DBA TIPS AND TRICKS
REDO LOG GENERATION

Damir Vadas
http://damir-vadas.blogspot.com

TEB Informatika d.o.o.
www.teb-informatika.hr

HROUG
Rovinj, 10/2011

Koncept

- Smisao i značenje Oracle redo log
- Analiza (kada, koliko, gdje i zašto)
  - Tko generira logove trenutačno
  - Zbirni pregled kada su nastali redo logovi, te koliko to iznosi u MB
  - Kada zbirno imamo najviše generiranja redo log datoteka (idealno van radnog vremena)?
  - Koji segmenti su generirali redo logove (da li postoje „neobična" generiranja redo log)?
  - Koji SQL je uzrok pojedinog generiranja redo logova
- Q and A

USPJEŠAN DBA

Onaj koji svojim radom uspije napraviti balans između fire fighter-a i proactive DBA.

Ova demonstracija pokušat će pokazati JOŠ jedan takav primjer.

Smisao i značenje redo log

- Sadrži opis promjena urađenih nad pojedinim blokom u bazi
- Najvažniji element prilikom recovery operacija (pored UNDO segmenta)
- Generiranja redo log datoteka rezultira:
  - Povećanje pojedinih objekata (vjerojatno)
  - Privremeno povećanje UNDO segmenta (SIGURNO)
  - Povećanje baze (vjerojatno)
  - Povećanje backup-a (archived redolog files) (SIGURNO)
  - Ako nema povećanja veličine objekata onda je riječ o gomili UPDATE ili ROLLBACK operacija koje su također zanimljive DBA

Smisao i značenje redo log

- **v$log_history** sadrži povijesne podatke čiji opseg je definiran:
  - **MAXLOGHISTORY**
    - **Fiksan** def. prilikom kreiranja baze,
    - Promjena iziskuje rekreiranja control file-a
    - Nema potrebe mijenjati-uvjetovan slijedećim parametrom
  - **CONTROL_FILE_RECORD_KEEP_TIME**
    - **dinamički** (`ALTER SYSTEM`)
    - Default 7 dana
    - Ja prakticiram 31 dan (povećanje veličine control file datoteke):

```
SQL> show parameter CONTROL_FILE_RECORD_KEEP_TIME

NAME_COL_PLUS_SHOW_PARAM       TYPE        VALUE_COL_PLUS_SHOW_PARAM
------------------------------ ----------- --------------------------
control_file_record_keep_time  integer     31
```

Tko generira logove trenutačno

- Omogućeno monitoriranje RAC i non RAC baza
- Upit se izvodi nad tri DBA view-a:
  - **gv$statname**
  - **gv$sesstat**
  - **gv$session**
- Ako se želi koncentrirati na prave Oracle usere, tada se postavi dodatni uvjet u vanjski „where clause”

```
and b.username is not null
```

Tko generira logove trenutačno

```
col machine for a15
col username for a10
col redo_MB for 999G990 heading "Redo |Size MB"
column sid_serial for a13;

select b.inst_id,
       lpad((b.SID || ',' || lpad(b.serial#,5)),11) sid_serial,
       b.username,
       machine,
       b.osuser,
       b.status,
       a.redo_mb
from (select n.inst_id, sid,
             round(value/1024/1024) redo_mb
        from gv$statname n, gv$sesstat s
        where n.inst_id=s.inst_id
              and n.name = 'redo size'
              and s.statistic# = n.statistic#
        order by value desc
     ) a,
     gv$session b
where b.inst_id=a.inst_id
  and a.sid = b.sid
  -- and b.username is not null (pravi oracle korisnici)
  and rownum <= 30
;
```

Tko generira logove trenutačno

INST_ID	SID_SERIAL	USERNAME	MACHINE	OSUSER	STATUS	Size MB
1	788, 1		iis1	oracle	ACTIVE	2,073
4	788, 1		iis4	oracle	ACTIVE	1,928
1	792, 1		iis1	oracle	ACTIVE	1,168
1	791, 1		iis1	oracle	ACTIVE	1,149
3	788, 1		iis3	oracle	ACTIVE	1,111
4	792, 1		iis4	oracle	ACTIVE	1,092
1	785, 1		iis1	oracle	ACTIVE	1,064
4	791, 1		iis4	oracle	ACTIVE	1,064
3	792, 1		iis3	oracle	ACTIVE	757
3	791, 1		iis3	oracle	ACTIVE	738
3	785, 1		iis3	oracle	ACTIVE	436
4	785, 1		iis4	oracle	ACTIVE	411
1	764, 4	SYS	iis1	oracle	ACTIVE	340
1	737,61477	DBSNMP	iis1	oracle	ACTIVE	117
3	703,33361	DBSNMP	iis3	oracle	ACTIVE	113
4	677,30159	DBSNMP	iis4	oracle	ACTIVE	86
4	795, 1		iis4	oracle	ACTIVE	81
1	795, 1		iis1	oracle	ACTIVE	77
4	794, 1		iis4	oracle	ACTIVE	76
3	795, 1		iis3	oracle	ACTIVE	75
1	794, 1		iis1	oracle	ACTIVE	74
3	794, 1		iis3	oracle	ACTIVE	70
1	645, 5393	ANPI	USR\APINTARIC	apintaric	INACTIVE	50
3	758, 5		iis3	oracle	ACTIVE	39
3	755, 2		iis3	oracle	ACTIVE	24
3	754, 2		iis3	oracle	ACTIVE	22
3	756, 2		iis3	oracle	ACTIVE	21
3	757, 3		iis3	oracle	ACTIVE	21
4	774, 5835	JAGO	CLT\JGOLUZA	jgoluza	INACTIVE	10
1	619,61303	LIMI	NIO\LMIHALIC	lmihalic	INACTIVE	9

30 rows selected.

Kada nastaju redo log

- Ideja je vidjeti globalnu sliku
- Preduvjet da se krene na temeljniju analizu
- Zavisi od **CONTROL_FILE_RECORD_KEEP_TIME** parametra

```
SQL> show parameter CONTROL_FILE_RECORD_KEEP_TIME

NAME_COL_PLUS_SHOW_PARAM       TYPE        VALUE_COL_PLUS_SHOW_PARAM
------------------------------ ----------- --------------------
control_file_record_keep_time  integer     31
```

Kada nastaju redo log

```
set pagesize 120;
set linesize 200;
col day for a8;
spool rl.txt
PROMPT Archive log distribution per hours on each day ...

select
  to_char(first_time,'YY-MM-DD') day,
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'00',1,0)),'999') "00",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'01',1,0)),'999') "01",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'02',1,0)),'999') "02",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'03',1,0)),'999') "03",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'04',1,0)),'999') "04",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'05',1,0)),'999') "05",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'06',1,0)),'999') "06",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'07',1,0)),'999') "07",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'08',1,0)),'999') "08",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'09',1,0)),'999') "09",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'10',1,0)),'999') "10",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'11',1,0)),'999') "11",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'12',1,0)),'999') "12",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'13',1,0)),'999') "13",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'14',1,0)),'999') "14",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'15',1,0)),'999') "15",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'16',1,0)),'999') "16",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'17',1,0)),'999') "17",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'18',1,0)),'999') "18",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'19',1,0)),'999') "19",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'20',1,0)),'999') "20",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'21',1,0)),'999') "21",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'22',1,0)),'999') "22",
  to_char(sum(decode(substr(to_char(first_time,'HH24'),1,2),'23',1,0)),'999') "23",
  COUNT(*) TOT
from v$log_history
group by to_char(first_time,'YY-MM-DD')
order by day
;
```

Kada nastaju redo log

DAY	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	TOT
11-01-14	0	0	23	16	17	16	16	16	22	39	23	18	22	18	18	18	22	18	19	16	19	16	16	17	425
11-01-15	24	23	17	18	21	16	16	16	22	18	20	19	19	18	18	21	20	18	20	16	18	16	16	17	447
11-01-16	40	39	43	24	17	16	16	16	22	18	18	21	21	18	19	19	22	18	19	17	18	16	16	16	509
11-01-17	25	23	18	16	17	16	16	16	22	18	21	18	20	18	18	18	24	18	18	17	20	16	16	17	446
11-01-18	25	23	21	16	18	17	16	17	21	43	18	19	26	18	20	40	21	18	20	16	18	16	16	17	500
11-01-19	24	24	22	16	17	16	16	16	23	18	19	18	23	19	18	19	19	18	22	16	21	16	16	17	453
11-01-20	24	24	19	16	17	16	16	16	22	18	19	18	24	18	18	19	70	18	19	18	17	16	16	17	495
11-01-21	24	23	20	16	17	16	16	16	22	18	18	20	48	25	46	57	22	18	18	17	20	16	16	17	546
11-01-22	24	27	22	16	17	16	16	16	22	18	18	19	21	18	18	21	20	18	18	17	20	16	16	17	451
11-01-23	24	19	19	16	17	16	16	16	22	18	18	18	21	18	20	18	21	18	20	17	18	16	16	17	439
11-01-24	24	21	19	16	17	16	16	16	22	18	20	18	21	18	21	18	21	18	19	17	17	16	16	17	442
11-01-25	24	23	18	16	17	16	16	16	43	18	18	18	25	18	115	41	23	18	19	17	20	16	16	18	589
11-01-26	23	22	19	16	17	16	16	16	22	18	21	18	30	18	43	18	20	18	22	16	21	16	16	16	478
11-01-27	21	23	23	16	17	16	16	16	22	18	18	20	22	20	40	97	155	145	155	93	109	116	164	214	1556
11-01-28	93	24	23	16	17	16	16	16	29	18	19	19	30	101	19	18	75	30	19	16	20	16	22	16	688
11-01-29	21	16	16	16	17	16	27	20	22	18	18	18	31	18	18	19	23	18	19	18	18	16	16	17	456
11-01-30	24	22	17	16	17	16	16	16	22	18	18	19	23	18	19	18	23	18	18	17	20	16	16	17	444
11-01-31	24	19	21	16	17	16	16	16	23	18	19	18	22	19	18	20	20	18	18	16	21	16	16	17	444
11-02-01	24	25	22	16	17	16	16	16	23	18	18	20	24	18	18	19	21	18	18	17	19	16	16	17	452
11-02-02	24	22	20	16	17	16	16	16	23	18	18	40	23	18	39	18	21	18	22	16	21	16	16	17	491
11-02-03	24	23	22	16	18	16	16	16	22	18	19	18	24	18	20	18	20	18	19	18	19	16	16	17	451
11-02-04	24	23	23	16	17	16	16	16	22	18	19	19	48	18	20	20	23	18	18	17	19	16	16	17	479
11-02-05	24	23	17	17	17	16	16	16	22	18	20	19	23	18	19	20	22	18	18	17	20	16	16	17	449
11-02-06	24	23	17	16	18	16	16	16	22	18	19	19	21	18	19	18	24	18	20	16	21	16	16	17	448
11-02-07	24	22	18	16	18	16	16	16	22	18	52	18	44	25	29	24	21	18	18	16	19	16	16	17	519
11-02-08	24	23	19	16	17	16	16	16	22	18	20	19	28	18	19	18	22	18	19	17	34	32	32	31	514
11-02-09	36	39	35	31	37	31	32	32	32	34	34	34	39	34	34	34	36	34	38	32	37	31	32	33	821
11-02-10	37	36	38	31	33	32	31	32	37	34	34	34	41	34	35	35	37	34	34	31	34	32	31	33	820
11-02-11	39	35	38	31	33	32	31	32	37	34	34	34	38	34	34	34	35	34	34	32	34	32	31	32	814
11-02-12	40	34	34	33	35	32	31	32	37	34	34	34	36	34	35	34	36	34	34	31	36	32	31	32	815
11-02-13	40	34	35	31	34	32	31	32	37	34	34	34	38	34	34	34	37	34	34	32	36	32	31	32	816
11-02-14	40	34	33	31	35	32	32	32	37	66	34	35	38	34	34	34	37	34	34	31	36	32	31	33	849
11-02-15	48	50	50	44	48	32	32	32	35	34	34	34	38	34	37	34	39	34	34	32	36	31	32	33	887
11-02-16	39	34	33	32	33	31	32	32	37	53	37	34	38	38	34	2	0	0	0	0	0	0	0	0	539

34 rows selected. <-

Koliko je to u MB?

- Ukupna veličina red log (i pripadnog archived log) NE MOŽE biti izračunata iz prethodnog upita (jer se ne događa red log switch kada je redo log datoteka puna). Evo boljeg upita:

```
select sum(value)/1048576 redo_MB
from sys.gv_$sysstat
where name = 'redo size';

   REDO_MB
----------
1074623,75
```

Koliko je to u MB?

- Ako želite grupiranje po instancama (RAC) onda je upit:

```
select inst_id,
       sum(value)/1048576 redo_MB
  from sys.gv_$sysstat
 where name = 'redo size'
group by inst_id;
```

INST_ID	REDO_MB
1	370325.298
2	4712.567
3	294457.100
4	405129.283

Koji segmenti su generirali redo logove

- Nakon što smo našli točke interesa, u ovom slučaju 11-01-28 13:00-11-01-28 14:00 , korisno je saznati koji su to segmenti uzrokovali generiranje
- Za to trebamo imati i "dba_hist" tablice koje su dio "Oracle AWR (Automated Workload Repository). Period čuvanja podataka sa:
  - ```
    DBMS_WORKLOAD_REPOSITORY.
        MODIFY_SNAPSHOT_SETTINGS
            (43200, 15);
    ```
  - 43200=30 dana
  - Interval između dva snapshota 15 min
    ```

Koji segmenti su generirali redo logove

```
SELECT to_char(begin_interval_time,'YY-MM-DD HH24') snap_time,
       dhso.object_name,
       sum(db_block_changes_delta) BLOCK_CHANGED
  FROM dba_hist_seg_stat dhss,
       dba_hist_seg_stat_obj dhso,
       dba_hist_snapshot dhs
  WHERE dhs.snap_id = dhss.snap_id
    AND dhs.instance_number = dhss.instance_number
    AND dhss.obj# = dhso.obj#
    AND dhss.dataobj# = dhso.dataobj#
    AND begin_interval_time BETWEEN to_date('11-01-28 13:00','YY-MM-DD HH24:MI')
                                AND to_date('11-01-28 14:00','YY-MM-DD HH24:MI')
  GROUP BY to_char(begin_interval_time,'YY-MM-DD HH24'),
           dhso.object_name
  HAVING sum(db_block_changes_delta) > 0
ORDER BY sum(db_block_changes_delta) desc ;
```

Koji segmenti su generirali redo logove

SNAP_TIME	OBJECT_NAME	BLOCK_CHANGED
11-01-28 13	USR_RACUNI_MV	1410112
11-01-28 13	TROK_TAB_RESEAU_I	734592
11-01-28 13	TROK_VOIE_I	638496
11-01-28 13	TROK_DATUM_ULAZA_I	434688
11-01-28 13	TROK_PAIEMENT_I	428544
11-01-28 13	D_DPX_VP_RAD	351760
11-01-28 13	TROK_SVE_OK_I	161472
11-01-28 13	I_DATPBZ_S002	135296
11-01-28 13	IDS2_DATUM_I	129904
11-01-28 13	IDS2_PZNBR	129632
11-01-28 13	IDS2_IDS1_FK_I	128848
11-01-28 13	IDS2_DATTRAN_I	127440
11-01-28 13	IDS2_DATSOC_I	127152
11-01-28 13	IDS2_VRSTA_PROD_I	122816

<-

Koji SQL je uzrok pojedinog generiranja

- Sada kada znamo kada, koliko i što, vrijeme je da saznamo ZAŠTO.
- Kako je u prethodnom primjeru **`USR_RACUNI_MV`** za red veličine osvojio prvo mjesto, normalno je da je to prvo što dolazi kao predmet daljnjeg proučavanja.
- Kao mali dodatak treba naglasiti da upiti tipa SELECT nisu predmet proučavanja budući da ne kreiraju promjene

Koji SQL je uzrok pojedinog generiranja

```
SELECT to_char(begin_interval_time,'YYYY_MM_DD HH24') WHEN,
       dbms_lob.substr(sql_text,4000,1) SQL,
       dhss.instance_number INST_ID,
       dhss.sql_id,
       executions_delta exec_delta,
       rows_processed_delta rows_proc_delta
  FROM dba_hist_sqlstat dhss,
       dba_hist_snapshot dhs,
       dba_hist_sqltext dhst
  WHERE upper(dhst.sql_text) LIKE '%USR_RACUNI_MV%'
    AND ltrim(upper(dhst.sql_text)) NOT LIKE 'SELECT%'
    AND dhss.snap_id=dhs.snap_id
    AND dhss.instance_number=dhs.instance_number
    AND dhss.sql_id=dhst.sql_id
    AND begin_interval_time BETWEEN to_date('11-01-28 13:00','YY-MM-DD HH24:MI')
                                AND to_date('11-01-28 14:00','YY-MM-DD HH24:MI')
;
```

Koji SQL je uzrok pojedinog generiranja

```
WHEN          SQL                                                inst_id       sql_id  exec_delta  rows
------------- -------------------------------------------------- ------- ------------- ---------- ------
2011_01_28 13 DECLARE                                                 1 duwxbg5d1dw0q           0      0
                job BINARY_INTEGER := :job;
                next_date DATE := :mydate;
                broken BOOLEAN := FALSE;
              BEGIN
                dbms_refresh.refresh('"TAB"."USR_RACUNI_MV"');
                :mydate := next_date;
                IF broken THEN :b := 1;
                ELSE :b := 0;
                END IF;
              END;
2011_01_28 13 delete from "TAB"."USR_RACUNI_MV"                       1 5n375fxu0uv89           0      0
```

U oba slučaja bilo je nemoguće odrediti broj redova koji su mijenjani zbog same operacije koja je izvođena nad objektom tipa `materialized view`!

Koji SQL je uzrok pojedinog generiranja

- Na isti način možemo pogledati i neki drugi segment, recimo NC_TRANSACTION_OK_T Oracle tablicu:

```
WHEN          SQL                                                inst_id       sql_id  exec_delta rows_proc_delta
------------- -------------------------------------------------- ------- ------------- ---------- --------------
2011_01_28 13 alter table TAB.NC_TRANSACTION_OK_T                      4 g5gvacc8ngnb8          0              0
              shrink space cascade
```

- Primjer DDL-a koji generira redo log!

Koji SQL je uzrok pojedinog generiranja

- Ako ste orijentirani na broj promjena, onda možete izvesti upiti kojima **inst_id** i **sql_id** nemaju značenje (izbačeni iz upita). Evo modificiranog upita za **Z_PLACENO** segment (klasična Oracle tablica) :

```
SELECT when, sql, SUM(sx) executions, sum (sd) rows_processed
FROM (
      SELECT to_char(begin_interval_time,'YYYY_MM_DD HH24') when,
             dbms_lob.substr(sql_text,4000,1) sql,
             dhss.instance_number inst_id,
             dhss.sql_id,
             sum(executions_delta) exec_delta,
             sum(rows_processed_delta) rows_proc_delta
        FROM dba_hist_sqlstat dhss,
             dba_hist_snapshot dhs,
             dba_hist_sqltext dhst
        WHERE upper(dhst.sql_text) LIKE '%Z_PLACENO%'
          AND ltrim(upper(dhst.sql_text)) NOT LIKE 'SELECT%'
          AND dhss.snap_id=dhs.snap_id
          AND dhss.instance_Number=dhs.instance_number
          AND dhss.sql_id = dhst.sql_id
          AND begin_interval_time BETWEEN to_date('11-01-25 14:00','YY-MM-DD HH24:MI')
                                      AND to_date('11-01-25 15:00','YY-MM-DD HH24:MI')
        GROUP BY to_char(begin_interval_time,'YYYY_MM_DD HH24'),
             dbms_lob.substr(sql_text,4000,1),
             dhss.instance_number,
             dhss.sql_id
)
group by when, sql;
```

Koji SQL je uzrok pojedinog generiranja

```
WHEN          SQL                                                                    exec_delta rows_proc_delta
------------- ---------------------------------------------------------------------- ---------- ---------------
2011_01_25 14 DELETE FROM Z_PLACENO                                                           4         7250031
2011_01_25 14 INSERT INTO Z_PLACENO(OBP_ID,MT_SIFRA,A_TOT)                                    4         7250830
              SELECT P.OBP_ID,P.MT_SIFRA,SUM(P.OSNOVICA)
                FROM (SELECT OPI.OBP_ID,
                             OPO.MT_SIFRA,
                             SUM(OPO.IZNKN) OSNOVICA
                        WHERE OPI.OBP_ID = OPO.OPI_OBP_ID
                          AND OPI.RBR = OPO.OPI_RBR
                          AND NVL(OPI.S_PRETPOREZA,'O') IN ( 'O','N','A','Z','S')
                        GROUP BY OPI.OBP_ID,OPO.MT_SIFRA
                      )
```

SQL temeljen na segmentu

- Ukoliko ne želite fokus na period već želite početi istraživanje preko segmenata kao početne točke. Za takvo što uporabite slijedeće (**USR_RACUNI_MV** segment je *hard kodiran*)

```
SELECT to_char(begin_interval_time,'YY-MM-DD HH24') snap_time,
       sum(db_block_changes_delta)
  FROM dba_hist_seg_stat dhss,
       dba_hist_seg_stat_obj dhso,
       dba_hist_snapshot dhs
  WHERE dhs.snap_id = dhss.snap_id
    AND dhs.instance_number = dhss.instance_number
    AND dhss.obj# = dhso.obj#
    AND dhss.dataobj# = dhso.dataobj#
    AND dhso.object_name = 'USR_RACUNI_MV'
  GROUP BY to_char(begin_interval_time,'YY-MM-DD HH24')
  ORDER BY to_char(begin_interval_time,'YY-MM-DD HH24');
```

SQL temeljen na segmentu

```
SNAP_TIME    SUM(DB_BLOCK_CHANGES_DELTA)
-----------  ---------------------------

...
11-01-28 11                      1224240
11-01-28 12                       702880
11-01-28 13                      1410112
11-01-28 14                       806416
11-01-28 15                      2008912
...
```

- Naglašeni redak po rezultatu je identičan prije dobivenom rezultatu na drugi način.

Credits

- Oracle dokumentacija (10g):
  - http://download.oracle.com/docs/cd/B19306_01/server.102/b14231/onlineredo.htm
- Detaljan opis:
  - http://damir-vadas.blogspot.com/2011/02/how-to-redo-logs-generation.html
- Google
- Moje 11 godišnje iskustvo u Oracle-u

Q and A