




Hewlett Packard
Enterprise

Installation guide

INSTALLING UBUNTU OS FOR MELLANOX OPENFABRICS ENTERPRISE LINUX DISTRIBUTION ON HPE INFINIBAND EDR/ETHERNET 100G ADAPTERS



CONTENTS

Introduction.....	3
About this document	3
System requirements.....	3
Installation process overview	3
Verify network adapter is installed.....	3
Download the Mellanox drivers.....	4
Install the OFED installation package.....	4
Post-installation actions.....	8
Additional links.....	11



INTRODUCTION

HPE and Mellanox® have aligned to enable end-to-end connectivity solutions for hyperconverged infrastructures. By leveraging the highest throughputs and lowest latency that Mellanox intelligent interconnect solutions provide, HPE customers can take advantage of data delivery speeds of up to 100Gb/s InfiniBand or 10/25/40/50/56/100Gb/s Ethernet. These speeds enable customers to run high-performance workloads such as real-time response and extreme computing.

About this document

This document is intended for users familiar with Linux® environments. This guide covers the steps for installing the Mellanox Linux OpenFabrics Enterprise Distribution (OFED), the software stack used to manage InfiniBand and Ethernet network adapter cards, specifically using the Mellanox ConnectX®-5 platform for HPE InfiniBand EDR/Ethernet 100Gb 2-port 841QSFP28 Adapter 872726-B21.

IMPORTANT

Superuser privileges are required for installation. Use the superuser (root) or use the sudo prefix if the sudo package is enabled.

SYSTEM REQUIREMENTS

The following table lists the system requirements to install the Mellanox InfiniBand drivers.

TABLE 1. Installation requirements

	Requirement
Operating system	Ubuntu 18.04
NIC	HPE InfiniBand EDR/Ethernet 100Gb 2-port 841QSFP28 Adapter 872726-B21

INSTALLATION PROCESS OVERVIEW

The high-level steps for installing the drivers are:

1. Verify the network adapter is present.
2. Download the Mellanox drivers (OFED file).
3. Install the OFED installation package.
4. Perform the post-installation actions.
5. Update the firmware.

Verify network adapter is installed

Issue the following command to verify the network adapter is installed:

```
user@ubuntu-gpu1:~$ lspci -v | grep Mellanox
```

Sample output:

```
user@ubuntu-gpu1:~$ lspci -v | grep Mellanox
```

```
86:00.0 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
```

```
86:00.1 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
```



Download the Mellanox drivers

The Mellanox drivers are available at: https://www.mellanox.com/page/products_dyn?product_family=26&mtag=linux_sw_drivers



Note: By downloading and installing MLNX_OFED package for Oracle Linux (OL) OS, you may be violating your operating system's support matrix. Please consult with your operating system support before installing.

Version (Current)	OS Distribution	OS Distribution Version	Architecture	Download/Documentation
4.7-1.0.0.1	Ubuntu	Ubuntu 19.04	x86_64	ISO: MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64.iso Size: 254M MD5SUM: bb8d5881f875045d9f59811725b808f5
	SLES	Ubuntu 18.04	ppc64le	
	RHEL/CentOS	Ubuntu 16.04	aarch64	
	OL	Ubuntu 14.04		
	Fedora			
	EulerOS			tgz: MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64.tgz Size: 253M MD5SUM: 000c62dbb0adbf3c0e592148180abb8
	Debian			SOURCES: MLNX_OFED_SRC-debian-4.7-1.0.0.1.tgz Size: 52M MD5SUM: 83e760ed658e9c35a6ae26970fd8e765

FIGURE 1. MLNX_OFED Download Center screen

Select the following options from the Mellanox Download Center:

- Version: 4.7-1.0.0.1
- OS Distribution: Ubuntu
- OS Distribution Version: Ubuntu 18.04
- Architecture: As applicable
- Download: tgz: MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64.tgz

Install the OFED installation package

Follow these steps to install the OFED package.

Installation of the OFED package is executed by running the installation script `mlnxofedinstall`. Refer to the MLNX_OFED Documentation Rev 4.7-1.0.0.1 for other installation methods.

1. Copy the downloaded OFED file to the `/tmp` directory.

```
user@ubuntu-gpu1:~$ cp MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64.tgz /tmp
```

2. Issue the following commands to extract the .tar file.

```
user@ubuntu-gpu1:~$ cd /tmp
user@ubuntu-gpu1:/tmp$ sudo tar -xzf ./MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64.tgz
```



3. Run the installation script (mlnxofedinstall) by following these commands:

```
user@ubuntu-gpu1:/tmp$ cd MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64/
```

```
user@ubuntu-gpu1:/tmp/MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64$ sudo ./mlnxofedinstall
```

Sample output:

```
Logs dir: /tmp/MLNX_OFED_LINUX.3524.logs
```

```
General log file: /tmp/MLNX_OFED_LINUX.3524.logs/general.log
```

The following list shows the MLNX_OFED_LINUX packages that you have chosen (some might have been added by the installer due to package dependencies):

```
ofed-scripts
mlnx-ofed-kernel-utils
mlnx-ofed-kernel-dkms
rshim-dkms
iser-dkms
srp-dkms
libibverbs1
ibverbs-utils
libibverbs-dev
libibverbs1-dbg
libmlx4-1
libmlx4-dev
libmlx4-1-dbg
libmlx5-1
libmlx5-dev
libmlx5-1-dbg
librxe-1
librxe-dev
librxe-1-dbg
libibumad
libibumad-static
libibumad-devel
ibacm
ibacm-dev
librdmacm1
librdmacm-utils
librdmacm-dev
mstflint
ibdump
libibmad
libibmad-static
libibmad-devel
libopensm
opensm
opensm-doc
libopensm-devel
infiniband-diags
infiniband-diags-compatible
mft
kernel-mft-dkms
libibcm1
libibcm-dev
perftest
ibutils2
libibdm1
cc-mgr
ar-mgr
dump-pr
ibsim
ibsim-doc
knem-dkms
knem
mxm
```



```
ucx
sharp
hcoll
openmpi
mpitests
libdapl2
dapl2-utils
libdapl-dev
srptools
mlnx-ethtool
mlnx-iproute2
```

This program will install the MLNX_OFED_LINUX package on your machine. Note that all other Mellanox, OEM, OFED, RDMA, or Distribution IB packages will be removed. Those packages are removed due to conflicts with MLNX_OFED_LINUX. Do not reinstall them.

```
Do you want to continue?[y/N]:y
```

```
Checking SW Requirements...
```

```
One or more required packages for installing MLNX_OFED_LINUX are missing.
```

```
Attempting to install the following missing packages:
```

```
libgfortran3 flex swig chrpath tk debhelper automake quilt gfortran libnl-route-3-200 m4 libltdl-dev dpatch
bison graphviz autoconf autotools-dev tcl
```

```
Removing old packages...
```

```
Installing new packages
```

```
Installing ofed-scripts-4.7...
```

```
Installing mlnx-ofed-kernel-utils-4.7...
```

```
Installing mlnx-ofed-kernel-dkms-4.7...
```

```
Installing rshim-dkms-1.8...
```

```
Installing iser-dkms-4.7...
```

```
Installing srp-dkms-4.7...
```

```
Installing libibverbs1-41mlnx1...
```

```
Installing ibverbs-utils-41mlnx1...
```

```
Installing libibverbs-dev-41mlnx1...
```

```
Installing libibverbs1-dbg-41mlnx1...
```

```
Installing libmlx4-1-41mlnx1...
```

```
Installing libmlx4-dev-41mlnx1...
```

```
Installing libmlx4-1-dbg-41mlnx1...
```

```
Installing libmlx5-1-41mlnx1...
```

```
Installing libmlx5-dev-41mlnx1...
```

```
Installing libmlx5-1-dbg-41mlnx1...
```

```
Installing librcxe-1-41mlnx1...
```

```
Installing librcxe-dev-41mlnx1...
```

```
Installing librcxe-1-dbg-41mlnx1...
```

```
Installing libibumad-43.1.1.MLNX20190905.1080879...
```

```
Installing libibumad-static-43.1.1.MLNX20190905.1080879...
```

```
Installing libibumad-devel-43.1.1.MLNX20190905.1080879...
```

```
Installing ibacm-41mlnx1...
```

```
Installing ibacm-dev-41mlnx1...
```

```
Installing librdmacm1-41mlnx1...
```

```
Installing librdmacm-utils-41mlnx1...
```

```
Installing librdmacm-dev-41mlnx1...
```

```
Installing mstflint-4.13.0...
```

```
Installing ibdump-5.0.0...
```

```
Installing libibmad-5.4.0.MLNX20190423.1d917ae...
```

```
Installing libibmad-static-5.4.0.MLNX20190423.1d917ae...
```

```
Installing libibmad-devel-5.4.0.MLNX20190423.1d917ae...
```

```
Installing libopensm-5.5.0.MLNX20190923.1c78385...
```

```
Installing opensm-5.5.0.MLNX20190923.1c78385...
```

```
Installing opensm-doc-5.5.0.MLNX20190923.1c78385...
```

```
Installing libopensm-devel-5.5.0.MLNX20190923.1c78385...
```

```
Installing infiniband-diags-5.4.0.MLNX20190908.5f40e4f...
```

```
Installing infiniband-diags-compatible-5.4.0.MLNX20190908.5f40e4f...
```

```
Installing mft-4.13.0...
```

```
Installing kernel-mft-dkms-4.13.0...
```



```
Installing libibcm1-41mlnx1...
Installing libibcm-dev-41mlnx1...
Installing perftest-4.4...
Installing ibutils2-2.1.1...
Installing libibdm1-1.5.7.1...
Installing cc-mgr-1.0...
Installing ar-mgr-1.0...
Installing dump-pr-1.0...
Installing ibsim-0.7mlnx1...
Installing ibsim-doc-0.7mlnx1...
Installing knem-dkms-1.1.3.90mlnx1...
Installing knem-1.1.3.90mlnx1...
Installing mxm-3.7.3112...
Installing ucx-1.7.0...
Installing sharp-2.0.0.MLNX20190922.a9ebf22...
Installing hcoll-4.4.2938...
Installing openmpi-4.0.2rc3...
Installing mpitests-3.2.20...
Installing libdapl2-2.1.10mlnx...
Installing dapl2-utils-2.1.10mlnx...
Installing libdapl-dev-2.1.10mlnx...
Installing srptools-41mlnx1...
Installing mlnx-ethtool-5.1...
Installing mlnx-iproute2-5.2.0...
Selecting previously unselected package mlnx-fw-updater.
(Reading database ... 148359 files and directories currently installed.)
Preparing to unpack ../mlnx-fw-updater_4.7-1.0.0.1_amd64.deb ...
Unpacking mlnx-fw-updater [4.7-1.0.0.1] ...
Setting up mlnx-fw-updater [4.7-1.0.0.1] ...
```

```
Added 'RUN_FW_UPDATER_ONBOOT=no' to /etc/infiniband/openib.conf
```

```
Attempting to perform Firmware update...
The firmware for this device is not distributed inside Mellanox driver: 86:00:0 [PSID: HPE0000000009]
To obtain firmware for this device, please contact your HW vendor.
```

```
Failed to update Firmware.
See /tmp/MLNX_OFED_LINUX.3524.logs/fw_update.log
Device [86:00:0]:
  86:00:0 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
  Link Width: x16
  PCI Link Speed: 8GT/s
```

```
Device [86:00:1]:
  86:00:1 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
  Link Width: x16
  PCI Link Speed: 8GT/s
```

```
Installation passed successfully
To load the new driver, run:
/etc/init.d/openibd restart
```

4. Load the new driver by running the following command:

```
user@ubuntugpu2:/tmp/MLNX_OFED_LINUX-4.7-1.0.0.1-ubuntu18.04-x86_64$ sudo /etc/init.d/openibd restart
```



Sample output:

```
Unloading HCA driver:           [ OK ]
Loading HCA driver and Access Layer: [ OK ]
```

Post-installation actions

After installation is complete, follow these steps:

1. Update IP settings if necessary by updating the `/etc/netplan/*.yaml` file.

- a. You can find a way to identify the logical name of the Mellanox Ethernet interface used to update the `*.yaml` file by running the command:

```
user@ubuntu-gpu1:~$ sudo lshw -class network
```

This command will list all Ethernet interfaces. The Mellanox Ethernet interface will have listed product: MT27800 Family [ConnectX-5] and vendor: Mellanox Technologies.

Sample output:

```
*-network:0
  description: Ethernet interface
  product: MT27800 Family [ConnectX-5]
  vendor: Mellanox Technologies
  physical id: 0
  bus info: pci@0000:86:00.0
  logical name: ens4f0
  version: 00
  serial: b8:83:03:58:d3:38
  width: 64 bits
  clock: 33MHz
  capabilities: pciexpress vpd msix pm bus_master cap_list ethernet physical autonegotiation
  configuration: autonegotiation=on broadcast=yes driver=mlx5_core driverversion=4.7-1.0.0
  duplex=full firmware=16.24.1000 [HPE0000000009] ip=10.20.100.125 latency=0 link=yes
  multicast=yes
```

- b. Identify the logical name and use this logical name to edit the network interfaces on the netplan `*.yaml` file.

Sample execution:

```
user@ubuntu-gpu1:/tmp$ cd /etc/netplan/
```

```
user@ubuntu-gpu1:/etc/netplan$ sudo vi /etc/netplan/01-netcfg.yaml
# This file describes the network interfaces available on your system
# For more information, see netplan(5).
```

```
network:
  version: 2
  renderer: networkd
  ethernets:
    eno1:
      dhcp6: no
      dhcp4: no
      addresses: [X.X.X.X/XX]
      gateway4: X.X.X.X
      nameservers:
        addresses: [X.X,X.X]
    ens4f0:
      dhcp6: no
      dhcp4: no
      addresses: [X.X.X.X/XX]
      gateway4: X.X.X.X
      nameservers:
        addresses: [X.X,X.X]
```

- a.) Reapply netplan if necessary

```
user@ubuntu-gpu1:/etc/netplan$ sudo netplan apply
```



2. Confirm activation by issuing the following command:

```
user@ubuntu-gpu1:~$ ibv_devinfo
```

Sample output:

```
hca_id: mlx5_1
  transport:          InfiniBand (0)
  fw_ver:            16.24.1000
  node_guid:         b883:03ff:ff58:d339
  sys_image_guid:    b883:03ff:ff58:d338
  vendor_id:         0x02c9
  vendor_part_id:    4119
  hw_ver:            0x0
  board_id:          HPE0000000009
  phys_port_cnt:     1
  Device ports:
    port:            1
                    state:          PORT_DOWN (1)
                    max_mtu:        4096 (5)
                    active_mtu:     1024 (3)
                    sm_lid:          0
                    port_lid:        0
                    port_lmc:        0x00
                    link_layer:      Ethernet

hca_id: mlx5_0
  transport:          InfiniBand (0)
  fw_ver:            16.24.1000
  node_guid:         b883:03ff:ff58:d338
  sys_image_guid:    b883:03ff:ff58:d338
  vendor_id:         0x02c9
  vendor_part_id:    4119
  hw_ver:            0x0
  board_id:          HPE0000000009
  phys_port_cnt:     1
  Device ports:
    port:            1
                    state:          PORT_ACTIVE (4)
                    max_mtu:        4096 (5)
                    active_mtu:     1024 (3)
                    sm_lid:          0
                    port_lid:        0
                    port_lmc:        0x00
                    link_layer:      Ethernet
```

Check the firmware version and update if needed. The latest firmware version when this document was published is 16.25.1020.

Identify the PCI address by using the following command:

```
user@ubuntu-gpu1:~$ lspci -v | grep Mellanox
```

Sample output:

```
86:00.0 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
86:00.1 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
```

3. Check the firmware version using the PCI address of the device by issuing the command:

```
~]# mstflint -d < device PCI address> q
```



Sample execution:

```
user@ubuntu-gpu1:~$ sudo mstflint -d 86:00.0 q
[sudo] password for user:
Image type:          FS4
FW Version:          16.24.1000
FW Release Date:     26.11.2018
Product Version:     16.24.1000
Rom Info:            type=UEFI version=14.17.11 cpu=AMD64
                    type=PXE version=3.5.603 cpu=AMD64
Description:         UID                GuidNumber
Base GUID:           b88303ffff58d338          8
Base MAC:            b8830358d338          8
Image VSD:           N/A
Device VSD:          N/A
PSID:                HPE00000000009
Security Attributes: secure-fw
```

- To update the firmware, download the firmware update from:
https://support.hpe.com/hpsc/swd/public/detail?swItemId=MTX_57d79dca3161495e963d90095e

- Install the firmware by issuing the command:

```
mstflint -d <device number of first device> -i <File path/<FW filename> burn
```

Sample execution:

```
user@ubuntu-gpu1:/tmp$ sudo tar -xzf ./fw-ConnectX5-rel-16.25.1020-872726-B21_Ax-UEFI-14.18.19-
FlexBoot-3.5.701.tar.gz
```

```
user@ubuntu-gpu1:/tmp$ sudo mstflint -d 86:00.0 -i /tmp/fw-ConnectX5-rel-16.25.1020-872726-
B21_Ax-UEFI-14.18.19-FlexBoot-3.5.701.signed.bin burn
```

```
Current FW version on flash: 16.24.1000
New FW version:             16.25.1020
```

```
Initializing image partition - OK
Writing Boot image component - OK
-I- To load new FW run mstfwreset or reboot machine.
```

- Reboot the server and confirm the firmware update.

```
ubuntu-gpu1:~$ sudo reboot
user@ubuntu-gpu1:~$ sudo mstflint -d 86:00.0 q
```

Sample output:

```
[sudo] password for user:
Image type:          FS4
FW Version:          16.25.1020
FW Release Date:     30.4.2019
Product Version:     16.25.1020
Rom Info:            type=UEFI version=14.18.19 cpu=AMD64
                    type=PXE version=3.5.701 cpu=AMD64
Description:         UID                GuidNumber
Base GUID:           b88303ffff58d338          8
Base MAC:            b8830358d338          8
Image VSD:           N/A
Device VSD:          N/A
PSID:                HPE00000000009
Security Attributes: secure-fw
```



ADDITIONAL LINKS

HPE and Mellanox

<https://www.mellanox.com/oem/hpe/>

Mellanox OFED documentation

<https://docs.mellanox.com/display/MLNXOFEDv471001>

HPE firmware download

https://support.hpe.com/hpsc/swd/public/detail?swItemId=MTX_57d79dca3161495e963d90095e

LEARN MORE AT

<https://h20195.www2.hpe.com/v2/GetPDF.aspx/c04154440.pdf>

Check if the document is available
in the language of your choice.



Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Share now



Get updates