

FAQ and Troubleshooting

Bitfusion Guide

Table of Contents

Can I use FlexDirect on my own hardware?	3
What is my performance going to be like?	3
“Your kernel may not have been built with NUMA support”	3
Running out of memory errors	3
Error establishing connection: Cannot allocate memory	3
Working with HTTP_PROXY settings	4
CUDA 9.0 “memory operations are not supported on this device”	4
CUDA_ERROR_PEER_ACCESS_UNSUPPORTED	5
Utility, nvidia-smi, not running	5
Error Message: could not find = char	5
Error Message: all CUDA-capable devices are busy or unavailable	5

Can I use FlexDirect on my own hardware?

Yes, it can be used both on-premise in your datacenter as well in public clouds like AWS, Azure, etc.

What is my performance going to be like?

Great question, it really depends on the model and instances you choose. We do recommend at least 10GbE networking for most use-cases. High-speed fabrics such as Infiniband and those with RDMA support will be necessary for multi-server scenarios. The best thing to do is to test it out yourself and contact us if you want us to dive deeper with you.

“Your kernel may not have been built with NUMA support”

When running with FlexDirect you may see the warning message, “Your kernel may not have been built with NUMA support.”. These messages have no impact on performance or accuracy of TensorFlow results. They are caused by TensorFlow looking for hardware properties through sysfs, and, of course, such information will not be available on a CPU node because it is using network-attached GPUs. The FlexDirect runtime performance benefits from NUMA optimizations when appropriate, so you can safely ignore these warnings.

Running out of memory errors

When running large models or batch sizes, frameworks such as TensorFlow can report out of memory errors:

Text

```
W tensorflow/core/common_runtime/gpu/gpu_bfc_allocator.cc:211] Ran out of memory trying
to allocate 877.38MiB. See logs for memory state

W tensorflow/core/kernels/cwise_ops_common.cc:56] Resource exhausted: OOM when allocating
tensor with shape[10000,23000]
```

These are legitimate errors. The application requires more memory than you have assigned or is available from the GPUs. Avoiding these issues can be a combination of one or more strategies:

- Reduce batch size
- Use a larger GPU size
- Increase model parallelism by splitting your model into smaller chunks

Error establishing connection: Cannot allocate memory

This error can occur if the system has a resource limit that is too restrictive. To avoid this issue increase the number of open files allowed with the `ulimit` command.

Text

```
$ ulimit -n 4096
# or
$ ulimit -n unlimited
```

Working with HTTP_PROXY settings

By default, the `http_proxy` and `https_proxy` environment variables are not honored by FlexDirect for communications between the client and server(s). This is by design, as in-cluster networking performance can potentially be reduced by web proxies. To force FlexDirect to use the system's proxy settings, use the `BF_USE_PROXY` environment variable either in your startup scripts or prior to launching any server or client:

Text

```
$ export BF_USE_PROXY=1
```

CUDA 9.0 “memory operations are not supported on this device”

CUDA 9.0, as of January 24, 2018, disables batch memory operations by default as an errata. These operations are mainly used for GPUDirect-enabled applications. Thus, it is recommended to enable this setting for best results. To re-enable, remove all NVIDIA modules and re-install with the `NVreg_EnableStreamMemOPs` parameter enabled:

Text

```
$ sudo rmmod nvidia nvidia_uvm nvidia_drm nvidia_modeset
$ sudo modprobe nvidia NVreg_EnableStreamMemOPs=1
```

Sometimes, a module cannot be removed because another application is using it. It can be difficult to determine what the specific application is. You may need to manually examine the list of running processes and kill likely candidates. There may desktop or graphical services running a known service often found in VMware environments is *lightdm*. Do some exploration to find which application is responsible. Desktop or other graphical services and applications are good candidates. You can see everything that is running with:

Text

```
$ ps auxf
# Examine process and, for example, note that “lightdm” is running, which uses the GPU
$ sudo kill <pid>
# Or
$ sudo systemctl stop <name> // e.g. lightdm
```

Then try again to uninstall-reinstall the nvidia module.

CUDA_ERROR_PEER_ACCESS_UNSUPPORTED

TensorFlow may emit an error, `CUDA_ERROR_PEER_ACCESS_UNSUPPORTED`, when it finds GPU pairs not connected by the PCIe and system topology. You may ignore these errors. The job of FlexDirect virtualization is to handle the necessary communication via the network (e.g., ethernet of InfiniBand). An example of the error message is here:

```
2018-09-05 20:42:10.049855: W tensorflow/core/common_runtime/gpu/gpu_device.
cc:1331] Unable to enable peer access between device ordinals 0 and 6, status: Internal:
failed to enable peer access from 0x55ef97c9fef0 to 0x55ef97cb2520: CUDA_ERROR_
PEER_ACCESS_UNSUPPORTED
```

Utility, `nvidia-smi`, not running

the Nvidia utility, `nvidia-smi`, is released with the Nvidia driver. The utility is often updated as well as the driver. An older `nvidia-smi` may not work with a later driver. For example, the version of `nvidia-smi` that comes with the 410 driver version, does not work with driver version 418.

Error message: could not find = char

This error message is sometimes seen near the beginning of the FlexDirect output. It may be ignored. It may be repeated several times:

```
could not find = char
```

Ultimately it comes from a third-party library, `ibverbs`. The best way to prevent unnecessary occurrences is to configure FlexDirect to explore and use only the network interfaces and transport mechanisms you want it to use. This can be configured is documented under [Advanced Networking Configuration](#).

Error Message: all CUDA-capable devices are busy or unavailable

If your attempt to run multiple applications on a GPU fails (or all but one of the applications fail) with an error message such as, `Cuda failure p2pBandwidthLatencyTest.cu:68: 'all CUDA-capable devices are busy or unavailable'`, then change the NVIDIA GPU compute mode setting from “Exclusive” to “Default.”

```
sudo nvidia-smi -c 0
```

The “Default” mode allows GPU sharing.

You can see the current compute mode with `nvidia-smi -q` (along with a lot of other information), e.g.,

```
Compute Mode           : Default
```

