Security Service Challenge 3

Romain Wartel, CERN IT

EGEE Operational Security Coordination Team

http://www.eu-egee.org/security/

GDB Meeting, 14/05/2008, CERN

Information Society and Media

Credits

MANY thanks to

- **Pål Anderssen**
- 

Security Service Challenge

- **Simulate (parts of) real incidents**
- **Objectives**
  - Improve the response from the sites and from the security contacts (training)
  - Ensure sufficient expertise is available at the sites and the ROCs
  - Ensure reporting channels are working
  - Check if communications between involved parties is working
  - Check if sufficient audit trails (logs) are available to investigate the activity
  - Check compliance with LCG/EGEE incident response policies/procedures
  - More details at http:/cern.ch/osct/ssc.html

- **Security Service Challenges (SSC) history**
  - SSC1: Trace a job (WN -> CE -> RB -> UI)
  - SSC2: Trace storage operations (creation, move, deletion, etc.)
  - SSC3: Treat this DN as malicious

- **SSC3 was run at most WLCG Tier1s.**

Security Service Challenge 3

- **SSC3 is a very real simulation of a security incident**
- **Same procedure followed for all challenged sites:**
  1. Submit a carefully crafted binary to the site
  2. Contact the site security team (GOCDB)
     "You are asked to follow the normal incident procedure, but you MUST_NOT take any collective action against the VO of the offending user.
     Consider any activity from the following user as malicious."
  3. Carefully monitor and assess the response of the site
- **Feedback from previous challenges: sites want to know how they performed**
  - Difficult to implement a detailed evaluation in a best effort environment

Security Service Challenge 3

Evaluation was based on three main areas:

- **Communication**
  - Acknowledge/Heads-up report sent to the CSIRT list
  - Alert sent to the affected VO Manager
  - Verify the responsible CA has been notified
  - Close-out report sent to the CSIRT list
- **Containment**
  - Found the malicious job and killed it
  - Suspended the user at the site
- **Forensics**
  - Discovery of initiating site (UI) and established contact with that site's CERT
  - Found evidence of malicious network traffic
  - Some analysis of the submitted binaries was performed
- **Plus some bonuses points for very fast response time, etc.**
- **Caveat 1: "***The scores are based on the quality and delay of the response of the sites. The scale of the various items used in the scoring and the procedure to determine what constitutes a satisfactory response are based on various parameters, including LCG/EGEE security policies, procedures, and experience from past incidents.*
*The best scores are not necessarily requirements but are generally perceived as good, appropriate practices in a best effort environment.***"**
- **Caveat 2: Debriefing in progress: sites may object on specific points**

CERN PROD

- **Total score: 31.5 / 27 PASS**
  - Communication: 14 / 12 PASS
  - Containment: 5 / 6 FAIL
  - Forensics: 8.55 / 9 FAIL
  - Bonus: + 8
- **Additional comments**
  - Some malicious processes have not been killed
  - Slightly long delay to obtain the name of the originating UI (at CERN too)
  - Note: LCG/EGEE security experts based at CERN did not intervene to help

RAL LCG2

- **Total score: 32.5 / 27 PASS**
  - Communication: 14 / 12 PASS
  - Containment: 5 / 6 FAIL
  - Forensics: 9.5 / 9 PASS
  - Bonus: + 10 (max)
- **Additional comments**
  - Very fast response
  - Some malicious processes have not been killed
  - Further jobs could be submitted via the CERN RB (user was suspended only at the local RB)
  - The analysis did not reveal the network traffic generated by the payload (problem now understood)
  - http://www.gridpp.rl.ac.uk/blog/2008/04/24/ssc_3-pipped-at-the-post/

- **Total score: 22.2 / 27 FAIL**
  - Communication: 7 / 12 FAIL
  - Containment: 6 / 6 PASS
  - Forensics: 7.25 / 9 FAIL
  - Bonus: + 5
- **Additional comments**
  - Malicious user was successfully banned and all processes terminated
  - Incomplete communication and forensics report

- **Total score: 21 / 27 FAIL**
  - Communication: 8 / 12 FAIL
  - Containment: 8 / 6 PASS
  - Forensics: 5 / 9 FAIL
  - Bonus: + 6
- **Additional comments**
  - Malicious user was successfully banned and all processes terminated
  - Incomplete communication (ex: the VO was not notified)
  - The origin of the job was not reported, some forensics information missing

FZK LCG2

- **Total score: 29 / 27 PASS**
  - Communication: 14 / 12 PASS
  - Containment: 8 / 6 PASS
  - Forensics: 5 / 9 FAIL
  - Bonus: + 8
- **Additional comments**
  - Malicious user was successfully banned and all processes terminated
  - Evidence was lost on the WN when the job was killed

- **Total score: 4 / 27 FAIL**
  - Communication: 0.5 / 12 FAIL
  - Containment: 1 / 6 FAIL
  - Forensics: 0.5 / 9 FAIL
  - Bonus: + 2
- **Additional comments**
  - Little information received and little action taken
  - Not the T1 results, SSC3 was submitted to 3 INFN sites to try to get better scores:
    - INFN-T1: 0 / 27
    - INFN Pisa: 0.66 /27
    - INFN Bologna: 4/27
  - Lots of efforts from the ROC Security Contact over the recent years/months to improve the situation.
    Security re-organised recently following several security incidents:

    http://indico.cern.ch/materialDisplay.py?sessionId=2&materialId=1&confId=29322
  - Suggestions?

NIKHEF PROD

- **Total score: 30.8 / 27 PASS**
  - Communication: 14 / 12 PASS
  - Containment: 4.8 / 6 FAIL
  - Forensics: 8 / 9 FAIL
  - Bonus: + 9
- **Additional comments**
  - Some malicious processes have not been killed
  - Off-site network traffic from the malicious binary was not identified

TAIWAN LCG2

- **Total score: 7.6 / 27 FAIL**
  - Communication: 0.6 / 12 FAIL
  - Containment: 1 / 6 FAIL
  - Forensics: 5 / 9 FAIL
  - Bonus: + 4
- **Additional comments**
  - Some malicious processes have not been killed and the malicious DN was not successfully banned
  - The nature of the malicious binary was not reported
  - Better understanding of incident EGEE response procedure needed (ongoing)

Scores

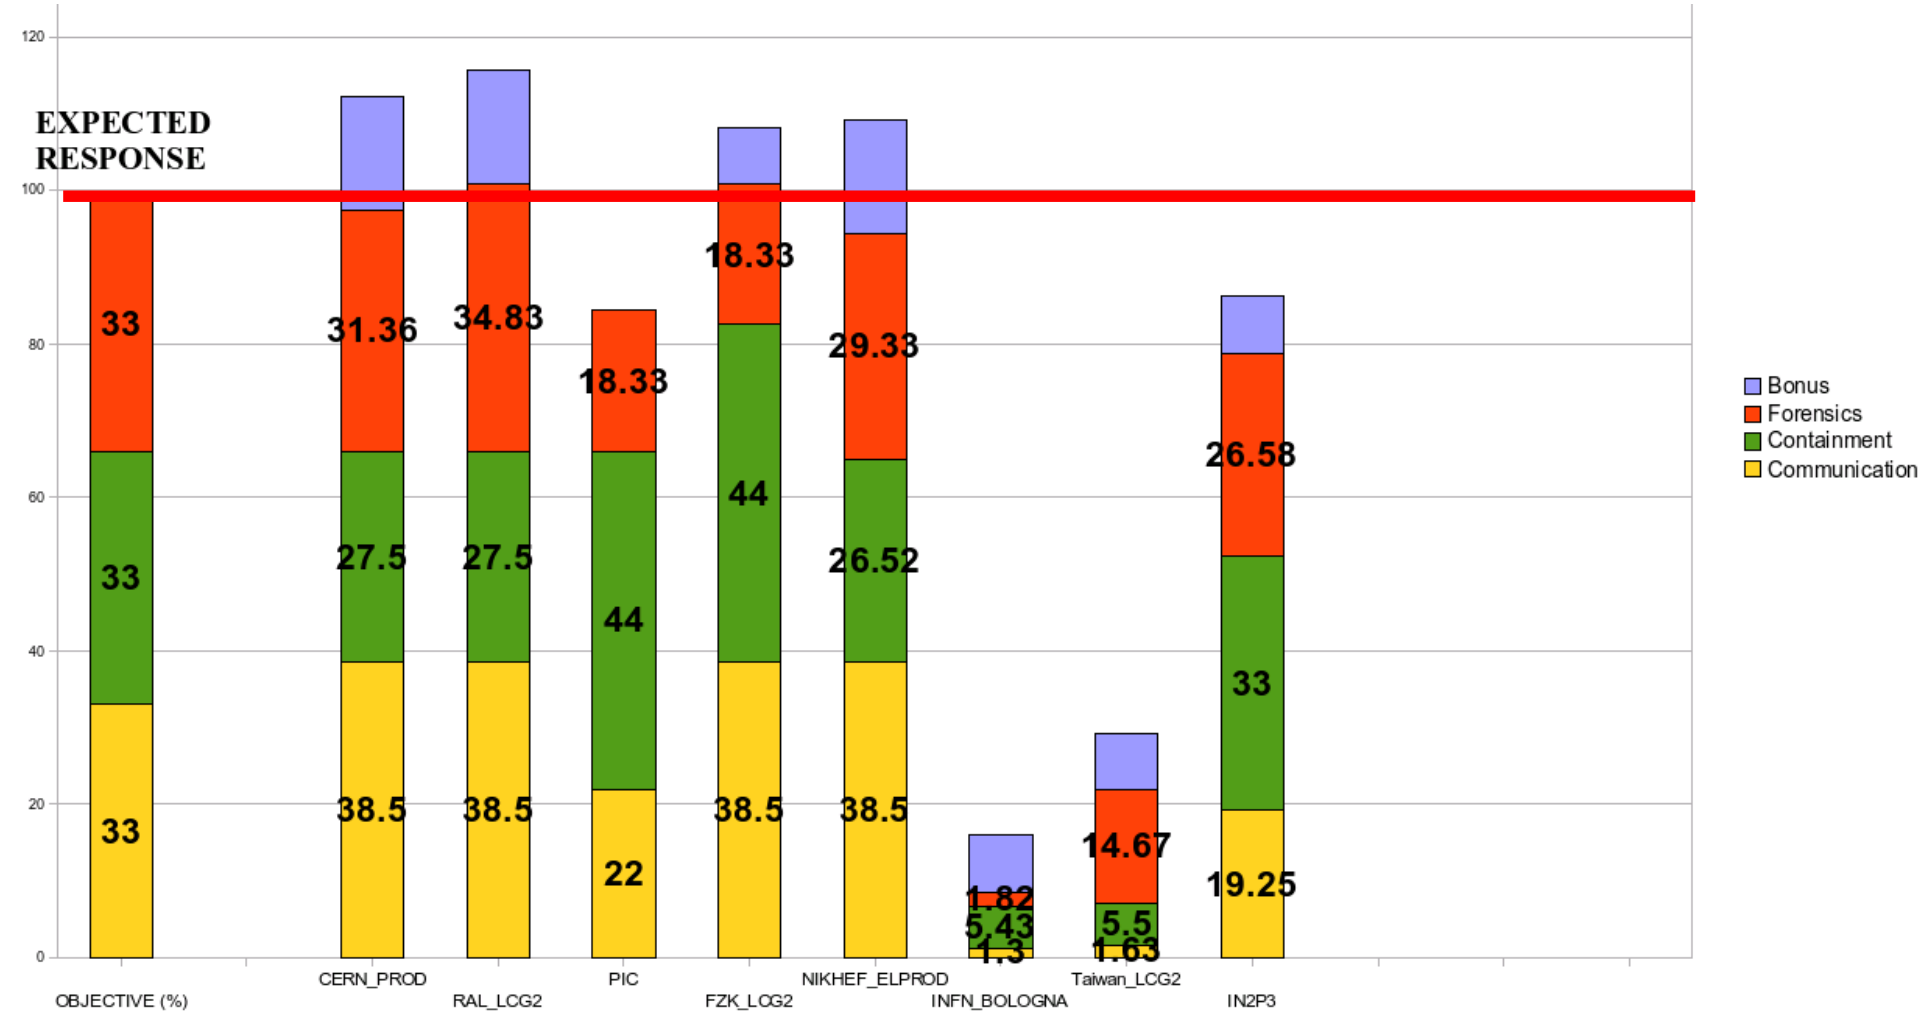
EXPECTED RESPONSE
120
100
80
60
40
20
0
33
33
33
31.36
27.5
38.5
34.83
27.5
38.5
18.33
44
22
18.33
44
38.5
29.33
26.52
38.5
1.82
5.43
1.3
14.67
5.5
1.63
26.58
33
19.25
Bonus
Forensics
Containment
Communication
OBJECTIVE (%)
CERN_PROD
RAL_LCG2
PIC
FZK_LCG2
NIKHEF_ELPROD
INFN_BOLOGNA
Taiwan_LCG2
IN2P3

Findings

- Heterogeneous results... but not so different from experience with real incidents
- Debriefing in progress, but overall feedback so far: very useful
- ~ 1/3 of the Tier1s unable to block the malicious DN
- ~ 2/3 of the Tier1s unable to kill the malicious processes
- Only 1 site killed all malicious processes *without* unplugging the WN
- Incident response procedure generally understood
- Need better mechanisms to suspend users
- Logging is often a problem: middleware logs, lack of central syslog
- Better security procedures are needed for the sites
- Need to reinforce security expertise at the sites

Security Incidents Statistics

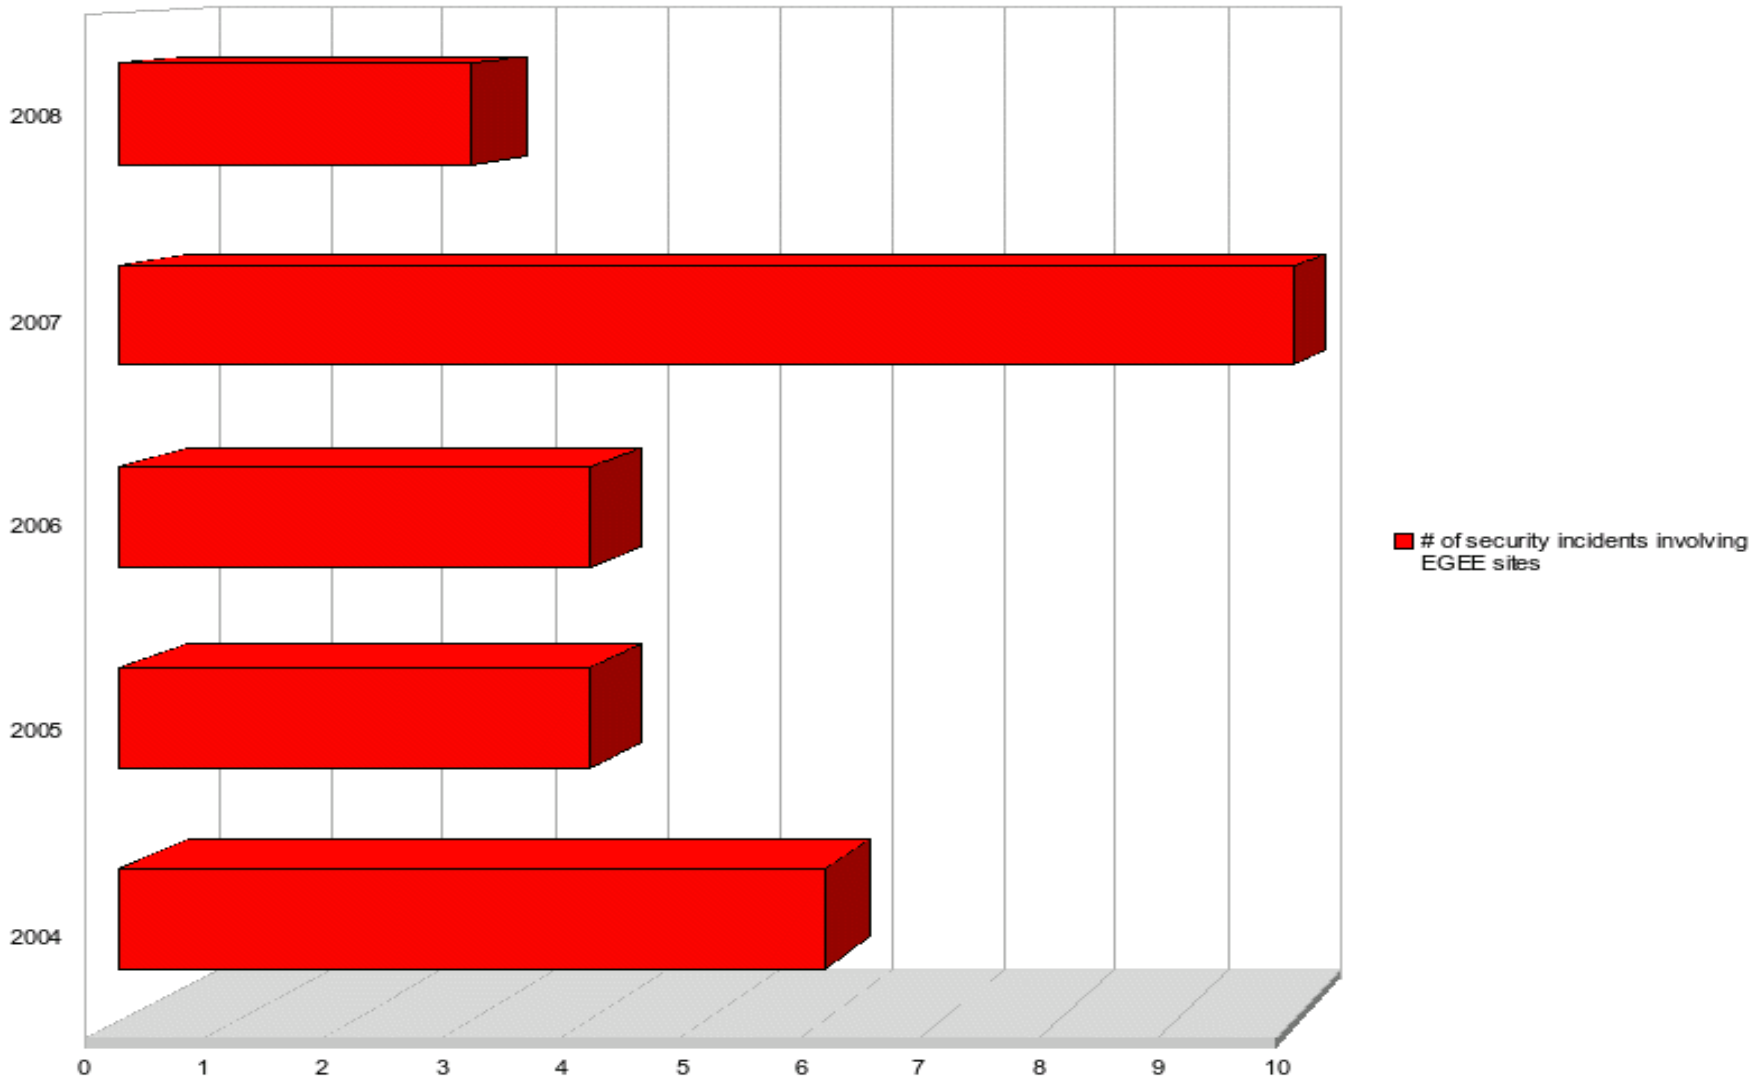

Future

- Challenge US Tier1s (in progress)
- Make the toolkit available to the ROCs
- Partnership with LHCb very fruitful: other VOs interested?
- Regional VOs needed to challenge Tier2s
- Adapt the OSCT training material to the SSC findings
- Long terms plan depends on efforts from the ROC (TSA 1.4.1)

Discussion