

Databases for Physics - update

CERN openlab II
Monthly Technical Review
11th December 2007

Dawid Wójcik





Monthly Technical Review

- Streams Optimizations
- New Streams 11g Features
- 32-bit Linux to 64-bit migration
- New HA setup for backups and monitoring
- Other
- Future plans

- Using rules
 - ATLAS: filter tables by prefix
 - rules on the capture side caused more overhead than on the propagation side
 - **complex rules converted to simple rules**

Complex Rule

```
condition => '( SUBSTR(:ddl.get_object_name(),1,7) IN ("COMP200", "OFLP200", "CMCP200",  
"TMCP200", "TBDP200", "STRM200")  
OR SUBSTR (:ddl.get_base_table_name(),1,7) IN ("COMP200", "OFLP200", "CMCP200",  
"TMCP200", "TBDP200", "STRM200") ) '
```

Avoid complex rules:

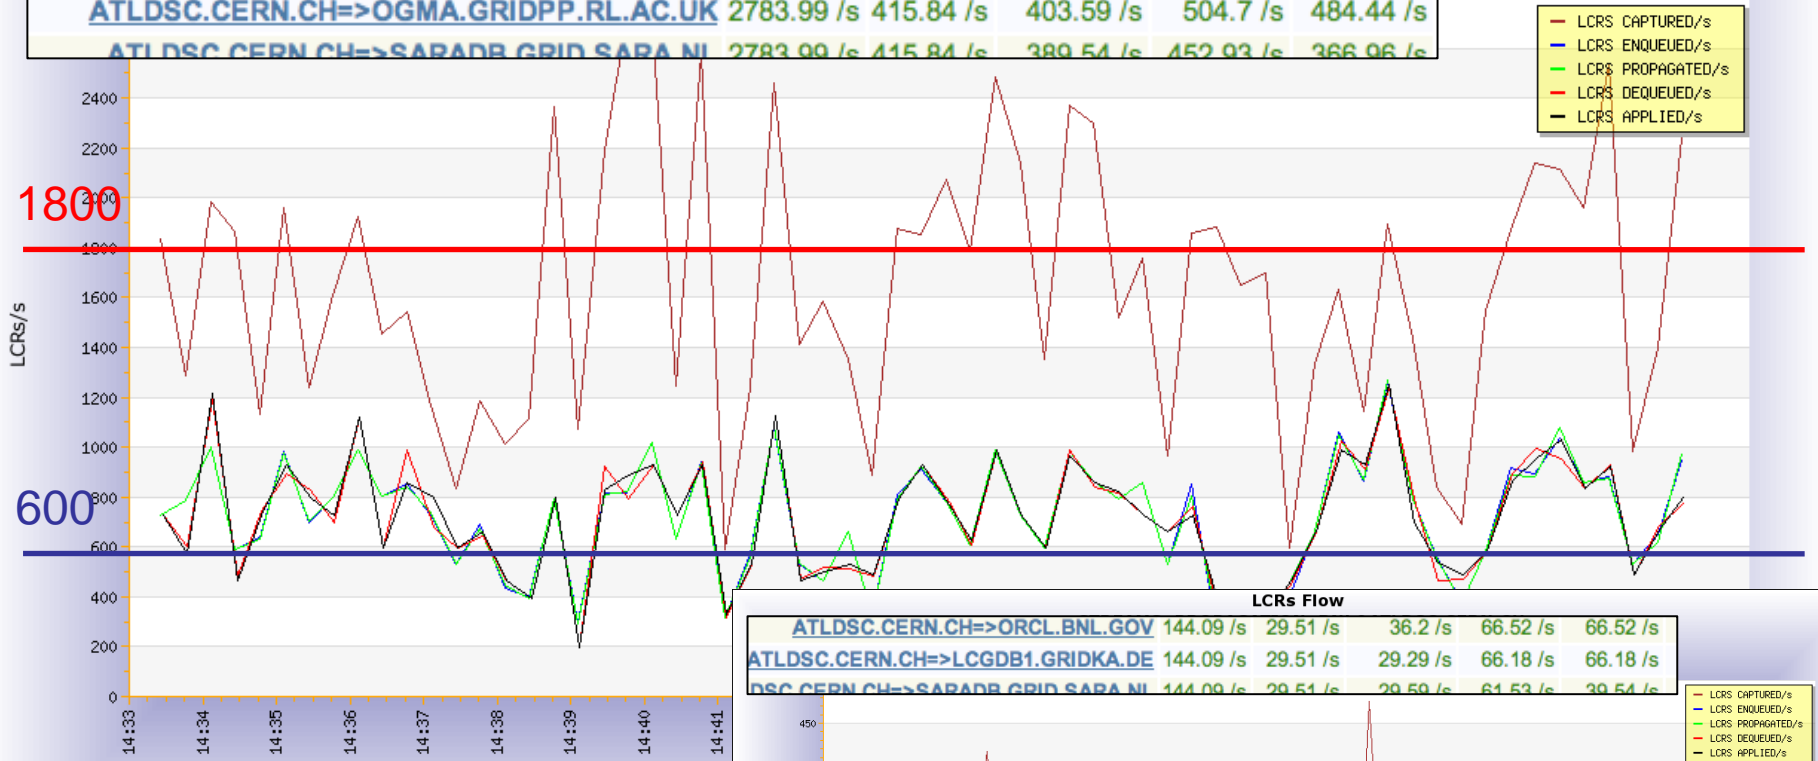
- LIKE
- Functions
- NOT

Simple Rule

```
condition => '(((ddl.get_object_name() >= "STRM200_A" and :ddl.get_object_name() <= "STRM200_Z") OR  
(:ddl.get_base_table_name() >= "STRM200_A" and :ddl.get_base_table_name() <= "STRM200_Z"))  
OR ((:ddl.get_object_name() >= "OFLP200_A" and :ddl.get_object_name() <= "OFLP200_Z") OR  
(:ddl.get_base_table_name() >= "OFLP200_A" and :ddl.get_base_table_name() <= "OFLP200_Z"))
```

LCRs Flow

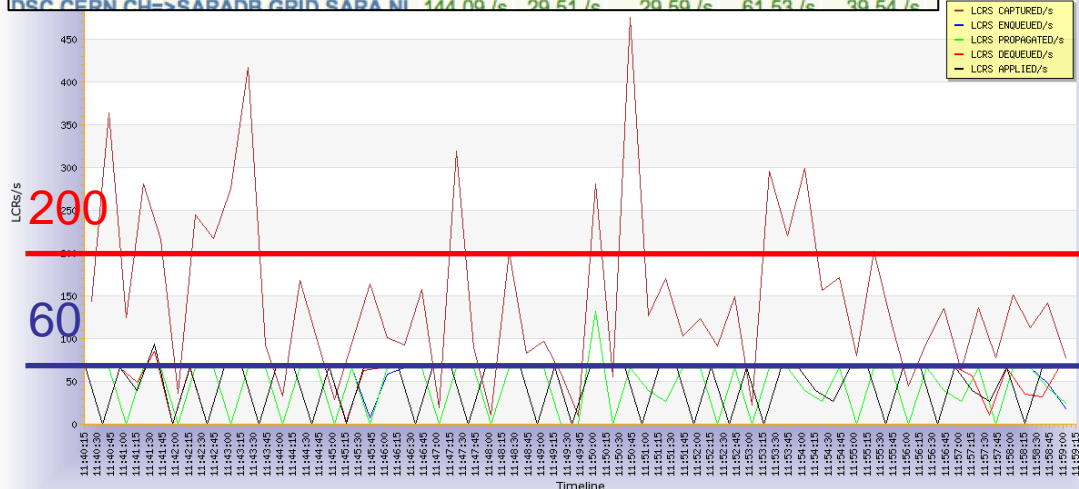
ATLDSC.CERN.CH=>LCGDB1.GRIDKA.DE	2783.99 /s	415.84 /s	413.94 /s	329.89 /s	320.84 /s
ATLDSC.CERN.CH=>ATLAS.DB1TIER1.NDGF.ORG	2783.99 /s	415.84 /s	406.71 /s	404.14 /s	388.76 /s
ATLDSC.CERN.CH=>OGMA.GRIDPP.RL.AC.UK	2783.99 /s	415.84 /s	403.59 /s	504.7 /s	484.44 /s
ATLDSC.CERN.CH=>SARADB.GRID.SARA.NL	2783.99 /s	415.84 /s	389.54 /s	452.93 /s	366.96 /s



Simple Rules

LCRs Flow

ATLDSC.CERN.CH=>ORCL.BNL.GOV	144.09 /s	29.51 /s	36.2 /s	66.52 /s	66.52 /s
ATLDSC.CERN.CH=>LCGDB1.GRIDKA.DE	144.09 /s	29.51 /s	29.29 /s	66.18 /s	66.18 /s
ATLDSC.CERN.CH=>SARADB.GRID.SARA.NL	144.09 /s	29.51 /s	29.50 /s	61.53 /s	39.54 /s



Complex Rules

- Combined Capture and Apply
 - capture sends LCRs directly to apply
 - source and target 11g databases
 - only 1 propagation job running
 - configuration detected automatically
 - big performance improvement
 - rate: 14.000 LCRs/sec (before 5.000 LCRs/sec)




CERN

LCRs Flow

STREAMS_PROPAGATE_DEST3D11@T11G64.CERN.CH

Generated on Tuesday 04th of December 2007 02:40:21 PM
Time range 10:36:24 06-11-2007 <-> 12:36:04 06-11-2007



- **Split/Merge of Streams**
 - maintains high performance for all replicas
 - automated, fast “catch-up” for unavailable replica

- **Cross-database LCR tracking**
 - trace Streams messages from start to finish in single view

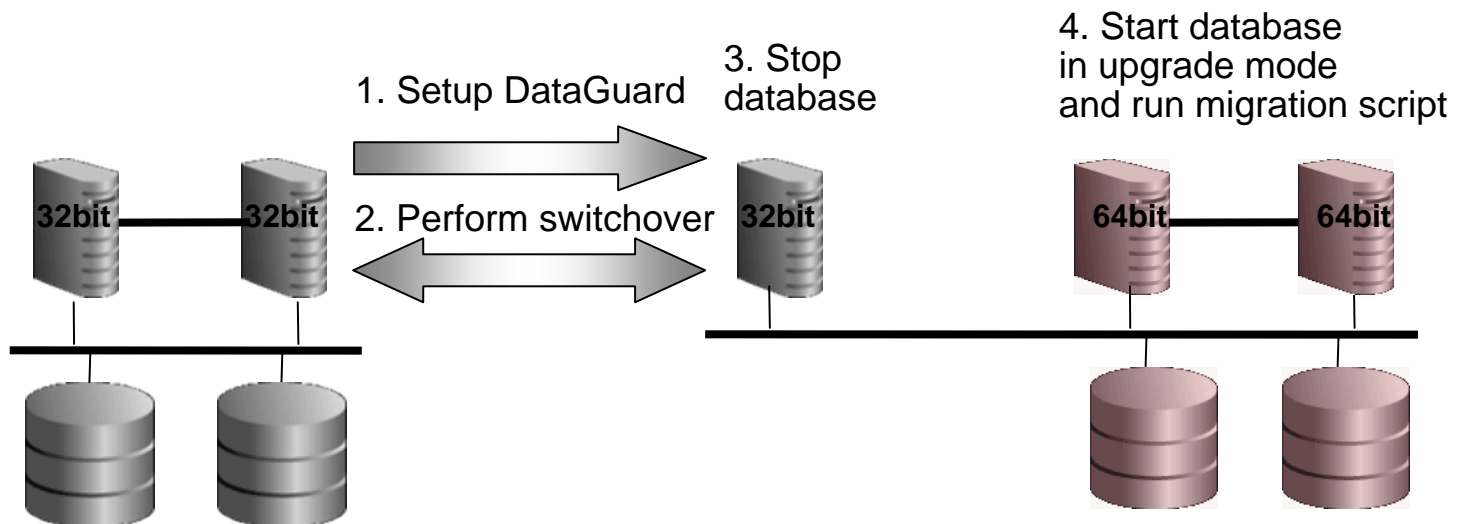
- **Source and Target data compare & converge**
 - compare rows in an object at 2 databases
 - converge objects in case of differences

32-bit Linux to 64-bit migration

- New hardware acquisitions next year
 - 60 SAN diskservers (16 disks x 400GB)
 - 35 mid-range servers (2 x Intel quad core, 16 GB RAM)
- Moving from 32-bit Linux to 64-bit Linux
 - Migration using Oracle DataGuard
 - minimum downtime required (independent of database size)
 - easy to rollback if something goes wrong

32-bit Linux to 64-bit migration

- Setup DataGuard
- Perform switchover
- Stop intermediate database
- Perform upgrade (*utlirp.sql*)



- Migrated backup/monitoring machine to cluster configuration (active-passive mode)
 - file system mounted on demand on the node that is active
 - cron jobs schedule backup, monitoring and PKI access management on active node
 - machine behind VIP to provide HA in case the current node fails
 - LDAP setup as cluster service for TNS names resolution

- **More PVSS tests**
 - space waste investigation (result of direct path inserts)
- **Following issues on production RACs**
 - ORA-00600 when updating more than 2 CLOB columns in one go
 - ORA-07445 when using nested joins in a query
 - ORA-00942 table or view does not exist – investigating
- **Presentations on WLCG Service Reliability Workshop**
 - DB application design issues
 - DB Monitoring tools
 - Streams service review

- More 11g features tests
 - streams testing
 - ASM
- HA setup for downstream capture (failover)
- HA setup for streams monitoring
- Big migration of services to 64-bit Linux using DataGuard
 - 35 new mid-range servers (280 cores) – 175% of current CPU power dedicated to databases (160 cores)

The feedback is circulated between the people involved.

Monica Marinucci Lopez
June Farmer
Graeme Kerr

Oracle EMEA

Management of the programmes
Management of the programmes
Technical liaison

Bjørn Engsig

Oracle Development

Primary Development Contact

Sverre Jarp
Juergen Knobloch
Maria Girone

CERN Openlab

Chief Technologist Officer
IT-PSS Group leader
IT-PSS-DP Team Leader

Q & A