
ACP – The Truth About Power Consumption Starts Here

Quad-Core AMD Opteron™ processors improve power efficiency with several key technologies:

Enhanced AMD PowerNow!™

Technology: For reduced power consumption by the entire processor.

Native quad-core technology enables enhanced power management across all four cores.

Independent Dynamic Core

Technology allows each core to vary its frequency, based on the specific needs of the system. This allows for more precise power management to reduce data center energy consumption and thereby reduce total cost of ownership (TCO).

Dual Dynamic Power Management™

allows each processor to maximize the power-saving benefits of Enhanced AMD PowerNow! technology without compromising performance. Dual Dynamic Power Management can reduce idle power consumption and allow for per-processor power management in multi-socket systems to decrease power consumption.

AMD CoolCore™ Technology: To reduce power consumption within each core AMD CoolCore Technology evaluates which parts of the die – the cores, the memory, or both – are needed to support currently running applications. It can cut power to unused transistor areas to reduce power consumption and lower heat generation.

AMD Smart Fetch Technology

helps reduce power consumption by allowing idle cores to enter a “halt” state, causing them to draw even less power during processing idle times.

For years, CIOs never had to worry about the cost of power – that financial cost was always buried somewhere in the facilities budget and companies rarely bothered to connect the dots between energy consumption and IT infrastructure.

As the world shifted from propriety expensive mainframe technologies to lower cost industry-standard servers, the demands of power on the data center grew and the relationship between energy costs and consumption became clearer, registering a “concern” for the CIO. The lower cost of industry-standard servers offered companies the ability to deploy more servers within their data center infrastructure, increasing the compute capabilities and included power and cooling requirements. But still, as long as there was no interruption to the delivery of power or limits set, CIOs remained unconcerned.

Then, industry-standard ultra-dense low cost 1U servers arrived, followed by even denser blade servers. Facilities couldn't supply more power as data center growth had outstripped the power capacity. There was no more room to expand power capacity, so some businesses had to stop growing their infrastructure. New IT deployments were put on hold and businesses started to be concerned about being uncompetitive based on their inability to react to changes in the market or new opportunities. Welcome to the crisis.

As companies scrambled to adapt to this new world where power mattered in the data center, getting a handle on exactly what needed to be done was clear – power consumption needed to be reduced. It's easy enough to say, but harder to practice because every generation of technology was more power-hungry than the last. Where do you begin, especially when you may not have visibility to your power consumption at the platform level? The simplest way is to take a system-level view of power consumption and address the different components in the platform on a per-device basis. Customers will quickly gravitate to the processor because this component represents a large percentage of the power consumed in the system. Not many customers have the tools to truly measure processor power consumption, at best they can only measure the full system.

The best way to measure a server's power consumption is the power meter, an inexpensive tool that is plugged into the wall, and then your device, like a server, can be plugged into the power meter. The meter displays the wattage drawn "at the wall" and allows you to analyze the power consumption under a variety of different utilization levels. As you would expect, a server sitting idle draws significantly less power than a server that is under a full load.

It is of little value to measure power consumption by only looking at the spec sheets for different components, adding the totals together, because these generally only report the maximum power consumption. This scenario is like the car with a speedometer that tells you the maximum speed is 150MPH – it's a maximum reading but it doesn't reflect daily usage. Unfortunately, AMD's own customer research indicates that nearly three out of four customers only use processor specification sheets to determine power levels. Based on the architectural differences between AMD and competitive processors, like AMD's integrated memory controller and AMD's use of energy-efficient DDR-2 memory, the numbers at the platform level can tell a very different story at the wall compared to the specification sheets.

As customers begin to seek more information regarding power consumption of processors as well as platforms, they are putting pressure on their vendors to deliver more detailed, and most importantly, more accurate power measurements to help them better determine the true power consumption. The days of simply being able to quote the maximum power from a design standpoint are gone forever – customers demand more accuracy in order to help them best plan for their infrastructures.

Based on these customer demands, AMD has developed a new metric to more accurately reflect power consumed by the processor as well as the memory controller (which is integrated into all AMD Opteron™ processors) during relevant peak workloads. The new Average CPU Power (ACP) metric is designed to give customers a more accurate idea of the power consumed by the processor, allowing customers to more realistically forecast their power budget assessments to estimate how much power might actually be consumed at the wall and more accurately allow them to plan their data center power and cooling infrastructure.

With AMD Opteron™ processors, the TDP, or Thermal Design Power, was used to represent the maximum power for the processor. Because this was an engineering design specification, it was significantly easier for AMD to report to customers. But that didn't help customers accurately assess the power required for the processor, because a processor

Integrated DDR2 DRAM Memory Controller: low-power memory to reduce power consumption. AMD's Integrated Memory Controller works exclusively with high bandwidth, energy-efficient DDR2 memory. It incorporates Memory RAS for increased fault tolerance to reduce system downtime and increase system reliability.

DDR2 memory can improve overall TCO by reducing memory power consumption almost 8 watts per DIMM compared to fully buffered DIMM memory technology.

with a 95W TDP may not break the 70W mark under extremely high workloads – just like a car with a 150MPH speedometer rarely hits speeds above 90MPH. Some customers were unnecessarily limiting the amount of growth within the rack based on an over-estimated power budget and potentially sacrificing data center efficiency. They demanded more accuracy, and that led AMD to develop ACP, a better way to think about processor power consumption.

With the introduction of the Quad-Core AMD Opteron™ processors, several new power management enhancements, including AMD Smart Fetch technology, AMD CoolCore™ technology, Dual Dynamic Power Management™, and Enhanced AMD PowerNow!™ technology with Independent Dynamic Core technology all help to reduce power consumption of the processor. But to further help customers more accurately model power consumption, ACP was developed to better match customer needs.

Breaking Down Processor Power Consumption

The Quad-Core AMD Opteron™ processors have many power efficiency enhancements to reduce overall power consumption. A new feature called Dual Dynamic Power Management™ (DDPM) provides the processors with additional power rails which are dedicated to the processor cores, HyperTransport™ technology links, and memory controller. Separating the power rails of the cores allows each of the cores to independently adjust frequency for the given workload while also allowing for voltage changes to the cores. Not only does this technology enable considerable power savings at the processor level during non-peak workloads, it also allows AMD to capture the power consumption of the individual sections of the processor during the power measurements. The ACP values for each power band include the power for the cores, memory controller, and HyperTransport™ links.

HOW POWER IS MEASURED

To accurately measure in-system power consumption by the processor, one must isolate the power consumed by the processor from the power consumed by the rest of the motherboard. Motherboards contain multiple voltage regulator modules dedicated to delivering power to the individual power rails of the processor, which satisfies this isolation requirement. However, the implementation of these voltage regulators typically does not allow for direct power measurement without modifications to the motherboard and these regulators. Thus, ACP measurements were taken on specially instrumented internal engineering server platforms. It should be noted that such instrumentation does not change the power requirements or characteristics of the processor itself. This instrumentation allowed for the direct measurement of all power rails supplying power to the processor while the system was running. Measurements were captured multiple times per second using a high precision digital multimeter.

PART SELECTION

In general, inherent process variations in semiconductor manufacturing lead to a normal distribution of manufacturing output vs. manufacturing targets. As it pertains to power consumption, one can think of this as processors being slightly colder (lower power) than target, on target, and slightly hotter (more power) than target. As with any normal distribution, the bulk of the chips are in the “on-target” region. To be conservative, parts used for ACP power measurement were sampled from the hotter side of the distribution.

WORKLOAD SELECTION

A suite of industry accepted server workloads was selected as a basis for determining ACP. These include floating point, integer, java, Web, memory bandwidth, and transactional workloads.

Using a suite of server workloads that represents the breadth of typical server applications allows for a better representation of enterprise server class workloads used by end customers. These workloads were Transaction Processing Performance Council (TPC-C), SPECcpu2006, SPECjbb2005, and STREAM. The geometric mean of measurements, taken during these workloads, is the ACP.

TEST CONDITIONS

Given the goal of representing typical power usage in real world conditions, environmental test conditions were chosen to reflect that aspect (room temp of 70°F, server's fan heat sink used, closed case, etc.) The power for the cores, memory controller, and HyperTransport™ links was logged multiple times per second throughout the entire duration of the workload tested, and the time-averaged power consumption for that workload was calculated.

The results across the suite of workloads are used to derive the ACP number. The ACP value for each processor power band is representative of the geometric mean for the entire suite of benchmark applications plus a margin based on AMD historical manufacturing experience.

So now that you can see how the ACP measurements are derived and how they are rooted in representative daily usages, it should be clear that this metric is a better tool for customers to assess data center power needs. By utilizing ACP, customers can make a more educated estimate of their true power consumption. TDP, representing the thermal design power, is a rudimentary indicator of consumption, and may leave customers overestimating their data center infrastructure needs.

In doing this, customers can potentially waste valuable data center space, resulting in un-optimized data centers that cost companies real dollars through their inefficiency.

Processor ACP Value	Power for Cores (% of ACP)	Power for memory controller (% of ACP)	Power for HyperTransport™ links (% of ACP)	Power at idle (w/ power management) (% of ACP)	Thermal Design Power Specification (TDP)
105W	80%	16%	4%	≤25%	137W

FIGURE 1: QUAD-CORE AMD OPTERON™ SE PROCESSORS (137W THEORETICAL MAXIMUM POWER)

Processor ACP Value	Power for Cores (% of ACP)	Power for memory controller (% of ACP)	Power for HyperTransport™ links (% of ACP)	Power at idle (w/ power management) (% of ACP)	Thermal Design Power Specification (TDP)
75W	75%	20%	5%	≤25%	115W

FIGURE 2: QUAD-CORE AMD OPTERON™ PROCESSORS (115W THEORETICAL MAXIMUM POWER)

Processor ACP Value	Power for Cores (% of ACP)	Power for memory controller (% of ACP)	Power for HyperTransport™ links (% of ACP)	Power at idle (w/ power management) (% of ACP)	Thermal Design Power Specification (TDP)
55W	69%	25%	6%	≤30%	79W

FIGURE 3: QUAD-CORE AMD OPTERON™ HE PROCESSORS (79W THEORETICAL MAXIMUM POWER)

Processor ACP Value	Power for Cores (% of ACP)	Power for memory controller (% of ACP)	Power for HyperTransport™ links (% of ACP)	Power at idle (w/ power management) (% of ACP)	Thermal Design Power Specification (TDP)
40W	61%	31%	8%	≤35%	60W

FIGURE 4: QUAD-CORE AMD OPTERON™ SE PROCESSORS (60W THEORETICAL MAXIMUM POWER)

ACP – THE NEW STANDARD?

The natural question becomes “is this a new industry standard measurement?” AMD believes the methodology behind ACP should become the new standard, as it includes both the processor cores and the memory controller. Today, companies can mask the true power consumption by fixating on the processor itself and neglecting to mention the power consumption of the external memory controller. Datacenter managers should measure power of the entire platform while running relevant workloads in order to get the most accurate estimate of platform-level power consumption. ACP is not a replacement metric for platform-level engineering, thermal and power specifications outlined in the AMD thermal and power datasheets should be referenced as guidance for platform-level thermal and power design.

SUMMARY

Customers are demanding a better methodology for determining the power consumption of systems in order to better plan their data center deployments. For years, TDP was the guideline that was employed, but as more and more power saving features are integrated into the processor, the delta between the TDP specification and the typical CPU power consumption has grown significantly.

AMD believes that ACP is a better way of thinking about typical CPU power that more accurately reflects the power consumption levels that customers can see in real life environments.

By choosing the “hot” parts and employing a rigorous testing regimen designed to err on the conservative side, ACP can provide a measurement above what customers expect to see in daily use but nowhere near as overstated as TDP has been in the past.