


Includes Release 2
Features!

Oracle10g New Features for DBA's

Presented by Dave Anderson of SkillBuilders
www.skillbuilders.com
888-803-5607
December 8, 2005


SKILLBUILDERS



Author: Dave Anderson, SkillBuilders

Release Date: October, 2004

Last Revision: January, 2005

3

Introducing Data Pump

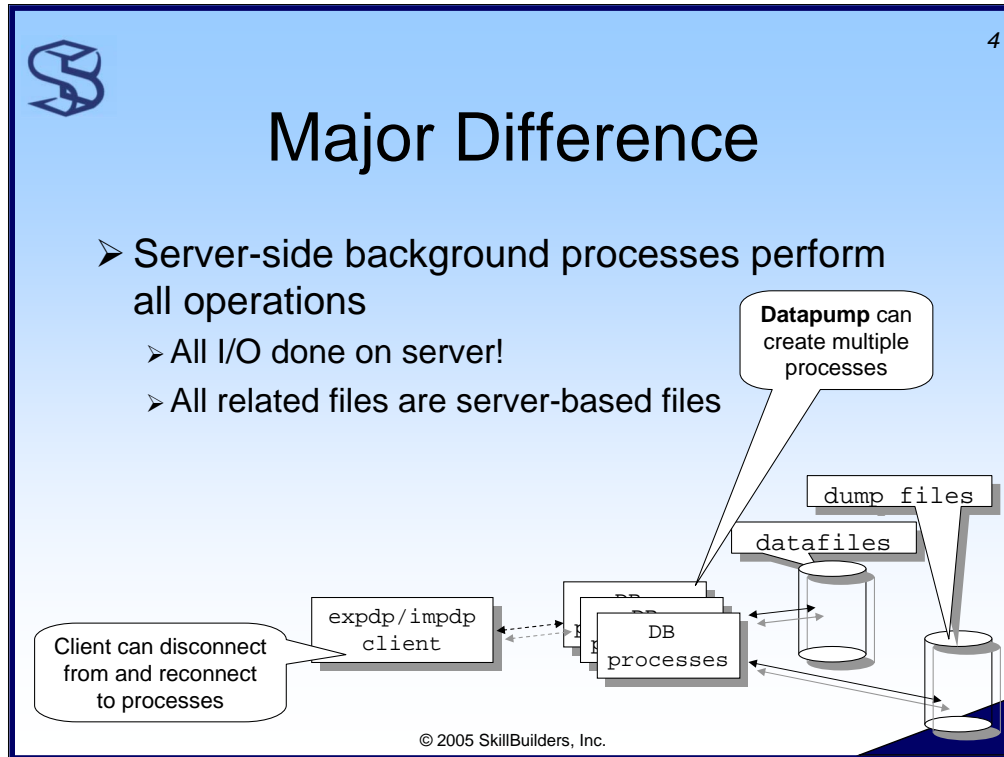
- Next generation export / import
- New utility for high-speed data movement
- Data moved with direct path API or external tables
 - 10g supports export with external tables
- Some similarities to original export / import
 - To / from OS “dump” files
 - Similar “modes”
 - FULL, SCHEMA, TABLE, etc.
 - Some same parameters

© 2005 SkillBuilders, Inc.

Datapump is considered the next generation of import / export utilities. Brand new with Oracle10g, its purpose is to provide better and faster movement of data into and out of the database.

There are some similarities to the original import and export utilities. (We will discuss the major differences in the following section in this lesson.):

- Like the original export, it can read database data and write to an OS dump file. And like the original import, it can read these proprietary format dump files and write them to the database.
- And it supports similar “modes”, which control the scope of data to be moved. (SCHEMA is similar to the old OWNER mode.)
- Today, you use the “imp” and “exp” host commands to invoke the original import and export utilities. Datapump simply adds the suffix “dp” to the command names.
- You will still use the concept of a parameter file (called a “PARFILE”).
- Finally, datapump supports a direct path mode of data movement, greatly increasing the speed of the operations. But, as you will learn, the utility chooses the mode, and it never uses the original, slow SQL mode!



One of the most important differences and improvements over the original utility is that datapump I/O operations run completely on the server. The datapump clients are lightweight clients that simply create calls to the datapump API, `DBMS_DATAPUMP`. So, even if you are connected to the database across a network, you will *not move* the data across the network.

Administrators and programmers will need to become familiar with the new datapump parameters. Some of these parameters simply replace existing exp/imp parameters, while others add new functionality. For example:

- The datapump parameter “`SQLFILE`” is similar to and replaces the “`SHOW`” parameter.
- The datapump parameter `NETWORK_LINK` let’s you export data from a *remote* database (writing to a dump file accessible to local server). This parameter is not supported in the original export utility.

You will see examples of these parameters as we progress through this lesson. Refer to the section “How Data Pump Export Parameters Map to Those of the Original Export Utility” in Chapters 2 and 3 of the **Oracle10g Database Utilities** manual for a comparison of old and new parameters.

Unlike the original import utility, `datapump import` can use the direct path technique, increasing the performance of import operations.

5

Export Example.

Must use a directory; R2 provides default directory DATA_PUMP_DIR

```
SQL> create directory datapump_dump as '/home/oracle';
Directory created.

SQL> grant read, write on directory datapump_dump to dave;
Grant succeeded.
```

Datapump creates a table in your schema with this job name; dropped at EOJ

```
# PARFILE - Test datapump parallel export
job_name=datapump_parallel
schemas=dave
parallel=2
directory=datapump_dump
dumpfile=datapump_parallel1%U.dmp, datapump_parallel2%U.dmp
logfile=datapump_parallel.log
```


Have 1 dump file for each DOP, preferably on separate disks.

© 2005 SkillBuilders, Inc.

This example demonstrates a simple schema-mode export operation, using parallel processing. Refer to the supplied script `datapump_export_parallel.par` for a working copy of this example. Steps include:

1. Creating a directory if one is not already available. The directory is used by datapump operations as the location for dump files, log files and SQLFILES (optionally created by datapump import to show the DDL in the dump file – like the SHOW parameter with original import). Different directories can be used. For example, to write the log file in a separate directory, code `LOGFILE=directory:filename`. Release 2 provides a default directory `DATA_PUMP_DIR` which, on my Linux server, points to `/u01/app/oracle/product/10.2.0/db_2/rdbms/log/`.
2. Creating a parameter file. This example uses the following parameters:
 - `JOB_NAME` – Optional parameter. The job name is used when monitoring the job (see `DBA_DATAPUMP_JOBS`) and reattaching to the job (see the `ATTACH` parameter on the `expdp` and `impdp` commands.) The `JOB_NAME` parameter is also important because datapump creates a table in your schema that helps control the datapump process. This table is dropped when the job completes successfully. Be sure to choose a name that does not conflict with an existing table name in your schema.

Notes for this slide continue on the next page...

6

...Export Example

```
LINUX> host expdp dave/dave@springsteen_linux3 parfile=datapump

Export: Release 10.1.0.2.0 - Production on Tuesday, 19 October,

Copyright (c) 2003, Oracle. All rights reserved.

Connected to: Oracle Database 10g Enterprise Edition Release 10.
With the Partitioning, OLAP and Data Mining options
FLASHBACK automatically enabled to preserve database integrity.
Starting "DAVE"."DATAPUMP_PARALLEL": dave/*****@springsteen_
Estimate in progress using BLOCKS method...

*****
Dump file set for DAVE.DATAPUMP_PARALLEL is:
/home/oracle/datapump_parallel101.dmp
/home/oracle/datapump_parallel201.dmp
Job "DAVE"."DATAPUMP_PARALLEL" successfully completed at 16:35
```

© 2005 SkillBuilders, Inc.

Here we see the actual expdp command invocation, and the first few and last few messages issued.

Supplemental Notes

Note the FLASHBACK-related message. I was hoping that this indicated expdp *always* created a consistent export, but, alas, Oracle support says this is not so. Support says to use the FLASHBACK parameter to get a consistent export; we will look at an example later in this lesson.

However, Oracle Datapump Development teams statement to me on 9-OCT-2004 concerning the flashback message is confusing:

“Currently this message simply means we will use flashback if internally detect a need for it...In future releases the message condition will be limited to schema and full db exports via expdp if minimal supplemental logging is present.”

7

Import Example

PARFILE=

```
# Test non-parallel datapump import
schemas=dave
dumpfile=datapump_noparallel.dmp
directory=datapump_dump
logfile=datapump_import.log
job_name=datapump_import
table_exists_action=replace
```

```
SQL> host impdp dave/dave@springsteen_linux3 parfile=datapump_i

Import: Release 10.1.0.2.0 - Production on Thursday, 14 0

Copyright (c) 2003, Oracle. All rights reserved.

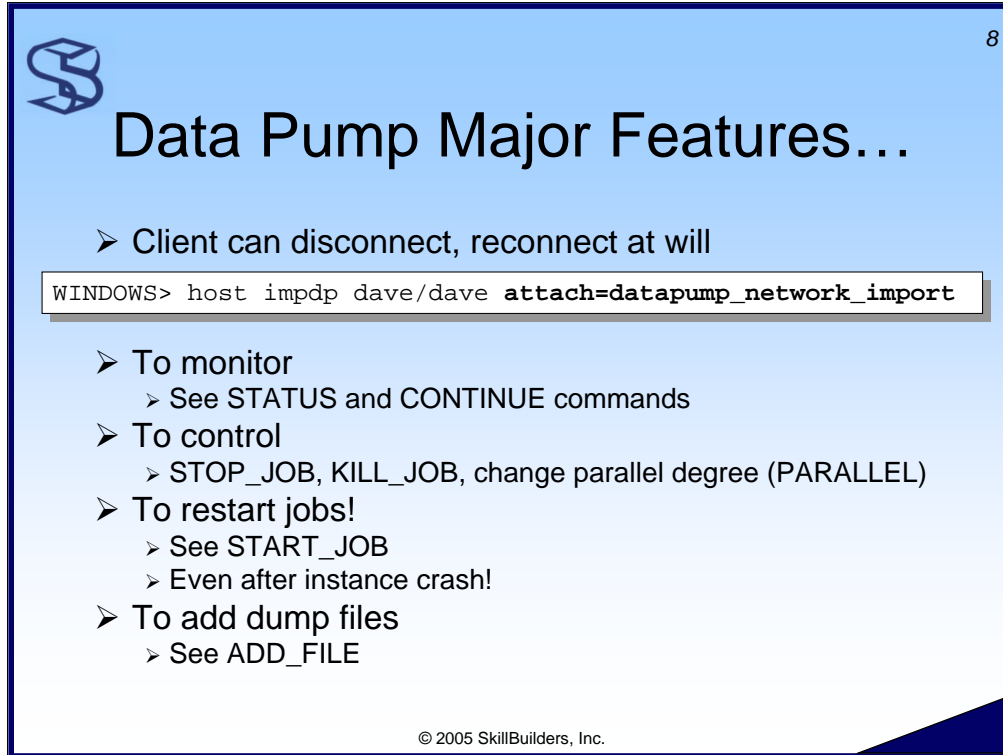
Connected to: Oracle Database 10g Enterprise Edition Rele
With the Partitioning, OLAP and Data Mining options
Master table "DAVE"."DATAPUMP_IMPORT" successfully loaded
Starting "DAVE"."DATAPUMP_IMPORT": dave/*****@springs
Processing object type SCHEMA_EXPORT/USER

© 2005 SkillBuilders, Inc.
```

This example demonstrates a simple import operation.

Also note the use of the TABLE_EXISTS_ACTION parameter; REPLACE means if the table exists, overwrite the data.

Refer to the supplied script datapump_import_noparallel.par for a working copy of this example.



The slide features a blue background with a white border. In the top left corner is the Oracle logo. In the top right corner is the number '8'. The main title is 'Data Pump Major Features...'. Below the title is a list of features. A screenshot of a Windows command prompt shows the command: 'host impdp dave/dave attach=datapump_network_import'. Below the screenshot is another list of features. At the bottom center is the copyright notice: '© 2005 SkillBuilders, Inc.'

Data Pump Major Features...

- Client can disconnect, reconnect at will

```
WINDOWS> host impdp dave/dave attach=datapump_network_import
```


- To monitor
 - See STATUS and CONTINUE commands
- To control
 - STOP_JOB, KILL_JOB, change parallel degree (PARALLEL)
- To restart jobs!
 - See START_JOB
 - Even after instance crash!
- To add dump files
 - See ADD_FILE

© 2005 SkillBuilders, Inc.

Like original import / export, datapump provides command-line client software (impdp and expdp) that is run at the OS prompt.

Unlike the original utility, datapump clients can disconnect from the job (i.e. the background process performing the work). The client can then be reconnected to the job – see the ATTACH parameter on the impdp command. Reason to reattach to a job might be:

- To monitor a long-running job. You may kick off a job, disconnect (ctrl+C followed by the EXIT_CLIENT command) and go to lunch. Upon your return, reconnect and check the status with the STATUS command.
- To control a job. You may determine that the job is taking too much resources from the server. In this case, you can dynamically decrease the degree of parallelization with the PARALLEL command. (DOP can also be dynamically increased.) Or, you might temporarily halt the job with STOP_JOB and restart it later with START_JOB.
- *Yes, datapump jobs that fail can be restarted – as long as they get past the “DEFINE” stage – even if the instance crashes. I’ve tested this (in a test environment) and it worked as advertised.*
- If a job runs fails due to an exhausting the space in the dump file, you can add dump files for the job with the ADD_FILE command and restart the job.



9

...Data Pump Major Features...

- Filtering during export and import
 - Data, metadata or both
 - See CONTENT parameter
 - Filter by object type, name
 - Wildcard support
 - See INCLUDE and EXCLUDE

```
include=schema: '='SCOTT' ', table: '='DEPT' "
```
 - Filter data
 - Supports WHERE, ORDER BY, more
 - Even on import
 - See QUERY

```
query=scott.dept: "WHERE dname = 'SALES' "
```

© 2005 SkillBuilders, Inc.

Datapump import and export utilities provide many techniques for filtering, i.e. what gets imported or exported. These include:

Using CONTENT parameter to import or export just the data, just the metadata, or both.

Using INCLUDE and / or EXCLUDE to control what type of objects are affected, and / or the names of objects affected by the operation. This is very powerful. For example, to exclude all of Dave's objects, code

```
exclude=schema: '='DAVE' "
```

To include only the table SCOTT.DEPT, code the parameter

```
include=schema: '='SCOTT' ', table: '='DEPT' "
```

Note that INCLUDE and EXCLUDE are mutually exclusive parameters.

Using the QUERY parameter to filter the data. For example, to load just the sales department, code:

```
query=scott.dept: "WHERE dname = 'SALES' "
```

Notes continue on the next page...

10

...Data Pump Major Features...

- Import remapping capabilities
 - Change name of datafiles
 - See REMAP_DATAFILE

```
remap_datafile='"+GROUP1/linux3/datafile/users.259.1": "C:\ORACLE\ORADATA\DAVE10G\DATAFILE\users_from_linux.dbf" '
```

- Change owner
 - See REMAP_SCHEMA
- Change tablespace
 - See REMAP_TABLESPACE


- Reuse existing datafiles
 - See REUSE_DATAFILES

```
remap_schema=dave:dave_from_linux
```

© 2005 SkillBuilders, Inc.

Datapump provides three import parameters for changing the attributes of objects in the dump file:

- REMAP_DATAFILE – Change the name or path of the datafile. In addition to simple renaming and path changes when importing to a different server, this provides a method of importing from different platforms. Refer to the supplied script `datapump_import_tablespace_remap.par` for an example of this parameter.
- REMAP_SCHEMA – Similar to the `TOUSER` parameter supplied with original import, this provides the ability to import into a different schema from which it came. Note that if import creates a new user the password is invalid (to prevent unauthorized logins) and must be changed before login.
- REMAP_TABLESPACE – Place all imported objects in the specified target tablespace.
- The `REUSE_DATAFILES` parameter, when set to “Y”, will reinitialize and reuse any existing datafiles when the `CREATE TABLESPACE` statement is executed. The default is “N”; if the datafile exists, in this case an error is generated from the `CREATE TABLESPACE`, but the job continues.



11

...Data Pump Major Features...

- Flashback support
 - Export at previous point in time
 - Get a *consistent* export
 - See FLASHBACK_SCN and FLASHBACK_TIME parameters


```
flashback_time="to_timestamp('16-10-2004 13:00:00',  
                             'DD-MM-YYYY HH24:MI:SS')"
```

© 2005 SkillBuilders, Inc.

Datapump supports flashback query.

On export, specify the SCN or time you want the data to be consistent with. Of course, all flashback query rules apply, e.g. sufficient UNDO must be available, you cannot flashback beyond last DDL time, etcetera.

Flashback is even supported on *network imports*! (More on network imports later in this lesson.)

12

...Data Pump Major Features

- Estimate space before exporting
 - See ESTIMATE
 - BLOCKS or STATISTICS
 - ESTIMATE_ONLY
 - Use “Y” to skip the actual export
- Table compression support
- R2 adds to expdp
 - COMPRESSION=METADATA_ONLY
 - SAMPLE
 - FILESIZE

© 2005 SkillBuilders, Inc.

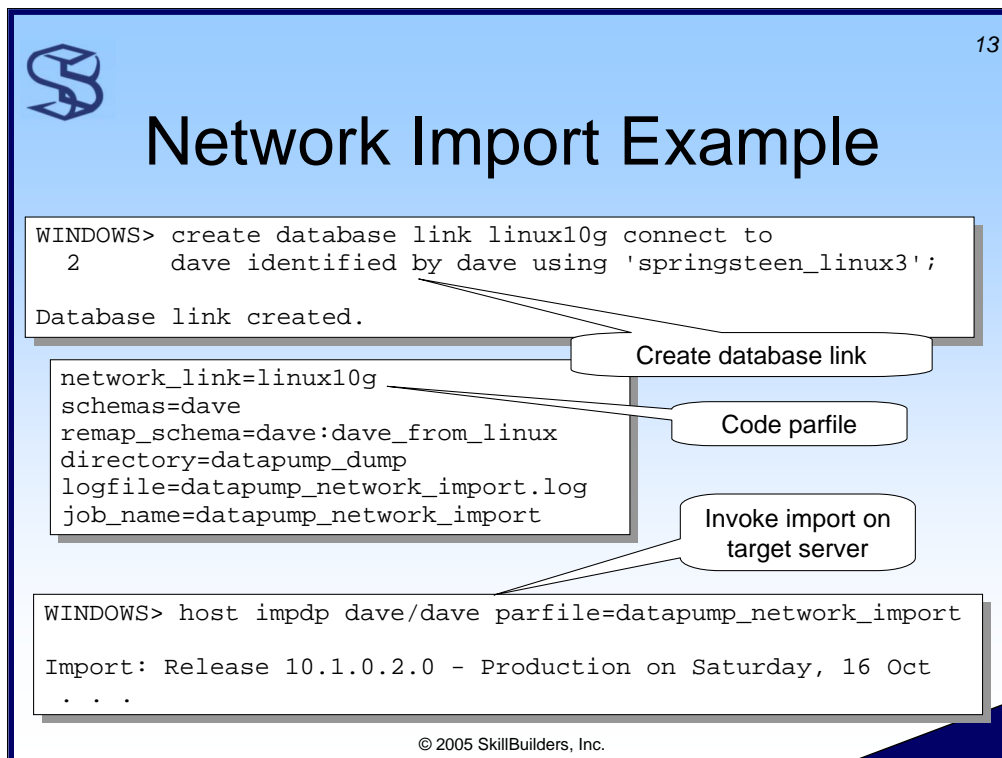
Datapump exports always estimate the space required for a dump file. The log will contain messages like:

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
. estimated "DAVE"."WEBHITS_DETAIL"          39 MB
. estimated "DAVE"."SALES_PER_DAY"          0 KB
Total estimation using BLOCKS method: 43.06 MB
```

There are two methods available for estimating the space: BLOCKS and STATISTICS. Blocks simply multiples “the number of database blocks used by the target objects times the appropriate block sizes.” Statistics uses the object’s statistics; you will want to be sure you recently updated statistics with DBMS_STATS package.

You can ask datapump to provide the estimate and skip the export. This is done with the ESTIMATE_ONLY=Y parameter. You will need to remove or comment out the DUMPFILE parameter if you use ESTIMATE_ONLY=Y.

Though not well documented, the **Oracle10g Utilities guide** does include this statement: “Use the new Data Pump Export and Import utilities to enable compression during import.” (Original import does not use direct path method – it uses INSERT – and thus imported data is not in compressed format.)



13

Network Import Example

```
WINDOWS> create database link linux10g connect to
2      dave identified by dave using 'springsteen_linux3';

Database link created.
```

Create database link

```
network_link=linux10g
schemas=dave
remap_schema=dave:dave_from_linux
directory=datapump_dump
logfile=datapump_network_import.log
job_name=datapump_network_import
```

Code parfile

Invoke import on target server

```
WINDOWS> host impdp dave/dave parfile=datapump_network_import

Import: Release 10.1.0.2.0 - Production on Saturday, 16 Oct
. . .
```

© 2005 SkillBuilders, Inc.

In this example we see the steps required for a network import operation. First, define a database link. Next create a parameter file. Code the database link name in the NETWORK_LINK parameter. In this example I request schema “dave” from the remote server; since I already have a user “dave” on the local server, I use the REMAP_SCHEMA parameter to import it into another schema. Lastly, issue the “impdp” command to invoke the network import.

Remember that if import creates a new user the password is invalid (to prevent unauthorized logins) and must be changed before login.

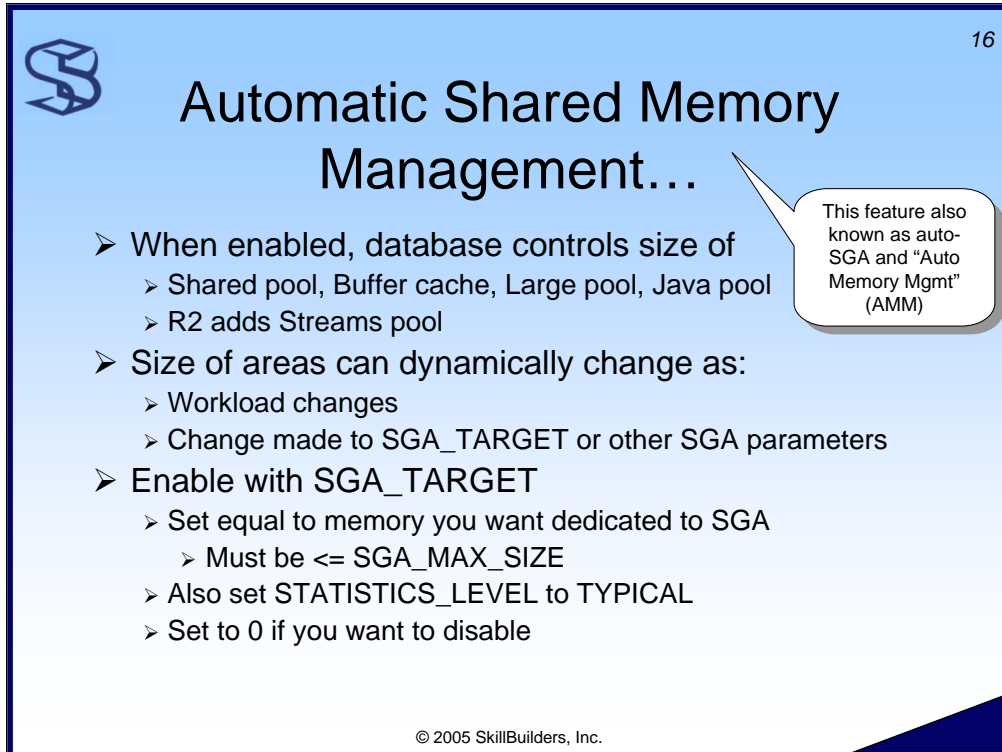
Refer to the supplied script datapump_network_import.par for a working copy of this example.



Data Pump Summary

- Next generation import / export utility
- Server processes move data
- Supports parallel operations





16

Automatic Shared Memory Management...

- When enabled, database controls size of
 - Shared pool, Buffer cache, Large pool, Java pool
 - R2 adds Streams pool
- Size of areas can dynamically change as:
 - Workload changes
 - Change made to SGA_TARGET or other SGA parameters
- Enable with SGA_TARGET
 - Set equal to memory you want dedicated to SGA
 - Must be <= SGA_MAX_SIZE
 - Also set STATISTICS_LEVEL to TYPICAL
 - Set to 0 if you want to disable

This feature also known as auto-SGA and "Auto Memory Mgmt" (AMM)

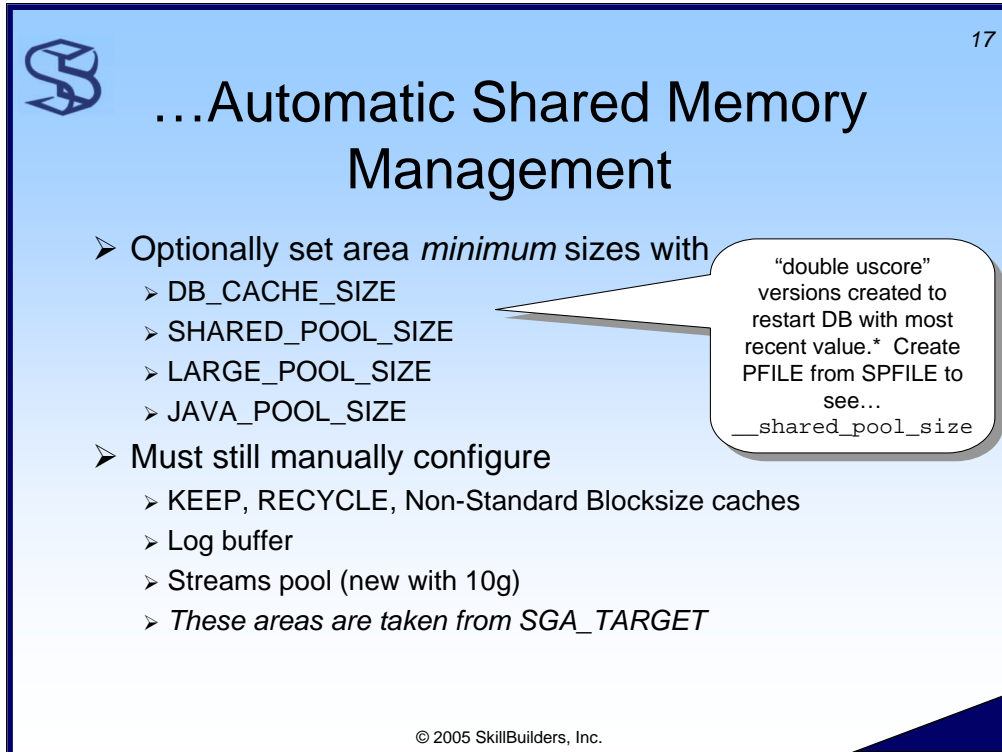
© 2005 SkillBuilders, Inc.

The automatic shared memory management (a.k.a. auto-SGA) feature simplifies SGA management by automatically adjusting size of SGA memory areas in response to the current workload. The areas managed by the auto-SGA feature are:

- Buffer cache (DB_CACHE_SIZE)
- Shared pool (SHARED_POOL_SIZE)
- Large pool (LARGE_POOL_SIZE)
- Java pool (JAVA_POOL_SIZE)

To enable the auto-SGA feature:

- Set the initialization parameter SGA_TARGET equal to the amount of memory you want to dedicate to the instance's SGA. The value must be less than or equal to SGA_MAX_SIZE (which is the upper limit) and should be less than the physical memory on your server (to avoid swapping of the SGA by the operating system). Set SGA_TARGET=0 to disable the auto-SGA feature.
- Set the initialization parameter STATISTICS_LEVEL=TYPICAL



17

...Automatic Shared Memory 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)
 - *These areas are taken from SGA_TARGET*

“double underscore” versions created to restart DB with most recent value.* Create PFILE from SPFILE to see...
`__shared_pool_size`


© 2005 SkillBuilders, Inc.

You can request minimums for the areas controlled by auto-SGA by setting the minimum value in the related initialization parameter.

The instance will remember the last value it used by putting double-underscore versions of the relevant initialization parameter in the SPFILE when shutting down the database. (* Source: **Oracle10g Concepts**, sort of. The concepts manual talks about “Persistence of Automatically Tuned Values” which says the database will remember the sizes of the areas.... However, it does not say how; my testing – including CREATE PFILE FROM SPFILE to view the text contents of the SPFILE – leads me to the conclusion that this is accomplished with the double-underscore parameters.)

The following areas are not controlled by Automatic SGA Management; however, they do count when calculating SGA_TARGET.

- Log buffer (LOG_BUFFER)
- Other buffer caches, such as KEEP (DB_KEEP_CACHE_SIZE), RECYCLE (DB_RECYCLE_CACHE_SIZE), and the non-standard block size caches (e.g. DB_2K_CACHE_SIZE,

18

Introducing AWR...

- Automatic Workload Repository
 - Collects and stores performance statistics
 - Automatically
 - Like an ongoing lightweight statspack
 - Don't have to replay workload to diagnose problem
- "AWR forms the foundation for all self-management functionality of Oracle Database"
 - Source: Concepts manual


© 2005 SkillBuilders, Inc.

The Automatic Workload Repository is an integrated repository of performance data, somewhat like an ongoing collection of statspack data. AWR collects:

- Object usage statistics
- Time model statistics (V\$SYS_TIME_MODEL and V\$SESS_TIME_MODEL views)
- System and session statistics (V\$SYSSTAT and V\$SESSTAT views)
- Top SQL (i.e. most resource intensive SQL)
- Recent session activity, called "Active Session History, or ASH" statistics

An important point to be made about the AWR is that the administrator does not have to replay (i.e. re-execute) an application in order to collect statistics for analysis. If AWR is enabled (and by default it is), the statistics have already be captured. The DBA simply needs to run reports and analyze the output. And if ADDM is used, the analysis will be automated too! (More on ADDM later.)

Refer to Chapter 5 of the **Oracle10g Database Performance Tuning Guide** for more information on AWR.

19

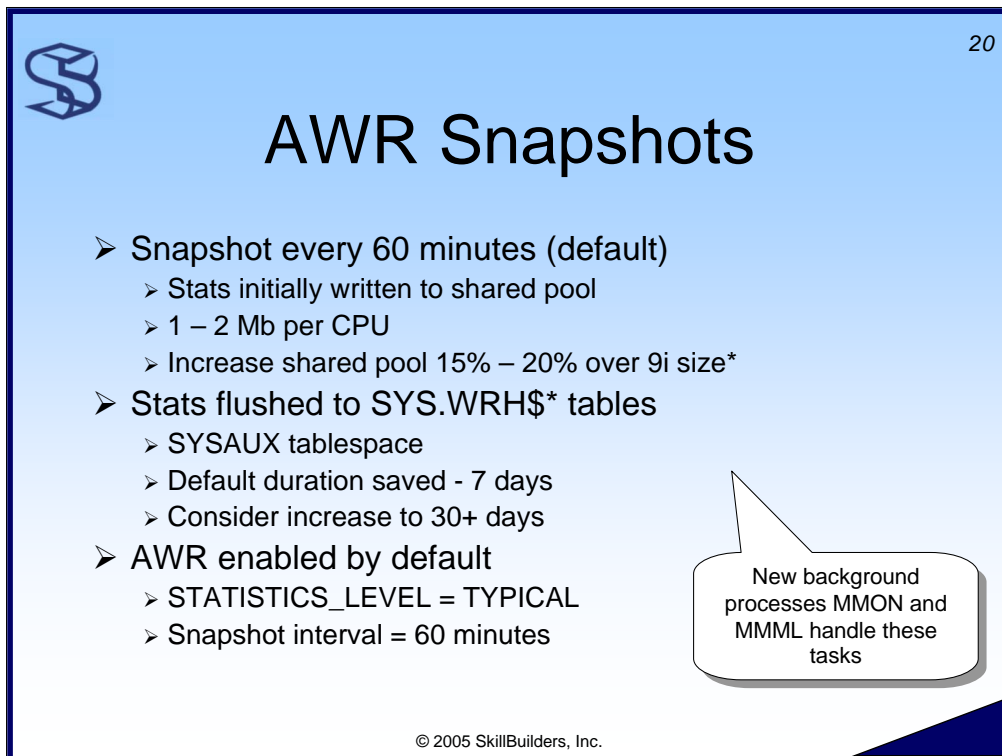
...Introducing AWR

- AWR data feeds “ADDM”
 - Analyze data and provide suggestions
 - A built-in tuning expert
 - More on this later...
- Replaces STATSPACK?
 - That’s the intention
 - You’ll be the judge

© 2005 SkillBuilders, Inc.

Statistics in the AWR are automatically analyzed by the Automatic Database Diagnostic Monitor (ADDM), which produces tuning recommendations. ADDM is presented in the next section of this lesson.

Oracle’s intention is for AWR to replace statspack. From my testing and research, I think this a good feature and it is likely that it will.



20

AWR Snapshots

- Snapshot every 60 minutes (default)
 - Stats initially written to shared pool
 - 1 – 2 Mb per CPU
 - Increase shared pool 15% – 20% over 9i size*
- Stats flushed to SYS.WRH\$* tables
 - SYSAUX tablespace
 - Default duration saved - 7 days
 - Consider increase to 30+ days
- AWR enabled by default
 - STATISTICS_LEVEL = TYPICAL
 - Snapshot interval = 60 minutes


New background processes MMON and MMML handle these tasks

© 2005 SkillBuilders, Inc.

AWR collects the statistics in the shared pool, so the size of your shared pool should be increased from your Oracle9i instance. (Source: Metalink Note PERFORMANCE TUNING USING 10g ADVISORS AND MANAGEABILITY FEATURES 276103.1) You can see the AWR statistics via querying some of the V\$ views.

AWR is enabled by default. It is enabled by setting the STATISTICS_LEVEL parameter equal to TYPICAL (or ALL) and the snapshot interval to some value (it is set to 60 minutes by default).

Every 60 minutes AWR flushes the statistics to the WRH\$_* tables in the SYSAUX tablespace. This gives us (and ADDM) a history of performance data to work with.



21

AWR Reporting...

➤ Query DBA_HIST_* views

```

1 select class, sum(wait_count), sum(time)
2 from dba_hist_waitstat
3 group by class
4* order by avg(time) desc
SQL> /

```

CLASS	SUM(WAIT_COUNT)	SUM(TIME)
undo header	50079	60966
data block	10022	31177
segment header	1689	6889
file header block	525	3353
[etcetera]		

© 2005 SkillBuilders, Inc.

Though you can query the SYS.WRH\$ tables directly, a new series of convenient dictionary views are available for querying the AWR repository. For example, the DBA_HIST_WAITSTAT view contains "historical block contention statistics." It can help detect what is causing waits in your database. This particular query might be useful to identify problems at a high level (like not having enough undo segments, or having "hot blocks" in the SGA).

An interesting point is that with Oracle10g there is now a historical view for many of the V\$ views. For example:

```
DAVE@linux3> desc v$waitstat
```

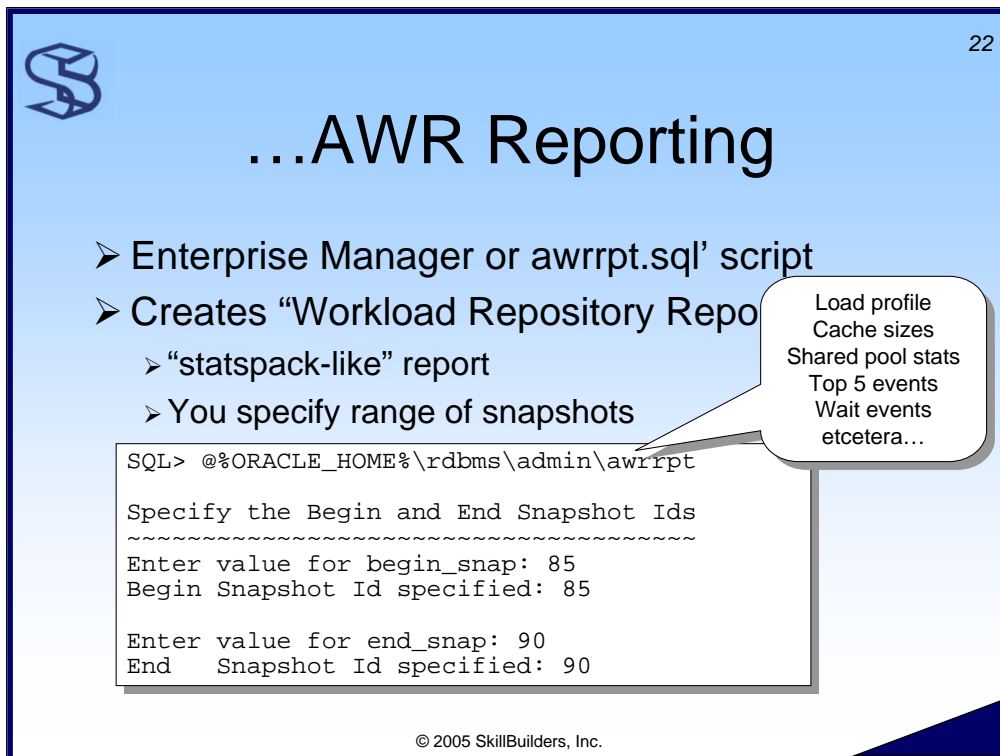
```
Name
```

```
-----
CLASS
COUNT
TIME
```

```
DAVE@linux3> desc dba_hist_waitstat
```

```
Name
```

```
-----
SNAP_ID
DBID
INSTANCE_NUMBER
CLASS
WAIT_COUNT
TIME
```



22

...AWR Reporting

- Enterprise Manager or awrrpt.sql' script
- Creates "Workload Repository Repo"
 - "statspack-like" report
 - You specify range of snapshots

```
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
Cache sizes
Shared pool stats
Top 5 events
Wait events
etcetera...


© 2005 SkillBuilders, Inc.

Beyond just querying the repository, there are many supplied reports that can be created. OEM provides a graphical user interface to access these reports. Within Enterprise Manager, the AWR report is accessed with these steps:

- Administration tab
- Automatic Workload Repository link (near bottom of Administration screen)
- Click on Snapshots link
- Select Beginning Snapshot
- Choose "View Report" in Actions drop-down
- Select "Go"
- Select Ending Snapshot
- Click "Go"
- Click "OK"

Or, use line-mode reports like shown in this slide.

Refer to Chapter 5 of the **Oracle10g Database Performance Tuning Guide** for more information on Oracle's awrrpt.sql script and AWR in general. The sample scripts supplied with this course contain several AWR reports. Refer to the supplied scripts with the prefix "awrrpt_".

23

Tuning Advisors...

➤ 10g R1 provides 7 advisors to assist with tuning

```
DAVE@linux3> select advisor_name
  2  from dba_advisor_definitions;

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

© 2005 SkillBuilders, Inc.

There are seven advisors available with Oracle10g. We have covered a brief introduction into ADDM; next, let's look at the SQL Tuning Advisor...




24

...Tuning Advisors

- 10g R2 adds an SGA Advisor
- Shows improvement in DB Time if SGA_TARGET is increased

```
SYS@linux3> desc v$sga_target_advice
Name
-----
SGA_SIZE
SGA_SIZE_FACTOR
ESTD_DB_TIME
ESTD_DB_TIME_FACTOR
ESTD_PHYSICAL_READS
```

© 2005 SkillBuilders, Inc.

25

Introduction to ADDM...


- Automatic Database Diagnostic Monitor
 - “ADDM” or “Adam”
- Provides tuning advice
 - Goal: Improve overall system throughput
- Uses AWR to produce “findings”
 - Tuning recommendations
 - Hardware changes, Configuration, Schema changes, Application changes
- R2 adds reporting on RMAN, Streams, AQ

© 2005 SkillBuilders, Inc.

The Automatic Database Diagnostic Monitor (ADDM) regularly analyzes data collected by AWR and makes tuning recommendations (called “findings”). As stated in Chapter 6 of the **Oracle10g Database Performance Tuning Guide**, ADDM reports on problems such as:

- CPU bottlenecks
- Undersized Memory Structures
- I/O capacity issues
- High load SQL statements
- High load PL/SQL executions

ADDM makes recommendations on how to solve a problem (see next pages).



26

...Introduction to ADDM

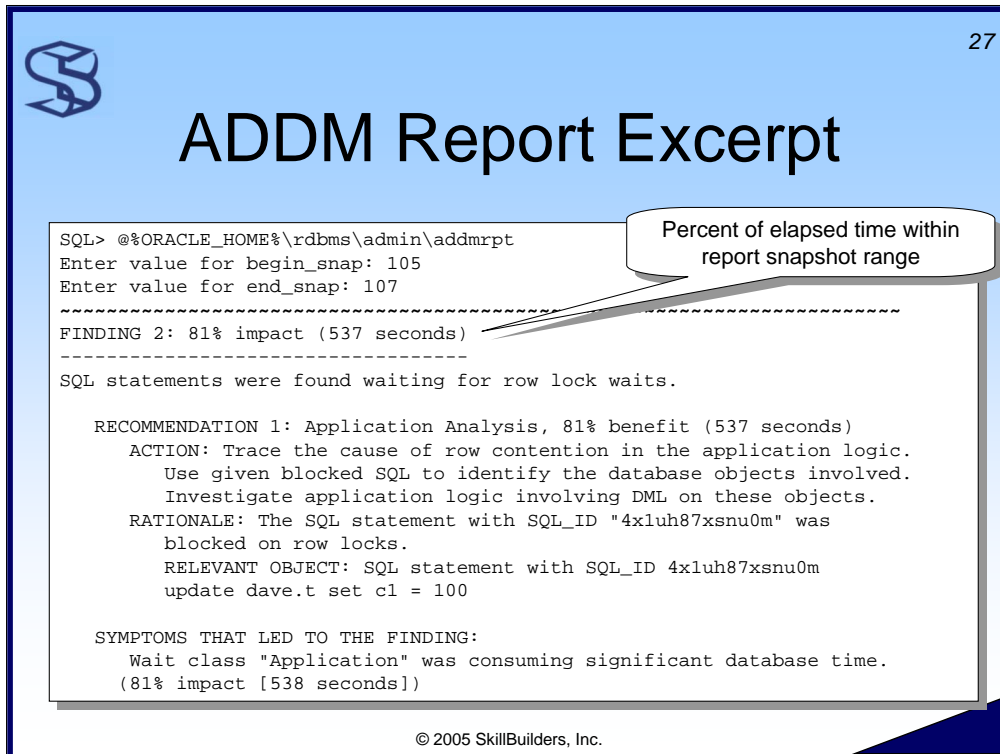
- Executes automatically after every snapshot
 - Stored in DBA_ADVISOR_* views
 - Can see from OEM
- Also run as-needed / on-demand
 - Use OEM or DBMS_ADVISOR package
 - You specify snapshot range
 - See supplied script
\$ORACLE_HOME/rdbms/admin/addmrpt.sql

OEM gives ability to implement suggestions with point and click ease

© 2005 SkillBuilders, Inc.

ADDM executes after every snapshot and stores recommendations in the data dictionary. The results are easily accessed via Oracle Enterprise Manager (OEM) or by querying the DBA_ADVISOR_FINDINGS and DBA_ADVISOR_RECOMMENDATIONS views.

On-demand analysis is easy with OEM or the supplied script addmrpt.sql.



The screenshot shows an Oracle SQL prompt with the following text:

```

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])

```

A callout box points to the finding line with the text: "Percent of elapsed time within report snapshot range".

© 2005 SkillBuilders, Inc.

Here is an excerpt from an ADDM Report. I intentionally forced an UPDATE statement to wait on a row lock during the time period reported on. We see that ADDM finds the blocked update and recommends that we uncover the cause of the contention.

A user needs "ADVISOR" privilege to execute the ADDMRPT.SQL script.

Query V\$SQLAREA to find more information about the SQL statement:

```

SQL> select sql_text from v$sqlarea
       2 where sql_id = '4x1uh87xsnu0m';

```

```

SQL_TEXT
-----

```

```

update dave.t set c1 = 100

```

Caution: If you select a range of snapshots that contain a database shutdown, this script will fail and your SQL*Plus session will unexpectedly terminate.

Caution: If you select a single snapshot (i.e. BEGIN_SNAP and END_SNAP are the same) this script will fail and your SQL*Plus session will unexpectedly terminate.

28

ADDM Recommendations

Oracle Enterprise Manager (SYS) - Performance Finding Details - Netscape

Database Control

Database: linux3 > Advisor Central > Automatic Database Diagnostic Monitor (ADDM) > Performance Finding Details Logged in As SYS

Performance Finding Details

Database Time (minutes) 5.67 Period Start Time Jan 11, 2005 12:30:31 PM Period Duration (minutes) 29.8
Task Owner SYS Task Name ADDM:2410307274_1_2947 Average Active Sessions 0.19

Finding SQL statements consuming significant database time were found.
Impact (minutes) 5.17
Impact (%) 91.27

Recommendations

Select Item(s) and... (Run SQL Tuning Advisor)

Select All | Select None | Show All Details | Hide All Details

Select Details	Category	Benefit (%)
<input checked="" type="checkbox"/>	Hide SQL Tuning	91.27

SQL Text DELETE FROM T WHERE C1 = B1
Action Run SQL Tuning Advisor on the SQL statement with SQL_ID "8ap4xkqja2b73". (Run Advisor Now)

Findings Path

Expand All | Collapse All

Findings

© 2005 SkillBuilders, Inc.

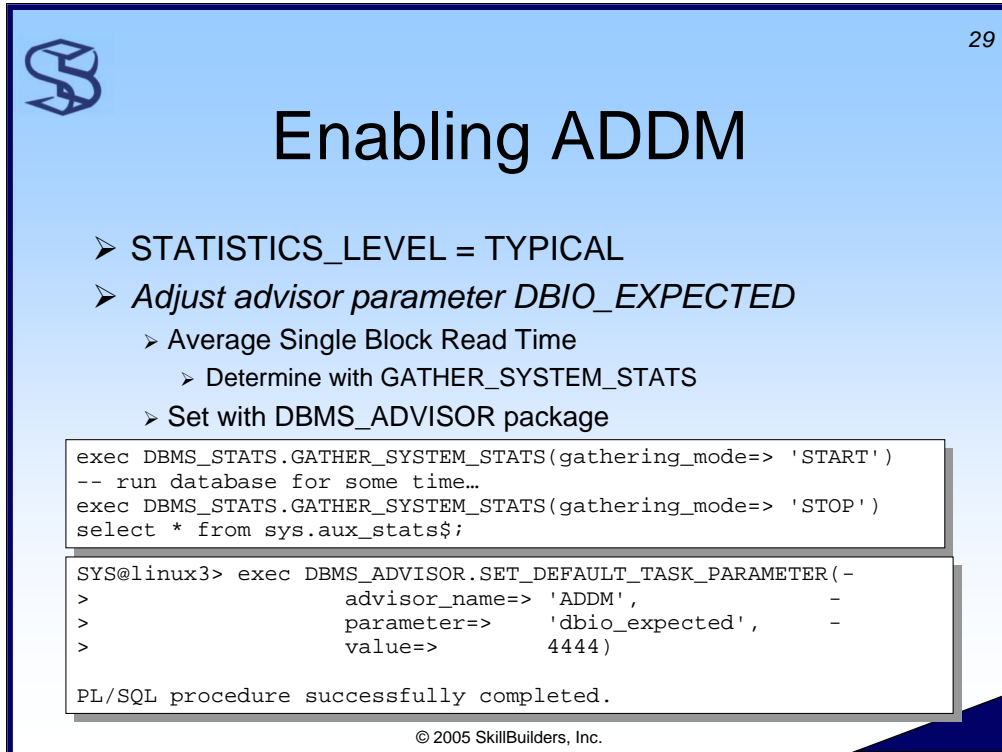
New Enterprise Manager with Diagnostics Pack has access to ADDM

New Enterprise Manager with Tuning Pack has access to SQL Tuning Advisor

Here we see a screen shot from OEM showing ADDM findings.

This “finding” was created by an automatic ADDM execution (ADDM analyzes each snapshot when the snapshot is created). Note that the finding says that the SQL statement consumed “significant database time” and that it recommends the “Advisor” be run.

Running the SQL Advisor may provide recommendations such as add an index or analyze the table. It will also provide a “implement” button allowing you to implement the recommendation with a simple mouse click. (It actually submits a job which runs the recommended SQL.) I will present more information about the Advisors later in this lesson.



29

Enabling ADDM

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

```
exec DBMS_STATS.GATHER_SYSTEM_STATS(gathering_mode=> 'START')
-- run database for some time...
exec DBMS_STATS.GATHER_SYSTEM_STATS(gathering_mode=> 'STOP')
select * from sys.aux_stats$;
```

```
SYS@linux3> exec DBMS_ADVISOR.SET_DEFAULT_TASK_PARAMETER(-
>          advisor_name=> 'ADDM',          -
>          parameter=>    'dbio_expected',  -
>          value=>        4444)
```

PL/SQL procedure successfully completed.

© 2005 SkillBuilders, Inc.


Also set STATISTICS_LEVEL=TYPICAL (or ALL) to enable ADDM. Setting to NONE will eliminate many of the critical statistics.

Also adjust the DBIO_EXPECTED Advisor parameter because ADDM I/O performance analysis relies on it. At the bottom of the ADDM report you will see:

“The analysis of I/O performance is based on the default assumption that the average read time for one database block is 10000 micro-seconds.”

The value should be the “average time it takes to read a single database block in microseconds. Oracle uses the default value of 10 milliseconds.” (Source: **Oracle10g Database Performance Tuning Guide**, Chapter 6, “Setting Up ADDM”) As demonstrated in the example shown above, the DBMS_STATS.GATHER_SYSTEM_STATS procedure will help find the single block read time. Basically, from user SYS, you start the OS system statistics gathering with GATHER_SYSTEM_STATS, run the database under “normal” conditions for a period of time, then stop the gathering. (See the 1st example in this slide.)

Notes for this slide continue on the next page...



30

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”


Jury is still out on these claims

© 2005 SkillBuilders, Inc.

The SQL Tuning Advisor is one of many new Advisors provided with Oracle10g. The SQL Tuning Advisor provides tuning recommendations for one or more SQL statements.

* Source: Metalink Note 262687.1

Refer to Chapter 13 of the **Oracle10g Database Performance Tuning Guide** for more information on the SQL Tuning Advisor.

31

...SQL Tuning Advisor

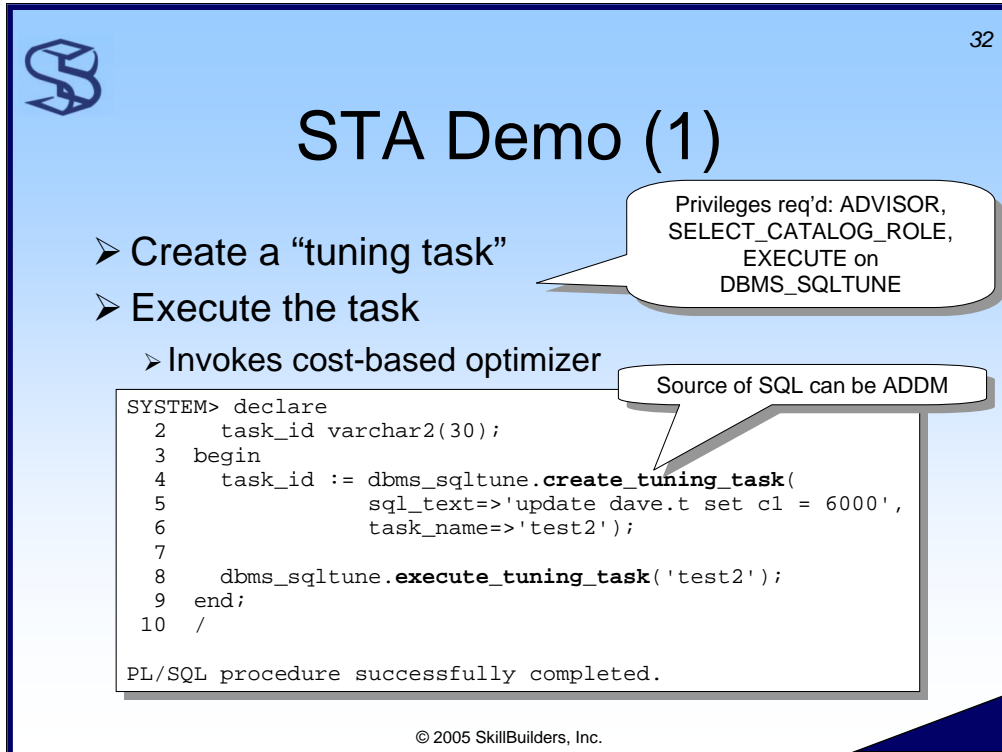
- Sources of SQL for input
 - ADDM shows high load statements
 - Snapshots / SQL Tuning Sets (STS)
 - Text
- Two interfaces
 - Enterprise Manager
 - Easy GUI
 - Launch with point and click
 - Implement recommendations with point and click
 - Submits job....
 - Supplied package DBMS_SQLTUNE

© 2005 SkillBuilders, Inc.

The SQL Tuning Advisor will accept SQL from several sources, including:

- ADDM. ADDM reports will identify high load SQL statements. With OEM, point and click to run the STA on those statements.
- SQL Tuning Sets (STS) – From snapshots we can create a tuning set – a series of SQL statements to be analyzed by the advisor.
- Simple text input.

Invoke the SQL Tuning Advisor through OEM or with the DBMS_SQLTUNE package. (The next page contains an example of the latter.)



32

STA 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.
```

© 2005 SkillBuilders, Inc.

Step 1 of using the SQL Tuning Advisor is to create a tuning task. Here we identify the SQL we want advice on. The source can be a literal (as shown above) or SQL identified by ADDM as needing tuning.

Step 2 involves executing the task created in step 1.

Refer to the supplied script `sql_advisor.sql` for a working copy of this example.

33

STA 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

© 2005 SkillBuilders, Inc.

In Step 3, we request the report with the `DBMS_SQLTUNE.REPORT_TUNING_TASK`. The top of the report has general information.

The screenshot shows a slide titled "STA Demo (3)" with a blue background and a white text area. The text area contains a sample STA report. Callouts point to specific parts of the report: "Findings" section identifies possible issues, "Recommendations" section gives possible solutions, and a "Clean up" section shows a SQL command to drop a tuning task.

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

Clean up

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

PL/SQL procedure successfully completed.

© 2005 SkillBuilders, Inc.

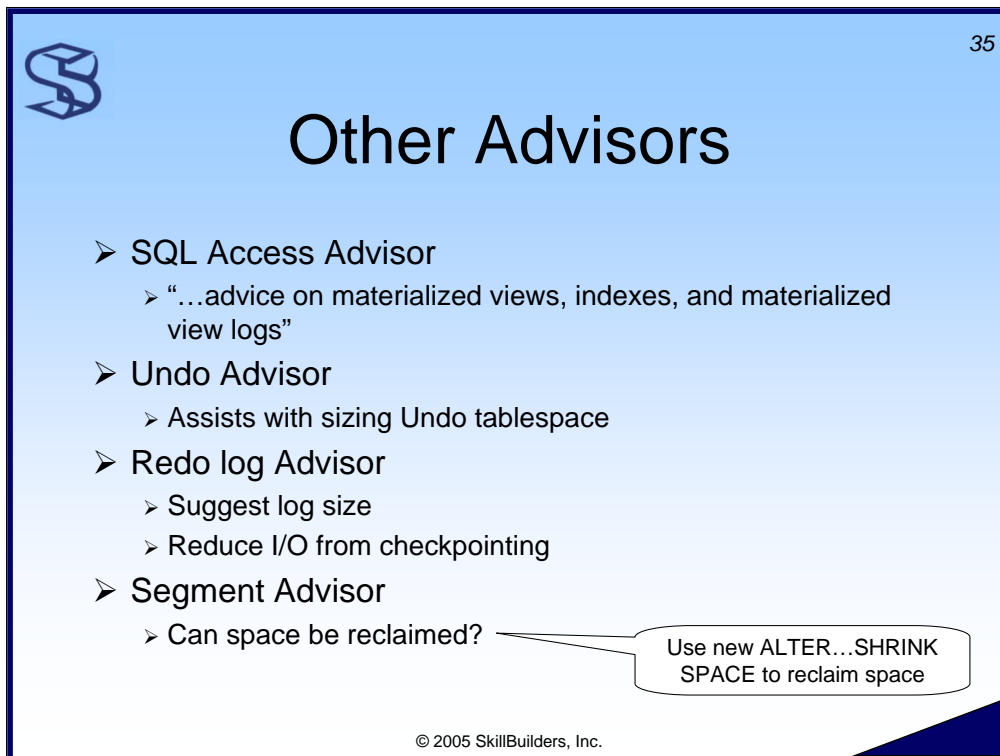
The main body of the SQL Tuning Advisor report contains findings, i.e. advice designed to make your SQL run faster. Note in the recommendations it even provides a possible solution!

The tasks are actually stored in the data dictionary, so clean up is a good idea. You can display the tasks with this query:

```
1 select owner, task_name
2 from DBA_ADVISOR_TASKS
3* where owner = 'SYSTEM';
```

OWNER	TASK_NAME
SYSTEM	SQL_TUNING_1091211216712
SYSTEM	TASK_793
SYSTEM	TASK_794
SYSTEM	TASK_2133
SYSTEM	TASK_2135
SYSTEM	test1
SYSTEM	TASK_2145

DBA_ADVISOR_LOGS also contains data on advisor tasks.



35

Other Advisors

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

Use new ALTER...SHRINK SPACE to reclaim space

© 2005 SkillBuilders, Inc.

Oracle10g provides several Advisors. "DBA_ADVISOR_DEFINITIONS displays the properties of all advisors in the database." For example:

```
DAVE@linux3> select advisor_name
  from dba_advisor_definitions;
```

```
ADVISOR_NAME
```

```
-----
```

```
ADDM
```

```
SQL Access Advisor
```

```
Undo Advisor
```

```
SQL Tuning Advisor
```

```
Segment Advisor
```

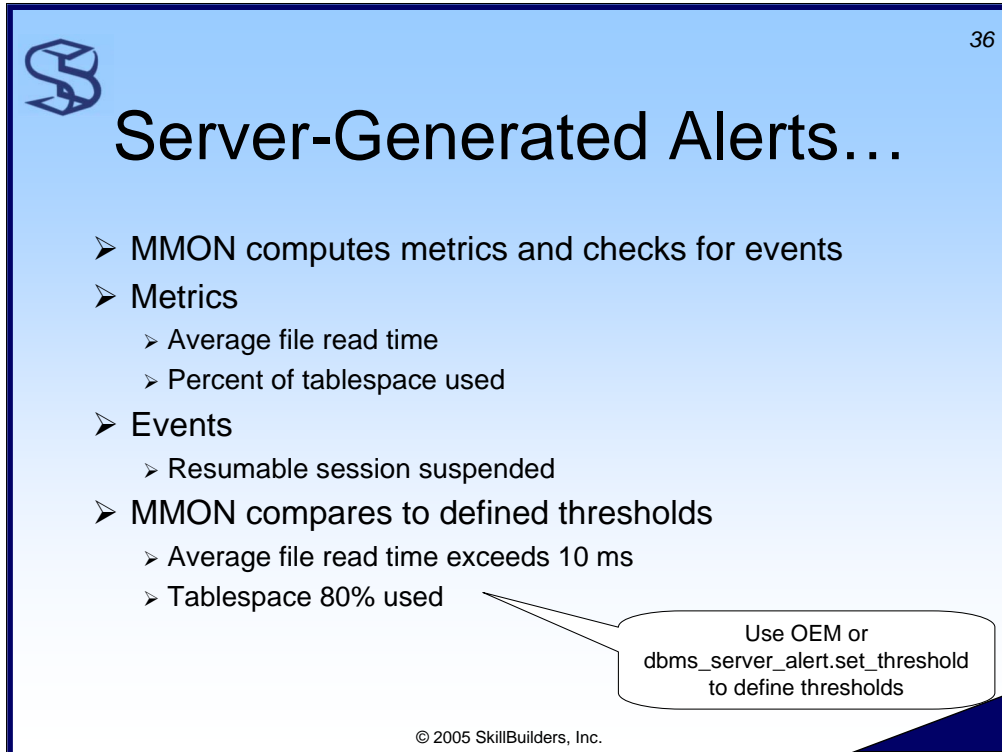
```
SQL Workload Manager
```

```
Tune MView
```

(1) Source: **Oracle10g Database Performance Guide**

(2) Source: **Oracle10g Administrator's Guide**

(3) The new "shrink" capability as described in the **SQL Reference**: "manually shrink space in a table, index-organized table, index, partition, subpartition, materialized view, or materialized view log." This can be done while the object is online.



36

Server-Generated Alerts...

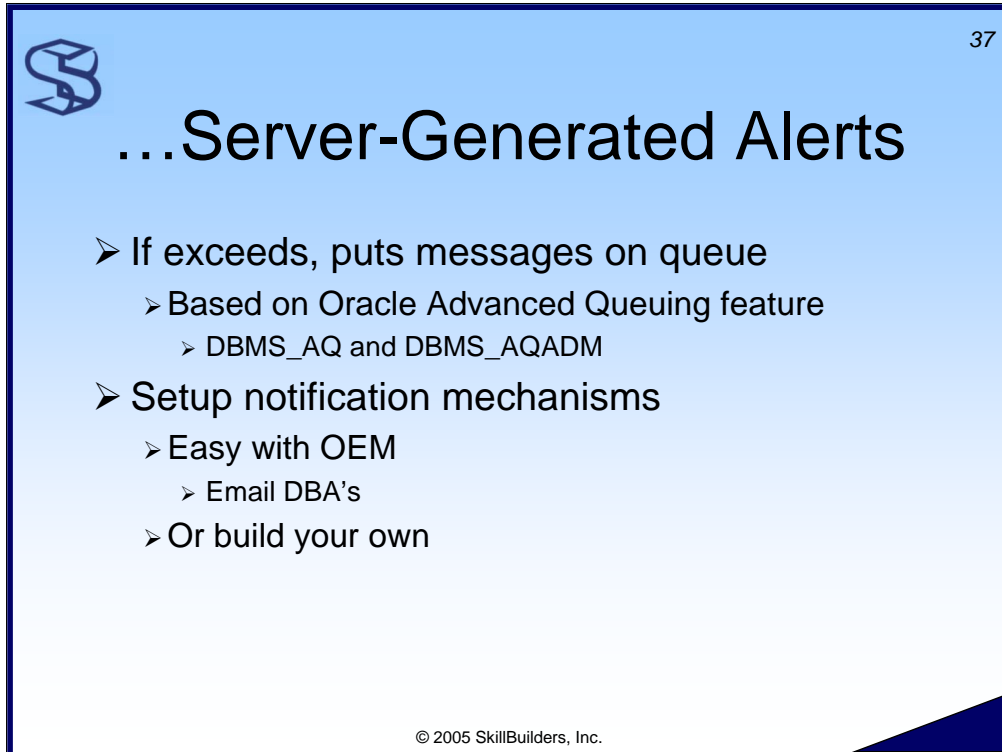
- MMON computes metrics and checks for events
- Metrics
 - Average file read time
 - Percent of tablespace used
- Events
 - Resumable session suspended
- MMON compares to defined thresholds
 - Average file read time exceeds 10 ms
 - Tablespace 80% used

Use OEM or
`dbms_server_alert.set_threshold`
to define thresholds

© 2005 SkillBuilders, Inc.

Every minute, MMON uses AWR-related statistics to compute “metrics” and check for exceeded thresholds and “events”. A metric, for example, is average file read time. An event, for example, is resumable session suspended.

MMON compares the metrics to see if they’ve exceeded defined thresholds. Thresholds are set with Enterprise Manager or with the supplied `DBMS_SERVER_ALERT` package.



37

...Server-Generated Alerts

- If exceeds, puts messages on queue
 - Based on Oracle Advanced Queuing feature
 - DBMS_AQ and DBMS_AQADM
- Setup notification mechanisms
 - Easy with OEM
 - Email DBA's
 - Or build your own

© 2005 SkillBuilders, Inc.

If a threshold is exceeded or an event has occurred, MMON puts an “alert” (i.e. message) on a predefined queue. (This queuing mechanism uses Oracle’s Advanced Queuing feature.)

We can then build notification mechanisms. Using Oracle Enterprise Manager it’s easy to email administrator’s if an alert occurs.


For more information on Server-Generated Alerts see Chapter 4 of the Oracle10g Database Administrator’s Guide (“Monitoring the Operation of Your Database”).

Supplemental Notes

I found this list of dictionary views related to alerts in Metalink Note Note:266970.1:

- DBA_THRESHOLDS lists the threshold settings defined for the instance.
- DBA_OUTSTANDING_ALERTS describes the outstanding alerts in the database.

Notes for this slide continue on the next page...

38

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 "Job Scheduler" feature.

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

© 2005 SkillBuilders, Inc.

10g implements automatic collection of statistics via the new job scheduler (refer to the lesson on the Job Scheduler for complete details).

Query the `DBA_SCHEDULER_JOBS` view to see the statistics collection job and all jobs defined in your database.

39

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.
```

© 2005 SkillBuilders, Inc.

Oracle recommends that statistics be collected on the dictionary in 10g. The GATHER_DICTIONARY_STATS procedure “Gathers statistics for dictionary schemas 'SYS', 'SYSTEM' and schemas of RDBMS components.” (Source: **Oracle10g PL/SQL Packages and Types Reference**.) You will need either SYSDBA or both the ANALYZE ANY DICTIONARY and the ANALYZE ANY system privileges to execute this privilege.

Another feature is the AUTO_DEGREE option for parallel operations – this puts the burden of finding the best degree of parallelization on Oracle.

40

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

© 2005 SkillBuilders, Inc.

RBO is still shipped with the database, but it is not supported.



41

Even 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
- COMPUTE STATISTICS now default on
 - CREATE INDEX
 - ALTER INDEX ... REBUILD
- Shared Server enhancements
 - MTS_ parameters obsolete
- Resource Manager Enhancements
- Flush buffer cache
 - Why?

```
SYSDBA> alter system flush buffer_cache;  
System altered.
```

© 2005 SkillBuilders, Inc.



42

OEM Enhancements...

- Now a browser-based UI
- Two versions
 - “Database Control”
 - For single instance
 - “Grid Control”
 - For RAC environments
- Management packs installed by *default*
 - Database Diagnostics Pack
 - Access to ADDM and AWR
 - Database Tuning Pack
 - SQL Tuning Advisor and more
 - Database Configuration Pack
 - Patching, cloning and more

© 2005 SkillBuilders, Inc.




43

...OEM Enhancements

- Packs cost more \$\$money\$\$!
- Without packs, limited usefulness
- Release 2 enhancements
 - Top Activity page
 - All wait classes and related statistics
 - “Memory Access Mode”
 - System-level calls to retrieve info from SGA
 - No SQL calls
 - Use for slow or hung database
 - “Historical View”
 - Session activity real-time or historical
 - Several OEM UI changes

© 2005 SkillBuilders, Inc.

44

PL/SQL Performance...

- New optimizing compiler
- Better
 - “integer performance”
 - “reuse of expression values”
 - “simplification of branching code”
 - “performance for some library calls”
 - “elimination of dead code”
- Looks for conversion possibilities
 - e.g. Convert FOR LOOPS into BULK COLLECT

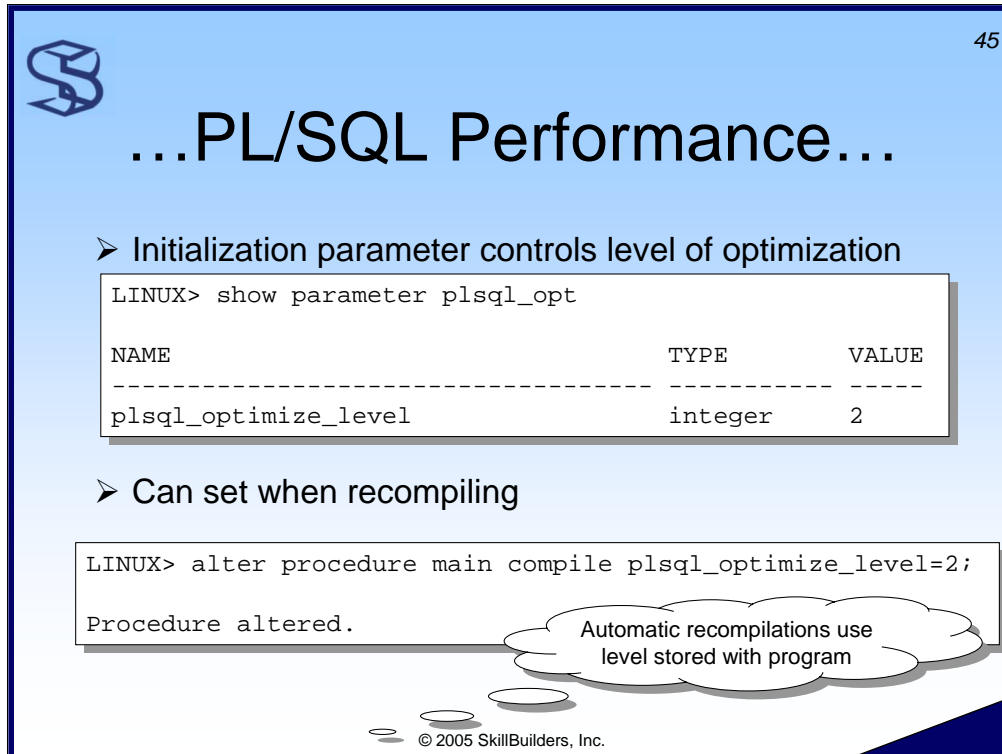
© 2005 SkillBuilders, Inc.

Oracle10g comes with a new PL/SQL compiler. This compiler is an “optimizing” compiler in that – unlike the 9i compiler – contains logic that, for example, makes better use of integers resulting in faster execution. This slide contains several quotes from the **Oracle10g PL/SQL User’s Guide and Reference** concerning the performance improvements made to the 10g PL/SQL compiler. Since these all refer to internal compiler code which we as customers are not privy to, I guess we must take this as fact and expect to see some performance improvement as we migrate applications to 10g.

The 10g compiler also looks for conversion opportunities to make the PL/SQL code execute faster. For example, FOR LOOPS that contain SELECT INTO statements (thus causing a single SQL statement to be executed each time through the loop) are converted to BULK COLLECT operations.

Supplemental Notes

Last year I saw the PL/SQL guru **Steve Feuerstein** talk about 10 PL/SQL features. Regarding performance he said, “in Oracle9i, PLS_INTEGER was faster...in 10g, it doesn’t matter; the database uses machine arithmetic for all integer types.”



45

...PL/SQL Performance...

- Initialization parameter controls level of optimization

```
LINUX> show parameter plsql_opt
```

NAME	TYPE	VALUE
-----	-----	-----
plsql_optimize_level	integer	2

- Can set when recompiling

```
LINUX> alter procedure main compile plsql_optimize_level=2;
Procedure altered.
```

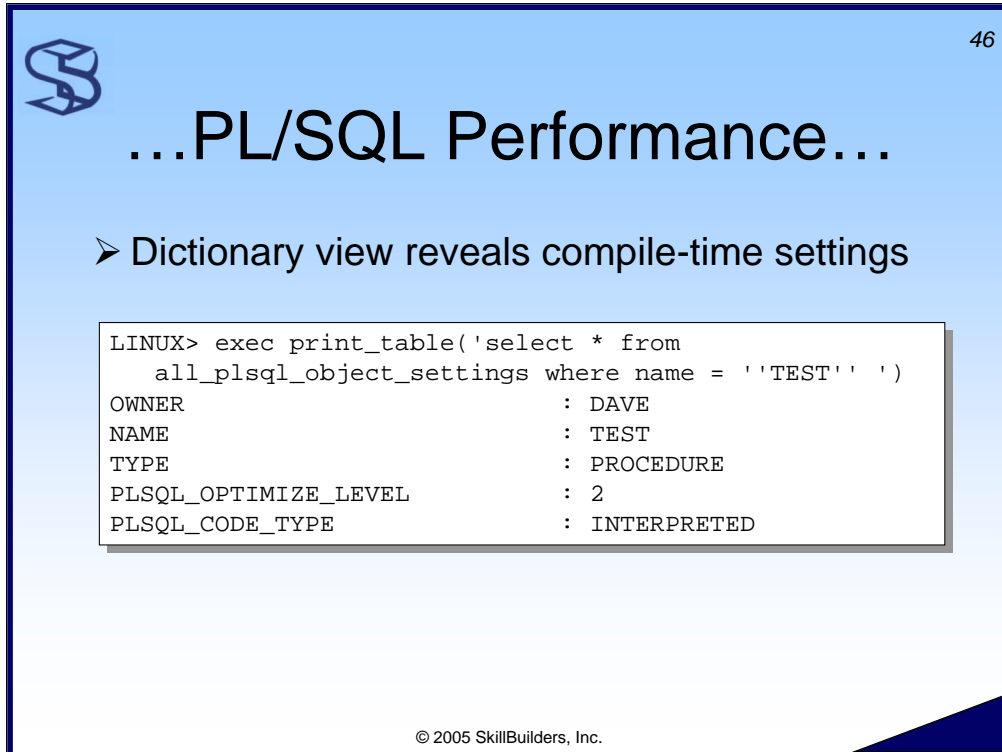
Automatic recom compilations use level stored with program

© 2005 SkillBuilders, Inc.

The PLSQL_OPTIMIZE_LEVEL parameter is a dynamic parameter that controls the level of optimization pursued by the compiler. It can be set at the system or session level. Set to one of the following values:

- 0 – Almost turns off all new optimization techniques. Oracle10g Reference manual says “code will run somewhat faster than it did in Oracle9i, use of level 0 will forfeit most of the performance gains.”
- 1 – Some optimization. **Oracle10g Reference** manual says “Applies a wide range of optimizations...but does not move source code out of its original source order.”
- 2 – Most aggressive optimization, which should result in best performance. However, some compilations may take too long; reverting to level 1 may decrease compilation time with only moderate loss of performance.

Note that the **Oracle10g Reference** manual says level 0 is the default but that may be incorrect; it was set to level 2 on both my 10g installations.



46

...PL/SQL Performance...

➤ Dictionary view reveals compile-time settings

```
LINUX> exec print_table('select * from
    all_plsql_object_settings where name = ''TEST'' ')
OWNER                               : DAVE
NAME                                : TEST
TYPE                                 : PROCEDURE
PLSQL_OPTIMIZE_LEVEL                 : 2
PLSQL_CODE_TYPE                       : INTERPRETED
```

© 2005 SkillBuilders, Inc.

ALL_PLSQL_OBJECT_SETTINGS contains the “compiler settings of stored objects accessible to the user”. (Source: **Oracle10g Reference** manual.)

The ALL_PLSQL_OBJECT_SETTINGS view deprecates (replaces) the ALL_STORED_SETTINGS view.

Complete output from the second example shown in this slide:

```
LINUX> create procedure test as
2  begin
3  null;
4  end;
5  /
```

Procedure created.

Notes for this slide continue on the next page...

47

Optimization Comparison...

```

ORACLE9i> begin
 2   for x in ( select * from all_objects where rownum < 10000)
 3   loop
 4     null;
 5   end loop;
 6 end;
 7 /

```

PL/SQL procedure successfully completed.

```

select *
from
  all_objects where rownum < 10000

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.06	0.07	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	10000	3.18	3.11	0	63481	0	9999
total	10002	3.24	3.18	0	63481	0	9999

© 2005 SkillBuilders, Inc.

Here we see a 10,000-iteration loop traced on an Oracle9i R2 database.

The trace reveals 10,000 fetches, i.e. 10,000 calls to the database, caused by the SQL statement “select * from all_objects where rownum < 10000”.

Now, let’s look at how 10g handles this code...

48

...Optimization Comparison

```

ORACLE10G> begin
2   for x in ( select * from all_objects where rownum < 10000)
3   loop
4     null;
5   end loop;
6 end;
7 /

```

10g compiler converts this into a bulk collect

PL/SQL procedure successfully completed.

```

SELECT *
FROM
  ALL_OBJECTS WHERE ROWNUM < 10000

```

Arraysize = 100

Great reduction in CPU, elapsed time and logical I/O

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	100	0.79	0.74	0	15812	0	9999
total	102	0.79	0.75	0	15812	0	9999

© 2005 SkillBuilders, Inc.

In Oracle10g, the optimizing compiler converted this to a BULK FETCH operation. The trace reveals that the arraysize must be 100; 10000 rows (rounded) divided by "x" = 100 fetches. "x" must be 100. Though I don't know of a way to control the arraysize, 100 is reasonable and should be fine. By looking at the reduction in CPU and elapsed time, we see that it is certainly a lot better than the 9i version.

Note that this was not entirely a fair test, in that the platform for my 10g instance was not exactly the same as my 9i instance. Nevertheless, you should see improvement.

Refer to the supplied script `plsql_optimizer.sql` for a working version of this demonstration (of course, you will need two databases to make the comparison; a 9i and a 10g).

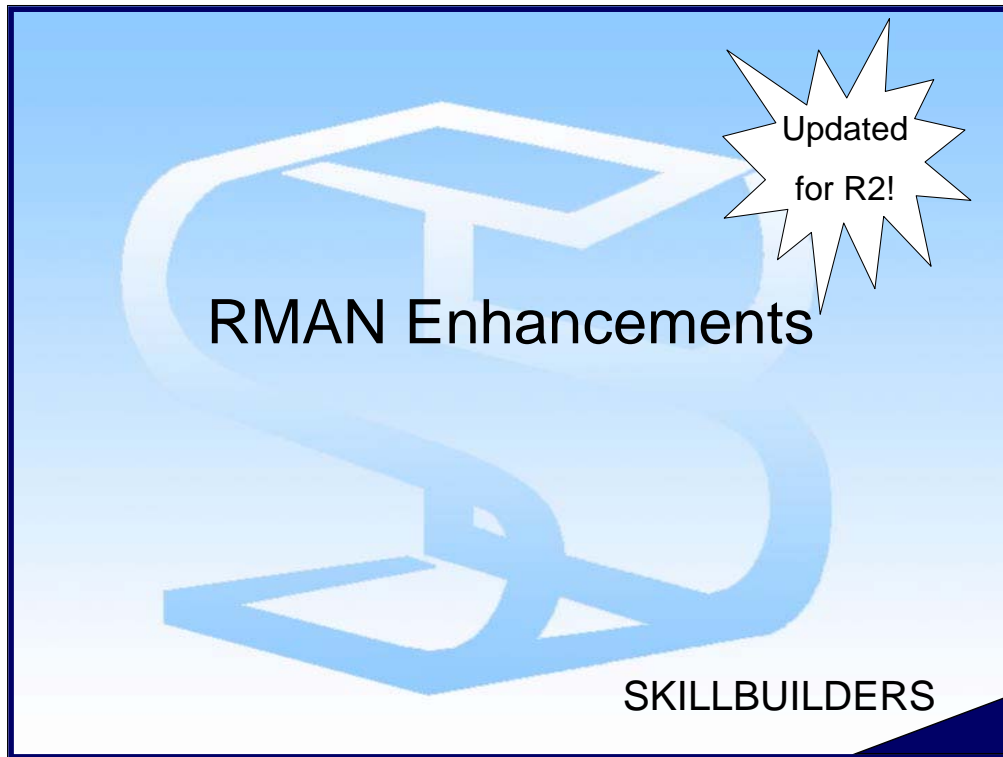


49

Performance & Tuning Feature Summary

- Automatic SGA Management
- Automatic Workload Repository
- Automatic Database Diagnostic Monitor
- New Advisors
- SQL Profiles
- Automatic Statistics Collection
- Wait Events Enhancements
- OEM
- New Hints
- Old Hints
- RBO Desupported
- Enterprise Manager Enhancements
- PL/SQL Enhancements


© 2005 SkillBuilders, Inc.



Author: Dave Anderson

Release Date: October, 2004

Last Revision: September, 2005



51

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!

Your results will vary...

© 2005 SkillBuilders, Inc.

RMAN now supports the creation of compressed backupset pieces. As you can see, my tests show significant reduction in piece size. Compression not only saves space but can save lots of time for network backups (due to the reduced number of bytes sent across the network).

Supplemental Notes

You can configure your RMAN environment so compression is the default. For example:

```

RMAN> configure device type disk backup type
2> to compressed backupset;
RMAN> configure device type sbt backup type
2> to compressed backupset;

```

The V\$BACKUP_PIECE view has a new column, COMPRESSED, which will contain “YES” for compressed backup pieces.

52

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

```

OMF controls file location

Level 0 duration

Level 3 duration


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

© 2005 SkillBuilders, Inc.

Oracle10g RMAN incremental backups are potentially much faster because the entire datafile does not have to be read to determine which blocks have changed. Instead, the database records (i.e. tracks) changed blocks in a file. When an incremental backup is taken, RMAN accesses the “block change tracking file” to determine which blocks to backup.

This behavior must be configured with the ALTER DATABASE ENABLE BLOCK CHANGE TRACKING statement as shown above. The default location of the block change tracking file is the Oracle Managed File DB_CREATE_FILE_DEST parameter. Use the “FILE” clause if you are not using OMF or you want to put the file in a location other than DB_CREATE_FILE_DEST.

Refer to the **Oracle10g SQL Reference** for more information on the ALTER DATABASE statement and the **Oracle10g Database Backup and Recovery Basics** manual for more information on this feature in general.



53

...Change Tracking

- Block change recorded in a file

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

STATUS	FILENAME	BYTES
ENABLED	+ASM_DISK_GROUP1/orcl/change_tracking/ctf.343.9	11599872

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

© 2005 SkillBuilders, Inc.

Change tracking is implemented by the creation of a file that records which blocks are changed. The Oracle documentation says that the initial size of the file is 10MB and grows in 10MB increments. The space *required* is 1/30,000 of the blocks tracked.

While the file size should not be an issue, Oracle does mention a performance hit related to the use of this feature. Test before implementing.


Refer to Chapter 4 of the **Oracle10g Database Backup and Recovery Basics** manual for more information on change tracking. Specifically, read the section entitled “Improving Incremental Backup Performance: Change Tracking.”

Supplemental Notes

Execute the following query to determine if change tracking is enabled for your database:

```
DAVE@linux3> select status from v$block_change_tracking;
```

```
STATUS
-----
ENABLED
```

54

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

© 2005 SkillBuilders, Inc.

The RMAN Incrementally Updated Backup feature is designed to limit the amount of redo you will need to apply during recovery operations – *thus reduce recovery time*.

The basic idea is to create image copies of your datafiles, then subsequently update the image copies with incremental backups, i.e. merge level 1 incremental backups into an existing image copy.

55

...Incrementally Updated Backups

Run this script daily

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

Tags must match

```

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> }
Starting 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]"

© 2005 SkillBuilders, Inc.


The script can be run daily. Each day this runs:

- The RECOVER COPY command updates all datafile image copies (with the tag "IC1") with the previous day's level 1 incremental. If it does not find an image copy to update or a level 1 to apply, the RECOVER command simply issues messages and successfully terminates. Note that an *identical* user-defined tag must be used on the RECOVER COPY command and the BACKUP FOR RECOVER OF COPY WITH TAG clause.
- The BACKUP command creates a new level 1 incremental. However, if a level 0 image copy does not exist (e.g. on the 1st run), the BACKUP command will create one.

Thus, at all times, you have available for recovery:

- Image copy up to a maximum of 48 hours old
- Level 1 incremental up to a maximum of 24 hours old
- Online and archive logs to support point-in-time or complete recovery.

Notes for this slide continue on the next page...



56

New V\$ Views

- R2 adds 18 new views
- V\$RMAN_BACKUP_JOB_DETAILS
 - Status, date/time, elapsed time, read rate, more
- V\$BACKUP_SET_DETAILS
 - Incremental level, number of pieces, size ,status
- V\$BACKUP_PIECE_DETAILS
 - File handles, status, size

© 2005 SkillBuilders, Inc.

Oracle10g adds 18 new backup and recovery-related views. Refer to the Reference manual for details on these views:

V\$BACKUP_ARCHIVELOG_DETAILS

V\$BACKUP_ARCHIVELOG_SUMMARY

V\$BACKUP_CONTROLFILE_DETAILS

V\$BACKUP_CONTROLFILE_SUMMARY

V\$BACKUP_COPY_DETAILS

V\$BACKUP_COPY_SUMMARY

V\$BACKUP_DATAFILE_DETAILS

V\$BACKUP_DATAFILE_SUMMARY

V\$BACKUP_PIECE_DETAILS

V\$BACKUP_SET_DETAILS

V\$BACKUP_SET_SUMMARY

V\$BACKUP_SPFILE_DETAILS

V\$BACKUP_SPFILE_SUMMARY

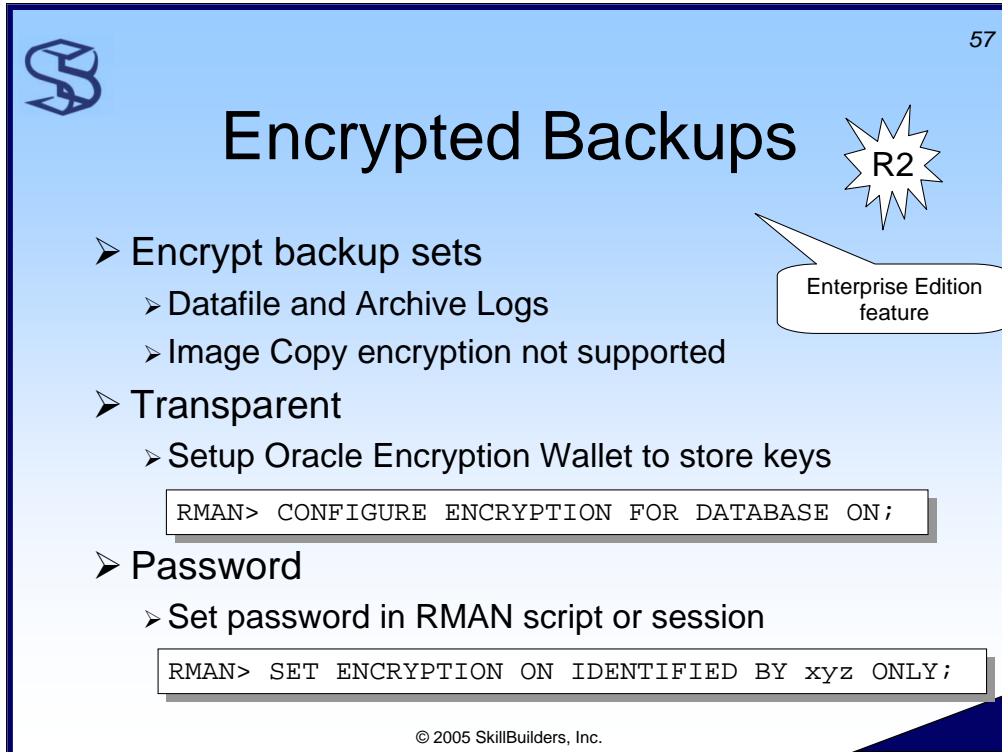
V\$RMAN_BACKUP_JOB_DETAILS

V\$RMAN_BACKUP_SUBJOB_DETAILS

V\$RMAN_BACKUP_TYPE

V\$FLASH_RECOVERY_AREA_USAGE

V\$RMAN_ENCRYPTION_ALGORITHMS



57

Encrypted Backups

- Encrypt backup sets
 - Datafile and Archive Logs
 - Image Copy encryption not supported
- Transparent
 - Setup Oracle Encryption Wallet to store keys
- Password
 - Set password in RMAN script or session

```
SQL> RMAN> CONFIGURE ENCRYPTION FOR DATABASE ON;
```

```
SQL> RMAN> SET ENCRYPTION ON IDENTIFIED BY xyz ONLY;
```

© 2005 SkillBuilders, Inc.

Another RMAN R2 feature (Enterprise Edition only) is support for encrypted backup sets (encrypting image copies is not supported).

Oracle supports three encryption techniques:

- Transparent – In this mode an encryption key is stored in the Oracle Encryption Wallet feature, which must be previously configured. I believe this (Oracle Encryption Wallet) also supports the R2 feature transparent column encryption. A password for encryption and decryption is *not* required. This is the default.
- Password – Password encryption requires that the DBA provide the same password on the backup and restore commands. Forgetting the password renders the backupset useless. This mode can be particularly useful when restoring on a remote system.
- Dual Mode – With this technique, either the password or the Wallet can be used to decrypt the backup.

Encryption of data requires CPU resources. I have not yet tested the overhead (Sept '05).

58

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"

© 2005 SkillBuilders, Inc.

The BACKUP DURATION parameter will do one of two things for you:

- If MINIMIZE LOAD is specified, RMAN will *reduce* resource consumption so that the job requires the amount of time specified. This is like the RATE parameter, but perhaps more intuitive.
- IF MINIMIZE LOAD is not specified, the time specified is the maximum time the BACKUP can execute. RMAN will *kill the job* when it reaches the specified time limit.

59

More Enhancements...


- Simplified recovery through incarnations
- Default backup location is flash recovery area
- Archive log deletion policy
- RESTORE DATABASE PREVIEW
 - Identify the backups that would be used
- Restore failover
 - If error reading backup, RMAN keeps trying different backup copies until all possible exhausted

© 2005 SkillBuilders, Inc.

Here is a brief description of other Oracle10g RMAN enhancements:

- Recovering through incarnations of a database with previous versions of RMAN was a complicated task. Oracle10g makes this much simpler – as easy as recovering to the current incarnation – as long as you restore or create a control file from the target incarnation (should not be a problem if you configure CONTROLFILE AUTOBACKUP ON). Refer to Chapter 8 of the **Backup and Recovery Advanced User's Guide** for details.
- Whereas the default backup location in Oracle9i RMAN was simply a local directory, the Oracle10g default is the flash recovery area. (The flash recovery area is covered in the “Flashback Enhancements” lesson of this course; it is configured with the DB_RECOVERY_FILE_DEST parameter.) This is a better choice because we can backup the flash recovery area to tape with one command (BACKUP RECOVERY AREA) and we can control the size of this area (see the DB_RECOVERY_FILE_DEST_SIZE parameter). Chapter 3 of the **Oracle Backup and Recovery Basics** manual is a good starting place for learning about the flash recovery area.

Notes for this slide continue on the next page...

60

...More Enhancements


- Channel failover for BACKUP
- COPY command deprecated
 - BACKUP can create image copies
- Default backup can be configured to create
 - Normal backupset
 - or compressed backupset
 - or disk-based image copy
- RMAN catalog tablespace created in SYSAUX
- More...

© 2005 SkillBuilders, Inc.

Additional RMAN enhancements, continued:


- RMAN now has the ability to try another channel if a channel fails during backup.
- The COPY command is deprecated (will eventually be desupported) because the BACKUP command can be used to create image copies. For example:

```
RMAN> backup as copy tablespace users;
```
- The CONFIGURE DEVICE TYPE command now supports configuring for normal backupset, compressed backup or image copy. (Note that the image copy option only works for DEVICE TYPE DISK.)
- To alleviate clutter in the SYSTEM tablespace, the RMAN catalog is now by default created in the new SYSAUX tablespace.

61

More R2 Enhancements


- Tempfiles automatically created during recovery
- Unused block compression now respects high-water mark
- RAC dynamic channel allocation
- Create incremental since last sync with Standby
- OEM enhancements



© 2005 SkillBuilders, Inc.

R2 also provides these features:

- After a whole (full) database restore and recovery, and at database open, Oracle will automatically create missing tempfiles. The original size, location, autoextend and maxsize attributes are used to recreate the tempfile.
- If backup compression is enabled, RMAN will backup only used blocks (blocks below the high-water mark). Previous releases of RMAN backed up any block that ever contained data – even if it was not currently used.
- RAC backups no longer require manual channel allocation for each node – dynamic channel allocation is now supported.
- We now have the ability to create an incremental backup as of from SCN. Only changes made at or after that SCN are backed up. This can be useful for updating a Standby Database with changes made to the primary since the last synchronization. See the Backup and Recovery Advanced User's Guide and the `BACKUP INCREMENTAL FROM SCN [scn] DATABASE` command.
- Oracle Enterprise Manager has been enhanced to support compression during cloning operations, RMAN scripts and backup job notification.

62

RMAN Summary

- RMAN now a viable enterprise-level backup and recovery utility
- Many useful new features
 - Compression
 - Change Tracking
 - Incrementally Updated Backups
 - Recover through incarnations
 - more...

© 2005 SkillBuilders, Inc.

With the new features Oracle10g introduces, I think many DBAs will now start to view RMAN as a viable tool for backup and recovery of production databases.


While I have tried to highlight the “most important” RMAN enhancements in this lesson, there are more and undoubtedly you will find enhancements not covered here that are important to you. So, refer to the section “What’s New in Backup and Recovery” in the **Backup and Recovery Advanced User’s Guide** for more information.



Author: Dave Anderson

Release Date: October 2004

Last Revision: January 2005

64


Introducing ASM...

- Automatic Storage Management
- Built-in file system and volume manager
 - “Vertical integration of both the file system and the volume manager”
 - Source: Oracle Corp
- Simplify administration
 - Manage small number of disk groups
 - Not thousands of files

© 2005 SkillBuilders, Inc.

Automatic Storage Management (ASM) is an integrated subsystem for disk management shipped with all editions of Oracle10g.

Its purpose is to simplify administration by eliminating the need for detailed management of large numbers of database files. The administrator instead focuses attention on a relatively small number of “disk groups”. A disk group is a new database object that represents a collection of disks and files.

65

...Introducing ASM...

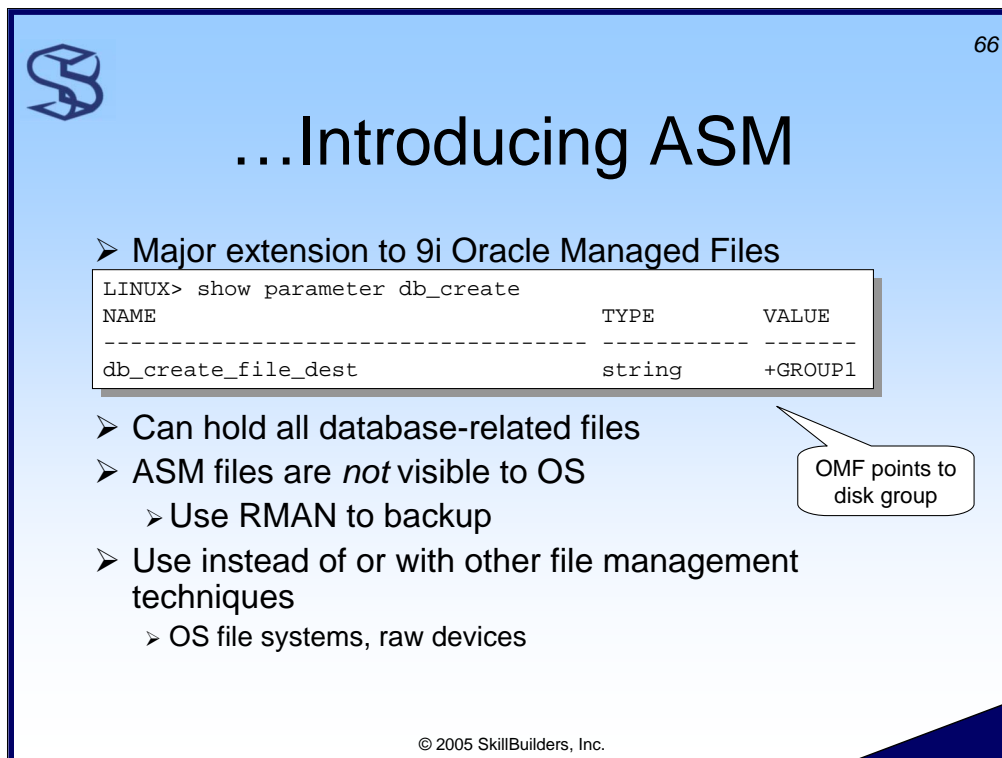
- Automated S.A.M.E.
 - Striping to balance I/O load (performance)
 - Mirroring (redundancy)
- Dynamic Parallel load balancing
 - Helps prevent hot spots
 - “IO spread evenly across available disks”
 - Prevents fragmentation
- Automatic online balancing
 - Add or remove disks online
- Technique for handling I/O errors

© 2005 SkillBuilders, Inc.

ASM implements the “Stripe and Mirror Everything” methodology, recommended by Oracle Corp for all Oracle databases. It provides several *automated* IO-related functions for database-related files, such as:

- Striping of database files across disks for better I/O performance
- Mirroring (2 or 3-way) of files for redundancy (protection against loss of files due the loss of a disk device)
- Load balancing while the database is open and active. This can be done in parallel.
- Support for reorganizing extents do the addition or removal of disks – while the database is open.

While ASM cannot fix I/O errors, ASM has a procedure for handling them, which in many cases – as long as you are mirroring – is transparent to applications. Consider this excerpt from Chapter 12 of the Administrator’s Guide: “ASM takes proactive measures with regards to IO errors. A permanent IO error is only signaled to the caller (Oracle IO process) after several retries in the device driver. If a permanent disk IO error is incurred during an Oracle write operation, then the affected disk is removed from the diskgroup by ASM, thus preventing more application failures. If the loss of a disk results in data loss, ASM will automatically dismount the diskgroup to protect the integrity of the diskgroup data.”



66

...Introducing ASM

- Major extension to 9i Oracle Managed Files

```
LINUX> show parameter db_create
NAME                                 TYPE                                 VALUE
-----                                 -                                 -
db_create_file_dest                  string                               +GROUP1
```

- Can hold all database-related files
- ASM files are *not* visible to OS
 - Use RMAN to backup
- Use instead of or with other file management techniques
 - OS file systems, raw devices


OMF points to disk group

© 2005 SkillBuilders, Inc.

ASM can technically be considered a major advance to the Oracle9i Oracle Managed Files feature. OMF creates and controls the placement of files. By setting the `DB_CREATE_FILE_DEST` parameter to point at an ASM disk group, we get:

- placement in the diskgroup
- creation without overly complex `DATAFILE` parameters
- striping
- mirroring
- rebalancing
- reorganization for new or dropped disks

ASM can be used instead of existing file management techniques or in conjunction with file system files or raw devices.

67

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

© 2005 SkillBuilders, Inc.

One of the most conspicuous changes a DBA needs to get comfortable with is the “ASM instance”. The ASM instance is started on the server to manage the ASM disk metadata, or so says the Oracle manual. (Given the background processes that start, I suspect it does more, like maybe even participate in the rebalance operations when a disk is added or removed.)

Each server can have one ASM instance that can be used by many “database instances.” Take note of the new term “database instance”. This describes a instance dependent on an ASM instance for its disk management.

68

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!

© 2005 SkillBuilders, Inc.

To connect to the ASM instance:

- Set the ORACLE_HOME if not already set.
- Set the ORACLE_SID to +ASM (the default ASM instance name created by DBCA).
- Always connect “AS SYSDBA” to an ASM instance. ASM instances do *not* have a dictionary nor database, so this is a requirement.


Stopping the ASM Instance

Use the SHUTDOWN command to stop the ASM Instance. If you have a database open, either shut it down first or use shutdown immediate (otherwise, SHUTDOWN will hang). Also, if Enterprise Manager is or was running, there might be open connections to the ASM instance. Therefore, do one of the following:

Stop OEM *services (process)* before shutting down the ASM instance:

```
[oracle@springsteen oracle]$ emctl stop dbconsole
```

Or use SHUTDOWN IMMEDIATE. **Warning:** *However, this will crash all databases using this ASM instance.*



69

ASM Disks

- ASM disks must be discovered
 - Give raw devices to ASM
 - Or install "ASMLIB"
- Parameter `ASM_DISKSTRING` to discover disks
 - Can limit to specific disks

```
SQL> select name, path, total_mb, free_mb,
2         reads, writes
3         from v$asm_disk;
```

NAME	PATH	TOTAL_MB	FREE_MB
GROUP1_0003	/dev/raw/raw4	34710	29769
GROUP1_0002	/dev/raw/raw3	34710	29852
GROUP1_0001	/dev/raw/raw2	34710	29213
GROUP1_0000	/dev/raw/raw1	32718	28416

ASM instance shows all available disks


Database shows disks in use by that database

© 2005 SkillBuilders, Inc.

When the ASM instance is started, the value of the `ASM_DISKSTRING` parameter controls what disks can be used by the ASM instance.

Querying the `V$ASM_DISK` view from the ASM instance shows the available disks.

The ASM-related `V$` views provide information about diskgroups, ASM disks, ASM files, etc. It is important to note that the views will return different results when queried from the ASM instance versus the database instance (i.e. the database using the ASM instance).



70

ASM 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
-----
```

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


© 2005 SkillBuilders, Inc.

Create diskgroups in the ASM instance. This example creates a group with NORMAL REDUNDANCY (which creates a 2-way mirror), so there are two failgroups. Each failgroup has 2 disks, so a file will have its extents balanced across the two disks.

Query V\$ASM_DISKGROUP to see the affect of the CREATE DISKGROUP statement. In the second example in this slide I have joined to V\$ASM_DISK to show the disk name, path and failgroup.

V\$ASM_DISKGROUP view contents –

- In an ASM instance, contains one row for every diskgroup defined to the ASM instance.
- In an database instance, contains one row for every diskgroup mounted to the database.



71

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
```

© 2005 SkillBuilders, Inc.

ASM provides easy scalability. Disks can be added or dropped while the database is active.

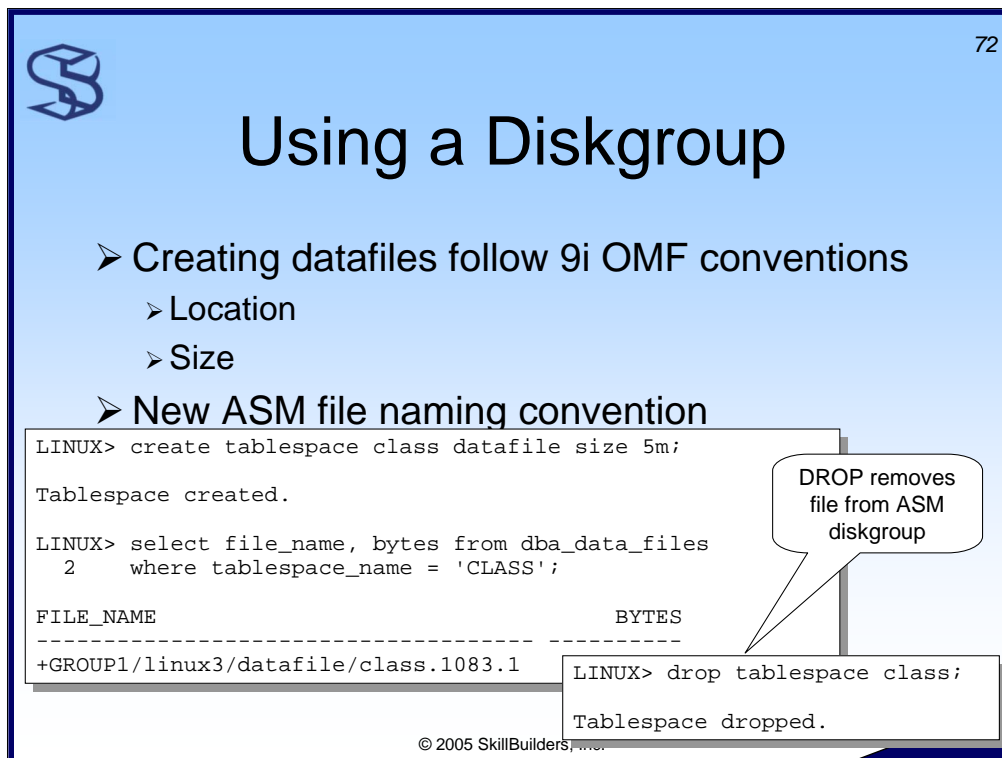
This example demonstrates removing a disk from a diskgroup. Querying V\$ASM_OPERATION shows information about the operation.

V\$ASM_OPERATION contains info about in-progress ASM operations such as rebalance operations when a device is added or dropped from a disk group or a hot-spot is being corrected. It can help the DBA determine if the rebalance operation is being hindered because the "POWER" is too low.

The "power" limit is used for tuning (speeding up or slowing down) rebalance operations. The value of 1 constrains the operation and causes it to consume less CPU and IO resources. The maximum value of 11 removes the constraints. The ASM instance initialization parameter ASM_POWER_LIMIT controls the power limit if the POWER parameter is not used on the statement. The parameter is dynamic and can be altered with ALTER SYSTEM:

```
+ASM> alter system set asm_power_limit=9 scope=memory;
```

Notes for this slide continue on the next page...



Using a Diskgroup

- Creating datafiles follow 9i OMF conventions
 - Location
 - Size
- New ASM file naming convention

```

LINUX> create tablespace class datafile size 5m;

Tablespace created.

LINUX> select file_name, bytes from dba_data_files
2      where tablespace_name = 'CLASS';

FILE_NAME                                BYTES
-----
+GROUP1/linux3/datafile/class.1083.1

```

DROP removes file from ASM diskgroup

```

LINUX> drop tablespace class;


Tablespace dropped.

```

© 2005 SkillBuilders, Inc.

Creating tablespaces is the same as it was with Oracle9i. The default datafile location is in the `DB_CREATE_FILE_DEST` parameter. The file name is generated by ASM.

Dropping an ASM tablespace always drops the associated ASM file. Use the `INCLUDING CONTENTS` option on the `DROP TABLESPACE` to force the drop even if the tablespace contains segments.

73

Summary ASM Benefits


- Easy Stripe and Mirror implementation
 - Oracle SQL-like commands to manage
 - Same commands across all platforms
 - Easier than cooked and LVM setups?
- Raw disk IO performance
- Scalable
- Great for tight budgets where 3rd party solutions not affordable

© 2005 SkillBuilders, Inc.

ASM provides an easy-to-implement S.A.M.E. implementation, without the cost of a EMC-like volume manager or in-depth knowledge of Unix raw file system management.

The early results on ASM performance look good. I recommend reading the available documents on ASM and – of course – doing your own testing.

ASM is scalable; we have seen in this lesson the ease of adding or removing disks.



Author: Dave Anderson

Release Date: October, 2004

Last Reviewed: January 2005

75

Predefined Variables

- New pre-defined variable support

```
SQL> set sqlprompt "_user'@'_connect_identifier> "  
DAVE@springsteen_linux3>
```

Callouts in the slide:
- "User name" points to `_user`
- "Net alias" points to `_connect_identifier`
- "Literal" points to the single quotes around `DAVE` in the output


- Also available
 - `_date`
 - `_privilege`

© 2005 SkillBuilders, Inc.

SQL*Plus predefines several new variables. (See the complete list of predefined variables with the DEFINE command.)

These include `_user` for the Oracle username and `_connect_identifier` for the Net alias used to connect. Also supported: `_date` (control format with ALTER SESSION SET NLS_DATE_FORMAT) and `_privilege` (for SYSDBA or SYSOPER, if used in the connect). Single quotes around a literal are only necessary if the literal directly precedes a variable name and thus would render the name unrecognizable.

The SET SQLPROMPT command is just one example of using these variables. The TTITLE and BTITLE commands also support these variables.



76

Recycle Bin

- SHOW RECYCLEBIN shows dropped tables and indexes

```
SQL> show recyclebin
ORIGINAL NAME      RECYCLEBIN NAME      OBJECT T
-----
CUSTOMER           BIN$5tDOyJMTz3ngMAAKMgAvzA==$0 TABLE
CUST_HISTORY       BIN$5tDOyJMpz3ngMAAKMgAvzA==$0 TABLE
```


- Clear recycle bin with PURGE RECYCLEBIN
 - (This is actually an SQL statement)

© 2005 SkillBuilders, Inc.

The SHOW RECYCLEBIN command shows dropped tables and dependent indexes. Refer to the lesson on 10g Flashback Enhancements for more information. This command is a shortcut for:

```
SELECT * FROM RECYCLEBIN;
SELECT * FROM USER_RECYCLEBIN;
```

The PURGE RECYCLEBIN SQL statement has many options. For example, you can purge a single table, all tables or the entire recycle bin. Refer to the **Oracle10g SQL Reference** for complete details.

77

SPOOL Enhancements

- SPOOL supports APPEND, CREATE, REPLACE options

```
SQL> save testfile
Created file testfile.sql
SQL>
SQL> select count(*) from dba_tables;

COUNT(*)
-----
          1570


SQL> save testfile append
Appended file to testfile.sql
```

© 2005 SkillBuilders, Inc.

The SPOOL command now supports a CREATE, APPEND and REPLACE options.

New SPOOL Syntax

```
SPOOL { file_name [.ext] [CRE[ATE]|REP[LACE]|APP[END]] | OFF | OUT }
```

78

Miscellaneous

- Quotes no longer required on OS startup using “/”

```
[oracle@springsteen oracle]$ sqlplus / as sysdba
```
- LOGIN.SQL re-executed with each CONNECT
 - GLOGIN.SQL too
 - Not just at startup
- DESCRIBE validates object
- R2 adds support for XQuery!
- R2 supports SET SERVEROUTPUT UNLIMITED

© 2005 SkillBuilders, Inc.

We used to have to enclose the “/ as sysdba” in quotes when used at the OS command prompt. No more.

LOGIN.SQL and GLOGIN are re-executed with each successful CONNECT.

The DESCRIBE command used to fail with a ORA-24372 if executed on an invalid object. Now, it validates the object and – if successful – displays the output.

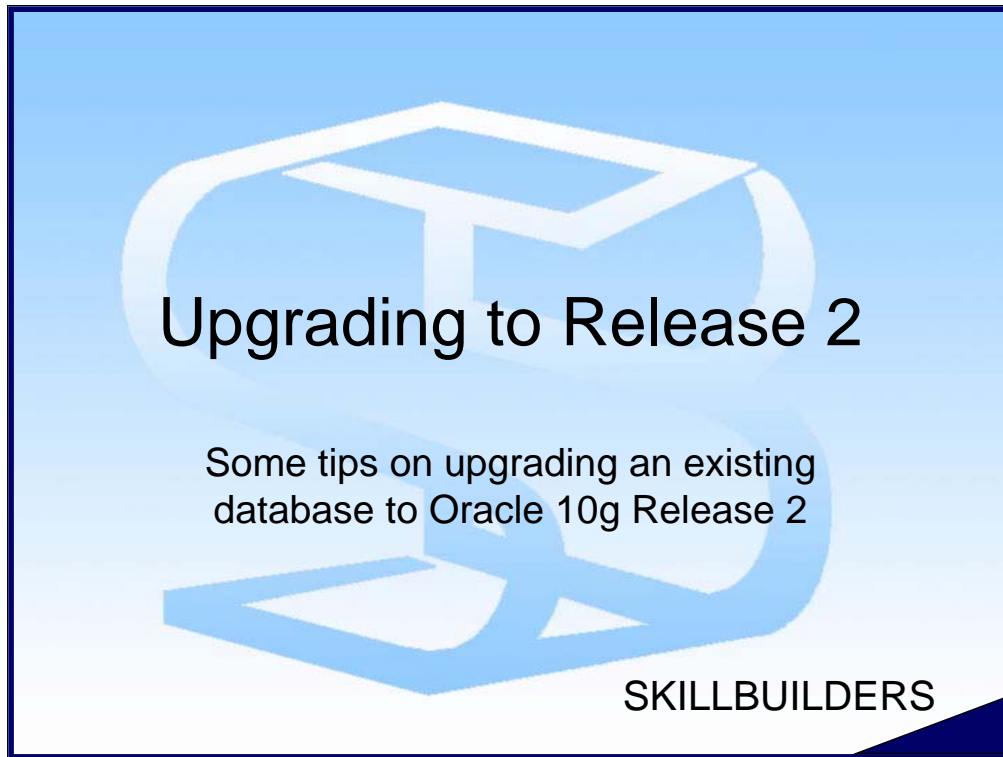


79

SQL*Plus Summary

- New predefined variables
- SHOW RECYCLEBIN
- SPOOL command enhancements
- LOGIN.SQL enhancements
- “/ AS SYSDBA”
- DESCRIBE command enhancements
- XQuery support

© 2005 SkillBuilders, Inc.



Author: Dave Anderson

Release Date: October, 2004

Last Reviewed: January 2005



81

Upgrading with DBUA

- I used Database Upgrade Assistant
 - DBUA is helpful GUI
- Upgrades existing database, in place
- Some gotcha's

© 2005 SkillBuilders, Inc.



82

Summary of Steps...

- Backup the operating system
- Validate the server host environment for the Release 2 installation
 - Download “validate” program from Metalink Note 283748.1
 - Check swap space, disk space, kernel parameters, and more
- Create the Release 2 Oracle Home
- Backup the database to be upgraded
- Install the 10g Release 2 software

© 2005 SkillBuilders, Inc.



83

...Summary of Steps

- Install available patches to Release 2 software
- Run the “Release 2 Upgrade Information Tool”
 - utlu102i.sql (R2Home/rdbms/admin)
 - Run on the source database
 - Checks init parameter, tablespace sizes and much more
- Start DBUA
 - Upgrade the source database
- Problems?
 - May have to restore the source database

© 2005 SkillBuilders, Inc.



84

Documents Required

- Quick Installation Guide and Installation Guide
- Database Readme
- Oracle Database Upgrade Guide.
- Oracle Database Administrator's Reference.
- 2-Day DBA
- Companion CD Installation Guide
- Metalink Notes:
 - Note 283748.1
 - Note 316889.1
 - Note 272322.1

© 2005 SkillBuilders, Inc.



Upgrade Summary

- Definitely some gotcha's
- I had to restore, start again
- Read my paper at
 - www.skillbuilders.com



86

Oracle 10g Classes

- Oracle10g New Features Overview
 - Jan 17, NYC
 - Jan 18, Melville
- Oracle10g New Features for Administrators
 - Hands-On!
 - Feb 7 – 10, Melville
 - Feb 14 – 17, NYC
- Oracle10g New Features for Developers
 - Hands-On!
 - Feb 28 – March 2, NYC
 - Feb 28 – March 2, Melville

© 2005 SkillBuilders, Inc.



87

The End

- Thanks for listening!
- Dave Anderson
- dave@skillbuilders.com
- www.skillbuilders.com

© 2005 SkillBuilders, Inc.