



Forum D-11: Flash in Data Centers

Tuesday, August 5

Tiering for Converged Flash Storage

Andy Mills, Co-founder/CEO Enmotus, Inc.

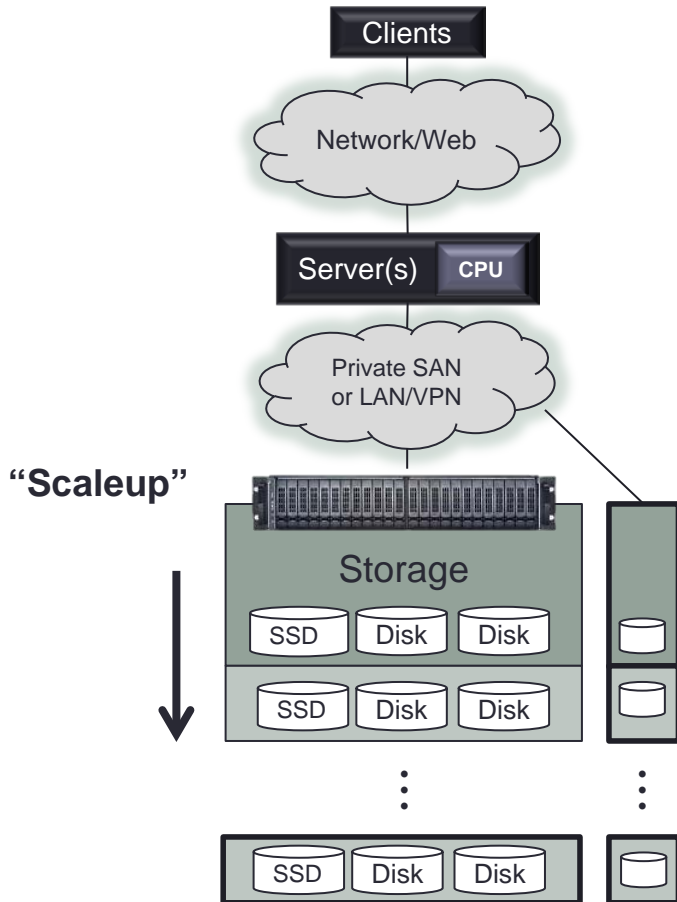
Data Center Trends



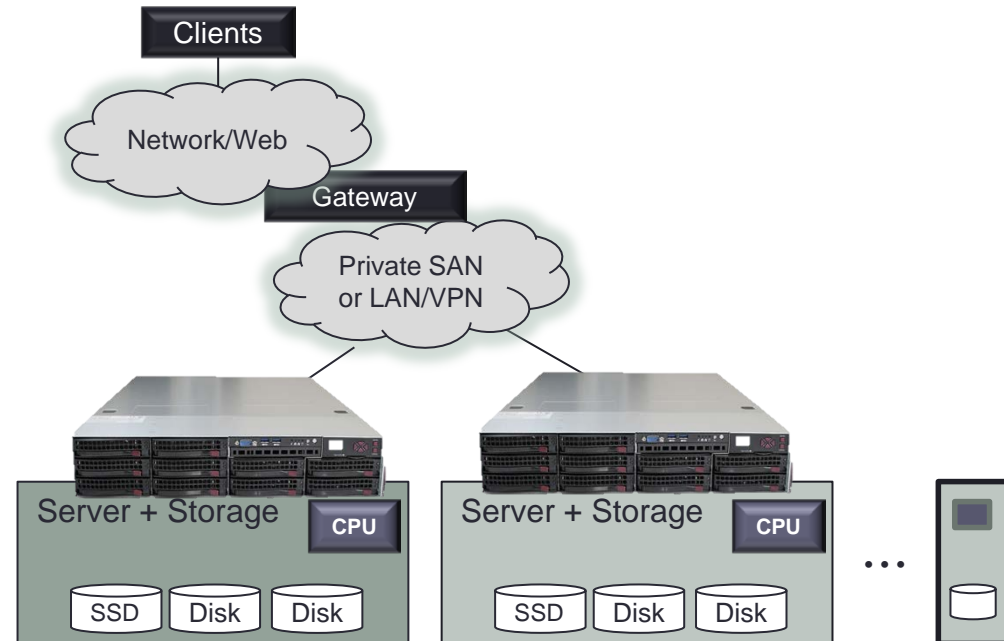
- **Convergence of Compute and Storage**
 - Big Data, Scaleout - networked commodity storage-servers
 - Collapsed SAN - all in a single node - looking more like DAS again
- **Software Defined Everything**
 - SDS - Software Defined Storage
- **Device Level Convergence**
 - Virtualization at device level within a node
 - Creation of virtual devices from flash, fast HDDs, slow HDDs

Scaleout vs Scaleup

Typical Enterprise Data Center



Distributed, Clustered or Hybrid Data Center



- Better suited to handle large unstructured data environments e.g. Web search, data analytics

Flash in Data Centers



**Memory Class
NVDIMM**

1 million IOPs+



PCIe SSDs

150-700K IOPs



**Small Form Factor
SAS, SATA**

20-90K IOPs



**Embedded
mSATA/M.2**

20-90K IOPs

(Hard Disk Drives range from 80-350 IOPs)

Flash Integration in Data Centers



- **Early Solutions Focused on Ease of Integration**
 - SAS/SATA SSDs emulate hard disk drives
 - Drop into existing or modified disk arrays
 - Preserve legacy controllers and storage management tools
- **Introduction of PCIe SSDs**
 - 20x+ improvement over first generation
 - Server based SSD acceleration/caching
 - New island of storage, new management tools
- **NVDIMM memory class**
 - Another new island of storage...

Converged Flash-Legacy Storage Generations



1st generation SSD Caching

- SSD caching in Disk Arrays
- Server side SAN caching
- Flash is look aside

2nd generation SSD Caching Mark II

- Increased intelligence to address application behavior
- Flash is look aside

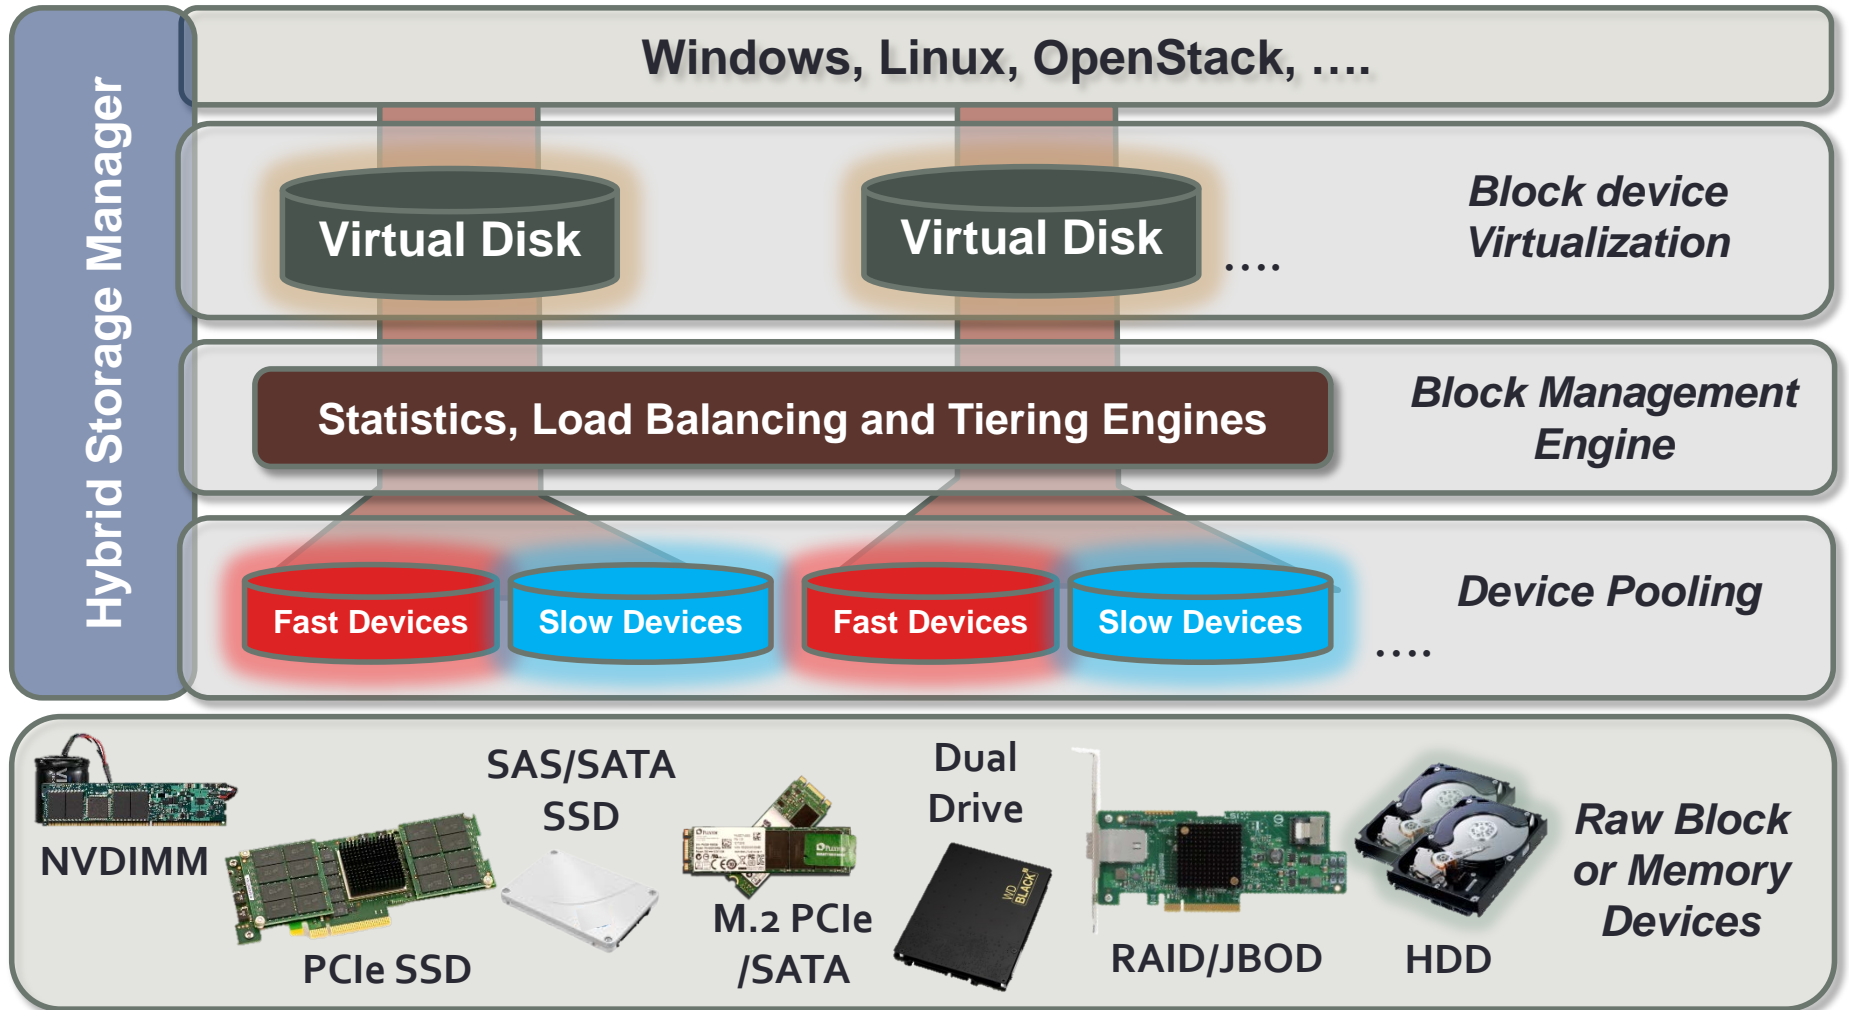
3rd generation Software Defined Storage

- Full virtualization model
- Flash is primary storage
- Higher performance, broader flash device support


Third Generation Flash Convergence Technology

- 1st/2nd Gen SSD caching challenges
 - Size restrictions
 - As cache capacity increases, performance tails off
 - Larger cache capacities do not contribute to usable storage pool
 - Effectiveness is increasingly hard to measure
- 3rd Gen SSD Virtualization and Auto Flash Tiering
 - Full virtualization abstracts several types of flash media
 - SSDs become usable as primary storage OR cache
 - Achieve full SSD performance for both reads AND writes
 - No capacity limits, any ratio of storage may be used

Software Defined Flash Convergence




Monitoring and Visual Mapping



- Visual at-at-glance tools are important to ensure active data is truly on the flash portion of the storage tier

Monitoring Flash Activity Levels vs. Hard Drives



- Ability to monitor activity over time to ensure most activity is being served off the fast tier

Wrap up

- Industry moving to third generation software defined storage that is flash friendly
- Fully virtualized, fully transparent
- Fully automated load balancing across all flash components and legacy storage
- Performance, improved visibility and management tools are key elements of the third generation