





# Oracle Real Application Clusters 10g on the HP platform



Rebecca Schlecht  
EMEA HP/Oracle Competence Center



© 2003 CTC Development Company, L.P.  
This presentation is subject to change without notice



## Agenda

- HP/Oracle partnership
- Oracle Real Application Clusters Overview
- Oracle Real Application Clusters 10g Architecture
  - Automatic Storage Management
  - Oracle Clusterware
  - Automatic Workload Management
  - Rolling Upgrades
- HP Parallel Database Clusters on Proliant
- HP/Oracle Joint Support Services

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved.



## HP/Oracle Partnership



- 20 years of collaborative partnership
- 80,000+ combined customers... \$6B business
- parity for all Oracle products
- \$6 billion business
- a market leading >40% of Oracle customers use HP systems
- 400 people dedicated to Oracle .... 13 worldwide technology & competency centers




 

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

## EMEA HP/Oracle CTC Cooperative Technology Center

- Located at HP in Boeblingen/Germany, Munich/Germany, Sophia Antipolis/France and Oracle Reading/UK
- Oracle and HP employees in one team
- Founded in spring 1994
- Technology consulting for partners / customers
- Evaluation and tests of new products/features
- Technology transfer to and from US labs
- Technical pre-sales assistance
- Support of customer specific Oracle benchmarks
- Demonstrating the latest Oracle solutions on HP at the *etc* (Enterprise Technology Center) in Reading/UK

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

## Oracle and HP are not only partners .....



- Oracle9i Real Application Clusters and CRM are developed, optimized and tested on the HP platform
- HP TruCluster is a core technology used in Oracle9i Real Application Clusters
- Oracle runs corporate CRM, ERP and e-mail systems on HP servers
- Oracle uses an HP infrastructure to host customers CRM instances
- Critical Oracle customer support runs on HP servers
- HP has standardized on Oracle databases
- Worldwide Informix to Oracle migration @ HP



4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.



## Agenda



- HP/Oracle partnership
- Oracle Real Application Clusters Overview
- Oracle Real Application Clusters 10g Architecture
  - Oracle Clusterware
  - Automatic Storage Management
  - Automatic Workload Management
  - Rolling Upgrades
- HP Parallel Database Clusters on Proliant
- HP/Oracle Joint Support Services

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.




ORACLE DATABASE 10<sup>g</sup> 


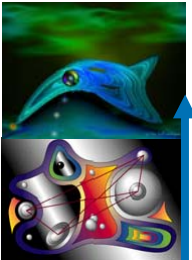
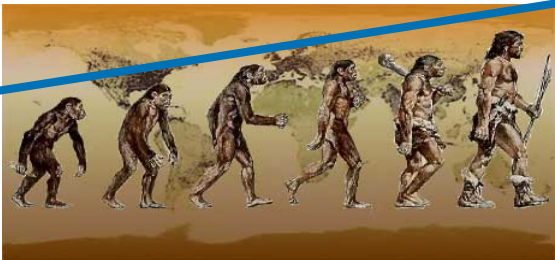
## Agenda

- HP/Oracle partnership
- Oracle Real Application Clusters Overview
- Oracle Real Application Clusters 10g Architecture
  - Automatic Storage Management
  - Oracle Clusterware
  - Automatic Workload Management
  - Rolling Upgrades
- HP Parallel Database Clusters on Proliant
- HP/Oracle Joint Support Services


4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

ORACLE DATABASE 10<sup>g</sup> 


## RAC Revolution & RAC Evolution




Oracle8i  
with OPS




Oracle9i Rel. 1  
with RAC




Oracle9i Rel. 2  
with RAC



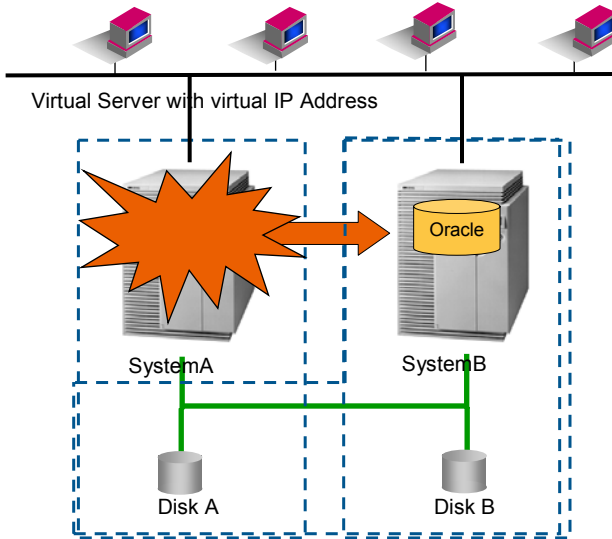
Oracle Database  
10g with RAC


4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

## Cold Failover Cluster with HP Serviceguard




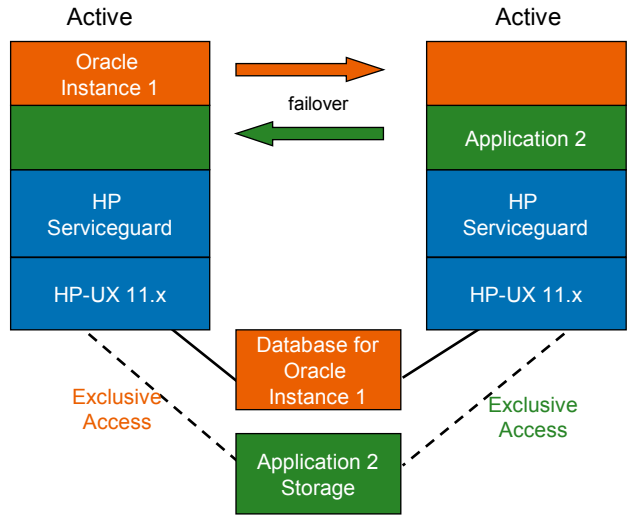
- Non-Shared Database
- Provides **16 node** failover solution
- Before Failover:
  - Virtual IP address and network name refers to system A
- After Failover:
  - Virtual IP address and network name refers to system B



4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

## hp Serviceguard Software Stack





4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 10



## Limits of Cold Failover

- Scalability of cluster is limited to scalability of one server
- Load cannot be distributed across all nodes in the cluster
- Cold failover is slow, as many time consuming tasks must be performed as part of failover
  - moving and mounting logical volumes
  - starting the oracle instance
  - opening the data files
- After failover, the instance caches are cold introducing a performance brownout

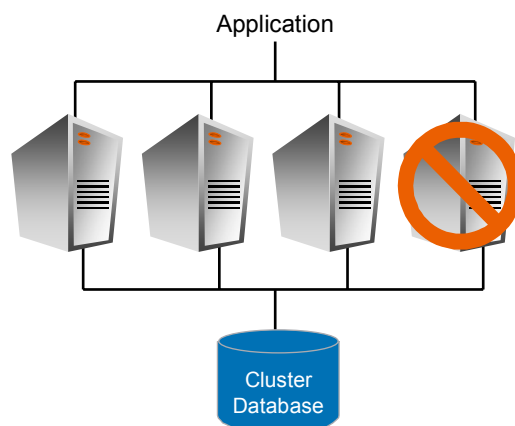
4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.



## Oracle9i Real Application Clusters

- High Availability - more computers = more reliable
- Scalability - more computers = more speed



4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.





## Oracle9i Real Application Clusters

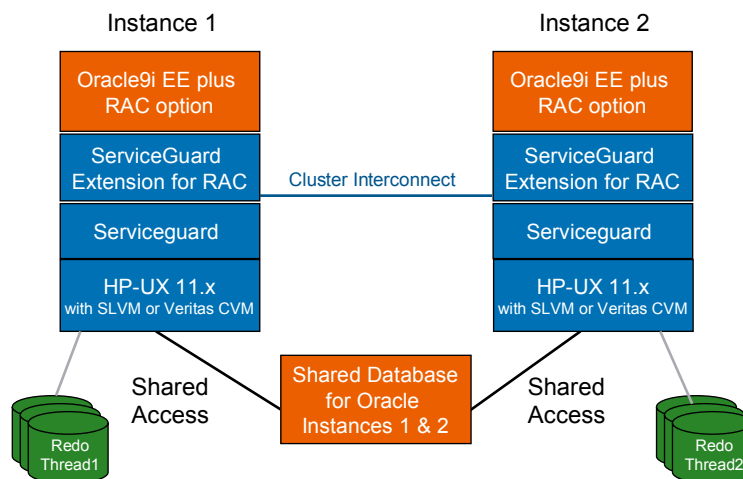
- All systems are active and access the same database
- Increase Scalability beyond one system
- Much faster failover and switchover
- Can pre-connect users to the secondary instance for fast restart without logon
- Load can be distributed
- Caches are warm on the secondary instance – better performance after failure
- Requires raw devices or CFS

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.

13

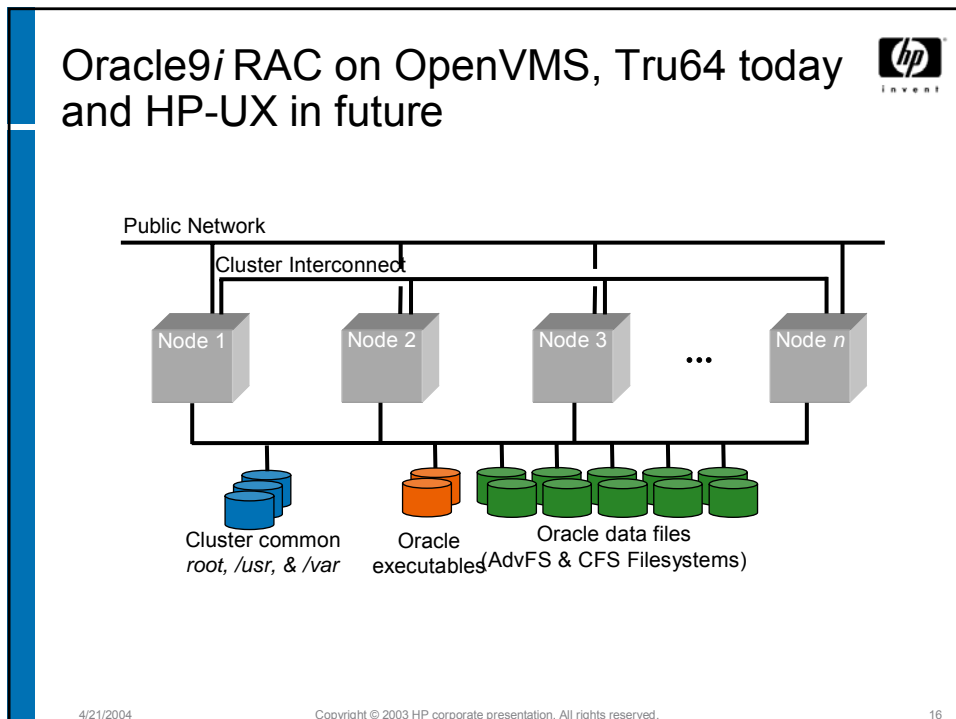
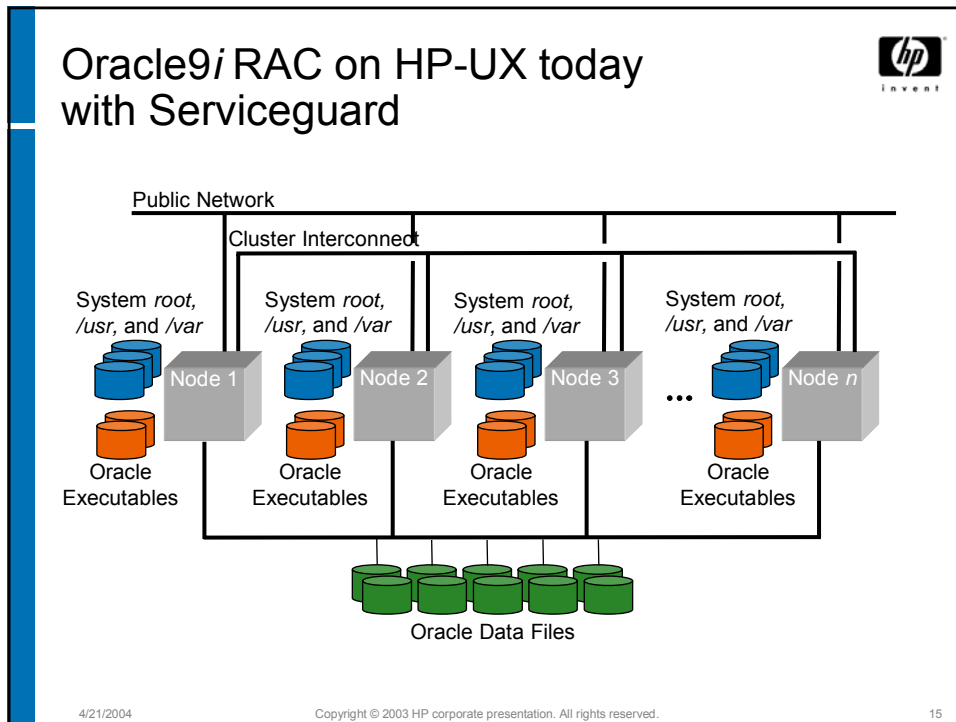
## Oracle9i RAC Software Stack on HP-UX



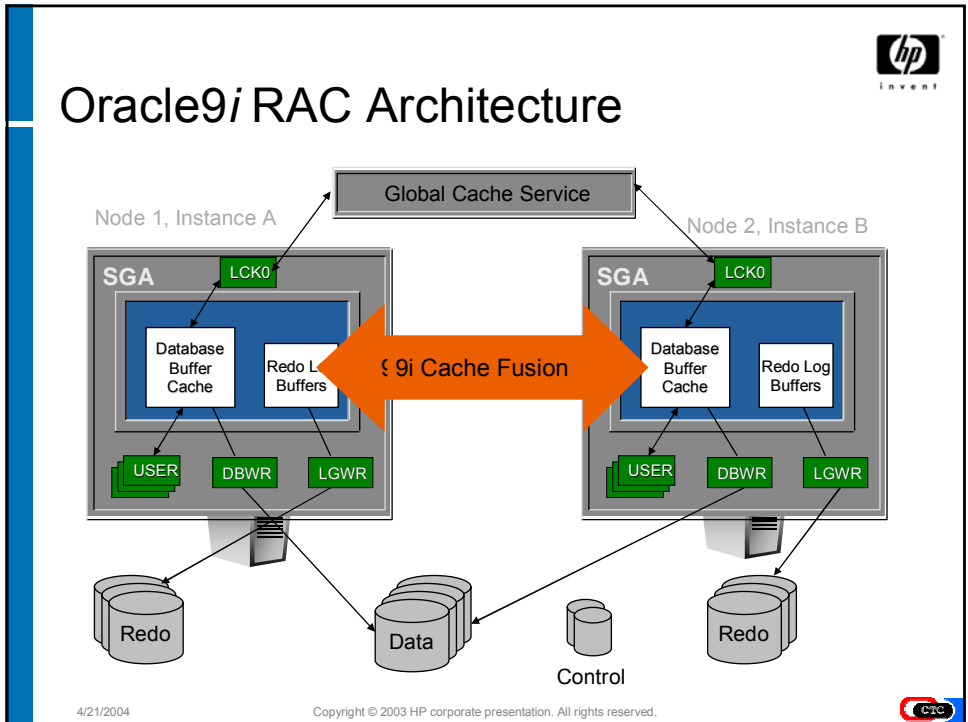
4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.









**Oracle9i RAC on HP**

- ✓ Fully engineered solutions
  - Benchmarking
    - Oracle Applications on HP-UX with scalability of 1.9
    - TPC-C Benchmark on Linux and Windows
    - SAP Benchmark on HP Tru64 with scalability of > 1.85)
  - Oracle versions certified on HP platforms
  - Pre-configured & pre-tested configurations
  - Best practices Blueprints
- ✓ Full Range of offerings
  - From ProLiant up through SuperDome
  - HP-UX, Tru64 Unix, Windows & Linux, OpenVMS
- ✓ References

PURCHASEPRO  
Where Business Takes Off.

Consejo General del Poder Judicial

ACXION

FreeMarkets

KITC 한국투자신탁증권

TOP TOY

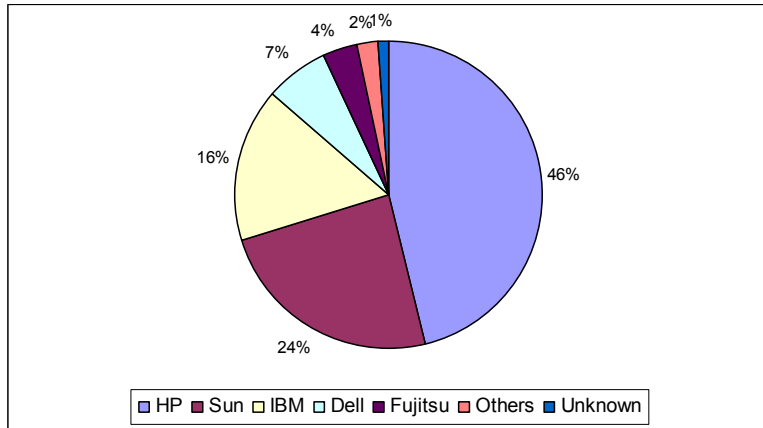
nor|dac

NAVISION  
The Way to Grow

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 18



## 9i RAC installations per platform



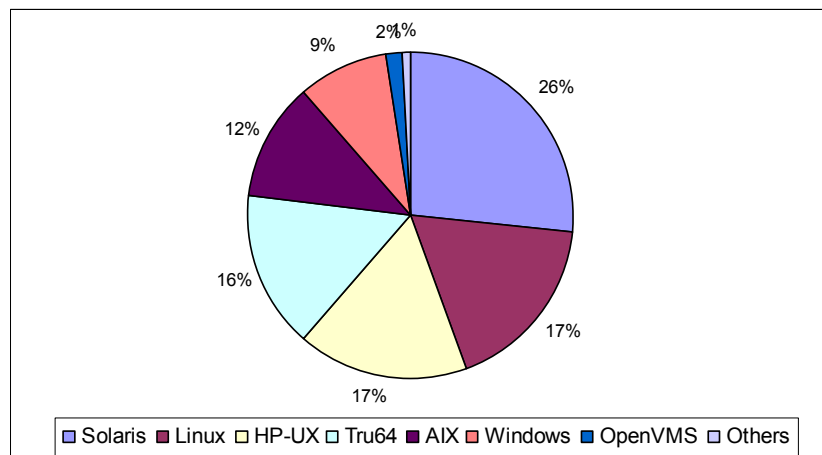
4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.

19



## 9i RAC Installations per OS



4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.


20

ORACLE DATABASE 10<sup>g</sup>  invent

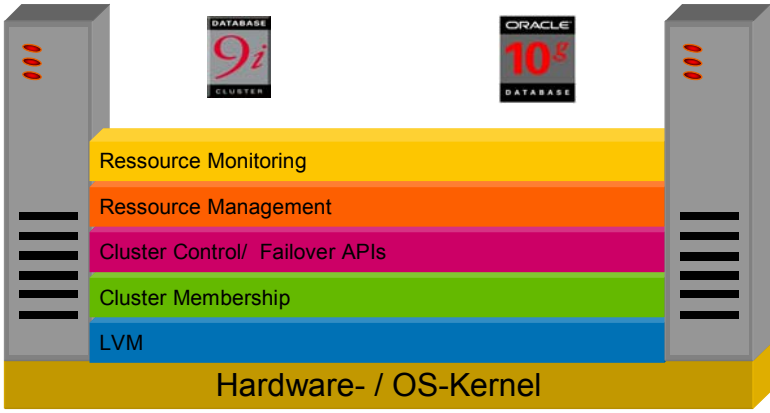
## Agenda

- HP/Oracle partnership
- Oracle Grid Computing
- Oracle Real Application Clusters Overview
- Oracle Real Application Clusters 10g Architecture
  - Automatic Storage Management
  - Oracle Clusterware
  - Automatic Workload Management
  - Rolling Upgrades
- HP Parallel Database Clusters on Proliant
- HP/Oracle Joint Support Services

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

ORACLE DATABASE 10<sup>g</sup>  invent


## Oracle10g RAC Architecture

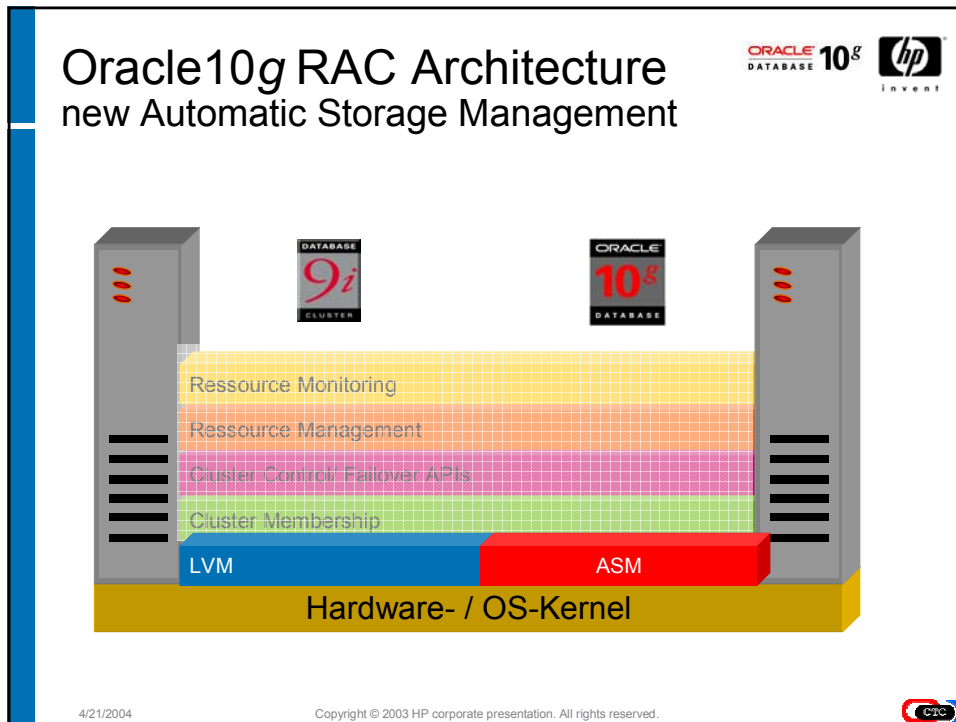


The diagram illustrates the Oracle 10g RAC architecture. It shows two server racks connected by a network. The architecture is layered as follows:

- Hardware- / OS-Kernel**: The base layer, represented by a yellow block.
- LVM**: Logical Volume Manager, represented by a blue block.
- Cluster Membership**: Represented by a green block.
- Cluster Control/ Failover APIs**: Represented by a pink block.
- Ressource Management**: Represented by an orange block.
- Ressource Monitoring**: Represented by a yellow block.

At the top, there are two server racks. The left rack is labeled 'DATABASE 9i CLUSTER' and the right rack is labeled 'ORACLE 10<sup>g</sup> DATABASE'. The racks are connected by a network, indicated by lines between them.

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

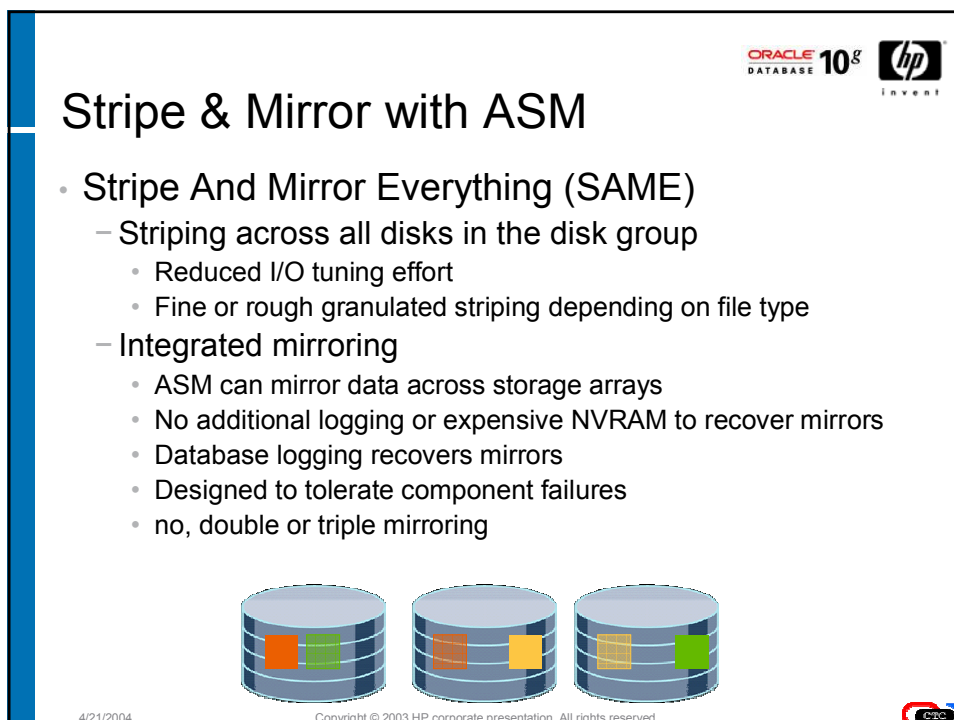
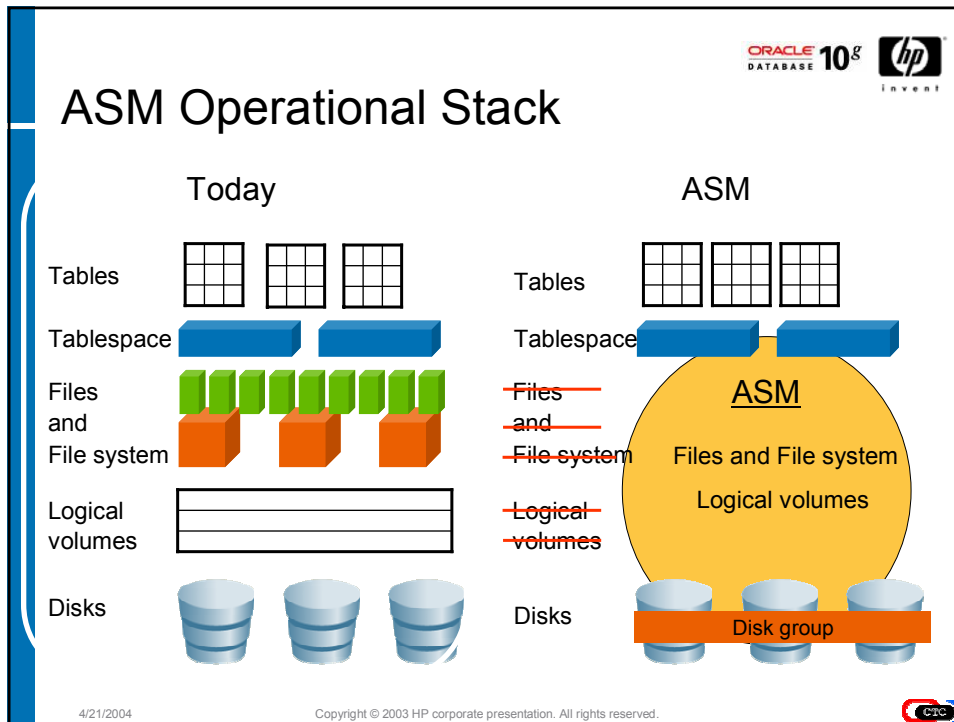



## What is ASM?

ASM = Automatic Storage Management

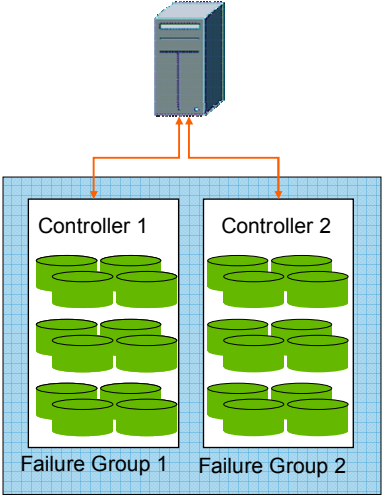
- A feature of Oracle Database 10g; no additional charge
- Portable, high performance cluster file system
- Data is spread across disks to balance load
- Integrated mirroring across disks
- No volumes: just a pool of storage

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved.




ORACLE 10<sup>g</sup> DATABASE 


## Disk Groups & Failure Groups



- Storage resources organized into **disk groups** and **failure groups**
- a disk group is a collection of disks managed as a logical unit in which ASM files are stored
- disk groups assigned by DBAs
- a failure group is a set of disks sharing a common resource whose failure needs to be tolerated
  - redundant copies of an extent are stored in separate failure groups
- failure groups manually assigned by DBAs or automatically by ASM


Disk Group

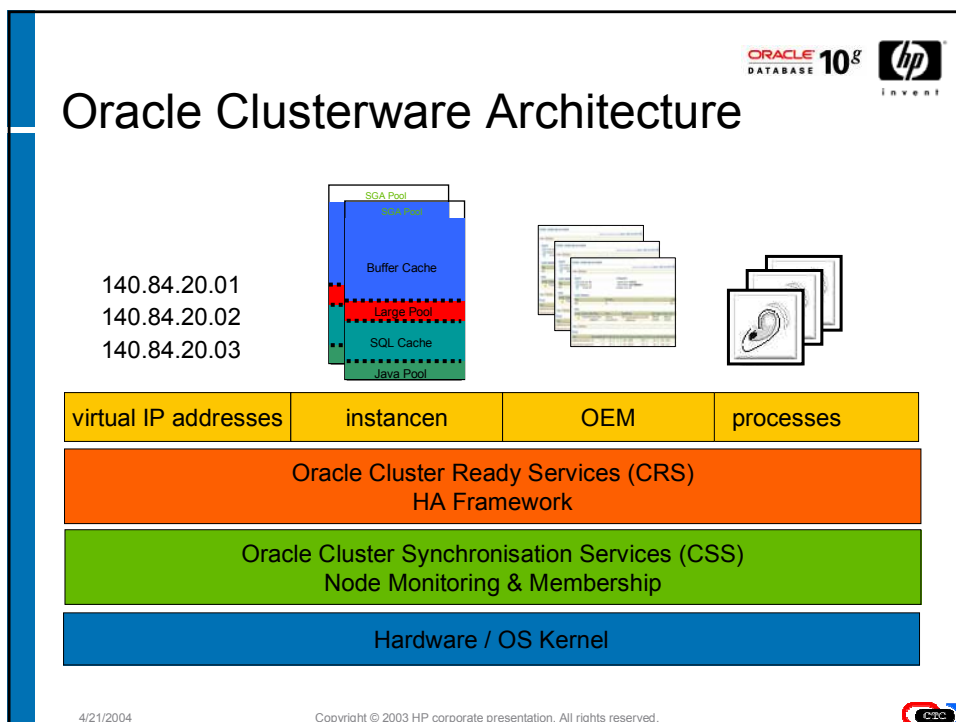
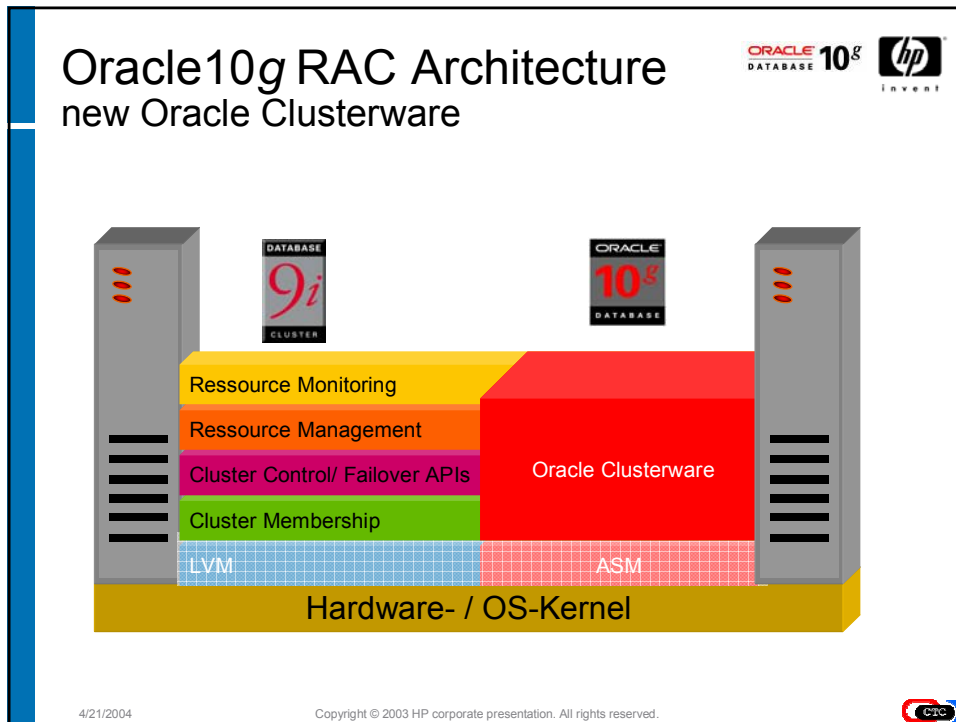
4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

ORACLE 10<sup>g</sup> DATABASE 

## ASM integrated with HP Storage

- ASM-Lib is an Open Interface extending ASM
- Vendor library for disk discovery and I/O
- Efficient I/O interface
- Data description allows end-to-end validation (HARD Initiative)
- HP is participating
- 1st HP ASM-Lib implementation available with XP and HP-UX





## How does Oracle Clusterware relate to HP Serviceguard?



- CRS provides many of the monitoring and failover capabilities seen in Serviceguard.
- However, **CRS is intended for use primarily with Oracle database services and resources.**
- HP Serviceguard is still recommended for high availability of 3rd party applications.
- Another key difference is that CRS does everything at a user level and, therefore, has some disadvantages when compared to Serviceguard.
- Serviceguard is required for any customer that is using (S)LVM or CVM for their shared storage.

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.



## Installation Options with RAC 10g on HP-UX



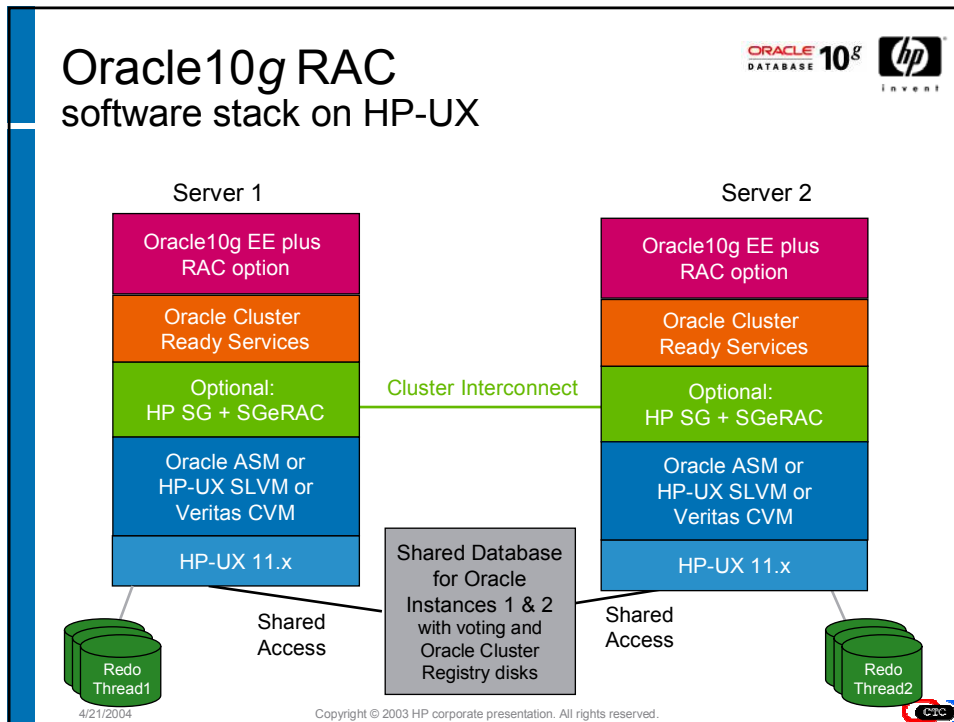
	HP Serviceguard & HP Serviceguard Extension for RAC	Oracle Clusterware
HP Shared Logical Volume Management	OK	NO
Oracle ASM	OK	OK

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.

32





## Automatic Workload Management

- Application workloads can be defined as Services
  - Individually managed and controlled
  - Services specified in TNS connect data
  - Assigned to instances during normal startup
  - On instance failure, automatic re-assignment
  - Service performance individually tracked
  - Rules specify automatic resource allocations
    - PREFERRED instances during normal conditions
    - AVAILABLE instances should failures occur

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved.

# Workload Management and New Service Concept

Customer Applications

`service_name = RAC`

Database Server

Shared Storage

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.

# Oracle 10g RAC Services

Connect to Services in using a service\_name

Customer Applications

`(CONNECT_DATA =`  
`(SERVICE_NAME = fi))`

Provide Services

`(SERVICE_NAME = fi)`

FI  
CR  
M

Database Server

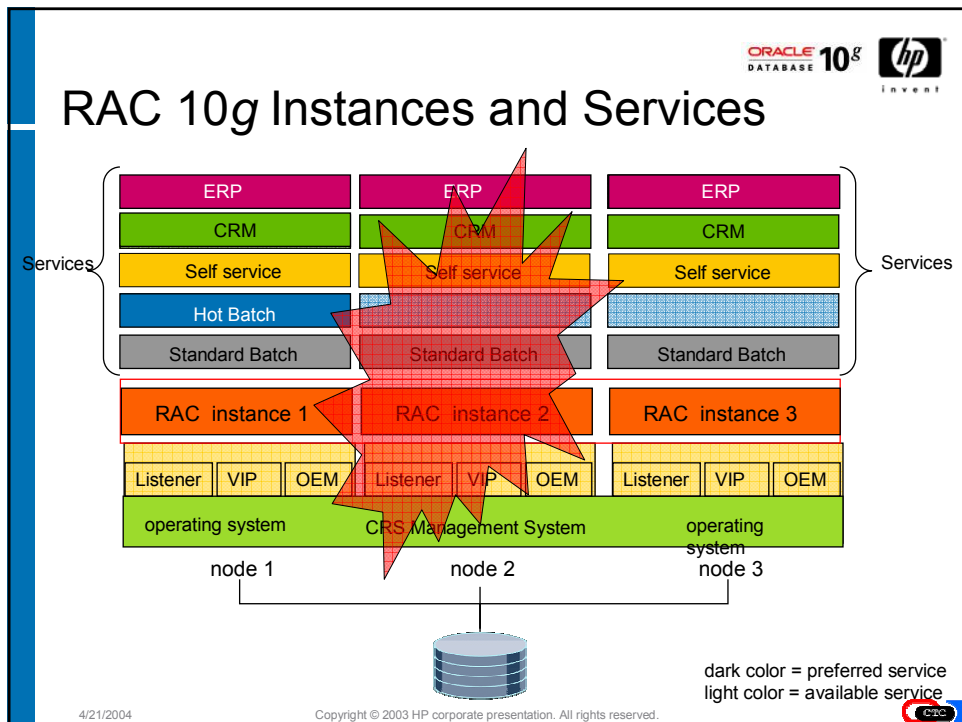
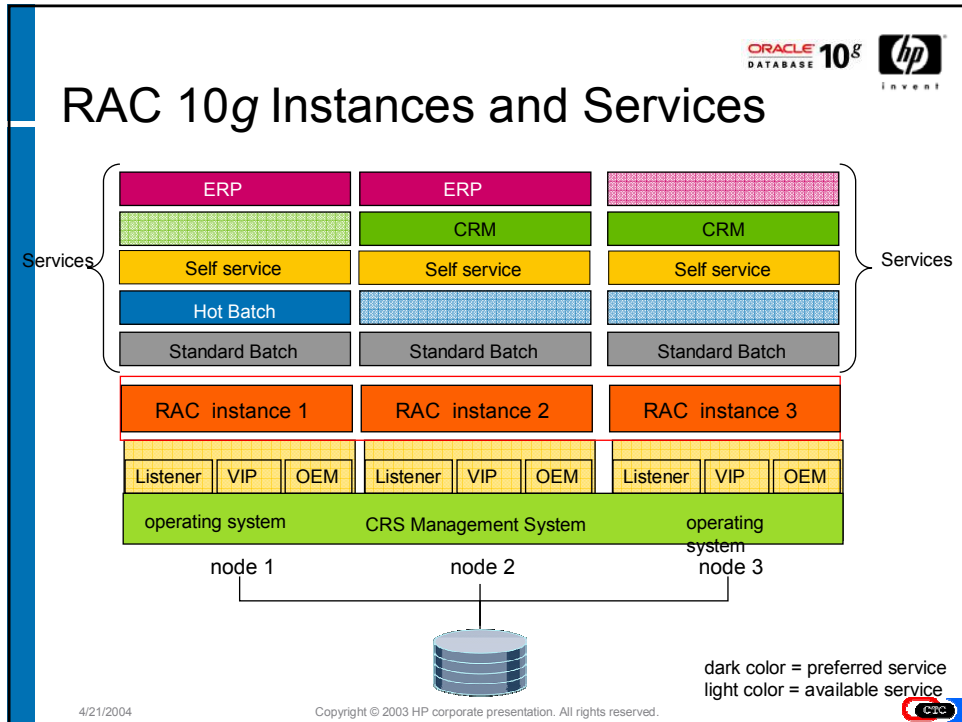
Shared Storage

CRM preferred service


CRM available service

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.



## Oracle RAC 10g Service Definition via DBCA




- DBCA is the recommended tool for configuring services

Expand the tree and click the Oracle database services from the left panel to see the service details in the right panel. Review the configuration details for each service.


Instance	Not Used	Preferred	Available
GRID1	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
GRID2	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
GRID3	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
GRID4	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>

TAF Policy:  None  Basic  Pre-connect

Buttons: Add, Remove, Cancel, Help, Back, Next, Finish

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

## Oracle RAC 10g Service Definition via SRVCTL utility






- svrctl utility can be used as an alternative:

```

$ svrctl add service -d <name> -s <service_name>
  [-r <preferred_list>] [-a <avail_list>]

$ svrctl add service -d ORACLE -s ERP
  -r GRID1,GRID2 -a GRID3,GRID4
$ svrctl add service -d ORACLE -s CRM
  -r GRID1 -a GRD2
$ svrctl add service -d ORACLE -s SS
  -r GRID1,GRID2,GRID3,GRID4
$ svrctl add service -d ORACLE -s HOT
  -r GRID1 -a GRID2, GRID3,GRID4
$ svrctl add service -d ORACLE -s STD
  -r GRID1,GRID2 -a GRID3,GRID4
    
```

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 






## Using Oracle RAC 10g Services

- Entry in tnsnames.ora:

```


ERP= (DESCRIPTION= (LOAD_BALANCE=on
  (ADDRESS= (PROTOCOL=TCP) (HOST=node1-vip) (PORT=1521))
  (ADDRESS= (PROTOCOL=TCP) (HOST=node2-vip) (PORT=1521))
  (ADDRESS= (PROTOCOL=TCP) (HOST=node3-vip) (PORT=1521))
  (ADDRESS= (PROTOCOL=TCP) (HOST=node4-vip) (PORT=1521))
)
(CONNECT_DATA= (SERVICE_NAME=ERP))
(FAILOVER MODE =
  (TYPE=SELECT)
  (METHOD=BASIC)
  (RETRIES=180)
  (DELAY =5)
)
)
    
```

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 




## Automatic Workload Management Enterprise Manager Controls

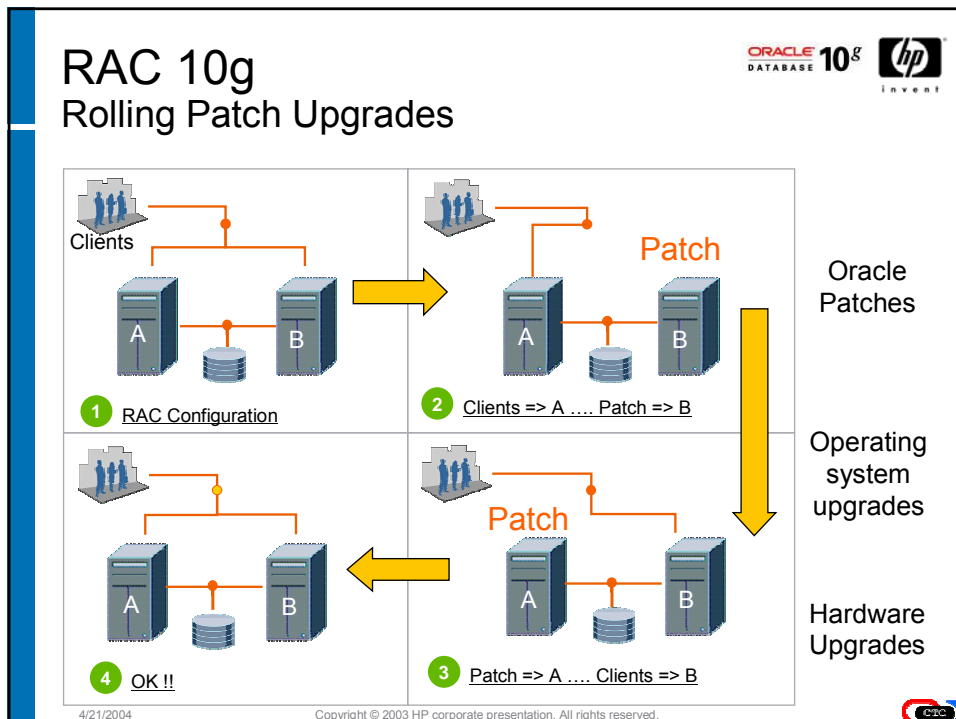
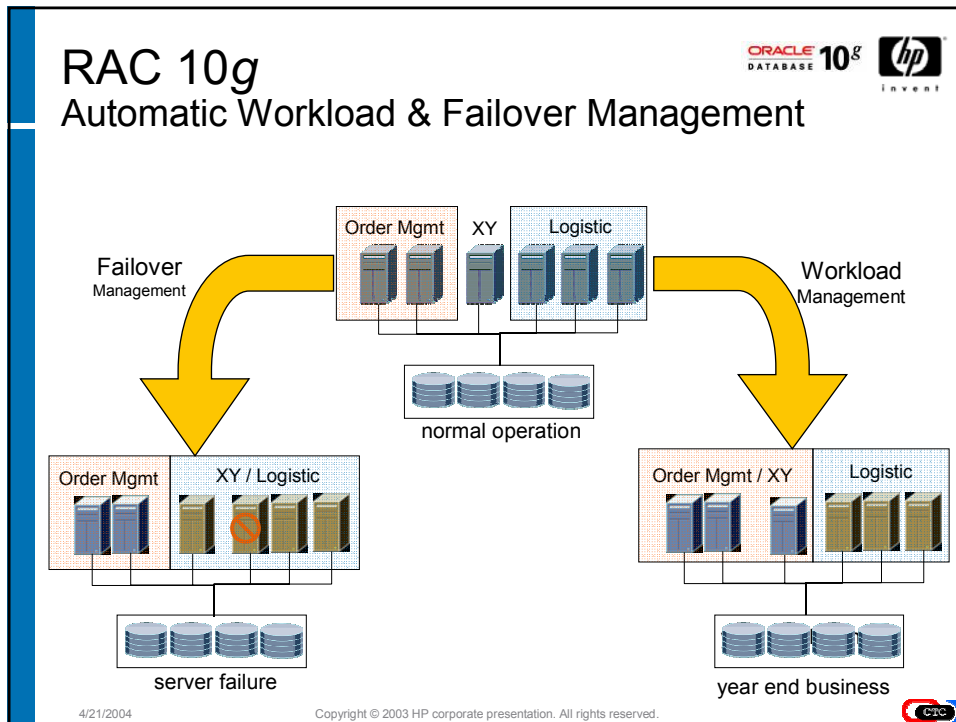
- Performs service operations
  - Start / Stop; Enable / Disable; Relocate
- View service status
  - Including automatic resource allocation rules

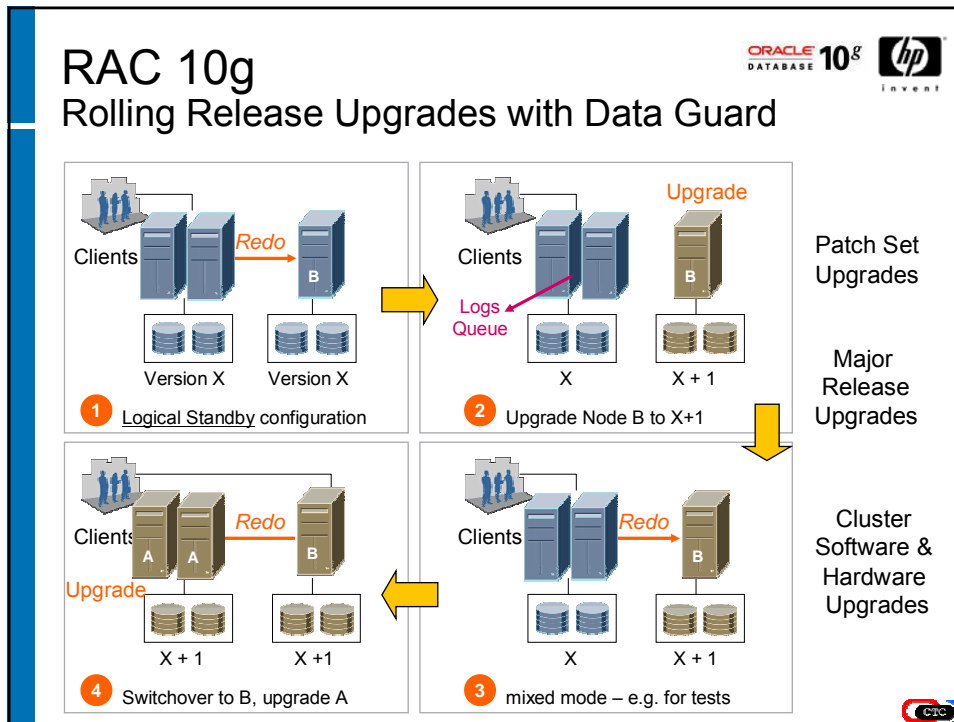


The screenshot shows the Oracle Enterprise Manager interface for Database Services. It displays a table of services with columns for Select Service Name, Running Instances, Preferred Instances, Available Instances, and FAI Policy. The table lists three services: sv2, sv3, and sv4, each with 'ezrac' as the running instance and 'ezrac2' as the preferred instance. The FAI Policy for sv2 and sv3 is 'NONE', and for sv4 it is 'BASIC'. Navigation buttons like 'Enable', 'Disable', 'Start', 'Stop', and 'Relocate' are visible above the table.

Copyright © 1996, 2003, Oracle. All rights reserved.  
About Oracle Enterprise Manager Version 4.1.0.0.0

4/21/04 



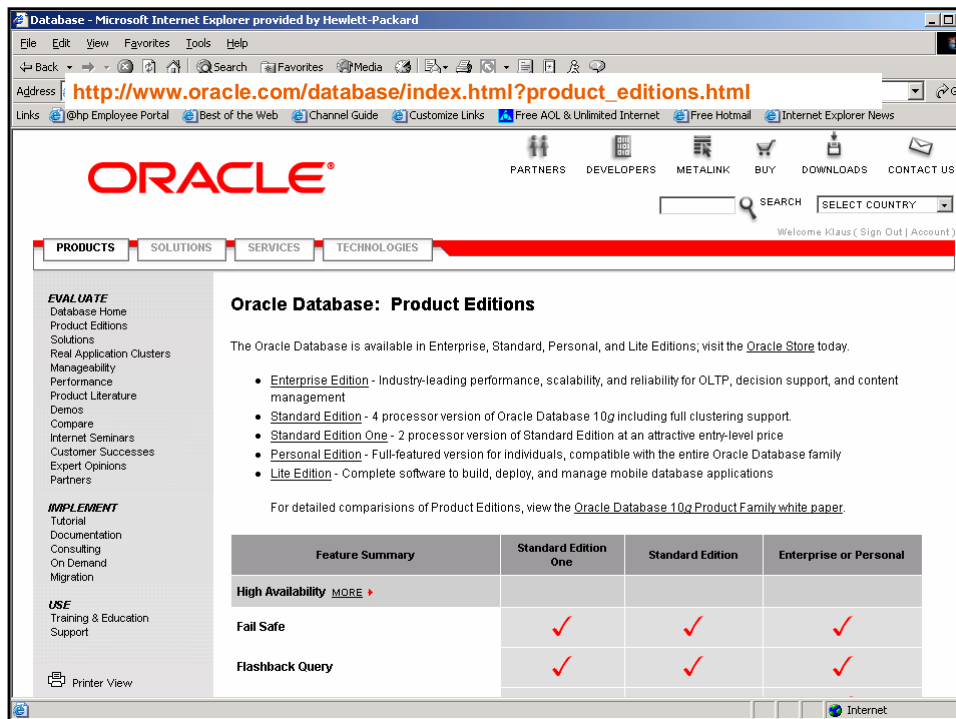
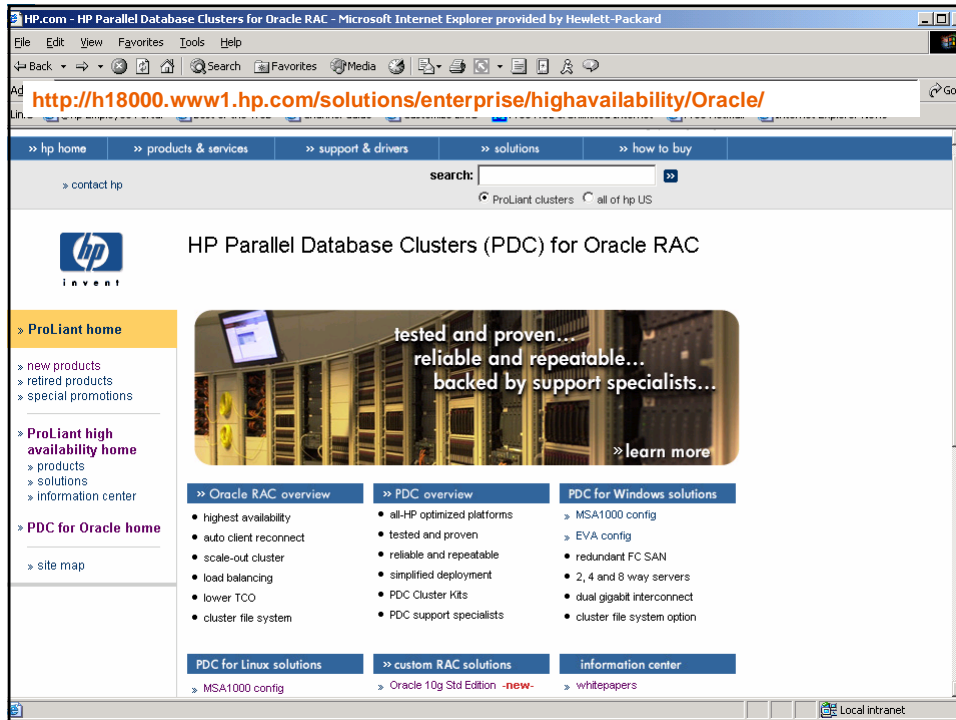


## Agenda

- HP/Oracle partnership
- Oracle Grid Computing
- Oracle Real Application Clusters Overview
- Oracle Real Application Clusters 10g Architecture
  - Automatic Storage Management
  - Oracle Clusterware
  - Automatic Workload Management
  - Rolling Upgrades
- HP Parallel Database Clusters on Proliant
- HP/Oracle Joint Support Services

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.





## Oracle 10g Standard Edition with RAC

### Sample configuration

- HP Proliant DL380G3 Packaged Cluster with MSA500
- Each server with
  - 2 Intel Xeon 2.8GHz CPU 2 GB Memory
  - internal storage: 2x 36 GB 15.000 RPM disks
  - management LAN for full control of the remote servers over the network via a standard browser
  - Operating System: Windows or Linux
- MSA500 with 10 disks



ORACLE 10<sup>g</sup>  
DATABASE

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.



## Agenda



- HP/Oracle partnership
- Oracle Grid Computing
- Oracle Real Application Clusters Overview
- Oracle Real Application Clusters 10g Architecture
  - Automatic Storage Management
  - Oracle Clusterware
  - Automatic Workload Management
  - Rolling Upgrades
- HP Parallel Database Clusters on Proliant
- HP/Oracle Joint Support Services

ORACLE 10<sup>g</sup> 

4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.




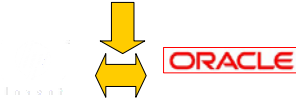



## HP/Oracle Joint Support Services

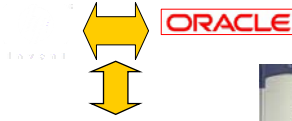
Simpler, faster problem resolution  
keeps your business on-line and your staff on-schedule

**Reactive**


Call either HP or Oracle for any interoperability problem


**Proactive**





HP and Oracle provide coordinated assessments and optimization






Together, we solve your problem.

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 






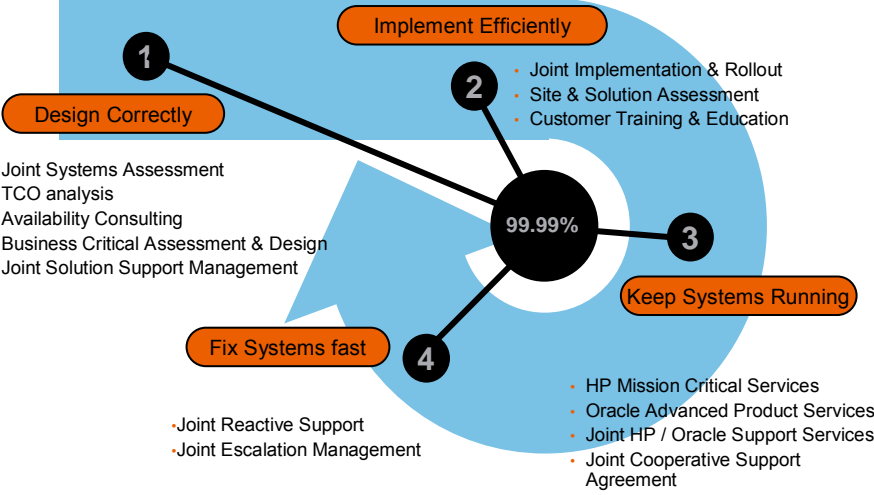
## HP/Oracle Joint Support Services

	Joint Support Services	
<ul style="list-style-type: none"> <li>• Service Centers</li> <li>• Monitoring Centers</li> <li>• Local APS Experts</li> <li>• Global Product Support organization</li> <li>• Support tools and connectivity</li> </ul>	<ul style="list-style-type: none"> <li>• Service &amp; Process simplification                             <ul style="list-style-type: none"> <li>– Designated-Point-of-Contact</li> <li>– Increased accountability</li> <li>– Synchronized problem resolution and escalation process</li> </ul> </li> <li>• Increased knowledge sharing and in-depth communication between main IT Suppliers</li> <li>• Knowledge and support of the complete solution and not only individual components                             <ul style="list-style-type: none"> <li>– End-to-end service-level</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Mission Critical Team</li> <li>• Monitoring Centers</li> <li>• Proactive onsite support</li> <li>• Reactive Support Team</li> <li>• Remote Support Technology</li> </ul>

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 

## HP/Oracle Joint Support Services Customer Approach



**1 Design Correctly**

- Joint Systems Assessment
- TCO analysis
- Availability Consulting
- Business Critical Assessment & Design
- Joint Solution Support Management

**2 Implement Efficiently**

- Joint Implementation & Rollout
- Site & Solution Assessment
- Customer Training & Education

**3 Keep Systems Running**


- HP Mission Critical Services
- Oracle Advanced Product Services
- Joint HP / Oracle Support Services
- Joint Cooperative Support Agreement

**4 Fix Systems fast**



- Joint Reactive Support
- Joint Escalation Management

4/21/2004


Copyright © 2003 HP corporate presentation. All rights reserved.

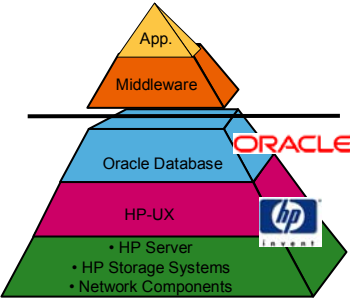


## Joint HP/Oracle support solution


- Main elements of the joint support:
  - HP and Oracle provide joint support to DaimlerChrysler for the operation of the QUIS-solution:
    - Joint Change-Management
    - Joint Assessments
    - Joint Planning and Consulting
    - HA Commitment on DB-Level
    - coordinated reactive processes
    - Single-Point-of-Contact
    - HP and Oracle work as one team
  - DaimlerChrysler has CSS-contract at HP and DAS contract at Oracle, integration via ISS-Modul






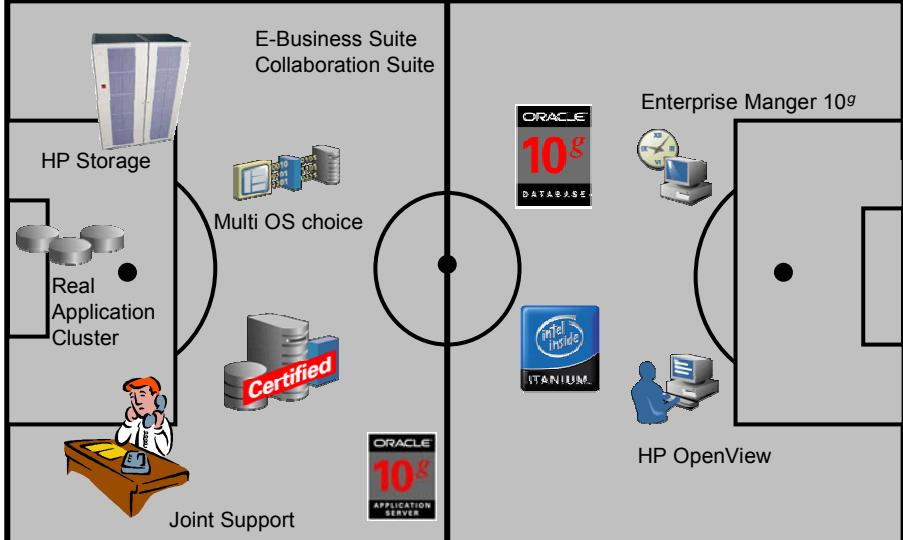
4/21/2004

Copyright © 2003 HP corporate presentation. All rights reserved.



**ORACLE** 10<sup>g</sup>   
DATABASE invent

## HP/Oracle – the perfect combination




The diagram is divided into three main sections. The left section features 'HP Storage' (server rack), 'Real Application Cluster' (server icons), and 'Joint Support' (a person at a desk). The center section lists 'E-Business Suite', 'Collaboration Suite', 'Multi OS choice', and 'Certified' (server icons). The right section includes 'ORACLE 10<sup>g</sup> DATABASE', 'Enterprise Manager 10<sup>g</sup>', 'Intel Inside TITANIUM', and 'HP OpenView' (person at computer). A central circle with a dot connects the three sections.

4/21/2004 Copyright © 2003 HP corporate presentation. All rights reserved. 55

**ORACLE**

# Q & A

  
invent