



Authentication and Authorisation for Research and Collaboration

Attribute Authority Operations

A fond of knowledge worth protecting

David Groep

AARC Community Policy WG



EUGridPMA 49 meeting

2020-05-14

Operational guideline landscape for - proxy or source - AAI components

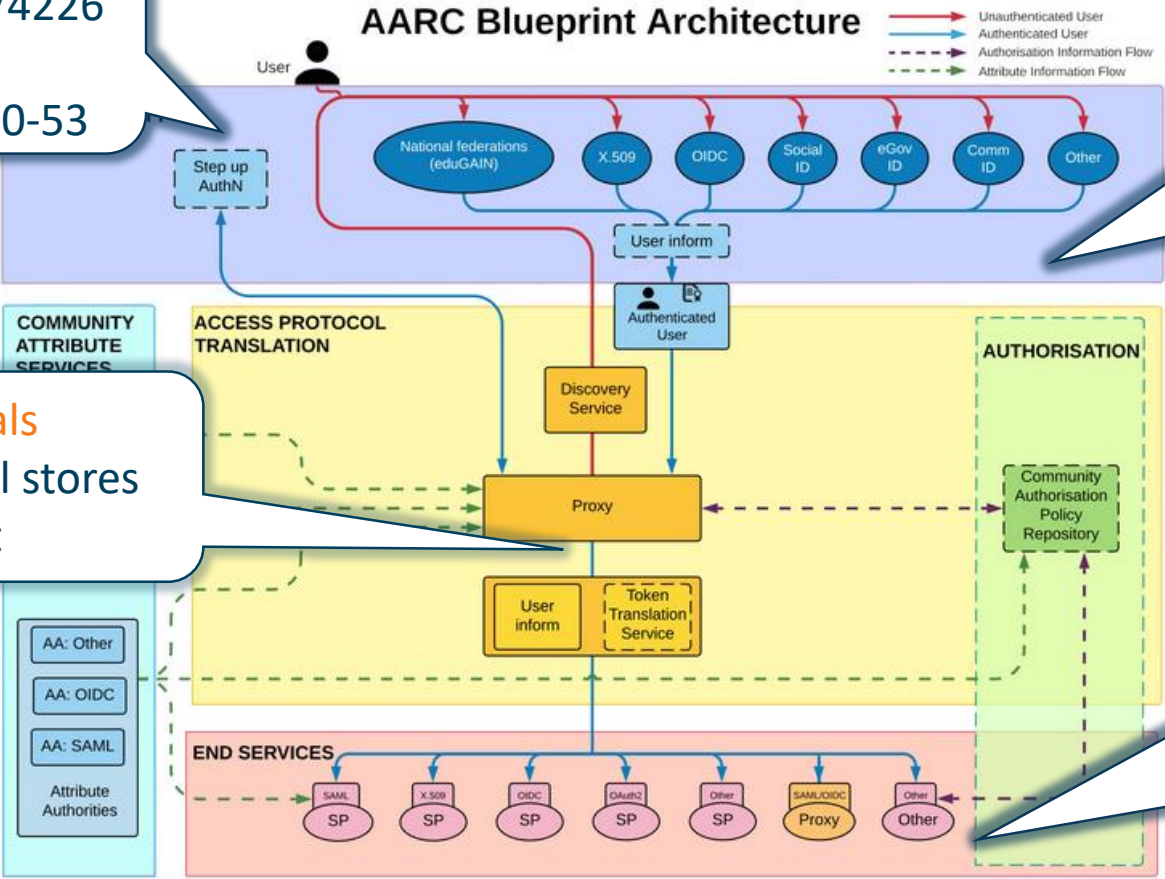
MFA
 RFC6238/4226
 FIPS140
 NISTSP800-53

Authentication/identity sources
 Sirtfi
 (eduGAIN) baselining
 IGTF AP Profiles
 NIST SP800-63
 eduGAIN sec. team workflow

Ephemeral credentials

- trusted credential stores
- protection at rest

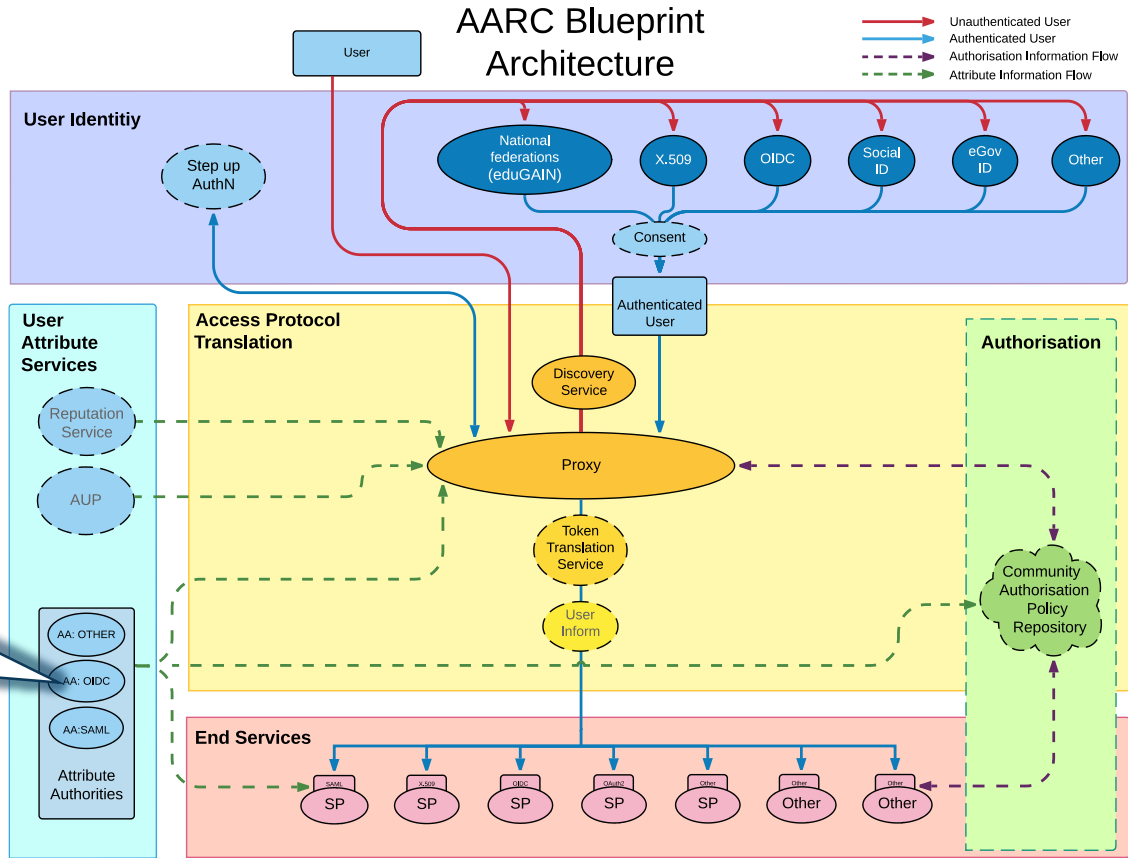
Service provider operations
 ISO27k
 Sirtfi
 Infrastructure response plans



Operational security focus in the BPA: beyond just the IdPs

Community membership management directories and attribute authorities

- integrity of membership
- identification, naming and traceability
- site and service security
- protection on the network
- assertion integrity



Guidelines for Secure Operation of Attribute Authorities and other issuers of access-granting statements (AARC-1048, in collaboration with IGTF AAOPS)

AARC-G048: keeping users & communities protected, moving across models




trusted delegation of response from communities to operators, and from services to communities in recognizing their assertions

Structured around concept of “**AA Operators**”, operating “**Attribute Authorities**” (technological entities), on behalf of, one or more, **Communities**

Guidelines for Secure Operation of Attribute Authorities and other issuers of access-granting statements

Publication Date: 2018-11-22
Authors: David Groep, David Ke Paetow, Maarten Kremers
Document Code: AARC-G048



3.3. Attribute Assertions

1. Assertions provided by an AA must be integrity-protected. They must be signed by the identified AA, or be transmitted over an integrity-protected channel where the server has been authenticated, and preferably both.

Push model

Where the protocol supports it, enable protection also of the messages conveyed over the established channel.
Good examples: SAML Attribute Query should enable message signing and use TLS.

Pull model

As a good example: LDAP should enable TLS protection of the channel

3.4.1. Key Management

1. A key used to protect assertions should be dedicated to assertion protection functions.

Push model

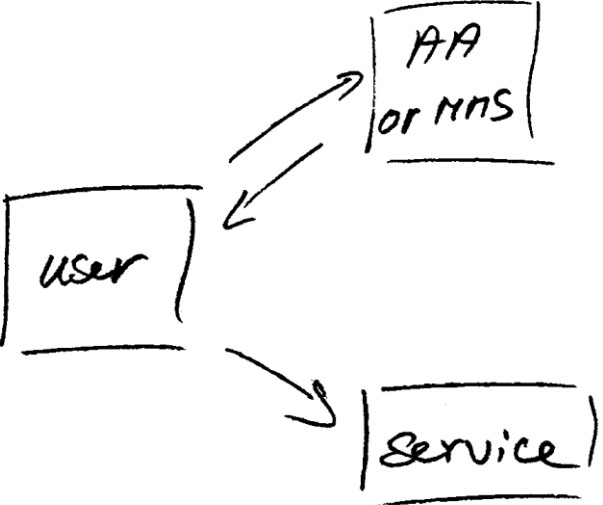
If the AA both signs assertions and provides functionality over protected channels, the keys used to sign assertions shall be different from those protecting those channels.

Pull model

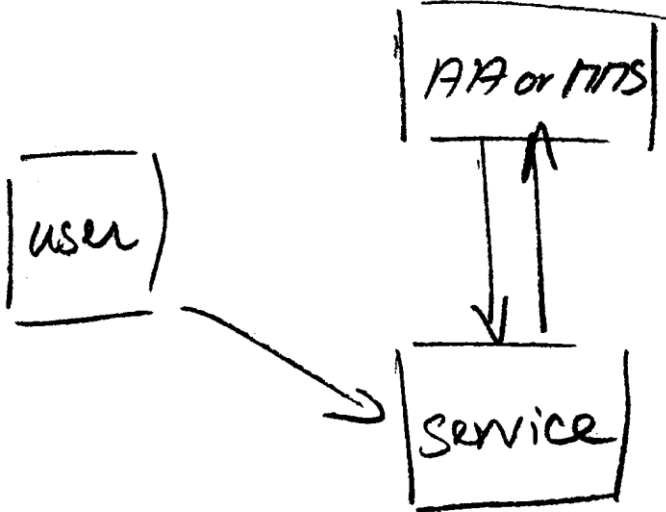
The key of the AA must be used solely for protecting connections to its protocol endpoint and ensure an integrity protected and mutually authenticated channel.

Protecting the community membership data and its proxy

- Intentionally targeted broader than just BPA-style communities, since operational security spans data centres and infrastructures using other forms of AA membership management
- PRACE: ‘pull model’ directory-based communities
- BPA: encourages ‘push model’ attribute-carrying service requests



*push model – the common BPA method
(e.g. SAML AttributeStatement, VOMS AC)*



*pull model – common when using directories
(e.g. LDAP in PRACE, GUMS in OSG)*

When the AA is in a managed environment ...

Many of the recommendations are already implemented 'implicitly'

- because common software implements it: e.g. signing SAML assertions and JWTs
- because a good data centre already has network monitoring and central logging in place
- because you signed up to Sirtfi (didn't you?) – so you collaborate in incident response
- because you have trained IT operations personnel looking after the service

And some are intuitive best practice

- like assigning a unique and lasting name to a group
- because implemented controls follow ought to be those that have been documented

Forward looking and specific requirements

Some controls are specific to AA operations and protect against current and future threats:

- minimum signing key length so that the community is not broken in the next few years (at least 112-bit symmetric, i.e. ≥ 2048 bit RSA keys)
- protect the key from data breaches, compromise, ransomware, and exfiltration by using HSM Hardware Security Modules or equivalent controls (and the HSMs you need are not that expensive, or you can even rent them in AWS...)

Or deal with commensurate incident response (you don't want just a big red button):

- re-issuance of attribute statement must be based on fresh data
- release them only in accordance with the community's policy and maximum life time
- require appropriate client authentication before releasing attributes to prevent data breaches
- for non-revocable tokens (like OAuth Access Tokens or PKIX 3820 proxies), limit life time < 24 hrs (for OIDC, these are anyway typically 15 minutes)

G048 AA Ops guidelines and AA hosting

Guideline was written with both physical and virtual deployment in mind

“An AA may be run in a virtual environment that has security requirements the same or better than required for the AA, and for all services running in this environment, and it must not leave this security context. **Any virtualization techniques employed** (including the hosting environment) **must not degrade the context** as compared to any secured physical setup. Only AA Operator designated personnel should have control over the virtualisation and security context of the AA.”

- if you can host it on-prem, the easiest solution is to host it on your security-service VM infrastructure (e.g. alongside your IdP, your AD, or your master LDAP servers) to limit guest compromise)
- If you run it in a cloud provider, select a provider that offers proper security and network controls, implement account role separation, and deploy the offered protections. E.g. in AWS you have *a lot* of controls available to do so. But Azure &co hve the same. – and rent a netHSM

Deployment guidance included ...

1. Assertions provided by an AA must be integrity-protected. They must be signed by the identified AA, or be transmitted over an integrity-protected channel where the server has been authenticated, and preferably both.

Push model

Where the protocol supports it, enable protection also of the messages conveyed over the established channel.

Good examples: SAML Attribute Query should enable message signing and use TLS.

Pull model

As a good example: LDAP should enable TLS protection of the channel

6. The network to which the AA system is connected must be highly protected and suitably monitored.

Service access should be protected by at least two distinct control layers not running the same software or operating system, and the AA system must not run any unnecessary services. The network should be monitored for anomalous events, such as detection of data exfiltration, credential probing, and brute-force attacks. It should preferably also be protected

Implementation of AA Operations Security guidelines

1. via AEGIS, major RPs and Infrastructures reviewed it in light of their the current (up-to-date) use cases and models
2. review showed the doc has some ‘hidden legacy’ (too specific in some cases, but then vague for other scenarios)
3. some acronyms were taken for granted (RP, OP, &c)
4. wording unintentionally conveyed that ‘physical on-prem’ was required
5. guideline put emphasis on community autonomy (and responsibility) which some infrastructures may want to see as a ‘business relationship’ with their community *but then, what ‘community’ means in a context is easy to mix up, as seen in the community membership management discussion in UK IRIS yesterday*

<https://wiki.geant.org/x/-YVgBw>

https://docs.google.com/document/d/1-hbqSpQegm7UaC_wupFzFMm19Q024UPkG-8Jwokkmzc

Discussion items: auditability, integrity, logging, and incident response

Section 3.7, item 2, **"The AA Operator must record and archive at least the follow for all its hosted AAs . . . (b) all issued attribute assertion . . ."**

GEANT (Christos): *This implies that the AA operator will have to keep in its logs all the personal data of the user. It makes more sense to log information about the identification of the (e.g. community identifier) and assertions that can be used for authorization, but not information like name, email address etc of the user.*

Section 3.6, item 6: **"The AA Operator must accept being audited following reasonable requests from a Community it serves and from relying parties [...]"**

GEANT (Christos): *This is something that is typically part of the contractual relationship between the AA operator and the community. Such a requirements has many assumption about the business relationships between the parties.*

On auditability and having the logs to ability to do incident response

2. The AA Operator must record and archive at least the following for all of its hosted AAs:
 - a) all requests for attributes
 - b) all issued attribute assertions
 - c) any configuration change to the AA relevant to the access control of the attribute repository
 - d) any change affecting the binding between subjects and attributes

Section 3.7, item 2, "The AA Operator must record and archive at least the follow for all its hosted AAs . . . (b) all issued attribute assertion . . ."

GEANT: *This implies that the AA operator will have to keep in its logs all the personal data of the user. It makes more sense to log information about the identification of the (e.g. community identifier) and assertions that can be used for authorization, but not information like name, email address etc of the user.*
EUDAT: *sees GDPR issues*

XSEDE (JimB): *Archiving the results of every LDAP query for 400 days seems excessive. (and besides, there are more absolute numbers in the document)*

Disaster recovery and response to compromise

Section 3.9 "The AA operator must have an adequate compromise and disaster recovery procedure and must be willing to disclose this to the hosted Communities or to either an assessor or all related relying parties"

GEANT: This is something that is typically part of the contractual relationship between the AA operator and the community. Such a requirements has many assumption about the business relationships between the parties. [echoed by EUDAT]

But beware of writing this in such a way that compromise and recovery leaves the relying parties in the cold! Importance of BC/DR and compromise handling is for RPs, who are typically left with absorbing all risk, cost, and liability, if not formally than at least in practice ... !

Section 3.3, item 5: "The AA Operator must only issue Attribute Assertions or release attributes to requesters in accordance with the Community policies."

GEANT (Christos): *This implies that a certain type relation between the community and the AA operator, which is not always the same. In the eduTEAMS Shared Service, the policies are owned and managed by GEANT for example.*

Section 3.6 "[publishing by AAOP] a web URL to a general information page about the community"

GEANT: *Again, this implies that the relationship of the relying party is with the AA operator and not with the community. This is correct in some case, but incorrect in others*

"not using project domain names unless due organizational care is taken"

XSEDE (JimB): Are we not allowed to use xede.org names?

Section 3.2, item 3 "The AA Operator must collect and publish the community documents for the benefits of Relying parties"

GEANT (Christos): *This recommendation assumes a certain model of operation where the Relying Parties have a relationship directly with the AA operator. This is not the case everywhere and actually in most of the cases that we know it is not like that.*

EUDAT (Sander): *I agree that there is not always a relationship between AA and RP.*

Section 3.3, item 2: "This may mean that AAs require client authentication, in addition to the encryption of the messages and the communication channel"

EUDAT (Sander): *IMHO client authentication should be recommended. This would enhance the data protection, too.*

About protections and isolation

1. An AA that issues attribute assertions must be a dedicated system, running no other services than those needed for the AA operations.
2. An AA may be run in a virtual environment that has security requirements the same or better than required for the AA, and for all services running in this environment, and it must not leave this security context. Any virtualization techniques employed (including the hosting environment) must not degrade the context as compared to any secured physical setup. Only AA Operator designated personnel should have control over the virtualisation and security context of the AA.
3. The AA must be located in a secure environment where access is controlled and limited to specific trained personnel.
4. The AA must be run with an intended continuous availability. Hosted Communities must be informed if AA Operator procedures change.

Section 3.4

GEANT: *The document seems to suggest that running services on dedicated systems is better than running services on a virtualised system and furthermore, that running service on your own data centre is more secure than using a public cloud. We believe this notion, might have been valid 5 years ago, but this does not reflect the current common practices. Furthermore, the statement "Only AA Operator designated personnel" seems to be in contradiction to running AA services on public clouds*

EUDAT: *This sounds like you need hardware for the AA. I do not agree with it. Of course they must be secured, but not dedicated hardware. This statement should be a dedicated subsection, e.g. server requirements.*

Enforcement by contract or by doing it: it is not clear enough in the doc

Section 3.5 "The AA Operator should document the physical site security controls and maintain them in a state consistent with the security requirements of the hosted Communities."

GEANT: *Again, to a reader it seems that the document implies that running AA service in your own physical location is better than running them on a public cloud. Of course there must be physical security control on the physical infrastructure, but the wording could be changed not to leave the read with the feeling that the rule is that they have to run their own infrastructure*

Yet also, by EUDAT: ***Is documentation available for communities and RPs?
This security controls are important to decide to trust the AA operator.***

"encouraged to consider using an HSM to store signing keys"

XSEDE (JimB): *is this a requirement or recommendation? Unclear and often infeasible.*

On explicit limits and specifications

Section 3.3, item 3:

"If an AA Operator issues Attribute Assertions containing a lifetime, this lifetime must be compliant with the Community policies, be no more than 24 hours, and the Attribute Assertion must not be valid beyond the validity period of the attributes it contains. The Community Management is responsible for the content of the Attribute Assertion, as issued, during its entire lifetime"

GEANT (Christos): *The document tends to be pretty high level, but at some points it becomes very specific. Where the requirements for no more than 24 hours is coming from? This should be something that should following the requirements of the communities and the relying parties.*
and EUDAT sees a potential conflict with community policies as well

Unclarity in the text and use case description

4. Relying Party obligations

1. If a Community uses AAs operated by multiple AA Operators then Relying Parties must assess each of the AA Operators individually.

EUDAT: *IMHO communities must ensure in this case, that they do not release contradictory attributes, specially because the RP must fetch all AAs.*

which indicates that the text is unclear. This was intended not model where they are augmenting, but as a multi-homes community with redundant AAs 😊

Is there an accreditation process for this document?

should infrastructures be mentioned as examples in the introduction?

Thank you
Any Questions?



<https://aarc-community.org>



© members of the AARC Community.
The work leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme and other sources.