



## **NVME-oF - NVM Express over Fabrics**

# Table of contents

NVME-oF

---

NVME-oF Target Offload

---

## NVME-oF

NVME-oF enables NVMe message-based commands to transfer data between a host computer and a target solid-state storage device or system over a network such as Ethernet, Fibre Channel, and InfiniBand. Tunneling NVMe commands through an RDMA fabric provides a high throughput and a low latency.

For information on how to configure NVME-oF, please refer to the [HowTo Configure NVMe over Fabrics](#) Community post.

### Note

The `--with-nvmf` installation option should **not** be specified, if `nvmf-tcp` kernel module is used. In this case, the native Inbox `nvme-tcp` kernel module will be loaded.

## NVME-oF Target Offload

### Note

This feature is only supported for ConnectX-5 adapter cards family and above.

NVME-oF Target Offload is an implementation of the new NVME-oF standard Target (server) side in hardware. Starting from ConnectX-5 family cards, all regular IO requests can be processed by the HCA, with the HCA sending IO requests directly to a real NVMe PCI device, using peer-to-peer PCI communications. This means that excluding connection management and error flows, no CPU utilization will be observed during NVME-oF traffic.

- For instructions on how to configure NVME-oF target offload, refer to [HowTo Configure NVME-oF Target Offload](#) Community post.

- For instructions on how to verify that NVME-oF target offload is working properly, refer to [Simple NVMe-oF Target Offload Benchmark Community post](#).

© Copyright 2024, NVIDIA. PDF Generated on 06/06/2024