



gLExec Accounting Solutions in OSG

Overview

- Pilot Jobs and gLExec in OSG
- Process Monitor orked by gLExec GUMS PlugIn
- Integration with the Gratia Accounting System

Mar 27, 2008

Middleware Security Group Meeting

Igor Sfiligoi and Gabriele Garzoglio
Computing Division, Fermilab



Context

- A pilot job is submitted to a cluster
- The pilot job lands on a worker node
- The pilot job calls gLExec to execute user jobs
 - The pilot job interacts with the site security system via gLExec
 - Important to account the used resources to the proper user / fqan



gLExec in OSG

- NIKHEF provides gLExec, LCAS and LCMAPS frameworks
- VO Services project provides the PRIMA/GUMS plugin
 - Contacts GUMS to get the correct mapping
 - Spawns a Monitor process for accounting



gLExec GUMS Plugin Monitor

- Process started by the GUMS plugin
- Logs user information to a local file
 - DN, VO, FQAN, parent processes
- Tracks the descendants of gLExec
 - Log info for long lived processes only (> 5min)
 - Periodically logs accounting info for CPU usage (Usr+Sys)



Process Tracking

- Use Condor-provided “procd”
 - Distributed via VDT
- Use a secondary GID of the processes for tracking
 - In this model, the system reserves a range of sGIDs dedicated to process tracking
 - gLExec will set an sGID before starting the user job
 - Only superuser can change list of sGIDs
 - Reliably tracks also daemonized processes



GRATIA Usage of the Records

- GRATIA probe at the WN parses the gLExec GUMS Monitor log file
- Collects:
 - DN, VO, FQAN
 - wallclock time, Usr and Sys CPU time
- Information is sent to the Gratia Collector



Summary

- OSG uses accounting information from gLExec
 - Accounting is provided by gLExec GUMS plugin...
 - ...it would be better to move it elsewhere
- Reliable process tracking based on sGID added recently
 - Can properly account for daemonized processes