DB & VM Backup and Recovery

- **mag. Sergej Rožman**; Abakus plus d.o.o.

- The latest version of this document is available at: http://www.abakus.si/

Real men don't use backups, they post their stuff on a public ftp server and let the rest of the world make copies.

Linus Torvalds

Abakus
As na disku.

DB & VM
Backup and Recovery

mag. Sergej Rožman

sergej.rozman@abakus.si

SrOUG

ORACLE Gold Partner

Sava
Mestna občina Ljubljana

GOODYEAR

Aerodrom Ljubljana

Mercator

Iskra
Iskra Sistemi

MERKUR

MESTNA OBČINA KOPER
COMUNE CITTÀ DI CAPODISTRIA

REPUBLIKA SLOVENIJA
MINISTRSTVO ZA FINANCE

BANKA SLOVENIJE
EVROSISTEM

KONTROLA ZRAČNEGA PROMETA SLOVENIJE

Abakus
As na disku.

DB & VM
Backup and Recovery

mag. Sergej Rožman

sergej.rozman@abakus.si

ORACLE Gold Partner

Sava
Mestna občina Ljubljana
GOODYEAR
Aerodrom Ljubljana
Mercator
Iskra
Iskra Sistemi
MERKUR
MESTNA OBČINA KOPER
COMUNE CITTÀ DI CAPODISTRIA
REPUBLIKA SLOVENIJA
MINISTRSTVO ZA FINANCE
BANKA SLOVENIJE
EVROSISTEM
KONTROLA ZRAČNEGA PROMETA SLOVENIJE

Abakus
As na disku.

DB & VM
Backup and Recovery

mag. Sergej Rožman

sergej.rozman@abakus.si

PALSIT INFOSEK 2013

ORACLE Gold Partner

Sava GOODYEAR Aerodrom Ljubljana Mercator Iskra Iskra Sistemi MERKUR

Mestna občina Ljubljana

MESTNA OBČINA KOPER
COMUNE CITTÀ DI CAPODISTRIA

REPUBLIKA SLOVENIJA
MINISTRSTVO ZA FINANCE

BANKA SLOVENIJE
EVROSISTEM

KONTROLA ZRAČNEGA PROMETA SLOVENIJE

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

Infrastructure:

- from 1995 GNU/Linux ***(18 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
MESTNA OBČINA KOPER
COMUNE CITTÀ DI CAPODISTRIA
Iskra
Iskra MIS
Aerodrom Ljubljana
KONTROLA ZRAČNEGA PROMETA SLOVENIJE
Gorenjska Banka
Banka s poslubom
GOODYEAR

Abakus plus d.o.o. - Kranj

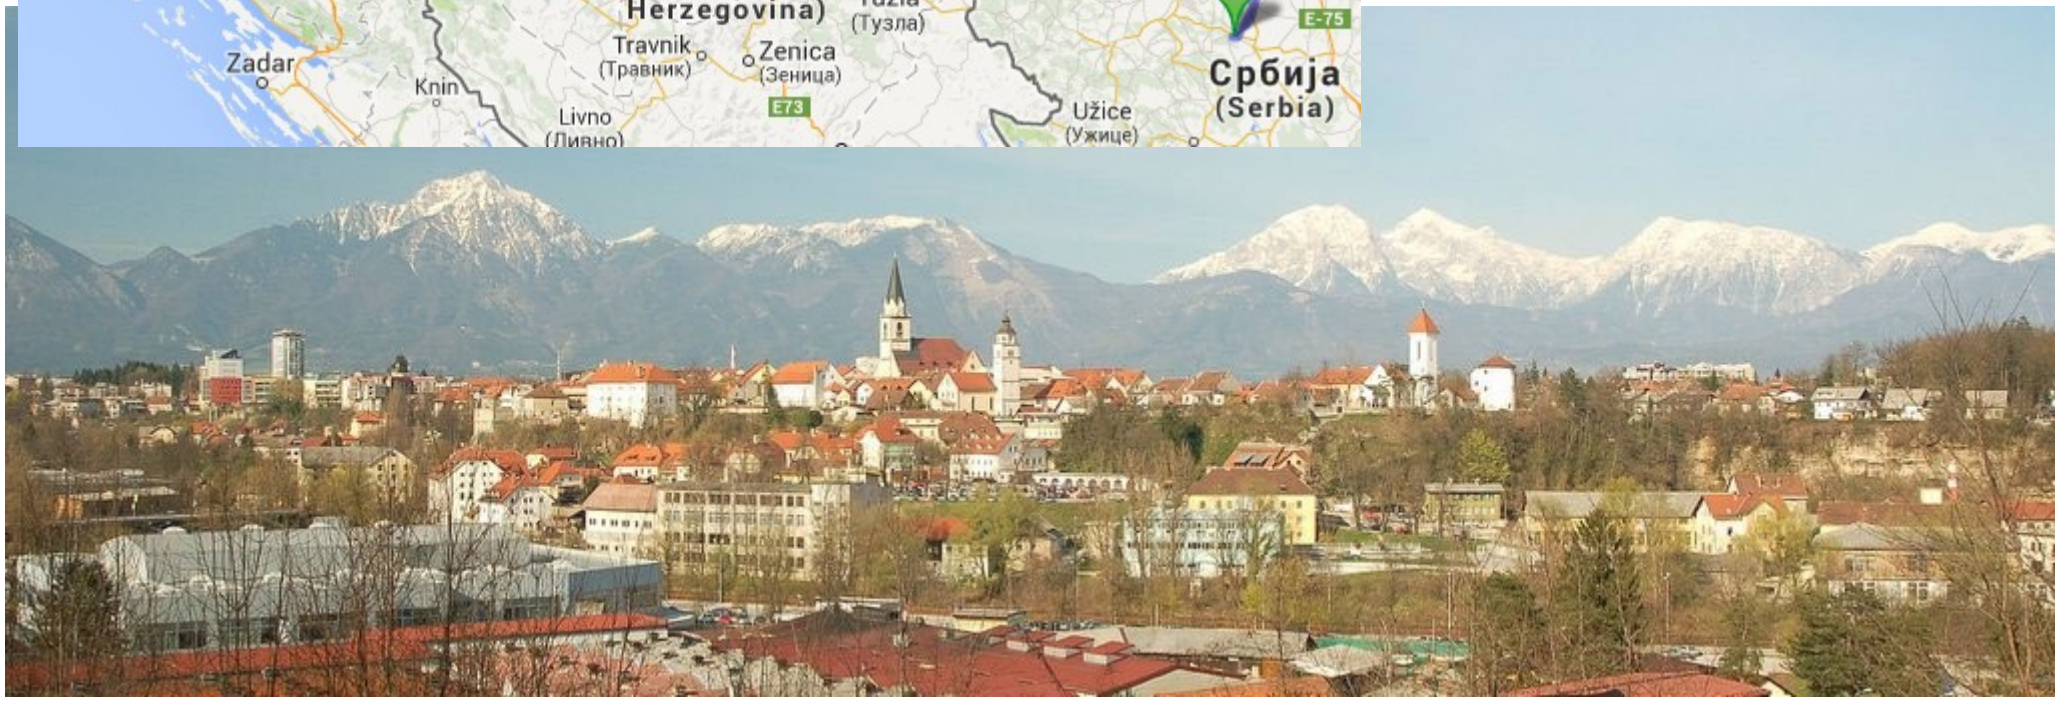

Backup and Recovery Best Practices

- 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.

Abaku
As na disku.

Classic Full/Incremental Backup Model

- backup takes long time (especially full)
- restore takes even longer (full + n × incremental)
- incremental backups not suitable for large files (DB, VMs)

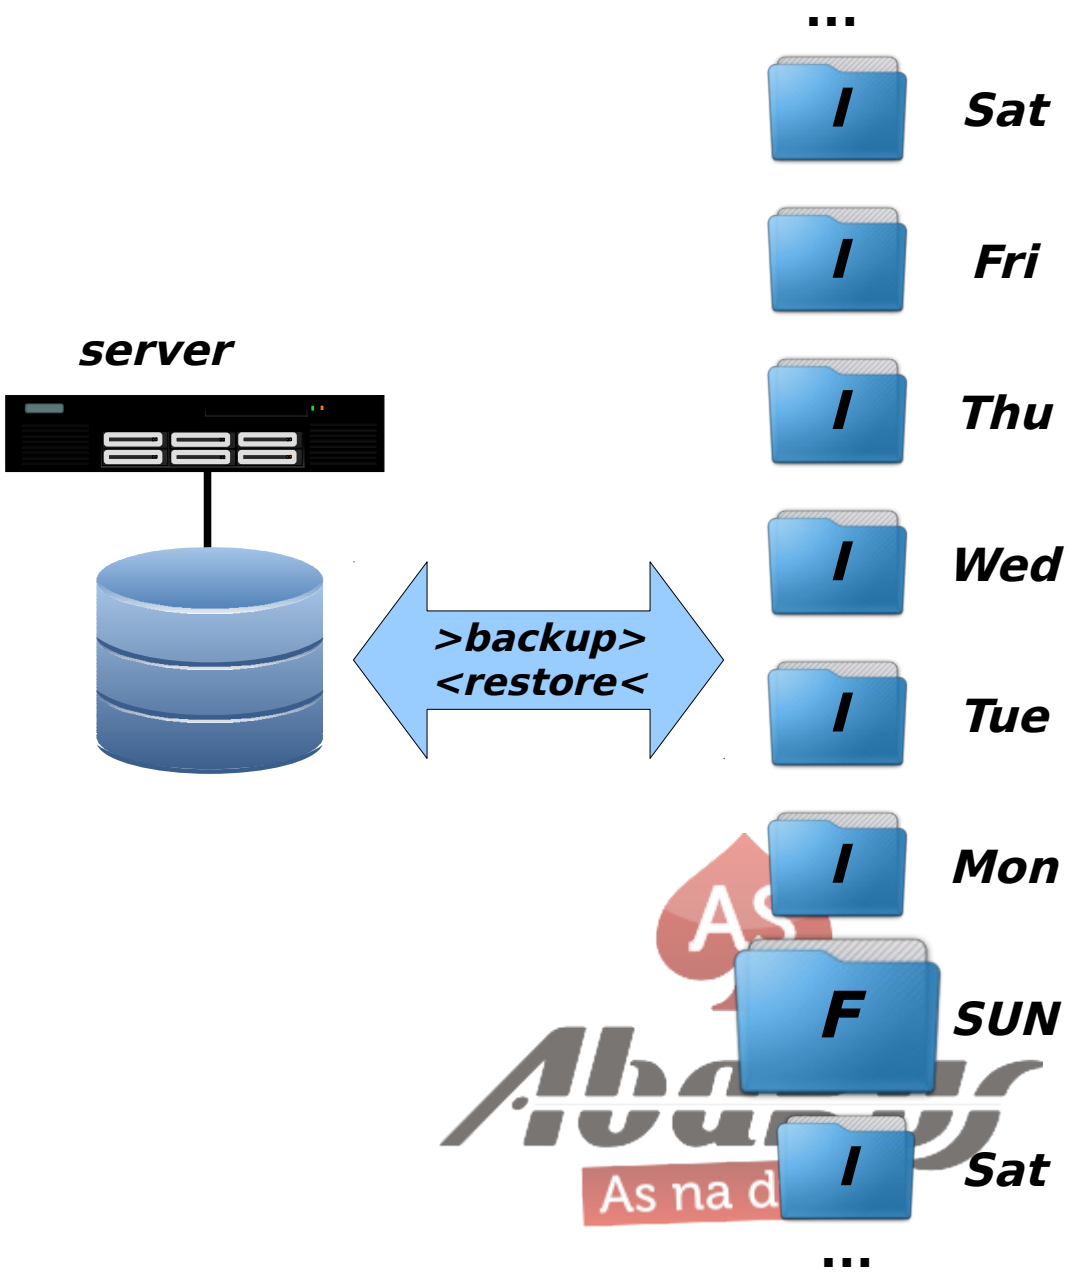

Status Board

Fact	DB	VM	Notes
BACKUP in no time	☒	☒	
no resources	☒	☒	
no disk space	☒	☒	or no tape space
RECOVER in no time	☒	☒	
no resources	☒	☒	
COPIES without errors	☒	☒	
consistent	☑	☑	if done right
DATA always available	☒	☒	☑ with autoloader
always in view	☒	☒	

Tape vs. Disk drive

Tape

- price:
  n×1000€ (drive) + <100€ / (cartridge)
  (LTO-6 native capacity 2.5 TB)
- no future compatibility
  (new drive needed)
- Is your data really on that tape?
- high throughput, slow access time

Disk

- price: >100€ / 3 TB SATA
- guaranted future compatibility
- WYSIWYG (if you see data, you can get data)
- moderate througput, fast access time

Abaku
As na disku.

Status Board (using disks)

Fact	DB	VM	Notes
BACKUP in no time	☒	☒	
no resources	☒	☒	
no disk space	☒	☒	
RECOVER in no time	☒	☒	
no resources	☒	☒	
COPIES without errors	☑	☑	with mirrored disks
consistent	☑	☑	
DATA always available	☑	☑	
always in view	☑	☑	

DB Backup Full/Incremental – Example

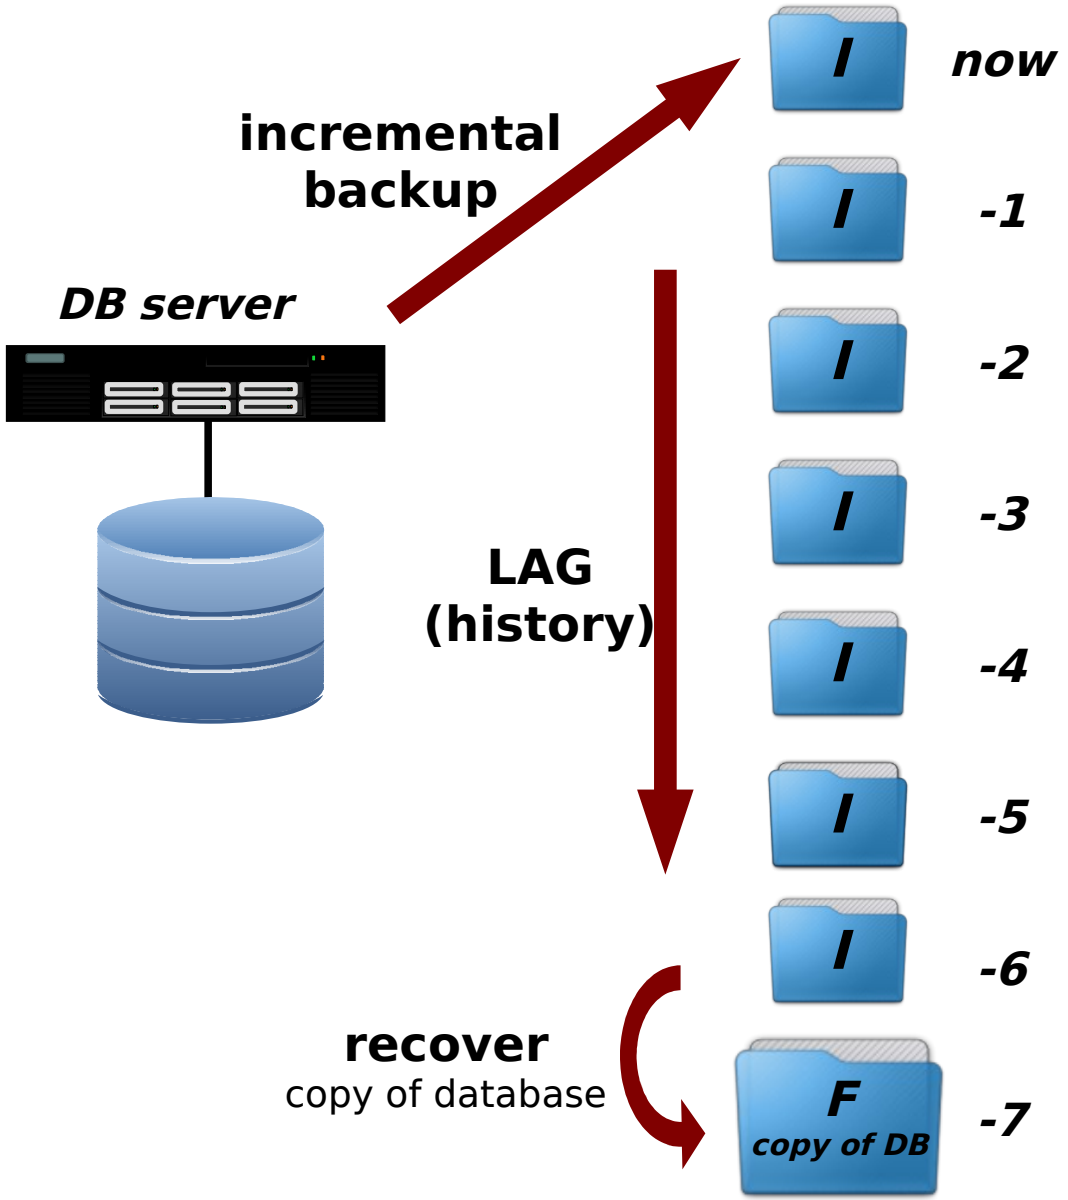

- somewhat optimized
  no Full backup except initial
- incremental backup optimized
  with Oracle Enterprise Edition
  (block change tracking)
- restore still takes long time

Abakus
As na disku.

Backup in »no time«

- backup takes no additional time
- backup needs no production resources
- **BUT, WHERE IS BACKUP HISTORY ?**

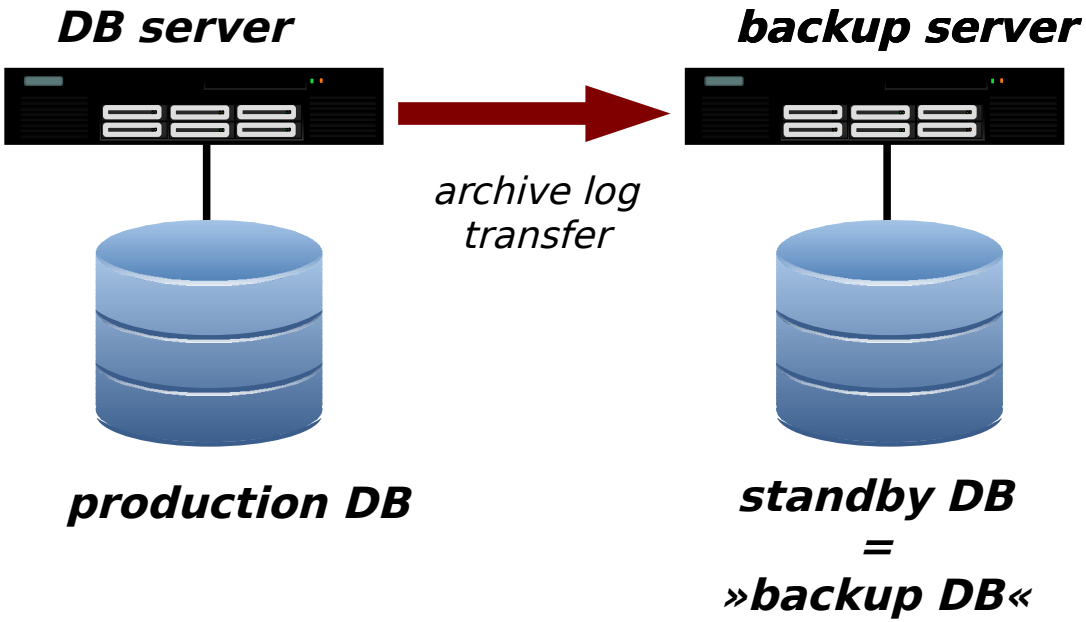

AS Abakus
As na disku.

… and history

- snapshot backup DB
- save snapshot

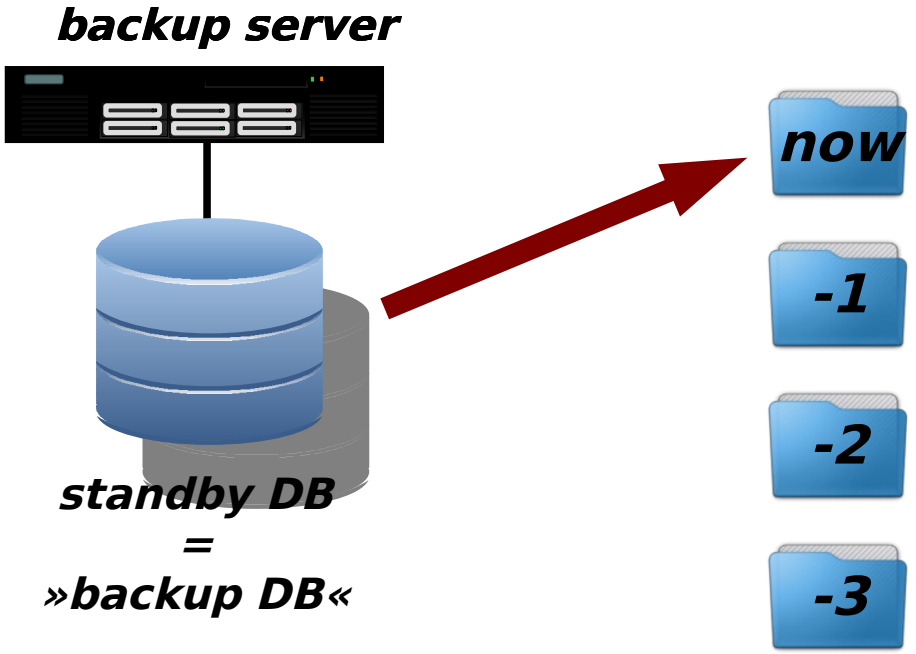

- **Time and resources are consumed exclusively on backup server**
- **Backup occupies a lot of disk space !?**
  (n × size of DB)

Deduplication

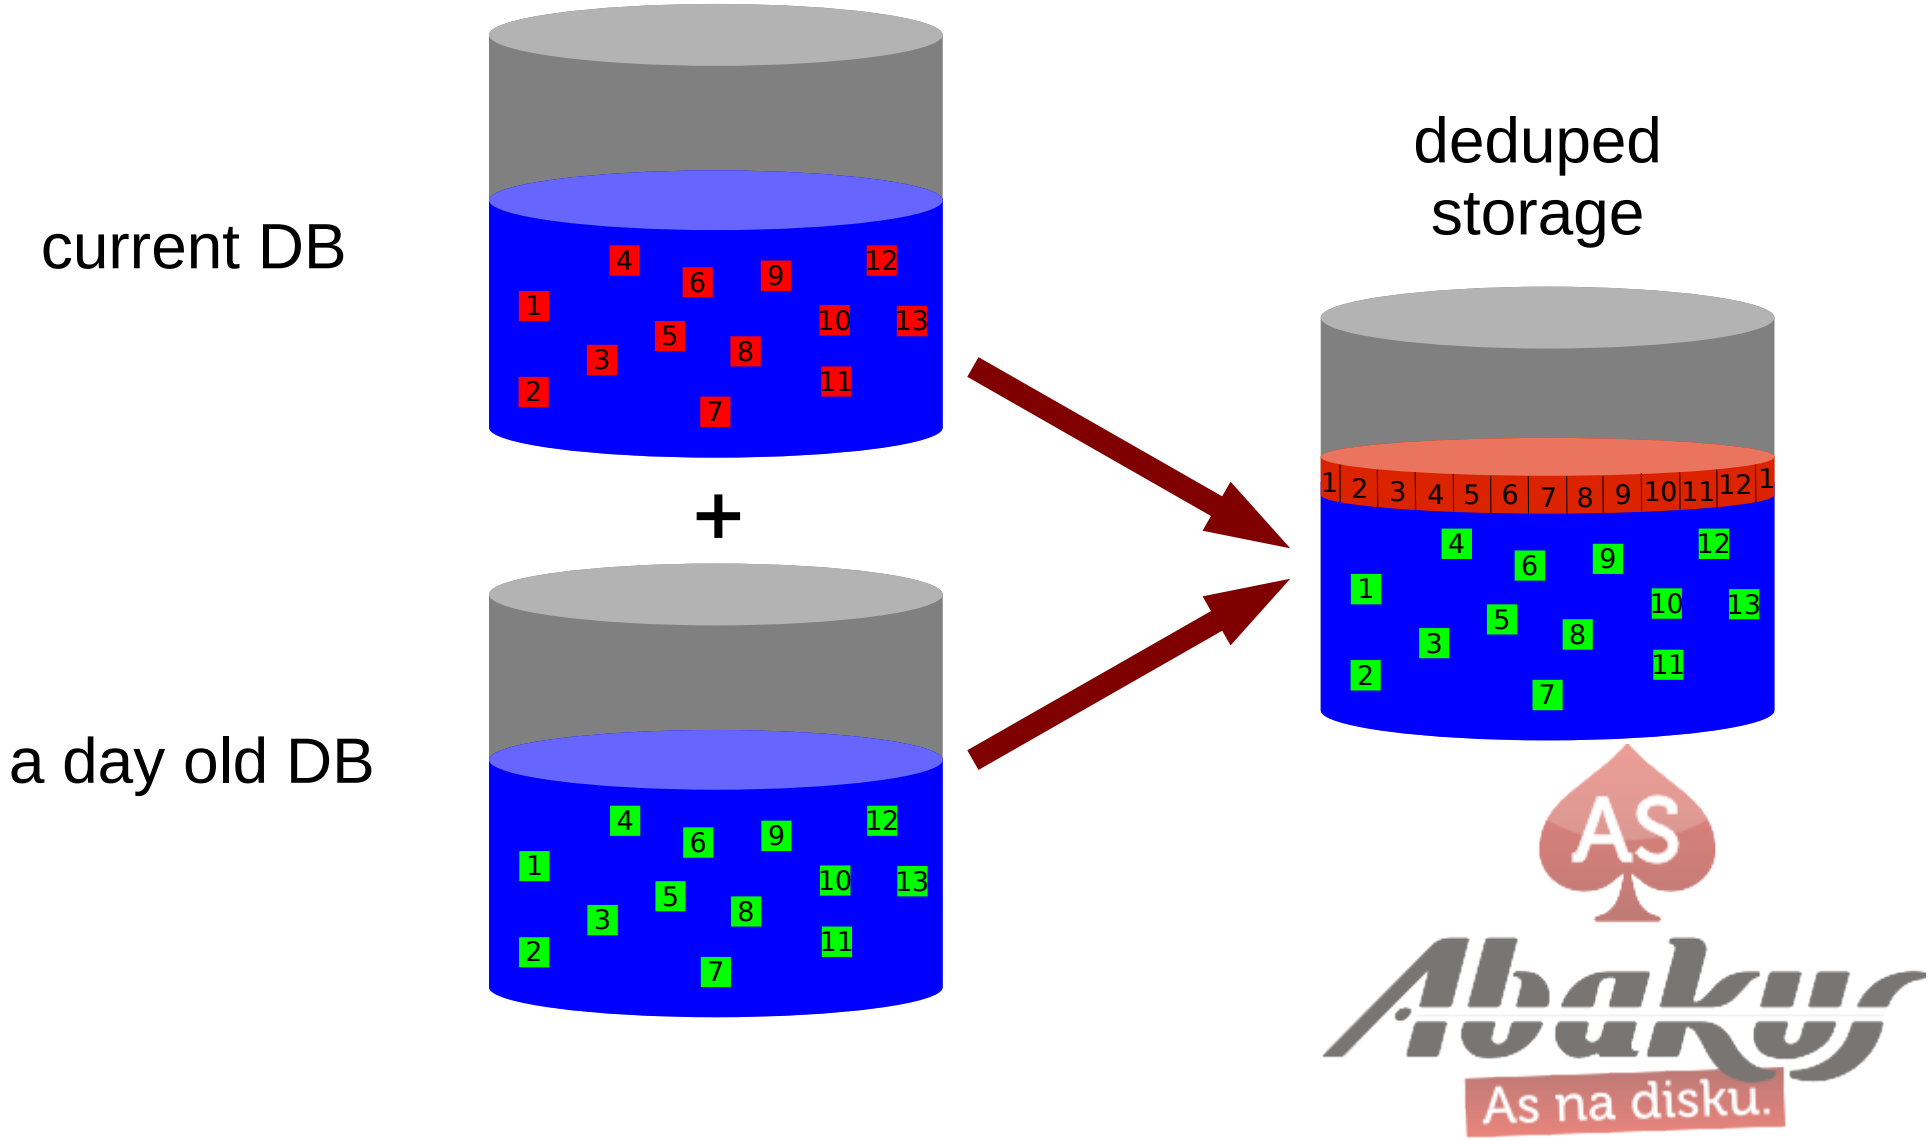

… and (almost) no disk space

- snapshot backup DB
- save snapshot
  to deduplicated area

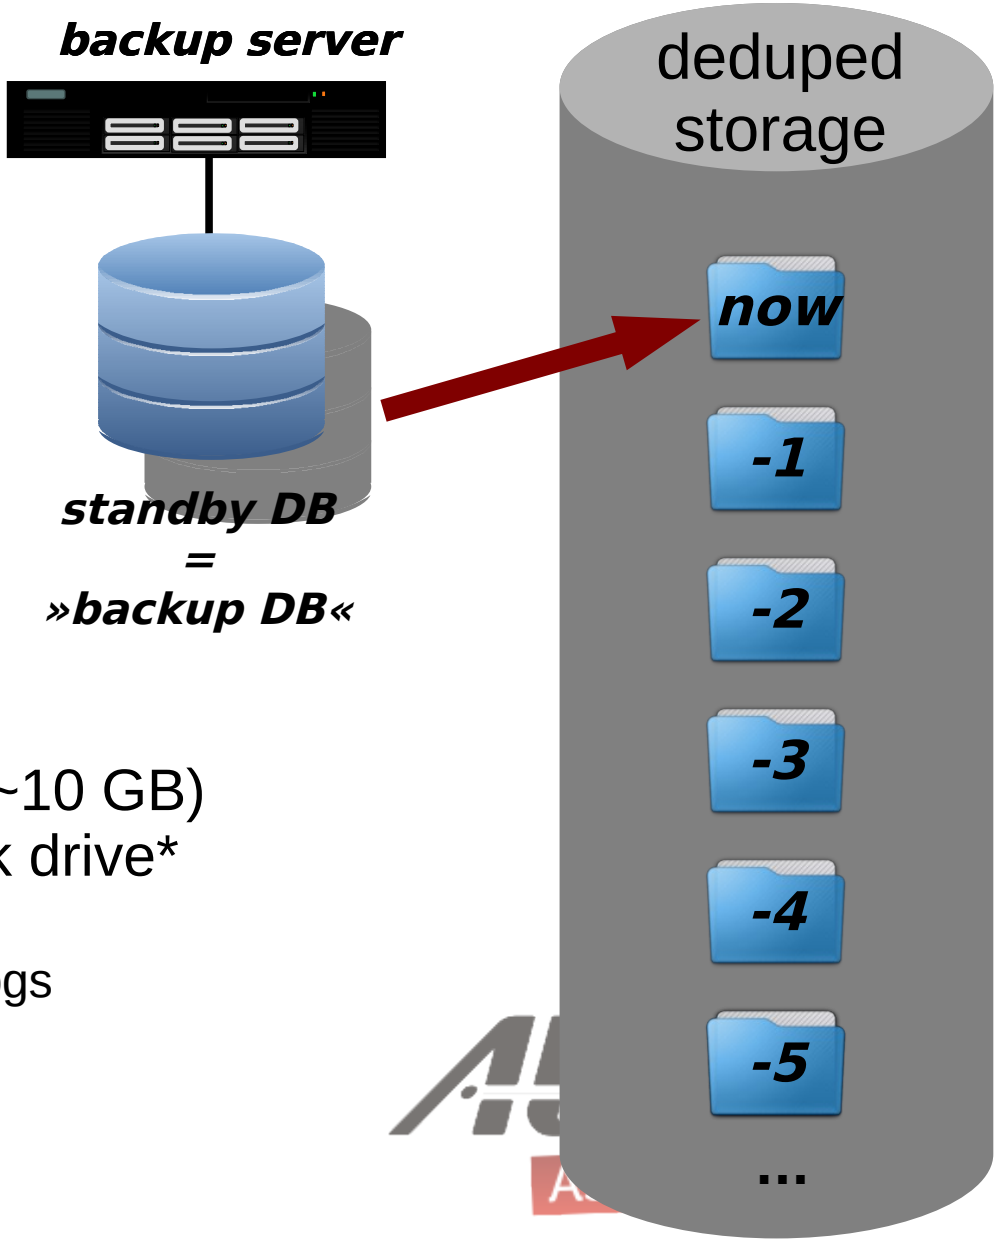

Example:

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

* not entirely true, did not consider archived logs

Status Board

Fact	DB	VM	Notes
BACKUP in no time	☑	☒	
no resources	☑	☒	
no disk space	☑	☑	dedupe plays well on VMs too
RECOVER in no time	☒	☒	
no resources	☒	☒	
COPIES without errors	☑	☑	
consistent	☑	☑	
DATA always available	☑	☑	
always in view	☑	☑	

Recover in »no time«

Lost or currupted datafile (or even whole DB)

- make snaphot of backup DB
- start SAN software – iSCSI, SRP, …
- export LUN with backup of lost datafile
- connect LUN to DB server
- offline corrupted datafile
- switch to backup datafile
- recover backup datafile
- online backup datafile

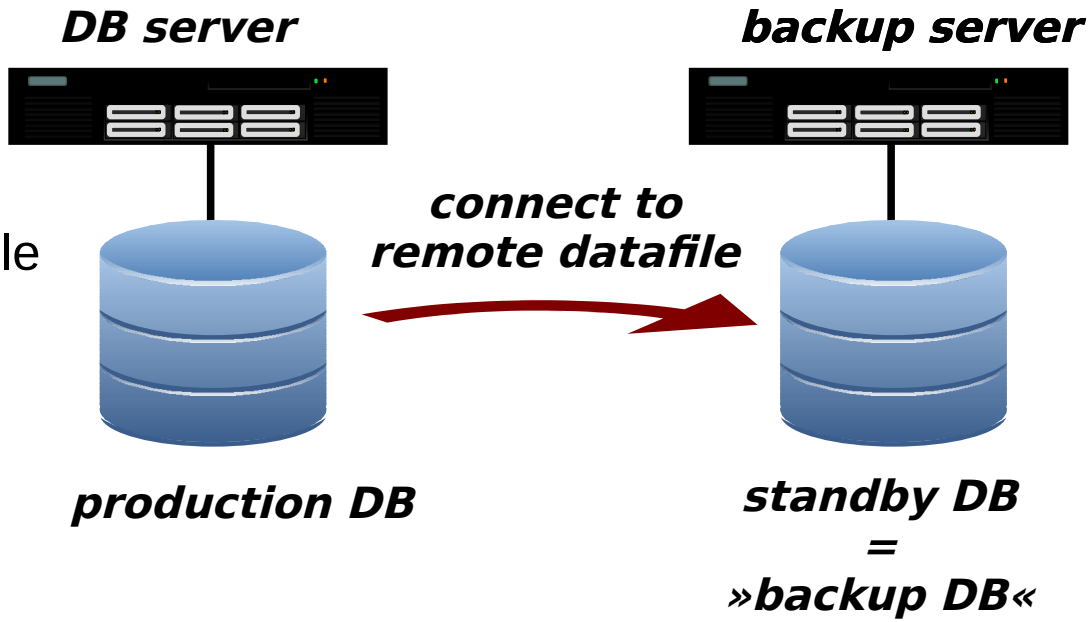

BACK IN BUSINESS IN NO TIME!*

* real restore in more appropriate time

Abakus
As na disku.

Performance

- test 1 (notebook with SSD, DB on VM):
  max_iops = **9.983**, latency = **8**, max_mbps = **251**
- test 2 (test DB, 10x 600 GB 15k FC):
  max_iops = **1.824**, latency = **11**, max_mbps = **280**
- test 3 (production DB, 30x 146 GB 15k FC):
  max_iops = **6.498**, latency = **10**, max_mbps = **455**
- test 4 (**Abakus SAN**, 16x SSD, Infiniband 40G):
  max_iops = **43.782**, latency = **0**, max_mbps = **1.727**

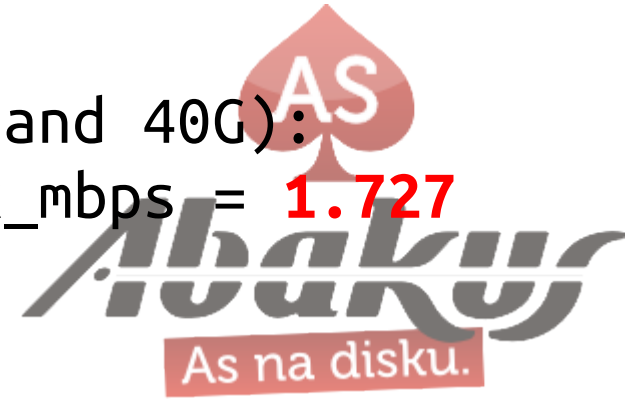

Restore/Access to Historical Data

- snapshot selected slot
- stop standby database
- switch active slot to snapshot
- start instance
- recover database until needed (optional)
- open database

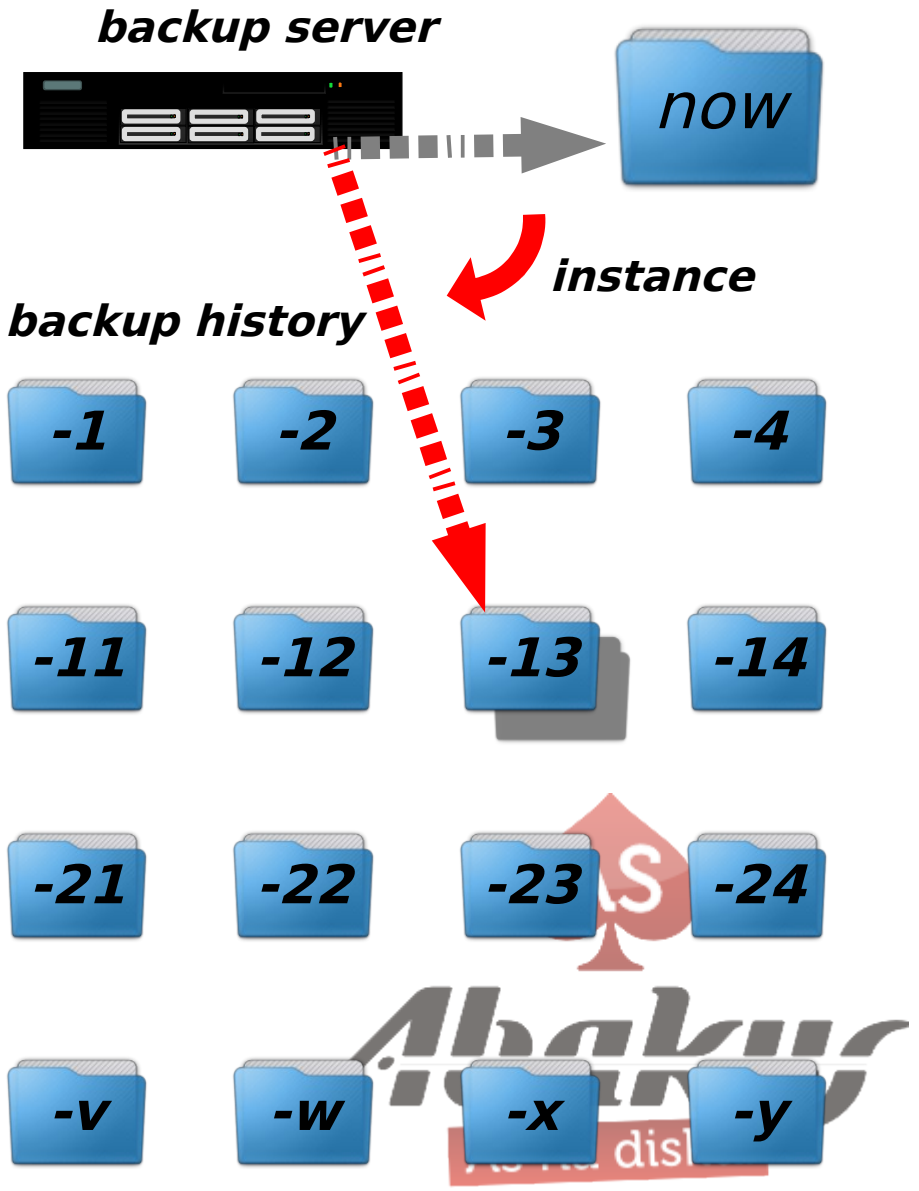

Status Board

Fact	DB	VM	Notes
BACKUP in no time	☑	☒	
no resources	☑	☒	
no disk space	☑	☑	
RECOVER in no time	☑	☑	
no resources	☑	☑	
COPIES without errors	☑	☑	
consistent	☑	☑	
DATA always available	☑	☑	
always in view	☑	☑	

Work in Progress

- general change block tracking at block device level (asynchronous and buffered)

  - will copy ONLY changed disk blocks to backup
  - suitable for VMs
  - will make possible to back up to remote site over slow link
  - no active instance – no license fee
- point in time recovery for VMs, (maybe)
- graphical user interface, (maybe)

Abakus
As na disku.

Status Board

Fact	DB	VM	Notes
BACKUP in no time	☑	☑	
no resources	☑	☑	
no disk space	☑	☑	
RECOVER in no time	☑	☑	
no resources	☑	☑	
COPIES without errors	☑	☑	
consistent	☑	☑	
DATA always available	☑	☑	
always in view	☑	☑	

Guidelines

- Open design, »everything is possible«.
  No automatic protection from »failure by design«. Plan cerefully!
- Can not have everything
  - remote back up is preferred, but makes restore more difficult
  - on line backups can be compromised (viruses, sabotage).
- Offline backups are still crucial for archival purposes
  (and if everything else fails).
- **Test procedures regularly!**

References

- Husnu Sensoy;
  How to Backup & Recovery Enormous Databases?
  (http://husnusensoy.files.wordpress.com/2009/12/enormous.pdf)

DB & VM Backup and Recovery

Questions

mag. Sergej Rožman

ABAKUS plus d.o.o.
Ljubljanska c. 24a
Kranj

e-mail: sergej.rozman@abakus.rs

phone: +386 4 287 11 14