
The New AMD Opteron™ Processor Core Technology

INTRODUCTION

When AMD introduced the first AMD Opteron™ processor in 2003, it featured an innovative design that became the standard for the x86 server industry, with an integrated memory controller, scalable core architecture and innovative power efficiency features. For the following 7 years, AMD has improved on this design, increasing from a single core, all the way up to 12 cores. Multiple dies now occupy the high-end processors and the performance has scaled significantly higher than what was originally introduced in 2003 – yet it has essentially remained in the same power and thermal envelope, with few socket/platform changes to help maintain the long-term stability of the servers.

BENEFITS FOR CUSTOMERS

In 2011, AMD introduced the next generation of AMD Opteron processors, based on a new core architecture that is designed specifically to enhance the scalability of enterprise applications. With up to 16-cores, the new AMD Opteron processors deliver up to 35% greater throughput¹ than existing 12-core AMD Opteron processors, while maintaining the same power and thermal ranges. Virtualization is enhanced with new features that deliver greater scalability and manageability. Database applications enjoy more simultaneous threads than ever before with up to 33% more cores², and HPC customers will find they can reach a high density of cores with far fewer nodes than ever before, allowing for the maximization of data center space. Finally, cloud customers will find that the high core count and massive memory scalability are an ideal complement for highly scalable public or private cloud deployments.

THE PROCESS TECHNOLOGY

This new processor features the latest technology from AMD's foundry partner, GLOBALFOUNDRIES. Based on a 32nm high-K metal gate design, the processor features a die size for an 8-core die that is slightly smaller than AMD's current 6-core server die. This means 33% more cores in essentially the same silicon budget as before. GLOBALFOUNDRIES is utilizing a gate-first approach in the fabrication of these products. The gate-first approach requires less-restrictive layout design rules, provides for a smaller die size, and offers other manufacturing benefits, while delivering on the power and performance needs of customers.

THE NEW AMD OPTERON PROCESSOR DESIGN PHILOSOPHY

The design goal for the new architecture was simple – drive greater scalability and higher density of cores by sharing some components of the processor die to maximize efficiency and keep other parts discrete to avoid bottlenecks. Many of today's modern processor designs include redundant components because of their multi-core design, and in many cases these redundancies merely consume more die space, which can increase cost and increase power consumption – without necessarily adding incremental value to the processor.



By moving to a new design, where two integer cores share a common front end (fetch and decode) along with a large shared L2 cache and an FPU complex, AMD has created a modular processor design that allows for a “building block” approach. These modules allow AMD to have a flexible design that serves as the blueprint for the new AMD Opteron 6200 and 4200 Series processors.

HEAVY THREADING

Each module features two integer cores capable of executing a single thread each. Thus, each module can handle two threads per cycle across each integer core. Previous AMD designs could also handle one thread per core and had only 3 pipelines, which shared ALU and AGU functions. Each integer core of the new AMD Opteron processor has four pipelines, 2 for ALU and 2 for AGU, enabling greater throughput for processing. The platform, the operating system and the applications do not see modules (these are mainly designations for the designers), but instead see only a pool of integer cores that can be utilized for processing.

THE FLEX FP

Floating point operations are about 10% of the processing in a typical server workload, with the other 90% being integer. The challenge with advancements in floating point technology is that they require increasingly greater amounts of silicon and power to solve problems that most workloads are not taking advantage of today. The Flex FP, a flexible floating point unit or (FPU), is designed to help address this inequity, by bringing next generation FPU capabilities in the way of support for 256-bit AVX instructions, yet still help minimize the die space and power impact for those that do not need AVX. As most applications still take advantage of 128 bit instructions, the Flex FP is designed to function as either a pair of 128 bit FPUs, or as a single 256 bit FPU.

In addition, the Flex FP features FMAC units that can handle an FMA (Fused Multiply Accumulate) instruction as well as the standard FMUL (multiply) and FADD (add). An FMA operation is far more powerful because it allows a calculation like $A = B \times C + D$ in one cycle. In a standard FPU, this would require 2 cycles, one for the multiply and one for the add instruction. In addition, the FMAC can handle both an FMUL and FADD, whereas the traditional designs have dedicated paths for those operations. Having to handle multiple FADDs or FMULs can be optimized by allowing the work to be split across both FMACs instead of having all of the operations stack up behind one, waiting for execution. With the new core architecture, AMD is utilizing the more powerful FMA4 operation which has four operands.

AMD Turbo CORE Technology

A new feature added to the new AMD Opteron™ processor architecture is AMD Turbo CORE technology, which allows processors to run above their base clock frequency provided there is additional power headroom available. When processors are rated for clock speed, it is typically done based on a worst case scenario for workloads. This results in a speed rating that is very conservative for the vast number of workloads that will be run, potentially leading to a lost opportunity in clock speed for less strenuous workloads. This is one reason why processors generally consume far less power than their rated TDP. AMD Turbo CORE technology captures this additional power headroom and turns it into higher clock frequency, allowing the processor to run 300-500MHz higher than the rated base frequency with all cores active and utilized. In environments where only half the cores are being utilized the frequency increase can be 1 GHz or more³.

AMD Turbo CORE technology is designed to monitor the workload as well and will not automatically drive to the highest clock frequency unless the workload demands it. So, just like AMD PowerNow!™ technology that can modulate the core frequency — driving it down in periods of low utilization – AMD Turbo CORE technology does the same, but in the opposite direction with the goal of squeezing extra clock speed out of the processor when loads are reaching their peak. AMD Turbo CORE technology works in conjunction with AMD CoolSpeed™ technology to help ensure that even when boosting frequencies, processors stay within safe operating ranges. Obviously, when increasing clock frequency, power consumption can go up, so customers must carefully consider their need for raw performance in conjunction with their power needs.

IMPROVED MEMORY CONTROLLER

When designing this new processor, AMD took the opportunity to redesign the memory controller, helping to increase the throughput by up to 35% over existing AMD Opteron processors¹. Having an integrated memory controller for over seven years has helped AMD learn much about the optimization of the memory controller and how to drive better efficiency for enterprise applications. Part of that throughput gain comes from optimizations to the memory controller – new algorithms and new ways to address reading and writing data that speed up access to information. In addition, the new AMD Opteron processors now support DDR3 1600MHz memory. Between the enhancements in memory controller circuitry and the faster memory we anticipate significantly higher throughput for memory-intensive environments like virtualization, HPC, database and business applications.

Two new memory standards that are now supported with the new AMD Opteron processors include LR DIMM (load reduced DIMM), which allow for greater capacities of memory to be installed in the server as well as the emerging 1.25V low voltage DIMMs. Over yesterday's LV DIMMs of 1.35V, the new 1.25V will lead to even more power efficiency for customers.

CACHE STRUCTURE

Each module includes two levels of cache, an L1 cache that is focused on execution and an L2 cache that is focused on being the “working area” for data being processed. There are two components in the L1 cache, instruction and data. The L1 instruction cache is 64K and it is shared by the two integer cores as there is a high probability that instructions needed by one core also will be needed by the other, increasing the efficiency by allowing that instruction to only have to reside in one place in the cache. Each core also has its own L1 data cache that is designed to hold the data tied directly to processing those instructions for the cycle. Beyond the L1 cache, each module has its own 2 MB L2 cache that is shared between 2 cores. If only one core is active and storing data, it can have access to all 2MB of the L2 cache. Outside of the modules, at the die level, an 8MB L3 cache is employed. For AMD Opteron 6200 Series processors, there are two die present inside of the processor, which means a total of 16MB of L3 cache. AMD Opteron 4200 Series processors have one die for a total of 8MB of L3 cache.

DEDICATED SCHEDULERS

Inside each module are three dedicated schedulers, one for each integer core and one to feed the Flex FP. The integer schedulers are 40-entry and the Flex FP scheduler is 60-entry. By having a dedicated scheduler for each integer core, the new AMD Opteron processor core architecture helps ensure that the 4 integer pipelines are being kept continually filled with instructions for maximum efficiency. Each integer core has control over its own scheduling so that there is virtually no bottleneck between the two dedicated threads that are executing in the module simultaneously. The Flex FP scheduler is a single entity because in AVX mode it needs to be able to send a single stream of 256-bit AVX operations through the FP pipes as well as handle dual 128-bit executions. In 128-bit mode, the extra entries in the Flex FP scheduler help ensure that the two 128-bit FMACs are receiving a constant stream of math instructions to execute. The key for great FP throughput is having a deeper scheduler, allowing for more out-of-order execution – this helps keep the FP pipelines full. FP code is more latency sensitive than integer code, so having a dedicated scheduler can give the full scheduler depth for the FPU for greater amounts of re-ordering. The scheduler depth on the new AMD Opteron processor FP scheduler is deeper than previous generations of AMD Opteron processors, allowing for maximum performance potential.

VIRTUALIZATION FEATURES

As virtualization continues to grow as part of a data center strategy, integrated virtualization features are an essential part of any processor design. Introduced in 2006, AMD Virtualization™ (AMD-V™) technology is an integrated set of virtualization features that help to boost performance and efficiency. Bit-level enhancements like Flush by ASID, Larger ASID space, and VMCB clean bits, along with a host of other hypervisor-level enhancements have been made to beef up the AMD-V support. In addition, Virtual Cache Partitioning, also supported by the new AMD Opteron processor, allows a portion of the L3 cache to be dedicated to a single module (thus creating a self-contained “system in a system” for better manageability).

In a virtualized environment – where core density and memory addressability are key performance drivers – the new AMD Opteron processor architecture delivers the greatest competitive advantage with its superior core count and greater scalability. For those customers that utilize the “1 VM per core” methodology for deployment, AMD-based platforms can support 60-100% more VMs than a comparable Intel-based platform⁴. AMD's work with leading virtualization providers like VMware, Microsoft, Xen, Citrix, Parallels and the Linux KVM community help ensure a highly supportable virtualization solution.

HPC ENVIRONMENTS

Customers who use high performance compute clusters in their business will find that the new AMD Opteron processor – based platforms deliver superb performance on these technical computing workloads. With 16 cores per processor and the ability to run 64 single-precision operations (or 32 double-precision operations) per processor per cycle, the new architecture excels in FLOPs intensive applications such as computational chemistry and molecular dynamic codes. The high memory bandwidth⁵ combined with 384 GB/socket memory addressability lends favorably to memory intensive applications such as fluid dynamics and weather modeling. In addition, some of the new instruction sets mentioned in this paper will allow those that are tuning code specifically for their applications to really harness the power of the architecture.

POWER EFFICIENCY FEATURES

AMD's continued focus on power efficiency brings new power enhancements to the new architecture. The new AMD Opteron 6200 and 4200 Series processors are designed around using power as a resource. At the core level, the new C6 power state is designed to power down a complete module when it has been in idle for a pre-determined amount of time. TDP Capping can help reduce core power consumption by enabling IT managers with a definable TDP setting that can be set to a reduced level through BIOS or dynamically through our Advanced Platform Management Link (APML). The new AMD Opteron™ processors also support both low power (1.35v) and ultra low power (1.25v) DDR3 memory. Ultra low power memory can offer reduced memory subsystem power by up to 16% compared to traditional 1.5v DDR3 memory.

All of these power features are wrapped up in 32nm Silicon on Insulator (SOI) architecture that has been tuned for power efficiency. The result: even though the new design features 33% more cores², higher core frequencies, support for 20% higher memory clock speed and the inclusion of AMD Turbo CORE boosting technology, the processors still operate in the same TDP ranges as current AMD Opteron 4100 and 6100 Series processors.

NEW SOFTWARE INSTRUCTIONS

AMD supports a series of new instructions and operations with the new AMD Opteron processor core that are designed to help drive greater performance and compatibility for high-powered applications:

INSTRUCTION	FUNCTION/BENEFIT
SSSE3	For media handling like video encoding and transcoding.
SSE4.1	Covers a range of applications with several new packed data operations including various specialized data movement instructions; dot product, min/max, compare and rounding operations for numeric processing, and further video encoding support.
SSE4.2	POPCNT (Population Count) instruction is applicable to bioinformatics algorithms; CRC32 is for accelerating CRC-based integrity checking of network or disk data transfers; string instructions are applicable to any text-intensive applications such as XML and HTML parsing.
AES, PCLMULQDQ	Accelerates any application that uses AES encryption, key uses being secure network transactions (internet and LANs), disk encryption such as Microsoft's BitLocker, and database encryption.
AVX	Provides 256-bit floating point as well as a significant performance boost for vector floating point applications, and a lesser boost for multimedia apps and non-vector FP-intensive apps.
FMA4	A further boost for numeric applications, particularly HPC-type applications, providing up to a 2x increase beyond AVX on AMD processors.
XOP	Provides enhancements for pack data operations, applicable to a variety of numeric and multimedia applications, including DSP-type applications.

PLATFORM COMPATIBILITY

When AMD designed the AMD Opteron 4000 and 6000 Series platforms, future market requirements and platform stability and compatibility were key design decisions. AMD anticipates that most, if not all, servers will have only a BIOS update in order to take advantage of this new processor. As few customers actually ever update their processors once the systems have been deployed, the real benefit is not from the socket compatibility, but rather from the fact that current generation and next generation platforms can live side-by-side, being managed the same way with the same BIOS, drivers and interfaces. Migrating software and virtual machines should be simple processes, since platform-level components like chipsets and peripherals are the same.

SUMMARY

The new AMD Opteron processor offers significant performance improvements in floating point and integer scheduling, memory throughput and the inclusion of new instruction sets. Additionally, the innovative modular architecture helps pack in more cores and more memory, making it ideal for virtualized environments and the natural choice for HPC and cloud workloads. As well as an added benefit, they're designed to save costs through efficient design and power efficiency. Nearly any data center or workload will benefit from this high-performing, flexible and efficient processor.

- 1 Based on AMD internal engineering performance estimates comparing 12-core AMD Opteron 6100 Series processor with 16-core AMD Opteron 6200 Series processor. SVR-16
- 2 Comparison of 12-core AMD Opteron 6100 Series processors and 16-core AMD Opteron 6200 Series processors. SVR-5
- 3 AMD Opteron 6200 Series processors experience all core boost of up to 500 MHz (P2 base to P1 boost state) and up to 1.3 GHz max turbo boost (half or fewer cores boost from P2 to P0 boost state). SVR-27
- 4 Based on comparison of 16-core AMD Opteron™ 6200 Series processor with 8-core Intel Xeon 7500 Series processor and 10-core Intel Xeon E7 Series processor using 1 VM per core loading rule. SVR-61
- 5 AMD Opteron 6200 Series supports up to 1.5 TB memory capacity in a four processor configuration using LR DIMMs. AMD Opteron 6100 Series supports up to 1TB memory capacity in a four processor configuration using RDIMMs. SVR-64

