Arch Linux

Index » Kernel & Hardware » [SOLVED] CUDA not working on linux-ck kernel

Pages: **1**

SilverMight 2018-07-29 18:05:21 #1

Member

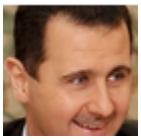

Registered: 2017-11-20
Posts: 25

Offline

Recently, after an update CUDA hasn't been working at all on the linux-ck kernel. Trying to use OBS with NVENC enabled doesn't work, using CUDA with TensorFlow doesn't work either, etc. etc. However, as soon as I boot into the standard linux kernel, everything works fine.

I'm using nvidia-dkms so the driver can work between both kernels. I do have linux-ck-headers and dkms seems to install fine to it, so nothing else seems wrong except the fact that no part of CUDA works when using the linux-ck kernel.

Any help is appreciated.

EDIT: Tried on linux-zen and it seems to work fine there, so a bit clueless on why it won't work on linux-ck

Last edited by SilverMight (2018-08-01 17:13:41)

inglor 2018-07-30 10:25:16 #2

Package Maintainer (PM)

Registered: 2008-07-22
Posts: 81

Offline

Are you using modprobed-db also? Maybe you are missing a module. If you can share the project I can try reproducing it locally.

Last edited by inglor (2018-07-30 10:26:14)

SilverMight 2018-07-30 15:49:59 #3

Member

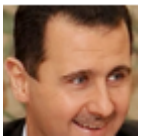

Registered: 2017-11-20
Posts: 25

Offline

> **inglor wrote:**
>
> Are you using modprobed-db also? Maybe you are missing a module. If you can share the project I can try reproducing it locally.

Don't believe so, however I have tried running nvidia-modprobe to no avail. I'll give that a shot.

The kernel can be found at https://aur.archlinux.org/packages/linux-ck/

inglor 2018-07-30 17:08:55 #4

Package Maintainer (PM)

Registered: 2008-07-22
Posts: 81

Offline

> **SilverMight wrote:**
>
> The kernel can be found at https://aur.archlinux.org/packages/linux-ck/

Sorry I wasn't clear. If you tell me the steps to reproduce it I can give it a go on my PC which I have linux-ck with DKMS and CUDA avalaible. This is why I was asking for a project.

> **SilverMight wrote:**
>
> [..]Trying to use OBS with NVENC enabled doesn't work, using CUDA with TensorFlow doesn't work either, etc. etc. However, as soon as I boot into the standard linux kernel, everything works fine. [..]

Is this coming from a project ?

SilverMight 2018-07-31 14:41:41 #5

Member

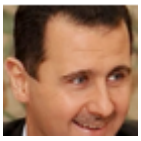

Registered: 2017-11-20
Posts: 25

Offline

My bad, yes. TensorFlow is pretty large so I'd recommend installing OBS (sudo pacman -S obs-studio), going to File -> Settings and then Output and then changing the recording encoder from Software to Hardware (NVENC), then hit start recording.

huyizheng 2018-07-31 15:11:12 #6

Member

Registered: 2018-05-15
Posts: 20

Offline

Same here.
When I use linux-ck-haswell in repo-ck. I can't run the deviceQuery in cuda samples:

```
$ cd "cuda sample's directory"
$ ./bin/x86_64/linux/release/deviceQuery
./bin/x86_64/linux/release/deviceQuery St

 CUDA Device Query (Runtime API) version

cudaGetDeviceCount returned 30
-> unknown error
Result = FAIL
```

However I have install nvidia-dkms and it works:

```
$ lsmod | grep nvidia
nvidia_drm             45056  10
nvidia_modeset       1093632  8 nvidia_dr
nvidia              14061568  825 nvidia_
drm_kms_helper        196608  2 nvidia_dr
drm                   466944  13 drm_kms_
ipmi_msghandler        57344  2 ipmi_devi
```

If I switch back to linux kernel then cuda works fine.

SilverMight 2018-07-31 20:51:18 #7

Member

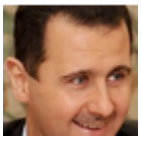

Registered: 2017-11-20
Posts: 25

Offline

> **huyizheng wrote:**
>
> Same here.
> When I use linux-ck-haswell in repo-ck. I can't run the deviceQuery in cuda samples:
>
> ```
> $ cd "cuda sample's directory"
> $ ./bin/x86_64/linux/release/device
> ./bin/x86_64/linux/release/deviceQu
>
>  CUDA Device Query (Runtime API) ve
>
> cudaGetDeviceCount returned 30
> -> unknown error
> Result = FAIL
> ```
>
> However I have install nvidia-dkms and it works:
>
> ```
> $ lsmod | grep nvidia
> nvidia_drm             45056  10
> nvidia_modeset       1093632  8 nvi
> nvidia              14061568  825 r
> drm_kms_helper        196608  2 nvi
> drm                   466944  13 dr
> ipmi_msghandler        57344  2 ipm
> ```
>
> If I switch back to linux kernel then cuda works fine.

Just tried that and got the same results as you.

```
./deviceQuery Starting...

 CUDA Device Query (Runtime API) version

cudaGetDeviceCount returned 30
-> unknown error
Result = FAIL
```

inglor 2018-07-31 21:32:22 #8

Package Maintainer (PM)

Registered: 2008-07-22
Posts: 81

Offline

Same here (linux-ck with nvidia-dkms and supposed to be working fine). Could it be that nvidia-dkms (the module) is build with gcc8 and Cuda only supports gcc7?

SilverMight 2018-07-31 21:45:16 #9

Member

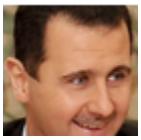

Registered: 2017-11-20
Posts: 25

Offline

> **inglor wrote:**
>
> Same here (linux-ck with nvidia-dkms and supposed to be working fine). Could it be that nvidia-dkms (the module) is build with gcc8 and Cuda only supports gcc7?

I don't think so, since it works fine on any other kernel except the -ck one.

inglor 2018-08-01 07:41:55 #10

Package Maintainer (PM)

Registered: 2008-07-22
Posts: 81

Offline

Enabled NUMA on the linux-ck kernel, recompile and works fine.

```
$ ./deviceQuery
./deviceQuery Starting...

 CUDA Device Query (Runtime API) version

Detected 1 CUDA Capable device(s)

Device 0: "GeForce GTX 1080"
  CUDA Driver Version / Runtime Version
  CUDA Capability Major/Minor version num
  Total amount of global memory:
  (20) Multiprocessors, (128) CUDA Cores/
  GPU Max Clock rate:
  Memory Clock rate:
  Memory Bus Width:
  L2 Cache Size:
  Maximum Texture Dimension Size (x,y,z)
  Maximum Layered 1D Texture Size, (num)
  Maximum Layered 2D Texture Size, (num)
  Total amount of constant memory:
  Total amount of shared memory per block
  Total number of registers available per

$ uname -a
Linux tiamat 4.17.11-1-ck #1 SMP PREEMPT
```

huyizheng 2018-08-01 09:31:05 #11

Member

Registered: 2018-05-15
Posts: 20

Offline

> **inglor wrote:**
>
> Enabled NUMA on the linux-ck kernel, recompile and works fine.
>
> ```
> $ ./deviceQuery
> ./deviceQuery Starting...
>
>  CUDA Device Query (Runtime API) ve
>
> Detected 1 CUDA Capable device(s)
>
> Device 0: "GeForce GTX 1080"
>   CUDA Driver Version / Runtime Ver
>   CUDA Capability Major/Minor versi
>   Total amount of global memory:
>   (20) Multiprocessors, (128) CUDA
>   GPU Max Clock rate:
>   Memory Clock rate:
>   Memory Bus Width:
>   L2 Cache Size:
>   Maximum Texture Dimension Size (x
>   Maximum Layered 1D Texture Size,
>   Maximum Layered 2D Texture Size,
>   Total amount of constant memory:
>   Total amount of shared memory per
>   Total number of registers availab
>
> $ uname -a
> Linux tiamat 4.17.11-1-ck #1 SMP PR
> ```

But in linux-ck's PKGBUILD it says that it's not recommend to enable this feature in single CPU platform.

progandy 2018-08-01 09:45:47 #12

Member

Registered: 2012-05-17
Posts: 5,091

Offline

With CUDA you are using the GPU as a processor, so it is not a single CPU platform anymore.

SilverMight 2018-08-01 17:14:37 #13

Member

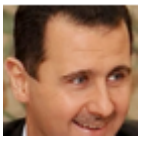

Registered: 2017-11-20
Posts: 25

Offline

Just compiled with NUMA, definitely fixed the issue. Thanks for the fix

graysky 2018-08-01 17:48:37 #14

Wiki Maintainer

From: :wq
Registered: 2008-12-01
Posts: 10,489

Website

Offline

I will comment the PKGBUILD for CUDA users and reference this discussion, thank you.

Pages: **1**

Index » Kernel & Hardware
» [SOLVED] CUDA not working on linux-ck kernel

Jump to
Kernel & Hardware Go

Atom topic feed

Powered by **FluxBB**