ATLAS responses to Ian Bird's questionnaire (in bold) on Security and Job Management

Security: we should re-visit the issues of glexec. Can we do this in a much simpler way by trusting the experiment frameworks? I know what your answers are … but this needs to be explored openly and discussed with the sites. I think we have shown that glexec is difficult to deploy. What might simpler alternatives be?

Currently ATLAS does not use gLexec in production, however many tests have been carried at T1s and significant work has been spent in adapting the ATLAS analysis framework (Panda) to use gLexec.

ATLAS runs "single payload" pilot jobs (pilot start, gets a user payload, runs, exits). Each job in Panda can be associated to a user. This has been proven both in SSC4 and SSC5.

In the current situation (without gLexec) the security risk consists of an ATLAS user (or someone who obtained the proxy of a user) stealing a pilot proxy (/atlas/Role=pilot) and doing something with it. In particular:

- They cannot submit another job (limited proxy on worker nodes not accepted by CEs)
- They can pull a different payload from Panda (and see which analysis another user is doing). It is internal to ATLAS to decide if this is a problem, but this is not a site security violation per se.
- They can store data and delete data into/from a cache area (scratchdisk). Such data is cleaned regularly and can be cleaned/made unavailable aggressively in case of a compromised credential.

ATLAS just concluded SSC5. The main lesson from SSC5 was that in the case of a security violation, the difficult aspect is to contain the incident quickly: declare the credential compromised, propagate the information and react at different levels. If ATLAS were to start using gLexec at **least** one site in production:

- It would be more difficult to compromise pilot credentials. However, containing the incident will be more complicated
    - In the current scenario, one needs to ban the pilot credential from all sites and from the Panda server
    - In case gLexec is in place and user proxies are stored in the MyProxy server, one would need to ban **in addition** all user proxies stored in MyProxy from all sites and from the Panda server (and from any other service like DDM) since the compromised pilot proxy could have fetched all ATLAS user proxy in MyProxy.
    - This can be alleviated with quite a bit of development at the ATLAS but not completely fixed.
    - A side note: security violations inevitably will happen. For example, gLexec protects the site from in case an ATLAS proxy gets stolen. But nothing protects ATLAS in case a pilot proxy (or a production proxy) gets stolen at a site (many ATLAS *FULL* proxies, delegated in CEs and WMSes).
- We would rely on a more complex system and more components which introduce more chances for security holes.

ATLAS operational experience with gLexec

- Significant work from many people has been employed to deploy gLexec and to use it. And a lot more work is needed.
- ATLAS experience with gLexec at T1s is that situation is "oscillatory": sites working at one point can stop working in a subsequent moment.
- gLexec introduces new failure modes in the analysis framework, with non negligible impact for users and support teams (ATLAS and WLCG).

The ATLAS conclusion is that:

- gLexec decreases the risk of a security violation, but increases the impact if this violation takes place
- gLexec introduces a non-negligible job inefficiency, will require still considerable development work and considerable operational effort from ATLAS and WLCG

So ATLAS would like to go directly to the next step:

- accept the risk that credentials can be compromised, whether gLexec is used or not
    - We consider the risk from ATLAS users to be minor, since ATLAS physicists are more interested in finding the Higgs than storing movies on scratchdisk in a T2 or disrupting their experiment's operations
- put effort in building a system which limits the impact.

In practical terms, this means:

- do not use gLexec
- accept that compromised DNs can be banned at any time by any site (and actually encourage it)
- integrate as soon as possible all ATLAS components with Argus, which should be the central place for credential banning
- encourage the different parts of the middleware to integrate with Argus as well

In this way, the ATLAS machinery will be able to:

- avoid a compromised pilot proxy from pulling jobs from Panda queues
- cancel all jobs submitted with compromised credentials
- avoid that a compromised user proxy can submit further payload into Panda
- avoid that a compromised user proxy can request data via ATLAS DDM

Security: what are the real needs for file access protection? There was a lot of "requirement" for ACLs on files, but as soon as these are used on a real system the overhead (particularly for reading many files) becomes significant. What is really needed? What is the simplest way to implement what is needed?

ATLAS does not require ACLs at the level of the single file. However, it requires ACLs at the level of space token or, equivalently, directory in the storage name space (in ATLAS one space token is associated with one directory tree). In particular:

- TAPE recall should be allowed only to /atlas/Role=production
- Writing into production spaces should be allowed only to /atlas/Role=production
- Writing into scratch spaces should be allowed only to /atlas users
- Reading from any space should be allowed only to /atlas users

ATLAS is considering if the last bullet (reading ATLAS files) should be relaxed, but this needs further discussion (at higher level than Computing Coordination)

Security: Romain has asked questions of you as input to the identity federation workshop. Is the use of X509 currently an issue for the experiments?

Not really an issue, but there are inconveniences. For example, at CERN the same user has a X509 identity and a AFS/Kerberos identity which complicates the analysis and data management workflow when Grid access has to be chained with local access. Having the possibility to map several user DNs to a user WLCG nickname and providing a X509toNickname (with different implementations corresponding to the local site security, for example a X509toKerberos for CERN) mapping service would be a big benefit.

Job management: pilot jobs are (almost) ubiquitous now. What is left that still needs a WMS?

The gLite WMS is not used for Production and is not a supported backend for analysis in ATLAS. Job submission to the CE happens only within the pilot factories and they all use Condor (the single WMS based factory has been decommissioned). Condor submission to the CE must work efficiently (but this is another thread). ATLAS still uses the WMS for the ATLAS Software Installation at sites. The SW installation framework is being moved to use Panda instead of direct submission. The timescale of this is a few months, including testing and a transition phase. From the ATLAS point of view:

- the WMS software can be frozen (development and certification)
- some instances of the WMS are still needed for the next 6 months: 2 instances at CERN and one in another site (CNAF) for fault tolerance.
- All ATLAS WMSes can be retired Jan 1st 2012

Job Management: can we simplify the needs for sites – I.e. reduce the complexity of a CE? Do we still need to require the CE to pass parameters to the batch system? (It has still not arrived….although promised). Don't your pilot frameworks do all this anyway?

Having a CE passing attributes to the Batch System is not an ATLAS requirement anymore since all the workload has been moved to pilot systems. The only use case one might foresee is passing a "whole node" request parameter. This today is managed by setting up dedicated queues. ATLAS is happy with this solution, as long as sites are content with it.

What is the intent with pilot factories? Will they be deployed at sites? If so, surely that replaces the CE completely (apart from a trivial mechanism to launch the factory at a site).

The ATLAS pilot factories are being modified to allow also direct submission to the batch system and being run by the site. ATLAS does not expect all sites to run a factory, but ones with enough manpower and with a strong ATLAS contact are willing to do it. A WLCG service to run pilots at the site (as a CE replacement) would be highly beneficial. ATLAS is very interested in discussing the specs with WLCG and middleware providers