
Subject: Re: Is there a stable OpenVZ kernel, and which should be fit for production

Posted by [Dariush Pietrzak](#) on Thu, 24 Nov 2011 12:15:39 GMT

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

> Just out of curiosity i use my kernel crash-test setup to test with
> "stress" and "bonnie". I simply use the OpenVZ-Kernel with two
> container (ubuntu-10.04) and let one run stress and the other
> bonnie. The load is at 15 but the machine is humming along since
> around 4 hours...

With such low load we also couldn't crash it in timely matter.
With lightly loaded machines we endured months without crash.

I use this:

```
stress -c 22 -i 24 -m 8 -d 20 --hdd-bytes 10G
```

and this:

```
while (true)
```

```
do
```

```
bonnie++ -d /fs/v/bonnie/ -c 8 -b -f -u root
```

```
echo next
```

```
done
```

in parallel, I don't even have to run it inside containers.

(test machine is single 4-core Xeon E5320, with 4G ram and two 146G raid 1s joined by lvm. With loadavg 50-80 we get crashes after few hours).

> Is it possible that your problem arise from the io devices used?

Possible, but unlikely, we first noticed crashed using FC devices, and then moved to testing on small P400i with 256M ram. One of the most affected machines used P410i controller, which is very similiar and the same generation as P400i.

I can re-test on FC again.

And while IO load seems to be neccessary to cause crash, resulting oops-es are similiar, very often account_system_time appears:

```
[38766.228063] panic occurred, switching back to text console
```

```
[38766.228063] BUG: scheduling while atomic: stress/1962/0x10000100
```

(this is identical to what we saw in production, only with 'java' instead of 'stress')

```
[38766.227505] BUG: unable to handle kernel paging request at 0000000000021300
```

```
[38766.227509] IP: [<ffffff81050ec4>] update_curr+0x154/0x200
```

```
[38766.227514] PGD 12c7b4067 PUD 12c7b5067 PMD 0
```

```
[38764.623677] BUG: unable to handle kernel paging request at 000000000001e440
```

```
[38764.623677] IP: [<ffffff814c8efe>] _spin_lock+0xe/0x30
```

[38764.599189] BUG: unable to handle kernel paging request at 0000000000019550

[38764.599189] IP: [<ffffffff8105674f>] account_system_time+0xaf/0x1f0

[1876.747809] BUG: unable to handle kernel paging request at 00000006000000bd

[1876.747815] IP: [<ffffffff8105a4fe>] select_task_rq_fair+0x32e/0xa20

[1515.270063] BUG: unable to handle kernel paging request at 00000004047118e0

[1515.270063] IP: [<ffffffff81050aad>] task_rq_lock+0x4d/0xa0

best regards, Eyck

--

Key fingerprint = 40D0 9FFB 9939 7320 8294 05E0 BCC7 02C4 75CC 50D9

Total Existence Failure
