Guidance for Notice Management by Proxies

Publication Date	2024-12-13
Authors:	David Groep, David Kelsey, Maarten Kremers, Nicolas Liampotis, Hannah Short, Arnout Terpstra, Catharina Vaendel
Document Code:	AARC-G083
Supported by:	the AARC TREE project
Publishing Organisation:	AARC Community
DOI:	10.5281/zenodo.14452340

Abstract

This guidelines document streamlines the presentation of user information notices (such as acceptable use policies or GDPR privacy notices) and how to support their presentation in infrastructures built on the AARC Blueprint Architecture (BPA) model. It provides sectoral recommendations based on four presentation models for notices and their expression by Notice Presentation Components, reflecting the diversity in business models in the infrastructures. The model recommendations include an aggregation-based representation based on machine-readable notice meta-data, a model based on common notices for infrastructures with central administrative coordination, a cascading model where presentation components can subsume responsibility on behalf of downstream service and data providers, as well as a fall-back scenario encouraging adoption of common presentation based on the WISE Baseline AUP model.

Since different research infrastructures deal with data of varying sensitivity levels, the model allows for a scalable level of control and verifiability, including the option of signalling acceptance of policies through attributes or claims. The guideline establishes a registry for notice identifiers and a resolver service for their meta-data, and pre-defines fundamental one-statement notice identifiers.

Table of Contents

1. Introduction 3
2. Objectives and Considerations 4
    Constructing notices and assigning responsibilities 4
    Stakeholders and their role 5
    General Data Protection considerations 6
        Notice management and protection of personal data 6
        Personal data and their presentation position in notices 8
        Access personal data and regulatory compliance 9
    Offline access and non-interactive (brokered) workflows 10
    Validation and compliance testing 10
3. Presentation models 11
    Machine-readable aggregated notices 11
    Common notice 11
    Cascading policy 12
    Coherent presentation 12
4. Recommendations 13
        Scenario-independent recommendations 13
        Requirements for each specific scenario 15
        Technical considerations 17
5. Notice meta-data and registry 18
    5.1 Policy identifiers for community purpose binding 19
    5.2 Relation to voPersonPolicyAgreement 20
    5.3 One-statement notices 20
    5.4 Meta-data document resolution 20
References 21
Appendix A Pre-registered identifiers 22
Appendix B Example meta-data document 23
    Example of a self-contained acceptable use policy 23
    Example of a community purpose binding statement for a community 24

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

1. Introduction

Services, data provides, and proxies must comply with laws and regulations, ensuring accountability and protecting both users and service providers. The open and collaborative nature of many research infrastructures (RIs) inherently introduces unpredictability regarding how individuals might use or, in some cases, misuse these systems. To manage this uncertainty, most services and infrastructures implement Acceptable Use Policies (AUPs) and privacy policies to inform users about permitted behaviour, safeguard against misuse and be transparent about the processing of personal data.

However, the reality for researchers working in such an ecosystem is that no single service operates in isolation. Research often requires leveraging multiple platforms, tools, and services, each with its own set of legal documents outlining terms of use, privacy considerations, and acceptable conduct. This fragmented landscape means that even something as simple as uploading a "HelloWorld.txt" file can be preceded by the daunting task of reading and agreeing to a maze of AUPs, privacy policies, and terms of service. This proliferation of documentation can be overwhelming, hindering productivity and creating confusion for users who are already navigating complex technical systems. Even worse, this practice results in users becoming resigned towards such documents, simply clicking ‘agree’ without reading them - which defies the whole purpose of presenting those documents to users in the first place.

Hence, the intent of this guideline is to

1. reduce the number of interstitial screens and acceptance ‘click-throughs’;
2. put in place a reference framework that enables proxy operators and service providers to rely on logged acceptance/consent information collected by peers operators in the workflow;
3. enhance consistency in the user experience for those cases where such inter-provider reliance is not sufficient for compliance or business reasons;
4. ensure end-users are aware of all notices and conditions that are applicable to their workflow(s);
5. support mechanisms to present updates to notices and conditions, from any party in the federation or workflow chain, with minimum end-user disturbance, while ensuring awareness of changes by the affected end-users.

This guideline is applicable both to community and infrastructure proxy [AARC-G045] operators (including hybrid proxies), as well as to service and data providers.

The guideline leverages existing underpinning components:

- WISE Baseline AUP [WISE-AUP]
- AARC-I044 Implementers Guide to the WISE Baseline Acceptable Use Policy [AARC-I044]
- GEANT Data Protection Code of Conduct (DP CoCo) v2 [COCOv2]

The considerations are reflected in a set of four presentation models, addressing the different types of workflows, using the composable WISE Baseline AUP as a foundational template.

2. Objectives and Considerations

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. This guidelines aims to achieve:

- the least number of user clicks and interstitial screens (Intent 1)
- to allow identification of a single point of presentation by assigning roles and responsibilities to a *notice presentation component*: a proxy operator, community, service provider, or data provider. It allows this component to act on behalf of connected services where appropriate*, i.e.,* proxy can take responsibility for its downstream services
- allow for service and data providers to have additional notices (terms and conditions) and have these presented by the proxy/notice presentation point so that it can be amalgamated with other notices but will definitely and demonstrably be shown (the additional elements in the Baseline AUP)
- that there shall be no need to show every possible notice to every user, clarifying the composition rules for the WISE Baseline AUP (discussed in AARC-I044 on notice composition), where - in absence of specific agreements - the noticeup-front notice can include connectable services (over just connected services for a community)
- address protocol specific requirements on user information, including *offline_access* requirements for OpenID Connect [OIDC-Core], on showing single aggregated notices (Intent 1)

Constructing notices and assigning responsibilities

The responsibility for presenting notices, for the content of the notice (or parts thereof), and responsibility for any processing of (access) personal data does not need to be vested in the same single entity. The equivalence of these roles was assumed in the construction model described in AARC-I044 for the WISE Baseline AUP. The introduction of multi-community and multi-tenancy proxies has changed that premise in two ways:

- the proxy does not equate the community, with multiple communities being hosted on the same proxy instance, or individuals registering first with the proxy in a generic mode (without specifying the community) and subsequently registering with one or more communities; and
- the proxy processing access personal data and the data controller (as meant in the EEA GDPR [GDPR]) of access personal data (elaborated below) are not the same entity, and the role of the proxy operator may change over time even for the same user (as the user takes on different roles during the registration life cycle in the proxy).

Stakeholders and their role

The **user** is the human end-user that is presented with notices before or during the execution of their workflows. The human user ought to be protected from a downpour of notice statements, as this both hampers usability but also lowers the attention paid to the notices presented (thereby lowering their effectiveness). Users organise themselves in structured **communities**.

The **notice presentation component (NPC)** owner-operator is the entity responsible for presenting notice(s) to the end-user. In a community environment there SHOULD be only a single point of interaction with the NPC (to meet Intent 1), but if only coherent presentation of notices can be achieved, there will be multiple NPCs and hence multiple owner-operators. In most cases, the NPC owner-operator will be the same entity as the one operating the (community) proxy as intended in the AARC BPA 2019, where the 'User notice' component implements the notice presentation component.

Service and data providers are responsible for their own risk assessment, and as a result are likely to stipulate terms and conditions to treat the identified risks. While the WISE Baseline AUP provides generic acceptable use conditions that intend to address the main common risks for service providers, there will be cases where additional terms and conditions apply (e.g. for accessing sensitive data or high-value services). These additional terms and conditions should be presentable in aggregated notices, even if the notice presentation component is not responsible for the content thereof.

Service and data providers often have bespoke notices, which have already been extensively reviewed by their own legal teams. These notices have to remain intact, and it is explicitly not the intent of this guideline to replace any such notices. In the context of this guideline, it is useful to be able to *identify* these policies, and allow them to be referenced, included, or augmented with other notices during presentation without changing their content in any way.

Leaving the definition and presentation of an AUP and other notices solely to the Community may be an overly complex task. A careful balance of service and data provider notices, proxy-defined notices, and community specific elements may enable the community to limit their role to defining the purpose of the community - with that purpose being incorporated into an amalgamated notice.

Hence, all parties should make sufficient notice information available for the notice presentation component to present a (composite) comprehensive notice. It is recommended that this follows the WISE Baseline AUP template, but where pre-existing notices are involved it may present (complementary) notices by reference or inclusion. When presented in an amalgamated way , this notice may contain elements from multiple stakeholders: the community, proxy operators, and service providers connected thereto. Specifically in case of machine-readable aggregated notices, having this information available during the initial connection by the user is essential, but this is similarly the case for common notices (one for the whole set of proxies and connected services) or notices that are combined in a non-automated way.

Terms and conditions on data should be conveyed by the service(s) that are delivering the data, and these requirements should be communicated (by service-specific means) if the

data is moved to another service or entity. Even if a service (or a proxy to which such a services connected) technically presents the notice, the responsibility rests with the data provider or data owner, who may delegate this task to the service delivering the data.

The user experience for presentation is left to the implementation of the notice presentation component and beyond the scope of this guideline.

General Data Protection considerations

Many communities, services, and proxies, as well as notice presentation components, will be operated in a way that must align with the General Data Protection Regulation (GDPR in the Euopean Economic Area) or equivalent local legislation. Depending on the structure of the community and the connected services and infrastructures, the basis on which personal data (specifically access personal data, as discussed below) is processed will vary depending on the relationship between user, community, and service. This guideline is intended for all notice presentations regardless of their GDPR legal basis, be it consent, performance of contract, legitimate interest, or otherwise.

The role of data controller for each processing may evolve over time, as the user enrols on a proxy platform, joins a community, or leaves a collaboration (as elaborated below). Such changes may be communited in-line on next-time access to a proxy or service that then engages the notice presentation component, or a change may initiate a proactive communication by the NPC (e.g. by email). To meet GDPR requirements, it should be clear at which time which specific notice has been shown. Hence a notice needs to be both versioned and its acceptance timestamped.

Ultimately, service and data providers are responsible for compliance and risk assessment for (access) personal data, like they are for other aspects of information security management. How to assess data processing risks has been discussed in complementary AARC guidelines, in particular AARC-G016 “Recommendations on the exchange of personal data in accounting data sharing” and AARC-G042 “Data Protection Impact Assessment – an initial guide for communities”.

Notice management and protection of personal data

The Implementers Guide to the WISE Baseline AUP (AARC-I044) recommended a presentation and aggregation model for the user presentation based on the premise that the notice presentation component itself holds an authoritative position. In the jurisdictions where GDPR applies, this means that the entity presenting the combined notice document (the entity managing the notice presentation component) either is, or is acting authoritatively on behalf of, the data *controller* – or for more than one controller at the same time. Since it is the intention of this guideline that this is the *only* entity with which the user will interact for the purposes of notice management (including any privacy notices incorporated into the main notice by reference), it has to take responsibility for presenting these privacy notices – and for recording their acceptance.

This guideline therefore explicitly recognised three different relationships between the community and the entity responsible for the notice presentation component (NPC):

- the community *is* by construction the same entity, or the same entity (organisation or administrative domain) is authoritative for both the notice presentation and for the community (a 'conflated' situation);
- the NPC is run *on behalf of* and *under the control of* the community (which may be represented by the organisation that is the host of the principal investigator). In terms of the GDPR, this means that the host organisation acts as a data controller, and the NPC operator is a data processor
- the NPC collects information itself, and offers to host communities on its platform, but means and purposes of the hosting platform and the NPC are entirely determined by the owner-operator of the proxy. Communities may use this service, and may even contribute resources to its maintenance and operations, but they have no direct say in how many resources are allocated to the operations or how its purpose evolves. In GDPR terms, the entity operating the NPC is itself the sole controller.

Hybrid forms are possible, in particular enrolment flows where the role of controller changes depending on the type of interaction (enrolment on the platform, joining a community) and during the user life cycle on the platform. An example of such a hybrid (evolving) controller change is shown in Figure 1.

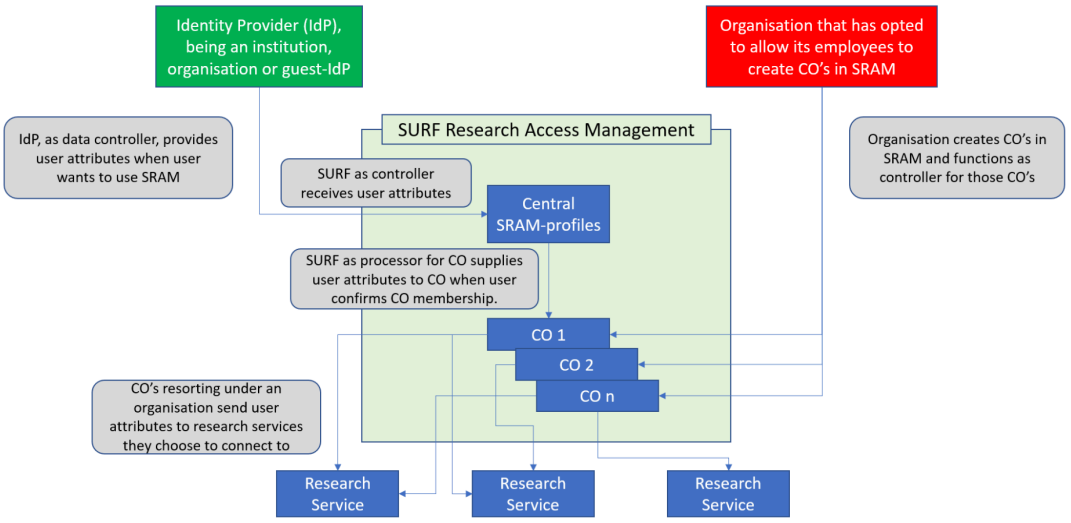

Figure 1: The SURF Research Access Management (SRAM) platform, a combined community-infrastructure proxy, is an example of a hybrid controller model. On enrolment, users join the SRAM platform for which SURF determined means and purpose of processing, and hence is the controller for the 'central profiles' managed on the platform. However, when users join a community (Collaborative Organisation), the principal investigator (acting as an employee of a host organisation, which has delegated management rights to the PI) determines which services connect to the community, how such services are provisioned and remunerated therefore, and what data is transferred.within the community. The PI's host organisation therefore becomes the controller (in GDPR sense) of the community data, even when this data is hosted in the same SRAM platform.

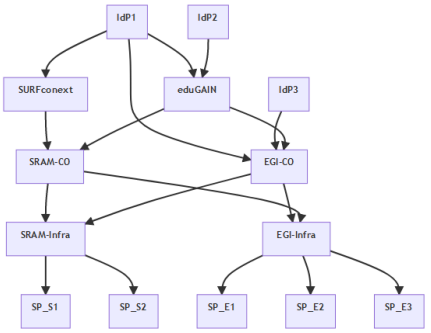

In some scenarios the operator of the community proxy (the proxy most suitable for presenting notices to users, since it holds the common personal identifier of the user across all 'downstream' proxies and services) and the controller(s) of the processing(s) involved are not directly connected by a bi-lateral contract. This could be because of the interposition of one or more infrastructure proxies, for example. In absence of (hierarchical) supply-chain communication, the controller would have no auditable proof that the relevant (art. 13 GDPR) information has been provided to the user, nor a way to signal to the presenting (community) proxy that the notice content has changed. Even if communication has been established out of band, propagating this information may be both time-consuming, cumbersome, or otherwise pose an unreasonable burden on the intermediaries. It is likely that in this case each Controller (at the point of first interaction with the user, likely at service provider level) will present an interstitial message and hence interrupt the user experience.

Personal data and their presentation position in notices

Infrastructures, including proxies, service providers, and data providers, may process different types of personal data:

- 'access personal data', i.e. personal data used *only for the purposes of enabling access to the Service* (as set out in the purpose limitation of the GEANT Data Protection Code of Conduct v2, appendix 1.B), including the activities listed in Appendix 3 "Purpose limitation and data minimisation" therein; and
- (research) content data, including personal data that form part of the content of data sets, (research) results, collections, &c, but that are not *in themselves* used to grant access.

The WISE Baseline AUP template includes specific placeholders for (references to) privacy notices. Entries (URLs) in this section are intended to refer to notices regarding '*access personal data'* specifically. Between interoperable infrastructures, (lists of) these privacy notices (of policy class 'privacy') should be used for access personal data, and *not* for content data.

Where content data contains personal data (for example social sciences survey results, clinical patient data, human genome data, &c), terms and conditions pertaining to that data, including privacy-related information, should be carried as 'supplementary terms and conditions' in the WISE Baseline AUP template (at the placeholders intended therefore, such as "*<insert additional numbered clauses here.>*" after the numbered clauses). Such notices and policies should also get their own unique identifier, separate from the processing of access personal data, even if the controller for both types of data is identical. Their processing class is 'conditions', rather than 'privacy', even if they only contain personal data protection information.

Access personal data will be transferred between proxies and between a proxy and a service provider as part of the user workflow. Such a transfer of access personal data to a Service Provider Organisation within the scope of the GEANT Data Protection Code of Conduct v2 usually constitutes a controller-to-controller transfer. The (aggregate) notice(s)

presented to the user must indicate the controllers that are responsible for each part of the processing, by including their privacy notices in the list, and designating the contact point for each processing (in the notice meta-data, these are the 'privacy_contact' endpoints, or, if these are not defined, the generic 'contact' endpoints). The list of controllers may be collected by other means and presented by the proxy or notice presentation point, as long as each processing is unambiguously associated with a notice and contact point.

Access personal data and regulatory compliance

Regulatory compliance, for instance with the General Data Protection Regulation (GDPR) in the European Economic Area, or with comparable legislation in other countries and regions, places specific requirements on the presentation of notices. In almost all cases, presenting a clear and unambiguous privacy notice is required, and this notice should meet specific semantic requirements.

In the scope of this guideline, we will consider privacy notices related to the processing of 'access personal data' as defined above. These have to be presented *at the time when personal data are obtained*. Also, it is standing practice that these notices should be *no more than two clicks away* from the point of data collection or from the main (web) presence. This constraint has to be implemented at the notice presentation point, regardless of the (sub)structure of the underlying mesh of proxies and service providers.

The advantage of amalgamated notices is that, once such a combined notice has been presented to the user ('data subject'), the data subject 'already has the information' that is relevant to the processing, and does not need to be informed again (GDPR Art. 13 par. 4). Hence, the request for affirmation by the user must be done *as early as possible*, when (most of) the information needed to construct the notice is available, and *be as encompassing as possible*, so that *all* relevant data processing information can be provided to the user.

This strongly suggests that, in the AARC BPA 2019 [BPA] composite model, the *Community Proxy* shall present the notice, including as much information as possible for connected service and data providers and subordinate proxies.

Updates to the notices, and (scheduled) changes to the processing such as a change of data controller or scope of the processing, may be communicated in different ways depending on the constituency. In practice, industry appears to converge on a dual-presentation model, where updates are sent out-of-band to a known contact address of the user, and *in addition* presented on the first subsequent use of the service. The latter satisfies the (real or perceived) need to record the affirmation by the user for audit purposes, which may be required depending on the basis for processing (for consent as a basis, this consent must be demonstrable).

Determining whether demonstrable affirmation is required for a processing is beyond the scope of this guideline. However, if a service or data provider or proxy connected to a community proxy expresses the need for demonstrable re-affirmation on changes, the community proxy is the most appropriate place to collect re-affirmation.

Where possible, updates to notices should be sent by out-of-band communications, if reliable contact information for all users has been obtained before. Sending such information

notices is in the legitimate interest of the community proxy as it aids in meeting the regulatory obligations of itself and its connected 'downstream' entities.

Communications about updates to privacy notices should be combined with other updates (to the AUP, the purpose of the community, or to terms and conditions), and - if re-affirmation to the AUP is to be collected periodically - be part of the applicable privacy notice in the conventional way (through the privacy notice links page).

Offline access and non-interactive (brokered) workflows

Activities that occur or continue to execute without user presence may require specific notices to be presented to the user to prevent abuse and unintended consequences. This in particular applies to OIDC flows that use the *offline_access* scope for requesting refresh tokens, as defined in the OpenID Connect Core specification section 11 (https://openid.net/specs/openid-connect-core-1_0.html#OfflineAccess). In these cases, the OP (including a proxy) either MUST or SHOULD have explicitly received or have consent for offline access, depending on the application type as stated in the OIDC Core specification. To ensure the user does not need to be presented with an interstitial consent page, this request for explicit consent MUST be presented as part of the initial set of notices if the user workflow(s) are, or likely are, requiring offline access. The service provider or proxy requiring offline access that is capable of presenting notice meta-data should signal this requirement in that way, and any service or proxy may (in lieu of or in addition) signal this by out-of-band mechanisms (e.g. in explicit agreements) for *single common notice* and *cascading notice* points in upstream proxies.

Validation and compliance testing

Validation and compliance testing of implementations of this guideline depend on the presentation model employed by the scrutinised proxy or service or data provider. Of these, neither the 'coherent presentation' model nor the 'single common notice' are automatically testable, and no accreditation body for such notices is designated by this guideline.

Notice presentation components that claim compliance with this guideline and select one of the machine-readable models can be tested and validated. The validation should consist of:

- verifying that the notice has been assigned a unique identifier (URI). This URI should preferably be listed in a well-known endpoint (to be assigned in a complementary guideline by the AARC Architecture working group).
- verifying that the URI is registered in the AEGIS notice registry
- verify that the notice meta-data (JSON) document is retrievable
- verify that the notice meta-data complies with the definition and semantics described in section 5 of this guideline.

The verification extends solely to the technical elements of notice presentation. It does not ensure either completeness or legal and regulatory compliance.

3. Presentation models

While the intent of this guideline favours a ‘click once’ approach, it is not practical to assume that all service providers and proxies can operate in the same manner. Differences in jurisdiction, sensitivity of the data and services provided, cultural and social elements, as well as historical context will decide to a large extent what can be achieved, but also what *should* be achieved (in the desired final state).

This guideline therefore defines a ‘waterfall model’ (decision tree) to select the appropriate notice presentation model, with four options. These options, ordered by scalability for interoperation and reduction of the number of user clicks, are:

Machine-readable aggregated notices

Acknowledging that a representative workflow for a community involves many stakeholders and is based on infrastructures wording across administrative domains, in order to improve user experience all relevant notices should be aggregated together and presented once as a common combined notice. Hence, the optimal model for a (new) community is to have a machine-readable notice suit where such aggregation can happen automatically.

- In this scenario, it can be signalled up and downstream which specific (lines of a) policy document(s) the user has already agreed to. Other parties in the chain can then decide whether that is enough or whether they want to present the user with more (lines of) policies for agreement
- This fulfils the largest number of intents, specifically Intents 1,2,4, and 5

While the “Common notice” model (listed next) is conceptually simpler and - at least at the start - more consistent for the user, not having such a single policy be machine-readable may limit the ability to mesh with other communities and new services.

Common notice

In highly coherent communities with common (delegated) management, all parties agree to a common policy set. This fully implements Intent 1:

- One single set of policies ‘to rule them all’ - combined by the proxy on behalf of all its connected participants and itself
- if all proxies and service and data providers adhere to the same (set of) policies, *by construction* the notice presentation is coherent and needs to be shown only once (since it is the same notice at each point).

Signalling may be supported, but existing management structures between the stakeholders allow to collect this formation as-needed.

Changes are commonly agreed to by the single management body, but achieving administrative coherency is complex to achieve in a reasonable amount of time.

Cascading policy

The ‘cascading’ model where the notice presentation reflects all known downstream providers and proxies, and where mechanisms are in place to signal updates upstream so that the notice is (manually or automatically) updated. This satisfies intents 1, 3, and 4.

- The proxy (each administrative domain) ensures all parties connected to it (belonging to that administrative domain) adhere to the same policies, determined by that administrative domain
- By doing so, users going to and through the proxies have to agree to only 1 (set of) polic(y)(ies)
- Any change in a policy by a downstream service must be accepted by the proxy and conveyed to the user if appropriate.
- No solution for multiple proxy scenarios. So this makes the problem smaller but does not necessarily solve it entirely

Signalling may be included, but needs an out-of-band *agreement management mechanism* to complement any technical information exchange.

Infrastructure proxies are natural aggregation points, and when infra proxies are part of the chain then they SHOULD absorb responsibility for their downstream service providers.

Coherent presentation

This presentation model does not *in itself* reduce the number of interstitial screens, but merely encourages a resognisable ‘look and feel’. This satisfies only Intents 3 (consistency) and 4 (awareness)

- The proxy presents all policy information in 1 place and manner to their users
- It is no ‘solution’ to the multiple-notice problem, but it is an improvement
- It does not address notice presentation reduction in multi-proxy scenarios

This is - in terms of user experience - the least-preferred model.

4. Recommendations

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. These identifiers MUST be unique, and SHOULD be registered with the maintaining registrar.

Where specific agreements exist or can be established between a group of service providers and proxies that could allow for a single **common notice**, such a single common notice MUST be used, and be presented in the same way by all entities which with a user may come into first contact. This is usually at the Community Proxy as per AARC BPA 2019. To enable interoperability with other infrastructures and service providers, an identifier and machine-readable meta-data for this notice SHOULD be provided to the registry.

When no single common notice can be agreed, notice presentation points MUST present **machine-readable aggregated notices** from the known connected entities (such as services and data providers) with which the user is expected to come into contact. This SHOULD follow the model recommended in AARC-I044, and it is RECOMMENDED to use the WISE Baseline AUP construction model to combine the (aggregated) purpose of the collection of services, its acceptable use (as per the WISE Baseline AUP), supplementary terms and conditions and specific notices, and a reference to (a list of) privacy statements.

This aggregate notice SHOULD be constructed automatically, using the machine-readable identifiers to construct the notice. If the notice presentation point is aggregatable (e.g. it can be meshed with other proxies or services), a notice identifier MUST be assigned to this combined notice, listing the included policies in the notice meta-data (includes_policy_uris).

When automatic aggregation is not feasible, the proxy or notice presentation point SHOULD present a **cascading policy**, and take responsibility for establishing agreements with all directly connected entities (proxies, service providers and data owners). Any connected downstream entity in this system MUST do the same, thereby establishing a chain of responsibility.

If none of the above is feasible, a proxy, service provider, or notice presentation point SHOULD use the WISE Baseline AUP template and follow the practice of AARC-I044 and thereby strive for **coherent presentation**.

Scenario-independent recommendations

- use of the WISE Baseline AUP is RECOMMENDED such that the number of different notices across the infrastructure can be reduced. The WISE Baseline model consists of
    - a purpose binding,
    - an acceptable use preamble, followed by
    - the immutable WISE Baseline AUP statements (being a limitative enumeration of permissible use), followed by
    - supplementary terms and conditions (e.g. constraints on usage of data and services, followed by
    - notices regarding the protection of the data sets and services that are provided, and service guarantees or provided capabilities), followed by

  - notices on processing of access personal data and the contact points of the responsible parties, and
  - the (list of) authorities responsible for the content of the notice.

- the notice presentation component MUST record, collect and retain time-stamped information related to notice presentation, its acceptance by a (human) end-user, and either preserve or be able to unambiguously describe the notice(s) presented. A list of registered notice identifiers (URIs) shall be sufficient information to unambiguously determine notice content and version.
  The notice presentation component SHOULD be able and willing to provide logged meta-data regarding this presentation (though not necessarily the documented evidence supporting it) to entities directly-connected downstream of the proxy or notice presentation component that have a legitimate need to verify this information in order to prevent repeat presentation of notices.

- a notice presentation component that records logs about notice presentations SHOULD collect and retain necessary user contact information (such as an email address) in order to contact the user (i) when the presented notices change, either directly or because notice elements that were used to construct an augmented notice were updated, (ii) in case of incidents affecting the user's record of notice presentation.

  The notice presentation component MAY inform the user pro-actively (e.g. by email) when it evaluates that the presented notices have materially changed or when the notice_refresh_period for any of the known included policies has expired. It can then autonomously re-present notices and record the updated presentation, and henceforth signal or continue to signal its presentations via *voPersonPolicyAgreement* (addressing Intent 5) [VOPERSON].
  If the user is not pro-actively requested to review updated or expired notices, such notices MUST be presented at the next interaction of the user with the proxy or notice presentation point.

- When, *within* a single notice presentation component, the formal responsible entity (such as a data controller) for a user either changes or one is added, this notice presentation component MUST inform the user of this change and, if it uses the aggregated notice model, construct a new notice indicating at least the new (or subsidiary) responsible entity.
  For cascading notice presentations, the list of responsible entities that would be presented to the user MUST be updated, and the user SHOULD be informed as soon as practical of the new or additional responsible entities (out of band or on the next interactive connection).

- a service provider or proxy that is intentionally connected up-stream to an AARC BPA proxy MUST inform all accepted and acknowledged upstream proxy operators regarding changes to (i) the processing of access personal data, (ii) the permissible purposes (WISE AUP purpose binding), and (iii) service-specific terms and conditions, which include changes to the processing of (personal) data managed by

or within the service or connected services.
For machine-readable policies, this information SHOULD be tagged with its 'policy class' (i.e. one of the above change types).

- A notice presentation component that has presented the immutable elements of the WISE Baseline AUP to the user, and that is capable of conveying accepted policies via *voPersonPolicyAgreement,* MUST include the WISE Baseline AUP notice URI (https://wise-community.org/wise-baseline-aup/v1/) in the asserted list of accepted policy agreements.
  This presentation component, if it publishes notice meta-data, MUST include this WISE Baseline AUP notice URI in the list of 'included_policy_uri' values.

- to facilitate exchange of update and revision information, a proxy operator MAY use the policy identifiers as described in section 5, and SHOULD make this information available via a well-known end-point of which the upstream proxy operators are informed.

- notice presentation components SHOULD aggregate pages for notices for presentation to the user, following the mechanism described in AARC-I044.

Requirements for each specific scenario

The use of **machine-readable aggregated notice** is recommended. In this case:

- The list of machine-readable policies MUST be the same for all users of the service and there SHOULD be one location to retrieve policy notice information per proxy, service or data provider.
  This guidelines acknowledges that this limits the flexibility of the aggregation by proxies and connected parties, but it will serve most cases. The advantage of having a single, static, endpoint for retrieval, and the ability to cache on a per-proxy (and per-presentation-point) basis outweighs the flexibility that would be offered by an API-based per-user or per-transaction retrieval of information. The latter both incurs too many round-trips, as well as complicates the reconstruction of the information that was presented to the user (needed to satisfy e.g. regulatory requirements for privacy notice presentation in some jurisdictions).
- When sending claims or attribute assertions towards downstream connected services and proxies, the proxy MUST include identifiers of all the policies to which the user has agreed, using the assigned policy identifiers and send that via the *voPersonPolicyAgreement* AVA or *voperson_policy_agreeement* claim (multivalued).
- When a service or proxy keeps persistent state about a user, and as part of a transaction receives a *voPersonPolicyAgreement* identifier from a trusted party, it SHOULD associate this information with the user state, and MUST NOT present notices that are (according to that receiving service provider or proxy) materially equivalent to notices that the user has already received.
- When notice presentation component is aware that it connects downstream services that may require offline_access (as intended in the OIDC Core specification), it MUST signal this explicitly, by way of augmenting the one-statement notice identifier defined herein to upstream identity providers and proxies by means of the '*augments_policy_uris'*.

If a proxy receives information (in any form or by any method) from a downstream service or proxy that offline_access is required, AND the proxy is capable of presenting notices to users, then this proxy MUST either retrieve from an upstream proxy the confirmation that a notice has been presented, OR inform the user that offline_access will be used downstream and include the one-statement policy identifiers for offline_access in its *'included_policy_uris'*.
This proxy SHOULD also signal this requirement upstream.

When a **common notice** is used:

- there shall be one authoritative entity that assumes responsibility for the common notice. That authority shall be vested in the notice presentation component responsible by all subordinate proxies and service and data providers in a suitable way, and (in GDPR jurisdictions) meet the requirements of all controllers of personal data.

When a **cascading notice** is presented:

- The list of machine-readable policies MUST be static at any point in time, MUST be the same for all users of the service.
- When sending claims or attribute assertions towards downstream connected services and proxies, the proxy SHOULD include identifiers of all the policies to which the user has agreed, using the assigned policy identifiers and send that via the *voPersonPolicyAgreement* AVA or *vo_person_policy_agreeement* claim (multivalued).
- When a service or proxy keeps persistent state about a user, and as part of a transaction receives a *voPersonPolicyAgreement* identifier from a trusted party, it SHOULD associate this information with the user state, and MUST NOT present notices that are (according to that receiving service provider or proxy) materially equivalent to notices that the user has already received.

For **coherent presentation**:

- The notice presentation component SHOULD assign a unique registered identifier to the notice presented to the user, and this identifier SHOULD be included in the *voPersonPolicyAgreement* attribute in any statements and claims that could include such an attribute or claim.

- When sending claims or attribute assertions towards downstream connected services and proxies, the proxy MAY include identifiers of all the policies to which the user has agreed, using the assigned policy identifiers and send that via the *voPersonPolicyAgreement* AVA or *vo_person_policy_agreeement* claim (multivalued).
- If the notice presentation component can positively confirm that the user has been presented (and confirmed acceptance of) other notices, and can obtain log information thereof, the list of notice URIs MAY be extended to include identifiers of any notices the user has already agreed with or has been presented with and confirmed. This attribute or claim MAY be sent to connected services and down-stream proxies.

- service and data providers, proxies, and notice presentation components SHOULD consult the list of *voPersonPolicyAgreement* URIs and SHOULD NOT present notices that have already been accepted or seen, to spare the user clicking too many policies.

Technical considerations

- This guideline establishes a registry for notice references (akin to the registry of RFC 6711 [RFC6177] for assurance) that can be conveyed by any such mechanism. The mechanism to exchange notice/policy requirements is out of scope of this guideline.
- This guideline proposes an errata for the voPerson schema maintained by REFEDS, extending the permissible semantics for the *voPersonPolicyAgreement* attribute to be a URI rather than only a URL. We acknowledge that, while such an errata is under consideration, implementations following this guideline may be using the voPerson schema outside of its stated definition.
- A service provider or proxy can optionally express a requirement on 'freshness' of acceptance to upstream notice presentation points through the 'notice_refresh_period' statement in the JSON meta-data document. Apart from acceptance of the policy itself, this is the only requirement that can be conveyed upstream through the notice meta-data statement.
- A service provider or proxy that requires, or is likely to require in all practical use cases, offline access in order to request OIDC refresh tokens (or equivalent mechanisms in other protocols), and which is notice meta-data capable, MUST include the one-statement notice URI "<NAMESPACE>:policy:notices:one-statement-notice:requires_offline_access" in the list of *augments_policy_uris* to signal the need for offline access for refresh tokens, and SHOULD provide an explanation for the need for offline access in its *description*. A notice presentation component incorporating such a service provider or proxy MUST present an explicit statement that offline access will be used (and present the description provided in its usual presentation location)

5. Notice meta-data and registry

This guideline establishes a registry of policy identifiers, specifically identifiers for acceptable use policies, terms and conditions, privacy notices, and compliance policies.

Identifiers registered as a result of this guideline MUST NOT refer to policies specifying a level of assurance - these must be registered with IANA in accordance with RFC6711.

Identifiers registered under this guideline require to:

- Be in the form of a URI,
- Be assigned a name, being a string uniquely and unambiguously identifying the notice for use in human presentation, and in protocols where URIs are not appropriate, and
- Include a resolvable http or https informational URL pointing to a JSON document containing additional structured information.

The JSON document returned at the informational URL SHALL include at least:

- **id** (required, single-value): string containing the URI of the identifier for the policy
- **aut** (recommended, single-value): URI identifying the authority governing this policy. It is recommended to use identifiers assigned by a recognised naming agency (such as a LEI-based URN) or long-term stable URL to the main web presence of an organisation. For privacy notices as meant in the EU GDPR (policy_class: “privacy#eu”), this SHOULD identify the data controller.
- **aut_name** (required, single-value): plain-text human-readable and disambiguating name of the authority (used in the WISE Baseline AUP preamble)
- **valid_from** (recommended, single-value): time from which this policy is in effect. This is expressed as Seconds Since the Epoch. When present, this value MUST increment whenever there is a minor change to the policy referring to this informational document. Note that major material changes SHOULD be assigned a new policy URI (id). Minor and major are defined discretionarily by the authority for the policy.
- **ttl** (optional, single-value): the time period after which this document SHOULD be retrieved again by consumers. This is expressed in seconds. In absence of this key, the document SHOULD NOT be retrieved more often than once a day; should be cached.
- **contacts** (required, multi-value), **security_contacts** (recommended, multi-value), **privacy_contacts** (recommended, multi-value): JSON arrays with one or more strings representing contact persons at the Entity. These MAY contain names, e-mail addresses, descriptions, phone numbers, etc. (incorporated from OpenID Federation 1.0) (used in the WISE Baseline AUP postscript)
- **policy_class** (required, single-value): string from the limitative enumeration (‘purpose’, ‘acceptable-use’, ‘conditions’, ’sla’, ’privacy’). The value ‘privacy’ MAY be qualified with a jurisdiction (e.g. ‘privacy#eu’). Jurisdictions SHALL use IANA ccTLD identifiers where possible. The jurisdiction value “eea” MAY be used to indicate the European Economic Area (e.g. “privacy#eea” will indicate a privacy policy in accordance with Regulation (EU) 2016/679 ‘GDPR’). Assigned subdomain names under the .int TLD MAY be used to indicate international organisations holding their

own jurisdiction (e.g. “privacy#cern.int” will indicate a policy in accordance with CERN’s OC11).
The policy_class ‘privacy’ applies to privacy policies governing service access data only (i.e. data used *for enabling access*, as meant in the REFEDS Data Protection Code of Conduct). Policies regarding privacy of the data processed in the service or in services connected to the proxy MUST be expressed in a ‘conditions’ policy_class.

- **notice_refresh_period** (optional, single-value): number of seconds after which this same notice has to be presented again to the same user, regardless of any earlier acceptance. Used to trigger periodic re-acceptance of e.g. acceptable use policies.
- **includes_policy_uris** (optional, multi-value): JSON array of policy URIs that are included in this policy and therefore implicitly fulfilled. Those policy URIs SHOULD be listed in the registry (which enabled automated composition by virtue of being machine-readable). This list MAY include also policies that are superseded by this policy, if the material content of deprecated policies is fully subsumed in this policy.
- **augments_policy_uris** (optional, multi-value): JSON array of policy URIs that are augmented by this policy, e.g. the WISE Baseline AUP itself. A presenting application MAY merge the presentation of this policy and any policies this policy augments.
- **policy_uri** (recommended, single-value): URL of the documentation of conditions and policies in human-readable form (incorporated from OpenID Federation 1.0)
- **description** (recommended): shortest plain-text human-readable description of the policy to be used for presentation in composite notices (used in the WISE Baseline AUP preamble).

All human-readable keys (aut_name, description) MAY be postfixed with a hash-sign followed by a locale code in RFC 4646 format (example: aut_name#nl_NL: “nationaal instituut”).

The registry follows the registration policy as set out in RC6711, with the authority of IANA subsumed by AEGIS, and the registered entities are policy identifiers rather than assurance profiles.

5.1 Policy identifiers for community purpose binding

The WISE Baseline AUP (https://wise-community.org/wise-baseline-aup/) model introduced the concept of ‘purpose binding’ (describe the stated goals and policies governing the intended use) for a community, agency, or infrastructure. Through this mechanism, the set of terms and conditions bind the person confirming the acceptance of an acceptable use policy based on the WISE Baseline AUP template to use Services only in a manner consistent with the purposes (and limitations) described in the community-specific preamble to the AUP.

In building and reading composite notices in an automated way, it may be advisable to retrieve the human-readable text of the ‘stated goals and intended use’ and associate a machine-readable identifier to this statement.

Normalised URIs defined in compliance with AARC-G069[1], including at least one <GROUP> and optionally made more specific by including zero or more <SUBGROUP>s, but not including any roles, are by construction valid policy identifiers. Such identifiers MAY be

[1] Guidelines for expressing group membership and role information

registered in the registry established under section 5, and be associated with an information URL pointing to the JSON document described therein.

The '{stated goals and policies governing the intended use}' in the WISE Baseline AUP template SHOULD be expressed as the description in the JSON document to which the information URL is pointing, and its policy_class MUST be 'purpose'.

Automated mechanisms to infer the information URL from a given AARC-G069 URI are beyond the scope of this document.

5.2 Relation to voPersonPolicyAgreement

The identifiers established by the registry can be used as values for the voPersonPolicyAgreement attribute in any information exchange.

5.3 One-statement notices

This document established the following pre-assigned identifiers for policies that may be used in augmentation and inclusion without needing explicit registration. Where notices need to include policies and requirements that are stipulated by these notice documents, the one-statement notices MUST be used. Appendix A pre-defined one-statement notice identifiers that are common and do not explicitly need to be notified to the registrar.

5.4 Meta-data document resolution

The resolution mechanism for meta-data JSON documents SHOULD be left to the AARC Architecture working group. The resolution set MAY be of the form of a HTTP GET request for a document at *https://nr.aarc-community.org/resolv/v1/<URL-encoded-URI>*. This URL MUST result in a 301 "Moved Permamently" response, and include a HTTP response "Location" header indicating the URL of the JSON meta-data document.

References

[WISE-AUP] WISE Information Security for E-infrastructures “WISE Baseline Acceptable Use Policy” https://wise-community.org/wise-baseline-aup/ (retrieved December 2024)

[AARC-I044] AARC Community “Implementers Guide to the WISE Baseline Acceptable Use Policy” https://aarc-community.org/guidelines/aarc-i044/

[COCOv2] REFEDS “REFEDS Data Protection Code of Conduct Entity Category” https://doi.org/10.5281/zenodo.6518055

[OIDC-Core] Sakimura *et al.* “OpenID Connect Core 1.0 incorporating errata set 2”, https://openid.net/specs/openid-connect-core-1_0.html

[GDPR] European Parliament and Council “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data” http://data.europa.eu/eli/reg/2016/679/oj

[AARC-G016] AARC Community “Recommendations on the exchange of personal data in accounting data sharing” https://aarc-community.org/guidelines/aarc-g016/

[AARC-G042] AARC Community “Data Protection Impact Assessment – an initial guide for communities” https://aarc-community.org/guidelines/aarc-g042/

[SRAM] SURF Cooperative “Easy and secure access to research applications for research collaborations” https://www.surf.nl/en/services/surf-research-access-management

[BPA] AARC Community “AARC Blueprint Architecture 2019” https://aarc-community.org/guidelines/aarc-g045/

[VOPERSON] REFEDS “voPerson schema specification” https://refeds.org/specifications/voperson

[RFC6711] Johansson, L. “An IANA Registry for Level of Assurance (LoA) Profiles” https://www.rfc-editor.org/rfc/rfc6711

Appendix A Pre-registered identifiers

The following notice URIs are defined by this guideline

- urn:geant:aarc:policy:notices:one-statement-notice:requires_offline_access
- https://wise-community.org/wise-baseline-aup/v1/

Appendix B Example meta-data document

The examples and resolver URLs are non-normative. Specifically, the resolver technical implementation and assignment of the resolver resolution end points is to be defined by an implementation guideline of the AARC Architecture working group.

Example of a self-contained acceptable use policy

This self-contained AUP implicitly also fulfils the requirements of the WLCG and EGI (joint security policy group) acceptable use policy, that are hence included policies:


```
{
  "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."
}
```


It might potentially then be retrievable from https://nr.aarc-community.org/resolv/v1/urn%3Adoi%3A10.60953%2F68611c23-ccc7-4199-96fe-74a7e6021815

This AUP is to be presented every 13 months to users, even if the policy itself does not change. A service that requires the user to have agreed to the ‘Joint SPG acceptable use policy’ (EGI document 2623) does not need to present that notice again, since the user already agreed to all relevant statements via the Nikhef AUP.

Example of a community purpose binding statement for a community


```
{
  "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"
}
```


It might potentially then be retrievable from https://nr.aarc-community.org/resolv/v1/https%3A%2F%2Foperations-portal.egi.eu%2Fvo%2Fview%2Fvoname%2Fxenon.biggrid.nl

Following AARC-G069, the identifier could have been auto-completed, once a namespace has been defined for BiG Grid communities. In that case, the identifier would have been "`urn:geant:nikhef.nl:projects:biggrid:group:xenon`", with associated resolver URL https://nr.aarc-community.org/resolv/v1/urn%3Ageant%3Anikhef.nl%3Aprojects%3Abiggrid%3Agroup%3Axenon.

Even if agreement to the community AUP is to be re-confirmed yearly, for example because the infrastructure requires re-acceptance of the AUP after 12 months, the purpose notice *by itself* does not need to be reaffirmed by the user. Hence, there is *no* notice_refresh_period included herein. That would be set for the combined AUP, that merges the WISE Baseline AUP and this purpose notice.