



Perfect Recall - Lossless Backup

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





BACKUP PLANS

Sometimes they're so good you want everything else to fail.

Abakus

As na disku.

Perfect Recall (Lossless Backup)

mag. **Sergej Rožman**

sergej.rozman@abakus.si

21. Strokovno srečanje SIOUG



SIOUG 2016

17. - 18. oktober 2016

ORACLE Gold Partner



Mestna občina Ljubljana



MESTNA OBČINA KOPER
COMUNE CITTA DI CAPODISTRIA



Aerodrom Ljubljana



REPUBLIKA SLOVENIJA
MINISTRSTVO ZA FINANCE



Mercator



Iskra

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)

Infrastructure:

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


Skills & Experience:

- from 1995 GNU/Linux (**~20 years of experience !**)
- Oracle on GNU/Linux: since RDBMS 7.1.5 & Forms 3.0 (**before Oracle !**)
- **>25 years of experience with High-Availability !**



Mestna občina Ljubljana





Customers





Goal

BACKUP POOL

production database



most recent copy



historical copies



...



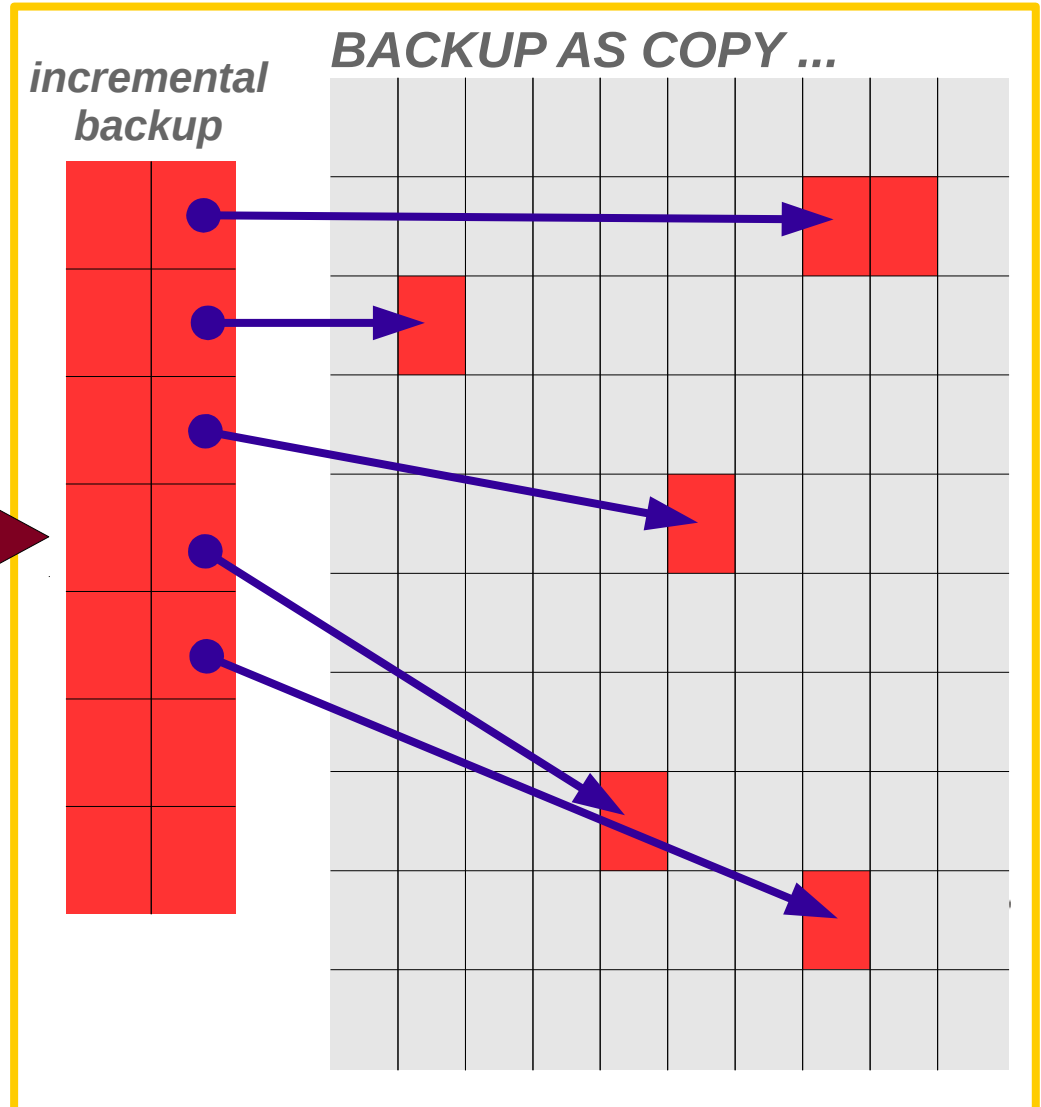



Backup - Traditional Way

Incrementally Updated Backups


- RECOVER COPY OF DATABASE ... **BACKUP POOL**

production database





Backup Full/Incremental



- somewhat optimized
no full backup except initial
- incremental backup optimized
with Oracle Enterprise Edition
(block change tracking)





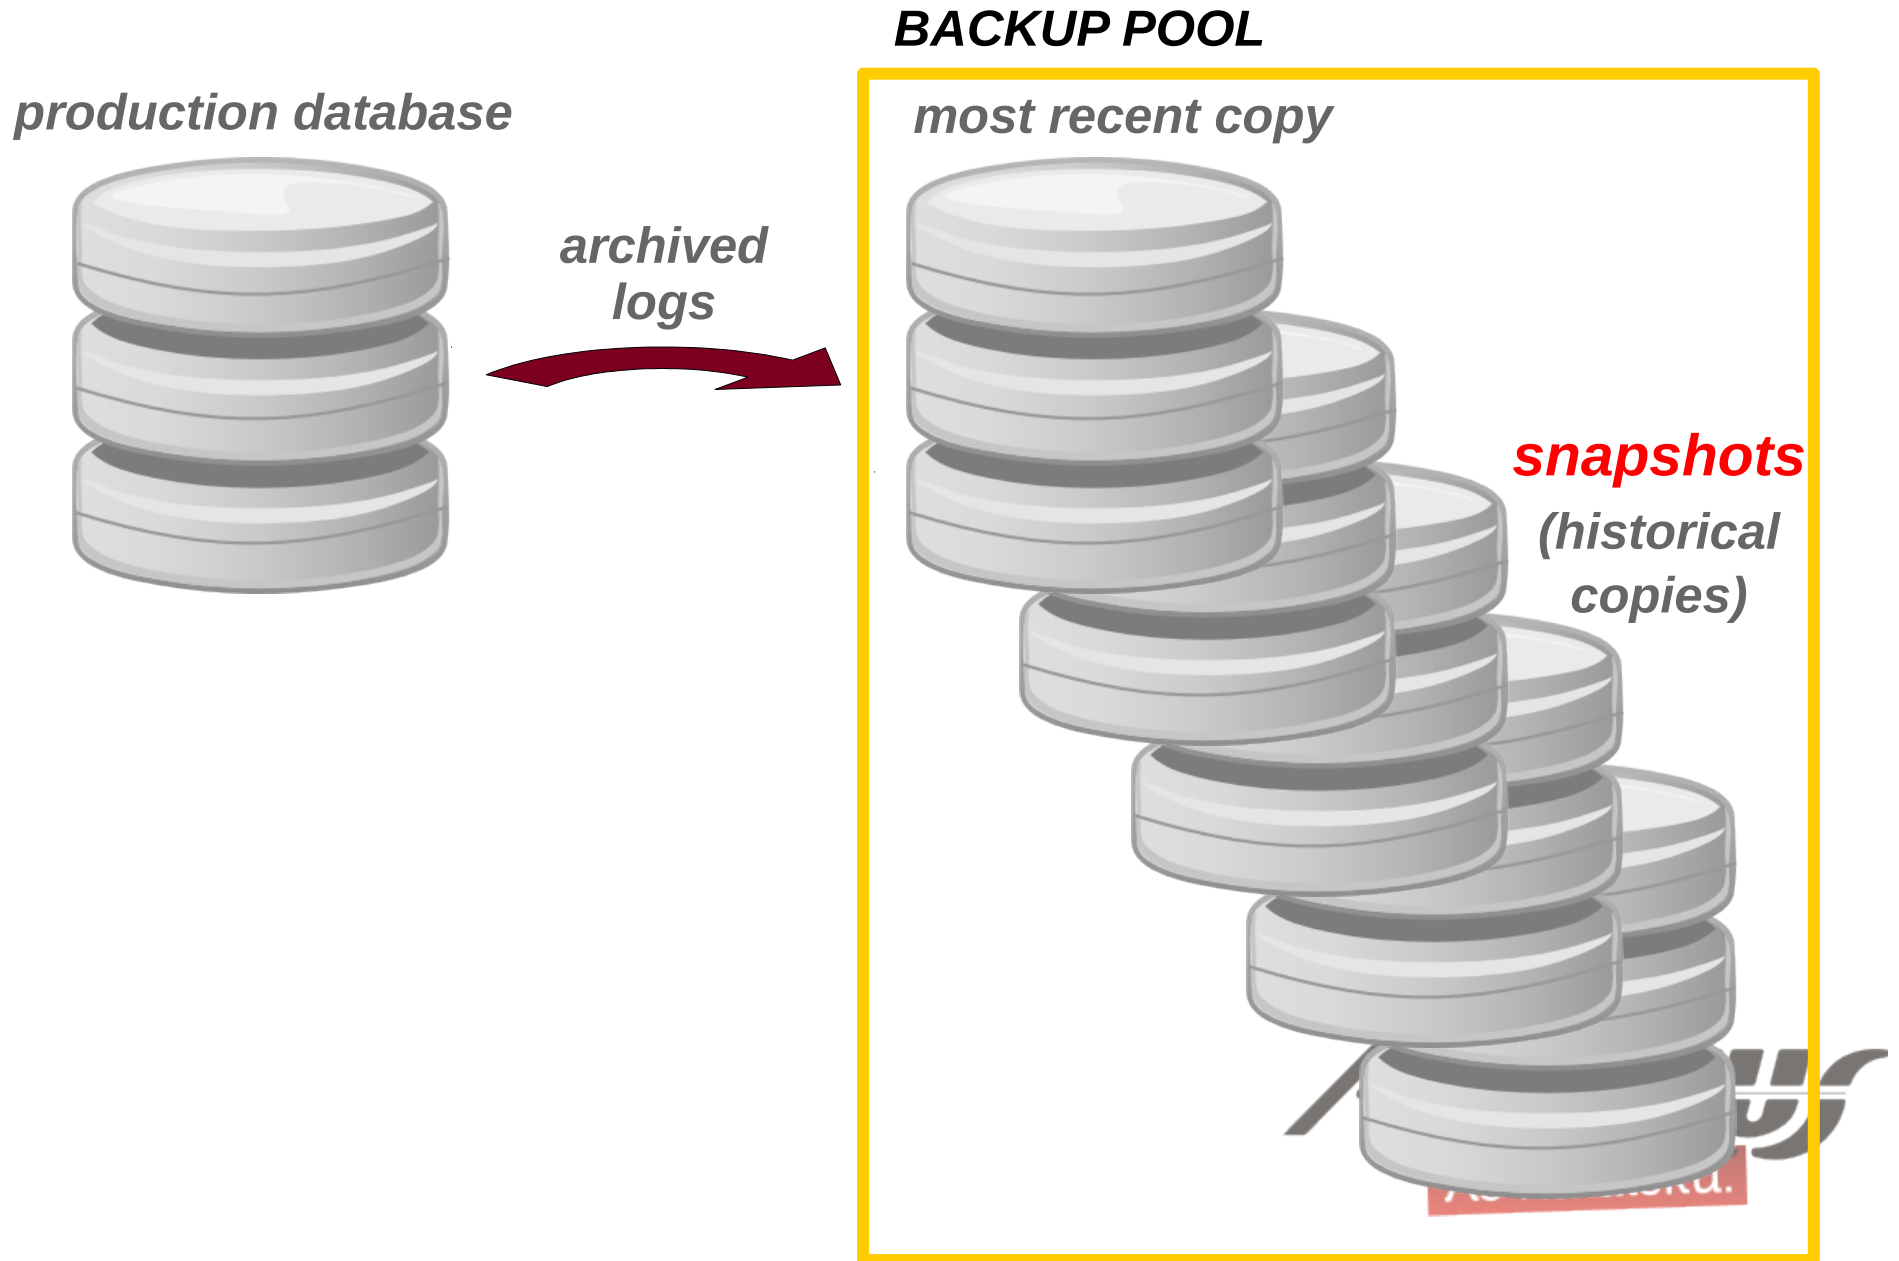
Pros & Cons

- ✓ simple
- ✓ standard (RMAN tool)
- ✓ optimized (especially with EE)
- ✓ possible to start an instance using backup datafiles and access to DB objects
- ✗ limited historical data
- ✗ disk space
- ✗ performance impact to production database





Backup Server





Implicit Deduplication

Examples

- 577 days of backup history
- 416,40 TB of data occupies 22 TB of physical disk space

Backup Server *backup*

563 backups provide **577 days** days of history until 2015-02-27.

234.06 TB of backup data is stored on 3.75 TB / **4.00 TB** physical volume.

Backup Server

398 backups provide **113 days** days of history since 2016-06-06.

416,40 TB of backup data is stored on 21,90 TB / **39,09 TB** physical volume.



Pros & Cons

- ✓ innovative
- ✓ optimized (even with SE)
- ✓ space efficient
- ✓ **near real-time backup**
- ✓ no interference with production database
- ✓ alternative use (DejaVu)
- ✗ potential minor data loss





Perfect Recall

What is a Perfect Recall?

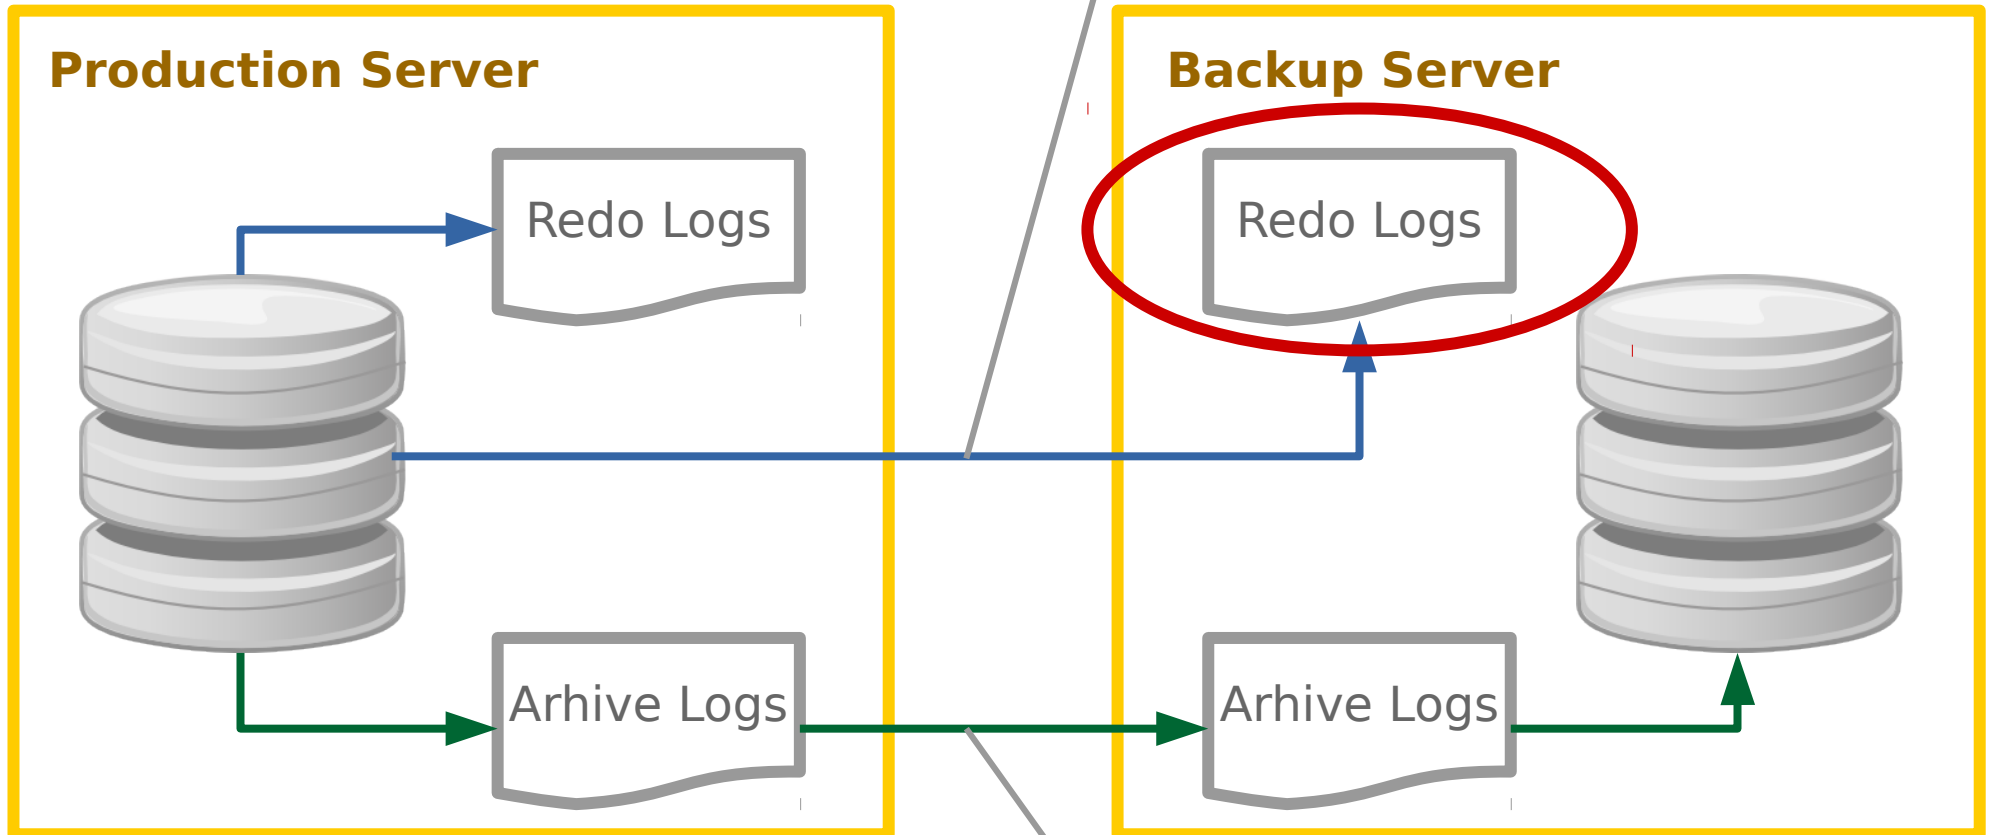
- **true real time backup**
- no data loss

Every committed transaction has already been backed up!

Perfect Recall



Each redo log group has one member located on disk provided by Backup Server



One archivelog destination is located on backup server



Redo Replication Options

- OS (Linux) replication using **DRBD** – network replication
 - ✗ very complex to set and maintain
 - ✓ asynchronous option
- Oracle replication on **NAS** – NFS
 - ✗ unstable
- Oracle replication on **SAN** – iSCSI, SRP





Hardware Requirements

Perfect recall is a hardware solution!

Requirements:

- high throughput
- low latency
- dedicated, local

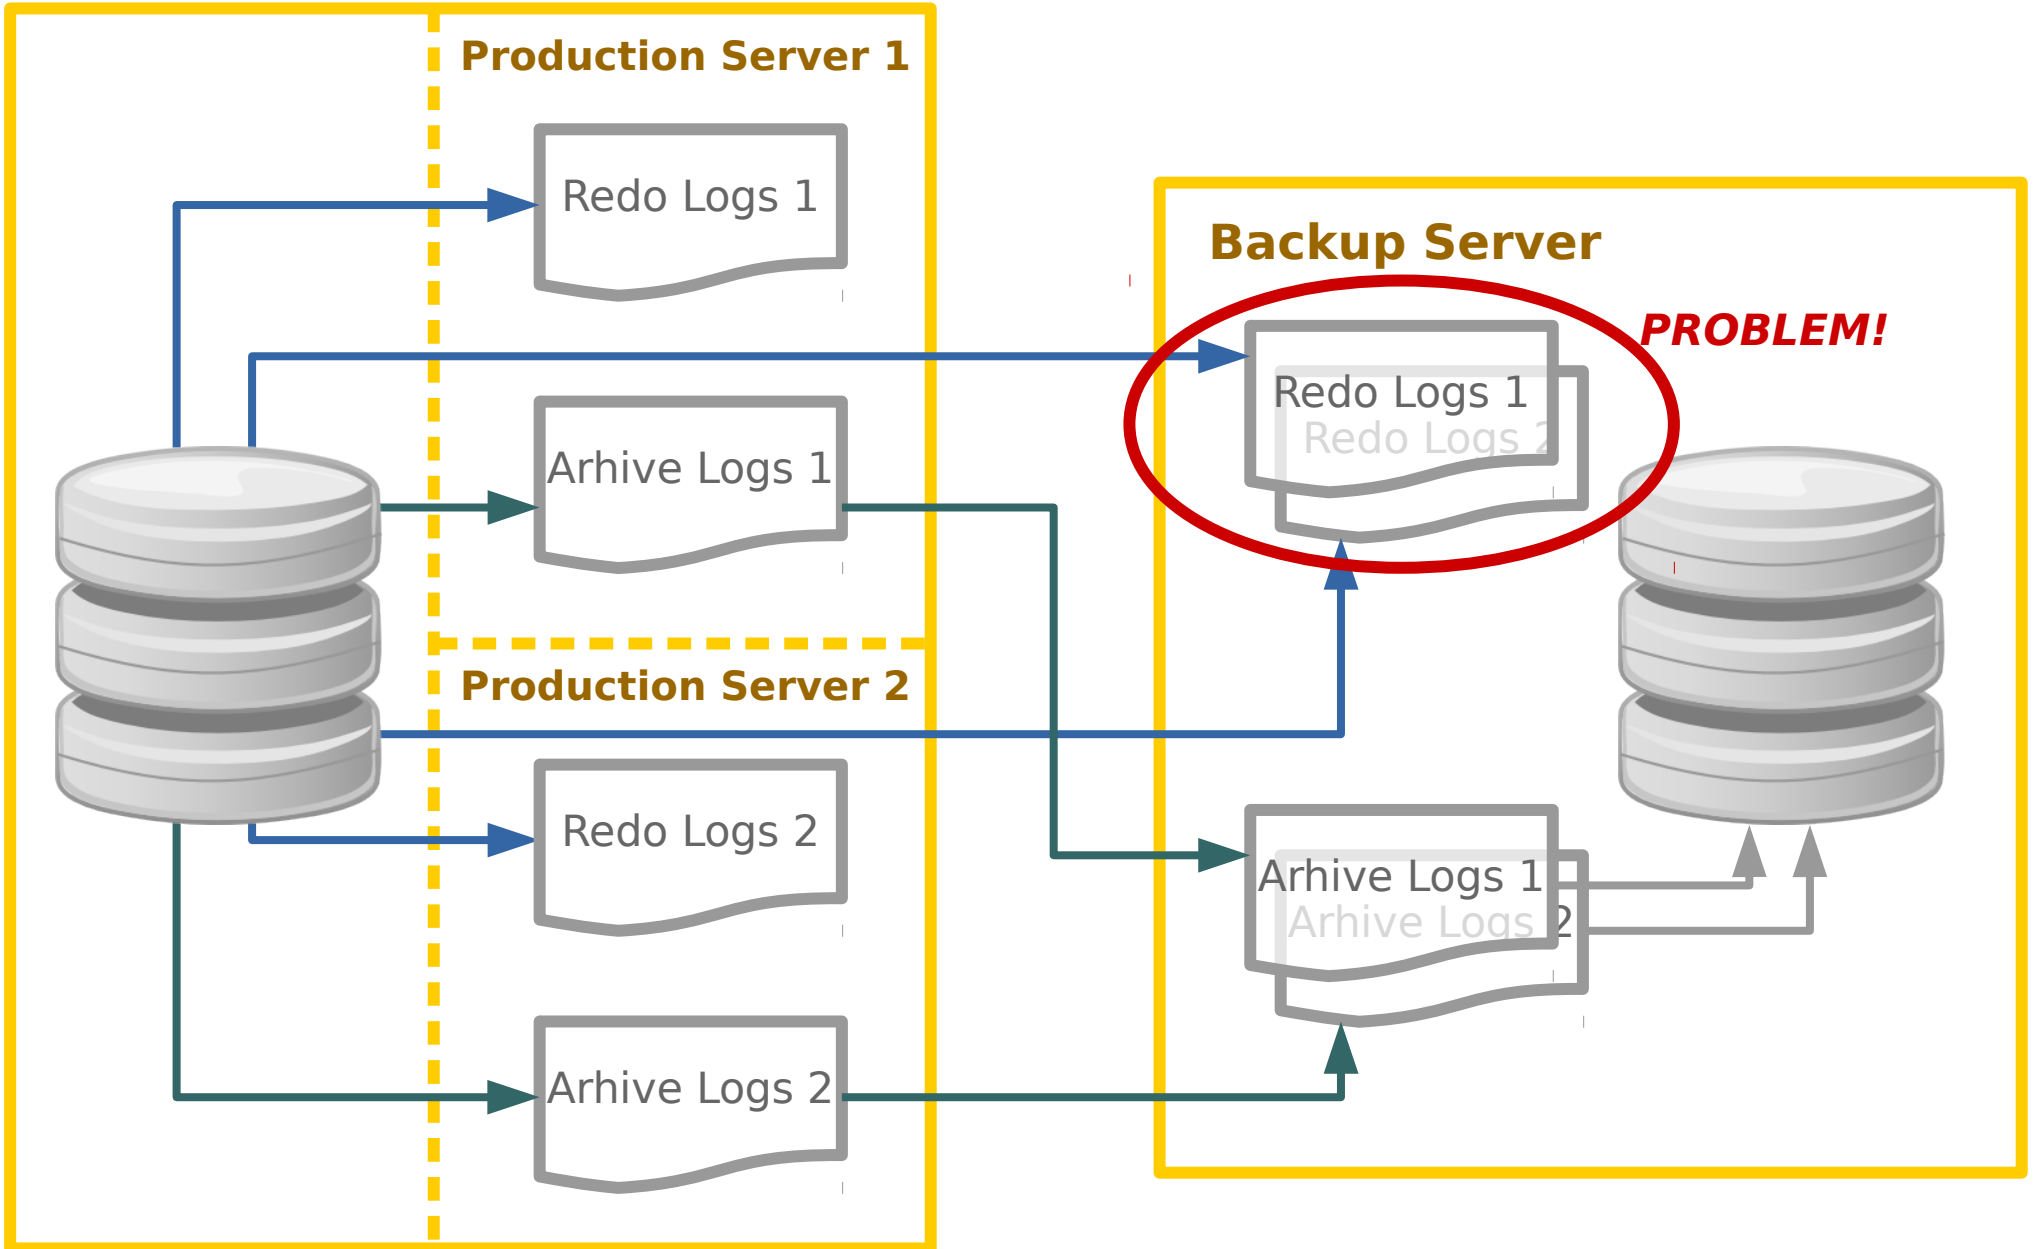
Options:

- 10 Gb ethernet
- QDR InfiniBand





Perfect Recall - RAC





RAC Redo Replication Options

- ~~OS (Linux) replication using **DRBD** – network replication~~
- Oracle replication on **NAS** – NFS
 - ✓ perfectly fits but
 - ✗ unstable
- Oracle replication on **SAN** – iSCSI, SRP
 - ✗ clustered file system
 - ASM (Oracle Automatic Storage Management) hack?





ASM References

- Bane Radulović;
ASM Support Guy
(<http://asmsupportguy.blogspot.si/2013/08/allocation-table.html>)





Restore/Recover

Resources

Status	Type	Name	Flash Date	First Date	Last Date
OFFLINE	database	DWSTAGE1	2014-10-02 18:10:13	2014-10-03 23:00:00	2014-06-20 23:00:00
OFFLINE	database	DWTARGET	2014-10-02 19:10:30	2014-10-03 23:00:00	2014-06-20 23:00:00
OFFLINE	database	EB	2014-10-02 18:10:07	2014-10-03 23:00:00	2014-06-20 23:00:00
OFFLINE	database	ORCL10	2014-10-02 21:10:49	2014-10-03 23:00:00	2014-06-20 23:00:00
OFFLINE	database	RIS10G	2014-10-02 23:10:16	2014-10-03 23:00:00	2014-06-20 23:00:00
	vm-incremental	san1		2014-10-03 18:00:00	2014-06-20 18:00:00

Oracle PITR Instance

Refresh

Backups

Path	Size	Taken	Actual	Status
/zbackup/ORCL10-2014-10-03-23-00				UNKNOWN
/zbackup/ORCL10-2014-10-02-23-00	861 G		21:10:49	OK
/zbackup/ORCL10-2014-10-01-23-00	860 G		21:10:42	OK
/zbackup/ORCL10-2014-09-30-23-00	860 G		21:09:22	OK
/zbackup/ORCL10-2014-09-29-23-00	859 G		21:09:44	OK
/zbackup/ORCL10-2014-09-28-23-00	858 G		05:09:04	OK
/zbackup/ORCL10-2014-09-27-23-00	858 G		21:09:59	OK
/zbackup/ORCL10-2014-09-26-23-00	858 G		21:09:23	OK
/zbackup/ORCL10-2014-09-25-23-00	858 G		19:09:28	OK
/zbackup/ORCL10-2014-09-24-23-00	857 G		19:09:45	OK
/zbackup/ORCL10-2014-09-23-23-00	857 G	2014-09-23 23:00:00	2014-09-23 19:09:47	OK
/zbackup/ORCL10-2014-09-22-23-00	857 G	2014-09-22 23:00:00	2014-09-22 19:09:53	OK
/zbackup/ORCL10-2014-09-21-23-00	857 G	2014-09-21 23:00:00	2014-09-21 19:09:46	OK
/zbackup/ORCL10-2014-09-20-23-00	857 G	2014-09-20 23:00:00	2014-09-20 19:09:38	OK
/zbackup/ORCL10-2014-09-19-23-00	857 G	2014-09-19 23:00:00	2014-09-19 19:09:55	OK
/zbackup/ORCL10-2014-09-18-23-00	857 G	2014-09-18 23:00:00	2014-09-18 19:09:50	OK
/zbackup/ORCL10-2014-09-17-23-00	857 G	2014-09-17 23:00:00	2014-09-17 19:09:47	OK
/zbackup/ORCL10-2014-09-16-23-00	857 G	2014-09-16 23:00:00	2014-09-16 19:09:53	OK
/zbackup/ORCL10-2014-09-15-23-00	857 G	2014-09-15 23:00:00	2014-09-15 19:09:10	OK
/zbackup/ORCL10-2014-09-14-23-00	857 G	2014-09-14 23:00:00	2014-09-14 19:09:32	OK

Details ✕

Resource is OFFLINE

Target:

Point-in-Time:

Log Level:

Log Files: [alert_pthibis.log](#), [bks_task_pthibis.log](#)



Conclusion

Traditional RMAN backup (Incrementally Updated Backups)

- ✓ simple
- ✓ standard
- ✓ optimized
- ✗ limited historical data
- ✗ disk space
- ✗ performance impact to production database
- ✗ potential minor data loss

Backup Server

- ✓ innovative
- ✓ optimized (even with SE)
- ✓ space efficient
- ✓ **near real-time backup**
- ✓ no interference with production database
- ✓ alternative use (DejaVu)
- ✗ potential minor data loss

Perfect Recall

- ✓ **true real time backup**
- ✓ no data loss
- ✗ hardware requirements
- ✗ somehow complicated (especially with RAC)






Prize Game Invitation

On your way out you are invited
to take a coupon for entering
Abakus plus prize game draw.





Perfect Recall

Thank You

mag. Sergej Rožman

ABAKUS plus d.o.o.

Ljubljanska c. 24a, Kranj, Slovenija

e-mail: sergej.rozman@abakus.si

phone: +386 4 287 11 14

