

CESNET Activities in JRA1

Aleš Křenek, on behalf of CESNET JRA1 team

- responsibility for *job tracking services*
 - grid job is the primary entity of interest
 - gather information (raw) on job life
 - provide information (cooked) on job state
 - store it for future reference
- **Logging and Bookkeeping (L&B)**
 - just in time
- **Job Provenance (JP)**
 - long term archive

- joint effort of JRA1, SA3, SA1, and applications
- tune services to match *acceptance criteria*
 - 10k jobs/day sustained rate for 5 days
 - no service restart required
 - less than 1 % failure rate
- dedicated WMS and L&B installations
- using production infrastructure (CE's and SE's, info system)
- developers' supervision
- fast bugfix turnaround (no formal process)
- acceptance criteria met well (16k jobs/day, 0.3 % failures)
- formal release (patch #1251) as final outcome
- L&B handles over 1M jobs/day :-)

- joint effort of JRA1, SA3, SA1, and applications
- tune services to match *acceptance criteria*
 - 10k jobs/day sustained rate for 5 days
 - no service restart required
 - less than 1 % failure rate
- dedicated WMS and L&B installations
- using production infrastructure (CE's and SE's, info system)
- developers' supervision
- fast bugfix turnaround (no formal process)
- acceptance criteria met well (16k jobs/day, 0.3 % failures)
- formal release (patch #1251) as final outcome
- L&B handles over 1M jobs/day :-)

- joint effort of JRA1, SA3, SA1, and applications
- tune services to match *acceptance criteria*
 - 10k jobs/day sustained rate for 5 days
 - no service restart required
 - less than 1 % failure rate
- dedicated WMS and L&B installations
- using production infrastructure (CE's and SE's, info system)
- developers' supervision
- fast bugfix turnaround (no formal process)
- acceptance criteria met well (16k jobs/day, 0.3 % failures)
- formal release (patch #1251) as final outcome
- L&B handles over 1M jobs/day :-)

- numerous non-trivial bugs fixed
- inter- and intra-service race conditions
 - production load is necessary to trigger them
- occurrence after several days/millions of requests
 - e.g. CRL reload in Globus GSI libraries
- hitting unforeseen limits
 - e.g. number of filedescriptors (not a leak)
 - `select()` can't handle `fd > 1024`
- **unique environment**
 - production load to trigger the problems
 - yet enough control to trace them down


- numerous non-trivial bugs fixed
- inter- and intra-service race conditions
 - production load is necessary to trigger them
- occurrence after several days/millions of requests
 - e.g. CRL reload in Globus GSI libraries
- hitting unforeseen limits
 - e.g. number of filedescriptors (not a leak)
 - `select()` can't handle `fd > 1024`
- **unique environment**
 - production load to trigger the problems
 - yet enough control to trace them down

- support “DAGless” job collections
- new WMS functionality, followed by L&B
- set of many jobs
 - common specification, inputs etc.
 - independent on each other
- master job “disappears” after break-up
- its state still to be reported
 - deduced from aggregate subjob state only
 - non-trivial logic due to resubmissions

- PBS
 - information gathered from all main components
 - almost complete state machine
 - experimental deployment in MetaCentre (CZ grid)
- Condor
 - experimental information gathering and state machine
- results presented at HPDC workshop

- JRA1-wide activity
- get rid of unnecessary external dependencies
- migrate to Etics builds
- disentagle internal dependencies
- not critical for L&B and JP
- good opportunity for internal cleanup
- definition of common logging format (JRA1-SA1 session on Friday)
- work in progress

- standardization activity
- attempt interoperability of participants' systems
- challenge workflow – medical image processing
- execution data captured by each system
- split into 3 parts and combined
- challenge queries to be answered
- presented at HPDC workshop



- JP used as “job catalogue”, little customization effort
 - computational chemistry application (small newcomer community)
 - Atlas experiment (large well-established community)
- overview and details of jobs
- application and middleware oriented views
- visualization of results
- support for new submission
- presented at User Forum and CHEP
- see poster #30 for details

