

E-Business Continuity with Oracle Parallel Fail Safe

Shishir Agarwal

Oracle

500 Oracle Parkway

Redwood Shores

California 94065

Tel: 650-607-2077

Email: shishir.agarwal@oracle.com

\$1800

The estimated cost of down time per second
for an e-business

7 Seconds

The amount of time it takes to click over to a competitor's web site

"Top 3" E-Business Objectives for IT Management

1. Availability
2. Availability
3. Performance

Downtime Bad News for e-business

BUSINESS

ZDNN TECH NEWS NOW

eBay outage going, going ... gone?

The largest online auctioneer says its outage has been traced back to Sun Microsystems software -- and that a fix is at hand.

By [Margaret Kane](#), ZDNN
June 11, 1999 7:22 AM PT

UPDATED 2:17 PM PT

BREAKING
June 11,
03:09p
Not all M

eBay officials say they expect the site back up and running "before the sun goes down."

Amazon.com Site Shut For 10 Hours Wednesday

SEATTLE (Reuters) - Buying books over the Internet took longer than usual on Wednesday morning when [Amazon.com Inc.](#)'s Web site went down for 10 hours.

Current quotes (delayed 20 mins.)

[AMZN](#) 102 5/8 -9 (-8.06%)

The popular online book and CD retailer blamed the interruption on some planned maintenance, but said the work lasted about twice as long as planned.

"It was for routine maintenance. We do it regularly," said an Amazon spokeswoman. She said the last time the site had been taken down was in January.

Amazon said it collected the e-mail addresses of all the customers who tried to visit the site during the interruption, and

AmeriTrade, Waterhouse Lose Online Trading

NEW YORK (Reuters) - Equipment failures at AmeriTrade and Waterhouse Securities stymied online traders for a short time Thursday, a day the market dropped more than 3 percent.

An AmeriTrade Holding Corp. representative said that the company experienced a "brief outage" in online trading after another failure in a part of its system that has caused problems for

replace the part of the system, and handle excessive volume like that said.

Downtime Adds Up

2.8 hours per week

\$39,950 per

hour
\$5,000,000 per

year

Source: MERIT Project

Downtime Costs Add Up

- **America Online**

August 1996 Outage: 24 hours
Maintenance/Human Error
Cost: \$3 million in rebates


- **E*Trade**

February 1999 through 3 March 1999
Four outages
Cost: 22 percent stock price hit on 5
February 1999

- **eBay**

June 1999 outage: 22 hours OS Failure
Cost: \$3 million to \$5 million revenue
hit
26% decline in stock price

Costs of Downtime



Source: Gartner Group and Contingency Planning Research

Database Availability

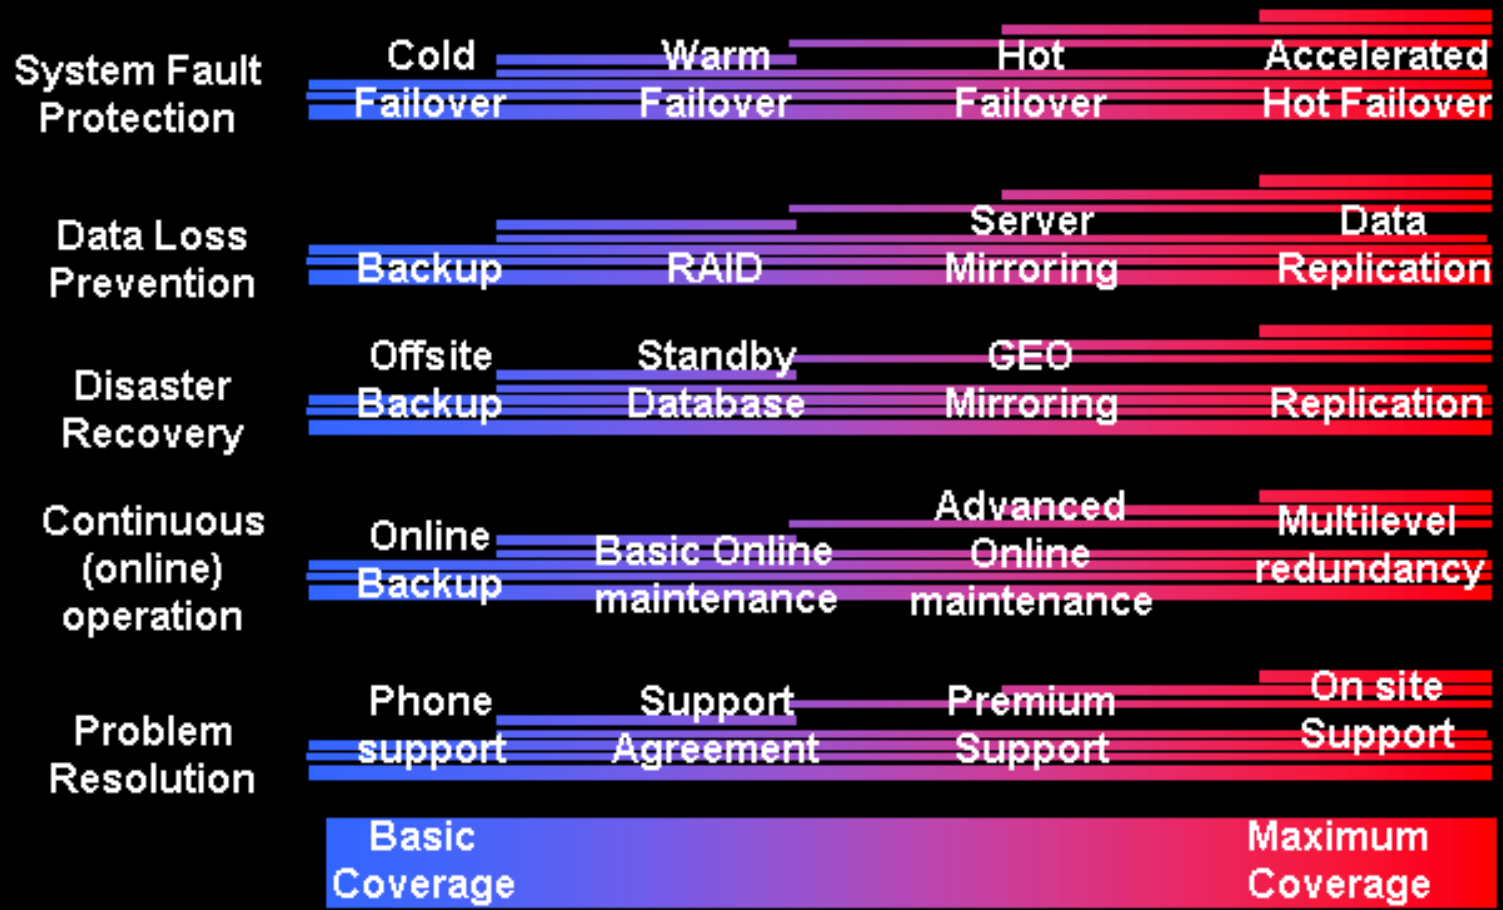
- The database is the most important component in the application stack
 - Database controls critical business data
 - Critical business data is irreplaceable
 - Database can mask lower level failures from applications
- Oracle database has been architected to ensure:
 - Application data availability
 - Application data consistency

Availability Substructure

The Means To The End

- There are three fundamental components of application availability
 - Reliability
 - The ability of the application to avoid or mitigate failures
 - Recoverability
 - The ability to reestablish application availability in a timely manner once a failure has occurred
 - Continuous Operation
 - The ability to perform maintenance and recovery operations while the application is in production, and online
- The Oracle8i architecture address all three components

High Availability Spectrum



Continuous Operation Spectrum

Oracle8i Advanced Replication

Online Index re-org, defrag, rebuilds, table moves

Parallel Everything (insert, update, load, index, join, sort ...)

Data Partitioning (range, hash, composite)

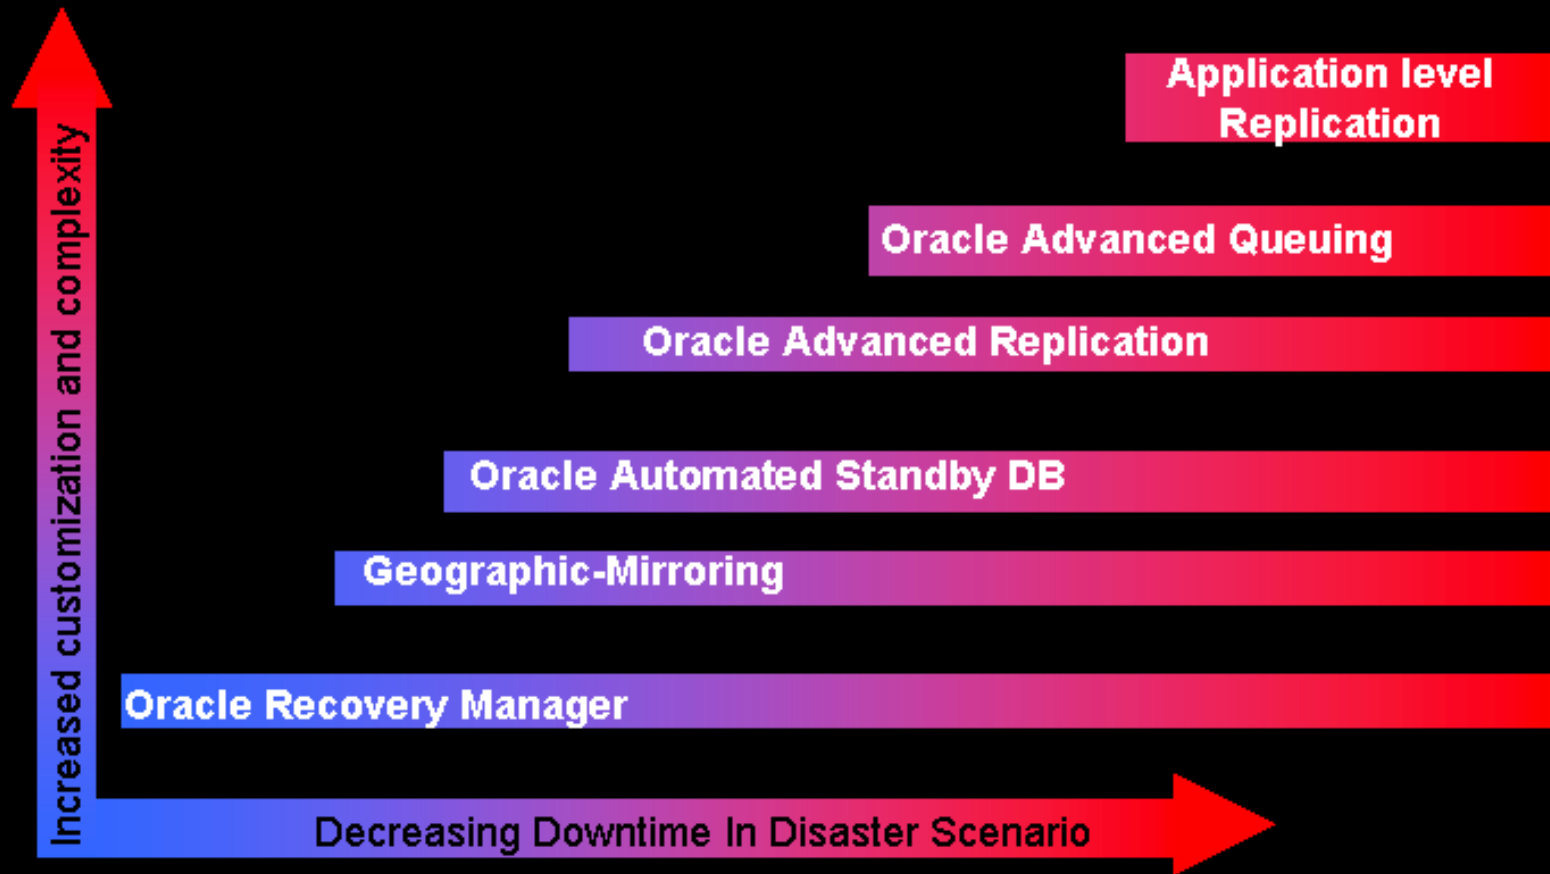
Online Backup – Oracle8i Recovery Manager (RMAN)

Oracle Enterprise Manager - Capacity Planning and Event Notification

Basic Online
Operations

Continuous
Operation

Disaster Recovery Spectrum



System Fault Recovery Spectrum



Oracle Parallel Fail Safe

Oracle Parallel Server

Oracle Fail Safe 3.0

Clustering

Oracle FastStart Recovery and Transparent Application Failover

RAID and hardware safeguards

Cold
Restart

Cold
Failover

Warm
Failover

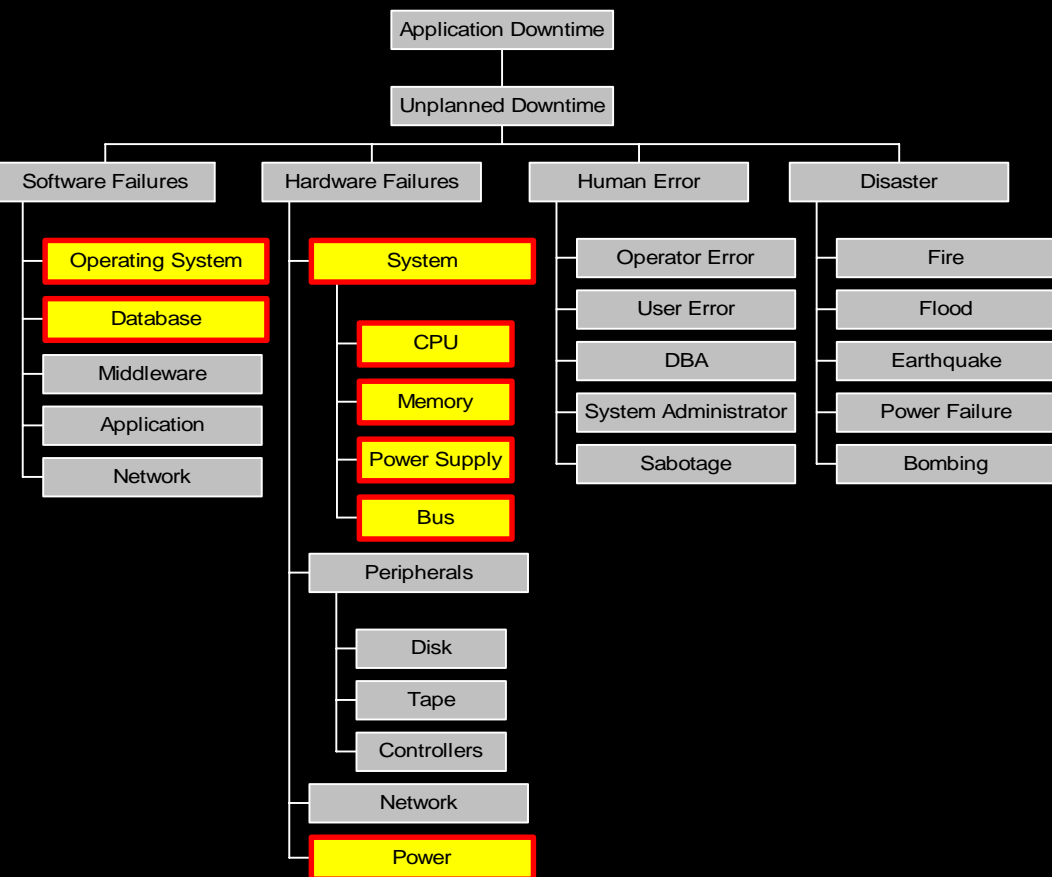
Hot
Failover

Accelerated
Hot Failover



System Fault or Crash

Causes Of Unplanned Application Downtime



- One of the most common causes of application downtime is a system fault or crash
 - Faults have many causes: hardware failure, power failure, operating system crash, Oracle crash, ...
- Oracle automatically recovers business data to a consistent state after a system fault by performing fault recovery
- Oracle8i introduces a new recovery architecture called **fast-start fault recovery** that makes recovery time rapid and predictable

Fast-Start™ Fault Recovery Architecture

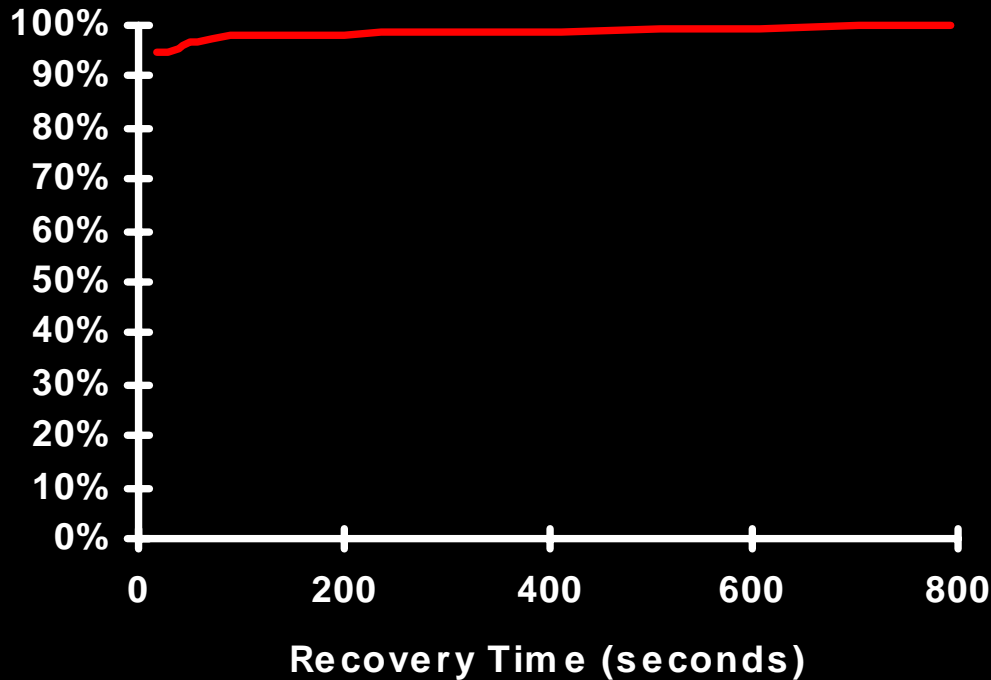
Enhanced
in 8i

- Fast-Start Checkpointing bounds the roll forward phase of recovery
 - Patented Fast-Start checkpointing algorithm continuously advances the recovery point as the database is updated
 - DBA specifies an upper bound on the time that roll forward recovery should take
 - Ensures smooth and fast opening
- Fast-Start Roll Back eliminates the delays associated with the roll back phase of recovery
 - New transactions begin immediately after roll forward completes
 - New transactions will roll back changes to the specific rows that block their progress
 - Other rows and transactions rolled back in parallel in the background
 - Long running transactions no longer effect recovery time

Near Instantaneous Fault Recovery

Enhanced
in 8i

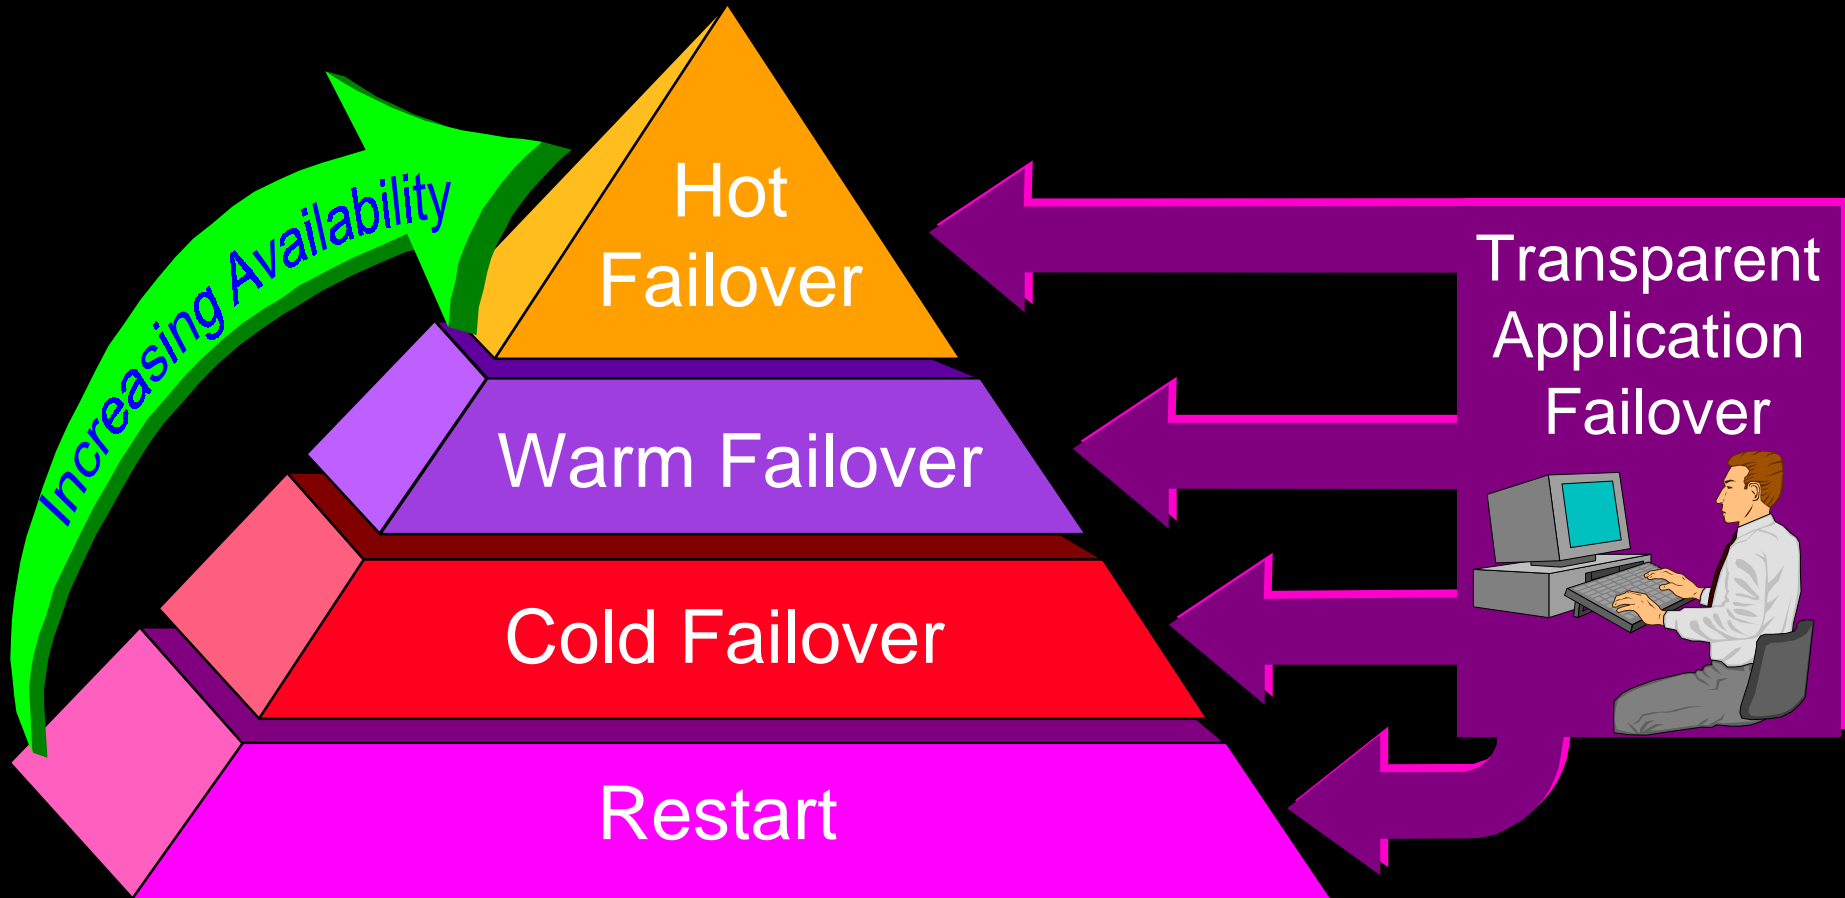
Fast-Start Recovery Time
versus Performance



- Worst case high OLTP workload with concurrent large batch operations
 - *Recovery time reduced from over 13 minutes to 19 seconds!*
- Administrator chooses appropriate recovery time to meet business objectives


Availability Hierarchy

Clustering Eliminates Single Points Of Failure



Bounding System Repair and Restart

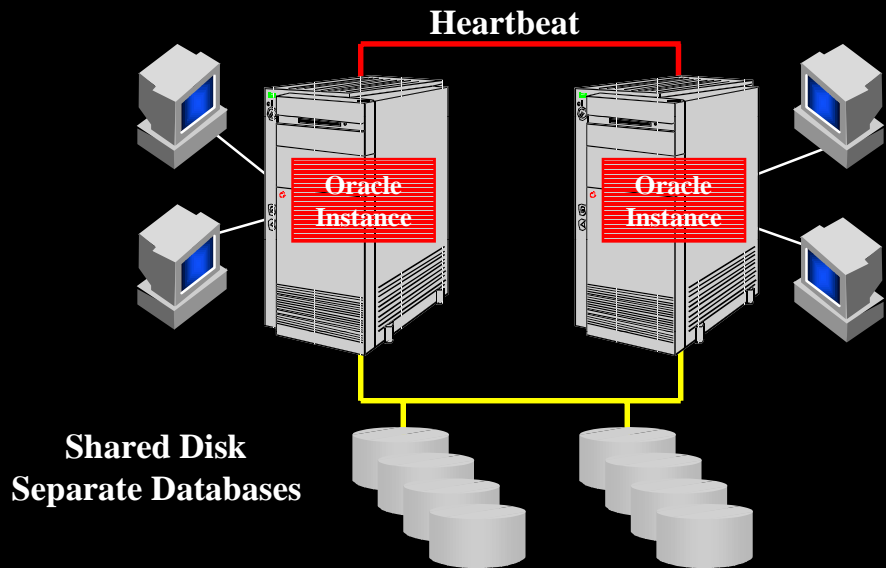
Recovery with Cold Failover



- Restart Oracle on a spare node in the cluster
- Only one system can mount a database at any one time
- Systems vendors provide cold failover solutions
- Oracle Fail Safe on Windows NT
- Remove time to restart or fix system from
no query

Oracle Fail Safe


Recovery with Cold Failover



- Simple automated failover solution for two node NT clusters
 - One node backs up another node
 - When a failure occurs, the database workload is moved to the surviving node
- Uses MS Cluster Server (AKA Wolfpack)
- Simple to administer
 - Oracle provided full function GUI for administration

Unix "Fail Safe-Like" Solutions

Recovery with Cold Failover



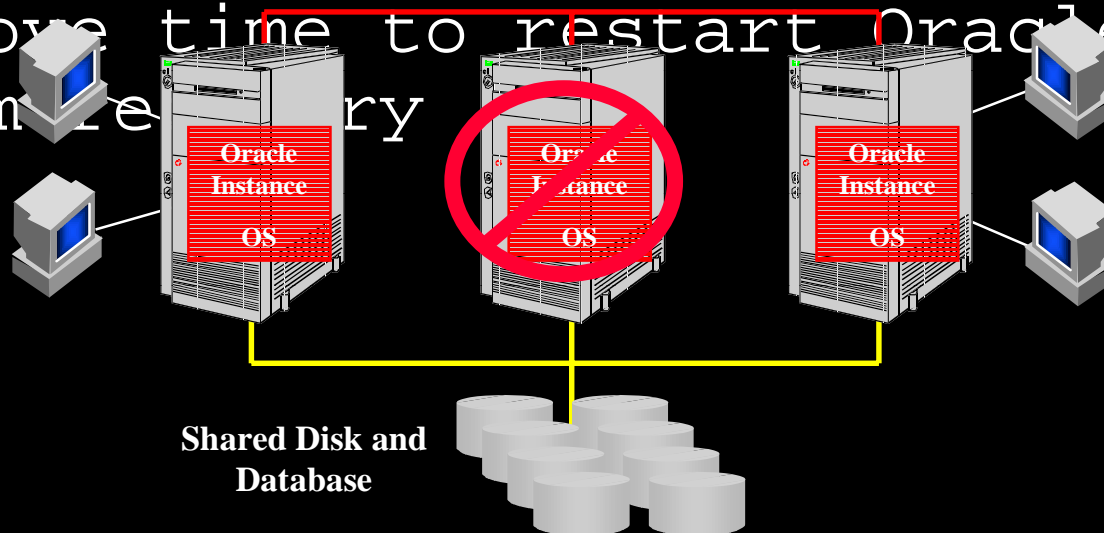
- Most Unix system suppliers offer Fail Safe-like solutions
 - IBM, HP, NCR, SGI, Sequent, Sun, etc.
- Functionality
 - One node backs up another node
 - Storage subsystem is shared between nodes
 - When a failure occurs, the database workload is moved to the surviving node
 - Implementation specific tuning and

Bounding System Repair and Restart

Recovery with Warm Failover

- Failover connections to a peer node in an Oracle Parallel Server cluster
- Pre-warmed buffer cache for better performance
- Remove time to restart Oracle from the recovery


Remove time to restart Oracle from the recovery



Oracle Parallel Server

Recovery with Warm Failover


Enhanced
in 8i



- Oracle Parallel Server provides the ability to build application platforms from multiple systems that are clustered together
- To avoid a single node failure, causing application downtime, use Oracle Parallel Server
 - Eliminates a node as single point of failure
- Allows applications

Bounding System Repair and Restart

Recovery with Hot Failover



- Preconnect to a peer node in an OPS cluster
- Use preconnected backup connection in the event of failure
- Removes time to connect from recovery
- Without preconnect, it can take several minutes to

Availability Demands Of Today's Applications

- Today's internet, OLTP, and Data Warehousing applications mandate that **everything be online all the time**
 - New knowledge workers have increased dependence on database applications
 - We have entered the self service world -- there are no longer intermediaries between in-house systems and customers
 - Every hour of the day, customers somewhere in the world will be accessing your systems
 - If you're applications and data are not available, you are not online, your customers know it, and you are not making²⁵

The Challenges of High Availability Data Servers

- Too many high availability options
- High availability currently complex to achieve. You must...
 - architect a configuration
 - implement the solution
 - perform extensive testing to verify...
 - components coexist
 - failure scenarios

High Availability Methodology

Prevent



- Prevent failures before they happen

Detect



- Detect obvious and subtle failures quickly

Capture




- Capture diagnostic information

Resume



- Resume services to end users quickly, and

Analyze



- Analyze failures with minimal disruption

- Analyze failures to prevent

Oracle Parallel Fail Safe™

- **Based on Oracle8 and Oracle8i**
 - Oracle's highest-end HA database configuration
 - Included with Oracle Parallel Server Option
- **Rapid and predictable recovery from system faults**
 - 10X improvement in server failover
- **Fast deployment**
 - Pre-configured, pre-tuned, pre-tested
- **Developed in conjunction with Hewlett-Packard**
 - Foundation platform for HP's 5Nines Program
 - Year of development, integration and testing
 - Pilot installations at several customer sites

Oracle Email Runs on OPFS

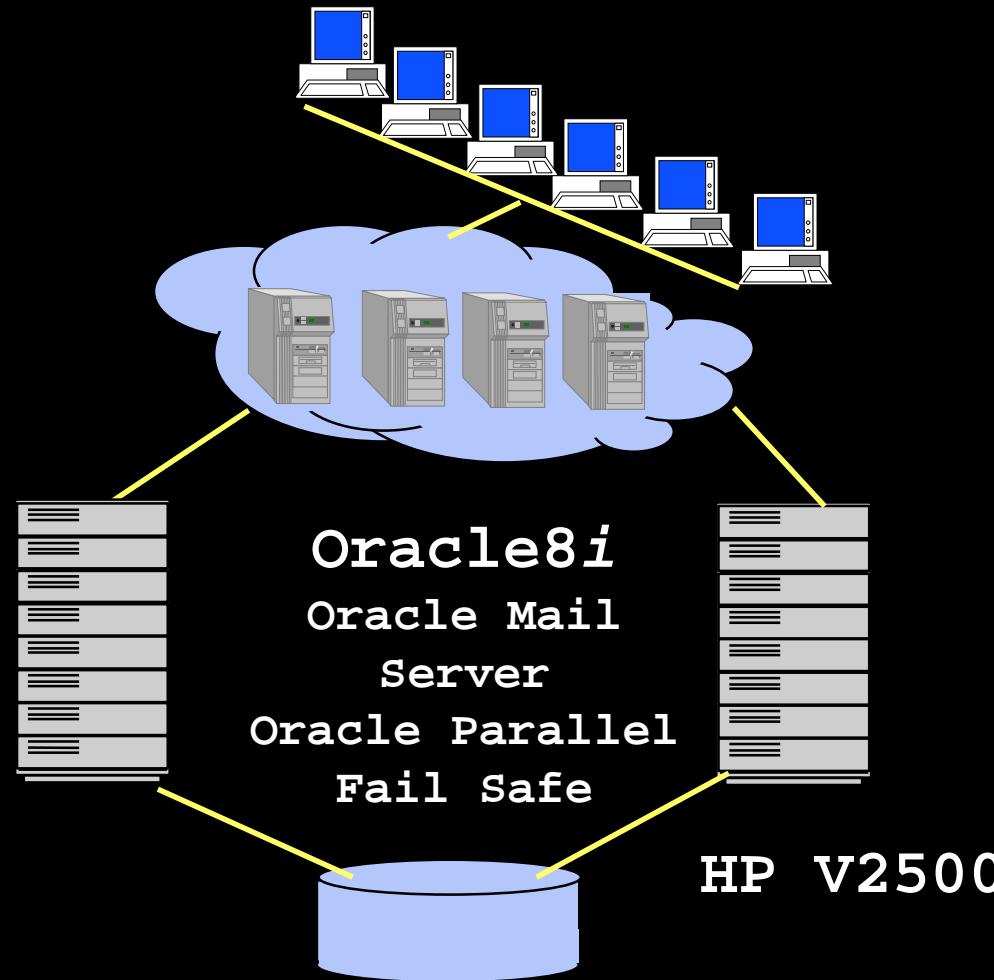
- Reliable and responsive Email services

- 40,000+ users worldwide
- Using Oracle/HP E-Business Continuity Platform

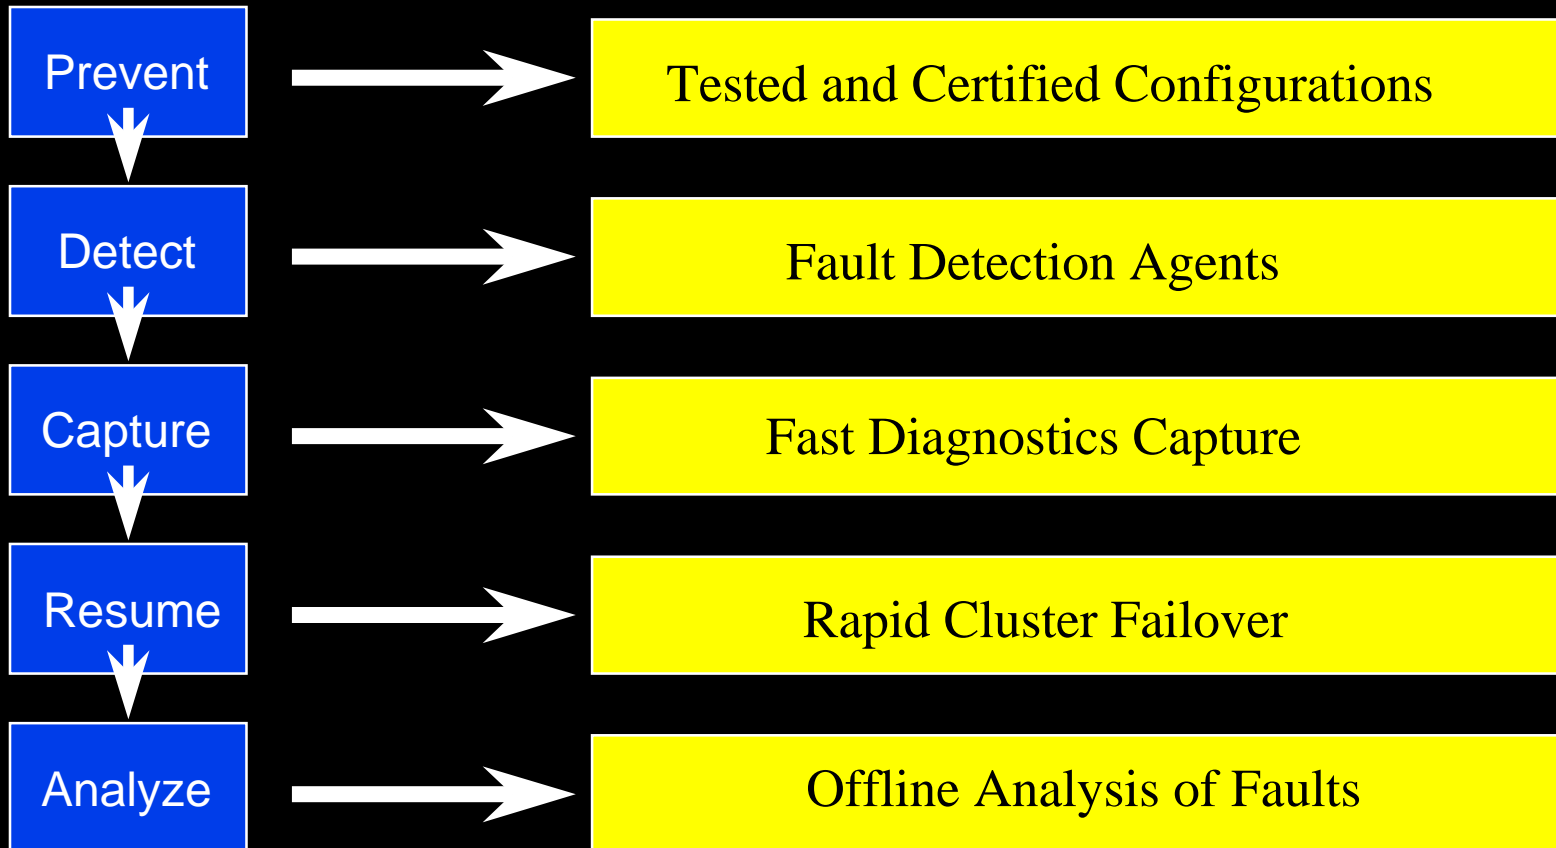
- Reduced costs

- Consolidated 97 servers to 2 HP V-class servers

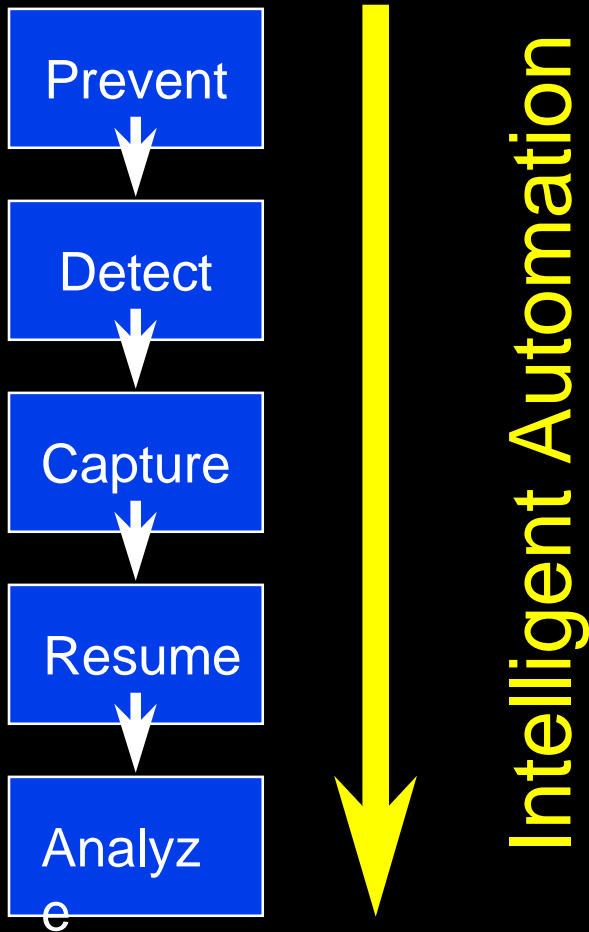
- Savings:
\$11M/year



Parallel Fail Safe Components




Intelligent Automation



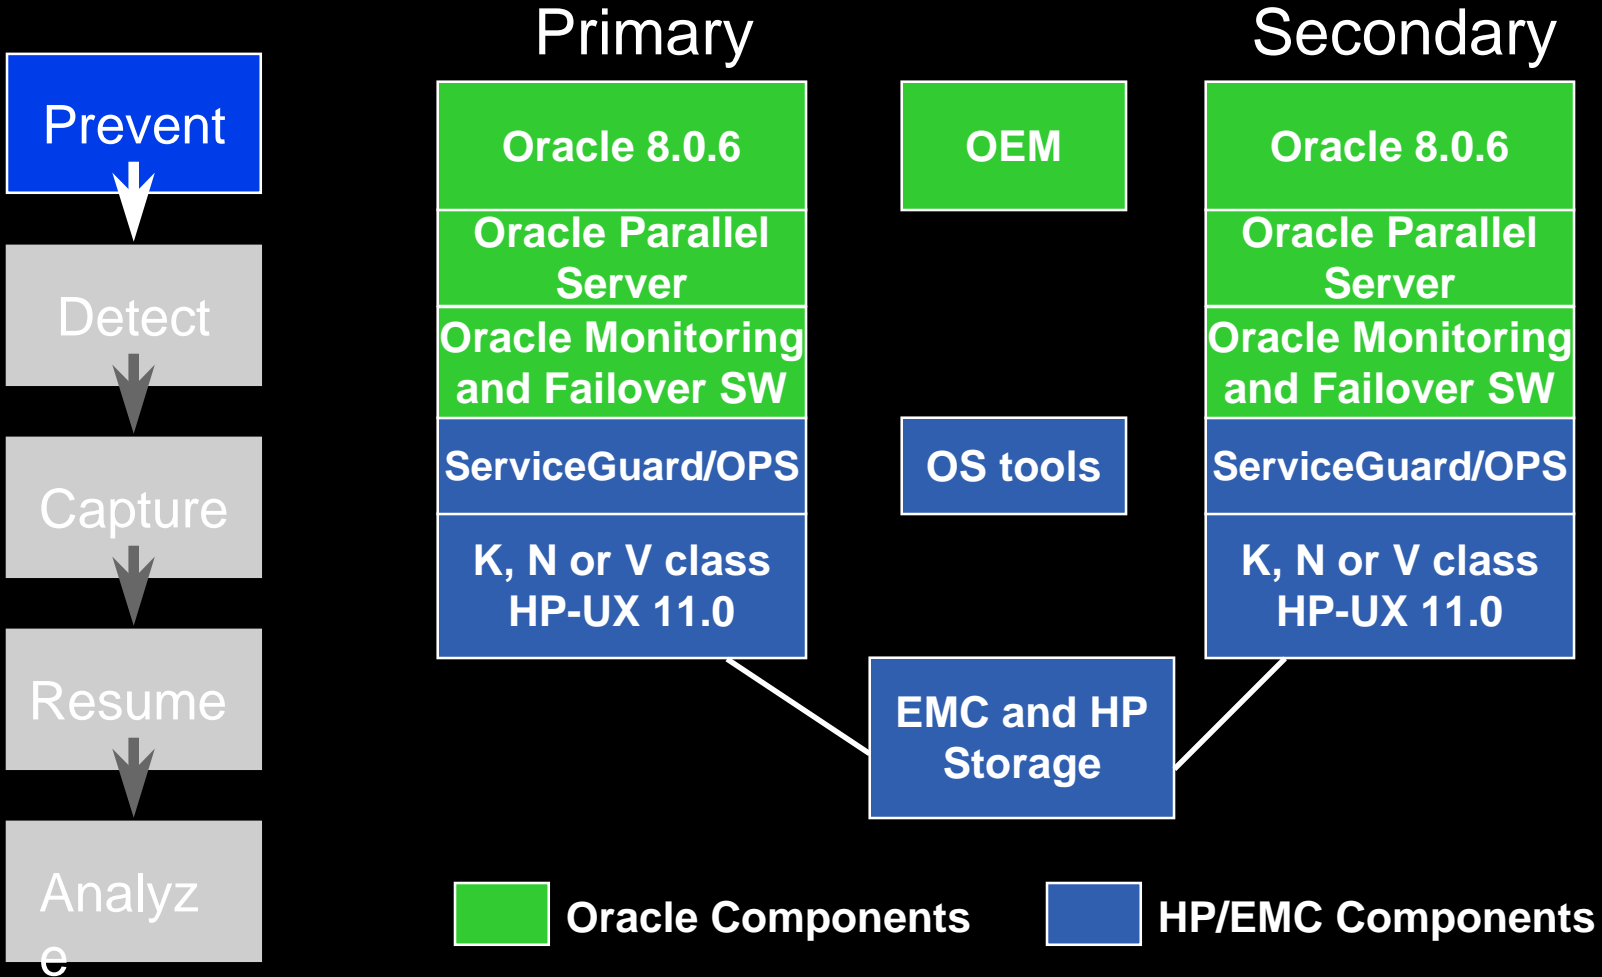
- All components provide intelligent automation for their functions
 - Provides faster, unattended operation
 - Reduces chance of human error impacting the processes

Preventing Downtime




- Oracle Parallel Fail Safe provides certified configurations that have been extensively tested to ensure maximum availability
 - Primary/Secondary cluster configurations provide no single point of failure
 - Oracle Parallel Server provides continuous service after faults
 - Configuration validated through comprehensive fault injection testing
- Graceful planned failovers minimize disruption due to planned maintenance

Parallel Fail Safe Configuration




Oracle Parallel Server provides Key Benefits




- Oracle Parallel Server provide the best possible protection from faults
 - Primary/Secondary mode simplifies configuration and operation
 - Faster and simpler than cluster failover
 - resources don't need to be failed over
 - fast internal detect on instance failure
 - No part of the system needs to be started for fail over
 - Failback is optional (no planned₃₄ downtime for reconnects)

Fault Detection Agents




- Detects hard failures
 - Hardware, OS, or Database crash
- Detects soft failures
 - Hangs
 - Degraded Performance
 - Partial Failures
- SQL level monitoring ensures end-to-end availability
- Automatically measures uptime
- Tight integration with OS cluster services

Capturing Diagnostic Data



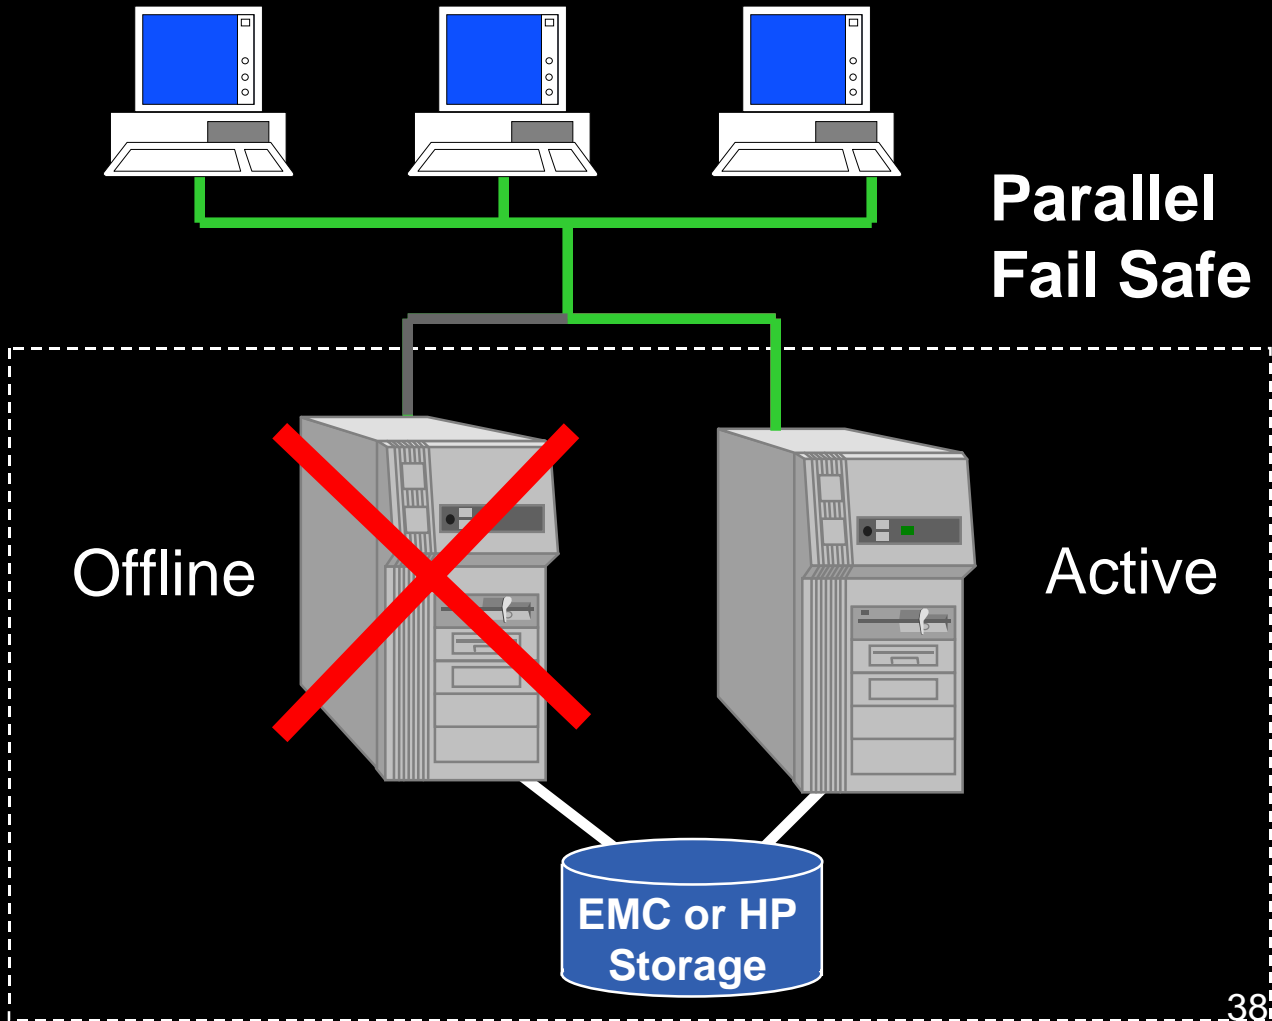

- In the event of a failure, the Oracle Parallel Fail Safe will automatically and quickly collect diagnostic data
- Intelligent agents capture relevant state only so recovery can proceed quickly
- The diagnostic data can often be used to perform root cause analysis on the failure to prevent its recurrence
- Reduces the need to recreate problems
- Capturing data allows for

Resuming Service




- In the event of a failure of the primary instance, Oracle Parallel Fail Safe will quickly switch over to the standby instance
 - No need to fail over volumes
 - No need to restart Oracle
 - No need for Oracle to re-open datafiles
 - Surviving instance will perform recovery
 - Recovery is fast and predictable
- Client Applications can automatically reconnect to the live server

Resuming Service




Fast-Start Checkpoints



- Oracle 8.0.6 provides Fast-start checkpointing
- Optimizes roll forward recovery by continually and incrementally advancing the checkpoint position
 - Roll forward is made smooth, rapid, and predictable
- Automatically adjusts the checkpoint write rate to meet a specified roll forward target




Analyze



- Oracle Parallel Fail Safe provides a holistic solution to quickly recover from problems and prevent recurrence in the future...
 - automatically capture diagnostics data when fault encountered, including logs, blocks, locks, process states
 - provide tools to analyze database hangs to determine the cause of the condition
 - analysis tools can send diagnostic information to the Oracle support center

- Analysis performed offline--after

OPFS Summary



- Oracle Parallel Fail Safe takes the complexity and risk out of building highly available data servers
 - Pre-packaged
 - Pre-configured
 - Pre-tested
- Parallel Fail Safe provides the highest possible availability for off-the-shelf applications
- Holistic solution addresses all aspects from prevention through analysis

ORACLE®

SOFTWARE POWERS THE INTERNET

Sign-up now for
FREE e-business
Continuity kit

<http://www.oracle.com/continuity>