

EPYC: Designed for Effective Performance

By Linley Gwennap
Principal Analyst

June 2017



www.linleygroup.com

EPYC: Designed for Effective Performance

By Linley Gwennap, Principal Analyst, The Linley Group

Measuring server-processor performance using clock speed (GHz) or even the traditional SPEC_int test can be misleading. AMD's new EPYC processor is designed to deliver strong performance across a wide range of server applications, meeting the needs of modern data centers and enterprises. These design capabilities include advanced branch prediction, data prefetching, coherent interconnect, and integrated high-bandwidth DRAM and I/O interfaces. AMD sponsored the creation of this white paper, but the opinions and analysis are those of the author. Trademark names are used in an editorial fashion and are the property of their respective owners.

Although many PC users can settle for “good enough” performance, data-center operators are always seeking more. Web searches demand more performance as the Internet continues to expand. Newer applications such as voice recognition (for services such as Alexa and Siri) and analyzing big data also require tremendous performance. Neural networks are gaining in popularity for everything from image recognition to self-driving cars, but training these networks can tie up hundreds of servers for days at a time. Processor designers must meet these greater performance demands while staying within acceptable electrical-power ratings.

Server processors are often characterized by core count and clock speed (GHz), but these characteristics provide only a rough approximation of application performance. As important as speed is, the amount of work that a processor can accomplish with each tick of the clock, a parameter known as instructions per cycle (IPC), is equally important. The CPU's internal design, or microarchitecture, determines both the clock rate and the IPC, so an effective design is crucial for building a powerful CPU.

Even the most powerful engine will starve without fuel, so a high-performance processor must include the right plumbing to efficiently deliver data to the CPUs when they need it. To keep the CPUs fed, large on-chip caches must be combined with multiple high-bandwidth memory (DRAM) channels. High-speed I/O connects the processor with storage and other servers in the network. Different applications require different mixes of CPU, memory, and I/O throughput, so flexibility is needed to support a broad range of applications.

AMD's next-generation EPYC processor meets the performance demands of current and emerging data-center applications. It uses the company's new “Zen” microarchitecture, which improves IPC by 52% over AMD's previous CPU generation and can reach turbo speeds of more than 3.0GHz. The processor packs up to 32 CPU cores, more than any current Xeon processor.¹ To feed these cores, the processor features up to eight independent DRAM channels. This performance is backed by up to 128 lanes of high-speed PCI Express (PCIe), which can connect to Ethernet and other I/O adapters as needed.

EPYC delivers this strong performance while meeting other data-center needs. AMD is the only company other than Intel that delivers x86-compatible processors, meaning they can run existing server software without modification or recompilation. The EPYC products also support a wide selection of reliability, availability, and serviceability

(RAS) features, so a data center can maximize uptime of its servers. These processors provide industry-standard memory and I/O interfaces and fit into typical server thermal limits, so they can be used in standard enclosures and racks. This compatibility simplifies the upgrade to EPYC.

The Zen CPU Design

AMD invested several years in developing the Zen microarchitecture, which provides many advances over the previous “Bulldozer” CPU family. Zen is a powerful CPU capable of sustained execution of four x86 instructions per cycle. Like Intel’s designs, Zen can execute two threads at a time, keeping the CPU busy even when one thread stalls. Each CPU has its own floating-point unit and its own 512KB cache, reducing sharing from the previous generation. These characteristics are important in achieving high performance, but Zen also includes smaller, less common features that keep the CPU operating at maximum efficiency.

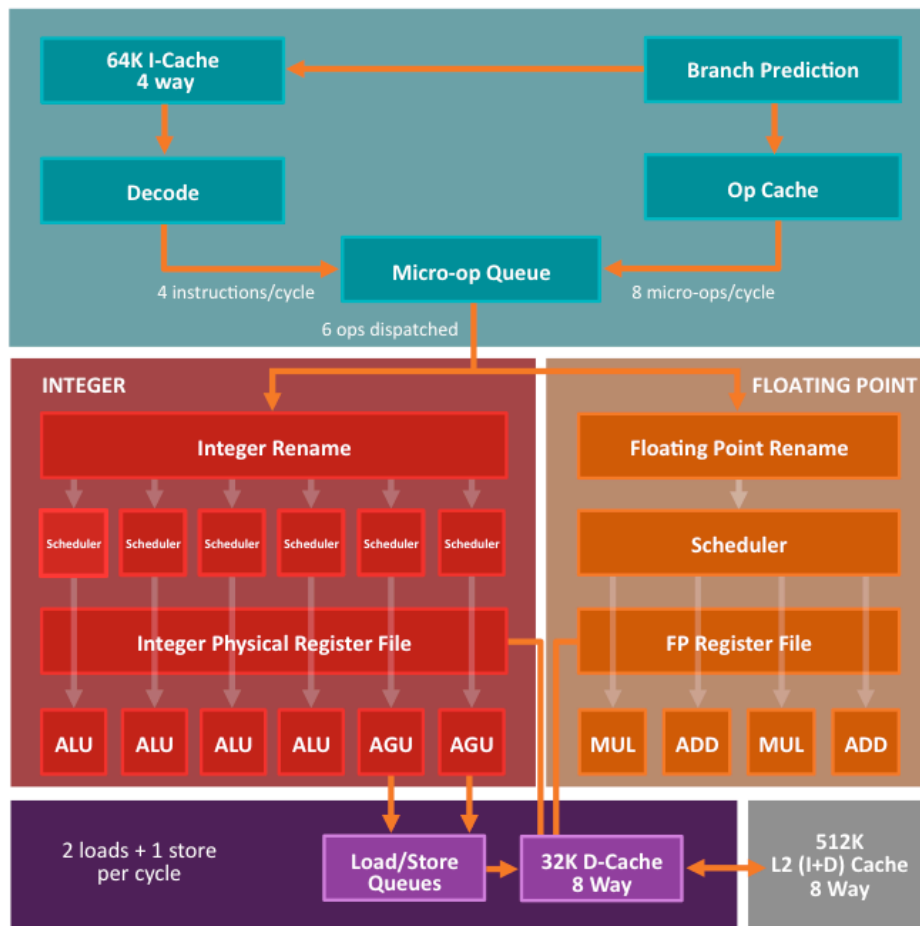


Figure 1. Zen CPU microarchitecture. The branch predictor and micro-op cache provide a significant performance boost. (Source: AMD)

The CPU fetches up to six instructions per cycle from the instruction cache, saving these instructions in a buffer for later decoding. This approach helps keep the four decoders busy even when the I-cache is unavailable for a few cycles. The decoded instructions

(called micro-ops) are queued before being dispatched to a set of 10 execution units. This wide array of execution units allows the CPU to execute almost any combination of instructions at one time, including four integer operations, four FP operations, and two load/store operations.

Zen is AMD's first CPU with an op cache. This small memory holds instructions that have already been decoded into micro-ops. Any time the CPU needs an instruction that is already in the micro-op cache, it avoids fetching and decoding that instruction, saving time (two clock cycles) and the energy (power) needed for the decoding logic. The cache, which holds 2,048 micro-ops, is particularly useful when the CPU repeatedly executes a loop of code, a frequent situation in many applications.

At other times, however, code branches to a new location, which not only bypasses the micro-op cache but can cause the entire CPU to stall for several cycles until the new instructions are loaded. To avoid this type of stall, most CPUs use branch prediction to prefetch the instructions it expects to need. This approach is only as good as the accuracy of the predictions, however. To improve accuracy, Zen uses an advanced prediction scheme based on perceptrons (simple neural networks). The new CPU also includes an indirect target array to handle the difficult case of dynamic indirect branches. Its larger branch-prediction tables also reduce the number of mispredictions. These changes combine to correctly predict most of the branches in typical code.

A similar problem occurs when the CPU can't immediately access data that it needs, causing it to stall for several cycles. Zen implements circuitry to prefetch data into the level-one (L1) and level-two (L2) caches. Since programs often access data in a linear sequence, these prefetchers typically load the subsequent memory addresses. Some programs, however, fetch data in a different pattern, such as every tenth word (for example, from a database that holds ten items for each person). The Zen CPU recognizes that the program has an unusual "stride" and prefetches data following that pattern. This approach ensures that the CPU has the data it needs as quickly as possible.

As a complete redesign, Zen offers many additional improvements over Bulldozer. For example, the L2 cache responds 50% faster, and the L3 cache responds 71% faster. This speedup means that, even if the data hasn't been prefetched, it gets back to the CPU much sooner. Because the L3 cache supports multiple cores, bandwidth is as important as speed. A massive 6x bandwidth improvement allows the L3 cache to deliver more data to more cores without overloading.

The EPYC Processor

AMD designed the EPYC processor to support up to 32 cores. All these CPUs operate coherently, meaning that software can easily share instructions and data among the cores using a single set of memory addresses. Each cluster of four CPUs shares an L3 cache that holds 8MB of data. A high-bandwidth coherent interconnect moves data among these clusters, so if one CPU needs data from another cluster, it can still share that data. In total, the processor includes 64MB of L3 cache, more than any current Xeon processor¹. The large L3 cache reduces the number of DRAM accesses, saving time and power.

EPYC: Designed for Effective Performance

Even with this cache, however, the processor must sometimes access DRAM. Some data-center applications access DRAM more often than others. For example, applications (e.g., big data) that don't fit into the L3 cache will frequently access DRAM. For these applications, EPYC supports eight memory channels, twice as many as any current Xeon processor and two more than we expect from Intel's next-generation Skylake-EP. Each channel supports the latest DDR4-2667 DIMMs with optional error correction (ECC). More channels deliver more memory bandwidth, helping to feed the 32 CPUs and keep them working efficiently. More channels also support more memory: up to 2TB (terabytes) per processor. This extra memory is good for in-memory databases and other applications with large working sets.

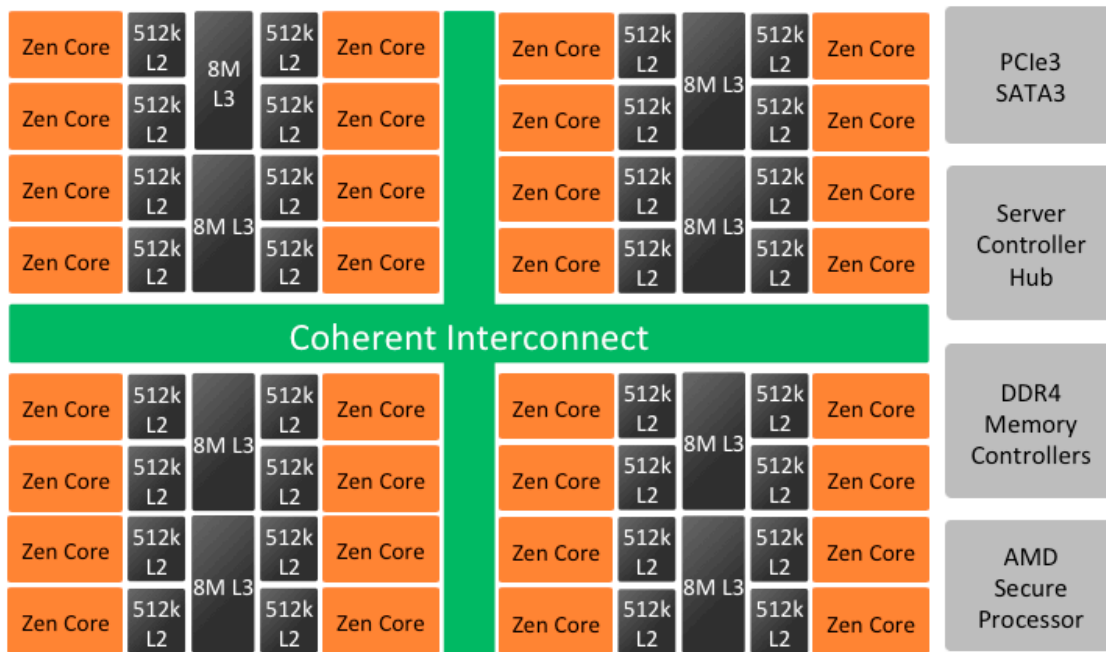


Figure 2. EPYC processor design. The coherent interconnect links up to 32 CPU cores in clusters of four, each with its own 8MB level-three cache. (Source: AMD)

Unlike most Xeon processors, EPYC is a complete system-on-a-chip (SoC) with both integrated DRAM controllers and integrated I/O controllers. This approach eliminates the cost and power of a south-bridge chipset, simplifying the motherboard. With 128 lanes of PCIe Gen3, a single EPYC processor has more than twice the I/O bandwidth of any current Xeon product¹. The greater bandwidth supports twice as many network adapters to connect to other servers in the rack and beyond. These ports can also connect to storage devices using the SATA or NVMe protocols.

Whereas many SoC processors are limited to single-socket designs, EPYC also supports the popular dual-socket configuration. In this configuration, two EPYC processors connect back-to-back using half of their PCIe lanes; these lanes employ a special protocol that maintains coherency across all of the CPUs in both sockets. This configuration still leaves a total of 128 PCIe lanes available for networking and storage. Together, the two sockets support up to 64 CPU cores and 16 DDR4 memory channels. Because of the coherent connection, the two processors appear to software as a single processor, avoiding any need to modify the code.

Processor Performance

Table 1 compares the performance and features of the top-of-the-line EPYC 7601 processor against Intel's flagship Xeon E5v4-2699A processor¹. It shows peak SPEC_rate2006 numbers measured on comparable dual-socket production systems from leading OEMs, as reported to SPEC.org by the system vendors. Both sets of numbers use the CPU vendor's optimizing compiler: ICC for Intel and Open64 for AMD. This comparison gives an advantage to Intel, which spends more resources on compiler tuning. A better comparison would use GCC, the compiler that most software developers rely on; AMD has measured GCC scores for EPYC, but Intel forbids its customers from publishing GCC scores, making such a comparison difficult.

	Xeon E5v4-2699A ²	EPYC 7601 ³
Cores/Threads	22/44	32/64
Base Clock Frequency	2.4GHz	2.2GHz
Boost Clock Frequency	3.6GHz	3.2GHz
Maximum Memory Bandwidth	77GB/s	170GB/s
Number of PCIe Lanes	40 Gen3	128 Gen3
Power (TDP)	145W + 7W*	180W
SPECint_rate2006 (dual socket)	1890 int	2360 int
SPECfp_rate2006 (dual socket)	1160 fp	1840 fp
SPECint_rate2006 per Watt	6.4 int/W	6.6 int/W
SPECfp_rate2006 per Watt	3.9 fp/W	5.1 fp/W

Table 1. EPYC performance versus Xeon's. The new AMD processor leaps ahead of Intel's Xeon processors on performance and power efficiency. *includes C612 south bridge. (Source: SPEC.org)

On the SPEC benchmark, the EPYC 7601 offers 25% greater integer performance and 59% greater floating-point performance than the Xeon E5v4-2999A, as Table 1 shows. (This advantage expands to 47% and 75% when measured using GCC, according to AMD.) Although the Xeon processor operates 9% faster, the EPYC chip's 45% more cores enable it to gain a performance advantage. On per-core performance, EPYC lags on integer workloads but pulls ahead on floating point. The FP advantage comes mainly from EPYC's greater memory bandwidth, which more than doubles that of the Xeon E5; many of the SPECfp tests have large data sets that don't fit into the on-chip cache. This advantage will help other applications that have large data sets.

At 180W, EPYC 7601 has a higher TDP rating as a result of having more cores, more memory controllers, and PCIe ports. But the EPYC chip integrates the south bridge, while Xeon E5 processors require an external south bridge such as the C612, which has a TDP rating of 7W. The combined TDP of the Xeon E5-2699 and the C612 is 152W, which is still smaller than that of the AMD processor. As the table shows, however, EPYC still comes out ahead in performance per watt, with 3% better integer efficiency and 31% better FP efficiency. These scores indicate that EPYC could significantly reduce electrical costs in large data centers.

In addition, the EPYC 7601 provides 122% more memory bandwidth and 60% more I/O than the Xeon E5-2699 processor. The SPEC CPU benchmark does not measure I/O speed, but applications that need the extra I/O capability will see a bigger advantage than shown here. These applications would require adding the cost and power of an external PCI switch to the Intel design, further reducing its efficiency. EPYC's extra I/O ports can also be configured for NVMe or SATA as needed.

Conclusion

The EPYC processor delivers high performance across a wide range of data-center applications. For scale-up applications, the Zen CPU is one of the most powerful designs available, capable of sustaining execution of four x86 instructions per cycle at speeds in excess of 3.0GHz. To maintain operation at this peak rate, the CPU uses features such as the micro-op cache, advanced branch prediction, and prefetching. When a stall does occur, dual threading allows the CPU to quickly switch to a second thread to fill the gap. The prefetcher works not just on streaming data but on variable strides, allowing it to accelerate many different data structures.

The EPYC design provides more memory bandwidth and more I/O bandwidth than any competing x86 processor¹. This bandwidth helps meet the data requirements of scale-up applications running on 32 cores. Some scale-out applications, however, have high memory or I/O demands but don't need as much CPU performance. EPYC allows system designers to combine lower core counts (e.g. 16 or 24) with the same high-bandwidth memory and I/O to effectively meet the needs of applications such as storage servers and in-memory databases.

AMD processors using the Bulldozer family can efficiently serve only certain data-center applications. With the new EPYC family, AMD can now address the broad data-center market. EPYC is fully compatible with existing server software and infrastructure while providing large gains in CPU, memory, and I/O performance across a wide range of server applications. It offers up to 25% better performance than Intel's flagship Xeon E5 processor¹ on the popular SPECint benchmark, with greater advantages on floating-point code and when using GCC. EPYC's power efficiency is better as well. Server-processor customers take note: AMD is back in the game.

Linley Gwennap is principal analyst at The Linley Group and editor-in-chief of Microprocessor Report. The Linley Group offers the most comprehensive analysis of microprocessor and SoC design. We analyze not only the business strategy but also the internal technology. Our in-depth reports also cover topics including server processors, embedded processors, IoT processors, and processor IP cores. For more information, see our web site at www.linleygroup.com.

¹Does not include Skylake-EP, which is not publicly announced at the time of publication.

²Huawei RH2288H v3, RHE Linux 7.2, ICC v16.0.0.101, 2x Xeon E5-2699 v4, 512GB (16 x 32GB 2Rx4 PC4-2400T-R), 1x 1000GB SATA, 7200 RPM

³Supermicro AS-1123US-TR4, Ubuntu 16.04, Open64 v4.5.2.1, 2x EPYC 7601, 512GB (16 x 32GB 2Rx4 PC4-2666), 1 x 500GB SSD