



Database and EventIndex Readiness for Run 2

Dario Barberis
Genoa University/INFN

Databases

- Database and Application infrastructure:
 - Maintaining support for over 300 schemas
 - All New DB and Application hardware – described before
 - PVSS (online supporting DCS):
 - Upgrade to latest version (major task: EN-ICE, ATLAS, IT-DB)
 - Also DBAs improved indexing, partitioning (volume reduced 50%)
- Database Distribution (Geometry, Trigger, Conditions, AMI, PVSS)
 - New schemas under replication to Tier-1s for Run 2:
 - 2 Trigger schemas (real data and MC)
 - Run 2 Conditions DB instance (CONDBR2) - 29 schemas
 - Replication Technology:
 - Oracle Streams --> Golden Gate for remote sites
 - Oracle Streams --> Active Data Guard (PVSS and Standby DBs)
 - Frontier: running smoothly and ready for Run 2 operations
 - excellent cooperation from sites
 - new versions deployed without effect on operations
- ADCR (ADC Databases):
 - smooth DQ2 data transfer to new Rucio system
 - exploit Standby DBs to reduce production DB load
 - exploit archive DB (ATLARC) for historic, read-only data
 - Current work: a lot of new client software: lots of work behind refining queries
 - addressing row contention, performance issues, tuning



Highlights from the DB applications area

- **ATONR:** Performed DB schema upgrades on all PVSS accounts with advantages in reduction of the future disk space usage by factor 2 plus range partitioning with uniform time interval.
- **ATONR and ATR:** new COOL folder instance CONDBR2 for the Run2 conditions data and new Run2 trigger DB setups.
- **ADCR database:**

PanDA monitor

PanDA server

Rucio modules

Rucio modules

Id	Name	CPU Usage	CPU Load
1	ADCR1	3%	2.05
2	ADCR2	23%	2.03
3	ADCR3	31%	6.48
4	ADCR4	1%	1.15

- PanDA server and PanDA monitor workloads served by different DB nodes of the Oracle cluster
- Started initiative for distributing the workload of the Rucio modules on different DB nodes

Advantages: 1) much better positioned in case of high workload
2) each DB instance data cache of 200GB can be used

New developments: Use Oracle's "Flashback archive" feature for keeping changes done by update and delete operations. Valuable archive which could have a role as real time machine.

Metadata, Conditions, Other

- Metadata:
 - AMI:
 - pyAMI5 and Tag Collector III: deployment
 - major changes to data sources: ADC DBs and transforms
 - Also changes following Metadata Review (eg: AMI Tags, etc)
 - In the works: association of datasets with publications/analysis and allowing corrected values for MC metadata to be added more easily
 - COMA and CTB (Cool Tag Browser):
 - Improved tools for Conditions management
 - AMI/COMA: Data Periods operational for Run 2
 - major changes to COMA loading workflow (underway)
- Conditions DB:
 - Ongoing activities:
 - provide feedback on ROOTv6/ COOLv3 issues
 - streamlining POOL file movement to online
 - DBRelease size: customized DBReleases for HPC
 - improvements to python based tools
- Other: implemented distribution of CP/Analysis calibration files over CVMFS



EventIndex Deployment Status (1)

- All major components exist and work satisfactorily:
 - Data Collection: Producer (stand-alone) transform and Consumer
 - Merged transform under development
 - Storage System: Data organisation in Hadoop and indexing in catalogue; also Trigger decoding interface
 - Query System: CLI and web interfaces.
 - Also EventLookup for pathena
 - Monitoring: System level monitoring OK; contents monitoring under development
- ✓ All Run 1 and Run2 Tier-0 production (RAW+ESD+AOD) indexed at Tier-0
 - Decided to run on all first-pass AODs at Tier-0 and last reprocessing version (whichever it may be for each run) on the Grid, starting from M7 and working backwards in time
 - In this way we emulated the 2015 data-taking conditions
 - EventIndex entries now routinely produced for all data15 at Tier-0
 - EI Prducer jobs run automatically in the Tier-0 workflow
 - ➔ Question: do we need to index DRAW/DESD/DAOD at Tier-0 too?
 - Still missing but coming soon: automatic import of Trigger decoding tables from COMA
 - Used to index events by trigger chain



EventIndex Deployment Status (2)

- Currently working on automation of the data flow and system interconnections
 - Discussion this week on integration with ADC tools (ProdSys-2 + Rucio + AMI):
 - Establish EI production as an additional production type
 - Automatically run EI production tasks in ProdSys-2 on all datasets of type **EVNT** and ***AOD***, projects **data15*** and **mc15***, that are "valid" and "closed"
 - Exact definition under discussion
 - Prototype this month
 - Initial operation with EVNT datasets to feed the Event Server
 - Workflow will be revised later this year to see if faster turn-around is possible
 - Other work in progress:
 - EventIndex production chain for simulated data exists but not completely tested (especially trigger decoding part)
 - We also need a discussion on which (past) simulated data is useful to index
 - All mc15? Or only newly produced datasets? Any Run 1 simulation needed?
 - Revision of user examples:
 - <https://twiki.cern.ch/twiki/bin/view/AtlasComputing/EventIndexTutorial>



Outlook

- We are ready for Run 2
- We are not scared
- Cross fingers!