gLExec Integration with the ATLAS PanDA Workload Management System

E Karavakis*, F Barreiro, S Campana*, K De, A Di Girolamo*, M Litmaath*,
T Maeno, R Medrano, P Nilsson, T Wenaus for the ATLAS Collaboration

Introduction

ATLAS user jobs are executed on Worker Nodes (WNs) by pilots sent to sites by pilot factories. This has greatly improved job reliability and although it has clear advantages, such as making the working environment homogeneous, the approach presents security and traceability challenges. Jobs inherit the identity of the pilot submitter and the payload jobs then execute directly under that same identity on a WN. This exposes the pilot environment to the payload, requiring any pilot 'secrets' such as the proxy to be hidden; it constrains the rights and identity of the user job to be identical to those of the pilot; and requires sites to take extra measures to achieve user traceability and user job isolation. To address these security risks, gLExec can be used to let the payloads for each user be executed under a different UNIX user id that uniquely identifies the ATLAS user.

Architecture

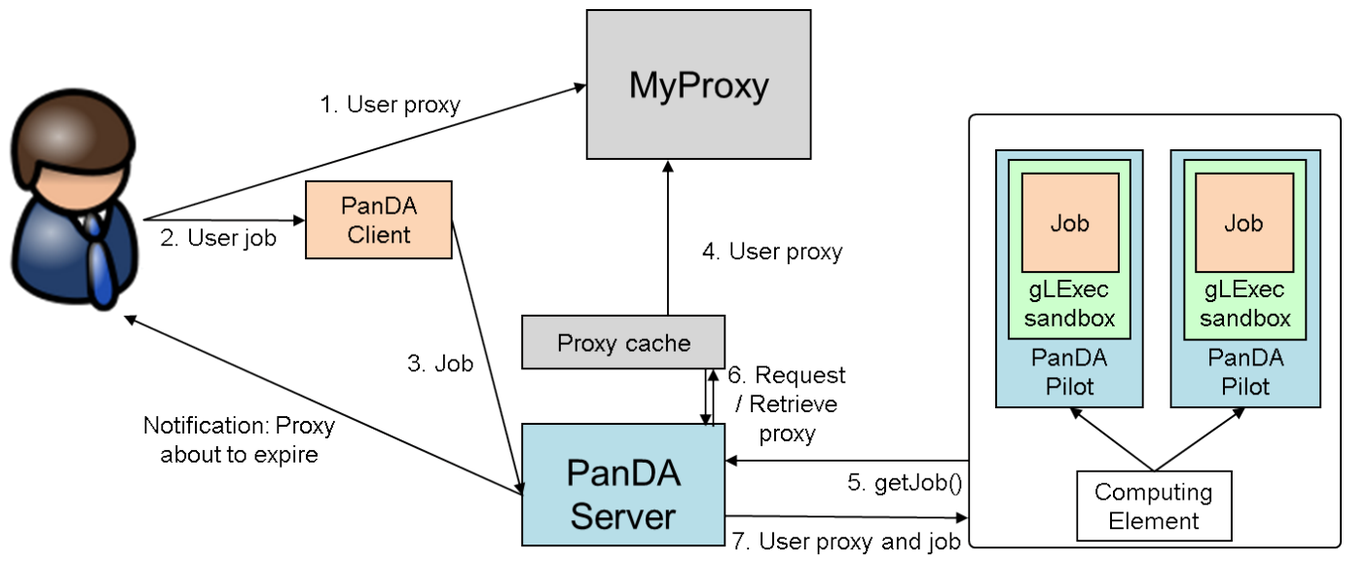

Proxy Delegation

- Users will be asked to upload a long-lived proxy on a MyProxy service every time they get a new user certificate
- The PanDA [1] host certificate is authorised to retrieve delegated proxies

Proxy Cache

- Provides an API to PanDA Server to retrieve proxies
- Simple mechanism to reduce the load on MyProxy
- Cronjob runs every 10 mins, gets 'active' users and uses the cache to check for proxies and their validity
- If proxy is there, it checks its validity and if valid for less than 3 days, it renews it. If not in the cache, it retrieves a plain proxy from MyProxy and adds ATLAS VOMS attributes with a validity of 4 days

PanDA Pilot

- Call to get a job from PanDA server returns the proxy along with the job if the gLExec flag is set for that site
- If flag is set, jobs will run under a gLExec sandbox. If not set, the pilot will work in normal mode without gLExec
- Pings gLExec infrastructure to check if it is OK (with a retry mechanism). If it is not OK, it falls back to the normal mode
- If the user's proxy is not provided the pilot can still run in gLExec mode by using the proxy that has started the pilot
- Switches back to the original environment when job finishes

✓ *This allows the site infrastructure to be tested*

✓ *Not all users have to upload their proxy before the site can be used with gLExec*

Conclusions

✓ The development work and the deployment campaign have been a big success so far as more than half of the ATLAS analysis sites were covered without affecting normal ATLAS activities thanks to the flexibility of the implemented system

✓ Ongoing effort that will be handed over to ATLAS Distributed Computing Operations eventually

✓ Some minor issues with specific sites are under investigation

Setting up gLExec queues

- Easy to set up a new gLExec queue or change an existing one by setting the gLExec flag in AGIS [2] to 'True' or 'test'

Deployment Campaign

- Campaign started in September 2014 with CERN-PROD and by slowly adding more Tier 1 sites
- Enabled in **61** analysis sites out of **94** sites with an online analysis queue (as of 5th of March 2015)
- **10** Tier 0+1 sites, **49** Tier 2 sites and **2** Tier 3 sites → *most worked "out of the box"*
- 18 GGUS tickets against sites to follow up issues on their gLExec infrastructure and configuration
- Sites are constantly tested with HammerCloud [3] stress and functional tests → *more than 1.2 million successful jobs!*

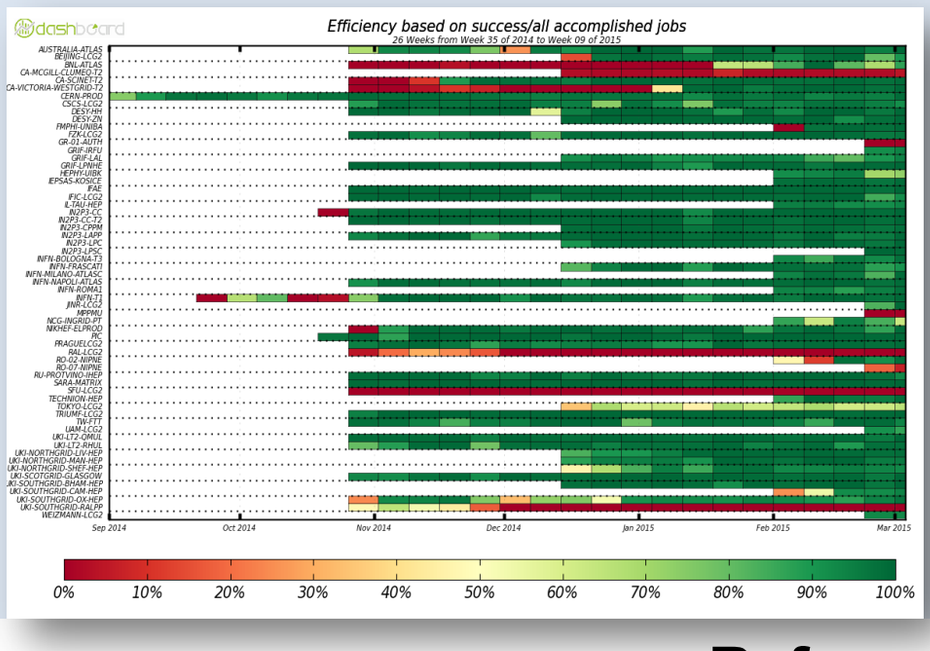

References

[1] CHEP ID #144, T. Maeno : The Future of PanDA in ATLAS Distributed Computing
[2] CHEP ID #168, A. Anisenkov : AGIS: Evolution of Distributed Computing information system for ATLAS
[3] CHEP ID #159, M. Boehler : Improved ATLAS HammerCloud Monitoring for local Site Administration

*CERN IT-SDC : Support for Distributed Computing

LARGE HADRON COLLIDER COMPUTING GRID WORLDWIDE