

Data replication workshop

Summary

The workshop structure

- Overview and status update session (Tuesday)
 - ATLAS
 - data replication deployment
 - IT-DB
 - databases deployment
 - replication technologies overview
 - plan for replication technologies evaluation
 - T1s
- Administration panel (Wednesday)
 - Processes administration (installation, configuration)
 - Troubleshooting
 - Monitoring
 - Data consistency verification

Main Oracle data **real-time** replication solutions

- Database **block-level** synchronization – physical replica
 - (Active) Data Guard
- **SQL-level** synchronization
 - Streams (since 2002)
 - GoldenGate (since 90s, bought by Oracle in 2009)

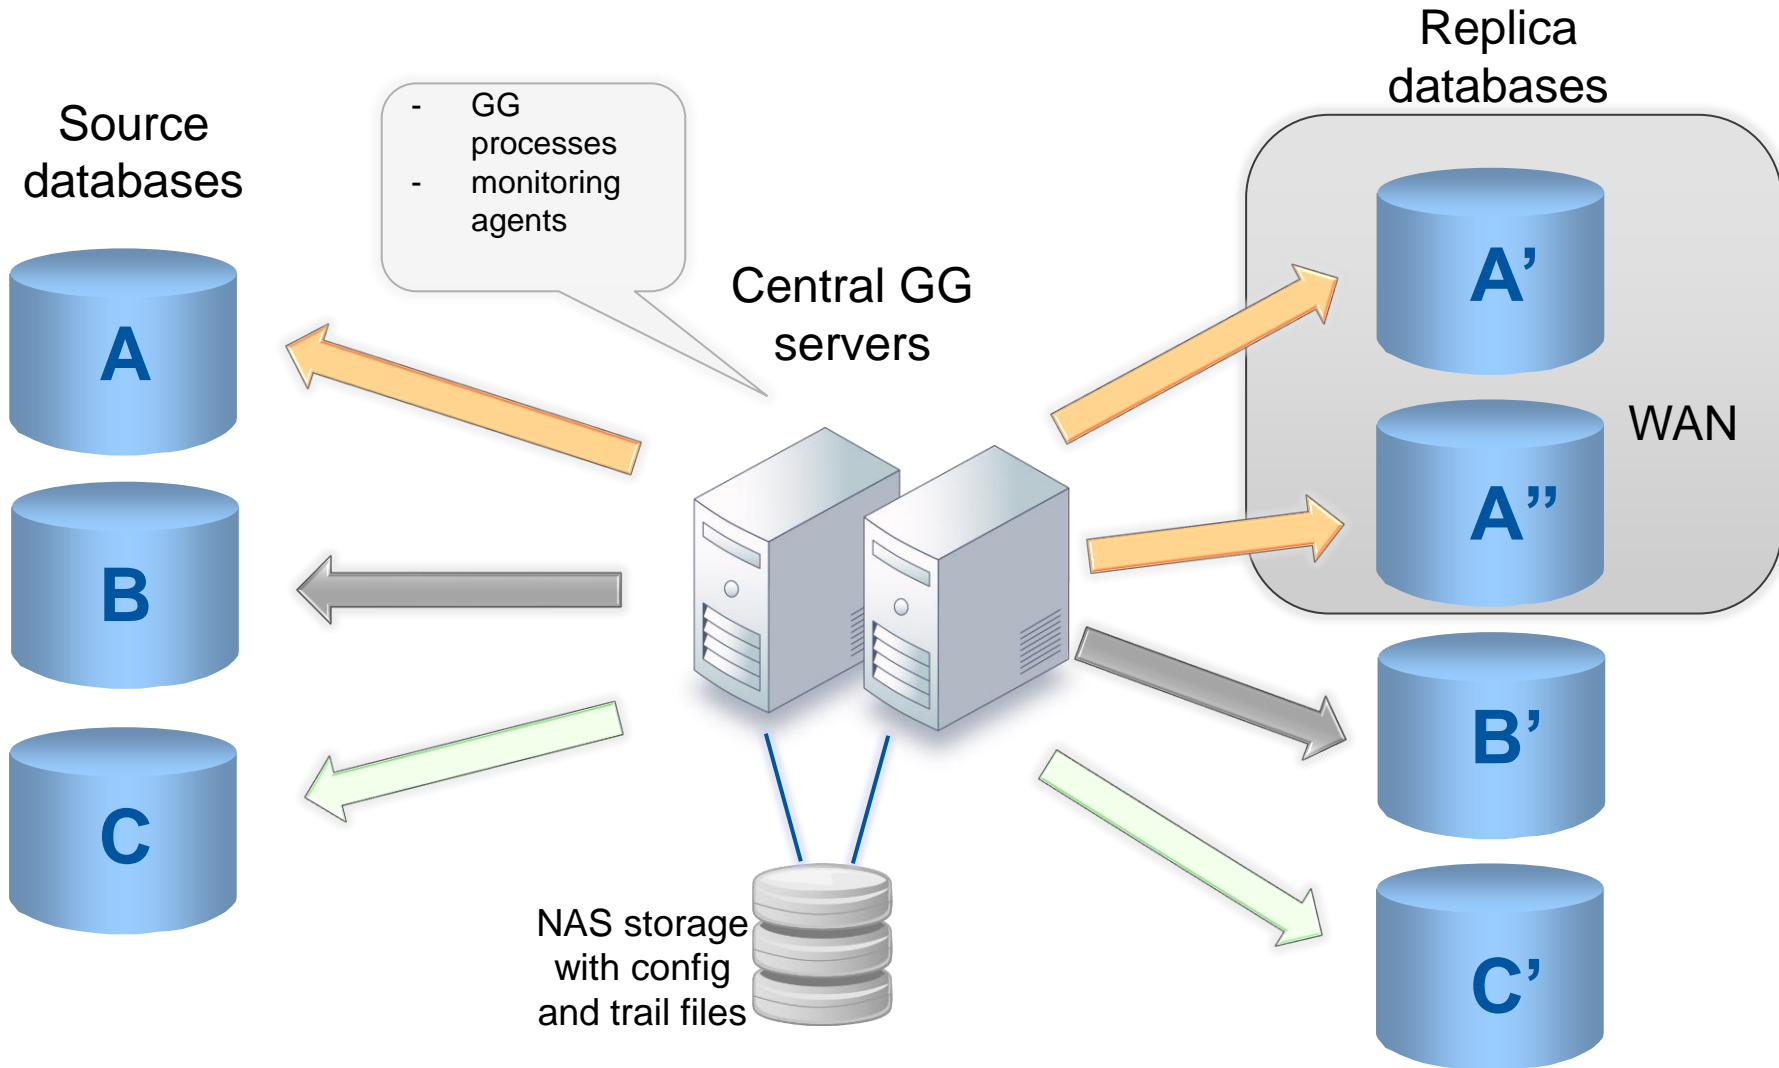
Replication service evolution

- **Conditions data:** preserve online->offline->T1s topology
 - Streams will be replaced with GoldenGate
 - New data folder for run2 (will be add to T1s replication later)
 - Amount of data for conditions data will not grow more than current rate
- **PVSS** replication will move from Streams to Active DataGuard
- **AMI**
 - GoldenGate instead of Streams
- **Muon calibration**
 - Decommissioning of Streams replica
 - Replication using home-made tools
 - New physical organization of data

T1s activities

- GoldenGate testing (T0->T1s)
 - Dedicated 11.2.0.4 database clusters
- Data Guard installation
 - Offloading primary databases
 - HA solution
- Minimal downtime migration using logical standby

Central GoldenGate installation



Progress of work

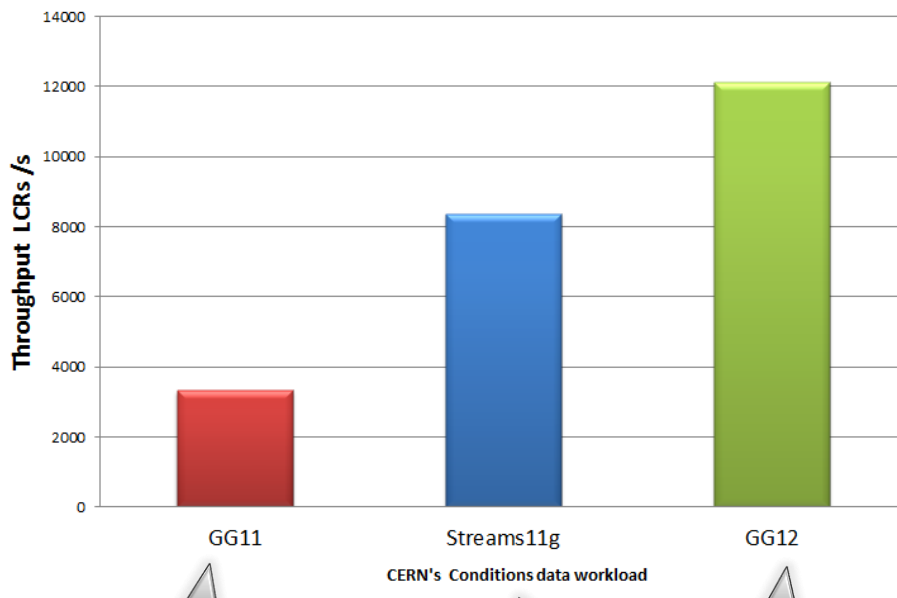
- Functional test
- Performance test
- Long distance tests (with T1s)
- Long term stability tests
- Tests with production data (ATLASCond)
- Optimize networking parameters
- Long term test with production data
 - Local and long distance



Performance measured @CERN

- Production workload (same hardware)

Average replication throughput for CERN production workload



Workload description:

- 5 days of ATLAS conditions data
- 675GB of redo volume
- 260k of transaction
- 18.9 M of row changes (LCRs)

No advantages from using **BatchSQL**
No apply **parallelism**

Bottleneck: IO & txns dependency

Bottleneck: IO & txns dependency

Action plan

1. Long term tests with production data (T0->T1s) **June**
2. ATONR & ATLR compatibility update **max: 1st week of July**
3. GoldenGate for online-offline (condDB) **week of 21st of July**
4. GoldenGate for offline-T1s **1st week of September**

Work to be done

- T0 and T1s (June – July)
 - Network tuning
 - Tests with production data
- T0
 - Installation of the GG production cluster (July)
- T1s
 - Copying data from production system
 - Firewall config for the GG production cluster



THANK YOU!

T1s

- BNL
 - Validation of migrations with physical standby
 - Hadoop for billing database
 - Plans upgrade to 11.2.0.4
- TRIUMF
 - DBs: Conditions, TAGs
 - Plans: RHEL6
- IN2P3
 - Pillar Axiom storage enables hybrid compression
 - 10x compression rate
 - Logical standby physical

ATLAS

- Intro, data set, and replication configuration
- PVSS -> ADG, move to IOT, sliding window 3 years
- Conditions -> preserver online->offline->T1s
 - New data folder for run2 (will be add to T1s replication later)
- AMI
 - GoldenGate instead of Streams
- Muon
 - No more streams
 - Replication using home-made tools
 - New physical organization
- Amount of data for conditions data will not grow more than current rate

IT update

- Software and Hardware storage evolution since 2010
 - Hardware: More CPU cores (8 cores and 32cores)
more RAM (16GB vs 128GB)
 - Software: 10.2.0.4 => 11.2.0.4 (and even 12.1.0.1)
- Active Data Guard widely used in production
-