

# Oracle Streams Replication to Tier-1

Dimitrov, Gancho  
Stonjek, Stefan  
Viegas, Florbela

# Replication Strategies at ATLAS

- Geometry database
  - Update frequency: several month
  - Replication with SQLite files
- Conditions database
  - Update frequency: seconds to hours
  - Replication with ORACLE streams
- Event data
  - Update frequency: 25 ns
  - Replication with DDM

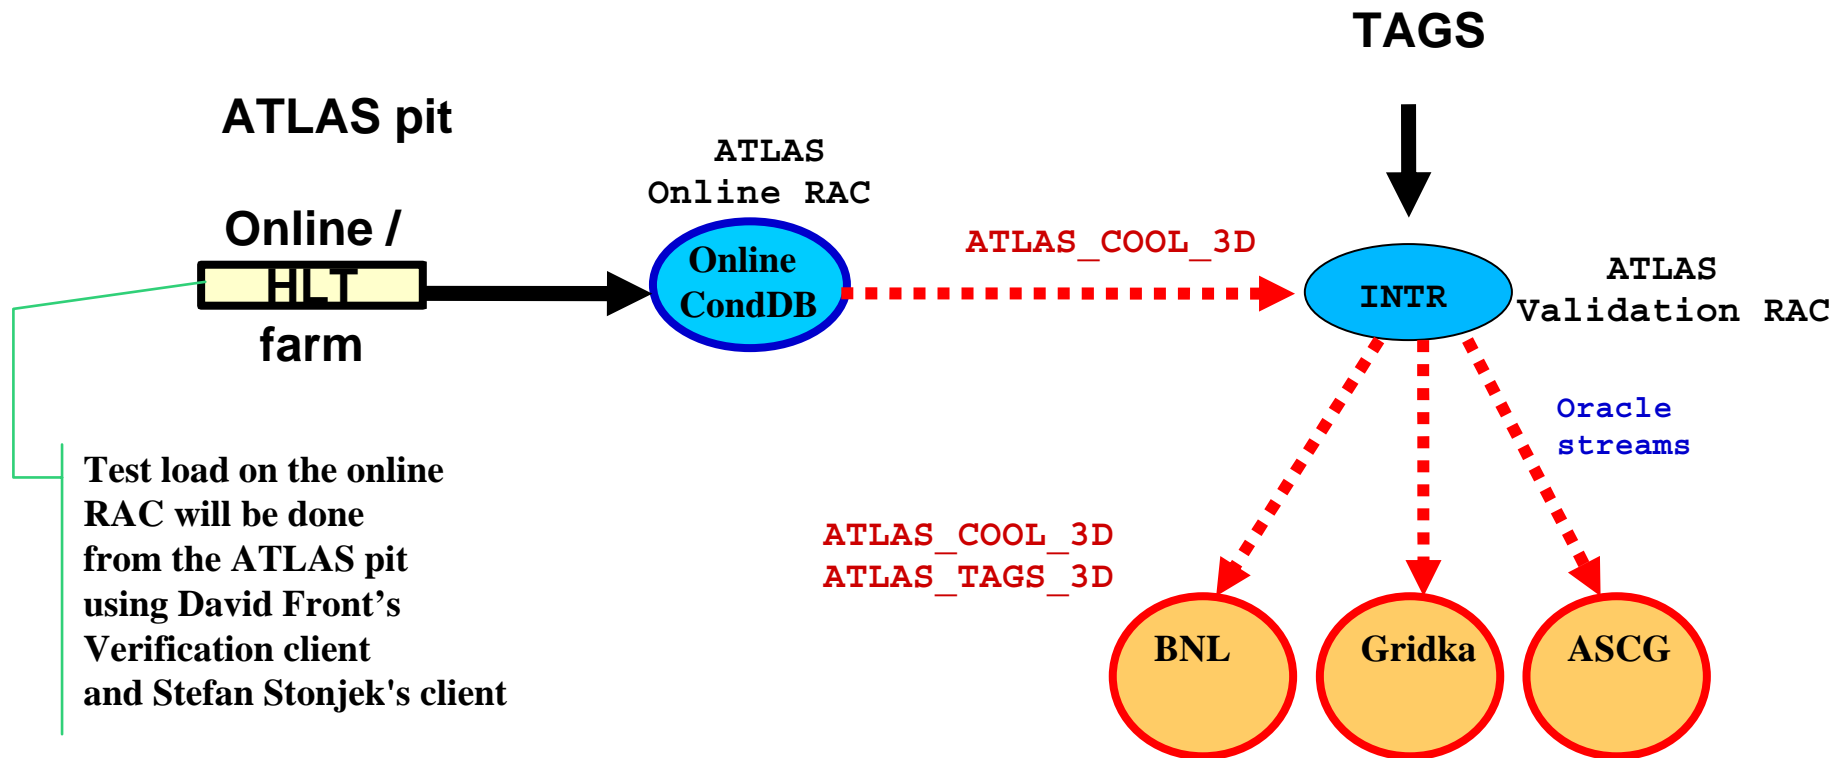
# Conditions Database Replication

- Conditions database mostly read by Tier-0/1
  - According to TDR Tier-2 will do mostly MC
- Writing to conditions database only at CERN
- Expected data volume: ~1TB/year
  
- TAGS database will have similar access pattern and data volume
  - Production will produce root files, inserted to database at CERN

# COOL Test scenario

- Folders: 500
- Channels: 200
  - 100 bytes per channel
- IOVs (Interval of validity): 300/day
  - 3 IOVs will be inserted at once (every 15 min)
- This is an estimate. We need better numbers soon!

# Testing Environment for Streams

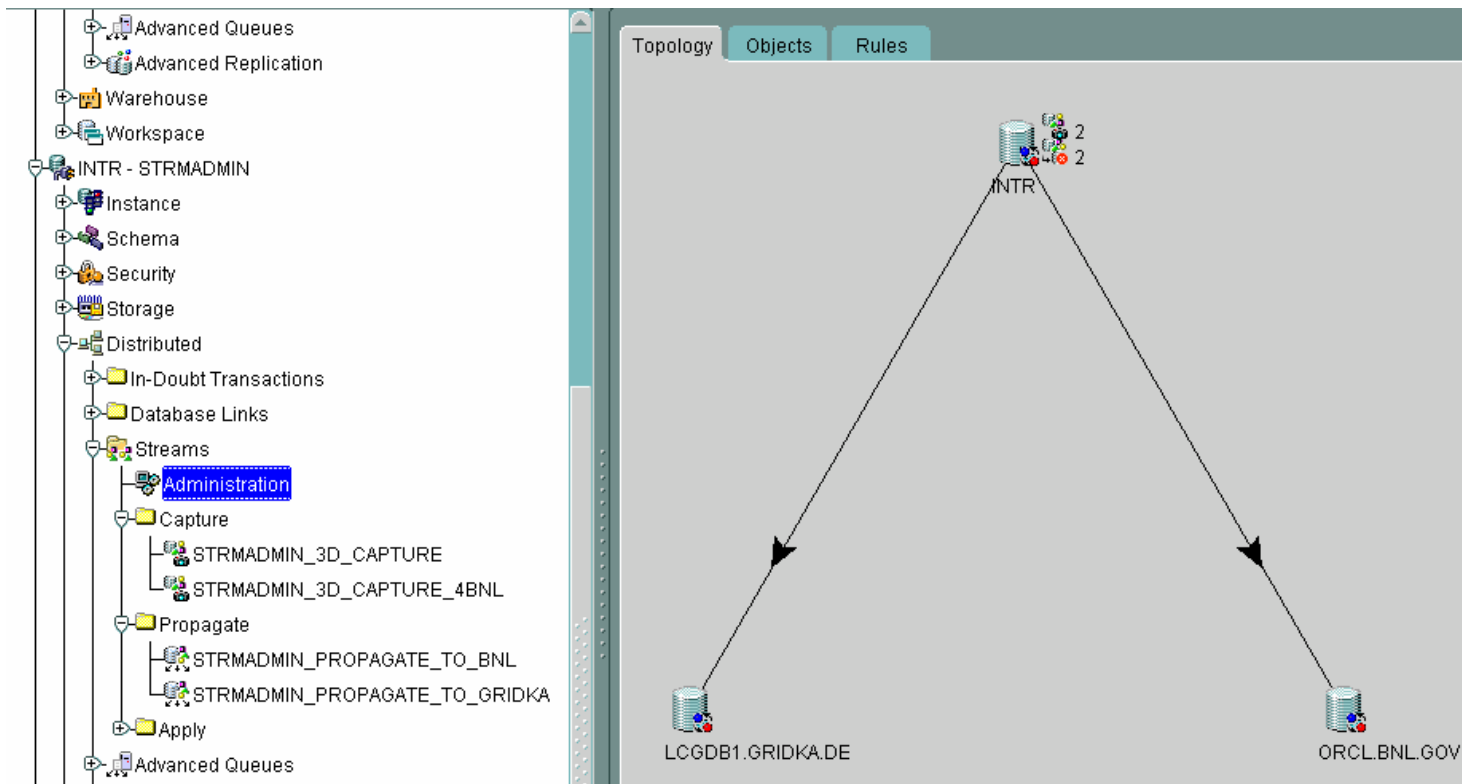


# Setup of Tests Environment

- Online database (ATONR)
  - Capture, propagation
- Offline database (INTR)
  - Apply, capture, propagation
- Tier-1 databases (BNL, GridKa, ASCG)
  - Apply processes
- Setup of statistics gathering at each point
- Schedule of tests, including coordination between COOL and TAGS tests

# Sites involved

- BNL and GridKa setup and started
  - baseline tests to these sites
- Should have ASCG within one week



# Test Criteria for Cond. and TAGS

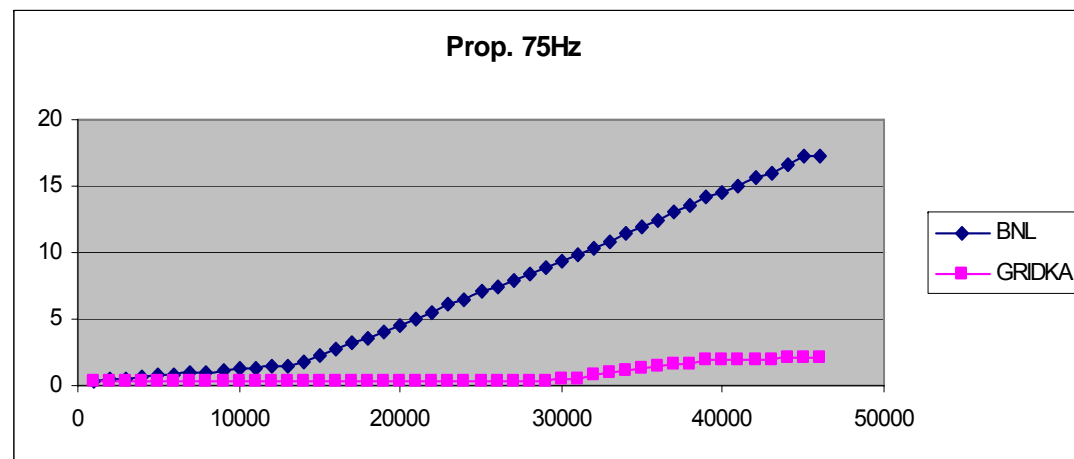
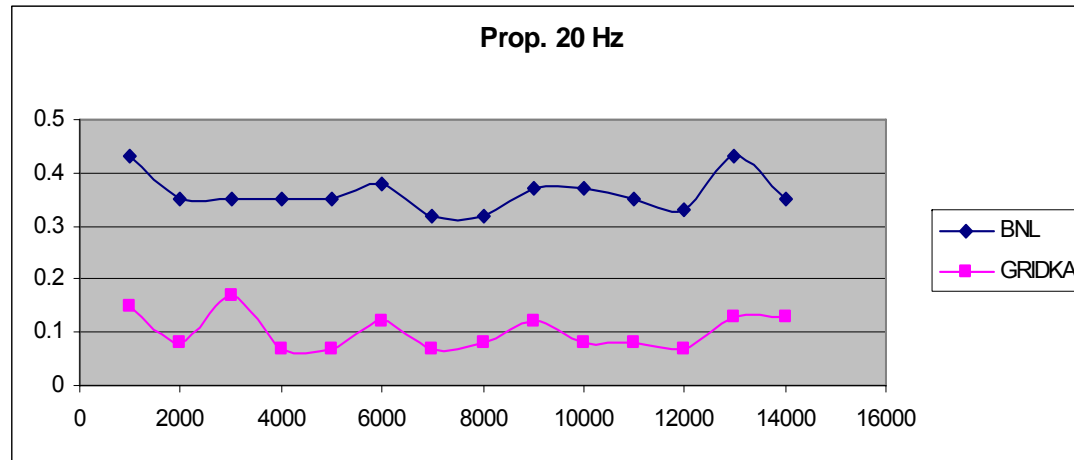
- Conditions database will have folders with different numbers of channels with different payloads.
  - Numbers needed!
- TAGS will write 1kByte/event
  - Same data volume for every event
- Easier to test



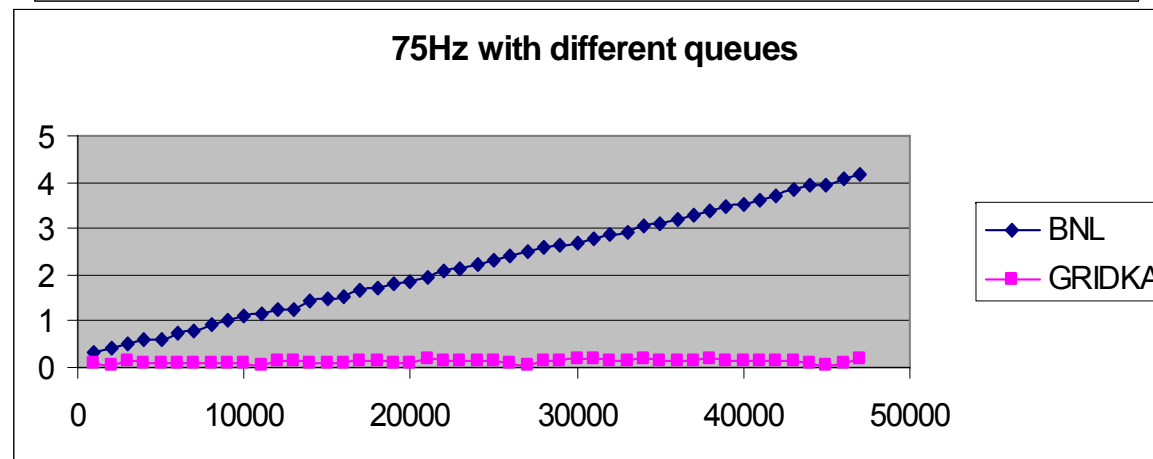
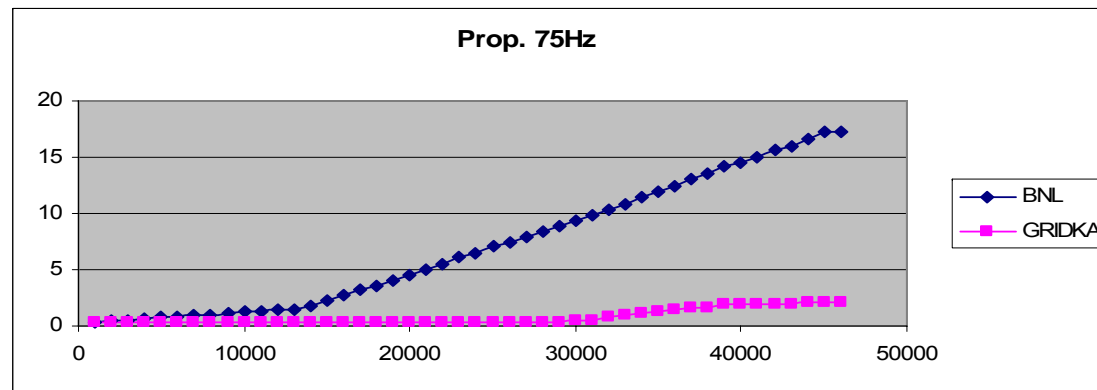
# Tests so far - TAGS

- According to plan test of TAGS database, we are in base line tests.
- Goal: establish 200Hz of pure inserts (no partition swapping, no indexes...)
- Results so far: GridKa can stand up to the rate, BNL not yet.
- Tests with 20Hz and 75Hz with single queue
- Tests with 75Hz and 200Hz with one queue for each site.
- Rate of commit: 1000 rows, size of transaction 1.3 Mbytes
- Rate of insertion : 200Hz = 15.6 Mb per minute

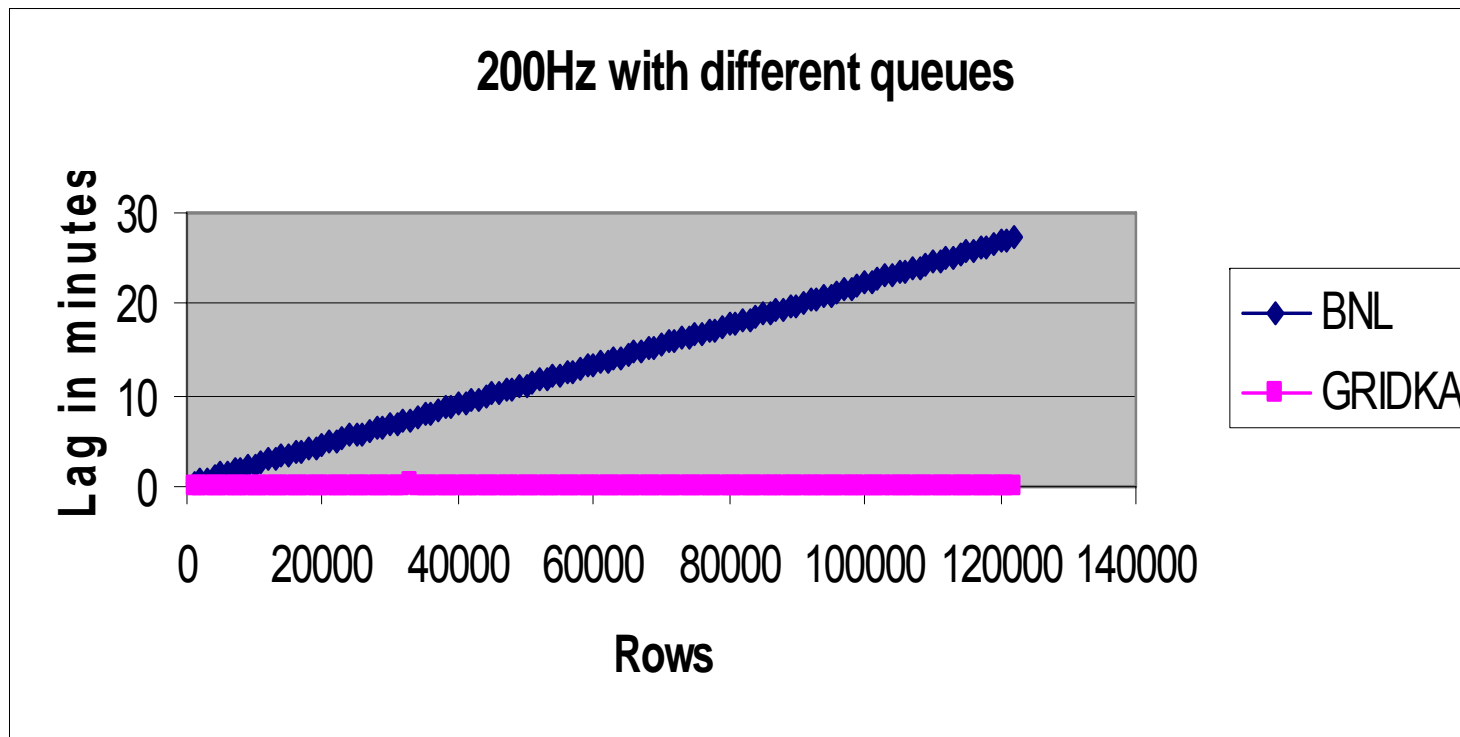
# Results with single queue



# Results with two queues



# Status at this time



# Status and Plan

## Tasks in hand:

- Include ASCG (Taiwan) in tests
- Improve performance of BNL
- Setup COOL performance stats gathering

## Next Tasks:

- Proceed with test plan for TAGS and COOL
- Gain thorough understanding of Streams administration, recovery from failure procedures

## Ultimate goal for TAGS:

**Get 200Hz rate in every site with streamlined architecture**