

Streams Service Review and Outlook

Distributed Database Workshop

PIC, 20th April 2009

Eva Dafonte Pérez

- Announce Interventions
- Tier0 Responsibilities
- New Split and Merge procedures
- Tier1 Responsibilities
- Streams Resynchronization
- Recent Problems or Interventions
- Recommended Patches
- 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

- New automated procedures
 - ORA-600 [KWQBMCRPTS101] error when dropping propagation fixed by Oracle
 - before it was needed to re-create all the streams components
 - “manual” intervention is avoided
 - scheduled downtime
 - new streams setup (queue, capture and propagation) is created in parallel to the main setup
 - unscheduled downtime
 - spilled LCRs removed from the main queue
 - execute resynchronize once the site is up again
- The procedures have been extended to all the database administrators at Tier0

- 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

- How to resynchronize a Tier1 site which is out of the Streams recovery window?

- Idea:


- use transportable tablespaces to move data faster

- tablespaces are copied from a “collaborative” Tier1
- “collaborative” Tier1 temporally unavailable
 - tablespaces need to be set to **read-only** while the files are copied

- complete re-instantiation using Streams

- 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

- MUON sites replication to CERN
 - master: 3 Tier2 sites (Rome, Munich, Michigan)
 - target: ATLAS offline
- AMI replication to CERN
 - master: Tier1 Lyon
 - target: ATLAS offline
- Resources:
 - currently 2 apply process @ATLAS offline
 - 4 more to be added!!
- Service level:
 - **problems must be addressed to the master side**

- Tier0 Responsibilities

- Initial Streams setup
- Add new schemas to the Streams environment
- Split & Merge
- Streams re-synchronization
- Analyze and test new features and optimizations
- Validate upgrades and patches
- Monitoring

Source database responsibilities



- ORA-01280: Fatal LogMiner Error + ORA-00600: [KRVRDCBMDDLSQL1]
 - caused by rebuild index operation using parallel option
 - ATLAS replication (conditions and PVSS)
 - **capture process cannot be restarted at the current SCN**
 - workaround proposed by Oracle: recreate capture using new dictionary after the index rebuild operation → **data loss!!**
 - **complete re-instantiation of the whole system**
- ORA-01372: Insufficient processes for specified LogMiner operation
 - one instance is down, number of parallel_max_servers is not enough
 - **increase parameter parallel_max_servers**

- Apply abortion: user error encountered while applying
 - ORA-04043: object ... does not exist
 - ATLAS replication (conditions and PVSS)
 - triggers, views, PL/SQL procedures, synonyms are not copied using transportable tablespaces
 - use datapump
 - "schema" triggers still not copied – manual creation
- GRANTs on views, PL/SQL procedures, functions and packages from owner to other accounts **do not get replicated**
 - “undocumented feature”, it is not a bug!

- ASGC (Taiwan)
 - January 09: system tablespace corruption
 - backups missing archived logs
 - complete re-instantiation using PIC
 - 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 will be needed – RAL volunteers as “collaborative” Tier1 site

- Power cut at RAL
 - caused corruption of the control file on the ATLAS database
 - complete recovery allowed to re-enable Streams without resynchronization
- Downtime at CNAF
 - scheduled for more than 5 days – out of the Streams recovery window
 - archived log files retention increased but not guaranteed in case of space pressure

- CREATE VIEW ON SCHEMA NOT IN STREAMS REPLICATED
 - the view references to a table in a replicated schema
 - same for synonyms, grants, ...
 - ATLAS and CMS replication
 - **apply aborts if schema does not exist in the apply side**
 - error might be safely ignored
- ORA-16146 REPORTED INTERMITTENTLY AND CAPTURE ABORTS + ORA-07445: [kghufree()+485]
 - related to change notification
 - CMS replication
 - **capture process aborts repeatedly**
 - manually clean `dbms_aqadm_sys.register_driver` jobs


- Connection problem to Gridka
 - propagation errors: ORA-12152: TNS:unable to send break message, ORA-03135: connection lost contact
 - apply side errors: ORA-00600 [knc1prcols:chrlen1], [kngo_kadadupk12:bad version AnyD], [OCIKCallPush: deprecated]
 - only affects LFC replication (even database is shared with LHCb)
 - propagation job is disabled after 16 unsuccessful connection attempts and cannot be restarted
 - workaround: recreate Gridka propagation
 - parallel streams setup to avoid data loss
 - diagnostic patch installed – waiting until the problem reproduces

- 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/>

Metalink note 437838.1

- 7579469 addresses performance improvement for capture process and logminer: merge label request on top of 10.2.0.4 for Bugs:
 - Bug 6599920 Capture aborts on LOB columns with ORA-26744 and ORA-26773
 - Bug 7345904 Streams capture slow processing direct path insert, high cpu for logmnr builder
 - Bug 6683178 High latencies in Streams capture, while capturing primary workload with a lot of DDL activities such as truncates of empty tables
 - Bug 6994160 Capture reader process constantly writing messages to trace file
 - Bug 6413089 Restarting a logminer session can be slow if the session has fallen behind
 - Bug 6650256 Parallel DDL (PDDL) transactions can cause logminer memory spill for Streams, or run slowly during adhoc log mining
- 7599054 in order to fix ORA-600 [KWQBMCRPTS101] when dropping propagation
- 5933656 Propagation ora-600 [KWQPCBK179], [1], [1369]
- 6827260 Excessive memory usage for lcr cache due to large freelists
- 7219752 ORA-26773 Malformed redo on capture of long
- 6452375 ORA-26687 No instantiation scn provided when drop child table
- 6640411 AQ propagation may fail after changing queue_to_queue=>>true
- 6838714 Apply process is slow after upgrading to 10.2.0.3
- 7033630 Apply aborting with ORA-600 [KNLQDQM2USR:4] after installing 10.2.0.4 patchset

- 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



