

Accounting and Traceability in Multi-Domain Service Provider Environments


Publication Date: 2019-04-06
Authors: Uros Stevanovic;David Groep;David Kelsey;Ian Neilson;Hannah Short

Document Code: AARC2-DNA3.2

Contractual Date: 31-03-2019
Actual Date: 06-04-2019
Grant Agreement No.: 730941
Work Package: WP3 (NA3)
Task Item: TNA3.2 Service-centric policies
Lead Partner: KIT

© GÉANT on behalf of the AARC project.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 730941 (AARC2).

This work is licensed under a Creative Commons Attribution 4.0 International License. 

Abstract

In this report, we address the service-centric policies that are applicable to the Blueprint Architecture (BPA) model proposed by AARC, how communities and generic e-Infrastructures can apply the SCI policy framework to their collective service operations, and in which way this also supports the exchange of accounting and traceability information.

The report is best appreciated in conjunction with the AARC policy guidelines and informational documents, specifically G042, G040, G021, the WISE SCI framework, and the AARC Policy Development Kit.



Table of Contents

Table of Contents.....	2
1. Introduction.....	3
2. Multi-Domain Service Provider Traceability.....	5
2.1. Attribute release and relation to the R&E federations.....	6
2.2. Privacy considerations.....	7
3. The Policy Development Kit.....	9
4. Security for Collaborating Infra-structures, and the WISE Community.....	11
4.1. SCI assessment methodology.....	12
5. Summary and conclusions.....	14
References.....	16

1. Introduction

Collaborating service providers are the mainstay of interoperating infrastructures. Managed ‘behind the proxy’ (also considered ‘southbound’ of the proxy), their characteristics are made opaque to identity federations and – especially for generic e-Infrastructures to a large extent – also to the community management services in the Blueprint Architecture (BPA). Yet their collective behaviour is key to enabling trust in the collective service providers, for example, when personal data like name and identifier attributes are provided to them.

At the same time, a significant amount of trust between the service providers within each infrastructure is needed to enable them to provide *collective* services: accounting data needs to be collected and aggregated, the service providers need to assess the risks in managing the personal data they collect as a result of offering their services, and they need to determine the base level policies which they will require their communities (other infrastructures and users) to meet or exceed.

Collectively, the set of (baseline) policies should be sufficient for the majority of service provider consortia – as well as preferably for all generic e-Infrastructures – to trust communities without having to present additional terms and conditions to each individual user. And despite the potential heterogeneity of the service providers, the requirements for accounting and traceability across the service providers should be aligned (so that collectively, and with cooperation of the proxy service itself, enough trust can be established between all parties).

To address this issue, a suite of *Service-centric policies* has been developed and their applicability to common (e-) Infrastructure use cases assessed:

- Are service providers able to collect and share the core set of user and community information (name, email, institutional relationship, and a single common identifier) between themselves in compliance with the pertinent regulatory framework, including GDPR (this has been discussed in AARC-G042, published mid-2018)?
- Can community information and attributes managed within the BPA proxy and the membership management service be shared with service providers in line with data minimisation principles of GDPR?
- Can a suite of template policies be defined that may – after a (usually light-weight) risk check by communities, infrastructures, of service provider consortia – be used to get a complete and gap-less policy suite for communities deploying BPA-compatible models?
- Can the resulting new policy frameworks, as well as the existing policy suites of research- and e-Infrastructures, be compared and ‘mapped’ to determine adequacy for interoperation? And which method of evaluation (auditor-based checks against single-domain standards or a peer-review of self-assessments based on inherently-federated policy models for collaborating infrastructures) is most appropriate to build trust?

On any of these topics, specific AARC guidelines and information documents have already been released. The aim of early release is both to provide timely input to collaborative



research organisations and infrastructures in deploying BPA architectures supported by interoperable policies, as well as to obtain feedback from actual policy implementations on the AARC guidelines themselves. In this report, we discuss the method by which this information was constructed, and describe how the compatibility assessment might be performed in the future.

This report should be read in conjunction with the applicable AARC guidelines and informational documents that are directly applicable to service provider collaborations:

- *AARC-G042* Data Protection Impact Assessment – an initial guide for communities
- *AARC-G040* Preliminary Policy Recommendations for the LS AAI (application to R&S and CoCo)
- *AARC-G021* Exchange of specific assurance information between Infrastructures
- *The AARC Policy Development Kit*, <https://aarc-project.eu/policies/policy-development-kit/>

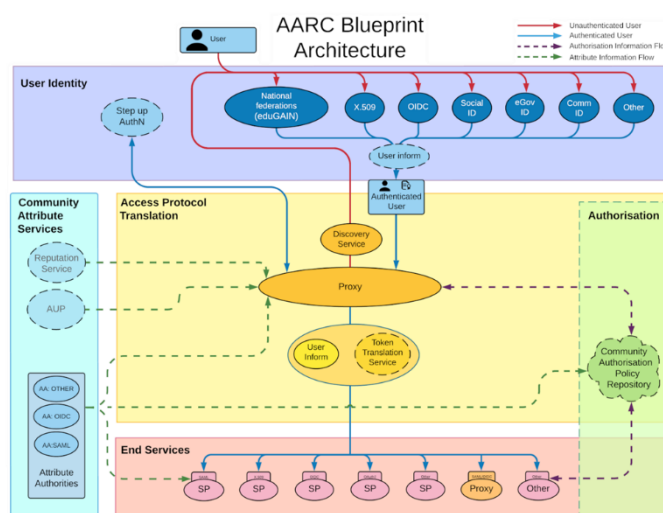
The work is part of the larger set of policies that includes also researcher-centric policies – discussed in the “Recommendations for e-Researcher-Centric Policies and Assurance” and its guidelines, as well as guidelines on operational security capabilities

2. Multi-Domain Service Provider Traceability

Collaborative compute and infrastructures for research, be they generic and serving multiple research communities, or domain specific to a (cluster) of research communities, by necessity require cooperative and collective actions by their service providers. These service providers have to expose collective properties towards the users and the community, and thus have to process information (personal data, attributes, authorization policies) both originating in the community as well as in peer infrastructures and service providers with whom they cooperate in order to provide the collective service. To do this effectively as well as in a way compatible with the regulatory environment in which they operate (of which the EU's General Data Protection Regulation, GDPR, is the most obvious one), they must have a coherent set of policies that allows them to establish mutual trust and permit the exchange of data to ensure access can be managed and traceability provided.

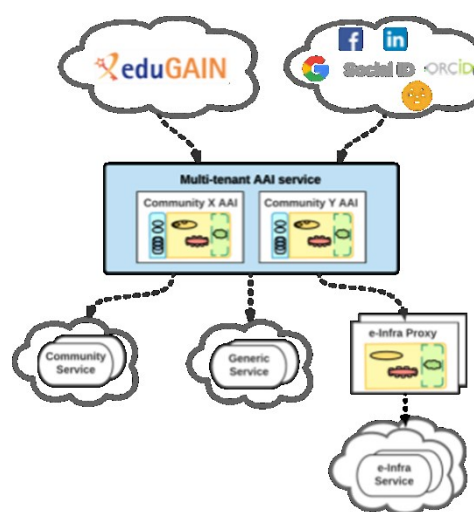
The most visible characteristics of the data exchange in a service-provider context are accounting and traceability. Both are most often associated with respectively billing (or similar management of allocations) and with operational security and incident response. Yet both also have a role in ensuring the integrity of the research data life cycle, in allowing service providers to demonstrably meet resource pledges, and provide transparency to both end-users and the communities to which they belong.

The AARC community has provided a set of guidelines and documents about the technical and organisational features within the Federated Identity Management (FIM) in the form of the AARC Blueprint Architecture (BPA) [AARC-G012, Figure 1]. Most visible aspect is the introduction of the Service Provider to Identity Provider proxy component (SP-IdP-Proxy). The proxy facilitates access of users coming from the Research and Education (R&E) identity federations to both Research Infrastructures (RI) and e-Infrastructures (EI) resources.



The proxy provides the ability to mitigate attribute deficiency in the information provided by identity sources in R&E federations, and it may further specialize authorization decisions, enable access across different technologies, token translation, linking accounts, and in general, expedite the interoperability between R&E federations and RI/EI infrastructures [FIM4Rv2]. And, as illustrated in the BPA graphic, all service providers within the constituency of the IdP-SP proxy will share the same source of attributes (the proxy), and may share additional properties such as the community authorisation policy and access to attribute authorities operated by the proxy.

Offering collective and coherent services by SPs from multiple administrative domains is not inherently limited to direct connections to a single proxy. The BPA model allows and supports cascading of multiple proxies where a group of service providers may be bound together more tightly (e.g. in a generic e-Infrastructure) and this ‘service provider collective’ is subsequently connected as a single service provider to another proxy instance. This is particularly appropriate for a *Community AAI*, which is “responsible for dealing with the complexity of using different identity providers with the required community services” [AARC-G045]. Additionally, the *Community AAI* enables the addition of the necessary attributes to facilitate proper authorization decisions made by service providers. Community services are not the only one that can be “connected” to the community AAI, but also generic services (e.g. RCauth.eu). For the generic e-Infrastructure ‘service provider collectives’ behind an e-Infra proxy, the complexity is abstracted at the proxy level, with the e-Infrastructure proxies potentially connected to multiple *Community AAIs*, yet allowing the e-Infrastructure services to be always connected only to a single, dedicated e-Infrastructure proxy.



With regards to the responsibility for (personal) data managed in the proxy (be it identifier state or attributes regarding group and role membership in the community), it should be kept in mind that a single organisation may choose to offer both e-Infra and community proxies ‘as a service’, and can do that either in a role of thin service provider under contract (as a processor) or as a managed service that they themselves support for the benefit of multiple collaborations (by defining means and purpose of the processing, and thus acting as controllers in the sense of GDPR).

2.1. Attribute release and relation to the R&E federations

The BPA in itself does not constitute the move towards a multi-domain environment, as also in eduGAIN the typical model of access is multi-lateral: users are provided access to service providers by authenticating at their home organisations, or Identity Providers (IdPs). The service providers (SPs) and IdPs typically belong to different administrative domains. However, in the conventional model access to attributes (including the identifiers necessary for collective access control to SPs) is regulated by the policies of R&E federations – whose policy frameworks emphasise regulating privacy, security, and in many cases limitation of attribute release. Policy initiatives such as REFEDS’ Research and Scholarship “implicit attribute release” model [R&S] and even the GEANT Data Protection Code of Conduct work [DPCOCO] have not seen adoption rates sufficient to satisfy the research and collaboration use cases, and other policy frameworks address different use cases (such as Sirtfi [SIRTFI] focussing on incident response).

With the introduction of proxies in the BPA, the model could appear to get slightly more complicated, since not all the SPs “behind” the proxy are necessarily under the same

administrative control as the proxy itself (and at the very least, are, as a matter of practice, “hidden” from IdPs). In recognition of potential trust issues of this model, the Sntctfi [SNCTFI] framework was developed to address the issue of “transitive trust”. Compliance with Sntctfi can ensure that all SPs “behind” a proxy follow the necessary provisions outlined in the Sntctfi framework, so that the proxy can assert all the needed “trust” information on their behalf (and therefore cater for its Infrastructure as a whole). Towards the R&E federations and IdPs, the proxy can then assert the requisite “R&S” and DPCoCo compliance, stimulating IdPs to release at least a basic set of attributes.

2.2. Privacy considerations

In the context of the European GDPR, personal data is “any information relating to an identified or identifiable natural person” (Art 4(1)) [GDPR]. When using federated identity management (FIM), both authentication and authorisation inevitably include personal data and its processing, where processing is defined as “any operation or set of operations which is performed on personal data” (Art 4(2)). This therefore means that we must consider current privacy regulations (i.e. GDPR) when talking about FIM.

In the guidance document on Data Protection Impact Assessment (DPIA) for communities [AARC-G042], privacy risks were considered in the context of a single proxy. In such a scenario, processing of personal data is already satisfying the key GDPR provisions for data processing, and the underlying risks are low when considering personal data emanating from the access to and use of the Infrastructure through federated identity means. Thus, in absence of aggravating conditions, this processing does not warrant conducting a DPIA as defined in the GDPR. The personal data considered in such a scenario are R&S attribute bundle [R&S] and authorisation information (such as group information, entitlements, and similar). Personal data such as email and personal name are considered Common Personal Data [CNIL-MAN], and as such do not require special consideration. The same can be said for other related personal data, or attributes, such as group information or entitlements. Such data are used to convey rights and roles the user may have in accessing or using a service, and therefore are still considered *common*. Naturally, such data are still personal, and should be treated with proper consideration, in line with previous guidelines and policy frameworks (such as Sirtfi, Sntctfi, DPCoCo).

One of the main provisions of the GDPR is ‘data protection by design and by default’, where it states - in Art. 25 - that “appropriate technical and organisational measures” need to be implemented. This should be done, however, taking into account “state of the art, cost of the implementation, nature, scope, context and purposes of processing as well as the risks”. One of the principles to achieve this is data minimisation, which is specifically mentioned in the GDPR. Article 5(c) says that that personal data shall be “adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed”. Also, Recital 39(7) states the same.

With the cascading of proxies in the Community AAI scenario, additional questions about data minimisation of attributes such as group, entitlements, roles, or other authorization attributes may arise. The underlying question of identification, i.e. authentication attributes that typically stem from IdPs were already considered in the guidance given in AARC-G042 [AARC-G042]. As mentioned, R&S attributes, in conjunction with proper policy frameworks

addressing security and data privacy and considering the privacy enhancing nature of using FIM, are already presenting only minimal risks, where additional minimisation of attributes would produce no meaningful benefits for the users, but will significantly hinder or block the operation of SPs in providing the collective coherent service that the user required from them in the first place.

Bearing in mind the purpose of the SP collective and the role of the IdP-SP-proxy in the BPA, one may assess the relevance and purpose of the processing of the data, both that coming from the IdPs as well as from community AAls and 'upstream' proxies, and forwarding that data towards services behind an (e-Infrastructure) proxy. The guiding principle for data minimization is not '*to process least amount of data conceivable*', but to limit processing to what is *relevant and necessary for the purpose*. In the multiple proxy scenario, the tendency in the approach is to release only data needed by an end service in order to for the user to use it. For example, for a storage service that requires entitlements for an authorisation decision, the proposition might be that such service would only need this entitlement information (as, e.g., defined in [AARC-G002]), in addition to authentication information. This approach is not what data minimisation mandates. If we use a definition of "accessing a service" from the new Data Protection Code of Conduct [DPCOCO], where "access" covers, among other things:

- **Authorisation** - groups, roles, entitlements, affiliation, all these information may be used for the decision
- **Identification** - typically information coming from the users' Home Organisation (HO)
- **Researcher unambiguity** - ensuring that the researcher is known (e.g. properly assigning the contribution)
- **Information Security** - safeguarding integrity, confidentiality, and availability of the service, which may include monitoring
- **Accounting and billing** - processing personal data to properly allocate resources to users or communities, logging, and similar (and 'billing' should be understood to include assessing resource usage against an allocation or pledge previously made)

As we can see, the information necessary in order to provide users "access to a service" may be comprehensive, and data minimisation principle does not hinder processing of such information. For example, the information that is necessary for accounting and billing may not be the same that is necessary for accessing an end service. The same can be said for researcher unambiguity (or uniqueness), where enough information, and of proper "quality", needs to be processed to ensure such functionality. The complexity of the multi-proxy scenario may require exhaustive information about the user. Due to the frameworks described by previous documents (AARC-G042 [AARC-G042], AARC-G040 [AARC-G040], and Policy Development Kit [POLICY-KIT]), and the frameworks in [SNCTFI], [SIRTFI], and [DPCOCO], the risks for the users in the multi-proxy environment are not increased in relation to the scenario described in [AARC-G042].

3. The Policy Development Kit

Accessing, using, and operating services for research in today's world, as a rule, is inherently distributed. Users are expecting to access resources not located in their Home Organization. In this complex environment, the question of trust for both users and resource providers, or Infrastructures, becomes paramount.

In order to regulate and facilitate this trust, a set of policies is necessary. These policies, which are essentially a set of documents, outline the operation and operational measures undertaken by the Infrastructure to properly provide services. While operating Infrastructure, arrangements need to be made for Data Protection, Membership Management, and Security Incident Response. The policies that are outlined in the Policy Development Kit (PDK) focus to take trust, assurance, and governance aspects into account and to provide a comprehensible set of documents to be adopted by relevant parties.

Policies are essential for operating the Infrastructure. They set the expectations and duties for the participants in the Infrastructure, from the management to the researchers, i.e. users, themselves. Conversely, a violation of the policy may be interpreted as a security incident, and may cause, and give grounds for, investigation to protect the Infrastructure. Policy decisions may or may not be enforced on a technical level; Infrastructure itself will have to define the permitted usage of their resources through a combination of technology and documentation.

The work is based on “A Trust Framework for Security Collaboration among Infrastructures” (SCI) [SCI], and more specifically “Scalable Negotiator for a Community Trust Framework in Federated Infrastructures” (Snctfi) [SNCTFI] framework. The target audience of these policies is the personnel responsible for the management, operations, and security of the Infrastructure. The policies outlined in the PDK are relying on additional policy frameworks, since they are introducing further necessary concepts. The policy frameworks are not policies themselves, i.e. they provide a conceptual structure within which actual policies are defined.

The policy development kit was developed with the explicit intention to be a set of living documents, whose contents can be taken (and adapted) by communities and generic e-Infrastructures. They support service providers in establishing coherency in a multi-domain environment, both by adoption in the infrastructure itself as well as by research and collaborative communities to ease their access to generic e-Infrastructure service providers from other domains. Providing an implementation model for Snctfi, the Scalable Negotiator for a Community Trust Framework in Federated Infrastructures, the PDK can ensure that the IdP-SP-proxy end-point exposed to an identity federation is capable of representing all the internal services with regard to their adoption of policies. The key policies needed for compliance with the framework are:

- Policies to stipulate requirements for Data Protection and Privacy
- Policies to regulate the behaviour of the management of collections of users
- Policies to coordinate the implementation of operational security practices and incident response

Templates provided by the PDK for these areas may be further specialised by the communities and service providers depending on, e.g. risk assessments or pre-existing trust within a particular community.

Given that the PDK is a living document, we refer for further information to the online resources at:

<https://aarc-project.eu/policies/policy-development-kit/>

To support the development kit, a promotional video and a *Moodle* training are also available (supported by the AARC2 NA2 activity). An impression of the PDK web site is shown below:

Get Started with Policies

A **Moodle course** is available to learn more about Policies for the AARC Blueprint Architecture and videos from this course are also available on the **AARC playlist** on YouTube GÉANTtv.

A **PDK promo video** is also available to share.

Supporting documents are available below for download.

Download Material

Show 100 entries Search:

Document	Who should complete the template?	Audience	Description	Link
Top Level Infrastructure Policy	Infrastructure Management	All Infrastructure Participants (abides by)	This policy template defines the roles of actors in the Research Infrastructure and binds the policy set together	Google Doc
Incident Response Procedure	Infrastructure Management & Security Contact	Infrastructure Security Contact, Services (abides by)	This template procedure provides a step-by-step breakdown of actions to take following a security incident.	Google Doc
Membership Management Policy	Infrastructure Management	Research Community (abides by)	This policy template defines how Research Communities should manage their members, including registration and expiration.	Google Doc
Acceptable Authentication Assurance	Infrastructure Management	Research Community, Services (abide by)	This is a placeholder for the Infrastructure to determine rules for the acceptable assurance profiles of user credentials.	Google Doc
Risk Assessment	Infrastructure Management, Services & Security Contact	Infrastructure Management (completes)	This table can be used as a starting point for identifying whether a full Data Protection Impact Assessment is required.	Google Doc
Policy on the Processing of Personal Data	Infrastructure Management & Data Protection Contact	Research Community, Services (abide by)	This document defines the obligations on Infrastructure Participants when processing personal data.	Google Doc
Privacy Policy	Infrastructure Management (for general	Users (view)	This can be used to document the data collected and processed by the	Google

4. Security for Collaborating Infrastructures, and the WISE Community

The WISE, Wise Information Security for Collaborating e-Infrastructures, community was born as the result of a workshop in October 2015, which was jointly organised by the GÉANT group SIG-ISM (Special Interest Group on Information Security Management) and SCI, the “Security for Collaboration among Infrastructures” group of staff from several large-scale distributed computing infrastructures [WISE].

As research infrastructures become increasingly linked, the need for coherent, common policies has been highlighted. Users of these infrastructures expect a unified policy landscape that limits the need for them to accept or view duplicate policy documents. Likewise, holding infrastructure service providers to a single policy set facilitates best practices and limits the need for edge cases and exceptions. To facilitate this, strong collaboration towards the development of joint policies is necessary, and hence the need for additional work on “Security for Collaborating Infrastructures” and also on the AARC2 policy development kit.

The Security for Collaborating Infrastructures (SCI) working group is a collaborative activity within the Wise Information Security for e-Infrastructures (WISE) trust community [SCI]. The aim of the SCI trust framework is to enable interoperation of collaborating Infrastructures in managing cross-infrastructure operational security risks. It also builds trust between Infrastructures by adopting policy standards for collaboration. The SCI framework is focusing on incident containment, ascertaining the causes of incidents, identifying affected parties, addressing data protection and risk management and understanding measures required to prevent an incident from reoccurring. In essence, it strives to “enable interoperation of collaborating infrastructures for the purpose of managing cross-Infrastructure operational security risks”.

The SCI framework has gone through two iterations, the last being released in 2017 [SCIv2]. All members of the AARC2 NA3 policy and best practices team were active participants of WISE and the SCI working group and indeed all are authors of the SCIv2 paper.

We should also note that two important trust frameworks based on SCI version 1, namely Sirtfi, addressing incident response in the federated identity world, and Snctfi, trust framework for services behind a BPA Proxy, have been co-developed and taken forward by teams including strong participation of AARC, AARC2 and all members of the NA3 policy team.

In June 2017 at the TNC17 conference in Linz, to coincide with the publication of SCI version 2, endorsement of SCI and WISE from supporting infrastructures was sought, resulting in signed statements from the following infrastructures: EGI, EUDAT, GÉANT, GridPP, HBP, MYREN, PRACE, SURF, WLCG, and XSEDE. Each of these infrastructures “welcomes the development of an information security community for the Infrastructures and underlines that the present activities by the research and e-Infrastructures should be continued and reinforced”.

The process of updating the SCI framework continues, to reflect changes in technology, culture, and to improve its relevance. As such, it provides inputs in following areas:

- **Operational Security (OS)** - Specifies criteria on how to properly establish and manage risks when operating an Infrastructure.
- **Incident Response (IR)** - Necessary procedures for proper behaviour in responding to security incidents.
- **Traceability (TR)** - Ensuring sufficient information is kept in logs to address security incidents (e.g. questions regarding who, when, how)
- **Participant Responsibilities (PR)** - Rules must be defined and enforced addressing the behaviour of Individual users, Collection of users (e.g. Communities), and Service providers.
- **Data Protection (DP)** - With the recent legal development addressing data protection is a front and centre issue, and therefore rules must exist to properly process personal data

4.1. SCI assessment methodology

Contrary to conventional information security assessment frameworks, the SCI model explicitly recognises the distribution of responsibility for information security management across multiple domains of authority and control. Thus, the structure of SCI is intentionally different from e.g. ISO 27001 and similar organisation-centric methodologies. This is also apparent from (preliminary, since they were based on SCI version 1) policy mappings that show that areas like asset management, systems management, and acquisition are deliberately not part of SCI. On the other hand, collaborative aspects such as responsibility for actions, common aims and purpose, and emphasis on communication are strengthened in SCI.

This emphasis on collaborative elements and communication are key enabling elements to support an SCI assessment methodology that leverages a mechanism common to research and research collaborations: employing transparency and peer-review based on 'open-book' self-assessments. Infrastructures that collaborate in service offerings to communities may use this mechanism to support the decision to exchange information, personal data, and operational security information based on trust established through such peer-assessment. Although potentially not applicable to ab-initio relationships between organisations (as has been argued in a different context for assurance frameworks in AARC-I050, one may have 'lost the grains of rice' that underlie so much of the shared trust in global research collaboration), the peer-reviewed self-assessment provides a scalable way to enable interoperation between the research infrastructures at the 'blueprint proxy' level.

In support of the SCI assessment framework, the SCI version 2 paper addresses the assessment of Infrastructure maturity against satisfying the requirements of the framework as follows:



“To evaluate the extent to which the requirements described in the SCI document are met, it is recommend that each Infrastructure assess the maturity of its implementation of each function or feature according to the following levels:

- Level 0: Not implemented for critical services;
- Level 1: Implemented for all critical services, but not documented;
- Level 2: Implemented and documented for all critical services;
- Level 3: Implemented, documented, and reviewed by a collaborating Infrastructure or by an independent external body;
- “Justifiable exclusion”: In the unlikely case that the function or feature is not relevant for the infrastructure.

In the interest of promoting trust, Infrastructures should make their maturity assessments available to collaborating Infrastructures. The documentation required for Levels 2 and 3 should either be publicly available or made available on request by a collaborating Infrastructure.”

In order to make the SCI framework more usable, the SCI working group with the support of AARC NA3 has prepared an initial assessment methodology, to be used to assess Infrastructure compliance with the framework. It takes into account whether and to which extent (using the levels described above) the requirements are fulfilled.

As a major contribution to this assessment work, the AARC2 NA3 policy and best practices team has produced a draft assessment spreadsheet [SCIV2-DOC], to be used by Infrastructures in self-assessment or for use by peer-review bodies in the assessment of others. This has been submitted to the WISE SCI working group for comment and for testing. The SCI group will in future be responsible for the maintenance and sustainability of this method of self-assessment and potential audit.

A snapshot of an excerpt of the assessment sheet is shown below. The latest version of the assessment model can be retrieved from <https://wiki.geant.org/display/WISE/SCIV2-WG+documents>.

	A	B	C	D	E	F	G	H	I	
1	Infrastructure Name:			<insert name>						
2	Prepared By:			<insert name>				On Date:	<insert date>	
3	Reviewed By:			<insert name>				On Date:	<insert date>	
4										
5	Operational Security [OS]			Maturity			Evidence	Version Number	Document Date	Document Page
6			Value		Σ		(Document Name and/or URL)			
7										
8	OS1 - Security Person/Team									
9	OS2 - Risk Management Process									
10	OS3 - Security Plan (architecture, policies, controls)				2.0					
11	OS3.1 - Authentication		3							
12	OS3.2 - Dynamic Response		1							
13	OS3.3 - Access Control									
14	OS3.4 - Physical and Network Security									
15	OS3.5 - Risk Mitigation									
16	OS3.6 - Confidentiality									
17	OS3.7 - Integrity and Availability	Q	1	1.0						
18	OS3.8 - Disaster Recovery									
19	OS3.9 - Compliance Mechanisms									
20	OS4 - Security Patching		1	1.0						
21	OS4.1 - Patching Process									
22	OS4.2 - Patching Records and Communication									
23	OS5 - Vulnerability Mgmt		1	0.7						
24	OS5.1 - Vulnerability Process									

5. Summary and conclusions

Distributed IT infrastructures for research that interoperate have to share a common basis for trust that enables them to provide collective service to collaborative research, and trust in any large-scale form of organisation relies on (documented) policies and the means to assess adherence to stated practices. Specifically in the context of security and data protection in infrastructures that rely on federated identity management, the security for collaboration amongst infrastructures (SCI) framework lists the areas in which an infrastructure or research community should define policies – and be prepared to be transparent in their implementation towards their peers.

The SCI framework addresses both policies naturally targeting service providers (operational security and incident response, traceability, responsibilities of communities - collections of users - and service providers, and data protection), as well as complementary areas targeting researchers (individual responsibilities, assurance considerations by communities).

In this report, we address the service-centric policies that are applicable to the Blueprint Architecture (BPA) model proposed by AARC, and how communities and generic e-Infrastructures can apply the SCI policy framework to their collective service operations.

The AARC Guidelines and informational documents developed in the service-centric policy activities, in conjunction with the application to the *community-first* AAI ‘cascading’ BPA model described, taken together provide a gap-less policy framework meeting the Snctfi requirements. This framework allows any BPA proxy both to assert towards the R&E federations REFEDS *Data Protection Code of Conduct* compliance as well as Research and Scholarship for those infrastructures that meet the purpose limitations described therein. At the same time, the policy development kit and guidelines provide the basis for the exchange of accounting and traceability information that is needed to provide *access to services* and share the data necessary for collective services offered by providers from different organisational domains.

In order to promote adoption of the framework and ease its implementation and assessment, the elements of service-centric policy are made available as individual AARC Guidelines and informational document. The complete picture of the work on service-centric policy therefore includes as well the following ancillary documents:

- *AARC-G042* Data Protection Impact Assessment – an initial guide for communities
- *AARC-G040* Policy Recommendations for the LS AAI (application to R&S and CoCo)
- *AARC-G021* Exchange of specific assurance information between Infrastructures
- *The AARC Policy Development Kit*, <https://aarc-project.eu/policies/policy-development-kit/>

The list above also exemplifies the need at times for guidance that brings together policy recommendations for a specific community and purpose. Since policy is generally perceived to be complex, providing a specific instantiation of the generic recommendations is useful not only for the infrastructure to which the specific recommendation is targeted, but more generally to serve as a reference for other communities as to how (elements of) the policy development kit may be applied.



The service-centric policies that have been furthered by the AARC and AARC2 projects have all been developed also in the context of existing sustainable groups and (international) collaborations, such as the WISE community (in the SCI working group), the Interoperable Global Trust Federation (for the peer-reviewed assessment methodology), and in joint efforts with the policy groups from the (e-)Infrastructures including EGI, EUDAT, GÉANT, PRACE, and XSEDE (in the USA). The outputs produced by AARC(2) will thus continue to be developed in these forums, so that the applicability to the continuously changing infrastructure landscape in Europe and in the world is ensured.

References

- AARC** <https://aarc-project.eu>
AARC-G002 <https://aarc-project.eu/guidelines/aarc-g002/>
AARC-G042 <https://aarc-project.eu/guidelines/aarc-g042/>
AARC-G045 <https://aarc-project.eu/guidelines/aarc-g045/>
BPA <https://aarc-project.eu/architecture/>
CNIL-MAN <https://www.cnil.fr/en/cnil-publishes-update-its-pia-guides>
ComAAI <https://aarc-project.eu/guidelines/aarc-g045/>
DPCOCO *GEANT Data Protection Code of Conduct version 2 (draft)*
<https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home>
eduGAIN <https://edugain.org/>
FIM4Rv2 <https://doi.org/10.5281/zenodo.1307551>
GDPR <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
POLICY-KIT <https://aarc-project.eu/policies/policy-development-kit/>
R&S <https://refeds.org/category/research-and-scholarship>
SCI <https://wise-community.org/sci/>
SCIV2 <https://wise-community.org/wp-content/uploads/2017/05/WISE-SCI-V2.0.pdf>
SCIV2-DOC <https://wiki.geant.org/display/WISE/SCIV2-WG+documents>
SIRTFI <https://refeds.org/sirtfi>
SNCTFI <https://www.igtf.net/snctfi/>
WISE “*WISE Information Security for Collaborating e- Infrastructures*”, Hannah Short et al., paper to be published in the proceedings of the Computing in High Energy Physics Conference, Sofia, Bulgaria, July 2018.