WLCG Group

Security TEG

WLCG Collaboration Workshop, NYC, 20th May 2012

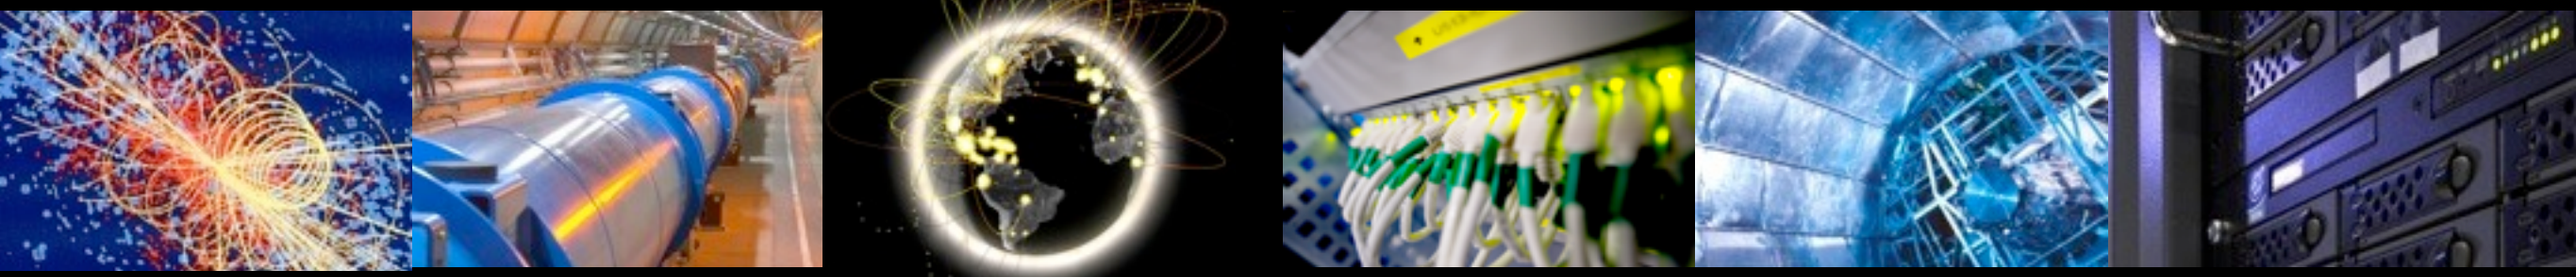

WLCG
Worldwide LHC Computing Grid

Summary

- Incidents happen on a regular basis, 10-12 per year
- Attacks continue to improve
  - More and more sophisticated
    - For example, Zeus Windows botnet used to steal HEP accounts
    - No easy or public mean to detect modern malware
  - No longer a side-effect of being connected to the Internet
    - State-of-the-art malware used against WLCG
    - Attackers being arrested for attacking WLCG resources
  - No reduction of the severity or # of incidents in the recent years
    - Yet most of them follow the same pattern
    - Needs to improve our tools and our practices
    - We have now built the necessary expertise and have experience
- The risk assessment should define the main directions
  - And be used as a reference to evaluate the efficiency of our tools
  - "Usability of security" TEG report another useful reference

Summary

- Important “technical areas” where work is needed
  - Fulfill traceability requirements on all services
    - Sufficient logging for middleware services
    - Improve the logging of WNs and UIs
    - Too many sites simply opt-out of incident response “no data, no investigation -> no work to be done!”
    - Prepare for future computing model (e.g. private clouds)
  - Enable appropriate security controls (AuthZ)
    - Need to incorporate identity federations
    - Enable convenient central banning
- Important “people” issues
  - Must improve our security patching and practices at the sites
  - Collaborate with external communities for incident response and policies
  - Building trust has proven extremely fruitful - needs to continue

Risk analysis

- highlighted the need for fine-grained traceability
  - Essential to contain, investigate incidents, prevents re-occurrence
- Aggravating factor for every risk:
  - Publicity and press impact arising from security incidents
- 11 separate risks identified and scored. Top risks:

Risk
Misused identities (“SSH”-type included)
Attack propagation between WLCG sites
Exploitation of a serious OS vulnerability
Threats originating from trust services
Negative publicity on a non-event
Insecure configuration leading to undesirable access
Insufficient protection of information leading to sensitive data leakage
Incidents on resources not bound by WLCG policies
Exploitation of a serious VO/middleware software vulnerability
Data removal/corruption/alteration
DoS from an external organisation

AAI on the worker nodes

- Several areas would need further discussion or are yet to be addressed, including:
  - The use and transport of credentials on the WN, including delegation, propagation, revocation and traceability.
  - The implementation of security controls (e.g. blocking/banning end users, credential revocation) and who should operate them.
  - The ownership of the traceability information. Is it OK to split the traceability information between VOs and sites?
  - The security implications of virtualization on the WN
  - The security implications of submitting jobs to external clouds
  - The network connectivity requirements of the experiments
  - The longer term future of the security model of the WN

- Discussion day dedicated to the WN on 12th June at CERN (pre-GDB day)