

## FreeNAS - Bug #26508

### Intel Optane 900p will not work in ESX passthrough

11/05/2017 02:12 PM - Thomas Rottig

<b>Status:</b> Closed: Third party to resolve	
<b>Priority:</b> Important	
<b>Assignee:</b> Alexander Motin	
<b>Category:</b> OS	
<b>Target version:</b>	
<b>Seen in:</b> 11.0-U4	<b>Needs Merging:</b> Yes
<b>Severity:</b> New	<b>Needs Automation:</b> No
<b>Reason for Closing:</b>	<b>Support Suite Ticket:</b> n/a
<b>Reason for Blocked:</b>	<b>Hardware Configuration:</b> Host: Xeon E3 v5 1270, 64 GB Ram VM: 4 Cores, 32 GB Ram
<b>Needs QA:</b> Yes	<b>ChangeLog Required:</b> No
<b>Needs Doc:</b> Yes	
<b>Description</b>	
Hi, just obtained a Intel 900p 280GB which I am trying to pass through to a FreeNas 11 U4 vm.  ESX is vmware -v1 VMware ESXi 6.5.0 build-6765664 VMware ESXi 6.5.0 Update 1  Bootup fails with error indicated in attached screenshot.  Same drive in a fresh physical installation is working without error Same VM without this drive is working without error Same drive is fine in ESX	
<b>Related issues:</b>	
Has duplicate FreeNAS - Bug #54240: Kernel Panic in nvme_qpair_reset()	<b>Closed</b>

### History

#### #1 - 11/06/2017 06:41 AM - Kris Moore

- Assignee changed from Release Council to Alexander Motin
- Priority changed from No priority to Nice to have
- Target version set to 11.2-BETA1

#### #2 - 11/06/2017 06:55 AM - Alexander Motin

- Status changed from Unscreened to 15

Could you try it with FreeNAS 11.1-RC1? It got some updated to the NVMe driver, so it may be fixed already. If it won't help, please show more logs printed before that, and may be output of `bt` command typed after that.

#### #3 - 11/06/2017 12:12 PM - Thomas Rottig

- File *freenas\_900p\_error2.PNG* added
- File *freenas\_900p\_error3.PNG* added

Alexander Motin wrote:

Could you try it with FreeNAS 11.1-RC1? It got some updated to the NVMe driver, so it may be fixed already. If it won't help, please show more logs printed before that, and may be output of `bt` command typed after that.

No difference, see additional info as requested attached.

**#4 - 11/07/2017 12:45 AM - Alexander Motin**

- Status changed from 15 to Investigation
- Priority changed from Nice to have to Important

**#5 - 11/08/2017 04:17 AM - Alexander Motin**

- Status changed from Investigation to 15

Could you try to boot latest FreeBSD 12-CURRENT snapshot in that VM:

<https://download.freebsd.org/ftp/snapshots/ISO-IMAGES/12.0/FreeBSD-12.0-CURRENT-amd64-20171030-r325156-bootonly.iso.xz> ? One of driver developers think there may be relevant fixes not merged to stable/11 branch yet.

**#6 - 11/08/2017 01:19 PM - Thomas Rottig**

- File *freenas\_900p\_error4.PNG* added
- File *freenas\_900p\_error5.PNG* added

Still no luck - see pics.

**#7 - 11/09/2017 08:16 AM - Alexander Motin**

- Status changed from 15 to Investigation

**#8 - 11/09/2017 09:26 AM - Alexander Motin**

- Status changed from Investigation to 15

Thomas, could you again reproduce it on FreeBSD HEAD, type in `show threads` command and look through for `nvme_ctrlr_identify()` or `nvme_ctrlr_start()` function names in call stacks in the list?

**#9 - 11/09/2017 01:02 PM - Thomas Rottig**

- File *freenas\_900p\_error6.PNG* added

Hi,

no nvme related calls visible, mostly `fork_trampoline` and `sched_switch`.

3 lines of `cpustop_handle` were the most exiting thing to be found - see screenshot.

Ran thrice to see if I had missed it :/

**#10 - 11/14/2017 05:37 AM - Thomas Rottig**

Any updates or next steps? Happy to run some further tests if you let me know what you need...

**#11 - 11/14/2017 07:14 AM - Alexander Motin**

Developer from Intel proposed to try such a patch:

```
diff --git a/sys/dev/nvme/nvme_ctrlr.c b/sys/dev/nvme/nvme_ctrlr.c
index b036eb6..47d9488 100644
--- a/sys/dev/nvme/nvme_ctrlr.c
+++ b/sys/dev/nvme/nvme_ctrlr.c
@@ -348,6 +348,7 @@ nvme_ctrlr_hw_reset(struct nvme_controller *ctrlr)
     DELAY(100*1000);

     nvme_ctrlr_disable(ctrlr);
+    DELAY(5000);
     return (nvme_ctrlr_enable(ctrlr));
 }
```

, but I see no clear logic in it, it looks more like a workaround.

Will you be able to build FreeBSD kernel with that patch or try patched kernel if I send it to you?

**#12 - 11/14/2017 08:10 AM - Thomas Rottig**

Haben's dabbled in Kernel building in a long while, so ready made might be easier. Ideally with some instructions to use, but I can Google those if need be...

**#13 - 11/18/2017 02:58 AM - Alexander Motin**

Here is FreeBSD HEAD kernel with the nvme patch: <https://www.dropbox.com/s/829p50uu92f67ap/kernel.tgz?dl=0>

Please try to drop it into /boot/ of installed FreeBSD VM and try add the NVMe device there.

**#14 - 11/18/2017 09:43 AM - Thomas Rottig**

- File *freenas\_900p\_error7.PNG* added

Hope I have done it correctly, but kernel and most files in the boot/kernel directory show updated timestamp so I assume so.

Result: negative, more or less same error.

Have tried Windows VM - no problem.

Also tried OmniOS 151024ce VM - no go either, so maybe I need to address to Intel instead...

**#15 - 11/20/2017 12:20 PM - Dru Lavigne**

- Status changed from 15 to Unscreened

**#16 - 11/20/2017 01:56 PM - Thomas Rottig**

Hi Dru,  
whats the effect of moving this to unscreened?  
Thanks

**#17 - 11/20/2017 02:10 PM - Dru Lavigne**

Thomas: it lets the developer know that the requested feedback was received.

**#18 - 11/20/2017 03:45 PM - Thomas Rottig**

Ah good to know, thanks:)

**#19 - 11/21/2017 08:40 AM - Sisyphe -**

Have you tried after installing Intel NVMe driver in ESXi?

v1.3.2.4 was published the 1st of November:

<https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI65-INTEL-INTEL-NVME-1324&productId=614>

**#20 - 11/21/2017 01:10 PM - Thomas Rottig**

- File *freenas\_900p\_error8.PNG* added

No I had not, the card had been identified as Optane properly (was on driver from May).  
Have tried now - no change.

```
esxcli system version get
  Product: VMware ESXi
  Version: 6.5.0
  Build: Releasebuild-6765664
  Update: 1
  Patch: 29
```

```
esxcli software vib list |grep nvme
intel-nvme      1.3.2.4-1OEM.650.0.0.4598673  INT  VMwareCertified  2017-11-21
nvme           1.2.0.32-4vmw.650.1.26.5969303  VMW   VMwareCertified  2017-07-29
vmware-esx-esxcli-nvme-plugin 1.2.0.10-1.26.5969303          VMware VMwareCertified  2017-07-29
```

**#21 - 11/27/2017 03:49 AM - Alexander Motin**

- Status changed from Unscreened to Screened

I've looked through the code and can not guess how can that happen. I need to see that myself and be able to debug it directly. Can you provide me some remote access to that VM for experiments? Otherwise it may take time for me to reproduce this issue. Or you may try to contact Intel or FreeBSD developers.

**#22 - 11/27/2017 12:38 PM - Thomas Rottig**

Hi,  
it should be possible to give you access to a (freenas) vm with the optane passed through. What do you need ? Gui or ssh?

Regards

**#23 - 11/27/2017 12:42 PM - Alexander Motin**

Considering we are talking about kernel panics, I'd prefer to have VM console for debugging.

**#24 - 11/27/2017 12:48 PM - Thomas Rottig**

You mean access to the ESX? Sorry not entirely clear what you need

**#25 - 11/27/2017 12:54 PM - Alexander Motin**

Thomas Rottig wrote:

You mean access to the ESX? Sorry not entirely clear what you need

Yes, I was thinking about ESX VM console.

**#26 - 11/27/2017 12:58 PM - Thomas Rottig**

Ok, web interface of the ESX then. Will need some time to set it up. Can you provide me an email for access details please?

**#27 - 11/27/2017 12:59 PM - Alexander Motin**

[mav@ixsystems.com](mailto:mav@ixsystems.com)

**#28 - 12/06/2017 06:02 PM - Sisyphe -**

Were you able to make progress on this issue? Thank you

**#29 - 12/07/2017 11:16 AM - Thomas Rottig**

Still working on supplying the test environment, sorry.

**#30 - 12/07/2017 01:25 PM - Sisyphe -**

Thank you for the update.

**#31 - 12/13/2017 12:37 PM - Thomas Rottig**

My preparations are done, sent the details to Alexander

**#32 - 12/19/2017 01:37 AM - Sisyphe -**

I've updated to Freenas 11.1 stable and I'm still seeing this issue.

Alexander, did you have the chance to look into this? Thanks.

**#33 - 12/19/2017 06:07 PM - Alexander Motin**

I'm sorry, not yet. I've been very busy recently, but I remember about this and hope to look on it nearest days,

**#34 - 12/21/2017 06:00 AM - Kris Moore**

- *Target version changed from 11.2-BETA1 to 11.3*

**#35 - 12/23/2017 11:02 AM - Alexander Motin**

- *Status changed from Screened to Closed: Third party to resolve*

- *Target version deleted (11.3)*

Hello Thomas,

Thank you for the provided access. After some experiments I found that kernel panic in this case is caused by the FreeBSD NVMe driver bug, which can not correctly handle NVMe command timeouts during early controller initialization phase. In particular, `nvme_ctrlr_cmd_set_num_queues()`, which is a real problem here. As I see, that command is the first one sent to the controller, so it may be not a problem related to specific command, but to command submission/handling in general. I've tried to disable MSI-X support in case it is interrupt delivery problem, but it didn't help. Unfortunately the problem seems to be nontrivial, and I have no time to dive there deeper right now.

Have you tried to repeat the experiment on real hardware, not inside a VM? It would give us idea whether it is a problem of the driver in general, or some issue related to VMware's NVMe pass-through.

In either case you should probably contact FreeBSD developers more actively working on this driver. It was actually written by people from Intel, so they may be familiar with both the driver and the hardware.

**#36 - 12/24/2017 01:58 AM - Thomas Rottig**

Hi Alexander,

thanks a lot for your effort.

Very early in the analysis I have run the same setup baremetal and it worked fine.

Given that a similar issue is present on other \*nix system (Solaris & variants) when used in ESX it seems that it is rather related to the virtualization part than the actual OS part. Unfortunately that card is not on VMWare HCL and I also have no support contrac. I think I will try to open a ticket with Intel, they are likely in the best position to work on this.

Thanks a lot,  
happy holidays,  
Thomas

**#37 - 01/04/2018 05:28 AM - Thomas Rottig**

Hi Alexander,

Opened a bug at Intel, not sure what will come out of it.  
Can you provide me with the name/contact details of the aforementioned Intel driver developer?  
The support team would like to reach out to him.

Thanks, regards Thomas

**#38 - 01/04/2018 07:12 AM - Alexander Motin**

I've contacted Jim Harris <[jim.harris@gmail.com](mailto:jim.harris@gmail.com)> and Warner Losh <[imp@FreeBSD.org](mailto:imp@FreeBSD.org)>.

**#39 - 01/11/2018 12:15 PM - Thomas Rottig**

Intel told me they don't support \*nix for this drive so will not support in this issue.

Not sure what else can be done.

**#40 - 01/11/2018 12:36 PM - Alexander Motin**

Have you written to some FreeBSD mailing list or the mentioned people?

**#41 - 01/11/2018 03:36 PM - Thomas Rottig**

No not yet. The issue is not limited to FreeBSD, it also hits OpenSolaris and variants. But might be worth a shot nevertheless, I'll give it a try.

**#42 - 01/30/2018 03:17 AM - Sisyphe -**

Hi Thomas,

Were you able to raise this issue to OpenSolaris or FreeBSD developers?

Thank you

**#43 - 01/30/2018 10:50 PM - Thomas Rottig**

Yes, Warner replied to my emails.

His latest suggestion was  
"There is a small chance <https://reviews.freebsd.org/D14053> fixes this. "  
But I have not had the time to investigate that - feel free to chime in if you can:)

**#44 - 01/31/2018 04:03 AM - Sisyphe -**

I can run some tests. I would however need some help to understand how to compile/get the nvme updated driver and install it on my system.

Thanks!

**#45 - 01/31/2018 06:24 AM - Alexander Motin**

Sisyphe - wrote:

I can run some tests. I would however need some help to understand how to compile/get the nvme updated driver and install it on my system.

FreeNAS has NVMe driver statically linked into the kernel. For that reason you'd have to build whole new kernel with modules at least. I am personally doing that on FreeBSD system according to regular FreeBSD guides, just taking FreeNAS kernel sources from (<https://github.com/freenas/os/tree/freenas/11.1-stable>) and configuration from repo (<https://github.com/freenas/build/blob/master/build/profiles/freenas/kernel/FREENAS.amd64>). Building whole FreeNAS image is much more time-consuming, though that process should also be documented somewhere. By default FreeNAS has no means for self-building, but there is special SDK train, which includes some parts as compilers, etc.

**#46 - 02/20/2018 11:42 PM - Wessel van Norel**

Thomas Rottig wrote:

Yes, Warner replied to my emails.

His latest suggestion was

"There is a small chance <https://reviews.freebsd.org/D14053> fixes this. "

But I have not had the time to investigate that - feel free to chime in if you can:)

Unfortunately I didn't find this issue before getting myself an Intel Optane 900P to passthrough via ESXi... Guess we should open an issue @FreeBSD since it's a kernel issue and not a FreeNAS issue. I've tested the latest nightly ISO FreeBSD-12.0-CURRENT-amd64-20180215-r329338-disc1.iso, and if the revision number in this nightly build is indeed the FreeBSD source revision number, then D14053 unfortunately doesn't fix things. When I only passthrough my Samsung 960 PRO, the ISO boots without errors. When I add the Optane 900P, it fails with the page fault.

**#47 - 02/21/2018 01:36 AM - Wessel van Norel**

I've created a bug @FreeBSD about this issue: [https://bugs.freebsd.org/bugzilla/show\\_bug.cgi?id=226086](https://bugs.freebsd.org/bugzilla/show_bug.cgi?id=226086). Since I guess that's the proper channel for it.

**#48 - 02/21/2018 04:31 AM - Dru Lavigne**

Thanks!

**#49 - 02/26/2018 06:42 AM - Wessel van Norel**

I've good and bad news.

The good news: I just updated the ESXi driver to the latest version: 1.3.2.8 (

<https://my.vmware.com/group/vmware/details?downloadGroup=DT-ESX65-INTEL-INTEL-NVME-1328&productId=614>) and with this version installed the FreeBSD 12 iso I tried before works. I also tried FreeBSD-12.0-CURRENT-amd64-20180125-r328383-disc1.iso, which AFAIK should not contain the <https://reviews.freebsd.org/D14053> fix, and also that one works.

The bad news: my FreeNAS 11.1 ISO still does not work. So I guess the remark from 4 months ago: "Could you try to boot latest FreeBSD 12-CURRENT snapshot in that VM:

<https://download.freebsd.org/ftp/snapshots/ISO-IMAGES/12.0/FreeBSD-12.0-CURRENT-amd64-20171030-r325156-bootonly.iso.xz> ? One of driver



developers think there may be relevant fixes not merged to stable/11 branch yet." is correct. Unfortunately (and understandable) the older FreeBSD 12 Snapshot builds are not available, so I can't find the exact version that contains the fix.

So I guess I've to install FreeNAS 11.1 without the device passthrough, update the kernel and then add the device in order to see if that fixes the problem for FreeNAS.

#### #50 - 02/26/2018 08:00 AM - Alexander Motin

I've merged many NVMe related commits into FreeBSD stable/11 and FreeNAS 11.1-U2, so you may try the last to see whether it is enough. IIRC it still does not include <https://reviews.freebsd.org/D14053>, since it was very new at that point, but I'll try to merge it soon also.

#### #51 - 02/26/2018 08:49 AM - Wessel van Norel

Alexander Motin wrote:

I've merged many NVMe related commits into FreeBSD stable/11 and FreeNAS 11.1-U2, so you may try the last to see whether it is enough. IIRC it still does not include <https://reviews.freebsd.org/D14053>, since it was very new at that point, but I'll try to merge it soon also.

~~Thanks for the quick reply. Unfortunately, FreeNAS 11.1 U2 doesn't work either. Working ISO's for me are at least:~~

~~FreeBSD-12.0-CURRENT-amd64-20180125-r328383-disc1.iso~~

~~FreeBSD-12.0-CURRENT-amd64-20180215-r329338-disc1.iso~~

~~The FreeBSD-12.0-CURRENT-amd64-20180125-r328383-disc1.iso is the oldest FreeBSD 12 iso that I could download. Should I try a FreeBSD 11 ISO as well, or is testing FreeNAS 11.1 U2 good enough for that?~~

See my remark below. Seems that FreeNAS 11.1-U2 (and even FreeNAS 11.1-U1) work depending on the virtual hardware configuration. I'll continue to investigate this when I've time.

#### #52 - 02/26/2018 11:36 AM - Wessel van Norel

I'm afraid the issue will be tougher to debug/resolve. After I said that I couldn't find older ISO images, I did find older VM-IMAGES. And here is where I'm getting worried about this issue. I downloaded the oldest VMDK that was available. Because of the format I'm not able to use the VMDK from the ESXi machine right away, I need to use VMWare Fusion to create a VM and then move it to the ESXi machine. First boot it without the passthrough: ok it works. Then boot it with the passthrough: no it didn't work. Ok, so perhaps the fix is not in that version. Next try to get one working version. So I download the VMDK that is of the same revision as the ISO that is working. But it didn't work.

Then I retry the VM with the ISO. That VM still works. Then I try the ISO on the VM that I created from the VMDK: it doesn't work. So I check the VMWare hardware configuration between the two and step by step make them the same. End result is that I've 2 vm's that look exactly the same from the ESXi web ui, except for the hard disk (since the FreeBSD build hard disk is 21GB in size and the one I made is 8GB in size) and still the new VM doesn't want to boot with the passed through Optane 900p.

So I had no clue what could be wrong. I've diffed the .vmx files and as far as I was able to determine nothing obviously is wrong there. The diff follows below, < is the not working VM, > is the working vm. The only things that I could imagine doing something were the *virtualHW.productCompatibility = "hosted"* and *acpi.smbiosVersion2.7 = "FALSE"*, so I removed those settings, and tried again, without success.

```
2c2,3
< acpi.smbiosVersion2.7 = "FALSE"
---
```

```

> RemoteDisplay.maxConnections = "-1"
> bios.bootRetry.delay = "10"
5,6c6,7
< displayName = "FreeBSDK-disks-test"
< ehci.pciSlotNumber = "36"
---
> displayName = "FreeBSD BugTesting"
> ehci.pciSlotNumber = "34"
9c10
< ethernet0.generatedAddress = "00:0c:29:23:26:30"
---
> ethernet0.generatedAddress = "00:0c:29:62:5d:a9"
12c13
< ethernet0.pciSlotNumber = "32"
---
> ethernet0.pciSlotNumber = "33"
14a16
> ethernet0.wakeOnPcktRcv = "FALSE"
19c21
< migrate.hostLog = "./FreeBSDK-disks-test-d3c8fced.hlog"
---
> migrate.hostLog = "./FreeBSD BugTesting-87d49086.hlog"
24c26
< nvram = "FreeBSDK-disks-test.nvram"
---
> nvram = "FreeBSD BugTesting.nvram"
62c64
< sata0.pciSlotNumber = "34"
---
> sata0.pciSlotNumber = "36"
78c80
< sched.swap.derivedName = "/vmfs/volumes/5a8c8dc2-66208b14-2858-ac1f6b17235e/FreeBSDK-disks-test/FreeBSDK-disks-test-d3c8fced.vswp"
---
> sched.swap.derivedName = "/vmfs/volumes/5a8c8dc2-66208b14-2858-ac1f6b17235e/FreeBSD BugTesting/FreeBSD BugTesting-87d49086.vswp"
83c85
< scsi0:0.fileName = "FreeBSDK-disks-test.vmdk"
---
> scsi0:0.fileName = "FreeBSD BugTesting.vmdk"
93,95c95,97
< tools.syncTime = "TRUE"
< tools.upgrade.policy = "upgradeAtPowerCycle"
< usb.pciSlotNumber = "35"
---
> tools.syncTime = "FALSE"
> tools.upgrade.policy = "manual"
> usb.pciSlotNumber = "32"
106,109c108,110
< uuid.bios = "56 4d cc 7b c6 70 63 53-df e2 94 e6 06 23 26 30"
< uuid.location = "56 4d cc 7b c6 70 63 53-df e2 94 e6 06 23 26 30"
< vc.uuid = "52 14 4b cd 05 81 e4 a9-82 ad 8d 77 42 41 11 68"
< virtualHW.productCompatibility = "hosted"
---
> uuid.bios = "56 4d 68 ef dd 9b 1e 0d-40 7d 8f 7f e9 62 5d a9"
> uuid.location = "56 4d 68 ef dd 9b 1e 0d-40 7d 8f 7f e9 62 5d a9"
> vc.uuid = "52 4b 0c 79 f2 93 e9 66-2c 17 27 0b c6 67 f7 69"
111,112c112,113
< vmci0.id = "102966832"
< vmci0.pciSlotNumber = "33"
---
> vmci0.id = "-379429463"
> vmci0.pciSlotNumber = "35"

```

The not working system hangs on the phase where the working system initialises the NVMe devices, but no output other than then page fault. I've tried creating another VM on the ESXi server and boot that with the ISO: that works. So it's not a fluke accidental working VM. But I'm confused by the VM that doesn't work. Should I reopen the FreeBSD 12 bug?

-- Edit

Seems I've figured out what difference is causing the issue: the order of the pciSlots. If I change the not working .vmx to contain the same pciSlotNumbers as the working .vmx the system boots from the .iso. Unfortunately I've updated all devices at the same time, so I guess I should

figure out what device needs to be in what order to break stuff again. And the other unfortunate part is that even though it now boots from the .iso, booting from the disk image itself still fails, with the note that this time it does not fail with a kernel panic, but the system just hangs (I waited a bit more than 5 minutes) on the NVMe initialisation.

I've also retested the FreeNAS-11.1 ISO's on the VM where the FreeBSD-12 ISO's worked, and the FreeNAS-11.1 ISO's also work. So my previous tests were not as good as I thought. VM with same hardware selected does not mean that it's actually the same. My apologies for that :(

So, it's possible that the Intel Optane 900P works, if you have the latests ESXi drivers installed & have your PCI devices in "the correct order". Now I've "only" to determine what is the correct order. And the other question will be how stable is this, since I find it strange that the PCI order breaks things.

**#53 - 03/23/2018 05:23 AM - Sisyphe -**

Hi,

Do you have an update on this issue? I did several tests but was not able to figure out a working order for PCI slots in the configuration of the VM.

Should the FreeBSD bug be reopened?

**#54 - 03/23/2018 05:31 AM - Wessel van Norel**

Unfortunately no update yet :( Have to make some time to get the system up and running, waste of resources it idling for about a month now... And not sure if to reopen it @ FreeBSD, since I do not understand how it can work (in my case) depending on the PCI slot ordering... It's really weird.

**#55 - 04/10/2018 07:55 AM - Ignacio Rocha**

Wessel van Norel wrote:

Unfortunately no update yet :( Have to make some time to get the system up and running, waste of resources it idling for about a month now... And not sure if to reopen it @ FreeBSD, since I do not understand how it can work (in my case) depending on the PCI slot ordering... It's really weird.

I'm also having the same problem with my Intel Optane 900p when I pass it through esxi. How did you managed to change the PCI slot ordering in the esx vm?

The problem is still active in the 11.1-U4.

I have also, per your recommendations, updated the Intel NVMe driver in esx.

If you need some screenshots or whatever, let me know

**#56 - 04/20/2018 12:40 PM - Jan Eagleman**

I tried to switch from "real hardware" to virtual today, and me too hit this bug. First off ESXi didnt detect the device correctly when using passthrough,

it didnt say its real name like this one for example: I210 Gigabit Network Connection. I then just forwarded the device since I knew it was the NVMe device. Then I tried to boot in a VM with and without the NVMe device and only with the Intel 900P attached it crashed FreeNAS.

Had the Intel 900P working for half a year on "real hardware". Only when using ESXi passthrough it crashes FreeNAS.

**#57 - 06/15/2018 04:56 AM - Sisyphe -**

- Severity set to New

I've updated to ESXI 6.7 and the issue persists.

@Wessel, should the FreeBSD bug be reopened? Should we have a bug opened at VMWare? Thanks.

**#58 - 06/15/2018 05:04 AM - Wessel van Norel**

Sisyphe - wrote:

I've updated to ESXI 6.7 and the issue persists.

@Wessel, should the FreeBSD bug be reopened? Should we have a bug opened at VMWare? Thanks.

I'm not sure. I've been swamped with work but I've tried to get in touch with Paul Braren about this issue. I wrote a comment on: <https://tinkertry.com/intel-optane-900p-should-be-great-for-home-lab-enthusiasts> but it became flagged as spam (it seems it finally came through, but no response yet :( ). He didn't respond on twitter either. I think the strangeness in the way I was able to solve it means it's more related to VMWare than to FreeBSD. And since his post was quite positive about the 900p at first perhaps he is able to put me in touch with the correct people @vmware to get into the internals of what is going wrong.

**#59 - 06/26/2018 06:25 AM - Sisyphe -**

Installed newly released Intel NVMe driver v1.4.0.1016, same issue.

@Norel, did you get a feedback from Paul Braren?

I would reconsider re-opening the FreeBSD ticket. Even if VMWare can support debugging the issue, I suppose the fix will need to be implemented in FreeBSD kernel, as it is also impacting other OpenBSD variants.

**#60 - 06/28/2018 01:24 PM - Wessel van Norel**

Sisyphe - wrote:

Installed newly released Intel NVMe driver v1.4.0.1016, same issue.

@Norel, did you get a feedback from Paul Braren?

@Sisyph: unfortunately no I did not. And since the card is officially not supported by VMWare I'm afraid that we will not get any help from them here. I hoped that the strangeness of the workaround I found would be enough trigger for them to be willing to look at it. Perhaps someone else can try to ping him on the question, someone with a more public profile.

I would reconsider re-opening the FreeBSD ticket. Even if VMWare can support debugging the issue, I suppose the fix will need to be implemented in FreeBSD kernel, as it is also impacting other OpenBSD variants.

My main question is: who's to blame for the issue. Is it FreeBSD or is it VMWare. Since it can be "fixed" in my situation by just a different PCI slot ordering in the VM image configuration, perhaps the problem is within the VMWare kernel and not in the FreeBSD kernel.

What I should have tried is a different OS and see if that works. Unfortunately I need the machine for other stuff right now and I'm not able to properly test things because of that.

**#61 - 07/06/2018 01:46 PM - Sisyph -**

- File *FreeBSD12.png* added

Retested with FreeBSD 12 snapshot (FreeBSD-12.0-CURRENT-amd64-20180628-r335760-bootonly)

System is not crashing and I get "nvme0: Missing interrupt" errors, so it appears an improvement was made to the driver. I would suggest to have the FreeBSD team to have a look again...

**#62 - 08/11/2018 03:03 PM - Sisyph -**

- File *optane900-freenas.png* added

I found a simple fix for this issue by adding the Optane 900P device ID to passthru.map :)

- ssh to ESXi

- edit /etc/vmware/passthru.map
- add following lines at the end of the file:  
# Intel Optane 900P  
8086 2700 d3d0 false
- restart hypervisor

I can now pass through the 900P to Freenas 11.1-U5 without issue:

Enjoy!

**#63 - 08/17/2018 02:32 PM - James McCoy**

Hi all. I was also about to buy and attempt this a few weeks ago before finding this thread at the last minute. Is this 'fix' stable, has anyone else managed to verify it on their system?  
Thanks.

**#64 - 09/29/2018 05:48 AM - Steve Higton**

James McCoy wrote:

Hi all. I was also about to buy and attempt this a few weeks ago before finding this thread at the last minute. Is this 'fix' stable, has anyone else managed to verify it on their system?  
Thanks.

I have just migrated a two way mirrored pool from two Samsung 950Pro drives to two Optane 900p drives. All seems fine so far...

The 950Pros were in passed through to a Freenas 11.1-U6 VM with no issues, running on ESXi 6.0. I split the pool, powered down, swapped a 950Pro for a 900p, powered the ESXi host backup and the Freenas VM crashed on boot in a very similar manner to the crash in the OP. I then tried adding the lines to passthru.map from post [#62](#) but this didn't help, same crash occurred. I then upgraded ESXi to 6.5U1 build 5969303 and the problem disappeared. I passed through the Optane 900p, added to the existing 950Pro, resilvered in a few minutes and all was fine. Another zpool split, power down, swap 950Pro for 900p, etc and now have a two way mirror of two 900p drives.

I only completed the migration about twenty minutes ago but fingers crossed all will remain OK.

**#65 - 10/18/2018 12:31 AM - James McCoy**

Steve Higton wrote:

James McCoy wrote:

Hi all. I was also about to buy and attempt this a few weeks ago before finding this thread at the last minute.  
Is this 'fix' stable, has anyone else managed to verify it on their system?  
Thanks.

I have just migrated a two way mirrored pool from two Samsung 950Pro drives to two Optane 900p drives. All seems fine so far...

The 950Pros were in passed through to a Freenas 11.1-U6 VM with no issues, running on ESXi 6.0. I split the pool, powered down, swapped a 950Pro for a 900p, powered the ESXi host backup and the Freenas VM crashed on boot in a very similar manner to the crash in the OP. I then tried adding the lines to passthru.map from post [#62](#) but this didn't help, same crash occurred. I then upgraded ESXi to 6.5U1 build 5969303 and the problem disappeared. I passed through the Optane 900p, added to the existing 950Pro, resilvered in a few minutes and all was fine. Another zpool split, power down, swap 950Pro for 900p, etc and now have a two way mirror of two 900p drives.

I only completed the migration about twenty minutes ago but fingers crossed all will remain OK.

I took the plunge and can confirm this has been running perfectly for over a month now on ESXi 6.7 and FreeNAS 11.1-U6.

**#66 - 10/29/2018 09:58 AM - Alexander Motin**

- Has duplicate Bug #54240: Kernel Panic in nvme\_qpair\_reset() added

**#67 - 03/31/2019 08:52 PM - Richard May**

This issue is still present in FreeNAS 11.2-U3 and VMware ESXi, 6.7.0, 13004448 with the Optane 800P Series. Editing passthru.map didn't help and the Intel NVMe driver update for ESXi has nothing to do with this particular Optane model (the PnP ID (8086/2522) is not listed in the VIB's XML files).

**#68 - 04/02/2019 02:32 PM - Cy Borg**

- File Screenshot 2019-04-02 at 22.23.58.png added

Richard May wrote:

This issue is still present in FreeNAS 11.2-U3 and VMware ESXi, 6.7.0, 13004448 with the Optane 800P Series. Editing passthru.map didn't help and the Intel NVMe driver update for ESXi has nothing to do with this particular Optane model (the PnP ID (8086/2522) is not listed in the VIB's XML files).

Login to the shell console of your ESXi and edit the .vmx config file for your FreeNAS. Assuming that Optane is the 1st passthru device that you have added [0], add

```
pciPassthru0.msiEnabled = "FALSE"
```

That solves the problem with physical IRQ sharing that seems to be the cause here.

Enjoy,  
CyBorg

## Files

---

freenas_900p_error.PNG	122 KB	11/05/2017	Thomas Rottig
freenas_900p_error3.PNG	143 KB	11/06/2017	Thomas Rottig
freenas_900p_error2.PNG	212 KB	11/06/2017	Thomas Rottig
freenas_900p_error5.PNG	128 KB	11/08/2017	Thomas Rottig
freenas_900p_error4.PNG	183 KB	11/08/2017	Thomas Rottig
freenas_900p_error6.PNG	274 KB	11/09/2017	Thomas Rottig
freenas_900p_error7.PNG	181 KB	11/18/2017	Thomas Rottig
freenas_900p_error8.PNG	173 KB	11/21/2017	Thomas Rottig
FreBSD12.png	87.8 KB	07/06/2018	Sisyph -
optane900-freenas.png	80.6 KB	08/11/2018	Sisyph -
Screenshot 2019-04-02 at 22.23.58.png	94.4 KB	04/02/2019	Cy Borg