

US ATLAS DDM Operations

Alexei Klimentov, BNL
US ATLAS Tier-2 Workshop
UTA, Dec 8th 2006

DDM Operations (US ATLAS)

- Coordination : Alexei Klimentov and Wensheng Deng
 - BNL : Wensheng Deng, Hironori Ito, Xin Zhao
 - GLTier2 : UM : Shawn Mckee,
 - NETier2 : BU : Saul Youssef,
 - WTier2 : Wei Yang
 - SWTier2 : Patrick McGuigan
 - OU : Horst Severini, Karthik Arunachalam
 - UTA : Patrick McGuigan, Mark Sosebee
 - MWTier2 : Dan Schrager
 - IU : Kristy Kallback-Rose, Dan Schrager,
 - UC : Robert Gardner, Greg Cross

DDM Operations (US ATLAS) Main Activities

- DDM Functional tests :
 - A.Klimentov, P.Nevski, H.Ito
- LRC stress test :
 - P.Salgado, S.Reddy, W.Deng (DB support : Y.Smirnov)
- LFC stress test :
 - A.Zaytsev and S.Pirogov (both from DDM central operations), AK, PN
- DDM development and deployment :
 - P.McGuigan, W.Deng, K.Kallback-Rose, H.Ito
- End users tools (development and support)
 - T.Maeno
- User's Support
 - X.Zhao, H.Ito, H.Severini, AK
- User's Feedback
 - N.Ozturk

US ATLAS DDM and Production Workshop at BNL

- Workshop action items and follow ups
- Priority list for DDM operations for Q1 2007

Action Items from DDM WS at BNL Sep 2006. 1/7

- Installation

- A2. DDM ops test-bed set up (coordinated by Wensheng and Patrick)
 - Central DB server prototype at BNL
 - Sites service VO box at UTA
- A3. DDM installation procedure (coordinated by Patrick)
 - e-mail from Patrick (Dec 5th)
 - “..Generic DDM installation procedure for all sites complete with pacman installation for 0.2.12
The most basic test of subscribing to data at BNL works without a problem with test site UTA_TEST1...”
 - DQ2 client (pacman version) Hiro (done)
- A11. DDM installation and deployment (coordinated by Alexei)
 - 0.2.12 will be in production at least up to Mar 2007
 - 0.2.12 validation done (Wensheng)
 - Time slot and scenario agreed with Kaushik
 - Installation on the first sites (BNL, BU, UTA) week 50 (Dec 11 – 18)

Action Items from DDM WS at BNL Sep 2006. 2/7

- A4. DDM operations WiKi, documentation and FAQ (coordinated by Horst) – done
- A5. DDM AFS and CVS repository (coordinated by Xin) – done
- A6. *Data Integrity* (coordinated by Kaushik) – see next slide
- A7. *DDM Monitoring* (coordinated by Alexei)
 - “classical” and ARDA monitoring
 - “classical” monitoring development is frozen
 - ARDA monitoring as one ATLAS will use in the future, all new features will be added to ARDA version only
 - Statistics of dataset subscription per site (Xin, Hiro) – in progress
 - Site Level DDM dashboard (Kristy) – in progress
 - Lack or inconsistencies in monitoring is the source of many mistakes, the situation is improved since May (when “classical” monitoring was brought into production)
 - ARDA takes responsibility for the monitoring and developing dashboard

Action Items from DDM WS at BNL Sep 2006. 3/7

- A8. DDM end-users tools (coordinated by Torre) – done, in production
 - Extend dq2_get package functionality and add possibility to register datasets and files in DQ2, LFC/LRC
- A9. DDM functional test (coordinated by Alexei, Pavel)
 - Tests scope
 - Using existing DDM tools for datasets registration, subscription and monitoring
 - Data transfer from CERN to Tier-1
 - Data exchange within the cloud
 - Large files data transfer
 - One of the problems observed with dccp command is fixed
 - Tier-1/Tier-1 data transfer
 - T1/foreignT2 access (Dan's proposal)
 - October and November tests are conducted
 - Tests are successful for all US ATLAS sites
 - Patches to 0.2.12 and features to be implemented in 0.3

Action Items from DDM WS at BNL Sep 2006. 4/7

- *A6. Data Integrity (coordinated by Kaushik)*
 - *Full chain to check data produced by Panda (on sites and at BNL)*
 - *Most typical problems*
 - *Replication failed because of computing facilities problems*
 - *Poor monitoring and notification*
 - *Manual resubscription is needed*
 - *“slow” data transfer (usually after outages)*
 - *Files are copied to BNL, but not registered in LRC*
 - *Subscription is not processed for a long time*
 - *‘0’ length and files with incorrect length*
 - *DDM operations is trying to set up central repository for data integrity scripts*
 - *More input is needed for OSG specifics part*

Action Items from DDM WS at BNL Sep 2006. 5/7

- *A10. LRC vs LFC evaluation (coordinated by Alexei)*
 - Set up mySQL cluster at CERN
 - Computers from CERN IT
 - Cluster installation (Pedro with help from Sasha and Yuri)
 - *Mirroring ATLAS LFC (Sudhamsh)*
 - Test suit for LFC – ready (Alexei, Pavel, A.Zaytsev, S.Pirogov)
 - *Tests will be conducted before Xmas*
- **US ATLAS and ATLAS Users support (Hiro, Xin, Alexei)**
 - **A12. AOD and ESD datasets consolidation at BNL**
 - Problems related to lack or inconsistent information and data accessibility from many clouds
 - **A13. Users support to transfer data to BNL and US sites**
 - via DDM ops Savannah requests

Action Items from DDM WS at BNL Sep 2006. 6/7

- DDM review and US ATLAS requirements and priority list (coordinated by Jim)
 - A8. DDM end-users tools
 - Very strong opposition from LCG sites for dq2_get extention
 - A10. catalogs evaluation (LRC vs LFC)
 - Very strong opposition from LCG sites even to consider LRC as a fallback solution for LFC
 - Support from many reviewers for LFC stress test
 - A13. US ATLAS requirements and priority list
 - A16. Auto fallback to alternative transfer mechanism in DQ2
 - According to DDM developers there is no alternative to FTS
- The coherent position of US ATLAS reviewers (Jim, Kaushik, myself). I believe it will be reflected in final document.

Action Items from DDM WS at BNL Sep 2006. 7/7

- Manpower issues (coordinated by Jim)
 - The situation in US cloud is better than for many other clouds, probably only LYON has the similar support for DDM operations
 - Each site has people working for DDM with very strong support from computing specialists (non-ATLAS members)
 - For many LCG sites the problem is not only in the lack of the experienced people working for DDM and computing. But very often in communication between ATLAS and Tiers experts.
 - A1, A13.
 - List of people in charge per Tier-2
 - Still missing names for some Tiers (Boston, AGLT2)
 - My estimation for Oct, Nov
 - 1 FTE for the development issues (Wensheng, Patrick, Sudhamsh, Pedro, Kristy, Hiro)
 - 1.8-2.0 FTE for the operational issues, users support and tests (including GRID SW, like FTS, etc)
- FYI : CMS Tier-2s (at least several actively running) dedicate 2 FTE per Tier-2 (2 computing specialists, 2 physicists)

DDM Operations priority list

- DDM 0.2.12 deployment on US ATLAS sites
- Request to DDM core team :
 - DQ2 subscription system
 - retrying and retransferring policy need to be more intelligent. All clouds (OSG and LCG) have similar problem with data transfer from Tier-2s to Tier-1 (recent case with 50 files for streaming jobs, when 4 files were transferred but not registered, followed by deletion (manually) of datasets, etc)
 - For LCG the problem can be more severe in the future because of operational model
 - TOKYO and BEIJING have files catalog in LYON, Melbourne in Taipei
 - Dataset location information
 - very often it is incorrect and it is the source of many mistakes
 - Monitoring
 - More predictable releases policy
 - Request to DDM developers to provide DQ2 0.3 for tests as early as possible
 - Central services stability
 - DB outages and computers saturation (2-3 times per week).
 - The problem persists since Jul.
 - mySQL DB will be replaced by ORACLE.
 - New catalogs structure in DQ2 0.3
 - Sites services
 - Hangs of site service, no way to get alarm from monitoring. Login on VO box and log files digging is required.

DDM Operations priority list (cont)

- DQ2/Panda
 - Policy in case of network or central services are not available
 - Files registration in TID datasets only after files are replicated to BNL and registered in LRC
- Minimize human intervention in DDM operations
 - Generic integrity check scripts for all ATLAS sites
 - Scripts from A6
 - More scripts to check dataset content on site
 - Kaushik's proposal : to get info for datasets produced by Panda using dataset name or task ID
- Recovering procedures
 - after failure of FTS, site services, network, etc
 - the worst case (TW and CAN clouds) 2-3 days of downtime
- LFC vs LRC evaluation
- 2007 Functional tests will address DDM performance issues