



# Network**Appliance**<sup>®</sup>

The evolution of storage.<sup>™</sup>

Christopher Slater  
Consulting Systems Engineer

Simplify Application Availability  
and Disaster Recovery





# Designs for Availability and Disaster Recovery

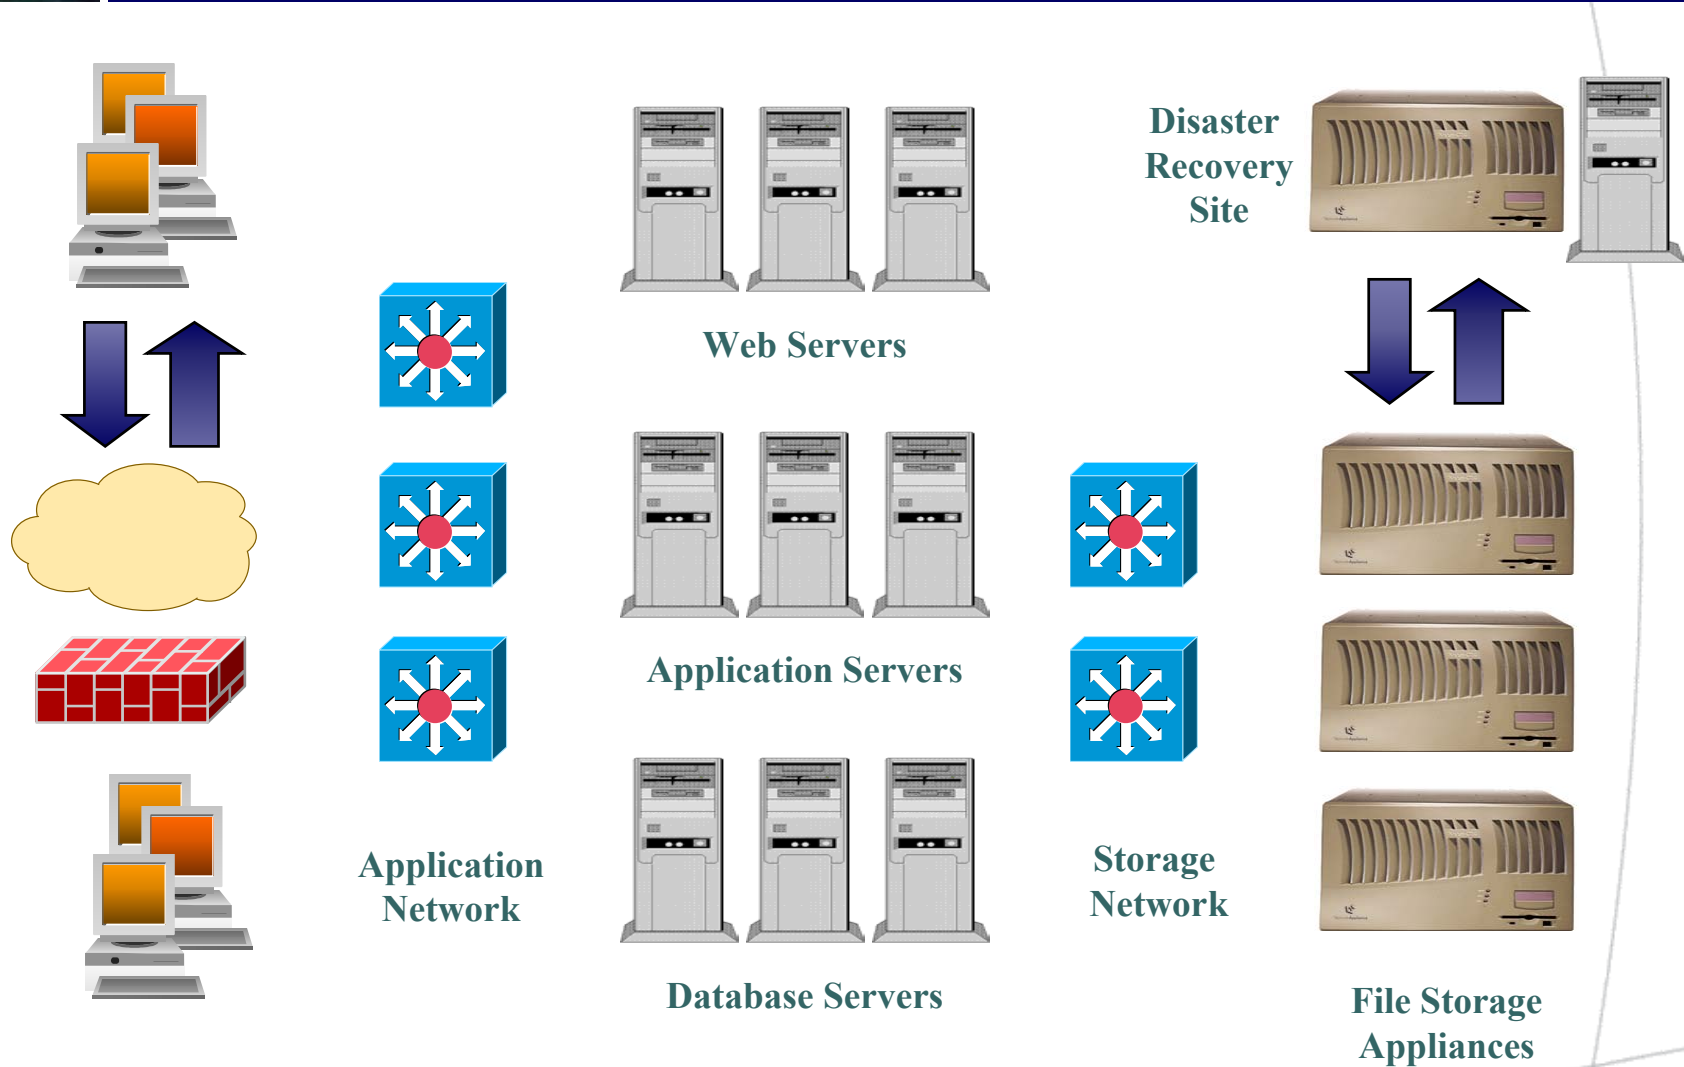
**The purpose of this session is to discuss designs for file system availability, application availability and database recovery**

**I will cover file systems, application clustering, data replication tools and database mirroring tools**

**Our goal is to provide maximum backup and recovery protection with minimal impact on your production environment**




# Data Center Architecture






# Single Site Data Center Model

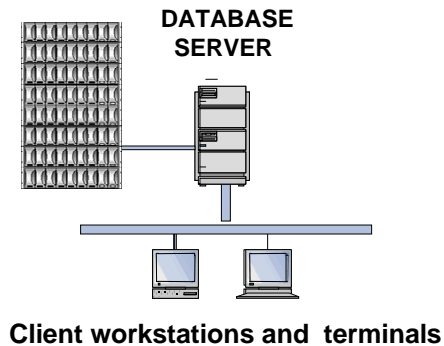


# File System Availability is Critical




# Use a Storage Appliance

## Direct Attached Storage




## Storage Appliance



- **Extra CPU in the architecture that is 100% dedicated to storage management**
- **Dramatic reduction in I/O latency is a key enabler for database environments**

# Unified Storage Appliance



## Heterogeneous Environments



# Network Appliance filer is robust

- **Network Appliance storage filer is highly available**
  - check sums protect each block of data
  - raid groups protect each volume
  - redundant power supplies, fans, FC-AL loops
  - documented 99.996% uptime
  
- **Designed to survive component failure**
  - “self healing” hardware continues to serve file system I/O requests even if storage sub-systems fail
  - auto-support feature automatically engages NetworkAppliance Global Support Center
  - clustered systems available for ultimate redundancy



# Simplify Your Database Storage

## A DBA has many responsibilities

Disk layout and grooming  
Backup design & scheduling  
Upgrade testing & scheduling  
Disaster recovery  
Database tuning  
Application tuning

## Select data storage that simplifies DBA tasks

*Simplify*

*Simplify*

*Simplify*

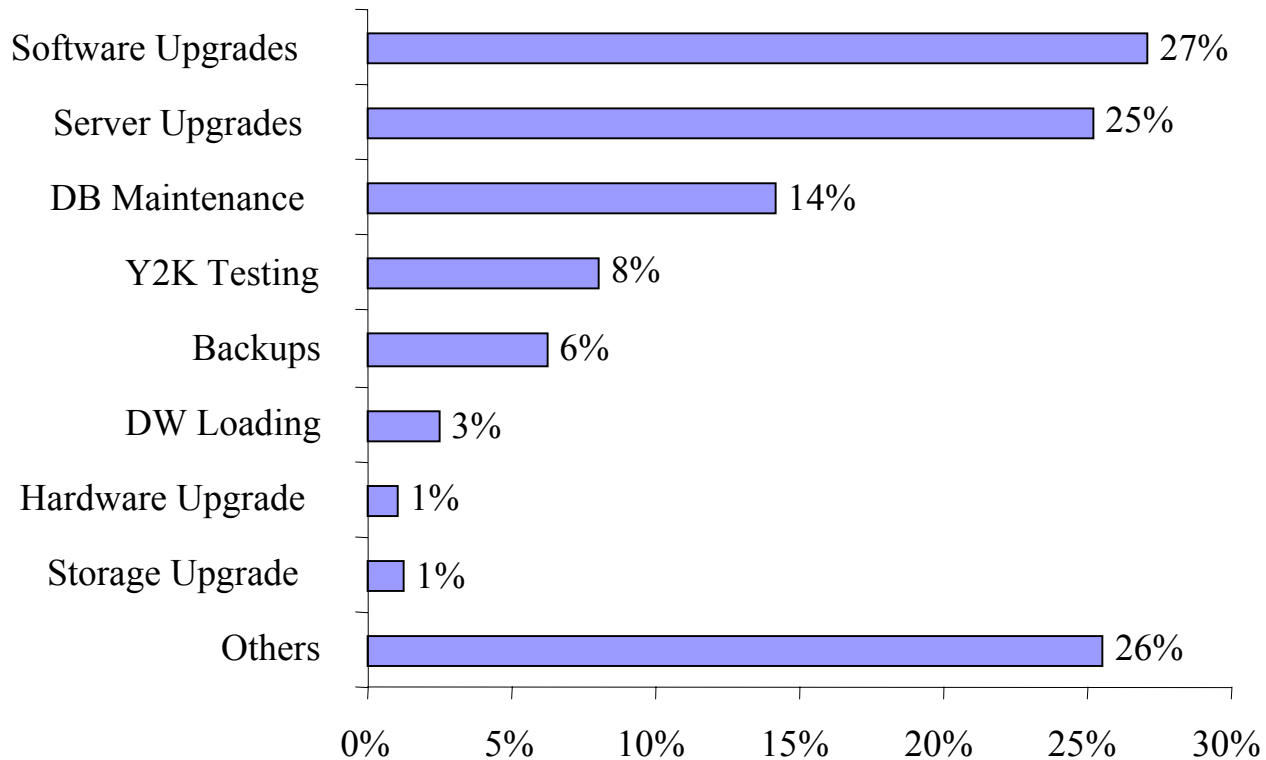
*Simplify*

*Allow focus*

*Allow focus*

# Reasons for Planned Downtime

## Key Reasons for Planned Downtimes...



Source: Strategic Research Corporation, 2000



# Reasons for Unplanned Downtime

- **Operator Errors - 40%**
  - not performing a required task
  - performing a task incorrectly
  - unanticipated mistakes made by inexperienced staff
- **Software Failures - 40%**
  - application and OS bugs
  - performance issues
  - software and server upgrades
- **Hardware Failures - 20%**
  - server, storage, network, and cabling failures
  - environmental factors - heating, cooling and power failure
  - physical disasters - terrorism, fire, flood, earthquake



# Origins of Unplanned Failures

## ➤ **Operators (40% failure rate)**

- System Administrators
- Database Administrators
- Software and Application Engineers

## ➤ **Software Components (40% failure rate)**

- Server OS
- Web server, App server, and Database server software
- File system, Volume management and backup software

## ➤ **Hardware Components (20% failure rate)**

- Firewalls, routers, switches, NICs, load balancers
- Web servers, Application servers, and Database servers
- Disks, power supplies, FC-AL network and switches



# Reduce Operator Errors

## ➤ **Types of Operator Errors**

- not performing a required task
- performing a task incorrectly
- unanticipated mistakes made by inexperienced staff

## ➤ **Improve Skills**

- Hire competent people
- Train your staff to improve competence
- Run fire drills to improve diagnostic and recovery skills

## ➤ **Operations Tactics**

- Document the existing system and procedures
- Increase the number of documented tasks
- Reduce the number of tasks that require specific people
- Automate processes wherever possible
- Employ change management processes

## ➤ **Types of Software Failures**

- application and OS bugs
- performance issues
- upgrades and changes

## ➤ **Failure Detection**

- product hangs or aborts operation
- end-user, operator, or HA tool recognizes failure

## ➤ **Recovery Options**

- operator manually restarts failed application
- HA tool restarts failed application
- remove upgrade and restart application

## ➤ **Types of Hardware Failures**

- server, storage, network, and cabling failures
- environmental factors - heating, cooling and power failure
- physical disasters - people, fire, flood, earthquake

## ➤ **Failure Detection**

- server hangs or aborts operation
- end-user, operator, or HA tool recognizes failure

## ➤ **Recovery Options**

- manually relocate applications to a healthy server
- HA tool redirects client access away from failed server
- repair server hardware and restart applications
- fail-over server applications to a disaster recovery site



# Create a Recovery Plan

## ➤ Ask yourself the following questions

- What is the goal of building an HA system?
- Which systems are business critical?
- What systems cannot be down?
- Which systems can be down and for how long?
- What is the minimum, and maximum, amount of time for fail-over and recovery?
- Can you afford the costs of building and supporting this level of high availability?
- Are you designing to expect all levels of HW and SW failures, including network cabling and power failures?
- Do you have a Disaster Recovery Timeline?

# Disaster Recovery Timeline



**1) Detect failure, choose corrective actions**

**2) Synchronize data**

**3) Recover applications**

**4) Redirect users to fail-over systems and/or site**

## ➤ **Application Clustering tools**

- Hewlett-Packard MC/ServiceGuard
- Integratus UHA
- Microsoft Cluster Services (MSCS)
- Oracle Fail Safe
- **Oracle 9i RAC**
- Sun Cluster
- Veritas Cluster Server (VCS)

## ➤ **Database Mirroring Tools**

- File system and volume mirroring tools
- Oracle Standby databases
- Oracle Advanced Replication
- Quest Shareplex

## Four designs for highly available systems

1. Single server with highly available file system
2. Multiple servers, application clustering solutions, and highly available file system
3. Multiple servers, application clustering, server hardware clustering, network redundancy, and highly available file system
4. Use database and file system mirroring tools to create a disaster recovery site



# Design #1: The File System

## ➤ Purpose

- A robust file system and storage solution is the foundation to your highly available system

## ➤ Decision

- Choose a file system that is based on storage hardware redundancy and offers 99.99%+ file system uptime

## ➤ Recovery

- File system tools (on-line backups, off-line backups and mirroring) provide file system recovery
- Manual process required to restart an application or application server

## ➤ **Create on line backups**

- Reduce downtime for cold database backups to minutes
- Reduce downtime for hot database backups to seconds

## ➤ **Reduce time to recover a database**

- Recover multi-terabyte file systems in seconds
- Peace of mind for major maintenance tasks like planned application upgrades and data warehouse rebuilds

## ➤ **Expand data access by mirroring**

- Duplicate an on-line production database to create a duplicate database for reporting, QA, development and backup
- Mirror to create a Disaster Recovery site
- Backup remote offices by mirroring back to the data center

## ➤ **200GB Database Recovery with Tape**

- Shutdown database, search for tapes
- Tape recovers 300GB at a rate of 60 GB/hour
- Startup database and replay archived redo log files
- Recovery time is 3 hours 20 minutes + log replay time

## ➤ **Database Recovery with Mirror**

- Shutdown database
- Switch to mirrored copy of the database
- Startup database and replay archived redo log files
- Recovery time for 200GB database is time to mount mirror + log replay time



# Design #2: Application Availability

## ➤ Purpose

- Provide automated recovery for application failures

## ➤ Decision

- What kind of application failures do you intend to correct?
- Do you want to automatically recover the database, or just the application server?

## ➤ Recovery

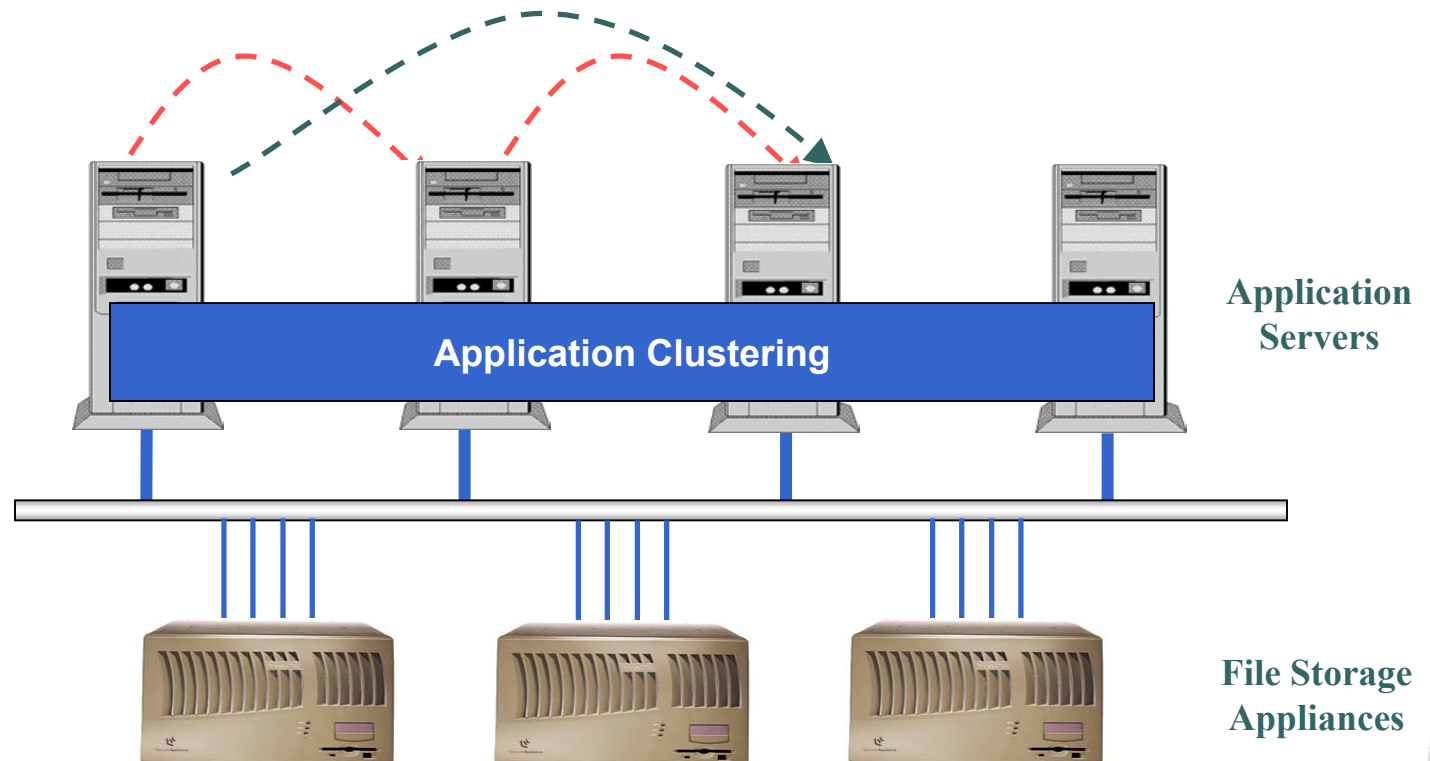
- Recovery should be automated for application servers
- Recovery of the database server can be automated
- Recovery of a corrupted database will probably involve your DBA

# Application Clustering Products

## Oracle 9i RAC

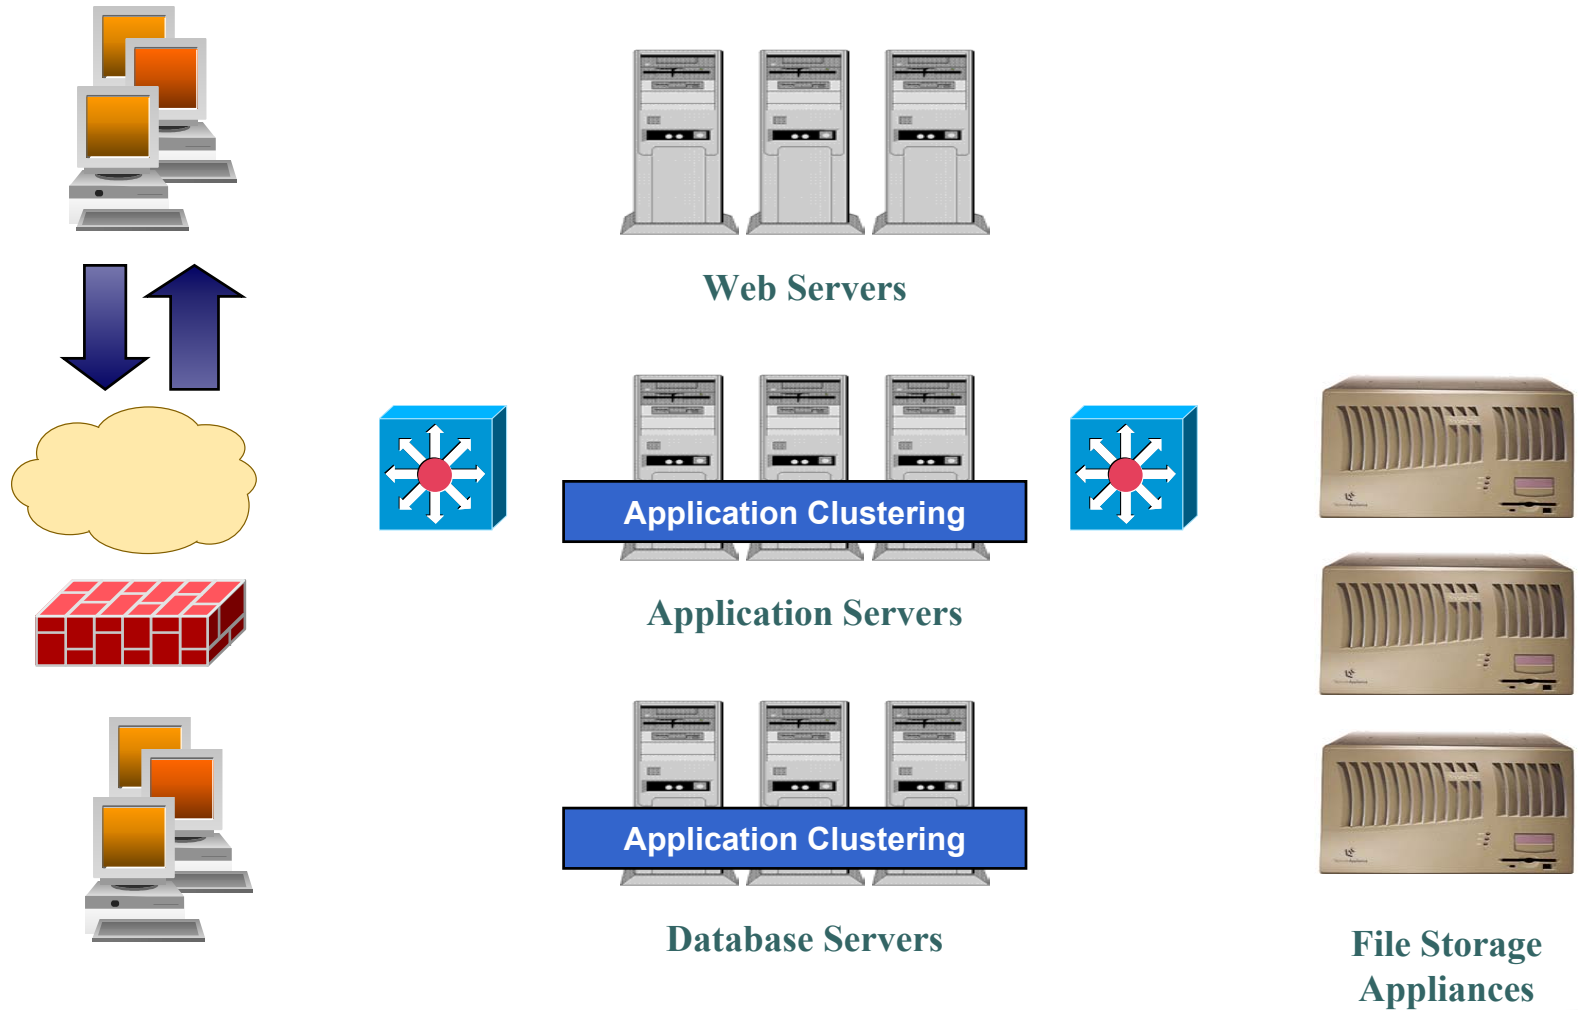
Hewlett Packard ServiceGuard

Integratus UHA - Veritas VCS



*Network Appliance Data Availability*

# Application Clustering





# Oracle, Sun E10K, and App Clustering


## Airline passenger / freight route mileage and profit tracking system



Clients



2 Sun E10K servers  
running 10 Oracle  
database instances  
with application  
clustering



File Storage  
Appliances  
containing 4TB of  
useable storage



# #3: Server and Network Availability

## ➤ Purpose

- Provide seamless fail-over to redundant network components, application servers, and database servers

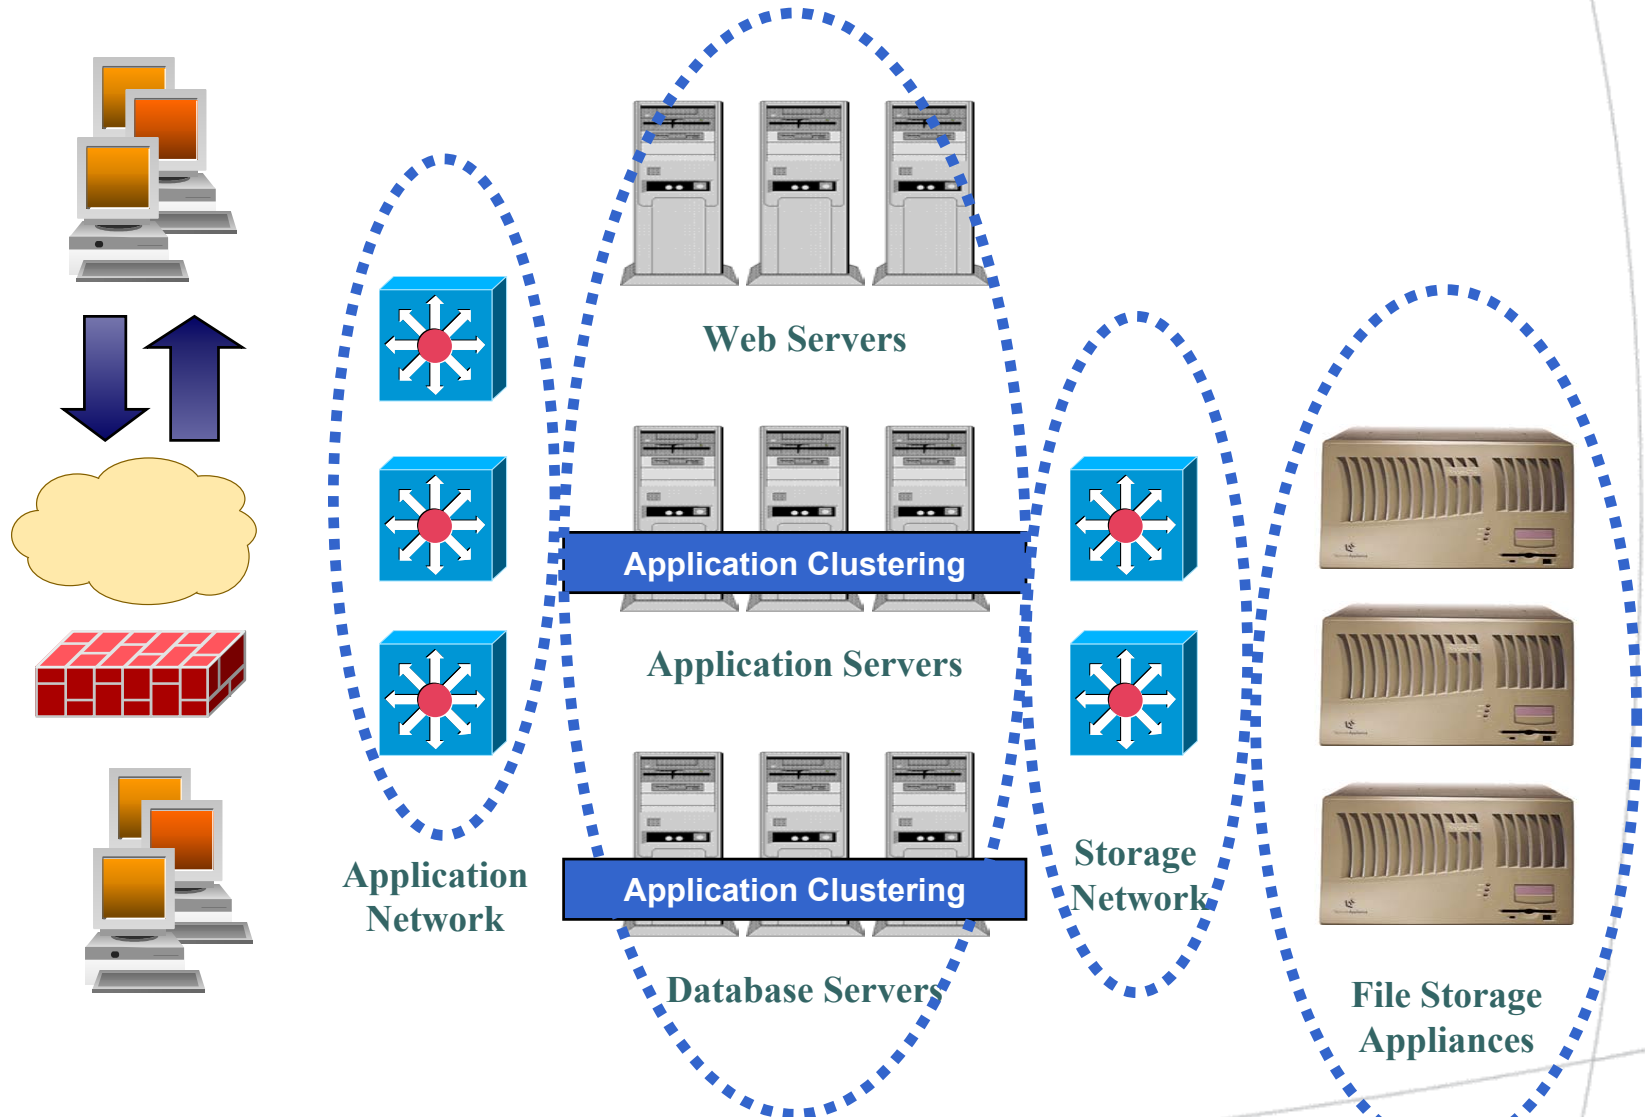
## ➤ Decision

- Which components automatically fail-over?
- Which components require clustering software?
- *Don't forget* - use power from separate power grids

## ➤ Recovery

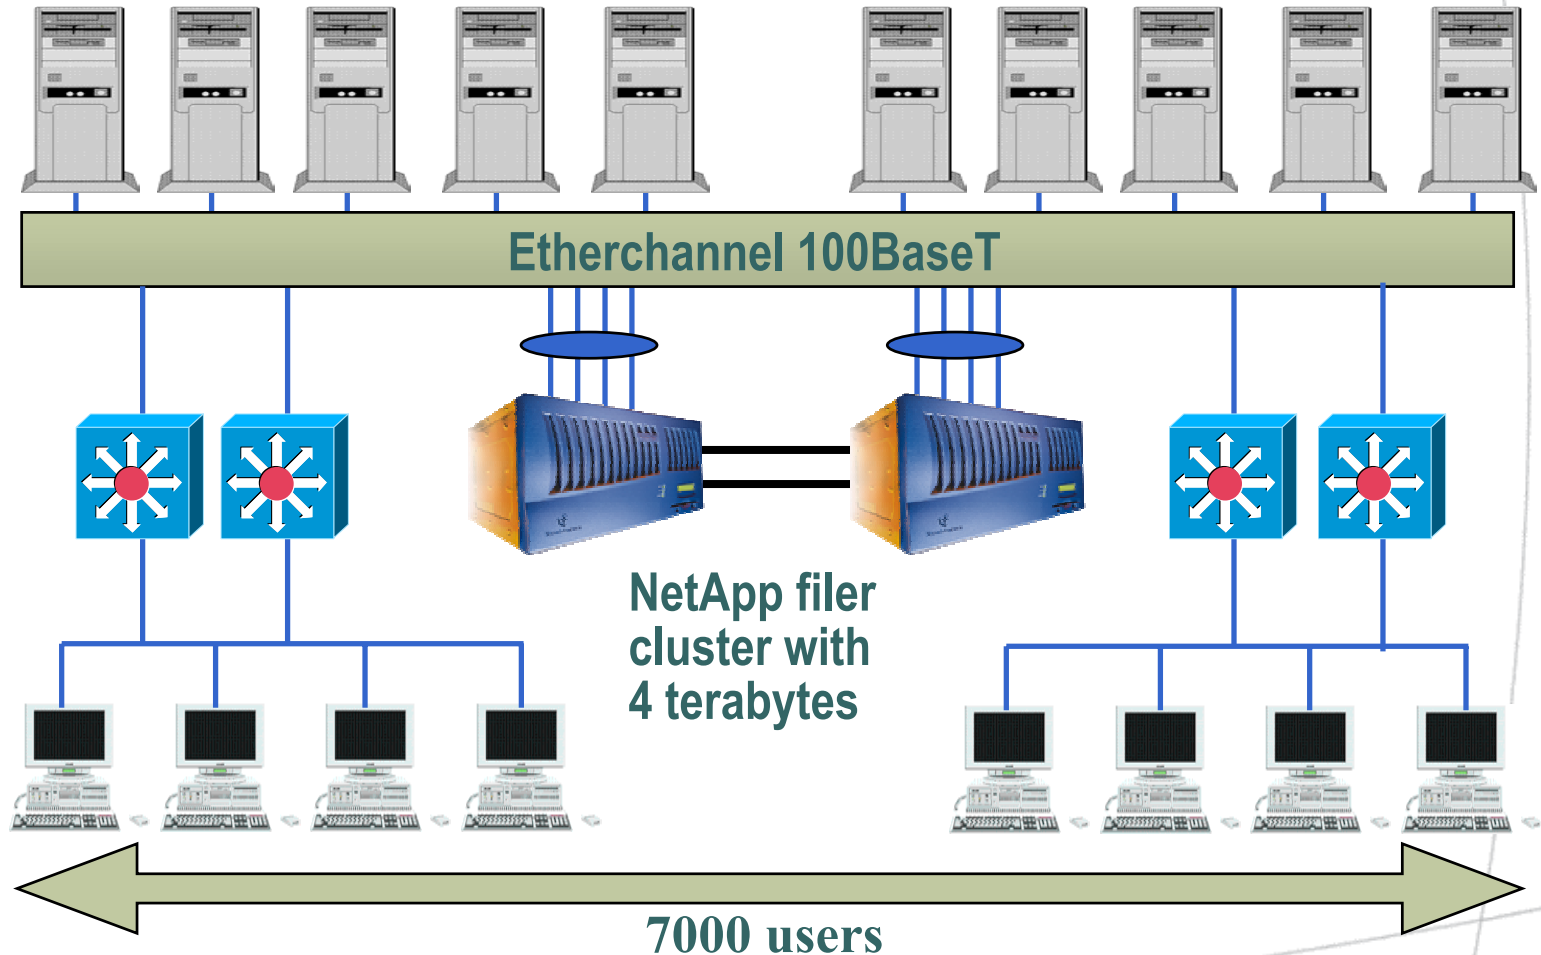
- Network and storage appliance fail-over is completely automated
- Use load balancers to redirect client access to surviving application servers
- Use database clustering tools to fail-over databases to surviving servers

# Hardware and Network Clustering



# SAP deployment at DBIS

15 SUN & HP servers running 15 Oracle instances to support SAP/R3





# Design #4: Data Center Availability

## ➤ Purpose

- Provide fail-over to a remote data center or site that can replace the local data center

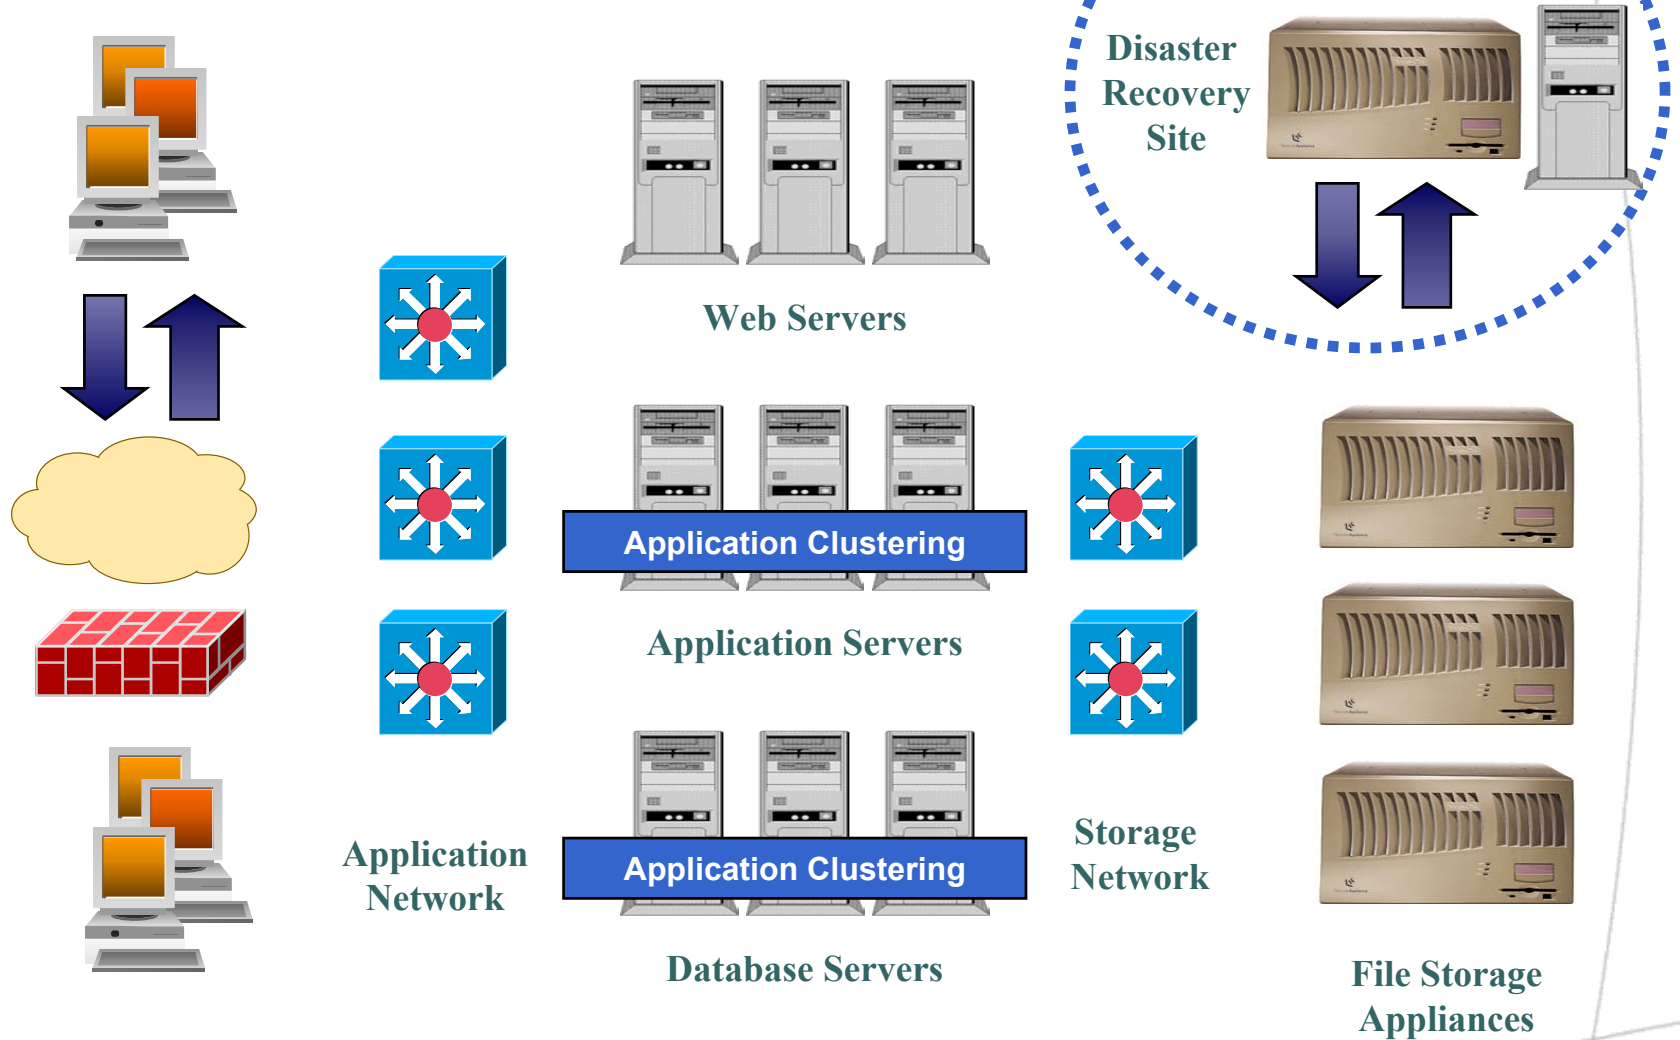
## ➤ Decision

- Do you duplicate the entire data center, or just critical portions of the systems in the main data center
- How do you mirror data to the remote site and how often?
- Do you use Synchronous or Asynchronous mirroring?

## ➤ Recovery

- Complete the task of synchronizing data
- Some portions of the application fail-over can be automated, but probably not database recovery
- Bring the remote data center on-line
- Redirect client access to that data center

# Mirroring for Recovery





# File System mirroring

## ➤ **Replicate the file system**

- Synchronous mirroring is used within a single data center
- Asynchronous replication is used between data centers
- Mirror the file system AND on-line recovery options
- The mirrored file system must be consistent at all times
- Mount the mirrored file system for read only access

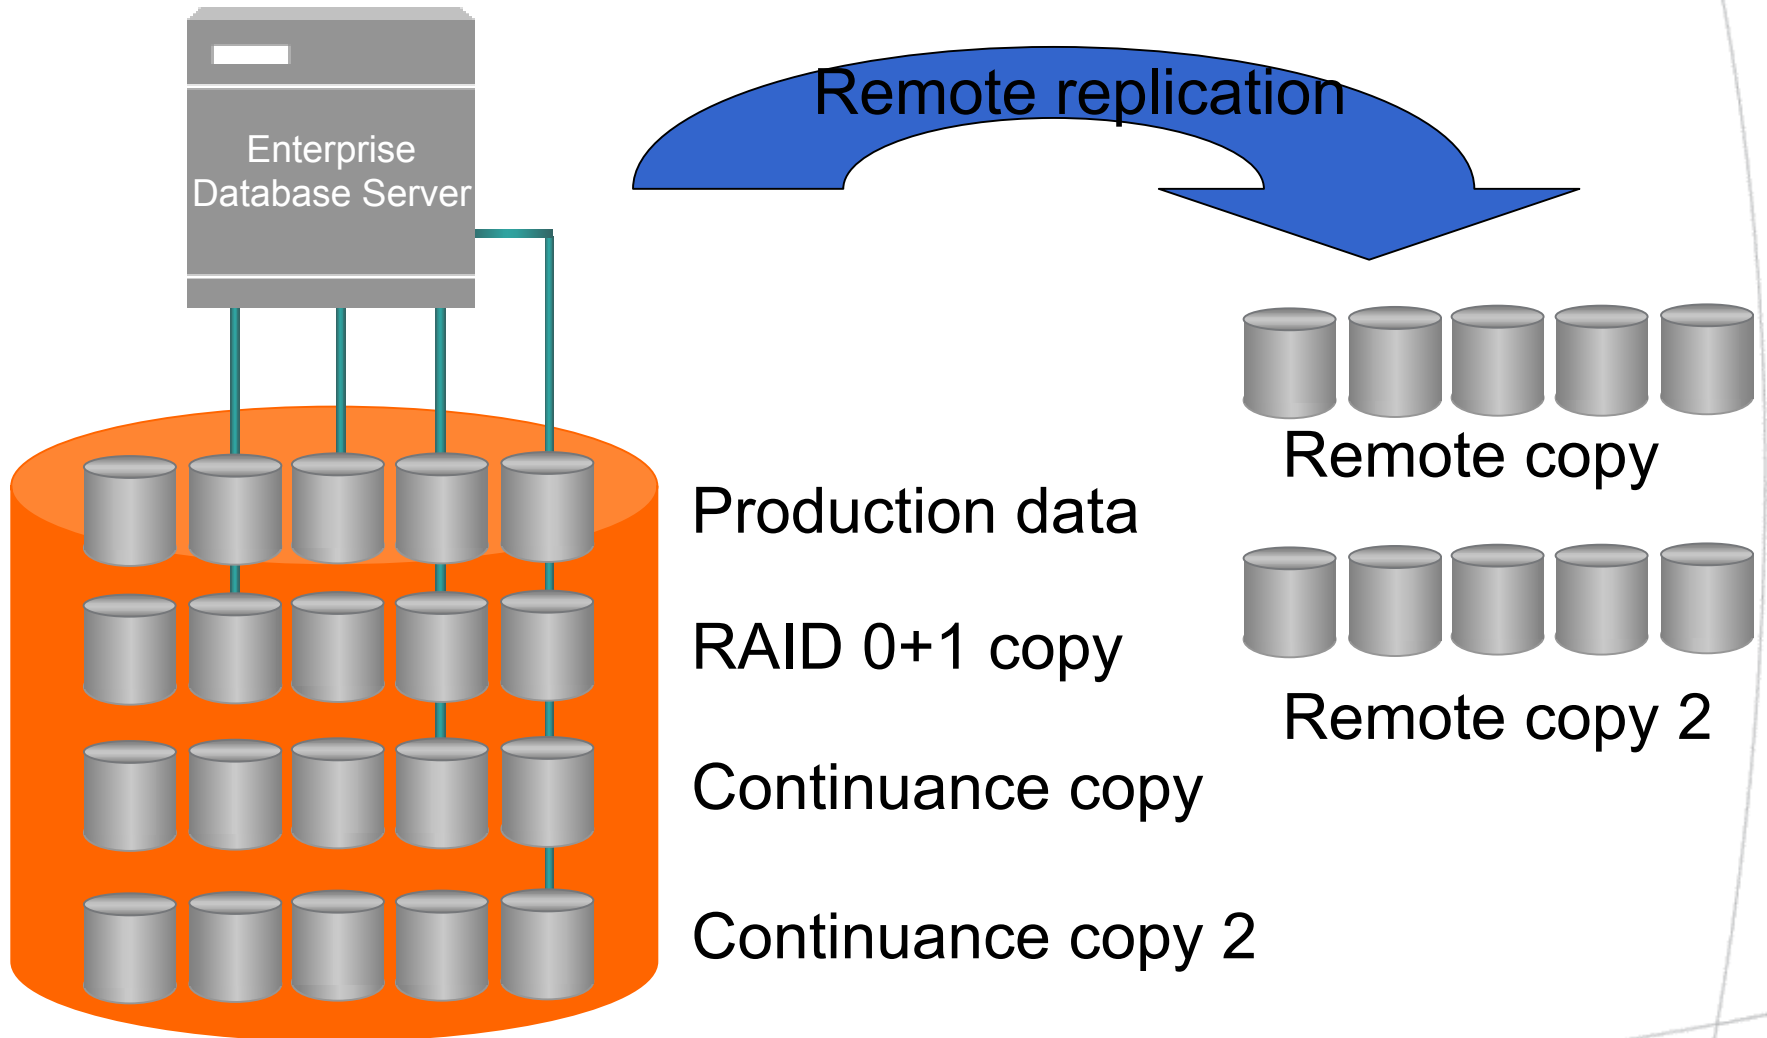
## ➤ **File system mirroring is a powerful tool**

- Create near line backup and recovery sites
- Create Disaster Recovery and QOS sites
- Provide hardware maintenance windows
- Mirror remote systems to a central location for backup
- Duplicate on-line data for QA testing and development systems
- Duplicate on-line data for data warehouse and reporting databases

## ➤ Database Replication Tools

- **File system and volume mirroring tools**
  - Database files are mirrored from an active database to a remote site
  - The mirror must be broken before write access is available, and then the database can be recovered
- **Oracle Standby databases**
  - Standby database applies transaction logs
  - Database is mounted, but not open
- **Oracle Replication options**
  - Basic replication makes read only access to tables that originate from a primary site
  - Advanced replication supports read and update table replicas at many sites

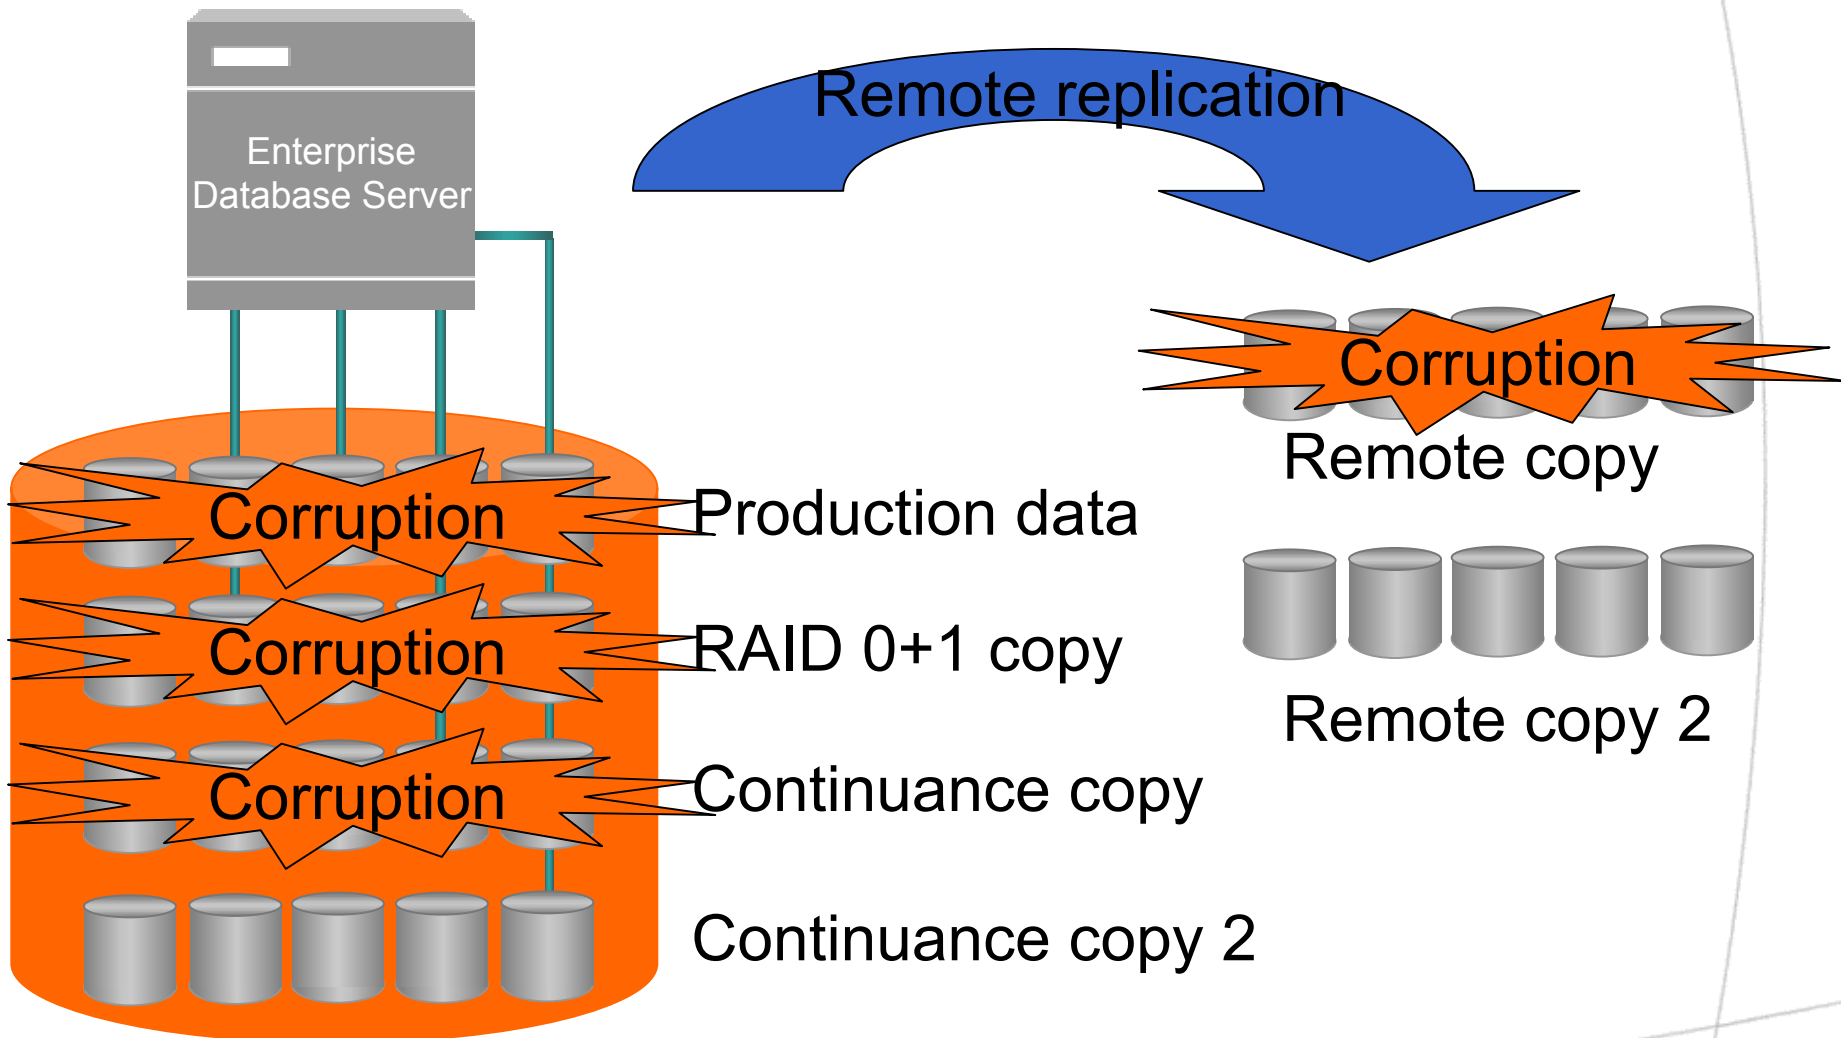
# Synchronous Replication Problem



*Single box eases management*

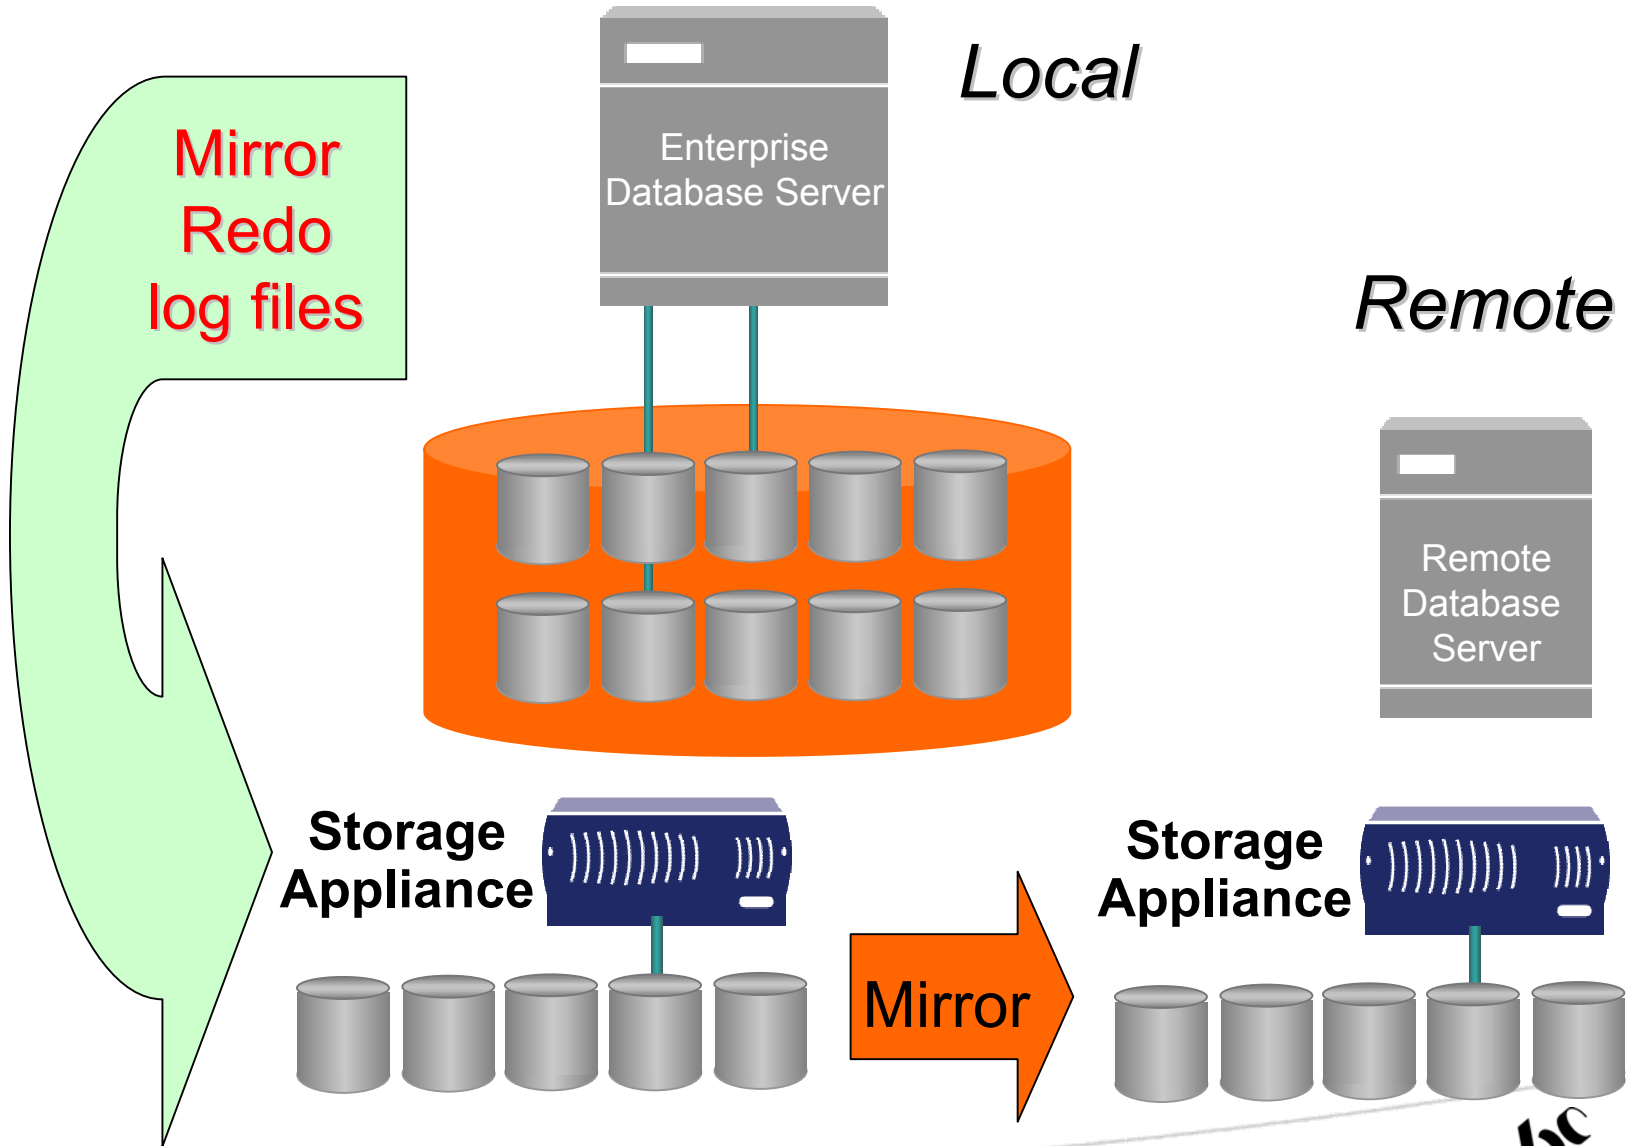


# Risk of corruption



*Single box eases management*

# Alternative Solution



- **In closing, here are some things you can do right now**
  - **Look for ways to reduce operator errors**
  - **Document application recovery options**
  - **Document server recovery options**
  - **Create a recovery time line**
  - **Hold practice recovery drills**



# Network**Appliance**<sup>®</sup>

The evolution of storage.<sup>™</sup>

Christopher Slater  
Consulting Systems Engineer

Simplify Application Availability  
and Disaster Recovery

