Security

LCG GDB

CERN, 6 June 2007

David Kelsey

STFC/RAL

d.p.kelsey@rl.ac.uk

Overview

- JSPG meeting (phone conf) on 24 May 07
- Topics discussed included
  - Grid Site Operations Policy
  - Grid Policy on Handling Logged Personal Information
- Today, I will cover
  - Update on policies
  - Trust between VOs and Sites
  - glexec – barriers to implementation

Grid Site Operations Policy

- **https://edms.cern.ch/document/726129**
  - **Draft V1.4, 19 Apr 2007**
- “Final” version
- Was approved by GDB at May 07 meeting
- Then seeking EGEE, OSG, NDGF etc approval
- OSG Executive Board wanted to split one item into two
  - JSPG discussed
  - Decided wasn’t just a simple split
- Ask OSG if willing to leave as V1.4 with promise to address in future update
  - Delayed asking EGEE PEB until OSG situation is clear

Grid Security Policy

- New top-level document
  - To replace very out of date LCG-specific version
- See https://edms.cern.ch/document/428008/4
- V5.6 (24 Apr 2007)
- Comments have been received
  - From OSG
  - From a few EGEE sites
- **Is everyone else happy?** More readers are welcome!
- JSPG will produce new version (26/27 June meeting)
- Aim for approval at July 07 GDB

“Logged Information” Policy

- Needed urgently for User-Level accounting
- “Grid Policy on the Handling of Logged Personal Information”
  - https://edms.cern.ch/document/840299
  - V0.3 – early draft – presented at May GDB
  - A general policy covering all types of logged data
    - With an appendix for each type of data

Logged information (2)

- JSPG discussed the scope of the policy on 24 May
  - OSG: Need to avoid restrictions on exposing operational data – essential for debugging
    - Peer review of resource usage is also important
  - Agreed that there is a difference between live information versus logged information
    - Dashboard versus Accounting data
- May have to match VO privacy requirements to Sites capabilities (but this is complex)

Logged information (3)

- Given that User-level Accounting is urgent
- JSPG agreed that
  - The general policy will take too long
  - So work on this in parallel with
  - User-level Accounting Policy
    - Promote the previous appendix to a full document
    - JSPG will discuss at June meeting
- Need to decide what to do with pilot jobs accounting
- We know EU sites are concerned about privacy issues related to user-level accounting
  - But… is exposure of the user DN OK in live monitoring such as the Dashboard?
    - As long as user has consented

VO/Site Trust

- A topic driven by OSG (within JSPG)
- Defining expectations/responsibilities on both sides
- Clearly relates to VO boxes, glexec on WN etc

- For this presentation…
- Start by looking at JSPG work on
  - VO Operations Policy
  - This will be finalised in June meeting
- Then consider glexec
- And come back to VO/Site Trust

VO Operations Policy

- Draft document exists (from OSG)
  - not yet in EDMS
- Similar form to the Site Operations Policy
  - VO Managers required to accept (and sign?)

Policy points are

1. Provide and maintain contact information. Respond to enquiries in timely way.
2. Comply with Security Policies. Do self assessment.
3. Maintain (or get someone else to maintain) a VO membership service. Keep user data up to date. Recognise that this is a critical service.

VO Operations Policy (2)

4. Provide a support channel for VO supplied software and services. Respond promptly, especially for security. Sites can disable VO if risk too high. VO must address license issues.
5. Statement on logged information. Due diligence and limited purposes.
6. Use of Grid is at own risk. No liability.
7. Can control access by users. Must comply with security incident handling policy.
8. Comply with Grid Operations procedures
9. Grid may block VO access
10. Dispute handling

glexec – reminder

- *a thin layer to change unix credentials based on grid identity and attribute information*

can be thought of as:

- **‘a replacement for the gatekeeper’**
- **‘a *grid* version of Apache’s suexec(8)’**
- **‘a program wrapper around LCAS/LCMAPS or SAZ/GUMS’**

Glexec – what it does

- **Input**

1. a certificate chain, possibly with VOMS extensions

2. a user program name & arguments to run

- **Action**

1. check authorization (LCAS, SAZ)

- user credentials, proper VOMS attributes, executable name

2. acquire local credentials (LCMAPS, GUMS)

- local (uid, gid) pair, possibly across a cluster

3. enforce the local credential on the process

Result

1. user program is run with the mapped credentials

Glexec – why?

variety in grid job submission systems is increasing

– need a common way of obtaining and enforcing site policy and credential mapping

– without the need to modify each and every system

- Used by gLite CE
  - – Change to user ID before submitting job
- For pilot jobs on WN's – use glexec
  - – To do proper authentication/authorisation
  - – To allow proper traceability
  - – User-level accounting

JSPG policy on Pilot jobs

- JSPG agreed this text last summer
- *we REQUIRE suitable auditing and traceability at the individual user level both on the WN and the VO Scheduler available on demand.*
- *Sites may hold the submitter of the Pilot Job responsible for all actions of that job.*
- *VO's should be aware that the controls to ban users will result in the blocking of the whole VO, instead of just one user.*
- JSPG has not yet discussed glexec on WN
  - But perhaps we do need a policy

glexec on WN - Barriers to implementation

Conflicting views

- (Some) sites wish for
  - proper authentication, authorisation of user submitting
  - Coping with legal requirements for traceability
- (Some) sites have concerns about glexec
  - (I have only seen comments from UK sites)
  - Credentials should be delegated not transferred
  - Uncontrolled use of suid
    - difficult for shared resources
  - Is implementation secure?
    - Modifying security code is dangerous
  - Dislike of pilot job model

suid-less glexec

- **Site admins can choose not to set the set-uid bit.**
  - – glexec will adapt its functionality
    - logging works
    - user banlist works
    - certificate chain checks work
    - mapping is disabled, verification only
      - – *needs root privileges*
  - – real user job gets run with **pilot job identity**
  - – The discrimination will only be expressed in the log files

Back to VO/Site Trust

FNAL policy for pilot jobs (glexec on WN – CDF)

- Pilot Jobs will only be acceptable from VOs whose trust relationships with Fermilab include authorization to use them.
- Pilot Jobs must use the system utility provided by Fermilab (currently glexec) to map the application and data files to the actual owner of the User Job. This system utility will perform the necessary callouts to authorization, accounting and user management services.
- The Pilot Job must respect the result of these policy decisions.
- The Pilot Job and the User Job must not attempt to circumvent job accounting or limits on placed system resources by the batch system.

FNAL Pilot job policy (2)

- A Pilot Job may launch multiple User Jobs in serial fashion, but must not attempt to maintain data files between jobs belonging to different users.
- When transferring a User Job into the worker node, the Pilot Job must use means at least as secure as the original job submission process.
- Fermilab reserves the right to terminate any batch jobs that appear to be operating beyond their authorization, including Pilot Jobs and User Jobs not in compliance with this policy. Other possible consequences include blacklisting of users or the VO. Fermilab expects any VO authorized to run Pilot Jobs to assure compliance by its users.

The way forward?

- N.b. JSPG has not recently discussed this
- JSPG should decide whether to produce a general pilot job policy – based on FNAL text
- Agreements between Sites and VOs suffer from scaling problems
  - Even if just the 4 LHC VOs
- There is better scaling if the Grid acts as broker
- We will not achieve a single approach here
  - Some sites will refuse to run glexec with suid
  - If they run suid-less, is this OK?
- We need a proper security review of the glexec code

Requests to GDB

- Please comment on Grid Security Policy (V5.6)
- Seeking Site views on (data privacy issues)
  - – User identities exposed in Expt Dashboards
    - i.e. dynamic info cf logged information
  - – User ID exposed in general operational data
    - Essential for debugging
- Pilot jobs/glexec policy
  - – Can we impose a policy for WLCG sites?
    - E.g. all sites deploy but site can choose suid or not
- Discussion

JSPG Meetings, Web etc

- Meetings - Agenda, presentations, minutes etc

http://agenda.cern.ch/displayLevel.php?fid=68

- JSPG Web site

http://proj-lcg-security.web.cern.ch/

- Membership of the JSPG mail list is closed, BUT
  - Requests to join stating reasons to D Kelsey
  - Volunteers to work with us are always welcome!
- Policy documents at

http://cern.ch/proj-lcg-security/documents.html