Deliverable D2.2 (AARC-I088):

AARC Policy Development Kit Revision

Publication Date	24-02-2026
Due Date	28-02-2026

Authors: David L. Groep (Nikhef NWO-I) (ed.), Liam Atherton (UKRI STFC RAL), Peter Bolha (Masaryk University, CESNET), 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-TREE D2.2 **(AARC-I088)**
Publishing Organisation:	Nikhef (NWO-I)
DOI:	10.5281/zenodo.18661187

Abstract

To facilitate the trust establishment between AAI participants using the Blueprint Architecture as their model, the AARC community designed a trust framework for proxies and Snctfi research services (AARC-I082). This trust framework identifies five distinct target audiences: external identity sources, the identity management, collaboration management, and service integration components, and site-local integrations and services.
However, each of these target audiences will have to address similar aspects of trust and of their mutual interactions. To facilitate this task, especially for the most prevalent use case of 'small-to-medium-sized' collaborations (a few tens up to a few hundred collaborators), AARC provides a Policy Development Kit (PDK) that brings together community good practice in terms of policy templates, notice templates, proven processes, and reference procedures that communities can adopt and extend as-needed.
The initial version of the PDK provided policy templates only, and its use required a level of expertise over which most communities did not avail. The revised 'version 2' of the PDK follows the target-audience approach of the trust framework, and is provided as an interactive set of pages with a common structure addressable through a clickable map. This new version is also dynamic, and allows continuous evolution also after the AARC TREE project concludes.

The AARC-TREE project is co-funded by the European Union under the HORIZON-INFRA-2023-DEV-01 call

Table of Contents

Executive Summary

Accessing, using, and operating AAI platforms and services that use federated access is inherently crossing administrative domains and organisational policy boundaries, as users access resources outside their home organisations. Whereas for trust within a single organisation there are many existing policy frameworks, these are lacking for the more complex collaborations, in particular for research activities where responsibilities are more distributed than in, for example, more conventional joint ventures.

A set of research collaboration policies is necessary to facilitate this trust in a scalable way, but crafting bespoke policy documents for each collaboration is time-consuming and complex. However, common patterns observed in collaborative research, such as those identified in the AARC Handbook and the Federated Identity Management for Research (FIM4R) communities, allow for a faster and easier pathway. This mirrors the layered approach of the AARC Blueprint Architecture (BPA) through separation of responsibilities to identified logical capabilities. For trust policies, these layers correspond to specific target audiences, where each organisation or group of organisations can take responsibility for part of the collaboration trust fabric.

The Policy Development Kit (PDK) is an on-line resource, contributing to the AARC Interoperability Framework, and containing policy guidelines, implementation suggestions, examples of procedures and processes, and template policies that can be adopted with minimal effort. This version is also 'dynamic' in that it allows continuous evolution also after the AARC TREE project concludes by virtue of allowing collaborative editing. While it is presented here through a written report, the Policy Development Kit itself can be found on-line:

https://aarc-community.org/policy/policy-development-kit/

As identified in the AARC informational guideline *Trust framework for proxies and Snctfi research services* (AARC-I082)[1], there is a foundational layer of 'research governance' that - for each collaboration - underpins the trust between the participant researchers, service providers, and infrastructures. While research governance *as such* is not purely a matter of authentication and authorization policy, adoption of federated AAI is frequently a first step towards structured collaboration, and hence the AAI policies are typically the first to initiate discussions on governance. To facilitate the discussions that arise, the Policy Development Kit contains some content and references to governance models and practices, even though research governance itself is out of scope of the AAI, and hence also out of scope of the PDK.

The policy development kit includes the *Snctfi* set of policies: 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. This serves a dual purpose: AAI platform providers can (self) assess their policy alignment and act as a trusted source of information for services connected to the infrastructure and site-local layers in the AARC BPA, and they can express to communities that the collaboration platform is aligned with the community good practice and (security) AARC Guidelines. For collaborations, Snctfi can serve as guidance when engaging their AAI platform operator(s) - as it is now common practice for mid-size and also larger collaborations to manage their collaboration with a specialised AAI platform operator.

The initial version of the PDK provided policy templates only, and its use required a level of expertise over which most communities did not avail. The revised 'version 2' of the PDK presented here follows the target-audience approach of the trust framework (AARC-I082) and is provided as an interactive set of pages with a common structure addressable through a clickable map. This specifically makes it easier for mid-sized communities to adopt and work with the PDK: it provides a structured approach where only those elements of the policy relevant to collaboration can be identified, and a set of practical steps guides new collaborations through the policy process.

[1] AARC Community (2025). Trust framework for proxies and Snctfi research services (AARC-I082),. https://doi.org/10.5281/zenodo.15506826 (https://aarc-community.org/guidelines/aarc-i082/)

This consolidated AARC Policy Development Kit and the Annex, representing the state of the PDK as of February 2026, has been simultaneously issued as Informational guideline AARC-I088[2].

1 The Policy Development Kit

"Policies are essential for setting expectations for participants in an Infrastructure, stretching from the Infrastructure management to the researchers themselves"
(AARC Policy Development Kit for Infrastructures, 2018)

In 2018 the AARC Community introduced the Policy Development Kit, the PDK, as a set of nine template policies covering collaboration membership management, privacy for access personal data, trustworthy service operations, and security incident response. A 'top level security' policy template was part of the toolkit to establish a chain of authority to allow intervention in case of security incidents - be they violations of the policies by real users or by external miscreants.

Policies are necessary to regulate and facilitate trust in a federated environment, and those policies are structurally different from those established for a single administrative domain. In case users, service providers, and management are all within the same organisation (or when service providers are managed under contract from that single organisation), there is ample reference material available. ISO, ENISA, NIST, and national cybersecurity centres and data protection authorities provide guidance and templates in abundance. However, for the federated case, inherently multi-lateral and without a-priori central management, no such ready guidance is available.

This is usually first identified when an authentication and authorisation infrastructure (AAI) is being established for a collaboration, or set up as a service for a collaboration hosting platform. Following the AARC Blueprint Architecture (BPA) the trust framework encompasses several layers, and the distribution of responsibility over these layers depends on the complexity (size) of the collaboration, the classification of the research data processed by the researchers, their field of research (e.g. is the research topic by itself sensitive), and whether or not generic AAI platform providers are engaged by the collaboration or service provider and data infrastructures.

The policy space that must be negotiated covers security measures, membership management, access data protection, acceptable use, and assurance. For both large and smaller collaborations, policy will appear a daunting or overly complex task if you start on your research collaboration journey.

The AARC 2019 Blueprint Architecture (BPA) introduced the 'community-first' approach to structuring the AAI interoperability layer. Using the community (collaboration) centric approach, the 2018 PDK ('version 1') consisted of 9 documents, with each policy template following a brief set of introductory remarks to establish scope and applicability of a policy document for the reader.

The nine templates were based on existing good practice and the then-novel developments in the AARC 'community first' BPA, and found ready use in larger, structured, collaborations that had strong central coherence and synergetic relations with their supporting e-infrastructures and service providers.

With the expanding use of the AARC BPA model, it was observed that a large number of mid-sized and smaller communities adopted an 'AAI-as-a-service' approach, leveraging specialist AAI platform providers as hosts for membership management, token translation, attribute aggregation, and identity linking. Analysis of the PDK version 1 also identified aspects where policy templates and procedures were too closely integrated (AARC-I082)[1].

[2] AARC Community *AARC-I088 Policy Development Kit Revision*, http://doi.org/10.5281/zenodo.18661187 (https://aarc-community.org/guidelines/aarc-i088/)

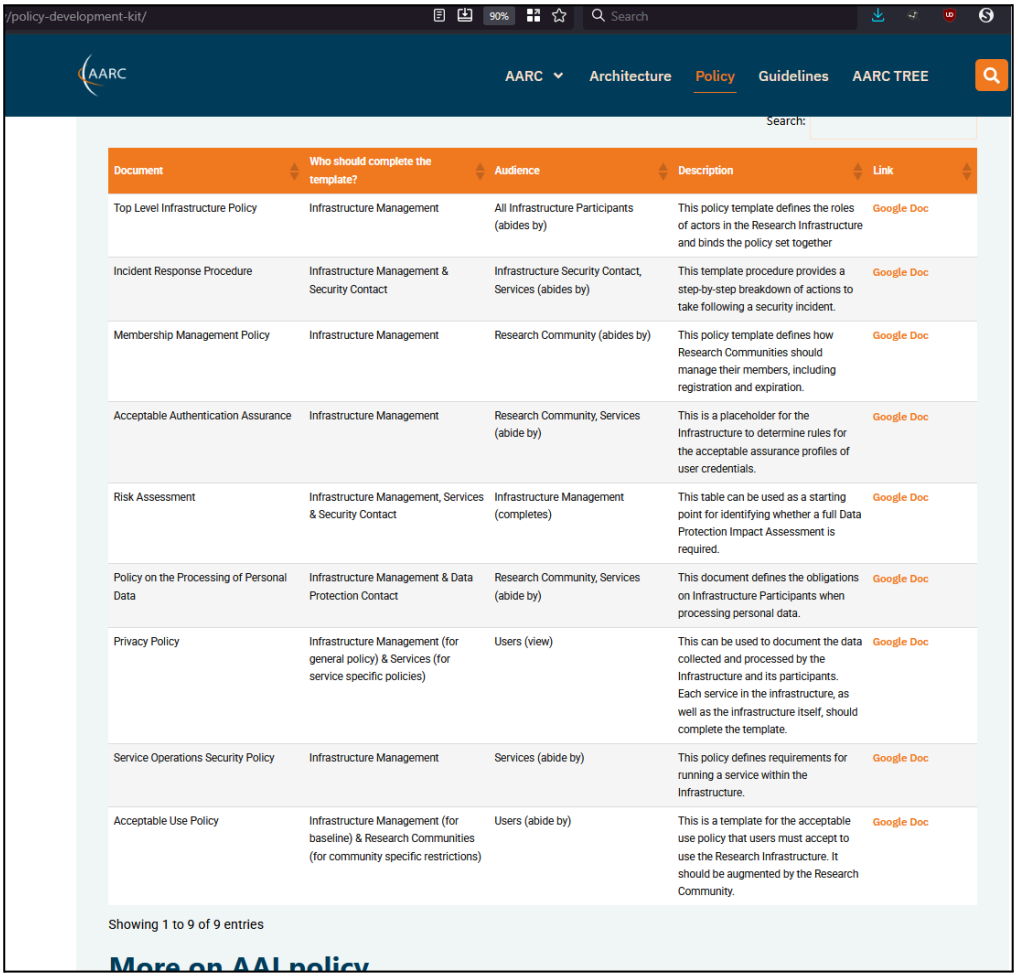

Figure 1: the original AARC Policy Development Kit documents assembled on the web site. The template documents themselves were hosted on a commercial cloud service to facilitate cloning by communities, bearing in mind that most collaboration at that time was based on sharing and collaborative editing of documents in that cloud service.

The trust framework was revised to be more suitable for mid-sized communities, better distinguish policies and procedures, and to identify target audiences that align with the layered platform architecture in the 2025 revision of the AARC BPA for enhanced effectiveness (AARC-G080)[3].

The revised Policy Development Kit (version 2) has now been released that incorporates the new trust framework, addresses the identified areas for improvement on version 1, and clarifies the *Snctfi* subset of policies and procedures that mid-sized communities can use to identify the most suited platform provider when hosting is considered.

[3] AARC Community *AARC-G080 AARC Blueprint Architecture 2025 - Initial Revision*, https://aarc-community.org/guidelines/aarc-g080/

2 Trust Framework and structure

In the AARC 'Trust Framework for proxies and Snctfi research services' [AARC-I082] the structure of an evolved Policy Development Kit was defined, taking into account the experience with the initial 2018 version, feedback from the Australian Access Federation in adopting the PDK for its national research computing infrastructure, and the intent of the AARC TREE project to make the policy guidance more easily accessible to mid-sized collaborations.

The trust framework developed in AARC-I082 therefore introduces the 'target audience' concept (only implicitly used in PDK version 1), matching the layers identified in the evolved AARC BPA2025, allowing mid-sized and early-stage collaborations to focus on those policies and procedures most relevant to them. Following AARC-I082, these are "*(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.*"

In order to achieve accessibility and broad adoption, the policies and procedures also align with existing community good practice to prevent duplication of effort. The AARC-I082 guidance explains:

"*[The] 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) 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.*"

The trust framework of AARC-I082 was used as the basis for the evolved Policy Development Kit version 2. The organisation of the trust framework, and hence the PDK version 2, is shown in Fig. 2.

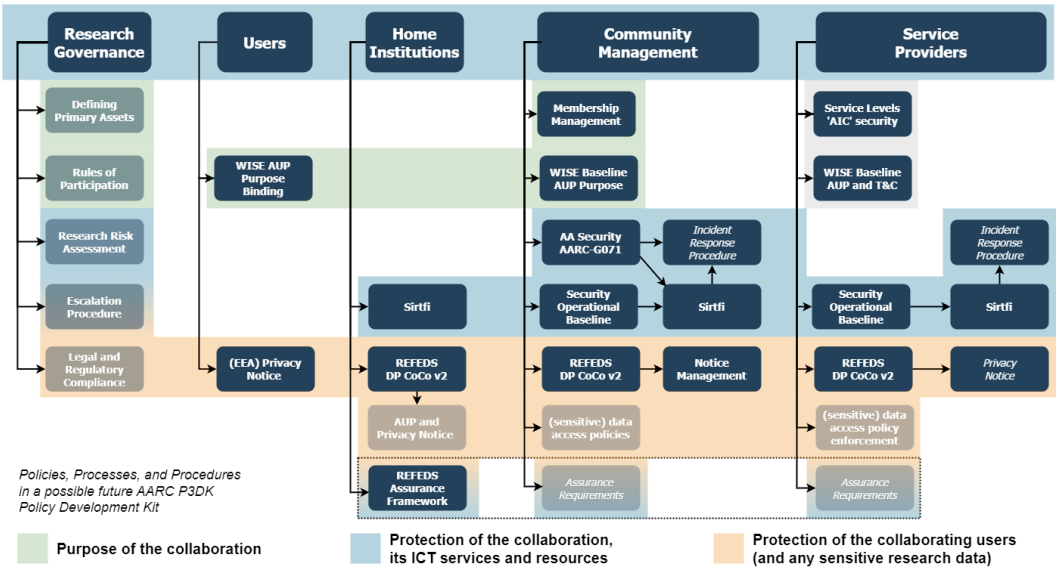

***Figure 2**: the structure of the Policy Development Kit, organised by target audience, and identifying the policies (roman type), procedures (italics) and their dependencies. The semi-opaque boxes indicate policy elements that are outside the scope of the AAI proper, but are relevant to the AAI policies. The Policy Development Kit provides information on all listed elements, though less extensively for the out-of-scope policy elements.*

3 The evolved Policy Development Kit

The evolved PDK ('version 2') is based on the structure of the trust framework described in AARC-I082 and the experience described therein on facilitating mid-sized and early-phase collaborations, but to a certain degree completeness of the PDK - cover of the entire trust framework - and the intent to make policy establishment a tractable proposition for mid-sized communities, pose conflicting requirements.

The PDK version 2 approaches this issue in three ways:

- by providing a minimal set of policies and practices - including fill-in templates - giving a low threshold to entry for mid-size and early-phase collaborations, specifically those that engage with an existing AAI platform provider;
- by using an audience-target based presentation of all trust elements and collaboration context (governance, federated authentication sources, service providers) for larger self-managed collaborations and for platform providers, including the relevant policy considerations, templates, and example procedures; and
- through presentation of the PDK as an on-line evolving resource, allowing new good practice and experience to be included continuously.

The presentation of the PDK version 2 as an on-line resource also enables a better user experience by a topical 'click-through' model for the trust framework, using the box-visualisation first used in AARC-I082 as an entry to the policy suite.

The entrypoint to the on-line resource is the AARC Community web site at

https://aarc-community.org/policy/policy-development-kit/

which links to the new PDK that is hosted on the AARC Wiki (for interactive use and continuous evolution) as well as retains links to the documents from the initial development kit (suitable for coordinated infrastructure-based collaborations).

The presentation of the policies and practices in the PDK version 2 on the AARC Wiki follows a common approach for each box: policy and process context ('why is this item relevant for this audience'), specific guidance ('what to do'), and in-depth references (the AARC guidelines and references to existing community good practice relevant to the policy or process).

Addressing mid-sized and early-phase collaborations, the evolved PDK includes a jump-starter of eight basic steps to quickly navigate the policy space and avoid the most common pitfalls. Tagged under 'Practical steps to getting started with Policies for a Research Collaboration', it lowers the barrier to entry by focussing on directly actionable items, and by assuming the use of a *Sncfti* aligned AAI platform provider for hosting the operational services and taking care of a large part of the trust fabric mechanisms (Fig. 3). By visual means ('folding' explanation boxes) it is more user friendly than the full trust framework.

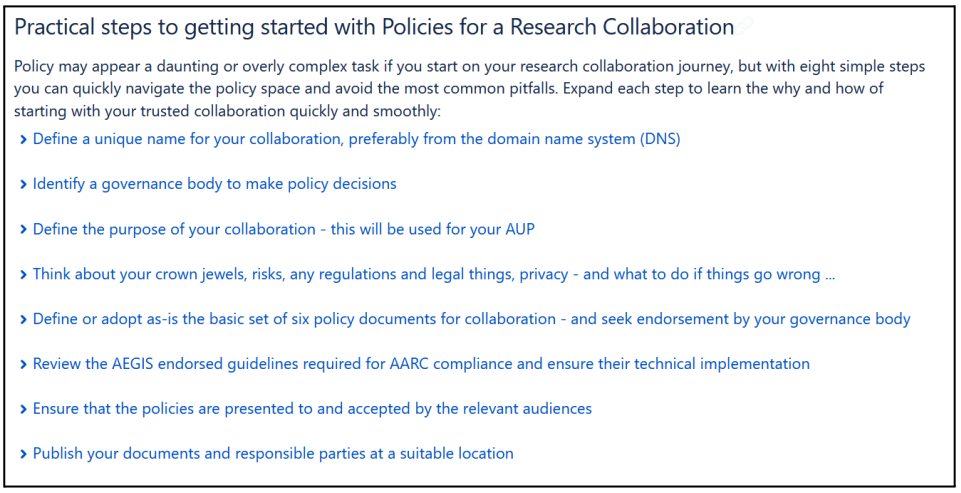

***Figure 3**: the jump-start policy set for small to mid-size and early-phase collaborations. Each step is foldable and contains a brief explanation on the 'reason why', a practical example, and links to the underlying resources. The eight step model is best used with a Snctfi aligned AAI platform provider.*

The 'folding' items contain compact guidance, also in a common structured 'why-how' format, as can be seen in the online version of the PDK (Fig. 4).

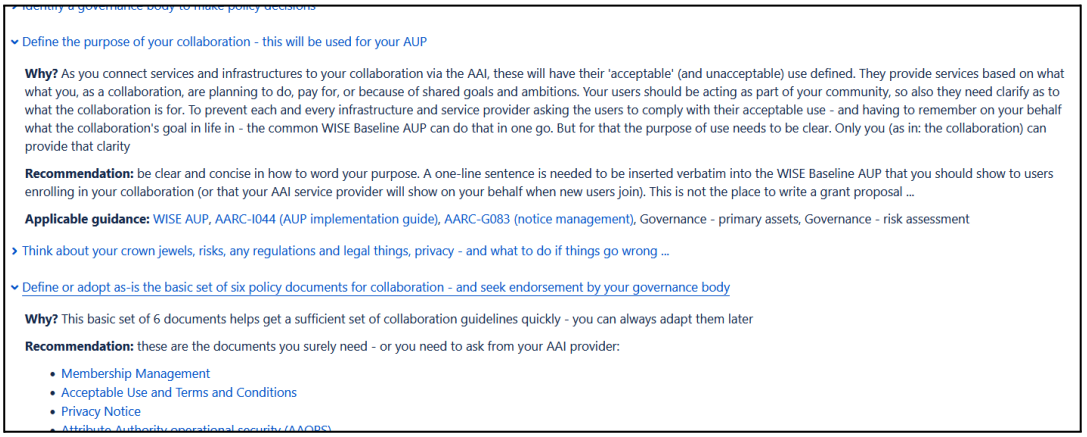

***Figure 4**: example of the policy guidance of the jump-start model. For each of the documents in the foundational set, a template is also provided as part of the PDK.*

For collaborations of any size or structure, six foundational policy statements are always needed. Of these, three are specific to each collaboration:

- how to manage membership - 'who is in, and why';
- acceptable use - the purpose of the collaboration in one sentence, needed to fill the placeholder of the WISE Baseline AUP and for requesting access to data and services; and
- a privacy notice - how does the collaboration intend to handle the access-personal data inherent in using services (a typical regulatory requirement, e.g. under the General Data Protection Regulation).

Two subsidiary policies and one procedure are also an almost universal prerequisite for using services:

- Operational cybersecurity: how does the collaboration (and collaboration management) react to IT security incidents, and how to share and receive information to mitigate incidents, and how to engage the collaboration in security readiness - complemented by a security incident response *procedure* with the specific details needed; and
- Operational integrity of the membership management services and the attribute authorities, important given the central role of the collaboration layer of the AARC BPA.

Collaborations that engage an AAI platform provider that aligns with *Sncfti* can rely on their provider to supply both operational cybersecurity as well as the attribute authority integrity, yet there remains a specific role for the collaboration management to act (suspend, inform, enforce) in case of security incidents and AUP violations.

By highlighting these six foundational documents, the PDK specifically aims to facilitate policy adoption for mid-sized communities that do not typically avail over the resources to write policy documents 'from scratch' based on the guidance of the AARC and community documents. Hence, for each of these a template is provided that can be copied and where placeholders can be filled in, as shown online and in Fig. 5.

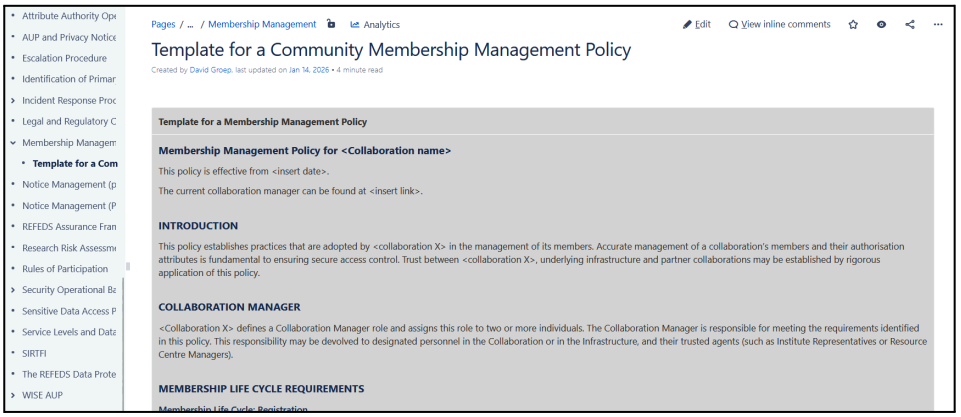

Figure 5*: for the basic set of six policies, the PDK provides ready-made templates for completion. The template shown is for the light-weight membership management policy, described in full in the informational AARC-I086.*

For large collaborations and the e-Infrastructures and research infrastructures that have specialised policy staff, the PDK on-line resource allows deep-linking to the individual policy and procedure guidance. This also facilitates re-use for training and support purposes.

A static rendition of the PDK version 2 is available in Appendix A, albeit without the interactive elements that makes the PDK version 2 easily accessible.

4 Conclusion and future evolution

The Policy Development Kit (PDK) evolution has resulted in a completely revised '*version 2*' of the PDK that implements suite of guidelines and templates for research and infrastructure communities following the new AARC Blueprint Architecture (BPA) and addressing the novel AAI platform structure that has emerged in the European Open Science Cloud (EOSC) and the broad adoption of managed AAI solutions for mid-sized collaborations.

The layered approach to AAI introduced in the evolved AARC BPA is reflected in the Policy Development Kit through the identification of specific audiences, and the role for 'hosted AAI platforms' identified as the recommended implementation path for AAIs in the AARC Handbook reflected in the updated *Snctfi* scope. Taking into account global feed-back on the incumbent policy suite, and more clearly separating policies from processes and procedures, the PDK version 2 has appeared as a good match to the EOSC AAI federation framework.

Policy guidance and good practice will never be 'complete', just as research and collaboration can never be 'complete', and just as the community best practice today is different from what was adequate guidance a decade ago. With thematic and national 'nodes' in the European Open Science Cloud, new collaborations, small and large, will emerge. The current PDK addresses research collaborations of the kind we see today: 'self-aware' to the extent that the collaboration reflects research granting practices and typical working groups in academia. In the future, the AAI may evolve to address the interaction of sole researchers with data and services - adding small- and micro-collaborations, expand with models where AAI is both centralised at the institutional level as well as devolved as in the current BPA collaboration-model, or incorporate education use cases alongside research at an equal footing. In the end, a lab course full of students is not fundamentally different from a short-lived collaboration.

As the context of collaboration-driven authentication and authorisation for research keeps evolving, architecture and policies will need to change. With continuous evolution in mind, the AARC PDK has been rendered as a collaborative set of curated web pages, so that the PDK can be revised and adapted to new circumstances, new operating conditions, and continuous alignment with the AARC BPA Architecture.

The continued maintenance of the PDK, and its contribution to the AARC Interoperability Framework, is ensured by the AARC Community, which is committed to ensuring evolution of the PDK.

Appendix A: Static Policy Development Kit v2

A static rendition of the Policy Development Kit version 2, as retrieved from the AARC on-line resources, is included on the following pages. The latest version of the Kit is available on-line at

https://aarc-community.org/policy/policy-development-kit/

The authors suggest that all users of the PDK use the on-line resource rather than relying on the static version included here.

Policy Development Kit

Accessing, using, and operating services for research in today's world, as a rule, is inherently distributed, where users access resources outside their home organisations. In this complex environment, the question of trust for users, resource providers, and infrastructures, becomes paramount.

A set of policy documents is necessary to regulate and facilitate this trust. These policies outline the operational measures undertaken by the infrastructure to properly provide services. The policies principally cover security measures, user management and data protection.Policy may appear a daunting or overly complex task if you start on your research collaboration journey, but with eight simple steps you can quickly navigate the policy space and avoid the most common pitfalls.

What is the Policy Development Kit?

The Policy Development Kit (PDK)) offers **policy templates** to provide a head start for Research Infrastructures that want to deploy the AARC Blueprint Architecture. The policies are there to provide a starting point, so that Research Infrastructures do not have to re-invent the wheel!

Alongside the updated AARC Blueprint Architecture 2025 the PDK has been revised to make it easier for collaborations to adopt and integrate with the new trust framework. The PDK now is an on-line interactive resource that can be viewed on the AARC Wiki pages. It is intended to be used in two scenarios:

- The practical steps listed above are insufficient to address the policy needs of your research infrastructures due to size or complexity
- The practical steps above are not relevant as you have a more specific role such as the operator of an Authentication Source

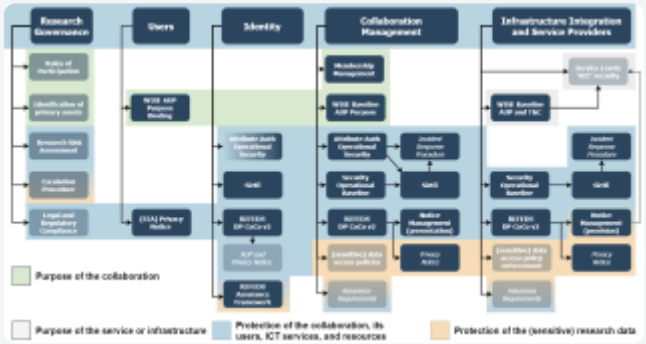

- 'Research governance' as a foundational area.
- 'Users' are (human) end-users who participate in a collaboration, are identified via
- 'Authentication Sources', i.e. external identity providers and the identity layer of the BPA, to be granted access by
- Collaboration Management', to
- 'Service and Infrastructure Providers'; in the BPA the infrastructure integration components, site-local integration components, and the actual service providers.

And be an early beneficiary of smoother trust and a new Snctfi, finding a provider of AAI platforms that already implements the key trust elements!

In addition to the templates, the AARC community over the years has also developed:

- Learn more about policies for the AARC Blueprint Architecture and videos from this course are also available on the AARC video playlist on YouTube GÉANT tv.
- A Training module on the GDPR Code of Conduct
- A Training Module on Policies for processing personal data
- The Policy guidelines offer more detailed advice on specific topics.
- The Snctfi framework- a 'Scalable Negotiator for a Community Trust Framework in Federated Infrastructures'.
- The Sirtfi framework to identify trusted and operationally secure partners in a federated authentication and authorisation environment.

▸ Policy Development Kit version 1 for larger and structured communities

aarc-contacts@lists.geant.org

Funded by the European Union

The work leading to these results has received funding from the European Union (GAP 101131237) and other sources. The contents of this publication is the sole responsibility of AARC and does not necessarily reflect the opinion of the European Union.

The AARC Policy Development Kit

Welcome to the Policy Development Kit!

The AARC Research Collaboration model is both the set of technical guidelines and interfaces in the AARC Blueprint Architecture (BPA) as well as the trust framework that helps research collaboration bridge across domains, sectors, and borders: the guidelines for end-to-end trust across the components for collaboration management, user privacy, identity assurance, and operational security. This Policy Development Kit (PDK) helps new and existing collaborations to build those trust relationships into their AAI and benefit from the combined experience of the infrastructures, collaborations, researchers and research managers, and trust and security engineers to quickly build that trust in our AARC connected world.

Practical steps to getting started with Policies for a Research Collaboration

Policy may appear a daunting or overly complex task if you start on your research collaboration journey, but with eight simple steps you can quickly navigate the policy space and avoid the most common pitfalls. Expand each step to learn the why and how of starting with your trusted collaboration quickly and smoothly:

Why? When your users connect to infrastructures and services, the services will need to identify the users as belonging to your group. And as you work together across sectors, you will want users with the same name but from different communities to work together. Similarly, if you use a shared AAI provider, for example based on the Snctfi guidelines, also there your collaboration should not be mixed up with others.

Recommendation: use a name that is almost certain to be unique globally, and pick a name that is not prone to changes, avoiding project naming for instance. The domain name system (DNS) is a good starting point, for example "he3epp.nikhef.nl" for a collaboration for studying the ^{3}He(*e*,*e*'*pp*) reaction at Nikhef, or "atlas.cern" for the global ATLAS collaboration located at CERN. Note that while the domains should be permanently assigned, you don't necessarily need a web site or email addresses with this domain. Uniqueness is enough. By using a DNS name, it fits easily in the 'scope' component of many AAI protocols like OpenID Connect and SAML.

Applicable guidance: AARC-G069, PDK Membership Management guidance
Why? While management of the user directory, or responding to compromised credentials and questions from your service providers may appear a technical task, these often involve decisions that deal with your 'primary assets': the research data you work with, the processes that keep your community together and wholesome. A handful of people working together can solve issues that arise in an ad-hoc fashion, but as soon as you grow a bit further, structured communication becomes essential. "Why was I suspended?" "Why can't I join your group ... you have a service I need (and by the way, I just want to use it, not contribute)?". You need a body to take up that authority. Unfortunately, the AAI is often the first time you hit these hard questions!

Recommendation: A principal investigator, research group chair, or faculty dean makes a good starting point for a local governance body. If your community becomes larger, write down rules of participation, draft a memorandum of understanding, or have a written collaboration agreement in place. Often there is already something there: many public research projects require having a collaboration or grant agreement. There is usually a governance structure in there, which can be re-used here. No need to re-invent the wheel within your community.

Applicable guidance: governance as such is out of scope of the PDK, but look at your project agreements, department model, or grant office to find a suitable and effective solution.
Why? As you connect services and infrastructures to your collaboration via the AAI, these will have their 'acceptable' (and unacceptable) use defined. They provide services based on what what you, as a collaboration, are planning to do, pay for, or because of shared goals and ambitions. Your users should be acting as part of your community, so also they need clarify as to what the collaboration is for. To prevent each and every infrastructure and service provider asking the users to comply with their acceptable use - and having to remember on your behalf what the collaboration's goal in life in - the common WISE Baseline AUP can do that in one go. But for that the purpose of use needs to be clear. Only you (as in: the collaboration) can provide that clarity

Recommendation: be clear and concise in how to word your purpose. A one-line sentence is needed to be inserted verbatim into the WISE Baseline AUP that you should show to users enrolling in your collaboration (or that your AAI service provider will show on your behalf when new users join). This is not the place to write a grant proposal ...

Applicable guidance: WISE AUP, AARC-I044 (AUP implementation guide), AARC-G083 (notice management), Governance - primary assets, Governance - risk assessment
Why? "Bad things can happen to good science" (1), and while you may not think of it at first, the data, ways of working, and collections created in your collaboration are valuable and deserve protection. External cybersecurity attacks of course come to mind, but in many cases inadvertent accidents happen and are at least as big a risk. Identifying your 'primary assets' (or the 'crown jewels' of the collaboration) helps you to identify where you need extra protections, and how to prevent deletion, changes, or loss of data ... and people. There may also be legal and regulatory reasons to apply controls through your AAI. They can be in the research data itself, like medical and patient data, dual-use goods and knowledge, commercially confidential data, or ethical reasons on human research or in the Nagoya Protocol.
And protect your own peers in the collaboration: they should know how their name, email address, or roles that are used in the AAI are protected. And for some sensitive or high-profile research, also names and contact info needs to be protected!

Recommendation:

- identify your crown jewels or primary assets. These are *not* computer things, nor your AAI, but research data, research processes, and knowledge.
- define your rules of participation and the escalation procedure in case of non-compliance. If you are dealing with sensitive subjects, or sensitive research, consider the risks and what measures in your AAI can help. If you use a hosted AAI, discuss the conditions and guarantees with your (Snctfi-ed) provider.
- Adopting the REFEDS Data Protection Code of Conduct - if it can apply to your research collaboration - to clarify privacy for access personal data (the personally identifiable information that results from your collaborators usin services, infrastructure, and the AAI itself)
- Do a (brief) risk assessment to check the impact of inadvertent or malicious events. Prioritise risks your crown jewels, and keep in mind that controls should not make your primary mission impossible!
- Does your collaboration work with human, societal data, or collects questionnaires? Is your research likely to be classed as dual-use or export restructured? Does the research, or your collaboration users, touch on knowledge safety? Is approval by medical/ethical commissions needed? Are you dealing with biodiversity or genetic resources subject to the Nagoya Protocol? Do a specific risk assessment or ask your institution for guidance.

Applicable guidance: REFEDS Data Protection Code of Conduct, (1) Open Science Cyber Risk Profile (Sean Peisert et al, TrustedCI), ITSRM2 (risk management), Privacy Notices
Why? This basic set of 6 documents helps get a sufficient set of collaboration guidelines quickly - you can always adapt them later

Recommendation: these are the documents you surely need - or you need to ask from your AAI provider:

- Membership Management
- Acceptable Use and Terms and Conditions
- Privacy Notice
- Security Operational Baseline
- Incident Response Procedure
- Attribute Authority operational security (AAOPS)

Applicable guidance: REFEDS privacy notice, UK-IRIS example privacy notice, EOSC, UK-IRIS security policies, AARC-I051 "federated incident response procedure"
Why? Assurance means both knowing if the person on the other side is indeed the same user that you know, but also includes *identity* assurance, verifying that the person is indeed the one they claim to be: name and affiliation being the most visible elements. How strong that assurance needs to be depends on the type of research and the collaboration risk assessment.
And for how long do you trust that the activity is still the intended one, and from the same user?

Recommendation: review the technical and policy guidelines endorsed by the AAI providers and infrastructures in AEGIS, the AARC Engagement Group for Infrastructures:

- Identify your assurance requirements, following AARC-G031 "evaluation and combination of the assurance of external identities".
- Identify suitable token lifetimes, using AARC-G081 "Recommendations for Token Lifetimes"

Applicable guidance: Assurance Requirements, AARC-G031 evaluation and combination of the assurance of external identities, AARC-G081 "Recommendations for Token Lifetimes"
Why? You realise you need to enforce a policy only once things do not 'go as planned' - and having the discussion on acceptance at that point is rather late. And how can users, for instance, know what they are allowed to do with the research data, or when to ask for additional roles and group membership from membership management?

Recommendation: the Policy Development Kit identifies five different 'audiences': governance, your users, the user home organisations and identity providers, the AAI management of the collaboration, and the infrastructures and service providers that control and host data, computing capacity, and the data transfer networks. Make sure all of then can access and understand your policies and processes, can work with you when you execute procedures for incident response, and engage with Sirtfi and security readiness exercises.
Notice Management helps to communicate with users in a coordinated way, and prevent needless pop-ups that interrupt their workflow. If you engage an AAI service provider, they may be able to help with communication.
If you used a DNS name for your community, and you can resolve the domain name to point to a web site, that is a great place to present your collaboration and provide questions & answers as well as contact details.

Applicable guidance: AARC-G083 on notice management, WISE Baseline AUP and AARC-I044, Privacy Notices, REFEDS DP CoCo v2, Membership Management
Why? Presenting policies and practices is one thing, but the AARC Blueprint Architecture also introduced a (chain of) AAI platforms or 'proxies' that augment, translate, or otherwise munge information about users and 'sources of authority'. Both for authentication sources and for service providers, it places intermediates in the chain of trust, and the longer the chain is, the more this trust will be diluted. Transparency through documentation can help retain that trust. And at the same time make it easier for the collaboration to engage with the users regarding the AAI. If identity is not bound to the user but to the user's home organisation (employer, university), the home organisation may be reluctant to make any claims for the authentication, even for trivial ones like name and email address (the 'personalised access' attributes that are foundational for research and scholarship). Or refuse to partake in authentication at all.

Recommendation: publish your policies, but especially your contact information, in a place where users, relying parties, and home organisations can find it. If you chose a DNS-based community name, and you can resolve the domain name to point to a web site, that is a good place to present this information. And if confidentiality is needed, you may have your own AAI to help you!

Applicable guidance: AARC-G071 (AAOPS), AARC-G083 on notice management, REFEDS Research and Scholarship, REFEDS Personalised Access,

Maturing your trusted collaboration policy and good practice

The following diagram summarises the full ecosystem of policies relevant for a research infrastructure. It is intended to be used in two scenarios:

1. The practical steps listed above are insufficient to address the policy needs of your research infrastructures due to size or complexity
2. The practical steps above are not relevant as you have a more specific role such as the operator of an Authentication Source

The Policy Development Kit (PDK) version 2 identifies five main target audiences, functionally following the **AARC 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. 'Authentication Sources', i.e. external identity providers and the identity layer of the BPA, to be granted access by
4. 'Collaboration Management', to
5. 'Service and Infrastructure 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 – that are endorsed by AEGIS and considered 'standards track'. Policies are endorsed by the organisation at the appropriate level of management, and express a commitment of adherence by the organisation's management. These are indicated in a roman font in the graphic below.

Indicated in *italics* in the diagram *processes and procedures,* being templates, are reference implementations where we assume these to be specialised for specific deployments.

The semi-opaque elements are relevant, but fall outside of the scope of the PDK, which targets the authentication and authorisation infrastructure. But even if, for example. identifying the 'why and what' of your research collaboration (your 'primary assets') may not be AAI per-se (and hence greyed-out), it is very useful to know that before embarking on your AAI journey!

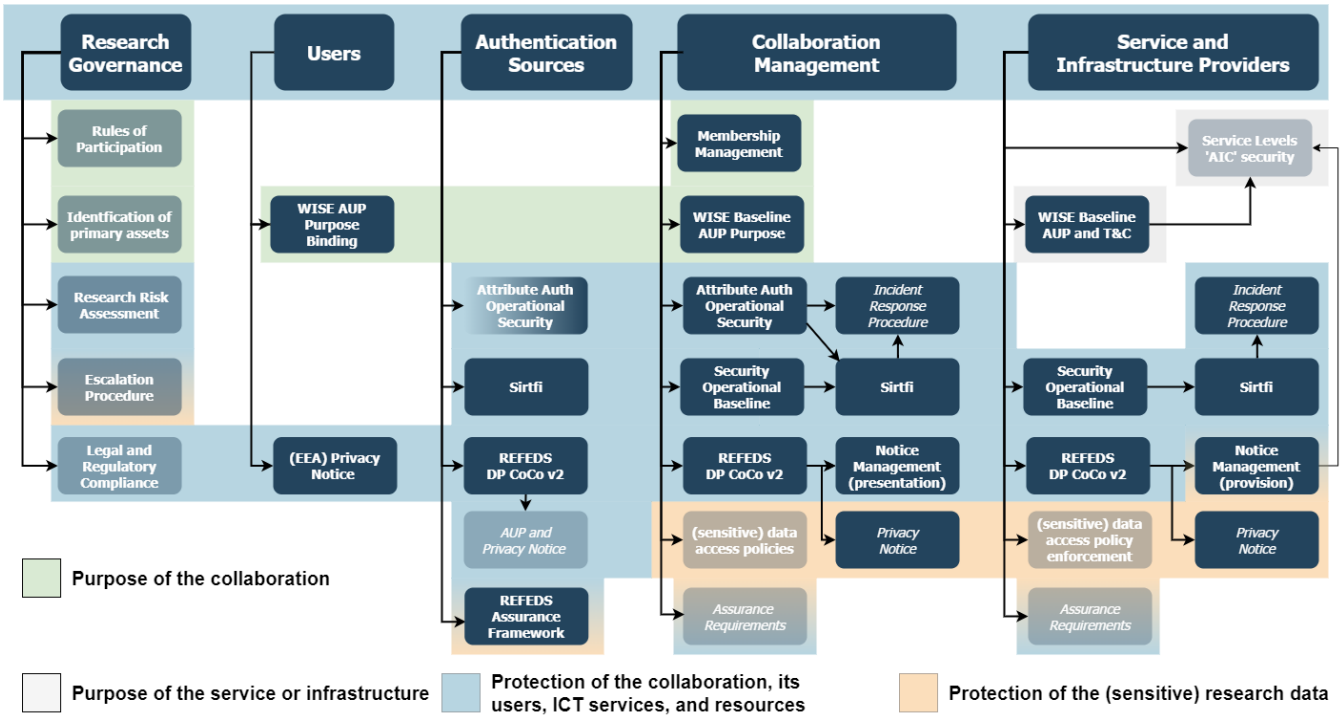

Snctfi, operational policies, and AAI service providers

Smaller and mid-sized communities may opt to offload some of the more complex aspects of authentication and authorisation to dedicated AAI service providers. And if you operate your own AAI core components, both your users and resource providers may want to have some assurance about the trust and security posture of your AAI platform. The *Snctfi* suite is the set of assessable and verifiable policies and procedures in the PDK that AAI platform providers can use to make the trustworthiness of their systems transparent to users and relying parties alike.

Like *Sirtfi* for security incident response, *Snctfi* provides a self-assessment framework, but having this assessment peer reviewed brings several benefits. For one, it increases the trust others have in your platform and your assessment, making it easier for 'as-a-service' operators to engage with new collaborations and infrastructures. And it brings advantages to yourself as well, as you can compare notes with your peers and become better together through shared learning.

AARC does not endorse any specific AAI platform or platform provider. By asking *Snctfi* specific information you can inform yourself about the suitability of the provider of your choice, and work with them to ensure your bases are covered by a secure, resilient, and interoperable AAI.

- Learn more about Snctfi in **AARC-I082** "Trust framework for proxies and Snctfi research services"

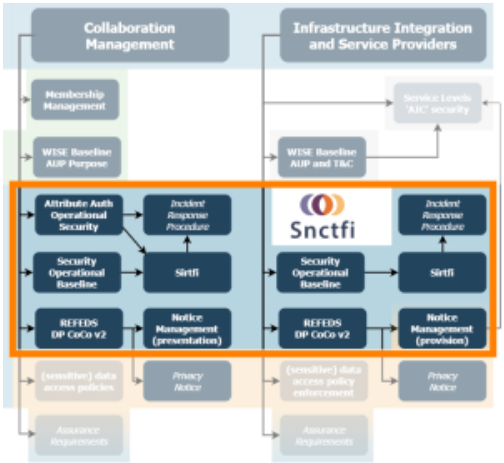

The Snctfi subset of assessable and verifiable policies and procedures in the PDK (from: AARC-i082)

Background to the Policy Development Kit

The first AARC Policy Development Kit, released in 2017,
comprised a set of nine reference documents (mostly templates) addressing the construction and operation of community AAIs in the original AARC "2019" Blueprint Architecture, based on the Community First Approach. A mix of policies and procedures, its primary audience was primarily larger-scale research collaborations, expected to review, augment, and specialise the templates for their own use. With the policy development kit being created prior to or in parallel to other work in the community at large, it duplicated some aspects (privacy in REFEDS DPCoCo, or incident response work parallel to the eduGAIN Security Handbook), while being overly complex for smaller collaborations. Work by the Australian Access Federation, the AARC Community, and in REFEDS, WISE, IGTF, and the e-Infrastructures helped restructure the PDK into the v2 model presented above.

An analysis of the improvements required on PDK v1 is included in the informational document **AARC-I082** "Trust framework for proxies and Snctfi research services" (doi:10.5281/zenodo.15506826)

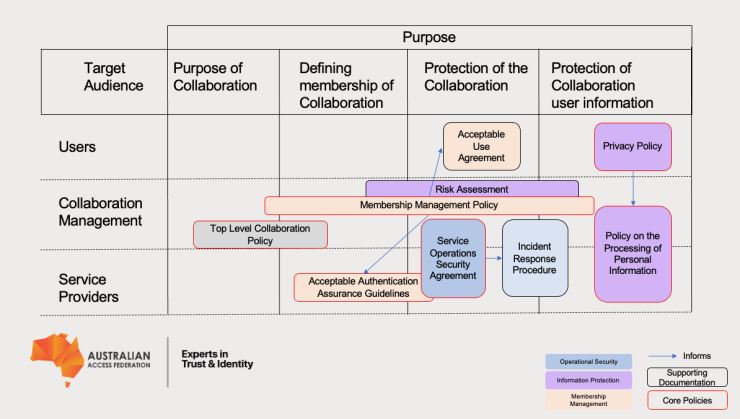

Analysis of the first (2017) version of the Policy Development Kit by the Australian Access Federation. The original set of templates, both policies and procedures, were primary targeting larger research collaborations.

(EEA) Privacy Notice

For those with data hosted in or users from countries covered by GDPR or UK GDPR a privacy notice is a requirement. The minimum requirements for this are:

- Identity of the data controller
    - Including contact information
- Purposes of data collection
- The legal basis used for processing of data
- The types of personal data being collected
- Who the data is being shared with
- How long the data is being kept for
- How individuals can exercise their rights over their own personal data
- How consent can be withdrawn
- Where data is transferred internationally, if this is outside of the EEA how the personal data is safeguarded must be covered

This is one of two documents in the PDK that MUST be presented and agreed to by all users.

Personal data is any data set that can be taken from or combined with any source that can be used to determine information about a natural person

There are eight data protection rules that each data controller must ensure are followed:

- Personal data must be processed legally and fairly.
- It must be collected for explicit and legitimate purposes and used accordingly.
- It must be adequate, relevant and not excessive in relation to the purposes for which it is collected and/or further processed.
- It must be accurate, and updated where necessary.
- Data controllers must ensure that data subjects can rectify, remove or block incorrect data about themselves.
- Data that identifies individuals (personal data) must not be kept any longer than strictly necessary.
- Data controllers must protect personal data against accidental or unlawful destruction, loss, alteration and disclosure, particularly when processing involves data transmission over networks. They shall implement the appropriate security measures.
- These protection measures must ensure a level of protection appropriate to the data

To use the explanation given by the Information Commissioner's Office, a data controller is "a person who (either alone or jointly or in common with other persons) determines the purposes for which and the manner in which any personal data are, or are to be, processed". A data controller is the responsible party that must ensure that all processing of personal data complies with the GDPR. Failure to do so may result in legal repercussions. Data processors, on the other hand, process personal data solely under the direction of a data controller, who decides what personal information will be kept and to what uses it may be put.

Templates of privacy notices

REFEDS DPCoCo v2 example

The REFEDS DPCoCo provides a tabular template for *service providers* to present their privacy notice. This 12-point notice ticks all the requirements of the GDPR in a way that is consistent and can (almost) be parsed by machines, although it is not very readable by people. The advantage of it is that all service providers that use the REFEDS DPCoCo template can be compares, and it makes it 'easier' to create combined notices (e.g. in the line of AARC-G083):

- https://wiki.refeds.org/display/CODE/Privacy+Notice+Template

WLCG Example

The Worldwide LHC Computing Grid (WLCG) notice is an example of a federated infrastructure where there is no single point of control and no single controller. It relies on the concept of controller-to-controller transfer of data and the fact that all parties (service providers and AAI platform) are bound by a common policy framework, overseen by the WLCG Management Board. However, a formally liable monitoring body cannot be identified - this is a very common case for research collaborations. It follows the 'BCR-like' model described in AARC-G016:

- http://wlcg-docs.web.cern.ch/wlcg-docs/?dir=policy/security

Jisc Example

The UK research and education organisation Jisc uses a privacy notice that emphasises readability and - through folding text sections - helps end-users understand how their data is used. Aimed at *external data subjects*, it targets the same audience type as many research collaborations, while also fulfilling all the GDPR and UK ICO requirements:

- https://www.jisc.ac.uk/website/privacy-notice

REFEDS DP CoCo Document development Guidance

The guidance on this page works along side the REFEDS Data Protection Code of Conduct which should be asserted in the privacy policy provided

Questions to ask yourself when defining this policy:

- Who or what is your Data Controller?

- Will your Research Community have a Data Protection Officer?
- Which information do you need to collect on the user? Is this minimised?
- Specific data collected by each service may vary. Can your Infrastructure provide a template statement for all services?

Resources

- GDPR - https://gdpr-info.eu/
- https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home
- AARC Guidance for exchange of personal information - https://aarc-community.org/guidelines/aarc-g016/
- AARC Data protection impact assessment - https://aarc-community.org/guidelines/aarc-g042/

Privacy Notice Template

Template for a Privacy Notice

Privacy Policy

This policy is effective from <insert date>.

Name of the Service	SHOULD be the same as mdui:DisplayName
Description of the Service	SHOULD be the same as mdui:Description
Data controller and a contact person	You may wish to include the Data Controller defined for the Infrastructure, rather than per-service
Data controller's data protection officer (if applicable)	
Jurisdiction and supervisory authority	The country in which the Service Provider is established and whose laws are applied. SHOULD be an ISO 3166 code followed by the name of the country and its subdivision if necessary for qualifying the jurisdiction. How to lodge a complaint to the competent Data protection authority: Instructions to lodge a complaint are available at...
Personal data processed and the legal basis	1. Personal data retrieved from your Home organisation: * your unique user identifier (SAML persistent identifier) * * your role in your Home Organisation (eduPersonAffiliation attribute) * * your name * * ... 2. Personal data gathered from yourself * Logfiles on the service activity* * Your profile * … * = the personal data is necessary for providing the Service. Other personal data is processed because you have consented to it. Please make sure the list A. matches the list of requested attributes in the Service Provider's SAML 2.0 metadata.
Purpose of the processing of personal data	Don't forget to describe also the purpose of the log files, if they contain personal data (they usually do)
Third parties to whom personal data is disclosed	Notice clause of the Code of Conduct for Service Providers. Are the 3rd parties outside EU/EEA or the countries or international organisations whose data protection EC has decided to be adequate? If yes, references to the appropriate or suitable safeguards.
How to access, rectify and delete the personal data and object to its processing	Contact the contact personal above. To rectify the data released by your Home Organisation, contact your Home Organisation's IT helpdesk.
Withdrawal of consent	If personal data is processed on user consent, how can he/she withdraw it?
Data portability	Can the user request his/her data be ported to another Service? How?
Data retention	When the user record is going to be deleted or anonymised? Remember, you cannot store user records infinitely. It is not sufficient that you promise to delete user records on request. Instead, consider defining an explicit period. Personal data is deleted on request of the user or if the user hasn't used the Service for 18 months
Data Protection Code of Conduct	Your personal data will be protected according to the Code of Conduct for Service Providers, a common standard for the research and higher education sector to protect your privacy

Assurance Requirements

In order to protect its assets, services providers, infrastructures, and collaboration alike needs to be able to authenticate, identify, and trace users granted access to their services. The authentication and identification must be sufficient to meet their requirements and match their risk appetite, legal and contractual obligations, and those of their suppliers. Collaborations in particular play a central role in the AARC Blueprint, since they bridge the user and infrastructure domains, and their components are often 'opaque' by translating both security mechanisms and policy domains.

Not having sufficient assurance about users may appear 'simpler' at first, but may be a long-term liability, as service providers and infrastructures need to protect themselves from harm in case of security incidents and then will suspend or terminate service to the collaboration. Infrastructures and service providers will more readily define acceptable assurance levels in order to meet their information security management system and risk assessment. Similarly, data providers, especially those providing access to sensitive data, will have specific policies in place to know exactly to whom they are providing access - patient records and the binding of permits for data access from the HDABs come to mind. Collaborations that access services must have assurance requirements in place as well, and these must meet or exceed those of their connected relying parties!

Defining assurance requirements based on risk

If, as a collaboration, you are using a shared AAI services platform, your AAI provider may have already put in place minimum assurance requirements and may offer 'standard' templates for step-up assurance, both for the authenticator (multi-factor) as well as for identity vetting (photo-ID, validation of data in authoritative databases, etc.). Ask your service provider about both minimum and optional assurance features: they could either be too high or too low for your collaboration use cases.

What to ask for

Express the collaboration and infrastructure requirements for identity vetting, freshness of affiliation, and 'uniqueness' of the identifiers you need, in terms of the REFEDS Assurance Framework (RAF) components. By doing so, as a collaboration you can leverage the collective effort to express assurance across many identity sources and peer collaborations - saving a lot of bi-lateral negotiations.

There are two 'collapsed' assurance levels, defined to be the most useful combinations for federated and collaborative use cases:

- Cappuccino: there is *a* unique identifier (try for subject-id first, then pairwise-id), medium-level identity assurance (comparable to eIDAS low and IGTF BIRCH and CEDAR levels), and the 'affiliation' attribute reflects the status no longer than one month in arrears. This level can be combined with single-factor authenticator strength (http://refeds.org/sfa); and
- Espresso: there is *a* unique identifier (try for subject-id first, then pairwise-id), high-level identity assurance (comparable to eIDAS substantial and Kantara IAF level 2), and the 'affiliation' attribute reflects the status no longer than one month in arrears. This level is best combined with multi-factor authenticators (http://refeds.org/mfa).

but you can also inspect the assurance components: ID, IAL, and ATP. If attribute freshness of one-day-or-better is required, inspection of ATP is necessary since none of the collapsed assurance level specifically requires this quality.

Resources

- REFEDS Assurance Framework
- TrustedCI Open Cyber-Security Risk Profile
- REFEDS Single-Factor Authentication (SFA) and Multi-Factor Authentication (MFA)
- AARC-I050 Comparison Guide to Identity Assurance Mappings for Infrastructures
- (older example) EGI Policy on Acceptable Authentication Assurance

Attribute Authority Operational Security

Attribute Authorities (AAs) play one of the most critical security roles in the infrastructure. The data they issue and information they assert must be highly trusted by the parties relying upon it. To that end, AARC recommends that certain practices be adopted by the operators of such services: AARC-G071 Guidelines for Secure Operation of Attribute Authorities. The requirements listed include best practices in encryption, hosting environments, logging and attribute management to name but a few.

Collaborations can either host their own AA, or - more commonly - engage the services of an AA operator or an AAI platform to host their collaboration structure. In this latter case, most of the onus of securing and operating the AA falls on the operator, and implementing the AA service is part of the *Snctfi* r equirements.

For the AA platform operator and those collaborations hosting their own service, to make safe authorisation decisions, Relying Parties need to be able to identify and trust the issuer or provider of an attribute assertion, and know to which Collaboration it pertains. In a typical scenario, a Collaboration designates one or more AA Operators to operate AAs, and informs Relying Parties of any related metadata necessary for Relying Parties to connect to or use the AA. The attributes are securely held by the AA and delivered on request to authorised Relying Parties, either directly or by way of the user. Authentication sources and Collaboration Management should abide by the minimum requirements and recommendations for the secure operation of Attribute Authorities [see AARC-G071 Resources], and similar services providing statements for obtaining access to Infrastructure services.

These attributes may be aggregated with identity assertions, such as delivered from a directory or group management system, or with attribute or capability tokens as asserted by an AARC BPA Proxy.

Stated compliance with the AAOPS guidelines may help to establish trust between the Collaboration and its AA, and Relying Parties. In the interest of scalability, these guidelines are intended to facilitate the assessment of AA Operators rather than individual AAs or Collaborations. The document does not provide guidance on the management (life cycle, technical implementation, exchange protocols etc.) of attributes nor the processes by which attributes are entered into the AA.

How should I adopt AARC-G071?

In AARC-G071 Guidelines for Secure Operation of Attribute Authorities the requirements are included in purple boxes, with additional information included around it. Importantly a distinction is made between "pull" and "push" attribute assertion.

- attribute authorities that permit binding of properties to entities by means of lookup in which the entity whose properties are sought is the key in the look-up (‘**pull model**’)
- attribute authorities that issue (usually integrity-protected and, optionally, confidentiality-protected) statements in which attributes are asserted (‘**pu sh model**’)

If you are running your own AAI should go through each requirement and analyse if your infrastructure supports it. In many cases, Research Communities choose to outsource their AAI to a trusted operator who should also support these requirements.

Recommendations:

- AA Operators should abide by the requirements of AARC-G071, to attain and maintain the trust of their relying parties.

Resources

- AARC-G071 Guidelines for Secure Operation of Attribute Authorities
- Review and self-assessment of AA operators complying with AARC-G071 (IGTF)
- AARC-I086 Membership Management Policy Development
- Light-weight Collaboration Management - explainer
- Template Membership Management Policy for Infrastructures and self-structured collaborations
- Template Membership Management Policy for light-weight and hosted collaborations

AUP and Privacy Notices for Authentication Sources

Almost all organisations will present their users with an AUP and Terms and Conditions. These usually include both professional and (limited) personal use, and may include explanations about monitoring and privacy.

Collaboration-friendly organisations are recommended to add the WISE Baseline AUP 10 points to their AUP, to facilitate access to shard resources and services that align with the WISE baseline.

As part of the released attributes or claims, they can then add the Baseline AUP identifier ("https://wise-community.org/wise-baseline-aup/v1/") to the voPersonPolicyAgreement set of values.

Resources

- WISE Baseline AUP
- Example AUP from UK-IRIS
- Organisational AUP example

Escalation Procedure

You realise you need to enforce a policy only once things do not 'go as planned' - and having the discussion on acceptance at that point is rather late. And how can users, for instance, know what they are allowed to do with the research data, or when to ask for additional roles and group membership from membership management?

Who decides when things go south?

The rules of participation define who has the authority to decide, and to suspend or cut-off participants, in case the 'regular' processes break down. While the initial thoughts may revolve around security incidents and who has been granted the 'power' to intervene and stop the incident, there are other times when you need that control, and that control can be both inclusive or exclusive. For example if by standard processes former employees of an organisation are removed from the groups and their roles revoked, there may be good reasons to keep a specific person 'in the collaboration' for a while longer. Why can then trigger a hardship clause in the rules of participation? And if there is no such clause, who is authorised to help? If a collaboration breaks up, or essential rules are violated, who can suspend access, and is there a process?

A hardship clause also helps to retain proper documentation for any exceptions, instead of policy being silently bypassed. Some examples are provided below.

Make sure all of them can access and understand your policies and processes, can work with you when you execute procedures for incident response, and engage with Sirtfi and security readiness exercises.

"Where the enrolment or approval process described leads to obvious injustice and hardship due to a implementation of the process being unduly hard, a **time-limited and documented** exception may be implemented after the explicit approval at the appropriate level of Management.
Any such exception must be followed up by immediate and concrete steps to address the deficiencies created within a limited time period commensurate with the discrepancy induced. The invocation of the hardship clause MUST not compromise the integrity or trustworthiness of the system with respect to any participants."

"In exceptional circumstances it may be necessary for participants to take emergency action in response to some unforeseen situation which may violate some aspect of this policy for the greater good of pursuing or preserving legitimate e-Infrastructure objectives. If such a policy violation is necessary, the exception should be minimised, documented, time-limited and authorised at the highest level of the Management commensurate with taking the emergency action promptly, and the details notified to the e-Infrastructure Security Officer at the earliest opportunity" (WLCG and EGI top-level policy)

How does escalation fit into an AARC BPA compliant infrastructure?

The AAI is used to grant and revoke access to resources, and the decisions taken during escalation have to be effectuated in the AAI at the appropriate layer. Often that is in the collaboration platform or at the infrastructure layer.

Resources

- eInfrastructure top level policy for EGI and WLCG

Identification of Primary Assets

Apart from enabling access for your collaborators, the AAI also has a role to play in protecting your collaboration and its data. "Bad things can happen to good science" (1), and while you may not think of it at first, the data, ways of working, and collections created in your collaboration are valuable and deserve protection. Identifying your 'primary assets' (or the 'crown jewels' of the collaboration, as MITRE would call them) helps you to identify where you need extra protections, and how to prevent deletion, changes, or loss of data ... and people. And protect your own peers in the collaboration: they should know how their name, email address, or roles that are used in the AAI are protected. And for some sensitive or high-profile research, also names and contact info needs to be protected!

There may also be legal and regulatory reasons to apply controls through your AAI. They can be in the research data itself, like medical and patient data, dual-use goods and knowledge, commercially confidential data, or ethical reasons on human research or in the Nagoya Protocol.

The challenge in risk management in AAI is to balance both aspects: enabling access and facilitating collaboration, making sure data and resources are *av ailable*, as well as protecting the *confidentiality and integrity* of data, resources, and users. Finding out where that balance is means you have to know your primary assets: "why are we, the collaboration, infrastructure, or service, here"?

What should I think of when identifying primary assets?

They are not what you may think they are! Primary assets are not computing thing, nor even AAI things. It's your research data, research processes, and the people and their knowledge.

Once you know your primary assets, you can proceed with your risk assessment, identify any secondary (or supporting) assets like ICT services, storage systems, and the AAI platform, and conceive controls to address exposure and limit the impact of any risks you identify.

Resources

- MITRE Crown Jewels Analysis
- ITSRM version 2 risk management
- ISO 27005 IT Risk Management Framework

Incident Response Procedure

A matter of *when*, not *if*. At some point you will face an cyber-security incident, or a security event that looks suspiciously like it. To contain, mitigate, and resolve the incident requires collaboration from everyone: resource providers, collaborations, users, and identity sources. The federated model of the AAI, in the AARC BPA as well as in eduGAIN more generally, means that many incidents touch all of these parties, and only by working together we can squash the incident. To do that effectively requires you have an *incident response procedure* that addresses your own security as well as the federation as a whole.

First response

- **Don't panic!** You are not the first one to be involved in a security incident, and by the time you observe the incident it has likely been ongoing for quite some time. On average: 3 months. So before you act: think!
    - look up your incident response procedure and have it handy so you do not forget any steps.
    - is this incident (potentially) going to be reported to law enforcement agencies (LEAs)? Then preserving *immutable* evidence is key: only work on copies and snapshots, and involve LEAs as early as possible. Otherwise, whatever evidence you collect may not hold up in court.
    - if you have an organisational security team (CSIRT, CERT, whatever), involve them! Not sure if you have one? Look up in Trusted Introducer or ask your Federation Operator.
- Think about physical safety: could the incident involve (remote) access to Operational Technology? SCADA systems controlling research equipment, lasers, beamlines, or BSL air scubbing facilities, for instance? Maybe at other sites? Make sure Occupational Health and Safety is alerted.
- What may appear as a malicious user could actually be a compromised identity, where the user is a victim as well. Do not jump to conclusions, but do share identifiers and names with identity providers and peer infrastructures, so the legitimate user can be warned and protected. But of course, there *can* be insider attacks and miscreants as well. Your investigations will find them...
- **Share information with peers in a trusted way**. Use the Traffic Light Protocol to ensure the audience knows what is expected without having to ask you every time. Could this incident potentially involve federated users? Involve your Federation Operator and (for services and infrastructures) the Collaboration platform and operator. They need to know to help you resolve the incident!
- If you find yourself overloaded with work, **find an incident coordinator in your infrastructure or at your federation**. Having a single focus point is tremendously helpful for keeping track of things.
- **For research collaborations**: you have a pivotal role in the AARC Blueprint as a hub for role and group management, and as the only entity that can link identifiers between services and identity providers! Cooperate, and be ready to suspend in case of doubt. The resource providers rely on your alertness and cooperation!

Response procedures

Every collaboration, infrastructure, and organisation is - to some extent - unique: who should be contacted, is an organisational security team available, what is the responsible chain of management. But the *structure* of incident response is mostly the same, regardless of the organisation. You should review the Security Incident Response Trust framework for Federated Identity (SIRTFI v2) and check whether you meet each of the basic steps. If you are in doubt about a requirement, that is a good place to start fixing it. Most common observation: lack of communication within your own organisation. This you can fix!

We provide a couple of examples from tried and tested infrastructures as well:

- UK IRIS
- eduGAIN Security Handbook (and other eduGAIN security resources)

Documenting your security contact information

"A clear statement of the policies and procedures of a CSIRT helps the constituent understand how best to report incidents and what support to expect afterwards. Will the CSIRT assist in resolving the incident? Will it provide help in avoiding incidents in the future?" "Constituent communities need to know exactly how their CSIRT will be working with other CSIRTs and organizations outside their constituency, and what information will be shared."

If you (an infrastructure, a collaboration, an identity source) have a means to publish information, do so:

1. publish a 'security.txt' file in its well-known location: `https://example.com/.well-known/security.txt`
2. write the brief description of the security team in RFC2350 format.

Well-known collaboration platform operators and security contacts

- B2ACCESS - security@eudat.eu
- CERN IAM services - computer.security@cern.ch
- CILogon - security@ncsa.illinois.edu
- eduGAIN - abuse@edugain.org
- EGI Checkin - abuse@egi.eu
- GEANT CoreAAI (MyAccessID, MyAcademicID, ...) - security@eduteams.org
- SURF SRAM - securityincident@surf.nl

Resources

- AARC I051 Guideline to Federated Security Incident Response for Research Collaboration
- SIRTFI version 2: federated security incident response and communications
- security contacts for federated identity providers and relying parties: infrastructures and services (REFEDS MET)
- UK IRIS
- eduGAIN Security Handbook

- eduGAIN security team
- Federation Operators list (eduGAIN)
- EGI CSIRT
- RFC 2350 Expectations for Computer Security Incident Response

Incident Response Procedure Template

This procedure should be shared with any service operator within your Research Collaboration.

Template for an Incident Response Procedure

Incident Response Procedure

The security contact for these procedures (Research Collaboration Security Officer) is: *{Security Officer Name + Contact Details (that must still reach Operators in the SO's absence)}*

Unless blank or otherwise stated, the following actions must be taken within the specified deadline as measured from the (suspected) incident first being reported.

Step	Action	Deadline
1. (Suspected) Discovery	Inform the local security team (if applicable)	WITHIN 4 HOURS
	Inform the Research Collaboration Security Officer	WITHIN 4 HOURS
	In collaboration with the Research Collaboration Security Officer, ensure that suspected affected external federated partners are informed	WITHIN 1 DAY
2. Containment	Isolate affected hosts (if feasible)	WITHIN 1 DAY
	Snapshot and/or suspend affected VMs	WITHIN 4 HOURS
	Disable affected appliances	WITHIN 4 HOURS
	Disable (suspected) detrimental user access	WITHIN 4 HOURS
3. Confirmation	Investigate to confirm whether the incident is real	WITHIN 1 DAY
4. Downtime Announcement	Announce applicable downtime. The reason e.g. "security operations in progress" should be stated.	WITHIN 1 DAY
5. Analysis	Collect evidence regarding the incident - as appropriate	
	Perform analysis of the incident - as appropriate	
	Follow-up on requests from security teams	WITHIN 4 HOURS of receipt
6. Debriefing	Prepare a post-mortem incident report in collaboration with the Research Collaboration Security Officer	WITHIN 1 MONTH
7. Normal Service Restoration	Restore normal service operation after incident handling is complete	

Post-mortem incident reports and the associated evidence and analysis must be retained for a period of at least 180 days after normal service operation.

Based on working hours: 1 day is 1 working day, 1 hour is 1 working hour

Legal and Regulatory Compliance

Most of your compliance will be determined by your local organisation and national rules and legislation. Items to keep in mind:

- export controls and dual-use restrictions
- processing of sensitive personal data
- knowledge safety
- medical and ethical requirements and mandatory reviews on human and animal research
- restrictions on GMOs and stem-cell research
- sanctions against individuals, organisations, and countries that are in effect in your jurisdiction

Please consult your local regulatory experts and legal council when in doubt.

How does compliance fit into an AARC BPA infrastructure?

Many of these controls can be at least partially implemented by access controls in the AARC platform, based on on attributes, roles, and group membership, and on the attributes released by the home organisations (authentication sources). Depending on the scope of the compliance rules, different layers may be the most appropriate place to enforce controls. E.g. a site-level or national infrastructure proxy is appropriate for national-level controls, whereas dual-use research might be more appropriately restricted at the community layer. Data on individuals, such as nationality, may only come from external government identity sources, potentially conveyed via another authentication source.

Membership Management

Getting your collaboration trusted as the ‘authoritative source of truth’ by authentication sources and service and infrastructure providers requires that your collaboration functions as intended, both today and in the future.

While collaboration management platform ‘downstream’, towards infrastructures and service providers, appears as an identity provider, it is – at least partially – making the actual identity opaque. The trust in the collaboration is based on its membership management and the adherence of its members to the purpose of the collaboration.

Similarly, towards the ‘identity’ layer - the part of the trust framework for authentication sources, possibly sourced from identity integration components or aggregators – the collaboration management should clarify that ‘access personal data’ is used in accordance with the identity provider requirements, in particular regarding minimisation of this personally identifiable information coming from the identity provider and its retention period.

In its basic form, collaboration management addresses *who is responsible for the collaboration* – the collaboration manager(s), and *what is the membership life cycle* – registration, assignment of roles, and group memberships, renewal, suspension, termination.

Large collaborations, and those that operate most of the registration process with specific, bespoke, processes, will need a more comprehensive ‘infrastructure-style’ membership management policy. It could include descriptions of a different enrolment flows, delegation of registration to a network of (home) organisations, or include review processes or a permit system for role assignment.

If you have your collaboration hosted on a platform

When you host your collaboration on a shared platform that offers its services to many communities, the platform usually defines a baseline for some operational aspects of membership management processes and handling ‘access personal data’. It can also help in makeing standard workflow available for collaboration managers, further easing this task.

And a collaboration platform provider will need to ensure the operational security of its platform and the publication of notices like the acceptable use policy and privacy notices. Since these elements are part of collaboration management, the collaboration should verify this capability, for example by reviewing the ‘Snctfi’ aspects of the policy development kit.

Membership Management Policy Development

The informational guideline AARC-I086 on membership management policy development, part of the AARC Policy Development Kit, provides membership management policy templates for use in both light-weight as well as composite (infrastructure) collaborations. The policy templates can be used by and should be adapted by collaborations before adoption. The light-weight template provides placeholder elements, such as the name of the collaboration, which can be filled in for seamless adoption. Adopting structured collaboration management facilitates trust by identity sources (ability to obtain more relevant authentication and identity attributes) and trust by infrastructures and service providers.

Using the PDK Membership Management templates

This informational guideline provides two ‘variants’ of a membership management policy template: one for light-weight collaborations, and a more extensive one for more ‘vertically integrated’ and composite collaborations. These are templates, in that each collaboration should review the proposed processes for suitability, and on adopting the policy fill in the placeholder elements, such as the name of the collaboration.

PDK v1 (infrastructures)	https://docs.google.com/document/d/1rVfpEGv_QIvf9V2gwtRS24kQfa3SYdbefAvaLnNpVa8
PDK v2 (collaborations)	https://docs.google.com/document/d/1mT6DhJxMg2APQlicE0UfjYjAfoZc1DqUJKQCZt_P4Q0

Resources

- Template for a membership policy
- AARC-I086 Membership Management Policy Development
- Light-weight Collaboration Management - explainer
- Template Membership Management Policy for Infrastructures and self-structured collaborations
- Template Membership Management Policy for light-weight and hosted collaborations

Template for a Community Membership Management Policy

Template for a Membership Management Policy

Membership Management Policy for <Collaboration name>

This policy is effective from <insert date>.

The current collaboration manager can be found at <insert link>.

INTRODUCTION

This policy establishes practices that are adopted by <collaboration X> in the management of its members. Accurate management of a collaboration's members and their authorisation attributes is fundamental to ensuring secure access control. Trust between <collaboration X>, underlying infrastructure and partner collaborations may be established by rigorous application of this policy.

COLLABORATION MANAGER

<Collaboration X> defines a Collaboration Manager role and assigns this role to two or more individuals. The Collaboration Manager is responsible for meeting the requirements identified in this policy. This responsibility may be devolved to designated personnel in the Collaboration or in the Infrastructure, and their trusted agents (such as Institute Representatives or Resource Centre Managers).

MEMBERSHIP LIFE CYCLE REQUIREMENTS

Membership Life Cycle: Registration

Membership Registration is the process by which an applicant joins the Collaboration and becomes a Member. Registration Data must be collected at the time of Registration, verified and stored in compliance with the Privacy Notice [ref]. Reasonable efforts must be spent to validate the data.

Membership Life Cycle: Assignment of Attributes

Assignment of attributes (such as group membership, entitlements, or roles) shall be the responsibility of the Collaboration Manager or of designated person(s).
Attributes shall be assigned only for as long as they are applicable.

Membership Life Cycle: Renewal

Membership Renewal is the process by which a User remains a member. Renewal procedures shall
* ensure that accurate Registration Data is maintained
* confirm continued eligibility of the User to use Resources assigned to the Collaboration
* confirm continued eligibility of the User to any attributes
* ensure the reaffirmation of acceptance of the Collaboration AUP

The maximum time span between Registration and Renewal, and between Renewals, shall be <INSERT RENEWAL TIMESPAN>. The User shall be able to correct and amend their Registration Data at any time.

Membership Life Cycle: Suspension

The Suspension of Collaboration membership is the temporary revocation of full or partial rights and of any attributes. Suspension is done by or on behalf of the Collaboration Manager.
A User should be suspended when the Collaboration Manager is presented with reasonable evidence that the member's identity or credentials have been used, with or without the user's consent, in breach of relevant Policy.
The Collaboration Manager must act on any requests for suspension without delay.
User's rights shall not be reinstated unless the Collaboration Manager has sent timely prior notification to all those who requested Suspension.

Membership Life Cycle: Termination

The Termination of Collaboration membership is the removal of a member from the Collaboration. Following Termination, the former member is no longer eligible to use Infrastructure Resources assigned to the Collaboration. The Collaboration must no longer assert membership or attributes for the former member.
In absence of overriding reasons, a request by the User for removal must be honoured.
The events that shall trigger possible termination of the User's membership of the Collaboration include:
* failure to complete a membership Renewal process within the allotted time
* end of participation of the User in the Collaboration

REGISTRATION DATA REQUIREMENTS

The Registration data for a User comprises verified information:
* family name(s)
* given name(s)
* the employing organisation name and address
* a professional email address
* unique and non-reassigned identifier(s) of the User and the source of authority of each identifier
* <Add or delete lines as required>
and is recommended to contain:
* professional contact telephone number so as to inform the User promptly during the investigation of security incidents and of lifecycle events
* other contact information, as voluntarily provided and maintained by the User.
The types of information recorded must be listed in the Privacy Notice

AUDIT AND TRACEABILITY REQUIREMENTS

The Collaboration records and maintains an audit log of all membership lifecycle transactions. This audit log is kept for a minimum period consistent with the Traceability and Logging Policies of all Infrastructures that provide resources to the Collaboration.
* Membership,
* assignment of or change to a member's attributes,
* membership renewal,
* membership suspension,
* membership termination or re-evaluation.
Each logged event should record the date and time, the originator, the details of the event, and whether or not it was approved. The identity of the person granting or refusing the request should be recorded, including any verification steps involved and other people consulted.

ACCEPTABLE USE POLICY REQUIREMENTS

Collaboration X defines an Acceptable Use Policy (AUP) [ref]. The AUP must be shown to all persons joining the Collaboration. Acceptance of the AUP by Collaboration members who act as responsible persons towards the Infrastructure must be an explicit action, must be recorded, and must be a prerequisite for registration in the Collaboration [ref]. The AUP should provide awareness on inappropriate actions by individual users that may affect the ability of other members to use an infrastructure.

Attribution

Notice Management (presentation)

The objective of notice management is to reduce the number of interactions (‘clicks’ and other ‘interruptions’) users face in order to achieve their actual desired objectives, both on first use and on subsequent use of the same services in a similar workflow. As proposed in AARC-G083, Guidance for Notice Management by Proxies, notice presenters (proxies, communities, and service and data providers) should provide machine-readable identifiers and associated registry meta-data for the notices they can present to users.

This feature needs technical support in the collaboration platform, and in most cases the platform operator will provide this as part of the AAI platform service. Only for collaborations implementing their own AAI service, the feature should be added explicitly.

How does notice presentation fit into an AARC BPA compliant infrastructure?

The collaboration platform is the consolidated entry point for all access to resources, usually across multiple infrastructures and services. Hence the collaboration layer is the most appropriate place to present this, and forestall interstitial notice screens for each infrastructure and service.

When a common notice can be constructed for all services and infrastructures that are supporting a collaboration, the collaboration can present that common notice once also without a technical notice presentation component.

Resources

- AARC-G083 Guidance for Notice Management by Proxies
- User interaction in a proof-of-concept implementation of a notice presentation component

File	Modified
PNG File 01_Registry-getPolicies.png	Jan 14, 2026 by Maarten Kremers
PNG File 02_Registry-getPolicy.png	Jan 14, 2026 by Maarten Kremers
PNG File 03_NPC-check.png	Jan 14, 2026 by Maarten Kremers
PNG File 04_NPC-interrupt.png	Jan 14, 2026 by Maarten Kremers
PNG File 05_NPC-interrupt-agree.png	Jan 14, 2026 by Maarten Kremers
PNG File 06_NPC-check-result.png	Jan 14, 2026 by Maarten Kremers

Download All

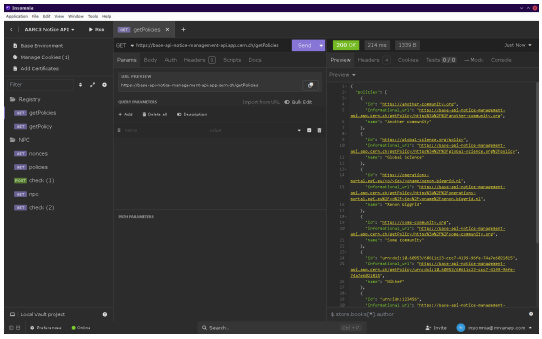

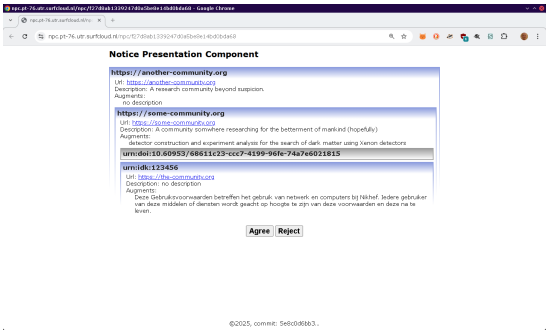

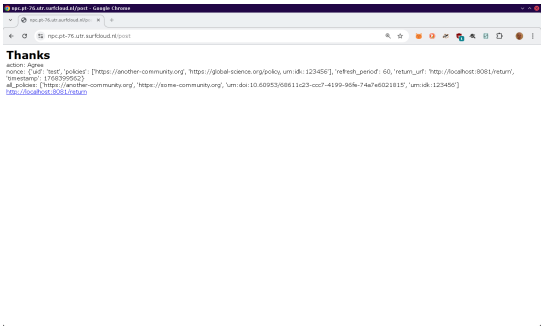

Notice Management (Provision)

The objective of notice management is to reduce the number of interactions (‘clicks’ and other ‘interruptions’) users face in order to achieve their actual desired objectives, both on first use and on subsequent use of the same services in a similar workflow. As proposed in AARC-G083, Guidance for Notice Management by Proxies:

> *All notice presenters (proxies, communities, and service and data providers) SHOULD provide machine-readable identifiers and associated registry meta-data for the notices they can present to users.*

Ideally Research Infrastructures should converge on the same policies, but that is often not feasible due to internal requirements (e.g. geographical, risk-tolerance etc.) In this case, a meta-data structure is proposed to clarify which additional policies are **augmented** by a policy, or **included** within it. Research Infrastructures are increasingly stacked and connected, meaning that researchers are interrupted to accept notifications of many policies to complete their work. Some of these policies are the same, and by publishing structured data about the policies it is possible to compute whether a new notification is required for a user to access a service.

A Notice Presentation element then picks up the metadata and processes it to minimise contact with the user.

What are these "Notifications"?

Concretely, the notifications that impact usability are acceptance of Acceptable Use Policies and Privacy Notices.

The WISE AUP is recommended by AARC as a common Acceptable Use Policy. No particular privacy notice is recommended by AARC due to diverging regional requirements.

What does notification meta-data look like?

Full details are available at AARC-G083. The following is an example of a fictional AUP at URL "https://www.nikhef.nl/aup/" required to access Nikhef. It includes a secondary policy "https://documents.egi.eu/document/2623", meaning that acceptance of this policy constitutes acceptance of the secondary policy.

```
{
    "id": "urn:doi:10.60953/68611c23-ccc7-4199-96fe-74a7e6021815",
    "aut": "https://www.nikhef.nl/",
    "aut_name": "Nikhef",
    "valid_from": 1649023200,
    "ttl": 604800,
    "contacts": [
        "helldesk@nikhef.nl",
        "information-security@nikhef.nl"
    ],
    "security_contacts": [
        "abuse@nikhef.nl"
    ],
    "privacy_contacts": [
        "privacy@nikhef.nl"
    ],
    "policy_class": "acceptable-use",
    "notice_refresh_period": 34214400,
    "includes_policy_uris": [
        "https://documents.egi.eu/document/2623"
    ],
    "policy_uri": "https://www.nikhef.nl/aup/",
    "description#nl_NL": "Deze Gebruiksvoorwaarden betreffen het gebruik van
    netwerk en computers bij Nikhef.Iedere gebruiker van deze middelen of
    diensten wordt geacht op hoogte te zijn van deze voorwaarden en deze na te
    leven.
    ",
    "description": "This Acceptable Use Policy governs the use of the Nikhef
    networking and computer services;all users of these services are expected
    to understand and comply to these rules.
    "
}
```

The following is an example of a fictional "purpose binding" policy at URL "https://operations-portal.egi.eu/vo/view/voname/xenon.biggrid.nl" that augments the WISE AUP with its description field "detector construction and experiment analysis for the search of dark matter using Xenon detectors". This means that a user would be shown the WISE AUP with the placeholder for purpose binding filled in with this description.

```
{
    "id": "https://operations-portal.egi.eu/vo/view/voname/xenon.biggrid.nl",
    "aut": "https://xenonexperiment.org/",
    "aut_name": "Xenon-nT collaboration",
    "valid_from": 1311890400,
    "ttl": 31557600,
    "contacts": [
        "grid.support@nikhef.nl",
    ],
    "security_contacts": [
        "vo-xenon-admins@biggrid.nl"
    ],
    "policy_class": "purpose",
    "augments_policy_uris": [
        "https://wise-community.org/wise-baseline-aup/v1/"
    ],
    "policy_uri": "https://operations-portal.egi.eu/vo/view/voname/xenon.biggrid.nl",
    "description": "detector construction and experiment analysis for the
    search of dark matter using Xenon detectors "
}
```


Resources

- AARC-G083 Guidance for Notice Management by Proxies https://aarc-community.org/guidelines/aarc-g083/

REFEDS Assurance Framework

What is RAF

To manage risk in federated access, Relying Parties (RPs) sometimes need more confidence in the identity- and attribute-related assertions made by an Identity Provider (IdP) and its underlying Credential Service Provider (CSP). The **R**EFEDS **A**ssurance **F**ramework (RAF) defines a pragmatic way to express this confidence in commonly-used federation protocols, so that RPs (or proxies) can make more informed access-control decisions and CSPs /IdPs can communicate what they actually do.

RAF focuses on identity and attribute assurance (e.g., uniqueness, identity proofing, attribute quality/freshness).
It *does not* define authentication strength; for that, use the REFEDS authentication profiles (SFA/MFA) alongside RAF.

RAF 2.0 components and profiles

RAF 2.0 splits assurance into independent components:

- Identifier Uniqueness
- Identity Assurance
- Attribute Assurance

To simplify consumption by RPs, RAF 2.0 also defines two **assurance profiles** (bundles of component requirements):

- RAF Cappuccino (moderate use cases)
- RAF Espresso (higher-risk use cases)

In RAF cappuccino: a unique identifier, medium-level identity assurance, and the 'affiliation' attribute reflects the status no longer than one month in arrears. This level can be combined with single-factor authenticator strength (http://refeds.org/sfa)
In RAF Espresso: there is a unique identifier, high-level identity assurance, and the 'affiliation' attribute reflects the status no longer than one month in arrears. This level is best combined with multi-factor authenticators (http://refeds.org/mfa).

Who should adopt RAF and why?

Identity Providers / Credential Service Providers (CSPs)

- To clearly communicate how user identities are established and managed (enrolment, proofing, lifecycle).
- To enable downstream RPs to apply risk-based access control using consistent, community-defined signals.

Federations and federation operators

- To provide a common “assurance vocabulary” that scales across many IdPs and RPs.
- To support consistent interpretation and smoother inter-federation interoperability.

Service Providers / Relying Parties (RPs)

- To request and interpret assurance information in a structured way, aligned to real risk and policy requirements.
- To avoid bilateral “questionnaires” and bespoke assurance mappings wherever possible.

Related work in AARC / trust-policy building blocks

If you are building an end-to-end trust posture for a collaboration/infrastructure, RAF is part on resolving the the Assurance Requirements

Resources

- REFEDS Assurance Framework (overview landing page): https://refeds.org/assurance
- RAF 2.0 specification (PDF): https://refeds.org/wp-content/uploads/2023/12/RAF-2.0-Final-version.pdf
- FAQ / Supporting Materials: RAF (identity assurance): https://wiki.refeds.org/pages/viewpage.action?pageId=31982150

Research Risk Assessment

"Bad things can happen to good science", as the Open Science Cyber Risk Profile acutely states. While you may not think of it at first, the data, ways of working, and collections created in your collaboration are valuable and deserve protection. External cybersecurity attacks of course come to mind, but in many cases inadvertent accidents happen and are at least as big a risk.

Working with sensitive and personal data

When the research data contains personal data, you may be required by regulation or law to perform specific risk assessments, like a Data Protection Impact Assessment (DPIA). The same holds true in case you work with human data and your research is subject to medical ethical guidelines.

Of course also the AAI itself will use and generate personal data as part of providing access to services. This is not the kind of data that usually leads to specific risks, as long as you follow the REFEDS Data Protection Code of Conduct. And of course if you engage platform providers for your AAI make sure these follow the REFEDS good practices as well.

Beware also of data-level access control, and how to work with replicated data. This usually needs a data access management system, such as the Resource Entitlement Management System (REMS), a tool for managing access rights to data and datasets, or the Data Passports in GA4GH.

When risk assessment is absolutely critical ...

Does your collaboration work with human, societal data, or collects questionnaires? Is your research likely to be classed as dual-use or export restructured? Does the research, or your collaboration users, touch on knowledge safety? Is approval by medical/ethical commissions needed? Are you dealing with biodiversity or genetic resources subject to the Nagoya Protocol? Do a specific risk assessment or ask your institution for guidance.

Resources

- Open Science Cyber Risk Profile by TrustedCI, Sean Peisart *et al.*)
- IT Security Risk Management v2 (EC framework)
- ISO 27005
- AARC-G042 Data Protection Impact Assessment – an initial guide for communities
- Global Alliance for Genomics and Health
- Resource Entitlement Management System (REMS)
- B2SAFE by EUDAT CDI

Rules of Participation

While management of the user directory, or responding to compromised credentials and questions from your service providers may appear a technical task, these often involve decisions that deal with your 'primary assets': the research data you work with, the processes that keep your community together and wholesome. A handful of people working together can solve issues that arise in an ad-hoc fashion, but as soon as you grow a bit further, structured communication becomes essential. "Why was I suspended?" "Why can't I join your group ... you have a service I need (and by the way, I just want to use it, not contribute)?". You need a body to take up that authority. Unfortunately, the AAI is often the first time you hit these hard questions!

Who is in and who is not? The role of the governance body

A principal investigator, research group chair, or faculty dean makes a good starting point for a local governance body. If your community becomes larger, write down rules of participation, draft a memorandum of understanding, or have a written collaboration agreement in place. Often there is already something there: many public research projects require having a collaboration or grant agreement. There is usually a governance structure in there, which can be re-used here. No need to re-invent the wheel within your community.

Resources

Governance comes in many shapes and forms, and how to organise participation in your federation and collaboration is really up to you. But look at your project agreements, department model, or grant office to find a suitable and effective solution, and some existing collaborations, large and small, may serve as inspiration!

- EC Horizon Model Grant Agreement (and annotated version)
- EOSC Rules of Participation
- CERN Convention
- LHCb memorandum of understanding 2002

Security Operational Baseline

Your resources, your research data, and your collaborators are valuable, yet continuously targeted by miscreants. Protecting these requires that everyone works together to protect them damage, disruption, and unauthorised use. The *security operational baseline* sets the bar to entry in most federated infrastructures, not only for the AAI but also for the protection of resources, data, and users. By adhering to a simple common baseline, the collaboration becomes easier and the expectations of all parties are clear. You will find that adherence to the security operational baseline is a prerequisite to entry for many of the ecosystems you want to join.

Adopting it as an e-Infrastructure or service provider

Users and research collaboration will be using your service(s) under the assumption that it is safe to do so, and - if you rely on others - that they can rely on you to manage your dependencies properly. Especially in 'cloud' scenarios, your supply chain in terms of both infrastructure and software is critical, and modern cybersecurity directives like Europe's NIS2 directive emphasise the importance of the supply chain.

The security baseline gives you the outline of the security measures that help you participate in federation and provide trustworthy services. It relies on SIR TFI, the Security Incident Response Trust Framework for Federated Identity, and helps identify, mitigate, and resolve security incidents in your service and in your peers. Remember: you will typically find a security incident quite a long time after the intrusion actually happened, so keeping logs is particularly important! And exercising these communications channels is needed to make sure they work in case of emergencies.

Adopting it as an authentication source or collaboration

The AARC BPA identifies the collaboration layer as a key control point for resource access: as a collaboration, you hold critical data on 'who did what when', and are the most effective place to control access to resources and protect the infrastructures you use. You may be notified of security events by infrastructures and resource providers, since they will often identify anomalous behaviour first and have a definite interesting in stopping the abuse. The collaboration and the home organisation are the only ones in the AARC BPA that can associate the identifiers at the infrastructure and service layer with actual users. You are an essential part in stooping the incident from spreading and causing more harm!

How does the Security Operational Baseline fit into an AARC BPA compliant infrastructure?

Operational security is of course much more than AAI, but identity plays a critical role: as users 'traverse' the layered AARC BPA, they will acquire additional identifiers, and are gaining access to resources by virtue of their role, group memberships, and capabilities. Each of the layers in the AARC BPA provides control points to mitigate incidents and prevent the impact from spreading.

Also keep in mind that each layer in the AARC BPA may keep *state* about the user sessions: even if an account is blocked at, for example, the identity layer, there may be sessions, tokens, or assertions active at the collaboration, infrastructure, or site-local layers. Therefore collaboration is essential, and reporting of security events and (potential) compromises to your partners is very important.

The twelve points to protect your resources and data

It is RECOMMENDED that all service providers follow these Baseline Requirements to achieve a sufficient level of security. These requirements *augment but do not replace* applicable security policies and obligations, nor any more specific security arrangements and service level agreements that may exist between participants. The baseline is specifically that: a set of minimal expectations between everyone in the infrastructure:

1. comply with the SIRTFI security incident response framework for structured and coordinated incident response
2. ensure that your Users agree to an Acceptable Use Policy (AUP) or Terms of Use, and that there is a means to contact each User.
3. promptly inform Users and other affected parties if action is taken to protect their Service, or the Infrastructure, by controlling access to their Service, and do so only
   for administrative, operational or security purposes.
4. honour the confidentiality requirements of information gained as a result of your Service's participation in the Infrastructure.
5. respect the legal and contractual rights of Users and others with regard to the personal data processed, and only use access personal data for administrative, operational, accounting, monitoring or security purposes.
6. retain system generated information (logs) in order to allow the reconstruction of a coherent and complete view of activity as part of a security incident (the 'who, what, where, when', and 'to whom'), for a minimum period of 180 days, to be used during the investigation of a security incident.
7. follow, as a minimum, generally accepted IT security best practices and governance, such as pro-actively applying secure configurations and security updates, and taking appropriate action in relation to security vulnerability notifications, and agree to participate in drills or simulation exercises to test Infrastructure resilience as a whole.
8. operate services and infrastructure in a manner which is not detrimental to the security of the Infrastructure nor to any of its Participants or Users.
9. collaborate in a timely fashion with others, specifically those with which there is a direct trust relationship, in the reporting and resolution of security events or incidents related to their participation in the infrastructure and those affecting the infrastructure as a whole.
10. honour the obligations on security collaboration and log retention (clauses 1, 6, and 9 above) for the period of 180 days after their Service is retired from the Infrastructure, including the retention of logs when physical or virtual environments are decommissioned.
11. not hold Users or other Infrastructure participants liable for any loss or damage incurred as a result of the delivery or use of the Service in the Infrastructure, except to the extent specified by law or any licence or service level agreement.
12. maintain an suppliers that ensures that engagement of such parties does not result in violation of this Security Baseline.

Resources

- AARC-G084 Security Operational Baseline
- SIRTFI, the Security Incident Response Trust framework for Federated Identity
- Network and Information Security directive 2.0 (NIS2)
- Incident Response Procedure

Security Operational Baseline Template

Introduction

To fulfil its mission and protect primary and secondary assets of any infrastructure and community, it is necessary to be protected from damage, disruption, and unauthorised use. This reference ‘Security Operational Baseline’ supports these goals by defining 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. It aims to establish a sufficient level of trust between all participants in an infrastructure to enable reliable and secure operation.

The Security Operational Baseline codifies current community good practice for protecting authentication providers, AAI platforms, and identity providers, participating in an AAI Federation. It is RECOMMENDED that all service providers follow these Baseline Requirements to achieve a sufficient level of security. These requirements augment but do not replace applicable security policies and obligations, nor any more specific security arrangements and service level agreements that may exist between participants.

Terminology

Terminology in this document follows conventional IT service management vocabulary, such as ITIL and FitSM, and the RFC 2119 key words. For clarification, we define the following specific terms.

Term	Definition
Service Provider	an organisation (or part of an organisation) that manages and delivers a service or services to customers
Identity Provider	a service that creates, maintains, and manages identity information for principals and provides authentication services to relying parties
AAI Platform	an authentication/authorization infrastructure (AAI) service or service component, identity, community, infrastructure, or local ‘proxy’ that augments, translates, or transposes authentication and authorization information, including the connected sources of access (AAI) attributes, as detailed in the AARC BPA 2025 (AARC-G080).
User	an individual that primarily benefits from and uses a service

This Guideline is accompanied by implementation recommendations and reference material. Links to these materials are provided at https://aarc-community.org/guidelines/aarc-g084/.

Security Baseline

To adhere to the Security Operational Baseline, you must:

1. comply with the SIRTFI security incident response framework for structured and coordinated incident response

2. ensure that your Users agree to an Acceptable Use Policy (AUP) or Terms of Use, and that there is a means to contact each User.

3. promptly inform Users and other affected parties if action is taken to protect their Service, or the Infrastructure, by controlling access to their Service, and do so only for administrative, operational or security purposes.

4. honour the confidentiality requirements of information gained as a result of your Service’s participation in the Infrastructure.

5. respect the legal and contractual rights of Users and others with regard to the personal data processed, and only use access personal data for administrative, operational, accounting, monitoring or security purposes.

6. retain system generated information (logs) in order to allow the reconstruction of a coherent and complete view of activity as part of a security incident (the ‘who, what, where, when’, and ‘to whom’), for a minimum period of 180 days, to be used during the investigation of a security incident.

7. follow, as a minimum, generally accepted IT security best practices and governance, such as pro-actively applying secure configurations and security updates, and taking appropriate action in relation to security vulnerability notifications, and agree to participate in drills or simulation exercises to test Infrastructure resilience as a whole.

8. operate services and infrastructure in a manner which is not detrimental to the security of the Infrastructure nor to any of its Participants or Users.

9. collaborate in a timely fashion with others, specifically those with which there is a direct trust relationship, in the reporting and resolution of security events or incidents related to their participation in the infrastructure and those affecting the infrastructure as a whole.

10. honour the obligations on security collaboration and log retention (clauses 1, 6, and 9 above) for the period of 180 days after their Service is retired from the Infrastructure, including the retention of logs when physical or virtual environments are decommissioned.

11. not hold Users or other Infrastructure participants liable for any loss or damage incurred as a result of the delivery or use of the Service in the Infrastructure, except to the extent specified by law or any licence or service level agreement.

12. maintain an agreement with representatives for individual service components and suppliers that ensures that engagement of such parties does not result in violation of this Security Baseline.

Sensitive Data Access Policies

Beware also of data-level access control, and how to work with replicated data. This usually needs a data access management system, such as the Resource Entitlement Management System (REMS), a tool for managing access rights to data and datasets, or the Data Passports in GA4GH.

When data access control is absolutely critical ...

Does your collaboration work with human, societal data, or collects questionnaires? Is your research likely to be classed as dual-use or export restructured? Does the research, or your collaboration users, touch on knowledge safety? Is approval by medical/ethical commissions needed? Are you dealing with biodiversity or genetic resources subject to the Nagoya Protocol?

Similarly, data providers, especially those providing access to sensitive data, will have specific policies in place to know exactly to whom they are providing access - patient records and the binding of permits for data access from the HDABs come to mind.

Ask your data provider or your institution for guidance.

Resources

- Global Alliance for Genomics and Health
- Resource Entitlement Management System (REMS)
- B2SAFE by EUDAT CDI

Service Levels and Data Classification (the "IAC" or "CIA" triad)

As a service provider or infrastructure, you have an idea what services you and your infrastructure partners offer, and to whom. And, as you are processing information, your service will contain valuable data. But what kind of data you are able to store in your infrastructure depends on the risk both you and your users are willing to absorb. And there may be regulatory requirements in place that prohibit you from storing certain classes or data, or that require you to secure of certify your services in a specific way.

Due to the diverse nature of data and services, the Policy Development Kit, while providing context, does not have specific recommendations, policy templates, or procedures. Public resources related to service management, risk assessment, and “AIC” (or “CIA”) information classification are provided below.

Service Levels

Service Management Systems, like the ISO 20 000 standards or FitSM, expect services to be “aligned to the needs and expectations of (potential) customers. Both the service provider and customer are aware of agreed service targets” (quoted from the FitSM standards). When working with a research collaboration, you may be faced with ‘unknown’ data and many implicit expectations, and making the users aware of what your service can and cannot do is essential for both security and functionality.

When engaging with users communities, be it directly or as part of an infrastructure, you should check whether the *purpose* for which the collaboration comes together matches the capabilities of your service(s). At times, you may want to explicitly clarify what your services may be used for, what their availability and reliability is, and any regulatory limitations users should be aware of. You do this by adding Terms and Conditions to the WISE Baseline AUP, after the 10 immutable clauses.

If you want to provide service guarantees to some of your collaborations, you can also add references to service level agreements at the end of the AUP presentation (“Applicable service levels agreements are located at: <URL>”).

Data classification: availability, integrity, confidentiality

A property of the information, rather than the infrastructure, the triad of availability, integrity, and confidentiality (seen in various permutations as well, like the “CIA Triad”), these different aspects of data security are foundational for both security as well as regulatory compliance.

As a service provider, you should make clear what classes of data may be processed by or stored in your services, and – depending on the sensitivity of the data – you should implement policy and technical controls to prevent the ’wrong’ type of data being sent to your services.

How important each aspect of the AIC/CIA triad is, depends on the (research) use case. Can the user accommodate delays in access, or being locked out for a longer time, as long as there is a near-absolute guarantee that the data will never be altered? It is essentially open, public data, but does it need to be available with 99.9999% availability?

Regulatory compliance also factors in here; for example personal electronic health data is especially protected, but also always available for primary use in healthcare. It scored “high” on all three aspects. Yet secondary use of this data, while still highly confidential, might be more lenient towards availability.

Work with your user communities to extract their data classification, preferably based on a risk assessment. In this assessment, users should take into account their ‘primary assets’ (data and processes that serve the purpose of their collaboration), and service providers should validate whether they have the technical and organisations measures in place to limit the risks to an acceptable level.

Beware of implicit expectations of users, and document the agreements as part of your Service Level Agreement (SLA) or Operational Level Agreements (OLA). Remember that in some cases you may need to (self) certify for compliance to legislation and implementing acts. This is especially common for personal and health data, but can also apply to dual-use and export-restricted knowledge.

Resources

- https://wise-community.org/wise-baseline-aup/
- https://www.itgovernance.co.uk/blog/what-is-the-cia-triad-and-why-is-it-important
- https://www.fitsm.eu/downloads/ (especially the Overview and Vocabulary)
- https://www.itgovernance.co.uk/blog/why-iso-27005-risk-management-is-key-to-iso-27001 (risk management and assets)
- https://www.trustedci.org/oscrp Open Science Cyber Risk Profile
- https://hora.surf.nl/index.php/BIV_classificaties (Higher Education and Research Reference Architecture, IAC classifications for objects (in Dutch))
- ...

SIRTFI

The **S**ecurity **I**ncident **R**esponse **T**rust Framework for **F**ederated **I**dentity (Sirtfi) aims to enable the coordination of incident response across federated organisations. This assurance framework comprises a list of assertions which an organisation can attest in order to be declared Sirtfi compliant. An official "trust mark" has been defined for SAML Service Providers and Identity Providers to allow organisations to technically advertise their adoption of the framework, and a similar approach is being planned for OpenID Federation.

Who should adopt Sirtfi and why?

Any federated organisation should adopt Sirtfi to improve their internal operational security and enable them to participate in federated security incident response. If you have not published a security contact in the identity federation then you will not be informed of any ongoing security incidents that may affect you.

How does Sirtfi fit into an AARC BPA compliant infrastructure?

In order to assert Sirtfi for your entire infrastructure it is important that you can ensure that the list of Sirtfi requirements is met for both the AAI layer and all connected services. Since Sirtfi requires that best practices in operational security are met, e.g. regular patching, you will need to adopt a policy for your infrastructure to ensure that these practices are followed by each service. A central operational security team for your infrastructure may play an important role in supporting these practices, such as handling communication during an incident and propagating information on vulnerabilities and breaches to downstream systems.

Sirtfi is referenced in several AARC policy guidelines such as

- Security Operational Baseline
- Incident Response Procedure
- Attribute Authority Operational Security

How to express your Sirtfi compliance?

Please visit https://wiki.refeds.org/display/SIRTFI/Guide+for+Federation+Participants

Resources

- REFEDS Sirtfi public page https://refeds.org/sirtfi
- REFEDS technical Sirtfi implementation guide https://wiki.refeds.org/display/SIRTFI/SIRTFI+Home

The REFEDS Data Protection Code of Conduct

Your collaboration may deal with personal data of two types, both of which deserve your protection:

1. personal data of the users *using* the infrastructure themselves, and (or)
2. personal data *contained in the research* objects processed within that infrastructure, or in the data exposed in the infrastructure.

All collaborations will have to manage "(1)", the ***access personal data***. This is the personal data, also known as personally identifiable information (PII), that comes from the authentication sources, is contained in the collaboration membership management service, and that is collected as part of access control to services, accounting, and security logging. In most cases it will contain (semi) public information like the name of the user, institutional email address, and the internet addresses hence the user originates. In a federated AAI, usually long-term credentials like passwords are only present in the authentication sources (home organisation, or the user's wallet), but more ephemeral credentials like ID tokens and structured JWTs may also carry information about the user. **This PDK article, and the REFEDS Data Protection Code of Conduct ('DPCoCo') deal with this type of *access personal data* (only)**.

When you are based in, collaborate with, or process data of live people in the European Union and the EEA, as well as in jurisdictions with similar statutes, you are required to project this data. You must follow 'Privacy by design, privacy by default', protect this data at rest and in transit, and have an identified reason for handling the personal data in the first place. The General Data Protection Regulation (GDPR) and any national implementation legislation describes what you should do, and what you cannot do. Even if you are not subject to GDPR, organisational policy may be in place (for quite some intergovernmental organisations), or your collaboration may want to collaborate with users in a 'GDPR country' ... and you will be affected by GDPR anyway.

The 'REFEDS Data Protection Code of Conduct' helps you align with GDPR requirements, and the example privacy notices help communicate *how* you meet these to both users (data subjects) and reviewers.

Some collaboration deal also with *personal data as part of the data they operate on*. Think of the data of real people in medical and health studies, social science surveys, and social geography. If your collaboration uses that type of data, and it is not fully anonymised (which is quite hard!), then you should take extra care when engaging service providers and infrastructures. Agreements and contracts should be in place before working with that data, and the safeguards should be firm enough to satisfy your organisation's risk appetite. This PDK article does not address the protection needed for research data - that is specific to the type of research data, and not an AAI issue. Of course, you can use the AAI to implement appropriate safeguards and technical and organisational measures to protect this sensitive data.

Resources

- Templates of data privacy notices
- https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home

WISE AUP

Acceptable use policy (AUP) and terms and conditions (T&C) are necessary instruments in the regulation of infrastructure access. And while for 'enterprise' sources they can be complex - as you have to define for example what is acceptable use for employees of organisational resources - for access to resources by collaborations and their users it can be far simpler. Basically, a collaboration access resources *for a particular specific purpose*: achieving the objectives of the common research goals. As such, the AUP can be simple: it binds the user to the 'purpose' for which the services and resources they use have been provided.

Yet, like with privacy notices, the reader is rather inclined to click through and proceed with the actual task at hand. Thus, to reduce the burden on the user and increase the likelihood that they will read the AUP, the number of times a user is presented with such notices must be kept to a minimum, preferably just a single time. Yet the notice should cover as much of the user's potential use of the infrastructure as possible: the more services and resources deem an AUP as sufficient for their policy purposes, the better it will be. This will allow users to use resources from multiple service and resource providers without the need to confirm acceptance of additional AUPs.

For what purpose, and under whose authority

> *This Acceptable Use Policy and Conditions of Use ("AUP") defines the rules and conditions that govern your access to and use (including transmission, processing, and storage of data) of the resources and services ("Services") as granted by* {community, agency, or infrastructure name} *for the purpose of* {describe the stated goals and policies governing the intended use}*.*

Remember that the PDK practical steps asked you to "identify a governance body to make policy decisions"? That's the name of the 'community, agency, or infrastructure name' that is authoritative for your AUP. it is here typically the name of your research collaboration, like in "*... as granted by the Harderwijk Herbal Research Collaboration HHRC for the purpose of ...*". The purpose of your collaboration was the next item in the priority list, and that text goes into the next placeholder. For example "... for the purpose of *identifying flowers and ever-green plants on your neighbourhood commons during the airing of BBC2 gardening TV programmes.*". So if somebody accesses the HHRC in order to locate where plant species could found, picked, and their generic composition taken in violation of the Nagoya protocol, that is clearly unacceptable use.

No more than 10 bullet points

... and of these points, none are controversial. So at a minimum include the 10 commandments from the WISE Baseline AUP. By doing so, *all* resource providers, infrastructures, and your peer collaborations or partner 'nodes' know they can accept your collaborators without further ado.

And if you as a collaboration, backed by the resource providers, can provide more guarantees to its users this is perfectly fine as well! Just add these terms and conditions, and state the advanced service levels, in the last section of the "AUP/T&C" in the spot indicated in the WISE Baseline AUP. But keep the Baseline AUP intact for interoperability. The basic set is merely these:

> *1. You shall only use the Services in a manner consistent with the purposes and limitations described above; you shall show consideration towards other users including by not causing harm to the Services; you have an obligation to collaborate in the resolution of issues arising from your use of the Services.*
> *2. You shall only use the Services for lawful purposes and not breach, attempt to breach, nor circumvent administrative or security controls.*
> *3. You shall respect intellectual property and confidentiality agreements*
> *4. You shall protect your access credentials (e.g. passwords, private keys or multi-factor tokens); no intentional sharing is permitted.*
> *5. You shall keep your registered information correct and up to date.*
> *6. You shall promptly report known or suspected security breaches, credential compromise, or misuse to the security contact stated below; and report any compromised credentials to the relevant issuing authorities.*
> *7. Reliance on the Services shall only be to the extent specified by any applicable service level agreements listed below. Use without such agreements is at your own risk.*
> *8. Your personal data will be processed in accordance with the privacy statements referenced below.*
> *9. Your use of the Services may be restricted or suspended, for administrative, operational, or security reasons, without prior notice and without compensation.*
> *10. If you violate these rules, you may be liable for the consequences, which may include your account being suspended and a report being made to your home organisation or to law enforcement.*

Now just add your contact address for information, security and privacy at the end and you are done. The actual privacy notices can be referenced and do not need to be included. It is fine to have these one more click away (so in total two clicks).

Resources

- WISE Baseline AUP
- Template for your Acceptable Use Policy based on the WISE AUP
- AARC-I044 Implementers Guide to the WISE Baseline Acceptable Use Policy
- AARC-G083 Notice Management by proxies

WISE AUP Template

Template for a WISE Baseline AUP

Acceptable Use Policy and Conditions of Use

This Acceptable Use Policy and Conditions of Use ("AUP") defines the rules and conditions that govern your access to and use (including transmission, processing, and storage of data) of the resources and services ("Services") as granted by {community, agency, or infrastructure name} for the purpose of {describe the stated goals and policies governing the intended use}.

<To further define and limit what constitutes acceptable use, the community, agency, or infrastructure may optionally add additional information, rules or conditions, or references thereto, here or at the placeholder below. These additions must not conflict with the clauses 1 to 10 below, whose wording and numbering must not be changed.>

1. You shall only use the Services in a manner consistent with the purposes and limitations described above; you shall show consideration towards other users including by not causing harm to the Services; you have an obligation to collaborate in the resolution of issues arising from your use of the Services.

2. You shall only use the Services for lawful purposes and not breach, attempt to breach, nor circumvent administrative or security controls.

3. You shall respect intellectual property and confidentiality agreements.

4. You shall protect your access credentials (e.g. passwords, private keys or multi-factor tokens); no intentional sharing is permitted.

5. You shall keep your registered information correct and up to date.

6. You shall promptly report known or suspected security breaches, credential compromise, or misuse to the security contact stated below; and report any compromised credentials to the relevant issuing authorities.

7. Reliance on the Services shall only be to the extent specified by any applicable service level agreements listed below. Use without such agreements is at your own risk.

8. Your personal data will be processed in accordance with the privacy statements referenced below.

9. Your use of the Services may be restricted or suspended, for administrative, operational, or security reasons, without prior notice and without compensation.

10. If you violate these rules, you may be liable for the consequences, which may include your account being suspended and a report being made to your home organisation or to law enforcement.

<Insert additional numbered clauses here.>

The administrative contact for this AUP is: {email address for the community, agency, or infrastructure name}

The security contact for this AUP is: {email address for the community, agency, or infrastructure security contact}

The privacy statements (e.g. Privacy Notices) are located at: {URL}

Applicable service level agreements are located at: <URLs>