ATLAS gLexec memo
(on behalf of ATLAS Software and Computing Coordination)

Currently ATLAS does not utilize gLexec in production, however many tests have been carried out at T1s and significant work has been spent in adapting the ATLAS analysis framework (Panda) to use gLexec. In this process, ATLAS discovered a variety of issues, which forced to reconsider the usage of gLexec.

1. Different Athena (the ATLAS simulation and reconstruction software) release versions behave differently in the presence of gLexec. In particular, some versions of Athena hang in the setup phase if gLexec is invoked, while they do not hang if gLexec is not invoked. The issue is not fully understood and it has been rather complicated to debug.
2. The gLexec mechanism does not provide a machinery to propagate user proxies to the worker nodes. Therefore, some external component should take care of this and in the ATLAS model this is the MyProxy server. A stolen pilot credential would therefore be able to fetch ATLAS users proxies. To avoid this, considerable development effort should be put in the ATLAS framework. In a scenario without gLexec, a pilot proxy stolen from a Worker Node by a malicious user would have very limited capabilities:
    a. The pilot proxy can NOT submit further jobs, since it is a limited proxy in the WN
    b. The pilot proxy can NOT write into or delete from production storage areas. The pilot proxy can copy data only to scratch endpoints at sites. Such data is automatically cleaned every 15 days.
    c. The pilot proxy can NOT fetch real users proxies.
    d. The pilot proxy can download other users payloads from the Panda server. It is internal ATLAS policy to decide whether this is a problem or not, but it is not a site security violation per se.

    Furthermore, should a site be compromised forcing a revocation of proxies present and potentially exposed at the site, with glExec in use the number of revocations will be much larger, directly affecting large numbers of analysis users, because individual user proxies will have been downloaded to the site for all jobs processed, rather than just the generic pilot proxy in non-glExec mode.
3. Several tests run at various sites have shown an unstable behavior of the sites coupled with the usage of gLexec. Invoking gLexec might work at a given period of time and not work anymore at a later stage. These instabilities are nowhere comparable with the “normal” instabilities of other Grid services at sites (in the sense that they are much more serious). This implies:
    a. The deployment and maintenance of gLexec by sites will be costy in terms of manpower for a prolonged period of time. In fact, the deployment of gLexec has been taking far longer than expected.
    b. The introduction of gLexec will also introduce new failure modes for the experiment which will imply a larger effort (again in term of manpower) for user support, central and cloud operations.

A further point is that glExec is not needed to provide full traceability of user activity at a site. The PanDA analysis system records authenticated identities for all jobs, correlated with local batch IDs, and makes this information available to site operators through monitoring tools.

Considering all the above, ATLAS believes that the amount of security that gLexec would bring to the infrastructure does not motivate the operational cost. In addition, while gLexec covers some vulnerabilities, it creates very serious ones in case of compromised credentials.

The conclusion is that ATLAS does not intend to invoke gLexec on the worker node and proposes the following scenario:

1. Any Grid site should ban the ATLAS pilot DN or if necessary the entire ATLAS VO in case there is a suspicion of compromised credentials or illegal usage of resources at

the site. The banning mechanism should eventually happen via ARGUS, so that ATLAS services can cooperate in the banning of the illegal activity from their side.

2. ATLAS will offer all needed forensics in case of a security violation even if suspect. In general ATLAS will offer forensics about user activity at sites at the request of the site. Tools and documentation can be provided so that the investigation can be carried on by the site itself. The traceability of ATLAS jobs through the Panda system has been demonstrated during the Security Challenges 4 and 5.