

M I C R O P R O C E S S O R

www.MPRonline.com

◆ THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE ◆

OPTERON 4100 IS COOL FOR CLOUDS

New AMD Server Processors Reduce Power in Data Centers

By Linley Gwennap and Tom R. Halfhill {8/9/10-01}

AMD has its head in the clouds. Its new Opteron 4100 server processors are intended for cloud-computing data centers that buy servers by the truckload. With prices starting at \$99 per chip and typical power as low as 32W, Opteron 4100 processors are challenging Intel's lowest-power Xeons in servers having one or two sockets. Although they can't match Xeon's most power-efficient models, they offer a less expensive alternative while still going easy on the electricity.

These prices and power levels may seem low for server processors, but at AMD, \$99 buys a quad-core chip running at 2.2GHz, and 32W represents 100% utilization for a server processor with six CPUs. In all, AMD has introduced nine new Opteron 4100 processors (available now), which are summarized in Table 1. Boards with one or two sockets (1P or 2P) represent about 95% of the server market, and that's where the 4100 series is aimed. It sits below the Opteron 6000 series, which targets higher-end servers with 2P and 4P configurations.

Cloud computing vaguely describes user applications running primarily on Internet servers, usually with remote file storage. (Google Docs is a prime example.) Essentially, it's client-server computing of the 1980s updated for the web. Cloud applications are often a free service supported by advertising or marketing gimmicks, so providers try to spend as little as possible on infrastructure and electric bills without driving away users with poor performance (see [MPR 11/24/08-02](#), "Azure Skies, Beautiful Clouds").

Direct and indirect costs related to electricity consumption can contribute 40% of the total cost of ownership (TCO) of a modern server (see [MPR 6/21/10-01](#), "New Processors Target Data Centers"). Cloud workloads tend to be more bursty than traditional server workloads, so idle power can be as important as active power. Consequently, cloud-computing providers need lots of cheap, power-efficient hardware.

They also need scalability, in case their services become popular. AMD is addressing those needs not only with

	4122	4130	4180	4184	4170 HE	4174 HE	4176 HE	4162 EE	4164 EE
# of CPUs	4 CPUs	4 CPUs	6 CPUs	6 CPUs	6 CPUs	6 CPUs	6 CPUs	6 CPUs	6 CPUs
CPU Speed	2.2GHz	2.6GHz	2.6GHz	2.8GHz	2.1GHz	2.3GHz	2.4GHz	1.7GHz	1.8GHz
L2 Cache	2MB	2MB	2MB	2MB	3MB	3MB	3MB	3MB	3MB
L3 Cache	6MB	6MB	6MB	6MB	6MB	6MB	6MB	6MB	6MB
North Bridge	2.2GHz	2.2GHz	2.2GHz	2.2GHz	2.2GHz	2.2GHz	2.2GHz	1.8GHz	1.8GHz
Platform	San Marino				Adelaide				
HT Interfaces	2xHyperTransport 3.1 (3x12.8GB/s)†				2xHyperTransport 1.0 (3x3.2GB/s)†				
C1E Sleep?	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Power (ACP)	75W	75W	75W	75W	50W	50W	50W	32W	32W
DP SPECint*	104‡	118‡	172‡	181	152‡	160‡	165	127‡	132
Price (1,000s)	\$99	\$125	\$188	\$316	\$174	\$255	\$377	\$316	\$698

Table 1. Key parameters for AMD's new 4100-series server processors. All are based on the 45nm "Lisbon" die with four or six active CPU cores. All fit AMD's 1,207-pin Socket C32 and are now available in volume. *Best published dual-processor SPECint2006_rate baseline result; †third link available for routing preference only. (Source: vendors and www.spec.org, except ‡The Linley Group estimate)

cheap, low-power processors, but also with system-logic chipsets that preserve hardware and software compatibility across the new Opteron 4000 and existing Opteron 6000 families. The company offers three SR5600-series system chipsets built on the same die, all with the same I/O interfaces, pinouts, and device drivers.

Opteron 4000 and Opteron 6000 designs will be upgradeable. Next year, AMD plans to improve the performance of both Opteron families by introducing a new CPU code-named Bulldozer. The new chips will be fabricated in a next-generation 32nm silicon-on-insulator (SOI) process that AMD is developing with GlobalFoundries. Single-die Bulldozer chips with six or eight CPUs (code-named Valencia) will drop into existing Opteron 4000 sockets. Twin-die Bulldozer chips with 12 or 16 CPUs (code-named Interlagos) will drop into existing Opteron 6000 sockets. Today's Opterons are manufactured in 45nm SOI, so the move to 32nm should reduce power consumption and cut AMD's manufacturing costs.

New Products Revisit Lisbon

All Opteron 4100 processors use the same six-CPU "Lisbon" die that appears in the Opteron 6000 products, known as Magny-Cours, announced in March. The 4100 series has a traditional single-die package; the 6000 series fits two die into one package, thus requiring a larger socket. Therefore, these processors are socket compatible within their respective series but not across both series. Using two die, Magny-Cours scales to 12 CPUs in a single package, as Figure 1 shows.

Lisbon is nearly identical to the "Istanbul" die in the Opteron 2400-series server processors introduced last year, but it delivers 6–7% better performance on SPECint at the same clock speed. This increase is primarily due to a shift from DDR2-800 to DDR3-1333 SDRAM, although Lisbon also includes minor CPU improvements and faster HyperTransport (HT) links. To reduce power, Lisbon has a new low-power sleep state, called C1E. It not only puts all CPUs to sleep, but it also reduces power to the memory controller

and the HT links, minimizing idle power consumption.

Lisbon differs in at least two ways from "Thuban," the newer die in AMD's Phenom II X6 desktop processors. Thuban introduced Turbo Mode, AMD's answer to Intel's Turbo Boost. To deliver a brief performance surge, one to three of Thuban's CPU cores can automatically jump to a higher clock speed (up to 500MHz above their normal frequency) while leaving the other CPUs idle, keeping the chip below its total thermal limit. This feature, however, is less important for server processors like the Opteron 4100 series that strive to reduce power consumption.

The other difference between Lisbon and Thuban is an improved dielectric in the latter die's metal layers, which reduces current leakage. Although GlobalFoundries manufactures both Lisbon and Thuban in similar 45nm SOI processes, the new dielectric hasn't yet spread through AMD's entire product line.

To maximize yield, AMD sells as quad-core processors those Lisbon chips in which one or two CPUs fail to meet specifications. Each CPU has a pair of 64KB L1 caches and a 512KB private L2 cache, so the quad-core chips have only 2MB of L2 cache instead of 3MB. All CPUs share a 6MB L3 cache, which is unaffected when some CPUs are disabled. As a result, each CPU in a quad-core processor can claim a larger share of the L3 cache.

Only two of the new Opteron 4100-series processors are quad-core chips: the 4122 and 4130, which are clocked at 2.2GHz and 2.6GHz, respectively. The only other distinguishing feature of these quad-core chips is that they cannot enter the C1E sleep state. Without C1E, the quad-core processors can't reach their lowest idle power. Otherwise, the 4122 and 4130 are identical to their six-CPU brethren, and they are the lowest-priced products in this series, having list prices of \$99 and \$125, respectively.

System designers can make boards that work with all 4100-series processors because these chips fit the same 1,207-pin Socket C32. This socket is much like AMD's older Socket F, which has the same pin count and similar pin assignments. The main difference is that Socket C32 supports dual-channel 1.5V DDR3 memory and 1.35V LV-DDR3 memory instead of DDR2 SDRAM. Each DDR3 channel is 64 bits wide, has ECC, and supports DRAM speeds up to 1.33GHz, so peak bandwidth is 21.3GB/s. Socket-F Opteron processors peaked at 12.8GB/s.

Platforms Make Power Tradeoffs

Opteron 4100-series chips are available in three power ranges: 75W, 50W, and 32W. To indicate their lower power level, the 50W products have an HE suffix, and 32W products have an EE suffix. The standard (no suffix) Opteron 4100 processors are typically used in AMD's general-purpose

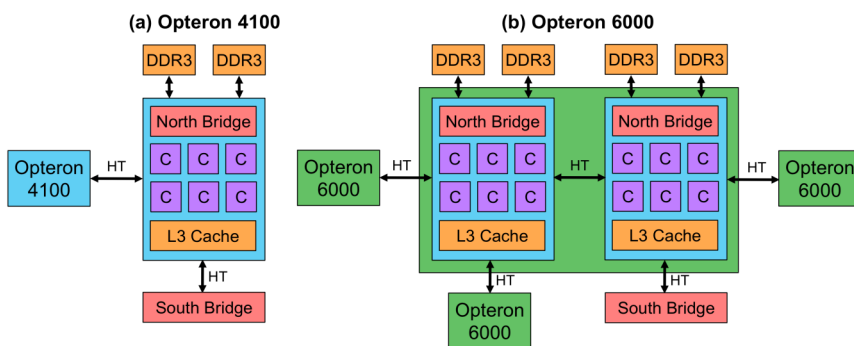


Figure 1. AMD Opteron 4100 and Opteron 6000. (a) The Opteron 4100 uses a single Lisbon die to support two sockets with up to six CPUs in each chip. (b) The Opteron 6000 combines two Lisbons in a single package. This product supports up to four sockets with up to 12 CPUs per socket.

“San Marino” platform, whereas the HE and EE processors usually appear in the low-power “Adelaide” platform. Either platform can have a 1P or 2P configuration.

Whereas the standard processors range from 2.2GHz to 2.8GHz, the HE and EE versions operate at 1.7–2.4GHz to save power. The EE processors further reduce power by underclocking the internal north-bridge logic at 1.8GHz; in the other models, this logic always runs at 2.2GHz.

The Lisbon die provides three HT 3.1 interfaces. For Adelaide, however, AMD saves power by underclocking these interfaces at the HT 1.0 frequency—800MHz instead of 3.2GHz. (AMD’s power ratings for Adelaide assume this underclocking.) HT doubles the data rate by using both edges of the clock signal, so HT 1.0 supports 1.6 billion transfers per second (1.6GT/s), and HT 3.1 supports 6.4GT/s. Although HT interfaces can be 32 bits wide, AMD processors use 16-bit pathways. Therefore, San Marino’s peak data rate is 12.8GB/s per link, and Adelaide’s peak data rate is 3.2GB/s per link. The Socket-F Opterons peaked at 9.6GB/s per link.

In a standard 2P configuration, only two HT connections are needed: one to connect to the system logic (south bridge) and the other to connect to the second processor, as Figure 1 shows. (The second processor often doesn’t need a south bridge, so it uses only one of its HT interfaces.) For this reason, AMD markets the Opteron 4100 as having only two HT interfaces. Because Lisbon supports a third HT interface for the 4P configuration in Magny-Cours, AMD decided to pin out the third interface for the Opteron 4100 as well. This interface is on the opposite side of the package, giving system designers more flexibility when laying out their boards.

The third HT channel can be used to increase the interchip bandwidth in a 2P configuration. Given HT 3.1’s higher speed, however, this configuration is largely unnecessary. AMD’s power ratings assume that the third channel is not used.

Opteron Powers Down

AMD’s average CPU power (ACP) ratings for Opteron 4100-series chips don’t include the system-logic chip, which adds 10W to 15W. In any case, AMD advises data-center customers to measure a server’s power consumption at the wall using their own power meters to calculate the required amount of computer-room cooling. (Intel makes the same recommendation, emphasizing system power over component power.)

AMD’s power ratings should be taken with a grain of salt. Rather than providing a specific power measurement for each product, the company places them into three broad categories (Intel uses a similar system). Thus, these power levels are more of a “not to exceed” rating than an indication of the exact power used by each product.

Within each category, certain chips with significantly different performance appear to use vastly different power.

For example, the four-CPU Opteron 4122 runs at 2.2GHz, and the six-CPU Opteron 4184 runs at 2.8GHz, yet both chips are rated at 75W. Because of manufacturing variations, however, the speed and power of any particular die can vary significantly. Only chips at the high end of the curve satisfy the 4184’s speed and power ratings; the scarcity of these chips determines the 4184’s higher price (\$316 versus \$99). Any functional part, by contrast, probably meets the minimal specifications of the 4122. This scheme allows AMD to maximize yield while charging a premium for its fastest and most power-efficient chips.

As Table 2 shows, the new Opteron products, despite their improvement, can’t match the power efficiency of Intel’s Westmere-based Xeon products. After adjusting AMD’s reported ACP rating to make it more comparable with Intel’s TDP (see the sidebar, “ACP vs. TDP”), we calculate that the Opteron 4164 EE, AMD’s most energy-efficient chip, delivers 1.65 SPECint_rate (baseline) per watt. By contrast, Intel’s lowest-power Xeon, the L5630, delivers 2.26 SPECint_rate per watt in a dual-processor configuration. The Xeon L5640 is even more efficient, but only slightly.

These SPECint per watt ratings, while crude, correspond to reported SPECpower scores. The Xeon L5640 is 41% more efficient than the Opteron 4164 EE on SPECpower and 27% more efficient on SPECint per watt. When running bursty cloud-computing workloads, higher performance allows the servers to complete their tasks and return to sleep state more quickly. Alternatively, higher performance reduces the number of servers required in the data center, which also saves power. Thus, servers using the newer Xeon processor will use considerably less electricity than servers using the energy-efficient AMD processor.

	AMD Opteron 4164 EE	Intel Xeon L5630	Intel Xeon L5640	Intel Xeon L5530
CPU Speed	1.8GHz	2.133GHz	2.267GHz	2.4GHz
# of CPUs	6 CPUs	4 CPUs	6 CPUs	4 CPUs
Total Cache	9MB	13MB	13.5MB	8MB
IC Process	45nm	32nm	32nm	45nm
Power	40W‡	40W TDP	60W TDP	60W TDP
DP SPECint*	132	181	276	194
SPEC/CPU	11.0/CPU	22.6/CPU	23.0/CPU	24.3/CPU
SPEC/W	1.65/W	2.26/W	2.30/W	1.62/W
SPECpower†	2,106	Not pub.	2,979	2,357
List Price	\$698	\$551	\$996	\$744

Table 2. AMD’s energy-efficient Opteron versus Intel’s low-power Xeon. The low-power Xeon processors outdo Opteron in both performance per CPU and, more importantly, performance per watt. *Best published dual-processor SPECint2006_rate baseline result; †best published dual-processor SPECpower_ssj2008 result in ssj_ops/W. ‡Our estimate of the Intel-equivalent TDP, which is 25% higher than AMD’s ACP rating of 32W; AMD quotes 35W TDP, an unusually small difference between ACP and TDP. (Source: vendors and www.spec.org, except ‡The Linley Group estimate)

Power Plays: ACP vs. TDP

AMD's power ratings differ from Intel's, making apples-to-apples comparisons difficult. Normally, AMD quotes average CPU power (ACP), which is essentially what engineers call a chip's typical power consumption. AMD measures ACP by running several server workloads (TPC-C, SPECcpu2006, SPECjbb2005, and STREAM) at 100% CPU utilization under typical temperature and voltage conditions. To the geometric average of CPU power consumed by these workloads, AMD adds a small safety margin.

According to AMD's own data, ACP is usually about 65% to 70% of maximum power. AMD measures maximum power by running a "thermal virus" program that simultaneously exercises every part of a processor that's running at maximum voltage and case temperature.

Intel instead quotes thermal design power (TDP) for its processors. According to Intel, TDP represents the worst-case power for any commercial software (excluding unnatural code, such as a thermal virus). In other words, where AMD takes an average across several applications, Intel takes the highest power across all applications. Intel's TDP ratings tend to be about 85% to 90% of maximum power, which it measures in a fashion similar to AMD's. Thus, we conclude that AMD's ACP must be adjusted upward by about 25% to be compared with Intel's TDP. AMD says our 25% adjustment is excessive. Even a 20% or 15% adjustment, however, wouldn't radically alter our conclusions. We will revisit the debate in future articles.

To sidestep this confusion, the industry has developed a standard power-efficiency benchmark called SPECpower_ssj2008. SPECpower is a system-level benchmark that tests the power consumed by the system (measured at the wall) when running the SPECssj workload at various levels of CPU utilization. The SPECpower rating is an average efficiency across all utilization levels. SPECssj is based on SPECjbb2005, which represents the workload of a Java server. SPECjbb, and thus SPECssj, stress the CPU, cache, and memory subsystems. Processors with larger caches fare better on this test. This situation often benefits Intel, which crams more cache

when running the same workloads.

This difference in power efficiency comes mainly from AMD's use of 45nm technology. As Table 2 shows, Intel's 45nm Xeon provides performance per watt similar to that of the 45nm Opteron processor. Intel's 32nm Westmere products have been in volume production since early this year, whereas AMD will not deliver 32nm processors until 1H11.

Opteron Has Price Advantage

This situation forces AMD to compete on price. At \$698, the Opteron 4164 EE is no bargain: it costs more than the Xeon L5630 while delivering 30% less performance at the same power. But the Opteron 4162 EE operates just a tick slower at 1.7GHz, and it carries less than half the list price. At \$316, this product is the least-expensive DP server processor available that is capable of at least 1.5 SPECint per watt. Compared with the L5506, Intel's least-expensive low-power Xeon, the Opteron 4162 EE delivers similar performance at much lower power, as Table 3 shows, and its list price is 25% less.

We've also matched the Opteron 4180 against Intel's low-cost Xeon E5504. In this comparison, the AMD processor delivers 29% more performance, albeit at slightly higher power, and at a 16% lower list price to boot. Overall, the Opteron 4180 has a 54% price/performance advantage. AMD offers less expensive versions, but even the \$99 product doesn't exceed the Opteron 4180's performance per dollar.

The Xeon 550x products are built in 45nm technology, and Intel disables Hyper-Threading on these models. These factors significantly reduce performance and performance per watt. Intel's least-expensive Hyper-Threaded product carries a list price of \$373, and its least-expensive 32nm product costs \$378. Except for the 4164 EE's exceedingly optimistic price (\$698), AMD has chosen to price Opteron 4100 chips below this threshold, so they will compete only against Intel's crippled server processors.

Under Intel's Price Umbrella

Despite considerable differences in microarchitecture, AMD's Lisbon CPU and Intel's Nehalem CPU ultimately deliver similar performance per watt. Although Intel's beefier CPUs deliver roughly twice the performance of AMD's, they also use roughly twice the power per core when manufactured in the same IC-process node. AMD is

	AMD Opteron 4162 EE	Intel Xeon L5506	AMD Opteron 4180	Intel Xeon E5504
CPU Speed	1.7GHz	2.13GHz	2.6GHz	2.0GHz
# of CPUs	6 CPUs	4 CPUs	6 CPUs	4 CPUs
Total Cache	9MB	4MB	9MB	4MB
IC Process	45nm	45nm	45nm	45nm
Power	40W (equiv)†	60W TDP	94W (equiv)†	80W TDP
DP SPECint*	127†	133	172†	133
SPEC/CPU	10.6/CPU	16.6/CPU	14.3/CPU	16.6/CPU
SPEC/W	1.59/W	1.11/W	0.91/W	0.83/W
List Price	\$316	\$423	\$188	\$224
SPEC/Dollar	0.201/\$	0.157/\$	0.457/\$	0.297/\$

Table 3. Comparison of low-priced versions of Opteron and Xeon. These low-end Xeon products do not support Hyper-Threading. *Best published dual-processor SPECint2006_rate baseline score. (Source: vendors and www.spec.org, except †The Linley Group estimate)

Price and Availability

AMD's Opteron 4100-series products are in production. List pricing for the nine models appears in Table 1. For additional information on these processors, [click here](#) for access

also able to pack more CPUs onto a single chip while still charging a lower price than Intel does.

The companies are also taking different manufacturing paths, but these paths are not as convergent. Whereas Intel continues to invest heavily in its industry-leading fab technology, AMD spun off its fabs to form GlobalFoundries, hoping that new foundry customers will increase the level of investment beyond what AMD could afford on its own. AMD's new strategy will take time to play out (see [MPR 11/24/08-01](#), "AMD's Fresh Start").

Meanwhile, the company is stuck with 45nm technology, whereas Intel is already using a 32nm process, which provides inherent advantages in die size, speed, and power.

Although Intel's 32nm Xeon processors provide a clear advantage over AMD's 45nm Opteron 4100 in performance per watt, Intel chooses to price these products at \$378 and higher. This pricing protects Intel's margins and allows it to maximize the revenue from its 32nm wafers. It also creates an umbrella under which AMD can safely compete.

At these lower price points, the Opteron 4100 offers clear advantages over the 45nm Xeon 55xx products. The AMD processors have a price/performance advantage of up to 54% over competing Xeon products and offer better performance per watt as well. Server designers willing to buy up to the Xeon 56xx will find better power efficiency, but designers seeking the best price/performance should look at the Opteron 4100. ♦

To subscribe to *Microprocessor Report*, access www.MPRonline.com or phone us at 408-945-3943.