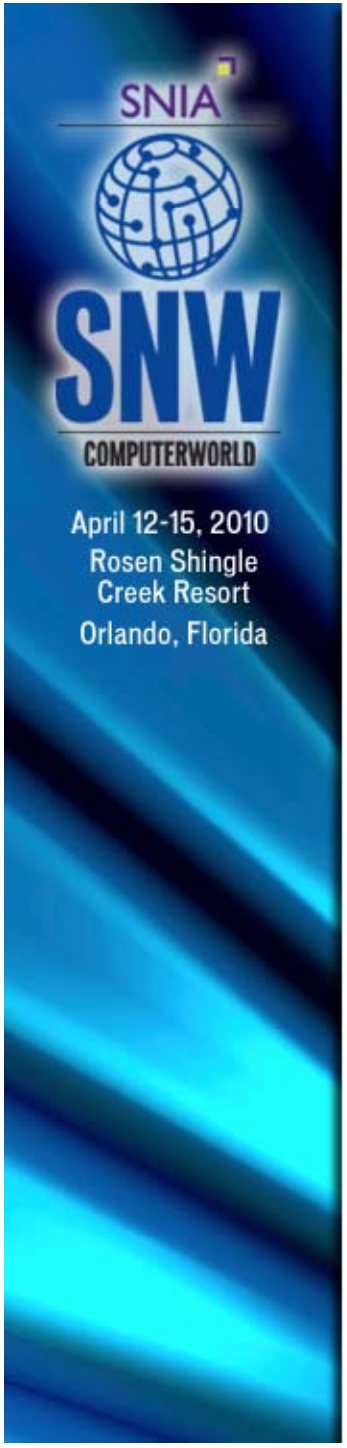

Data Protection
Best Practices for Databases

Ashish Ray

Senior Director of Product Management

Database High Availability

Oracle

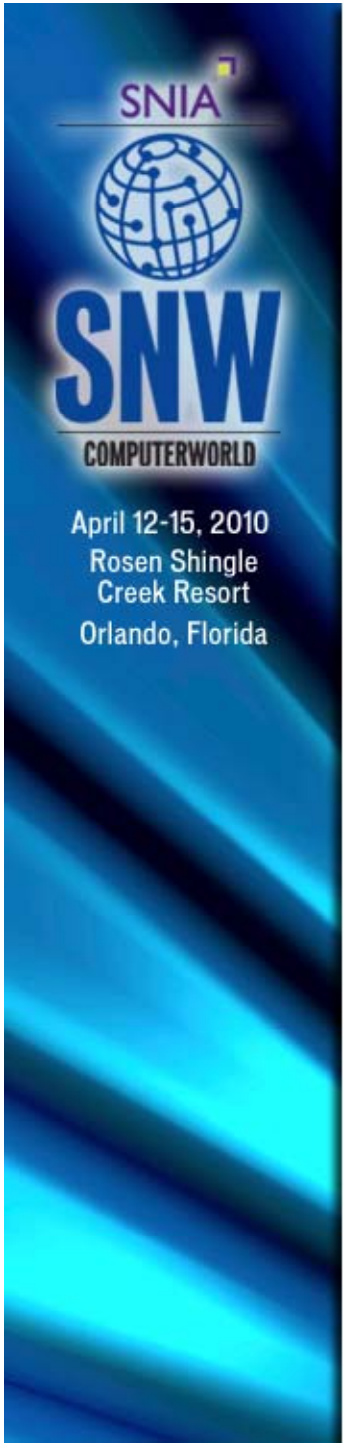

Agenda

- Business Problem
- Overview of Data Protection Solutions
  - Storage-centric Solutions
  - Database-integrated Solutions
- Evaluation Framework & Summary

Business Problem

Reduce Run-time Costs
- Utilize all systems resources – no idle components
- Keep it simple, automated, integrated

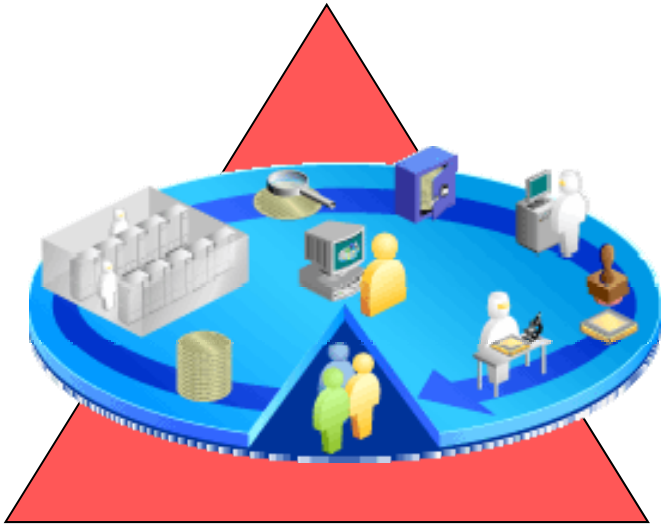

Data Protection – RPO
- Minimize data loss after outage
- Isolate data faults so that valid data can still be accessed

High Availability – RTO
- Minimize unplanned & planned downtime
- Ensure full-stack application availability

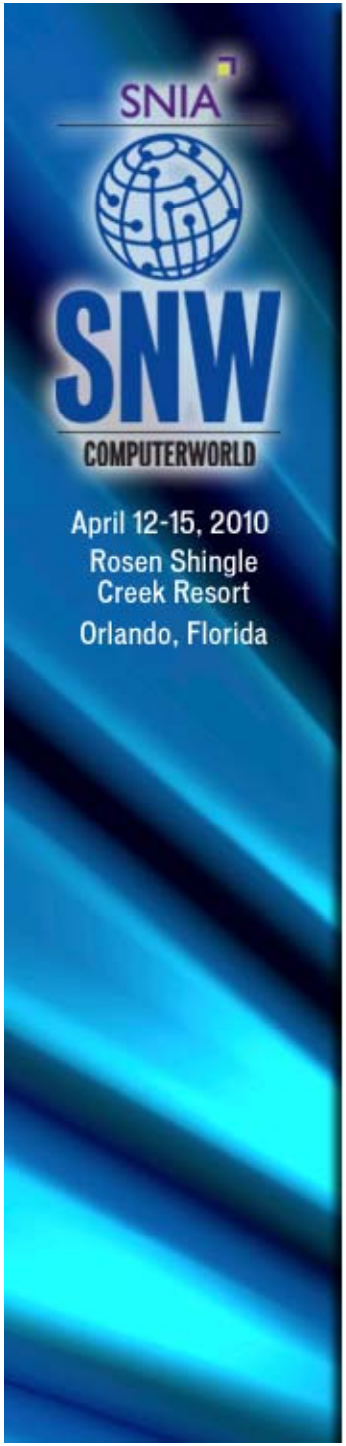

What is a Disaster?

- Headline-grabbing catastrophic events
  - Fire, earthquake, tsunami, flood, hurricane, …
- And . . . more mundane and frequent events
  - Faulty system components – server, network, storage, software, …
  - Data corruptions
  - Backup/recovery of bad data
  - Wrong batch job
  - Bad hardware & software installations / upgrades / patching
  - Operator errors
  - Power outages

  …

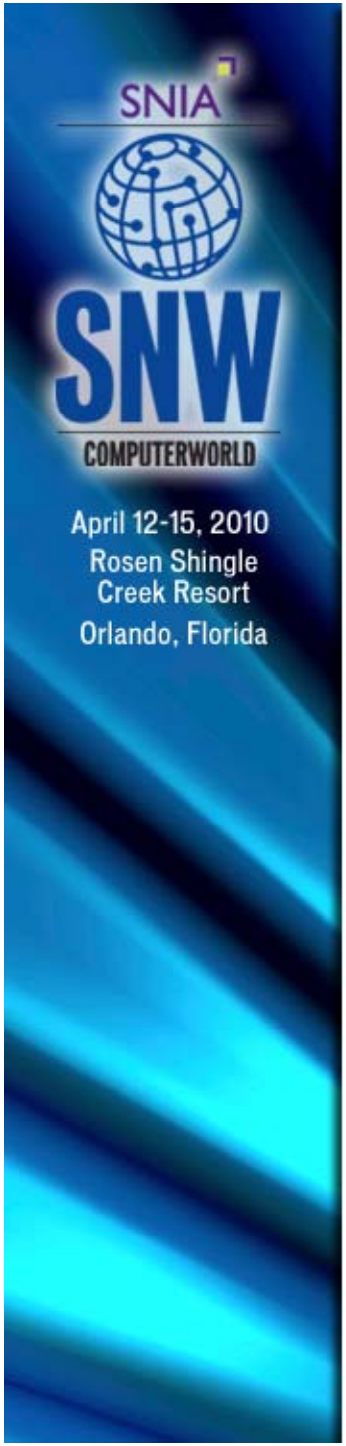

Murphy's Law?

From a recent email

Subject: Data Guard ..

Date: Fri, 09 Oct 2009 17:42:44

XYZ is still trying to catch up after their data center got hit by lightning. [...]

From a Customer Report with Oracle Support

- **n node RAC … disks failed and no longer available … Backup was running when disks died.**
- **Database has only been backed up ONCE many years ago … Customer's attempt at a backup strategy. Unfortunately disks and database died.**

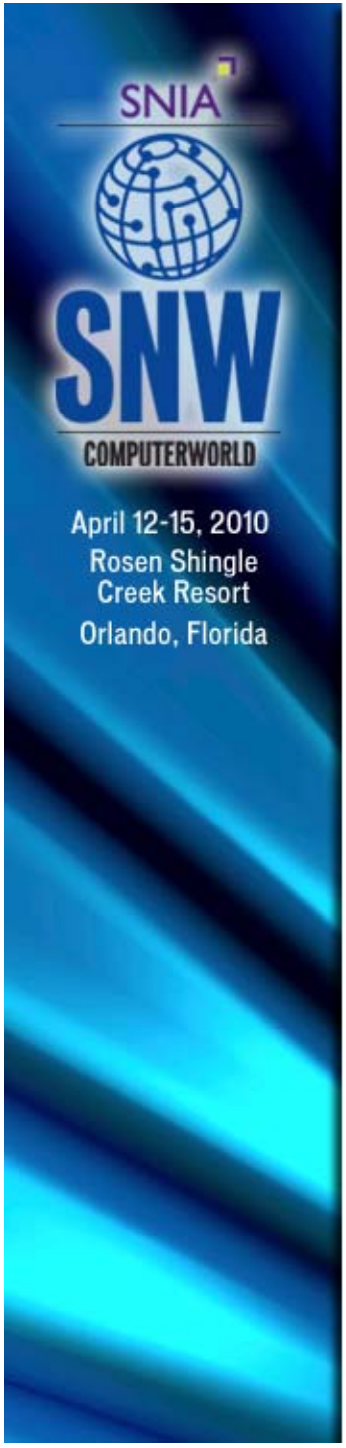

Real-life Disaster

Financial Services Company

- Errors observed in the alert.log of the production database:

```
Errors in file /opt/app/oracle/admin/dg/bdump/dg1.trc:
ORA-01186 : file 93 failed verification tests
ORA-01122 : database file 93 failed verification check
ORA-01110 : data file 93: '/dbmnt/db01/oradata/dg/arch05.dg'
ORA-01251 : Unknown File Header Version read for file number 93
```

```
ORA-01251 - Corrupted file header. This could be caused due to
missed read or write or hardware problem or process external
to oracle overwriting the information in file header.
```

- Affected: primary customer facing applications for trade transaction confirmation, new accounts, and customer account information

Data Corruptions

Severe Impact on Data Availability

- Any component in the systems stack can fail and cause data corruptions*
  - Software – applications, middleware, database, …
  - Hardware – disk drives, disk controllers, HBAs, memory, …
  - Network – routers, switches, cables, …
  - Operational – human errors, bad installs & upgrades, …
- Data corruptions can be disastrous
- Very hard to debug and diagnose

**** "Hard Disk Drives – the Good, the Bad & the Ugly", ACM Queue, Sep/Oct 2007, http://queue.acm.org/detail.cfm?id=1317403***

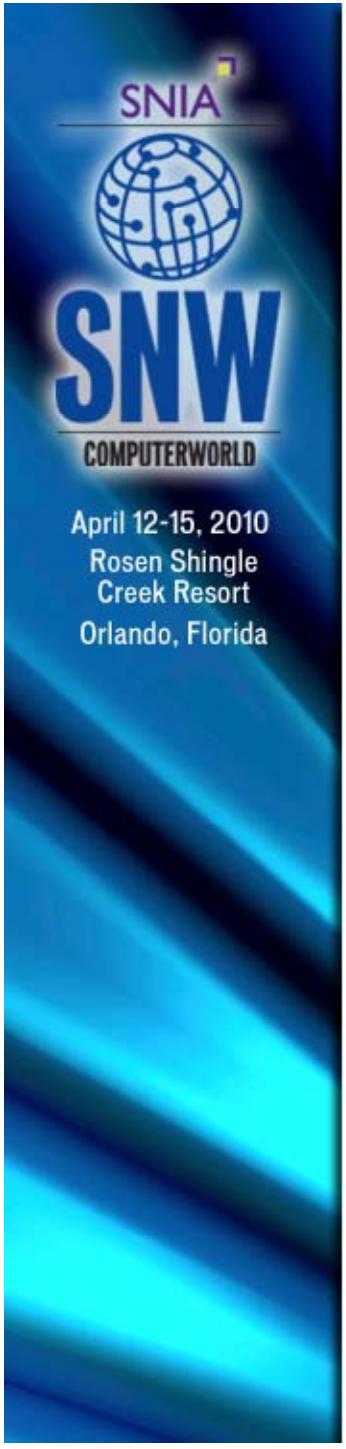

Agenda

- Business Problem
- Overview of Data Protection Solutions
  - Storage-centric Solutions
  - Database-integrated Solutions
- Evaluation Framework & Summary

Data Protection Solution Spectrum

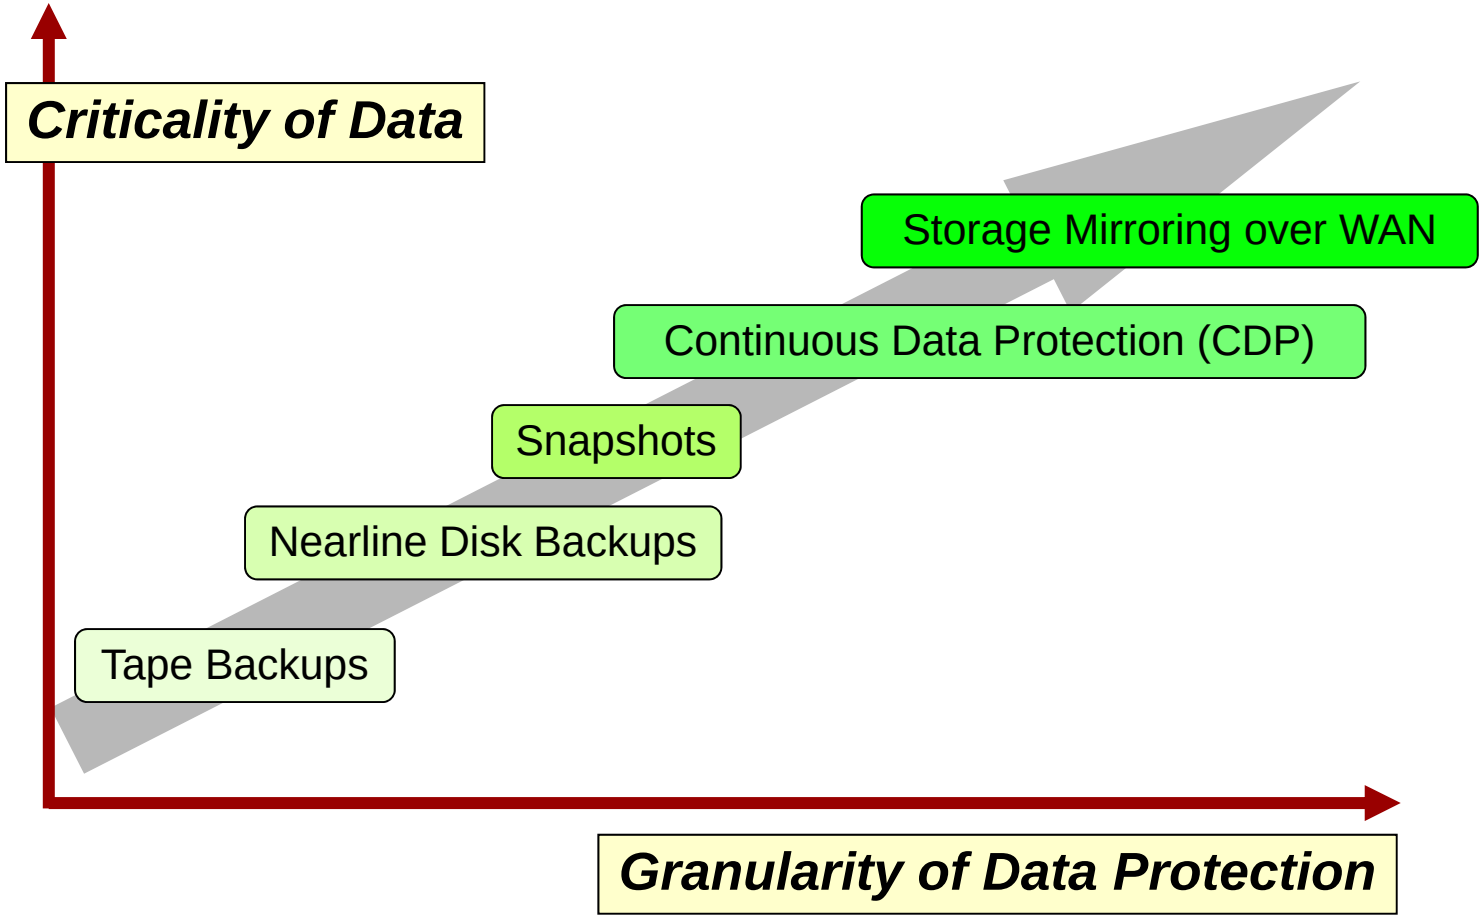

Data Protection at Which Layer?

Simplicity vs. Value Trade-Offs

Application Layer

- Data: Business applications
- Application failover
- Consistency across multiple databases
- Solution profile: logical

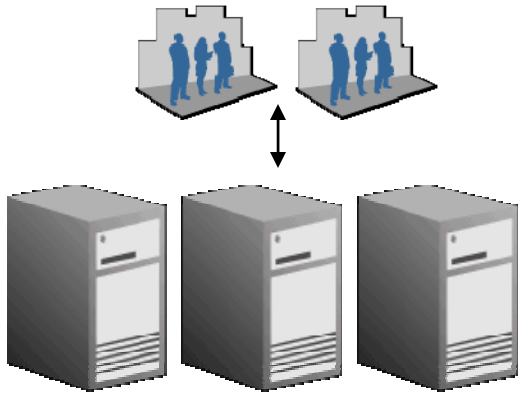

Application-level sophistication

Database Layer

- Data: Tables, transactions, schemas
- Database failover
- Consistency across multiple storage arrays
- Solution profile: logical + physical

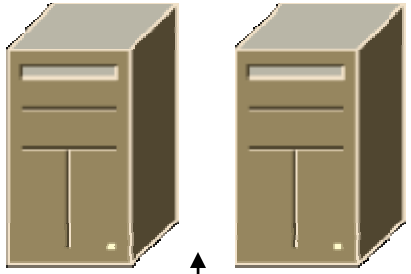

Balance between optimization and performance

Storage Layer

- Data: Bits, bytes
- Storage array takeover
- Consistency within storage array
- Solution profile: physical

Bit-level simplicity

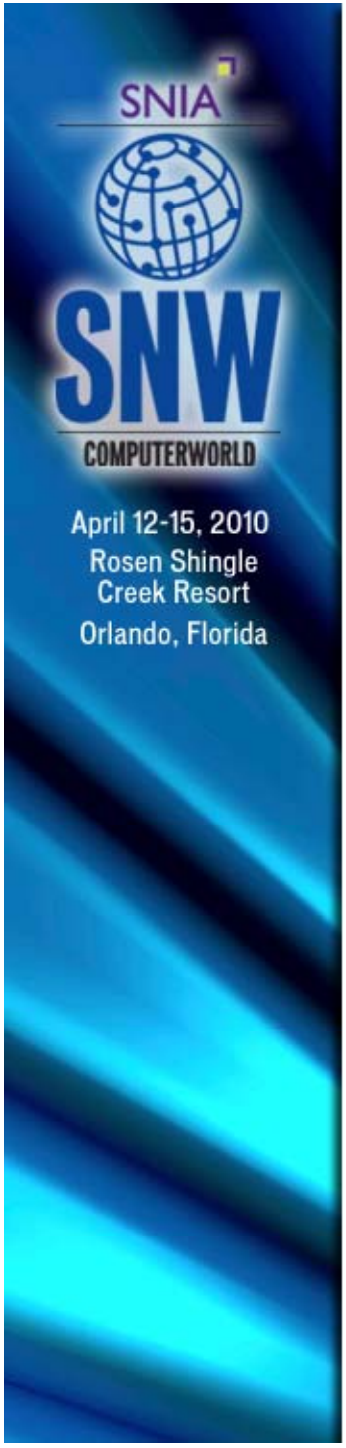

Agenda

- Business Problem
- Overview of Data Protection Solutions
  - Storage-centric Solutions
  - Database-integrated Solutions
- Evaluation Framework & Summary

Backup & Recovery

Data Protection 101

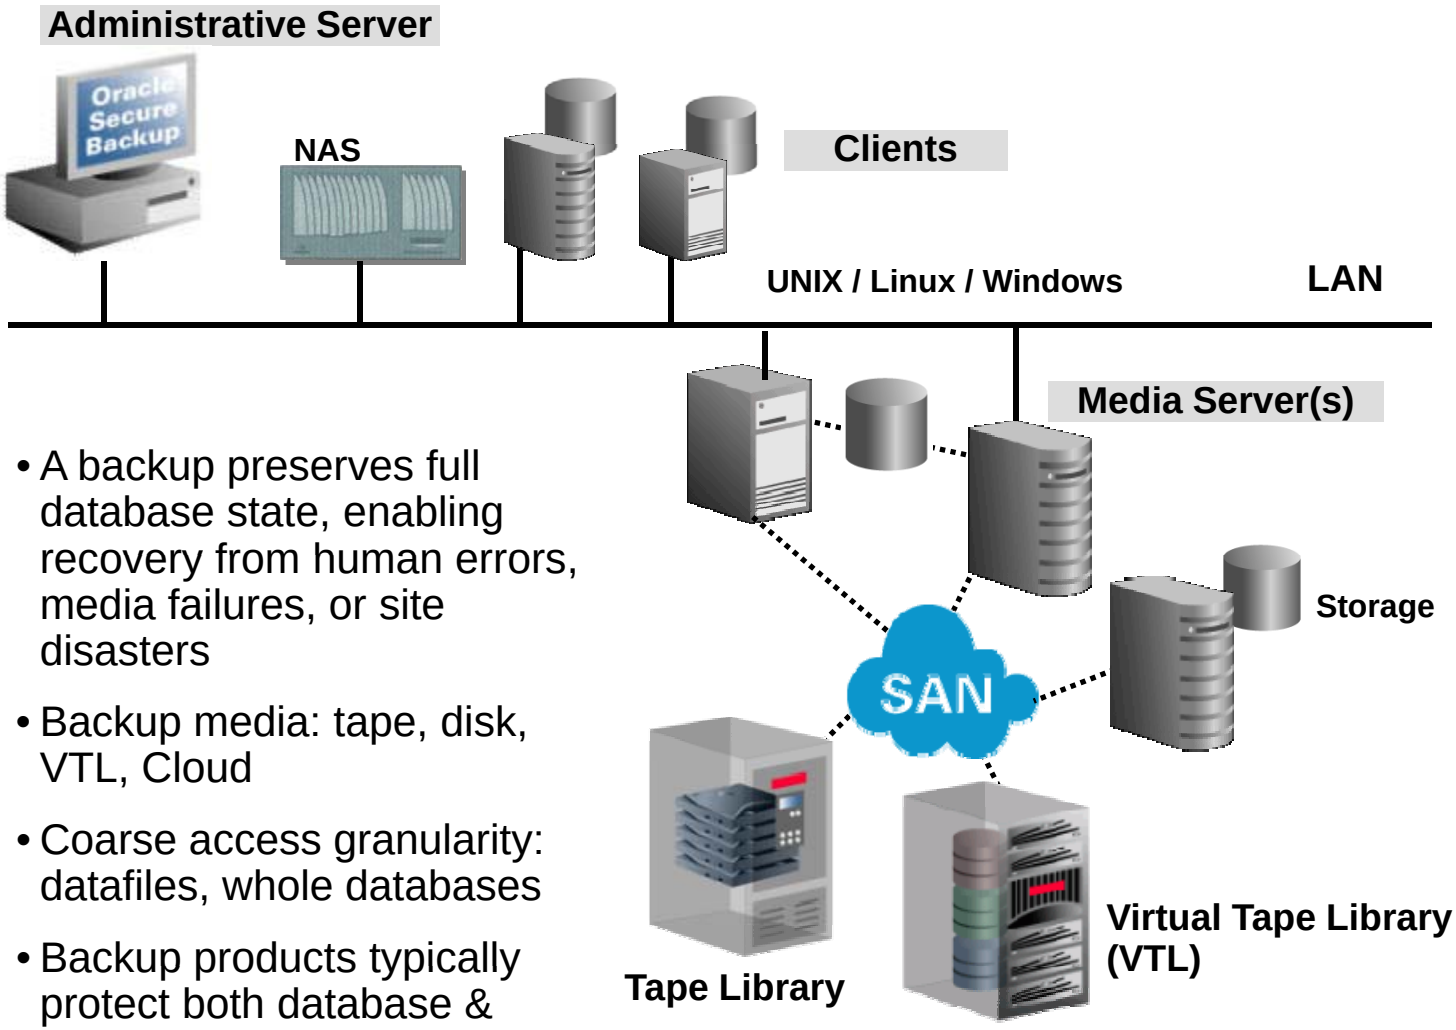

- A backup preserves full database state, enabling recovery from human errors, media failures, or site disasters
- Backup media: tape, disk, VTL, Cloud
- Coarse access granularity: datafiles, whole databases
- Backup products typically protect both database & non-database data

VTL & Deduplication

Optimize Backup & Recovery

- VTL (Virtual Tape Library): Software or Appliance that makes a disk array emulate a tape library
  - Improves backup & recovery performance
  - Minimal disruption to existing tape backup infrastructure
- Deduplication: Replace redundant data with pointers to shared copy
  - Lowers storage costs by reducing capacity requirements
  - Can be done inline or post-process
  - Mileage varies for deduplication of database blocks

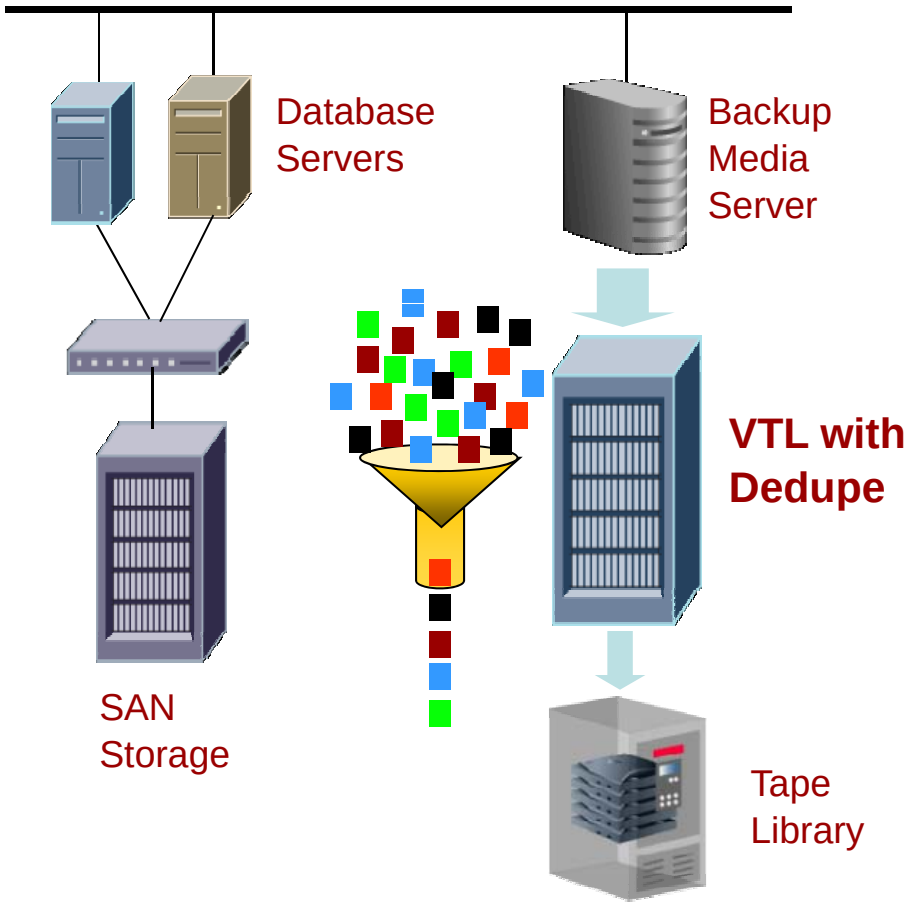

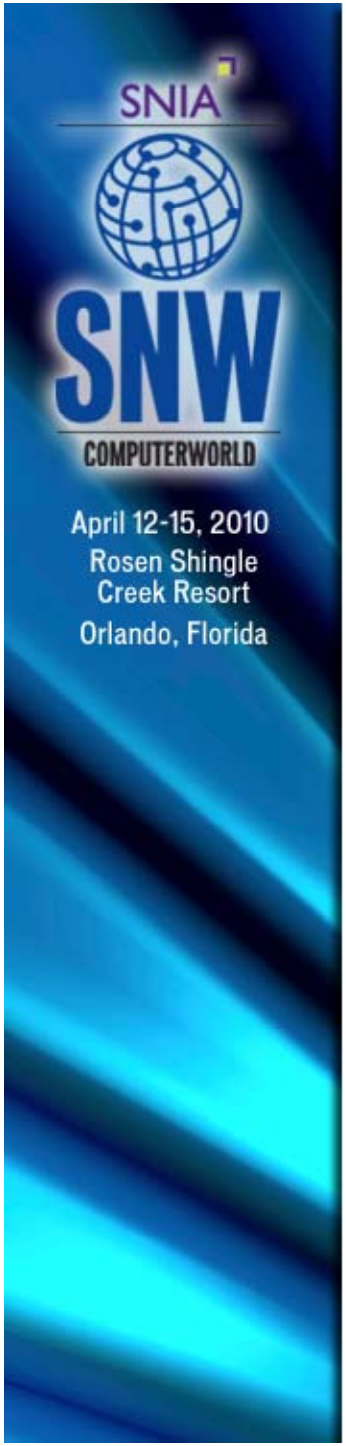

Storage-level Snapshots

Application-agnostic Storage-state Undo

- Preserves disk state at specific time
  - Upon later error, system can be set back to that point
- Quiesce application, copy metadata state, then begin “branching” writes to snapshot storage
  - Branch via copy-on-write, redirect-on-write, split-mirror
  - Extra write overhead often absorbed by storage hardware
- Low space and setup cost
  - Easy to take many snapshots
  - Easy to restore to point prior but near error
  - Split-mirrors protect from media failures, but increase space costs

Storage-level Snapshots

Example: Copy-on-Write

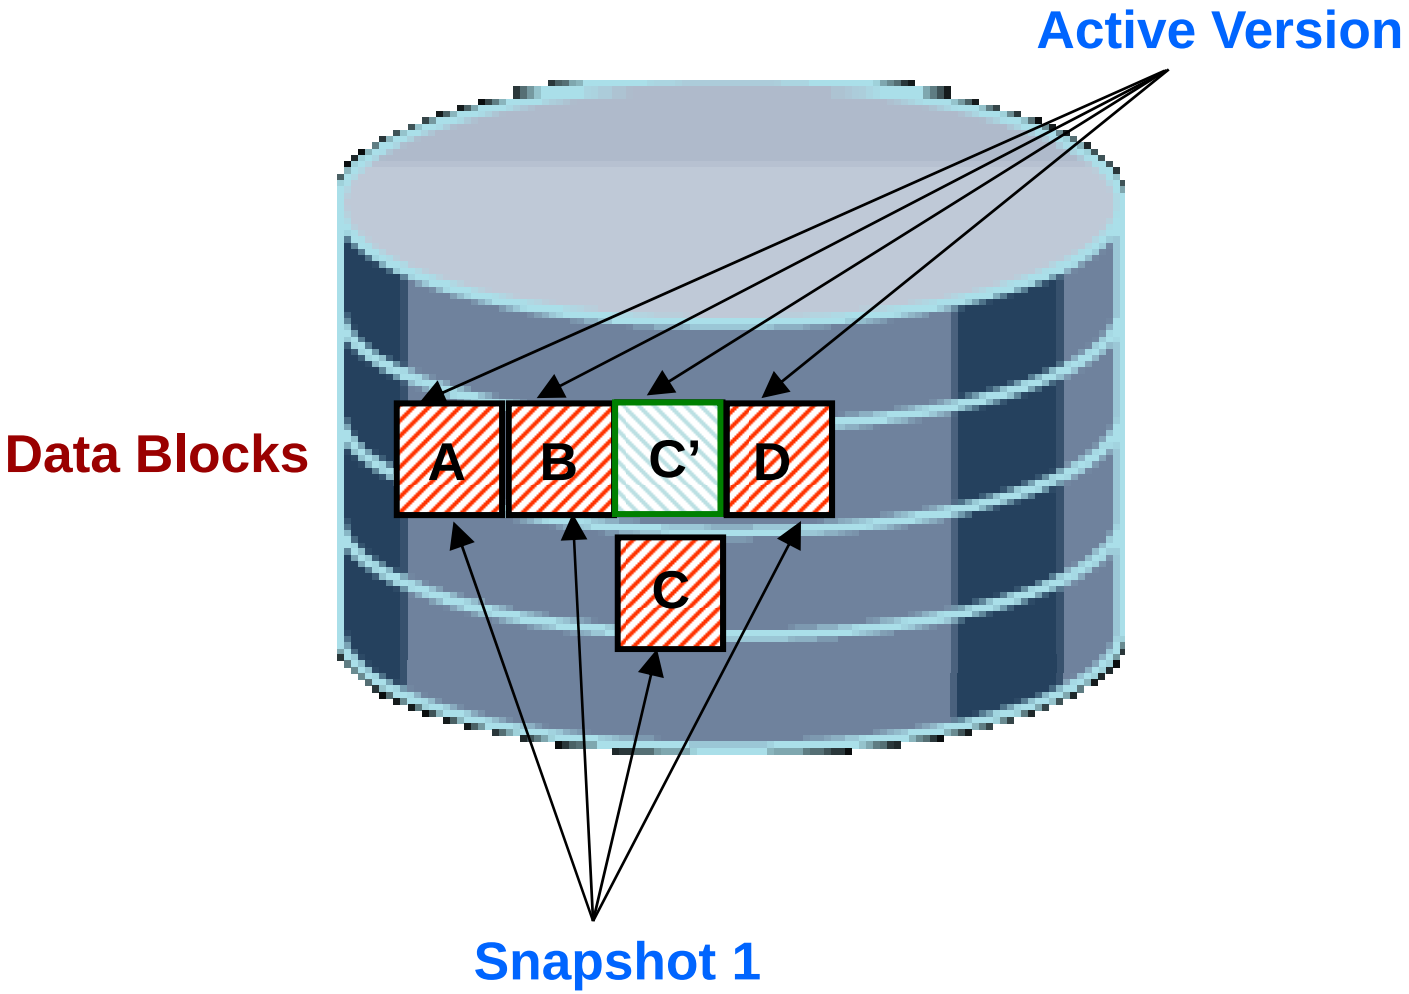

Storage-level Snapshots

Example: Redirect-on-Write

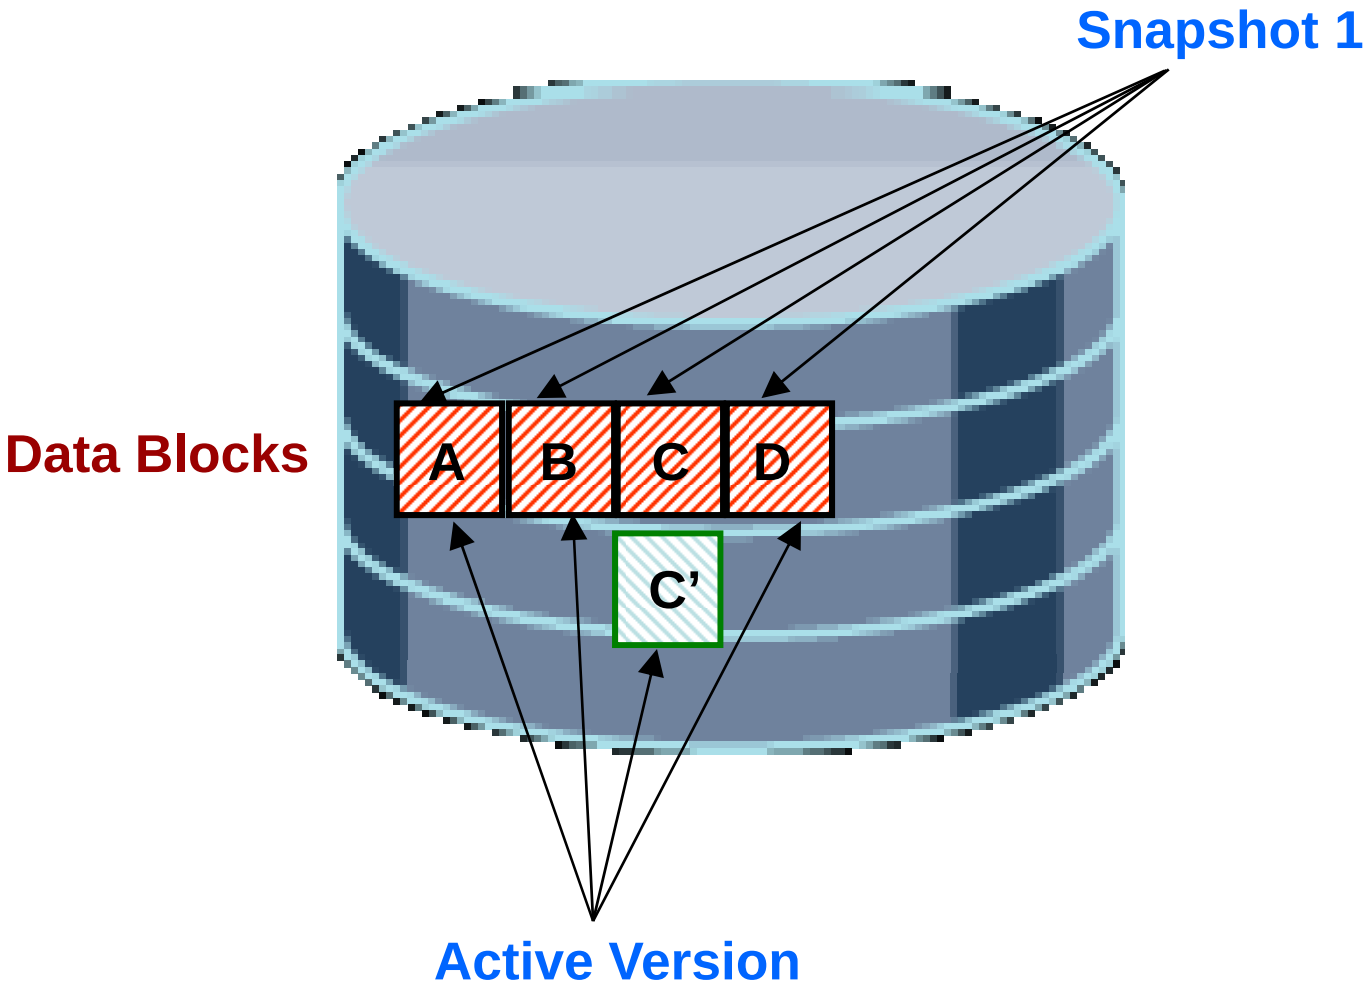

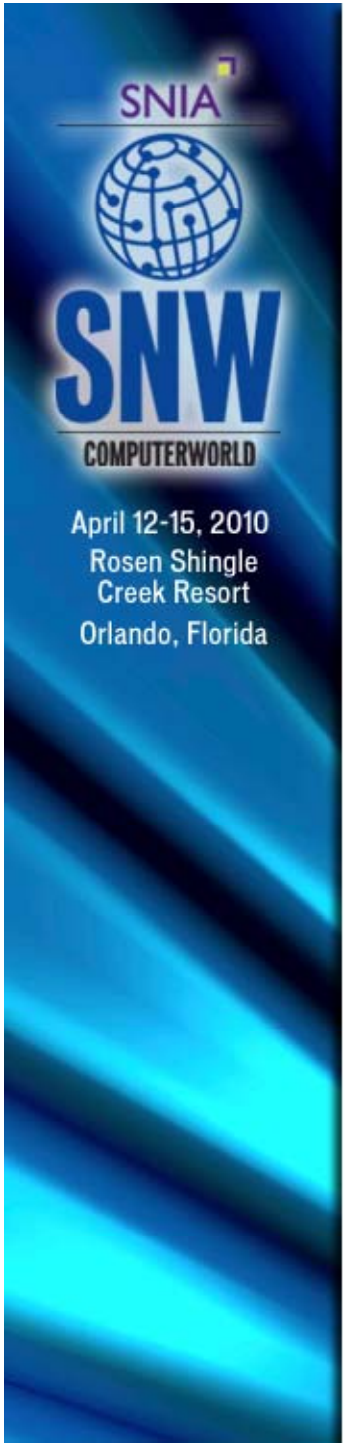

Continuous Data Protection (CDP)

Create a Separate Copy of Every Disk Write

- Split writes: each time data are written to disk, a copy is sent to a separate location (asynchronously)
  - Similar to snapshots, but every disk update is captured
  - Application agnostic, but at recovery time it is complicated to choose the consistent point to go back to
  - CDP may be space-efficient by saving byte- or block-level differences rather than entire write
  - High-frequency snapshots sometimes considered CDP
- Write may be split at host (driver) or at switch
  - Unsynchronized CDP of 2+ filesystems may cause corruption, so CDP infrastructure failures require careful management
- Works best for simple filesystems and manual recovery

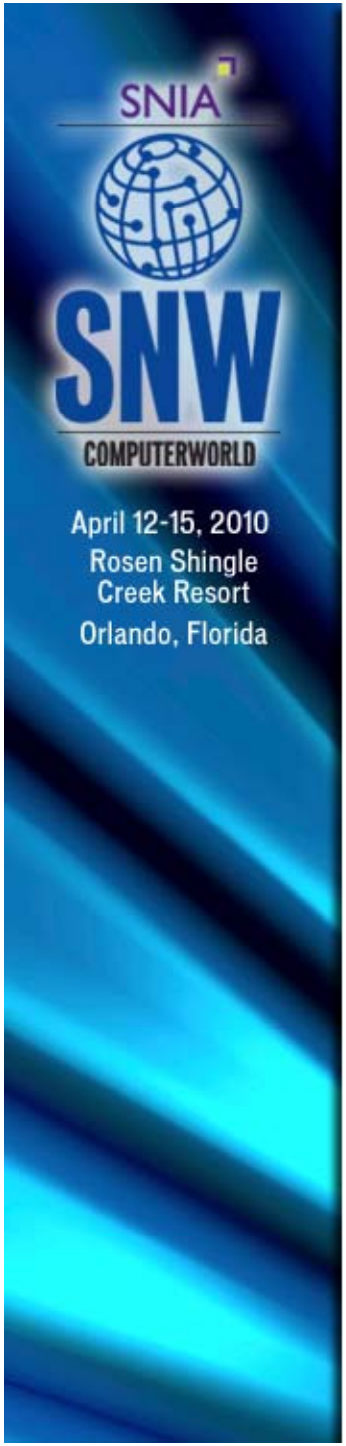

Storage Mirroring

Block-level Disaster Recovery

- Storage array controllers at the primary site send changes to a similar storage array (mirror) at the secondary site
  - As I/Os occur at the primary server, data is written to the cache of the source array, and placed in a queue
  - The link adapter dequeues and moves data to the mirrored array
  - Supports synchronous/asynchronous writes
  - Protocols supported: ESCON, FICON, Fibre Channel, IP – controlled by specialized adapters loaded with appropriate microcode
- Target arrays unavailable for data access

Storage Mirroring Configuration

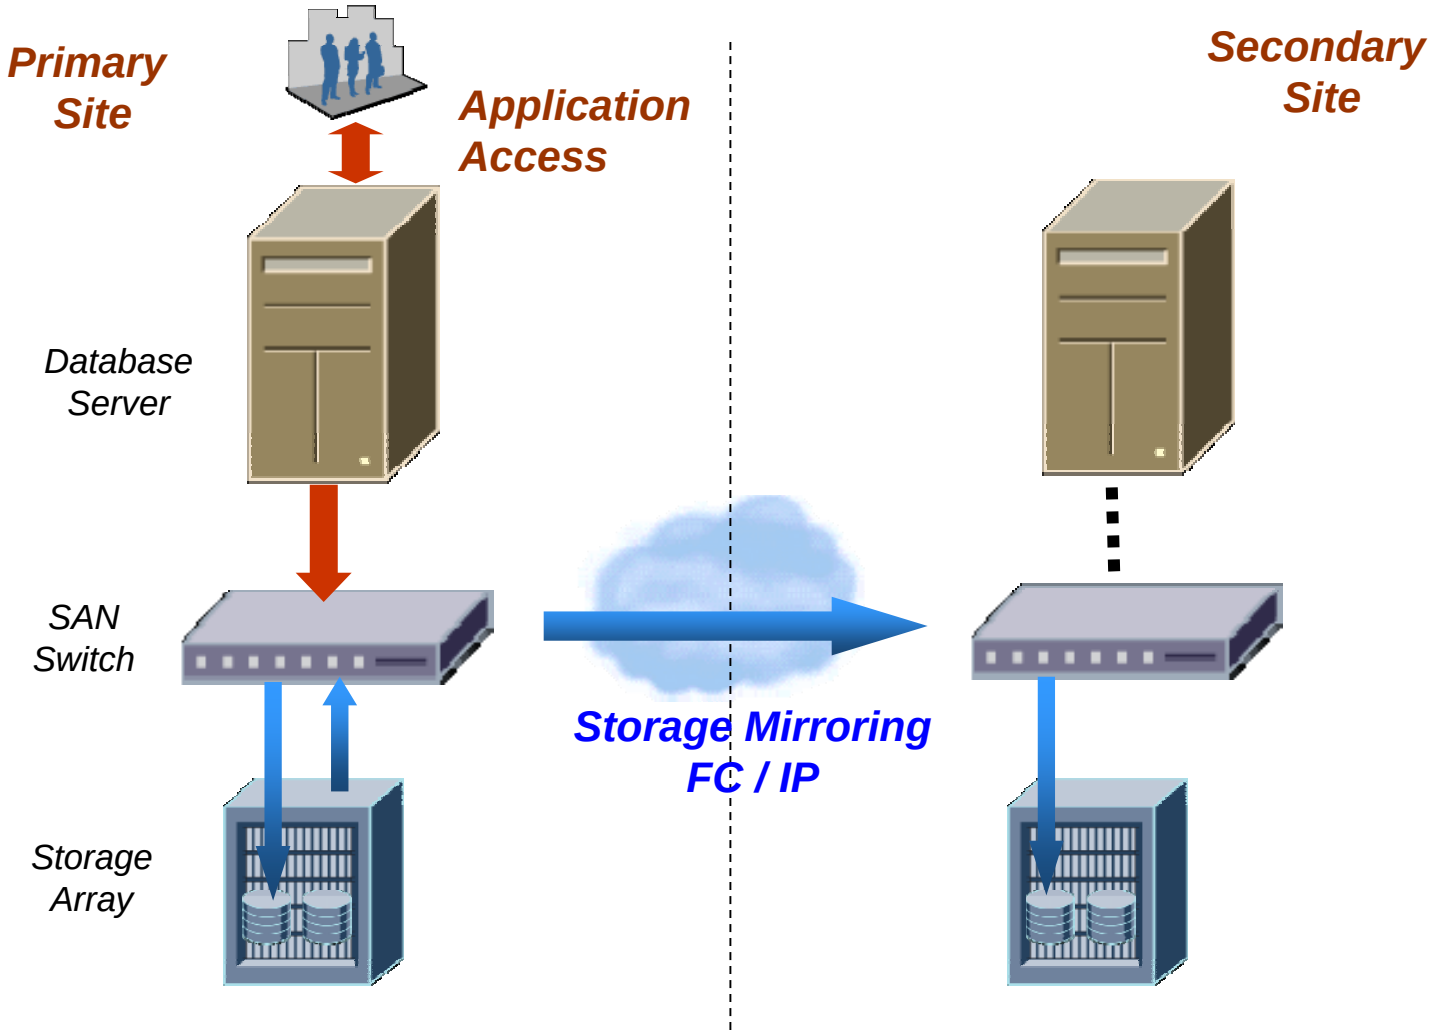

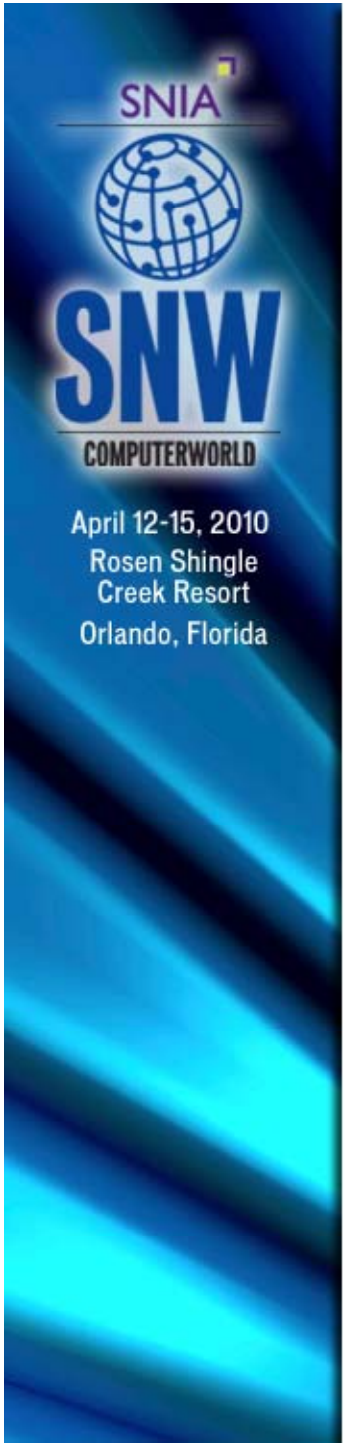

Agenda

- Business Problem
- Overview of Data Protection Solutions
  - Storage-centric Solutions
  - Database-integrated Solutions
- Evaluation Framework
- Summary

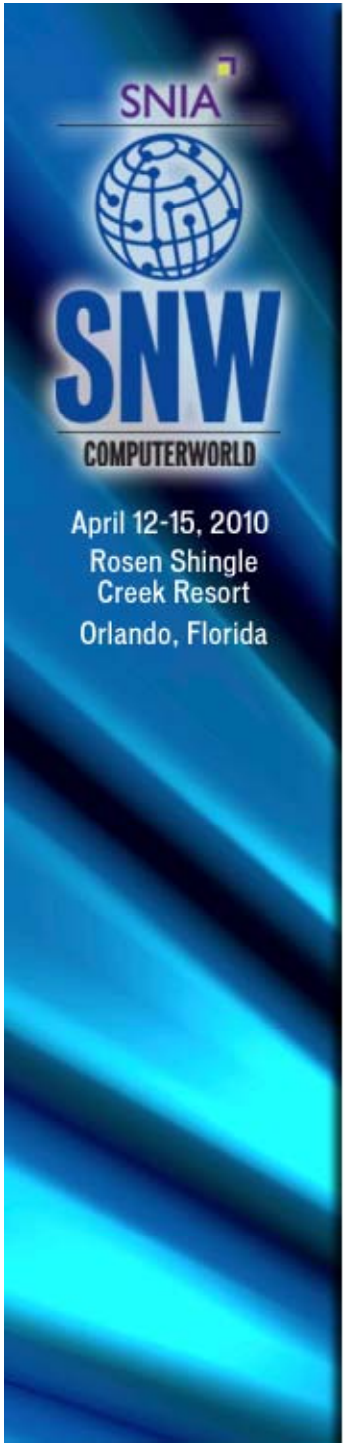

Database Integrated Data Protection

- Popular relational databases (RDBMS) have various levels of data protection capabilities available with the database kernel
  - E.g. RDBMS: Oracle, DB2, SQL Server, MySQL, …
  - E.g. Capabilities: Backup & Recovery, CDP, Deduplication, Mirroring, …
- Data Protection strategy is a bit different for "cloud-based databases"
  - E.g. BigTable, Cassandra, SimpleDB, mongoDB, …
  - Much of the infrastructure heavy-lifting shifted to software layer on top
    - Driven by unique scalability and data model demands

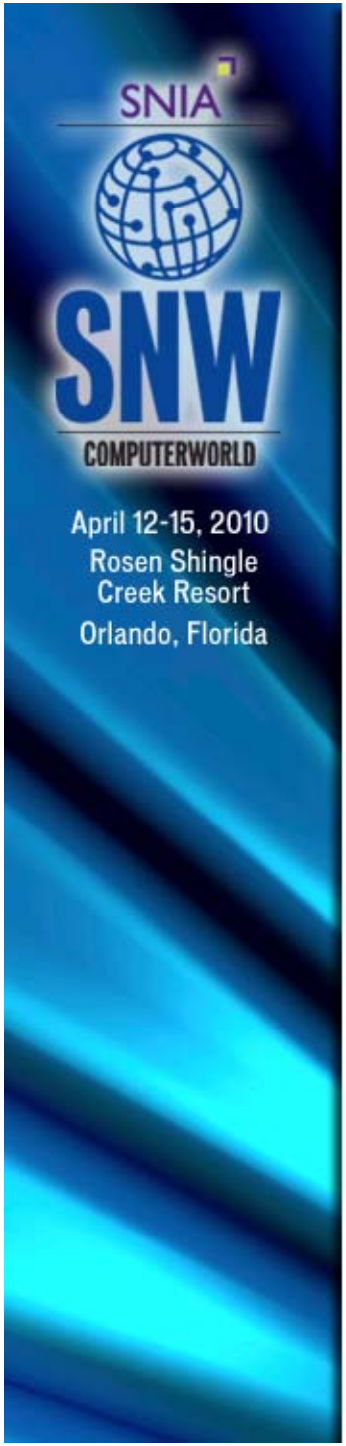

RDBMS Usage Among Audience

- Show of hands – major database deployment within your organization:
  - Oracle
  - SQL Server
  - DB2
  - MySQL
  - Informix
  - Sybase
  - Others:

Database Integrated Backup & Recovery

Online, Granular Operations

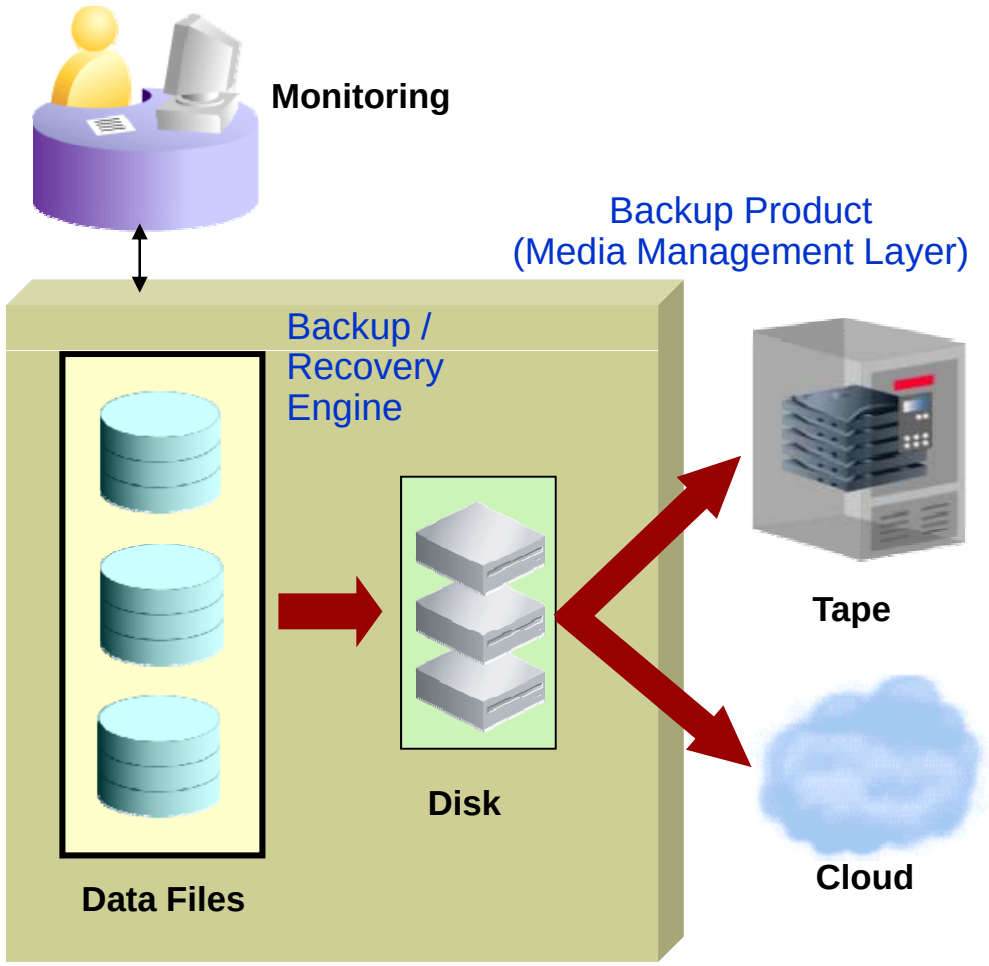

- Most databases provide an integrated backup & recovery engine
  - Online backup
  - Block validation
  - Online block recovery
  - Native encryption
  - Block compression
  - Incremental block-level backup (implicit deduplication)
- Enables integrated backup technologies:
  - Disk
  - Tape
  - Cloud
- Near-line database-optimized compression
  - No de-compress needed
  - Improves performance

Database Integrated CDP

With Consistency

- Some databases have integrated CDP mechanisms
  - View ‘good’ data as of a point-in-time
  - Track changes on disk in a database-optimized manner
  - Granularity: table, transaction, database
  - Simply rewind data changes & keep database transactionally consistent
  - Independent of database size, dependent on extent of changes to unwind
  - Provided through extension of DDLs and/or standard management interface

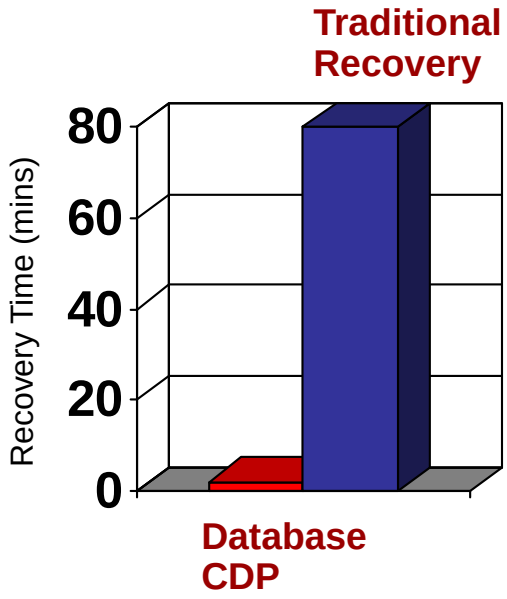

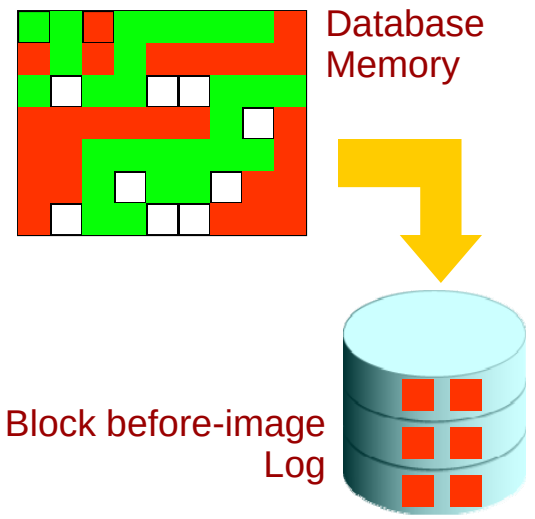

Database Integrated Disaster Recovery

Leverage DR Investment

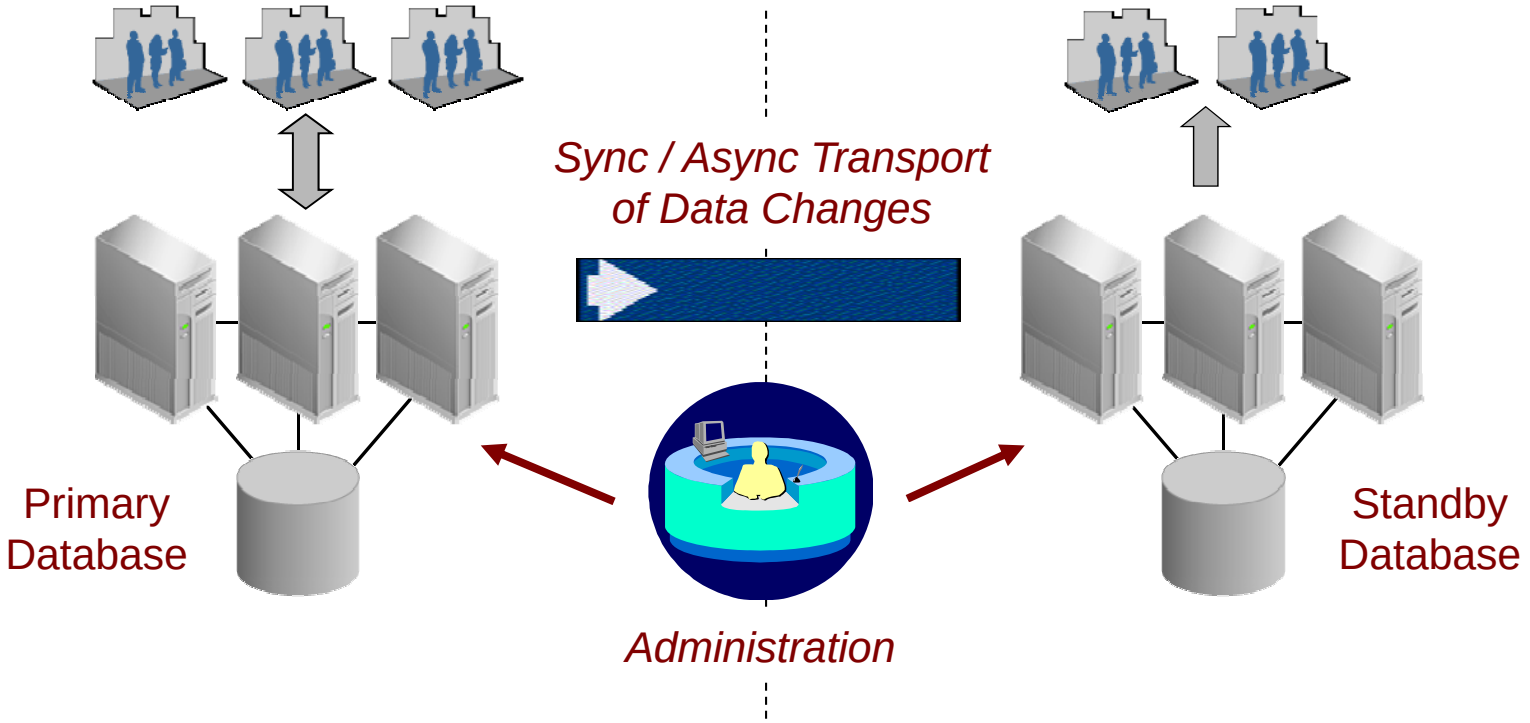

- Significantly enhanced from early-generation log shipping technologies
- Capability to leverage standby database for some processing (queries, backups)
- Transactional consistency maintained throughout

Dealing with Data Corruptions

- Remember the data corruption “disaster”?

```
ORA-01251 - Corrupted file header. This could be caused due to
missed read or write or hardware problem or process external to
oracle overwriting the information in file header.
```

- Checksums maintained at database block level can detect corruptions end-to-end as the block traverses the I/O path

  - Memory
  - OS
  - Device Driver
  - HBA / NIC
  - Fibre Channel / TCP/IP
  - SAN / NAS
  - Disk Controller
  - Firmware

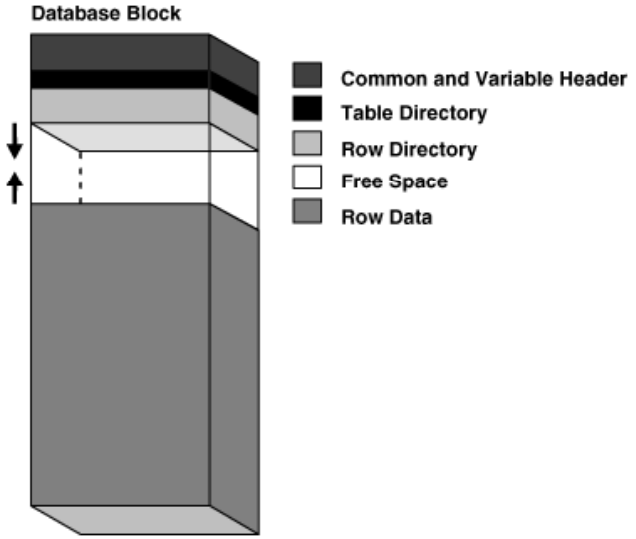

Ref. “Silent Corruptions”, CERN, http://fuji.web.cern.ch/fuji/talk/2007/kelemen-2007-C5-Silent_Corruptions.pdf

Ref. “Hard Disk Drives – the Good, the Bad & the Ugly”, ACM Queue, Sep/Oct 2007, http://queue.acm.org/detail.cfm?id=1317403

Database Integrated Corruption Check

End-to-End Validation

- Some databases have integrated checks to detect and repair corruptions
  - Detects corruptions in data and redo log blocks using checksum validation
  - Detects data block corruptions using semantic checks
  - Detects writes acknowledged, but actually lost by the I/O subsystem
- Various levels of checks can be configured by the administrator
  - Choose the desired protection level
  - Can be configured for data blocks / data + index blocks
- Specific technologies provide additional validation
  - Validate blocks while doing backup & recovery
  - Validate blocks using mirrored copies
  - Validate blocks while synchronizing standby database

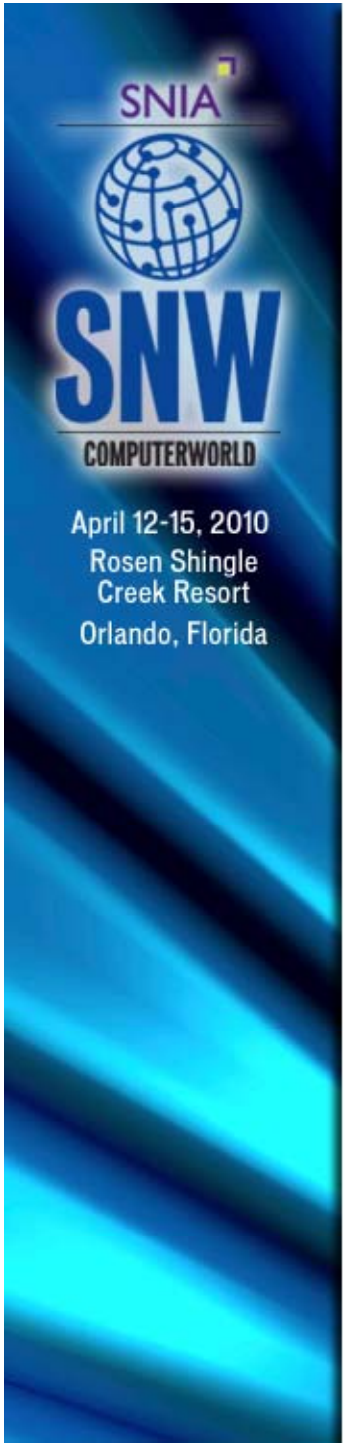

Agenda

- Business Problem
- Overview of Data Protection Solutions
  - Storage-centric Solutions
  - Database-integrated Solutions
- Evaluation Framework & Summary

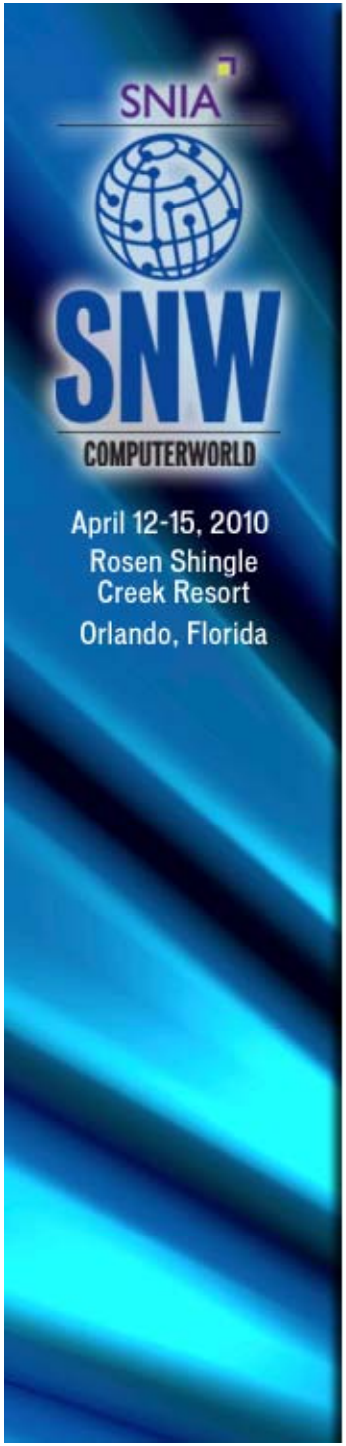

Evaluation Framework

- Measure effectiveness across following criteria:
  1. Protection from various component failures
  2. Minimizing application downtime
  3. Utilization of all resources
  4. Reduction of runtime costs
  5. Support of technology mix

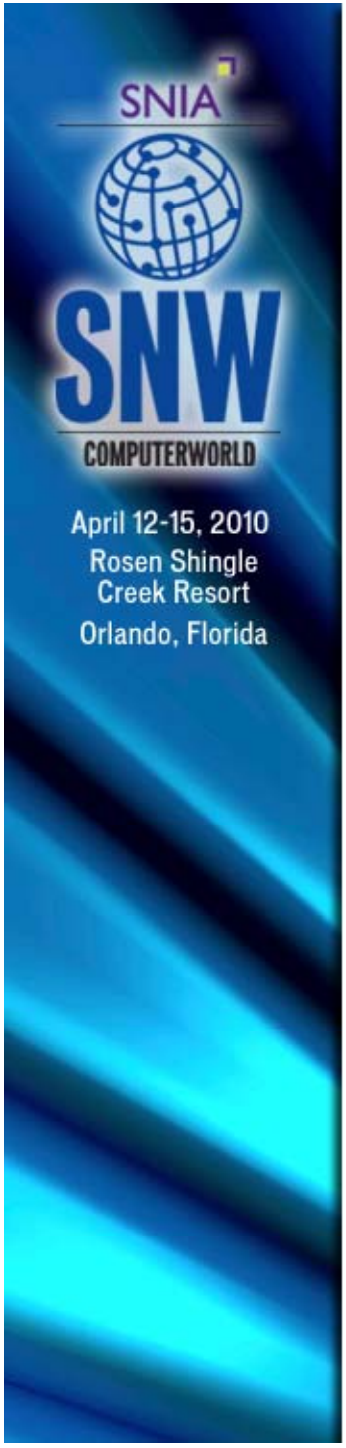

Handling Failures

Attributes	Storage-Centric Protection	Database-Integrated Protection
Server failures	Not Applicable	Excellent (through built-in Clustering)
Storage failures	Excellent (RAID, Mirroring)	Excellent (integrated volume management, replication)
Site failures	Excellent (remote mirroring)	Excellent (replication)
Data corruptions	Limited (within storage array, no detection of lost writes)	Excellent (end-to-end corruption detection)
Logical / transaction failures / human errors	Limited (through storage snapshots)	Excellent (through granular recovery of business objects, e.g. tables, transactions)

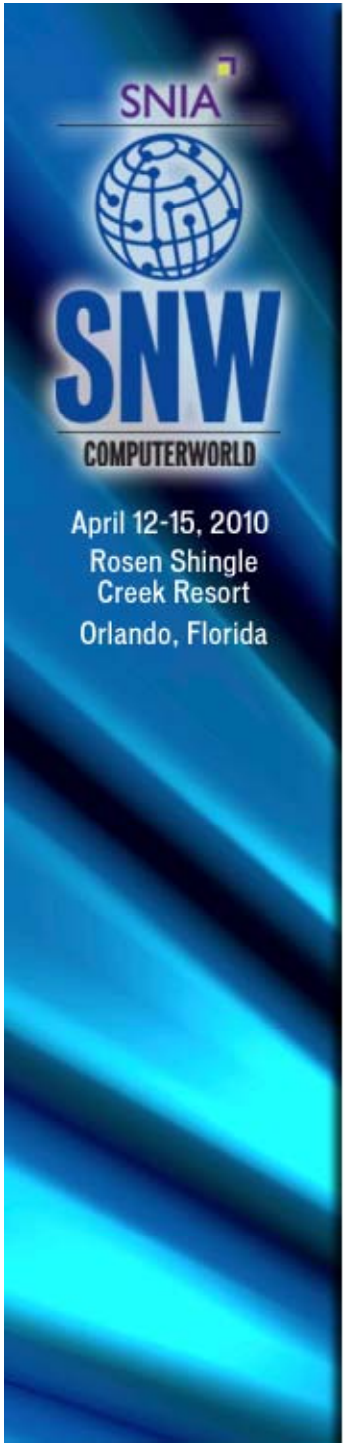

Minimizing Application Downtime

Attributes	Storage-Centric Protection	Database-Integrated Protection
Application-integrated failover	Missing	Very Good (through application notification mechanisms available via mid-tier client libraries such as JDBC)
Transactional consistency	Limited (e.g. database crash recovery needs to be run after storage array takeover)	Excellent (transactional consistency always maintained)
Consistency between database and non-database data	Excellent (through “consistency groups”)	Limited (data protection technologies within the database typically limited to database data)

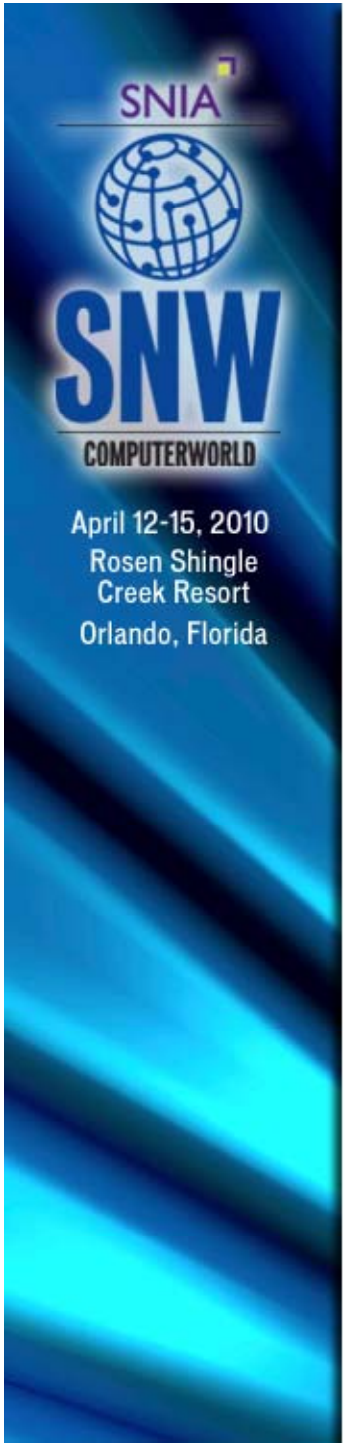

Utilizing Available Resources

Attributes	Storage-Centric Protection	Database-Integrated Protection
Local storage	Excellent (different RAID levels, compression, deduplication)	Excellent (integrated volume management, database optimized compression, deduplication)
Snapshots and clones	Excellent (high-performance read-only snapshots, writable snapshots, minimal storage requirements)	Good (incremental backups, standby databases, read-only or read-write “views”)
Disaster site	Limited (mirrored volumes are offline, clones can be created but require additional storage)	Excellent (database replication enables standby database to be accessed for query offloads, backups)

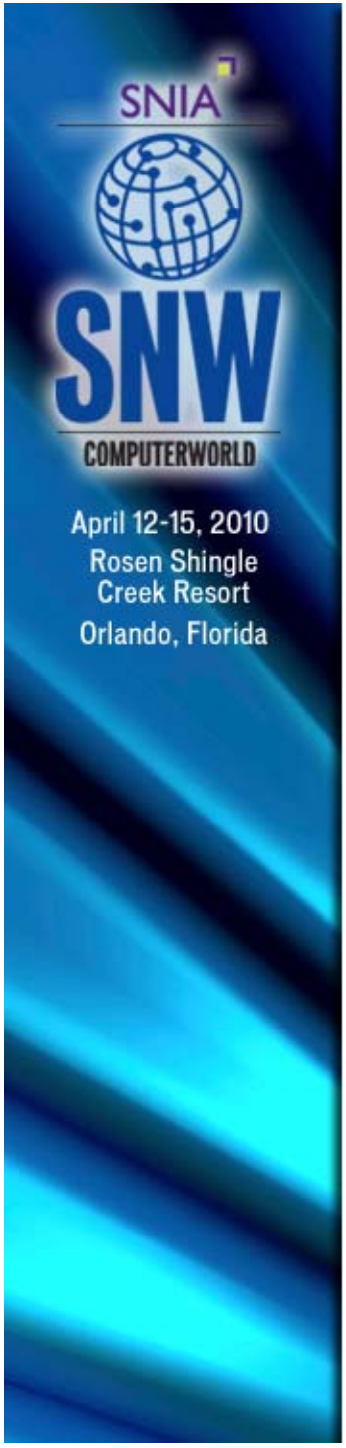

Reducing Runtime Costs

Attributes	Storage-Centric Protection	Database-Integrated Protection
Managing backup & recovery	Excellent (based on snapshot-based backup and recovery)	Good (requires space management of backups and logfiles)
Managing disaster recovery	Good (simple commands available to invoke takeover but app integration lacking)	Excellent (extension of DDLs and standard database management interface, very good app integration)
Automation	Very Good (automated snapshot schedules, reliance on scripts for automated disaster recovery)	Excellent (automated backup space management, automatic failover for disaster recovery)
Monitoring and Management	Very Good (additional CLI / GUI-based management interfaces)	Excellent (extension of standard database management interfaces)

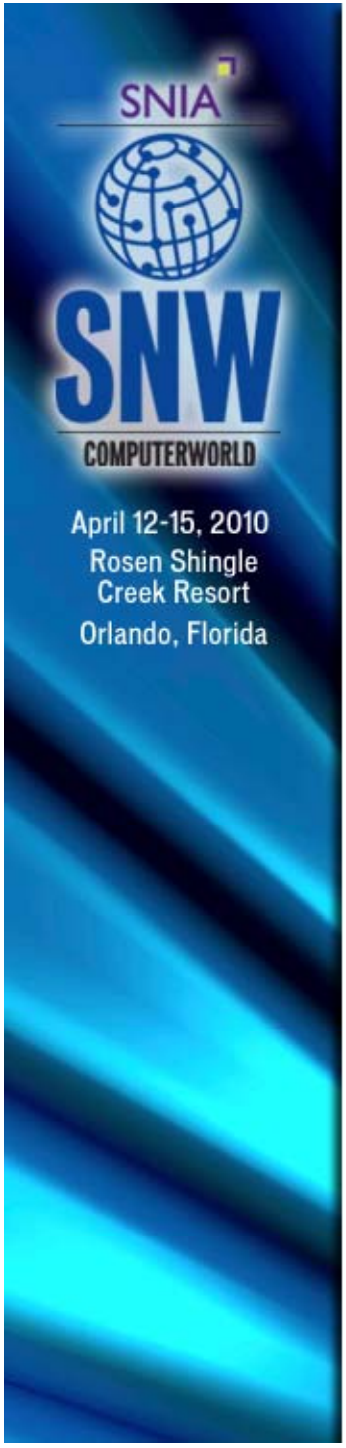

Supporting Technology Mix

Attributes	Storage-Centric Protection	Database-Integrated Protection
Mix of Applications	Excellent (data protection solutions agnostic of applications that are running on the storage array)	Excellent (data protection solutions agnostic of applications that are accessing the database)
Mix of Databases	Excellent (same set of data protection solutions applies to all databases resident in the storage array)	Limited (data protection solutions typically applicable to the particular database)
Mix of Storage Technologies	Limited (data protection solutions typically applicable to the particular storage array)	Excellent (data protection solutions agnostic of the underlying storage technologies)

So … Is There A Silver Bullet?

The Answer Is ….

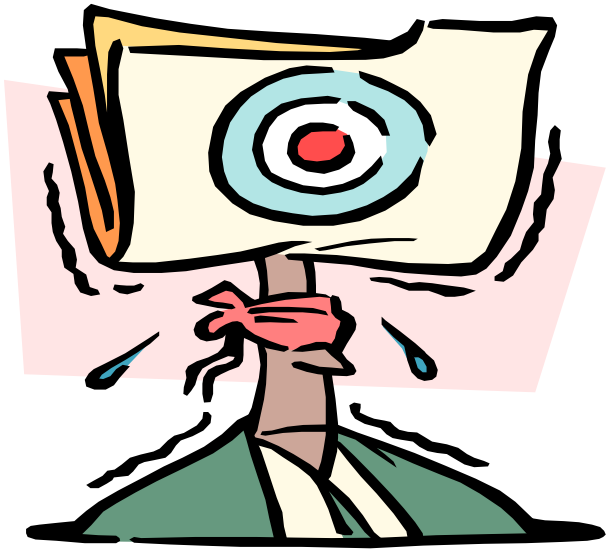

IT DEPENDS!

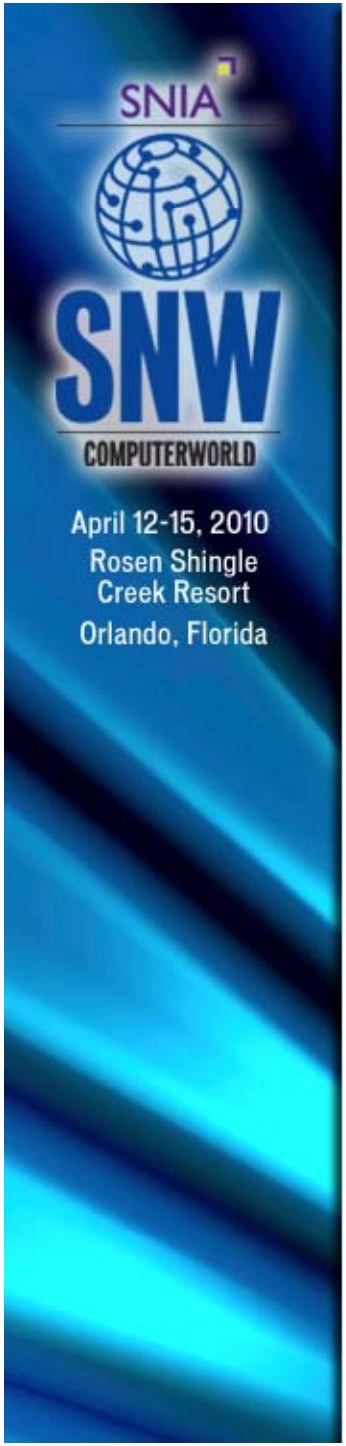

Summary

Database integrated solutions

- Comprehensive HA and data protection, with application integration
- Transactional consistency always maintained
- Applicable only to the particular database
- Supports a mix of underlying storage technologies

Storage centric solutions

- Comprehensive capabilities for read-only & writeable snapshots, backups & clone
- Insufficient application integration, corruption protection and resource utilization
- Applicable only to the particular storage array
- Supports a mix of databases in that array

Takeaway

Choose solutions that have the best balance across:

- ✓ Protection from various component failures
- ✓ Minimizing application downtime
- ✓ Utilization of all resources
- ✓ Reduction of runtime costs
- ✓ Support of technology mix