

# x86\_86 SMP megaraid\_mbox hangups and panics.

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-04/msg01879.html>

---

- *From:* dormando <[dormando@xxxxxxxxxx](mailto:dormando@xxxxxxxxxx)>
  - *Date:* Tue, 11 Apr 2006 17:33:45 -0700
- 

Hey,

I had originally sent this to linux-scsi, but was told to try the maintainer/kernel list instead.

Having hangs and kernel panics trying to boot AMD64 SMP with an LSI MegaRaid 320-1 card using megaraid\_mbox driver. I'm trying to boot a monolithic vanilla 2.6.16.1 64-bit SMP on a SuperMicro Opteron server running a dualcore AMD 270 CPU and 8G of RAM.

Most of the time the server hits: "megaraid: probe new device" – with the device information, then hangs and starts the 180 second countdown: "megaraid: wait for FW to boot [blah]"  
After which I get a VFS panic for not having a root disk.

If it does not hit this, there is an immediate kernel panic somewhere in megaraid\_ack\_sequence. There are two panics for the two different times megaraid\_ack\_sequence is called in the driver. The top level seems to be in the megaraid\_isr function.

One trace looks generally like:

```
hrtimer_run_queues, megaraid_isr, handle_IRQ_event, __do_IRQ, do_IRQ, default_idle, ret_from_intr,  
thread_return, default_idle, cpu_idle.  
RIP megaraid_ack_sequence+298, RSP
```

The other one ends the same way, starts differently. Easy enough to find in the code.

I've tried five identical machines and they all do the same thing. So here's the breakdown of what I narrowed: (unless otherwise specified, all "does not work" has the same symptoms described above).

2.6.15.7 64-bit SMP – does not work

2.6.16.1 64-bit MSI/NUMA disabled – does not work.

2.6.16.1 64-bit ACPI disabled – does not work.

2.6.16.1 32-bit SMP – works every time. (then panics against my 64-bit OS ;)

2.6.16.1 64-bit UP – works every time.

2.6.16.1 64-bit SMP with megaraid\_mbox/mm compiled as modules – Boots all the way sometimes, mostly hangs or panics.

I tried changing the clock values, idle=poll, acpi=off, and twiddled the iommu bits without any luck. So it's looking like an x86-64 SMP specific timing problem with the driver. 32-bit SMP does not appear to be affected.

All related BIOS/firmwares have been upgraded to their latest available versions. Below are an lspci from a working machine, and a cut dmesg. All of the kernel configs were just about identical except for the changes

## x86\_86 SMP megaraid\_mbox hangups and panics.

noted above.

Hope I'm not making an idiot out of myself, but I've spent two weeks twiddling bits and hardware with no luck. If anyone needs more information about the system/setup and what I've tried, there are tons, just ask.  
-Dormando

### LSPCI:

```
0000:00:06.0 PCI bridge: Advanced Micro Devices [AMD] AMD-8111 PCI (rev 07)
0000:00:07.0 ISA bridge: Advanced Micro Devices [AMD] AMD-8111 LPC (rev 05)
0000:00:07.1 IDE interface: Advanced Micro Devices [AMD] AMD-8111 IDE (rev 03)
0000:00:07.2 SMBus: Advanced Micro Devices [AMD] AMD-8111 SMBus 2.0 (rev 02)
0000:00:07.3 Bridge: Advanced Micro Devices [AMD] AMD-8111 ACPI (rev 05)
0000:00:0a.0 PCI bridge: Advanced Micro Devices [AMD] AMD-8131 PCI-X Bridge (rev 13)
0000:00:0a.1 PIC: Advanced Micro Devices [AMD] AMD-8131 PCI-X APIC (rev 01)
0000:00:0b.0 PCI bridge: Advanced Micro Devices [AMD] AMD-8131 PCI-X Bridge (rev 13)
0000:00:0b.1 PIC: Advanced Micro Devices [AMD] AMD-8131 PCI-X APIC (rev 01)
0000:00:18.0 Host bridge: Advanced Micro Devices [AMD] K8 NorthBridge
0000:00:18.1 Host bridge: Advanced Micro Devices [AMD] K8 NorthBridge
0000:00:18.2 Host bridge: Advanced Micro Devices [AMD] K8 NorthBridge
0000:00:18.3 Host bridge: Advanced Micro Devices [AMD] K8 NorthBridge
0000:01:03.0 RAID bus controller: LSI Logic / Symbios Logic MegaRAID (rev 01)
0000:02:05.0 Ethernet controller: Broadcom Corporation NetXtreme BCM5704 Gigabit Ethernet (rev 10)
0000:02:05.1 Ethernet controller: Broadcom Corporation NetXtreme BCM5704 Gigabit Ethernet (rev 10)
0000:03:00.0 USB Controller: Advanced Micro Devices [AMD] AMD-8111 USB (rev 0b)
0000:03:00.1 USB Controller: Advanced Micro Devices [AMD] AMD-8111 USB (rev 0b)
0000:03:04.0 VGA compatible controller: ATI Technologies Inc Rage XL (rev 27)
```

### DMESG:

```
Bootdata ok (command line is root=/dev/sda1 ro )
Linux version 2.6.16.1gaiadb (root@3-18) (gcc version 3.3.5 (Debian 1:3.3.5-13)) #1 SMP Thu Apr 6
17:36:34 PDT 2006
BIOS-provided physical RAM map:
BIOS-e820: 0000000000000000 - 000000000009fc00 (usable)
BIOS-e820: 000000000009fc00 - 00000000000a0000 (reserved)
BIOS-e820: 00000000000e8000 - 0000000000100000 (reserved)
BIOS-e820: 0000000000100000 - 000000007fff0000 (usable)
BIOS-e820: 000000007fff0000 - 000000007ffff000 (ACPI data)
BIOS-e820: 000000007ffff000 - 0000000080000000 (ACPI NVS)
BIOS-e820: 00000000ff780000 - 0000000100000000 (reserved)
BIOS-e820: 0000000100000000 - 0000000280000000 (usable)
ACPI: RSDP (v000 ACPIAM ) @ 0x000000000000f97b0
ACPI: RSDT (v001 A M I OEMRSDT 0x01000604 MSFT 0x00000097) @ 0x000000007fff0000
ACPI: FADT (v002 A M I OEMFACP 0x01000604 MSFT 0x00000097) @ 0x000000007fff0200
ACPI: MADT (v001 A M I OEMAPIC 0x01000604 MSFT 0x00000097) @ 0x000000007fff0380
ACPI: OEMB (v001 A M I OEMBIOS 0x01000604 MSFT 0x00000097) @ 0x000000007ffff040
ACPI: DSDT (v001 H8DA8 H8DA8010 0x00000000 INTL 0x02002026) @ 0x0000000000000000
Scanning NUMA topology in Northbridge 24
Number of nodes 1
Node 0 MemBase 0000000000000000 Limit 0000000280000000
```

x86\_86 SMP megaraid\_mbox hangups and panics.

## x86\_86 SMP megaraid\_mbox hangups and panics.

NUMA: Using 63 for the hash shift.  
Using node hash shift of 63  
Bootmem setup node 0 0000000000000000-0000000280000000  
On node 0 totalpages: 2059484  
DMA zone: 2228 pages, LIFO batch:0  
DMA32 zone: 505896 pages, LIFO batch:31  
Normal zone: 1551360 pages, LIFO batch:31  
HighMem zone: 0 pages, LIFO batch:0  
ACPI: Local APIC address 0xfee00000  
ACPI: LAPIC (acpi\_id[0x01] lapic\_id[0x00] enabled)  
Processor #0 15:1 APIC version 16  
ACPI: LAPIC (acpi\_id[0x02] lapic\_id[0x01] enabled)  
Processor #1 15:1 APIC version 16  
ACPI: LAPIC (acpi\_id[0x03] lapic\_id[0x82] disabled)  
ACPI: LAPIC (acpi\_id[0x04] lapic\_id[0x83] disabled)  
ACPI: IOAPIC (id[0x02] address[0xfec00000] gsi\_base[0])  
IOAPIC[0]: apic\_id 2, version 17, address 0xfec00000, GSI 0-23  
ACPI: IOAPIC (id[0x03] address[0xfebfe000] gsi\_base[24])  
IOAPIC[1]: apic\_id 3, version 17, address 0xfebfe000, GSI 24-27  
ACPI: IOAPIC (id[0x04] address[0xfebff000] gsi\_base[28])  
IOAPIC[2]: apic\_id 4, version 17, address 0xfebff000, GSI 28-31  
ACPI: INT\_SRC\_OVR (bus 0 bus\_irq 0 global\_irq 2 dfl dfl)  
ACPI: INT\_SRC\_OVR (bus 0 bus\_irq 0 global\_irq 2 dfl dfl)  
ACPI: IRQ0 used by override.  
ACPI: IRQ2 used by override.  
ACPI: IRQ9 used by override.  
Setting APIC routing to flat  
Using ACPI (MADT) for SMP configuration information  
Allocating PCI resources starting at 88000000 (gap: 80000000:7f780000)  
Checking aperture...  
CPU 0: aperture @ c000000 size 32 MB  
Aperture from northbridge cpu 0 too small (32 MB)  
No AGP bridge found  
Your BIOS doesn't leave a aperture memory hole  
Please enable the IOMMU option in the BIOS setup  
This costs you 64 MB of RAM  
Mapping aperture over 65536 KB of RAM @ c000000  
Built 1 zonelists  
Kernel command line: root=/dev/sda1 ro Initializing CPU#0  
PID hash table entries: 4096 (order: 12, 131072 bytes)  
time.c: Using 1.193182 MHz WALL PIT GTOD PIT/TSC timer.  
time.c: Detected 1994.357 MHz processor.  
Console: colour VGA+ 80x25  
Dentry cache hash table entries: 1048576 (order: 11, 8388608 bytes)  
Inode-cache hash table entries: 524288 (order: 10, 4194304 bytes)  
Memory: 8154112k/10485760k available (4266k kernel code, 234044k reserved, 1856k data, 252k init)  
Calibrating delay using timer specific routine.. 3995.21 BogoMIPS (lpj=19976058)  
Security Framework v1.0.0 initialized  
Capability LSM initialized  
Mount-cache hash table entries: 256  
CPU: L1 I Cache: 64K (64 bytes/line), D cache 64K (64 bytes/line)

x86\_86 SMP megaraid\_mbox hangups and panics.

CPU: L2 Cache: 1024K (64 bytes/line)  
CPU 0(2) -> Node 0 -> Core 0  
Using local APIC timer interrupts.  
result 12464730  
Detected 12.464 MHz APIC timer.  
Booting processor 1/2 APIC 0x1  
Initializing CPU#1  
Calibrating delay using timer specific routine.. 3988.74 BogoMIPS (lpj=19943722)  
CPU: L1 I Cache: 64K (64 bytes/line), D cache 64K (64 bytes/line)  
CPU: L2 Cache: 1024K (64 bytes/line)  
CPU 1(2) -> Node 0 -> Core 1  
Dual Core AMD Opteron(tm) Processor 270 stepping 02  
CPU 1: Syncing TSC to CPU 0.  
CPU 1: synchronized TSC with CPU 0 (last diff 0 cycles, maxerr 488 cycles)  
Brought up 2 CPUs  
testing NMI watchdog ... OK.  
migration\_cost=349  
checking if image is initramfs... it is  
Freeing initrd memory: 5579k freed  
DMI 2.3 present.  
NET: Registered protocol family 16  
ACPI: bus type pci registered  
PCI: Using configuration type 1  
ACPI: Subsystem revision 20060127  
ACPI: Interpreter enabled  
ACPI: Using IOAPIC for interrupt routing  
ACPI: PCI Root Bridge [PCI0] (0000:00)  
PCI: Probing PCI hardware (bus 00)  
Boot video device is 0000:03:04.0  
ACPI: PCI Interrupt Routing Table [\_SB\_.PCI0.\_PRT]  
ACPI: PCI Interrupt Routing Table [\_SB\_.PCI0.PCI1.\_PRT]  
ACPI: PCI Interrupt Routing Table [\_SB\_.PCI0.GOLA.\_PRT]  
ACPI: PCI Interrupt Routing Table [\_SB\_.PCI0.GOLB.\_PRT]  
ACPI: PCI Interrupt Link [LNKA] (IRQs 3 4 5 6 7 9 10 11 12 14 15) \*0, disabled.  
ACPI: PCI Interrupt Link [LNKB] (IRQs 3 4 5 6 7 9 10 \*11 12 14 15)  
ACPI: PCI Interrupt Link [LNKC] (IRQs 3 4 5 6 7 \*9 10 11 12 14 15)  
ACPI: PCI Interrupt Link [LNKD] (IRQs 3 4 5 6 7 9 \*10 11 12 14 15)  
SCSI subsystem initialized  
usbcore: registered new driver usbfs  
usbcore: registered new driver hub  
PCI: Using ACPI for IRQ routing  
PCI: If a device doesn't work, try "pci=routeirq". If it helps, post a report  
PCI-DMA: Disabling AGP.  
PCI-DMA: aperture base @ c000000 size 65536 KB  
PCI-DMA: using GART IOMMU.  
PCI-DMA: Reserving 64MB of IOMMU area in the AGP aperture  
PCI: Bridge: 0000:00:06.0  
IO window: b000-bfff  
MEM window: fca00000-feafffff  
PREFETCH window: disabled.  
PCI: Bridge: 0000:00:0a.0

x86\_86 SMP megaraid\_mbox hangups and panics.

IO window: disabled.  
MEM window: fc900000–fc9fffff  
PREFETCH window: ff500000–ff5fffff  
PCI: Bridge: 0000:00:0b.0  
IO window: disabled.  
MEM window: fc800000–fc8fffff  
PREFETCH window: ff400000–ff4fffff  
IA32 emulation \$Id: sys\_ia32.c,v 1.32 2002/03/24 13:02:28 ak Exp \$  
VFS: Disk quotas dquot\_6.5.1  
Dquot–cache hash table entries: 512 (order 0, 4096 bytes)  
Installing knfsd (copyright (C) 1996 okir@xxxxxxxxxxxxx).  
fuse init (API version 7.6)  
SGI XFS with ACLs, security attributes, realtime, large block/inode numbers, no debug enabled  
Initializing Cryptographic API  
io scheduler noop registered  
io scheduler anticipatory registered (default)  
io scheduler deadline registered  
io scheduler cfq registered  
PCI: MSI quirk detected. pci\_msi\_quirk set.  
PCI: MSI quirk detected. pci\_msi\_quirk set.  
Real Time Clock Driver v1.12ac  
hw\_random: AMD768 system management I/O registers at 0x5000.  
hw\_random hardware driver 1.0.0 loaded  
Linux apgart interface v0.101 (c) Dave Jones  
serio: i8042 AUX port at 0x60,0x64 irq 12  
serio: i8042 KBD port at 0x60,0x64 irq 1  
Serial: 8250/16550 driver \$Revision: 1.90 \$ 4 ports, IRQ sharing enabled  
serial8250: ttyS0 at I/O 0x3f8 (irq = 4) is a 16550A  
serial8250: ttyS1 at I/O 0x2f8 (irq = 3) is a 16550A  
RAMDISK driver initialized: 16 RAM disks of 65536K size 1024 blocksize  
loop: loaded (max 8 devices)  
nbd: registered device at major 43  
Intel(R) PRO/1000 Network Driver – version 6.3.9–k4–NAPI  
Copyright (c) 1999–2005 Intel Corporation.  
Ethernet Channel Bonding Driver: v3.0.1 (January 9, 2006)  
bonding: Warning: either miimon or arp\_interval and arp\_ip\_target module parameters must be specified,  
otherwise bonding will not detect link failures! see bonding.txt for details.  
tg3.c:v3.49 (Feb 2, 2006)  
GSI 16 sharing vector 0xA9 and IRQ 16  
ACPI: PCI Interrupt 0000:02:05.0[A] –> GSI 26 (level, low) –> IRQ 16  
eth0: Tigon3 [partno(BCM95704A6) rev 2100 PHY(5704)] (PCIX:100MHz:64–bit) 10/100/1000BaseT  
Ethernet 00:30:48:57:3d:4e  
eth0: RXcsums[1] LinkChgREG[0] MIirq[0] ASF[1] Split[0] WireSpeed[1] TSOcap[0] eth0:  
dma\_rwctrl[769f4000] dma\_mask[64–bit]  
GSI 17 sharing vector 0xB1 and IRQ 17  
ACPI: PCI Interrupt 0000:02:05.1[B] –> GSI 27 (level, low) –> IRQ 17  
eth1: Tigon3 [partno(BCM95704A6) rev 2100 PHY(5704)] (PCIX:100MHz:64–bit) 10/100/1000BaseT  
Ethernet 00:30:48:57:3d:4f  
eth1: RXcsums[1] LinkChgREG[0] MIirq[0] ASF[0] Split[0] WireSpeed[1] TSOcap[1] eth1:  
dma\_rwctrl[769f4000] dma\_mask[64–bit]  
tun: Universal TUN/TAP device driver, 1.6

x86\_86 SMP megaraid\_mbox hangups and panics.

tun: (C) 1999–2004 Max Krasnyansky <maxk@xxxxxxxxxxxx>  
Uniform Multi–Platform E–IDE driver Revision: 7.00alpha2  
ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx  
AMD8111: IDE controller at PCI slot 0000:00:07.1  
AMD8111: chipset revision 3  
AMD8111: not 100% native mode: will probe irqs later  
AMD8111: 0000:00:07.1 (rev 03) UDMA133 controller  
ide0: BM–DMA at 0xffa0–0xffa7, BIOS settings: hda:prio, hdb:prio  
ide1: BM–DMA at 0xffa8–0xffaf, BIOS settings: hdc:prio, hdd:prio  
Probing IDE interface ide0...  
Probing IDE interface ide1...  
Probing IDE interface ide0...  
Probing IDE interface ide1...  
ide–floppy driver 0.99.newide  
megaraid cmm: 2.20.2.6 (Release Date: Mon Mar 7 00:01:03 EST 2005)  
megaraid: 2.20.4.7 (Release Date: Mon Nov 14 12:27:22 EST 2005)  
megaraid: probe new device 0x1000:0x1960:0x1000:0x0520: bus 1:slot 3:func 0  
GSI 18 sharing vector 0xB9 and IRQ 18  
ACPI: PCI Interrupt 0000:01:03.0[A] –> GSI 29 (level, low) –> IRQ 18  
megaraid: fw version:[1L37] bios version:[G119]  
scsi0 : LSI Logic MegaRAID driver  
scsi[0]: scanning scsi channel 0 [Phy 0] for non–raid devices  
scsi[0]: scanning scsi channel 1 [virtual] for logical drives  
Vendor: MegaRAID Model: LD0 RAID5 50030R Rev: 1L37  
Type: Direct–Access ANSI SCSI revision: 02  
megasas: 00.00.02.04 Fri Feb 03 14:31:44 PST 2006  
3ware Storage Controller device driver for Linux v1.26.02.001.  
3ware 9000 Storage Controller device driver for Linux v2.26.02.005.  
ipr: IBM Power RAID SCSI Device Driver version: 2.1.2 (February 8, 2006)  
libata version 1.20 loaded.  
SCSI device sda: 716861440 512–byte hdwr sectors (367033 MB)  
sda: Write Protect is off  
sda: Mode Sense: 00 00 00 00  
sda: asking for cache data failed  
sda: assuming drive cache: write through  
SCSI device sda: 716861440 512–byte hdwr sectors (367033 MB)  
sda: Write Protect is off  
sda: Mode Sense: 00 00 00 00  
sda: asking for cache data failed  
sda: assuming drive cache: write through  
sda: sda1 sda2 sda3 sda4 < sda5 >  
sd 0:1:0:0: Attached scsi disk sda  
sd 0:1:0:0: Attached scsi generic sg0 type 0  
[junk cut from this point]

–  
To unsubscribe from this list: send the line "unsubscribe linux–kernel" in  
the body of a message to majordomo@xxxxxxxxxxxxxxx  
More majordomo info at <http://vger.kernel.org/majordomo–info.html>  
Please read the FAQ at <http://www.tux.org/lkml/>