

VMware vSphere NVMe features in 2023 and beyond

Arvind Jagannath

Sr Product Line Manager, vSphere Platform

Murali Rajagopal

Sr Dir of Technology, OCTO

VMware High-Level Updates

- NVMe[®]/FC continues to be dominant
- NVMe/TCP betters NVMe/RDMA in terms of adoption

Continued focus on

- Scale and Performance
- Discovery
- Security
- Operations (Boot etc.)


- VMware Compatibility Guide - “VMware’s HCL”

https://www.vmware.com/resources/compatibility/search.php?deviceCategory=io&details=1&keyword=nvme&page=3&display_interval=500&sortColumn=Partner&sortOrder=Asc

- ~2420 NVMe listings supporting the latest and greatest HBA/Firmware
 - (Up from 1900 in 2022!!)
 - 163 for NVMe/FC, 688 for NVMe/RDMA
- Supporting a large ecosystem of IO vendors, Array vendors, OEMs
 - NVMe Flash/SSDs (**Storage** vendors)
 - **HBAs** (RAID, Tri-mode, RDMA, FC), NIC (supporting NVMe/TCP)
 - **OEMs** supporting Direct-attached NVMe or NVMe-oF
 - **Array vendors** supporting (NVMe-oF /RDMA /FC and /TCP)
- **62 storage arrays certified with NVMe/TCP including NetApp & Pure!!**
- **148 for NVMe/FC**
- **49 for NVMe/RDMA**

vSphere Released Features

Released (vSphere 8.0/8.1) NVMe Features



Discovery


- **NVMe-oF Discovery**
 - TP 8002 - Full support
- **NVMe-oF Central Discovery Controller (CDC)**
 - TP 8010 – incl. support for frame format
- **Automated Discovery of IP Discovery Controllers (NVMe/TCP)**
 - TP 8009 - Full support
 - mDNS daemon
 - ESXi traps & processes mDNS broadcast messages post NVMe adapter creation
 - CDC/DDC supported

Protocol and Stack

- **NVMe-oF Reservation support for clustered VMDK (WSFC)**
 - FC only
- **NVMe-oF Scale Improvements (VMFS based)**
 - 32 → 256 NS, 128 → 2K Paths
- **Increased max paths per NVMeoF namespace from 8 to 32**
 - Limit is still 4K Paths
- **Storage DRS (Distributed Resource Scheduler) for NVMe/TCP**
- **E2E NVMe support**
- **Max Data Transfer for non-IO Commands (MDTS)**
 - TP 4040; untested

vSphere NVMe Roadmap

vSphere 8.0U2 and beyond....



Security

- NVMe-oF In-Band Authentication
 - TP 8006
- Unique Discovery Controller Identification
 - TP 8013 and TP 8017
- Authentication Verification Entity (AVE)
 - TP 8019
- TLS 1.3 for NVMe/TCP
 - TP 8011
 - TP 8018
- NVMe-oF Security Configurations
 - TP 8025

Protocol and Stack

- NVMe/TCP Reservation support for clustered VMDK (WSFC)
 - WEAR only
- Boot Over NVMe/TCP
 - TP 8012
 - TP 4146 - Support for consistent HostID and HostNQN values
- MetroCluster NVMe
 - Multi NVM subsystem
 - TP 4034 Dispersed Namespaces
- Namespace Symbolic Name - TP 4145
- CDC symbolic names for hosts and subsystems - TP 8020
- Advanced Command Retry Enable (ACRE) - TP 4033
- Boot
 - Authentication keys provisioning via BMC (Redfish Interface)
 - Support for encrypted boot
- Copy Across Namespaces – TP 4130
- Live Migration - TP 4159 (PCIe only) - early spec

Questions?