

Implementing incremental forever strategy for Oracle database backups

Ilmar Kerm
@ilmarkerm

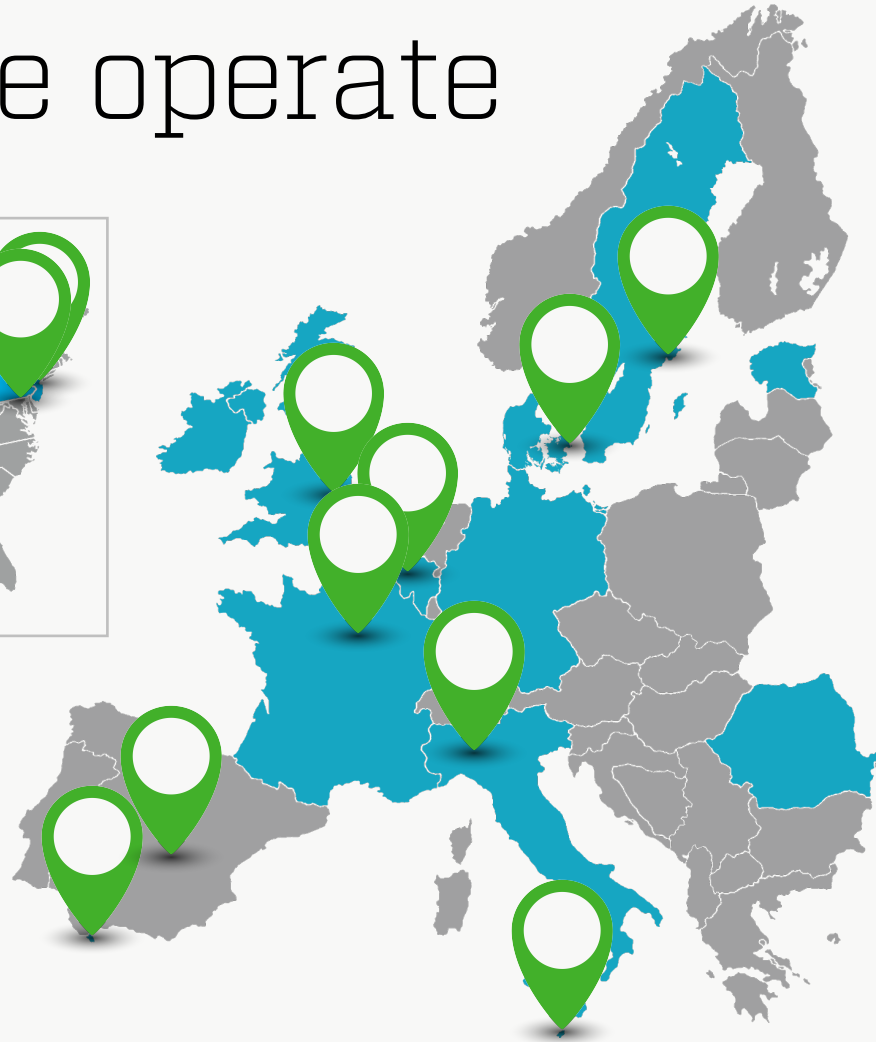

Our brands

A portfolio of
brands within
sportsbook,
casino and
games





Where we operate



Our Offices

Antwerp | Copenhagen | Darwin | Gibraltar | London | Madrid
Malta | Milan | New Jersey | New York | Paris | Stockholm | Sydney

Our Licenses

Australia | Belgium | Denmark | Estonia | France | Germany | Gibraltar | Ireland
Italy | Malta | New Jersey | Pennsylvania | Romania | Sweden | United Kingdom



Outline

- Who am I
- Problem with backups
- Implementing image copy backups and how to recover
- Beyond recovery
- Comparison with Oracle Zero Data Loss Recovery Appliance
- Free code!!



WHOAMI: Ilmar Kerm

- Senior database administrator in Kindred Group
- Working in IT since 2000
- Working with Oracle database since 2003
- Working at Kindred since 2007/2015
- From LAMP developer to Oracle DBA
- Oracle ACE

Blog: <https://ilmarkerm.eu>

Email: ilmar.kerm@kindredgroup.com

Twitter: [@ilmarkerm](https://twitter.com/ilmarkerm)



The problem with backups





The problems we have today with backups

- 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!
- Backups are rarely used
- Are we actually able to restore our backups?



What solutions are there?

- Storage snapshot backup
 - RMAN supports it since 12c
 - You are not protected against storage failures
 - Still need a separate archivelog backup
- Data Guard/DBVisit standby
 - Disaster recovery solution, not a backup
 - No history
 - Standby database needs to be licensed
- 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
- A classic but **flawed response** to the question "**Do you have a backup?**" is "**We have something even better than a backup—replication!**" Replication provides many benefits, including locality of data and protection from a site-specific disaster, but it can't protect you from many sources of data loss. Datastores that automatically sync multiple replicas guarantee that a corrupt database row or errant delete are pushed to all of your copies, likely before you can isolate the problem.
 - Raymond Blum and Rhandeev Singh, Google Site Reliability Engineering, ch26



Even if you have
standby database you
still need backups!

RMAN - Recovery Manager for Oracle Database



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

- 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?

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



Restoring to image copy

- **RMAN> SWITCH DATABASE TO COPY;**
 - Just a quick dictionary update
- **RMAN> RECOVER DATABASE;**
 - To apply archivelogs
- DONE 😊



What RMAN features are supported with image copies?



- Block recovery
- Duplicate database
- Block change tracking to speed up incremental refresh
- Skip read-only tablespaces during refresh
- 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?

- Yes!
- 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 about archive logs!

- Otherwise hot image copy is useless
- You can have multiple archive log 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 storage and databases

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



Daily backup

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

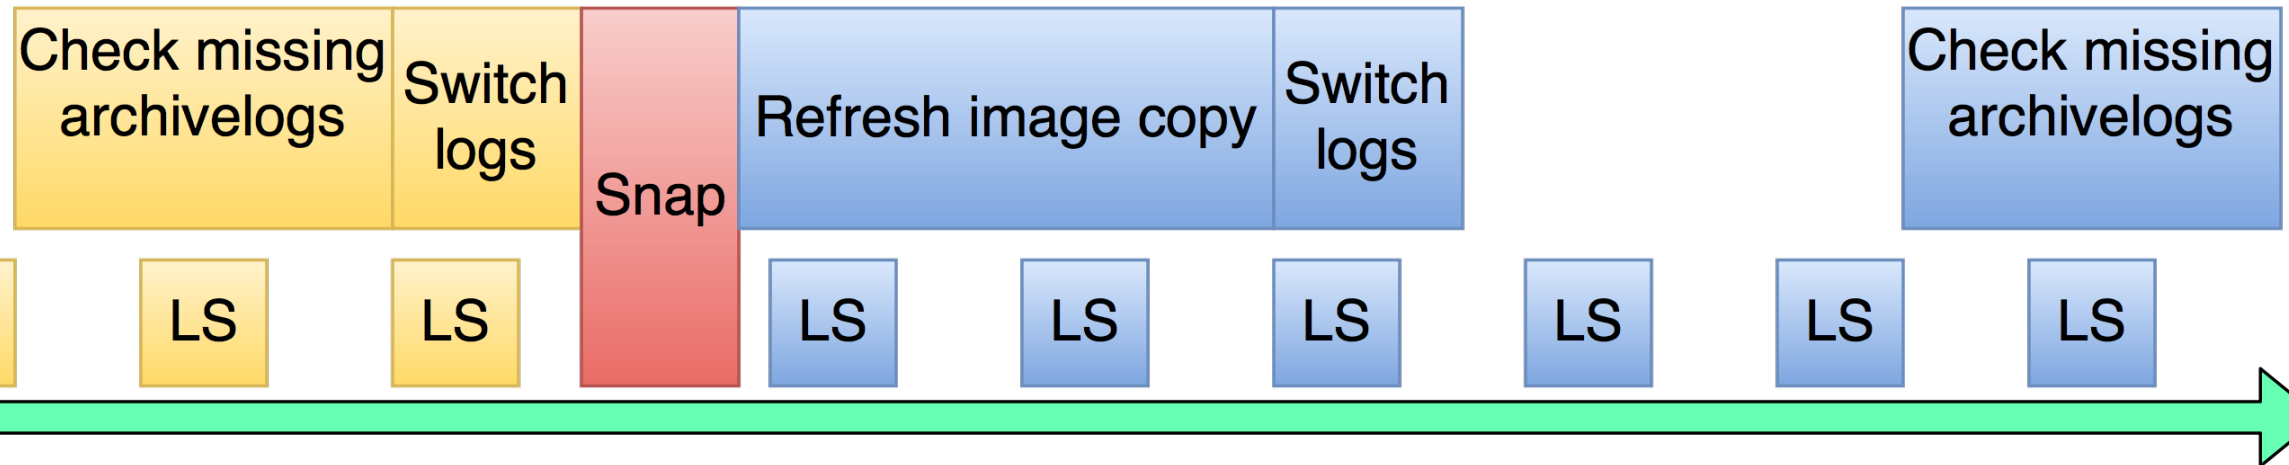


Extra check for archive logs

- 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
- 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
```


Prerequisites

- Smart capable storage!
- 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
- Optional: Block change tracking
 - Enterprise Edition feature



Checklist for storage capabilities

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



List of storage systems with required capabilities



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



Compression ratio

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



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
```



Scheduling

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

Scheduling - Jenkins



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

The screenshot shows the Jenkins web interface for the 'Backup' view. The left sidebar contains navigation links: New Item, People, Build History, Edit View, Delete View, Manage Jenkins, My Views, and Credentials. The main content area displays a table of backup jobs with columns for status, name, last success, last failure, last duration, and next launch. Below the table, there are sections for 'Build Queue' (showing no builds), 'Build Executor Status' (showing 8 executors, with 1 running 'Backup SWE1CDB1' and 7 idle), and 'Next Executions' (listing upcoming jobs and their scheduled times).

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	🔗

Icon: [S](#) [M](#) [L](#)

Legend [RSS for all](#) [RSS for failures](#) [RSS for just latest builds](#)

Build Queue
No builds in the queue.

Build Executor Status

Executor	Build	Status
1	Backup SWE1CDB1	#43 🚫
2	Idle	
3	Idle	
4	Idle	
5	Idle	
6	Idle	
7	Idle	
8	Backup CDB2	#45 🚫

Next Executions

Job	Next Execution
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

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



Advantages

- Never take full backup again
- You can also skip archive log backups
- Each snapshot is self-contained
- Retention time is not limited to continuous recovery window
- Quicker restore process
- Backups become “useful”
 - ... you can easily go beyond simple backup and restore process



Disadvantages

- Needs more features from storage
 - Does not work with simple JBOD
 - Need to make sure storage has the needed features before buying it
- Recovery needs more OS and storage knowledge
- 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
 - 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

- RMAN does not automatically delete incremental backups used for image copy refresh.
 - Need to delete them manually
- Not backing up archivelogs means they are not automatically deleted by policy.
- Error checking
 - All scripts need to check RMAN and SQLPlus exit codes for failures
- Not all image copy is available for an incremental refresh
 - **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 (unless max 4 times per year)
- Do not run automatic restores on primary database nodes, clusters or servers that have access to primary DB storage!



Database virtualisation

- Reminder: Each snapshot is a self-contained KEEP backup
- It is possible to clone a snapshot as many times as needed
 - Minimal storage overhead
- Some use cases
 - Give developers and testers a fresh production database copy
 - To create new Goldengate replica database
 - To create new Data Guard standby
 - Testing database upgrades
 - Transportable tablespace source

How does it compare to Oracle ZDLRA Appliance?



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

Free code and summary



Scripts to fully manage image copy backups with snapshots



- Backup, configuration, scheduling, retention
- Also includes automatic backup testing script!
- Currently supports Netapp, Oracle ZFS SA and SoftNAS storage systems
 - Easy to extend with other storage technologies
- <https://github.com/unibet/oracle-imagecopy-backup>
- Free to use, Apache 2.0 license





How can I play with this?

- Download Oracle ZFS SA Simulator
 - It's a VirtualBox image
 - <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/sun-simulator-1368816.html>





Thank You

ilmar.kerm@kindredgroup.com

@ilmarkerm