

Trust framework for proxies and Snctfi research services

Publication Date: 2025-05-30
Authors: David L. Groep (Nikhef, Maastricht University) (ed.), Liam Atherton (UKRI STFC RAL), Peter Bolha (Masaryk University, CESNET), Amber Daniels (AAF), Fahame Emamjome (AAF), Diana Gudu (KIT), Marcus Hardt (KIT), David Kelsey (UKRI STFC RAL), Maarten Kremers (SURF), Mischa Sallé (Nikhef), Hannah Short (CERN), Arnout Terpstra (SURF)

Document Code: **AARC-I082**

DOI: <https://doi.org/10.5281/zenodo.15506826>

Abstract

To provide trust across the layered architecture of the AARC Blueprint (BPA), end-to-end trust across the components for collaboration management, user privacy, identity assurance, and operational security must be provided. This document sets out the overview of trust relationships in the AARC BPA, building on the body of guidelines under development in the AARC community and related coordination bodies for research and education identity federation: REFEDS, IGTF, and WISE. Reviewing the state of the policy landscape and the effectiveness of the first Policy Development Kit (PDK), we propose a new structure for policy organisation based on identified target audiences: external identity sources, the identity, collaboration management, and infrastructure integration components, and site-local integrations and services. Research community governance is discussed as far as it affects authentication and authorisation. Based on deployment experience with the PDK this framework distinguishes more clearly between policies and the processes and procedures that implement such policies. While this Trust Framework provides the structure for the revised Policy Development Kit, it intentionally does not provide the policies and procedures themselves, but identifies the smallest set of distinct guidelines (policies, good practices, procedures) necessary to cover the trust, security, and operational interactions.

Copyright

© Members of the AARC community.
This work is licensed under a Creative Commons Attribution 4.0 License.



Table of Contents

Introduction.....	3
Context and Background Information.....	4
Policy Development Kit.....	7
Analysis of the first generation Policy Development Kit.....	7
Current Composition of the PDK.....	7
Identified Limits and Observations.....	7
Insight from Community Feedback.....	8
Amalgamated terms and definitions for a policy development kit.....	8
Trust Framework Structure of a new PDK.....	10
Guidelines by target audience.....	14
Research Governance.....	14
Users and collaboration purpose.....	15
Identity.....	16
Collaboration Management.....	18
Federated operational security guidance.....	19
Security Operational Baseline and Sirtfi.....	20
Attribute authority and assertion issuer operations.....	20
Data Protection Code of Conduct best practice.....	21
Assurance requirements and acceptable assurance.....	22
Infrastructure Integrations and Service Providers.....	22
Trust qualification through Sncftfi.....	24
Procedural and implementation guidance.....	26
Evolution of the policy development kit.....	28
Acknowledgements.....	28

Introduction

The AARC Blueprint Architecture [G045, G080] introduces components that connect users, external identity providers, identity integrators, collaboration management, service providers, and collections and combinations of these roles. As these components can change or obscure identity and authentication data, it's sometimes difficult to trace back the information they provide to its original source. This brings challenges to trust(ing the user/authentication): which link in the 'chain' asserts which information and how trustworthy are they?

In order to provide trust across the entire chain of AAI components, this document sets out an overview of all available trust-related documents available within AARC. The document starts with a detailed description of the problem space, followed by a (brief) review of earlier work and concludes with guidance on how to structure application of policies and AARC guidelines.

It is important to make the distinction between *policies* and the *processes and procedures* that implement such policies. *Policies* refer to documents where explicit approval by management (at the appropriate level) is advisable. They should therefore be both unambiguous and temporally stable. The *processes* (and the *procedures* that specify activities performed in a process) implementing the policies can have more agility and should be seen as templates, as they depend on specific organisational processes, and will need adapting to changing conditions (like new adversaries in threat scenarios).

In the Policy Development Kit (PDK) version 2, we provide *policies* as AARC Guidelines [AARC-GLS], and *process and procedure* templates as AARC Informational documents. Hence, while the trust framework described in this document provides the *structure* for the next version of the Policy Development Kit, it – intentionally – does not at this stage provide the policies, processes, and procedures themselves.

Context and Background Information

Trust along the entire authentication, authorisation, and attribute chain is essential. This trust must cover:

- Upstream trust
Service providers must trust the attributes released by each AAI component in the AAI layers above them, and the sources behind it.
- Sideways trust
Communities and projects must trust that all components correctly manage user data, group memberships, and account linking.
- Downstream trust
Identity providers and assurance sources must trust that the components handle personal data properly and respect responsibilities related to legal liability.




Figure 1: Mesh of linked authentication sources (top) to service providers (bottom). Collaboration management components and infrastructure integration components are cross-connected to multiple authentication sources. The example depicted here is a mesh of collaborative organisations (COs) of SURF (Dutch collaborative infrastructure) and EGI (European computing service for research), using identity providers from global (eduGAIN) and national (SURFconext) identity federations, connecting to service providers (SPs) in each, linked through two infrastructure integration components, incidentally also from SURF and EGI. More complex composition scenarios are possible [Kremers2023].

In simpler cases - with one amalgamated component (identity and collaboration management level components) or two separate components (collaboration management and infrastructure integration) - these responsibilities are clearly divided. The infrastructure integration components represent their connected service providers, while the collaboration management components handle trust on behalf of the users and identity sources.

The first version of the Policy Development Kit, the more simple structure of the blueprint architecture, consisted of more composite components, called 'proxies' throughout both the architecture as well as in the policy development kit and the trust frameworks, policies, and procedures. In the new Blueprint Architecture, these have been logically separated, and appear as layers and components. We will use these same terms in this trust framework from hereon.

In more complex mesh scenarios, as shown in Figure 1, responsibilities become harder to separate. Other systems, for example, SURFconext [SURFconext] or eduGAIN [eduGAIN], although outside the scope of the AARC BPA, introduce extra layers affecting trust relationships between users and services.

This document aims to define the minimum set of guidelines, including policies, practices, and procedures, needed to establish and manage these trust relationships. Some are already available (e.g., on attribute authority operational security, commonly known as "AAOPS" [AARC-G071]), while others are still needed to fill existing gaps.

For infrastructure integration components, i.e. cross domain capabilities, and hence excluding any Site-local Integration capabilities, trust and security towards service providers should include:

- Ensuring integrity of the component and its attribute services following the "AAOPS" guidelines [AARC-G071].
- Timely alerts about compromised credentials. This is only partially supported by current solutions like Sirtfi v2 [Sirtfi]. Although automated communications regarding compromised accounts were proposed, these have not been widely adopted, and sharing of threat intelligence such as MISP [MISP] and the Security Event Tokens [RFC8417] are not yet widely used for sharing identity events. This requires coordinators in incident response, and addressing of 'supply-chain' wide response capabilities in line with NIS 2.
- Proper handling of privacy and access data, ensuring legal requirements and user agreements are respected as discussed in 'Guidance for Notice Management by Proxies' [AARC-G083].
- Clear communication of standards compliance from infrastructure integration components to upstream identity sources - similar to using Entity Category tags.

For the community 'sideways' trust, requirements include:

- Ensuring collaboration management operations align with the community expectations for access data management ("AAOPS" [AARC-G071]).
- Assurance of authentication quality, following the 'Guideline on the exchange of specific assurance information between Infrastructures' [AARC-G021] and the 'Guidelines for the evaluation and combination of the assurance of external identities' [AARC-G031].
- Participation in multilateral incident response, both upstream and downstream.
- Proper management of user data and consent for downstream sharing, and presenting clear terms and acceptable use policies, following the 'Guidance for Notice Management by Proxies' [AARC-G083]).

For an identity component, i.e., authentication sources including identity providers, assurance providers, including cases where these are combined with a collaboration platform, requirements include:

- Verifying they operate securely and meet the standards required by all identity sources involved: a security baseline and community good practice on access personal data protection.
- Ensuring identity components enforce assurance levels and uniqueness as claimed, or use compensating controls.
- Meeting Sirtfi expectations for prompt incident communication. Sirtfi's intent is to reduce the spread and impact of federated incidents, this becomes more difficult with long chains of AAI intermediary components, or when assertions are validated offline over extended periods, for which guidance is provided in the 'Site requirements for Grid Authentication, Authorization and Accounting' [GFD-I.032].

There are several requirements common to all roles, as also shown in Figure 2, the trust framework structure:

- Improve usability: for users to facilitate the access flow, for service providers, collaboration managers, and AAI operators to lower barriers to entry and usage, for communities to attract and retain users and perform their primary aims
- Reduce complexity and duplication to enhance efficiency and adoption.

The resulting guidelines, the policies and the associated procedure templates put together, form a cohesive framework that form an updated version of the Policy Development Kit (PDK). A subset of reviewable adopted policies and practices can be used by communities and services that source their AAI infrastructure from platform providers, evolving *Snctfi*, the scalable, negotiated community trust framework for federated infrastructure.

Policy Development Kit

Analysis of the first generation Policy Development Kit

The policy and good practice landscape for trusted AAI operations has evolved over the past decade to include a range of guidelines, both in the AARC Policy Development Kit (PDK) as well as in related communities and infrastructures. We list these either as inspiration for improved guidelines in the PDK, or as reference to existing work that can be repurposed for the suite of guidelines in this 'AARC Trust framework for proxies and Snctfi research services'.

Current Composition of the PDK

The first version of the Policy Development Kit [PDKv1] consists of nine documents, each addressing a specific aspect of practice or policy:

- Top-Level Security Policy
- Incident Response Procedure
- Membership Management Policy
- Acceptable Authentication Assurance Policy
- Risk Assessment Guidance
- Processing of Personal Data Policy
- Privacy Policy
- Service Operations Security Policy
- Acceptable Use Policy

Each document is intended to allow for research collaborations to adapt them based on need and governance structure. The current list reflects a mix of policy and procedure, which has led to some confusion and redundancy.

Identified Limits and Observations

A review of the current PDK has highlighted areas for improvement. Reducing redundancy, regional bias, lack of clarity around guidance documents.

There are documents in the PDK with scope or content overlapping with existing guidance and governance. The Risk Assessment largely reiterates existing external IT security risk assessment guidelines. Similarly, the Membership Management Policy content reiterates existing points, allowing for a reduction and refactoring into a much shorter community management policy. The Top Level Policy could be more effectively presented as a concise web page or explanatory document, with a separate document created for the granting of authority to specific roles and functions.

The Personal Data document in the current PDK is heavily influenced by the General Data Protection Regulation [GDPR]. While useful for infrastructures operating within or with the European Economic Area (EEA), this provides a challenge for global federations. For example, within WLCG [WLCG] and EGI [EGI] this extends GDPR aligned practice to non-EEA countries, this approach may not align with local data protection norms for non-EEA adjacent federations.

Some policies in the PDK would benefit from improved clarity, particularly in an operational context. For example, the Service Operations Security Agreement would benefit from accompanying procedural guidance. Conversely the Incident Response Procedure is overly detailed, leading to difficulty in practical application.

Insight from Community Feedback

In the development work done for the PDK version 2, several research communities provided feedback to AARC. These use cases provided insight into how the PDK can be applied across different governance models, operational frameworks and national environments.

The Australian Access Federation [AAF], in collaboration with fellow National Collaborative Research Infrastructure Strategy (NCRIS) funded facilities, coordinated a Trust and Identity (T&I) Policy Working Group (PWG). The T&I PWG reviewed the AARC PDK both in the context of their own policy suites but also within the collaborations they are currently operating or considering to operate.

After the first round of review and consultation, a majority of the AAF PWG members agreed that version 1 of AARC PDK is a valuable baseline suitable for the Australian research landscape. However, the following points were identified:

- Terminology differences between the European Union and Australian context, specifically different use of terminology such as “personal data”, “community”, “infrastructure”, and “policy”. AAF has been working on developing definitions based on the feedback from the community.
- Different governance structure across research facilities in Australia: many of the research facilities and services sit within a parent organisation or span across multiple organisations. Some are Businesses Ltd by Guarantee, others are full business Ltd by shares.
- The PWG also highlighted the need for visualisation of policy sets and how they are related.

Further, through incubator programs, the T&I PWG identified that, possibly due to different funding governance approaches in EU compared to Australia, it is difficult to define the entities that are responsible for governance and maintenance of policy for a research collaboration. Accordingly, the T&I PWG considers developing different implementation approaches for PDK according to the agreed governance structure for each collaboration, a concept used in this new trust framework and scheduled for PDK version 2.

Amalgamated terms and definitions for a policy development kit

As discussed from the feedback above as well as work done in creating the new AARC Blueprint Architecture (BPA) the updated terms are listed below, aligned with their identical meaning in the 2025 Blueprint Architecture.

Term	Updated PDK terminology, augmented with BPA 2025 concepts
Community	A group of Users, organised with a common purpose, and jointly granted access to the Collaboration. It may act as the interface between individual Users and the Collaboration.
Community Manager	A nominated individual responsible for the management of the members (Users) of a Community and/or the Collaboration. In research collaborations this is often a principal investigator (PI) or somebody designated by the PI.
Collaboration	The bounded collection of universities, laboratories, institutions or similar entities, which adhere to the Collaboration Policies. The Collaboration offers research infrastructure to the Community. In earlier versions of the terminology, also 'community' and 'virtual organisation (VO)' have also been used in this sense.
Collaboration Management	The collection of the various boards, committees, groups and/or individuals mandated to oversee and control the Collaboration. <i>Implementation note: where this is a governing body, 'Management' should be customised to the Collaboration. For example, 'the ABC Steering Committee (ABC-SC)'.</i>
Collaboration Policies	The set of policies governing the management, operations and security of the Collaboration as approved by Management.
Infrastructure	The IT hardware, software, networks, data, facilities, processes and any other elements that together are required to develop, test, deliver, monitor, control or support services.
Infrastructure Integration	Infrastructure Integration components connect multiple administrative domains. These "classic" infrastructure integrators represent large infrastructures that are operated in one common way. Examples are e.g. the EGI / WLCG infrastructure.
Site-local Integration	Site-local integration capabilities connect multiple services at one site, i.e. within one administrative domain. They enforce site-specific policies and settings, e.g. by adjusting attributes so that site services can integrate federated users. From the (northbound) AAI and policy perspective, Site-local Integration components are not different from a service.

Trust Framework Structure of a new PDK

The trust framework (and hence the Policy Development Kit (PDK) version 2) identifies five main target audiences, functionally following the BPA 2025 hierarchy and identifying (1) 'Research governance' as a foundational area. (2) 'Users' are (human) end-users who participate in a collaboration, are identified via (3) 'identity', i.e. external identity providers and the identity layer of the BPA, to be granted access by (4) 'collaboration management', to (5) 'infrastructure integration and service providers'; in the BPA the infrastructure integration components, site-local integration components, and the actual service providers.

Policies in PDK version 2 are standards to which adherence can be asserted and that can be assessed and validated – for example as trust marks – and that are endorsed by AEGIS and considered 'standards track'. Policies also are endorsed by the organisation at the appropriate level of management, and express a commitment of adherence by the organisation's management.

The *processes and procedures*, being templates, are reference implementations where we assume these to be specialised for specific deployments. For example, while *Sirtfi* stipulates that "a security incident response capability exists" when dealing with federated security incidents, and hence is a *policy* in terms of the PDK version 2. It does not specify how the incident response capability is to be provided. To support organisations at different strata of AARC BPA in providing such capability, the PDK contains informational 'templates' that can be customised by an organisation depending on their own context. Thus the template for *Incident Response Procedures* is to provide community good practice examples for operational processes. For example, from the eduGAIN Security Handbook, from communities, e.g. the Worldwide LHC Computing Grid, and from infrastructure and service providers, e.g. EGI CSIRT [EGI-CSIRT].

Policies and procedures evolve, and what today are considered appropriate security measures may and will have changed in a few years. Hence, we foresee that the Policy Development Kit v2 will be a continuously evolving suite of Guidelines and Informational documents. The 'first version' of this new PDK that will comprise all in-scope components is foreseen early in 2026, but will continue to evolve thereafter, under the management of AEGIS [AEGIS] and the AARC community.

Figure 2 shows the 5 target audiences and the required or applicable policies and processes in a schematic overview.

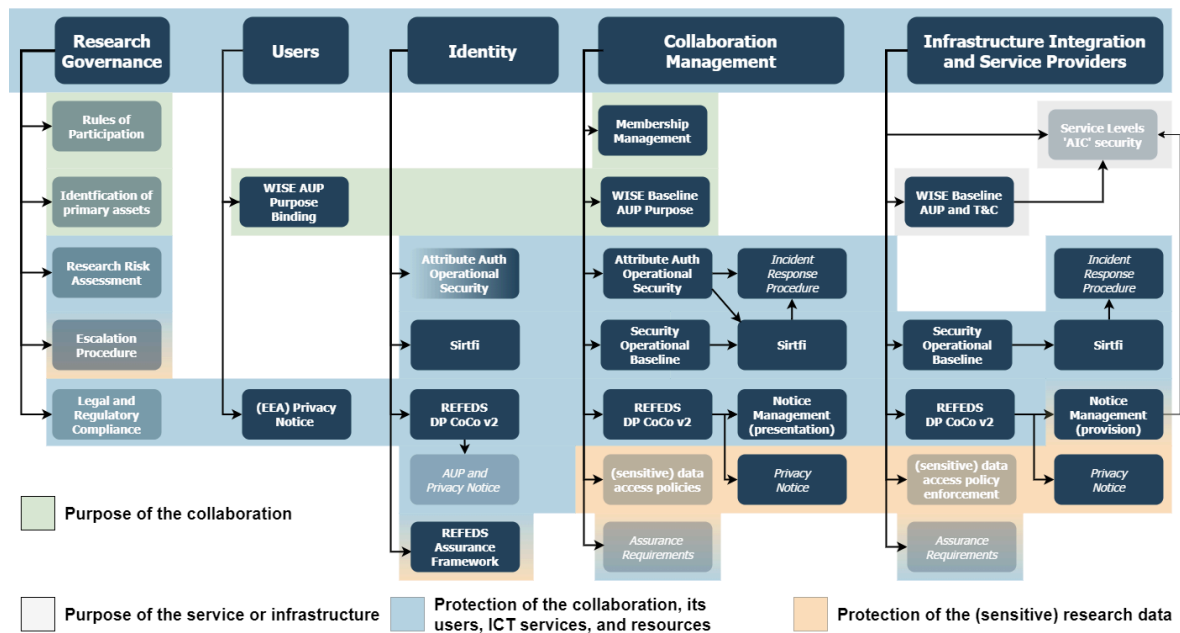


Figure 2: Trust Framework and structure of the AARC Policy Development Kit (PDK) version 2.0. The five main audiences are indicated on top, with the PDK policies relevant to that audience featuring below each (hence some policies and procedures are duplicated). The background colour coding indicated the policy domain – with some policies and procedures addressing straddling the interface between protecting the collaboration resources as well as their users and the research data itself. Shaded elements are relevant policies that are not in scope of the AARC PDK proper since they extend into either non-AAI policy (governance) or are discipline specific (protection of research data rather than the AAI), but are often a prerequisite for adoption of AAI policies. The policies and procedures shown are discussed in the text.

Bearing in mind the scope of the AARC BPA itself, the PDK does not attempt to provide definitive guidance for any and all aspects of collaborative access, although it does identify the elements. In particular, and as discussed in the per-audience sections, research governance, i.e. rules of participation, actually setting the research purposes, is out of scope of the PDK, and hence greyed-out in the trust framework diagram. Similarly, we acknowledge the need for data access policies, especially for sensitive data, and the higher-assurance requirements that go along with the handling of such sensitive data. However, there are many more aspects associated with handling sensitive data, such as permit requirements, review by ethics committees, knowledge safety reviews, that are usually discipline-specific, for example the 'Health Data Access Bodies' (HDABs) for the European Health Data Space [EHDS]. Hence, we do not provide policies for data access, leaving those to discipline specific expert groups, while identifying the need for such policies in this trust framework. Lastly, we do not provide guidelines for *organisational* AUPs and privacy notices, those usually being linked closely to employment or student enrolment terms and conditions in the user home organisation.

Where feasible, the Trust Framework builds on existing standards and practices. Where for services and their providers a significant body of standards exists (such as the ISO 27000 series for IT security [ISO27k], and ISO 20000 for IT service management [ISO20k]), there is

far less guidance for *federated* infrastructures, in particular those featuring multi-lateral relationships.

For these, specific guidance has been developed over the years. The AARC community in 2015 initiated *Sirtfi*, the Security Incident Response Trust framework for Federated Identity. The GÉANT project [GEANT], and later REFEDS [REFEDS], addressed protection of access personal data in the *REFEDS Data Protection Code of Conduct* [DPCoCov2], revised over the 2016-2022 period to address release of personal data used to access services. And the AARC community in collaboration with the interoperable global trust federation [IGTF] created guidelines for security operation of attribute authorities and 'translation' components ("AAOPS"). Similarly associations of service providers and the e-Infrastructures coalesced around good practices for operational security for environments where attacks spread laterally across multiple organisations.

These 'federated' guidelines address federation-specific elements of trust that are not captured by the usual standards and practices; those are either addressing requirements bespoke to a single organisation, e.g. credential management, media handling, personnel hiring policies, backups, or physical & building security, or at best address information exchange in highly structured information sharing communities that are under some form of sector-specific, and conceptually centralised control, like ISO 27010 from the ISO 27000 series [ISO27k].

This Trust framework, and the elements thereof that constitute *Snctfi* as discussed later, intentionally builds on existing guidance and policies that have already been agreed or implemented in the global federated identity ecosystem and they are indicated in the respective audience-specific sections. In doing so, this Trust framework intends to:

- facilitate adoption of the AARC Trust Framework and Snctfi by leveraging existing implementations and organisational alignment;
- build on community consensus, usually established through a multi-stakeholder process that builds on several years of consultation processes;
- ensure continued evolution of the Policy Development Kit that implements this Trust framework, by devolving future responsibility for their content to sustained community groups (REFEDS, IGTF, eduGAIN, WISE [WISE]) and infrastructure and service providers, such as the European Open Science Cloud [EOSC]; and
- be scalable, acknowledging that small- to mid-sized communities will typically engage the services of specialised AAI service providers that already implement common federated standards. The Federated Identity Management for Research [FIM4R] community highlighted the challenges faced by these small- and mid-sized communities of between a handful to a few hundred collaborators.

It may be noted that in establishing policies for authentication and authorization, fundamental questions are often identified that are much broader than what should be considered 'AAI' or 'security' policies. These often concern matters of collaboration governance, compliance, escalation and remediation. This is well recognised in formal risk management frameworks, such as ISO31000, ISO27005, ITSRM² [ITSRM2], which all start with the identification of

‘primary assets’, identifying the ‘raison d’être’ of the collaboration and the environment in which it operates. Many of the design decisions of an AAI, as well as the policies and their source of authority, can only be answered once the governance model, identification of primary assets, and rules of participation are clarified.

In this trust framework, we acknowledge the need for clear collaboration governance, especially for collaborations where the number of participants makes it hard to rely solely on ‘implicit’ or ‘social’ trust. While the size of collaborations where ‘social’ trust is workable is left here to social science researchers, we note that once the number of collaborators gets closer to a hundred rather than to ten individuals, getting clarity on what the collaboration governance is becomes hard or even impossible. For small to mid-sized communities, a ‘principal investigator’, with just a few collaborators, will usually be able to answer such questions.

Taken together, the ‘governance’ and ‘federation’ aspects of a collaboration or infrastructure, when documented appropriately, in fact constitute a ‘top-level’ policy that gives authority to the more specific trust & identity and security policies in this Trust framework. The need for a formalised ‘top-level’ policy is more apparent in large, structured collaborations and infrastructures, but lack of such a policy may hinder trust and security even in smaller collaborations when interpersonal relationships no longer suffice to deal with policy escalations. By construction, no single ‘top-level’ policy is foreseen in the AARC PDK – such a policy needs to be derived from the collaboration governance structure, the identified primary assets, legal and regulatory compliance requirements, and (ultimate) the risk assessment and ‘risk appetite’, e.g. risk averseness in a more bureaucratic organisations, or a higher risk-acceptance in research organisations. In the governance section of the Trust framework we will hence only identify the elements that may feature in a ‘top-level’ policy.

Guidelines by target audience

For small- and mid-size collaborations that do not process particularly sensitive research data, policy suites can be very concise. As discussed below for the *Snctfi* policy set, many of the operational trust and security policies are implemented centrally by the AAI service providers, and as long as the collaboration engages with 'Snctfi' compliant providers, they can limit themselves to the collaboration-specific 'purpose' aspects of collaboration management: setting the purpose of the collaboration (typically a one-liner), and how to manage members, including designing a person, either a principal investigator or their research support engineer, to grant and remove group memberships and roles.

Research Governance

Defining documented *Rules of Participation* applies primarily to larger collaborations, where trust can no longer be based on inter-human relationships and informal understanding. In small to mid-sized collaboration, the principal investigator(s) provide the authority for (AAI) policies to be set and processes to be accepted. In larger collaborations, the source of authority for policies should be clearly defined, as collaboration management roles will evolve over time or roles are reassigned to different individuals. For all AAI policies, either directly adopted by a collaboration or used as selection criteria by the collaboration to choose an AAI provider, the source of authority is determined through the (informal) rules of participation. Apart from the role in establishing authority over AAI policies, the rules of participation are out of scope for the Policy Development Kit.

The main reason for a group of people to work together is, at least to some extent, established at the start of collaboration. This allows the collaboration to identify the *primary assets* in terms of information, e.g. data sets, survey results, lab logbooks, publications, research algorithms, etc., and 'business' processes, such as project execution, scientific methods, pathways to impact for the research, etc.. In small to mid-sized collaborations, these are often taken 'for granted' and not made explicit, but lack of clarity of these primary assets may lead to confusion later, when security or administrative controls intended to protect secondary assets, e.g. IT systems, are enforced that however in themselves jeopardize the primary assets or strategic mission of the organisation as a whole. Hence, we recommend that the 'primary assets', usually just a handful, are identified by the collaboration. These will subsequently inform the *purpose* of the collaboration, the one-line statement to be used in – for example – an acceptable use policy (AUP) based on the WISE Baseline AUP template [WISE-AUP].

When the primary assets include sensitive or particularly valuable data, be it from the perspective of the collaboration or from its legal and regulatory environment, performing a structured *Risk Assessment* may be called for. There are several risk assessment frameworks available for IT systems, which mostly follow the same structure. When embedded in a larger framework, e.g. also including physical safety or protection of

intellectual property and knowledge, frameworks such as ISO 31000 can be usefully employed. For risks relating primarily to information processing, ISO 27005, ITSRM², or OCTAVE [OCTAVE] are more appropriate. The risk assessment specifically informs the security policies, possibly beyond the security operational baseline, assurance requirements requiring certain elements of the REFEDS Assurance Framework [RAF], incident response timeliness and commensurate suspension controls, and data authorization policies.

When research is subject to specific *legal and regulatory requirements*, these should be reviewed in the context of the risk assessment, making sure that the AAI or the providers of AAI services meet these requirements. Some requirements, like the protection of access personal data in the European Economic Area, are permanent features of the regulatory landscape for which AAI reference policies exist – in particular the REFEDS Data Protection Code of Conduct (DPCoCo) [DPCoCov2]. While the REFEDS DPCoCo addresses the use of access personal data, i.e. the personal data released through the AAI in order to access services, it does intentionally not deal with personal data that could be contained in research data itself. If a collaboration deals with personal or otherwise sensitive data, specific data access policies should be defined, and their enforcement be implemented in the connected infrastructures and service providers. In the AARC BPA, such data access controls form the authorization subsystem, where the policy should be defined centrally by the collaboration in a ‘policy administration point’. Even when data itself is not sensitive, research collaborations will usually generate new data, including publications, FAIR data sets, for which policies and practices may be required. The nature of these will depend on the jurisdiction(s) involved and contractual practices, memoranda of understanding, or similar instruments. Who holds rights over the generated research data, as distinct from access data generated by the AAI, is relevant for collaborations, but outside the scope of the Policy Development Kit.

Lastly, *escalation procedures* are needed to ensure that policies are effectively enforced. These may be defined either informally, via discretionary authority by a principal investigator for small collaborations, or form part of the rules of participation. To which extent these need explicit definition is beyond the scope of the Policy Development Kit.

Users and collaboration purpose

End users join a collaboration, either through an invitation process initiated by or on behalf of a collaboration manager, e.g. a principal investigator, or by enrolment. Prior to joining a collaboration, the user may have been enrolled or requested to join the AAI hosting platform used by the collaboration. The user may also have been enrolled earlier therein, having joined other collaborations on a multi-tenant platform.

Apart from any governance-related understandings between the collaboration and the user, the enrolment establishes a *purpose binding*: what is the intent of the collaboration, as derived from implicit primary assets. In jurisdictions where GDPR applies, it also starts the processing of ‘access’ personal data, since personal data of the user will be processed for collaboration management. Those responsible for the collaboration should establish the role of data controller for the ‘access’ personal data, which in most cases will be the home organisation, i.e. the employer, of the collaboration manager or principal investigator.

However, this is by no means required: the collaboration manager may also act under the direction of another designated legal, or natural, entity. Designating the data controller is necessary to present the privacy notice, mandatory in at least the EEA, and hence should be done before a collaboration starts enrolling members.

In case a collaboration is hosted on a multi-tenant AAI platform, during enrolment in a collaboration the data controller, as meant in the EU GDPR, will likely change from the platform provider to the hosting institution of the collaboration manager. In these cases, the *Privacy Notice* will have to be presented to the user indicating this change of controllership. Since at such a moment also the binding to the collaboration *purpose* takes place, the collaboration management component shall present the user with an updated Notice at that point.

The PDK recommends that the notice consist of the WISE Baseline AUP including both the new purpose binding and an updated designation of the data controller as part of the privacy notice section of the WISE Baseline AUP. At this time, additional terms and conditions may be added to the notice, such as conditions on sensitive research data or of personal data contained therein.

Responsibility for notice management and signalling, in accordance with the AARC guideline on Notice Management, lies with the collaboration, but likely uses facilities of the AAI platform provider.

Identity

At the authentication layer, the BPA identifies both external identity sources, including home organisation identity providers, social identity, or citizen identity, as well as identity integration that supports identity enrollment and linking across multiple authentication sources. The latter is the logical layer where a non-re-assigned identifier is assigned to an entity, irrespective of the source of identity for a user. Other attributes of the user will depend on the upstream authentication source, such as affiliation, freshness, and assurance level.

The identity capability does not have a trust role for the collaboration's purpose and its membership, groups, and roles – and hence in its role as an identity layer does not augment the acceptable use policy, and – although it has an implicit acceptable usage – such usage is fully defined by the collaboration management component or collaboration AAI platform provider where users may enrol prior to joining a collaboration. Any notice presentation is thus logically assigned to the collaboration management capability or platform.

The home organisation, as an upstream authentication source, which may include citizen or social authentication, of the user *will* have an acceptable use policy, typically setting many more terms and conditions than the WISE Baseline AUP. It is likely to address organisational enterprise use, any allowed personal use of the organisations' resources, and sanctions in case of non-compliance. However, while a home organisation may choose to base its acceptable use on the WISE Baseline AUP, or to materially include it, that is not commonly

the case. Hence the AUP will almost certainly have to be presented, for the first time, by the collaboration management component or platform.

The fact that the policies for acceptable use, terms & conditions are set independently, raises the possibility of inconsistencies between the research governance and the home organisation. Examples thereof would be the ownership and access conditions to research data and foreground intellectual property (IP) created in the collaboration, the use of background IP by collaborators, copyright, and any moral rights established by the user. The rights that the individual user, rather than for example their employer, has over the foreground is highly dependent on jurisdiction and employment conditions. Such relations should be worked out as part of research governance and are outside the scope of this PDK Trust framework, but may be noted in the (AUP) Notice managed by the collaboration.

The identity capability, either as a stand-alone platform component or integrated in the collaboration management component in the BPA 2025, may combine or link more than one 'upstream' source of authentication to a user. In the BPA this layer has no logical membership management function and hence aligns more closely with the 'Authentication Sources' in the trust framework. In linking accounts from upstream sources, the identity component must comply with the account linking [AARC-G009] and observe the architectural guidelines on inference and expression of affiliation information ([AARC-G057] and [AARC-G025]).

This is specifically important when the collaboration or infrastructures (service providers) have assurance requirements expressed as REFEDS Assurance Framework qualities and assurance profiles. Since in the AARC BPA the identity component is opaque to the collaboration and infrastructures, the identity component must make sure the assurance components are either met by the instantaneous authentication source or provide its own step-up mechanisms - be it for the authenticator strength or identity vetting requirements. While some sources may provide good authenticator strength, they may lack identity-vetting assurance, e.g. most social identity sources. Upstream sources should be disclosed via the relevant AARC architectural standards to provide the necessary transparency to collaborations and infrastructure service providers.

The identity capability is stateful due to its role in account linking, and it will keep records of user interaction. As an "issuer of statements for entities", it provides a control point and a traceability capability in case of security incident response. The operational security requirements of the "AAOPS" guideline [AARC-G071] apply to the identity component in as far as they concern attribute assertions, operational environment, key management, network, site security, assessment & review and privacy elements, i.e. the 'AAS', 'OE', 'KM', 'NET', 'SITE', 'AR', and 'PC' requirements.

Sirtfi similarly applies to this component. While some aspects of Sirtfi can be covered by an identity component itself taking active measures, such as emergency suspension, it is worth noting that security incident resolution depends on the ability to trace back to the initiating entity: the impersonated user or miscreant. Upstream authentication sources should be aligned with Sirtfi, such that communication to the ultimate source is assured in a timely

manner. Special care should be taken when self-signup authentication sources are connected as external identity providers to the identity component in the BPA. Adding such open 'social' identity sources must be disclosed by the identity component provided to collaborations and the connected infrastructures and service providers.

Any authentication source being stateful and handling access-personal data for human end-users means that REFEDS Data Protection Code of Conduct (DPCoCo) must be adhered to. Adherence to DPCoCo may, to a limited extent, and dependent on the user home organisations, aid with the release of the *personalised* attribute bundle.

The collaboration or the infrastructure service providers set the identity and authentication assurance requirements on the authentication sources, based on their own policies, and following this trust framework determined by the risk assessment for access to sensitive data and legal and regulatory compliance. In this trust framework, it should be expressed in terms of the REFEDS Assurance Framework, as discussed under collaboration management trust below, hence the authentication sources must support the REFEDS Assurance Framework. This similarly holds for the REFEDS authenticator profiles i.e. single factor [SFA] and multifactor authentication [MFA].

Collaboration Management

The collaboration management in the 2019 version of the Blueprint Architecture [AARC-G045] has a single component identifying both the linking of identities – which in the new BPA '2025' [AARC-G080] is part of the identity capability – as well as the collaboration specific aspects: attribute enrichment and authorisation. The 'identity' layer being considered part of the trust framework for authentication sources, possibly sourced from identity integration components or aggregators, the collaboration trust is based on its membership management and the adherence of its members to the purpose of the collaboration.

The collaboration management capability itself also needs protection. The relevant guidelines for trust are common between all operators: operational security of the attribute authority and its protocol translation and attribute enrichment mechanism, the security operational baseline of the service itself, Sirtfi, and adherence to the REFEDS Data Protection code of conduct.

Notice management, an indirect requirement of the data protection code of conduct, it a function of collaboration management, both in terms of its own operation, but also because collaboration management is the single component that all users of a collaboration, and hence all its service and infrastructure providers, will see at some point of their membership life cycle. The Guideline on notice presentation [AARC-G083] therefore places a special responsibility on the collaboration management component as the place where notices should be presented to users; they may also be presented in other places, but at least in collaboration management. This does not mean that the collaboration itself must run such a presentation service. Like for the other elements, the collaboration may source this service component from a specialised provider.

Federated operational security guidance

The operational security guidelines apply horizontally across the BPA components of collaboration management and infrastructure integration. While they are equally applicable also to authentication sources and site-local integration and services, in these two cases 'conventional' information security management standards already cover the majority of operational security. The exception to this is *Sirtfi*, which deals specifically with the communication of security incidents in a federated context, and the mutual assumptions and collaboration that parties in a federated AAI exchange make. The *Sirtfi* hence apply to any federated entity operator.

The collaboration and infrastructure integration both convey trust, and act as a usually opaque source of augmented claims or attributes. As such, it avails itself over sensitive material, i.e. signing keys and membership databases, and at the same time 'hides' any upstream sources. This makes these components, as aggregators, more 'valuable' than a single authentication source or single service, thereby changing its risk profile.

In this trust framework, we separate their trust elements in two parts:

- a common 'security operational baseline' that applies to any service, including end-services and sites. The baseline 'defines minimum expectations and requirements of the behaviour of those offering services to users and communities, and of those providing access to services or assembling service components', and focusses on the *ability* of service provider organisations to meaningfully participate in security collaboration. It targets loosely coupled federations of service providers across administrative domains, but who offer services that collaborations will compose into higher-level service offerings, e.g. to compose workflows in a 'research pipeline'.
The security operational baseline includes *Sirtfi* by reference.
- Guidance for secure operation of the attribute authorities and issuers of statements for entities, addressing the specific elements of cross-service collaborative trust: naming, assertion integrity, expression of authorisation information, and conveyance of claims and attributes.

The Policy Development Kit version 2 implementing the trust framework will contain specific guidance for each: the security baseline [AARC-G084], evolving from the *Site Security Policy* in PDK version 1 and building on additional experience in national and European EOSC nodes; and the 'Guideline for Secure Operation of Attribute Authorities and issuers of statements for entities' ("AAOPS" [AARC-G071]), evolving from the IGTF assurance profiles for cross-domain trust, as well as taking into account the possibility for both *push* (SAML POST, PKI attribute certificate flows) and *pull* (LDAP, API, and OIDC flows) of authorisation information in the typical AAA flows identified in the generic AAA architecture [RFC2903]. The "AAOPS" guideline will be re-assessed in the evolution of the Policy Development Kit version 2.

Both Sirtfi and the AAOPS guideline presuppose that a security incident response mechanism is in place, and that a documented incident response procedure exists. Such a procedure is community good practice, and in many cases mandated for regulatory compliance. The procedure must be adjusted to fit the service and infrastructure, the organisation, and the local environment. An annotated template, based on existing infrastructures and the example in the first version of the AARC Policy Development Kit, is scheduled for inclusion also in the second version of the PDK.

Security Operational Baseline and Sirtfi

The Security Operational Baseline (AARC-G084) provides a reference set of minimum expectations and requirements of the behaviour of those offering services to users, communities, and other participants in the distributed BPA ecosystem, and of the behaviour of those providing access to services or assembling service components. It establishes a level of trust between all Participants in the Infrastructure for reliable and secure operation based on a list of requirements, augmented with references to current good practice for software vulnerability management, configuration, and checklists such as OWASP [OWASP] and the NIST security software development framework [NIST-SSDF].

Trust in the security response capability is founded on *Sirtfi*, the Security Incident Response Trust Framework for Federated Identity [Sirtfi]. Recognising that tracing security incidents through a chain of AAI components emphasises *traceability* and *timeliness of response*, periodic communications tests and security exercises are part of the baseline. In practice, most providers that have mature security processes in place will meet these requirements.

The supplementary aspects deal with collaborative incident response and proactive data sharing. Since identifiers used for access control will change as the activities move between AAI layers and platforms, ensuring account linkage and the logging of such linking is essential. During incident response, these must be accessible to the relevant incident response teams in near-real time. In this aspect, the security operational baseline augments Sirtfi: while Sirtfi version 2 requires proactive notification of affected federation partners, bi-lateral transmission of information is too slow for federated incident response. Federated infrastructures, such as eduGAIN, have hence initiated security coordination teams that aid in managing the information flow. AAI operators and collaboration platforms shall also take part in such coordination, e.g. by affiliating themselves to a collaborating infrastructure, as indicated in the Security Operational Baseline.

Attribute authority and assertion issuer operations

The 'Guideline for Secure Operation of Attribute Authorities and Issuers of statements for entities' ("AAOPS") addresses the hybrid model acting as a relying party as well as a source of authority, while at the same time managing collaboration data. In its role as membership directory, it retains access-personal data that needs to be protected, generates signed assertions for which confidential key material is needed, and it can re-generate the identifiers that are needed for traceability in incident response.

Importantly, when used in a collaboration platform, it is the single place where a component holds information about the hosted communities on the platform, and is thus the linking pin between infrastructure and service providers and the collaboration. The guidance thus also includes publication and meta-data responsibilities.

The AAOPS guideline also includes Sirtfi by reference, but augments it with a broader trust context: naming, attribute management and release, key management, site and network security, privacy, meta-data publication, and business continuity or disaster recovery. It also places obligations on the relying parties to maintain trust: the need to validate and verify integrity, honour lifetime limitations on assertions, and maintain the balance between lifetime and the ability to revoke compromised assertions.

Data Protection Code of Conduct best practice

A collaboration can deal with personal data of two types: (1) personal data of the users using the infrastructure themselves, or (2) personal data contained in the research objects processed within that infrastructure, or in the data exposed in the infrastructure.

The REFEDS Data Protection Code of Conduct “describes an approach to meet the requirements of the EU GDPR in federated identity management”, and is applicable to first type of data: the personal data used for the purposes of access to service providers in a federated infrastructure, referred to as ‘access personal data’.

Handling of access personal data throughout the chain of BPA components should comply with the REFEDS DPCoCo so that users are consistently informed of the processing of their data, the exposure of their data is minimized, but that at the same time a coherent service can be delivered that comprises service providers and infrastructures from many independent administrative domains. While the collaboration management layer is either operated by the collaboration itself, principally the primary hosting organisation of the collaboration, or sourced from a platform provider, the other entities in the federation are typically independent controllers, as meant in the GDPR. They do not share information between them, yet in order to provide a collective service to the user, all services must agree on the identifier(s) used between them for their users. It is thus necessary that the user identifiers are omnidirectional and persistent (next to their basic quality of being non-reassigned). At least a ‘personalized’ identifier (the SAML Subject ID or the OIDC sub claim) is necessary for providing collective services for research that spans multiple administrative domains.

REFEDS DPCoCo, while normatively scoped to the European Economic Area, provides a de-facto global best practice for dealing with access-personal data. This trust framework uses DPCoCo as a common ground for the collaboration management layer, infrastructure integration, and site-local integration and services, such that the authentication sources can relay access personal data of their users without having to do a pairwise trust review. The Policy Development Kit version 2 will include by reference REFEDS DPCoCo.

A set of informational practices and a Privacy Notice Template accompany REFEDS DPCoCo. The template lays out a mechanism for automatically populating the contents of a notice based on meta-data, but in whatever way the notice is presented it should meet local

regulations. Privacy notice should also be exposed through the Notice Management framework [AARC-G083], such that it can be included by reference in amalgamated notices in the collaboration management layer.

Research data may also contain personal data: this is the 'second type' of personal data, contained in or exposed through the connected services and infrastructures. When collaborations processes personal data, in many cases this is subject to legal, research integrity, intellectual property, or export compliance conditions. In the EEA this would be GDPR, for data regarding humans it could include required approval by a medical-ethical committee, non-disclosure agreements with partners, et cetera. The policies, processes, and procedures that are applicable to handling sensitive data are outside the scope of the Policy Development Kit, yet they will place specific assurance requirements on the security of service providers, collaboration management providers, and the identity sources.

Assurance requirements and acceptable assurance

The logical separation of identity layer and collaboration management layer introduced in the BPA 2025 shifts some of the aspects of trust that were hitherto in scope of the collaboration management component towards the identity component of the trust framework, particularly regarding the provision of assurance information.

The corresponding minimal requirements on assurance set by the collaboration come from the risk assessment conducted as part of research governance, and will determine the 'acceptance assurance' as provided by the authentication source(s). Since also the service and infrastructure providers will have such minimal assurance requirements (to mitigate their own risks), these will have to be combined with those of the collaboration before being compared to the provided assurance level by the authentication source(s). The trust framework recommends that these requirements be expressed in terms of the assurance components identified in the REFEDS Assurance Framework. For example: if the conclusion from the combined risk assessment is that affiliation 'freshness' should be 'approximately something around three weeks', it is preferred that this is mapped to the REFEDS affiliation freshness 'assurance/ATP/ePA-1m' (one month), rather than defining a bespoke requirement.

Infrastructure Integrations and Service Providers

For service providers, site-local integrations, and infrastructure integration components to be trusted by their users, the 'upstream' collaboration management, with the research communities they serve, and the identity sources, they need to align with most of the policies that are applicable also to the collaboration management layer. Only if the infrastructure or site-local integration is stateless, and managed as a local service within the same administrative domain, the relevant parts of AAOPS, i.e. the subset identifier for the identity capability, are implicit in the operational security of the services.

A specific service provider and any site-local integrations operate within a single administrative domain, making enterprise-oriented security specifications applicable to their

service management. This is a well-known area with extensive guidance already available, including ISO standards, NIS 2 [NIS2], the maturity models of many national research and educational networks, and applicable best practice, compatible with an open research environment and mindful of the primary assets of the organisations involved, i.e. research, open science, can be followed.

Service providers should review, based on their own risk assessment, whether the WISE Baseline AUP foundation is sufficient for their purpose. Adopting the Baseline AUP without additional clauses or requirements ensures a smooth workflow for users, as no additional review of terms and conditions is necessary beyond what the user already agreed to when enrolling in the collaboration. If specific terms and conditions are necessary, for example because the service or data provider gives access to sensitive data, or has usage and access restrictions that cannot be inferred from the community and user attributes received, the WISE Baseline AUP provides anchor points for these to be added. This also applies to 'promises' to the user, reflecting any service level agreements in place.

The service provider should provide notice meta-data for the AUP and any additional terms and conditions in line with the guideline on Notice Management [AARC-G083].

Trust qualification through Sncfti

The first version of the 'Scalable Negotiator for a Community Trust Framework in Federated Infrastructures' [Sncfti1] was established to enhance the trust in the research collaboration AAls towards the national research and education (R&E) federations and eduGAIN. Building on the structure of the WISE Security for Collaboration among Infrastructures (SCI) framework it establishes a policy framework that allows determination of the 'interoperable trust' of AARC Blueprint components and the community of services behind these. For example, an otherwise opaque infrastructure integration component, i.e. where an identity federation cannot directly observe in federation meta-data the services that are integrated in an infrastructure, would be able to express to the R&E federations that it has an internally-consistent policy set.

A Sncfti compliant protocol-translation component under version 1 could make statements about all its constituent services and resource providers, and that it would abide by best practices in the R&E community, such as adherence to the REFEDS Data Protection Code of Conduct, REFEDS Research and Scholarship (R&S) entity category, and Sirtfi. Sncfti relied on a peer-reviewed self-assessment model and transparency of policy through documentation – a model that is also used for the IGTF identity trust fabric and for the attribute authority operations "AAOPS" policy.

The same set of policies and the policy transparency, with self-assessment and where relevant supported by peer review, proves equally valuable towards collaboration management. Over the years, multiple AARC BPA compliant providers of collaboration infrastructure have emerged that offer 'AAI as a platform', with several providers of both identity aggregation services, collaboration management services, and infrastructure integrations. While the AEGIS mechanism is in place for architectural interoperability, there is no equivalent system for assessing suitability when engaging AAI platform providers. Since collaborations are heterogeneous and will have different requirements on their providers concerning resilience, assurance levels, scope of the service, and costs, strict interoperability is neither in scope of standardisation and likely not required. Yet to support collaborations in making appropriate choices, alignment with a specific verifiable subset of the AARC PDK policies and any supporting procedures. Similar to Sncfti version 1, this goal can be achieved by self-assessment, and where appropriate peer review, and statements from the respective providers to the research governance, community, and collaborations on which they can base their choice.

The revised version of Sncfti (version 2) comprises the *assessable and verifiable subset of policies and procedures that an AAI platform provider can assert* when engaging with both collaboration management as well as with service and infrastructure providers.

In this trust framework, and hence in version 2 of the Policy Development Kit, this corresponds to the aspects of collaboration management that can be outsourced:

- Attribute authority operation and the operation of the collaboration management components “AAOPS” (AARC-G071)
- The security operational baseline for the collaboration services offered by the platform (AARC-G084)
- Adherence to the REFEDS DPCoCo version 2
- Implementation of the Guideline on Notice Management (AARC-G083)

For site-local integrations and their connected services, where the assertions are consumed within the same administrative domain, attribute authority operations do not require specific attention beyond the information security management system by the provider.

However, infrastructure integration, site-local integration, and service provider should all support the provision of notice management information since this is consumed by the collaboration management layer in the presentation of aggregated notices.

Through the Security Operational Baseline, the AAI providers that assert Snctfi also confirm that the entities exposed either participate in the collaboration themselves, or that the AAI provider has controls in place to mitigate the information security risks in accordance with their published policies. In this aspect, Snctfi addresses the supply-chain risk inherent in opaque protocol translation components, in line with the concepts from, for example, NIS 2.0

The diagram below illustrates the subset of policies and their associated procedures that comprise the Snctfi version 2 subset of this trust framework.

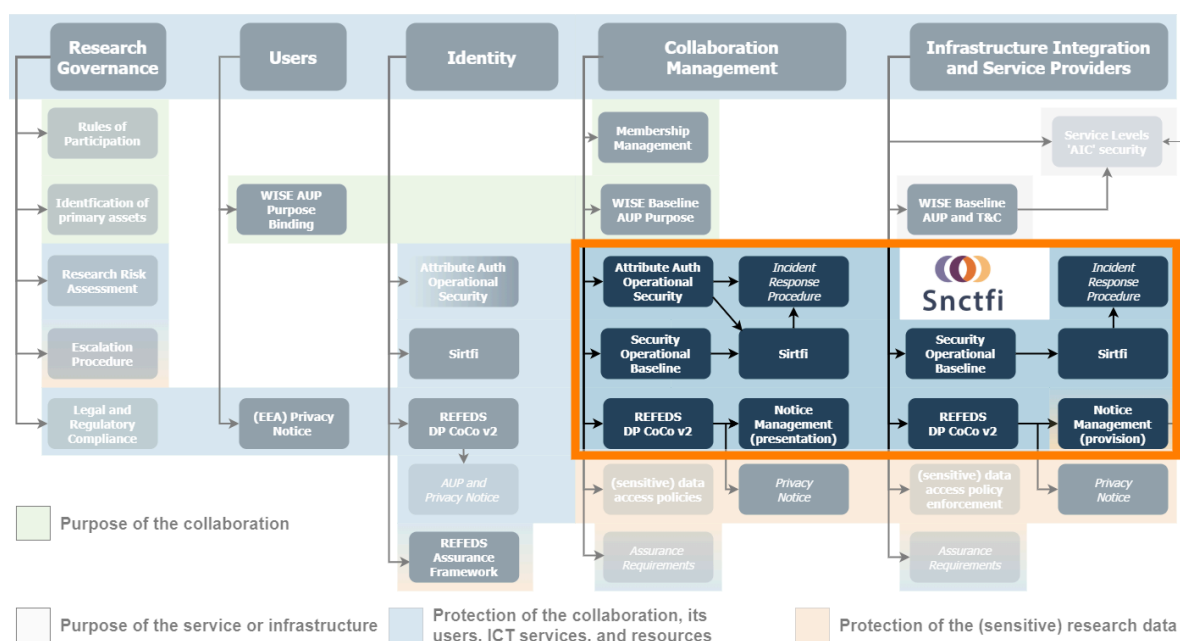


Figure 3: The ‘Snctfi’ sub set of accessible and verifiable policies in the Policy Development Kit version 2 structure that AAI platform providers can assert, overlaid on the structure of the PDK version 2 trust framework.

Procedural and implementation guidance

While this Trust framework emphasises policy, allowing endorsement by the appropriate level of management, their effectiveness depends on the processes and procedures that implement them. The majority of the policies also indicates for which aspects procedures should be in place: Sirtfi mentions notifications ‘should also follow the security procedures of any federations to which your organisation belongs’, i.e. IR3 from the Sirtfi requirements, that a defined ‘process is used to manage vulnerabilities’, OS2 therein, and the security baselines implies process definitions as generally accepted IT security best practices.

Other trust components, such as notice management and the protection of access-personal data, will result in processes and procedures that help the researchers and collaborations to improve the trust by the users, i.e. researchers, in the ecosystem.

An operational implementation of this trust framework therefore requires that processes and procedures are in place, but – by its nature – should not prescribe what these processes and procedures should be. The Policy Development Kit version 2 will therefore provide reference examples – in line with the AARC Compendium work – of community good practices in these areas:

- Security Incident Response: building on the eduGAIN Security Handbook and the Incident Response procedures of the larger research and e-infrastructures, since these are typically the most complete. This is required both for the security baseline as well as for Sirtfi.
- Privacy notices; also known as privacy ‘policies’
- Purpose bindings based on examples from research infrastructures

Collaborations and infrastructures that process restricted or sensitive data will need to add specific processes to implement consistency between membership management and data access rights. Data access rights are currently managed in a bespoke way - systems like the Resource Entitlement Management System [REMS] often used in the life sciences, but also the built-in systems in research data management repositories (Invenio, Djehuty) all have their own way of managing either role-based or ‘magic URL’ based access.

When consolidated in the AARC BPA model, the management of these policies should use administrative interfaces of a Policy Administration Point, in the BPA 2025 the ‘PAP’ component. For managing these data policies, higher assurance may be required. If the data access system is distributed, both automated mechanisms and processes must be in place to ensure the policy decision (PDP) and enforcement (PEP) points implement the data access policies set in the PAP. A reference process or procedure for these will not be part of the Policy Development Kit, but we refer to the AARC Compendium to identify those research infrastructures handling sensitive data.

We do note that the collaboration, usually together with the data providers, must describe policies and practices for this PAP-PDP-PEP authorisation chain, and that – at least logically – the responsibilities lies with the collaboration and infrastructure AAI layer where the data

access policies are defined in the PAP, since it needs to ensure consistency between the PAP and the enforcing PEPs in any downstream service provider. Protections on the PAP, PDP and PEP should consider the same operational security requirements identified in the “AAOPS” guideline [AARC-G071], where the PAP - as the source of authority - is the most crucial component.

Evolution of the policy development kit

The Trust framework outlined herein has identified the primary elements needed for federated collaborations when they engage following the AARC Blueprint Architecture. For some of the policies, existing guidance for federated identity management has been developed that can be usefully applied directly, in particular the *Sirtfi* security incident response framework as augmented by the eduGAIN Security Handbook, as well as the REFEDS Data Protection Code of Conduct and the WISE Baseline Acceptable Use Policy. In other cases, specific guidance is required - such as the Security Operational Baseline, coherent Notice presentation, and the operation of trustworthy attribute authorities and AAI components – or reference procedures can be usefully shared.

The Policy Development Kit version 2 will provide the components of the Trust framework that are in-scope for the AAI proper, both existing as well as AARC-initiated guidance. The AARC specific policies will be in the form of AARC Guidelines that can be adopted by AEGIS and the AARC community and can be used as-is for (self) assessment and verification. Procedures will be provided in the form of AARC Informational templates, 'commented' or annotated, which may be used as a reference by collaborations, AAI service providers, and infrastructures and their services, for implementation of the policies.

Some of these Guidelines have already been developed and are deployed, are in the AEGIS review process or are currently being incorporated in operational federations. Yet even the already adopted ones will see continued evolution, and this evolution will continue for as long as federated access for research collaboration will thrive. AARC-G071, i.e. secure operation of attribute authorities and of issuers of assertions and statements, AARC-G083, coherence presentation of notices and reduction of interstitial pages, and AARC-G084, security operational baseline for services and infrastructures, will all be part of the evolved policy development kit.

The Policy Development Kit version 2, whose release is foreseen in early 2026, will thus both be an implementation of this Trust framework as well as the starting point for a sustained policy evolution process, indicating which policies and guidelines are appropriate to be dealt with first for new communities.

Acknowledgements

This work has been co-supported by the AARC TREE project, which is funded by the European Union under the HORIZON-INFRA-2023-DEV-01 call with grant number 101131237.



We also gratefully acknowledge the contributions by the membership of the Policy Working Group (PWG) of the Australian Access Federation AAF.

References

- AAF** Australian Access Federation, <https://aaf.edu.au/>, visited May 2025
- AEGIS** AARC Community 'AARC Engagement Group for Infrastructures', <https://aarc-community.org/about/aegis/>, visited May 2025
- AARC-G009** AARC Project 'Account linking and LoA elevation use cases and common practices for international research collaboration', 2017, <https://aarc-community.org/guidelines/aarc-g009/>
- AARC-G021** AARC Consortium and Applnt members, 'Guideline on the exchange of specific assurance information between Infrastructures (AARC-G021)'. Zenodo, Feb. 15, 2018. doi: <https://doi.org/10.5281/zenodo.1173558>, <https://aarc-community.org/guidelines/aarc-g021/>
- AARC-G025** AARC Consortium Partners and Applnt members, 'Guidelines for expressing affiliation information (AARC-G025)'. Zenodo, Apr. 08, 2019. doi: <https://doi.org/10.5281/zenodo.3700927>, <https://aarc-community.org/guidelines/aarc-g025/>
- AARC-G031** AARC Consortium Partners and Applnt members, 'Guidelines for the evaluation and combination of the assurance of external identities (AARC-G031)'. Zenodo, May 18, 2018. doi: <https://doi.org/10.5281/zenodo.1308682>, <https://aarc-community.org/guidelines/aarc-g031/>
- AARC-G045** AARC Community members and Applnt members, 'AARC Blueprint Architecture 2019 (AARC-G045)'. Zenodo, Nov. 06, 2019. doi: <https://doi.org/10.5281/zenodo.3672785>, <https://aarc-community.org/guidelines/aarc-g045/>
- AARC-G057** D. L. Groep, C. Kanellopoulos, M. Linden, 'Inferring and constructing origin-affiliation information across infrastructures (AARC-G057)'. Zenodo, Jan. 12, 2021. doi: <https://doi.org/10.5281/zenodo.4433649>, <https://aarc-community.org/guidelines/aarc-g057/>
- AARC-G071** D. L. Groep *et al.*, 'Guidelines for Secure Operation of Attribute Authorities and issuers of statements for entities'. Zenodo, Apr. 11, 2022. doi: <https://doi.org/10.5281/zenodo.5927799>, <https://aarc-community.org/guidelines/aarc-g071/>
- AARC-G080** AARC Community members and Applnt members, 'AARC Blueprint Architecture 2025 - Initial Revision (AARC-G080)' (under review) <https://aarc-community.org/guidelines/aarc-g080/>
- AARC-G083** David L. Groep, David Kelsey, Maarten Kremers, Nicolas Liampotis, Hannah Short, Anrout Terpstra, Catharina Vaendel 'Guidance for Notice Management by Proxies' AARC-G083, doi: <https://doi.org/10.5281/zenodo.14452340>, <https://aarc-community.org/guidelines/aarc-g083/>
- AARC-G084** D.L. Groep (ed.) 'Security Operational Baseline for Proxies and Services'. AARC Community, Apr. 01, 2025. doi: <https://doi.org/10.5281/zenodo.15119296>, <https://aarc-community.org/guidelines/aarc-g084/>, visited May 2025

AARC-GLS	AARC Community members, 'AARC Guidelines', https://aarc-community.org/guidelines/ , visited May 2025
DPCoCov2	REFEDS group 'Data Protection Code of Conduct Home' 2022, https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home , visited May 2025
eduGAIN	eduGAIN, official website, https://edugain.org/ , visited May 2025
EGI	EGI federation 'EGI: Advanced Computing Services for Research', https://egi.eu/ , visited May 2025
EGI-CSIRT	EGi federation 'EGi Computer Security Incident Response Team', https://csirt.egi.eu , visited May 2025
EHDS	Council of the European Union, European Parliament, 'Regulation (EU) 2025/327 of the European Parliament and of the Council of 11 February 2025 on the European Health Data Space and amending Directive 2011/24/EU and Regulation (EU) 2024/2847 (Text with EEA relevance)', 2025 Official Journal L 327, 1-96, http://data.europa.eu/eli/reg/2025/327/oj , visited May 2025
EOSC	EOSC Association 'European Open Science Cloud', https://eosc.eu/ , visited May 2025
FIM4R	Federated Identity Management for Research community 'Federated Identity Management for Research Collaborations', June 2018, doi: https://doi.org/10.5281/zenodo.1307551 , https://fim4r.org/ , visited May 2025
GDPR	Council of the European Union, European Parliament, 'Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)', 2016 Official Journal L 119, 1-88, http://data.europa.eu/eli/reg/2016/679/oj , visited May 2025
GEANT	GÉANT 'GN4 projects', 2016-2022, https://geant.org/projects/ , visited May 2025
GFD-I.032	Shawn Mullen, Matt Crawford, Markus Lorch, Dane Skow 'Site Requirements for Grid Authentication, Authorization and Accounting', SEC SAAAR-RG, Open Grid Forum, 2004, https://ogf.org/documents/GFD.32.pdf
IGTF	'Interoperable Global Trust Federation (IGTF)', 2005, https://igtf.net/ , visited May 2025
ISO20k	International Organization for Standardization 'ISO/IEC 20000 Information technology — Service management', https://www.iso.org/standard/70636.html , visited May 2025
ISO27k	International Organization for Standardization 'ISO/IEC 27000 family Information security management', https://www.iso.org/standard/iso-iec-27000-family , visited May 2025
ITSRM2	European Commission DG-DIGIT 'IT Security Risk Management Methodology' in 'JRC Technical Report Information security in the age of EU - Institutions digitalisation, a landscape analysis', https://commission.europa.eu/document/download/650e25cf-cf41-412a-80b6-628c2aed50e2_en , visited May 2025

Kremers2023	Maarten Kremers, 'SRAM community cross-service AAI', in 'Abingdon 59th EUGridPMA+ Meeting Summary', October 2023, https://eugridpma.org/minutes/59 , visited May 2025
MISP	Cynthia Wagner, Alexandre Dulaunoy, Gérard Wagener, Andras Iklody 'MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform', in Proc. 2016 ACM Workshop on Information Sharing and Collaborative Security, pp. 49-56, ACM 2016, https://www.misp-project.org/ , visited May 2025
MFA	REFEDS 'REFEDS MFA Profile', 2018, https://refeds.org/profile/mfa
NCRIS	Australian Government Department of Education 'National Collaborative Research Infrastructure Strategy', 2023, https://www.education.gov.au/ncris , visited May 2025
NIS2	Council of the European Union, European Parliament, 'Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive) (Text with EEA relevance)', 2022 Official Journal L 333, 1-73, http://data.europa.eu/eli/dir/2022/2555/oj , visited May 2025
NIST-SSDP	National Institute of Standards and Technology 'Secure Software Development Framework', https://csrc.nist.gov/projects/ssdf , visited May 2025
OCTAVE	Richard A. Caralli, James F. Stevens, Lisa R. Young, William R. Wilson 'Introducing OCTAVE Allegro: Improving the Information Security Risk Assessment Process', Software Engineering Institute Carnegie Mellon University, 2007, https://insights.sei.cmu.edu/documents/786/2007_005_001_14885.pdf , visited May 2025
OWASP	'Open Worldwide Application Security Project' https://owasp.org/ , visited May 2025
PDKv1	Hannah Short, Uros Stevanovic (eds.) 'Policy Development Kit', AARC Community, 2019, https://aarc-community.org/policy/policy-development-kit/ , visited May 2025
RAF	REFEDS 'REFEDS Assurance Framework version 2.0', 2023, https://refeds.org/assurance , visited May 2025
REFEDS	Research and Education FEDerations group (REFEDS), https://refeds.org/ , visited May 2025
REMS	J. Leinonen, M. Hakalaand T. Nyrönen, REMS – Resource Entitlement Management System. Zenodo, 2024. doi: https://doi.org/10.5281/zenodo.13692099 , https://github.com/CSCfi/remss
RFC2903	C. de Laat <i>et al.</i> 'Generic AAA Architecture', RFC editor, RFC 2903, http://www.rfc-editor.org/rfc/rfc2903.txt
RFC8417	P. Hunt (ed.) <i>et al.</i> 'Security Event Token (SET)', RFC editor, RFC 8417, http://www.rfc-editor.org/rfc/rfc8417.txt
SFA	REFEDS 'REFEDS Single Factor Authentication Profile', 2018, https://refeds.org/profile/sfa

Sirtfi	REFEDS, AARC Community 'A Security Incident Response Trust Framework for Federated Identity (Sirtfi) version 2', https://refeds.org/sirtfi , visited May 2025
Snctfi1	AARC Community, members of the IGTF 'Snctfi - the Scalable Negotiator for a Community Trust Framework in Federated Infrastructures', 2017, https://www.igtf.net/snctfi/ , visited May 2025
SURFconext	SURF co-operative 'SURFconext: Secure access everywhere with one set of credentials', https://www.surf.nl/en/services/identity-access-management/surfconext , visited May 2025
WISE	WISE Community 'Wise Information Security for E-infrastructures (WISE)', https://wise-community.org/ , visited May 2025
WISE-AUP	Members of the WISE Community SCI working group 'WISE Baseline AUP', 2019, https://wise-community.org/wise-baseline-aup/ , visited May 2025
WLCG	WLCG Collaboration, 'The Worldwide LHC Computing Grid (WLCG)', https://home.cern/science/computing/grid , visited May 2025