

ORACLE®

SOFTWARE POWERS THE INTERNET™

Database Storage Management Techniques and Technology

Paul Tsien
Product Manager
Systems Technology Group

Oracle, High Availability, Storage

- Recent trends
 - “Everything” is on the Web and “everything” needs to be stored somewhere
 - Every business is eBusiness
 - eCommerce (B2B, B2C) is everywhere
 - 7x24 is becoming a must
 - Storage is hot

Why High Availability Is Important?

Time to lose a customer on
the internet

30 Seconds

All customers will become
internet customers

**E-business cost of downtime
per second**

\$1800

**All businesses will
become E-businesses**

Why Storage Is Important?


Database cannot exist without
storage

Database Resides On Storage!

Managing data and storage is what
Oracle does

Storage Spending Is Growing Fast!

Storage Spending



Gartner Group Forecast

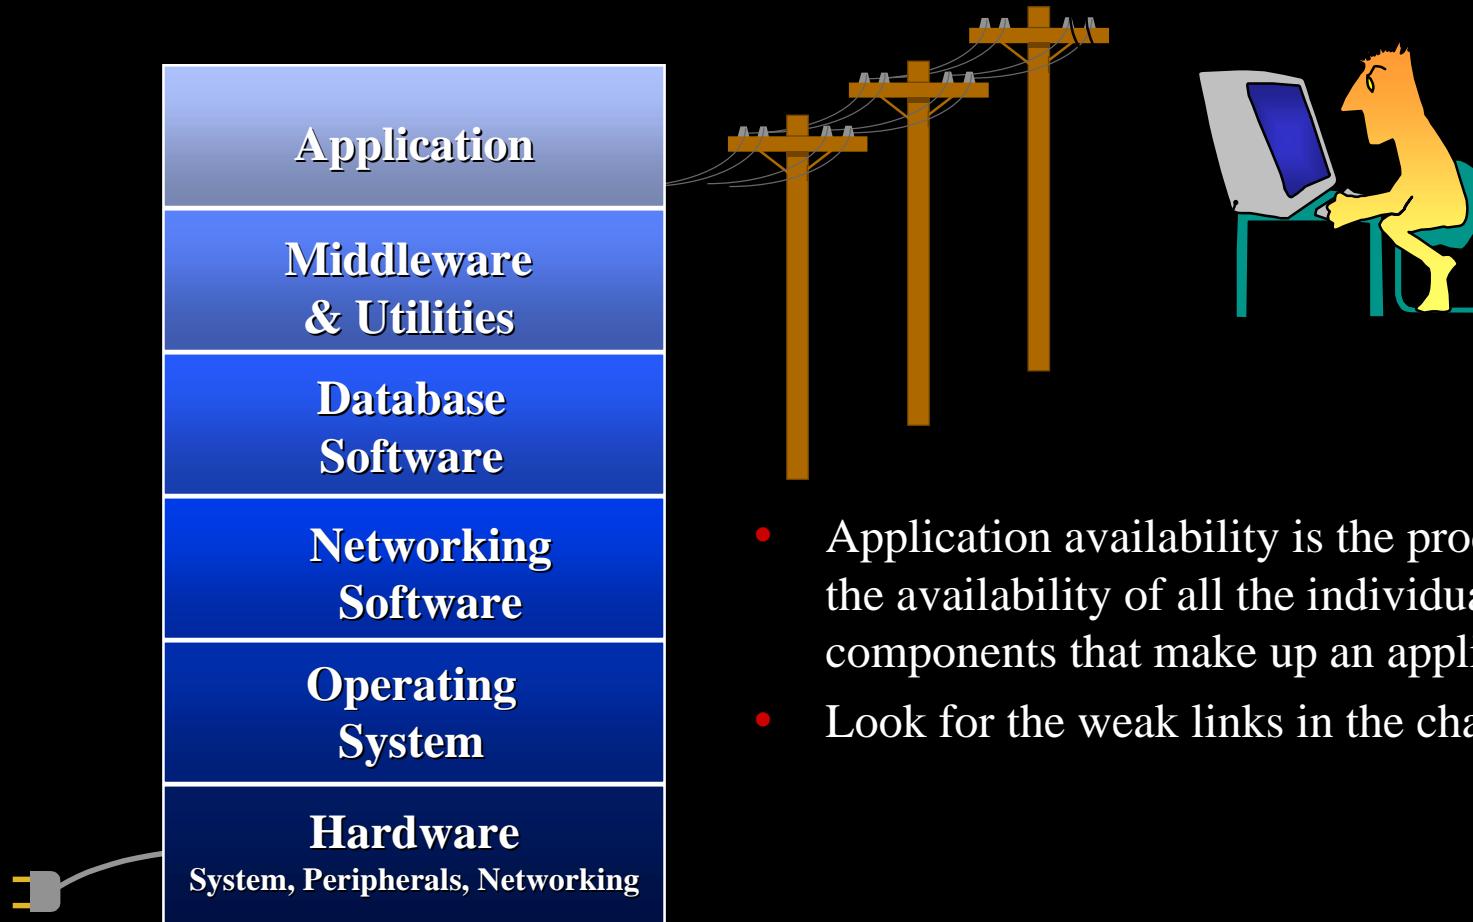
ORACLE

Hardware and software can
be replaced!

Data is Irreplaceable!

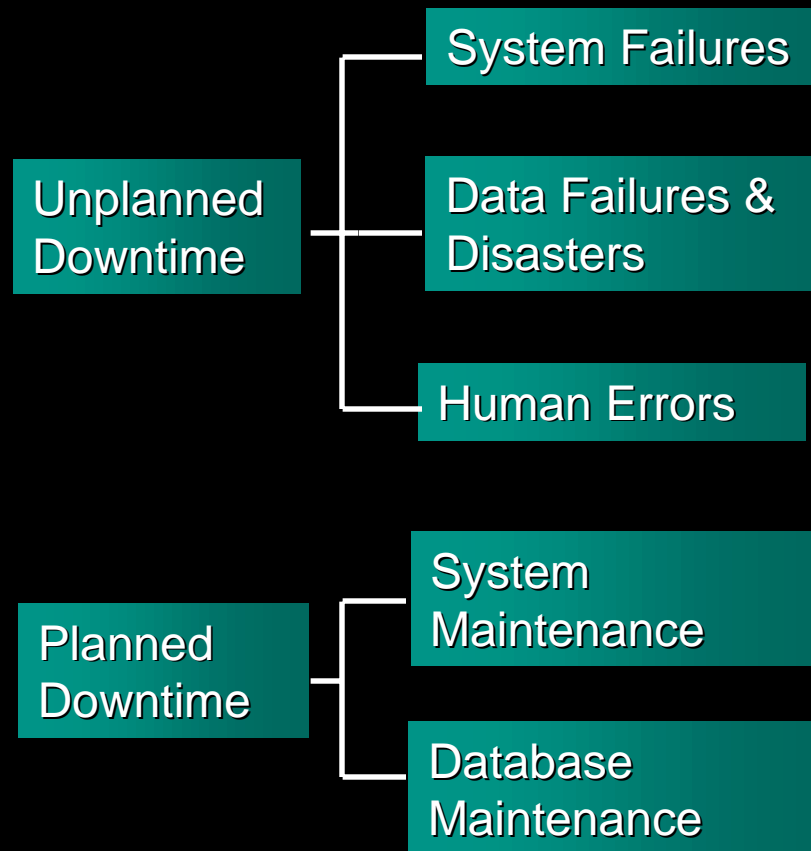
Customer information,
sales records, etc.

Application Versus Component Availability



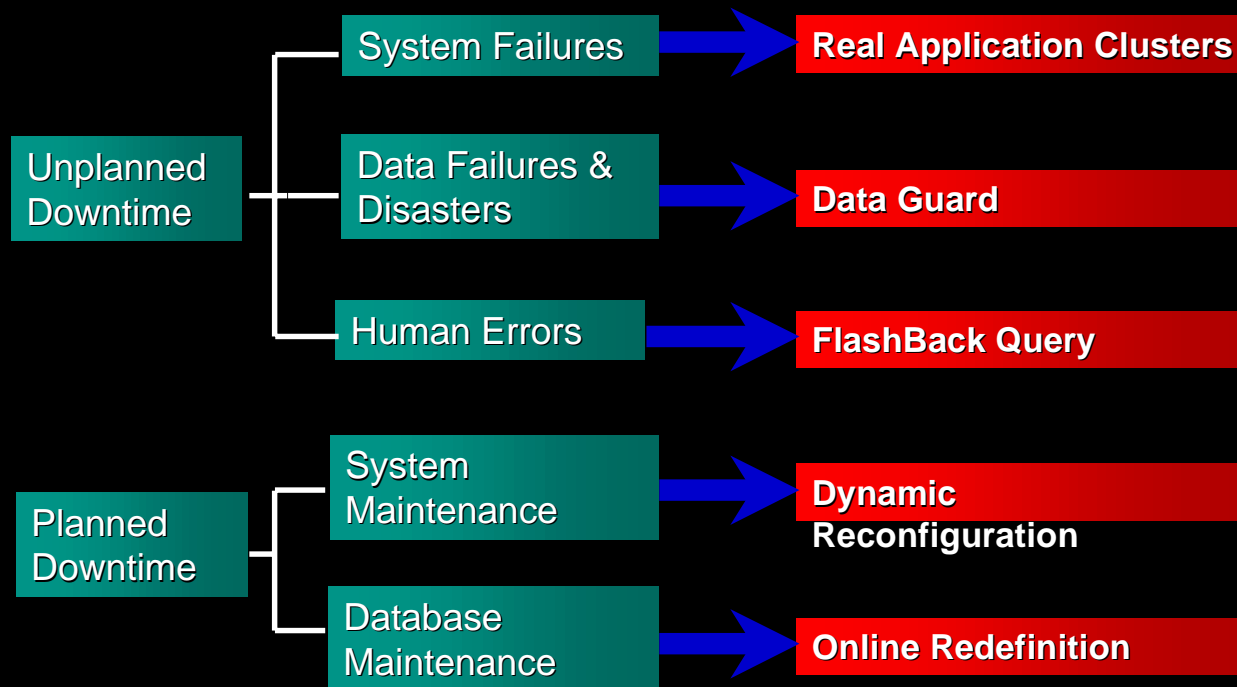
- Application availability is the product of the availability of all the individual components that make up an application
- Look for the weak links in the chain

Oracle And High Availability



We Have The Solutions!

- Oracle9i is architected for **Continuous Data Availability**
- Oracle9i handles all downtime scenarios



Oracle And Storage

- Storage was simple in the past
- Oracle deals with storage primarily from file perspective
 - Files system
 - Raw Partition
- Mostly SCSI direct attached disks (JBOD) or tape drives
- Oracle assumes storage is reliable (mostly)

Storage Is More Complex Now

- Sophisticated RAID systems with large amount of cache
- Storage appliances
- Fibre Channel, iSCSI, InfiniBand technologies
- Advanced volume managers and file systems
- Solid state disks
- Intelligence and cache are built into subsystems, controllers, hubs, switches, or even disks
- Snapshot and geo mirroring technologies

Storage Is Becoming More Independent and Autonomous

- Computing environment is heterogeneous
- Network Attached Storage (NAS)
- Storage Area Network (SAN)
- Storage Virtualization
- Storage as utility

Oracle And The New Storage Environment

- Oracle is working with industry standard organizations and leading storage vendors to deliver the best database/storage solutions in the industry!
 - Data integrity and availability
 - Interoperability
 - Performance

Oracle Storage Objectives

- Deliver world class database/storage solutions!
- Leverage Oracle's storage and I/O knowledge to enhance combined Oracle/storage solutions
 - Only Oracle knows how data will be accessed
 - ODM, OMF, BMR, etc.
- Use partners' storage technologies to complement Oracle capabilities
 - Remote mirroring, server-less backup, etc.
- Make storage management simple and easy
- Make Oracle and storage integration simple and easy
- Recommend and validate storage configurations

Oracle Storage Initiatives

- Oracle storage technology requirements
- Oracle storage management features
- Oracle storage APIs
- Oracle storage programs
- Oracle storage configuration projects


Oracle Storage Technology Requirements

- Cache and Non-Volatile RAM (NVRAM)
- Write ordering
- Atomic writes
- Other requirements

Oracle Storage Technology Requirements - Cache and Non-Volatile RAM (NVRAM)

- Oracle treats storage subsystem cache and NVRAM as persistent (as part of the disk)
- If content of cache or NVRAM is lost, all storage associated with that cache or NVRAM must be considered lost
 - Battery failure
 - Accidental cache board swap
 - Software corruption of cache
 - Unlikely but they do happen!

Storage Subsystem



Oracle Storage Technology Requirements - Write Ordering

- In the database remote mirroring case, Oracle write order must be maintained (across multiple logical volumes and physical boxes)
- For example remote log write must happen before remote data write

Local Writes

Block A

Block B

Time

Remote Writes

Block A must be written before Block B on the remote site

Oracle Storage Technology Requirements - Atomic Writes

- Oracle requires 512 byte (a physical sector) atomic writes
- Size of an Oracle log block
- For example, after a power failure, Oracle needs the log for recovery
- Not a problem for direct attached storage

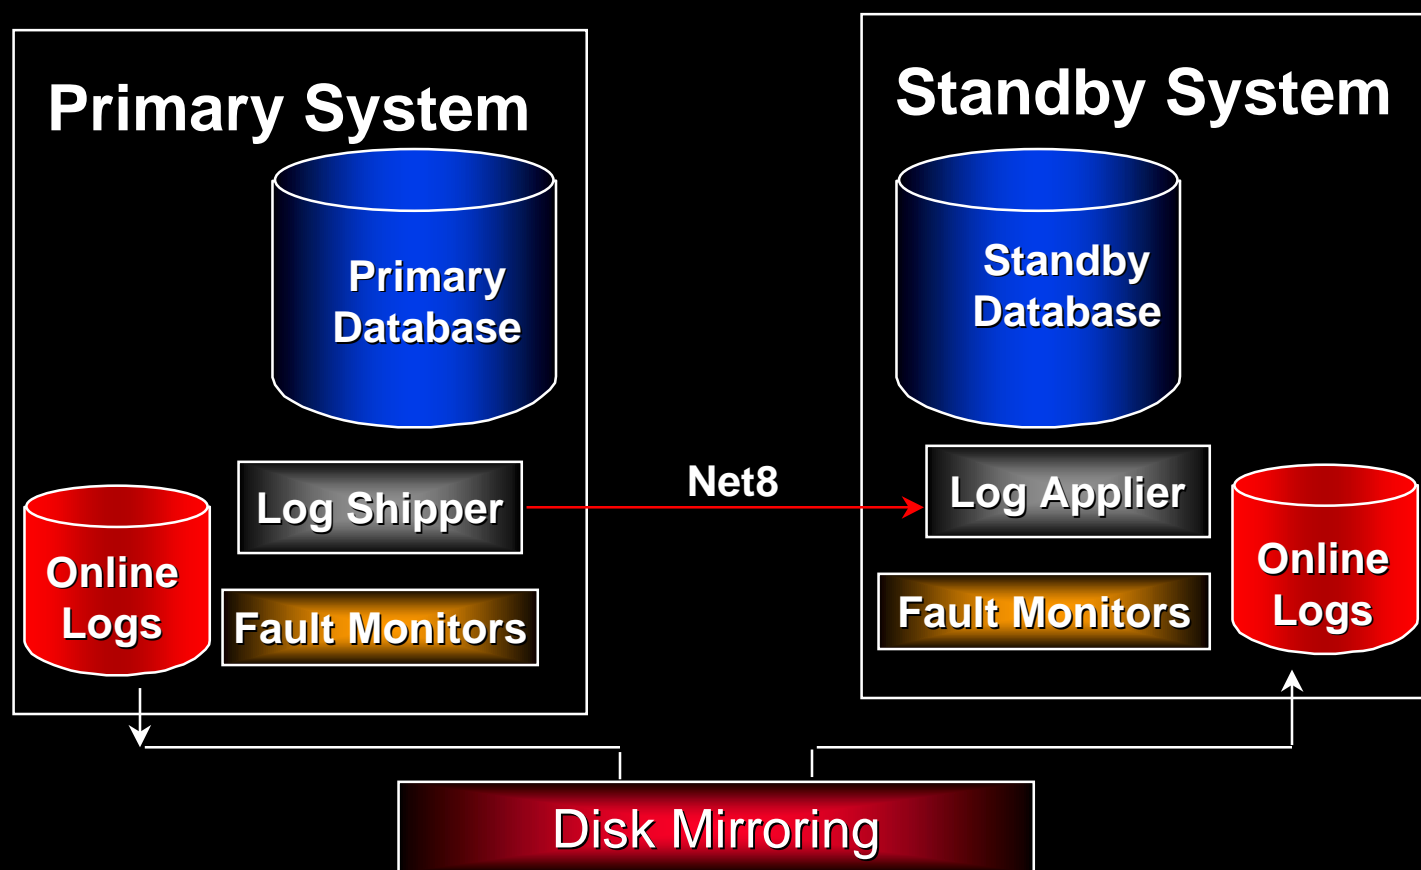
Oracle Storage Technology Requirements - Others

- Many implementation requirements
 - Oracle must be in hot backup mode for fuzzy backup
 - Can not use NFS client cache with OPS (no stale cache)
 - Async I/O implementation needs to be robust
 - Etc.

Oracle Storage Management Features

- Oracle Data Guard
- Automatic space management features
- Online data reorganization and redefinition
- Oracle Managed Files
- Recovery Manager (RMAN)
- Oracle Enterprise Manager (OEM)
- Other storage management features

Oracle Storage Management Features - Zero Data Loss with Oracle8i Data Guard



Oracle8i Data Guard maintains complete copy of Primary Database

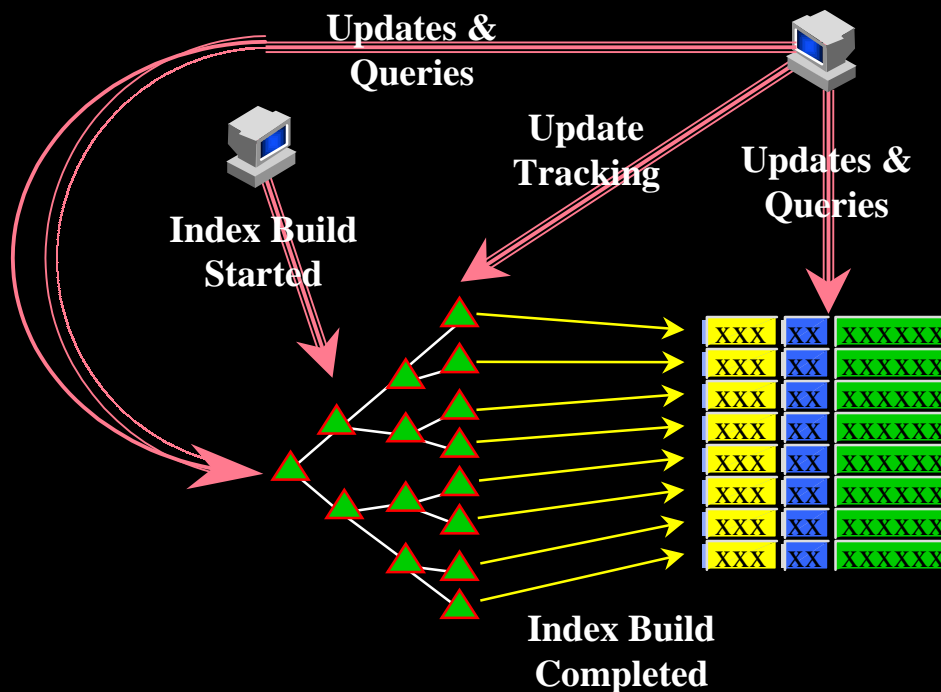
Oracle Storage Management Features - Automatic Space Management

- Oracle manages space automatically with locally managed tablespaces (inter-segments)
 - Ease of use, elimination of fragmentation
- Automatic segment space management (intra-segments)
 - Dynamic space affinity to instances
- Automatic undo management

Oracle Storage Management Features - Online Reorganization

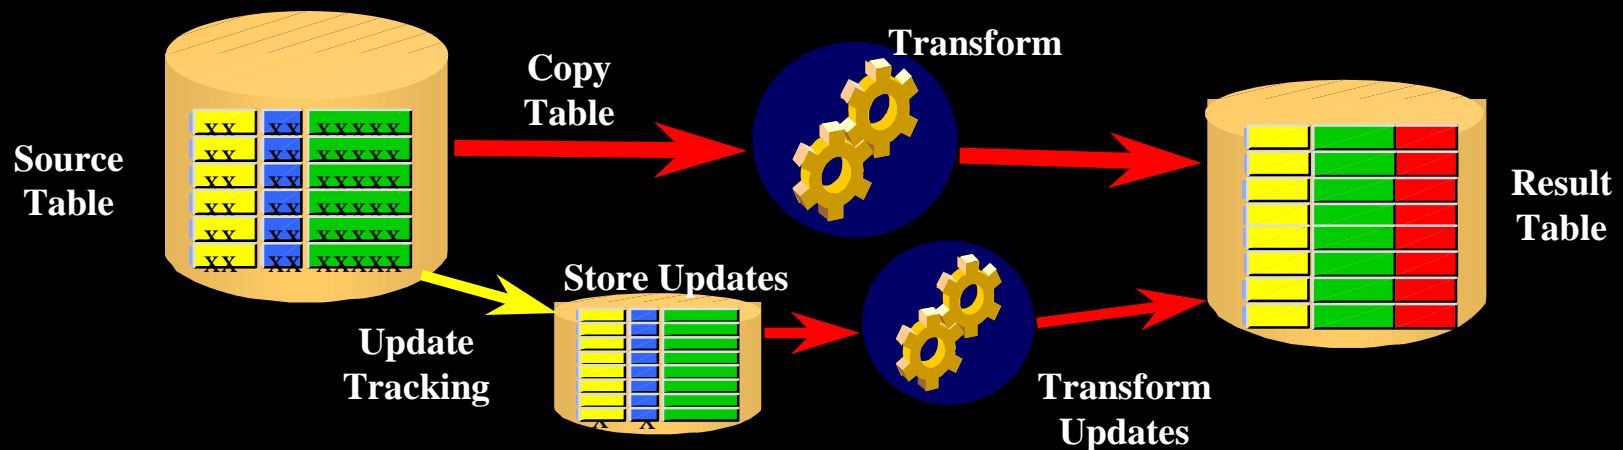
- Unlimited Online Index Operations

- Create new index for any type of table or index
 - Secondary indexes on IOTs
 - Reverse Key, functional, compressed indexes
- Move or reorg existing index
- Move or reorg existing Index Organized Table



Oracle Storage Management Features - Online Table Redefinition

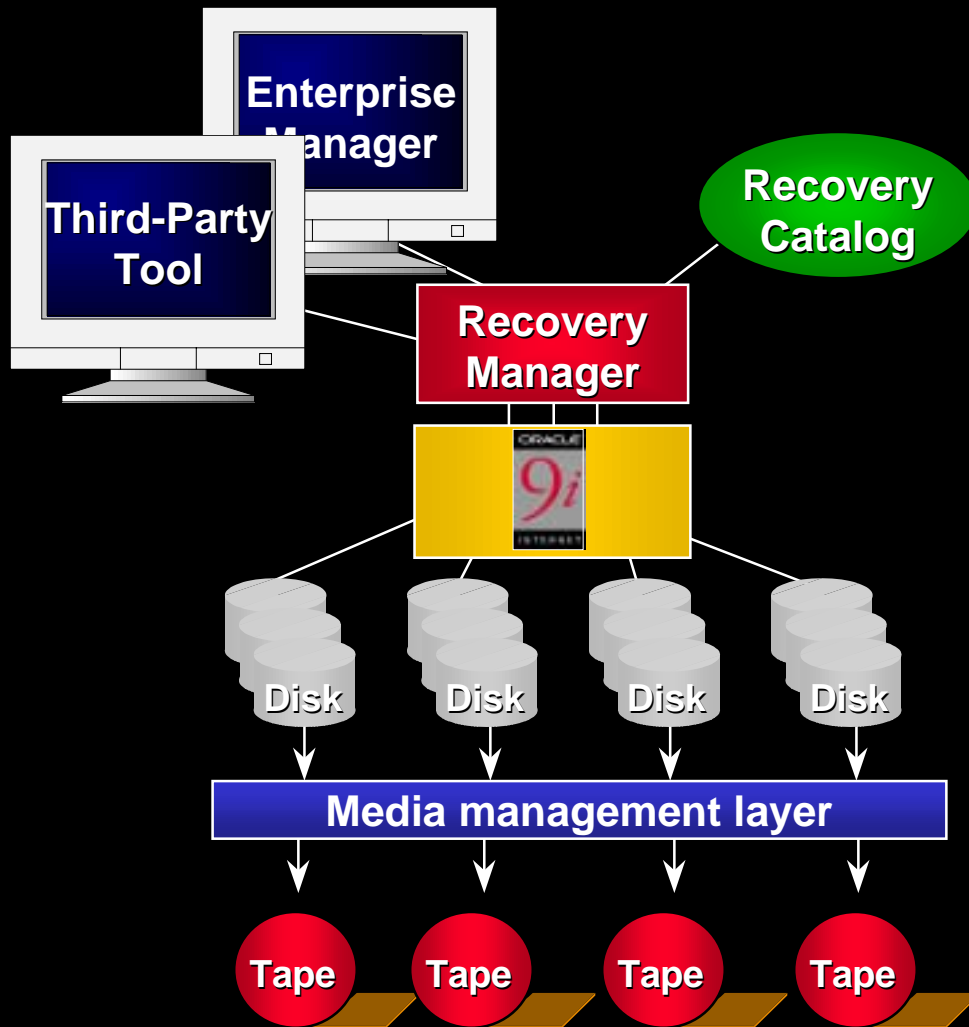
- Reducing planned outages
 - Table contents are copied to a new table
 - Eliminates fragmentation
 - Change location, table type, partitioning, parameters
 - Contents can be transformed as they are copied
 - Allows changes to columns, types, sizes, etc.
 - Updates and Queries can continue uninterrupted



Oracle Storage Management Features - Oracle Managed Files (OMF)

- Simplifies the administration of Oracle databases files
- Administrators specifies in terms of database objects
- Oracle automatically creates and deletes files as required
- Administrators can specify different locations for data files and online redo log/control files
- Third party applications need not embed OS specific file names in their scripts

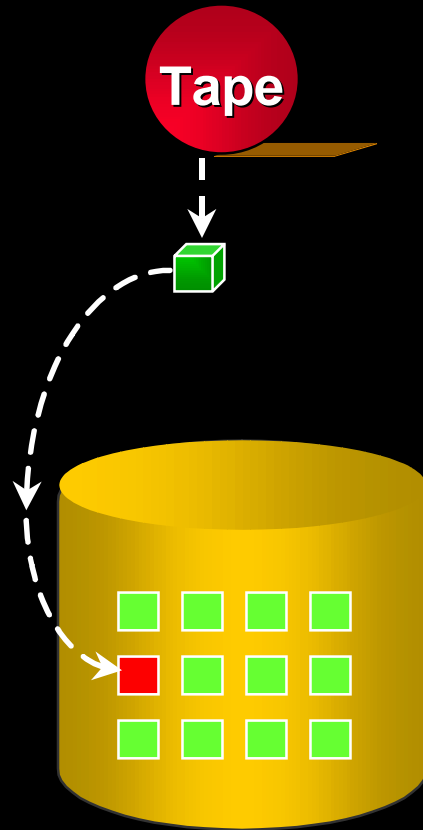
Oracle Storage Management Features - Recovery Manager (RMAN)



• Recovery Manager is a tool that:

- Manages the backup, restore, and recovery process
- Creates and maintains backup policies
- Catalogs all backup and recovery activities
- Many optimizations for improved availability and performance

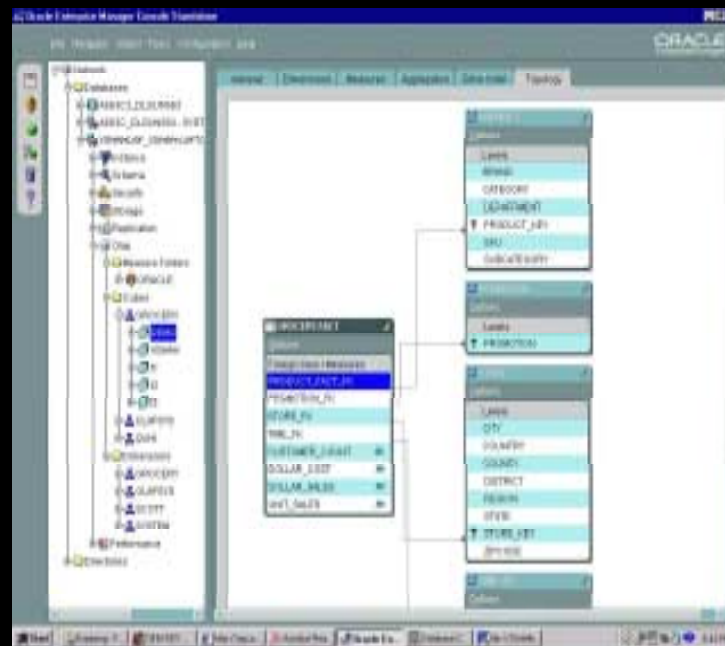
Oracle Storage Management Features - Block Media Recovery



- Increased data availability
 - Datafile remains online
 - Only blocks requiring recovery are inaccessible
 - Maximize data availability
- Reduce mean time to recover
 - Selectively restore and recover damaged blocks
 - Minimal I/O needed with redo only applied to damaged blocks
- Managed by Recovery Manager

Oracle Storage Management Features - Oracle Enterprise Management (OEM)

- Oracle GUI front end to space management, storage management, Data Guard, RMAN, etc.



Oracle Storage Management Features - Others

- HSM support
 - Read_only_open_delayed option
- Read-only tablespaces
 - Can be on WORM or CD-ROM
- Partitioning
 - Improve performance, availability and recoverability

Oracle Storage Management Features - Others

- Transportable tablespaces
 - Mechanism to quickly move a tablespace from one Oracle database to another
 - Auto parts branch store receiving monthly catalog as CDs
- Heterogeneous data transfer
 - Import/Export
 - Replication

Oracle Storage APIs

- Media Management/Proxy Copy API
- Oracle Mirror Resilvering
- Oracle Disk Manager (ODM)
- Other Storage APIs

Oracle Storage APIs - Media Management/Proxy Copy API

- Also known as System Backup to Tape (SBT) API
- Defines the interface between Oracle's software (i.e. RMAN) and non-Oracle media management software packages
- API defines functions including:
 - Create backup files
 - Write to and read from backup media
 - Search and remove backup files
 - Server-less backup

Oracle Storage APIs - Oracle Mirror Resilvering

- Use for software disk mirroring
- How to recover mirror inconsistency? (For example, after a power failure)
- OS may need to copy the whole mirror
- Catch up of a large mirror may take a long time
- Leverage Oracle redo log to speed up resilvering

Oracle Storage APIs - Oracle Disk Manager (ODM)

- Special interface from Oracle to ‘files’
 - Raw performance with file system manageability
- Reduce system overhead by not requiring file open per each Oracle process
- Special locking modes prevent errors
- Batched async I/O issue and completion
- Allows dynamic resizing of files
- Can pass usage hints, priorities
- Automatic generation of unique file name - OMF
- Support Oracle mirror resilvering

Oracle Storage Programs

- Back Solutions Program (BSP)
- Oracle Storage Compatibility Program (OSCP)
- Other Storage Programs

Oracle Storage Programs - Backup Solutions Program (BSP)

- Recovery Manager - RMAN
 - Oracle's integrated backup and recovery tool
 - SBT interface - interface to tertiary storage
 - Proxy Copy - server-less backup
 - Oracle provides SDK to certify backup product interoperability with RMAN
 - Enterprise backup integration with RMAN
 - Web page of supported configurations

Oracle Storage Programs - Oracle Storage Compatibility Program (OSCP)

- Validate Compatibility of “Specialized” storage solutions with Oracle
 - Oracle over NFS - NAS
 - Remote Mirroring - primarily for log files
 - Snapshot technologies
- Process - architecture review, test kit, usage guide, web page of supported configurations
- A validation program only
 - Vendors self-test Oracle supplied OSCP test suites

Oracle Storage Configuration Projects

- Standard storage configuration
 - Stripe And Mirror Everywhere (SAME)
- High availability standard configuration and validation
 - Oracle Parallel Fail Safe (OPFS)
 - HP Server with HP SureStore XP or EMC disk array
 - Oracle8i Data Guard
 - Push button remote fail over and fail back to protect against corruptions, human errors and disasters

Standard Storage Configuration - SAME

- Standard methodology for configuring storage for Oracle databases
 - Optimal storage configuration made easy
- S.A.M.E. = Stripe And Mirror Everything
 - Specific details and methodology described in white paper
- Storage subsystem improvements can enable wide spread usage of the SAME
- Joint testing and white paper on specific storage stacks

Summary

- Oracle and partners continue to deliver world class database/storage solutions!
 - Oracle storage technology requirements
 - Oracle storage management features
 - Oracle storage APIs
 - Oracle storage programs
 - Oracle storage configuration projects

Useful URLs

- Oracle high availability and storage management technical white papers
 - <http://technet.oracle.com/deploy/availability/>
- Oracle Storage Compatibility Program (OSCP)
 - <http://www.oracle.com/ip/deploy/database/storage/>
- Oracle Backup Products
 - <http://www.oracle.com/ip/deploy/database/features/recovery/index.html?content.html>



Q U E S T I O N S
A N S W E R S