



ArcSDE for Oracle Administration

Presented by Mark Harris and
Jill Silvertand



ESRI®

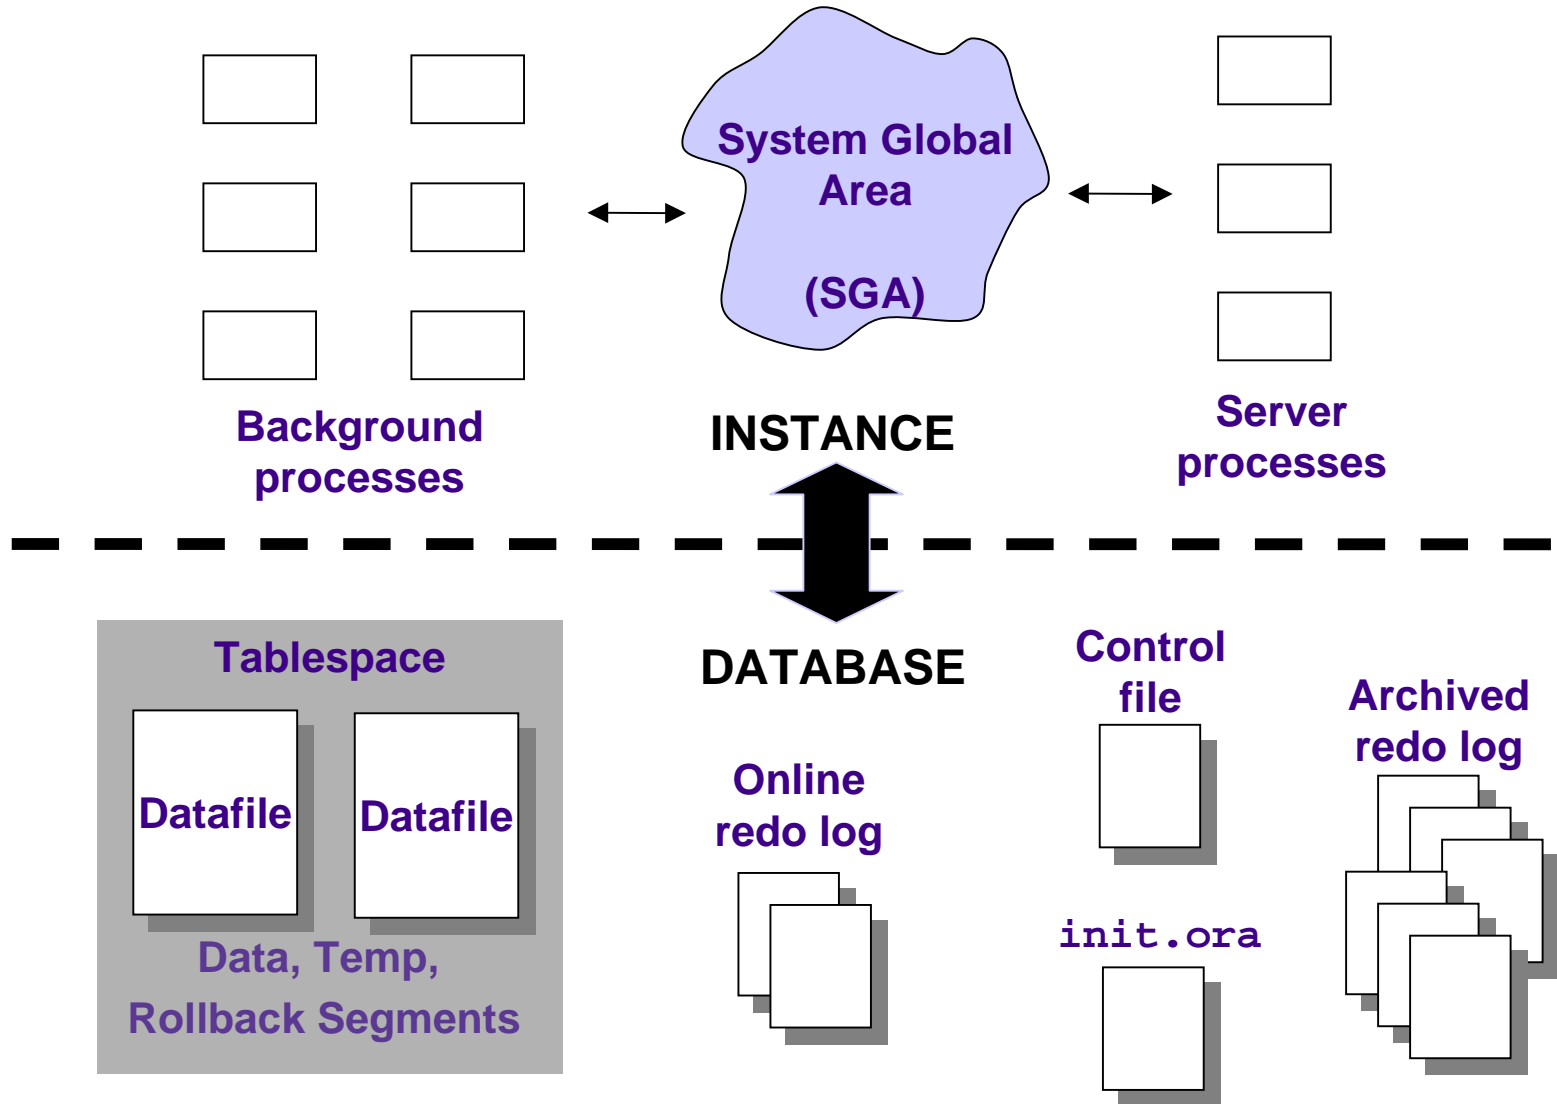
Outline

- **Basic Oracle Architecture**
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index




ESRI

Basic Oracle Architecture



ESRI

Oracle Instance




- **Dedicated server processes**
 - Created when a user connects to the Oracle server
 - Interaction between the Oracle server and the client
 - Uses an area of memory called the Program Global Area (PGA)



ESRI®

Oracle Instance

- In a multi-threaded environment, several user processes share one server process



ESRI

Background Processes

- **SMON** - System monitor; performs instance recovery at startup
- **PMON** - Process monitor; performs process recovery when a user process fails
- **DBWn** - Database writer; writes changed data to the data files



ESRI

Background Processes

- **LGWR** - Log Writer; records data changes in the redo log files
- **CKPT** - Checkpoint; updates database status information
- **ARCn** - Archiver; copies online redo log file to archive directory



ESRI

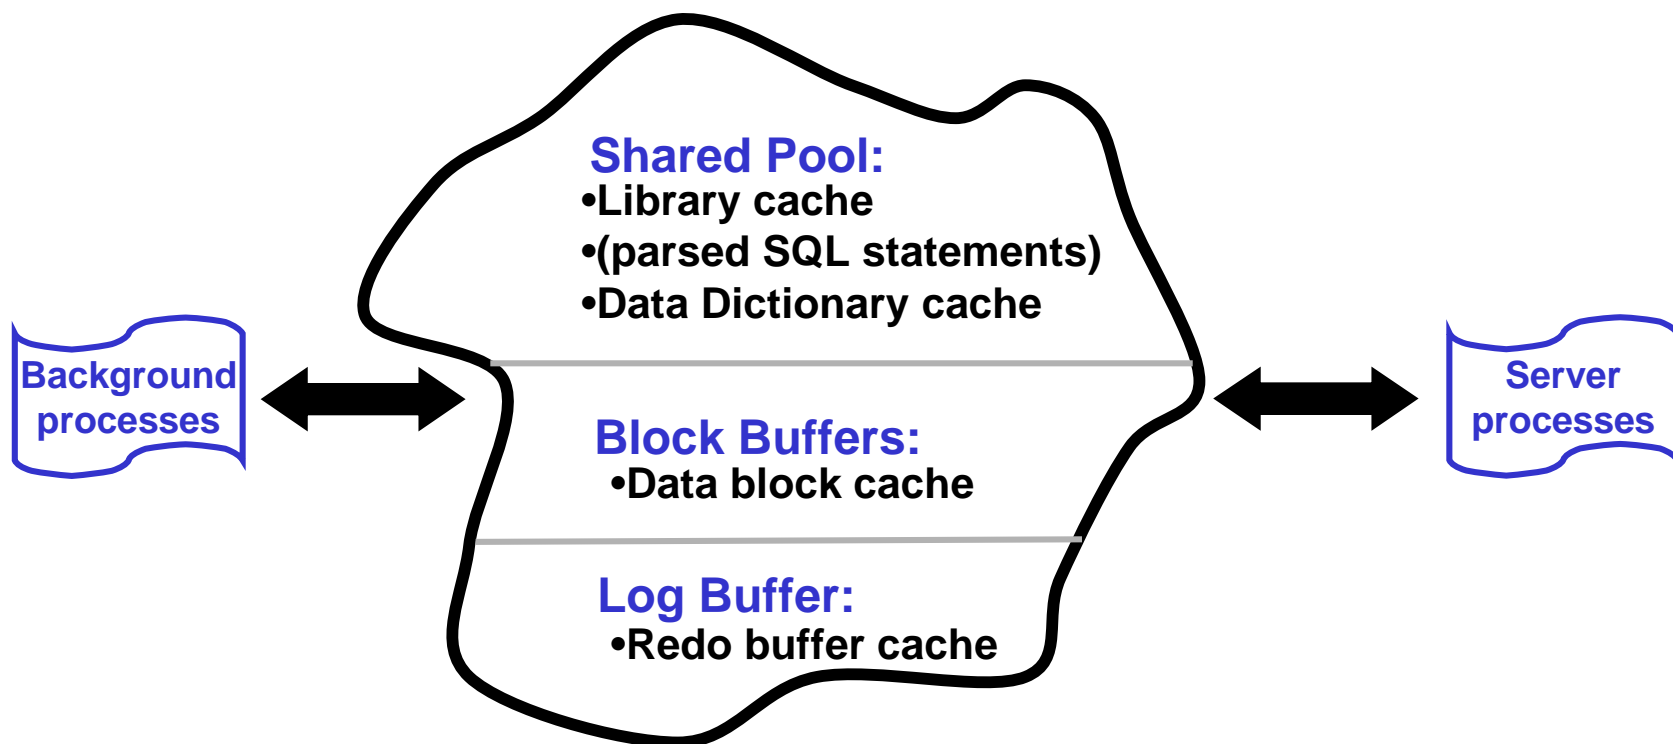
Outline

- Basic Oracle Architecture
- **The SGA**
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

System Global Area (SGA)



ESRI

Sizing the SGA

- **Consider the active processes, including:**
 - **Operating System**
 - **Oracle background processes**
 - **Oracle server processes**
 - **ArcSDE giomgr and gsrvr processes**
 - **Other applications**



ESRI

Shared Pool

- Set with the `init.ora` parameter `shared_pool_size`
- The number of reloads should be kept to less than 1%.

```
select sum(pins) / (sum(pins) + sum(reloads)) * 100  
from v$librarycache;
```



ESRI

Log Buffer

- **Set by the init.ora parameter log_buffer**
- **Size set as a multiple of your O/S block size**
- **Set to 512 KB or 128 KB * number CPU**



ESRI

Log buffer

- Redo log space requests should be 0

```
select value  
from v$sysstat  
where name = 'redo log space requests';
```



ESRI

Block Buffer

- Set by the init.ora parameter `db_block_buffer`
- Should have a buffer cache hit ratio of 95% or better

Buffer cache hit ratio=

$$\left(1 - \left(\frac{\text{physical reads}}{\text{db_block_gets} + \text{consistent_gets}} \right) \right) * 100$$

SELECT name, value FROM v\$sysstat

WHERE name IN ('db block gets', 'consistent gets', 'physical reads');



ESRI


Outline

- Basic Oracle Architecture
- The SGA
- **The Oracle Database**
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index




ESRI®

Oracle Database




Data Files


System
Tables
Indexes
Rollback Segments
Temporary Segments




Redo Log Files



Control Files




Init.ora
File



Archived Redo
Log Files



ESRI




init.ora

init.ora

- Stores configuration parameters
- Only read at instance startup
 - Re-start server for new setting to take effect
 - Some parameters set dynamically with **ALTER SESSION** or **ALTER SYSTEM**
 - Settings accessible from the **V\$PARAMETER**



ESRI



init.ora

init.ora

- **Default location**

- **UNIX:**

- `$ORACLE_HOME/dbs/init<sid>.ora`

- **NT:**

- `%ORACLE_HOME%\database\init<sid>.ora`




ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - **The control files**
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI




control files

Control files

- **Store information about the database**
 - Database name
 - Redo log member and datafile locations
 - Internal system information
 - Current log sequence number
 - Checkpoint information



ESRI®



control files

Control files


- **Mirror on a different device to protect from media failure**
 - Difficult to recover if you lose all of them
 - Control files are referenced in `init.ora`

```
control_files =
```

```
( "d:\Oracle\oradata\sde8\control01ctl",  
  "e:\Oracle\oradata\sde8\control02ctl" )
```



ESRI



control files

Control files

- **Backup immediately after structural changes**
- **Allocate 10MB of space if using RAW device**
- **Set CONTROL_FILE_RECORD_KEEP_TIME parameter to control reuse**




ESRI®

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - **The online redo log files**
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®




redo log files

The redo log files

- **Stores changes to database contents**
 - Allows Oracle to redo and rollback transactions in the event of a recovery
- **Must have at least two groups**
- **Checkpoints at log switch**
- **Each group may have multiple members (mirrors)**



ESRI



redo log files

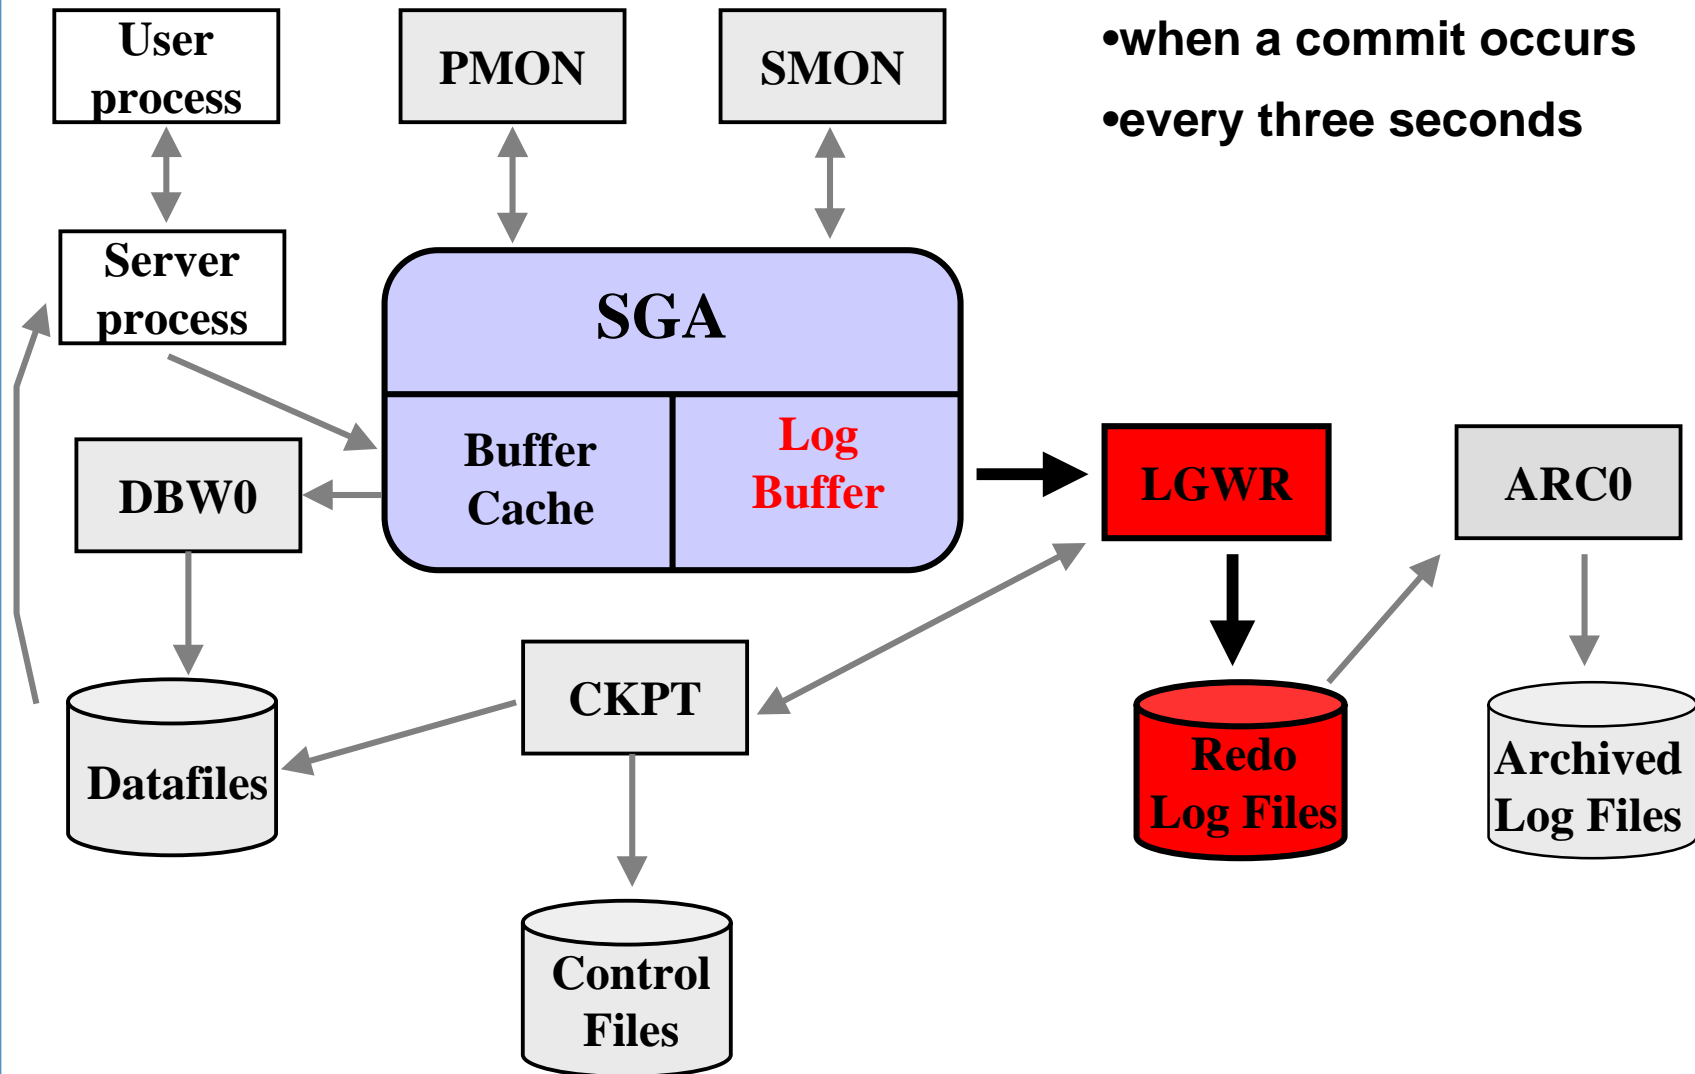
The redo log

- **NOARCHIVELOG** mode
 - Stores changes between checkpoints in online redo log
 - Enables recovery from instance failure
- **ARCHIVELOG** mode
 - Copies online redo log to archived redo log on switch
 - Enables recovery from media failure



ESRI

Log flush



Log flush occurs:

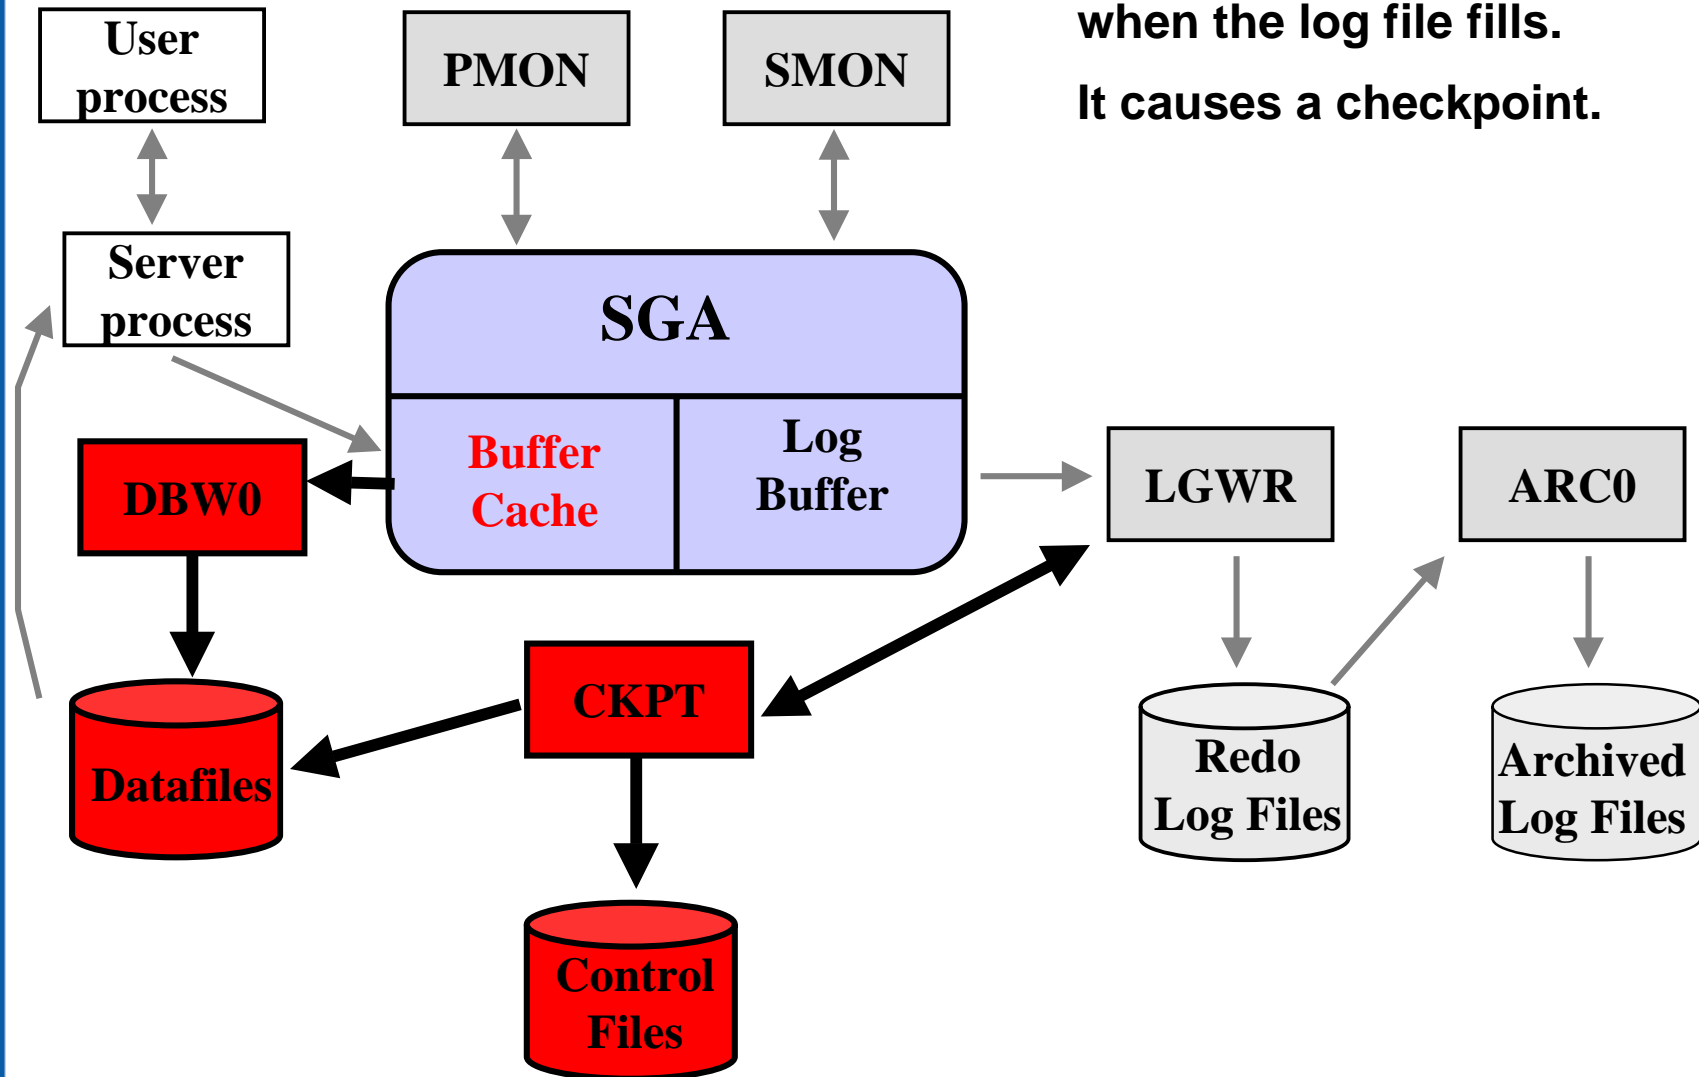
- when 1/3 full
- when a commit occurs
- every three seconds



ESRI

Log switch

Log switch occurs when the log file fills. It causes a checkpoint.



The archived redo log

- **ARCHIVELOG mode**
 - Copies online redo log to archived redo log after a log switch
 - Enables recovery from media failure
- The **ARCn** process is used to copy the redo log to the archive location



ESRI

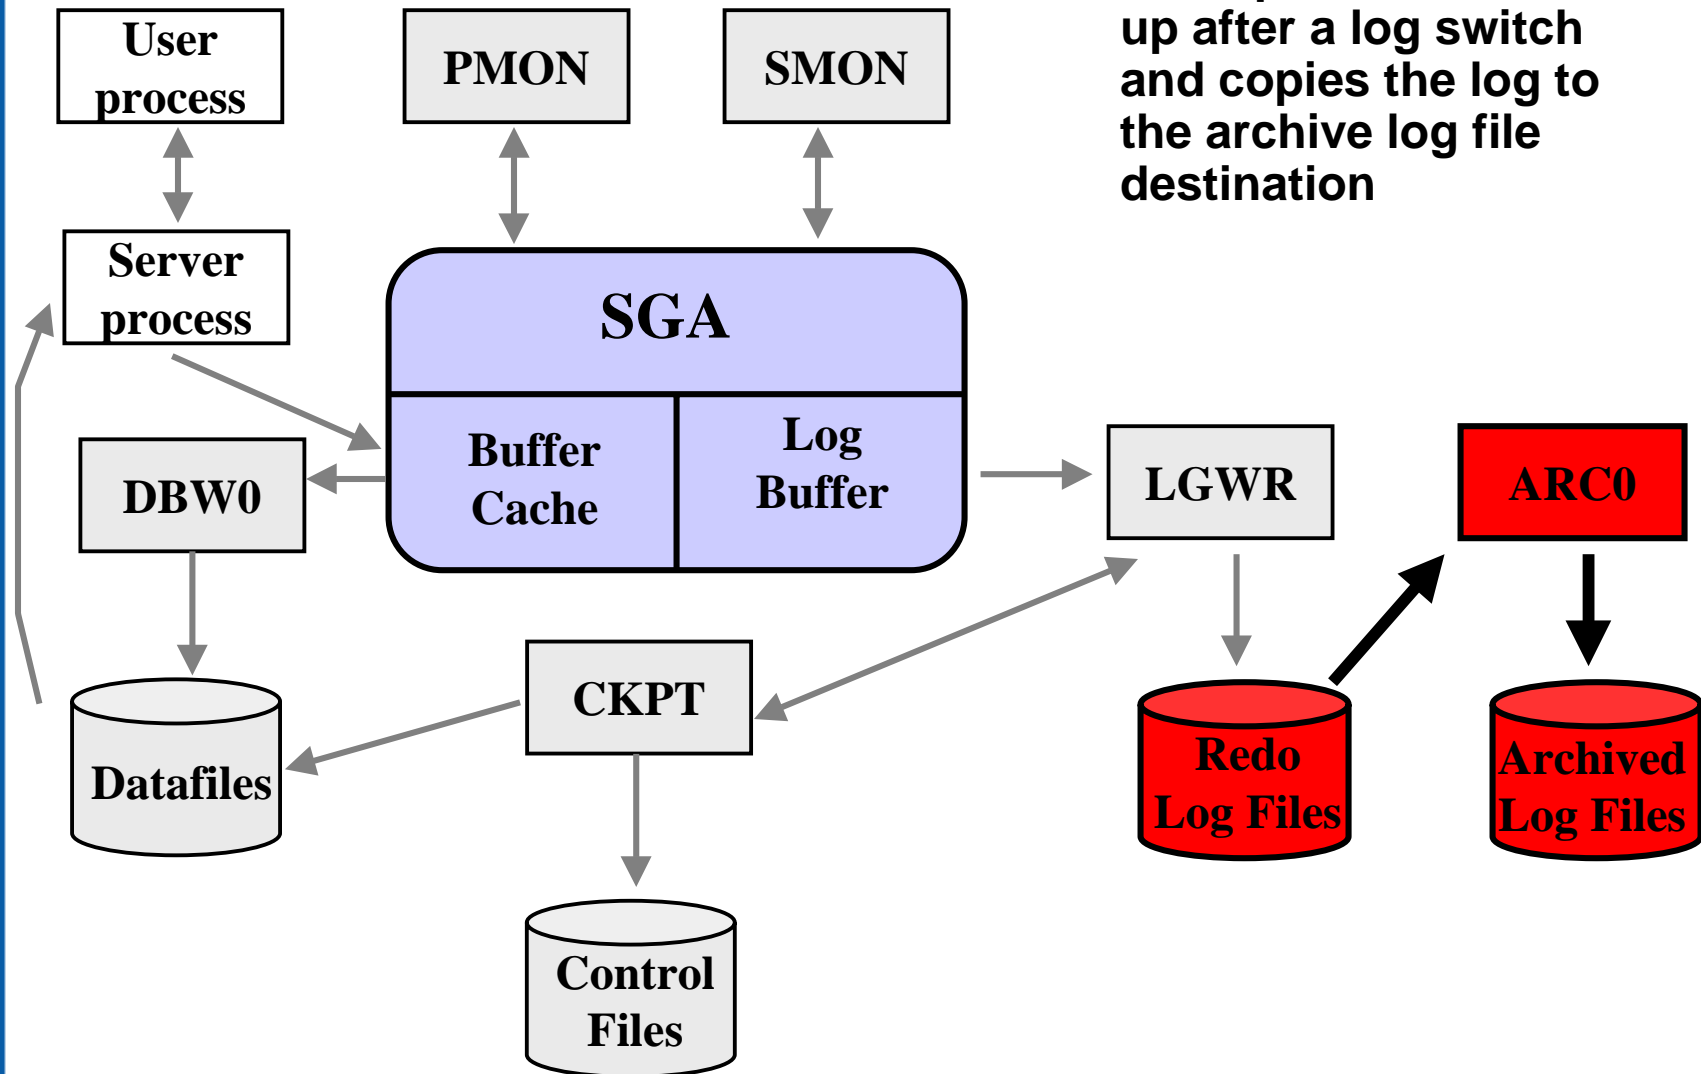
The archived redo log

- **Archive log files allow you to:**
 - Perform online backups
 - Perform point in time recovery
 - Maintain a standby database
- **Store archive logs separately**



ESRI

Log archive



•ARC process wakes up after a log switch and copies the log to the archive log file destination



ESRI

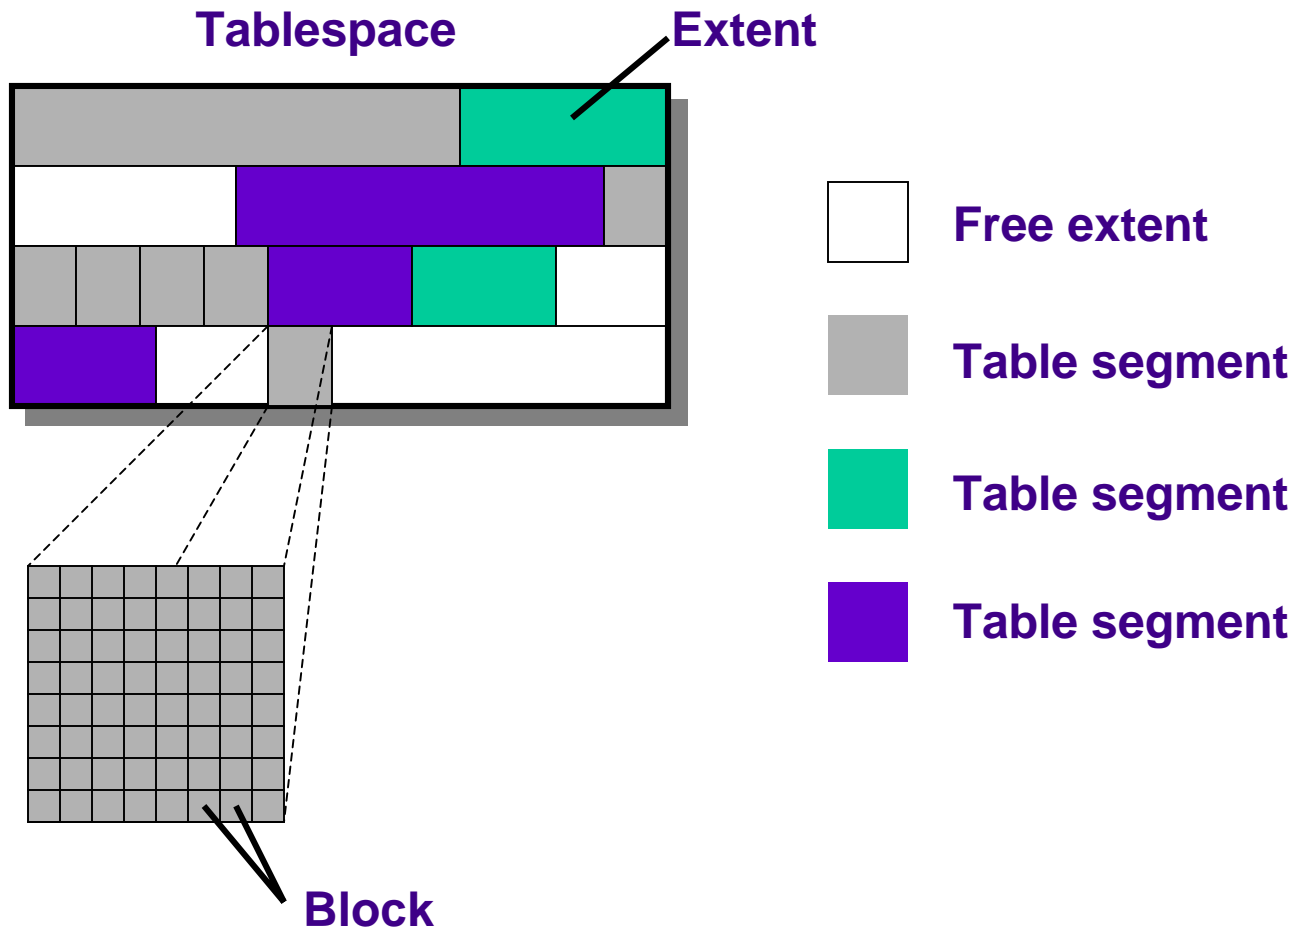
Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - **Storage basics**
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

Storage basics



ESRI

Data files

- **Physical storage of data on disk**
- **Store:**
 - Rollback & temporary segments
 - Oracle data dictionary
 - ArcSDE data dictionary
 - Tables and indexes



ESRI

Tablespaces

- **Tablespaces are logical storage structures for data**
 - **DBA can alter the physical storage without impacting users**



ESRI

Tablespaces

- **Assigned one or more physical datafiles**
 - Add datafiles to increase available space
- **Reduce disk I/O contention by separating datafiles on different disks**



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - **System Tablespace**
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI

SYSTEM tablespace

- **Data dictionary**
- **Store properties of the database**
 - Information about all database objects
- **Do not load other data into the SYSTEM tablespace**



ESRI

SYSTEM tablespace

- **Dynamic performance views report volatile system information**
 - Initialization parameters, memory statistics, redo log groups, etc.
 - Re-set when instance is shut down
 - Named v\$*



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - **Rollback segments**
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

Rollback segments

- **Stores before image of modified data blocks**
- **Provides read consistency in multi-user environment**
- **Allocate space to avoid having the rollbacks shrink**



ESRI

Rollback & Temporary segments

- Will fragment disk space
- Store segments in separate tablespaces
- Reduce fragmented free space by making extents the same size



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - **Temporary segments**
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

Temporary segments

- **Allocated and freed by Oracle**
- **Used for sorts larger than `sort_area_size`**
- **Oracle allocates `sort_area_size` to each table specified in a join**
- **UNIX does not release sort area until disconnect**



ESRI

Use CREATE TEMPORARY TABLESPACE

```
CREATE TEMPORARY TABLESPACE temp TEMPFILE  
'/azteca8/oradata/case_study/temp01.dbf'  
UNIFORM SIZE 2M AUTOEXTEND ON NEXT 1M MAXSIZE  
UNLIMITED;
```

- **Locally-managed**
 - Not logged
 - No need to coalesce free extents
- Does not need to be recovered



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - **The SDE Tablespace**
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

Creating the SDE tablespace

- **Stores the ArcSDE system tables**
 - Includes the ArcInfo geodatabase tables
 - At ArcSDE 8.1 for Oracle 8 the I/O is low because most data stays cached
 - At ArcSDE 8.1 for Oracle 8i the I/O increases because locks and sequences are stored in the SDE tables
 - Name tablespace **SDE** for convenience



ESRI

Creating the SDE tablespace

```
CREATE TABLESPACE sde
  DATAFILE <'drive:\path\filename.dbf' >
  SIZE 40M
  DEFAULT STORAGE
  (
    initial 40k
    next 40k
    minextents 1
    maxextents 500
    pctincrease 0
  )
;
```



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - **Disk I/O contention**
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI

Creating tablespaces for feature classes and indexes

- Determine the required space
- Place tables and indexes in different datafiles on different disks
- Size and frequency of use will influence distribution of tablespaces
- Distribute feature class tablespaces in file system for better performance



ESRI


Datafile distribution

- **Aim is to reduce disk I/O contention**
- **Distribute files that are accessed frequently or simultaneously onto different devices**



ESRI

Disk I/O




High Priority

- Redo log files mirrored and on separate disks
- Large redo log file to delay checkpoint
- Index tablespaces and table tablespaces on separate disks



ESRI

Disk I/O




Medium Priority

- Rollback tablespace on different disk than the table tablespace
- Temporary tablespace on different disk than the index tablespace
- ArcSDE for Oracle 8i the SDE tablespace stores SDE locks and sequences



ESRI

Disk I/O



Low Priority

- Rollback and temporary tablespaces separated on different disks
- **SYSTEM** tablespace can be placed with hot files
- Archive log directory on separate disk
- Control file mirrored and on separate disks



ESRI

Example (4 disks):

- **C:**
 - Oracle software
 - ArcSDE software
 - Pagefile
- **D:**
 - SYSTEM tablespace
 - SDE tablespace
 - Control file mirror
 - Redo log
- **E:**
 - Data tablespace
 - Temporary tablespace
 - Control file mirror
 - Data tablespace
- **F:**
 - Redo log
 - Rollback tablespace
 - Archive log directory
 - Index tablespace

More disks will directly impact the configuration and performance of your system.



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - **Creating users for SDE**
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

Creating the SDE user

- ArcSDE requires an Oracle user called **SDE**
 - `giomgr` connects as **SDE**
 - Default tablespace to **SDE**



ESRI

Creating the SDE user

```
CREATE USER sde
  IDENTIFIED BY < password >
  DEFAULT TABLESPACE < sde_tablespace >
  TEMPORARY TABLESPACE < temporary_tablespace >
  QUOTA UNLIMITED ON < sde_tablespace >
  QUOTA UNLIMITED ON < temporary_tablespace >
;
GRANT create session, alter session,
  create indextype, create operator,
  create procedure, create sequence,
  create synonym, create table, create trigger,
  create type, create view
TO sde
;
```



ESRI

Creating users

- **Users will own, edit and view the data**
- **Assign a default tablespace to each user**
- **Assign a temporary tablespace for sorting**
 - **SYSTEM tablespace is used by default; this results in a fragmented SYSTEM tablespace**



ESRI

Creating users

- **Assign system privileges**
 - create session
 - create table
 - create trigger
 - create index
 - create sequence
 - unlimited tablespace



ESRI®

Create users

- **Users must be granted all privileges for first ArcSDE connection, then you may revoke all but create session.**
 - **Must be able to create logoff trigger, sde logfiles and their indexes**
- **Assign profiles to prevent users from consuming all available space**



ESRI

General users

USER	SYSTEM PRIVILEGES	ArcSDE PRIVILEGES
Data owner/creator	CREATE TABLE, CREATE INDEX, CREATE SEQUENCE, CREATE SESSION	SELECT, INSERT UPDATE, DELETE
Data operator/editor	CREATE SESSION	SELECT, INSERT, UPDATE, DELETE
Data browser	CREATE SESSION	SELECT



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- Tuning the Spatial Index



ESRI

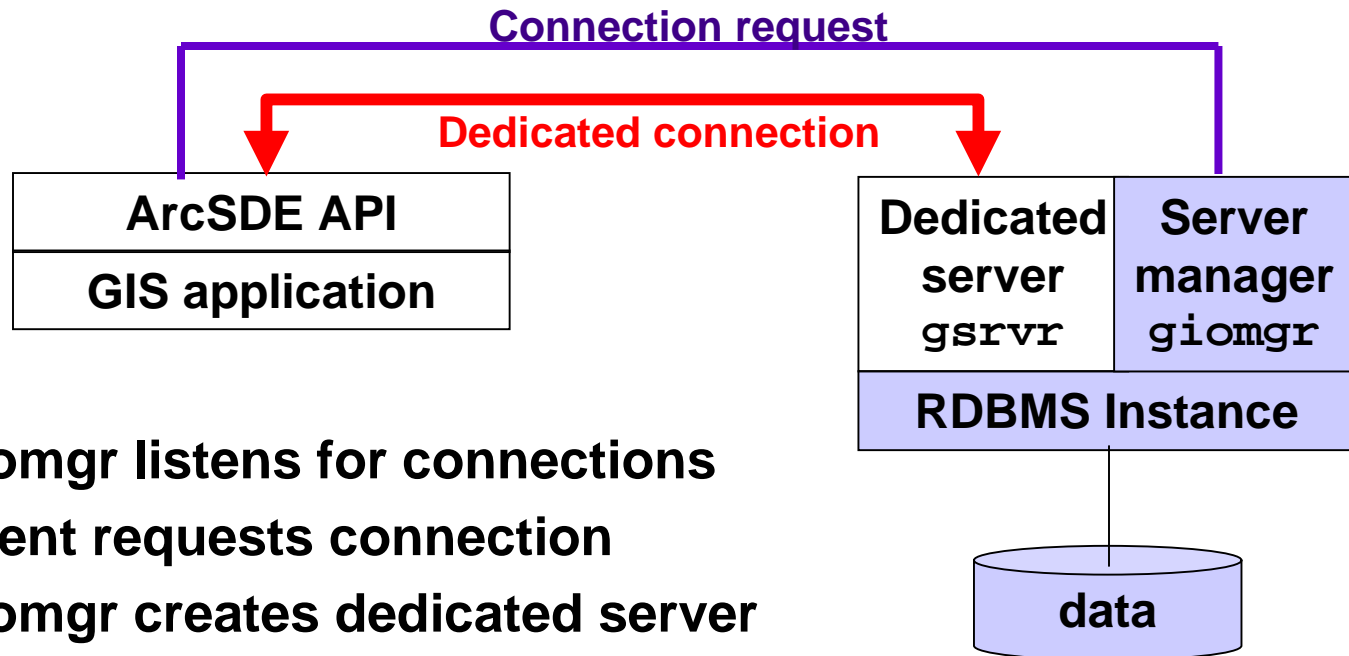
ArcSDE Three-Tiered Architecture

- **Client-server architecture**
- **Communicates through a TCP/IP network protocol**
- **Giomgr is the ArcSDE server process which listens over a dedicated port**
- **Establish port number during installation of ArcSDE**



ESRI

Three-tiered components



- Giomgr listens for connections
- Client requests connection
- Giomgr creates dedicated server
 - Connects to RDBMS as client user
- Dedicated server processes data with client
- Client terminates session, halts gsvr



ESRI

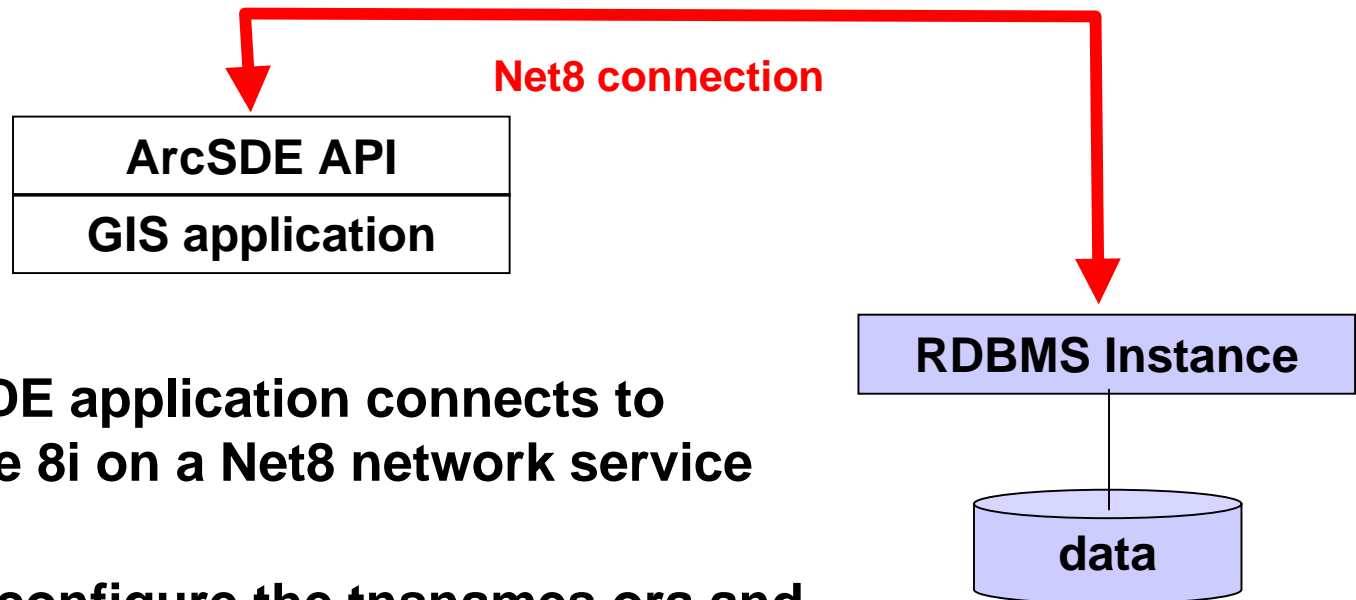
ArcSDE Two-Tiered Architecture

- Oracle 8i only
- Also called direct connect
- ArcSDE client connects directly to Oracle 8i through a Net8 network service name
- Gsrvr process has been moved into the ArcSDE client



ESRI

Two-tiered components



- ArcSDE application connects to Oracle 8i on a Net8 network service name
- Must configure the tnsnames.ora and listener.ora



ESRI®

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- **Configuring Oracle for loading data**
- **Configuring ArcSDE for loading**
- The DBTUNE table
- **Tuning the Spatial Index**



ESRI

Configuring Oracle for loading

- **Increase the size of the redo log files to delay checkpoint**
- **The database should be offline to general public**



ESRI

Rollback Segment

- Oracle records the before image of the data in the rollback segment
- If the rollback segment cannot allocate more space before the commit, the database will rollback the transaction



ESRI

Rollback Segment

- **Cannot specifically set the rollback segment before loading with ArcSDE utilities**
- **Set rollback segment extents to 2 to 3 MB**
- **Take the default smaller rollback segments off line**



ESRI

Temporary Segment

- Small sorts occur in the sort area
- Allocate 2 times the largest index size to the temporary tablespace
- Increase the **SORT_AREA_SIZE** during index creation



ESRI

Altering the Redo Log

- **NOARCHIVELOG** mode for initial loading
 - Eliminates archiving I/O
 - Data can be reloaded from source in the event of a media failure.



Altering the Redo Log

- **Increase redo log group size**
 - **Reduce checkpoints and delay the log switch by increasing the size of the redo log members**



Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- **Configuring ArcSDE for loading**
- The DBTUNE table
- Tuning the Spatial Index



ESRI®

Configure ArcSDE for loading

- Use the `giomgr.defs` file to alter ArcSDE configuration for data loading
- ArcSDE uses *transmission buffers* to support asynchronous I/O and streamline network usage



ArcSDE Transmission Buffers



- **Buffers are allocated on the server side and the client side**
 - **The size of the transmission buffer is set with `MAXBUFFSIZE` and `MINBUFFSIZE`**
- **Larger buffers allow higher throughput increasing performance during bulk loading**



ESRI

ArcSDE Transmission Buffers

- If the server is waiting, the buffer will gather up to **MINBUFFSIZE** of data to send to the server
- If the server is busy the client will gather up to **MAXBUFFSIZE** of data before sending the buffer to the client
- Reduces I/O by batching a minimum of amount of data



ESRI

ArcSDE Transmission Buffers

- High transmission buffers increase performance during bulk loading
- Do not use high buffers for normal query operations
 - ArcSDE will assign the `MAXBUFSIZE` amount of memory to each stream and may exhaust available memory



ESRI

Autocommit

- **The giomgr.defs default commit interval is 1000 records**
- **Increase autocommit interval to 5000 to reduce the log flush frequency**
- **Log flush occurs:**
 - **when the log buffer is 1/3 full**
 - **when ever a commit occurs**
 - **every 3 seconds**



Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- **The DBTUNE table**
- Tuning the Spatial Index



ESRI

DBTUNE table

- **The DBTUNE table is used to define Oracle storage parameters for the object class**
- **Use sdedbtune tool to manage the DBTUNE table**
- **A keywords contain CREATE TABLE and CREATE INDEX parameters**



ESRI

dbtune.sde (8.0.2)

- ArcSDE 8.0.2 and prior each dbtune parameter mapped to an Oracle parameter

```
##DEFAULTS  
  
INDEX_TABLESPACE <index_tspace>  
  
A_TBLSP           <data_tspace>  
A_INIT            131072  
A_NEXT            65536  
A_IX1_INIT        65536  
A_IX1_NEXT        65536  
A_MINX            1  
A_MAXX            100  
A_PCTI            0  
A_ITRANS          2  
A_MAXTRS          255  
A_PCTFREE         5  
A_PCTUSD          90
```



ESRI

dbtune file (8.1)

- At ArcSDE 8.1 parameters map to Oracle table and index objects

```
##DEFAULTS
```

```
UI_TEXT      `ESRI DEFAULTS`
```

```
COMMENT      `June 21,00`
```

```
B_STORAGE    `Tablespace busines
```

```
B_INDEX_1    `Tablespace indexes
```

```
F_STORAGE    `Tablespace feature
```

```
F_INDEX_1    `Tablespace indexes
```

```
S_STORAGE    `Tablespace spatial
```

```
S_INDEX_1    `Tablespace indexes
```

```
S_INDEX_2    `Tablespace indexes
```

```
A_STORAGE    `Tablespace delta
```

```
A_INDEX_1    `Tablespace indexes
```

```
D_STORAGE    `Tablespace delta
```

```
D_INDEX_1    `Tablespace indexes
```

```
END
```



ESRI

sdedbtune

- **Imports a dbtune file into a DBTUNE table**
- **Exports the DBTUNE table to a dbtune file**
- **Allows you to update the DBTUNE table by editing the dbtune file**



ESRI

Upgrading to DBTUNE table

- **ArcSDE 8.1 setupora81 admin tool automatically reads the dbtune.sde file into the DBTUNE table**
- **Reads in old style, writes out new style**



ESRI

Outline

- Basic Oracle Architecture
- The SGA
- The Oracle Database
 - The control files
 - The online redo log files
 - Storage basics
 - Rollback segments
 - Temporary segments
 - The SDE Tablespace
 - Disk I/O contention
 - Creating users for SDE
- ArcSDE architecture
- Configuring Oracle for loading data
- Configuring ArcSDE for loading
- The DBTUNE table
- **Tuning the Spatial Index**



ESRI


Spatial index grid size

- **Used to speed up spatial data retrieval**
- **Goals:**
 - **Minimize features / grid cell**
 - **Minimize grid cell / feature**
- **Start with 3X average feature extent**
- **Can change at any time**




ESRI

Estimating a grid size



Average feature extent:
S-table is too large



3X feature extent:
Balanced tiles and features

Balance the size of the S table versus the selectivity of the S table



ESRI

Average feature extent

- To obtain the average feature extent

```
SELECT AVG(emaxx - eminx),  
        AVG(emaxy - eminy)  
FROM F<n>;
```



ESRI®

Questions?

- **Please don't forget to fill out the survey before you leave**



ESRI



ESRI