



64-bit Intel® Xeon™ Processor with 2MB L2 Cache

Thermal/Mechanical Design Guidelines

February 2005

Order Number: 306250-001



Notice: This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information. Contact your local Intel sales office or your distributor to obtain the latest specification before placing your product order.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The 64-bit Intel® Xeon™ Processor with 2MB L2 Cache processor may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available upon request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at <http://www.intel.com>.

Intel, Pentium and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © 2004-2005, Intel Corporation. All rights reserved.

* Other brands and names may be claimed as the property of others.



Contents

1	Introduction.....	7
1.1	Objective	7
1.2	Scope	7
1.3	References.....	7
1.4	Definition of Terms.....	8
2	Thermal/Mechanical Reference Design	11
2.1	Mechanical Requirements	11
2.1.1	Processor Mechanical Parameters.....	11
2.1.2	64-bit Intel® Xeon™ Processor with 2MB L2 Cache Package.....	12
2.1.3	64-bit Intel® Xeon™ Processor with 2MB L2 Cache Considerations.....	13
2.2	Thermal Requirements	14
2.2.1	Thermal Profile	14
2.2.2	TCONTROL Definition.....	16
2.2.3	Dual Thermal Profile Concept for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache.....	17
2.2.4	Performance Targets.....	18
2.2.5	Altitude.....	19
2.3	Characterizing Cooling Solution Performance Requirements.....	19
2.3.1	Fan Speed Control	19
2.3.2	Processor Thermal Characterization Parameter Relationships.....	20
2.3.3	Chassis Thermal Design Considerations	23
2.4	Thermal/Mechanical Reference Design Considerations	23
2.4.1	Heatsink Solutions.....	23
2.4.2	Thermal Interface Material.....	24
2.4.3	Summary	25
2.4.4	Assembly Overview of the Intel Reference Thermal Mechanical Design.....	25
2.4.5	Thermal Solution Performance Characteristics.....	27
2.4.6	Thermal Profile Adherence.....	29
2.4.7	Components Overview	31
2.4.8	Boxed Active Thermal Solution for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache.....	34
A	Mechanical Drawings	39
B	Safety Requirements.....	61
C	Quality and Reliability Requirements	63
C.1	Intel Verification Criteria for the Reference Designs	63
C.1.1	Reference Heatsink Thermal Verification	63
C.1.2	Environmental Reliability Testing	63
C.1.3	Material and Recycling Requirements.....	65
D	Enabled Suppliers Information	67
D.1	Supplier Information.....	67
D.1.1	Intel Enabled Suppliers.....	67
D.1.2	Additional Suppliers.....	68
E	Processor Thermal Management Logic and Thermal Monitor Features.....	69

E.1	Thermal Management Logic and Thermal Monitor Feature	69
E.1.1	Processor Power Dissipation	69
E.1.2	Thermal Monitor Implementation	69
E.1.3	Operation and Configuration	71
E.1.4	Thermal Monitor 2	71
E.1.5	System Considerations	73
E.1.6	Operating System and Application Software Considerations	73
E.1.7	Legacy Thermal Management Capabilities	73
E.1.8	On-Die Thermal Diode	73
E.1.9	Cooling System Failure Warning	76

Figures

2-1	64-bit Intel® Xeon™ Processor with 2MB L2 Cache Mechanical Drawing	12
2-2	Thermal Profile Diagram	15
2-3	TCONTROL and Thermal Profile Interaction	16
2-4	Dual Thermal Profile Diagram	17
2-5	Thermal Profiles A and B for the 64-bit intel® Xeon™ Processor with 2MB L2 Cache	18
2-6	TCONTROL and Fan Speed Control	20
2-7	Processor Thermal Characterization Parameter Relationships	22
2-8	Exploded View of CEK Thermal Solution Components	26
2-9	2U+ CEK Heatsink Thermal Performance	28
2-10	1U CEK Heatsink Thermal Performance	29
2-11	2U+ CEK Thermal Adherence to 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Thermal Profile A	30
2-12	1U CEK Thermal Adherence to 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Thermal Profile B	31
2-13	Isometric View of the 2U+ CEK Heatsink	32
2-14	Isometric View of the 1U CEK Heatsink	32
2-15	CEK Spring Isometric View	34
2-16	Isometric View of CEK Spring Attachment to the Base Board	34
2-17	Boxed Active CEK Heatsink Solutions with PWM/T-diode control (Representation Only)	35
2-18	Fan Cable Connection (Active CEK)	36
A-1	2U CEK Heatsink (Sheet 1 of 4)	40
A-2	2U CEK Heatsink (Sheet 2 of 4)	41
A-3	2U CEK Heatsink (Sheet 3 of 4)	42
A-4	2U CEK Heatsink (Sheet 4 of 4)	43
A-5	CEK Spring (Sheet 1 of 3)	44
A-6	CEK Spring (Sheet 2 of 3)	45
A-7	CEK Spring (Sheet 3 of 3)	46
A-8	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 1 of 6)	47
A-9	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 2 of 6)	48
A-10	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 3 of 6)	49
A-11	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 4 of 6)	50
A-12	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 5 of 6)	51
A-13	Baseboard Keepout Footprint Definition and Height Restrictions for Enabling	



	Components (Sheet 6 of 6).....	52
A-14	1U CEK Heatsink (Sheet 1 of 4)	53
A-15	1U CEK Heatsink (Sheet 2 of 4)	54
A-16	1U CEK Heatsink (Sheet 3 of 4)	55
A-17	1U CEK Heatsink (Sheet 4 of 4)	56
A-18	Active CEK Thermal Solution Volumetric (Sheet 1 of 3)	57
A-19	Active CEK Thermal Solution Volumetric (Sheet 2 of 3)	58
A-20	Active CEK Thermal Solution Volumetric (Sheet 3 of 3)	59
C-1	Random Vibration PSD	64
C-2	Shock Acceleration Curve	64
E-1	Thermal Sensor Circuit	70
E-2	Concept for Clocks under Thermal Monitor Control.....	71
E-3	Thermal Monitor 2 Frequency and Voltage Ordering.....	72
E-4	On-Die Thermal Diode Sensor Time Delay	75

Tables

1-1	Reference DocumentsIntroduction.fm	7
1-2	Terms and Descriptions	8
2-1	Processor Mechanical Parameters	11
2-2	Intel Reference Heatsink Performance Targets for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache.....	19
2-3	Fan Speed Control, TCONTROL and TDIODE Relationship.....	20
2-4	CEK Heatsink Thermal Mechanical Characteristics.....	33
2-5	Recommended Thermal Grease Dispense Weight	33
2-6	Fan Specifications (Boxed 4-wire PWM/T-diode Heatsink Solution)	36
2-7	Fan Cable Connector Pin Out (Active CEK)	36
A-1	Mechanical Drawing List	39
D-1	Suppliers for the 64-bit Intel(R) Xeon(TM) Processor with 2MB L2 Cache Intel Reference Solution	67
D-2	Processor Enabling Solutions – Additional Suppliers	68

Revision History

Document Number	Revision Number	Description	Date
306250	-001	Initial release of the document.	February 2005

§



1 Introduction

1.1 Objective

To describe the reference thermal solution and design parameters required for 64-bit Intel® Xeon™ Processor with 2MB L2 Cache. It is also the intent of this document to comprehend and demonstrate the processor cooling solution features and requirements. Furthermore, this document provides an understanding of the processor thermal characteristics, and discusses guidelines for meeting the thermal requirements imposed on the entire life of the processor. The thermal/mechanical solutions described in this document are intended to aid component and system designers in the development and evaluation of processor compatible thermal/mechanical solutions.

1.2 Scope

The thermal/mechanical solutions described in this document pertain only to a solution(s) intended for use with the 64-bit Intel Xeon Processor with 2MB L2 Cache in 1U, 2U, 2U+ and workstation form factors systems. This document contains the mechanical and thermal requirements of the processor cooling solution. In case of conflict, the data in the *64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet* supersedes any data in this document. Additional information is provided as a reference in the appendix section(s).

1.3 References

Material and concepts available in the following documents may be beneficial when reading this document.

Table 1-1. Reference DocumentsIntroduction.fm (Sheet 1 of 2)

Document	Comment
<i>European Blue Angel Recycling Standards</i>	http://www.blauer-engel.de
<i>Intel® Xeon™ Processor with 800 MHz System Bus Thermal/Mechanical Design Guidelines</i>	http://developer.intel.com/
<i>mPGA604 Socket Design Guidelines</i>	Note
<i>64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet</i>	Note
<i>64-bit Intel® Xeon™ Processor with 2MB L2 Cache Mechanical Models (in IGES and ProE* format)</i>	Available electronically
<i>64-bit Intel® Xeon™ Processor with 2MB L2 Cache Cooling Solution Mechanical Models (in IGES and ProE* format)</i>	Available electronically

Table 1-1. Reference DocumentsIntroduction.fm (Sheet 2 of 2)

64-bit Intel® Xeon™ Processor with 2MB L2 Cache Thermal Models (in Flotherm* and Icepak*)	Available electronically
64-bit Intel® Xeon™ Processor with 2MB L2 Cache Cooling Solution Thermal Models (in Flotherm* and Icepak* format)	Available electronically
Thin Electronics Bay Specification (A Server System Infrastructure (SSI) Specification for Rack Optimized Servers)	http://www.ssiforum.com

NOTE: Contact your Intel field sales representative for the latest revision and order number of this document.

1.4 Definition of Terms

Table 1-2. Terms and Descriptions (Sheet 1 of 2)

Term	Description
Bypass	Bypass is the area between a passive heatsink and any object that can act to form a duct. For this example, it can be expressed as a dimension away from the outside dimension of the fins to the nearest surface.
FMB	Flexible Motherboard Guideline: an estimate of the maximum value of a processor specification over certain time periods. System designers should meet the FMB values to ensure their systems are compatible with future processor releases.
FSC	Fan Speed Control
IHS	Integrated Heat Spreader: 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.
mPGA604	The surface mount Zero Insertion Force (ZIF) socket designed to accept the 64-bit Intel Xeon Processor with 2MB L2 Cache.
Offset	A value programmed into each processor during manufacturing that can be obtained by reading the IA_32_TEMPERATURE_TARGET MSR. This is a static and a unique value. Refer to the IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide for further details.
P _{MAX}	The maximum power dissipated by a semiconductor component.
Ψ_{CA}	Case-to-ambient thermal characterization parameter (psi). A measure of thermal solution performance using total package power. Defined as $(T_{CASE} - T_{LA}) / \text{Total Package Power}$. Heat source should always be specified for Ψ measurements.
Ψ_{CS}	Case-to-sink thermal characterization parameter. A measure of thermal interface material performance using total package power. Defined as $(T_{CASE} - T_S) / \text{Total Package Power}$.
Ψ_{SA}	Sink-to-ambient thermal characterization parameter. A measure of heatsink thermal performance using total package power. Defined as $(T_S - T_{LA}) / \text{Total Package Power}$.
T _{CASE}	The case temperature of the processor, measured at the geometric center of the topside of the IHS.
T _{CASE-MAX}	The maximum case temperature as specified in a component specification.
TCC	Thermal Control Circuit: Thermal monitor uses the TCC to reduce the die temperature by using clock modulation when the die temperature is very near its operating limits.
T _{CONTROL}	A processor unique value, which defines the lower end of the thermal profile and is targeted to be used in fan speed control mechanisms.
TDP	Thermal Design Power: Thermal solution should be designed to dissipate this target power level. TDP is not the maximum power that the processor/chipset can dissipate.
Thermal Monitor	A feature on the processor that can keep the processor's die temperature within factory specifications under nearly all conditions.

Table 1-2. Terms and Descriptions (Sheet 2 of 2)

Thermal Profile	Line that defines case temperature specification of a processor at a given power level.
TIM	Thermal Interface Material: The thermally conductive compound between the heatsink and the processor case. This material fills the air gaps and voids, and enhances the transfer of the heat from the processor case to the heatsink.
T_{LA}	The measured ambient temperature locally surrounding the processor. The ambient temperature should be measured just upstream of a passive heatsink or at the fan inlet for an active heatsink.
T_{SA}	The system ambient air temperature external to a system chassis. This temperature is usually measured at the chassis air inlets.
U	A unit of measure used to define server rack spacing height. 1U is equal to 1.75 in, 2U equals 3.50 in, etc.

§



2 Thermal/Mechanical Reference Design

2.1 Mechanical Requirements

The mechanical performance of the processor cooling solution must satisfy the requirements described in this section.

2.1.1 Processor Mechanical Parameters

Table 2-1. Processor Mechanical Parameters

Parameter	Minimum	Maximum	Unit	Notes
Volumetric Requirements and Keepouts				Refer to drawings in Appendix A
Heatsink Mass		1000 2.2	g lbs	
Static Compressive Load	44 10	222 50	N lbf	1,2,3,4
	44 10	288 65	N lbf	1,2,3,5
Dynamic Compressive Load	N/A NA	222 N + 0.45 kg * 100 G 50 lbf (static) + 1 lbm * 100 G	N lbf	1,3,4,6,7
	N/A NA	288 N + 0.45 kg * 100 G 65 lbf (static) + 1 lbm * 100 G	N lbf	1,3,5,6,7
Transient	N/A	445	N	1,3,8
		100	lbf	


NOTE: In the case of a discrepancy, the most recent processor datasheet supersedes targets listed in the above table.

1. These specifications apply to uniform compressive loading in a direction perpendicular to the IHS top surface.
2. This is the minimum and maximum static force that can be applied by the heatsink and retention solution to maintain the heatsink and processor interface.
3. These parameters are based on limited testing for design characterization. Loading limits are for the package only and do not include the limits of the processor socket.
4. This specification applies for thermal retention solutions that allow baseboard deflection.
5. This specification applies for thermal retention solutions that prevent baseboard deflection.
6. Dynamic loading is defined as an 11 ms duration average load superimposed on the static load requirement.
7. Experimentally validated test condition used a heatsink mass of 1 lbm (~0.45 kg) with 100 G acceleration measured at heatsink mass. The dynamic portion of this specification in the product application can have flexibility in specific values, but the ultimate product of mass times acceleration should not exceed this validated dynamic load (1 lbm x 100 G = 100 lb).
8. Transient loading is defined as a 2 second duration peak load superimposed on the static load requirement, representative of loads experienced by the package during heatsink installation.

2.1.2 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Package

The 64-bit Intel Xeon Processor with 2MB L2 Cache processor is packaged using the flip-chip micro pin grid array 4 (FC-mPGA4) package technology. Please refer to the *64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet* for detailed mechanical specifications. The 64-bit Intel Xeon Processor with 2MB L2 Cache Mechanical drawing, [Figure 2-1](#), provides the mechanical information for 64-bit Intel Xeon Processor with 2MB L2 Cache. The stack up height of the processor in the socket is shown in [Appendix A](#). The drawing is superseded with the drawing in the processor datasheet, should there be any conflicts.

Figure 2-1. 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Mechanical Drawing



The package includes an integrated heat spreader (IHS). The IHS transfers the non-uniform heat from the die to the top of the IHS, out of which the heat flux is more uniform and spread over a larger surface area (not the entire IHS area). This allows more efficient heat transfer out of the package to an attached cooling device. The IHS is designed to be the interface for contacting a heatsink. Details can be found in the *64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet*.

The processor connects to the baseboard through a 604-pin surface mount, zero insertion force (ZIF) socket. A description of the socket can be found in the *mPGA604 Socket Design Guidelines*.

The processor package has mechanical load limits that are specified in the processor datasheet and in [Table 2-1](#). These load limits should not be exceeded during heatsink installation, removal, mechanical stress testing, or standard shipping conditions. For example, when a compressive static load is necessary to ensure thermal performance of the Thermal Interface Material (TIM) between the heatsink base and the IHS, it should not exceed the corresponding specification given in the processor datasheet.

The heatsink mass can also add additional dynamic compressive load to the package during a mechanical shock event. Amplification factors due to the impact force during shock must be taken into account in dynamic load calculations. The total combination of dynamic and static compressive load should not then exceed the processor compressive dynamic load specified in the datasheet and in [Table 2-1](#) during a vertical shock. It is not recommended to use any portion of the processor substrate as a mechanical reference or load-bearing surface in either static or dynamic compressive load conditions.

2.1.3 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Considerations

An attachment mechanism must be designed to support the heatsink since there are no features on the mPGA604 socket to directly attach a heatsink. In addition to holding the heatsink in place on top of the IHS, this mechanism plays a significant role in the robustness of the system in which it is implemented, in particular:

- Ensuring thermal performance of the TIM applied between the IHS and the heatsink. TIMs, especially ones based on phase change materials, are very sensitive to applied pressure: the higher the pressure, the better the initial performance. TIMs such as thermal greases are not as sensitive to applied pressure. Refer to [Section 2.4.2](#) for information on tradeoffs made with TIM selection. Designs should consider possible decrease in applied pressure over time due to potential structural relaxation in enabled components.
- Ensuring system electrical, thermal, and structural integrity under shock and vibration events. The mechanical requirements of the attach mechanism depend on the weight of the heatsink and the level of shock and vibration that the system must support. The overall structural design of the baseboard and system must be considered when designing the heatsink attach mechanism. Their design should provide a means for protecting mPGA604 socket solder joints as well as preventing package pullout from the socket.

Note: The load applied by the attachment mechanism must comply with the package specifications, along with the dynamic load added by the mechanical shock and vibration requirements, as identified in [Section 2.1.1](#).

A potential mechanical solution for heavy heatsinks is the direct attachment of the heatsink to the chassis pan. In this case, the strength of the chassis pan can be utilized rather than solely relying on the baseboard strength. In addition to the general guidelines given above, contact with the baseboard surfaces should be minimized during installation in order to avoid any damage to the baseboard.

The Intel reference design for 64-bit Intel Xeon Processor with 2MB L2 Cache is using such a heatsink attachment scheme. Refer to [Section 2.4](#) for further information regarding the Intel reference mechanical solution.

2.2 Thermal Requirements

The operating thermal limits of the processor are defined by the Thermal Profile. The intent of the Thermal Profile specification is to support acoustic noise reduction through fan speed control and ensure the long-term reliability of the processor. This specification requires that the temperature at the center of the processor IHS, known as (T_{CASE}) remains within a certain temperature specification. Compliance with the T_{CASE} specification is required to achieve optimal operation and long-term reliability (See [Appendix B](#) for Case Temperature definition and measurement methods).

To ease the burden on thermal solutions, the Thermal Monitor feature and associated logic have been integrated into the silicon of the processor. One feature of the Thermal Monitor is the Thermal Control Circuit (TCC). When active, the TCC lowers the processor temperature by reducing the power consumed by the processor. This is done by changing the duty cycle of the internal processor clocks, resulting in a lower effective frequency. When active, the TCC turns the processor clocks off and then back on with a predetermined duty cycle.

By taking advantage of the Thermal Monitor features, system designers may reduce thermal solution cost by designing to the Thermal Design Power (TDP) instead of maximum power. TDP should be used for processor/chipset thermal solution design targets. TDP is not the maximum power that the processor/chipset can dissipate. TDP is based on measurements of processor power consumption while running various high power applications. This data set is used to determine those applications that are interesting from a power perspective. These applications are then evaluated in a controlled thermal environment to determine their sensitivity to activation of the thermal control circuit. This data set is then used to derive the TDP targets published in the processor datasheet. The Thermal Monitor can protect the processor in rare workload excursions above TDP. Therefore, thermal solutions should be designed to dissipate this target power level.

The relationship between TDP to the thermal profile, and thermal management logic and thermal monitor features, is discussed in the sections to follow. The thermal management logic and thermal monitor features are discussed in extensive detail in [Appendix E](#).

2.2.1 Thermal Profile

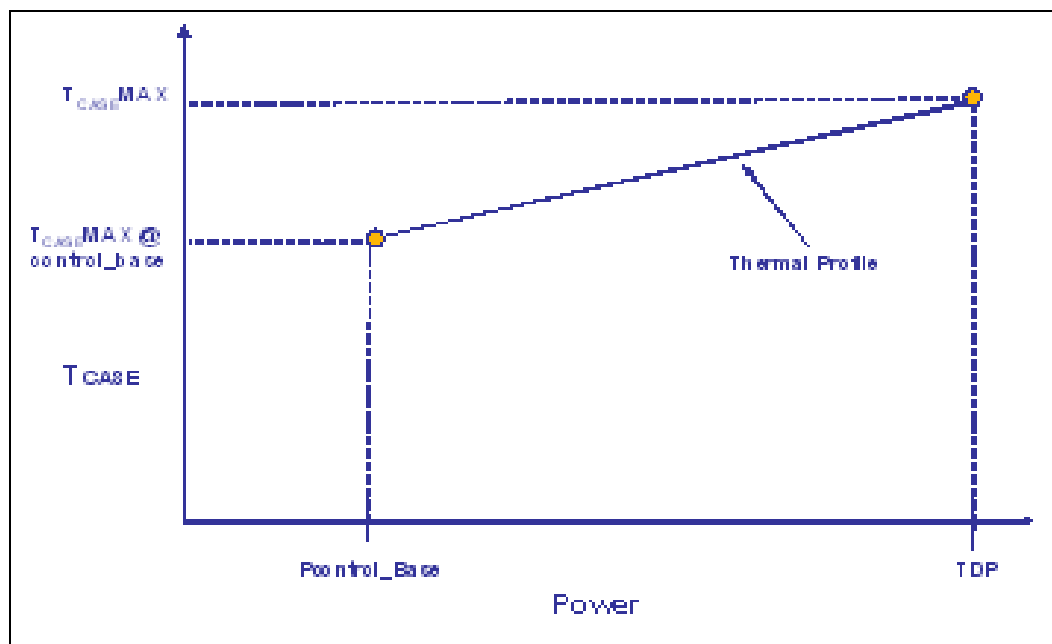
The thermal profile is a linear line that defines the relationship between a processor's case temperature and its power consumption as shown in [Figure 2-2](#). The equation of the thermal profile is defined as:

Equation 2-1. $y = ax + b$

Where:

y	=	Processor case temperature, T_{CASE} (°C)
x	=	Processor power consumption (W)
a	=	Case-to-ambient thermal resistance, Ψ_{CA} (°C/W)
b	=	Processor local ambient temperature, T_{LA} (°C)

Figure 2-2. Thermal Profile Diagram



The higher end point of the Thermal Profile represents the processor's TDP and the associated maximum case temperature ($T_{CASEMAX}$). The lower end point of the Thermal Profile represents the power value ($P_{CONTROL_BASE}$) and the associated case temperature ($T_{CASEMAX@P_{CONTROL_BASE}}$ for the lowest possible theoretical value of $T_{CONTROL}$ (see [Section 2.2.2](#)). The slope of the Thermal Profile line represents the case-to-ambient resistance of the thermal solution with the y-intercept being the local processor ambient temperature. The slope of the Thermal Profile is constant between $P_{CONTROL_BASE}$ and TDP, which indicate that all frequencies of a processor defined by the Thermal Profile, will require the same heatsink case-to-ambient resistance.

In order to satisfy the Thermal Profile specification, a thermal solution must be at or below the Thermal Profile line for the given processor when its diode temperature is greater than $T_{CONTROL}$ (refer to [Section 2.2.2](#)). The Thermal Profile allows the customers to make a trade-off between the thermal solution case-to-ambient resistance and the processor local ambient temperature that best suits their platform implementation (refer to [Section 2.2.3](#)). There can be multiple combinations of thermal solution case-to-ambient resistance and processor local ambient temperature that can meet a given Thermal Profile. If the case-to-ambient resistance and the local ambient temperature are known for a specific thermal solution, the Thermal Profile of that solution can easily be plotted against the Thermal Profile specification. As explained above, the case-to-ambient resistance represents the slope of the line and the processor local ambient temperature represents the y-axis intercept. Hence the T_{CASE} values of a specific solution can be calculated at the TDP and $P_{CONTROL_BASE}$ power levels. Once these points are determined, they can be joined by a line, which represents the Thermal Profile of the specific solution. If that line stays at or below the Thermal Profile specification, then that particular solution is deemed as a compliant solution.

2.2.2 T_{CONTROL} Definition

T_{CONTROL} is a temperature specification based on a temperature reading from the processor's thermal diode. T_{CONTROL} defines the lower end of the Thermal Profile line for a given processor, and it can be described as a trigger point for fan speed control implementation. The value for T_{CONTROL} is calibrated in manufacturing and configured for each processor individually. For the 64-bit Intel Xeon Processor with 2MB L2 Cache, the T_{CONTROL} value is obtained by reading a processor model specific register (MSR) and adding this offset value to a base value. The equation for calculating T_{CONTROL} is:

$$\text{Equation 2-2. } T_{\text{CONTROL}} = T_{\text{CONTROL_BASE}} + \text{Offset}$$

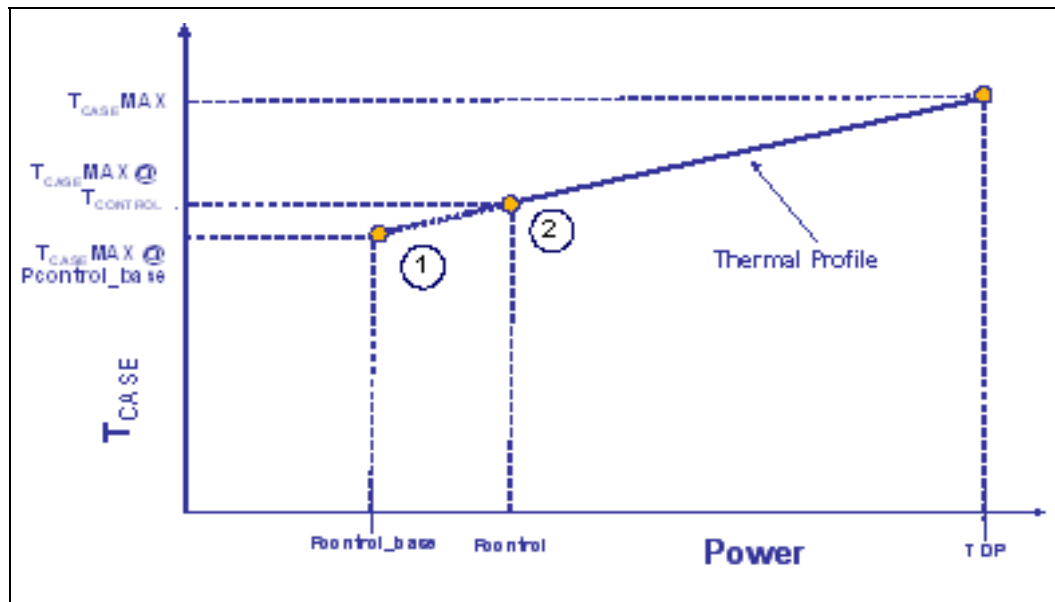
Where:

T_{CONTROL_BASE}=A fixed base value defined for a given processor generation as published in the processor datasheet.

Offset= A value programmed into each processor during manufacturing that can be obtained by reading the IA32_TEMPERATURE_TARGET MSR. This is a static and a unique value. Refer to the *IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide* for further details.

The T_{CONTROL_BASE} value for the 64-bit Intel Xeon Processor with 2MB L2 Cache is 50°C. The Offset value, which depends on several factors (i.e. leakage current), can be any number between 0 and (T_{CASE}^{MAX} - T_{CONTROL_BASE}). Figure 2-3 depicts the interaction between the Thermal Profile and T_{CONTROL} for an Offset value that is greater than 0 (i.e. T_{CONTROL} greater than T_{CONTROL_BASE}).

Figure 2-3. T_{CONTROL} and Thermal Profile Interaction



Since T_{CONTROL} is a processor diode temperature value, an equivalent T_{CASE} temperature must be determined to plot the T_{CASE}^{MAX} @ T_{CONTROL} point on the Thermal Profile graph. Location 1 on the Thermal Profile represents a T_{CASE} value corresponding to an Offset of 0 (the theoretical minimum for the given processor family). Any Offset value greater than 0 moves the point where the Thermal Profile must be met upwards, as shown by location 2 on the graph. If the diode temperature is less than T_{CONTROL}, then the case temperature is permitted to exceed the Thermal


Profile, but the diode temperature must remain at or below $T_{CONTROL}$. In other words, there is no T_{CASE} specification for the processor at power levels less than $P_{control}$. The thermal solution for the processor must be able to keep the processor's T_{CASE} at or below the T_{CASE} values defined by the Thermal Profile between the T_{CASE}^{MAX} @ $T_{CONTROL}$ and T_{CASE}^{MAX} points at the corresponding power levels.

Refer to [Section 2.3.1](#) for the implementation of the $T_{CONTROL}$ value in support of fan speed control (FSC) design to achieve better acoustic performance.

2.2.3 Dual Thermal Profile Concept for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache

The 64-bit Intel Xeon Processor with 2MB L2 Cache is designed to go into various form factors, including the volumetrically constrained 1U and custom blade form factors. Due to certain limitations of such form factors (i.e. airflow, thermal solution height), it is very challenging to meet the thermal requirements of the processor. To mitigate these form factor constraints, Intel has developed a dual Thermal Profile specification, shown in [Figure 2-4](#).

Figure 2-4. Dual Thermal Profile Diagram



The Thermal Profile A is based on Intel's 2U+ air cooling solution. Designing to Thermal Profile A ensures that no measurable performance loss due to Thermal Control Circuit (TCC) activation is observed in the processor. It is expected that TCC would only be activated for very brief periods of time when running a worst-case real world application in a worst-case thermal condition. These brief instances of TCC activation are not expected to impact the performance of the processor. A worst case real world application is defined as a commercially available, useful application which dissipates a power equal to, or above, the TDP for a thermally relevant timeframe. One example of a worst-case thermal condition is when a processor local ambient temperature is at or above 43 °C for 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Thermal Profile A.

Thermal Profile B supports volumetrically constrained platforms (i.e. 1U, blades, etc.), and is based on Intel's 1U air cooling solution. Because of the reduced capability represented by such thermal solutions, designing to Thermal Profile B results in an increased probability of TCC activation and an associated measurable performance loss. Refer to [Appendix E](#) for more details on

the Thermal Monitor features. Measurable performance loss is defined to be any degradation in the processor’s performance greater than 1.5%. The 1.5% number is chosen as the baseline since the run-to-run variation in a given performance benchmark is typically between 1 - 2%.


Although designing to Thermal Profile B results in increased T_{CASE} temperatures compared to Thermal Profile A at a given power level, both of these Thermal Profiles ensure that Intel’s long-term processor reliability requirements are satisfied. In other words, designing to Thermal Profile B does not impose any additional risk to Intel’s long-term reliability requirements. Thermal solutions that exceed Thermal Profile B specification are considered incompliant and will adversely affect the long-term reliability of the processor.

Refer to the *64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet* or [Section 2.2.4](#) for the Thermal Profile A and Thermal Profile B specifications. [Section 2.4](#) of this document also provides details on the 2U+ and 1U Intel reference thermal solutions that are designed to meet the 64-bit Intel Xeon Processor with 2MB L2 Cache Thermal Profile A and Thermal Profile B respectively.

2.2.4 Performance Targets

The Thermal Profile specifications for this processor are published in the *64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet*. These Thermal Profile specifications are shown as a reference in the subsequent discussions.

Figure 2-5. Thermal Profiles A and B for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache



Note: The thermal specifications shown in this graph are for reference only. Refer to the *64-bit Intel® Xeon™ Processor with 2MB L2 Cache Datasheet* for the Thermal Profile specifications. In case of conflict, the data information in the datasheet supersedes any data in this figure.

Table 2-2 describes thermal performance targets for the processor cooling solution enabled by Intel.

Table 2-2. Intel Reference Heatsink Performance Targets for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache

Thermal Solution Type	Target Thermal Profile	T _{LA} Assumption (°C)	TDP (W)	Thermal Performance Target, Ψ _{ca} (Mean + 3σ) (°C/W)
2U+ Form Factor	Thermal Profile A	40°C	110	0.293
1U Form Factor	Thermal Profile B	40°C	110	0.346

2.2.5 Altitude


The reference heatsink solutions will be evaluated at sea level (0 meters). The system designers who need to account for altitude effects in the overall system thermal design must make sure that either Thermal Profile A or B specification for the processor is met at the targeted altitude.

2.3 Characterizing Cooling Solution Performance Requirements

2.3.1 Fan Speed Control

Fan speed control (FSC) techniques to reduce system level acoustic noise are a common practice in server designs. The fan speed is one of the parameters that determine the amount of airflow provided to the thermal solution. Additionally, airflow is proportional to a thermal solution’s performance, which consequently determines the T_{CASE} of the processor at a given power level. Since the T_{CASE} of a processor is an important parameter in the long-term reliability of a processor, the FSC implemented in a system directly correlates to the processor’s ability to meet the Thermal Profile and hence the long-term reliability requirements. For this purpose, the parameter called T_{CONTROL} as explained in Section 2.2.2, is to be used in FSC designs to ensure that the long-term reliability of the processor is met while keeping the system level acoustic noise down. Figure 2-6 depicts the relationship between T_{CONTROL} and FSC methodology.

Figure 2-6. $T_{CONTROL}$ and Fan Speed Control



Once the $T_{CONTROL}$ value is determined as explained earlier, the thermal diode temperature reading from the processor can be compared to this $T_{CONTROL}$ value. A fan speed control scheme can be implemented as described in Table 2-3 without compromising the long-term reliability of the processor.

Table 2-3. Fan Speed Control, $T_{CONTROL}$ and T_{DIODE} Relationship

Condition	FSC Scheme
$T_{DIODE} \leq T_{CONTROL}$	FSC can adjust fan speed to maintain $T_{DIODE} = T_{CONTROL}$ (low acoustic region).
$T_{DIODE} > T_{CONTROL}$	FSC should adjust fan speed to keep T_{CASE} at or below the Thermal Profile specification (increased acoustic region).

There are many different ways of implementing fan speed control, including FSC based on processor ambient temperature, FSC based on processor thermal diode temperature (T_{DIODE}) or a combination of the two. If FSC is based only on the processor ambient temperature, low acoustic targets can be achieved under low ambient temperature conditions. However, the acoustics cannot be optimized based on the behavior of the processor temperature. If FSC is based only on the thermal diode, sustained temperatures above $T_{CONTROL}$, drives fans to maximum RPM. If FSC is based both on ambient and thermal diode, ambient temperature can be used to scale the fan RPM controlled by the thermal diode. This would result in an optimal acoustic performance. Regardless of which scheme is employed, system designers must ensure that the Thermal Profile specification is met when the processor diode temperature exceeds the $T_{CONTROL}$ value for a given processor.

2.3.2 Processor Thermal Characterization Parameter Relationships

The idea of a “thermal characterization parameter”, Ψ (psi), is a convenient way to characterize the performance needed for the thermal solution and to compare thermal solutions in identical conditions (heating source, local ambient conditions). A thermal characterization parameter is

convenient in that it is calculated using total package power, whereas actual thermal resistance, θ (theta), is calculated using actual power dissipated between two points. Measuring actual power dissipated into the heatsink is difficult, since some of the power is dissipated via heat transfer into the socket and board. Be aware, however, of the limitations of lumped parameters such as Ψ when it comes to a real design. Heat transfer is a three-dimensional phenomenon that can rarely be accurately and easily modeled by lump values.

The case-to-local ambient thermal characterization parameter value (Ψ_{CA}) is used as a measure of the thermal performance of the overall thermal solution that is attached to the processor package. It is defined by the following equation, and measured in units of °C/W:

$$\text{Equation 2-3. } \Psi_{CA} = (T_{CASE} - T_{LA}) / TDP$$

Where:

- Ψ_{CA} = Case-to-local ambient thermal characterization parameter (°C/W).
- T_{CASE} = Processor case temperature (°C).
- T_{LA} = Local ambient temperature in chassis at processor (°C).
- TDP = TDP dissipation (W) (assumes all power dissipates through the integrated heat spreader (IHS)).

The case-to-local ambient thermal characterization parameter of the processor, Ψ_{CA} , is comprised of Ψ_{CS} , the TIM thermal characterization parameter, and of Ψ_{SA} , the sink-to-local ambient thermal characterization parameter:

$$\text{Equation 2-4. } \Psi_{CA} = \Psi_{CS} + \Psi_{SA}$$

Where:


- Ψ_{CS} = Thermal characterization parameter of the TIM (°C/W).
- Ψ_{SA} = Thermal characterization parameter from heatsink-to-local ambient (°C/W).

Ψ_{CS} is strongly dependent on the thermal conductivity and thickness of the TIM between the heatsink and IHS.

Ψ_{SA} is a measure of the thermal characterization parameter from the bottom of the heatsink to the local ambient air. Ψ_{SA} is dependent on the heatsink material, thermal conductivity, and geometry. It is also strongly dependent on the air velocity through the fins of the heatsink.

Figure 2-7 illustrates the combination of the different thermal characterization parameters.

Figure 2-7. Processor Thermal Characterization Parameter Relationships



2.3.2.1 Example

The cooling performance, Ψ_{CA} , is then defined using the principle of thermal characterization parameter described above:

- Define a target case temperature $T_{CASE-MAX}$ and corresponding TDP at a target frequency, F , given in the processor datasheet.
- Define a target local ambient temperature at the processor, T_{LA} .

Since the processor thermal specifications ($T_{CASE-MAX}$ and TDP) can vary with the processor frequency, it may be important to identify the worse case (lowest Ψ_{CA}) for a targeted chassis (characterized by T_{LA}) to establish a design strategy such that a given heatsink can cover a given range of processor frequencies.

The following provides an illustration of how one might determine the appropriate performance targets. The example power and temperature numbers used here are not related to any Intel processor thermal specifications, and are for illustrative purposes only.

Assume the datasheet TDP is 85 W and the case temperature specification is 68 °C for a given frequency. Assume as well that the system airflow has been designed such that the local processor ambient temperature is 45°C. Then the following could be calculated using equation (2-3) from above for the given frequency:

$$\text{Equation 2-5. } \Psi_{CA} = (T_{CASE} - T_{LA}) / TDP = (68 - 45) / 85 = 0.27 \text{ } ^\circ\text{C/W}$$

To determine the required heatsink performance, a heatsink solution provider would need to determine Ψ_{CS} performance for the selected TIM and mechanical load configuration. If the heatsink solution were designed to work with a TIM material performing at $\Psi_{CS} \leq 0.05 \text{ } ^\circ\text{C/W}$, solving for equation (2-4) from above, the performance of the heatsink would be:

$$\text{Equation 2-6. } \Psi_{SA} = \Psi_{CA} - \Psi_{CS} = 0.27 - 0.05 = 0.22 \text{ } ^\circ\text{C/W}$$

If the local processor ambient temperature is assumed to be 40°C, the same calculation can be carried out to determine the new case-to-ambient thermal resistance:

Equation 2-7. $\Psi_{CA} = (T_{CASE} - T_{LA}) / TDP = (68 - 40) / 85 = 0.33 \text{ }^\circ\text{C/W}$

It is evident from the above calculations that, a reduction in the local processor ambient temperature has a significant positive effect on the case-to-ambient thermal resistance requirement.

2.3.3 Chassis Thermal Design Considerations

2.3.3.1 Chassis Thermal Design Capabilities and Improvements

One of the critical parameters in thermal design is the local ambient temperature assumption of the processor. Keeping the external chassis temperature fixed, internal chassis temperature rise is the only component that can affect the processor local ambient temperature. Every degree gained at the local ambient temperature directly translates into a degree relief in the processor case temperature.

Given the thermal targets for the processor, it is extremely important to optimize the chassis design to minimize the air temperature rise upstream to the processor (T_{rise}), hence minimizing the processor local ambient temperature.

The heat generated by components within the chassis must be removed to provide an adequate operating environment for both the processor and other system components. Moving air through the chassis brings in air from the external ambient environment and transports the heat generated by the processor and other system components out of the system. The number, size and relative position of fans, vents and other heat generating components determine the chassis thermal performance, and the resulting ambient temperature around the processor. The size and type (passive or active) of the thermal solution and the amount of system airflow can be traded off against each other to meet specific system design constraints. Additional constraints are board layout, spacing, component placement, and structural considerations that limit the thermal solution size.

In addition to passive heatsinks, fan heatsinks and system fans, other solutions exist for cooling integrated circuit devices. For example, ducted blowers, heat pipes and liquid cooling are all capable of dissipating additional heat. Due to their varying attributes, each of these solutions may be appropriate for a particular system implementation.

To develop a reliable, cost-effective thermal solution, thermal characterization and simulation should be carried out at the entire system level, accounting for the thermal requirements of each component. In addition, acoustic noise constraints may limit the size, number, placement, and types of fans that can be used in a particular design.

2.4 Thermal/Mechanical Reference Design Considerations

2.4.1 Heatsink Solutions

2.4.1.1 Heatsink Design Considerations

To remove the heat from the processor, three basic parameters should be considered:

- **The area of the surface on which the heat transfer takes place** - Without any enhancements, this is the surface of the processor package IHS. One method used to improve thermal performance is by attaching a heatsink to the IHS. A heatsink can increase the

effective heat transfer surface area by conducting heat out of the IHS and into the surrounding air through fins attached to the heatsink base.

- **The conduction path from the heat source to the heatsink fins** - Providing a direct conduction path from the heat source to the heatsink fins and selecting materials with higher thermal conductivity typically improves heatsink performance. The length, thickness, and conductivity of the conduction path from the heat source to the fins directly impact the thermal performance of the heatsink. In particular, the quality of the contact between the package IHS and the heatsink base has a higher impact on the overall thermal solution performance as processor cooling requirements become strict. Thermal interface material (TIM) is used to fill in the gap between the IHS and the bottom surface of the heatsink, and thereby improves the overall performance of the thermal stackup (IHS-TIM-Heatsink). With extremely poor heatsink interface flatness or roughness, TIM may not adequately fill the gap. The TIM thermal performance depends on its thermal conductivity as well as the pressure load applied to it. Refer to [Section 2.4.2](#) for further information on the TIM between the IHS and the heatsink base.
- **The heat transfer conditions on the surface on which heat transfer takes place** - Convective heat transfer occurs between the airflow and the surface exposed to the flow. It is characterized by the local ambient temperature of the air, T_{LA} , and the local air velocity over the surface. The higher the air velocity over the surface, the resulting cooling is more efficient. The nature of the airflow can also enhance heat transfer via convection. Turbulent flow can provide improvement over laminar flow. In the case of a heatsink, the surface exposed to the flow includes the fin faces and the heatsink base.

An active heatsink typically incorporates a fan that helps manage the airflow through the heatsink.

Passive heatsink solutions require in-depth knowledge of the airflow in the chassis. Typically, passive heatsinks see slower air speed. Therefore these heatsinks are typically larger (and heavier) than active heatsinks due to the increase in fin surface required to meet a required performance. As the heatsink fin density (the number of fins in a given cross-section) increases, the resistance to the airflow increases: it is more likely that the air will travel around the heatsink instead of through it, unless air bypass is carefully managed. Using air-ducting techniques to manage bypass area is an effective method for maximizing airflow through the heatsink fins.

2.4.2 Thermal Interface Material

TIM application between the processor IHS and the heatsink base is generally required to improve thermal conduction from the IHS to the heatsink. Many thermal interface materials can be pre-applied to the heatsink base prior to shipment from the heatsink supplier and allow direct heatsink attach, without the need for a separate TIM dispense or attach process in the final assembly factory.

All thermal interface materials should be sized and positioned on the heatsink base in a way that ensures the entire processor IHS area is covered. It is important to compensate for heatsink-to-processor attach positional alignment when selecting the proper TIM size.

When pre-applied material is used, it is recommended to have a protective application tape over it. This tape must be removed prior to heatsink installation.

The TIM performance is susceptible to degradation (i.e. grease breakdown) during the useful life of the processor due to the temperature cycling phenomena. For this reason, the measured T_{CASE} value of a given processor can decrease over time depending on the type of TIM material.

2.4.3 Summary

In summary, considerations in heatsink design include:

- The local ambient temperature T_{LA} at the heatsink, airflow (CFM), the power being dissipated by the processor, and the corresponding maximum T_{CASE} . These parameters are usually combined in a single lump cooling performance parameter, Ψ_{CA} (case to air thermal characterization parameter). More information on the definition and the use of Ψ_{CA} is given in [Section 2.4](#) and [Section 2.3.2](#).
- Heatsink interface (to IHS) surface characteristics, including flatness and roughness.
- The performance of the TIM used between the heatsink and the IHS.
- Surface area of the heatsink.
- Heatsink material and technology.
- Development of airflow entering and within the heatsink area.
- Physical volumetric constraints placed by the system.

2.4.4 Assembly Overview of the Intel Reference Thermal Mechanical Design

The reference design heatsinks that meet the 64-bit Intel Xeon Processor with 2MB L2 Cache thermal performance targets are called the Common Enabling Kit (CEK) heatsinks, and are available in 1U, 2U & 2U+ sizes. Each CEK consists of the following components:


- Heatsink (with captive standoff and screws)
- Thermal Interface Material (TIM-2)
- CEK Spring

2.4.4.1 Geometric Envelope

The baseboard keepout zones on the primary and secondary sides and height restrictions under the enabling component region are shown in detail in [Appendix A](#). The overall volumetric keep in zone encapsulates the processor, socket, and the entire thermal/mechanical enabling solution.

2.4.4.2 Assembly Drawing

Figure 2-8. Exploded View of CEK Thermal Solution Components



The CEK reference thermal solution is designed to extend air-cooling capability through the use of larger heatsinks with minimal airflow blockage and bypass. CEK retention solution can allow the use of much heavier heatsink masses compared to the legacy limits by using a load path directly attached to the chassis pan. The CEK spring on the secondary side of the baseboard provides the necessary compressive load for the thermal interface material. The baseboard is intended to be isolated such that the dynamic loads from the heatsink are transferred to the chassis pan via the stiff screws and standoffs. This reduces the risk of package pullout and solder-joint failures.

The baseboard mounting holes for the CEK solution are at the same location as the hole locations used for previous Intel Xeon processor with 800 MHz system bus thermal solution. However, CEK assembly requires 10.16 mm [0.400 in.] large diameter holes to compensate for the CEK spring embosses.

The CEK solution is designed and optimized for a baseboard thickness range of 1.57 – 2.31 mm. [0.062-0.093 in.]. While the same CEK spring can be used for this board thickness range, the heatsink standoff height is different for a 1.57 mm [0.062 in.] thick board than it is for a 2.31 mm. [0.093 in.] thick board. In the heatsink assembly, the standoff protrusion from the base of the heatsink needs to be 0.6 mm. [0.024 in.] longer for a 2.31 mm [0.093 in.] thick board, compared to a 1.57 mm [0.062 in.] thick board. If this solution is intended to be used on baseboards that fall outside of this range, then some aspects of the design, including but not limited to the CEK spring

design and the standoff heights, may need to change. Therefore, system designers need to evaluate the thermal performance and mechanical behavior of the CEK design on baseboards with different thicknesses.

Refer to [Appendix A](#) for drawings of the heatsinks and CEK spring. The screws and standoffs are standard components that are made captive to the heatsink for ease of handling and assembly.

Contact your Intel field sales representative for an electronic version of mechanical and thermal models of the CEK (Pro Engineer*, IGES and Icepak*, Flotherm* formats). Pro Engineer*, Icepak* and Flotherm* models are available on Intel Business Link (IBL).

Note: Intel reserves the right to make changes and modifications to the design as necessary.

Note: The thermal mechanical reference design for the 64-bit Intel Xeon Processor with 2MB L2 Cache was verified according to the Intel validation criteria given in [Appendix C.1](#). Any thermal mechanical design using some of the reference components in combination with any other thermal mechanical solution needs to be fully validated according to the customer criteria. Also, if customer thermal mechanical validation criteria differ from the Intel criteria, the reference solution should be validated against the customer criteria.

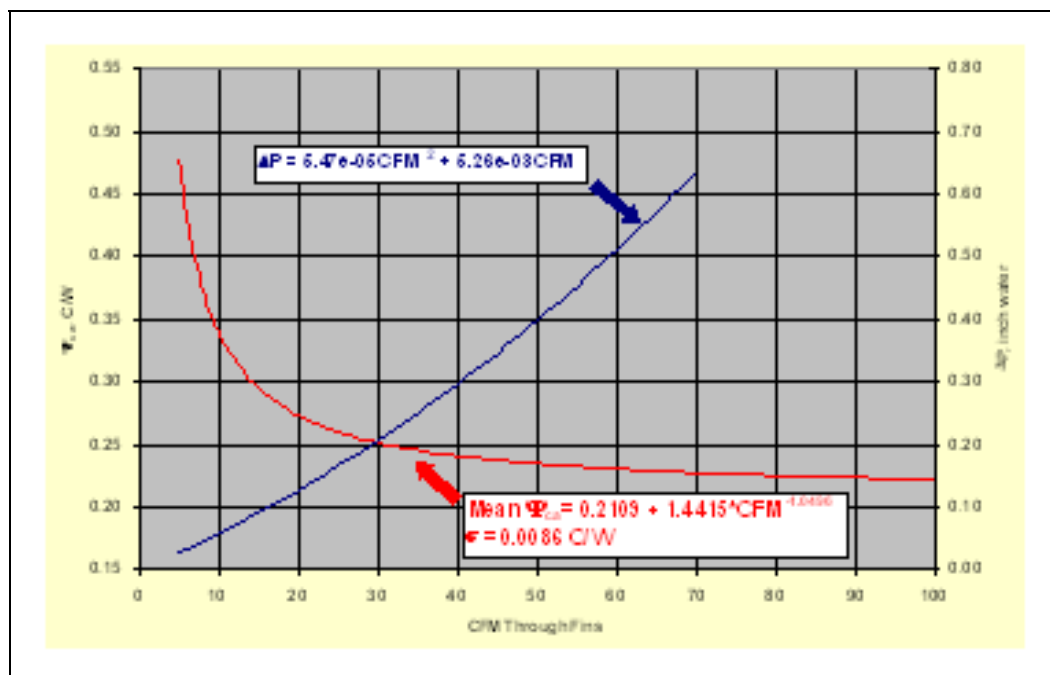
2.4.4.3 Structural Considerations of CEK

As Intel explores methods of keeping thermal solutions within the air-cooling space, the mass of the thermal solutions is increasing significantly. Due to the flexible nature (and associated large deformation) of baseboard-only attachments, Intel reference solutions, such as CEK, are now commonly using direct chassis attach (DCA) as the mechanical retention design. The mass of the new thermal solutions is large enough to require consideration for structural support and stiffening on the chassis.

2.4.5 Thermal Solution Performance Characteristics

The optimization of the CEK heatsinks for thermal performance has been completed. [Figure 2-9](#) and [Figure 2-11](#) show the performance of the 2U+ and 1U passive heatsinks, respectively. These figures show the thermal performance and the pressure drop through fins of the heatsink versus the airflow provided. The best-fit equations for these curves are also provided to make it easier for users to determine the desired value without any error associated with reading the graph.

Figure 2-9. 2U+ CEK Heatsink Thermal Performance



If other custom heatsinks are intended for use with the 64-bit Intel Xeon Processor with 2MB L2 Cache, they must support the following interface control requirements to be compatible with the reference mechanical components:


- **Requirement 1:** Heatsink assembly must stay within the volumetric keep-in.
- **Requirement 2:** Maximum mass and center of gravity.

Current maximum heatsink mass is 1000 grams [2.2 lbs] and the maximum center of gravity 38.1 mm [1.5 in.] above the bottom of the heatsink base.

- **Requirement 3:** Maximum and minimum compressive load.

Any custom thermal solution design should meet the loading specification as documented within this document, and should refer to the datasheet for specific details on package loading specifications.

Figure 2-10. 1U CEK Heatsink Thermal Performance



2.4.6 Thermal Profile Adherence

The 2U+ CEK Intel reference thermal solution is designed to meet the Thermal Profile A for the 64-bit Intel Xeon Processor with 2MB L2 Cache. From Table 2-2, the three-sigma (mean+3sigma) performance of the thermal solution is computed to be 0.293 °C/W and the processor local ambient temperature (T_{LA}) for this thermal solution is 40 °C. Hence, the Thermal Profile equation for this thermal solution is calculated as:

Equation 2-8. $y = 0.293x + 40$

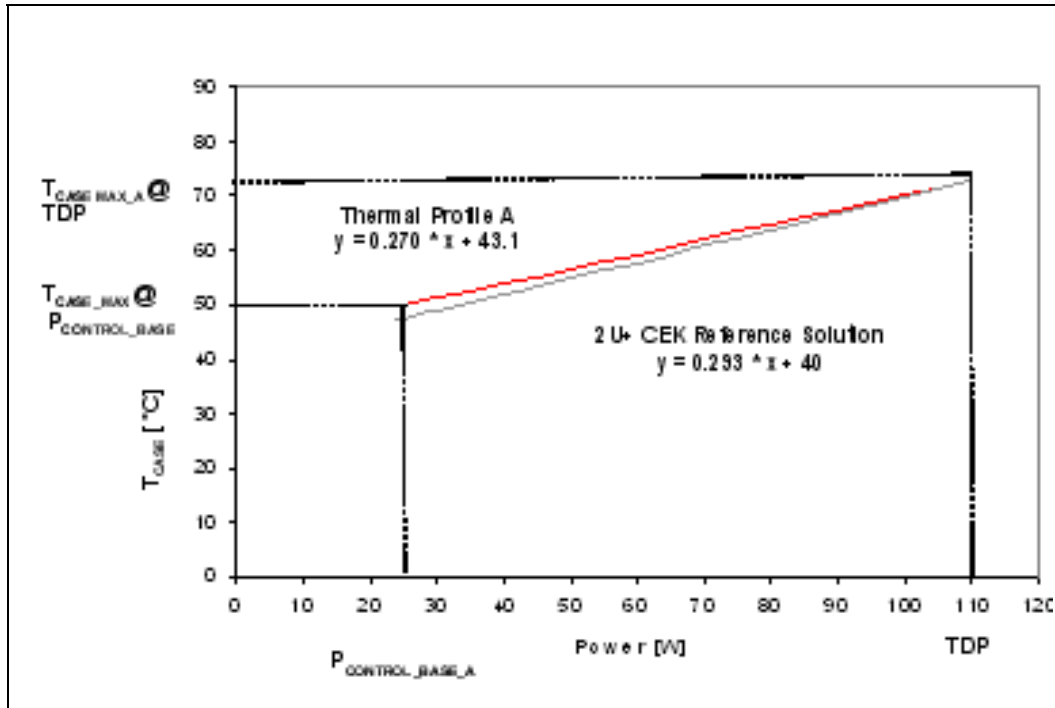
where,

y = Processor T_{CASE} value (°C)

x = Processor power value (W)

Figure 2-11 below shows the comparison of this reference thermal solution's Thermal Profile to the 64-bit Intel Xeon Processor with 2MB L2 Cache Thermal Profile A specification. The 2U+ CEK solution meets the Thermal Profile A with a 2.1 °C margin at the lower end (Pcontrol_base_A) and a 0.7 °C margin at the upper end (TDP). By designing to Thermal Profile A, it is ensured that no measurable performance loss due to TCC activation is observed under the given environmental conditions.

Figure 2-11. 2U+ CEK Thermal Adherence to 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Thermal Profile A



The 1U CEK Intel reference thermal solution is designed to meet the Thermal Profile B for the 64-bit Intel Xeon Processor with 2MB L2 Cache. From Table 2-2, the three-sigma (mean+3sigma) performance of the thermal solution is computed to be 0.346 °C/W and the processor local ambient temperature (T_{LA}) for this thermal solution is 40 °C. Hence, the Thermal Profile equation for this thermal solution is calculated as:

Equation 2-9. $y = 0.346x + 40$

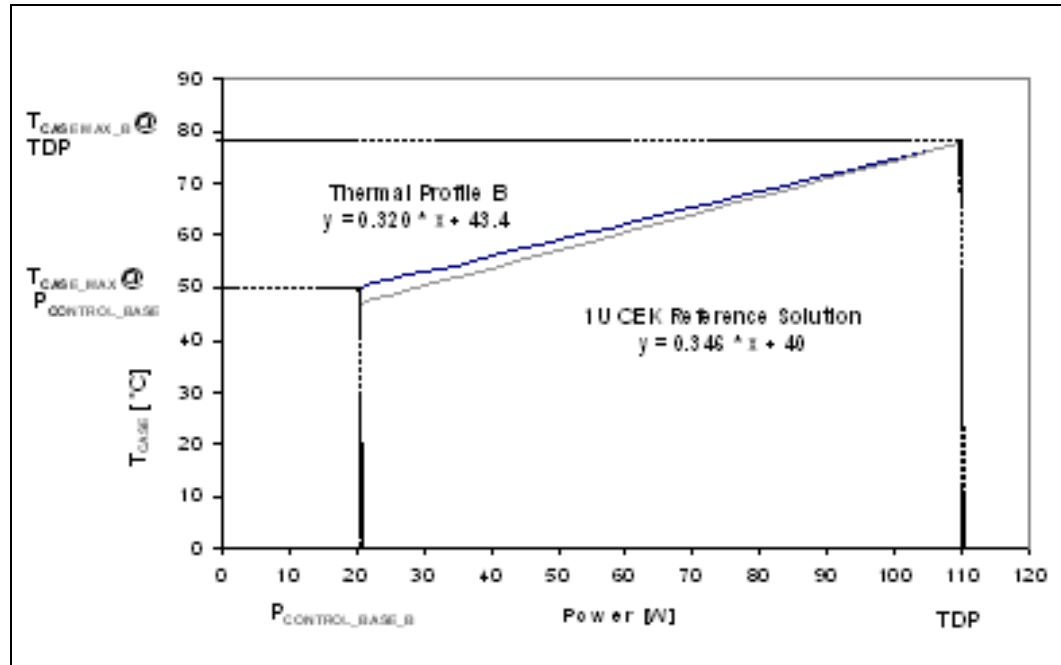
where,

y = Processor T_{CASE} value (°C)

x = Processor power value (W)

Figure 2-12 below shows the comparison of this reference thermal solution’s Thermal Profile to the 64-bit Intel Xeon Processor with 2MB L2 Cache Thermal Profile B specification. The 1U CEK solution meets the Thermal Profile B with a 2.4 °C margin at the lower end ($P_{control_base_B}$) and a 0.9 °C margin at the upper end (TDP). However, as explained in Section 2.2.3, designing to Thermal Profile B results in increased TCC activation and measurable performance loss for the processor. In this case, it is estimated that up to 5% of all the processors in a population that utilizes the 1U CEK reference solution may see TCC activation that results in a measurable performance loss of >1.5% when running an application that consumes power equivalent to TDP.


Figure 2-12. 1U CEK Thermal Adherence to 64-bit Intel® Xeon™ Processor with 2MB L2 Cache Thermal Profile B




2.4.7 Components Overview

2.4.7.1 Heatsink with Captive Screws and Standoffs

The CEK reference heatsink uses snapped-fin technology for its design. It consists of a copper base and copper fins with Shin-Etsu* G751 thermal grease as the TIM. The mounting screws and standoffs are also made captive to the heatsink base for ease of handling and assembly as shown in Figure 2-13 and Figure 2-14 for the 2U+ and 1U heatsinks, respectively.

Figure 2-13. Isometric View of the 2U+ CEK Heatsink

Note: Refer to [Appendix A](#) for more detailed mechanical drawings of the heatsink.

Figure 2-14. Isometric View of the 1U CEK Heatsink

Note: Refer to [Appendix A](#) for more detailed mechanical drawings of the heatsink.

The function of the standoffs is to provide a bridge between the chassis and the heatsink for attaching and load carrying. When assembled, the heatsink is rigid against the top of the standoff, and the standoff is rigid to a chassis standoff with the CEK spring firmly sandwiched between the two. In dynamic loading situations the standoff carries much of the heatsink load, especially in lateral conditions, when compared to the amount of load transmitted to the processor package. As such, it is comprised of steel. The distance from the bottom of the heatsink to the bottom of the standoff is 10.26 mm [0.404 in.] for a board thickness of 1.57 mm [0.062 in.]. The standoff will need to be modified for use in applications with a different board thickness, as defined in Chapter 2.4.4.2.

The function of the screw is to provide a rigid attach method to sandwich the entire CEK assembly together, activating the CEK spring under the baseboard, and thus providing the TIM preload. A screw is an inexpensive, low profile solution that does not negatively impact the thermal performance of the heatsink due to air blockage. Any fastener (i.e. head configuration) can be used as long as it is of steel construction; the head does not interfere with the heatsink fins, and is of the correct length of 20.64 mm [0.8125 in.].

Although the CEK heatsink fits into the legacy volumetric keep-in, it has a larger footprint due to the elimination of retention mechanism and clips used in the older enabled thermal/mechanical components. This allows the heatsink to grow its base and fin dimensions, further improving the thermal performance. A drawback of this enlarged size and use of copper for both the base and fins is the increased weight of the heatsink. The CEK heatsink is estimated to weigh twice as much as previous heatsinks used with Intel Xeon processors. However, the new retention scheme employed by CEK is designed to support heavy heatsinks (approximately up to 1000 grams) in cases of shock, vibration and installation as explained in [Appendix C](#). Some of the thermal and mechanical characteristics of the CEK heatsinks are shown in [Table 2-4](#).

Table 2-4. CEK Heatsink Thermal Mechanical Characteristics

Size	Height mm [in.]	Weight (kg) [lbs]	Target Airflow Through Fins (CFM)	Mean Ψ_{ca} ($^{\circ}\text{C}/\text{W}$)	Standard Deviation Ψ_{ca} ($^{\circ}\text{C}/\text{W}$)	Pressure Drop (in H_2O)
2U+	50.8 [2.00]	1.0 [2.2]	22	0.267	0.0086	0.142
1U	26.4 [1.04]	0.53 [1.2]	15	0.324	0.0073	0.384

2.4.7.2 Thermal Interface Material (TIM-2)

A TIM must be applied between the package and the heatsink to ensure thermal conduction. The CEK reference design uses Shin-Etsu* G751 thermal grease.

The recommended grease dispense weight to ensure full coverage of the processor IHS is given below. For an alternate TIM, full coverage of the entire processor IHS is recommended.

Table 2-5. Recommended Thermal Grease Dispense Weight

Processor	Recommended Thermal Grease	Dispense Weight (mg)
64-bit Intel Xeon Processor with 2MB L2 Cache	Shin-Etsu* G751	400

Additional recommendations include recalibrating the dispense controller settings after any two hour pause in grease dispense. The grease should be dispensed just prior to heatsink assembly to prevent any degradation in material performance. Finally, the thermal grease should be verified to be within its recommended shelf life before use.

The CEK reference solution is designed to apply a compressive load of up to 222 N [50 lbf] on the TIM to improve the thermal performance.

2.4.7.3 CEK Spring

The CEK spring, which is attached on the secondary side of the baseboard, is made from 0.80 mm [0.0315 in.] thick 301 stainless steel half hard. Any future versions of the spring will be made from a similar material. The CEK spring has four embosses which, when assembled, rest on the top of the chassis standoffs. The CEK spring is located between the chassis standoffs and the heatsink standoffs. The purpose of the CEK spring is to provide compressive preload at the TIM interface

when the baseboard is pushed down upon it. This spring does not function as a clip of any kind. The two tabs on the spring are used to provide the necessary compressive preload for the TIM when the whole solution is assembled. The tabs make contact on the secondary side of the baseboard. In order to avoid damage to the contact locations on the baseboard, the tabs will be insulated with a 0.127 mm [0.005 in.] thick Kapton* tape (or equivalent). [Figure 2-15](#) shows an isometric view of the CEK spring design.

Figure 2-15. CEK Spring Isometric View


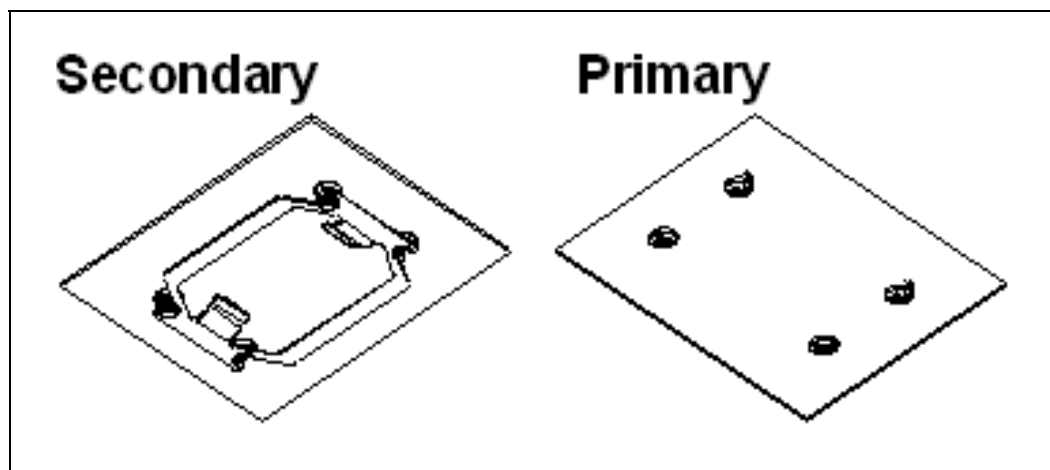


Figure 2-16. Isometric View of CEK Spring Attachment to the Base Board



Please refer to [Appendix A](#) for more detailed mechanical drawings of the CEK spring. Also, the baseboard keepout requirements shown in [Appendix A](#) must be met to use this CEK spring design.

2.4.8 Boxed Active Thermal Solution for the 64-bit Intel® Xeon™ Processor with 2MB L2 Cache

- In addition to the 1U and 2U passive CEK heatsinks, Intel is developing an active heatsink solution. This heatsink solution is primarily designed to be used in a pedestal chassis where sufficient air inlet space is present and side directional airflow is not an issue. All three heatsinks will be offered as part of boxed 64-bit Intel Xeon Processor with 2MB L2 Cache products. These solutions are intended for system integrators who build systems from components available through distribution channels. The retention solution used for these


products is called the Common Enabling Kit, or CEK. The CEK base is compatible with all three heatsink solutions.

[Figure 2-17](#) provides a representation of the active CEK solution. This design is based on a 4-pin PWM/T-diode controlled active fan heatsink solution. This new solution is being offered to help provide better control over pedestal chassis acoustics. This is achieved through accurate measurement of processor temperature through the processor's temperature diode (T-diode). Fan RPM is modulated through the use of an ASIC located on the serverboard, that sends out a PWM control signal to the 4th pin of the connector labeled as **Control**.

This heatsink solution also requires a constant +12 V supplied to pin 2 and does not support variable voltage control or 3-pin PWM control. If no PWM signal is detected on the 4th pin this heatsink solution will revert back to thermistor control mode, supporting both the 4-wire PWM and standard 3-wire ambient air control methods.

The solution is still under development at this time. Intel may make changes to specification and product descriptions at any time, without notice. The active heatsink solution will not exceed a mass of approximately 1150 grams. Note that this is per processor, so a dual processor system will have up to approximately 2300 grams total mass in the heatsinks. This large mass will require a minimum chassis stiffness to be met in order to withstand force during shock and vibration. Please refer to *Chassis Strength and Stiffness Measurement and Improvement Guidelines, for Direct Chassis Attach Solutions* for more details on chassis requirements.

Figure 2-17. Boxed Active CEK Heatsink Solutions with PWM/T-diode control (Representation Only)



Clearance is required around the heatsink to ensure unimpeded airflow for proper cooling. The physical baseboard keepout requirements for the active solution are the same as the passive CEK solution shown in [Appendix A](#). Refer to [Figure A-18](#) through [Figure A-20](#) for additional details on the active CEK thermal solution volumetrics.

2.4.8.1 Fan Power Supply

The active heatsink includes a fan, which requires a +12 V power supply. Platforms must provide a matched fan power header to support the boxed processor. [Table 2-6](#) contains specifications for the input and output signals at the heatsink fan connector.

The fan outputs a SENSE signal, an open-collector output, which pulses at a rate of two pulses per fan revolution. A baseboard pull-up resistor provides V_{CC} to match the baseboard-mounted fan speed monitor requirements, if applicable. Use of the SENSE signal is optional. If the SENSE signal is not used, pin 3 of the connector should be tied to GND.

It is recommended that a 4 pin fan header be used on the baseboard, in addition to, a control ASIC that can send a PWM signal to the active fan heatsink solution on the 4th pin, at a nominal 25 KHz frequency. If a 3-pin CPU fan header is used instead, the active fan heatsink solution will revert back to an automatic ambient air temperature control mode.

The fan power header on the baseboard must be positioned to allow the fan heatsink power cable to reach it. The fan power header identification and location must be documented in the supplier’s 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 2-6. Fan Specifications (Boxed 4-wire PWM/T-diode Heatsink Solution)

Description	Min	Typ Steady	Max Steady	Max Startup	Unit	Notes
+12V: 12 Volt Fan Power Supply	10.8	12	12	13.2	V	
IC: Fan Current Draw	N/A	1.5	1.25	1.5	A	
SENSE: SENSE Frequency	2	2	2	2	Pulses per fan revolution	Note

Note: System board should pull this pin up to V_{CC} with a resistor.

Figure 2-18. Fan Cable Connection (Active CEK)




Table 2-7. Fan Cable Connector Pin Out (Active CEK)

Pin Number	Signal	Color
1	Ground (Constant)	Black
2	Power (+12V)	Yellow
3	Signal: 2 pulses per revolution	Green
4	Control	Blue

2.4.8.2 Systems Considerations Associated with the Active CEK

This heatsink was designed to help pedestal chassis users to meet the thermal processor requirements without the use of chassis ducting. It may be necessary to implement some form of chassis air guide or air duct to meet the T_{LA} temperature of 40 °C depending on the pedestal chassis layout. Also, while the active heatsink solution is designed to mechanically fit into a 2U chassis, it may require additional space at the top of the heatsink to allow sufficient airflow into the heatsink fan. Therefore, additional design criteria may need to be considered if this heatsink is used in a 2U rack mount chassis, or in a chassis that has drive bay obstructions above the inlet to the fan heatsink.

Thermal Profile A should be used to help determine the thermal performance of the platform. The primary recommended control method for this solution is using pulse width modulation control. This control method requires the motherboard provide the correct PWM duty cycle to the active fan heatsink solution to properly follow the thermal profile. If no PWM signal is detected the active heatsink solution will default back to a thermistor controlled mode and the fan will automatically adjust fan RPM to meet the thermal profile.

It is critical to supply a constant +12 V to the fan header so that the active CEK heatsink solution can operate properly. If a system board has a jumper setting to select either a constant +12 V power to the fan header or a variable voltage, it is strongly recommended that the jumper be set by default to the constant +12 V setting.

It is recommended that the ambient air temperature outside of the chassis be kept at or below 35 °C. The air passing directly over the processor heatsink should not be preheated by other system components. Meeting the processor's temperature specification is the responsibility of the system integrator.

2.4.8.3 Boxed Processor Contents

A direct chassis attach method must be used to avoid problems related to shock and vibration, due to the weight of the heatsink required to cool the processor. The board must not bend beyond specification in order to avoid damage. The boxed processor contains the components necessary to solve both issues. The boxed processor will include the following items:

- 64-bit Intel Xeon Processor with 2MB L2 Cache
- Unattached (Active or Passive) Heatsink
- 4 screws, 4 springs, and 4 heatsink standoffs (all captive to the heatsink)
- Thermal Interface Material (TIM)
- Certificate of Authenticity (COA) and Manual
- Intel Inside® Logo

The other items listed in [Figure 2-8](#) that are required to complete this solution will be shipped with either the chassis or boards. They are as follows:

- CEK Spring (supplied by baseboard vendors)
- Chassis standoffs (supplied by chassis vendors)



A Mechanical Drawings

The mechanical drawings included in this appendix refer to the thermal mechanical enabling components for the 64-bit Intel® Xeon™ Processor with 1MB L2 Cache.

Note: Intel reserves the right to make changes and modifications to the design as necessary.

Table A-1. Mechanical Drawing List

Drawing Description	Figure Number
"2U CEK Heatsink (Sheet 1 of 4)"	Figure A-1
"2U CEK Heatsink (Sheet 2 of 4)"	Figure A-2
"2U CEK Heatsink (Sheet 3 of 4)"	Figure A-3
"2U CEK Heatsink (Sheet 4 of 4)"	Figure A-4
"CEK Spring (Sheet 1 of 3)"	Figure A-5
"CEK Spring (Sheet 2 of 3)"	Figure A-6
"CEK Spring (Sheet 3 of 3)"	Figure A-7
"Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 1 of 6)"	Figure A-8
"Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 2 of 6)"	Figure A-9
"Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 3 of 6)"	Figure A-10
"Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 4 of 6)"	Figure A-11
"Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 5 of 6)"	Figure A-12
"Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 6 of 6)"	Figure A-13
"1U CEK Heatsink (Sheet 1 of 4)"	Figure A-14
"1U CEK Heatsink (Sheet 2 of 4)"	Figure A-15
"1U CEK Heatsink (Sheet 3 of 4)"	Figure A-16
"1U CEK Heatsink (Sheet 4 of 4)"	Figure A-17
"Active CEK Thermal Solution Volumetric (Sheet 1 of 3)"	Figure A-18
"Active CEK Thermal Solution Volumetric (Sheet 2 of 3)"	Figure A-19
"Active CEK Thermal Solution Volumetric (Sheet 3 of 3)"	Figure A-20

Figure A-1. 2U CEK Heatsink (Sheet 1 of 4)




Figure A-3. 2U CEK Heatsink (Sheet 3 of 4)




Figure A-4. 2U CEK Heatsink (Sheet 4 of 4)




Figure A-6. CEK Spring (Sheet 2 of 3)




Figure A-7. CEK Spring (Sheet 3 of 3)




Figure A-8. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 1 of 6)




Figure A-10. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 3 of 6)




Figure A-12. Baseboard Keepout Footprint Definition and Height Restrictions for Enabling Components (Sheet 5 of 6)





Figure A-15. 1U CEK Heatsink (Sheet 2 of 4)





B Safety Requirements

Heatsink and attachment assemblies shall be consistent with the manufacture of units that meet the safety standards:

1. UL Recognition-approved for flammability at the system level. All mechanical and thermal enabling components must be a minimum UL94V-2 approved.
2. CSA Certification. All mechanical and thermal enabling components must have CSA certification.
3. Heatsink fins must meet the test requirements of UL1439 for sharp edges.

§





C **Quality and Reliability Requirements**

C.1 **Intel Verification Criteria for the Reference Designs**

C.1.1 **Reference Heatsink Thermal Verification**

The Intel reference heatsinks will be verified within specific boundary conditions using a TTV and the methodology described in the *Intel[®] Xeon[™] Processor Family Thermal Test Vehicle User's Guide*.

The test results, for a number of samples, are reported in terms of a worst-case mean + 3 σ value for thermal characterization parameter using real processors (based on the TTV correction offset).

C.1.2 **Environmental Reliability Testing**

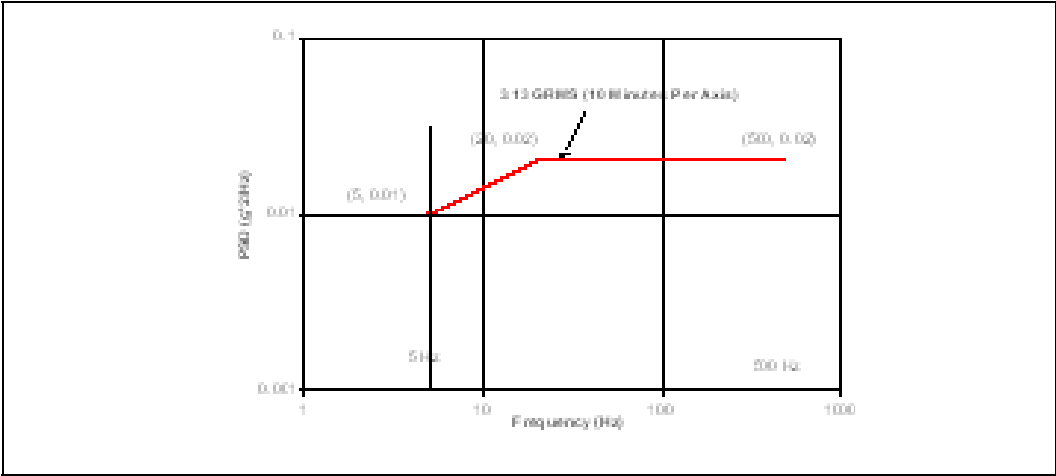
C.1.2.1 **Structural Reliability Testing**

Structural reliability tests consist of unpackaged, board-level vibration and shock tests of a given thermal solution in assembled state, as well as long-term reliability testing (temperature cycling, bake test). The thermal solution should be capable of sustaining thermal performance after these tests are conducted; however, the conditions of the tests outlined here may differ from the customers' system requirements.

C.1.2.2 **Random Vibration Test Procedure**

- Duration: 10 min/axis, 3 axes
- Frequency Range: 5 Hz to 500 Hz
- Power Spectral Density (PSD) Profile: 3.13 G RMS (refer to [Figure C-1](#)).

Figure C-1. Random Vibration PSD

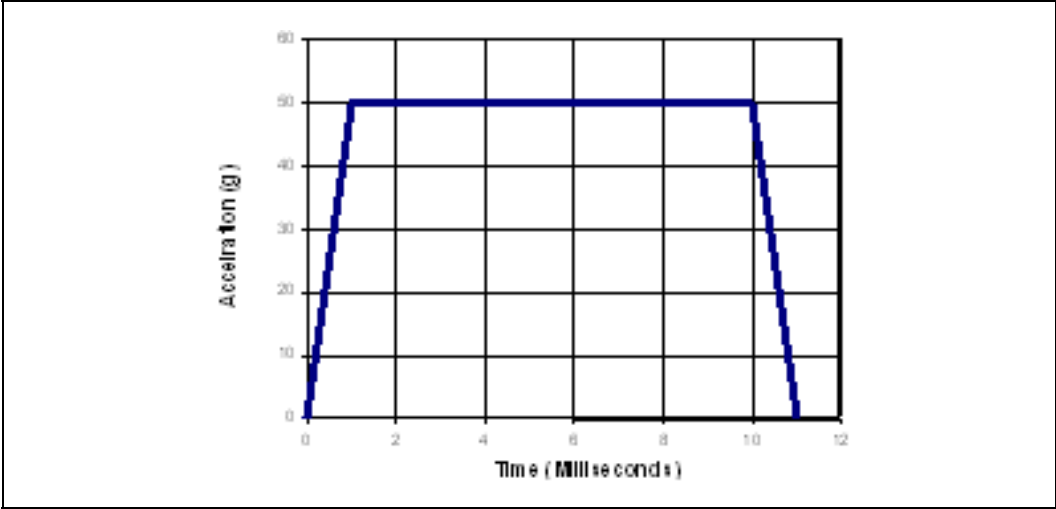


C.1.2.3 Shock Test Procedure

Recommended performance requirement for a baseboard:

- Quantity: 3 drops for + and – directions in each of 3 perpendicular axes (i.e. total 18 drops).
- Profile: 50 G trapezoidal waveform, 11 ms duration, 4.32 m/sec minimum velocity change.
- Setup: Mount sample board on test fixture.

Figure C-2. Shock Acceleration Curve



C.1.2.4 Recommended Test Sequence

Each test sequence should start with components (i.e. baseboard, heatsink assembly, etc.) that have not been previously submitted to any reliability testing.

The test sequence should always start with a visual inspection after assembly, and BIOS/Processor/memory test. The stress test should be then followed by a visual inspection and then BIOS/Processor/memory test.

C.1.2.5 Post-Test Pass Criteria

The post-test pass criteria are:

1. No significant physical damage to the heatsink and retention hardware.
2. Heatsink remains seated and its bottom remains mated flatly against the IHS surface. No visible gap between the heatsink base and processor IHS. No visible tilt of the heatsink with respect to the retention hardware.
3. No signs of physical damage on baseboard surface due to impact of heatsink.
4. No visible physical damage to the processor package.
5. Successful BIOS/Processor/memory test of post-test samples.
6. Thermal compliance testing to demonstrate that the case temperature specification can be met.

C.1.2.6 Recommended BIOS/Processor/Memory Test Procedures

This test is to ensure proper operation of the product before and after environmental stresses, with the thermal mechanical enabling components assembled. The test shall be conducted on a fully operational baseboard that has not been exposed to any battery of tests prior to the test being considered.

Testing setup should include the following components, properly assembled and/or connected:

- Appropriate system baseboard.
- Processor and memory.
- All enabling components, including socket and thermal solution parts.

The pass criterion is that the system under test shall successfully complete the checking of BIOS, basic processor functions and memory, without any errors. *Intel PC Diags* is an example of software that can be utilized for this test.

C.1.3 Material and Recycling Requirements

Material shall be resistant to fungal growth. Examples of non-resistant materials include cellulose materials, animal and vegetable based adhesives, grease, oils, and many hydrocarbons. Synthetic materials such as PVC formulations, certain polyurethane compositions (e.g. polyester and some polyethers), plastics which contain organic fillers of laminating materials, paints, and varnishes also are susceptible to fungal growth. If materials are not fungal growth resistant, then MIL-STD-810E, Method 508.4 must be performed to determine material performance.

Material used shall not have deformation or degradation in a temperature life test.

Any plastic component exceeding 25 grams must be recyclable per the European Blue Angel recycling standards.

§



D Enabled Suppliers Information

D.1 Supplier Information

D.1.1 Intel Enabled Suppliers

The Intel reference solutions have been verified to meet the criteria outlined in [Section D.1](#). Customers can purchase the Intel reference thermal solution components from the suppliers listed in [Table D-1](#).

To meet Thermal Profile B in 1U rack optimized servers, the Intel 64-bit Intel Xeon Processor with 2MB L2 Cache Reference Design (CEK604-1U-02) was optimized for thermal performance, while maintaining similar pressure drop characteristics as the *Intel® Xeon™ Processor with 800 MHz System Bus* Reference Design. Performance optimization involved modifications to the base thickness and fin density. For additional details, please refer to the 64-bit Intel Xeon Processor with 2MB L2 Cache drawings in Appendix A, and to the *Intel® Xeon™ Processor with 800 MHz System Bus Thermal/Mechanical Design Guidelines*.

Table D-1. Suppliers for the 64-bit Intel(R) Xeon(TM) Processor with 2MB L2 Cache Intel Reference Solution

Assembly	Component	Description	Development Suppliers	Supplier Contact Info
CEK604-2U-01 (for 2U, 2U+)	CEK Heatsink	Copper Fin, Copper Base	Fujikura* CNDA 36187 (stacked fin)	Mechatronics* Steve Carlson 800-453-4569 x205 steve@mechatronics.com
			Furukawa* CNDA 65755 (crimped fin)	Furukawa America* Katsu Mizushima (408) 232-9306 katsumizushima@mindspring.com
	Thermal Interface Material	Grease	Shin-Etsu G751 CNDA 75610	Donna Hartigan (480) 893-8898
	CEK Spring	Stainless Steel 301, Kapton* Tape on Spring Fingers	ITW Fastex* CNDA 78538	Ron Schmidt (847) 299-2222 rschmidt@itwfastex.com
AVC* CNDA 2085011			Felicia Lee 886-2-22996390 x144 felicia@avc.com.tw	
Foxconn* CNDA 11251				
CEK604-1U-02 (for 1U)	CEK Heatsink	Copper Fin, Copper Base	Fujikura* CNDA 36187 (stacked fin)	Mechatronics* Steve Carlson 800-453-4569 x205 steve@mechatronics.com
	Thermal Interface Material	See CEK604-2U-01		
	CEK Spring	See CEK604-2U-01		

D.1.2 Additional Suppliers

As mentioned in Appendix E.1.1, the Intel 64-bit Intel Xeon Processor with 2MB L2 Cache Reference Design (CEK604-1U-02) was optimized for thermal performance to meet Thermal Profile B in 1U rack optimized servers. Similar optimizations can be achieved with heatsinks designs from multiple suppliers of Intel Xeon processor with 800 MHz System Bus heatsinks.

Two suppliers have submitted samples and thermal test results demonstrating designs that meet the 64-bit Intel Xeon Processor with 2MB L2 Cache thermal performance requirements. Intel has performed limited thermal testing to confirm that the solutions meet thermal performance requirements. The solutions have not been thoroughly tested by Intel, and may or may not meet the 64-bit Intel Xeon Processor with 2MB L2 Cache enabling solution mechanical performance and quality and reliability requirements. Customers should contact the suppliers for further information on enabling solution performance, quality, and reliability. This supplier list has been summarized in Table E-2.

Table D-2. Processor Enabling Solutions – Additional Suppliers

Assembly	Component	Description	Development Suppliers	Supplier Contact Info
CEK604-1U-02 (for 1U)	CEK Heatsink	Copper Fin, Copper Base	Aavid Thermalloy* CNDA 36187 (stacked fin)	Scott Martineau 603-223-1710 martineau@aavid.com
			Heat Technology, Inc.* CNDA (folded fin)	Gil Dillon (978) 422-7100 gdillon@heattechnology.com

§

E Processor Thermal Management Logic and Thermal Monitor Features

E.1 Thermal Management Logic and Thermal Monitor Feature

E.1.1 Processor Power Dissipation

An increase in processor operating frequency not only increases system performance, but also increases the processor power dissipation. The relationship between frequency and power is generalized in the following equation: $P = CV^2F$ (where P = power, C = capacitance, V = voltage, F = frequency). From this equation, it is evident that power increases linearly with frequency and with the square of voltage. In the absence of power saving technologies, ever increasing frequencies will result in processors with power dissipations in the hundreds of watts. Fortunately, there are numerous ways to reduce the power consumption of a processor, and Intel is aggressively pursuing low power design techniques. For example, decreasing the operating voltage, reducing unnecessary transistor activity, and using more power efficient circuits can significantly reduce processor power consumption.

An on-die thermal management feature called Thermal Monitor is available on the 64-bit Intel Xeon Processor with 2MB L2 Cache. It provides a thermal management approach to support the continued increases in processor frequency and performance. By using a highly accurate on-die temperature sensing circuit and a fast acting temperature control circuit, the processor can rapidly initiate thermal management control. The Thermal Monitor can reduce cooling solution cost, by allowing designs to target TDP instead of maximum processor power.

E.1.2 Thermal Monitor Implementation

On the 64-bit Intel Xeon Processor with 2MB L2 Cache, the Thermal Monitor is integrated into the processor silicon. The Thermal Monitor includes:

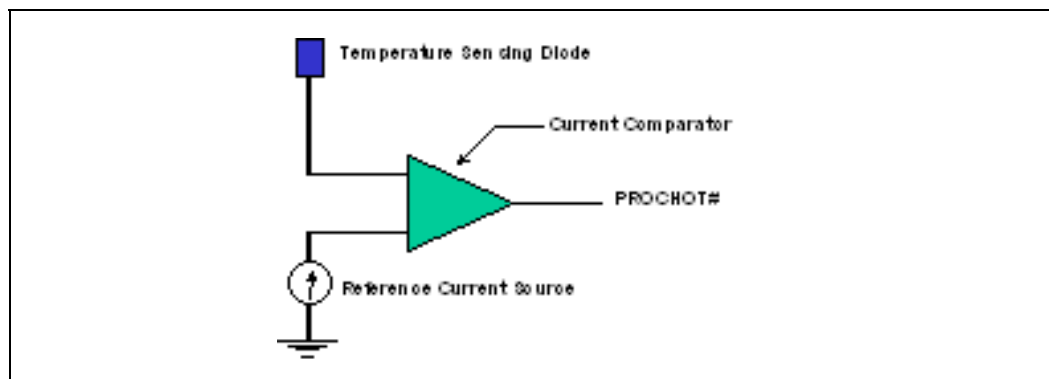
An on-die temperature sensing circuit.

- An external output signal (PROCHOT#) that indicates the processor has reached its maximum operating temperature.
- An external input signal (FORCEPR#) that allows the platform to force a power reduction by the processor by activating the TCC.
- A TCC that can reduce processor temperature by rapidly reducing power consumption when the on-die temperature sensor indicates that it has reached the maximum operating point.
- Registers to determine the processor thermal status.

The processor temperature is determined through an analog thermal sensor circuit comprised of a temperature sensing diode, a factory calibrated reference current source, and a current comparator (see [Figure E-1](#)). A voltage applied across the diode induces a current flow that varies with

temperature. By comparing this current with the reference current, the processor temperature can be determined. The reference current source corresponds to the diode current when at the maximum permissible processor operating temperature. Processors are calibrated during manufacturing on a small sample set. Once configured, the processor temperature at which the PROCHOT# signal is asserted (trip point) is not re-configurable.

Figure E-1. Thermal Sensor Circuit




The PROCHOT# signal is available internally to the processor as well as externally. External indication of the processor temperature status is provided through the bus signal PROCHOT#. When the processor temperature reaches the trip point, PROCHOT# is asserted. When the processor temperature is below the trip point, PROCHOT# is de-asserted. Assertion of the PROCHOT# signal is independent of any register settings within the processor. It is asserted any time the processor die temperature reaches the trip point. The point where the TCC activates is set to the same temperature at which the processor is tested and at which PROCHOT# asserts.

The TCC portion of the Thermal Monitor must be enabled for the processor to operate within specifications. The Thermal Monitor's TCC, when active, lowers the processor temperature by reducing the power consumed by the processor. This is done by changing the duty cycle of the internal processor clocks, resulting in a lower effective frequency. When active, the TCC turns the processor clocks off and then back on with a predetermined duty cycle. The duty cycle is processor specific, and is fixed for a particular processor. The maximum time period the clocks are disabled is $\sim 3 \mu\text{s}$, and is frequency dependent. Higher frequency processors will disable the internal clocks for a shorter time period. Figure E-2 illustrates the relationship between the internal processor clocks and PROCHOT#.

Performance counter registers, status bits in model specific registers (MSRs), and the PROCHOT# output pin are available to monitor and control the Thermal Monitor behavior. Details regarding the use of these registers will be described in the *IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide*.

Figure E-2. Concept for Clocks under Thermal Monitor Control



E.1.3 Operation and Configuration

To maintain compatibility with previous generations of processors, which have no integrated thermal logic, the TCC portion of Thermal Monitor is disabled by default. During the boot process, the BIOS must enable the TCC; or a software driver may do this after the operating system has booted. **Thermal Monitor or Thermal Monitor 2 feature must be enabled for the processor to remain within specification.**

The TCC feature can be configured and monitored in a number of ways. OEMs are expected to enable the TCC while using various registers and outputs to monitor the processor thermal status. The TCC is enabled by the BIOS setting a bit in an MSR (model specific register). Enabling the TCC allows the processor to maintain a safe operating temperature without the need for special software drivers or interrupt handling routines. When the TCC has been enabled, processor power consumption will be reduced within a few hundred clock cycles after the thermal sensor detects a high temperature, i.e. PROCHOT# assertion. The TCC and PROCHOT# transition to inactive once the temperature has been reduced below the thermal trip point, although a small time-based hysteresis has been included to prevent multiple PROCHOT# transitions around the trip point. External hardware can monitor PROCHOT# and generate an interrupt whenever there is a transition from active-to-inactive or inactive-to-active. PROCHOT# can also be configured to generate an internal interrupt which would initiate an OEM supplied interrupt service routine. Regardless of the configuration selected, PROCHOT# will consistently indicate the thermal status of the processor.

The TCC can also be activated manually using an “on-demand” mode.

E.1.4 Thermal Monitor 2

The 64-bit Intel Xeon Processor with 2MB L2 Cache also supports an enhanced TCC that works in conjunction with the existing Thermal Monitor logic. This capability is known as Thermal Monitor 2. This improved TCC provides a more efficient means for limiting the processor temperature by reducing the power consumption within the processor.

Note: Not all 64-bit Intel Xeon Processor with 2MB L2 Cache are capable of supporting Thermal Monitor 2. Details on which processor frequencies support Thermal Monitor 2 are provided in the *NDA Preliminary 64-bit Intel® Xeon™ Processor with 800 MHz System Bus (1MB and 2MB L2*

Cache Versions) Specification Update when available. For more details also refer to IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide.

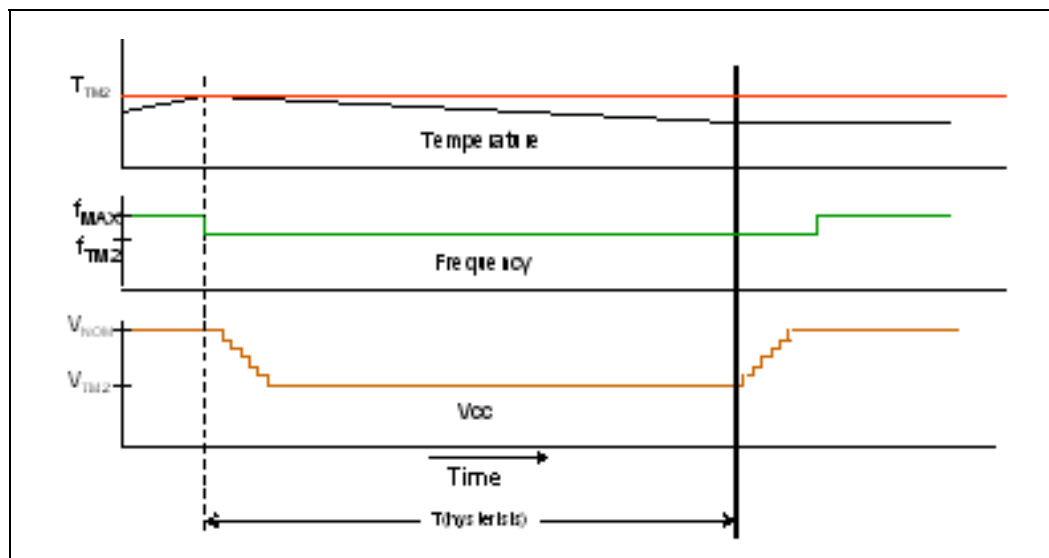
When Thermal Monitor 2 is enabled, and a high temperature situation is detected, the enhanced TCC will be activated. The enhanced TCC causes the processor to adjust its operating frequency (bus-to-core multiplier) and input voltage identification (VID) value. This combination of reduced frequency and the lowering of VID results in a reduction in processor power consumption.

A processor enabled for Thermal Monitor 2 includes two operating points, each consisting of a specific operating frequency and voltage. The first operating point represents the normal operating condition for the processor. The second operating point consists of both a lower operating frequency and voltage.

When the TCC is activated, the processor automatically transitions to the new frequency. This transition occurs very rapidly (on the order of 5 microseconds). During the frequency transition, the processor is unable to service any bus requests, and consequently, all bus traffic is blocked during the frequency transition. Edge-triggered interrupts will be latched and kept pending until the processor resumes operation at the new frequency.

Once the new operating frequency is engaged, the processor will transition to the new core operating voltage by issuing a new VID code to the voltage regulator. The voltage regulator must support dynamic VID changes in order to support Thermal Monitor 2. During the voltage change, it will be necessary to transition through multiple VID codes to reach the target operating voltage. Each step will be one VID table entry (i.e. 12.5 mV steps). The processor continues to execute instructions during the voltage transition. Operation at the lower voltage reduces both the dynamic and leakage power consumption of the processor. Once the processor has sufficiently cooled, and the time based hysteresis period has expired, the operating frequency and voltage transition back to the normal system operating point. Transition of the VID code will occur first, in order to insure proper operation once the processor reaches its normal operating frequency. Refer to Figure E-3 for an illustration of this ordering.

Figure E-3. Thermal Monitor 2 Frequency and Voltage Ordering



E.1.5 System Considerations

The Thermal Monitor feature may be used in a variety of ways, depending upon the system design requirements and capabilities.

- **Intel requires the TCC to be enabled for all 64-bit Intel Xeon Processor with 2MB L2 Cache-based systems.** At a minimum, the TCC provides an added level of protection against processor thermal solution failure.

A system designed to meet the TDP and T_{CASE} targets published in the processor datasheet greatly reduces the probability of real applications causing the TCC to activate under normal operating conditions. Systems that do not meet these specifications could be subject to more frequent activation of the TCC depending upon ambient air temperature and application power profile. Moreover, if a system is significantly under designed, there is a risk that the Thermal Monitor feature will not be capable of maintaining a safe operating temperature and the processor could shutdown and signal THERMTRIP#.

For information regarding THERMTRIP#, refer to [Section E.1.8.1](#) and to the processor datasheet.

E.1.6 Operating System and Application Software Considerations

The Thermal Monitor feature and its TCC work seamlessly with ACPI compliant operating systems and those utilizing hardware based timing routines. The Thermal Monitor feature is transparent to application software since the processor bus snooping, ACPI timer, and interrupts are active at all times.

Activation of the TCC during a non-ACPI aware operating system boot process may result in incorrect calibration of operating system software timing loops. This is also the case with operating systems that utilize execution based timing routines. The BIOS must disable the TCC prior to boot and then the operating system or BIOS must enable the TCC after the operating system boot process completes. Refer to the *IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide* for specific programming details.

Intel has worked with the major operating system vendors to ensure support for non-execution based operating system calibration loops and ACPI support for the Thermal Monitor feature.

E.1.7 Legacy Thermal Management Capabilities

In addition to Thermal Monitor, the 64-bit Intel Xeon Processor with 2MB L2 Cache supports the same thermal management features originally introduced with the Intel® Pentium® III Xeon™ processor. These features include the on-die thermal diode and THERMTRIP# signal for indicating catastrophic thermal failure.

E.1.8 On-Die Thermal Diode

There are two independent thermal diodes in the 64-bit Intel Xeon Processor with 2MB L2 Cache. One is the on-die thermal diode and the other is in the temperature sensor used for the Thermal Monitor and for THERMTRIP#. The Thermal Monitor's temperature sensor and the on-die thermal diode are independent and isolated devices with no direct correlation to one another. Circuit constraints and performance requirements prevent the Thermal Monitor's temperature sensor and the on-die thermal diode from being located at the same place on the silicon. The temperature distribution across the die may result in significant temperature differences between

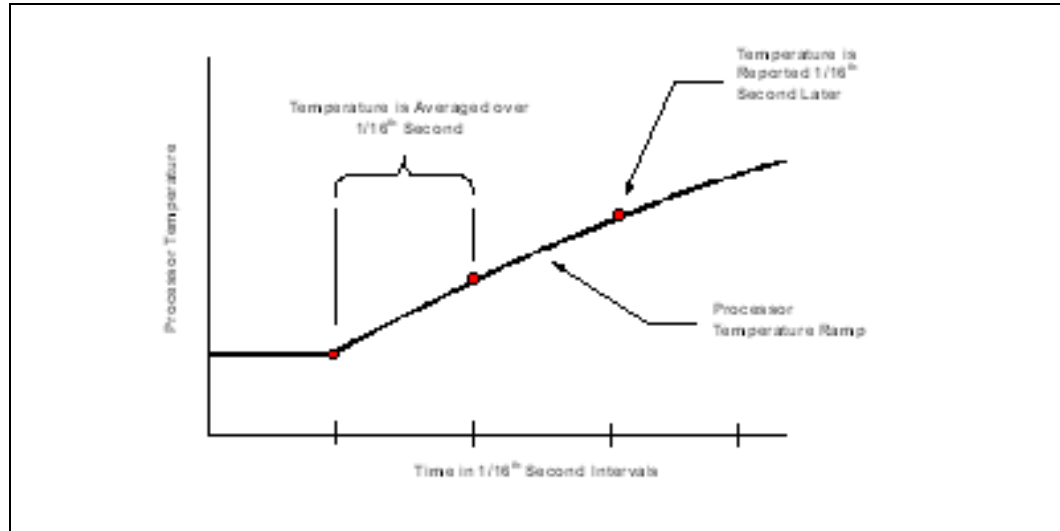
the on-die thermal diode and the Thermal Monitor's temperature sensor. This temperature variability across the die is highly dependent on the application being run. As a result, it is not possible to predict the activation of the TCC by monitoring the on-die thermal diode.

System integrators that plan on using the thermal diode for system or component level fan control need to be aware of the potential for rapid changes in processor power consumption as the executing workload changes. Variable performance thermal solutions that fail to react quickly to changing workloads may experience TCC activation or worst yet, result in automatic shutdown via THERMTRIP# (refer to [Section E.1.8.1](#) for more information on THERMTRIP). One example of this situation is as follows: A fan control scheme slows the fans such that the processor is operating very near the thermal trip point while executing a relatively low power workload. The start of a higher power application creates a sudden increase in power consumption and elevates the temperature of the processor above the trip point, causing the TCC to activate. The power reduction resulting from TCC activation slows the rate of temperature increase, but is not sufficient to clamp the temperature, due to inadequate thermal solution performance at reduced fan speed. As a result, the temperature continues to slowly increase. The fan is then sped up to compensate for the change in processor workload but reacts too slowly to prevent the processor from shutting down due to THERMTRIP# activation.

High temperature change rates on-die can also limit the ability to accurately measure the on-die thermal diode temperature. As a result, the on-die thermal diode should not be relied upon to warn of processor cooling system failure or predict the onset of the TCC. An illustration of this is as follows. Many thermal diode sensors report temperatures a maximum of 8 times per second. Within the $1/8^{\text{th}}$ (0.125 sec.) second time period, the temperature is averaged over $1/16^{\text{th}}$ of a second. In a scenario where the silicon temperature ramps at $50^{\circ}\text{C}/\text{sec}$, or approximately $6^{\circ}\text{C}/0.125$ sec, the processor will be $\sim 4.5^{\circ}\text{C}$ above the temperature reported by the thermal sensor. Change in diode temperature averaged over $1/16^{\text{th}}$ seconds = $\sim 1.5^{\circ}\text{C}$; temperature reported $1/16^{\text{th}}$ second later at $1/8^{\text{th}}$ second when the actual processor temperature would be 6°C higher (see [Figure E-4](#)).

The on-die thermal diode can be used with an external device (thermal diode sensor) to monitor long-term temperature trends. By averaging this data information over long time periods (hours/days vs. min/sec), it may be possible to derive a trend of the processor temperature. Analysis of this information could be useful in detecting changes in the system environment that may require attention. Design characteristics and usage models of the thermal diode sensors are described in datasheets available from the thermal diode sensor manufacturers.

Figure E-4. On-Die Thermal Diode Sensor Time Delay



E.1.8.1 THERMTRIP# Signal Pin

In the event of a catastrophic cooling failure, the processor will automatically shut down when the silicon temperature has reached its operating limit. At this point the system bus signal THERMTRIP# signal goes active and power must be removed from the processor. THERMTRIP# stays active until RESET# has been initiated. THERMTRIP# activation is independent of processor activity and does not generate any bus cycles.

E.1.8.2 FORCEPR# Signal Pin

The 64-bit Intel Xeon Processor with 2MB L2 Cache provides a means for system hardware to force activation of the TCC. One possible usage model would be to use this capability to protect the voltage regulator from overheating in order to avoid a catastrophic shutdown. Refer to the appropriate platform design guidelines and voltage regulator design guidelines for implementation details. The use of the FORCEPR# signal pin requires that BIOS code enable the signal's recognition via an MSR. Refer to *IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide* for details on enabling this feature.

E.1.9 Cooling System Failure Warning

§If desired, the system may be designed to cool the maximum processor power. In this situation, it may be useful to use the PROCHOT# signal as an indication of cooling system failure. Messages could be sent to the system administrator to warn of the cooling failure, while the TCC would allow the system to continue functioning or allow a graceful system shutdown. If no thermal management action is taken, the silicon temperature may exceed the operating limits, causing THERMTRIP# to activate and shut down the processor. Regardless of the system design requirements or thermal solution ability, the Thermal Monitor feature must still be enabled to ensure proper processor operation.

§