Unit[19:16] (see Section 2.5.2.6.13 for details on Package Power SKU Unit)

For example, using this formula, a control time value of 0x0A will correspond to a '1-second' time window. A valid range for the value of the 'Control Time Window' in Figure 2-19 that can be programmed into bits [23:17] is 250 mS - 40 seconds.

From a DRAM power management standpoint, all post-boot DRAM power management activities (also referred to as 'DRAM RAPL' or 'DRAM Running Average Power Limit') should be managed exclusively through a single interface like PECI or alternatively an inband mechanism. If PECI is being used to manage DRAM power budgeting activities, BIOS should lock out all subsequent inband DRAM power limiting accesses by setting bit 31 of the DRAM_POWER_LIMIT MSR or DRAM_PLANE_POWER_LIMIT CSR to '1'.

Figure 2-19. DRAM Power Limit Data

31 : 24	23 : 17	16	15	14 : 0
RESERVED	Control Time Window	RESERVED	DRAM Power Limit Enable	DRAM Power Limit

DRAM_POWER_LIMIT Data

2.5.2.6.10 DRAM Power Limit Performance Status Read

This service allows the PECI host to assess the performance impact of the currently active DRAM power limiting modes. The read return data contains the sum of all the time durations for which each of the DIMMs has been operating in a low power state. This information is tracked by a 32-bit counter that wraps around. The unit for time is determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.11. The DRAM performance data does not account for stalls on the memory interface.

In general, for the purposes of DRAM RAPL, the DRAM power management entity should use PECI accesses to DRAM energy and performance status in conjunction with the power limiting feature to budget power between the various memory sub-systems in the server system.

Figure 2-20. DRAM Power Limit Performance Data

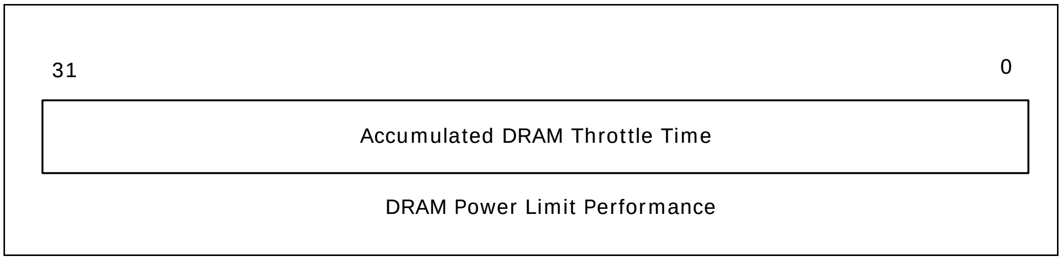

2.5.2.6.11 CPU Thermal and Power Optimization Capabilities

Table 2-8 provides a summary of the processor power and thermal optimization capabilities that can be accessed over PECI.

Note: **The Index values referenced in Table 2-8 are in decimal format.**

Table 2-8 also provides information on alternate inband mechanisms to access similar or equivalent information for register reads and writes where applicable. The user should consult the appropriate *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3* or *Intel® Xeon® Processor E5 Product Family Datasheet Volume Two* for exact details on MSR or CSR register content.

Table 2-8. RdPkgConfig() & WrPkgConfig() CPU Thermal and Power Optimization Services Summary (Sheet 1 of 4)

Service	Index Value (decimal)	Parameter Value (word)	RdPkgConfig() Data (dword)	WrPkgConfig() Data (dword)	Description	Alternate Inband MSR or CSR Access
Package Identifier Read	00	0x0000	CPUID Information		Returns processor-specific information including CPU family, model and stepping information.	Execute CPUID instruction to get processor signature
		0x0001	Platform ID		Used to ensure microcode update compatibility with processor.	MSR 17h: IA32_PLATFORM_ID
		0x0002	PCU Device ID		Returns the Device ID information for the processor Power Control Unit.	CSR: DID
		0x0003	Max Thread ID		Returns the maximum 'Thread ID' value supported by the processor.	MSR: RESOLVED_CORES_MASK CSR: RESOLVED_CORES_MASK
		0x0004	CPU Microcode Update Revision		Returns processor microcode and PCU firmware revision information.	MSR 8Bh: IA32_BIOS_SIGN_ID
		0x0005	MCA Error Source Log		Returns the MCA Error Source Log	CSR: MCA_ERR_SRC_LOG
Package Power SKU Unit Read	30	0x0000	Time, Energy and Power Units	N/A	Read units for power, energy and time used in power control registers.	MSR 606h: PACKAGE_POWER_SKU_UNIT CSR: PACKAGE_POWER_SKU_UNIT

Table 2-8. RdPkgConfig() & WrPkgConfig() CPU Thermal and Power Optimization Services Summary (Sheet 2 of 4)

Service	Index Value (decimal)	Parameter Value (word)	RdPkgConfig() Data (dword)	WrPkgConfig() Data (dword)	Description	Alternate Inband MSR or CSR Access
Package Power SKU Read	28	0x0000	Package Power SKU[31:0]	N/A	Returns Thermal Design Power and minimum package power values for the processor SKU.	MSR 614h: PACKAGE_POWER_SKU CSR: PACKAGE_POWER_SKU
Package Power SKU Read	29	0x0000	Package Power SKU[64:32]	N/A	Returns the maximum package power value for the processor SKU and the maximum time interval for which it can be sustained.	MSR 614h: PACKAGE_POWER_SKU CSR: PACKAGE_POWER_SKU
"Wake on PECI" Mode Bit Write / Read	05	0x0001 - Set 0x0000 - Reset	N/A	"Wake on PECI" mode bit	Enables package pop-up to C2 to service PECI PCIConfig() accesses if appropriate.	N/A
"Wake on PECI" Mode Bit Write / Read	05	0x0000	"Wake on PECI" mode bit	N/A	Read status of "Wake on PECI" mode bit	N/A
Accumulated Run Time Read	31	0x0000	Total reference time	N/A	Returns the total run time.	MSR 10h: IA32_TIME_STAMP_COUNTER
Package Temperature Read	02	0x00FF	Processor package Temperature	N/A	Returns the maximum processor die temperature in PECI format.	MSR 1B1h: IA32_PACKAGE_THERM_STATUS
Per Core DTS Temperature Read	09	0x0000-0x0007 (cores 0-7) 0x00FF - System Agent	Per core DTS maximum temperature	N/A	Read the maximum DTS temperature of a particular core or the System Agent within the processor die in relative PECI temperature format	MSR 19Ch: IA32_THERM_STATUS
Temperature Target Read	16	0x0000	Processor T_{jmax} and $T_{CONTROL}$	N/A	Returns the maximum processor junction temperature and processor $T_{CONTROL}$.	MSR 1A2h: TEMPERATURE_TARGET CSR: TEMPERATURE_TARGET
Package Thermal Status Read / Clear	20	0x0000	Thermal Status Register	N/A	Read the thermal status register and optionally clear any log bits. The register includes status and log bits for TCC activation, PROCHOT_N assertion and Critical Temperature.	MSR 1B1h: IA32_PACKAGE_THERM_STATUS
Thermal Averaging Constant Write / Read	21	0x0000	Thermal Averaging Constant	N/A	Reads the Thermal Averaging Constant	N/A
Thermal Averaging Constant Write / Read	21	0x0000	N/A	Thermal Averaging Constant	Writes the Thermal Averaging Constant	N/A

Table 2-8. RdPkgConfig() & WrPkgConfig() CPU Thermal and Power Optimization Services Summary (Sheet 3 of 4)

Service	Index Value (decimal)	Parameter Value (word)	RdPkgConfig() Data (dword)	WrPkgConfig() Data (dword)	Description	Alternate Inband MSR or CSR Access
Thermally Constrained Time Read	32	0x0000	Thermally Constrained Time	N/A	Read the time for which the processor has been operating in a lowered power state due to internal TCC activation.	N/A
Current Limit Read	17	0x0000	Current Limit per power plane	N/A	Reads the current limit on the VCC power plane	CSR: PRIMARY_PLANE_CURRENT_ CONFIG_CONTROL
Accumulated Energy Status Read	03	0x0000 - VCC 0x00FF - CPU package	Accumulated CPU energy	N/A	Returns the value of the energy consumed by just the VCC power plane or entire CPU package.	MSR 639h: PP0_ENERGY_ STATUS CSR: PP0_ENERGY_STATUS MSR 611h: PACKAGE_ENERGY_STATUS CSR: PACKAG_ENERGY_STATUS
Power Limit for the VCC Power Plane Write / Read	25	0x0000	N/A	Power Limit Data	Program power limit for VCC power plane	MSR 638h: PP0_POWER_LIMIT CSR: PP0_POWER_LIMIT
Power Limit for the VCC Power Plane Write / Read	25	0x0000	Power Limit Data	N/A	Read power limit data for VCC power plane	MSR 638h: PP0_POWER_LIMIT CSR: PP0_POWER_LIMIT
Package Power Limits For Multiple Turbo Modes	26	0x0000	N/A	Power Limit 1 Data	Write power limit data 1 in multiple turbo mode.	MSR 610h: PACKAGE_POWER_LIMIT CSR: PACKAGE_POWER_LIMIT
Package Power Limits For Multiple Turbo Modes	27	0x0000	N/A	Power Limit 2 Data	Write power limit data 2 in multiple turbo mode.	MSR 610h: PACKAGE_POWER_LIMIT CSR: PACKAGE_POWER_LIMIT
Package Power Limits For Multiple Turbo Modes	26	0x0000	Power Limit 1 Data	N/A	Read power limit 1 data in multiple turbo mode.	MSR 610h: PACKAGE_POWER_LIMIT CSR: PACKAGE_POWER_LIMIT
Package Power Limits For Multiple Turbo Modes	27	0x0000	Power Limit 2 Data	N/A	Read power limit 2 data in multiple turbo mode.	MSR 610h: PACKAGE_POWER_LIMIT CSR: PACKAGE_POWER_LIMIT
Package Power Limit Performance Status Read	08	0x00FF - CPU package	Accumulated CPU throttle time	N/A	Read the total time for which the processor package was throttled due to power limiting.	CSR: PACKAGE_RAPL_PERF_STATUS
Efficient Performance Indicator Read	06	0x0000	Number of productive processor cycles	N/A	Read number of productive cycles for power budgeting purposes.	N/A
ACPI P-T Notify Write & Read	33	0x0000	N/A	New p-state equivalent of P1 used in conjunction with package power limiting	Notify the processor PCU of the new p-state that is one state below the turbo frequency as specified through the last ACPI Notify	N/A

Table 2-8. RdPkgConfig() & WrPkgConfig() CPU Thermal and Power Optimization Services Summary (Sheet 4 of 4)

Service	Index Value (decimal)	Parameter Value (word)	RdPkgConfig() Data (dword)	WrPkgConfig() Data (dword)	Description	Alternate Inband MSR or CSR Access
ACPI P-T Notify Write & Read	33	0x0000	New p-state equivalent of P1 used in conjunction with package power limiting	N/A	Read the processor PCU to determine the p-state that is one state below the turbo frequency as specified through the last ACPI Notify	N/A
Caching Agent TOR Read	39	Cbo Index, TOR Index, Bank#; Read Mode	Caching Agent (Cbo) Table of Requests (TOR) data; Core ID & associated valid bit	N/A	Read the Cbo TOR data for all enabled cores in the event of a 3-strike timeout. Can alternatively be used to read 'Core ID' data to confirm that IERR was caused by a core timeout	N/A
Thermal Margin Read	10	0x0000	Thermal margin to processor thermal profile or load line	N/A	Read margin to processor thermal load line	N/A

2.5.2.6.12 Package Identifier Read

This feature enables the PECI host to uniquely identify the PECI client processor. The parameter field encodings shown in Table 2-8 allow the PECI host to access the relevant processor information as described below.

- **CPUID data**: This is the equivalent of data that can be accessed through the CPUID instruction execution. It contains processor type, stepping, model and family ID information as shown in Figure 2-21.

Figure 2-21. CPUID Data

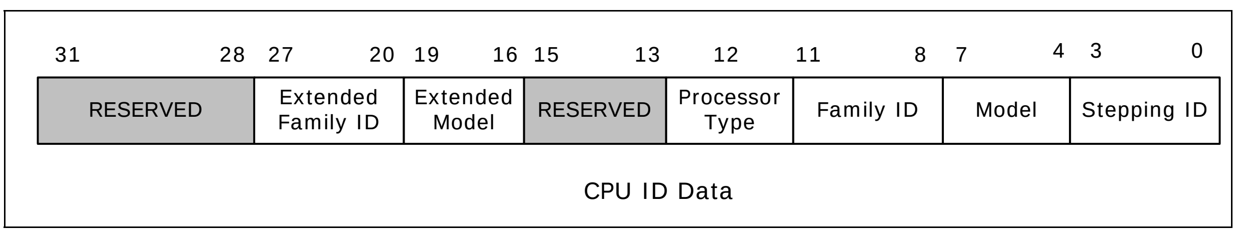

- **Platform ID data**: The Platform ID data can be used to ensure processor microcode updates are compatible with the processor. The value of the Platform ID or Processor Flag[2:0] as shown in Figure 2-22 is typically unique to the platform type and processor stepping. Refer to the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3* for more information.

Figure 2-22. Platform ID Data

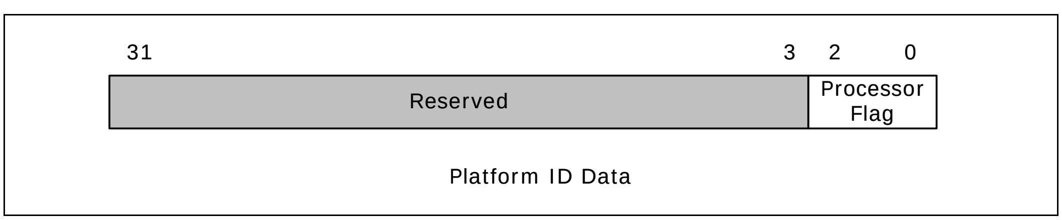

- **PCU Device ID**: This information can be used to uniquely identify the processor power control unit (PCU) device when combined with the Vendor Identification register content and remains constant across all SKUs. Refer to the appropriate register description for the exact processor PCU Device ID value.

Figure 2-23. PCU Device ID

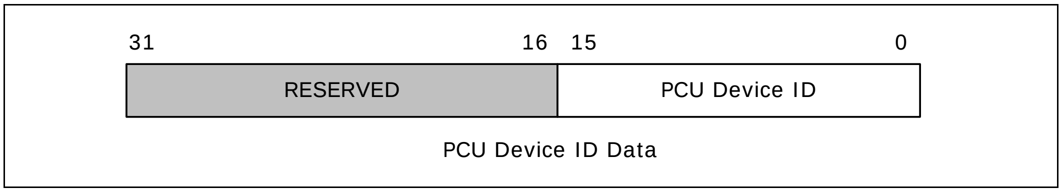

- **Max Thread ID**: The maximum Thread ID data provides the number of supported processor threads. This value is dependent on the number of cores within the processor as determined by the processor SKU and is independent of whether certain cores or corresponding threads are enabled or disabled.

Figure 2-24. Maximum Thread ID

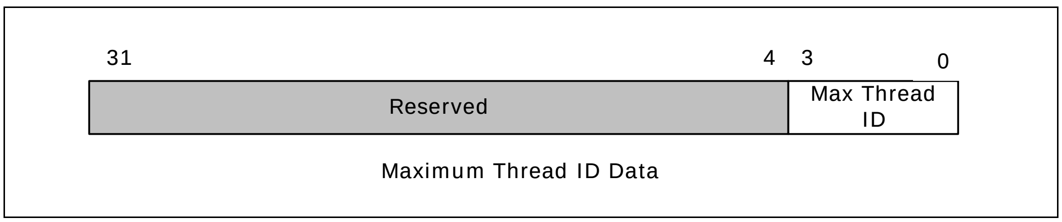

- **CPU Microcode Update Revision**: Reflects the revision number for the microcode update and power control unit firmware updates on the processor sample. The revision data is a unique 32-bit identifier that reflects a combination of specific versions of the processor microcode and PCU control firmware.

Figure 2-25. Processor Microcode Revision

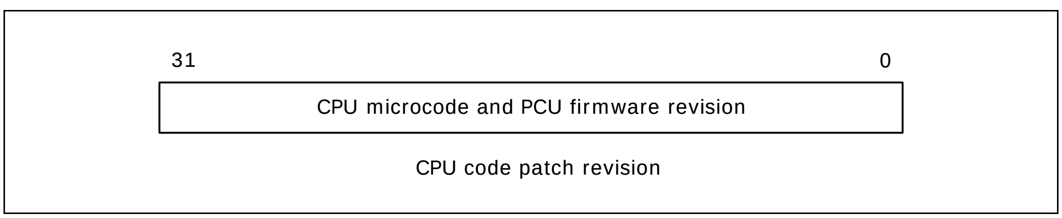

- **Machine Check Status**: Returns error information as logged by the MCA Error Source Log register. See Figure 2-26 for details. The power control unit will assert the relevant bit when the error condition represented by the bit occurs. For example, bit 29 will be set if the package asserted MCERR, bit 30 is set if the package asserted IERR and bit 31 is set if the package asserted CAT_ERR_N. The CAT_ERR_N may be used to signal the occurrence of a MCERR or IERR.

Figure 2-26. Machine Check Status

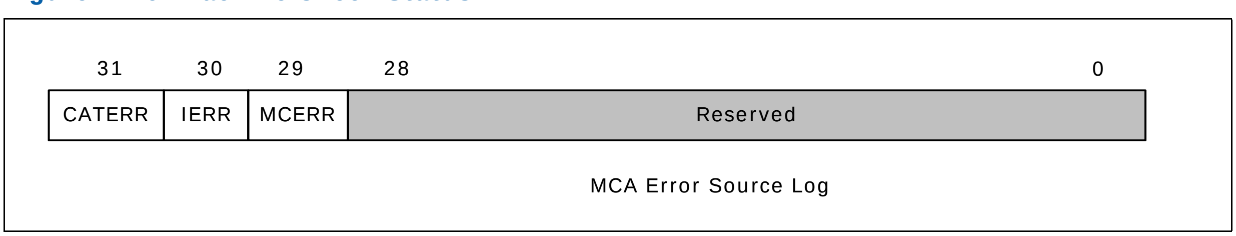

2.5.2.6.13 Package Power SKU Unit Read

This feature enables the PECI host to read the units of time, energy and power used in the processor and DRAM power control registers for calculating power and timing parameters. In Figure 2-27, the default value of the power unit field [3:0] is 0011b, energy unit [12:8] is 10000b and the time unit [19:16] is 1010b. Actual unit values are calculated as shown in Table 2-9.

Figure 2-27. Package Power SKU Unit Data

31 – 20	19 – 16	15 – 13	12 – 8	7 – 4	3 – 0
Reserved	Time Unit	Reserved	Energy Unit	Reserved	Power Unit

Table 2-9. Power Control Register Unit Calculations

Unit Field	Value Calculation	Default Value
Time	$1s / 2^{TIME\ UNIT}$	$1s / 2^{10} = 976\ \mu s$
Energy	$1J / 2^{ENERGY\ UNIT}$	$1J / 2^{16} = 15.3\ \mu J$
Power	$1W / 2^{POWER\ UNIT}$	$1W / 2^{3} = 1/8\ W$

2.5.2.6.14 Package Power SKU Read

This read allows the PECI host to access the minimum, Thermal Design Power and maximum power settings for the processor package SKU. It also returns the maximum time interval or window over which the power can be sustained. If the power limiting entity specifies a power limit value outside of the range specified through these settings, power regulation cannot be guaranteed. Since this data is 64 bits wide, PECI facilitates access to this register by allowing two requests to read the lower 32 bits and upper 32 bits separately as shown in Table 2-8. Power units for this read are determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.13.

'Package Power SKU data' is programmed by the PCU firmware during boot time based on SKU dependent power-on default values set during manufacturing. The TDP package power specified through bits [14:0] in Figure 2-28 is the maximum value of the 'Power Limit1' field in Section 2.5.2.6.26 while the maximum package power in bits [46:32] is the maximum value of the 'Power Limit2' field.

The minimum package power in bits [30:16] is applicable to both the 'Power Limit1' & 'Power Limit2' fields and corresponds to a mode when all the cores are operational and in their lowest frequency mode. Attempts to program the power limit below the minimum power value may not be effective since BIOS/OS, and not the PCU, controls disabling of cores and core activity.

The 'maximum time window' in bits [54:48] is representative of the maximum rate at which the power control unit (PCU) can sample the package energy consumption and reactively take the necessary measures to meet the imposed power limits. Programming too large a time window runs the risk of the PCU not being able to monitor and take timely action on package energy excursions. On the other hand, programming too small a time window may not give the PCU enough time to sample energy information and enforce the limit. The minimum value of the 'time window' can be obtained by reading bits [21:15] of the PWR_LIMIT_MISC_INFO CSR using the PECI RdPCIConfigLocal() command.

Figure 2-28. Package Power SKU Data

Package Power SKU (upper bits)

63–55	54–48	47	46–32
Reserved	Maximum Time Window	Reserved	Maximum Package Power

Package Power SKU (lower bits)

31	30–16	15	14–0
Reserved	Minimum Package Power	Reserved	TDP Package Power

2.5.2.6.15 "Wake on PECI" Mode Bit Write / Read

Setting the "Wake on PECI" mode bit enables successful completion of the WrPCIConfigLocal(), RdPCIConfigLocal(), WrPCIConfig() and RdPCIConfig() PECI commands by forcing a package 'pop-up' to the C2 state to service these commands if the processor is in a low-power state. The exact power impact of such a 'pop-up' is determined by the product SKU, the C-state from which the pop-up is initiated and the negotiated PECI bit rate. A 'reset' or 'clear' of this bit or simply not setting the "Wake on PECI" mode bit could result in a "timeout" response (completion code of 0x82) from the processor indicating that the resources required to service the command are in a low power state.

Alternatively, this mode bit can also be read to determine PECI behavior in package states C3 or deeper.

2.5.2.6.16 Accumulated Run Time Read

This read returns the total time for which the processor has been executing with a resolution of 1 mS per count. This is tracked by a 32-bit counter that rolls over on reaching the maximum value. This counter activates and starts counting for the first time at RESET_N de-assertion.

2.5.2.6.17 Package Temperature Read

This read returns the maximum processor die temperature in 16-bit PECI format. The upper 16 bits of the response data are reserved. The PECI temperature data returned by this read is the 'instantaneous' value and not the 'average' value as returned by the PECI GetTemp() described in Section 2.5.2.3.

Figure 2-29. Package Temperature Read Data

31–16	15	14–6	5–0
RESERVED	Sign Bit	PECI Temperature (Integer Value)	PECI Temperature (Fractional Value)

2.5.2.6.18 Per Core DTS Temperature Read

This feature enables the PECI host to read the maximum value of the DTS temperature for any specific core within the processor. Alternatively, this service can be used to read the System Agent temperature. Temperature is returned in the same format as the Package Temperature Read described in Section 2.5.2.6.17. Data is returned in relative PECI temperature format.

Reads to a parameter value outside the supported range will return an error as indicated by a completion code of 0x90. The supported range of parameter values can vary depending on the number of cores within the processor. The temperature data returned through this feature is the instantaneous value and not an averaged value. It is updated once every 1 mS.

2.5.2.6.19 Temperature Target Read

The Temperature Target Read allows the PECI host to access the maximum processor junction temperature (T_{jmax}) in degrees Celsius. This is also the default temperature value at which the processor thermal control circuit activates. The T_{jmax} value may vary from processor part to part to reflect manufacturing process variations. The Temperature Target read also returns the processor $T_{CONTROL}$ value. $T_{CONTROL}$ is returned in standard PECI temperature format and represents the threshold temperature used by the thermal management system for fan speed control.

Figure 2-30. Temperature Target Read

31 24	23 16	15 8	7 0
RESERVED	Processor Tjmax	$T_{CONTROL}$	RESERVED

2.5.2.6.20 Package Thermal Status Read / Clear

The Thermal Status Read provides information on package level thermal status. Data includes:

- Thermal Control Circuit (TCC) activation
- Bidirectional PROCHOT_N signal assertion
- Critical Temperature

Both status and sticky log bits are managed in this status word. All sticky log bits are set upon a rising edge of the associated status bit and the log bits are cleared only by Thermal Status reads or a processor reset. A read of the Thermal Status word always includes a log bit clear mask that allows the host to clear any or all of the log bits that it is interested in tracking.

A bit set to ‘0’ in the log bit clear mask will result in clearing the associated log bit. If a mask bit is set to ‘0’ and that bit is not a legal mask, a failing completion code will be returned. A bit set to ‘1’ is ignored and results in no change to any sticky log bits. For example, to clear the TCC Activation Log bit and retain all other log bits, the Thermal Status Read should send a mask of 0xFFFFFFFD.

Figure 2-31. Thermal Status Word

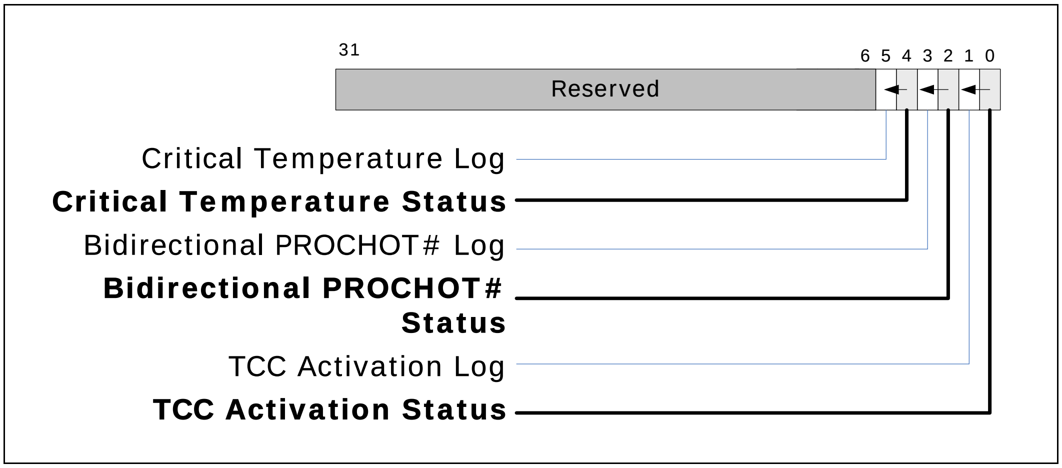

2.5.2.6.21 Thermal Averaging Constant Write / Read

This feature allows the PECI host to control the window over which the estimated processor PECI temperature is filtered. The host may configure this window as a power of two. For example, programming a value of 5 results in a filtering window of 2^5 or 32 samples. The maximum programmable value is 8 or 256 samples. Programming a value of zero would disable the PECI temperature averaging feature. The default value of the thermal averaging constant is 4 which translates to an averaging window size of 2^4 or 16 samples. More details on the PECI temperature filtering function can be found in Section 2.5.7.3.

Figure 2-32. Thermal Averaging Constant Write / Read

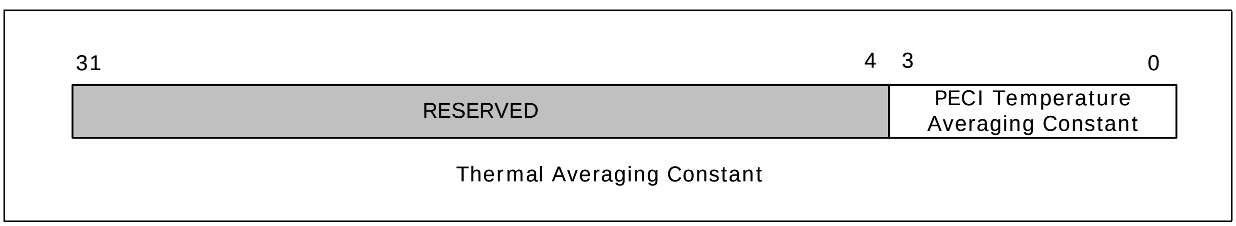

2.5.2.6.22 Thermally Constrained Time Read

This features allows the PECI host to access the total time for which the processor has been operating in a lowered power state due to TCC activation. The returned data includes the time required to ramp back up to the original P-state target after TCC activation expires. This timer does not include TCC activation as a result of an external assertion of PROCHOT_N. This is tracked by a 32-bit counter with a resolution of 1mS per count that rolls over or wraps around. On the processor PECI clients, the only logic that can be thermally constrained is that supplied by VCC.

2.5.2.6.23 Current Limit Read

This read returns the current limit for the processor VCC power plane in 1/8A increments. Actual current limit data is contained only in the lower 13 bits of the response data. The default return value of 0x438 corresponds to a current limit value of 135A.

Figure 2-33. Current Config Limit Read Data

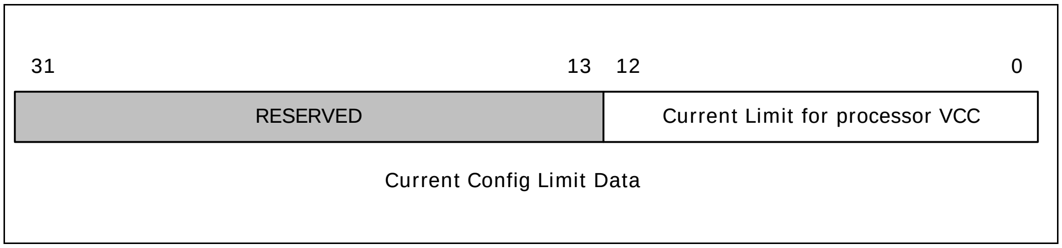

2.5.2.6.24 Accumulated Energy Status Read

This service can return the value of the total energy consumed by the entire processor package or just the logic supplied by the VCC power plane as specified through the parameter field in Table 2-8. This information is tracked by a 32-bit counter that wraps around and continues counting on reaching its limit. Energy units for this read are determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.13.

While Intel requires reading the accumulated energy data at least once every 16 seconds to ensure functional correctness, a more realistic polling rate recommendation is once every 100mS for better accuracy. This feature assumes a 150W processor. In general, as the power capability decreases, so will the minimum polling rate requirement.

When determining energy changes by subtracting energy values between successive reads, Intel advocates using the 2's complement method to account for counter wrap-arounds. Alternatively, adding all 'F's ('0xFFFFFFFF') to a negative result from the subtraction will accomplish the same goal.

Figure 2-34. Accumulated Energy Read Data

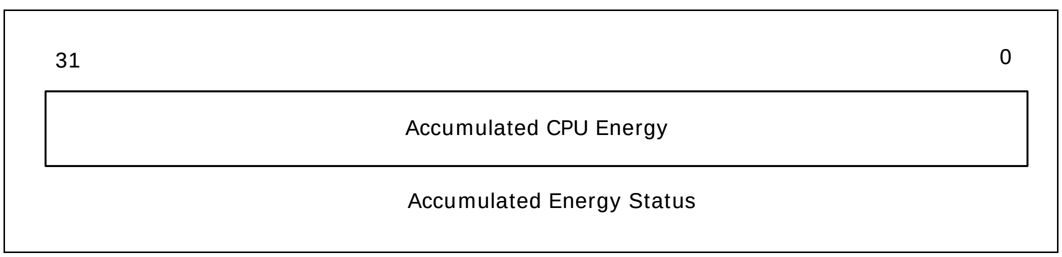

2.5.2.6.25 Power Limit for the VCC Power Plane Write / Read

This feature allows the PECI host to program the power limit over a specified time or control window for the processor logic supplied by the VCC power plane. This typically includes all the cores, home agent and last level cache. The processor does not support power limiting on a per-core basis. Actual power limit values are chosen based on the external VR (voltage regulator) capabilities. The units for the Power Limit and Control Time Window are determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.13.

Since the exact VCC plane power limit value is a function of the platform VR, this feature is not enabled by default and there are no default values associated with the power limit value or the control time window. The Power Limit Enable bit in Figure 2-35 should be set to activate this feature. The Clamp Mode bit is also required to be set to allow the cores to go into power states below what the operating system originally requested. In general, this feature provides an improved mechanism for VR protection compared to the input PROCHOT_N signal assertion method. Both power limit enabling and initialization of power limit values can be done in the same command cycle. Setting a power limit for the VCC plane enables turbo modes for associated logic. External VR protection is guaranteed during boot through operation at safe voltage and frequency. All RAPL parameter values including the power limit value, control time window, clamp mode and enable bit will have to be specified correctly even if the intent is to change just one parameter value when programming over PECI.

The usefulness of the VCC power plane RAPL may be somewhat limited if the platform has a fully compliant external voltage regulator. However, platforms using lower cost voltage regulators may find this feature useful. The VCC RAPL value is generally expected to be a static value after initialization and there may not be any use cases for dynamic control of VCC plane power limit values during run time. BIOS may be ideally used to read the VR (and associated heat sink) capabilities and program the PCU with the power limit information during boot. No matter what the method is, Intel recommends exclusive use of just one entity or interface, PECI for instance, to manage VCC plane power limiting needs. If PECI is being used to manage VCC plane power limiting activities, BIOS should lock out all subsequent inband VCC plane power limiting accesses by setting bit 31 of the PP0_POWER_LIMIT MSR and CSR to ‘1’.

The same conversion formula used for DRAM Power Limiting (see Section 2.5.2.6.9) should be applied for encoding or programming the ‘Control Time Window’ in bits [23:17].

Figure 2-35. Power Limit Data for VCC Power Plane

31 – 24	23 – 17	16	15	14 – 0
RESERVED	Control Time Window	Clamp Mode	Power Limit Enable	VCC Plane Power Limit

VCC Power Plane Power Limit Data

2.5.2.6.26 Package Power Limits For Multiple Turbo Modes

This feature allows the PECI host to program two power limit values to support multiple turbo modes. The operating systems and drivers can balance the power budget using these two limits. Two separate PECI requests are available to program the lower and upper 32 bits of the power limit data shown in Figure 2-36. The units for the Power Limit and Control Time Window are determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.13 while the valid range for power limit values are determined by the Package Power SKU settings described in Section 2.5.2.6.14. Setting the Clamp Mode bits is required to allow the cores to go into power states below what the operating system originally requested. The Power Limit Enable bits should be set to enable the power limiting function. Power limit values, enable and clamp mode bits can all be set in the same command cycle. All RAPL parameter values including the power limit value, control time window, clamp mode and enable bit will have to be specified correctly even if the intent is to change just one parameter value when programming over PECI.

Intel recommends exclusive use of just one entity or interface, PECI for instance, to manage all processor package power limiting and budgeting needs. If PECI is being used to manage package power limiting activities, BIOS should lock out all subsequent inband package power limiting accesses by setting bit 31 of the PACKAGE_POWER_LIMIT MSR and CSR to ‘1’. The ‘power limit 1’ is intended to limit processor power consumption to any reasonable value below TDP and defaults to TDP.

‘Power Limit 1’ values may be impacted by the processor heat sinks and system air flow. Processor ‘power limit 2’ can be used as appropriate to limit the current drawn by the processor to prevent any external power supply unit issues. The ‘Power Limit 2’ should always be programmed to a value (typically 20%) higher than ‘Power Limit 1’ and has no default value associated with it.

Though this feature is disabled by default and external programming is required to enable, initialize and control package power limit values and time windows, the processor package will still turbo to TDP if ‘Power Limit 1’ is not enabled or initialized. ‘Control Time Window#1’ (Power_Limit_1_Time also known as Tau) values may be programmed to be within a range of 250 mS-40 seconds. ‘Control Time Window#2’ (Power_Limit_2_Time) values should be in the range 3 mS-10 mS.

The same conversion formula used for the DRAM Power Limiting feature (see Section 2.5.2.6.9) should be applied when programming the ‘Control Time Window’ bits [23:17] for ‘power limit 1’ in Figure 2-36. The ‘Control Time Window’ for ‘power limit 2’ can be directly programmed into bits [55:49] in units of mS without the aid of any conversion formulas.

Figure 2-36. Package Turbo Power Limit Data

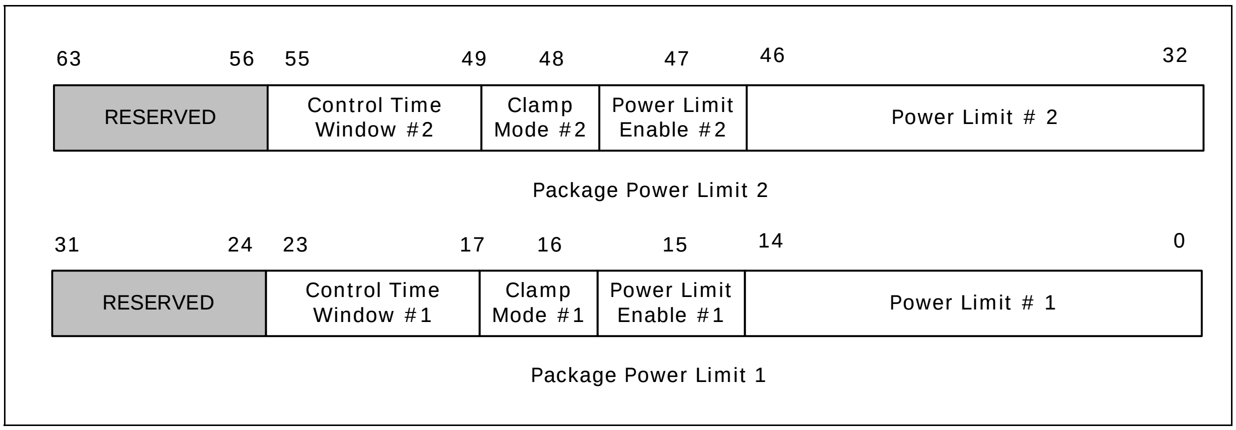

2.5.2.6.27 Package Power Limit Performance Status Read

This service allows the PECI host to assess the performance impact of the currently active power limiting modes. The read return data contains the total amount of time for which the entire processor package has been operating in a power state that is lower than what the operating system originally requested. This information is tracked by a 32-bit counter that wraps around. The unit for time is determined as per the Package Power SKU Unit settings described in Section 2.5.2.6.13.

Figure 2-37. Package Power Limit Performance Data

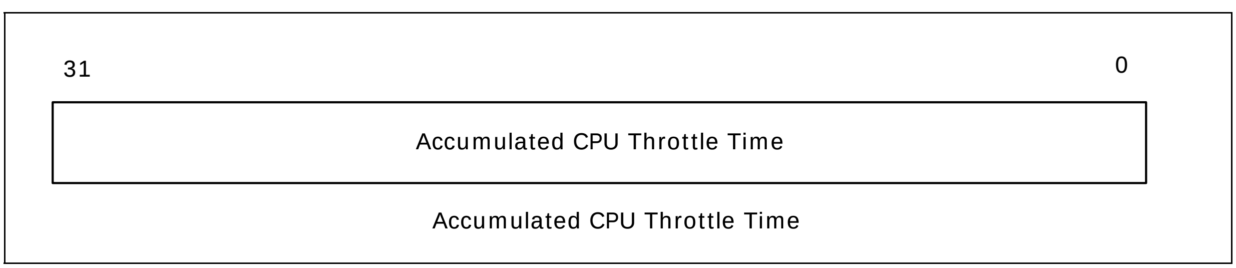

2.5.2.6.28 Efficient Performance Indicator Read

The Efficient Performance Indicator (EPI) Read provides an indication of the total number of productive cycles. Specifically, these are the cycles when the processor is engaged in any activity to retire instructions and as a result, consuming energy. Any power management entity monitoring this indicator should sample it at least once every 4 seconds to enable detection of wraparounds. Refer to the processor *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3*, for details on programming the Energy/Performance Bias (MSR_MISC_PWR_MGMT) register to set the 'Energy Efficiency' policy of the processor.

Figure 2-38. Efficient Performance Indicator Read

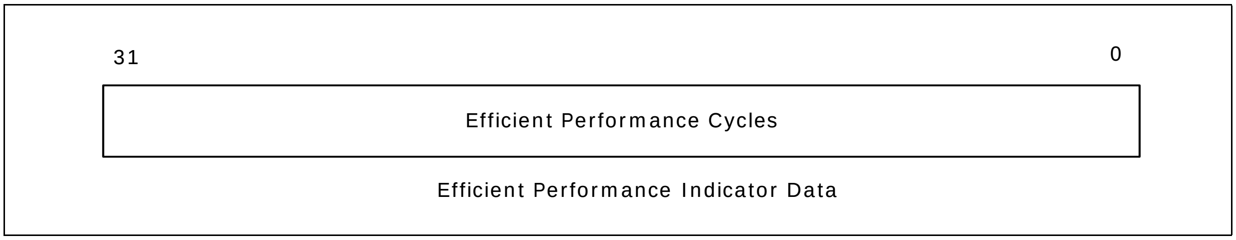

2.5.2.6.29 ACPI P-T Notify Write & Read

This feature enables the processor turbo capability when used in conjunction with the PECI package RAPL or power limit. When the BMC sets the package power limit to a value below TDP, it also determines a new corresponding turbo frequency and notifies the OS using the 'ACPI Notify' mechanism as supported by the _PPC or performance present capabilities object. The BMC then notifies the processor PCU using the PECI 'ACPI P-T Notify' service by programming a new state that is one p-state below the turbo frequency sent to the OS via the _PPC method.

When the OS requests a p-state higher than what is specified in bits [7:0] of the PECI ACPI P-T Notify data field, the CPU will treat it as request for P0 or turbo. The PCU will use the IA32_ENERGY_PERFORMANCE_BIAS register settings to determine the exact extent of turbo. Any OS p-state request that is equal to or below what is specified in the PECI ACPI P-T Notify will be granted as long as the RAPL power limit does not impose a lower p-state. However, turbo will not be enabled in this instance even if there is headroom between the processor energy consumption and the RAPL power limit.

This feature does not affect the Thermal Monitor behavior of the processor nor is it impacted by the setting of the power limit clamp mode bit.

Figure 2-39. ACPI P-T Notify Data

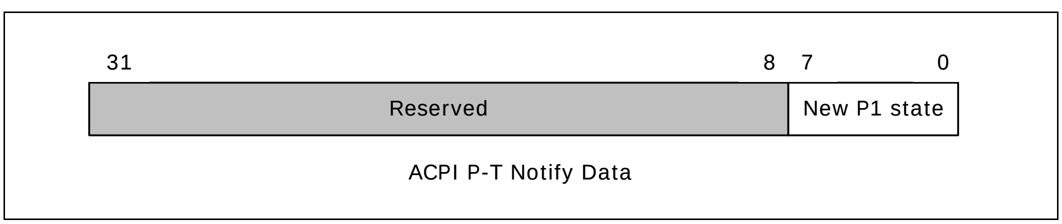

2.5.2.6.30 Caching Agent TOR Read

This feature allows the PECI host to read the Caching Agent (Cbo) Table of Requests (TOR). This information is useful for debug in the event of a 3-strike timeout that results in a processor IERR assertion. The 16-bit parameter field is used to specify the Cbo index, TOR array index and bank number according to the following bit assignments.

- Bits [1:0] - Bank Number - legal values from 0 to 2
- Bits [6:2] - TOR Array Index - legal values from 0 to 19
- Bits [10:7] - Cbo Index - legal values from 0 to 7
- Bit [11] - Read Mode - should be set to ‘0’ for TOR reads
- Bits [15:12] - Reserved

Bit[11] is the Read Mode bit and should be set to ‘0’ for TOR reads. The Read Mode bit can alternatively be set to ‘1’ to read the ‘Core ID’ (with associated valid bit as shown in Figure 2-40) that points to the first core that asserted the IERR. In this case bits [10:0] of the parameter field are ignored. The ‘Core ID’ read may not return valid data until at least 1 mS after the IERR assertion.

Figure 2-40. Caching Agent TOR Read Data

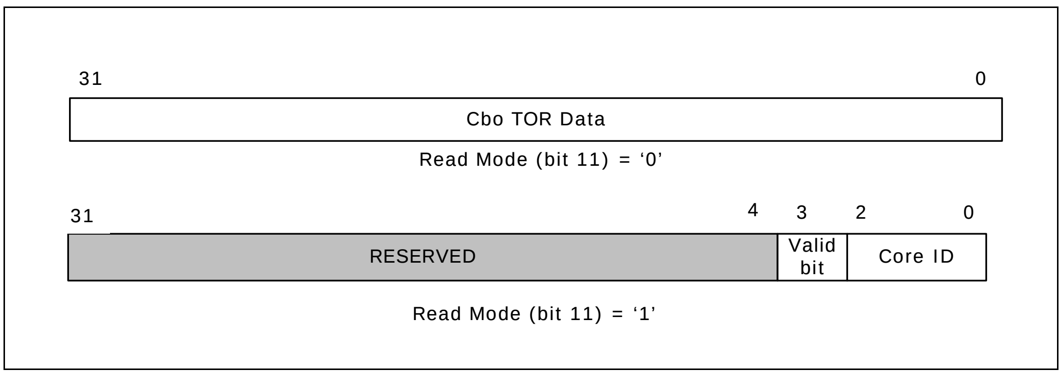

Note: Reads to caching agents that are not enabled will return all zeroes. Refer to the debug handbook for details on methods to interpret the crash dump results using the Cbo TOR data shown in Figure 2-40.

2.5.2.6.31 Thermal Margin Read

This service allows the PECI host to read the margin to the processor thermal profile or load line. Thermal margin data is returned in the format shown in Figure 2-41 with a sign bit, an integer part and a fractional part. A negative thermal margin value implies that the processor is operating in violation of its thermal load line and may be indicative of a need for more aggressive cooling mechanisms through a fan speed increase or other means. This PECI service will continue to return valid margin values even when the processor die temperature exceeds T_{jmax}.

Figure 2-41. DTS Thermal Margin Read

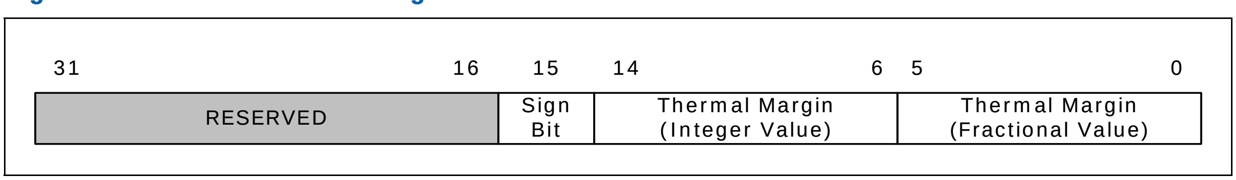

2.5.2.7 RdIAMSR()

The RdIAMSR() PECI command provides read access to Model Specific Registers (MSRs) defined in the processor's Intel® Architecture (IA). MSR definitions may be found in the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3*. Refer to Table 2-11 for the exact listing of processor registers accessible through this command.

2.5.2.7.1 Command Format

The RdIAMSR() format is as follows:

Write Length: 0x05

Read Length: 0x09 (qword)

Command: 0xb1

Description: Returns the data maintained in the processor IA MSR space as specified by the 'Processor ID' and 'MSR Address' fields. The Read Length dictates the desired data return size. This command supports only qword responses. All command responses are prepended with a completion code that contains additional pass/fail status information. Refer to Section 2.5.5.2 for details regarding completion codes.

2.5.2.7.2 Processor ID Enumeration

The 'Processor ID' field that is used to address the IA MSR space refers to a specific logical processor within the CPU. The 'Processor ID' always refers to the same physical location in the processor silicon regardless of configuration as shown in the example in Figure 2-42. For example, if certain logical processors are disabled by BIOS, the Processor ID mapping will not change. The total number of Processor IDs on a CPU is product-specific.

'Processor ID' enumeration involves discovering the logical processors enabled within the CPU package. This can be accomplished by reading the 'Max Thread ID' value through the RdPkgConfig() command (Index 0, Parameter 3) described in Section 2.5.2.6.12 and subsequently querying each of the supported processor threads. Unavailable processor threads will return a completion code of 0x90.

Alternatively, this information may be obtained from the RESOLVED_CORES_MASK register readable through the RdPCIConfigLocal() PECI command described in Section 2.5.2.9 or other means. Bits [7:0] and [9:8] of this register contain the 'Core Mask' and 'Thread Mask' information respectively. The 'Thread Mask' applies to all the enabled cores within the processor package as indicated by the 'Core Mask'. For the processor PECI clients, the 'Processor ID' may take on values in the range 0 through 15.

Figure 2-42. Processor ID Construction Example

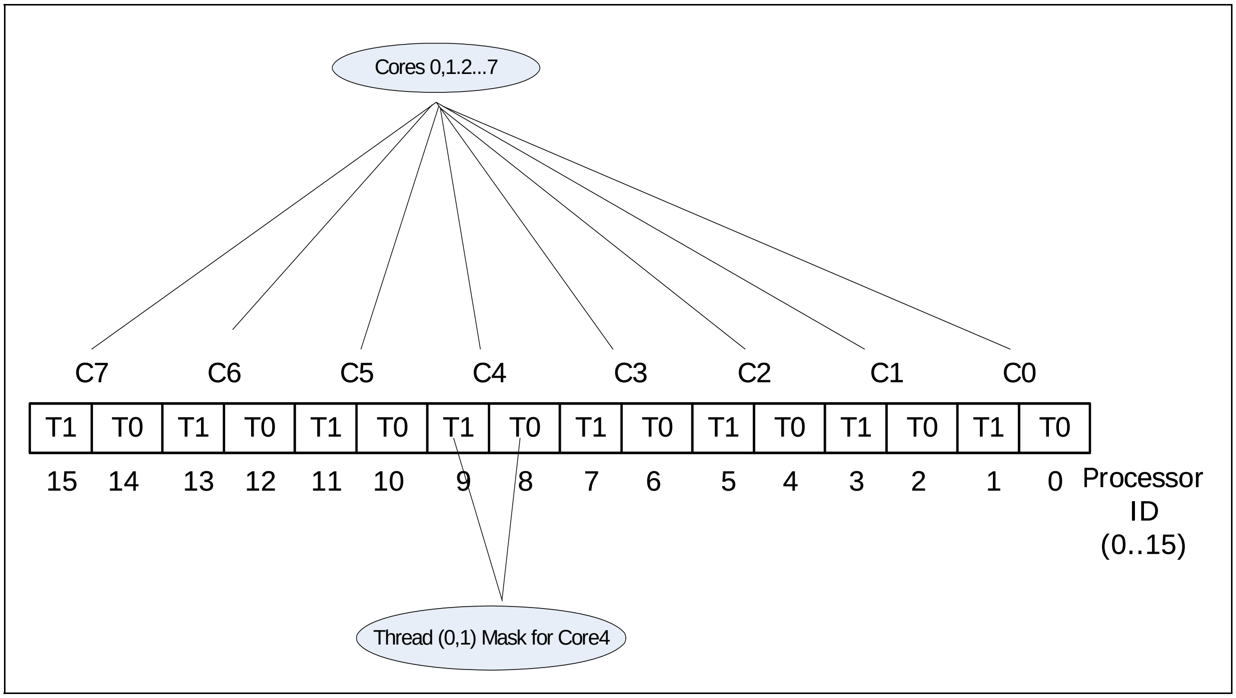

Figure 2-43. RdIAMSR()

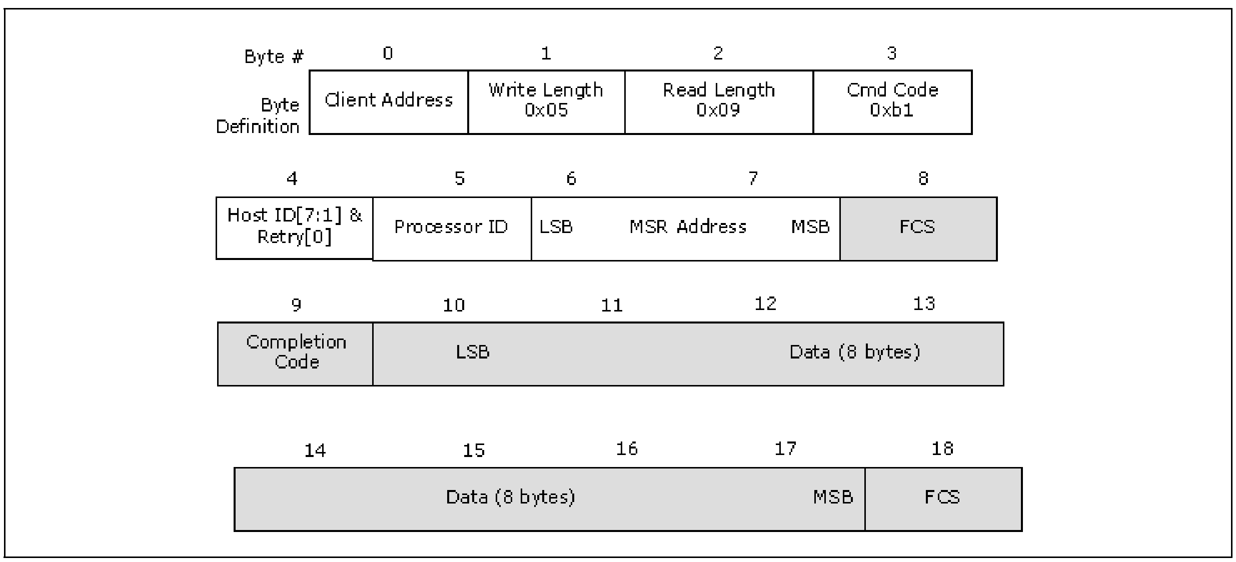

Note: The 2-byte MSR Address field and read data field defined in Figure 2-43 are sent in standard PECI ordering with LSB first and MSB last.

2.5.2.7.3 Supported Responses

The typical client response is a passing FCS, a passing Completion Code and valid data. Under some conditions, the client's response will indicate a failure.

Table 2-10. RdIAMSR() Response Definition

Response	Meaning
Bad FCS	Electrical error
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
CC: 0x40	Command passed, data is valid.
CC: 0x80	Response timeout. The processor was not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor is not able to allocate resources for servicing this command at this time. Retry is appropriate.
CC: 0x82	The processor hardware resources required to service this command are in a low power state. Retry may be appropriate after modification of PECI wake mode behavior if appropriate.
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.

2.5.2.7.4 RdIAMSR() Capabilities

The processor PECI client allows PECI RdIAMSR() access to the registers listed in Table 2-11. These registers pertain to the processor core and uncore error banks (machine check banks 0 through 19). Information on the exact number of accessible banks for the processor device may be obtained by reading the IA32_MCG_CAP[7:0] MSR (0x0179). This register may be alternatively read using a RDMSR BIOS instruction. Please consult the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3* for more information on the exact number of cores supported by a particular processor SKU. Any attempt to read processor MSRs that are not accessible over PECI or simply not implemented will result in a completion code of 0x90.

PECI access to these registers is expected only when in-band access mechanisms are not available.

Table 2-11. RdIAMSR() Services Summary (Sheet 1 of 2)

Process or ID (byte)	MSR Address (dword)	Meaning	Process or ID (byte)	MSR Address (dword)	Meaning	Processor ID (byte)	MSR Address (dword)	Meaning
0x0-0xF	0x0400	IA32_MC0_CTL	0x0-0xF	0x041B	IA32_MC6_MISC	0x0-0xF	0x0436	IA32_MC13_ADDR
0x0-0xF	0x0280	IA32_MC0_CTL2	0x0-0xF	0x041C	IA32_MC7_CTL	0x0-0xF	0x0437	IA32_MC13_MISC
0x0-0xF	0x0401	IA32_MC0_STATUS	0x0-0xF	0x0287	IA32_MC7_CTL2	0x0-0xF	0x0438	IA32_MC14_CTL
0x0-0xF	0x0402	IA32_MC0_ADDR	0x0-0xF	0x041D	IA32_MC7_STATUS	0x0-0xF	0x028E	IA32_MC14_CTL2
0x0-0xF	0x0403	IA32_MC0_MISC[1]	0x0-0xF	0x041E	IA32_MC7_ADDR	0x0-0xF	0x0439	IA32_MC14_STATUS
0x0-0xF	0x0404	IA32_MC1_CTL	0x0-0xF	0x041F	IA32_MC7_MISC	0x0-0xF	0x043A	IA32_MC14_ADDR
0x0-0xF	0x0281	IA32_MC1_CTL2	0x0-0xF	0x0420	IA32_MC8_CTL	0x0-0xF	0x043B	IA32_MC14_MISC
0x0-0xF	0x0405	IA32_MC1_STATUS	0x0-0xF	0x0288	IA32_MC8_CTL2	0x0-0xF	0x043C	IA32_MC15_CTL
0x0-0xF	0x0406	IA32_MC1_ADDR	0x0-0xF	0x0421	IA32_MC8_STATUS	0x0-0xF	0x028F	IA32_MC15_CTL2
0x0-0xF	0x0407	IA32_MC1_MISC	0x0-0xF	0x0422	IA32_MC8_ADDR	0x0-0xF	0x043D	IA32_MC15_STATUS
0x0-0xF	0x0408	IA32_MC2_CTL[2]	0x0-0xF	0x0423	IA32_MC8_MISC	0x0-0xF	0x043E	IA32_MC15_ADDR

Table 2-11. RdIAMSR() Services Summary (Sheet 2 of 2)

Processor ID (byte)	MSR Address (dword)	Meaning	Processor ID (byte)	MSR Address (dword)	Meaning	Processor ID (byte)	MSR Address (dword)	Meaning
0x0-0xF	0x0282	IA32_MC2_CTL2	0x0-0xF	0x0424	IA32_MC9_CTL	0x0-0xF	0x043F	IA32_MC15_MISC
0x0-0xF	0x0409	IA32_MC2_STATUS	0x0-0xF	0x0289	IA32_MC9_CTL2	0x0-0xF	0x0440	IA32_MC16_CTL
0x0-0xF	0x040A	IA32_MC2_ADDR[2]	0x0-0xF	0x0425	IA32_MC9_STATUS	0x0-0xF	0x0290	IA32_MC16_CTL2
0x0-0xF	0x040B	IA32_MC2_MISC[2]	0x0-0xF	0x0426	IA32_MC9_ADDR	0x0-0xF	0x0441	IA32_MC16_STATUS
0x0-0xF	0x040C	IA32_MC3_CTL	0x0-0xF	0x0427	IA32_MC9_MISC	0x0-0xF	0x0442	IA32_MC16_ADDR
0x0-0xF	0x0283	IA32_MC3_CTL2	0x0-0xF	0x0428	IA32_MC10_CTL	0x0-0xF	0x0443	IA32_MC16_MISC
0x0-0xF	0x040D	IA32_MC3_STATUS	0x0-0xF	0x028A	IA32_MC10_CTL2	0x0-0xF	0x0444	IA32_MC17_CTL
0x0-0xF	0x040E	IA32_MC3_ADDR	0x0-0xF	0x0429	IA32_MC10_STATUS	0x0-0xF	0x0291	IA32_MC17_CTL2
0x0-0xF	0x040F	IA32_MC3_MISC	0x0-0xF	0x042A	IA32_MC10_ADDR	0x0-0xF	0x0445	IA32_MC17_STATUS
0x0-0xF	0x0410	IA32_MC4_CTL	0x0-0xF	0x042B	IA32_MC10_MISC	0x0-0xF	0x0446	IA32_MC17_ADDR
0x0-0xF	0x0284	IA32_MC4_CTL2	0x0-0xF	0x042C	IA32_MC11_CTL	0x0-0xF	0x0447	IA32_MC17_MISC
0x0-0xF	0x0411	IA32_MC4_STATUS	0x0-0xF	0x028B	IA32_MC11_CTL2	0x0-0xF	0x0448	IA32_MC18_CTL
0x0-0xF	0x0412	IA32_MC4_ADDR[2]	0x0-0xF	0x042D	IA32_MC11_STATUS	0x0-0xF	0x0292	IA32_MC18_CTL2
0x0-0xF	0x0413	IA32_MC4_MISC[2]	0x0-0xF	0x042E	IA32_MC11_ADDR	0x0-0xF	0x0449	IA32_MC18_STATUS
0x0-0xF	0x0414	IA32_MC5_CTL	0x0-0xF	0x042F	IA32_MC11_MISC	0x0-0xF	0x044A	IA32_MC18_ADDR
0x0-0xF	0x0285	IA32_MC5_CTL2	0x0-0xF	0x0430	IA32_MC12_CTL	0x0-0xF	0x044B	IA32_MC18_MISC
0x0-0xF	0x0415	IA32_MC5_STATUS	0x0-0xF	0x028C	IA32_MC12_CTL2	0x0-0xF	0x044C	IA32_MC19_CTL
0x0-0xF	0x0416	IA32_MC5_ADDR	0x0-0xF	0x0431	IA32_MC12_STATUS	0x0-0xF	0x0293	IA32_MC19_CTL2
0x0-0xF	0x0417	IA32_MC5_MISC	0x0-0xF	0x0432	IA32_MC12_ADDR	0x0-0xF	0x044D	IA32_MC19_STATUS
0x0-0xF	0x0418	IA32_MC6_CTL	0x0-0xF	0x0433	IA32_MC12_MISC	0x0-0xF	0x044E	IA32_MC19_ADDR
0x0-0xF	0x0286	IA32_MC6_CTL2	0x0-0xF	0x0434	IA32_MC13_CTL	0x0-0xF	0x0179	IA32_MCG_CAP
0x0-0xF	0x0419	IA32_MC6_STATUS	0x0-0xF	0x028D	IA32_MC13_CTL2	0x0-0xF	0x017A	IA32_MCG_STATUS
0x0-0xF	0x041A	IA32_MC6_ADDR	0x0-0xF	0x0435	IA32_MC13_STATUS	0x0-0xF	0x0178	IA32_MCG_CONTAIN

Notes:

1. The IA32_MC0_MISC register details will be available upon implementation in a future processor stepping.
2. The MCi_ADDR and MCi_MISC registers for machine check banks 2 & 4 are not implemented on the processors. The MCi_CTL register for machine check bank 2 is also not implemented.
3. The PECI host must determine the total number of machine check banks and the validity of the MCi_ADDR and MCi_MISC register contents prior to issuing a read to the machine check bank similar to standard machine check architecture enumeration and accesses.
4. The information presented in Table 2-11 is applicable to the processor only. No association between bank numbers and logical functions should be assumed for any other processor devices (past, present or future) based on the information presented in Table 2-11.
5. Intel® Xeon® processor E5-2400 product family banks 7 & 8 corresponding to QPI[1] and iMC[0] are not available on this processor. Reading any registers within these banks will return all '0's.
6. The processor machine check banks 4 through 19 reside in the processor uncore and hence will return the same value independent of the processor ID used to access these banks.
7. The IA32_MCG_STATUS, IA32_MCG_CONTAIN and IA32_MCG_CAP are located in the uncore and will return the same value independent of the processor ID used to access them.
8. The processor machine check banks 0 through 3 are core-specific. Since the processor ID is thread-specific and not core-specific, machine check banks 0 through 3 will return the same value for a particular core independent of the thread referenced by the processor ID.
9. PECI accesses to the machine check banks may not be possible in the event of a core hang. A warm reset of the processor may be required to read any sticky machine check banks.
10. Valid processor ID values may be obtained by using the enumeration methods described in Section 2.5.2.7.2.
11. Reads to a machine check bank within a core or thread that is disabled will return all zeroes with a completion code of 0x90.
12. For SKUs where Intel QPI is disabled or absent, reads to the corresponding machine check banks will return all zeros with a completion code of 0x40.

2.5.2.8 RdPCIConfig()

The RdPCIConfig() command provides sideband read access to the PCI configuration space maintained in downstream devices external to the processor. PECI originators may conduct a device/function/register enumeration sweep of this space by issuing reads in the same manner that the BIOS would. A response of all 1's may indicate that the device/function/register is unimplemented even with a 'passing' completion code. Alternatively, reads to unimplemented registers may return a completion code of 0x90 indicating an invalid request. Responses will follow normal PCI protocol.

PCI configuration addresses are constructed as shown in Figure 2-44. Under normal in-band procedures, the Bus number would be used to direct a read or write to the proper device. Actual PCI bus numbers for all PCI devices including the PCH are programmable by BIOS. The bus number for PCH devices may be obtained by reading the CPUBUSNO CSR. Refer to the *Intel® Xeon® Processor E5 Product Family Datasheet Volume Two* document for details on this register.

Figure 2-44. PCI Configuration Address

31 28	27 20	19 15	14 12	11 0
Reserved	Bus	Device	Function	Register

PCI configuration reads may be issued in byte, word or dword granularities.

2.5.2.8.1 Command Format

The RdPCIConfig() format is as follows:

Write Length: 0x06

Read Length: 0x05 (dword)

Command: 0x61

Description: Returns the data maintained in the PCI configuration space at the requested PCI configuration address. The Read Length dictates the desired data return size. This command supports only dword responses with a completion code on the processor PECI clients. All command responses are prepended with a completion code that includes additional pass/fail status information. Refer to Section 2.5.5.2 for details regarding completion codes.

Figure 2-45. RdPCIConfig()

Byte #	0	1	2	3
Byte Definition	Client Address	Write Length 0x06	Read Length 0x05	Cmd Code 0x61

4	5	6	7	8	9
Host ID[7:1] & Retry[0]	LSB	PCI Configuration Address		MSB	FCS

10	11	12	13	14	15
Completion Code	LSB	Data (4 bytes)		MSB	FCS

Note: The 4-byte PCI configuration address and read data field defined in Figure 2-45 are sent in standard PECI ordering with LSB first and MSB last.

2.5.2.8.2 Supported Responses

The typical client response is a passing FCS, a passing Completion Code and valid data. Under some conditions, the client's response will indicate a failure.

The PECI client response can also vary depending on the address and data. It will respond with a passing completion code if it successfully submits the request to the appropriate location and gets a response.

Table 2-12. RdPCIConfig() Response Definition

Response	Meaning
Bad FCS	Electrical error
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
CC: 0x40	Command passed, data is valid.
CC: 0x80	Response timeout. The processor was not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor is not able to allocate resources for servicing this command at this time. Retry is appropriate.
CC: 0x82	The processor hardware resources required to service this command are in a low power state. Retry may be appropriate after modification of PECI wake mode behavior if appropriate.
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.

2.5.2.9 RdPCIConfigLocal()

The RdPCIConfigLocal() command provides sideband read access to the PCI configuration space that resides within the processor. This includes all processor IIO and uncore registers within the PCI configuration space as described in the *Intel® Xeon® Processor E5 Product Family Datasheet Volume Two* document.

PECI originators may conduct a device/function enumeration sweep of this space by issuing reads in the same manner that the BIOS would. A response of all 1's may indicate that the device/function/register is unimplemented even with a 'passing' completion code. Alternatively, reads to unimplemented or hidden registers may return a completion code of 0x90 indicating an invalid request. It is also possible that reads to function 0 of non-existent IIO devices issued prior to BIOS POST may return all '0's with a passing completion code. PECI originators can access this space even prior to BIOS enumeration of the system buses. There is no read restriction on accesses to locked registers.

PCI configuration addresses are constructed as shown in Figure 2-46. Under normal in-band procedures, the Bus number would be used to direct a read or write to the proper device. PECI reads to the processor IIO devices should specify a bus number of '0000' and reads to the rest of the processor uncore should specify a bus number of '0001' for bits [23:20] in Figure 2-46. Any request made with a bad Bus number is ignored and the client will respond with all '0's and a 'passing' completion code.

Figure 2-46. PCI Configuration Address for local accesses

23 20	19 15	14 12	11 0
Bus	Device	Function	Register

2.5.2.9.1 Command Format

The RdPCIConfigLocal() format is as follows:

Write Length: 0x05

Read Length: 0x02 (byte), 0x03 (word), 0x05 (dword)

Command: 0xe1

Description: Returns the data maintained in the PCI configuration space within the processor at the requested PCI configuration address. The Read Length dictates the desired data return size. This command supports byte, word and dword responses as well as a completion code. All command responses are prepended with a completion code that includes additional pass/fail status information. Refer to Section 2.5.5.2 for details regarding completion codes.

Figure 2-47. RdPCIConfigLocal()

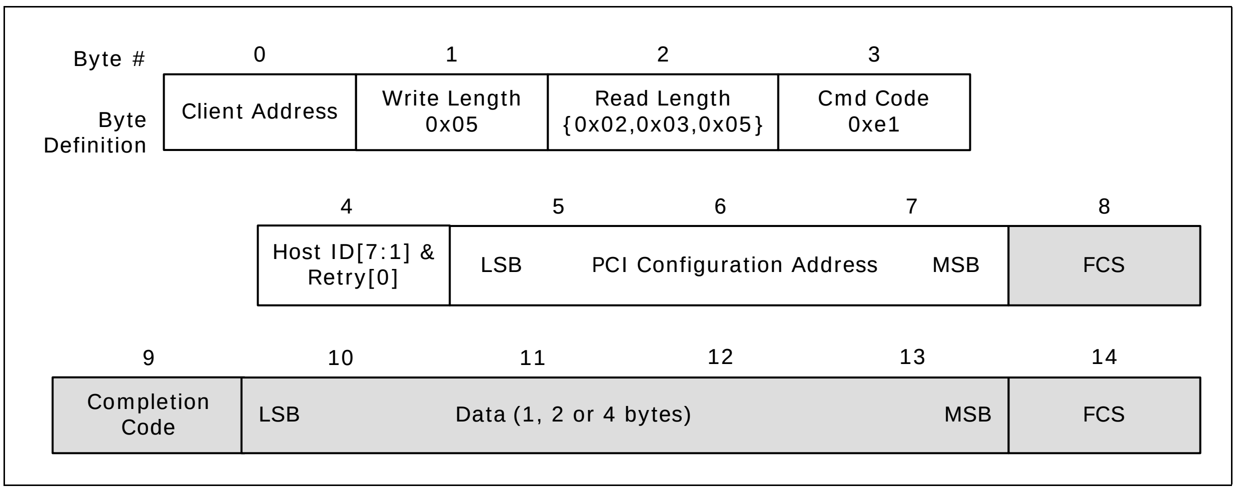

Note: The 3-byte PCI configuration address and read data field defined in Figure 2-47 are sent in standard PECI ordering with LSB first and MSB last.

2.5.2.9.2 Supported Responses

The typical client response is a passing FCS, a passing Completion Code and valid data. Under some conditions, the client’s response will indicate a failure.

The PECI client response can also vary depending on the address and data. It will respond with a passing completion code if it successfully submits the request to the appropriate location and gets a response.

Table 2-13. RdPCIConfigLocal() Response Definition (Sheet 1 of 2)

Response	Meaning
Bad FCS	Electrical error
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
CC: 0x40	Command passed, data is valid.
CC: 0x80	Response timeout. The processor was not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor is not able to allocate resources for servicing this command at this time. Retry is appropriate.

Table 2-13. RdPCIConfigLocal() Response Definition (Sheet 2 of 2)

Response	Meaning
CC: 0x82	The processor hardware resources required to service this command are in a low power state. Retry may be appropriate after modification of PECI wake mode behavior if appropriate.
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.

2.5.2.10 WrPCIConfigLocal()

The WrPCIConfigLocal() command provides sideband write access to the PCI configuration space that resides within the processor. PECI originators can access this space even before BIOS enumeration of the system buses. The exact listing of supported devices and functions for writes using this command on the processor is defined in Table 2-19. The write accesses to registers that are locked will not take effect but will still return a completion code of 0x40. However, write accesses to registers that are hidden will return a completion code of 0x90.

Because a WrPCIConfigLocal() command results in an update to potentially critical registers inside the processor, it includes an Assured Write FCS (AW FCS) byte as part of the write data payload. In the event that the AW FCS mismatches with the client-calculated FCS, the client will abort the write and will always respond with a bad write FCS.

PCI Configuration addresses are constructed as shown in Figure 2-46. The write command is subject to the same address configuration rules as defined in Section 2.5.2.9. PCI configuration writes may be issued in byte, word or dword granularity.

2.5.2.10.1 Command Format

The WrPCIConfigLocal() format is as follows:

Write Length: 0x07 (byte), 0x08 (word), 0x0a (dword)

Read Length: 0x01

Command: 0xe5

AW FCS Support: Yes

Description: Writes the data sent to the requested register address. Write Length dictates the desired write granularity. The command always returns a completion code indicating pass/fail status. Refer to Section 2.5.5.2 for details on completion codes.

Figure 2-48. WrPCIConfigLocal()

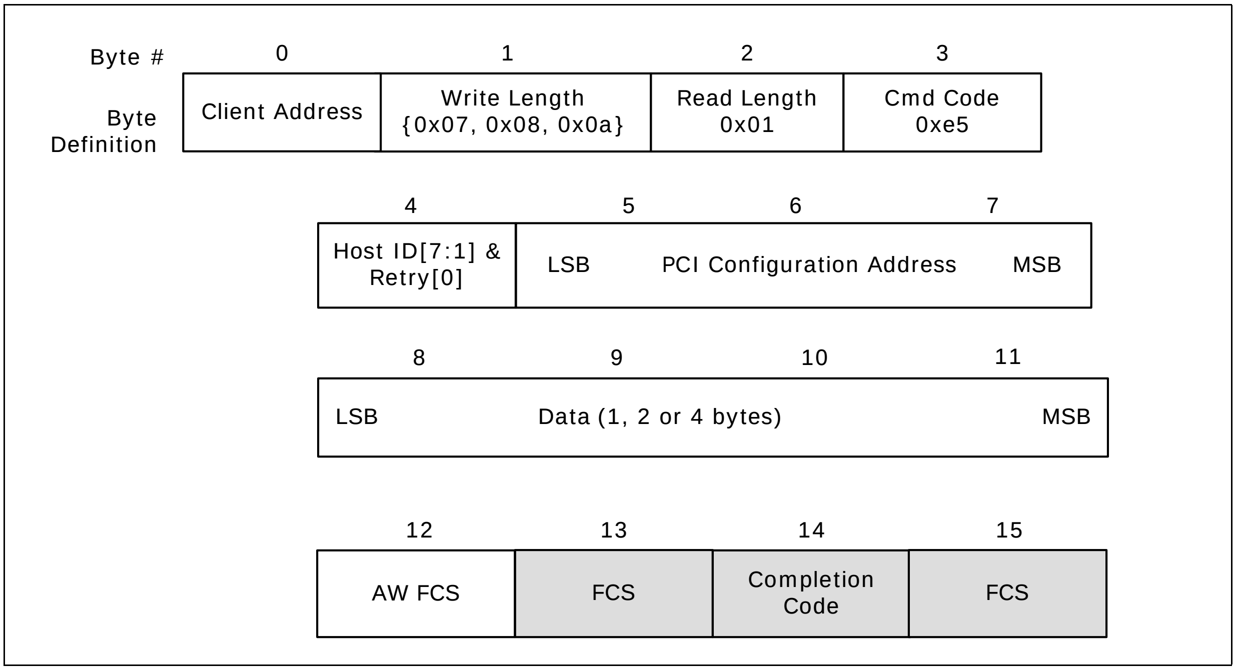

Note: The 3-byte PCI configuration address and write data field defined in Figure 2-48 are sent in standard PECI ordering with LSB first and MSB last.

2.5.2.10.2 Supported Responses

The typical client response is a passing FCS, a passing Completion Code and valid data. Under some conditions, the client's response will indicate a failure.

The PECI client response can also vary depending on the address and data. It will respond with a passing completion code if it successfully submits the request to the appropriate location and gets a response.

Table 2-14. WrPCIConfigLocal() Response Definition

Response	Meaning
Bad FCS	Electrical error or AW FCS failure
Abort FCS	Illegal command formatting (mismatched RL/WL/Command Code)
CC: 0x40	Command passed, data is valid.
CC: 0x80	Response timeout. The processor was not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor is not able to allocate resources for servicing this command at this time. Retry is appropriate.
CC: 0x82	The processor hardware resources required to service this command are in a low power state. Retry may be appropriate after modification of PECI wake mode behavior if appropriate.
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.

2.5.2.10.3 WrPCIConfigLocal() Capabilities

On the processor PECI clients, the PECI WrPCIConfigLocal() command provides a method for programming certain integrated memory controller and IIO functions as described in Table 2-15. Refer to the *Intel® Xeon® Processor E5 Product Family Datasheet Volume Two* for more details on specific register definitions. It also enables writing to processor REUT (Robust Electrical Unified Test) registers associated with the Intel QPI, PCIe* and DDR3 functions.

Table 2-15. WrPCIConfigLocal() Memory Controller and IIO Device/Function Support

Bus	Device	Function	Offset Range	Description
0000	0-5	0-7	000-FFFh	Integrated I/O (IIO) Configuration Registers
0001	15	0	104h-127h	Integrated Memory Controller MemHot Registers
0001	15	0	180h-1AFh	Integrated Memory Controller SMBus Registers
0001	15	1	080h-0CFh	Integrated Memory Controller RAS Registers (Scrub/Spare)
0001	16	0, 1, 4, 5	104h-18Bh 1F4h-1FFh	Integrated Memory Controller Thermal Control Registers
0001	16	2, 3, 6, 7	104h-147h	Integrated Memory Controller Error Registers

2.5.3 Client Management

2.5.3.1 Power-up Sequencing

The PECI client will not be available when the PWRGOOD signal is de-asserted. Any transactions on the bus during this time will be completely ignored, and the host will read the response from the client as all zeroes. PECI client initialization is completed approximately 100 µS after the PWRGOOD assertion. This is represented by the start of the PECI Client “Data Not Ready” (DNR) phase in Figure 2-49. While in this phase, the PECI client will respond normally to the Ping() and GetDIB() commands and return the highest processor die temperature of 0x0000 to the GetTemp() command. All other commands will get a ‘Response Timeout’ completion in the DNR phase as shown in Table 2-16. All PECI services with the exception of core MSR space accesses become available ~500 µS after RESET_N de-assertion as shown in Figure 2-49. PECI will be fully functional with all services including core accesses being available when the core comes out of reset upon completion of the RESET microcode execution.

In the event of the occurrence of a fatal or catastrophic error, all PECI services with the exception of core MSR space accesses will be available during the DNR phase to facilitate debug through configuration space accesses.

Table 2-16. PECI Client Response During Power-Up (Sheet 1 of 2)

Command	Response During ‘Data Not Ready’	Response During ‘Available Except Core Services’
Ping()	Fully functional	Fully functional
GetDIB()	Fully functional	Fully functional
GetTemp()	Client responds with a ‘hot’ reading or 0x0000	Fully functional
RdPkgConfig()	Client responds with a timeout completion code of 0x81	Fully functional
WrPkgConfig()	Client responds with a timeout completion code of 0x81	Fully functional
RdIAMSR()	Client responds with a timeout completion code of 0x81	Client responds with a timeout completion code of 0x81

Table 2-16. PECI Client Response During Power-Up (Sheet 2 of 2)

Command	Response During 'Data Not Ready'	Response During 'Available Except Core Services'
RdPCIConfigLocal()	Client responds with a timeout completion code of 0x81	Fully functional
WrPCIConfigLocal()	Client responds with a timeout completion code of 0x81	Fully functional
RdPCIConfig()	Client responds with a timeout completion code of 0x81	Fully functional

In the event that the processor is tri-stated using power-on-configuration controls, the PECI client will also be tri-stated. Processor tri-state controls are described in Section 7.3, "Power-On Configuration (POC) Options".

Figure 2-49. The Processor PECI Power-up Timeline()

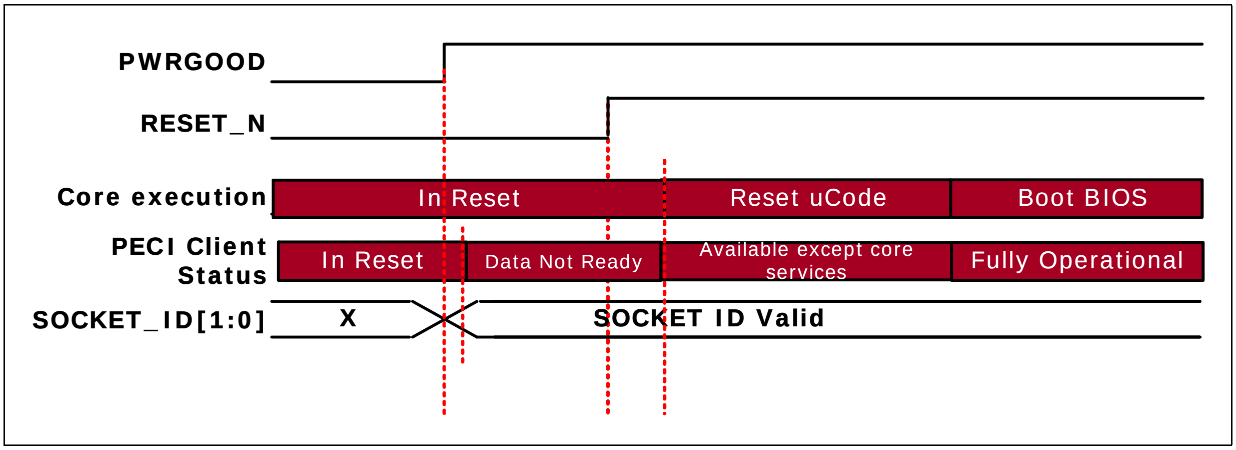

2.5.3.2 Device Discovery

The PECI client is available on all processors. The presence of a PECI enabled processor in a CPU socket can be confirmed by using the Ping() command described in Section 2.5.2.1. Positive identification of the PECI revision number can be achieved by issuing the GetDIB() command. The revision number acts as a reference to the PECI specification document applicable to the processor client definition. Please refer to Section 2.5.2.2 for details on GetDIB response formatting.

2.5.3.3 Client Addressing

The PECI client assumes a default address of 0x30. The PECI client address for the processor is configured through the settings of the SOCKET_ID[1:0] signals. Each processor socket in the system requires that the two SOCKET_ID signals be configured to a different PECI addresses. Strapping the SOCKET_ID[1:0] pins results in the client addresses shown in Table 2-17. These package strap(s) are evaluated at the assertion of PWRGOOD (as depicted in Figure 2-49). Refer to the appropriate Platform Design Guide (PDG) for recommended resistor values for establishing non-default SOCKET_ID settings.

The client address may not be changed after PWRGOOD assertion, until the next power cycle on the processor. Removal of a processor from its socket or tri-stating a processor will have no impact to the remaining non-tri-stated PECI client addresses. Since each socket in the system should have a unique PECI address, the SOCKET_ID strapping is required to be unique for each socket.

Table 2-17. SOCKET ID Strapping

SOCKET_ID[1] Strap	SOCKET_ID[0] Strap	PECI Client Address
Ground	Ground	0x30
Ground	V_{TT}	0x31
V_{TT}	Ground	0x32
V_{TT}	V_{TT}	0x33

2.5.3.4 C-states

The processor PECI client may be fully functional in most core and package C-states.

- The Ping(), GetDIB(), GetTemp(), RdPkgConfig() and WrPkgConfig() commands have no measurable impact on CPU power in any of the core or package C-states.
- The RdIAMSR() command will complete normally unless the targeted core is in a C-state that is C3 or deeper. The PECI client will respond with a completion code of 0x82 (see Table 2-22 for definition) for RdIAMSR() accesses in core C-states that are C3 or deeper.
- The RdPCIConfigLocal(), WrPCIConfigLocal(), and RdPCIConfig() commands will not impact the core C-states but may have a measurable impact on the package C-state. The PECI client will successfully return data without impacting package C-state if the resources needed to service the command are not in a low power state.
  — If the resources required to service the command are in a low power state, the PECI client will respond with a completion code of 0x82 (see Table 2-22 for definition). If this is the case, setting the “Wake on PECI” mode bit as described in Section 2.5.2.6 can cause a package ‘pop-up’ to the C2 state and enable successful completion of the command. The exact power impact of a pop-up to C2 will vary by product SKU, the C-state from which the pop-up is initiated and the negotiated PECI bit rate.

Table 2-18. Power Impact of PECI Commands vs. C-states

Command	Power Impact
Ping()	Not measurable
GetDIB()	Not measurable
GetTemp()	Not measurable
RdPkgConfig()	Not measurable
WrPkgConfig()	Not measurable
RdIAMSR()	Not measurable. PECI client will not return valid data in core C-state that is C3 or deeper
RdPCIConfigLocal()	May require package ‘pop-up’ to C2 state
WrPCIConfigLocal()	May require package ‘pop-up’ to C2 state
RdPCIConfig()	May require package ‘pop-up’ to C2 state

2.5.3.5 S-states

The processor PECI client is always guaranteed to be operational in the S0 sleep state.

- The Ping(), GetDIB(), GetTemp(), RdPkgConfig(), WrPkgConfig(), RdPCIConfigLocal() and WrPCIConfigLocal() will be fully operational in S0 and S1. Responses in S3 or deeper states are dependent on POWERGOOD assertion status.
- The RdPCIConfig() and RdIAMSR() responses are guaranteed in S0 only. Behavior in S1 or deeper states is indeterminate.
- PECI behavior is indeterminate in the S3, S4 and S5 states and responses to PECI originator requests when the PECI client is in these states cannot be guaranteed.

2.5.3.6 Processor Reset

The processor PECI client is fully reset on all RESET_N assertions. Upon deassertion of RESET_N where power is maintained to the processor (otherwise known as a 'warm reset'), the following are true:

- The PECI client assumes a bus Idle state.
- The Thermal Filtering Constant is retained.
- PECI SOCKET_ID is retained.
- GetTemp() reading resets to 0x0000.
- Any transaction in progress is aborted by the client (as measured by the client no longer participating in the response).
- The processor client is otherwise reset to a default configuration.

The assertion of the CPU_ONLY_RESET signal does not reset the processor PECI client. As such, it will have no impact on the basic PECI commands, namely the Ping(), GetTemp() and GetDIB(). However, it is likely that other PECI commands that utilize processor resources being reset will receive a 'resource unavailable' response till the reset sequence is completed.

2.5.3.7 System Service Processor (SSP) Mode Support

Sockets in SSP mode have limited PECI command support. Only the following PECI commands will be supported while in SSP mode. Other PECI commands are not guaranteed to complete in this mode.

- Ping
- RdPCIConfigLocal
- WrPCIConfigLocal (all uncore and IIO CSRs within the processor PCI configuration space will be accessible)
- RdPkgConfig (Index 0 only)

Sockets remain in SSP mode until the "Go" handshake is received. This is applicable to the following SSP modes.

2.5.3.7.1 BMC INIT Mode

The BMC INIT boot mode is used to provide a quick and efficient means to transfer responsibility for uncore configuration to a service processor like the BMC. In this mode, the socket performs a minimal amount of internal configuration and then waits for the BMC or service processor to complete the initialization.

2.5.3.7.2 Link Init Mode

In cases where the socket is not one Intel QPI hop away from the Firmware Agent socket, or a working link to the Firmware Agent socket cannot be resolved, the socket is placed in Link Init mode. The socket performs a minimal amount of internal configuration and waits for complete configuration by BIOS.

2.5.3.8 Processor Error Handling

Availability of PECI services may be affected by the processor PECI client error status. Server manageability requirements place a strong emphasis on continued availability of PECI services to facilitate logging and debug of the error condition.

- Most processor PECI client services are available in the event of a CAT_ERR_N assertion though they cannot be guaranteed.
- The Ping(), GetDIB(), GetTemp(), RdPkgConfig() and WrPkgConfig() commands will be serviced if the source of the CAT_ERR_N assertion is not in the processor power control unit hardware, firmware or associated register logic. Additionally, the RdPCIConfigLocal() and WrPCIConfigLocal() commands may also be serviced in this case.
- It is recommended that the PECI originator read Index 0/Parameter 5 using the RdPkgConfig() command to debug the CAT_ERR_N assertion.
    - — The PECI client will return the 0x91 completion code if the CAT_ERR_N assertion is caused by the PCU hardware, firmware or associated logic errors. In such an event, only the Ping(), GetTemp() and GetDIB() PECI commands may be serviced. All other processor PECI services will be unavailable and further debug of the processor error status will not be possible.
    - — If the PECI client returns a passing completion code, the originator should use the response data to determine the cause of the CAT_ERR_N assertion. In such an event, it is also recommended that the PECI originator determine the exact suite of available PECI client services by issuing each of the PECI commands. The processor will issue ‘timeout’ responses for those services that may not be available.
    - — If the PECI client continues to return the 0x81 completion code in response to multiple retries of the RdPkgConfig() command, no PECI services, with the exception of the Ping(), GetTemp() and GetDIB(), will be guaranteed.
- The RdIAMSR() command may be serviced during a CAT_ERR_N assertion though it cannot be guaranteed.

2.5.3.9 Originator Retry and Timeout Policy

The PECI originator may need to retry a command if the processor PECI client responds with a ‘response timeout’ completion code or a bad Read FCS. In each instance, the processor PECI client may have started the operation but not completed it yet. When the 'retry' bit is set, the PECI client will ignore a new request if it exactly matches a previous valid request.

The processor PECI client will not clear the semaphore that was acquired to service the request until the originator sends the ‘retry’ request in a timely fashion to successfully retrieve the response data. In the absence of any automatic timeouts, this could tie up shared resources and result in artificial bandwidth conflicts.

2.5.3.10 Enumerating PECI Client Capabilities

The PECI host originator should be designed to support all optional but desirable features from all processors of interest. Each feature has a discovery method and response code that indicates availability on the destination PECI client.

The first step in the enumeration process would be for the PECI host to confirm the Revision Number through the use of the GetDIB() command. The revision number returned by the PECI client processor always maps to the revision number of the PECI specification that it is designed to. The Minor Revision Number as described in Table 2-2 may be used to identify the subset of PECI commands that the processor in question supports for any major PECI revision.

The next step in the enumeration process is to utilize the desired command suite in a real execution context. If the Write FCS response is an Abort FCS or if the data returned includes an “Unknown/Invalid/Illegal Request” completion code (0x90), then the command is unsupported.

Enumerating known commands without real, execution context data, or attempting undefined commands, is dangerous because a write command could result in unexpected behavior if the data is not properly formatted. Methods for enumerating write commands using carefully constructed and innocuous data are possible, but are not guaranteed by the PECI client definition.

This enumeration procedure is not robust enough to detect differences in bit definitions or data interpretation in the message payload or client response. Instead, it is only designed to enumerate discrete features.

2.5.4 Multi-Domain Commands

The processor does not support multiple domains, but it is possible that future products will, and the following tables are included as a reference for domain-specific definitions.

Table 2-19. Domain ID Definition

Domain ID	Domain Number
0b01	**0**
0b10	1

Table 2-20. Multi-Domain Command Code Reference

Command Name	Domain 0 Code	Domain 1 Code
GetTemp()	**0x01**	0x02
RdPkgConfig()	**0xa1**	0xa2
WrPkgConfig()	**0xa5**	0xa6
RdIAMSR()	**0xb1**	0xb2
RdPCIConfig()	**0x61**	0x62
RdPCIConfigLocal()	**0xe1**	0xe2
WrPCIConfigLocal()	**0xe5**	0xe6

2.5.5 Client Responses

2.5.5.1 Abort FCS

The Client responds with an Abort FCS under the following conditions:

- The decoded command is not understood or not supported on this processor (this includes good command codes with bad Read Length or Write Length bytes).
- Assured Write FCS (AW FCS) failure. Under most circumstances, an Assured Write failure will appear as a bad FCS. However, when an originator issues a poorly formatted command with a miscalculated AW FCS, the client will intentionally abort the FCS in order to guarantee originator notification.

2.5.5.2 Completion Codes

Some PECI commands respond with a completion code byte. These codes are designed to communicate the pass/fail status of the command and may also provide more detailed information regarding the class of pass or fail. For all commands listed in Section 2.5.2 that support completion codes, the definition in the following table applies. Throughout this document, a completion code reference may be abbreviated with 'CC'.

An originator that is decoding these commands can apply a simple mask as shown in Table 2-21 to determine a pass or fail. Bit 7 is always set on a command that did not complete successfully and is cleared on a passing command.

Table 2-21. Completion Code Pass/Fail Mask

0xxx xxxxb	Command passed
1xxx xxxxb	Command failed

Table 2-22. Device Specific Completion Code (CC) Definition

Completion Code	Description
0x40	Command Passed
CC: 0x80	Response timeout. The processor was not able to generate the required response in a timely fashion. Retry is appropriate.
CC: 0x81	Response timeout. The processor was not able to allocate resources for servicing this command. Retry is appropriate.
CC: 0x82	The processor hardware resources required to service this command are in a low power state. Retry may be appropriate after modification of PECI wake mode behavior if appropriate.
CC: 0x83-8F	Reserved
CC: 0x90	Unknown/Invalid/Illegal Request
CC: 0x91	PECI control hardware, firmware or associated logic error. The processor is unable to process the request.
CC: 0x92-9F	Reserved

Note: The codes explicitly defined in Table 2-22 may be useful in PECI originator response algorithms. Reserved or undefined codes may also be generated by a PECI client device, and the originating agent must be capable of tolerating any code. The Pass/Fail mask defined in Table 2-21 applies to all codes, and general response policies may be based on this information. Refer to Section 2.5.6 for originator response policies and recommendations.

2.5.6 Originator Responses

The simplest policy that an originator may employ in response to receipt of a failing completion code is to retry the request. However, certain completion codes or FCS responses are indicative of an error in command encoding and a retry will not result in a different response from the client. Furthermore, the message originator must have a response policy in the event of successive failure responses. Refer to Table 2-22 for originator response guidelines.

Refer to the definition of each command in Section 2.5.2 for a specific definition of possible command codes or FCS responses for a given command. The following response policy definition is generic, and more advanced response policies may be employed at the discretion of the originator developer.

Table 2-23. Originator Response Guidelines

Response	After 1 Attempt	After 3 Attempts
Bad FCS	Retry	Fail with PECI client device error.
Abort FCS	Retry	Fail with PECI client device error if command was not illegal or malformed.
CC: 0x8x	Retry	The PECI client has failed in its attempts to generate a response. Notify application layer.
CC: 0x9x	Abandon any further attempts and notify application layer	N/A
None (all 0's)	Force bus idle (drive low) for 1 mS and retry	Fail with PECI client device error. Client may not be alive or may be otherwise unresponsive (for example, it could be in RESET).
CC: 0x4x	Pass	N/A
Good FCS	Pass	N/A

2.5.7 DTS Temperature Data

2.5.7.1 Format

The temperature is formatted in a 16-bit, 2's complement value representing a number of 1/64 degrees centigrade. This format allows temperatures in a range of ±512° C to be reported to approximately a 0.016° C resolution.

Figure 2-50. Temperature Sensor Data Format

MSB Upper nibble					MSB Lower nibble					LSB Upper nibble					LSB Lower nibble			
S	x	x	x		x	x	x	x		x	x	x	x		x	x	x	x
Sign	Integer Value (0-511)											Fractional Value (~0.016)						

2.5.7.2 Interpretation

The resolution of the processor's Digital Thermal Sensor (DTS) is approximately 1°C, which can be confirmed by a RDMSR from the IA32_THERM_STATUS MSR where it is architecturally defined. The MSR read will return only bits [13:6] of the PECI temperature sensor data defined in Figure 2-50. PECI temperatures are sent through a configurable low-pass filter prior to delivery in the GetTemp() response data. The output of this filter produces temperatures at the full 1/64°C resolution even though the DTS itself is not this accurate.

Temperature readings from the processor are always negative in a 2's complement format, and imply an offset from the processor T_{jmax} (PECI = 0). For example, if the processor T_{jmax} is 100°C, a PECI thermal reading of -10 implies that the processor is running at approximately 10°C below T_{jmax} or at 90°C. PECI temperature readings are not reliable at temperatures above T_{jmax} since the processor is outside its operating range and hence, PECI temperature readings are never positive.

The changes in PECI data counts are approximately linear in relation to changes in temperature in degrees centigrade. A change of '1' in the PECI count represents roughly a temperature change of 1 degree centigrade. This linearity is approximate and cannot be guaranteed over the entire range of PECI temperatures, especially as the offset from the maximum PECI temperature (zero) increases.

2.5.7.3 Temperature Filtering

The processor digital thermal sensor (DTS) provides an improved capability to monitor device hot spots, which inherently leads to more varying temperature readings over short time intervals. Coupled with the fact that typical fan speed controllers may only read temperatures at 4Hz, it is necessary for the thermal readings to reflect thermal trends and not instantaneous readings. Therefore, PECI supports a configurable low-pass temperature filtering function that is expressed by the equation:

$$T_N = (1-\alpha) * T_{N-1} + \alpha * T_{SAMPLE}$$

where T_N and T_{N-1} are the current and previous averaged PECI temperature values respectively, T_{SAMPLE} is the current PECI temperature sample value and the variable 'α' = $1/2^X$, where 'X' is the 'Thermal Averaging Constant' that is programmable as described in Section 2.5.2.6.21.

2.5.7.4 Reserved Values

Several values well out of the operational range are reserved to signal temperature sensor errors. These are summarized in Table 2-24.

Table 2-24. Error Codes and Descriptions

Error Code	Description
0x8000	General Sensor Error (GSE)
0x8001	Reserved
0x8002	Sensor is operational, but has detected a temperature below its operational range (underflow)
0x8003-0x81ff	Reserved

§

3 Technologies

3.1 Intel® Virtualization Technology (Intel® VT)

Intel® Virtualization Technology (Intel® VT) makes a single system appear as multiple independent systems to software. This allows multiple, independent operating systems to run simultaneously on a single system. Intel VT comprises technology components to support virtualization of platforms based on Intel architecture microprocessors and chipsets.

- **Intel® Virtualization Technology (Intel® VT) for Intel® 64 and IA-32 Intel® Architecture (Intel® VT-x)** adds hardware support in the processor to improve the virtualization performance and robustness. Intel VT-x specifications and functional descriptions are included in the Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B and is available at http://www.intel.com/products/processor/manuals/index.htm
- **Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)** adds processor and uncore implementations to support and improve I/O virtualization performance and robustness. The Intel VT-d spec and other Intel VT documents can be referenced at http://www.intel.com/technology/virtualization/index.htm.

3.1.1 Intel VT-x Objectives

Intel VT-x provides hardware acceleration for virtualization of IA platforms. Virtual Machine Monitor (VMM) can use Intel VT-x features to provide improved reliable virtualized platform. By using Intel VT-x, a VMM is:

- **Robust:** VMMs no longer need to use para-virtualization or binary translation. This means that they will be able to run off-the-shelf OS's and applications without any special steps.
- **Enhanced:** Intel VT enables VMMs to run 64-bit guest operating systems on IA x86 processors.
- **More reliable:** Due to the hardware support, VMMs can now be smaller, less complex, and more efficient. This improves reliability and availability and reduces the potential for software conflicts.
- **More secure:** The use of hardware transitions in the VMM strengthens the isolation of VMs and further prevents corruption of one VM from affecting others on the same system.

3.1.2 Intel VT-x Features

The processor core supports the following Intel VT-x features:

- Extended Page Tables (EPT)
  - — hardware assisted page table virtualization
  - — eliminates VM exits from guest OS to the VMM for shadow page-table maintenance
- Virtual Processor IDs (VPID)
  - — Ability to assign a VM ID to tag processor core hardware structures (for example, TLBs)
  - — This avoids flushes on VM transitions to give a lower-cost VM transition time and an overall reduction in virtualization overhead.
- Guest Preemption Timer
  - — Mechanism for a VMM to preempt the execution of a guest OS after an amount of time specified by the VMM. The VMM sets a timer value before entering a guest
  - — The feature aids VMM developers in flexibility and Quality of Service (QoS) guarantees
- Descriptor-Table Exiting
  - — Descriptor-table exiting allows a VMM to protect a guest OS from internal (malicious software based) attack by preventing relocation of key system data structures like IDT (interrupt descriptor table), GDT (global descriptor table), LDT (local descriptor table), and TSS (task segment selector).
  - — A VMM using this feature can intercept (by a VM exit) attempts to relocate these data structures and prevent them from being tampered by malicious software.
- Pause Loop Exiting (PLE)
  - — PLE aims to improve virtualization performance and enhance the scaling of virtual machines with multiple virtual processors
  - — PLE attempts to detect lock-holder preemption in a VM and helps the VMM to make better scheduling decisions

3.1.3 Intel VT-d Objectives

The key Intel VT-d objectives are domain-based isolation and hardware-based virtualization. A domain can be abstractly defined as an isolated environment in a platform to which a subset of host physical memory is allocated. Virtualization allows for the creation of one or more partitions on a single system. This could be multiple partitions in the same operating system, or there can be multiple operating system instances running on the same system – offering benefits such as system consolidation, legacy migration, activity partitioning or security.

3.1.3.1 Intel VT-d Features Supported

The processor supports the following Intel VT-d features:

- Root entry, context entry, and default context
- Support for 4-K page sizes only
- Support for register-based fault recording only (for single entry only) and support for MSI interrupts for faults

— Support for fault collapsing based on Requester ID

- Support for both leaf and non-leaf caching
- Support for boot protection of default page table
  — Support for non-caching of invalid page table entries
- Support for hardware based flushing of translated but pending writes and pending reads upon IOTLB invalidation.
- Support for page-selective IOTLB invalidation.
- Support for ARI (Alternative Requester ID - a PCI SIG ECR for increasing the function number count in a PCIe* device) to support IOV devices.
- Improved invalidation architecture
- End point caching support (ATS)
- Interrupt remapping

3.1.4 Intel Virtualization Technology Processor Extensions

The processor supports the following Intel VT Processor Extensions features:

- Large Intel VT-d Pages
  — Adds 2 MB and 1 GB page sizes to Intel VT-d implementations
  — Matches current support for Extended Page Tables (EPT)
  — Ability to share CPU's EPT page-table (with super-pages) with Intel VT-d
  — Benefits:
    - Less memory foot-print for I/O page-tables when using super-pages
    - Potential for improved performance - Due to shorter page-walks, allows hardware optimization for IOTLB
- Transition latency reductions expected to improve virtualization performance without the need for VMM enabling. This reduces the VMM overheads further and increase virtualization performance.

3.2 Security Technologies

3.2.1 Intel® Trusted Execution Technology

Intel® Trusted Execution Technology (Intel® TXT) defines platform-level enhancements that provide the building blocks for creating trusted platforms.

The Intel TXT platform helps to provide the authenticity of the controlling environment such that those wishing to rely on the platform can make an appropriate trust decision. The Intel TXT platform determines the identity of the controlling environment by accurately measuring and verifying the controlling software.

Another aspect of the trust decision is the ability of the platform to resist attempts to change the controlling environment. The Intel TXT platform will resist attempts by software processes to change the controlling environment or bypass the bounds set by the controlling environment.

Intel TXT is a set of extensions designed to provide a measured and controlled launch of system software that will then establish a protected environment for itself and any additional software that it may execute.

These extensions enhance two areas:

- The launching of the Measured Launched Environment (MLE).
- The protection of the MLE from potential corruption.

The enhanced platform provides these launch and control interfaces using Safer Mode Extensions (SMX).

The SMX interface includes the following functions:

- Measured/Verified launch of the MLE.
- Mechanisms to ensure the above measurement is protected and stored in a secure location.
- Protection mechanisms that allow the MLE to control attempts to modify itself.

For more information refer to the *Intel® Trusted Execution Technology Software Development Guide.* For more information on Intel Trusted Execution Technology, see http://www.intel.com/technology/security/

3.2.2 Intel Trusted Execution Technology – Server Extensions

- Software binary compatible with Intel Trusted Execution Technology Server Extensions
- Provides measurement of runtime firmware, including SMM
- Enables run-time firmware in trusted session: BIOS and SSP
- Covers support for existing and expected future Server RAS features
- Only requires portions of BIOS to be trusted, for example, Option ROMs need not be trusted
- Supports S3 State without teardown: Since BIOS is part of the trust chain

3.2.3 Intel® Advanced Encryption Standard Instructions (Intel® AES-NI)

These instructions enable fast and secure data encryption and decryption, using the Intel® AES New Instructions (Intel® AES-NI), which is defined by FIPS Publication number 197. Since Intel AES-NI is the dominant block cipher, and it is deployed in various protocols, the new instructions will be valuable for a wide range of applications.

The architecture consists of six instructions that offer full hardware support for Intel AES-NI. Four instructions support the Intel AES-NI encryption and decryption, and the other two instructions support the Intel AES-NI key expansion. Together, they offer a significant increase in performance compared to pure software implementations.

The Intel AES-NI instructions have the flexibility to support all three standard Intel AES-NI key lengths, all standard modes of operation, and even some nonstandard or future variants.

Beyond improving performance, the Intel AES-NI instructions provide important security benefits. Since the instructions run in data-independent time and do not use lookup tables, they help in eliminating the major timing and cache-based attacks that threaten table-based software implementations of Intel AES-NI. In addition, these instructions make AES simple to implement, with reduced code size. This helps reducing the risk of inadvertent introduction of security flaws, such as difficult-to-detect side channel leaks.

3.2.4 Execute Disable Bit

Intel's Execute Disable Bit functionality can help prevent certain classes of malicious buffer overflow attacks when combined with a supporting operating system.

- Allows the processor to classify areas in memory by where application code can execute and where it cannot.
- When a malicious worm attempts to insert code in the buffer, the processor disables code execution, preventing damage and worm propagation.

3.3 Intel® Hyper-Threading Technology

The processor supports Intel® Hyper-Threading Technology (Intel® HT Technology), which allows an execution core to function as two logical processors. While some execution resources such as caches, execution units, and buses are shared, each logical processor has its own architectural state with its own set of general-purpose registers and control registers. This feature must be enabled via the BIOS and requires operating system support. For more information on Intel Hyper-Threading Technology, see http://www.intel.com/products/ht/hyperthreading_more.htm.

3.4 Intel® Turbo Boost Technology

Intel® Turbo Boost Technology is a feature that allows the processor to opportunistically and automatically run faster than its rated operating frequency if it is operating below power, temperature, and current limits. The result is increased performance in multi-threaded and single threaded workloads. It should be enabled in the BIOS for the processor to operate with maximum performance.

3.4.1 Intel® Turbo Boost Operating Frequency

The processor's rated frequency assumes that all execution cores are running an application at the thermal design power (TDP). However, under typical operation, not all cores are active. Therefore most applications are consuming less than the TDP at the rated frequency. To take advantage of the available TDP headroom, the active cores can increase their operating frequency.

To determine the highest performance frequency amongst active cores, the processor takes the following into consideration:

- The number of cores operating in the C0 state.
- The estimated current consumption.
- The estimated power consumption.
- The die temperature.

Any of these factors can affect the maximum frequency for a given workload. If the power, current, or thermal limit is reached, the processor will automatically reduce the frequency to stay with its TDP limit.

Note: Intel Turbo Boost Technology is only active if the operating system is requesting the P0 state. For more information on P-states and C-states refer to Section 4, "Power Management".

3.5 Enhanced Intel SpeedStep® Technology

The processor supports Enhanced Intel SpeedStep® Technology as an advanced means of enabling very high performance while also meeting the power-conservation needs of the platform.

Enhanced Intel SpeedStep Technology builds upon that architecture using design strategies that include the following:

- **Separation between Voltage and Frequency Changes**. By stepping voltage up and down in small increments separately from frequency changes, the processor is able to reduce periods of system unavailability (which occur during frequency change). Thus, the system is able to transition between voltage and frequency states more often, providing improved power/performance balance.
- **Clock Partitioning and Recovery**. The bus clock continues running during state transition, even when the core clock and Phase-Locked Loop are stopped, which allows logic to remain active. The core clock is also able to restart more quickly under Enhanced Intel SpeedStep Technology.

For additional information on Enhanced Intel SpeedStep Technology see Section 4.2.1.

3.6 Intel® Intelligent Power Technology

Intel® Intelligent Power Technology conserves power while delivering advanced power-management capabilities at the rack, group, and data center level. Providing the highest system-level performance per watt with “Automated Low Power States” and “Integrated Power Gates”. Improvements to this processor generation are:

- Intel Network Power Management Technology
- Intel Power Tuning Technology

For more information on Intel Intelligent Power Technology, see this link http://www.intel.com/technology/intelligentpower/.

3.7 Intel® Advanced Vector Extensions (Intel® AVX)

Intel® Advanced Vector Extensions (Intel® AVX) is a new 256-bit vector SIMD extension of Intel Architecture. The introduction of Intel AVX starts with the 2nd Generation Intel(r) Core(TM) Processor Family. Intel AVX accelerates the trend of parallel computation in general purpose applications like image, video, and audio processing, engineering applications such as 3D modeling and analysis, scientific simulation, and financial analysts.

Intel AVX is a comprehensive ISA extension of the Intel 64 Architecture. The main elements of Intel AVX are:

- Support for wider vector data (up to 256-bit) for floating-point computation.
- Efficient instruction encoding scheme that supports 3 operand syntax and headroom for future extensions.
- Flexibility in programming environment, ranging from branch handling to relaxed memory alignment requirements.
- New data manipulation and arithmetic compute primitives, including broadcast, permute, fused-multiply-add, and so forth.

The key advantages of Intel AVX are:

- **Performance** - Intel AVX can accelerate application performance via data parallelism and scalable hardware infrastructure across existing and new application domains:
  - — 256-bit vector data sets can be processed up to twice the throughput of 128-bit data sets.
  - — Application performance can scale up with number of hardware threads and number of cores.
  - — Application domain can scale out with advanced platform interconnect fabrics, such as Intel QPI.
- **Power Efficiency** - Intel AVX is extremely power efficient. Incremental power is insignificant when the instructions are unused or scarcely used. Combined with the high performance that it can deliver, applications that lend themselves heavily to using Intel AVX can be much more energy efficient and realize a higher performance-per-watt.
- **Extensibility** - Intel AVX has built-in extensibility for the future vector extensions:
  - — OS context management for vector-widths beyond 256 bits is streamlined.
  - — Efficient instruction encoding allows unlimited functional enhancements:
    - Vector width support beyond 256 bits
    - 256-bit Vector Integer processing
    - Additional computational and/or data manipulation primitives.
- **Compatibility** - Intel AVX is backward compatible with previous ISA extensions including Intel® SSE4:
  - — Existing Intel SSE applications/library can:
    - Run unmodified and benefit from processor enhancements
    - Recompile existing Intel SSE intrinsic using compilers that generate Intel AVX code
    - Inter-operate with library ported to Intel AVX
  - — Applications compiled with Intel AVX can inter-operate with existing Intel SSE libraries.

3.8 Intel® Dynamic Power Technology (Intel® DPT)

Intel® Dynamic Power Technology (Intel® DPT) (Memory Power Management) is a platform feature with the ability to transition memory components into various low power states based on workload requirements. The Intel® Xeon® processor E5-2400 product family platform supports Dynamic CKE (hardware assisted) and Memory Self Refresh (software assisted). For further details refer to the *ACPI Specifications for Memory Power Management* document.

§

4 Power Management

This chapter provides information on the following power management topics:

- ACPI States
- System States
- Processor Core/Package States
- Integrated Memory Controller (IMC) and System Memory States
- Direct Media Interface Gen 2 (DMI2)/PCI Express* Link States
- Intel QuickPath Interconnect States

4.1 ACPI States Supported

The ACPI states supported by the processor are described in this section.

4.1.1 System States

Table 4-1. System States

State	Description
G0/S0	Full On
G1/S3-Cold	Suspend-to-RAM (STR). Context saved to memory
G1/S4	Suspend-to-Disk (STD). All power lost (except wakeup on PCH).
G2/S5	Soft off. All power lost (except wakeup on PCH). Total reboot.
G3	Mechanical off. All power removed from system.

4.1.2 Processor Package and Core States

Table 4-2 lists the package C-state support as: 1) the shallowest core C-state that allows entry into the package C-state, 2) the additional factors that will restrict the state from going any deeper, and 3) the actions taken with respect to the Ring Vcc, PLL state and LLC.

Table 4-3 lists the processor core C-states support.

Table 4-2. Package C-State Support (Sheet 1 of 2)

Package C-State	Core States	Limiting Factors	Retention and PLL-Off	LLC Fully Flushed	Notes[1]
PC0 - Active	CC0	N/A	No	No	2
PC2 - Snoopable Idle	CC3-CC7	• PCIe*/PCH and Remote Socket Snoops • PCIe*/PCH and Remote Socket Accesses • Interrupt response time requirement • DMI Sidebands • Configuration Constraints	VccMin Freq = MinFreq PLL = ON	No	2

Table 4-2. Package C-State Support (Sheet 2 of 2)

Package C-State	Core States	Limiting Factors	Retention and PLL-Off	LLC Fully Flushed	Notes[1]
PC3 - Light Retention	at least one Core in C3	• Core C-state • Snoop Response Time • Interrupt Response Time • Non Snoop Response Time	Vcc = retention PLL = OFF	No	2,3,4
PC6 - Deeper Retention	CC6-CC7	• LLC ways open • Snoop Response Time • Non Snoop Response Time • Interrupt Response Time	Vcc = retention PLL = OFF	No	2,3,4

Notes:
1. Package C7 is not supported.
2. All package states are defined to be "E" states - such that they always exit back into the LFM point upon execution resume
3. The mapping of actions for PC3, and PC6 are suggestions - microcode will dynamically determine which actions should be taken based on the desired exit latency parameters.
4. CC3/CC6 will all use a voltage below the VccMin operational point; The exact voltage selected will be a function of the snoop and interrupt response time requirements made by the devices (PCIe* and DMI) and the operating system.

Table 4-3. Core C-State Support

Core C-State	Global Clock	PLL	L1/L2 Cache	Core VCC	Context
CC0	Running	On	Coherent	Active	Maintained
CC1	Stopped	On	Coherent	Active	Maintained
CC1E	Stopped	On	Coherent	Request LFM	Maintained
CC3	Stopped	On	Flushed to LLC	Request Retention	Maintained
CC6	Stopped	Off	Flushed to LLC	Power Gate	Flushed to LLC
CC7	Stopped	Off	Flushed to LLC	Power Gate	Flushed to LLC

4.1.3 Integrated Memory Controller States

Table 4-4. System Memory Power States (Sheet 1 of 2)

State	Description
Power Up/Normal Operation	CKE asserted. Active Mode, highest power consumption.
CKE Power Down	Opportunistic, per rank control after idle time: • Active Power Down (APD) (default mode) — CKE de-asserted. Power savings in this mode, relative to active idle state is about 55% of the memory power. Exiting this mode takes 3 – 5 DCLK cycles. • Pre-charge Power Down Fast Exit (PPDF) — CKE de-asserted. DLL-On. Also known as Fast CKE. Power savings in this mode, relative to active idle state is about 60% of the memory power. Exiting this mode takes 3 – 5 DCLK cycles. • Pre-charge Power Down Slow Exit (PPDS) — CKE de-asserted. DLL-Off. Also known as Slow CKE. Power savings in this mode, relative to active idle state is about 87% of the memory power. Exiting this mode takes 3 – 5 DCLK cycles until the first command is allowed and 16 cycles until first data is allowed. • Register CKE Power Down: — IBT-ON mode: Both CKE's are de-asserted, the Input Buffer Terminators (IBTs) are left "on". — IBT-OFF mode: Both CKE's are de-asserted, the Input Buffer Terminators (IBTs) are turned "off".

Table 4-4. System Memory Power States (Sheet 2 of 2)

State	Description
Self-Refresh	CKE de-asserted. In this mode, no transactions are executed and the system memory consumes the minimum possible power. Self refresh modes apply to all memory channels for the processor. • IO-MDLL Off: Option that sets the IO master DLL off when self refresh occurs. • PLL Off: Option that sets the PLL off when self refresh occurs. In addition, the register component found on registered DIMMs (RDIMMs) is complemented with the following power down states: — Clock Stopped Power Down with IBT-On — Clock Stopped Power Down with IBT-Off

4.1.4 DMI2/PCI Express* Link States

Table 4-5. DMI2/PCI Express* Link States

State	Description
L0	Full on – Active transfer state.
L1	Lowest Active State Power Management (ASPM) - Longer exit latency.

Note: L1 is only supported when the DMI2/PCI Express* port is operating as a PCI Express* port.

4.1.5 Intel QuickPath Interconnect States

Table 4-6. Intel QPI States

State	Description
L0	Link on. This is the power on active working state,
L0p	A lower power state from L0 that reduces the link from full width to half width
L1	A low power state with longer latency and lower power than L0s and is activated in conjunction with package C-states below C0.

4.1.6 G, S, and C State Combinations

Table 4-7. G, S and C State Combinations

Global (G) State	Sleep (S) State	Processor Core (C) State	Processor State	System Clocks	Description
G0	S0	C0	Full On	On	Full On
G0	S0	C1/C1E	Auto-Halt	On	Auto-Halt
G0	S0	C3	Deep Sleep	On	Deep Sleep
G0	S0	C6/C7	Deep Power Down	On	Deep Power Down
G1	S3	Power off		Off, except RTC	Suspend to RAM
G1	S4	Power off		Off, except RTC	Suspend to Disk
G2	S5	Power off		Off, except RTC	Soft Off
G3	N/A	Power off		Power off	Hard off

4.2 Processor Core/Package Power Management

While executing code, Enhanced Intel SpeedStep Technology optimizes the processor's frequency and core voltage based on workload. Each frequency and voltage operating point is defined by ACPI as a P-state. When the processor is not executing code, it is idle. A low-power idle state is defined by ACPI as a C-state. In general, lower power C-states have longer entry and exit latencies.

4.2.1 Enhanced Intel SpeedStep Technology

The following are the key features of Enhanced Intel SpeedStep Technology:

- Multiple frequency and voltage points for optimal performance and power efficiency. These operating points are known as P-states.
- Frequency selection is software controlled by writing to processor MSRs. The voltage is optimized based on temperature, leakage, power delivery loadline and dynamic capacitance.
  - If the target frequency is higher than the current frequency, V_{CC} is ramped up to an optimized voltage. This voltage is signaled by the SVID Bus to the voltage regulator. Once the voltage is established, the PLL locks on to the target frequency.
  - If the target frequency is lower than the current frequency, the PLL locks to the target frequency, then transitions to a lower voltage by signaling the target voltage on the SVID Bus.
  - All active processor cores share the same frequency and voltage. In a multi-core processor, the highest frequency P-state requested amongst all active cores is selected.
  - Software-requested transitions are accepted at any time. The processor has a new capability from the previous processor generation, it can preempt the previous transition and complete the new request without waiting for this request to complete.
- The processor controls voltage ramp rates internally to ensure glitch-free transitions.
- Because there is low transition latency between P-states, a significant number of transitions per second are possible.

4.2.2 Low-Power Idle States

When the processor is idle, low-power idle states (C-states) are used to save power. More power savings actions are taken for numerically higher C-states. However, higher C-states have longer exit and entry latencies. Resolution of C-states occurs at the thread, processor core, and processor package level. Thread level C-states are available if Hyper-Threading Technology is enabled. Entry and exit of the C-States at the thread and core level are shown in Figure 4-2.

Figure 4-1. Idle Power Management Breakdown of the Processor Cores

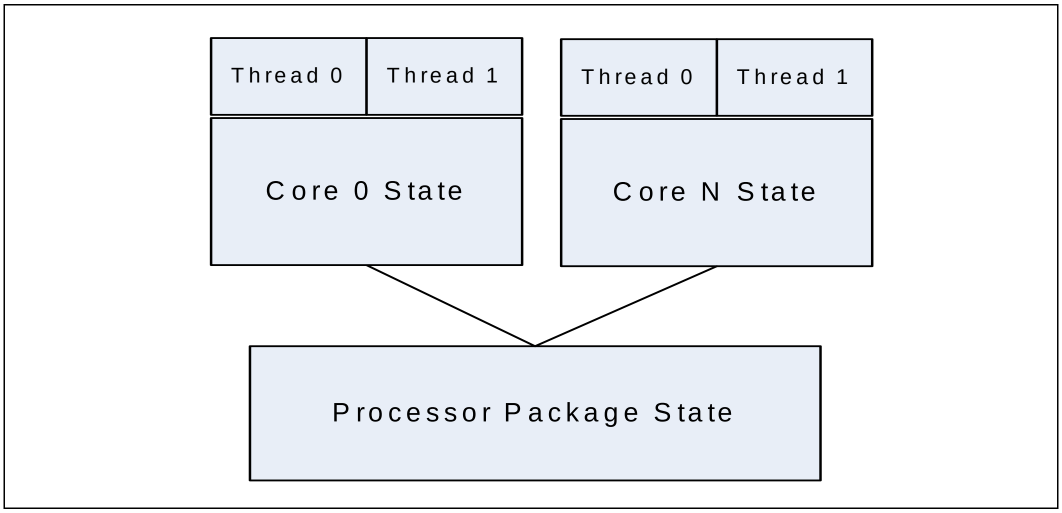

Figure 4-2. Thread and Core C-State Entry and Exit

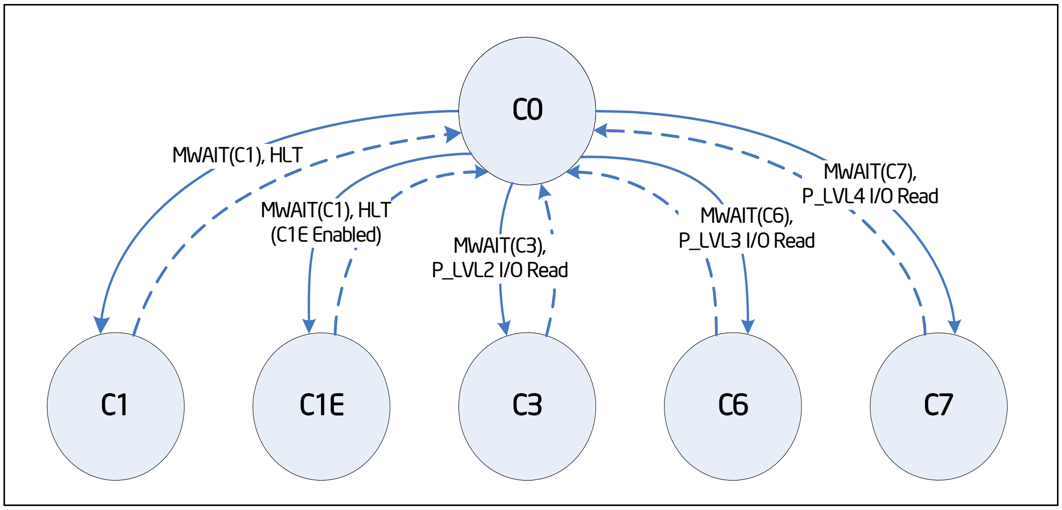

While individual threads can request low power C-states, power saving actions only take place once the core C-state is resolved. Core C-states are automatically resolved by the processor. For thread and core C-states, a transition to and from C0 is required before entering any other C-state.

4.2.3 Requesting Low-Power Idle States

The core C-state will be C1E if all actives cores have also resolved a core C1 state or higher.

The primary software interfaces for requesting low power idle states are through the MWAIT instruction with sub-state hints and the HLT instruction (for C1 and C1E). However, software may make C-state requests using the legacy method of I/O reads from the ACPI-defined processor clock control registers, referred to as P_LVLx. This method of requesting C-states provides legacy support for operating systems that initiate C-state transitions via I/O reads.

For legacy operating systems, P_LVLx I/O reads are converted within the processor to the equivalent MWAIT C-state request. Therefore, P_LVLx reads do not directly result in I/O reads to the system. The feature, known as I/O MWAIT redirection, must be enabled in the BIOS. To enable it, refer to the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3*.

Note: The P_LVLx I/O Monitor address needs to be set up before using the P_LVLx I/O read interface. Each P-LVLx is mapped to the supported MWAIT(Cx) instruction as follows.

Table 4-8. P_LVLx to MWAIT Conversion

P_LVLx	MWAIT(Cx)	Notes
P_LVL2	MWAIT(C3)	The P_LVL2 base address is defined in the PMG_IO_CAPTURE MSR, described in the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3.*
P_LVL3	MWAIT(C6)	C6. No sub-states allowed.
P_LVL4	MWAIT(C7)	C7. No sub-states allowed.

The BIOS can write to the C-state range field of the PMG_IO_CAPTURE MSR to restrict the range of I/O addresses that are trapped and emulate MWAIT like functionality. Any P_LVLx reads outside of this range does not cause an I/O redirection to MWAIT(Cx) like request. They fall through like a normal I/O instruction.

Note: When P_LVLx I/O instructions are used, MWAIT substates cannot be defined. The MWAIT substate is always zero if I/O MWAIT redirection is used. By default, P_LVLx I/O redirections enable the MWAIT 'break on EFLAGS.IF' feature which triggers a wakeup on an interrupt even if interrupts are masked by EFLAGS.IF.

4.2.4 Core C-states

The following are general rules for all core C-states, unless specified otherwise:

- A core C-State is determined by the lowest numerical thread state (for example, Thread 0 requests C1E while Thread 1 requests C3, resulting in a core C1E state). See Table 4-7.
- A core transitions to C0 state when:
  — an interrupt occurs.
  — there is an access to the monitored address if the state was entered via an MWAIT instruction.
- For core C1/C1E, and core C3, an interrupt directed toward a single thread wakes only that thread. However, since both threads are no longer at the same core C-state, the core resolves to C0.
- An interrupt only wakes the target thread for both C3 and C6 states. Any interrupt coming into the processor package may wake any core.

4.2.4.1 Core C0 State

The normal operating state of a core where code is being executed.

4.2.4.2 Core C1/C1E State

C1/C1E is a low power state entered when all threads within a core execute a HLT or MWAIT(C1/C1E) instruction.

A System Management Interrupt (SMI) handler returns execution to either Normal state or the C1/C1E state. See the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3* for more information.

While a core is in C1/C1E state, it processes bus snoops and snoops from other threads. For more information on C1E, see Section 4.2.5.2, "Package C1/C1E".

4.2.4.3 Core C3 State

Individual threads of a core can enter the C3 state by initiating a P_LVL2 I/O read to the P_BLK or an MWAIT(C3) instruction. A core in C3 state flushes the contents of its L1 instruction cache, L1 data cache, and L2 cache to the shared L3 cache, while maintaining its architectural state. All core clocks are stopped at this point. Because the core's caches are flushed, the processor does not wake any core that is in the C3 state when either a snoop is detected or when another core accesses cacheable memory.

4.2.4.4 Core C6 State

Individual threads of a core can enter the C6 state by initiating a P_LVL3 I/O read or an MWAIT(C6) instruction. Before entering core C6, the core will save its architectural state to a dedicated SRAM. Once complete, a core will have its voltage reduced to zero volts. In addition to flushing core caches core architecture state is saved to the uncore. Once the core state save is completed, core voltage is reduced to zero. During exit, the core is powered on and its architectural state is restored.

4.2.4.5 Core C7 State

Individual threads of a core can enter the C7 state by initiating a P_LVL4 I/O read to the P_BLK or by an MWAIT(C7) instruction. Core C7 and core C7 substate are the same as Core C6. The processor does not support LLC flush under any condition.

4.2.4.6 C-State Auto-Demotion

In general, deeper C-states such as C6 or C7 have long latencies and have higher energy entry/exit costs. The resulting performance and energy penalties become significant when the entry/exit frequency of a deeper C-state is high. In order to increase residency in deeper C-states, the processor supports C-state auto-demotion.

There are two C-State auto-demotion options:

- C6/C7 to C3
- C3/C6/C7 To C1

The decision to demote a core from C6/C7 to C3 or C3/C6/C7 to C1 is based on each core's immediate residency history. Upon each core C6/C7 request, the core C-state is demoted to C3 or C1 until a sufficient amount of residency has been established. At that point, a core is allowed to go into C3/C6 or C7. Each option can be run concurrently or individually.

This feature is disabled by default. BIOS must enable it in the PMG_CST_CONFIG_CONTROL register. The auto-demotion policy is also configured by this register. See the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3* for C-state configurations.

4.2.5 Package C-States

The processor supports C0, C1/C1E, C2, C3, and C6 power states. The following is a summary of the general rules for package C-state entry. These apply to all package C-states unless specified otherwise:

- A package C-state request is determined by the lowest numerical core C-state amongst all cores.
- A package C-state is automatically resolved by the processor depending on the core idle power states and the status of the platform components.
  - — Each core can be at a lower idle power state than the package if the platform does not grant the processor permission to enter a requested package C-state.
  - — The platform may allow additional power savings to be realized in the processor.
- For package C-states, the processor is not required to enter C0 before entering any other C-state.

The processor exits a package C-state when a break event is detected. Depending on the type of break event, the processor does the following:

- If a core break event is received, the target core is activated and the break event message is forwarded to the target core.
  - — If the break event is not masked, the target core enters the core C0 state and the processor enters package C0.
  - — If the break event is masked, the processor attempts to re-enter its previous package state.
- If the break event was due to a memory access or snoop request.
  - — But the platform did not request to keep the processor in a higher package C-state, the package returns to its previous C-state.
  - — And the platform requests a higher power C-state, the memory access or snoop request is serviced and the package remains in the higher power C-state.

The package C-states fall into two categories: independent and coordinated. C0/C1/C1E are independent, while C2/C3/C6 are coordinated.

Starting with the 2nd Generation Intel(r) Core(TM) Processor Family, package C-states are based on exit latency requirements which are accumulated from the PCIe* devices, PCH, and software sources. The level of power savings that can be achieved is a function of the exit latency requirement from the platform. As a result, there is no fixed relationship between the coordinated C-state of a package, and the power savings that will be obtained from the state. Coordinated package C-states offer a range of power savings which is a function of the guaranteed exit latency requirement from the platform.

There is also a concept of Execution Allowed (EA), when EA status is 0, the cores in a socket are in C3 or a deeper state, a socket initiates a request to enter a coordinated package C-state. The coordination is across all sockets and the PCH.

Table 4-9 shows an example of a dual-core processor package C-state resolution. Figure 4-3 summarizes package C-state transitions with package C2 as the interim between PC0 and PC1 prior to PC3 and PC6.

Table 4-9. Coordination of Core Power States at the Package Level

Package C-State		Core 1 C0	Core 1 C1	Core 1 C3	Core 1 C6
Core 0	C0	C0	C0	C0	C0
	C1	C0	C1[1]	C1[1]	C1[1]
	C3	C0	C1[1]	C3	C3
	C6	C0	C1[1]	C3	C6

Notes:

1. The package C-state will be C1E if all actives cores have resolved a core C1 state or higher.

Figure 4-3. Package C-State Entry and Exit

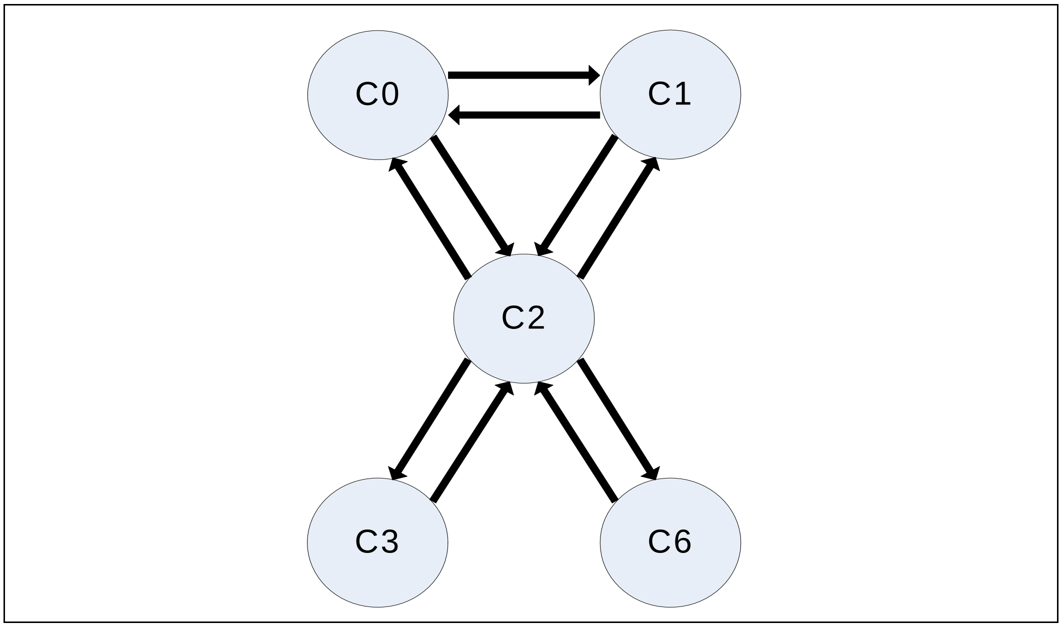

4.2.5.1 Package C0

The normal operating state for the processor. The processor remains in the normal state when at least one of its cores is in the C0 or C1 state or when the platform has not granted permission to the processor to go into a low power state. Individual cores may be in lower power idle states while the package is in C0.

4.2.5.2 Package C1/C1E

No additional power reduction actions are taken in the package C1 state. However, if the C1E substate is enabled, the processor automatically transitions to the lowest supported core clock frequency, followed by a reduction in voltage. Autonomous power reduction actions which are based on idle timers, can trigger depending on the activity in the system.

The package enters the C1 low power state when:

- At least one core is in the C1 state.
- The other cores are in a C1 or lower power state.

The package enters the C1E state when:

- All cores have directly requested C1E via MWAIT(C1) with a C1E sub-state hint.
- All cores are in a power state lower that C1/C1E but the package low power state is limited to C1/C1E via the PMG_CST_CONFIG_CONTROL MSR.
- All cores have requested C1 using HLT or MWAIT(C1) and C1E auto-promotion is enabled in POWER_CTL.

No notification to the system occurs upon entry to C1/C1E.

4.2.5.3 Package C2 State

Package C2 state is an intermediate state which represents the point at which the system level coordination is in progress. The package cannot reach this state unless all cores are in at least C3.

The package will remain in C2 when:

- it is awaiting for a coordinated response
- the coordinated exit latency requirements are too stringent for the package to take any power saving actions

If the exit latency requirements are high enough the package will transition to C3 or C6 depending on the state of the cores.

4.2.5.4 Package C3 State

A processor enters the package C3 low power state when:

- At least one core is in the C3 state.
- The other cores are in a C3 or lower power state, and the processor has been granted permission by the platform.
- L3 shared cache retains context and becomes inaccessible in this state.
- Additional power savings actions, as allowed by the exit latency requirements, include putting Intel QPI and PCIe* links in L1, the uncore is not available, further voltage reduction can be taken.

In package C3, the ring will be off and as a result no accesses to the LLC are possible. The content of the LLC is preserved.

4.2.5.5 Package C6 State

A processor enters the package C6 low power state when:

- At least one core is in the C6 state.
- The other cores are in a C6 or lower power state, and the processor has been granted permission by the platform.
- L3 shared cache retains context and becomes inaccessible in this state.

- Additional power savings actions, as allowed by the exit latency requirements, include putting Intel QPI and PCIe* links in L1, the uncore is not available, further voltage reduction can be taken.

In package C6 state, all cores have saved their architectural state and have had their core voltages reduced to zero volts. The LLC retains context, but no accesses can be made to the LLC in this state, the cores must break out to the internal state package C2 for snoops to occur.

4.2.6 Package C-State Power Specifications

The table below lists the processor package C-state power specifications for various processor SKUs.

Table 4-10. Package C-State Power Specifications

TDP SKUs	C1E (W)	C3 (W)	C6 (W)
8-Core / 6-Core			
95W (8-core)	46	20 (35W @ 2.0/2.1 GHz)	15
95W (6-core)	50	20 (35W @ 1.9 GHz)	15 (21W @ 1.9 GHz)
70W (8-core) & LV70W-8C	43	20	13
60W (6-core) & LV60W-6C	34	16	13
50W (8-core)	43	20	13
4-Core / 2-Core			
80W (4-core)	37	27	16
80W (2-core, 1 Socket)	42	24 (30W @ 2.6 GHz)	14 (21W @ 2.6 GHz)
50W (4-core) & LV50W-4C	34	16	13
LV40W-24C (2-core, 1 Socket)	34	16	13

Notes:
1. Package C1E power specified at Tcase = 60°C.
2. Package C3/C6 power specified at Tcase = 50°C.

4.3 System Memory Power Management

The DDR3 power states can be summarized as the following:

- Normal operation (highest power consumption).
- CKE Power-Down: Opportunistic, per rank control after idle time. There may be different levels.
  - — Active Power-Down.
  - — Precharge Power-Down with Fast Exit.
  - — Precharge power Down with Slow Exit.
- Self Refresh: In this mode no transaction is executed. The DDR consumes the minimum possible power.

4.3.1 CKE Power-Down

The CKE input land is used to enter and exit different power-down modes. The memory controller has a configurable activity timeout for each rank. Whenever no reads are present to a given rank for the configured interval, the memory controller will transition the rank to power-down mode.

The memory controller transitions the DRAM to power-down by de-asserting CKE and driving a NOP command. The memory controller will tri-state all DDR interface lands except CKE (de-asserted) and ODT while in power-down. The memory controller will transition the DRAM out of power-down state by synchronously asserting CKE and driving a NOP command.

When CKE is off the internal DDR clock is disabled and the DDR power is significantly reduced.

The DDR defines three levels of power-down:

- Active power-down.
- Precharge power-down fast exit.
- Precharge power-down slow exit.

4.3.2 Self Refresh

The Power Control Unit (PCU) may request the memory controller to place the DRAMs in self refresh state. Self refresh per channel is supported. The BIOS can put the channel in self-refresh if software remaps memory to use a subset of all channels. Also processor channels can enter self refresh autonomously without PCU instruction when the package is in a package C0 state.

4.3.2.1 Self Refresh Entry

Self refresh entrance can be either disabled or triggered by an idle counter. The idle counter always clears with any access to the memory controller and remains clear as long as the memory controller is not drained. As soon as the memory controller is drained, the counter starts counting, and when it reaches the idle-count, the memory controller will place the DRAMs in self refresh state.

Power may be removed from the memory controller core at this point. But V_{CCD} supply (1.5 V or 1.35 V) to the DDR IO must be maintained.

4.3.2.2 Self Refresh Exit

Self refresh exit can be either a message from an external unit or as reaction for an incoming transaction.

4.3.2.3 DLL and PLL Shutdown

Self refresh, according to configuration, may be a trigger for master DLL shut-down and PLL shut-down. The master DLL shut-down is issued by the memory controller after the DRAMs have entered self refresh.

The PLL shut-down and wake-up is issued by the PCU. The memory controller gets a signal from PLL indicating that the memory controller can start working again.

4.3.3 DRAM I/O Power Management

Unused signals are tristated to save power. This includes all signals associated with an unused memory channel.

The I/O buffer for an unused signal should be tristated (output driver disabled), the input receiver (differential sense-amp) should be disabled. The input path must be gated to prevent spurious results due to noise on the unused signals (typically handled automatically when input receiver is disabled).

4.4 DMI2/PCI Express* Power Management

Active State Power Management (ASPM) support using L1 state.

§

5 Thermal Management Specifications

5.1 Package Thermal Specifications

The processor requires a thermal solution to maintain temperatures within operating limits. Any attempt to operate the processor outside these limits may result in permanent damage to the processor and potentially other components within the system, see Section 7.7.1, "Storage Conditions Specifications". Maintaining the proper thermal environment is key to reliable, long-term system operation.

A complete solution includes both component and system level thermal management features. Component level thermal solutions can include active or passive heatsinks attached to the processor integrated heat spreader (IHS). Typical system level thermal solutions may consist of system fans combined with ducting and venting.

This section provides data necessary for developing a complete thermal solution. For more information on designing a component level thermal solution, refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

5.1.1 Thermal Specifications

To allow optimal operation and long-term reliability of Intel processor-based systems, the processor must remain within the minimum and maximum case temperature (T_{CASE}) specifications as defined by the applicable thermal profile. Thermal solutions not designed to provide sufficient thermal capability may affect the long-term reliability of the processor and system. For more details on thermal solution design, please refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

The processors implement a methodology for managing processor temperatures which is intended to support acoustic noise reduction through fan speed control and to assure processor reliability. Selection of the appropriate fan speed is based on the relative temperature data reported by the processor's Platform Environment Control Interface (PECI) as described in Section 2.5, "Platform Environment Control Interface (PECI)".

If the DTS value is less than $T_{CONTROL}$, then the case temperature is permitted to exceed the Thermal Profile, but the DTS value must remain at or below TCONTROL.

For T_{CASE} implementations, if DTS is greater than TCONTROL, then the case temperature must meet the T_{CASE} based Thermal Profiles.

For DTS implementations:

- T_{CASE} thermal profile can be ignored during processor run time.
- If DTS is greater than Tcontrol then follow DTS thermal profile specifications for fan speed optimization.

The temperature reported over PECI is always a negative value and represents a delta below the onset of thermal control circuit (TCC) activation, as indicated by PROCHOT_N (see Section 7, "Electrical Specifications"). Systems that implement fan speed control must be designed to use this data. Systems that do not alter the fan speed need to guarantee the case temperature meets the thermal profile specifications.

Some processor SKUs support two thermal profiles; refer to Table 5-1 for a summary of the planned SKUs and their supported thermal profiles. Thermal profiles ensure adherence to Intel reliability requirements. With single thermal profile, it is expected that the Thermal Control Circuit (TCC) would be activated for very brief periods of time when running the most power intensive applications. Additionally, utilization of a thermal solution that does not meet this Thermal Profile will violate the thermal specifications and may result in permanent damage to the processor. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for details on system thermal solution design, thermal profiles and environmental considerations.

For Embedded Servers, Communications and storage markets Intel has plan SKU's that support Thermal Profiles with nominal and short-term conditions designed to meet NEBS level 3 compliance. For these SKU's operation at either the nominal or short-term thermal profiles should result in virtually no TCC activation. Thermal Profiles for these SKU's are found in Section 5.1.3, "Embedded Server Thermal Profiles".

Intel recommends that complete thermal solution designs target the Thermal Design Power (TDP). The Adaptive Thermal Monitor feature is intended to help protect the processor in the event that an application exceeds the TDP recommendation for a sustained time period. To ensure maximum flexibility for future requirements, systems should be designed to the Flexible Motherboard (FMB) guidelines, even if a processor with lower power dissipation is currently planned. **The Adaptive Thermal Monitor feature must be enabled for the processor to remain within its specifications.**

5.1.2 T_{CASE} and DTS Based Thermal Specifications

To simplify compliance to thermal specifications at processor run time, the processor has added a Digital Thermal Sensor (DTS) based thermal specification. Digital Thermal Sensor reports a relative die temperature as an offset from TCC activation temperature. T_{CASE} thermal based specifications are used for heat sink sizing and DTS based specs are used for acoustic and fan speed optimizations. For the processor family, firmware (for example, BMC or other platform management devices) will have DTS based specifications for all SKUs programmed by the customer. 8-core and 6-core SKUs may share T_{CASE} thermal profiles but they will have separate T_{DTS} based thermal profiles. See Table 5-1 for the T_{CASE} and DTS SKU summary.

The processor fan speed control is managed by comparing DTS thermal readings via PECI against the processor-specific fan speed control reference point, or Tcontrol. Both Tcontrol and DTS thermal readings are accessible via the processor PECI client. At a one time readout only, the Fan Speed Control firmware will read the following:

- TEMPERATURE_TARGET MSR
- Tcontrol via PECI - RdPkgConfig()
- TDP via PECI - RdPkgConfig()
- Core Count - RdPCIConfigLocal()

DTS PECI commands will also support DTS temperature data readings. Please see Section 2.5.7, "DTS Temperature Data" for PECI command details.

Also, refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/ Mechanical Design Guide* for details on DTS based thermal solution design considerations.

Table 5-1. SKU Table Summary

TDP SKUs	Tcase	DTS
8-Core / 6-Core		
95W (8-Core)	Figure 5-1	Figure 5-2
95W (6-Core)	Figure 5-1	Figure 5-3
70W (8-Core)	Figure 5-8	Figure 5-9
60W (6-Core)	Figure 5-10	Figure 5-11
50W (8-Core)	Figure 5-12	Figure 5-13
4-Core / 2-Core		
80W (4-Core)	Figure 5-4	Figure 5-5
80W (2-Core, 1 Socket only)	Figure 5-6	Figure 5-7
50W (4-Core)	Figure 5-14	Figure 5-15

5.1.2.1 95W Thermal Specifications

Table 5-2. Server SKU Platform (95W) Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	95	5	See Figure 5-1 and Table 5-3	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-1. Server SKU Platform (8-Core / 6-Core 95W) T_{CASE} Thermal Profile

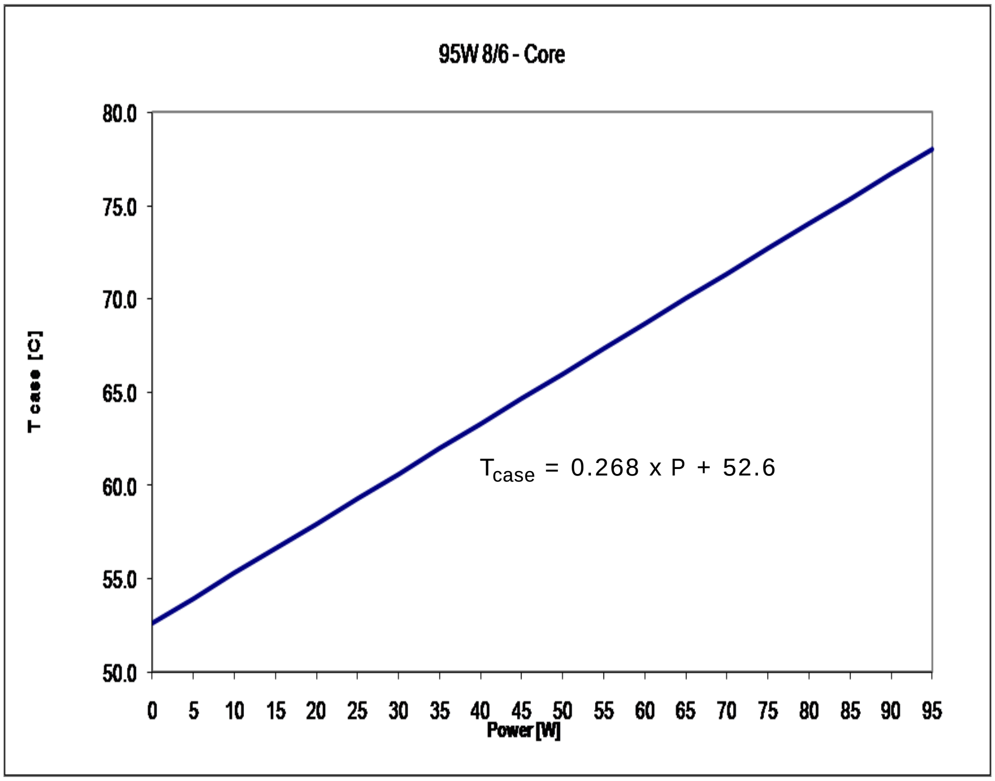

Notes:

1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-3 for discrete points that constitute the thermal profiles.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The 8-core and 6-core 95W SKU have the same T_{CASE} based thermal profile. The difference between the 8-core and the 6-core T_{CASE} is very insignificant.
4. Please refer to Table 5-3 and Figure 5-2 and Figure 5-3 for the DTS based thermal profile for the 8-core and 6-core 95W SKU.
5. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-2. Server SKU Platform (8-Core 95W) DTS Based Thermal Profile

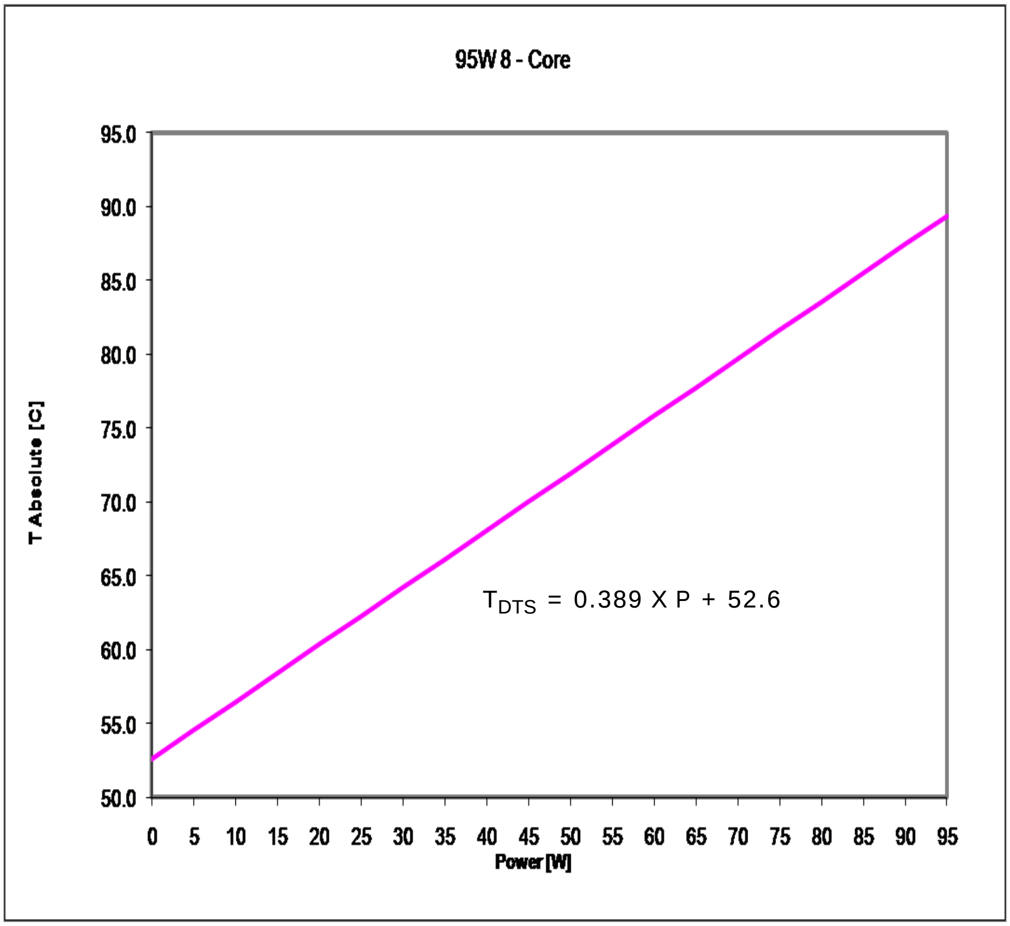

Figure 5-3. Server SKU Platform (6-Core 95W) DTS Based Thermal Profile

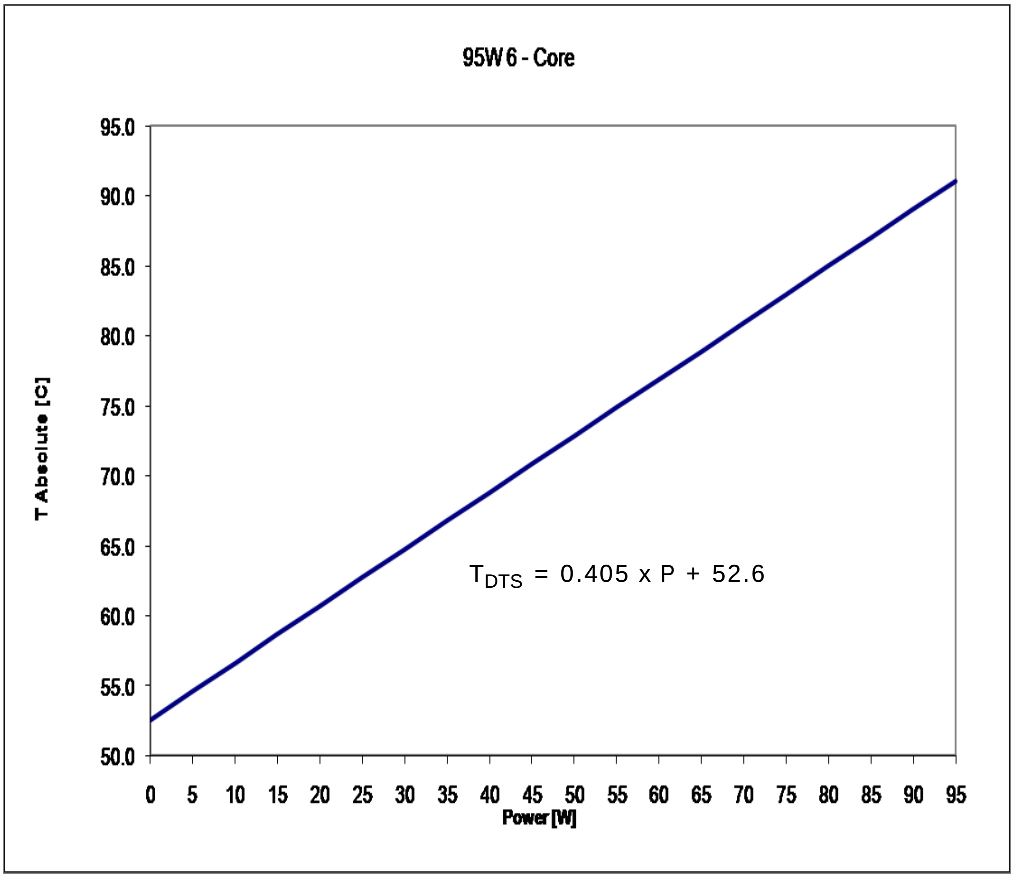

Table 5-3. Server SKU Platform (8-Core / 6-Core 95W) T_{CASE} and DTS Thermal Profiles (Sheet 1 of 2)

Power (W)	Maximum T_{CASE} (°C) 8 / 6 Core	Maximum DTS (°C) 8 Core	Maximum DTS (°C) 6 Core
0	52.6	52.6	52.6
5	53.9	54.5	54.6
10	55.3	56.5	56.7
15	56.6	58.4	58.7
20	58.0	60.3	60.7
25	59.3	62.3	62.7
30	60.6	64.2	64.8
35	62.0	66.1	66.8
40	63.3	68.1	68.8
45	64.7	70.0	70.8
50	66.0	72.0	72.9
55	67.3	73.9	74.9
60	68.7	75.8	76.9

Table 5-3. Server SKU Platform (8-Core / 6-Core 95W) T_{CASE} and DTS Thermal Profiles (Sheet 2 of 2)

Power (W)	Maximum T_{CASE} (°C) 8 / 6 Core	Maximum DTS (°C) 8 Core	Maximum DTS (°C) 6 Core
65	70.0	77.8	78.9
70	71.4	79.7	81.0
75	72.7	81.6	83.0
80	74.0	83.6	85.0
85	75.4	85.5	87.0
90	76.7	87.4	89.1
95	78.0	90.0	92.0

5.1.2.2 80W Thermal Specifications

Table 5-4. Server SKU Platform (4-Core 80W) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	80	5	See Figure 5-4; Table 5-5	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-4. Server SKU Platform (4-Core 80W) T_{CASE} Thermal Profile

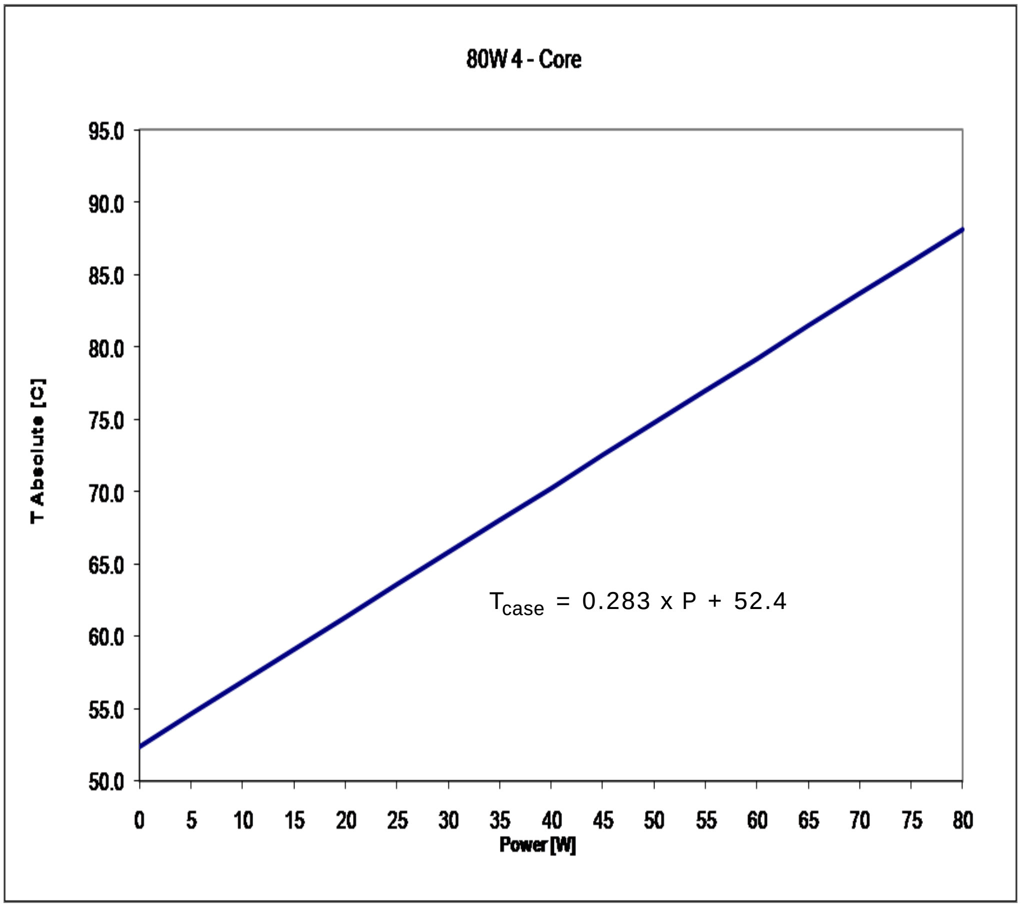

Notes:

1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-5 for discrete points that constitute thermal profiles.
2. Implementation of the this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss.
3. The 4-core and 2-core 80W SKU have the same T_{CASE} based thermal profile. The difference between the 4-core and the 2-core T_{CASE} is very insignificant.
4. Please refer to Table 5-5 and Figure 5-5 and for the DTS based thermal profiles for the 4-core 80W SKU.
5. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-5. Server SKU Platform (4-Core 80W) DTS Based Thermal Profile

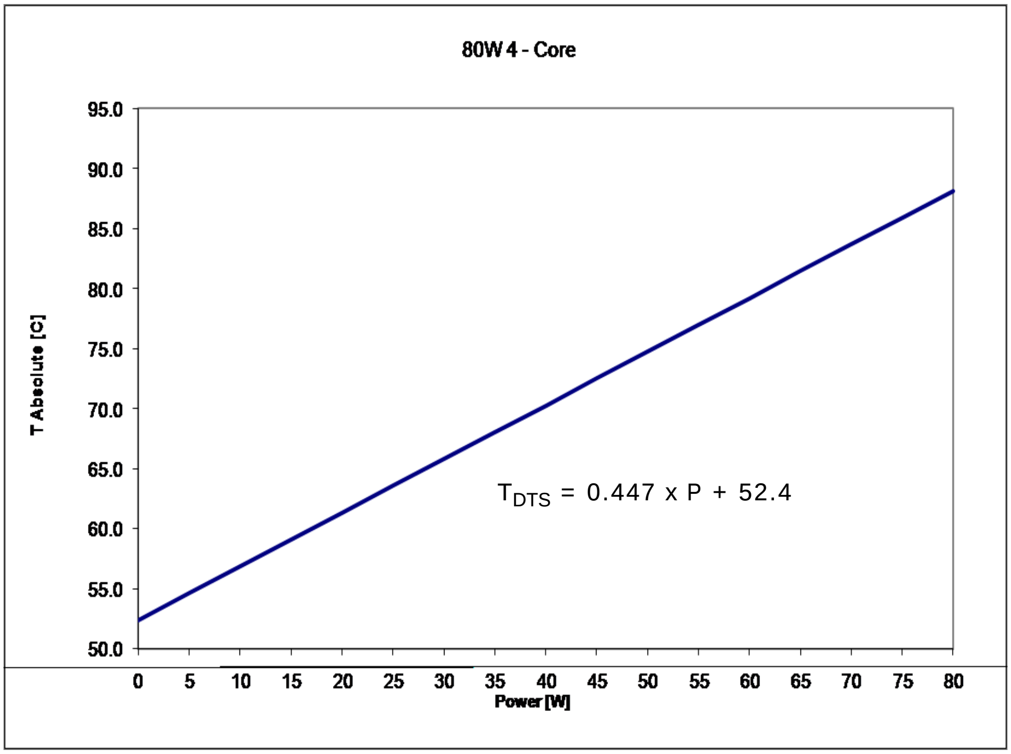

Table 5-5. Server SKU Platform (4-Core 80W) T_{CASE} and DTS Thermal Profiles (Sheet 1 of 2)

Power (W)	Maximum T_{CASE} (°C) 4 Core	Maximum DTS (°C) 4 Core
0	52.4	52.4
5	53.8	54.6
10	55.2	56.9
15	56.6	59.1
20	58.1	61.3
25	59.5	63.6
30	60.9	65.8
35	62.3	68.0
40	63.7	70.3
45	65.1	72.5
50	66.6	74.8
55	68.0	77.0
60	69.4	79.2
65	70.8	81.5
70	72.2	83.7
75	73.6	85.9

Table 5-5. Server SKU Platform (4-Core 80W) T_{CASE} and DTS Thermal Profiles (Sheet 2 of 2)

Power (W)	Maximum T_{CASE} (°C) 4 Core	Maximum DTS (°C) 4 Core
80	75.0	89.0

Table 5-6. Server SKU Platform (2-Core 80W) One Socket T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	80	5	See Figure 5-6; Table 5-7	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product familymay be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-6. Server SKU Platform (2-Core 80W) One Socket T_{CASE} Thermal Profile

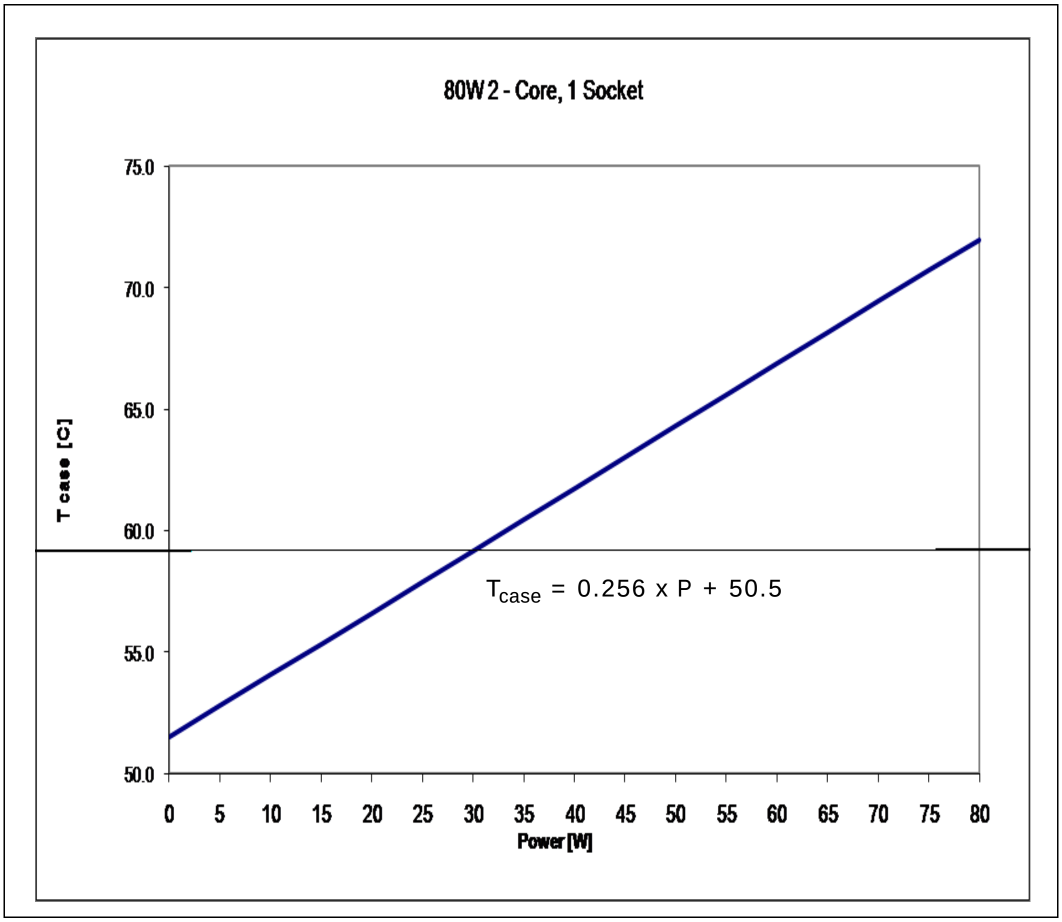

Notes:

1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-7 for discrete points that constitute thermal profiles.
2. Implementation of the this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss.
3. Please refer to Table 5-7 and Figure 5-7 for the DTS based thermal profile.
4. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-7. Server SKU Platform (2-Core 80W) One Socket DTS Based Thermal Profile.

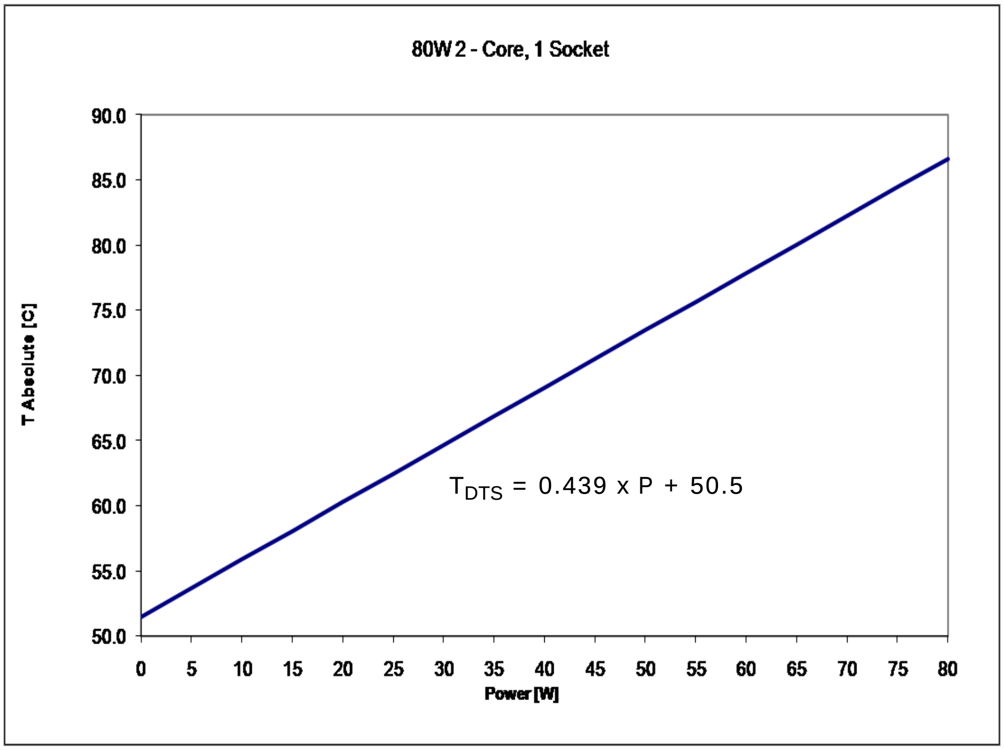

Table 5-7. Server SKU Platform (2-Core 80W) 1-Socket T_{CASE} and DTS Based Thermal Profiles (Sheet 1 of 2)

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
0	50.5	50.5
5	51.8	52.7
10	53.1	54.9
15	54.3	57.1
20	55.6	59.3
25	56.9	61.5
30	58.2	63.7
35	59.5	65.9
40	60.7	68.1
45	62.0	70.3
50	63.3	72.5
55	64.6	74.6

Table 5-7. Server SKU Platform (2-Core 80W) 1-Socket T_{CASE} and DTS Based Thermal Profiles (Sheet 2 of 2)

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
60	65.9	76.8
65	67.1	79.0
70	68.4	81.2
75	69.7	83.4
80	71.0	86.0

5.1.2.3 70W Thermal Specifications

Table 5-8. Server SKU Platform (8-Core 70W) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	70	5	See Figure 5-8; Table 5-9	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-8. Server SKU Platform (8-Core 70W) T_{CASE} Thermal Profile

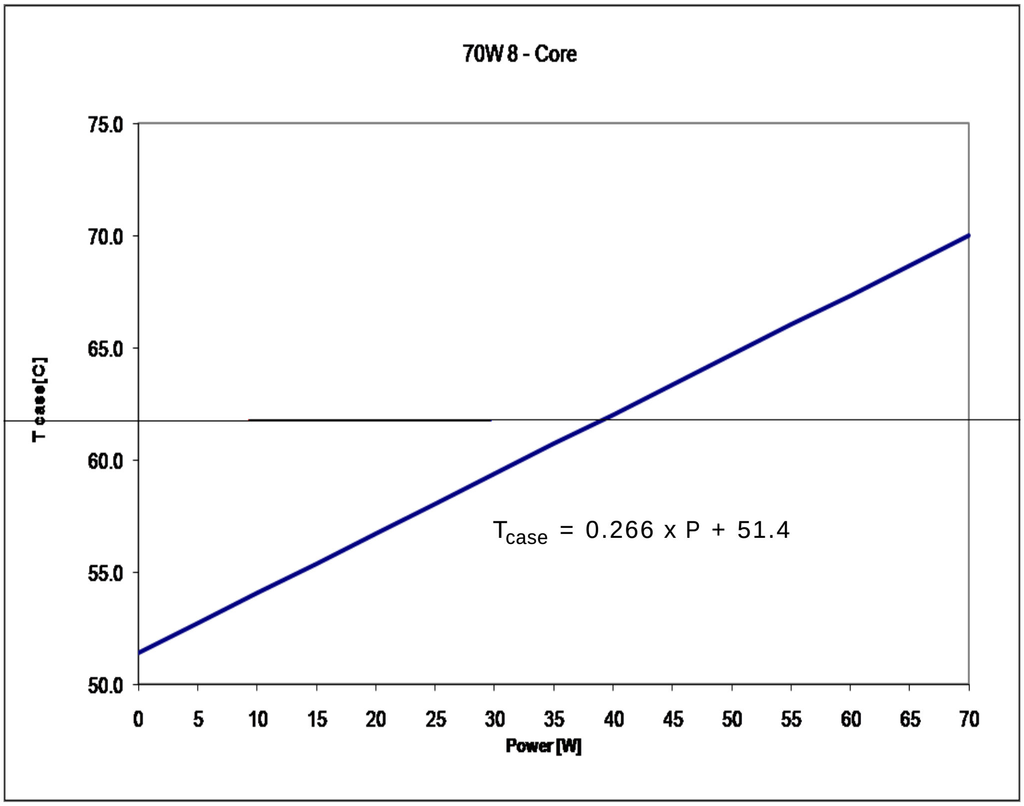

Notes:

1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-9 for discrete points that constitute thermal profiles.
2. Implementation of the this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss.
3. Please refer to Table 5-9 and Figure 5-9 for the DTS based thermal profile.
4. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-9. Server SKU Platform (8-Core 70W) DTS Based Thermal Profile

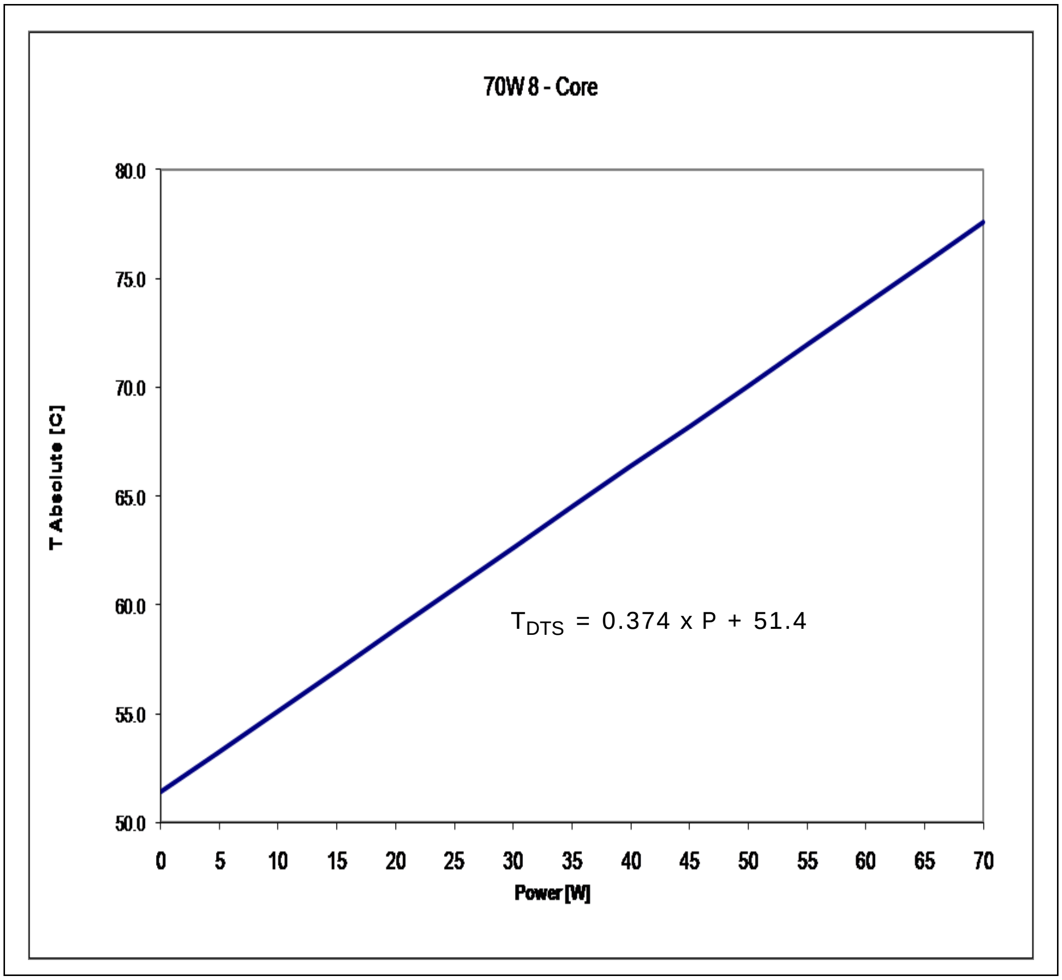

Table 5-9. Server SKU Platform (8-Core 70W) T_{CASE} and DTS Bases Thermal Profile (Sheet 1 of 2)

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
0	51.4	51.4
5	52.7	53.3
10	54.1	55.1
15	55.4	57.0
20	56.7	58.9
25	58.1	60.8
30	59.4	62.6
35	60.7	64.5
40	62.0	66.4
45	63.4	68.2
50	64.7	70.1
55	66.0	72.0

Table 5-9. Server SKU Platform (8-Core 70W) T_{CASE} and DTS Bases Thermal Profile (Sheet 2 of 2)

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
60	67.4	73.8
65	68.7	75.7
70	70.0	78.0

5.1.2.4 60W Thermal Specification

Table 5-10. Server SKU Platform (6-Core 60W) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	60	5	See Figure 5-10; Table 5-11	1, 2, 3, 4, 5

Note:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-10. Server SKU Platform (6-Core 60W) T_{CASE} Thermal Profile

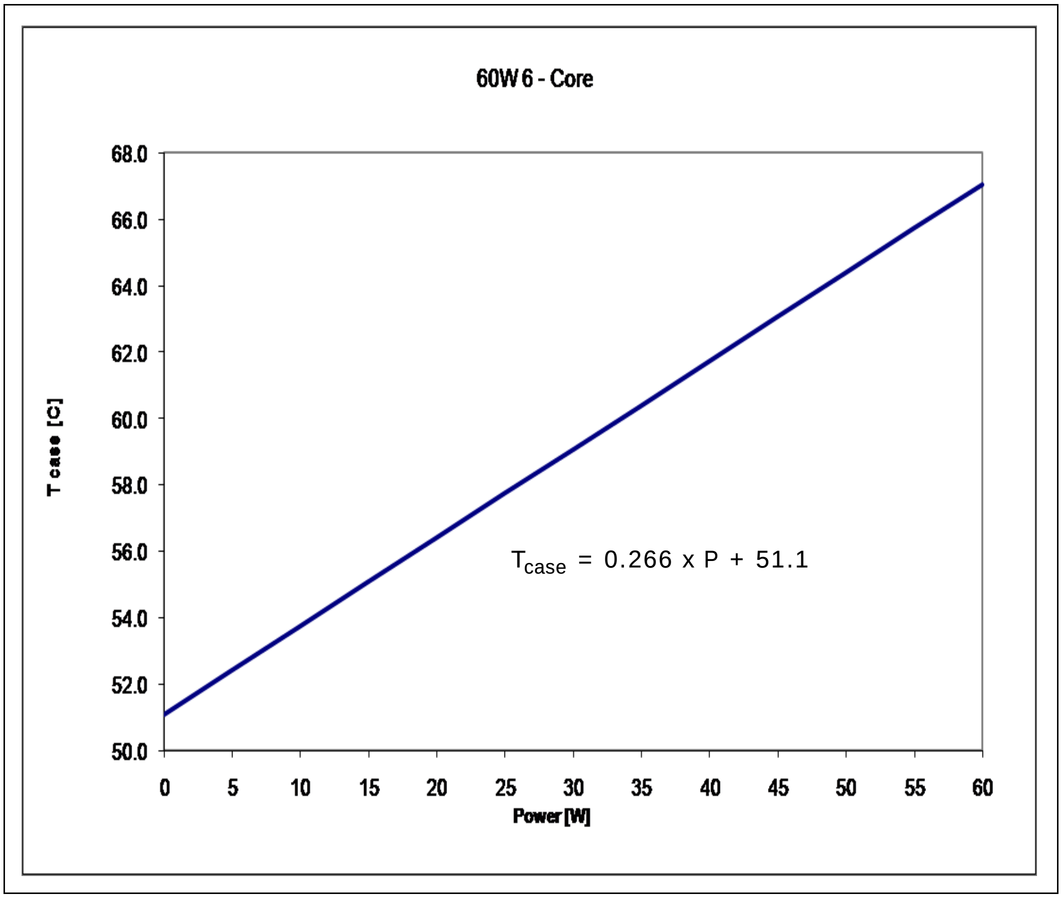

Note:

1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-13 for discrete points that constitute thermal profiles.
2. Implementation of the this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss.
3. Please refer to Table 5-11 and Figure 5-11 for the DTS based thermal profile.
4. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-11. Server SKU Platform (6-Core 60W) DTS Based Thermal Profile

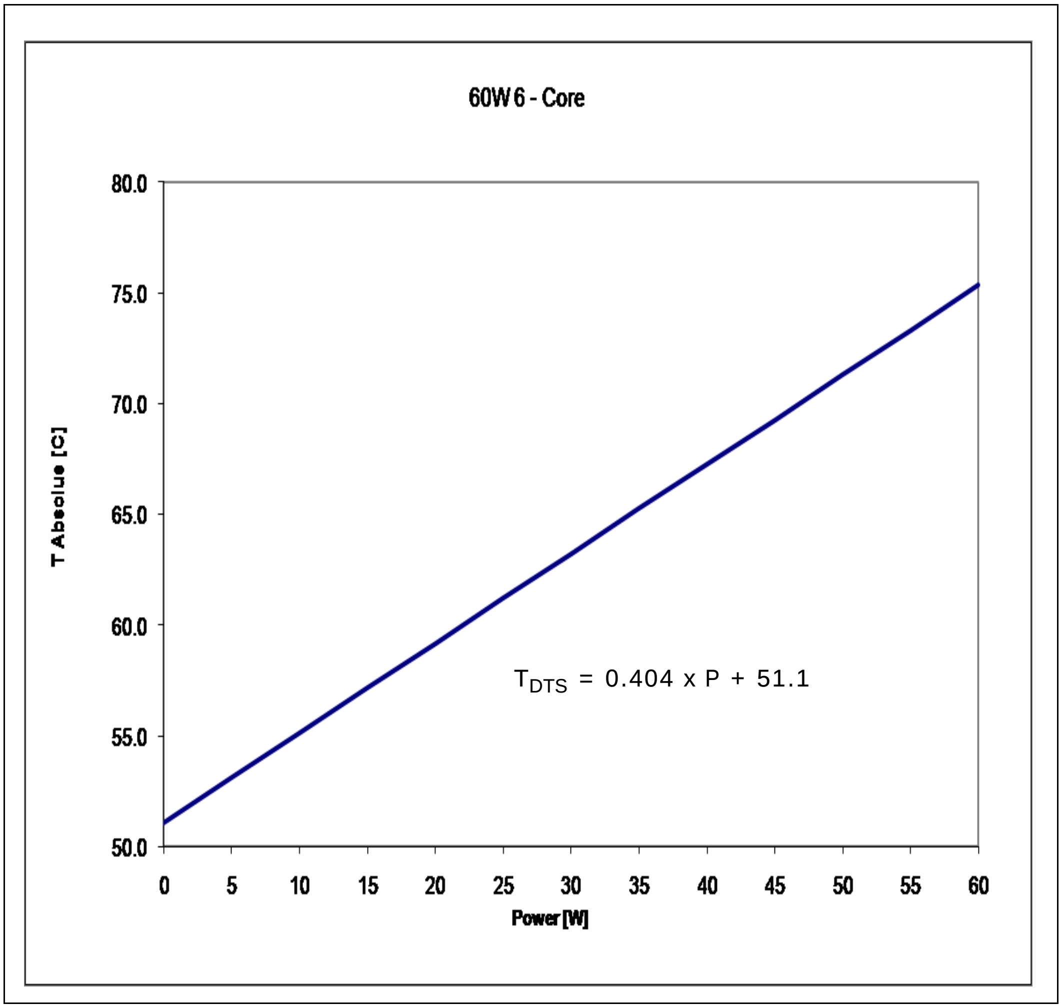

Table 5-11. Server SKU Platform (6-Core 60W) T_{CASE} and DTS Based Thermal Profiles (Sheet 1 of 2)

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
0	51.1	51.1
5	52.4	53.1
10	53.8	55.1
15	55.1	57.2
20	56.4	59.2
25	57.8	61.2
30	59.1	63.2
35	60.4	65.2
40	61.7	67.3
45	63.1	69.3
50	64.4	71.3
55	65.7	73.3

Table 5-11. Server SKU Platform (6-Core 60W) T_{CASE} and DTS Based Thermal Profiles (Sheet 2 of 2)

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
60	67.0	76.0

5.1.2.5 50W 8-Core Thermal Specification

Table 5-12. Server SKU Platform (8-Core 50W) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	50	5	See Figure 5-12; Table 5-13	1, 2, 3, 4, 5

Note:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product familymay be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-12. Server SKU Platform (8-Core 50W) T_{CASE} Thermal Profile.

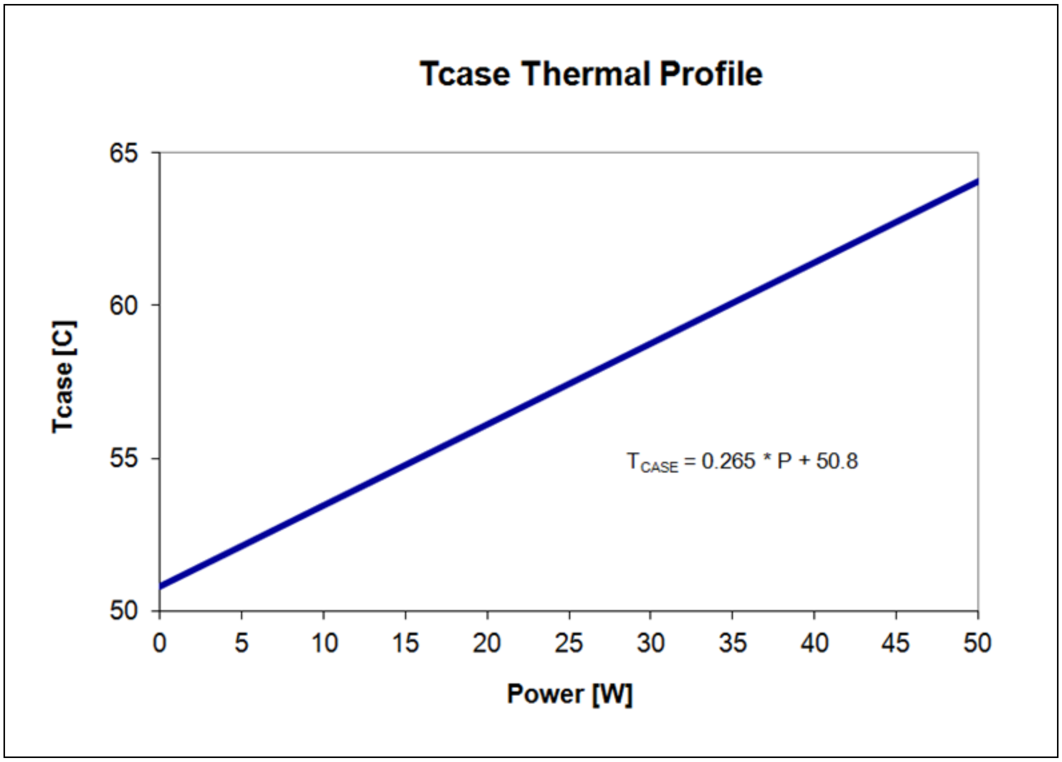

Note:
1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-13 for discrete points that constitute thermal profiles.
2. Implementation of the this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss.
3. Please refer to Table 5-13 and Figure 5-13 for the DTS based thermal profile.

4. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-13. Server SKU Platform (8-Core 50W) DTS Based Thermal Profile.

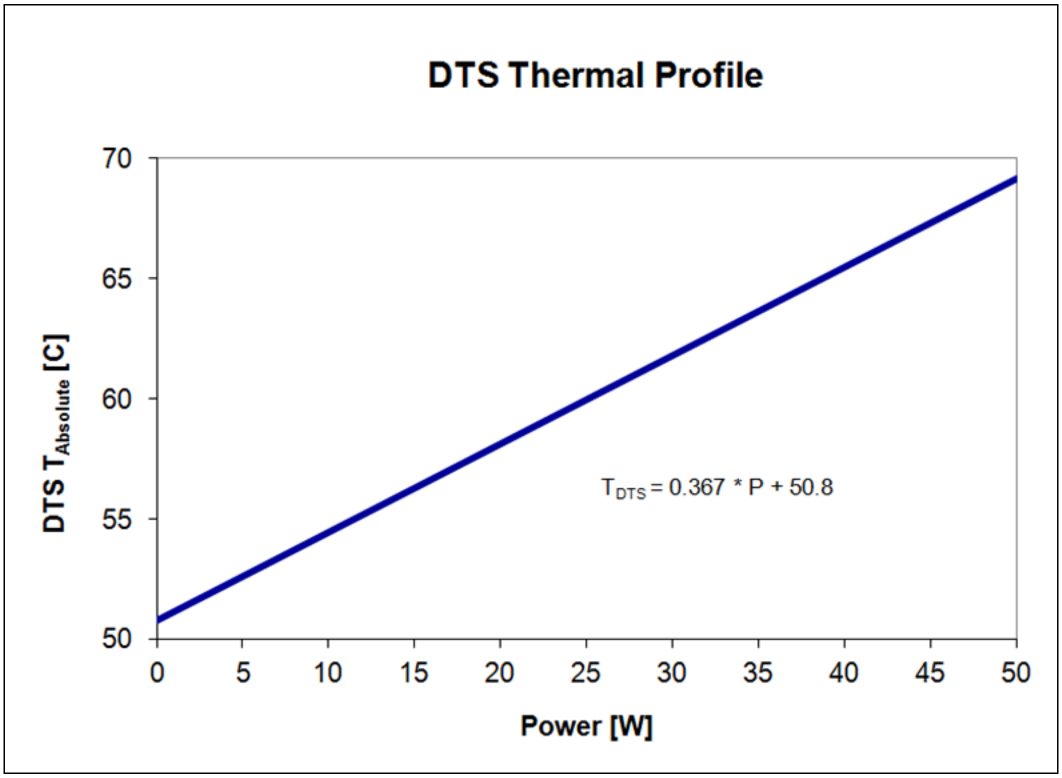

Table 5-13. Server SKU Platform (8-Core 50W) T_{CASE} and DTS Based Thermal Profiles

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
0	51	51
5	52	53
10	54	55
15	55	57
20	56	59
25	58	61
30	59	62
35	60	64
40	62	66
45	63	67
50	64	69

5.1.2.6 50W 4-Core Thermal Specification

Table 5-14. Server SKU Platform (4-Core 50W) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	50	5	See Figure 5-14; Table 5-15	1, 2, 3, 4, 5

Note:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product familymay be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-14. Server SKU Platform (4-Core 50W) T_{CASE} Thermal Profile.

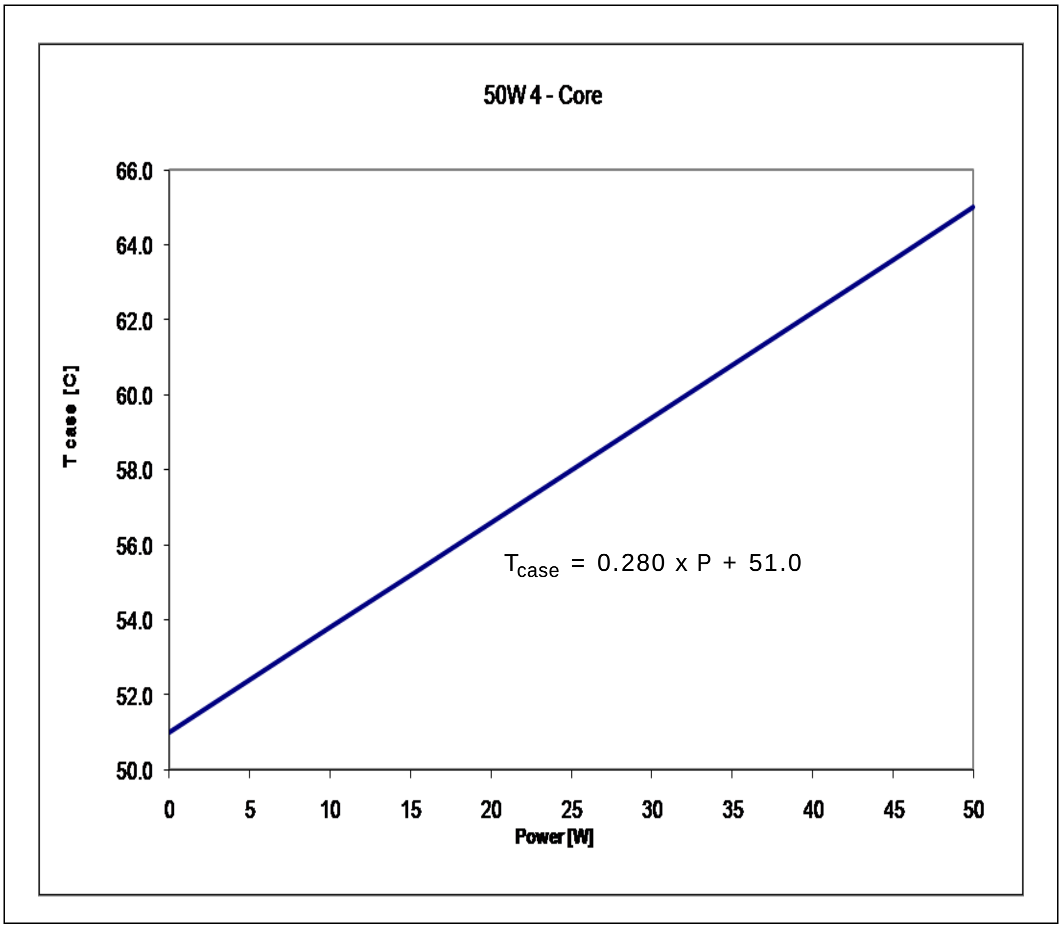

Note:
1. This Thermal profile is representative of a volumetrically unconstrained platform. Please refer to Table 5-15 for discrete points that constitute thermal profiles.
2. Implementation of the this Thermal Profile should result in virtually no TCC activation. Furthermore, utilization of thermal solutions that do not meet this Thermal Profile will result in increased probability of TCC activation and may incur measurable performance loss.
3. Please refer to Table 5-15 and Figure 5-15 for the DTS based thermal profile.

4. Some processor units may be tested to lower TDP and the IA32_TEMPERATURE_TARGET MSR will be aligned to that lower TDP.

Figure 5-15. Server SKU Platform (4-Core 50W) DTS Based Thermal Profile.

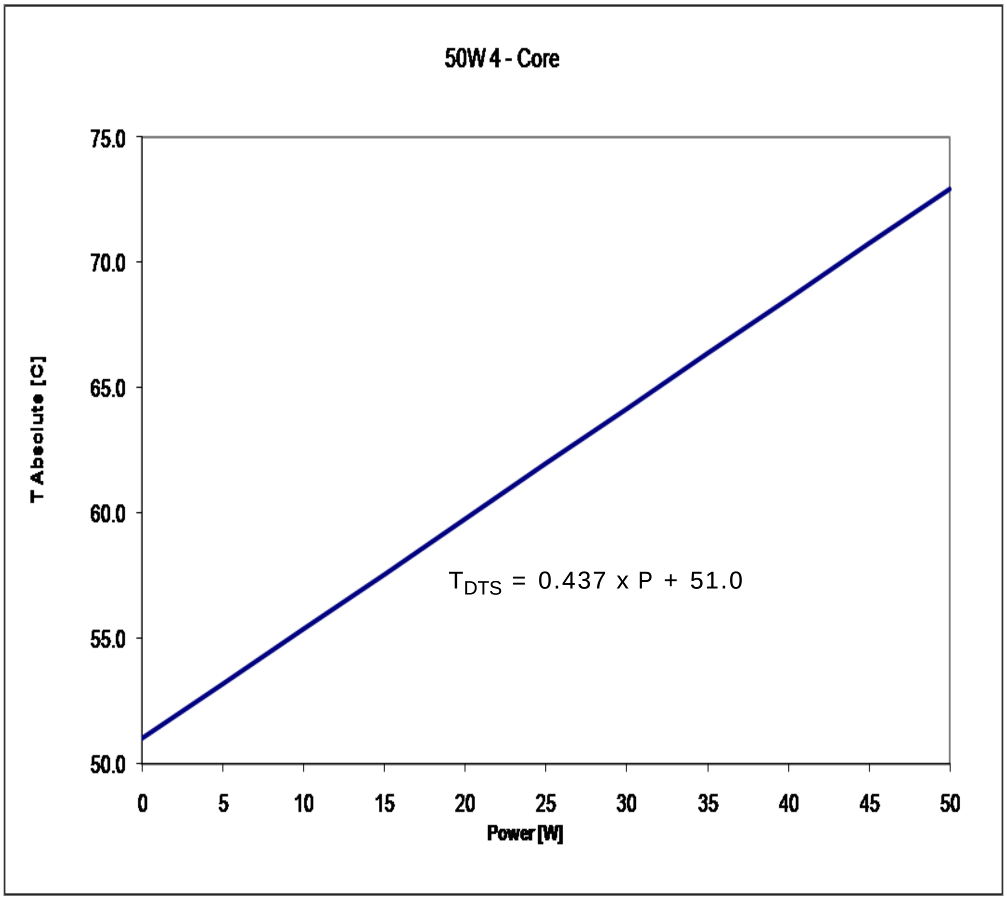

Table 5-15. Server SKU Platform (4-Core 50W) T_{CASE} and DTS Based Thermal Profiles

Power (W)	Maximum T_{CASE} (°C)	Maximum DTS (°C)
0	51.0	51.0
5	52.4	53.2
10	53.8	55.4
15	55.2	57.6
20	56.6	59.8
25	58.0	62.0
30	59.4	64.2
35	60.8	66.4
40	62.2	68.6
45	63.6	70.8
50	65.0	73.0

5.1.3 Embedded Server Thermal Profiles

Embedded server SKU's target operation at higher case temperatures and/or NEBS thermal profiles for embedded communications server and storage form factors. The thermal profiles in this section pertain only to those specific SKU's. Network Equipment Building System (NEBS) is the most common set of environmental design guidelines applied to telecommunications equipment in the United States.

Digital Thermal Sensor (DTS) based thermal profiles are also provided for each Embedded server SKU. The thermal solution is expected to be developed in accordance with the Tcase thermal profile. Operational compliance monitoring of thermal specifications and fan speed modulation may be done via the DTS based thermal profile. The slope of a DTS profile assumes full fan speed which is not required over much of the power range. At most power levels on embedded SKU's, temperatures of the nominal profile are less than Tcontrol as indicated by the blue shaded region in each DTS profile graph. As a further simplification, operation at DTS temperatures up to Tcontrol is permitted at all power levels. Compliance to the DTS profile is required for any temperatures exceeding Tcontrol.

Table 5-16. Embedded Server Elevated Tcase SKU Summary Table

TDP SKUs	Tcase Spec	DTS Based Spec
LV70W-8C (8-Core)	Figure 5-16	Figure 5-17
LV60W-6C (6-Core)	Figure 5-18	Figure 5-19
LV50W-4C (4-Core)	Figure 5-20	Figure 5-21
LV40W-24C (2-Core 1S)	Figure 5-22	Figure 5-23

5.1.3.1 Embedded LV70W-8C Thermal Specifications

Table 5-17. Embedded Server SKU (LV70W-8C) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	70	5	See Figure 5-16; Table 5-18	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-16. Embedded Server SKU (LV70W-8C) T_{CASE} Thermal Profile

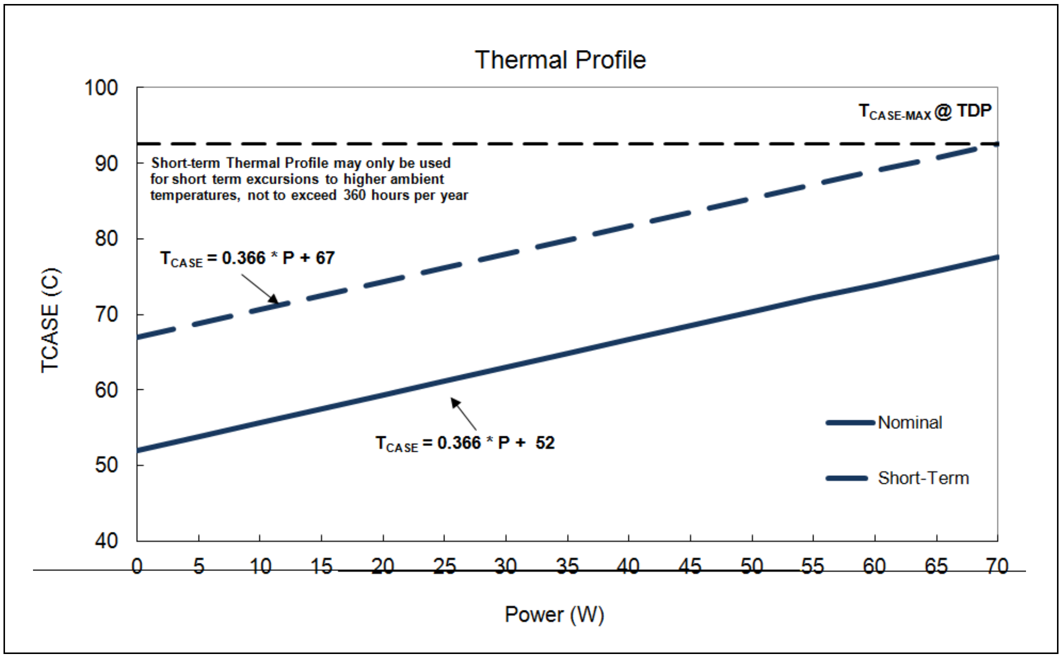

Notes:

1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-18 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Figure 5-17. Embedded Server SKU (LV70W-8C) DTS Thermal Profile

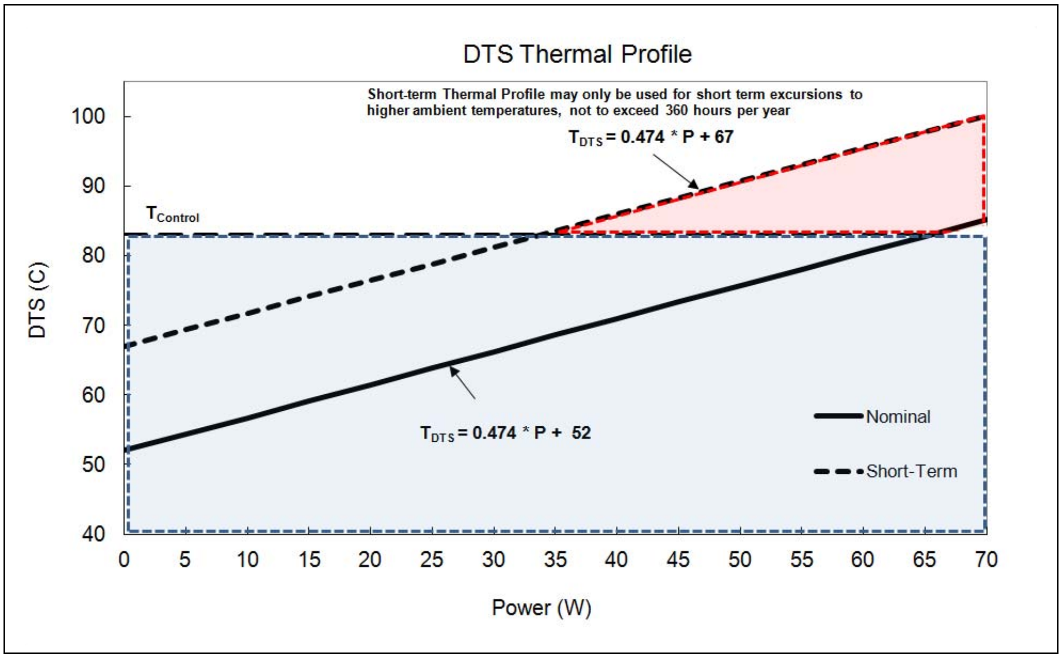

Notes:
1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-18 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance. As indicated by the blue shaded region, operation at DTS temperatures up to Tcontrol is permitted at all power levels.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Table 5-18. Embedded Server SKU (LV70W-8C) Thermal Profiles (Sheet 1 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
0	52	67	52	67
5	54	69	54	69
10	56	71	57	72
15	57	72	59	72
20	59	74	61	74
25	61	76	64	76
30	63	78	66	78
35	65	80	69	80
40	67	82	71	82
45	68	83	73	83
50	70	85	76	85
55	72	87	78	87

Table 5-18. Embedded Server SKU (LV70W-8C) Thermal Profiles (Sheet 2 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
60	74	89	80	95
65	76	91	83	98
70	78	93	85	100

5.1.3.2 Embedded LV60W-6C Thermal Specifications

Table 5-19. Embedded Server SKU (LV60W-6C) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	60	5	See Figure 5-18; Table 5-20	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product familymay be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-18. Embedded Server SKU (LV60W-6C) T_{CASE} Thermal Profile

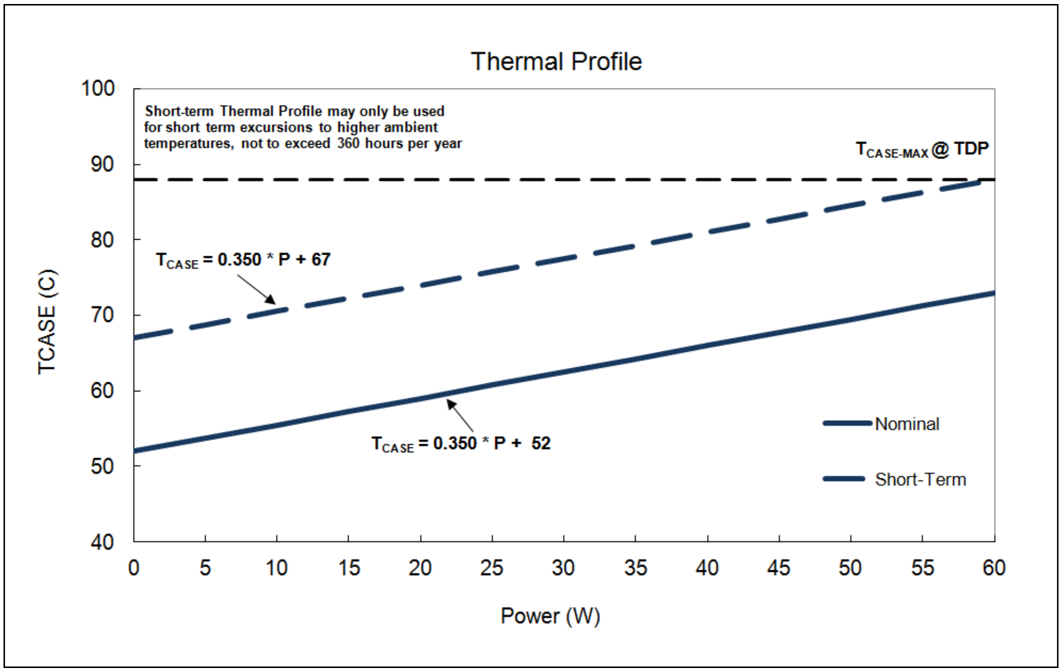

Note:
1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-20 for discrete points that constitute the thermal profile.

2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Figure 5-19. Embedded Server SKU (LV60W-6C) DTS Thermal Profile

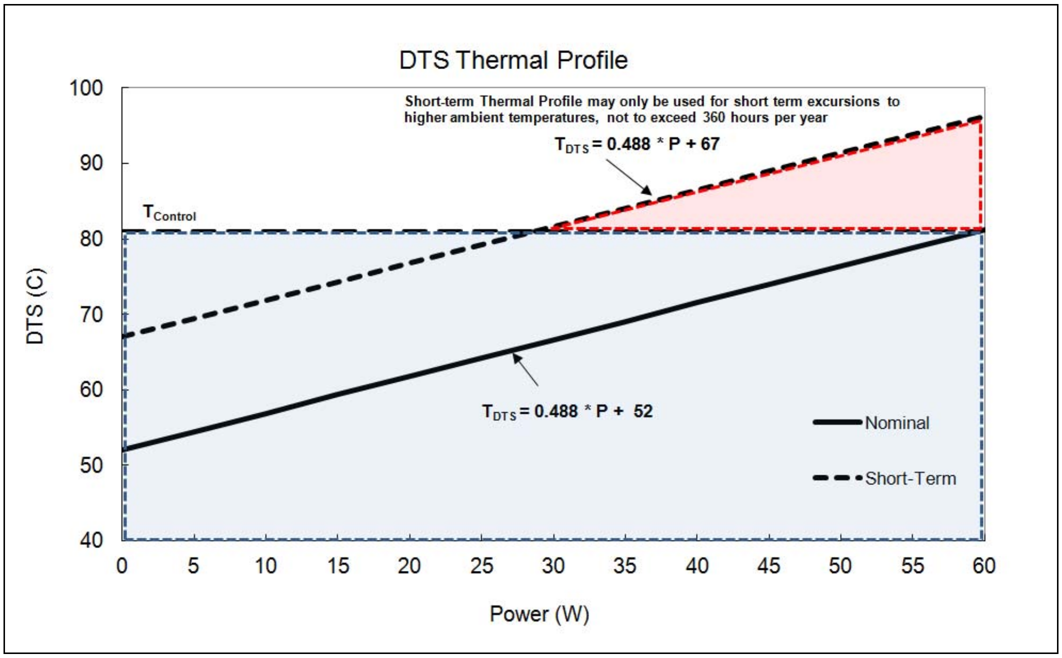

Notes:
1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-20 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance. As indicated by the blue shaded region, operation at DTS temperatures up to Tcontrol is permitted at all power levels.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Table 5-20. Embedded Server SKU (LV60W-6C) Thermal Profiles (Sheet 1 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
0	52	67	52	67
5	54	69	54	69
10	56	71	57	72
15	57	72	59	74
20	59	74	62	77
25	61	76	64	79
30	63	78	67	82

Table 5-20. Embedded Server SKU (LV60W-6C) Thermal Profiles (Sheet 2 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
35	64	79	69	84
40	66	81	72	87
45	68	83	74	89
50	70	84	76	91
55	71	86	79	94
60	73	88	81	96

5.1.3.3 Embedded LV50W-4C Thermal Specifications

Table 5-21. Embedded Server SKU (LV50W-4C) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	50	5	See Figure 5-20; Table 5-22	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-20. Embedded Server SKU (LV50W-4C) T_{CASE} Thermal Profile

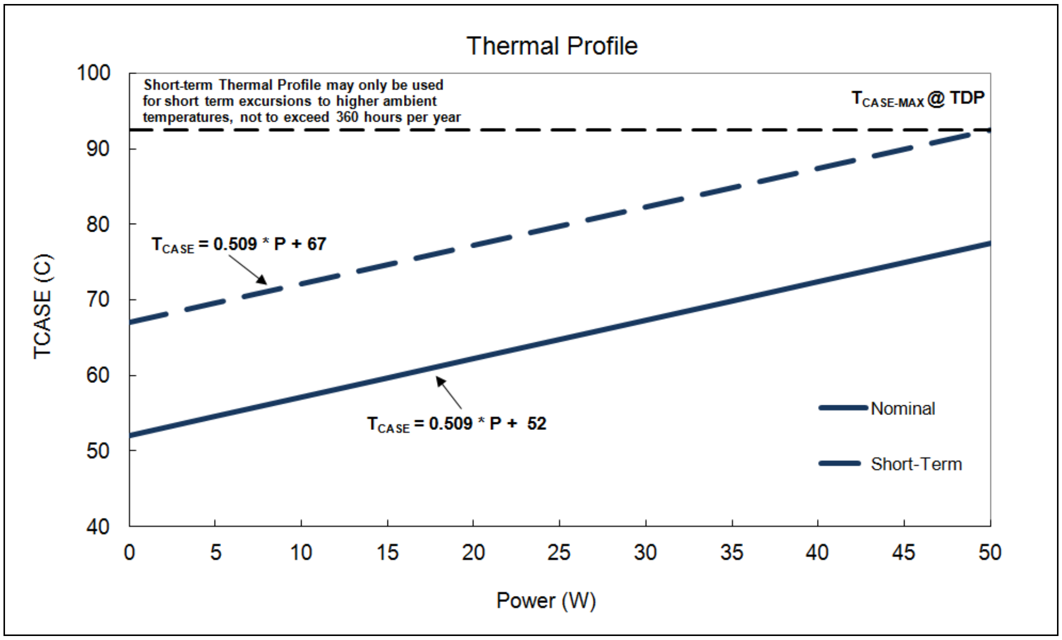

Note:

1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-22 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Figure 5-21. Embedded Server SKU (LV50W-4C) DTS Thermal Profile

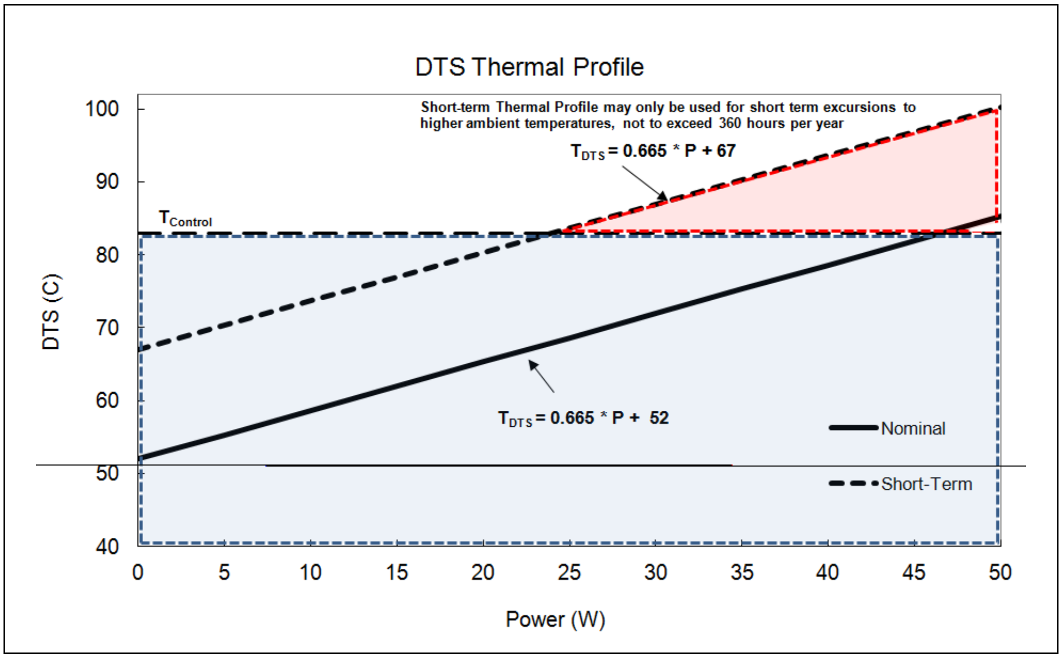

Notes:
1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-22 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance. As indicated by the blue shaded region, operation at DTS temperatures up to Tcontrol is permitted at all power levels.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Table 5-22. Embedded Server SKU (LV50W-4C) Thermal Profiles (Sheet 1 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	**Nominal**	**Short Term**	**Nominal**	**Short Term**
0	52	67	52	67
5	55	70	55	70
10	57	72	59	74
15	60	75	62	77
20	62	77	65	80

Table 5-22. Embedded Server SKU (LV50W-4C) Thermal Profiles (Sheet 2 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
25	65	80	69	84
30	67	82	72	87
35	70	85	75	90
40	72	87	79	94
45	75	90	82	97
50	77	92	85	100

5.1.3.4 Storage SKU LV40W-24C Thermal Specifications

Table 5-23. Storage SKU (LV40W-24C) T_{CASE} Thermal Specifications

Core Frequency	Thermal Design Power (W)	Minimum T_{CASE} (°C)	Maximum T_{CASE} (°C)	Notes
Launch to FMB	40	5	See Figure 5-22; Table 5-24	1, 2, 3, 4, 5

Notes:
1. These values are specified at V_{CC_MAX} for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CC} exceeds V_{CC_MAX} at specified ICC. Please refer to the electrical loadline specifications in Chapter 7.
2. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum T_{CASE}.
3. These specifications are based on final silicon characterization.
4. Power specifications are defined at all VIDs found in Table 7-3. The Intel® Xeon® processor E5-2400 product family may be delivered under multiple VIDs for each frequency.
5. FMB, or Flexible Motherboard, guidelines provide a design target for meeting all planned processor frequency requirements.

Figure 5-22. Storage SKU (LV40W-24C) T_{CASE} Thermal Profile

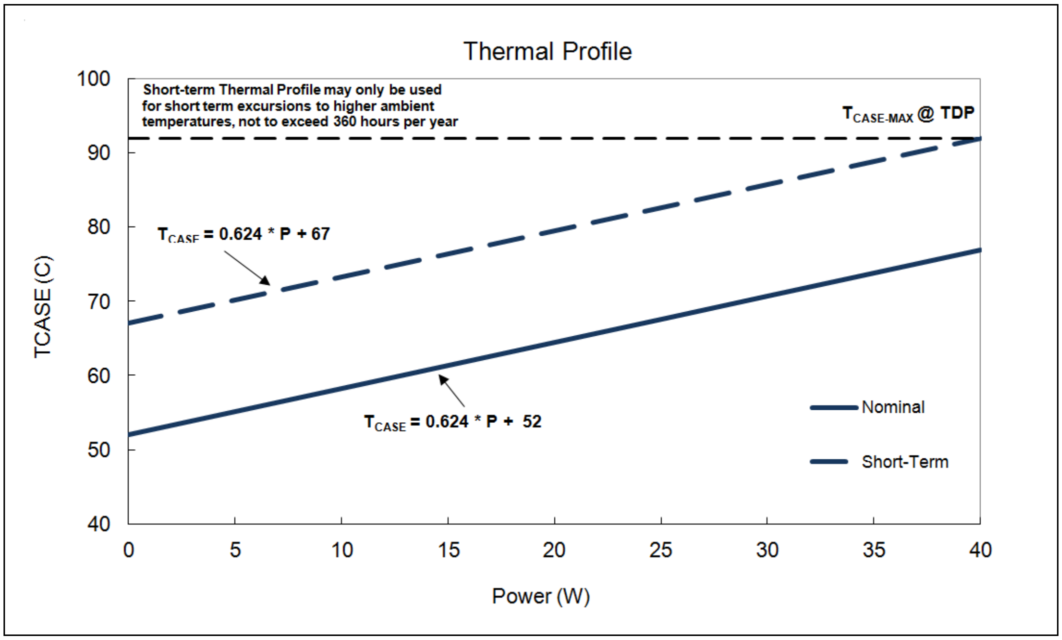

Note:

1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-24 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Figure 5-23. Storage SKU (LV40W-24C) DTS Thermal Profile

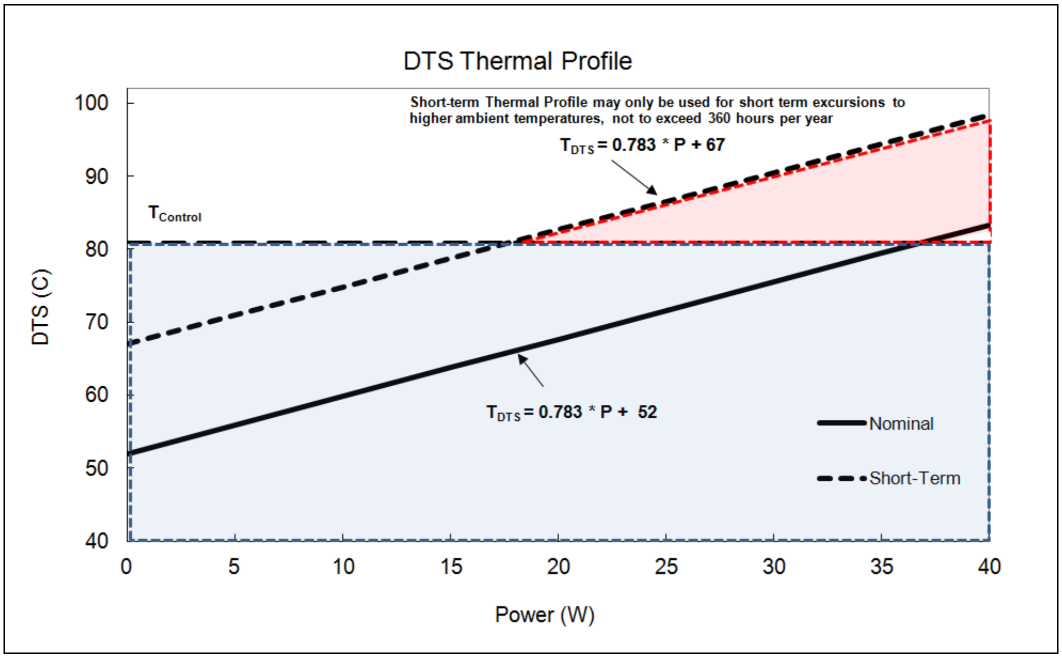

Notes:
1. The Thermal Profile is representative of a volumetrically constrained platform. Please refer to Table 5-24 for discrete points that constitute the thermal profile.
2. Implementation of this Thermal Profile should result in virtually no TCC activation. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for system and environmental implementation details.
3. The Nominal Thermal Profile must be used for all normal operating conditions or for products that do not require NEBS Level 3 compliance. As indicated by the blue shaded region, operation at DTS temperatures up to Tcontrol is permitted at all power levels.
4. The Short-Term Thermal Profile may only be used for short-term excursions to higher ambient operating temperatures, not to exceed 96 hours per instance, 360 hours per year, and a maximum of 15 instances per year, as compliant with NEBS Level 3. Operation at the Short-Term Thermal Profile for durations exceeding 360 hours per year violate the processor thermal specifications and may result in permanent damage to the processor.

Table 5-24. Storage SKU (LV40W-24C) Thermal Profiles (Sheet 1 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
0	52	67	52	67
5	55	70	56	71
10	58	73	60	75
15	61	76	64	79
20	64	79	68	83

Table 5-24. Storage SKU (LV40W-24C) Thermal Profiles (Sheet 2 of 2)

	Maximum T_{CASE} (°C)		Maximum DTS (°C)	
Power (W)	Nominal	Short Term	Nominal	Short Term
25	68	83	72	87
30	71	86	75	90
35	74	89	79	94
40	77	92	83	98

5.1.4 Thermal Metrology

The minimum and maximum case temperatures (T_{CASE}) are specified in Table 5-2 through Table 5-13. and are measured at the geometric top center of the processor integrated heat spreader (IHS). Figure 5-24 illustrates the location where T_{CASE} temperature measurements should be made. For detailed guidelines on temperature measurement methodology, refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

Figure 5-24. Case Temperature (T_{CASE}) Measurement Location

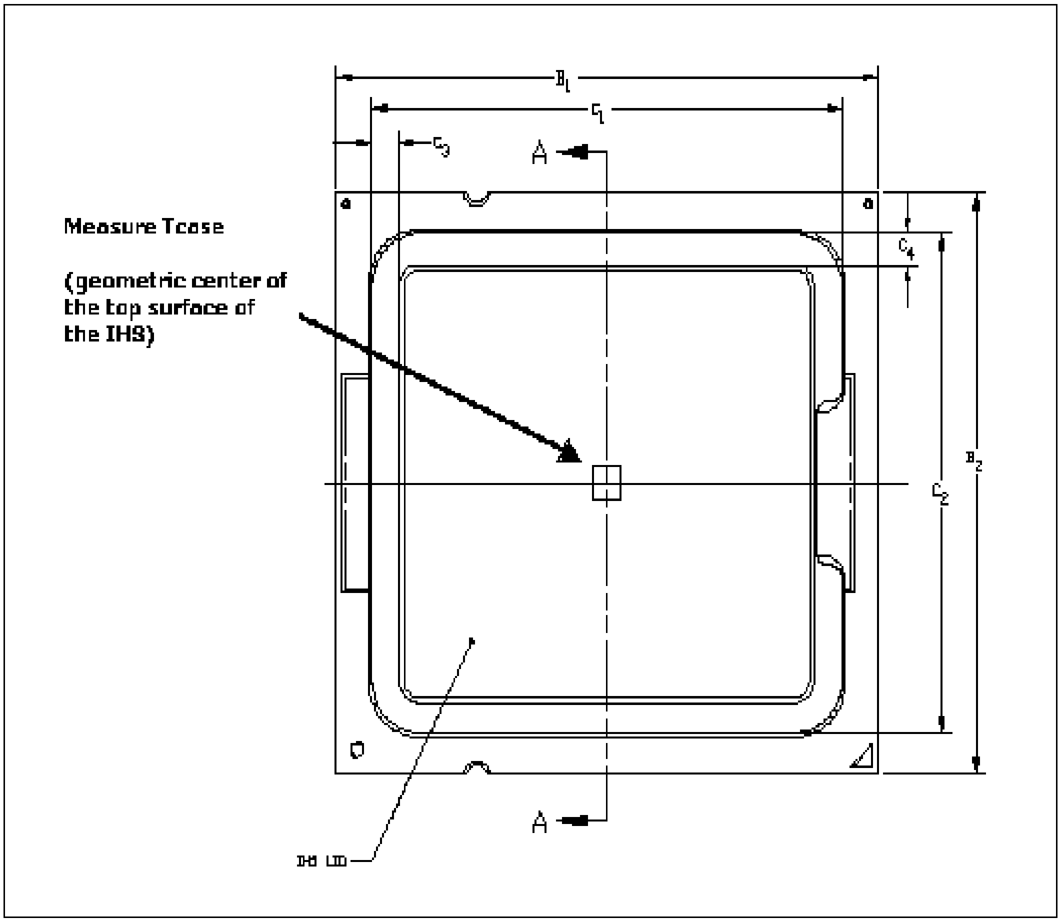

Notes:
1. Figure is not to scale and is for reference.
2. B1: Max = 45.07 mm, Min = 44.93 mm
3. B2: Max = 42.57 mm, Min = 42.43 mm

4. C1: Max = 39.1 mm, Min = 38.9 mm
5. C2: Max = 36.6 mm, Min = 36.4 mm
6. C3: Max = 2.3 mm, Min = 2.2 mm
7. C4: Max = 2.3 mm, Min = 2.2 mm

5.2 Processor Core Thermal Features

5.2.1 Processor Temperature

A new feature in the processor is a software readable field in the TEMPERATURE_TARGET MSR register that contains the minimum temperature at which the TCC will be activated and PROCHOT_N will be asserted. The TCC activation temperature is calibrated on a part-by-part basis and normal factory variation may result in the actual TCC activation temperature being higher than the value listed in the register. TCC activation temperatures may change based on processor stepping, frequency or manufacturing efficiencies.

5.2.2 Adaptive Thermal Monitor

The Adaptive Thermal Monitor feature provides an enhanced method for controlling the processor temperature when the processor silicon reaches its maximum operating temperature. Adaptive Thermal Monitor uses Thermal Control Circuit (TCC) activation to reduce processor power via a combination of methods. The first method (Frequency/ SVID control) involves the processor adjusting its operating frequency (via the core ratio multiplier) and input voltage (via the SVID signals). This combination of reduced frequency and voltage results in a reduction to the processor power consumption. The second method (clock modulation) reduces power consumption by modulating (starting and stopping) the internal processor core clocks. The processor intelligently selects the appropriate TCC method to use on a dynamic basis. BIOS is not required to select a specific method.

The Adaptive Thermal Monitor feature must be enabled for the processor to be operating within specifications. Snooping and interrupt processing are performed in the normal manner while the TCC is active.

With a properly designed and characterized thermal solution, it is anticipated that the TCC would be activated for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be so minor that it would be immeasurable. An under-designed thermal solution that is not able to prevent excessive activation of the TCC in the anticipated ambient environment may cause a noticeable performance loss, and in some cases may result in a T_C that exceeds the specified maximum temperature which may affect the long-term reliability of the processor. In addition, a thermal solution that is significantly under-designed may not be capable of cooling the processor even when the TCC is active continuously. Refer to the *Intel® Xeon® Processor E5-2400 Product Family Thermal / Mechanical Design Guide* for information on designing a compliant thermal solution.

The duty cycle for the TCC, when activated by the Thermal Monitor, is factory configured and cannot be modified. The Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines.

5.2.2.1 Frequency/SVID Control

The processor uses Frequency/SVID control whereby TCC activation causes the processor to adjust its operating frequency (via the core ratio multiplier) and VCC input voltage (via the SVID signals). This combination of reduced frequency and voltage results in a reduction to the processor power consumption.

This method includes multiple operating points, each consisting of a specific operating frequency and voltage. The first operating point represents the normal operating condition for the processor. The remaining points consist of both lower operating frequencies and voltages. When the TCC is activated, the processor automatically transitions to the new lower operating frequency. This transition occurs very rapidly (on the order of microseconds).Once the new operating frequency is engaged, the processor will transition to the new core operating voltage by issuing a new SVID code to the VCC voltage regulator. The voltage regulator must support dynamic SVID steps to support this method. During the voltage change, it will be necessary to transition through multiple SVID codes to reach the target operating voltage. Each step will be one SVID table entry (see Table 7-3). The processor continues to execute instructions during the voltage transition. Operation at the lower voltages reduces the power consumption of the processor.

A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the operating frequency and voltage transition back to the normal system operating point via the intermediate SVID/frequency points. Transition of the SVID code will occur first, to insure proper operation once the processor reaches its normal operating frequency. Refer to Figure 5-25 for an illustration of this ordering.

Figure 5-25. Frequency and Voltage Ordering

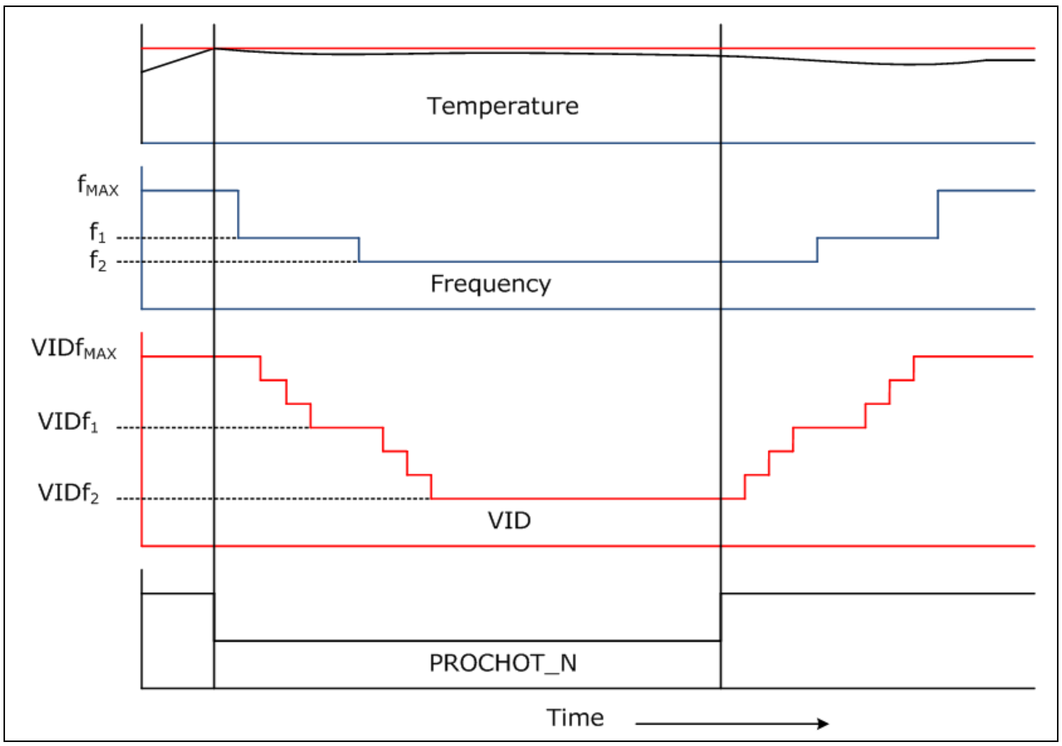

5.2.2.2 Clock Modulation

Clock modulation is performed by alternately turning the clocks off and on at a duty cycle specific to the processor (factory configured to 37.5% on and 62.5% off for TM1). The period of the duty cycle is configured to 32 microseconds when the TCC is active. Cycle times are independent of processor frequency. A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the TCC goes inactive and clock modulation ceases. Clock modulation is automatically engaged as part of the TCC activation when the Frequency/SVID targets are at their minimum settings. It may also be initiated by software at a configurable duty cycle.

5.2.3 On-Demand Mode

The processor provides an auxiliary mechanism that allows system software to force the processor to reduce its power consumption. This mechanism is referred to as "On-Demand" mode and is distinct from the Adaptive Thermal Monitor feature. On-Demand mode is intended as a means to reduce system level power consumption. Systems must not rely on software usage of this mechanism to limit the processor temperature. If bit 4 of the IA32_CLOCK_MODULATION MSR is set to a '1', the processor will immediately reduce its power consumption via modulation (starting and stopping) of the internal core clock, independent of the processor temperature. When using On-Demand mode, the duty cycle of the clock modulation is programmable via bits 3:0 of the same IA32_CLOCK_MODULATION MSR. In On-Demand mode, the duty cycle can be programmed from 6.25% on / 93.75% off to 93.75% on / 6.25% off in 6.25% increments. On-Demand mode may be used in conjunction with the Adaptive Thermal Monitor; however, if the system tries to enable On-Demand mode at the same time the TCC is engaged, the factory configured duty cycle of the TCC will override the duty cycle selected by the On-Demand mode.

5.2.4 PROCHOT_N Signal

An external signal, PROCHOT_N (processor hot), is asserted when the processor core temperature has reached its maximum operating temperature. If Adaptive Thermal Monitor is enabled (note it must be enabled for the processor to be operating within specification), the TCC will be active when PROCHOT_N is asserted. The processor can be configured to generate an interrupt upon the assertion or de-assertion of PROCHOT_N.

The PROCHOT_N signal is bi-directional in that it can either signal when the processor (any core) has reached its maximum operating temperature or be driven from an external source to activate the TCC. The ability to activate the TCC via PROCHOT_N can provide a means for thermal protection of system components.

As an output, PROCHOT_N will go active when the processor temperature monitoring sensor detects that one or more cores has reached its maximum safe operating temperature. This indicates that the processor Thermal Control Circuit (TCC) has been activated, if enabled. As an input, assertion of PROCHOT_N by the system will activate the TCC, if enabled, for all cores. TCC activation due to PROCHOT_N assertion by the system will result in the processor immediately transitioning to the minimum frequency and corresponding voltage (using Freq/SVID control). Clock modulation is not activated in this case. The TCC will remain active until the system de-asserts PROCHOT_N.

PROCHOT_N can allow voltage regulator (VR) thermal designs to target maximum sustained current instead of maximum current. Systems should still provide proper cooling for the VR, and rely on PROCHOT_N as a backup in case of system cooling failure. The system thermal design should allow the power delivery circuitry to operate within its temperature specification even while the processor is operating at its Thermal Design Power.

With a properly designed and characterized thermal solution, it is anticipated that PROCHOT_N will be asserted for very short periods of time when running the most power intensive applications. An under-designed thermal solution that is not able to prevent excessive assertion of PROCHOT_N in the anticipated ambient environment may cause a noticeable performance loss. Refer to the appropriate platform design guide and for details on implementing the bi-directional PROCHOT_N feature.

5.2.5 THERMTRIP_N Signal

Regardless of whether Adaptive Thermal Monitor is enabled, in the event of a catastrophic cooling failure, the processor will automatically shut down when the silicon has reached an elevated temperature (refer to the THERMTRIP_N definition in Section 6, "Signal Descriptions"). At this point, the THERMTRIP_N signal will go active and stay active. THERMTRIP_N activation is independent of processor activity and does not generate any Intel QuickPath Interconnect transactions. If THERMTRIP_N is asserted, all processor supplies (VCC, VTTA, VTTD, VSA, VCCPLL, VCCD) must be removed within the timeframe provided. The temperature at which THERMTRIP_N asserts is not user configurable and is not software visible.

5.2.6 Integrated Memory Controller (IMC) Thermal Features

5.2.6.1 DRAM Throttling Options

The Integrated Memory Controller (IMC) has two, independent mechanisms that cause system memory throttling:

- Open Loop Thermal Throttling (OLTT) and Hybrid OLTT (OLTT_Hybrid)
- Closed Loop Thermal Throttling (CLTT) and Hybrid CLTT (CLTT_Hybrid)

5.2.6.1.1 Open Loop Thermal Throttling (OLTT)

Pure energy based estimation for systems with no BMC or Intel ME. No memory temperature information is provided by the platform or DIMMs. The CPU is informed of the ambient temperature estimate by the BIOS or by a device via the PECI interface. DIMM temperature estimates and bandwidth control are monitored and managed by the PCU on a per rank basis.

5.2.6.1.2 Hybrid Open Loop Thermal Throttling (OLTT_Hybrid)

Temperature information is provided by the platform (for example, BMC or Intel® Management Engine (Intel® ME)) through PECI and the PCU interpolates gaps with energy based estimations.

5.2.6.1.3 Closed Loop Thermal Throttling (CLTT)

The processor periodically samples temperatures from the DIMM TSoD devices over a programmable interval. The PCU determines the hottest DIMM rank from TSoD data and informs the integrated memory controller for use in bandwidth throttling decisions.

5.2.6.2 Hybrid Closed Loop Thermal Throttling (CLTT_Hybrid)

The processor periodically samples temperature from the DIMM TSoD devices over a programmable interval and interpolates gaps or the BMC/Intel® ME samples a motherboard thermal sensor in the memory subsection and provides this data to the PCU via the PECI interface. This data is combined with an energy based estimations calculated by the PCU. When needed, system memory is then throttled using CAS bandwidth control. The processor supports dynamic reprogramming of the memory thermal limits based on system thermal state by the BMC or Intel ME.

5.2.6.3 MEM_HOT_C1_N and MEM_HOT_C23_N Signal

The processor includes new bi-directional memory thermal status signals useful for manageability schemes. Each signal presents and receives thermal status for a pair of memory channels (channel 1 and channels 2 and 3).

- Input Function: The processor can periodically sense the MEM_HOT_{C1/C23}_N signals to detect if the platform is requesting a memory throttling event. Manageability hardware could drive this signal due to a memory voltage regulator thermal or electrical issue or because of a detected system thermal event (for example, fan is going to fail) other system devices are exceeding their thermal target. The input sense period of these signals are programmable, 100 us is the default value. The input sense assertion time recognized by the processor is programmable, 1 us is the default value. If the sense assertion time is programmed to zero, then the processor ignores all external assertions of MEM_HOT_{C1/C23}_N signals (in effect they become outputs).
- Output Function: The output behavior of the MEM_HOT_{C1/C23}_N signals supports Level mode. In this mode, MEM_HOT_{C1/C23}_N event temperatures are programmable via TEMP_OEM_HI, TEMP_LOW, TEMP_MID, and TEMP_HI threshold settings in the iMC. In Level mode, when asserted, the signal indicates to the platform that a BIOS-configured thermal threshold has been reached by one or more DIMMs in the covered channel pair.

5.2.6.4 Integrated Dual SMBus Master Controllers for SMI

The processor includes two integrated SMBus master controllers running at 100 KHz for dedicated PCU access to the serial presence detect (SPD) devices and thermal sensors (TSoD) on the DIMMs. Each controller is responsible for a pair of memory channels and supports up to four SMBus slave devices. Note that clock-low stretching is not supported by the processor. To avoid design complexity and minimize package C-state transitions, the SMBus interface between the processor and DIMMs must be connected.

The SMBus controllers for the system memory interface support the following SMBus protocols/commands:

- Random byte Read
- Byte Write
- I^2C* Write to Pointer Register
- I^2C Present Pointer Register Word Read
- I^2C Pointer Write Register Read.

Refer to the *System Management Bus (SMBus) Specification, Revision 2.0* for standing timing protocols and specific command structure details.

§

6 Signal Descriptions

This chapter describes the processor signals. They are arranged in functional groups according to their associated interface or category.

6.1 System Memory Interface Signals

Table 6-1. Memory Channel DDR1, DDR2, DDR3

Signal Name	Description
DDR{1/2/3}_BA[2:0]	Bank Address. Defines the bank which is the destination for the current Activate, Read, Write, or Precharge command.
DDR{1/2/3}_CAS_N	Column Address Strobe.
DDR{1/2/3}_CKE[3:0]	Clock Enable.
DDR{1/2/3}_CLK_DN[3:0] DDR{1/2/3}_CLK_DP[3:0]	Differential clocks to the DIMM. All command and control signals are valid on the rising edge of clock.
DDR{1/2/3}_CS_N[7:0]	Chip Select. Each signal selects one rank as the target of the command and address.
DDR{1/2/3}_DQ[63:00]	Data Bus. DDR3 Data bits.
DDR{1/2/3}_DQS_DP[17:00] DDR{1/2/3}_DQS_DN[17:00]	Data strobes. Differential pair, Data/ECC Strobe. Differential strobes latch data/ECC for each DRAM. Different numbers of strobes are used depending on whether the connected DRAMs are x4,x8. Driven with edges in center of data, receive edges are aligned with data edges.
DDR{1/2/3}_ECC[7:0]	Check bits. An error correction code is driven along with data on these lines for DIMMs that support that capability
DDR{1/2/3}_MA[15:00]	Memory Address. Selects the Row address for Reads and writes, and the column address for activates. Also used to set values for DRAM configuration registers.
DDR{1/2/3}_MA_PAR	Odd parity across Address and Command.
DDR{1/2/3}_ODT[3:0]	On Die Termination. Enables DRAM on die termination during Data Write or Data Read transactions.
DDR{1/2/3}_PAR_ERR_N	Parity Error detected by Registered DIMM (one for each channel).
DDR{1/2/3}_RAS_N	Row Address Strobe.
DDR{1/2/3}_WE_N	Write Enable.

Table 6-2. Memory Channel Miscellaneous

Signal Name	Description
DDR_RESET_C1_N DDR_RESET_C23_N	System memory reset: Reset signal from processor to DRAM devices on the DIMMs. DDR_RESET_C1_N is used for memory channel 1 while DDR_RESET_C23_N is used for memory channels 2 and 3.
DDR_SCL_C1 DDR_SCL_C23	SMBus clock for the dedicated interface to the serial presence detect (SPD) and thermal sensors (TSoD) on the DIMMs. DDR_SCL_C1 is used for memory channel 1 while DDR_SCL_C23 is used for memory channels 2 and 3.
DDR_SDA_C1 DDR_SDA_C23	SMBus data for the dedicated interface to the serial presence detect (SPD) and thermal sensors (TSoD) on the DIMMs. DDR_SDA_C1 is used for memory channel 1 while DDR_SDA_C23 is used for memory channels 2 and 3.
DDR_VREFDQRX_C1 DDR_VREFDQRX_C23	Voltage reference for system memory reads. DDR_VREFDQRX_C1 is used for memory channel 1 while DDR_VREFDQRX_C23 is used for memory channels 2 and 3.
DDR_VREFDQTX_C1 DDR_VREFDQTX_C23	Voltage reference for system memory writes. DDR_VREFDQTX_C1 is used for memory channel 1 while DDR_VREFDQTX_C23 is used for memory channels 2 and 3. These signals are not connected.
DDR{1/23}_RCOMP[2:0]	System memory impedance compensation. Impedance compensation must be terminated on the system board using a precision resistor. See the appropriate Platform Design Guide (PDG) for implementation details.
DRAM_PWR_OK_C1 DRAM_PWR_OK_C23	Power good input signal used to indicate that the VCCD power supply is stable for memory channel 1 and channels 2 & 3.

6.2 PCI Express* Based Interface Signals

Note: PCI Express* Ports 1 and 3 Signals are receive and transmit differential pairs.

Table 6-3. PCI Express* Port 1 Signals

Signal Name	Description
PE1A_RX_DN[3:0] PE1A_RX_DP[3:0]	PCIe* Receive Data Input
PE1B_RX_DN[7:4] PE1B_RX_DP[7:4]	PCIe* Receive Data Input
PE1A_TX_DN[3:0] PE1A_TX_DP[3:0]	PCIe* Transmit Data Output
PE1B_TX_DN[7:4] PE1B_TX_DP[7:4]	PCIe* Transmit Data Output

Table 6-4. PCI Express* Port 3 Signals (Sheet 1 of 2)

Signal Name	Description
PE3A_RX_DN[3:0] PE3A_RX_DP[3:0]	PCIe* Receive Data Input
PE3B_RX_DN[7:4] PE3B_RX_DP[7:4]	PCIe* Receive Data Input
PE3C_RX_DN[11:8] PE3C_RX_DP[11:8]	PCIe* Receive Data Input
PE3D_RX_DN[15:12] PE3D_RX_DP[15:12]	PCIe* Receive Data Input

Table 6-4. PCI Express* Port 3 Signals (Sheet 2 of 2)

Signal Name	Description
PE3A_TX_DN[3:0] PE3A_TX_DP[3:0]	PCIe* Transmit Data Output
PE3B_TX_DN[7:4] PE3B_TX_DP[7:4]	PCIe* Transmit Data Output
PE3C_TX_DN[11:8] PE3C_TX_DP[11:8]	PCIe* Transmit Data Output
PE3D_TX_DN[15:12] PE3D_TX_DP[15:12]	PCIe* Transmit Data Output

Table 6-5. PCI Express* Miscellaneous Signals

Signal Name	Description
PE_RBIAS	This input is used to control PCI Express* bias currents. A 50 ohm 1% tolerance resistor must be connected from this land to VSS by the platform. PE_RBIAS is required to be connected as if the link is being used even when PCIe* is not used. Refer to the appropriate Platform Design Guide (PDG) for further details.
PE_RBIAS_SENSE	Provides dedicated bias resistor sensing to minimize the voltage drop caused by packaging and platform effects. PE_RBIAS_SENSE is required to be connected as if the link is being used even when PCIe* is not used. Refer to the appropriate Platform Design Guide (PDG) for further details.
PE_VREF_CAP	PCI Express* voltage reference used to measure the actual output voltage and comparing it to the assumed voltage. A 0.01uF capacitor must be connected from this land to VSS.
PEHPSCL	PCI Express* Hot-Plug SMBus Clock: Provides PCI Express* hot-plug support via a dedicated SMBus interface. Requires an external general purpose input/output (GPIO) expansion device on the platform.
PEHPSDA	PCI Express* Hot-Plug SMBus Data: Provides PCI Express* hot-plug support via a dedicated SMBus interface. Requires an external general purpose input/output (GPIO) expansion device on the platform.

Note: Refer to the appropriate Platform Design Guide (PDG) for additional implementation details.

6.3 DMI2/PCI Express* Port 0 Signals

Table 6-6. DMI2 and PCI Express* Port 0 Signals

Signal Name	Description
DMI_RX_DN[3:0] DMI_RX_DP[3:0]	DMI2 Receive Data Input
DMI_TX_DP[3:0] DMI_TX_DN[3:0]	DMI2 Transmit Data Output

6.4 Intel QuickPath Interconnect Signals

Table 6-7. Intel QPI Port Signals

Signal Name	Description
QPI_CLKRX_DN/DP	Reference Clock Differential Input. These pins provide the PLL reference clock differential input. The Intel QPI forward clock frequency is half the Intel QPI data rate.
QPI_CLKTX_DN/DP	Reference Clock Differential Output. These pins provide the PLL reference clock differential input. The Intel QPI forward clock frequency is half the Intel QPI data rate.
QPI_DRX_DN/DP[19:00]	Intel QPI Receive data input.
QPI_DTX_DN/DP[19:00]	Intel QPI Transmit data output.

Table 6-8. Intel QPI Miscellaneous Signals

Signal Name	Description
QPI_RBIAS	This input is used to control Intel QPI bias currents. QPI_RBIAS is required to be connected as if the link is being used even when Intel QPI is not used. Refer to the appropriate Platform Design Guide (PDG) for further details.
QPI_RBIAS_SENSE	Provides dedicated bias resistor sensing to minimize the voltage drop caused by packaging and platform effects. QPI_RBIAS_SENSE is required to be connected as if the link is being used even when Intel QPI is not used. Refer to the appropriate Platform Design Guide (PDG) for further details.
QPI_VREF_CAP	Intel QPI voltage reference used to measure the actual output voltage and comparing it to the assumed voltage. Refer to the appropriate Platform Design Guide (PDG) for further details.

Note: Refer to the appropriate Platform Design Guide (PDG) for additional implementation details.

6.5 PECI Signal

Table 6-9. PECI Signals

Signal Name	Description
PECI	PECI (Platform Environment Control Interface) is the serial sideband interface to the processor and is used primarily for thermal, power and error management. Details regarding the PECI electrical specifications, protocols and functions can be found in the Platform Environment Control Interface Specification.

6.6 System Reference Clock Signals

Table 6-10. System Reference Clock (BCLK) Signals

Signal Name	Description
BCLK{0/1}_D[N/P]	Reference Clock Differential input. These pins provide the PLL reference clock differential input into the processor. Both 100MHz BCLK0 and BCLK1 from the same clock source provide the required reference clock inputs to the various PLLs inside the CPU.

6.7 JTAG and TAP Signals

Table 6-11. JTAG and TAP Signals

Signal Name	Description
BPM_N[7:0]	Breakpoint and Performance Monitor Signals: I/O signals from the processor that indicate the status of breakpoints and programmable counters used for monitoring processor performance. These are 100 MHz signals.
EAR_N	External Alignment of Reset, used to bring the processor up into a deterministic state. This signal is pulled up on the die, refer to Table 7-6 for details.
PRDY_N	Probe Mode Ready is a processor output used by debug tools to determine processor debug readiness.
PREQ_N	Probe Mode Request is used by debug tools to request debug operation of the processor.
TCK	TCK (Test Clock) provides the clock input for the processor Test Bus (also known as the Test Access Port).
TDI	TDI (Test Data In) transfers serial test data into the processor. TDI provides the serial input needed for JTAG specification support.
TDO	TDO (Test Data Out) transfers serial test data out of the processor. TDO provides the serial output needed for JTAG specification support.
TMS	TMS (Test Mode Select) is a JTAG specification support signal used by debug tools.
TRST_N	TRST_N (Test Reset) resets the Test Access Port (TAP) logic. TRST_N must be driven low during power on Reset.

Note: Refer to the appropriate Platform Design Guide (PDG) for Debug Port implementation details.

6.8 Serial VID Interface (SVID) Signals

Table 6-12. SVID Signals

SVIDALERT_N	Serial VID alert.
SVIDCLK	Serial VID clock.
SVIDDATA	Serial VID data out.

6.9 Processor Asynchronous Sideband and Miscellaneous Signals

Table 6-13. Processor Asynchronous Sideband Signals (Sheet 1 of 4)

Signal Name	Description
BIST_ENABLE	BIST Enable Strap. Input which allows the platform to enable or disable built-in self test (BIST) on the processor. This signal is pulled up on the die, refer to Table 7-6 for details.

Table 6-13. Processor Asynchronous Sideband Signals (Sheet 2 of 4)

Signal Name	Description
BMCINIT	BMC Initialization Strap. Indicates whether Service Processor Boot Mode should be used. Used in combination with FRMAGENT and SOCKET_ID inputs. • 0: Service Processor Boot Mode Disabled. Example boot modes: Local PCH (this processor hosts a legacy PCH with firmware behind it), Intel QPI Link Boot (for processors one hop away from the FW agent), or Intel QPI Link Init (for processors more than one hop away from the firmware agent). • 1: Service Processor Boot Mode Enabled. In this mode of operation, the processor performs the absolute minimum internal configuration and then waits for the Service Processor to complete its initialization. The socket boots after receiving a “GO” handshake signal via a firmware scratchpad register. This signal is pulled down on the die, refer to Table 7-6 for details.
CAT_ERR_N	Indicates that the system has experienced a fatal or catastrophic error and cannot continue to operate. The processor will assert CAT_ERR_N for nonrecoverable machine check errors and other internal unrecoverable errors. It is expected that every processor in the system will wire-OR CAT_ERR_N for all processors. Since this is an I/O land, external agents are allowed to assert this land which will cause the processor to take a machine check exception. This signal is sampled after PWRGOOD assertion. On the processor, CAT_ERR_N is used for signaling the following types of errors: • Legacy MCERR's, CAT_ERR_N is asserted for 16 BCLKs. • Legacy IERR's, CAT_ERR_N remains asserted until warm or cold reset.
CPU_ONLY_RESET	Resets all the processors on the platform without resetting the DMI2 links.
ERROR_N[2:0]	Error status signals for integrated I/O (IIO) unit: • 0 = Hardware correctable error (no operating system or firmware action necessary) • 1 = Non-fatal error (operating system or firmware action required to contain and recover) • 2 = Fatal error (system reset likely required to recover)
FRMAGENT	Bootable Firmware Agent Strap. This input configuration strap used in combination with SOCKET_ID to determine whether the socket is a legacy socket, bootable firmware agent is present, and DMI links are used in PCIe* mode (instead of DMI2 mode). The firmware flash ROM is located behind the local PCH attached to the processor via the DMI2 interface.This signal is pulled down on the die, refer to Table 7-6 for details.
MEM_HOT_C1_N MEM_HOT_C23_N	Memory throttle control. MEM_HOT_C1_N and MEM_HOT_C23_N signals have two modes of operation – input and output mode. Input mode is externally asserted and is used to detect external events such as VR_HOT# from the memory voltage regulator and causes the processor to throttle the appropriate memory channels. Output mode is asserted by the processor known as level mode. In level mode, the output indicates that a particular branch of memory subsystem is hot. MEM_HOT_C1_N is used for memory channel 1 while MEM_HOT_C23_N is used for memory channels 2 & 3.
PMSYNC	Power Management Sync. A sideband signal to communicate power management status from the Platform Controller Hub (PCH) to the processor.
PROCHOT_N	PROCHOT_N will go active when the processor temperature monitoring sensor detects that the processor has reached its maximum safe operating temperature. This indicates that the processor Thermal Control Circuit has been activated, if enabled. This signal can also be driven to the processor to activate the Thermal Control Circuit. This signal is sampled after PWRGOOD assertion. If PROCHOT_N is asserted at the deassertion of RESET_N, the processor will tristate its outputs.

Table 6-13. Processor Asynchronous Sideband Signals (Sheet 3 of 4)

Signal Name	Description
PWRGOOD	Power Good is a processor input. The processor requires this signal to be a clean indication that BCLK, VTTA/VTTD, VSA, VCCPLL, and VCCD supplies are stable and within their specifications. “Clean” implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until they come within specification. The signal must then transition monotonically to a high state. PWRGOOD can be driven inactive at any time, but clocks and power must again be stable before a subsequent rising edge of PWRGOOD. PWRGOOD transitions from inactive to active when all supplies except VCC are stable. VCC has a VBOOT of zero volts and is not included in PWRGOOD indication in this phase. However, for the active to inactive transition, if any CPU power supply (VCC, VTTA/VTTD, VSA, VCCD, or VCCPLL) is about to fail or is out of regulation, the PWRGOOD is to be negated. The signal must be supplied to the processor; it is used to protect internal circuits against voltage sequencing issues. It should be driven high throughout boundary scan operation. ***Note:*** VCC has a Vboot setting of 0.0V and is not included in the PWRGOOD indication and VSA has a Vboot setting of 0.9V. Refer to the *VR12/IMVP7 Pulse Width Modulation Specification*.
RESET_N	Asserting the RESET_N signal resets the processor to a known state and invalidates its internal caches without writing back any of their contents. Note some PLL, Intel QuickPath Interconnect and error states are not effected by reset and only PWRGOOD forces them to a known state.
RSVD	RESERVED. All signals that are RSVD must be left unconnected on the board. Refer to Section 7.1.10, “Reserved or Unused Signals” for details.
SAFE_MODE_BOOT	Safe mode boot Strap. SAFE_MODE_BOOT allows the processor to wake up safely by disabling all clock gating, this allows BIOS to load registers or patches if required. This signal is sampled after PWRGOOD assertion. The signal is pulled down on the die, refer to Table 7-6 for details.
SOCKET_ID[1:0]	Socket ID Strap. Socket identification configuration straps for establishing the PECI address, Intel QPI Node ID, and other settings. This signal is used in combination with FRMAGENT to determine whether the socket is a legacy socket, bootable firmware agent is present, and DMI links are used in PCIe* mode (instead of DMI2 mode). Each processor socket consumes one Node ID, and there are 128 Home Agent tracker entries. This signal is pulled down on the die, refer to Table 7-6 for details.
TEST[4:0]	Test[4:0] must be individually connected to an appropriate power source or ground through a resistor for proper processor operation. Refer to the appropriate Platform Design Guide (PDG) for additional implementation details.
THERMTRIP_N	Assertion of THERMTRIP_N (Thermal Trip) indicates one of two possible critical over-temperature conditions: One, the processor junction temperature has reached a level beyond which permanent silicon damage may occur and Two, the system memory interface has exceeded a critical temperature limit set by BIOS. Measurement of the processor junction temperature is accomplished through multiple internal thermal sensors that are monitored by the Digital Thermal Sensor (DTS). Simultaneously, the Power Control Unit (PCU) monitors external memory temperatures via the dedicated SMBus interface to the DIMMs. If any of the DIMMs exceed the BIOS defined limits, the PCU will signal THERMTRIP_N to prevent damage to the DIMMs. Once activated, the processor will stop all execution and shut down all PLLs. To further protect the processor, its core voltage (VCC), VTTA, VTTD, VSA, VCCPLL, VCCD supplies must be removed following the assertion of THERMTRIP_N. Once activated, THERMTRIP_N remains latched until RESET_N is asserted. While the assertion of the RESET_N signal may de-assert THERMTRIP_N, if the processor's junction temperature remains at or above the trip level, THERMTRIP_N will again be asserted after RESET_N is de-asserted. This signal can also be asserted if the system memory interface has exceeded a critical temperature limit set by BIOS. This signal is sampled after PWRGOOD assertion.
TXT_AGENT	Intel TXT Platform Enable Strap. 0 = Default. The socket is not the Intel TXT Agent. 1 = The socket is the Intel TXT Agent. In non-Scalable DP platforms, the legacy socket (identified by SOCKET_ID[1:0] = 00b) with Intel TXT Agent should always set the TXT_AGENT to 1b. On Scalable DP platforms the Intel TXT AGENT is at the Node Controller. Refer to the Platform Design Guide for more details. This signal is pulled down on the die, refer to Table 7-6 for details.

Table 6-13. Processor Asynchronous Sideband Signals (Sheet 4 of 4)

Signal Name	Description
TXT_PLTEN	Intel TXT Platform Enable Strap. 0 = The platform is not Intel TXT enabled. All sockets should be set to zero. Scalable DP (sDP) platforms should choose this setting if the Node Controller does not support Intel TXT. 1 = Default. The platform is Intel TXT enabled. All sockets should be set to one. In a non-Scalable DP platform this is the default. When this is set, Intel TXT functionality requires user to explicitly enable Intel TXT via BIOS setup. This signal is pulled up on the die, refer to Table 7-6 for details.

Table 6-14. Miscellaneous Signals

Signal Name	Description
IVT_ID_N	This output can be used by the platform to determine if the installed processor is a future processor planned for the Intel® Xeon® processor E5-2400 product family-based Platform. There is no connection to the processor silicon for this signal. This signal is also used by the VCCPLL and VTT rails to switch their output voltage to support future processors.
SKTOCC_N	SKTOCC_N (Socket occupied) is used to indicate that a processor is present. This is pulled to ground on the processor package; there is no connection to the processor silicon for this signal.

6.10 Processor Power and Ground Supplies

Table 6-15. Power and Ground Signals (Sheet 1 of 2)

Signal Name	Description
VCC	Variable power supply for the processor cores, lowest level caches (LLC), ring interface, and home agent. It is provided by a VRM/EVRD 12.0 compliant regulator for each CPU socket. The output voltage of this supply is selected by the processor, using the serial voltage ID (SVID) bus. ***Note:*** VCC has a Vboot setting of 0.0V and is not included in the PWRGOOD indication. Refer to the *VR12/IMVP7 Pulse Width Modulation Specification*.
VCC_SENSE VSS_VCC_SENSE	VCC_SENSE and VSS_VCC_SENSE provide an isolated, low impedance connection to the processor core power and ground. These signals must be connected to the voltage regulator feedback circuit, which insures the output voltage (that is, processor voltage) remains within specification. Please see the applicable platform design guide for implementation details.
VSA_SENSE VSS_VSA_SENSE	VSA_SENSE and VSS_VSA_SENSE provide an isolated, low impedance connection to the processor system agent (VSA) power plane. These signals must be connected to the voltage regulator feedback circuit, which insures the output voltage (that is, processor voltage) remains within specification. Please see the applicable platform design guide for implementation details.
VTTD_SENSE VSS_VTTD_SENSE	VTTD_SENSE and VSS_VTTD_SENSE provide an isolated, low impedance connection to the processor I/O power plane. These signals must be connected to the voltage regulator feedback circuit, which insures the output voltage (that is, processor voltage) remains within specification. Please see the applicable platform design guide for implementation details.

Table 6-15. Power and Ground Signals (Sheet 2 of 2)

Signal Name	Description
VCCD	Variable power supply for the processor system memory interface. Provided by two VRM/EVRD 12.0 compliant regulators per CPU socket. VCCD is used for memory channels 1, 2, and 3. The valid voltage of this supply (1.50 V or 1.35 V) is configured by BIOS after determining the operating voltages of the installed memory. ***Note:*** The processor must be provided VCCD for proper operation, even in configurations where no memory is populated. A VRM/EVRD 12.0 controller is recommended, but not required.
VCCPLL	Fixed power supply (1.V) for the processor phased lock loop (PLL).
VSA	Variable power supply for the processor system agent units. These include logic (non-I/O) for the integrated I/O controller, the integrated memory controller (iMC), the Intel QPI agent, and the Power Control Unit (PCU). The output voltage of this supply is selected by the processor, using the serial voltage ID (SVID) bus. ***Note:*** VSA has a Vboot setting of 0.9V. Refer to the *VR12/IMVP7 Pulse Width Modulation Specification*.
VSS	Processor ground node.
VTTA VTTD	Combined fixed analog and digital power supply for I/O sections of the processor Intel QPI interface, Direct Media Interface Gen 2 (DMI2) interface, and PCI Express* interface. These signals will also be referred to as VTT. Please see the appropriate Platform Design Guide (PDG)for implementation details.

§

7 Electrical Specifications

7.1 Processor Signaling

The processor includes 1356 lands, which utilize various signaling technologies. Signals are grouped by electrical characteristics and buffer type into various signal groups. These include DDR3 (Reference Clock, Command, Control, and Data), PCI Express*, DMI2, Intel QuickPath Interconnect, Platform Environmental Control Interface (PECI), System Reference Clock, SMBus, JTAG and Test Access Port (TAP), SVID Interface, Processor Asynchronous Sideband, Miscellaneous, and Power/Other signals. Refer to Table 7-5 for details.

Detailed layout, routing, and termination guidelines corresponding to these signal groups can be found in the applicable platform design guide (Refer to Section 1.7, "Related Documents").

Intel strongly recommends performing analog simulations of all interfaces. Please refer to Section 1.7, "Related Documents" for signal integrity model availability.

7.1.1 System Memory Interface Signal Groups

The system memory interface utilizes DDR3 technology, which consists of numerous signal groups. These include: Reference Clocks, Command Signals, Control Signals, and Data Signals. Each group consists of numerous signals, which may utilize various signaling technologies. Please refer to Table 7-5 for further details. Throughout this chapter the system memory interface maybe referred to as DDR3.

7.1.2 PCI Express* Signals

The PCI Express* Signal Group consists of PCI Express* ports 1and 3, and PCI Express* miscellaneous signals. Please refer to Table 7-5 for further details.

7.1.3 DMI2/PCI Express* Signals

The Direct Media Interface Gen 2 (DMI2) sends and receives packets and/or commands to the PCH. The DMI2 is an extension of the standard PCI Express* Specification. The DMI2/PCI Express* Signals consist of DMI2 receive and transmit input/output signals and a control signal to select DMI2 or PCIe* 2.0 operation for port 0. Please refer to Table 7-5 for further details.

7.1.4 Intel QuickPath Interconnect (Intel QPI)

The processor provides one Intel QPI port for high speed serial transfer between other processors. The port consists of two uni-directional links (for transmit and receive). A differential signaling scheme is utilized, which consists of opposite-polarity (DP, DN) signal pairs.

7.1.5 Platform Environmental Control Interface (PECI)

PECI is an Intel proprietary interface that provides a communication channel between Intel processors and chipset components to external system management logic and thermal monitoring devices. The processor contains a Digital Thermal Sensor (DTS) that reports a relative die temperature as an offset from Thermal Control Circuit (TCC) activation temperature. Temperature sensors located throughout the die are implemented as analog-to-digital converters calibrated at the factory. PECI provides an interface for external devices to read processor temperature, perform processor manageability functions, and manage processor interface tuning and diagnostics. Please refer to Section 2.5, "Platform Environment Control Interface (PECI)" for processor specific implementation details for PECI.

The PECI interface operates at a nominal voltage set by V_{TTD}. The set of DC electrical specifications shown in Table 7-16 is used with devices normally operating from a V_{TTD} interface supply.

7.1.5.1 Input Device Hysteresis

The PECI client and host input buffers must use a Schmitt-triggered input design for improved noise immunity. Please refer to Figure 7-1 and Table 7-16.

Figure 7-1. Input Device Hysteresis

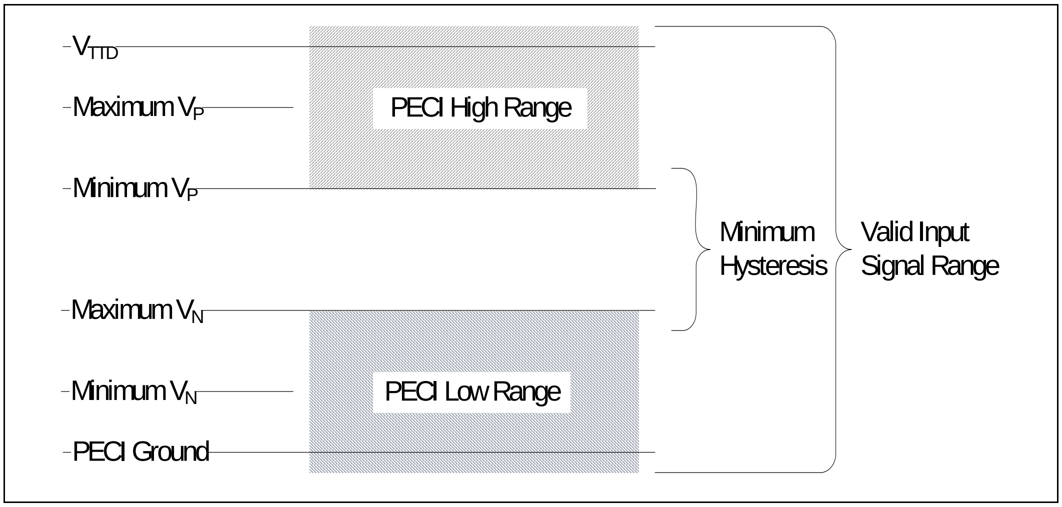

7.1.6 System Reference Clocks (BCLK{0/1}_DP, BCLK{0/1}_DN)

The processor core, processor uncore, Intel® QuickPath Interconnect link, PCI Express* and DDR3 memory interface frequencies) are generated from BCLK{0/1}_DP and BCLK{0/1}_DN signals. There is no direct link between core frequency and Intel QuickPath Interconnect link frequency (for example, no core frequency to Intel QuickPath Interconnect multiplier). The processor maximum core frequency, Intel QuickPath Interconnect link frequency and DDR memory frequency are set during manufacturing. It is possible to override the processor core frequency setting using software. This permits operation at lower core frequencies than the factory set maximum core frequency.

The processor core frequency is configured during reset by using values stored within the device during manufacturing. The stored value sets the lowest core multiplier at which the particular processor can operate. If higher speeds are desired, the appropriate ratio can be configured via the IA32_PERF_CTL MSR (MSR 199h); Bits [15:0].

Clock multiplying within the processor is provided by the internal phase locked loop (PLL), which requires a constant frequency BCLK{0/1}_DP, BCLK{0/1}_DN input, with exceptions for spread spectrum clocking. DC specifications for the BCLK{0/1}_DP, BCLK{0/1}_DN inputs are provided in Table 7-17. These specifications must be met while also meeting the associated signal quality specifications outlined in Section 7.9.

7.1.6.1 PLL Power Supply

An on-die PLL filter solution is implemented on the processor. Refer to Table 7-12 for DC specifications and to the applicable platform design guide for decoupling and routing guidelines.

7.1.7 JTAG and Test Access Port (TAP) Signals

Due to the voltage levels supported by other components in the JTAG and Test Access Port (TAP) logic, Intel recommends the processor be first in the TAP chain, followed by any other components within the system. Please refer to the *Intel® Xeon® Processor E5-2400 Product Family BSDL (Boundary Scan Description Language)* for more details. A translation buffer should be used to connect to the rest of the chain unless one of the other components is capable of accepting an input of the appropriate voltage. Two copies of each signal may be required with each driving a different voltage level.

7.1.8 Processor Sideband Signals

The processor include asynchronous sideband signals that provide asynchronous input, output or I/O signals between the processor and the platform or Platform Controller Hub. Details can be found in Table 7-5 and the applicable platform design guide.

All Processor Asynchronous Sideband input signals are required to be asserted/deasserted for a defined number of BCLKs in order for the processor to recognize the proper signal state. Refer to Section 7.9 for applicable signal integrity specifications.

7.1.9 Power, Ground and Sense Signals

Processors also include various other signals including power/ground and sense points. Details can be found in Table 7-5 and the applicable platform design guide.

7.1.9.1 Power and Ground Lands

All V_{CC}, V_{CCPLL}, V_{SA}, V_{CCD}, V_{TTA}, and V_{TTD} lands must be connected to their respective processor power planes, while all V_{SS} lands must be connected to the system ground plane. Refer to the applicable platform design guide for decoupling, voltage plane and routing guidelines for each power supply voltage.

For clean on-chip power distribution, processors include lands for all required voltage supplies. These are listed in Table 7-1.

Table 7-1. Power and Ground Lands

Power and Ground Lands	Number of Lands	Comments
V_{CC}	135	Each V_{CC} land must be supplied with the voltage determined by the SVID Bus signals. Table 7-3 Defines the voltage level associated with each core SVID pattern.Table 7-12, Figure 7-2represent V_{CC} static and transient limits. VCC has a VBOOT setting of 0.0V.
V_{CCPLL}	2	Each V_{CCPLL} land is connected to a 1.0 V supply, power the Phase Lock Loop (PLL) clock generation circuitry. An on-die PLL filter solution is implemented within the processor.
V_{CCD}	16	Each V_{CCD} land is connected to a switchable 1.50 V and 1.35 V supply, provide power to the processor DDR3 interface. These supplies also power the DDR3 memory subsystem. V_{CCD} is also controlled by the SVID Bus.
V_{TTA}	9	V_{TTA} lands must be supplied by a fixed 1.0V supply.
V_{TTD}	18	V_{TTD} lands must be supplied by a fixed 1.0V supply.
V_{SA}	23	Each V_{SA} land must be supplied with the voltage determined by the SVID Bus signals, typically set at 0.965V. VSA has a VBOOT setting of 0.9V.
V_{SS}	353	Ground

7.1.9.2 Decoupling Guidelines

Due to its large number of transistors and high internal clock speeds, the processor is capable of generating large current swings between low and full power states. This may cause voltages on power planes to sag below their minimum values if bulk decoupling is not adequate. Large electrolytic bulk capacitors (C_{BULK}), help maintain the output voltage during current transients, for example coming out of an idle condition. Care must be taken in the baseboard design to ensure that the voltages provided to the processor remain within the specifications listed in Table 7-11. Failure to do so can result in timing violations or reduced lifetime of the processor. For further information, refer to the appropriate Platform Design Guide (PDG).

7.1.9.3 Voltage Identification (VID)

The Voltage Identification (VID) specification for the $V_{CC,}$ V_{SA}, V_{CCD} voltage are defined by the *VR12/IMVP7 Pulse Width Modulation Specification*. The reference voltage or the VID setting is set via the SVID communication bus between the processor and the voltage regulator controller chip. The VID settings are the nominal voltages to be delivered to the processor's $V_{CC,}$ V_{SA}, V_{CCD} lands. Table 7-3 specifies the reference voltage level corresponding to the VID value transmitted over serial VID. The VID codes will change due to temperature and/or current load changes in order to minimize the power and to maximize the performance of the part. The specifications are set so that a voltage regulator can operate with all supported frequencies.

Individual processor VID values may be calibrated during manufacturing such that two processor units with the same core frequency may have different default VID settings.

The processor uses voltage identification signals to support automatic selection of $V_{CC,}$ V_{SA}, and V_{CCD} power supply voltages. If the processor socket is empty (SKTOCC_N high), or a “not supported” response is received from the SVID bus, then the voltage regulation circuit cannot supply the voltage that is requested, the voltage regulator must disable itself or not power on. Vout MAX register (30h) is programmed by the processor to set the maximum supported VID code and if the programmed VID code is higher than the VID supported by the VR, then VR will respond with a “not supported” acknowledgement. See the *VR12/IMVP7 Pulse Width Modulation Specification* for further details.

7.1.9.3.1 SVID Commands

The processor provides the ability to operate while transitioning to a new VID setting and its associated processor voltage rails ($V_{CC,}$ V_{SA}, and V_{CCD}). This is represented by a DC shift. It should be noted that a low-to-high or high-to-low voltage state change may result in as many VID transitions as necessary to reach the target voltage. Transitions above the maximum specified VID are not supported. The processor supports the following VR commands:

- SetVID_fast (20mV/µs for V_{CC},10mV/µs for V_{SA}/V_{CCD}),
- SetVID_slow (5mV/µs for V_{CC}, 2.5mV/µs for V_{SA}/V_{CCD}), and
- Slew Rate Decay (downward voltage only and it’s a function of the output capacitance’s time constant) commands. Table 7-3 and Table 7-20 includes SVID step sizes and DC shift ranges. Minimum and maximum voltages must be maintained as shown in Table 7-11.

The VRM or EVRD utilized must be capable of regulating its output to the value defined by the new VID. The *VR12/IMVP7 Pulse Width Modulation Specification* contains further details.

Power source characteristics must be guaranteed to be stable whenever the supply to the voltage regulator is stable.

7.1.9.3.2 SetVID Fast Command

The SetVID-fast command contains the target VID in the payload byte. The range of voltage is defined in the VID table. The VR should ramp to the new VID setting with a fast slew rate as defined in the slew rate data register. Typically 10 to 20 mV/µs depending on platform, voltage rail, and the amount of decoupling capacitance.

The SetVID-fast command is preemptive, the VR interrupts its current processes and moves to the new VID. The SetVID-fast command operates on 1 VR address at a time. This command is used in the processor for package C6 fast exit and entry.

7.1.9.3.3 SetVID Slow Command

The SetVID-slow command contains the target VID in the payload byte. The range of voltage is defined in the VID table. The VR should ramp to the new VID setting with a “slow” slew rate as defined in the slow slew rate data register. The SetVID_Slow is 1/4 slower than the SetVID_fast slew rate.

The SetVID-slow command is preemptive, the VR interrupts its current processes and moves to the new VID. This is the instruction used for normal P-state voltage change. This command is used in the processor for the Intel Enhanced SpeedStep Technology transitions.

7.1.9.3.4 SetVID-Decay Command

The SetVID-Decay command is the slowest of the DVID transitions. It is only used for VID down transitions. The VR does not control the slew rate, the output voltage declines with the output load current only.

The SetVID- Decay command is preemptive, that is, the VR interrupts its current processes and moves to the new VID.

7.1.9.3.5 SVID Power State Functions: SetPS

The processor has three power state functions and these will be set seamlessly via the SVID bus using the SetPS command. Based on the power state command, the SetPS commands sends information to VR controller to configure the VR to improve efficiency, especially at light loads. For example, typical power states are:

- PS0(00h): Represents full power or active mode
- PS1(01h): Represents a light load 5 A to 20 A
- PS2(02h): Represents a very light load <5 A

The VR may change its configuration to meet the processor's power needs with greater efficiency. For example, it may reduce the number of active phases, transition from CCM (Continuous Conduction Mode) to DCM (Discontinuous Conduction Mode) mode, reduce the switching frequency or pulse skip, or change to asynchronous regulation. For example, typical power states are 00h = run in normal mode; a command of 01h= shed phases mode, and an 02h=pulse skip.

The VR may reduce the number of active phases from PS0 to PS1 or PS0 to PS2 for example. There are multiple VR design schemes that can be used to maintain a greater efficiency in these different power states, please work with your VR controller suppliers for optimizations.

The SetPS command sends a byte that is encoded as to what power state the VR should transition to.

If a power state is not supported by the controller, the slave should acknowledge with command rejected (11b)

Note the mapping of power states 0-n will be detailed in the *VR12/IMVP7 Pulse Width Modulation Specification*.

If the VR is in a low power state and receives a SetVID command moving the VID up then the VR exits the low power state to normal mode (PS0) to move the voltage up as fast as possible. The processor must re-issue low power state (PS1 or PS2) command if it is in a low current condition at the new higher voltage. See Figure 7-2 for VR power state transitions.

Figure 7-2. VR Power-State Transitions

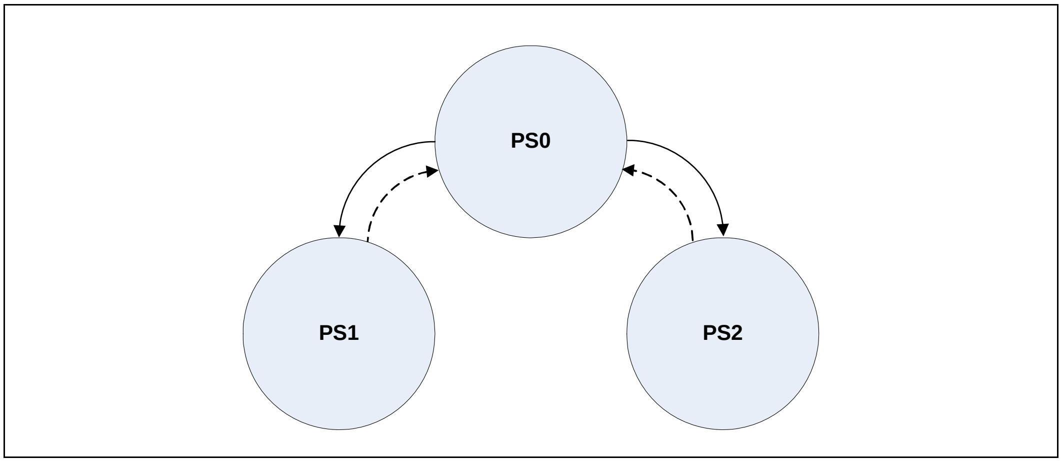

7.1.9.3.6 SVID Voltage Rail Addressing

The processor addresses 4 different voltage rail control segments within VR12 (VCC, VCCD, and VSA). The SVID data packet contains a 4-bit addressing code:

Table 7-2. SVID Address Usage

PWM Address (HEX)	Processor
00	V_{cc}
01	V_{sa}
02	V_{CCD}
03	N/A
04	N/A
05	N/A

Notes:
1. Check with VR vendors for determining the physical address assignment method for their controllers.
2. VR addressing is assigned on a per voltage rail basis.
3. Dual VR controllers will have two addresses with the lowest order address, always being the higher phase count.
4. For future platform flexibility, the VR controller should include an address offset, as shown with +1 not used.

Table 7-3. VR12.0 Reference Code Voltage Identification (VID) Table (Sheet 1 of 2)

HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD
00	0.00000	55	0.67000	78	0.84500	9B	1.02000	BE	1.19500	E1	1.37000
33	0.50000	56	0.67500	79	0.85000	9C	1.02500	BF	1.20000	E2	1.37500
34	0.50500	57	0.68000	7A	0.85500	9D	1.03000	C0	1.20500	E3	1.38000
35	0.51000	58	0.68500	7B	0.86000	9E	1.03500	C1	1.21000	E4	1.38500
36	0.51500	59	0.69000	7C	0.86500	9F	1.04000	C2	1.21500	E5	1.39000
37	0.52000	5A	0.69500	7D	0.87000	A0	1.04500	C3	1.22000	E6	1.39500
38	0.52500	5B	0.70000	7E	0.87500	A1	1.05000	C4	1.22500	E7	1.40000
39	0.53000	5C	0.70500	7F	0.88000	A2	1.05500	C5	1.23000	E8	1.40500
3A	0.53500	5D	0.71000	80	0.88500	A3	1.06000	C6	1.23500	E9	1.41000
3B	0.54000	5E	0.71500	81	0.89000	A4	1.06500	C7	1.24000	EA	1.41500
3C	0.54500	5F	0.72000	82	0.89500	A5	1.07000	C8	1.24500	EB	1.42000
3D	0.55000	60	0.72500	83	0.90000	A6	1.07500	C9	1.25000	EC	1.42500
3E	0.55500	61	0.73000	84	0.90500	A7	1.08000	CA	1.25500	ED	1.43000
3F	0.56000	62	0.73500	85	0.91000	A8	1.08500	CB	1.26000	EE	1.43500
40	0.56500	63	0.74000	86	0.91500	A9	1.09000	CC	1.26500	EF	1.44000
41	0.57000	64	0.74500	87	0.92000	AA	1.09500	CD	1.27000	F0	1.44500
42	0.57500	65	0.75000	88	0.92500	AB	1.10000	CE	1.27500	F1	1.45000
43	0.58000	66	0.75500	89	0.93000	AC	1.10500	CF	1.28000	F2	1.45500
44	0.58500	67	0.76000	8A	0.93500	AD	1.11000	D0	1.28500	F3	1.46000
45	0.59000	68	0.76500	8B	0.94000	AE	1.11500	D1	1.29000	F4	1.46500
46	0.59500	69	0.77000	8C	0.94500	AF	1.12000	D2	1.29500	F5	1.47000
47	0.60000	6A	0.77500	8D	0.95000	B0	1.12500	D3	1.30000	F6	1.47500
48	0.60500	6B	0.78000	8E	0.95500	B1	1.13000	D4	1.30500	F7	1.48000
49	0.61000	6C	0.78500	8F	0.96000	B2	1.13500	D5	1.31000	F8	1.48500

Table 7-3. VR12.0 Reference Code Voltage Identification (VID) Table (Sheet 2 of 2)

HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD	HEX	VCC, VSA, VCCD
4A	0.61500	6D	0.79000	90	0.96500	B3	1.14000	D6	1.31500	F9	1.49000
4B	0.62000	6E	0.79500	91	0.97000	B4	1.14500	D7	1.32000	FA	1.49500
4C	0.62500	6F	0.80000	92	0.97500	B5	1.15000	D8	1.32500	FB	1.50000
4D	0.63000	70	0.80500	93	0.98000	B6	1.15500	D9	1.33000	FC	1.50500
4E	0.63500	71	0.81000	94	0.98500	B7	1.16000	DA	1.33500	FD	1.51000
4F	0.64000	72	0.81500	95	0.99000	B8	1.16500	DB	1.34000	FE	1.51500
50	0.64500	73	0.82000	96	0.99500	B9	1.17000	DC	1.34500	FF	1.52000
51	0.65000	74	0.82500	97	1.00000	BA	1.17500	DD	1.35000		
52	0.65500	75	0.83000	98	1.00500	BB	1.18000	DE	1.35500		
53	0.66000	76	0.83500	99	1.01000	BC	1.18500	DF	1.36000		
54	0.66500	77	0.84000	9A	1.01500	BD	1.19000	E0	1.36500		

Notes:
1. 00h = Off State
2. VID Range HEX 01-32 are not used by the processor.
3. For VID Ranges supported see Table 7-12.
4. VCCD is a fixed voltage of 1.35 V or 1.5 V.

7.1.10 Reserved or Unused Signals

All Reserved (RSVD) signals must not be connected. Connection of these signals to V_{CC}, V_{TTA}, V_{TTD}, V_{CCD}, V_{CCPLL}, V_{SS}, or to any other signal (including each other) can result in component malfunction or incompatibility with future processors. See Chapter 8, "Processor Land Listing," for a land listing of the processor and the location of all Reserved signals.

For reliable operation, always connect unused inputs or bi-directional signals to an appropriate signal level. Unused active high inputs should be connected through a resistor to ground (V_{SS}). Unused outputs maybe left unconnected; however, this may interfere with some Test Access Port (TAP) functions, complicate debug probing, and prevent boundary scan testing. A resistor must be used when tying bi-directional signals to power or ground. When tying any signal to power or ground, a resistor will also allow for system testability. Resistor values should be within ± 20% of the impedance of the baseboard trace, unless otherwise noted in the appropriate platform design guidelines.

7.2 Signal Group Summary

Signals are grouped by buffer type and similar characteristics as listed in Table 7-5. The buffer type indicates which signaling technology and specifications apply to the signals.

Table 7-4. Signal Description Buffer Types (Sheet 1 of 2)

Signal	Description
Analog	Analog reference or output. May be used as a threshold voltage or for buffer compensation
Asynchronous[1]	Signal has no timing relationship with any system reference clock.
CMOS	CMOS buffers: 1.0 V or 1.5 V tolerant
DDR3	DDR3 buffers: 1.5 V and 1.35 V tolerant

Table 7-4. Signal Description Buffer Types (Sheet 2 of 2)

Signal	Description
DMI2	Direct Media Interface Gen 2 signals. These signals are compatible with PCI Express* 2.0 and 1.0 Signaling Environment AC Specifications.
Intel QPI	Current-mode 6.4 GT/s and 8.0 GT/s forwarded-clock Intel QuickPath Interconnect signaling
Open Drain CMOS	Open Drain CMOS (ODCMOS) buffers: 1.0V tolerant
PCI Express*	PCI Express* interface signals. These signals are compatible with PCI Express* 3.0 Signalling Environment AC Specifications and are AC coupled. The buffers are not 3.3-V tolerant. Refer to the PCIe* specification.
Reference	Voltage reference signal.
SSTL	Source Series Terminated Logic (JEDEC SSTL_15)

Notes:
1. Qualifier for a buffer type.

Table 7-5. Signal Groups (Sheet 1 of 3)

Differential/Single Ended	Buffer Type	Signals[1]
DDR3 Reference Clocks[2]		
Differential	SSTL Output	DDR{1/2/3}_CLK_D[N/P][3:0]
DDR3 Command Signals[2]		
Single ended	SSTL Output	DDR{1/2/3}_BA[2:0] DDR{1/2/3}_CAS_N DDR{1/2/3}_MA[15:00] DDR{1/2/3}_MA_PAR DDR{1/2/3}_RAS_N DDR{1/2/3}_WE_N
	CMOS1.5v Output	DDR_RESET_C{1/23}_N
DDR3 Control Signals[2]		
Single ended	CMOS1.5v Output	DDR{1/2/3}_CS_N[7:0] DDR{1/2/3}_ODT[3:0] DDR{1/2/3}_CKE[3:0]
	Reference Output	DDR_VREFDQTX_C{1/23}
	Reference Input	DDR_VREFDQRX_C{1/23} DDR{1/23}_RCOMP[2:0]
DDR3 Data Signals[2]		
Differential	SSTL Input/Output	DDR{1/2/3}_DQS_D[N/P][17:00]
Single ended	SSTL Input/Output	DDR{1/2/3}_DQ[63:00] DDR{1/2/3}_ECC[7:0]
	SSTL Input	DDR{1/2/3}_PAR_ERR_N
DDR3 Miscellaneous Signals[2]		
Single ended	CMOS1.5v Input	DRAM_PWR_OK_C{1/23}

Table 7-5. Signal Groups (Sheet 2 of 3)

Differential/Single Ended	Buffer Type	Signals[1]
PCI Express* Port 1 & 3Signals		
Differential	PCI Express* Input	PE1A_RX_D[N/P][3:0] PE1B_RX_D[N/P][7:4] PE3A_RX_D[N/P][3:0] PE3B_RX_D[N/P][7:4] PE3C_RX_D[N/P][11:8] PE3D_RX_D[N/P][15:12]
Differential	PCI Express* Output	PE1A_TX_D[N/P][3:0] PE1B_TX_D[N/P][7:4] PE3A_TX_D[N/P][3:0] PE3B_TX_D[N/P][7:4] PE3C_TX_D[N/P][11:8] PE3D_TX_D[N/P][15:12]
PCI Express* Miscellaneous Signals		
Single ended	Analog Input	PE_RBIAS_SENSE
	Reference Input/Output	PE_RBIAS PE_VREF_CAP
DMI2/PCI Express* Signals		
Differential	DMI2 Input	DMI_RX_D[N/P][3:0]
	DMI2 Output	DMI_TX_D[N/P][3:0]
Intel® QuickPath Interconnect (QPI) Signals		
Differential	Intel QPI Input	QPI1_DRX_D[N/P][19:00] QPI1_CLKRX_D[N/P]
	Intel QPI Output	QPI1_DTX_D[N/P][19:00] QPI1_CLKTX_D[N/P]
Single ended	Analog Input	QPI_RBIAS_SENSE
	Analog Input/Output	QPI_RBIAS
Platform Environmental Control Interface (PECI)		
Single ended	PECI	PECI
System Reference Clock (BCLK{0/1})		
Differential	CMOS1.0v Input	BCLK{0/1}_D[N/P]
SMBus		
Single ended	Open Drain CMOS Input/ Output	DDR_SCL_C{1/23} DDR_SDA_C{1/23} PEHPSCL PEHPSDA
JTAG & TAP Signals		
Single ended	CMOS1.0V Input	TCK, TDI, TMS, TRST_N
	CMOS1.0V Input/Output	PREQ_N
	CMOS1.0V Output	PRDY_N
	Open Drain CMOS Input/ Output	BPM_N[7:0] EAR_N
	Open Drain CMOS Output	TDO
Serial VID Interface (SVID) Signals		

Table 7-5. Signal Groups (Sheet 3 of 3)

Differential/Single Ended	Buffer Type	Signals[1]
Single ended	CMOS1.0v Input	SVIDALERT_N
	Open Drain CMOS Input/ Output	SVIDDATA
	Open Drain CMOS Output	SVIDCLK
Processor Asynchronous Sideband Signals		
Single ended	CMOS1.0v Input	BIST_ENABLE BMCINIT FRMAGENT PWRGOOD PMSYNC RESET_N SAFE_MODE_BOOT SOCKET_ID[1:0] TXT_AGENT TXT_PLTEN
	Open Drain CMOS Input/ Output	CAT_ERR_N CPU_ONLY_RESET MEM_HOT_C{1/23}_N PROCHOT_N
	Open Drain CMOS Output	ERROR_N[2:0] THERMTRIP_N
Miscellaneous Signals		
N/A	Output	IVT_ID_N SKTOCC_N
Power/Other Signals		
	Power / Ground	V_{CC}, V_{TTA}, V_{TTD}, V_{CCD}, V_{CCPLL}, V_{SA} and V_{SS}
	Sense Points	VCC_SENSE VSS_VCC_SENSE VSS_VTTD_SENSE VTTD_SENSEVSA_SENSE VSS_VSA_SENSE

Notes:

1. Refer to Section 6, "Signal Descriptions" for signal description details.
2. DDR{1/2/3} refers to DDR3 Channel 1, DDR3 Channel 2 and DDR3 Channel 3.

Table 7-6. Signals with On-Die Termination (Sheet 1 of 2)

Signal Name	Pull Up /Pull Down	Rail	Value	Units	Notes
DDR{1}_PAR_ERR_N	Pull Up	VCCD	65	Ω	
DDR{2/3}_PAR_ERR_N	Pull Up	VCCD	65	Ω	
BMCINIT	Pull Down	VSS	2K	Ω	1
FRMAGENT	Pull Down	VSS	2K	Ω	1
TXT_AGENT	Pull Down	VSS	2K	Ω	1
SAFE_MODE_BOOT	Pull Down	VSS	2K	Ω	1
SOCKET_ID[1:0]	Pull Down	VSS	2K	Ω	1
BIST_ENABLE	Pull Up	VTT	2K	Ω	1
TXT_PLTEN	Pull Up	VTT	2K	Ω	1

Table 7-6. Signals with On-Die Termination (Sheet 2 of 2)

Signal Name	Pull Up / Pull Down	Rail	Value	Units	Notes
EAR_N	Pull Up	VTT	2K	Ω	2

Notes:
1. Please refer to the applicable platform design guide to change the default states of these signals.
2. Refer to Table 7-19 for details on the R_{ON} (Buffer on Resistance) value for this signal.

7.3 Power-On Configuration (POC) Options

Several configuration options can be configured by hardware. The processor samples its hardware configuration at reset, on the active-to-inactive transition of RESET_N, or upon assertion of PWRGOOD (inactive-to-active transition). For specifics on these options, please refer to Table 7-7.

The sampled information configures the processor for subsequent operation. These configuration options cannot be changed except by another reset transition of the latching signal (RESET_N or PWRGOOD).

Table 7-7. Power-On Configuration Option Lands

Configuration Option	Land Name	Notes
Output tri state	PROCHOT_N	1
Execute BIST (Built-In Self Test)	BIST_ENABLE	2
Enable Service Processor Boot Mode	BMCINIT	3
Enable Intel TXT Platform	TXT_PLTEN	3
Power-up Sequence Halt for ITP configuration	EAR_N	3
Enable Bootable Firmware Agent	FRMAGENT	3
Enable Intel TXT Agent	TXT_AGENT	3
Enable Safe Mode Boot	SAFE_MODE_BOOT	3
Configure Socket ID	SOCKET_ID[1:0]	3

Notes:
1. Output tri-state option enables Fault Resilient Booting (FRB), for FRB details see Section 7.4. The signal used to latch PROCHOT_N for enabling FRB mode is RESET_N.
2. BIST_ENABLE is sampled at RESET_N de-assertion and CPU_ONLY_RESET de-assertion (on the falling edge).
3. This signal is sampled after PWRGOOD assertion.

7.4 Fault Resilient Booting (FRB)

The processor supports both socket and core level Fault Resilient Booting (FRB), which provides the ability to boot the system as long as there is one processor functional in the system. One limitation to socket level FRB is that the system cannot boot if the legacy socket that connects to an active PCH becomes unavailable since this is the path to the system BIOS. See Table 7-8 for a list of output tri-state FRB signals.

Socket level FRB will tri-state processor outputs via the PROCHOT_N signal. Assertion of the PROCHOT_N signal through RESET_N de-assertion will tri-state processor outputs. Note, that individual core disabling is also supported for those cases where disabling the entire package is not desired.

The processor extends the FRB capability to the core granularity by maintaining a register in the uncore so that BIOS or another entity can disable one or more specific processor cores.

Table 7-8. Fault Resilient Booting (Output Tri-State) Signals

Output Tri-State Signal Groups	Signals
Intel QPI	QPI0_CLKTX_DN[1:0] QPI0_CLKTX_DP[1:0] QPI0_DTX_DN[19:00] QPI0_DTX_DP[19:00] QPI1_CLKTX_DN[1:0] QPI1_CLKTX_DP[1:0] QPI1_DTX_DN[19:00] QPI1_DTX_DP[19:00]
SMBus	DDR_SCL_C1 DDR_SDA_C1 DDR_SCL_C23 DDR_SDA_C23 PEHPSCL PEHPSDA
JTAG & TAP	TDO
Processor Sideband	CAT_ERR_N ERROR_N[2:0] BPM_N[7:0] PRDY_N THERMTRIP_N PROCHOT_N PECI
SVID	SVIDCLK

7.5 Mixing Processors

Intel supports and validates two processor configurations only in which all processors operate with the same Intel QuickPath Interconnect frequency, core frequency, power segment, and have the same internal cache sizes. Mixing components operating at different internal clock frequencies is not supported and will not be validated by Intel. Combining processors from different power segments is also not supported.

Note: Processors within a system must operate at the same frequency per bits [15:8] of the FLEX_RATIO MSR (Address: 194h); however this does not apply to frequency transitions initiated due to thermal events, Extended HALT, Enhanced Intel SpeedStep Technology transitions signal. Please refer to the *Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM) Volumes 1, 2, and 3* for details on the FLEX_RATIO MSR and setting the processor core frequency.

Not all operating systems can support dual processors with mixed frequencies. Mixing processors of different steppings but the same model (as per CPUID instruction) is supported provided there is no more than one stepping delta between the processors, for example, S and S+1.

S and S+1 is defined as mixing of two CPU steppings in the same platform where one CPU is S (stepping) = CPUID.(EAX=01h):EAX[3:0], and the other is S+1 = CPUID.(EAX=01h):EAX[3:0]+1. The stepping ID is found in EAX[3:0] after executing the CPUID instruction with Function 01h.

Details regarding the CPUID instruction are provided in the *AP-485, Intel® Processor Identification and the CPUID Instruction* application note. Also refer to the *Intel® Xeon® Processor E5 Product Family Specification Update*

7.6 Flexible Motherboard Guidelines (FMB)

The Flexible Motherboard (FMB) guidelines are estimates of the maximum values the processor will have over certain time periods. The values are only estimates and actual specifications for future processors may differ. Processors may or may not have specifications equal to the FMB value in the foreseeable future. System designers should meet the FMB values to ensure their systems will be compatible with future processors.

7.7 Absolute Maximum and Minimum Ratings

Table 7-9 specifies absolute maximum and minimum ratings. At conditions outside functional operation condition limits, but within absolute maximum and minimum ratings, neither functionality nor long-term reliability can be expected. If a device is returned to conditions within functional operation limits after having been subjected to conditions outside these limits, but within the absolute maximum and minimum ratings, the device may be functional, but with its lifetime degraded depending on exposure to conditions exceeding the functional operation condition limits.

Although the processor contains protective circuitry to resist damage from Electro-Static Discharge (ESD), precautions should always be taken to avoid high static voltages or electric fields.

Table 7-9. Processor Absolute Minimum and Maximum Ratings

Symbol	Parameter	Min	Max	Unit
V_{CC}	Processor core voltage with respect to Vss	-0.3	1.4	V
V_{CCPLL}	Processor PLL voltage with respect to Vss	-0.3	2.0	V
V_{CCD}	Processor IO supply voltage for DDR3 (standard voltage) with respect to V_{SS}	-0.3	1.85	V
V_{CCD}	Processor IO supply voltage for DDR3L (low Voltage) with respect to V_{SS}	-0.3	1.7	V
V_{SA}	Processor SA voltage with respect to V_{SS}	-0.3	1.4	V
V_{TTA} V_{TTD}	Processor analog IO voltage with respect to V_{SS}	-0.3	1.4	V

Notes:
1. For functional operation, all processor electrical, signal quality, mechanical, and thermal specifications must be satisfied.
2. Overshoot and undershoot voltage guidelines for input, output, and I/O signals are outlined in Section 7.9.5. Excessive overshoot or undershoot on any signal will likely result in permanent damage to the processor.

7.7.1 Storage Conditions Specifications

Environmental storage condition limits define the temperature and relative humidity limits to which the device is exposed to while being stored in a Moisture Barrier Bag. The specified storage conditions are for component level prior to board attach (see notes in Table 7-10 for post board attach limits).

Table 7-10 specifies absolute maximum and minimum storage temperature limits which represent the maximum or minimum device condition beyond which damage, latent or otherwise, may occur. The table also specifies sustained storage temperature, relative humidity, and time-duration limits. These limits specify the maximum or minimum device storage conditions for a sustained period of time. At conditions outside sustained limits, but within absolute maximum and minimum ratings, quality & reliability may be affected.

Table 7-10. Storage Condition Ratings

Symbol	Parameter	Min	Max	Unit
$T_{absolute\ storage}$	The minimum/maximum device storage temperature beyond which damage (latent or otherwise) may occur when subjected to for any length of time.	-25	125	°C
$T_{sustained\ storage}$	The minimum/maximum device storage temperature for a sustained period of time.	-5	40	°C
$T_{short\ term\ storage}$	The ambient storage temperature (in shipping media) for a short period of time.	-20	85	°C
$RH_{sustained\ storage}$	The maximum device storage relative humidity for a sustained period of time.	60% @ 24		°C
$Time_{sustained\ storage}$	A prolonged or extended period of time; typically associated with sustained storage conditions Unopened bag, includes 6 months storage time by customer.	0	30	months
$Time_{short\ term\ storage}$	A short period of time (in shipping media).	0	72	hours

Notes:
1. Storage conditions are applicable to storage environments only. In this scenario, the processor must not receive a clock, and no lands can be connected to a voltage bias. Storage within these limits will not affect the long-term reliability of the device. For functional operation, please refer to the processor case temperature specifications.
2. These ratings apply to the Intel component and do not include the tray or packaging.
3. Failure to adhere to this specification can affect the long-term reliability of the processor.
4. Non-operating storage limits post board attach: Storage condition limits for the component once attached to the application board are not specified. Intel does not conduct component level certification assessments post board attach given the multitude of attach methods, socket types and board types used by customers. Provided as general guidance only, Intel board products are specified and certified to meet the following temperature and humidity limits (Non-Operating Temperature Limit: -40°C to 70°C & Humidity: 50% to 90%, non condensing with a maximum wet bulb of 28°C).
5. Device storage temperature qualification methods follow JEDEC High and Low Temperature Storage Life Standards: *JESD22-A119* (low temperature) and *JESD22-A103* (high temperature).

7.8 DC Specifications

DC specifications are defined at the processor pads, unless otherwise noted. DC specifications are only valid while meeting specifications for case temperature (T_{CASE} specified in Section 5), clock frequency, and input voltages. Care should be taken to read all notes associated with each specification.

7.8.1 Voltage and Current Specifications

Table 7-11. Voltage Specification

Symbol	Parameter	Voltage Plane	Min	Typ	Max	Unit	Notes[1]
V_{CC} VID	V_{CC} VID Range		0.6		1.35	V	2, 3
V_{CC}	Core Voltage (Launch - FMB)	V_{CC}	See Table 7-13 and Figure 7-3			V	3, 4, 7, 8, 12, 14, 18
V_{VID_STEP} (Vcc, Vsa, Vccd)	VID step size during a transition			5.0		mV	10
V_{CCPLL}	PLL Voltage	V_{CCPLL}	$0.955{*}V_{CCPLL\_TYP}$	1.8	$1.045{*}V_{CCPLL\_TYP}$	V	11, 12, 13, 17
V_{CCD}	I/O Voltage for DDR3 (Standard Voltage)	V_{CCD}	$0.95{*}V_{CCD\_TYP}$	1.5	$1.05{*}V_{CCD\_TYP}$	V	11, 13, 14, 16, 17
V_{CCD}	I/O Voltage for DDR3L (Low Voltage)	V_{CCD}	$0.95{*}V_{CCD\_TYP}$	1.35	$1.075{*}V_{CCD\_TYP}$	V	11, 13, 14, 16, 17
V_{TT} (V_{TTA}, V_{TTD})	Uncore Voltage (Launch - FMB)	V_{TT}	$0.957{*}V_{TT\_TYP}$	1.05	$1.043{*}V_{TT\_TYP}$	V	3, 5, 9, 12, 13
V_{SA_VID}	V_{SA} VID Range	V_{SA}	0.6	0.965	1.20	V	2, 3, 14, 15
V_{SA}	System Agent Voltage (Launch - FMB)	V_{SA}	V_{SA_VID} - 0.064	V_{SA_VID}	V_{SA_VID} + 0.064	V	3, 6, 12, 14, 19

Notes:

1. Unless otherwise noted, all specifications in this table apply to all processors. These specifications are based on final silicon characterization.
2. Individual processor VID values may be calibrated during manufacturing such that two devices at the same speed may have different settings.
3. Voltages are targets only. A variable voltage source should exist on systems in the event that a different voltage is required.
4. The V_{CC} voltage specification requirements are measured across the remote sense pin pairs (VCC_SENSE and VSS_VCC_SENSE) on the processor package. Voltage measurement should be taken with a DC to 100 MHz bandwidth oscilloscope limit (or DC to 20 MHz for older model oscilloscopes), using a 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of the ground wire on the probe should be less than 5 mm to ensure external noise from the system is not coupled in the scope probe.
5. The $V_{TTA,}$ and V_{TTD} voltage specification requirements are measured across the remote sense pin pairs (VTTD_SENSE and VSS_VTTD_SENSE) on the processor package. Voltage measurement should be taken with a DC to 100 MHz bandwidth oscilloscope limit (or DC to 20 MHz for older model oscilloscopes), using a 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of the ground wire on the probe should be less than 5 mm to ensure external noise from the system is not coupled in the scope probe.
6. The V_{SA} voltage specification requirements are measured across the remote sense pin pairs (VSA_SENSE and VSS_VSA_SENSE) on the processor package. Voltage measurement should be taken with a DC to 100 MHz bandwidth oscilloscope limit (or DC to 20 MHz for older model oscilloscopes), using a 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of the ground wire on the probe should be less than 5 mm to ensure external noise from the system is not coupled in the scope probe.
7. For the Intel® Xeon® processor E5-2400 product family processor refer to Table 7-13 and corresponding Figure 7-3. The processor should not be subjected to any static V_{CC} level that exceeds the V_{CC_MAX} associated with any particular current. Failure to adhere to this specification can shorten processor lifetime.
8. Minimum V_{CC} and maximum I_{CC} are specified at the maximum processor case temperature (T_{CASE}) shown in Chapter 5, "Thermal Management Specifications". I_{CC_MAX} is specified at the relative V_{CC_MAX} point on the V_{CC} load line. The processor is capable of drawing I_{CC_MAX} for up to 5 seconds. Refer to Figure 7-4 for further details on the average processor current draw over various time durations.
9. The processor should not be subjected to any static $V_{TTA,}$ V_{TTD} level that exceeds the V_{TT_MAX} associated with any particular current. Failure to adhere to this specification can shorten processor lifetime.
10. This specification represents the V_{CC} increase or decrease due to each VID transition, see Section 7.1.9.3, "Voltage Identification (VID)".
11. Baseboard bandwidth is limited to 20 MHz.
12. FMB is the flexible motherboard guidelines. See Section 7.4, "Fault Resilient Booting (FRB)" for FMB details.
13. DC + AC + Ripple specification
14. For Power State Functions see Section 7.1.9.3.5.
15. V_{SA_VID} does not have a loadline, the output voltage is expected to be the VID value.
16. V_{CCD} tolerance at processor pins. Tolerance for VR at remote sense is ±3.3%*V_{CCD}.

17. The V_{CCPLL}, V_{CCD} voltage specification requirements are measured across vias on the platform. Choose V_{CCPLL}, V_{CCD} vias close to the socket and measure with a DC to 100MHz bandwidth oscilloscope limit (or DC to 20MHz for older model oscilloscopes), using 1.5 pF maximum probe capacitance, and 1M Ω minimum impedance. The maximum length of the ground wire on the probe should be less than 5 mm to ensure external noise from the system is not coupled in the scope probe.
18. V_{CC} has a Vboot setting of 0.0V and is not included in the PWRGOOD indication. Refer to the *VR12/IMVP7 Pulse Width Modulation Specification*.
19. V_{SA} has a Vboot setting of 0.9V. Refer to the *VR12/IMVP7 Pulse Width Modulation Specification*.

Table 7-12. Processor Supply Current Specifications

Parameter and Definition	Processor TDP / Core count	TDC (A)	Max (A)	Notes[1]
I_{TT} I/O Termination Supply, Processor Current on V_{TTA}/V_{TTD}	All Intel® Xeon® Processor E5-2400 Product Families	16	20	2,3,6
I_{SA} System Agent Supply, Processor Current on V_{SA}		18	19	
I_{CCPLL} PLL Supply, Processor Current on V_{CCPLL}		2	2	
I_{CCD} Memory Controller DDR3 Supply, Processor Current on V_{CCD}		5	7	
I_{CCD_S3} Memory Controller Supply, Processor Current on V_{CCD} while in System S3 Standby State		--	1	4,5
I_{CC} Core Processor Supply, Processor Current on V_{CC}	95W 8-core / 6-core	106	135	2,3,6
	80W 4-core	80	85	
	80W 2-core 1S	70	80	
	70W 8-core and LV70W-8C	90	110	
	60W 6-core and LV60W-6C	75	90	
	50W 8-core	65	80	
	50W 4-core and LV50W-4C	65	80	
	LV40W-24C (2-core 1S)	40	50	

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processors. These specifications are based on silicon characterization.
2. Launch to FMB, See Section 7.6, "Flexible Motherboard Guidelines (FMB)" for details.
3. TDC (Thermal Design Current) is the sustained (DC equivalent) current that the processor is capable of drawing indefinitely and should be used for the voltage regulator thermal assessment. The voltage regulator is responsible for monitoring its temperature and asserting the necessary signal to inform the processor of a thermal excursion. Please refer to the *VR12/IMVP7 Pulse Width Modulation Specification* for further details.
4. Specification is at T_{CASE} = 50 °C. Characterized by design (not tested).
5. I_{CCD} specifications are current draw on V_{CCD} of processor only and do not include current consumption by memory devices.
6. Minimum VCC and maximum ICC are specified at the maximum processor case temperature (T_{CASE}) shown in Section 5, "Thermal Management Specifications". ICC_MAX is specified at the relative VCC_MAX point on the VCC load line. The processor is capable of drawing ICC_MAX for up to 5 seconds. Refer to Section 7-4, "Load Current Versus Time" for details on processor current draw over various durations.

Table 7-13. Processor VCC Static and Transient Tolerance (Sheet 1 of 2)

I_{CC} [A]	V_{CC_MAX} [V]	V_{CC_TYP} [V]	V_{CC_MIN} [V]	Notes
0	VID + 0.015	VID - 0.000	VID - 0.015	1,2,3,4
5	VID + 0.009	VID - 0.006	VID - 0.021	1,2,3,4
10	VID + 0.003	VID - 0.013	VID - 0.028	1,2,3,4
15	VID - 0.004	VID - 0.019	VID - 0.034	1,2,3,4
20	VID - 0.010	VID - 0.025	VID - 0.040	1,2,3,4

Table 7-13. Processor VCC Static and Transient Tolerance (Sheet 2 of 2)

I_{CC} [A]	V_{CC_MAX} [V]	V_{CC_TYP} [V]	V_{CC_MIN} [V]	Notes
25	VID - 0.016	VID - 0.031	VID - 0.046	1,2,3,4
30	VID - 0.023	VID - 0.038	VID - 0.053	1,2,3,4
35	VID - 0.029	VID - 0.044	VID - 0.059	1,2,3,4
40	VID - 0.035	VID - 0.050	VID - 0.065	1,2,3,4
45	VID - 0.041	VID - 0.056	VID - 0.071	1,2,3,4
50	VID - 0.048	VID - 0.063	VID - 0.078	1,2,3,4
55	VID - 0.054	VID - 0.069	VID - 0.084	1,2,3,4
60	VID - 0.060	VID - 0.075	VID - 0.090	1,2,3,4
65	VID - 0.066	VID - 0.081	VID - 0.096	1,2,3,4
70	VID - 0.073	VID - 0.088	VID - 0.103	1,2,3,4
75	VID - 0.079	VID - 0.094	VID - 0.109	1,2,3,4
80	VID - 0.085	VID - 0.100	VID - 0.115	1,2,3,4
85	VID - 0.091	VID - 0.106	VID - 0.121	1,2,3,4
90	VID - 0.098	VID - 0.113	VID - 0.128	1,2,3,4
95	VID - 0.104	VID - 0.119	VID - 0.134	1,2,3,4
100	VID - 0.110	VID - 0.125	VID - 0.140	1,2,3,4
105	VID - 0.116	VID - 0.131	VID - 0.146	1,2,3,4
110	VID - 0.123	VID - 0.138	VID - 0.153	1,2,3,4
115	VID - 0.129	VID - 0.144	VID - 0.159	1,2,3,4
120	VID - 0.135	VID - 0.150	VID - 0.165	1,2,3,4
125	VID - 0.141	VID - 0.156	VID - 0.171	1,2,3,4
130	VID - 0.148	VID - 0.163	VID - 0.178	1,2,3,4
135	VID - 0.154	VID - 0.169	VID - 0.184	1,2,3,4

Notes:
1. The loadline specification includes both static and transient limits.
2. This table is intended to aid in reading discrete points on graph in Figure 7-3.
3. The loadlines specify voltage limits at the die measured at the Vcc_sense and Vss_Vcc_sense lands. Voltage regulation feedback for voltage regulator circuits must also be taken from processor Vcc_sense and Vss_Vcc_sense lands. Refer to the *VR12/IMVP7 Pulse Width Modulation Specification* for loadline guidelines and VR implementation details.
4. The Icc ranges extend to IccMax of the target processor as follows:
   - 0-135 A for Intel® Xeon® processor E5-2400 product family Processor (95W)
   - 0-85 A for Intel® Xeon® processor E5-2400 product family Processor (80W)
   - 0-80 A for Intel® Xeon® Processor E5-1400 (80W, 1 Socket)
   - 0-110 A for Intel® Xeon® Processor E5-2400 (70W and LV70W-8C)
   - 0-90 A for Intel® Xeon® Processor E5-2400 (60W and LV60W-6C)
   - 0-80 A for Intel® Xeon® Processor E5-2400 (50W and LV50W-4C)
   - 0-50 A for LV40W-24C

Figure 7-3. V_{CC} Static and Transient Tolerance Loadlines

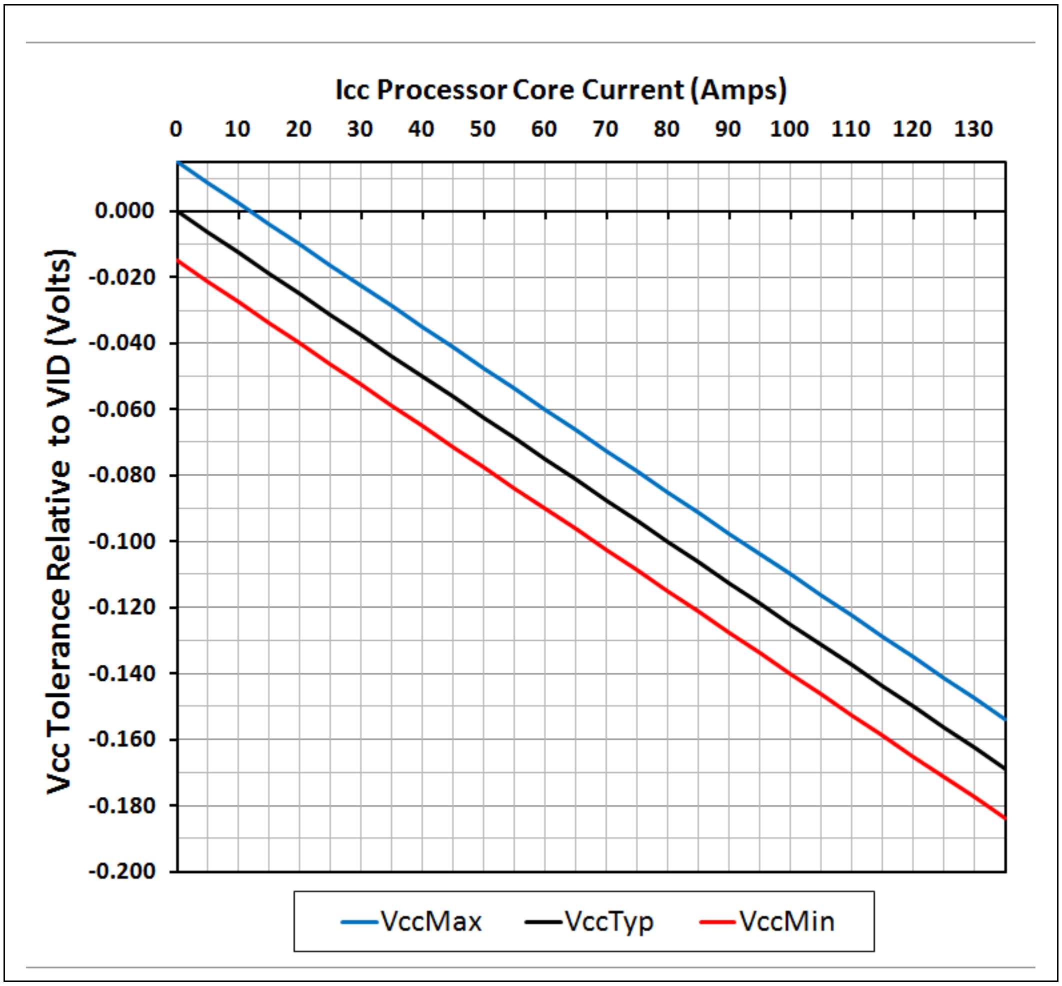

7.8.2 Die Voltage Validation

Core voltage (V_{CC}) overshoot events at the processor must meet the specifications in Table 7-14 when measured across the VCC_SENSE and VSS_VCC_SENSE lands. Overshoot events that are < 10 ns in duration may be ignored. These measurements of processor die level overshoot should be taken with a 100 MHz bandwidth limited oscilloscope.

Figure 7-4. Load Current Versus Time

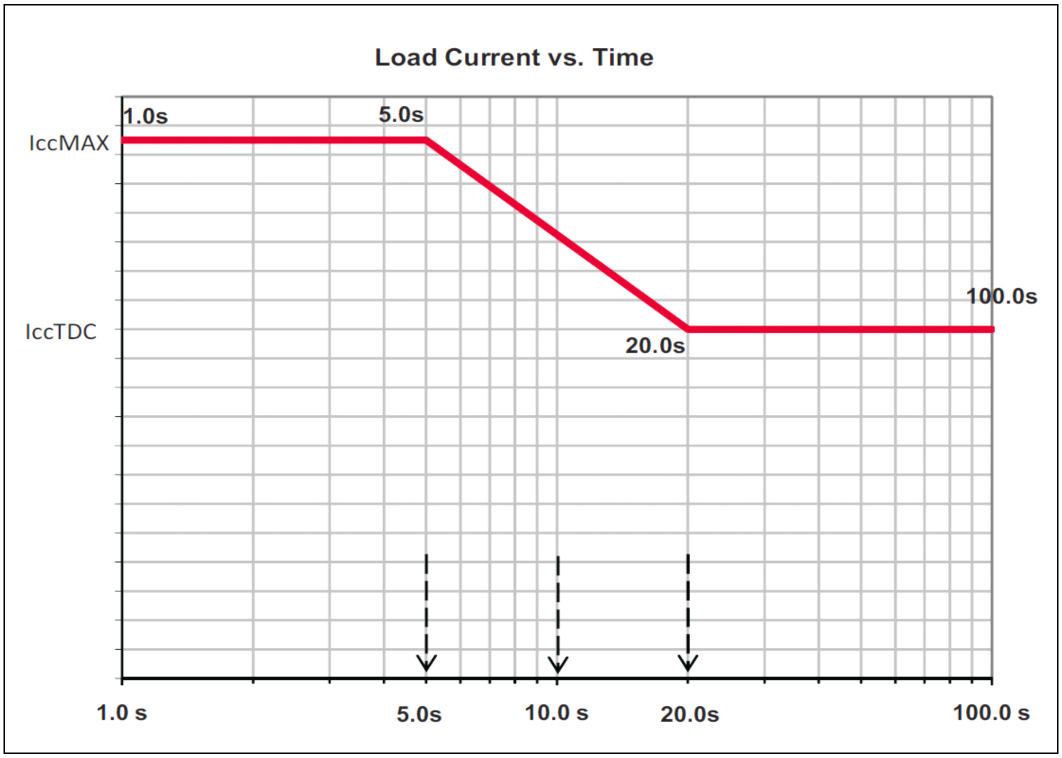

Notes:
1. The peak current for *any* 5 second sample does not exceed Icc_max.
2. The average current for *any* 10 second sample does not exceed the Y value at 10 seconds.
3. The average current for *any* 20 second period or greater does not exceed Icc_tdc.
4. Turbo performance may be impacted by failing to meet durations specified in this graph. Ensure that the platform design can handle peak and average current based on the specification.
5. Processor or voltage regulator thermal protection circuitry should not trip for load currents greater than I_{CC_TDC}.
6. Not 100% tested. Specified by design characterization.

7.8.2.1 V_{CC} Overshoot Specifications

The processor can tolerate short transient overshoot events where V_{CC} exceeds the VID voltage when transitioning from a high-to-low current load condition. This overshoot cannot exceed VID + V_{OS_MAX} (V_{OS_MAX} is the maximum allowable overshoot above VID). These specifications apply to the processor die voltage as measured across the VCC_SENSE and VSS_VCC_SENSE lands.

Table 7-14. V_{CC} Overshoot Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Max	Units	Figure	Notes
V_{OS_MAX}	Magnitude of V_{CC} overshoot above VID		65	mV	7-5	

Table 7-14. V_{CC} Overshoot Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Max	Units	Figure	Notes
T_{OS_MAX}	Time duration of V_{CC} overshoot above VccMAX value at the new lighter load		25	µs	7-5	

Figure 7-5. V_{CC} Overshoot Example Waveform

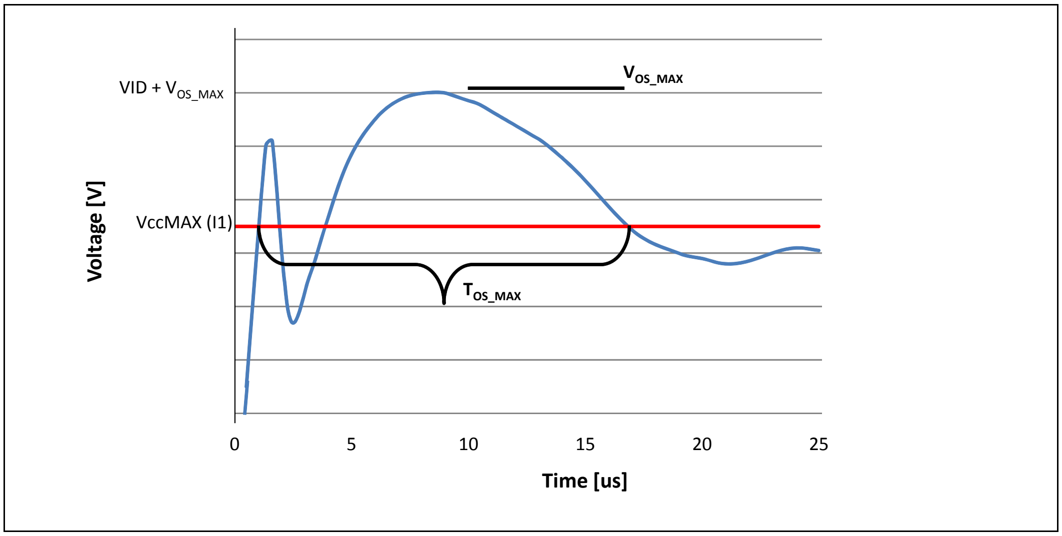

Notes:
1. V_{OS_MAX} is the measured overshoot voltage.
2. T_{OS_MAX} is the measured time duration above VccMAX(I1).
3. Istep: Load Release Current Step, for example, I2 to I1, where I2 > I1.
4. VccMAX(I1) = VID - I1*RLL + 15 mV

7.8.3 Signal DC Specifications

DC specifications are defined at the processor pads, unless otherwise noted. DC specifications are only valid while meeting specifications for case temperature (T_{CASE} specified in Section 5, “Thermal Management Specifications”), clock frequency, and input voltages. Care should be taken to read all notes associated with each specification.

Table 7-15. DDR3 and DDR3L Signal DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Typ	Max	Units	Notes[1]
I_{IL}	Input Leakage Current	-500		+500	uA	10
Data Signals						
V_{IL}	Input Low Voltage			$0.43*V_{CCD}$	V	2, 3
V_{IH}	Input High Voltage	$0.57*V_{CCD}$			V	2, 4, 5
R_{ON}	DDR3 Data Buffer On Resistance	21		31	Ω	6
Data ODT	On-Die Termination for Data Signals	45 90		55 110	Ω	8
PAR_ERR_N ODT	On-Die Termination for Parity Error Signals	59		72	Ω	
Reference Clock Signals, Command, and Data Signals						

Table 7-15. DDR3 and DDR3L Signal DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Typ	Max	Units	Notes[1]
V_{OL}	Output Low Voltage		$(V_{CCD} / 2)*(R_{ON} / (R_{ON}+R_{VTT\_TERM}))$		V	2, 7
V_{OH}	Output High Voltage		$V_{CCD} - ((V_{CCD} / 2)*(R_{ON}/(R_{ON}+R_{VTT\_TERM}))$		V	2, 5, 7
Reference Clock Signal						
R_{ON}	DDR3 Clock Buffer On Resistance	21		31	Ω	6
Command Signals						
R_{ON}	DDR3 Command Buffer On Resistance	16		24	Ω	6
R_{ON}	DDR3 Reset Buffer On Resistance	25		75	Ω	6
$V_{OL\_CMOS1.5v}$	Output Low Voltage, Signals DDR_RESET_ C{1/23}_N			$0.2*V_{CCD}$	V	1,2
$V_{OH\_CMOS1.5v}$	Output High Voltage, Signals DDR_RESET_ C{1/23}_N	$0.9*V_{CCD}$			V	1,2
$I_{IL\_CMOS1.5v}$	Input Leakage Current	-100		+100	µA	1,2
Control Signals						
R_{ON}	DDR3 Control Buffer On Resistance	21		31	Ω	6
DDR1_RCOMP[0]	COMP Resistance	128.7	130	131.3	Ω	9,12
DDR1_RCOMP[1]	COMP Resistance	25.839	26.1	26.361	Ω	9,12
DDR1_RCOMP[2]	COMP Resistance	198	200	202	Ω	9,12
DDR23_RCOMP[0]	COMP Resistance	128.7	130	131.3	Ω	9,12
DDR23_RCOMP[1]	COMP Resistance	25.839	26.1	26.361	Ω	9,12
DDR23_RCOMP[2]	COMP Resistance	198	200	202	Ω	9,12
DDR3 Miscellaneous Signals						
V_{IL}	Input Low Voltage DRAM_PWR_OK_C{01/23}			0.55*VCCD - 0.2	V	2, 3, 11, 13
V_{IH}	Input High Voltage DRAM_PWR_OK_C{01/23}	0.55*VCCD + 0.3			V	2, 4, 5, 11, 13

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. The voltage rail V_{CCD} which will be set to 1.50 V or 1.35 V nominal depending on the voltage of all DIMMs connected to the processor.
3. V_{IL} is the maximum voltage level at a receiving agent that will be interpreted as a logical low value.
4. V_{IH} is the minimum voltage level at a receiving agent that will be interpreted as a logical high value.
5. V_{IH} and V_{OH} may experience excursions above V_{CCD}. However, input signal drivers must comply with the signal quality specifications. Refer to Section 7.9.
6. This is the pull down driver resistance. Refer to processor signal integrity models for I/V characteristics. Reset drive does not have a termination.
7. R_{VTT_TERM} is the termination on the DIMM and not controlled by the processor. Please refer to the applicable DIMM datasheet.
8. The minimum and maximum values for these signals are programmable by BIOS to one of the pairs.
9. COMP resistance must be provided on the system board with 1% resistors. See the applicable platform design guide for implementation details. DDR1_RCOMP[2:0] and DDR23_RCOMP[2:0] resistors are terminated to VSS.
10. Input leakage current is specified for all DDR3 signals.
11. DRAM_PWR_OK_C{1/23} must have a maximum of 30 ns rise or fall time over VCCD * 0.55 +300 mV and -200 mV and the edge must be monotonic.
12. The DDR1/23_RCOMP error tolerance is ± 15% from the compensated value.

13. DRAM_PWR_OK_C{1/23}: Data Scrambling must be enabled for production environments. Disabling Data scrambling may be used for debug and testing purposes only. Operating systems with Data Scrambling off will make the configuration out of specification.

Table 7-16. PECI DC Specifications

Symbol	Definition and Conditions	Min	Max	Units	Figure	Notes[1]
V_{In}	Input Voltage Range	-0.150	V_{TT}	V		
$V_{Hysteresis}$	Hysteresis	0.100 * V_{TT}		V		
V_N	Negative-edge threshold voltage	0.275 * V_{TT}	0.500 * V_{TT}	V	7-1	2
V_P	Positive-edge threshold voltage	0.550 * V_{TT}	0.725 * V_{TT}	V	7-1	2
I_{SOURCE}	High level output source V_{OH} = 0.75 * V_{TT}	-6.0		mA		
I_{Leak+}	High impedance state leakage to V_{TTD} (V_{leak} = V_{OL})	50	200	µA		3
C_{Bus}	Bus capacitance per node	N/A	10	pF		4,5
V_{Noise}	Signal noise immunity above 300 MHz	0.100 * V_{TT}	N/A	V_{p-p}		

Notes:
1. V_{TTD} supplies the PECI interface. PECI behavior does not affect V_{TTD} min/max specification
2. It is expected that the PECI driver will take into account, the variance in the receiver input thresholds and consequently, be able to drive its output within safe limits (-0.150 V to 0.275*V_{TTD} for the low level and 0.725*V_{TTD} to V_{TTD}+0.150 V for the high level).
3. The leakage specification applies to powered devices on the PECI bus.
4. One node is counted for each client and one node for the system host. Extended trace lengths might appear as additional nodes.
5. Excessive capacitive loading on the PECI line may slow down the signal rise/fall times and consequently limit the maximum bit rate at which the interface can operate.

Table 7-17. System Reference Clock (BCLK{0/1}) DC Specifications

Symbol	Parameter	Signal	Min	Max	Unit	Figure	Notes[1]
$V_{BCLK\_diff\_ih}$	Differential Input High Voltage	Differential	0.150	N/A	V	7-7	
$V_{BCLK\_diff\_il}$	Differential Input Low Voltage	Differential		-0.150	V	7-7	
V_{cross} (abs)	Absolute Crossing Point	Single Ended	0.250	0.550	V	7-6 7-8	2, 4, 7
V_{cross}(rel)	Relative Crossing Point	Single Ended	0.250 + 0.5*(VH_{avg} - 0.700)	0.550 + 0.5*(VH_{avg} - 0.700)	V	7-6	3, 4, 5
ΔV_{cross}	Range of Crossing Points	Single Ended	N/A	0.140	V	7-9	6
V_{TH}	Threshold Voltage	Single Ended	Vcross - 0.1	Vcross + 0.1	V		
I_{IL}	Input Leakage Current	N/A		1.50	µA		8
C_{pad}	Pad Capacitance	N/A	0.9	1.	pF		

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies. These specifications are specified at the processor pad.
2. Crossing Voltage is defined as the instantaneous voltage value when the rising edge of BCLK{0/1}_DN is equal to the falling edge of BCLK{0/1}_DP.
3. V_{Havg} is the statistical average of the VH measured by the oscilloscope.
4. The crossing point must meet the absolute and relative crossing point specifications simultaneously.
5. V_{Havg} can be measured directly using “Vtop” on Agilent* and “High” on Tektronix oscilloscopes.
6. V_{CROSS} is defined as the total variation of all crossing voltages as defined in Note 3.
7. The rising edge of BCLK{0/1}_DN is equal to the falling edge of BCLK{0/1}_DP.
8. For Vin between 0 and Vih.

Table 7-18. SMBus DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Max	Units	Notes
V_{IL}	Input Low Voltage		0.3*V_{TT}	V	

Table 7-18. SMBus DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Max	Units	Notes
V_{IH}	Input High Voltage	0.7*VTT		V	
V_{OL}	Output Low Voltage		$0.2*V_{TT}$	V	
V_{OH}	Output High Voltage		$V_{TT(max)}$	V	
R_{ON}	Buffer On Resistance		14	Ω	
I_L	Leakage Current Signals DDR_SCL_C{1/23}, DDR_SDA_C{1/23}	-100	+100	µA	
I_L	Leakage Current Signals PEHPSCL, PEHPSDA		+900	µA	

Table 7-19. JTAG and TAP Signals DC Specifications

Symbol	Parameter	Min	Max	Units	Notes
V_{IL}	Input Low Voltage		$0.3*V_{TT}$	V	
V_{IH}	Input High Voltage	$0.7*V_{TT}$		V	
V_{OL}	Output Low Voltage (R_{TEST} = 500 ohm)		$0.12*V_{TT}$	V	
V_{OH}	Output High Voltage (R_{TEST} = 500 ohm)	$0.88*V_{TT}$		V	
R_{ON}	Buffer On Resistance Signals BPM_N[7:0], TDO, EAR_N		14	Ω	
I_{IL}	Input Leakage Current Signals PREQ_N, TCK, TDI, TMS, TRST_N	-50	+50	µA	
I_{IL}	Input Leakage Current Signals BPM_N[7:0], TDO, EAR_N (R_{TEST} = 50 ohm)		+900	µA	
I_O	Output Current Signal PRDY_N (R_{TEST} = 500 ohm)	-1.50	+1.50	mA	
	Input Edge Rate Signals: BPM_N[7:0], EAR_N, PREQ_N, TCK, TDI, TMS, TRST_N	0.05		V/ns	1

Note:
1. These signals are measured between VIL and VIH.
2. The signal edge rate must be met or the signal must transition monotonically to the asserted state.

Table 7-20. Serial VID Interface (SVID) DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Typ	Max	Units	Notes
V_{TT}	CPU I/O Voltage	VTT - 3%	1.0	VTT + 3%	V	
V_{IL}	Input Low Voltage Signals SVIDDATA, SVIDALERT_N			$0.3*V_{TT}$	V	1
V_{IH}	Input High Voltage Signals SVIDDATA, SVIDALERT_N	$0.7*V_{TT}$			V	1
V_{OH}	Output High Voltage Signals SVIDCLK, SVIDDATA			$V_{TT(max)}$	V	1

Table 7-20. Serial VID Interface (SVID) DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Typ	Max	Units	Notes
R_{ON}	Buffer On Resistance Signals SVIDCLK, SVIDDATA			14	Ω	2
I_{IL}	Input Leakage Current Signals SVIDCLK, SVIDDATA			±900	μA	3,4
I_{IL}	Input Leakage Current Signal SVIDALERT_N			±500	μA	3,4

Notes:
1. V_{TT} refers to instantaneous V_{TT}.
2. Measured at 0.31*V_{TT}
3. Vin between 0V and V_{TT}
4. Refer to the appropriate Platform Design Guide (PDG) for routing design guidelines.

Table 7-21. Processor Asynchronous Sideband DC Specifications

Symbol	Parameter	Min	Max	Units	Notes
	Input Edge Rate Signals: CAT_ERR_N, MEM_HOT_C{1/23}_N, PMSYNC, PROCHOT_N, PWRGOOD, RESET_N	0.05		V/ns	4,5
CMOS1.05v Signals					
$V_{IL\_CMOS1.05v}$	Input Low Voltage		0.3*V_{TT}	V	1,2
$V_{IH\_CMOS1.05v}$	Input High Voltage	0.7*V_{TT}		V	1,2
V_{IL_MAX}	Input Low Voltage Signal PWRGOOD		0.320	V	1,2,5,
V_{IH_MIN}	Input High Voltage Signal PWRGOOD	0.640		V	1,2,5
$V_{OL\_CMOS1.05v}$	Output Low Voltage		0.12*V_{TT}	V	1,2
$V_{OH\_CMOS1.05v}$	Output High Voltage	0.88*V_{TT}		V	1,2
$I_{IL\_CMOS1.05v}$	Input Leakage Current		±50	mA	1,2
$I_{O\_CMOS1.05v}$	Output Current (R_{TEST} = 500 ohm)		±1.50	mA	1,2
A_{NM_Rise}	Non-Monotonicity Amplitude, Rising Edge Signal PWRGOOD		0.135	V	5
A_{NM_Fall}	Non-Monotonicity Amplitude, Falling Edge Signal PWRGOOD		0.165	V	5
Open Drain CMOS (ODCMOS) Signals					
V_{IL_ODCMOS}	Input Low Voltage		0.3*V_{TT}	V	1,2
V_{IH_ODCMOS}	Input High Voltage	0.7*V_{TT}		V	1,2
V_{OH_ODCMOS}	Output High Voltage Signals: CAT_ERR_N, ERROR_N[2:0], THERMTRIP_N, PROCHOT_N, CPU_ONLY_RESET		$V_{TT(max)}$	V	1,2
I_{OL}	Output Leakage Current, Signal MEM_HOT_C{1/23}_N		±100	mA	3
I_{OL}	Output Leakage Current (R_{TEST} = 50 ohm)		±900	mA	3
R_{ON}	Buffer On Resistance Signals: CAT_ERR_N, CPU_ONLY_RESET, ERROR_N[2:0], MEM_HOT_C{1/23}_N, PROCHOT_N, THERMTRIP_N		14	W	1,2

Notes:

1. These specifications This table applies to the processor sideband and miscellaneous signals specified in Table 7-5.
2. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
3. For Vin between 0 and Voh.For Vin between 0 and Voh.
4. PWRGOOD Non Monotonicity duration (T_{NM}) time is maximum 1.3 ns.
5. These are measured between VIL and VIH. If the edge rate specification is not met, make sure there is a monotonic edge and the edge rate is not lower than the edge rate specification for the monotonic edges. The monotonic input edge rate is 0.02 V/ns.
6. The waveform could be non-monotonic when measured at the land (near the socket at the bottom side of via) but not when observed at the pad during simulation. The waveform measured at the land could violate specifications defined at the pad. Customers could measure the land timings on their boards and then use the package length information found in the Model Usage Guidelines (MUG) which comes with the I/O model to correlate the results to the specification at the pad.

Table 7-22. Miscellaneous Signals DC Specifications

Symbol	Parameter	Min	Typical	Max	Units	Notes
IVT_ID_N Signal						
$V_{O\_ABS\_MAX}$	Output Absolute Max Voltage		1.10	1.80	V	1
I_O	Output Current			0	µA	1, 3
SKTOCC_N Signal						
$V_{O\_ABS\_MAX}$	Output Absolute Max Voltage		3.30	3.50	V	1
I_{OMAX}	Output Max Current			1	mA	2

Notes:
1. For specific routing guidelines, see the appropriate Platform Design Guide (PDG) for details.
2. See the appropriate Platform Design Guide (PDG) for details.
3. IVT_ID_N land is a no connect on die.

7.8.3.1 PCI Express* DC Specifications

The processor DC specifications for the PCI Express* are available in the *PCI Express Base Specification - Revision 3.0*. This document will provide only the processor exceptions to the *PCI Express Base Specification - Revision 3.0*.

7.8.3.2 DMI2/PCI Express* DC Specifications

The processor DC specifications for the DMI2/PCI Express* are available in the *PCI Express Base Specification 2.0 and 1.0*. This document will provide only the processor exceptions to the *PCI Express Base Specification 2.0 and 1.0*.

7.8.3.3 Intel QuickPath Interconnect DC Specifications

Intel QuickPath Interconnect specifications are defined at the processor lands. Please refer to the appropriate platform design guidelines for specific implementation details. In most cases, termination resistors are not required as these are integrated into the processor silicon.

7.8.3.4 Reset and Miscellaneous Signal DC Specifications

For a power-on Reset, RESET_N must stay active for at least 3.5 millisecond after V_{CC} and BCLK{0/1} have reached their proper specifications. RESET_N must not be kept asserted for more than 100 ms while PWRGOOD is asserted. RESET_N must be held asserted for at least 3.5 millisecond before it is deasserted again. RESET_N must be held asserted before PWRGOOD is asserted. This signal does not have on-die termination and must be terminated on the system board.

Figure 7-6. BCLK{0/1} Differential Clock Crosspoint Specification

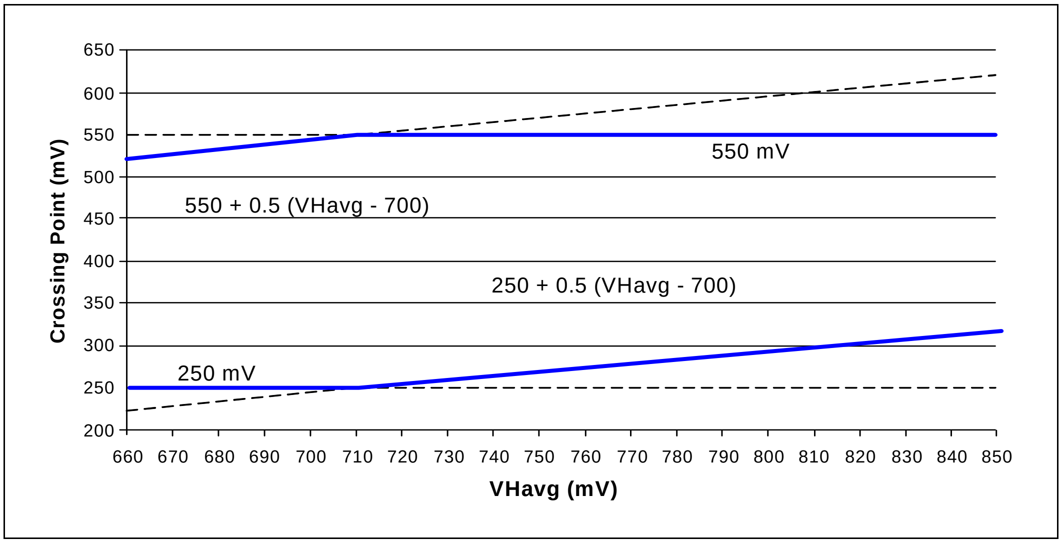

Figure 7-7. BCLK{0/1} Differential Clock Measurement Point for Ringback

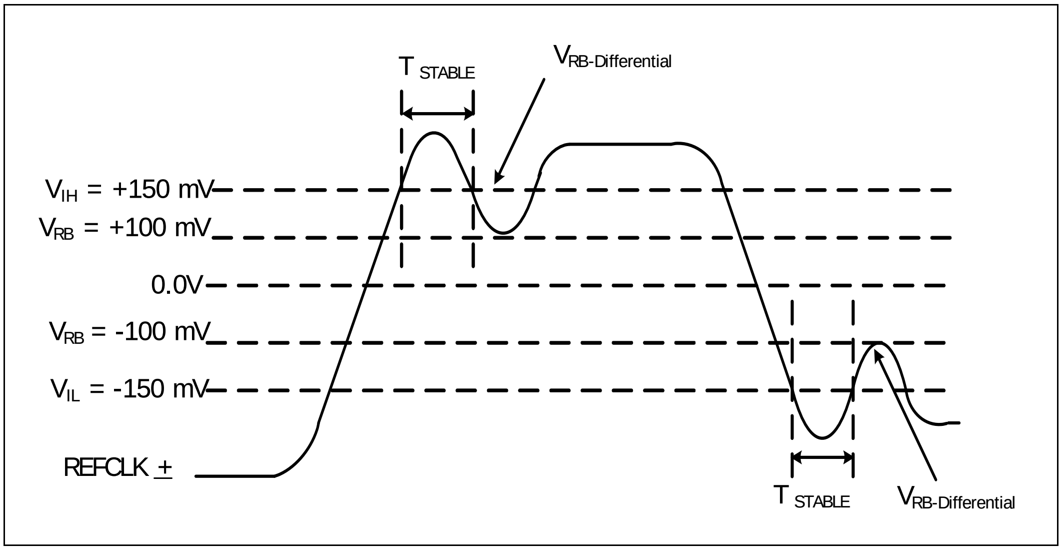

Figure 7-8. BCLK{0/1} Single Ended Clock Measurement Points for Absolute Cross Point and Swing

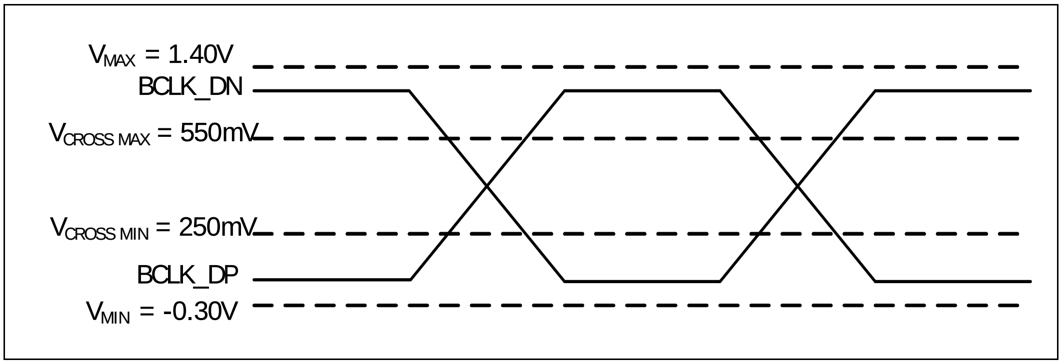

Figure 7-9. BCLK{0/1} Single Ended Clock Measurement Points for Delta Cross Point

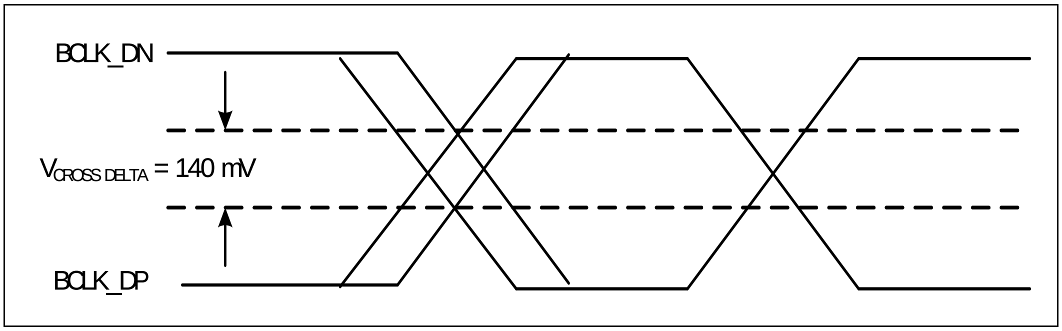

7.9 Signal Quality

Data transfer requires the clean reception of data signals and clock signals. Ringing below receiver thresholds, non-monotonic signal edges, and excessive voltage swings will adversely affect system timings. Ringback and signal non-monotonicity cannot be tolerated since these phenomena may inadvertently advance receiver state machines. Excessive signal swings (overshoot and undershoot) are detrimental to silicon gate oxide integrity, and can cause device failure if absolute voltage limits are exceeded. Overshoot and undershoot can also cause timing degradation due to the build up of inter-symbol interference (ISI) effects.

For these reasons, it is crucial that the designer work towards a solution that provides acceptable signal quality across all systematic variations encountered in volume manufacturing.

This section documents signal quality metrics used to derive topology and routing guidelines through simulation. All specifications are specified at the processor die (pad measurements).

Specifications for signal quality are for measurements at the processor core only and are only observable through simulation. Therefore, proper simulation is the only way to verify proper timing and signal quality.

7.9.1 DDR3 Signal Quality Specifications

Various scenarios for the DDR3 Signals have been simulated to generate a set of layout guidelines which are available in the appropriate Platform Design Guide (PDG).

7.9.2 I/O Signal Quality Specifications

Signal Quality specifications for PCIe* Signals are included as part of the PCIe* DC specifications. Various scenarios have been simulated to generate a set of layout guidelines which are available in the appropriate Platform Design Guide (PDG).

7.9.3 Intel QuickPath Interconnect Signal Quality Specifications

Signal Quality specifications for Differential Intel® QuickPath Interconnect Signals are included as part of the Intel QuickPath Interconnect signal quality specifications. Various scenarios have been simulated to generate a set of layout guidelines which are available in the appropriate Platform Design Guide (PDG).

7.9.4 Input Reference Clock Signal Quality Specifications

Overshoot/Undershoot and Ringback specifications for BCLK{0/1}_D[N/P] are found in Table 7-23. Overshoot/Undershoot and Ringback specifications for the DDR3 Reference Clocks are specified by the DIMM.

7.9.5 Overshoot/Undershoot Tolerance

Overshoot (or undershoot) is the absolute value of the maximum voltage above or below V_{SS}, see Figure 7-10. The overshoot/undershoot specifications limit transitions beyond V_{CCD} or V_{SS} due to the fast signal edge rates. The processor can be damaged by single and/or repeated overshoot or undershoot events on any input, output, or I/O buffer if the charge is large enough (that is, if the over/undershoot is great enough). Determining the impact of an overshoot/undershoot condition requires knowledge of the magnitude, the pulse direction, and the activity factor (AF). Permanent damage to the processor is the likely result of excessive overshoot/undershoot.

Baseboard designs which meet signal integrity and timing requirements and which do not exceed the maximum overshoot or undershoot limits listed in Table 7-23 will insure reliable IO performance for the lifetime of the processor.

Table 7-23. Processor I/O Overshoot/Undershoot Specifications

Signal Group	Minimum Undershoot	Maximum Overshoot	Overshoot Duration	Undershoot Duration	Notes
Intel QuickPath Interconnect	-0.2 * VTT	1.2 * VTT	39 ps	15 ps	1,2
DDR3	-0.2 * V_{CCD}	1.2 * V_{CCD}	0.25*T_{CH}	0.1*T_{CH}	1,2,3
System Reference Clock (BCLK{0/1})	-0.3V	1.15V	N/A	N/A	1,2
PWRGOOD Signal	-0.420V	VTT + 0.28	N/A	N/A	4

Notes:
1. These specifications are measured at the processor pad.
2. Refer to Figure 7-10 for description of allowable Overshoot/Undershoot magnitude and duration.
3. TCH is the minimum high pulse width duration.
4. For PWRGOOD DC specifications see Table 7-21.

7.9.5.1 Overshoot/Undershoot Magnitude

Overshoot/Undershoot magnitude describes the maximum potential difference between a signal and its voltage reference level. For the processor, both overshoot and undershoot magnitude are referenced to V_{SS}. It is important to note that the overshoot and undershoot conditions are separate and their impact must be determined independently.

The pulse magnitude and duration, and activity factor must be used to determine if the overshoot/undershoot pulse is within specifications.

7.9.5.2 Overshoot/Undershoot Pulse Duration

Overshoot/undershoot pulse duration describes the total amount of time that an overshoot/undershoot event exceeds the overshoot/undershoot reference voltage. The total time could encompass several oscillations above the reference voltage. Multiple overshoot/undershoot pulses within a single overshoot/undershoot event may need to be measured to determine the total pulse duration.

Note: Oscillations below the reference voltage cannot be subtracted from the total overshoot/undershoot pulse duration.

7.9.5.3 Activity Factor

Activity factor (AF) describes the frequency of overshoot (or undershoot) occurrence relative to a clock. Since the highest frequency of assertion of any common clock signal is every other clock, an AF = 0.1 indicates that the specific overshoot (or undershoot) waveform occurs every other clock cycle.

The specification provided in the table shows the maximum pulse duration allowed for a given overshoot/undershoot magnitude at a specific activity factor. Each table entry is independent of all others, meaning that the pulse duration reflects the existence of overshoot/undershoot events of that magnitude ONLY. A platform with an overshoot/undershoot that just meets the pulse duration for a specific magnitude where the AF < 0.1, means that there can be no other overshoot/undershoot events, even of lesser magnitude (note that if AF = 0.1, then the event occurs at all times and no other events can occur).

7.9.5.4 Reading Overshoot/Undershoot Specification Tables

The overshoot/undershoot specification for the processor is not a simple single value. Instead, many factors are needed to determine the over/undershoot specification. In addition to the magnitude of the overshoot, the following parameters must also be known: the width of the overshoot and the activity factor (AF). To determine the allowed overshoot for a particular overshoot event, the following must be done:

1. Determine the signal group a particular signal falls into.
2. Determine the magnitude of the overshoot or the undershoot (relative to VSS).
3. Determine the activity factor (How often does this overshoot occur?).
4. Next, from the appropriate specification table, determine the maximum pulse duration (in nanoseconds) allowed.
5. Compare the specified maximum pulse duration to the signal being measured. If the pulse duration measured is less than the pulse duration shown in the table, then the signal meets the specifications.

Undershoot events must be analyzed separately from overshoot events as they are mutually exclusive.

7.9.5.5 Determining if a System Meets the Overshoot/Undershoot Specifications

The overshoot/undershoot specifications listed in the table specify the allowable overshoot/undershoot for a single overshoot/undershoot event. However most systems will have multiple overshoot and/or undershoot events that each have their own set of parameters (duration, AF and magnitude). While each overshoot on its own may meet the overshoot specification, when you add the total impact of all overshoot events, the system may fail. A guideline to ensure a system passes the overshoot and undershoot specifications is shown below.

1. If only one overshoot/undershoot event magnitude occurs, ensure it meets the over/undershoot specifications in the following tables, OR
2. If multiple overshoots and/or multiple undershoots occur, measure the worst case pulse duration for each magnitude and compare the results against the AF = 0.1 specifications. If all of these worst case overshoot or undershoot events meet the

specifications (measured time < specifications) in the table (where AF= 0.1), then the system passes.

Table 7-24. Processor Sideband Signal Group Overshoot/Undershoot Tolerance

Absolute Maximum Overshoot (V)	Absolute Maximum Undershoot (V)	Pulse Duration (ns) AF=0.1	Pulse Duration (ns) AF=0.01
1.3335 V	0.2835 V	3 ns	5 ns
1.2600 V	0.210 V	5 ns	5 ns

Figure 7-10. Maximum Acceptable Overshoot/Undershoot Waveform

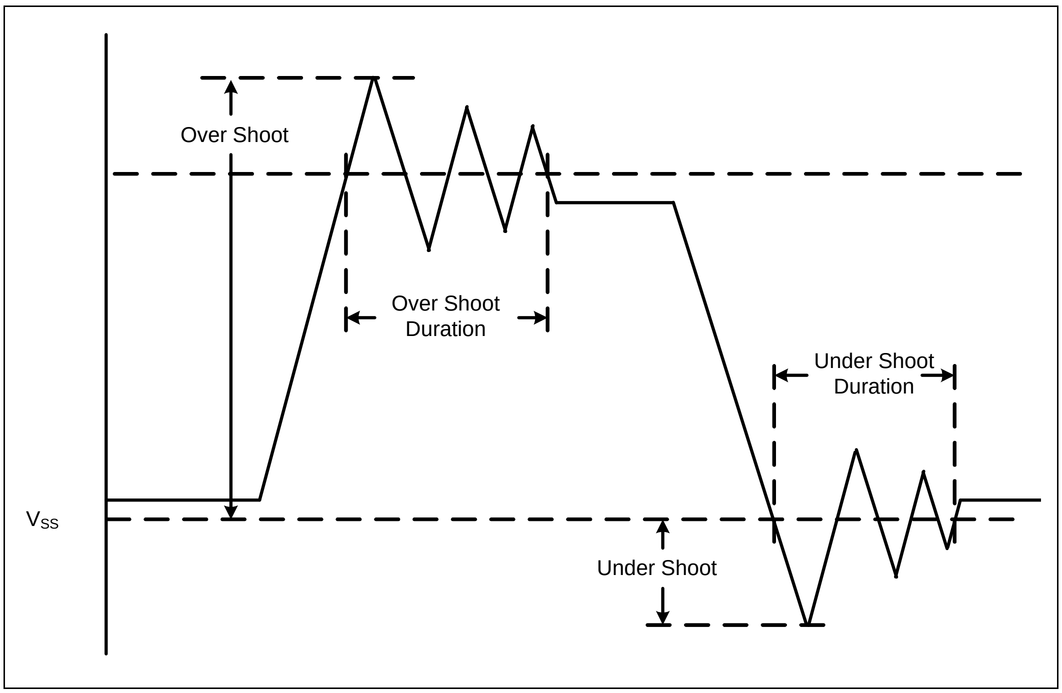

§

8 Processor Land Listing

This chapter provides sorted land list in Section 8.1 and Section 8.2. Table 8-1 is a listing of all Intel® Xeon® processor E5-2400 product family lands ordered alphabetically by land name. Table 8-2 is a listing of all processor lands ordered by land number.

8.1 Listing by Land Name

Table 8-1. Land Name (Sheet 1 of 37)

Land Name	Land Number	Buffer Type	Direction
BCLK0_DN	AP13	CMOS	I
BCLK0_DP	AR13	CMOS	I
BCLK1_DN	AM30	CMOS	I
BCLK1_DP	AN30	CMOS	I
BIST_ENABLE	AF4	CMOS	I
BMCINIT	AF1	CMOS	I
BPM_N[0]	AL10	ODCMOS	I/O
BPM_N[1]	AL11	ODCMOS	I/O
BPM_N[2]	AN9	ODCMOS	I/O
BPM_N[3]	AN11	ODCMOS	I/O
BPM_N[4]	AP10	ODCMOS	I/O
BPM_N[5]	AP11	ODCMOS	I/O
BPM_N[6]	AN10	ODCMOS	I/O
BPM_N[7]	AM11	ODCMOS	I/O
CAT_ERR_N	AT6	ODCMOS	I/O
CPU_ONLY_RESET	AM3	ODCMOS	I/O
DDR_RESET_C1_N	L26	CMOS_1.5V	O
DDR_RESET_C23_N	C29	CMOS_1.5V	O
DDR_SCL_C1	W7	ODCMOS	I/O
DDR_SCL_C23	V40	ODCMOS	I/O
DDR_SDA_C1	W8	ODCMOS	I/O
DDR_SDA_C23	V41	ODCMOS	I/O
DDR_VREFDQRX_C1	N10	DC	I
DDR_VREFDQRX_C23	M36	DC	I
DDR_VREFDQTX_C1	W4	DC	O
DDR_VREFDQTX_C23	V37	DC	O
DDR1_BA[0]	L17	SSTL	O
DDR1_BA[1]	L18	SSTL	O
DDR1_BA[2]	J24	SSTL	O

Table 8-1. Land Name (Sheet 2 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR1_CAS_N	K13	SSTL	O
DDR1_CKE[0]	K25	SSTL	O
DDR1_CKE[1]	J26	SSTL	O
DDR1_CKE[2]	J25	SSTL	O
DDR1_CKE[3]	H26	SSTL	O
DDR1_CLK_DN[0]	H17	SSTL	O
DDR1_CLK_DN[1]	G19	SSTL	O
DDR1_CLK_DN[2]	H16	SSTL	O
DDR1_CLK_DN[3]	H18	SSTL	O
DDR1_CLK_DP[0]	J17	SSTL	O
DDR1_CLK_DP[1]	H19	SSTL	O
DDR1_CLK_DP[2]	J16	SSTL	O
DDR1_CLK_DP[3]	J18	SSTL	O
DDR1_CS_N[0]	J15	SSTL	O
DDR1_CS_N[1]	L15	SSTL	O
DDR1_CS_N[2]	L11	SSTL	O
DDR1_CS_N[3]	K12	SSTL	O
DDR1_CS_N[4]	K15	SSTL	O
DDR1_CS_N[5]	H14	SSTL	O
DDR1_CS_N[6]	L12	SSTL	O
DDR1_CS_N[7]	J12	SSTL	O
DDR1_DQ[00]	AC34	SSTL	I/O
DDR1_DQ[01]	AC35	SSTL	I/O
DDR1_DQ[02]	W35	SSTL	I/O
DDR1_DQ[03]	W36	SSTL	I/O
DDR1_DQ[04]	AD34	SSTL	I/O
DDR1_DQ[05]	AD35	SSTL	I/O
DDR1_DQ[06]	Y35	SSTL	I/O
DDR1_DQ[07]	Y36	SSTL	I/O

Table 8-1. Land Name (Sheet 3 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR1_DQ[08]	T35	SSTL	I/O
DDR1_DQ[09]	T36	SSTL	I/O
DDR1_DQ[10]	N36	SSTL	I/O
DDR1_DQ[11]	N35	SSTL	I/O
DDR1_DQ[12]	U35	SSTL	I/O
DDR1_DQ[13]	U36	SSTL	I/O
DDR1_DQ[14]	P34	SSTL	I/O
DDR1_DQ[15]	N34	SSTL	I/O
DDR1_DQ[16]	K35	SSTL	I/O
DDR1_DQ[17]	K36	SSTL	I/O
DDR1_DQ[18]	K33	SSTL	I/O
DDR1_DQ[19]	L33	SSTL	I/O
DDR1_DQ[20]	L36	SSTL	I/O
DDR1_DQ[21]	L35	SSTL	I/O
DDR1_DQ[22]	J34	SSTL	I/O
DDR1_DQ[23]	J33	SSTL	I/O
DDR1_DQ[24]	L31	SSTL	I/O
DDR1_DQ[25]	K31	SSTL	I/O
DDR1_DQ[26]	L27	SSTL	I/O
DDR1_DQ[27]	K27	SSTL	I/O
DDR1_DQ[28]	L32	SSTL	I/O
DDR1_DQ[29]	K32	SSTL	I/O
DDR1_DQ[30]	L28	SSTL	I/O
DDR1_DQ[31]	K28	SSTL	I/O
DDR1_DQ[32]	K9	SSTL	I/O
DDR1_DQ[33]	L9	SSTL	I/O
DDR1_DQ[34]	K5	SSTL	I/O
DDR1_DQ[35]	L5	SSTL	I/O
DDR1_DQ[36]	K10	SSTL	I/O
DDR1_DQ[37]	L10	SSTL	I/O
DDR1_DQ[38]	K6	SSTL	I/O
DDR1_DQ[39]	L6	SSTL	I/O
DDR1_DQ[40]	N8	SSTL	I/O
DDR1_DQ[41]	P8	SSTL	I/O
DDR1_DQ[42]	V9	SSTL	I/O
DDR1_DQ[43]	V8	SSTL	I/O
DDR1_DQ[44]	N9	SSTL	I/O
DDR1_DQ[45]	P9	SSTL	I/O

Table 8-1. Land Name (Sheet 4 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR1_DQ[46]	U9	SSTL	I/O
DDR1_DQ[47]	U8	SSTL	I/O
DDR1_DQ[48]	AA9	SSTL	I/O
DDR1_DQ[49]	AA8	SSTL	I/O
DDR1_DQ[50]	AE9	SSTL	I/O
DDR1_DQ[51]	AE8	SSTL	I/O
DDR1_DQ[52]	Y9	SSTL	I/O
DDR1_DQ[53]	Y8	SSTL	I/O
DDR1_DQ[54]	AD9	SSTL	I/O
DDR1_DQ[55]	AD8	SSTL	I/O
DDR1_DQ[56]	Y1	SSTL	I/O
DDR1_DQ[57]	AA3	SSTL	I/O
DDR1_DQ[58]	AE1	SSTL	I/O
DDR1_DQ[59]	AE2	SSTL	I/O
DDR1_DQ[60]	Y3	SSTL	I/O
DDR1_DQ[61]	Y2	SSTL	I/O
DDR1_DQ[62]	AD2	SSTL	I/O
DDR1_DQ[63]	AD3	SSTL	I/O
DDR1_DQS_DN[00]	AA35	SSTL	I/O
DDR1_DQS_DN[01]	P35	SSTL	I/O
DDR1_DQS_DN[02]	H33	SSTL	I/O
DDR1_DQS_DN[03]	L29	SSTL	I/O
DDR1_DQS_DN[04]	K7	SSTL	I/O
DDR1_DQS_DN[05]	T9	SSTL	I/O
DDR1_DQS_DN[06]	AC9	SSTL	I/O
DDR1_DQS_DN[07]	AC2	SSTL	I/O
DDR1_DQS_DN[08]	H29	SSTL	I/O
DDR1_DQS_DN[09]	AB35	SSTL	I/O
DDR1_DQS_DN[10]	R36	SSTL	I/O
DDR1_DQS_DN[11]	K34	SSTL	I/O
DDR1_DQS_DN[12]	K30	SSTL	I/O
DDR1_DQS_DN[13]	L8	SSTL	I/O
DDR1_DQS_DN[14]	R8	SSTL	I/O
DDR1_DQS_DN[15]	AB8	SSTL	I/O
DDR1_DQS_DN[16]	AC3	SSTL	I/O
DDR1_DQS_DN[17]	G30	SSTL	I/O
DDR1_DQS_DP[00]	AA36	SSTL	I/O
DDR1_DQS_DP[01]	P36	SSTL	I/O

Table 8-1. Land Name (Sheet 5 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR1_DQS_DP[02]	H34	SSTL	I/O
DDR1_DQS_DP[03]	K29	SSTL	I/O
DDR1_DQS_DP[04]	L7	SSTL	I/O
DDR1_DQS_DP[05]	T8	SSTL	I/O
DDR1_DQS_DP[06]	AC8	SSTL	I/O
DDR1_DQS_DP[07]	AC1	SSTL	I/O
DDR1_DQS_DP[08]	G29	SSTL	I/O
DDR1_DQS_DP[09]	AB34	SSTL	I/O
DDR1_DQS_DP[10]	R35	SSTL	I/O
DDR1_DQS_DP[11]	L34	SSTL	I/O
DDR1_DQS_DP[12]	L30	SSTL	I/O
DDR1_DQS_DP[13]	K8	SSTL	I/O
DDR1_DQS_DP[14]	R9	SSTL	I/O
DDR1_DQS_DP[15]	AB9	SSTL	I/O
DDR1_DQS_DP[16]	AB3	SSTL	I/O
DDR1_DQS_DP[17]	H30	SSTL	I/O
DDR1_ECC[0]	G32	SSTL	I/O
DDR1_ECC[1]	G31	SSTL	I/O
DDR1_ECC[2]	H27	SSTL	I/O
DDR1_ECC[3]	G27	SSTL	I/O
DDR1_ECC[4]	H31	SSTL	I/O
DDR1_ECC[5]	H32	SSTL	I/O
DDR1_ECC[6]	H28	SSTL	I/O
DDR1_ECC[7]	G28	SSTL	I/O
DDR1_MA[00]	K19	SSTL	O
DDR1_MA[01]	L20	SSTL	O
DDR1_MA[02]	K20	SSTL	O
DDR1_MA[03]	L21	SSTL	O
DDR1_MA[04]	M22	SSTL	O
DDR1_MA[05]	K22	SSTL	O
DDR1_MA[06]	L22	SSTL	O
DDR1_MA[07]	L23	SSTL	O
DDR1_MA[08]	K23	SSTL	O
DDR1_MA[09]	L25	SSTL	O
DDR1_MA[10]	M18	SSTL	O
DDR1_MA[11]	M24	SSTL	O
DDR1_MA[12]	J20	SSTL	O
DDR1_MA[13]	L13	SSTL	O

Table 8-1. Land Name (Sheet 6 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR1_MA[14]	J22	SSTL	O
DDR1_MA[15]	K24	SSTL	O
DDR1_MA_PAR	J19	SSTL	O
DDR1_ODT[0]	G15	SSTL	O
DDR1_ODT[1]	H13	SSTL	O
DDR1_ODT[2]	J14	SSTL	O
DDR1_ODT[3]	G13	SSTL	O
DDR1_PAR_ERR_N	J21	SSTL	I
DDR1_RAS_N	K17	SSTL	O
DDR1_RCOMP[0]	J4	Analog	I
DDR1_RCOMP[1]	N7	Analog	I
DDR1_RCOMP[2]	M4	Analog	I
DDR1_WE_N	K14	SSTL	O
DDR2_BA[0]	E17	SSTL	O
DDR2_BA[1]	G18	SSTL	O
DDR2_BA[2]	F26	SSTL	O
DDR2_CAS_N	F13	SSTL	O
DDR2_CKE[0]	D27	SSTL	O
DDR2_CKE[1]	E28	SSTL	O
DDR2_CKE[2]	E27	SSTL	O
DDR2_CKE[3]	D28	SSTL	O
DDR2_CLK_DN[0]	E19	SSTL	O
DDR2_CLK_DN[1]	H21	SSTL	O
DDR2_CLK_DN[2]	G20	SSTL	O
DDR2_CLK_DN[3]	F21	SSTL	O
DDR2_CLK_DP[0]	E20	SSTL	O
DDR2_CLK_DP[1]	G21	SSTL	O
DDR2_CLK_DP[2]	F20	SSTL	O
DDR2_CLK_DP[3]	F22	SSTL	O
DDR2_CS_N[0]	G16	SSTL	O
DDR2_CS_N[1]	G14	SSTL	O
DDR2_CS_N[2]	F11	SSTL	O
DDR2_CS_N[3]	J11	SSTL	O
DDR2_CS_N[4]	E15	SSTL	O
DDR2_CS_N[5]	E14	SSTL	O
DDR2_CS_N[6]	G11	SSTL	O
DDR2_CS_N[7]	H11	SSTL	O
DDR2_DQ[00]	AC39	SSTL	I/O

Table 8-1. Land Name (Sheet 7 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR2_DQ[01]	AC38	SSTL	I/O
DDR2_DQ[02]	W39	SSTL	I/O
DDR2_DQ[03]	W38	SSTL	I/O
DDR2_DQ[04]	AD37	SSTL	I/O
DDR2_DQ[05]	AC37	SSTL	I/O
DDR2_DQ[06]	Y39	SSTL	I/O
DDR2_DQ[07]	Y38	SSTL	I/O
DDR2_DQ[08]	T39	SSTL	I/O
DDR2_DQ[09]	T38	SSTL	I/O
DDR2_DQ[10]	M38	SSTL	I/O
DDR2_DQ[11]	M39	SSTL	I/O
DDR2_DQ[12]	U39	SSTL	I/O
DDR2_DQ[13]	U38	SSTL	I/O
DDR2_DQ[14]	N39	SSTL	I/O
DDR2_DQ[15]	N38	SSTL	I/O
DDR2_DQ[16]	J38	SSTL	I/O
DDR2_DQ[17]	J39	SSTL	I/O
DDR2_DQ[18]	H36	SSTL	I/O
DDR2_DQ[19]	G36	SSTL	I/O
DDR2_DQ[20]	K38	SSTL	I/O
DDR2_DQ[21]	K39	SSTL	I/O
DDR2_DQ[22]	H37	SSTL	I/O
DDR2_DQ[23]	G37	SSTL	I/O
DDR2_DQ[24]	D39	SSTL	I/O
DDR2_DQ[25]	C39	SSTL	I/O
DDR2_DQ[26]	C37	SSTL	I/O
DDR2_DQ[27]	E37	SSTL	I/O
DDR2_DQ[28]	E38	SSTL	I/O
DDR2_DQ[29]	E39	SSTL	I/O
DDR2_DQ[30]	D37	SSTL	I/O
DDR2_DQ[31]	B37	SSTL	I/O
DDR2_DQ[32]	D9	SSTL	I/O
DDR2_DQ[33]	E9	SSTL	I/O
DDR2_DQ[34]	E5	SSTL	I/O
DDR2_DQ[35]	D5	SSTL	I/O
DDR2_DQ[36]	D10	SSTL	I/O
DDR2_DQ[37]	E10	SSTL	I/O
DDR2_DQ[38]	D6	SSTL	I/O

Table 8-1. Land Name (Sheet 8 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR2_DQ[39]	E6	SSTL	I/O
DDR2_DQ[40]	G9	SSTL	I/O
DDR2_DQ[41]	H9	SSTL	I/O
DDR2_DQ[42]	G5	SSTL	I/O
DDR2_DQ[43]	H5	SSTL	I/O
DDR2_DQ[44]	G10	SSTL	I/O
DDR2_DQ[45]	H10	SSTL	I/O
DDR2_DQ[46]	G6	SSTL	I/O
DDR2_DQ[47]	H6	SSTL	I/O
DDR2_DQ[48]	P6	SSTL	I/O
DDR2_DQ[49]	P5	SSTL	I/O
DDR2_DQ[50]	V6	SSTL	I/O
DDR2_DQ[51]	V5	SSTL	I/O
DDR2_DQ[52]	N6	SSTL	I/O
DDR2_DQ[53]	N5	SSTL	I/O
DDR2_DQ[54]	U6	SSTL	I/O
DDR2_DQ[55]	U5	SSTL	I/O
DDR2_DQ[56]	AA6	SSTL	I/O
DDR2_DQ[57]	AA5	SSTL	I/O
DDR2_DQ[58]	AE6	SSTL	I/O
DDR2_DQ[59]	AE5	SSTL	I/O
DDR2_DQ[60]	Y6	SSTL	I/O
DDR2_DQ[61]	Y5	SSTL	I/O
DDR2_DQ[62]	AD6	SSTL	I/O
DDR2_DQ[63]	AD5	SSTL	I/O
DDR2_DQS_DN[00]	AA39	SSTL	I/O
DDR2_DQS_DN[01]	P39	SSTL	I/O
DDR2_DQS_DN[02]	G39	SSTL	I/O
DDR2_DQS_DN[03]	C38	SSTL	I/O
DDR2_DQS_DN[04]	D7	SSTL	I/O
DDR2_DQS_DN[05]	G7	SSTL	I/O
DDR2_DQS_DN[06]	T6	SSTL	I/O
DDR2_DQS_DN[07]	AC6	SSTL	I/O
DDR2_DQS_DN[08]	E31	SSTL	I/O
DDR2_DQS_DN[09]	AB38	SSTL	I/O
DDR2_DQS_DN[10]	R38	SSTL	I/O
DDR2_DQS_DN[11]	H38	SSTL	I/O
DDR2_DQS_DN[12]	A39	SSTL	I/O

Table 8-1. Land Name (Sheet 9 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR2_DQS_DN[13]	E8	SSTL	I/O
DDR2_DQS_DN[14]	H8	SSTL	I/O
DDR2_DQS_DN[15]	R5	SSTL	I/O
DDR2_DQS_DN[16]	AB5	SSTL	I/O
DDR2_DQS_DN[17]	D32	SSTL	I/O
DDR2_DQS_DP[00]	AA38	SSTL	I/O
DDR2_DQS_DP[01]	P38	SSTL	I/O
DDR2_DQS_DP[02]	G38	SSTL	I/O
DDR2_DQS_DP[03]	B38	SSTL	I/O
DDR2_DQS_DP[04]	E7	SSTL	I/O
DDR2_DQS_DP[05]	H7	SSTL	I/O
DDR2_DQS_DP[06]	T5	SSTL	I/O
DDR2_DQS_DP[07]	AC5	SSTL	I/O
DDR2_DQS_DP[08]	D31	SSTL	I/O
DDR2_DQS_DP[09]	AB39	SSTL	I/O
DDR2_DQS_DP[10]	R39	SSTL	I/O
DDR2_DQS_DP[11]	H39	SSTL	I/O
DDR2_DQS_DP[12]	B39	SSTL	I/O
DDR2_DQS_DP[13]	D8	SSTL	I/O
DDR2_DQS_DP[14]	G8	SSTL	I/O
DDR2_DQS_DP[15]	R6	SSTL	I/O
DDR2_DQS_DP[16]	AB6	SSTL	I/O
DDR2_DQS_DP[17]	E32	SSTL	I/O
DDR2_ECC[0]	E33	SSTL	I/O
DDR2_ECC[1]	F34	SSTL	I/O
DDR2_ECC[2]	E29	SSTL	I/O
DDR2_ECC[3]	D29	SSTL	I/O
DDR2_ECC[4]	E35	SSTL	I/O
DDR2_ECC[5]	E34	SSTL	I/O
DDR2_ECC[6]	E30	SSTL	I/O
DDR2_ECC[7]	D30	SSTL	I/O
DDR2_MA[00]	D19	SSTL	O
DDR2_MA[01]	E22	SSTL	O
DDR2_MA[02]	H22	SSTL	O
DDR2_MA[03]	G23	SSTL	O
DDR2_MA[04]	F23	SSTL	O
DDR2_MA[05]	H23	SSTL	O
DDR2_MA[06]	E24	SSTL	O

Table 8-1. Land Name (Sheet 10 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR2_MA[07]	G24	SSTL	O
DDR2_MA[08]	D24	SSTL	O
DDR2_MA[09]	H24	SSTL	O
DDR2_MA[10]	F18	SSTL	O
DDR2_MA[11]	E25	SSTL	O
DDR2_MA[12]	F25	SSTL	O
DDR2_MA[13]	E13	SSTL	O
DDR2_MA[14]	D26	SSTL	O
DDR2_MA[15]	G26	SSTL	O
DDR2_MA_PAR	E18	SSTL	O
DDR2_ODT[0]	F16	SSTL	O
DDR2_ODT[1]	F12	SSTL	O
DDR2_ODT[2]	H12	SSTL	O
DDR2_ODT[3]	E12	SSTL	O
DDR2_PAR_ERR_N	G25	SSTL	I
DDR2_RAS_N	F17	SSTL	O
DDR2_WE_N	F15	SSTL	O
DDR23_RCOMP[0]	H35	Analog	I
DDR23_RCOMP[1]	E36	Analog	I
DDR23_RCOMP[2]	L38	Analog	I
DDR3_BA[0]	B16	SSTL	O
DDR3_BA[1]	D17	SSTL	O
DDR3_BA[2]	C27	SSTL	O
DDR3_CAS_N	B15	SSTL	O
DDR3_CKE[0]	B28	SSTL	O
DDR3_CKE[1]	B29	SSTL	O
DDR3_CKE[2]	A28	SSTL	O
DDR3_CKE[3]	C28	SSTL	O
DDR3_CLK_DN[0]	B19	SSTL	O
DDR3_CLK_DN[1]	A20	SSTL	O
DDR3_CLK_DN[2]	A18	SSTL	O
DDR3_CLK_DN[3]	C21	SSTL	O
DDR3_CLK_DP[0]	C19	SSTL	O
DDR3_CLK_DP[1]	B20	SSTL	O
DDR3_CLK_DP[2]	B18	SSTL	O
DDR3_CLK_DP[3]	C20	SSTL	O
DDR3_CS_N[0]	D16	SSTL	O
DDR3_CS_N[1]	B13	SSTL	O

Table 8-1. Land Name (Sheet 11 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR3_CS_N[2]	B11	SSTL	O
DDR3_CS_N[3]	D11	SSTL	O
DDR3_CS_N[4]	A15	SSTL	O
DDR3_CS_N[5]	B14	SSTL	O
DDR3_CS_N[6]	C11	SSTL	O
DDR3_CS_N[7]	D12	SSTL	O
DDR3_DQ[00]	AC41	SSTL	I/O
DDR3_DQ[01]	AC42	SSTL	I/O
DDR3_DQ[02]	W42	SSTL	I/O
DDR3_DQ[03]	W43	SSTL	I/O
DDR3_DQ[04]	AD42	SSTL	I/O
DDR3_DQ[05]	AD43	SSTL	I/O
DDR3_DQ[06]	Y41	SSTL	I/O
DDR3_DQ[07]	W41	SSTL	I/O
DDR3_DQ[08]	U42	SSTL	I/O
DDR3_DQ[09]	T43	SSTL	I/O
DDR3_DQ[10]	P41	SSTL	I/O
DDR3_DQ[11]	N43	SSTL	I/O
DDR3_DQ[12]	U41	SSTL	I/O
DDR3_DQ[13]	U43	SSTL	I/O
DDR3_DQ[14]	P43	SSTL	I/O
DDR3_DQ[15]	P42	SSTL	I/O
DDR3_DQ[16]	L43	SSTL	I/O
DDR3_DQ[17]	L42	SSTL	I/O
DDR3_DQ[18]	H43	SSTL	I/O
DDR3_DQ[19]	H41	SSTL	I/O
DDR3_DQ[20]	M41	SSTL	I/O
DDR3_DQ[21]	L41	SSTL	I/O
DDR3_DQ[22]	J43	SSTL	I/O
DDR3_DQ[23]	H42	SSTL	I/O
DDR3_DQ[24]	F41	SSTL	I/O
DDR3_DQ[25]	E41	SSTL	I/O
DDR3_DQ[26]	C41	SSTL	I/O
DDR3_DQ[27]	B41	SSTL	I/O
DDR3_DQ[28]	F42	SSTL	I/O
DDR3_DQ[29]	F43	SSTL	I/O
DDR3_DQ[30]	C42	SSTL	I/O
DDR3_DQ[31]	D41	SSTL	I/O

Table 8-1. Land Name (Sheet 12 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR3_DQ[32]	A9	SSTL	I/O
DDR3_DQ[33]	B9	SSTL	I/O
DDR3_DQ[34]	B5	SSTL	I/O
DDR3_DQ[35]	B4	SSTL	I/O
DDR3_DQ[36]	A10	SSTL	I/O
DDR3_DQ[37]	B10	SSTL	I/O
DDR3_DQ[38]	A6	SSTL	I/O
DDR3_DQ[39]	B6	SSTL	I/O
DDR3_DQ[40]	E3	SSTL	I/O
DDR3_DQ[41]	D2	SSTL	I/O
DDR3_DQ[42]	G1	SSTL	I/O
DDR3_DQ[43]	G2	SSTL	I/O
DDR3_DQ[44]	C3	SSTL	I/O
DDR3_DQ[45]	D3	SSTL	I/O
DDR3_DQ[46]	F3	SSTL	I/O
DDR3_DQ[47]	G3	SSTL	I/O
DDR3_DQ[48]	J2	SSTL	I/O
DDR3_DQ[49]	K1	SSTL	I/O
DDR3_DQ[50]	M3	SSTL	I/O
DDR3_DQ[51]	N3	SSTL	I/O
DDR3_DQ[52]	J3	SSTL	I/O
DDR3_DQ[53]	J1	SSTL	I/O
DDR3_DQ[54]	M2	SSTL	I/O
DDR3_DQ[55]	M1	SSTL	I/O
DDR3_DQ[56]	R2	SSTL	I/O
DDR3_DQ[57]	R1	SSTL	I/O
DDR3_DQ[58]	V1	SSTL	I/O
DDR3_DQ[59]	V3	SSTL	I/O
DDR3_DQ[60]	P1	SSTL	I/O
DDR3_DQ[61]	R3	SSTL	I/O
DDR3_DQ[62]	U1	SSTL	I/O
DDR3_DQ[63]	V2	SSTL	I/O
DDR3_DQS_DN[00]	AB41	SSTL	I/O
DDR3_DQS_DN[01]	R43	SSTL	I/O
DDR3_DQS_DN[02]	K41	SSTL	I/O
DDR3_DQS_DN[03]	D43	SSTL	I/O
DDR3_DQS_DN[04]	A7	SSTL	I/O
DDR3_DQS_DN[05]	F2	SSTL	I/O

Table 8-1. Land Name (Sheet 13 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR3_DQS_DN[06]	L1	SSTL	I/O
DDR3_DQS_DN[07]	T3	SSTL	I/O
DDR3_DQS_DN[08]	B32	SSTL	I/O
DDR3_DQS_DN[09]	AB43	SSTL	I/O
DDR3_DQS_DN[10]	R41	SSTL	I/O
DDR3_DQS_DN[11]	K43	SSTL	I/O
DDR3_DQS_DN[12]	E42	SSTL	I/O
DDR3_DQS_DN[13]	B8	SSTL	I/O
DDR3_DQS_DN[14]	E1	SSTL	I/O
DDR3_DQS_DN[15]	L3	SSTL	I/O
DDR3_DQS_DN[16]	T1	SSTL	I/O
DDR3_DQS_DN[17]	B33	SSTL	I/O
DDR3_DQS_DP[00]	AA41	SSTL	I/O
DDR3_DQS_DP[01]	R42	SSTL	I/O
DDR3_DQS_DP[02]	J41	SSTL	I/O
DDR3_DQS_DP[03]	D42	SSTL	I/O
DDR3_DQS_DP[04]	B7	SSTL	I/O
DDR3_DQS_DP[05]	F1	SSTL	I/O
DDR3_DQS_DP[06]	L2	SSTL	I/O
DDR3_DQS_DP[07]	U3	SSTL	I/O
DDR3_DQS_DP[08]	A32	SSTL	I/O
DDR3_DQS_DP[09]	AB42	SSTL	I/O
DDR3_DQS_DP[10]	T41	SSTL	I/O
DDR3_DQS_DP[11]	K42	SSTL	I/O
DDR3_DQS_DP[12]	E43	SSTL	I/O
DDR3_DQS_DP[13]	A8	SSTL	I/O
DDR3_DQS_DP[14]	E2	SSTL	I/O
DDR3_DQS_DP[15]	K3	SSTL	I/O
DDR3_DQS_DP[16]	T2	SSTL	I/O
DDR3_DQS_DP[17]	B34	SSTL	I/O
DDR3_ECC[0]	C34	SSTL	I/O
DDR3_ECC[1]	C33	SSTL	I/O
DDR3_ECC[2]	B30	SSTL	I/O
DDR3_ECC[3]	A30	SSTL	I/O
DDR3_ECC[4]	C35	SSTL	I/O
DDR3_ECC[5]	B35	SSTL	I/O
DDR3_ECC[6]	B31	SSTL	I/O
DDR3_ECC[7]	A31	SSTL	I/O

Table 8-1. Land Name (Sheet 14 of 37)

Land Name	Land Number	Buffer Type	Direction
DDR3_MA[00]	C17	SSTL	O
DDR3_MA[01]	D21	SSTL	O
DDR3_MA[02]	D22	SSTL	O
DDR3_MA[03]	C22	SSTL	O
DDR3_MA[04]	D23	SSTL	O
DDR3_MA[05]	C24	SSTL	O
DDR3_MA[06]	C23	SSTL	O
DDR3_MA[07]	D25	SSTL	O
DDR3_MA[08]	B24	SSTL	O
DDR3_MA[09]	B25	SSTL	O
DDR3_MA[10]	A17	SSTL	O
DDR3_MA[11]	A25	SSTL	O
DDR3_MA[12]	B26	SSTL	O
DDR3_MA[13]	D14	SSTL	O
DDR3_MA[14]	C26	SSTL	O
DDR3_MA[15]	A27	SSTL	O
DDR3_MA_PAR	C18	SSTL	O
DDR3_ODT[0]	D15	SSTL	O
DDR3_ODT[1]	C12	SSTL	O
DDR3_ODT[2]	C14	SSTL	O
DDR3_ODT[3]	C13	SSTL	O
DDR3_PAR_ERR_N	A26	SSTL	I
DDR3_RAS_N	A16	SSTL	O
DDR3_WE_N	C16	SSTL	O
DMI_RX_DN[0]	AJ34	PCIEX	I
DMI_RX_DN[1]	AH35	PCIEX	I
DMI_RX_DN[2]	AG34	PCIEX	I
DMI_RX_DN[3]	AF35	PCIEX	I
DMI_RX_DP[0]	AJ33	PCIEX	I
DMI_RX_DP[1]	AH34	PCIEX	I
DMI_RX_DP[2]	AG33	PCIEX	I
DMI_RX_DP[3]	AF34	PCIEX	I
DMI_TX_DN[0]	AM35	PCIEX	O
DMI_TX_DN[1]	AL36	PCIEX	O
DMI_TX_DN[2]	AK35	PCIEX	O
DMI_TX_DN[3]	AJ36	PCIEX	O
DMI_TX_DP[0]	AM36	PCIEX	O
DMI_TX_DP[1]	AL37	PCIEX	O

Table 8-1. Land Name (Sheet 15 of 37)

Land Name	Land Number	Buffer Type	Direction
DMI_TX_DP[2]	AK36	PCIEX	O
DMI_TX_DP[3]	AJ37	PCIEX	O
TXT_PLTEN	AN8	CMOS	I
DRAM_PWR_OK_C1	Y10	CMOS_1.5V	I
DRAM_PWR_OK_C23	AD40	CMOS_1.5V	I
EAR_N	AH3	ODCMOS	I/O
ERROR_N[0]	AL34	ODCMOS	O
ERROR_N[1]	AM34	ODCMOS	O
ERROR_N[2]	AL33	ODCMOS	O
FRMAGENT	AF3	CMOS	I
IVT_ID_N	V34		O
TXT_AGENT	AE4	CMOS	I
MEM_HOT_C1_N	W10	ODCMOS	I/O
MEM_HOT_C23_N	N41	ODCMOS	I/O
PE_RBIAS	AL31	PCIEX3	I/O
PE_RBIAS_SENSE	AL32	PCIEX3	I
PE_VREF_CAP	AM32	PCIEX3	I/O
PE1A_RX_DN[0]	AW41	PCIEX3	I
PE1A_RX_DN[1]	AV39	PCIEX3	I
PE1A_RX_DN[2]	AU38	PCIEX3	I
PE1A_RX_DN[3]	AT37	PCIEX3	I
PE1A_RX_DP[0]	AV41	PCIEX3	I
PE1A_RX_DP[1]	AU39	PCIEX3	I
PE1A_RX_DP[2]	AT38	PCIEX3	I
PE1A_RX_DP[3]	AR37	PCIEX3	I
PE1A_TX_DN[0]	AE40	PCIEX3	O
PE1A_TX_DN[1]	AG37	PCIEX3	O
PE1A_TX_DN[2]	AF38	PCIEX3	O
PE1A_TX_DN[3]	AG39	PCIEX3	O
PE1A_TX_DP[0]	AE39	PCIEX3	O
PE1A_TX_DP[1]	AG36	PCIEX3	O
PE1A_TX_DP[2]	AF39	PCIEX3	O
PE1A_TX_DP[3]	AG40	PCIEX3	O
PE1B_RX_DN[4]	AY40	PCIEX3	I
PE1B_RX_DN[5]	BA39	PCIEX3	I
PE1B_RX_DN[6]	AY38	PCIEX3	I
PE1B_RX_DN[7]	AW37	PCIEX3	I
PE1B_RX_DP[4]	AW40	PCIEX3	I

Table 8-1. Land Name (Sheet 16 of 37)

Land Name	Land Number	Buffer Type	Direction
PE1B_RX_DP[5]	AY39	PCIEX3	I
PE1B_RX_DP[6]	AW38	PCIEX3	I
PE1B_RX_DP[7]	AV37	PCIEX3	I
PE1B_TX_DN[4]	AF41	PCIEX3	O
PE1B_TX_DN[5]	AG42	PCIEX3	O
PE1B_TX_DN[6]	AH41	PCIEX3	O
PE1B_TX_DN[7]	AJ42	PCIEX3	O
PE1B_TX_DP[4]	AF42	PCIEX3	O
PE1B_TX_DP[5]	AG43	PCIEX3	O
PE1B_TX_DP[6]	AH42	PCIEX3	O
PE1B_TX_DP[7]	AJ43	PCIEX3	O
PE3A_RX_DN[0]	AU36	PCIEX3	I
PE3A_RX_DN[1]	AT35	PCIEX3	I
PE3A_RX_DN[2]	AU34	PCIEX3	I
PE3A_RX_DN[3]	AT33	PCIEX3	I
PE3A_RX_DP[0]	AT36	PCIEX3	I
PE3A_RX_DP[1]	AR35	PCIEX3	I
PE3A_RX_DP[2]	AT34	PCIEX3	I
PE3A_RX_DP[3]	AR33	PCIEX3	I
PE3A_TX_DN[0]	AH38	PCIEX3	O
PE3A_TX_DN[1]	AJ39	PCIEX3	O
PE3A_TX_DN[2]	AK38	PCIEX3	O
PE3A_TX_DN[3]	AL39	PCIEX3	O
PE3A_TX_DP[0]	AH39	PCIEX3	O
PE3A_TX_DP[1]	AJ40	PCIEX3	O
PE3A_TX_DP[2]	AK39	PCIEX3	O
PE3A_TX_DP[3]	AL40	PCIEX3	O
PE3B_RX_DN[4]	AU32	PCIEX3	I
PE3B_RX_DN[5]	AV31	PCIEX3	I
PE3B_RX_DN[6]	AU30	PCIEX3	I
PE3B_RX_DN[7]	AV29	PCIEX3	I
PE3B_RX_DP[4]	AT32	PCIEX3	I
PE3B_RX_DP[5]	AU31	PCIEX3	I
PE3B_RX_DP[6]	AT30	PCIEX3	I
PE3B_RX_DP[7]	AU29	PCIEX3	I
PE3B_TX_DN[4]	AM38	PCIEX3	O
PE3B_TX_DN[5]	AN39	PCIEX3	O
PE3B_TX_DN[6]	AP38	PCIEX3	O

Table 8-1. Land Name (Sheet 17 of 37)

Land Name	Land Number	Buffer Type	Direction
PE3B_TX_DN[7]	AR39	PCIEX3	O
PE3B_TX_DP[4]	AM39	PCIEX3	O
PE3B_TX_DP[5]	AN40	PCIEX3	O
PE3B_TX_DP[6]	AP39	PCIEX3	O
PE3B_TX_DP[7]	AR40	PCIEX3	O
PE3C_RX_DN[10]	AY34	PCIEX3	I
PE3C_RX_DN[11]	AW33	PCIEX3	I
PE3C_RX_DN[8]	AY36	PCIEX3	I
PE3C_RX_DN[9]	AW35	PCIEX3	I
PE3C_RX_DP[10]	AW34	PCIEX3	I
PE3C_RX_DP[11]	AV33	PCIEX3	I
PE3C_RX_DP[8]	AW36	PCIEX3	I
PE3C_RX_DP[9]	AV35	PCIEX3	I
PE3C_TX_DN[10]	AM41	PCIEX3	O
PE3C_TX_DN[11]	AN42	PCIEX3	O
PE3C_TX_DN[8]	AK41	PCIEX3	O
PE3C_TX_DN[9]	AL42	PCIEX3	O
PE3C_TX_DP[10]	AM42	PCIEX3	O
PE3C_TX_DP[11]	AN43	PCIEX3	O
PE3C_TX_DP[8]	AK42	PCIEX3	O
PE3C_TX_DP[9]	AL43	PCIEX3	O
PE3D_RX_DN[12]	AY32	PCIEX3	I
PE3D_RX_DN[13]	BA31	PCIEX3	I
PE3D_RX_DN[14]	AY30	PCIEX3	I
PE3D_RX_DN[15]	BA29	PCIEX3	I
PE3D_RX_DP[12]	AW32	PCIEX3	I
PE3D_RX_DP[13]	AY31	PCIEX3	I
PE3D_RX_DP[14]	AW30	PCIEX3	I
PE3D_RX_DP[15]	AY29	PCIEX3	I
PE3D_TX_DN[12]	AP41	PCIEX3	O
PE3D_TX_DN[13]	AR42	PCIEX3	O
PE3D_TX_DN[14]	AT41	PCIEX3	O
PE3D_TX_DN[15]	AU42	PCIEX3	O
PE3D_TX_DP[12]	AP42	PCIEX3	O
PE3D_TX_DP[13]	AR43	PCIEX3	O
PE3D_TX_DP[14]	AT42	PCIEX3	O
PE3D_TX_DP[15]	AU43	PCIEX3	O
PECI	AN37	PECI	I/O

Table 8-1. Land Name (Sheet 18 of 37)

Land Name	Land Number	Buffer Type	Direction
PEHPSCL	AR30	ODCMOS	I/O
PEHPSDA	AP30	ODCMOS	I/O
PMSYNC	AT3	CMOS	I
PRDY_N	AR11	CMOS	O
PREQ_N	AR10	CMOS	I/O
PROCHOT_N	AH7	ODCMOS	I/O
PWRGOOD	AK6	CMOS	I
QPI_RBIAS	AK13	Analog	I/O
QPI_RBIAS_SENSE	AL13	Analog	I
QPI_VREF_CAP	AM13	Intel QPI	I/O
QPI1_CLKRX_DN	AL5	Intel QPI	I
QPI1_CLKRX_DP	AL6	Intel QPI	I
QPI1_CLKTX_DN	AY9	Intel QPI	O
QPI1_CLKTX_DP	BA9	Intel QPI	O
QPI1_DRX_DN[00]	AT4	Intel QPI	I
QPI1_DRX_DN[01]	AT1	Intel QPI	I
QPI1_DRX_DN[02]	AR5	Intel QPI	I
QPI1_DRX_DN[03]	AR2	Intel QPI	I
QPI1_DRX_DN[04]	AP4	Intel QPI	I
QPI1_DRX_DN[05]	AP1	Intel QPI	I
QPI1_DRX_DN[06]	AP8	Intel QPI	I
QPI1_DRX_DN[07]	AN2	Intel QPI	I
QPI1_DRX_DN[08]	AN6	Intel QPI	I
QPI1_DRX_DN[09]	AM1	Intel QPI	I
QPI1_DRX_DN[10]	AM5	Intel QPI	I
QPI1_DRX_DN[11]	AL2	Intel QPI	I
QPI1_DRX_DN[12]	AK5	Intel QPI	I
QPI1_DRX_DN[13]	AK1	Intel QPI	I
QPI1_DRX_DN[14]	AJ6	Intel QPI	I
QPI1_DRX_DN[15]	AJ2	Intel QPI	I
QPI1_DRX_DN[16]	AH5	Intel QPI	I
QPI1_DRX_DN[17]	AH1	Intel QPI	I
QPI1_DRX_DN[18]	AG6	Intel QPI	I
QPI1_DRX_DN[19]	AG2	Intel QPI	I
QPI1_DRX_DP[00]	AT5	Intel QPI	I
QPI1_DRX_DP[01]	AT2	Intel QPI	I
QPI1_DRX_DP[02]	AR6	Intel QPI	I
QPI1_DRX_DP[03]	AR3	Intel QPI	I

Table 8-1. Land Name (Sheet 19 of 37)

Land Name	Land Number	Buffer Type	Direction
QPI1_DRX_DP[04]	AP5	Intel QPI	I
QPI1_DRX_DP[05]	AP2	Intel QPI	I
QPI1_DRX_DP[06]	AP7	Intel QPI	I
QPI1_DRX_DP[07]	AN3	Intel QPI	I
QPI1_DRX_DP[08]	AN5	Intel QPI	I
QPI1_DRX_DP[09]	AM2	Intel QPI	I
QPI1_DRX_DP[10]	AM4	Intel QPI	I
QPI1_DRX_DP[11]	AL3	Intel QPI	I
QPI1_DRX_DP[12]	AK4	Intel QPI	I
QPI1_DRX_DP[13]	AK2	Intel QPI	I
QPI1_DRX_DP[14]	AJ5	Intel QPI	I
QPI1_DRX_DP[15]	AJ3	Intel QPI	I
QPI1_DRX_DP[16]	AH4	Intel QPI	I
QPI1_DRX_DP[17]	AH2	Intel QPI	I
QPI1_DRX_DP[18]	AG5	Intel QPI	I
QPI1_DRX_DP[19]	AG3	Intel QPI	I
QPI1_DTX_DN[00]	AW14	Intel QPI	O
QPI1_DTX_DN[01]	AY13	Intel QPI	O
QPI1_DTX_DN[02]	AU13	Intel QPI	O
QPI1_DTX_DN[03]	AW12	Intel QPI	O
QPI1_DTX_DN[04]	AT12	Intel QPI	O
QPI1_DTX_DN[05]	AY11	Intel QPI	O
QPI1_DTX_DN[06]	AU11	Intel QPI	O
QPI1_DTX_DN[07]	AW10	Intel QPI	O
QPI1_DTX_DN[08]	AT10	Intel QPI	O
QPI1_DTX_DN[09]	AW8	Intel QPI	O
QPI1_DTX_DN[10]	AU9	Intel QPI	O
QPI1_DTX_DN[11]	AY7	Intel QPI	O
QPI1_DTX_DN[12]	AT8	Intel QPI	O
QPI1_DTX_DN[13]	BA4	Intel QPI	O
QPI1_DTX_DN[14]	AU7	Intel QPI	O
QPI1_DTX_DN[15]	AY3	Intel QPI	O
QPI1_DTX_DN[16]	AW5	Intel QPI	O
QPI1_DTX_DN[17]	AW2	Intel QPI	O
QPI1_DTX_DN[18]	AV4	Intel QPI	O
QPI1_DTX_DN[19]	AV1	Intel QPI	O
QPI1_DTX_DP[00]	AY14	Intel QPI	O
QPI1_DTX_DP[01]	BA13	Intel QPI	O

Table 8-1. Land Name (Sheet 20 of 37)

Land Name	Land Number	Buffer Type	Direction
QPI1_DTX_DP[02]	AV13	Intel QPI	O
QPI1_DTX_DP[03]	AY12	Intel QPI	O
QPI1_DTX_DP[04]	AU12	Intel QPI	O
QPI1_DTX_DP[05]	BA11	Intel QPI	O
QPI1_DTX_DP[06]	AV11	Intel QPI	O
QPI1_DTX_DP[07]	AY10	Intel QPI	O
QPI1_DTX_DP[08]	AU10	Intel QPI	O
QPI1_DTX_DP[09]	AY8	Intel QPI	O
QPI1_DTX_DP[10]	AV9	Intel QPI	O
QPI1_DTX_DP[11]	BA7	Intel QPI	O
QPI1_DTX_DP[12]	AU8	Intel QPI	O
QPI1_DTX_DP[13]	BA5	Intel QPI	O
QPI1_DTX_DP[14]	AV7	Intel QPI	O
QPI1_DTX_DP[15]	AY4	Intel QPI	O
QPI1_DTX_DP[16]	AW6	Intel QPI	O
QPI1_DTX_DP[17]	AW3	Intel QPI	O
QPI1_DTX_DP[18]	AV5	Intel QPI	O
QPI1_DTX_DP[19]	AV2	Intel QPI	O
RESET_N	AU2	CMOS	I
RSVD	AP37		
RSVD	AT39		
RSVD	AU6		
RSVD	AV42		
RSVD	AR8		
RSVD	AR9		
RSVD	N4		
RSVD	AR7		
RSVD	AN34		
RSVD	AN33		
RSVD	AN36		
RSVD	AP36		
RSVD	AH11		
RSVD	AG10		
RSVD	AK10		
RSVD	AM9		
RSVD	AK11		
RSVD	AL9		
RSVD	AJ10		

Table 8-1. Land Name (Sheet 21 of 37)

Land Name	Land Number	Buffer Type	Direction
RSVD	AG9		
RSVD	AG11		
RSVD	AH10		
RSVD	AF11		
RSVD	AK9		
RSVD	AH9		
RSVD	AJ8		
RSVD	AM8		
RSVD	AH8		
RSVD	AG8		
RSVD	AK8		
RSVD	AY41		
RSVD	AV38		
RSVD	AR31		
RSVD	AP31		
RSVD	BA38		
RSVD	AP32		
RSVD	AN32		
RSVD	AY6		
RSVD	AW13		
RSVD	AT14		
RSVD	AU14		
RSVD	AM14		
RSVD	AN14		
RSVD	BA6		
RSVD	AU3		
RSVD	AU4		
RSVD	AL30		
RSVD	M20		
RSVD	M19		
RSVD	L16		
RSVD	M12		
RSVD	M21		
RSVD	AF7		
RSVD	AE7		
RSVD	B2		
SAFE_MODE_BOOT	AU40	CMOS	I
SKTOCC_N	AM31		O

Table 8-1. Land Name (Sheet 22 of 37)

Land Name	Land Number	Buffer Type	Direction
SOCKET_ID[0]	AW11	CMOS	I
SVIDALERT_N	AL7	CMOS	I
SVIDCLK	AM7	ODCMOS	O
SVIDDATA	AK7	ODCMOS	I/O
TCK	AR12	CMOS	I
TDI	AM12	CMOS	I
TDO	AK12	ODCMOS	O
TEST0	H4		O
TEST1	C4		O
TEST2	N33		O
TEST3	G40		O
TEST4	AK32		I
THERMTRIP_N	AG7	ODCMOS	O
TMS	AP12	CMOS	I
TRST_N	AL12	CMOS	I
VCC	AA11	PWR	
VCC	AA33	PWR	
VCC	AB11	PWR	
VCC	AB33	PWR	
VCC	AC11	PWR	
VCC	AD11	PWR	
VCC	AE11	PWR	
VCC	AK15	PWR	
VCC	AK16	PWR	
VCC	AK18	PWR	
VCC	AK19	PWR	
VCC	AK21	PWR	
VCC	AK22	PWR	
VCC	AK24	PWR	
VCC	AK25	PWR	
VCC	AK27	PWR	
VCC	AK28	PWR	
VCC	AL15	PWR	
VCC	AL16	PWR	
VCC	AL18	PWR	
VCC	AL19	PWR	
VCC	AL21	PWR	
VCC	AL22	PWR	

Table 8-1. Land Name (Sheet 23 of 37)

Land Name	Land Number	Buffer Type	Direction
VCC	AL24	PWR	
VCC	AL25	PWR	
VCC	AL27	PWR	
VCC	AL28	PWR	
VCC	AM15	PWR	
VCC	AM16	PWR	
VCC	AM18	PWR	
VCC	AM19	PWR	
VCC	AM21	PWR	
VCC	AM22	PWR	
VCC	AM24	PWR	
VCC	AM25	PWR	
VCC	AM27	PWR	
VCC	AM28	PWR	
VCC	AN15	PWR	
VCC	AN16	PWR	
VCC	AN18	PWR	
VCC	AN19	PWR	
VCC	AN21	PWR	
VCC	AN22	PWR	
VCC	AN24	PWR	
VCC	AN25	PWR	
VCC	AN27	PWR	
VCC	AN28	PWR	
VCC	AP15	PWR	
VCC	AP16	PWR	
VCC	AP18	PWR	
VCC	AP19	PWR	
VCC	AP21	PWR	
VCC	AP22	PWR	
VCC	AP24	PWR	
VCC	AP25	PWR	
VCC	AP27	PWR	
VCC	AP28	PWR	
VCC	AR15	PWR	
VCC	AR16	PWR	
VCC	AR18	PWR	
VCC	AR19	PWR	

Table 8-1. Land Name (Sheet 24 of 37)

Land Name	Land Number	Buffer Type	Direction
VCC	AR21	PWR	
VCC	AR22	PWR	
VCC	AR24	PWR	
VCC	AR25	PWR	
VCC	AR27	PWR	
VCC	AR28	PWR	
VCC	AT15	PWR	
VCC	AT16	PWR	
VCC	AT18	PWR	
VCC	AT19	PWR	
VCC	AT21	PWR	
VCC	AT22	PWR	
VCC	AT24	PWR	
VCC	AT25	PWR	
VCC	AT27	PWR	
VCC	AT28	PWR	
VCC	AU15	PWR	
VCC	AU16	PWR	
VCC	AU18	PWR	
VCC	AU19	PWR	
VCC	AU21	PWR	
VCC	AU22	PWR	
VCC	AU24	PWR	
VCC	AU25	PWR	
VCC	AU27	PWR	
VCC	AU28	PWR	
VCC	AV15	PWR	
VCC	AV16	PWR	
VCC	AV18	PWR	
VCC	AV19	PWR	
VCC	AV21	PWR	
VCC	AV22	PWR	
VCC	AV24	PWR	
VCC	AV25	PWR	
VCC	AV27	PWR	
VCC	AV28	PWR	
VCC	AW15	PWR	
VCC	AW16	PWR	

Table 8-1. Land Name (Sheet 25 of 37)

Land Name	Land Number	Buffer Type	Direction
VCC	AW18	PWR	
VCC	AW19	PWR	
VCC	AW21	PWR	
VCC	AW22	PWR	
VCC	AW24	PWR	
VCC	AW25	PWR	
VCC	AW27	PWR	
VCC	AW28	PWR	
VCC	AY15	PWR	
VCC	AY16	PWR	
VCC	AY18	PWR	
VCC	AY19	PWR	
VCC	AY24	PWR	
VCC	AY25	PWR	
VCC	AY27	PWR	
VCC	AY28	PWR	
VCC	BA15	PWR	
VCC	BA16	PWR	
VCC	BA18	PWR	
VCC	BA19	PWR	
VCC	BA24	PWR	
VCC	BA25	PWR	
VCC	BA27	PWR	
VCC	BA28	PWR	
VCC	R11	PWR	
VCC	R33	PWR	
VCC	T11	PWR	
VCC	T33	PWR	
VCC	U11	PWR	
VCC	U33	PWR	
VCC	V11	PWR	
VCC	V33	PWR	
VCC	W11	PWR	
VCC	W33	PWR	
VCC	Y11	PWR	
VCC	Y33	PWR	
VCC_SENSE	P11		O
VCCD	E21	PWR	

Table 8-1. Land Name (Sheet 26 of 37)

Land Name	Land Number	Buffer Type	Direction
VCCD	E26	PWR	
VCCD	F14	PWR	
VCCD	F19	PWR	
VCCD	F24	PWR	
VCCD	G12	PWR	
VCCD	G17	PWR	
VCCD	G22	PWR	
VCCD	H15	PWR	
VCCD	H20	PWR	
VCCD	H25	PWR	
VCCD	J13	PWR	
VCCD	K11	PWR	
VCCD	K16	PWR	
VCCD	K18	PWR	
VCCD	K26	PWR	
VCCPLL	AR36	PWR	
VCCPLL	AV36	PWR	
VSA	AD39	PWR	
VSA	AE35	PWR	
VSA	AE36	PWR	
VSA	AE37	PWR	
VSA	AE38	PWR	
VSA	AE41	PWR	
VSA	AE42	PWR	
VSA	AE43	PWR	
VSA	AF33	PWR	
VSA	AF36	PWR	
VSA	AF37	PWR	
VSA	AG41	PWR	
VSA	AH37	PWR	
VSA	AJ41	PWR	
VSA	AL38	PWR	
VSA	AL41	PWR	
VSA	AM37	PWR	
VSA	AN41	PWR	
VSA	AR41	PWR	
VSA	AT40	PWR	
VSA	AV40	PWR	

Table 8-1. Land Name (Sheet 27 of 37)

Land Name	Land Number	Buffer Type	Direction
VSA	AV43	PWR	
VSA	AY42	PWR	
VSA_SENSE	AE33		O
VSS	A14	GND	
VSS	A19	GND	
VSS	A24	GND	
VSS	A29	GND	
VSS	A4	GND	
VSS	A40	GND	
VSS	A41	GND	
VSS	A5	GND	
VSS	AA10	GND	
VSS	AA34	GND	
VSS	AA37	GND	
VSS	AA4	GND	
VSS	AA40	GND	
VSS	AA7	GND	
VSS	AB10	GND	
VSS	AB36	GND	
VSS	AB37	GND	
VSS	AB4	GND	
VSS	AB40	GND	
VSS	AB7	GND	
VSS	AC10	GND	
VSS	AC33	GND	
VSS	AC36	GND	
VSS	AC4	GND	
VSS	AC40	GND	
VSS	AC43	GND	
VSS	AC7	GND	
VSS	AD1	GND	
VSS	AD10	GND	
VSS	AD33	GND	
VSS	AD36	GND	
VSS	AD38	GND	
VSS	AD4	GND	
VSS	AD41	GND	
VSS	AD7	GND	

Table 8-1. Land Name (Sheet 28 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	AE10	GND	
VSS	AE3	GND	
VSS	AF10	GND	
VSS	AF2	GND	
VSS	AF40	GND	
VSS	AF43	GND	
VSS	AF5	GND	
VSS	AF6	GND	
VSS	AF8	GND	
VSS	AG1	GND	
VSS	AG35	GND	
VSS	AG38	GND	
VSS	AG4	GND	
VSS	AH33	GND	
VSS	AH36	GND	
VSS	AH40	GND	
VSS	AH43	GND	
VSS	AH6	GND	
VSS	AJ1	GND	
VSS	AJ11	GND	
VSS	AJ35	GND	
VSS	AJ38	GND	
VSS	AJ4	GND	
VSS	AJ9	GND	
VSS	AK14	GND	
VSS	AK17	GND	
VSS	AK20	GND	
VSS	AK23	GND	
VSS	AK26	GND	
VSS	AK29	GND	
VSS	AK31	GND	
VSS	AK37	GND	
VSS	AK40	GND	
VSS	AK43	GND	
VSS	AL1	GND	
VSS	AL14	GND	
VSS	AL17	GND	
VSS	AL20	GND	

Table 8-1. Land Name (Sheet 29 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	AL23	GND	
VSS	AL26	GND	
VSS	AL29	GND	
VSS	AL35	GND	
VSS	AL4	GND	
VSS	AL8	GND	
VSS	AM10	GND	
VSS	AM17	GND	
VSS	AM20	GND	
VSS	AM23	GND	
VSS	AM26	GND	
VSS	AM29	GND	
VSS	AM33	GND	
VSS	AM40	GND	
VSS	AM43	GND	
VSS	AM6	GND	
VSS	AN1	GND	
VSS	AN12	GND	
VSS	AN13	GND	
VSS	AN17	GND	
VSS	AN20	GND	
VSS	AN23	GND	
VSS	AN26	GND	
VSS	AN29	GND	
VSS	AN31	GND	
VSS	AN38	GND	
VSS	AN4	GND	
VSS	AP14	GND	
VSS	AP17	GND	
VSS	AP20	GND	
VSS	AP23	GND	
VSS	AP26	GND	
VSS	AP29	GND	
VSS	AP40	GND	
VSS	AP43	GND	
VSS	AP6	GND	
VSS	AP9	GND	
VSS	AR1	GND	

Table 8-1. Land Name (Sheet 30 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	AR14	GND	
VSS	AR17	GND	
VSS	AR20	GND	
VSS	AR23	GND	
VSS	AR26	GND	
VSS	AR29	GND	
VSS	AR32	GND	
VSS	AR34	GND	
VSS	AR38	GND	
VSS	AR4	GND	
VSS	AT11	GND	
VSS	AT13	GND	
VSS	AT17	GND	
VSS	AT20	GND	
VSS	AT23	GND	
VSS	AT26	GND	
VSS	AT29	GND	
VSS	AT31	GND	
VSS	AT43	GND	
VSS	AT9	GND	
VSS	AU1	GND	
VSS	AU17	GND	
VSS	AU20	GND	
VSS	AU23	GND	
VSS	AU26	GND	
VSS	AU33	GND	
VSS	AU35	GND	
VSS	AU37	GND	
VSS	AU41	GND	
VSS	AU5	GND	
VSS	AV10	GND	
VSS	AV12	GND	
VSS	AV14	GND	
VSS	AV17	GND	
VSS	AV20	GND	
VSS	AV23	GND	
VSS	AV26	GND	
VSS	AV3	GND	

Table 8-1. Land Name (Sheet 31 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	AV6	GND	
VSS	AV8	GND	
VSS	AW1	GND	
VSS	AW17	GND	
VSS	AW20	GND	
VSS	AW23	GND	
VSS	AW26	GND	
VSS	AW29	GND	
VSS	AW31	GND	
VSS	AW39	GND	
VSS	AW4	GND	
VSS	AW42	GND	
VSS	AY17	GND	
VSS	AY20	GND	
VSS	AY26	GND	
VSS	AY33	GND	
VSS	AY35	GND	
VSS	AY37	GND	
VSS	AY5	GND	
VSS	B12	GND	
VSS	B17	GND	
VSS	B27	GND	
VSS	B3	GND	
VSS	B36	GND	
VSS	B40	GND	
VSS	B42	GND	
VSS	BA10	GND	
VSS	BA12	GND	
VSS	BA14	GND	
VSS	BA17	GND	
VSS	BA20	GND	
VSS	BA26	GND	
VSS	BA3	GND	
VSS	BA32	GND	
VSS	BA40	GND	
VSS	BA8	GND	
VSS	C10	GND	
VSS	C15	GND	

Table 8-1. Land Name (Sheet 32 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	C2	GND	
VSS	C25	GND	
VSS	C30	GND	
VSS	C31	GND	
VSS	C32	GND	
VSS	C36	GND	
VSS	C40	GND	
VSS	C43	GND	
VSS	C5	GND	
VSS	C6	GND	
VSS	C7	GND	
VSS	C8	GND	
VSS	C9	GND	
VSS	D1	GND	
VSS	D13	GND	
VSS	D18	GND	
VSS	D20	GND	
VSS	D33	GND	
VSS	D34	GND	
VSS	D35	GND	
VSS	D36	GND	
VSS	D38	GND	
VSS	D4	GND	
VSS	D40	GND	
VSS	E11	GND	
VSS	E16	GND	
VSS	E23	GND	
VSS	E4	GND	
VSS	E40	GND	
VSS	F10	GND	
VSS	F27	GND	
VSS	F28	GND	
VSS	F29	GND	
VSS	F30	GND	
VSS	F31	GND	
VSS	F32	GND	
VSS	F33	GND	
VSS	F35	GND	

Table 8-1. Land Name (Sheet 33 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	F36	GND	
VSS	F37	GND	
VSS	F38	GND	
VSS	F39	GND	
VSS	F4	GND	
VSS	F40	GND	
VSS	F5	GND	
VSS	F6	GND	
VSS	F7	GND	
VSS	F8	GND	
VSS	F9	GND	
VSS	G33	GND	
VSS	G34	GND	
VSS	G35	GND	
VSS	G4	GND	
VSS	G41	GND	
VSS	G42	GND	
VSS	G43	GND	
VSS	H1	GND	
VSS	H2	GND	
VSS	H3	GND	
VSS	H40	GND	
VSS	J10	GND	
VSS	J23	GND	
VSS	J27	GND	
VSS	J28	GND	
VSS	J29	GND	
VSS	J30	GND	
VSS	J31	GND	
VSS	J32	GND	
VSS	J35	GND	
VSS	J36	GND	
VSS	J37	GND	
VSS	J40	GND	
VSS	J42	GND	
VSS	J5	GND	
VSS	J6	GND	
VSS	J7	GND	

Table 8-1. Land Name (Sheet 34 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	J8	GND	
VSS	J9	GND	
VSS	K2	GND	
VSS	K21	GND	
VSS	K37	GND	
VSS	K4	GND	
VSS	K40	GND	
VSS	L14	GND	
VSS	L19	GND	
VSS	L24	GND	
VSS	L37	GND	
VSS	L39	GND	
VSS	L4	GND	
VSS	L40	GND	
VSS	M10	GND	
VSS	M11	GND	
VSS	M13	GND	
VSS	M16	GND	
VSS	M23	GND	
VSS	M27	GND	
VSS	M28	GND	
VSS	M32	GND	
VSS	M33	GND	
VSS	M34	GND	
VSS	M35	GND	
VSS	M37	GND	
VSS	M40	GND	
VSS	M42	GND	
VSS	M43	GND	
VSS	M5	GND	
VSS	M6	GND	
VSS	M7	GND	
VSS	M8	GND	
VSS	M9	GND	
VSS	N1	GND	
VSS	N2	GND	
VSS	N37	GND	
VSS	N40	GND	

Table 8-1. Land Name (Sheet 35 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	N42	GND	
VSS	P10	GND	
VSS	P2	GND	
VSS	P3	GND	
VSS	P33	GND	
VSS	P37	GND	
VSS	P4	GND	
VSS	P40	GND	
VSS	P7	GND	
VSS	R10	GND	
VSS	R34	GND	
VSS	R37	GND	
VSS	R4	GND	
VSS	R40	GND	
VSS	R7	GND	
VSS	T10	GND	
VSS	T34	GND	
VSS	T37	GND	
VSS	T4	GND	
VSS	T40	GND	
VSS	T42	GND	
VSS	T7	GND	
VSS	U10	GND	
VSS	U2	GND	
VSS	U34	GND	
VSS	U37	GND	
VSS	U4	GND	
VSS	U40	GND	
VSS	U7	GND	
VSS	V10	GND	
VSS	V35	GND	
VSS	V36	GND	
VSS	V38	GND	
VSS	V39	GND	
VSS	V4	GND	
VSS	V42	GND	
VSS	V43	GND	
VSS	V7	GND	

Table 8-1. Land Name (Sheet 36 of 37)

Land Name	Land Number	Buffer Type	Direction
VSS	W1	GND	
VSS	W2	GND	
VSS	W3	GND	
VSS	W34	GND	
VSS	W37	GND	
VSS	W40	GND	
VSS	W5	GND	
VSS	W6	GND	
VSS	W9	GND	
VSS	Y34	GND	
VSS	Y37	GND	
VSS	Y4	GND	
VSS	Y40	GND	
VSS	Y7	GND	
VSS_VCC_SENSE	N11		O
VSS_VSA_SENSE	AE34		O
VSS_VTTD_SENSE	AK33		O
VTTA	AT7	PWR	
VTTA	AW7	PWR	
VTTA	AW9	PWR	
VTTA	AY2	PWR	
VTTA	AN35	PWR	
VTTA	AP33	PWR	
VTTA	AP34	PWR	
VTTA	AP35	PWR	
VTTA	AV34	PWR	
VTTD	AF9	PWR	
VTTD	AJ7	PWR	
VTTD	AK3	PWR	
VTTD	AK30	PWR	
VTTD	AN7	PWR	
VTTD	AP3	PWR	
VTTD	AV30	PWR	
VTTD	AV32	PWR	
VTTD	BA30	PWR	
VTTD	M14	PWR	
VTTD	M15	PWR	
VTTD	M17	PWR	

Table 8-1. Land Name (Sheet 37 of 37)

Land Name	Land Number	Buffer Type	Direction
VTTD	M25	PWR	
VTTD	M26	PWR	
VTTD	M29	PWR	
VTTD	M30	PWR	
VTTD	M31	PWR	
VTTD_SENSE	AK34		O

8.2 Listing by Land Number

Table 8-2. Land Number (Sheet 1 of 37)

Land Number	Land Name	Buffer Type	Direction
A10	DDR3_DQ[36]	SSTL	I/O
A14	VSS	GND	
A15	DDR3_CS_N[4]	SSTL	O
A16	DDR3_RAS_N	SSTL	O
A17	DDR3_MA[10]	SSTL	O
A18	DDR3_CLK_DN[2]	SSTL	O
A19	VSS	GND	
A20	DDR3_CLK_DN[1]	SSTL	O
A24	VSS	GND	
A25	DDR3_MA[11]	SSTL	O
A26	DDR3_PAR_ERR_N	SSTL	I
A27	DDR3_MA[15]	SSTL	O
A28	DDR3_CKE[2]	SSTL	O
A29	VSS	GND	
A30	DDR3_ECC[3]	SSTL	I/O
A31	DDR3_ECC[7]	SSTL	I/O
A32	DDR3_DQS_DP[08]	SSTL	I/O
A39	DDR2_DQS_DN[12]	SSTL	I/O
A4	VSS	GND	
A40	VSS	GND	
A41	VSS	GND	
A5	VSS	GND	
A6	DDR3_DQ[38]	SSTL	I/O
A7	DDR3_DQS_DN[04]	SSTL	I/O
A8	DDR3_DQS_DP[13]	SSTL	I/O
A9	DDR3_DQ[32]	SSTL	I/O

Table 8-2. Land Number (Sheet 2 of 37)

Land Number	Land Name	Buffer Type	Direction
AA10	VSS	GND	
AA11	VCC	PWR	
AA3	DDR1_DQ[57]	SSTL	I/O
AA33	VCC	PWR	
AA34	VSS	GND	
AA35	DDR1_DQS_DN[00]	SSTL	I/O
AA36	DDR1_DQS_DP[00]	SSTL	I/O
AA37	VSS	GND	
AA38	DDR2_DQS_DP[00]	SSTL	I/O
AA39	DDR2_DQS_DN[00]	SSTL	I/O
AA4	VSS	GND	
AA40	VSS	GND	
AA41	DDR3_DQS_DP[00]	SSTL	I/O
AA5	DDR2_DQ[57]	SSTL	I/O
AA6	DDR2_DQ[56]	SSTL	I/O
AA7	VSS	GND	
AA8	DDR1_DQ[49]	SSTL	I/O
AA9	DDR1_DQ[48]	SSTL	I/O
AB10	VSS	GND	
AB11	VCC	PWR	
AB3	DDR1_DQS_DP[16]	SSTL	I/O
AB33	VCC	PWR	
AB34	DDR1_DQS_DP[09]	SSTL	I/O
AB35	DDR1_DQS_DN[09]	SSTL	I/O
AB36	VSS	GND	
AB37	VSS	GND	
AB38	DDR2_DQS_DN[09]	SSTL	I/O

Table 8-2. Land Number (Sheet 3 of 37)

Land Number	Land Name	Buffer Type	Direction
AB39	DDR2_DQS_DP[09]	SSTL	I/O
AB4	VSS	GND	
AB40	VSS	GND	
AB41	DDR3_DQS_DN[00]	SSTL	I/O
AB42	DDR3_DQS_DP[09]	SSTL	I/O
AB43	DDR3_DQS_DN[09]	SSTL	I/O
AB5	DDR2_DQS_DN[16]	SSTL	I/O
AB6	DDR2_DQS_DP[16]	SSTL	I/O
AB7	VSS	GND	
AB8	DDR1_DQS_DN[15]	SSTL	I/O
AB9	DDR1_DQS_DP[15]	SSTL	I/O
AC1	DDR1_DQS_DP[07]	SSTL	I/O
AC10	VSS	GND	
AC11	VCC	PWR	
AC2	DDR1_DQS_DN[07]	SSTL	I/O
AC3	DDR1_DQS_DN[16]	SSTL	I/O
AC33	VSS	GND	
AC34	DDR1_DQ[00]	SSTL	I/O
AC35	DDR1_DQ[01]	SSTL	I/O
AC36	VSS	GND	
AC37	DDR2_DQ[05]	SSTL	I/O
AC38	DDR2_DQ[01]	SSTL	I/O
AC39	DDR2_DQ[00]	SSTL	I/O
AC4	VSS	GND	
AC40	VSS	GND	
AC41	DDR3_DQ[00]	SSTL	I/O
AC42	DDR3_DQ[01]	SSTL	I/O
AC43	VSS	GND	
AC5	DDR2_DQS_DP[07]	SSTL	I/O
AC6	DDR2_DQS_DN[07]	SSTL	I/O
AC7	VSS	GND	
AC8	DDR1_DQS_DP[06]	SSTL	I/O
AC9	DDR1_DQS_DN[06]	SSTL	I/O
AD1	VSS	GND	
AD10	VSS	GND	
AD11	VCC	PWR	
AD2	DDR1_DQ[62]	SSTL	I/O
AD3	DDR1_DQ[63]	SSTL	I/O

Table 8-2. Land Number (Sheet 4 of 37)

Land Number	Land Name	Buffer Type	Direction
AD33	VSS	GND	
AD34	DDR1_DQ[04]	SSTL	I/O
AD35	DDR1_DQ[05]	SSTL	I/O
AD36	VSS	GND	
AD37	DDR2_DQ[04]	SSTL	I/O
AD38	VSS	GND	
AD39	VSA	PWR	
AD4	VSS	GND	
AD40	DRAM_PWR_OK_C23	CMOS_1.5 V	I
AD41	VSS	GND	
AD42	DDR3_DQ[04]	SSTL	I/O
AD43	DDR3_DQ[05]	SSTL	I/O
AD5	DDR2_DQ[63]	SSTL	I/O
AD6	DDR2_DQ[62]	SSTL	I/O
AD7	VSS	GND	
AD8	DDR1_DQ[55]	SSTL	I/O
AD9	DDR1_DQ[54]	SSTL	I/O
AE1	DDR1_DQ[58]	SSTL	I/O
AE10	VSS	GND	
AE11	VCC	PWR	
AE2	DDR1_DQ[59]	SSTL	I/O
AE3	VSS	GND	
AE33	VSA_SENSE		O
AE34	VSS_VSA_SENSE		O
AE35	VSA	PWR	
AE36	VSA	PWR	
AE37	VSA	PWR	
AE38	VSA	PWR	
AE39	PE1A_TX_DP[0]	PCIEX3	O
AE4	TXT_AGENT	CMOS	I
AE40	PE1A_TX_DN[0]	PCIEX3	O
AE41	VSA	PWR	
AE42	VSA	PWR	
AE43	VSA	PWR	
AE5	DDR2_DQ[59]	SSTL	I/O
AE6	DDR2_DQ[58]	SSTL	I/O
AE7	RSVD		
AE8	DDR1_DQ[51]	SSTL	I/O

Table 8-2. Land Number (Sheet 5 of 37)

Land Number	Land Name	Buffer Type	Direction
AE9	DDR1_DQ[50]	SSTL	I/O
AF1	BMCINIT	CMOS	I
AF10	VSS	GND	
AF11	RSVD		
AF2	VSS	GND	
AF3	FRMAGENT	CMOS	I
AF33	VSA	PWR	
AF34	DMI_RX_DP[3]	PCIEX	I
AF35	DMI_RX_DN[3]	PCIEX	I
AF36	VSA	PWR	
AF37	VSA	PWR	
AF38	PE1A_TX_DN[2]	PCIEX3	O
AF39	PE1A_TX_DP[2]	PCIEX3	O
AF4	BIST_ENABLE	CMOS	I
AF40	VSS	GND	
AF41	PE1B_TX_DN[4]	PCIEX3	O
AF42	PE1B_TX_DP[4]	PCIEX3	O
AF43	VSS	GND	
AF5	VSS	GND	
AF6	VSS	GND	
AF7	RSVD		
AF8	VSS	GND	
AF9	VTTD	PWR	
AG1	VSS	GND	
AG10	RSVD		
AG11	RSVD		
AG2	QPI1_DRX_DN[19]	Intel QPI	I
AG3	QPI1_DRX_DP[19]	Intel QPI	I
AG33	DMI_RX_DP[2]	PCIEX	I
AG34	DMI_RX_DN[2]	PCIEX	I
AG35	VSS	GND	
AG36	PE1A_TX_DP[1]	PCIEX3	O
AG37	PE1A_TX_DN[1]	PCIEX3	O
AG38	VSS	GND	
AG39	PE1A_TX_DN[3]	PCIEX3	O
AG4	VSS	GND	
AG40	PE1A_TX_DP[3]	PCIEX3	O
AG41	VSA	PWR	

Table 8-2. Land Number (Sheet 6 of 37)

Land Number	Land Name	Buffer Type	Direction
AG42	PE1B_TX_DN[5]	PCIEX3	O
AG43	PE1B_TX_DP[5]	PCIEX3	O
AG5	QPI1_DRX_DP[18]	Intel QPI	I
AG6	QPI1_DRX_DN[18]	Intel QPI	I
AG7	THERMTRIP_N	ODCMOS	O
AG8	RSVD		
AG9	RSVD		
AH1	QPI1_DRX_DN[17]	Intel QPI	I
AH10	RSVD		
AH11	RSVD		
AH2	QPI1_DRX_DP[17]	Intel QPI	I
AH3	EAR_N	ODCMOS	I/O
AH33	VSS	GND	
AH34	DMI_RX_DP[1]	PCIEX	I
AH35	DMI_RX_DN[1]	PCIEX	I
AH36	VSS	GND	
AH37	VSA	PWR	
AH38	PE3A_TX_DN[0]	PCIEX3	O
AH39	PE3A_TX_DP[0]	PCIEX3	O
AH4	QPI1_DRX_DP[16]	Intel QPI	I
AH40	VSS	GND	
AH41	PE1B_TX_DN[6]	PCIEX3	O
AH42	PE1B_TX_DP[6]	PCIEX3	O
AH43	VSS	GND	
AH5	QPI1_DRX_DN[16]	Intel QPI	I
AH6	VSS	GND	
AH7	PROCHOT_N	ODCMOS	I/O
AH8	RSVD		
AH9	RSVD		
AJ1	VSS	GND	
AJ10	RSVD		
AJ11	VSS	GND	
AJ2	QPI1_DRX_DN[15]	Intel QPI	I
AJ3	QPI1_DRX_DP[15]	Intel QPI	I
AJ33	DMI_RX_DP[0]	PCIEX	I
AJ34	DMI_RX_DN[0]	PCIEX	I
AJ35	VSS	GND	
AJ36	DMI_TX_DN[3]	PCIEX	O

Table 8-2. Land Number (Sheet 7 of 37)

Land Number	Land Name	Buffer Type	Direction
AJ37	DMI_TX_DP[3]	PCIEX	O
AJ38	VSS	GND	
AJ39	PE3A_TX_DN[1]	PCIEX3	O
AJ4	VSS	GND	
AJ40	PE3A_TX_DP[1]	PCIEX3	O
AJ41	VSA	PWR	
AJ42	PE1B_TX_DN[7]	PCIEX3	O
AJ43	PE1B_TX_DP[7]	PCIEX3	O
AJ5	QPI1_DRX_DP[14]	Intel QPI	I
AJ6	QPI1_DRX_DN[14]	Intel QPI	I
AJ7	VTTD	PWR	
AJ8	RSVD		
AJ9	VSS	GND	
AK1	QPI1_DRX_DN[13]	Intel QPI	I
AK10	RSVD		
AK11	RSVD		
AK12	TDO	ODCMOS	O
AK13	QPI_RBIAS	Analog	I/O
AK14	VSS	GND	
AK15	VCC	PWR	
AK16	VCC	PWR	
AK17	VSS	GND	
AK18	VCC	PWR	
AK19	VCC	PWR	
AK2	QPI1_DRX_DP[13]	Intel QPI	I
AK20	VSS	GND	
AK21	VCC	PWR	
AK22	VCC	PWR	
AK23	VSS	GND	
AK24	VCC	PWR	
AK25	VCC	PWR	
AK26	VSS	GND	
AK27	VCC	PWR	
AK28	VCC	PWR	
AK29	VSS	GND	
AK3	VTTD	PWR	
AK30	VTTD	PWR	
AK31	VSS	GND	

Table 8-2. Land Number (Sheet 8 of 37)

Land Number	Land Name	Buffer Type	Direction
AK32	TEST4		I
AK33	VSS_VTTD_SENSE		O
AK34	VTTD_SENSE		O
AK35	DMI_TX_DN[2]	PCIEX	O
AK36	DMI_TX_DP[2]	PCIEX	O
AK37	VSS	GND	
AK38	PE3A_TX_DN[2]	PCIEX3	O
AK39	PE3A_TX_DP[2]	PCIEX3	O
AK4	QPI1_DRX_DP[12]	Intel QPI	I
AK40	VSS	GND	
AK41	PE3C_TX_DN[8]	PCIEX3	O
AK42	PE3C_TX_DP[8]	PCIEX3	O
AK43	VSS	GND	
AK5	QPI1_DRX_DN[12]	Intel QPI	I
AK6	PWRGOOD	CMOS	I
AK7	SVIDDATA	ODCMOS	I/O
AK8	RSVD		
AK9	RSVD		
AL1	VSS	GND	
AL10	BPM_N[0]	ODCMOS	I/O
AL11	BPM_N[1]	ODCMOS	I/O
AL12	TRST_N	CMOS	I
AL13	QPI_RBIAS_SENSE	Analog	I
AL14	VSS	GND	
AL15	VCC	PWR	
AL16	VCC	PWR	
AL17	VSS	GND	
AL18	VCC	PWR	
AL19	VCC	PWR	
AL2	QPI1_DRX_DN[11]	Intel QPI	I
AL20	VSS	GND	
AL21	VCC	PWR	
AL22	VCC	PWR	
AL23	VSS	GND	
AL24	VCC	PWR	
AL25	VCC	PWR	
AL26	VSS	GND	
AL27	VCC	PWR	

Table 8-2. Land Number (Sheet 9 of 37)

Land Number	Land Name	Buffer Type	Direction
AL28	VCC	PWR	
AL29	VSS	GND	
AL3	QPI1_DRX_DP[11]	Intel QPI	I
AL30	RSVD		
AL31	PE_RBIAS	PCIEX3	I/O
AL32	PE_RBIAS_SENSE	PCIEX3	I
AL33	ERROR_N[2]	ODCMOS	O
AL34	ERROR_N[0]	ODCMOS	O
AL35	VSS	GND	
AL36	DMI_TX_DN[1]	PCIEX	O
AL37	DMI_TX_DP[1]	PCIEX	O
AL38	VSA	PWR	
AL39	PE3A_TX_DN[3]	PCIEX3	O
AL4	VSS	GND	
AL40	PE3A_TX_DP[3]	PCIEX3	O
AL41	VSA	PWR	
AL42	PE3C_TX_DN[9]	PCIEX3	O
AL43	PE3C_TX_DP[9]	PCIEX3	O
AL5	QPI1_CLKRX_DN	Intel QPI	I
AL6	QPI1_CLKRX_DP	Intel QPI	I
AL7	SVIDALERT_N	CMOS	I
AL8	VSS	GND	
AL9	RSVD		
AM1	QPI1_DRX_DN[09]	Intel QPI	I
AM10	VSS	GND	
AM11	BPM_N[7]	ODCMOS	I/O
AM12	TDI	CMOS	I
AM13	QPI_VREF_CAP	Intel QPI	I/O
AM14	RSVD		
AM15	VCC	PWR	
AM16	VCC	PWR	
AM17	VSS	GND	
AM18	VCC	PWR	
AM19	VCC	PWR	
AM2	QPI1_DRX_DP[09]	Intel QPI	I
AM20	VSS	GND	
AM21	VCC	PWR	
AM22	VCC	PWR	

Table 8-2. Land Number (Sheet 10 of 37)

Land Number	Land Name	Buffer Type	Direction
AM23	VSS	GND	
AM24	VCC	PWR	
AM25	VCC	PWR	
AM26	VSS	GND	
AM27	VCC	PWR	
AM28	VCC	PWR	
AM29	VSS	GND	
AM3	CPU_ONLY_RESET	ODCMOS	I/O
AM30	BCLK1_DN	CMOS	I
AM31	SKTOCC_N		O
AM32	PE_VREF_CAP	PCIEX3	I/O
AM33	VSS	GND	
AM34	ERROR_N[1]	ODCMOS	O
AM35	DMI_TX_DN[0]	PCIEX	O
AM36	DMI_TX_DP[0]	PCIEX	O
AM37	VSA	PWR	
AM38	PE3B_TX_DN[4]	PCIEX3	O
AM39	PE3B_TX_DP[4]	PCIEX3	O
AM4	QPI1_DRX_DP[10]	Intel QPI	I
AM40	VSS	GND	
AM41	PE3C_TX_DN[10]	PCIEX3	O
AM42	PE3C_TX_DP[10]	PCIEX3	O
AM43	VSS	GND	
AM5	QPI1_DRX_DN[10]	Intel QPI	I
AM6	VSS	GND	
AM7	SVIDCLK	ODCMOS	O
AM8	RSVD		
AM9	RSVD		
AN1	VSS	GND	
AN10	BPM_N[6]	ODCMOS	I/O
AN11	BPM_N[3]	ODCMOS	I/O
AN12	VSS	GND	
AN13	VSS	GND	
AN14	RSVD		
AN15	VCC	PWR	
AN16	VCC	PWR	
AN17	VSS	GND	
AN18	VCC	PWR	

Table 8-2. Land Number (Sheet 11 of 37)

Land Number	Land Name	Buffer Type	Direction
AN19	VCC	PWR	
AN2	QPI1_DRX_DN[07]	Intel QPI	I
AN20	VSS	GND	
AN21	VCC	PWR	
AN22	VCC	PWR	
AN23	VSS	GND	
AN24	VCC	PWR	
AN25	VCC	PWR	
AN26	VSS	GND	
AN27	VCC	PWR	
AN28	VCC	PWR	
AN29	VSS	GND	
AN3	QPI1_DRX_DP[07]	Intel QPI	I
AN30	BCLK1_DP	CMOS	I
AN31	VSS	GND	
AN32	RSVD		
AN33	RSVD		
AN34	RSVD		
AN35	VTTA	PWR	
AN36	RSVD		
AN37	PECI	PECI	I/O
AN38	VSS	GND	
AN39	PE3B_TX_DN[5]	PCIEX3	O
AN4	VSS	GND	
AN40	PE3B_TX_DP[5]	PCIEX3	O
AN41	VSA	PWR	
AN42	PE3C_TX_DN[11]	PCIEX3	O
AN43	PE3C_TX_DP[11]	PCIEX3	O
AN5	QPI1_DRX_DP[08]	Intel QPI	I
AN6	QPI1_DRX_DN[08]	Intel QPI	I
AN7	VTTD	PWR	
AN8	TXT_PLTEN	CMOS	I
AN9	BPM_N[2]	ODCMOS	I/O
AP1	QPI1_DRX_DN[05]	Intel QPI	I
AP10	BPM_N[4]	ODCMOS	I/O
AP11	BPM_N[5]	ODCMOS	I/O
AP12	TMS	CMOS	I
AP13	BCLK0_DN	CMOS	I

Table 8-2. Land Number (Sheet 12 of 37)

Land Number	Land Name	Buffer Type	Direction
AP14	VSS	GND	
AP15	VCC	PWR	
AP16	VCC	PWR	
AP17	VSS	GND	
AP18	VCC	PWR	
AP19	VCC	PWR	
AP2	QPI1_DRX_DP[05]	Intel QPI	I
AP20	VSS	GND	
AP21	VCC	PWR	
AP22	VCC	PWR	
AP23	VSS	GND	
AP24	VCC	PWR	
AP25	VCC	PWR	
AP26	VSS	GND	
AP27	VCC	PWR	
AP28	VCC	PWR	
AP29	VSS	GND	
AP3	VTTD	PWR	
AP30	PEHPSDA	ODCMOS	I/O
AP31	RSVD		
AP32	RSVD		
AP33	VTTA	PWR	
AP34	VTTA	PWR	
AP35	VTTA	PWR	
AP36	RSVD		
AP37	RSVD		
AP38	PE3B_TX_DN[6]	PCIEX3	O
AP39	PE3B_TX_DP[6]	PCIEX3	O
AP4	QPI1_DRX_DN[04]	Intel QPI	I
AP40	VSS	GND	
AP41	PE3D_TX_DN[12]	PCIEX3	O
AP42	PE3D_TX_DP[12]	PCIEX3	O
AP43	VSS	GND	
AP5	QPI1_DRX_DP[04]	Intel QPI	I
AP6	VSS	GND	
AP7	QPI1_DRX_DP[06]	Intel QPI	I
AP8	QPI1_DRX_DN[06]	Intel QPI	I
AP9	VSS	GND	

Table 8-2. Land Number (Sheet 13 of 37)

Land Number	Land Name	Buffer Type	Direction
AR1	VSS	GND	
AR10	PREQ_N	CMOS	I/O
AR11	PRDY_N	CMOS	O
AR12	TCK	CMOS	I
AR13	BCLK0_DP	CMOS	I
AR14	VSS	GND	
AR15	VCC	PWR	
AR16	VCC	PWR	
AR17	VSS	GND	
AR18	VCC	PWR	
AR19	VCC	PWR	
AR2	QPI1_DRX_DN[03]	Intel QPI	I
AR20	VSS	GND	
AR21	VCC	PWR	
AR22	VCC	PWR	
AR23	VSS	GND	
AR24	VCC	PWR	
AR25	VCC	PWR	
AR26	VSS	GND	
AR27	VCC	PWR	
AR28	VCC	PWR	
AR29	VSS	GND	
AR3	QPI1_DRX_DP[03]	Intel QPI	I
AR30	PEHPSCL	ODCMOS	I/O
AR31	RSVD		
AR32	VSS	GND	
AR33	PE3A_RX_DP[3]	PCIEX3	I
AR34	VSS	GND	
AR35	PE3A_RX_DP[1]	PCIEX3	I
AR36	VCCPLL	PWR	
AR37	PE1A_RX_DP[3]	PCIEX3	I
AR38	VSS	GND	
AR39	PE3B_TX_DN[7]	PCIEX3	O
AR4	VSS	GND	
AR40	PE3B_TX_DP[7]	PCIEX3	O
AR41	VSA	PWR	
AR42	PE3D_TX_DN[13]	PCIEX3	O
AR43	PE3D_TX_DP[13]	PCIEX3	O

Table 8-2. Land Number (Sheet 14 of 37)

Land Number	Land Name	Buffer Type	Direction
AR5	QPI1_DRX_DN[02]	Intel QPI	I
AR6	QPI1_DRX_DP[02]	Intel QPI	I
AR7	RSVD		
AR8	RSVD		
AR9	RSVD		
AT1	QPI1_DRX_DN[01]	Intel QPI	I
AT10	QPI1_DTX_DN[08]	Intel QPI	O
AT11	VSS	GND	
AT12	QPI1_DTX_DN[04]	Intel QPI	O
AT13	VSS	GND	
AT14	RSVD		
AT15	VCC	PWR	
AT16	VCC	PWR	
AT17	VSS	GND	
AT18	VCC	PWR	
AT19	VCC	PWR	
AT2	QPI1_DRX_DP[01]	Intel QPI	I
AT20	VSS	GND	
AT21	VCC	PWR	
AT22	VCC	PWR	
AT23	VSS	GND	
AT24	VCC	PWR	
AT25	VCC	PWR	
AT26	VSS	GND	
AT27	VCC	PWR	
AT28	VCC	PWR	
AT29	VSS	GND	
AT3	PMSYNC	CMOS	I
AT30	PE3B_RX_DP[6]	PCIEX3	I
AT31	VSS	GND	
AT32	PE3B_RX_DP[4]	PCIEX3	I
AT33	PE3A_RX_DN[3]	PCIEX3	I
AT34	PE3A_RX_DP[2]	PCIEX3	I
AT35	PE3A_RX_DN[1]	PCIEX3	I
AT36	PE3A_RX_DP[0]	PCIEX3	I
AT37	PE1A_RX_DN[3]	PCIEX3	I
AT38	PE1A_RX_DP[2]	PCIEX3	I
AT39	RSVD		

Table 8-2. Land Number (Sheet 15 of 37)

Land Number	Land Name	Buffer Type	Direction
AT4	QPI1_DRX_DN[00]	Intel QPI	I
AT40	VSA	PWR	
AT41	PE3D_TX_DN[14]	PCIEX3	O
AT42	PE3D_TX_DP[14]	PCIEX3	O
AT43	VSS	GND	
AT5	QPI1_DRX_DP[00]	Intel QPI	I
AT6	CAT_ERR_N	ODCMOS	I/O
AT7	VTTA	PWR	
AT8	QPI1_DTX_DN[12]	Intel QPI	O
AT9	VSS	GND	
AU1	VSS	GND	
AU10	QPI1_DTX_DP[08]	Intel QPI	O
AU11	QPI1_DTX_DN[06]	Intel QPI	O
AU12	QPI1_DTX_DP[04]	Intel QPI	O
AU13	QPI1_DTX_DN[02]	Intel QPI	O
AU14	RSVD		
AU15	VCC	PWR	
AU16	VCC	PWR	
AU17	VSS	GND	
AU18	VCC	PWR	
AU19	VCC	PWR	
AU2	RESET_N	CMOS	I
AU20	VSS	GND	
AU21	VCC	PWR	
AU22	VCC	PWR	
AU23	VSS	GND	
AU24	VCC	PWR	
AU25	VCC	PWR	
AU26	VSS	GND	
AU27	VCC	PWR	
AU28	VCC	PWR	
AU29	PE3B_RX_DP[7]	PCIEX3	I
AU3	RSVD		
AU30	PE3B_RX_DN[6]	PCIEX3	I
AU31	PE3B_RX_DP[5]	PCIEX3	I
AU32	PE3B_RX_DN[4]	PCIEX3	I
AU33	VSS	GND	
AU34	PE3A_RX_DN[2]	PCIEX3	I

Table 8-2. Land Number (Sheet 16 of 37)

Land Number	Land Name	Buffer Type	Direction
AU35	VSS	GND	
AU36	PE3A_RX_DN[0]	PCIEX3	I
AU37	VSS	GND	
AU38	PE1A_RX_DN[2]	PCIEX3	I
AU39	PE1A_RX_DP[1]	PCIEX3	I
AU4	RSVD		
AU40	SAFE_MODE_BOOT	CMOS	I
AU41	VSS	GND	
AU42	PE3D_TX_DN[15]	PCIEX3	O
AU43	PE3D_TX_DP[15]	PCIEX3	O
AU5	VSS	GND	
AU6	RSVD		
AU7	QPI1_DTX_DN[14]	Intel QPI	O
AU8	QPI1_DTX_DP[12]	Intel QPI	O
AU9	QPI1_DTX_DN[10]	Intel QPI	O
AV1	QPI1_DTX_DN[19]	Intel QPI	O
AV10	VSS	GND	
AV11	QPI1_DTX_DP[06]	Intel QPI	O
AV12	VSS	GND	
AV13	QPI1_DTX_DP[02]	Intel QPI	O
AV14	VSS	GND	
AV15	VCC	PWR	
AV16	VCC	PWR	
AV17	VSS	GND	
AV18	VCC	PWR	
AV19	VCC	PWR	
AV2	QPI1_DTX_DP[19]	Intel QPI	O
AV20	VSS	GND	
AV21	VCC	PWR	
AV22	VCC	PWR	
AV23	VSS	GND	
AV24	VCC	PWR	
AV25	VCC	PWR	
AV26	VSS	GND	
AV27	VCC	PWR	
AV28	VCC	PWR	
AV29	PE3B_RX_DN[7]	PCIEX3	I
AV3	VSS	GND	

Table 8-2. Land Number (Sheet 17 of 37)

Land Number	Land Name	Buffer Type	Direction
AV30	VTTD	PWR	
AV31	PE3B_RX_DN[5]	PCIEX3	I
AV32	VTTD	PWR	
AV33	PE3C_RX_DP[11]	PCIEX3	I
AV34	VTTA	PWR	
AV35	PE3C_RX_DP[9]	PCIEX3	I
AV36	VCCPLL	PWR	
AV37	PE1B_RX_DP[7]	PCIEX3	I
AV38	RSVD		
AV39	PE1A_RX_DN[1]	PCIEX3	I
AV4	QPI1_DTX_DN[18]	Intel QPI	O
AV40	VSA	PWR	
AV41	PE1A_RX_DP[0]	PCIEX3	I
AV42	RSVD		
AV43	VSA	PWR	
AV5	QPI1_DTX_DP[18]	Intel QPI	O
AV6	VSS	GND	
AV7	QPI1_DTX_DP[14]	Intel QPI	O
AV8	VSS	GND	
AV9	QPI1_DTX_DP[10]	Intel QPI	O
AW1	VSS	GND	
AW10	QPI1_DTX_DN[07]	Intel QPI	O
AW11	SOCKET_ID[0]	CMOS	I
AW12	QPI1_DTX_DN[03]	Intel QPI	O
AW13	RSVD		
AW14	QPI1_DTX_DN[00]	Intel QPI	O
AW15	VCC	PWR	
AW16	VCC	PWR	
AW17	VSS	GND	
AW18	VCC	PWR	
AW19	VCC	PWR	
AW2	QPI1_DTX_DN[17]	Intel QPI	O
AW20	VSS	GND	
AW21	VCC	PWR	
AW22	VCC	PWR	
AW23	VSS	GND	
AW24	VCC	PWR	
AW25	VCC	PWR	

Table 8-2. Land Number (Sheet 18 of 37)

Land Number	Land Name	Buffer Type	Direction
AW26	VSS	GND	
AW27	VCC	PWR	
AW28	VCC	PWR	
AW29	VSS	GND	
AW3	QPI1_DTX_DP[17]	Intel QPI	O
AW30	PE3D_RX_DP[14]	PCIEX3	I
AW31	VSS	GND	
AW32	PE3D_RX_DP[12]	PCIEX3	I
AW33	PE3C_RX_DN[11]	PCIEX3	I
AW34	PE3C_RX_DP[10]	PCIEX3	I
AW35	PE3C_RX_DN[9]	PCIEX3	I
AW36	PE3C_RX_DP[8]	PCIEX3	I
AW37	PE1B_RX_DN[7]	PCIEX3	I
AW38	PE1B_RX_DP[6]	PCIEX3	I
AW39	VSS	GND	
AW4	VSS	GND	
AW40	PE1B_RX_DP[4]	PCIEX3	I
AW41	PE1A_RX_DN[0]	PCIEX3	I
AW42	VSS	GND	
AW5	QPI1_DTX_DN[16]	Intel QPI	O
AW6	QPI1_DTX_DP[16]	Intel QPI	O
AW7	VTTA	PWR	
AW8	QPI1_DTX_DN[09]	Intel QPI	O
AW9	VTTA	PWR	
AY10	QPI1_DTX_DP[07]	Intel QPI	O
AY11	QPI1_DTX_DN[05]	Intel QPI	O
AY12	QPI1_DTX_DP[03]	Intel QPI	O
AY13	QPI1_DTX_DN[01]	Intel QPI	O
AY14	QPI1_DTX_DP[00]	Intel QPI	O
AY15	VCC	PWR	
AY16	VCC	PWR	
AY17	VSS	GND	
AY18	VCC	PWR	
AY19	VCC	PWR	
AY2	VTTA	PWR	
AY20	VSS	GND	
AY24	VCC	PWR	
AY25	VCC	PWR	

Table 8-2. Land Number (Sheet 19 of 37)

Land Number	Land Name	Buffer Type	Direction
AY26	VSS	GND	
AY27	VCC	PWR	
AY28	VCC	PWR	
AY29	PE3D_RX_DP[15]	PCIEX3	I
AY3	QPI1_DTX_DN[15]	Intel QPI	O
AY30	PE3D_RX_DN[14]	PCIEX3	I
AY31	PE3D_RX_DP[13]	PCIEX3	I
AY32	PE3D_RX_DN[12]	PCIEX3	I
AY33	VSS	GND	
AY34	PE3C_RX_DN[10]	PCIEX3	I
AY35	VSS	GND	
AY36	PE3C_RX_DN[8]	PCIEX3	I
AY37	VSS	GND	
AY38	PE1B_RX_DN[6]	PCIEX3	I
AY39	PE1B_RX_DP[5]	PCIEX3	I
AY4	QPI1_DTX_DP[15]	Intel QPI	O
AY40	PE1B_RX_DN[4]	PCIEX3	I
AY41	RSVD		
AY42	VSA	PWR	
AY5	VSS	GND	
AY6	RSVD		
AY7	QPI1_DTX_DN[11]	Intel QPI	O
AY8	QPI1_DTX_DP[09]	Intel QPI	O
AY9	QPI1_CLKTX_DN	Intel QPI	O
B10	DDR3_DQ[37]	SSTL	I/O
B11	DDR3_CS_N[2]	SSTL	O
B12	VSS	GND	
B13	DDR3_CS_N[1]	SSTL	O
B14	DDR3_CS_N[5]	SSTL	O
B15	DDR3_CAS_N	SSTL	O
B16	DDR3_BA[0]	SSTL	O
B17	VSS	GND	
B18	DDR3_CLK_DP[2]	SSTL	O
B19	DDR3_CLK_DN[0]	SSTL	O
B2	RSVD		
B20	DDR3_CLK_DP[1]	SSTL	O
B24	DDR3_MA[08]	SSTL	O
B25	DDR3_MA[09]	SSTL	O

Table 8-2. Land Number (Sheet 20 of 37)

Land Number	Land Name	Buffer Type	Direction
B26	DDR3_MA[12]	SSTL	O
B27	VSS	GND	
B28	DDR3_CKE[0]	SSTL	O
B29	DDR3_CKE[1]	SSTL	O
B3	VSS	GND	
B30	DDR3_ECC[2]	SSTL	I/O
B31	DDR3_ECC[6]	SSTL	I/O
B32	DDR3_DQS_DN[08]	SSTL	I/O
B33	DDR3_DQS_DN[17]	SSTL	I/O
B34	DDR3_DQS_DP[17]	SSTL	I/O
B35	DDR3_ECC[5]	SSTL	I/O
B36	VSS	GND	
B37	DDR2_DQ[31]	SSTL	I/O
B38	DDR2_DQS_DP[03]	SSTL	I/O
B39	DDR2_DQS_DP[12]	SSTL	I/O
B4	DDR3_DQ[35]	SSTL	I/O
B40	VSS	GND	
B41	DDR3_DQ[27]	SSTL	I/O
B42	VSS	GND	
B5	DDR3_DQ[34]	SSTL	I/O
B6	DDR3_DQ[39]	SSTL	I/O
B7	DDR3_DQS_DP[04]	SSTL	I/O
B8	DDR3_DQS_DN[13]	SSTL	I/O
B9	DDR3_DQ[33]	SSTL	I/O
BA10	VSS	GND	
BA11	QPI1_DTX_DP[05]	Intel QPI	O
BA12	VSS	GND	
BA13	QPI1_DTX_DP[01]	Intel QPI	O
BA14	VSS	GND	
BA15	VCC	PWR	
BA16	VCC	PWR	
BA17	VSS	GND	
BA18	VCC	PWR	
BA19	VCC	PWR	
BA20	VSS	GND	
BA24	VCC	PWR	
BA25	VCC	PWR	
BA26	VSS	GND	

Table 8-2. Land Number (Sheet 21 of 37)

Land Number	Land Name	Buffer Type	Direction
BA27	VCC	PWR	
BA28	VCC	PWR	
BA29	PE3D_RX_DN[15]	PCIEX3	I
BA3	VSS	GND	
BA30	VTTD	PWR	
BA31	PE3D_RX_DN[13]	PCIEX3	I
BA32	VSS	GND	
BA38	RSVD		
BA39	PE1B_RX_DN[5]	PCIEX3	I
BA4	QPI1_DTX_DN[13]	Intel QPI	O
BA40	VSS	GND	
BA5	QPI1_DTX_DP[13]	Intel QPI	O
BA6	RSVD		
BA7	QPI1_DTX_DP[11]	Intel QPI	O
BA8	VSS	GND	
BA9	QPI1_CLKTX_DP	Intel QPI	O
C10	VSS	GND	
C11	DDR3_CS_N[6]	SSTL	O
C12	DDR3_ODT[1]	SSTL	O
C13	DDR3_ODT[3]	SSTL	O
C14	DDR3_ODT[2]	SSTL	O
C15	VSS	GND	
C16	DDR3_WE_N	SSTL	O
C17	DDR3_MA[00]	SSTL	O
C18	DDR3_MA_PAR	SSTL	O
C19	DDR3_CLK_DP[0]	SSTL	O
C2	VSS	GND	
C20	DDR3_CLK_DP[3]	SSTL	O
C21	DDR3_CLK_DN[3]	SSTL	O
C22	DDR3_MA[03]	SSTL	O
C23	DDR3_MA[06]	SSTL	O
C24	DDR3_MA[05]	SSTL	O
C25	VSS	GND	
C26	DDR3_MA[14]	SSTL	O
C27	DDR3_BA[2]	SSTL	O
C28	DDR3_CKE[3]	SSTL	O
C29	DDR_RESET_C23_N	CMOS_1.5 V	O
C3	DDR3_DQ[44]	SSTL	I/O

Table 8-2. Land Number (Sheet 22 of 37)

Land Number	Land Name	Buffer Type	Direction
C30	VSS	GND	
C31	VSS	GND	
C32	VSS	GND	
C33	DDR3_ECC[1]	SSTL	I/O
C34	DDR3_ECC[0]	SSTL	I/O
C35	DDR3_ECC[4]	SSTL	I/O
C36	VSS	GND	
C37	DDR2_DQ[26]	SSTL	I/O
C38	DDR2_DQS_DN[03]	SSTL	I/O
C39	DDR2_DQ[25]	SSTL	I/O
C4	TEST1		O
C40	VSS	GND	
C41	DDR3_DQ[26]	SSTL	I/O
C42	DDR3_DQ[30]	SSTL	I/O
C43	VSS	GND	
C5	VSS	GND	
C6	VSS	GND	
C7	VSS	GND	
C8	VSS	GND	
C9	VSS	GND	
D1	VSS	GND	
D10	DDR2_DQ[36]	SSTL	I/O
D11	DDR3_CS_N[3]	SSTL	O
D12	DDR3_CS_N[7]	SSTL	O
D13	VSS	GND	
D14	DDR3_MA[13]	SSTL	O
D15	DDR3_ODT[0]	SSTL	O
D16	DDR3_CS_N[0]	SSTL	O
D17	DDR3_BA[1]	SSTL	O
D18	VSS	GND	
D19	DDR2_MA[00]	SSTL	O
D2	DDR3_DQ[41]	SSTL	I/O
D20	VSS	GND	
D21	DDR3_MA[01]	SSTL	O
D22	DDR3_MA[02]	SSTL	O
D23	DDR3_MA[04]	SSTL	O
D24	DDR2_MA[08]	SSTL	O
D25	DDR3_MA[07]	SSTL	O

Table 8-2. Land Number (Sheet 23 of 37)

Land Number	Land Name	Buffer Type	Direction
D26	DDR2_MA[14]	SSTL	O
D27	DDR2_CKE[0]	SSTL	O
D28	DDR2_CKE[3]	SSTL	O
D29	DDR2_ECC[3]	SSTL	I/O
D3	DDR3_DQ[45]	SSTL	I/O
D30	DDR2_ECC[7]	SSTL	I/O
D31	DDR2_DQS_DP[08]	SSTL	I/O
D32	DDR2_DQS_DN[17]	SSTL	I/O
D33	VSS	GND	
D34	VSS	GND	
D35	VSS	GND	
D36	VSS	GND	
D37	DDR2_DQ[30]	SSTL	I/O
D38	VSS	GND	
D39	DDR2_DQ[24]	SSTL	I/O
D4	VSS	GND	
D40	VSS	GND	
D41	DDR3_DQ[31]	SSTL	I/O
D42	DDR3_DQS_DP[03]	SSTL	I/O
D43	DDR3_DQS_DN[03]	SSTL	I/O
D5	DDR2_DQ[35]	SSTL	I/O
D6	DDR2_DQ[38]	SSTL	I/O
D7	DDR2_DQS_DN[04]	SSTL	I/O
D8	DDR2_DQS_DP[13]	SSTL	I/O
D9	DDR2_DQ[32]	SSTL	I/O
E1	DDR3_DQS_DN[14]	SSTL	I/O
E10	DDR2_DQ[37]	SSTL	I/O
E11	VSS	GND	
E12	DDR2_ODT[3]	SSTL	O
E13	DDR2_MA[13]	SSTL	O
E14	DDR2_CS_N[5]	SSTL	O
E15	DDR2_CS_N[4]	SSTL	O
E16	VSS	GND	
E17	DDR2_BA[0]	SSTL	O
E18	DDR2_MA_PAR	SSTL	O
E19	DDR2_CLK_DN[0]	SSTL	O
E2	DDR3_DQS_DP[14]	SSTL	I/O
E20	DDR2_CLK_DP[0]	SSTL	O

Table 8-2. Land Number (Sheet 24 of 37)

Land Number	Land Name	Buffer Type	Direction
E21	VCCD	PWR	
E22	DDR2_MA[01]	SSTL	O
E23	VSS	GND	
E24	DDR2_MA[06]	SSTL	O
E25	DDR2_MA[11]	SSTL	O
E26	VCCD	PWR	
E27	DDR2_CKE[2]	SSTL	O
E28	DDR2_CKE[1]	SSTL	O
E29	DDR2_ECC[2]	SSTL	I/O
E3	DDR3_DQ[40]	SSTL	I/O
E30	DDR2_ECC[6]	SSTL	I/O
E31	DDR2_DQS_DN[08]	SSTL	I/O
E32	DDR2_DQS_DP[17]	SSTL	I/O
E33	DDR2_ECC[0]	SSTL	I/O
E34	DDR2_ECC[5]	SSTL	I/O
E35	DDR2_ECC[4]	SSTL	I/O
E36	DDR23_RCOMP[1]	Analog	I
E37	DDR2_DQ[27]	SSTL	I/O
E38	DDR2_DQ[28]	SSTL	I/O
E39	DDR2_DQ[29]	SSTL	I/O
E4	VSS	GND	
E40	VSS	GND	
E41	DDR3_DQ[25]	SSTL	I/O
E42	DDR3_DQS_DN[12]	SSTL	I/O
E43	DDR3_DQS_DP[12]	SSTL	I/O
E5	DDR2_DQ[34]	SSTL	I/O
E6	DDR2_DQ[39]	SSTL	I/O
E7	DDR2_DQS_DP[04]	SSTL	I/O
E8	DDR2_DQS_DN[13]	SSTL	I/O
E9	DDR2_DQ[33]	SSTL	I/O
F1	DDR3_DQS_DP[05]	SSTL	I/O
F10	VSS	GND	
F11	DDR2_CS_N[2]	SSTL	O
F12	DDR2_ODT[1]	SSTL	O
F13	DDR2_CAS_N	SSTL	O
F14	VCCD	PWR	
F15	DDR2_WE_N	SSTL	O
F16	DDR2_ODT[0]	SSTL	O

Table 8-2. Land Number (Sheet 25 of 37)

Land Number	Land Name	Buffer Type	Direction
F17	DDR2_RAS_N	SSTL	O
F18	DDR2_MA[10]	SSTL	O
F19	VCCD	PWR	
F2	DDR3_DQS_DN[05]	SSTL	I/O
F20	DDR2_CLK_DP[2]	SSTL	O
F21	DDR2_CLK_DN[3]	SSTL	O
F22	DDR2_CLK_DP[3]	SSTL	O
F23	DDR2_MA[04]	SSTL	O
F24	VCCD	PWR	
F25	DDR2_MA[12]	SSTL	O
F26	DDR2_BA[2]	SSTL	O
F27	VSS	GND	
F28	VSS	GND	
F29	VSS	GND	
F3	DDR3_DQ[46]	SSTL	I/O
F30	VSS	GND	
F31	VSS	GND	
F32	VSS	GND	
F33	VSS	GND	
F34	DDR2_ECC[1]	SSTL	I/O
F35	VSS	GND	
F36	VSS	GND	
F37	VSS	GND	
F38	VSS	GND	
F39	VSS	GND	
F4	VSS	GND	
F40	VSS	GND	
F41	DDR3_DQ[24]	SSTL	I/O
F42	DDR3_DQ[28]	SSTL	I/O
F43	DDR3_DQ[29]	SSTL	I/O
F5	VSS	GND	
F6	VSS	GND	
F7	VSS	GND	
F8	VSS	GND	
F9	VSS	GND	
G1	DDR3_DQ[42]	SSTL	I/O
G10	DDR2_DQ[44]	SSTL	I/O
G11	DDR2_CS_N[6]	SSTL	O

Table 8-2. Land Number (Sheet 26 of 37)

Land Number	Land Name	Buffer Type	Direction
G12	VCCD	PWR	
G13	DDR1_ODT[3]	SSTL	O
G14	DDR2_CS_N[1]	SSTL	O
G15	DDR1_ODT[0]	SSTL	O
G16	DDR2_CS_N[0]	SSTL	O
G17	VCCD	PWR	
G18	DDR2_BA[1]	SSTL	O
G19	DDR1_CLK_DN[1]	SSTL	O
G2	DDR3_DQ[43]	SSTL	I/O
G20	DDR2_CLK_DN[2]	SSTL	O
G21	DDR2_CLK_DP[1]	SSTL	O
G22	VCCD	PWR	
G23	DDR2_MA[03]	SSTL	O
G24	DDR2_MA[07]	SSTL	O
G25	DDR2_PAR_ERR_N	SSTL	I
G26	DDR2_MA[15]	SSTL	O
G27	DDR1_ECC[3]	SSTL	I/O
G28	DDR1_ECC[7]	SSTL	I/O
G29	DDR1_DQS_DP[08]	SSTL	I/O
G3	DDR3_DQ[47]	SSTL	I/O
G30	DDR1_DQS_DN[17]	SSTL	I/O
G31	DDR1_ECC[1]	SSTL	I/O
G32	DDR1_ECC[0]	SSTL	I/O
G33	VSS	GND	
G34	VSS	GND	
G35	VSS	GND	
G36	DDR2_DQ[19]	SSTL	I/O
G37	DDR2_DQ[23]	SSTL	I/O
G38	DDR2_DQS_DP[02]	SSTL	I/O
G39	DDR2_DQS_DN[02]	SSTL	I/O
G4	VSS	GND	
G40	TEST3		O
G41	VSS	GND	
G42	VSS	GND	
G43	VSS	GND	
G5	DDR2_DQ[42]	SSTL	I/O
G6	DDR2_DQ[46]	SSTL	I/O
G7	DDR2_DQS_DN[05]	SSTL	I/O

Table 8-2. Land Number (Sheet 27 of 37)

Land Number	Land Name	Buffer Type	Direction
G8	DDR2_DQS_DP[14]	SSTL	I/O
G9	DDR2_DQ[40]	SSTL	I/O
H1	VSS	GND	
H10	DDR2_DQ[45]	SSTL	I/O
H11	DDR2_CS_N[7]	SSTL	O
H12	DDR2_ODT[2]	SSTL	O
H13	DDR1_ODT[1]	SSTL	O
H14	DDR1_CS_N[5]	SSTL	O
H15	VCCD	PWR	
H16	DDR1_CLK_DN[2]	SSTL	O
H17	DDR1_CLK_DN[0]	SSTL	O
H18	DDR1_CLK_DN[3]	SSTL	O
H19	DDR1_CLK_DP[1]	SSTL	O
H2	VSS	GND	
H20	VCCD	PWR	
H21	DDR2_CLK_DN[1]	SSTL	O
H22	DDR2_MA[02]	SSTL	O
H23	DDR2_MA[05]	SSTL	O
H24	DDR2_MA[09]	SSTL	O
H25	VCCD	PWR	
H26	DDR1_CKE[3]	SSTL	O
H27	DDR1_ECC[2]	SSTL	I/O
H28	DDR1_ECC[6]	SSTL	I/O
H29	DDR1_DQS_DN[08]	SSTL	I/O
H3	VSS	GND	
H30	DDR1_DQS_DP[17]	SSTL	I/O
H31	DDR1_ECC[4]	SSTL	I/O
H32	DDR1_ECC[5]	SSTL	I/O
H33	DDR1_DQS_DN[02]	SSTL	I/O
H34	DDR1_DQS_DP[02]	SSTL	I/O
H35	DDR23_RCOMP[0]	Analog	I
H36	DDR2_DQ[18]	SSTL	I/O
H37	DDR2_DQ[22]	SSTL	I/O
H38	DDR2_DQS_DN[11]	SSTL	I/O
H39	DDR2_DQS_DP[11]	SSTL	I/O
H4	TEST0		O
H40	VSS	GND	
H41	DDR3_DQ[19]	SSTL	I/O

Table 8-2. Land Number (Sheet 28 of 37)

Land Number	Land Name	Buffer Type	Direction
H42	DDR3_DQ[23]	SSTL	I/O
H43	DDR3_DQ[18]	SSTL	I/O
H5	DDR2_DQ[43]	SSTL	I/O
H6	DDR2_DQ[47]	SSTL	I/O
H7	DDR2_DQS_DP[05]	SSTL	I/O
H8	DDR2_DQS_DN[14]	SSTL	I/O
H9	DDR2_DQ[41]	SSTL	I/O
J1	DDR3_DQ[53]	SSTL	I/O
J10	VSS	GND	
J11	DDR2_CS_N[3]	SSTL	O
J12	DDR1_CS_N[7]	SSTL	O
J13	VCCD	PWR	
J14	DDR1_ODT[2]	SSTL	O
J15	DDR1_CS_N[0]	SSTL	O
J16	DDR1_CLK_DP[2]	SSTL	O
J17	DDR1_CLK_DP[0]	SSTL	O
J18	DDR1_CLK_DP[3]	SSTL	O
J19	DDR1_MA_PAR	SSTL	O
J2	DDR3_DQ[48]	SSTL	I/O
J20	DDR1_MA[12]	SSTL	O
J21	DDR1_PAR_ERR_N	SSTL	I
J22	DDR1_MA[14]	SSTL	O
J23	VSS	GND	
J24	DDR1_BA[2]	SSTL	O
J25	DDR1_CKE[2]	SSTL	O
J26	DDR1_CKE[1]	SSTL	O
J27	VSS	GND	
J28	VSS	GND	
J29	VSS	GND	
J3	DDR3_DQ[52]	SSTL	I/O
J30	VSS	GND	
J31	VSS	GND	
J32	VSS	GND	
J33	DDR1_DQ[23]	SSTL	I/O
J34	DDR1_DQ[22]	SSTL	I/O
J35	VSS	GND	
J36	VSS	GND	
J37	VSS	GND	

Table 8-2. Land Number (Sheet 29 of 37)

Land Number	Land Name	Buffer Type	Direction
J38	DDR2_DQ[16]	SSTL	I/O
J39	DDR2_DQ[17]	SSTL	I/O
J4	DDR1_RCOMP[0]	Analog	I
J40	VSS	GND	
J41	DDR3_DQS_DP[02]	SSTL	I/O
J42	VSS	GND	
J43	DDR3_DQ[22]	SSTL	I/O
J5	VSS	GND	
J6	VSS	GND	
J7	VSS	GND	
J8	VSS	GND	
J9	VSS	GND	
K1	DDR3_DQ[49]	SSTL	I/O
K10	DDR1_DQ[36]	SSTL	I/O
K11	VCCD	PWR	
K12	DDR1_CS_N[3]	SSTL	O
K13	DDR1_CAS_N	SSTL	O
K14	DDR1_WE_N	SSTL	O
K15	DDR1_CS_N[4]	SSTL	O
K16	VCCD	PWR	
K17	DDR1_RAS_N	SSTL	O
K18	VCCD	PWR	
K19	DDR1_MA[00]	SSTL	O
K2	VSS	GND	
K20	DDR1_MA[02]	SSTL	O
K21	VSS	GND	
K22	DDR1_MA[05]	SSTL	O
K23	DDR1_MA[08]	SSTL	O
K24	DDR1_MA[15]	SSTL	O
K25	DDR1_CKE[0]	SSTL	O
K26	VCCD	PWR	
K27	DDR1_DQ[27]	SSTL	I/O
K28	DDR1_DQ[31]	SSTL	I/O
K29	DDR1_DQS_DP[03]	SSTL	I/O
K3	DDR3_DQS_DP[15]	SSTL	I/O
K30	DDR1_DQS_DN[12]	SSTL	I/O
K31	DDR1_DQ[25]	SSTL	I/O
K32	DDR1_DQ[29]	SSTL	I/O

Table 8-2. Land Number (Sheet 30 of 37)

Land Number	Land Name	Buffer Type	Direction
K33	DDR1_DQ[18]	SSTL	I/O
K34	DDR1_DQS_DN[11]	SSTL	I/O
K35	DDR1_DQ[16]	SSTL	I/O
K36	DDR1_DQ[17]	SSTL	I/O
K37	VSS	GND	
K38	DDR2_DQ[20]	SSTL	I/O
K39	DDR2_DQ[21]	SSTL	I/O
K4	VSS	GND	
K40	VSS	GND	
K41	DDR3_DQS_DN[02]	SSTL	I/O
K42	DDR3_DQS_DP[11]	SSTL	I/O
K43	DDR3_DQS_DN[11]	SSTL	I/O
K5	DDR1_DQ[34]	SSTL	I/O
K6	DDR1_DQ[38]	SSTL	I/O
K7	DDR1_DQS_DN[04]	SSTL	I/O
K8	DDR1_DQS_DP[13]	SSTL	I/O
K9	DDR1_DQ[32]	SSTL	I/O
L1	DDR3_DQS_DN[06]	SSTL	I/O
L10	DDR1_DQ[37]	SSTL	I/O
L11	DDR1_CS_N[2]	SSTL	O
L12	DDR1_CS_N[6]	SSTL	O
L13	DDR1_MA[13]	SSTL	O
L14	VSS	GND	
L15	DDR1_CS_N[1]	SSTL	O
L16	RSVD		
L17	DDR1_BA[0]	SSTL	O
L18	DDR1_BA[1]	SSTL	O
L19	VSS	GND	
L2	DDR3_DQS_DP[06]	SSTL	I/O
L20	DDR1_MA[01]	SSTL	O
L21	DDR1_MA[03]	SSTL	O
L22	DDR1_MA[06]	SSTL	O
L23	DDR1_MA[07]	SSTL	O
L24	VSS	GND	
L25	DDR1_MA[09]	SSTL	O
L26	DDR_RESET_C1_N	CMOS_1.5 V	O
L27	DDR1_DQ[26]	SSTL	I/O
L28	DDR1_DQ[30]	SSTL	I/O

Table 8-2. Land Number (Sheet 31 of 37)

Land Number	Land Name	Buffer Type	Direction
L29	DDR1_DQS_DN[03]	SSTL	I/O
L3	DDR3_DQS_DN[15]	SSTL	I/O
L30	DDR1_DQS_DP[12]	SSTL	I/O
L31	DDR1_DQ[24]	SSTL	I/O
L32	DDR1_DQ[28]	SSTL	I/O
L33	DDR1_DQ[19]	SSTL	I/O
L34	DDR1_DQS_DP[11]	SSTL	I/O
L35	DDR1_DQ[21]	SSTL	I/O
L36	DDR1_DQ[20]	SSTL	I/O
L37	VSS	GND	
L38	DDR23_RCOMP[2]	Analog	I
L39	VSS	GND	
L4	VSS	GND	
L40	VSS	GND	
L41	DDR3_DQ[21]	SSTL	I/O
L42	DDR3_DQ[17]	SSTL	I/O
L43	DDR3_DQ[16]	SSTL	I/O
L5	DDR1_DQ[35]	SSTL	I/O
L6	DDR1_DQ[39]	SSTL	I/O
L7	DDR1_DQS_DP[04]	SSTL	I/O
L8	DDR1_DQS_DN[13]	SSTL	I/O
L9	DDR1_DQ[33]	SSTL	I/O
M1	DDR3_DQ[55]	SSTL	I/O
M10	VSS	GND	
M11	VSS	GND	
M12	RSVD		
M13	VSS	GND	
M14	VTTD	PWR	
M15	VTTD	PWR	
M16	VSS	GND	
M17	VTTD	PWR	
M18	DDR1_MA[10]	SSTL	O
M19	RSVD		
M2	DDR3_DQ[54]	SSTL	I/O
M20	RSVD		
M21	RSVD		
M22	DDR1_MA[04]	SSTL	O
M23	VSS	GND	

Table 8-2. Land Number (Sheet 32 of 37)

Land Number	Land Name	Buffer Type	Direction
M24	DDR1_MA[11]	SSTL	O
M25	VTTD	PWR	
M26	VTTD	PWR	
M27	VSS	GND	
M28	VSS	GND	
M29	VTTD	PWR	
M3	DDR3_DQ[50]	SSTL	I/O
M30	VTTD	PWR	
M31	VTTD	PWR	
M32	VSS	GND	
M33	VSS	GND	
M34	VSS	GND	
M35	VSS	GND	
M36	DDR_VREFDQRX_C23	DC	I
M37	VSS	GND	
M38	DDR2_DQ[10]	SSTL	I/O
M39	DDR2_DQ[11]	SSTL	I/O
M4	DDR1_RCOMP[2]	Analog	I
M40	VSS	GND	
M41	DDR3_DQ[20]	SSTL	I/O
M42	VSS	GND	
M43	VSS	GND	
M5	VSS	GND	
M6	VSS	GND	
M7	VSS	GND	
M8	VSS	GND	
M9	VSS	GND	
N1	VSS	GND	
N10	DDR_VREFDQRX_C1	DC	I
N11	VSS_VCC_SENSE		O
N2	VSS	GND	
N3	DDR3_DQ[51]	SSTL	I/O
N33	TEST2		O
N34	DDR1_DQ[15]	SSTL	I/O
N35	DDR1_DQ[11]	SSTL	I/O
N36	DDR1_DQ[10]	SSTL	I/O
N37	VSS	GND	
N38	DDR2_DQ[15]	SSTL	I/O

Table 8-2. Land Number (Sheet 33 of 37)

Land Number	Land Name	Buffer Type	Direction
N39	DDR2_DQ[14]	SSTL	I/O
N4	RSVD		
N40	VSS	GND	
N41	MEM_HOT_C23_N	ODCMOS	I/O
N42	VSS	GND	
N43	DDR3_DQ[11]	SSTL	I/O
N5	DDR2_DQ[53]	SSTL	I/O
N6	DDR2_DQ[52]	SSTL	I/O
N7	DDR1_RCOMP[1]	Analog	I
N8	DDR1_DQ[40]	SSTL	I/O
N9	DDR1_DQ[44]	SSTL	I/O
P1	DDR3_DQ[60]	SSTL	I/O
P10	VSS	GND	
P11	VCC_SENSE		O
P2	VSS	GND	
P3	VSS	GND	
P33	VSS	GND	
P34	DDR1_DQ[14]	SSTL	I/O
P35	DDR1_DQS_DN[01]	SSTL	I/O
P36	DDR1_DQS_DP[01]	SSTL	I/O
P37	VSS	GND	
P38	DDR2_DQS_DP[01]	SSTL	I/O
P39	DDR2_DQS_DN[01]	SSTL	I/O
P4	VSS	GND	
P40	VSS	GND	
P41	DDR3_DQ[10]	SSTL	I/O
P42	DDR3_DQ[15]	SSTL	I/O
P43	DDR3_DQ[14]	SSTL	I/O
P5	DDR2_DQ[49]	SSTL	I/O
P6	DDR2_DQ[48]	SSTL	I/O
P7	VSS	GND	
P8	DDR1_DQ[41]	SSTL	I/O
P9	DDR1_DQ[45]	SSTL	I/O
R1	DDR3_DQ[57]	SSTL	I/O
R10	VSS	GND	
R11	VCC	PWR	
R2	DDR3_DQ[56]	SSTL	I/O
R3	DDR3_DQ[61]	SSTL	I/O

Table 8-2. Land Number (Sheet 34 of 37)

Land Number	Land Name	Buffer Type	Direction
R33	VCC	PWR	
R34	VSS	GND	
R35	DDR1_DQS_DP[10]	SSTL	I/O
R36	DDR1_DQS_DN[10]	SSTL	I/O
R37	VSS	GND	
R38	DDR2_DQS_DN[10]	SSTL	I/O
R39	DDR2_DQS_DP[10]	SSTL	I/O
R4	VSS	GND	
R40	VSS	GND	
R41	DDR3_DQS_DN[10]	SSTL	I/O
R42	DDR3_DQS_DP[01]	SSTL	I/O
R43	DDR3_DQS_DN[01]	SSTL	I/O
R5	DDR2_DQS_DN[15]	SSTL	I/O
R6	DDR2_DQS_DP[15]	SSTL	I/O
R7	VSS	GND	
R8	DDR1_DQS_DN[14]	SSTL	I/O
R9	DDR1_DQS_DP[14]	SSTL	I/O
T1	DDR3_DQS_DN[16]	SSTL	I/O
T10	VSS	GND	
T11	VCC	PWR	
T2	DDR3_DQS_DP[16]	SSTL	I/O
T3	DDR3_DQS_DN[07]	SSTL	I/O
T33	VCC	PWR	
T34	VSS	GND	
T35	DDR1_DQ[08]	SSTL	I/O
T36	DDR1_DQ[09]	SSTL	I/O
T37	VSS	GND	
T38	DDR2_DQ[09]	SSTL	I/O
T39	DDR2_DQ[08]	SSTL	I/O
T4	VSS	GND	
T40	VSS	GND	
T41	DDR3_DQS_DP[10]	SSTL	I/O
T42	VSS	GND	
T43	DDR3_DQ[09]	SSTL	I/O
T5	DDR2_DQS_DP[06]	SSTL	I/O
T6	DDR2_DQS_DN[06]	SSTL	I/O
T7	VSS	GND	
T8	DDR1_DQS_DP[05]	SSTL	I/O

Table 8-2. Land Number (Sheet 35 of 37)

Land Number	Land Name	Buffer Type	Direction
T9	DDR1_DQS_DN[05]	SSTL	I/O
U1	DDR3_DQ[62]	SSTL	I/O
U10	VSS	GND	
U11	VCC	PWR	
U2	VSS	GND	
U3	DDR3_DQS_DP[07]	SSTL	I/O
U33	VCC	PWR	
U34	VSS	GND	
U35	DDR1_DQ[12]	SSTL	I/O
U36	DDR1_DQ[13]	SSTL	I/O
U37	VSS	GND	
U38	DDR2_DQ[13]	SSTL	I/O
U39	DDR2_DQ[12]	SSTL	I/O
U4	VSS	GND	
U40	VSS	GND	
U41	DDR3_DQ[12]	SSTL	I/O
U42	DDR3_DQ[08]	SSTL	I/O
U43	DDR3_DQ[13]	SSTL	I/O
U5	DDR2_DQ[55]	SSTL	I/O
U6	DDR2_DQ[54]	SSTL	I/O
U7	VSS	GND	
U8	DDR1_DQ[47]	SSTL	I/O
U9	DDR1_DQ[46]	SSTL	I/O
V1	DDR3_DQ[58]	SSTL	I/O
V10	VSS	GND	
V11	VCC	PWR	
V2	DDR3_DQ[63]	SSTL	I/O
V3	DDR3_DQ[59]	SSTL	I/O
V33	VCC	PWR	
V34	IVT_ID_N		O
V35	VSS	GND	
V36	VSS	GND	
V37	DDR_VREFDQTX_C23	DC	O
V38	VSS	GND	
V39	VSS	GND	
V4	VSS	GND	
V40	DDR_SCL_C23	ODCMOS	I/O
V41	DDR_SDA_C23	ODCMOS	I/O

Table 8-2. Land Number (Sheet 36 of 37)

Land Number	Land Name	Buffer Type	Direction
V42	VSS	GND	
V43	VSS	GND	
V5	DDR2_DQ[51]	SSTL	I/O
V6	DDR2_DQ[50]	SSTL	I/O
V7	VSS	GND	
V8	DDR1_DQ[43]	SSTL	I/O
V9	DDR1_DQ[42]	SSTL	I/O
W1	VSS	GND	
W10	MEM_HOT_C1_N	ODCMOS	I/O
W11	VCC	PWR	
W2	VSS	GND	
W3	VSS	GND	
W33	VCC	PWR	
W34	VSS	GND	
W35	DDR1_DQ[02]	SSTL	I/O
W36	DDR1_DQ[03]	SSTL	I/O
W37	VSS	GND	
W38	DDR2_DQ[03]	SSTL	I/O
W39	DDR2_DQ[02]	SSTL	I/O
W4	DDR_VREFDQTX_C1	DC	O
W40	VSS	GND	
W41	DDR3_DQ[07]	SSTL	I/O
W42	DDR3_DQ[02]	SSTL	I/O
W43	DDR3_DQ[03]	SSTL	I/O
W5	VSS	GND	
W6	VSS	GND	
W7	DDR_SCL_C1	ODCMOS	I/O
W8	DDR_SDA_C1	ODCMOS	I/O
W9	VSS	GND	
Y1	DDR1_DQ[56]	SSTL	I/O
Y10	DRAM_PWR_OK_C1	CMOS_1.5 V	I
Y11	VCC	PWR	
Y2	DDR1_DQ[61]	SSTL	I/O
Y3	DDR1_DQ[60]	SSTL	I/O
Y33	VCC	PWR	
Y34	VSS	GND	
Y35	DDR1_DQ[06]	SSTL	I/O
Y36	DDR1_DQ[07]	SSTL	I/O

Table 8-2. Land Number (Sheet 37 of 37)

Land Number	Land Name	Buffer Type	Direction
Y37	VSS	GND	
Y38	DDR2_DQ[07]	SSTL	I/O
Y39	DDR2_DQ[06]	SSTL	I/O
Y4	VSS	GND	
Y40	VSS	GND	
Y41	DDR3_DQ[06]	SSTL	I/O
Y5	DDR2_DQ[61]	SSTL	I/O
Y6	DDR2_DQ[60]	SSTL	I/O
Y7	VSS	GND	
Y8	DDR1_DQ[53]	SSTL	I/O
Y9	DDR1_DQ[52]	SSTL	I/O

§

9 Package Mechanical Specifications

The processor is packaged in a Flip-Chip Land Grid Array (FCLGA10)package that interfaces with the baseboard via an 1356 socket. The package consists of a processor mounted on a substrate land-carrier. An integrated heat spreader (IHS) is attached to the package substrate and core and serves as the mating surface for processor component thermal solutions, such as a heatsink. Figure 9-1 shows a sketch of the processor package components and how they are assembled together. Refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for complete details on the 1356 socket.

The package components shown in Figure 9-1 include the following:

1. Integrated Heat Spreader (IHS)
2. Thermal Interface Material (TIM)
3. Processor core (die)
4. Package substrate
5. Capacitors

Figure 9-1. Processor Package Assembly Sketch

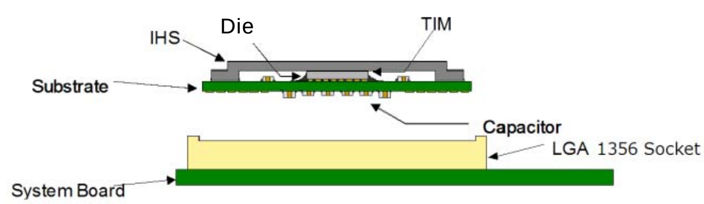

Note:
1. Socket and baseboard are included for reference and are not part of processor package.

9.1 Package Mechanical Drawing

The package mechanical drawings are shown in . The drawings include dimensions necessary to design a thermal solution for the processor. These dimensions include:

1. Package reference with tolerances (total height, length, width, and so forth)
2. IHS parallelism and tilt
3. Land dimensions
4. Top-side and back-side component keep-out dimensions
5. Reference datums

6. All drawing dimensions are in millimeters (mm).
7. Guidelines on potential IHS flatness variation with socket load plate actuation and installation of the cooling solution is available in the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

Figure 9-2. Processor Package Drawing Sheet 1 of 2

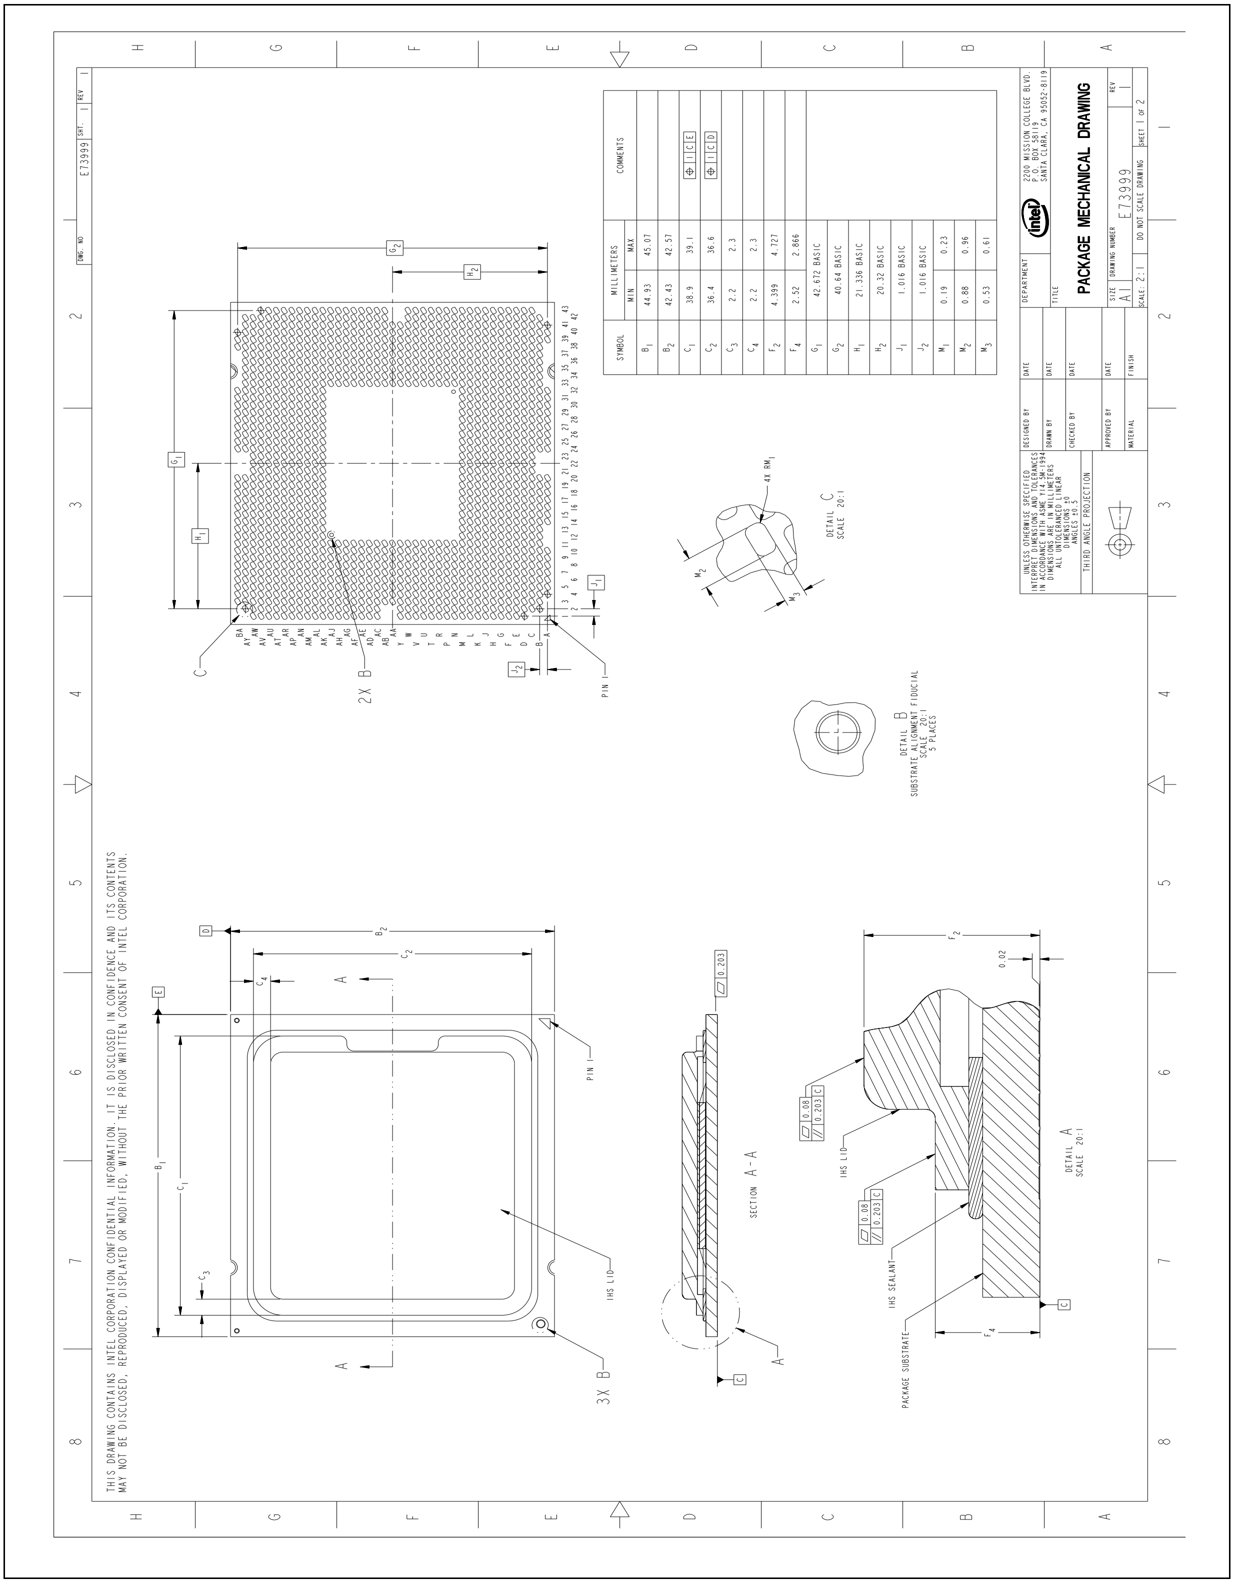

Figure 9-3. Processor Package Drawing Sheet 2 of 2

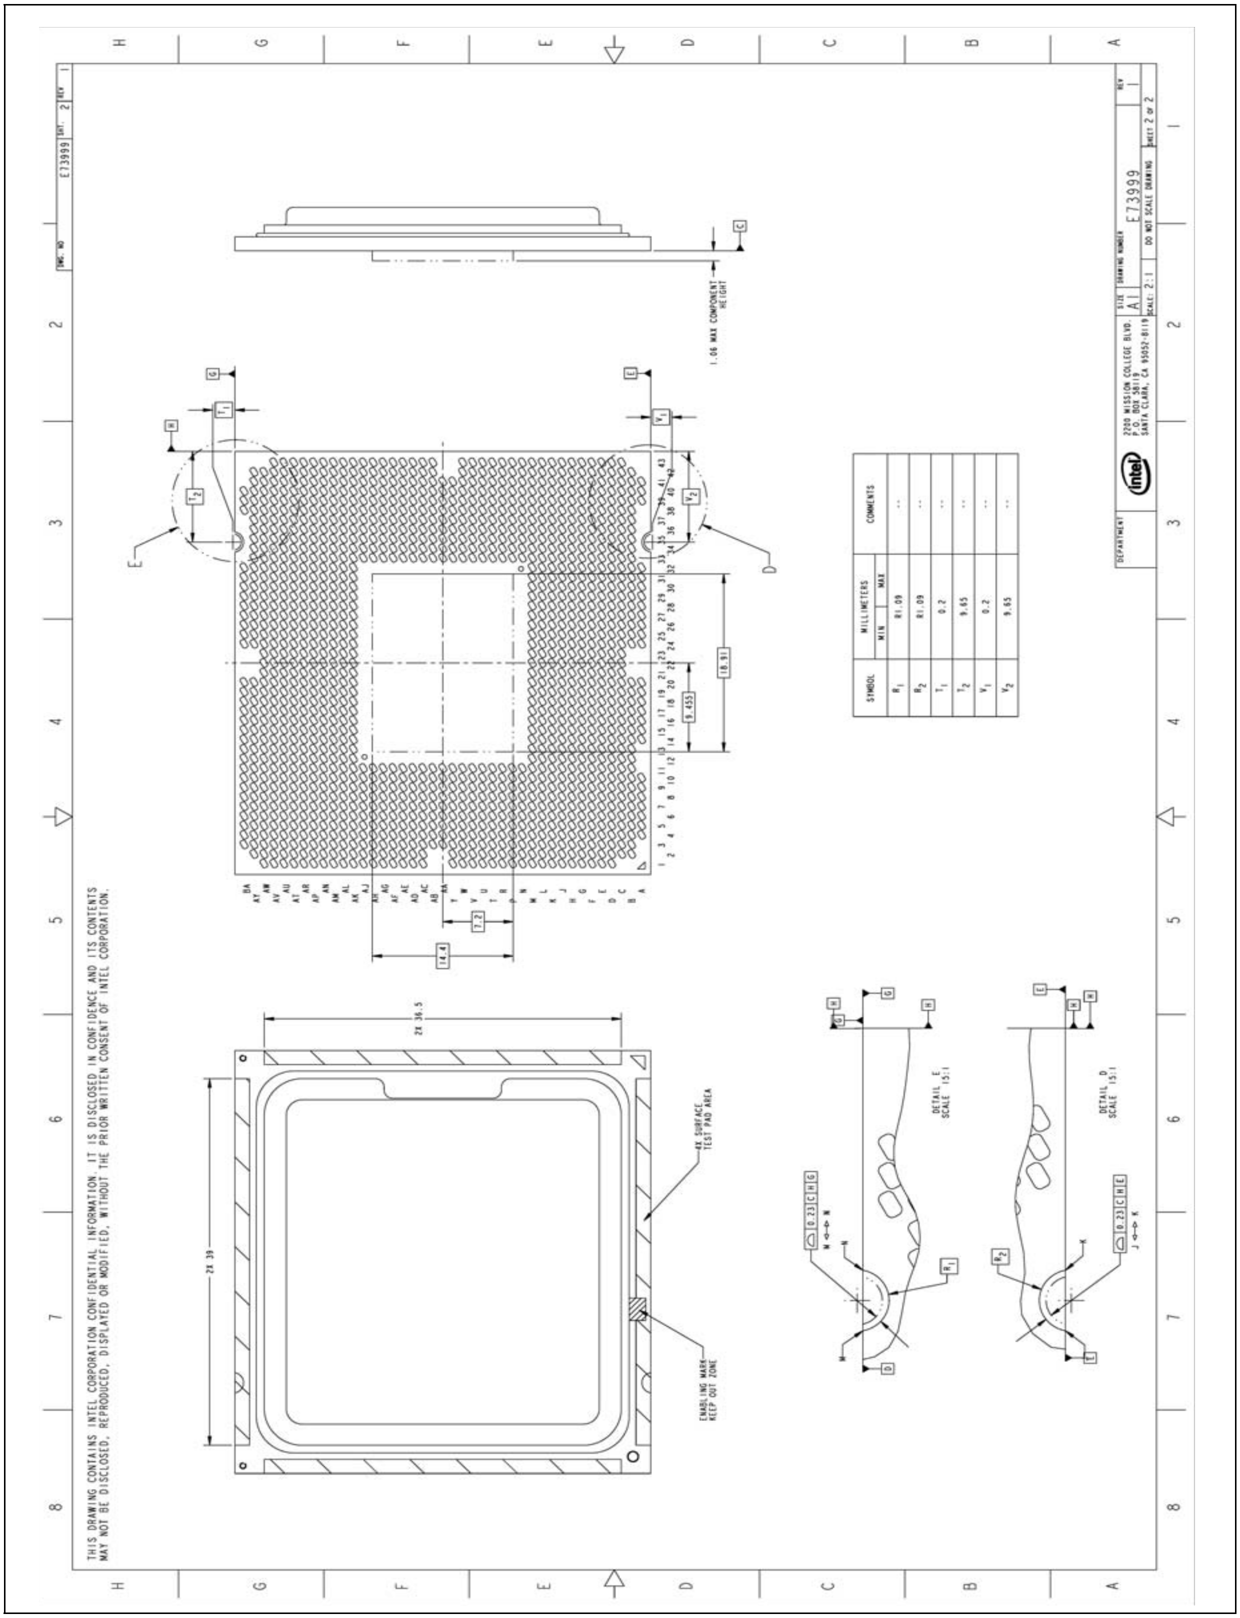

9.2 Processor Component Keep-Out Zones

The processor may contain components on the substrate that define component keep-out zone requirements. A thermal and mechanical solution design must not intrude into the required keep-out zones. Do not contact the Test Pad Area with conductive material. Decoupling capacitors are typically mounted to either the topside or land-side of the package substrate. See for keep-out zones. The location and quantity of package capacitors may change due to manufacturing efficiencies but will remain within the component keep-in.

9.3 Package Loading Specifications

Table 9-1 provides load specifications for the processor package. These maximum limits should not be exceeded during heatsink assembly, shipping conditions, or standard use condition. Exceeding these limits during test may result in component failure. The processor substrate should not be used as a mechanical reference or load-bearing surface for thermal solutions.

Table 9-1. Processor Loading Specifications

Parameter	Maximum	Notes
Static Compressive Load	890 N [200 lbf]	1, 2, 3, 5
Dynamic Load	1779 N [400lbf] [max static compressive + dynamic load]	1, 3, 4, 5

Notes:
1. These specifications apply to uniform compressive loading in a direction normal to the processor IHS.
2. This is the maximum static force that can be applied by the heatsink and Independent Loading Mechanism (ILM).
3. These specifications are based on limited testing for design characterization. Loading limits are for the package constrained by the limits of the processor socket.
4. Dynamic loading is defined as an 11 ms duration average load superimposed on the static load requirement.
5. See *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for minimum socket load to engage processor within socket.

9.4 Package Handling Guidelines

Table 9-2 includes a list of guidelines on package handling in terms of recommended maximum loading on the processor IHS relative to a fixed substrate. These package handling loads may be experienced during heatsink removal.

Table 9-2. Package Handling Guidelines

Parameter	Maximum Recommended	Notes
Shear	70 lbs	
Tensile	25 lbs	
Torque	35 in.lbs	

9.5 Package Insertion Specifications

The processor can be inserted into and removed from an 1356 socket 15 times. The socket should meet the 1356 requirements detailed in the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

9.6 Processor Mass Specification

The typical mass of the processor is currently 35 grams. This mass [weight] includes all the components that are included in the package.

9.7 Processor Materials

Table 9-3 lists some of the package components and associated materials.

Table 9-3. Processor Materials

Component	Material
Integrated Heat Spreader (IHS)	Nickel Plated Copper
Substrate	Halogen Free, Fiber Reinforced Resin
Substrate Lands	Gold Plated Copper

9.8 Processor Markings

Figure 9-4 shows the topside markings on the processor. This diagram is to aid in the identification of the processor.

Figure 9-4. Processor Top-Side Markings

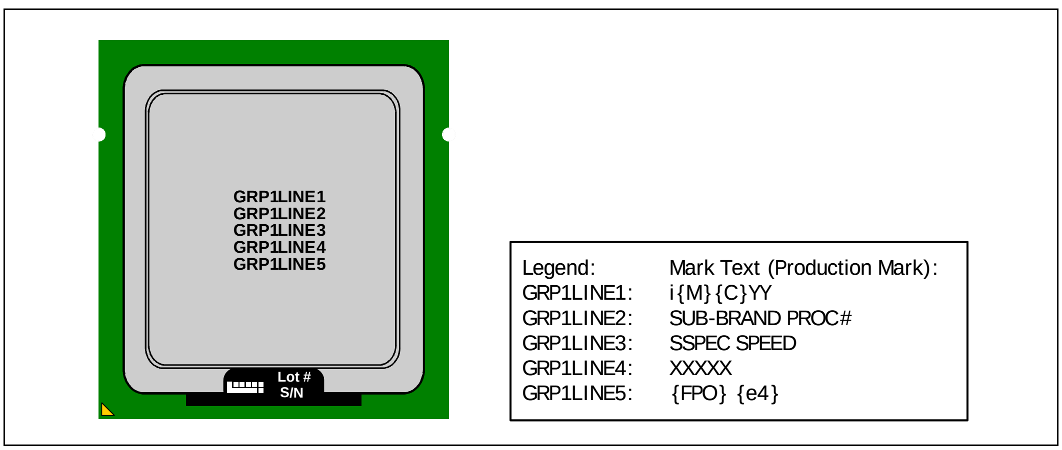

Notes:
1. XXXXX = Country of Origin
2. SPEED Format = X.XX GHz and no rounding

§

10 Boxed Processor Specifications

10.1 Introduction

Intel boxed processors are intended for system integrators who build systems from components available through distribution channels. The Intel® Xeon® processor E5-2400 product family (LGA1356) processors will be offered as Intel boxed processors, however the thermal solutions will be sold separately.

Boxed processors will not include a thermal solution in the box. Intel will offer boxed thermal solutions separately through the same distribution channels. Please reference Section 10.1.1 - Section 10.1.4 for a description of Boxed Processor thermal solutions.

10.1.1 Available Boxed Thermal Solution Configurations

Intel will offer three different Boxed Heat Sink solutions to support LGA1356 Boxed Processors

- Boxed Intel® Thermal Solution STS100C (Order Code BXSTS100C): A Passive / Active Combination Heat Sink Solution that is intended for processors with a TDP up to 95W in a pedestal or 2U+ chassis with appropriate ducting.
- Boxed Intel® Thermal Solution STS100A (Order Code BXSTS100A): An Active Heat Sink Solution that is intended for processors with a TDP of 95W or lower in pedestal chassis.
- Boxed Intel® Thermal Solution STS100P (Order Code BXSTS100P): A 25.5 mm Tall Passive Heat Sink Solution that is intended for processors with a TDP of 95W or lower in 1U, or 2U chassis with appropriate ducting. Check with Blade manufacturer for compatibility.

10.1.2 Intel Thermal Solution STS100C (Passive / Active Combination Heat Sink Solution)

The STS100C, based on a 2U passive heat sink with a removable fan, is intended for use with processors with TDP's up to 95W. This heat pipe-based solution is intended to be used as either a passive heat sink in a 2U or larger chassis, or as an active heat sink for pedestal chassis and will provide improved acoustic performance when compared to the STS100A. Figure 10-1 and Figure 10-2 are representations of the heat sink solution. Although the active combination solution with the removable fan installed mechanically fits into a 2U keepout, its use has not been validated in that configuration.

The STS100C in the active fan configuration is primarily designed to be used in a pedestal chassis where sufficient air inlet space is present. The STS100C with the fan removed, as with any passive thermal solution, will require the use of chassis ducting and are targeted for use in rack mount or ducted pedestal servers. The retention solution used for these products is called Unified Retention System (URS).

Figure 10-1. STS100C Passive/Active Combination Heat Sink (with Removable Fan)

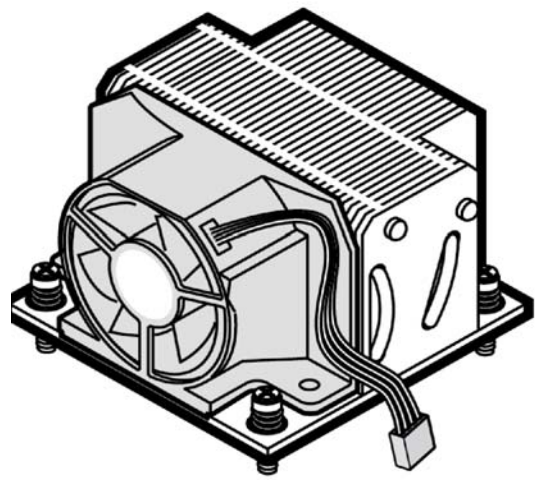

Figure 10-2. STS100C Passive/Active Combination Heat Sink (with Fan Removed)

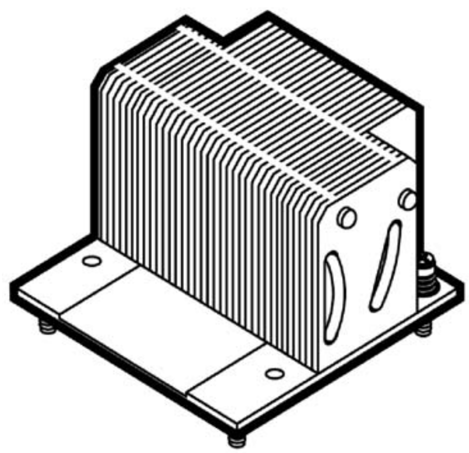

10.1.3 Intel Thermal Solution STS100A (Active Heat Sink Solution)

The STS100A in the active fan configuration is primarily designed to be used in a pedestal chassis where sufficient air inlet space is present (see Figure 10-3). The STS100A with the fan removed, as with any passive thermal solution, will require the use of chassis ducting and is targeted for use in rack mount or ducted pedestal servers. The retention solution used for these products is called Unified Retention System (URS).

The STS100C and STS100A utilize a fan capable of 4-pin pulse width modulated (PWM) control. Use of a 4-pin PWM controlled active thermal solution helps customers meet acoustic targets in pedestal platforms through the baseboard's ability to directly control the RPM of the processor heat sink fan. See Section 10.3 for more details on fan speed control. Also see Section 2.5, "Platform Environment Control Interface (PECI)" for more on the PWM and PECI interface along with Digital Thermal Sensors (DTS).

Figure 10-3. STS100A Active Heat Sink

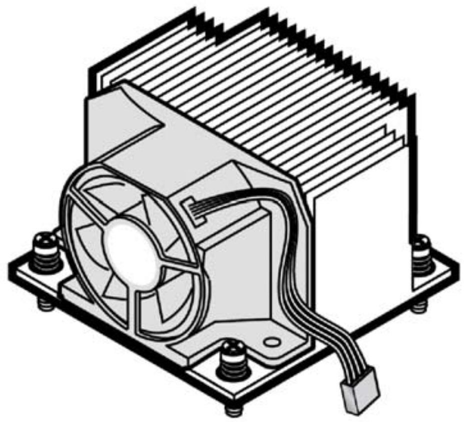

10.1.4 Intel Thermal Solution STS100P (Boxed 25.5 mm Tall Passive Heat Sink Solution)

The STS100Pis available for use with boxed processors that have TDP's of 95W and lower. The 25.5 mm Tall passive solution is designed to be used in SSI Blades, 1U, and 2U chassis where ducting is present. The use of a 25.5 mm Tall heatsink in a 2U chassis is recommended to achieve a lower heatsink T_{LA} and more flexibility in system design optimization. Figure 10-4 is a representation of the heat sink solution. The retention solution used for these products is called Unified Retention System (URS).

Figure 10-4. STS100P 25.5 mm Tall Passive Heat Sink

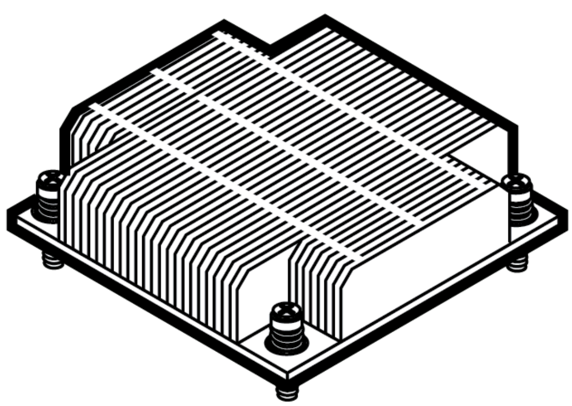

10.2 Mechanical Specifications

This section documents the mechanical specifications of the boxed processor solution.

10.2.1 Boxed Processor Heat Sink Dimensions and Baseboard Keepout Zones

The boxed processor and boxed thermal solutions will be sold separately. Clearance is required around the thermal solution to ensure unimpeded airflow for proper cooling. Baseboard keepout zones are Figure 10-5 - Figure 10-8. Physical space requirements and dimensions for the boxed processor and assembled heat sink are shown in Figure 10-9 and Figure 10-10. Mechanical drawings for the 4-pin fan header and 4-pin connector used for the active fan heat sink solution are represented in Figure 10-11and Figure 10-12.

None of the heat sink solutions exceed a mass of 550 grams. Note that this is per processor, a dual processor system will have up to 1100 grams total mass in the heat sinks. See Section 9.6 for details on the processor mass test.

Figure 10-5. Boxed Processor Motherboard Keepout Zones (1 of 4)

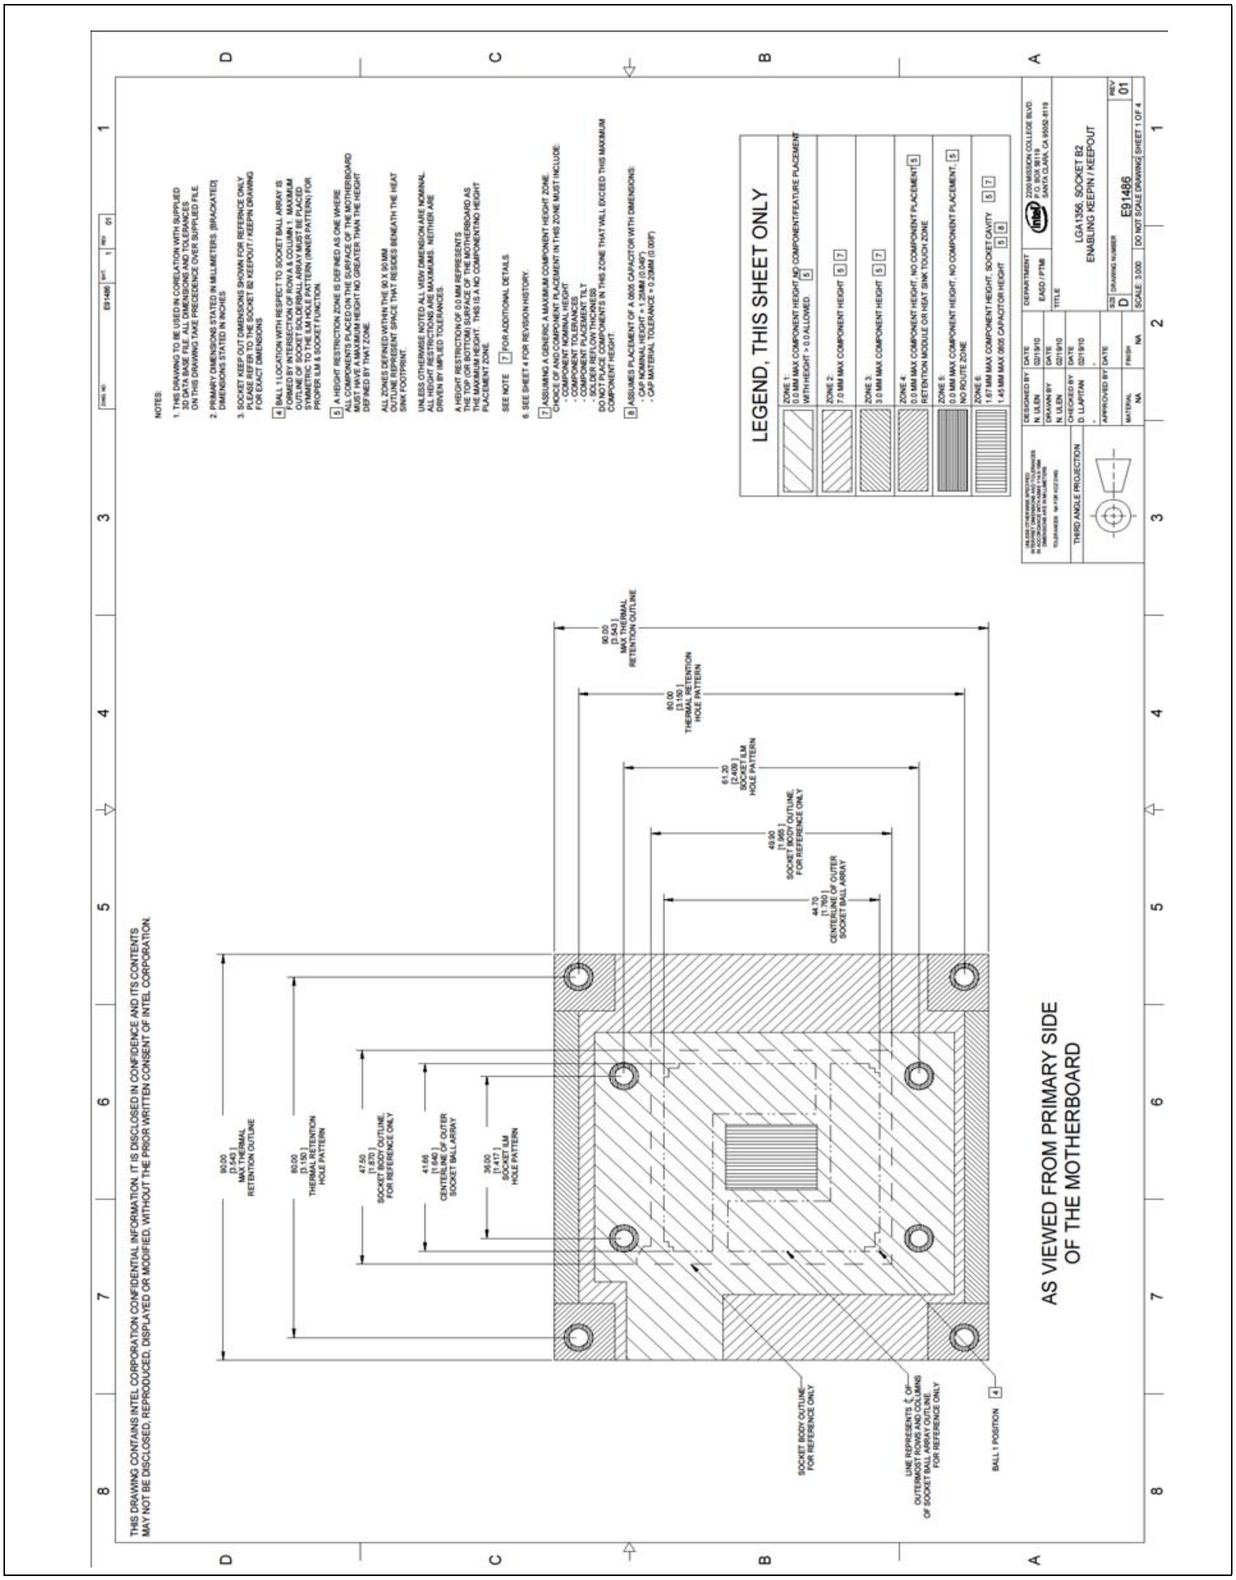

Figure 10-6. Boxed Processor Motherboard Keepout Zones (2 of 4)

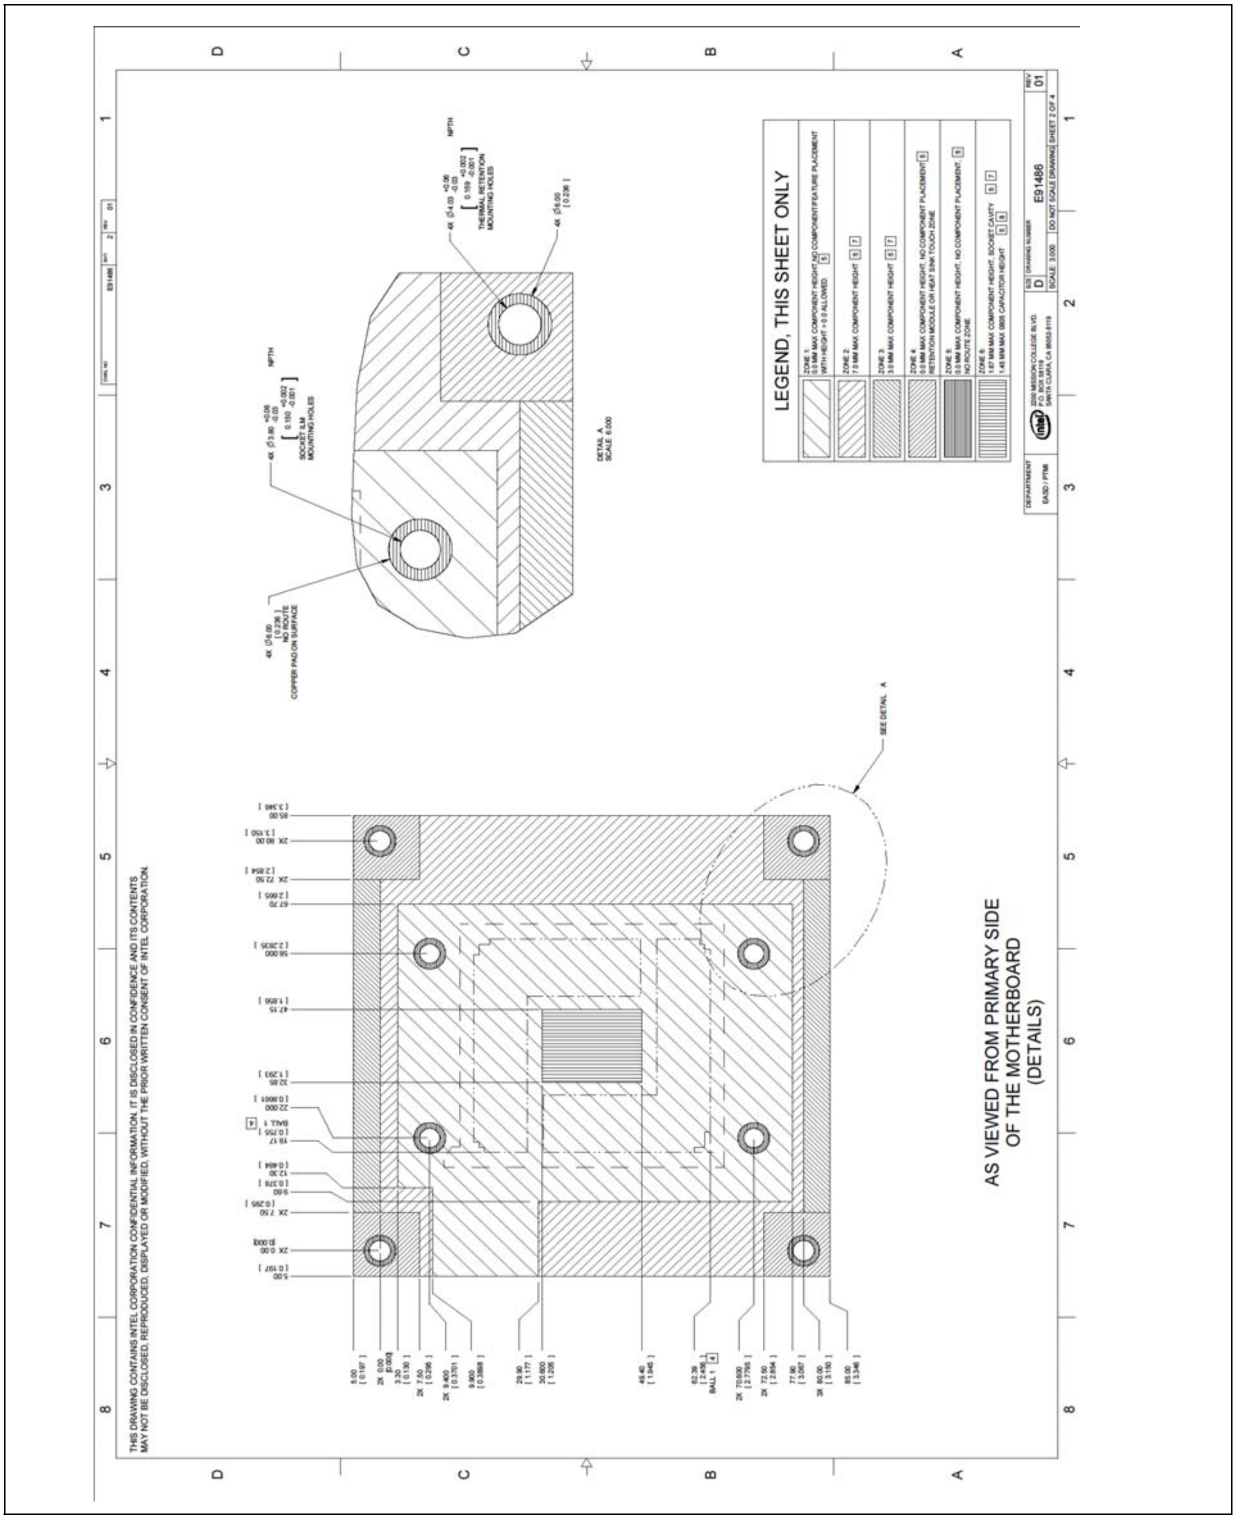

Figure 10-7. Boxed Processor Motherboard Keepout Zones (3 of 4)

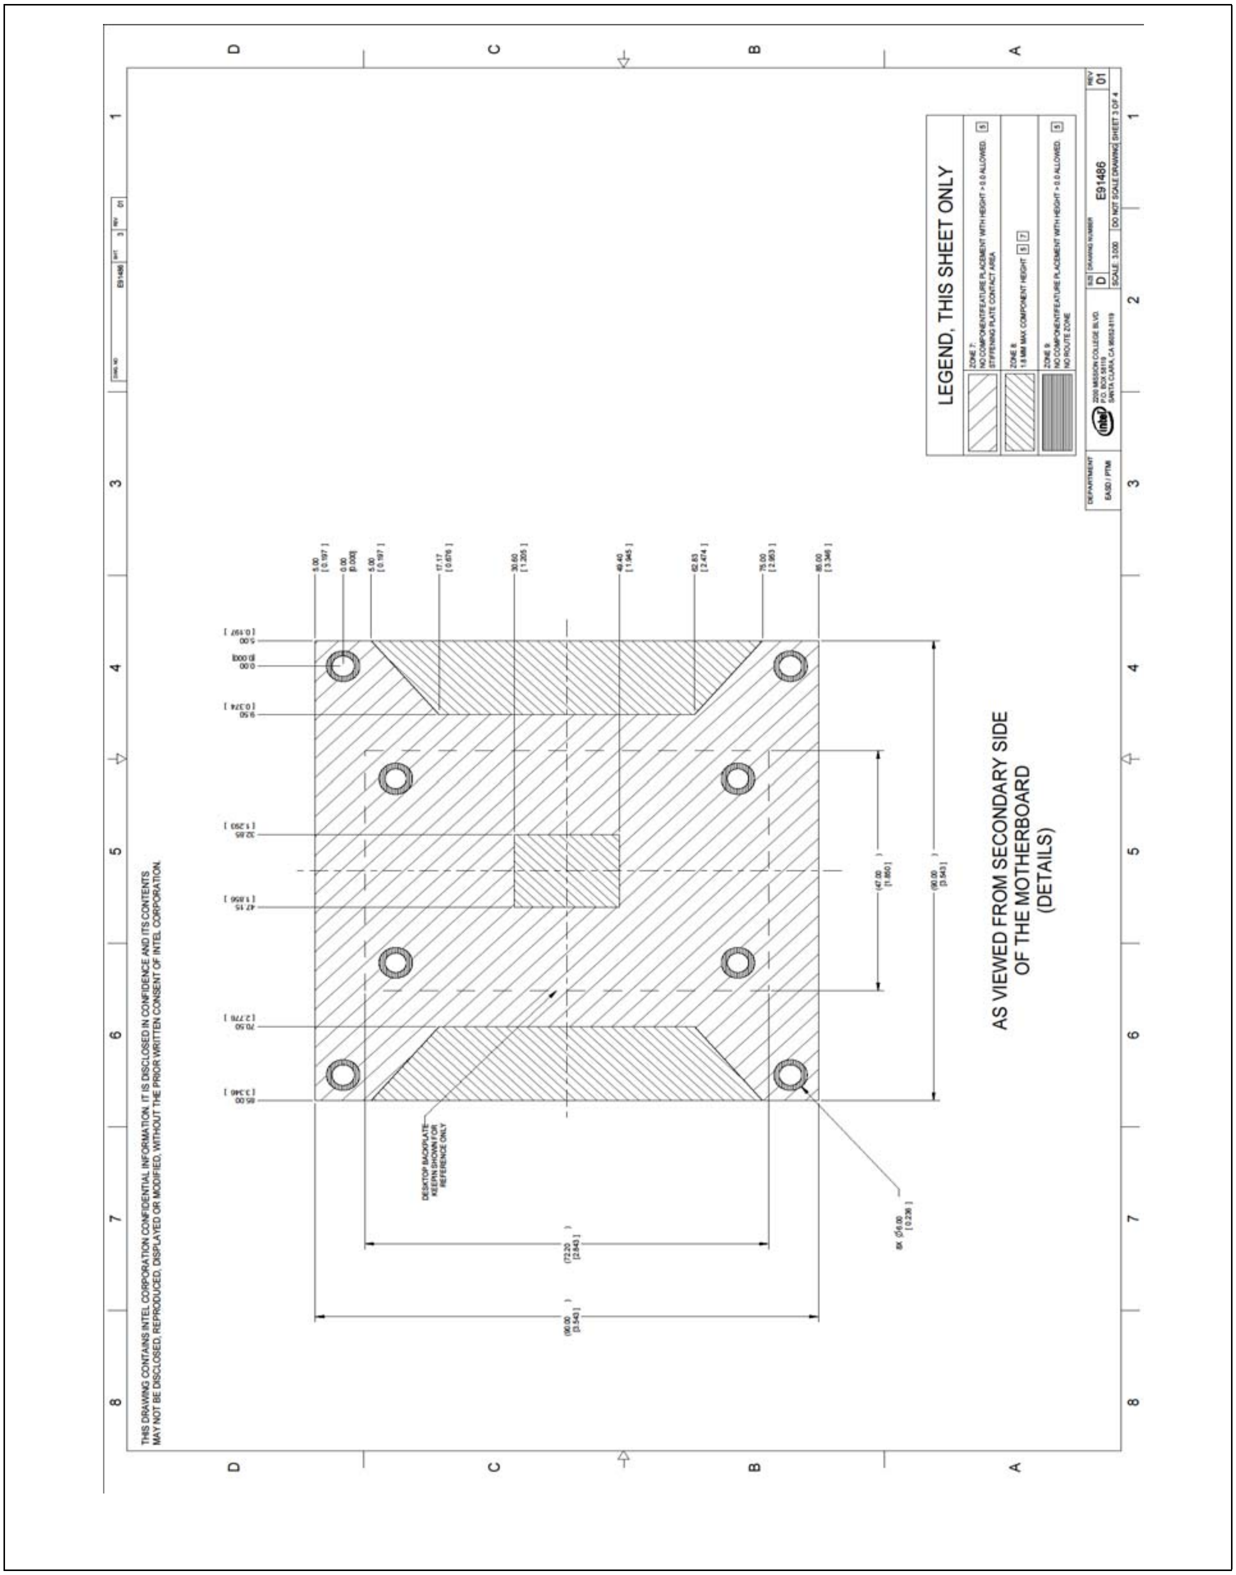

Figure 10-8. Boxed Processor Motherboard Keepout Zones (4 of 4)

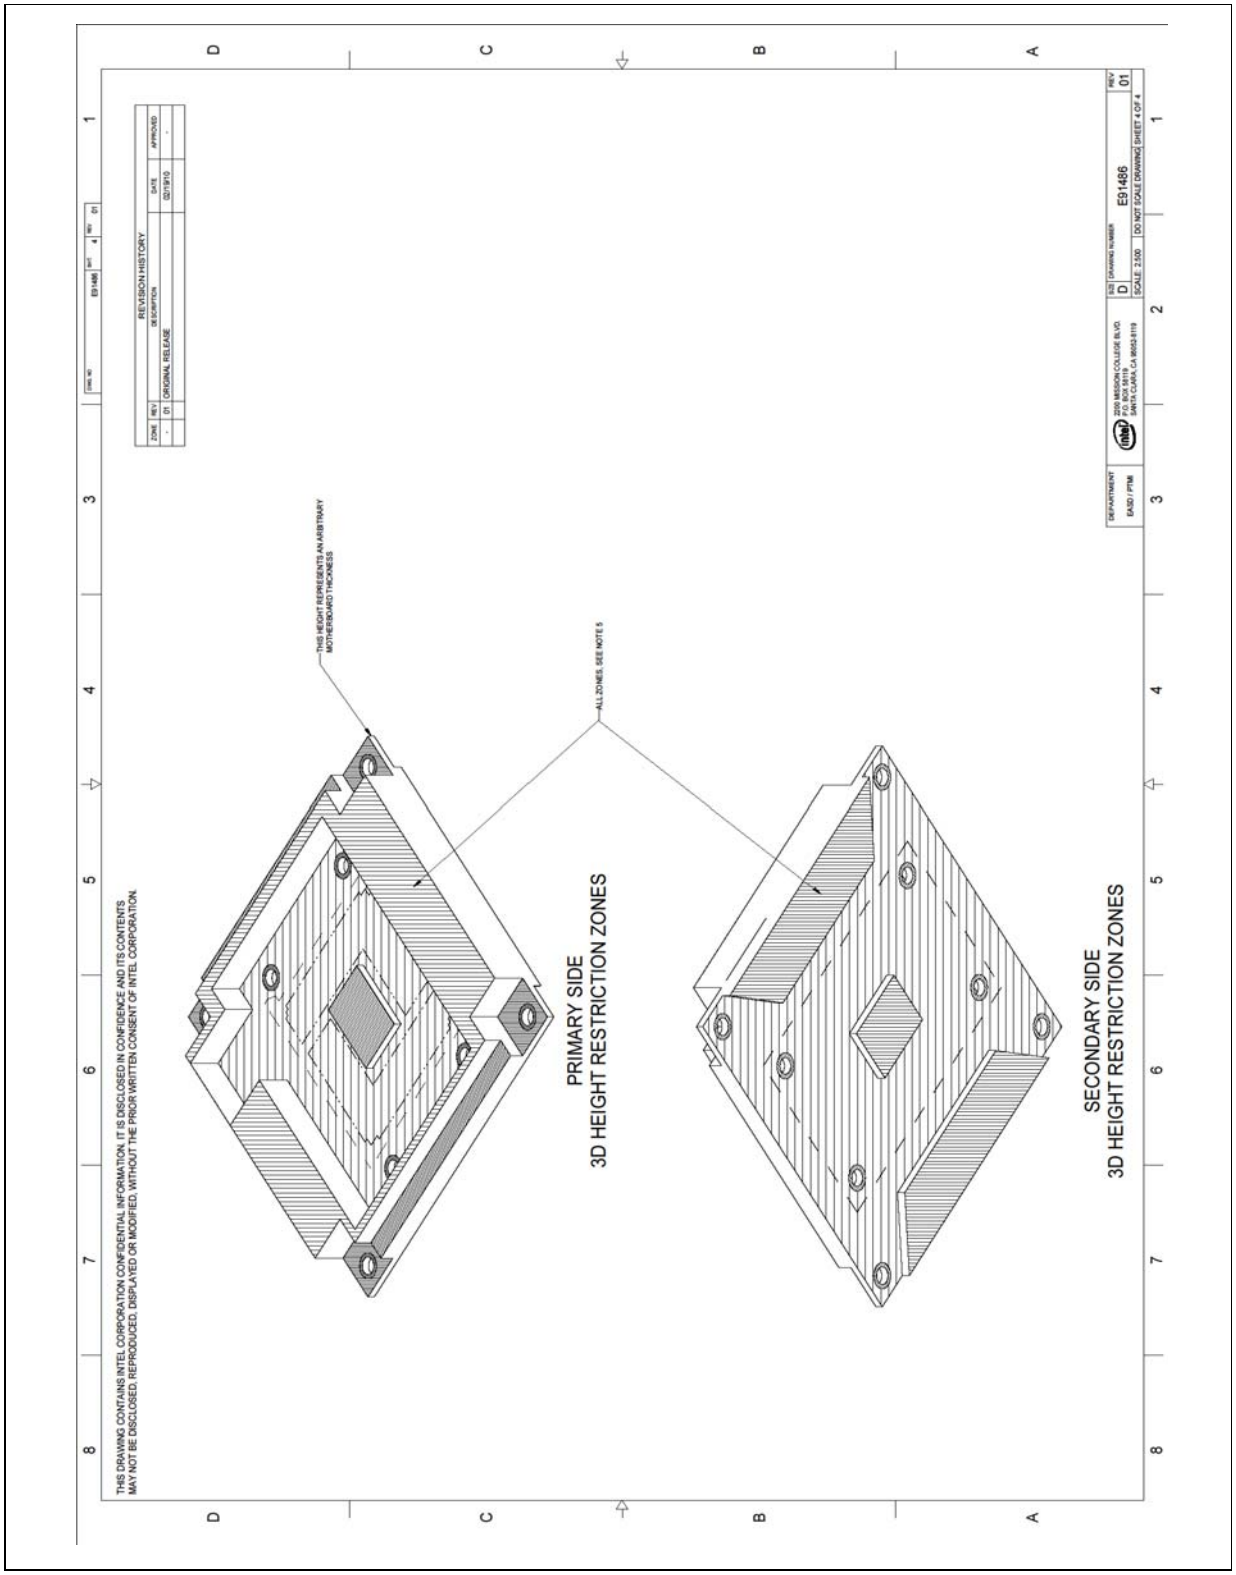

Figure 10-9. Boxed Processor Heat Sink Volumetric (1 of 2)

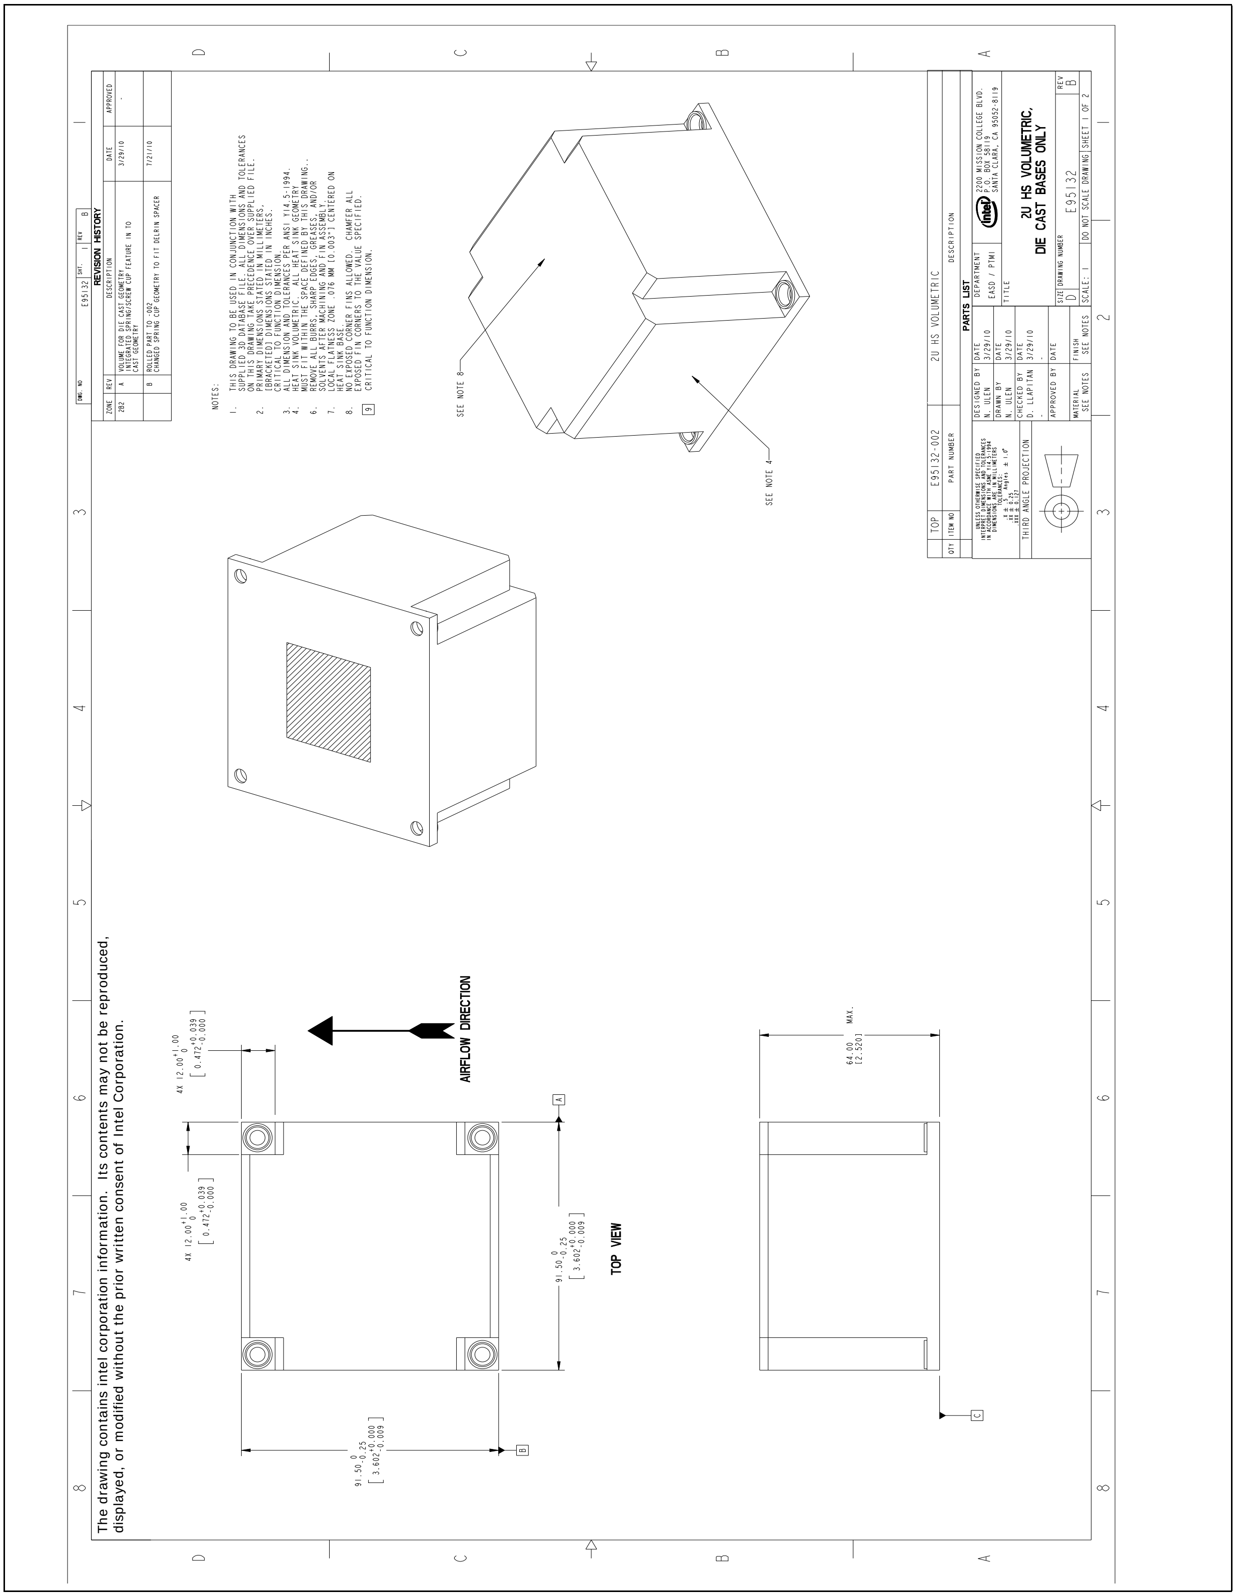

Figure 10-10.Boxed Processor Heat Sink Volumetric (2 of 2)

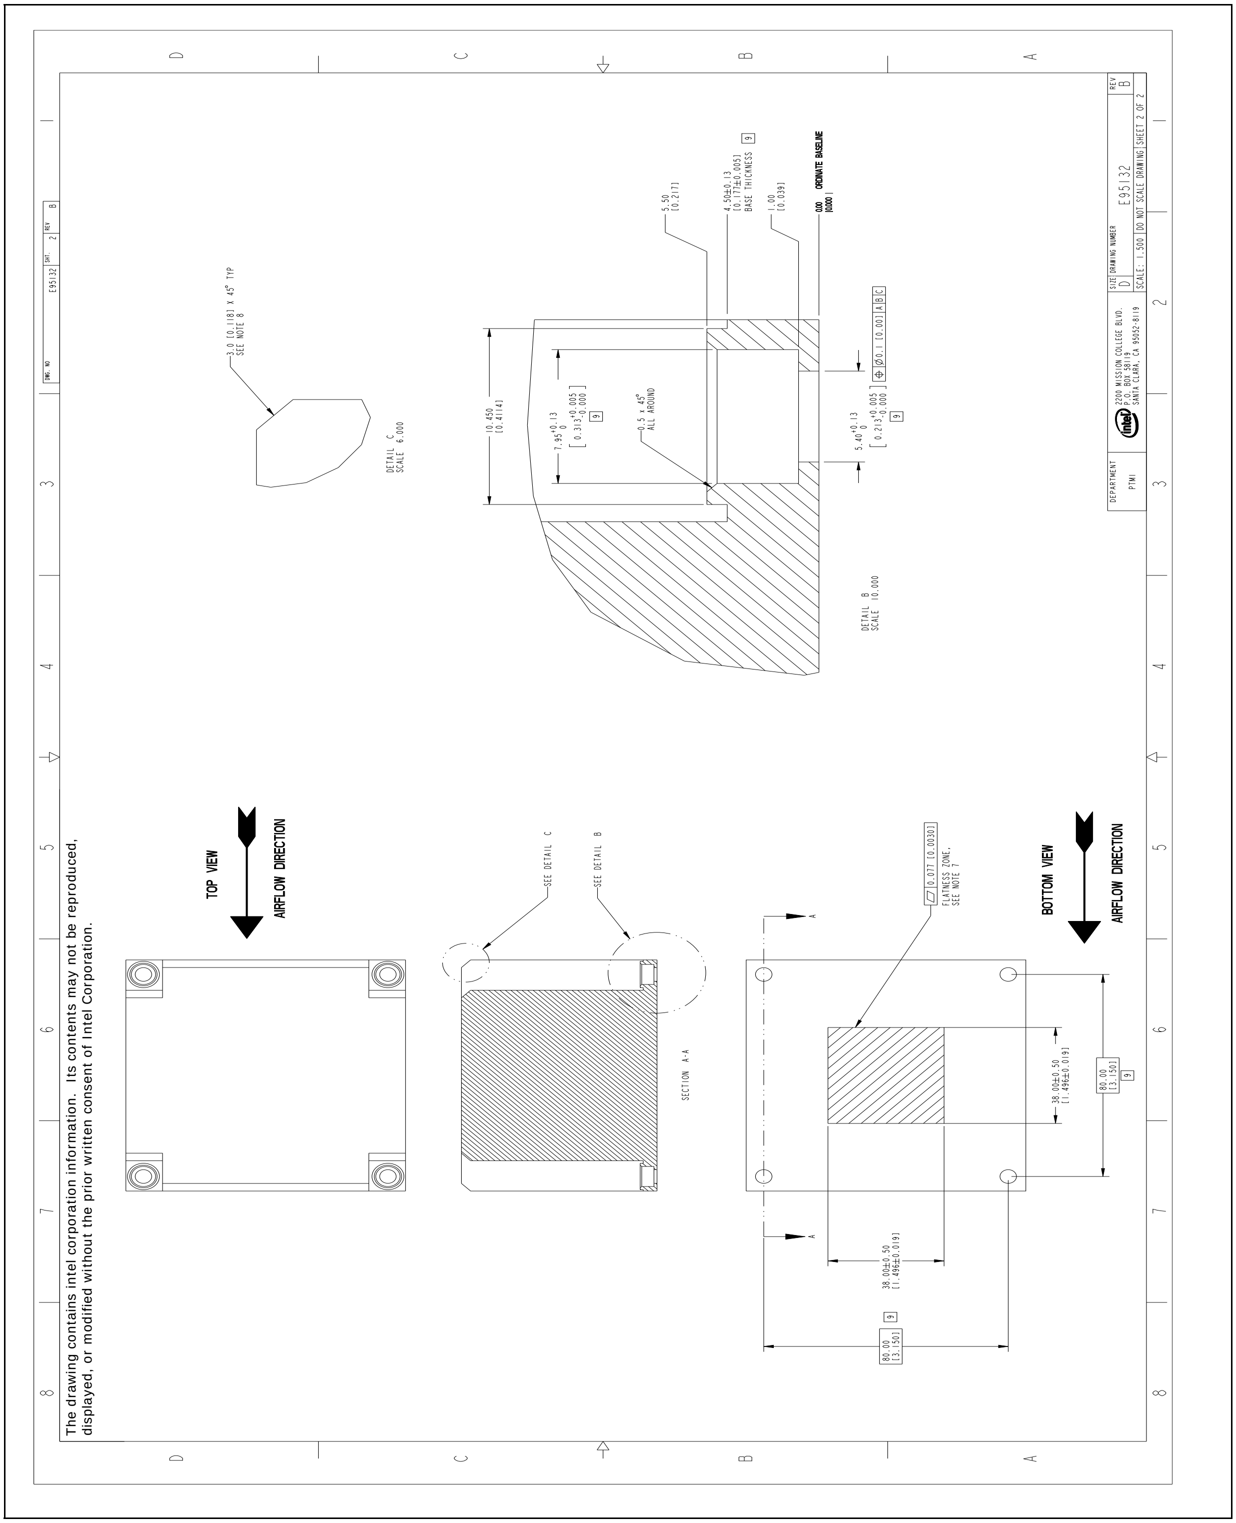

Figure 10-11.4-Pin Fan Cable Connector (For Active Heat Sink)

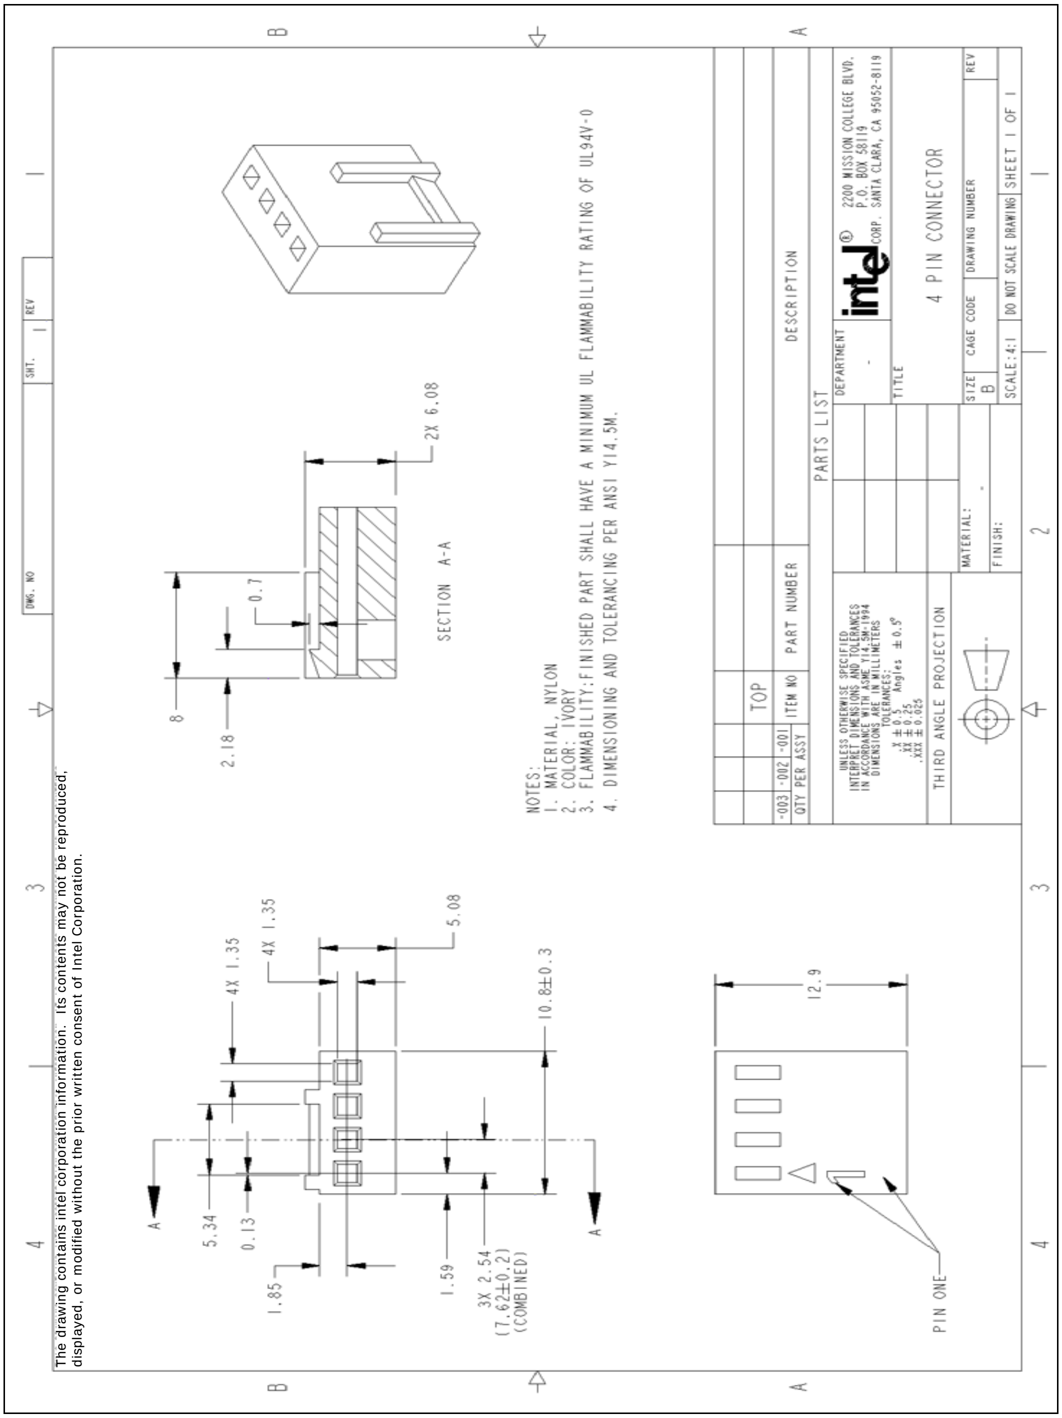

Figure 10-12. 4-Pin Base Baseboard Fan Header (For Active Heat Sink)

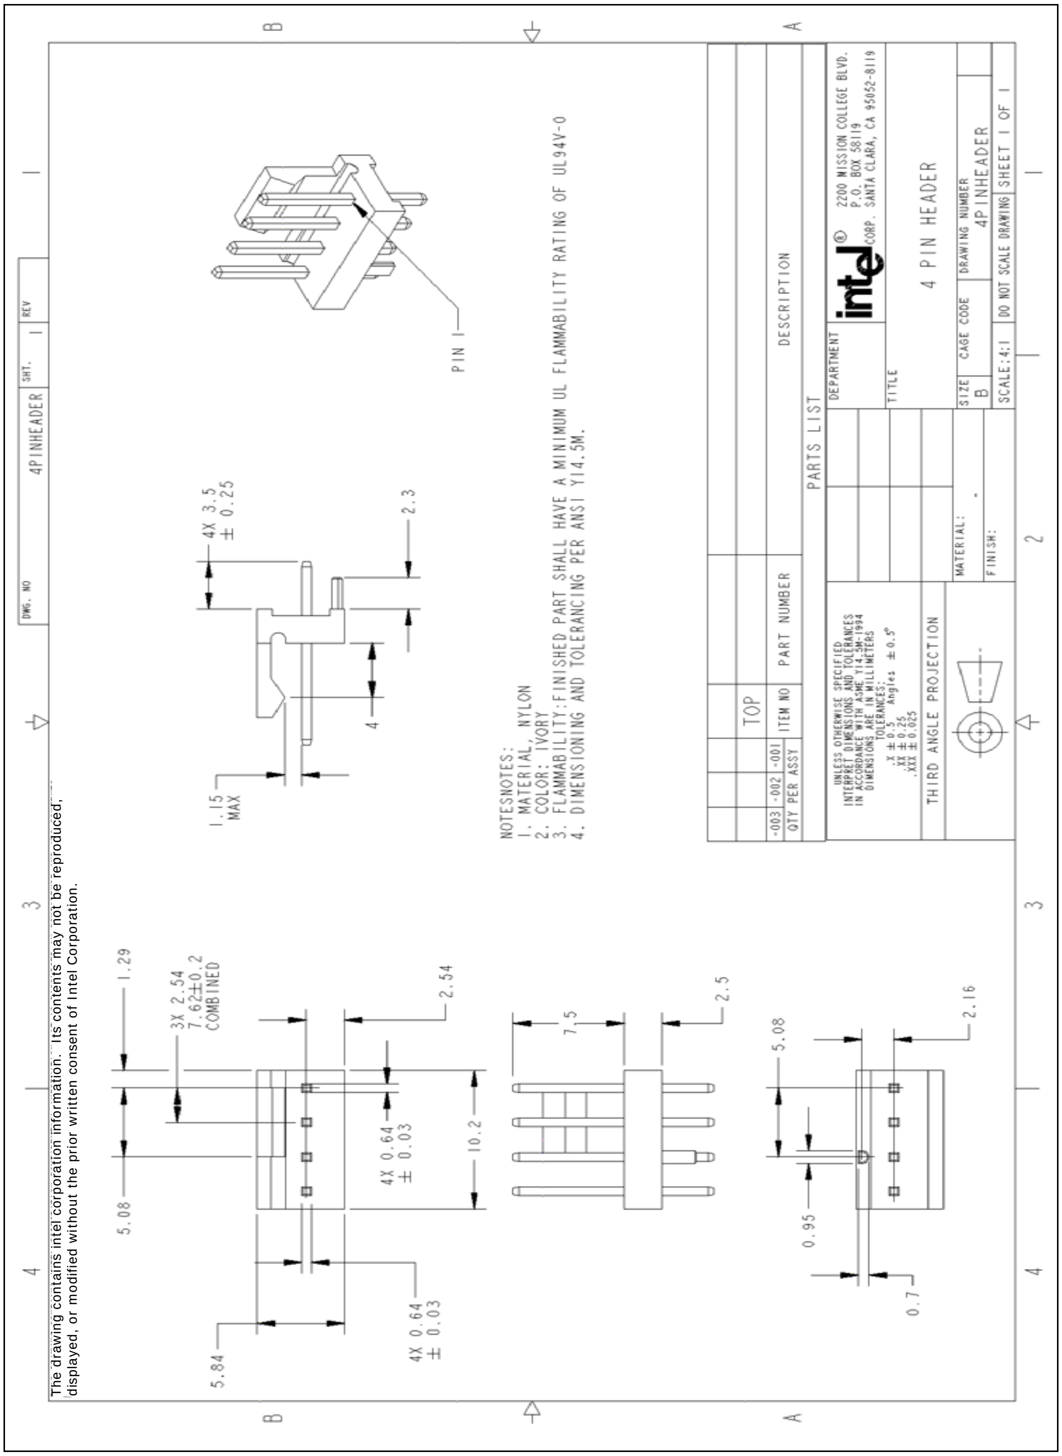

10.2.2 Boxed Processor Retention Mechanism and Heat Sink Support (URS)

Baseboards designed for use by a system integrator should include holes that are in proper alignment with each other to support the boxed processor. Refer to Figure 10-5 through Figure 10-8 for LGA1356 mounting hole dimensions.

LGA1356 Unified Retention System (URS) and the Unified Backplate Assembly. The URS is designed to extend air-cooling capability through the use of larger heat sinks with minimal airflow blockage and bypass. URS retention transfers load to the baseboard via the Unified Backplate Assembly. The URS spring, captive in the heatsink, provides the necessary compressive load for the thermal interface material. For specific design details on the URS and the Unified Backplate please refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide*.

All components of the URS heat sink solution will be captive to the heat sink and will only require a Phillips screwdriver to attach to the Unified Backplate Assembly. When installing the URS the screws should be tightened until they will no longer turn easily. This should represent approximately 8 inch-pounds of torque. More than that may damage the retention mechanism components.

10.3 Fan Power Supply [STS100C and STS100A]

The 4-pin PWM controlled thermal solution is being offered to help provide better control over pedestal chassis acoustics. This is achieved through more accurate measurement of processor die temperature through the processor's Digital Thermal Sensors. Fan RPM is modulated through the use of an ASIC located on the baseboard that sends out a PWM control signal to the 4th pin of the connector labeled as Control. This thermal solution requires a constant +12 V supplied to pin 2 of the active thermal solution and does not support variable voltage control or 3-pin PWM control. See Figure 10-13 and Table 10-1 for details on the 4-pin active heat sink solution connectors.

The fan power header on the baseboard must be positioned to allow the fan heat sink power cable to reach it. The fan power header identification and location must be documented in the suppliers platform documentation, or on the baseboard itself. The baseboard fan power header should be positioned within 177.8 mm [7 in.] from the center of the processor socket.

Table 10-1. PWM Fan Frequency Specifications For 4-Pin Active Thermal Solution

Description	Min Frequency	Nominal Frequency	Max Frequency	Unit
PWM Control Frequency Range	21,000	25,000	28,000	Hz

Figure 10-13. Fan Cable Connector Pin Out For 4-Pin Active Thermal Solution

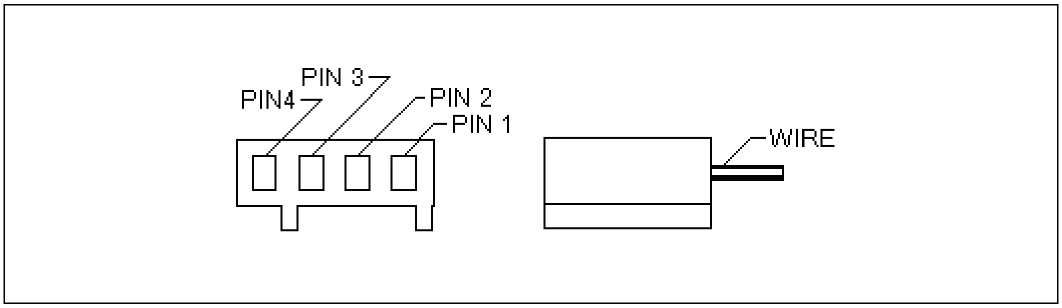

10.3.1 Boxed Processor Cooling Requirements

As previously stated the boxed processor will have three thermal solutions available. Each configuration will require unique design considerations. Meeting the processor's temperature specifications is also the function of the thermal design of the entire system, and ultimately the responsibility of the system integrator. The processor temperature specifications are found in Section 5, "Thermal Management Specifications" of this document.

10.3.1.1 STS100C(Passive / Active Combination Heat Sink Solution)

The active configuration of the combination solution is designed to help pedestal chassis users to meet the thermal processor requirements without the use of processor chassis ducting. However, it is strongly recommended to implement some form of air duct to meet memory cooling and processor T_{LA} temperature requirements. Use of the active configuration in a 2U rackmount chassis is not recommended.

In the passive configuration it is assumed that a chassis duct will be implemented.

For a list processor and thermal solution boundary conditions, such as Psi_{ca}, T_{LA}, airflow, flow impedance, etc, see Table 10-2. It is recommended that the ambient air temperature outside of the chassis be kept at or below 35 $^{\circ}$C. Meeting the processor's temperature specification is the responsibility of the system integrator.This thermal solution is for use with processor SKUs no higher than 95W (6 and 8 Core) or 80W (2 and 4 core).

10.3.1.2 STS100A (Active Heat Sink Solution) (Pedestal only)

The active configuration of the combination solution is designed to help pedestal chassis users to meet the thermal processor requirements without the use of processor chassis ducting. It is strongly recommended to implement some form of air duct to meet memory cooling and processor T_{LA} temperature requirements. Use of the active configuration in a 2U rackmount chassis is not recommended.

In the passive configuration it is assumed that a chassis duct will be implemented.

For a list processor and thermal solution boundary conditions, such as Psi_{ca}, T_{LA}, airflow, flow impedance, etc, Table 10-2. It is recommended that the ambient air temperature outside of the chassis be kept at or below 35 $^{\circ}$C. Meeting the processor's temperature specification is the responsibility of the system integrator.

This thermal solution is for use with processor SKUs no higher than 95W(6 and 8 Core), 80W (2 and 4 Core).

10.3.1.3 STS100P (25.5mm Tall Passive Heat Sink Solution) (Blade + 1U + 2U Rack)

This passive solution is intended for use in SSI Blade, 1U or 2U rack configurations. It is assumed that a chassis duct will be implemented in all configurations.

For a list processor and thermal solution boundary conditions, such as Psi_{ca}, T_{LA}, airflow, flow impedance, etc, see Table 10-2. It is recommended that the ambient air temperature outside of the chassis be kept at or below 35 °C. Meeting the processor's temperature specification is the responsibility of the system integrator.

This thermal solution is for use with processor SKUs no higher than 95W (6 and 8 Core), 80W (4 Core), or 80W (2 Core).

Note: Please refer to the *Intel® Xeon® Processor E5-2400 Processor Product Family Thermal/Mechanical Design Guide* for detailed mechanical drawings of the STS100P.

Table 10-2. Server Thermal Solution Boundary Conditions

TDP	Thermal Solution	ΨCA^2 (°C/W)	T_{LA} [1] (°C)	Airflow [3] (CFM) (inch of H_2O)	Delta P	Heatsink Volumetric[4] (mm)
95W-8/6 Core (Pedestal)	STS100A (with fan) STS100C (with fan)	8Core/6Core 0.268/0.272 0.167/0.171	53 62	Max RPM	NA	90x90x64
95W-8/6 Core (2U)	STS100A (without fan) STS100C (without fan)	0.298/0.302 0.173/0.177	50 62	26 26	0.070 0.140	90x90x64
95W-8/6 Core (1U)	STS100P	0.303/0.307	49	9.7	0.196	90x90x25.5
80W-4/2 Core (Pedestal)	STS100A (with fan) STS100C (with fan)	4Core/2Core 0.281/0.291 0.180/0.190	53 61	Max RPM	NA	90x90x64
80W-4/2 Core (2U)	STS100A (without fan) STS100C (without fan)	0.311/0.321 0.186/0.196	50 60	26 26	0.070 0.140	90x90x64
80W-4/2 Core (1U)	STS100P	0.316/0.326	49	9.7	0.196	90x90x25.5
70W-8 Core (Pedestal)	STS100A (with fan) STS100C (with fan)	0.267 0.166	52 59	Max RPM	NA	90x90x64
70W-8 Core (2U)	STS100A (without fan) STS100C (without fan)	0.297 0.172	50 59	26 26	0.070 0.140	90x90x64
70W-8 Core (1U)	STS100P	0.302	49	9.7	.196	90x90x25.5
60W-6 Core (Pedestal)	STS100A (with fan) STS100C (with fan)	0.268 0.167	52 58	Max RPM	NA	90x90x64
60W-6 Core (2U)	STS100A (without fan) STS100C (without fan)	0.298 0.173	50 57	26 26	0.070 0.140	90x90x64
60W-6 Core (1U)	STS100P	0.303	49	9.7	.196	90x90x25.5
45W-4 Core (Pedestal)	STS100A (with fan) STS100C (with fan)	0.276 0.175	50 55	Max RPM	NA	90x90x64
45W-4 Core (2U)	STS100A (without fan) STS100C (without fan)	0.306 0.181	49 54	26 26	0.070 0.140	90x90x64
45W-4 Core (1U)	STS100P	0.0311	49	9.7	.196	90x90x25.5

Notes:
1. Local ambient temperature of the air entering the heatsink or fan. System ambient and altitude are assumed 35°C and sea level.
2. Max target (mean + 3 sigma) for thermal characterization parameter.
3. Airflow through the heatsink fins with zero bypass. Max target for pressure drop (dP) measured in inches H2O.
4. See Table 10-2 for detailed dimensions. Dimensions of heatsinks do not include socket or processor.

10.4 Boxed Processor Contents

The Boxed Processor and Boxed Thermal Solution contents are outlined below.

Boxed Processor

- Intel® Xeon® processor E5-2400 product family
- Installation and warranty manual
- Intel Inside Logo

Boxed Thermal Solution

- Thermal solution assembly
- Thermal interface material (pre-applied)
- Installation and warranty manual

§