Balancing HDD and SSD in Your Deployment

Presenter: Odie Killen, Vice President of Engineering at Viking Enterprise Solutions

FMS the Future of Memory and Storage

Overview

- Typical Deployment Needs
- Tier Make Up
- NVMeoF Everywhere
- NVMeoF Everywhere – Advantages
- NVMeoF Everywhere – Challenges
- Flash Options – Tier 0
- Flash Options – Tier 1
- Flash Options – Tier 2
- Flash Options – Tier 3
- Wrap Up

Typical Deployment Needs

- A enterprise deployment consists of several tiers, each serving a specific need
- Typically 4 tiers
  - (0) Mission Critical; high performance workloads that are latency critical (100% uptime)
  - (1) Hot Data; Used daily for business operations, but not latency critical (>99.999% uptime)
  - (2) Warm Data; Must have access, but accessed rarely (<99% uptime)
  - (3) Cold Data; Archived or rarely used data (<99% uptime)
- Each tier has a function and that function dictates type of storage needed
  - Lower tiers can be less reliable and are less expensive
- Flash can be used in all tiers, however flash is not needed for all tiers
  - As capacity requirements increase, flash quickly becomes prohibitive in lower tiers

Tier Make Up

- Most of the stored data, is located in the bottom tiers
  - Very cost sensitive, suggests that HDDs are not going away any time soon
- Regulatory and government requirements drive larger content in lowest storage tier
  - HIPPA, Email Storage, etc.
- Pay for Performance
  - Not just an HR concept
  - Tier 0 is the most expensive tier in terms of $/GB
  - Tier 3 is the lest expensive tier in terms of $/GB
  - Tier 0 through 3 are similar in cost in terms of $/IOPs
- We can simplify tier management and optimize cost
  - NVMeoF Everywhere

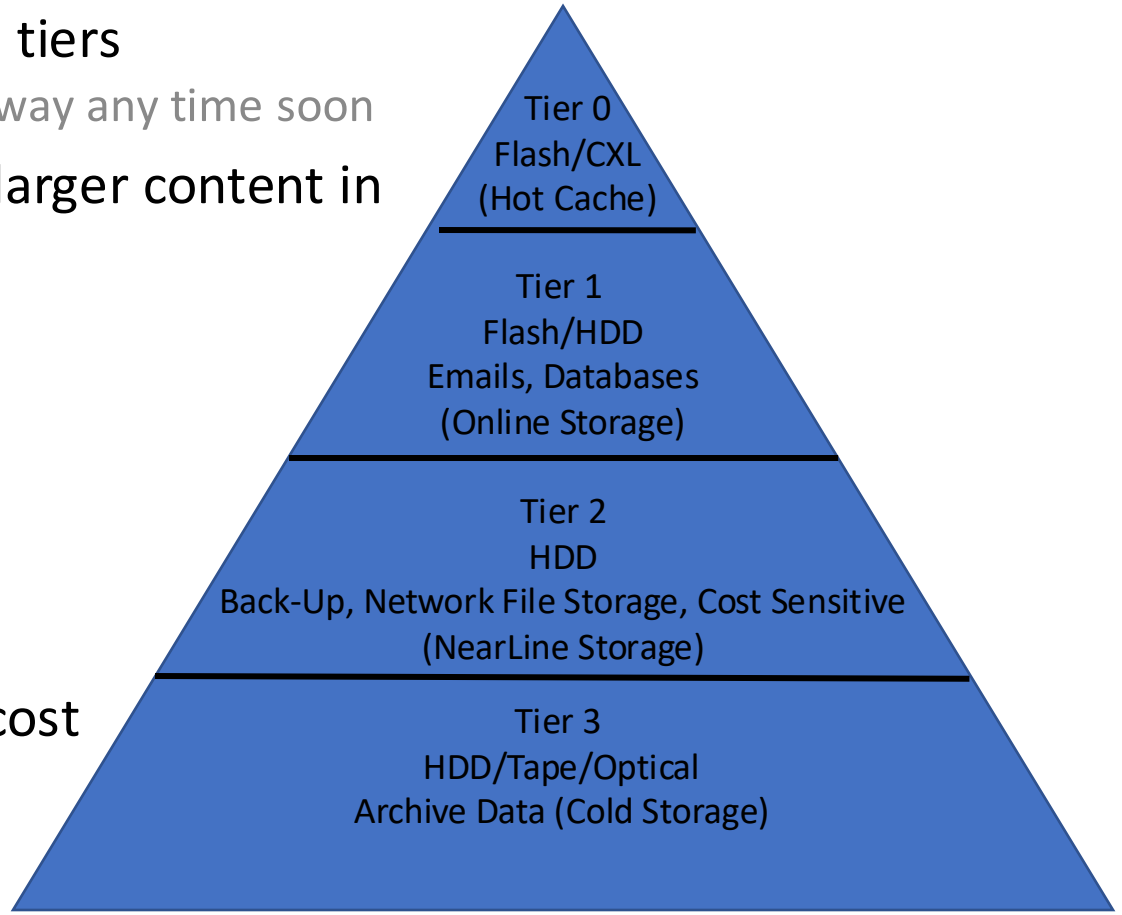

NVMeoF Everywhere

- 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

NVMEoF Everywhere – 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

NVMEoF Everywhere – Challenges

- 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 enabling the elimination of SAS (or other protocol) from the enterprise

Flash Options – Tier 0

- Tier 0 is uniquely positioned for CXL and the promise of shared memory
  - Shortens latency to access data shared between multiple CPUs
  - Simplifies SW and improves performance
  - CXL needs to continue to mature before it will take hold in this space
- Currently implemented with NVMe based SSDs (includes Storage Class Memory)
  - Provide best in class performance (latency and throughput)
  - Gold standard for solutions today
  - Light weight NVMe protocol makes CXL the only contender to this technology in this tier
  - Dual ported SSDs provide access to data from two hosts
  - Switch fabrics can extend access to more than 2 hosts
  - Single ported SSDs offer lower acquisition cost solutions
- CXL is promising, but NVMe SSDs are still preferred in this space

Flash Options – Tier 1

- This tier was historically implemented with 2.5” HDDs
  - 2.5” HDDs had better throughput and lower latency than 3.5” HDDs
  - Less relevant in this space now due to flash cost and capacity
- Emergence of flash, especially SAS SSD, greatly disrupted this space
- SAS is still the best acquisition cost option in this space
  - Lower acquisition cost versus NVMe
  - Multiple SAS SSD vendors in market place
  - Acceptable performance for this tier
  - Legacy SW and applications leverage this interface
  - Some legacy FC Hosts in this space also, but SAS SSDs are optimal storage media
- This tier greatly benefits from a NVMeoF Everywhere approach with backend SAS SSDs

Flash Options – Tier 2

- Commonly implemented as a Near Line storage tier
- Bifurcation point for HDD versus Flash
- Flash is not advised in this tier
- Benefits of flash are not required for this space
  - Typically attached to a slower performing and heavily subscribed network
  - Data is not required to be immediately accessible
  - Longer latency time is acceptable
  - Not running mission critical
  - Larger tier so more sensitive to cost per GB
- A balanced topology will not include Flash in this tier
- This tier greatly benefits from a NVMeoF Everywhere approach with backend SAS HDDs

Flash Options – Tier 3

- NONE
- This tier is often referred to as cold storage
- Majority of data is written once and may never be read
- SATA HDDs and LTO dominate this space
- Limited optical play, with emerging options coming
- Flash should not be considered here
- This tier greatly benefits from a NVMeoF Everywhere approach with backend SATA HDDs

Wrap Up

- Flash has its place and is not replaceable in certain tiers (0 and 1)
- A balanced enterprise will include HDDs along with Flash
- This balanced approach drives several different storage interfaces into enterprise
  - NVMe or CXL based flash (Tier 0)
  - SAS SSDs or lower cost/performance NVMe SSDs (Tier 1)
  - SAS and SATA HDDs (Tier 2/3)
  - LTO or Optical storage solutions (Tier 3)
- Managing the properly tiered enterprise is challenging due to various interconnects
  - Enterprise management could be greatly simplified with an NVMeoF Everywhere approach
  - Enables a common pane of glass approach to enterprise management