WISE Community: Security Communication Challenges Coordination WG (SCCC-WG)

Introduction and background

Maintaining trust between different infrastructures and domains depends largely on predictable responses by all parties involved. Many frameworks – e.g. SCI and Sirtfi – and groups such as the coordinated e-Infrastructures, the IGTF, and REFEDS, all promote mechanisms to publish security contact information, and have either explicit or implicit expectations on their remit, responsiveness, and level of confidentiality maintained. However, it is a well-recognised fact that data that is not verified becomes stale: security contact information that is appropriate at time of enrolment in an infrastructure may later bounce, or have different 'characteristics'.

One of the ways to ensure contact details are maintained is to 'exercise' these contacts regularly and compare their performance against the expectations or requirements, in what is usually called 'communications challenges'. However, with many distinct stakeholders interested in ensuring correctness of these contact details, it is likely that uncoordinated challenges have a detrimental effect on responsiveness: tests are duplicated, follow each other too closely in time, or measure the same aspect of contact responsiveness in different (and thus potentially confusing) ways. This is likely to 'overload' the targets of these challenges, resulting in disengagement and understandable ill will to participate in the future even in case of real incidents

It is also likely that the many different stakeholders value different 'aspects' of communications (timeliness, investigative capability, confidentiality, ability to take action), and have different expectations. Yet the way of measuring these aspects may be very similar for each of these stakeholders – it is just their expectations that are different. Running similar communications challenges by different stakeholders concurrently – where the observed characteristics are the same but its interpretation different – would again unduly waste resources at the communications end-point. Testing for 7-day responsiveness and a response within 4-business hours is the same kind of test, and two infrastructures, each interested in one of these options, should likely do a single combined challenge.

Yet there are also different types of challenges, and the collaborating infrastructures may each have expertise in different aspects of such challenges. Some groups run email-responsiveness challenges as a matter of course (e.g. TF-CSIRT, SURFcert), others specialise in either 'white-box' or 'black-box' challenges where also the operational capability and the ability to react to actionable intelligence is an integral part of the challenge (e.g. the EGI Security Service Challenges). These different aspects of communications challenges should be taken into account, and both provide significant value to all WISE infrastructures if the detailed results are shared, yet also these different types must still be coordinated in time, because the operational challenges implicitly also measure responsiveness (and thus provide valuable data on that point as well). Doing the latter shortly after the former would again be annoying for the targets of the challenge.

At the same time, it should be recognised that – in absence of a strong coordination and information sharing mechanism between trust groups and e-infrastructure stakeholders – testing by *one* party

does not improve or foster trust in the communications ability of those tested for *another* party. For this trust to be transient, those who are conducting a communications challenge must share also details with their infrastructure peers, not just a compliance statement. The reason for this is the necessarily different *interpretation* of the results. For example, where for the infrastructure conducting a communications challenge a 7-day response time may be adequate, some of its peers may expect (in an actual case) a response within four business hours. An unqualified 'compliance statement' many by the first infrastructure thus has limited value to the second infrastructure – who may want to act differently (e.g. implement its own mitigating measures) had it known that the response could take as long as 7 days.

Work items

The SCCC WG will address the following aspects of security communications challenge (CC) coordination:

- Coordination of 'CCs recipient groups' among participating infrastructures
  making sure that targets are not overload by coinciding or overlapping challenges, for example by designating a lead infrastructure for each category of targeted entities
- Transitivity of trust in CC results between infrastructures
  for example by specifying the level of disclosure detail for CCs between trusted infrastructures, by using an SCI evaluation framework approach to it, or by coordination of testing and success criteria.
  How can requests for CCs between infrastructures be handled, e.g. in response to changing needs or a changed risk assessments; or as remediation after an incident in which communications did not meet expectation.
- Definition of CC models and classification
  the 'depth' of the CC testing is a balance between the level of trust gained (more profound testing and good results gives more trust) and expediency (asking the recipient to respond to a mail or click a link consumes less resources than requesting forensic investigation of a simulated incident of deliberately unknown nature).
- Frequency of CCs
  simple communications challenges are often performed one or several times per year (e.g. for TF-CSIRT, by SURFcert for the SURFconext federation, EGI Operations on their sites). Complex challenges are less frequent (e.g., the 'black-box traceability' trials of the EGI Security Service Challenges tale place once every 1-2 years). Following a CC model classification, propose an appropriate frequency for each class.

The SCCC-WG should thereafter become a standing interest group in the WISE-community that maintains a timetable of planned CCs (to prevent overlap), provides a lightweight mechanism to request and coordinate CCs, and promotes the sharing of results with qualified peer infrastructures.