



USING IMAGE COPIES FOR ORACLE DATABASE BACKUPS

ILMAR KERM

@ILMARKERM

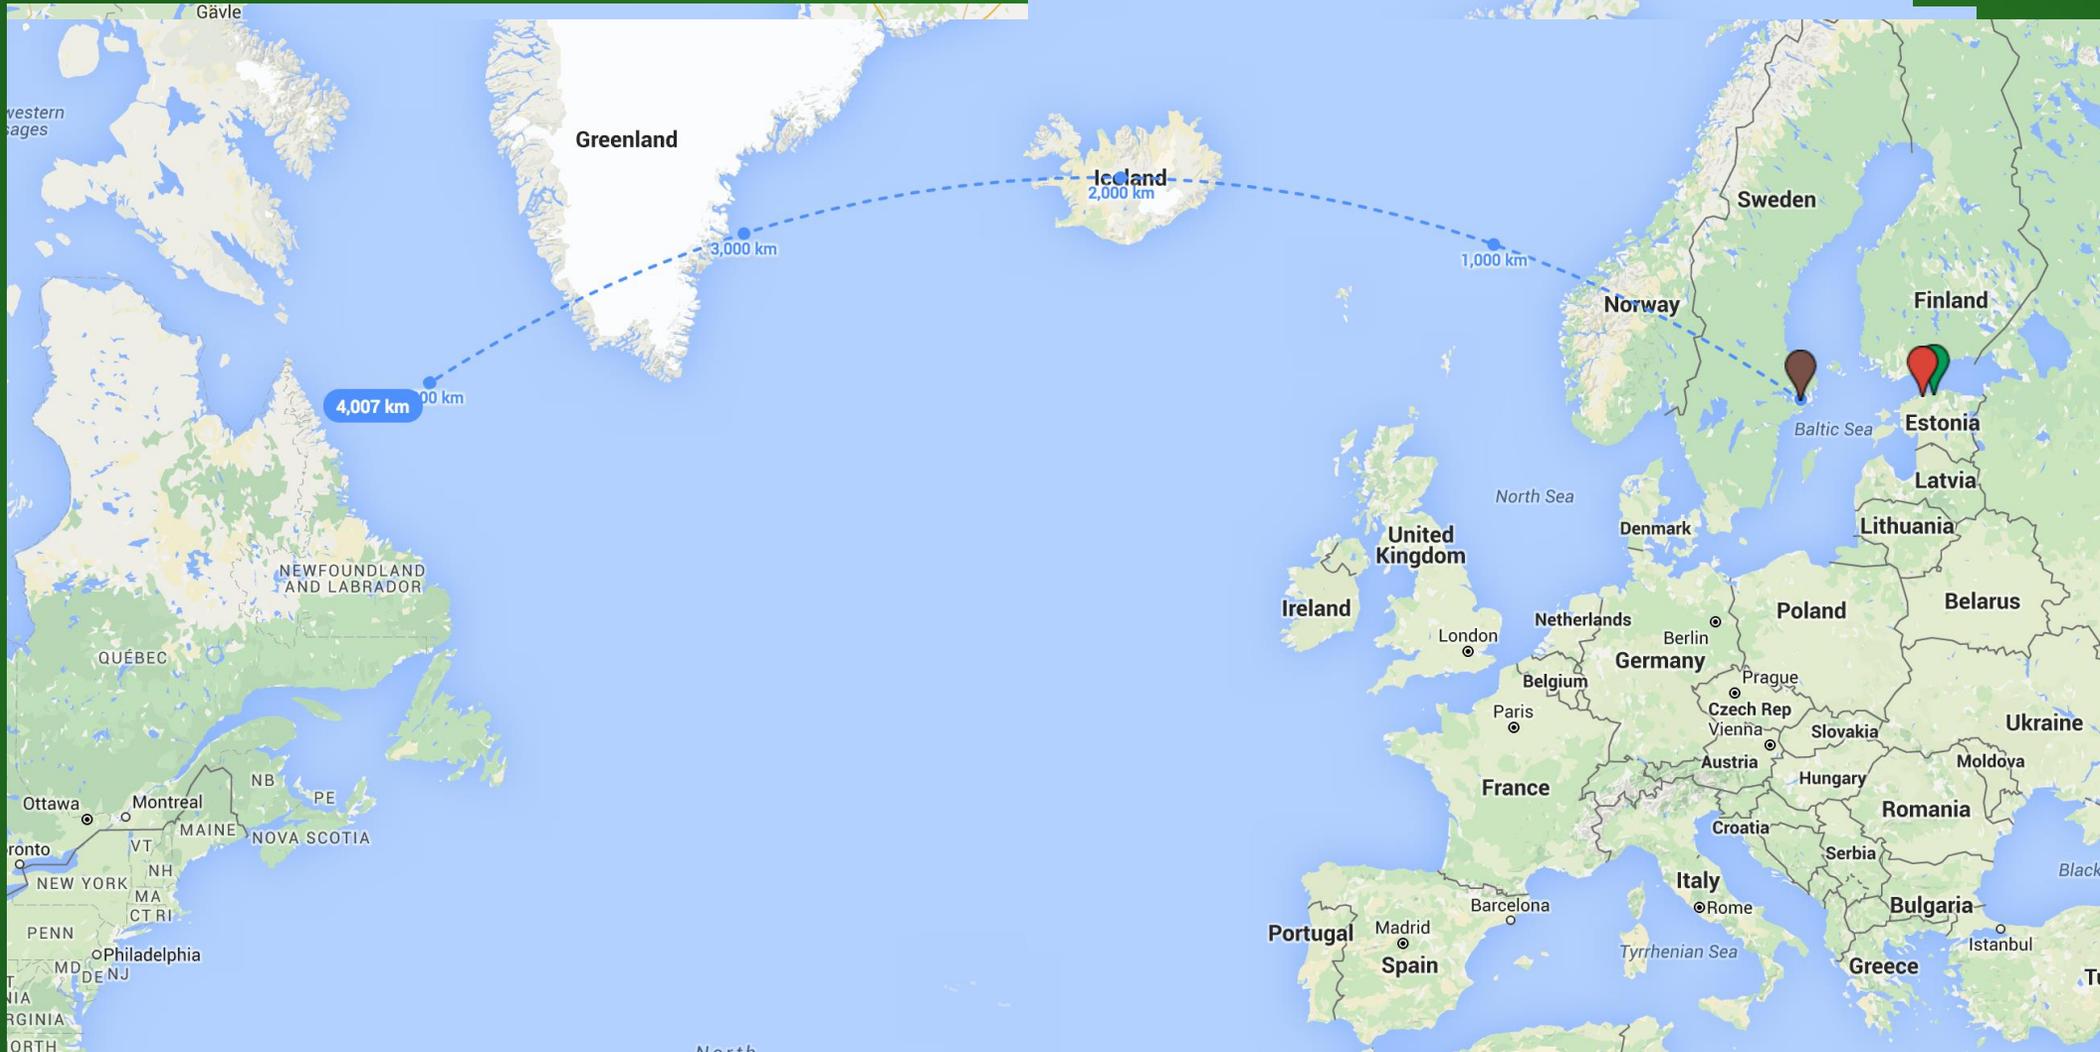
OUTLINE

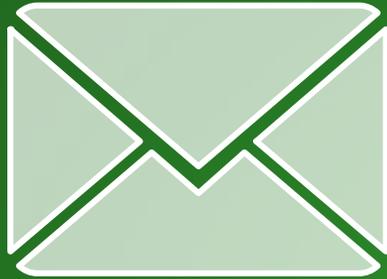
- U WHO AM I**
- U THE PROBLEM WE HAVE TODAY**
- U WHAT SOLUTIONS ARE ON THE MARKET?**
- U USING IMAGE COPIES**
- U RECOVERY SCENARIO**
- U BEYOND RECOVERY**
- U COMPARISON WITH ORACLE ZDLRA**

ILMAR KERM

- u Senior database administrator in Unibet
- u Working in IT since 2000
- u Working with Oracle database since 2003
- u Working in Unibet since 2007/2015
- u From LAMP developer to Oracle DBA
- u President of Oracle User Group Estonia
- u Oracle ACE Associate
- u Blog: <https://ilmarkerm.eu>

WHERE DO I COME FROM?





ILMAR.KERM@UNIBET.COM



@ILMARKERM

UNIBET GROUP

PROUD TO BE ONE OF EUROPE'S BIGGEST ONLINE GAMBLING COMPANIES

OVER 15.3 MILLION CUSTOMER IN 100 COUNTRIES

OVER 1.1 MILLION ACTIVE PLAYERS

OFFICES WORLDWIDE: LONDON, STOCKHOLM, MALTA, GIBRALTAR, ESTONIA,
PARIS AND SYDNEY

FOUNDED 1997 AND NOW EMPLOYS OVER 1200 PEOPLE

OUR VALUES ARE FRIENDLY, PASSIONATE AND EXPERT

WE'RE A WORLD CLASS TECH BUSINESS - WE HANDLE MORE THAN 400
TRANSACTIONS PER SECOND. DURING 2015 WE MADE OVER 1200 RELEASES
TO OUR LIVE ENVIRONMENT WITH ZERO HOURS OF PLANNED DOWNTIME

CAREERS: [HTTPS://TEAM.UNIBET.COM](https://team.unibet.com)





SOCIALLY RESPONSIBLE OPERATOR



SPORTS BETTING OPERATOR



SOCIALLY RESPONSIBLE OPERATOR

THE PROBLEM WE HAVE TODAY

- ⌚ Databases are getting larger
 - ⌚ Backing up tens of TB will take days!
- ⌚ RTO is getting smaller
 - ⌚ Fully restoring tens of TB can also take equally as long!
- ⌚ Are we actually able to restore our backups?

WHAT SOLUTIONS ARE THERE?

- u Storage snapshot backup
 - RMAN supports it since 12c
 - You are not protected against storage failures
 - Still need a separate archivelog backup
- u Data Guard/DBVisit standby
 - Disaster recovery solution, not a backup
 - No history
- u Take a backup from standby database
 - Just offloading the work
 - Backup and recovery still take the same time

WHY DO YOU NEED BACKUPS?

- History
 - Recover from logical/human errors
- Long term archival of some fixed state
- Protection against main storage failure/corruption

Even if you have standby database you still need backups!

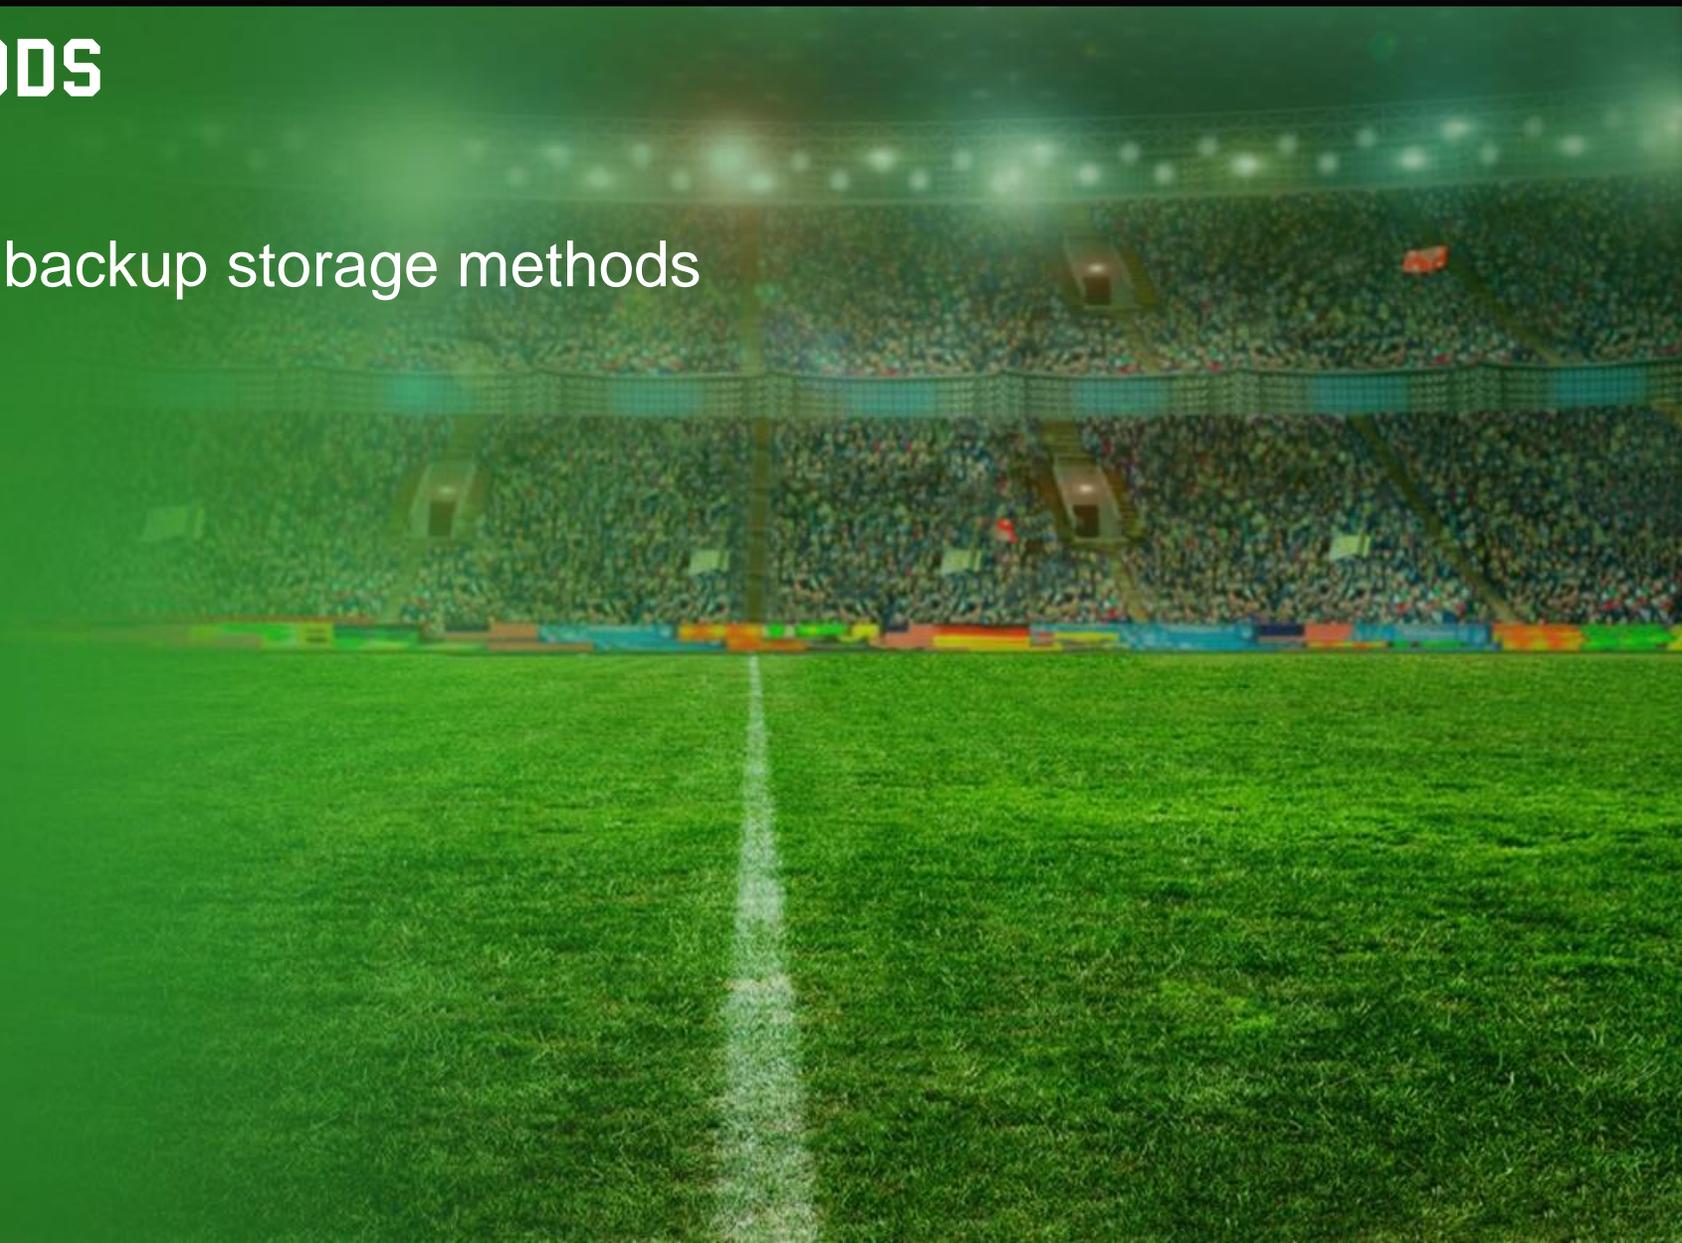
WOULD IT BE GREAT, IF RMAN ALREADY HAD A SOLUTION?

RMAN – Oracle Recovery Manager

- ⊍ Fully supported by Oracle
- ⊍ Basic essential DBA knowledge
- ⊍ No GUI, pure “simple” CLI that is easily scriptable
- ⊍ 3rd party backup tools just execute RMAN internally
 - Just skip the 3rd party agent 😊

RMAN BACKUP METHODS

- ① RMAN has two types of backup storage methods
 - ❖ Backup sets
 - ❖ Image copies



RMAN BACKUP METHODS

BACKUP SETS

DEFAULT OPTION

ORACLE PROPRIETARY BACKUP
FORMAT

CAN STORE DATAFILE, ARCHIVELOG,
CONTROLFILE, SPFILE

INCREMENTAL BACKUPS

CAN BE COMPRESSED, ENCRYPTED

VARIOUS OPTIMIZATIONS

HISTORY AND AUTOMATIC RETENTION

MULTISECTION

IMAGE COPIES

JUST EXACT COPIES OF
DATAFILES, ARCHIVED LOG
FILES OR CONTROL FILE

NO COMPRESSION

NO ENCRYPTION

NO HISTORY

NO OPTIMIZATIONS

CAN BE INCREMENTALLY
REFRESHED



DO WE HAVE A CLEAR WINNER?

- ❶ Compression/encryption is expensive
 - Oracle license per CPU core/socket
- ❷ Restoring incremental backup sets is slow
 - Restore full and apply incrementals
 - To make it faster take more full backups?



RESTORING TO IMAGE COPY

- ① RMAN> SWITCH DATABASE TO COPY ;
 - Just a quick dictionary update
- ① RMAN> RECOVER DATABASE ;
 - To apply archive logs

DONE 😊

WHAT RMAN FEATURES ARE SUPPORTED WITH IMAGE COPIES?

- ⓪ Block recovery
- ⓪ Duplicate database
- ⓪ Block change tracking to speed up incremental refresh
- ⓪ Skip read-only tablespaces during refresh
- ⓪ Not supported: Undo optimization

HOW TO CREATE IMAGE COPIES?

- BACKUP AS COPY DATABASE ;
- BACKUP AS COPY TABLESPACE SYSAUX ;
- BACKUP AS COPY DATAFILE 3 ;
- CONFIGURE DEVICE TYPE DISK PARALLELISM 2 BACKUP TYPE TO COPY ;

HOW TO INCREMENTALLY REFRESH IMAGE COPIES?

```
run {  
    BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG  
        'image_copy_backup' DATABASE;  
    RECOVER COPY OF DATABASE WITH TAG 'image_copy_backup';  
    DELETE NOPROMPT BACKUPSET TAG 'image_copy_backup';  
}  
run {  
    RECOVER COPY OF DATABASE WITH TAG 'image_copy_backup'  
        UNTIL TIME 'SYSDATE-2';  
    BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG  
        'image_copy_backup' DATABASE;  
}
```

ARE WE MISSING A LOT OF FEATURES?

- u Yes!
- u But storage today is much smarter than it used to be
 - ❖ It can snapshot and clone the snapshots
 - ❖ It can compress, deduplicate and encrypt

DON'T FORGET ARCHIVELOGS!

- Ⓢ Otherwise hot image copy is useless
- Ⓢ You can have multiple archivelog destinations in the DB
 - Mandatory and optional
- Ⓢ Store them together with image copies
 - Then snapshot will include them

SETTING ARCHIVE LOG DESTINATIONS

```
ALTER SYSTEM SET log_archive_dest_1=  
    'LOCATION=USE_DB_RECOVERY_FILE_DEST MANDATORY' SCOPE=both;
```

```
ALTER SYSTEM SET log_archive_dest_2=  
    'LOCATION=/nfs/backup/oem/archivelog/  
VALID_FOR=(ONLINE_LOGFILE,PRIMARY_ROLE)' SCOPE=both;
```

**PUTTING IT ALL TOGETHER
A SUGGESTED BACKUP
SCHEME**



PREPARE THE STORAGE AND DATABASES

- ④ Each database has its own dedicated backup filesystem
- ④ Database secondary optional archive log destination is in the same backup filesystem
 - ❖ `LOG_ARCHIVE_DEST_2` parameter
- ④ Complete the initial full image copy backup

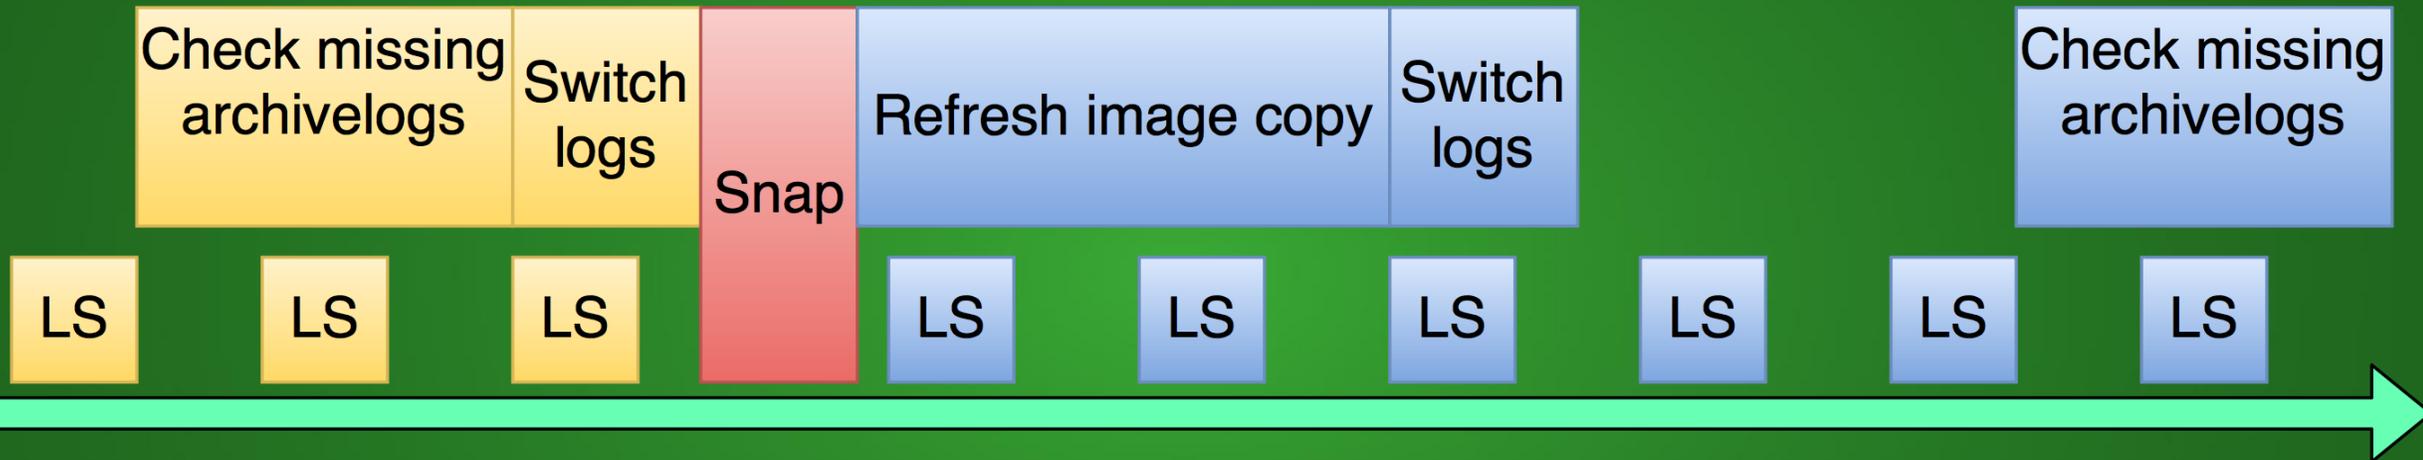
DAILY BACKUP

- ④ Check if there are any gaps in the optional archive log destination
 - ❖ A quick data dictionary lookup
- ④ Switch and archive all online log threads
 - ❖ **ALTER SYSTEM ARCHIVE LOG CURRENT;**
- ④ Create snapshot of the backup area
- ④ Incrementally refresh image copy
- ④ Switch and archive all online log threads

EXTRA CHECK FOR ARCHIVELOGS

- ④ You can run the check for missing archive logs multiple times per day
 - ❖ Remember, the destination was optional!
 - ❖ One missing archive log stops recovery
 - ❖ A quick data dictionary lookup
- ④ In case a missing log is found, copy it
 - ❖ **BACKUP FORCE AS COPY ARCHIVELOG SEQUENCE 12345 THREAD 2;**

HOW IT LOOKS LIKE



BACKUP RETENTION

```
[2015-09-30 00:01:34 UTC] total=144GB unique=20GB clones=0 valid
[2015-10-31 01:01:09 UTC] total=248GB unique=14GB clones=0 valid
[2015-11-23 08:22:34 UTC] total=299GB unique=62GB clones=0 valid
[2015-12-31 01:00:23 UTC] total=93GB unique=16GB clones=0 valid
[2016-01-31 01:00:26 UTC] total=95GB unique=13GB clones=0 valid
[2016-02-07 01:00:40 UTC] total=96GB unique=4GB clones=0 dropped
[2016-02-08 01:00:28 UTC] total=95GB unique=1GB clones=0 valid
[2016-02-09 01:00:20 UTC] total=95GB unique=2GB clones=0 valid
[2016-02-11 01:00:36 UTC] total=97GB unique=6GB clones=0 valid
[2016-02-12 01:00:21 UTC] total=97GB unique=5GB clones=0 valid
[2016-02-13 01:00:32 UTC] total=97GB unique=5GB clones=0 valid
[2016-02-14 01:00:23 UTC] total=97GB unique=1GB clones=0 valid
[2016-02-15 01:00:23 UTC] total=95GB unique=846MB clones=0 valid
[2016-02-16 01:00:20 UTC] total=95GB unique=1GB clones=0 valid
[2016-02-17 01:00:21 UTC] total=96GB unique=6GB clones=0 valid
```

BACKUP RETENTION – SNAPSHOT SIZE OVER TIME

```
[2016-03-31 00:44:58 UTC] total=2TB unique=424GB
[2016-04-30 00:45:57 UTC] total=2TB unique=244GB
[2016-05-31 00:45:37 UTC] total=2TB unique=214GB
[2016-06-29 00:46:41 UTC] total=3TB unique=246GB
[2016-07-31 00:48:50 UTC] total=3TB unique=234GB
[2016-08-31 01:34:45 UTC] total=3TB unique=150GB
...
[2016-09-29 00:50:38 UTC] total=3TB unique=30GB
[2016-09-30 00:49:18 UTC] total=3TB unique=30GB
[2016-10-01 00:50:39 UTC] total=3TB unique=30GB
[2016-10-02 00:52:16 UTC] total=3TB unique=25GB
[2016-10-03 00:49:33 UTC] total=3TB unique=24GB
[2016-10-04 00:49:58 UTC] total=3TB unique=21GB
[2016-10-05 00:51:02 UTC] total=3TB unique=22GB
```

BACKUP RETENTION – SNAPSHOT SIZE OVER TIME

```
[2016-03-31 01:44:01 UTC] total=10TB unique=1TB
[2016-04-30 01:43:29 UTC] total=12TB unique=2TB
[2016-05-31 01:42:35 UTC] total=12TB unique=2TB
[2016-06-29 01:44:44 UTC] total=14TB unique=3TB
[2016-07-31 01:45:05 UTC] total=15TB unique=3TB
[2016-08-31 01:43:04 UTC] total=18TB unique=6TB
...
[2016-09-30 01:45:16 UTC] total=16TB unique=398GB
[2016-10-01 01:44:02 UTC] total=16TB unique=422GB
[2016-10-02 01:43:28 UTC] total=16TB unique=434GB
[2016-10-03 01:43:11 UTC] total=16TB unique=415GB
[2016-10-04 01:45:11 UTC] total=15TB unique=372GB
[2016-10-05 01:44:23 UTC] total=15TB unique=347GB
[2016-10-06 01:43:06 UTC] total=15TB unique=346GB
```

PREREQUISITES

- ④ Smart capable storage!
- ④ NFS
 - Direct NFS enabled on database side
 - Enables the use of soft-mounted volumes for backup
 - If not in use, NFS must be hard mounted
 - “ORA-27054: NFS not mounted with proper option”
 - Optional, but highly recommended for high availability
- ④ Optional: Block change tracking
 - Enterprise Edition feature

CHECKLIST FOR STORAGE CAPABILITIES

- u Snapshots
- u Creating clones from snapshots
 - u Thin clones
- u Compression
 - u Maybe also deduplication
- u NFS
- u (REST?) API for remotely calling storage functions
- u Optional: encryption

LIST OF STORAGE SYSTEMS WITH REQUIRED CAPABILITIES

- u Oracle ZFS Storage Appliance
- u Netapp FAS8000
- u IBM Spectrum
- u HDS HNAS
- u Various open source products:
FreeNAS? Nexenta?
- u Audience: do you know more?

COMPRESSION RATIO

- u Examples are from Oracle ZFS Storage Appliance
- u Archivelogs
 - u Production (GZIP): 3.36x, 3.10x, 2.24x
 - u Test (GZIP): 4.11x
 - u Test (LZJB): 2.12x
- u Data files
 - u Production (GZIP + LZJB)*: 2.85x, 2.22x, 2.71x
 - u Test (GZIP): 7.73x
 - u Test (LZJB): 3.70x
 - u Test (RMAN compressed backupset): 15.63x
- u Redo per day / daily snapshot size **
 - u Production (GZIP): 2.7-5.4x, 14-16x, 3.0-4.2x

SCHEDULING

- u Each database can schedule it's own backup
 - u External job from DBMS_SCHEDULER
- u If backup fails, job will also fail
 - u Make sure backup scripts check for RMAN exit code
- u Scheduler job status is convenient for monitoring
- u In 12c DBMS_SCHEDULER can execute RMAN scripts directly

SCHEDULING - JENKINS

- u Popular tool for continuous integration and continuous delivery
- u Simple GUI gives good overview of all backup jobs
- u Jobs are executed over SSH
- u Email alerts
- u Full console output
- u Free!
- u <https://jenkins.io/>

The screenshot shows the Jenkins web interface for a 'Backup' job. The main area displays a table of job runs with the following columns: S (Status), W (Weather icon), Name, Last Success, Last Failure, Last Duration, and Next Launch. The table lists several jobs including 'Autorestore SWE1OC1', 'Autorestore SWE1OC2', 'Backup ARCHPT01', 'Backup CDB2', 'Backup PTESTD1', 'Backup QADB01RAC', 'Backup QASTD1', 'Backup SWE1ARCH', 'Backup SWE1CDB1', 'Backup SWE1IM', and 'Backup UNIQA01'. Each row includes a status icon (blue sphere for success, yellow sun for failure) and a link to the job's details.

S	W	Name ↓	Last Success	Last Failure	Last Duration	Next Launch
●	☁	Autorestore SWE1OC1	18 hr - #43	1 day 18 hr - #42	2 hr 27 min	
●	☁	Autorestore SWE1OC2	15 hr - #56	3 days 3 hr - #53	1 hr 4 min	
●	☀	Backup ARCHPT01	10 hr - #38	4 days 4 hr - #33	1 min 19 sec	
●	☀	Backup CDB2	3 hr 35 min - #44	6 hr 57 min - #43	9 min 21 sec	
●	☀	Backup PTESTD1	5 hr 31 min - #40	4 days 5 hr - #35	7 min 56 sec	
●	☀	Backup QADB01RAC	8 hr 50 min - #44	4 days 8 hr - #39	7 min 12 sec	
●	☀	Backup QASTD1	13 hr - #39	1 mo 0 days - #8	3 min 26 sec	
●	☀	Backup SWE1ARCH	7 hr 54 min - #39	4 days 7 hr - #34	2 min 29 sec	
●	☀	Backup SWE1CDB1	14 hr - #42	4 days 14 hr - #37	12 min	
●	☀	Backup SWE1IM	10 hr - #46	4 days 10 hr - #41	4 min 34 sec	
●	☀	Backup UNIQA01	11 hr - #38	4 days 11 hr - #33	1 min 38 sec	

Build Queue: No builds in the queue.

Build Executor Status:

- 1 [Backup SWE1CDB1](#) #43
- 2 Idle
- 3 Idle
- 4 Idle
- 5 Idle
- 6 Idle
- 7 Idle
- 8 [Backup CDB2](#) #45

Next Executions:

- [Autorestore SWE1OC1](#) 06/10/2016 18:16
- [Autorestore SWE1OC2](#) 06/10/2016 21:05
- [Backup SWE1CDB1](#) 06/10/2016 22:54
- [Backup QASTD1](#) 06/10/2016 23:26

RESTORE PROCESS FLOW

- u Clone snapshot you need
 - u Or the latest area
- u Mount the clone and catalog the files
- u Switch datafiles
- u Recover – apply archive logs up to the requested point in time
- u You have now an option to open the database
- u Move datafiles to production storage
 - u Move only SYSTEM, open, and then take each file one by one
 - u Before opening, move all datafiles
 - u Undo tablespaces can be recreated
 - u 12c can move datafiles online, including SYSTEM

ADVANTAGES

- ⓪ Never take full backup again
- ⓪ You can also skip archive log backups
- ⓪ Each snapshot is self-contained
- ⓪ Retention time is not limited to continuous recovery window
- ⓪ Quicker restore process

DISADVANTAGES

- ⓪ Needs more features from storage
 - ⓪ Does not work with simple JBOD
 - ⓪ Need to make sure storage has the needed features before buying it
- ⓪ Recovery needs more OS and storage knowledge
- ⓪ If the required recovery PIT is during the time image copy was being refreshed, need extra steps to recover.

RECOVERY POINT AND RECOVERY TIME ESTIMATIONS

- ⌚ Recovery Point – how much data can you lose
 - ⌚ Set ARCHIVE_LAG_TARGET parameter
- ⌚ Recovery time
 - ⌚ The time needed to apply archivelogs
 - ⌚ Up to 12.2 it is still single threaded and can take hours

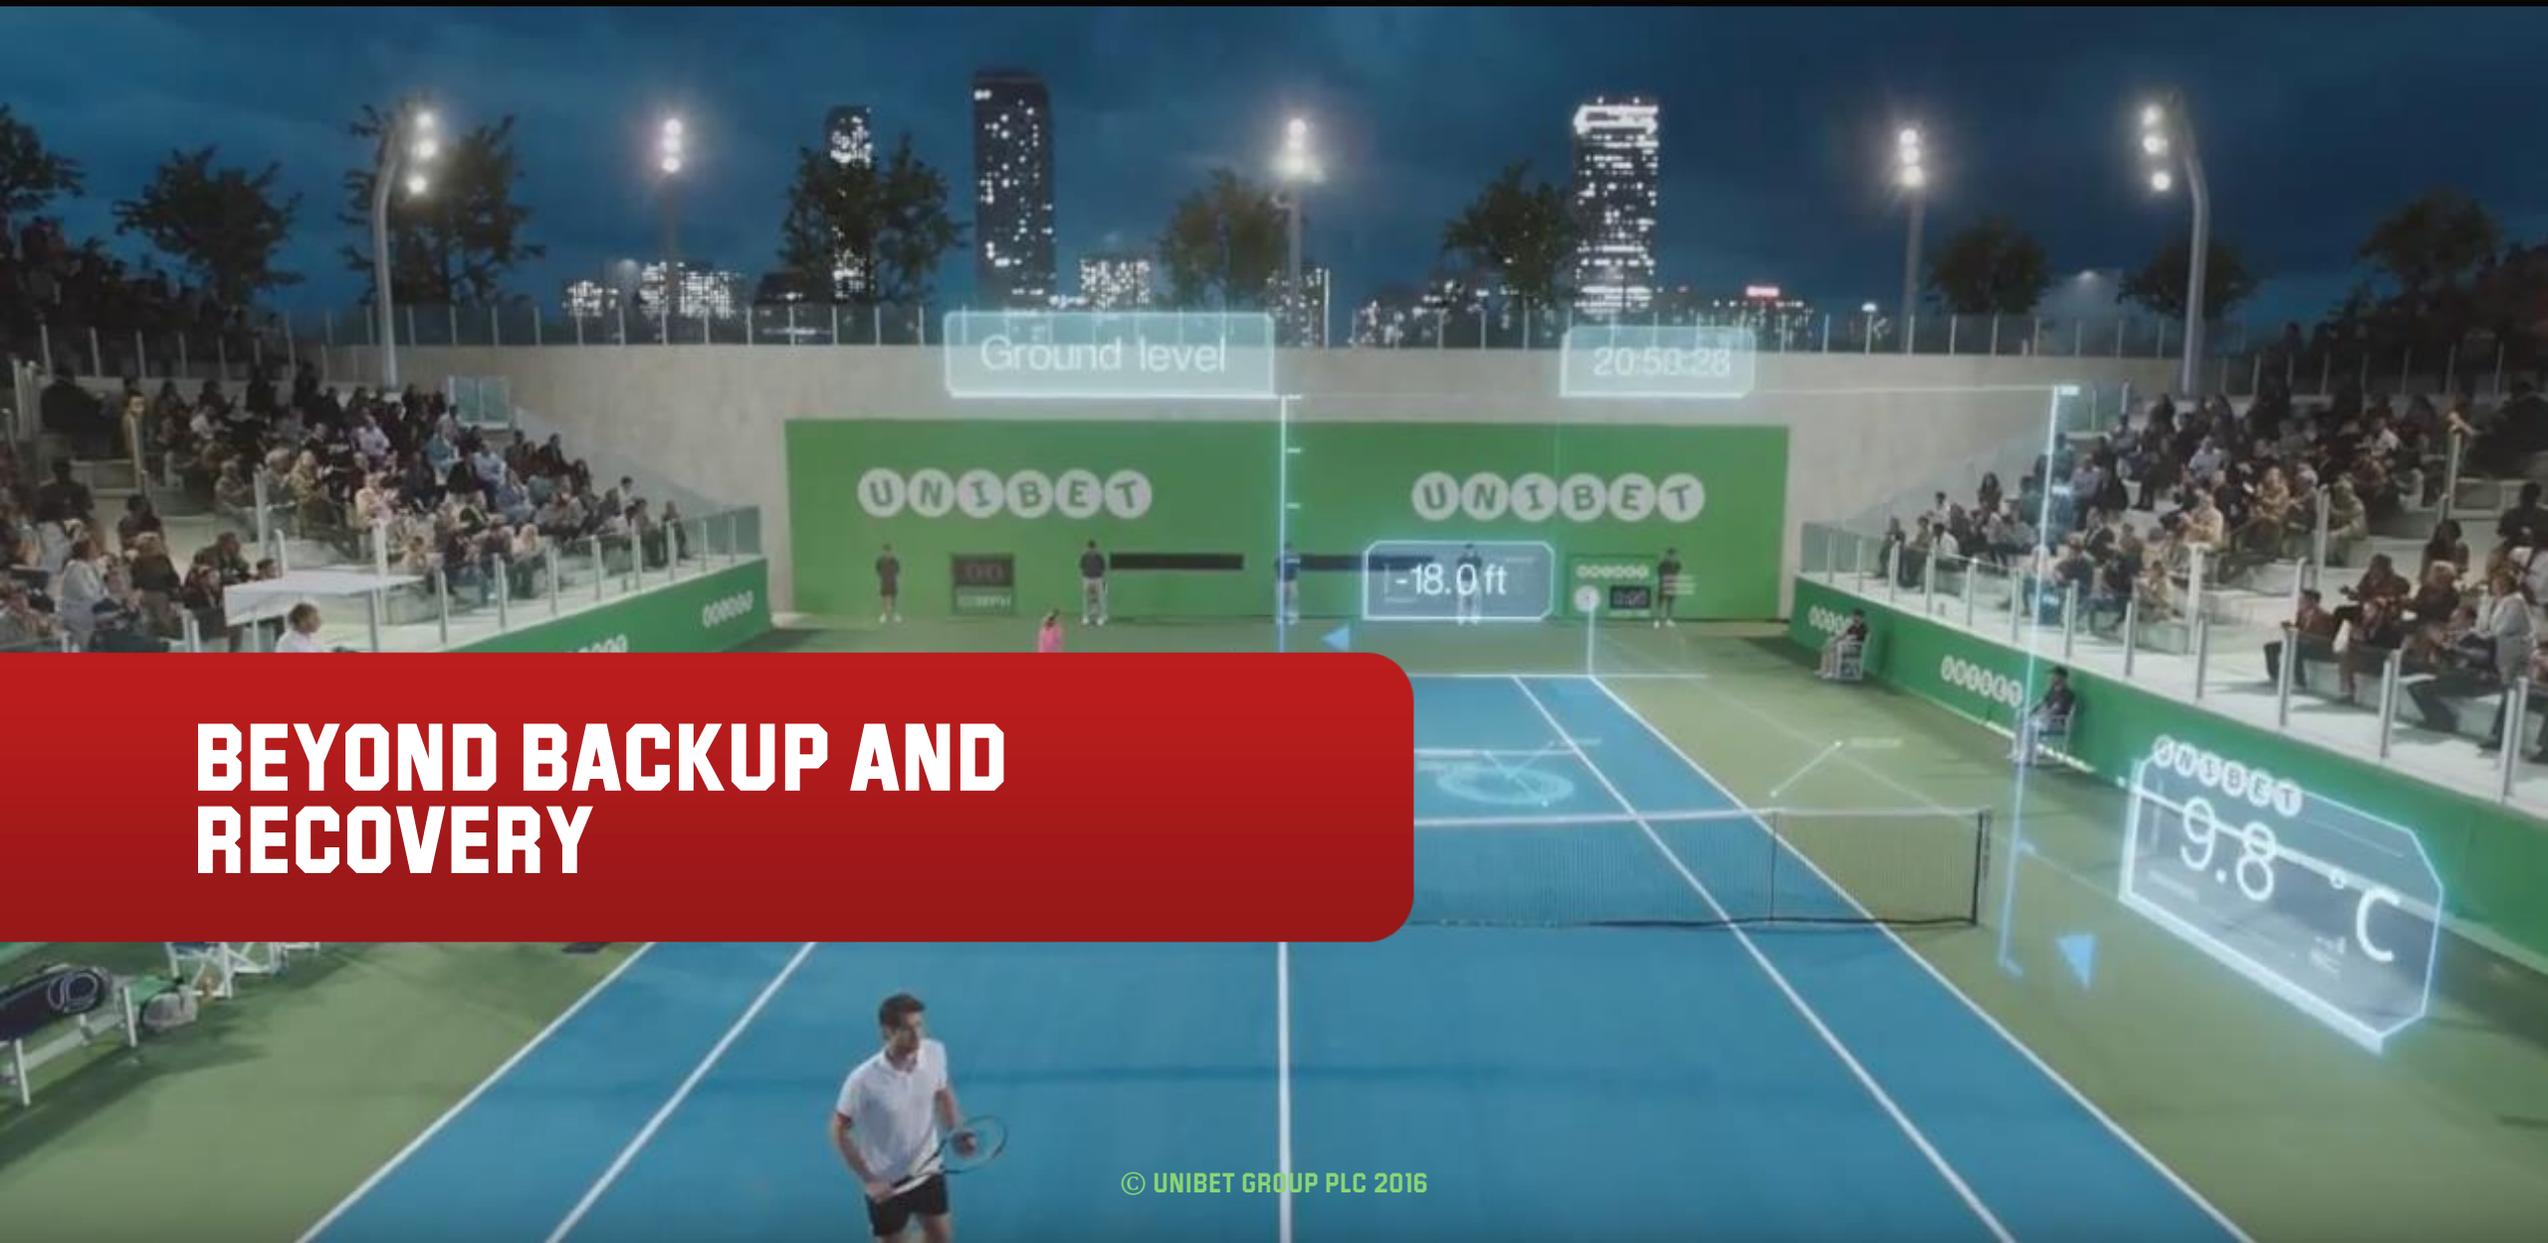
If your RPO and/or RTO are much higher than this, you need a disaster recovery solution IN ADDITION to backups

DO YOU STILL NEED RMAN CATALOG DB?

- ⌚ RMAN catalog is used for
 - ⌚ Keeping backup catalog longer and KEEP backup records
 - ⌚ Storing scripts
 - ⌚ Making some advanced recovery scenarios easier
- ⌚ RMAN catalog is only needed
 - ⌚ If you have read-only tablespaces
 - ⌚ It could be useful for centralised backup reporting

PRACTICAL TIPS

- ⓪ RMAN does not automatically delete incremental backups used for image copy refresh.
 - ⓪ Need to delete them manually
- ⓪ Not backing up archivelogs means they are not automatically deleted by policy.
- ⓪ Error checking
 - ⓪ All scripts need to check RMAN and SQLPlus exit codes for failures
- ⓪ Not all image copy is available for an incremental refresh
 - ⓪ `CATALOG DATAFILECOPY '/nfs/backup/orcl/data.dbf'`
`LEVEL 0 TAG 'image_copy_backup';`



BEYOND BACKUP AND RECOVERY

AUTOMATIC BACKUP VERIFICATION

- ⌚ You have to make restore tests
 - ⌚ Backup media can be corrupted
 - ⌚ Archivelog may be damaged?
 - ⌚ Deduplication database damaged?
- ⌚ This can be automated
 - ⌚ Daily recover
 - ⌚ Monthly full validation?
- ⌚ Requires extra database node to be licensed
- ⌚ Do not run automatic restores on primary database nodes, clusters or servers that have access to primary DB storage!

DATABASE VIRTUALISATION

- ⓪ Reminder: Each snapshot is a self-contained KEEP backup
- ⓪ It is possible to clone a snapshot as many times as needed
 - ⓪ Give developers and testers a fresh production database copy
 - ⓪ Minimal storage overhead

HOW DOES IT COMPARE TO ORACLE ZDLRA APPLIANCE?

- u + No need to take a full backup again
- u + No need to restore incremental backups – „synthetic full“ backups
- u -/+ No data loss
 - u - Cannot do no data loss, but remember ARCHIVE_LAG_TARGET
 - u + No need to take archivelog backups separately
- u + Nice GUI (using Jenkins)
- u - Integration with OEM for centralized control and monitoring
 - u Possible with some programming skills 😊
- u +/- Constant internal backup checks
- u - Push backups to tape
- u +/- Replicate backups to remote site
- u NB! ZDLRA cannot create KEEP backups internally!

SUMMARY



HOW CAN I PLAY WITH THIS?

- u Download Oracle ZFS SA Simulator
 - u It's a VirtualBox image
 - u <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/sun-simulator-1368816.html>
- u Download scripts to fully manage image copy backups with snapshots
 - u Includes autorestore script!
 - u Currently supports ZFS SA, but Netapp is coming next
 - u Easy to extend with other storage technologies
 - u <https://github.com/unibet/oracle-imagecopy-backup>
 - u Free to use, Apache 2.0 license



FINAL THOUGHTS

- ⓪ You still need backups, standby database does not save you from all disasters.
- ⓪ RMAN is a very powerful tool and every DBA must know it.
- ⓪ Using 3rd party backup tool does not release you from knowing RMAN 😊

Q&A



U N I B E T Group