

NVMe-oF Training

Let MindShare bring “NVMe-oF” to life for you

MindShare's NVMe-oF (Non-Volatile Memory Express – over Fabrics) course begins with a brief review of NVMe basics, discusses the forces driving the migration of NVMe into network fabrics, followed by an overview of the different fabrics. Finally, the details of NVMe over fabrics are described.

You Will Learn:

- An overview of the storage fabrics in use today, and their strengths and weaknesses
- The changes needed for NVMe to work in a fabric environment
- How an NVMe host discovers and accesses drives on the fabric
- How NVMe-oF provides a combination of attractive features for a storage fabric

Who Should Attend?

This course is hardware-oriented, but is suitable for both hardware and software engineers because the registers used to control the hardware are described in detail. The course is ideal for RTL-, chip-, system- or system board-level design engineers who need a broad understanding of NVMe-oF.

Course Length: 1 Day

Course Outline:

- NVMe Background (as much as needed)
 - NVMe basics: registers, queue management, commands and completions
 - HCI shared-system-memory model
 - NVMe over PCIe
- Introduction to NVMe-oF
 - Motivation for moving NVMe to fabrics
 - Improve storage network performance compared to legacy protocols
 - Avoid limitations of PCIe, extend NVMe protocol to the fabric
 - Optimize access latency across the network
 - Differences from base NVMe
 - Initialization Background
- Fabric Discovery
 - Discovery Service, Discovery Controllers
 - Discovery Mechanism
- NVMe-oF Initialization
 - Fabrics Commands
 - Device Setup
 - Transport mapping/binding RDMA
 - Introduction to RDMA (Remote Direct Memory Access)
 - Protocol Translation
 - RDMA Examples
- RDMA Concepts
 - Motivation: Kernel bypass – reduce CPU involvement
 - Zero-copy operation when possible
 - RNIC (RDMA-aware Network Interface Card)
 - Pinned remote memory and storage
 - Command-based Flow Control
 - Read and Write Command Examples
- NVMe-oF Details
 - Controller Initialization
 - Connections and Associations
 - Static vs. Dynamic Controller model

- Properties instead of registers
- Authentication
- Encapsulation
- Accessing RDMA Queue pairs
- Data transfer: Command, Response, Data, Flow Control
- Fabrics Commands and Responses
- RNIC Details
 - Software Layers
 - RDMA Verbs
 - Memory Regions and Windows
- Ethernet overview
 - TCP/IP: addressing, layers, packets
 - RDMA models added to it: RoCE and iWARP.
 - Overview of iWARP
 - Added to TCP/IP, so packets are routable and don't need DCB
 - Messages
 - Layers
 - Advantages and disadvantages
 - Overview of RoCE
 - Version 2 packets include routing information, version 1 did not
 - Use of Data Center Bridging (DCB) to manage routing
 - Flow control
 - Soft RoCE
 - Advantages and disadvantages

Recommended Prerequisites:

Previous exposure to NVMe architecture is highly recommended.

Understanding the operation of a high-speed serial transport like PCIe is highly recommended.

Course Material:

A downloadable PDF version of the presentation slides.