VERITAS™

BUSINESS WITHOUT INTERRUPTION™

Improving Oracle Database Performance

Robert Spurzem

Product Marketing Manager

VERITAS Software Corporation

Mountain View, CA

THE DATA AVAILABILITY COMPANY™

Agenda

- **Challenges of Storage Management in the Internet World**
- **Storage Management for Oracle Databases in the Internet World**
  - Overview
  - Building a Foundation for Performance
  - Replication for Off-Host Processing and Disaster Recovery
  - Clustering for Continuous Data Availability
  - Efficient Data Management for Data Protection
  - Summary

Evolution of the Application Market

The ‘70s:
Basic Accounting

Automating manual bookkeeping tasks.

The ‘80s:
Enterprise Applications

Information on every asset of the enterprise operation. Storage requirements grow

The ‘90s:
Data Analysis

Analysis of business-critical information. Storage requirements grow “ten-fold again”!

What’s next?

Storage requirements *will* grow “ten-fold again”!

The ‘90s/’00s:
WEB

Proliferation of multimedia data. Storage requirements grow “ten-fold *yet* again”!

E-commerce Configuration

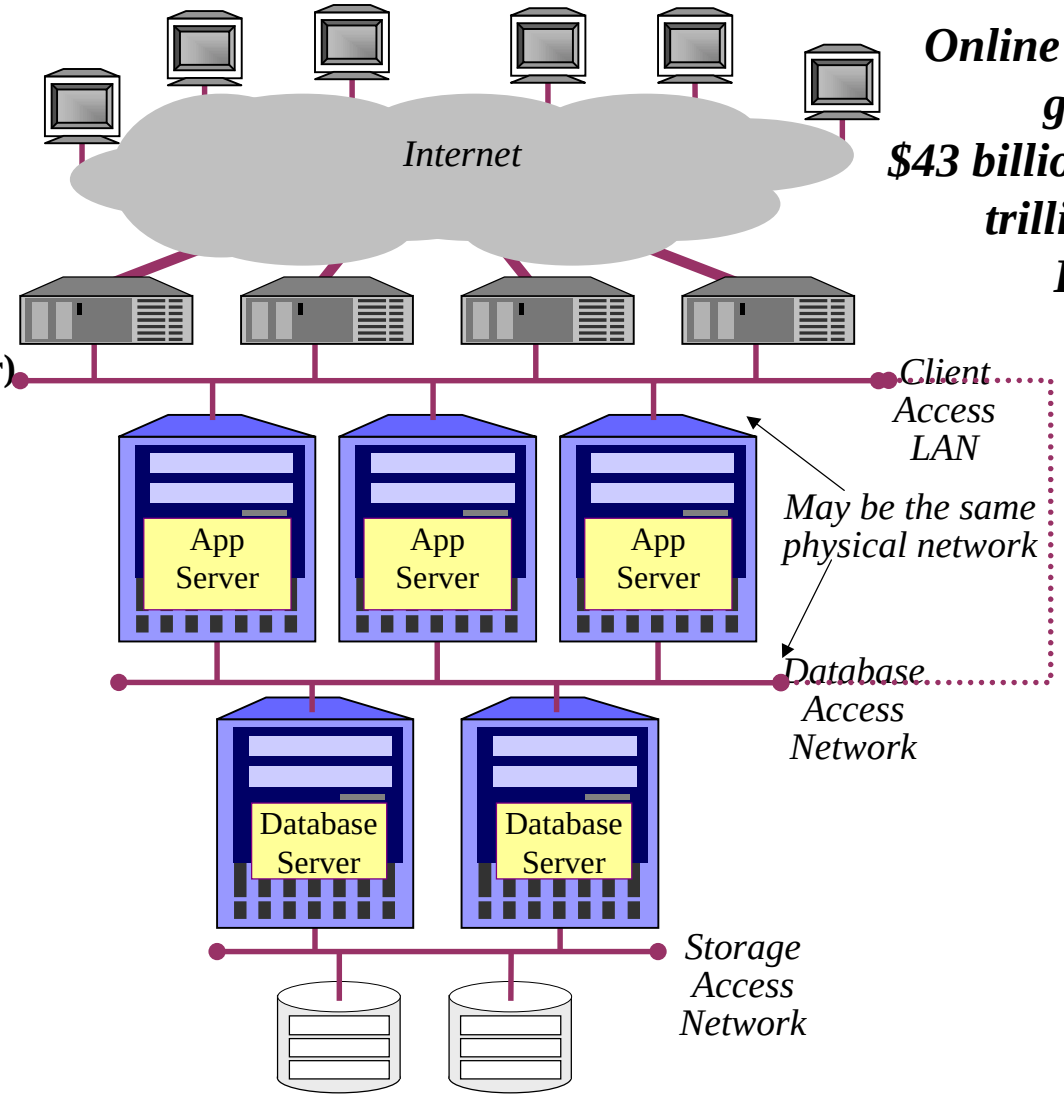

Online Commerce will grow from $43 billion in 1998 to $2.7 trillion in 2004 - Forrester

Challenges in the Internet World

- Every internet user needs high performance access to information
- Volume of data generated by each new major web application has expanded exponentially
- Need to be assured that data is always available - 24x365

Storage Fuels Applications - Databases Drive them

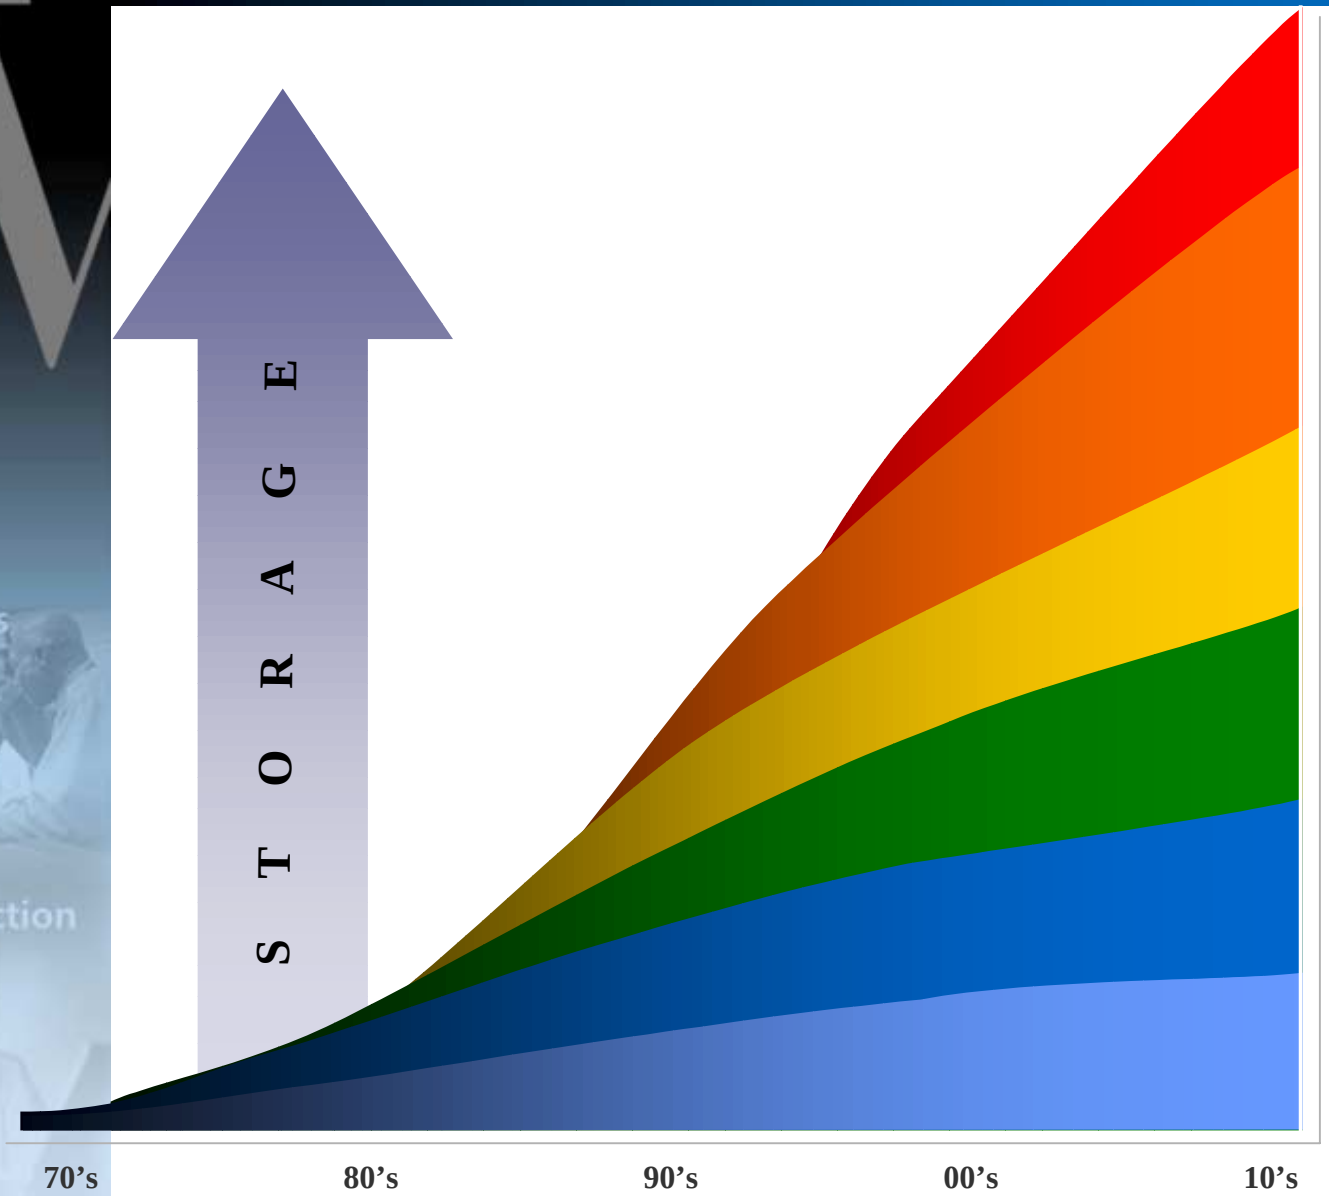

Storage Fuels Applications - Databases Drive them

As Storage Increases Exponentially so Does the Need for Storage Management

S T O R A G E

Storage Management Curve

- *Voice/data integration*
- *Wireless*
- *Video-on-demand*

WEB

Data Analysis

Enterprise Applications

Client/Server

Basic Accounting

70's 80's 90's 00's 10's

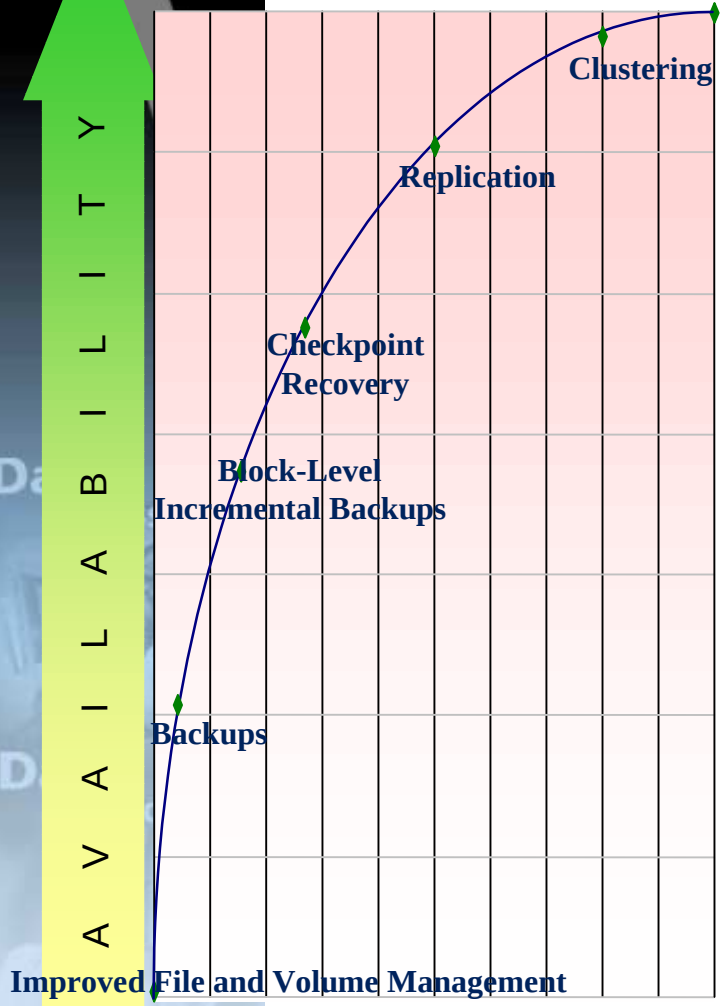

% uptime	Annual downtime	Annual Internet Shopping Cost	Annual Brokerage System Cost
99.9999	30 seconds	$950	$53,750
99.999	5 minutes	$9,417	$537,500
99.99	52 minutes	$98,000	$5,590,000
99.9	8.75 hours	$988,750	$56 million
99.5	43.7 hours	$5 million	$280 million
99.0	87.6 hours	$10 million	$560 million
98.0	180+ hours	$20+ million	$1+ billion
95.0	450+ hours	$50+ million	$3+ billion

Storage Management Objectives

- **Configuring and maintaining storage resources to:**
- **Maximize I/O performance**
  - Is I/O done in an optimal manner for my application?
  - Is my storage configured optimally?
- **Minimize loss of access to data: “no time for down-time”**
- **Meet service level agreements**
- **Increase overall system reliability**

Overview of Storage Management Components

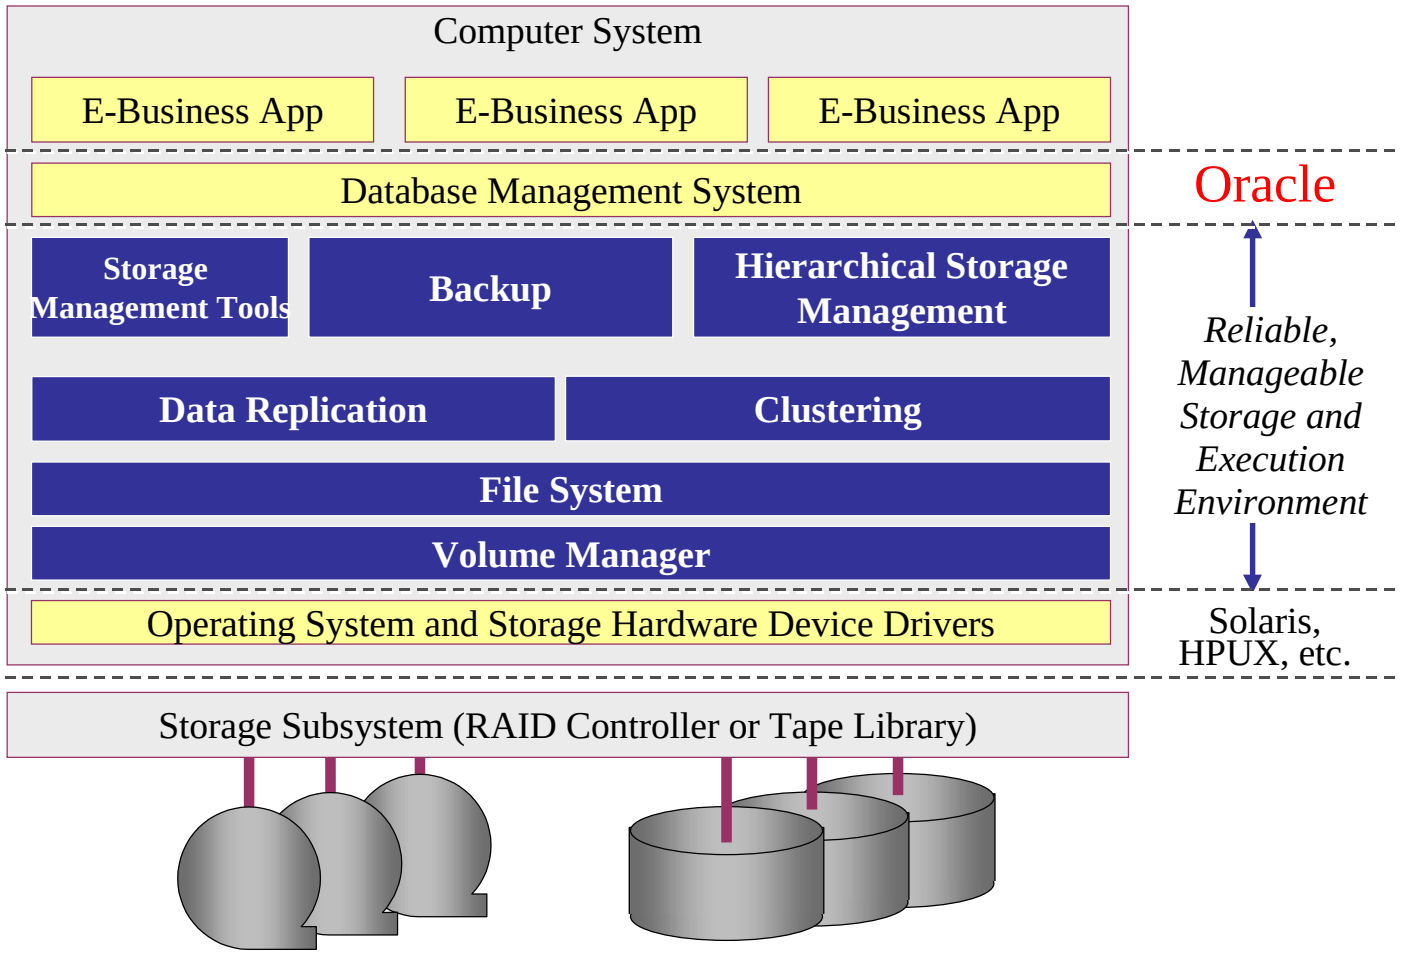

Storage Is the Foundation of a Database System

- **Database functions to:**
  - Store, update, and retrieve structured information
- **Good storage design helps:**
  - Meets rigorous performance and availability requirements of production databases

VERITAS™

BUSINESS WITHOUT INTERRUPTION™

Building a Foundation for Performance

Performance Can Be Everything

- **Need to meet service level agreements**
  - Elegant designs can fail in production because of performance problems
  - Administrators spend hours monitoring and balancing I/O loads to achieve better performance

What Makes up the Foundation?

- **Volume management**
  - Move data between devices for load balancing
  - Reorganize storage layout
  - Dynamic multi-pathing
  - Hot sparing
  - RAID functionality
  - Fast recovery mechanisms
- **File system**
  - Journaling / fast recovery, extent-based
  - Online growth & shrinkage
  - Online reorganization / de-fragmentation
- **Perform all disk & file management operations while applications are running**

Logical Volumes As a Highly Performing Platform for Databases

- **Resolving hot spots can be:**
  - Time and resource consuming
  - Analysis not guaranteed to work
    - Past behavior doesn't predict future behavior
  - May not be possible

- **Creating logical volumes of physical disks can balance I/O loads**
  - Through data striping

RAID Functionality

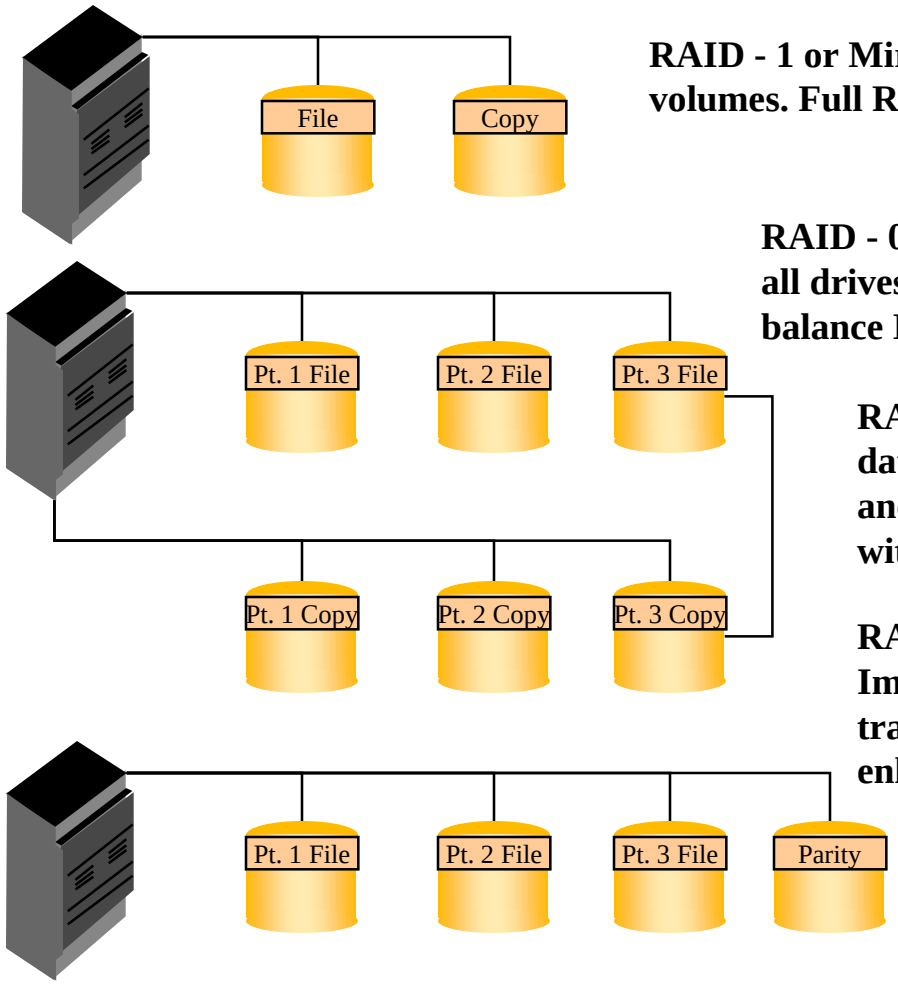

RAID - 1 or Mirroring: Exact duplicate volumes. Full Redundancy.

RAID - 0 or Striping: Stripes data across all drives. No Redundancy., but helps to balance I/O load

RAID - 0 + 1 or Mirrored Stripes: Stripes data across all drives. Mirror with another RAID 0 column. Full Redundancy with Full Performance.

RAID - 1 + 0 or Striped Mirrors: Improves performance by parallel data transfer with benefits of mirroring and enhanced redundancy

RAID - 5: Stripes data across all drives with parity. Limited Redundancy. Doesn't require duplicate storage investment.

Volume Manager Increases Performance Through Proper Storage Configuration

- **RAID functionality without specialized hardware**
  - Minimization of I/O bottlenecks
  - Maximum throughput & bandwidth of I/O
- **Online Relayout**
  - Ability to change configuration ONLINE

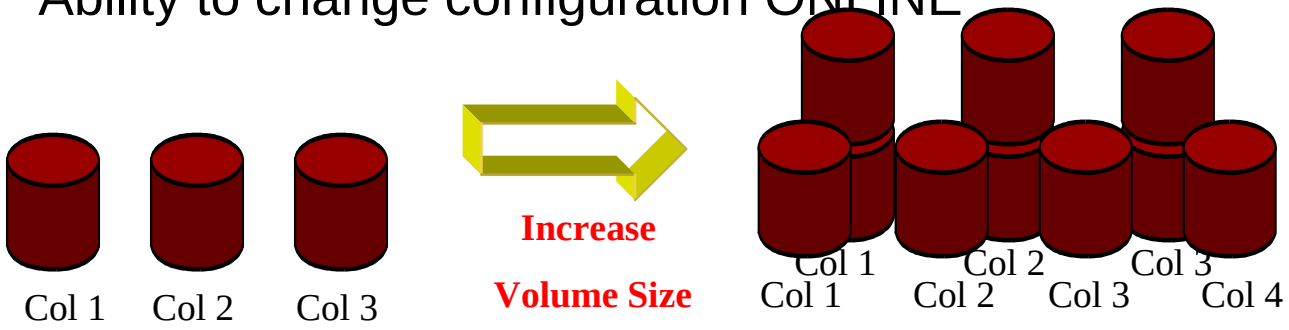

  - Enables volume growth by adding columns
    - *Incremental growth*
  - Ability to convert a RAID5 volume to a striped volume ONLINE
  - Allows ONLINE changing of stripe unit size

File Systems As a Highly Performing Platform for Databases

- Best of both worlds: raw I/O performance and ease of file system administration

Oracle

Raw Disks		Conventional File System	
Advantages	**Disadvantages**	**Advantages**	**Disadvantages**
Performance ▾ Best OLTP performance **Data Reliability** ▾ No potential for data loss with data written to disks directly	**Manageability** ▾ "Invisible" local disks for tablespaces ▾ Difficult to grow tablespaces	**Manageability** ▾ Tablespace organization ▾ System-wide backup and restore policies ▾ Easy to create and expand files for Oracle	**Performance** ▾ Lower database throughput (30-9 of raw disk) ▾ High CPU overhead Data Reliability ▾ Potential data loss

Journaled File Systems

Journaled File Systems if not closed properly simply check the status of all blocks that were previously being updated. This is known via the entries in the intent log.

A journaled file system's *fsck,* or file system check, is totally dependent on the amount of entries in the journal.

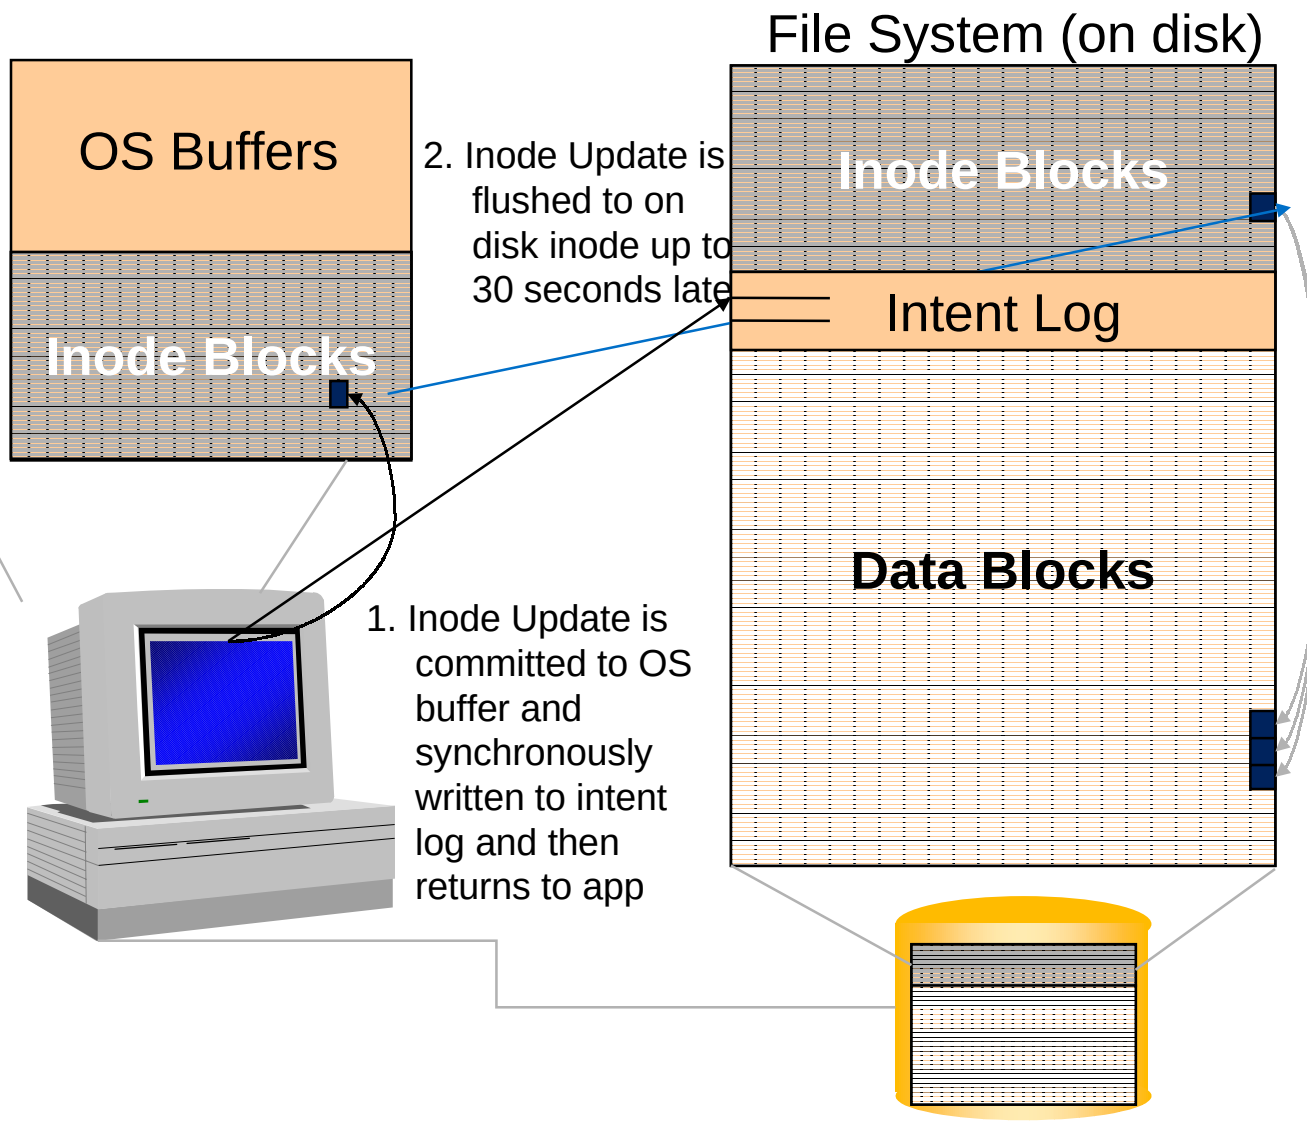

Increased Performance for Large I/O

- **Extent-based allocation can accelerate I/O by increasing the I/O transfer size in the File System**
- **Optimal for multimedia files commonly found on the Internet**

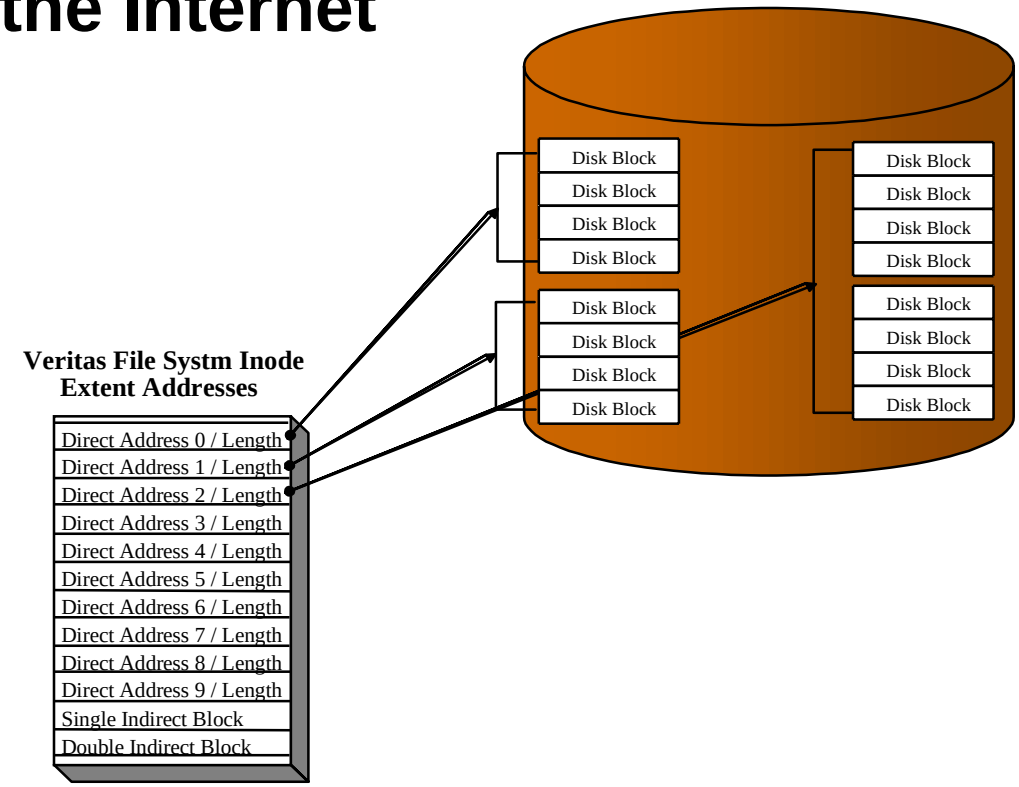

Raw Device Performance
Quick I/O for Databases

- **Break-through VERITAS File System interface technology**
- **Presents regular VERITAS File System files to Oracle as raw character devices**

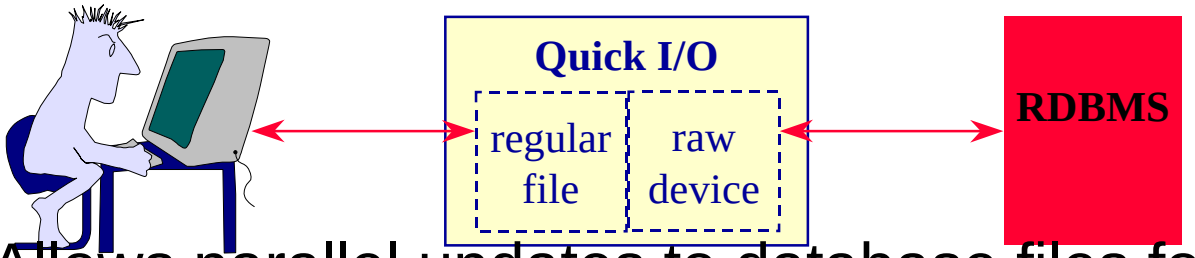

  - Allows parallel updates to database files for increased throughput
  - Takes advantage of kernel async I/O (Solaris' KAIO)
  - Oracle handles locking for data integrity
- **Eliminates traditional UNIX file system overhead**

Quick I/O Features

- **Direct (non-buffered I/O)**
  - Buffered reads in case of Cached Quick I/O
- **Bypasses the file system layer's locking**
- **Supports asynchronous I/O**
- **Easy Online Management to turn Cache on/off for Quick I/O files**
- **Multiple Database instances can co-exist with selective caching for specific datafiles**
- **Minimum changes to existing configuration**
- **No change to the OS or Database version levels**

Quick I/O Performance

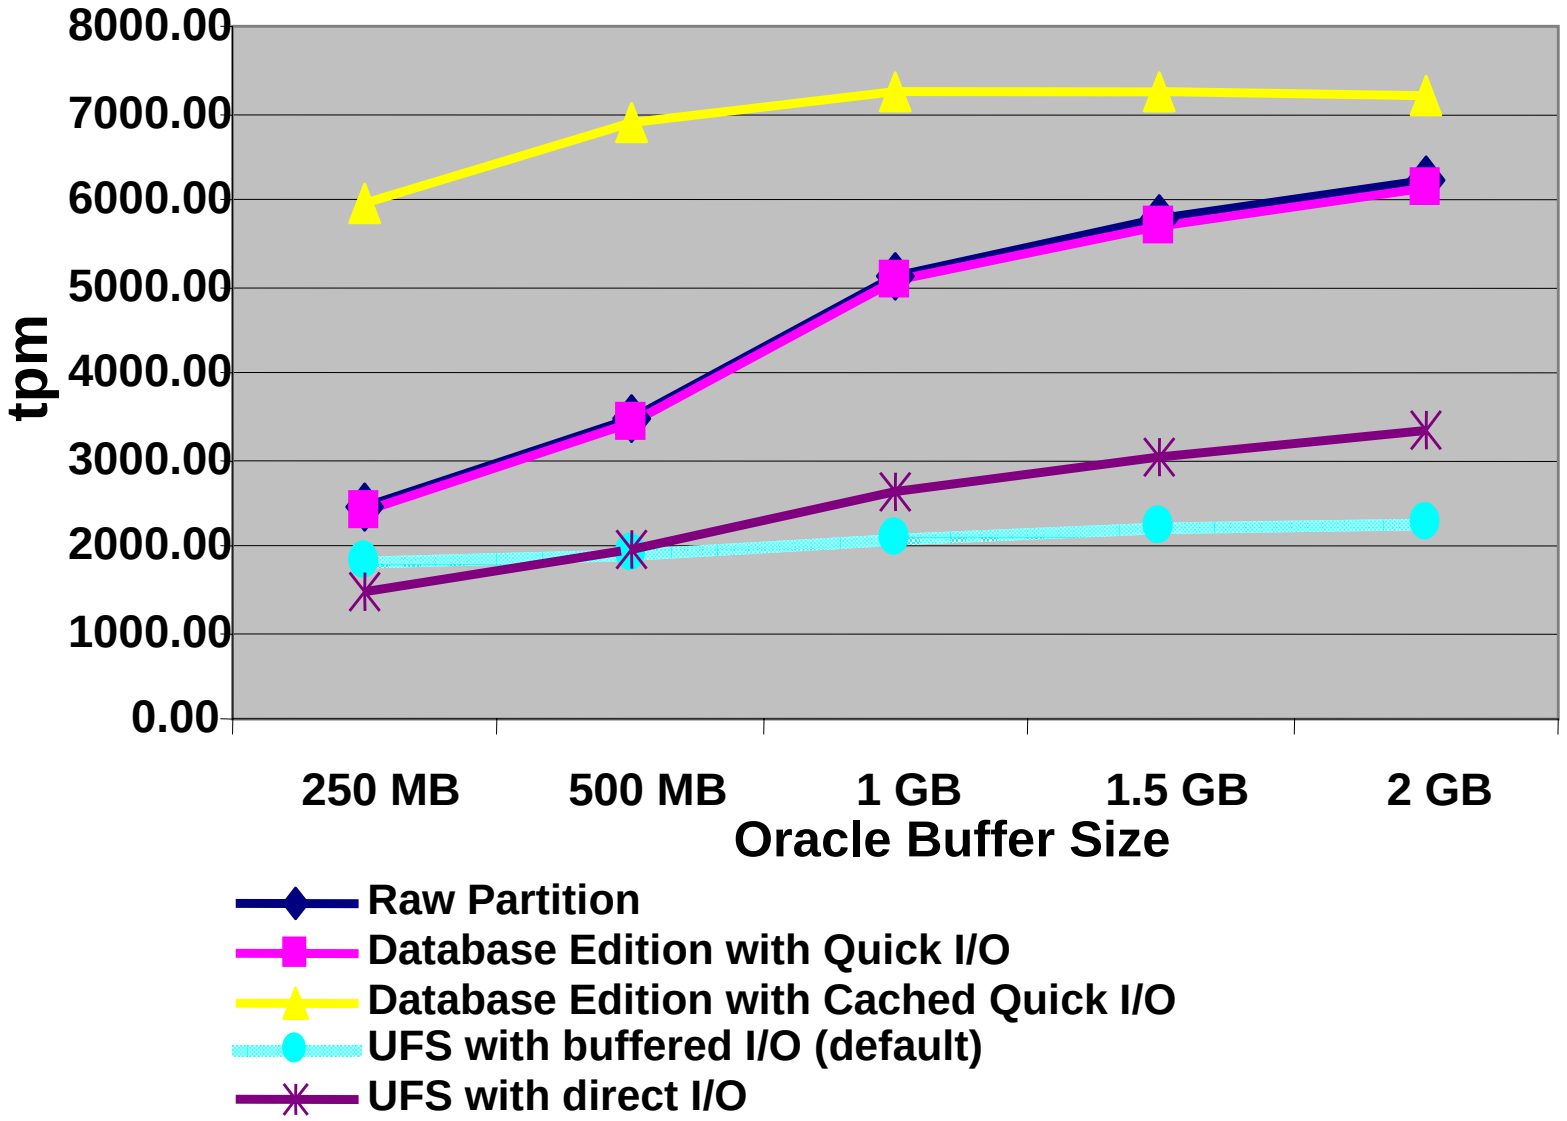

VERITAS™

BUSINESS WITHOUT INTERRUPTION™

Replication for Off-host Processing and Disaster Recovery

Why use Storage Replication?

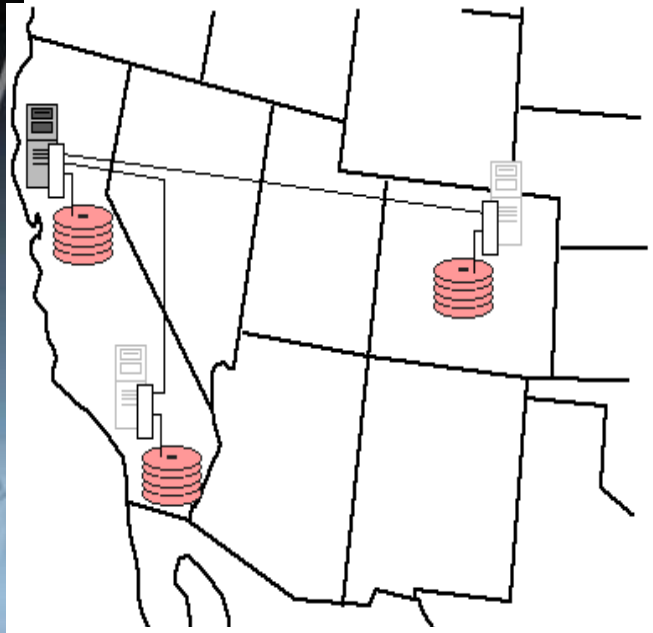

- 1 - Site Disaster Recovery
- 2 - Information Distribution
- 3 - Off-host processing or backup
- 4 - Wide-Area HA (migration)
  (Clustering section)

Disaster Recovery

- Replicate critical information to a disaster-safe location
- Replicate data over WAN
- Current data at remote hot site
- Replicate entire environment (not just logs)
- Data ALWAYS consistent (integrity)

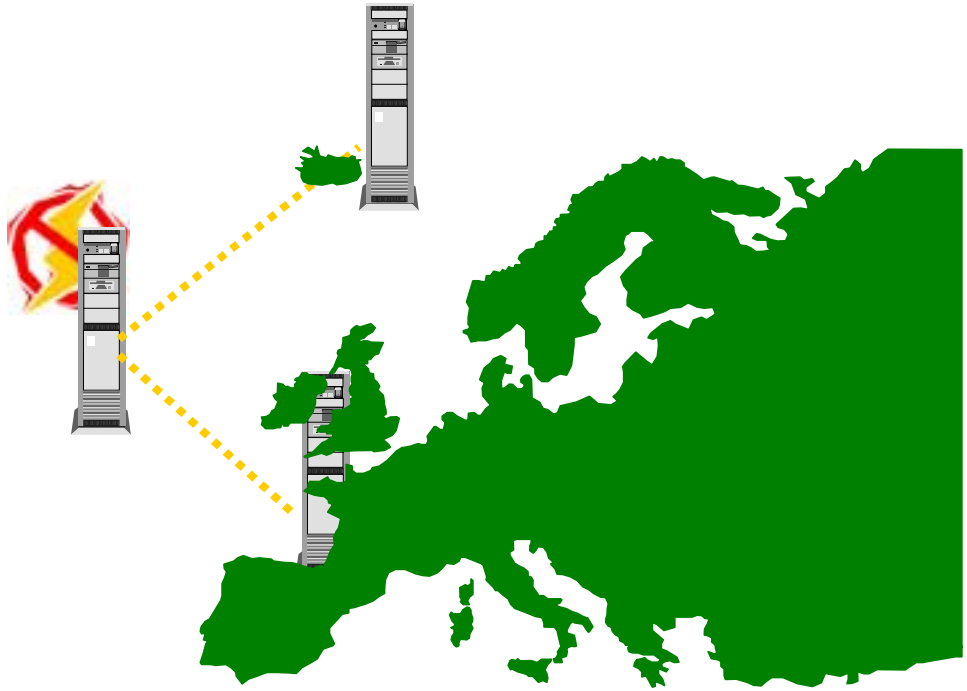

Information Distribution

- **Efficiently keep web data synchronized**
  - One to many (one source, many targets)
  - Keep multiple Web or information servers totally in sync
  - User controllable bandwidth usage

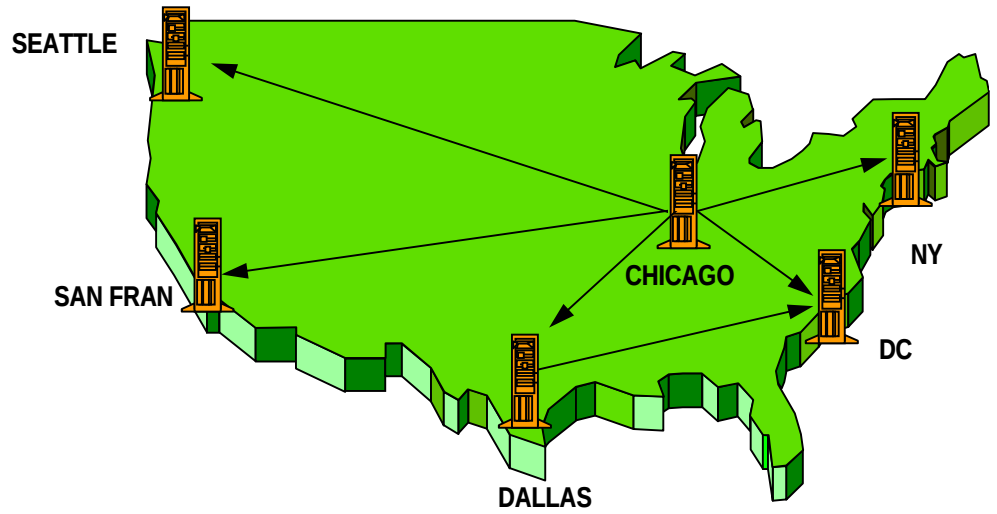

Off-host Processing
Block-Level Replication

- Off-load processing of maintenance tasks to secondary server
- Backups or any decision support can then be performed on data on second system
- Low CPU and network usage due to only replicating changed data blocks

Higher Performance on Primary Server

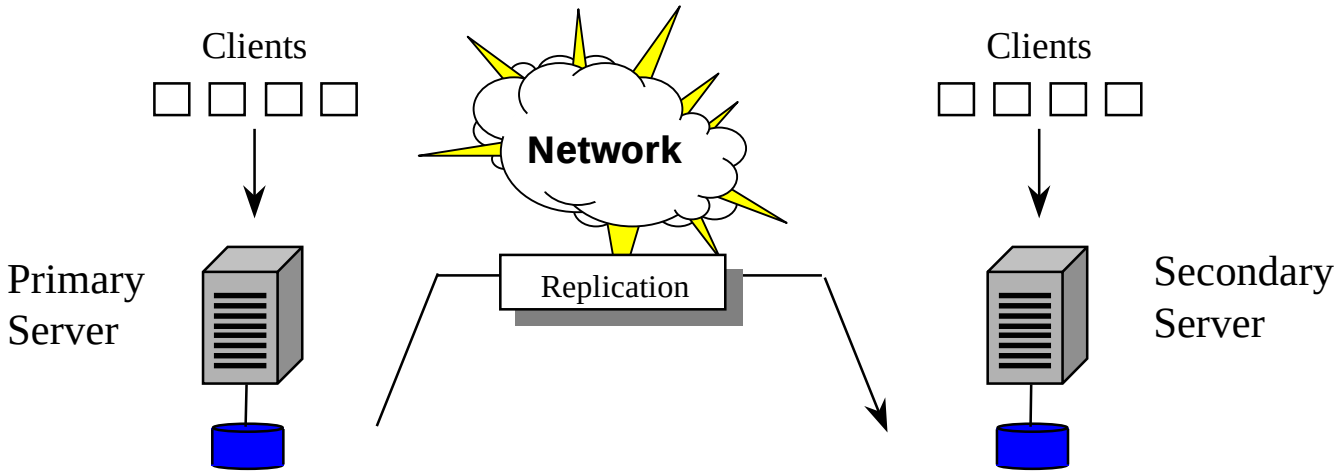

VERITAS™

BUSINESS WITHOUT INTERRUPTION™

Clustering for Continuous Data Availability

Business Without Interruption

- **Database environments increasingly complex**

Increased Risk to Data & Bottom Line

Storage Management Software Helps Mitigate Risk

Clusters Help Solve E-business Problems

- Take over for failed server or crashed application
- Load balance among servers
- Connect clients to alternate data paths if network links fail
- Recovery from site disasters
- Help meet service level agreements

High Availability

- Why?
  - **Service resumes with minimal impact**
  - **Administrator fixes problem in own time rather than with users shouting!**

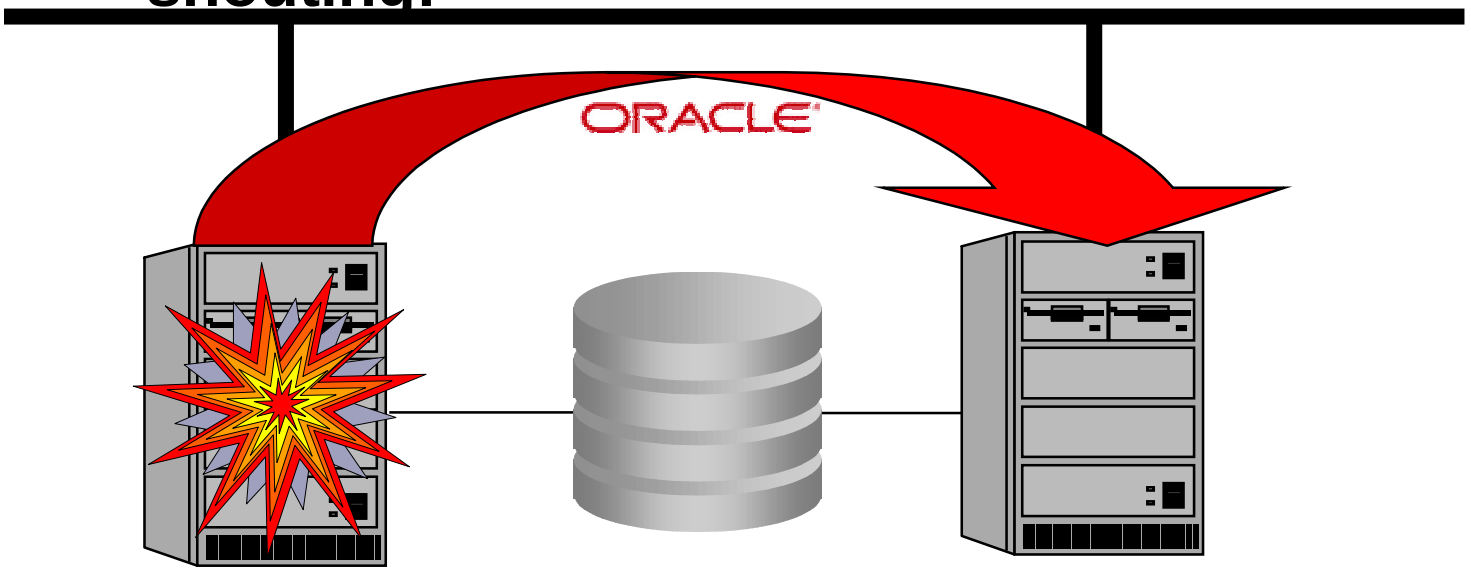

Clustering

- **Cluster: a type of parallel or distributed system that:**
  - Consists of a collection of interconnected whole computers,
  - And is used as a single, unified computing resources

Pfister - In Search of Clusters

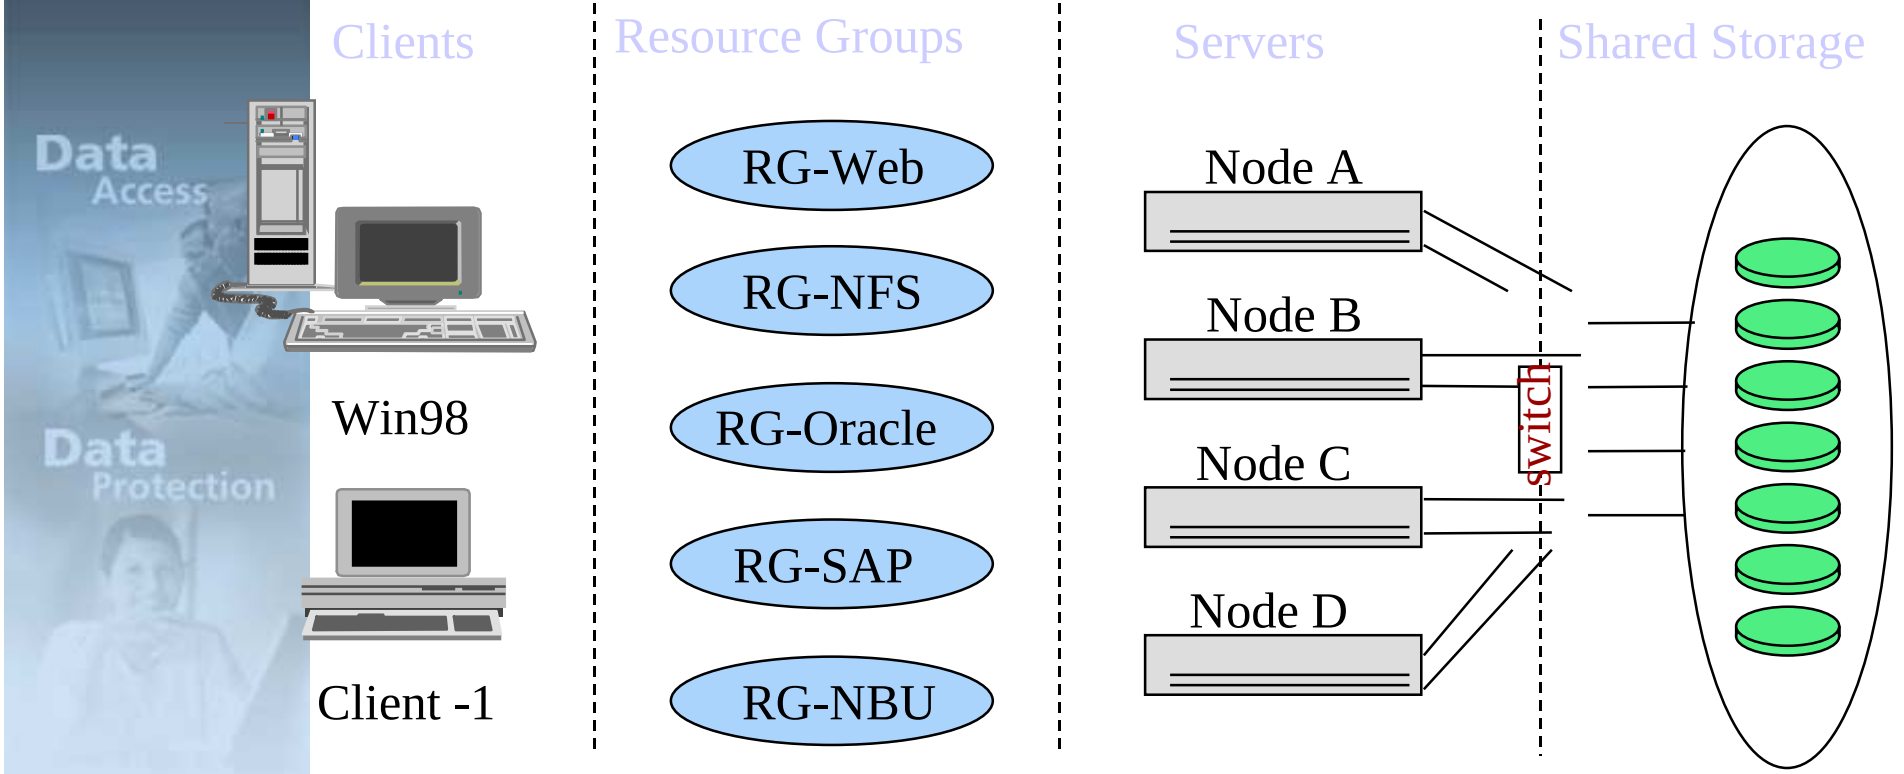

Wide Area Failover/migration

- Monitor and control of multi-heterogeneous clusters at multiple sites
- Disaster recovery between sites (data replicated in real-time)
- Ability to trigger action on a cluster based on events on other clusters
- Follow-the-sun migration

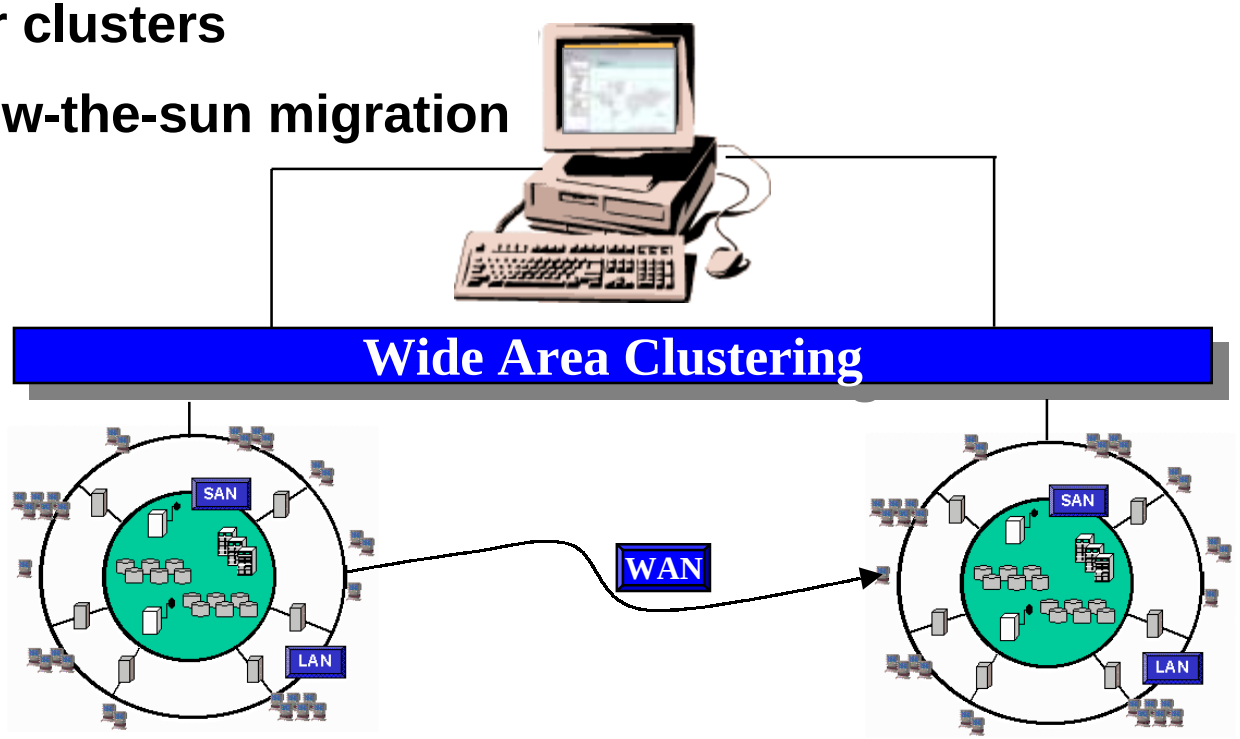

VERITAS™

BUSINESS WITHOUT INTERRUPTION™

Efficient Data Management for Data Protection

Backup Policies for E-Business

- **When to Back Data Up**
  - Windows shrinking, data growing
- **What Data to Back Up**
  - Seldom vs. frequently changing data
- **Where to Back Data Up**
  - local, remote, SAN
- **Optimize Flexibility, Reliability, and Utilization of Resources to Better Meet Service Level Agreements**

Protection Techniques - Backup & Restore

- **Can use recovery manager as the interface to many backup and restore techniques**
- **Online backup**
  - Requires efficient creation of consistent, point-in-time image (snapshot)
    - To backup to other media (disk or tape) & discard
    - To maintain as an online copy than can be restored (rollback)
- **Incremental backup**
  - Only copy data that has changed
- **Fast restoration of a consistent, point-in-time image**

Snapshot Technologies

- **Snapshot = instant-in-time, online image of a database**
  - Used for online backup and rollback
- **Snapshot technologies**
  - Mirror break-off
    - High storage costs (= size of database)
    - Easy to access mirror from other server
    - Examples: EMC Symmetrix, VxVM volume manager (future)
  - Copy-on-write
    - Low storage costs (= size of changed data)
    - Access by other server through cluster file system (future)
    - Example: VxFS file system storage checkpoints

Off-host / Online Backup by Separate Backup Server Using a Split Mirror

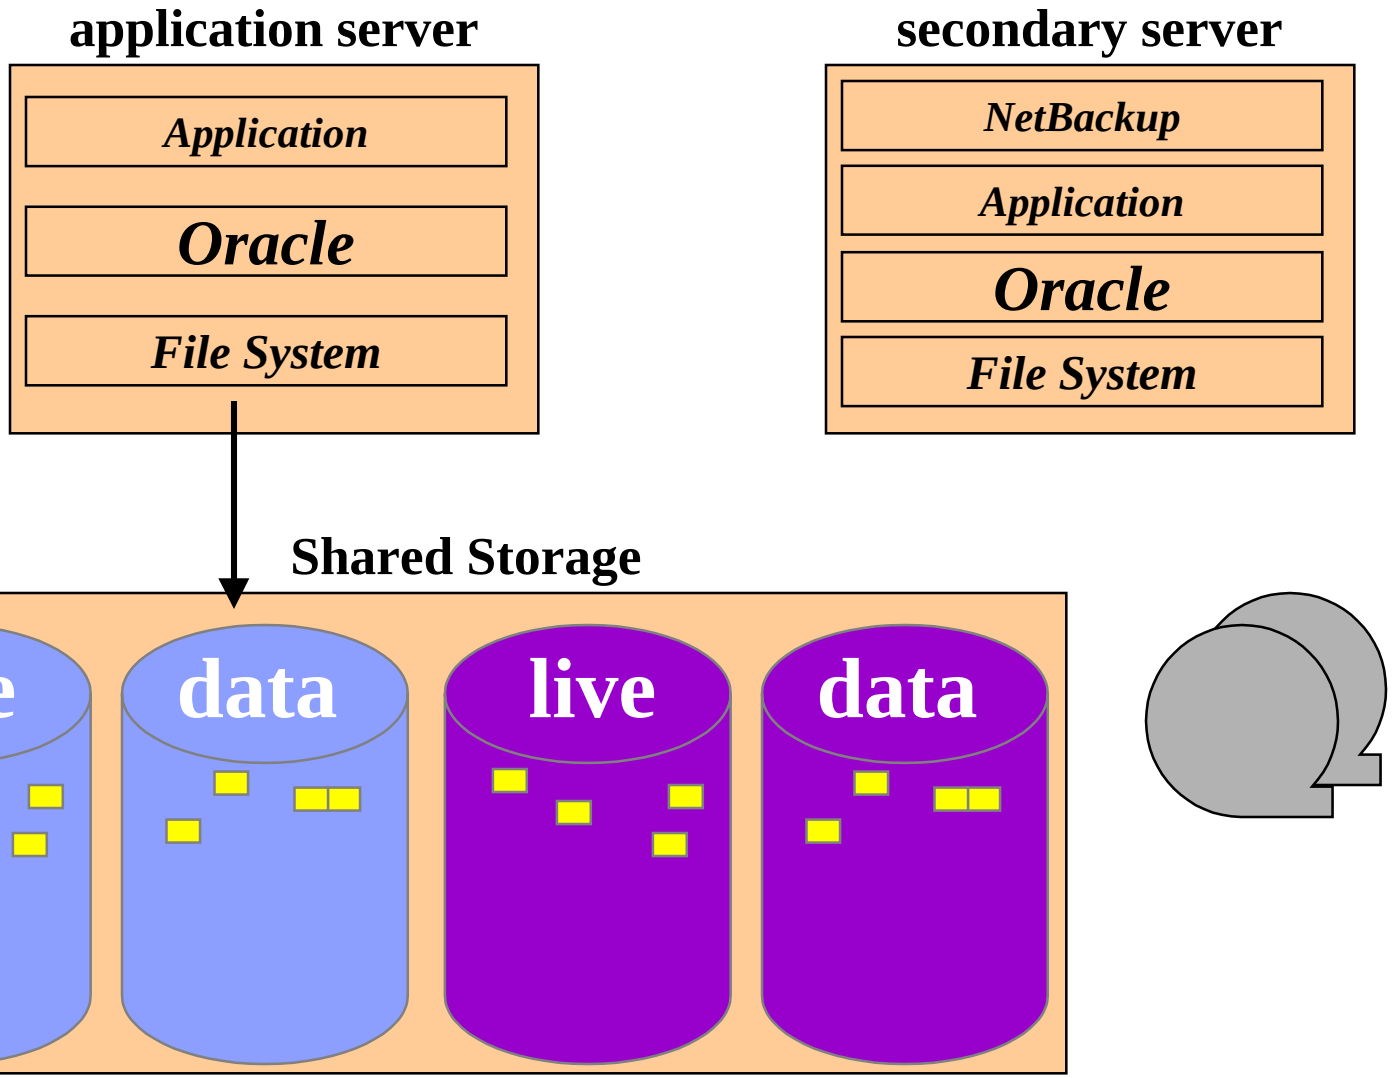

Off-host / Online Backup by Separate Backup Server Using a Split Mirror

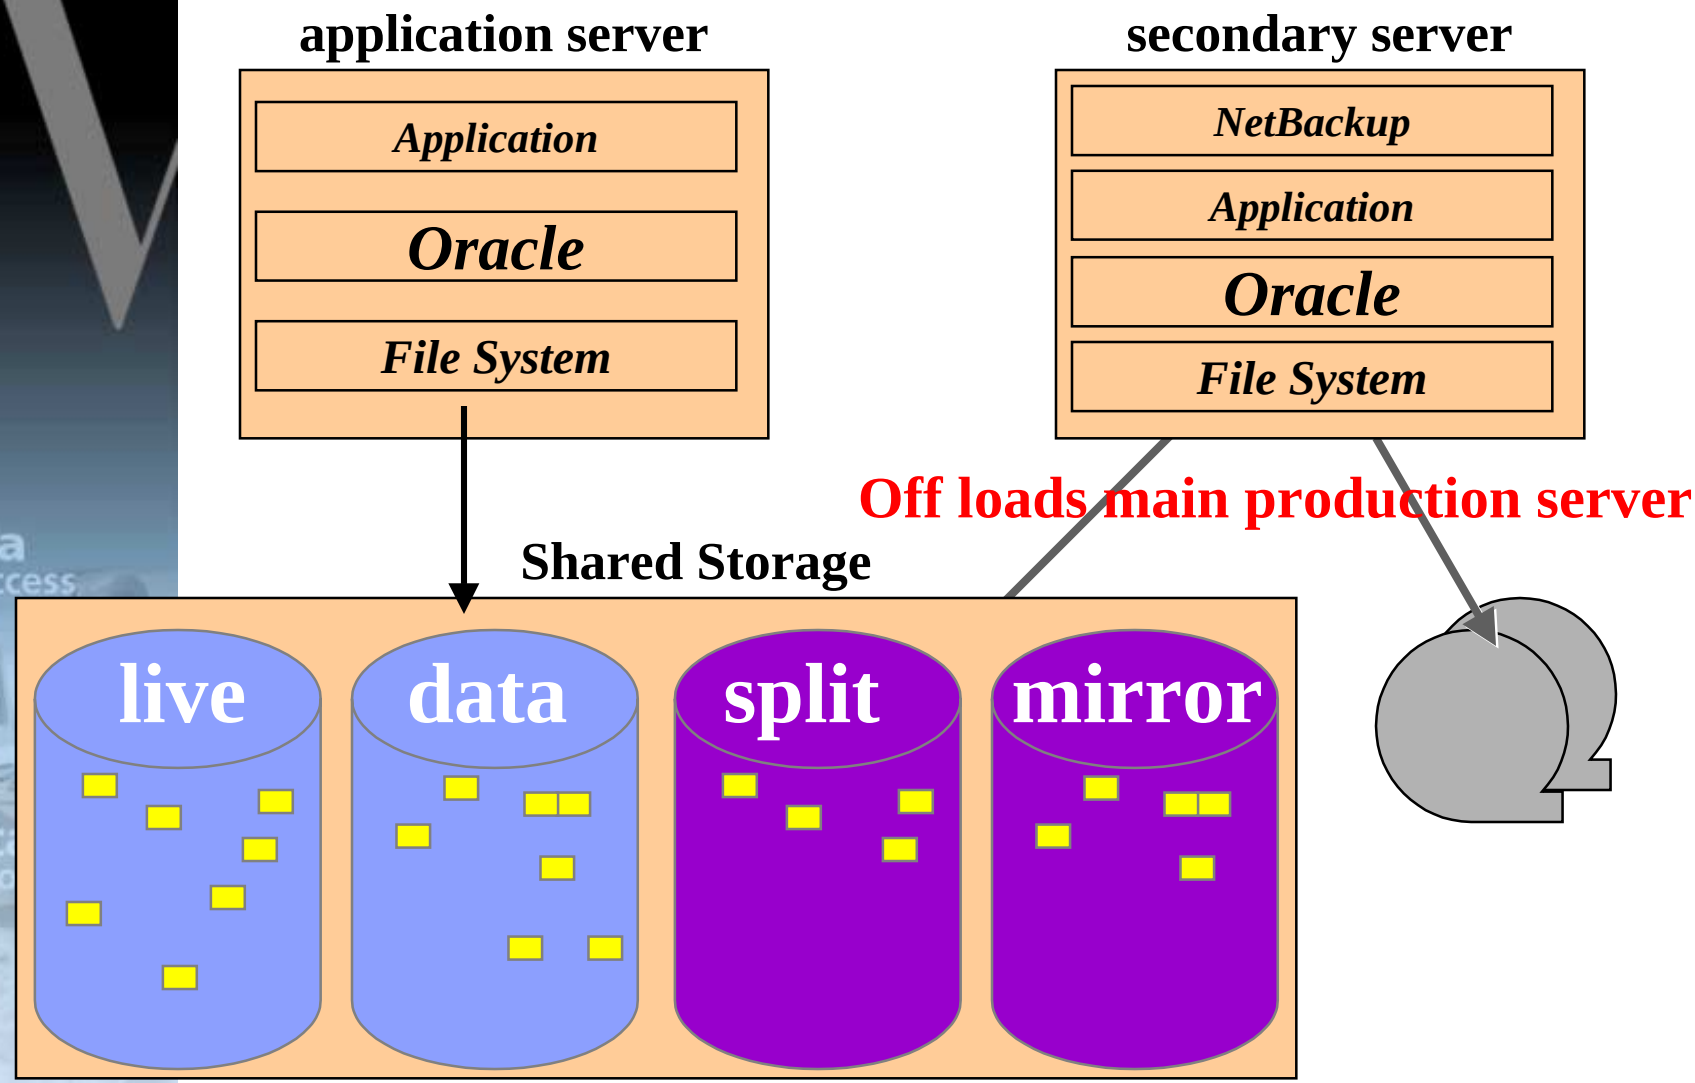

VxFS Storage Checkpoints

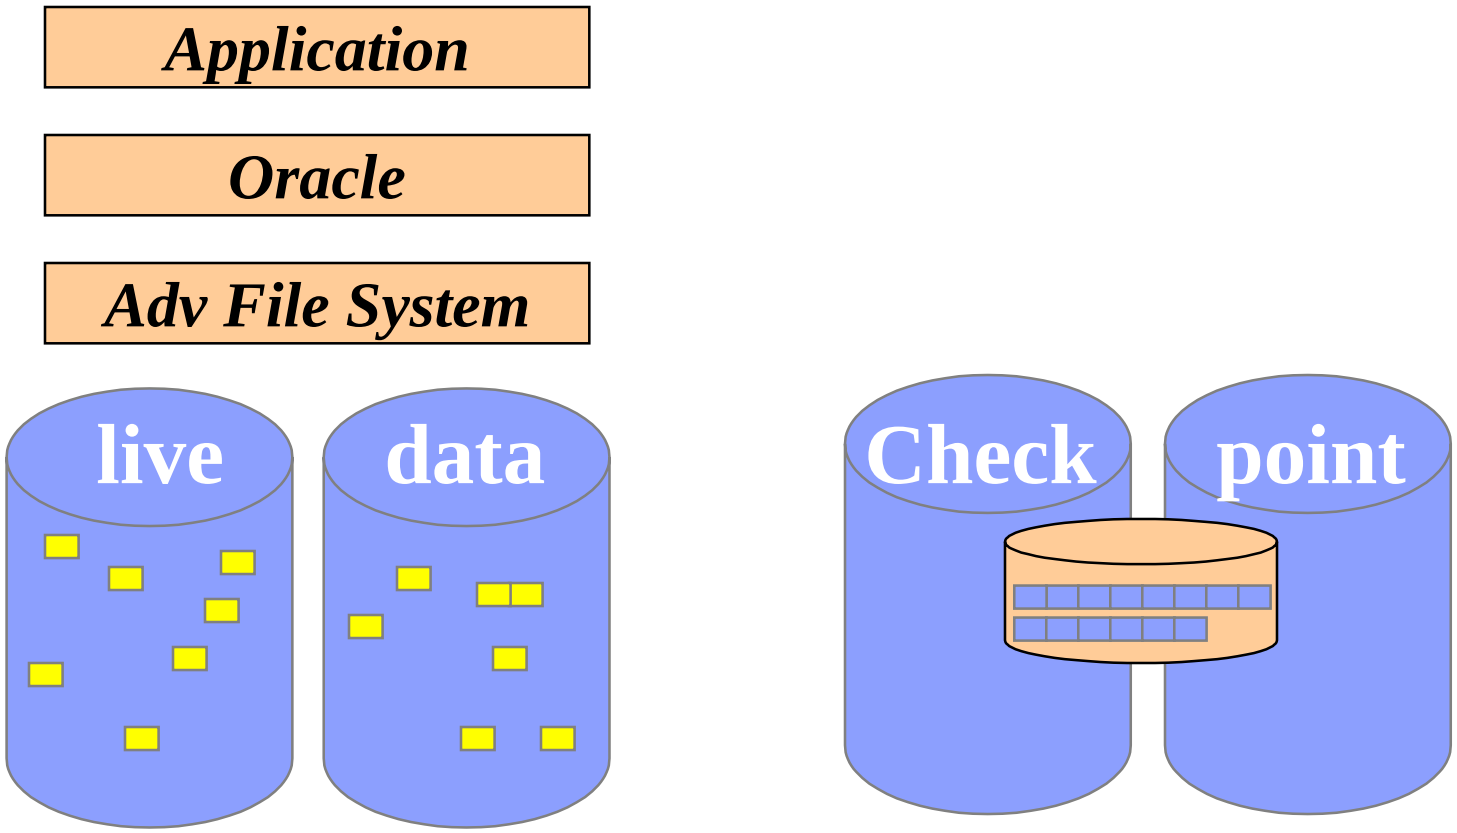

A Storage Checkpoint is a stable snapshot of a Database….

….that only uses a small amount of storage

Online Full Backup

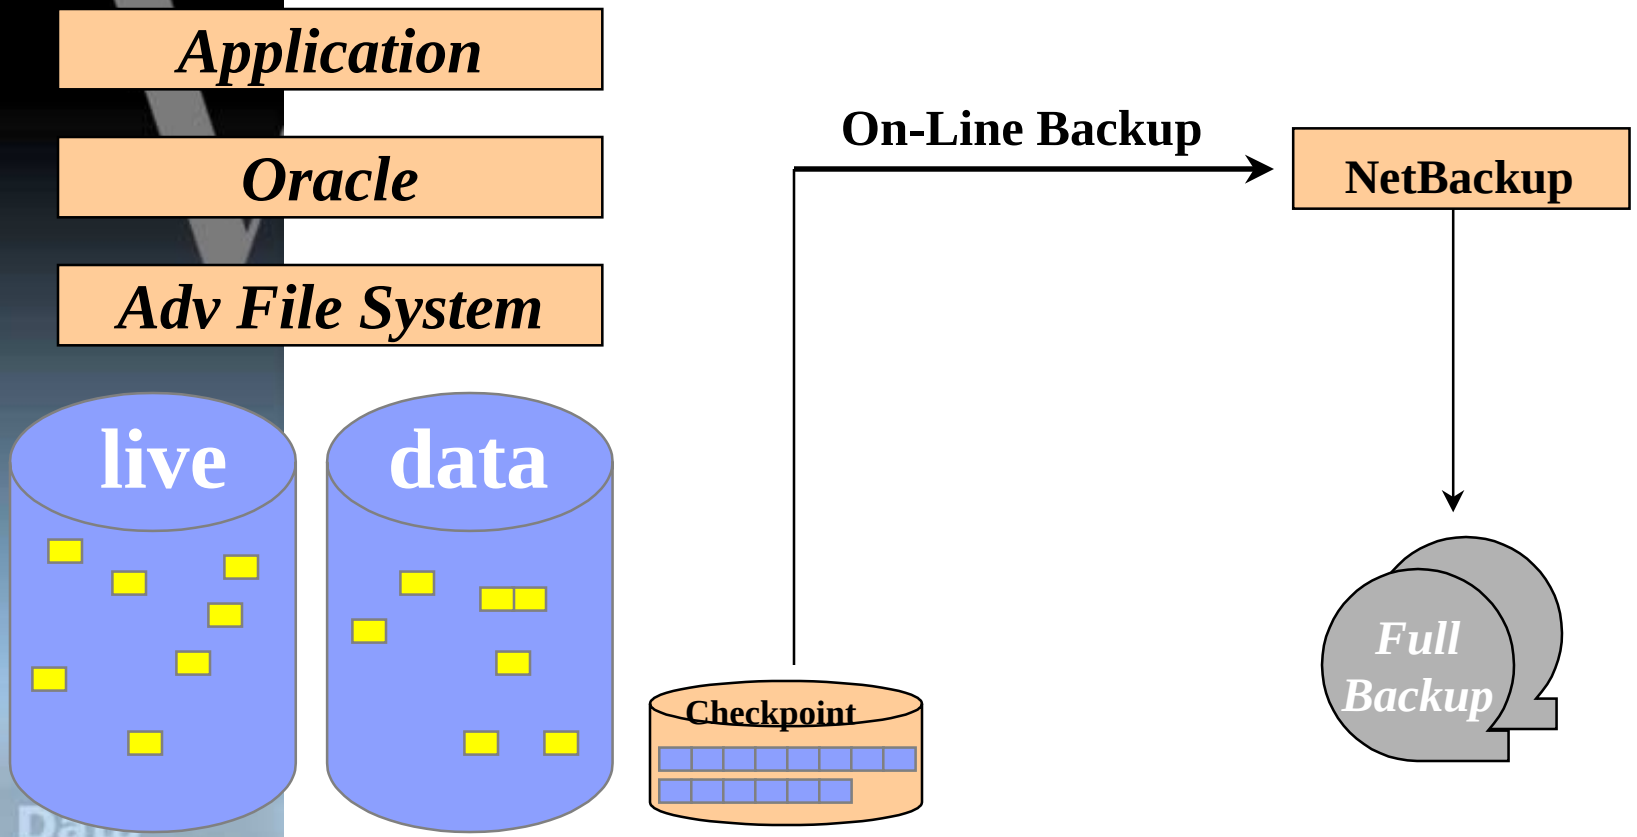

Online Incremental Backup for Oracle

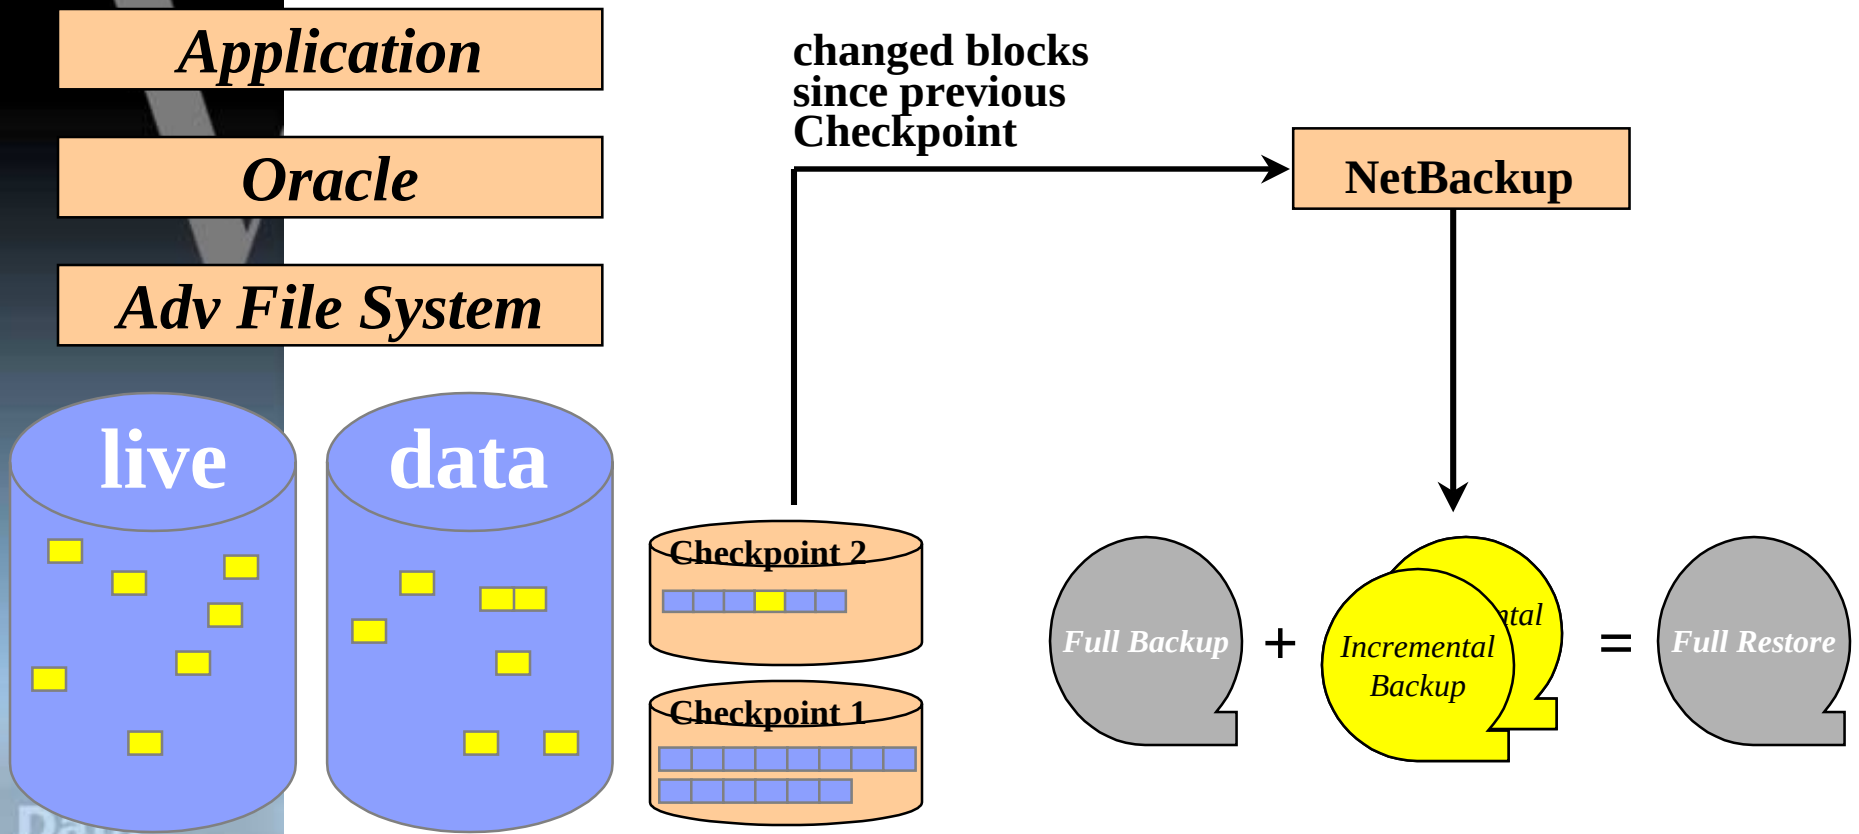

VxFS Storage Checkpoint & Storage Rollback

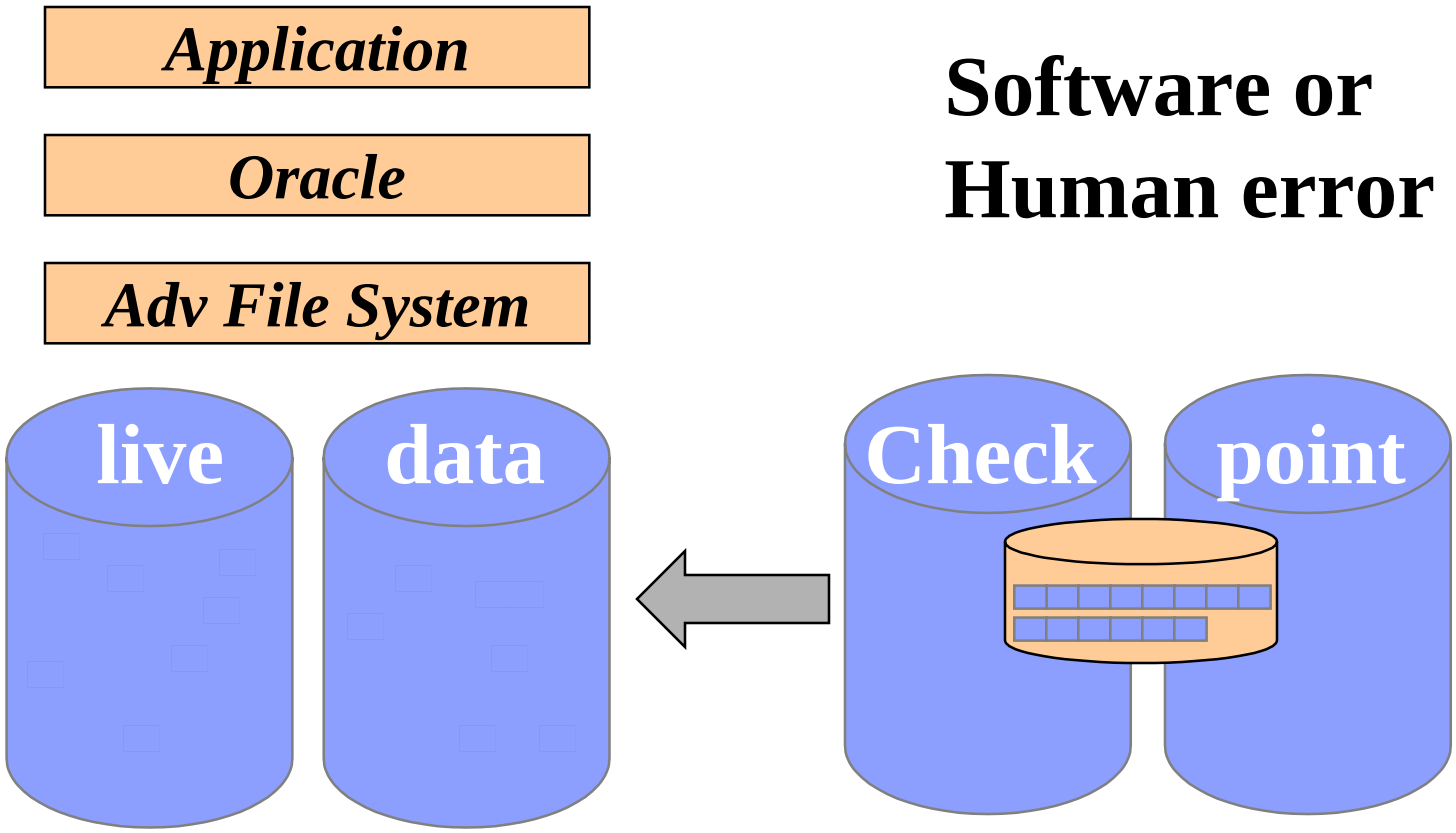

Instantaneous recovery from logical data corruption

Restore Techniques

- **Full restore from tape backup**
  - Very slow restore
  - Low frequency backup -> long log replay
- **Changed data restore**
  - From incremental backup on disk or tape
    - Slow, but protects against media failures
  - From broken-off mirror
    - Instantaneous, but expensive in storage costs
    - Resynchronization in background- resource intensive
  - Rollback from a snapshot
    - Fast recovery from logical data corruption
  - Restore to more recent image means shorter log replay

Storage Management

- **Integration of Data Management, Clustering & Replication, Disk & File Management, Databases, and Hardware Subsystems to Increase**
  - Performance
  - Availability
  - Manageability

of Oracle Databases for E-Business