

Information Means The World.

Increase the Business Value of IT

Datalink helps you store, manage, and protect one of your organization's most important assets information.

From strategy to support, we work closely with you to maximize the value IT delivers to your business.

How Innovations in Storage Change Your Oracle Playing Field



Ari Kaplan Datalink Database Practice & IOUG President Emeritus akaplan@datalink.com 312-399-0079



Agenda

- Market information
- Read-only Snapshots: Enterprise Backup and Recovery
- Writeable Snapshots: Cloning for Development and Testing
- Disaster Recovery: storage-based replication
- Cost savings: tiered architecture / ILM
- Performance and high-availability: RAID-DP and aggregates
- Storage security / encryption
- Virtual Tape Libraries
- De-duplication



IOUG Survey Overview

- 57% reported database growth has impinged upon available storage resources
- 60% report that the lack of available storage has impacted the performance of their databases
- 43% delayed the rollout of an application within the past two years because of lack of storage resources
- 31% are managing 1TB+ database, up from 13% the previous year



Database Statistics

GROWTH OF DATABASE SIZE



The WinterCorp survey found that the world's largest databases have posted an annual compounded growth rate of approximately 75% since 1995.

- The Winter Group 2006



IOUG Survey Details

FIGURE 4: Projected Growth in Storage Needs (Changes over the next 12 months)

	Structured(1)	Unstructured(2)	
Increase by more than 100%	8%	5%	
Increase by 50-100%	12%	12%	
Increase by 25-50%	26%	17%	
Increase by 10-25%	28%	20%	
Increase by 5-10%	13%	10%	
Increase by 1-5%	5%	7%	
Stay the same	3%	7%	
Decrease	0.3%	0.5%	
Don't know/unsure	3%	19%	

(1) Traditional relational database files (2) Email, images, video, other unstructured data



IOUG Survey Details

FIGURE 6: What's Driving Demand for Increased Storage?

Increased business/customer transactions
More input from new systems devices/clients
More storage needed to meet data compliance
Data warehousing project
More storage to increase systems performance
More storage needed for online/e-business initiatives
More replication to secondary or backup sites
Other
Don't know/unsure





Datalink / Storage IO Survey, Aug 09

What are customer challenges in their virtual server environment?

- •Backup and recovery 31.2% indicated this was a challenge
- •Performance issues 25% indicated this was a challenge
- •Contention among networks, storage & servers 18.8% indicated this was a challenge
- •Difficult to manage-12.5% indicated this was a challenge
- •None—31.2% had no challenges
- •All of the above—12.5%

The takeaway here is that ~70% have identified that they have issues in their virtual server environment.



Datalink / Storage IO Survey, Aug 09

Which types of virtualization are you currently leveraging?

- •Server 75%
- •Storage 37.5%
- •Desktop 31.2%
- •Tape 18.8%
- •None 25 %



Issues With Database-Only Solutions

- May be impossible to meet RTO and RPO objectives
- Single storage points-of-failure exist
- May be unable to effectively scale to performance and size needs cost-effectively
- Time to provision new databases is prohibitive or too costly, slowing application lifecycles



The Answer: Combine Storage and Database Solutions

- Meet RPO and RTO objectives with snapshots and storage replication
- Eliminate periods of single points of disk failure with RAID-DP
- Ensure performance and cost are effective with increasing database sizes through aggregates and tiered architectures
- Provision new databases in minutes and without costly storage with snapshot cloning



DBA Challenges

Backup Issues

- Scalability: physically copying 2 TB of data to tape or disk is time consuming
- **Cost**: expensive to purchase 200 TB of storage to perform physical image backups of 200 TB databases; costly to even purchase hardware to test backups
- **Performance:** keeping large databases in hot backup mode negatively affects the performance of high-transaction systems (inserts, updates, deletes)
- **Complexity of systems:** multiple databases, interlinked systems, different database versions, RMAN/non-RMAN, RAC, ASM, etc.
- Manageability: setting up, managing, and testing backups is often difficult

Recovery Issues

- **Manageability:** human errors, lost data, inconsistent data, physical failures, and corruption can require restores; recovering to a consistent point-in-time is a manual and daunting process
- **Performance:** how do you recover a 2 TB database in 15 minutes?

Disaster Recovery / Replication

How do you architect your database and surrounding environment for DR? With no data loss? With a 15-minute failover timeframe?

Data Growth

- **Cost:** DBAs tend to put storage on a single class of storage without archiving or tiering considerations
- **Performance:** system response time is 5 seconds now—what happens when data triples in size?
- Manageability: Getting additional storage from non-DBA groups is often a political process

Development and Testing

- **Cost:** Purchasing 20 TB of storage to get several image copies of production to test and development is costly
- **Scalability:** Providing 5, 10, 15 or more copies to development and testing teams is unrealistic
- Manageability: managing the cloning process can take 25% or more of a DBA's time



RPO and RTO



Recovery Point – Database-only is typically days or hours of lost data in disaster, or too costly. Combined with storage can be faster, heterogeneous, and with less cost.

Recovery Time – Database-only can be days or hours. Combined with storage can be minutes.

Backup Window – Database-only is typically days for terabyte sizes. Combined with storage can be seconds.



Read-Only Snapshots: Enterprise Backup and Recovery



Simpler Backup and Recovery





Local Backup and Recovery Backup

Copies per Hour

Minutes

Tape 300 GB at 60 GB/hr = 5 hrs (Best Case) 1 1 1 1 1 1 0 1 2 3 4 5 6 Time to Backup (Hours)

Restore

Disk-Based



Restore from Snapshot Copy

Redo Logs (Shorter with More Frequent Backups)

Таре		Resto	re from Tape	Redo Logs	Redo Logs			
	3-0 HOUIS		l	I		I	I	
		0	1	2	3	4	5	6
		Time to	Recovery (H	lours)				



Storage Backup Methods: Snapshots

Instant backup and recovery of large data sets using a sophisticated, scalable, and fail-safe pointer system of storage blocks. Snapshots represent a frozen view of data taken at a specific point-in-time. Data and entire environments can be restored to a known stable point prior to the event that caused the disruption or corruption.



Snapshot.1 file system version of DB.DBF is composed of blocks A, B & C'



Snapshot for Backup/DR validation

- Snapshots
 - Point-in-time copy
 - Created in a few seconds
 - No performance penalty
- TPC-C published with 5 active snapshots





Writeable Snapshots: Cloning for Development and Testing



Using Snapshot Cloning to Accelerate Application Development and Upgrade Cycle



- Quickly reconfigure multiple test, development, QA, DW, auditing, staging environments
- Rapid restore from online
 Snapshot copies
- Allows multiple test
 environments
- Accelerates test cycles
- Helps deliver new applications quickly
- Use database cloning procedures to clone application environments



Before: Large E-Business Suite

Production



Challenges

- Copies consume lots of disk
 - < 10% data differences for each instance
- Copies take a lot of time
 - Slower time to market



After, with Snapshot Clones



Solution

- Instantaneous copies
- Low storage overhead



- Faster TTM
- Higher quality
- Lower cost



Disaster Recovery: Storage-Based Replication



Database Replication Methods: Data Guard

What:

- Replicates databases from one data center to another
- Performs backups from the standby database instead of the production database
- · Some vendors have both physical and logical versions





Oracle Physical Files

All of these should be backed up!

•Datafiles - data, index, temporary, rollback, system

- •Online and archived redo logs
- Control files
- Database vendor executables and patches
- •Exports, backupset files
- •Auditing files
- Database Parameter files
- Database and application alert logs
- Password files
- •Single sign-on files
- •LOB or BFILE or library storage structures
- •External tables
- •Home-grown scripts



Data Guard: replication

PROS:

- Synchronous mode
- Logical mode

CONS:

- Not supported by earlier versions of databases
- Does not replicate source code or non-Oracle systems
- Disables NOLOGGING-type modes
- Usually an extra database license cost
- Standby database must be online



Array-Based Replication

What:

 Replication between two storage arrays that sends storage layer blocks to a standby site whenever there is a storage change at the primary site

PROS:

- Fills the gaps of database replication by replicating non-database systems, source code, and external objects
- More frequent backups with more copies kept online
- Works best in conjunction with database vendor replication solutions by jointly reducing replication traffic and reducing or eliminating single points of failure
- Significantly improves RTO and RPO
- Centralization of remote backup
- Replication to off-site location lowers tape media management costs for offsite tape rotation

<u>CONS:</u>

- Source and target may need to be the same vendor, unless using heterogeneous replication solutions
- There is a cost associated with purchasing and maintaining a third-party array-based replication solution. An ROI analysis will demonstrate that this cost can be justified. The point at which it is recouped will depend on the value of the data



Array-based replication for Databases



- DB volumes on primary and standby site are SnapMirrored
- Remote site's DB is kept in recovery mode
- When disaster strikes primary site, it's a matter of breaking the SnapMirror and bringing up the remote standby site's DB in query/update mode



Cost Savings: Tiered Architecture / ILM



Tiered Architecture / ILM

Database options for tiering data

- Tablespaces on specific tiers of disks
- Table and index partitioning tiers of disk relate to specific partitions
- Triggers, procedures, packages for migrating data
- ILM solutions from third-parties
 - Define the Data Classes
 - Create Storage Tiers for the Data Classes
 - Create Data Access and Migration Policies
 - Define and Enforce Compliance Policies



Cost savings example: 94% reduction in storage costs!

Skip Navigation Print Logout Lifecycle Setup Lifecycle Management Compliance & Security User: ADMIN Reports Help Logical Storage Tiers Lifecycle Definitions Lifecycle Tables Preferences Back to Lifecycle Tables Cost Summary for Managed Table TL.PART_ORDERS Shows the storage costs and savings associated with the current lifecycle table. Lifecycle Definition General Orders Q Without Lifecycle Management Single-Tier With Lifecycle Management Multi-Tier Cost Savings Report Storage Size (Gb) 8,636 Storage Size (Gb) 8,636 Storage Cost \$35,266 Shows the cost if all data was Average Cost Per Gb Cost Per Gb \$72.00 \$4.08 held on the most expensive storage Storage Cost \$621,826 Storage Cost Savings \$586,560 (94%) Compared to cost of distributing across storage Filter Options devices defined for the lifecycle Lifecycle Logical Cost Per Multi-Tier Multi-Tier Multi-Tier Partitions Stages Storage Tier Gb Cost Savings Compressed Size (Gb) Note Compressed Compression is included High Performance \$72.00 6 \$439 0 of 1 0 of 1 because this table is already Low Cost \$14.00 31 \$427 \$1,770 81% 1 of 1 0 of 5 being managed by the ILM Online Archive \$4.00 8,600 \$34,399 \$584,790 94% 0 of 0 0 of 2 Assistant Totals 8.636 \$35.266 \$586.560 1 - 3

User: ADMIN_Language: en-us

ORACLE ILM Assistant



Performance and High Availability: RAID-DP



RAID-DP

- RAID-DP is an advanced, cost-effective disk failure/error protection solution protecting information in the event of a double disk or media failure within a single RAID group
- RAID-DP is based on RAID4 adding a diagonal parity calculation to enhance overall performance versus competitive double parity adaptations based on RAID6





How is RAID-DP Unique?

Protection

- 'Traditional' single-parity-drive RAID group no longer provides enough protection
 - Reasonably-sized RAID groups (e.g. 8 drives) are exposed to data loss during reconstruction
- RAID-DP's double disk-failure protection does what RAID5 and RAID1/0 cannot
 - Reduces RISK: Limits exposure to same RAID group second disk failure or non-recoverable media error

Cost

- RAID 1 is too costly for widespread use
 - Mirroring doubles the cost of storage
 - Not affordable for all data
- RAID-DP exceeds RAID1/0 protection levels without the associated doubling of capacity an cost
- RAID-DP has a comparable operational cost to RAID 4

Performance

- Optimized for performance
- Reduces RAID group rebuild time



RAID-DP vs other RAIDs

RAID

Protects against single disk failure P





- >10,000 times more secure than single-parity RAID
- More reliable than mirroring for double-disk failure
- 13% Parity Overhead vs
 50% Overhead w/ Mirror^(*)
- 75% more usable capacity than competitive offerings^(*)
 - * Comparing 14d + 2p vs. 8+8 mirror



The Cost of Data Availability & Protection

Compare RAID-DP[™] to RAID10

Count the drives needed for 2TB useable storage using 144GB disk drives



A Simpler Database Architecture







RAID-DP RAID 1/0 RAID 5



Aggregates: Improved Performance



Storage Security vs Database Security



Encryption by Storage Vendors

What:

 Unencrypted data opens up the company to several vulnerabilities—whether it's a disgruntled employee or consultant stealing the tape or the tape being lost in the warehouse or while being transported off-site.

In a simplified example of a NAS environment, DataFort is shown on the IP network, between departmental clients and the file server or NAS appliance.



DataFort can also be installed in clusters for high availability.

Encryption by Database

1) Database solutions

Database Encryption:

- Store data encrypted within database tables and, thus, also encrypted on tape
- Data may be unencrypted in indexes, temp tablespace, flashback, audit, and elsewhere

Secure Backup:

- Some database vendors encrypt backups before they are written to tape
- By encrypting at the database level, an organization reduces much of its exposure. This method is also free (for one direct-attached storage device)
- Typically, the DBA is solely responsible for retaining the key and there are inherent weaknesses with this strategy

2) Storage vendor encryption

- Encryption at the storage level
- Multi-master key management that prevents single points-of-failure for key loss
- No auto-destruct capabilities
- By using a hardware and software appliance, backup streams can be encrypted with negligible impact on performance



VTL (Virtual Tape Libraries)



Today's Backup and Restore Issues

"What are the most common causes of a backup failure?" --Percent of All Users (multiple responses accepted), N = 222





Source: Enterprise Strategy Group



Inline Configuration



Backup software manages virtual libraries, VTL console manages the physical library



VTL Features and Differentiators

In addition to meeting capacity and performance requirements, the appropriate solution may consider one or more of these variables:

- Inline or Parallel Configuration
- Software or Appliance-based
- Replication of virtual tapes to DR site
- High Availability (Active-Active)
- De-duplication
- Backup Software Integration / Management
- Physical Tape Encryption
- Tape Stacking / Shredding
- Shadow Images
- NDMP Support
- OEM/Product Maturity and Install Base
- Virtual Devices/Media or File System shares
- AS400 and/or MF compatibility
- Available Budget





De-duplication



Other Backup Methods: Deduplication

<u>What</u>

- Deduplication works with VTLs by detecting redundant data streams during the backup process and sending pointers to that data (versus actual streams of blocks of data)when duplicate streams are detected
- Network dedupe for WAN optimization
- Primary storage dedupe for data "at rest"





Deduplication - Under the Hood

BACKUP DATA	LOGICAL	ESTIMATED REDUCTION	PHYSICAL
FRIDAY FULL	1 TB	2- 4x	250 GB
Monday Incr	100 GB	7-10x	10 GB
Tuesday Incr	100 GB	7-10x	10 GB
Wednesday Incr	100 GB	7-10x	10 GB
Thursday Incr	100 GB	7-10x	10 GB
2 nd FRIDAY FULL	1 TB	50-60x	18 GB
TOTAL	2.4 TB	7.8x	308 GB

- = Unique variable segments (4KB-12KB)
- = Redundant data segments
- = Compressed unique segments





Other Backup Methods: De-duplication

<u>PROS:</u>

- Can dramatically reduce the backup stream size—typically 5-20 times—and bandwidth required for backup
- 10-40% lower acquisition cost than tape
- RMAN: when set up properly, de-duplication can work well with multiple database backup channels

<u>CONS:</u>

Initial costs



Summary, Q & A Ari Kaplan akaplan@datalink.com 312-399-0079



Existing Oracle Backup Methods: Pros and Cons



Oracle Database Architecture





Image from Oracle Corp

Oracle Backup Methods

- Cold / Hot Backups
- Export / Import
- Data Pump
- RMAN





Oracle Backup Methods: Oracle Flashback

- Oracle's flashback area allows for recovering a table (or database to any point-in-time in the past) by storing images of data online
- Flashback Database
 - Recover to point-in-time from Flash Recovery Area
- Flashback Table
 - FLASHBACK TABLE LAOUG_audience, free_gift_list TO TIMESTAMP (06-FEB-2007, 13:25:00);





Oracle Backup Methods: Oracle Flashback

Flashback Drop

 DROP command puts object into a "recycle bin" for quick recovery





Oracle Backup Methods: Oracle Flashback

PROS

- This method provides online backup and recovery, eliminating the need to recover from tape and saving valuable recovery time and management effort
- Extremely fast recovery of tables with simple commands

CONS

- Requires a significant amount of flashback area online, taking up lots of storage
- Can't recover from media failure
- Can't undo operations such as shrink datafile
- Can only flashback to the oldest SCN in the FRA
- Only works with newer releases of Oracle



Database and Virtualization: Enterprise Grid



- Scalable everything: performance, availability, capacity, etc.
- Single view; single point of control

