



Challenges of Implementing Automated Tiering in Enterprise Servers

Special Session on New Approaches to Enterprise Server
Solutions

Andy Mills
CEO Enmotus Inc


- Review of challenges of adopting SSD-hard drive tiering in general purpose virtual servers
- Direct attach and primary storage I/O focus
- Topics covered
 - Macro/Micro problems that tiering solves
 - Challenges of integrating into virtual server environments
 - Introduction to MicroTiering™ - a high performance DAS approach to SSD-HDD tiering

Macro Problem

- Over the next decade:
 - 10x more servers (virtual and physical)
 - 50x more information managed by enterprise datacenters
 - 75x more files in the data center
- <1.5x more IT professionals in the same timeframe
- Full and complete automation of data management is becoming crucial to balance the work per IT professional load
- Automated data movement primarily available to SAN users
 - No easy to use performance solution for server centric storage users
 - Next generation Hadoop-like architectures not well served

Data Tiering

- Part of the broad suite of SAN virtualization trends
- Data Tiering automatically moves **frequently accessed** data to the faster storage layer and **least accessed** to the slower, lowest cost layer
- Why the need to tier storage?
 - Too costly to put 100% of the data on the faster, most expensive storage
 - Exaggerated by virtualized and cloud systems by files, changing too frequently

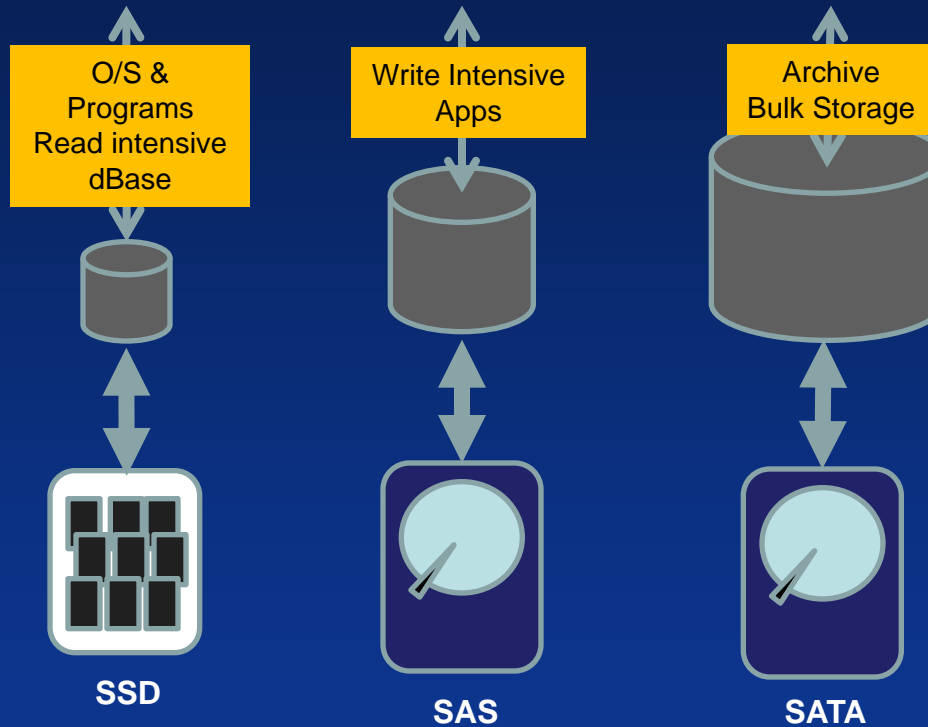


Server Key Trends

- More than 50% of all servers are now virtualized*
 - Represents around 18-20% of new annual physical server shipments*
- CPU utilization targets are changing (again....)
 - Pre virtualization – target 30-40% max
 - Post virtualization – 90%+ utilization
- VM has a large impact on traditional server storage I/O
 - Increasing trend toward hardware accelerated I/O
 - Cannot assume spare CPU cycles are available for software RAID and other caching functions

DAS Islands Inside Servers

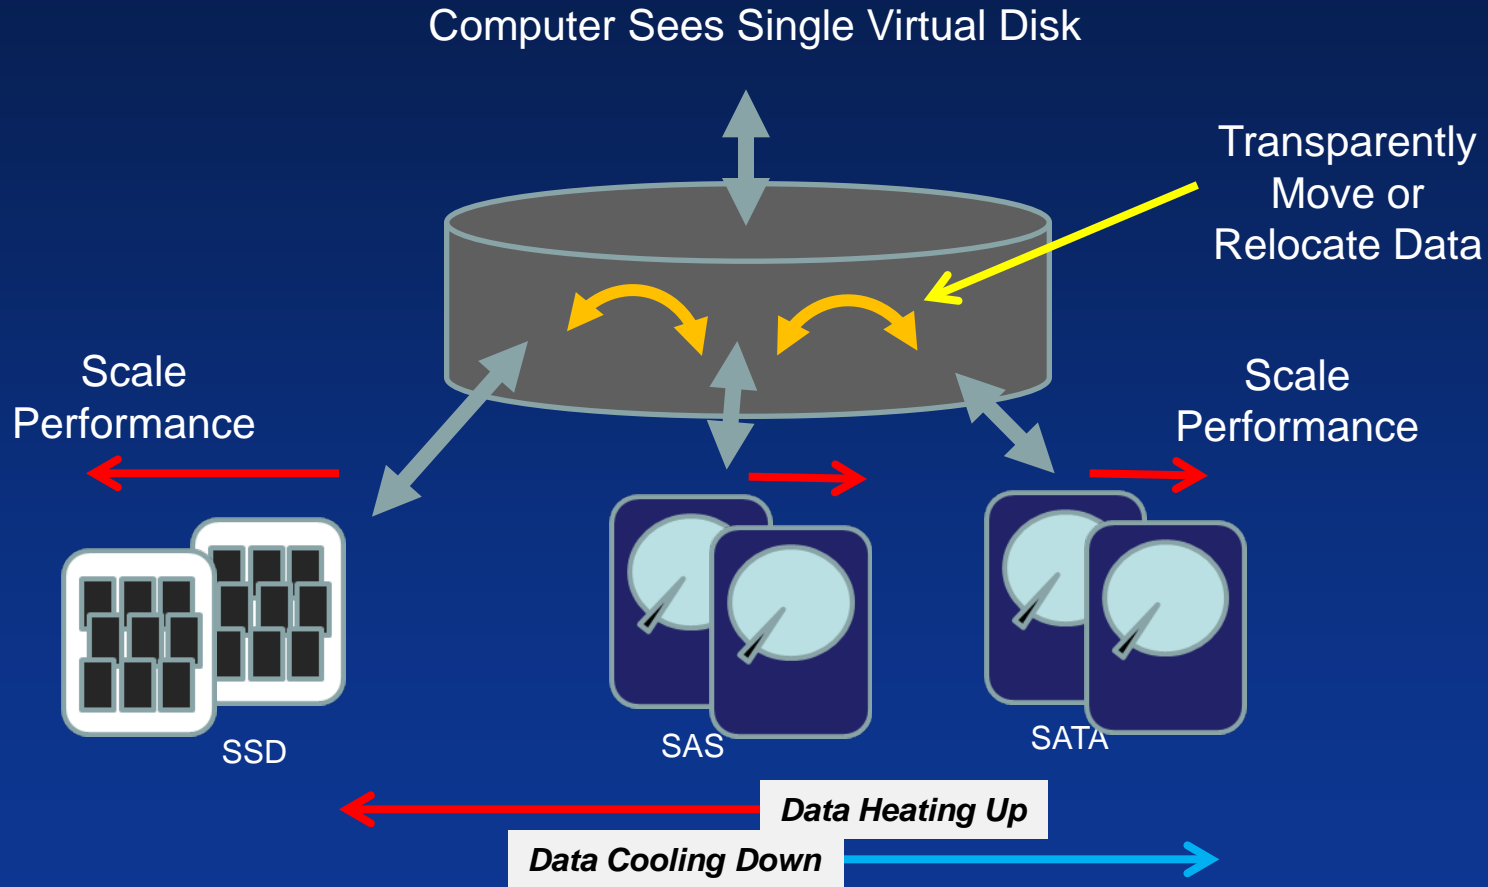
Computer Sees Individual Disks



- *No Single Media Fits All Applications*
- *Data becomes "hot and cold" over time*
- *Manually optimize data location*

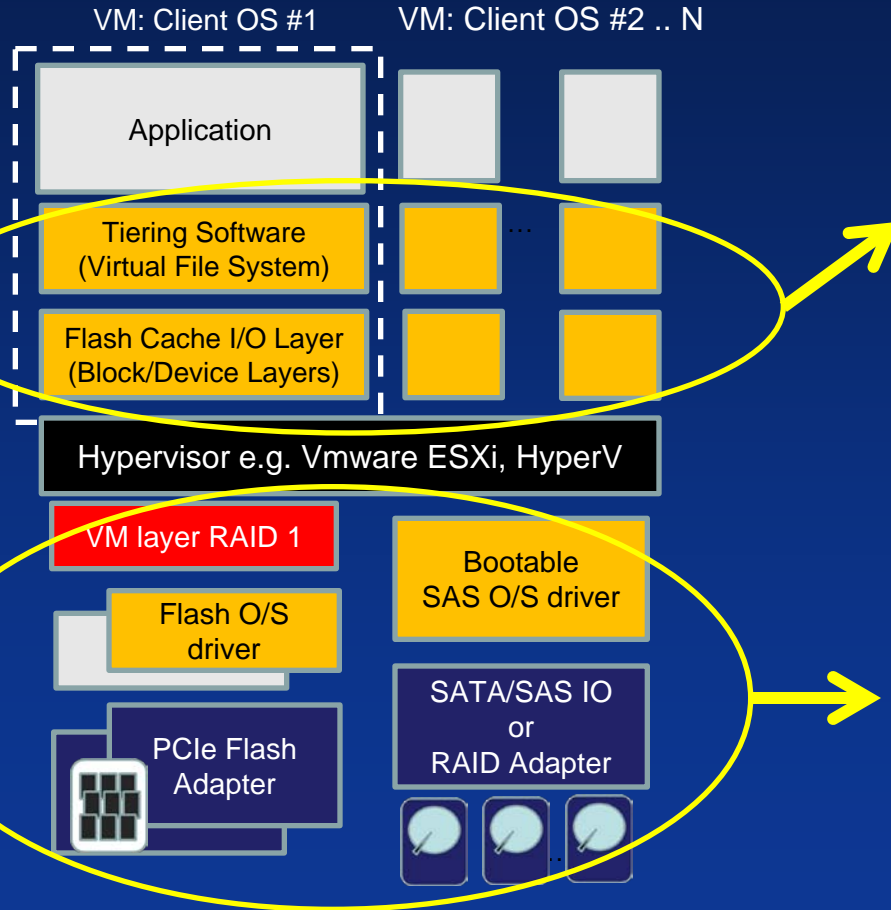
<i>Capacity:</i>	200G	600G	3TB
<i>Cost:</i>	\$11-40/GB	\$0.90/GB	\$0.08/GB
<i>Read IOPS:</i>	25,000+	400	130
<i>Write IOPS:</i>	2,500	360	130

Storage Virtualization



Auto-tiering transparently matches data-blocks to the appropriate media based on frequency of access and access patterns

Hypervisors and Storage I/O



 Not generally available today

Storage I/O above the Hypervisor

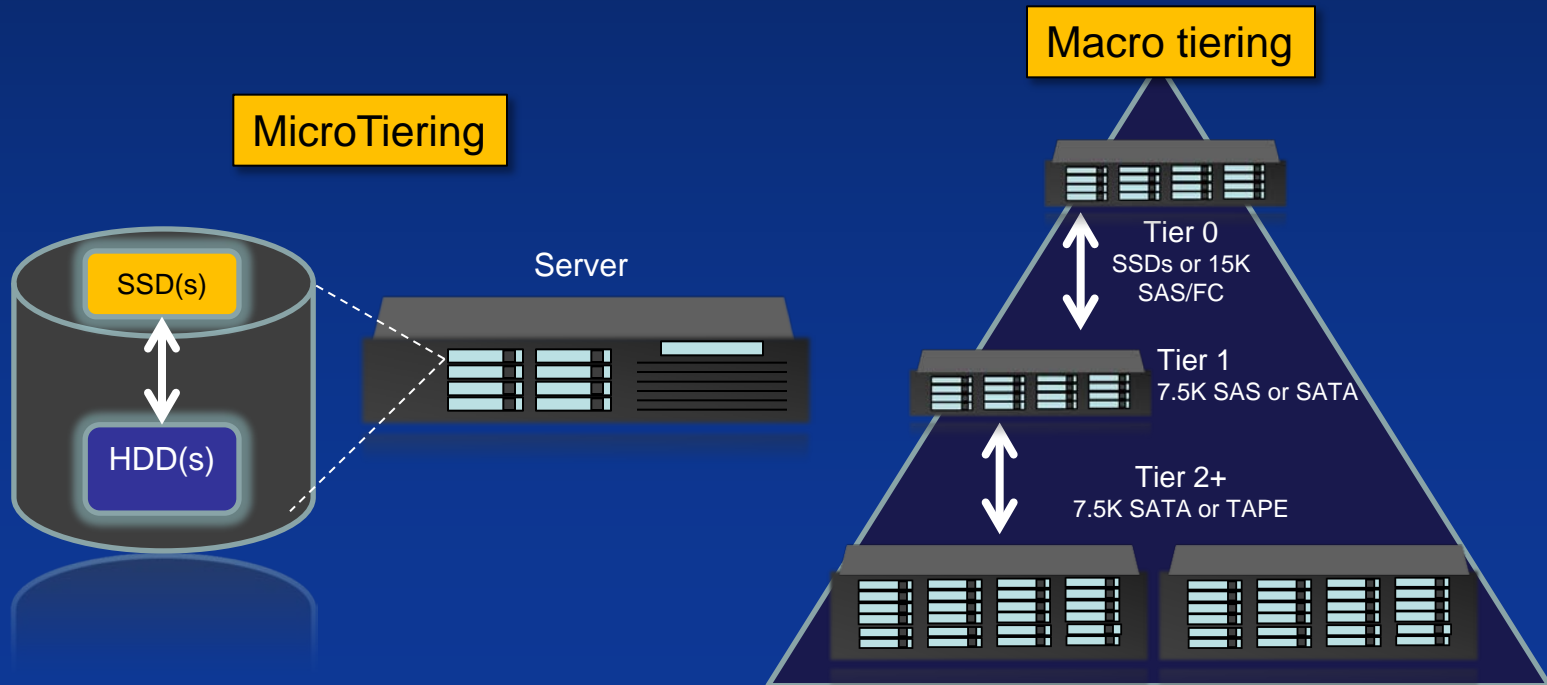
- Everything done on a per VM basis
- ‘N’ copies of File or virtual file system plus tiering and/or caching
- No ultra performance access to raw disk devices
- Severely limits performance of any “block” based utilities in VM

Storage I/O below the Hypervisor


- Better access to raw blocks
- Requires several block layer drivers to work, including a VM-RAID solution
- Ability to be VM agnostic and provide shared functionality across all VMs

MicroTiering™

- High performance automated data tiering at the server level
- Transparently moves data to the optimum storage device inside the server
- 100% hardware accelerated and bootable
- Integrated SSD and SAS/SATA storage I/O replacing two adapters in one

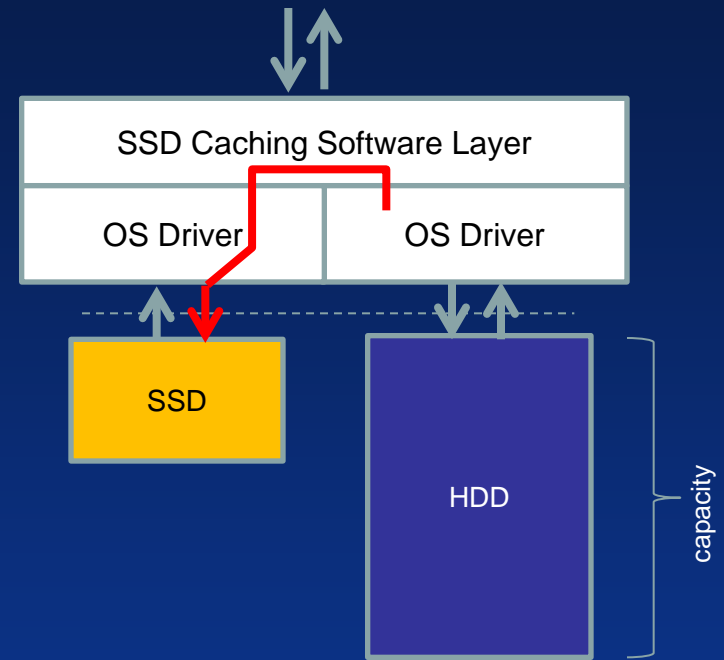



MicroTiering™ Operation



- Faster “hot” data relocated to SSD i.e. not cached
- Less frequently accessed data moved to HDD
- Application sees nearly all of the combined capacity
- Supports multiple striped or redundant SSD or HDD sets

MicroTiering™ vs. Caching

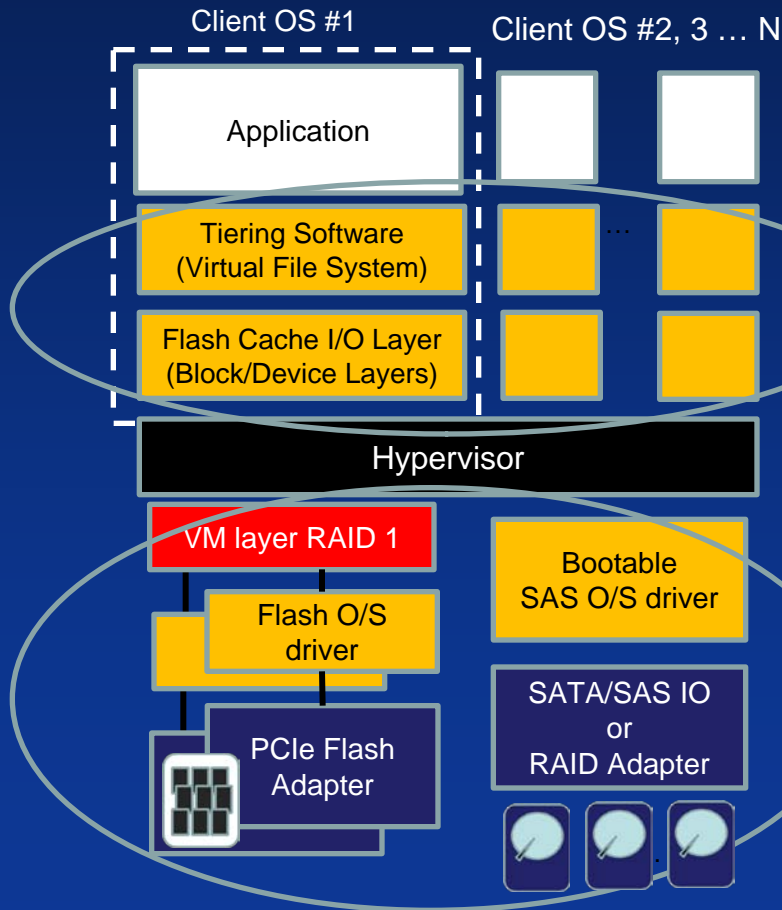


- Accelerates read and write I/O
- All capacity is visible
- Virtual drive avoids multiple software/driver layers in the host
- Built-in support for VMs

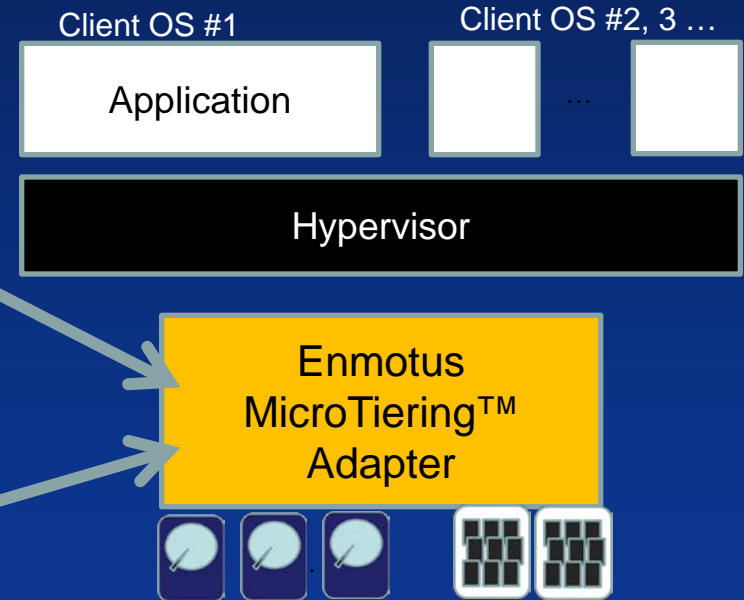
- Typically accelerates read I/O only
- SSD capacity is invisible
- Requires system level software to operate above base driver level
- Hard to support in VMs eg. VMware ESXi

MicroTiering™ Adapters

Conventional Approach




MicroTiering™ Approach




MicroTiering™ Applications

Traditional Enterprise Shared Storage



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



- Classic shared SAN network
- Data protection relies on SAN devices and multi-path network connections

- Emerging distributed grid storage leveraging approaches developed for web applications
- Data protection based on replication across multiple nodes (node= server + DAS storage)

- Challenges exist in virtualized server SSD adoption
- Limited options inside servers for high performance primary storage tiering for boot and data volumes
- New method required to support tiering without adding significant software layers
- MicroTiering™ effectively solves the problem