



# Life of a transaction

Vit Spinka

# Agenda

- Look at data in redo log
- Simple transaction examples
- And with rollbacks
- And distributed transactions
- Q&A

# Vit Spinka

- Working with Oracle Database since *8i*
  - Oracle Certified Master
  - Principal developer of Dbvisit Replicate
  - ... which gets its data by parsing Oracle redo logs
- 
- @vitspinka
  - [vit.spinka@dbvisit.com](mailto:vit.spinka@dbvisit.com)
  - This presentation for download at <http://vitspinka.cz/download.html>



# Dbvisit



- HQ in New Zealand, US subsidiary, partners throughout the world
- Used in 80+ Countries
- Database Replication is our playground
- Worldwide leader in DR solutions for Oracle Standard Edition
- Product Engineers with “real world” DBA Experience
- Regular presenters at Oracle events such as OOW, Collaborate and NZOUG
- Passionate about Oracle Technology



Trusted in 80+ countries. . .

CDLQ ALDO ARUP Dbvisit THE SMART ALTERNATIVE Mercedes-Benz UK FUELS at&t Pythian love your data AGFA HealthCare IBM Foosle Roamware fiserv. Volkswagen Amazon web services Alcatel-Lucent Verizon NHS Engaging Networks Orion Health Imation

. . . By 800+ companies.

# Approach of this session

- Focus on redo logs
- We look at what Oracle stores on disk for the transactions
- We look at several common examples of transactions
  
- Oracle keeps more info about transactions than it stores in redo
- But redo is a good start to understand what is going on
- And usually more than enough
- (This is true for many other Oracle stuff)

# Changes in redo

- All persistent changes are written to redo
- Undo is written, too – Oracle may need it to recover after instance crash
- Let's look at different transactions
- To see what and where Oracle stores about them
- Oracle 12c multitenant used



# Transactions overview

- Simple heap table with one index (primary key)
- **DESCRIBE THE TABLE**
- Insert, then commit
- Add rollback, rollback to savepoint
- Kill session
- Abort instance and crash recovery
- Detached transaction
- Autonomous transaction
- Global transactions

# Case #1: insert, commit

- Start transaction (allocate undo header): 5.2 (undo tbs)
- Allocated in undo/rollback segment
- Undo for the insert: 5.1 (undo tbs)
- Redo for the insert: 11.2 (data tbs)
- (and again 5.1, 10.\* for index)
- Session information 5.19/5.20 (not in tbs)
- Commit (undo tbs)

CHANGE #1 **CON\_ID:1** TYP:0 **CLS:19 AFN:4** DBA:0x01000090 OBJ:4294967295

SCN:0x0000.0040fe2b SEQ:1 OP:5.2 ENC:0 RBL:0

ktudh redo: slt: 0x000b sqn: 0x00000d9f flg: 0x0052 siz: 136 fbi: 0

uba: 0x01000a4b.01b0.1e pxid: 0x0000.000.00000000

CHANGE #2 **CON\_ID:1** TYP:0 **CLS:20 AFN:4** DBA:0x01000a4b OBJ:4294967295

SCN:0x0000.0040fe2a SEQ:1 OP:5.1 ENC:0 RBL:0

ktudb redo: siz: 136 spc: 4168 flg: 0x0012 seq: 0x01b0 xid: 0x0002.00b.00000d9f

ktubl redo: slt: 11 rci: 0 opc: 11.1 [objn: 104325 objd: 104325 tsn: 3]

KDO Op code: DRP row dependencies Disabled

itli: 1 ispac: 0 maxfr: 4858

tabn: 0 slot: 0(0x0)

CHANGE #3 **CON\_ID:3** TYP:0 **CLS:1 AFN:10** DBA:0x028168ec OBJ:104325 SCN:0x0000.0040fe9b

SEQ:2 OP:11.2 ENC:0 RBL:0

KDO Op code: IRP row dependencies Disabled

itli: 1 ispac: 0 maxfr: 4858

tabn: 0 slot: 0(0x0) size/delt: 10

fb: --H-FL-- lb: 0x1 cc: 2

null: --

col 0: [ 1] 80

col 1: [ 4] 74 78 74 30

# Case #1: insert, commit

- Multiple ids indicates where the info is stored:
- CON\_ID: 12c multitenant. 1=CDB\$ROOT, that's the undo (and redo for internal stuff, too)
- CLS: block class. 1=data block. >16: undo. Odd is undo header, even is undo block. (Each undo segment gets it's own pair of ids.)
- AFN: absolute file number (see v\$datafile)
- DBA: data block address

CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20 ENC:0  
session number = 33  
serial number = 9  
transaction name =  
version 202375424  
audit sessionid 4294967295  
Client Id =  
login username = SYS

CHANGE #1 **CON\_ID:1** TYP:0 **CLS:19 AFN:4** DBA:0x01000090 OBJ:4294967295  
SCN:0x0000.0040fe9b SEQ:1 OP:5.4 ENC:0 RBL:0  
ktucm redo: slt: 0x000b sqn: 0x00000d9f srt: 0 sta: 9 flg: 0x2 ktucf redo: uba:  
0x01000a4b.01b0.1f ext: 4 spc: 3936 fbi: 0

# Case #2: insert, rollback

- Start transaction (allocate undo header): 5.2 (undo tbs)
- Undo for the insert: 5.1 (undo tbs)
- Redo for the insert: 11.2 (data tbs)
- (and again 5.1, 10.\* for index)
- Session information 5.19/5.20 (not in tbs)
- Rollback of the insert (and index) – 5.6, 5.11 (5.11 includes undo header)
- Rollback (undo tbs) – the transaction is empty now, so it's really just commit with a flag

CHANGE #1 CON\_ID:1 TYP:0 CLS:19 AFN:4 DBA:0x01000090 OBJ:4294967295  
SCN:0x0000.004026ee SEQ:1 OP:5.2 ENC:0 RBL:0

ktudh redo: slt: 0x0021 sqn: 0x00000d8f flg: 0x0052 siz: 136 fbi: 0  
uba: 0x01000e9e.01af.34 pxid: 0x0000.000.00000000

CHANGE #2 CON\_ID:1 TYP:0 CLS:20 AFN:4 DBA:0x01000e9e OBJ:4294967295  
SCN:0x0000.004026ed SEQ:1 OP:5.1 ENC:0 RBL:0

ktudb redo: siz: 136 spc: 1860 flg: 0x0012 seq: 0x01af xid: 0x0002.021.00000d8f  
ktubl redo: slt: 33 rci: 0 opc: 11.1 [objn: 104321 objd: 104321 tsn: 3]

KDO Op code: DRP row dependencies Disabled

itli: 1 ispac: 0 maxfr: 4858

tabn: 0 slot: 0(0x0)

CHANGE #3 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.0040271a  
SEQ:2 OP:11.2 ENC:0 RBL:0

KDO Op code: IRP row dependencies Disabled

itli: 1 ispac: 0 maxfr: 4858

tabn: 0 slot: 0(0x0) size/delt: 11

fb: --H-FL-- lb: 0x1 cc: 2

null: --

col 0: [ 2] c1 02

col 1: [ 4] 74 78 74 31

CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20 ENC:0

CHANGE #1 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168db OBJ:104322 SCN:0x0000.0040271a  
SEQ:1 OP:10.3 ENC:0 RBL:0

index redo (kdxlpu): purge leaf row (count = 2)

CHANGE #2 CON\_ID:1 TYP:0 CLS:20 AFN:4 DBA:0x01000e9e OBJ:4294967295  
SCN:0x0000.0040271a SEQ:2 OP:5.6 ENC:0 RBL:0

ktubu redo: slt: 33 rci: 52 opc: 10.22 objn: 104322 objd: 104322 tsn: 3

Undo type: Regular undo Undo type: User undo done

CHANGE #1 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.0040271a  
OP:11.3

KDO Op code: DRP row dependencies Disabled

CHANGE #2 CON\_ID:1 TYP:0 CLS:19 AFN:4 DBA:0x01000090 OBJ:4294967295  
SCN:0x0000.0040271a SEQ:1 OP:5.11 ENC:0 RBL:0

ktubu redo: slt: 33 rci: 0 opc: 11.1 objn: 104321 objd: 104321 tsn: 3

Undo type: Regular undo Undo type: User undo done **Begin trans**

CHANGE #1 CON\_ID:1 TYP:0 CLS:19 AFN:4 DBA:0x01000090 OBJ:4294967295  
SCN:0x0000.0040271e SEQ:1 OP:5.4 ENC:0 RBL:0

ktucm redo: slt: 0x021 sqn: 0x00000d8f srt: 0 sta: 9 flg: 0x4 **rolled back transaction**



# Case #3: insert, savepoint, insert, rollback to savepoint, commit

- Start transaction (allocate undo header): 5.2 (undo tbs)
- Session information 5.19/5.20 (not in tbs)\*
- Undo for the insert: 5.1 (undo tbs), Redo for the insert: 11.2 (data tbs)
- *nothing in the redo about the savepoint – that is not written to disk*
- Undo for the insert: 5.1 (undo tbs), Redo for the insert: 11.2 (data tbs)
- Rollback of the insert (and index) – 5.6, 5.11 (5.11 includes undo header)
- Commit (undo tbs)

CHANGE #1 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168db OBJ:104322 SCN:0x0000.00402721  
SEQ:2 **OP:10.3** ENC:0 RBL:0

CHANGE #2 CON\_ID:1 TYP:0 CLS:34 AFN:4 DBA:0x010006d2 OBJ:4294967295  
SCN:0x0000.00402721 SEQ:4 **OP:5.6** ENC:0 RBL:0

CHANGE #1 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.00402721  
SEQ:2 **OP:11.3** ENC:0 RBL:0

CHANGE #2 CON\_ID:1 TYP:0 CLS:34 AFN:4 DBA:0x010006d2 OBJ:4294967295  
SCN:0x0000.00402721 SEQ:5 **OP:5.6** ENC:0 RBL:0

End of rollback

CHANGE #1 CON\_ID:1 TYP:0 CLS:33 AFN:4 DBA:0x01000100 OBJ:4294967295  
SCN:0x0000.00402721 SEQ:1 OP:5.4 ENC:0 RBL:0

ktucm redo: slt: 0x0019 sgn: 0x00000d7f srt: 0 sta: 9 flg: 0x2 ktucf redo: uba:  
0x010006d2.0260.18 ext: 6 spc: 5146 fbi: 0

# Case #4: insert, killed session

- Start transaction (allocate undo header): 5.2 (undo tbs)
- Session information 5.19/5.20 (not in tbs)
- Undo for the insert: 5.1 (undo tbs), Redo for the insert: 11.2 (data tbs)
- *nothing in the redo about the kill*
- *it's really just a transaction rollback*
- Rollback of the insert (and index) – 5.6, 5.11
- Rollback transaction (undo tbs)
  
- Large transactions: delayed block cleanout (done during some other change)

CHANGE #1 CON\_ID:1 TYP:0 CLS:27 AFN:4 DBA:0x010000d0 OBJ:4294967295  
SCN:0x0000.004026c2 SEQ:1 OP:5.2 ENC:0 RBL:0

CHANGE #2 CON\_ID:1 TYP:1 CLS:28 AFN:4 DBA:0x010002e4 OBJ:4294967295  
SCN:0x0000.00402723 SEQ:1 OP:5.1 ENC:0 RBL:0

CHANGE #3 CON\_ID:3 TYP:2 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.00402722  
SEQ:1 OP:11.2 ENC:0 RBL:0

CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20

kill session

CHANGE #1 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.00402723  
SEQ:1 OP:11.3 ENC:0 RBL:0

KDO Op code: DRP row dependencies Disabled

CHANGE #2 CON\_ID:1 TYP:0 CLS:27 AFN:4 DBA:0x010000d0 OBJ:4294967295  
SCN:0x0000.00402723 SEQ:1 OP:5.11 ENC:0 RBL:0

CHANGE #1 CON\_ID:1 TYP:0 CLS:27 AFN:4 DBA:0x010000d0 OBJ:4294967295  
SCN:0x0000.00402729 SEQ:1 OP:5.4 ENC:0 RBL:0

ktucm redo: slt: 0x0008 sgn: 0x0000ca7 srt: 0 sta: 9 flg: 0x4

rolled back transaction

CHANGE #2 CON\_ID:3 TYP:2 CLS:1 AFN:8 DBA:0x0041b90a OBJ:8 SCN:0x0000.00402713 SEQ:1  
OP:11.5 ENC:0 RBL:0

KTB Redo

op: 0x11 ver: 0x01

compat bit: 4 (post-11) padding: 1

op: F xid: 0x0004.005.00000b4e uba: 0x010008d1.0244.20

**Block cleanout record, scn: 0x0000.00402717 ver: 0x01 opt: 0x02, entries follow...**

**itli: 2 flg: (opt=2 whr=1) scn: 0x0000.00402713**

KDO Op code: URP row dependencies Disabled

xtype: XA flags: 0x00000000 bdba: 0x0041b90a hdba: 0x004000c0

itli: 2 ispac: 0 maxfr: 4863

tabn: 2 slot: 13(0xd) flag: 0x6c lock: 2 ckix: 13

ncol: 16 nnew: 16 size: 0

col 0: [ 2] c1 06

col 1: [ 2] c1 09

col 2: [ 2] c1 02

col 3: [ 2] c1 09

col 4: [ 2] c1 02

col 5: [ 6] c5 16 30 31 25 2e

...

# Case #5: insert, instance abort

- Start transaction (allocate undo header): 5.2 (undo tbs)
- Session information 5.19/5.20 (not in tbs)\*
- Undo for the insert: 5.1 (undo tbs), Redo for the insert: 11.2 (data tbs)
- *no rollback of change*
- cleanup of the data block (block cleanout)
- cleanup of undo header
- (was the block written to disk at all?)
  
- Oracle does as little work as possible

CHANGE #1 CON\_ID:1 TYP:0 CLS:35 AFN:4 DBA:0x01000110 OBJ:4294967295  
SCN:0x0000.004026c7 SEQ:1 OP:5.2 ENC:0 RBL:0

CHANGE #2 CON\_ID:1 TYP:0 CLS:36 AFN:4 DBA:0x010017c2 OBJ:4294967295  
SCN:0x0000.004026c6 SEQ:7 OP:5.1 ENC:0 RBL:0

CHANGE #3 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.00402729  
SEQ:1 OP:11.2 ENC:0 RBL:0

CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.19

shutdown  
abort

CHANGE #1 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028168d4 OBJ:104321 SCN:0x0000.0040272d  
SEQ:1 OP:4.1 ENC:0 RBL:0

Block cleanout record, scn: 0x0000.00407877 ver: 0x01 opt: 0x01, entries follow...  
itli: 1 flg: (opt=2 whr=4) scn: 0x0000.00402722

CHANGE #1 CON\_ID:1 TYP:0 CLS:35 AFN:4 DBA:0x01000110 OBJ:4294967295  
SCN:0x0000.0040272d SEQ:1 OP:5.8 ENC:0 RBL:0

ktumr redo: slt: 1

# Case #6: detach a single transaction



- A single session can have multiple independent transactions opened at the same time
- `OCIDetachTran()`
- Really nothing special – looks like any two independent transactions
- Only `sid/serial#` show it's the same session



CHANGE #1 CON\_ID:1 TYP:0 CLS:35 AFN:4 DBA:0x01000110 OBJ:4294967295 SEQ:1 OP:5.2  
CHANGE #2 CON\_ID:1 TYP:0 CLS:36 AFN:4 DBA:0x010017a6 OBJ:4294967295 SEQ:5 OP:5.1  
CHANGE #3 CON\_ID:3 TYP:2 CLS:1 AFN:10 DBA:0x028046d6 OBJ:104486 SEQ:1 OP:11.5  
CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20 ENC:0  
session number = 272  
serial number = 871

CHANGE #1 CON\_ID:1 TYP:0 CLS:29 AFN:4 DBA:0x010000e0 OBJ:4294967295 SEQ:1 OP:5.2  
CHANGE #2 CON\_ID:1 TYP:0 CLS:30 AFN:4 DBA:0x0100039e OBJ:4294967295 SEQ:2 OP:5.1  
CHANGE #3 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028046d6 OBJ:104486 SEQ:1 OP:11.5  
CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20 ENC:0  
session number = 272  
serial number = 871

CHANGE #1 CON\_ID:1 TYP:0 CLS:35 AFN:4 DBA:0x01000110 OBJ:4294967295 SEQ:1 OP:5.4  
ktucm redo: slt: 0x001d sqn: 0x00000c85 srt: 0 sta: 9 flg: 0x2 ktucf redo: uba:  
0x010017a6.024d.06 ext: 3 spc: 7070 fbi: 0  
CHANGE #1 CON\_ID:1 TYP:0 CLS:29 AFN:4 DBA:0x010000e0 OBJ:4294967295 SEQ:1 OP:5.4  
ktucm redo: slt: 0x000c sqn: 0x00000b59 srt: 0 sta: 9 flg: 0x2 ktucf redo: uba:  
0x0100039e.01c4.0d ext: 2 spc: 5408 fbi: 0

# Case #7: autonomous transaction

- A single session can have multiple *nested* transactions opened at the same time
- Again, looks like normal transaction in the DB
- pxid is empty

REDO RECORD - Thread:1 RBA: 0x0000ef.0000001e.009c LEN: 0x01e8 VLD: 0x01 CON\_UID:  
2392015286

SCN: 0x0000.00446c2a SUBSCN: 10 08/16/2014 14:17:05

CHANGE #1 CON\_ID:1 TYP:0 CLS:27 AFN:4 DBA:0x010000d0 OBJ:4294967295

SCN:0x0000.00446bca SEQ:1 OP:5.2 ENC:0 RBL:0

ktudh redo: slt: 0x0007 sqn: 0x00000ce5 flg: 0x0052 siz: 136 fbi: 0

uba: 0x010002f1.028b.09 pxid: 0x0000.000.00000000

pdbid:2392015286

REDO RECORD - Thread:1 RBA: 0x0000ef.00000020.0010 LEN: 0x0218 VLD: 0x05 CON\_UID:  
2392015286

SCN: 0x0000.00446c2c SUBSCN: 1 08/16/2014 14:17:05

CHANGE #1 CON\_ID:1 TYP:0 CLS:33 AFN:4 DBA:0x01000100 OBJ:4294967295

SCN:0x0000.00446bf2 SEQ:1 OP:5.2 ENC:0 RBL:0

ktudh redo: slt: 0x000e sqn: 0x00000db3 flg: 0x0052 siz: 136 fbi: 0

uba: 0x01000aeb.026a.29 pxid: 0x0000.000.00000000

pdbid:2392015286

(rest omitted for brevity)

# Case #8: distributed transaction

- Two sessions work on branches of the same transaction
- Two-phase commit
- Shows as single transaction in database, but with two undo headers, two session information
- Usually application adds more transaction info, that's stored, too (external name, internal name)

CHANGE #1 CON\_ID:1 TYP:0 CLS:25 AFN:4 DBA:0x010000c0 OBJ:4294967295 SEQ:1 OP:5.2  
ktudh redo: slt: 0x000c sqn: 0x00000d24 flg: 0x0052 siz: 288 fbi: 0  
uba: 0x0100083e.0244.34 pxid: 0x0000.000.00000000

CHANGE #2 CON\_ID:1 TYP:0 CLS:26 AFN:4 DBA:0x0100083e OBJ:4294967295 SEQ:4 OP:5.1  
ktudb redo: siz: 288 spc: 400 flg: 0x1012 seq: 0x0244 rec: 0x34  
xid: 0x0005.00c.00000d24

CHANGE #3 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028046d6 OBJ:104486 SEQ:1 OP:11.5

CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20 ENC:0  
session number = 44, serial number = 2211

CHANGE #1 CON\_ID:1 TYP:0 CLS:25 AFN:4 DBA:0x010000c0 OBJ:4294967295 SEQ:1 OP:5.2  
ktudh redo: slt: 0x000c sqn: 0x00000000 flg: 0x000a siz: 132 fbi: 0  
uba: 0x0100083f.0244.01 pxid: 0x0000.000.00000000

CHANGE #2 CON\_ID:1 TYP:1 CLS:26 AFN:4 DBA:0x0100083f OBJ:4294967295 SEQ:1 OP:5.1  
ktudb redo: siz: 132 spc: 110 flg: 0x100a seq: 0x0244 rec: 0x01  
xid: 0x0005.00c.00000d24

CHANGE #3 CON\_ID:3 TYP:0 CLS:1 AFN:10 DBA:0x028046d6 OBJ:104486 SEQ:1 OP:11.5

CHANGE #4 MEDIA RECOVERY MARKER CON\_ID:3 SCN:0x0000.00000000 SEQ:0 OP:5.20 ENC:0  
session number = 264, serial number = 2957

CHANGE #1 CON\_ID:1 TYP:0 CLS:26 AFN:4 DBA:0x0100083f OBJ:4294967295  
SCN:0x0000.0044f63b SEQ:1 OP:5.1 ENC:0 RBL:0  
ktudb redo: siz: 292 spc: 7930 flg: 0x1024 seq: 0x0244 rec: 0x03  
xid: 0x0005.00c.00000d24

Dumping kcocv element #2 size 256:

7F744FE8B060 01000100 00000000 00000904 6D656401 [.....dem]  
7F744FE8B070 6E78746F 6D656420 0000326F 00000000 [otxn demo2.....]

CHANGE #1 CON\_ID:1 TYP:0 CLS:25 AFN:4 DBA:0x010000c0 OBJ:4294967295  
SCN:0x0000.0044f63c SEQ:1 OP:5.12 ENC:0 RBL:0  
ktust redo: slt: 12 sgn: 0x00000d24 sta: 2 cfl: 0x22

CHANGE #1 CON\_ID:1 TYP:0 CLS:25 AFN:4 DBA:0x010000c0 OBJ:4294967295  
SCN:0x0000.0044f643 SEQ:1 OP:5.4 ENC:0 RBL:0  
ktucm redo: slt: 0x000c sgn: 0x00000d24 srt: 0 sta: 3 flg: 0x0

# Case #9: discrete transaction

- This is really Oracle 7 (1993), then happily forgotten
- `alter system set "_discrete_transactions_enabled"=TRUE scope=spfile;`
- `dbms_transaction.begin_discrete_transaction;`
  
- Does nothing today... the redo looks the same, even the restrictions for discrete transaction don't apply anymore

QA





THE SMART ALTERNATIVE

Twitter: @dbvisit

Blog: [blog.dbvisit.com](http://blog.dbvisit.com)

Forum: [www.dbvisit.com/forums](http://www.dbvisit.com/forums)