

## Re: Opteron box and 4Gb memory

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-11/msg01547.html>

---

- *From:* "J.A. Magallón" <[jamagallon@xxxxxxx](mailto:jamagallon@xxxxxxx)>
  - *Date:* Tue, 6 Nov 2007 00:03:19 +0100
- 

On Mon, 5 Nov 2007 13:50:40 -0500, [lsorense@xxxxxxxxxxxxxxxxxxxx](mailto:lsorense@xxxxxxxxxxxxxxxxxxxx) (Lennart Sorensen) wrote:

On Mon, Nov 05, 2007 at 07:45:21PM +0100, J.A. Magallón wrote:

Well, problem solved...

I'm going to kill all pc assemblers in the world... Someone should teach them to learn manuals before assembling anything but a power chord.

The memory was not paired, so the motherboard was not interleaving the access.

With no inter-node but with inter-module interleaving, and a couple 1Gb sticks

for each processor now I get something like:

```
cicely:~/bn> bn
name: cicely.cps.unizar.es
arch: x86-64
proc: 4 x x86_64 @ 2200 MHz
ram: 3555 Mb
os: unx, Linux, 2.6.23.1-desktop-1mdv
cc: gcc-4.3.0
vector size : 8 x 1024 x 1024
allocation: 0.02 ms
int scl add: ..... 60.56 ms, 138.52 Mips | 62.96 Mips /GHz
int scl mul: ..... 59.34 ms, 141.36 Mips | 64.26 Mips /GHz
flt scl add: ..... 59.01 ms, 142.16 Mflops | 64.62 Mflops/GHz
flt vec add: ..... 14.79 ms, 567.06 Mflops | 257.75 Mflops/GHz
flt scl mul: ..... 59.02 ms, 142.12 Mflops | 64.60 Mflops/GHz
flt vec mul: ..... 14.82 ms, 566.19 Mflops | 257.36 Mflops/GHz
total: 5019.86 ms
```

Much better, but not like the other opteron box.

My processors are higher than Rev E0, because the BIOS does not let me choose

the 'software' hole. If I activate the 'hardware hole', I see all the memory I can:

Re: Opteron box and 4Gb memory

```
cicely:~/bn> free
total used free shared buffers cached
Mem: 3640628 214496 3426132 0 21240 84184
-/+ buffers/cache: 109072 3531556
Swap: 4200988 0 4200988
```

3.64 Gb. The rest is eaten by the graphics card, as I could read in the AMD site. Don't know if mem=4096 to boot the kernel would help, even if it is possible (don't think so, as it looks like a BIOS mis-feature).  
The ram is DDR 400.

The video card is stealing 300MB of ram? What for? What does the mtrr and e820 map look like with the hardware hole enabled?

(correction, this was not AMD website but SuperMicro's)

I just said the same... Board is a SuperMicro H8DCE. From the FAQ section at supermicro:

Question

I have an H8DCE motherboard with 4 x 1GB DIMMS installed but the amount of memory displayed in the BIOS is 3.865GB and in 64-bit XP is 3.76GB. I have the hardware memory hole option enabled in BIOS (rev 1.1a). How I can get the board to count the full 4GB of memory?

Answer

The total available size depends on the PCI-e card you are using; some high-end cards may occupy more memory. For example, with a Quadro FX4500 on the H8DCE with the memory hole enabled, 4GB memory will show up as 3728MB in BIOS and 3.64GB in Windows. For some low-end PCI-e VGA cards, it may show up as 4048MB in BIOS.

Why ? Who knows...

Chipset is all nVidia. I have a GeForce 8800GTX with 768 Mb. It eats up 400Mb.

This are my settings:

BIOS-provided physical RAM map:

```
BIOS-e820: 0000000000000000 - 000000000009fc00 (usable)
BIOS-e820: 000000000009fc00 - 00000000000a0000 (reserved)
BIOS-e820: 00000000000e6000 - 0000000000100000 (reserved)
BIOS-e820: 0000000000100000 - 0000000009ffd000 (usable)
BIOS-e820: 0000000009ffd000 - 0000000009ffde00 (ACPI data)
BIOS-e820: 0000000009ffde00 - 00000000a0000000 (ACPI NVS)
BIOS-e820: 00000000fec00000 - 00000000fec01000 (reserved)
BIOS-e820: 00000000fee00000 - 00000000fee01000 (reserved)
BIOS-e820: 00000000ff780000 - 0000000100000000 (reserved)
BIOS-e820: 0000000100000000 - 0000000147000000 (usable)
```

Re: Opteron box and 4Gb memory

```
cicely:~# cat /proc/mtrr
reg00: base=0x100000000 (4096MB), size=1024MB: write-back, count=1
reg01: base=0x140000000 (5120MB), size= 64MB: write-back, count=1
reg02: base=0x144000000 (5184MB), size= 32MB: write-back, count=1
reg03: base=0x146000000 (5216MB), size= 16MB: write-back, count=1
reg04: base=0x00000000 ( 0MB), size=2048MB: write-back, count=1
reg05: base=0x80000000 (2048MB), size= 512MB: write-back, count=1
```

This is with BIOS set for MTRR=Discrete.

With MTRR=Continuous, the mtrr's are simpler, a full range and a non-usable hole. Which is better for Linux ? Many separate usable zones or one big zone and an un-usable hole ?

BTW, mtrr formatting should be set to 0x%013lx000, to get them aligned with nowadays memory amounts and similar to e820 map, 16 hex digits...

Anyways, can I trust what dmidecode says ? I installed the ram as the board manual said in banks 1A+1B (not 2A+2B) for each processor, but this program says this:

```
BANK0 64Mb BANK4 64Mb
BANK1 64Mb BANK5 64Mb
BANK2 1024Mb BANK6 1024Mb
BANK3 1024Mb BANK7 1024Mb
```

I would always have thought that BANK0 would be slot 1A in first processor, but it looks like not...  
And where do the 64 Mb blocks come from ?

Well if you have 4 sticks of 1GB, then I would hope they are installed as a pair for each CPU so that both CPUs can have dual channel ram directly connected.

That's how they are plugged. The strange thing is that I filled the 1st and 2nd slot (by mobo manual numbering), but dmidecode (BIOS?) thinks they are 3rd and 4th. I don't know if it matters, probably not, but who knows...  
This s\*\*\*t of PC architecture is a Pandora's Box.

I have no idea where the 64Mb comes from.

Nor I do...

Re: Opteron box and 4Gb memory

Re: Opteron box and 4Gb memory

—  
J.A. Magallon <jamagallon()ono!com> \ Software is like sex:

\ It's better when it's free

Mandriva Linux release 2008.1 (Cooker) for i586

Linux 2.6.23-jam01 (gcc 4.2.2 20070909 (4.2.2-0.RC.1mdv2008.0)) SMP PREEMPT

09 F9 11 02 9D 74 E3 5B D8 41 56 C5 63 56 88 C0

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>