



Integrating SSDs into Virtual Servers

Opportunities and Challenges


Andy Mills, Enmotus, Inc.

- Market and Technology Trends
- Virtual Servers and Storage I/O
- Observed Performance Numbers
- SSD Caching and MicroTiering

Data warehousing,
database servers



Data Intensive, Web Servers
Cloud/Grid Clusters
(Social Networks and Emerging Enterprise)



Virtualized Hosting/Cloud
Computing



Small
Medium
Business



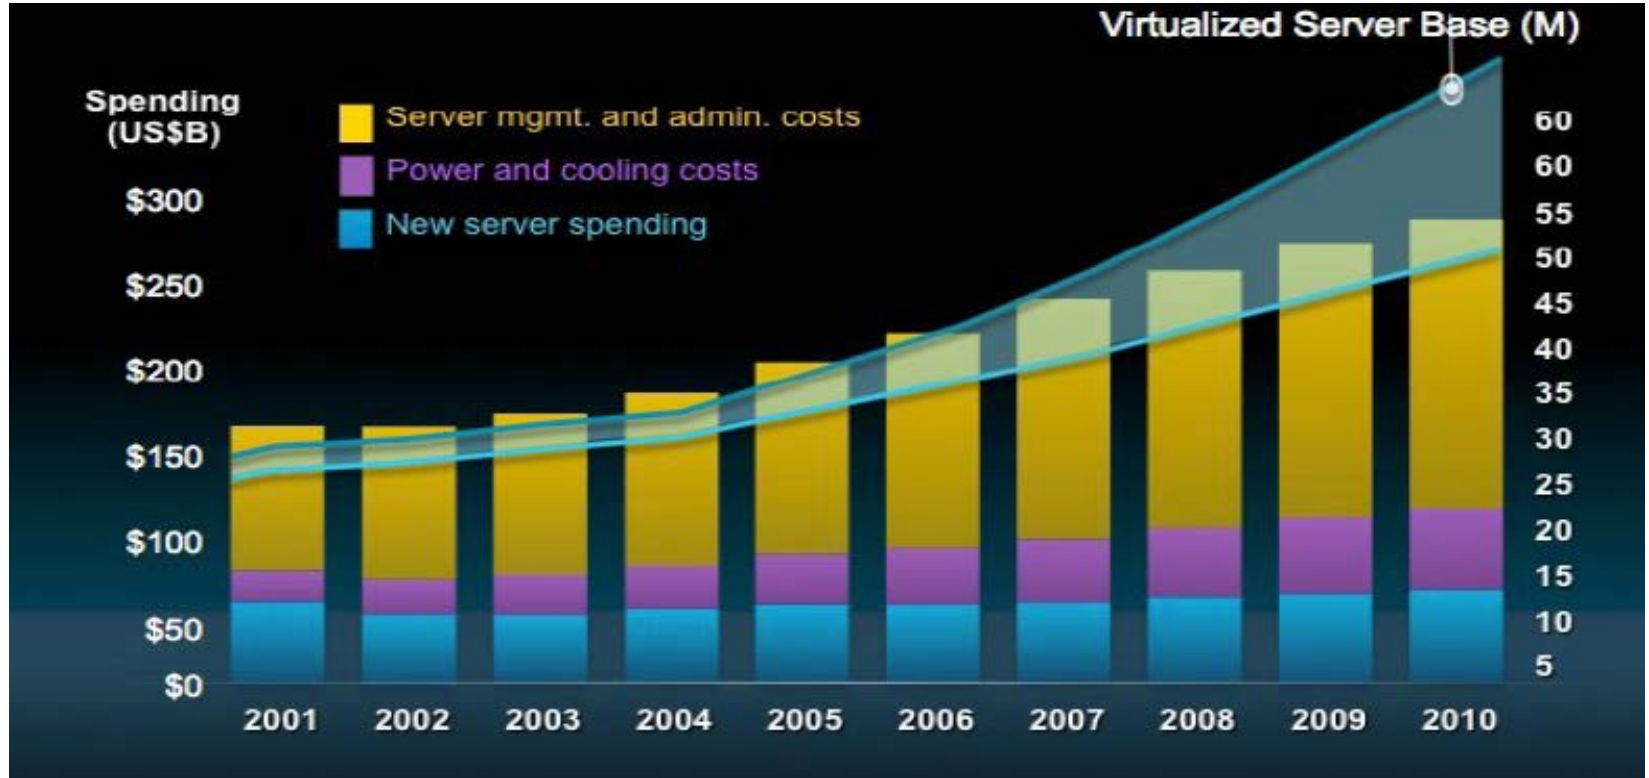
Servers

- ~8M servers annually, 23% unit growth, 11% revenue (IDC)
- ~1.8m virtual machine enabled physical servers ship annually
- ~1.3m database servers ship annually
- 700K-1M servers new social network server build-outs
- ~\$12BN spent on storage management software

SSDs/Flash Storage

- Server-storage performance gap widening
 - 20-100x+ raw performance gains
 - 2-10x+ transaction performance gains
 - Reduction in power - 100-1000%
 - Drives trend back to DAS storage

Server Trends




Virtualized servers are growing rapidly and management costs along with them...

Source: Cisco , IDC, Server Summit 2011
<http://www.serverdesignsummit.com/>

Issues for Virtual Servers

- Virtual server environments behave very differently creating the infamous IO blender effect
- SSDs are migrating back into the server creating several integration and reliability issues
- Utility based computing and virtual machines have created a fluid computing environment making the setup more dynamic and unstructured
- Growing trend toward smaller, lower cost clustered virtualized servers that are SAN-less
- Storage virtualization strategies needed for these emerging systems and less sophisticated users


IO Blender Effect



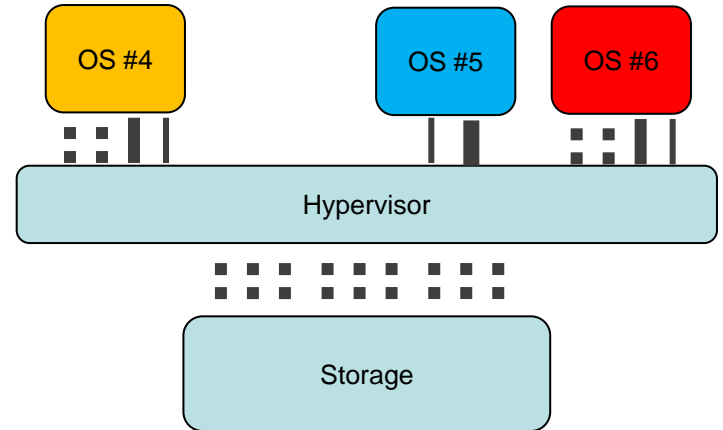
Sequential Streams are turned to Random
Ideally would like the SSD to be servicing the highest activity OS

Live Machine Migration

Physical Server A



Physical Server B

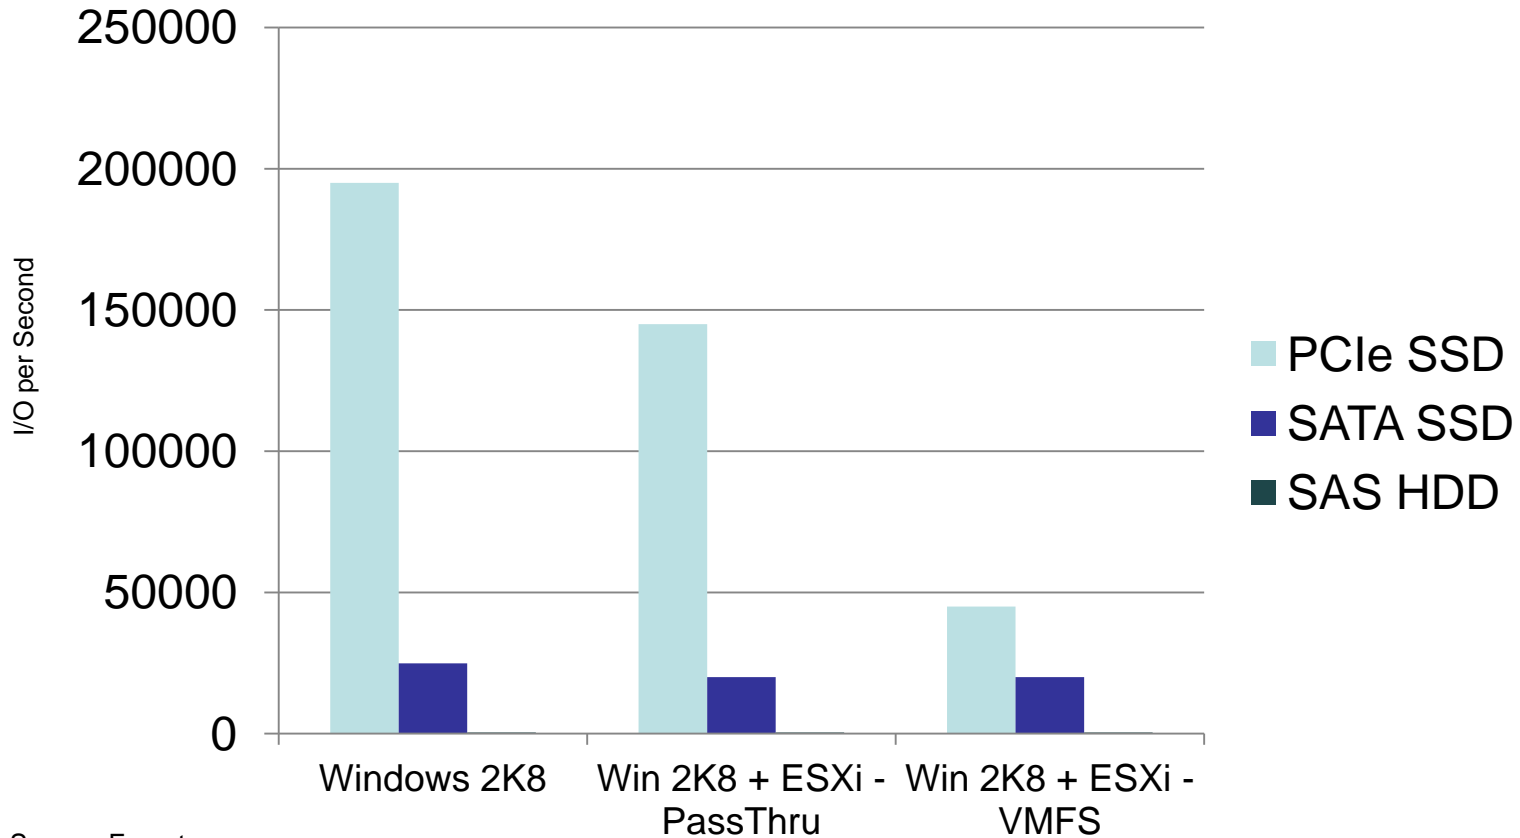


Moving or Creating a New VM Changes the Storage IO Balance

SSD Integration – Today's Options

- **Software Based Caching and Tiering**
 - Client OS only e.g. EMC Fast
 - Hypervisor + Client Driver e.g. FlashSoft/Sandisk
 - Hypervisor e.g. Proximal
 - SSD type agnostic e.g. SATA/SAS SSD or PCIe SSD
- **Hardware Accelerated**
 - RAID adapters with SSD caching
 - PCIe SSDs – new class of storage
 - Intelligent Storage IO Processors – MicroTiering
- **Hybrid Hardware-Software**
 - PCIe SSDs with host based flash management - FusionIO

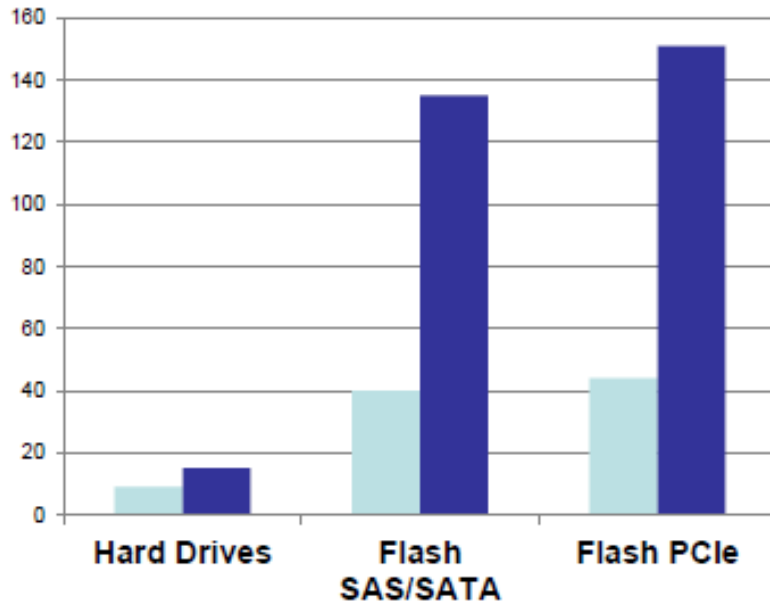
Raw SSD Performance



Source: Enmotus

Dell R515, Iometer 2006, 512B Random Read IOPs

Application Performance



TPM (in 000s) using DBT-2 at 1,000 warehouses

SchoonerSQL based on 5.1.44
Stock MySQL 5.1.44

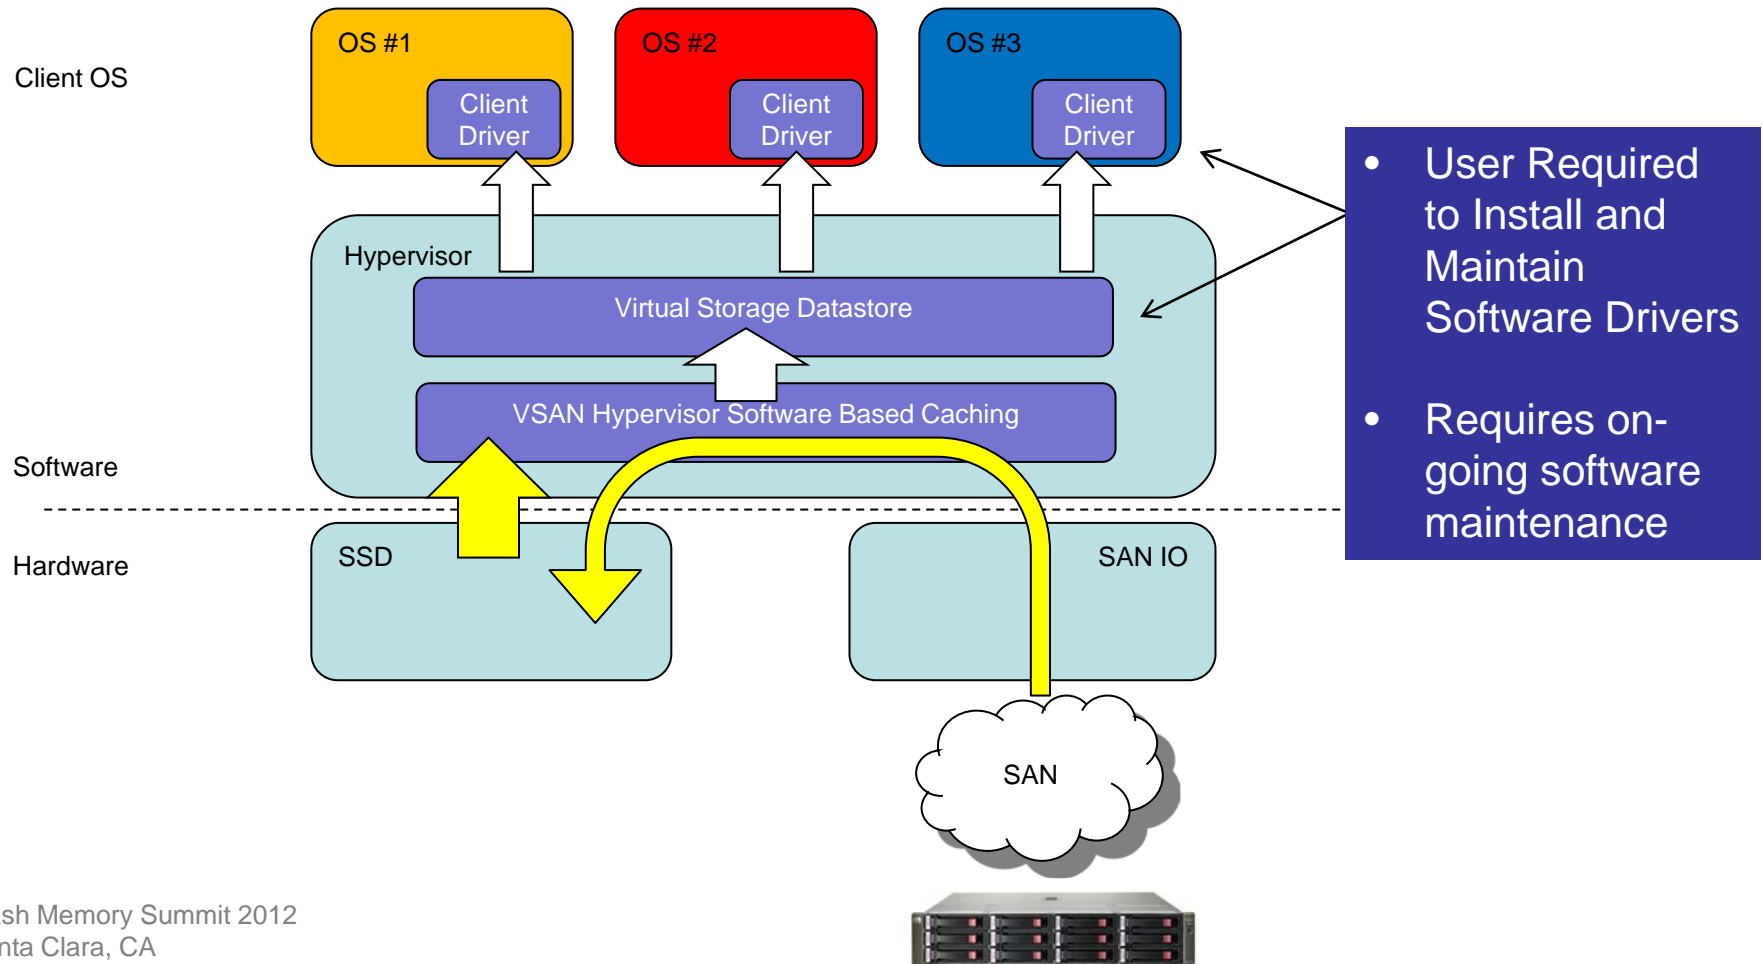


Source: Schooner Information Technology

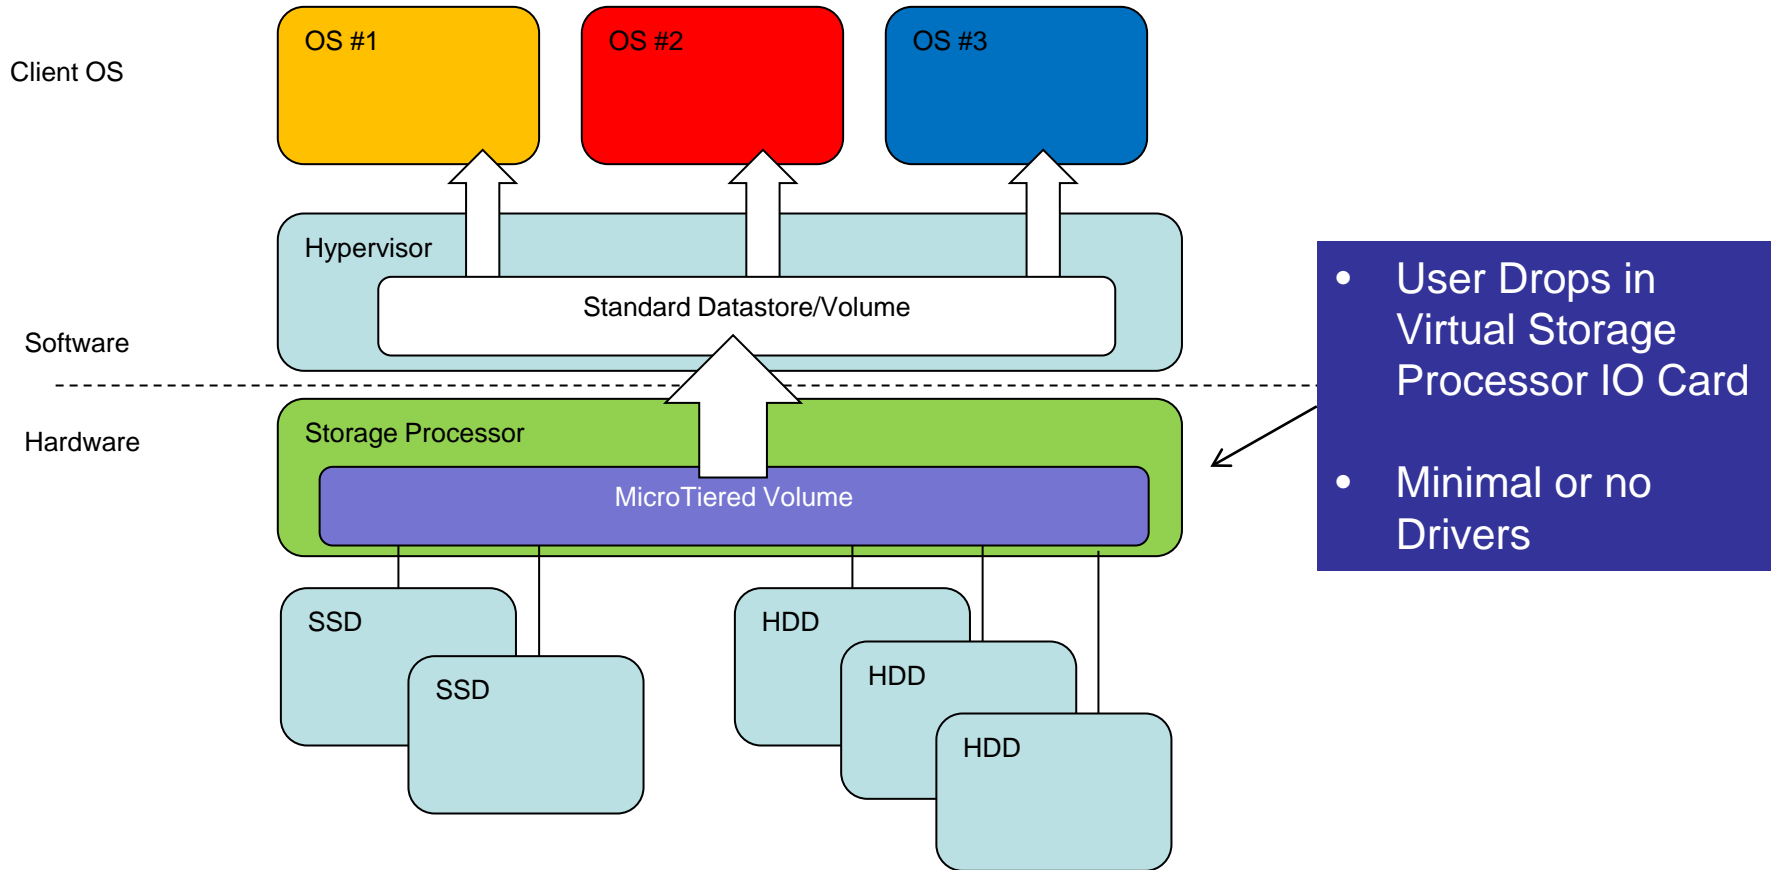
- 4-5x increase from HDD to SAS/SATA SSD
- Only 10% incremental when SSD performance is more than doubled
- 7-8x possible with application and operating system stack enhancements
- Conclusion – ease of use and transparent integration of SSDs most more important than raw IOPs for the next 2-3 years

Traditional applications remain a bottleneck for the next few years

Today's Approach – Caching, SAN



Hardware Accelerated Tiering



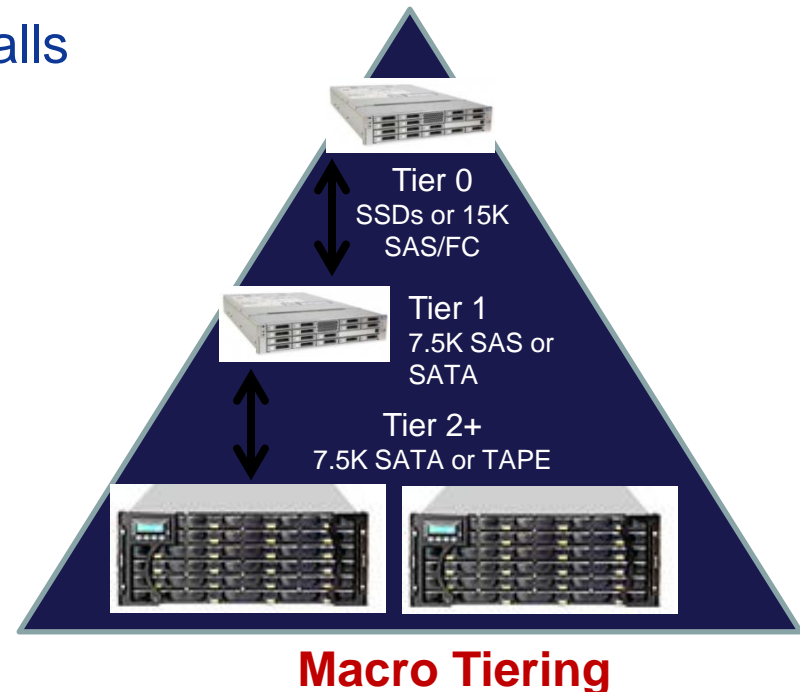
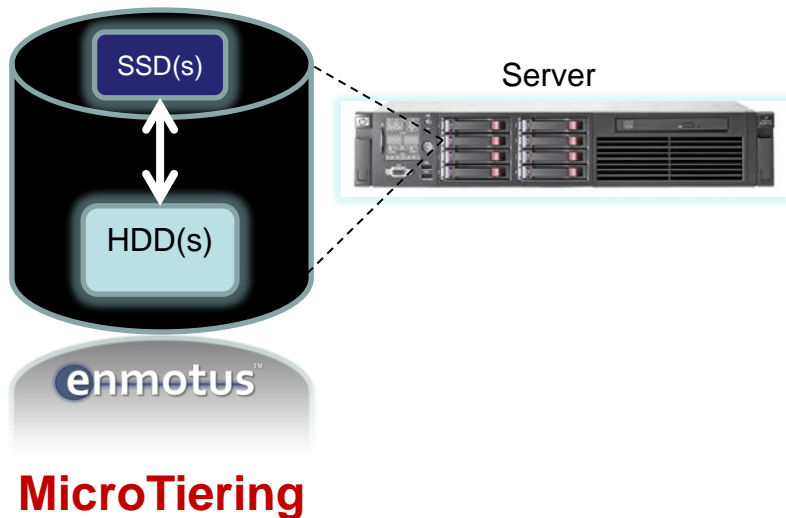
SSDs transparently tiers beneath the standard hypervisor datastore and are load balanced across all OSES

Pros and Cons

- Software based caching
 - SSD is a cache, usually accelerating read only
 - Able to accelerate existing SAN block storage systems
 - Expertise required to maintain drivers at multiple levels and hypervisor levels
 - OS and Hypervisor dependent, requires knowledge
- Hardware accelerated tiering
 - SSD is a cache or true data tier (depending on vendor) accelerating both reads and writes
 - Limited to DAS but scales to 200TB+ virtual volumes
 - Easy to deploy – set and forget
 - Bare metal deployment, OS and Hypervisor Agnostic
 - Less sensitive to OS and Hypervisor “versioning”

Enmotus MicroTiering™

- A new class of high performance SSD optimized tiering for non-SAN and hybrid SSD-HDD tiering applications
- 100% transparent automated tiering at sub-hypervisor level
- No hypervisor or client software required
- Legacy volume migration support for upgrade markets
- Bare metal deployment for new installs



Thank you!