



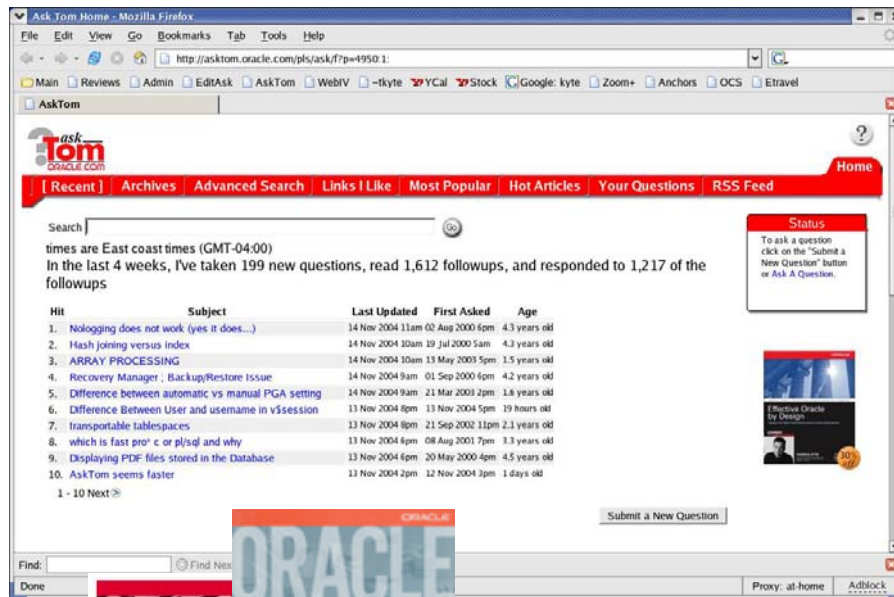
Top 10 Favorite New Features in 10gR2

Tom Kyte

ORACLE®

ORACLE®

Who am I




- Been with Oracle for over 12 years
- User of Oracle for over 17 years
- The “Tom” behind AskTom in Oracle Magazine
www.oracle.com/oramag
- Expert Oracle Database Architecture: *9i and 10g Programming Techniques and Solutions*
- Effective Oracle by Design
- Beginning Oracle
- Expert One on One Oracle

Top Ten Oracle *10g Release 1* Things

Online Segment Shrink

- Table fills over time
- You delete rows
- Lots of whitespace
 - You want to get it back
 - You full scan and want it smaller
- In the past
 - Alter table move, Alter index rebuild
 - Export/Import
 - Offline
- Not any more...
- Demo [shrink.sql](#)



Less Invalidations

- Change tables without
 - Recompiling Stored Procedures
 - Recompiling Views
- Compiled stored objects can depend on the synonym
- Uses synonyms
- I used to use loopback database links and views, this is infinitely *easier* and faster

Less invalidations

```
SQL> create table t1 ( x int );
```

```
SQL> create table t2 ( x int, y int );
```

```
SQL> create synonym t for t1;
```

```
SQL> create or replace procedure p
```

```
2 as
```

```
3 begin
```

```
4   for x in (select * from t) loop null; end loop;
```

```
8 end;
```

```
9 /
```

```
SQL> create or replace view v
```

```
2 as
```

```
3 select x from t
```

```
4 /
```

Less invalidations

```
ORA9IR2> select status from user_objects  
          where object_name in( 'P', 'V' );
```

```
STATUS  
-----  
VALID  
VALID
```

```
ORA9IR2> create or replace synonym t for t2;  
Synonym created.
```

```
ORA9IR2> select status from user_objects  
          where object_name in( 'P', 'V' );
```

```
STATUS  
-----  
INVALID  
INVALID
```


Less invalidations

```
ORA10G> select status from user_objects  
         where object_name in( 'P', 'V' );
```

```
STATUS  
-----  
VALID  
VALID
```

```
ORA10G> create or replace synonym t for t2;  
Synonym created.
```

```
ORA10G> select status from user_objects  
         where object_name in( 'P', 'V' );
```

```
STATUS  
-----  
VALID  
VALID
```

Case insensitive...

- Sorts
- Compares
- 100% transparent to applications
- Can still use indexes

ci.sql

Online Redefine

- Enhanced Online Table Redefinition
 - Easy cloning of indexes, grants, constraints, etc.
 - Down to 4 easy steps
 1. Create new, empty table
 2. Start redef (initial copy)
 3. Copy_Table_Dependents, instead of manually indexing, altering, etc
 4. Finish redef
 - Convert from long to LOB online
 - Allow unique index instead of primary key

Redef.sql

SQL Profiles

- Tune SQL without tuning SQL
- Looks at 'how the query is used'
 - First rows vs All rows for example
- Based on execution history
- Will also store additional statistics relating to cross object relationships
- For example.....

SQL Profiles

```
ops$tkyte@ORA10G> create or replace procedure p
 2  as
 3      cursor c1
 4      is
 5      select object_id, object_name
 6          from sqlprof
 7          order by object_id;
 9      l_object_id  sqlprof.object_id%type;
10      l_object_name sqlprof.object_name%type;
11  begin
12      open c1;
13      for i in 1 .. 10
14      loop
15          fetch c1 into l_object_id, l_object_name;
16          exit when c1%notfound;
17          -- .....
18      end loop;
19  end;
20  /
```

Procedure created.

SQL Profiles

```
SELECT OBJECT_ID, OBJECT_NAME FROM SQLPROF ORDER BY OBJECT_ID
```

call	count	cpu	elapsed	query	current	rows
Parse	1	0.00	0.00	0	0	0
Execute	1	0.00	0.00	0	0	0
Fetch	10	0.07	0.10	659	0	10
total	12	0.07	0.10	659	0	10

Misses in library cache during parse: 1

Optimizer mode: ALL_ROWS

Parsing user id: 410 (recursive depth: 1)

Rows	Row Source Operation
10	SORT ORDER BY (cr=659 pr=0 pw=0 time=101152 us)
47487	TABLE ACCESS FULL SQLPROF (cr=659 pr=0 pw=0 time=47604 us)

SQL Profiles

```
ops$tkyte@ORA10G> declare
 2             l_sql_id v$sql.sql_id%type;
 3  begin
 4
 5      select sql_id  into l_sql_id
 6             from v$sql
 7            where sql_text = 'SELECT OBJECT_ID, OBJECT_NAME
 8                               FROM SQLPROF ORDER BY OBJECT_ID';
 9
10      dbms_output.put_line(
11          sys.dbms_sqltune.create_tuning_task
12          ( sql_id      => l_sql_id,
13            task_name => 'sqlprof_query' ) || ' in place...' );
14      dbms_sqltune.execute_tuning_task
15      ( task_name => 'sqlprof_query' );
16  end;
17  /
PL/SQL procedure successfully completed.
```

SQL Profiles

```
ops$tkyte@ORA10G> SELECT DBMS_SQLTUNE.REPORT_TUNING_TASK( 'sqlprof_query')
2      FROM DUAL;
```

```
DBMS_SQLTUNE.REPORT_TUNING_TASK( 'SQLPROF_QUERY' )
```

```
-----
GENERAL INFORMATION SECTION
```

```
...
```

```
-----
SQL ID      : 3zfpas86satsm3
SQL Text:   SELECT OBJECT_ID, OBJECT_NAME FROM SQLPROF ORDER BY OBJECT_ID
```

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

```
-----
1- SQL Profile Finding (see explain plans section below)
```

```
-----
A potentially better execution plan was found for this statement.
```

```
Recommendation (estimated benefit: 99.45%)
```

```
-----
Consider accepting the recommended SQL profile.
execute :profile_name := dbms_sqltune.accept_sql_profile(task_name =>
                    'sqlprof_query')
```


SQL Profiles

EXPLAIN PLANS SECTION

1- Original With Adjusted Cost

Plan hash value: 1044598349

Id	Operation	Name	Rows	Bytes	TempSpc	Cost (%CPU)	Time
0	SELECT STATEMENT		47487	1391K		546 (3)	00:00:07
1	SORT ORDER BY		47487	1391K	3736K	546 (3)	00:00:07
2	TABLE ACCESS FULL	SQLPROF	47487	1391K		151 (2)	00:00:02

2- Using SQL Profile

Plan hash value: 337606071

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		10	300	3 (0)	00:00:01
1	TABLE ACCESS BY INDEX ROWID	SQLPROF	47487	1391K	3 (0)	00:00:01
2	INDEX FULL SCAN	SQLPROF_PK	10		2 (0)	00:00:01

DBMS_ADVANCED_REWRITE


- Fix that unfixable SQL statement
 - For correctness
 - For performance
- When you cannot otherwise touch it

rewrite.sql

ASM – Automatic Storage Management

- Database file system
 - Striping
 - Mirroring
 - Rebalancing
- No more moving files from point A to point B – ever
- Fundamentally changes the way you look at storage
- It's a performance thing
- It's a manageability thing
- It's an *availability* thing




ASM – Automatic Storage Management



- Choices for disk group redundancy
 - External: defers to hardware mirroring
 - Normal: 2-way mirroring
 - High: 3-way mirroring




ASM – Automatic Storage Management

- Mirror at extent level
- Mix primary & mirror extents on each disk





ASM – Automatic Storage Management

- Mirror at extent level
- Mix primary & mirror extents on each disk

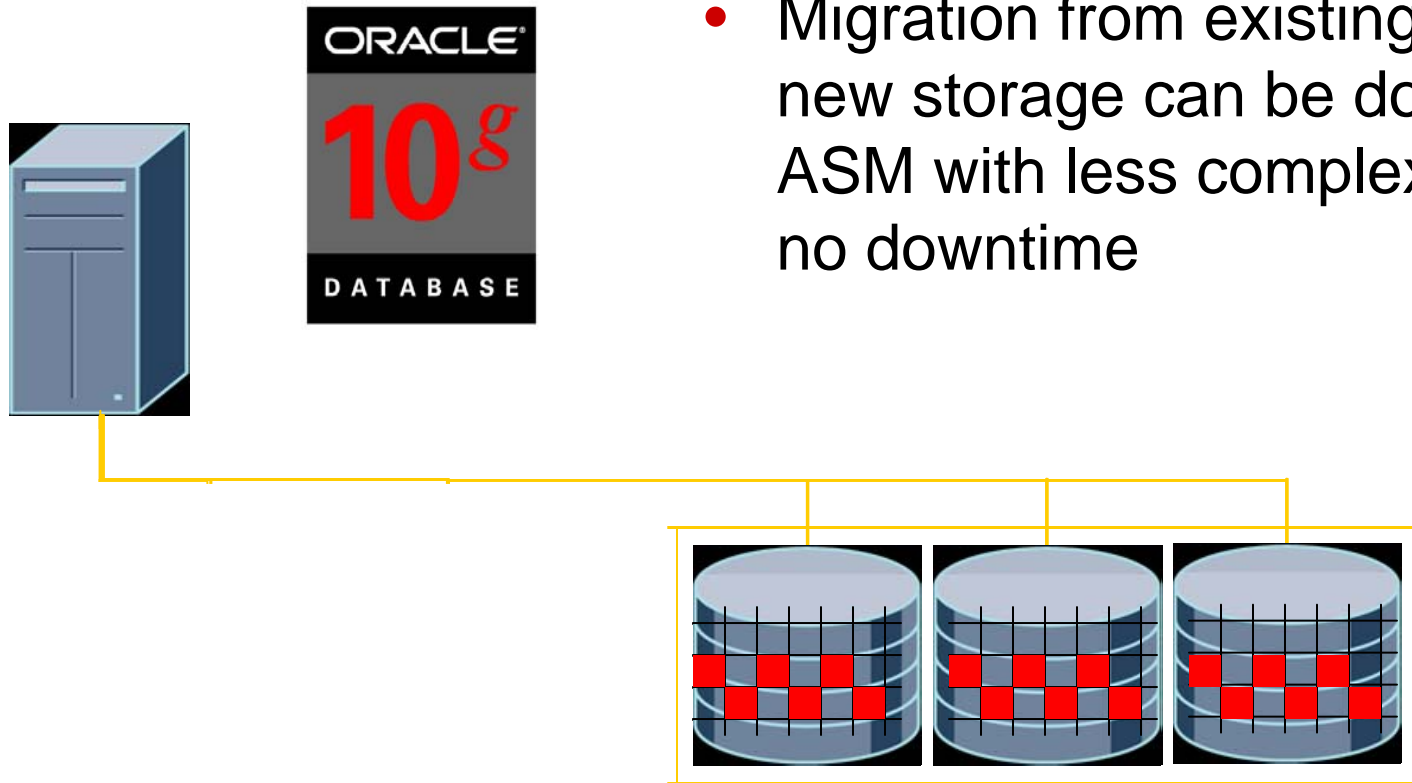


ASM – Automatic Storage Management

- No hot spare disk required
 - Just spare capacity
 - Failed disk load spread among survivors
 - Maintains balanced I/O load



ASM – Automatic Storage Management



PL/SQL Performance

- Out of the box – no changes, it is faster
- It's all about PL/SQL.
 - Optimizing compiler
 - *Bad_code.sql (on next slides)*
- Implicit Array Fetching
 - *Af.sql (on following slides)*

PL/SQL Performance

```
SQL> create or replace procedure P
  2  as
  3    a positiven := 1 /* binary_integer subtypes */;
  4    b positiven := 1;
  5    c positiven := 1;
  6    d positiven:=1;
  7    t0 integer;
  8  begin
  9    t0 := Dbms_Utility.Get_Time();
 10    for j in 1..1000000
 11    loop
 12      b := j /* pls_integer to binary_integer conversion */;
 13      d := 42 /* constant assignment within a loop */;
 14      c := d + b /* can be combined... */;
 15      a := b + c /* ...c not used except here */;
 16    end loop;
 17    Dbms_Output.Put_Line(Dbms_Utility.Get_Time()-t0);
 18  end P;
 19  /
```

Procedure created.

PL/SQL Performance

```
ops$tkyte@ORA9IR2> exec p
```

```
204
```

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

```
ops$tkyte@ORA9IR2> exec p
```

```
208
```

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

```
ops$tkyte@ORA9IR2> exec p
```

```
211
```

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

```
ops$tkyte@ORA10GR2> exec p
```

```
43
```

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

```
ops$tkyte@ORA10GR2> exec p
```

```
45
```

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

```
ops$tkyte@ORA10GR2> exec p
```

```
42
```

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

PL/SQL Performance

```
begin
    for x in
        ( select *
          from big_table.big_table
          where rownum <= 10000 )
    loop
        null;
    end loop;
end;
```

PL/SQL Performance

```
declare
    type array is table of big_table%rowtype;
    l_data array;
    cursor c is
        select * from big_table where rownum <= 1000;
begin
    open c;
    loop
        fetch c bulk collect into l_data limit 100;
        for i in 1 .. l_data.count
            loop
                null;
            end loop;
        exit when c%notfound;
    end loop;
    close c;
end;
```

PL/SQL Performance 9i

```
SELECT * FROM BIG_TABLE.BIG_TABLE WHERE ROWNUM <= 10000
```

call	count	cpu	elapsed	query	rows
Parse	1	0.01	0.00	0	0
Execute	1	0.00	0.00	0	0
Fetch	10001	0.15	0.17	10005	10000
total	10003	0.16	0.17	10005	10000

PL/SQL Performance 10g

```
SELECT * FROM BIG_TABLE.BIG_TABLE WHERE ROWNUM <= 10000
```

call	count	cpu	elapsed	query	rows
Parse	1	0.00	0.00	0	0
Execute	1	0.00	0.00	0	0
Fetch	101	0.05	0.07	152	10000
total	103	0.05	0.07	152	10000

PL/SQL Performance

- Native Floats
 - IEEE floating point support in SQL and PLSQL
 - $\pi / 4 = 1 - 1/3 + 1/5 - 1/7 + 1/9...$
 - *Bindbl.sql*
- Predicable fixed width datatype
- Scientific applications will truly appreciate this
- But remember – numbers are very “accurate”, floats/doubles “not as accurate”
- But also remember – 11 bit exponent gives larger range of values

Flashback

- Flashback Query
 - 9iR1 – primitive
 - Had to open flashback cursors before doing any DML
 - It worked, but was not “easy”
 - 9iR2 – sophisticated
 - No need to open cursors before doing modifications
 - Can flashback in SQL, no packages needed
 - Can flashback in modifications
 - It worked and was much easier
 - o Could join the table with itself as of N minutes ago
 - o Put update a set of rows, put them back as they were N minutes ago

Flashback

- In 10g
 - Flashback Query
 - Flashback Table
 - Flashback Row History
 - Flashback Drop
 - Flashback Database

Flashback – Flashback Row History

- Instead of “show me the data as of”, you can say “show me all versions of the data between”

```
Select ename, sal
      from emp versions between timestamp a and b
      where ename = 'SCOTT'
```

ENAME	SAL
-----	-----
SCOTT	3000
SCOTT	3300
SCOTT	3630
...	

Flashback – Flashback Row History

- See related information about each row
 - SCN range the row was “valid for”
 - Time range (approx) the row was valid for
 - The transaction id (XID) that modified the row
 - The operation (I/U/D) that was performed on the row

```
select ename, sal,  
       versions_operation,  
       versions_starttime,  
       versions_endtime,  
       versions_startscn,  
       versions_endscn,  
       versions_xid  
from emp versions between timestamp &A and &B  
where empno = 7788  
order by versions_startscn nulls first
```

Fb7.sql

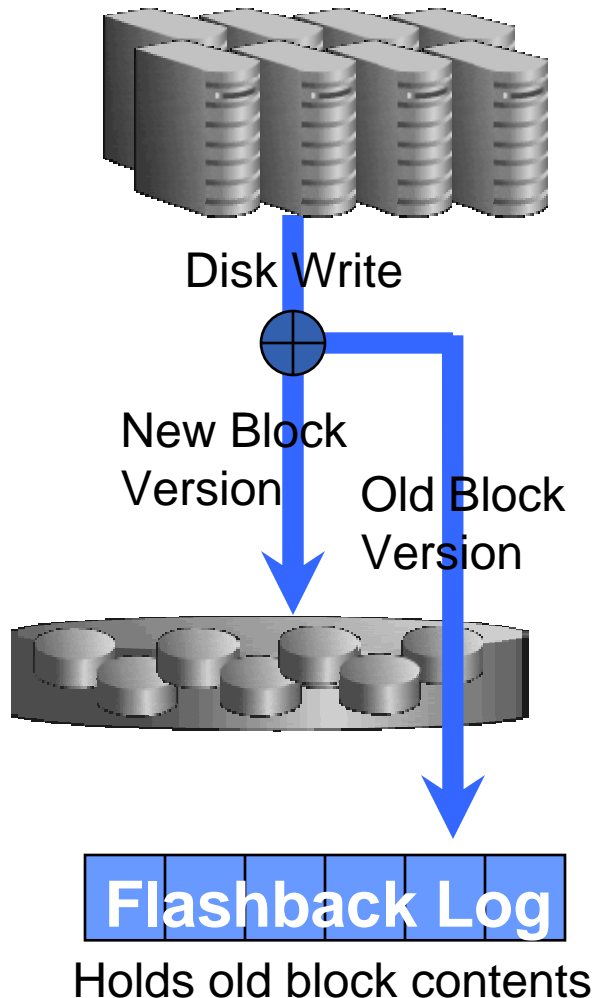
Flashback – Flashback Drop Table

- Uses a recycle bin metaphor
 - Table/indexes/triggers/constraints are renamed
 - Extents are *not* deallocated immediately
 - Whether you can undrop something depends entirely on how “full” your tablespace was!
 - This is useful for “whoops” recovery right after the fact
 - You may be able to undrop a table months after the fact
 - You may *not* be able to undrop a table seconds after the fact!

Flashback – Flashback Table

- Multiple table T's may exist in the recycle bin
 - HTML/DB demo story
- You should rename indexes et.al. after the fact
- RI (foreign keys from child tables) are not “undropped”
 - You had to specify “cascade constraints”
 - The fkeys are not really part of “this table” anyway

Flashback Database



- A new strategy for point in time recovery
- Flashback Log captures old versions of changed blocks
 - Think of it as a continuous backup
 - Replay log to restore DB to time
 - Restores just **changed** blocks
- It's **fast** - recover in minutes, not hours
- It's **easy** - single command restore
`SQL> Flashback Database to '2:05 PM'`

“Rewind” button for the Database


Top Ten Things about Oracle Database 10g Release 2

And a couple more

Top Ten Things about Oracle Database 10g Release 2

Conditional Compilation

- Like `#ifdef` in C
- Leave debug code in production
- Version independent code
- Best practices in unit testing
- *Quick Demo...* [cond.sql](#)



Full Database Transports


- Use RMAN to effectively back up platform X and restore to Y.
- Must have same byte order

```
SQL> select endian_format, platform_name
2      from v$transportable_platform
3      order by endian_format
4      /
```

ENDIAN_FORMAT	PLATFORM_NAME
Big	HP-UX IA (64-bit) Solaris[tm] OE (32-bit) Apple Mac OS HP-UX (64-bit) IBM zSeries Based Linux AIX-Based Systems (64-bit) Solaris[tm] OE (64-bit) IBM Power Based Linux
Little	Solaris Operating System (x86) Microsoft Windows 64-bit for AMD Linux 64-bit for AMD Microsoft Windows IA (64-bit) HP Open VMS HP Tru64 UNIX Linux IA (32-bit) Microsoft Windows IA (32-bit) Linux IA (64-bit)

XQuery

- Native implementation of XQuery in the Database (<http://www.w3.org/TR/xquery/>)
- A programming language of sorts, just to query XML documents
- Nearly 100 functions and operators
- Can do standalone XQuery or mix it with SQL
- *Quick Demo...* [xquery.sql](#)



SGA Direct Attach

- For when the database is not responsive
- Or “the system is very slow”
- We can now diagnose the problem, without having to ‘bounce’



Transparent Data Encryption

PALO ALTO Data stolen from children's psychiatric center
Diana Walsh, Chronicle Staff Writer *Tuesday, September 20, 2005*

A computer tape that stored **psychiatric records, evaluations and Social Security numbers of thousands of children** treated at the highly regarded Children's Health Council in Palo Alto was stolen two weeks ago, the executive director of the agency said Monday.

The backup tape, **which also contained payroll data on hundreds of current and former employees and credit card information from parents of patients**, and a computer hard drive were taken from a locked computer server room shortly after Labor Day, according to Stephen Joffe, executive director of Children's Health Council.

"We're all victims in this," said Joffe, whose personal payroll data are among the information on the missing tape. "We're really concerned. It's a terrible thing that someone would take this information."


Transparent Data Encryption

- Revolutionary new feature
- Keyword – *transparent*
- Protect data at rest
- Protects your data in the event of theft
- Not about preventing access, not access control
 - Grant is access control
 - VPD is access control
- *Quick Demo...* [tde.sql](#)




AWR Data is Transportable

- DBMS_SWRF_INTERNAL
 - AWR_EXTRACT
 - AWR_LOAD
 - MOVE_TO_AWR
- Review performance metrics, without affecting performance
- Consolidate many AWR information from many sources into a single database



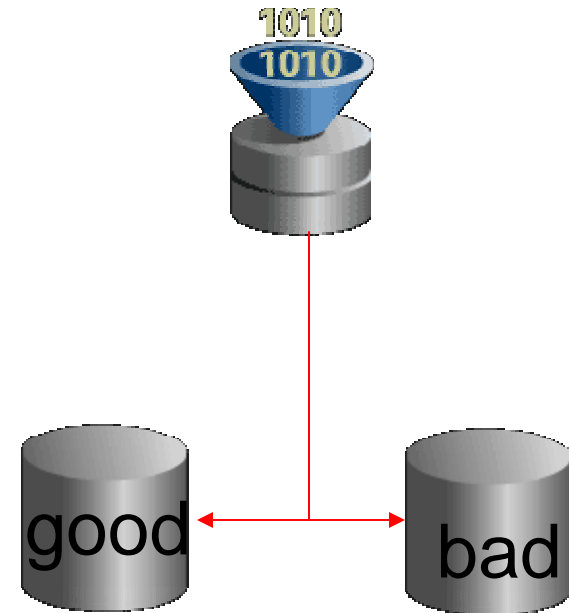
Move



Consolidate


LOG ERRORS

- This will change the way ETL is done.
- Modifications are no longer “*all for one, and one for all*”
- Errors can be captured in a logging table, for easy resolution afterwards
- Like a bad file from sqlldr, only better.
- *Quick Demo...* [logerrors.sql](#)



Audit in XML

- Configure auditing information to be records in XML
 - `audit_file_dest = '/auditfs'`
 - `audit_trail = xml`
- `V$XML_AUDIT_TRAIL`
 - Query these files relationally
 - Or XQuery
 - Or Xpath



Legacy OS Audit Trail

```
Audit file /home/ora10gr2/rdbms/audit/ora_14971.aud
Oracle Database 10g Enterprise Edition Release 10.2.0.1.0 - Production
With the Partitioning, OLAP and Data Mining options
ORACLE_HOME = /home/ora10gr2
System name:      Linux
Node name:        me
Release:          2.6.9-11.EL
Version:          #1 Fri May 20 18:17:57 EDT 2005
Machine:          i686
Instance name:    ora10gr2
Redo thread mounted by this instance: 0 <none>
Oracle process number: 0
Unix process pid: 14971, image: oracle@me
```

```
Fri Aug 19 17:56:13 2005
ACTION : 'CONNECT'
DATABASE USER: '/'
PRIVILEGE : SYSDBA
CLIENT USER: tkyte
CLIENT TERMINAL: Not Available
STATUS: 0
```

XML OS Audit Trail


The screenshot displays the Altova XMLSpy application window titled "ora_3340.xml". The main editor shows the following XML content:

```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <Audit
3   xmlns="http://xmlns.oracle.com/oracleas/schema/dbserver_audittrail-10_2.xsd"
4   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
5   xsi:schemaLocation="http://xmlns.oracle.com/oracleas/schema/dbserver_audittrail-10_2.xsd"
6
7   >
8     <Version>10.2</Version>
9     <AuditRecord>
10       <Audit_Type>1</Audit_Type>
11       <Session_Id>276</Session_Id>
12       <StatementId>1</StatementId>
13       <EntryId>1</EntryId>
14       <Extended_Timestamp>2005-09-20T11:44:38.171000</Extended_Timestamp>
15       <DB_User>SYSTEM</DB_User>
16       <OS_User>CVM-LAP03\Cary Millsap</OS_User>
17       <Userhost>HOTSOS\CVM-LAP03</Userhost>
18       <OS_Process>804:3340</OS_Process>
19       <Terminal>CVM-LAP03</Terminal>
20       <Instance_Number>0</Instance_Number>
21       <Action>100</Action>
22       <Returncode>0</Returncode>
23       <Comment_Text>Authenticated by: DATABASE</Comment_Text>
24     </AuditRecord>
25   </Audit>
```

The interface includes a Project Explorer on the left with a tree view containing folders like "Org", "Expe", "Inter", "Purc", "Indu", "XML", and "Tam". The bottom status bar shows "XMLSpy v2005 rel. 3 U Registered to Cary Millsap (private) ©1998-2005 Altova GmbH" and "Ln 5, Col 92".

Online Transportable TS

- Transporting introduced in 8i
- 9i added multiple block size support
- 10g R1 added cross platform
- 10g R2 – makes it an online operation
 - No need to make tablespace readonly
 - RMAN does it (don't have to use RMAN for backup)
 - Uses backups
 - Can be done as of a point in time



Honorable Mentions

- Autotrace – Quick Demo [autotrace.sql](#)
- DBMS_OUTPUT – Quick Demo [dbmsoutput.sql](#)
- Restore Points
- Automatic Failover to standby
- Datapump Metadata Compression
- Secure Backup

Questions

Better Trace Files

```
select count(*) from all_objects
```

Elapsed times include waiting on following events:

Event waited on	Times	Max. Wait	Total Waited
-----	Waited	-----	-----
SQL*Net message to client	2	0.00	0.00
db file sequential read	78	0.02	0.32
SQL*Net message from client	2	0.00	0.00

```
$ grep 'db file sequential read' 10gr1.trc | head -4
```

```
WAIT #5: nam='db file sequential read' ela= 20506 p1=1 p2=6746 p3=1
```

```
WAIT #9: nam='db file sequential read' ela= 17400 p1=1 p2=57783 p3=1
```

```
WAIT #5: nam='db file sequential read' ela= 23100 p1=1 p2=25604 p3=1
```

```
WAIT #10: nam='db file sequential read' ela= 6944 p1=1 p2=5867 p3=1
```

```
$ grep 'db file sequential read' 10gr2.trc | head -4
```

```
WAIT #7: ... file#=1 block#=6348 blocks=1 obj#=348 tim=1100647249561141
```

```
WAIT #7: ... file#=1 block#=339 blocks=1 obj#=348 tim=1100647249574644
```

```
WAIT #9: ... file#=1 block#=1847 blocks=1 obj#=348 tim=1100647249616388
```

```
WAIT #9: ... file#=1 block#=1846 blocks=1 obj#=348 tim=1100647249620503
```


Asynchronous Commit

- COMMIT WORK WRITE BATCH NOWAIT
- COMMIT WORK
- Useful for data load applications that want to checkpoint.
- Removes the `log file sync` waits
- Not for OLTP!
- *Quick Demo... `./sync, ./async`*



The logo for the Suncoast Oracle Users' Group is centered in a white rounded rectangle. It features the text "SUNCOAST" in blue, "ORACLE" in white on a red rectangular background, and "USERS' GROUP" in blue, all in a bold, sans-serif font. A large, yellow, stylized "C" shape is positioned behind the text.

SUNCOAST
ORACLE
USERS' GROUP

Monthly
4th Thursday
6pm – 8pm

IBM Center
Rocky Point

[WWW.SOUUG.NET](http://www.souug.net)

ORACLE®