

Increasing Storage Performance

High Performance MicroTiering for Server DAS Storage

Andy Mills

President/CEO, Enmotus

andy.mills@enmotus.com

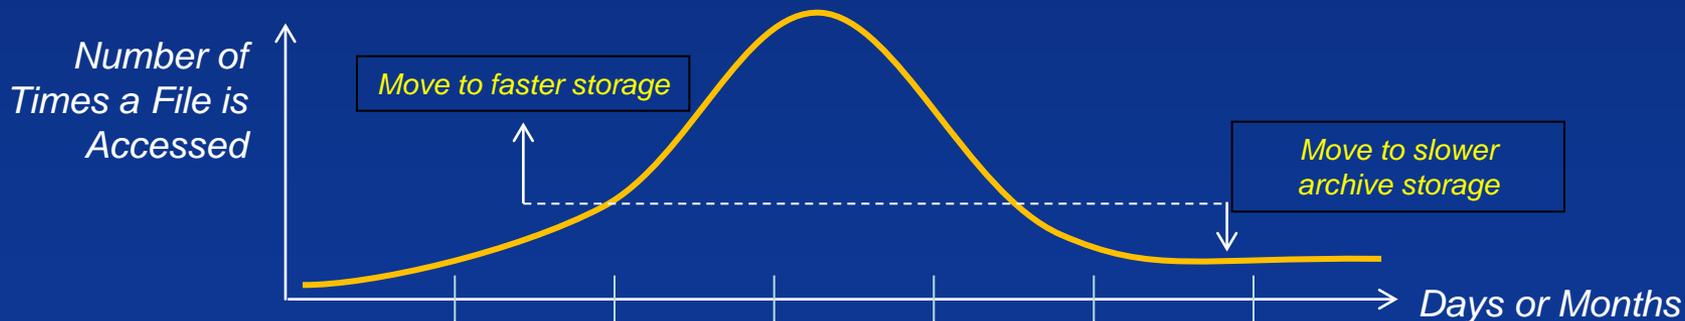
- Review of challenges of adopting SSD-hard drive tiering in general purpose virtual servers
- Direct attach and primary storage I/O focus
- Topics covered
 - Macro/Micro problems that tiering solves
 - Challenges of integrating into virtual server environments
 - Introduction to MicroTiering™ - a high 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
- Movement of SSD tier inside the server
 - No easy-to-use performance solution for server centric storage users
 - Next generation Hadoop-like architectures not well served

Data Tiering

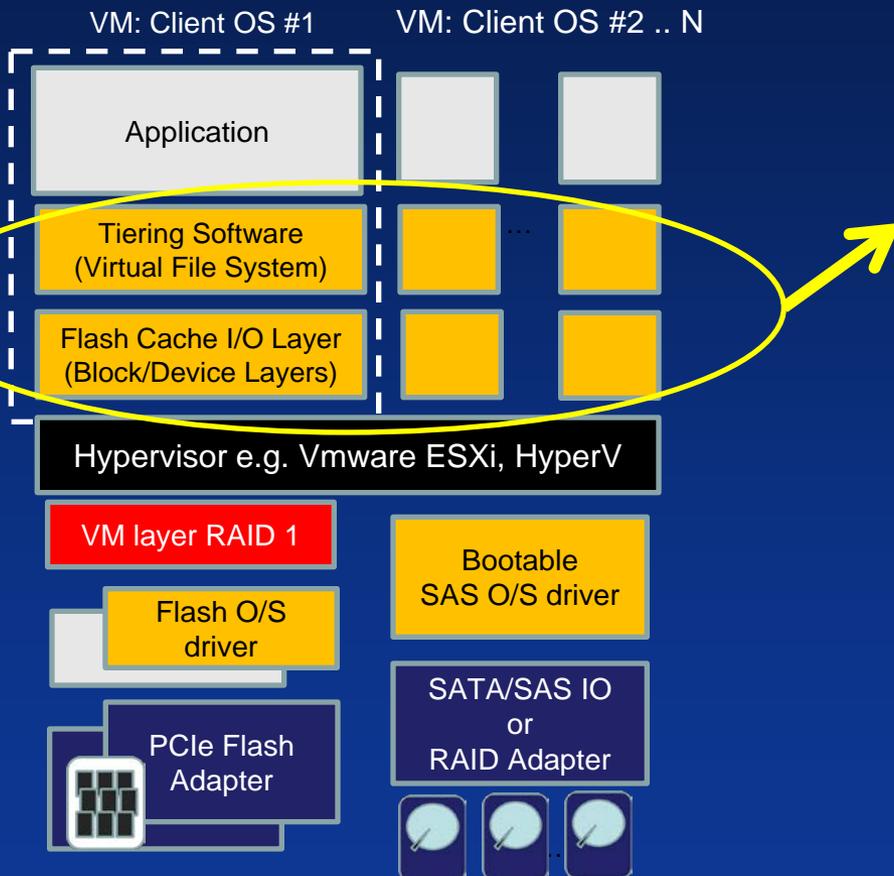
- Data Tiering common place is networked storage appliances and/or specialized server file systems
- Data Tiering automatically moves **frequently accessed** data to the faster storage layer and **least accessed** to the slower, lowest cost layer
- Why the need to tier storage?
 - Too costly to put 100% 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

Hypervisors and Storage I/O Today



Storage I/O above the Hypervisor

- Everything done on a per VM basis
- Each VM client has to be responsible for it's own data tiering
- No ultra performance access to raw disk devices
- Severely limits performance of any “block” based utilities in VM

Storage I/O below the Hypervisor

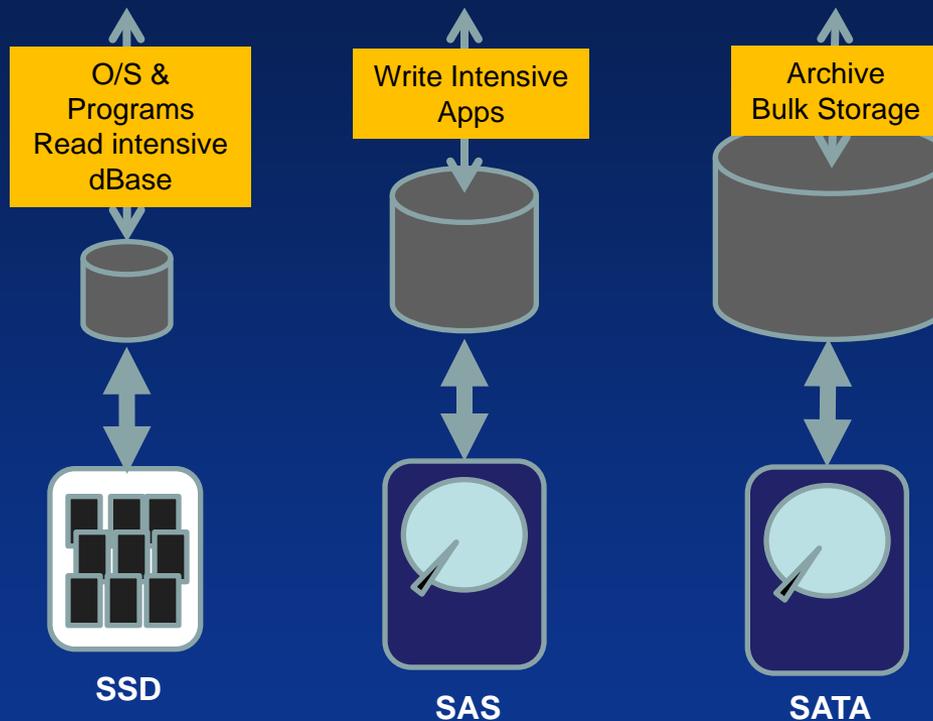
- Better access to raw blocks
- Requires several block layer drivers to work, including a VM-RAID solution
- Ability to be VM agnostic and provide shared functionality across all VMs



Not generally available today

DAS Islands Inside Servers

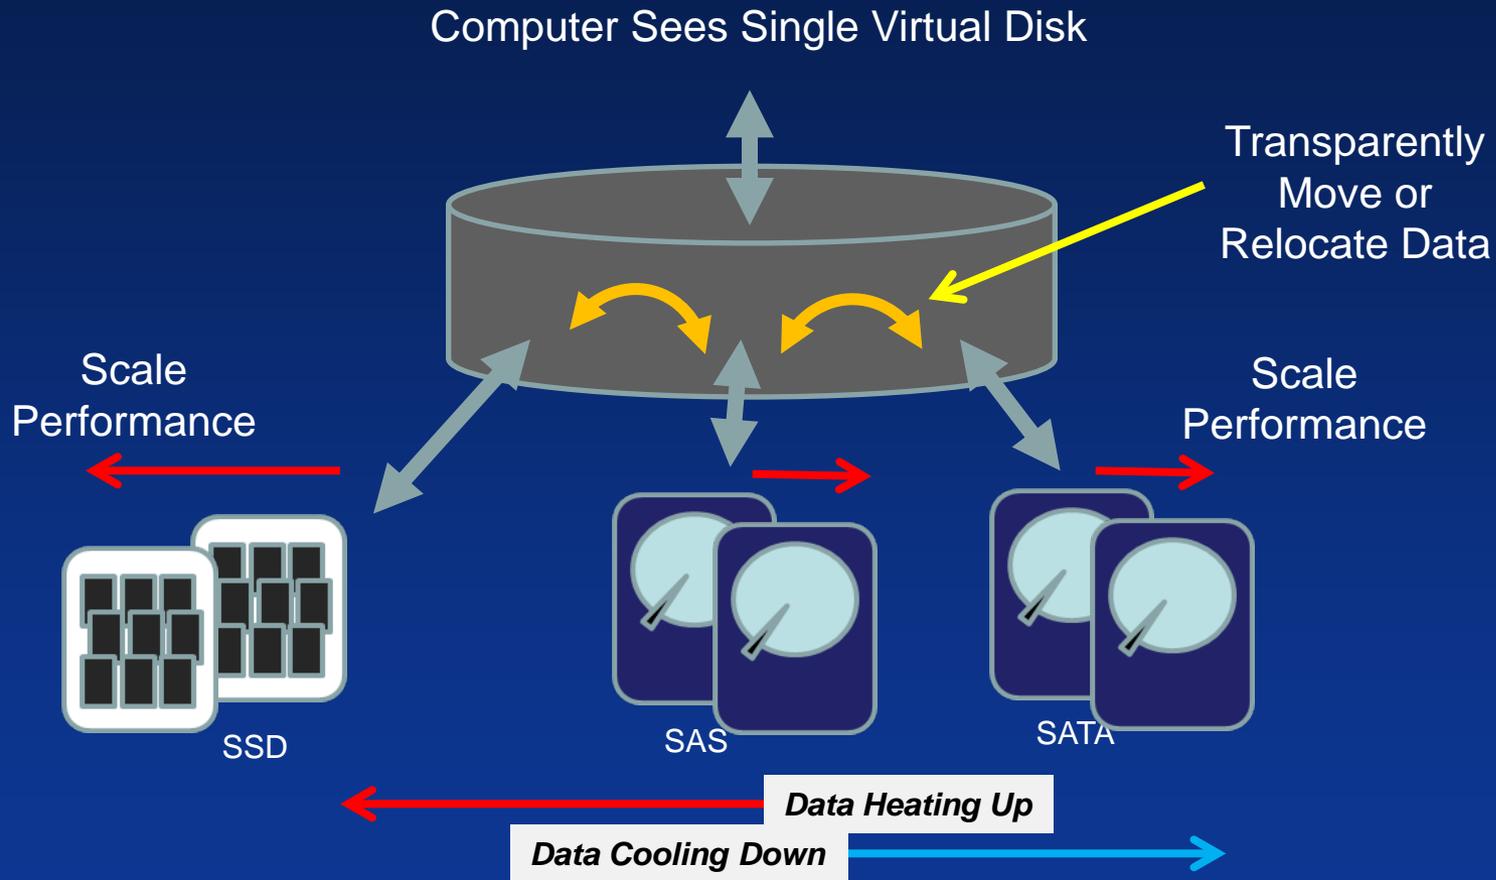
Computer Sees Individual Disks



Capacity:	200G	600G	3TB
Cost:	\$11-40/GB	\$0.90/GB	\$0.08/GB
IOPS:	25,000+	280	80-100

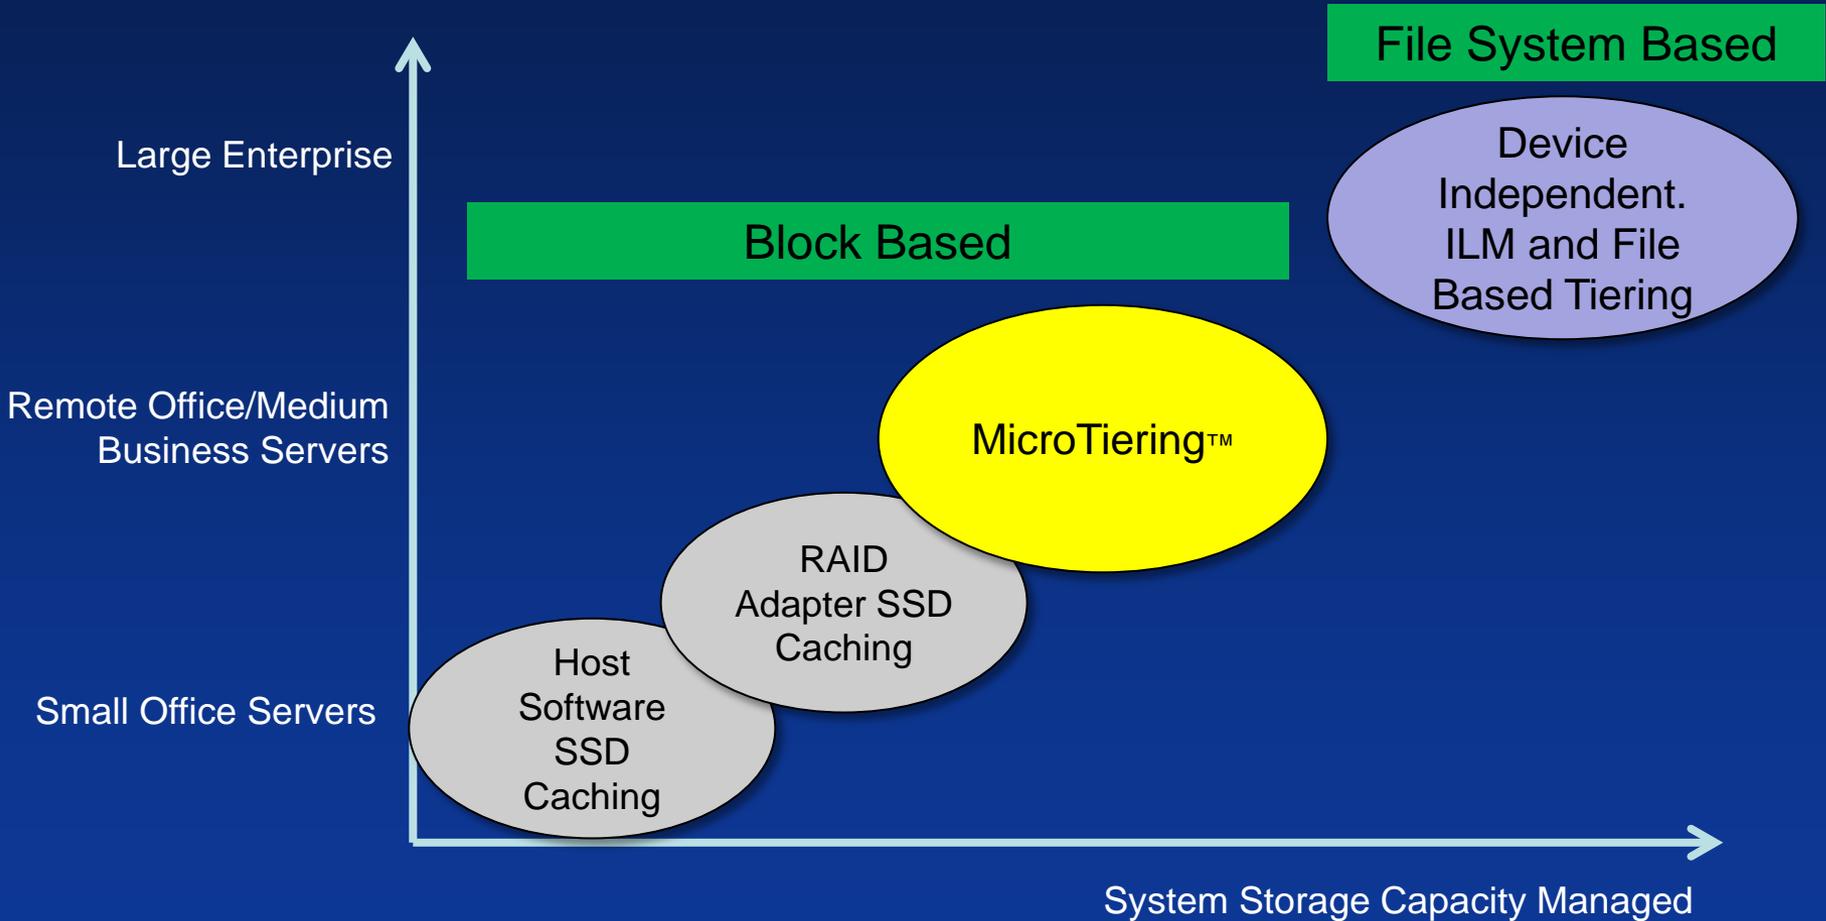
- *No Single Media Fits All Applications*
- *Data becomes "hot and cold" over time*
- *Manually optimize data location*

DAS Virtualization



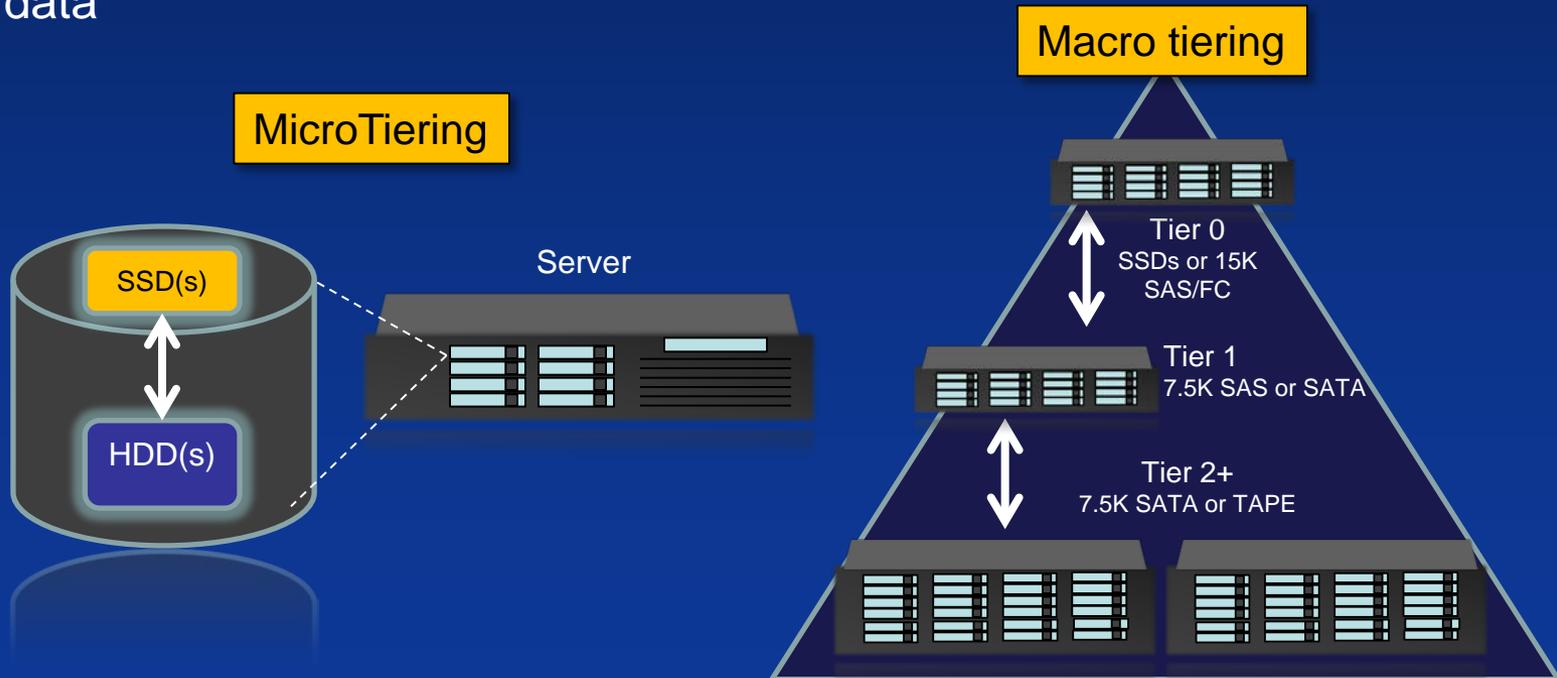
Auto-tiering transparently matches data-blocks to the appropriate media based on frequency of access and access patterns

Server Caching and Tiering

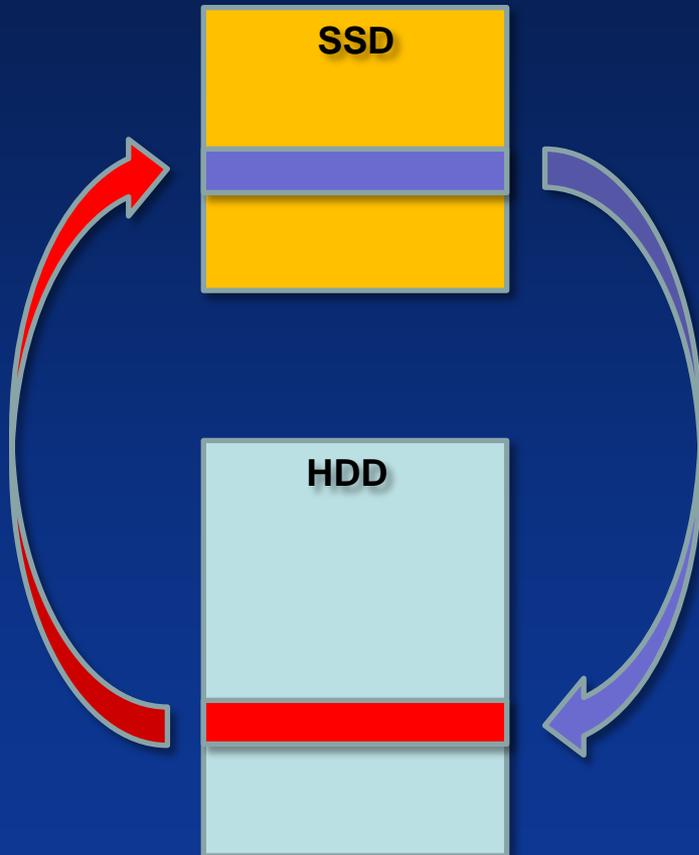


MicroTiering™

- High performance automated data tiering at the primary DAS 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
- Provides full read and write performance of SSDs to host volume for hot data

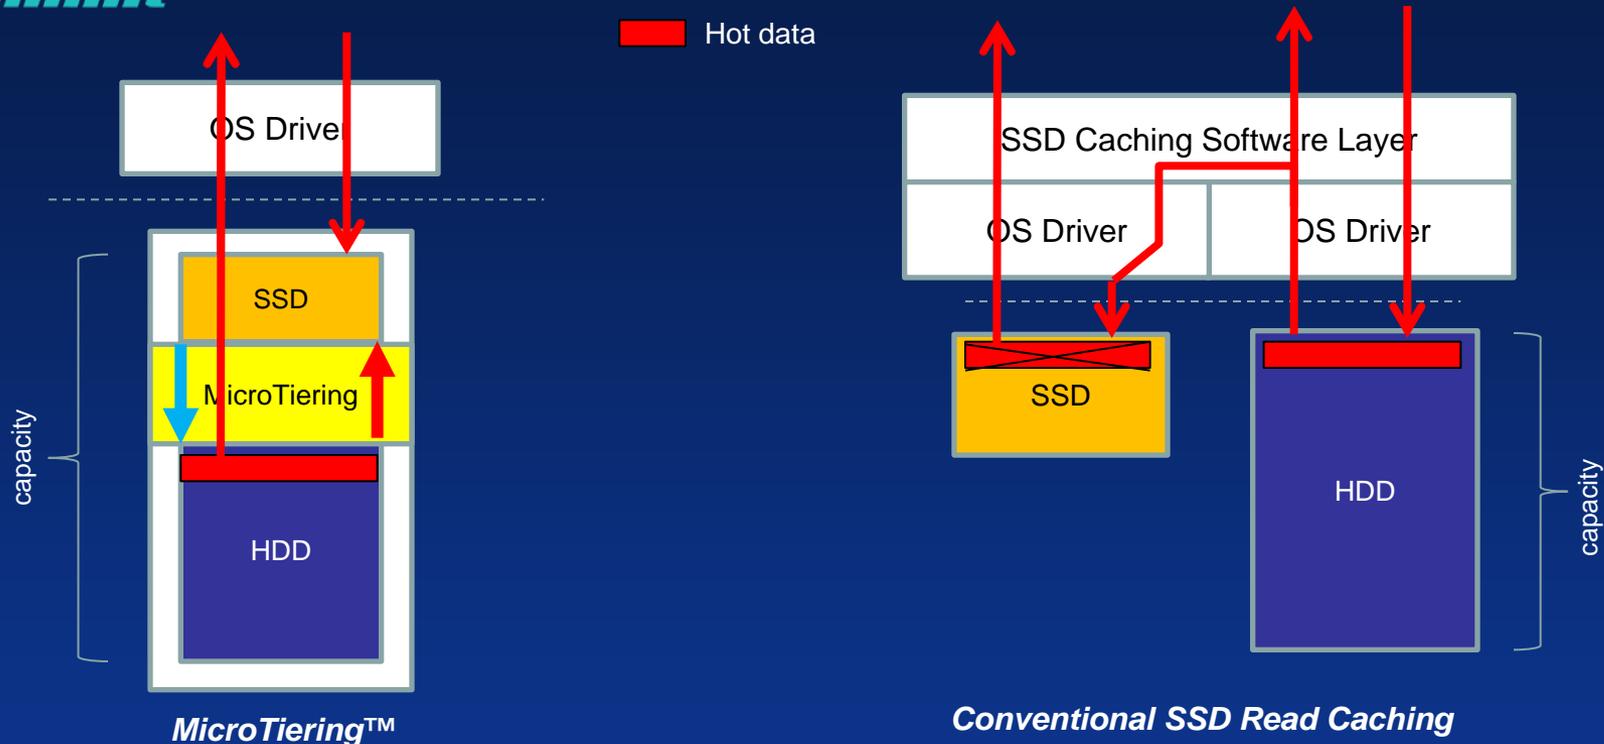


MicroTiering Basic Operation



- “Hot” data on HDD tier is moved to SSD tier
- “Cold” data on SSD tier is moved to HDD tier
- No large incremental space or pool needs to be reserved on the disk(s)
- Non-traditional algorithms employed - full Read and Write SSD speed available to host
- Real time scanning of statistics used for other move based decisions

MicroTiering™ vs. Caching

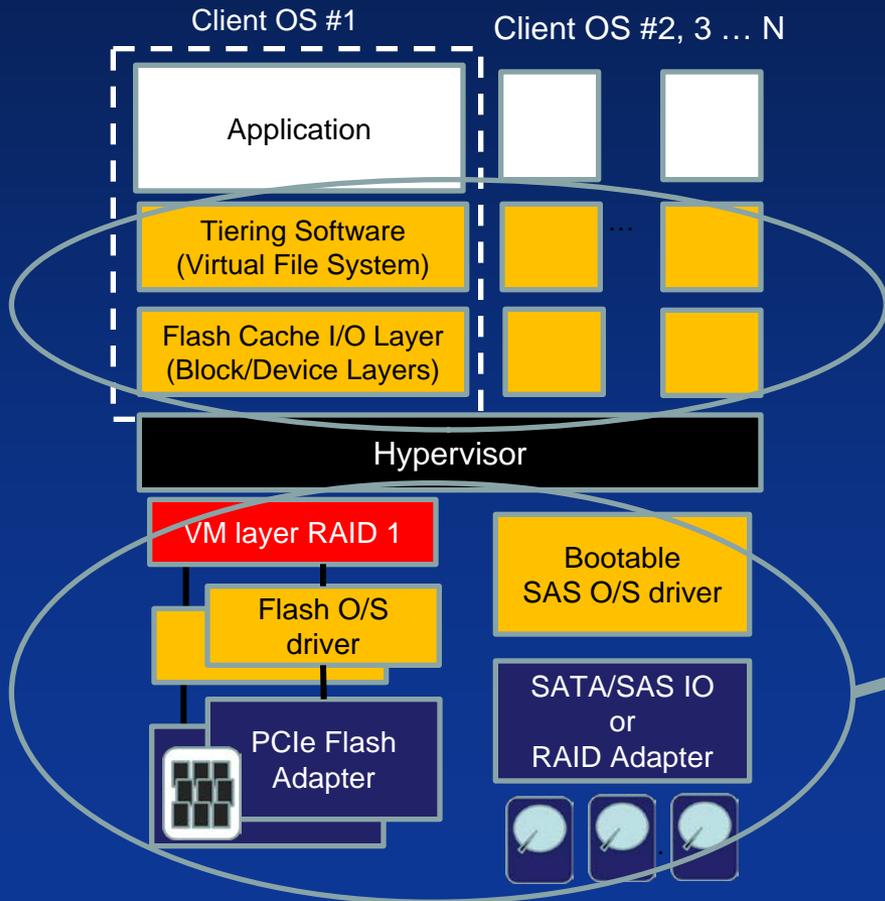


- Moves data to SSD
- Accelerates read and write I/O
- All capacity is visible
- Virtual drive avoids multiple software/driver layers in the host
- Easy integration with hypervisors

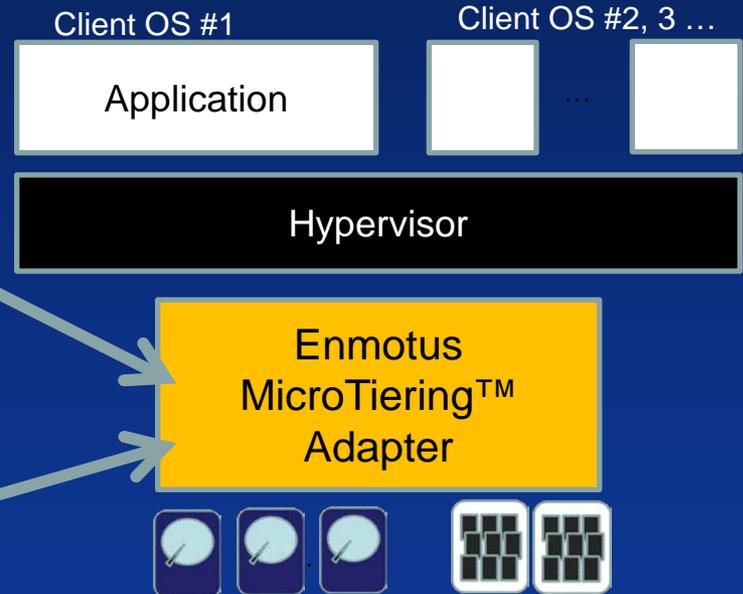
- Duplicates data in SSD
- Typically accelerates read I/O only
- SSD capacity is invisible
- Requires system level software to operate above base driver level
- Custom approaches for hypervisors

MicroTiering™ Adapters

Conventional Approach

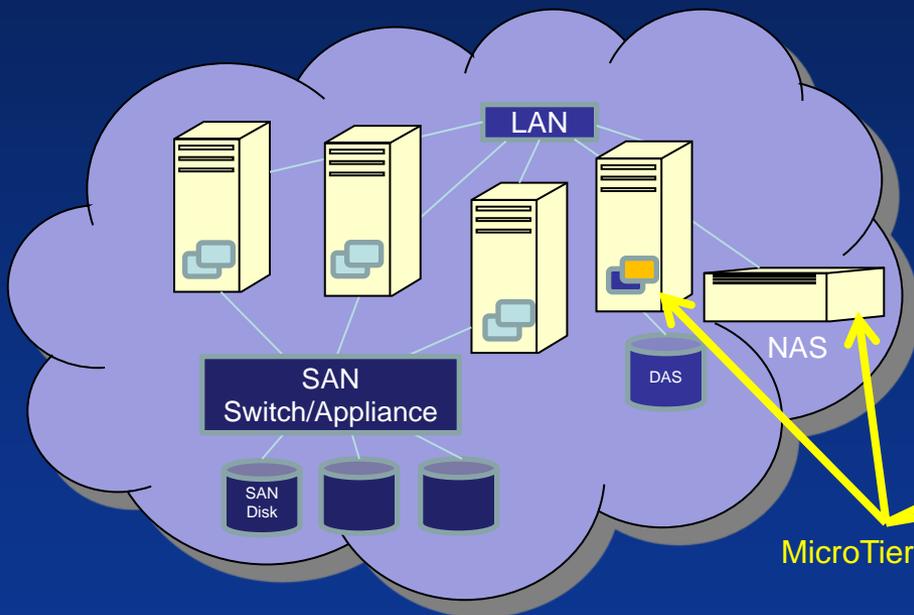


MicroTiering™ Approach

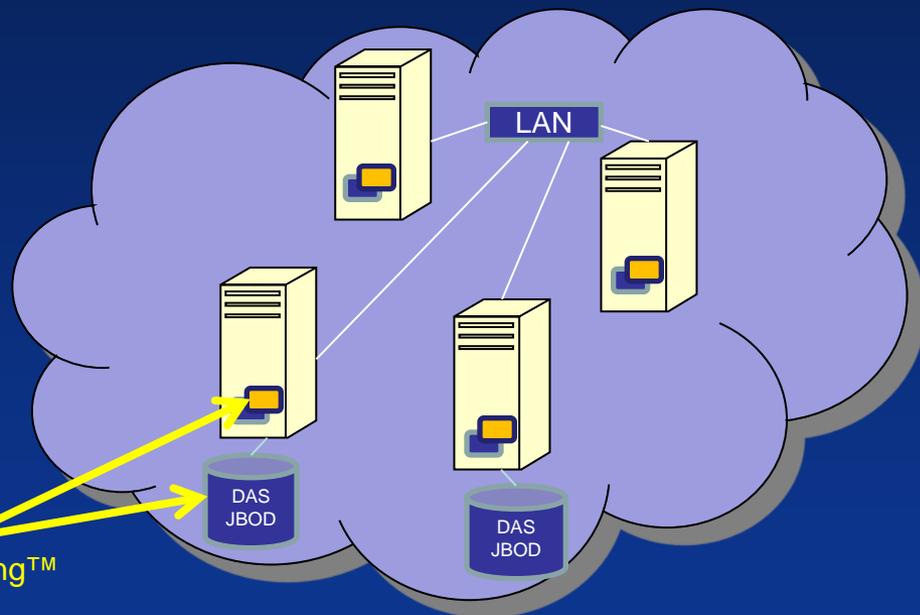


MicroTiering™ Applications

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
- MicroTiering a component of the server, NAS or open storage-server

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

- Challenges exist with SSD caching or tiering in virtualized servers
 - Has led to highly custom caching solutions for hypervisors
 - Hypervisors don't see full benefit of SSD writes
 - Forces separate boot drive for smaller virtual server environments
- MicroTiering™ solves several issues
 - No client software required
 - No hypervisor caching layers required
 - Fully auto-balancing amongst existing and newly migrated VMs
 - Full read and write performance exposed to host/hypervisors for hot data