
Subject: *SOLVED* Seg fault with quotaugidlimit
Posted by [MaFL](#) on Fri, 15 Jun 2007 13:23:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

I'm getting seg faults during vps start if I activate the internal quota.
Not 100% sure when it started, but it was not with the development kernels of 2.6.18.

Tested on 2 different hosts.
Filesystem is Reiser.

Vps starts and the quota inside seems to work also.

vzctl start 10

```
Starting VE ...
Running: /usr/sbin/vzquota show 10
Running: /usr/sbin/vzquota on 10 -r 0 -b 10240100 -B 11264100 -i 100100 -l 110100 -e 604800 -n
604800 -s 1 -u 200
Mounting root: /data1/vz/root/10 /data1/vz/private/10
VE is mounted
Set iptables mask 0x000017bf
Set features mask 0000000000000000/0000000000000000
Adding IP address(es): ...
Running: /usr/lib/vzctl/scripts/vps-net_add
Running VE script: /etc/vz/dists/scripts/debian-add_ip.sh
Setting CPU units: 20000
Configure meminfo: 256000
Set hostname: ...
Running VE script: /etc/vz/dists/scripts/debian-set_hostname.sh
Running VE script: /etc/vz/dists/scripts/set_dns.sh
File resolv.conf was modified
Setting quota ugidlimit: 200
Running VE script: /etc/vz/dists/scripts/set_ugid_quota.sh
Running: /usr/sbin/vzquota stat 10 -f
Running: /usr/sbin/vzquota stat 10 -f -t
Segmentation fault
```

"vzquota stat 10 -f -t" from the commandline works.

Any ideas?

Tnx
Matt

Subject: Re: Seg fault with quotauidlimit
Posted by [rickb](#) on Sat, 16 Jun 2007 00:56:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi, I was using reiser+openvz last year, as I know its a superior filesystem from personal experience. But, I ran into problems like this which were unsupported by openvz since it pushes ext3. Since switching to ext3, I haven't had any quota or filesystem problems.

So, I would say either move to ext3 or start a filesystem revolution and get openvz moving in the direction that will support it.

Rick

Subject: Re: Seg fault with quotauidlimit
Posted by [nksupport](#) on Sat, 16 Jun 2007 06:37:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Do /proc/user_beancounters look OK for that VE?

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Sat, 16 Jun 2007 17:38:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Rick, we do not officially support reiserfs for the only reason - stability/maintaiablity. However, vzquota should work fine with reiser and AFAIR we were fixing all the bugs people reported.

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Sat, 16 Jun 2007 17:40:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

can you please do the following and report the result:

start your VE and do the following:

```
# strace -f -o out /usr/sbin/vzquota stat 10 -f -t  
# dmesg >> out
```

and show me the out file plz

Subject: Re: Seg fault with quotauidlimit
Posted by [MaFL](#) on Sat, 16 Jun 2007 19:09:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

UBCs are fine.
It's not vps specific.

```
strace -f -o out /usr/sbin/vzquota stat 10 -f -t
resource      usage      softlimit   hardlimit   grace
1k-blocks     839273    10240100   11264100
inodes        39257     100100     110100
```

User/group quota: -,active
Ugids: loaded 30, total 30, limit 200
Ugid limit was exceeded: no

User/group grace times and quotafile flags:
type block_exp_time inode_exp_time dqf_flags
user 0h
group 0h

User/group objects:

| ID | type | resource | usage | softlimit | hardlimit | grace | status |
|----|-------|-----------|--------|-----------|-----------|--------|--------|
| 0 | user | 1k-blocks | 766109 | 0 | 0 | loaded | |
| 0 | user | inodes | 36114 | 0 | 0 | loaded | |
| 0 | group | 1k-blocks | 748558 | 0 | 0 | loaded | |
| 0 | group | inodes | 34784 | 0 | 0 | loaded | |
| 1 | user | 1k-blocks | 5 | 0 | 0 | loaded | |
| 1 | user | inodes | 3 | 0 | 0 | loaded | |
| 1 | group | 1k-blocks | 5 | 0 | 0 | loaded | |
| 1 | group | inodes | 3 | 0 | 0 | loaded | |
| 2 | user | 1k-blocks | 1733 | 0 | 0 | loaded | |
| 2 | user | inodes | 45 | 0 | 0 | loaded | |
| 2 | group | 1k-blocks | 2539 | 0 | 0 | loaded | |
| 2 | group | inodes | 56 | 0 | 0 | loaded | |
| 4 | group | 1k-blocks | 982 | 0 | 0 | loaded | |
| 4 | group | inodes | 4 | 0 | 0 | loaded | |
| 5 | group | 1k-blocks | 21 | 0 | 0 | loaded | |
| 5 | group | inodes | 34 | 0 | 0 | loaded | |
| 6 | user | 1k-blocks | 608 | 0 | 0 | loaded | |
| 6 | user | inodes | 28 | 0 | 0 | loaded | |
| 8 | user | 1k-blocks | 50 | 0 | 0 | loaded | |
| 8 | user | inodes | 15 | 0 | 0 | loaded | |
| 8 | group | 1k-blocks | 62 | 0 | 0 | loaded | |
| 8 | group | inodes | 14 | 0 | 0 | loaded | |
| 9 | user | 1k-blocks | 1 | 0 | 0 | loaded | |
| 9 | user | inodes | 1 | 0 | 0 | loaded | |
| 9 | group | 1k-blocks | 1 | 0 | 0 | loaded | |
| 9 | group | inodes | 4 | 0 | 0 | loaded | |
| 40 | group | 1k-blocks | 123 | 0 | 0 | loaded | |

| | | | | | | |
|------|-------|-----------|-------|---|---|--------|
| 40 | group | inodes | 36 | 0 | 0 | loaded |
| 42 | group | 1k-blocks | 73 | 0 | 0 | loaded |
| 42 | group | inodes | 6 | 0 | 0 | loaded |
| 43 | group | 1k-blocks | 401 | 0 | 0 | loaded |
| 43 | group | inodes | 4 | 0 | 0 | loaded |
| 50 | group | 1k-blocks | 23373 | 0 | 0 | loaded |
| 50 | group | inodes | 1902 | 0 | 0 | loaded |
| 100 | group | 1k-blocks | 20026 | 0 | 0 | loaded |
| 100 | group | inodes | 1542 | 0 | 0 | loaded |
| 101 | group | 1k-blocks | 33 | 0 | 0 | loaded |
| 101 | group | inodes | 3 | 0 | 0 | loaded |
| 103 | user | 1k-blocks | 20026 | 0 | 0 | loaded |
| 103 | user | inodes | 1542 | 0 | 0 | loaded |
| 103 | group | 1k-blocks | 9907 | 0 | 0 | loaded |
| 103 | group | inodes | 6 | 0 | 0 | loaded |
| 104 | user | 1k-blocks | 7591 | 0 | 0 | loaded |
| 104 | user | inodes | 596 | 0 | 0 | loaded |
| 104 | group | 1k-blocks | 3432 | 0 | 0 | loaded |
| 104 | group | inodes | 483 | 0 | 0 | loaded |
| 105 | group | 1k-blocks | 29666 | 0 | 0 | loaded |
| 105 | group | inodes | 363 | 0 | 0 | loaded |
| 106 | group | 1k-blocks | 1 | 0 | 0 | loaded |
| 106 | group | inodes | 1 | 0 | 0 | loaded |
| 500 | user | 1k-blocks | 81 | 0 | 0 | loaded |
| 500 | user | inodes | 12 | 0 | 0 | loaded |
| 500 | group | 1k-blocks | 81 | 0 | 0 | loaded |
| 500 | group | inodes | 12 | 0 | 0 | loaded |
| 1000 | user | 1k-blocks | 9847 | 0 | 0 | loaded |
| 1000 | user | inodes | 5 | 0 | 0 | loaded |
| 1001 | user | 1k-blocks | 3623 | 0 | 0 | loaded |
| 1001 | user | inodes | 517 | 0 | 0 | loaded |
| 1002 | user | 1k-blocks | 29604 | 0 | 0 | loaded |
| 1002 | user | inodes | 379 | 0 | 0 | loaded |

dmesg has no failure

Linux version 2.6.18-028stab035.1-ovz-smp (tsd@debian.systs.org) (gcc version 4.1.2 20061115 (prerelease) (Debian 4.1.1-21)) #1 SMP Wed Jun 13 22:08:06 CEST 2007

BIOS-provided physical RAM map:

BIOS-e820: 0000000000000000 - 000000000009f400 (usable)
 BIOS-e820: 000000000009f400 - 00000000000a0000 (reserved)
 BIOS-e820: 00000000000e4000 - 0000000000100000 (reserved)
 BIOS-e820: 0000000000100000 - 000000007fff0000 (usable)
 BIOS-e820: 000000007fff0000 - 000000007ffff000 (ACPI data)
 BIOS-e820: 000000007ffff000 - 0000000080000000 (ACPI NVS)
 BIOS-e820: 00000000fee00000 - 00000000fee01000 (reserved)
 BIOS-e820: 00000000ffb00000 - 0000000100000000 (reserved)

1151MB HIGHMEM available.

896MB LOWMEM available.

found SMP MP-table at 000ff780
On node 0 totalpages: 524272
DMA zone: 4096 pages, LIFO batch:0
Normal zone: 225280 pages, LIFO batch:31
HighMem zone: 294896 pages, LIFO batch:31
DMI 2.3 present.
ACPI: RSDP (v000 ACPIAM) @ 0x000f87e0
ACPI: RSDT (v001 A M I OEMRSDT 0x01000626 MSFT 0x00000097) @ 0x7fff0000
ACPI: FADT (v002 A M I OEMFACP 0x01000626 MSFT 0x00000097) @ 0x7fff0200
ACPI: MADT (v001 A M I OEMAPIC 0x01000626 MSFT 0x00000097) @ 0x7fff0390
ACPI: OEMB (v001 A M I AMI_OEM 0x01000626 MSFT 0x00000097) @ 0x7fff040
ACPI: DSDT (v001 STP1A STP1A042 0x00000042 INTL 0x02002026) @ 0x00000000
ACPI: PM-Timer IO Port: 0x808
ACPI: Local APIC address 0xfe00000
ACPI: LAPIC (acpi_id[0x01] lapic_id[0x00] enabled)
Processor #0 15:2 APIC version 20
ACPI: LAPIC (acpi_id[0x02] lapic_id[0x01] enabled)
Processor #1 15:2 APIC version 20
ACPI: LAPIC_NMI (acpi_id[0x01] dfl dfl lint[0x1])
ACPI: LAPIC_NMI (acpi_id[0x02] dfl dfl lint[0x1])
ACPI: IOAPIC (id[0x02] address[0xfec00000] gsi_base[0])
IOAPIC[0]: apic_id 2, version 32, address 0xfec00000, GSI 0-23
ACPI: IOAPIC (id[0x03] address[0xfec10000] gsi_base[24])
IOAPIC[1]: apic_id 3, version 32, address 0xfec10000, GSI 24-47
ACPI: INT_SRC_OVR (bus 0 bus_irq 0 global_irq 2 dfl dfl)
ACPI: INT_SRC_OVR (bus 0 bus_irq 10 global_irq 10 high level)
ACPI: IRQ0 used by override.
ACPI: IRQ2 used by override.
ACPI: IRQ10 used by override.
Enabling APIC mode: Flat. Using 2 I/O APICs
Using ACPI (MADT) for SMP configuration information
Allocating PCI resources starting at 88000000 (gap: 80000000:7ee00000)
Detected 2793.182 MHz processor.
Built 1 zonelists. Total pages: 524272
Kernel command line: root=/dev/md0 ro
mapped APIC to fffd000 (fee00000)
mapped IOAPIC to fffc000 (fec00000)
mapped IOAPIC to fffb000 (fec10000)
Enabling fast FPU save and restore... done.
Enabling unmasked SIMD FPU exception support... done.
Initializing CPU#0
CPU 0 irqstacks, hard=c05f3000 soft=c05eb000
PID hash table entries: 4096 (order: 12, 16384 bytes)
Console: colour VGA+ 80x25
Dentry cache hash table entries: 131072 (order: 7, 524288 bytes)
Inode-cache hash table entries: 65536 (order: 6, 262144 bytes)
Memory: 2067524k/2097088k available (3604k kernel code, 28164k reserved, 1088k data, 296k
init, 1179584k highmem)

Checking if this processor honours the WP bit even in supervisor mode... Ok.
Calibrating delay using timer specific routine.. 5588.38 BogoMIPS (lpj=2794191)
Mount-cache hash table entries: 512
CPU: After generic identify, caps: bfebfbff 00000000 00000000 00000000 00004400 00000000
00000000
CPU: After vendor identify, caps: bfebfbff 00000000 00000000 00000000 00004400 00000000
00000000
CPU: Trace cache: 12K uops, L1 D cache: 8K
CPU: L2 cache: 512K
CPU: Physical Processor ID: 0
CPU: After all inits, caps: bfebfbff 00000000 00000000 00000080 00004400 00000000 00000000
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#0.
CPU0: Intel P4/Xeon Extended MCE MSR (12) available
CPU0: Thermal monitoring enabled
Compat vDSO mapped to fffe000.
Checking 'hlt' instruction... OK.
Freeing SMP alternatives: 20k freed
ACPI: Core revision 20060707
Page beancounter hash is 131072 entries.
CPU0: Intel(R) Pentium(R) 4 CPU 2.80GHz stepping 05
Booting processor 1/1 eip 2000
CPU 1 irqstacks, hard=c05f4000 soft=c05ec000
Initializing CPU#1
Calibrating delay using timer specific routine.. 5585.25 BogoMIPS (lpj=2792628)
CPU: After generic identify, caps: bfebfbff 00000000 00000000 00000000 00004400 00000000
00000000
CPU: After vendor identify, caps: bfebfbff 00000000 00000000 00000000 00004400 00000000
00000000
CPU: Trace cache: 12K uops, L1 D cache: 8K
CPU: L2 cache: 512K
CPU: Physical Processor ID: 0
CPU: After all inits, caps: bfebfbff 00000000 00000000 00000080 00004400 00000000 00000000
Intel machine check architecture supported.
Intel machine check reporting enabled on CPU#1.
CPU1: Intel P4/Xeon Extended MCE MSR (12) available
CPU1: Thermal monitoring enabled
CPU1: Intel(R) Pentium(R) 4 CPU 2.80GHz stepping 05
Total of 2 processors activated (11173.63 BogoMIPS).
ENABLING IO-APIC IRQs
..TIMER: vector=0x31 apic1=0 pin1=2 apic2=-1 pin2=-1
checking TSC synchronization across 2 CPUs: passed.
Brought up 2 CPUs
migration_cost=159
checking if image is initramfs... it is
Freeing initrd memory: 3265k freed
NET: Registered protocol family 16
ACPI: bus type pci registered

PCI: PCI BIOS revision 2.10 entry at 0xf0031, last bus=3
PCI: Using configuration type 1
Setting up standard PCI resources
ACPI: Interpreter enabled
ACPI: Using IOAPIC for interrupt routing
ACPI: PCI Root Bridge [PCI0] (0000:00)
PCI: Probing PCI hardware (bus 00)
PCI quirk: region 0800-087f claimed by ICH4 ACPI/GPIO/TCO
PCI quirk: region 0480-04bf claimed by ICH4 GPIO
PCI: Ignoring BAR0-3 of IDE controller 0000:00:1f.1
Boot video device is 0000:03:00.0
PCI: Firmware left 0000:03:01.0 e100 interrupts enabled, disabling
PCI: Transparent bridge - 0000:00:1e.0
ACPI: PCI Interrupt Routing Table [_SB_.PCI0._PRT]
ACPI: PCI Interrupt Routing Table [_SB_.PCI0.P0P5._PRT]
ACPI: PCI Interrupt Routing Table [_SB_.PCI0.P0P4._PRT]
ACPI: PCI Interrupt Link [LNKA] (IRQs 3 4 *5 6 9 11 12 14 15)
ACPI: PCI Interrupt Link [LNKB] (IRQs 3 4 5 6 9 *11 12 14 15)
ACPI: PCI Interrupt Link [LNKC] (IRQs 3 4 5 6 *9 11 12 14 15)
ACPI: PCI Interrupt Link [LNKD] (IRQs 3 4 5 6 *9 11 12 14 15)
ACPI: PCI Interrupt Link [LNKE] (IRQs 3 4 5 6 9 11 12 14 15) *0, disabled.
ACPI: PCI Interrupt Link [LNKF] (IRQs 3 4 5 6 9 *11 12 14 15)
ACPI: PCI Interrupt Link [LNKG] (IRQs 3 4 5 6 9 11 12 14 15) *0, disabled.
ACPI: PCI Interrupt Link [LNKH] (IRQs *7)
SCSI subsystem initialized
PCI: Using ACPI for IRQ routing
PCI: If a device doesn't work, try "pci=routeirq". If it helps, post a report
PCI: Bridge: 0000:00:03.0
IO window: a000-afff
MEM window: fc300000-fc3fffff
PREFETCH window: disabled.
PCI: Bridge: 0000:00:1c.0
IO window: b000-bfff
MEM window: fc400000-fc5fffff
PREFETCH window: fc100000-fc1fffff
PCI: Bridge: 0000:00:1e.0
IO window: c000-cfff
MEM window: fc600000-fe6fffff
PREFETCH window: 88000000-880fffff
PCI: Setting latency timer of device 0000:00:1e.0 to 64
NET: Registered protocol family 2
IP route cache hash table entries: 32768 (order: 5, 131072 bytes)
TCP established hash table entries: 131072 (order: 8, 1048576 bytes)
TCP bind hash table entries: 65536 (order: 7, 524288 bytes)
TCP: Hash tables configured (established 131072 bind 65536)
TCP reno registered
highmem bounce pool size: 64 pages
VFS: Disk quotas dquot_6.5.1

Dquot-cache hash table entries: 1024 (order 0, 4096 bytes)
Initializing Cryptographic API
io scheduler noop registered
io scheduler anticipatory registered
io scheduler deadline registered
io scheduler cfq registered (default)
pci_hotplug: PCI Hot Plug PCI Core version: 0.5
Real Time Clock Driver v1.12ac
Serial: 8250/16550 driver \$Revision: 1.90 \$ 4 ports, IRQ sharing disabled
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 16384K size 1024 blocksize
Compaq SMART2 Driver (v 2.6.0)
HP CISS Driver (v 3.6.10)
Uniform Multi-Platform E-IDE driver Revision: 7.00alpha2
ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx
ICH5: IDE controller at PCI slot 0000:00:1f.1
ACPI: PCI Interrupt 0000:00:1f.1[A] -> GSI 18 (level, low) -> IRQ 16
ICH5: chipset revision 2
ICH5: not 100% native mode: will probe irqs later
 ide0: BM-DMA at 0xfc00-0xfc07, BIOS settings: hda:prio, hdb:prio
 ide1: BM-DMA at 0xfc08-0xfc0f, BIOS settings: hdc:prio, hdd:prio
Probing IDE interface ide0...
Probing IDE interface ide1...
Probing IDE interface ide0...
Probing IDE interface ide1...
Loading iSCSI transport class v1.1-646.<6>ACPI: PCI Interrupt 0000:02:01.0[A] -> GSI 27 (level, low) -> IRQ 17
scsi0 : Adaptec AIC79XX PCI-X SCSI HBA DRIVER, Rev 3.0
 <Adaptec AIC7901 Ultra320 SCSI adapter>
 aic7901: Ultra320 Wide Channel A, SCSI Id=7, PCI-X 50-66Mhz, 512 SCBs

 Vendor: FUJITSU Model: MAP3735NC Rev: 0108
 Type: Direct-Access ANSI SCSI revision: 03
target0:0:0: asynchronous
scsi0:A:0:0: Tagged Queuing enabled. Depth 32
target0:0:0: Beginning Domain Validation
target0:0:0: wide asynchronous
target0:0:0: FAST-160 WIDE SCSI 320.0 MB/s DT IU QAS RTI PCOMP (6.25 ns, offset 127)
target0:0:0: Ending Domain Validation
 Vendor: FUJITSU Model: MAP3735NC Rev: 0107
 Type: Direct-Access ANSI SCSI revision: 03
target0:0:1: asynchronous
scsi0:A:1:0: Tagged Queuing enabled. Depth 32
target0:0:1: Beginning Domain Validation
target0:0:1: wide asynchronous
target0:0:1: FAST-160 WIDE SCSI 320.0 MB/s DT IU QAS RTI PCOMP (6.25 ns, offset 127)
target0:0:1: Ending Domain Validation

Adaptec aacraid driver (1.1-5[2409]-mh2)
QLogic Fibre Channel HBA Driver
Emulex LightPulse Fibre Channel SCSI driver 8.1.9
Copyright(c) 2004-2006 Emulex. All rights reserved.
megaraid cmm: 2.20.2.7 (Release Date: Sun Jul 16 00:01:03 EST 2006)
megaraid: 2.20.4.9 (Release Date: Sun Jul 16 12:27:22 EST 2006)
megasas: 00.00.03.01 Sun May 14 22:49:52 PDT 2006
GDT-HA: Storage RAID Controller Driver. Version: 3.05
GDT-HA: Found 0 PCI Storage RAID Controllers
3ware Storage Controller device driver for Linux v1.26.02.001.
3ware 9000 Storage Controller device driver for Linux v2.26.02.007.
libata version 2.00 loaded.
ata_piix 0000:00:1f.2: version 2.00
ata_piix 0000:00:1f.2: MAP [P0 -- P1 --]
ACPI: PCI Interrupt 0000:00:1f.2[A] -> GSI 18 (level, low) -> IRQ 16
PCI: Setting latency timer of device 0000:00:1f.2 to 64
ata1: SATA max UDMA/133 cmd 0xE400 ctl 0xE002 bmdma 0xD400 irq 16
ata2: SATA max UDMA/133 cmd 0xDC00 ctl 0xD802 bmdma 0xD408 irq 16
scsi1 : ata_piix
ATA: abnormal status 0x7F on port 0xE407
scsi2 : ata_piix
ata2.00: ATA-8, max UDMA7, 976773168 sectors: LBA48 NCQ (depth 0/32)
ata2.00: ata2: dev 0 multi count 16
ata2.00: configured for UDMA/133
Vendor: ATA Model: SAMSUNG HD501LJ Rev: CR10
Type: Direct-Access ANSI SCSI revision: 05
SCSI device sda: 143571316 512-byte hdwr sectors (73509 MB)
sda: Write Protect is off
sda: Mode Sense: b3 00 00 08
SCSI device sda: drive cache: write back
SCSI device sda: 143571316 512-byte hdwr sectors (73509 MB)
sda: Write Protect is off
sda: Mode Sense: b3 00 00 08
SCSI device sda: drive cache: write back
sda: sda1 sda2 sda3
sd 0:0:0:0: Attached scsi disk sda
SCSI device sdb: 143571316 512-byte hdwr sectors (73509 MB)
sdb: Write Protect is off
sdb: Mode Sense: b3 00 00 08
SCSI device sdb: drive cache: write back
SCSI device sdb: 143571316 512-byte hdwr sectors (73509 MB)
sdb: Write Protect is off
sdb: Mode Sense: b3 00 00 08
SCSI device sdb: drive cache: write back
sdb: sdb1 sdb2 sdb3
sd 0:0:1:0: Attached scsi disk sdb
SCSI device sdc: 976773168 512-byte hdwr sectors (500108 MB)
sdc: Write Protect is off

sdc: Mode Sense: 00 3a 00 00
SCSI device sdc: drive cache: write back
SCSI device sdc: 976773168 512-byte hdwr sectors (500108 MB)
sdc: Write Protect is off
sdc: Mode Sense: 00 3a 00 00
SCSI device sdc: drive cache: write back
sdc: sdc1
sd 2:0:0:0: Attached scsi disk sdc
serio: i8042 AUX port at 0x60,0x64 irq 12
serio: i8042 KBD port at 0x60,0x64 irq 1
mice: PS/2 mouse device common for all mice
md: linear personality registered for level -1
md: raid0 personality registered for level 0
md: raid1 personality registered for level 1
md: raid10 personality registered for level 10
raid6: int32x1 761 MB/s
raid6: int32x2 824 MB/s
raid6: int32x4 750 MB/s
raid6: int32x8 531 MB/s
raid6: mmxx1 2242 MB/s
raid6: mmxx2 2863 MB/s
raid6: sse1x1 1261 MB/s
raid6: sse1x2 2496 MB/s
raid6: sse2x1 2140 MB/s
raid6: sse2x2 3152 MB/s
raid6: using algorithm sse2x2 (3152 MB/s)
md: raid6 personality registered for level 6
md: raid5 personality registered for level 5
md: raid4 personality registered for level 4
raid5: automatically using best checksumming function: pIII_sse
pIII_sse : 3624.000 MB/sec
raid5: using function: pIII_sse (3624.000 MB/sec)
md: multipath personality registered for level -4
md: md driver 0.90.3 MAX_MD_DEVS=256, MD_SB_DISKS=27
md: bitmap version 4.39
device-mapper: ioctl: 4.7.0-ioctl (2006-06-24) initialised: dm-devel@redhat.com
device-mapper: multipath: version 1.0.4 loaded
device-mapper: multipath round-robin: version 1.0.0 loaded
device-mapper: multipath emc: version 0.0.3 loaded
TCP bic registered
NET: Registered protocol family 1
NET: Registered protocol family 10
IPv6 over IPv4 tunneling driver
Starting balanced_irq
Using IPI Shortcut mode
Freeing unused kernel memory: 296k freed
Time: tsc clocksource has been installed.
ACPI: Processor [CPU1] (supports 8 throttling states)

usbcore: registered new driver usbfs
usbcore: registered new driver hub
USB Universal Host Controller Interface driver v3.0
ACPI: PCI Interrupt 0000:00:1d.0[A] -> GSI 16 (level, low) -> IRQ 18
PCI: Setting latency timer of device 0000:00:1d.0 to 64
uhci_hcd 0000:00:1d.0: UHCI Host Controller
uhci_hcd 0000:00:1d.0: new USB bus registered, assigned bus number 1
uhci_hcd 0000:00:1d.0: irq 18, io base 0x0000e800
usb usb1: configuration #1 chosen from 1 choice
hub 1-0:1.0: USB hub found
hub 1-0:1.0: 2 ports detected
Intel(R) PRO/1000 Network Driver - version 7.1.9-k4
Copyright (c) 1999-2006 Intel Corporation.
e100: Intel(R) PRO/100 Network Driver, 3.5.10-k2-NAPI
e100: Copyright(c) 1999-2005 Intel Corporation
ACPI: PCI Interrupt 0000:00:1d.1[B] -> GSI 19 (level, low) -> IRQ 19
PCI: Setting latency timer of device 0000:00:1d.1 to 64
uhci_hcd 0000:00:1d.1: UHCI Host Controller
uhci_hcd 0000:00:1d.1: new USB bus registered, assigned bus number 2
uhci_hcd 0000:00:1d.1: irq 19, io base 0x0000ec00
usb usb2: configuration #1 chosen from 1 choice
hub 2-0:1.0: USB hub found
hub 2-0:1.0: 2 ports detected
ACPI: PCI Interrupt 0000:01:01.0[A] -> GSI 18 (level, low) -> IRQ 16
PCI: Setting latency timer of device 0000:01:01.0 to 64
e1000: 0000:01:01.0: e1000_probe: (PCI:33MHz:32-bit) 00:04:23:c7:48:c8
e1000: eth0: e1000_probe: Intel(R) PRO/1000 Network Connection
ACPI: PCI Interrupt 0000:03:01.0[A] -> GSI 17 (level, low) -> IRQ 20
ACPI: PCI Interrupt 0000:00:1d.7[D] -> GSI 23 (level, low) -> IRQ 21
PCI: Setting latency timer of device 0000:00:1d.7 to 64
ehci_hcd 0000:00:1d.7: EHCI Host Controller
ehci_hcd 0000:00:1d.7: new USB bus registered, assigned bus number 3
ehci_hcd 0000:00:1d.7: debug port 1
PCI: cache line size of 128 is not supported by device 0000:00:1d.7
ehci_hcd 0000:00:1d.7: irq 21, io mem 0xfe7ffc00
ehci_hcd 0000:00:1d.7: USB 2.0 started, EHCI 1.00, driver 10 Dec 2004
usb usb3: configuration #1 chosen from 1 choice
hub 3-0:1.0: USB hub found
hub 3-0:1.0: 4 ports detected
e100: eth1: e100_probe: addr 0xfe6fe000, irq 20, MAC addr 00:04:23:C7:48:C9
usb 2-2: new full speed USB device using uhci_hcd and address 2
md: md0 stopped.
md: bind<sdb1>
md: bind<sda1>
raid1: raid set md0 active with 2 out of 2 mirrors
md: md1 stopped.
md: bind<sdb2>
md: bind<sda2>

raid1: raid set md1 active with 2 out of 2 mirrors
usb 2-2: new full speed USB device using uhci_hcd and address 3
Attempting manual resume
kjournald starting. Commit interval 5 seconds
EXT3-fs: mounted filesystem with ordered data mode.
usb 2-2: configuration #1 chosen from 1 choice
shpchp: Standard Hot Plug PCI Controller Driver version: 0.4
usbcore: registered new driver hiddev
input: Peppercon AG MultiDevice as /class/input/input0
input: USB HID v1.01 Keyboard [Peppercon AG MultiDevice] on usb-0000:00:1d.1-2
input: Peppercon AG MultiDevice as /class/input/input1
input: USB HID v1.01 Mouse [Peppercon AG MultiDevice] on usb-0000:00:1d.1-2
usbcore: registered new driver usbhid
drivers/usb/input/hid-core.c: v2.6:USB HID core driver
ACPI: PCI Interrupt 0000:00:1f.3[B] -> GSI 17 (level, low) -> IRQ 20
Adding 979956k swap on /dev/sda3. Priority:-1 extents:1 across:979956k
Adding 979956k swap on /dev/sdb3. Priority:-2 extents:1 across:979956k
EXT3 FS on md0, internal journal
loop: loaded (max 8 devices)
ReiserFS: dm-0: found reiserfs format "3.6" with standard journal
ReiserFS: dm-0: using writeback data mode
ReiserFS: dm-0: journal params: device dm-0, size 8192, journal first block 18, max trans len 1024, max batch 900, max commit age 30, max trans age 30
ReiserFS: dm-0: checking transaction log (dm-0)
ReiserFS: dm-0: Using r5 hash to sort names
SGI XFS with ACLs, large block numbers, no debug enabled
SGI XFS Quota Management subsystem
XFS mounting filesystem sdc1
Ending clean XFS mount for filesystem: sdc1
e1000: eth0: e1000_watchdog: NIC Link is Up 100 Mbps Full Duplex
tun: Universal TUN/TAP device driver, 1.6
tun: (C) 1999-2004 Max Krasnyansky <maxk@qualcomm.com>
ACPI: Power Button (FF) [PWRF]
ACPI: Power Button (CM) [PWRB]
ACPI: Sleep Button (CM) [SLPB]
Installing knfsd (copyright (C) 1996 okir@monad.swb.de).
ip_tables: (C) 2000-2006 Netfilter Core Team
eth0: no IPv6 routers present
NET: Registered protocol family 17
ip_contrack version 2.4 (8192 buckets, 65536 max) - 208 bytes per contrack
VE: 52: started
VE: 53: started
VE: 104: started
VE: 10: started
VE: 10: stopped
VE: 10: started
VE: 10: stopped
VE: 10: started

VE: 53: stopped
VE: 63: started
VE: 53: started
VE: 10: stopped
VE: 10: started
VE: 10: stopped
VE: 10: started

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Mon, 18 Jun 2007 08:54:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you please apply the following patch:

```
--- ./arch/i386/mm/fault.c.ve5555    2007-06-08 17:29:35.000000000 +0400
+++ ./arch/i386/mm/fault.c    2007-06-18 12:53:50.000000000 +0400
@@ -475,6 +475,10 @@ bad_area_nosemaphore:
         if (is_prefetch(regs, address, error_code))
             return;

+         printk( "%s[%d]: segfault at %08lx rip %08lx rsp %08lx error %lx\n",
+                 tsk->comm, tsk->pid, address, regs->eip,
+                 regs->esp, error_code);
+
         tsk->thread.cr2 = address;
         /* Kernel addresses are always protection faults */
         tsk->thread.error_code = error_code | (address >= TASK_SIZE);
```

to your kernel, recompile it and report the output after Segmentation fault. My guess is that it is not vzquota tool at all. Anyway, we need an address where it (or something else crashes)

Subject: Re: Seg fault with quotauidlimit
Posted by [MaFL](#) on Wed, 20 Jun 2007 23:04:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tnx for your effort!

vzctl[10838]: segfault at 000000a8 rip b7f1525c rsp bfa5f63c error 4

Subject: Re: Seg fault with quotaugidlimit
Posted by [dev](#) on Thu, 21 Jun 2007 06:29:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

ok. so now we know that it was vzctl actually, not vzquota.

We need to find out what ip address b7f1525c corresponds to.

1. is printed address always the same?
2. what vzctl version do you use? compiled yourself or from openvz.org?
3. I need all vzctl binaries

Subject: Re: Seg fault with quotaugidlimit
Posted by [falcon](#) on Sun, 24 Jun 2007 11:18:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

I've the same problem with Second-Level quota:

```
Starting VE ...  
VE is mounted  
Setting CPU units: 1000  
Set hostname: test  
Setting quota ugidlimit: 500  
Segmentation fault
```

I am using Gentoo. Kernel is from sys-kernel/openvz-sources testing (2.6.18-028stab035). vzctl is 3.0.16, also from gentoo portage, filesystem is ext3. Quotas don't work inside the VE, but since this is the first time I am using second-level quotas, I don't know if this is my fault:

```
test:~# quotacheck -fvagum  
quotacheck: Scanning /dev/simfs [/] done  
quotacheck: Checked 824 directories and 7538 files  
quotacheck: Cannot turn user quotas off on /dev/simfs: Invalid argument  
Kernel won't know about changes quotacheck did.  
quotacheck: Cannot turn group quotas off on /dev/simfs: Invalid argument  
Kernel won't know about changes quotacheck did.  
test:~# exit
```

Perhaps I will apply your patch, migrate the running VEs and tell you the addresses in the dmesg tomorrow, but I don't know wheter I will have enough time.

Which binaries do you exactly need? I attached a list with the files (and directories) the sys-cluster/vzctl ebuild installed.
I can also create a tarball ("binary package") which includes all files of this ebuild.

Thanks,
falcon

File Attachments

1) [files](#), downloaded 360 times

Subject: Re: Seg fault with quotaugidlimit
Posted by [MaFL](#) on Sun, 24 Jun 2007 18:31:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote:We need to find out what ip address b7f1525c corresponds to.

1. is printed address always the same?
2. what vzctl version do you use? compiled yourself or from openvz.org?
3. I need all vzctl binaries

1. yes
2. vzctl version 3.0.16-5dso1 from debian.systs.org
3. attached vzctl and libs

File Attachments

1) [vzctl.tar.gz](#), downloaded 337 times

Subject: Re: Seg fault with quotaugidlimit
Posted by [dev](#) on Tue, 26 Jun 2007 12:53:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

argh... vzctl libraries are stipped :/

1. is it possible to get an access to the node to try out debugging with gdb?
2. if not, can you please run vzctl command under gdb as:
./gdb --args vzctl start 10
then, gdb should catch vzctl when it receives SIGSEGV
and type in gdb then:
disassemble \$eip

Thanks in advance!

Subject: Re: Seg fault with quotaugidlimit
Posted by [MaFL](#) on Tue, 26 Jun 2007 18:26:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

...
File resolv.conf was modified
Setting quota ugidlimit: 200

Running VE script: /etc/vz/dists/scripts/set_ugid_quota.sh
Running: /usr/sbin/vzquota stat 10 -f
Running: /usr/sbin/vzquota stat 10 -f -t

Program received signal SIGSEGV, Segmentation fault.
0xb7f0425c in quota_set@plt () from /usr/lib/libvzctl-0.0.2.so
(gdb) disassemble \$eip
Dump of assembler code for function quota_set@plt:
0xb7f0425c <quota_set@plt+0>: jmp *0xa8(%ebx)
0xb7f04262 <quota_set@plt+6>: push \$0x138
0xb7f04267 <quota_set@plt+11>: jmp 0xb7f03fdc <_init+24>
End of assembler dump.

If you still need access I have to check how can I realize it...

Subject: Re: Seg fault with quotaugidlimit
Posted by [dev](#) on Wed, 27 Jun 2007 07:52:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

this looks damn strange :/ since SIGSEGV happens at plt@ relocation, which should not happen normally at all...

1. do you use vzctl 3.0.16-5dso1 from debian.systs.org as well?
 2. does the same bug happens if vzctl 3.0.17 from git.openvz.org or .src.tgz from download.openvz.org is compiled and installed?
 3. yes, looks like I need an access :/
-

Subject: Re: Seg fault with quotaugidlimit
Posted by [MaFL](#) on Wed, 27 Jun 2007 14:29:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

1. yes.. I'm running vzutils from debian.systs.org on 2 boxes, both have the problem (falcon has the segfault with gentoo)

2. vzctl 3.0.17 from gid
Program received signal SIGSEGV, Segmentation fault.
0xb7f2325c in quota_set@plt () from /usr/local/lib/libvzctl-0.0.2.so
(gdb) disassemble \$eip
Dump of assembler code for function quota_set@plt:
0xb7f2325c <quota_set@plt+0>: jmp *0xa8(%ebx)
0xb7f23262 <quota_set@plt+6>: push \$0x138
0xb7f23267 <quota_set@plt+11>: jmp 0xb7f22fdc <_init+24>
End of assembler dump.
(gdb)

@falcon: Is your machine productive?

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Wed, 27 Jun 2007 14:31:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

ok. let's dump in gdb the following:

type:

1. bt
 2. info registers
-

Subject: Re: Seg fault with quotauidlimit
Posted by [falcon](#) on Wed, 27 Jun 2007 15:47:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

@MaFL, yes, it is, but I think I can temporarily move all VMs to other HNs.

Subject: Re: Seg fault with quotauidlimit
Posted by [MaFL](#) on Wed, 27 Jun 2007 15:53:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

i think it's not needed anymore.

I just found an empty to play with.

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Wed, 27 Jun 2007 15:55:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

Which I'm thankful for very much! investigating right now.

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Wed, 27 Jun 2007 17:26:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Fixed it. Patch can be found at:

http://bugzilla.openvz.org/show_bug.cgi?id=635

Subject: Re: Seg fault with quotauidlimit
Posted by [dev](#) on Wed, 27 Jun 2007 17:27:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Fixed it. Patch can be found at:
http://bugzilla.openvz.org/show_bug.cgi?id=635

Subject: Re: Seg fault with quotauidlimit
Posted by [MaFL](#) on Wed, 27 Jun 2007 18:24:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tnx!!

Subject: Re: *SOLVED* Seg fault with quotauidlimit
Posted by [falcon](#) on Sat, 30 Jun 2007 16:17:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thank you very much!

Will you release a new vzctl or will I have to patch it myself?

Subject: Re: *SOLVED* Seg fault with quotauidlimit
Posted by [dev](#) on Mon, 02 Jul 2007 07:11:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sure we will release the fix in next version, but since Kir is at Linux Syposium it will take some time :/
