

http://www.internetnews.com/ent-news/article.php/10794_3699846_2

Memory: The Overlooked Power Issue

Before you deck that virtualization server with 64GB of RAM, you might want to think about your electric bill.

September 17, 2007

By Andy Patrizio: [More stories by this author.](#)

There has been considerably emphasis on the subject of power efficiency and microprocessors, mostly because Intel and AMD have been beating the drum of lower power consumption louder than Keith Moon in [his prime](#).

However, there is an overlooked element inside that refrigerator-sized humming black plastic and metal box that is sucking up considerable amounts of power and generating heat: the memory.

Memory sticks are plastic and silicon and don't, for the most part, require a heat sink like the CPU. However, they draw a measurable amount of power, and in this new era of 64-bit computing, they are becoming the hidden electric bill.

The great limiting factor of 32-bit computing was its 4GB memory limit. Now that we are firmly in the era of [64-bit computing](#), that limitation has been smashed, consigned to the ash heap along with single core CPUs and the AGP bus. A 64-bit machine can easily handle 64GB, 128GB or more of memory, up to 16 petabytes.

With the mania around virtualization and people running five, 10 or 20 virtual servers on one box, that means a lot of memory is needed in these beasts. Instead of the days where a server had a pair of 2GB memory sticks, they now have as many as 32 memory sticks in them. Many four-socket motherboards have eight memory banks per socket.

A memory stick, or DIMM ([define](#)), can consume up to 12 watts. Multiplied by 32, and suddenly you've got what could be the single biggest power draw in the server next to storage, depending on the kind of memory you use. What has always been an afterthought in system purchases is going to become a major concern for the power constrained.

DDR, or dual data rate memory, is used in desktops, laptops and video cards. Currently, computers are using DDR2, the second generation of the technology, while DDR3, designed for faster speed and lower power, is under development. AMD uses DDR2 memory in its Opteron-based servers.

However, Intel uses fully buffered DIMM, or FBDIMM, in its Xeon-based servers. FBDIMM has a chip smack in the middle of the stick called an Advanced Memory Buffer (AMB), which is not used in regular DDR2 sticks.

The AMB is a serial interface that increases the bandwidth of memory and makes it easier to put eight sticks in a bank without degradation of performance. In a bank of eight DIMM slots, memory would degrade in performance once you go beyond four sticks. It's also very good for accessing large amounts of sequential memory and offers error correction that DDR doesn't have.

The drawback, though, is its power draw. An FBDIMM can consume anywhere from five to eight watts more power than a DDR2 stick, which adds up quickly as you add more memory. The reason for the high power draw is the AMB never really has a chance to power down when under low workloads. A benchmarking lab called Neal Nelson Benchmark Laboratories ran a series of tests on power consumption.

It found that while idle, a dual processor, dual core Xeon server consumed 119.3 watts of power while a dual core, two-processor Opteron server drew only 66.7 watts, a 44 percent power savings for the Opteron-based machine. Both machines had 8GB of memory.

Under a full load of 500 simulated users, the Xeon system drew 145.5 watts while the Opteron drew 134.8 watts, giving the Opteron a 7.8 percent advantage. Intel declined to comment on the test results.

Another series of tests was done by the hobbyist site Anandtech, and it found that the FBDIMMs in a Xeon server consumed 862 percent more power than the DDR2 in an AMD server.

AMD won't hesitate to point out this disparity. "When you run your datacenter, you're worried not just about peak loads but what are the loads throughout the whole range from idle to peak," said Brent Kerby, Opteron product manager. "If you look at idle measurements on a stick of DDR2, it's considerably lower than on a stick of FBDIMM because they always have to power that AMB chip at all times."

Intel acknowledges this issue and says it will address it in future designs. "Our primary motivation for FBDIMM was to have capacity without sacrificing performance. That was the benefit of the buffer," said Scott Huck, competitive architect in the server marketing team at Intel.

"You can keep stacking FBDIMMs and not lose performance. We were trying to address the capacity without having to make a trade off in performance. It was an issue of looking at where the industry was going. Virtualization was all the rage and it would need both performance and capacity and there was no room for a trade off," he said.

Both AMD and IBM chose DDR2 over FBDIMM, for different reasons, but now both are noticing the difference in the power draw. (IBM made a custom motherboard chip for its X-Architecture servers that does the work of the AMB chip, so it says it doesn't need FBDIMM in its Xeon servers).

"We have to think about what it does for overall system efficiency. Then we have to think about what it does for overall power and then how it affects overall system design," Margaret Lewis, director of commercial marketing for AMD told *InternetNews.com*.

"Along the lines in evaluating where things were going, we decided DDR memory would be more power efficient, so that was the decision we made even though it was counter to the trends in the industry," she added.

Jay Bretzmann, manager of product marketing for IBM's X servers, added "power was one reason we didn't go with FBDIMM, but the price premium was the big one. It used to be much more expensive, but that gap has been closing since then. But power is one issue that's been persistent," he said.

With the typical utilization rates today for unvirtualized servers at 7 to 12 percent, it's not very efficient for these servers to keep their memory revved up. Even with virtualization, Bretzmann points out servers will be hitting the 30 to 50 percent utilization mark. But if memory can't dial back under a low load, it's a lot of power for nothing.

Some like it hot

There is a secondary issue with FBDIMM, and that's heat. FBDIMMs need heat dissipaters because they build heat rather quickly. If the memory components are referenced randomly, as is often the case, then the heat on an unbuffered module is spread evenly throughout the stick.

But on an FBDIMM, all memory access has to go through the AMB, and heat collects around it very quickly under a heavy load and is not spread around like a DDR stick. "All that heat collects in that one spot. As that one component gets too hot, the Intel processor recognizes it and throttles down the performance because it's getting too hot," said Bob Merritt, an analyst for Semico Research.

Huck said FBDIMM is a "1.0" design right now, the first iteration, and that going forward there is a concerted effort to dial down on the AMB in an idle state. Work in that area might also benefit other serial interfaces within the PC, he said.

"We're working on determining when no data is being passed through the interface and shutting down power on that data path," he said. "That same type of technology can be used on any other serial interface, be it PCI Express or other types of serial interfaces within the system, so when you are not transferring data, the bus can go to a low power mode."

Intel is holding its annual Developer's Forum this week and promises some news in this space on Wednesday. For now, memory power draw is a major differentiator between AMD and Intel. The question is will it affect buying decisions.

"I don't know if it's become a buying issue," said Merritt. "Power consumption is becoming more of a problem, and how you avoid that is becoming a concern, but there are lots of ways to measure that."

He added that if Intel doesn't see a loss of sales as a result of its memory choice, then it will stick with it, but if there is a loss of sales due to memory power draw, then it can always counter with better price or price/performance.

Bretzmann added "It's not front and center in people's minds yet, but one of the mantras from IBM will be 'go green.' We will continually build green features into our servers," and low memory power consumption will be a key element, he said.