Implementing a Modern Backup Architecture

Oracle's Tiered Data Protection Strategy

Fred Moore, President
Horison, Inc.
www.horison.com

Introduction

Fail-proof data backup and recovery is more critical to an organization's survival than ever before as few businesses can survive for any period of time without their IT function. With so much reliance on electronic data, an organization can lose millions of dollars associated with lost data including its competitive advantage and credibility as in the case of security breaches. Implementing a *cost-effective* fail-proof backup and recovery capability enables an organization to protect itself from data loss and costly downtime that result from hardware or software failure, power failure, natural disaster, intrusion, or human error. With our focus on backup architectures, it's important to understand the various roles of backup. We can define "backup" as simply a tool or method for executing two primary functions in the data center:

Business continuity - to provide a local copy of data to be used should an application or infrastructure component fail or data become corrupted. For business continuity, fast initial access time is critical. Therefore, disk is the preferred choice, enabling fast backup and restore of small chunks of data such as system and user files, emails, and incremental backups.

Disaster recovery - to provide a copy of data which can be maintained off-site and which can be restored from another location, should the primary data center facility no longer be available. For disaster recovery, fast data transfer time and high availability are critical. Therefore, tape is the optimal choice, enabling fast backup and restore of large chunks of data (large databases, servers, or even entire data centers).

The optimum data protection solution enables both business continuity and disaster recovery by taking advantage of both disk and tape technology. In addition, tape technology has surpassed disk in reliability and remains the most cost effective means of protecting long-term archival data (compliance, full backups, video, scientific, medical imaging, etc.).

Business Continuity and Disaster Recovery

It's important to distinguish between a business continuity event and a disaster recovery event as they are different. Business continuity is the activity performed by an organization to ensure that critical business functions will be available to customers, suppliers, regulators, and other entities that must have access to those functions. These activities can include many daily chores such as file and data base backups, project management, change control, or help desk services. Business continuity is not something implemented at the time of a disaster - business continuity refers to those activities performed *daily* to maintain quality of service, consistency, and recoverability. Disaster recovery (DR) is the process, policies and procedures that are related to preparing for recovery which are vital to an

organization after a natural or human-induced disaster that may impact the entire data center rather than specific files. While business continuity involves planning for keeping all aspects of a business functioning in the midst of disruptive events, disaster recovery focuses on the IT or technology systems that support business functions.

A data recovery event involves restoring the most current image copy from disk or tape, then applying any applicable log data to get to a valid recovery point. With tape, when a backup image expires, the tape cartridge is returned to the backup scratch pool for reuse. With disk, when a backup image expires, the image is erased from disk and the disk space is freed up for more backups. In a well-run organization, the back-out event is planned but rarely executed. Unplanned downtime, by its nature, is a surprise and can have numerous causes including hardware failures, software errors, user errors, poor maintenance, or a system upgrade as shown in the chart below.

Type of Storage Disruption	% of Incidents	Possible Protection Solutions
Hardware, network, or system failure	44%	Mirroring, RAID, backup copies, transaction logs, clusters, failover, virtualization
Human error, equipment theft	32%	Snapshots, CDP, surveillance, facilities security
Software and coding errors	14%	Encryption, CDP, replication
Intrusion – security, virus, web scams	7%	Firewalls, authentication, anti-virus, encryption
Natural disasters, power, flood, building damages, fire….	3%	Off-site compute facilities, emergency power, UPS

Source: Data from a variety of industry sources

The recovery process can depend on the cause of the outage. That's why you need to perform detection and analysis to determine the cause and scope of the failure. For example, a software program failure could result in one record being impacted or multiple databases being corrupted. If the recovery process is manual, the system operator may not be familiar with the recovery options and the process can be time consuming. An RTO (Recovery Time Objective) must be established for critical applications and files at a minimum and it can be different for certain databases or applications, depending on business needs.

Understanding the Impact of Downtime and Data Loss

Before updating the recovery processes for your organization, it's important to understand just how costly unplanned as well as planned downtime can be. Planned downtime includes various activities such as database maintenance (performing copies and reorgs) and required system upgrades. Planned downtime typically doesn't result in a recovery event unless the reason for the downtime requires the change to be backed out (e.g., a structural change is reversed or abandoned).

As shown below, the cost of an hour of downtime may be damaging and the cost of several days of downtime could exceed a billion dollars. Remember these are industry averages; you should calculate *your* specific cost of downtime per hour for business resumption and disaster recovery purposes. It's no wonder that implementing the most cost-effective data protection strategy that meets the availability requirements of a given application has become today's most important IT strategy.

Estimated Cost of Downtime per Hour	
Energy	>$5.0 million
Telecommunications	>$4.0 million
Finance	>$2.8 million
Retail	>$2.2 million
Transportation	>$1.4 million
Healthcare	>$1.4 million

Source: Horison, Inc. and a variety of Industry Sources

RTO (Recovery Time Objective) – How long can businesses successfully operate without this data? Once the problem is discovered, RTO is the time required to recover from a data loss event and return to service. In other words, this requires classifying data or an application by its criticality or value to the business and determining how long the business can survive without having this data available. RTO is a key metric as there is constant business pressure to reduce the length of time it takes to recover a critical file or an application.

RPO (Recovery Point Objective) – The desired amount of time between data protection events.

Keep in mind that RTO and RPO are user defined policies based on the criticality or business value of the data being protected. Defining the RTO and RPO are key components of a successful data classification plan. While the user doesn't have much control over how long it takes for the problem to be discovered, they do have control over what means (technology solution) from which to recover and the level of redundancy. In optimal data protection architectures, there's always a tape safety net in case a recovery from disk fails (regardless of the data's value).In addition, an HSM allows you to set policies to take advantage of the economics and superior reliability of tape where it makes the most sense from a value/RTO-RPO perspective. The following table provides suggestions for which types of recoveries should optimally be supported by disk and which ones should be supported by tape.

Data Classification by Value	Mission-critical	Vital	Sensitive	Non-critical
Avg. Data Distribution (varies)	15%	20%	25%	40%
Availability index	99.999%	99.99%	99.9%	99%
Downtime minutes/year	5.256	52.56	525.6	5256
Typical RTO Typical RPO				
30 min 0 min.	Disk	Disk	Disk	Disk
2 hours 15 mins.	Disk	Disk	Disk	Disk
12-24 hours 2-6 hours	Disk	Disk	Tape	Tape
>1 day 12-24 hours	Disk/Tape	Disk/Tape	Tape	Tape
> 1 week > 1 day	Tape	Tape	Tape	Tape

Note: $ impact of data loss = min/year of downtime x lost revenue/min.

Planning for a Modern Data Protection Strategy

In its most basic form, we can define "backup" as simply a tool or method for executing business continuity processes (with disk) or disaster recovery processes (with tape). The challenges of effective backup are numerous given the range of application availability requirements, but so are the options. Effective data protection plans address both business continuity and disaster recovery by defining operational procedures, implementing hardware redundancy, and practicing/testing the recovery processes. With mounting pressure to reduce the amount of time required and the amount of storage consumed, many new and improved backup/recovery methods are now available. These offer a variety of choices depending on what operating system is used, the type of storage technologies used, when and how the data is maintained, if compression, encryption or WORM is used, and if any additional geographic locations are involved.

Backing up and later restoring potentially huge amounts of data in the least disruptive manner is becoming increasingly difficult given the tremendous amount of digital data growth. Disk has become the preferred business continuity backup target for smaller data files demanding the fastest RTO (Recovery Time Objective), while tape is the preferred backup choice for large files and disaster recovery. An all-disk data protection and archiving solution is an increasingly expensive option. Recently published studies indicate that the 5-year TCO (Total Cost of Ownership) for disk ranges up to 15X higher than tape for backup and archiving. The initial acquisition cost or purchase price per GB is also much lower for tape. Remember - data that is not being used should not consume energy. To best address business continuity and disaster recovery, the optimal solution has emerged and deploys a tiered storage approach using **both** disk and tape.

What Are the Best Protection Options Available?

A variety of data protection schemes exist and as expected, higher levels of data protection cost more to implement. Software errors, human errors, natural disasters, increasingly common power failures, building damages, and destructive intrusion such as worms, viruses', scams, and spy-ware have turned data protection into a complex data management process. Data requires different types of protection depending on its criticality and the point at which it resides in its lifecycle. For example, mirroring data on disk that has reached archival status is unnecessarily expensive given the data is seldom used. Some businesses have tried to go "tapeless" for their entire data protection architecture; however this is the most expensive option and yields a sizeable gap in their disaster recovery strategy. Disaster recovery processes are better served with the high availability and fast transfer time of modern tape technology.

The tremendous advances in the past decade for reliable, low-cost, high-capacity, high-speed magnetic tape storage has been among the key factors helping tape expand its foothold in building optimal data protection and archive strategies that combine disk and tape. Not everyone is aware that tape has become more reliable, has a faster data rate, is less expensive to own and operate, has a longer media life (30 years), and has a higher capacity than disk. The chart below summarizes the myriad of backup and recovery options available today with disk and tape combining to deliver the optimal solution.

Disaster Recovery Options	Description	Data types supported	Data compressed on backup?	Storage devices used	Failover and/or restore
Full backup	Straightforward process, maximum storage consumption, used on all applications, time consuming	All data supported, used for hot sites	Yes	Tape	Restore
Mirror (disk)	Doubles disk costs, only protects from disk HW failures, does not protect against intrusion, data corruption or human error	All data types supported, used for hot-sites	No, one-for-one copy	Disk option only, can mirror remotely	Failover
Business Continuity Options	**Description**	**Data types supported**	**Data compressed on backup?**	**Storage devices used**	**Failover and/or restore**
Incremental	Reduces backup window, all incremental copies used in recovery making recovery longer than differential	Application specific, not supported by all applications	Yes	Tape, disk	Restore
Differential	Reduces recovery time, only last differential copy used in recovery making backup time longer than incremental	Application specific, not supported by all applications	Yes	Tape, disk	Restore
De-duplication	Fast backup, reduces size of backup recovery load, reduces disk storage requirements and expense but is more costly than tape, compute intensive requiring costly appliance	Application and backup software **neutral**, avoid using for non-duplicate data	Yes, provides further data reduction via eliminating redundant data	Disk only	Reconstruct data first from metadata, then restore
Snapshot	Very fast, does not protect against disk failures, does protect against intrusion or corruption, can take time to determine the RPO	Application specific, supported by most applications	No	Disk only	Restore
Continuous Data Protection (CDP)	Is often application specific, can take time to determine the recovery point, fast recovery, protects against intrusion and corruption	Application specific, not supported by all applications	No	Disk only	Restore
VTL (Virtual Tape Library)	Virtualized disk array that appears as tape library with multiple tape drives	All data types supported	No	Low-cost disk	Restore

Business Continuity *and* DR	Description	Data types supported	Data compressed on backup?	Storage devices used	Failover and/or restore
Tiered storage	Optimizes RTO and costs matching backup/recovery and archive solutions with application requirements	All data types supported	Yes	Disk and tape together	Optimal failover and restore

Source: Horison, Inc.

What Are the Tradeoffs of a "Disk Only" Data Protection Strategy?

From the chart above, several straight-forward data protection options are available. Choosing a disk-only data protection strategy often requires deduplication to be evaluated in the process. Before investing in deduplication, it's important to really understand what you are getting into. Deduplication works best with large amounts of duplicate data that do not significantly change from day-to-day. Claims of 70%, 80% or even 90% space reductions in the size of the backup load from using deduplication are initially appealing, but are frequently misleading as these claims are seldom attained unless the user first *turns off* compression and incremental or differential backup processes. This process returns the data to its original size, which enables deduplication to deliver a bigger space reduction. Thus, in order to attain the large space reduction claims, the process often involves significant user intervention.

While data deduplication running on backup servers can potentially save considerable time and the amount of disk space required for backups, what impact does data deduplication have on data recovery operations? For some reason, this question is not often asked but the answer is it actually can slow things down a little – or a lot. When it comes to data recovery, deduplication introduces an additional single-threaded, time consuming step in which indexes or other pointer structures are used to first reconstruct or "re-hydrate" de-duplicated data before it can be used. Currently, the only way to get a copy of your data off of a deduped array is to first re-hydrate it and then to write it back out in its original size. That means if you have data that's deduped at 30:1, roughly 1.5PB of data can be backed up and stored on a smaller 50TB array. You still have to write out all 1.5PB somewhere else when the data is recovered, either to a 1.5PB disk array or 1.5PB of tape. Having to re-hydrate the data prior to restore defeats much of the purpose of having a deduped device in the first place. Furthermore, the high cost of the deduplication engine can offset the cost savings from reducing the number of backup disks and is an additional expense consideration.

Network data recovery poses further considerations. When recovering data over a LAN from a backup server with data deduplication, the LAN speeds can be the limiting factor and not data re-hydration. Recovery over a WAN is another matter, however. While data deduplication significantly reduces the load on the WAN for backups (after the first backup or deduplication scan), a huge amount of data must still be sent over the WAN for a recovery. This problem gets some relief when there are two data deduplication engines operating in concert at each end, however deduplication engines are very costly and published TCO studies indicate that their cost alone will usually offset any disk savings.

It is only when protecting a large data set for disaster recovery purposes that these considerations could be the difference between recovering successfully with tape or not at all with disk. Today's smart tape compression techniques select the best algorithm and typically yield 2-3x compression ratios and allow customers to restore the data on any system, reducing the degree of lock-in inherent with deduplication. In most cases, none of the above potential shortfalls are a major issue with relatively small data sets. For disaster recovery purposes, spending thousands of dollars to protect the data on tape or possibly millions of dollars to protect it on disk becomes an easy choice for most businesses.

Implementing a Modern Data Protection Architecture with Oracle Storage Solutions

The tiered storage architecture is widely supported by numerous software vendors including Oracle, though few vendors offer *all* storage tiers and management software. The primary purpose of tiered storage software, regardless of vendor, is to automate policy-driven movement of data between the tiers. Oracle is one of the few companies to offer a complete tiered storage solution that supports a full

selection of storage hardware tiers 0-3 along with open systems HSM software.. With a full range of disk and tape offerings coupled with StorageTek Storage Archive Manager, Oracle has in place the key ingredients to position itself as the leading provider of tiered storage solutions for cost-effective business continuity and disaster recovery.

SAM-QFS

StorageTek Storage Archive Manager (SAM) with its Quick File System (QFS) software is key to Oracle's tiered storage strategy. QFS is an open source, POSIX compliant file system. Together, SAM-QFS provides policy-based tiered storage management and shared file access in Open Systems environments. It is this software that integrates Oracle's tiered storage hierarchy making it an "active" tiered storage solution. Without policy-based software to integrate and manage the storage tiers, the tiers become "islands" of labor-intensive storage to manage and some of the up-front cost savings can be lost to increased management costs. For example, files that have been backed up on disk for a month or more can be seamlessly migrated to tape using SAM-QFS. When an Oracle integrated tiered storage stack of disk and tape hardware are combined using SAM-QFS, the foundation for a highly cost effective and scalable tiered storage backup architecture can be built. The complete range of Oracle's tier 0 through tier 3 storage products is highlighted below with SAM-QFS providing policy-based movement of data between the storage tiers. From these building blocks you can implement the data protection solution of the future.

Oracle Tiered Storage - Building Optimized Storage

Integrated Storage Systems

Tier 0 (SSD) **Sun F5100 Flash Array, Sun SSD, Sun F20 PCIe Card**	
Tier 1 (Primary Disk) **SAN Pillar Axiom 600**	
Tier 2 (Secondary Disk) **Sun Storage 2500M-2 Fibre Channel Array, NAS Sun ZFS Storage Appliance, VLE, VSM**	
Tier 3 (Archival Storage) **Libraries – SL8500, SL3000, SL150** **Drives – STK T10000, STK LTO**	

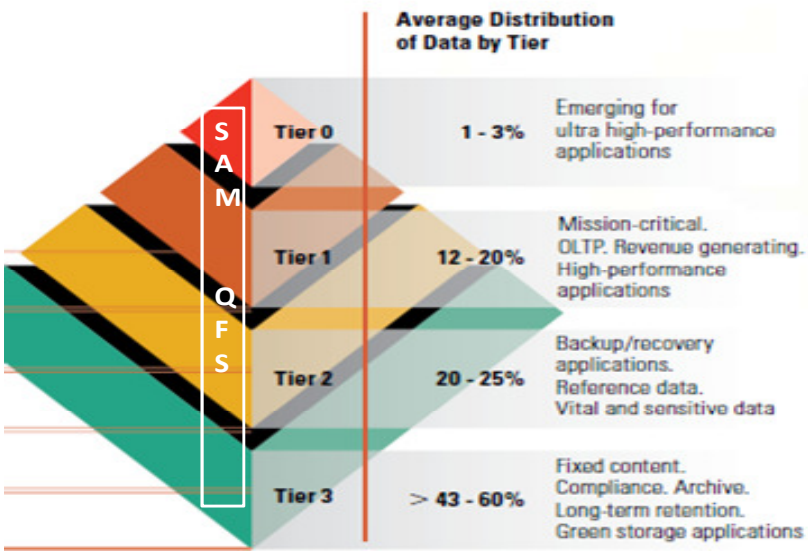

Source: Horison, Inc.

Moving data to newer storage technologies can be time very consuming and is often performed as a manual task in many businesses. An additional benefit of an automated tiered storage environment is that it can relieve much of this burden. Specifically, StorageTek Storage Archive Manager software can

move data from older technology onto new technology. Old devices and media may be retired non-disruptively, automating much of what has historically been a costly and time-consuming manual process.

Oracle Tape and Libraries

Effective use of tape is a key component that makes Oracle's tiered data protection solution so cost-effective. Oracle remains in the leadership position with the entire StorageTek tape product family. The StorageTek T10000C tape drive provides a native cartridge capacity of up to 5.5 TB, making it the largest capacity tape cartridge ever announced. The StorageTek T10000C has a native data rate of up to 252 MB/sec., offering significant improvement in recovery times and making disk drives the limiting factor in large-scale recovery performance. Uncorrectable bit error rate for the StorageTek T10000C tape drive is three orders of magnitude lower than the most reliable Fibre Channel disk drive and offers a media life of 30 years or more. Encryption and WORM are available with the StorageTek T10000C for further protection of data at rest. Oracle's StorageTek SL8500 modular library system can scale to a capacity of more than one exabyte ($1x10^{18}$). The Linear Tape File System (LTFS) format is also available for StorageTek T10000C and LTO tape drives, providing faster disk-like access capabilities for tape files. Relative to disk technology, tape technology progress has been significant on many fronts; tape has a higher capacity than disk, a faster data rate than disk, and is more reliable than disk.

Sun ZFS Backup Appliance – High Performance Business Continuity

Oracle's Sun ZFS Backup Appliance is specifically tuned for business continuity to deliver the fastest backup and recovery for Oracle engineered systems. The Sun ZFS Backup Appliance uses native high-bandwidth InfiniBand interconnects between Oracle storage devices and Oracle's Exadata database servers, delivering up to 20 terabytes per hour full backup and up to 9.4 terabytes per hour full restore throughputs. The Sun ZFS Backup Appliance is embedded with storage efficiency features like Thin Provisioning, Snapshot, Clone copies, In-Line De-Duplication and Oracle's HCC (Hybrid Columnar Compression) technology. The unique HCC technology can yield a 3X to 5X reduction in storage footprint and backup load for customers with existing NAS-based Oracle Databases for database archives, OLTP, data warehousing, or mixed workloads. The Sun ZFS Backup Appliance coupled with Oracle tape provides an unprecedented solution for delivering both business continuity and disaster recovery.

Summary

Data protection has become the most critical IT discipline as most businesses in the modern world can no longer survive without their IT function. As a result, we are beginning to see the next generation of data protection solutions appear as legacy processes become increasingly burdensome, expensive, and unreliable. Today's optimum data protection solution ensures business continuity (using disk), disaster recovery (using tape), and provides the most cost-effective means of protecting long-term storage data (using tape). Oracle's tiered data protection solutions are engineered and integrated with user defined policy management software to take advantage of the performance/cost/availability tradeoffs between tier 0-3 storage devices. Building the data protection solution of the future is attainable – now is the time to develop a solid and sustainable game plan. Oracle is one of the few companies to offer a complete, active, tiered data protection solution by providing a broad range of storage hardware coupled with robust StorageTek Storage Archive Manager software. Data protection is not an option – it is a requirement for survival. It's time to implement the optimal data protection strategy and now you know how to do it.