Subject: *NOT SOLVED* 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable

Posted by tchipman on Mon, 29 Oct 2007 14:13:55 GMT

View Forum Message <> Reply to Message

Hi,

Just want to report this "open issue" I have got, with some slight debug progress. If I have more notes (later today or tomorrow) I'll update this thread in case it is of interest to others. Not positive this is a "bug report" yet.

I have a 32-bit dual xeon system (Tyan server-class motherboard) with a 64-bit PCI slot 3ware Raid card (model 8506, takes 8 x SATA disks)

There appears to be some problem with this 3ware card with certain kernel versions, ie,

(1) stock RHEL/CentOS 5.0 kernel is BAD to use on this card, flat out (you end up with disk failure, dmesg errors logged abundantly after significant disk IO, approx as follows:

3w-xxxx: scsi0: Character ioctl (0x1f) timed out, resetting card.

sd 0:0:0:0: rejecting I/O to offline device

the raid array then flags itself as "rebuild from scratch" state and you lose everything in the array

- (2) stock RHEL/CentOS 4.5 which is booted without "noapic" and "acpi=off" kernel flags -- also is unstable. By adding these flags, it becomes stable.
- (3) OpenVZ stock production OpenVZ kernel booted for either CentOS 5 or 4.5 install .. will be unstable, regardless of use of the noapci/acpi=off flags. kernel used for test identified itself as "vmlinuz-2.6.9-023stab044.11-smp"

This machine has been reinstalled about 6 times in the last week for testing. Managed to run solid over the weekend it seems and I'm doing one last cycle of sustained disk activity (500 runs of bonnie) -- if it survives this, then I'm willing to believe that scenarion (2) above is OK (stock CentOS 4.5 kernel with the acpi-noapic flags)

Assuming it survives, I'll then try install of OpenVZ "RHEL Stock" kernel rather than the standard OVZ production kernel, and see if that runs more stable, or if that is also a lost cause.

I suspect based on observations so far,

- -there is something new in RHEL5 which breaks 3ware controller hardware
- -this feature may also be present in the default build of the OVZ kernel for CentOS 4.X lineage as well?

For anyone who wishes to see, there is a thread open in the CentOS bug-tracker on this 3ware hardware issue. The URL is,

http://bugs.centos.org/view.php?id=2186

(I have been posting to that thread, as have a number of other 3ware customers using CentOS 4.X or 5.X)

Lots of fun,

Tim Chipman

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by vaverin on Thu, 15 Nov 2007 13:28:11 GMT

View Forum Message <> Reply to Message

Hello Tim,

Unfortunately it is still unclear what happened on your node. Could you please provide us the serial console logs?

thank you, Vasily Averin

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by tchipman on Sat, 17 Nov 2007 13:39:28 GMT

View Forum Message <> Reply to Message

Hi,

If you look at the URL,

http://bugs.centos.org/view.php?id=2186

you can see the type of error messages that are visible in dmesg/messages logfile when the problem takes place.

Again, to reiterate, this problem appears to be "normal" for the stock RHEL5.0 kernel, both 32 and 64 bit. Possibly it is how the default 3ware driver works on the newer kernel?

Alas I haven't got any further material to provide; the machine is now running stable since a few weeks, using CentOS 4.5 vanilla install with noapic and acpi=off options passed to kernel at boot. (and it is in semi-production now, so I'm not at liberty to reboot / crash the machine with impunity)

I was not able to get a (recent) openvz kernel to run stable on this machine.

Thanks for the followup,

Tim

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by vaverin on Mon, 19 Nov 2007 07:53:09 GMT

View Forum Message <> Reply to Message

Hi Tim,

I'm sorry but information in CentOs bug is not enough for investigation. I need to know exact kernel version (currently I do not understand is it x86 or x86_64?) and see all kerenel messages.

Could you please boot our latest kernel, collect all kernel messages by using remote console and send it to us?

Because we have disk-related failures, the error messages are visible in dmesg only and cannot be saved on the local disk. Therefore to collect error log messages it's required to configure remote console:

http://wiki.openvz.org/Remote_console_setup

Then it would be great to check latest vendor kernels (2.6.9-63 for RHEL4 and 2.6.18-53 for RHEL5). If you still have some troubles with these kernels -- I can help you to collect all required information and send bugreport to Red Hat and 3Ware.

Thank you, Vasily Averin

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by tchipman on Mon, 26 Nov 2007 16:09:34 GMT View Forum Message <> Reply to Message

Hi, just to followup / clarify,

I cannot any longer change / reboot this machine for testing. It is in production, and any tests will (1) cause downtime, (2) likely corrupt and destroy all data on the HDDs, resulting a full-reinstall, something I don't cherish given the number of reinstalls done to get the machine operational.

Just to give you more concrete reference info from my notes,

- the problem happend for certain with the kernel available in the openvz yum repo ~1 month ago, which was:

- the type of error message logged in dmesg was,

```
--paste--
scsi0 (2:0): rejecting I/O to offline device
EXT3-fs error (device sdb1): ext3 find entry: reading directory #110313495 offset 0
scsi0 (2:0): rejecting I/O to offline device
EXT3-fs error (device sdb1): ext3 find entry: reading directory #110313495 offset 0
scsi0 (2:0): rejecting I/O to offline device
EXT3-fs error (device sdb1): ext3 find entry: reading directory #110313492 offset 0
--endpaste--
```

I don't have any further exact notes to provide, unfortunately.

I did get this same general type of problem with all of the following,

- * CentOS 5.0 default install 32 bit or OpenVZ yum repo kernel
- * CentOS 4.5 default install 32-bit or OpenvZ yum repo kernel
- * the only current solid working setup I have on this hardware, is CentOS 4.5 32-bit default kernel, booted with "noapic" and "apci=off" kernel parameters passed via grub at boot time.

I haven't tested CentOS 5.0 with the kernel options being passed, as I had a decently working config for the machine, and I was sick and tired of re-re-re-installing this system... didn't feel like doing another pair of installs just to test "maybe if it works" for CentOS 5.

Sorry I cannot provide more info. Possibly / hopefully this is slightly of more use.

Most likely, if the problem is "common" then other folks with 3ware gear .. will report the problem soon enough .. and they can then provide "live and up to date" debug info. (?)

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by vaverin on Tue, 27 Nov 2007 05:28:21 GMT

View Forum Message <> Reply to Message

Hi Tim,

tchipman wrote on Mon, 26 November 2007 19:09 --paste-scsi0 (2:0): rejecting I/O to offline device scsi0 (2:0): rejecting I/O to offline device EXT3-fs error (device sdb1): ext3 find entry: reading directory #110313495 offset 0 scsi0 (2:0): rejecting I/O to offline device EXT3-fs error (device sdb1): ext3 find entry: reading directory #110313495 offset 0 scsi0 (2:0): rejecting I/O to offline device EXT3-fs error (device sdb1): ext3_find_entry: reading directory #110313492 offset 0

These messages shows the result of problem. We can see that disks are not work and any operations are rejected.

But these messages said nothing about causes of this problem.

Usually first error messages are most informative. To save these messages you can use remote console. Without such error messages we cannot investigate the problem.

Thank you, Vasily Averin

--endpaste--

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by tchipman on Tue, 04 Dec 2007 13:18:09 GMT

View Forum Message <> Reply to Message

Sorry that I didn't have enough info captured to help diagnose properly. As I indicated in previous message, the machine is in production now so I can't break it for testing any further. Additionally I know from previous breaks, it will corrupt all HDD contents (full reinstall from scratch) which I'm not in a position to do either.

So, this topic must be closed I think, and maybe if someone else has similar problem in future they will find this thread and proceed appropriately.

Thanks for your interest in this topic,

--Tim

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by mcluver on Fri, 21 Mar 2008 18:49:30 GMT

View Forum Message <> Reply to Message

Hi Tim,

I have not been receiving the exact same errors messages as you, but it sounds like I do have nearly the exact same hardware setup and configuration.

Here is what I have been getting quite frequently:

kernel: attempt to access beyond end of device

kernel: sda2: rw=0, want=25781269064, limit=2433349485

kernel: attempt to access beyond end of device

kernel: sda2: rw=0, want=25781269064, limit=2433349485

I am using the 3ware Escalade 8506-8 (8 x 250GB WD SATA Drives), along with a Tyan server grade motherboard (Tyan Tiger i7501 (S2723)), with CentOS 4.4 installed in an OpenVZ HA cluster (using DRBD between two exactly matching servers).

I have had them in production since February 2007 in this configuration. Since, I have had to remove and completely wipe and reinstall them several times because of inevitable partition table and data corruption.

I am currently in the process of testing out whether using Debian etch with this hardware and software combination is more stable. I see that your last post was in Dec. 2007, have you made any discoveries or headway since your last post?

Has anyone else been experiencing the effects of this horrible hardware combination?

Best regards,

Matthew Cluver

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by vaverin on Sat, 22 Mar 2008 09:00:57 GMT

View Forum Message <> Reply to Message

It looks like following issue:

http://www.3ware.com/KB/article.aspx?id=15243&cNode=6I1C 6S

We have fixed this issue already and I expect Kir will release new kernel soon.

thank you, Vasily Averin

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by mcluver on Mon, 24 Mar 2008 17:17:07 GMT

View Forum Message <> Reply to Message

Hi Vasily,

Unfortunately I am not running a 64-bit processor, nor more than 4GB of RAM, do you believe this bug fix will apply?

Regards,

Matt

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by vaverin on Tue, 25 Mar 2008 05:15:31 GMT

View Forum Message <> Reply to Message

Matt,

your system should not be affected and situation looks unclear. errors points to the memory corruption but its reasons are unclear for me. There I can give You the generall recommendation only.

Usually in such cases we recommend to re-check the memory on the node, but because Your system was in production I do not think that we really will find some memory defects.

Then I would like to advise you to upgrade BIOS on Your 3ware controller and on Your motherboard. -- it should exclude possible internal controller's and motherboard's troubles.

Next step -- IF you have Intel Cpu It makes sense to update microcode: (micorocode service can upload new firmware into your cpus, and fix already known CPU-related issues). Latest micocode version can be downloaded from Intel web site:

http://downloadcenter.intel.com/default.asp

If it will not fix issue on your node I need to get first kernel error messages, you can catch it by

using remote console: http://wiki.openvz.org/Remote console setup

Also it would be great to find some way to reproduce this issue.

thank you, Vasily Averin

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by TheWiseOne on Sat, 29 Mar 2008 17:22:00 GMT

View Forum Message <> Reply to Message

vaverin wrote on Sat, 22 March 2008 04:00lt looks like following issue: http://www.3ware.com/KB/article.aspx?id=15243&cNode=6l1C 6S

We have fixed this issue already and I expect Kir will release new kernel soon.

Any ETA on when a kernel with this fix will be released? I've experienced no issues, but as someone with 64-bit 4GB+ RAM and 3Ware 8000 series servers it is a bit stressful to know the possibility is out there.

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by vaverin on Sun, 30 Mar 2008 04:58:49 GMT

View Forum Message <> Reply to Message

Matt,

due some strange reasons our QA has not tested ovz kernel together with vz kernel. therefore ovz kernel is under testing now and I hope it will be released on the next week.

But you can use vz kernel 028stab053.10 on openvz node.

thank you, Vasily Averin

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by tchipman on Sun, 30 Mar 2008 08:44:39 GMT

View Forum Message <> Reply to Message

Hi Matt,

Just to reply to your question: I ended up deploying the system in question in a totally different config due to the stability problems.

It ended up being deployed thus,

- -no hardware raid, used the 3ware card as JBOD (ugh) only
- -CentOS stock install, no OpenVZ at all
- -software raid config for OS (mirrors) and data (raid5) slices

In this setup, the systems has been 100% solid and working since deployment, which dates back to my previous post in the thread.

Since the requirements for virtualization on this machine were not strict ones (ie, non-virtualized config was tolerable) and the stability issues were a problem, that is how the machine ended up. Clearly not an optimal use of the resource. However, given the history of the hardware (acquired via ebay, absurd low price) and the (very basic) requirements of the project, investing further time in debugging, reinstalling, etc. was not merited.

Sorry I don't have better news for you on this. From the sound of other postings in this thread since your query, it seems maybe things have improved / that other options are now available, which is nice..

---Tim

Subject: Re: *NOT SOLVED* 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable

Posted by TheWiseOne on Sun, 30 Mar 2008 12:38:46 GMT View Forum Message <> Reply to Message

My systems are rock solid, I was just saying it was scary that this bug was not patched yet.

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by TheWiseOne on Tue, 08 Apr 2008 15:52:55 GMT View Forum Message <> Reply to Message

vaverin wrote on Sat, 22 March 2008 04:00lt looks like following issue: http://www.3ware.com/KB/article.aspx?id=15243&cNode=6l1C 6S

We have fixed this issue already and I expect Kir will release new kernel soon.

Any ETA on when this kernel will be released that includes this patch?

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by khorenko on Wed, 09 Apr 2008 12:05:41 GMT

Hello Matt,

the kernel with the fix is going to be released tomorrow.

Thank you for your patience.

--

Konstantin

Subject: Re: 3ware 8506 - CentOS 4.5 or 5.0 - OpenVZ Stock Kernel - Unstable Posted by mcluver on Sun, 25 Apr 2010 20:38:45 GMT

View Forum Message <> Reply to Message

The reason I am posting here again about this topic is because I have switched to Debian since my previous posts and have still been running into the very same issues as described here.

As well as here in detail: http://marc.info/?l=linux-kernel&m=105752554030895&w =2

Upon researching for way, way, way too long -- I believe I have finally found the exact problem (from the horse's mouth) that is causing this kernel panic, data loss, controller failure.

It is a problem between the 3ware Escalade 7506 and 8506 cards and a specific chip on certain Tyan Motherboards, specifically the E7501 Intel chipset, which is present on my Tiger i7501 (S2723) motherboard.

--

FROM 3WARE (http://tyan.com/archive/support/html/f_s2885.html):

Why can't I get my 3ware Raid card to work correctly with this motherboard?

3ware and Tyan have worked together to identify the issue between 3ware 7506 ATA and 8506 SATA RAID controllers and certain chipsets. 3Ware has a technical note available on their knowledge base:

https://www.3ware.com/kbadmin/attachments/TM900-0045-00%20Re v%20A_P.pdf

Tyan will be working with 3Ware to validate the next generation of 8506/9500 RAID controllers on our motherboards.

--

The PDF suggests that the solution is as simple enough as changing the 1U riser card being used to Model# PCITX8-3R manufactured by Adex Electronics, Inc. (http://www.adexelec.com). Too bad I have gone through two years of intermittent hell to finally get this solution, I hope it works...

Best regards,

Matthew Cluver

File Attachments
1) TM900-0045-00 Rev A_P.pdf, downloaded 436 times