



Right Way for Losing Data

- **mag. Sergej Rožman; Abakus plus d.o.o.**
- The latest version of this document is available at:
<http://www.abakus.si/>



Right Way for Losing Data

mag. Sergej Rožman

sergej.rozman@abakus.si

19. Strokovno srečanje SIOUG



SIOUG 2014

13.-14. oktober 2014

ORACLE Gold Partner



Mestna občina Ljubljana



MESTNA OBČINA KOPER
COMUNE CITTA DI CAPODISTRIA



Aerodrom Ljubljana



REPUBLIKA SLOVENIJA
MINISTRSTVO ZA FINANCE



Mercator

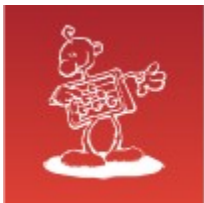


Iskra
IskraSistemi



BANKA
SLOVENIJE
EVROSISTEM





Abakus plus d.o.o.

ORACLE Gold Partner

History

- from 1992, ~20 employees

Applications:

- special (DB – Newspaper Distribution, FIS – Flight Information System)
- **ARBITER – the ultimate tool in audit trailing**
- **APPM - Abakus Plus Performance Monitoring Tool**

Services:

- DBA, OS administration , programming (MediaWiki, Oracle)
- networks (services, VPN, QoS, security)
- open source, monitoring (Nagios, OCS, Wiki)

Hardware:

- servers, **SAN storage**, firewalls, **backup servers**

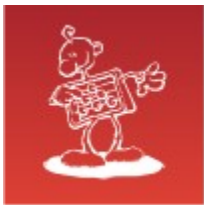
Infrastructure:

- from 1995 GNU/Linux (**~20 years of experience !**)
- Oracle on GNU/Linux: since RDBMS 7.1.5 & Forms 3.0 (**before Oracle !**)
- **>20 years of experience with High-Availability !**



Mestna občina Ljubljana





Once Upon a Time ...

1





Once Upon a Time ...

- 4 GB common disk drive/file system size limit
- OracleDB (ver. 7 & 8) had »a weak control« of the datafile size limit (at start time only)
- autoextendable datafile could exceed over the »weak limit«
- the database worked until shutdown
- after that ...





Once Upon a Time ...

```
SQL> startup
```

```
ORACLE instance started.
```

```
Total System Global Area 130248864 bytes
Fixed Size                  73888 bytes
Variable Size              123449344 bytes
Database Buffers           6553600 bytes
Redo Buffers                172032 bytes
Database mounted.
```

```
(Error stack:)*
```

```
ORA-01122: database file <datafile-name> failed verification check
ORA-19510: failed to set size of 524288 blocks for file
           <datafile-name> (blocksize=8192)
ORA-27059: skgfrsz: could not reduce file size
ORA-27072: skgfdisp: I/O error
```

* The fact that human memory is not like a DVD does mean that memories may not be completely accurate






Once Upon a Time ...

»Solution«

- hack the datafile size and open the DB
- BE AFRAID of STALE PENDING OPERATION
- create new DB and transfer data from old DB

- limit max datafile sizes





How to move OMF in ASM

2





OMF – Oracle Managed Files

- introduced in Oracle 9i
- eliminate the need to directly manage the operating system files (really?)
- introduces »funny« names
(/oradata/ARBITER/datafile/o1_mf_ts_test_b21jyq0c_.dbf)

- important parameters

```
db_create_file_dest  
db_create_online_log_dest_1
```

```
string /oradata  
string /oradata
```





ASM – Automatic Storage Management

- introduced in Oracle 10i
- Oracle volume manager
- introduces even more »funny« names
(+DATA/arbiter/datafile/sysaux.674.847463079)





How to move/rename datafile

Usual way

- take the datafile offline
- move the datafile at the OS level
- rename the datafile in the database
- recover the datafile and put it back online





How to rename OMF

- YOU SHOULD NOT !
- it doesn't make sense





How to move OMF

- it should be possible
- OMF: maybe Oracle will do the OS part by himself?





How to move OMF

- Oracle won't do anything by himself !
- you move the OMF the same way as ordinary datafile





How to move OMF in ASM

OMF: »OS part« can't be done in ASM:

```
SQL> ALTER DISKGROUP DATA RENAME ALIAS  
'+DATA/test/datafile/users.271.858525959' TO  
'+DATA/ttt/datafile/users.271.858525959';
```

ERROR at line 1:

ORA-15032: not all alterations performed

ORA-15177: cannot operate on system aliases





How to move OMF in ASM

DB part – Oracle ≥10.2:

```
SQL> alter database rename file  
'+DATA/test/datafile/users.271.858525959' to  
'+DATA/ttt/datafile/users.271.858525959';
```

ERROR at line 1:

ORA-01511: error in renaming log/data files

ORA-01141: error renaming data file 4 - new file

'+DATA/ttt/datafile/users.271.858525959' not found

ORA-01110: data file 4: '+DATA/test/datafile/users.271.858525959'

ORA-17503: ksfdopn:2 Failed to open file

+DATA/ttt/datafile/users.271.858525959

ORA-15173: entry 'users.271.858525959' does not exist in directory
'datafile'





How to move OMF in ASM

but ...

Oracle 10.1:

```
SQL> alter database rename file  
'+ARCH/test/datafile/users.7515.858523367' to  
'+ARCH/ttt/datafile/users.7515.858523367';
```

Database altered.





How to move OMF in ASM

Oracle 10.1:

```
SQL> alter database rename file  
'+ARCH/test/datafile/users.7515.858523367' to  
'+ARCH/ttt/datafile/users.7515.858523367';
```

Database altered.

```
SQL> alter database open;
```

ERROR at line 1:

```
ORA-01157: cannot identify/lock data file 4 - see DBWR trace file  
ORA-01110: data file 4: '+ARCH/ttt/datafile/users.7515.858523367'
```





How to move OMF in ASM

Oracle 10.1:

```
SQL> alter database rename file  
'+ARCH/test/datafile/users.7515.858523367' to  
'+ARCH/tnt/datafile/users.7515.858523367';
```

Database altered.


```
SQL> alter database open;
```

ERROR at line 1:

```
ORA-01157: cannot identify/lock data file 4 - see DBWR trace file  
ORA-01110: data file 4: '+ARCH/tnt/datafile/users.7515.858523367'
```

Oracle! Where is my datafile???





Mixed version issue - part 1

3





ASM

- Oracle ASM used raw devices (or ASMLib) until ver. 10.2.0.2
- On Linux raw devices were deprecated and scheduled for removal at one point, because the O_DIRECT flag can be used instead
- Remember that ASM does not perform database IO as the database still performs its own IO just as always





ASM – what if ...

Mixed versions (supported configuration)

- Oracle ASM ver. 11.2
- Oracle DB ver. 10.1
- **Caution:** ASM supports block devices but DB does not !
- Remember that ASM does not perform database IO as the database still performs its own IO just as always





ASM – what if ...

```
SQL> alter diskgroup DATA add  
      failgroup SAN1 disk /oradata/+ASM/san1-disk4;
```

Diskgroup altered.


ASM alertlog:

```
WARNING: RDBMS client (XX) with version (10.1.0.5.0) is not compatible  
with a pending operation
```

DB alertlog:

```
ORA-00376: file 5 cannot be read at this time  
ORA-01110: data file 5: '+DATA/test/datafile/undotbs2.269.696087027'
```





Mixed version issue - part 2

4





ASM – what if ... (again)

Mixed versions (supported configuration)

- Oracle ASM ver. 12.1
- Oracle DB ver. 11.2
- **Caution:**
ASM supports large disk drive (>2TB) devices but DB does not !
- Remember that ASM does not perform database IO as the database still performs its own IO just as always





ASM – what if ... (again)

Mixed versions (legal & supported condition)

- works fine until data reaches 2 TB limit on the individual disk drive
- even for quite a long time (weeks, months)





ASM – what if ...

ASM alertlog:

```
Errors in file /oradmin/diag/asm/+asm/+ASM2/trace/+ASM2_arb0_9392.trc:  
ORA-15196: invalid ASM block header [kfc.c:28346] [endian_kfbh][2147483655]  
[243] [83 != 1]  
ORA-15196: invalid ASM block header [kfc.c:28346] [endian_kfbh] [2147483655]  
[243] [83 != 1]  
NOTE: cache repaired a corrupt block: group=3(DATA) dsk=7 blk=243 on disk 7  
from disk=7 (DATA_0007) incarn=3491241260 au=11 blk=243 count=1
```

DB alertlog:

```
Reading datafile '+DATA/arb/datafile/users.265.841516141' for corruption at  
rdba: 0x3481c1bf (file 210, block 115135)  
Read datafile mirror 'DATA_0001' (file 210, block 115135) found same corrupt  
data (no logical check)  
Read datafile mirror 'DATA_0006' (file 210, block 115135) found valid data  
  
Read of datafile '+DATA/arb/datafile/ts_prod.268.841517003' (fno 288) header  
failed with ORA-01208  
Rereading datafile 288 header from mirror side 'DATA_0002' failed with  
ORA-01208  
Rereading datafile 288 header from mirror side 'DATA_0006' found valid data  
Repaired corruption in datafile 288 header
```





ASM – what if ...

Recover:


Output:

```
ORA-00600: internal error code, arguments: [3020], [210],  
[692130], [881495970], [], [], [], [], [], [], [], []  
ORA-10567: Redo is inconsistent with data block (file# 210,  
block# 692130, file offset is 1374961664 bytes)  
ORA-10564: tablespace USERS  
ORA-01110: data file 210:  
'/oradata/ARB/datafile/users.11936.803735271'  
ORA-10561: block type 'TRANSACTION MANAGED DATA BLOCK', data  
object# 3744293
```

DB alertlog:

```
WARNING! Recovering data file 395 from a fuzzy backup. It  
might be an online backup taken without entering the begin  
backup command.
```





End of destruction

Begin of construction





Recover

Some interesting readings

- [1] Terry Sutton; Database Specialists, Inc.;
Recovering an Oracle Database with Missing Archived Logs
(http://www.dbspecialists.com/files/presentations/missing_logs.html)
- [2] Vladimir's blog;
Some things that I've found to recover database without redo
(<http://myoracledbablogon.blogspot.com/2010/09/all-that-ive-found-to-recover-database.html>)
- [3] Doc ID: Note:30681.1;
EVENT: ADJUST_SCN - Quick Reference





Recover

General instructions from [1]

1. **Taking a cold backup** of what you have now.
2. **Restoring the lost datafile from a backup** and applying the archived redo logs that you do have.
3. Setting an undocumented instance parameter which will allow you to open the database in its current state (`_allow_resetlogs_corruption=TRUE`).
4. Doing exports and selects to retrieve what data you can from the problem tablespace.
5. **Restoring the entire database from the cold backup taken earlier.**
6. Taking the damaged datafile offline.
7. Doing exports and selects to retrieve additional data not salvaged in step 4.
8. **Restoring again from the cold backup.**
9. Dropping the problem tablespace.
10. Recreating the problem tablespace.
11. Rebuilding the data in the problem tablespace with the data extracted in steps 4 and 7.





Backup Server



As na disku.




Backup and Recovery Best Practices

- Backup takes no time!
no resources needed & no disk space;
- Recover takes no time as well!
no resources needed;
- Copies are without errors and consistent;
- Data is always available & always in view.





DB Backup in »no time«





Historical copies occupy (almost) no disk space

... by using data deduplication

- snapshot
- save snapshot **to deduplicated area**

Example:

- DB size 1 TB
- 1% changed/added data per day (~10 GB)
- ~200 days backup fits on 3 TB disk drive






Recovery in »no time« – scenario 1

Recovery as a Service (RaaS)

- services are offered directly from the backup server

**BACK IN BUSINESS
IN NO TIME!***

* real restore in more appropriate time






Alternative use

- BI – business analysis purposes
- reporting
- development & test
- **testing of recovery procedures**



»activated backup DB«






Recovery in »no time« – scenario 2

Recovery as an Infrastructure

(change role to SAN storage)


- backup server works as a SAN storage infrastructure to production servers



**BACK IN BUSINESS
IN NO TIME!***

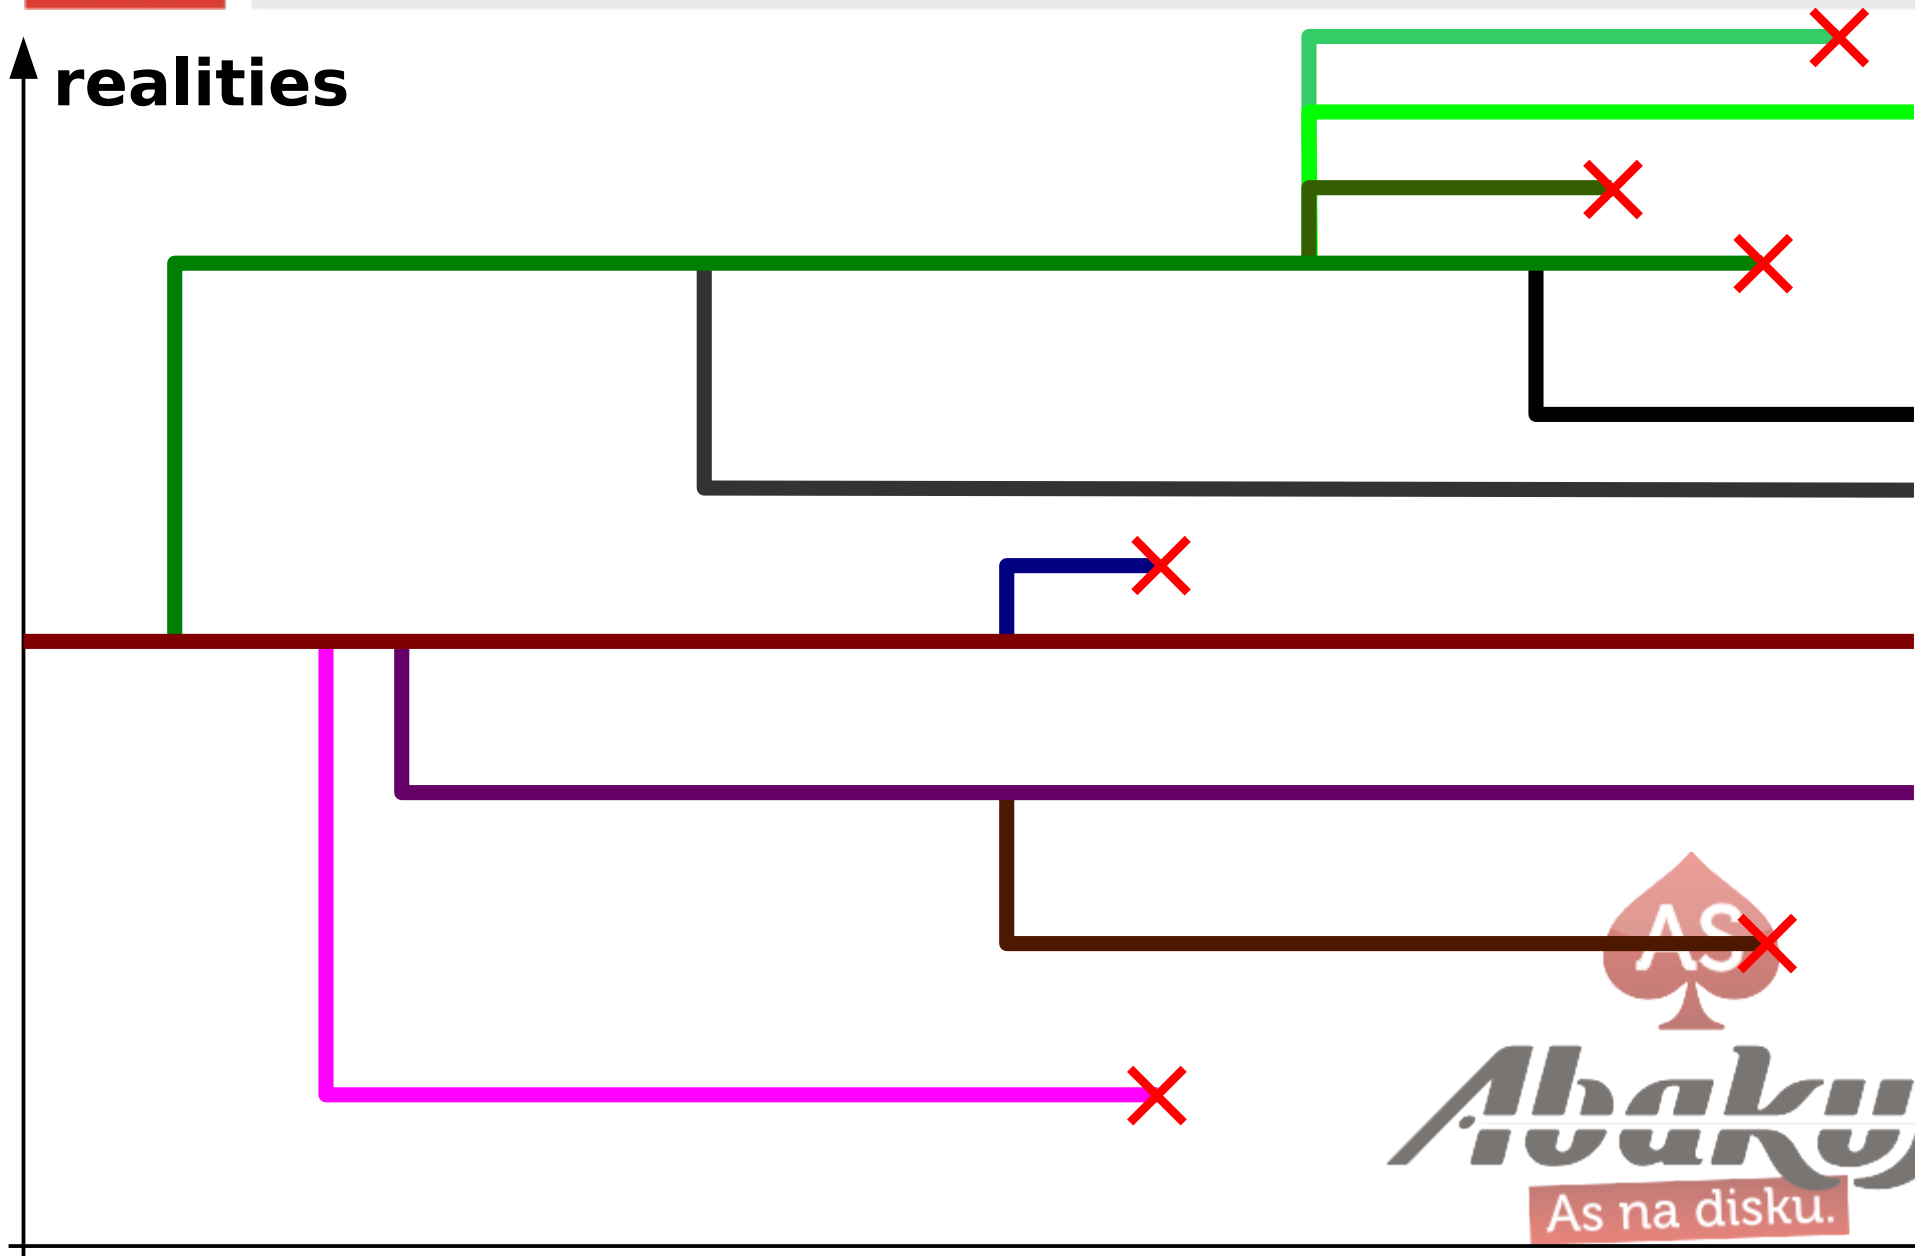
* real restore in more appropriate time





Abakus Time Warp

↑ realities



time



Invitation


Ask my colleague to help recreating your data from scratch



- **SIOUG 2014**
- hall C; 15:00

- **Jure Kajzer, ABAKUS plus d.o.o.**
- **Forensic analysis of Oracle log files**





Right Way for Losing Data

Thank You

mag. Sergej Rožman

ABAKUS plus d.o.o.

Ljubljanska c. 24a, Kranj

e-mail: sergej.rozman@abakus.si

phone: +386 4 287 11 14

