

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 155

Type: poster presentation

gLExec integration with the ATLAS PanDA workload management system

The ATLAS Experiment at the Large Hadron Collider has collected data during Run 1 and is ready to collect data in Run 2. The ATLAS data are distributed, processed and analysed at more than 130 grid and cloud sites across the world. At any given time, there are more than 150,000 concurrent jobs running and about a million jobs are submitted on a daily basis on behalf of thousands of physicists within the ATLAS collaboration. The Production and Distributed Analysis (PanDA) workload management system has proved to be a key component of ATLAS and plays a crucial role in the success of the large-scale distributed computing as it is the sole system for distributed processing of Grid jobs across the collaboration since October 2007.

ATLAS user jobs are executed on worker nodes by pilots sent to the sites by pilot factories. This pilot architecture has greatly improved job reliability and although it has clear advantages, such as making the working environment homogeneous by hiding any potential heterogeneities, the approach presents security and traceability issues distinct from standard batch jobs for which the submitter is also the payload owner. Jobs initially inherit the identity of the pilot submitter, typically a robot certificate with very limited rights. By default the payload jobs then execute directly under that same identity on a Worker Node. 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 the pilot; and it requires sites to take extra measures to achieve user traceability and user job isolation.

To address these security risks, the gLExec tool and framework can be used to let the payloads for each user be executed under a different UNIX user identity that uniquely identifies the ATLAS user.

This presentation describes the recent improvements and evolution of the security model within the ATLAS PanDA system, including improvements in the PanDA pilot, in the PanDA server and their integration with MyProxy, a credential caching system that entitles a person or a service to act in the name of the issuer of the credential. Finally, we will present results from ATLAS user jobs running with gLExec and give an insight into future deployment plans.

Primary authors: DI GIROLAMO, Alessandro (C); Dr KARAVAKIS, Edward (CERN); DE, Kaushik (University of Texas at Arlington (US)); LITMAATH, Maarten (CERN); NILSSON, Paul (Brookhaven National Laboratory (US)); MEDRANO LLAMAS, Ramon (CERN); MAENO, Tadashi (Brookhaven National Laboratory (US)); Dr WEINAUS, Torre (Brookhaven National Laboratory (US))

Co-authors: BARREIRO MEGINO, Fernando Harald (CERN); Dr CAMPANA, Simone (CERN)

Presenter: Dr CAMPANA, Simone (CERN)

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing