NVMeoF Everywhere

Presenter: Odie Killen, VP of Engineering at Viking Enterprise Solutions

FMS the Future of Memory and Storage

Overview

- NVMeoF Everywhere
- Advantages
- Disadvantages
- CPU Based HDD Solution
- Sample Solution
- Native NVMe HDD Solution
- Sample Solution
- Ethernet Connected HDD
- Path Forward

NVMeoF Everywhere

- NVMeoF is a fabric protocol that allows the Ethernet interface to connect devices together
  - NVMe is a lighter protocol than others, resulting in better utilization of the CPU cycles and available system BW
  - Lower latency and improved management and provisioning of storage
  - Better enablement to disaggregate compute and storage
- Originally targeted to connect SSDs to PCIe busses, but can be extended to connect other devices, such as rotating HDDs to PCIe busses
- Higher cost of acquisition, however, lower operating costs and simplified enterprise management
- For storage to migrate to NVMe everywhere, there must be a HDD solution in this space

Advantages

- Common interface throughout the enterprise
- Leverages existing Ethernet connectivity
  - Eliminates multiple protocol interfaces within infrastructure
  - Reduced infrastructure cost
  - Simplifies management of infrastructure
- We all know how to manage Ethernet based devices
- Enables a common pane of glass approach to enterprise management
- More efficient utilization of the BW and available resources

Disadvantages

- Not all device types are available with NVMeoF enabled connectivity
- Not all storage device tiers are available in NVMe
  - HDD tier used for online, nearline and offline/cold storage is not natively NVMe
  - HDD tiers can be made to support NVMe front end, but this adds significant acquisition cost
- Ways to enable NVMe capable HDD-based solutions
  - Add a compute node with NVMe adapter running in target mode
  - Add a native NVMe front end to HDD-based IOMs
  - Ethernet connected HDDs
- Both approaches increase the acquisition cost of the solution while enabaling the elimination of SAS (or other protocol) from the enterprise

CPU Based HDD Solution

- Simplest approach is to utilize a storage server with a NVMe NIC to provide native NVMe connection
- Not a preferred solution
  - Much higher cost due to CPU complex cost
  - Requires target mode SW to be written for CPU and NIC
  - Way too much CPU power for the task, so not an efficient use of solution
  - CPU complex takes a lot of space, resulting in lower density drive solution
- Will work, but not a cost effective solution

Sample Solution

- 4U Chassis with 60-66 HDDs
- CPU based solution (x86 class CPU)
- Support for native NVMe connection to HDDDs
- Advanced capabilities of CPUs are often under-utilized

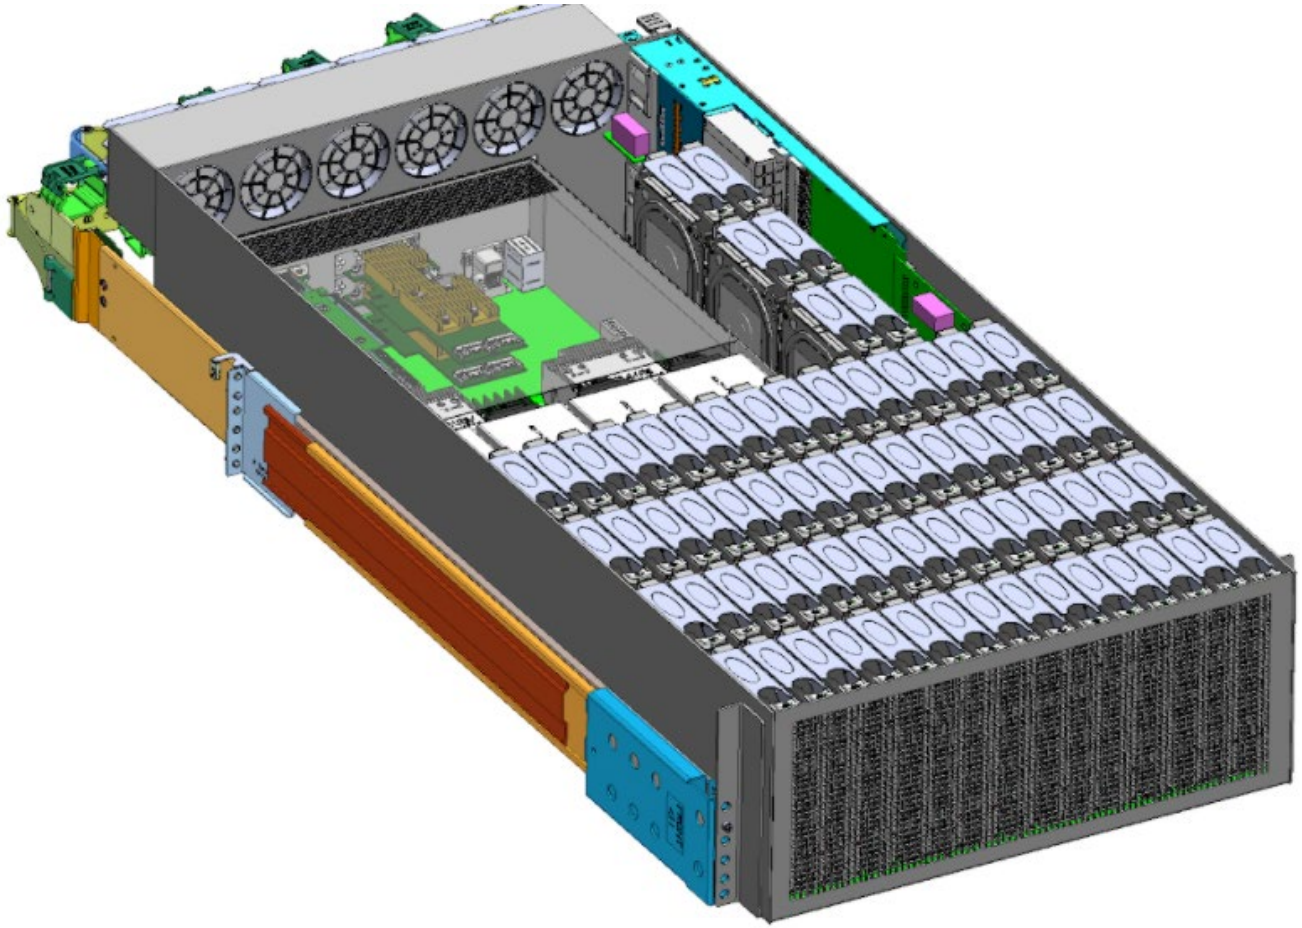

Native NVME HDD Solution

- Incorporate a NVMe NIC and SAS HBA into an enclosure
- Provides a native NVMe front end interface for enclosure, while supporting SAS HDDs
- Provides the most cost effective solution
  - NIC silicon is much cheaper than CPU
  - More effective use of space, resulting in higher density enclosure
  - SW for NIC and translation between NVMe/SAS is still required
  - Lower power consumption
- This is the most cost effective solution for enabling NVMe attached HDDs today

Sample Solution

- Full High Availability solution
- 4U-102 HDDs
- Native NVME connection
- Advanced features of NIC are readily available
- Lower cost of acquisition

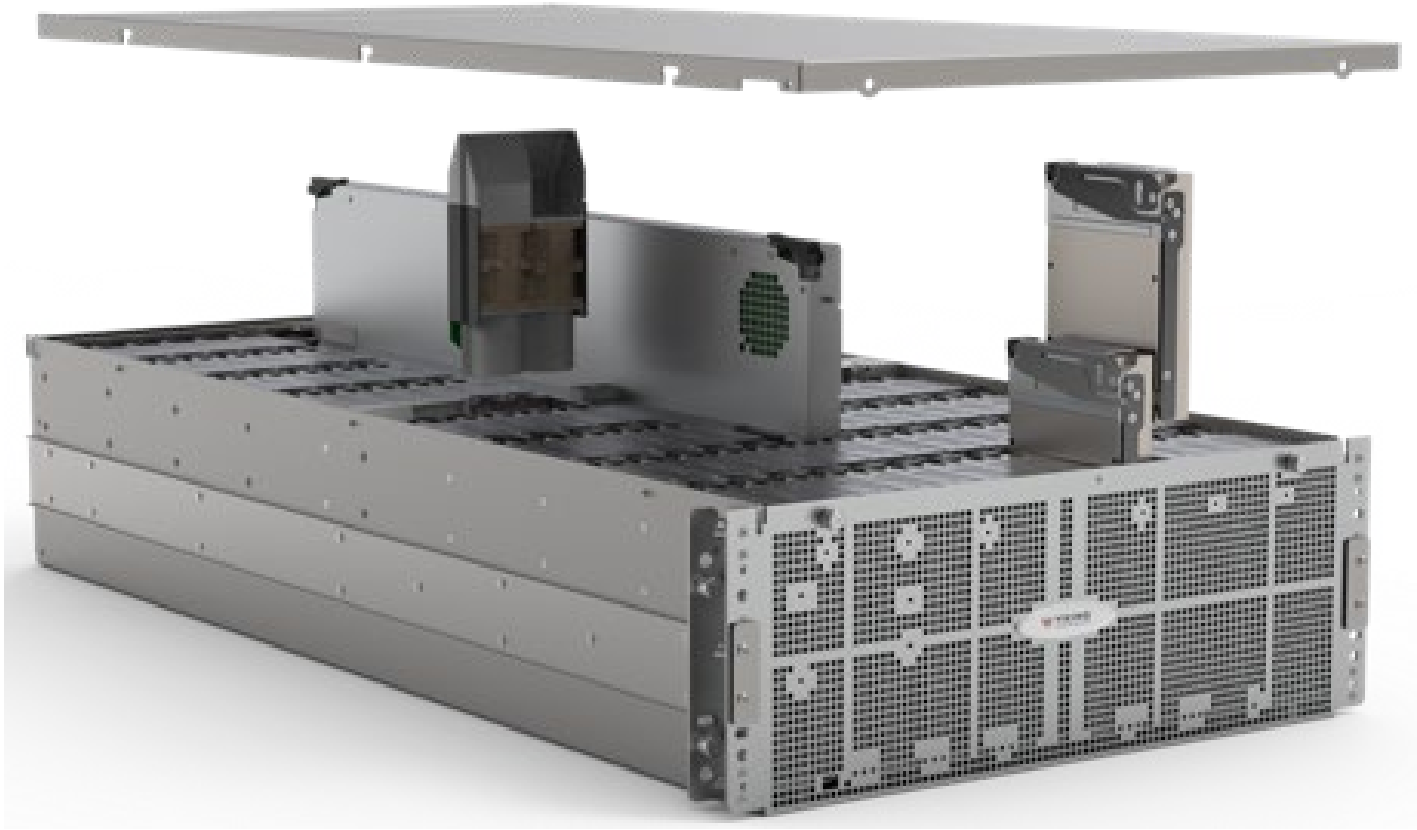

Ethernet Connected HDDs

- Concept has been around for a number of years
- HDD interface is replaced with a native Ethernet or NVMe front end
- Simplifies enclosure design and reduces new SW content
- Large scale adoption may be limited unless 2 of the 3 HDD vendors adopt this approach
- This approach may take flight in the future, however, historically this has not been successful

Path Forward

- Push the enterprise towards NVMeoF and Ethernet as the common fabric connection
- Reduce/Eliminate SAS and other protocols from the enterprise
- Leverage exiting tools to implement a single pane of glass management approach for the enterprise
- Adoption of NVMe enabled HDD solutions for storage tiers
- Migrate the entire storage tier to NVMeoF