

System Requirements

Bitfusion Guide

| | |
|---|---|
| Table of Contents | |
| System Requirements | 3 |
| Hardware Requirements | 3 |
| Networking | 4 |
| Verifying System Health | 5 |
| Software Dependencies | 6 |
| Ubuntu | 6 |
| CentOS and RedHat | 7 |
| List CUDA Libraries | 7 |
| Verify the CUDA Installation | 7 |
| CUDA Installation on CPU Nodes/VMs/Containers | 7 |

System Requirements

FlexDirect can be installed for different modes of operation. Some modes include, in whole or in part, the abilities of other modes. This one document lists the requirements for each mode.

| FLEXDIRECT CLIENTS | FLEXDIRECT ANALYTICS | FLEXDIRECT SERVER | FLEXDIRECT MANAGER |
|--|--|--|--|
| <ul style="list-style-type: none"> • Ubuntu LTS 16.04 • CentOS 7 or • RHEL 7.4+ NVIDIA Toolkit <ul style="list-style-type: none"> • CUDA 7.5 • CUDA 8 • CUDA 9 • CUDA 10 | <ul style="list-style-type: none"> • Ubuntu LTS 16.04 • CentOS 7 or • RHEL 7.4+ | <ul style="list-style-type: none"> • Ubuntu LTS 16.04 • CentOS 7 or • RHEL 7.4+ | <ul style="list-style-type: none"> • Ubuntu LTS 16.04 • CentOS 7 or • RHEL 7.4+ |
| | NVIDIA GPU driver version 372 or higher | NVIDIA GPU driver version 372 or higher | NVIDIA GPU driver version 372 or higher |
| | | Note: Can also run Applications under FlexDirect Client , in which case NVIDIA Toolkit prerequisites are needed. | Note: Can also run Applications under FlexDirect Client , in which case NVIDIA Toolkit prerequisites are needed. |

Hardware Requirements

| FLEXDIRECT CLIENTS | FLEXDIRECT ANALYTICS | FLEXDIRECT SERVER | FLEXDIRECT MANAGER |
|--|---|--|--|
| • N/A | • Any generation CUDA-enabled NVIDIA GPU(s) | • Any generation CUDA-enabled NVIDIA GPU(s) | • Any generation CUDA-enabled NVIDIA GPU(s) |
| 10Gbs or more <ul style="list-style-type: none"> • Ethernet (TCP/IP) • RoCE or • Infiniband | Network Connectivity | 10Gbs or more <ul style="list-style-type: none"> • Ethernet (TCP/IP) • RoCE or • Infiniband | 10Gbs or more <ul style="list-style-type: none"> • Ethernet (TCP/IP) • RoCE or • Infiniband |
| System memory on the client CPU machines should be 1.5x the total aggregate GPU memory across all the GPUs in the largest GPU machine. | N/A | N/A | N/A |

Networking

FlexDirect makes use of several ports or ranges of ports for process to process communication. Please ensure that your firewalls do not block the ports below. Please ensure there are no port conflicts with other applications. Some of the ranges can be modified with command-line arguments.

| FLEXDIRECT CLIENTS | FLEXDIRECT ANALYTICS | FLEXDIRECT SERVER | FLEXDIRECT MANAGER |
|---|--|---|--|
| 56001 <i>Server Communication</i> <ul style="list-style-type: none"> • Outbound to: <ul style="list-style-type: none"> – Server nodes – Manager nodes | 56001 <i>Server Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Analytics nodes – Manager nodes • Outbound to: <ul style="list-style-type: none"> – Analytics nodes – Manager nodes <p>Note: Override with <code>--srs_port 56008</code>, for example.</p> | 56001 <i>Server Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Client nodes <p>Note: Override with <code>--srs_port 56008</code>, for example.</p> | 56001 <i>Server Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Client nodes – Analytics nodes – Manager nodes • Outbound to: <ul style="list-style-type: none"> – Analytics nodes – Manager nodes <p>Note: Override with <code>--srs_port 56008</code>, for example.</p> |
| 55001-55100 <i>Dispatcher Communication</i> <ul style="list-style-type: none"> • Outbound to: <ul style="list-style-type: none"> – Server nodes – Manager nodes | 55001-55100 <i>Dispatcher Communication</i> N/A | 55001-55100 <i>Dispatcher Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Client nodes | 55001-55100 <i>Dispatcher Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Client nodes |
| 45201-46225 <i>CUDA Communication</i> <ul style="list-style-type: none"> • Outbound to: <ul style="list-style-type: none"> – Server nodes – Manager nodes | 45201-46225 <i>CUDA Communication</i> N/A | 45201-46225 <i>CUDA Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Client nodes | 45201-46225 <i>CUDA Communication</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Client nodes |
| 54000 <i>Web Server</i> N/A | 54000 <i>Web Server</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Anyone with a browser <p>Note: Override with <code>--web_port 12345</code>, for example.</p> | 54000 <i>Web Server</i> N/A | 54000 <i>Web Server</i> <ul style="list-style-type: none"> • Inbound from: <ul style="list-style-type: none"> – Anyone with a browser <p>Note: Override with <code>--web_port 12345</code>, for example.</p> |

| FLEXDIRECT CLIENTS | FLEXDIRECT ANALYTICS | FLEXDIRECT SERVER | FLEXDIRECT MANAGER |
|--|---|--|---|
| <p>54001 <i>Statistics Collection</i> N/A</p> | <p>54001 <i>Statistics Collection</i></p> <ul style="list-style-type: none"> • Inbound/ Outbound – localhost only <p>Note: Override with <code>--collection_port 54321</code>, for example.</p> | <p>54001 <i>Statistics Collection</i> N/A</p> | <p>54001 <i>Statistics Collection</i></p> <ul style="list-style-type: none"> • Inbound/ Outbound – localhost only <p>Note: Override with <code>--collection_port 54321</code>, for example.</p> |

- Ensure your machines have ingress/egress internet access for access to downloads, licensing, etc.
- Check for other networking policies, such as outbound proxy or internal DNS, static versus dynamic IPs. Consult with the Bitfusion team if you are using dynamic IPs or have any proxies set up.

Verifying System Health

Once you have installed FlexDirect, you can validate your environment for the best results by running the FlexDirect health check.

```

Shell

# Assumes flexdirect is installed; see installation guide.
flexdirect health
    
```

The health check will run checks on the nodes of your cluster appropriate to their hardware (e.g., GPUs) or to their configuration (e.g. RoCE). So your output may differ from that shown below. Checks try to find settings or problems that will limit or prevent the high-bandwidth, low-latency communication needed for the best performance. The results should be self explanatory. Checks will be performed on the local host and on all configured GPU servers. Shown below is the output from the check on the local node.

Output

```

$ flexdirect health
Health report for server: 192.168.10.41:56001
=====
Software Versions checks:
=====
[PASS   ] Check library dependency
[PASS   ] Check OFED Version: 4.3
[PASS   ] Check CUDA version >= 7050: Current version: 9020
[PASS   ] Check GPU driver version >= 367.0: Current version: 396.37

System Resources checks:
=====
[PASS   ] Check flexdirect install
[PASS   ] Check external connectivity to Bitfusion license server. Ignore for
on-premise installations.
[PASS   ] Check shadow memory 192076MB and total gpu mem 64640MB

Performance checks:
=====
[PASS   ] Check MTU Size: hi-speed interfaces MTU >= 4K
[PASS   ] Check mem ops
[PASS   ] Check ulimit -n >= 4096: 4096
[MARGINAL] Check multinode support: GPU direct not supported
[MARGINAL] Check nv_peer_mem: nv_peer_mem module not loaded
[SKIPPED] Cannot perform PCIe diagnosis without root access. Run the binary with sudo
to get diagnosis

Stability checks:
=====
[PASS   ] Check Network Errors/Drops: found no errors or packet drops
[PASS   ] Check IB physically up
[PASS   ] Check IB SM up
[PASS   ] Check MAD agent registration errors
[PASS   ] Check GPU API mismatch
[PASS   ] Check GPU Xid errors
[PASS   ] Check temperature <= 100.00: current temp: 42.00
[PASS   ] Check ECC errors
[MARGINAL] Check RDMA: TX pause parameter is off
           : RX pause parameter is off

PASS: 18
MARGINAL: 3
FATAL: 0
nvmSystemGetNVMLVersion: 9.396.37
nvmSystemGetDriverVersion: 396.37
CUDA device capability Major/Minor version number: 7.0
cudaDriverVersion: 9020
cudaDriverGetVersion: 9020
cudaRuntimeGetVersion: 9000
cudnnGetVersion: 7201
cudnnGetCudartVersion: 9020

```

Software Dependencies

When you run the default FlexDirect Installer from above, it downloads and installs all the software dependencies for you. If, however, you would like to install all the software dependencies manually, here is the list.

Ubuntu

Ubuntu Dependencies

```
$ apt -y install uuid libjson-c2 librdmacm1 libprocps4 procps
```

CentOS and RedHat

Ubuntu Dependencies

```
$ yum install -y json-c-devel librdmacm libibverbs libuuid proc-ng-devel
```

List CUDA Libraries

CUDA libs might be missing if your installation is a minimal one. Copy/paste the files from /usr/local/cuda/lib64:

Shell

```
$ ls /usr/local/cuda/lib64
```

Verify the CUDA installation

The easiest way to do this will be to install the CUDA samples. [Here is the RPM for CUDA samples](#). It is the RPM for RHEL 7 and CUDA 9.1.

After installing the samples, make sure deviceQuery can be run either directly (if it is a GPU machine) or with [FlexDirect](#) (if it is a Client CPU machine)

Schedule a test job using CUDA deviceQuery:

Shell

```
$ cd /usr/local/cuda/samples/1_Uutilities/deviceQuery
$ make all
$ flexdirect run -n 1 ./deviceQuery
$ flexdirect run -n 7 ./deviceQuery
```

CUDA Installation on CPU Nodes/VMs/Containers

In order to run GPU applications with FlexDirect on CPU nodes, VMs, or in containers which do not have direct access to physical GPU hardware, you still need to install CUDA—but there is no need to install any Nvidia driver components. The example below illustrates how to install CUDA 9.1 in a few easy steps. The CUDA version that you install on the Client should match the CUDA version which is running on the server.

Shell

```
$ wget http://developer.download.nvidia.com/compute/cuda/repos/ubuntu1604/x86_64/cuda-repo-ubuntu1604_9.1.85-1_amd64.deb
$ sudo dpkg -i cuda-repo-ubuntu1604_9.1.85-1_amd64.deb
$ sudo apt-key adv --fetch-keys http://developer.download.nvidia.com/compute/cuda/repos/ubuntu1604/x86_64/7fa2af80.pub
$ sudo apt-get update
$ sudo apt-get install -y cuda-toolkit-9-1
$ rm cuda-repo-ubuntu1604_9.1.85-1_amd64.deb
```

