

Best Practices for Implementing Oracle9i RAC on HP-UX 11i

**Gretchen Stewart, HP
Subbu Iyer, VERITAS Software**



Agenda

- A powerful combination – VERITAS and HP
- Oracle9i Real Application Clusters (RAC)
- HP's Two Solutions for RAC
- VERITAS Database Edition/Advanced Cluster™ (DBE/AC)
 - Cluster Volume Management
 - Cluster File System
 - Oracle Disk Manager and Database Acceleration
 - VERITAS Cluster Server™
 - Optimizations for Oracle9i RAC

VERITAS & HP

A Powerful Combination



- Committed to a strong relationship
 - Highest levels of corporate relationship
 - HP's leadership in servers and storage
 - VERITAS' leadership in storage management software

- Technology and services
 - Improve your effectiveness through joint engineering, integration, and support
 - Highest return on IT – cost savings

- Financial stability

HP & VERITAS


A History of Success



- **1989:** HP resells the entire, multi-platform VERITAS product line
- **1993:** HP OEMs VERITAS File System as Online JFS
- **2000:**
 - Lite versions of VERITAS File System (JFS) and VERITAS Volume Manager included with HP-UX
 - Full version of VERITAS Volume Manager available
 - Full version of VERITAS File System included in HP's packaged software bundles
- **2001:** HP OEMs VERITAS Cluster Volume Manager
- **2002:**
 - Integrated Support and Services Agreement for VERITAS products
 - Refreshed OEM agreement covering future releases of HP-UX and VERITAS file system and volume manager products
- **2003:** HP & VERITAS collaborate to provide VERITAS Database Edition Advanced Cluster for Oracle9i Real Application Clusters
 - Beta started in April 2003, Generally available August 2003

VERITAS #1 in Non-Array Based Storage Software

All Software, All Platforms, New License Revenue Only



Total New License Revenue = \$3.2 Billion

Source: Gartner Dataquest "2002 Storage Management Software Market Share" (April 15, 2003)

HP-UX 11i #1 Momentum Continues

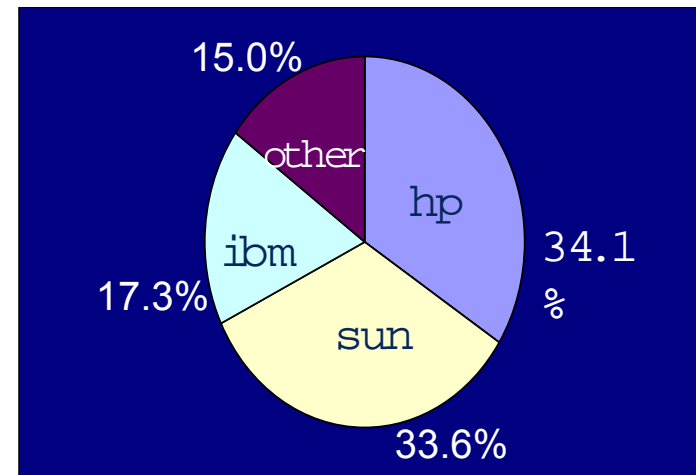
HP-UX has more than 10 years proven expertise in mission-critical environments

The HP-UX ecosystem

- A \$20B business
- Over 2 million units active
- #1 UNIX® market share by revenue last 2 qtrs

Aggressive investment in HP-UX 11i

- Major component of HP's overall R&D budget
- Over 2000 software engineers focused on the HP-UX ecosystem
 - Over 1300 HP-UX patents filed in 2002



#1 UNIX market share - 2H02 per IDC

HP-UX 11i – #1 UNIX even better



- #1 scalability
- #1 reliability, availability and serviceability
- #1 systems management
- #1 internet and web application services
- #1 directory and security services

Source: D.H. Brown 2002 UNIX Function Review Report

"clearly reflecting hp's increased investment in its Unix product line, HP-UX moves to the head of the class for UNIX operating systems functions. HP-UX occupies the top spot in every studied category, with a particularly strong lead in internet and web application services, and an impressive surge forward in the intensely competitive RAS category"

AND a FEW More..... #1 UNIX – 2002 VARBusiness Annual Report Card
#1 UNIX – 2002 Butler Group OS Review

HP and Oracle long standing partnership



technology partnership
for Oracle9i[®] RAC
development




ORACLE[®]

- HP is the market share leader of Oracle9i RAC implementation
- first company to provide rapid deployment (9 days)
- selected components of HP TruCluster technology are licensed for use in Oracle9i RAC
- HP has the closest development relationship for Oracle9i RAC on UNIX[®], Linux and Windows
- scalability, stress, performance and leading benchmark testing is done on HP infrastructures
- HP played a major role in cache fusion design discussions
- HP and Oracle have cooperative support agreement

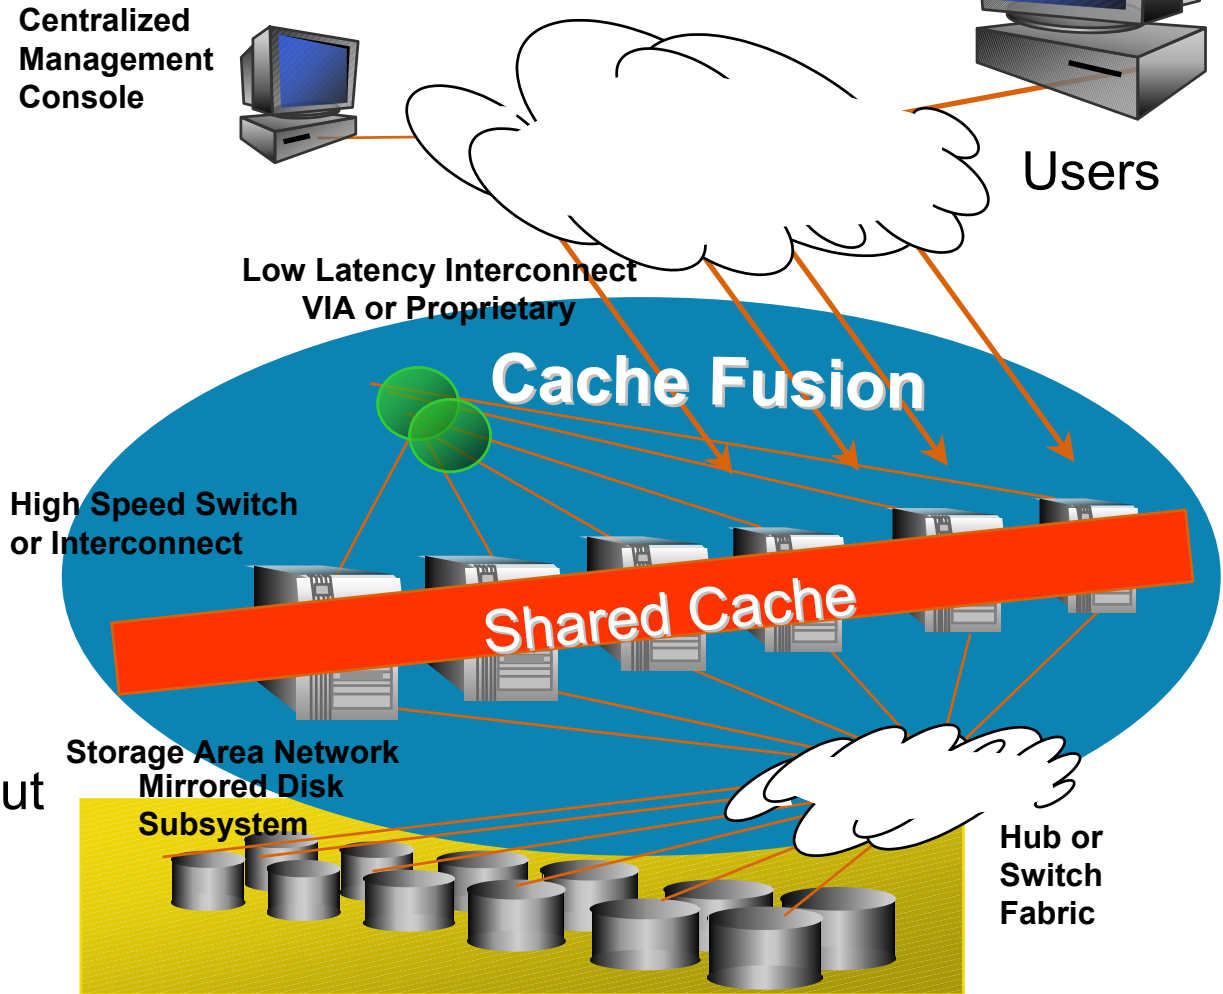
Oracle9i Real Application Clusters

- Leading parallel database in the market
- Single database, shared by multiple Oracle instances on a cluster of servers
 - Optimized for scalability and availability



Oracle9i RAC Architecture

- Full Cache Fusion
 - Cache-to-cache data shipping
 - Shared cache eliminates slow I/O
 - Enhanced IPC
- Allows flexible and transparent deployment
 - Scale-up/Scale-out

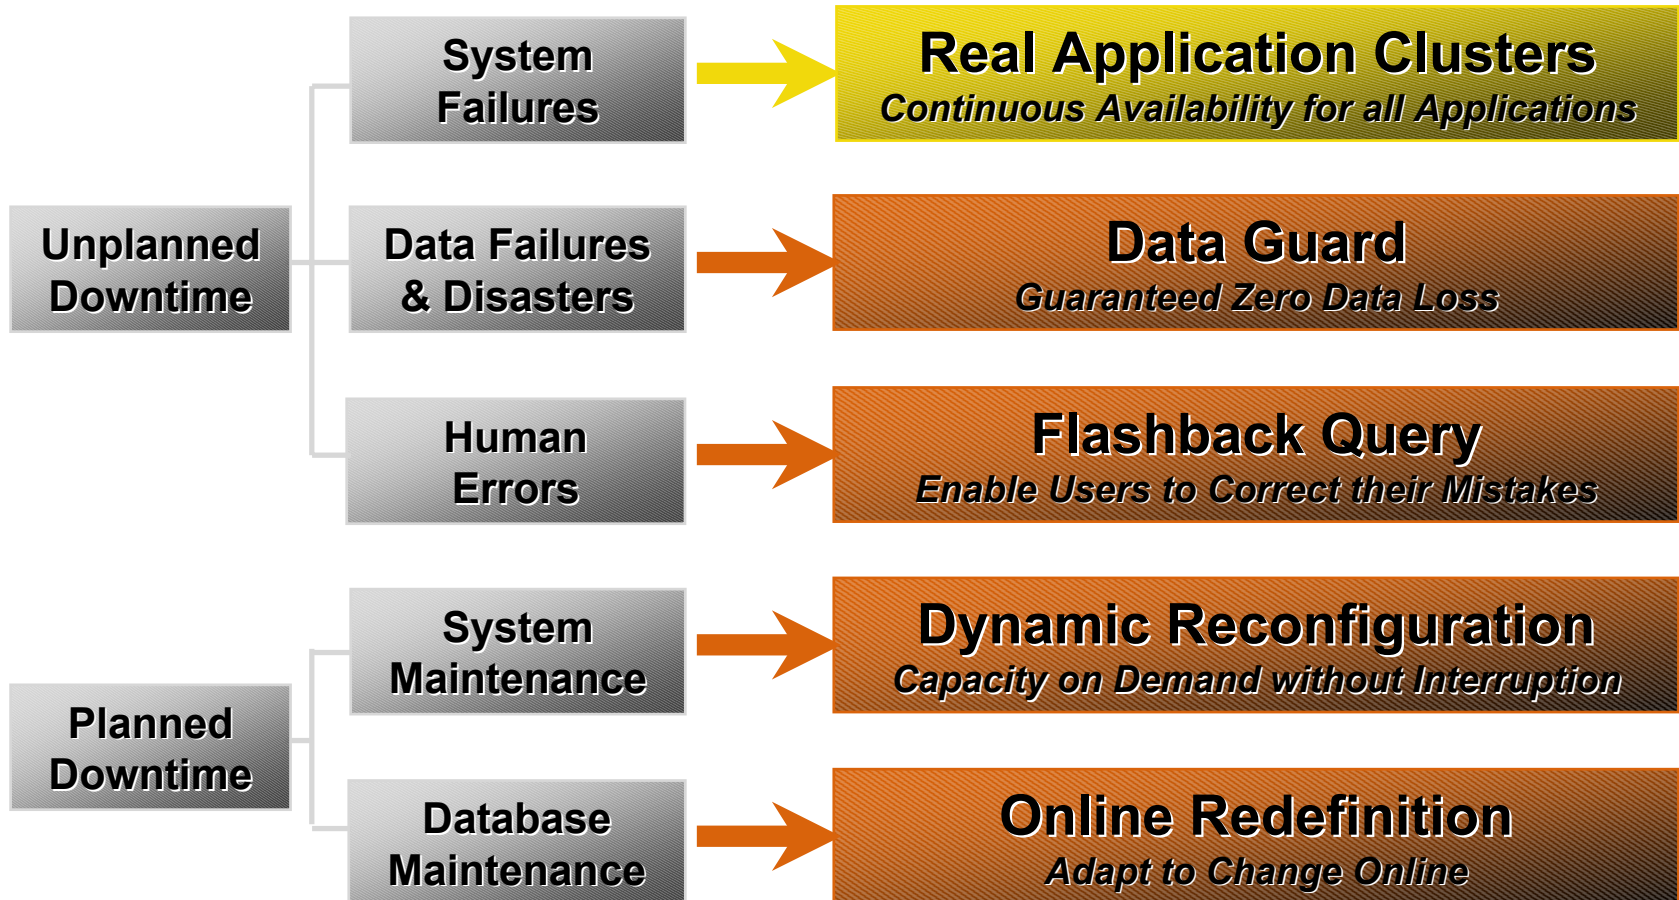


Drive and Exploit Industry Advances in Clustering

Oracle9i RAC: Availability and Scalability Features

- Real Application Clusters overcomes the limitations of failover solutions
 - No need to failover volume groups/file systems
- Reduced time to recovery
 - Concurrent resource reconfiguration and instance (cache) recovery
 - Global Cache Service coordination local buffer cache and remote block transfers
 - Deferred recovery operations
- Fast/reliable detection of node/network failure
- Reducing Scheduled Downtime
 - The Most complete 'On-line' maintenance
 - Full Cache Fusion
 - Enhanced IPC

Oracle9i Database – ensures your business information is always available



Two solutions for Oracle9i RAC – we are committed to flexibility in customer choice

- HP Serviceguard Solution – SGeRAC
 - Simple solution for existing Serviceguard customers
 - Low initial acquisition cost
 - “Raw” device access offers the best performance possible

- VERITAS Database Edition / Advanced Cluster
 - Cluster File System with near “raw” performance and simplified management
 - Scalability with high availability
 - Ease of data manipulation on disk

- **Both solutions provide the infrastructure for RAC**

Oracle9i RAC Challenges

How do I protect against database corruption?

How do I build an infrastructure that can handle data growth of 50% to 200% per year?

How do I reduce ongoing management costs?

How do I backup my RAC database?



Technology Components Needed for RAC

- **Oracle RAC Requires**
 - **Shared concurrent access to storage**
 - Cluster Volume Manager
 - Cluster File System
 - Oracle Disk Manager Support


 - **Cluster Membership**
 - Who is in the cluster?
 - Who is joining the cluster?
 - Who has left the cluster?

 - **Communications**
 - Intra instance messaging
 - Cluster State
 - Cache Fusion

Database Edition / Advanced Cluster (DBE/AC)



Oracle9i RAC



VERITAS Database
Edition / Advanced
Cluster for Oracle9i



HP-UX Server

- Database Edition / Advanced Cluster is a complete solution for RAC
- Optimized for RAC
- Tested with RAC
- Certified by Oracle

Database Edition/Advanced Cluster (DBE/AC)

Oracle9i RAC

RAC Extensions

Cluster Server

Database Accelerator (ODM)

Cluster File System

Cluster Volume Manager

HP-UX Server

- Based on core VERITAS technologies
 - Robust features & stability
 - Reuse VERITAS knowledge
- Specific enhancements for the RAC environment


Cluster Volume Manager (CVM)



Cluster Volume Management

- Concurrent access to volumes from multiple nodes
- Extension of VERITAS Volume Manager™ technology
- Flexible Striping and Mirroring
- Online Grow/Shrink for maximum uptime and utilization
- Dynamic Multi-Pathing between servers and storage

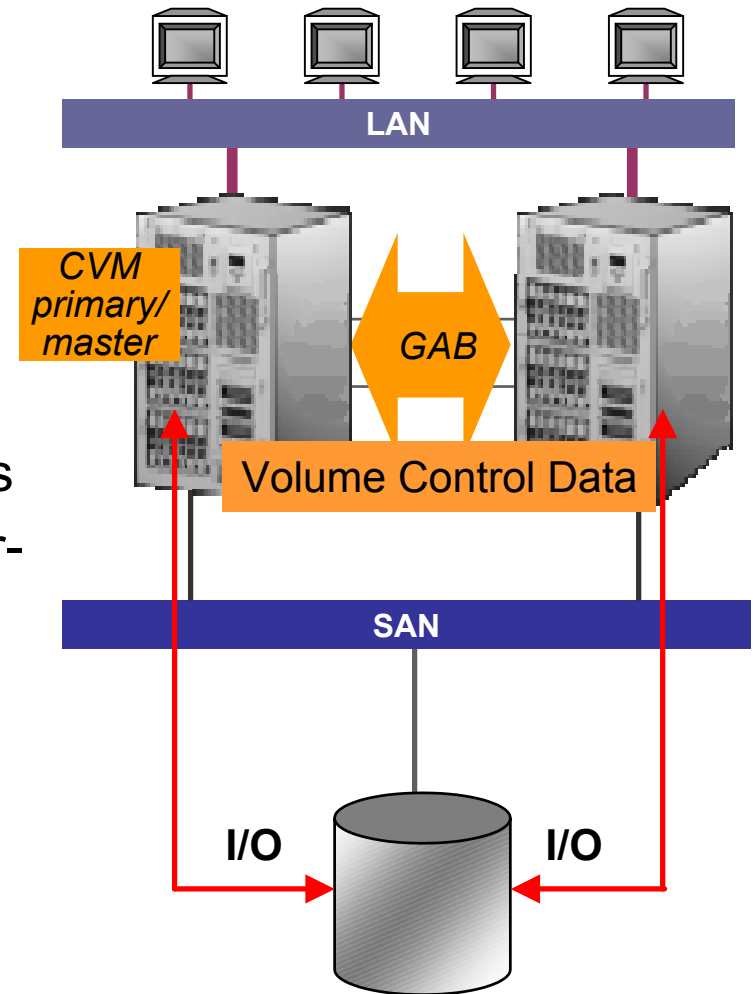
Oracle9i RAC



HP-UX Server

CVM Architecture

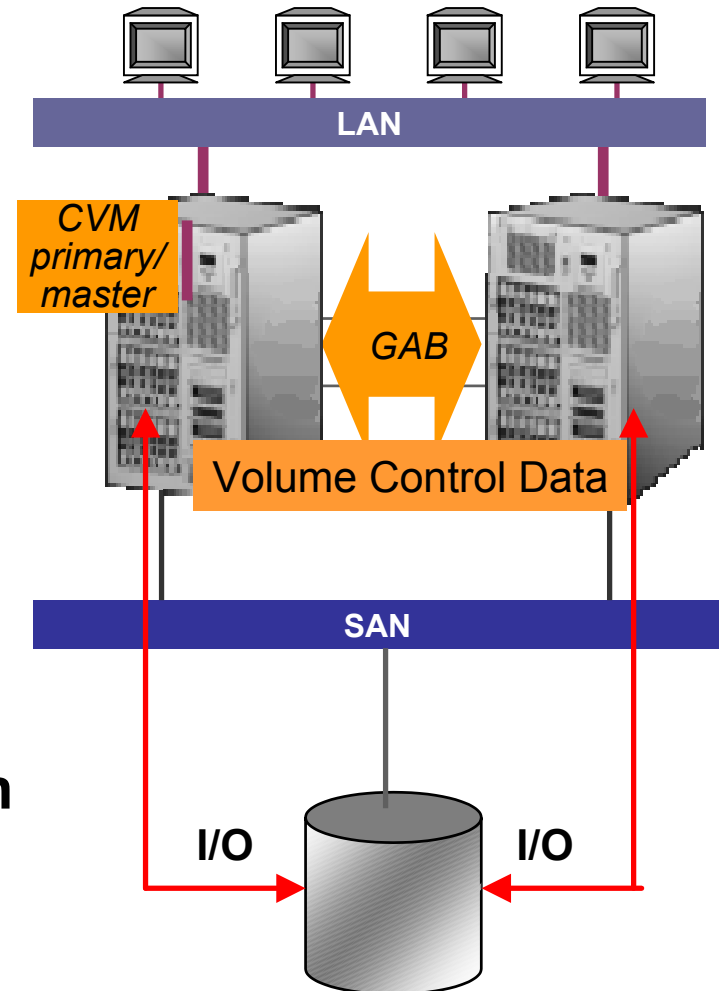
- **Primary/Secondary design for logical volume management**
 - Primary handles all configuration changes
 - Primary role fails over if node crashes or leaves
 - Any node has the ability to switch roles
 - Configuration changes handled cluster-wide



CVM Features

- **Full VxVM capabilities**
 - Clustered Dynamic Multi-pathing
 - Add disks online
 - Online volume growth/shrink
 - Online volume relayout
 - Fast Mirror Resync
 - Disk group split/join
 - Works across multiple arrays

- **Complete industry standard VERITAS volume management in a cluster environment**



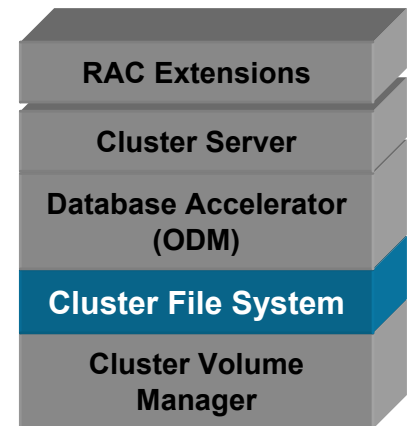
Cluster File System (CFS)



VERITAS Cluster File System

- Concurrent access to file systems from multiple nodes
- All nodes can directly read/write
- Same Look & Feel as JFS/OJFS
 - JFS/OJFS are VxFS
 - CFS is licensable feature of VxFS
 - No retraining needed!
- Simplified deployment of RAC
 - Single \$ORACLE_HOME
 - No NFS cross mounts for recovery
- Lower administration costs


Oracle9i RAC



HP-UX Server

CFS Architecture

- **Primary/Secondary Design for metadata I/O and coherency locking**
 - Primaries are assigned on a per-filesystem basis
- **Global Lock Manager (GLM) locks are used for cluster wide cache coherency**
- **Each node performs file data I/O direct to the storage**
- **The Primary role fails over if the node crashes or leaves**
 - Automatic
 - Fail-over order can be specified



CFS Features

■ Full Cluster File system

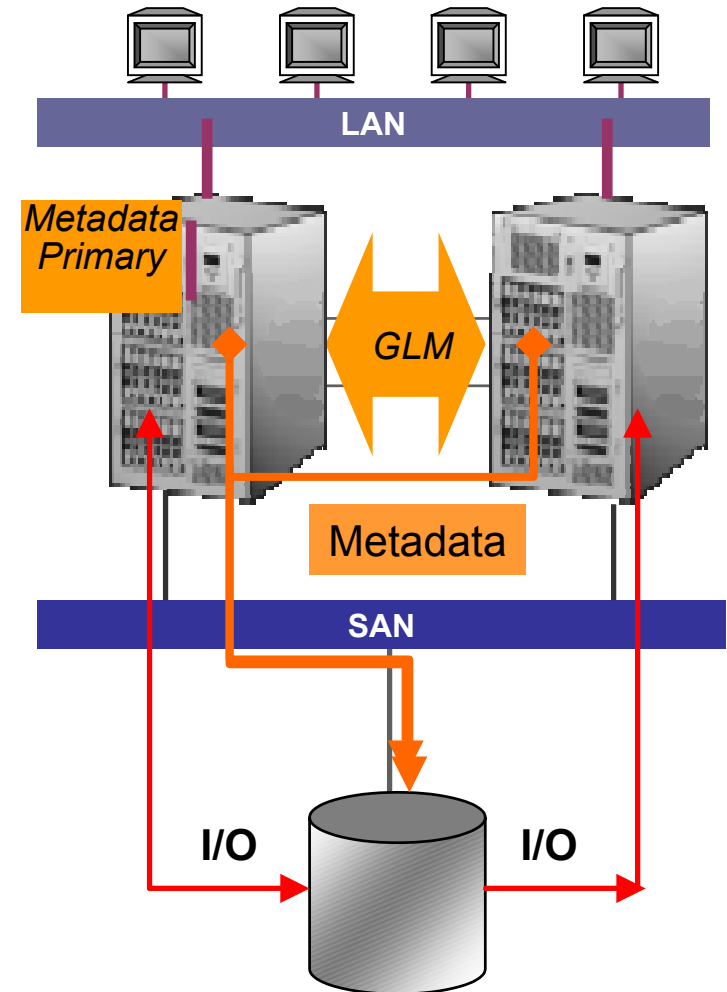
- Not a proxy file system!
- All nodes can directly write
- Journalled for rapid recovery
- Extent based for maximum performance
- OJFS/VxFS with cluster license enabled

■ Full function CFS for all file system uses

- Oracle datafiles and logs
- ORACLE_HOME
- Application binaries and data

■ File system locking bypassed for data files by ODM interface


- CFS locking only used to add, delete or extend a data file
- RAC maintains all locking with GCS/GES



Cluster File System with ODM

- All Oracle processes use single I/O library (ODM)
- I/O API defined by Oracle
- VERITAS provides its own library for Oracle process to call
 - No performance penalty of using a file system instead of raw device
 - Supports Oracle Managed Files (OMF)
 - Support for auto_extend tablespaces

Oracle9i RAC




HP-UX Server

VERITAS Performance Testing

- VERITAS CFS performance on HP with Oracle9i RAC
 - Raw device performance
 - CFS manageability with no performance penalty

Oracle9i RAC on VERITAS CFS vs. RAW




User Load Scaling Test with DBE/AC,
Oracle RAC 9.2, HP-UX 11i, 2 node cluster

Mission Critical Availability

- Leverages industry leading heterogeneous clustering technology
- Manage cluster as a single entity
- Inter-node messaging for state and configuration information
 - Global Atomic Broadcast (GAB)
 - Low Latency Transport (LLT)

Oracle9i RAC



HP-UX Server

VERITAS Optimizations for RAC

- Libraries for Cluster Management, Cache Fusion, and Oracle Disk Manager
- Enhanced Inter-node link performance
- Increased reliability by fencing off crashed node
 - Protects against “Split Brain”

Oracle9i RAC

RAC Extensions

Cluster Server

Database Accelerator
(ODM)


Cluster File System

Cluster Volume
Manager

HP-UX Server

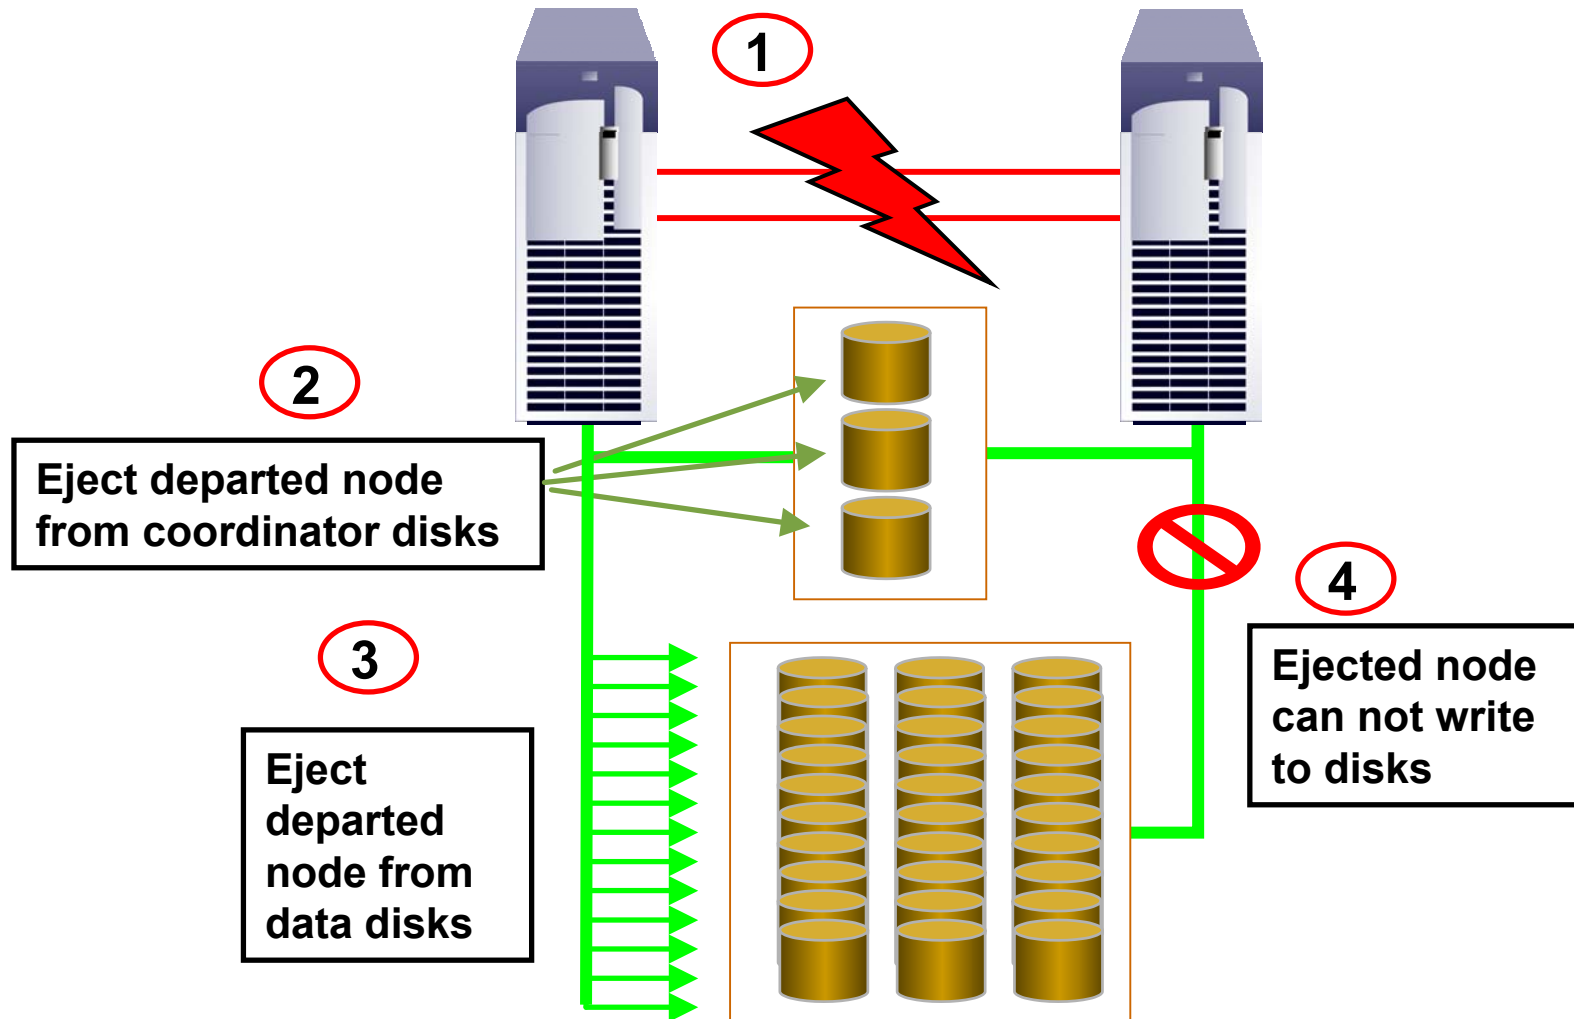
Understanding Split Brain

- Split Brain can cause database corruption
- One node must survive, the other must be shutdown
- VERITAS I/O Fencing handles all scenarios



**Network failure or system failure?
System failure or system hang?**

I/O Fencing: How Does it Work



Upgrade Considerations

Oracle8*i* OPS to Oracle9*i* RAC with DBE/AC on HP-UX 11i



■ From Oracle8*i* OPS to Oracle9*i* RAC

- Upgrade to HP-UX 11i
 - AR0902 or later
- Upgrade to Oracle9*i* RAC, Release 2 (9.2.0.2) or later
- SCSI-III fencing capable storage mandatory
- Install and use DBE/AC
 - Configure and use Veritas CVM and CFS
 - HP OJFS is Veritas File System (VxFS)
 - Veritas CFS is licensable feature of VxFS
 - SLVM conversion to Veritas CVM mandatory
 - In-place conversion from LVM to VxVM
 - Configure and use Veritas Cluster Server (VCS)

What does VERITAS add to RAC?

- **Lower Total Cost of Ownership**
 - Manageability of a Cluster File System

- **High Performance Infrastructure**
 - Cluster File System at near raw device speed
 - High performance inter-node connection
 - Faster more resilient I/O

- **Increased Availability and Resiliency**
 - Flexible storage protection to keep database online
 - Protection against cluster component failures

VERITAS Database Edition Family for HP-UX 11i



- Database Edition for Oracle
 - Building robust **single instance** implementations
 - Tested and certified with MC/Serviceguard

- Database Edition for Oracle / HA
 - Building robust **single instance failover** implementation
 - Includes VERITAS Cluster Server technology

- Database Edition for Oracle / Advanced Cluster
 - Building a scalable and highly available **parallel** environment, supporting Oracle9i RAC

VERITAS & HP

A Powerful Combination



- Committed to a strong relationship
 - Highest level corporate relationship
 - HP's leadership in servers and storage
 - VERITAS' leadership in storage software

- Technology and services
 - Improve your effectiveness through joint engineering, integration, and support
 - Highest return on IT – cost savings

- Financial stability



HP WORLD 2003

Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

