Cloud Security
Issues in Distributed Infrastructures

Ian Collier
UK Tier 1 Centre, RAL

Journées Sécurité France-Grilles
10th September 2014

- Thanks and apologies to Sven Gabriel & Vincent Brillaut who I borrow from heavily
- Not discussing the need for traceability
- Not reviewing or discussing all issues
- Evolving area

Motivation

When things go wrong we must be able to answer these questions about any problematic activity:

- **Who**
- did **What**
- **When** did they do it,
- and **Where**

Current Situation

- Grid sites in our collaborations currently know what to do in order to be able to answer ***who***, ***what***, ***where*** & ***when***
  - Archiving central logs
  - How to link remote users to specific processes on local systems
- And in addition we know
  - How to suspend credentials when we suspect problems
  - Incorporating new components
  - All the things Leif discussed help which help us manage systems we (sites) control
  - etc., etc.

What does cloud computing change?

Conceptually not so much

WLCG/EGI Risk assessment:

- Mostly applies to cloud (some new threats)
- Most important identified asset: Trust
- Most dangerous threat: Misused identities
- Focuses on traceability for:
  - Incident containment
  - Preventing reoccurence/repetition of incidents

- As Romain said **traceability** is the key thing we must maintain.

Clouds do move some system configuration & security responsibility to VOs

Virtual Machine Image Endorsement

EGI Security Policy for the endorsement and operation of Virtual Machine images*

- 2 roles:
  - Endorser: Certify VM Image
  - VM Operator: Root access on the VM
- Security requirements for both roles
- Users are not endorsers:

  An Endorser should be one of a limited number of authorised and trusted individuals appointed either by the Infrastructure Organisation, a VO or a resource centre

*https://documents.egi.eu/public/ShowDocument?docid=771

Virtual Machine Image Endorsement

Scenarios (not exhaustive)

- endorser/operator = site: current situation (mostly)
- endorser = VO: could provide more flexibility
- operator = VO: could provide technical debugging
- endorser/operator = end user: not foreseen useful

Traceability

EGI 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

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

Traceability

Virtualization mostly introduces new possibilities:

- Logging requirements not changed/impacted:
  - Every action/every user
  - Forwarded to a central server
- New logs are almost certainly required (policy extension?):
  - Which endorsed VM is running?
  - Who is operating it (Site/VO) ?
- User compartmentalization:
  - Similar to glexec? (one UID per user)
    - May be easier to implement than glexec
  - Re-instantiate VM for each user (not job)
    - Perfect easy compartmentalization
    - High impact for unique short jobs

Traceability

Complete root access for user is dangerous:

- Endorsed VM:
  - Contains up-to-date software (by policy)
  - Contains secured configuration (by policy)
  - Can include protections/logging...
- User in full-control can:
  - break configuration (maliciously or by error)
  - disable logging (maliciously or by error)
  - falsify data (non-trusted logs)
- Simple accountability/traceability: user responsible
- Difficult detailed incident analysis
- VM cannot be re-used by different users

No identified reason for such situation: should be highly discouraged

Traceability

Complete user control -> no security

- Unendorsed VM :
  - Could be vulnerable (not patched, outdated...)
  - Could be badly configured (no logs, anonymous access...)
  - Could be fully-encrypted (no forensics possible)
- User in full-power:
  - Can falsify data (non-trusted logs)
- Simple accountability/traceability: user responsible
- Potentially impossible incident analysis
- VM cannot be re-used by different users

No identified reason for such situation: should be highly discouraged

Traceability

- VM creation/deletion is easy (could be VO/user initialized)
- VM lifetime foreseen shorter than current WN
- If trusted operator/endorser:
  - Application logs centrally kept
  - More system logs probably needed
  - Unknown/modified file preservation would help forensics
- If non-trusted operator/endorser:
  - Application logs (central) are not trustworthy
  - System logs (central) are not trustworthy
  - VM disk MUST be preserved after deletion

Policy extension required?

Limiting exposure to vulnerabilities

Three scenarios (may be more):

- Probe every VM for vulnerabilities:
  - Much more work than now (who would do it?)
    - Extremely diverse security contacts
- Limit VM lifetime:
  - Vulnerability window restricted (automatic)
  - How long (soft/hard limits ?) ?
    - Hours ?
    - 2-3days?
    - Week(s) ?
    - Month(s) ?
- Even with Trusted endorser/operator:
  - Identify vulnerable VM in trusted VM store
    - Contact *all* VM operators (who are they, who would do this?)
    - Kill switch to be implemented (who will implement it?)

Incident Response

- Need well defined security contacts
- Require root access on VM for:
  - Site admin ?
  - EGI/OSG security team, WLCG security officer ?
- VM freezing/isolation (could break jobs):
  - Who is authorized to do it?
  - What is the procedure (under which circumstances ?) ?
- Forensic analysis of suspect VM images:
  - Who is authorized to do it?
  - Procedure (under which circumstances ?) ?
  - Private data protection ?

Incident Response

- How to ban a user:
  - From site/VO operated VM ?
  - From cloud system (user-operated VM) ?
- How to ban a cloud provider (site) ?
- How to ban a problematic VM image (from the VM store):
  - For newly created VMs ?
  - Killing running VMs ?

Incident Response

- Some documents may need to be revisited/ extended:
  - Risk assessment (new threats)
  - Traceability requirement (new layer, VM deletion)
  - Incident procedures
- Will probably need an extended Acceptable Use Policy (AUP) which all operators and final users need to abide by a potentially :
  - Recognizing liability
  - Allowing security teams to intervene

More Cloud Specific Issues

Hypervisor containment might be broken:

- Require separated hypervisor clusters for:
  - Infrastructure ?
  - Worker Nodes (Site/VO operated) ?
  - Untrusted VM (End User operated) ?
- Require physical host for critical infrastucture
- Hypervisor traceability needed:
  - VM traceability (On which hyperviser each VM is)
  - System & audit central logs

Cloud Issues

- Incident response procedure?
- Abuse detection (IDS not available in same way) ?
- Security incident costs, e.g. Amazon agreement*:

> *If we or our affiliates are obligated to respond to a third party subpoena or other compulsory legal order or process described above, you will also reimburse us for reasonable attorneys' fees, as well as our employees' and contractors' time and materials spent responding to the third party subpoena or other compulsory legal order or process at our then-current hourly rates.*

*https://aws.amazon.com/agreement/

Next steps

- You'll notice this was mostly questions
- We should answer them *before* workflows become too firmly established.
  - Easier to build in early than to add on afterwards