PROCEEDINGS OF SCIENCE

Coordinating Operational Security in evolving distributed IT-Infrastructures

Vincent Brillault[a], Linda Cornwall[b], Nuno Dias[c], Tobias Dussa[d], Sophie Ferry[e], Sven Gabriel*[f], David Groep[f], David Kelsey[b], Daniel Kouril[g], Barbara Krasovec[h], Ian Neilson[b], and Fyodor Yarochkin[i]

[a] *CERN, Geneva, Switzerland*

[b] *Science and Technology Facilities Council (STFC) – Rutherford Appleton Laboratory, Didcot, UK*

[c] *Laboratorio de Instrumentacao e Fisica de Particulas (LIP), Lisbon, Portugal*

[d] *Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany*

[e] *CEA Saclay, Gif-Sur-Yvette, France*

[f] *NIKHEF, Amsterdam, the Netherlands*

[g] *CESNET, Prague, Czech Republic*

[h] *Jožef Stefan Institute, Ljubljana, Slovenia*

[i] *Academia Sinica, Taipei, Taiwan*

E-mail: csirt@egi.eu

Operational Security in scientific distributed IT-Infrastructures such as EGI is challenging. Existing computation frameworks are continually being extended, and new technologies implemented, thereby expanding the potential attack surface and exposing new risks.

In this rapidly evolving environment new security policies have to be developed, and existing policies and procedures have to be constantly updated to meet new requirements.

To efficiently enforce these new policies, the security monitoring infrastructure has to be further developed to cover all elements of the evolving infrastructure. Finally the incident response (IR) tool set has to be extended to be able to efficiently handle security incidents affecting new technologies.

In this paper we discuss EGI-CSIRTs strategy for expanding its portfolio to provide all aspects of operational security in a Cloud environment, whilst maintaining its current capabilities. The paper describes the developments associated with a Virtual Machine Endorsement Policy and related technical aspects to allow a provision of a trustworthy set of Virtual Machine Images (VMI) to the user community by means of an Application-DataBase.

VMIs with vulnerable configurations have already involved in incidents handled by EGI-CSIRTs Incident Response Task Force (IRTF). In dealing with these incidents it became apparent that the existing procedures and tools, which were otherwise successfully applied to IR in EGI, exposed deficiencies when applied to the EGI Federated Cloud (EGI-FedCloud) services. This understanding triggered the development of central User- and Virtual Machine-Management frameworks deployed in EGI-FedCloud. The status of these tools and the integration with the existing IR tools is discussed.

In EGI, security policies and procedures are tested in Security Service Challenges (SSCs) which are designed to verify that they do, in practice, help with security operations to prevent and respond to incidents. An SSC addressing EGI-FedCloud services and IR procedures will is described.

International Symposium on Grids and Clouds 2017 -ISGC 2017-
5-10 March 2017
Academia Sinica, Taipei, Taiwan

[*] Speaker.

1. Introduction

EGI[1], European Grid Infrastructure, is a federated e-infrastructure set up to provide advanced computing services for research and science. It was founded in 2010 with the mission to create and deliver open source solutions for science and research infrastructures, to invite new disciplines to use the federated infrastructure and to manage development of software with the cooperation of the national grid infrastructures (NGI-s). Software and middleware development was unified and standardized to some extent, under the aegis of EMI middleware project, so in terms of Grid Middleware and Operations, the infrastructure was rather homogeneous, supporting grid and high-throughput computing (HTC) centres. EGI today consists of more than 400 sites organized in 40 National Grid Initiatives (NGIs).

In order to use the EGI infrastructure users are organized in *Virtual Organisations* (VOs). The VOs are in general supported by a particular subset of the Resource Centres (RC) forming EGI. Access control to the infrastructure is based on X.509 certificates, with which the users register with a VO. One can say that users are granted access to the infrastructure through their VOs. Most of the VOs operate their own Workload-Management-Systems (WMS) which produce logging information relevant for incident response. Therefore VOs are an independent service provider in EGI and an important partner for the EGIs Computer Security Incident Response Team (EGI-CSIRT).

In recent years, a lot of interest and development has been directed to the usage of private and public clouds for scientific workloads and to the integration of grid and cloud. Clouds offer seemingly infinite resources and they meet various end-user demands with the possibility of customizing the execution environments for users' workloads. RCs adopted various different approaches to cloud implementation. Some RCs run jobs within the virtual machines, some offer scientific private clouds within the EGI Federated Cloud, or they offer an interoperable, agile infrastructure, combining both physical and virtual resources. EGI has had to adapt to the changing requirements of user communities by integrating new technologies ranging from new Workload Management Systems (WMS), developed by the VOs, to new computing technologies of cloud computing. The policies, procedures and incident response tools need to adapt to constant development and changes in the infrastructure are constantly reviewed and updated to reflect the new security requirements resulting from these changes.

Operational Security, can only be organized within a proper policy framework, which not only defines acceptable usage and access to the infrastructure, but also allows the CSIRT to enforce a certain software patch level at the RCs and requires the RCs to actively contribute to incident Response. Of course, this set of policies has to be regularly updated and extended to cover the relevant parts of an evolving infrastructure.

Backed by the necessary policy framework, EGI CSIRT [1] maintains operational security for the EGI Infrastructure. This includes responding to computer security incidents affecting the infrastructure (see Section 3). This is carried out by co-ordinating the incident handling activities across the NGIs/EIROs, RCs, VOs, and, where applicable, interacting with partner Infrastructure CSIRTs and CSIRT communities with which EGI-CSIRT has a trust relationship. If needed, RCs are provided with expert level forensics support for the incident investigation and resolution. EGI-CSIRT

[1] `https://www.egi.eu/`

also provides preventive (see Section 5) and educational services such as security monitoring, vulnerability assessment (see Section 4), advisories to mitigate risks due to vulnerabilities and security training. Effective incident prevention reduces the number of incidents the CSIRT has to handle, but still the existing incident handling procedures have to be tested in Security Service Challenges (SSCs) (see Section 6.2) to ensure that the remaining incidents are correctly handled.

The rest of this paper discusses particular areas in more details, highlighting recent developments and achievements.

Critical Vulnerability handling for an infrastructure at the scale of EGI needs to be highly automated. This is achieved with the help of a framework that manages the exchange of information between the above mentioned tools, see Section 3.2.

Based on the results of these SSCs, the incident response procedures and tools are further developed, and specialized training prepared and given at events such as ISGC, DI4R and EGI conferences (see Subsection: 6.2).

2. Development of the Trustmodel and Policies

The development of new security policies in EGI is performed by the EGI Security Policy Group (SPG) [2] The procedures used by SPG to draft, develop and maintain EGI Security Policies and to agree these with the various EGI stakeholders are documented in [3]. Responding to new requirements from EGI and other infrastructures, the development brings new policies or updates to handle the usage scenarios as they evolve in EGI. Architectural assumptions are validated through testing in partnership with user communities under realistic production conditions, and support provided on security issues in close coordination with VO-, NGI- and, local security teams located at the RCs within EGI.

Several new security policies have been produced and many others have been revised and updated recently. These changes were made to properly address issues related to the evolution of EGI services and technology and to mitigate risks identified in recent security risk analyses. It is important to consult widely on the new documents and to take feedback into account. This consultation included presentations and discussions at the regular series of EGI conferences and also in the more frequent meetings of the management boards and with representatives of users and operators.

An overview of approved security policies is in [4] An updated version of the *EGI Acceptable Use Policy* was produced and circulated widely for comments. This version was generalised to include all EGI service offerings (HTC, Clouds, EGI Science Applications on-demand Infrastructure, a.k.a. Long Tail of Science, etc.). At the same time, wording was changed to require appropriate acknowledgement of the use of resources and support received in publications. It also addressed issues of liability.

A revised *Security Policy for the Endorsement and Operation of Virtual Machine Images* was produced using input from a better understanding of the usage of virtual machines in EGI FedCloud. This revised policy included changes to responsibilities and trust to better fit the EGI FedCloud.

A new draft policy and guidelines document entitled *The EGI Science Applications on-demand Security Policy* was produced. This policy (previously known as the Long Tail of Science scoped security policy) aims to enable a low-entry-barrier service to be offered to a wide range of research users in Europe and their collaborators world-wide, by any Resource Centre organisation that elects to do so. Offering such EGI Access Services in alignment with this policy, the RC should not negatively affect the security or change the security risk of any other Resource Centre or any other part of EGI. The document also provides guidelines on the implementation of security procedures and controls. A version of the new EGI Acceptable Use Policy (AUP) specific to the EGI Access Service was also produced and adopted.

Another policy worked on was a complete revision of the policy on Data Protection. This new policy ensures that data collected as a result of the use of the EGI (or other relevant infrastructure) is processed fairly and lawfully by Infrastructure participants. Some of these data, for example those relating to user registration, monitoring and accounting contain “personal data” as defined by the European Union (Directive 95/46/EC). The collection and processing of personal data is subject to restrictions aimed at protecting the privacy of individuals. The policy explicitly does not cover Personal Data relating to third parties included in datasets provided by the End User or the research community to which they belong as part of their research activity. Examples of such data are medical datasets which may contain Personal Data.

Security policy development work included a full revision of the top-level Security Policy document, making it more general and more obviously applicable to all current and new EGI services.

A policy on Acceptable Authentication Assurance was produced and adopted in February 2017. This policy is an update of the old security policy *Approval of Certification Authorities*. It was updated to cover the current Interoperable Global Trust Federation (IGTF) levels of assurance and other changes. This policy defines the approved authentication assurance sources.

During 2016, all remaining old security policies, except for those related to VO management, were updated to make them more general and to use new glossary terms for the current and evolving EGI services. There was no change to the policy content of the documents. The specific policies updated were:

- The VO Portal Policy
- The Policy on e-Infrastructure Multi-User Pilot Jobs
- The Security Traceability and Logging Policy
- The Security Incident Response Policy

The full list of currently adopted security policies is maintained on the EGI policies and procedures wiki at:

- `https://wiki.egi.eu/wiki/SPG:Documents`
- `https://wiki.egi.eu/wiki/Policies_and_Procedures`

2.1 The evolution of trust and policy in AAI and federated identity management

The Interoperable Global Trust Federation[5] (IGTF) is the primary source of identity assurance level specifications in use within the EGI Infrastructure, and EGI maintains a dedicated liaison membership in the IGTF to support its policy and engagement evolution with the EGI user communities. The IGTF is a joint effort that permits global federation of identities and trust, aligning identity assurance requirements also for PRACE (the Partnership for Advanced Computing in Europe), XSEDE and the Open Science Grid in the US, HPCI (the High Performance Computing Infrastructure) in Japan, and a large number of national e-Infrastructures.

The EGI-IGTF liaison function is visible to the resource centres and users primarily by way of the single *trust anchor distribution*: a set of roots of trust that all meet or exceed defined minimum requirements. This distribution remains a key responsibility of the liaison function, which also supports the IGTF in this role, and a number of releases were distributed to the EGI Infrastructure so far. Yet the scope and range of the trust anchors is continuously evolving. Based on a global user requirements analysis with the Research and e-Infrastructures, the IGTF introduced multiple assurance profiles: combinations of assurance elements, structured according to the OGF CAOPS-WG Authentication Service Profile [6], that combine identity assurance elements into a limited set of combined profiles[2] that match specific risk profiles. Three of these (named “ASPEN”, “BIRCH”, and “CEDAR”) all correspond to approximately the same level, but use different underlying authentication technologies: respectively local site, R&E (inter) federation, and end-user-based. In addition an ‘identifier-only’ trust assurance level (“DOGWOOD”) was introduced. This provides persistent non-reassigned identifiers to all entities, alongside a revocation-, freshness-, and traceability capability, but does not convey identity information (names, affiliation) by itself.

The key benefit of this approach is that it allows distribution of authentication responsibilities between identity providers, research communities, and resource centres. Several structured user communities (WLCG, ELIXIR, but also the EGI Access platform) by themselves collect identity data, and require from the authenticator no more than the guarantee of uniqueness (non-reassignment). This brings a new level of flexibility to the EGI trust management system: users are not required to register twice, and the enrolment flow can be simplified. It introduces a combined assurance model, where the combination of identity data from the home organisation (typically a federated organisation in an existing national R&E identity federation) with identity data held by the structured user community provides sufficient means to answer the basic security questions who, what, where, and when, and permits access control decisions to be made. The model was piloted with the four user communities from the LHC experiment by way of a specific trust anchor (the “CERN WLCG IOTA CA”) using existing mechanisms in the *Policy on Approval of Certification Authorities*, pending the necessary software support by the EGI technology providers. Following successful testing, this mechanism was formalised in the *Policy on Acceptable Authentication Assurance* which lays down the requirements on the joint assurance level. The work in EGI and the IGTF was carried out in in close conjunction with the AARC project and the REFEDS community. The evolution of the assurance trust framework in particular permits the integration of the RCauth.eu AARC pilot service, which provides a CILogon-like token translation service for Europe.

The pilot service was connected to the EGI trust framework in a controlled way, since secure introduction of the new class of ‘identifier-only’ trust services requires simultaneous deployment and configuration of authorization software components. This is technically achieved by the in-

[2] `https://www.igtf.net/ap/authn-assurance/`

troduction of a supplementary trust anchor package (“egi-policy-cam”, reflecting the combined assurance/adequacy model name) and made available to the resource centres for controlled testing. Wider adoption will be recommended to the service administrators once authorization software support has been fully deployed. Meanwhile, the mechanism is used by selected user communities in EGI (WLCG, ELIXIR) and in some constituent NGIs (e.g. in NGI-NL) based on the EGI IGTF package distribution.

Other trust issues related to federated identity included input to the work on the Sirtfi activity of REFEDs[3], which is building a trust framework for security incident response in the identity federations in collaboration with the AARC project.

3. Incident Response Task Force

Within the EGI CSIRT, day-to-day security operations are handled by the Incident Response Task Force (IRTF). These security operations not only include the timely reaction to security incidents, but also include activities to prevent incidents. The highly distributed nature of both the Infrastructure and IRTF adds a layer of complexity to these operations, which can only be resolved by ensuring that proper procedures are in place and that these can easily be followed when time is of the essence.

In recent years, IRTF has maintained a high level of operational security by:

- Coordinating incidents on the whole EGI Infrastructure, with other communities, and providing forensic support.
- Handling critical vulnerabilities and increasing the Infrastructure’s resilience against standard attacks, by keeping the installed software up-to-date.
- Revising existing policies and practices or developing new ones to address new technologies, services and threats.

3.1 IRTF organization and duties

IRTF is a small team of just over half a dozen security experts distributed over several countries and multiple organizations. Each team member only works a fraction of their time for IRTF duties, there are no members working in the IRTF full-time. As a result, IRTF duties are organized on shifts: each week, one IRTF member will take the role of the EGI Security Officer on duty while another will stay on stand-by, capable of taking over if EGI Security Officer on duty cannot fulfil their duty. Other members of EGI CSIRT usually do not participate in the task force, except for long term tasks or in case of emergency. This concentration of duties on one individual every few weeks means that, even if half of the working week is spend on duty, less than 10 percent of their annual work time is spend on IRTF duties. However, this also means that tasks must be properly handed over to the next person on duty every week.

IRTF duties rely on tools, all crucial to coordinate actions distributed across organizations, geographical location and time:

- A wiki which contains all internal procedures, howtos and meeting minutes

[3] `https://refeds.org/sirtfi`

- A ticketing system used to track every incident and every interaction with EGI constituents. This system was extended by a custom tools developed within IRTF to automatically create and send uniform tickets to multiple RCs.
- A mailing list, used to discuss within the team or with other teams
- An audio-call system, used weekly to coordinate the team's actions
- jabber, a low latency communication system to coordinate urgent actions with the team members.

One of the critical points for these shifts of different EGI Security Officer on duty is the weekly handover which takes place over an audio-call. During this meeting, the officer on duty from the previous week will briefly explain what happened during their shift, raising attention in particular to any issue still open, giving the opportunity to the team member that is taking over the shift to ask questions. Depending on the size or criticality of ongoing issues, this meeting is also the opportunity to split a large incident and distribute the resulting tasks or to quickly reach an agreement on the next steps to be taken. More importantly, the outgoing duty officer will, if needed, report any trouble or doubt they had, allowing the rest of the team to clarify the situation or plan an update of the tool or procedure in question.

3.2 Procedures and tools evolution within IRTF

IRTF relies for the most part on two procedures which have to be followed by all members of the EGI Infrastructure: the EGI CSIRT Security Incident Handling Procedure, also known as SEC01[4], and the EGI CSIRT Critical Vulnerability Handling, also known as SEC03[5].

The EGI CSIRT Security Incident Handling Procedure, was modified in 2015 (and approved in 2016), following two main objectives. First, the introduction of cloud technologies introduced new capabilities (e.g. taking full memory and disk snapshots of a live VM) that could be used to improve forensic data retention, but also introduced opaque indirection layers which had to be dealt with properly. This procedure update introduced new steps, taking advantage of these new capabilities but also requiring more data to be reported, to ensure full traceability. Secondly, while Virtual Organisations (VOs), and more generally users, have always been part of successful incident handling, their participation was not written into the procedure. Worse, the new virtualisation layer completely blinded sites, which became unable to observe or investigate user actions. As a result, this update included VOs and users directly in the new procedure.

At the same time as this public procedure was modified and while it was used to resolve incidents in 2016 and 2017, the internal procedures for IRTF have also been adapted. Each incident that impacted the EGI Infrastructure during this period was later followed by a debriefing discussion, during which any imperfections of the incident handling could be identified and improvements discussed within the team. While most of these modifications were minor, not leading to major changes in the procedures, the accumulation of these changes and the enhancement of internal documentation have clarified the role, duties and capabilities of response duty coordinators, thereby improving the response capabilities of IRTF.

[4] https://wiki.egi.eu/wiki/SEC01

[5] https://wiki.egi.eu/wiki/SEC03

In addition, the EGI CSIRT provides its constituents with guidelines to investigate and resolve incidents (as part of SEC01). In order to improve their immediate response and decrease forensic data losses, the EGI forensics guidelines[6] have been updated and greatly improved. It now includes detailed steps to collect and analyse forensic information in the context of computer security and also provides configuration advice to increase the range of forensic data that can be collected and decrease data loss and alteration.

Concerning the prevention of incidents, in particular the resolution of critical vulnerabilities that could expose the EGI Infrastructure to a large-scale incident, IRTF relies on the EGI-CSIRT Critical Vulnerability Handling procedure. This procedure, well tested for conventional grid resources, did not require any extension to support the evolving Infrastructure. The procedure was, however, updated in 2015 to a more readable and accessible wiki format, indicating each step in readable tables and following the same model as for other EGI procedures. Moreover, the different timescales for each step have been clarified. The complicated and sometimes incompatible measures of time (7 calendar days versus 3 working days) have been replaced with straightforward incremental delays. While such a change might slightly increase the vulnerability patching delays, clear and predictable deadlines allow RCs to properly align their own procedures. Last but not least, an additional step has been added to the procedure: running the EGI diagnostic tool, pakiti, manually on each of the affected nodes after their patching. This new step, which has been accompanied with a simplification of the deployment of pakiti on worker nodes, allows faster and simpler confirmation of the resolution of the vulnerability.

In response to the new concept of Virtual Appliances[7], which introduced certified configurations and was not covered by the existing procedures, an extension to SEC03 has been developed. This new procedure, which required new communication endpoints and new interaction, has not yet been put into practice. The elements missing for its operational deployment have been identified and are currently being worked on by the responsible teams.

Procedures often have to rely on automated tools which can perform or simplify the various operations required by either a member of IRTF or more generally of the EGI Infrastructure. IRTF developed internal tools in order to decrease the number of repetitive tasks and to standardise the different classes of messages sent. The EGI CSIRT has also adapted itself to the new EGI security dashboard, which exposed more information to EGI constituents but also added another source of information. Such a transition required extra care, as discrepancies between the sources and reports appeared and had to be identified and fixed.

In order to improve the rapid isolation and containment of security incidents, the EGI CSIRT has been pushing for the wider deployment of a central emergency suspension solution. Manual suspension and reinstatement of identities involved in an incident by every EGI service operator has proven itself to be unreliable and leads to delays or partial results. A central solution has been deployed with a master record being distributed through each National Grid Infrastructure (NGI). In order to monitor the effectiveness of this solution, a nagios monitoring probe has been developed and used internally by the IRTF. It is currently under review before publication to be included in EGI's standard operational tests. While the initial deployment will be only for NGIs, it could easily

[6] https://wiki.egi.eu/wiki/Forensic

[7] See https://wiki.appdb.egi.eu/main:faq:what_is_the_egi_applications_database_appdb

be extended to sites that expose their internal servers to the monitoring system.

3.3 Security incidents handled by IRTF in 2016

In 2015, IRTF coordinated three incidents, including a serious one spannning several sites and countries. In 2016, the number of incidents more than doubled, again with a large incident spanning several universities and continents. Unfortunately, for most of these incidents, their cause seemed to indicate that much of the improvement made over the previous years has been lost due to new technologies.

While industry has been using virtualization for quite some time now, resulting in a relatively stable technology, the use of virtualization within EGI only started to take off very recently. The deployment of this new technology resulted in a fundamental shift in computing models. In the more traditional *Grid* computing, system configuration and maintenance were the responsibility of system administrators who have been, over the years, trained and exposed to dealing with security problems. As a result, the security level of the infrastructure was raised, with basic security errors and mistakes almost completely disappearing. In the new *Cloud* computing model, users are directly configuring and maintaining their Virtual Machines (VMs). With this their responsibility evolved dramatically, from issuing simple jobs contained on worker nodes, to running full VMs, exposed to the internet. Unfortunately, this shift to Cloud model has not been accompanied with enough help or training for users which resulted in basic configuration errors. Each of these errors, coupled with the default exposure of VMs to the internet, lead to security incidents that had to be handled by IRTF. These errors included:

- Use of very weak password for administrative access
- Shared file systems exposed with no access control to the internet

In recent years, the EGI CSIRT has tried to push for a different model for *Cloud* computing: provide users with pre-configured solutions that have been vetted by security-aware administrators. Unfortunately, such restrictions are seen by many users as unnecessary hurdles with too many restriction, leading to the current situation. It is even more worrying to discover that some of the incidents were due to *orchestrators*, new systems intended to simplify the management of VMs for users, but currently lacking thorough configuration audit or best practices. The EGI CSIRT will continue to try to push for a better cloud experience for users: we should keep the knowledge that we have accumulated over the years and not restart from scratch.

Virtualization is not the sole source of incidents. The largest incident that happened in 2016, dubbed Venom, was detected in *Grid* computing. This incident was first discovered by a system administrator who was investigating poor performance of a server. After the initial report of anomalous processes, the local security team was able to extract malicious binaries. These binaries, analysed by IRTF [7], revealed a multiple backdoor used by the attacker to maintain control over the system. IRTF was unable to identify, from this system, either the initial vector of the attack or any malicious actions of the attacker, but the analysis of the network traffic pointed to other systems, which appeared to be compromised. IRTF was able to contact owners of some of these systems, confirm that they were compromised and identify further systems for investigation. After the publication of the analysis and the notification of our contacts, it appears that the attacked targeted the astrophysics community more than EGI communities. More than 25 systems were found compromised by other organizations, but unfortunately neither the attackers nor their goals were identified. This incident revealed issues in the coordination of such an incident between different communities, especially on the topics of data sharing and investigation coordination. The EGI CSIRT will continue to try to work with other communities and build the trust relationship that is fundamental to the exchange of information concerning such incidents.

4. Vulnerability Handling

The purpose of the EGI Software Vulnerability Group (SVG) is to minimize the risk to the EGI infrastructure arising from software vulnerabilities. This is an important activity as some of the threats with the highest risk value concern threats due to software vulnerabilities. The biggest way in which this has been enacted has been through handling software vulnerabilities which are reported and are relevant to the EGI infrastructure. Handling software vulnerabilities is an important part of incident prevention, reducing the security risks to the infrastructure.

Towards the end of 2015 the EGI SVG carried out a major revision of the Software Vulnerability Issue handling procedure[8] to address the evolving EGI services. The main purpose of this document is to describe the EGI Software vulnerability Group issue handling procedure, including how to report a vulnerability, which steps are carried out, and the responsibilities of the various parties involved. All types of software vulnerability which are relevant to EGI are handled, which includes software vulnerabilities both ‘discovered’ in software, usually in software developed by persons collaborating with EGI to enable the secure sharing of resources, as well as vulnerabilities announced by software providers. In addition, it briefly describes other strategies for minimizing the risk to the EGI infrastructure due to vulnerabilities.

Previously the main focus of SVG was on handling vulnerabilities in Grid Middleware, and EGI CSIRT handled vulnerabilities in the Linux operating system. In recent years the proliferation of different types of software on the infrastructure has meant that EGI has had to revise the procedure and strategy for handling vulnerabilities. In EGI-Engage it was decided to have one group to handle all vulnerabilities, and members of the IRTF members of SVG. The IRTF therefore also sees all information on vulnerabilities reported to SVG, and if they wish to take urgent action to protect sites then they may. The major revision of the EGI Software vulnerability issue handling took this into account. In addition, the revision took into account the reduction in homogeneity of the EGI infrastructure, and changing technology including the EGI FedCloud.

The Vulnerability issue handling procedure allows that anyone may report a vulnerability, by e-mail to `report-vulnerability@egi.eu` This may be to report a vulnerability discovered in software, or to alert SVG to a publicly announced vulnerability which may be both relevant and a concern to EGI. SVG, along with the reporter and if appropriate the technology provider, investigate the relevance and effect of the vulnerability in EGI. If the issue is valid and relevant, a risk assessment is carried out where the vulnerability is placed in one of four risk categories: *Critical*, *High*, *Moderate* or *Low*. If the vulnerability has not been fixed, a target date is set according to risk:

- *Critical* – a special process is carried out according to the circumstance,

[8] `https://documents.egi.eu/public/ShowDocument?docid=2538`

- *High* – 6 weeks
- *Moderate* – 4 months
- *Low* – 1 year.

This target date is the date by which software free from the vulnerability should be available for installation in all appropriate repositories. This ranking allows the prioritization for the timely fix of software vulnerabilities. An advisory is issued:

- If EGI SVG is the main handler of vulnerabilities concerning this software, regardless of the risk
  - when it is fixed, or
  - on the Target date if it is not fixed by then.
- If the issue is assessed as *High* or *Critical* risk
- If the EGI SVG considers it useful to alert sites

Since the advent of the EGI FedCloud, consideration of vulnerabilities in software enabling the EGI FedCloud as well as software included in virtual machines has had to be considered. Cloud enabling software is handled in a very similar way to Grid enabling software. Software in Virtual Appliances presents new challenges, which are not fully addressed in the current procedure. If a critical vulnerability is found related to a Virtual Appliance, then this needs to be updated urgently by those responsible for the virtual appliance.

A wider range of software and technology is being used in EGI than in the past, and it is not reasonable to expect there to be expertise on all the technology used in EGI within the software vulnerability group. SVG cannot control what software is used in the EGI infrastructure. Hence SVG relies more on the software providers to analyse vulnerabilities in cases where there is not the expertise in the group, and vulnerabilities are handled more from a procedural point of view.

One addition to help with the greater proliferation of software is to ask those who develop or select software to consider security and maintainability. To help this SVG produced a Software Security Checklist[9], which at present consists of 10 points which people should consider when developing or selecting software to avoid some of the common problems from which vulnerabilities arise or which make it difficult to address if they do.

The number of issues handled during the 2 years since the start of EGI-Engage has been tracked. Between 1st March 2015 and 31st March 2017, 98 vulnerabilities have been reported and handled by SVG. During that time 55 advisories have been issued including for 12 which were assessed as *Critical* risk and 23 assessed as *High* risk. The types of software where issues are reported have changed. Of the 98 potential vulnerabilities handled 19 concerned Grid Middleware. 16 concerned cloud enabling software, and 10 concerned the Linux kernel.

[9] `https://wiki.egi.eu/wiki/SVG:Software_Security_Checklist`

5. Security Monitoring

In order to obtain information about the state of the infrastructure, the EGI CSIRT uses a security monitoring framework that collects and evaluates data from EGI sites. The information from security monitoring enables the EGI CSIRT to check basic security characteristics of services that EGI sites make available to the users.

For the purpose of security monitoring, the EGI CSIRT maintains a set of security probes to collect individual characteristics. The probe suite consists of both general probes and probes focused on particular problems that were identified during operations. Typical representatives of the former category are probes to reveal expired certificate revocation lists or dangerous access rights on files. The latter category is typically represented by probes checking the existence of precautions to mitigate security vulnerabilities announced by the EGI CSIRT.

The execution of security probes and basic processing of results is performed by the central EGI security monitoring service. This service operates a standard EGI monitoring suite that utilizes Nagios as the engine to schedule and execute the probes and aggregate the results produced. The engine runs the probes on regular basis so that every EGI site provides results at least once a day, usually more often. The probes are run as normal compute jobs submitted to the tested site so that they exploit as much as possible the standard interfaces commonly used by users. The data collected therefore refer to a situation that could be seen by a malicious user or an attacker who managed to impersonate a legitimate user. The drawback of this approach is that that the EGI security monitoring cannot obtain information from internal services that are not accessible by Grid users.

EGI CSIRT's operational experience has shown that security vulnerabilities are very often used by attackers to obtain unauthorized access to systems and therefore pose a significant risk to our computing infrastructure. Thus, it is essential that all resource providers in the EGI update their systems with the latest security fixes as soon as they are available. However, the experience gathered over several years of operations in the EGI and its predecessors clearly shows that ensuring a homogeneous level of security across multiple, often heterogeneous, resources is challenging. This is especially the case when applying software updates, which require technical expertise and significant coordination efforts, and in many cases service downtime is inevitable. Unfortunately, failure to promptly apply security updates remains one of the main causes of security incidents affecting EGI's computing infrastructure. The central EGI security monitoring service controls the engine which runs the security probes. In order to gain visibility of security patching status across EGI, the Pakiti monitoring system[8] has been developed.

Pakiti is a client/server solution with the server collecting information about installed software packages reported by the clients running on particular nodes of the infrastructure. The server evaluates the information and makes the results of the evaluation available for further checking. The Pakiti client is part of the standard security probe set and is used to send reports from the monitored sites to the EGI Pakiti server.

Due to the large number of resources joining the EGI e-Infrastructure it is becoming increasingly challenging for the EGI CSIRT to follow up all identified security issues. To solve this problem and scale up the operational capability, a security dashboard has been developed, which allows resource providers' security officers and their NGI operations staff to access the monitoring results, and therefore to handle the issues directly. The dashboard aggregates the data produced by different security monitoring components and provides interfaces to its visualization. Access to the collected data is subject to strict access control so that sensitive information is only available to those that need it. The security dashboard was developed as a specific module of the common EGI Operations portal and EGI CSIRT believes that the handling of some security issues, such as patching known vulnerabilities, should be incorporated with current (non-security) issue handling procedure, to significantly reduce the overall operation cost.

5.1 Cloud assessment

Current utilization of cloud introduces new possibilities for how vulnerabilities can be exposed. A commonly seen pattern is an image of a virtual machine that exposes vulnerabilities which can be easily detected once the virtual machine is instantiated and connected to the Internet. In order to detect common vulnerabilities, the Secant framework for image assessment has been developed.

Secant runs as a service that periodically checks for new images available in a repository and performs their security assessment. When a new image becomes available in the system, it is taken by Secant and checked for security vulnerabilities. In order to perform the security checks, Secant instantiates a virtual machine from the appliance that is being verified and performs two phases of security checks. During the first phase, Secant launches a series of external scans that try to detect vulnerabilities exposed by the machine to the Internet. Following these tests, and if the machine supports it, Secant runs a series of internal probes on the virtual machine which check security properties of the installed software. Both internal and external probes are modular and new tests can be easily added when needed. After the probes are executed, Secant processes the results and generates the assessment report.

6. Security Service Challenges and Trainings

6.1 Security Service Challenges

The incident response capabilities of security teams active in EGI at all levels (Resource Center, NGI, Global/project level) are tested with Security Service Challenges (SSCs). Here EGI CSIRT creates a realistic scenario of an incident spreading in the Infrastructure and lets the teams use their incident response tools. Not only the effectiveness of the deployed centralized incident response tools are tested but also the response procedures and EGI policies are checked to confirm that they support an efficient incident response.

Plans for an enhanced security challenge framework for the EGI FedCloud were agreed by the EGI CSIRT. The framework has been developed and is ready and running and has been tested at a number of sites. More work, however, is required in the usage of contextualisation and configuration of EGI FedCloud sites before a full challenge can be run.

6.2 Security Training

The necessary skills for proper incident response are usually beyond the experience of system administrators, in particular in specialised environments. In EGI it is crucial to have a deep understanding of the technology to be able to use the available information for a more complete incident response. Security training is therefore of vital importance.

The training courses have been developed and offered to the participants are in 3 major categories:

Defensive training. The participants are administrators of a virtual Grid Site and have to defend against attacks performed by the trainers. The focus of this training is to detect anomalies in the system, understand the origin (attack vector) and communicate the results to trainers. The basis for these exercises are real incidents handled by EGI-CSIRT. The emphasis here is also to improve the forensic skills of the participants.

Offensive training. Here the participant takes the role of an attacker and they are asked to attack a provided virtualised grid site with known vulnerabilities. The scope here is to demonstrate how attackers operate and how easy it is to compromise a system with information available on the internet. The target audience here are also EGI FedCloud users that have to manage virtual machines.

Role Play training. Here the participants are presented with a “what if” situation and role-play possible incident response problems when new technology is part of an incident. Here one can look at collaboration and communication issues, not only of the security teams, but also involve management, press officer and others in the process.

During EGI-Engage to date 10 security training events have been run with 20 participants each on average.

7. Information Security Management - collaboration with other e-Infrastructures and Research Infrastructures

The EGI-Engage project takes a leading role in the coordination of Information Security management for e-Infrastructures. Following on from initial discussions at the EGI Conference in Lisbon in May 2015 we (EGI Security and a PDO from the GEANT Amsterdam office) decided to pursue the possibility of enabling the Security for Collaborating Infrastructures (SCI) activity, chaired by the EGI Security Coordinator, to meet jointly with the newly formed SIG-ISM activity of GEANT. This SIG acts as an information exchange forum to discuss standards and best practices for Information Security in the NRENs. It was recognised that there could be great benefits from the e-Infrastructures and NREN communities working closer together on security topics of common interest. Further discussions in June 2015 created a small programme committee for what became ”The First WISE meeting” in Barcelona on 20-22 October 2015. Approximately fifty people in total attended made up of representatives of the EU e-Infrastructures (EGI, EUDAT, GEANT and PRACE) together with representatives of many NRENs, participants from the USA (XSEDE, NCSA, CTSC) and communities like LIGO, HEP/CERN, Human Brain Project and others.

WISE stands for ”Wise Information Security for Collaborating E-infrastructure” and is a global trust community where security experts share information and work together, creating collaboration among different e-infrastructures. WISE provides a framework of standards, guidelines, and practices to promote the protection of critical infrastructure.

The first meeting was a success and it was agreed that we would continue to work together, meet face to face twice a year and in the meantime make actual progress in a number of working groups. The EGI Security Coordinator is co-leading a working group called ”SCI-V2WG” which will take the SCI Version 1 document forward to include more stakeholders, e.g. NRENs, in the defined Trust Framework. EGI will wherever possible also be active in other working groups; Risk Analysis, Security in Big Data/Open Data, Review and Audit, training and awareness. Members of the EGI security team will also continue to serve on the WISE steering committee.

Further meetings of WISE were subsequently held at the XSEDE meeting in the USA in July 2016, at the Digital Infrastructures for Research Conference in Krakow in Sep 2016 and then hosted by Nikhef in Amsterdam in March 2017.

Through its broad representation in the steering committee, EGI can benefit from better alignment of security practices and increased input into its risk assessment, training and trust programmes. The construction of comparative policy frameworks allows easy movement of researchers and data across multiple Infrastructures, whilst the construction of the human network through periodic WISE workshops (twice-yearly) establishes communications channels also for operational security activities.

Acknowledgment

The authors acknowledge the support and collaboration of many other colleagues in their respective institutes, NGIs, national CSIRTs, VOs and IT Infrastructures, together with the funding received by these from many different sources. These include but are not limited to the following:

a.EGI acknowledges the funding and support received from the European Commission and the many National Grid Initiatives and other members. The EGI-Engage project is co-funded by the European Commission (grant number 654142).

b. The Worldwide LHC Computing Grid (WLCG) project is a global collaboration of more than 170 computing centres in 42 countries, linking up national and international grid infrastructures. Funding is acknowledged from many national funding bodies and we acknowledge the support of several operational infrastructures including EGI, OSG and NDGF/NeIC. The work reported here has been carried out in collaboration with the EGI-Engage project (Horizon 2020) under Grant number 654142.

References

[1] EGI CSIRT. Egi csirt public wiki. `https://wiki.egi.eu/wiki/EGI_CSIRT:Main_Page`.

[2] David Kelsey. Security Policy Group (SPG) – Terms of Reference. `https://documents.egi.eu/public/ShowDocument?docid=64`, 2011.

[3] David Kelsey and David Groep. Security Policies within EGI. `https://documents.egi.eu/public/ShowDocument?docid=210`, 2010.

[4] David Kelsey. EGI Approved Security Policies. `https://wiki.egi.eu/wiki/SPG:Documents`.

[5] J. Astalos, R. Cecchini, B. Coghlan, R. Cowles, U. Epting, T. Genovese, J. Gomes, D. Groep, M. Gug, A. Hanushevsky, M. Helm, J. Jensen, C. Kanellopoulos, D. Kelsey, R. Marco, I. Neilson, S. Nicoud, D. O'Callaghan, D. Quesnel, I. Schaeffner, L. Shamardin, D. Skow, M. Sova, A. Wäänänen, P. Wolniewicz, and W. Xing. *International Grid CA Interworking, Peer Review and Policy Management Through the European DataGrid Certification Authority Coordination Group*, pages 285–295. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

[6] Christos Kanellopoulos and David Groep. Framework for writing Authentication Profiles for use in Grids. `https://redmine.ogf.org/dmsf_files/29?download=`, 2009.

[7] EGI CSIRT. Analysis of the VENOM Linux rootkit. `https://wiki.egi.eu/wiki/Venom_Rootkit`, 2017.

[8] Michal Procházka, Daniel Kouřil, Romain Wartel, Christos Kanellopoulos, and Christos Triantafyllidis. A Race for Security: Identifying Vulnerabilities on 50 000 Hosts Faster than Attackers. In *Proceedings of the International Symposium on Grids and Clouds (ISGC) 2011*, Taipei, Taiwan, 2011. Academia Sinica.