Backup configuration best practices

Jacek Wojcieszuk, CERN IT-DM

Distributed Database Operations Workshop

November 26th, 2009

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Outline

- Types of threats
- Available tools
- Desired configuration – Maximum Availability Architecture
- How to avoid data loss and minimize downtime without implementing MMA:
  - On tape backups
  - On disk backups
  - RMAN configuration and other hints
  - Backup validation
- Conclusions

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Types of threats

- Oracle instance failure
  - Usually due to a failure of an Oracle process
- Media failure
  - Disk failure, RAID controller failure, etc.
- Physical data corruption
- Human error
  - In most cases accidentally deleted/updated data
  - Caused either by a user or a DBA
- Disaster
  - Fire, flood, earthquake, plane crash, overvoltage, etc.

Available tools

- Oracle offers many tools that help to backup data and address failures:
  - Recovery Manager (RMAN)
  - Data Guard
  - Export/Import
  - Data Pump
  - Streams
- Oracle supports using OS and hardware features for taking backups
  - snapshots
  - cp command
- Each tool has its strong an weak points

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Oracle Maximum Availability Architecture (MAA)

- Oracle's best practices blueprint
- Goal: to achieve the optimal high availability architecture at the lowest cost and complexity
- Helps to minimize impact of different types of unplanned and planned downtimes
- Is based on such Oracle products/features like:
  - RAC
  - ASM
  - RMAN
  - Flashback
  - Data Guard
- http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN implementation of MAA

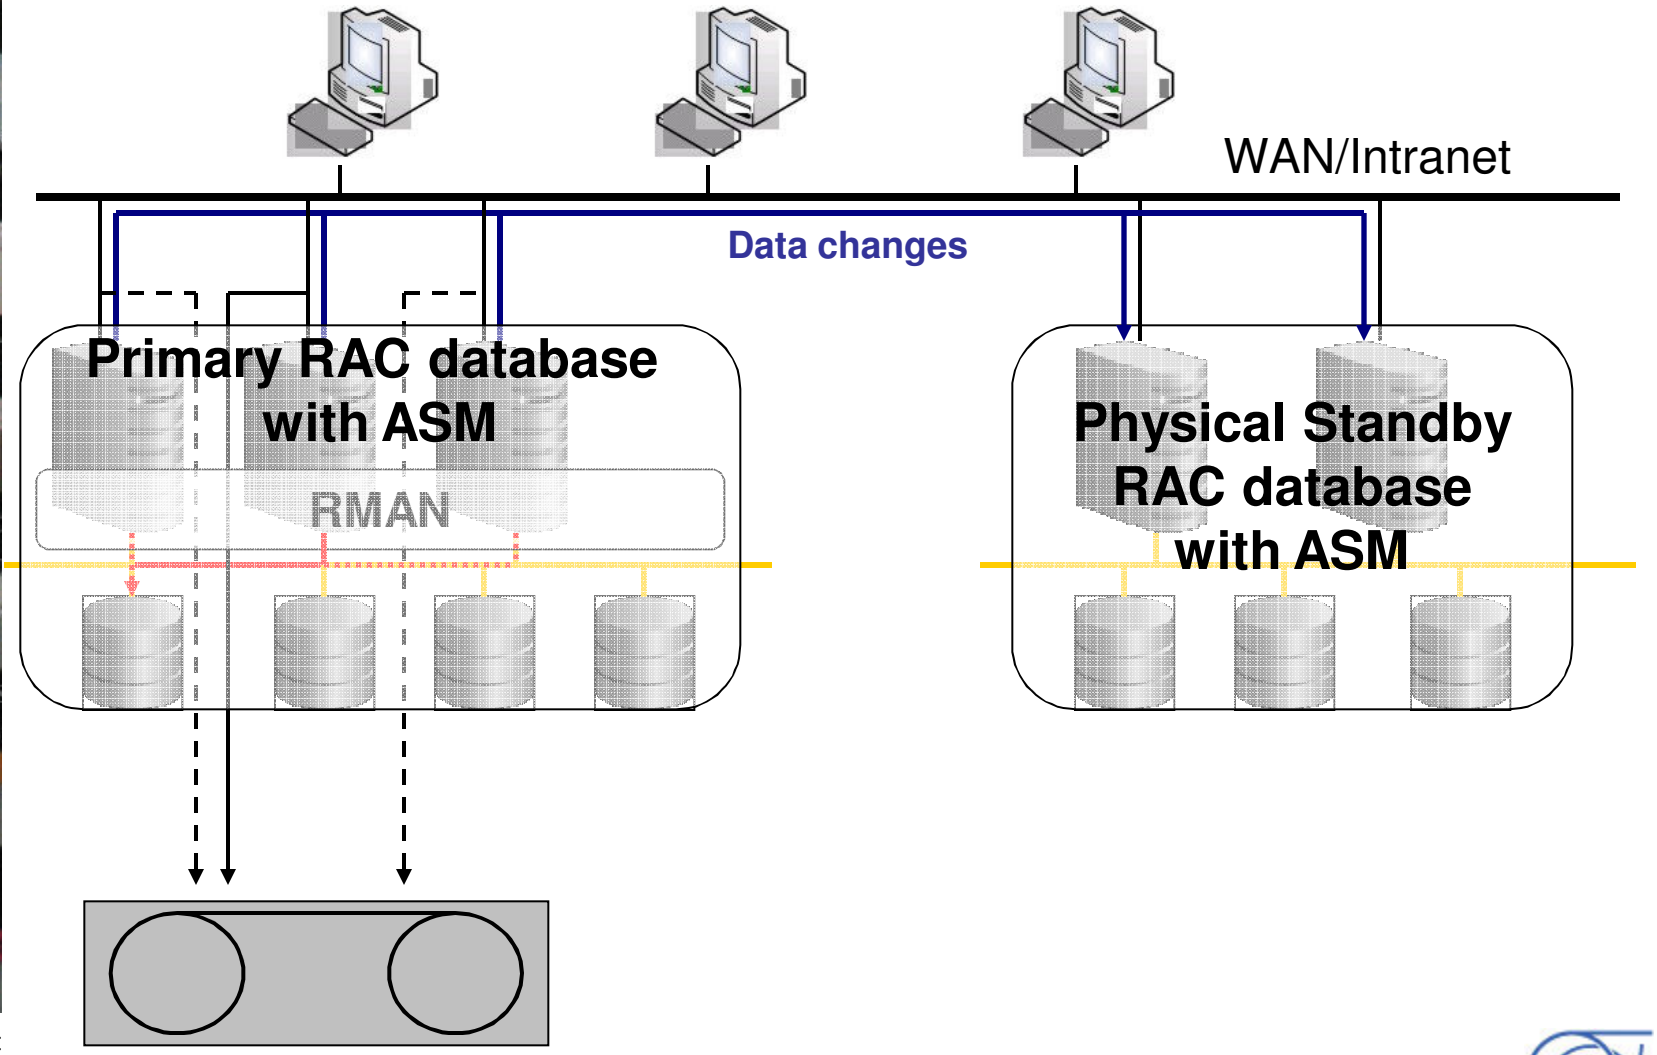

Failure handling

CERN IT Department

Failure	Recovery	Downtime
Oracle instance failure	Not needed - RAC keeps the database available	0
Media failure	Not needed - ASM keeps data healthy	0
Small physical data corruption	RMAN block media recovery using on-disk or on-tape backup	Database: 0 Affected application: few hours
Wide-range physical data corruption	• Switchover to the standby database • RMAN full database restore using on-disk backup	<1 hour with Data Guard <1 hour with on-disk backup
Human error	• RMAN + DataPump using on-disk backup • Standby DB + DataPump • RMAN + DataPump using on-tape backup	Database: 0 Affected application: usually few hours
Disaster	• Switchover to the standby database (if available) • RMAN full database restore using on-tape backups	<1 hour with Data Guard **Hours or days in case of restore from tapes**

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

When Standy DB cannot be put in place

- Tape backups to ensure recoverability
- On-disk image copy for availability
- Simplify and automatize backup procedures
- Verify your backups on regular basis
- Practice restore and recovery

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Tape backups

- Still the fundamental way of protecting databases against all types of failures
- Despite the associated cost they have many advantages:
  - Tapes can be easily taken offsite
  - Backups once properly stored on tapes are quite reliable
  - If configured properly can be very fast

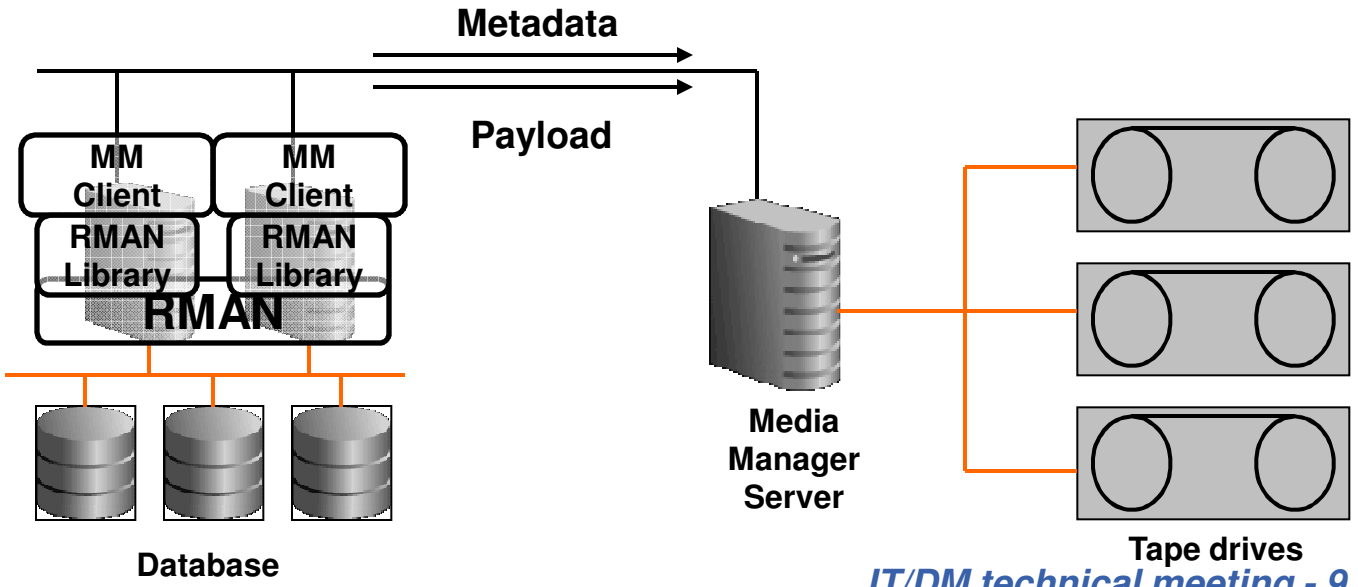

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Tape backups – recommended backup strategy

CERN IT Department

- Incremental backup strategy example:
  - Full backups every two weeks

```
backup force tag 'full_backup_tag' incremental level 0 check logical
    database plus archivelog;
```

  - Incremental cumulative every 3 days

```
backup force tag 'incr_backup_tag' incremental level 1 cumulative for
    recover of tag 'last_full_backup_tag' database plus archivelog;
```

  - Daily incremental differential backups

```
backup force tag 'incr_backup_tag' incremental level 1 for recover of
    tag 'last_full_backup_tag' database plus archivelog;
```

  - Hourly archivelog backups

```
backup tag 'archivelog_backup_tag' archivelog all;
```

  - Monthly automatic test restore
    - See tomorrow's presentation

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Super fast backups/recoveries to/from tapes

- Modern tape drives can archive data with the speed up to 200MB/s compressed
  - When backups send over 1 Gb network only part of this huge bandwith can be used
- Tivoli Storage Manager supports so-called LAN-free backup configuration where:
  - Backup data flows to tape drives directly over SAN
  - Media Management Server used only to register backups
  - Very good performance observed during tests (up to 400 MB/s for 2 RMAN channels and 2 tape drives)

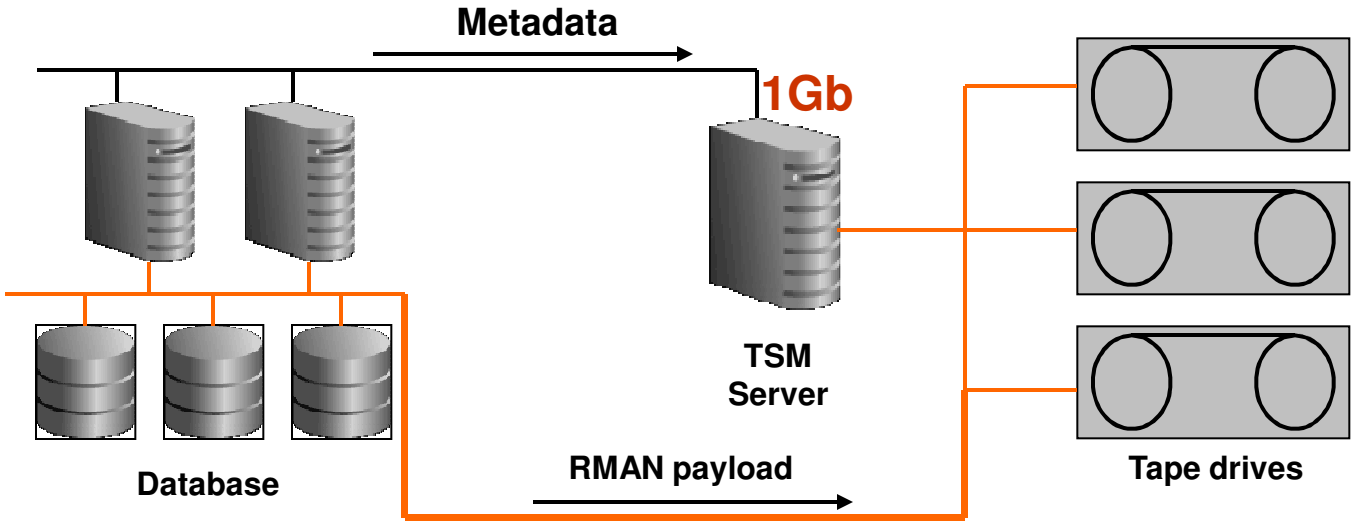

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

On disk image copy

- Even when running extremely fast tape backups are not optimal to handle certain types of failures:
  - Wide-scale physical corruption
  - Logical corruption
  - In both cases time to recover proportional to the database size
- On disk image copy can be very useful:
  - in case original datafiles are not usable anymore database can be switched to it
  - if lagging behind the database can be used to address logical corruptions
- **On disk image copy alone does not provide enough safety for the data**

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

On disk image copy – recommended strategy

- Image copy created in the beginning of database existence

  backup force tag ‘image_copy_tag’ as copy database;

- Daily incremental backups for recovery of copy to disk:

  backup force tag ‘backup_tag' incremental level 1 for recover of copy with tag ‘image_copy_tag ' database;

  – This may interfere with incremental backup strategy if implemented in parallel. More details later on

- Daily updates of the copy using incremental backups:

  recover copy of database with tag image_copy_tag‘’ until time ‘copy_lag’;

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Iterference between 2 described backup strategies

- Both discussed backup strategies cannot smootly coexist if configured as shown earlier

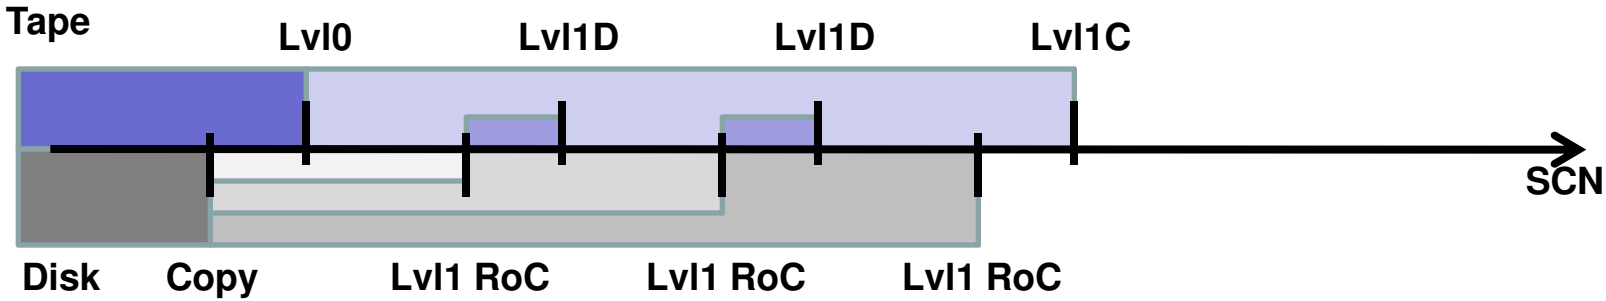

- To workaround this problem one can:
  - Use incremental backups sent to tapes to update the on-disk image copy
  - Take a level 1 backup for recovery of copy each time there is a full backup to tapes and store it on tapes too

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Backup&Recovery tuning - hints and tips

- Configuring RMAN properly is esential:
  - Helps to simplify backup scripts
    - The simplest are your backup scripts the more robust your backups will be
  - Using 'recovery window' for backup retention policy makes point-in-time recovery more predictible
  - Controlfile autobackups are usually highly desired
  - Enabling backup optimization helps to decrease archivelog backup volume
    - One has to use force option of the backup command to be sure that read-only datafiles are backed up
  - Default channel configuration makes things clearer
  - Setting 'maxopenfiles' property of an RMAN channel helps to optimize utilization of IO subsystem:

```
CONFIGURE CHANNEL DEVICE TYPE SBT MAXOPENFILES 4;
```

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Exempalry RMAN config

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 62 DAYS;
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE DEFAULT DEVICE TYPE TO 'SBT';
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT TO '%F';
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F';
CONFIGURE DEVICE TYPE 'SBT' PARALLELISM 2 BACKUP TYPE TO BACKUPSET;
CONFIGURE DEVICE TYPE DISK PARALLELISM 2 BACKUP TYPE TO COPY;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE SBT TO 1;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE 'SBT' TO 1;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1;
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' MAXOPENFILES 4 PARMS
    'BLKSIZE=1048576,ENV=(TDPO_OPTFILE=/opt/tivoli/tsm/client/oracle/bin64/tdpo.opt)';
CONFIGURE MAXSETSIZE TO 200 G;
CONFIGURE ENCRYPTION FOR DATABASE OFF;
CONFIGURE ENCRYPTION ALGORITHM 'AES128';
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE;
CONFIGURE SNAPSHOT CONTROLFILE NAME TO
    '/ORA/dbs01/oracle/product/10.2.0/rdbms/dbs/snapcf_d3r2.f';
```

Backup&Recovery tuning - hints and tips (2)

- Recovery catalog can be very helpful in case of complicated recoveries
  - Keeps track of all taken backups while the controlfile does it for certain number of days
  - If you can't afford it then better write down dbid of you DBs
- Block change tracking feature speeds up incremental backups by orders of magnitude
  - Make sure the `_bct_bitmaps_per_file` parameter is set to something bigger than the maximum number of incremental backups between 2 subsequent full backups

Don't trust your backups, don't trust RMAN

- Verify you backups on regular basis:
  - `restore validate` - the command reads contents of all needed backups without restoring anything
  - `restore preview` – shows which backups would be used for the recovery
- Keep checking if incremental backups were taken properly:
  - `report need backup days X` list datafiles that would need more than X days of recovery with archivelogs
- Setup automatic test point-in-time recoveries
  - The ultimate source of truth

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Don’t trust too much yourself, neither

- Practice recovery and journal all the steps and findings
  - Each recovery is different
  - RMAN syntax is not intuitive
  - There are still bugs here and there

Conclussions

- Backup&Recovery is a challenging task
- Typically relying on a single solution is not enough when high availability is important:
  - Different recovery solutions are optimal for handling different types of failures
- Only a synchronized standby DB gives full confidence
- Properly architectured, tuned and validated backups can give enough confidence if standby DB not feasible

Thank you

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it