

# Oracle Database 11g New Features




Charles Kim

Practice Manager

VISKosity Technologies, LLC

<http://www.viscosity.com>



Leaders in Oracle Vision and Knowledge



# Updates and Additional Information

- Look for updates at <http://www.viscosity.com>
- Look for additional Oracle Database 11g New Features presentations from other local Oracle User Groups
- Look for RAC and ASM presentations
- Visit the DBAExpert Blog Site: <http://blog.dbaexpert.com>


➤ Download  
Viscosity's RAC  
Pocket Reference  
Guide

➤ Download  
Viscosity's ASM  
Pocket Reference  
Guide



# Audience Poll

- Who is running 11g today in Production? Development?
- Who goes from terminal release to terminal release of Oracle?
- Who is just upgrading to 10gR2 right now?
- Who is still running Oracle 9i?
- Who is considering upgrading to 11g within the next 6 months? 12 months?
- Who is waiting for 11gR2 before considering the upgrade?





# Charles' 10 New Features

1. Real Applications Testing (RAT)
2. Enhanced Security out-of-the-box
3. Active Standby Database
4. Snapshot Standby
5. Direct NFS
6. SecureFiles
7. Result Cache
8. RMAN Recovery Advisor
9. Total Recall - Flashback Data Archive
10. Additions to Extended Composite Partitioning

☰ Editions	MORE \$\$\$
Overview	
Enterprise Edition	
☰ Enterprise Options	
In-Memory DB Cache	
Real Application Testing	
Advanced Compression	
Total Recall	
Active Data Guard	
Real App. Clusters	
Manageability	
Partitioning	
Content Database	
OLAP	
Data Mining	
Database Vault	
Advanced Security	
Label Security	
Spatial	
Standard Edition	
Standard Edition One	
Express Edition	





# Sam's Top 10 New Features

1. Database replay
2. SQL Performance Analyzer
3. Database Recovery Advisor
4. Flashback data archive and the flashback transaction feature
5. Partitioning enhancements such as interval, system, reference, virtual partitioning and the new composite partitioning schemes
6. Active database duplication
7. New Diagnostic Framework, ADR, Support Bench, Incident management and incident packaging etc..
8. SecureFiles
9. New security features, including the password and audit related features
10. SQL Plan Management

# Oracle RACSIG - Officers of the Groups



- Erik Peterson – Oracle Liaison
- Rich Niemiec – Acting President & IOUG Liaison
- Carl Dudley - VP & EOUG Liaison
- Julian Dyke - VP & UKOUG Liaison
- TBD – VP & ASOUG Liaison
- TBD - VP & AIOUG Liaison
- Dan Norris – Events Chair
- Sergio Del Rio – Secretary & Hosting



# About the Group

- With the **popularity of RAC & Grid** its increasing use by the Oracle community, we introduced a by users, for users resource for Oracle RAC information in **2003**.
- The Oracle Real Application Clusters Special Interest Group (RAC SIG) is comprised of people that want to **exchange knowledge** about Oracle RAC.
- We have arranged **panel discussions, technical presentations, and live web seminars** to provide answers to your technical questions about Oracle RAC.
- We maintain a **website** where you can find technical whitepapers, presentation materials, and best practices documentation in addition to a forum where anyone can read and post questions and answers.

# About the Group

- Oracle Corporation and the RAC SIG work together to bring you **live webcasts** with a schedule that started in July 2004 and providing one to two technical sessions per month.
- The past web seminar sessions (which are excellent) were recorded and are **available for on-demand playback** any time. To view webcasts, register with our site and provide some basic information about your Oracle environment.
- To be included on the **list server**, you must be a **Member of IOUG** and subscribe to the RAC SIG mailing list. Your IOUG profile contains a section where you can indicate what SIGs you belong to.



# We are Global

**RAC SIG Event @  
Oracle World  
X-Treme Ev**

**RAC SIG Event @  
IOUG – LIVE  
Toronto - April  
2004**

**RAC SIG Event @  
UKOUG  
November 2004**

**RAC SIG Event @  
Collaborate 2007  
New Forums Launch!**

**RAC SIG Launch  
EOUG @ Oracle World  
Paris - October 2003**

**RAC SIG Launch –@  
AIOUG, Bangalore  
India - October 2007**

**RAC SIG Discussion Forum  
Launch  
September 2003**

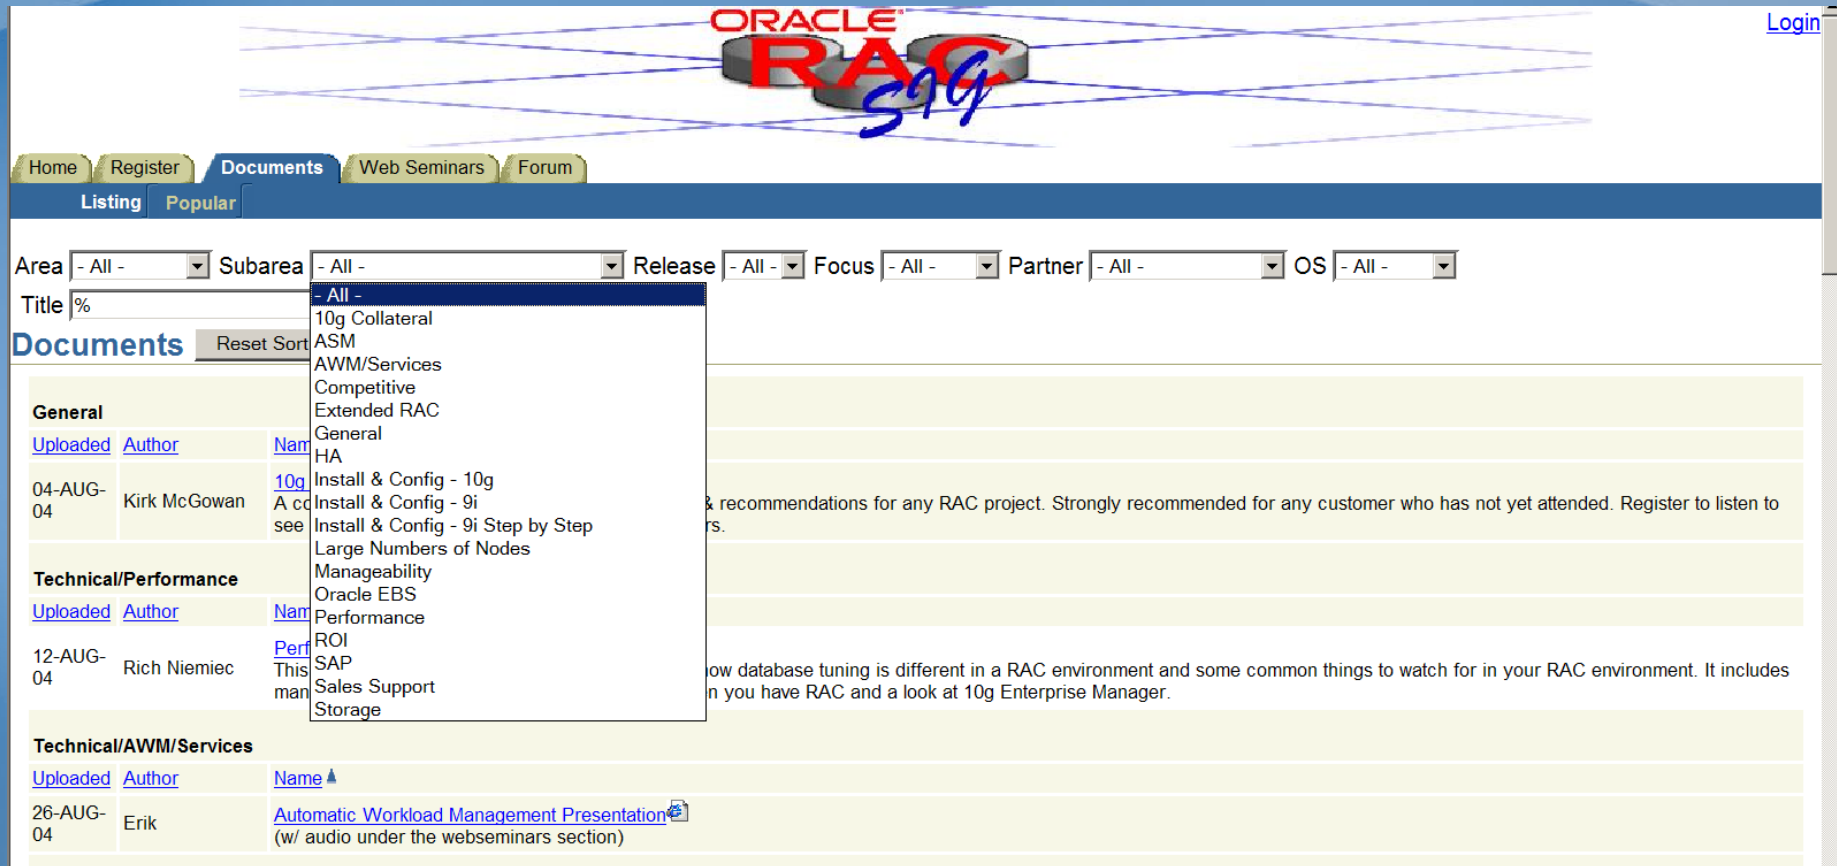
**Logo and Web database Launch  
May 2004**

**RACSIG website uses Oracle HTMLDB**

**RAC SIG Launch  
AUSOUG Conference  
Sydney & Melbourne  
October 2004**

# Website – www.oracleracsig.org

- Documents – Search for what you're looking for...



The screenshot shows the Oracle RAC SIG website. At the top is the Oracle RAC logo with '599' handwritten in blue. Below the logo is a navigation bar with links: Home, Register, Documents (selected), Web Seminars, and Forum. Under 'Documents', there are sub-links for Listing and Popular. A search filter bar is visible with dropdowns for Area, Subarea, Release, Focus, Partner, and OS. The 'Area' dropdown is open, showing a list of categories including General, Technical/Performance, and Technical/AWM/Services. The main content area displays a list of documents with columns for Uploaded, Author, and Name. The first document is '10g Collateral' by Kirk McGowan, dated 04-AUG-04. The second document is 'Automatic Workload Management Presentation' by Erik, dated 26-AUG-04.

ORACLE<sup>®</sup>  
**RAC**  
599

Home Register Documents Web Seminars Forum

Listing Popular

Area - All - Subarea - All - Release - All - Focus - All - Partner - All - OS - All -

Title %

**Documents** Reset Sort

**General**

Uploaded Author Name

04-AUG-04 Kirk McGowan 10g  
A coll  
see

**Technical/Performance**

Uploaded Author Name

12-AUG-04 Rich Niemiec Perf  
This man

**Technical/AWM/Services**

Uploaded Author Name

26-AUG-04 Erik Automatic Workload Management Presentation (w/ audio under the webseminars section)

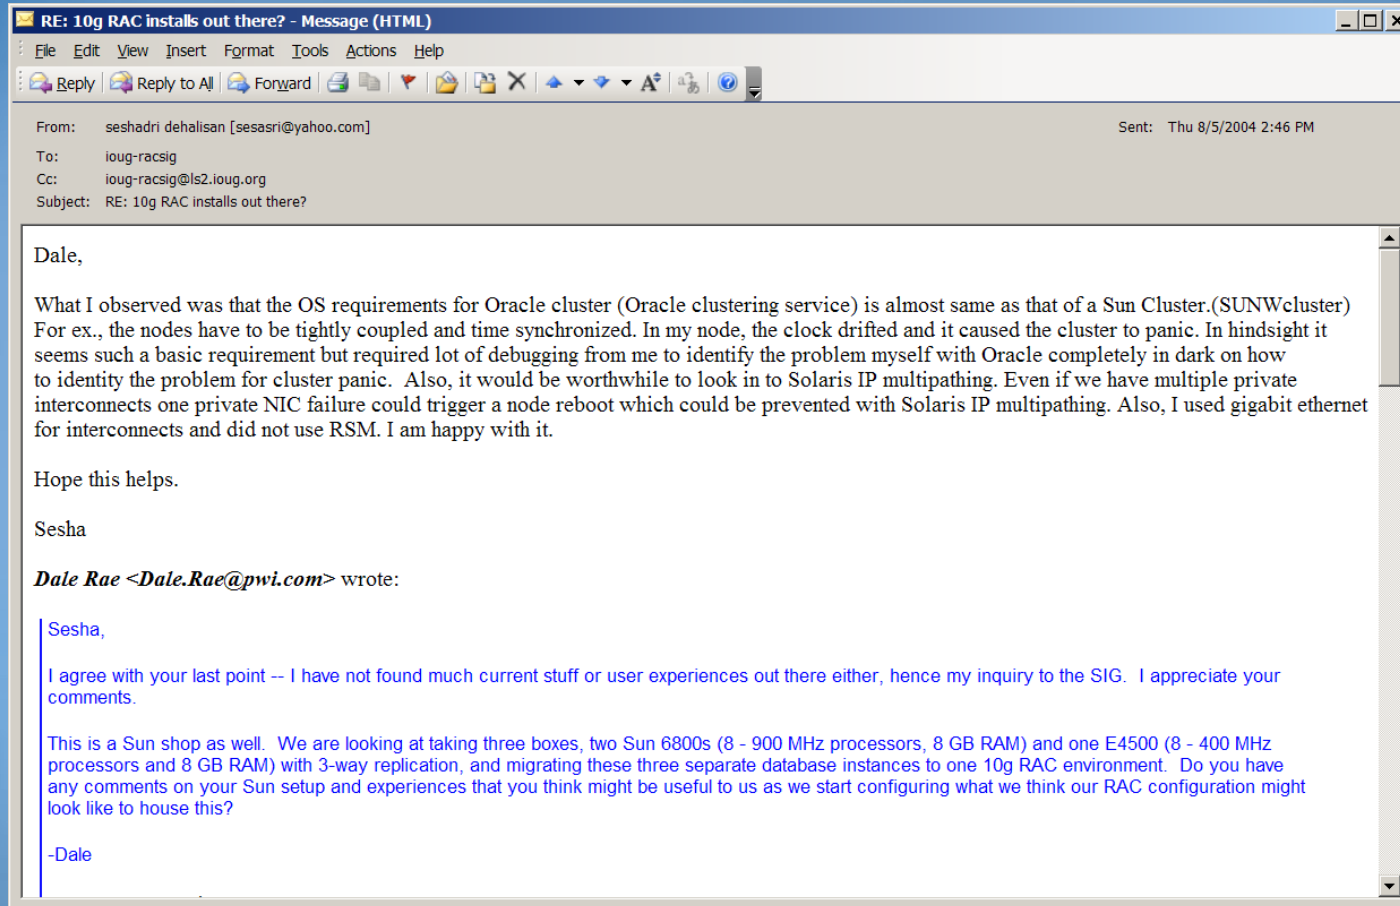
10g Collateral  
ASM  
AWM/Services  
Competitive  
Extended RAC  
General  
HA  
Install & Config - 10g  
Install & Config - 9i  
Install & Config - 9i Step by Step  
Large Numbers of Nodes  
Manageability  
Oracle EBS  
Performance  
ROI  
SAP  
Sales Support  
Storage

& recommendations for any RAC project. Strongly recommended for any customer who has not yet attended. Register to listen to  
rs.

ow database tuning is different in a RAC environment and some common things to watch for in your RAC environment. It includes  
n you have RAC and a look at 10g Enterprise Manager.

# ListServe

- Great email exchange when you register:



# We need your help!



- Tell everyone you know about the group!
- We need volunteers to help with our website
  - Application Express - APEX (Used to be HTML DB)
  - Website Administrators & Content Moderators
  - Document Librarians
  - Book Reviewers
  - Representatives in local & regional groups
  - Web Seminar Presenters
- Help us grow and share the knowledge in the forum
- We need your input to make us better (events/services)





# Desupport

The following components are deprecated:

- iSQL\*Plus
- Oracle Workflow
- Oracle Enterprise Manager Java Console
- Oracle Data Mining Scoring Engine
- Raw storage support (installer only)



# Redhat / OEL

## Linux RPM Requirements

Additional RPMs from Oracle Database 10g R2

Can use up2date to install RPMs:

### 1. elfutils-libelf-devel

```
elfutils-libelf-0.97.1-5.i386.rpm: ##### Done.  
elfutils-libelf-devel-0.97.1-5.i386.rpm: ##### Done.  
elfutils-0.97.1-5.i386.rpm: ##### Done.
```

### 2. libaio-devel

### 3. unixODBC


### 4. unixODBC-devel

For complete list of all RPMs  
required for Oracle Database  
11g, visit  
<http://blogs.dbaexpert.com>

# Linux Recipes for Oracle DBAs



Look for it at:  
Collaborate 2009







# Installation

- Set ORACLE\_BASE
  - Underscore Init Parameter
  - Default directory for DIAG\_DEST
- Software installation – 3.3GB for base installation
- Change in default oraInventory location
- SQL\*Developer as part of Oracle Database 11g
- APEX included
- OWB included

**ORACLE 11g**  
DATABASE

### Specify Inventory directory and credentials

You are starting your first installation on this host. As part of this install, you need to specify a directory for installer files. This is called the "inventory directory". Within the inventory directory, the installer automatically sets up subdirectories for each product to contain inventory data and will consume typically 150 Kilobytes per product.

Enter the full path of the inventory directory:

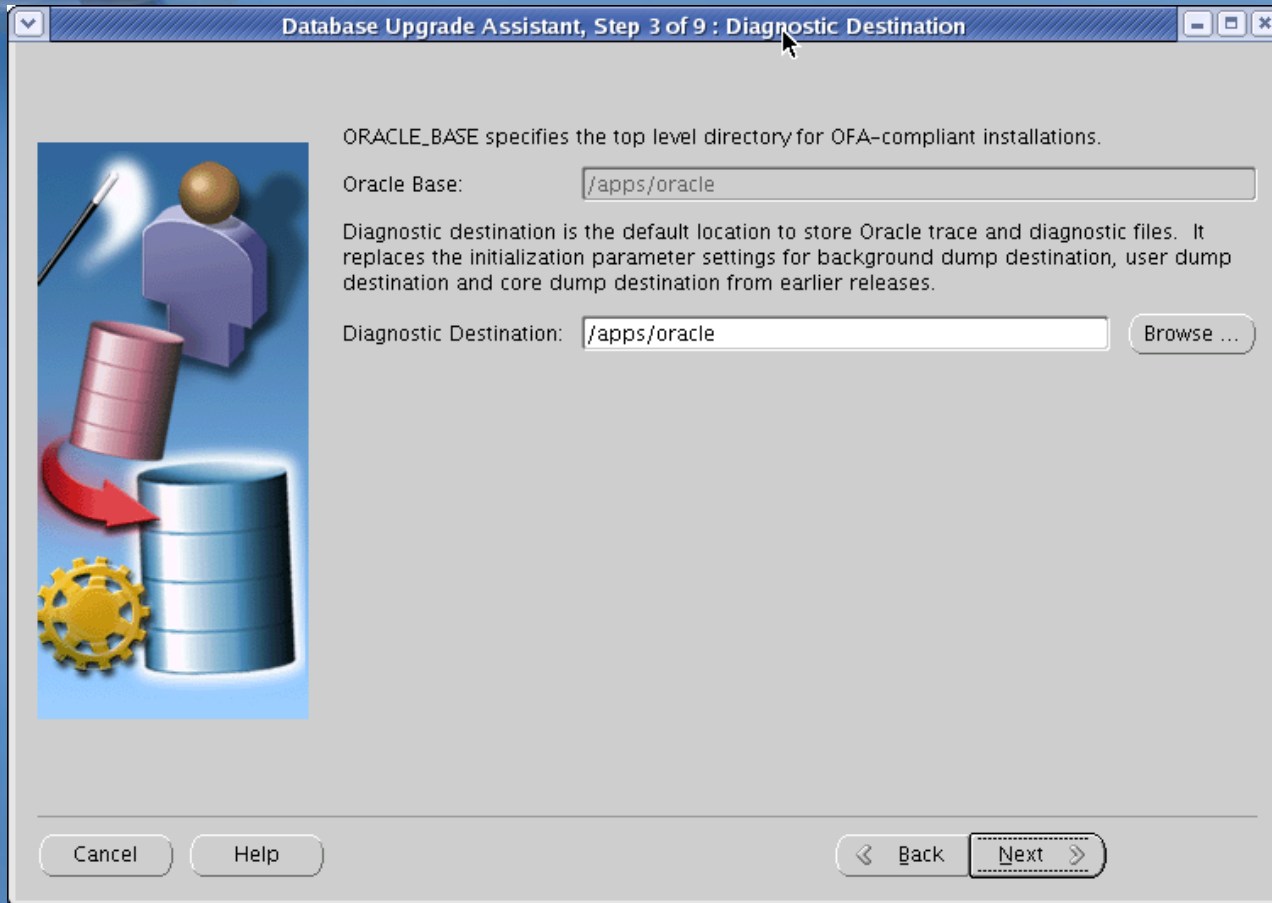
You can specify an Operating System group that has write permission to the above inventory directory. You can leave the field blank if you want to perform the above operations as a Superuser.

Specify Operating System group name:





# dbua - Upgrade



Need to backup before upgrade

Consider using dbua instead of the manual upgrade

DBUA has been tried and tested thoroughly since 10gR1 by IOUG Beta Testers and Beta Program

DBUA upgrade eBusiness Suite flawlessly without errors

**How many of you are using DBUA for upgrades?**



# Upgrade

Database Upgrade Assistant, Step 4 of 9 : Move Database Files

Specify if you want to move the database files during the upgrade process.

☐ Do Not Move Database Files as Part of Upgrade

☒ Move Database Files during Upgrade

Select the storage mechanism you would like to use for the database.

☐ File System  
Use the File System for Database storage.

☒ Automatic Storage Management (ASM)

Automatic Storage Management simplifies database storage administration and optimizes database layout for I/O performance. To use this option, an ASM instance should exist on this host with configured disk groups, otherwise, please use DBCA to create an ASM instance and necessary disk groups on this host then re-start the DBUA.

Cancel Help Back Next



# Recommended upgrade path

Databases newer than 9.2.0.4 can upgrade directly to 11g

Source Database	Target Database
9.2.0.4 or higher	11.1.0.6
10.1.0.2.0 or higher	11.1.0.6
10.2.0.1.0 or higher	11.1.0.6

1. 7.3.3 (or lower) ➤ 7.3.4 ➤ 9.2.0.8 ➤ 11.1
2. 8.0.5 (or lower) ➤ 8.0.6 ➤ 9.2.0.8 ➤ 11.1
3. 8.1.7 (or lower) ➤ 8.1.7.4 ➤ 9.2.0.8 ➤ 11.1
4. 9.0.1.3 (or lower) ➤ 9.0.1.4 ➤ 9.2.0.8 ➤ 11.1
5. 9.2.0.3 (or lower) ➤ 9.2.0.8 ➤ 11.1



# Manual Upgrade Scripts

Three basic scripts:

1. `utlu111i.sql`: The pre-upgrade script.
2. `catupgrd.sql`: The actual upgrade script and is similar to the script in previous releases. The major change now is that it has been restructured to support parallel upgrades of the database.
3. `utlu111s.sql`: This is the upgrade status utility script, which you invoke after completing the database upgrade.





# Clusterware Upgrade

**Same process as upgrading from 10.2.0.1 to 10.2.0.3**

- **Shutdown CRS**

Modify /etc/inittab and comment out last three lines

init q

`$ORA_CRS_HOME/bin/crsctl stop crs`

- **Execute pre-update script from the unzipped software/clusterware/upgrade directory**

`./preupdate.sh -crshome $ORA_CRS_HOME -crsuser  
oracrs`

Best Practice for RAC:  
oracrs, oraasm and  
oradb



# Complete CRS Upgrade

- Install new software binaries on all the RAC nodes
  - ./runInstaller
- From \$ORA\_CRS\_HOME/install directory
  - ./rootupgrade
- Check upgrade status by querying activeversion of the CRS
  - crsctl query crs activeversion

## For Rolling Upgrades

-updateNodeList  
"CLUSTER\_NODES=node1"  
-local \$ORA\_CRS\_HOME



# Key RAC Improvements

No more need for raw devices

- Voting Disk and OCR can be on block devices

RAC Remote Kill

- alter system kill session 'sid, serial#', instance\_id';

Global ADDM report

- See CRS health reports at the entire cluster level

Upgrade Concern

- OCR size doubles during the upgrade
- ocrconfig -replace ocr file\_or\_disk
- ocrconfig -replace ocrmirror file\_or\_disk

Linux

- No more hangcheck-timer → replaced by oproc

```
$ cat /etc/oracle/ocr.loc
ocrconfig_loc=/dev/raw/raw21
ocrmirrorconfig_loc=/dev/raw/raw22
local_only=FALSE

ocrconfig -showbackup
crsctl stop crs
Modify ocr.loc file
ocrconfig -restore <filename>
```



# Hybrid Technology Stack

Possible RAC Topology with  
Oracle Database 11g and 10g

11.1.0.6  
CRS

11.1.0.6  
ASM

10.2.0.3 RDBMS

Possible ASM Topology with  
Oracle Database 11g and 10g

11.1.0.6 ASM

10.2.0.3 RDBMS





# Early Adopters

Where can you introduce Oracle Database 11g to your company?

You start at the DBA Sandbox then slowly promote it to:

1. Infrastructure databases
  - RMAN Repository
  - Grid Control Repository
2. Development
3. Internal databases that does not support external customers (i.e. the IT Web Portal Database, Openview Database, etc.)

Note: Consider participating in the Oracle 11g Release 2 Beta Program



# Customer Case Study – Client standardized on ASM for non-RAC

## 11g ASM with 11g Grid Control / RMAN Repository

Client is not committed to 11.1.0.6 YET but wants to leverage the latest infrastructure technologies

1. ASM 11.1.0.6
2. Grid Control 10gR4 (Certified with Oracle Database 11g)
  - Upgrade path is to install 10gR1 GC on 10.2.0.x database for OMS
  - Apply necessary DST patches
  - Upgrade the database to 11.1.0.6 using DBUA
  - Upgrade GC to 10gR4
3. Oracle Warehouse Builder from 11.1.0.6
4. Application Express from 11.1.0.6
5. Other production databases → 10.2.0.3 Database with 11.1.0.6 ASM



# Basic Minor Thing you will notice

- Change ORACLE\_SID using oraenv in 11g

```
[oracle@rac34 dbs]$ . oraenv
```

```
ORACLE_SID = [DUMMY] ? VISK4
```



The Oracle base for

```
ORACLE_HOME=/apps/oracle/product/11.1.0/
```

```
DB is /apps/oracle
```



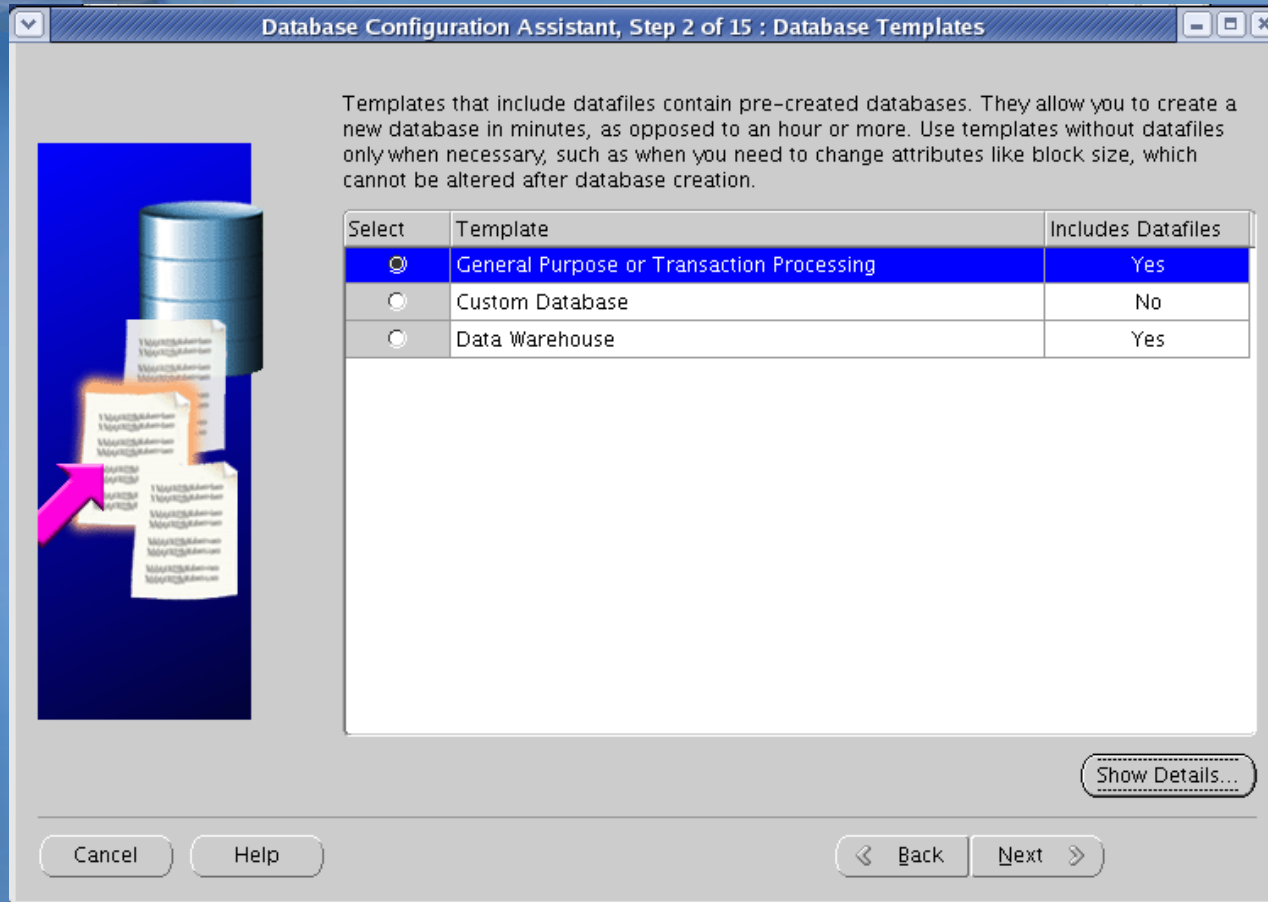
# Database Maintenance







# dbca

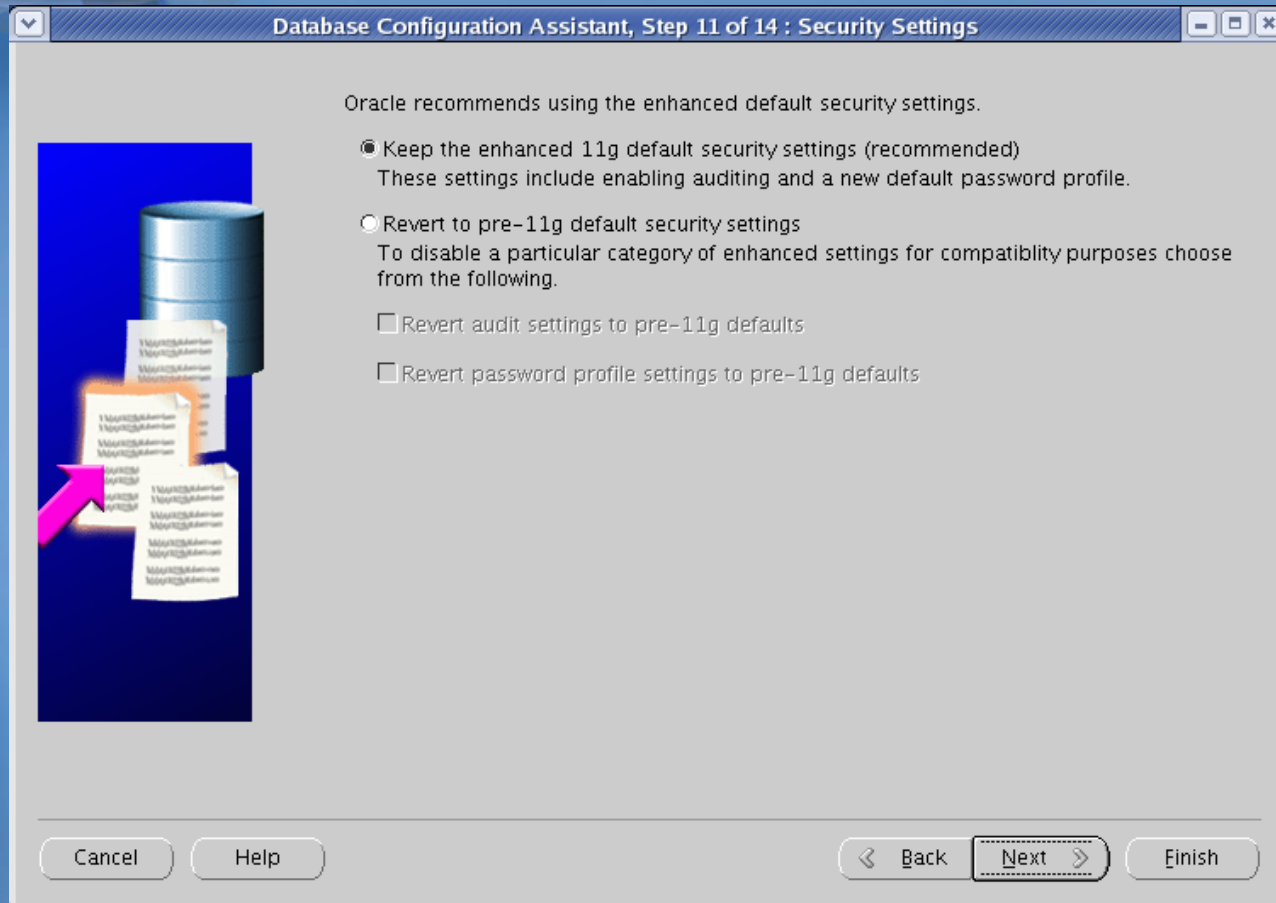


- XML DB is installed
- Data Mining schema is part of catproc.sql
- Data Mining does not appear in DBA\_REGISTRY

How many people use DBCA? Scripts to create databases?



# 11g Enhanced Security



- Default is 11g Enhanced Security

- Can use DBCA to switch back to 10g based security



# Create PFILE / SPFILE From Memory

```
SQL> create spfile='/tmp/spfileVISK4' from memory;
```


File created.

```
SQL> create pfile='/tmp/initVISK4.ora' from memory;
```

File created.



# DDL\_LOCK\_TIMEOUT



- **In one window:**

```
SQL> update emp set sal=sal*1.1 where empno=7934;  
1 row updated.
```

**BUT HE GOES OUT FOR LUNCH AND DOES NOT COMMIT**

- In another window, we set the DDL\_LOCKOUT\_TIMEOUT parameter to 20 seconds:

```
SQL> alter session set ddl_lock_timeout=20;  
Session altered.
```

- **DDL\_LOCK\_TIMEOUT can be set at the system level:**

```
SQL> Alter system set ddl_lock_timeout = 120  
scope = both;  
System altered.
```





# DDL\_LOCK\_TIMEOUT

Next, we will set time on and timing on to validate our tests:

```
SQL> set time on timing on
```

Now, we will finally perform our DDL:

```
14:45:41 SQL> alter table emp add column ssn;  
alter table emp add column ssn
```

\*

ERROR at line 1:

ORA-00054: resource busy and acquire with NOWAIT specified or timeout expired

**Elapsed: 00:00:20.01**



# Locking of Tables

SQL> set time on timing on

00:14:23 SQL> lock table scott.emp2 in exclusive mode wait 10;

lock table scott.emp2 in exclusive mode wait 10

\*

ERROR at line 1:

ORA-00054: resource busy and acquire with NOWAIT specified or timeout expired

Elapsed: 00:00:10.01

00:14:36 SQL> lock table scott.emp2 in exclusive mode nowait;

lock table scott.emp2 in exclusive mode nowait

\*

ERROR at line 1:

ORA-00054: resource busy and acquire with NOWAIT specified or timeout expired

Elapsed: 00:00:00.01

In the other window,  
We give A everyone a  
10% raise

Update emp2 set salary



# Default Value

```
1 select empno, phone_choice,  
2        compute_choice  
3 from emp2  
4 * where rownum < 10  
SQL> /
```

EMPNO	PHONE	COMPUTE
10001	iPhone	Mac Air
10002	iPhone	Mac Air
..		
10007	iPhone	Mac Air
10008	iPhone	Mac Air
10009	iPhone	Mac Air

9 rows selected.

```
1* alter table emp2 add  
   phone_choice varchar2(30)  
   default 'iPhone' not null  
SQL> /
```

Table altered.

Elapsed: 00:00:00.46

```
SQL> alter table emp2 add  
   compute_choice  
   varchar2(30)  
   default 'Mac Air';
```

Table altered.

Elapsed: 00:00:33.60



# Read-Only Tables



```
SQL> create table edba_alerts  
      (alert_id number, description varchar2(1000));  
Table created.
```

```
SQL> alter table edba_alerts read only;  
Table altered.
```

```
SQL> alter table edba_alerts read write;  
Table altered.
```

*Oracle does allow you to  
drop a read only table:*

```
SQL> drop table edba_alerts;  
Table dropped.
```





# Read-Only Tables

- Oracle does not allow inserts/updates or deletes

```
SQL> insert into edba_alerts values  
(1, 'Users are calling the helpdesk  
and reporting that their screens  
are hung. Possibly the archive  
destination ran out of space  
again')
```

\*

```
ERROR at line 1:  
ORA-12081: update operation not  
allowed on table  
"RODBA"."EDBA_ALERTS"
```

- Oracle does not allow you to add another column:

```
SQL> alter table edba_alerts add  
lasted_updated date;  
alter table edba_alerts add  
lasted_updated date
```

\*

```
ERROR at line 1:  
ORA-12081: update operation  
not allowed on table  
"RODBA"."EDBA_ALERTS"
```



# Temporary Tablespace - Create

1\* create temporary tablespace temp2 tempfile '+data' size 1g, '+data' size 1g;  
Tablespace created.

1\* select file#,name,bytes/1024/1024 mb from v\$tempfile;

FILE#	NAME	MB
1	+DATA/visk/tempfile/temp01.dbf	34
2	+DATA/visk/tempfile/temp2.269.650271065	1024
3	+DATA/visk/tempfile/temp2.270.650271065	1024



# Temporary Tablespace Management

```
SQL> alter tablespace temp2 shrink space keep 500m;
```

Tablespace altered.

```
SQL> select file#,name,bytes/1024/1024 mb from v$tempfile;
```

FILE#	NAME	MB
-----		
1	+DATA/visk/tempfile/temp01.dbf	34
2	+DATA/visk/tempfile/temp2.269.650271065	1.0625
3	+DATA/visk/tempfile/temp2.270.650271065	499.9375

```
SQL> alter database tempfile '+DATA/visk/tempfile/temp2.270.650271065' resize  
100m;
```

Database altered.



# Restore Point Enhancements

```
SQL> select current_scn from v$database;
```

```
CURRENT_SCN
```

```
-----
```

```
1685790
```

```
SQL> create restore point rp_24MAR08 as of scn 1650000;
```

Restore point created.

```
1 create restore point rp_24MAR08_2 as of timestamp
```

```
2* to_timestamp('24-MAR-08 12.51.51.500000')
```

```
SQL> /
```

Restore point created.





# Restore Point #2

```
SQL> create restore point rp_25MAR08 preserve;  
Restore point created.
```

```
1* select name, time, scn, preserved from v$restore_point;
```

NAME	TIME	SCN	PRE
-----	-----	-----	----
RP_25MAR08	25-MAR-08 07.21.39. 0000000000 AM	1685942	YES
RP_24MAR08	25-MAR-08 07.15.19. 0000000000 AM	1650000	NO
RP_24MAR08_2	25-MAR-08 07.20.06. 0000000000 AM	1634672	NO



# Automatic Memory Management (AMM)

Let Oracle do it for you

- Init parameter addresses SGA and PGA :
  - memory\_target
  - memory\_max\_target



# Additional Key Database Features

- Automatic Memory Management
  - memory\_target
  - memory\_max\_target
- Database resident connection pooling
  - dbms\_connection\_pool.start\_pool();
  - dbms\_connection\_pool.stop\_pool();
- Minimum invalidation of dependent objects
- Automated database maintenance tasks
  - Default resource manger plan
  - New maintenance windows



# Flashback Data Archive (FDA)

**First, grant the flashback archive administrator:**

```
1* grant flashback  
archive administer to scott;
```

Grant succeeded.

**Connecting as Scott create flashback archive:**

```
1 create flashback archive flash_archive_1  
2 tablespace fa1  
3 quota 100m  
4* retention 1 year;
```

Flashback archive created.





# FDA – Setup at Table Level

```
SQL> alter table emp flashback archive flash_archive_1;  
Table altered.
```

```
SQL> select  empno, ename, sal from emp where empno=7902;  
      EMPNO ENAME          SAL  
-----  
      7902 FORD            3000
```

```
SQL> update emp set sal=3000*1.06 where  empno=7902;  
1 row updated.
```

```
SQL> select  empno, ename, sal from emp where empno=7902;  
      EMPNO ENAME          SAL  
-----  
      7902 FORD            3180
```





# FDA – Query Back in Time

**Check Ford's Salary as of 10 minutes ago:**

```
1  select empno, ename, sal
2  from emp
3  as of timestamp
4  (systimestamp - interval
   '15' minute)
5* where empno=7902
SQL> /
```

EMPNO	ENAME	SAL
7902	FORD	3000

As of timestamp can be:

#1:

*'50' SECOND*

*'50' DAY*

*'12' MONTH*

#2:

as of timestamp

to\_timestamp ('03/27/2008

18:30:00','mm/dd/yyyy hh24:mi:ss')



# FDA – Don't Try To Go Too Far Back

```
1 select empno, ename, sal
2 from emp
3 as of timestamp
4 (systimestamp - interval '1' year)
5* where empno=7902
SQL> /
where empno=7902
*
```

ERROR at line 5:

ORA-08180: no snapshot found based on specified time

ORA-06512: at "SYS.TIMESTAMP\_TO\_SCN", line 1



# FDA – Explain Plan

## **Perform an explain plan**

```
1  explain plan for
2  select empno, ename, sal
3  from emp
4  as of timestamp
5  (systimestamp - interval '15' minute)
6* where empno=7902;
```

Explained.

```
1* select plan_table_output
   from table(dbms_xplan.display
('plan_table',null,'serial'));
```



# FDA – Explain Plan Output

PLAN\_TABLE\_OUTPUT

Plan hash value: 2886175892

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		2	66	7 (0)	00:00
1	VIEW		2	66	7 (0)	00:00
2	UNION-ALL					
* 3	FILTER					
4	PARTITION RANGE SINGLE		1	59	3 (0)	00:00
* 5	TABLE ACCESS FULL	<b>SYS_FBA_HIST_69515</b>	1	59	3 (0)	00:00

PLAN\_TABLE\_OUTPUT

* 6	FILTER					
7	NESTED LOOPS OUTER		1	2060	4 (0)	00:00
* 8	TABLE ACCESS BY INDEX ROWID	EMP	1	45	2 (0)	00:00
* 9	INDEX UNIQUE SCAN	PK_EMP	1		1 (0)	00:00
* 10	VIEW		1	2015	2 (0)	00:00
* 11	TABLE ACCESS FULL	<b>SYS_FBA_TCRV_69515</b>	1	2028	3 (0)	00:00



# FDA - Put the data back as it was 15 minutes ago

```
1  update emp
2  set sal =
3      (select sal from emp
4       as of timestamp
5        (systimestamp - interval '15' minute)
6       where empno=7902)
SQL> /

1 row updated.
```





# FDA Maintenance

- 1 alter flashback archive flash\_archive\_1
- 2 **purge before timestamp**
- 3\* (systimestamp - interval '10' minute)

SQL> /

Flashback archive altered.

Note: You can also purge

- **purge before scn scn#**
- purge all

- 1 alter flashback archive flash\_archive\_1
- 2 **add tablespace trans\_d**
- 3\* quota 100m

SQL> /

Flashback archive altered.

- 1 alter flashback archive flash\_archive\_1
  - 2\* **modify retention** 15 month
- SQL> /

Flashback archive altered.

SQL> alter table emp **no** flashback archive;

Table altered.

- 1\* **drop** flashback archive flash\_archive\_1
- SQL> /

Flashback archive dropped.



# FDA Default and History Table

```
SQL> alter flashback archive flash_archive_1 set default;
```

Flashback archive altered.

```
1* select * from dba_flashback_archive_tables
SQL> /
```

TABLE_NAME	OWNER_NAME	FLASHBACK_ARCHIVE_NAME	ARCHIVE_TABLE_NAME
-----	-----	-----	-----
EMP	SCOTT	FLASH_ARCHIVE_1	SYS_FBA_HIST_69515



# FDA Retention and Quota

```
SQL> select flashback_archive_name, retention_in_days  
       from dba_flashback_archive;
```

```
FLASHBACK_ARCHIVE_NAME  RETENTION_IN_DAYS
```

```
-----  
FLASH_ARCHIVE_1        365
```


```
SQL> select flashback_archive_name,  
       tablespace_name, quota_in_mb  
       2  from dba_flashback_archive_ts;
```

```
FLASHBACK_ARCHIVE_NAME  TABLESPACE_NAME  QUOTA_IN_MB
```

```
-----  
FLASH_ARCHIVE_1        FA1              100  
FBA_UNLIMITED          FA1
```



# Gather Stats – Auto-Sampling



Accuracy is supposed to be comparable to COMPUTE  
ESTIMATE\_PERCENT parameter of GATHER\_TABLE\_STATS

```
1  begin
2    dbms_stats.gather_table_stats(null,
3      'EMP2',
4      estimate_percent => dbms_stats.auto_sample_size);
5* end;
```

SQL> /

PL/SQL procedure successfully completed.

Elapsed: 00:00:02.64



# Gather Stats - Pending

```
SQL> exec dbms_stats.set_table_prefs('SCOTT','EMP2','PUBLISH','FALSE');
```

PL/SQL procedure successfully completed.

```
select dbms_stats.get_prefs (
ownname=>'SCOTT',
tablename=>'EMP2',
pname=>'PUBLISH') PUBLISHED
from dual
SQL> /
```

PUBLISHED

-----

FALSE





# Gather Stats – Query Pending

```
SQL> exec dbms_stats.gather_table_stats('SCOTT','EMP2');  
PL/SQL procedure successfully completed.
```

```
1  select table_name, num_rows,  
2         to_char(last_analyzed,'mm/dd/yy hh24:mi:ss')  
3*  from user_tab_pending_stats  
SQL> /
```

TABLE_NAME	NUM_ROW	TO_CHAR(LAST_ANAL
-----	-----	-----
EMP2	269990	03/22/08 19:53:26



# Gather Stats – Query Pending

```
SQL> select table_name, last_analyzed, num_rows, blocks  
2 from user_tables where table_name='EMP2';
```

TABLE_NAME	LAST_ANAL	NUM_ROWS	BLOCKS
EMP2			

```
SQL> alter session set optimizer_use_pending_statistics=true;  
Session altered.
```


```
SQL> exec dbms_stats.publish_pending_stats('SCOTT','EMP2');  
PL/SQL procedure successfully completed.
```

```
SQL> select table_name, last_analyzed, num_rows, blocks  
2 from user_tables where table_name='EMP2';
```

TABLE_NAME	LAST_ANAL	NUM_ROWS	BLOCKS
EMP2	22-MAR-08	269990	1759



# Oracle Advanced Compression Option



Licensed as a separate option – includes:

- Compression of OLTP data
- Compression of unstructured data with SecureFiles
- Compression using the new RMAN backup compressed configuration (ZLIB)
- Compression of redo data during transmission to the physical standby database
- Compression of Datapump export data
- Oracle claims compression to be 3.5 to 1 ratio



# Oracle Advanced Compression Option

1 Create table compressed\_emp2

2 **compress for all operations**

3\* as select \* from emp2

SQL> /

Table created.

SQL> alter table emp2 compress;

Table altered.

SQL> alter table emp2 nocompress;

Table altered.

*Note:*

*Only new data will be compressed  
Existing data will remain uncompressed*





# Invisible Indexes

```
create table books (book_id number, name varchar2(255), author varchar2(255), isbn varchar2(255));  
create index BOOK_ISBN_INVISIBLE_IDX on books(isbn) invisible;
```

```
insert into books (book_id, name, author, isbn) values (2,'RMAN Recipes','Darl Kuhn','1590598512');  
insert into books (book_id, name, author, isbn) values (1,'Oracle ASM','Nitin  
Vengurlekar','0071496076');  
insert into books (book_id, name, author, isbn) values (3,'Oracle 11g New Features','Charles  
Kim','1590599101');
```

```
1* select index_name,VISIBILITY  
from user_indexes  
where index_name='BOOK_ISBN_INVISIBLE_IDX'  
SQL> /
```

INDEX_NAME	VISIBILITY
-----	-----
BOOK_ISBN_INVISIBLE_IDX	INVISIBLE







# Invisible Indexes –

## optimizer\_use\_invisible\_indexes = TRUE

```
SQL> explain plan for select /*+ index(books BOOK_ISBN_INVISIBLE_IDX) */  
name, author from books where isbn='0071496076';
```

Explained.

```
SQL> select * from table(dbms_xplan.display);
```

```
PLAN_TABLE_OUTPUT
```

```
-----  
-  
Plan hash value: 3295519971  
-----  
| Id | Operation | Name | Rows | Bytes | Cost (%) |  
-----  
| 0 | SELECT STATEMENT | | 1 | 387 | 276 |  
| 1 | TABLE ACCESS BY INDEX ROWID | BOOKS | 1 | 387 | 276 |  
|* 2 | INDEX RANGE SCAN | BOOK_ISBN_INVISIBLE_IDX | 1 | | |  
00:00:01 |
```



# Invisible Indexes –

## `optimizer_use_invisible_indexes=FALSE;`

```
SQL> explain plan for select /*+ index(books BOOK_ISBN_INVISIBLE_IDX) */ name, author  
from books where isbn='0071496076';
```

Explained.

```
SQL> select * from table(dbms_xplan.display);
```

```
PLAN_TABLE_OUTPUT
```

```
-----  
Plan hash value: 2688610195  
-----
```

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		1	387	3 (0)	00:00
* 1	TABLE ACCESS FULL	BOOKS	1	387	3 (0)	00:00

```
-----
```

```
Predicate Information (identified by operation id):  
-----
```



# Invisible Indexes

```
SQL> alter index BOOK_ISBN_INVISIBLE_IDX visible;  
Index altered.
```

```
SQL> select index_name, VISIBILITY  
from user_indexes  
where index_name='BOOK_ISBN_INVISIBLE_IDX';
```

INDEX_NAME	VISIBILITY
-----	-----
BOOK_ISBN_INVISIBLE_IDX	VISIBLE

## Note:

- If you rebuild an invisible index, it will become visible. You will need to manually set it to invisible.



# Invisible Indexes

```
SQL> exec dbms_stats.gather_index_stats(user,'BOOK_ISBN_INVISIBLE_IDX');  
BEGIN dbms_stats.gather_index_stats(user,'BOOK_ISBN_INVISIBLE_IDX'); END;  
*
```

ERROR at line 1:

ORA-00904: : invalid identifier

ORA-06512: at "SYS.DBMS\_STATS", line 14037

ORA-06512: at "SYS.DBMS\_STATS", line 14060

ORA-06512: at line 1

**-- Bug No. 6344547**

```
SQL> alter session set optimizer_use_invisible_indexes=TRUE;  
Session altered.
```

```
SQL> exec dbms_stats.gather_index_stats(user,'BOOK_ISBN_INVISIBLE_IDX');  
PL/SQL procedure successfully completed.
```



# Virtual Columns

```
create table tech_books  
(book_id number,  
author varchar2(255),  
price number,  
discounted_price generated always as(price*.75),  
isbn number)  
/
```

*Can be written using the  
optional VIRTUAL keyword:  
• tech\_price generated  
always as(price\*.75)  
virtual*

```
SQL> desc tech_books
```

Name	Null?	Type
-----		
BOOK_ID		NUMBER
AUTHOR		VARCHAR2 (255)
PRICE		NUMBER
DISCOUNTED_PRICE		NUMBER
ISBN		NUMBER





# Virtual Columns

```
SQL> select * from tech_books;
```

BOOK_ID	AUTHOR	PRICE	DISCOUNT PRICE	ISBN
2	Darl Kuhn	49.99	37.4925	1590598512
1	Nitin Vengurlekar	39.99	29.9925	71496076
3	Charles Kim	39.99	29.9925	1590599101

## Rules:

1. Virtual column is a derived expression
2. Can't create a virtual column as a LOB, RAW or User-defined TYPE
3. Can be derived from columns of the same table
4. Can be derived from constants
5. Can use SQL or user-defined PL/SQL functions
6. Can create an index or partition on a virtual column
7. Can not be on Index Organized and External Tables
8. Indexes against virtual columns are equivalent to function-based indexes



# Cannot apply DML Against Virtual Column

Cannot attempt to insert into a virtual column:

```
SQL> insert into tech_books (book_id, author, price,  
    discounted_price)
```

```
2 values
```

```
3 (4,'Sam Alapati',69.99,55);
```

```
insert into tech_books (book_id, author, price,  
    discounted_price)
```

```
*
```

ERROR at line 1:

ORA-54013: INSERT operation disallowed on virtual columns



# Cannot have virtual columns on virtual column

Cannot have virtual columns based on other virtual columns:

- 1 create table discounted\_books
- 2 (book\_id number,
- 3 author varchar2(255),
- 4 price number,
- 5 tech\_price generated always as(price\*.75),
- 6 discounted\_price as (tech\_price\*.5)
- 7\* )

SQL> /

tech\_price generated always as(price\*.75),  
\*

ERROR at line 5:

ORA-54012: virtual column is referenced in a column expression



# Check Virtual Column

```
SQL> COLUMN data_default FORMAT A50
SQL> SELECT column_name, data_default
       2 FROM user_tab_columns
       3 WHERE table_name='TECH_BOOKS';
```

COLUMN_NAME	DATA_DEFAULT
-----	
BOOK_ID	
AUTHOR	
PRICE	
DISCOUNTED_PRICE	"PRICE"*.75
ISBN	



# DBMS\_COMPARISON

Can compare data between:

1. Tables
2. Single table views
3. MV
4. Synonyms that point to tables, single-table views, and MVs
5. Local and remote database

*•SharePlex Compare and Repair  
•Intended for Streams*

## Requirements

1. Local Database has to be 11.1
2. Remote database has to be 10.1
3. At least one index must be available on the objects being compared
4. Preference is NUMBER datatype
5. Same database character set





# DBMS\_COMPARISON

Restrictions - Does not support:

1. LONG
2. LONG RAW
3. CLOB
4. NCLOB
5. BLOB
6. BFILE
7. User-Defined Types
8. Oracle Supplied Types  
such as XMLTypesTables

As SYS:

```
SQL> grant execute on dbms_comparison to scott;  
Grant succeeded.
```

```
SQL> create table emp_visk as select * from emp;  
Table created.
```

```
1 alter table emp_visk  
2* add constraint emp_visk_pk primary key (empno)  
SQL> /  
Table altered.
```

```
SQL> update emp_visk set sal=sal*1.1;  
14 rows updated.
```

```
SQL> commit;  
Commit complete.
```



# DBMS\_COMPARISON.CREATE\_COMPARISON

BEGIN

DBMS\_COMPARISON.CREATE\_COMPARISON

( comparison\_name => 'compare\_emp\_to\_emp\_visk'

, schema\_name => 'scott'

, object\_name => 'emp'

, dblink\_name => null

, remote\_schema\_name=>'scott'

, remote\_object\_name=>'EMP\_VISK'

);

END;

/

PL/SQL procedure successfully completed.



# DBMS\_COMPARISON.COMPARE

```
DECLARE
  consistent  BOOLEAN;
  scan_info  DBMS_COMPARISON.COMPARISON_TYPE;
BEGIN
  consistent := DBMS_COMPARISON.COMPARE
    (comparison_name => 'compare_emp_to_emp_visk'
    ,scan_info      => scan_info
    ,perform_row_dif => TRUE);
  DBMS_OUTPUT.PUT_LINE('Scan ID: ' || scan_info.scan_id);
  IF consistent=TRUE THEN
    DBMS_OUTPUT.PUT_LINE('No differences were found.');
```

Scan ID: 1  
Differences were found.

PL/SQL procedure successfully completed.

```
ELSE
  DBMS_OUTPUT.PUT_LINE('Differences were found.');
```

Scan ID: 1  
Differences were found.

PL/SQL procedure successfully completed.

```
END IF;
END;
SQL> /
```



# Do Some Inserts and Create Another Comparison

```
1 insert into emp (empno, ename, job, mgr,  
    hiredate, sal, deptno)
```

```
2 values
```

```
3* (8000, 'CKIM', 'SR. DBA', 7902, sysdate,  
    90000, 20)
```

```
SQL> /
```

1 row created.

```
1 insert into emp (empno, ename, job, mgr,  
    hiredate, sal, deptno)
```

```
2 values
```

```
3* (8001, 'SALAPATI', 'SR. DBA', 7902, sysdate,  
    90000, 20)
```

```
SQL> /
```

1 row created.

```
SQL> commit;
```

Commit complete.

```
DECLARE
```

```
    consistent  BOOLEAN;
```

```
    scan_info  DBMS_COMPARISON.COMPARISON_TYPE;
```

```
BEGIN
```

```
    consistent := DBMS_COMPARISON.COMPARE
```

```
        ( comparison_name => 'compare_emp_to_emp_visk2'
```

```
        , scan_info      => scan_info
```

```
        , perform_row_dif => TRUE
```

```
        );
```

```
    DBMS_OUTPUT.PUT_LINE('Scan ID: ' || scan_info.scan_id);
```

```
    IF consistent=TRUE THEN
```

```
        DBMS_OUTPUT.PUT_LINE('No differences were found.');
```

```
    ELSE
```

```
        DBMS_OUTPUT.PUT_LINE('Differences were found.');
```

```
    END IF;
```

```
    END;
```

```
SQL> /
```

Scan ID: 4

Differences were found.

PL/SQL procedure successfully completed.



# Comparison Scan Summary

```
SELECT c.owner, s.scan_id ,c.comparison_name
      ,c.schema_name ,c.object_name
      ,s.current_dif_count DIFF_COUNT
FROM   dba_comparison c
      ,dba_comparison_scan_summary s
WHERE  c.comparison_name = s.comparison_name
AND    s.scan_id in (1,4)
SQL> /
```

OWNER	SCAN_ID	COMPARISON_NAME	SCHEMA	OBJECT	DIFF_COUNT
-----	-----	-----	-----	-----	-----
SCOTT	4	COMPARE_EMP_TO_EMP_VISK2	SCOTT	EMP	16
SYS	1	COMPARE_EMP_TO_EMP_VISK	SCOTT	EMP	14





# Local and Remote Index Values

```
SELECT s.comparison_name,  
       c.column_name  
       ,r.index_value  
,case when r.local_rowid is null  
       then 'No'  
       else 'Yes'  
end local_rowid  
,case when r.remote_rowid is null  
       then 'No'  
       else 'Yes'  
end remote_rowid  
FROM dba_comparison_columns c,  
     dba_comparison_row_dif r  
     ,dba_comparison_scan s  
WHERE c.comparison_name in  
      ('COMPARE_EMP_TO_EMP_VISK',  
       'COMPARE_EMP_TO_EMP_VISK2')  
AND r.scan_id = s.scan_id  
AND s.last_update_time > sysdate - 1  
AND r.status= 'DIF'  
AND c.index_column= 'Y'  
AND c.comparison_name = r.comparison_name  
order by s.comparison_name, r.index_value  
SQL> /
```

COMPARISON_NAME	COLUMN_NAME	INDEX_VAL	LOC	REM
COMPARE_EMP_TO_EMP_VISK	EMPNO	7369	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7499	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7521	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7566	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7654	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7698	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7782	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7788	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7839	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7844	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7876	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7900	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7902	Yes	Yes
COMPARE_EMP_TO_EMP_VISK	EMPNO	7934	Yes	Yes
COMPARE_EMP_TO_EMP_VISK2	EMPNO	7369	Yes	Yes
[. .]				
COMPARE_EMP_TO_EMP_VISK2	EMPNO	7934	Yes	Yes
COMPARE_EMP_TO_EMP_VISK2	EMPNO	8000	Yes	No
COMPARE_EMP_TO_EMP_VISK2	EMPNO	8001	Yes	No

30 rows selected.



# Converge – Synchronize the Tables

```
DECLARE
  scan_info  DBMS_COMPARISON.COMPARISON_TYPE;
BEGIN
  DBMS_COMPARISON.CONVERGE
    ( comparison_name => 'COMPARE_EMP_TO_EMP_VISK'
    , scan_id         => 1
    , scan_info       => scan_info
    , converge_options => DBMS_COMPARISON.CMP_CONVERGE_LOCAL_WINS);
  DBMS_OUTPUT.PUT_LINE('Local Rows Merged: ' || scan_info.loc_rows_merged);
  DBMS_OUTPUT.PUT_LINE('Remote Rows Merged: ' || scan_info.rmt_rows_merged);
  DBMS_OUTPUT.PUT_LINE('Local Rows Deleted: ' || scan_info.loc_rows_deleted);
  DBMS_OUTPUT.PUT_LINE('Remote Rows Deleted: ' || scan_info.rmt_rows_deleted);
END;
SQL> /
Local Rows Merged: 0
Remote Rows Merged: 14
Local Rows Deleted: 0
Remote Rows Deleted: 0
PL/SQL procedure successfully completed.
```



# Recheck

```
DECLARE
v_comp BOOLEAN;
BEGIN
v_comp := DBMS_COMPARISON.RECHECK
    (comparison_name=>'COMPARE_EMP_TO_EMP_VISK2',
    scan_id=>4);
IF v_comp = TRUE then
    DBMS_OUTPUT.PUT_LINE('No Deltas Found During Recheck');
ELSE
    DBMS_OUTPUT.PUT_LINE('Deltas Found During Recheck');
END IF;
END;
SQL> /
Deltas Found During Recheck
```

Rows are inserted after the last comparison

```
1 insert into emp_visk (empno, ename,
job, mgr, hiredate, sal, deptno)
2 values
3* (8002, 'DKUHN', 'SR. DBA', 7902,
sysdate, 95000, 20)
SQL> /
1 row created.
```

```
SQL> update emp_visk set sal=99000
where empno=8000;
1 row updated.
```

```
SQL> commit;
Commit complete.
```

PL/SQL procedure successfully completed.



# Maintenance

```
BEGIN
```

```
DBMS_COMPARISON.PURGE_COMPARISON(comparison_name=>'COMPARE_EMP_TO_EMP_VISK');
```

```
END;
```

```
/
```

PL/SQL procedure successfully completed.

```
SQL> exec dbms_comparison.drop_comparison('COMPARE_EMP_TO_EMP_VISK');
```

PL/SQL procedure successfully completed.

```
SQL> exec dbms_comparison.drop_comparison('COMPARE_EMP_TO_EMP_VISK2');
```

PL/SQL procedure successfully completed.






# ADR (Automatic Diagnostic Repository) Key Directories

```
gc.dbaexpert.com:/apps/oracle/diag
```

```
VISK > ls -ltr
```

```
total 72
```

```
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 ofm
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 netcman
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 lsnrctl
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 diagtool
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 crs
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 clients
drwxrwxr-x 2 oracle oinstall 4096 Mar 3 16:36 asm
drwxrwxr-x 3 oracle oinstall 4096 Mar 4 18:15 rdbms
drwxrwxr-x 3 oracle oinstall 4096 Mar 6 19:16 tnslnsr
```



## Useful cd command:

```
cd $ORACLE_BASE/diag/rdbms/`echo $ORACLE_SID | tr '[:upper:]' '[:lower:]'`/$ORACLE_SID/rdbms
alias rdbms='cd $ORACLE_BASE/diag/rdbms/`echo $ORACLE_SID | tr '[:upper:]' '[:lower:]'`/$ORACLE_SID'
```





# Diag RDBMS

gc.dbaexpert.com:/apps/oracle/diag/rdbms/visk/VISK

+ASM > ls -ltr

total 104

```
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 sweep
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 stage
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 incpkg
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 incident
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 hm
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 cdump
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 metadata
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 lck
drwxr-x--- 2 oracle oinstall 4096 Mar  4 18:15 alert
drwxr-x--- 2 oracle oinstall 4096 Mar  6 19:14 ir
drwxr-x--- 2 oracle oinstall 20480 Mar  8 15:40 trace
```



# v\$diag\_info

```
1* select * from v$diag_info
SQL> /
```

INST_	NAME	VALUE
-----	-----	-----
1	Diag Enabled	TRUE
1	ADR Base	/apps/oracle
1	ADR Home	/apps/oracle/diag/rdbms/visk/VISK
1	Diag Trace	/apps/oracle/diag/rdbms/visk/VISK/trace
1	Diag Alert	/apps/oracle/diag/rdbms/visk/VISK/alert
1	Diag Incident	/apps/oracle/diag/rdbms/visk/VISK/incident
1	Diag Cdump	/apps/oracle/diag/rdbms/visk/VISK/cdump
1	Health Monitor	/apps/oracle/diag/rdbms/visk/VISK/hm
1	Default Trace File	/apps/oracle/diag/rdbms/visk/VISK/trace/VISK_ora_23558.trc
1	Active Problem Count	0
1	Active Incident Count	0



# Old and New Mapping

Let's start with a new initialization parameter:

```
SQL> show parameter diag
```

NAME	TYPE	VALUE
-----	-----	-----
diagnostic_dest	string	/apps/oracle

Data	Deprecated Init.ora Parameter	ADR location as of 11g
ADR_HOME		\${diagnostic_dest}/diag/rdbms/lowerSID/upperSID
Core Dump	CORE_DUMP_DEST	\$ADR_HOME/cdump
Alert log data	BACKGROUND_DUMP_DEST	\$ADR_HOME/trace/alert_upperSID.log
Alert Log XML file		\$ADR_HOME/alert/log.xml
Background process trace	BACKGROUND_DUMP_DEST	\$ADR_HOME/trace
User process trace	USER_DUMP_DEST	\$ADR_HOME/trace



# log.xml

**This is how the log.xml file looks like:**

```
rac34:/apps/oracle/diag/rdbms/visk4/VISK4/alert
VISK4 > tail -10 log.xml
  module='' pid='6596'>
    <txt>Thread 1 advanced to log sequence 15
  </txt>
</msg>
<msg time='2008-03-19T10:20:26.128-05:00'
org_id='oracle' comp_id='rdbms'
  client_id='' type='UNKNOWN' level='16'
  module='' pid='6596'>
    <txt>  Current log# 3 seq# 15 mem# 0:
+FRA/visk4/onlineelog/group_3.259.649601669
  </txt>
</msg>
```



# Key ADR Directories

- alert – For the alert log in XML format
- cdump – For core dump files. This is equivalent of the core\_dump\_dest from the previous versions
- hm - For Health Monitor runs checks to store many components and output
- incident – For all incidents dumps. Each incident is stored in a separate subdirectory.
- incpkg – For package incidents support files
- metadata – For metadata about problems, incidents, packages, etc.
- trace – For user traces and background traces + the traditional text version of the alert log





# adrci

```
VISK > adrci
```

```
ADRCI: Release 11.1.0.6.0 - Beta on Tue Mar 18 08:40:51 2008
```

```
Copyright (c) 1982, 2007, Oracle. All rights reserved.
```

```
ADR base = "/apps/oracle"
```

## **adrci command help**

```
VISK > adrci -help
```

```
Syntax:
```

```
adrci [-help] [script=script_filename]  
      [exec = "one_command [;one_command;...]" ]
```

Options	Description	(Default)
---------	-------------	-----------

-----

--

script	script file name	(None)
--------	------------------	--------

help	help on the command options	(None)
------	-----------------------------	--------

exec	exec a set of commands	(None)
------	------------------------	--------

-----



# adrci -Base and Home

```
adrci> show base
```

```
ADR base is "/apps/oracle"
```

```
adrci> show homes
```

```
ADR Homes:
```

```
diag/rdbms/visk4/VISK4
```

```
diag/clients/user_unknown/host_411310321_11
```

```
diag/clients/user_oracle/host_1204296188_11
```

```
diag/asm/+asm/+ASM
```

```
diag/tnslsnr/rac34/listener
```

**Set Home to the database:**

```
adrci> set home diag/rdbms/visk4/VISK4
```



# adrci – Retention Policy

## Show retention policy:

```
adrci> show control
```

```
ADR Home = /apps/oracle/diag/rdbms/visk4/VISK4:
```

```
*****
```

ADRID	SHORTP_POLICY	LONGP_POLICY	LAST_MOD_TIME	LAST_AUTOPRG_TI
ME	LAST_MANUPRG_TIME	ADRDIR_VERSION	ADRSCHM_VERSION	ADRSCH
MV_SUMMARY	ADRALERT_VERSION	CREATE_TIME		

2813981488	<b>720</b>	<b>8760</b>	2008-03-17 12:51:39.378653 -	
05:00	1	2	0 1	2008-03-17 12:51:39.378653 -05:00

```
1 rows fetched
```

- Basically, the short policy=720 (1 month - The incident files and dumps retention policy - ( Default is one month))
- And the long policy=8760 (1 Year - The incident metadata retention policy ( default is 1 year ) )
- Specified in hours

Now, let's change the short and long retention policies:

- `adrci> set control (shortp_policy=336)` ← 2 Weeks
- `adrci> set control (longp_policy=4380)`



# adrci – show alert

**Show alert (no arguments) - shows the alert log in the default editor:**

**adrci> show alert**

ADR Home = /apps/oracle/diag/rdbms/visk4/VISK4:

\*\*\*\*\*

Output the results to file: /tmp/alert\_6731\_30868\_VISK4\_1.ado

**Show alert without setting your home - Prompts for a number that represents your home:**

**adrci> show alert**

Choose the alert log from the following homes to view:

1: diag/rdbms/visk4/VISK4

2: diag/clients/user\_unknown/host\_411310321\_11

3: diag/clients/user\_oracle/host\_1204296188\_11

4: diag/asm/+asm/+ASM

5: diag/tnslsnr/rac34/listener

Q: to quit

Please select option: 1

Output the results to file: /tmp/alert\_7068\_30865\_VISK4\_1.ado

Please select option: Q



# adrci – show alert #2

**Show last 30 lines of the alert log:**

```
adrci> show alert -tail 30
```

**Do an equivalent of a tail -f alert\$SID.log file:**

```
adrci> show alert -tail -f
```

Control-c to break

**Show just ORA- messages:**

```
adrci> set home diag/rdbms/visk4/VISK4
```

```
adrci> show alert -p "message_text like '%ORA-%'"
```

ADR Home = /apps/oracle/diag/rdbms/visk4/VISK4:

\*\*\*\*\*

**Something cool –**

**display your alert in IE or Firefox:**

```
adrci> set editor mozilla
```

```
adrci> show alert
```

Put it back:

```
adrci> set editor vi
```





# adrci – Script

## Scripting it:

```
adrci exec="show homes; set home diag/rdbms/visk4/VISK4; show alert -tail 100"
```

```
alias tail100='adrci exec="show homes; set home diag/rdbms/visk4/VISK4; show alert -tail 100"'
```

## Now put it in an ADR script file:

```
adrci script=adr_tail100.txt
```

```
VISK4 > cat adr_tail100.txt
```

```
show homes; set home diag/rdbms/visk4/VISK4; show alert -tail 100
```

### adrci with here documents

```
VISK > cat adrci.ksh
```

```
adrci << EOF
```

```
set hompath diag/rdbms/visk/VISK
```

```
show alert -tail 1000
```

```
exit
```

```
EOF
```



# adrci – show tracefile

**adrci> show tracefile**

```
diag/rdbms/visk4/VISK4/trace/VISK4_vktm_6578.trc
diag/rdbms/visk4/VISK4/trace/VISK4_arc2_16052.trc
..
diag/rdbms/visk4/VISK4/trace/VISK4_vktm_15450.trc
diag/rdbms/visk4/VISK4/trace/VISK4_m000_6814.trc
```

**Equivalent to ls -lt:**

**adrci> show tracefile -rt**

```
19-MAR-08 09:20:26 diag/rdbms/visk4/VISK4/trace/VISK4_arc1_6634.trc
19-MAR-08 09:20:26 diag/rdbms/visk4/VISK4/trace/alert_VISK4.log
..
19-MAR-08 09:09:28 diag/rdbms/visk4/VISK4/trace/VISK4_lgwr_6596.trc
19-MAR-08 09:09:20 diag/rdbms/visk4/VISK4/trace/VISK4_vktm_6578.trc
19-MAR-08 09:09:19 diag/rdbms/visk4/VISK4/trace/VISK4_ora_6558.trc
18-MAR-08 12:47:27 diag/rdbms/visk4/VISK4/trace/VISK4_ora_7242.trc
18-MAR-08 12:47:27 diag/rdbms/visk4/VISK4/trace/VISK4_vktm_6860.trc
```



# adrci – Problem vs. Incident

## **Problem:**

- A critical error in the database.
- ADR records a problem with a problem key
- There is a one to many relationship between a problem and incidents.

## **Incidents:**

- An incident is a single occurrence of a problem.
- ADR identifies incidents by a incident ID.
- Not all incidents are captured.
- Oracle, by default, performs flood control by limiting the number of dumps per hour to 5.



# Adrci – show problem for all homes

**Show problem without setting your home - get single report for all homes:**

adrci> show problem

ADR Home = /apps/oracle/diag/rdbms/visk4/VISK4:

\*\*\*\*\*

0 rows fetched

ADR Home = /apps/oracle/diag/clients/user\_unknown/host\_411310321\_11:

\*\*\*\*\*

0 rows fetched

ADR Home = /apps/oracle/diag/clients/user\_oracle/host\_1204296188\_11:

\*\*\*\*\*

0 rows fetched

ADR Home = /apps/oracle/diag/asm/+asm/+ASM:

\*\*\*\*\*

0 rows fetched

ADR Home = /apps/oracle/diag/tnslsnr/rac34/listener:

\*\*\*\*\*

0 rows fetched



# Show Problem / Incident

**adrci> show problem**

ADR Home = /apps/oracle/diag/clients/user\_oracle/host\_282291381\_11:

\*\*\*\*\*

PROBLEM_ID	PROBLEM_KEY	LAST_INCIDENT	LASTINC_TIME
------------	-------------	---------------	--------------

1	oci 24550 [3]	1	2008-02-27 15:34:18.974522 -06:00
---	---------------	---	-----------------------------------

1 rows fetched

**adrci> show incident**

ADR Home = /apps/oracle/diag/clients/user\_oracle/host\_282291381\_11:

\*\*\*\*\*

INCIDENT_ID	PROBLEM_KEY	CREATE_TIME
-------------	-------------	-------------

1	oci 24550 [3]	2008-02-27 15:34:18.974522 -06:00
---	---------------	-----------------------------------

1 rows fetched





# adrci - IPS

**adrci> show incident**

ADR Home = /apps/oracle/diag/clients/user\_oracle/host\_282291381\_11:

\*\*\*\*\*

INCIDENT_ID	PROBLEM_KEY	CREATE_TIME
1	oci 24550 [3]	2008-02-27 15:34:18.974522 -06:00

1 rows fetched

adrci> set home diag/clients/user\_oracle/host\_282291381\_11

adrci> ips create package incident 1

Created package 1 based on incident id 1, correlation level typical

adrci> ips add incident 1 package 1

Added incident 1 to package 1

adrci> ips generate package 1 in /tmp

Generated package 1 in file /tmp/oci245503\_20080319153442\_COM\_2.zip,  
mode complete



# Contents of Package

```
/tmp > unzip -l oci245503_20080319153442_COM_2.zip
```

```
Archive:  oci245503_20080319153442_COM_2.zip
```

Length	Date	Time	Name
-----	----	----	----
0	02-27-08	15:34	diag/clients/user_oracle/host_282291381_11/
106736	02-27-08		
15:34			diag/clients/user_oracle/host_282291381_11/incident/incdir_1/ora_24864_3086436032_i1.trc
308	02-27-08		
15:34			diag/clients/user_oracle/host_282291381_11/trace/ora_24864_3086436032.trc
206	02-27-08		
15:34			diag/clients/user_oracle/host_282291381_11/trace/ora_24864_3086436032.trm
8040	03-19-08		
15:35			diag/clients/user_oracle/host_282291381_11/alert/log.xml
465	02-27-08		
15:34			diag/clients/user_oracle/host_282291381_11/trace/sqlnet.log
...			
0	02-27-08	15:34	diag/clients/user_oracle/host_282291381_11/trace/
539	03-19-08	15:39	metadata.xml
-----			-----
232206			67 files



# Health Monitor Check

Two kinds:

1. Reactive - Automatically executed by the diagnostic infrastructure
2. Manually by the DBA

1\* select name  
from V\$HM\_CHECK;

NAME

```
-----  
HM Test Check  
DB Structure Integrity Check  
Data Block Integrity Check  
Redo Integrity Check  
Logical Block Check  
Transaction Integrity Check  
Undo Segment Integrity Check  
All Control Files Check  
CF Member Check  
All Datafiles Check  
Single Datafile Check  
Log Group Check  
Log Group Member Check  
Archived Log Check  
Redo Revalidation Check  
IO Revalidation Check  
Block IO Revalidation Check  
Txn Revalidation Check  
Failure Simulation Check  
Dictionary Integrity Check  
21 rows selected.
```



# Health Monitor Check – Manual Run

## Health Monitor Run:

```
adrci> show hm_run -p "run_id=1"
```

```
*****
```

```
HM RUN RECORD 1
```

```
*****
```

RUN_ID	1
RUN_NAME	HM_RUN_1
CHECK_NAME	DB Structure Integrity Check
NAME_ID	2
MODE	2
START_TIME	2008-03-17 12:54:17.240082 -05:00
RESUME_TIME	<NULL>
END_TIME	2008-03-17 12:54:43.528348 -05:00
MODIFIED_TIME	2008-03-17 12:54:43.528348 -05:00
TIMEOUT	0
FLAGS	0
STATUS	5
SRC_INCIDENT_ID	0
NUM_INCIDENTS	0
ERR_NUMBER	0
REPORT_FILE	<NULL>

```
1 rows fetched
```

## Manually execute a health monitor run:

```
SQL> exec dbms_HM.RUN_CHECK  
      ('Dictionary Integrity Check');
```

```
PL/SQL procedure successfully completed.
```



# Run Report

```
adrci> create report hm_run hm_run_1
```

```
SQL> set long 1000000
```

```
1* select dbms_hm.get_run_report('HM_RUN_1') from dual;  
DBMS_HM.GET_RUN_REPORT('HM_RUN_1')
```

---

## Basic Run Information

Run Name	: HM_RUN_1
Run Id	: 1
Check Name	: DB Structure Integrity Check
Mode	: REACTIVE
Status	: COMPLETED
Start Time	: 2008-03-17 12:54:17.240082 -05:00
End Time	: 2008-03-17 12:54:43.528348 -05:00
Error Encountered	: 0

..

Message : System datafile 1:  
'+DATA/visk4/datafile/system.256.649601515'needs media

..





# adrci Report in XML Format

**adrci> show report hm\_run HM\_RUN\_1**

```
<?xml version="1.0" encoding="US-ASCII"?>
<HM-REPORT REPORT_ID="HM_RUN_1">
  <TITLE>HM Report: HM_RUN_1</TITLE>
  <RUN_INFO>
    <CHECK_NAME>DB Structure Integrity
    Check</CHECK_NAME>
    <RUN_ID>1</RUN_ID>
    <RUN_NAME>HM_RUN_1</RUN_NAME>
    <RUN_MODE>REACTIVE</RUN_MODE>
    <RUN_STATUS>COMPLETED</RUN_STATUS>
    <RUN_ERROR_NUM>0</RUN_ERROR_NUM>
  <SOURCE_INCIDENT_ID>0</SOURCE_INCIDENT_ID> 05:00</SOURCE_INCIDENT_ID>
  ...
  ...
  ...
```

```
<NUM_INCIDENTS_CREATED>0</NUM_INCIDENTS_CREATED>
  <RUN_START_TIME>2008-03-17 12:54:17.240082 -
  05:00</RUN_START_TIME>
  <RUN_END_TIME>2008-03-17 12:54:43.528348 -
  05:00</RUN_END_TIME>
  </RUN_INFO>
  <RUN_PARAMETERS/>
  <RUN-FINDINGS>
    <FINDING>
      <FINDING_NAME>System datafile is old</FINDING_NAME>
      <FINDING_ID>2</FINDING_ID>
      <FINDING_TYPE>FAILURE</FINDING_TYPE>
      <FINDING_STATUS>OPEN</FINDING_STATUS>
      <FINDING_PRIORITY>CRITICAL</FINDING_PRIORITY>
      <FINDING_CHILD_COUNT>0</FINDING_CHILD_COUNT>
      <FINDING_CREATION_TIME>2008-03-17 12:54:42.631861 -
      05:00</FINDING_CREATION_TIME>
      <FINDING_MESSAGE>System datafile 1:
      '+DATA/visk4/datafile/system.256.649601515' needs media
      recovery</FINDING_MESSAGE>
      <FINDING_MESSAGE>Database cannot be
      opened</FINDING_MESSAGE>
    </FINDING>
```



# Real Application Testing – a.k.a RAT (Database Replay)

- Capture production workload including concurrency, transaction characteristics, and performance overhead
- Replay workload in another environment that is a copy of production
- Replay simulates production timing, concurrency, and think time
- Oracle Enterprise Manager or Command Line Option





# Real Application Testing - RAT

Mitigate as much risk as possible for testing:

1. Database upgrades or patches
2. Database initialization changes
3. Schema changes
4. Application upgrades
5. Operating System upgrades
6. Storage changes
7. Configuration changes such as changing to RAC or migrating to ASM
8. Interconnect changes



# RAT Presetup

Consider creating this directory as part of OFA setup

Create Oracle Directory:

```
CREATE OR REPLACE DIRECTORY RAT_DIR AS  
    '/apps/oracle/admin/DBATOOLS/replay';
```

Set Filters to capture only a subset of database workload or to exclude user sessions:

```
SQL> BEGIN  
  2  DBMS_WORKLOAD_CAPTURE.ADD_FILTER (fname => 'SCOTT_FILTER',  
  3                                     fattribute => 'USER',  
  4                                     fvalue => 'SCOTT');  
  5  END;  
  6  /
```

PL/SQL procedure successfully completed.





# RAT – Capture Workload

## Capture Workload

```
1 BEGIN
2 DBMS_WORKLOAD_CAPTURE.start_capture
3 (name => 'PROD_CAPTURE',
4  dir=>'RAT_DIR',
5  duration => NULL);
6* END;
SQL> /
PL/SQL procedure successfully completed.
```

- ☐ Captures all requests by external clients
- ☐ Captured in binary format

### Note:

1. Duration is in seconds
2. A null duration indicates that a capture will continue until stopped
3. Duration is optional (default of NULL)





# RAT – Finish Capture

Export AWR data if  
planning to run  
AWR Compare  
Period report

## Finish Capture

```
BEGIN
```

```
    DBMS_WORKLOAD_CAPTURE.finish_capture();
```

```
END;
```

Note:

- ☐ Database Capture is back ported to 10.2.0.4
- ☐ As of today, databases have to be upgraded to 10.2.0.4 to take advantage of Database Capture



# RAT – Process Capture

## Workload Processing

BEGIN

```
DBMS_WORKLOAD_REPLAY.process_capture('RAT_DIR');
```

END;

Note:

- ☐ The process\_capture has a single parameter for the directory name of the captured workload
- ☐ Processing converts data and metadata into replay files
- ☐ Processing database must be on the same version as the capture database
- ☐ Resource intensive process



# Database Replay - Initialize

```
BEGIN
  DBMS_WORKLOAD_REPLAY.initialize_replay
    (replay_name => 'PROD_REPLAY',
     replay_dir  => 'RAT_DIR');
END;
```

## Note:

Loads the metadata into tables to be used for replay



# Database Replay - Prepare

```
1  BEGIN
2      DBMS_WORKLOAD_REPLAY.prepare_replay(
3          synchronization => TRUE);
4*  END;
SQL> /
```

PL/SQL procedure successfully completed.

## Note:

- ☐ Synchronization of TRUE (required), the COMMIT order will be preserved during the replay



# Database Replay – Start Clients

```
wrc mode=calibrate replaydir=/apps/oracle/admin/DBATOOLS/replay
```

```
wrc system/<SYSTEM_PW> mode=replay replaydir=/apps/oracle/admin/DBATOOLS/replay
```

Workload Replay Client: Release 11.1.0.6.0 - Production on Wed Mar 02 00:31:41 2008

Copyright (c) 1982, 2007, Oracle. All rights reserved.

Report for Workload in: /apps/oracle/admin/DBATOOLS/replay

-----  
Recommendation:

Consider using at least 1 clients divided among 1 CPU(s).

Workload Characteristics:

- max concurrency: 1 sessions
- total number of sessions: 7

Assumptions:

- 1 client process per 50 concurrent sessions
- 4 client process per CPU
- think time scale = 100
- connect time scale = 100
- synchronization = TRUE





# Database Replay – Start Replay

```
BEGIN  
    DBMS_WORKLOAD_REPLAY.start_replay;  
END;  
/
```

Note:

To Cancel the replay:

```
BEGIN  
    DBMS_WORKLOAD_REPLAY.cancel_replay;  
END;
```



# Database Replay – Views

Views:

- ☐ DBA\_WORKLOAD\_CAPTURES
- ☐ DBA\_WORKLOAD\_REPLAYS
- ☐ DBA\_WORKLOAD\_CONNECTION\_MAP
- ☐ DBA\_WORKLOAD\_FILTERS
- ☐ DBA\_WORKLOAD\_REPLAY\_DIVERGENCE
- ☐ V\$WORKLOAD\_REPLAY\_THREAD



# Backup and Recovery

- RMAN recover ... block
- RMAN substitution variables
- RMAN compression
- Multi-section Backups
- Backup Shredding
- Data Repair Advisor





# RMAN – recover ... block

## New Syntax and Enhancements

- recover ... block
- Searches flashback logs first for good version of corrupt block

## New Syntax Example:

- recover datafile 13 block 10 datafile 20 block 11;

## Previous Blockrecover command:

- BLOCKRECOVER DATAFILE 13 BLOCK 10 DATAFILE 20 BLOCK 11 from backupset;



# Example of how we did RMAN substitution prior to 11g

```
export RMAN_BACKUP_LEVEL=$SH/rman_backup.sql
```

```
cat $RMAN_BACKUP_LEVEL |sed
```

```
-e "s/###_DATE_###/$ORADATE/g" \
```

```
-e "s/###_ORACLE_SID_###/$ORACLE_SID/g" \
```

```
-e "s/###_BACKUP_LEVEL_###/$BACKUP_LEVEL/g" \
```

```
-e "s/###_sqlspfile_###/$SPFILE_BACKUP_SYNTAX/g" \
```

```
> $RMAN_SCRIPT
```





# RMAN Substitution - Simple

```
CONNECT TARGET /  
BACKUP DATABASE TAG '&1';  
BACKUP ARCHIVELOG ALL TAG '&2';  
EXIT;
```

```
rman @/tmp/backup.sql USING DB_27MAR08  
ARCH_25MAR08
```



# RMAN Substitution Variables

```
cat back.ksh
export param1=$1
export param2=$2
rman target / cmdfile=/tmp/back2.sql using $param1 $param2
```

**VISK > ./back.ksh db\_27mar arch\_27mar**

Recovery Manager: Release 11.1.0.6.0 - Production on Wed Mar 26 23:38:25 2008  
Copyright (c) 1982, 2007, Oracle. All rights reserved.  
connected to target database: VISK (DBID=354790782)

```
RMAN> BACKUP DATABASE TAG 'db_27mar';
2> BACKUP ARCHIVELOG ALL TAG 'arch_27mar';
3> EXIT;
```

Starting backup at 26-MAR-08

using target database control file instead of recovery catalog

allocated hannel: ORA\_DISK\_1

```
cat back2.sql
BACKUP DATABASE TAG '&1';
BACKUP ARCHIVELOG ALL TAG '&2';
EXIT;
```



# Execute RMAN script

Within RMAN:

```
rman target /
```

```
RMAN> @back2.sql db_27mar arch_27mar
```

```
RMAN> BACKUP DATABASE TAG 'db_27mar';
```

Starting backup at 26-MAR-08

using target database control file instead of recovery catalog

allocated channel: ORA\_DISK\_1

channel ORA\_DISK\_1: SID=105 device type=DISK



# RMAN Script with Variable Substitution

```
create script backup_db
{
  backup tag &1
  format '/apps/oracle/admin/&2/bkups /%d.%s.%p.%t.DB' (database)
  plus archivelog;
}
```

```
RMAN> run
{
  execute script backup_db
  using db_mar27 VISK;
}
```



# RMAN Compression

- RMAN> show all;

RMAN configuration parameters for database with db\_unique\_name VISK4 are:

CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default

CONFIGURE BACKUP OPTIMIZATION OFF; # default

CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default

CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default

CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default

CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default

CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default

CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default

CONFIGURE MAXSETSIZE TO UNLIMITED; # default

CONFIGURE ENCRYPTION FOR DATABASE OFF; # default

CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default

**CONFIGURE COMPRESSION ALGORITHM 'BZIP2'; # default**

**CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default**

CONFIGURE SNAPSHOT CONTROLFILE NAME TO

'/apps/oracle/product/11.1.0/DB/dbs/snapcf\_VISK4.f'; # default





# RMAN Compression

```
1 select algorithm_name, algorithm_description,  
2         database_compatibility Compatibility, is_default  
3* from v$rman_compression_algorithm;
```

ALGORITHM_NAME	ALGORITHM_DESCRIPTION	COMPATIBILITY	IS_
-----	-----	-----	--
ZLIB	optimized for speed	11.1.0.0.0	NO
BZIP2	optimized for maximum compression	11.1.0.0.0	YES

## ZLIB

1. Part of Oracle Advanced Compression Option
2. Optimized for CPU efficiency
3. Larger ZIP files than BZIP2 – sacrifice compression for better performance
4. COMPATIBILITY must be set to 11.1.0 to use ZLIB



# RMAN Compression

Set compression back to the new 11g compression:

```
RMAN> configure compression algorithm 'zlib';
```

using target database control file instead of recovery catalog

new RMAN configuration parameters:

```
CONFIGURE COMPRESSION ALGORITHM 'zlib';
```

new RMAN configuration parameters are successfully stored



# RMAN Multisection Backups

**RMAN> backup section size 500m tablespace trans\_d;**

```
Starting backup at 18-MAR-08
using target database control file instead of recovery catalog
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=124 device type=DISK
channel ORA_DISK_1: starting full datafile backup set
channel ORA_DISK_1: specifying datafile(s) in backup set
input datafile file number=00006
name=+DATA/visk4/datafile/trans_d.275.649684237
backing up blocks 1 through 64000
```

...

```
backing up blocks 64001 through 128000
backing up blocks 128001 through 192000
backing up blocks 256001 through 262144
```



# RMAN Multisection Backups

**RMAN> backup as compressed backupset section size 500m  
tablespace trans\_d;**

```
Starting backup at 18-MAR-08
using channel ORA_DISK_1
channel ORA_DISK_1: starting compressed full datafile backup
set
channel ORA_DISK_1: specifying datafile(s) in backup set
input datafile file number=00006
name=+DATA/visk4/datafile/trans_d.275.649684237
backing up blocks 1 through 64000
channel ORA_DISK_1: starting piece 1 at 18-MAR-08
channel ORA_DISK_1: finished piece 1 at 18-MAR-08
```



# RMAN Multisection Backups

```
SQL> select file#, section_size  
       from v$backup_datafile;
```

FILE#	SECTION_SIZE
1	0
2	0
3	0
5	0
4	0
0	0
0	0
6	64000
6	64000

9 rows selected.

```
SQL> select pieces,multi_section  
       from v$backup_set;
```

PIECES	MUL
1	NO
1	NO
1	NO
1	NO
5	YES
5	YES

6 rows selected.





# RMAN Backup Shredding

Delete the encryption key of a backup to render the encrypted backup inaccessible

To enable it:

```
RMAN> configure encryption external key  
storage on;
```

To use backup shredding:

```
RMAN> delete force backup;
```



# RMAN Network-Based Database Duplication

## Also called Active Database Duplication

- Both source and target must be of same OS
- Oracle Net has to be setup for both target and duplicate database (even if the source and target is on the same host)
- Must have the same sysdba password enforced through the password file
- Source database can be open or in mounted state
- If the source database is open, it must be in archivelog mode
- If the source database is in mounted mode, it must have been shutdown clean before it was mounted
- By default, the password file will not be copied so you must specify the password file clause inside the duplicate command



# Merge Recovery Catalogs

Connect to the destination recovery catalog

```
$ rman
```

```
RMAN> connect catalog rman/rman@rmanprod
```

```
RMAN> import catalog rman1/rman1@rmantest;
```

```
Starting import catalog at 08-APR-07
```

```
connected to source recovery catalog database
```

```
import validation complete
```

```
database unregistered from the source recovery  
catalog
```

```
Finished import catalog at 08-APR-07
```

```
RMAN>
```



# Merge Recovery Catalogs

Automatically unregisters the databases in the source catalog

- `RMAN> import catalog rman1/rman1@rmantest no unregister;`

Merge specific DBID or DB\_NAME

- `RMAN> import catalog rman1/rman1@rmantest dbid=354790782;`
- `RMAN> import catalog rman1/rman1@rmantest db_name=VISK1;`



# Virtual Private Catalogs - Grant

```
RMAN> grant register database to vrman1;  
Grant succeeded.
```

```
RMAN> grant catalog for database visk1 to vrman2;  
Grant succeeded.
```

```
RMAN> grant catalog for database visk to vrman2;  
Grant succeeded.
```

## Note:

- “Grant Register Database” implicitly grants “grant catalog for database” databases registered by the user





# Virtual Private Catalogs

```
RMANP > rman catalog vrman2/vrman123@rmanp
```

```
Recovery Manager: Release 11.1.0.6.0 - Production on Sat  
May 3 10:27:06 2008
```

```
Copyright (c) 1982, 2007, Oracle. All rights reserved.  
connected to recovery catalog database
```

```
RMAN> create virtual catalog;
```

```
found eligible base catalog owned by RMAN
```

```
created virtual catalog against base catalog owned by RMAN
```



# Manage Virtual Catalogs

```
RMANP > rman catalog rman/rman123@rmanp target /
```

```
RMAN> revoke register database from vrman1;  
Revoke succeeded.
```

```
RMAN> revoke catalog for database visk from vrman2;  
Revoke succeeded.
```

```
RMAN> revoke catalog for database visk1 from vrman2;  
Revoke succeeded.
```

```
RMAN> revoke all privileges from vrman1;  
Revoke succeeded.
```



# DRA - List Failure

RMAN> list failure;

```
using target database control file instead of  
List of Database Failures
```

=====

Failure ID	Priority	Status	Time Detected	Summary
-----	-----	-----	-----	-----
182	HIGH	OPEN	20-MAR-08	One or more non-system datafiles are missing





# DRA – List Failure Detail

RMAN> list failure 182 detail;

## List of Database Failures

=====

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

182	HIGH	OPEN	20-MAR-08	One or more non-system datafiles are missing
-----	------	------	-----------	--

Impact: See impact for individual child failures

List of child failures for parent failure ID 182

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

210	HIGH	OPEN	20-MAR-08	Datafile 8: '/u30/oradata/VISK/transd_01.dbf' is missing
-----	------	------	-----------	---

Impact: Some objects in tablespace TRANS\_D might be unavailable



# DRA - Getting Failure Advise

## RMAN> advise failure;

List of Database Failures

=====

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

182	HIGH	OPEN	20-MAR-08	One or more non-system datafiles are missing
-----	------	------	-----------	--

analyzing automatic repair options; this may take some time

..

analyzing automatic repair options complete

Mandatory Manual Actions

=====

no manual actions available

Optional Manual Actions

=====

1. If file /u30/oradata/VISK/transd\_01.dbf was unintentionally renamed or moved, restore it

Automated Repair Options

=====

Option Repair Description

-----

1 Restore and recover datafile 8

Strategy: The repair includes complete media recovery with no data loss

Repair script: /apps/oracle/diag/rdbms/visk/VISK/hm/reco\_3536057071.hm





# DRA – Repairing the Failure

**RMAN> repair failure;**

Strategy: The repair includes complete media recovery with no data loss

Repair script: /apps/oracle/diag/rdbms/visk/VISK/hm/reco\_3536057071.hm

contents of repair script:

```
# restore and recover datafile
sql 'alter database datafile 8 offline';
restore datafile 8;
recover datafile 8;
sql 'alter database datafile 8 online';
```

Do you really want to execute the above repair (enter YES or NO)? yes  
executing repair script

sql statement: alter database datafile 8 offline

Starting restore at 20-MAR-08  
using channel ORA\_DISK\_1

creating datafile file number=8 name=/u30/oradata/VISK/transd\_01.dbf  
restore not done; all files read only, offline, or already restored  
Finished restore at 20-MAR-08

..  
..

sql statement: alter database datafile 8 online  
repair failure complete



# DRA – List Failure

**The failure is no longer listed after the repair action:**

```
RMAN> list failure;
```

no failures found that match specification

## List Closed Failures:

```
RMAN> list failure closed;
```

List of Database Failures

=====

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

105	CRITICAL	CLOSED	08-MAR-08	System datafile 1: '+DATA/visk/datafile/system.257.648832243' needs media recovery
102	CRITICAL	CLOSED	08-MAR-08	Control file needs media recovery
185	HIGH	CLOSED	20-MAR-08	Datafile 8: '/u30/oradata/VISK/transd_01.dbf' is missing Impact: Some objects in tablespace TRANS_D might be unavailable
123	HIGH	CLOSED	08-MAR-08	Datafile 4: '+DATA/visk/datafile/users.260.648832243' needs media recovery Impact: Some objects in tablespace USERS might be unavailable
117	HIGH	CLOSED	08-MAR-08	Datafile 3: '+DATA/visk/datafile/undotbs1.259.648832243' needs media recovery Impact: Some objects in tablespace UNDOTBS1 might be unavailable
111	HIGH	CLOSED	08-MAR-08	Datafile 2: '+DATA/visk/datafile/sysaux.258.648832243' needs media recovery Impact: Some objects in tablespace SYSAUX might be unavailable
108	HIGH	CLOSED	08-MAR-08	One or more non-system datafiles need media recovery



# List and Change High Priority Failures

RMAN> list failure high;

List of Database Failures

=====

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

182	HIGH	OPEN	20-MAR-08	One or more non-system datafiles are missing
-----	------	------	-----------	--

## Change failure priority from High to Low

RMAN> change failure 182 priority low;

List of Database Failures

=====

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

182	HIGH	OPEN	20-MAR-08	One or more non-system datafiles are missing
-----	------	------	-----------	--

Do you really want to change the above failures (enter YES or NO)? yes  
changed 1 failures to LOW priority



# List Low Priority Failures

```
RMAN> list failure low;
```

List of Database Failures

=====

Failure ID	Priority	Status	Time Detected	Summary
------------	----------	--------	---------------	---------

-----

182	LOW	OPEN	20-MAR-08	One or more non-system datafiles are missing
-----	-----	------	-----------	--

## *Note:*

Once the priority is changed to a low priority, the failure will no longer show up on the list failure command.



# Dynamic View to list failure advise

```
1 select advise_id, rank, message
2* from v$ir_manual_checklist
SQL> /
```

ADVISE_ID	RANK	MESSAGE
201	0	If file /u30/oradata/VISK/transd_01.dbf was unintentionally renamed or moved, restore it





# Advise Before Repair

RMAN> repair failure preview;

using target database control file instead of recovery catalog

RMAN-00571: =====

RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====

RMAN-00571: =====

RMAN-03002: failure of repair command at 03/20/2008 07:47:52

RMAN-06954: REPAIR command must be preceded by ADVISE command in same session



# Repair Failure Preview

**Execute the repair in PREVIEW mode to see what DATA Recovery Advisor will do:**

RMAN> repair failure preview;

Strategy: The repair includes complete media recovery with no data loss

Repair script:

/apps/oracle/diag/rdbms/visk/VISK/hm/reco\_395520741.hm

contents of repair script:

```
# restore and recover datafile  
sql 'alter database datafile 8 offline';  
restore datafile 8;  
recover datafile 8;  
sql 'alter database datafile 8 online';
```



# V\$IR\_REPAIR

**The results of the repair failure command are captured in the V\$IR\_REPAIR view:**

```
1  select repair_id, advise_id, summary, rank
2*  from v$ir_repair
```

```
SQL> /
```

REPAIR_ID	ADVISE_ID	SUMMARY	RANK
202	201	NO DATA LOSS	1
267	266	NO DATA LOSS	1



# Validate Database – dbverify

**RMAN> validate database;**

Starting validate at 21-MAR-08

using channel ORA\_DISK\_1

channel ORA\_DISK\_1: starting validation of datafile

channel ORA\_DISK\_1: specifying datafile(s) for validation

input datafile file number=00008 name=/u30/oradata/VISK/transd\_01.dbf

..

input datafile file number=00007 name=+DATA/visk/datafile/visk\_aes256.268.649637901

channel ORA\_DISK\_1: validation complete, elapsed time: 00:00:15

List of Datafiles

=====

File	Status	Marked Corrupt	Empty Blocks	Blocks Examined	High SCN
------	--------	----------------	--------------	-----------------	----------

----	-----	-----	-----	-----	-----
------	-------	-------	-------	-------	-------

3	OK	0	15	3840	1379409
---	----	---	----	------	---------

File Name: +DATA/visk/datafile/undotbs1.259.648863163

Block Type	Blocks Failing	Blocks Processed
------------	----------------	------------------

-----	-----	-----
-------	-------	-------

Data	0	0
------	---	---

Index	0	0
-------	---	---

Other	0	3825
-------	---	------

..

Control File OK	0	594
-----------------	---	-----

Finished validate at 21-MAR-08





# Validate at a Granular Level

RMAN> validate tablespace VISK\_AES256;

RMAN> validate datafile 7;

RMAN> validate datafile 7 block 100;

## Validate options:

1. Recovery area
2. Recovery files
3. Spfile
4. Tablespace
5. Controlfilecopy
6. Backupset





# Archivelog Retention Policy

## 10g

- delete noprompt archivelog until time 'sysdate -2' backed up 2 times to device type disk;
- backup ... delete input;

## 11g

- Configure archivelog deletion policy to backuped 2 times to disk;
- Configure archivelog delete policy to none;



# sqlplus

## Errorlogging:

- set errorlogging on
- show errorlogging
- select username, statement, message from sperrorlog;

sqlplus/admin/glogin.sql

- Is blank now

Previously: sqlplus / as sysdba And / as sysoper

- sqlplus / as sysasm

## Enhanced BLOB support

- Can query and print tables with BLOB and BFILE datatype

SQL>	desc	sperrorlog	Null?	Type
Name				
-----				
USERNAME				VARCHAR2 (256)
TIMESTAMP				TIMESTAMP (6)
SCRIPT				VARCHAR2 (1024)
IDENTIFIER				VARCHAR2 (256)
MESSAGE				
STATEMENT				
				CLOB
				CLOB



**PL/SQL**

**PL/SQL**

**PL/SQL**

**PL/SQL**

**PL/SQL**



# Dynamic SQL

- Native dynamic SQL supports CLOB > 32K characters
- DBMS\_SQL.PARSE() gains a CLOB overload
- REF CURSOR can be converted to a DBMS\_SQL cursor
- DBMS\_SQL cursor can be converted to a REF CURSOR
- DBMS\_SQL supports collections and object types
- DBMS\_SQL allows bulk binds using user-defined collection types



# Finer Grained Dependency Tracking

Significantly less invalidations

- Create or replace synonym
- Create or replace view
- Add or drop or modify a column to a table
- Add to a specification





# PL/SQL Continue

Can skip to the next iteration in a loop

```
BEGIN
```

```
FOR I in 1..10 LOOP
```

```
  dbms_output.put_line('I=' || to_char(I));
```

```
  IF ( I in (5,7) ) THEN
```

```
    continue;
```

```
  END IF;
```

```
  dbms_output.put_line('Number  
is:' || to_char(I));
```

```
END LOOP;
```

```
END;
```

```
/
```

```
SQL> /
```

```
I=1
```

```
Number is:1
```

```
I=2
```

```
Number is:2
```

```
I=3
```

```
Number is:3
```

```
I=4
```

```
Number is:4
```

```
I=5
```

```
I=6
```

```
Number is:6
```

```
I=7
```

```
I=8
```

```
Number is:8
```

```
I=9
```

```
Number is:9
```

```
I=10
```

```
Number is:10
```



# PL/SQL Continue

Written in another way:

```
BEGIN
  FOR I in 1..10 LOOP
    dbms_output.put_line('I='||to_char(I));
    continue when ( I in (5,7) );
    dbms_output.put_line('Number is:'||to_char(I));
  END LOOP;
END;
/
```



# PL/SQL Sequences Without Dual

## Old Convention

```
SQL> create sequence trans_seq;
Sequence created.

SQL> declare
2  cursor c1 is
3  select trans_seq.nextval
4  from dual;
5  v_seq NUMBER;
6  begin
7  open c1; fetch c1 into v_seq;
  close c1;
8  dbms_output.put_line('Sequence:
  ||to_char(v_seq));
9  end;
10 /
Sequence: 1
```

## New Convention

```
1  declare
2  v_seq NUMBER;
3  begin
4  v_seq := trans_seq.nextval;
5  dbms_output.put_line('Sequence:
  ||to_char(v_seq));
6* end;
SQL> /
Sequence: 2

PL/SQL procedure successfully
  completed.
```



# SIMPLE\_INTEGER

- Range of  $-2,147,483,648$  to  $2,147,483,647$
- Variables of this datatype wrap from smallest to largest and from largest to smallest
- Cannot be NULL
- SIMPLE\_INTEGER elapsed time(seconds): 1.60
- PLS\_INTEGER elapsed time(seconds): 1.75
- NUMBER elapsed time(seconds): 3.84

```
alter package si_demo compile plsql_code_type=native;
```

- SIMPLE\_INTEGER elapsed time(seconds): 0.10
- PLS\_INTEGER elapsed time(seconds): 0.80
- NUMBER elapsed time(seconds): 3.22



# PL/SQL Native Compilation

No more C-Compilers

Faster and easier to implement

```
SQL> alter session set plsql_code_type=native;  
Session altered.
```

```
SQL> alter procedure edba_write_alert compile;  
Procedure altered.
```

**How to check if it is really natively compiled:**

```
1 select name, type, plsql_code_type, plsql_optimize_level  
2 from dba_plsql_object_settings  
3* where plsql_code_type = 'NATIVE';
```

NAME	TYPE	PLSQL_CODE_T	PLSQL_OPTIMIZE_LEVEL
EDBA_WRITE_ALERT	PROCEDURE	NATIVE	

*Compile at database level:*  
*plsql\_code\_type=native*  
*plsql\_optimize\_level=3*  
*startup upgrade;*  
*@dbsupgnv.sql -- dbmsupgin.sql to go back*  
*@utlrp.sql*





# Regular Expression Enhancements

- The REGEXP\_INSTR and REGEXP\_SUBSTR functions include a new SUBEXPR parameter to select the *nth subexpression* in the regular expression being evaluated.

```
SELECT REGEXP_INSTR('All Good Dogs Go Go To Heaven...' -- source
,'(go) (go)' -- regular expression with two subexpressions
,1 -- starting position to begin evaluation
,1 -- match occurrence
,0 -- 0 = return position of start, 1 = position after end
,'i' -- case insensitive search
,2 -- which subexpression to return position
) Dup_position from dual;
```

DUP\_POSITION

-----



# REGEXP\_COUNT

- Count the number of matches a regular expression
- In this example: Count the numbers in the source text

```
SELECT REGEXP_COUNT('We have plans to hire 2 of 9 DBAs', --  
    source text  
'[0-9]' -- regular expression  
,1 -- starting position  
, 'i' -- ignore case  
) NUMBER_COUNT  
FROM dual  
/
```

NUMBER\_COUNT

-----



# REGEXP\_COUNT

- In this example: Count the CAPITAL Letters in the source text
- Ignore case option is NOT used

```
SELECT REGEXP_COUNT('We have plans to hire 2 of 9 DBAs', --  
    source text  
    '[A-Z]' -- regular expression  
    ,1 -- starting position  
    ) CAPS_COUNT  
FROM dual  
/
```

```
CAPS_COUNT
```

```
-----
```

```
4
```



# SQL Arguments for Function Calls

Name => value is now supported

```
SELECT
    VENDORS_PKG.vendor_Id (vendor_name, organization_id)
FROM VENDORS;
```

```
SELECT
    VENDORS_PKG.vendor_Id (vendor_name, pOrgId=>organiza
        tion_id)
FROM VENDORS;
```

```
SELECT
    VENDORS_PKG.vendor_Id (pVendorName=>vendor_name, pOr
        gId=>organization_id)
FROM VENDORS;
```



# RESULT\_CACHE

## Initialization Parameters

SQL> show parameter result\_cache

NAME	TYPE	VALUE
-----	-----	-----
client_result_cache_lag	big integer	3000
client_result_cache_size	big integer	0
result_cache_max_result	integer	5
result_cache_max_size	big integer	<b>1536K</b>
result_cache_mode	string	MANUAL
result_cache_remote_expiration	integer	<b>0</b>

Note:

- Exercise caution setting the result\_cache\_remote\_expiration parameter
- Result\_cache\_remote\_expiration is in minutes





# Result Cache Example

SQL> alter system set result\_cache\_max\_size=200m;  
System altered.

```
1  set autot traceonly
2  set timing on
3  select /*+ result_cache */
      empno,
      ename,
      job
    from emp2
4* order by empno;
```

You can add the result\_cache hint to:

1. SQL Queries
2. Subqueries
3. Inline Views



# Result Cache Example

## Execution Plan

-----  
Plan hash value: 2441141433  
-----

Id	Operation	Name	Rows	Bytes	TempSpc	Cost (%CPU)	Time
-----							
--							
0	SELECT STATEMENT		48554	1232K		739 (1)	00:00:09
1	RESULT CACHE	4skzh0m6swvmv4xxgm3cdh643g					
2	SORT ORDER BY		48554	1232K	3448K	739 (1)	00:00:09
3	TABLE ACCESS FULL	EMP2	48554	1232K		377 (1)	00:00:05
-----							

Result Cache Information (identified by operation id):  
-----

```
1 - column-count=3; dependencies=(SCOTT.EMP2); name="select /*+ result_cache */ empno, ename, job
  from emp2 order by empno"
```



# Result Cache – Execution #1

## Statistics

---

4	recursive calls
0	db block gets
<b>1435</b>	<b>consistent gets</b>
<b>1379</b>	<b>physical reads</b>
0	redo size
2566705	bytes sent via SQL*Net to client
66409	bytes received via SQL*Net from client
6001	SQL*Net roundtrips to/from client
<b>1</b>	<b>sorts (memory)</b>
0	sorts (disk)
90000	rows processed



# Result Cache – Execution #2

## Statistics

---

```
0 recursive calls
0 db block gets
0 consistent gets
0 physical reads
0 redo size
2566705 bytes sent via SQL*Net to client
66409 bytes received via SQL*Net from client
6001 SQL*Net roundtrips to/from client
0 sorts (memory)
0 sorts (disk)
90000 rows processed
```



# PL/SQL Function Result Cache

## Notes:

- Results from function calls can be cached in the SGA
- Future invocation of the function call with the same arguments will return values from the cache
- Keyword `RESULT_CACHE` is used as part of the function/package specification
- Keyword `RELIES_ON()` can be utilized to cause cache entries to become invalid if table or view is updated
- Cached values are available to all sessions that have execute privileges on the function





# PL/SQL Function Result Cache – Example #1

```
1  create or replace function count_job
2    (p_job in varchar2)
3    return number
4    result_cache
5    as
6    cursor c1 is
7    select count(*)
8    from emp2
9    where job=p_job;
10  v_count simple_integer := 0;
11  begin
12    open c1; fetch c1 into v_count; close c1;
13    -- Pause for 1 second to slow the function down.
14    dbms_lock.sleep(1);
15    return v_count;
16* end;
SQL> /
```

Function created.



# PL/SQL Function Result Cache – Example #1 Result

```
SQL> select count_job('SALESMAN') from dual;
```

```
COUNT_JOB('SALESMAN')
```

```
-----
```

```
205698
```

```
Elapsed: 00:00:01.50
```

```
SQL> insert into emp2 (empno, ename, job)
      values (9999, 'SANCHEZ', 'SALESMAN');
```

```
1 row created.
```

```
Elapsed: 00:00:00.10
```

```
SQL> commit;
```

## AFTER THE INSERT

```
SQL> select count_job('SALESMAN')
      from dual;
```

```
COUNT_JOB('SALESMAN')
```

```
-----
```

```
205698
```

```
Elapsed: 00:00:00.00
```



# PL/SQL Function Result Cache – Example #2

```
1  create or replace function count_job
2  (p_job in varchar2)
3  return number
4  result_cache
5  relies_on (emp2)
6  as
7  cursor c1 is
8  select count(*)
9  from emp2
10 where job=p_job;
11 v_count simple_integer := 0;
12 begin
13 open c1; fetch c1 into v_count; close c1;
14 return v_count;
15* end;
SQL> /
Function created.
```



# PL/SQL Function Result Cache – Example #2 Result

```
SQL> select count_job('SALESMAN') from dual;
```

```
COUNT_JOB('SALESMAN')
```

```
-----
```

```
205698
```

```
Elapsed: 00:00:01.09
```

```
SQL> /
```

```
COUNT_JOB('SALESMAN')
```

```
-----
```

```
205698
```

```
Elapsed: 00:00:00.00
```

```
SQL> insert into emp2 (empno, ename, job)
      values (9999, 'SANCHEZ', 'SALESMAN');
```

```
1 row created.
```

```
Elapsed: 00:00:00.00
```

```
SQL> commit;
```

## AFTER THE INSERT

```
SQL> select count_job('SALESMAN')
      from dual;
```

```
COUNT_JOB('SALESMAN')
```

```
-----
```

```
205699
```

```
Elapsed: 00:00:01.09
```



# Result Cache Memory Report

```
SQL> set serveroutput on
```

```
SQL> exec dbms_result_cache.memory_report
```

```
Result Cache Memory Report
```

```
[Parameters]
```

```
Block Size = 1024 bytes
```

```
Maximum Cache Size = 393216 bytes (384 blocks)
```

```
Maximum Result Size = 19456 bytes (19 blocks)
```

```
[Memory]
```

```
Total Memory = 13412 bytes [0.017% of the Shared Pool]
```

```
... Fixed Memory = 10560 bytes [0.013% of the Shared Pool]
```

```
... State Object Pool = 2852 bytes [0.004% of the Shared Pool]
```

```
... Cache Memory = 0 bytes (0 blocks) [0.000% of the Shared Pool]
```

```
PL/SQL procedure successfully completed.
```





# Disable Result Cache

```
SQL> exec dbms_result_cache.memory_report;
```

ResultCacheMemoryReport

Cache is disabled.

PL/SQL procedure successfully completed.

Toggle Result Cache parameters at the system level:

```
SQL> alter system set result_cache_mode = manual;
```

Force Result Cache at the Database Level:

```
SQL> alter system set result_cache_mode = force;
```

Last Option - **Auto**:

Based on frequency of execution, cost of execution and how much underlying objects are changing



# Result Cache

## Out of the box settings:

- .25% of memory\_target
- .5% of sga\_target
- 1% of shared\_pool\_size
- Max allocation cannot be > 75% of shared\_pool

## Override Hints:

`/*+ result_cache +/`

`/*+ no_result_cache +/`

*Much more to come on Result  
Cache !!!!!*

*Stay tuned to:  
Viscosity.com*



# Triggers – Ordering and Disabled

Use the follows clause for Trigger Ordering

```
CREATE TRIGGER B4I_EMP2_AUD  
BEFORE INSERT ON EMP2  
FOLLOWS B4I_EMP2_SAL,B4I_EMP2_SEC  
WHEN ...
```

Use the ENABLED or DISABLED clause to create triggers

```
CREATE TRIGGER B4I_EMP2_AUD  
BEFORE INSERT ON EMP2  
FOLLOWS B4I_EMP2_SAL,B4I_EMP2_SEC  
DISABLED  
WHEN ...
```



# Compound Triggers

```
CREATE OR REPLACE TRIGGER
    emp2_compound_trigger
FOR INSERT ON emp2
COMPOUND TRIGGER

v_num NUMBER;
v_char VARCHAR2(30);

-- Before Statement Level
BEFORE STATEMENT IS
BEGIN
    dbms_output.put_line('BEFORE
        STATEMENT');
END BEFORE STATEMENT;
```

```
-- Before Row Level
BEFORE EACH ROW IS
BEGIN
    dbms_output.put_line('BEFORE ROW');
END BEFORE EACH ROW;
-- After Statement Level
AFTER STATEMENT IS
BEGIN
    dbms_output.put_line('AFTER
        STATEMENT');
END AFTER STATEMENT;
-- After Row Level
AFTER EACH ROW IS
BEGIN
    dbms_output.put_line('AFTER ROW');
END AFTER EACH ROW;
END emp2_compound_trigger;
/
```

Trigger created.



# Datapump Compression

```
VISK > expdp full=yes userid='"/ as sysdba' "  
> dumpfile=data_pump_dir:full.dmp.compress compression=all
```

Export: Release 11.1.0.6.0 - Production on Monday, 24 March, 2008 6:43:01

Copyright (c) 2003, 2007, Oracle. All rights reserved.

```
Connected to: Oracle Database 11g Enterprise Edition Release 11.1.0.6.0 - Production  
With the Partitioning, OLAP, Data Mining and Real Application Testing options  
Starting "SYS"."SYS_EXPORT_FULL_02": full=yes userid='/'***** AS SYSDBA'  
dumpfile=data_pump_dir:full.dmp.compress compression=all  
Estimate in progress using BLOCKS method...  
Processing object type DATABASE_EXPORT/SCHEMA/TABLE/TABLE_DATA
```





# Datapump Compression

```
expdp userid=scott/tiger dumpfile=emp2.dmp tables=emp2  
directory=data_pump_dir
```

```
-rw-r----- 1 oracle oinstall 16158720 Mar 24 07:06 emp2.dmp
```

```
-rw-r----- 1 oracle oinstall 1009386 Mar 24 07:06 emp2.dmp.gz
```

```
expdp userid=scott/tiger dumpfile=emp2.compress.dmp tables=emp2 \  
directory=data_pump_dir compression=all
```

```
-rw-r----- 1 oracle oinstall 1437696 Mar 24 07:08 emp2.compress.dmp
```



# Datapump Encryption

- encryption = {all | data\_only | encrypted\_columns\_only | metadata\_only | none}

**Note:**

Encrypted\_columns\_only → only the encrypted columns are written to the dump file

- Oracle supports three algorithms:
  1. AES128 - Default
  2. AES192
  3. AES256

Three ENCRYPTION\_MODE options:

1. dual
2. password
3. transparent



# Datapump Encryption

- expdp userid=scott/tiger dumpfile=tde.dmp tables=emp2  
directory=data\_pump\_dir parallel=2 encryption\_mode=transparent  
encryption=all
- expdp userid=scott/tiger dumpfile=emp2.enc.dmp tables=emp2  
directory=data\_pump\_dir parallel=2 encryption\_mode=password  
encryption\_password=oracle123
- expdp userid=scott/tiger dumpfile=emp2.aes256.dmp tables=emp2  
directory=data\_pump\_dir parallel=2 encryption\_mode=password  
encryption\_password=oracle123 encryption\_algorithm=aes256
- expdp userid=scott/tiger dumpfile=emp2.dual.dmp tables=emp2  
directory=data\_pump\_dir parallel=2 encryption\_mode=dual encryption=all  
encryption\_password=oracle123 encryption\_algorithm=aes256

NOTE: The password parameter is required for dual or password mode



# Trying to import a password protected dumpfile without the password

```
VISK > impdp userid=scott/tiger dumpfile=emp2.enc.dmp tables=emp2  
directory=data_pump_dir
```

Import: Release 11.1.0.6.0 - Production on Monday, 24 March, 2008 7:23:49

Copyright (c) 2003, 2007, Oracle. All rights reserved.

Connected to: Oracle Database 11g Enterprise Edition Release 11.1.0.6.0 -  
Production

With the Partitioning, OLAP, Data Mining and Real Application Testing options

**ORA-39002: invalid operation**

**ORA-39174: Encryption password must be supplied.**



# Successful Import With the Password

```
SQL> drop table emp2;  
Table dropped.
```

```
VISK > impdp userid=scott/tiger dumpfile=emp2.enc.dmp tables=emp2  
directory=data_pump_dir encryption_password=oracle123
```

```
Starting "SCOTT"."SYS_IMPORT_TABLE_01": userid=scott/*****  
dumpfile=emp2.enc.dmp tables=emp2 directory=data_pump_dir  
encryption_password=*****
```

```
Processing object type TABLE_EXPORT/TABLE/TABLE
```

```
Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
```

```
.. imported "SCOTT"."EMP2" 15.33 MB 269990 rows
```

```
Processing object type TABLE_EXPORT/TABLE/STATISTICS/USER_PREF_STATISTICS
```

```
Job "SCOTT"."SYS_IMPORT_TABLE_01" successfully completed at 07:25:18
```





# Reuse a Dump File

ORA-39001: invalid argument value

ORA-39000: bad dump file specification

ORA-31641: unable to create dump file

`"/apps/oracle/admin/VISK/dpdump/emp2.dual.dmp"`

ORA-27038: created file already exists

Additional information:

What can you do now? Use the REUSE\_DUMPFILES parameter

```
expdp userid=scott/tiger dumpfile=emp2.dual.dmp tables=emp2  
directory=data_pump_dir parallel=2 encryption_mode=dual  
encryption=all encryption_password=oracle123  
encryption_algorithm=aes256 reuse_dumpfiles=y
```



# Datapump to Remap Data

## First create a package

```
create or replace package emp_sal
is
function remap_sal (p_sal number) return
    number;
end;
/
create or replace package body emp_sal
is
function remap_sal
(p_sal number) return number
as
v_sal number := 75000;
begin
return v_sal;
end;
end;
/
```

## Export with the remap\_data option:

```
expdp userid=scott/tiger dumpfile=emp2.dual.dmp
tables=emp2 directory=data_pump_dir parallel=2
encryption_mode=dual encryption=all
encryption_password=oracle123
encryption_algorithm=aes256
remap_data=scott.emp2.sal:emp_sal.remap_sal
reuse_dumpfiles=y
```

## After the import:

```
SQL> select ename, sal from emp2
where rownum<10;
```

ENAME	SAL
-----	-----
SMITH	75000
ALLEN	75000
WARD	75000
JONES	75000
[...]	



# Datapump External Tables

```
create directory dbadir as '/tmp/dba';
create directory tmpdir as '/tmp';

create table emp2_external
organization external
(
  type oracle_datapump
  default directory dbadir
  access parameters
  (logfile tmpdir:emp2 compression enabled encryption enabled)
  location ('emp2_external.dmp')
)
as select * from emp2;
```

Table created.



# Application Express

V3.0 is shipped with Oracle Database 11g

V3.1 is available on OTN as of Feb 2008

Apex configuration is simple:

```
SQL> @?/apex/apxconf.sql
```

```
PORT
```

```
-----
```

```
8080
```



# Application Express Setup Output

Enter values below for the XDB HTTP listener port and the password for the Application Default values are in brackets [ ].

Press Enter to accept the default value.

Enter a password for the ADMIN user [ ] oracle123

Enter a port for the XDB HTTP listener [ 8080] 8080

...changing HTTP Port

PL/SQL procedure successfully completed.

PL/SQL procedure successfully completed.

Session altered.

...changing password for ADMIN

PL/SQL procedure successfully completed.

Now you can unlock the anonymous user:

```
SQL> alter user anonymous account unlock;
```

```
User altered.
```





# Application Express

## Access APEX:

- [http://rac11.dbaexpert.com:8080/apex/apex\\_admin](http://rac11.dbaexpert.com:8080/apex/apex_admin)
- Application Express does not need HTTP Server
- Runs with the Embedded PL/SQL Gateway (EPG)
- If upgrading the database from 10g to 11g, you must manually install and configure APEX



# APEX Enhancements

- Provides the ability to print a PDF for report regions.
- Renders up to 18 different flash charts and converts scalable vector graphics to flash.
- Migrate Access applications to Application Express using the Application Migration Workshop.
- Provides drag-and-drop layout pages for development productivity.
- Provides new item types such as date picker, pop-up list of values, shuttle, and new HTML text editor areas.
- Renders calendars by providing built-in wizards for daily, weekly, and monthly views.



# APEX Enhancements #2

- It compares database objects in two separate schemas.
- It compares differences between two different applications.
- It facilitates bookmark pages within the application.
- It performs page and region caching for application performance.
- Find items for CSS and images.
- Defines rules for password expiration, forces strong passwords, locks accounts, and forces password changes upon first login.
- Manage workspace effectively by configuring sizes for new workspace and schema requests and additional space for existing workspace.



# Copyright Information

- Neither Viskosity Technologies nor the authors guarantee this document to be error-free. Please provide comments and feedback to [Charles@viskosity.com](mailto:Charles@viskosity.com)
- Apress Copyright 2007. This document or any portions of this document may not be reproduced without the expressed written consent from the authors of the Oracle Database 11g New Features for DBAs and Developers book.
- Contact Information:
  - Charles Kim – [charles@viskosity.com](mailto:charles@viskosity.com)





# Viscosity Technology Service Offerings

VISKosity professionals provide end-to-end consulting services in the following areas:

- Real Application Clusters (RAC) - 9i, 10g, 10gR2 and 11g
  - RAC Assessment
  - Infrastructure Readiness
- EBusiness Suite Administration (11i and 12i)
- Data Warehousing
- Business Intelligence
- Database Security
- Enterprise Systems Management (ITIL and ITSM)
- Remote Database Administration

Contact us:  
Email: [info@Viscosity.com](mailto:info@Viscosity.com)  
Phone: 214.502.1231  
URL: <http://viscosity.com>





# Partners and Associations

TRIORA Group

<http://www.trioragroup.com>



Specialists in EBusiness Suite Implementations

Contact: [rob.mcmillan@trioragroup.com](mailto:rob.mcmillan@trioragroup.com)

Novara Solutions

Oracle EBusiness Suite

Contact: Vic Roselli