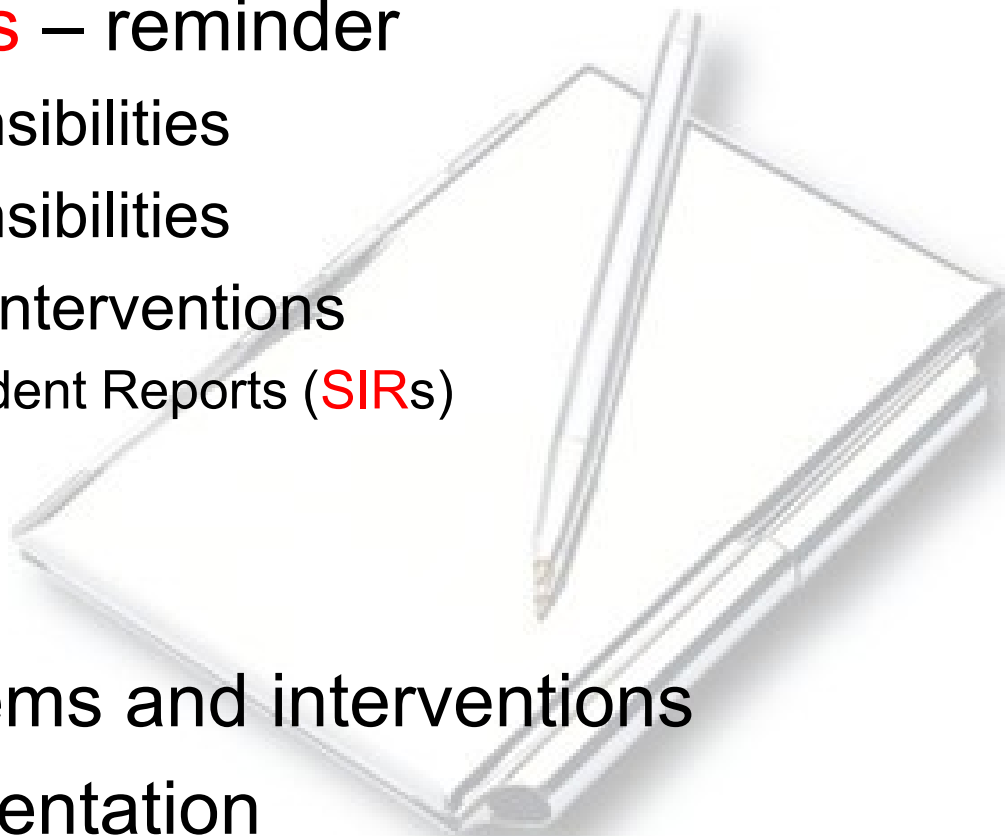


# Streams Status Report

Distributed Database Operations Workshop  
CERN, 17th November 2010

Dawid Wójcik

- 3D **Operations** – reminder
  - **Tier0** Responsibilities
  - **Tier1** Responsibilities
  - **Announcing** Interventions
    - Service Incident Reports (**SIRs**)
- Procedures
- Monitoring
- Recent problems and interventions
- Useful documentation





- Initial Streams setup
- Adding new schemas to the Streams environment
- Split & Merge procedures
- Streams resynchronization
  - Split & merge
  - **Coordination** role between voluntary site and affected one
- Analyze and test new features and optimizations
- Validate upgrades and patches
- Monitoring
  - Oracle Enterprise Manager for 3D
    - Agent installation and configuration is **Tier1 responsibility**
  - StreamMon



- **Interventions**
  - **Before – announce**
  - **After – check and re-enable Streams processes**
    - **Apply**
    - **Propagation** – use “STRMPROP\_<TIER1>” account to connect to the downstream database (i.e. STRMPROP\_PIC for PIC Tier1)
    - Enable the **capture** process when site is **split**
- **Maintain the 3D OEM operational**
  - **Check agents status** (recommender version is 10.2.0.5)
  - **Configure targets**
- **Streams resynchronization**
  - Collaborate with voluntary Tier1 on re-synchronizing streams
  - Update Tier0 on resynchronization progress

- Announce DB-related interventions
  - Schedule new intervention using [3D Twiki](#)
  - 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** – **contact Tier0**
  - Problem description, progress and expected duration
- Report regularly
  - Feel free to ask for a specific help from Tier0 and Tier1 DBAs





- **Service Incident Reports (SIRs)**
  - Unexpected database downtime > 4 hours
  - Unexpected streams outages caused by hardware or software problems
    - **Create** SIR for site **resynchronization**
    - **Create** SIR for any **streams outage** caused by Oracle bug or not known issue
    - Do not create SIR for apply problems after node reboots and 'user' errors caused by modifying read-only data
  - Log SIRs to:
    - <https://twiki.cern.ch/twiki/bin/view/PDBService/ServiceIncidentReportT1>
  - Template:
    - <https://twiki.cern.ch/twiki/bin/view/PDBService/TemplateSIR>





- Procedures
  - See <https://twiki.cern.ch/twiki/bin/view/PDBService/ServiceDocs>
- Streams **resynchronization**
  - <https://twiki.cern.ch/twiki/bin/view/PDBService/TransportableTablespacesforScalableResynchronization>
  - Resynchronize a Tier1 site which is out of the Streams recovery window using transportable tablespaces
  - For Tier1-perspective and experience see talk from Alexander & Carmine
  - Please remember to **instantiate** streamed schemas after copying to destination!

- Recovery procedures in Streams environment

- <https://twiki.cern.ch/twiki/bin/view/PDBService/3DStreamsRecovery>

- **Recovery is not transparent** in Streams environment

- **Always notify Tier0** when attempting **recovery**
    - Changes that had already being applied but lost by the recovery must be re-applied
      - Split required and new capture with different SCN needs to be created

- **If replication is re-enabled without notifying Tier0**

- Apply aborts, SCN information is updated, last correct SCN applied is lost
    - **Streams re-configuration is not longer possible**





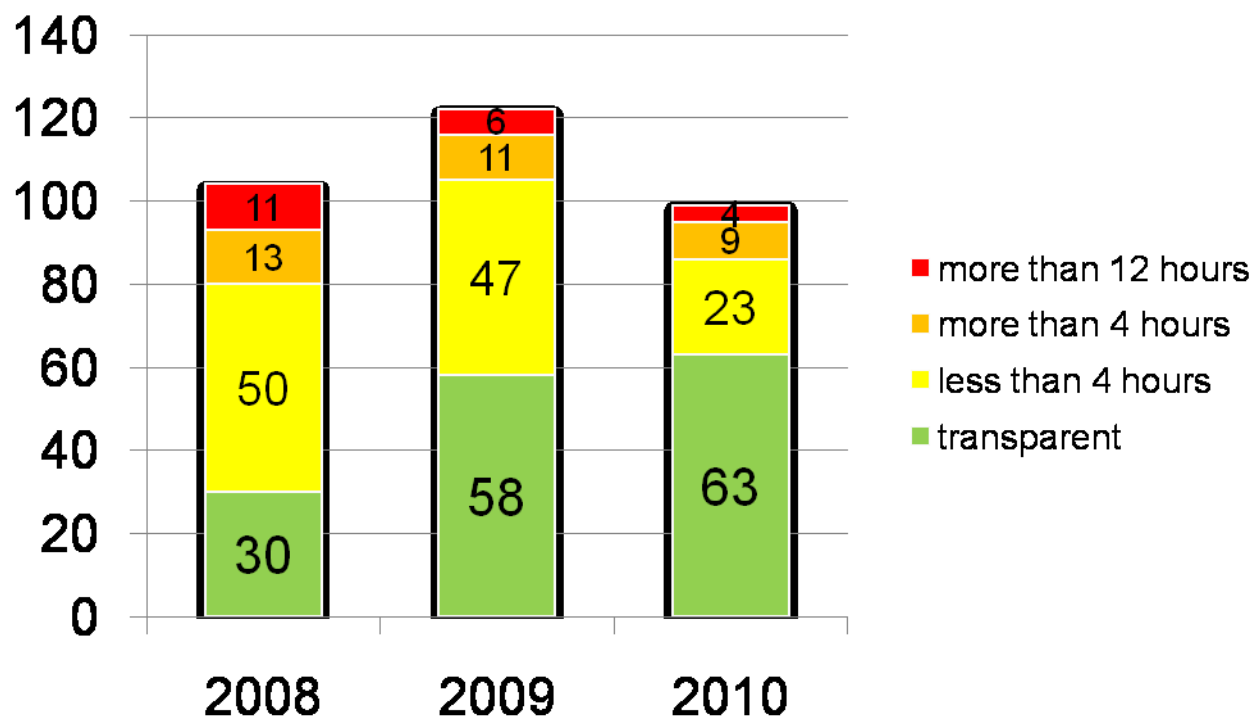
- **OMS 3D**
  - **Cleanup required**
    - 33% targets seen as unknown (132 targets)
    - 3% targets seen as down (12)
    - **Tier1s affected**: CNAF, IN2P3, MICHIGAN, RAL, INFN, SARA, TRIUMF
  - For the moment we stay on OMS 10.2.0.5
- **StreamMon** – new features
  - Better streams time statistics (capture time, replication delay, time since last LCR)
  - DB PGA monitoring
  - Sessions metrics (total, sessions/s)
  - Schema statistics – LCRs per schema



- Weekly reports for Tier1s
  - Provide
    - CPU and IO usage statistics (per domain/user)
    - Service usage
    - Simple session statistics
    - Password expiration
    - Space usage
  - Installation instructions
    - <https://oms3d.cern.ch:1159/admin/t1reports.php?install>
  - Access
    - <https://oms3d.cern.ch:1159/admin/t1reports.php>
  - Please install if you haven't done it yet
    - Not obligatory though



- More transparent interventions at Tier1s in 2010 (64% of total, compared to 48% in 2009)
- Decreased total number of interventions and interventions requiring more downtime



- RAL – 20<sup>th</sup> of July 2010

- Multipath reconfiguration intervention
- After the intervention only few LCRs were applied before the ORA-01403 was hit.

It looked like another instance of the bug where the LCRs are consumed but not physically applied on the database. A service request (3-1950544531) was open with oracle under the CERN support ID.

- No LCRs were applied until the database was resynchronized on 21<sup>st</sup> July 2010





- SARA – 18<sup>th</sup> of August 2010
  - DB crashed with datafile corruption
  - Extensive analysis conducted at SARA found potential storage problems
  - DB has been later restored to different hardware
  - RAL helped SARA to recover using transportable tablespaces
  - 8<sup>th</sup> of September – SARA replication is back
- SARA – 26<sup>th</sup> of October 2010
  - Apply failed with ORA-01403: no data found
  - Some data was found to be missing from the time of last resynchronization with RAL (between export and import)
  - Resynchronized selectively from Tier0
  - Problem traced down to **missing instantiation** at destination
    - **Transportable tablespaces procedure needs an update**



- ASGC – 24<sup>th</sup> of September
  - Site-wide power cut, after half an hour abnormal DB shutdown as UPS's battery run out
  - Several problems during data recovery
  - Decision taken to re-synchronize from full backup from RAL
  - Problems with transferring the backup to ASGC
  - Network is not fast enough on normal links, decided to put backup in Castor and ship data via FTS
  - 17<sup>th</sup> October – ATLAS decided to drop conditions DB at ASGC and use Frontier instead





- **PSU July 2010 (PSU4)**
  - Two issues not caught during PSU4 **validation**
    - **Non-rollingness** of the patch (bug 9827988) – patched **instance hangs** in mount state, **affecting the whole cluster**
      - Oracle Support recently confirmed PSU4 is not rolling if installing, in a fully supported configuration, with shared ASM and Oracle home
    - ORA-7445 and high contention during logoff with multiple sessions in one process (bug 6196748) – **whole cluster unresponsive**
      - affecting COOL and other applications using multiple sessions in one server process
      - Fixed by one off patch, tested; **fixed in 10.2.0.5**
  - We are now pushing for more comprehensive validation and gained much more experience

- 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/>
- Ask questions to grid-service-databases mailing list or use Tier1 Service Coordination Meeting to raise any outstanding DB-related issues



