
Subject: IBM RAID controller re-mount lock

Posted by [maximilian.csuk](#) on Wed, 05 Aug 2009 14:35:44 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear OpenVZ community,

We have a system from IBM (System x3400) with an IBM ServeRAID 8K-I controller on which we are running Debian-Lenny 4.1.2-25. The kernel is `openvz-kernel 2.6.26-2-openvz-686`.

The containers are located on a separate `ext3`-partition and are mounted like this:
`/dev/drbd0` on `/mnt/hq-drbd-stor01` type `ext3` (rw)

In particular, we experience a lockup when we try to un-mount and remount the container's partition.

The following steps work:

```
vzctl chkpnt 77
vzctl stop 77
/etc/init.d/openvz stop
/etc/init.d/openvz start
vzctl restore 77
```

But these don't work:

```
vzctl chkpnt 77
vzctl stop 77
/etc/init.d/openvz stop
umount /dev/drbd0
mount /dev/drbd0
/etc/init.d/openvz start
vzctl restore 77
```

The restore produces the lockup.

The output until it hangs:

```
Restoring VE ...
Starting VE ...
VE is mounted
  undump...
Adding IP address(es): 10.20.60.20
Setting CPU units: 1000
Configure meminfo: 65536
```

Nothing more is displayed after this.

The host shows 100% CPU (system-)usage, load goes up to about 10. From our understanding,

a container's processes run inside a kernel-process on the host. The container's currently active process (as far as we can tell, it's a different one on any try) takes all the CPU from the host.

OpenVZ-version is 036test001.

We already tried changing the RAID-controller driver in use from aacraid-1.1.5.2449 to aacraid-1.1.5.2467 but without any luck.

Then we tried using USB-disks instead of the dev/drbd0 where the restore worked even after re-mounting, so the issue must lie within the RAID controller.

We also tried to run a DRBD device on two USB-disks instead of using ServeRAID attached disks this constellation also works, so we can assume that DRDB is working correctly

Is this a known issue? If yes, (how) can it be solved?

Is there a way to debug this and provide you and us with more information?

Thank you for your time.,
Maximilian Csuk
ipcenter.at

Subject: Re: IBM RAID controller re-mount lock
Posted by [maratrus](#) on Thu, 06 Aug 2009 09:00:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

- let's increase the verbose level of vzctl utility. In /etc/vz/vz.conf please set LOG_LEVEL and VERBOSE variables both to 10.

-

Quote:

OpenVZ-version is 036test001.

Sorry, I didn't catch this statement.

-

Quote:

openvz-kernel 2.6.26-2-openvz-686.

Do you have a chance to check the stable rhel5-2.6.18 or at least 2.6.18 kernels? I'm not sure that you do, but nevertheless ...

<http://download.openvz.org/kernel/>

- Are there any suspicious messages in dmesg or in logs?

-

Quote:

The host shows 100% CPU (system-)usage, load goes up to about 10. From our understanding, a

container's processes run inside a kernel-process on the host. The container's currently active process (as far as we can tell, it's a different one on any try) takes all the CPU from the host.

Do you have a chance to get the pid of the process that consumes 100% cpu? What is the output of "/proc/pid/wchan"?

-

It may be useful to gather alt-sysrq-* outputs:

alt-sysrq-

- 1) "m" - for memory info dump
- 2) "p" - for registers - several times, please, twice the number of CPUs
- 3) "a" - for scheduler stat - 3 times
- 4) "w" - another scheduler info - 3 times
- 5) "t" - for all processes calltraces. Warning - this is a resource consuming operation. At least - twice.

Subject: Re: IBM RAID controller re-mount lock

Posted by [maximilian.csuk](#) on Fri, 07 Aug 2009 09:07:26 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote: let's increase the verbose level of vzctl utility. In /etc/vz/vz.conf please set LOG_LEVEL and VERBOSE variables both to 10.

Done.

Quote: Do you have a chance to get the pid of the process that consumes 100% cpu? What is the output of "/proc/pid/wchan"?

Yes. The output is "0".

Quote: Are there any suspicious messages in dmesg or in logs?

I am not sure, but I guess not. Please have a look:

/var/log/vzctl:

```
2009-08-07T09:40:01+0200 vzctl : VE 77 : VE start in progress...
2009-08-07T09:40:01+0200 vzctl : VE 77 : Restoring completed successfully
2009-08-07T09:41:28+0200 vzctl : VE 77 : Setting up checkpoint...
2009-08-07T11:41:28+0400 vzctl : VE 77 :      suspend...
2009-08-07T11:41:28+0400 vzctl : VE 77 :      dump...
2009-08-07T11:41:29+0400 vzctl : VE 77 :      kill...
2009-08-07T09:41:29+0200 vzctl : VE 77 : Running: /usr/lib/vzctl/scripts/vps-net_del
2009-08-07T09:41:29+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota stat 77 -f
2009-08-07T09:41:29+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota off 77
2009-08-07T09:41:29+0200 vzctl : VE 77 : VE is unmounted
2009-08-07T09:41:29+0200 vzctl : VE 77 : Checkpointing completed successfully
```

```
2009-08-07T09:41:36+0200 vzctl : VE 77 : Restoring VE ...
2009-08-07T09:41:36+0200 vzctl : VE 77 : Starting VE ...
2009-08-07T09:41:36+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota show 77
2009-08-07T09:41:36+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota on 77 -r 0 -b 1048676 -B
1153124 -i 200100 -l 220100 -e 0 -n 0 -s 0
2009-08-07T09:41:36+0200 vzctl : VE 77 : Mounting root: /var/lib/vz/root/77 /var/lib/vz/private/77
2009-08-07T09:41:36+0200 vzctl : VE 77 : VE is mounted
2009-08-07T09:41:36+0200 vzctl : VE 77 : Adding IP address(es): 10.20.60.20
2009-08-07T09:41:36+0200 vzctl : VE 77 : Running: /usr/lib/vzctl/scripts/vps-net_add
2009-08-07T09:41:38+0200 vzctl : VE 77 : Setting CPU units: 1000
2009-08-07T09:41:38+0200 vzctl : VE 77 : Configure meminfo: 65536
2009-08-07T09:41:38+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota stat 77 -f
2009-08-07T09:41:38+0200 vzctl : VE 77 : Running: vzquota setlimit 77 -b 1048576 -B 1153024 -i
200000 -l 220000 -e 0 -n 0
2009-08-07T09:41:38+0200 vzctl : VE 77 : VE start in progress...
2009-08-07T09:41:38+0200 vzctl : VE 77 : Restoring completed succesfully
2009-08-07T09:42:11+0200 vzctl : VE 77 : Setting up checkpoint...
2009-08-07T11:42:11+0400 vzctl : VE 77 : suspend...
2009-08-07T11:42:11+0400 vzctl : VE 77 : dump...
2009-08-07T11:42:11+0400 vzctl : VE 77 : kill...
2009-08-07T09:42:11+0200 vzctl : VE 77 : Running: /usr/lib/vzctl/scripts/vps-net_del
2009-08-07T09:42:11+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota stat 77 -f
2009-08-07T09:42:11+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota off 77
2009-08-07T09:42:11+0200 vzctl : VE 77 : VE is unmounted
2009-08-07T09:42:11+0200 vzctl : VE 77 : Checkpointing completed succesfully
2009-08-07T09:42:12+0200 vzctl : VE 77 : Restoring VE ...
2009-08-07T09:42:12+0200 vzctl : VE 77 : Starting VE ...
2009-08-07T09:42:12+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota show 77
2009-08-07T09:42:12+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota on 77 -r 0 -b 1048676 -B
1153124 -i 200100 -l 220100 -e 0 -n 0 -s 0
2009-08-07T09:42:12+0200 vzctl : VE 77 : Mounting root: /var/lib/vz/root/77 /var/lib/vz/private/77
2009-08-07T09:42:12+0200 vzctl : VE 77 : VE is mounted
2009-08-07T09:42:12+0200 vzctl : VE 77 : Adding IP address(es): 10.20.60.20
2009-08-07T09:42:12+0200 vzctl : VE 77 : Running: /usr/lib/vzctl/scripts/vps-net_add
2009-08-07T09:42:14+0200 vzctl : VE 77 : Setting CPU units: 1000
2009-08-07T09:42:14+0200 vzctl : VE 77 : Configure meminfo: 65536
2009-08-07T09:42:14+0200 vzctl : VE 77 : Running: /usr/sbin/vzquota stat 77 -f
2009-08-07T09:42:14+0200 vzctl : VE 77 : Running: vzquota setlimit 77 -b 1048576 -B 1153024 -i
200000 -l 220000 -e 0 -n 0
```

Nothing more after this...

In dmesg, shortly before the lockup, it says:

```
[74677.049342] CT: 77: started
```

But then, there are more and more messages showing that the particular pid is hanging. In this try, it was sendmail-mta, but others have been the cause too (like apache):

```
[75007.309371] BUG: soft lockup - CPU#3 stuck for 61s! [sendmail-mta:13227]
```

```

[75007.309371] Modules linked in: vzethdev vznetdev simfs vzrst vzcpt tun vzmon xt_tcpudp
xt_length ipt_ttl xt_tcpmss xt_TCPMSS iptable_mangle iptable_filter xt_multiport xt_limit xt_dscp
ipt_REJECT ip_tables x_tables vxdquota vzdev drbd cn bridge ipv6 bonding loop iTCO_wdt
i2c_i801 rng_core pcpkr shpchr parport_pc parport i2c_core i5000_edac edac_core pci_hotplug
container button joydev evdev ext3 jbd mbcache usbhid hid ff_memless usb_storage sg sr_mod
cdrom 8139too ata_piix ata_generic libata dock ses enclosure sd_mod 8139cp mii
ide_pci_generic ide_core ehci_hcd uhci_hcd usbcore tg3 e1000e aacraid scsi_mod thermal
processor fan thermal_sys [last unloaded: simfs]
[75007.309371]
[75007.309371] Pid: 13227, comm: sendmail-mta Not tainted (2.6.26-2-openvz-686 #1
036test001)
[75007.309371] EIP: 0060:[<c013d31f>] EFLAGS: 00000257 CPU: 3
[75007.309371] EIP is at up_read+0x7/0xf
[75007.309371] EAX: f7996bf8 EBX: f156121c ECX: 00000000 EDX: 00000001
[75007.309371] ESI: f7996bf8 EDI: f7996bc0 EBP: 00000090 ESP: f148fc2c
[75007.309371] DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
[75007.309371] CR0: 8005003b CR2: b7856000 CR3: 37c3e000 CR4: 000006f0
[75007.309371] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[75007.309371] DR6: ffff0ff0 DR7: 00000400
[75007.309371] [<c011b3f4>] ? do_page_fault+0x50e/0x8c6
[75007.309371] [<c0131599>] ? run_timer_softirq+0x30/0x1d1
[75007.309371] [<c012dff2>] ? __do_softirq+0x9b/0x14e
[75007.309371] [<c011aee6>] ? do_page_fault+0x0/0x8c6
[75007.309371] [<c02c5672>] ? error_code+0x72/0x78
[75007.309371] [<c02c00d8>] ? init_cyrix+0x168/0x46b
[75007.309371] [<c0161b69>] ? file_read_actor+0x2b/0xcc
[75007.309371] [<c016398d>] ? generic_file_aio_read+0x2e2/0x4da
[75007.309371] [<c0184346>] ? do_sync_read+0xbf/0xfe
[75007.309371] [<c013a2c4>] ? autoremove_wake_function+0x0/0x2d
[75007.309371] [<f9191224>] ? file_pread+0x3c/0x66 [vzrst]
[75007.309371] [<f9190f5d>] ? _rst_get_object+0x85/0x98 [vzrst]
[75007.309371] [<f9195d1d>] ? do_rst_mm+0x1642/0x1b40 [vzrst]
[75007.309371] [<f919634b>] ? rst_mm_complete+0x130/0x1f5 [vzrst]
[75007.309371] [<f918fe58>] ? hook+0xa13/0x19db [vzrst]
[75007.309371] [<c01db7ed>] ? cic_free_func+0x4b/0x53
[75007.309371] [<c012cd42>] ? do_exit+0x642/0x646
[75007.309371] [<c0197941>] ? mntput_no_expire+0x13/0xd2
[75007.309371] [<c01079c7>] ? __switch_to+0xa5/0x125
[75007.309371] [<c0125314>] ? schedule_tail+0x17/0x44
[75007.309371] [<c01087fa>] ? ret_from_fork+0x6/0xe
[75007.309371] [<f918f445>] ? hook+0x0/0x19db [vzrst]
[75007.309371] [<f918f445>] ? hook+0x0/0x19db [vzrst]
[75007.309371] [<c0109597>] ? kernel_thread_helper+0x7/0x10
[75007.309371] =====

```

Sorry for the long output, but this is what the sysrq-calls gave me:

Output of sysrq-m:

```
[ 739.006053] Show Memory
[ 739.006053] Mem-info:
[ 739.006053] DMA per-cpu:
[ 739.006053] CPU 0: hi: 0, btch: 1 usd: 0
[ 739.006053] CPU 1: hi: 0, btch: 1 usd: 0
[ 739.006053] CPU 2: hi: 0, btch: 1 usd: 0
[ 739.006053] CPU 3: hi: 0, btch: 1 usd: 0
[ 739.006053] Normal per-cpu:
[ 739.006053] CPU 0: hi: 186, btch: 31 usd: 132
[ 739.006053] CPU 1: hi: 186, btch: 31 usd: 55
[ 739.006053] CPU 2: hi: 186, btch: 31 usd: 156
[ 739.006053] CPU 3: hi: 186, btch: 31 usd: 88
[ 739.006053] HighMem per-cpu:
[ 739.006053] CPU 0: hi: 186, btch: 31 usd: 156
[ 739.006053] CPU 1: hi: 186, btch: 31 usd: 64
[ 739.006053] CPU 2: hi: 186, btch: 31 usd: 112
[ 739.006053] CPU 3: hi: 186, btch: 31 usd: 41
[ 739.006053] Active:7699 inactive:4989 dirty:53 writeback:0 unstable:0
[ 739.006053] free:1535325 slab:2934 mapped:1615 pagetables:136 bounce:0
[ 739.006053] DMA free:12852kB min:68kB low:84kB high:100kB active:0kB inactive:0kB
present:16240kB pages_scanned:0 all_unreclaimable? no
[ 739.006053] lowmem_reserve[]: 0 872 7089 7089
[ 739.006053] Normal free:798948kB min:3740kB low:4672kB high:5608kB active:12220kB
inactive:5760kB present:893200kB pages_scanned:0 all_unreclaimable? no
[ 739.006053] lowmem_reserve[]: 0 0 49735 49735
[ 739.006053] HighMem free:5329500kB min:512kB low:7180kB high:13852kB active:18576kB
inactive:14196kB present:6366080kB pages_scanned:0 all_unreclaimable? no
[ 739.006053] lowmem_reserve[]: 0 0 0 0
[ 739.006053] DMA: 3*4kB 3*8kB 5*16kB 2*32kB 4*64kB 3*128kB 1*256kB 1*512kB 1*1024kB
1*2048kB 2*4096kB = 12852kB
[ 739.006053] Normal: 107*4kB 315*8kB 236*16kB 170*32kB 91*64kB 47*128kB 25*256kB
7*512kB 5*1024kB 7*2048kB 182*4096kB = 798916kB
[ 739.006053] HighMem: 3*4kB 19*8kB 41*16kB 112*32kB 83*64kB 44*128kB 12*256kB
5*512kB 2*1024kB 3*2048kB 1294*4096kB = 5329396kB
[ 739.006053] 10767 total pagecache pages
[ 739.006053] Swap cache: add 0, delete 0, find 0/0
[ 739.006053] Free swap = 979956kB
[ 739.006053] Total swap = 979956kB
[ 739.024690] 1835008 pages of RAM
[ 739.024690] 1605632 pages of HIGHMEM
[ 739.024690] 279573 reserved pages
[ 739.024690] 10533 pages shared
[ 739.024690] 0 pages swap cached
[ 739.024690] 53 pages dirty
[ 739.024690] 0 pages writeback
[ 739.024690] 1615 pages mapped
```



```

[ 739.024690] 2934 pages slab
[ 739.024690] 136 pages pagetables
[ 739.024690] Top 10 caches:
[ 739.024690] ext3_inode_cache   : size  461760 objsize   512
[ 739.024690] sysfs_dir_cache     : size  434176 objsize   44
[ 739.024690] radix_tree_node    : size  413696 objsize  288
[ 739.024690] filp                : size  450560 objsize   192
[ 739.024690] shmem_inode_cache  : size  552960 objsize  468
[ 739.024690] inode_cache        : size  688128 objsize   360
[ 739.024690] size-2048          : size  2878720 objsize  2048
[ 739.024690] size-512           : size  540800 objsize   512
[ 739.024690] size-64            : size  524288 objsize   64
[ 739.024690] dentry             : size  774144 objsize  144

```

Output sysrg-p:

```

[ 961.824746] Show Regs
[ 961.825125] ----- IPI show regs -----
[ 961.825125] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 961.825125] EIP: 0060:[<c01079c7>] EFLAGS: 00000046 CPU: 1
[ 961.825125] EIP is at __switch_to+0xa5/0x125
[ 961.825125] EAX: 00000000 EBX: f796ae60 ECX: 00000018 EDX: 00000000
[ 961.825125] ESI: 00000001 EDI: c4f21c80 EBP: f796ac10 ESP: f5cd7eac
[ 961.825125] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 961.825125] CR0: 8005003b CR2: b7f8c000 CR3: 379f8000 CR4: 000006f0
[ 961.825125] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 961.825125] DR6: ffff0ff0 DR7: 00000400
[ 961.825125] [<c02c4058>] ? schedule+0x7a6/0x886
[ 961.825125] [<c012a9fc>] ? do_syslog+0x129/0x493
[ 961.825125] [<c013a2c4>] ? autoremove_wake_function+0x0/0x2d
[ 961.825125] [<c01ba565>] ? kmsg_read+0x0/0x36
[ 961.825125] [<c01b311d>] ? proc_reg_read+0x58/0x6b
[ 961.825125] [<c01b30c5>] ? proc_reg_read+0x0/0x6b
[ 961.825125] [<c0184abf>] ? vfs_read+0x81/0x11e
[ 961.825125] [<c0184f53>] ? sys_read+0x3c/0xa2
[ 961.825125] [<c0108972>] ? syscall_call+0x7/0xb
[ 961.825125] =====
[ 961.825130] ----- IPI show regs -----
[ 961.825130] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 961.825130] EIP: 0060:[<c010df2a>] EFLAGS: 00000246 CPU: 2
[ 961.825130] EIP is at mwait_idle+0x2f/0x3d
[ 961.825130] EAX: 00000000 EBX: c010defb ECX: 00000000 EDX: 00000000
[ 961.825130] ESI: 00000002 EDI: 00000000 EBP: 00000000 ESP: f749ffac
[ 961.825130] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 961.825130] CR0: 8005003b CR2: b7f413b0 CR3: 35cb9000 CR4: 000006f0
[ 961.825130] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 961.825130] DR6: ffff0ff0 DR7: 00000400

```

```
[ 961.825130] [<c0107627>] cpu_idle+0xab/0xcb
[ 961.825130] =====
[ 961.824745] ----- IPI show regs -----
[ 961.824745] Pid: 4094, comm: syslogd Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 961.824745] EIP: 0060:[<c02c55f8>] EFLAGS: 00000046 CPU: 3
[ 961.824745] EIP is at __kprobes_text_start+0x0/0x8
[ 961.824745] EAX: 00000b90 EBX: 00000000 ECX: b7dfc000 EDX: 00002000
[ 961.824745] ESI: c4e684f4 EDI: f4d7dd28 EBP: 00000b90 ESP: f4d7dcb8
[ 961.824745] DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
[ 961.824745] CR0: 8005003b CR2: b7dfc000 CR3: 37904000 CR4: 000006f0
[ 961.824745] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 961.824745] DR6: ffff0ff0 DR7: 00000400
[ 961.824745] [<c0161b69>] ? file_read_actor+0x2b/0xcc
[ 961.824745] [<c016398d>] ? generic_file_aio_read+0x2e2/0x4da
[ 961.824745] [<c0184346>] ? do_sync_read+0xbf/0xfe
[ 961.824745] [<c013a2c4>] ? autoremove_wake_function+0x0/0x2d
[ 961.824745] [<f919f224>] ? file_pread+0x3c/0x66 [vzrst]
[ 961.824745] [<f919ef5d>] ? _rst_get_object+0x85/0x98 [vzrst]
[ 961.824745] [<f91a3d1d>] ? do_rst_mm+0x1642/0x1b40 [vzrst]
[ 961.824745] [<f91a434b>] ? rst_mm_complete+0x130/0x1f5 [vzrst]
[ 961.824745] [<f919de58>] ? hook+0xa13/0x19db [vzrst]
[ 961.824745] [<c01db7ed>] ? cic_free_func+0x4b/0x53
[ 961.824745] [<c012cd42>] ? do_exit+0x642/0x646
[ 961.824745] [<c01079c7>] ? __switch_to+0xa5/0x125
[ 961.824745] [<c0125314>] ? schedule_tail+0x17/0x44
[ 961.824745] [<c01087fa>] ? ret_from_fork+0x6/0xe
[ 961.824745] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 961.824745] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 961.824745] [<c0109597>] ? kernel_thread_helper+0x7/0x10
[ 961.824745] =====
```

```
[ 1044.357052] Show Regs
[ 1044.357052] ----- IPI show regs -----
[ 1044.357052] Pid: 4094, comm: syslogd Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1044.357052] EIP: 0060:[<c02c55f8>] EFLAGS: 00000046 CPU: 3
[ 1044.357052] EIP is at __kprobes_text_start+0x0/0x8
[ 1044.357052] EAX: 00000b90 EBX: 00000000 ECX: b7dfc000 EDX: 00002000
[ 1044.357052] ESI: c4e684f4 EDI: f4d7dd28 EBP: 00000b90 ESP: f4d7dcb8
[ 1044.357052] DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
[ 1044.357052] CR0: 8005003b CR2: b7dfc000 CR3: 37904000 CR4: 000006f0
[ 1044.357052] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1044.357052] DR6: ffff0ff0 DR7: 00000400
[ 1044.357052] [<c0161b69>] ? file_read_actor+0x2b/0xcc
[ 1044.357052] [<c016398d>] ? generic_file_aio_read+0x2e2/0x4da
[ 1044.357052] [<c0184346>] ? do_sync_read+0xbf/0xfe
[ 1044.357052] [<c013a2c4>] ? autoremove_wake_function+0x0/0x2d
```



```

[ 1044.357052] [<f919f224>] ? file_pread+0x3c/0x66 [vzrst]
[ 1044.357052] [<f919ef5d>] ? _rst_get_object+0x85/0x98 [vzrst]
[ 1044.357052] [<f91a3d1d>] ? do_rst_mm+0x1642/0x1b40 [vzrst]
[ 1044.357052] [<f91a434b>] ? rst_mm_complete+0x130/0x1f5 [vzrst]
[ 1044.357052] [<f919de58>] ? hook+0xa13/0x19db [vzrst]
[ 1044.357052] [<c01db7ed>] ? cic_free_func+0x4b/0x53
[ 1044.357052] [<c012cd42>] ? do_exit+0x642/0x646
[ 1044.357052] [<c01079c7>] ? __switch_to+0xa5/0x125
[ 1044.357052] [<c0125314>] ? schedule_tail+0x17/0x44
[ 1044.357052] [<c01087fa>] ? ret_from_fork+0x6/0xe
[ 1044.357052] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 1044.357052] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 1044.357052] [<c0109597>] ? kernel_thread_helper+0x7/0x10
[ 1044.357052] =====
[ 1044.357361] ----- IPI show regs -----
[ 1044.357361] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1044.357361] EIP: 0060:[<c010df2a>] EFLAGS: 00000246 CPU: 2
[ 1044.357361] EIP is at mwait_idle+0x2f/0x3d
[ 1044.357361] EAX: 00000000 EBX: c010defb ECX: 00000000 EDX: 00000000
[ 1044.357361] ESI: 00000002 EDI: 00000000 EBP: 00000000 ESP: f749ffac
[ 1044.357361] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 1044.357361] CR0: 8005003b CR2: b7be9195 CR3: 379f4000 CR4: 000006f0
[ 1044.357361] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1044.357361] DR6: ffff0ff0 DR7: 00000400
[ 1044.357361] [<c0107627>] cpu_idle+0xab/0xcb
[ 1044.357361] =====
[ 1044.357361] ----- IPI show regs -----
[ 1044.357361] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1044.357361] EIP: 0060:[<c02c3f82>] EFLAGS: 00000003 CPU: 1
[ 1044.357361] EIP is at schedule+0x6d0/0x886
[ 1044.357361] EAX: c03cb040 EBX: f796ac10 ECX: 00000001 EDX: fc303755
[ 1044.357361] ESI: f5cfcd80 EDI: f5c17540 EBP: f748ce70 ESP: f7493f58
[ 1044.357361] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 1044.357361] CR0: 8005003b CR2: b7f1c4e0 CR3: 379f4000 CR4: 000006f0
[ 1044.357361] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1044.357361] DR6: ffff0ff0 DR7: 00000400
[ 1044.357361] [<c013eb18>] ? getnstimeofday+0x37/0xbc
[ 1044.357361] [<c013cfac>] ? ktime_get+0xd/0x21
[ 1044.357361] [<c01421d8>] ? tick_nohz_stop_idle+0x19/0x45
[ 1044.357361] [<c01427fb>] ? tick_nohz_restart_sched_tick+0x28/0x11c
[ 1044.357361] [<c010defb>] ? mwait_idle+0x0/0x3d
[ 1044.357361] [<c0107642>] ? cpu_idle+0xc6/0xcb
[ 1044.357361] =====

```

```

[ 1099.329424] Show Regs
[ 1099.325423] ----- IPI show regs -----

```

```

[ 1099.325423] Pid: 4094, comm: syslogd Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1099.325423] EIP: 0060:[<c02c55f8>] EFLAGS: 00000046 CPU: 3
[ 1099.325423] EIP is at __kprobes_text_start+0x0/0x8
[ 1099.325423] EAX: 00000b90 EBX: 00000000 ECX: b7dfc000 EDX: 00002000
[ 1099.325423] ESI: c4e684f4 EDI: f4d7dd28 EBP: 00000b90 ESP: f4d7dcb8
[ 1099.325423] DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
[ 1099.325423] CR0: 8005003b CR2: b7dfc000 CR3: 37904000 CR4: 000006f0
[ 1099.325423] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1099.325423] DR6: ffff0ff0 DR7: 00000400
[ 1099.325423] [<c0161b69>] ? file_read_actor+0x2b/0xcc
[ 1099.325423] [<c016398d>] ? generic_file_aio_read+0x2e2/0x4da
[ 1099.325423] [<c0184346>] ? do_sync_read+0xbf/0xfe
[ 1099.325423] [<c013a2c4>] ? autoremove_wake_function+0x0/0x2d
[ 1099.325423] [<f919f224>] ? file_pread+0x3c/0x66 [vzrst]
[ 1099.325423] [<f919ef5d>] ? _rst_get_object+0x85/0x98 [vzrst]
[ 1099.325423] [<f91a3d1d>] ? do_rst_mm+0x1642/0x1b40 [vzrst]
[ 1099.325423] [<f91a434b>] ? rst_mm_complete+0x130/0x1f5 [vzrst]
[ 1099.325423] [<f919de58>] ? hook+0xa13/0x19db [vzrst]
[ 1099.325423] [<c01db7ed>] ? cic_free_func+0x4b/0x53
[ 1099.325423] [<c012cd42>] ? do_exit+0x642/0x646
[ 1099.325424] [<c01079c7>] ? __switch_to+0xa5/0x125
[ 1099.325424] [<c0125314>] ? schedule_tail+0x17/0x44
[ 1099.325424] [<c01087fa>] ? ret_from_fork+0x6/0xe
[ 1099.325424] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 1099.325424] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 1099.325424] [<c0109597>] ? kernel_thread_helper+0x7/0x10
[ 1099.325424] =====
[ 1099.329424] ----- IPI show regs -----
[ 1099.329424] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1099.329424] EIP: 0060:[<c010df2a>] EFLAGS: 00000246 CPU: 2
[ 1099.329424] EIP is at mwait_idle+0x2f/0x3d
[ 1099.329424] EAX: 00000000 EBX: c010defb ECX: 00000000 EDX: 00000000
[ 1099.329424] ESI: 00000002 EDI: 00000000 EBP: 00000000 ESP: f749ffac
[ 1099.329424] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 1099.329424] CR0: 8005003b CR2: b7f75000 CR3: 35cb9000 CR4: 000006f0
[ 1099.329424] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1099.329424] DR6: ffff0ff0 DR7: 00000400
[ 1099.329424] [<c0107627>] cpu_idle+0xab/0xcb
[ 1099.329424] =====
[ 1099.329424] ----- IPI show regs -----
[ 1099.329429] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1099.329432] EIP: 0060:[<c011fbb0>] EFLAGS: 00000092 CPU: 1
[ 1099.329438] EIP is at hrtick_start_fair+0x1e/0x12c
[ 1099.329440] EAX: f748ce70 EBX: 00000000 ECX: f796ac38 EDX: f5cd6000
[ 1099.329443] ESI: f796ac10 EDI: c4f23fa0 EBP: f796ac38 ESP: f7493f2c
[ 1099.329446] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 1099.329448] CR0: 8005003b CR2: b7f4b1c0 CR3: 379f8000 CR4: 000006f0
[ 1099.329451] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000

```

```
[ 1099.329454] DR6: ffff0ff0 DR7: 00000400
[ 1099.329467] [<c0120f8e>] ? pick_next_task_fair+0x4d/0x53
[ 1099.329473] [<c02c3bda>] ? schedule+0x328/0x886
[ 1099.329478] [<c013eb18>] ? getnstimeofday+0x37/0xbc
[ 1099.329484] [<c013cf78>] ? ktime_get_ts+0x22/0x49
[ 1099.329489] [<c013cfac>] ? ktime_get+0xd/0x21
[ 1099.329494] [<c01421d8>] ? tick_nohz_stop_idle+0x19/0x45
[ 1099.329500] [<c01427fb>] ? tick_nohz_restart_sched_tick+0x28/0x11c
[ 1099.329505] [<c010defb>] ? mwait_idle+0x0/0x3d
[ 1099.329511] [<c0107642>] ? cpu_idle+0xc6/0xcb
[ 1099.329520] =====
```

```
[ 1911.608858] Show Regs
[ 1911.608858] ----- IPI show regs -----
[ 1911.608858] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1911.608858] EIP: 0060:[<c010df2a>] EFLAGS: 00000246 CPU: 2
[ 1911.608858] EIP is at mwait_idle+0x2f/0x3d
[ 1911.608858] EAX: 00000000 EBX: c010defb ECX: 00000000 EDX: 00000000
[ 1911.608858] ESI: 00000002 EDI: 00000000 EBP: 00000000 ESP: f749ffac
[ 1911.608858] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 1911.608858] CR0: 8005003b CR2: b7f413b0 CR3: 35cb9000 CR4: 000006f0
[ 1911.608858] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1911.608858] DR6: ffff0ff0 DR7: 00000400
[ 1911.608858] [<c0107627>] cpu_idle+0xab/0xcb
[ 1911.608858] =====
[ 1911.607256] ----- IPI show regs -----
[ 1911.607256] Pid: 4222, comm: sendmail-mta Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1911.607256] EIP: 0060:[<c02c55f8>] EFLAGS: 00000046 CPU: 3
[ 1911.607256] EIP is at __kprobes_text_start+0x0/0x8
[ 1911.607256] EAX: 00000ec0 EBX: 00000000 ECX: b7857000 EDX: 00001000
[ 1911.607256] ESI: c4e92c2c EDI: f4859d28 EBP: 00000ec0 ESP: f4859cb8
[ 1911.607256] DS: 007b ES: 007b FS: 00d8 GS: 0033 SS: 0068
[ 1911.607256] CR0: 8005003b CR2: b7857000 CR3: 378bc000 CR4: 000006f0
[ 1911.607257] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1911.607257] DR6: ffff0ff0 DR7: 00000400
[ 1911.607257] [<c0161b69>] ? file_read_actor+0x2b/0xcc
[ 1911.607257] [<c016398d>] ? generic_file_aio_read+0x2e2/0x4da
[ 1911.607257] [<c0184346>] ? do_sync_read+0xbf/0xfe
[ 1911.607257] [<c013a2c4>] ? autoremove_wake_function+0x0/0x2d
[ 1911.607257] [<f919f224>] ? file_pread+0x3c/0x66 [vzrst]
[ 1911.607257] [<f919ef5d>] ? _rst_get_object+0x85/0x98 [vzrst]
[ 1911.607257] [<f91a3d1d>] ? do_rst_mm+0x1642/0x1b40 [vzrst]
[ 1911.607257] [<f91a434b>] ? rst_mm_complete+0x130/0x1f5 [vzrst]
[ 1911.607257] [<f919de58>] ? hook+0xa13/0x19db [vzrst]
[ 1911.607257] [<c01079c7>] ? __switch_to+0xa5/0x125
[ 1911.607257] [<c0125314>] ? schedule_tail+0x17/0x44
```

```

[ 1911.607257] [<c01087fa>] ? ret_from_fork+0x6/0xe
[ 1911.607257] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 1911.607257] [<f919d445>] ? hook+0x0/0x19db [vzrst]
[ 1911.607257] [<c0109597>] ? kernel_thread_helper+0x7/0x10
[ 1911.607257] =====
[ 1911.608858] ----- IPI show regs -----
[ 1911.608858] Pid: 0, comm: swapper Not tainted (2.6.26-2-openvz-686 #1 036test001)
[ 1911.608858] EIP: 0060:[<c02c3fae>] EFLAGS: 00000007 CPU: 1
[ 1911.608858] EIP is at schedule+0x6fc/0x886
[ 1911.608858] EAX: 379f8000 EBX: f796ac10 ECX: 00000001 EDX: c4f20040
[ 1911.608858] ESI: f5c171c0 EDI: f5c17540 EBP: f748ce70 ESP: f7493f58
[ 1911.608858] DS: 007b ES: 007b FS: 00d8 GS: 0000 SS: 0068
[ 1911.608858] CR0: 8005003b CR2: b7f8c000 CR3: 379f8000 CR4: 000006f0
[ 1911.608858] DR0: 00000000 DR1: 00000000 DR2: 00000000 DR3: 00000000
[ 1911.608858] DR6: ffff0ff0 DR7: 00000400
[ 1911.608858] [<c013eb18>] ? getnstimeofday+0x37/0xbc
[ 1911.608858] [<c013cfac>] ? ktime_get+0xd/0x21
[ 1911.608858] [<c01421d8>] ? tick_nohz_stop_idle+0x19/0x45
[ 1911.608858] [<c01427fb>] ? tick_nohz_restart_sched_tick+0x28/0x11c
[ 1911.608858] [<c010defb>] ? mwait_idle+0x0/0x3d
[ 1911.608858] [<c0107642>] ? cpu_idle+0xc6/0xcb
[ 1911.608858] =====

```

Sysrq-a (3 times):

```

[ 1966.916085] SysRq HELP : loglevel0-8 reBoot Crashdump tErm Full debuG klll saK aLlcpus
showMem Nice powerOff showPc show-all-timers(Q) unRaw Sync showTasks Unmount
shoW-blocked-tasks
[ 2528.787849] SysRq HELP : loglevel0-8 reBoot Crashdump tErm Full debuG klll saK aLlcpus
showMem Nice powerOff showPc show-all-timers(Q) unRaw Sync showTasks Unmount
shoW-blocked-tasks
[ 2530.345032] SysRq HELP : loglevel0-8 reBoot Crashdump tErm Full debuG klll saK aLlcpus
showMem Nice powerOff showPc show-all-timers(Q) unRaw Sync showTasks Unmount
shoW-blocked-tasks

```

Sysrq-w:

```

Sched Debug Version: v0.07, 2.6.26-2-openvz-686 #1
[ 2065.319754] now at 1932715.499105 msecs
[ 2065.319754] .sysctl_sched_latency           : 60.000000
[ 2065.319754] .sysctl_sched_min_granularity           : 12.000000
[ 2065.319754] .sysctl_sched_wakeup_granularity       : 30.000000
[ 2065.319754] .sysctl_sched_child_runs_first         : 0.000001
[ 2065.319754] .sysctl_sched_features                 : 895
[ 2065.319754]
[ 2065.319754] cpu#0, 2327.615 MHz

```

```

[ 2065.319754] .nr_running          : 1
[ 2065.319754] .load                : 1024
[ 2065.319754] .nr_switches        : 157692
[ 2065.319754] .nr_load_updates    : 368147
[ 2065.319754] .nr_uninterruptible : 4294967286
[ 2065.319754] .jiffies            : 408179
[ 2065.319754] .next_balance       : 0.408179
[ 2065.319754] .curr->pid          : 4097
[ 2065.319754] .clock              : 1932716.006896
[ 2065.319754] .cpu_load[0]        : 0
[ 2065.319754] .cpu_load[1]        : 0
[ 2065.319754] .cpu_load[2]        : 0
[ 2065.319754] .cpu_load[3]        : 0
[ 2065.319754] .cpu_load[4]        : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[0]:
[ 2065.319754] .exec_clock         : 0.000000
[ 2065.319754] .MIN_vruntime       : 0.000001
[ 2065.319754] .min_vruntime       : 22937.847487
[ 2065.319754] .max_vruntime       : 0.000001
[ 2065.319754] .spread             : 0.000000
[ 2065.319754] .spread0            : 0.000000
[ 2065.319754] .nr_running         : 0
[ 2065.319754] .load               : 0
[ 2065.319754] .nr_spread_over     : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[0]:
[ 2065.319754] .exec_clock         : 0.000000
[ 2065.319754] .MIN_vruntime       : 0.000001
[ 2065.319754] .min_vruntime       : 22937.847487
[ 2065.319754] .max_vruntime       : 0.000001
[ 2065.319754] .spread             : 0.000000
[ 2065.319754] .spread0            : 0.000000
[ 2065.319754] .nr_running         : 1
[ 2065.319754] .load               : 1024
[ 2065.319754] .nr_spread_over     : 0
[ 2065.319754]
[ 2065.319754] runnable tasks:
[ 2065.319754]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2065.319754] -----
[ 2065.319754] R      bash 4097    22879.611338  205 120      0      0
0.000000  0.000000  0.000000
[ 2065.319754]
[ 2065.319754] cpu#1, 2327.615 MHz
[ 2065.319754] .nr_running          : 2
[ 2065.319754] .load                : 2048
[ 2065.319754] .nr_switches        : 155809

```

```

[ 2065.319754] .nr_load_updates      : 351066
[ 2065.319754] .nr_uninterruptible      : 1259
[ 2065.319754] .jiffies                  : 408179
[ 2065.319754] .next_balance            : 0.408179
[ 2065.319754] .curr->pid               : 2754
[ 2065.319754] .clock                   : 1932716.506691
[ 2065.319754] .cpu_load[0]             : 0
[ 2065.319754] .cpu_load[1]             : 0
[ 2065.319754] .cpu_load[2]             : 0
[ 2065.319754] .cpu_load[3]             : 0
[ 2065.319754] .cpu_load[4]             : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[1]:
[ 2065.319754] .exec_clock              : 0.000000
[ 2065.319754] .MIN_vruntime            : 0.000001
[ 2065.319754] .min_vruntime            : 35218.302793
[ 2065.319754] .max_vruntime            : 0.000001
[ 2065.319754] .spread                  : 0.000000
[ 2065.319754] .spread0                 : 12280.455306
[ 2065.319754] .nr_running              : 0
[ 2065.319754] .load                    : 0
[ 2065.319754] .nr_spread_over          : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[1]:
[ 2065.319754] .exec_clock              : 0.000000
[ 2065.319754] .MIN_vruntime            : 35158.302793
[ 2065.319754] .min_vruntime            : 35218.302793
[ 2065.319754] .max_vruntime            : 35158.302793
[ 2065.319754] .spread                  : 0.000000
[ 2065.319754] .spread0                 : 12280.455306
[ 2065.319754] .nr_running              : 2
[ 2065.319754] .load                    : 2048
[ 2065.319754] .nr_spread_over          : 0
[ 2065.319754]
[ 2065.319754] runnable tasks:
[ 2065.319754]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2065.319754] -----
[ 2065.319754] R   rsyslogd 2754   35158.302794   649 120      0      0
0.000000   0.000000   0.000000
[ 2065.319754]   rsyslogd 4254   35158.302793   484 120      0      0
0.000000   0.000000   0.000000
[ 2065.319754]
[ 2065.319754] cpu#2, 2327.615 MHz
[ 2065.319754] .nr_running              : 0
[ 2065.319754] .load                    : 0
[ 2065.319754] .nr_switches             : 115926
[ 2065.319754] .nr_load_updates        : 358812

```



```

[ 2065.319754] .nr_uninterruptible      : 4294965985
[ 2065.319754] .jiffies                          : 408179
[ 2065.319754] .next_balance                       : 0.408179
[ 2065.319754] .curr->pid                          : 0
[ 2065.319754] .clock                              : 1932717.001283
[ 2065.319754] .cpu_load[0]                       : 0
[ 2065.319754] .cpu_load[1]                       : 0
[ 2065.319754] .cpu_load[2]                       : 0
[ 2065.319754] .cpu_load[3]                       : 0
[ 2065.319754] .cpu_load[4]                       : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[2]:
[ 2065.319754] .exec_clock                         : 0.000000
[ 2065.319754] .MIN_vruntime                       : 0.000001
[ 2065.319754] .min_vruntime                       : 15277.589318
[ 2065.319754] .max_vruntime                       : 0.000001
[ 2065.319754] .spread                             : 0.000000
[ 2065.319754] .spread0                            : -7660.258169
[ 2065.319754] .nr_running                         : 0
[ 2065.319754] .load                               : 0
[ 2065.319754] .nr_spread_over                     : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[2]:
[ 2065.319754] .exec_clock                         : 0.000000
[ 2065.319754] .MIN_vruntime                       : 0.000001
[ 2065.319754] .min_vruntime                       : 15277.589318
[ 2065.319754] .max_vruntime                       : 0.000001
[ 2065.319754] .spread                             : 0.000000
[ 2065.319754] .spread0                            : -7660.258169
[ 2065.319754] .nr_running                         : 0
[ 2065.319754] .load                               : 0
[ 2065.319754] .nr_spread_over                     : 0
[ 2065.319754]
[ 2065.319754] runnable tasks:
[ 2065.319754]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2065.319754] -----
[ 2065.319754]
[ 2065.319754] cpu#3, 2327.615 MHz
[ 2065.319754] .nr_running          : 4
[ 2065.319754] .load                : 184788
[ 2065.319754] .nr_switches        : 35136
[ 2065.319754] .nr_load_updates    : 361726
[ 2065.319754] .nr_uninterruptible : 66
[ 2065.319754] .jiffies            : 408179
[ 2065.319754] .next_balance       : 0.408188
[ 2065.319754] .curr->pid          : 4222
[ 2065.319754] .clock              : 1932717.502539

```

```

[ 2065.319754] .cpu_load[0]           : 184788
[ 2065.319754] .cpu_load[1]           : 184788
[ 2065.319754] .cpu_load[2]           : 184788
[ 2065.319754] .cpu_load[3]           : 184788
[ 2065.319754] .cpu_load[4]           : 184788
[ 2065.319754]
[ 2065.319754] cfs_rq[3]:
[ 2065.319754] .exec_clock             : 0.000000
[ 2065.319754] .MIN_vruntime          : 0.000001
[ 2065.319754] .min_vruntime          : 874283.510611
[ 2065.319754] .max_vruntime          : 0.000001
[ 2065.319754] .spread                : 0.000000
[ 2065.319754] .spread0               : 851345.663124
[ 2065.319754] .nr_running            : 1
[ 2065.319754] .load                  : 1024
[ 2065.319754] .nr_spread_over        : 0
[ 2065.319754]
[ 2065.319754] cfs_rq[3]:
[ 2065.319754] .exec_clock             : 0.000000
[ 2065.319754] .MIN_vruntime          : 874163.510611
[ 2065.319754] .min_vruntime          : 874283.510611
[ 2065.319754] .max_vruntime          : 874223.510611
[ 2065.319754] .spread                : 60.000000
[ 2065.319754] .spread0               : 851345.663124
[ 2065.319754] .nr_running            : 3
[ 2065.319754] .load                  : 7266
[ 2065.319754] .nr_spread_over        : 0
[ 2065.319754]
[ 2065.319754] runnable tasks:
[ 2065.319754]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2065.319754] -----
[ 2065.319754]   watchdog/3  14    -4.329575    3  0          0          0
0.000000      0.000000      0.000000
[ 2065.319754]   events/3   18   874223.510611   738 115          0          0
0.000000      0.000000      0.000000
[ 2065.319754]   kblockd/3  57   874163.510611    36 115          0          0
0.000000      0.000000      0.000000
[ 2065.319754] R sendmail-mta 4222 1430610.256048    4 120          0          0
0.000000      0.000000      0.000000
[ 2065.319754]

```

```

[ 2574.805414] Sched Debug Version: v0.07, 2.6.26-2-openvz-686 #1
[ 2574.805414] now at 2441403.983109 msecs
[ 2574.805414] .sysctl_sched_latency      : 60.000000
[ 2574.805414] .sysctl_sched_min_granularity : 12.000000

```

```

[ 2574.805414] .sysctl_sched_wakeup_granularity      : 30.000000
[ 2574.805414] .sysctl_sched_child_runs_first                : 0.000001
[ 2574.805414] .sysctl_sched_features                        : 895
[ 2574.805414]
[ 2574.805414] cpu#0, 2327.615 MHz
[ 2574.805414] .nr_running                                  : 2
[ 2574.805414] .load                                       : 2048
[ 2574.805414] .nr_switches                               : 237719
[ 2574.805414] .nr_load_updates                           : 495313
[ 2574.805414] .nr_uninterruptible                        : 17
[ 2574.805414] .jiffies                                    : 535351
[ 2574.805414] .next_balance                              : 0.535351
[ 2574.805414] .curr->pid                                  : 4097
[ 2574.805414] .clock                                       : 2441404.006918
[ 2574.805414] .cpu_load[0]                               : 0
[ 2574.805414] .cpu_load[1]                               : 0
[ 2574.805414] .cpu_load[2]                               : 0
[ 2574.805414] .cpu_load[3]                               : 0
[ 2574.805414] .cpu_load[4]                               : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[0]:
[ 2574.805414] .exec_clock                                 : 0.000000
[ 2574.805414] .MIN_vruntime                              : 0.000001
[ 2574.805414] .min_vruntime                              : 25591.135747
[ 2574.805414] .max_vruntime                              : 0.000001
[ 2574.805414] .spread                                     : 0.000000
[ 2574.805414] .spread0                                    : 0.000000
[ 2574.805414] .nr_running                                : 0
[ 2574.805414] .load                                       : 0
[ 2574.805414] .nr_spread_over                            : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[0]:
[ 2574.805414] .exec_clock                                 : 0.000000
[ 2574.805414] .MIN_vruntime                              : 25531.135781
[ 2574.805414] .min_vruntime                              : 25591.135747
[ 2574.805414] .max_vruntime                              : 25531.135781
[ 2574.805414] .spread                                     : 0.000000
[ 2574.805414] .spread0                                    : 0.000000
[ 2574.805414] .nr_running                                : 2
[ 2574.805414] .load                                       : 2048
[ 2574.805414] .nr_spread_over                            : 0
[ 2574.805414]
[ 2574.805414] runnable tasks:
[ 2574.805414]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2574.805414] -----
[ 2574.805414]      rsyslogd 4363  25531.135781    5 120          0          0
0.000000      0.000000      0.000000

```

```

[ 2574.805414] R      bash 4097   25531.135762   262 120         0         0
0.000000      0.000000      0.000000
[ 2574.805414]
[ 2574.805414] cpu#1, 2327.615 MHz
[ 2574.805414] .nr_running      : 1
[ 2574.805414] .load            : 1024
[ 2574.805414] .nr_switches    : 229295
[ 2574.805414] .nr_load_updates : 478238
[ 2574.805414] .nr_uninterruptible : 2262
[ 2574.805414] .jiffies        : 535351
[ 2574.805414] .next_balance   : 0.535352
[ 2574.805414] .curr->pid      : 2754
[ 2574.805414] .clock          : 2441404.506646
[ 2574.805414] .cpu_load[0]   : 0
[ 2574.805414] .cpu_load[1]   : 0
[ 2574.805414] .cpu_load[2]   : 0
[ 2574.805414] .cpu_load[3]   : 0
[ 2574.805414] .cpu_load[4]   : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[1]:
[ 2574.805414] .exec_clock     : 0.000000
[ 2574.805414] .MIN_vruntime   : 0.000001
[ 2574.805414] .min_vruntime   : 36700.432858
[ 2574.805414] .max_vruntime   : 0.000001
[ 2574.805414] .spread        : 0.000000
[ 2574.805414] .spread0       : 11109.297111
[ 2574.805414] .nr_running     : 0
[ 2574.805414] .load          : 0
[ 2574.805414] .nr_spread_over : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[1]:
[ 2574.805414] .exec_clock     : 0.000000
[ 2574.805414] .MIN_vruntime   : 0.000001
[ 2574.805414] .min_vruntime   : 36700.432858
[ 2574.805414] .max_vruntime   : 0.000001
[ 2574.805414] .spread        : 0.000000
[ 2574.805414] .spread0       : 11109.297111
[ 2574.805414] .nr_running     : 1
[ 2574.805414] .load          : 1024
[ 2574.805414] .nr_spread_over : 0
[ 2574.805414]
[ 2574.805414] runnable tasks:
[ 2574.805414]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2574.805414] -----
[ 2574.805414] R      rsyslogd 2754   36640.432859   710 120         0         0
0.000000      0.000000      0.000000
[ 2574.805414]

```

```

[ 2574.805414] cpu#2, 2327.615 MHz
[ 2574.805414] .nr_running      : 0
[ 2574.805414] .load                : 0
[ 2574.805414] .nr_switches        : 180593
[ 2574.805414] .nr_load_updates    : 485984
[ 2574.805414] .nr_uninterruptible : 4294964955
[ 2574.805414] .jiffies            : 535351
[ 2574.805414] .next_balance       : 0.535351
[ 2574.805414] .curr->pid          : 0
[ 2574.805414] .clock              : 2441401.001475
[ 2574.805414] .cpu_load[0]        : 0
[ 2574.805414] .cpu_load[1]        : 0
[ 2574.805414] .cpu_load[2]        : 0
[ 2574.805414] .cpu_load[3]        : 0
[ 2574.805414] .cpu_load[4]        : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[2]:
[ 2574.805414] .exec_clock         : 0.000000
[ 2574.805414] .MIN_vruntime       : 0.000001
[ 2574.805414] .min_vruntime       : 15350.377933
[ 2574.805414] .max_vruntime       : 0.000001
[ 2574.805414] .spread             : 0.000000
[ 2574.805414] .spread0            : -10240.757814
[ 2574.805414] .nr_running         : 0
[ 2574.805414] .load               : 0
[ 2574.805414] .nr_spread_over     : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[2]:
[ 2574.805414] .exec_clock         : 0.000000
[ 2574.805414] .MIN_vruntime       : 0.000001
[ 2574.805414] .min_vruntime       : 15350.377933
[ 2574.805414] .max_vruntime       : 0.000001
[ 2574.805414] .spread             : 0.000000
[ 2574.805414] .spread0            : -10240.757814
[ 2574.805414] .nr_running         : 0
[ 2574.805414] .load               : 0
[ 2574.805414] .nr_spread_over     : 0
[ 2574.805414]
[ 2574.805414] runnable tasks:
[ 2574.805414]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2574.805414] -----
[ 2574.805414]
[ 2574.805414] cpu#3, 2327.615 MHz
[ 2574.805414] .nr_running      : 4
[ 2574.805414] .load            : 184788
[ 2574.805414] .nr_switches    : 35136
[ 2574.805414] .nr_load_updates : 488898

```

```

[ 2574.805414] .nr_uninterruptible      : 66
[ 2574.805414] .jiffies                : 535351
[ 2574.805414] .next_balance           : 0.535356
[ 2574.805414] .curr->pid              : 4222
[ 2574.805414] .clock                  : 2441405.459833
[ 2574.805414] .cpu_load[0]           : 184788
[ 2574.805414] .cpu_load[1]           : 184788
[ 2574.805414] .cpu_load[2]           : 184788
[ 2574.805414] .cpu_load[3]           : 184788
[ 2574.805414] .cpu_load[4]           : 184788
[ 2574.805414]
[ 2574.805414] cfs_rq[3]:
[ 2574.805414] .exec_clock            : 0.000000
[ 2574.805414] .MIN_vruntime          : 0.000001
[ 2574.805414] .min_vruntime          : 874283.510611
[ 2574.805414] .max_vruntime          : 0.000001
[ 2574.805414] .spread                : 0.000000
[ 2574.805414] .spread0               : 848692.374864
[ 2574.805414] .nr_running            : 1
[ 2574.805414] .load                  : 1024
[ 2574.805414] .nr_spread_over        : 0
[ 2574.805414]
[ 2574.805414] cfs_rq[3]:
[ 2574.805414] .exec_clock            : 0.000000
[ 2574.805414] .MIN_vruntime          : 874163.510611
[ 2574.805414] .min_vruntime          : 874283.510611
[ 2574.805414] .max_vruntime          : 874223.510611
[ 2574.805414] .spread                : 60.000000
[ 2574.805414] .spread0               : 848692.374864
[ 2574.805414] .nr_running            : 3
[ 2574.805414] .load                  : 7266
[ 2574.805414] .nr_spread_over        : 0
[ 2574.805414]
[ 2574.805414] runnable tasks:
[ 2574.805414]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2574.805414] -----
[ 2574.805414] watchdog/3  14    -4.329575    3  0          0          0
0.000000    0.000000    0.000000
[ 2574.805414] events/3    18    874223.510611  738 115         0          0
0.000000    0.000000    0.000000
[ 2574.805414] kblockd/3   57    874163.510611   36 115         0          0
0.000000    0.000000    0.000000
[ 2574.805414] R sendmail-mta 4222 1939299.790290   4 120         0          0
0.000000    0.000000

```


Sched Debug Version: v0.07, 2.6.26-2-openvz-686 #1

```
[ 2612.520474] now at 2479099.757691 msecs
[ 2612.520474] .sysctl_sched_latency           : 60.000000
[ 2612.520474] .sysctl_sched_min_granularity             : 12.000000
[ 2612.520474] .sysctl_sched_wakeup_granularity         : 30.000000
[ 2612.520474] .sysctl_sched_child_runs_first          : 0.000001
[ 2612.520474] .sysctl_sched_features                   : 895
[ 2612.520474]
[ 2612.520474] cpu#0, 2327.615 MHz
[ 2612.520474] .nr_running                             : 1
[ 2612.520474] .load                                    : 1024
[ 2612.520474] .nr_switches                            : 239641
[ 2612.520474] .nr_load_updates                        : 504737
[ 2612.520474] .nr_uninterruptible                     : 17
[ 2612.520474] .jiffies                                 : 544775
[ 2612.520474] .next_balance                           : 0.544775
[ 2612.520474] .curr->pid                              : 4097
[ 2612.520474] .clock                                   : 2479100.006869
[ 2612.520474] .cpu_load[0]                            : 0
[ 2612.520474] .cpu_load[1]                            : 0
[ 2612.520474] .cpu_load[2]                            : 32
[ 2612.520474] .cpu_load[3]                            : 784
[ 2612.520474] .cpu_load[4]                            : 2203
[ 2612.520474]
[ 2612.520474] cfs_rq[0]:
[ 2612.520474] .exec_clock                             : 0.000000
[ 2612.520474] .MIN_vruntime                           : 0.000001
[ 2612.520474] .min_vruntime                           : 25711.135749
[ 2612.520474] .max_vruntime                           : 0.000001
[ 2612.520474] .spread                                  : 0.000000
[ 2612.520474] .spread0                                 : 0.000000
[ 2612.520474] .nr_running                             : 0
[ 2612.520474] .load                                    : 0
[ 2612.520474] .nr_spread_over                         : 0
[ 2612.520474]
[ 2612.520474] cfs_rq[0]:
[ 2612.520474] .exec_clock                             : 0.000000
[ 2612.520474] .MIN_vruntime                           : 0.000001
[ 2612.520474] .min_vruntime                           : 25711.135749
[ 2612.520474] .max_vruntime                           : 0.000001
[ 2612.520474] .spread                                  : 0.000000
[ 2612.520474] .spread0                                 : 0.000000
[ 2612.520474] .nr_running                             : 1
[ 2612.520474] .load                                    : 1024
[ 2612.520474] .nr_spread_over                         : 0
[ 2612.520474]
[ 2612.520474] runnable tasks:
[ 2612.520474]      task  PID      tree-key  switches  prio  exec-runtime  sum-exec
```

sum-sleep

```
[ 2612.520474] -----  
[ 2612.520474] R      bash 4097  25651.135750  264 120      0      0  
0.000000      0.000000      0.000000  
[ 2612.520474]  
[ 2612.520474] cpu#1, 2327.615 MHz  
[ 2612.520474] .nr_running      : 2  
[ 2612.520474] .load            : 2048  
[ 2612.520474] .nr_switches     : 229920  
[ 2612.520474] .nr_load_updates : 487662  
[ 2612.520474] .nr_uninterruptible : 2264  
[ 2612.520474] .jiffies         : 544775  
[ 2612.520474] .next_balance    : 0.544776  
[ 2612.520474] .curr->pid       : 2754  
[ 2612.520474] .clock           : 2479100.506684  
[ 2612.520474] .cpu_load[0]     : 0  
[ 2612.520474] .cpu_load[1]     : 0  
[ 2612.520474] .cpu_load[2]     : 32  
[ 2612.520474] .cpu_load[3]     : 784  
[ 2612.520474] .cpu_load[4]     : 2203  
[ 2612.520474]  
[ 2612.520474] cfs_rq[1]:  
[ 2612.520474] .exec_clock      : 0.000000  
[ 2612.520474] .MIN_vruntime    : 0.000001  
[ 2612.520474] .min_vruntime    : 36700.811092  
[ 2612.520474] .max_vruntime    : 0.000001  
[ 2612.520474] .spread          : 0.000000  
[ 2612.520474] .spread0         : 10989.675343  
[ 2612.520474] .nr_running      : 0  
[ 2612.520474] .load            : 0  
[ 2612.520474] .nr_spread_over  : 0  
[ 2612.520474]  
[ 2612.520474] cfs_rq[1]:  
[ 2612.520474] .exec_clock      : 0.000000  
[ 2612.520474] .MIN_vruntime    : 36644.471365  
[ 2612.520474] .min_vruntime    : 36700.811092  
[ 2612.520474] .max_vruntime    : 36644.471365  
[ 2612.520474] .spread          : 0.000000  
[ 2612.520474] .spread0         : 10989.675343  
[ 2612.520474] .nr_running      : 2  
[ 2612.520474] .load            : 2048  
[ 2612.520474] .nr_spread_over  : 0  
[ 2612.520474]  
[ 2612.520474] runnable tasks:  
[ 2612.520474]      task PID      tree-key switches prio  exec-runtime  sum-exec  
sum-sleep  
[ 2612.520474] -----  
[ 2612.520474] R      rsyslogd 2754  36644.427551  711 120      0      0
```

```

0.000000      0.000000      0.000000
[ 2612.520474] rsyslogd 4363 36644.471365 7 120 0 0
0.000000      0.000000      0.000000
[ 2612.520474]
[ 2612.520474] cpu#2, 2327.615 MHz
[ 2612.520474] .nr_running      : 0
[ 2612.520474] .load            : 0
[ 2612.520474] .nr_switches    : 181276
[ 2612.520474] .nr_load_updates : 495408
[ 2612.520474] .nr_uninterruptible : 4294964953
[ 2612.520474] .jiffies        : 544775
[ 2612.520474] .next_balance   : 0.544775
[ 2612.520474] .curr->pid      : 0
[ 2612.520474] .clock          : 2479101.001347
[ 2612.520474] .cpu_load[0]    : 0
[ 2612.520474] .cpu_load[1]    : 0
[ 2612.520474] .cpu_load[2]    : 32
[ 2612.520474] .cpu_load[3]    : 784
[ 2612.520474] .cpu_load[4]    : 2203
[ 2612.520474]
[ 2612.520474] cfs_rq[2]:
[ 2612.520474] .exec_clock     : 0.000000
[ 2612.520474] .MIN_vruntime   : 0.000001
[ 2612.520474] .min_vruntime   : 15350.377933
[ 2612.520474] .max_vruntime   : 0.000001
[ 2612.520474] .spread        : 0.000000
[ 2612.520474] .spread0       : -10360.757816
[ 2612.520474] .nr_running     : 0
[ 2612.520474] .load          : 0
[ 2612.520474] .nr_spread_over : 0
[ 2612.520474]
[ 2612.520474] cfs_rq[2]:
[ 2612.520474] .exec_clock     : 0.000000
[ 2612.520474] .MIN_vruntime   : 0.000001
[ 2612.520474] .min_vruntime   : 15350.377933
[ 2612.520474] .max_vruntime   : 0.000001
[ 2612.520474] .spread        : 0.000000
[ 2612.520474] .spread0       : -10360.757816
[ 2612.520474] .nr_running     : 0
[ 2612.520474] .load          : 0
[ 2612.520474] .nr_spread_over : 0
[ 2612.520474]
[ 2612.520474] runnable tasks:
[ 2612.520474]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2612.520474] -----
[ 2612.520474]
[ 2612.520474] cpu#3, 2327.615 MHz

```

```

[ 2612.520474] .nr_running          : 4
[ 2612.520474] .load                : 184788
[ 2612.520474] .nr_switches        : 35136
[ 2612.520474] .nr_load_updates    : 498322
[ 2612.520474] .nr_uninterruptible : 66
[ 2612.520474] .jiffies            : 544775
[ 2612.520474] .next_balance       : 0.544828
[ 2612.520474] .curr->pid          : 4222
[ 2612.520474] .clock              : 2479101.502575
[ 2612.520474] .cpu_load[0]        : 184788
[ 2612.520474] .cpu_load[1]        : 184788
[ 2612.520474] .cpu_load[2]        : 184788
[ 2612.520474] .cpu_load[3]        : 184788
[ 2612.520474] .cpu_load[4]        : 184788
[ 2612.520474]
[ 2612.520474] cfs_rq[3]:
[ 2612.520474] .exec_clock         : 0.000000
[ 2612.520474] .MIN_vruntime       : 0.000001
[ 2612.520474] .min_vruntime       : 874283.510611
[ 2612.520474] .max_vruntime       : 0.000001
[ 2612.520474] .spread             : 0.000000
[ 2612.520474] .spread0            : 848572.374862
[ 2612.520474] .nr_running         : 1
[ 2612.520474] .load               : 1024
[ 2612.520474] .nr_spread_over     : 0
[ 2612.520474]
[ 2612.520474] cfs_rq[3]:
[ 2612.520474] .exec_clock         : 0.000000
[ 2612.520474] .MIN_vruntime       : 874163.510611
[ 2612.520474] .min_vruntime       : 874283.510611
[ 2612.520474] .max_vruntime       : 874223.510611
[ 2612.520474] .spread             : 60.000000
[ 2612.520474] .spread0            : 848572.374862
[ 2612.520474] .nr_running         : 3
[ 2612.520474] .load               : 7266
[ 2612.520474] .nr_spread_over     : 0
[ 2612.520474]
[ 2612.520474] runnable tasks:
[ 2612.520474]      task PID      tree-key switches prio  exec-runtime  sum-exec
sum-sleep
[ 2612.520474] -----
[ 2612.520474]   watchdog/3  14    -4.329575     3  0         0         0
0.000000    0.000000    0.000000
[ 2612.520474]   events/3   18  874223.510611   738 115         0         0
0.000000    0.000000    0.000000
[ 2612.520474]   kblockd/3  57  874163.510611   36 115         0         0
0.000000    0.000000    0.000000
[ 2612.520474] R sendmail-mta 4222 1976995.797668    4 120         0         0

```

0.000000

0.000000

Sysrq-t gives the same output as sysrq-w.

Does this help in any way?

Thanks for your help!

Subject: Re: IBM RAID controller re-mount lock
Posted by [maratrus](#) on Thu, 13 Aug 2009 07:55:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thank you very much for the provided information and sorry for delay.
I'd recommend you to file a bug into bugzilla.
<http://bugzilla.openvz.org/>
