AARC LoA Survey Responses for EGI

Document identifier	EGI-SCG-AARCLoASPresponses-v1.0a (draft)
Document Link	*Not assigned*
Last Modified	20/10/2015
Version	1.0a
Policy Group Acronym	SCG
Policy Group Name	Security Coordination Group
Contact Person	David Groep
Document Type	N/A
Document Status	DRAFT
Approved by	N/A
Approved Date	N/A

TABLE OF CONTENTS

COPYRIGHT NOTICE

AUTHORS LIST

	Name	Partner/Activity/ Organisation/Function	Date
From	David Groep	Nikhef	15/10/2015

DELIVERY SLIP

	Body	*Date*
Reviewed by:	EGI Security Coordination Group (informational)	2015-10-19
Reviewed by:		
Reviewed by:		
Approved by:		

DOCUMENT LOG

Issue	*Date*	*Comment*	*Author/Partner*
v.1	20151020	Initial version	DLG
v.1.0a	20151020	Add reference to risk assessment basis and differentiated responsibilities	DLG
…			
v.n			

TERMINOLOGY

A complete project glossary is provided at the following page: http://www.egi.eu/about/glossary/

APPLICATION AREA

This is an informational document.

POLICY/PROCEDURE AMENDMENT PROCEDURE

N/A.

1 Introduction

The European e-Infrastructure EGI is a publicly funded e-infrastructure put together to give scientists access to more than 530,000 logical CPUs, 200 PB of disk capacity and 300 PB of tape storage to drive research and innovation in Europe. The infrastructure provides both high throughput computing and cloud compute/storage capabilities. Resources are provided by about 350 resource centres who are distributed across 56 countries in Europe, the Asia-Pacific region, Canada and Latin America.

EGI is coordinated by EGI.eu, a not-for-profit foundation created to manage the infrastructure on behalf of its participants: National Grid Initiatives (NGIs) and European Intergovernmental Research Organisations (EIROs). The foundation is governed by a Council of participant countries and institutions.

2 Questions on the research infrastructures/communities

The EGI infrastructure supports a wide range of research communities: some very large and data intensive (it is providing services to communities like the LHC experiments or earth sciences) as well as smaller communities that may start nationally and then expand across countries (like biomedical engineering and medical imaging).

EGI is operated as a federated trans-national infrastructure, organised around communities that span multiple countries – it is therefore complex to describe exactly which communities are served, and how these are distributed across countries. There are also many global communities that are so supported, as well as national collaborations that are registered with EGI for to enable resource sharing [http://operations-portal.egi.eu/vo/search]

For most of the large and structured communities, EGI leverages the IGTF trust fabric, implemented through X.509 technology. It today accepts identity assertions issued under the IGTF Classic, MICS and SLCS authentication profiles [https://documents.egi.eu/document/83], that provide end-entity traceability.

The EGI assurance needs are loosely based on a E2E chain risk assessment [https://documents.egi.eu/document/863 (under revision since LToS and IaaS are not yet included)]. Based on a risk-assessment EGI provides for ‘differentiated and redistributed’ responsibilities, leading to two noticeable exceptions. Using credentials issued to automated agents (“Robots”) for whom a designated individual is responsible, access to the infrastructure is also granted to:

- ‘science gateway’ services that offer restricted or pre-defined access to resources, or offer access to the general public, a ‘Portal policy’ [https://documents.egi.eu/document/80] allows alternative and/or lower-assurance end-user credentials (or even anonymous access) when off-set by compensatory controls limiting the kind of work to be executed

- On designated resources (usually IaaS cloud systems) EGI-vetted end-users that have quantitatively limited needs (“Long Tail of Science” users), and on systems that have additional security controls in place as per a specific policy [https://wiki.egi.eu/wiki/SPG:Drafts:LToS_Service_Scoped_Security_Policy], any workload is accepted on a per-user traceable basis. The user is not issued a specific credential, but a science gateway or portal will generate user-specific ‘labelled’ delegation credentials (“per-user sub-proxies”) that are subsequently used by the portal to access the resources

Due to lack of software support, EGI is not at the moment capable of moving the ‘traceability’ control from the identity provider (IGTF) to the community (VO or science gateway) or a per-community basis. It therefore does not at the moment support the IGTF “IOTA” authentication profile. Such support, permitting access control decisions based on a combination of “VO && CA” (plus local policies) is foreseen but not in use today.

For the science gateways and Long Tail of Science (LToS) systems, EGI either encourages (portals) or actively pushes for the use of federated AAI to authenticate the end-users to the portal. When used without additional assurance profiles, such credentials are considered “authenticated” but not strongly authenticated, and (unless augmented by vetting in a special EGI user management portal as for the LToS service) do not permit full-service access to the infrastructure (they can be used for portals up to data management level, but not for running arbitrary compute jobs).

Technical reasons have traditionally favoured PKI and ‘proxy (RFC3820) based solutions for resource access. RFC3820 proxies allow client-initiated delegation and authentication for non-web access. Ironically, this makes PKI the ‘simple and easy’ way of granting access to resources in EGI, despite its known end-user usability challenges. For all non-web access, EGI relies on PKI. This is however orthogonal to any LoA issues, since the nominal ‘second factor’ apparently provided by end-user PKI is not in fact used: the PKI credentials can be obtained by end-users through a variety of means including single-factor username password with quality controls (e.g. through the Trusted Certificate Service, the DFN AAI SLCS, or CILogon).

Who are your end users (who need to log in to your services):

- researchers with a Home Organisation (that operates or potentially operates an IdP)? *YES*
- citizen scientists? *Not yet, but science gateways may be offering services to the genral public including citizen-scientists – subject to the VO portal policy (allow 'canned' jobs and parameter-sweeping portals)*
- students with a Home Organisation (that operates or potentially operates an IdP)? *Yes (there can be students in the research groups as well)*
- else/what? *Includes research from SME and private sector pre-competitive R&D (just a small fraction). There are ongoing experiments with a per-for-use model for general access, and more heterogeneity may exist at the NGI level.*

If you are a research community *EGI is not in itself a research community, but supports a broad spectrum of different cases*.

- is affiliation of a researcher (user) with your community typically longer lived than any organizational affiliation or employment, or does community membership stem primarily from organizational affiliation? *Movement of researchers within the same user community (VO) but between employers/host organisations is common. The VO community management infrastructure (mostly VOMS) has mechanisms to permit multiple identity credentials to be lonked to a person's VO membership entry.*
- do you consider yourself also as a source of (identity) assurance for your community members? *For the LToS service, the EGI User Management Portal (UMP) performs additional identity verification and eligibility checks for its users, based on (usually remote) checks by community engagement officers at the NGI level. The check is mostly around workload eligibility (scientific relevance), with remote but lowish-quality ID checking added (usually remote). Checking requirements do not conform to known levels, but are described in the LToS service security policy.*

3 Questions on Identity and Authentication

User's "network identity" distinguishes him/her from other users of the SP.

3.1 Identity concept

How important is it for you that

- all user identities (accounts in the Home Organisation) belongs to an individual person (i.e. there are no shared accounts like "libraryuser1". Any robot/automated agent is traceable to a named person)? *Yes, mandatory. Robots are in use, and must either be tracable to an individual named person, or to a designated organisational group with compensatory incident response requirements [as per https://www.eugridpma.org/guidelines/robot].*

- and all users are traceable (i.e. the Home Organization knows who they are and can reach them)? *Mandatory, as per IGTF ASPEN, BIRCH or CEDAR. For LToS, checks are done on a yearly basis or as soon as the resource allocation runs out.*
- and the Home Organisation is willing to collaborate with you if you think their user misbehaves in your service? *Required. For authentication identity services EGI requires this as per the accepted IGTF LoAs. For the LToS service, where this requirement cannot be put on the upstream IdPs in eduGAIN, this is implemented through the additional EGI User Management Portal*
- that you (as an SP) can block him/her from your service? *Mandatory. EGI expects stable unique identifiers and does not anticicapte and cannot accommodate rapidly changing identifiers for a single user. EGI has an emergency suspension mechanism in place to instantly (1-6hrs) block access to any identifier.*

- user identifiers are persistent i.e. a user account is not re-assigned (re-cycled) to another person over time? *Mandatory The EGI resources and data stores need a specific non-reassigned identifier. Data may be stored for 50+ years. EGI can accommodate multiple identifiers for the same person, as long as these changes are not frequent. Re-assigning identifiers would cause inadvertent changes of data ownership, resource allocations, accounting, and incident response processes.*
- user identifiers are shared by multiple SPs i.e. if you have 2 SPs, do they both receive the same user identifier when the same user logs in to the two services? *Mandatory, e.g. the same information will be needed on different EGI systems as EGI uses a cross-domain cross-service brokering architecture.*

3.2 Initial proof of identity

How important is it for you that

- the Home Organization has a documented identity vetting process (whatever it is) in English and you can study it? *For full access to the infrastructure, EGI requires the IGTF ASPEN, BIRCH, or CEDAR processes including its disclosure clauses that have been co-defined by EGI. For some restructured science gateways/portals, vetting requiremetns may be relaxed as long as the work does not leave persistent state in the infrastructure and compensatory controls are in place (see VO portal policy)*
- each Home Organisation has a machine-readable tag that indicates how the organization carries out identity proofing and the tag is from a well-defined international vocabulary? *Following PRACE here: quite desirable. e.g. IGTF has different profiles and PRACE can rely on them. EGI can evolve the vocabulary, through.*

- each user in a Home Organisation has the above tag and different end users in the same organization can have different tags (depending how their identity was initially proofed)? *Following PRACE: Desirable, there may be use cases*
- the identity proofing is done face-to-face based on a government photo-ID or equivalent? *Vetting requirements depend on the level of access to the infrastructure granted. While "F2F" has been easiest to describe, alternative models are supported.*

For the VO portals and LToS, the access may be either based on lower vetting assurance (restricted service) or based on existing relationships (LToS). It should be noted that, with EGI accepting the IGTF ASPEN and BIRCH LoA levels, it supports 'time-delayed' or other vetting levels based on ongoing business relationships, mediated 'out-of-band' behind provisioning systems (TCS, CILogon, &c).

3.3 On-line authentication

- Are password-based authentication good enough for you? *De facto yes for the moment. Although the PKI credentials used in EGI are technically a two-factor system, the certificates can be obtained also via password-based systems where password quality and account current-ness are controlled (TCS, DFN SLCS AAI).*
- Should passwords have some kind of quality floor? (What kind of quality floor?) *Yes, minimum length and a combination of lower and uppercase characters (of European languages), digits, and non-alphanumeric characters. For PKI credentials, it is a (non-enforceable) requirement to use at least 12 characters.*
- Do you need two factor authentication? (What kind of?) Are you willing to share its costs?. *Depends on the use case, level of access, and the community served. At the moment, EGI has not received concrete requests from the communities, and for the resource centres two-factor is not at the moment an issue, since the mutli-stage approach to getting end-entity PKI credentials is apparently a great barrier also for attackers (there are easier ways into the system for attackers). Cost participation is not foreseen. IGTF certificates in use currently.*

3.4 Step-up authentication as a service

Step-up authentication means that the user first authenticates with a password, and subsequently with a second factor (such as by a one-time password delivered to his/her cellphone). Step-up authentication could be delivered to research communities as a service.

Would you like to make use of step-up authentication

- if it costs you money? *Not EGI itself, but it's user communtiies may – and they would have to bear the cost. Implementation at the resource centres would need to be negotiated by the VOs.*
- if it costs you work (for instance, you need to operate one or several registration authorities where your community's users come to show their photo-ID and you record their cellphone number)? *EGI is used to expending effort on PKI, a move to a different technology that is not more complex but does offer true two-factor would justify work on the infrastructure.*

4 Questions on user attributes

Besides an identifier, the Home Organisation's Identity Provider is able to deliver also other attributes of the person that logs in.

4.1 Freshness of user accounts and attributes

Many Home Organisations close the user account when an individual departs (e.g. researcher changes his/her employer). Closing the account closes also federated access to your SP. However, some organisations keep the accounts open (e.g. to serve alumni etc).

- Do you expect that user accounts are closed as a user departs? How promptly? *EGI expects the possibility for authentication towards the infrastructure to be stopped when tracability to the individual is lost. For most services it is however not the only mechanism that EGI has to terminate access: that potential is also vested in the community membership services (VOMS) and – for the LToS service – in the User Management Portal. It is furthermore important that users who's authenticator is based on a home organisation but who, whilst changing home organisation, remain active in their community. For those users it will be important to have a period in which they can associate any new identifiers/authenticators with their existing VO membership. EGI does expect any attributes beyond the non-reassigned identifier to accurately reflect reality, although at the moment it does not rely on these. All EGI resources have local authorization engines that can make policy decisions based on attributes (VO membership, user identifier, 'identity provider' in the sense today of 'PKI credential issuing CA'). Although by policy account close should be prompt, the multiple controls in practice allow lenience and EGI today accommodates up to 13 months grace period for person-bound credentials – the risk is off-set by expulsion from VO membership – although traceability is lost, the actual entity remains the same and therefore authentication is correct unless the account/credential is compromised.*

- Do you expect that user's role attributes (e.g. eduPersonAffiliation="faculty") value is updated as an individual departs? How promptly? *Yes, although a status attribute like ePAffiliation would likely only be used in a binary decision (its value of "member" being absent or present being sufficient). For LoA-expressing attributes (ePAssurance on a per-account basis), if these were available, its value should be updated promptly.*

4.2 Quality/provenance of user data

In larger universities the IdP/IdM gathers users' attributes from several registries (payroll system, CRIS (current research information system), student registry) with varying data quality. Some attributes can even be self-asserted by the user him/herself.

- Is it important for you to know the quality/provenance of the user data on the attribute level? What attributes? On what level of granularity? *Quality of the data is important, and any attributes provided must be reasonably reliable. Contact attributes: email and phone number. Self-asserted values are still good enough if these are also used by the organsiation itself – no selfasserted values should be used for things like self-service credential reset.*

4.3 Population and release of attributes

- What are the key attributes Home Organisations should populate for their end users and release to your SP?

  For access to the full EGI services, the attributes needed are

  Required: { givenName and sn (Surname) } or {commonName}
  eduPersonUniqueID or a non-re-assigned eduPersonPrincipalName (not eduPersonTargetedID because it can be different for different SPs)

  Requested: mail (address), schacHomeOrgansiation (name of organisation)

  All other attributes are collected either by the Community/VO, or by the LToS User Management Portal. If reliable contact details (telephone, mail, schacHomeOrg) would be available, these would be used by the LToS UMP and by portals and not likely be re-verified. Any long-term structured VO (LCG, ESR, &c) would re-collect these data anyway.

5 Questions on audits

- Is it enough for you that a Home Organisation self-asserts that it complies with a certain LoA level? *Yes, as long as it is clear what is self-asserted, and those specifications are public. It is desirable if (high-level) operational and policy practices for the organisation are public and correspond to known (community) policy levels. There must be external requirements that are specific and comprehensive enough.*
- Should some external body have some enforcement rights (e.g. Home identity federation can remove "compliant" tag from the Home Organisation if there are doubts that a Home Organisation fails its LoA level)? *Tags endorsed by the federation should reflect the best knowledge the federation can reasonably obtain about compliance. Negative self-evaluation should lead to exclusion.*

- Are internal periodic self-assessments needed? Should these be reviewed (or open to review) by e.g. the Home identity federation or federation peers? *Self-assessments are today required, as well as their peer review (including the possibility of reviewers from major relying parties like EGI in the review process). The frequency of the self-assessments and their peer review are open to discussion (yearly is desirable, in practice EGI is accommodating up to a 3-year peer reviewed assessment cycle).*

- Are internal audits needed where the auditors are from an independent organization unit? *EGI does not insist on internal independent audits as long as the policies and practices are public and the results of self-assessments made available to qualified EGI peers.*

- Are external audits needed? Are you willing to share their costs? *EGI does not insist on external audits, under the same conditions as for internal ones. It has no current interest in funding external audits.*