APPENDIX A

Technology and Terms

AC	Alternating current.
ACNT	A term commonly used to describe the IA32_APERF MSR used to calculate average frequency over a user-defined time window.
ACPI	The Advanced Configuration and Power Interface is used by BIOS to expose platform power management capabilities.
APIC	Advanced Programmable Interrupt Controller.
AR	The application ratio is used to describe the CPU logic switching rate of a workload.
ASHRAE	American Society of Heating, Refrigerating, and Air-Conditioning Engineers.
ASPM	Active State Power Management is a feature used to manage the power of PCIe links.
Avoton	Codename for Atom C2000-series SoC that follows Centerton.
AVX	Advanced vector extensions are integer and floating-point instructions used to improve performance.
Bin	A term used to describe the increase in frequency between any two P-states P*n* and P*n*-1.
BIOS	Basic Input/Output System refers to the firmware used to initialize a server.
BMC	The baseboard management controller is a dedicated microcontroller that provides remote monitoring and management functionality.
CC0/CC1/CC3/CC6	Describes a specific core-level C-state.
Centerton	Codename for Atom S1200-series processor.

(*continued*)

CKE	The clock enable signal is commonly used to identify rank-level power down modes for memory.
CLST	Closed Loop System Throttling is a power management feature that enables hardware protection using Node Manager and a PMBus power supply.
CLTT	Closed Loop Thermal Throttling utilizes temperature monitoring to manage memory thermal throttling.
CPI	Cycles per instruction is a basic performance metric.
CPU	Central processing unit.
CPUID	CPU identification instruction used to discover processor type and features.
CRAC	Computer room air conditioner.
CRAH	Computer room air handler.
CSR	A control and status register frequently used for power management monitoring or control.
C-state	An idle state where the processor has halted execution of instructions.
DC	Direct current.
DPC	The DIMMs per channel is a ratio used to describe the memory population of a platform.
DRAM	Dynamic Random Access Memory.
D-state	A low-power idle state for devices (PCIe, SATA).
DTSMAX	The maximum allowed temperature of the processor.
Dynamic Switching	A processor power management feature that automatically switches a platform to performance mode when capacity is high.
EBS	Event-based sampling is a monitoring technique that allows operators to associate power and performance events with the specific modules, functions, and lines of code that caused them.
ECC	Memory error correction used to provide protection from both transient errors and device failures.
EDP	The electrical design point
EEE	Energy efficient Ethernet is a low power mode that reduces PHY power.

EN	Identifies a platform or processor as Entry Level (Xeon E3 for one-socket servers).
Energy Perf Bias	A model-specific register used to control how aggressively power management features will be used.
Entry Latency	The time it takes to transition from an active to idle state (typically measured in microseconds).
EP	Identifies a platform or processor as efficient performance (Xeon E5 for two- to four-socket servers).
EPA	Environmental Protection Agency, responsible for the Energy Star program.
EX	Identifies a platform or processors as expandable (Xeon E7 for 4-socket and larger servers).
Exit Latency	The time it takes to transition from an idle to active state (typically measured in microseconds).
FIVR	A Fully Integrated Voltage Regulator is a high-current switching regulator integrated into the processor.
G-state	A global state that identifies the overall power state of a platform.
Haswell	Codename for the Xeon E5 v3 processor that follows Ivy Bridge.
HDD	A hard disk drive is a traditional spinning hard drive.
HIS	Integrated heat spreader.
HLT	Halt instruction used by an operating system to enter a C1 state.
HPC	High performance computing.
HSC	Hot swap controller.
HT	Hyper-Threading technology is Intel's implementation of simultaneous multithreading.
I/O	Input/output is used to describe capabilities for communication such as DDR, PCIe, and coherent interconnects such as QPI.
IA	The Intel Architecture term is commonly used to identify a hardware feature unique to Intel products.
IB	InfiniBand is a low-latency and high-throughput communications link frequently used in high performance computing.

(*continued*)

ICCMAX	The maximum current power delivery a platform can supply.
IMON	A voltage regulator current monitor used to measure power.
IPMI	The Intelligent Platform Management Interface is a specification and operating system independent interface for remote management.
ITD	Inverse temperature dependence.
Ivy Bridge	Codename for the Xeon E5 v2 processor that follows Sandy Bridge.
LDO	Low-dropout regulators used to provide variable voltages across cores in a processor with a single input voltage.
Linpack	An HPC benchmark derived from a collection of Fortran linear algebra routines.
L-state	A low-power idle state for interconnects (PCIe, DMI, QPI).
MBVR	A motherboard voltage regulator.
MCNT	A term commonly used to describe the IA32_MPERF MSR used to calculate average frequency over a user-defined time window.
MCP	A multi-chip package is where multiple chips are integrated together in the same package.
ME	The (Intel) Management Engine in the Platform Controller Hub used for monitoring, power capping, and hardware protection.
MMIO	Memory mapped I/O.
MSR	A model-specific register frequently used for power management monitoring or control.
MWAIT	A Monitor Wait instruction used by an operating system to enter a C1 or deeper C-state.
Nehalem	Codename for the Xeon 5500 processor.
NTB	Non-transparent bridging is a support technology used to create non-coherent interconnects between nodes using PCIe.
NUMA	Non-uniform memory access allocation provides contiguous memory regions for each processor's local memory.

NVMe	Non-Volatile Memory Express is a specification for directly connecting SSDs on PCIe that provides lower latency and higher performance than SAS and SATA.
OEM	Original equipment manufacturer.
OLTP	Online transaction processing.
OLTT	Open loop thermal throttling utilizes a static bandwidth limit to manage memory thermal throttling.
OS	Operating system.
OSPM	Operating system power management is a term commonly used to describe operating system power management policies and device drivers.
P1 frequency	Represents the CPU base frequency, guaranteed frequency, or the marked frequency of a CPU.
Path Length	Path length is a basic performance metric that measures the average number of instructions it takes to complete a single unit of work.
PC0/PC1/PC2/ PC3/PC6	Describes a specific package-level C-state.
PCH	Also known as South Bridge, Platform Controller Hub is a chipset connected to the processor that integrates many features that would otherwise require discrete controllers (such as storage, network, USB, management, and legacy).
PCIe	Peripheral Component Interconnect Express is a high-speed serial communication bus.
PCM	Performance Counter Monitor is a set of stand-alone tools used to collect core and uncore power and performance events.
PCPS	Per-core P-states allows individual cores that can each operate at their own frequency and voltage independent of what P-state other cores are in.
PCU	The power control unit is an internal microcontroller used to facilitate CPU power management.
PECI	The Platform Environment Control Interface is an interface for management controllers to communicate with the CPU.
PL1/PL2/PL3	A power level indicates a specific power limit used for power capping and power delivery protection.

(*continued*)

P-Limit (I/O)	A power management feature that allows the uncore to autonomously increase its uncore P-state to improve PCIe performance.
P-Limit (perf)	A power management feature that allows an idle socket to increase its uncore P-state to improve snoop and memory latency.
PLL	Phase-locked loop used to drive clocks.
PMBus	Power Management Bus is an open standard protocol used for power management of power supplies.
PMIC	A power management integrated circuit is applied to integrated circuits that have multiple power conversion controllers in one small package.
P*n* frequency	Represents the lowest frequency P-state or the most energy efficient frequency.
P-state	A performance state is an active state that represents a fixed frequency and voltage operating point.
PSU	Power supply unit.
PUE	Power usage effectiveness is defined as the ratio of the total energy use by the datacenter to that of the energy used by the IT equipment.
PWM	Pulse-width modulation.
QPI	QuickPath Interconnect is used for multi-socket communication.
RAPL	Running Average Power Limit is a power management feature used to maximize performance while meeting a specific thermal or power constraint.
RDTSC	Read time stamp counter instruction used by software to measure time.
Sandy Bridge	Codename for the Xeon E5 processor that follows Westmere.
SAS	Serial attached SCSI is a common protocol for connecting disks to a storage controller.
SATA	Serial ATA is a common protocol for connecting disks to a storage controller.
SEL	The System Event Log is a centralized event log used by management firmware.

Self-refresh	A low power memory power state where the DIMM itself is responsible for handling refresh.
SKU	The stock keeping unit term is commonly used to identify a CPU by its specific features (microarchitecture, core count, frequency, TDP).
SmaRT	Smart Ride Through is a technology that allows a server to function through momentary loss of AC power.
SMBus	The System Management Bus enables lightweight communication between platform devices.
SMT	Simultaneous multithreading.
SoC	A system on a chip is the coupling of the CPU with special-function hardware components in the same die.
SPEC	The Standard Performance Evaluation Corporation creates and maintains server benchmarks.
SRAM	Static random access memory.
SSD	Solid-state disk drive is a high-performance hard drive that stores data in flash memory chips.
S-state	A sleep state that powers down most platform components.
SVID	Serial VID is a serial communication bus between the processor package and the voltage regulator controllers.
TC0/TC1/TC3/TC6	Describes a specific thread-level C-state.
TCO	Total cost of ownership is a metric that estimates both the direct and indirect costs of a system.
TDP	A thermal design point specifies the amount of power that the CPU can consume, and therefore the amount of heat that the platform must be able to remove in order to avoid thermal throttling conditions.
THERMTRIP	A term used to describe the catastrophic trip temperature that, when exceeded, will result in immediate hardware shutdown.
TIM	Thermal interface material fills the air gaps between the component being cooled and a heat sink.
Tj	The junction temperature describes the internal temperature of the die.
TPC	The Transaction Processing Performance Council creates and maintains server benchmarks.

(*continued*)

TSC	Time Stamp Counter.
TSOD	Thermal sensor on die is a thermal sensor used in memory to measure temperature.
T-state	An active state where core execution is duty-cycled at a fixed interval for thermal, electrical, or power reasons.
Turbo frequency	Represents opportunistic frequency range about the CPU base frequency.
UFS	Uncore Frequency Scaling describes a power management feature that allows the uncore to maintain its own P-state.
UMA	Uniform memory access allocation interleaves every other cache line across each processor's local memory.
Uncore	A term commonly used to describe processor on-die logic outside of the cores.
USB	Universal Serial Bus.
Vmin	The minimum voltage used for the lowest frequency P-state.
VMM	Virtual machine monitor.
VR	Voltage regulator.
Vret	The retention voltage required to maintain state in a circuit.
VT	Virtualization technology is a term commonly used to describe technologies used to improve performance in a virtualized environment.
Westmere	Codename for the Xeon 5600 processor that follows Nehalem.
Xeon E3	Processor type used in one-socket servers for workloads with low compute requirements.
Xeon E5	Processor type used in two-socket servers for most workloads.
Xeon E7	Processor type used in 2-socket to 256-socket servers for mission critical and scale-up workloads.
Xeon Phi	Coprocessors used to accelerate workstation and cluster performance typically used in HPC.