

CZ Cluster Summary

Aleš Křenek, on behalf of the CESNET team

Bug fixes

- only minor issues on 3.0
- all critical fixes ported to 3.1

Migration to gSoap 2.7.9b

- multiple versions in `security.gsoap-plugin`
 - motivation: migration between gsoap versions difficult
 - encapsulation of gsoap library – more versions can co-exist
 - all lib sources in CVS – not clean but avoids multiple dependencies
 - build-time gsoap dependence – only for examples
- L&B and JP migration
 - more difficult than expected
 - almost done this week

Bug fixes

- only minor issues on 3.0
- all critical fixes ported to 3.1

Migration to gSoap 2.7.9b

- multiple versions in `security.gsoap-plugin`
 - motivation: migration between gsoap versions difficult
 - encapsulation of gsoap library – more versions can co-exist
 - all lib sources in CVS – not clean but avoids multiple dependencies
 - build-time gsoap dependence – only for examples
- L&B and JP migration
 - more difficult than expected
 - almost done this week

Non-DAG collections

- support ready in L&B, available at preview testbed
- reporting aggregate collection state
 - not declared to be urgent
 - something reasonable should be reported either
- events coming from WMS being evaluated

L&B database purge

- getting more and more urgent
- reported as configuration bug #23608, detailed info provided
- already overgrown database
 - standard (on-line) approach does not work smoothly (timeout)
 - database is not shrinken
 - work on rescue (off-line) script in progress

Non-DAG collections

- support ready in L&B, available at preview testbed
- reporting aggregate collection state
 - not declared to be urgent
 - something reasonable should be reported either
- events coming from WMS being evaluated

L&B database purge

- getting more and more urgent
- reported as configuration bug #23608, detailed info provided
- already overgrown database
 - standard (on-line) approach does not work smoothly (timeout)
 - database is not shrinken
 - work on rescue (off-line) script in progress

Build problems

- turned to be trivial (libtool needs “gcc” named compiler)
- hopefully fixed now

Broken functionality in 3.1

- difficult to reproduce
- combination with specific Globus and VOMS versions
- tracked down to be a name clash between these
- fixed in VOMS

- slow progress, mostly due to Etics issues
- two planned stages
 - conservative: prune 3.1 dependencies with minimal L&B code changes
 - more aggressive: split modules, reshuffle code, etc.
- details in Jan's talk tomorrow
- overall deadline (June) still realistic

Motivation

- standing requests for additional job types
 - advance reservations
 - Condor
 - PBS (*META Center* – Czech grid)
- unified view on job state
 - convenient for users
 - aggregate information from multiple domains
 - ▶ gLite + Condor + PBS + data transfer

Changes in L&B core

- generic support for multiple job types
 - using `job_type` field in registration event
 - job-type specific switch to call state machine
- “grey” jobs
 - registration needn't be the first event anymore
 - put events in quarantine, invisible for queries
 - either validated by registration or deleted after timeout
- extended authorization
 - Kerberos, UNIX identity, ...
 - currently grid-mapfile
 - foreseen mechglue (multiple mechanisms under GSS)

Job-type specific events and states

- defined in .T files
- except of state machine everything is generated automatically

Changes in L&B core

- generic support for multiple job types
 - using `job_type` field in registration event
 - job-type specific switch to call state machine
- “grey” jobs
 - registration needn't be the first event anymore
 - put events in quarantine, invisible for queries
 - either validated by registration or deleted after timeout
- extended authorization
 - Kerberos, UNIX identity, ...
 - currently grid-mapfile
 - foreseen mechglue (multiple mechanisms under GSS)

Job-type specific events and states

- defined in .T files
- except of state machine everything is generated automatically

Log cruncher

- we cannot rely on instrumenting third-party code anymore
- new component, lightweight counterpart to gLite Log Monitor
- generic part
 - read log files, handle their rotation
 - keep persistent state for recovery
- job-type specific modules
 - detect events in the log
 - crate and send L&B events

Current state

- most L&B core changes done
- PBS support works quite well
 - correct event ordering is still an issue
- work on Condor support started
 - more difficult log file parsing
- paper submitted for HPDC grid monitoring workshop

Log cruncher

- we cannot rely on instrumenting third-party code anymore
- new component, lightweight counterpart to gLite Log Monitor
- generic part
 - read log files, handle their rotation
 - keep persistent state for recovery
- job-type specific modules
 - detect events in the log
 - crate and send L&B events

Current state

- most L&B core changes done
- PBS support works quite well
 - correct event ordering is still an issue
- work on Condor support started
 - more difficult log file parsing
- paper submitted for HPDC grid monitoring workshop

- followup of successful Challenge 1
- focus on co-operation
 - challenge workflow is split into 3 parts
 - participants exchange provenance data on the parts
 - the same queries should be answered
- first stage (define interchange format) done in February
- to be concluded at HPDC workshop in June
- participation in Provenance Challenges linked in NA5 list of standardization activities

- pilot application in computational chemistry
 - interaction of protein with multiple small molecules, eg. potential drugs
- many related jobs, Job Provenance used to keep the track
- focus on application – user attributes
- supported search for interesting patterns

jp_docking_proto

Trajectory XY: Room for your buttons and other controls

| | | | | | | |
|------|--|----|--|-----|--|-----|
| 1000 | | 5 | | 2 | | 1 |
| 1100 | | 3 | | | | 3 |
| 1200 | | 3 | | 3 | | 1 |
| 1300 | | | | | | |
| 1400 | | 5 | | | | |
| 1500 | | | | 2 | | |
| 1600 | | | | | | 1 |
| 1700 | | | | | | |
| 1800 | | | | 5 | | 78 |
| 1900 | | | | | | |
| 2000 | | | | | | |
| 2100 | | 85 | | 193 | | 77 |
| 2200 | | | | 196 | | |
| 2300 | | | | 1 | | 167 |
| 2400 | | | | 206 | | |
| 2500 | | 89 | | 104 | | 188 |
| 2600 | | | | 121 | | 220 |
| 2700 | | 23 | | | | |

Job Radius Flex. ...

Refresh Resubmit Selected New

| Parameter | Value |
|-----------|-------|
| | |

Show Jobs
 Show Metrics
 Show Annotations

eGEE Room for other buttons, icons, ads, etc. :-)

Quit

Status

- main goals
 - have a local L&B and JP build
 - define internal procedures for configuration management
 - start gLite restructuring
- first attempts around EGEE conference (September)
- systematic effort since November
- approx. 40 person-days spent
- rather tiny results
 - we are able to build current 3.1 configurations
 - we have cumbersome procedures for config management
- **good news either – there is some progress**

Reasons

- lack of documentation
- instability of Ethics client
 - the issues are blocking despite being solved fairly quickly
 - 0.7.3 (quite recent) is the first quasi-stable version
- terrible inefficiency of Ethics client
 - getting better, 0.7.3 is already tolerable
- bad support for configuration management
- orientation on integration, not friendly to developers
 - I'd be glad if this statement is reputed at this meeting

Reasons

- lack of documentation
- instability of Ethics client
 - the issues are blocking despite being solved fairly quickly
 - 0.7.3 (quite recent) is the first quasi-stable version
- terrible inefficiency of Ethics client
 - getting better, 0.7.3 is already tolerable
- bad support for configuration management
- orientation on integration, not friendly to developers
 - **I'd be glad if this statement is reputed at this meeting**

Etics client inefficiency

- communication with Etics server
 - order of minutes in this phase when checking out L&B
 - maybe too much “ping-pong”, streaming bigger chunks preferred
 - bandwidth is not problem, latency is (we've got 20ms)
- etics-checkout --merge can take up to 1 hour!
- strace witness on etics-* commands
 - ...
 - say 'Done!'
 - do nothing usefull for upto several mins, consuming 100% 3GHz CPU
 - call `exit(2)`
- **very annoying when done dozens times a day**

Configuration management

- `etics-tag` does not work at all
- `etics-configuration` is the most buggy part of client
- editing `.ini` is tedious and error prone
 - difficult to read format (eg. order of properties)
 - multiple occurrences of version information
 - web interface is not an option for routine usage
- lack of configuration history (CVS-like), including comments
- missing support for “branch” configuration
 - including easy “create release config now” functionality
 - support for merging branch configurations
- unclear semantics of property inheritance, configuration overriding etc.
 - we faced false circular dependencies detected
 - we don't know whether it is feature or bug
- local dependencies (not from repository) welcome

Proposed tutorial

- start with local build of a specific (last known working) configuration of a subsystem
- fix a repoted bug, check the fix, and commit into CVS
- let another developer checkout the current code from the branch, and augment the bug fix eventually
- create new tags and configurations, assuming the bug fix spans multiple components