10g New Features for DBAs

September 21, 2004

NYOUG General Meeting

New York

Dave Anderson

SKILLBUILDERS

Agenda

- A brief introduction to:
  - Flashback Enhancements
  - Auto Storage Management
  - Job Scheduler
  - RMAN Enhancements
  - Performance & Tuning Enhancements

Flashback Enhancements

SKILLBUILDERS

Flashback Evolution

- 9i provided
  - Session-level flashback
    - DBMS_FLASHBACK
  - Statement (and sub-statement) flashback
    - SQL “AS OF” clause

9i R1

9i R2

Can compare table to a previous version of *itself*

```
SQL> select a.lastname, a.total_purchase, b.total_purchase
  2  from sales a , sales AS OF timestamp(sysdate - 1) b
  3  where a.cust_no = b.cust_no
  4  and a.total_purchase != nvl(b.total_purchase, 0);

LASTNAME     TOTAL_PURCHASE TOTAL_PURCHASE
------------ -------------- --------------
ANDERSON                55000
DASWANI                 55000
```

Amount yesterday

Amount today

10g Enhancements

- Flashback Version Query
  - "VERSIONS BETWEEN" clause
- Flashback Transaction Query
  - FLASHBACK_TRANSACTION_QUERY view
- Flashback Table
  - Recover table to previous point-in-time
  - Recover from DROP
- Flashback Database
  - Recover database to previous PIT

Flashback Setup

- Flashback query and flashback table (sometimes) require *AUM* and *sufficient UNDO*

```
LINUX> show parameter undo_

NAME                                 TYPE        VALUE
------------------------------------ ----------- ---------
undo_management                      string      AUTO
undo_retention                       integer     172800
undo_tablespace                      string      UNDOTBS1

  1  select file_name, bytes
  2  from dba_data_files
  3* where tablespace_name = 'UNDOTBS1'

FILE_NAME                                                BYTES
-------------------------------------------------- ----------
+ASM_DISK_GROUP1/orcl/datafile/undotbs1.265.1      209715200
+ASM_DISK_GROUP1/orcl/datafile/undotbs1.295.17     104857600
```

48 hours

Flashback Version Query…

- Show all “versions” of data between two
  - SCN’s
  - Timestamps
- “Version” is transaction-based

…Flashback Version Query…

```
DAVE@LINUX> select c2, versions_starttime, versions_endtime,
  2             versions_startscn , versions_endscn,
  3             versions_operation, versions_xid
  4  from system.test
  5      versions between timestamp
  6                       to_timestamp('11-SEP-04 12.02.00.000000000 PM',
  7                                    'dd-mon-yy hh.mi.ss.ff PM')
  8                       and systimestamp
  9  where c1 = '1'
 10  order by versions_startscn nulls first;

C2   VERSIONS_STARTTIME          VERSIONS_ENDTIME          VERSIONS_STARTSCN
---- ------------------------- ------------------------- ----------------
a                              11-SEP-04 12.01.55 PM
b    11-SEP-04 12.01.55 PM     11-SEP-04 12.01.55 PM             988714
c    11-SEP-04 12.01.55 PM     11-SEP-04 12.02.04 PM             988716
d    11-SEP-04 12.02.04 PM     11-SEP-04 12.02.17 PM             988719
x    11-SEP-04 12.02.17 PM                                       988730
```

New pseudocolumns

NULL in ENDTIME…this is the current value

…Flashback Version Query

- Users need privileges to flashback

Privilege to flashback query, flashback version query and flashback table

```
LINUX> grant flashback, select on system.test to dave;

Grant succeeded.
```

Flashback_Transaction_Query

- Mine (audit) undo records for details on changes
  - Easier than log miner utility
- Get XID from flashback version query

```
LINUX> exec print_table('select logon_user, undo_sql -
>                       from flashback_transaction_query -
>                       where xid = ''07000100590200000'' ')
LOGON_USER  : SYSTEM
UNDO_SQL    : delete from "SYSTEM"."TEST"
              where ROWID = 'AAALquAABAAAaIiAAA'
```

Get XID from flashback version query

```
LINUX> grant select any transaction to dave;

Grant succeeded.
```

User needs this privilege

Flashback Table…

- Turn back hands of time for one or more tables
- Recover from DROP TABLE
  - Uses “recycle bin”
- Recover from application or user changes
  - Uses UNDO
- Implemented with FLASHBACK TABLE statement

…Flashback Table

- Prepare
  - Enable row movement on table
  - Grant privileges to desired user(s)

```
SYSTEM@LINUX> alter table test enable row movement;

Table altered.

SYSTEM@LINUX> grant flashback on test to dave;

Grant succeeded.

SYSTEM@LINUX> grant alter, select, update, delete,
  2           insert on system.test to dave;

Grant succeeded.
```

Flashback Table: Undo Changes

```
DAVE@LINUX> select dbms_flashback.get_system_change_number
from dual;

         X
----------
   1057181

DAVE@LINUX> flashback table system.test to timestamp
  2         systimestamp - interval '30' minute;
```

Can undo the flashback if you know SCN

```
DAVE@LINUX> flashback table system.test to scn 1057181;

Flashback complete.
```

Flashback Table: Recover from DROP

- Dropped tables and dependent objects aren't really dropped
- Renamed
  - "BIN*uniquestring*"
- This is the "recycle bin" concept
- Objects actually dropped when
  - Space needed
    - Could be right away, might be there "forever"
  - Tablespace needed
  - Recycle bin purged
  - PURGE clause used on DROP

Flashback Table: Recover from DROP

- Recover from DROP TABLE or unwanted application changes

```
SQL> drop table t purge;

Table dropped.
```

```
SQL> drop table big;

Table dropped.
```

Table remains in segment unless PURGE used

Table is recorded in “recycle bin”

```
SQL> show recyclebin
ORIGINAL NAME    RECYCLEBIN NAME
---------------- ------------------------------
BIG              BIN$Kq3scQsDQP+jkI0atdsqqQ==$0
```

```
 OBJECT TYPE  DROP TIME
------------ --------------------
TABLE        2004-08-19:11:40:00
```

Easy and fast to recover table

```
SQL> flashback table big to before drop;

Flashback complete.
```

Repeat Un-Drop

```
SQL> show recyclebin
ORIGINAL NAME    RECYCLEBIN NAME                OBJECT TYPE  DROP TIME
---------------- ------------------------------ ------------ -------------------
T                BIN$KbsjpLROTs6+eGOGqifipQ==$0 TABLE        2004-09-19:14:59:30
T                BIN$y0AR2dYXSnC0bRn+vJTF0Q==$0 TABLE        2004-08-19:16:52:54

SQL> flashback table t to before drop;

Flashback complete.

SQL> show recyclebin
ORIGINAL NAME    RECYCLEBIN NAME                OBJECT TYPE  DROP TIME
---------------- ------------------------------ ------------ -------------------
T                BIN$y0AR2dYXSnC0bRn+vJTF0Q==$0 TABLE        2004-08-19:16:52:54
SQL> rename t to t_old;

Table renamed.

SQL> flashback table t to before drop;

Flashback complete.

SQL> show recyclebin
SQL>
```

Purging Recycle Bin

```
SQL> show recyclebin
ORIGINAL NAME    RECYCLEBIN NAME                 OBJECT TYPE  DROP TIME
---------------- ------------------------------ ------------ -------------------
T2               BIN$cfmTndAnQk+junRG6hCngg==$0 TABLE        2004-09-19:14:59:28
T_AUDIT          BIN$Fx+nQHXcR4qgr4fSc18QiQ==$0 TABLE        2004-08-23:15:52:29
T_AUDIT          BIN$mPxGfYd8QJ+sLvSNPLszZg==$0 TABLE        2004-08-23:15:51:13
T_AUDIT          BIN$7sKuyR6YR7KM4Bc7mzewiQ==$0 TABLE        2004-08-23:15:19:19
T_AUDIT          BIN$ZD1sowFXTt6T7uFsmDt4aw==$0 TABLE        2004-08-23:13:50:33
T_AUDIT          BIN$STuYtaPCQcC0eEhgZd+HRw==$0 TABLE        2004-08-23:13:49:26
T_AUDIT          BIN$Em072EijQIilk8I5QhHFrA==$0 TABLE        2004-08-23:13:48:08
T_AUDIT          BIN$nuPfG+vCT8OsZDi4lJ0jdA==$0 TABLE        2004-08-23:13:05:05

SQL> purge recyclebin;

Recyclebin purged.

SQL> show recyclebin
SQL>
```

Flash Recovery Area…

- Optional storage area for backup-related files
  - Online and archive logs
  - RMAN backups
    - Default location for RMAN backups if configured
  - Flashback logs
    - Required for FLASHBACK DATABASE
- "Automates management of backup-related files"
  - Convenient directory structure
  - Auto-delete obsolete files when space needed
- Keep on separate disk from datafiles

Next topic is FB Database…

…Flash Recovery Area…

- Set space limit
  - Maximum space dedicated to flashback area
- Set retention limit
  - How far back can we flashback database?

```
NAME                                 VALUE
------------------------------ ------------------------------
db_recovery_file_dest          /mnt/mickeymantle
db_recovery_file_dest_size     10G
db_flashback_retention_target  1440
log_archive_dest_1             LOCATION=USE_DB_RECOVERY_FILE_DEST
```

Limit size

How long to keep flash logs

Write archive logs to flash recovery area

…Flash Recovery Area…

- New V$ view

```
SQL> exec print_table('select * from v$recovery_file_dest')
NAME                           : /mnt/mickeymantle/
SPACE_LIMIT                    : 10737418240
SPACE_USED                     :  2485816832
SPACE_RECLAIMABLE              :    19995648
NUMBER_OF_FILES                :          35
-----------------
```

How much space can be made available through delete of “obsolete, redundant or low priority files”

…Flash Recovery Area

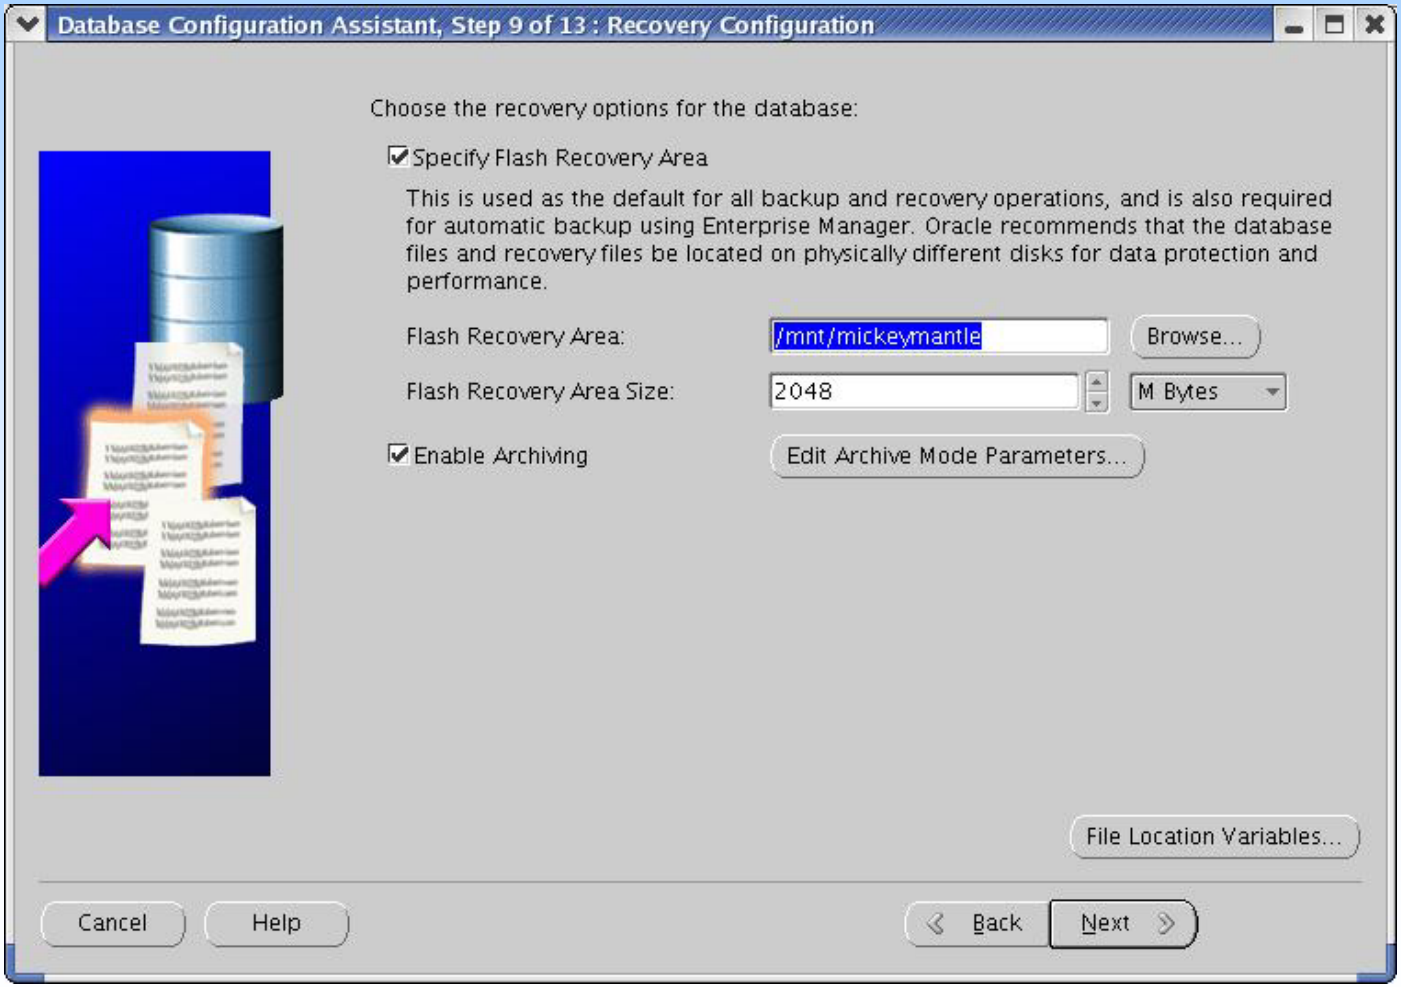

Flashback Database…

- Alternative to point-in-time recovery
  - Very easy
  - Can flashback many times
    - Until you OPEN RESETLOGS
- Use *flashback logs* to recover database
  - Contain changed blocks
  - Ongoing creation of new logs to capture changes
  - Written to “Flash Recovery Area”
  - Transactions dictate frequency / size of logs
  - Open questions
    - Exactly what initiates log write? What’s threshold?
  - Automatic deletion of obsolete logs

…Flashback Database

- Overhead
  - Initial tests show measurable increase
  - Stay tuned…
- Performance versus point-in-time recovery
  - Often faster
  - Reapply changed blocks versus restore then recover

Flashback Database: Setup

- Configure flash recovery area
  - Destination, size limit, retention
- Configure database

Starts recovery writer background process (RVWR)

```
SQL> alter database flashback on;

Database altered.

SQL> select flashback_on from v$database;
FLA
---
YES
```

Flashback Database: Example

```
SQL> startup mount

SQL> alter database datafile
  2> '+ASM_DISK_GROUP1/orcl/datafile/users.268.1' offline;

Database altered.

SQL> flashback database to scn 1427369;

Flashback complete.

SQL> alter database open read only;

Database altered.
```

Optional step - offline files are not flashed back

Flashback to SCN or timestamp. Need SYSDBA privilege.

Can open READ ONLY to check things out

Can then re-mount, flashback to different scn or time

```
SQL> alter database open resetlogs;

Database altered.
```

Resetlogs deletes old FB logs

NOW DO FULL BACKUP!

Flashback Summary

- Flashback Query
  - Session or sub-statement level
  - Query data at a previous point-in-time
- Flashback Version Query
  - Show changes made by transactions
  - Details about transaction
- Flashback Transaction Query
  - Access UNDO records, including undo-SQL
- Flashback Table
  - “Reset” table to previous point-in-time
  - Recover from dropped table
- Flashback Database
  - Point-in-time recovery for entire database
  - Must configure database for this…

Automatic Storage Management

ASM

SKILLBUILDERS

Introducing ASM

- Built-in file system and volume manager
  - "Vertical integration of both the file system and the volume manager"
    - Oracle Corp
- Automated S.A.M.E.
  - Striping to balance I/O load
    - Even with database open
  - 2 or 3-way mirroring
- Major extension to Oracle Managed Files (9i)
- Can hold all database-related files
  - Not Oracle executables
- Supports Real Application Clusters

The ASM Instance…

- A completely separate instance
  - See parameter INSTANCE_TYPE=ASM
- Each physical Oracle server has one ASM instance
  - Can be used by many database instances
  - Failure of ASM instance causes failure of dependents
- Starts background process to manage ASM disk metadata
  - Requires approximately 100MB
- ASM instance has no database
  - No datafiles, control file, log files
  - No need to backup ASM

…The ASM Instance…

- DBCA will either
  - Create new ASM instance
    - Configure diskgroups
  - Allow use of existing ASM instance and diskgroups
    - Can configure new diskgroups

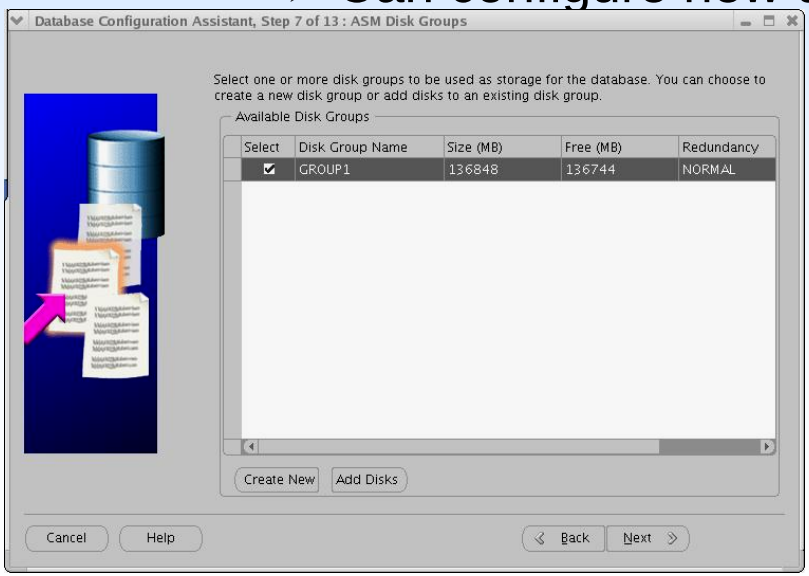

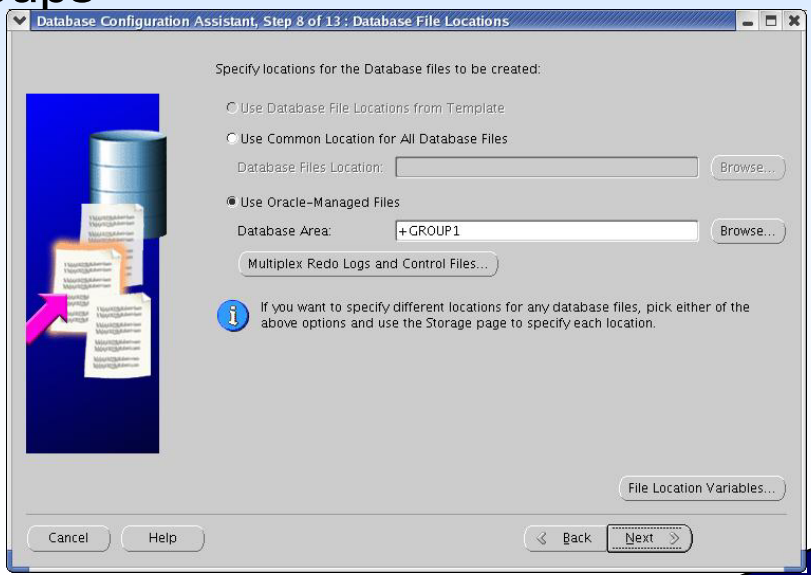

…The ASM Instance

- Must start ASM instance before dependent database instances
- Stop ASM instance last
  - *Dependent database instances will crash if ASM instance shut down first*
- Handles IO errors
  - Remove disk if suffers write errors
  - Read from secondary disk for read errors

Starting "+ASM" Instance

- Start ASM first
  - Mounts all diskgroups in ASM_DISKGROUPS parameter
- Shutdown ASM last

```
[oracle@springsteen oracle]$ ORACLE_SID=+ASM
[oracle@springsteen oracle]$ sqlplus / as sysdba
. . .
SQL> startup
ASM instance started

Total System Global Area  100663296 bytes
Fixed Size                   777616 bytes
Variable Size              99885680 bytes
Database Buffers                  0 bytes
Redo Buffers                      0 bytes
ASM diskgroups mounted
SQL>
```

Default ASM instance name

Must (always) connect as SYSDBA

Mounts disks; no database to open!

ASM Disks

- ASM disks must be discovered
  - Linux
    - Give raw devices to ASM
    - Or install “ASMLIB”
  - Windows
    - Assign unformatted logical partitions to ASM?
- Parameter ASM_DISKSTRING to discover disks
  - Can limit to specific disks
  - Query V$ASM_DISK to see available disks

ASM Diskgroups…

- Create diskgroups from ASM disks
  - Like a volume or storage group
- Used by one or more database instances
- Striping for balanced IO load across disks
  - Files extents “distributed equally across all … disks in the diskgroup”*

…ASM Diskgroups

- Mirroring
  - External
    - Use EMC or other storage product
  - Normal
    - 2-way
  - High
    - 3-way
- Failgroups provide ASM (internal) mirroring

Creating Diskgroup

```
+ASM> create diskgroup group1 normal redundancy
  2  failgroup fgroup1 disk '/dev/raw/raw1', '/dev/raw/raw2'
  3  failgroup fgroup2 disk '/dev/raw/raw3', '/dev/raw/raw4';

Diskgroup created.

SQL> select * from v$asm_diskgroup;

GROUP_NUMBER NAME         SECTOR_SIZE BLOCK_SIZE
------------ ------------ ----------- ----------
           1 GROUP1               512       4096

ALLOCATION_UNIT_SIZE STATE     TYPE     TOTAL_MB      FREE_MB
-------------------- -------   ------   --------      -------
             1048576 MOUNTED   NORMAL     136848       136744
```

```
+ASM> select g.name as group_name, d.name as disk_name, d.path, d.failgro
  2  from v$asm_diskgroup g, v$asm_disk d
  3  where g.group_number = d.group_number

GROUP_NAME DISK_NAME        PATH              FAILGROUP
---------  ---------------- ----------------- -------
GROUP1     GROUP1_0003      /dev/raw/raw4     FGROUP2
GROUP1     GROUP1_0002      /dev/raw/raw3     FGROUP2
GROUP1     GROUP1_0001      /dev/raw/raw2     FGROUP1
GROUP1     GROUP1_0000      /dev/raw/raw1     FGROUP1
```

Altering Diskgroup

- Easily add or remove disks while database open
- Automatic rebalancing starts

```
+ASM> alter diskgroup group1 drop disk GROUP1_0002;

Diskgroup altered.

+ASM> select * from v$asm_operation
  2  /

GROUP_NUMBER OPERA STAT      POWER     ACTUAL       SOFAR   EST_WORK   EST_RATE
------------ ----- ---- ---------- ---------- ---------- ---------- ----------
           1 REBAL RUN           1          1        181       1275        338

EST_MINUTES
-----------
          3
```

Database Instance

- Databases can mix
  - ASM files, OS files, raw
- Can setup OMF DB_CREATE_* parameters
- ASM files are not visible to OS
- Backup with RMAN

```
SQL> select name from v$datafile;
NAME
----------------------------------------
+GROUP1/orcl/datafile/sysaux.256.3
+GROUP1/orcl/datafile/system.258.3
+GROUP1/orcl/datafile/undotbs1.2
+GROUP1/orcl/datafile/users.265.
```

Can create ASM aliases for ASM FQFN's

```
SQL> select member from v$logfile;
MEMBER
----------------------------------------
+GROUP1/orcl/onlinelog/group_3.264.3
+GROUP1/orcl/onlinelog/group_2.266.3
+GROUP1/orcl/onlinelog/group_1.267.3
```

Some ASM Benefits

- Easy S.A.M.E. implementation
- Raw disk IO performance
- Easy to add and remove disks
- Oracle SQL-like commands to manage
  - Same commands across all platforms

ASM Summary

- Built-in file system and volume manager
- Provides automated S.A.M.E.
- Will it - in some environments - phase out 3rd party products?
  - You'll be the judge
- Certainly attractive to small shops at least
- Resources
  - Must read Oracle White Paper :
    - Oracle Database10g Automatic Storage Management Technical Best Practices by Nitin Vengurlekar
  - Oracle docs
  - Metalink notes

Job Scheduler

SKILLBUILDERS

Job Scheduler - Concepts

- Replacement for DBMS_JOB
  - More capabilities
  - Can still use DBMS_JOB
    - Migrate over time
  - Uses same coordinator process
- Implemented via DBMS_SCHEDULER package
- DBCA automatically creates two jobs
  - Scheduler log purge
  - Daily DBMS_STATS

Capability Overview…

- Schedule jobs to run one or more times
  - PL/SQL
  - OS scripts
- Powerful and flexible scheduling capability
- Pass parameters at run time
- Control resource usage
  - Establish job “classes”
    - Group jobs with same characteristics into a class
  - Prioritize with resource plans (see Resource Manager)
  - Set “windows”
    - Different resource plans at different times

…Capability Overview

- Define “programs”
  - Pre-defined reusable OS executable
- Monitor job status and progress
  - Job details recorded in job log
  - Control
    - Amount of logging by job class
    - Lifespan of entries

```
SQL> exec dbms_scheduler.set_scheduler_attribute –
>         ('log_history', '45')

PL/SQL procedure successfully completed.
```

Keep log data for 45 days

- Define number of failures that dictate “broken” job
- RAC supported
- Export scheduler objects with Data Pump

Create Schedule & Job

```
begin
dbms_scheduler.create_schedule (
  schedule_name      => 'backup_schedule',
  start_date         => SYSTIMESTAMP,
  repeat_interval    => 'FREQ=DAILY; BYHOUR=01',
  comments           => 'Schedule for daily backups');
end;
/
```

Run daily at 1 AM

```
begin
dbms_scheduler.create_job (
   job_name             =>  'system.daily_backup',
   schedule_name        =>  'dave.backup_schedule',
   job_type             =>  'executable',
   job_action           =>  '/home/oracle/full_backup.sh',
   enabled              =>  true,
   comments             =>  'Daily full backup');
end;
/
```

Runs with privs of owner.

Run OS command

Specify path

OS Scripts

- Fully qualify paths
- Set permissions
- Set environment variables

full_backup.sh

```
#/bin/sh
exec >/tmp/log.txt 2>&1
export ORACLE_HOME=/u01/app/oracle/product/10.1.0/db_2
export PATH=$ORACLE_HOME/bin:$PATH
rman target sys/dave@orcl cmdfile /home/oracle/full_backup.rman
```

full_backup.rman

```
backup database plus archivelog tag='daily full backup';
exit;
```

Stored Jobs

```
LINUX> exec print_table('select owner, job_name, program_owner, program_name, -
>                schedule_owner, schedule_name, start_date, repeat_interval, -
>                last_start_date, next_run_date, comments -
>                from dba_scheduler_jobs')
OWNER                          : SYS
JOB_NAME                       : PURGE_LOG
PROGRAM_OWNER                  : SYS
PROGRAM_NAME                   : PURGE_LOG_PROG
SCHEDULE_OWNER                 : SYS
SCHEDULE_NAME                  : DAILY_PURGE_SCHEDULE
START_DATE                     : 02-SEP-04 03.00.00.000000 AM -04:00
REPEAT_INTERVAL                :
LAST_START_DATE                : 07-SEP-04 03.00.00.437528 AM -04:00
NEXT_RUN_DATE                  : 08-SEP-04 03.00.00.000000 AM -04:00
COMMENTS                       : purge log job
-----------------
OWNER                          : SYS
JOB_NAME                       : GATHER_STATS_JOB
PROGRAM_OWNER                  : SYS
PROGRAM_NAME                   : GATHER_STATS_PROG
SCHEDULE_OWNER                 : SYS
SCHEDULE_NAME                  : MAINTENANCE_WINDOW_GROUP
.  .  .
```

Stored Schedules

```
SQL> exec print_table('select * from dba_scheduler_schedules')
OWNER                          : SYS
SCHEDULE_NAME                  : DAILY_PURGE_SCHEDULE
START_DATE                     :
REPEAT_INTERVAL                : freq=daily;byhour=3;byminute=0;bysecon
END_DATE                       :
COMMENTS                       :
-----------------
OWNER                          : DAVE
SCHEDULE_NAME                  : BACKUP_SCHEDULE
START_DATE                     : 06-SEP-04 02.38.31.174000 PM -04:00
REPEAT_INTERVAL                : FREQ=DAILY; BYHOUR=01
END_DATE                       :
COMMENTS                       : Schedule for daily backups
-----------------
```

Other Dictionary Views

- DBA_SCHEDULER_RUNNING_JOBS
  - Information on active jobs
- DBA_SCHEDULER_JOB_LOG
  - One row for each job recorded in log
- DBA_SCHEDULER_JOB_RUN_DETAILS
  - Historical data on jobs

Miscellaneous

- Drop jobs

Need MANAGE SCHEDULER system privilege

```
SQL> exec dbms_scheduler.drop_job('daily_backup');

PL/SQL procedure successfully completed.
```

- Run job now

Need to own job or have ALTER JOB privilege

```
SQL> exec dbms_scheduler.run_job('daily_backup')

PL/SQL procedure successfully completed.

SQL> exec dbms_scheduler.run_job(job_name=>'gather_stats_job', -
>         use_current_session=>false)

PL/SQL procedure successfully completed.
```

Runs in another session

Privileges

- SCHEDULER_ADMIN
  - All powerful
- MANAGE SCHEDULER
  - Create, alter and drop windows, classes, window groups, purge logs
- CREATE JOB, CREATE ANY JOB
  - Create jobs, schedules and programs in schema
- EXECUTE ANY PROGRAM
  - Execute any program
- EXECUTE ANY CLASS
  - Run your jobs in any class

Scheduler Summary

- Better job scheduler than DBMS_JOB
- Integrated with Resource Manager

RMAN Enhancements

SKILLBUILDERS

Compressed Backups

- "Real" compressed backupset pieces

```
RMAN> backup as compressed backupset full database plus archivelog;


Starting backup at 08-SEP-04
current log archived
```

- If using ASM, get piece size from ASM instance

```
+ASM> select a.name, f.blocks, f.bytes
2 from v$asm_alias a, v$asm_file f
3 where a.file_number=f.file_number
4 and type = 'BACKUPSET'
5* order by f.type, a.name
+ASM> /

NAME                                     BLOCKS      BYTES
---------------------------------------- ---------- ----------
annnf0_TAG20040908T162645_0.284.5            12115    6202880
nnndf0_TAG20040908T161717_0.343.13          136003  557068288
```

Savings of 123888 blocks!

Tiny database; your results will vary…

Change Tracking…

- Change tracking eliminates need to read entire file during incremental backup

```
RMAN> sql 'alter database enable block change tracking';

sql statement: alter database enable block change tracking

RMAN> backup incremental level 0 database;

Starting backup at 08-SEP-04
channel ORA_DISK_1: backup set complete, elapsed time: 00:02:45

-- later, after normal database use. . .

RMAN> backup incremental level 3 database;

Starting backup at 08-SEP-04
channel ORA_DISK_1: backup set complete, elapsed time: 00:00:03
```

Level 0 duration

Level 3 duration

Level 3 duration w/o change tracking
time: 00:00:46

…Change Tracking

- Block change recorded in a file

```
LINUX> select * from v$block_change_tracking;

STATUS   FILENAME                                                BYTES
-------- ---------------------------------------------------- ----------
ENABLED  +ASM_DISK_GROUP1/orcl/changetracking/ctf.343.9    11599872
```

- File size estimate 1/30,000 of db size
- Caution
  - Oracle doc mentions performance hit
  - Nothing specific
  - Must Test

Incrementally Updated Backups…

- Apply (merge) incremental backups into image copies
- Reduce recovery time
  - At most 24 hours of redo to apply if run daily
- Reduce number of times image copies are taken
  - Optimally, just once

…Incrementally Updated Backups

Run this script daily

Apply any available level 1 incrementals to the image copy with the tag “ic1”

```
RMAN> run {
2>  recover copy of database with tag 'ic1';
3>  backup
4>      incremental level 1 tag 'level_1_for_image_copy_update'
5>      for recover of copy with tag 'ic1'
6>      database;
7>  }
Star    recover at 09-SEP-04
```

Create a level 1 incremental to be used to merge into image copy

Also can add:

“as compressed backupset”

“plus archivelog [delete input]”

DURATION Parameter

- Control backup rate
- More flexible and intuitive than RATE parameter

```
RMAN> backup duration 0:30 minimize load tablespace users;

Starting backup at 03-AUG-04
```

Maximum time allotted is 30 minutes

OPTIONAL: "minimize load" says "take the entire 30 minutes"

More Enhancements…

- Archive log deletion policy
- Default backup can be configured to create
  - Normal backupset
  - Compressed backupset
  - Image copy
- Default backup location is flash recovery area (if configured)
- RMAN catalog tablespace created in SYSAUX
- RESTORE . . . PREVIEW [SUMMARY]
  - Identify the backups that would be used
- Simplified recovery through incarnations

…More Enhancements

- Restore failover
  - If error reading backup, RMAN keeps trying different backup copies until all possible exhausted
- Channel failover
- COPY command deprecated
  - BACKUP can create image copies
- CATALOG command to catalog user-made copies of backups
- More…

10g Performance and Tuning Features

SKILLBUILDERS

What's New?

- Automatic SGA (memory) Management
- Automatic Workload Repository
- Automatic Database Diagnostic Monitor
- Automatic Statistics Collection
- New Advisors
  - SQL Tuning Advisor
- SQL Profiles
- PL/SQL Enhancements
  - Optimizing compiler
- Data Pump Utility
- etcetera

Automatic SGA Management…

- When enabled, database controls size of
  - Shared pool, Buffer cache, Large pool, Java pool
- Enable with SGA_TARGET
  - Set equal to memory you want dedicated to SGA
    - Must be <= SGA_MAX_SIZE
    - Dynamic parameter
  - Set = 0 to disable
  - Also set STATISTICS_LEVEL to TYPICAL
- Size of areas can dynamically change as:
  - Workload changes
  - Change made to SGA_TARGET or other SGA parameters

This feature also called “Auto Memory Mgmt” or “AMM”

…Automatic SGA Management

- Optionally set area *minimum* sizes with
  - DB_CACHE_SIZE
  - SHARED_POOL_SIZE
  - LARGE_POOL_SIZE
  - JAVA_POOL_SIZE
- Must still manually configure
  - KEEP, RECYCLE, Non-Standard Blocksize caches
  - Log buffer
  - Streams pool (new with 10g)
    - See STREAMS_POOL_SIZE parameter
  - *These areas are taken from SGA_TARGET*

"double uscore" versions created to restart DB with most recent value.* Create PFILE from SPFILE to see… `__shared_pool_size`

Auto SGA Demo…

```
SQL> select component, current_size
  2  from v$sga_dynamic_components;

COMPONENT                 CURRENT_SIZE
------------------------- ------------
shared pool                   50331648
large pool                     4194304
java pool                      4194304
streams pool                         0
DEFAULT buffer cache          104857600
```

```
SQL> select name, value
  2  from v$parameter
  3  where name like 'sga_%'

NAME                 VALUE
-------------------- -------
sga_max_size         209715200
sga_target           163577856

SQL> select sum(current_size)
  2  from v$sga_dynamic_components;

SUM(CURRENT_SIZE)
-----------------
       1635778560
```

…Auto SGA Demo

```
SQL> select * from v$sga_dynamic_free_memory;

CURRENT_SIZE
------------
    41943040

SQL> alter system set sga_target=200m;

System altered.

SQL> select sum(current_size
  2  from v$sga_dynamic_compo

SUM(CURRENT_SIZE)
-----------------
        205520896
```

Buffer cache given more memory

```
SQL> select component, c        t_size
  2  from v$sga_dynamic_c        nents;

COMPONENT                       RENT_SIZE
------------------------------  ---------
shared pool                      50331648
large pool                        4194304
java pool                         4194304
streams pool                            0
DEFAULT buffer cache            146800640
```

AWR…

- Automatic Workload Repository
- Repository of performance statistics
  - Don't have to replay workload to diagnose problem!
- "AWR forms the foundation for all self-management functionality of Oracle Database" *Concepts manual*
  - Feeds the "auto" stuff
- Replaces STATSPACK?
  - You'll be the judge
- Used by:
  - The "Advisors"
  - ADDM
  - DBA
  - Auto SGA?

```
SYSDBA> select advisor_name from
dba_advisor_definitions;

ADVISOR_NAME
-----------------------------
ADDM
SQL Access Advisor
Undo Advisor
SQL Tuning Advisor
Segment Advisor
SQL Workload M
Tune MView
```

We will discuss Advisors later…

…AWR…

- SGA / shared pool contains statistics
  - Snapshot every 60 minutes
  - 1 – 2 Mb per CPU
  - Increase shared pool 15% – 20% over 9i size*
  - Query V$SYSSTAT, V$SEGMENT_STATISTICS, etcetera
- Written to SYSAUX tablespace
  - Default write interval every 30 minutes
  - Default duration saved - 7 days
    - Maybe increase to 30 days?
  - Query DBA_HIST_* and WRM$_* tables
    - DBA_HIST_ACTIVE_SESS_HISTORY
    - DBA_HIST_WAITSTAT

See new background processes MMON and MMML that handle these tasks

Controlling AWR

- Control with DBMS_WORKLOAD_REPOSITORY package

```
SYSDBA> exec dbms_workload_repository.modify_snapshot_settings( -
>       retention=>43200, -
>       interval=>30)

PL/SQL procedure successfully comple

SYSDBA> select * from dba_hist_wr_control
  2  /

      DBID SNAP_INTERVAL     RETENTION
---------- ----------------- -----------------
3652163603 +00000 00:30:00.0 +00030 00:00:00.0
```

Retention = 43200 minutes = 30 days

Can set interval to 0 (1 year), effectively turning off snapshots.

AWR Reporting

- Use scripts or Enterprise Manager
- AWRRPT.SQL creates text or HTML report
- "statspack-like" report
  - *For specified time period*

```
SQL> @%ORACLE_HOME%\rdbms\admin\awrrpt

Specify the Begin and End Snapshot Ids
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
Enter value for begin_snap: 85
Begin Snapshot Id specified: 85

Enter value for end_snap: 90
End   Snapshot Id specified: 90
```

- Load profile
- In-use parameters
- Cache sizes and stats
- Shared pool stats
- Top 5 events
- Wait events
- SQL statement stats
- Tablespace IO
- File IO
- etcetera…

AWR Baselines

- Create baselines for comparisons
- Create during typical performance period

```
SQL> select snap_id, begin_interval_time from dba_hist_snapshot
  2  where to_char(begin_interval_time, 'mmddyyyy') = '07282004';

   SNAP_ID BEGIN_INTERVAL_TIME
---------- ---------------------------------
        91 28-JUL-04 03.46.49.000 PM
        84 28-JUL-04 01.00.31.991 AM
 . . .
        93 28-JUL-04 07.00.37.603 PM
        94 28-JUL-04 09.00.34.111 PM

SQL> exec dbms_workload_repository.create_baseline(start_snap_id=>85, -
>         end_snap_id=>95, baseline_name=>'weekday load')

PL/SQL procedure successfully completed.
```

Find start / end ranges

Specify start / end ranges

ADDM

- Automatic Database Diagnostic Monitor
  - Some say “Adam”
- Uses AW Repository to produce “findings”
  - Tuning recommendations
  - Goal: Improve overall system throughput
- Executes automatically after every snapshot
  - DBA can use to tune database
- Use OEM or DBMS_ADVISOR to configure and perform on-demand analysis
  - See supplied script $ORACLE_HOME/rdbms/admin/addmrpt.sql

ADDM Report Excerpt

```
SQL> @%ORACLE_HOME%\rdbms\admin\addmrpt
Enter value for begin_snap: 105
Enter value for end_snap: 107
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
FINDING 2: 81% impact (537 seconds)
-----------------------------------
SQL statements were found waiting for row lock waits.

   RECOMMENDATION 1: Application Analysis, 81% benefit (537 seconds)
      ACTION: Trace the cause of row contention in the application logic.
         Use given blocked SQL to identify the database objects involved.
         Investigate application logic involving DML on these objects.
      RATIONALE: The SQL statement with SQL_ID "4x1uh87xsnu0m" was
         blocked on row locks.
         RELEVANT OBJECT: SQL statement with SQL_ID 4x1uh87xsnu0m
         update dave.t set c1 = 100

   SYMPTOMS THAT LED TO THE FINDING:
      Wait class "Application" was consuming significant database time.
     (81% impact [538 seconds])
```

Percent of elapsed time within report snapshot range

Enabling ADDM…

- _ADDM_AUTO_ENABLE = TRUE to enable
  - FALSE to turn off
    - Oracle does not recommend
    - Must do your research on ramifications
  - Enabled by default

```
SQL> select ksppinm,ksppdesc,ksppstvl
  2  from x$ksppi,x$ksppcv
  3  where x$ksppi.indx = x$ksppcv.indx
  4  and x$ksppi.ksppinm like '_addm%';

KSPPINM                        KSPPDESC                  KSPPSTVL
------------------------------ ------------------------- --------
_addm_auto_enable              governs whether ADDM      TRUE
                                gets run automatica
                               lly after every AWR
                               snapshot
```

…Enabling ADDM

- STATISTICS_LEVEL = TYPICAL
- Adjust advisor parameter DBIO_EXPECTED*
  - Average Single Block Read Time
  - Set with DBMS_ADVISOR package

SQL Tuning Advisor…

- STA identifies issues and provides recommendations
  - Missing stats
  - Different plan
  - Different access paths and object
  - Restructuring SQL
- Oracle claims
  - "automates entire tuning process"
  - "replaces manual SQL tuning"

Could be a bit overstated…

…SQL Tuning Advisor

- Two interfaces
  - Enterprise Manager
  - API with supplied PL/SQL package
    - DBMS_SQLTUNE
- Three step process
  - Create tuning “task”
    - Stored persistently in data dictionary
  - Execute tuning task
    - Executes the optimizer in “Plan Tuning Analysis mode”
    - Gathers additional info
    - Creates a “profile” when warranted (more later)
  - Display report
    - Contains “findings” and possible solutions

SQL Tuning Advisor Demo (1)

- Create a “tuning task”
- Execute the task
  - Invokes cost-based optimizer

Privileges req’d: ADVISOR, SELECT_CATALOG_ROLE, EXECUTE on DBMS_SQLTUNE

Source of SQL can be ADDM

```
SYSTEM> declare
  2    task_id varchar2(30);
  3  begin
  4    task_id := dbms_sqltune.create_tuning_task(
  5                 sql_text=>'update dave.t set c1 = 6000',
  6                 task_name=>'test2');
  7
  8    dbms_sqltune.execute_tuning_task('test2');
  9  end;
 10  /

PL/SQL procedure successfully completed.
```

SQL Tuning Advisor Demo (2)

- Display the report

```
SYSTEM> set long 1000
SYSTEM> set linesize 120
SYSTEM> select dbms_sqltune.report_tuning_task('test2') from dual;

DBMS_SQLTUNE.REPORT_TUNING_TASK('TEST2')
--------------------------------------------------------------------------------
GENERAL INFORMATION SECTION
-------------------------------------------
Tuning Task Name   : test2
Scope              : COMPREHENSIVE
Time Limit(seconds): 1800
Completion Status  : COMPLETED
Started at         : 08/02/2004 17:49:29
Completed at       : 08/02/2004 17:49:30

--------------------------------------------------------------------------------
```

1st section of report gives general information

SQL Tuning Advisor Demo (3)

```
-------------------------------------------------------------------------------
FINDINGS SECTION (1 finding)
-------------------------------------------------------------------------------

1- Statistics Finding
---------------------
  Table "DAVE"."T" was not analyzed.

  Recommendation

DBMS_SQLTUNE.REPORT_TUNING_TASK('TEST2')
-------------------------------------------------------------------------------
 Consider collecting optimizer statistics for this table.
 execute dbms_stats.gather_table_stats(ownname => 'DAVE', tabname => 'T',
      estimate_percent => DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR
```

"Findings" section identifies possible issues

"Recommendations" section gives possible solutions

Clean up

```
SYSTEM> exec dbms_sqltune.drop_tuning_task('test2');

PL/SQL procedure successfully completed.
```

SQL Profiles

- Additional statistics and information for CBO
- Changes normal decision of CBO at run-time
  - Sort of like a hint
  - Without specifying the hint
  - Useful for packaged apps?
    - e.g. Peoplesoft, SAP
- Created - *when warranted* - by SQL Tuning Advisor
  - When run in COMPREHENSIVE mode
  - Stored in dictionary
- Don't confuse with stored outlines
  - Profile can change over time
  - How? Details???

SQL Profile Demo (1)

- Run STA in comprehensive mode on query

```
SQL Text: select /*+ full(big) */ * from big where object_id = 1

-------------------------------------------------------------------------------
FINDINGS SECTION (1 finding)
-------------------------------------------------------------------------------

1- SQL Profile Finding (see explain plans section below)
--------------------------------------------------------
  A potentially better execution plan was found for this statement.

  Recommendation (estimated benefit: 99.66%)
  ------------------------------------------
    Consider accepting the recommended SQL profile
```

Forcing a full scan for this test

Mystery: how to see the better execution plan – *before* accepting?

STA created a profile

SQL Profile Demo (2)

- Accept the profile

```
DAVE> var x varchar2(30)
DAVE> begin
  2  :x := dbms_sqltune.accept_sql_profile
       (task_name=>'profile_test',
       name=>'profile_test_profile');
  3  end;
  4  /

PL/SQL procedure successfully completed.
```

Doc says this is a procedure – NOT.

Can limit use by session with "CATEGORY" attributes

SQL Profile Demo (3)

- Test – is profile used?

```
DAVE> set autotrace traceonly
DAVE> select /*+ full(big) */ * from big where object_id = 1;

no rows selected


Execution Plan
----------------------------------------------------------
   0      SELECT STATEMENT Optimizer=ALL_ROWS (Cost=2 Card=1
   1    0   TABLE ACCESS (BY INDEX ROWID) OF 'BIG' (TABLE) (C

   2    1     INDEX (RANGE SCAN) OF 'BIG_IDX' (INDEX) (Cost=1
```

YES. Profile is used.

```
DAVE> exec dbms_sqltune.drop_sql_profile('profile_test_profile')

PL/SQL procedure successfully completed.
```

Cleanup

Other Advisors

- SQL Access Advisor (1)
  - "…advice on materialized views, indexes, and materialized view logs"
- Undo Advisor (2)
  - Assists with sizing Undo tablespace
- Redo log Advisor (2)
  - Suggest log size
  - Reduce I/O from checkpointing
- Segment Advisor (2)
  - Fragmentation?
  - Can space be reclaimed?

Other notable undo features:
- Set retention = 0. Oracle will manage
- Eliminate "snapshot too old" errors
  - See *RETENTION GUARANTEE* on CREATE UNDO TABLESPACE
  - Can cause DML to fail

Automatic Optimizer Statistics

- DB installation creates a scheduled daily job

```
SQL> select last_start_date, last_run_duration
  2  from dba_scheduler_jobs
  3  where job_name = 'GATHER_STATS_JOB';

31-JUL-04 10.00.03.904000 AM -04:00
+000000000 00:01:25.814000
```

Part of new “Scheduler” feature. See dbms_scheduler package and DBA_SCHEDULER_JOBS

- Collects only if stats *stale* or missing
- Set STATISTICS_LEVEL = TYPICAL
  - See new auto-monitoring feature
  - Extension of 9i Table Monitoring feature

More Optimizer Stats Features

- Collect data dictionary statistics

```
SYSDBA> exec dbms_stats.gather_dictionary_stats

PL/SQL procedure successfully completed.
```

- Auto parallel for DBMS_STATS

```
DAVE> exec dbms_stats.gather_schema_stats(user, -
>       cascade=>TRUE, -
>       method_opt=> 'FOR ALL INDEXED COLUMNS', -
>       options=>'GATHER STALE', -
>       degree=>dbms_stats.auto_degree)

PL/SQL procedure successfully completed.
```

Rule-Based Optimization

- Still exists, but not supported
- OPTIOMIZER_MODE considerations
  - CHOOSE or RULE causes ALERT.LOG warning
  - ALL_ROWS is the default
- CHOOSE and RULE hints not supported
- Migrate existing apps to CBO
  - Helpful documentation:
    - Oracle Database Upgrade Guide
    - Metalink Doc 189702.1
    - Oracle Database Performance Tuning Guide

Finding Wait Events…

- Improved V$ views related to wait events
- V$SESSION_WAIT_CLASS
  - Wait “classes” for grouping

```
DAVE> select wait_class, sum(total_waits), sum(time_waited)
  2  from v$session_wait_class
  3  group by wait_class order by sum(time_waited) desc;

WAIT_CLASS           SUM(TOTAL_WAITS) SUM(TIME_WAITED)
-------------------- ---------------- ----------------
Idle                           524213        213386445
System I/O                     125248            51796
User I/O                         7393            14291
Other                           15009            10487
Configuration                  320966              460
Concurrency                     75820              222
Commit                             51               29 . . .
```

…Finding Wait Events

- V$SESSION_WAIT
  - New columns eliminate need to join for cause of wait
- V$ACTIVE_SESSION_HISTORY keeps history of wait events for session
  - Reduces problem of hard-to-capture, transient wait events
  - Reduces re-runs to find problems
- More views to research…

More Features

- Self-Tuning Checkpointing
  - Set FAST_START_MTTR_TARGET > 0
- Enhanced trace
  - See the DBMS_MONITOR package
- Sorted Hash Clusters
  - Rows kept in sort order
- Default dynamic sampling level now 2, not 1
  - More aggressive than 9i default
- Auto UNDO Tuning
- Shared Server enhancements
  - MTS_ parameters obsolete
- Resource Manager Enhancements
- Flush buffer cache
  - Why?

```
SYSDBA> alter system flush buffer_cache;

System altered.
```

Watch Out For…

- ASM instance
  - Load on small servers
- OEM DBControl
  - Load on small servers
  - Many SYSMAN sessions
- Bug?:

ASM instance killed performance of my test platform (W2000, 512M memory)

OEM Windows service sometimes fails to start on my machine

Will not be able to start DB. Backup SPFILE first!!!
`SYSDBA> create pfile from spfile;`

```
SQL> alter system set sga_target=100m scope=both;
System altered.

SQL> show parameter sga_
NAME TYPE VALUE
------------------------------ ----------- --------
sga_max_size                   big integer 200M
sga_target                     big integer 16M
```

…Watch Out For

- Performance hit or usefulness of all the automatic stuff?
  - I'm *not* saying you won't benefit, but *test*, especially
    - AWR (repository) and ADDM (diagnostics)
    - Automatic optimizer statistics collection
    - Automatic SGA Management
    - New OEM

Performance & Tuning Feature Summary

- At the start, we may be concerned about all new automatic stuff
  - Overhead?
  - Losing control?
- Still some mysteries
  - Lots of testing to do
- Visit www.skillbuilders.com
  - Download complete presentation (with some notes)
  - Download demo scripts
  - I'll publish more papers as I learn more

AWR, Snapshots, Scheduled jobs, etcetera

We didn’t have time for…

- Default permanent tablespace
- SYSAUX Tablespace
- Bigfiles
- Rename tablespace
- Temporary tablespace groups
- Drop database
- Shrink objects

The End

- Flashback
- ASM
- Job Scheduler
- RMAN
- Tuning
- More….

- Thanks for listening!
- Please call SkillBuilders for your training needs:
  - Oracle
  - Web Services
  - XML
  - Java
- 888-803-5607