

SCI: A Trust Framework for Security Collaboration among Infrastructures

Authors

K. Chadwick (FNAL), I. Gaines (FNAL), D. Groep (Nikhef), U. Kaila (CSC), C. Kanellopoulos (GRNET), D. Kelsey (STFC), J. Marsteller (PSC), R. Niederberger (FZ-Juelich), V. Ribaillier (IDRIS), R. Wartel (CERN), W. Weisz (University of Vienna), J. Wolfrat (SURFsara)

Abstract

The Security for Collaborating Infrastructures (SCI) group is a collaborative activity of information security officers from several large-scale distributed computing infrastructures, including EGI, OSG, PRACE, EUDAT, CHAIN, WLCG, and XSEDE. SCI is developing a framework to enable interoperation of collaborating Grids with the aim of managing cross-Grid operational security risks and to build trust and develop policy standards for collaboration especially in cases where we cannot just share identical security policy documents.

Table of Contents

1	Introduction	2
2	Operational Security [OS]	3
3	Incident Response [IR]	3
4	Traceability [TR].....	4
5	Participant Responsibilities [PR].....	4
5.1	Individual Users	4
5.2	Collections of Users.....	5
5.3	Resource Providers and Service Operators	5
6	Legal Issues and Management procedures [LI]	6
7	Protection and processing of Personal Data/Personally Identifiable Information [DP].....	6
8	Acknowledgements	6

1 Introduction

In recent years we have seen the implementation of a variety of infrastructures supporting distributed computing environments and sharing of resources. Each such infrastructure consists of distributed computing and data resources, users (who may be organised into separate user communities), and a set of policies and procedures. Examples of such infrastructures include computing grids and/or clouds, as well as cooperating computing facilities managed by different organisations.

Even when such an infrastructure considers itself to be decoupled from other infrastructures, it is in fact subject to many of the same threats and vulnerabilities as other infrastructures because of the use of common software and technologies. Moreover, there may be users who take part in more than one infrastructure and are thus potential vectors that can spread infection from one infrastructure to another. Finally, one infrastructure may want to extend rights to use its resources to users who are enrolled in a different infrastructure. In each of these situations, the infrastructures can benefit from working together and sharing information on security issues.

Security in a distributed collaborative environment is governed by the same principles that apply to a local centrally managed system, but complicated by the diversity of sites (both in terms of hardware and software systems and in terms of local policies and practices that apply), and by the lack of a centralized management hierarchy that can "order" certain operations to be performed in specific ways.

Governing principles include:

- The management of risk; both to mitigate the most likely occurring and dangerous risks, and to take counter measures that are commensurate with the scale of the involved risks
- Containing the impact of a security incident while keeping services operational, but in certain cases this may require identifying and fixing a security vulnerability before re-enabling user access
- Identifying the cause of incidents and understanding what measures must be taken to prevent them from re-occurring
- Identifying users, hosts and services, and controlling their access to resources, all of which must be sufficiently robust and commensurate to the value of the resources and the level of risk and must comply with the regulatory environment

In this document we lay out a series of numbered requirements in six areas (operational security, incident response, traceability, participant responsibilities, legalities, and data protection) that each infrastructure should address as part of promoting trust between infrastructures.

To evaluate the extent to which the requirements described in this document are met, we recommend that each infrastructure assess the maturity of its implementation according to the following levels:

Level 0: Function or feature not implemented

Level 1: Function or feature exists, is operationally implemented but not documented

Level 2: Function or feature is comprehensively documented and operationally implemented

Level 3: Function or feature implemented, documented, and reviewed by an independent external body

We encourage openness and transparency in the documentation and for Levels 2 and 3 we recommend that wherever possible such documents should be made available to collaborating infrastructures as a way of promoting trust.

2 Operational Security [OS]

Retaining operational availability and integrity is the most urgent and visible aspect of security. Each of the collaborating infrastructures must therefore have the following:

- [OS1] A security model addressing issues such as authentication, authorisation, access control, confidentiality, integrity and availability, together with compliance mechanisms ensuring its implementation
- [OS2] A process that ensures that security patches in operating system and application software are applied in a timely manner, and that patch application is recorded and communicated to the appropriate contacts
- [OS3] A process to manage vulnerabilities (including reporting and disclosure) in any software distributed within the infrastructure. This process must be sufficiently dynamic to respond to changing threat environments
- [OS4] The capability to detect possible intrusions and protect the infrastructure against significant and immediate threats on the infrastructure
- [OS5] The capability to regulate the access of authenticated users
- [OS6] The capability to identify and contact authenticated users, service providers and resource providers
- [OS7] The capability to enforce the implementation of the security policies, including an escalation procedure, and the powers to require actions as deemed necessary to protect resources from or contain the spread of an incident

3 Incident Response [IR]

The management of risk is fundamental to the operation of any Infrastructure. Identifying the cause of incidents is essential to prevent them from re-occurring. In addition, it is a goal to contain the impact of an incident while keeping services operational. For response to incidents to be acceptable this needs to be commensurate with the scale of the problem.

It is imperative that every infrastructure has an organized approach to addressing and managing events that threaten the security of resources, data and overall project integrity.

Each infrastructure must have the following:

- [IR1] Security contact information for all service providers, resource providers and communities together with expected response times for critical situations
- [IR2] A formal Incident Response procedure. This must address: roles and responsibilities, identification and assessment of an incident, minimizing damage, response & recovery strategies, communication tools and procedures
- [IR3] The capability to collaborate in the handling of a security incident with affected service and resource providers, communities, and infrastructures

- [IR4] Assurance of compliance with information sharing restrictions on incident data obtained during collaborative investigations. If no information sharing guidelines are specified, incident data will only be shared with site-specific security teams on a need to know basis, and will not be redistributed further without prior approval

4 Traceability [TR]

The minimum level of traceability for the Infrastructure is to be able to identify the source of all actions (executables, file transfer, etc.) together with the individual¹ initiating the actions. In addition, sufficiently fine-grained controls, such as blocking the originating user, system or service and monitoring to detect abnormal behaviour, are necessary for keeping services operational. It is essential to be able to understand the cause and to fix any problems before re-enabling access for the user.

The aim is to be able to answer the basic questions "who, what, where, when and how" concerning any incident. This requires retaining all relevant information, including accurate timestamps and the digital identity of the initiator, sufficient to identify, for each service instance, and for every security event including at least the following: connect, authenticate, authorize (including identity changes) and disconnect.

Each infrastructure must provide the following:

- [TR1] Traceability of service usage, by the production and retention of appropriate logging data, to identify the source of all actions as defined above
- [TR2] A specification of the data retention period, consistent with local, national and international regulations and policies
- [TR3] A specification of the controls that the resource provider implements to achieve the goals of [TR1]

5 Participant Responsibilities [PR]

All participants in a group of collaborating infrastructures need to rely on appropriate behavior by various actors in both their own and other infrastructures. We separate these responsibilities into behavior expected of:

- Individual users
- Collections of users
- Resource providers and service operators

Each infrastructure must ensure that the various participants are aware that they have these responsibilities.

5.1 Individual Users

Each infrastructure must provide:

- [PRU1] An Acceptable Use Policy (AUP). The AUP must at least address the following areas: defined acceptable use, non-acceptable use, user registration, protection and use

¹ For software agents initiating actions there must be a human individual responsible for all actions of the agent

of credentials, data protection and privacy, Intellectual Property Rights (IPR), disclaimers, liability and sanctions

- [PRU2] A process to ensure that all users are aware of, and accept the requirement to abide by, the AUP
- [PRU3] Communication to their users of any additional restrictions or requirements on acceptable use that arise out of new collaborative partnerships

5.2 Collections of Users

A Collection of users is a group of individuals organised around a common purpose jointly granted access to the Infrastructure. It may serve as an entity which acts as the interface between the individual users and each Infrastructure. In general the members of the Collection will not need to separately negotiate with Resource Providers or Infrastructures.

Examples of Collections of users include: User groups, Virtual Organisations, Research Communities, Virtual Research Communities, Projects, Science gateways, and geographically organised communities.

Each infrastructure must have:

- [PRC1] A process to ensure that all Collections of users using their infrastructure are aware of, and accept the need to abide by, various policy requirements
- [PRC2] Policies and procedures regulating the individual user registration and membership management (registration, renewal, suspensions, removal, and banning). At a minimum these must address the accuracy of contact information both for initial collection and periodic renewal

Collections of users must:

- [PRC3] be aware that they will be held responsible for actions by an individual member of the collection which in turn may reflect on the ability of other members to utilise the infrastructure
- [PRC4] ensure a way of identifying the individual user responsible for an action
- [PRC5] keep appropriate logs of membership management actions² sufficient to participate in security incident response
- [PRC6] define their common aims and purposes and make this available to the Infrastructure and/or Resource Providers to allow them to make decisions on resource allocation

5.3 Resource Providers and Service Operators

The Infrastructure must have policies and procedures in place to ensure that Resource Providers and Service Operators understand and agree to abide by expected standards of behaviour, including:

- [PRR1] vulnerability patching

² Examples include but are not limited to: Registration or renewal in a membership system, dynamic authorisation such as acquisition of VOMS attributes, authentication to a Science Gateway or portal, job submission or file transfer initiated by the Collection on behalf of an individual user

- [PRR2] incident reporting
- [PRR3] physical and network security
- [PRR4] confidentiality and integrity of data
- [PRR5] retention of appropriate logs

6 Legal Issues and Management procedures [LI]

Infrastructures, resource providers, service providers and collections of users must have policies and procedures, appropriately communicated to all participants, that address legal issues including but not limited to the following:

- [LI1] Intellectual Property Rights clarifying the rights and obligations of the participants
- [LI2] Liability responsibilities and disclaimers to make the participants aware of their obligations
- [LI3] Software licensing clarifying the rights and obligations of the participants
- [LI4] Dispute handling and escalation procedures
- [LI5] Data Protection responsibilities (also see the next section)
- [LI6] Any additional regulations such as export controls, ethical use, externally imposed data protection and/or access control requirements
-

7 Protection and processing of Personal Data/Personally Identifiable Information [DP]

Infrastructures, resource providers, service providers and collections of users must have policies and procedures addressing the protection of individuals with regard to the processing of their personal data (PII) collected as a result of their participation in the infrastructure, including but not limited to:

- [DP1] Accounting Data
- [DP2] User Registration Data
- [DP3] Monitoring Data
- [DP4] Logging Data
- [DP5] Data owned by or produced by Users or Collections of Users

8 Acknowledgements

<awaiting input from each of the authors>