

NVMe over Fabrics (NVMe-oF) Training

Let MindShare bring “NVMe over Fabrics (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 NVMe-oF will work over such as Ethernet, InfiniBand and Fibre Channel. 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
- How NVMe protocol will work over Ethernet, InfiniBand and Fibre Channel
- An overview of Ethernet including RoCE and iWARP, InfiniBand and Fibre Channel Protocols

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
 - NVMe basics: registers, queue management, commands and completions
 - HCI shared-system-memory model
 - NVMe over PCIe
- Why move NVMe to fabrics?
 - Access to storage networks of thousands of SSDs
 - Avoid protocol translation to get to those networks
 - Optimize access latency across the network
 - Allow systems to work using only remote storage (no local storage)
- Improving network access with RDMA (Remote DMA)
 - Kernel bypass – reduce CPU involvement
 - Zero-copy operation
 - Reliable, lossless delivery
 - Credit-based flow control
 - Use of common primitives
 - Pinned remote memory and storage
 - Encapsulation avoids need for protocol translation
- Ethernet overview – popular, high speed, inexpensive
 - TCP/IP: addressing, layers, packets
 - Advantages and disadvantages compared to other fabrics
 - RDMA not natively supported by Ethernet, so two competing models were developed to add support for it: RoCE and iWARP. Both require RDMA-capable NICs (RNICs).
 - 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

- Overview of iWARP
 - Added to TCP/IP, so packets are routable and don't need DCB
 - Messages
 - Layers
- InfiniBand overview
 - Designed for RDMA
 - Messages
 - Discovery
 - Packets
 - Routers
 - Physical layer encoding schemes
 - Changes to support NVMe
- Fibre Channel overview
 - Addressing
 - Initialization
 - Discovery
 - Packets
 - Physical layer encoding schemes
 - Changes to support NVMe over FC (NVMe-oFC)
- NVMe over Fabrics (NVMe-oF) Details
 - Transport mapping/binding: Fibre Channel or RDMA
 - Discovery Service
 - Connections
 - Properties
 - Authentication
 - Capsules
 - Queue pairs
 - Data transfer: Command, Response, Data, Flow Control

Recommended Prerequisites:

Previous exposure to NVMe architecture is highly recommended.

Course Material:

A downloadable PDF version of the presentation slides