COMPAQ
Inspiration Technology

ORACLE COMPAQ
The Oracle Compaq Alliance

Compaq TruCluster for Oracle9i RAC

Clusters for the real world.

gilgu.yoon@compaq.com
Alpha Solutions Consulting Team
Technical Consulting & Pre-sales
May, 2002

Agenda

- The Oracle/Compaq Partnership
- Huron Project
- Certified Configurations
- Compaq Database Utility Solution With Oracle 9i RAC
- TruClutser Technology
- Compaq's Unique Advantage for Oracle9i RAC (Why Compaq?)
- 9i RAC references on Compaq TruCluster
- Summary

COMPAQ
hp & compaq. the new power of invention

Compaq/Oracle Partnership

Long History Together

- First commercial SQL database (PDP, 1979)
- First 32-bit database (VAX)
- First Oracle Parallel Server development (1991) on VAX
- First Oracle TPC benchmark
- First 64-bit version of Oracle (v7.3, 1995)
- First 64-bit OPS (1996)
- First 100,000 tpmC on AlphaServer System (1996)
- First VLM on NT & OpenVMS
- First 64-bit version of Oracle8i (1998)

COMPAQ
hp & compaq. the new power of invention

Huron Project & Certified Configuration

Cluster Inventor

Driving Radical Simplification

Michael Capellas

To Create Clustered Database Solution
Highly Scalable
Highly Available
Highly Manageable
SIMPLE to Grow, Manage, and Maintain

Making the Complex Simple

Larry Ellison

Clustered Database Inventor

IDC Server Clustering Forecast, 2000–2005

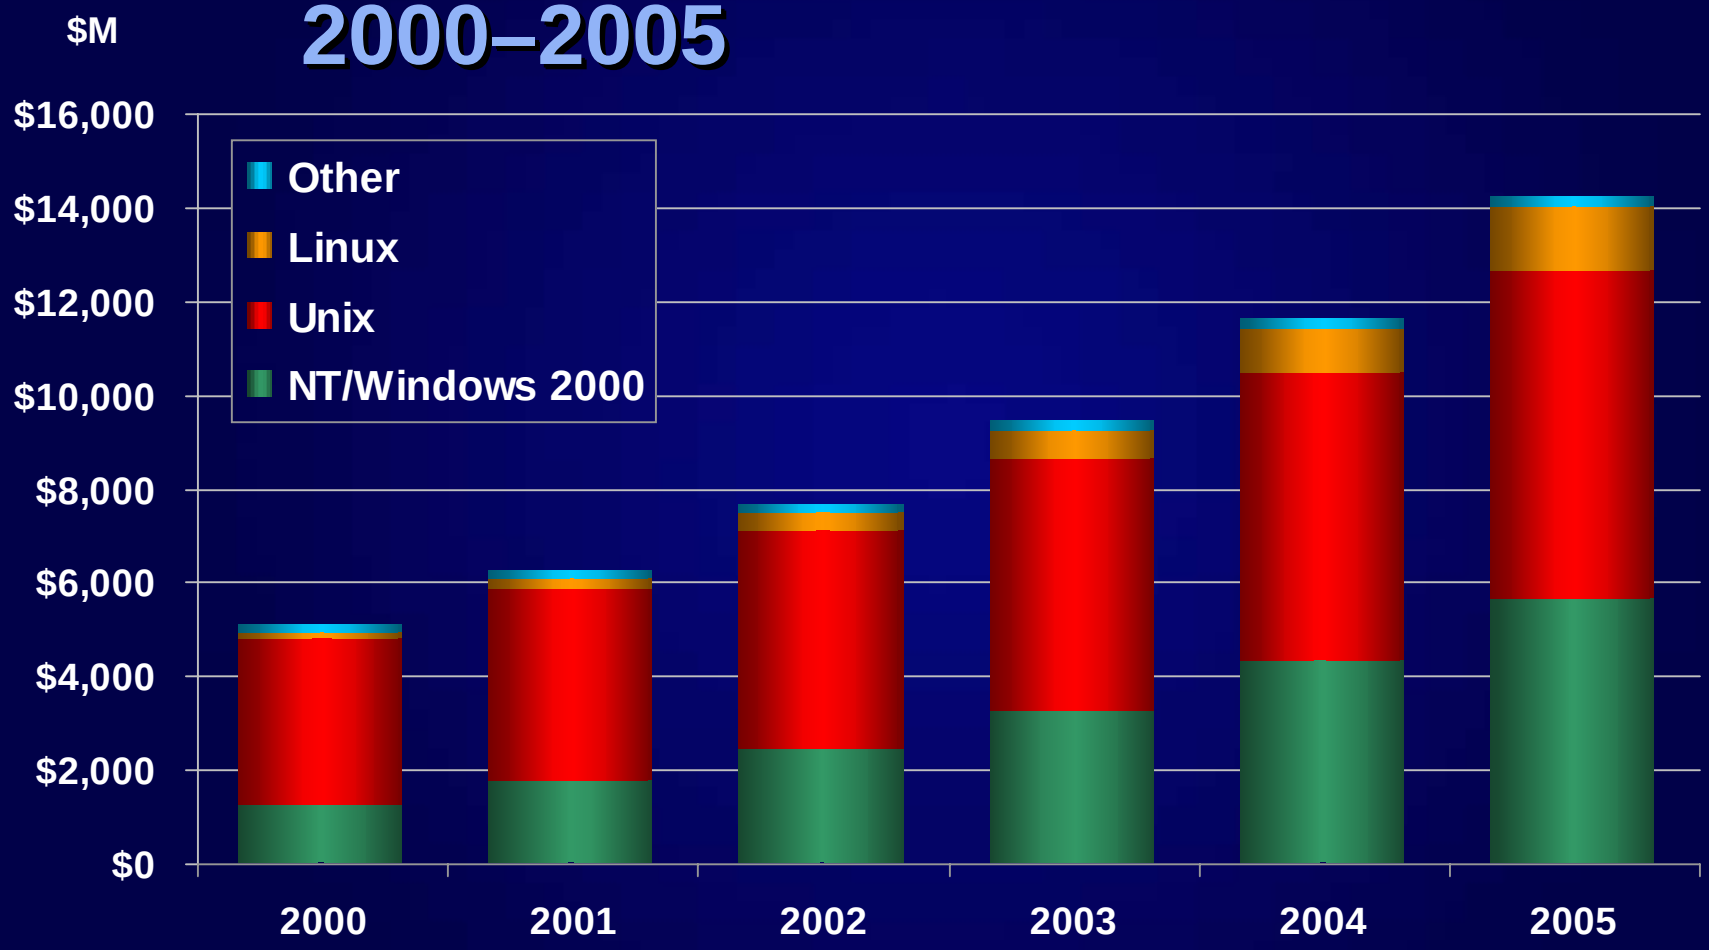

Server Clusters Take Off: What Does it Mean For IT?
Jean S. Bozman, Matthew Eastwood
IDC Telebriefing #25380 - August 2001

Compaq and Oracle Partnership

- ***Joint engineering* agreement** to integrate key *Tru64* UNIX clustering technology into Oracle9i Real Application Clusters (RAC)
  - Oracle is creating portable clusterware for all Oracle9i Platforms
  - Compaq is the lead development partner for RAC
  - *Tru64* UNIX is a reference platform for RAC
- Significant multi-year ***market development partnership*** for clustered Oracle solutions on Compaq *Tru64* UNIX

Oracle and Compaq Partnership Goal: Broaden adoption of Clustered Databases

- Why?
  - Clustered databases are the best choice for a

    scalable, manageable, always-on

    database environment that is

    easy to buy, install, maintain and grow
- How?
  - Implement with no single point of failure
  - Make it simple

Oracle's Any Vendor, Any Hardware Architecture …

- Start with SMP or NUMA
- Any Vendor: Sun, HP, Compaq, IBM, …
- Add cluster nodes on demand

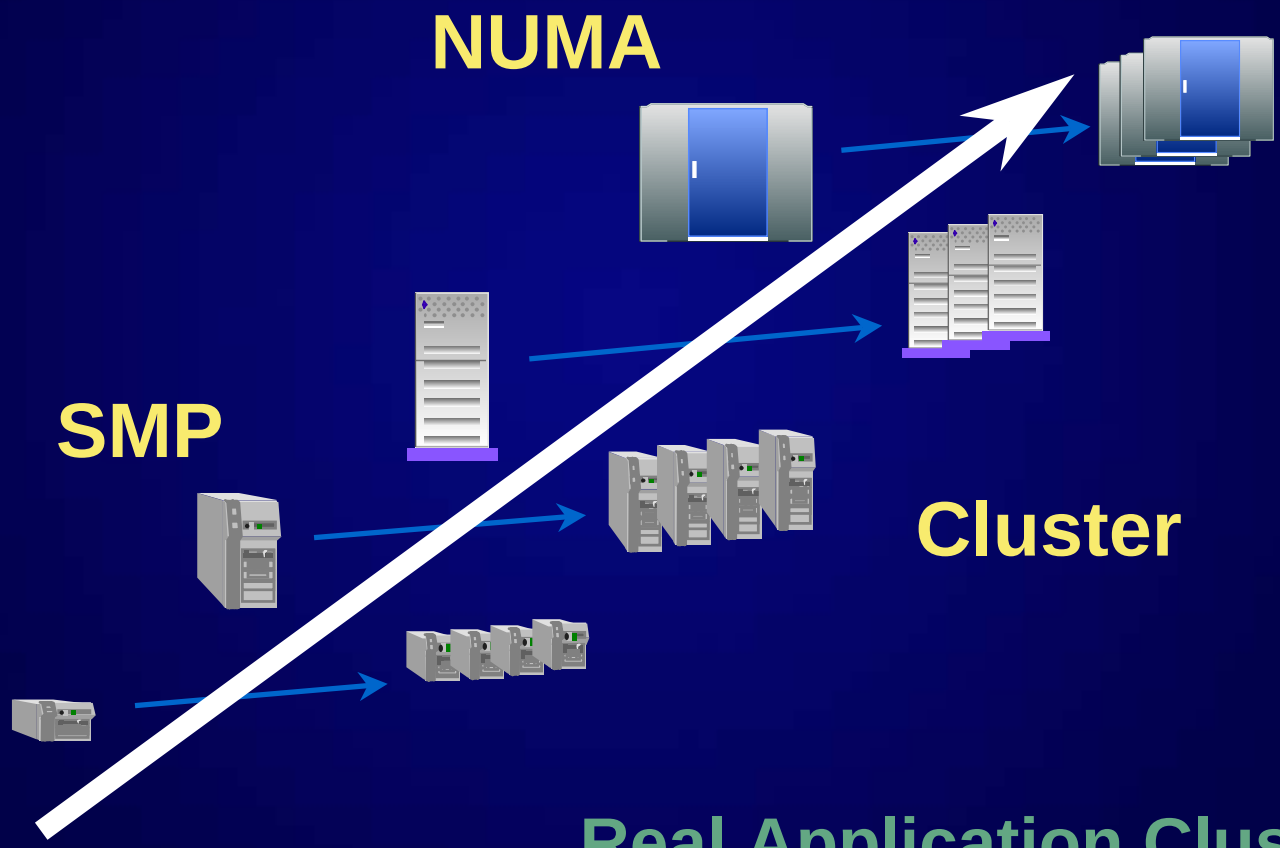

Real Application Clusters

…Is available now with Compaq *Tru64* UNIX and AlphaServer Systems

- Start with SMP or NUMA
- Save management costs
- Minimize database disruption
- Add cluster nodes on demand, in under 20 minutes

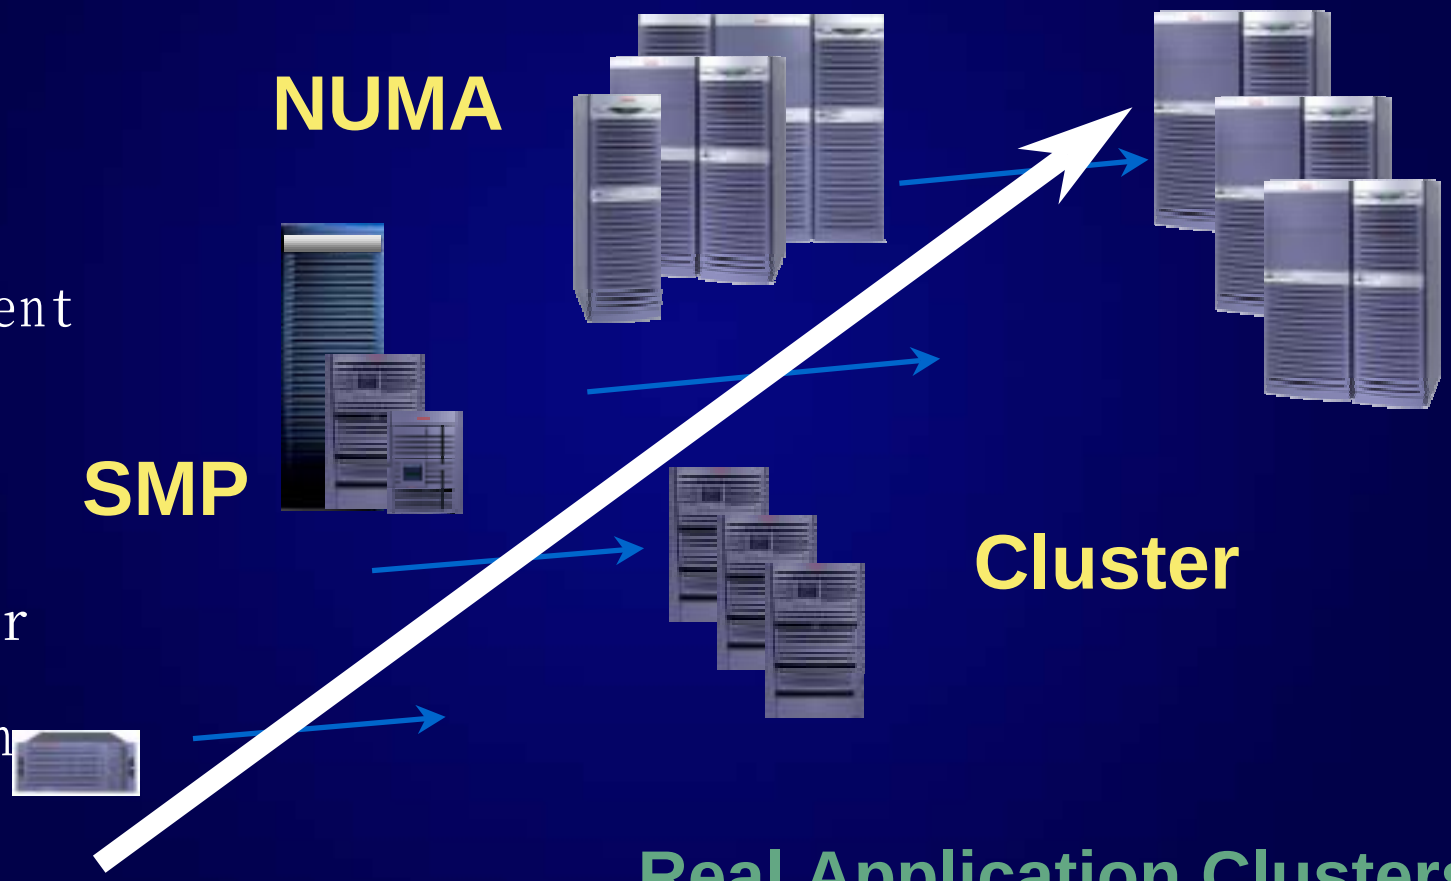

Oracle 9i Real Application Clusters

- Grow your DATA
- Grow your USERS
- Grow PROCESSING POWER through SMP and/or clusters
- Reduce management costs and complexity

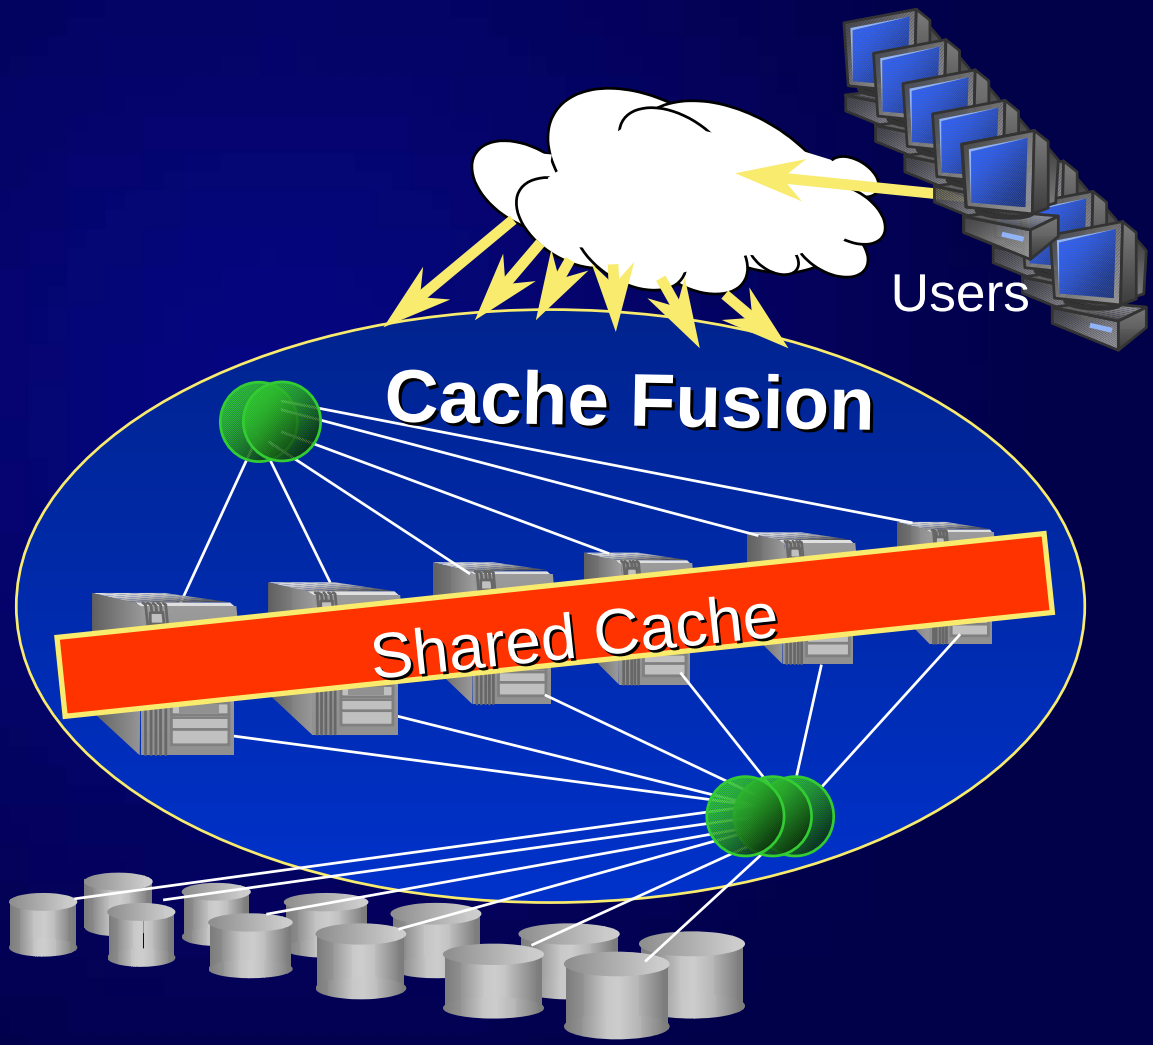

Oracle9i Portable Clusterware Components

Oracle9*i* Real Application Clusters

COMPAQ *TruCluster* Technology

- Cache Fusion, IPC Interconnect
- Oracle OSDs
- Cluster Application Availability Facility
- Cluster Management (SysMan)
- Cluster MIB and SNMP Agent
- Event Manager (EVM)
- Distributed Lock Management
- Connection Manager

TruCluster Server _Tru64_ UNIX

- Single System Image
- Cluster (Shared) File System
- Single Security Environment
- Global Device Names

COMPAQ
hp & compaq. the new power of invention

Certified Configuration

Database Utility

- Dual rack configuration
- Significant expansion: 8 nodes, 4.6TB
- Simple to add a node, storage, or support more users or workload
- New version of Oracle9i Real Application Clusters with scalability & manageability enhancements

Certified Configuration

- Compaq & Oracle preinstalled, pretuned, and stress-tested
- TP application performance optimized and tuned to site specifications
- Racks of clustered database systems
- Configured for expansion
- Single part number
- 9 day delivery

Standard Products

- Oracle*9i* Real Application Clusters
- Tru64 UNIX & TruCluster Server
- AlphaServer Systems
- Storage
- Services

Compaq/Oracle Certified Configurations

- Pre-installed, pre-tuned and stress-tested solutions on *AlphaServer* and *ProLiant* hardware
- Most efficient deployment vehicle with lowest Total Cost of Ownership (TCO)
- Comprehensive infrastructure framework services for complete stack management
- Serviced jointly by Oracle and Compaq, with a single point of contact

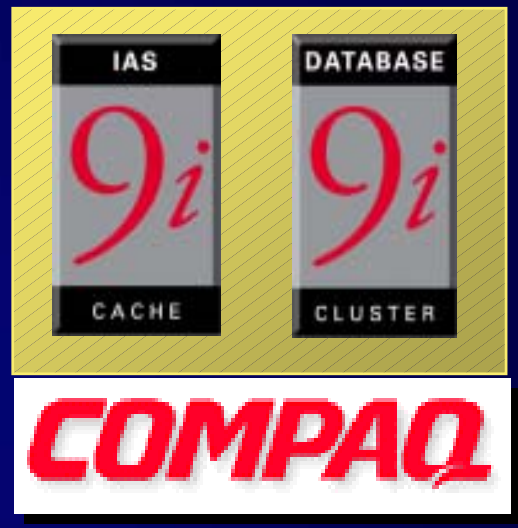

Certified Configurations !

Oracle9*i* Real Application Clusters*

- 2-Node AlphaServer Cluster

Includes:

- Dual ES40 Cluster
- Tru64 UNIX V5.1
- Oracle9*i* Real Application Clusters Preinstalled

Oracle9*i* Real Application Clusters*

- Select *ProLiant* servers

Includes:

- Flexible memory and processor configurations
- SuSE Linux, Windows 2000
- Oracle9*i* Real Application Clusters Preinstalled

Oracle e-Business Suite

- ProLiant DL580
- Oracle e-Business Suite Preinstalled

Oracle 9*i*AS

- ProLiant DL360 or DL580
- Oracle9*i*AS Preinstalled

**Also available with Oracle8i OPS*

Compaq Certified Configurations for Oracle 9i

No-Single Point of Failure eBusiness Scale-Out Solution

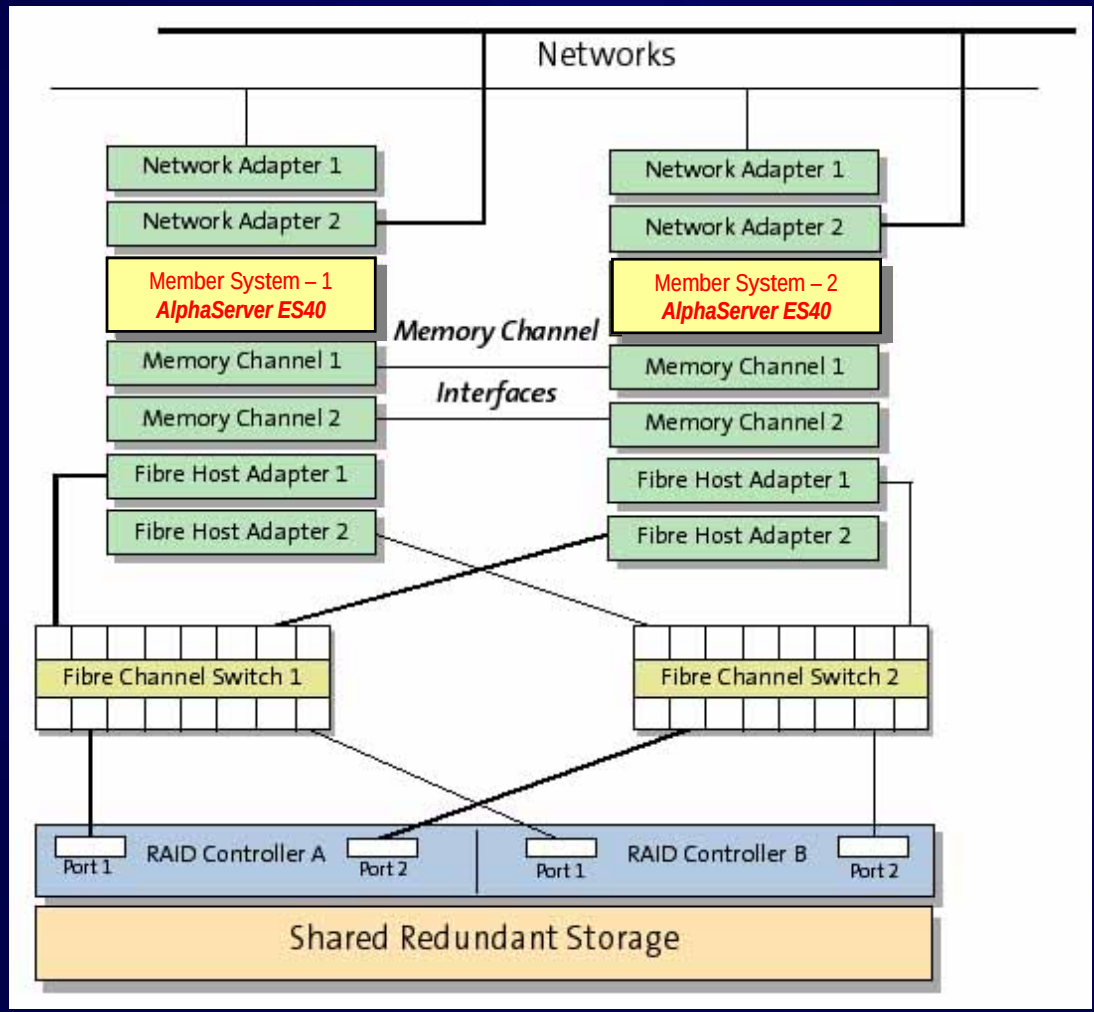

Compaq Clusters Roadmap
for Oracle9*i* Real Application Clusters

SuSE Linux

ProLiant DL580

Windows 2000

Select ProLiant Servers with RA4x00 or RA8000 Storage

TruCluster Server

ES40 Entry ***ES40*** ***ES45***

Higher Levels of Availability, Scalability & Management

- Full Range of Certified Configurations for Customer Choice: Linux, Windows 2000, *Tru64* UNIX
- All systems are pre-configured, pre-tested and certified
- Offering a range of availability, scalability and management

COMPAQ
hp & compaq. the new power of invention

Compaq Database Utility Solution With Oracle 9i RAC

Compaq TruClusters - Continues to set the pace!

- **Compaq Database Utility Solution - Available Today!**
  - Dynamic expansion of Oracle9i Real Application Clusters, including disk space and number of systems
  - Highly available
  - Manageability enhancements

December 3, 2001

Database Utility solution
with Oracle9i Real Application Clusters

Single Part Number

- **2 ES45s**
- **Tru64 UNIX TruClusters**
- **Oracle9i RAC**
- **Oracle Certified**
- **Ready-To-Go**
- **9 day delivery**

Key Features of the Utility

- *Build highly available systems from smaller building blocks and scale up to new levels*
- *Dynamic expansion: database performance & size can increase without taking the database offline*
- *No single point of failure*
- *Reduced management costs that extend from Database to the System*
- *Increased IT productivity*

Base Configuration

- **Pre-defined operating environment**
- **Pre-defined configuration allows for optimization**
- **Smart installation & upgrade specific to Utility configuration**

Ease of Expansion

- **Purchase what you need when you need it without disruption - CPUs, Memory, Storage**
- **Single command adds cluster nodes and storage**

TruCluster Technology

Compaq and Clusters
20th Anniversary

Cluster Inventor

- First DIGITAL Cluster shipped (OpenVMS)– May 1981
- First 64-bit UNIX Cluster shipped –1994
- First *true* single system image UNIX cluster – August 1999

AlphaServers are used in the most challenging genome sequencing tasks at public and private centers worldwide

Celera Genomics, Sanger Centre,
Whitehead Institute, GeneProt, etc…

"The race to map the human genome is over, and the winner is … Compaq.
At least from a computing standpoint."

David Einstein, Forbes News

..and Require Innovative Solutions

Many Systems

Separate Environments	Inefficient use Of resources	High Availabiilty With complexity

Single System Image

One Single environment

The Essence of Computing

Meeting the Challenge of High Availability

Causes of Downtime

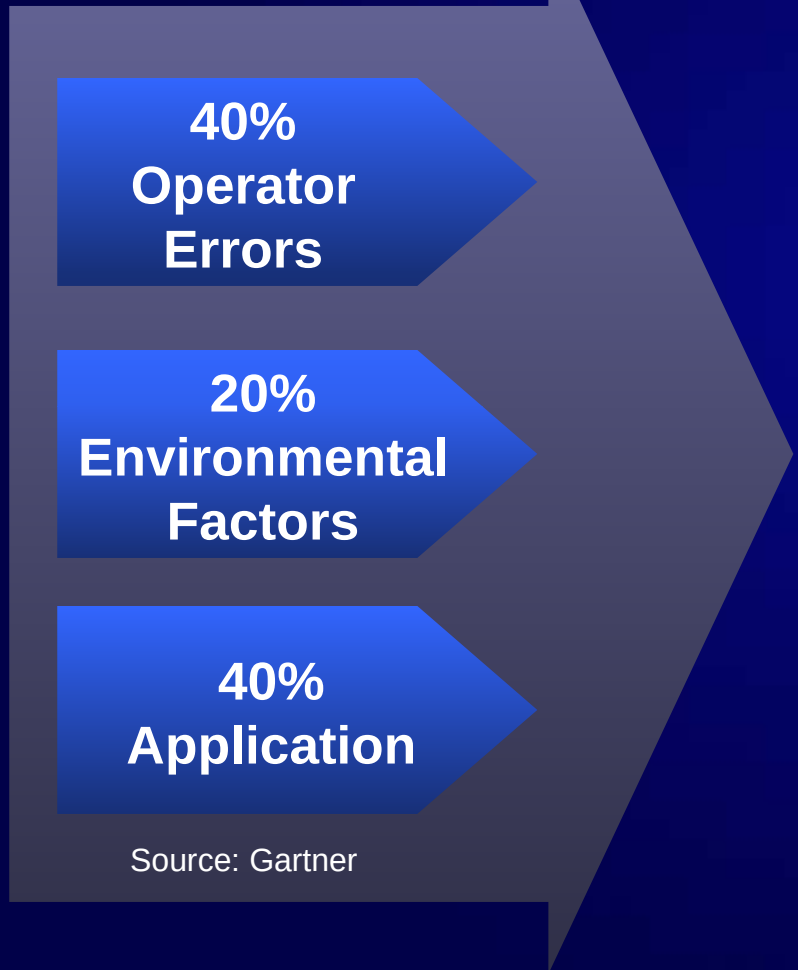

Source: Gartner

Tru64 UNIX on AlphaServer Features

- Automation
- Multipathing load balancing
- Alternate pathing
- On Line Add and Remove (OLAR)
- Online maintenance and upgrades
- Multi-user patch
- Single System Image
- Cluster File System with Shared Root
- HA Services and Uptime Guarantees 99.999%

Cluster Alias

- Cluster appears as single system to network
- Single host name to clients
- Quicker failure recovery
- Dynamic load balancing
- Network services
- Efficient forwarding over cluster interconnect

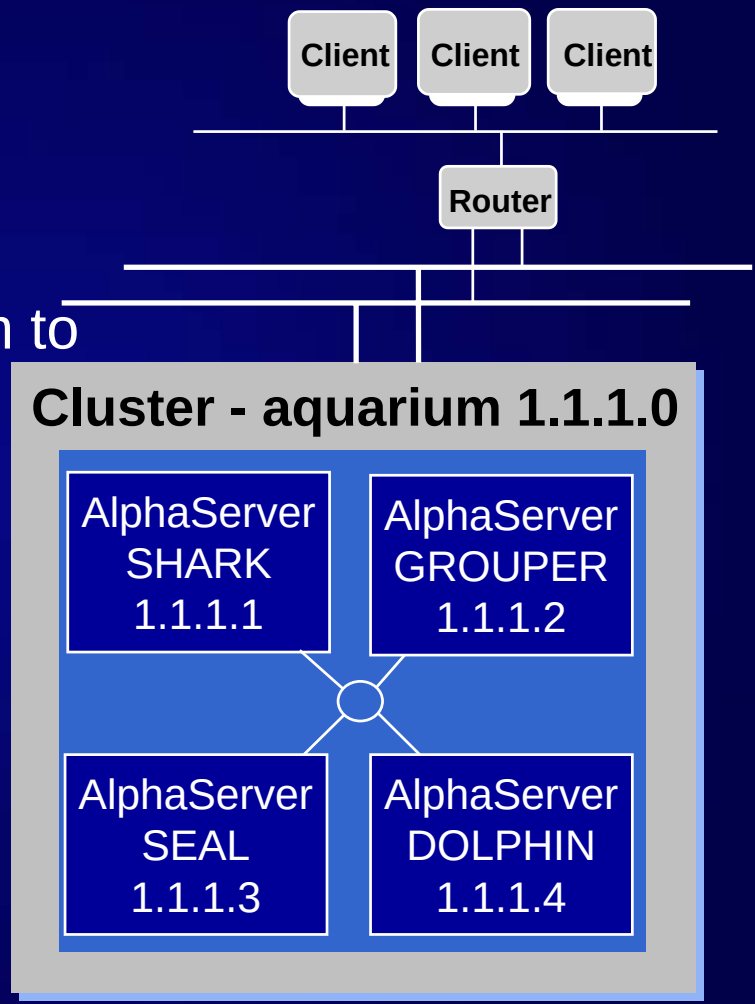

Cluster Application Availability (CAA)

- Provides application failover or restart within the cluster
- Application and resource dependencies
- Application profile determines failover policy and dependencies
- Tru64 UNIX system services configured with CAA
  - NFS rpc.lockd, BIND, DHCP

SysMan Station: View of: earth.unx.dec.com as: fvc

File Action Tools View Options Windows Help

Views

- AdvFS_Filesys
- CAA_Applicati
- CAA_Applicati
- Hardware
- Mounted_Files
- Physical_Files

Hardware View

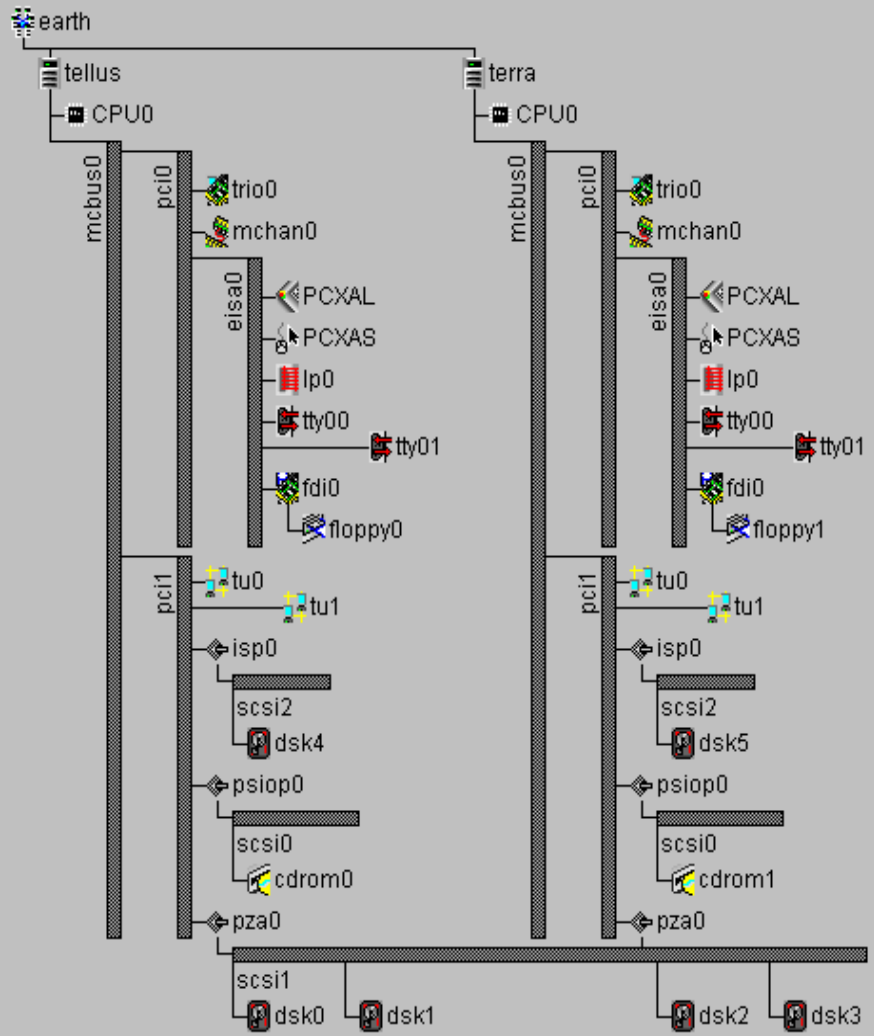

Easy to manage Storage - Online!

Data Protection

- Log-based File System
- On-line Backup and cloning
- Data Mirroring
- Multi-Platform Backup and Archive
- Data Recovery Utility

Data Availability

- On-line File Management
- On-line Capacity Balancing
- Multi-path / Multi-bus
- SAP/Oracle Database Backup
- Hot Sparing
- Software or Hardware RAID

Performance

- Direct I/O
- Large I/O size
- Multi-Disk Support, online file migration between disks
- On-line Defragment
- Fast On-line Backup
- Smooth Synch
- Directory Hashing

Usability

- Graphical User Interface
- Automated File and Disk Management
- Automated Media Management

TruCluster Server – The Straightest Path to Oracle9*i* Real Application Clusters

Scalability

Scale up and out at the speed of now
Growth to match your business velocity

Availability

Your business data is always available
Completely eliminate your system downtime

Manageability

Slash your management time and effort
Continue to reduce them dramatically

Clusters for the real world.

TruCluster Server Clusters are unique

Separate Systems

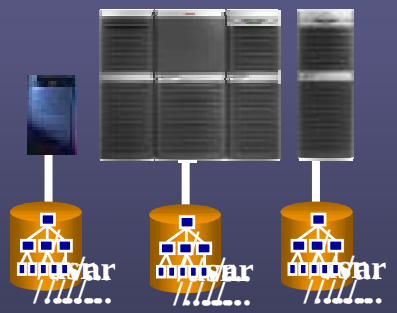

- Manage as many
- Complex

Single System View

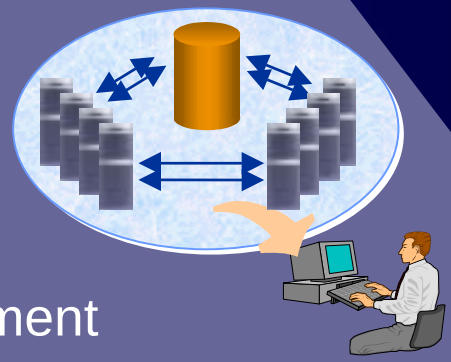

Simplified Management

Multiple nodes for increased scalability

Redundant nodes for availability

Paradigm shift: Single System Personality

Trust the inventors of clusters to make UNIX clusters simple.

TruCluster V5 File System

Cluster file system, shared root, single system image

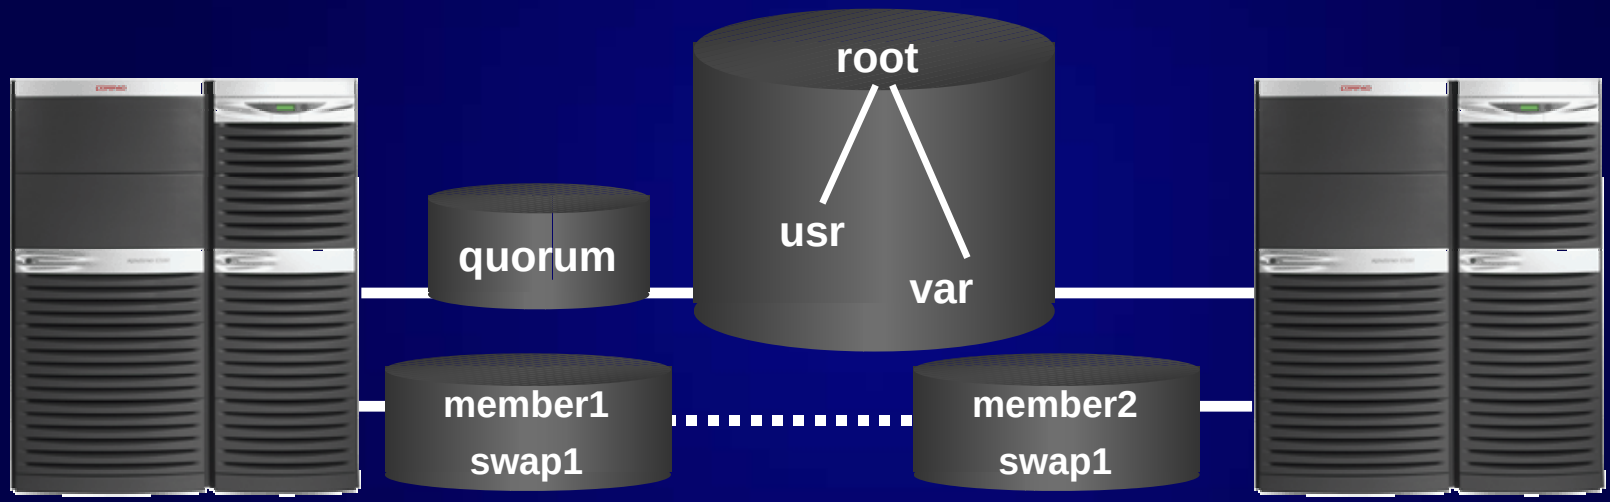

The cluster file system is shared by all cluster nodes.
All nodes see the same data.
The cluster manages as a single system, even as you add nodes.
Adding a node takes about 15 minutes.

Scalability

Cost-effective Scaling with *TruCluster* Server

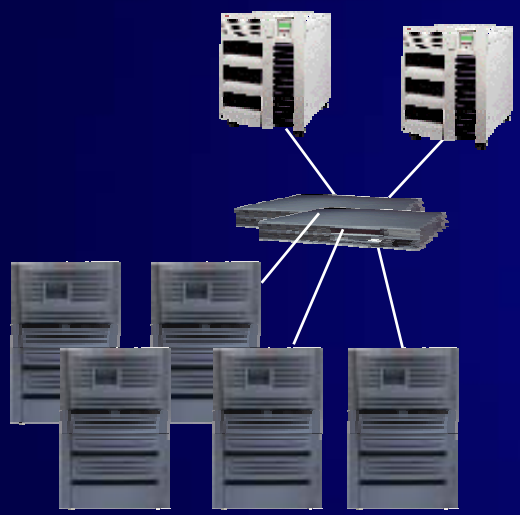

- Add processors and storage independently, when you need it
- Use cost-effective mid-range systems for “big box” performance
- Increase performance for all subscribers to the cluster
- Buy the best price/performance versus being locked into old constraints
- Budget predictably

Scalability

Only *Tru64* UNIX Has No Hard Limits

1TeraByte database files used to be exceptional.

Tru64 UNIX on 64-bit *AlphaServer* systems allows larger single files, and larger file systems, than any other UNIX

Vendors' UNIX	Max File System Size	Max File Size
HP-UX 11i	2TB	2TB
AIX 4.3	1TB	64GB
Solaris 8	1TB	1TB
Tru64 UNIX 5.1	16TB	16TB

Now, your data can grow with your business.

Availability

Oracle9*i* Real Application Clusters Data Protection At Every Level

All nodes see all the data

- More Computers = More Reliable
- With TruCluster Server, all nodes do real work
- All remaining nodes can take on the work of a failed node, then balance the load.

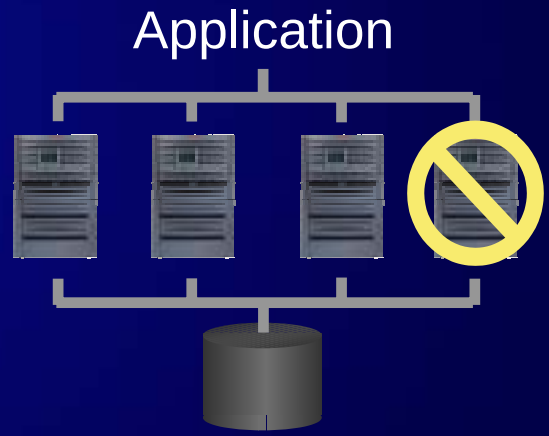

Oracle9*i* Real Application Cluster

Availability

TruCluster Server Scores in High Availability

#1 in Clustering RAS

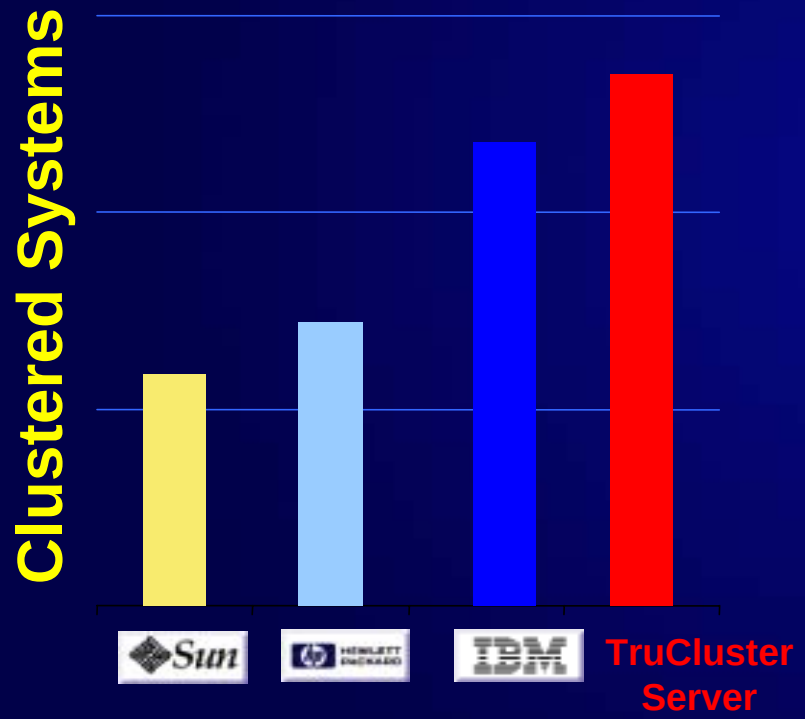

Source: DH Brown, March 2000, "Competitive Analysis of UNIX Cluster Functionality"

#2 in Single System RAS

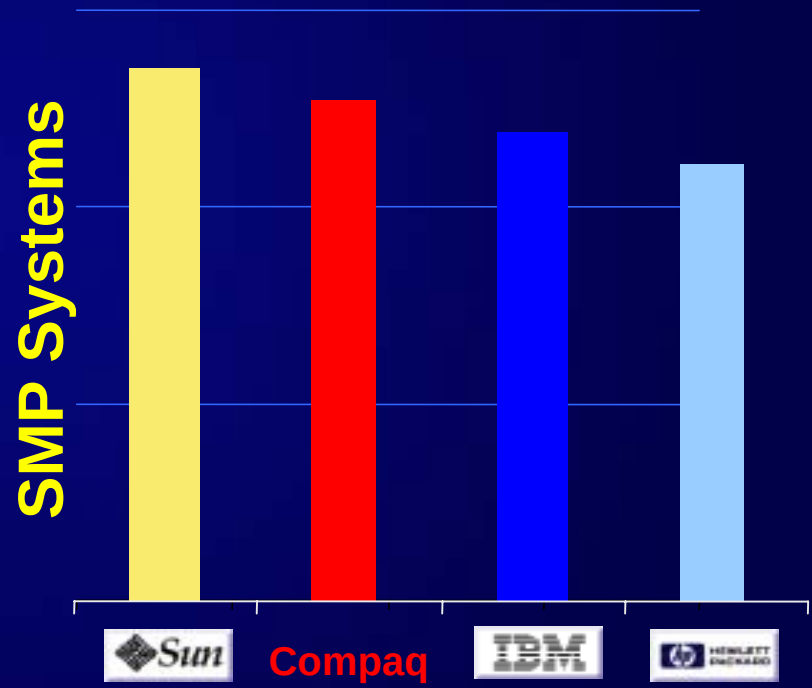

Source: DH Brown, May 2001, "Single System High Availability"

Ranked #1 in High Availability RAS

- D.H. Brown HA report confirms Tru64 UNIX, AlphaServer and Storage
  - **First** in Cluster RAS
  - Second in Single system RAS
  - **First Overall**
    Best in Class

Composite Score
Cluster + Single System RAS

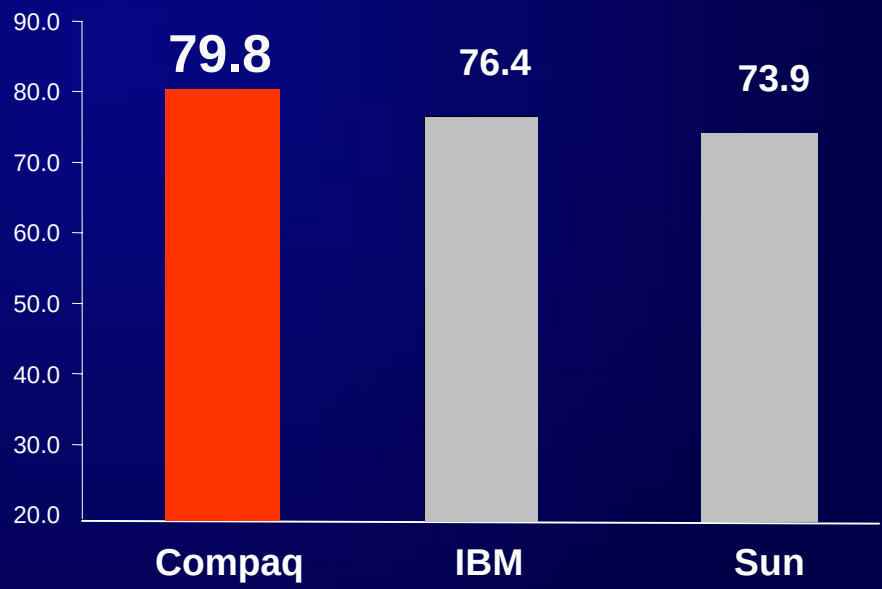

DH Brown High Availability,7/01

So clusters provide resilience against failure..

Because the unexpected happens!

Manageability

Cost Savings Vs Single Systems and Traditional UNIX Clusters

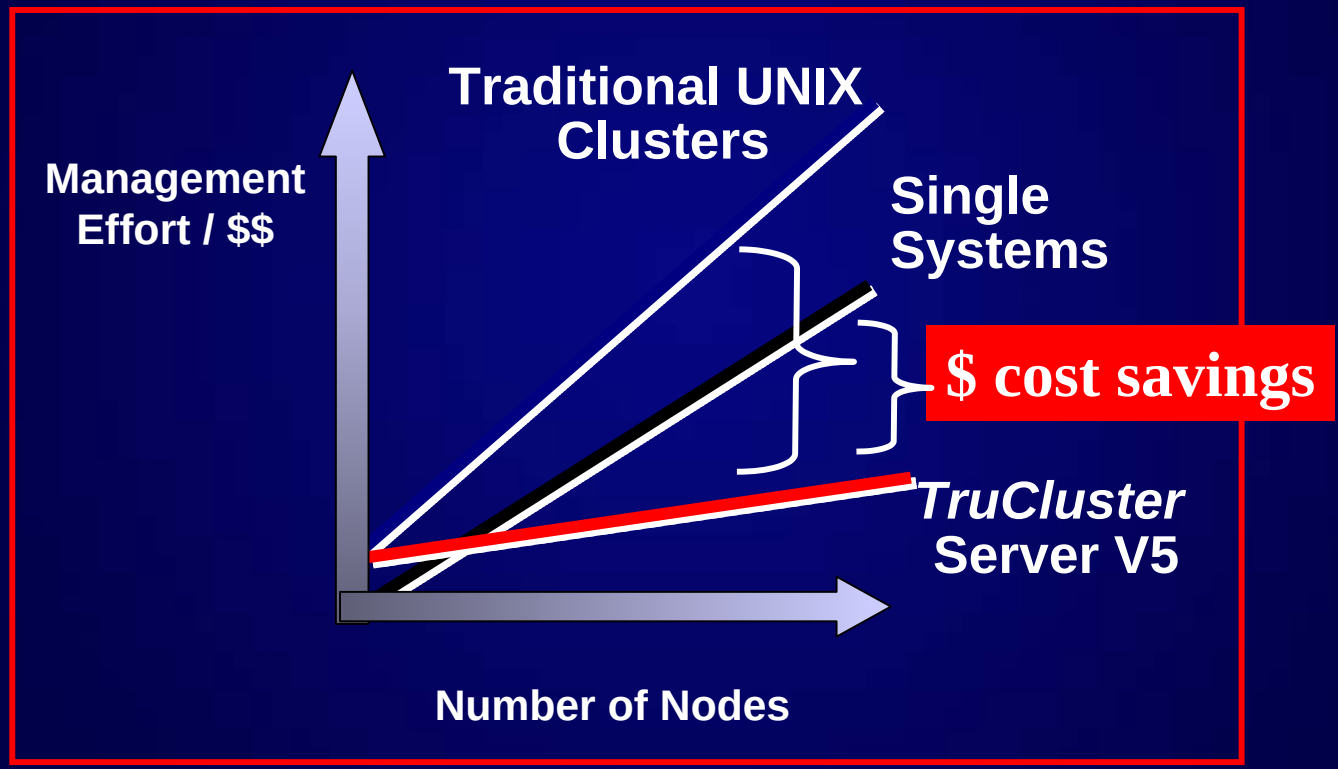

Manageability

Low TCO — Unique on *TruCluster* Server

Reduced management costs
-- up to 82%

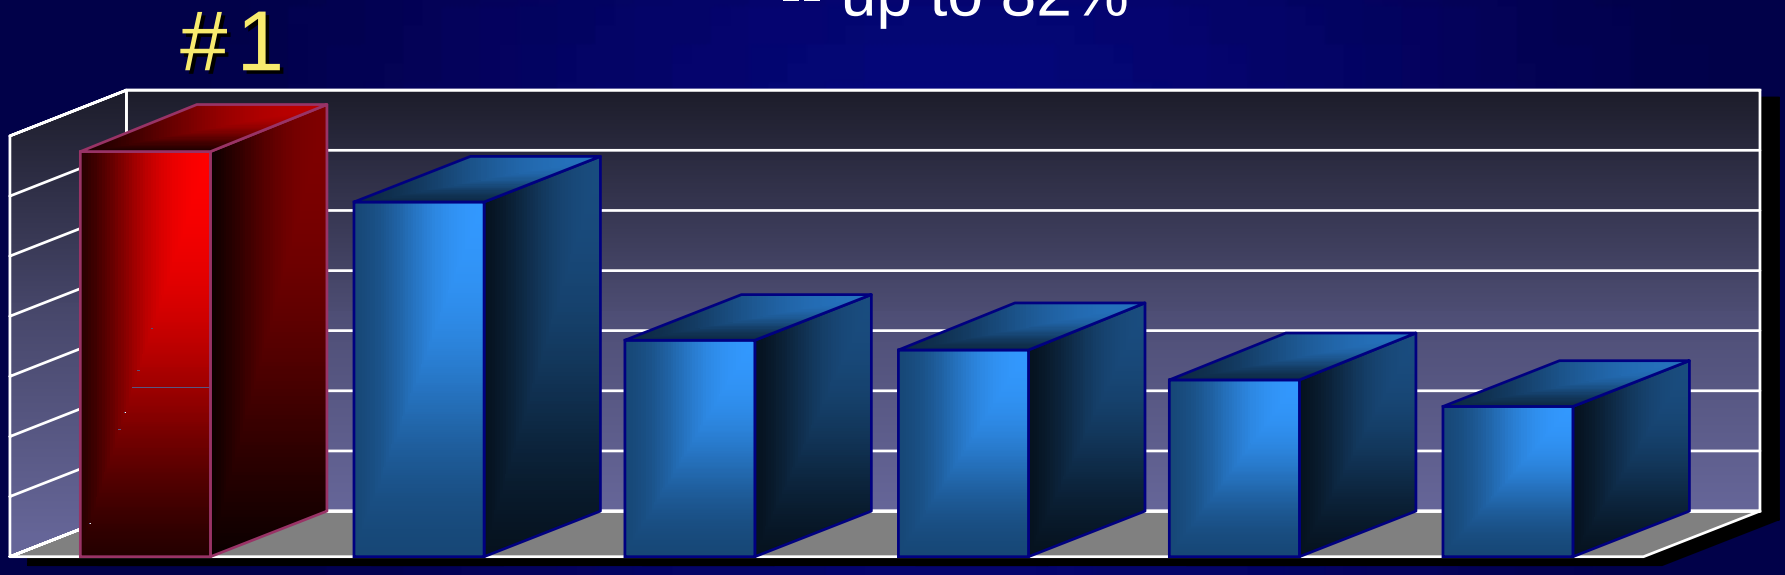

(Maximum score = 65% Average score = 48.4%)

Maximum score = 65% D. H. Brown, March 2000
Compaq placed 1st in the Competitive Analysis of UNIX Cluster Functionality, March 2001 – 2nd year in a row

Compaq's Unique Advantage for Oracle 9i RAC (Why Compaq?)

Together Tru64 UNIX and Oracle 9i RAC Deliver *Unique* Leadership Database Clusters

- **Scale –**
  - Tru64 UNIX #1 in scalability (DH Brown, 3/2001)
  - High speed memory (low latency) cluster interconnect optimizes message passing in database
  - Performance – cluster wide file system, direct I/O, RDG service
- **Management**
  - Extends database management efficiencies with Cluster File System “manage as one” abilities
- **Availability**
  - Reference configuration platform
  - Better integration with other features in the cluster
- **Cost of Ownership**
  - Lowest TCO with TruCluster Server SSI management

The 3 Database Challenges
#1: Scalability

- Scale to millions of users
  - Transparently
  - All types of applications
- Scale workloads without limits
- Grow storage easily

Real Application Clusters

Tru64 UNIX is #1 in Scalability

DH Brown, 3/2001

Unique Features

- Cluster Interconnect (Memory Channel, RDG S/W)
- 256 GB physical memory
- Large file systems up to 16 TB
- Cluster File System with shared root, single system image
- Dynamically grow/contract in size (advfs)
- Two types of interconnect: Memory Channel and LAN connect

DB Capabilities

- Industry leading performance -- optimized message passing with a high bandwidth low latency interconnect
- Bigger in-memory database caches
- Grow database easily without retuning or bringing database down
- Rapid deployment as you add a new node in as little as 15 minutes
- Add new drives/change drives and integrate into Oracle db proactively before running out of disk drive space

Tru64 UNIX on AlphaServer: Industry Leading Performance

Unique Features

- 256 GB large memory (V VLM)
- Asynchronous concurrent I/O to the file system with (near) performance of raw disk
- Memory channel interconnect and reliable data grams (RDG)

DB Capabilities

- Bigger in-memory database caches
- Bypass O/S file cache and issues I/O requests asynchronously
- Minimize block modification contention (Oracle 9i Cache Fusion)

TruCluster V5.1A CFS File I/O

Most file system reads are cached and direct, some metadata served

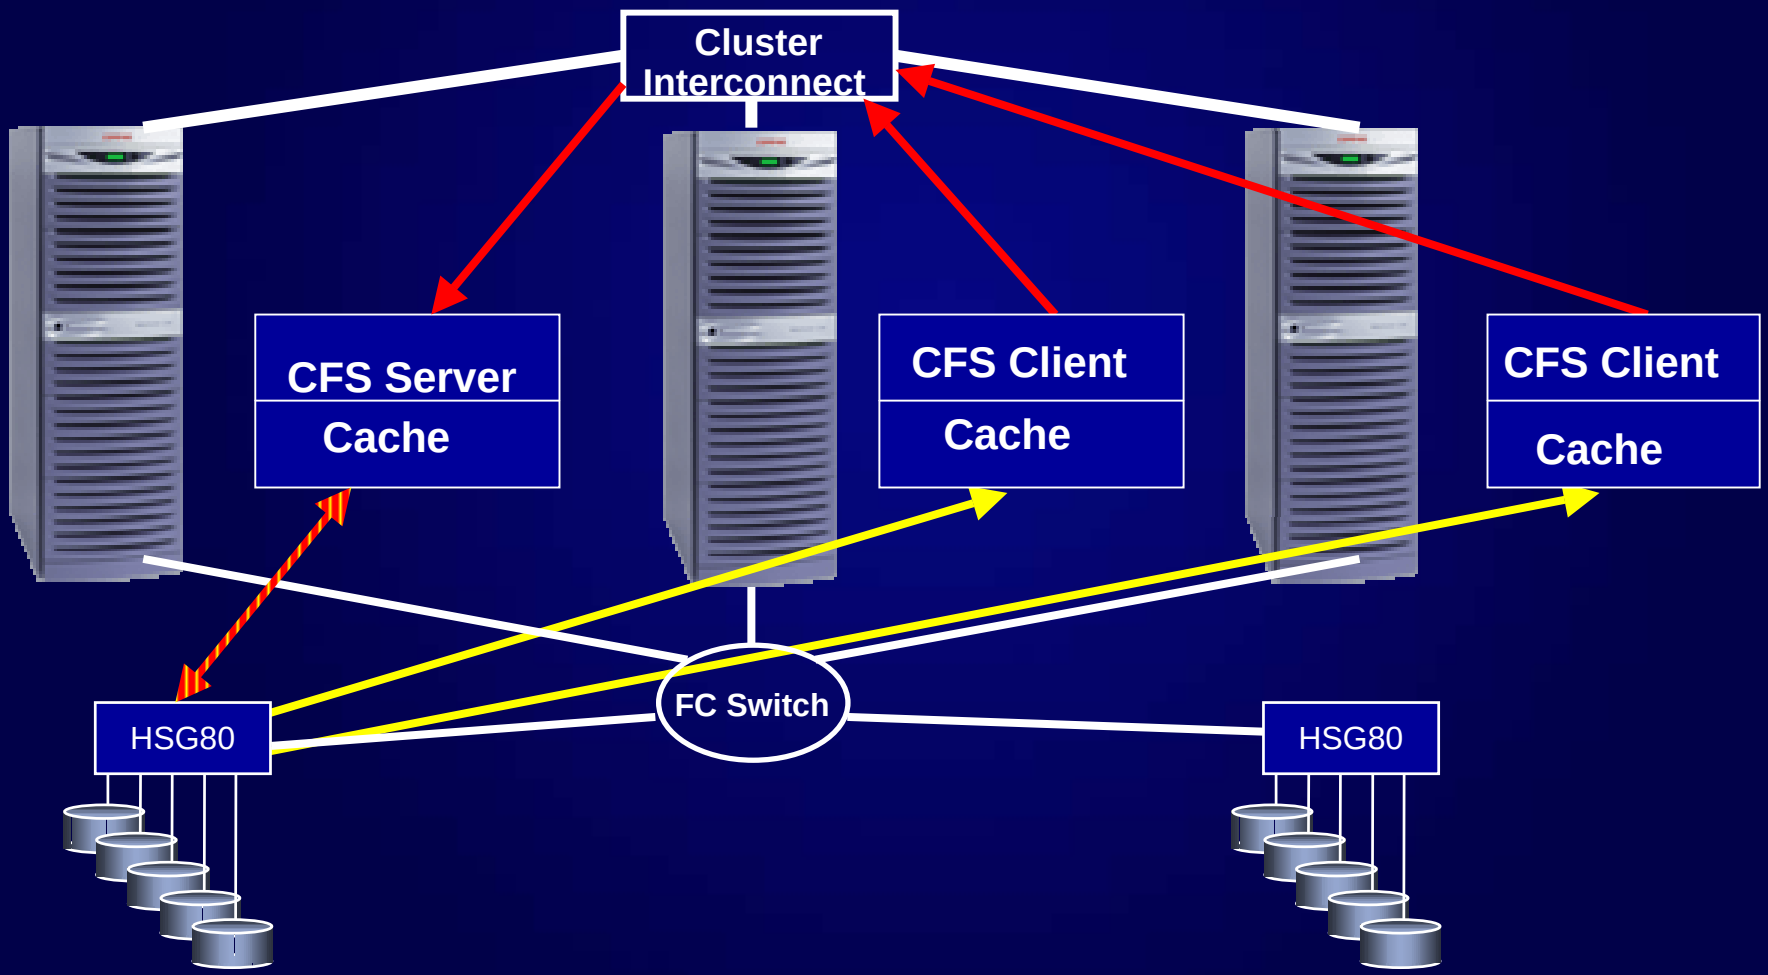

Memory Channel Cluster Interconnect

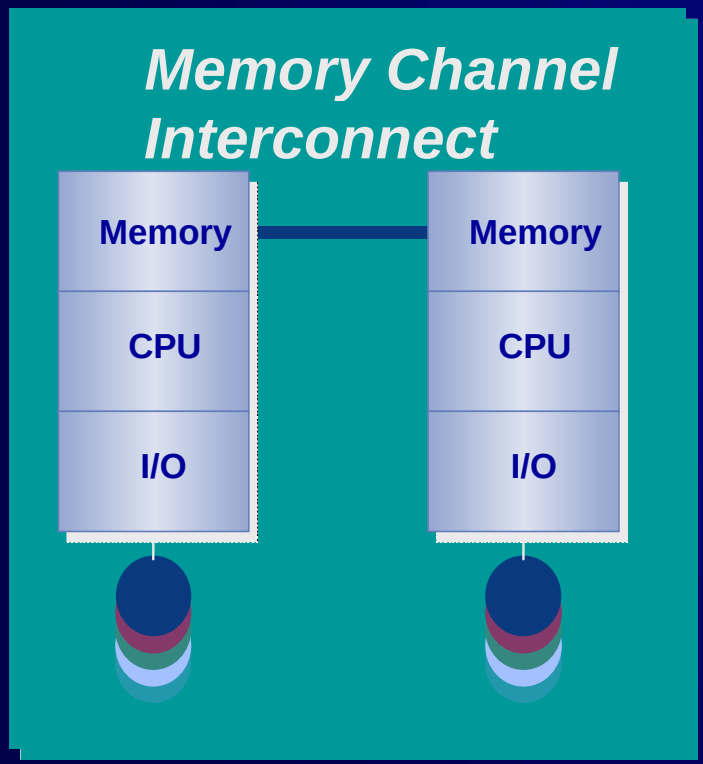

High-Speed **direct memory connection from system to system**

Very low cluster I/O overhead

- ***No*** **protocols**
- ***No*** **specialized programming**
- ***Very low*** **system overhead**
- ***Low*** **latency**
- ***High*** **bandwidth**

The 3 Database Challenges
#2: Availability

- Be available all the time
  - Design for fault tolerance
  - Provide fast and reliable fault recovery
  - Eliminate maintenance downtime
- Support mission-critical business operations

Real Application Clusters

Highest Availability

Unique Features

- **CAA better integrated with Tru64 UNIX (kernel)**
- **Cluster alias allows multiple layers of redundancy**
- **Fully dynamic load balanced multi-path I/O on a per device granularity**
- **Multipath I/O bundled in base O/S at no additional cost**

DB Capabilities

- Transparent failover of 9i components enables planned maintenance without impact to client and faster failover
- Transparent handling of node failures
- Transparent handling of storage failover
- Transparent availability of file systems

The 3 Database Challenges
#3: Management

- Create one virtual system to configure and manage
- Single system image for the database integrated with the cluster
- Integrated Tru64 UNIX Management Tools

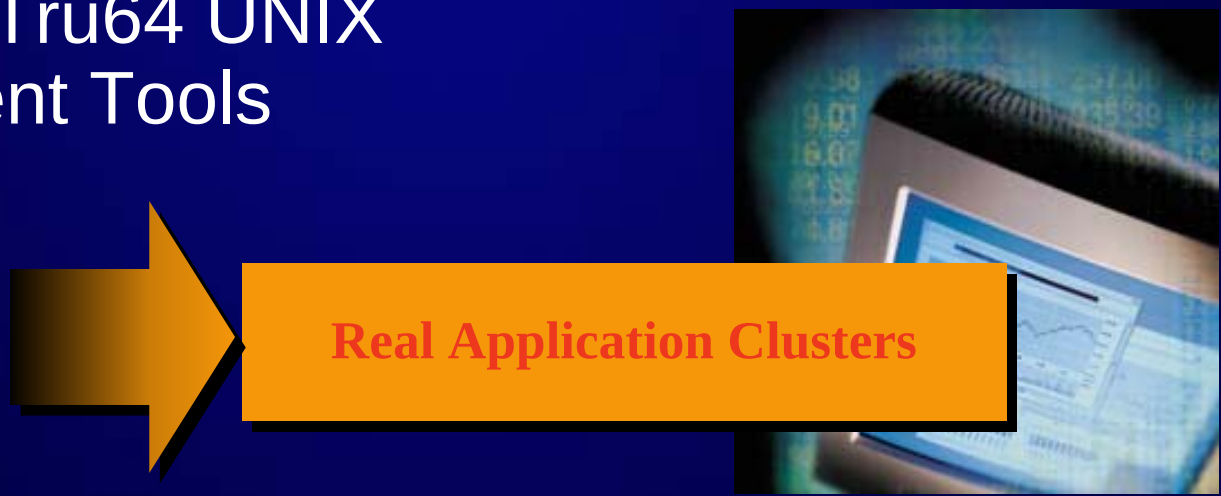

Easiest Management

Unique Features

- CFS single system image with shared root
- Cluster wide storage; single name space
- Direct I/O across cluster delivers performance of raw disk and is simpler to manage with a cluster file system
- AdvFS features: add, move, reduce volumes and file system sizes dynamically (ONLY Tru64 UNIX can do all 3!)

DB Capabilities

- Easier to manage when data is in a file system instead of raw disk (raw disk performance without management challenges)
- Basic administration of RAC is identical to Oracle db in non-clustered environment
- Manage from anywhere "once" (Other cluster management requires multiple steps on multiple nodes -- manually)
- Install once; Easy to get up and running quickly
- Online resize of database

Single System Management

LOWEST Management Costs -
Install: Compaq 79% less than HP & IBM, 81% less than SUN
Manage: Compaq 19% less than HP & IBM, 33% less than SUN

	Hours Per year			
	Compaq	HP	IBM	Sun
Add 6 nodes	5	8	8	10
Plan & Config 8 nodes	0	16	16	16
Monitor & tune	416	520	520	520
Small Pgm Testing	120	180	180	300
Respond to Failovers	200	240	240	320
User Accounts	26	35	35	52
Other special Projects	36	60	60	108
Meetings	210	210	210	210
Total Hours Saved	**1013**	**1269**	**1269**	**1536**

INSTALL

MANAGE

TruCluster Server 5.0 Slashes Cost of Administration by Streamlining Operations
D.H.Brown June 2000 http://tru64unix.compaq.com/cluster/DHB-Cluster-Admin.pdf

The Analysts Agree

Giga Information Group

"It is actually less work to manage [a TruCluster of systems] than to manage separate systems"

Illuminata

"Compaq's Tru64 UNIX is the standout in this pack, the only one to address RAS in its full profile"

D.H. Brown

Compaq placed 1st in the Competitive Analysis of UNIX Cluster Functionality, March 2001 – 2nd year in a row

"A Compaq sysadmin would take 127 days, an HP or IBM sysadmin would take 159 days, and a Sun sysadmin would take 192 days, to complete these selected tasks"

Manageability

TruCluster Server

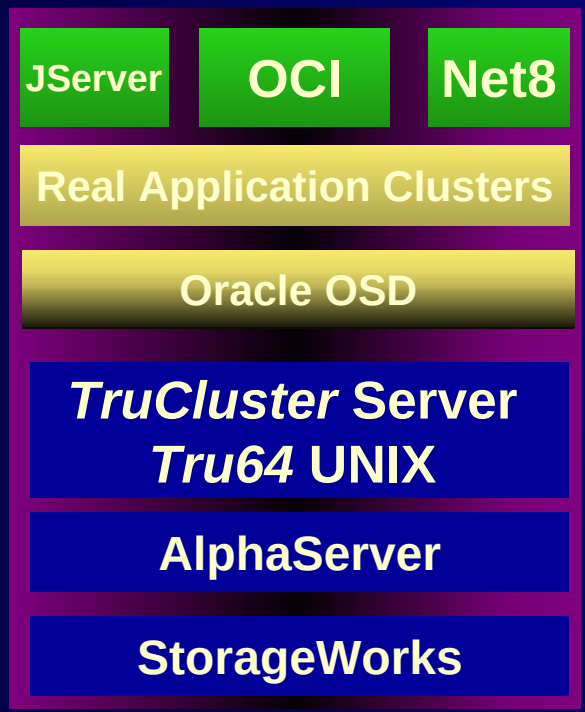

Sun Veritas SANPoint cluster stack

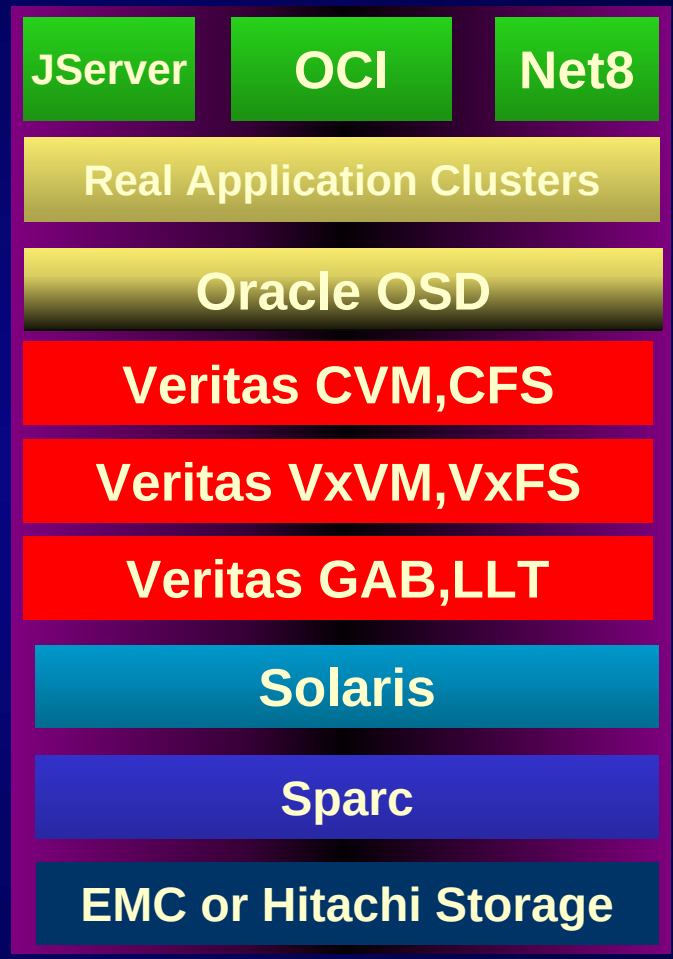

Which would you rather manage?

Moving from Oracle8*i* or Oracle9*i* to Oracle9*i* Real Application Clusters

On TruCluster Server -

- Recompile Oracle to recognize Real Application Clusters
- Tune parameters
- Leave your database where it is.
- Manage as usual.

On Sun, HP, etc

- Recompile Oracle to recognize Real Application Clusters
- Tune parameters
- Backup your database
- Restore it to raw partitions
- Manage raw partitions, in addition to the system, from then on.
- No autoextend

It's easier with *TruCluster* Server

AlphaServer Product Family

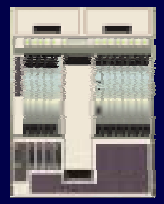

AlphaServer
TS Series

New
SC45

AlphaServer
SC Series

DS10L

DS10

uni

DS20E

dual

AlphaServer DS Series

New

ES40 ES45

4-way

AlphaServer
ES Series

GS80/160/320
8, 16,
32-way

AlphaServer
GS Series

AlphaServer GS320 –
a record-breaking
230,533 tpmC
June 2001
with *Tru64* UNIX

Top Single System TPCC Results

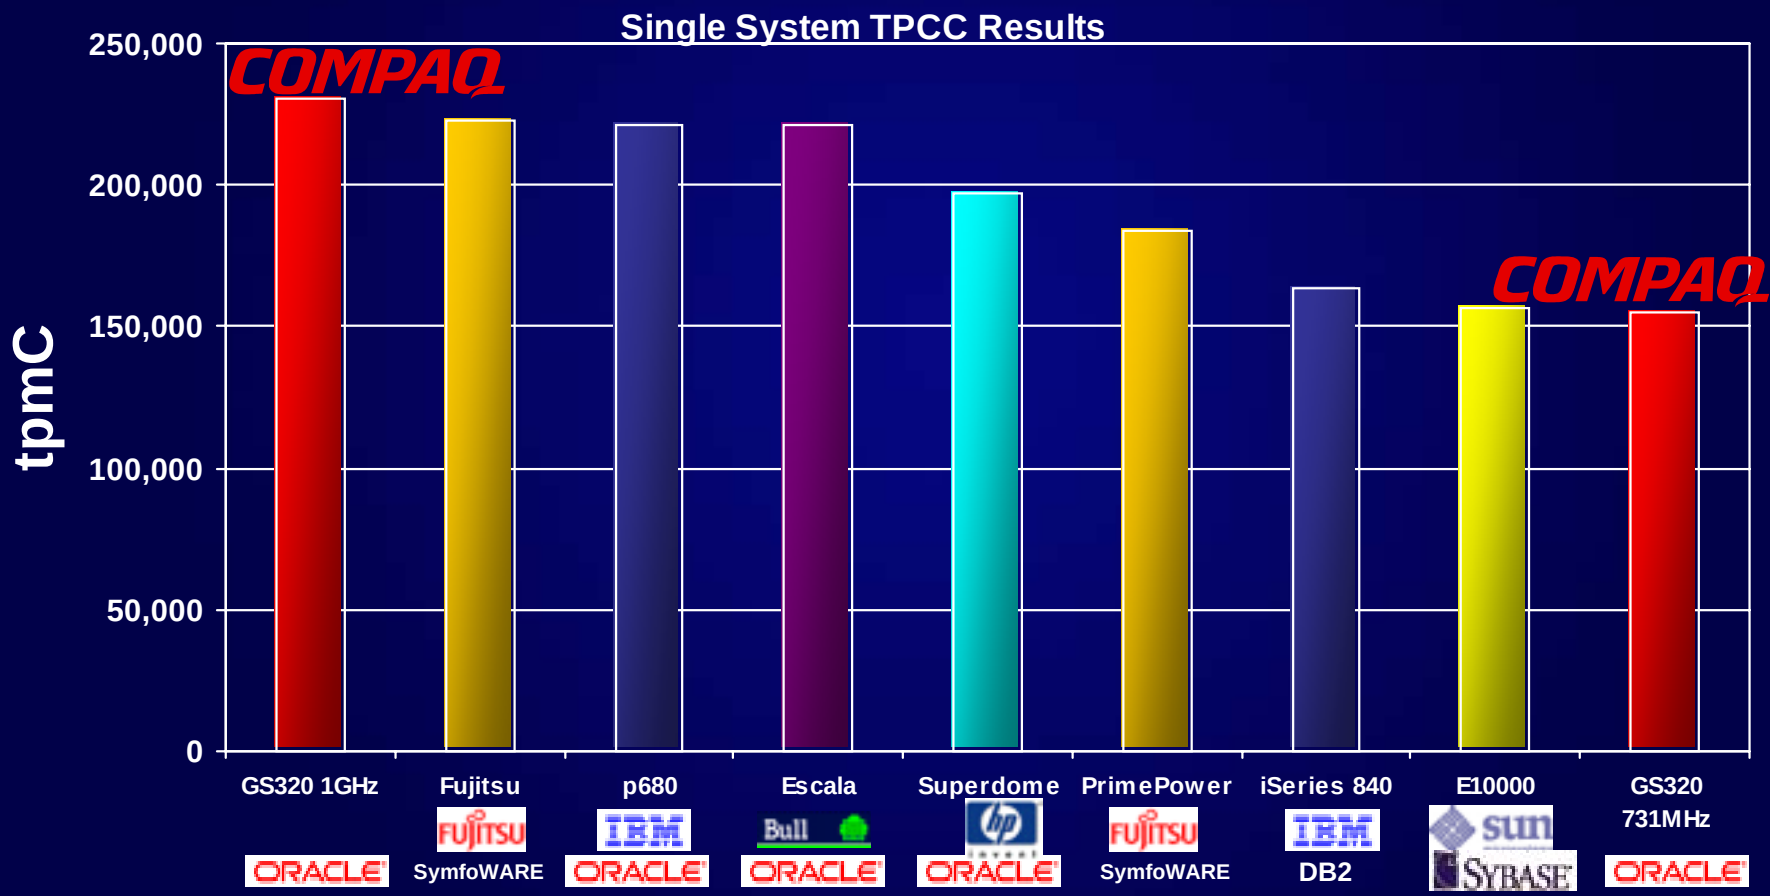

Updated 5/9/2001

Compaq Certifies Its Storage with Oracle Test Suites

Oracle Backup Program

Oracle Storage Compatability Program

Compaq performance, reliability, manageability

Leadership in Open SANs

Compaq Data Replication Technology

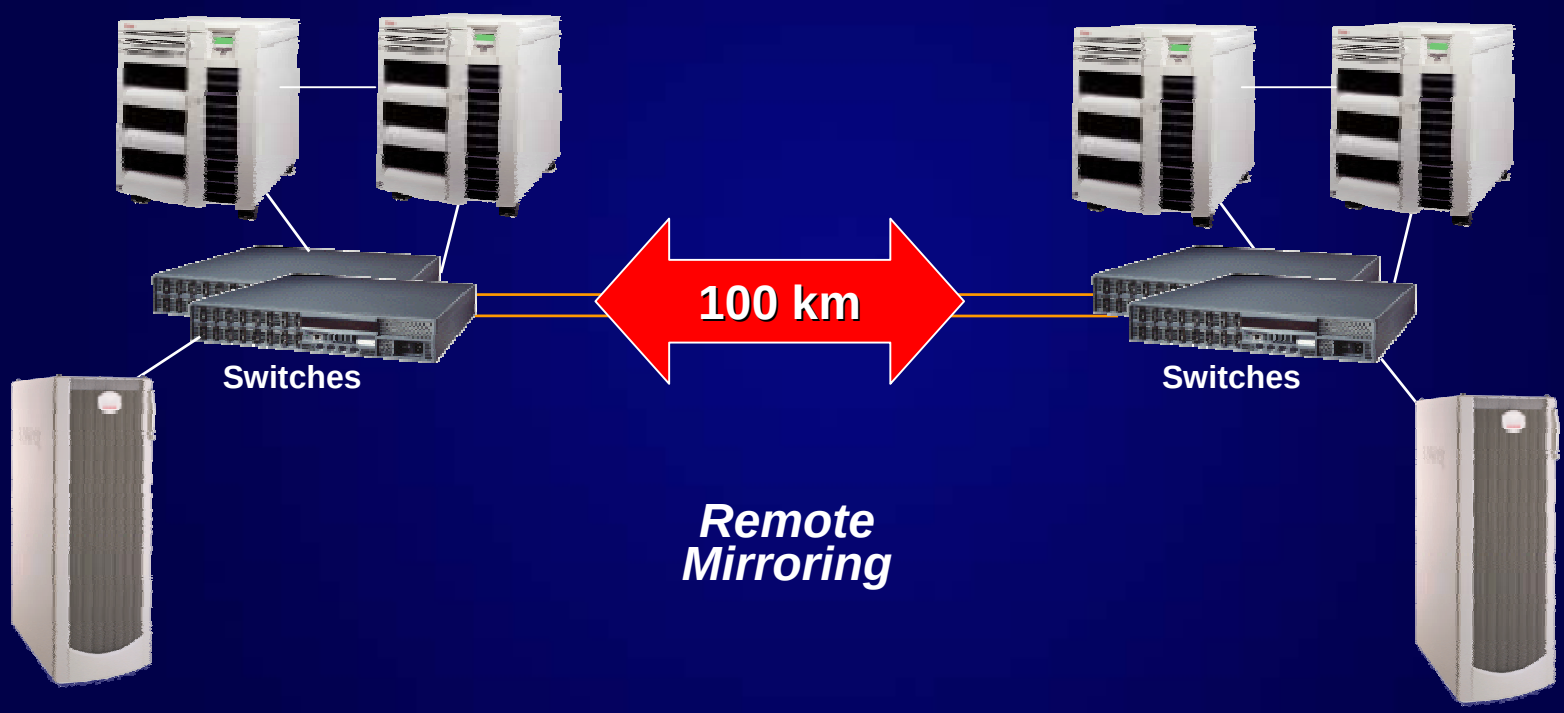

Preserves system performance for on-line transaction processing and application throughput

Industry SAN Offering Assessment

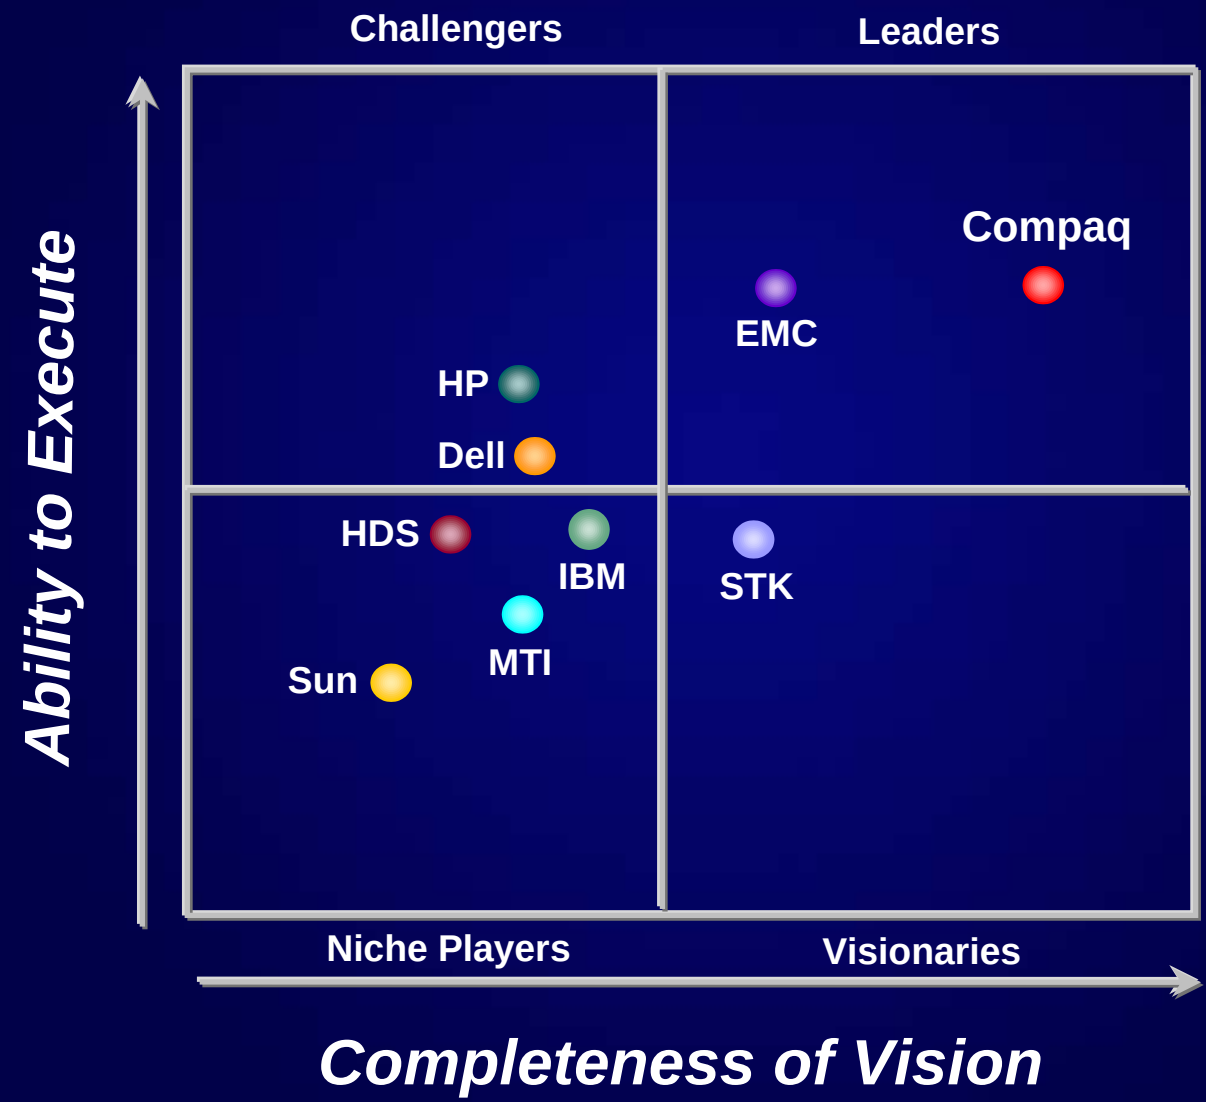

Source: GartnerGroup September 2000

SANs Deliver for Your Enterprise Database

- Flexibility
  - Scalability (capacity, performance, accessibility)
- Service Levels
  - Data and application availability
  - Throughput performance
- Efficiency
  - Centralized management
  - Server and storage consolidation

Backup terabytes of data in hours versus days, without taking the database offline

Growth with Flexibility

AlphaServer Database Engines

Scale Up

Upgrades

SMP

Capacity on Demand
Partitioning
Configure/Reconfigure storage
Add nodes/balance workload
Remote management and control
Resource management and reporting

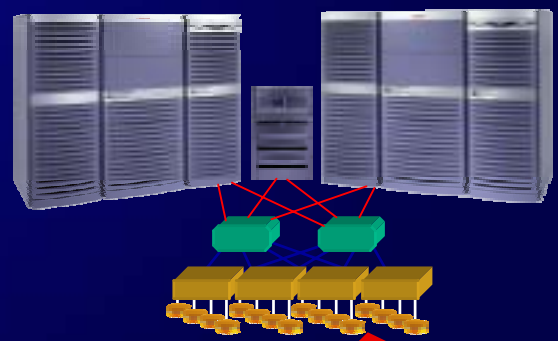

Clusters

Virtual storage

Scale Out

Single system image

Compaq / Oracle Joint Architecture Services Modules

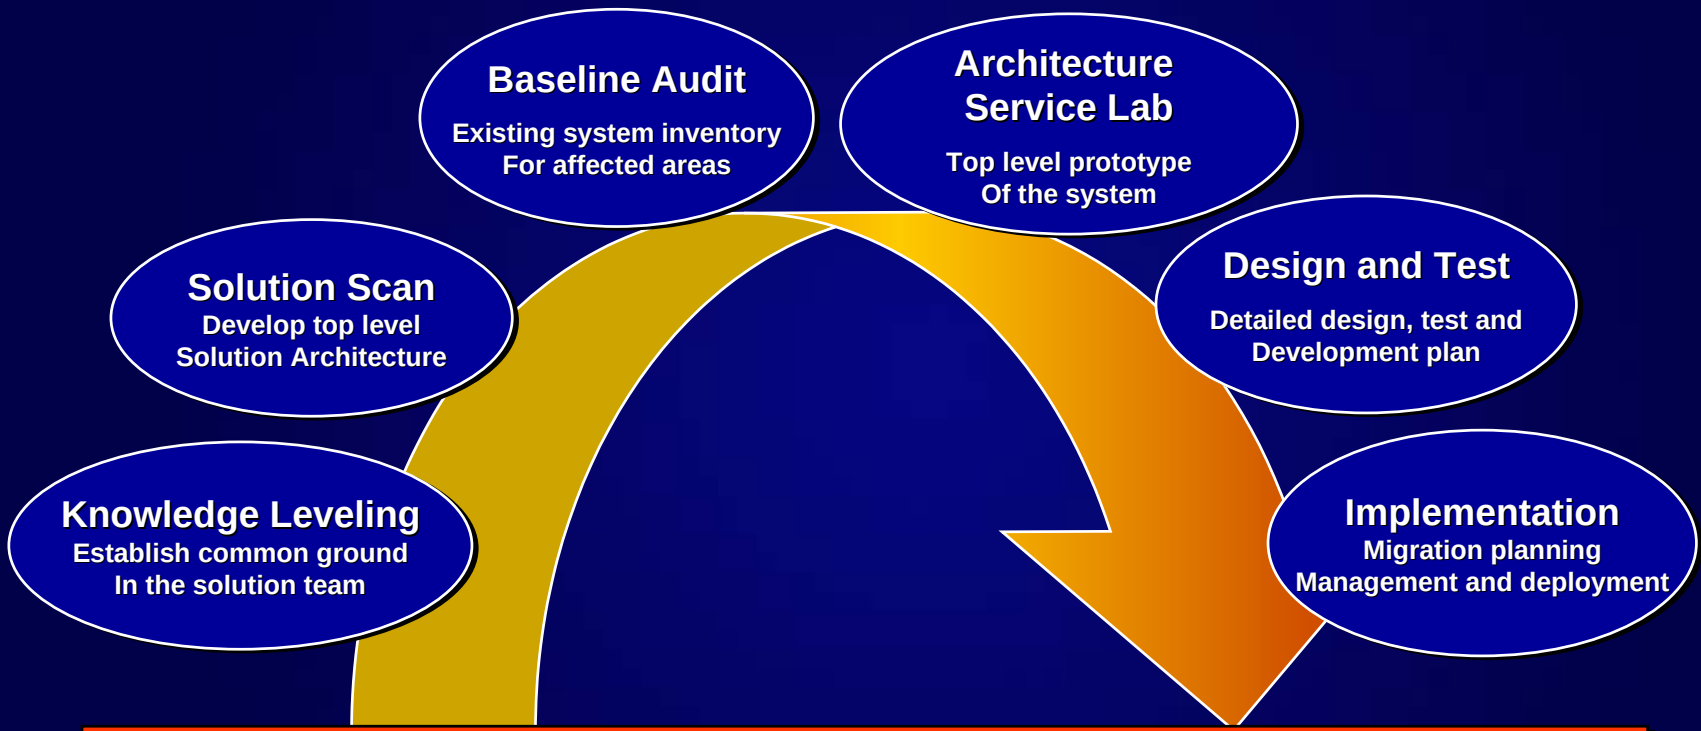

Oracle9*i* Real Application Clusters Consulting

Implementation, migration, and integration

Integrated Compaq / Oracle On-going Support

Benefits to you

Simplest call escalation

No finger-pointing

Quickest resolution

Reliable answers at the solution level

Added value at no added cost

Simpler, faster problem resolution keeps your business on-line and your staff on-schedule

Other vendors approach clustering this way…

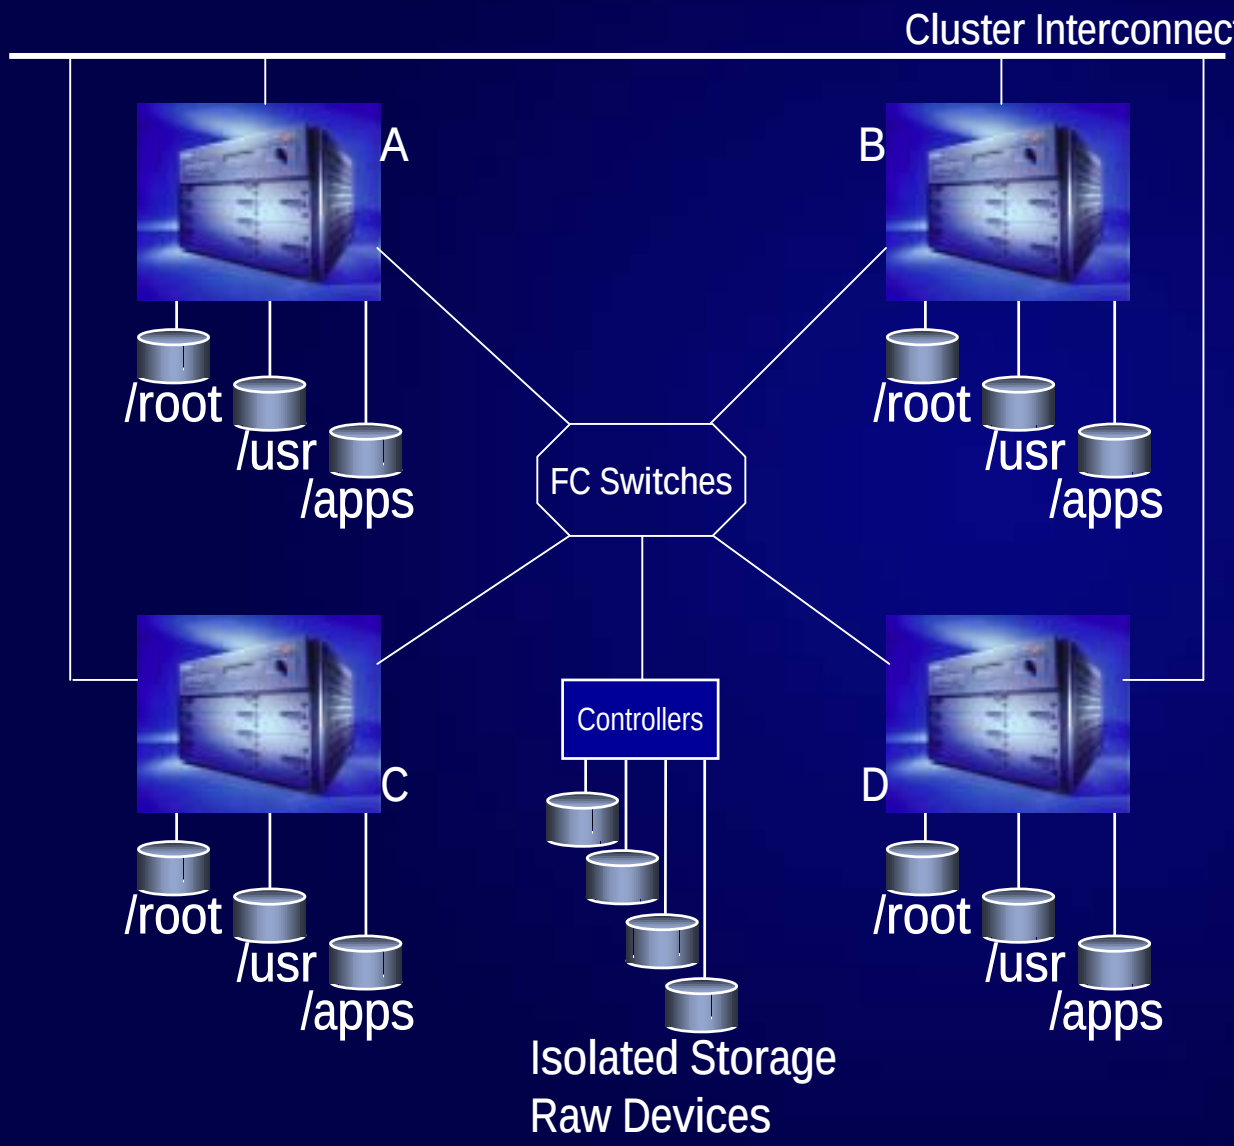

A – Install UNIX & patches
Install the apps & patches
Configure the storage

B– Install UNIX & patches
Install the apps & patches
Configure the storage

C – Install UNIX & patches
Install the apps & patches
Configure the storage

D – Install UNIX & patches
Install the apps & patches
Configure the storage

E – Install UNIX & patches
Install the apps & patches
Configure the storage

F – Install UNIX & patches
Install the apps & patches
Configure the storage

G – Install UNIX & patches
Install the apps & patches
Configure the storage

Tru64 UNIX Clustering works this way…

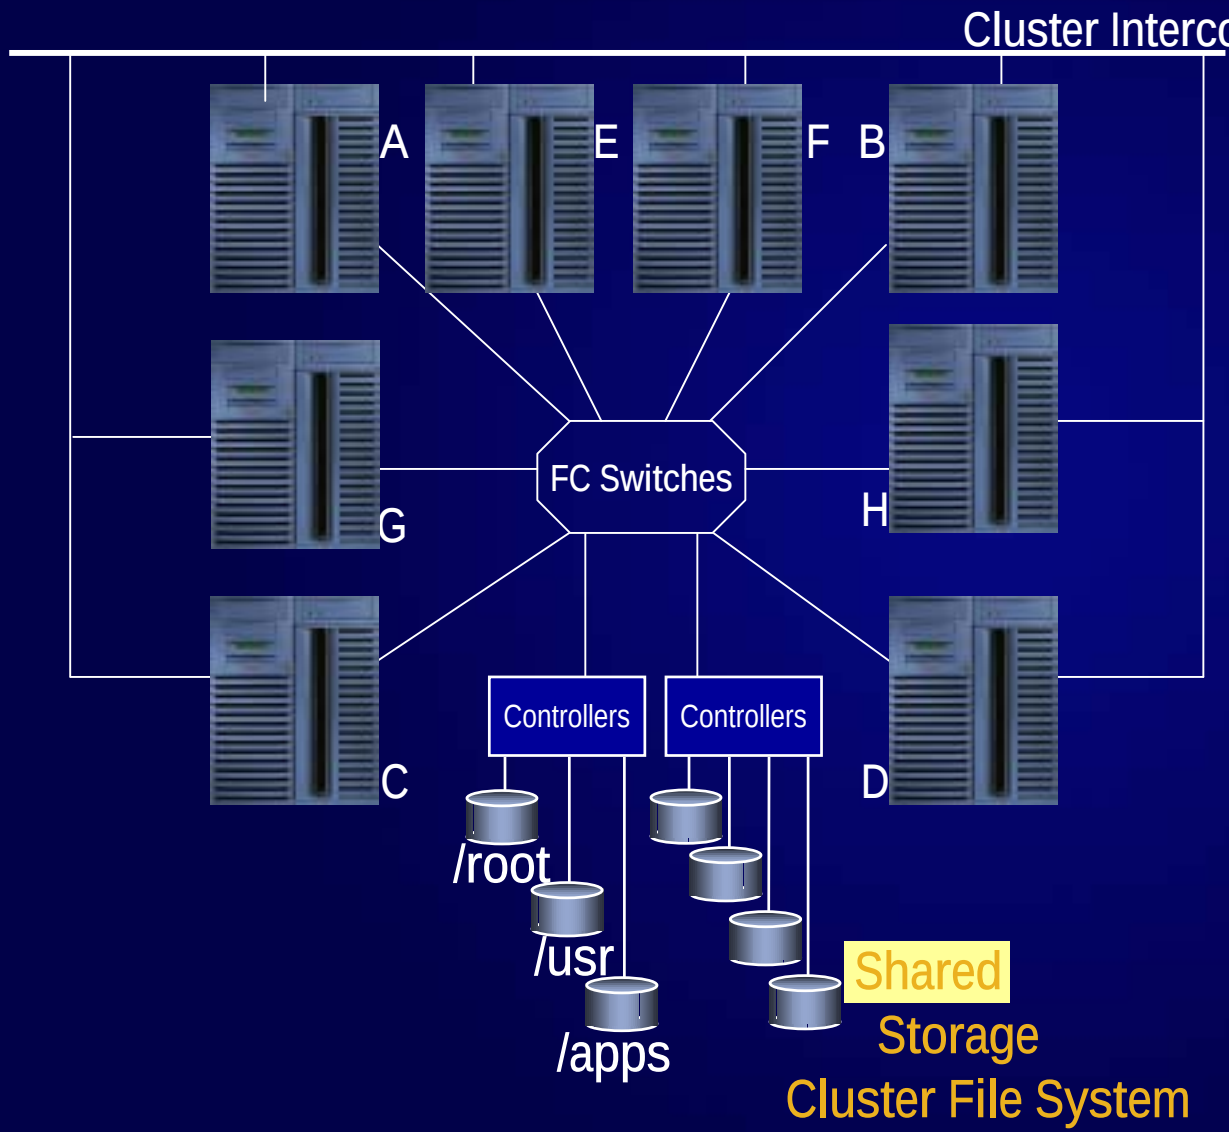

A - Install OS & patches
- Configure the storage
- Run clu_create
- Install apps & patches
- Run clu_add_member for each node

Why Oracle Chose Compaq Tru64 UNIX Clusters: *Delivering Unique Clustering Capabilities*

- Cluster ***Experience*** – 20 years
  - Key technologies for database scalability
- ***Unique*** Single system image
  - Reduces system management costs
- True Cluster file system
  - Simplifies system management

WHO HELPS ORACLE, THE BEST IN E-BUSINESS SOFTWARE, GET EVEN BETTER?

COMPAQ.

COMPAQ
hp & compaq. the new power of invention

9i RAC references on Compaq TruCluster

COMPAQ
hp & compaq. the new power of invention

Compaq Catches TIGERS.

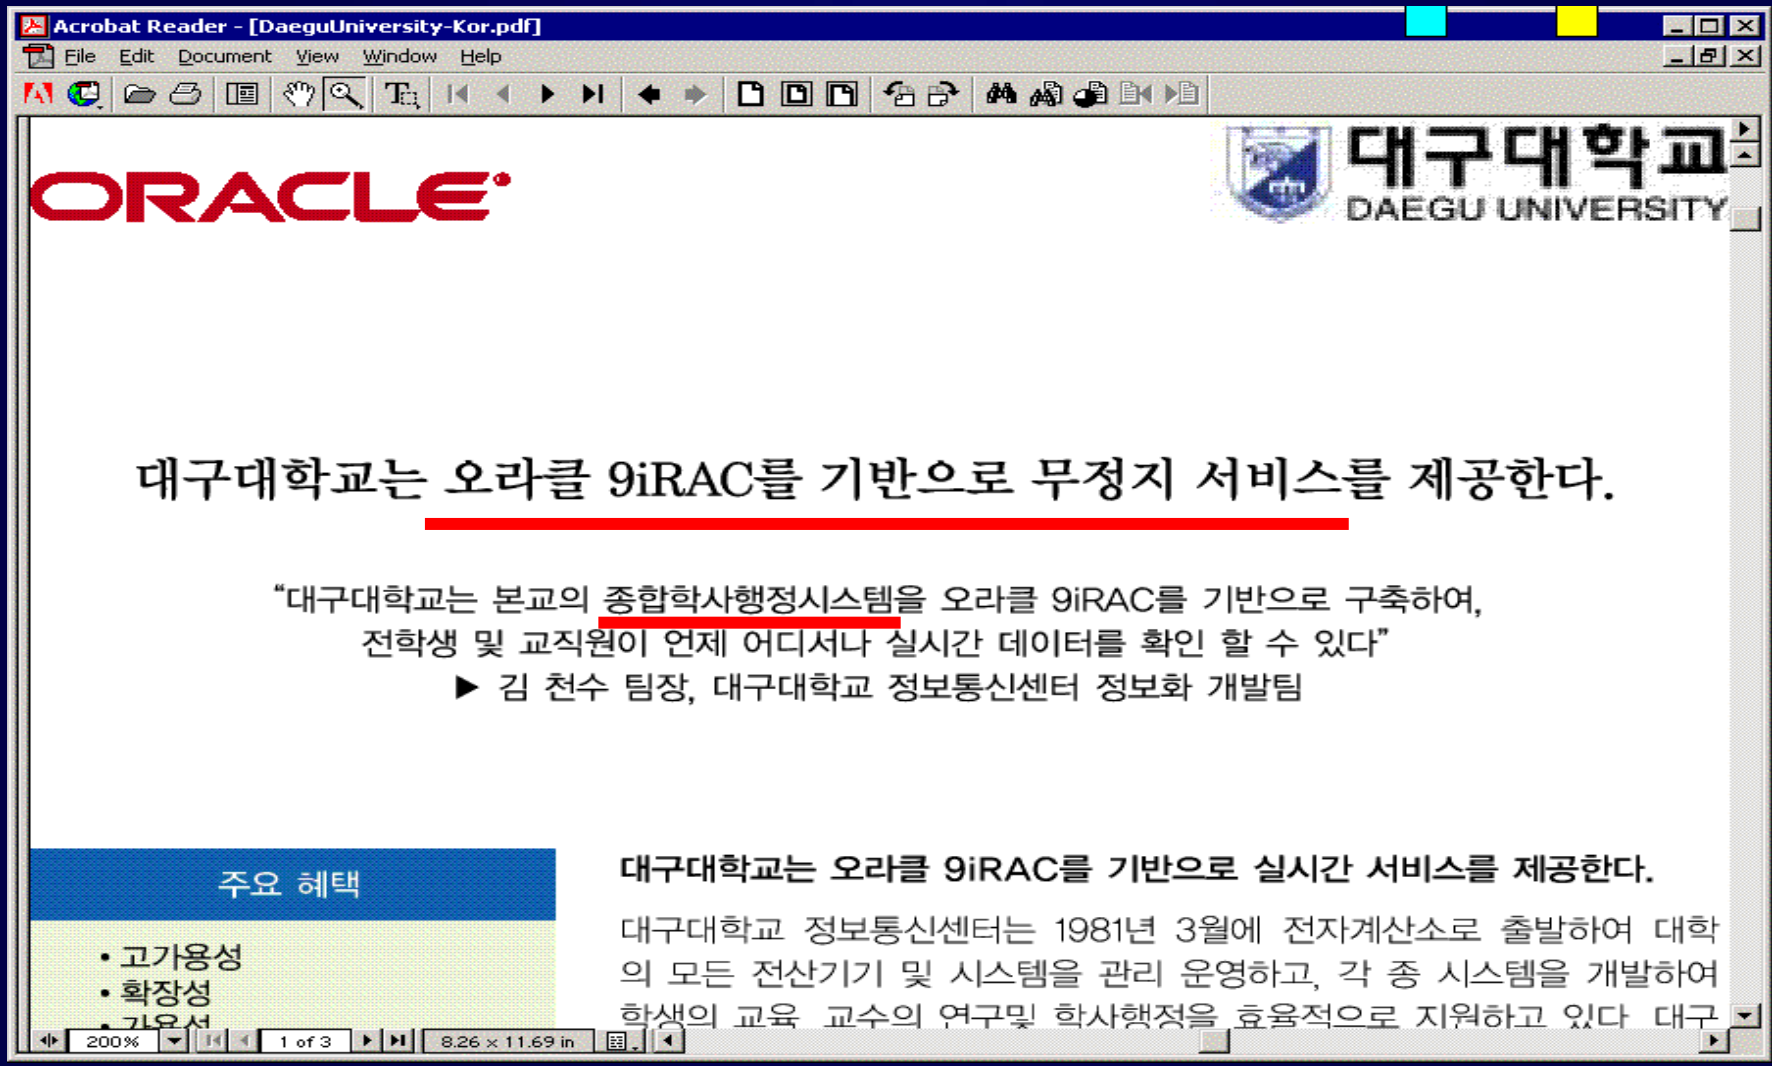

COMPAQ
hp & compaq. the new power of invention

Acrobat Reader - [DaeguUniversity-Kor.pdf]

File Edit Document View Window Help

데이터베이스 크기
20 GB

하드웨어
컴팩 ES40

운영체계
Tru64 유닉스

CPU
각각 4 CPU

메모리
2 노드
A 노드=4GB, B 노드=6GB

이것은 1996년 10월에 시
1997년 5월에 오라클 8i s
버 환경에서 윈도우 프로
학사행정시스템은 국내 최
대학정보시스템 개발의 좋

대구대학교는 학생들과 교
하여, 웹서버 환경에서 수
사행정관리시스템에 있는
9iRAC으로 업그레이드 하
약 두 달간에 마이그레이션
업 후, 변경된 환경하에서
동시에 클라이언트 환경에

300% 1 of 3 8.26 x 11.69 in

대구대학교 정보통신센터의 소속 직원 수는 일반직 20명, 조교 5명으로 총 25명으로 구성되어 있고, 시스템 개발 및 운영인원은 8명의 인원으로 최대한의 생산성을 유지하기 위하여 불철주야로 급변하는 대학 환경에 대한 지원을 원활히 하고 있다. 또한 데이터베이스 관리자(DBA)는 현재 2명이 담당하고 있지만, 현업의 업무가 별도로 주어져있어 데이터베이스 관리는 부가적인 업무임에도 오라클 9iRAC의 우수한 유지보수 기능으로 별다른 어려움 없이 원활한 관리가 되고 있다.

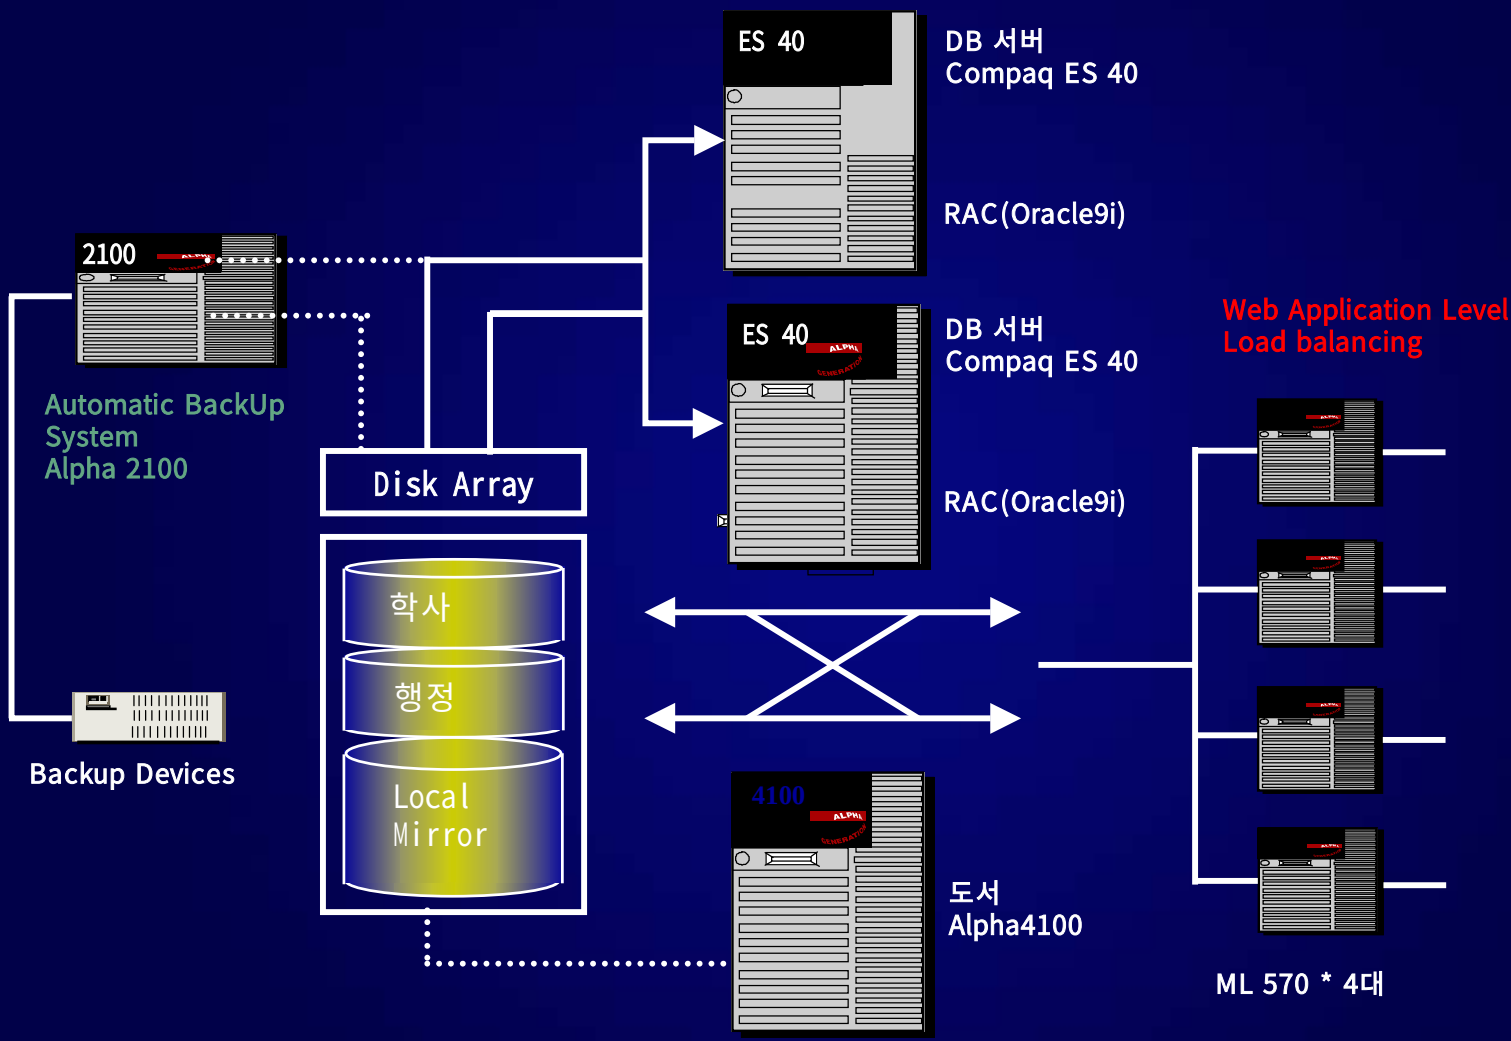
ES 40
DB 서버
Compaq ES 40
RAC(Oracle9i)
2100
ES 40
DB 서버
Compaq ES 40
RAC(Oracle9i)
Web Application Level
Load balancing
Automatic BackUp
System
Alpha 2100
Disk Array
학사
행정
Local
Mirror
Backup Devices
4100
도서
Alpha4100
ML 570 * 4대

COMPAQ
hp & compaq. the new power of invention

Chicago 911 System
Tru64 UNIX, TruCluster & Oracle

COMPAQ OneTeam ORACLE

www.oraclecompaqclusters.com

COMPAQ