

Streams Service Review

Distributed Database Workshop

CERN, 27th November 2009

Eva Dafonte Pérez

- Announce Interventions
- Tier0 Responsibilities
- Tier1 Responsibilities
- Streams Resynchronization Review
- Recovery Scenarios Review
- Recent Problems or Interventions
- Summary

- Announce interventions
 - schedule new intervention using 3D wiki
 - submit EGEE broadcasts
 - register outages in the CIC portal
 - long interventions: contact Tier0 to analyze if it is necessary to split the Streams setup
- Unplanned downtime: update Tier0
 - problem description, progress and expected duration
- Report regularly

- Initial Streams setup
- Add new schemas to the Streams environment
- Split & Merge – new procedures in place
- Streams resynchronization
- Analyze and test new features and optimizations
- Validate upgrades and patches
- Monitoring

- Announce interventions
- Maintain the 3d OEM operational
 - check agents status
 - configure targets
- After an intervention: check and re-enable Streams processes
 - Use “STRMPROP_<TIER1>” account to connect to the downstream database
 - i.e. STRMPROP_PIC for Tier1 PIC
 - enable the propagation job
 - enable the capture process when site is split
- Streams resynchronization

- Service Incident Reports
 - unexpected database downtime > 4 hours

<https://twiki.cern.ch/twiki/bin/view/PDBService/ServiceIncidentReportT1>
- Recent **service changes**:
 - AMI replication to CERN
 - master: Tier1 Lyon
 - target: ATLAS offline
 - Tier1 Lyon acts as a Tier0 site for the AMI Streams setup

- Resynchronize a Tier1 site which is out of the Streams recovery window using **Transportable Tablespaces**
- Tasks are **shared between Tier0 and Tier1 sites**
 - Tier0:
 - split Tier1 sites
 - Streams operations: configuration, stop and re-start
 - Tier 1:
 - copy tablespaces from the “collaborative” site
- **Exercised:**
 - Taiwan using collaborative sites:
 - RAL – June 09
 - BNL – October 09

- Recovery is not transparent for the Streams setup
 - notify always to Tier0
- Recovery on destination
 - changes that had already being applied but lost by the recovery must be re-applied
 - split destination
 - configure Streams to re-capture changes from the last SCN applied at the destination
 - if replication is re-enabled without notifying
 - apply aborts, SCN information is updated, last correct SCN applied is lost
 - Streams re-configuration is not longer possible

- ASGC (Taiwan)
 - February 09: fire incident in ASGC data centre
 - database reallocated after several weeks
 - service unavailable – listener problem?
 - removed from the ATLAS Streams setup after 1 month
 - **incomplete recovery due to control file corruption**
 - data loss – impossible to re-start Streams
 - complete re-instantiation needed – RAL was volunteer as “collaborative” Tier1 site
 - September 09: after an unscheduled intervention, data got corrupted
 - **no valid backups found**
 - complete re-instantiation needed – BNL was volunteer as “collaborative” Tier1 site

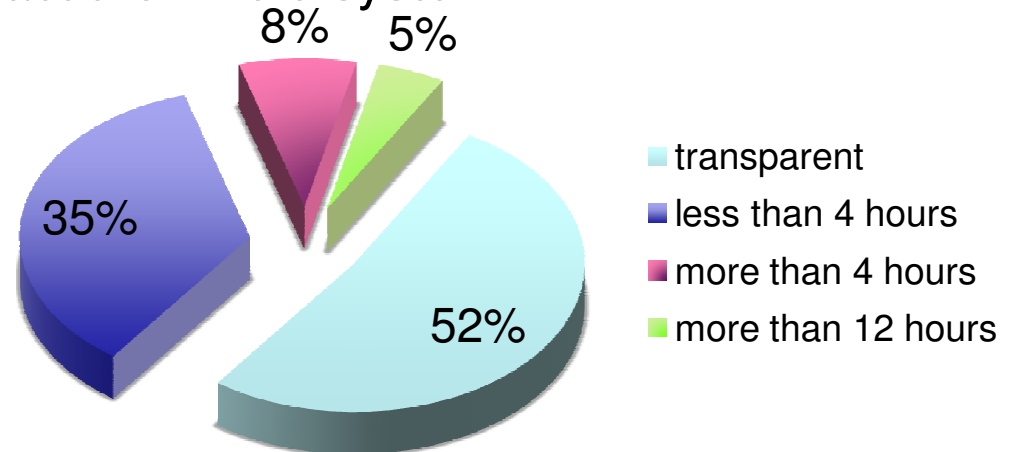
- Hardware problems after migration @RAL
 - caused data corruption
 - point-in-time recovery performed
 - Streams resynchronization for ATLAS and LHCb-LFC
- After 1st node reboot, database could not open @GRIDKA
 - due to high load
 - ATLAS database inaccessible
 - fixed after 2nd node reboot
 - real cause?

- ATLDSC and LHCB DSC downstream databases migration to new hw
 - smoothly
- Tier1 database migrations to new hw
 - RAL, SARA and NDGF
 - using **Data Guard** technology (see talk later)
 - transparent for RAL and NDGF
 - some problems at SARA, related to the connection string change

- Fatal Logminer error reported by ATLDSC capture
 - caused by ORA-07445: [ksfdetch()+622]
 - when restarting capture, after logminer checkpoint table shrink
 - patch suggested by Oracle support included in PSU 10.2.0.4.2
 - re-try capture re-start
- Apply dequeues but not applies LCRs
 - several Tier1 sites affected for LHCb and LFC
 - under investigation


- 3d wiki – Streams operations manual
<https://twiki.cern.ch/twiki/bin/view/PSSGroup/ServiceDocs>
- Overview for Troubleshooting Streams Performance Issues (metalink note 730036.1)
- Streams monitoring
<https://oms3d.cern.ch:1159/streams/main>
- Streams health check report
 - metalink note 273674.1
- 3d OEM
<http://oms3d.cern.ch:4889/em/console/>

- Keep the monitoring operational
 - spot problems quickly, understand bottlenecks, ...
- Coordination with Tier0
 - complex streams environments where the activity at one point might impact the whole system




- **Feedback!!!**
 - and collaboration to improve the documentation and the service

DM Questions



- Steps at Tier0
 - split “collaborative” Tier1 site
 - temporary Streams setup for Tier1 site to be resynchronized



- new dictionary
- clone capture process
- include apply rule set to avoid LCRs to be applied
- (only one “test” schema is replicated)

- Steps at Tier1
 - coordination between “collaborative” and “resynchronized” Tier1s
 - share connection strings
 - share streams administrator account !
 - create database links between databases
 - create directories pointing to datafiles and grant access
 - ask Tier0 to stop replication for both sites
 - “collaborative”: ensure tablespaces are read-only
 - alter tablespace ... read only;
 - “resynchronized”: remove tablespaces and datafiles

- Steps at Tier1
 - transfer datafiles from “collaborative” to “resynchronized”
 - dbms_file_transfer
 - parallel sessions
 - “resynchronized”: import tablespaces metadata
 - change tablespaces back to read-write
 - ask Tier0 to re-enable Streams
 - **apply rules must be dropped first!**
- Streams will recover the backlog produced during the operation automatically
 - merge all streaming to be done by Tier0