



# Improvements in data access efficiency for jobs

*how the WLCG Technical Forum  
might contribute*

Maarten Litmaath

Andrea Sciabà

CERN

EGEE'09

21-25 September, 2009

Barcelona



# Data access efficiencies need improving

- Lower failure rates, better CPU-wallclock ratios
- Depends on which experiments are supported by the site
  - Maybe OK for one, not so good for another
- Depends on the site layout
  - Network infrastructure
  - Usually cannot be adapted in the short term
- Depends on the access protocols and their usage
  - Read-ahead good for sequential, bad for random access
  - RFIIO read ahead buffer size cannot (yet) be set by client code
    - One value per WN set by admin
  - Different experiments need different buffer sizes
    - Due to different event and processing models
  - Phase out some protocols in favour of others?
    - Probably not feasible in the short term



# Current status

- ATLAS are determining the best access method per site
  - Activity mostly driven by clouds that want to improve
  - Central steering through HammerCloud team
    - Global tests foreseen early Oct.
  - Best method recorded per site in central configuration
  - CERN T3 testing: see next slide
- CMS intend to investigate improvements early Oct.
  - Taking note of ATLAS results
  - Each CMS site configures the protocol to be used by jobs
- LHCb ran tests at T0 and T1 sites in May
  - High failure rates at some T1 sites
  - dCache tuning (DCACHE\_RAHEAD=1, DCACHE\_RA\_BUFFER=100k)
  - CPU/wallclock can rise to 70-80 % after tuning
    - For successful jobs



# ATLAS tests of CERN T3 CASTOR

- Preliminary recommendations for ATLAS by Max Baak
- Xrootd and RFIO read-ahead very inefficient
  - Lots of unnecessary data transferred, saturating NIC
  - Large spread in job times
- RFIO buffer size cannot (yet) be changed by client
  - Do not use RFIO
- Xrootd works stable and fast with files present on WN disk
  - No buffering
- Two approaches to deal with files residing only on CASTOR
  - Access them remotely through Xrootd, no buffering
  - Copy them to the WN using the intelligent FileStager



# How can WLCG Technical Forum help?

- Discuss issues for improvement between WLCG stakeholders
  - Input on a common WLCG position to EGEE, EGI, OSG, ...
- Also longer term needs w.r.t. services and middleware
  - Sustainability and evolution of the existing middleware in the light of changing technologies and experience
- Can we think (again) of common solutions in areas where existing practice is weak?
- Needs to represent all the stakeholders
  - Experiments, sites, grid projects, ...
  - Bring in the appropriate experts depending on the topic
- Does not take decisions
  - Should produce clear documents for discussion in the GDB and potential agreement in the MB



# Membership

- Experiments
  - Various representatives per experiment
- Sites
  - Tier 0
  - All Tier 1s
  - Significant number of Tier 2s with different setups
- Infrastructures
  - EGEE/EGI, OSG, ARC
- Experts, developers
  - May also be brought in as needed
- Further details on TWiki page
  - <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTechnicalForum>



# Data Management

- Efficient, scalable data access by jobs ← Here we are!
  - Local vs. remote
  - Protocols
  - Throttling
  - T3 farms vs. T2 load
- ACLs
- Quotas
- SRM
- Xrootd
- GPFS, Lustre, NFSv4, Hadoop, REDDNet, ...
  - “file” protocol
  - Clouds
- Issues specific to some implementation(s)
  - BeStMan, CASTOR, dCache, DPM, StoRM



# Job Management

- CREAM
- WMS
- ARC
- Condor-G,-C, GT4
- MyProxy failover
- Pilot jobs
  - Glexec
  - Frameworks
- Virtualization
- Clouds
- Shared SW area scalability
  - ALICE: BitTorrent
- PROOF





# And more...

- Security
  - Vulnerabilities
  - Consistency
- Information system
  - Fail-over
  - GLUE 2.0
- Monitoring
  - Jobs
  - Consistency, consolidation
- Accounting
  - Messaging system
  - Storage



# How would it work?

- Mailing lists members were asked to report their top 3 issues for possible discussion
  - See following pages
- The issues will be categorized and ranked on the TWiki page
  - Many top issues may not be improvable in the short term
- TWiki workspaces will be set up for promising items
- Each item has a small working group to deliver a short report
  - Mailing list members can contribute to any WG
- Draft versions are announced on the mailing list regularly
- Consensus need not be unanimous
  - “Group A/B/C prefer this, group D/E/F prefer that”
- Reports are presented in the GDB and/or MB



# Collected top 3 candidates (1)

- DPM: short and long-term future
- Error messages
- Logging
- Fault tolerance
- SRM scalability
- Rate of I/O errors
- dCache administration and troubleshooting
- Pilot jobs push problems from users to sysadmins
- dCache and DPM: different ports, different security layers
- gLite UI: difficult for users to install



## Collected top 3 candidates (2)

- WMS: Condor-G not VOMS-aware
  - WMS: VOViews should be decisive where present
  - Wrong info published by sites making them attractive
  - Shared area at many sites
  - dCache: optimization
  - dCache: DCAP service instability & client recovery -- fixed?
  - Local protocols: client vs. server vs. application versions
  - Temporary unavailability of T0D1 files
- 
- SRM vs. data access patterns outside HEP
  - Backward/forward compatibility of clients
  - Documentation, APIs and standardization, e.g. for Java clients



## Collected top 3 candidates (3)

- Dodgy packaging, requiring workarounds in YAIM
- Support for multiple SubClusters per CE
- Tight connection of CE to Torque installation
- Benchmarks: SW must be bought, results debatable
- YAIM: example site-info.def does not show new variables
- Reliability of jobs, random failures, much better in ARC
- YAIM: understanding configuration problems is difficult, much better in ARC
- ETICS: difficult to get package sources, ARC does this right



## Collected top 3 candidates (4)

- Virtual machines: agreement on remotely generated images
- SRMv2.2: some complexities/issues due to lack of a standard for synchronizing the SE and experiment catalogues
- WMS/LB: operation has required a lot of effort
- WMS/LB: should be stateless and load-balanceable
- Change strategy for sites: stability  $\leftrightarrow$  upgrades
- Upgrade rollbacks often not possible
- Pool accounts: scalability
- IPv6



# Collected top 3 candidates (5)

- Documentation
- Configuration complexity, customization
  - Site upgrades require long downtimes
  - SE downtime implies CE downtime
- Service stability and failover
  - Services should protect themselves against misuse
  - Failover can allow for transparent upgrades
- Collaboration and communication between gLite, EGEE, deployment, experiments: avoid duplication and surprises
  - Sometimes it seems we are too ambitious with what we ask from the sites



# Conclusions

- Data access: need to further investigate performance and collect recommendations for sites and experiments
  - Improve exchange of information among experiments on tests
- The recently created WLCG Technical Forum
  - Is a forum open to experts, VOs, sites, developers
  - to identify the most important problems with the current middleware
  - to give recommendations on how to solve these problems in the short to medium term
  - To propose longer term strategies
- How should it work?
  - Rely on the expertise of the participants
  - Form working groups on selected topics
  - Conduct dedicated investigations when available information is insufficient