



**Driving Innovation
Through the Information
Infrastructure**

SPRING 2011



Managing VM sprawl in the Cloud

Vish Mulchand

Director, Solutions Marketing
Hewlett Packard StorageWorks

Two Major Tectonic Shifts in IT are Underway


Physically Dedicated

Secure, Multi-tenant



Internal Data Center

**Distributed
Computing
To
Utility Computing**



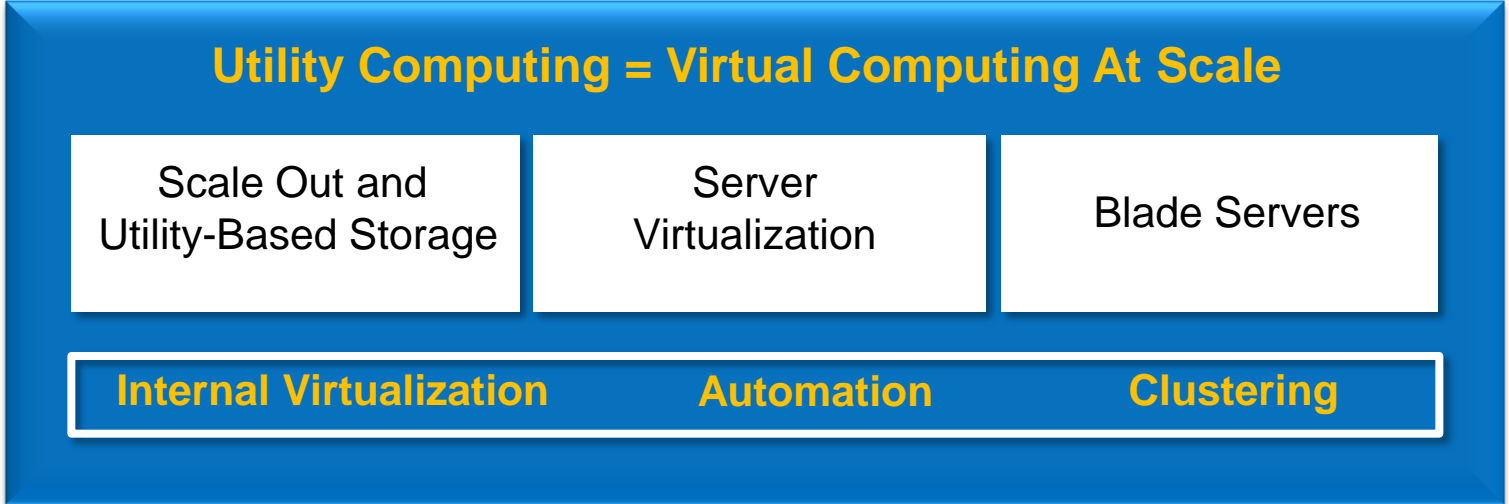
Cloud



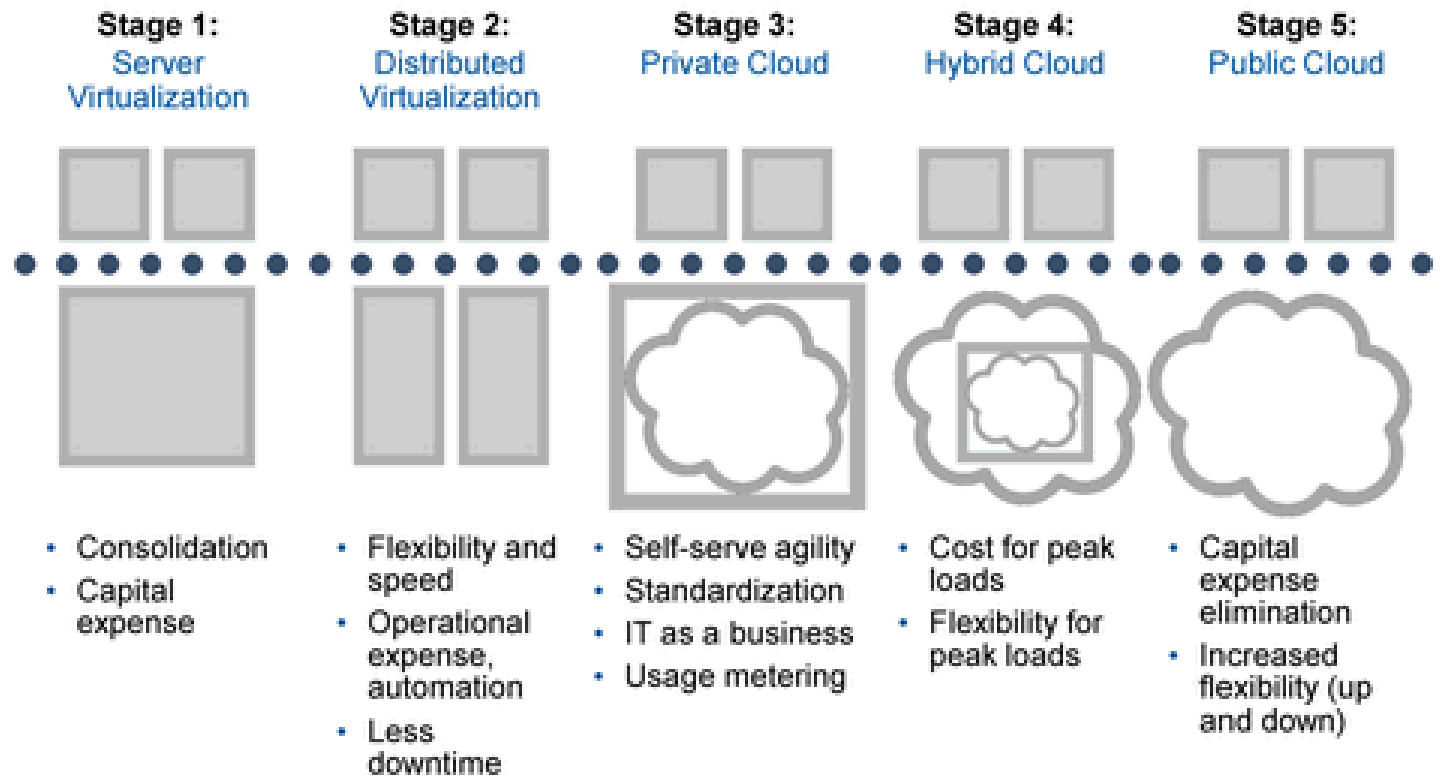
**Deliver
Enterprise IT
As
A Utility Service**



Delivery of IT as a Utility Service



Server Virtualization is the first step to a cloud environment



Source: Gartner (March 2011)

Source: Gartner, The Road Map From Virtualization to Cloud Computing, Thomas J. Bittman March 2011

What is VM Sprawl?

- Large number of virtual machines consuming resources but not being used
- Typically occurs when orgs advance to distributed virtualization ¹
 - Processes need to fundamentally change to handle the speed that virtualization enables
 - Enterprises that do not virtualization-enable their management processes can find themselves dealing with VM sprawl.
 - Rapid increase in storage requirements, leading to unhealthy virtualized environments.

1. Source: Gartner, The Road Map From Virtualization to Cloud Computing, Thomas J. Bittman March 2011

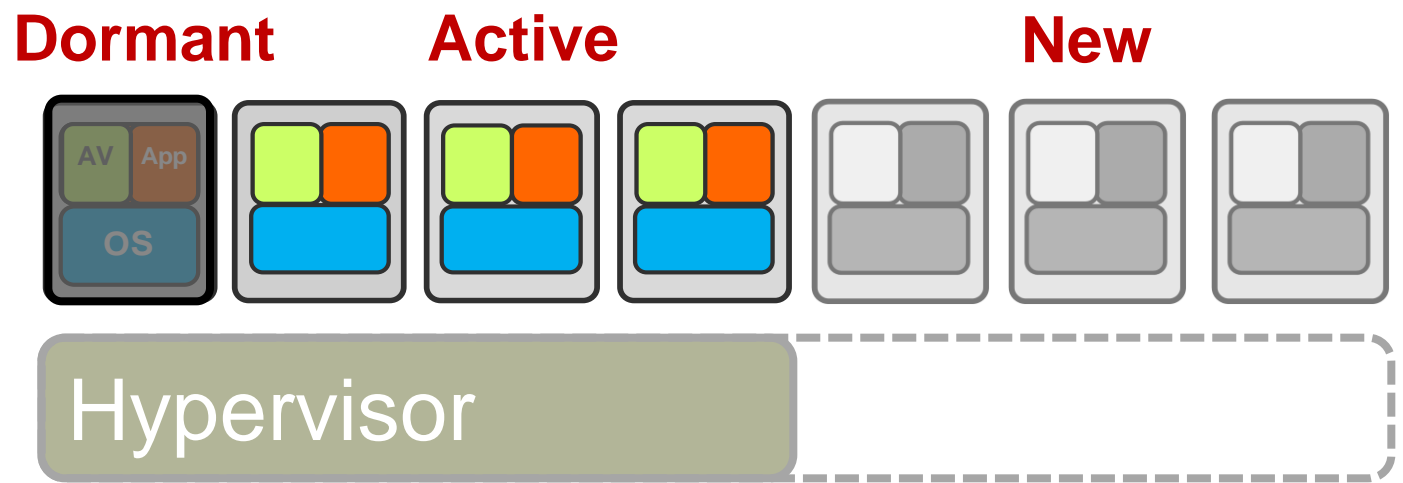


Is VM sprawl really an issue?

- External Service Provider
 - Issue or revenue opportunity -- question sustainability
 - Pricing by # of VMs keep VM Sprawl in check
 - Potential for this to be an issue with growth in the public cloud
- Enterprises
 - More likely - especially if distributed processes not in lockstep
 - Currently not a top server virtualization issue but can potentially be if left unchecked

What are the issues with VM sprawl?

- Security
- Down rev level OS Images
- Wasted storage capacity





Techniques to Cope with VM Sprawl

- Identifying orphaned VMs
 - Orphaned VMs are not really idle
 - Consuming memory and CPU cycles
 - Consuming disk capacity,
 - Add complexity to data protection
 - Variety of software tools available to help track and automate VM activity and state\
- Application Stacking
 - Fundamentally about reducing # of OS images
 - Harder to in X86 server environments
- Advanced Storage Techniques to mitigate risk associated with VM sprawl



Advanced storage techniques to cope with VM Sprawl

- Thin Provisioning
- Thin Reclamation
- Tiered Storage
- Array based snapshots

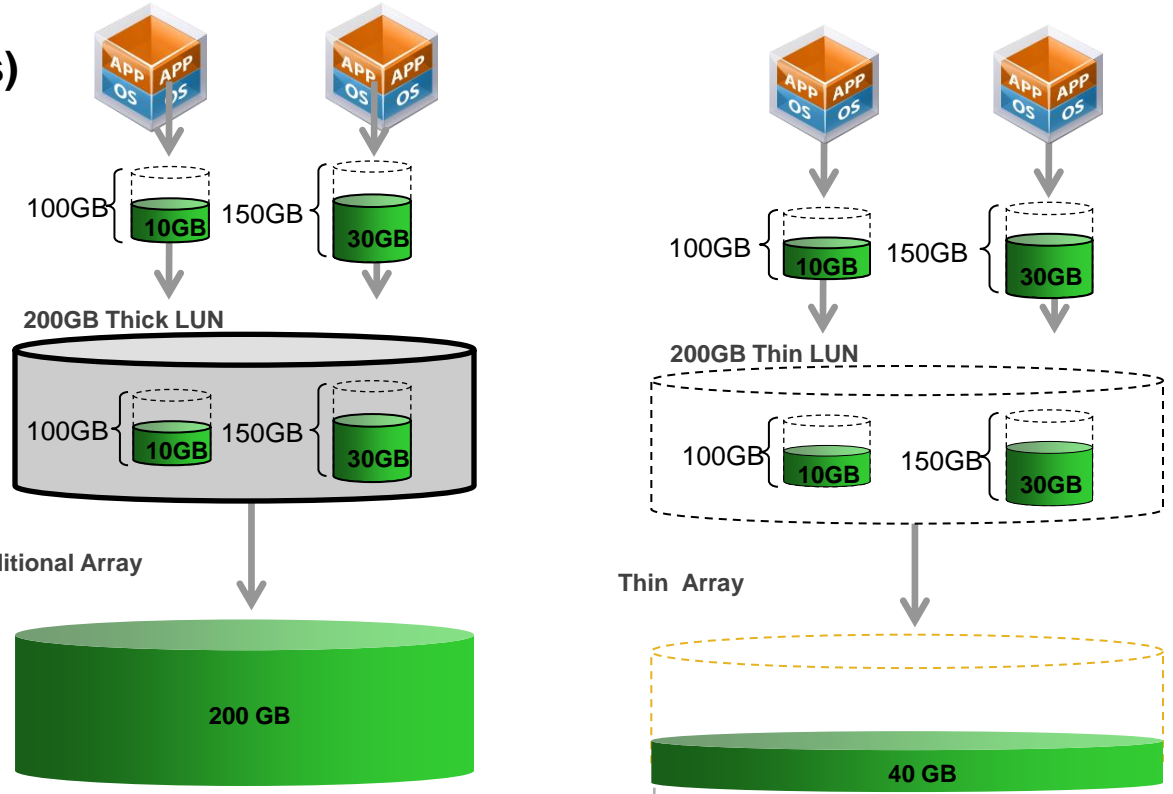
Array based Thin Provisioning drives capacity reductions

Virtual Machines (VMs)

Thin Virtual Disks

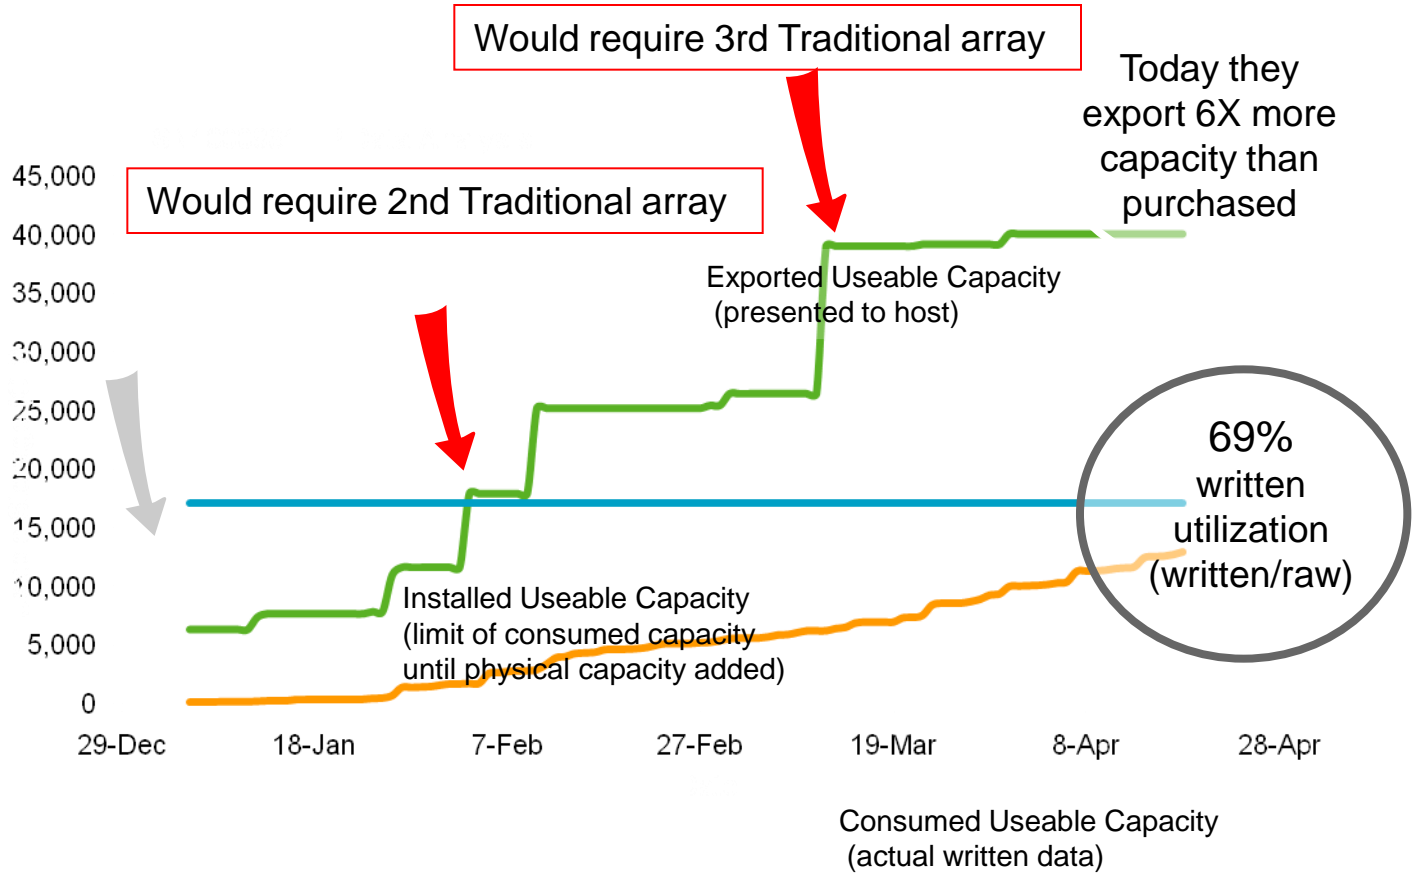
Datastore

LUN



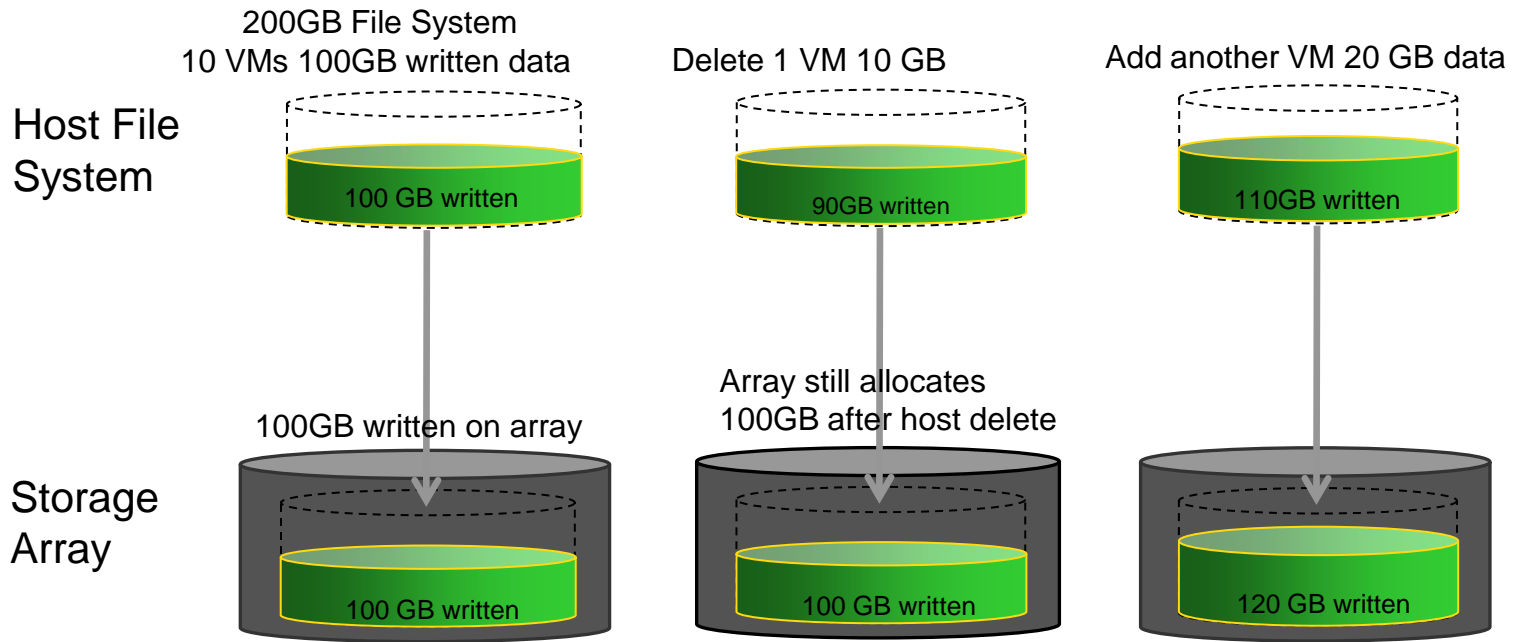
Over provisioned VMs:	250 GB	250 GB
Physically Allocated:	200 GB	40 GB
Capacity Savings:	50GB	210 GB

Thin Provisioning Case Study: Global Investment Bank



\$3M Saved in Storage Purchases

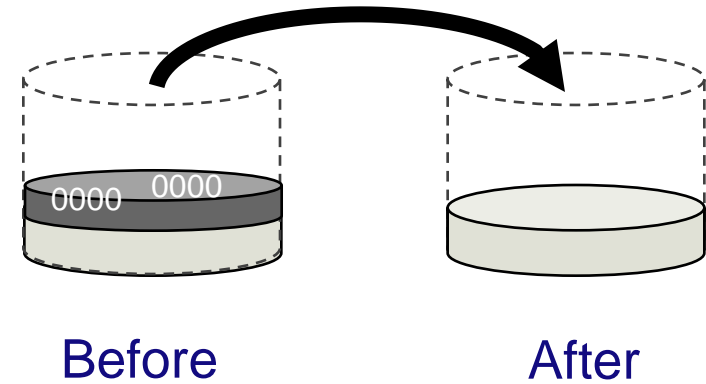
The challenge with reclaiming storage space with Thin Provisioned VMs



Deleted space on a host file system does not result in deleted space on the storage array

Considerations for reclaiming capacity from Thin Provisioned volumes

- Useful when there are a large number of deleted orphaned VMs
- Non-disruptive and application-transparent “re-thinning”
- Returns space to thin provisioned volumes and to free pool for reuse
- Simplicity – any special software or services required.
- Practical and efficient

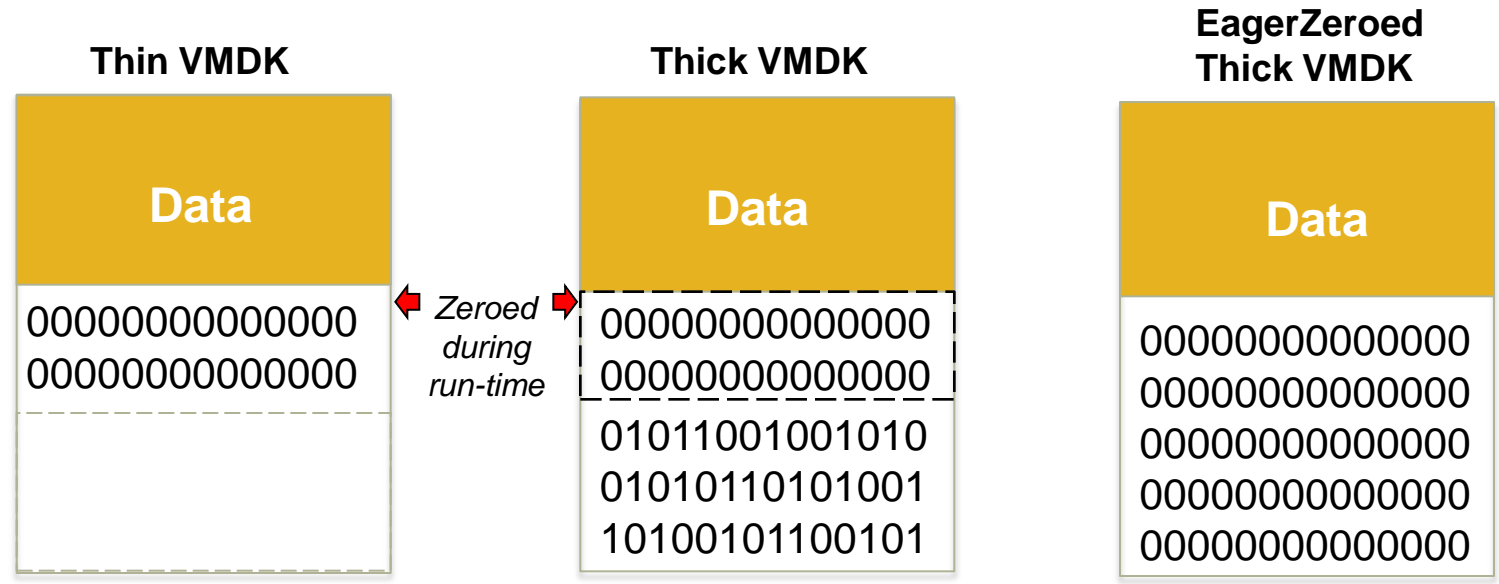


Case Study: Large online media company reclaim storage from Thin Volumes

- VMware Thick provisioned VMDKs.
- Deploy Thin Provisioning at the array level
- Large number of VM deletions and creations – 40 GB VMs, 50 to 60 creates/deletes per day
- Uses VMware EZT and zero detection to reclaim space

Understanding Virtual Machine disk file formats

- VMware VMFS supports three virtual machine disk (VMDK) formats
- Zeroing of space required to ensure data security

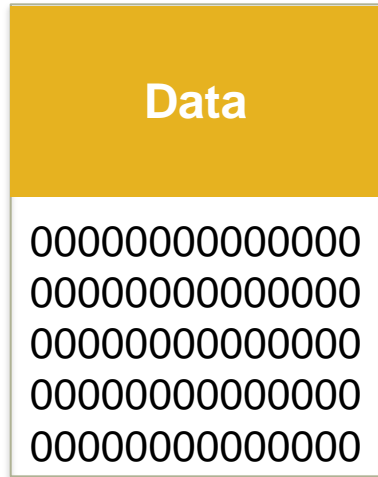


	Thin	Thick	EZT
SPACE ALLOCATION	On Write	On Create	On Create
ZEROING	Run-time	Run-Time	Create-Time

Performance and reclamation capabilities without capacity tradeoffs

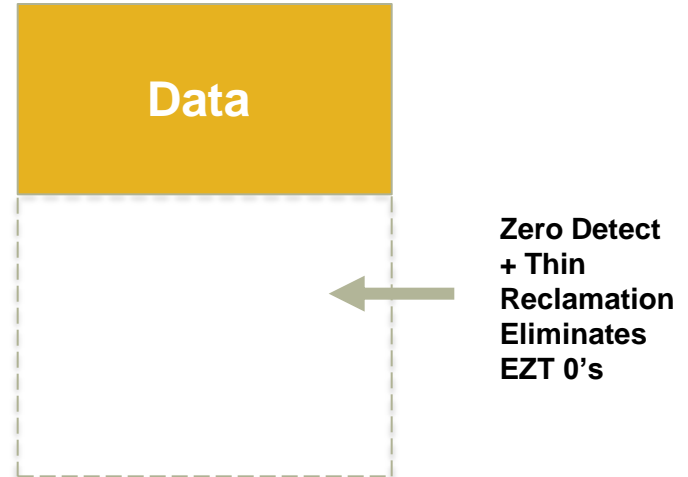
- EZT and Zero detect allow for capacity efficiency
- Storage reclamation use case with creating and deleting large EZT VM

EZT on older arrays



20 GB VMDK (4 GB Data, 16GB Zeros) consume 20GB capacity

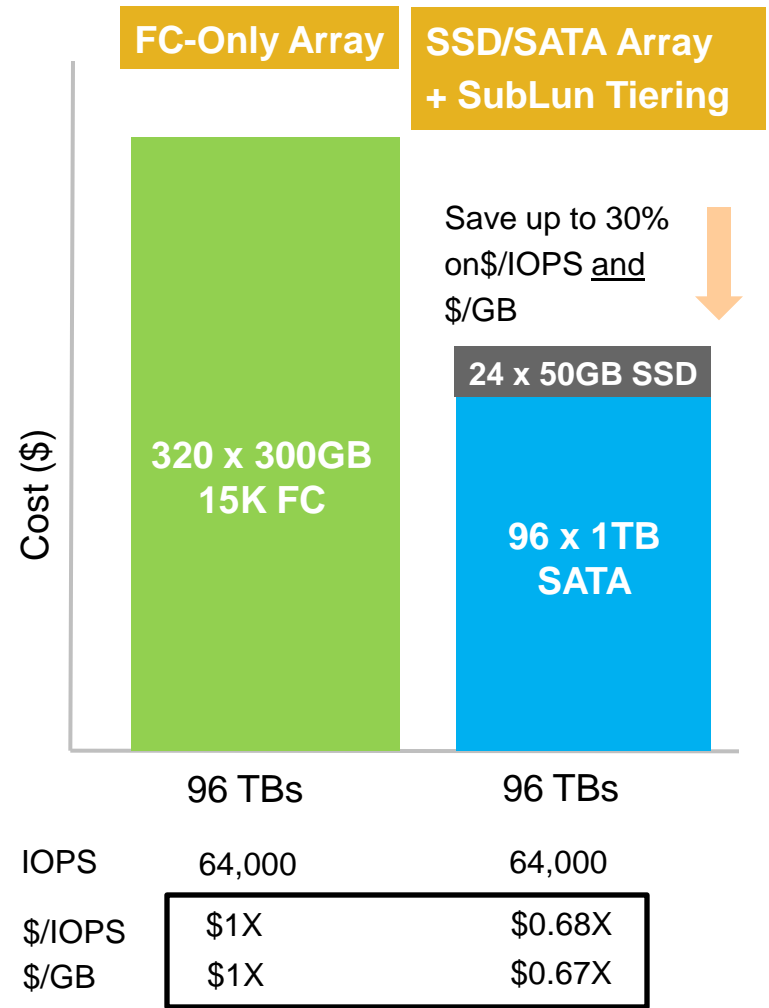
EZT on current array



20 GB VMDK (4 GB Data, 16GB Zeros) consume 4 GB capacity

Automated Storage Tiering

- Sub-volume tiering migrates portion of a LUN to lower cost
 - (I/O associated with orphaned VMs for example)
- Lower overall storage cost and mitigates risk of VM sprawl
- Take advantage of this capability with no additional configuration or tuning in the VM or application layer
- Cost optimization occurs automatically via defined policies in the storage array





Summary

- VM sprawl can negate the benefits of Server Virtualization
- Process improvement, appropriate software tools and techniques crucial to keep VM sprawl in check
- Don't underestimate the impact of storage utilization and inefficiency with VM Sprawl