



# Effective Data Management in a Tiered Storage Solution

Steve Aaker – Principal Product Manger  
Damon Clark- Sr. Principal Software Engineer  
Oracle

August 14, 2013  
Session: 13764



# Agenda

- Challenges Storage IT organizations face
- Getting jobs done at the lowest cost
  - Right tool for the job
  - Oracle Solutions
- Summary

## IT Challenges- Balance Requirements to Cost

Faster RTO's

Shorter RPO's

Integration With Existing  
Infrastructure

Shorter Backup Windows

Data Growth

Versus



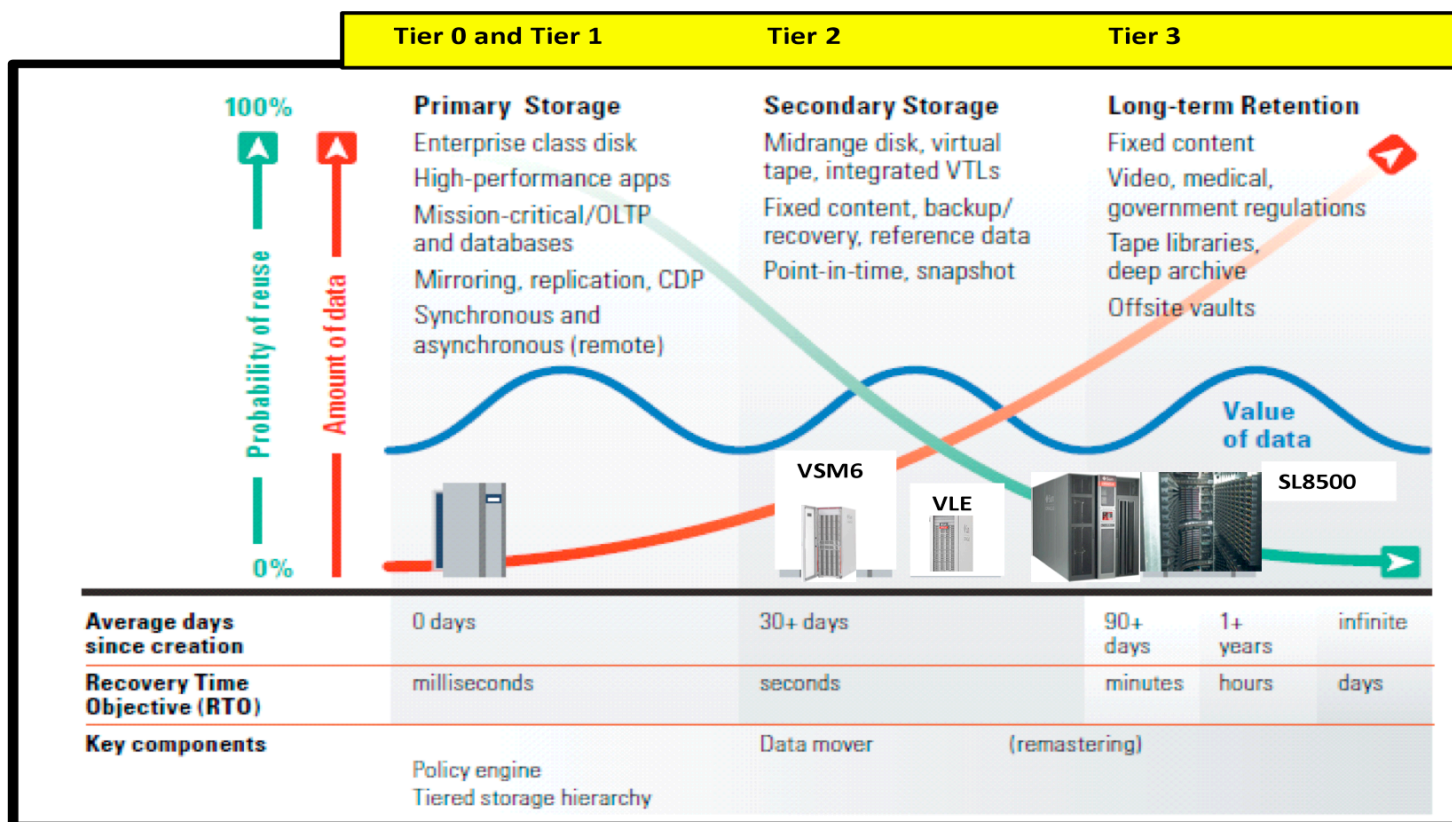
Reduced Cost

**Murphy's Law:  
Anything that can go wrong  
will go wrong.**


# Lifecycle Data Protection – Data Usage

## Lifecycle Data Protection with VSM6

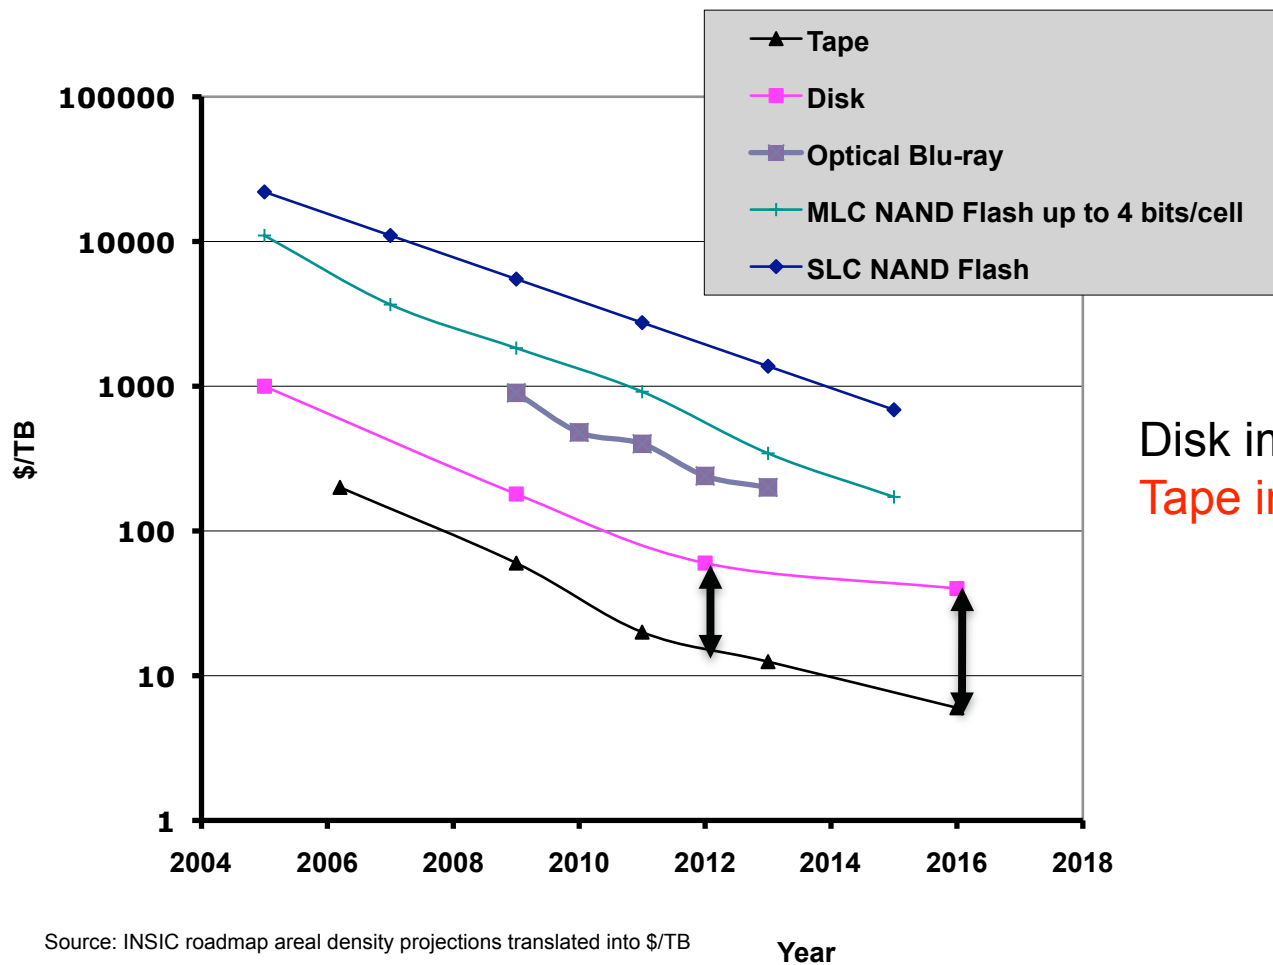
*Integrates storage virtualization, disk, and tape*



Source: Horison Information Strategies



# Cost Trends



Disk improvements: ~20%/yr  
Tape improvements: ~40%/yr

# It's Not Only About Cost/TB



	Disk	Tape
Total cost of ownership (archive) <sup>1</sup>	15X	1X
Total cost of ownership (backup) <sup>2</sup>	2-5X	1X
Max shelf life (bit rot)	10 years	30 years
Best practices for data migration to new technology	3-5 years	8-12 years
Uncorrected Bit Error Rate, (probability average 1 error in x TB)	10 <sup>-14</sup> ( ~10's of TB)	10 <sup>-19</sup> (~1 million TB)
Power and cooling <sup>1</sup>	>200X	1X
Labor (TB managed per storage admin) <sup>3</sup>	100's	1000's

“The cost of energy alone for the average disk-based (archive) solution exceeds the entire TCO of the average tape-based solution.”<sup>1</sup>


<sup>1</sup> The Clipper Group, “In Search of the Long-term Archiving Solution”

<sup>2</sup> Enterprise Strategy Group, Inc. “A Comparative TCO Study: VTLs and Physical Tape Solution”

<sup>3</sup> Moore, F. Horison Information Strategies, “Tiered Storage Takes Center Stage”

Complete your sessions evaluation online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)





**SHARE**  
Technology • Connections • Results

# Tiered Storage

- Definition
  - Taking advantage of two or more storage technologies to build a better storage solution
    - Get the benefits of each technology
    - Avoid the pain points of each technology
- Storage Technologies
  - Solid State Disk
    - High Cost however higher IOPs, Random Access, Fast
  - High speed Disk (15000 RPM)
  - SAS/SATA Disk (7200 RPM)
    - RAID, Log Structured File Systems
  - Physical Tape
    - Lowest Cost with high capacity lower performance

Complete your sessions evaluation online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)




# Tier 1 – High Performance

- Provides repository for active data
  - Large primary disk system
  - Fast, immediate access
  - Replication
    - Asynchronous, Synchronous
    - 1-1, 1-many, many-many
  - Non-disruptive maintenance and upgrades
  - High availability



## Middle tier – Processing Layer


- Economical disk
  - Fast recall
  - Larger nearline storage
- More efficiently utilize very large tape repository
  - Reduce tape recall
- Offload peripheral activities from Tier 1
  - Frees Tier 1 system to handle service requests
  - Copy/replication to DR site
  - Deduplication
    - Performance precludes use at Tier 1
    - Most applicable at Tier 2 – low throughput/recall impact
    - Move de-duped data to remote site



# Tape tier – The Massive Repository


- Very large storage capacity
- Cost
- Encryption
- Tape automation removes the manual interaction
- Long term protection
- Infinitely scalable, grows as needed
  - Low impact, inexpensive growth
    - Add media to increase capacity
    - Add drives to increase throughput
  - Long term data retention
- With new technologies can grow up to 1EB and beyond in a single library

# Simple Tiered Virtual Tape with Oracle's StorageTek Virtual Storage Manager (VSM)




- Provides seamless integration of hardware and software
  - Centralized management for all parts of storage system
- Performance and cost of storage closely matched to type of data

- **Current data – High Performance Disk 10TB – 307PB**
- **Recent data – Economical disk 330TB–338PB**
- **Archival – Tape – 1000PB+**



**SHARE**  
Technology • Connections • Results


# Large Capacity Data Storage – **Batch/Backup**



...




- Disk-only
- Benefits:
  - No removable media




**SHARE**  
Technology • Connections • Results

# Large Capacity Data Storage – **Batch/Backup**



- Disk-only
- Benefits:
  - No removable media

# Large Capacity Data Storage – Long term backup and/or archive



System z

FICON



Optional



- Disk and tape system
- Benefits:
  - Fast access for most recent data
  - Large capacity repository
  - Persistent, long term storage

# Best Practice: “Temporary” loss of data access

Failure of storage device, infrastructure, software, people

## Risk Profile

- The risk *seems* small
  - Mainframe devices boast 99.999% (five nines) – approx 5 minutes outage a year but...
    - Data access can require multiple components to be available – data access could be less than 5 nines
    - One component on the data path with less than five nines significantly changes the outage probability



## Response

- Plan for High Availability (HA)
  - No single point of failure
  - Continuous access to data
- Can include
  - More hardware
  - Higher cost
  - Longer Job times

# Best Practice: Protection Against Temporary Outages

## High Availability with Virtual Tape

### Virtualization of Tape makes simple HA possible

- Many copies of tape data can exist in:
  - multiple locations
  - multiple storage technologies




### HA tape solutions need to

- Global access to all data
- No single point of failure
- Ease of fail-over and fail-back




# Examples: Simple Tiered Virtual Tape with Oracle's StorageTek Virtual Storage Manager (VSM)



- Policy Management can be used to secure multiple back-end copies in multiple locations and storage tiers

# Examples: Simple Tiered Virtual Tape HA with VSM


Multiple Copies of Data in Multiple Locations



- Policy Management can be used to secure multiple back-end copies in multiple locations
- VSM is a “global” solution
- All VSM systems can access all back-end storage data
- No single point of failure
  - All z/OS hosts can access all VSM systems
  - All VSM systems can access all back-end data
- Fail-over & fail-back is trivial
- Solution can scale massively

# Examples: Virtual Tiered Tape HA with VSM


## Multiple Copies of Data at Two Sites



- Within a metro area 2 sites can provide further HA capabilities
- As before all hosts can access all data
  - Second host is optional
  - Single host must access both VSMs

# Examples: Virtual Tiered Tape HA with VSM


Data at Both Sites Can be Synchronized With Clustering



- Within a metro area 2 sites can provide further HA capabilities
- As before all hosts can access all data
- Now the virtual tape solution at both sites can be kept synchronised
  - Data created on one site is synchronously replicated to the other site
- Fail-over & fail-back are trivial events
- Active – Active
- Hosts can preference local VSM

# Examples: Virtual Tiered Tape HA with VSM

Data at Both Sites Can be Synchronized With Clustering



- Within a metro area 2 sites can provide further HA capabilities
- As before all hosts can access all data
- Now the virtual tape solution at both sites can be kept synchronised
  - Data created on one site is synchronously replicated to the other site
- Fail-over & fail-back are trivial events
- Active – Active
- Hosts can preference local VSM
- VSMs can preference local Real tape drives or Virtual tape drives

# Challenge: “Catastrophic” Event

Long term or permanent disruption to entire facility

## Risk Profile

- Probability is extremely low
- Impact is extremely high



## Response

- Plan for Disaster Recovery (DR)
  - Data & Metadata in multiple geographically dispersed locations
    - Metadata includes Tape Catalog, Tape Management System
  - Adequate compute, network, personnel and storage resources to maintain business
  - A Tested and Validated DR plan in place

# Protection Against Catastrophic Events

## Disaster Recovery and Business Continuance with Virtual Tape


Multiple copies of tape data can be distributed in multiple sites

### Key strategic questions

1. What are the demands of the business on the data?
  - Recovery Point & Recovery Time Objectives (RPO, RTO)
2. What is the plan to re-establish data access?
  - How to get data back to a known and consistent point?
  - How to simplify the recovery plan?
    - Assume it is not an existing IT employee performing the recovery
  - How to verify the recovery plans?

# Examples: Using Tiered Virtual Tape for Disaster Recovery

Let Virtual Tape Manage the Data Movement




- Host writes once and the virtual tape solution makes all the copies.



# Examples: Using Tiered Virtual Tape for Disaster Recovery

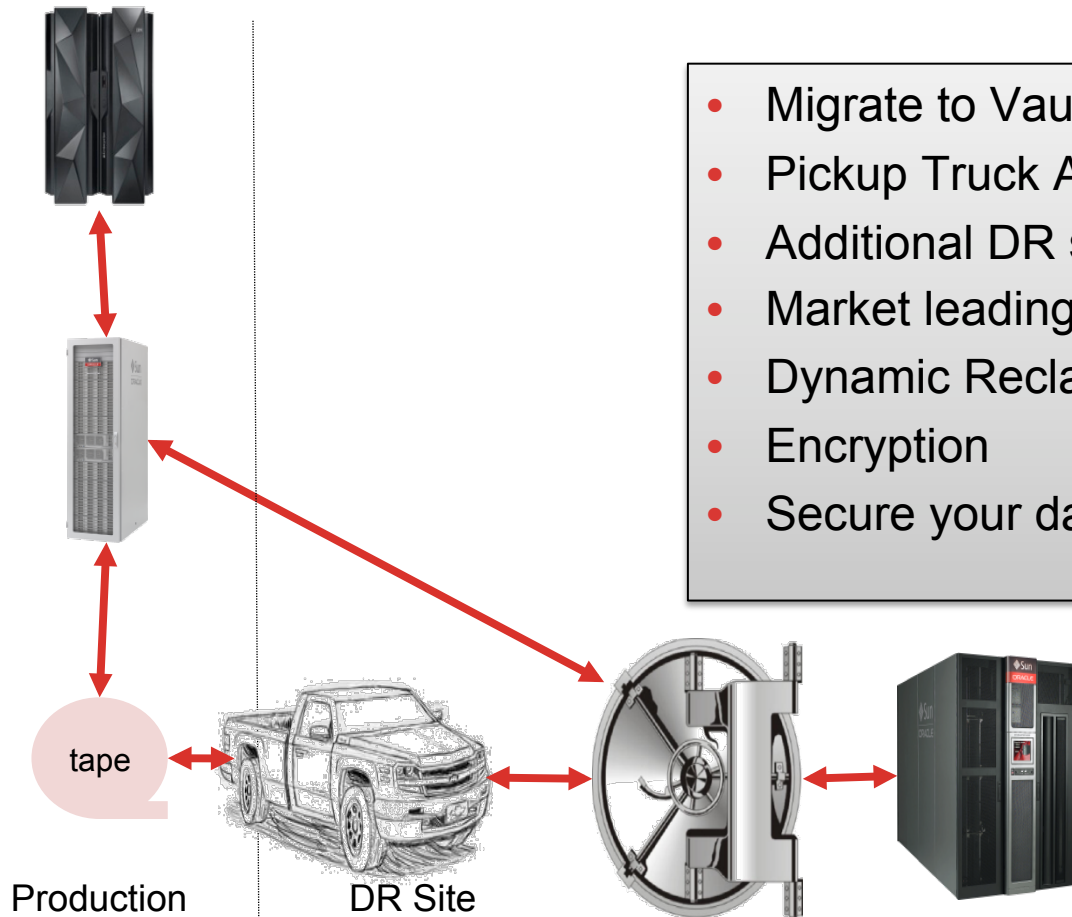
Let Virtual Tape Manage the Data Movement



- Host writes once and the virtual tape solution makes all the copies.
- Replication (Copy)
  - VSM->VSM
  - Synchronous
  - Asynchronous
- Migrate to Virtual Library Extension (VLE)
  - Tier 2
- Migrate to Real Tape
  - Tier 3

# Examples: Using Tiered Virtual Tape for Disaster Recovery

Driving Physical Tape Offsite is Another Alternative



- Migrate to Vault
- Pickup Truck Access Method
- Additional DR sites
- Market leading tape capacity is now 8+ TB
- Dynamic Reclaim (partitioned tape)
- Encryption
- Secure your data

# Vaulting is Acceptable for Long RTO Data

Supported in Virtual Tape as an EXPORT/IMPORT Type Function

## Pros

- Beautifully simple and easy to understand
- Generates consistent data
  - All data and metadata represents a well known point in time
  - Cost effective bulk data transportation.
- Compliance

## Cons

- Security risk unless data is encrypted
- Difficult to test without additional cost


# Why Tiered Storage?



- Not a new concept. The fundamental concept has been around since the '70's
- What has changed?
  - Virtual Tape
    - High capacity/economical disk
    - Very high capacity tape
  - Economics of storage technology
  - Storage software has evolved
  - Business Application software now comes with embedded tiered-storage capabilities
  - Optimized Solutions Offerings

# Summary

- Match data type and usage patterns closely to the type of storage
  - Design with maximum flexibility in mind
- Plan for more growth
- Include technology migration in your plans
  - Disruptive/non-disruptive
- Consider “peripheral” factors
  - Expected “shelf life” of data on medium
  - Expected lifetime of solution
  - Cooling and power consumption



**SHARE**  
Technology • Connections • Results





**SHARE**  
Technology • Connections • Results

# Back-up

# VSM Cross TapePlex Replication

## Active vs. Inactive CPU at Disaster Recovery Site


### Active DR CPU

- Synchronous copy (PPRC) production CDS or output of DRMON to DR site
- Use DRMON or VTV Report to determine what data was not replicated

### Inactive DR CPU

- Use LPAR at Production site to host HSC/VTCS (MULTMODE)
- Remote connect to DR CDS or synchronously copy (PPRC) DR CDS to DR site.
- Synchronous copy (PPRC) production CDS or send output of DRMON to DR site.
- Need extended FICON control path from HSC/VTCS LPAR to DR site VSM (ECAM OVER IP with VSM6 will address this)
- Use DRMON or VTV Report to determine what data was not replicated
- After DR event, power on CPU, use DR CDS





# VSM Remote Migration

## Remote Real Tape Drives vs. Remote VLE for Disaster Recovery

### Remote Real Tape Drives

- Need extended FICON for remote RTDs
- Synchronous copy (PPRC) production CDS
- Optionally send output of DRMON to DR site
- After DR event, power on VSM and CPU, use copy of production CDS
  
- Use DRMON output or VTV Report to determine what data was not migrated

### Remote VLE

- Uses IP for control and data path
- Synchronous copy (PPRC) production CDS or send output of DRMON to DR site
- After DR event, power on VSM and CPU, use copy of production CDS or run VLE Audit to create new CDS
- Use DRMON output or VTV Report to determine what data was not migrated
- All DR Data must fit in VLEs