17-04-2019

Deliverable DNA3.3
Accounting and Traceability in Multi-Domain Service Provider Environments

Deliverable DNA3.3

Contractual Date:	31-03-2019
Actual Date:	17-04-2019
Grant Agreement No.:	730941
Work Package	WP3 (NA3)>
Task Item:	TNA3.2 Service-centric policies
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	KIT
Document ID:	AARC2-DNA3.2
Authors:	Uros Stevanovic (KIT), David Groep (Nikhef), David Kelsey (STFC), Ian Neilson (STFC), Hannah Short (CERN)

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 730941 (AARC2).

Abstract

This report details the service-centric policies that apply 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 how this supports the exchange of accounting and traceability information.
The report is complemented by 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 Figures

Executive Summary

Collaborating service providers are the mainstay of interoperating research infrastructures. Most of the research collaboration that offer services implement an architecture that follows the AARC Blueprint Architecture and therefore makes usage of an IdP-SP proxy. The proxy provides a single-entry point for services but at the same time makes them opaque both to identity federations as well as to the e-infrastructures which only see the proxy. However, the services collective behaviour is key to enabling trust in collective service providers, for example, when personal data like name and identifier attributes is provided to them.

This document presents the work carried out by the AARC Policy team (NA3) to define a suite of service-centric policies to address accounting and traceability requirements and align them across service providers so that collectively and with cooperation of the proxy service itself, enough trust can be established between all parties, namely service providers, research infrastructures, e-infrastructures and eduGAIN.

The major outcome of the NA3 team in the area of service-centric policies, is the Policy Development Kit (PDK) described in section 3. The PDK aims to support research infrastructures in adopting or enhancing a policy set that regulates the operation and use of an Authentication and Authorisation Infrastructure (AAI) in line with the AARC Blueprint Architecture. The policy development kit offers an online training as well as a set of living documents, whose contents can be used (and adapted) by research communities and generic e-Infrastructures. The documents support service providers in establishing coherency in a multi-domain environment; they offer templates to address different aspects and to adopt policy guidelines, such as security incident, data protection risk assessment, privacy policy, incident response procedures, etc.

This document also reports on the contributions of NA3 team to WISE initiative, the 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.

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 to the community management services in the Blueprint Architecture (BPA), especially for generic e-Infrastructures. However, their collective behaviour is key to enabling trust in collective service providers, for example, when personal data like name and identifier attributes is provided to them.

A significant amount of trust is needed between the service providers within each infrastructure to enable them to provide collective services: accounting data needs to be collected and aggregated, the service providers need to assess the risks involved 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 require their communities (other infrastructures and users) to meet.

Collectively, the set of (baseline) policies should be sufficient for the majority of service provider consortia – and preferably for all generic e-Infrastructures – to trust communities without having to present additional terms and conditions to each individual user. Despite the potential diversity of service providers, accounting and traceability requirements across them 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, a suite of service-centric policies has been developed and their applicability to common infrastructure use cases assessed. The policies have been designed to offer an answer, supported by a guideline or document template, to some common questions research infrastructures faces when deploying an AAI that follows the AARC BPA, namely:

- 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, comprehensive 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? Which evaluation method (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 been released. The early release aims to:

- Provide timely input to collaborative research organisations and infrastructures in deploying BPA architectures supported by interoperable policies
- Obtain feedback from actual policy implementations on the AARC guidelines themselves.

This report discusses the method used to construct this information and describes how the compatibility assessment might be performed in the future. The 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-G042].
- AARC-G040 Preliminary Policy Recommendations for the LS AAI (application to R&S and CoCo) [AARC-G040].
- AARC-G021 Exchange of specific assurance information between Infrastructures [AARD-G021].
- The AARC Policy Development Kit [POLICY-KIT].

The work is part of the larger set of policies that also includes researcher-centric policies. These are discussed in Deliverable DNA3.4 *Recommendations for e-Researcher-Centric Policies and Assurance*, which also includes guidelines on operational security capabilities

2 Multi-Domain Service Provider Traceability

Collaborative computing and research infrastructures by necessity require cooperative and collective actions by their service providers, irrespective of whether they are generic, serving multiple research communities, or domain specific to a cluster of research communities. These service providers have to expose collective properties to the users and the community, and thus have to process information (personal data, attributes and authorisation policies) that originates both from the community, and from peer infrastructures and the service providers they cooperate with to provide the collective service. To be able to do this in an effective way that is compatible with their regulatory environment (of which the EU's General Data Protection Regulation (GDPR) is the most obvious one), they need a coherent set of policies that allows them to establish mutual trust and permits the exchange of data, while ensuring access is managed and traceability provided.

The most visible characteristics of data exchange in a service-provider context are accounting and traceability. Both are most often associated with billing (or similar allocation management), operational security and incident response. However, 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 their communities.

The AARC community has provided a set of guidelines and documents detailing the technical and organisational features of federated identity management (FIM) in the form of the AARC BPA [AARC-G012]. The most pertinent aspect of the BPA is the introduction of the Service Provider to Identity Provider proxy component (SP-IdP-Proxy). The proxy gives Research and Education (R&E) identity federation users access to both Research Infrastructures (RI) and e-Infrastructures (EI) resources.

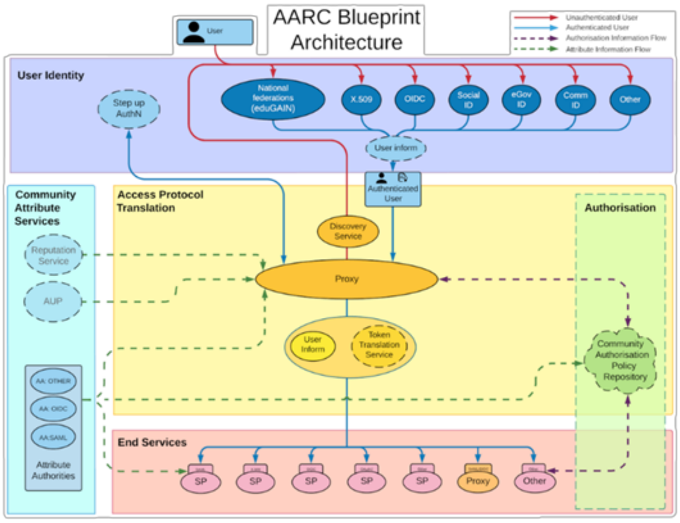

Figure 2.1: AARC Blueprint Architecture

The proxy also provides the ability to mitigate attribute deficiency of information provided by identity sources in R&E federations, and can furthermore make authorisation decisions, enable access across different technologies, translate tokens, link accounts, and, in general, expedite the interoperability between R&E federations and RI/EI infrastructures [FIM4Rv2]. As shown in Figure 2.1, 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 the cascading of multiple proxies, where a group of service providers may be bound together more tightly (e.g. in a generic e-Infrastructure). This 'service provider collective' is subsequently connected as a single service provider to another proxy instance. This is particularly appropriate for a community Authentication and Authorisation Infrastructure (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 attributes to be added that are required to facilitate proper authorisation decisions made by service providers. Aside from community services, generic services (e.g. RCauth.eu) can be 'connected' to the community AAI. 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, but allowing the e-Infrastructure services to be always connected only to a single, dedicated e-Infrastructure proxy.

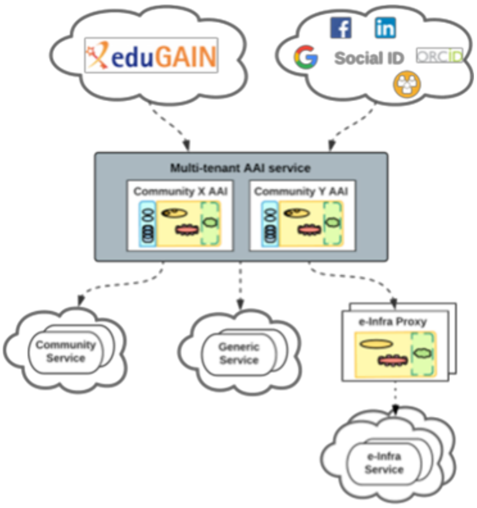

Figure 2.2: AARC BPA Community first view

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 as a thin service provider under contract (a processor) or as a managed service that they themselves support for the benefit of multiple collaborations (by defining the 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 itself does not constitute a move towards a multi-domain environment, as in eduGAIN the typical model of access is multi-lateral: users are provided access to SPs by authenticating at their home organisations, or IdPs. The 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 the regulation of privacy, security, and in many cases limitation of attribute release. Policy initiatives such as the REFEDS Research and Scholarship (R&S) implicit attribute release model [R&S] and even the GÉANT Data Protection Code of Conduct work [DPCOCO] have not seen sufficient adoption rates 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 Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (Snctfi) [SNCTFI] framework was developed to address the issue of transitive trust. Compliance with Snctfi can ensure that all SPs behind a proxy follow the necessary provisions outlined in the Sntcfi 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, inducing 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)) [GDPR]. Therefore, current privacy regulations (i.e. GDPR) must be considered when talking about FIM.

In the guidance document on Data Protection Impact Assessment (DPIA) for communities [AARC-G042], privacy risks are considered in the context of a single proxy. In such a scenario, the processing of personal data already satisfies the key GDPR provisions for data processing, and the underlying risks are low for the personal data required for accessing the infrastructure through federated identity. Therefore, this processing does not require a DPIA to be conducted. The personal data considered in such a scenario comprises R&S attribute bundle [R&S] and authorisation information (such as group information, entitlements and similar). Personal data such as email and names 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 is used to convey rights and roles the user may have in accessing or using a service, and therefore still considered common. Naturally, such data is still personal, and should be treated with proper consideration, in line with existing guidelines and policy frameworks (such as Sirtfi, Snctfi, DPCoCo).

One of the main provisions of the GDPR is data protection by design and by default. Article 25 states that "appropriate technical and organisational measures" need to be implemented [GDPR]. This should be done taking into account "state of the art, cost of the implementation, nature, scope, context and purposes of processing as well as the risks" [GDPR]. One of the principles to achieve this is data minimisation, which is specifically mentioned in the GDPR. Article 5(c) states that that personal data shall be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed" [GDPR]. 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 authorisation 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 would significantly hinder or block the operation of SPs in providing the collective, coherent service the user requires.

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

- **Authorisation** - groups, roles, entitlements and affiliation 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 (billing should be understood to include assessing resource usage against an allocation or pledge previously made).

This demonstrates that the information necessary to give users access to a service may be comprehensive, and the data minimisation principle does not hinder the processing of such information. For example, information necessary for accounting and billing may not be the same as for accessing an end service. The same can be said for researcher unambiguity (or uniqueness), where enough information 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 in previous documents ([AARC-G042], [AARC-G040], [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 research services is inherently distributed. Users expect to access resources that are not located in their home organisation. In this complex environment, the question of trust for both users and resource providers, or Infrastructures, is of paramount importance.

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 the infrastructure, arrangements need to be made covering data protection, membership management and security incident response. The policies that are outlined in the Policy Development Kit (PDK) consider trust, assurance and governance, and provide a comprehensive set of documents to be adopted by relevant parties.

Policies are essential for operating the infrastructure. They set expectations and define the duties of infrastructure users, from management to researchers. A policy violation may be interpreted as a security incident and may trigger an investigation to protect the infrastructure. Policy decisions may or may not be enforced on a technical level; the infrastructure itself has to define the permitted usage of their resources through a combination of technology and documentation.

The work that led to the PDK, is based on “A Trust Framework for Security Collaboration among Infrastructures” [SCI] and, more specifically, “Scalable Negotiator for a Community Trust Framework in Federated Infrastructures” [SNCTFI]. The target audience of these policies is the personnel responsible for the management, operation and security of the infrastructure. The policies outlined in the PDK rely on additional policy frameworks, since they introduce further necessary concepts. The policy frameworks are not policies themselves but provide a conceptual structure within which actual policies are defined.

The policy development kit was developed with the explicit intention to offer an online training as well as a set of living documents, whose contents can be used (and adapted) by communities and generic e-Infrastructures. The documents support service providers in establishing coherency in a multi-domain environment, by adopting them in the infrastructure itself and by research and collaborative communities using them to ease their access to generic e-Infrastructure service providers from other domains. Providing an implementation model for Snctfi, 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 data protection and privacy requirements.
- Policies to regulate 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 customised 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, up-to-date information is available online [POLICY-KIT]. To support the development kit, a promotional video and a Moodle training course are also available (supported by the AARC2 NA2 activity).

4 Security for Collaborating Infrastructures and the WISE Community

The Wise Information Security for Collaborating e-Infrastructures (WISE) community was formed as a result of a workshop in October 2015, which was jointly organised by GÉANT's Special Interest Group on Information Security Management (SIG-ISM) and the Security for Collaboration among Infrastructures (SCI) 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. Therefore, additional work on security for collaborating infrastructures and also on the AARC2 policy development kit is needed.

The SCI working group is a collaborative activity within the 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 focuses on incident containment, ascertaining the causes of incidents, identifying affected parties, addressing data protection and risk management, and understanding the 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 all are authors of the SCIv2 paper.

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

In June 2017n at the TNC17 conference in Linz, to coincide with the publication of SCI version 2, endorsement of SCI and WISE was sought from supporting infrastructures, resulting in a signed statement from EGI, EUDAT, GÉANT, GridPP, HBP, MYREN, PRACE, SURF, WLCG and XSEDE. Each welcomed "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 to the following areas:

- **Operational Security (OS)** - How to establish and manage risks when operating an infrastructure.
- **Incident Response (IR)** - Procedures for responding to security incidents.
- **Traceability (TR)** - How to log relevant information for addressing security incidents (e.g. questions regarding who, when, how).
- **Participant Responsibilities (PR)** - Rules that must be defined and enforced addressing the behaviour of individual users, collections of users (e.g. communities) and service providers.
- **Data Protection (DP)** - Rules that are required to properly process personal data

4.1 SCI Assessment Methodology

Unlike conventional information security assessment frameworks, the SCI model explicitly recognises the distribution of information security management responsibilities 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 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 is key to supporting an SCI assessment methodology that leverages a mechanism common to research and research collaborations: employing transparency and peer-review based on self-assessments. Infrastructures that collaborate in offering services 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-assessments. Although potentially not applicable to an initial relationships between organisations (as has been argued in a different context for assurance frameworks in AARC-I050), 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."

To make the SCI framework more usable, the SCI working group, with the support of AARC NA3, prepared an initial assessment methodology, to be used to assess infrastructure compliance with the framework. It considers whether and to which extent the requirements are fulfilled (using the levels described above).

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 for self-assessment or by peer-review bodies for the assessment of others. This has been submitted to the WISE SCI working group for comment and testing. The SCI group will in future be responsible for the maintenance and sustainability of this method of self-assessment and potential auditing.

Figure 4.1 shows an excerpt of the assessment sheet. The latest version of the assessment model is available online [SCIv2-DOC].

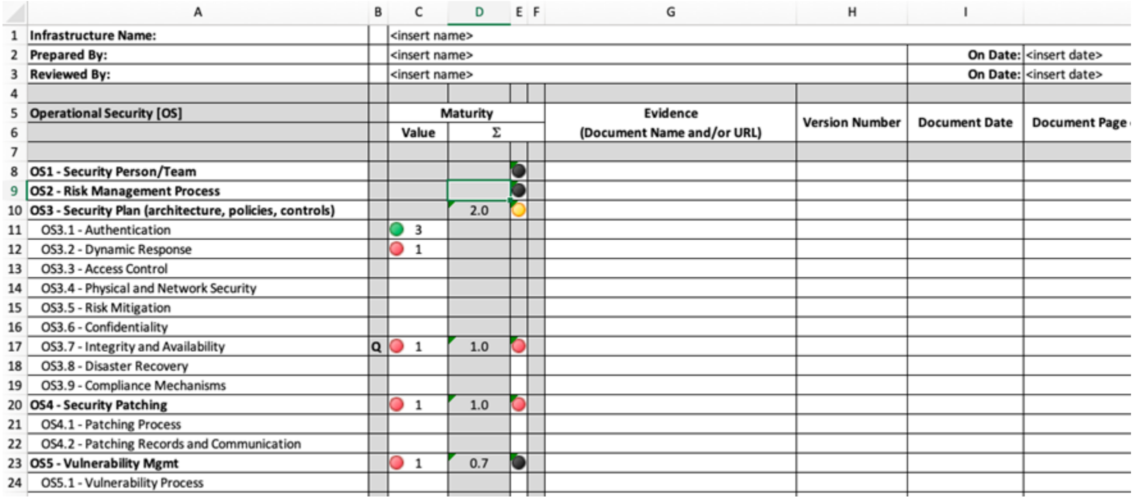

	A	B	C	D	E	F	G	H	I	
1	**Infrastructure Name:**		<insert name>							
2	**Prepared By:**		<insert name>						**On Date:**	<insert date>
3	**Reviewed By:**		<insert name>						**On Date:**	<insert date>
4										
5	**Operational Security [OS]**		**Maturity**				**Evidence**	**Version Number**	**Document 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									

Figure 4.1: Assessment sheet

5 Conclusions

Distributed, inter-operational IT research infrastructures must share a common trust basis that enables them to provide a collective service to collaborative research. Trust in any large-scale form of organisation relies on documented policies and the means to assess adherence to stated practices. In the context of security and data protection in infrastructures that rely on federated identity management, the SCI framework lists the areas where an infrastructure or research community should define policies and be prepared to be transparent in their implementation towards their peers.

The SCI framework addresses policies aimed at service providers (regarding operational security, incident response, traceability, participant responsibilities and data protection) and complementary areas that target researchers (regarding individual responsibilities and assurance considerations by research communities).

The AARC guidelines and informational documents developed by the service-centric policy activities, together with the application of the community-first AAI 'cascading' BPA model described, provide a comprehensive policy framework that meets the Snctfi requirements. Snctfi framework allows any proxy in a BPA-compliant architecture to assert compliance with the REFEDS Data Protection Code of Conduct and Research and Scholarship Entity Category. 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 required by collective services offered by providers from different organisational domains.

To promote the adoption of the framework and make its implementation and assessment easier, the entire elements developed by the NA3 in the area of service-centric policy are made available as individual AARC guidelines and informational documents. The service-centric policy work therefore includes the following ancillary documents:

- AARC-G042 Data Protection Impact Assessment – an initial guide for communities [AARC-G042].
- AARC-G040 Policy Recommendations for the LS AAI (application to R&S and CoCo) [AARC-G040].
- AARC-G021 Exchange of specific assurance information between Infrastructures [AARC-G021].
- The AARC Policy Development Kit [POLICY-KIT].

This list also addresses the need for policy recommendations for a specific community and purpose. Since policy is generally perceived to be complex, providing specific recommendations is not only useful for the infrastructure the specific recommendation targets, but also serves as a reference for other communities as to how the policy development kit may be applied.

The service-centric policies that have been developed by the AARC and AARC2 projects have all been created 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. 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-G002]	https://aarc-project.eu/guidelines/aarc-g002/
[AARC-G012]	https://aarc-project.eu/guidelines/aarc-g012/
[AARC-G021]	https://aarc-project.eu/guidelines/aarc-g021/
[AARC-G040]	https://aarc-project.eu/guidelines/aarc-g040/
[AARC-G042]	https://aarc-project.eu/guidelines/aarc-g042/
[AARC-G045]	https://aarc-project.eu/guidelines/aarc-g045/
[CNIL-MAN]	https://www.cnil.fr/en/cnil-publishes-update-its-pia-guides
[DPCOCO]	https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home
[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

Glossary

AAI	Authentication and Authorisation Infrastructure
BPA	Blueprint Architecture
DPIA	Data Protection Impact Assessment
GDPR	General Data Protection Regulation
EI	e-Infrastructures
EU	European Union
FIM	Federated Identity Management
HO	Home Organisation
IdP	Identity Provider
PDK	Policy Development Kit
R&E	Research and Education
R&S	Research and Scholarship
RI	Research Infrastructures
SCI	Security for Collaboration among Infrastructures
SIG-ISM	Special Interest Group on Information Security Management
SP	Service Provider
WISE	Wise Information Security for Collaborating e-Infrastructures