

SSDs: Practical Ways to Accelerate Virtual Servers

Session B-101, Increasing Storage Performance

Andy Mills


CEO Enmotus

- 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

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



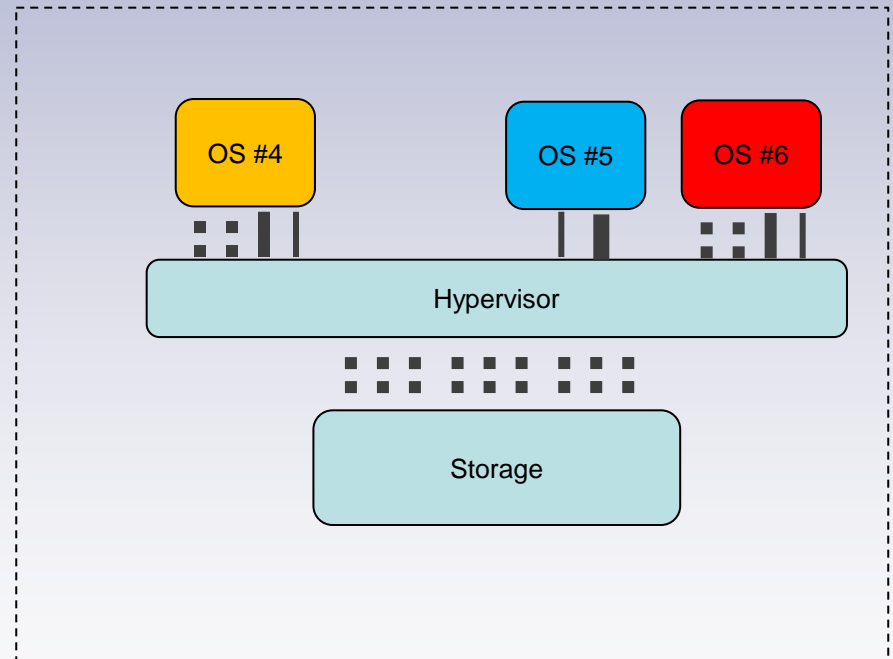
Sequential Streams are turned to Random
Ideally would like the SSD to be servicing the highest activity OS

Live Machine Migration

Physical Server A



Physical Server B

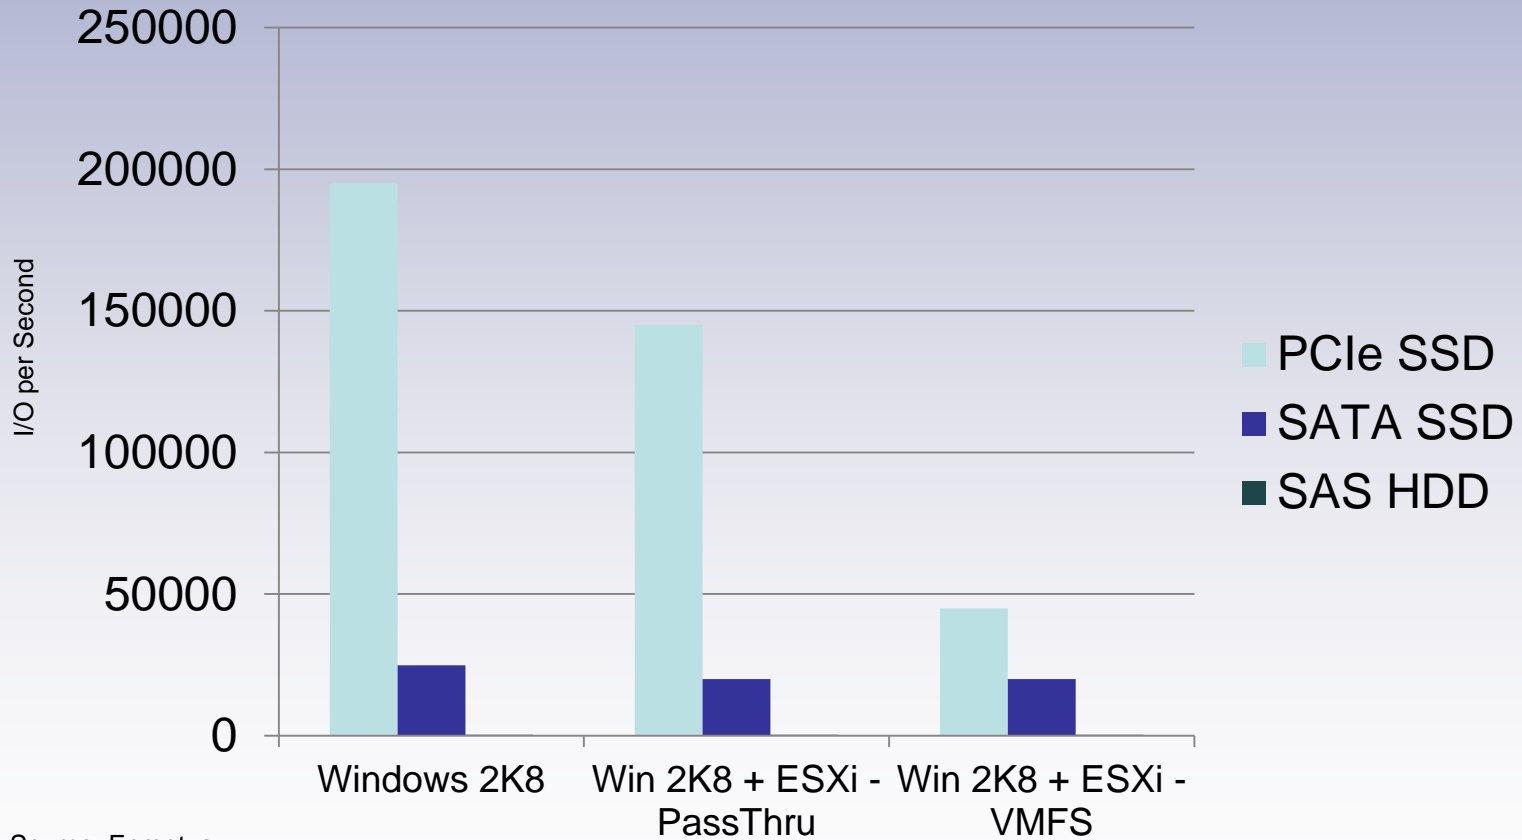


Moving or Creating a New VM Changes the Storage IO Balance

SSD Integration – Today's Options

- Software Based Caching and Tiering
 - Client OS only
 - Hypervisor + Client Driver
 - Hypervisor
 - SSD type agnostic e.g. SATA/SAS SSD or PCIe SSD
- Hardware Accelerated
 - RAID adapters with SSD caching
 - PCIe Hardware Accelerated SSDs
 - Intelligent Storage IO Processors – MicroTiering
- Hybrid Hardware-Software
 - PCIe SSDs with software based flash management

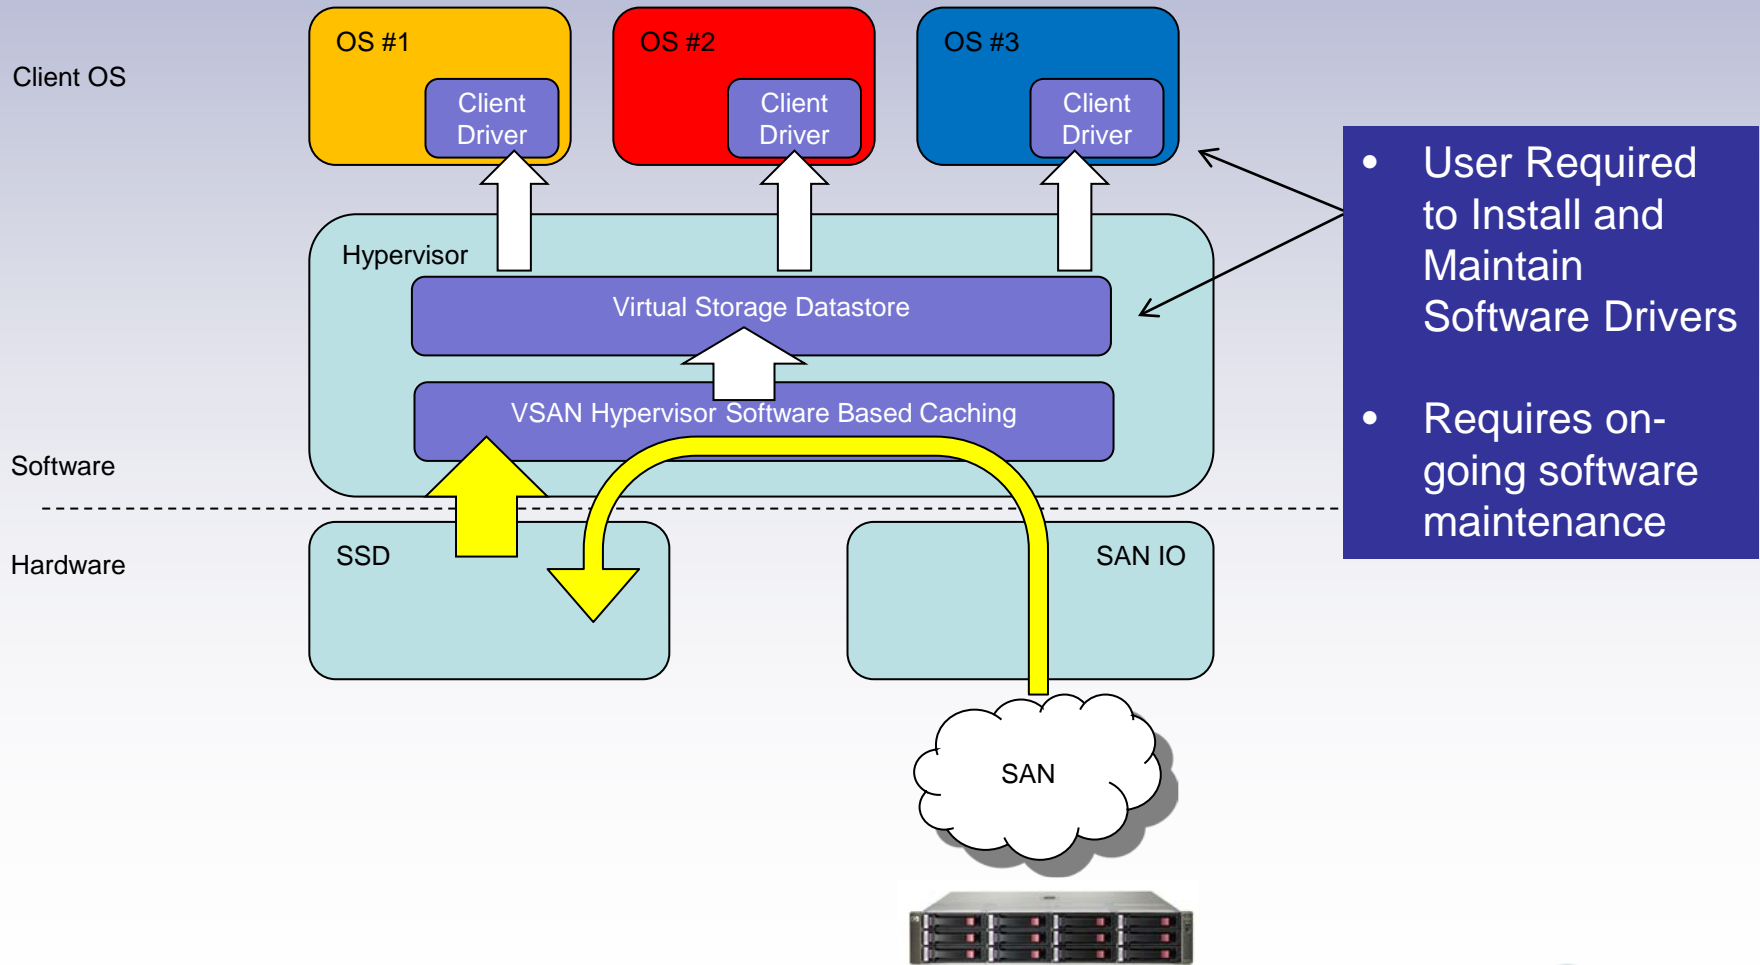
Raw SSD Performance



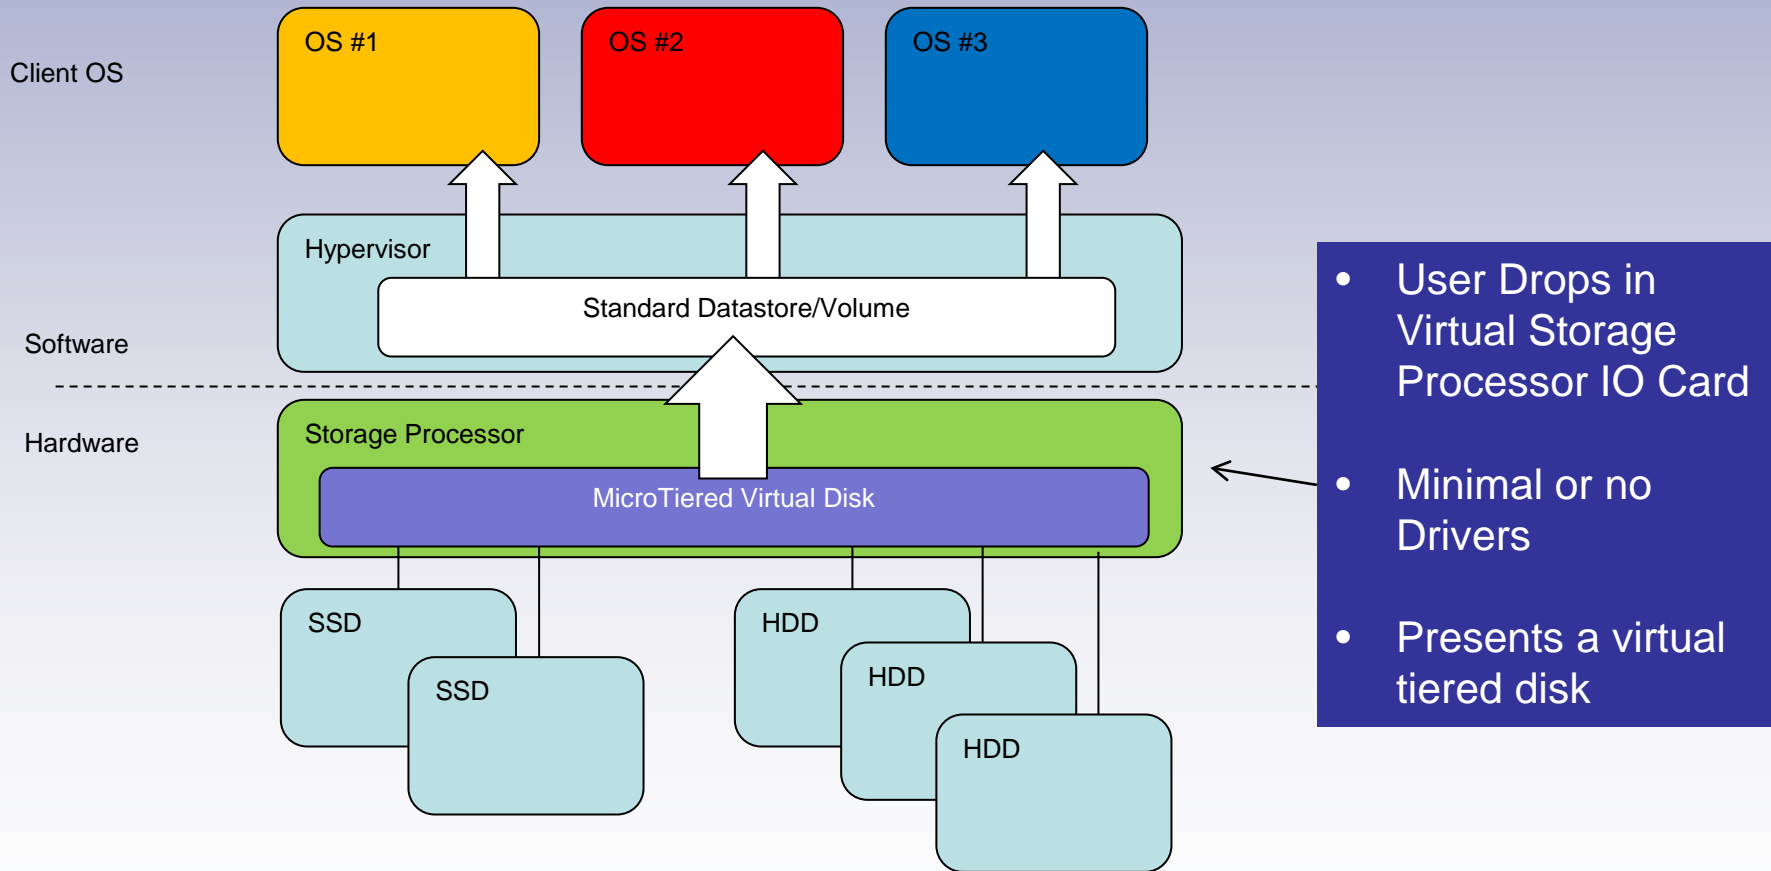
Source: Enmotus

Dell R515, Iometer 2006, 512B Random Read IOPs

Today's Approach – Caching, SAN



Hardware Accelerated Tiering



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

Software based caching

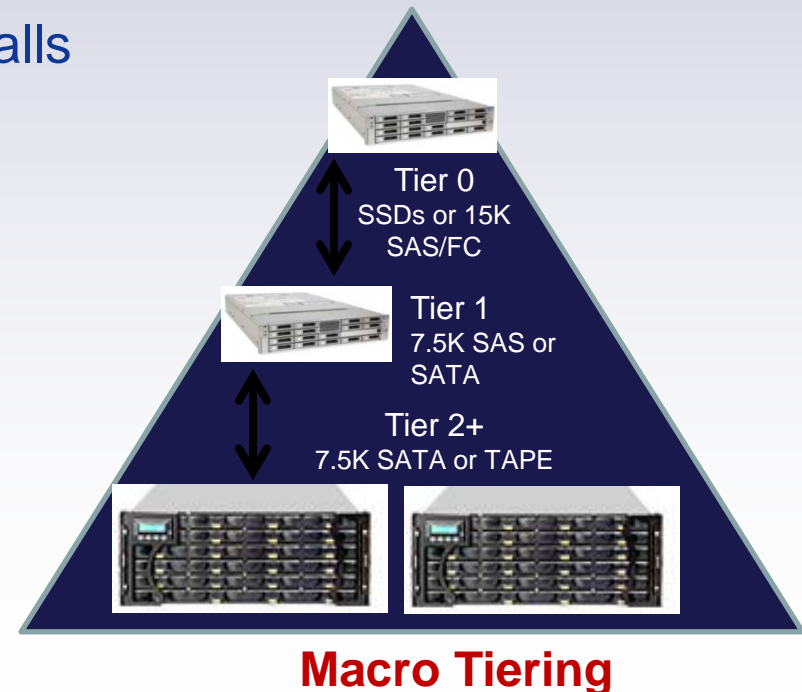

- SSD used as a cache, usually accelerating read only
- Write-back cache modes existing but very high level SSD wear
- No incremental hardware, works with existing storage
- Higher maintenance with drivers at OS and hypervisor levels
- OS and Hypervisor dependent

Hardware accelerated tiering

- Presents a single virtual tiered disk to the host
- SSD is a cache or true data tier (depending on vendor) accelerating both reads and writes
- Easier to deploy – set and forget
- Bare metal deployment, OS and Hypervisor Agnostic
- Less sensitive to OS and Hypervisor “versioning”

Enmotus MicroTiering™

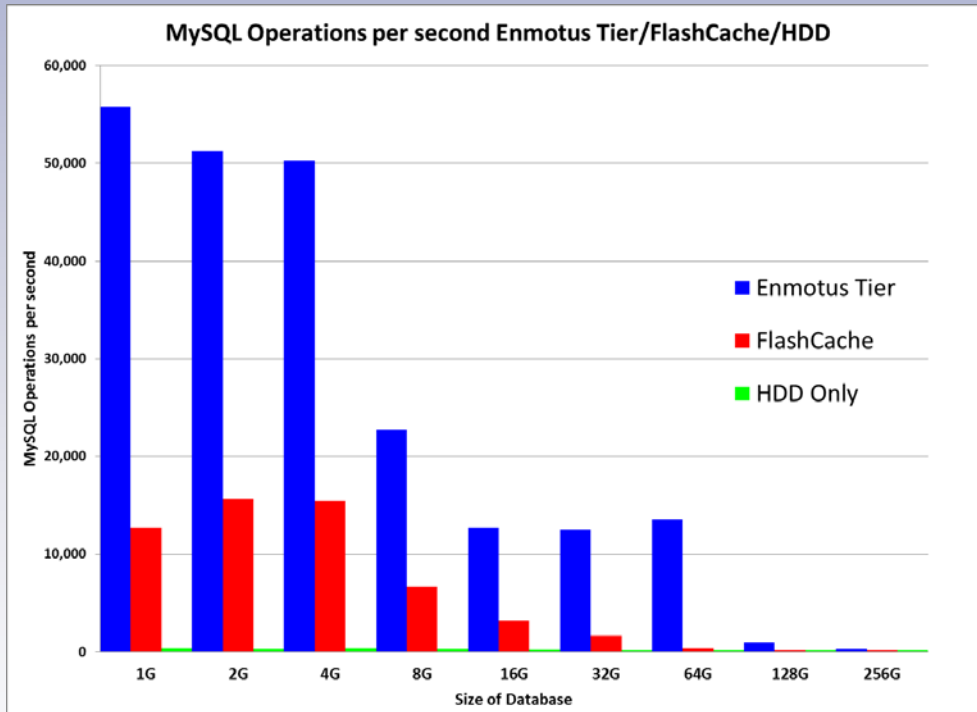
- A new class of high performance SSD optimized tiering for non-SAN and hybrid SSD-HDD tiering applications
- 100% transparent automated tiering at sub-hypervisor level
- No hypervisor or client software required
- Legacy volume migration support for upgrade markets
- Bare metal deployment for new installs



Tiering vs. Caching

	Tiering	Caching
Performance - reads	Near native SSD speeds for data in fast tier consistently	Near native SSD speeds for cached data, lower when cache is near full or thrashing
Performance - writes	Near native SSD speeds	Write back cache mode only
Overhead Wear Leveling	Near zero e.g. <1 % incremental writes measured	Generates significant number of overhead writes to SSD consistently
Migration to Fast Tier	Policy driven, moved chunks at a time in the background after the IO event	Moved into SSD as part of the read IO event, and all write events in write back cache mode
Length of Time in SSD Tier	Stays until displaced by more important data, persistent through power cycle	Removed once cache fills or at every scheduled inactivity flush event

Tiering vs, Caching



Enmotus Tiering Software
MSI Z77A-GD55, Intel i7 3770 3.5Ghz
8GBytes DDR3
2TB HDD Toshiba, 120GB SSD SanDisk Extreme
SysBench 0.4.12
128 files & Block size 16Kb
Read/Write ratio = 1.50

Tiering data provides several benefits over SSD caching

- Faster primarily due to write performance
- Significantly lower SSD wear out, often negligible

Virtual servers benefit more from tiering than caching due to faster writes and more intelligent data movement

Thank you!