
Subject: [migration/0] 100% CPU

Posted by [berlo](#) on Fri, 19 Aug 2011 19:14:44 GMT

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

hi,

some days ago i moved some my nodes to CentOS 6. After migration i see that centos 5 node still stable, but migrated centos 6 module random crash. i have some screen during the crash. I do not think that is something in cron because crash are randon and in different node.

All nodes went offline with this status:

```
top - 19:26:07 up 1 day, 5:35, 2 users, load average: 10.59, 2.90, 1.30
Tasks: 985 total, 5 running, 976 sleeping, 0 stopped, 4 zombie
Cpu(s): 0.7%us, 26.0%sy, 0.0%ni, 0.0%id, 73.2%wa, 0.0%hi, 0.1%si, 0.0%st
Mem: 8296140k total, 6580416k used, 1715724k free, 460436k buffers
Swap: 10403832k total, 14496k used, 10389336k free, 3686296k cached
```

```
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
 3 root RT 0 0 0 0 S 100.0 0.0 1:10.82 [migration/0]
28697 root 20 0 3184 1716 888 R 1.6 0.0 5:55.67 top -c
16168 root 20 0 3172 1752 888 S 1.3 0.0 14:04.70 top -c
 9766 105 20 0 364m 18m 4940 S 0.6 0.2 10:01.93 /usr/bin/python
/usr/lib/checker/Checker.py -c /etc/checker.conf
21873 root 20 0 74848 32m 7264 S 0.6 0.4 1:20.45 ./hlds_i686 -console -game cstrike
-master -secure -pingboost 3 +ip 213.92.118.171 +sys_ticrate 20 -heapsize 100 +exec
xhlds1.cfg?
12217 apache 20 0 35444 1916 592 S 0.3 0.0 4:59.34
/var/www/html/cast/files/linux/sc_serv temp/8000_1313668473.conf
12580 apache 20 0 42144 1780 564 S 0.3 0.0 4:58.78
/var/www/html/cast/files/linux/sc_serv temp/8002_1313668474.conf
12662 apache 20 0 38704 2132 752 S 0.3 0.0 5:13.15
/var/www/html/cast/files/linux/sc_serv temp/8004_1313668479.conf
13124 33 20 0 40600 11m 3920 D 0.3 0.1 0:01.80 /usr/sbin/apache2 -k start
13874 apache 20 0 35444 1828 564 S 0.3 0.0 4:58.83
/var/www/html/cast/files/linux/sc_serv temp/8006_1313668490.conf
14201 33 20 0 39832 11m 3940 D 0.3 0.1 0:01.20 /usr/sbin/apache2 -k start
14382 apache 20 0 35944 2156 756 S 0.3 0.0 5:17.00
/var/www/html/cast/files/linux/sc_serv temp/8010_1313668497.conf
15512 apache 20 0 35944 2252 752 S 0.3 0.0 5:30.24
/var/www/html/cast/files/linux/sc_serv temp/8008_1313668517.conf
15907 65534 20 0 1173m 45m 44m D 0.3 0.6 0:30.02 /usr/sbin/varnishd -P
/var/run/varnishd.pid -a :80 -T localhost:6082 -f /etc/varnish/default.vcl -S /etc/varnish/secret -s
file,/var/lib/varnish/vps-it/varnish_storage.bin,1G
21882 root 20 0 74848 32m 7268 S 0.3 0.4 1:19.77 ./hlds_i686 -console -game cstrike
-master -secure -pingboost 3 +ip 213.92.118.171 +sys_ticrate 20 -heapsize 100 +exec xhlds3.cfg
21883 root 20 0 74848 32m 7348 S 0.3 0.4 1:19.73 ./hlds_i686 -console -game cstrike
-master -secure -pingboost 3 +ip 213.92.118.171 +sys_ticrate 20 -heapsize 100 +exec
```

xhlds2.cfg?

```
1 root    20  0 2828 1356 1192 S 0.0 0.0 0:01.05 /sbin/init
2 root    20  0  0  0  0 S 0.0 0.0 0:00.00 [kthreadd]
4 root    20  0  0  0  0 R 0.0 0.0 0:00.57 [ksoftirqd/0]
5 root    RT  0  0  0  0 S 0.0 0.0 0:00.00 [migration/0]
6 root    RT  0  0  0  0 R 0.0 0.0 0:00.04 [watchdog/0]
7 root    RT  0  0  0  0 S 0.0 0.0 0:00.42 [migration/1]
8 root    RT  0  0  0  0 S 0.0 0.0 0:00.00 [migration/1]
9 root    20  0  0  0  0 S 0.0 0.0 0:01.69 [ksoftirqd/1]
10 root   RT  0  0  0  0 S 0.0 0.0 0:00.04 [watchdog/1]
11 root   RT  0  0  0  0 S 0.0 0.0 0:01.03 [migration/2]
12 root   RT  0  0  0  0 S 0.0 0.0 0:00.00 [migration/2]
13 root   20  0  0  0  0 S 0.0 0.0 0:01.72 [ksoftirqd/2]
14 root   RT  0  0  0  0 S 0.0 0.0 0:00.05 [watchdog/2]
15 root   RT  0  0  0  0 S 0.0 0.0 0:01.01 [migration/3]
16 root   RT  0  0  0  0 S 0.0 0.0 0:00.00 [migration/3]
17 root   20  0  0  0  0 S 0.0 0.0 0:01.68 [ksoftirqd/3]
18 root   RT  0  0  0  0 S 0.0 0.0 0:00.04 [watchdog/3]
19 root   20  0  0  0  0 R 0.0 0.0 0:00.24 [events/0]
20 root   20  0  0  0  0 S 0.0 0.0 0:09.13 [events/1]
21 root   20  0  0  0  0 S 0.0 0.0 0:01.79 [events/2]
```

i know that process [migration/0] is a kernel thread that move thread between cpu's but i don't know what cause this situation.

Anyone had this problem or know how to solve or debug it?

configuration:

```
# uname -a
```

```
Linux node82 2.6.32-042stab031.1 #1 SMP Fri Aug 12 21:21:55 MSD 2011 i686 i686 i386
GNU/Linux
```

```
# lspci
```

```
00:00.0 Host bridge: Intel Corporation 5000X Chipset Memory Controller Hub (rev 12)
00:02.0 PCI bridge: Intel Corporation 5000 Series Chipset PCI Express x4 Port 2 (rev 12)
00:03.0 PCI bridge: Intel Corporation 5000 Series Chipset PCI Express x4 Port 3 (rev 12)
00:04.0 PCI bridge: Intel Corporation 5000 Series Chipset PCI Express x8 Port 4-5 (rev 12)
00:05.0 PCI bridge: Intel Corporation 5000 Series Chipset PCI Express x4 Port 5 (rev 12)
00:06.0 PCI bridge: Intel Corporation 5000 Series Chipset PCI Express x8 Port 6-7 (rev 12)
00:07.0 PCI bridge: Intel Corporation 5000 Series Chipset PCI Express x4 Port 7 (rev 12)
00:10.0 Host bridge: Intel Corporation 5000 Series Chipset FSB Registers (rev 12)
00:10.1 Host bridge: Intel Corporation 5000 Series Chipset FSB Registers (rev 12)
00:10.2 Host bridge: Intel Corporation 5000 Series Chipset FSB Registers (rev 12)
00:11.0 Host bridge: Intel Corporation 5000 Series Chipset Reserved Registers (rev 12)
```

00:13.0 Host bridge: Intel Corporation 5000 Series Chipset Reserved Registers (rev 12)
00:15.0 Host bridge: Intel Corporation 5000 Series Chipset FBD Registers (rev 12)
00:16.0 Host bridge: Intel Corporation 5000 Series Chipset FBD Registers (rev 12)
00:1c.0 PCI bridge: Intel Corporation 631xESB/632xESB/3100 Chipset PCI Express Root Port 1 (rev 09)
00:1d.0 USB Controller: Intel Corporation 631xESB/632xESB/3100 Chipset UHCI USB Controller #1 (rev 09)
00:1d.1 USB Controller: Intel Corporation 631xESB/632xESB/3100 Chipset UHCI USB Controller #2 (rev 09)
00:1d.2 USB Controller: Intel Corporation 631xESB/632xESB/3100 Chipset UHCI USB Controller #3 (rev 09)
00:1d.3 USB Controller: Intel Corporation 631xESB/632xESB/3100 Chipset UHCI USB Controller #4 (rev 09)
00:1d.7 USB Controller: Intel Corporation 631xESB/632xESB/3100 Chipset EHCI USB2 Controller (rev 09)
00:1e.0 PCI bridge: Intel Corporation 82801 PCI Bridge (rev d9)
00:1f.0 ISA bridge: Intel Corporation 631xESB/632xESB/3100 Chipset LPC Interface Controller (rev 09)
00:1f.1 IDE interface: Intel Corporation 631xESB/632xESB IDE Controller (rev 09)
01:00.0 RAID bus controller: LSI Logic / Symbios Logic MegaRAID SAS 1078 (rev 04)
02:00.0 PCI bridge: Broadcom EPB PCI-Express to PCI-X Bridge (rev c3)
03:00.0 Ethernet controller: Broadcom Corporation NetXtreme II BCM5708 Gigabit Ethernet (rev 12)
04:00.0 PCI bridge: Intel Corporation 6311ESB/6321ESB PCI Express Upstream Port (rev 01)
04:00.3 PCI bridge: Intel Corporation 6311ESB/6321ESB PCI Express to PCI-X Bridge (rev 01)
05:00.0 PCI bridge: Intel Corporation 6311ESB/6321ESB PCI Express Downstream Port E1 (rev 01)
05:01.0 PCI bridge: Intel Corporation 6311ESB/6321ESB PCI Express Downstream Port E2 (rev 01)
06:00.0 PCI bridge: Broadcom EPB PCI-Express to PCI-X Bridge (rev c3)
07:00.0 Ethernet controller: Broadcom Corporation NetXtreme II BCM5708 Gigabit Ethernet (rev 12)
0e:0d.0 VGA compatible controller: ATI Technologies Inc ES1000 (rev 02)

cat /etc/redhat-release
CentOS Linux release 6.0 (Final)

thank you

Subject: Re: [migration/0] 100% CPU
Posted by [bjdea1](#) on Tue, 23 Aug 2011 14:01:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

I think this might be the similar to what I have seen on my server node with centos 6

2.6.32-042stab031.1 kernel when the server crashed twice in about 2 weeks. Same mention of migration issue in /var/log/messages with process getting stuck in single CPU for 67 seconds and causes crash.

kernel: [633492.036001] BUG: soft lockup - CPU#0 stuck for 67s! [migration/0:3]

Subject: Re: [migration/0] 100% CPU

Posted by [berlo](#) on Tue, 23 Aug 2011 14:04:41 GMT

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

hi,

yes is same issue. I opened a bug request and received this:

```
# echo 0 > /proc/sys/kernel/sched_cpulimit_nr_balance
```

Now are two days that i do not have crash. Usually server crash every 1-2 days, but i will wait some days to mark as resolved.

Have you tried this solution?

Subject: Re: [migration/0] 100% CPU

Posted by [bjdea1](#) on Tue, 23 Aug 2011 14:26:01 GMT

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

Ok I am trying this, hope it works, thanks.

Subject: Re: [migration/0] 100% CPU

Posted by [deziweb](#) on Fri, 26 Aug 2011 12:07:59 GMT

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

We have experienced the same issue after upgrading some hosts to CentOS 6: a proces migration/0 eating up 100% CPU. Load rises within 15 minutes from 2 to 300+ and eventually the system crashes.

I've now set the proc/sys/kernel/sched_cpulimit_nr_balance to 0.

Does anyone know if this resolves the issue?

Subject: Re: [migration/0] 100% CPU

Posted by [berlo](#) on Mon, 29 Aug 2011 10:37:49 GMT

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

hi,
is a workaround, but this operation solve the issue.

Developer proposed a patch, i tested it and seems working great. Read here: http://bugzilla.openvz.org/show_bug.cgi?id=1954

Regards

Subject: Re: [migration/0] 100% CPU
Posted by [deziweb](#) on Thu, 15 Sep 2011 09:38:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

Since we've changed it manually by setting it to 0 it has worked fine for us as well. However, I'm still running Kernel RHEL6 042stab033.1 and when I do a reboot of the server, it's set back to 4 again, which will cause it crash again in a very short time.

I've now upgraded to kernel vzkernel-2.6.32-042stab036.6. but I'm not sure if it's corrected already in this kernel. Can somebody confirm this please?

Thank you.
