Security: Traceablity & Liability

Ian Collier
RAL Tier 1
Pre-GDB 9th September 2014

Today's talk

- Following on from Vincent's talk in January & recent EGI CSIRT discussions
- Not discussing the need for traceability
- Not reviewing or discussing all issues
- Discussing next steps in achieving/maintaining traceability.

Motivation

To adequately respond to incidents we need to be able to answer these questions about any problematic activity:

- Who
- What
- When
- Where

Current Grid Security Traceability and Logging Policy

- Idea: understand and prevent incidents*
- Requirements:
  - Grid software MUST produce application logs:
    - Source of any action
    - Initiator of any action
  - Logs MUST be collected centrally
  - Logs MUST be kept 180 days
- Sites currently know what to do in order to be able to answer who, what, where & when

*https://edms.cern.ch/document/428037 & https://documents.egi.eu/document/81

New territory 1

- We cannot implement traceability in exactly the same way
- Sites can log observable behaviour
  - VM launched at such and such a time
  - Network connection to such and such an adress at a certain time
  - Etc.
- Sites can no longer see
  - Credential used to run workload(s) inside VMs
  - Detailed application logs from within past VMs
- CAN isolate running VMs for analysis

New territory 2

- VOs will have to participate in incident response to provide the missing information.
- Are VOs going to maintain detailed central application logs and retain them?
- Could sites provide a central syslog service for VMs run at their site?
  - But that would not help for public cloud work
  - Perhaps just for some nodes
- Many more issues and questions

Next steps

- We need to address this *before* workflows become too firmly established.
  - Easier to build in early than to add on afterwards
- Working group (sites and VOs) to
  - Test different approaches to filling traceability gaps
  - Update guidelines
  - Disseminate