Configuring NVMe over iWARP RDMA Fabrics

Quick Start Guide

NVMe over Fabrics (NVMe-oF) utilizes the existing NVMe technology and network interconnect like iWARP (RDMA over TCP) to connect subsystems with flash-based storage devices across datacenters.

Installation

Follow the steps mentioned below to install NVMe-oF supported kernel and related components on both target and initiator machines:

1. Download the driver package from http://service.chelsio.com/
2. Untar the tarball and change your working directory to *Chelsio-NVMeoF-x.x.x.x*

```
[root@host~]# tar zxvfm Chelsio-NVMeoF-x.x.x.x.tar.gz
[root@host~]# cd Chelsio-NVMeoF-x.x.x.x
```

3. On target, run the following command:

```
[root@host~]# ./setup.sh -t
```

   On initiator, run the following command:

```
[root@host~]# ./setup.sh -i
```

4. Reboot the system to installed 4.8+ (RC2) kernel.

Driver Loading

Follow the steps mentioned below on both target and initiator machines:

1. Load the following modules:

```
[root@host~]# modprobe iw_cxgb4
[root@host~]# modprobe rdma_ucm
```

2. Bring up the Chelsio interface(s):

```
[root@host~]# ifconfig ethX x.x.x.x up
```

3. Mount *configfs* by running the below command:

```
[root@host~]# mount -t configfs none /sys/kernel/config
```

4. On target, run the following commands:

```
[root@host~]# modprobe null_blk
[root@host~]# modprobe nvmet
[root@host~]# modprobe nvmet-rdma
```

   On initiator, run the following commands:

```
[root@host~]# modprobe nvme
[root@host~]# modprobe nvme-rdma
```

Copyright 2016. Chelsio Communications Inc. All rights reserved.

Configuration

Target

1. The following commands will configure target using nvmetcli with a LUN:

```
[root@host~]# nvmetcli
/> cd subsystems
/subsystems> create nvme-ram0
/subsystems> cd nvme-ram0/namespaces
/subsystems/n...m0/namespaces> create nsid=1
/subsystems/n...m0/namespaces> cd 1
/subsystems/n.../namespaces/1> set device path=/dev/ram1
/subsystems/n.../namespaces/1> cd ../..
/subsystems/nvme-ram0> set attr allow_any_host=1
/subsystems/nvme-ram0> cd namespaces/1
/subsystems/n.../namespaces/1> enable
/subsystems/n.../namespaces/1> cd ../../../..
/> cd ports
/ports> create 1
/ports> cd 1/
/ports/1> set addr adrfam=ipv4.
/ports/1> set addr trtype=rdma
/ports/1> set addr trsvcid=4420
/ports/1> set addr traddr=102.1.1.102
/ports/1> cd subsystems
/ports/1/subsystems> create nvme-ram0
```

2. Save the target configuration to a file,

```
/ports/1/subsystems> saveconfig /root/nvme-target_setup
/ports/1/subsystems> exit
```

3. Clear the targets,

```
[root@host~]# nvmetcli clear
```

Initiator

1. Discover the target

```
[root@host~]# nvme discover -t rdma -a <target_ip> -s 4420
```

2. Connect to target

   Connecting to a specific target:

```
[root@host~]# nvme connect -t rdma -a <target_ip> -s 4420 -n
<target_name>
```

   Connecting to all targets configured on a portal:

```
[root@host~]# nvme connect-all -t rdma -a <target_ip> -s 4420
```

3. List the connected targets. Format and mount the NVMe disks shown below.

```
[root@host~]# nvme list
```

4. Disconnect from the target and unmount the disk:

```
[root@host~]# nvme disconnect -d <nvme_disk_name>
```

NOTE: *nvme_disk_name* is the name of the device (E.g.: nvme0n1) and not the device path.

Copyright 2016. Chelsio Communications Inc. All rights reserved.