

# DB

# Database Services

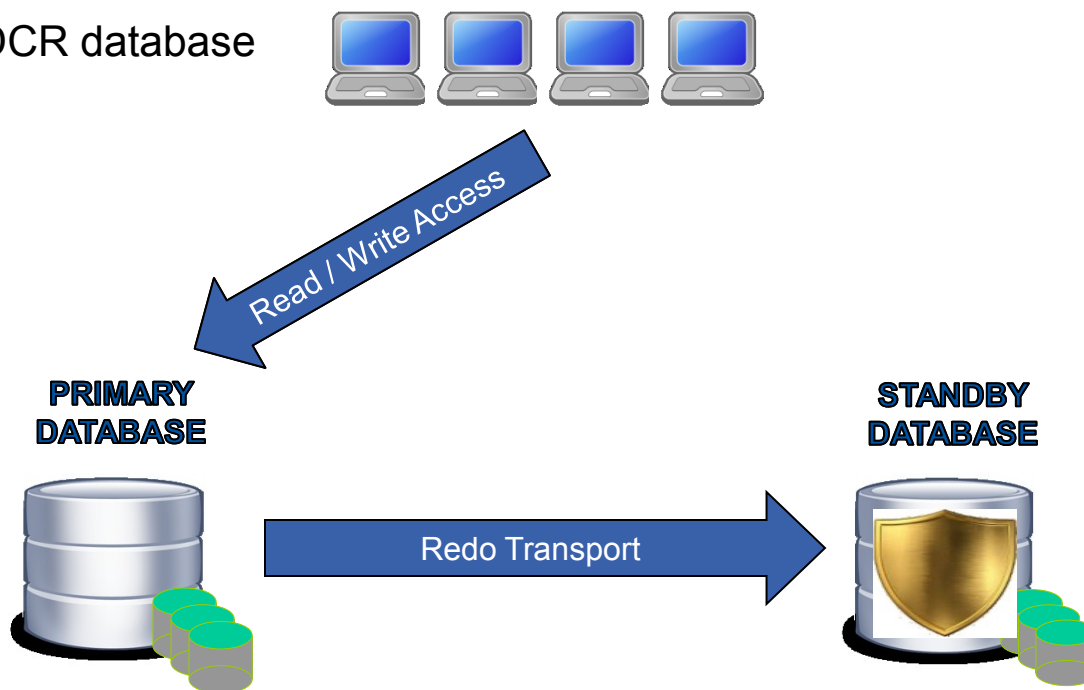
CERN  
IT  
Department

## Atlas standby database tests February 2011

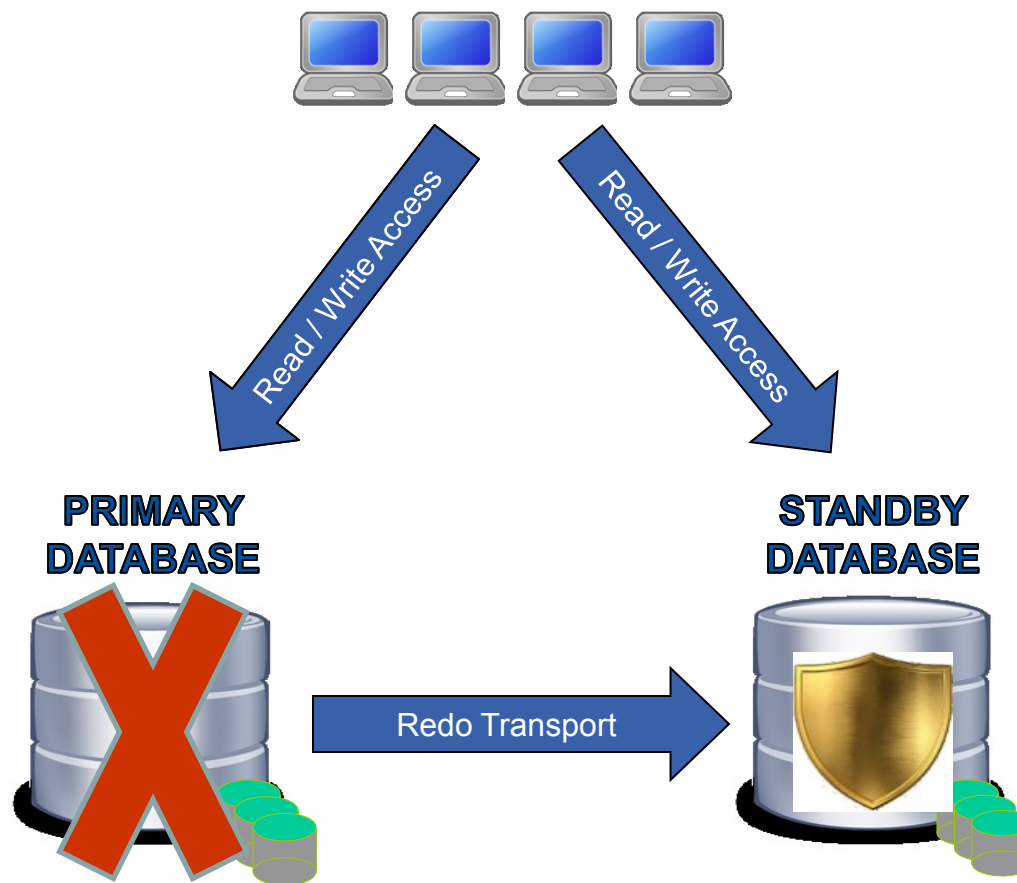
*Marcin Blaszczyk, IT-DB*  
[marcin.blaszczyk@cern.ch](mailto:marcin.blaszczyk@cern.ch)

- Standby databases for ATLAS
- Failover and Switchover
- Test of standby switchover – February 17<sup>th</sup> 2011
- Conclusions

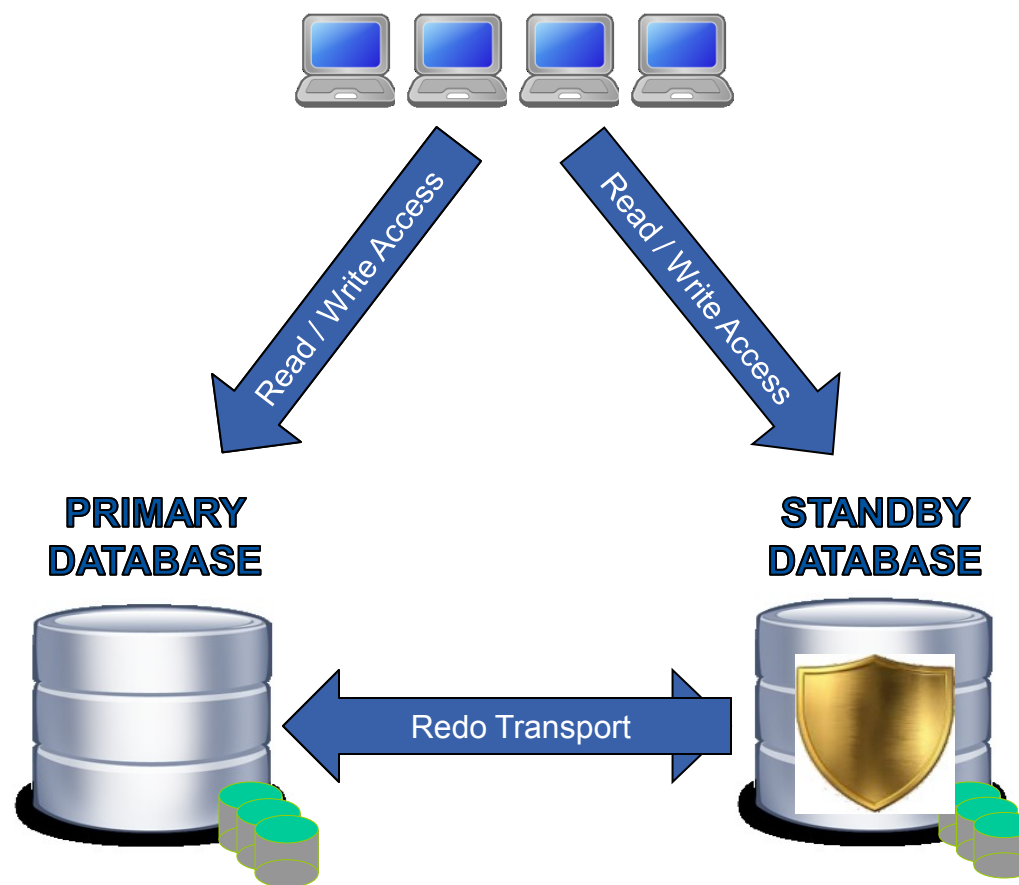
- Standby database is a **copy of production** database that can be used for disaster protection
  - **Dedicated** physical standby database for:
    - ATONR database
    - ATLR database
    - ADCR database



- All ATLAS standby databases:
  - Installed on new hardware provisioned in 2010
    - Quadcore servers and high-capacity disks
      - This has **increased resources on standby DBs** comparing to previous standby setups
      - Provided good compromise cost/performance in case of switchover operation
  - Are located in Safehost **outside CERN** campus
    - Reduce risk in case of disaster recovery
  - Asynchronous transport mode (no influence on primary database performance)







- Advantages:
  - Failover **minimizes downtime – it's faster** than full database recovery from backups
  - **no reconfiguration** is needed for users and applications
- Real life scenarios from other LHC experiments:
  - LHCb online database failover
    - August 2010
    - Reason: power cut in LHCb pit
  - CMS offline database failover
    - March 2011
    - Reason: Electrical issue with storages in CC

- Scenario
  - Tests performed on ATONR cluster to validate **disaster recovery scenario and infrastructure**
  - Coordinated by Luca Canali (IT-DB), Gancho Dimitrov, Florbela Tique Aires Viegas (ATLAS), Rainer Bartoldus (ATLAS Online DB coordinator)
  - Performed during technical stop on 17<sup>th</sup> of February 2011
- First phase:
  - Standby has been opened in **read only** mode for testing while primary database was running
  - Several tests performed regarding connectivity checks for online systems
- Second phase
  - **Full switchover**
  - All applications have been successfully reconnected to primary database while working on standby hardware
  - **Switch back** to original hardware



- General outcome:
  - Tests were **successful**, switchover scenario has been tested and validated
  - Standby database has been working fine **handling production load** for around 2 hours after switchover.
  - We are able to do switchover / failover in **~30 minutes**
- Issues encountered during test:
  - **DNS local caching** can caused some client-specific connectivity problems
    - RDB manager restart solves this problem
  - connection problems encountered for **COOL & CORAL reconnecting** after switch back to original hardware
    - We believe this was a one-off issue
    - Fixed with a service restart in that particular occurrence

- In case of Primary database lost:
  - It's feasible to perform a failover in **around 30 minutes**
    - Determining that failover is the only and best option can be time consuming
  - Due to asynchronous transport mode transaction lost is possible but limited to seconds
  - **No reconfiguration** on client side is needed
  - **Full database access** is guaranteed immediately after switchover
    - Global connection descriptors use aliases instead of physical machine names – **all changes on DNS level**
    - DNS local caching can cause some client-specific connectivity problems, connection checks needed after of Failover / Switchover

- *Luca Canali (IT-DB)*

<http://phydb.web.cern.ch/phydb/>

Thank You!



Questions?

[phydb.support@cern.ch](mailto:phydb.support@cern.ch)

[marcin.blaszczyk@cern.ch](mailto:marcin.blaszczyk@cern.ch)

[luca.canali@cern.ch](mailto:luca.canali@cern.ch)