

CZ Cluster Summary

Aleš Křenek, on behalf of CESNET JRA1 team

- this talk
 - L&B maintenance and development
 - support non-gLite jobs in L&B
 - preview testbed
- following talk
 - Etics build and related issues
 - gLite restructuring
 - presented by Jan Pospíšil
- common logging format
 - talks and discussion on Tuesday
 - presented by Steve Fisher
- JP status and plans on Wednesday
 - demo at EGEE User Forum
 - 2nd Provenance Challenge
 - the status and plans

- not many in general, sw is stable
- bug #26811: request for backport 3.1 WS interface fixes to 3.0
 - done, patch to be released
 - good news – there are WS interface users
- patch #1153: “shared” socket between mysql and syslog
 - stupid bug, serious consequences – causes L&B to hang
 - bad news – 7 weeks between fix and production release
 - still room for improvements in software release process
- VOMS corrupts SSL context
 - SSL connection established correctly
 - VOMS server of client proxy unknown to L&B server
 - following SSL I/O fails
 - not easy to reproduce deterministically, still not completely understood

- not many in general, sw is stable
- bug #26811: request for backport 3.1 WS interface fixes to 3.0
 - done, patch to be released
 - good news – there are WS interface users
- patch #1153: “shared” socket between mysql and syslog
 - stupid bug, serious consequences – causes L&B to hang
 - bad news – **7 weeks** between fix and production release
 - still room for improvements in software release process
- VOMS corrupts SSL context
 - SSL connection established correctly
 - VOMS server of client proxy unknown to L&B server
 - following SSL I/O fails
 - not easy to reproduce deterministically, still not completely understood

- not many in general, sw is stable
- bug #26811: request for backport 3.1 WS interface fixes to 3.0
 - done, patch to be released
 - good news – there are WS interface users
- patch #1153: “shared” socket between mysql and syslog
 - stupid bug, serious consequences – causes L&B to hang
 - bad news – **7 weeks** between fix and production release
 - still room for improvements in software release process
- VOMS corrupts SSL context
 - SSL connection established correctly
 - VOMS server of client proxy unknown to L&B server
 - following SSL I/O fails
 - not easy to reproduce deterministically, still not completely understood

- prototype implementation didn't handle job resubmission
 - unclear semantics at design time
- current view: subjob state mapped into class:

DONE	done/OK
CLEARED	cleared
ABORTED	aborted
RUNNING	running
REST	anything else, including Done/Failed
- collection state from subjob histogram only (in this order):

Submitted	no events of subjobs logged yet
Running	$\text{hist}[\text{RUNNING}] > 0$
Cleared	$\text{hist}[\text{CLEARED}] == N$
Done/OK	$\text{hist}[\text{DONE}] + \text{hist}[\text{CLEARED}] == N$
Aborted	$\text{hist}[\text{ABORTED}] + \text{hist}[\text{DONE}] + \text{hist}[\text{CLEARED}] == N$
Waiting	otherwise
- **is this semantics acceptable?**

- design fault in handling attributes of Event and JobStatus classes in C++ API
- adding support for non-gLite jobs \Rightarrow backward ABI incompatibility of several 3.1 releases
- non-trivial API revision to avoid similar problems in future
 - rearrangement of struct and enum's
 - backward incompatibility of ABI
 - no need to change any client (besides recompilation)
 - database format preserved (no upgrade procedure required)
- being done thoroughly
 - usual “next week” phenomenon
 - high priority task

- extend LB to R-GMA export (VOMS attributes and UI hostname)
 - features requested in bugs #21575, #21578
 - L&B and WMPProxy (AFAIK) part is ready
 - LCGMon schema upgrade required – **to be coordinated**
- filedescriptor “leak” in L&B proxy connections
 - not a true leak
 - L&B context instance not used for long time by client
 - still keeps open socket, despite closed by server-side slave process pre-emption
 - to be fixed in L&B library by sharing the socket among multiple contexts
- L&B sequence code incompatibility
 - “Cannot restart logmonitor” thread
 - dagless collection support required adding new field
 - incompatible with long-lived jobs
 - fixed in recent 3.1 release

- extend LB to R-GMA export (VOMS attributes and UI hostname)
 - features requested in bugs #21575, #21578
 - L&B and WMPProxy (AFAIK) part is ready
 - LCGMon schema upgrade required – **to be coordinated**
- filedescriptor “leak” in L&B proxy connections
 - not a true leak
 - L&B context instance not used for long time by client
 - still keeps open socket, despite closed by server-side slave process pre-emption
 - to be fixed in L&B library by sharing the socket among multiple contexts
- L&B sequence code incompatibility
 - “Cannot restart logmonitor” thread
 - dagless collection support required adding new field
 - incompatible with long-lived jobs
 - fixed in recent 3.1 release

- extend LB to R-GMA export (VOMS attributes and UI hostname)
 - features requested in bugs #21575, #21578
 - L&B and WMPProxy (AFAIK) part is ready
 - LCGMon schema upgrade required – **to be coordinated**
- filedescriptor “leak” in L&B proxy connections
 - not a true leak
 - L&B context instance not used for long time by client
 - still keeps open socket, despite closed by server-side slave process pre-emption
 - to be fixed in L&B library by sharing the socket among multiple contexts
- L&B sequence code incompatibility
 - “Cannot restart logmonitor” thread
 - dagless collection support required adding new field
 - incompatible with long-lived jobs
 - fixed in recent 3.1 release

- semi-parallel job registration to L&B proxy and L&B server
 - $2 \times$ speedup for large collections theoretically
 - job owner determined from SSL \Rightarrow race condition if L&B proxy path is faster
 - solution: “already logged directly” flag at L&B event, L&B proxy won't forward
- certificate DN string representation
 - reports on authorization failures
 - different SSL library version provide different DN strings
 - solution: use official X509 representation and some canonical string form
 - longer term (serious compatibility problem)

- semi-parallel job registration to L&B proxy and L&B server
 - $2 \times$ speedup for large collections theoretically
 - job owner determined from SSL \Rightarrow race condition if L&B proxy path is faster
 - solution: “already logged directly” flag at L&B event, L&B proxy won't forward
- certificate DN string representation
 - reports on authorization failures
 - different SSL library version provide different DN strings
 - solution: use official X509 representation and some canonical string form
 - longer term (serious compatibility problem)

- interlogger spool files
 - at least two `sync()` calls per logged event
 - non-trivial I/O load on WMS node, future bottleneck
 - changed structure: one file per destination (rather than job) – done
 - new API call to log multiple events at time – to be done
 - ▶ necessary to preserve semantics “stored locally when API returns”
 - must be used by WMS to get full performance gain
 - ▶ subjob registration
 - ▶ bulk matchmaking
 - ▶ ...

- motivation and detail reported in March in Catania
 - requested functionality (Condor, *META Centre* Czech national Grid)
 - foreseen unified view on jobs in multiple “universes”
 - specific events and state machine for each “universe”
- PBS support in pre-production state
 - deployed in *META Centre* (approx. 1000 CPU's)
- Condor support
 - testing instance, able to handle typical job life
- presentation at HPDC workshop on Grid monitoring next week

- glite CE torque + WN
 - skurut67-4.cesnet.cz, skurut67-5.cesnet.cz
- CREAM CE
 - skurut67-5.cesnet.cz
- “experimental” 3.1 WMS
 - skurut67-7.cesnet.cz
- current 3.1 L&B server (operational)
 - skurut68-1.cesnet.cz – fast purge & JP import
 - skurut68-2.cesnet.cz – normal purge & JP import
- JP Primary Storage (to appear soon)
 - umbar.ics.muni.cz
- JP Index Server (to appear)
 - skurut68-1, skurut68-2
 - elementary deployment (see my JP talk), data from these machines

- gLite restructuring
- common loggin format
- Job Provenance
 - Provenance Challenge
 - elementary deployment
 - stress tests (handle part of Atlas traffic)
 - support applications (UF demo follow-up)
- multiple job types in L&B and JP
 - extending the set (CREAM, BLAH, . . . , but also service state)
 - links between the universes
 - ▶ gLite job x waited for data tranfer A first, it became PBS job z via CREAM job y , it failed due to CE C failure, and it was resubmitted as Condor job w . . .
 - further cleanup of L&B internals