

ORACLE®

11 things about Oracle Database 11g Release 2

Thomas Kyte http://asktom.oracle.com/





1 Do it yourself Parallelism

Incrementally modify a table in parallel

- Used to do this manually all of the time
 - Search for 'diy parallel' on asktom...
 - Spent part of a chapter on 'how to' in Expert Oracle Database Architecture
- I split by rowid ranges
 - Split table into N equi-sized, non-overlapping chunks
 - Create a job passing in the low and high rowids for each range
 - Job would process "where rowid between :lo and :hi"
- Or by primary key ranges using NTILE()
- DBMS_PARALLEL_EXECUTE automates both approaches and makes it easy (and more functional)

Divp.sal



2 Analytics are the coolest thing to happen to SQL since the keyword SELECT



More Analytics!

- Long awaited LISTAGG
 - First did STRAGG in 9iR2 with user defined aggregates
 - Oracle Database 10g gave us a sys_connect_by_path 'trick'
 - Oracle Database 11g Release 2 makes it 'easy'



SQL> select deptno, 2 substr(3 max(sys connect by path(ename, '; ')), 4 3) enames 5 from (select deptno, 6 7 ename, row number() 8 9 over 10 (partition by deptno order by ename) rn 11 12 from emp DEPTNO ENAMES 13) 10 CLARK; KING; MILLER start with rn = 114 20 ADAMS; FORD; JONES; connect by prior deptno = deptno 15 SCOTT; SMITH 16 and prior rn+1 = rn17 group by deptno 30 ALLEN; BLAKE; 18 order by deptno JAMES; MARTIN; 19 / TURNER; WARD

SQL> select deptno,

2		<pre>listagg(ename, '; ')</pre>
3		within group
4		(order by ename) enames
5	from	emp
6	group	by deptno
7	order	by deptno
8	/	

DEPTNO ENAMES

10 CLARK; KING; MILLER 20 ADAMS; FORD; JONES; SCOTT; SMITH

_ _ _ _ _ _ _ _ _ _ _ _ _ _ _

30 ALLEN; BLAKE; JAMES; MARTIN; TURNER; WARD



SQL> select deptno,

2	ename,
3	row_number()
4	over (partition by deptno
5	order by ename) rn,
6	first_value(ename)
7	over (partition by deptno
8	order by ename) "1st ename",
9	nth_value(ename,3)
10	over (partition by deptno
11	order by ename) "3rd ename",
12	last_value(ename)
13	over (partition by deptno
14	order by ename
15	rows between current row
16	and unbounded following) "last ename"
17	from emp
18	order by deptno, ename
19	/

		DEP	TNO	ENAME	RN	1st e	3rd ena	last en
			10	CLARK	1	CLARK		MILLER
SQL>	select	deptno,		KING	2	CLARK		MILLER
2		ename,		MTT.T.ER	2	CT.ARK	MTT.T.ER	MTT.T.FR
3		row_number()			5		MIDDER	мтынык
4		over (partition by deptno						
5		order by ename) rn,	~~	1011/2	-	1011/2		a)/T ====
6		first_value(ename)	20	ADAMS	T	ADAMS		SMITH
8		order by ename) "1st ename",		FORD	2	ADAMS		SMITH
9		nth_value(ename,3)		TONEC	С	שעעעע	TONEC	CMTTI
10		over (partition by deptno		JONES	3	ADAMS	JONES	SMITH
11		order by ename) "3rd ename",		SCOTT	4	ADAMS	JONES	SMITH
12		last_value(ename)			_		••••	
13		over (partition by deptno		SMITH	5	ADAMS	JONES	SMITH
14		order by ename						
15		rows between current row						
17	from	emp	30	ALLEN	1	ALLEN		WARD
18	order	by deptno, ename		BLAKE	2	ALLEN		WARD
19	/				•			
				JAMES	3	ALLEN	JAMES	WARD
				MARTIN	4	ALLEN	JAMES	WARD
				TURNER	5	ALLEN	JAMES	WARD
				MADD	6	אד ד דיאד	тамра	
				WARD	Ø	АГГСИ	Oamed	WARD



3 Execute on a directory



External Tables can run code now

- External tables allow for a preprocessor
 - Program is run when you SELECT from external table
 - The 'location' is passed to the script/executable
 - The executable does whatever it wants and writes to stdout
 - Stdout is treated as the input file
- We need a way to control who can do what
- GRANT EXECUTE ON DIRECTORY handles that

EXECUTE and PREPROCESSOR

ops\$tkyte%ORA11GR2> CREATE or replace DIRECTORY load_dir

- 2 AS '/mnt/hgfs/docs/Presentations/Seminar/11gr2'
- 3 /

Directory created.

ops\$tkyte%ORA11GR2> CREATE or replace DIRECTORY exec_dir

- 2 AS '/mnt/hgfs/docs/Presentations/Seminar/11gr2'
- 3 /

Directory created.

EXECUTE and PREPROCESSOR

ops\$tkyte%ORA11GR2> CREATE TABLE EMP ET

- 2 (
- 3 "EMPNO" NUMBER(4),
- "ENAME" VARCHAR2(10), 4
- 5 "JOB" VARCHAR2(9),
- 6 "MGR" NUMBER(4),
- 7 "HIREDATE" DATE,
- 8 "SAL" NUMBER(7,2),
- 9 "COMM" NUMBER(7,2),
- 10 "DEPTNO" NUMBER(2)
- 11)

17

18

20

19

21) 22 /

)

Table created.

- 12 ORGANIZATION external
- (TYPE oracle_loader 13
- DEFAULT DIRECTORY load dir 14
- ACCESS PARAMETERS 15

- 16

location ('emp.dat.gz')

- (RECORDS DELIMITED BY NEWLINE

FIELDS TERMINATED BY "|" LDRTRIM

- - preprocessor exec dir: 'run gunzip.sh'

EXECUTE and PREPROCESSOR

ops\$tkyte%ORA11GR2> !file emp.dat.gz
emp.dat.gz: gzip compressed data, was "emp.dat", from Unix, last
modified: Wed Oct 7 12:48:53 2009

ops\$tkyte%ORA11GR2> !cat run_gunzip.sh
#!/bin/bash

/usr/bin/gunzip -c \$*

ops\$tkyte%ORA11GR2> select empno, ename from emp_et where rownum <= 5;

EMPNO ENAME ----- -----7369 SMITH 7499 ALLEN 7521 WARD 7566 JONES 7654 MARTIN



EXECUTE and PREPROCESSOR, interesting idea...

```
ops$tkyte%ORA11GR2> CREATE TABLE 1s
 2 (
 3
      line varchar2(255)
 4)
 5 ORGANIZATION external
 6 ( TYPE oracle loader
 7
    DEFAULT DIRECTORY load_dir
     ACCESS PARAMETERS
 8
 9
     ( RECORDS DELIMITED BY NEWLINE
10
            preprocessor exec_dir:'run_ls.sh'
     FIELDS TERMINATED BY "|" LDRTRIM
11
12)
13
      location ( 'run_ls.sh')
14)
15 /
```

Table created.

ORACLE

EXECUTE and PREPROCESSOR, interesting idea...

ops\$tkyte%ORA11GR2> select * from ls;

LINE

11 things about 11gr2.ppt
diyp.sql
ebr.old.sql
ebr.sql
emp.ctl
emp.dat.gz
EMP_ET_26122.log
emp_et.sql
LS_26122.log
run_gunzip.sh
run_1s.sh

11 rows selected.







- ANSI SQL replacement for connect by
- Can be
 - Easier to understand than connect by
 - Unless of course, you have been using connect by for 22 years in which case it looks confusing

ops\$tkyte%ORA11GR2> with emp_data(ename,empno,mgr,1)

```
2 as
   (select ename, empno, mgr, 1 lvl from emp where mgr is null
 3
 4 union all
 5
    select emp.ename, emp.empno, emp.mgr, ed.l+1
 6
    from emp, emp data ed
 7
    where emp.mgr = ed.empno
 8
   )
 9
   SEARCH DEPTH FIRST BY ename SET order by
10
   select 1,
          lpad('*',2*1,'*')||ename nm
11
12
   from emp data
13
   order by order by
14
   /
```



L NM

- 1 **KING
- 2 ****BLAKE
- 3 *****ALLEN
- 3 *****JAMES
- 3 *****MARTIN
- 3 *****TURNER
- 3 *****WARD
- 2 ****CLARK
- 3 *****MILLER
- 2 ****JONES
- 3 *****FORD
- 4 ******SMITH
- 3 *****SCOTT
- 4 ******ADAMS

14 rows selected.

```
ops$tkyte%ORA11GR2> with data(r)
  2 as
  3
     (select 1 r from dual
  4 union all
  5 select r+1 from data where r < 5
  6
     )
     select r, sysdate+r
  7
  8
       from data;
         R SYSDATE+R
         1 08-OCT-09
         2 09-OCT-09
         3 10-OCT-09
         4 11-OCT-09
         5 12-OCT-09
```



- ANSI SQL replacement for connect by
- Can be
 - Easier to understand than connect by
 - Unless of course, you have been using connect by for 22 years in which case it looks confusing
 - <u>Used to solve Sudoku puzzles!</u>



5 Improved Time Travel



Improved Time Travel

- Flashback Data Archive
 - Query data as of 5 days, 5 weeks, 5 months, 5 years whatever in the past
 - http://www.oracle.com/technology/oramag/oracle/08jul/o48totalrecall.html
 - Article by Jonathan Gennick on this feature for more info
- How does it work...



How Does Flashback Data Archive Work?

- Primary source for history is the undo data
- History is stored in automatically created history tables inside the archive
- Transactions and its undo records on tracked tables marked for archival
 - Undo records not recycled until history is archived
- History is captured asynchronously by new background process (fbda)
 - Default capture interval is 5 minutes
 - Capture interval is self-tuned based on system activities
 - Process tries to maximize undo data reads from buffer cache for better performance
 - INSERTs do not generate history records



Oracle Database 11g Release

Total Recall Schema Evolution Support

- Alter base table history table automatically adjusts
 - Drop, Rename, Modify Column
 - Drop, Truncate Partition
 - Rename, Truncate Table
- Flashback query supported across DDL changes



- Complex DDL changes (e.g. table split) accommodated
 - Associate/Diassociate history table via DBMS_FLASHBACK_ARCHIVE package



6 You've got Mail



- As files arrive in some directory
 - An event is generated
 - And your code can be invoked to deal with it...



ops\$tkyte%ORA11GR2> begin

dbr	ms_schedule:	c.create	e_	credential(
	$credential_$	_name =:	>	'watch_credential',
	username	=;	>	'tkyte',
	password	=;	>	`foobar');
end;				
/				
	dbı end; /	<pre>dbms_scheduler credential_ username password end; /</pre>	<pre>dbms_scheduler.create credential_name == username == password == end; /</pre>	<pre>dbms_scheduler.create_ credential_name => username => password => end; /</pre>



```
ops$tkyte%ORA11GR2> create or replace directory MY_FILES as
'/home/tkyte/files'
2 /
```

Directory created.

```
ops$tkyte%ORA11GR2> create table files
```

```
2 (
3 file_name varchar2(100),
4 loaded timestamp,
5 contents clob
6 );
```

Table created.



```
ops$tkyte%ORA11GR2> create or replace procedure process files
     (p payload in sys.scheduler filewatcher result)
  2
  3 is
  4
         l clob clob;
  5
         l bfile bfile;
  6
    begin
  7
         insert into files
         (loaded, file_name, contents )
  8
  9
         values (p payload.file timestamp,
          p_payload.directory_path || '/' || p_payload.actual_file_name,
 10
 11
          empty_clob()
 12
         ) returning contents into 1 clob;
 13
         l bfile := bfilename( 'MY_FILES', p_payload.actual_file_name );
 14
 15
         dbms lob.fileopen( l bfile );
 16
         dbms lob.loadfromfile( l clob, l bfile, dbms lob.getlength(l bfile) );
 17
         dbms_lob.fileclose( l_bfile );
 18 end;
 19
    1
```

Procedure created.

ops\$tkyte%ORA11GR2> begin

2	dbms_scheduler.create_	_pro	ogram(
3	program_name	=>	'file_watcher',
4	program_type	=>	<pre>'stored_procedure',</pre>
5	program_action	=>	'Process_Files',
6	number_of_arguments	=>	1,
7	enabled	=>	false);
8	dbms_scheduler.define_	_met	adata_argument(
9	program_name	=>	'file_watcher',
10	metadata_attribute	=>	<pre>'event_message',</pre>
11	argument_position	=>	1);
12	dbms_scheduler.enable	('f:	ile_watcher');
13	end;		
14	/		

ops\$tkyte%ORA11GR2> begin

2	dbms_scheduler.creat	te_f	ile_watcher(
3	file_watcher_name	=>	'my_file_watcher',
4	directory_path	=>	<pre>'/home/tkyte/files',</pre>
5	file_name	=>	'*',
6	credential_name	=>	<pre>'watch_credential',</pre>
7	destination	=>	null,
8	enabled	=>	<pre>false);</pre>
9	end;		
10	1		

```
ops$tkyte%ORA11GR2> begin
```

2	dbms_scheduler.create_job(
3	job_name	=>	'my_file_job',		
4	program_name	=>	'file_watcher',		
5	event_condition	=>	<pre>'tab.user_data.file_size > 10',</pre>		
6	queue_spec	=>	'my_file_watcher',		
7	auto_drop	=>	false,		
8	enabled	=>	false);		
10	end;				
11	/				

PL/SQL procedure successfully completed.
ops\$tkyte%ORA11GR2> exec
 dbms_scheduler.enable('my_file_watcher,my_file_job');



ops\$tkyte%ORA11GR2> select * from files;

FILE_NAME	LOADED	CONTENTS		
/home/tkyte/files/file4.txt	07-OCT-09 07.37.22.000000 PM	hello world, ho		
		w are you		
		hello world, ho		
		w are you		
		hello world, ho		
		w are you		
		hello world, ho		
		w are you		







- Segments (tables, indexes, etc) normally allocate an initial extent
- They might be small, but they exist
- If you do something "small" (or fast) over and over a lot – it gets "big" (or slow)
- Many third party applications create thousands of tables
 - And then use 100 of them
- Deferred segment creation allows us to put off initial extent allocation until the first row is put into a segment.

SQL> alter session set

2 deferred_segment_creation=false; Session altered.

SQL> select segment name, SOL> create table t1 2 2 (x int extent id, 3 constraint t1_pk 3 bytes 4 primary key, 4 from user_extents 5 order by segment name; 5 y int 6 constraint t1 y 7 unique, SEGMENT NAM EXTENT ID BYTES 8 z clob 65536 т1 9) 0 65536 T1_PK 0 10 lob(z)11 store as t1_z_lob T1 Y 65536 0 T1_Z_LOB 65536 12 (index t1 z lobidx); 0 T1 Z LOBIDX 65536 Table created. 0

SQL> alter session set

2 deferred_segment_creation=true; Session altered.

No Change!

SQL>	create table <i>t2</i>
2	(x int
3	constraint <i>t2_pk</i>
4	primary key,
5	y int
6	constraint $t2_y$
7	unique,
8	z clob
9)
10	lob(z)
11	store as <i>t2_z_lob</i>
12	(index $t2_z$ _lobidx);
Table	e created.

SQL>	select	seg	gment_name,	,
2		ext	tent_id,	
3		by	tes	
4	from	use	er_extents	
5	order	by	segment_na	ame;
SEGM	ENT_NAM	E	KTENT_ID	BYTES
т1			0	65536
T1_PI	ĸ		0	65536
T1_Y			0	65536
T1_Z_	LOB		0	65536
T1_Z	LOBIDX		0	65536

SQL> insert into t2 values (1, 2, 'hello world'); 1 row created.

SQL> select segment_name,

- 2 extent_id,
- 3 bytes
- 4 from user_extents
- 5 order by segment_name;

SEGMENT_NAM	EXTENT_ID	BYTES
Т1	0	65536
Т1_РК	0	65536
T1_Y	0	65536
T1_Z_LOB	0	65536
T1_Z_LOBIDX	0	65536
Т2	0	65536
Т2_РК	0	65536
T2_Y	0	65536
T2_Z_LOB	0	65536
T2_Z_LOBIDX	0	65536
10 rows sele	cted.	



8 Flash Cache



Oracle Database 11g Release 2

Reduce I/O bandwidth requirement with Flash Cache

- A transparent extension of the database buffer cache using solid-state disk (SSD) technology
 - SSD acts as a Level 2 cache (SGA is Level 1)
 - Faster than disk (100x faster for reads)
 - Cheaper than memory (\$50 per gigabyte)
 - Large capacity (hundreds of gigabytes per flash disk)
- Fewer drives and better performance
 - For I/O throughput, users often use hundreds of drives today
 - Flash enables I/O throughput without all the drives
 - Large jobs complete faster





Flash Cache

How it works



Flash Cache





9 Parallel Improved



Automated Degree of Parallelism How it works



Parallel Statement Queuing

How it works



In-Memory Parallel Execution

How it works





10 Edition-based Redefinition



Yes, this is here twice But only because It is the *killer feature* Of Oracle Database 11g Release 2 It is worth 2 features

10+Edition-based Redefinition!

Online Application Upgrade

Edition-based redefinition

- Code changes are installed in the privacy of a new edition
- Data changes are made safely by writing only to new columns or new tables not seen by the old edition
- An editioning view exposes a different projection of a table into each edition to allow each to see just its own columns
- A crossedition trigger propagates data changes made by the old edition into the new edition's columns, or (in hot-rollover) vice-versa







DEMONSTRATION

Edition-based Redefinition

ebr.sql





How to get there



What are my upgrade paths?

Predictable performance post-upgrade





search.oracle.com



or oracle.com



