



Building Oracle 10g RAC in a Single PC

Shaojing Guo
The Washington Post
Dec. 9, 2005

Topics

- Oracle RAC architecture
- Virtual machine
- Setup (Hardware/Software)
- Oracle 10g RAC installation
- Cluster admin utilities
- “Demo”

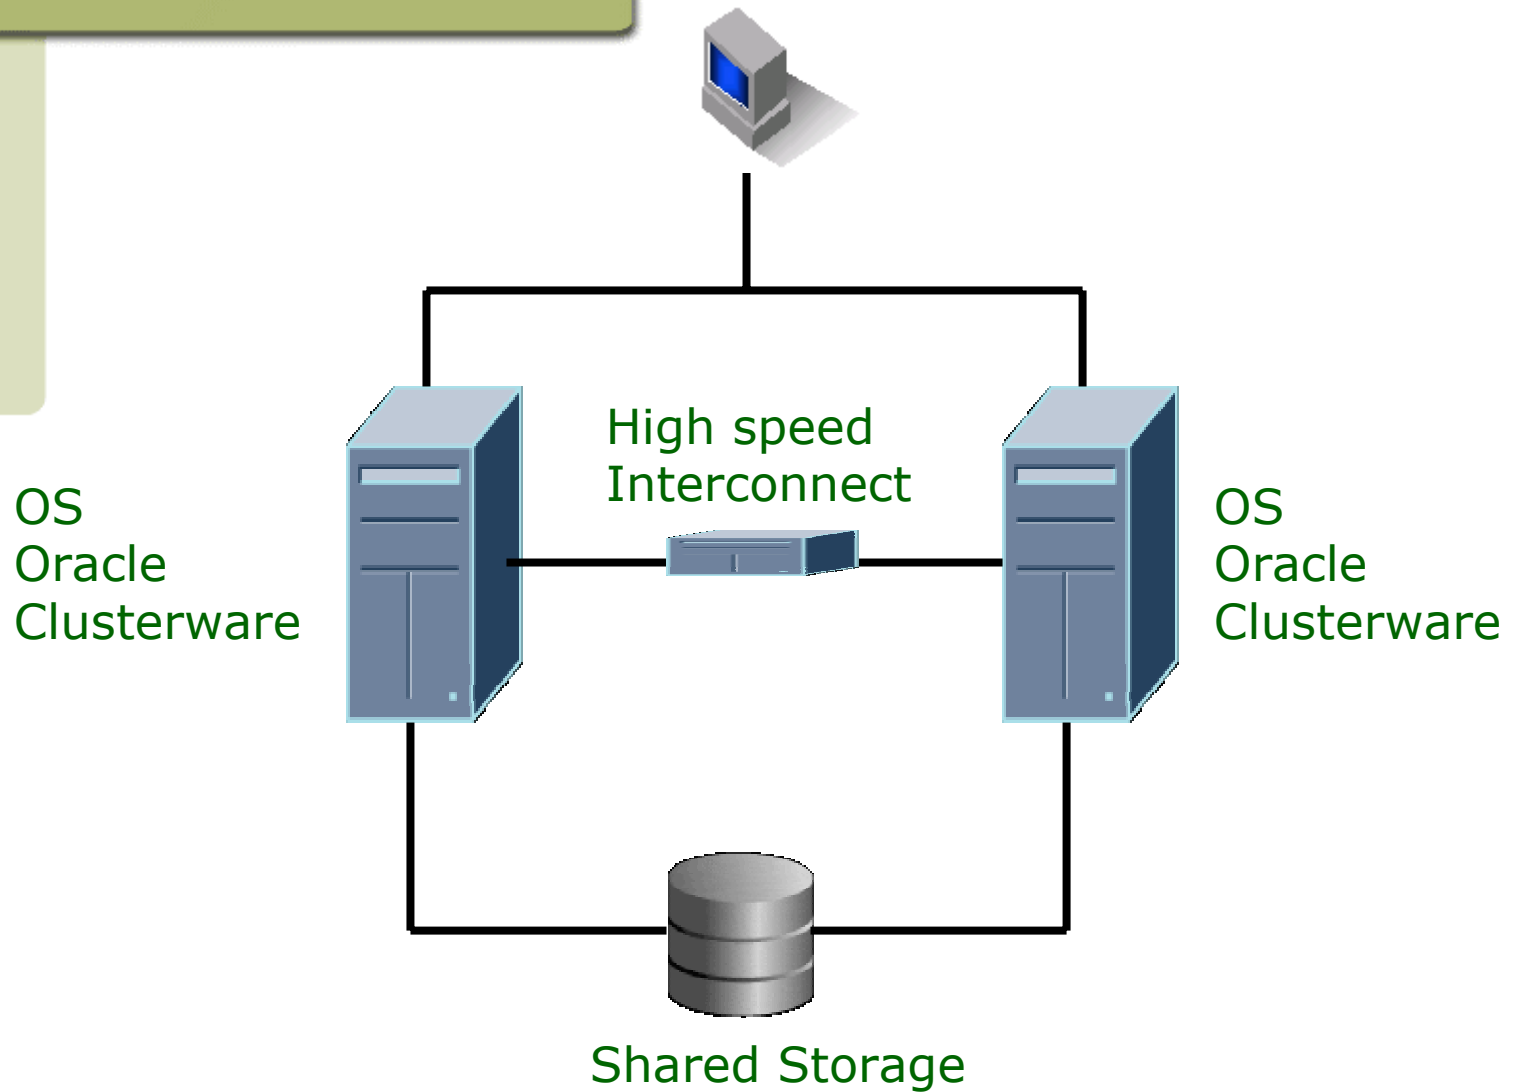
Introduction

- First released in Oracle 6
- Oracle 6, 7, 8, 8i – Oracle Parallel Server (OPS)
- Oracle 9i, 10g – Oracle Real Application Cluster (RAC)
- 16% of new Oracle customers are RAC customers *
- 20% of RAC deals are from repeat customers *

RAC Benefits

- High availability
- Scalability
- Low Cost (Comparing with big SMP box)
- SMP box is not big enough

RAC Architecture



Shared Storage

- RAW device
- Oracle Automatic Storage Management - ASM
- Oracle Cluster File System - OCFS (Windows, Linux)
- Cluster File System provided by other vendors

ORACLE Clusterware

- Oracle CRS - Cluster Ready Service
- Provides Cluster Management Services
- Incorporates clustering technology from DEC's True64 True Cluster
- Started with 10g, Oracle provides CRS w/o additional cost

RAC is great, but...

- It's not cheap to set up
 - At least two computers
 - Interconnect
 - Shared storage
- Solution - Virtual Machine!
 - Only one computer is needed
 - Multiple Virtual Linux machines can share same storage, easily!
 - Downloadable software

What is Virtual Machine?

A program running on a computer that creates a self-contained operating environment and presents the appearance to the user of a different computer.

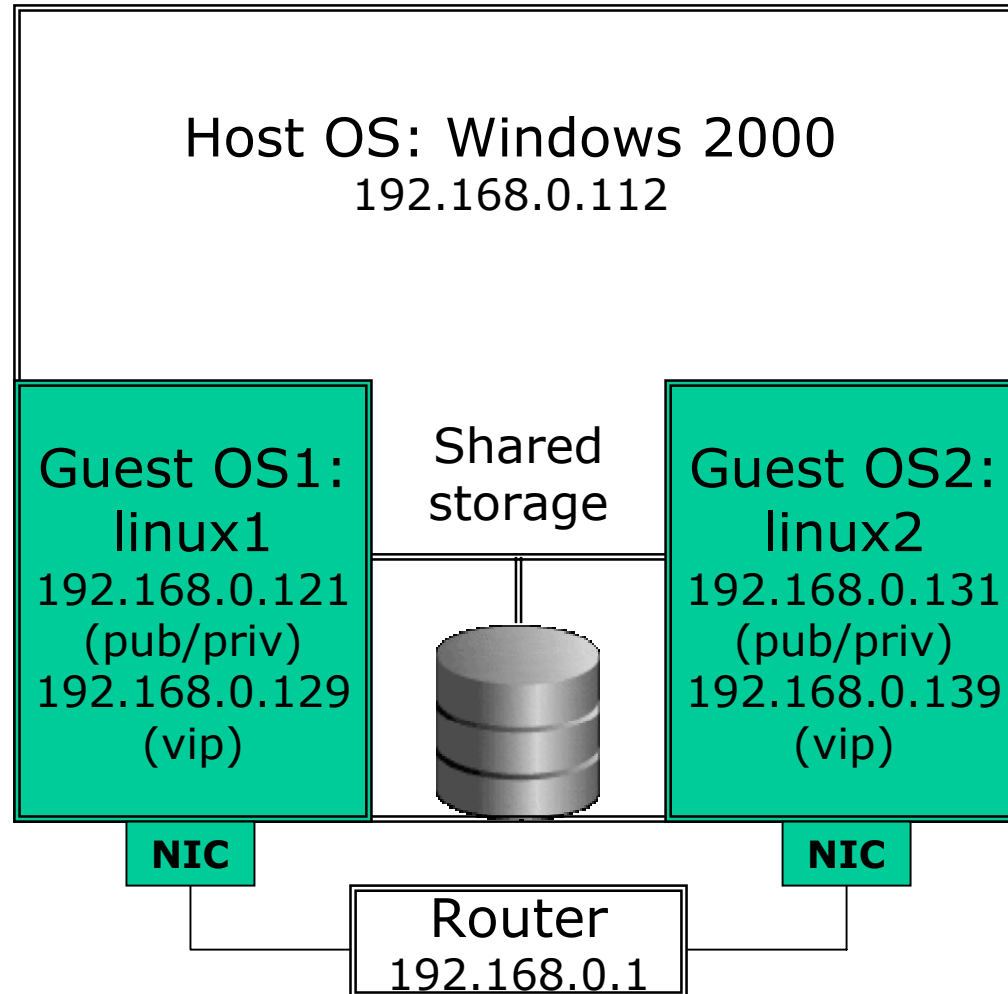
Benefits:

Partitioning

Isolation

Encapsulation

Virtual Machine



RAC Setup (Hardware)

- One PC
- AMD Athlon XP processor (1.3 GHz)
- 1.0 GB RAM
- 40 GB internal HD (Disk1)
- 6 GB internal HD (Disk2) - shared storage
- CDROM and Floppy

RAC Setup (Software)

- Windows 2000 Operation System (host system)
- VMWare Workstation 4
- SUSE Enterprise Linux Server (SELS) 8.0 SP3
- Oracle 10g (10.1.0.3) Cluster Ready Service for x86 Linux
- Oracle 10g (10.1.0.3) Database for x86 Linux

RAC Installation (1.1)

- Disk partition for OS and RAC database

Partition (Disk1)		Size	Contents
Disk1: Primary Partition (hda1)		20GB	Windows 2000 OS (host)
Disk1: Primary Partition (hda2)		9GB	linux1 (Virtual)
Disk1: Primary Partition (hda3)		9GB	linux2 (Virtual)
Disk1: Extended Partition	(hda4)	1GB	swap for linux1
	(hda5)	1GB	swap for linux2

Partition (Disk2)		Size	Contents
Disk2: Primary Partition (hdb1)		200MB	Quorum disk for Oracle Cluster Registry
Disk2: Primary Partition (hdb2)		100MB	Voting disk for Cluster Synchronization Services
Disk2: Primary Partition (hdb3)		~1GB	unused

RAC Installation (1.2)

- Disk partition (Continued)

Partition (Disk2)		Size	Contents
Disk2: Extended Partition	(hdb4)	700MB	unused
	(hdb5)	800MB	sysaux tablespace
	(hdb6)	160MB	example tablespace
	(hdb7)	120MB	users tablespace
	(hdb8)	500MB	system tablespace
	(hdb9)	500MB	undo tablespace for instance 1
	(hdb10)	500MB	undo tablespace for instance 2
	(hdb11)	120MB	redo member 1 for instance 1
	(hdb12)	120MB	redo member 2 for instance 1
	(hdb13)	120MB	redo member 1 for instance 2
	(hdb14)	120MB	redo member 2 for instance 2
	(hdb15)	110MB	controlfile for instance 1
	(hdb16)	110MB	controlfile for instance 2
	(hdb17)	8M	spfile
(hdb18)	8M	pwdfile	
(hdb19)	250MB	temp tablespace	
(hdb20)	700MB	unused	

RAC Installation (2)

- Install VMWare 4 on Windows
 - <http://www.vmware.com>
- Create two virtual Linux machines
 - <http://www.novel.com/linux/suse>

	linux1	linux2
Memory	360 MB	360 MB
Location (/)	hda2	hda3
Swap	hda4	hda5
NIC	Bridged	Bridged
CDROM	Yes	Yes

Favorites

- Linux1
- Linux2

```
Linux1 Linux2
oracle@priv1:/home/oracle> uname -a
Linux linux1 2.4.21-138-athlon #1 Fri Oct 31 00:47:49 UTC 2003 i686 unknown
oracle@priv1:/home/oracle> ssh oracle@linux2 uname -a
Linux linux2 2.4.21-138-athlon #1 Fri Oct 31 00:47:49 UTC 2003 i686 unknown
oracle@priv1:/home/oracle> srvctl status database -d RAC1
Instance RAC11 is running on node linux1
Instance RAC12 is running on node linux2
oracle@priv1:/home/oracle> srvctl status asm -n linux1
ASM instance +ASM1 is running on node linux1.
oracle@priv1:/home/oracle> srvctl status asm -n linux2
ASM instance +ASM2 is running on node linux2.
oracle@priv1:/home/oracle> _
```

Linux2 - [Ctrl-Alt-F2] - VMware Workstation

File Edit Power Snapshot View Windows Help

Snapshot Revert

Linux1 Linux2

Shell - Konsole

Session Edit View Settings Help

```

$ hostname
linux2
$ ssh oracle@linux1 hostname
linux1
$ crs_stat2
HA Resource          Target          State
-----
ora.RAC1.DB1.RAC11.srv OFFLINE        OFFLINE
ora.RAC1.DB1.RAC12.srv OFFLINE        OFFLINE
ora.RAC1.DB1.cs       OFFLINE        OFFLINE
ora.RAC1.RAC11.inst   ONLINE         ONLINE on linux1
ora.RAC1.RAC12.inst   ONLINE         ONLINE on linux2
ora.RAC1.db           ONLINE         ONLINE on linux1
ora.linux1.ASM1.asm   ONLINE         ONLINE on linux1
ora.linux1.LISTENER_LINUX1.lsnr ONLINE         ONLINE on linux1
ora.linux1.gsd        ONLINE         ONLINE on linux1
ora.linux1.ons        ONLINE         ONLINE on linux1
ora.linux1.vip        ONLINE         ONLINE on linux1
ora.linux2.ASM2.asm   ONLINE         ONLINE on linux2
ora.linux2.LISTENER_LINUX2.lsnr ONLINE         ONLINE on linux2
ora.linux2.gsd        ONLINE         ONLINE on linux2
ora.linux2.ons        ONLINE         ONLINE on linux2
ora.linux2.vip        ONLINE         ONLINE on linux2
$

```

New Shell

2 3 4 Shell - Konsole

2005-11-29

 You do not have VMware Tools installed.

RAC Installation (3)

- Apply SELS 8 SP3
 - Upgrade Linux kernel
 - Configure kernel parameter for Oracle
- Configure network settings
 - /etc/hosts:

```
192.168.0.121      priv1  linux1
192.168.0.129      vip1
192.168.0.131      priv2  linux2
192.168.0.139      vip2
```

RAC Installation (4)

- Create Oracle account on both nodes
- Create RAW bindings on both nodes
 - Edit /etc/raw file: raw1:hdb1
 - chown oracle:dba /dev/raw/raw1
 - /etc/init.d/raw start|status
- Configure RAC notes for remote access on both nodes
 - ssh-keygen to create RSA keys
 - On linux1:
 - # su - oracle
 - \$ ssh oracle@linux2 hostname
 - linux2

RAC Installation (5)

Parameters	Node 1	Node 2
RAC node name (public/private)	linux1/priv1	linux2/priv2
RAC node IP	192.168.0.121	192.168.0.131
RAC VIP	192.168.0.129	192.168.0.139
ORACLE_BASE	/opt/oracle/product	/opt/oracle/product
CRS_HOME	/opt/oracle/product/9ir2	/opt/oracle/product/9ir2
ORACLE_HOME	/opt/oracle/product/10g	/opt/oracle/product/10g
DB_NAME	RAC1	RAC1
INSTANCE_NAME	RAC11	RAC12

RAC Installation (6)

- Install Oracle 10g Cluster Ready Service

Cluster Configuration

Specify the cluster name. For each node in the cluster, specify the public name (the host name), and the private name, to be used to interconnect the nodes within the cluster. The private name cannot be the same as the public name, but can be an IP address.

Cluster Name :

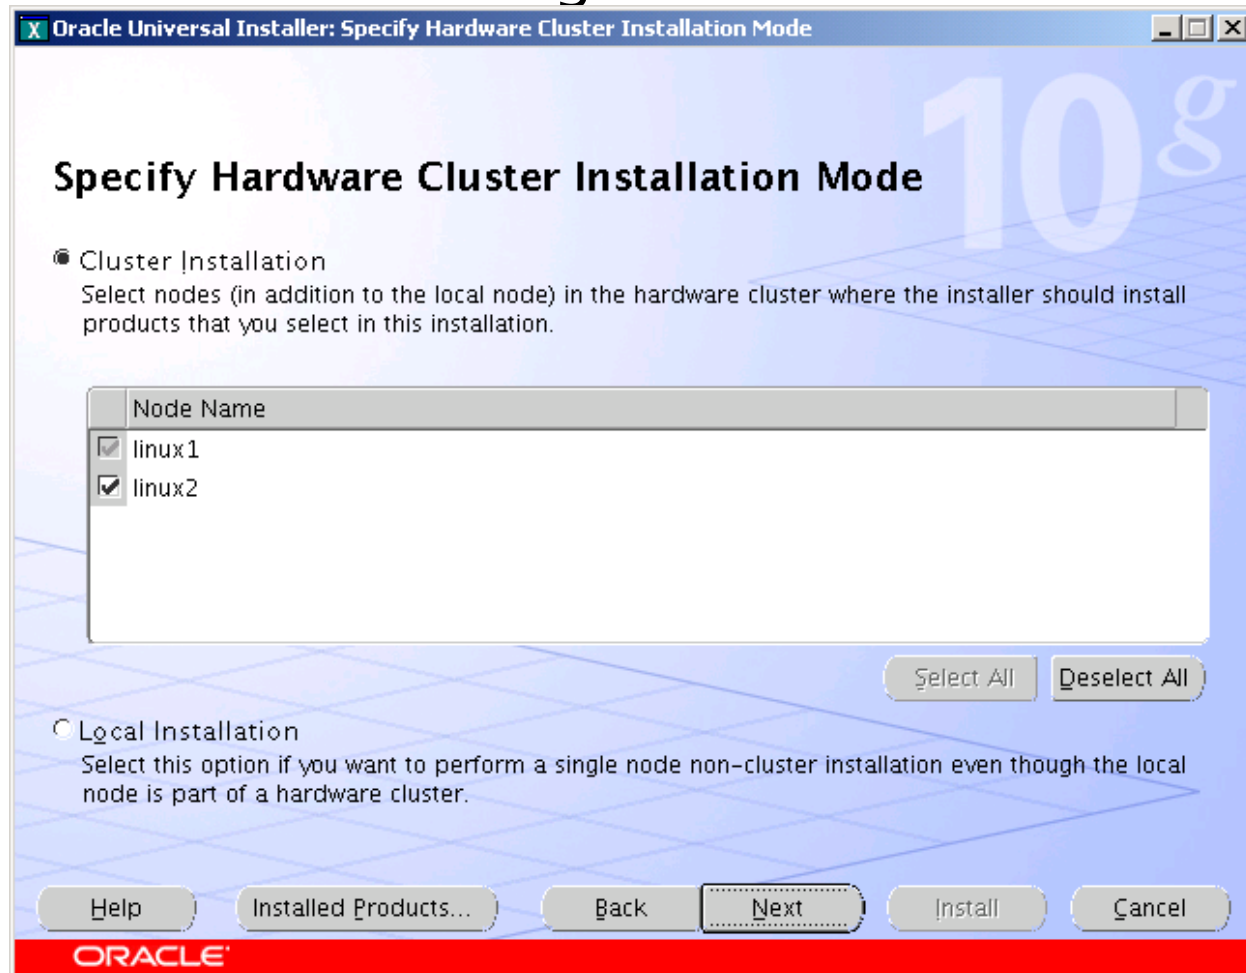
Public Node Name	Private Node Name
linux1	priv1
linux2	priv2

Help Installed Products... Back Next Install Cancel

ORACLE

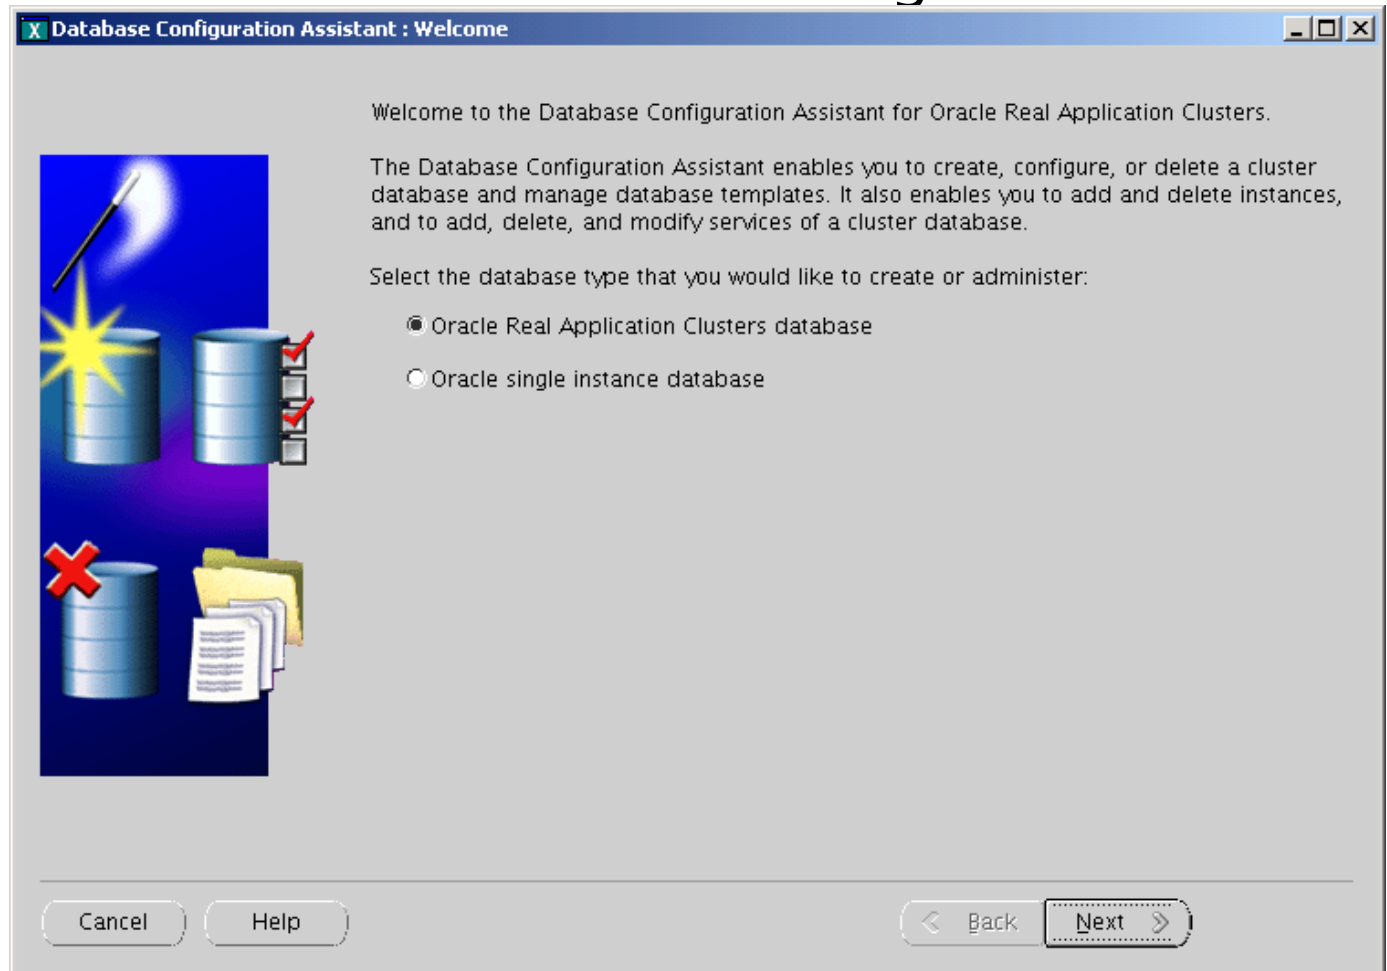
RAC Installation (7)

- Install Oracle 10g with RAC



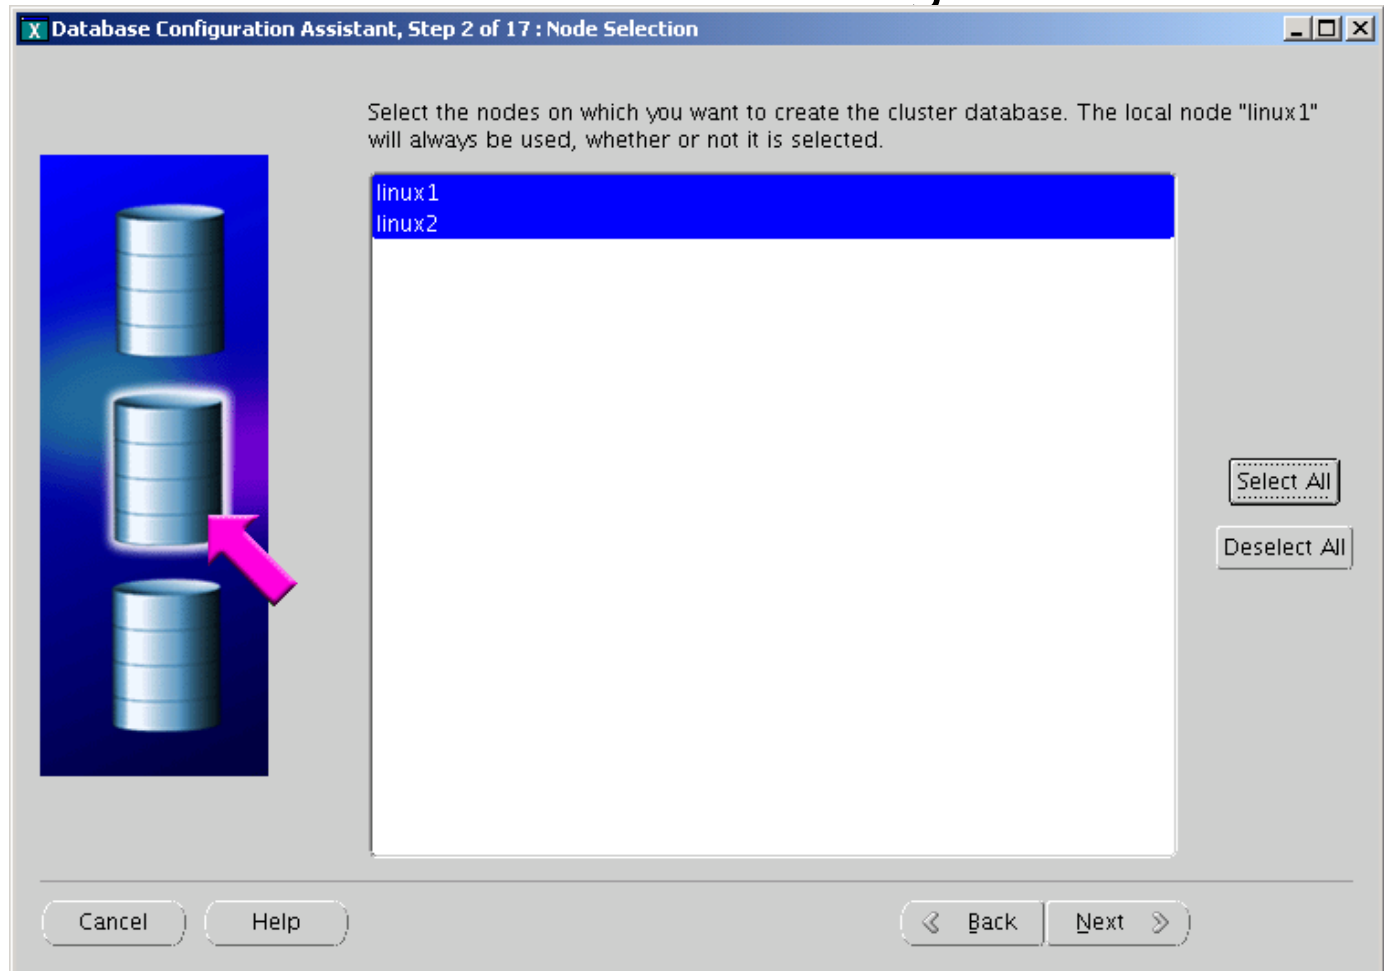
RAC Installation (8.1)

- Create RAC database using DBCA



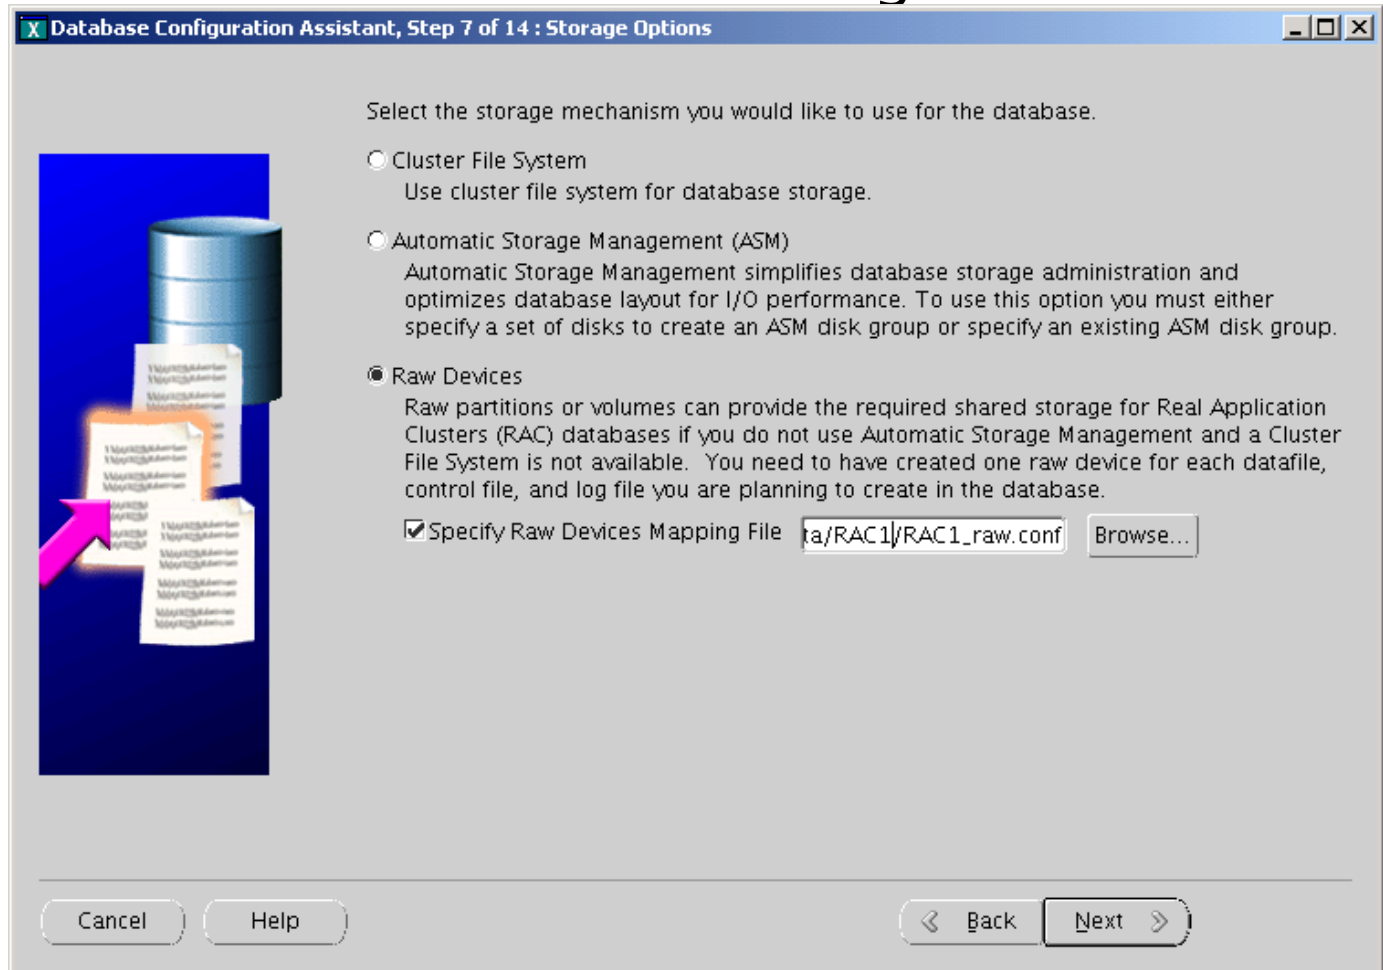
RAC Installation (8.2)

- Create RAC database using DBCA



RAC Installation (8.3)

- Create RAC database using DBCA



RAC Installation (8.4)

- Create RAC database using DBCA

Raw Devices Mapping file rac1_raw.conf

```
sysaux=/dev/raw/raw5  
example=/dev/raw/raw6  
users=/dev/raw/raw7  
system=/dev/raw/raw8  
undotbs1=/dev/raw/raw9  
undotbs2=/dev/raw/raw10  
redo1_1=/dev/raw/raw11  
redo1_2=/dev/raw/raw12  
redo2_1=/dev/raw/raw13  
redo2_1=/dev/raw/raw14  
control1=/dev/raw/raw15  
control2=/dev/raw/raw16  
spfile=/dev/raw/raw17  
pwdfile=/dev/raw/raw18  
temp=/dev/raw/raw19
```

Cluster Admin Utility (1)

- Start|Stop CRS
- Run as root

```
/etc/init.d/init.crs start           (10.1.0.4 and above)  
/etc/init.d/init.crs stop  
/etc/init.d/init.crs enable  
/etc/init.d/init.crs disable
```

Cluster Admin Utility (2)

- SRVCTL
 - Start or stop database/instance/application
 - Manage configuration

```
$ srvctl start database -d RAC1
```

```
$ srvctl stop instance -d RAC1 -i RAC12
```

```
$ srvctl status database -d RAC1
```

```
Instance RAC11 is running on linux1
```

```
Instance RAC11 is running on linux1
```

Cluster Admin Utility (3)

- `$ORACLE_HOME/bin/crs_stat`

NAME=ora.RAC1.RAC11.inst
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux1

NAME=ora.RAC1.RAC12.inst
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux2

NAME=ora.linux1.ASM1.asm
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux1

NAME=ora.linux2.ASM2.asm
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux2



NAME=ora.linux1.LISTENER_LINUX1.lsnr
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux1

NAME=ora.linux2.LISTENER_LINUX2.lsnr
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux2

NAME=ora.linux1.vip
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux1

NAME=ora.linux2.vip
TYPE=application
TARGET=ONLINE
STATE=ONLINE on linux2

...

Testing RAC with TAF

- TAF – Transparent Application Failover

```
DB3.WORLD =  
  (DESCRIPTION =  
    (ADDRESS = (PROTOCOL = TCP)(HOST = 192.168.0.129)(PORT =  
      1521))  
    (ADDRESS = (PROTOCOL = TCP)(HOST = 192.168.0.139)(PORT =  
      1521))  
    (LOAD_BALANCE = yes)  
    (CONNECT_DATA =  
      (SERVER = DEDICATED)  
      (SERVICE_NAME = RAC1.WORLD)  
      (FAILOVER_MODE =  
        (TYPE = SELECT)  
        (METHOD = BASIC)  
        (RETRIES = 120)  
        (DELAY = 4)  
      )  
    )  
  )  
)
```

Demo

- Server crashes during SELECT
- Server crashes during INSERT
 - Transaction is committed;
 - Transaction is not committed;



Questions?