

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

ORACLE

Typical Job Requirement

Our client is looking for someone with Strong Oracle DBA experience. They are implementing **Oracle Applications** and any experience in Oracle Apps is a huge plus.

This person must have **5-7 years Oracle DBA experience**. Must have **proven track record** in all areas of Oracle Database Administration.

Expertise in **RAC (Real Application Cluster)** is a must. Experience in implementing Oracle RAC on Windows and **Linux environments** is required.

Must have significant database performance tuning experience and must be able to run the **Oracle Enterprise Manager** and **Statspack** to show the metrics to the customer.

Looking for a **team player** with excellent communication and documentation skills.

State of the Oracle Database Industry

- Retire pension/gold watch
- BSCS, BSEE, MIS
- MCSE, Novell Certified
- OCP
- Oracle Masters
- Data Guard
- Linux and Oracle
- Real Application Clusters
- Oracle10g Grid
- Before 1980s
- 1980s
- Early 1990s
- Late 1990s
- 2000
- Now!
- Now!
- Now!
- Staying ahead.

Why should I choose Real Application Clusters?

Why Enhance My Career?

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

Grid Control

Manage multiple systems with one system:

- **Oracle Application Server 10g**
- **Oracle Database 10g**
- **Oracle Collaboration Suite 10g**
- **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

Hardware

- Nodes
- Interconnect
- Shared storage for your data

Operating Software

- Cluster Software - Cluster Manager
- Oracle Software
- Oracle Database and Instances

Nodes and Their Components

A node has the following main components:

- **CPUs**
- **Memory**
- **Interconnect**
- **Storage for OS, Cluster and Oracle Software**
- **Operating System Software - OS**
- **Cluster Software or Cluster Manager**

Interconnect

High Speed network connection for:

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

- **Infiniband**
- **Fast Ethernet**
- **Gigabit Ethernet**

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

- **Have been in use for a long time**
- **Bypasses the Operating System buffer cache**
- **Can be used in 9i or 10g**
- **Difficult to manage**
- **High performance**

RAW Device

On RAW Devices

- Data files
- Redo log files
- Control files
- Oracle Cluster Ready files
- Voting Disk

On Local File System

- Archive log files
- Oracle Home
- Alert ,Trace files
- External Tables
- Voting Disk

Cluster File System

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

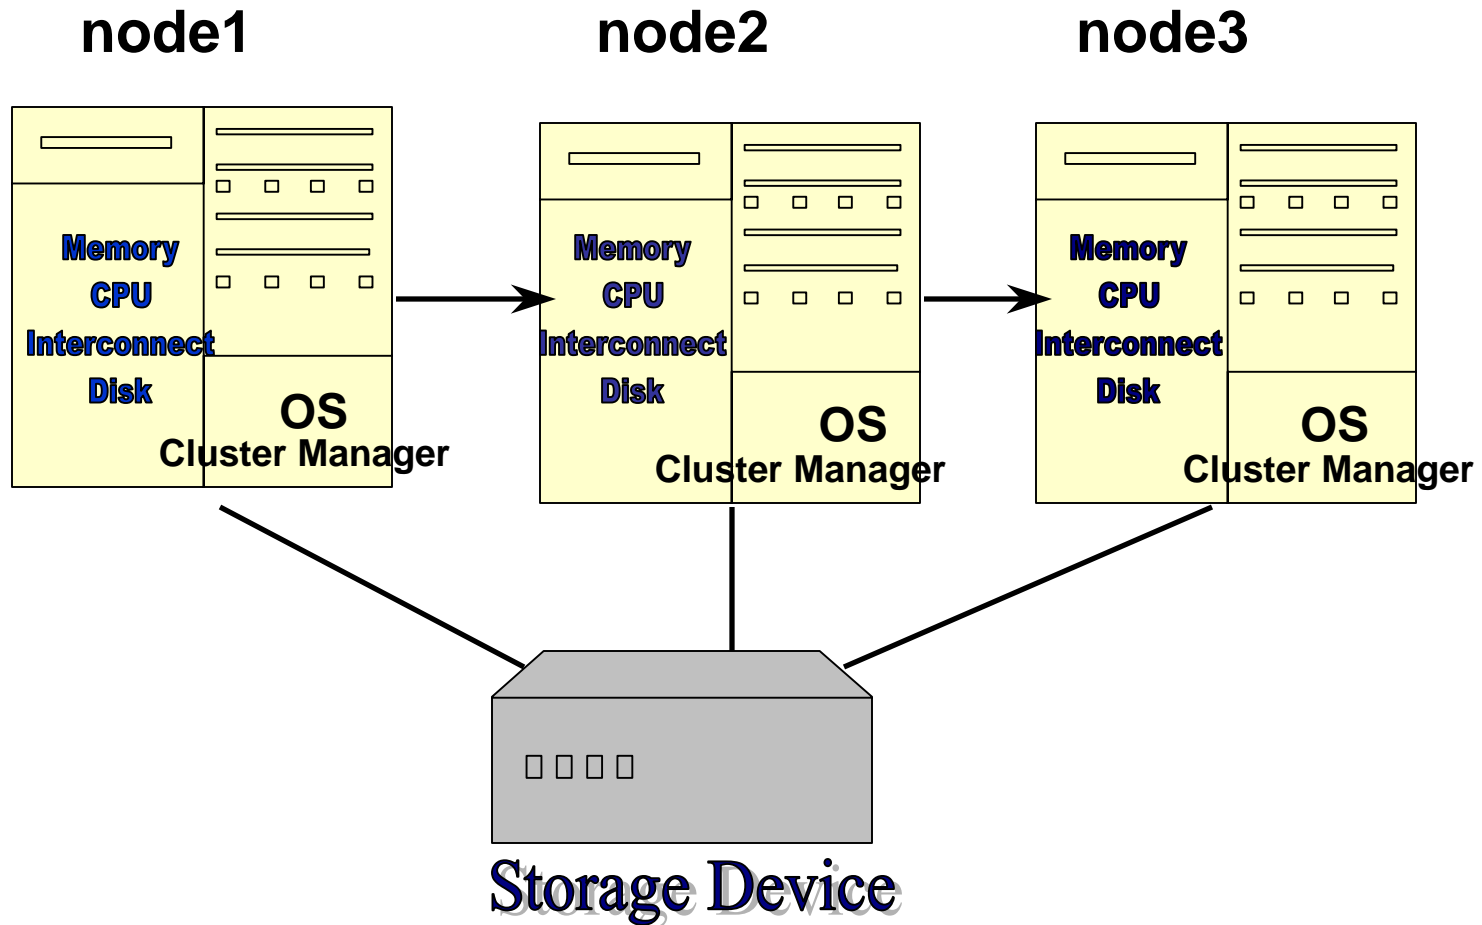
Cluster File System

- **Data files**
- **Redo log files**
- **Control files**
- **Oracle Cluster Ready files**
- **Archive log files**
- **Oracle Home**
- **Alert ,Trace files**
- **External Tables**

Automatic Storage Management

- **New star on the block for 10g**
- **Ease of Administration**
- **Should get almost the same performance of RAW**
- **Eliminates the need for cluster file system**
- **Eliminates the need for volume management**
- **No Support for:**
 - **External Tables**
 - **Transportable Tables**
 - **No Export/Import output file**

Nodes and Their Components



Operating System

- **Linux**
- **Windows**
- **Solaris**
- **HP**
- **IBM AIX**
- **Redhat and United Linux**
- **2000 and 2003**
- **Cluster 3.x(raw) and Veritas DBE(CFS)**
- **Itanium, Tru64 and TruCluster**
- **4.2(raw) and 5.1(both)**

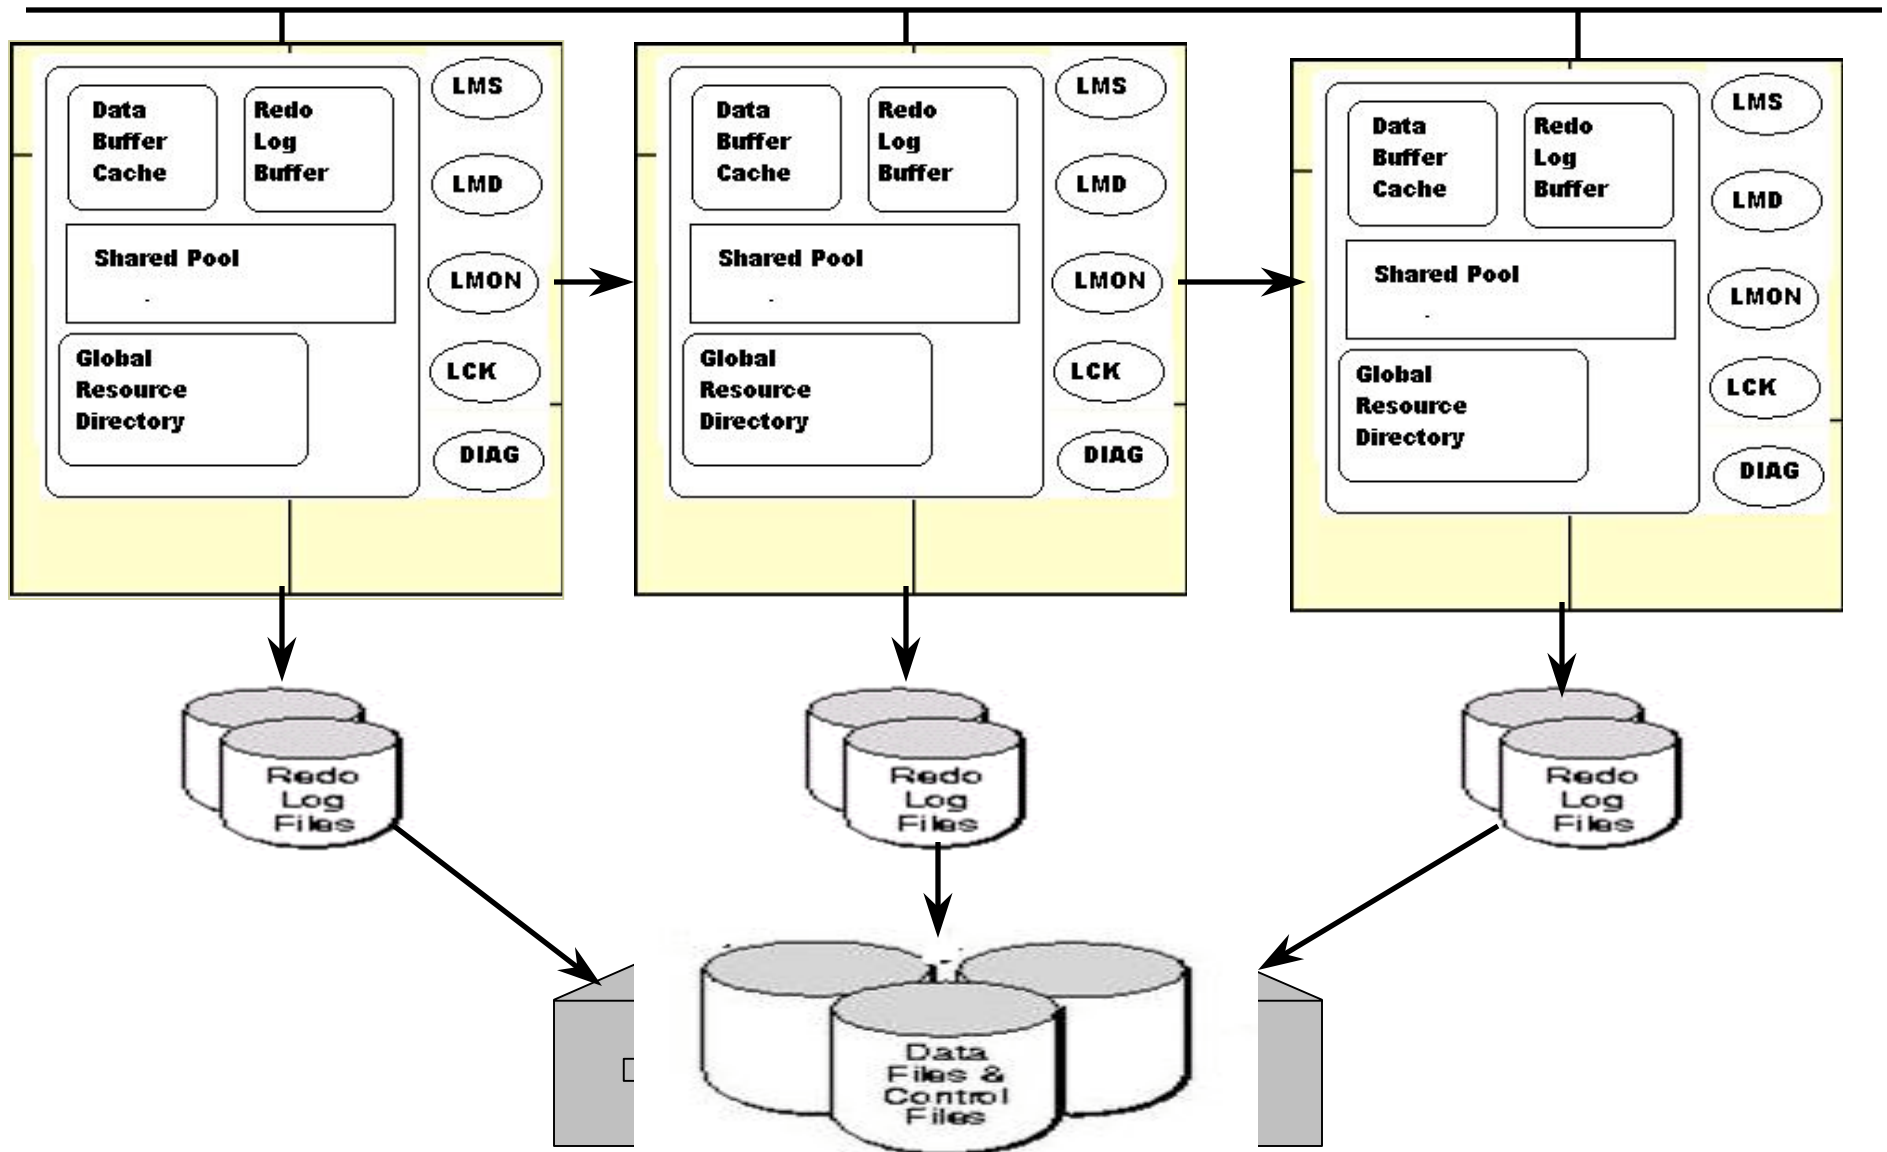
Check Oracle certification Matrix and vendor suppliers to choose right combination of RAC.

Cluster Software

- **Cluster Ready Services**
- **Sun 3.x or Veritas DBE**
- **HACMP**
- **Cluster services**
- **Windows and Linux**
- **Sun Solaris**
- **IBM**
- **HP**

Real Application Clusters-Specific Instance Processes

Interconnect Communication



ORACLE

RAC Instance

Global Resource Directory – Tracking status of oracle blocks.

Additional Background processes:

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

Global Resource Directory

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

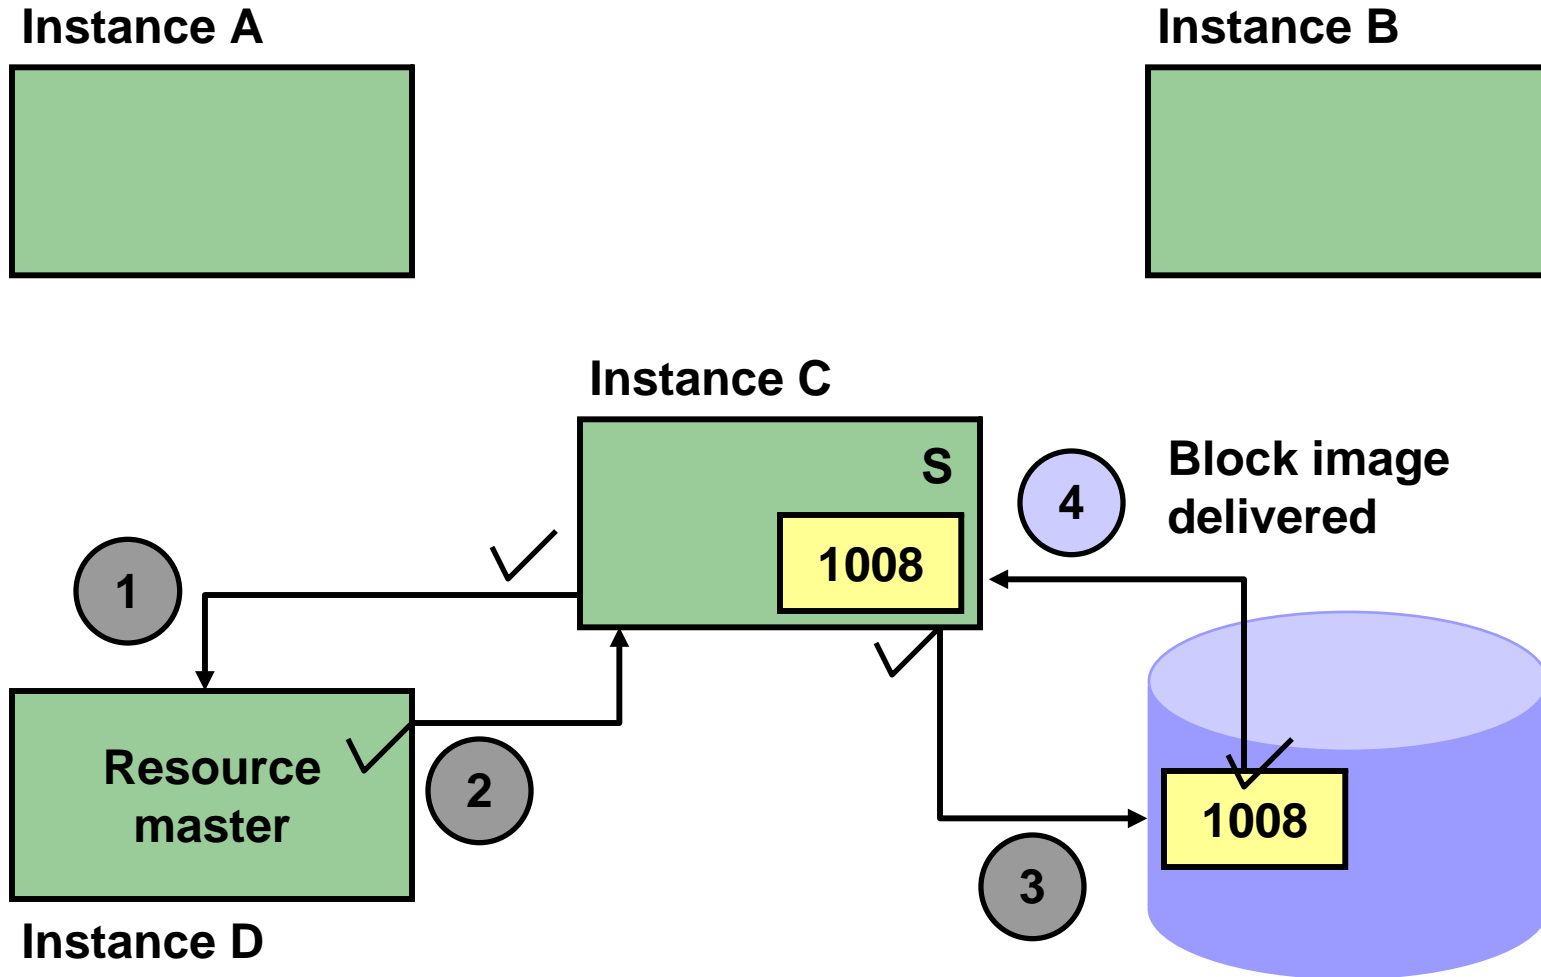
Cache Fusion

- 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.

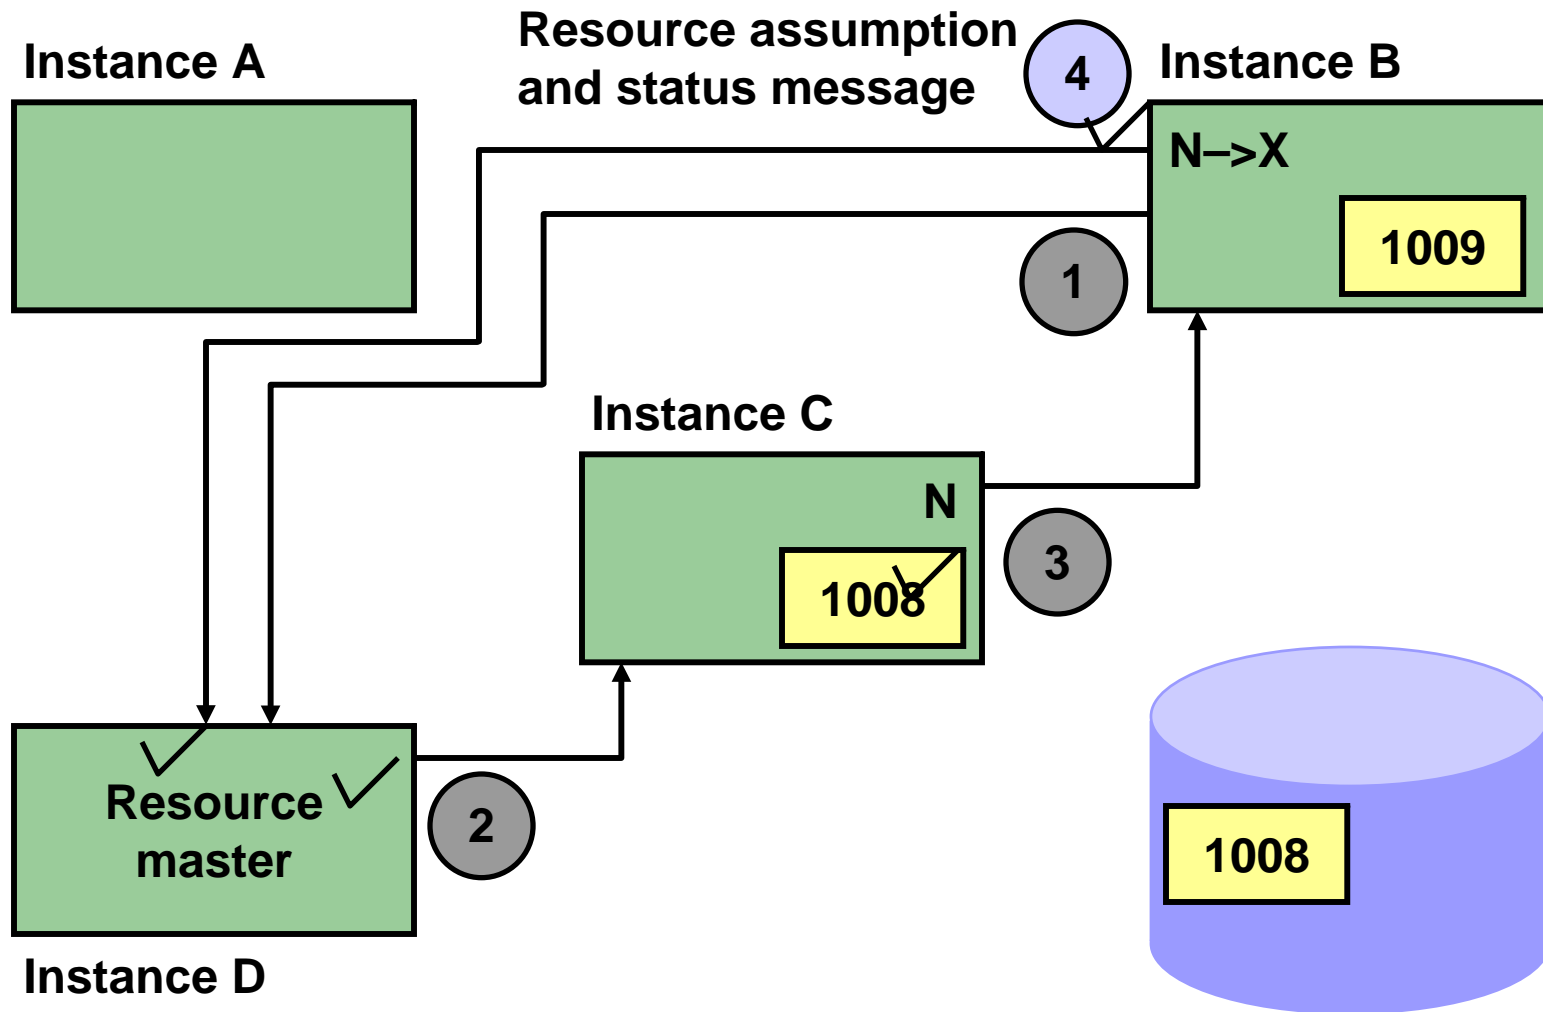
Cache Fusion Scenarios

- **Read with no Transfer**
- **Read/Write**
- **Write/Write**
- **Write/Read**

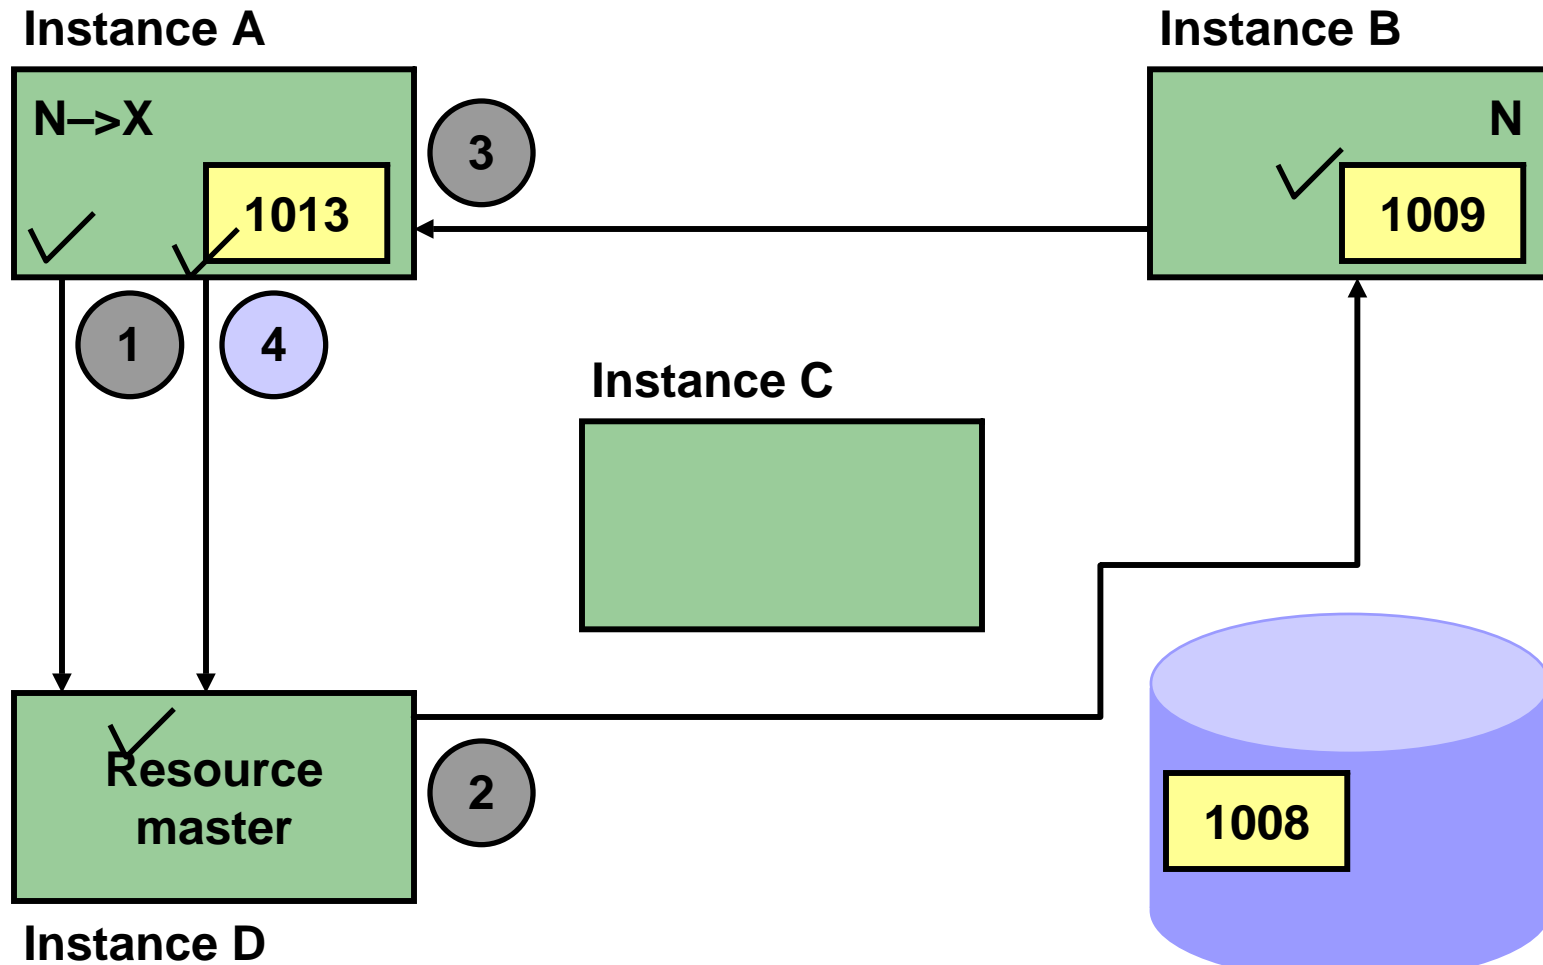
Example 1: Read with No Transfer



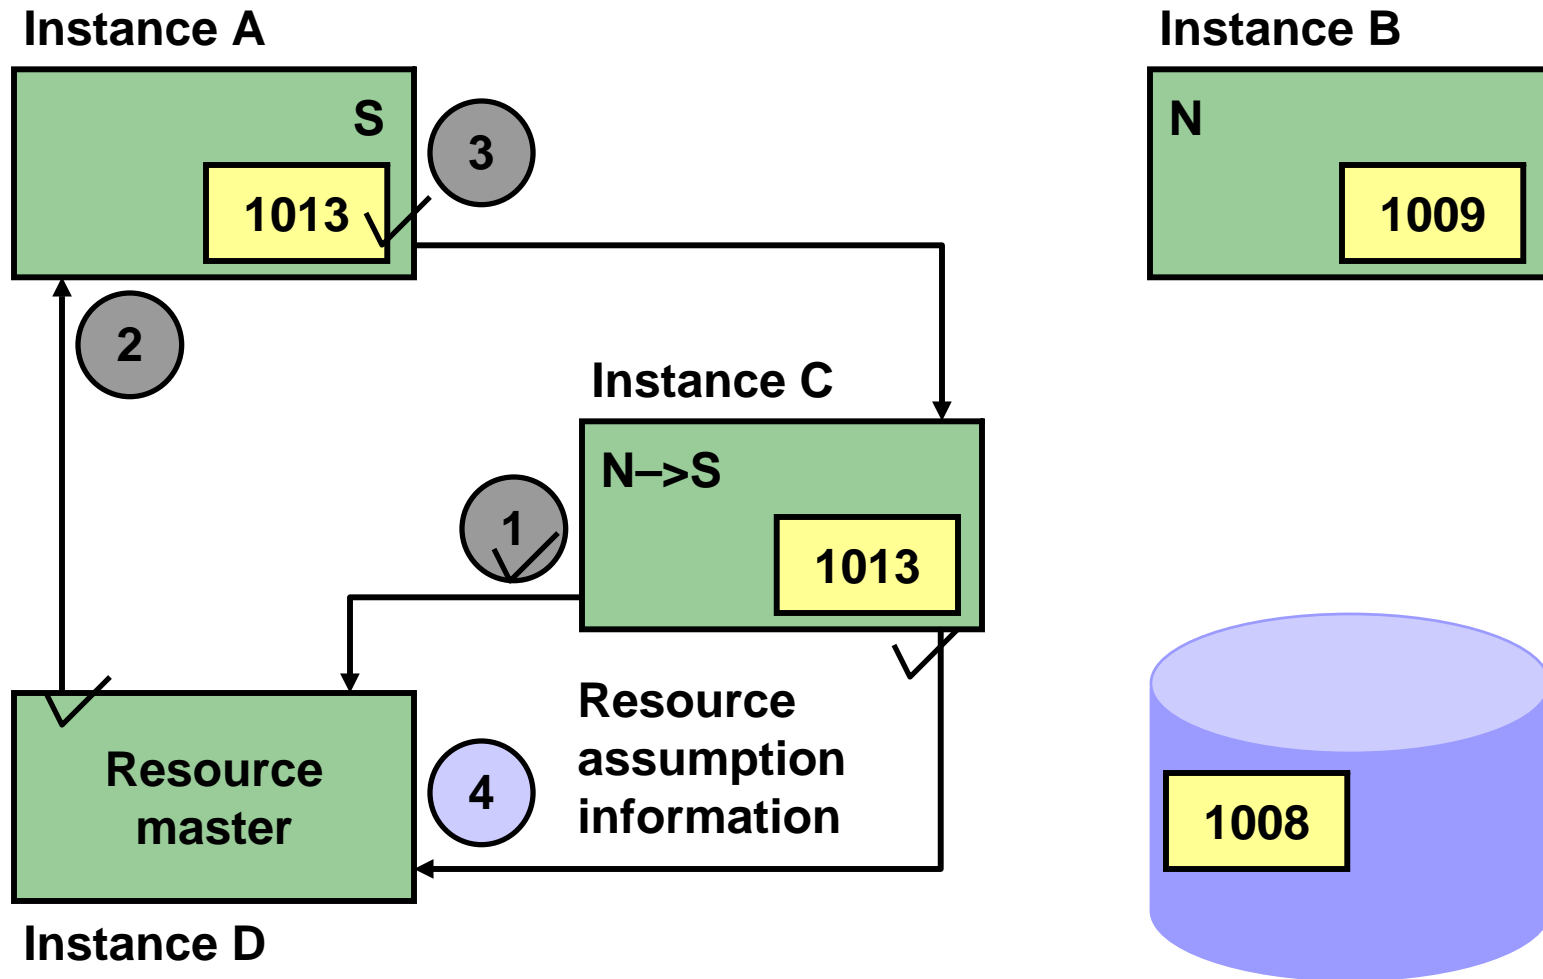
Example 2: Read to Write Transfer



Example 3: Write to Write Transfer

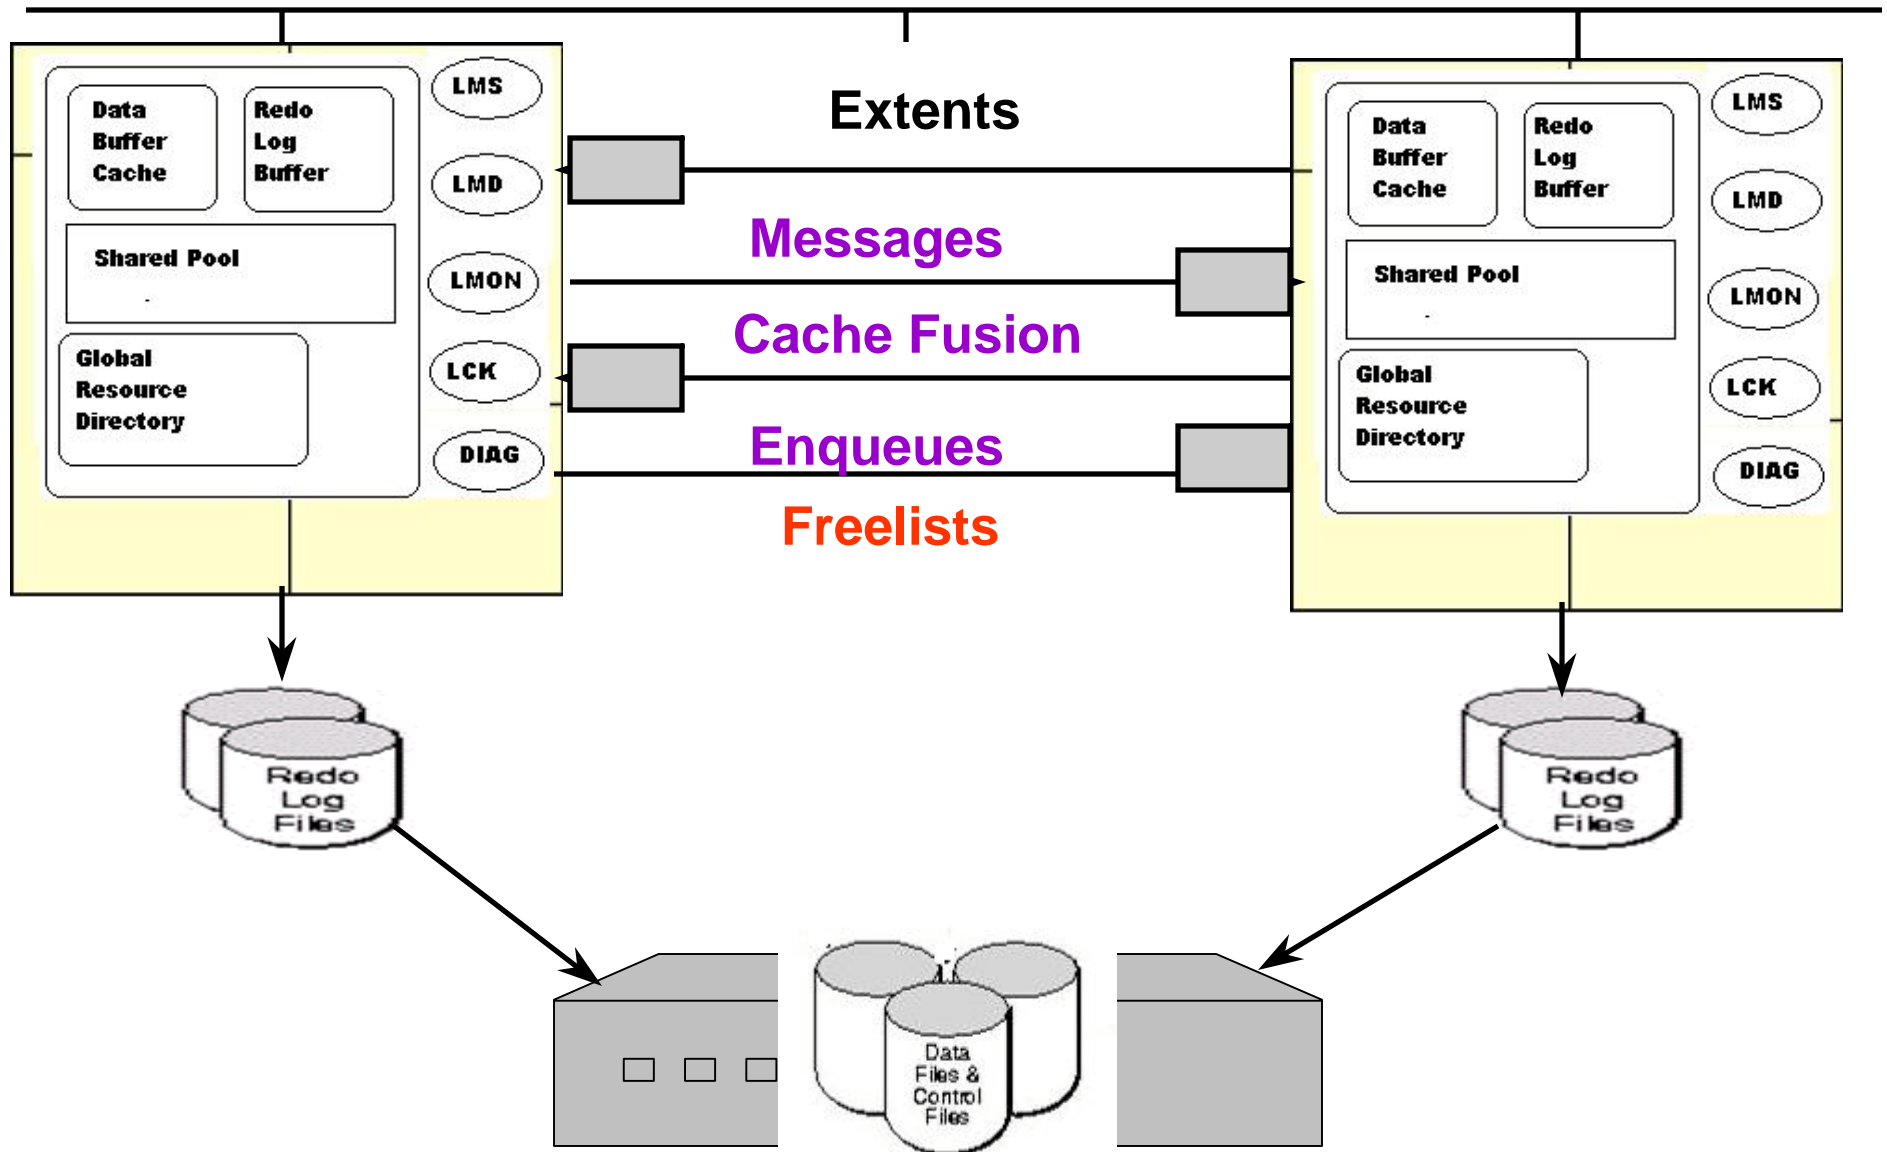


Example 4: Write to Read Transfer



Real Application Clusters-Specific Instance Processes

Interconnect Communication

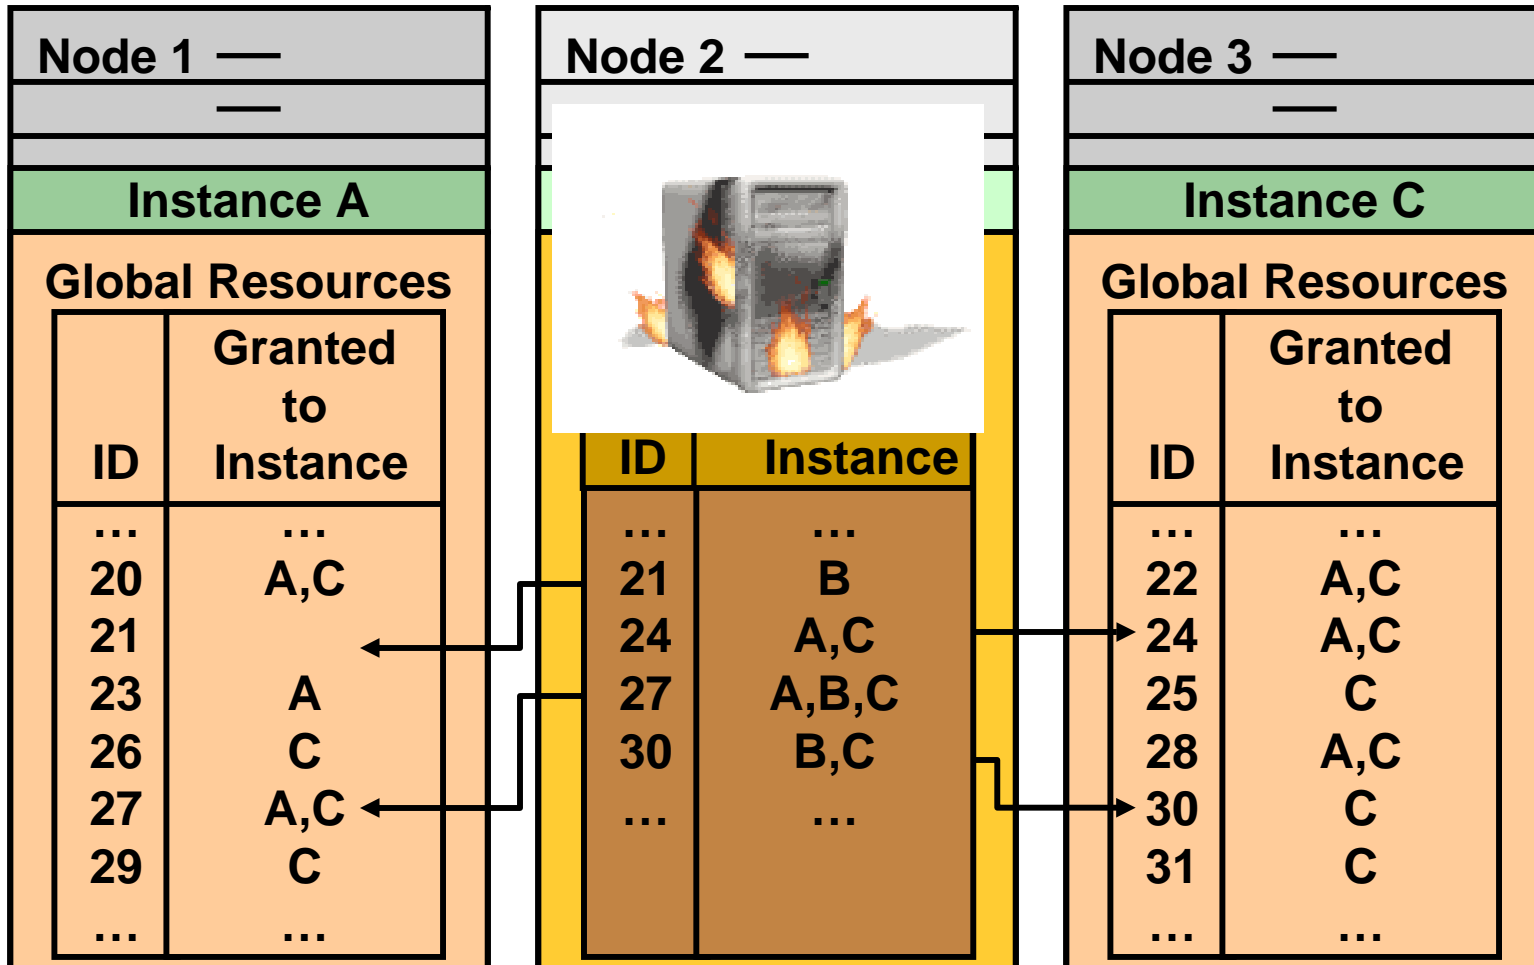


ORACLE

Cluster Reorganization: Example

Node 1 —		Node 2 —		Node 3 —	
—		—		—	
Instance A		Instance B		Instance C	
Global Resources		Global Resources		Global Resources	
ID	Granted to Instance	ID	Granted to Instance	ID	Granted to Instance
...
20	A,B,C	21	B	22	A,C
23	A,B	24	A,C	25	C
26	C	27	A,B,C	28	A,B,C
29	B,C	30	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

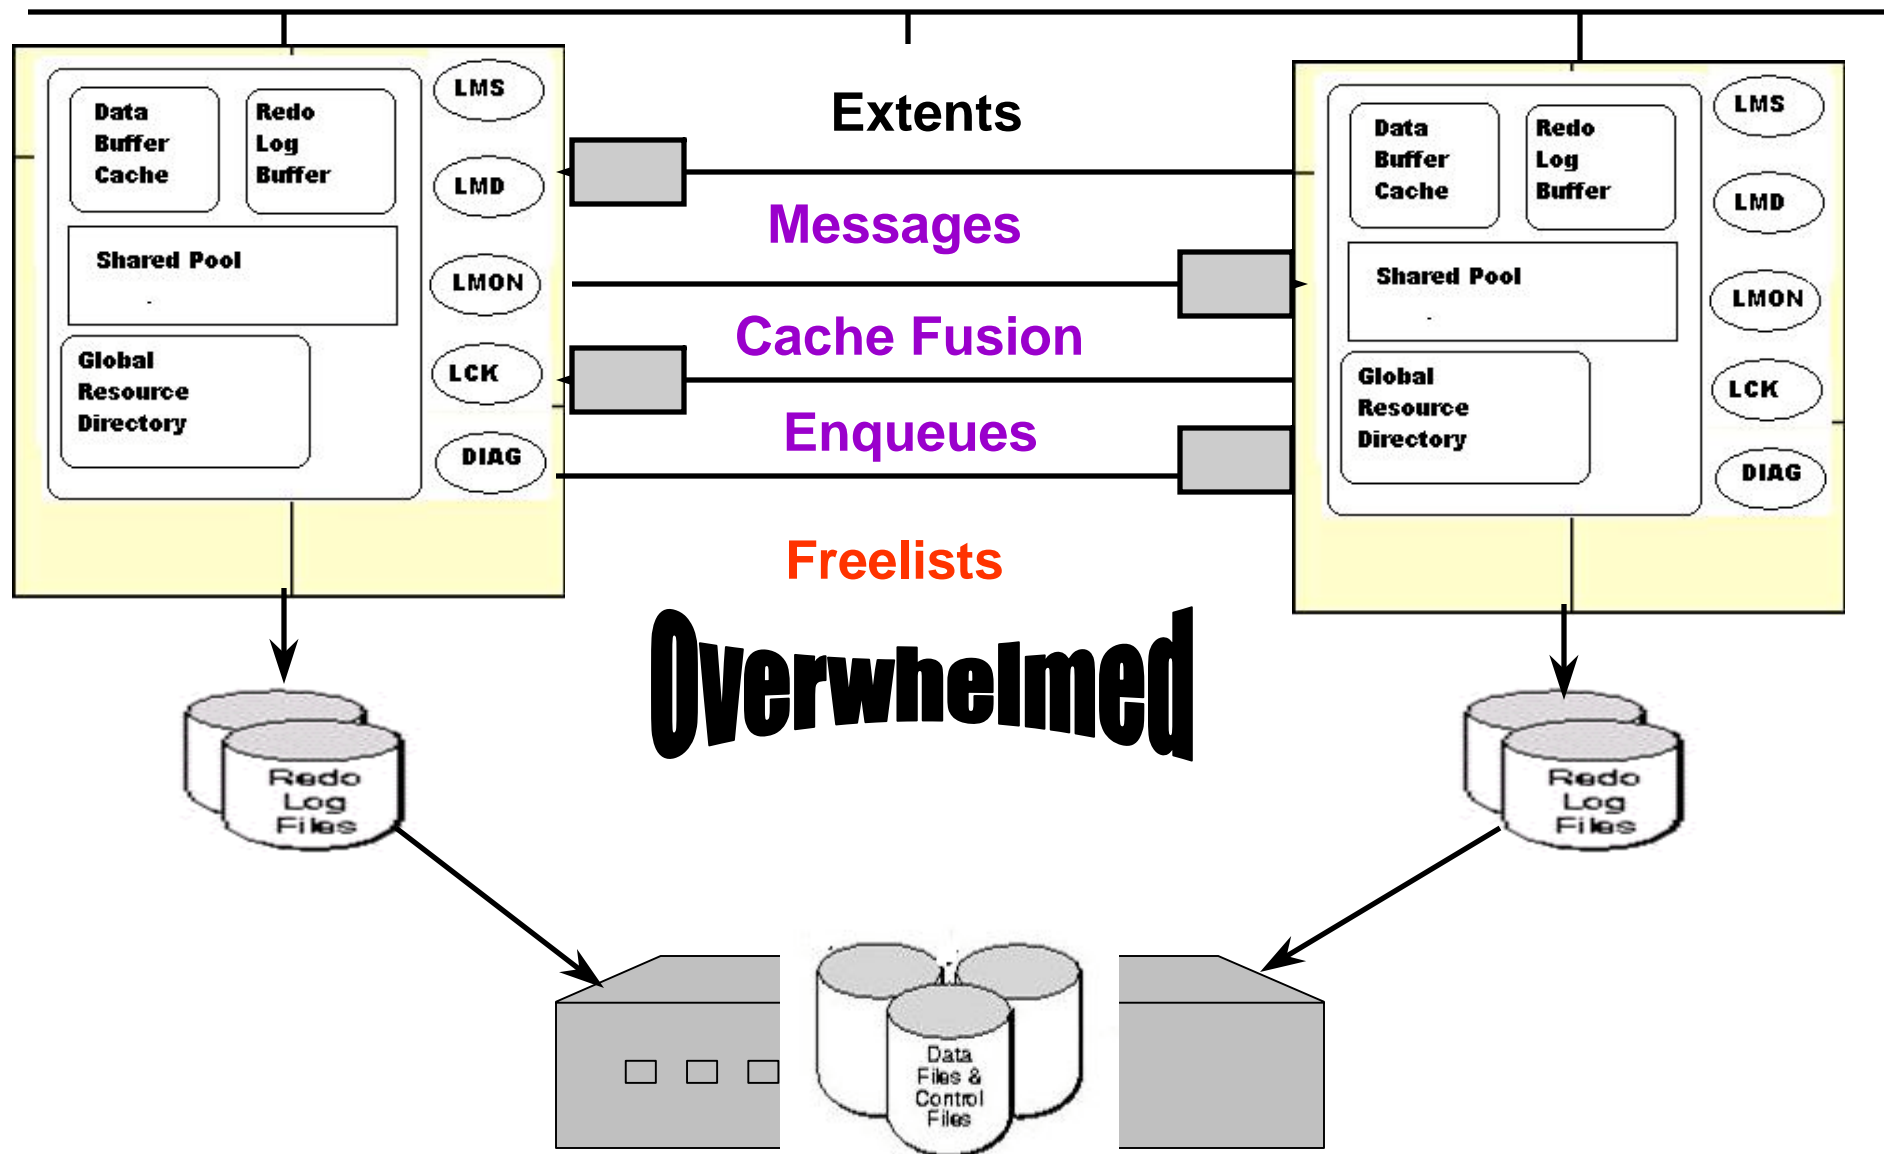
Freelist Management

Distribute inserts for B-tree

Autoextend for datafile increase on RAW devices


Real Application Clusters-Specific Instance Processes

Interconnect Communication

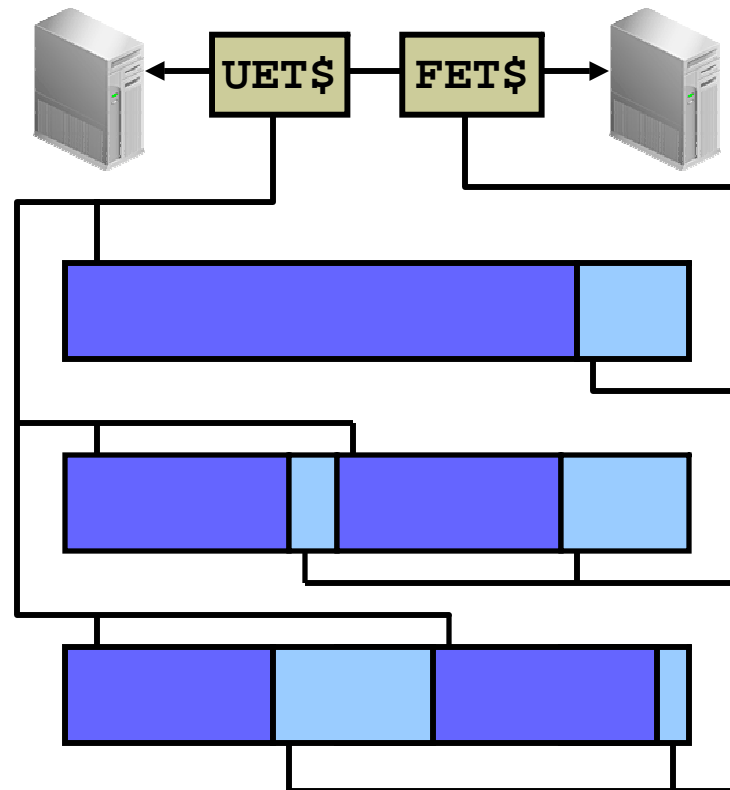


ORACLE

Extent Management Options



Local management is controlled by bitmaps in the data files.



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

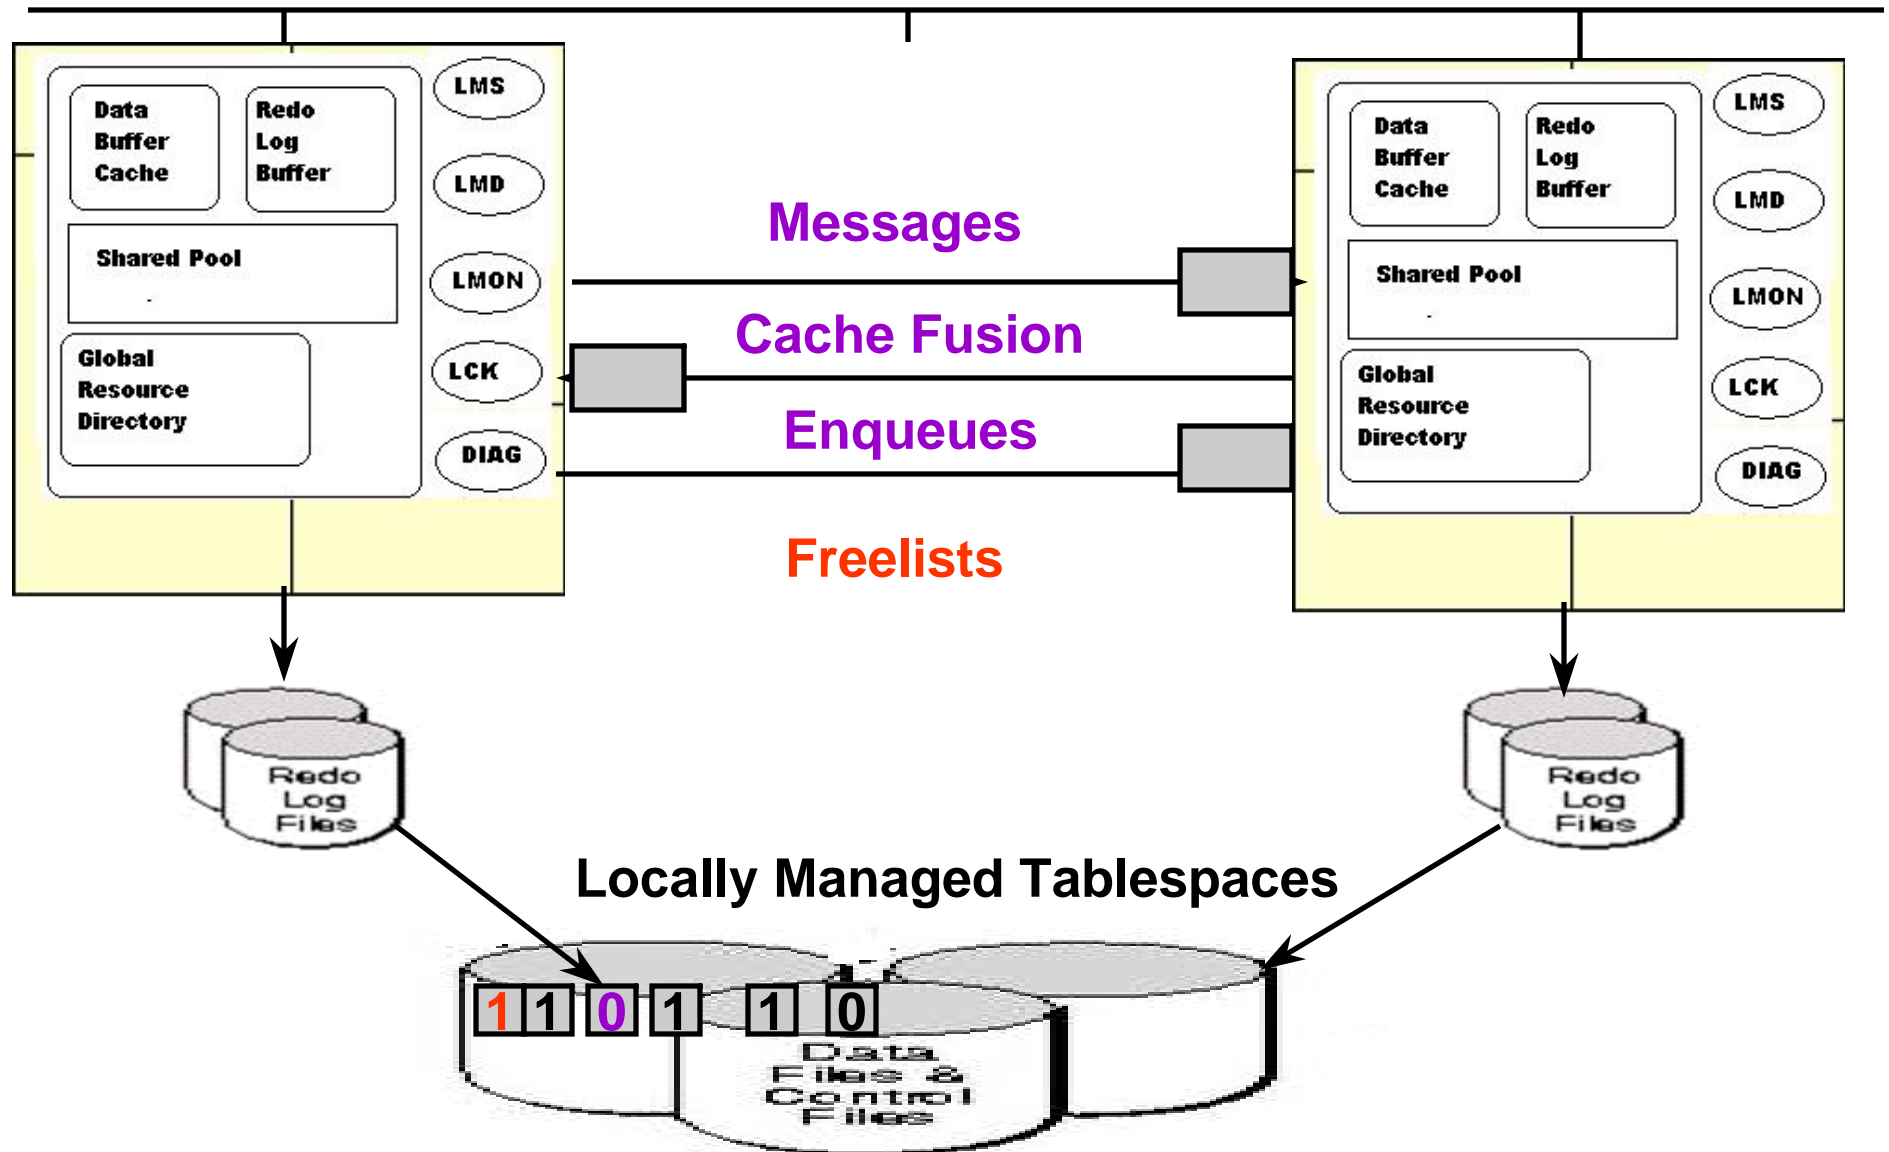
Locally Managed Tablespaces

Locally managed tablespaces:

- Are recommended by Oracle Corporation
- 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



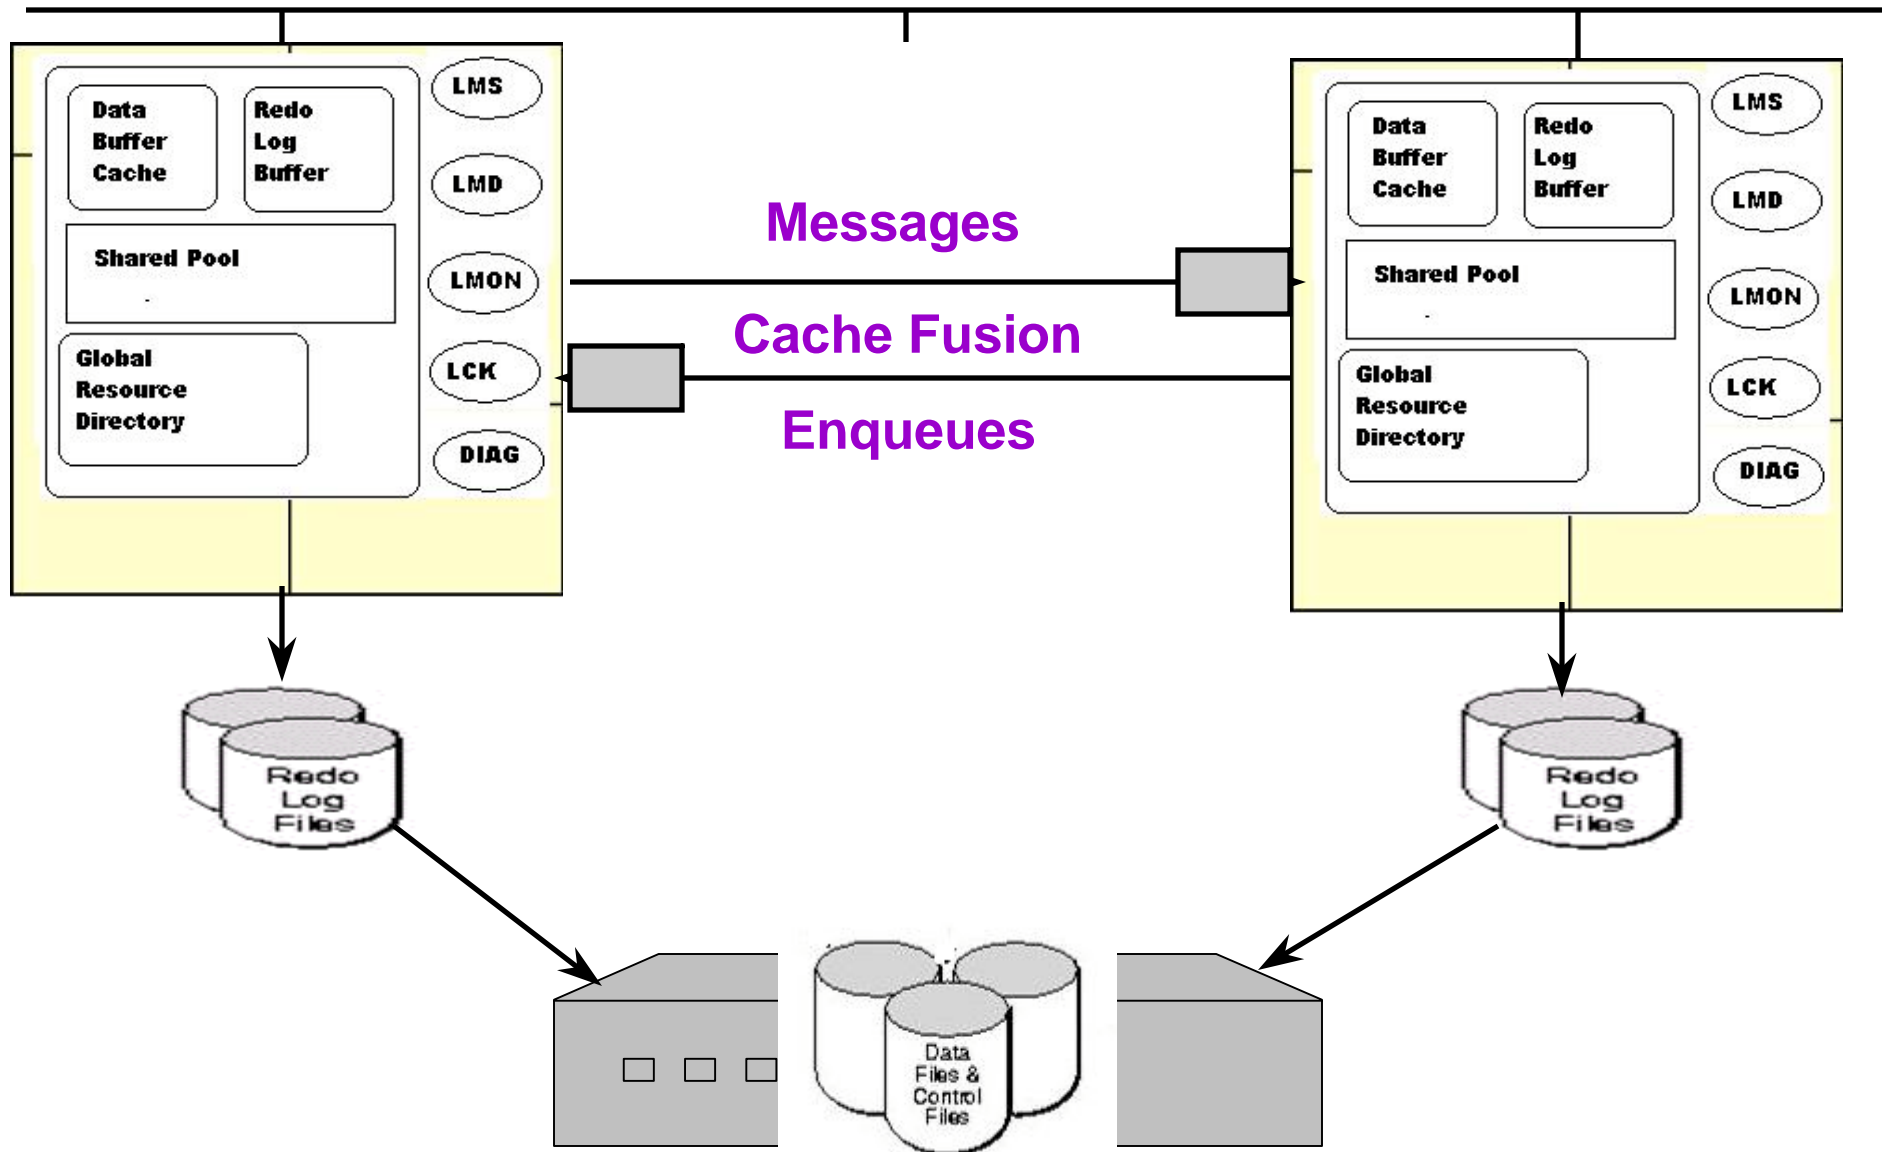
ORACLE

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

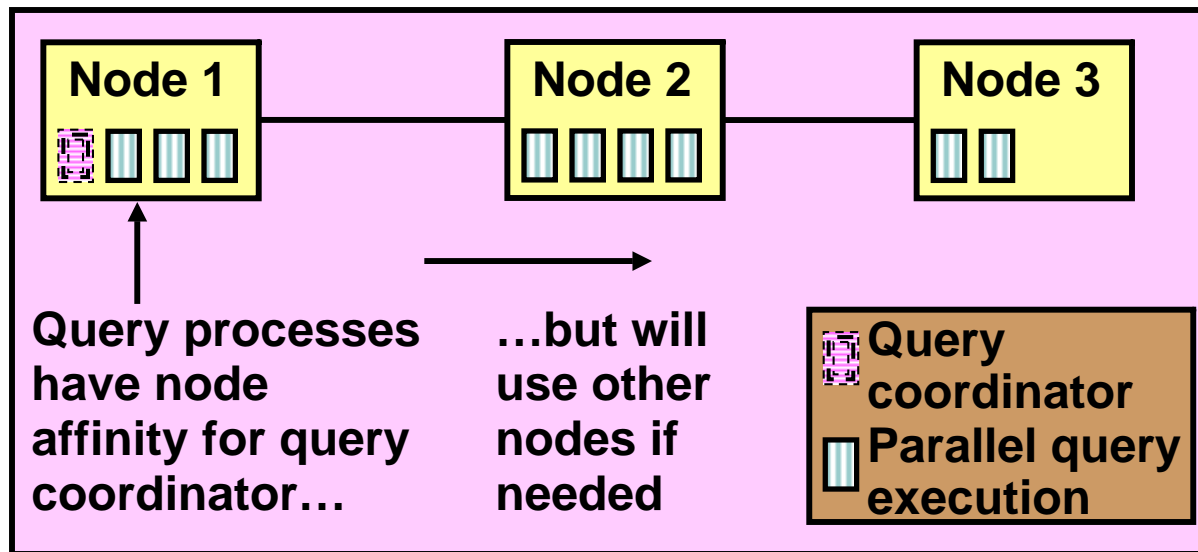



ORACLE

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

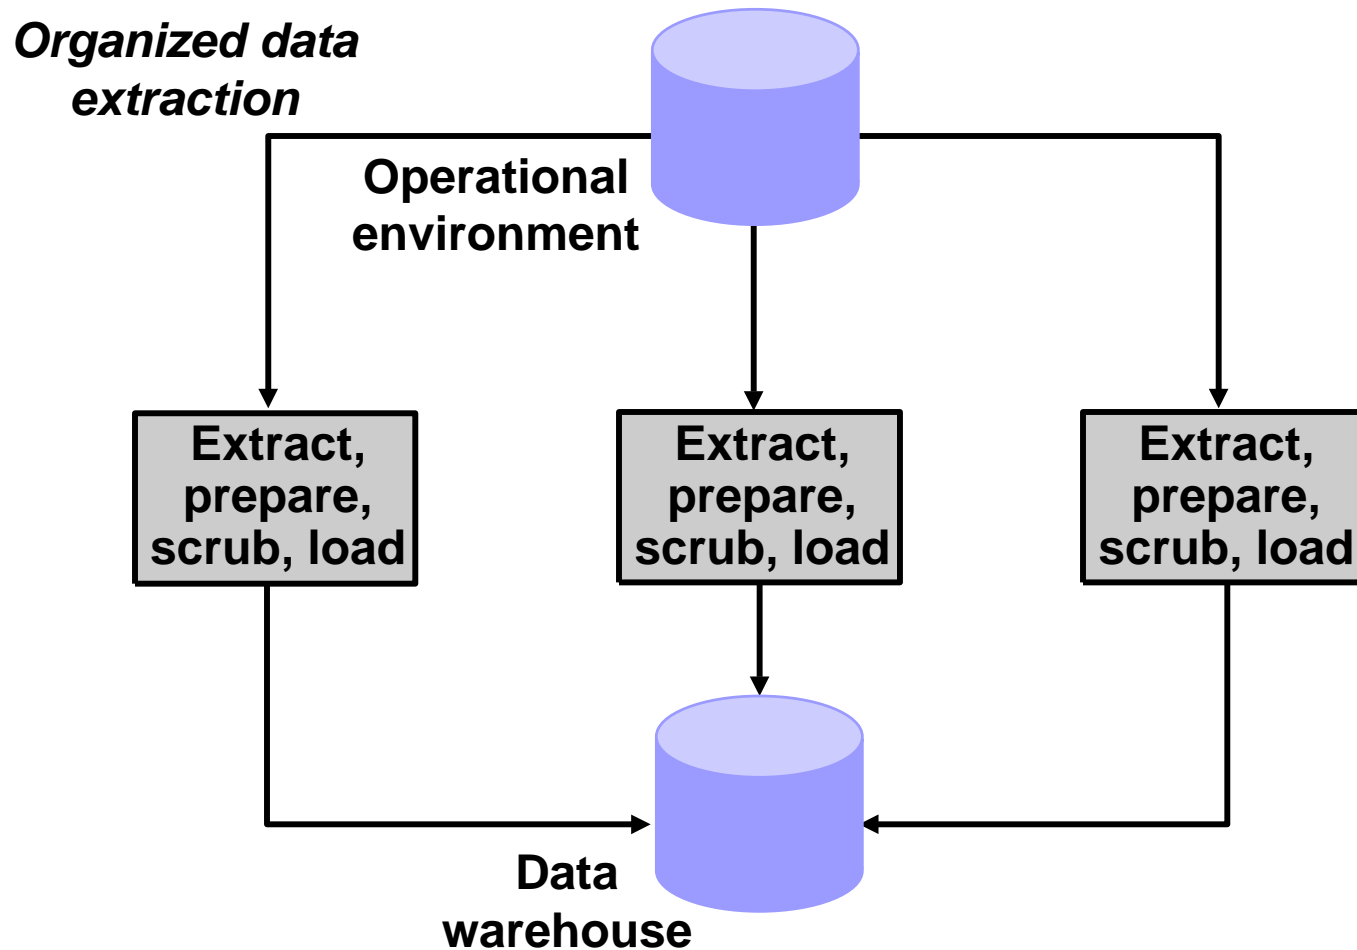
- **Query-intensive databases include:**
 - Online analytic processing (OLAP) servers
 - Decision support systems (DSS)
 - Data warehouses
- **Such databases are characterized by:**
 - Large amounts of data
 - Extensive query access
 - Scheduled batch loads to refresh or replace data
- **High amounts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

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

Data Warehouse Procedure

- **Prepare**
 - **Extract**
 - **Scrub**
 - **Load**
- **Instance 1**
 - **Instance 2**
 - **Instance 3**
 - **Instance 4**

Data Loading



Typical Job Requirement

Senior Oracle DBA (Guru):

Senior Oracle DBA with the following experience:

- * Informix to Oracle migration experience
- * **Oracle RAC implementations**
- * **Data Guard implementations**
- * Oracle fail-safe (MS cluster required)
- * Oracle Partitioning
- * **Linux/Unix, Windows**
- * Peoplesoft Apps experience
- * Peoplesoft is currently running on the Informix DB and the client needs the application migrated to Oracle DB.
- * **Good communication** and documentation skills are a must.

Start Date: ASAP

Location: Atlanta, GA

If interested, please provide below information.

- 1. Availability**
- 2. Hourly rate**
- 3. Best time and telephone numbers to reach**
- 4. Could you commit for long term project.**
- 5. Email Formatted resume.**

ORACLE

Oracle Job Ideas

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

Why should I choose Real Application Clusters?

- **The Future Role of the Oracle DBA**
 - **RAC,**
 - **Data Guard,**
 - **Advanced RMAN and**
 - **Grid Control Knowledge are essential.**