

Paul Nilsson

EVENT SERVICE PILOT PLANS

Event Service issues, plan and timeline: pre-production on the grid

Currently Knows ES Issues

- “Checksum problem” -> AthenaMP issue
 - Object store transfers failed with checksum error
 - Initially believed to be a bug with xrscp or the CERN object store
 - Turned out to be a problem with AthenaMP that was not finished with the output file after pilot handover
 - Peter van Gemmeren is looking into it
 - Cannot run merge trf successfully until this has been resolved
 - Trf fails the file corruption test
- Pilot is currently staging in TAG and EVGEN files
 - Not exactly following the Event Service model but no choice at the moment..
 - Need/want to switch to Event Index as soon as possible/available (otherwise complicated with HammerCloud tests) – TAG file not needed any more
 - If AthenaMP can directly access EVGEN file pilot don't need to stage it in (Vakho: no need to stage in EVGEN file any more)
- Checksum test currently not possible at CERN Object Store
 - Coming soon: “we can update to a newer XRootD version and configure server side check summing. We can also store the checksum as an attribute in the object store and make it retrievable” (Andreas Peters)
 - (Incidentally, HTTP access at end of August - useful for log files, easily reachable from the monitor)

Immediate Plans

- Finish complete Event Service test job suite as soon as possible
 - Initial event service test job runs well – except for mentioned output file problem
 - Improve error handling
 - For various problems that can happen with AthenaMP, Token Extractor, Event Index and stage-out failures
 - Eventually extend to use job recovery pilot feature, i.e. place output file in special directory and leave it for a later pilot to recover it – some complications to job recovery algorithm expected
 - Merge step pilot updates done, final tests pending AthenaMP update
 - Token Extractor update now available (since Aug 17), implement update in TE handling in pilot immediately (~tomorrow)
 - Start using Event Index as soon as service is available

Immediate Plans II

- Wen Guan is up to speed on running Event Service test jobs
 - Wen is following a work plan outlined a couple of weeks ago
 - Currently focusing on testing at BNL (Paul is running tests at UTA), but also on other MCORE queues
 - Test jobs write/read files from BNL S3 Object Store
 - First step towards running tests on Amazon spot market
 - Implemented checksum on server side (stored as metadata), already tested and verified
 - Test event service on the Amazon spot market as early as possible to identify any issues but also to be able to declare that the event service is working on the spot market
 - Amazon S3 cloud storage (s3.amazonaws.com) already setup by Hiro and tested by Wen
- Discuss steps needed for Event Service on HPC:s (today-tomorrow)
 - Toy Yoda from Vakho, rank 0-1 communications (with Tadashi)

(Slightly) Longer Term Plans

- Towards pre-production and production
 - Wen will help with scaling up tests
 - From current simple tests to pre-production
 - Use Pilot Release Candidate Test Framework for pre-production event service testing. This includes involving the HammerCloud team to help setup event service test jobs
 - Want/Need to use Event Index, don't want jobs to have any input files
 - Site info need to be updated for many sites; especially schedconfig.copyprefixin (needed for SURL to TURL conversion)
 - Timescale ~1 month (pending mentioned bugs and event index availability)
- Bring the event service workflow into US HPCs
 - Requires development of new pilot module RunJobHPCEvent (Paul/Wen in collaboration with Vakho/Danila/Tadashi/..)
 - Code will be based on existing RunJobEvent module and Vakho's toy Yoda code (to be developed)
 - Test on BNL Blue Gene?
 - Hard deadline for demonstration in November