Intel® Xeon® Processor E5-2600 V2 Product Family Technical Overview

Contents

1. Executive Summary

The Intel® Xeon® processor E5-2600 V2 product family, codenamed "Ivy Bridge EP", is a 2-socket platform based on Intel's most recent microarchitecture. Ivy Bridge is the 22-nanometer shrink of the Intel® Xeon® processor E5-2600 (codenamed "Sandy Bridge EP") microarchitecture. This product brings additional capabilities for data centers: more cores and more memory bandwidth. As a result, platforms based on the Intel Xeon processor E5-2600 V2 product family will yield up to 50% improvement in performance[1] compared to the previous generation "Sandy Bridge EP".

2. Introduction

The Intel Xeon processor E5-2600 V2 product family is based on Ivy Bridge EP microarchitecture, an enhanced version of the Sandy Bridge EP microarchitecture (http://software.intel.com/en-us/articles/intel-xeon-processor-e5-26004600-product-family-technical-overview). The platform supporting the Intel Xeon processor E5-2600 V2 product family is named "Romley." This paper discusses the new features available in the Intel Xeon processor E5-2600 V2 product family compared to the Intel Xeon processor E5-2600 product family. Each section includes information about what developers need to do to take advantage of new features for improving application performance and security.

3. Intel Xeon processor E5-2600 V2 product family enhancements

Some of the new features that come with the Intel Xeon processor E5-2600 V2 product family include:

1. 22-nm process technology
2. Security: Intel® Secure Key (DRNG)
3. Security: Intel® OS Guard (SMEP)
4. Intel® Advanced Vector Extensions (Intel® AVX): Float 16 Format Conversion
5. Virtualization: APIC Virtualization (APICv)
6. PCI Express* (PCIe): Support for atomic operation, x16 Non Transparent Bridge

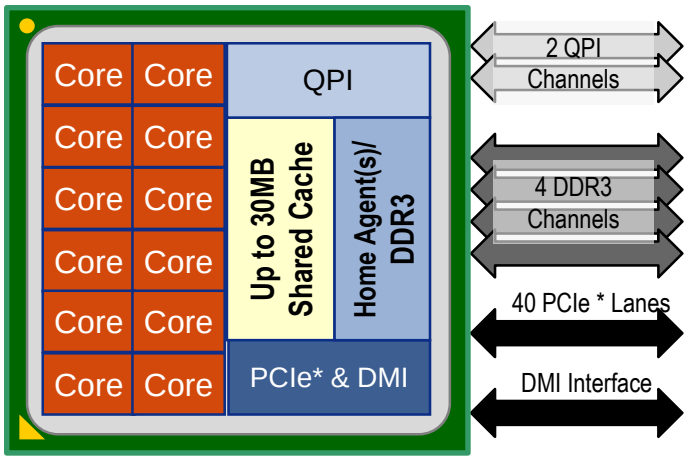

Figure 1. The Intel® Xeon® processor E5-2600 V2 product family Microarchitecture

Figure 1 shows a block diagram of the Intel Xeon processor E5-2600 V2 product family microarchitecture. All processors in the family have up to 12 cores (compared to 8 cores in its predecessor), which bring additional computing power to the table. They also have 50% additional cache (30 MB) and more memory bandwidth. With the 22-nm process technology, the Intel Xeon processor E5-2600 V2 product family has less idle power and is capable of delivering 25% more performance[2] while consuming less power compared to the earlier version.

Table 1 shows a comparison of the Intel Xeon processor E5-2600 V2 product family features compared to its predecessor, the Intel Xeon processor E5-2600.

Table 1. Comparison of the Intel® Xeon® processor E5–2600 product family to the Intel® Xeon® processor E5–2600 V2 product family

Feature	Intel® Xeon® E5-2600 (Sandy Bridge-EP)	Intel® Xeon® E5-2600 v2 (Ivy Bridge-EP)
QPI Speed (GT/s)	8.0, 7.2 and 6.4	
Addressability	46 bits physical, 48 bits virtual	
Cores	Up to 8	**Up to 12**
Threads Per Socket	Up to 16 threads	**Up to 24 threads**
Last-level Cache (LLC)	Up to 20 MB	**Up to 30 MB**
Intel® Turbo Boost Technology[1]	Yes	
Memory Population	4 channels of up to 3 RDIMMs, 3 LRDIMMs or 2 UDIMMs	
Max Memory Speed	Up to 1600	**Up to 1866**
Memory RAS	ECC, Patrol Scrubbing, Demand Scrubbing, Sparring, Mirroring, Lockstep Mode, x4/x8 SDDC	
PCIe* Lanes / Controllers/Speed (GT/s)	40 / 10 (PCIe* 3.0 at 8 GT/s)	
TDP (W)	150 (Workstation only), 130, 115, 95, 80, 70, 60, 50	
Idle Power Targets (W)	15W or higher, 12W for LV SKUs	**10.5W or higher, 7.5W for LV SKUs**

[1] Requires a system with Intel® Turbo Boost Technology. Intel Turbo Boost Technology and Intel Turbo Boost Technology 2.0 are only available on select Intel® processors. Consult your PC manufacturer. Performance varies depending on hardware, software, and system configuration. For more information, visit http://www.intel.com/go/turbo

The rest of this paper discusses some of the main enhancements in this product family.

3.1 Intel® Secure Key (DRNG)

Intel Secure Key (Digital Random Number Generator: DRNG) is a hardware approach to high-quality and high-performance entropy and random number generation. The entropy source is thermal noise within the silicon.

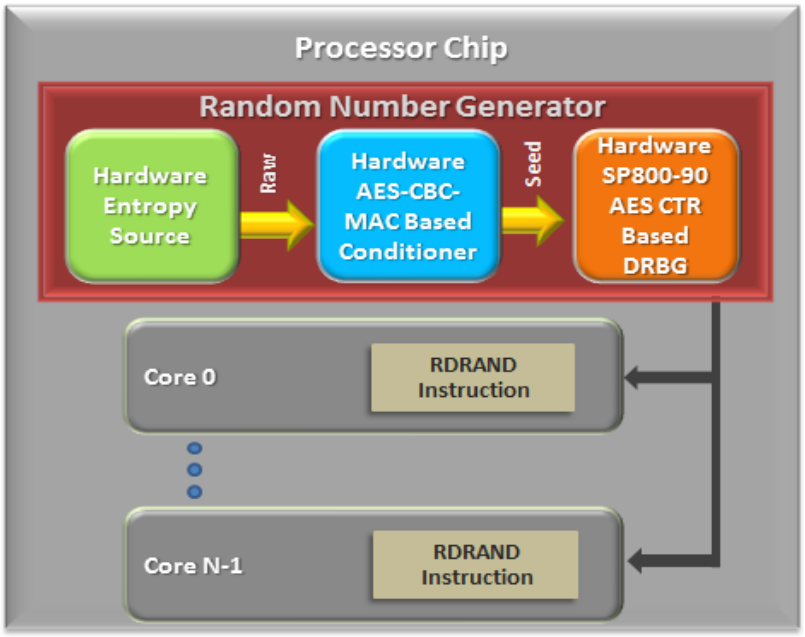

Figure 2. Digital Random Number Generator using RDRAND instruction

Figure 2 shows a block diagram of the Digital Random Number Generator. The entropy source outputs a random stream of bits at the rate of 3 GHz that is sent to the conditioner for further processing. The conditioner takes pairs of 256-bit raw entropy samples generated by the entropy source and reduces them to a single 256-bit conditioned entropy sample. This is passed to a deterministic random bit generator (DRBG) that spreads the sample into a large set of random values, thus increasing the amount of random numbers available by the module. DRNG is compliant with ANSI X9.82, NIST, and SP800-90 and certifiable to FIPS-140-2.

Since DRNG is implemented in hardware as a part of the processor chip, both the entropy source and DRBG execute at processor clock speeds. There is no system I/O required to obtain entropy samples and no off-chip bus latencies to slow entropy transfer. DRNG is scalable enough to support heavy server application workloads and multiple VMs.

DRNG can be accessed through a new instruction named RDRAND. RDRAND takes the random value generated by DRNG and stores it in a 16-bit or 32-bit destination register (size of the destination register determines size of the random value). RDRAND can be emulated via CPUID.1.ECX[30] and is available at all privilege levels and operating modes. Performance of RDRAND instruction is dependent on the bus infrastructure; it varies between processor generations and families.

Software developers can use the RDRAND instruction either through cryptographic libraries (OpenSSL* 1.0.1) or through direct application use (assembly functions). Intel® Compiler (starting with version 12.1), Microsoft Visual Studio* 2012, and GCC* 4.6 support the RDRAND instruction.

Microsoft Windows* 8 uses the DRNG as an entropy source to improve the quality of output from its cryptographically secure random number generator. Linux* distributions based on the 3.2 kernel use DRNG inside the kernel for random timings. Linux distributions based on the 3.3 kernel use it to improve the *quality* of random numbers coming from /dev/random and /dev/urandom, but not the *quantity*. That being said, Red Hat Fedora* Core 18 ships with the rngd daemon enabled by default, which will use DRNG to increase *both* the quality and quantity of random numbers in /dev/random and /dev/urandom.

For more details on DRNG and RDRAND instruction, refer to the Intel DRNG Software Implementation Guide.

3.2 Intel® OS Guard (SMEP)

Intel OS Guard (Supervisor Mode Execution Protection: SMEP) prevents execution out of untrusted application memory while operating at a more privileged level. By doing this, Intel OS Guard helps prevent Escalation of Privilege (EoP) security attacks. Intel OS Guard is available in both 32-bit and 64-bit operating modes and can be enumerated via CPUID.7.0.EBX[7].

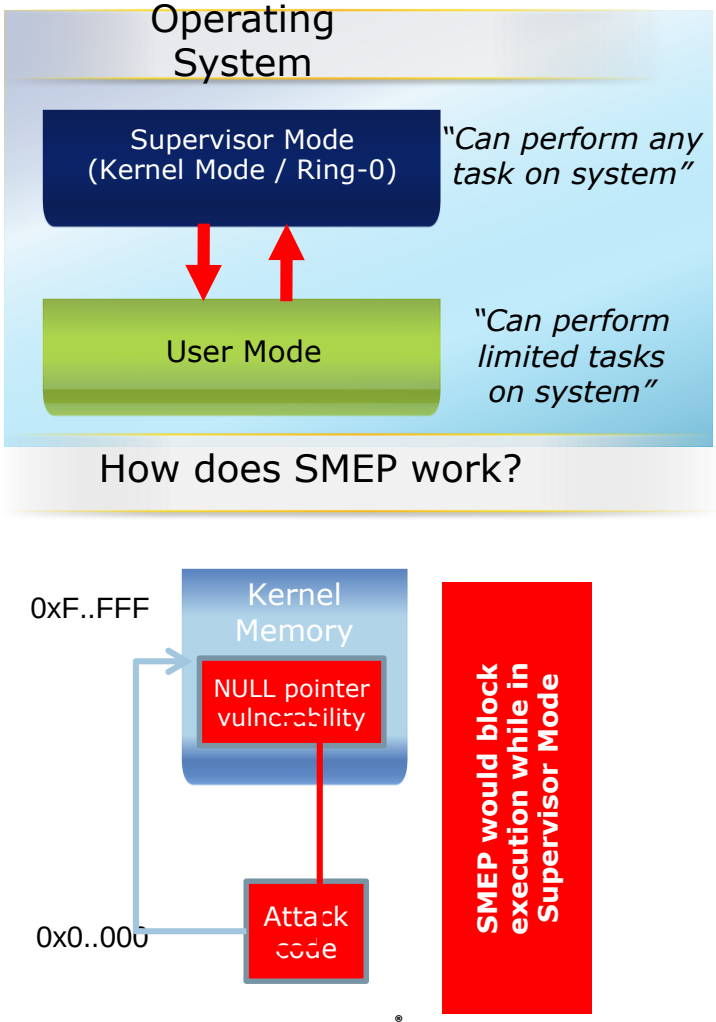

Figure 3. Pictorial description of Intel® OS Guard operation

Support for Intel OS Guard needs to be in the operating system (OS) or Virtual Machine Monitor (VMM) you are using. Please contact your OS or VMM providers to determine which versions include this support. No changes are required in the BIOS or application level to use this feature.

3.3 Intel® Advanced Vector Extensions (Intel® AVX): Float 16 Format Conversion

The "Sandy Bridge" microarchitecture introduced Intel AVX, a new-256 bit instruction set extension to Intel® SSE designed for applications that are floating-point (FP) intensive. The "Ivy Bridge" microarchitecture enhances this with the addition of float 16 format conversion instructions.

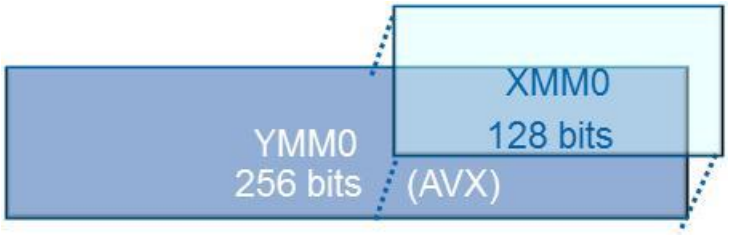

Figure 4. Intel® Advanced Vector Extensions Instruction Format

Intel Xeon processor E5-2600 V2 product family supports half-precision (16-bit) floating-point data types. Half-precision floating-point data types provide 2x more compact data representation than single-precision (32-bit) floating-point data format, but sacrifice data range and accuracy. In particular, half-floats may provide better performance than 32-bit floats when the 32-bit float data does not fit into the L1 cache. This format is widely used in graphics and imaging applications to reduce dataset size and memory bandwidth consumption.

Because the half-precision floating-point format is a storage format, the only operation performed on half-floats is conversion to and from 32-bit floats. The Intel Xeon processor E5-2600 V2 product family introduces two half-float conversion instructions: **vcvtps2ph** for converting from 32-bit float to half-float (4x speedup compared to alternative Intel AVX code implementation), and **vcvtph2ps** for converting from half-float to 32-bit float (2.5x speedup compared to alternative Intel AVX implementation). A developer can utilize these instructions without writing assembly by using the corresponding intrinsics instructions: _mm256_cvtps_ph for converting from 32-bit float to half-float, and _mm256_cvtph_ps for converting from half-float to 32-bit float (_mm_cvtps_ph and _mm_cvtph_ps for 128-bit vectors).

The compilers that support these instructions include Intel Compiler (starting with version 12.1), Visual Studio 2012, and GCC 4.6. To direct the Intel Compiler to produce the conversion instructions for execution on Intel Xeon processor E5-2600 V2 product family (or later), a developer can either compile the entire application with the –xCORE-AVX-I flag (/QxCORE-AVX-I on Windows), or use the Intel®-specific optimization pragma with target_arch=CORE-AVX-I for the individual function(s).

For more details on half precision floating point instructions, refer to:
http://software.intel.com/en-us/articles/performance-benefits-of-half-precision-floats

3.4 Advanced Programmable Interrupt Controller (APIC) Virtualization (APICv)

A significant amount of performance overhead in machine virtualization is due to Virtual Machine (VM) exits. Every VM exit can cause a penalty of approximately 2,000 – 7,000 CPU cycles (see Figure 5), and a significant portion of these exits are for APIC and interrupt virtualization. Whenever a guest operating system tries to read an APIC register, the VM has to exit and the Virtual Machine Monitor (VMM) has to fetch and decode the instruction.

The Intel Xeon processor E5-2600 V2 product family introduces support for APIC virtualization (APICv); in this context, the guest OS can read most APIC registers without requiring VM exits. Hardware and microcode emulate (virtualize) the APIC controller, thus saving thousands of CPU cycles and improving VM performance.

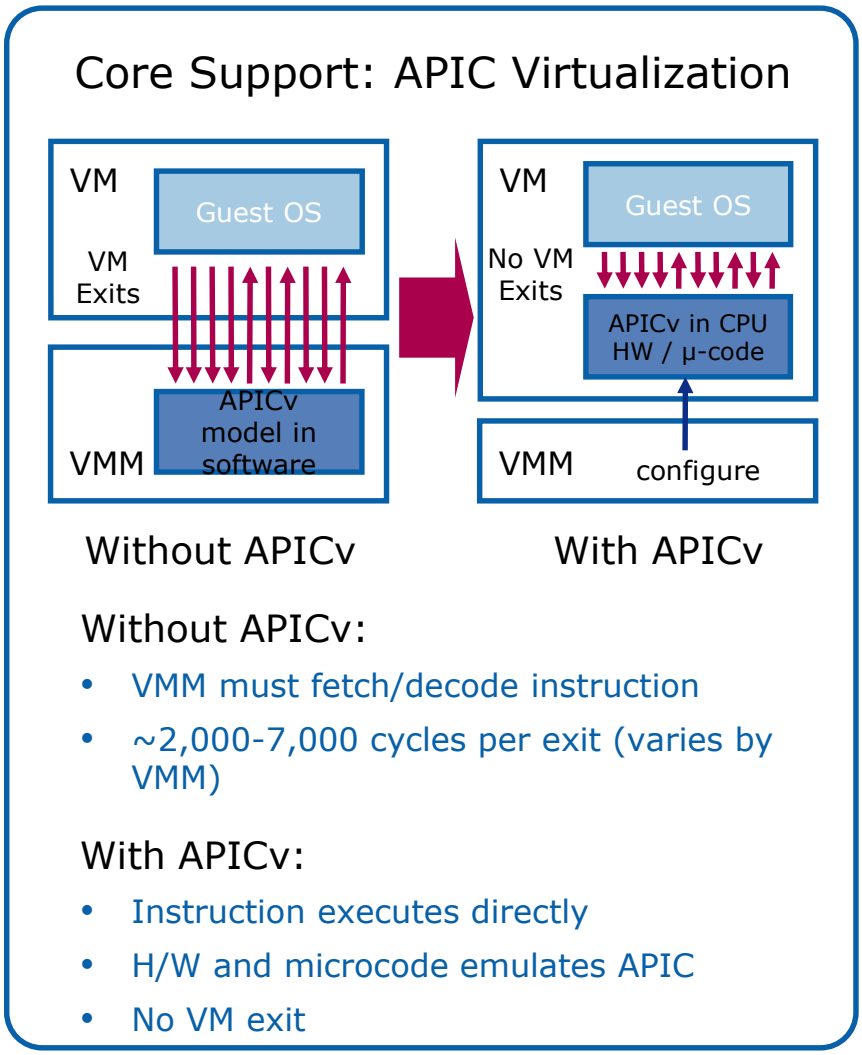

Figure 5. APIC Virtualization

This feature must be enabled at the VMM layer: please contact your VMM supplier for their roadmap on APICv support. No application-level changes are required to take advantage of this feature.

3.5 PCI Express Enhancements

The Intel Xeon processor E5-2600 V2 product family supports PCIe atomic operations (as a completer). Today, message-based transactions are used for PCIe devices, and these use interrupts that can experience long latency, unlike CPU updates to main memory that use atomic transactions. An Atomic Operation (AtomicOp) is a single PCIe transaction that targets a location in memory space, reads the location's value, potentially writes a new value back to the location, and returns the original value. This "read-modify-write" sequence to the location is performed atomically. This is a new operation added per PCIe Specification 3.0. FetchAdd, Swap, and CAS (Compare and Swap) are the new atomic transactions.

The benefits of atomic operations include:

- Lower overhead for synchronization
- Lock-free statistics (e.g. counter updates)
- Performance enhancement for device drivers

The Intel Xeon processor E5-2600 V2 product family also supports X16 non transparent bridge. All these contribute to better I/O performance.

These PCIe features are inherently transparent and require no application changes.

For more details on these PCIe features, refer to:

- PCI-SIG ENGINEERING CHANGE NOTICE
- **PCIe**® Protocol Updates - PCI-SIG

4. Conclusion

In summary, the Intel Xeon processor E5-2600 V2 product family combined with the Romley platform provides many new and improved features that could significantly change your performance and power experience on enterprise platforms. Developers can make use of most of these new features without making any changes to their applications.

About the Author

Sree Syamalakumari is a software engineer in the Software & Service Group at Intel Corporation. Sree holds a Master's degree in Computer Engineering from Wright State University, Dayton, Ohio.

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

[1] Baseline Configuration and Score on SPECVirt_sc2013* benchmark: Platform with two Intel® Xeon® Processor E5-2690, 256GB memory, RHEL 6.4(KVM). Baseline source as of July 2013. Score: 624.9 @ 37 VMs. New Configuration: IBM System x3650 M4* platform with two Intel® Xeon® Processor E5-2697 v2, 512GB memory, RHEL 6.4(KVM). Source as of Sept. 2013. Score: 947.9 @ 57 VMs. For more information go to http://www.intel.com/performance

[2] Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance