WLCG Cloud Traceability Working Group

Managing Risk and the Emergence of Cloud Computing

Ian Collier
ian.collier@stfc.ac.uk
STFC RAL Tier 1, UK

HEPiX, Oxford, 24th March 2015

Emerging clouds

- Across our infrastructures the emergence of private, public and federated cloud resources changes many aspects of the way distributed computing works.

- Cloud resources and interfaces bring changes to workflows

  - Sometimes removing complexity for users (that is the aim)

  - Changing things for providers - some things are easier - some things perhaps not

- Clouds also introduce new software components and new workflows.

  - And new ways for things to go wrong

Without a trace?

- Management of risk is fundamental to the operation of any distributed computing infrastructure.
- Identifying the cause of incidents is essential to prevent them from re-occurring.
- In addition, we need to contain the impact of incidents while keeping services running.
- Our response to incidents also needs to be appropriate to the seriousness or scale of the issue.
- The **minimum** level of traceability for distributed computing infrastructures is to be able to identify:
  - the source of all actions (executables, file transfers, pilot jobs, portal jobs, etc)
  - and the individual who initiated them.

Currently

- We know how to do all those things in traditional grid based distributed computing infrastructures
  - Sites log in detail from the execution environment (worker nodes) and from CEs, batch systems etc to central loggers
  - Obtain granular authorisation & traceability in multi user pilot jobs with glexec (although this is not implemented universally)
  - Argus provides the fine grained authorisation required ***and*** allows us to centrally suspend compromised or suspicious credentials.
- Not just a technical problem - many years hard work mean:
  - We have developed *and agreed* incident response procedures
  - With clearly identify contact points
  - which help established trust relationships
  - And facilitate both the analysis of and response to problematic activities.

Change in landscape

- As conditions change you need different measures in order manage risk or feel comfortable.
- Sites no longer have the same control of or access to the execution environment
- VOs are developing & maintaining VM images
  - need to mange vulnerabilities previously managed by sites
- Site central logs no longer only data source required
  - VOs already log workflows for debugging & other purposes
  - How do we bring those logs into traceability IR process?

1958 Photo: M. T. Millett, NSIDC Glacier Photograph Collection

1985 Photo: M. M. Miller

2003 Photo: M. J. Beedle

2010 Photo: M. J. Beedle

WLCG Cloud Traceability Working Group

- WLCG set up cloud traceability working group to investigate these issues
  - Focus on practical work testing options for maintaining traceability
  - Many WLCG sites and all 4 LHC VOs represented
  - Has met face to face once
  - Work has begun

Areas of interest

- At F2F meeting we began by identifying areas of interest
  - Hypervisor & netflow logging
  - Logging from VMs to site syslog
  - Quarantining VMs
  - Increase of VO role in maintaining traceability
  - Giving sites root access to VMs for analysis
  - Policy evolution

Logging Issues

Logging Issues

- Increase logging of externally observable behaviour
  - Hypervisor & Cloud management framework
  - Network activity & flows (neglected until now)
- Within WLCG VM images are somewhat well controlled by VOs and there is a degree of trust
  - User and ‘supervisor’ roles are well separated.
  - It should be relatively easy to connect VMs to central loggers at sites but need standardised hooks in VM images
- Aggregation of and cross checking between multiple sources is vital
- Improved tools for storing, aggregating & searching increasingly important
- Potential for changes to VO workflow logging in order to better support traceability

Netflow & Hypervisor Logging

- Network flow logs until now not available to site admins at many sites.
- Survey experience of sites that do have access to hardware level network flow monitoring
- There may be different requirements for acceptable retention policies (more identifying information)
- Investigate approaches to network flow monitoring on hypervisors
  - Some possible approaches discussed
  - Need to test especially for any performance impact
  - Would not require access to network hardware (problematic at some sites)
- Formalise recommendations for logging from hypervisors (and cloud management frameworks) of what user or VO instantiate what VMs on which hypervisors etc.

Syslog

- Provide remote syslog service for running VMs
  - This does not happen at all at most public cloud providers, but should be straightforward to implement
  - Should be aware that in some incidents these logs may be unreliable
  - Will need improved frameworks for managing & searching high volumes of logs (eg ELK, much referenced elsewhere)
- Test creating VM images with hooks for sending syslog to loggers provided by the site
  - Compare using machine/job features and site contextualisation via, for example, cloud-init

Quarantining VM images

One new huge advantage of virtualisation is that we can more readily capture VM images for forensic examination.

- Easy for a running VM - but what if an attacker deliberately uses short lived VMs?
- Want VM images retained after VM shutdown
- Deferred deletion, for a tunable period, would be ideal
- Has a cost in storage occupied
- Some cloud platforms already do this
- Will investigate implementing this in others - specifically OpenNebula & OpenStack using Ceph

QUARANTINE
ZOMBIE OUTBREAK
RESTRICTED AREA
AUTHORIZED PERSONNEL ONLY
This area is QUARANTINED as a Class 3 Zombie Infestation Site.
No one shall enter or leave this area without written permission of the local health authority.

VOs as IR partners

- Already for some grid jobs we’d need to go to VOs to find out what user ran some jobs. (So that we can suspend just that user.)
- We know in WLCG that VOs already log workflows extensively to support debugging & workload management. We don’t know yet if more detailed logs are needed to provide full traceability.
- Rather than attempt an up front gap analysis, working group decided to approach this by running traceability service challenges and using these to identify any gaps
    - payloads and detailed challenge methodologies are currently being developed
- To some extent this is an opportunity to formally recognise the existing reality that we need the active participation of VOs in order to maintain traceability.

Policy & best practise

- We can use the results from all these areas of practical investigation to develop:
  - Updated policies setting out requirements for running these new forms of distributed computing infrastructures not only without compromising traceability but even improving it.
  - Best practise recommendations for how to gather additional logging information and how to configure management frameworks and VM images.
- While this work is focussed in the already well developed WLCG collaboration the policy and best practise we produce can provide a model for the many emerging cloud & virtualisation based distributed computing infrastructures.

I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices
I will adopt Best Practices

Area of debate

- Somewhat ‘uncertain’ if this approach will work
- For at least some sites/resource providers it remains essential to be able to directly suspend any possibly compromised credential
  - this would mean turning off entire VO
- Is it realistic for VOs to respond as quickly to
  - A. Identify problematic user
  - B. Suspend that user
- Alternative would be something like glexec within cloud frameworks.

Summary

- Practical work described here just looks at identifying & filling the traceability gaps
  - in one collaboration
- VOs will almost certainly become more formally involved in incident response processes
  - The well established - and large - LHC VOs are a good place to start
  - Can provide an example for other federated cloud infrastructure
- New policies will need to be agreed
- We will need to test that the implementation of those policies works through service challenges
  - and work to improve the areas where it does not

Questions

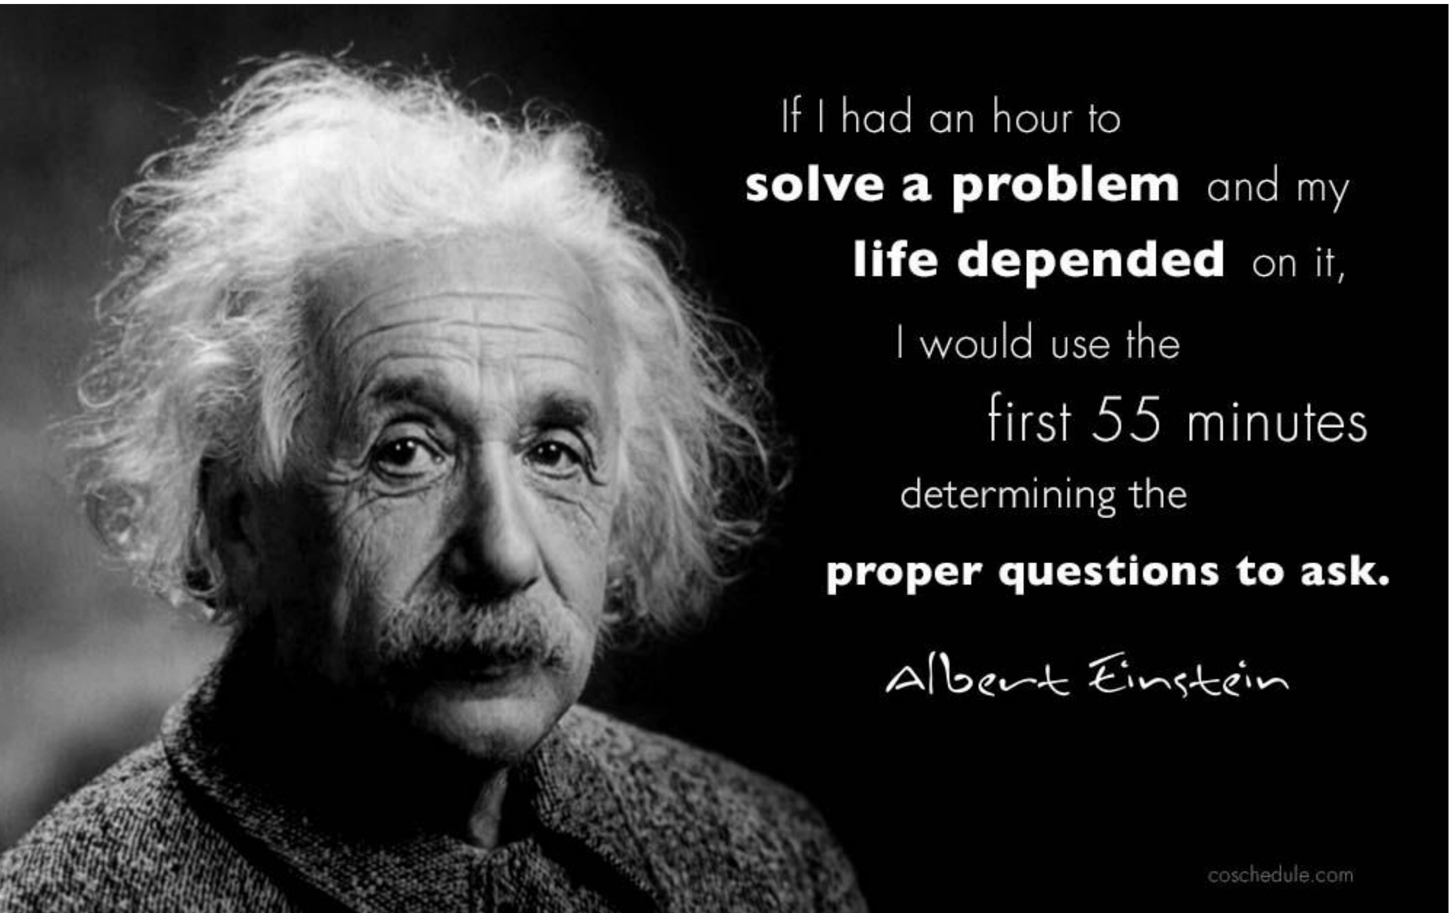

If I had an hour to **solve a problem** and my **life depended** on it, I would use the first 55 minutes determining the **proper questions to ask.**

Albert Einstein

coschedule.com