



Flash Memory Summit

Benefits and Use Cases for NVMe-oF

Mellanox NVMe SNAP Use Case
Oren Duer | Mellanox Technologies



Seamless Disaggregation

Motivation

Move local NVMe drives to centralized location

- No local disks needed (disk-less)
- Grow storage or compute independently
- Higher performance per node
- Immediate CAPX saving
- Lower MTBF

Problem

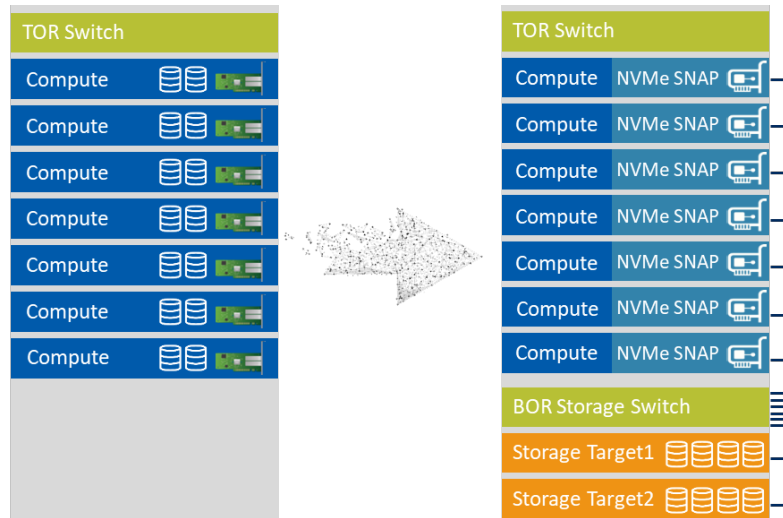
Requires software changes

- RDMA software stack
- NVMe-oF drivers – limited OS support
- Different management

Solution

NVMe SNAP

- Compute nodes see NVMe local drives
- Zero software changes
- Supported on all OSs
- Latency as local NVMe drive
- Bandwidth up to network available (100Gbps and above)

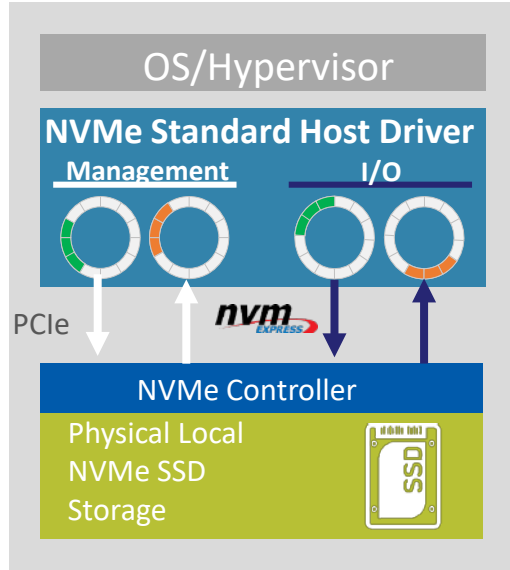




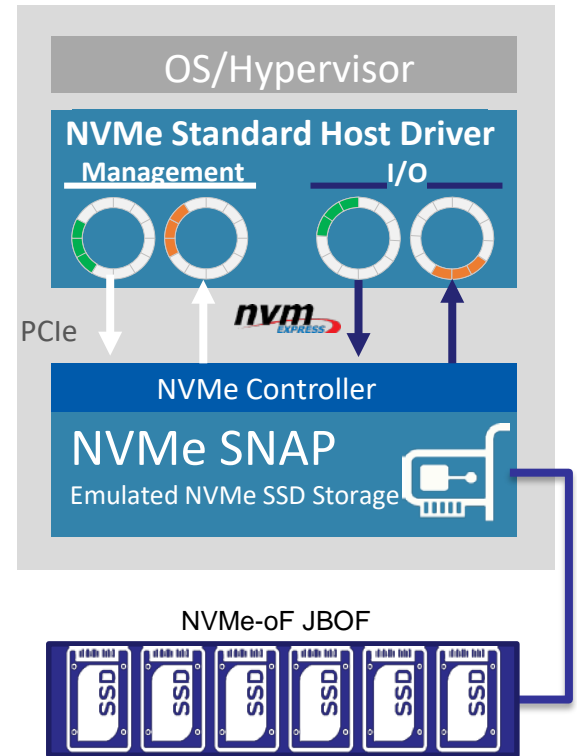
NVMe SNAP

- Emulated NVMe PCI drives
- OS agnostic
- Software defined
- Hardware accelerated
- Bootable
- NVMe SRIOV support

Host Server – local disk



Host Server – NVMe SNAP



Remote Storage



NVMe SNAP internals

SPDK advantages

- Efficient memory management
- Zero-copy all the way
- Full polling
- Multi queues, multi threads, lockless
- Well defined APIs: vBdev, Bdev drivers...

NVMe emulation SDK

- Handle NVMe registers and admin

Customer's proprietary code

- **BDEV**: for proprietary storage network protocols
- **vBDEV**: for per-io routing decisions, RADIs, etc

