# Lose your data
## without ever noticing it

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

# Rubber Duck Debugging

https://en.wikipedia.org/wiki/Rubber_duck_debugging

»The name Rubber Duck Debugging is a reference to a story in the book The Pragmatic Programmer in which a programmer would carry around a rubber duck and debug their code by forcing themselves to explain it, line-by-line, to the duck.«

Abakus
As na disku.

# Lose your data
## without ever noticing it

**mag. Sergej Rožman**

sergej.rozman@abakus.si

Make IT
2019

ORACLE Gold Partner

Hidria GOODYEAR DUNLOP SAVA TIRES Mercator Iskra MERKUR

Mestna občina Ljubljana

MESTNA OBČINA KOPER
COMUNE CITTÀ DI CAPODISTRIA

REPUBLIKA SLOVENIJA
MINISTRSTVO ZA OBRAMBO

ZAVOD ZA ŠPORT RS PLANICA

KONTROLA ZRAČNEGA PROMETA SLOVENIJE

# Abakus plus d.o.o.

ORACLE Gold Partner

## History

- from 1992, ~20 employees

## Applications:

- **DejaVu** – virtual DBs in real-time
- **APPM** - Abakus Plus Performance Monitoring tool
- **ARBITER** – the ultimate tool in audit trailing

## Services:

- DBA, OS administration , programming
- networks (services, VPN, QoS, security)

## Hardware:

- servers, **CEPH clusters**, firewalls, **backup servers**

## Experience:

- from 1995 GNU/Linux ***(~25 years of experience !)***
- Oracle on GNU/Linux: since RDBMS 7.1.5 & Forms 3.0 ***(before Oracle !)***
- **~30 years of experience with High-Availability !**

BANKA SLOVENIJE · Gorenjska Banka · LON · SKB SOCIETE GENERALE GROUP · GENERALI Zavarovalnica · NLB Vita Življenjska zavarovalnica

# Losing Data – Reloaded

Some old stuff

- SIOUG 2014: Right Way for Losing Data

# 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

## 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'
```

## Oracle! Where is my datafile???

# ASM – what if ...

## Mixed versions (supported configuration)

- Oracle ASM ver. 19.4
- 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 /dev/sdd;

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'
```

- Oracle should probably force using raw devices when DATABASE_COMPATIBILITY <= 10.2.0.2

# ASM – what if ... (again)

## Mixed versions (supported configuration)

- Oracle ASM ver. 19.4
- 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 ...

## Recovery:

### 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.
```

- Oracle should probably force limiting disk drives to 2TB when DATABASE_COMPATIBILITY <= 11.2

# Database Lost Write Corruption

## A data block lost write:

- I/O subsystem acknowledges the completion of the block write.
- While in fact the write did not occur in the persistent storage.

# Lost Write Causes

- Faulty disk & controller
  (Sergej Rožman, SSD & Oracle @ SIOUG 2017)
- Faulty memory
- Faulty network components
- Volume manager
- NFS
- Other reasons
  (bug in some code, **system crash, power outage**)

# Lost Write Detection

Common tools usually do not detect lost writes

- DBVerify finds no corruptions
- RMAN> BACKUP VALIDATE DATABASE
  shows no errors
- ANALYZE command finds nothing
- DB_BLOCK_CHECKING parameter is at no use
- No errors in logs.

An old version of the block remains and is perfectly correct for DBV, RMAN, and checksum.

# Special Case: Fractured Block

**»Partialy« lost write**
one part of the block was stored and the other part is missing

- Leads to a block corruption – detectable corruption with errors in logs (ORA-600, ...)
- Database restored from wrongly invoked hot backup.
- Database activated from a snapshot.

Not a subject of this presentation.

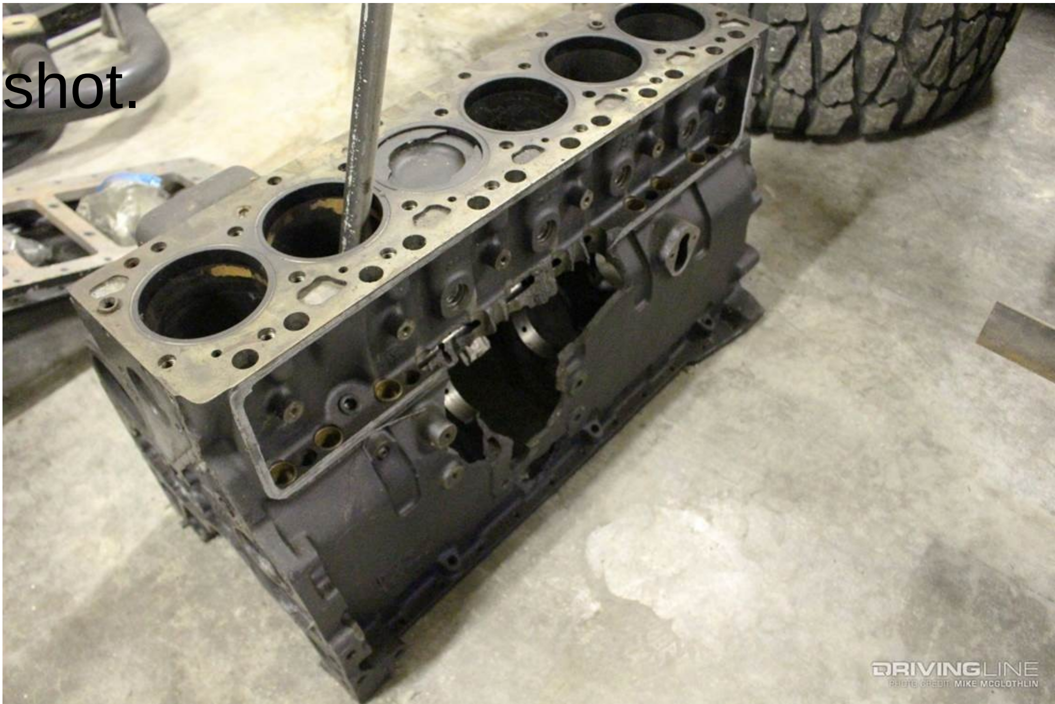

PHOTO CREDIT: MIKE MCGLOTHLIN

# Backup is Useless

- Lost write is usually not detected on time.
- All backup copies contain corrupted data as well.
- An error can show up months after lost write event or not at all.

Backup is useless (AGAIN?)

# Lost Write Detection

- The database is logically corrupted – errors in logs (index entry for nonexistent table row or vice versa, broken foreign key relation, …)
- The data is evidently incorrect
  (a million on your bank account is missing)
- **We have something to compare with!**
- … or not at all?!

# Lost Write Protection?

- Database lost write detection functionality introduced in Oracle 11.1 – parameter DB_LOST_WRITE_PROTECT
- Version 12.2 – DBCOMP is PL/SQL package. EE only! Deprecates the parameter DB_LOST_WRITE_PROTECT.

Syntax:

```
DBMS_DBCOMP.DBCOMP (
    datafile IN varchar2,
    outputfile IN varchar2,
    block_dump IN boolean);
```

reference: https://community.oracle.com/docs/DOC-1023009

# We have something to compare with

- Lost write detection is based on comparing the database with physical stand-by database (DataGuard)
- Lost write detection **on write operations** are performed even without DB_LOST_WRITE_PROTECT

```
ORA-00600: internal error code, arguments: [3020], [2], [129], [8388737], [], [], [], [], [], [], [], []
ORA-10567: Redo is inconsistent with data block (file# 2, block# 129, file offset is 1056768 bytes)
```

- With DB_LOST_WRITE_PROTECT detections are performed **on read operations** as well

```
ORA-00752: recovery detected a lost write of a data block
ORA-10567: Redo is inconsistent with data block (file# 2, block# 129, file offset is 1056768 bytes)
```

# Which block is correct?

Options

- Primary data block is correct, stand-by block is lost.
- Stand-by data block is correct, primary block is lost.
- Data block with highest SCN is correct.
- Both blocks are wrong

Recovery of a lost block is a chalange.
Solution may vary from case to case.

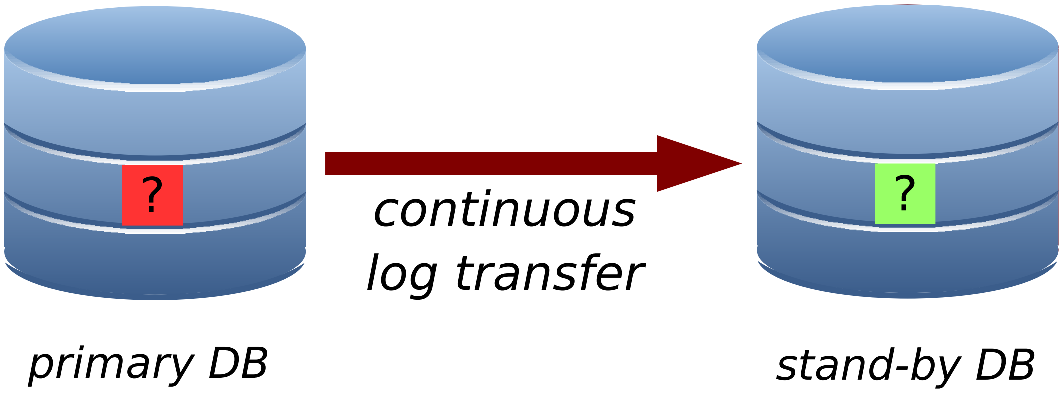

# Abakus Backup Server

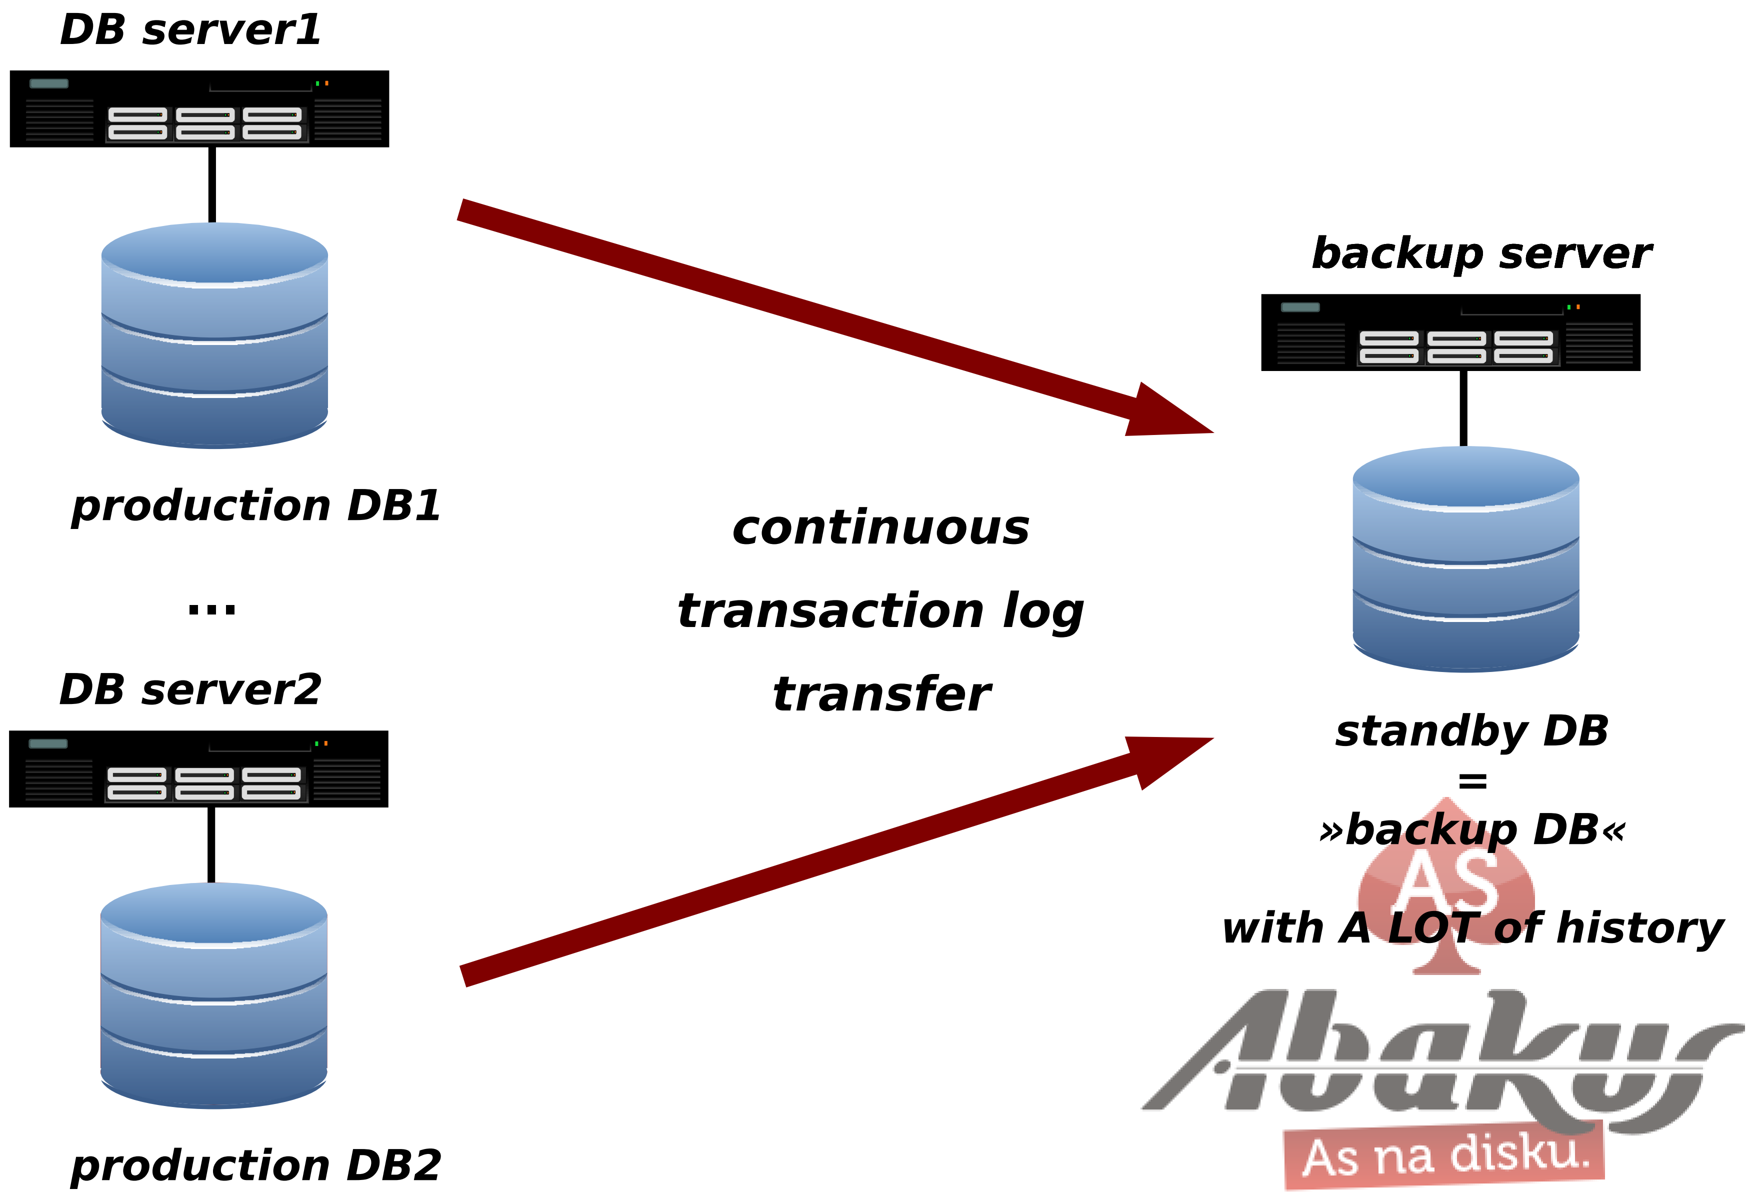

DB server1

production DB1

...

DB server2

production DB2

continuous

transaction log

transfer

backup server

standby DB

=

»backup DB«

with A LOT of history

# Historical copies

**… 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

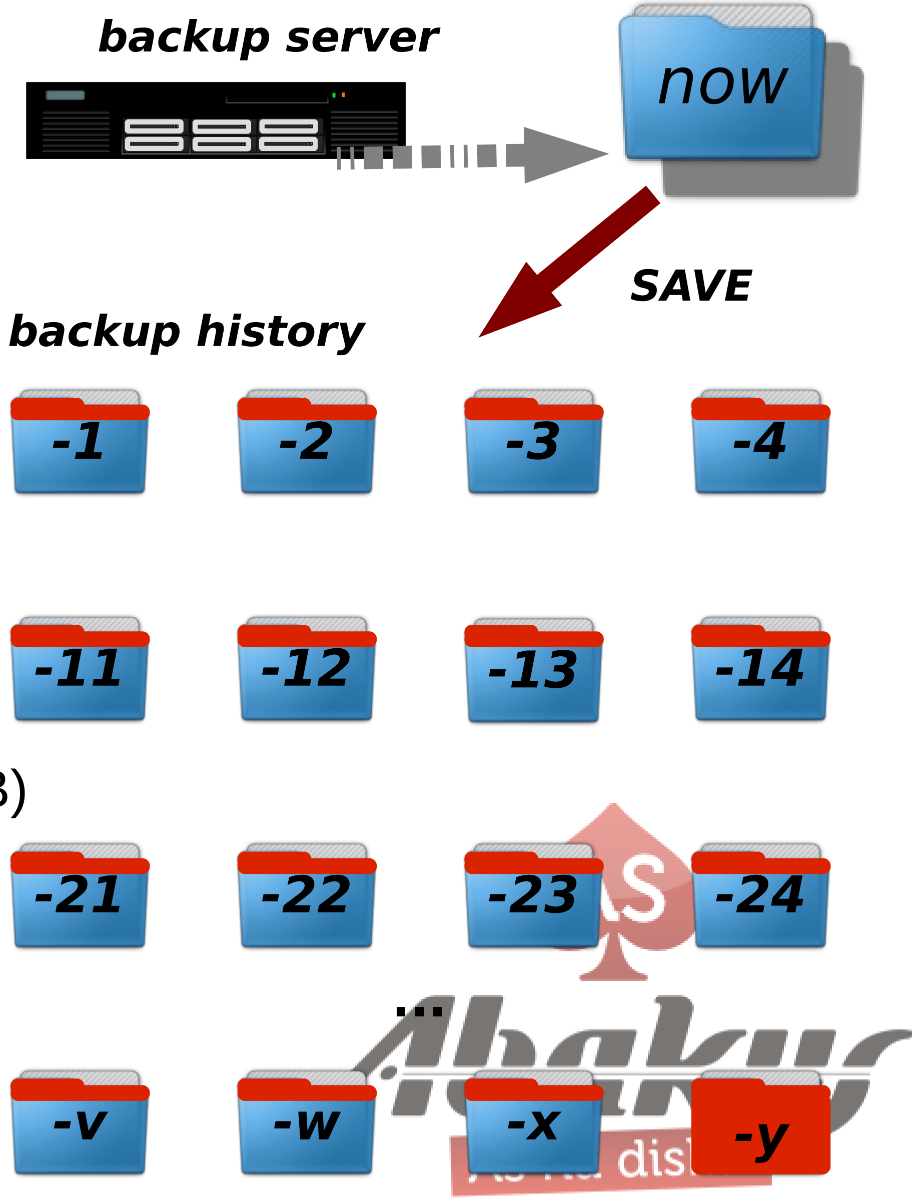

# 18c New Lost Write Protection?

- New feature: comparing inside the database

```
SQL> create bigfile tablespace TS_LOST_WRITE_SHADOW datafile
        size 10M lost write protection;
SQL> alter tablespace users enable lost write protection;
```

- Detection

```
SQL> SELECT * FROM test_table;

ERROR at line 1:
ORA-65478: shadow lost write protection - found lost write
```

- Enterprise Edition feature!

# Conclusion

- Lost writes do happen.
- For lost writes detection use physical stand-by database. Other types of replication (SAN, mirror) do not comply.
- Common tools do not detect lost writes.
- Recovery of a lost block is a chalange.

# Lose your data
## without ever noticing it

# 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

NOVA BANKA MILENIJUM OSIGURANJE CENTROSINERGIJA ANDRITZ futuraplus