

Darron Clark
Senior Oracle Instructor
darron.clark@oracle.com
908-420-3053



Shulman Fleming & PARTNERS

WORKFORCE ARCHITECTS

[Apply for this job](#)[Email to a friend](#)

Title: Oracle DBA / 9i-10G and RAC with Linux

Skills: Oracle DBA / 9i-10G and RAC with Linux

Date: 4-9-2005

Location: NYC, NY

Area code: 212

Tax term: FULLTIME

Pay rate: Open

Length:

Position ID: Qdba

Dice ID: shulman

Job description:

We need an Oracle DBA who can work in a 24x7 environment where you will be the only DBA on staff. Responsibilities will include installation, performance tuning, high availability (RAC) and backups/monitoring. Should be a good scripter and know some programming. You must know Oracle TAR system.

Travel required: none


Telecommute: no

[Apply for this job](#)[Email to a friend](#)

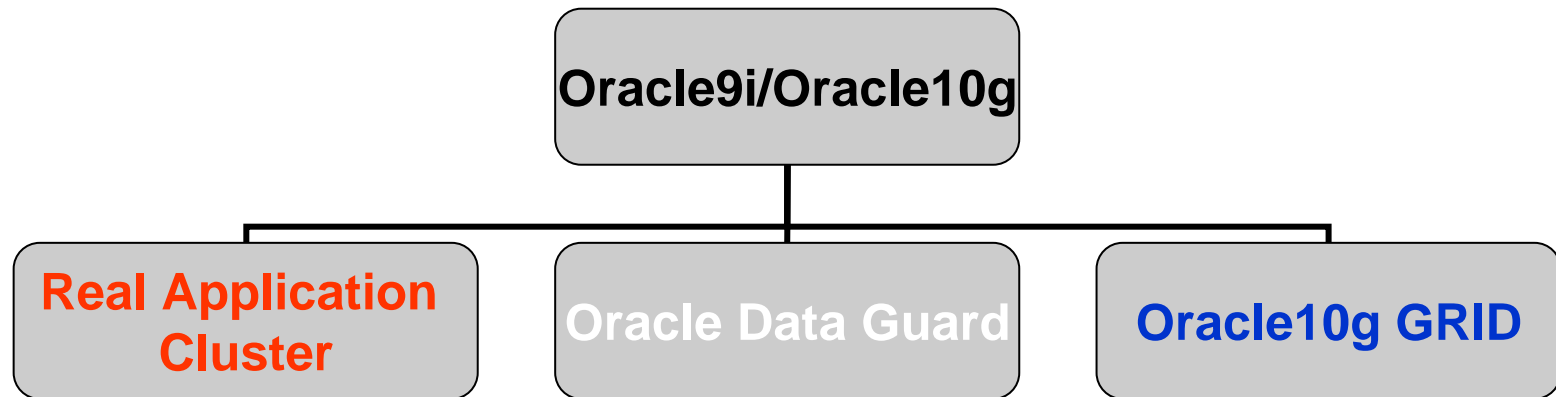
State of the Oracle Database Industry

- | | |
|------------------------------|-------------------|
| 1. Retire pension/gold watch | 1. Before 1980s |
| 2. BSCS, BSEE, MIS | 2. 1980s |
| 3. MCSE, Novell Certified | 3. Early 1990s |
| 4. OCP | 4. Late 1990s |
| 5. Oracle Masters | 5. 2000 |
| 6. Data Guard | 6. Now! |
| 7. Linux and Oracle | 7. Now! |
| 8. Real Application Clusters | 8. Now! |
| 9. Oracle10g Grid | 9. Staying ahead. |

Which direction should I go?



Which direction should I go?



Why should I choose Real Application Clusters?

Why Enhance My Career?

- 1. Resilience against “bench” and “downtime”**
- 2. Easier transition to new environments**
- 3. Enhanced performance and service response times**

Grid Control

Manage multiple systems with one system:

- 1. Oracle Application Server 10g**
- 2. Oracle Database 10g**
- 3. Oracle Collaboration Suite 10g**
- 4. Oracle Enterprise Manager 10g**

What is RAC?

In Real Application Clusters environments, all active instances can concurrently execute transactions against a shared database.

What is RAC?

Technical Overview

Let's Build our own RAC System

Hardware

1. Nodes
2. Interconnect
3. Shared storage for your data

Operating Software

1. Cluster Software - Cluster Manager
2. Oracle Software
3. Oracle Database and Instances


Nodes and Their Components

A node has the following main components:


- 1. CPUs**
- 2. Memory**
- 3. Interconnect**
- 4. Storage for OS, Cluster and Oracle Software**
- 5. Operating System Software - OS**
- 6. Cluster Software or Cluster Manager**

Nodes


node1



node2



node3



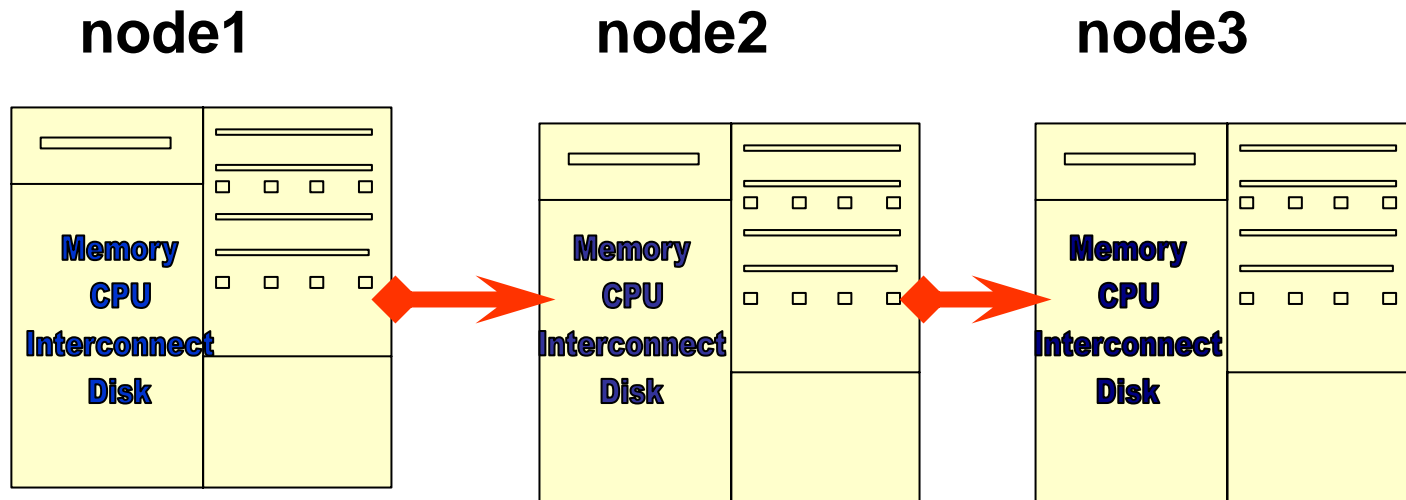
Interconnect

High Speed network connection for:

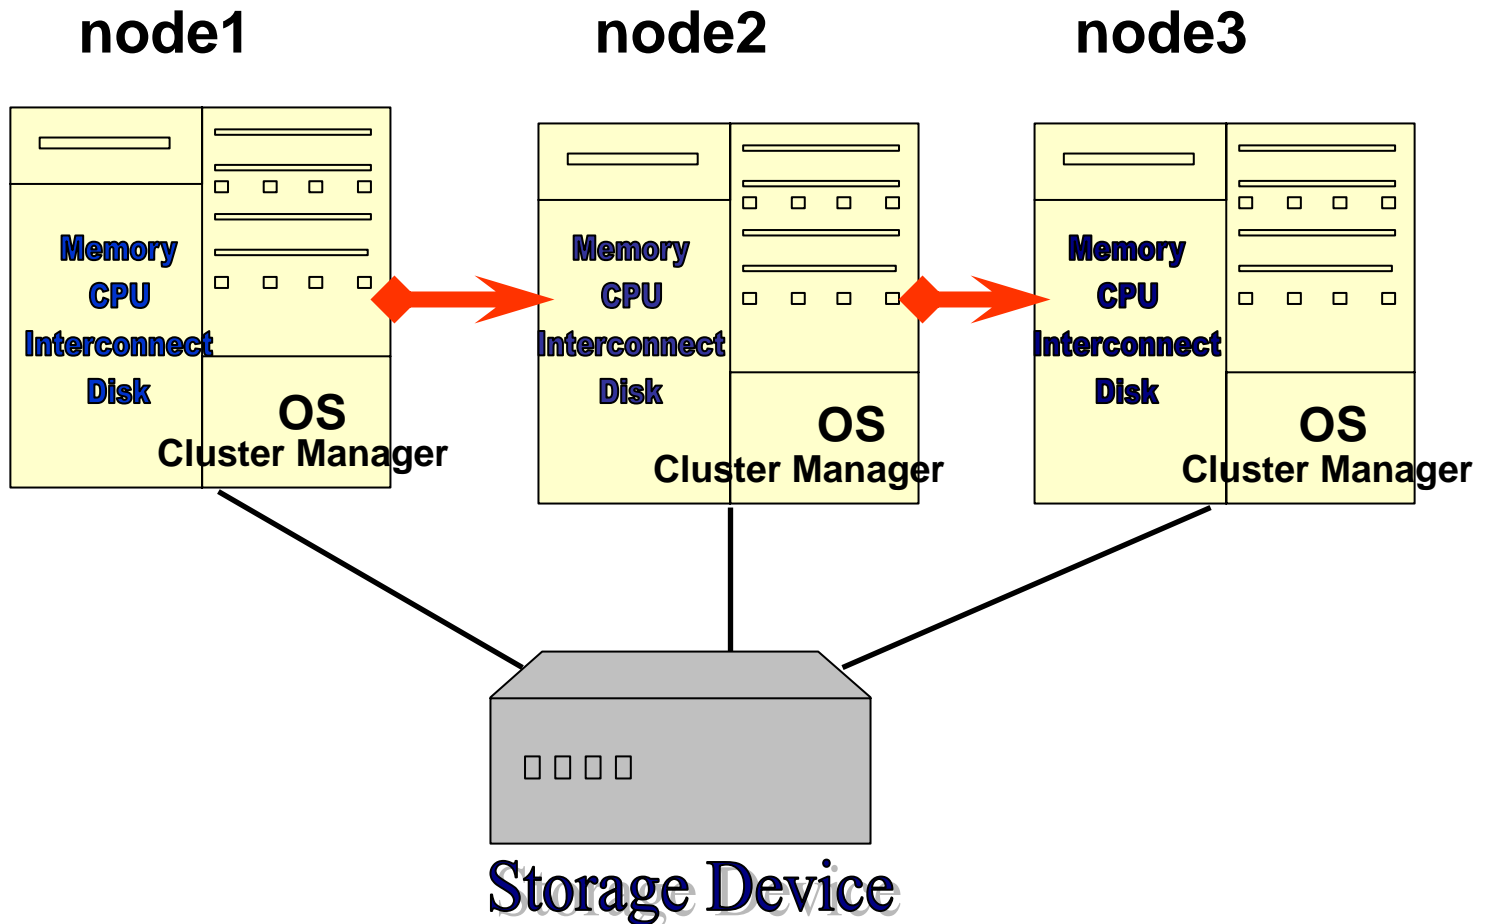
- **Monitoring, message**
- **Cache transfer**
- **Lock transfer**
- **Extents information**
- **Freelists information**

1. **Infiniband**
2. **Fast Ethernet**
3. **Gigabit Ethernet**

Nodes and Interconnects



Nodes and Their Components



Shared Storage

This is where the data and control files will reside

- **RAW**
- **Cluster file system**
- **Network Attached Storage**
- **Storage Area Network**
- **RAID**
- **Automatic Storage Management – 10g**

RAW Device

1. Have been in use for a long time
2. Bypasses the Operating System buffer cache
3. Can be used in 9i or 10g
4. **Difficult to manage**
5. High performance
6. **AUTOEXTEND forbidden**

RAW Device

On RAW Devices

1. Data files
2. Redo log files
3. Control files
4. Oracle Cluster Ready files
5. Voting Disk

On Local File System

1. Archive log files
2. Oracle Home
3. Alert ,Trace files
4. External Tables
5. Voting Disk

Cluster File System

1. Provides a shared file system for all cluster nodes.
2. Can share datafiles on Oracle Home in same storage area.
3. **Simple to manage**
4. Use of Oracle Managed files
5. **AUTOEXTEND feature okay**
6. Voting Disk

Cluster File System

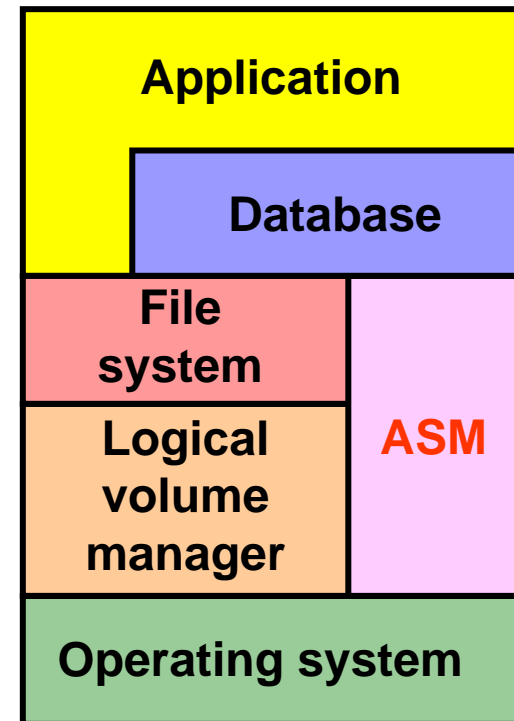
1. **Data files**
2. **Redo log files**
3. **Control files**
4. **Oracle Cluster Ready files**
5. **Archive log files**
6. **Oracle Home**
7. **Alert ,Trace files**
8. **External Tables**

Automatic Storage Management

1. New star on the block for 10g
2. Ease of Administration
3. Should get almost the same performance of RAW
4. Eliminates the need for cluster file system
5. Eliminates the need for volume management
6. No Support for:
 1. External Tables
 2. Transportable Tables
 3. No Export/Import output file

What Is Automatic Storage Management?

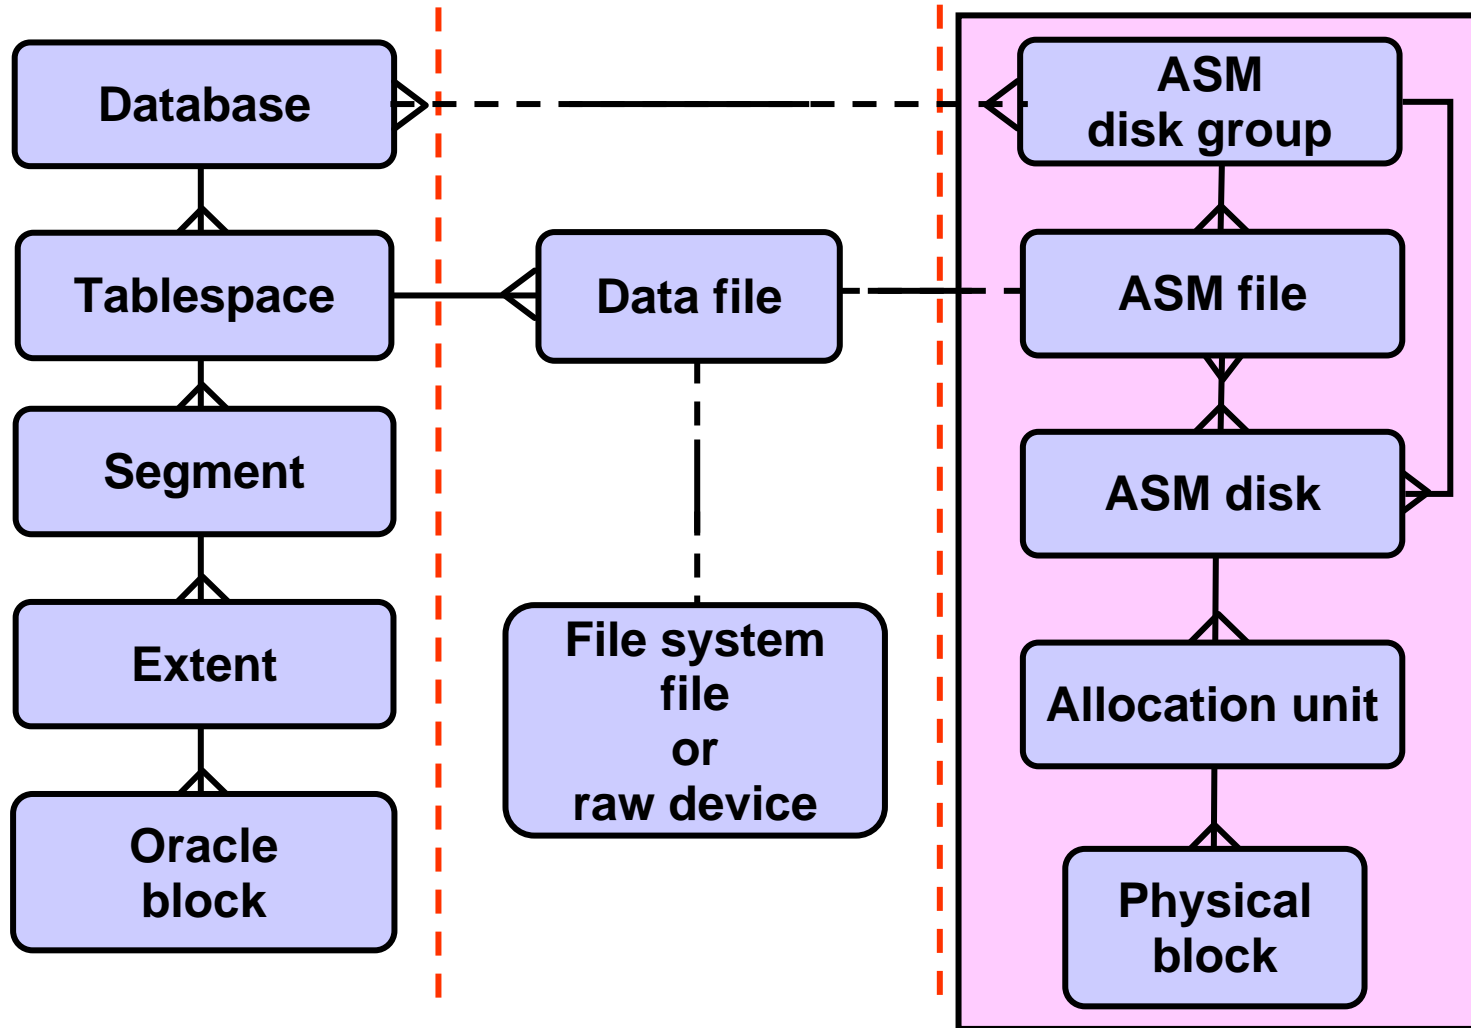
- Is a purpose-built cluster file system and volume manager
- Manages Oracle database files
- Spreads data across disks to balance load
- Provides integrated mirroring across disks
- Solves many storage management challenges



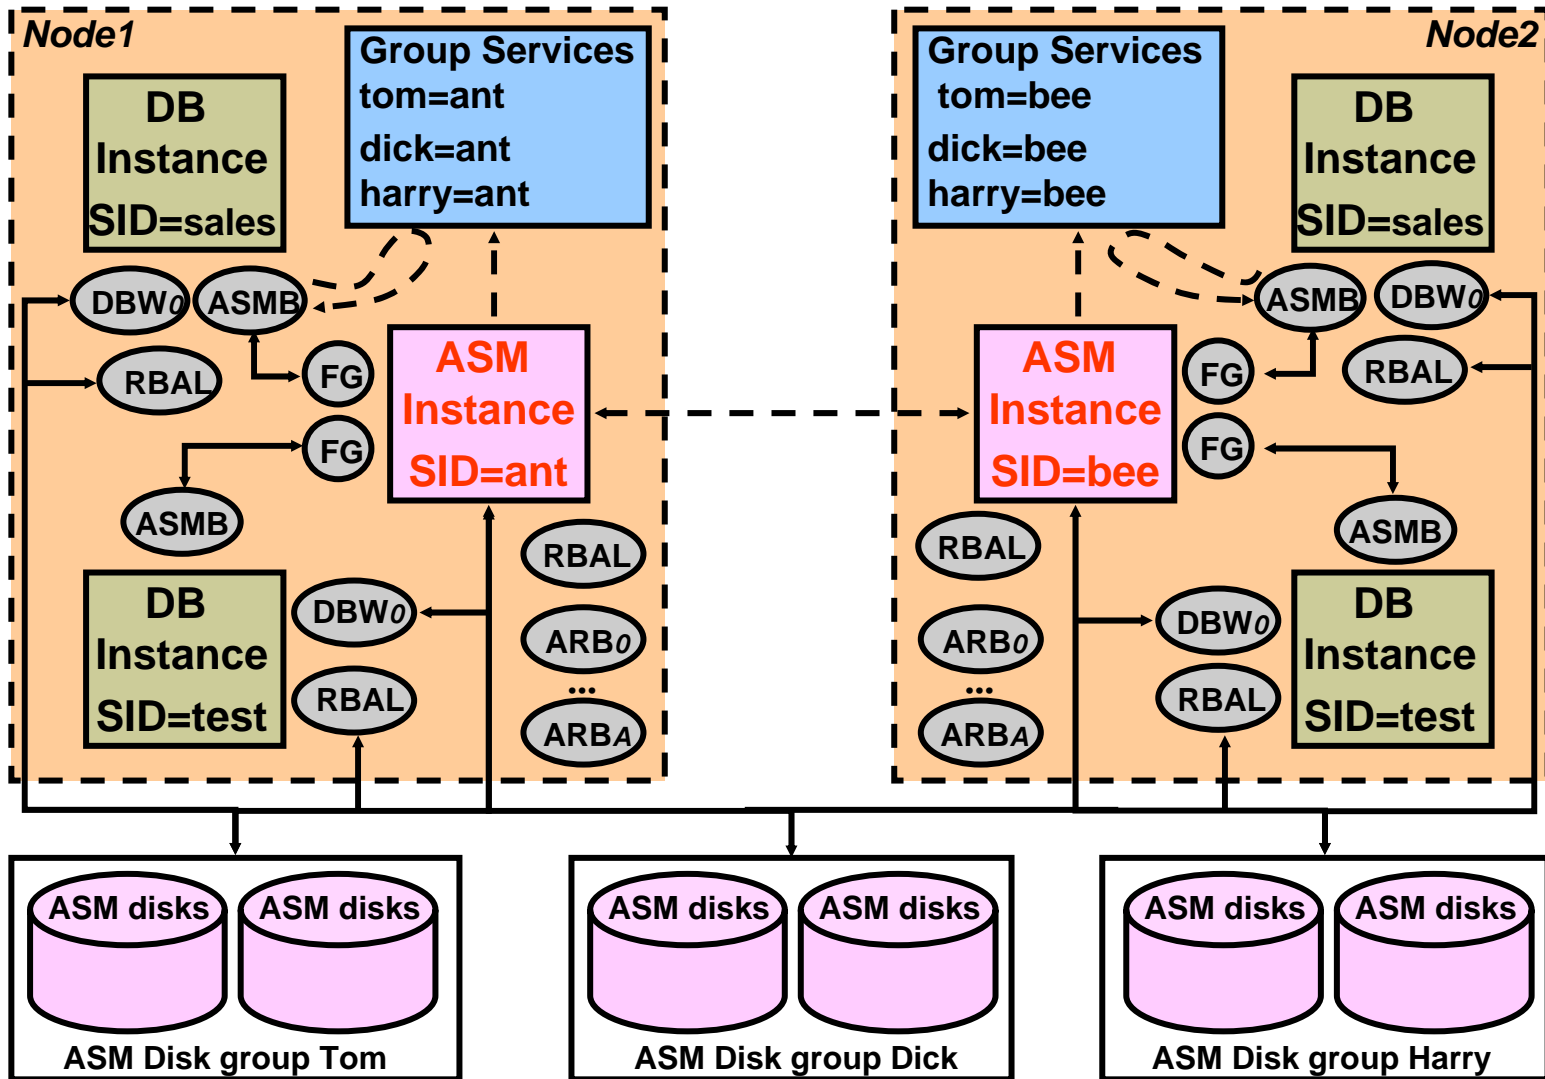
ASM: Key Features and Benefits

- **Stripes files rather than logical volumes**
- **Enables online disk reconfiguration and dynamic rebalancing**
- **Provides adjustable rebalancing speed**
- **Provides redundancy on a file basis**
- **Supports only Oracle files**
- **Is cluster aware**
- **Is automatically installed as part of the base code set**

ASM: New Concepts

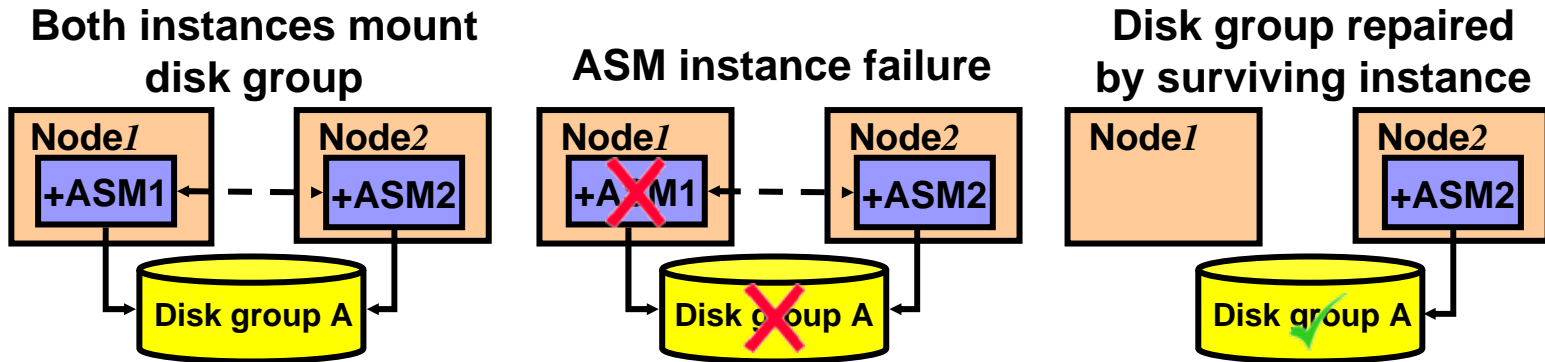


ASM: General Architecture

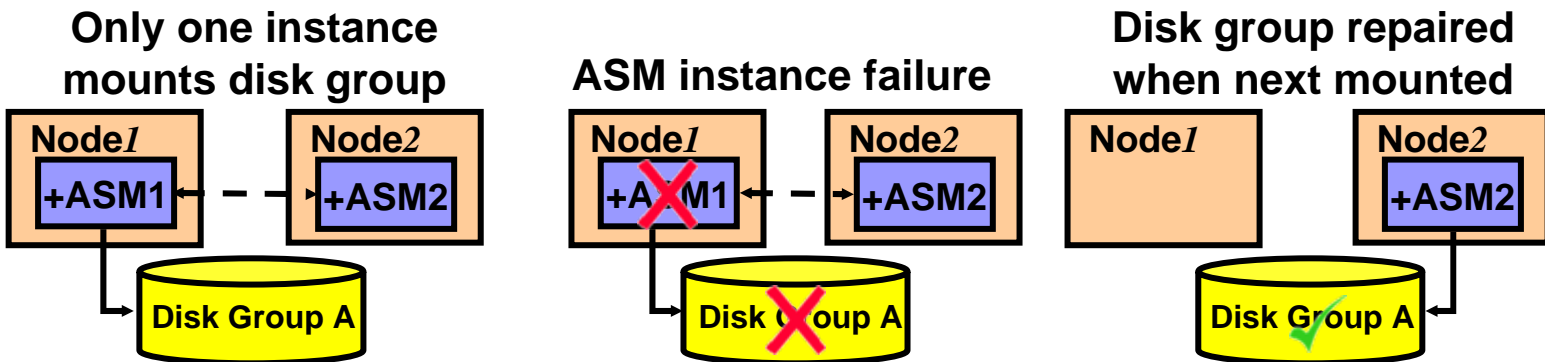


ASM Instance and Crash Recovery in RAC


ASM instance recovery



ASM crash recovery



Nodes and Their Components



Operating System

1. Linux
2. Windows
3. Solaris
4. HP
5. IBM AIX

1. Redhat and United Linux
2. 2000 and 2003
3. Cluster 3.x(raw) and Veritas DBE(CFS)
4. Itanium, Tru64 and TruCluster
5. 4.2(raw) and 5.1(both)

Check Oracle certification Matrix and vendor suppliers to choose right combination of RAC.

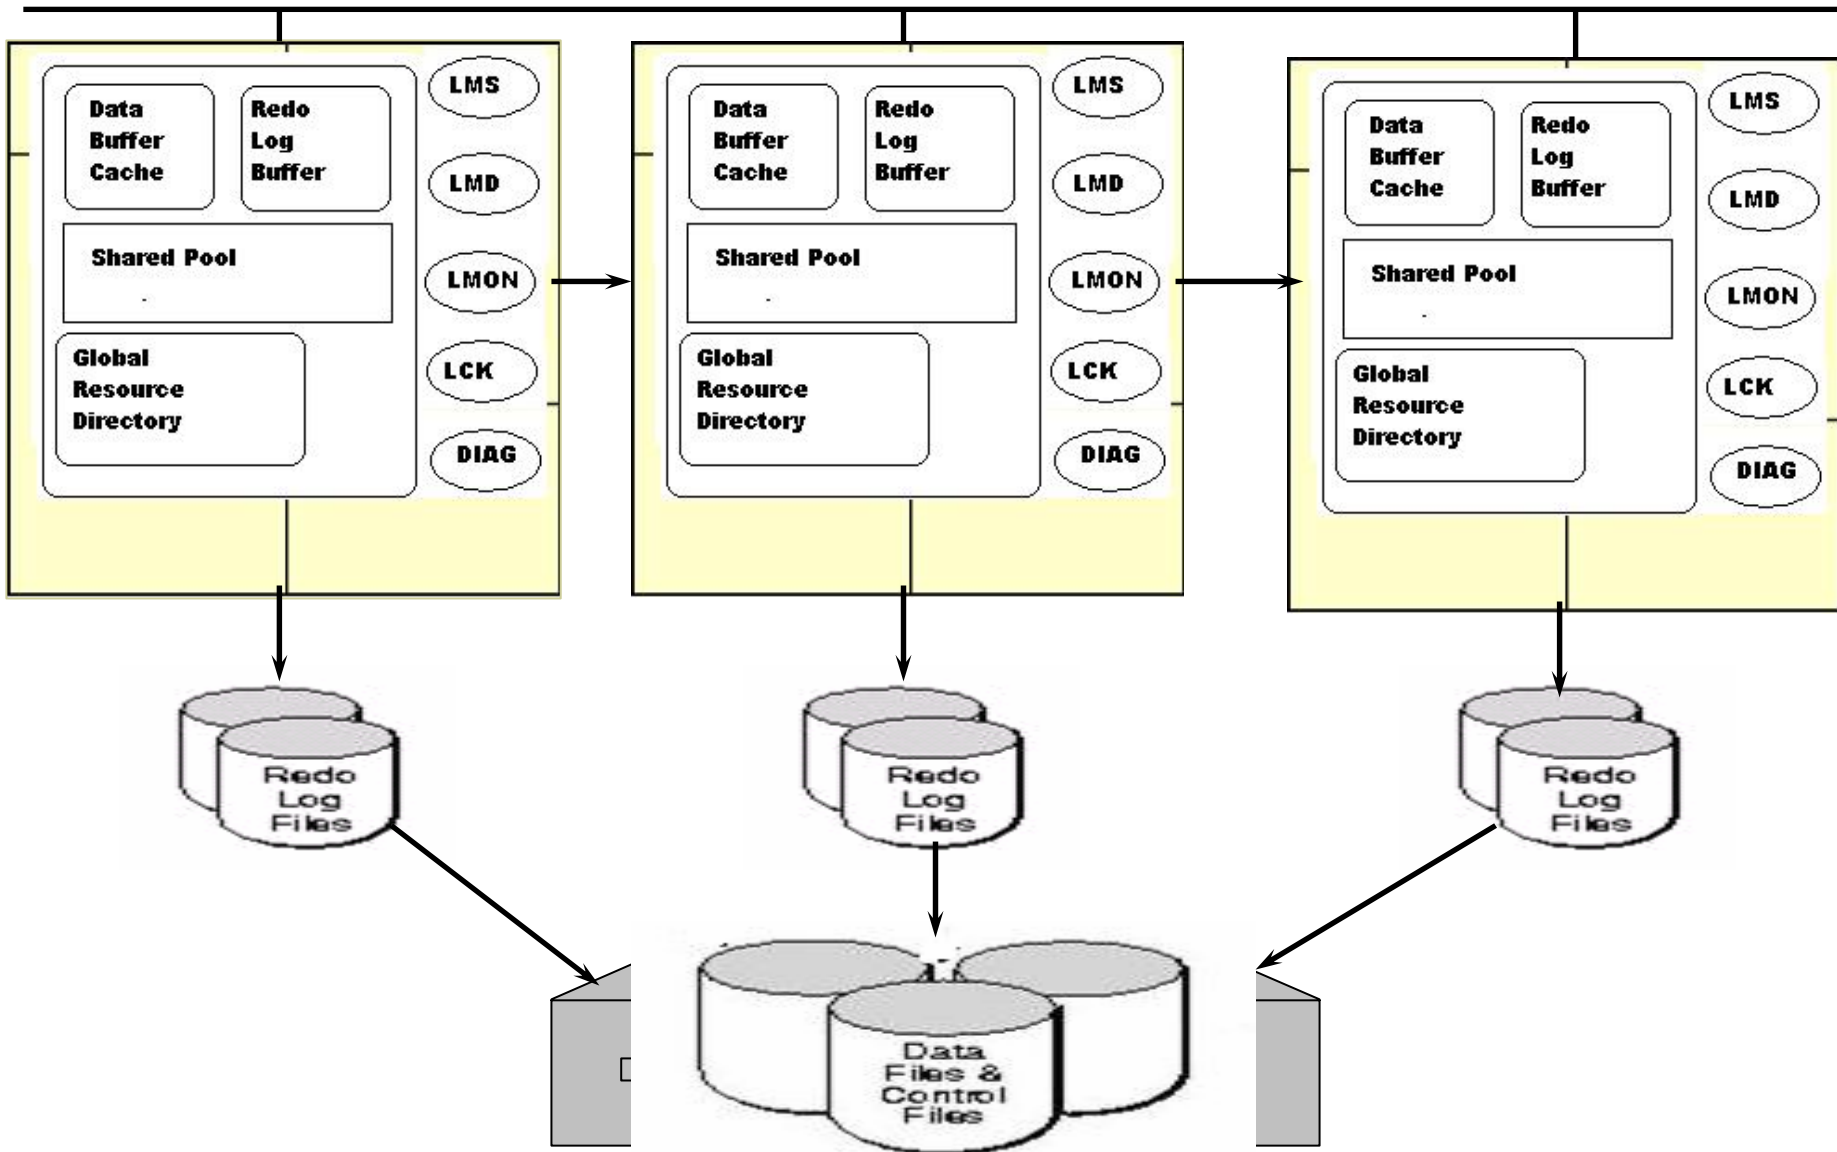
Cluster Software

- | | |
|---------------------------|----------------------|
| 1. Cluster Ready Services | 1. Windows and Linux |
| 2. Sun 3.x or Veritas DBE | 2. Sun Solaris |
| 3. HACMP | 3. IBM |
| 4. Cluster services | 4. HP |

With Oracle10g the Cluster Software is provided with the RAC software.

Real Application Clusters-Specific Instance Processes

Interconnect Communication



RAC Instance

Global Resource Directory – Tracking status of oracle blocks.

Additional Background processes:

- 1. LMON - Monitor Instance status**
- 2. LMSn - Cache Fusion management**
- 3. LCK - Enqueue requests**
- 4. LMD - Enqueue management**
- 5. DIAG - Health of RAC instances**
 - PMON restart if it dies.**

Global Resource Directory

1. Data Block Identifier
2. Location of most current version
3. Mode of the data block
 - Null (N), Shared (S), or Exclusive (X)
4. The Role of the data block
5. Information of Buffer caches on multiple nodes

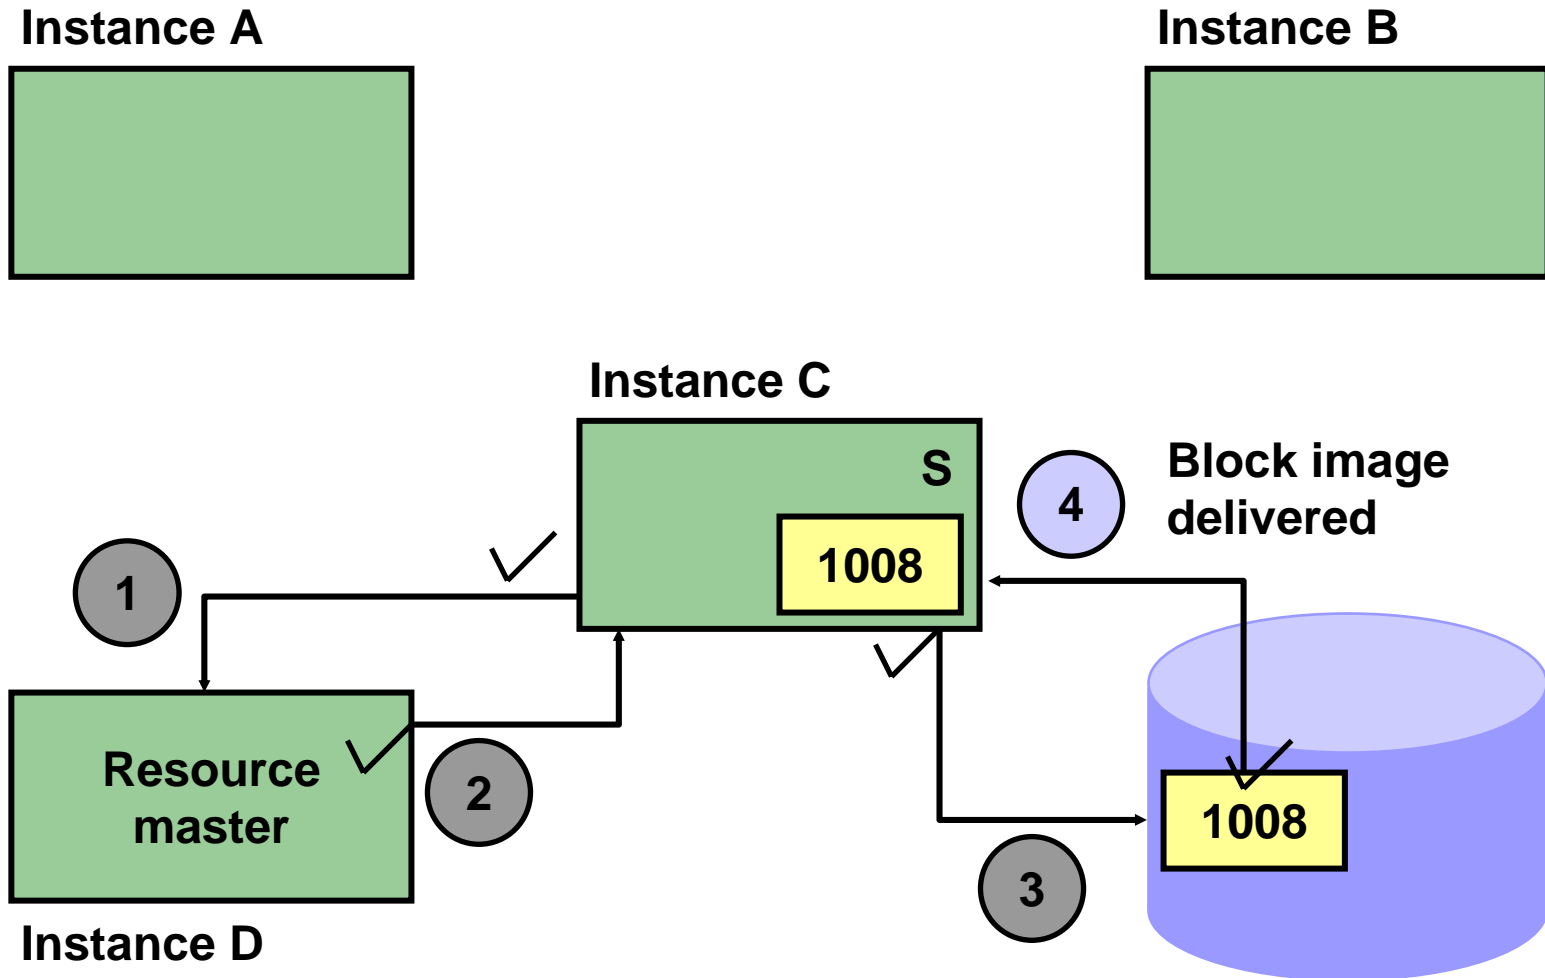
Cache Fusion

- Cache Fusion helps **provide transparent scalability** in a Real Application Clusters database.
- The algorithms **enable transportation of block images between instances.**
- Cache Fusion services **track the current location and status of resources.**
- Directory structures in the SGA of each instance store the resource information.

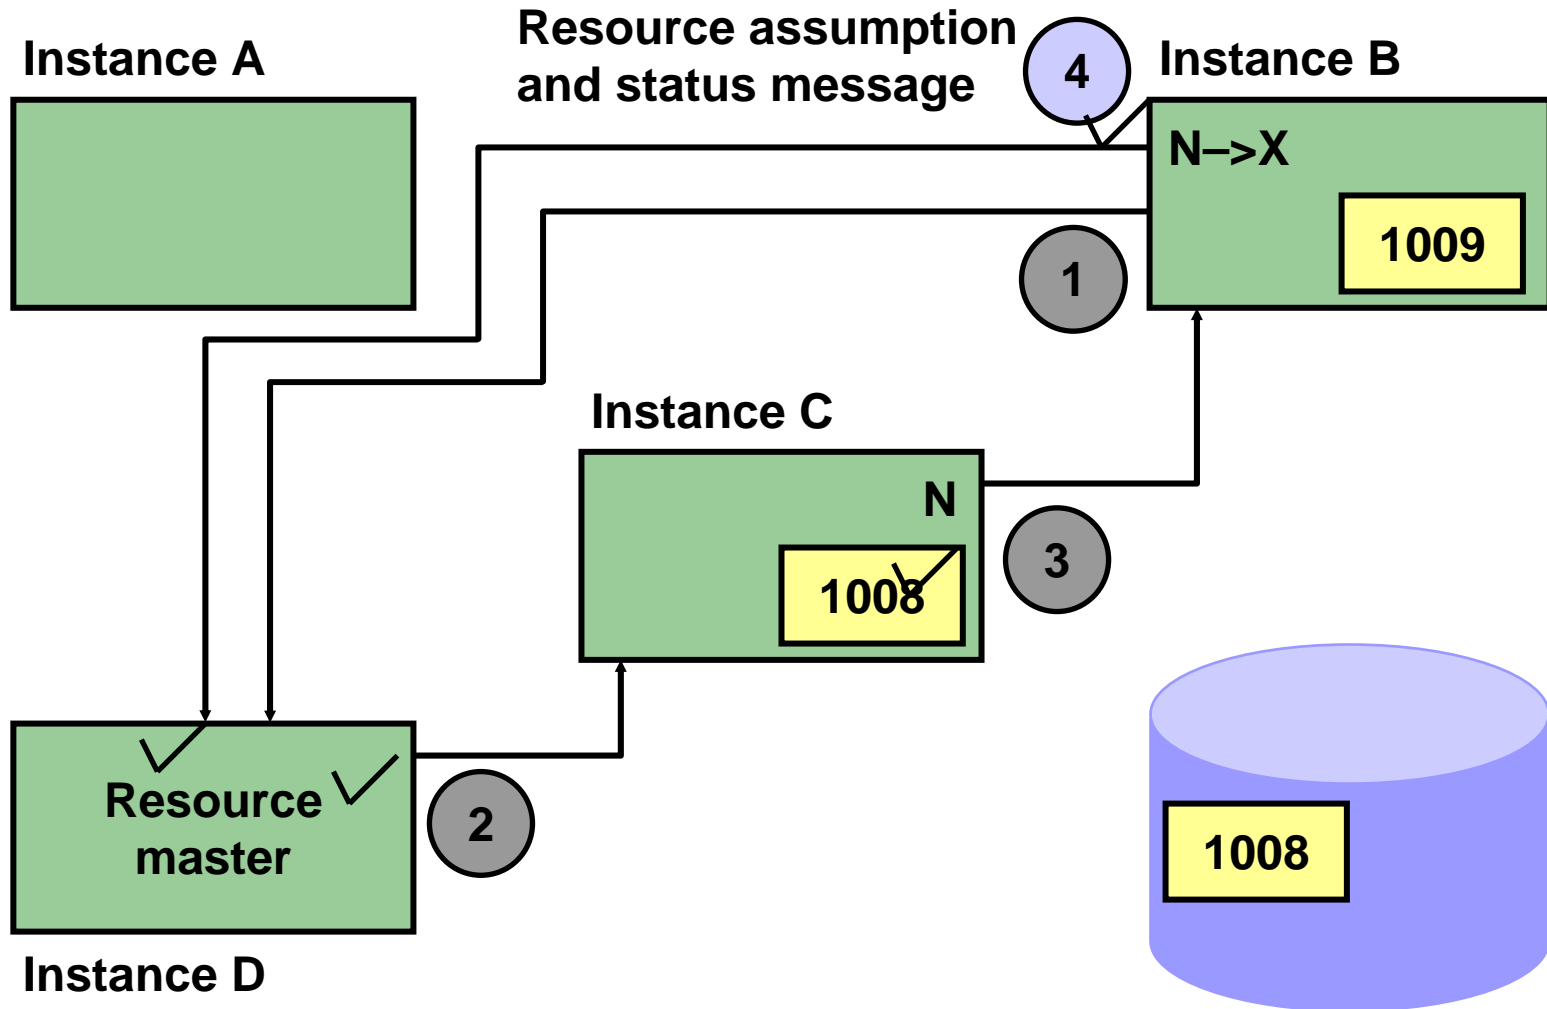
Cache Fusion Scenarios

- 1. Read with no Transfer**
- 2. Read/Write**
- 3. Write/Write**
- 4. Write/Read**

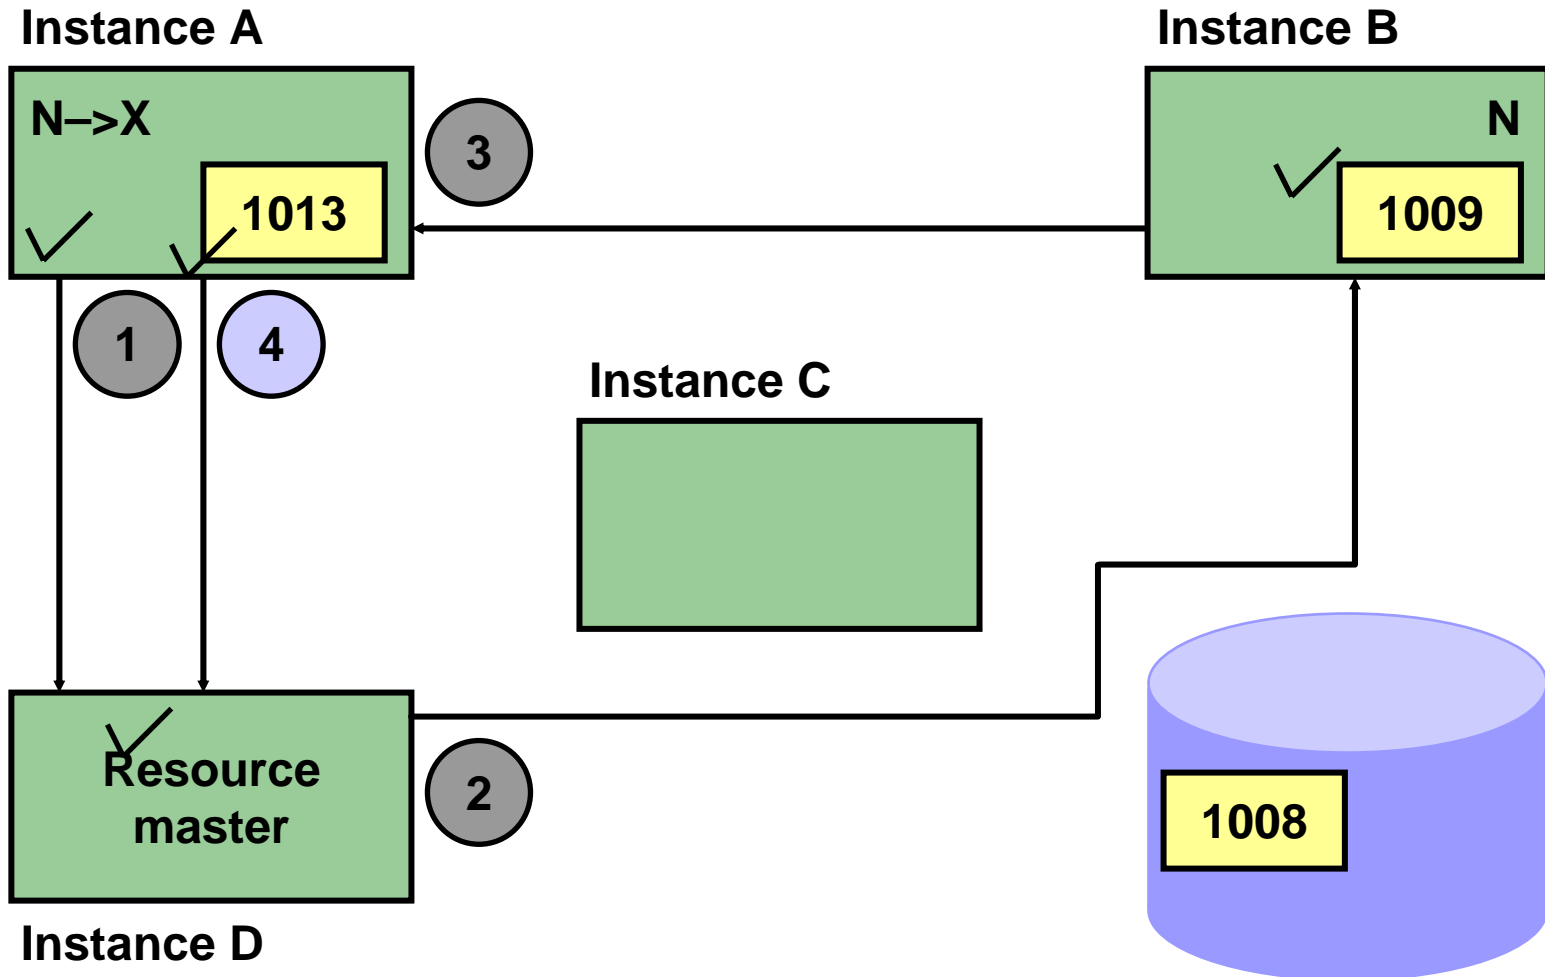
Example 1: Read with No Transfer



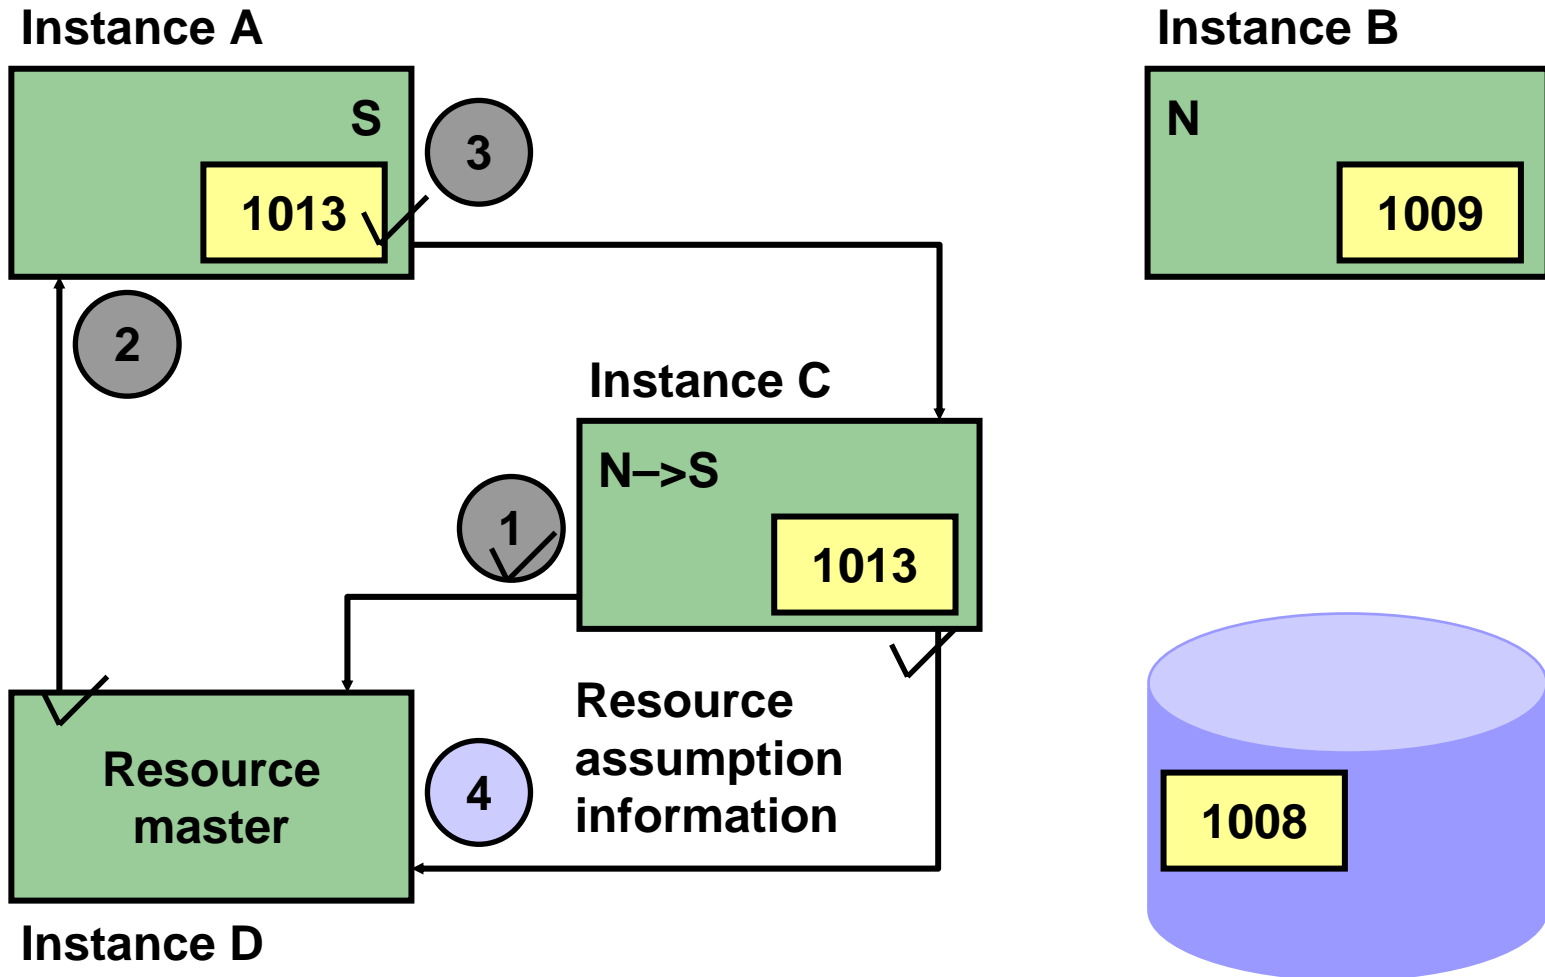
Example 2: Read to Write Transfer



Example 3: Write to Write Transfer

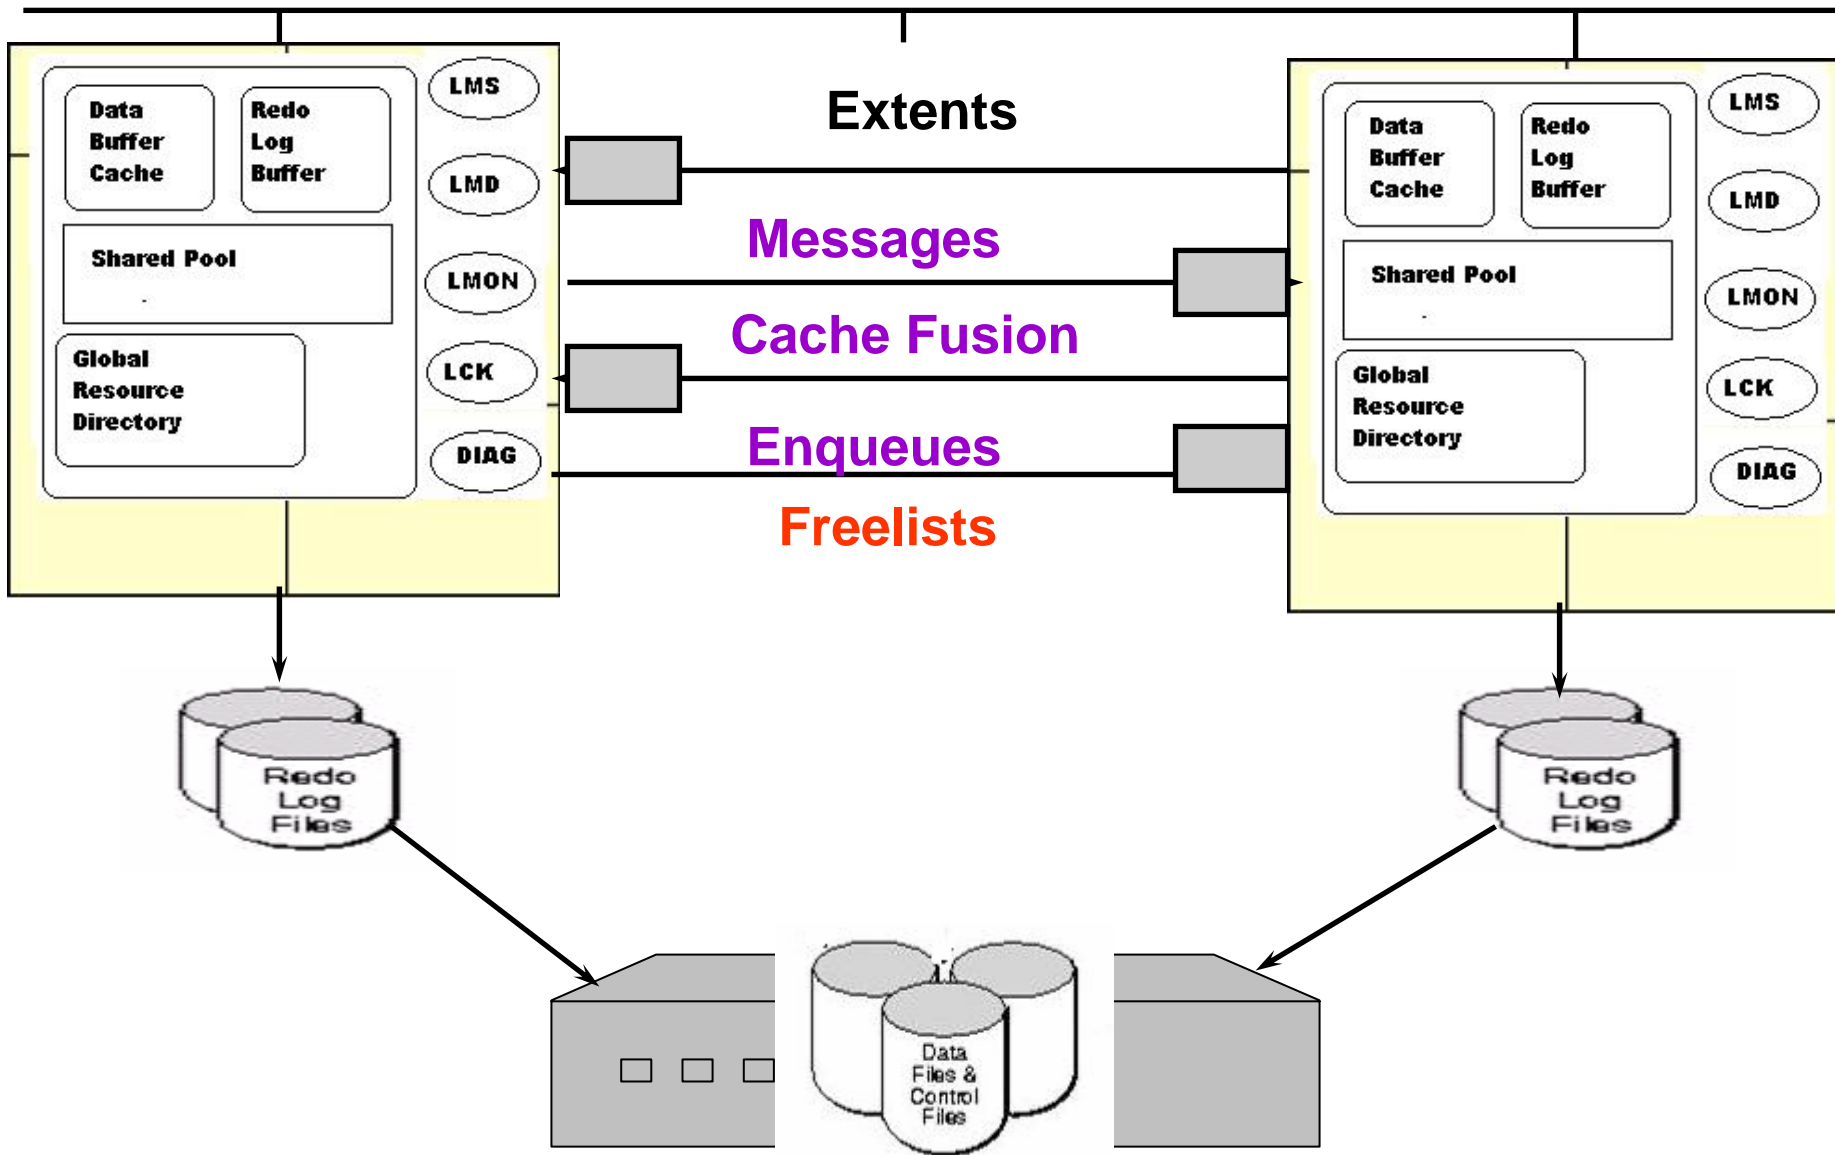


Example 4: Write to Read Transfer



Real Application Clusters-Specific Instance Processes

Interconnect Communication



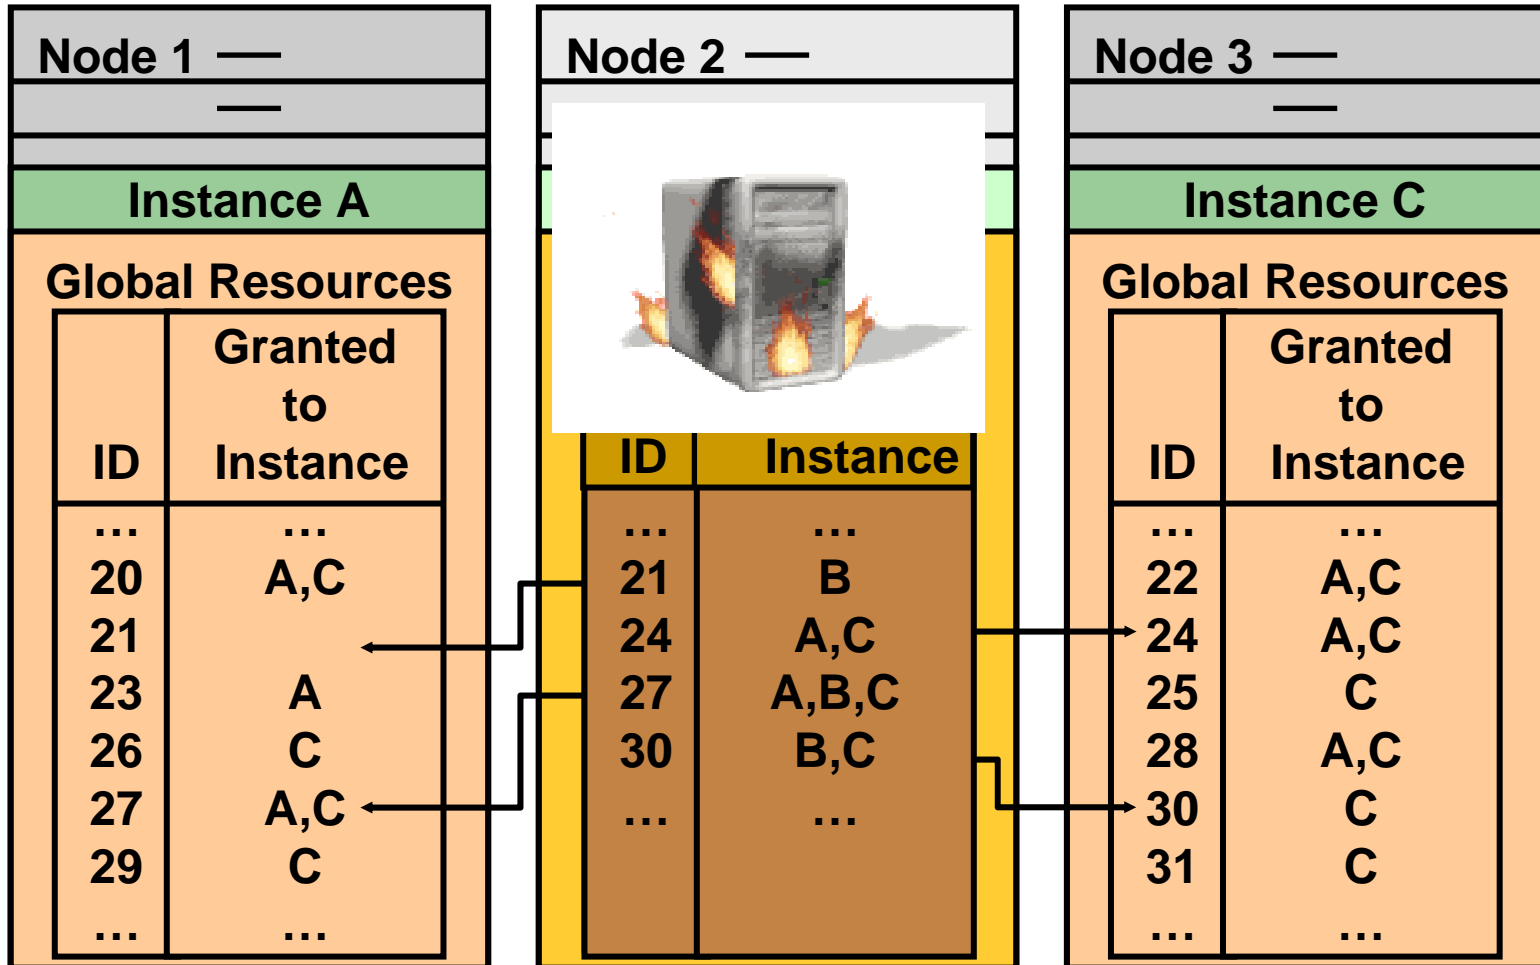
Cluster Reorganization: Example

Node 1 —	
Instance A	
Global Resources	
ID	Granted to Instance
...	...
20	A,B,C
23	A,B
26	C
29	B,C
...	...

Node 2 —	
Instance B	
Global Resources	
ID	Granted to Instance
...	...
21	B
24	A,C
27	A,B,C
30	B,C
...	...

Node 3 —	
Instance C	
Global Resources	
ID	Granted to Instance
...	...
22	A,C
25	C
28	A,B,C
31	B,C
...	...

Cluster Reorganization: Example



Best Practices

Do's

Local Management

***Automatic Segment Space
Management***

Localize bulk insert for B-tree

**Automatic Storage
Management 10g**

Don'ts

Dictionary Management

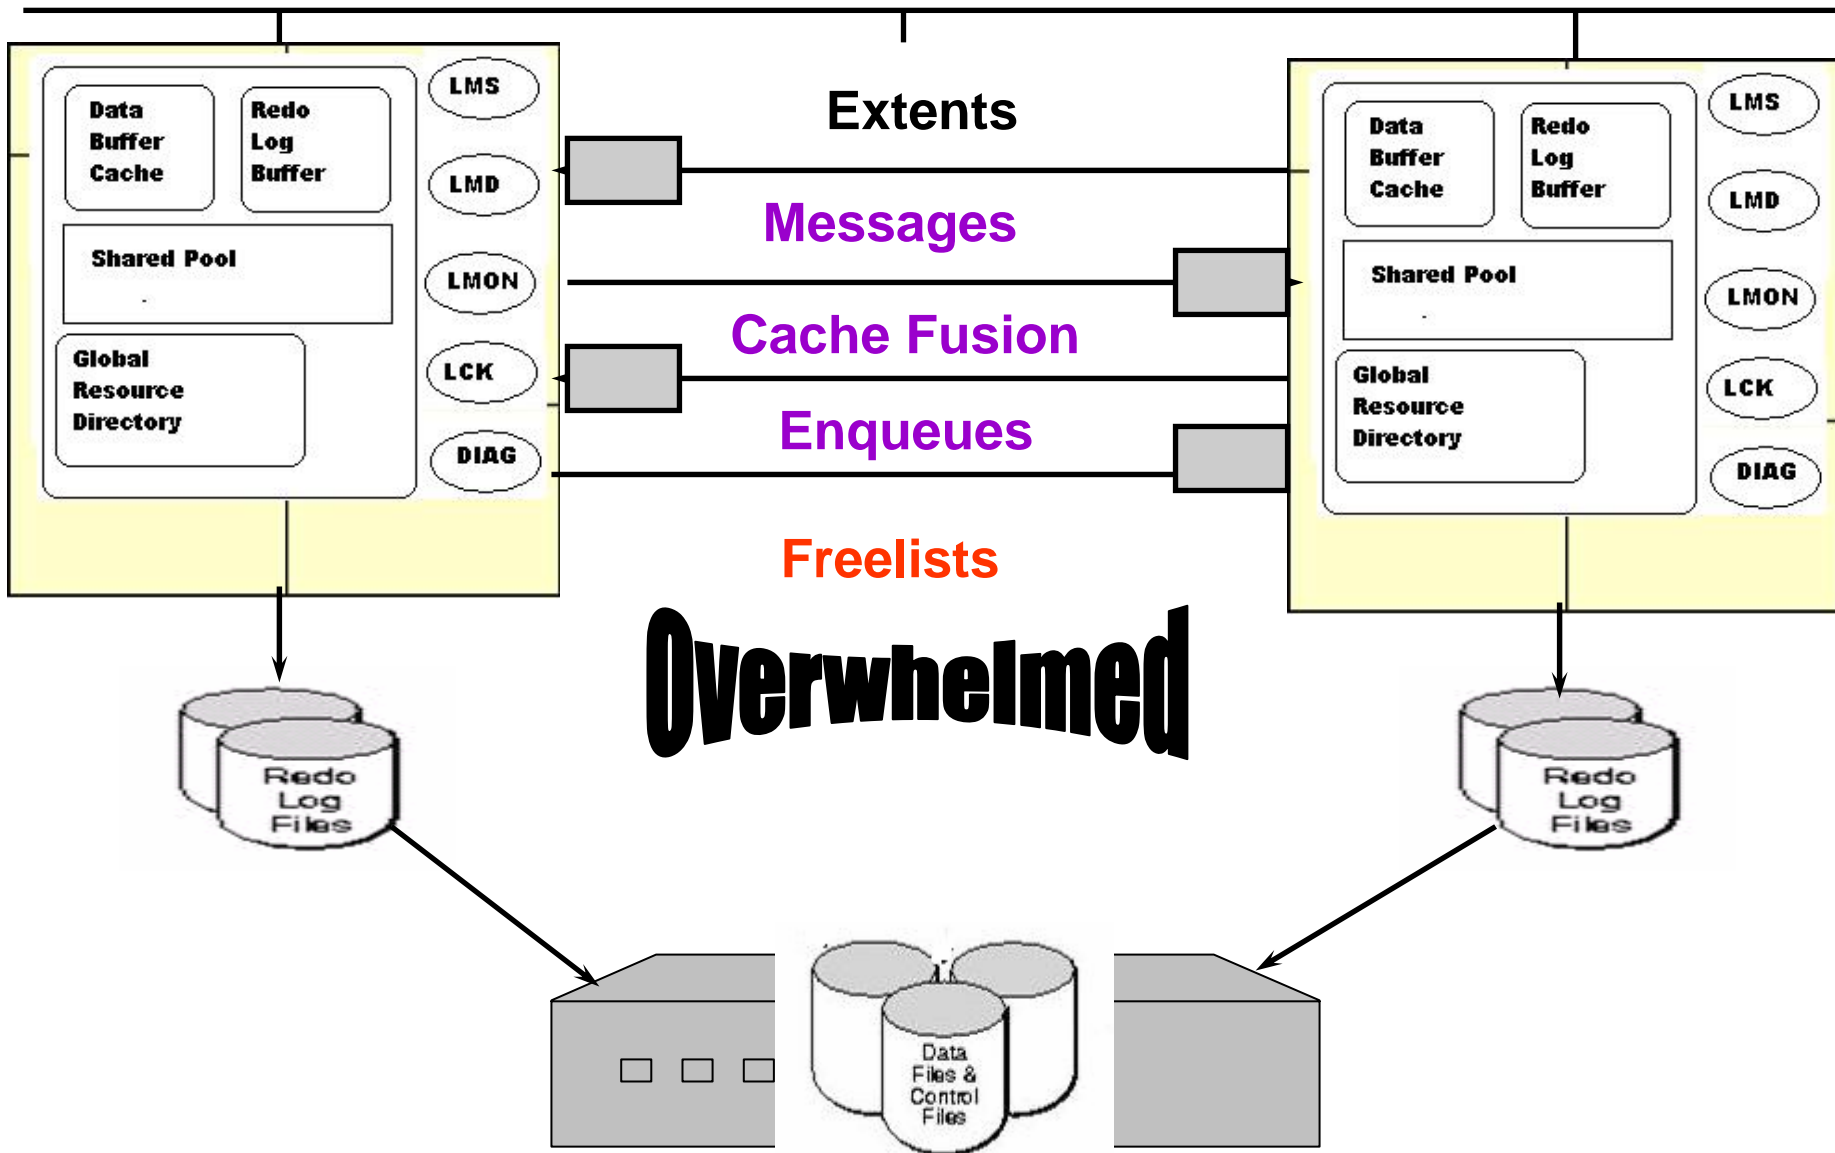
Freelist Management

Distribute inserts for B-tree


**Autoextend for datafile
increase on RAW devices**

Real Application Clusters-Specific Instance Processes


Interconnect Communication



Extent Management Options



Local management is controlled by bitmaps in the data files.



Dictionary management is controlled by the UET\$ and FET\$ tables.

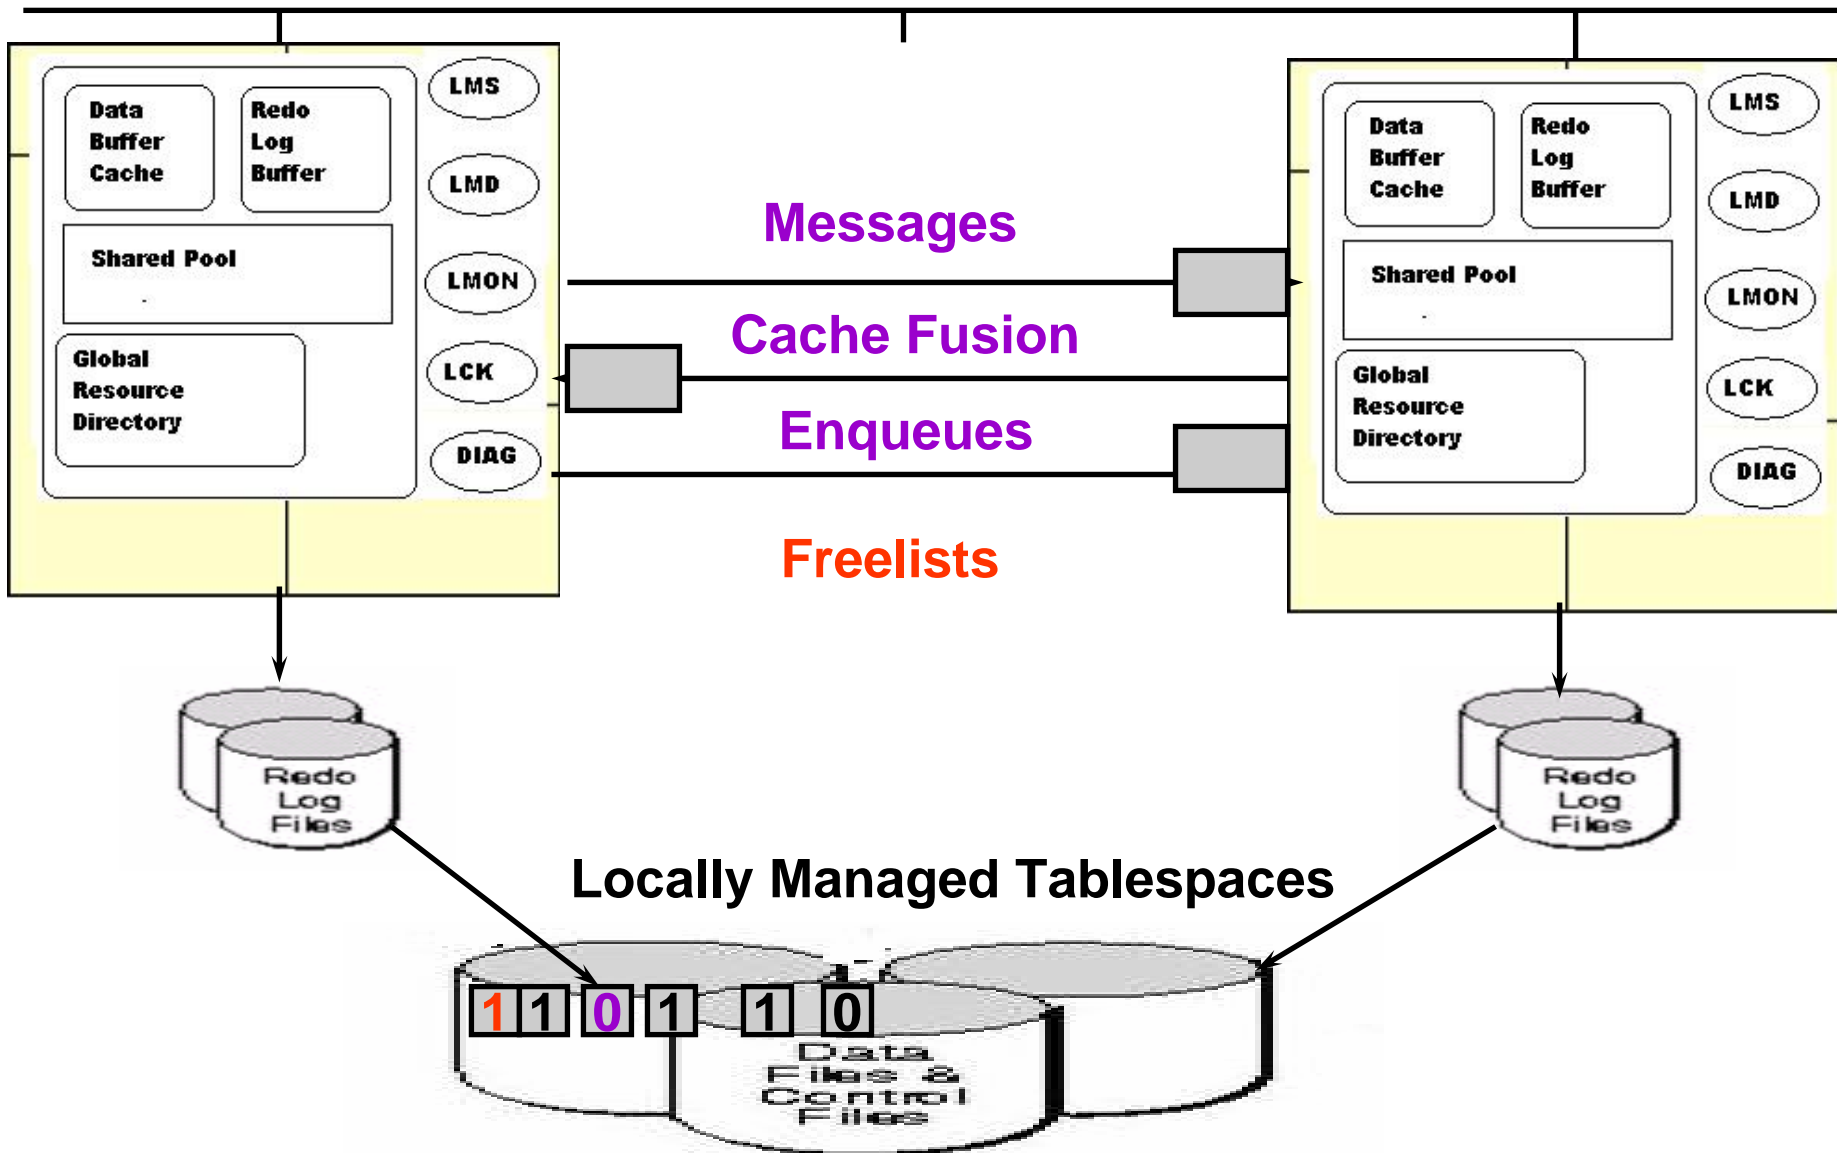
Locally Managed Tablespaces

Locally managed tablespaces:

- Are recommended by Oracle Corporation
- Pretend as though dictionary managed does not exist
- Avoid contention between instances for a small number of blocks (in UET\$ and FET\$) during extent management
- Remove fragmentation potential when different-sized extents share a tablespace
- Enable automatic segment free space management

Real Application Clusters-Specific Instance Processes

Interconnect Communication

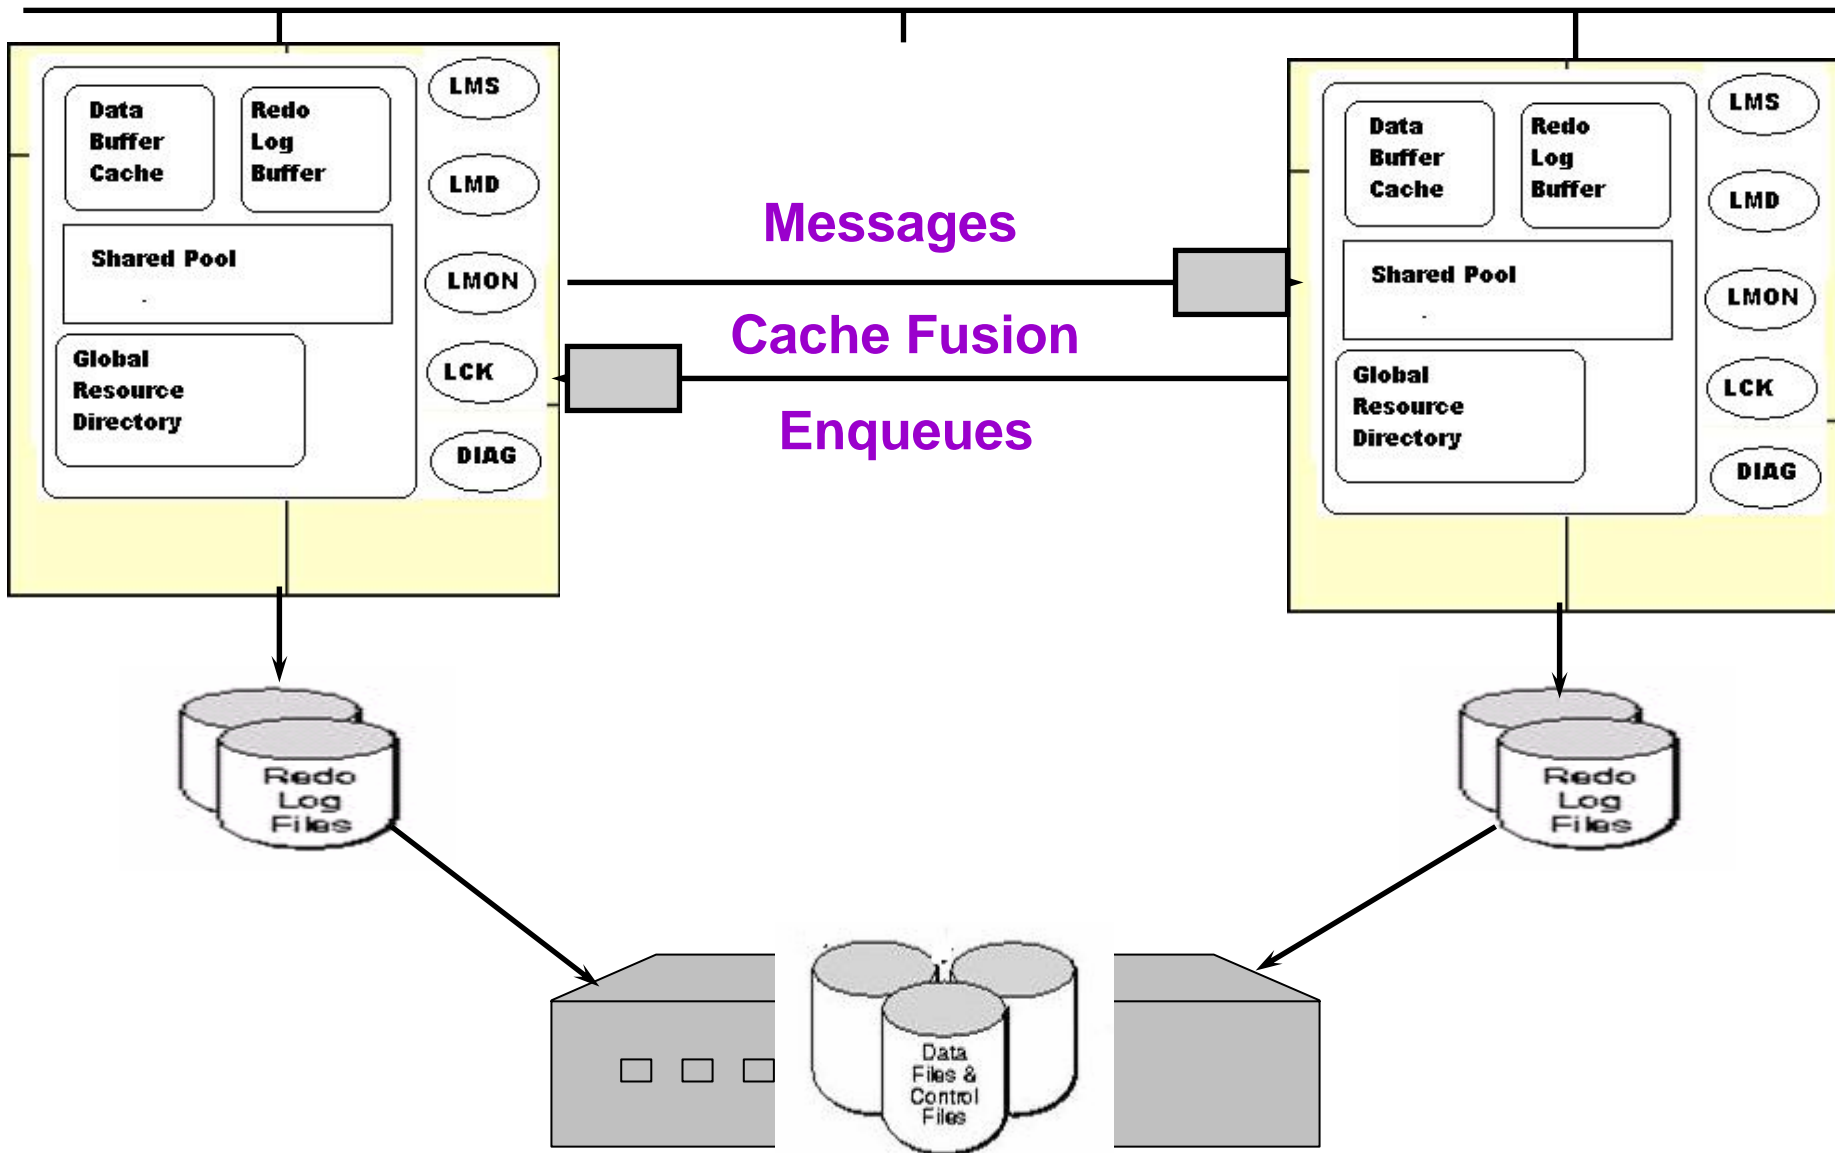


Automatic Segment-Space Management

- **Bitmap blocks are stored throughout a segment** using automatic space management.
- Each bitmap block **contains space-availability information** for a distinct subset of blocks.
- Only boundary condition changes in a block's free space availability are recorded.
 - **Only one bit needs to be changed** to record a change.
 - **These changes are fast and cause little contention.**
- Bitmap blocks are allocated to a session requiring free space based on the following:
 - **Instance number to avoid inter-instance contention**
 - **Session ID to avoid inter-session contention**

Real Application Clusters-Specific Instance Processes

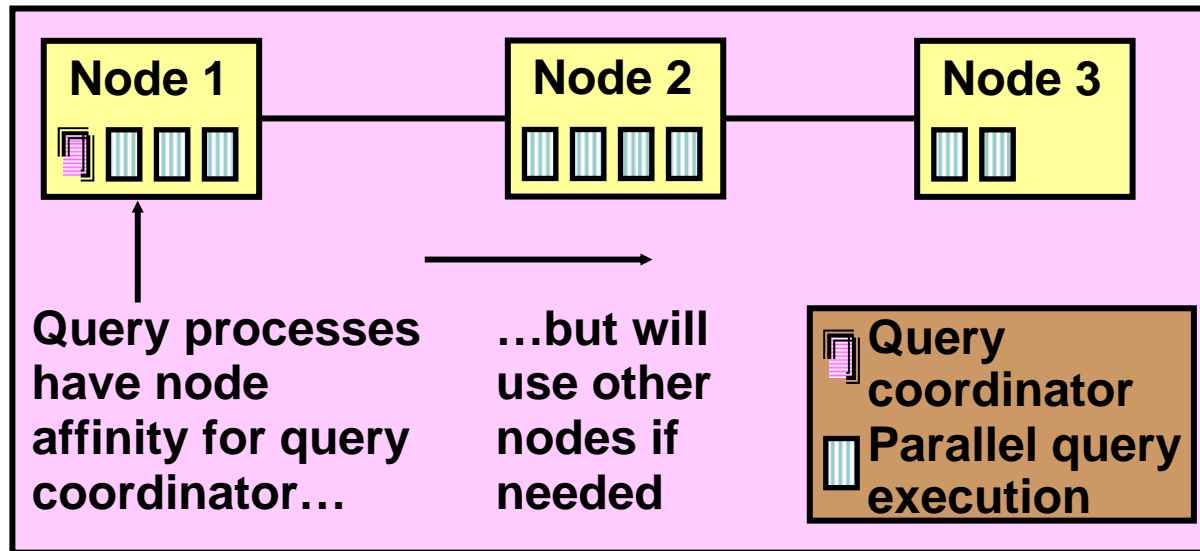

Interconnect Communication



Comparison of Free Space Management Methods

Method	Key Benefit	Main Drawback
Automatic segment-space management	Virtually no additional work once created	Upgraded database may require conversion to locally managed tablespaces
Manual free list group assignment	Can control extent location as well as instance access	Requires constant monitoring to avoid out-of-space errors
Automatic free list group assignment	Easier free list group management than manual method	No control over which instances acquire blocks

Adaptive Parallel Query



Query-Intensive Database Issues

- **Query-intensive databases include:**
 - Online analytic processing (OLAP) servers
 - Decision support systems (DSS)
 - Data warehouses
- **Such databases are characterized by:**
 - Large amounts of data
 - Extensive query access
 - Scheduled batch loads to refresh or replace data
- **High amounts of parallelism benefit processing.**

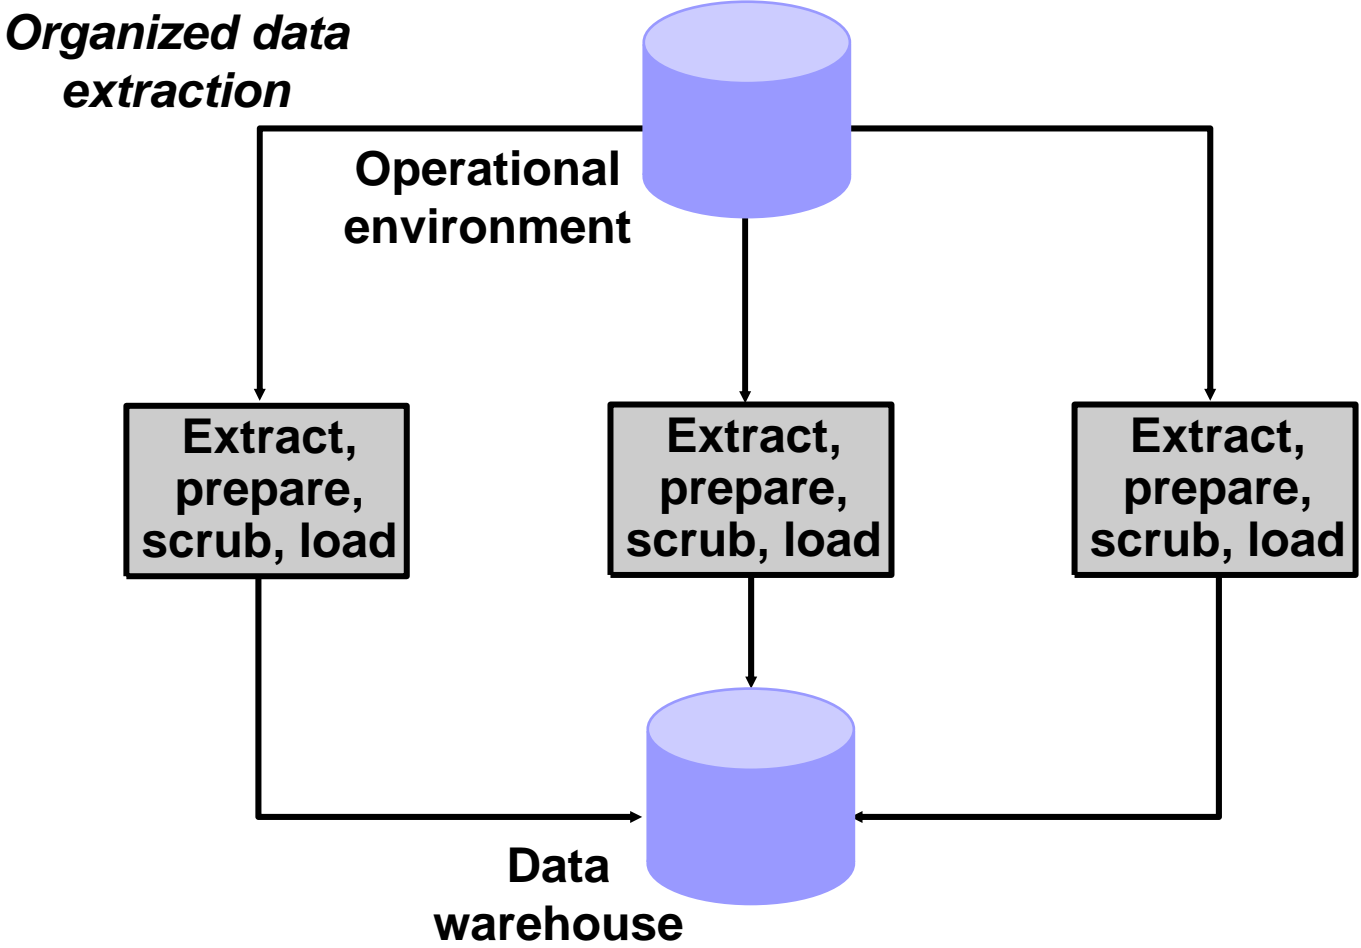
Note: In this lesson, the term *data warehouse* refers to any type of query-intensive database.

Data Warehouse Procedure

1. Prepare
2. Extract
3. Scrub
4. Load

1. Instance 1
2. Instance 2
3. Instance 3
4. Instance 4

Data Loading





[Apply for this job](#)

[Email to a friend](#)

Title: Oracle DBA Director
Skills: Please see description below for details

Date: 4-10-2005
Location: Sherman Oaks, CA
Area code: 91403

Tax term: FULLTIME
Pay rate: 100000-140000 yearly
Length: Perm

Position ID: ITSHC57927
Dice ID: kforcecx

Job description:

6 DBAs report to this position. You will lead team of DBAs, win loyalty and respect of your team, and be able to communicate and sell up to executive management. You must be able to problem solve and ask the right kinds of questions. Utilize your great expertise, personality, and communication skills.

You must have large server farm experience within a data warehouse environment. You must be knowledgeable with Oracle 10G installation and support. Oracle Applications 11i support and RAC experience is also required. This is a hands-on Director role responsible for 150 servers.

Travel required: none
Telecommute: no

[Apply for this job](#)

[Email to a friend](#)

To place this job in your Saved Jobs folder, [log in](#) or click [here](#) to register.

Oracle Job Ideas

1. **Attend Seminar on:**
 - **Writing Cover Letters**
 - **Writing Resumes**
 - **Conducting telephone interview**
 - **Closing that perfect Oracle Job**
 - **Be flexible (travel or relocation)**
2. **www.hitechblast.com**
3. **www.monster.com**
4. **www.computerjobs.com**
5. **www.dice.com**

Why should I choose Real Application Clusters?

1. The Future Role of the Oracle DBA

- RAC,
- Data Guard,
- Advanced RMAN and
- Grid Control Knowledge are essential.