Attacking disjoint federations the old way

Vincent BRILLAULT

What worked before still works…
What about the future?

EGI-Engage is co-funded by the Horizon 2020 Framework Programme
of the European Union under grant number 654142

Disclaimer

- Some of the incidents presented occurred recently
  - As this forum is public, most details cannot be shared
  - In particular, no name or date, to protect the victims
- Presenting from my own/EGI's point of view
  - Not directly affected, angle of perception different
  - Not all details were shared by everyone
    Only the attackers know everything...
- No EGI resources affected by the most visible attack

Recent incidents

2-3 recent incidents affecting multiple places

- Exact number sometimes difficult to tell:
  - Never concluding by arrests: are they coming back?
  - Two incidents affecting the same servers

  $\rightarrow$ Some initial reports can be misleading

- Affecting everything (as far as we know)
  - Academic places, HPCs, private sites, etc
  - Large range of operating systems & versions

Successful attack methods

- One successful methodology: steal credentials
  - Remote vulnerabilities rare, monitored & patched
  - Users/Admins credentials & logins less monitored
- Different approach (most seen this year)
  - Replace SSH/SSHD with malicious variant
    - Evolution: replace SSH/SSHD library
  - Malicious PAM modules (passwords only)
  - Collect local private keys (SSH keys only)

Spreading around

Aren't these different systems/disjoint federations?

- Single user needed to jump over between places
  - Same credential, targets in shell history/known hosts
  - Malicious logging of credential & target
- Privilege escalation for next steps?
  - Non privileged access can still expose more systems
  - Some attacks without any trace of exploit:
    - Connecting directly as root
    - Connection as user with sudo privileges
  - Exploitation of (old?) vulnerabilities

Exploiting vulnerabilities

- Hard to prove what vulnerability was used
  - In most cases, no trace left (exploits cleaned up)
  - Even if suspicious file found, was it the one used?
- One case clear: undisclosed vulnerability in GPFS
- Doubts about several other cases:
  - *Exploit* files found on some systems
  - Seems tailored to specific kernel (offsets)
  - Included strings, reverse engineering identified CVE:
    - Different known kernel memory issues
    - Theoretical privileged escalations, never proven

Traces left by the attack

How to detect it? What can you see afterwards?

- Connections from unknown/suspicious location
  - Only if attacker do not proxy through infected nodes
  - Only if logs stored remotely: often cleaned up
- Filesystem metadata (atime/mtime/ctime)
  - Useful in most cases but sometime fully cleaned up
  - Seen in these cases:
    - Same atime most of ~/.ssh/* for all users
    - Suspicious ctime on ssh, sshd, configuration

Going forward: Cloud & Federations

- Clouds: no fundamental changes
  - In most cases just another layer, nothing else
  - Sometime just more credentials to be lost...
- Federations: federated access, AAI
  - Possible game changer!
    - Main attack vector: unsupervised reused credentials
  - Delegated credential: lateral movement difficult
  - Main credential compromised: all service exposed...
  - Central management: block, revocation, renewal easy

Going forward: Protecting ourself for now

- Monitor remote accesses
  - People usually connect from few places
  - Collect all connection logs to a central place
  - Notify them of any new, different activity
- Patch *non-critical* kernel vulnerabilities
  - Someone might be able to use *theoretical* exploits
  - Assessments usually not updated to critical...
- Share anonymous compromised credentials?
  - Sharing fingerprint of compromised SSH keys

Any question?