Sustainability models for the AARC CILogon-like TTS Pilot and RCauth.eu

Publication Date: *DRAFT 20160506-04*
Grant Agreement No.: 653965
Work Package: NA3

Authors: David Groep, Nikhef

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

Abstract

In this study we have surveyed the various sustainability models that could accompany a production deployment of the AARC CILogon-like TTS pilot system. This system and its close derivatives are being piloted in major RIs and e-Infrastructures, and can bridge the gap that currently exists between institutional user identities, R&E identity federations, and the non-web compute, cloud, and data services and brokering that use client certificates, such as those in EGI, PRACE, and cloud-based data transfer services of e.g. Globus Online.

Table of Contents

1 About the AARC CILogon "TTS" pre-pilot

The AARC project is running a pilot with a bridging AAI solution based on the CILogon [CILogon] model to enable resources that use conventional identity and attribute certificates for access control to be used by researchers using exclusively federated credentials. While certificate-based access is effective for many non-web (command-line) and brokered-access (delegation) use cases, exposing this technology to a wide user base is seen as a significant barrier. In this pilot a set of mutually-interconnected third-party software components is composed to hide the technical details of certificate-based access.

It combines authentication using SAML-based identities such as provided by eduGAIN [SAML, EDUGAIN], public-key authentication certificates (PKIX) such as those coordinated by the IGTF [IGTF], the use of VOMS [VOMS] community membership management statements, and the OpenID Connect [OIDC] authentication protocol, used by many light-weight web applications (e.g. Globus Online [GO] and science gateways). Technical information is provided in more detail in [CILPPW].

Using the AARC CILogon-like Token Translation Service "TTS" pilot technology, infrastructures such as EGI and ELIXIR can implement AAI controls for their existing resources and services with SAML-based credentials in an end-user friendly way.

The pilot aims at demonstrating that a production service could be built based on these components that fulfills the following criteria:

- Ability to serve a large pan-European user base without national restrictions, and without having to rely on specific national participation exclusively for this service – thus serving the needs of cross-national user communities that have a large but sparsely distributed user base

- Use existing resources and e-Infrastructure services without the needs for security model changes at the resource centre or national level - i.e. transparent to the service providers ("relying parties") in the current federated non-web HTC, HPC, or IaaS cloud platforms that use client certificate based access controls

- Allow integration of this system into science gateways and portals with minimal effort, using only light-weight industry-standard protocols

- Permit the use of the VOMS community membership service attributes for group and role management when the portals and science gateways access the e-Infrastructure

- Concentrate service elements that require significant operational expertise or that expose confidentiality and privacy risks in a few places in the architecture, allowing these to be operated by expert partners whilst permitting the user community developing gateways and portals to focus on their research domain (and not burden these groups with the need to care for security-sensitive service components)

1.1 Elements of the AARC 'CILogon-like TTS' pilot

The elements of the CILogon TTS pilot are schematically shown in the figure below.

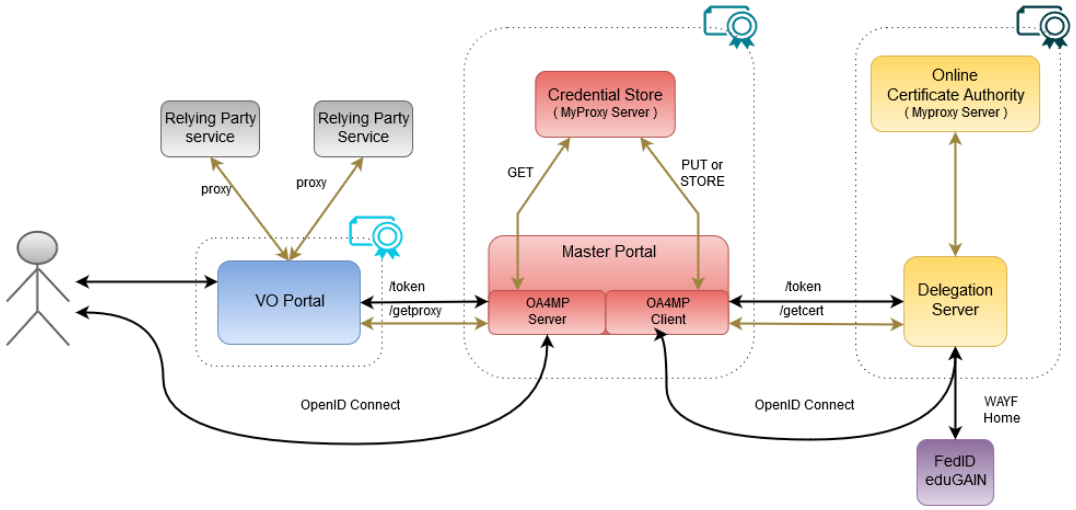

The different colours in the figure indicate the various separately identifiable components in the chain of authentication:

- in blue the VO portal (science gateway) serving the research community topical experts with domain-specific tools;
- in red the Master Portal hosting the credential repository on behalf of the end-users – taking care of the management of client credentials and for augmenting them (on request and with consent of the user) with community attributes and membership data;
- in yellow the Delegation Server and Online CA responsible for creating trusted credentials and a control point for classifying identity assurance levels and enforcing long-term persistence and non-reassignment of identifiers[1] in the system (and providing a revocation and correction mechanism for incident recovery); and
- in purple the general-purpose federated ID mechanism such as the one provided by e.g. eduGAIN via the national identity federations.

Each of these elements is a separate entity, and they can all be deployed in distinct administrative domains as long as sufficient bi-lateral trust between these domains can be established. The result in all cases is a system that provides a trust chain between the user (hidden behind the purple 'FedID' service), and the relying party (RP) services (the grey boxes) that can scale to arbitrary size based on assurance 'filters' that can be applied at the delegation service (for classification), in the master portal (when augmenting identity with community attributes), and at the resource level (making decisions based on a combination of both identity and community-provided assurance).

[1] Non-reassignment and traceability of end-user credentials were identified in the AARC survey of resource provider requirements on assurance (and documented as a baseline assurance profile in AARC's MNA3.1 document) as being of primary importance to research and collaboration, but are not guaranteed qualities of the identification provided at the moment through federated mechanisms such as eduGAIN.

2 Component usability and sensitivity considerations

Although the software and service elements composed in the AARC CILogon-like TTS pilot could theoretically be deployed and operated by a single entity for the whole world, or many times over by each team of collaborating researchers, neither extreme is an attractive solution. Each element has its own characteristics that need to be considered:

- The **VO portal or science gateway** is directly facing the researcher and end-user: it should integrate closely with the domain workflow and be 'intuitive' to its user. This can only be achieved in close collaboration with research domain experts, and this by nature results in a very distributed process. Many instances will emerge in practice. This component has therefore been chosen such that is does not contain any long-lived user credentials and does not have to store sensitive or personally-identifiable information. This makes compliance with data protection rules for user credentials straightforward. The VO portal can use user-originating credentials to communicate with the e-Infrastructure 'relying party' (RP) services, and will act towards them as if it were the user. With this delegated user credential also community membership assertions and other attributes will be sent. All of these are short-lived, and can be obtained by the VO portal only as a result of a specific user action – incidents are therefore limited to credentials being active in the portal during the time of an incident only. There may of course be other confidential or sensitive domain (research) data in the portal – such aspects are out of scope for this study.

  Apart from having to establish bilateral trust with a credential repository of choice (to host the user credentials), there are no specific considerations for this service element.

- The **Master Portal** is a bridging component between the identity service, any community assertion services (VOMS, not shown in figure), and the VO portals. It uses secure bilateral protocols to exchange information with both the VO portals and with the Delegation Service. Yet in order to fulfil its role effectively it also acts as a credential repository, and will hold long-term credentials for a (potentially large) set of users. This credential repository ensures that end-users can obtain command-line access to their credentials (for complete non-web use cases), and it provides a credential caching feature both for load reduction on the Delegation Server as well as for providing fault resilience for temporary outages. The credential repository is an element that will also permit refreshing of user credentials for long-running scientific workflows, even when the user is not physically present to re-authenticate.

  The Credential Repository annexed to the Master Portal is a service with a significant security profile: it potentially permits bulk access to large numbers of user credentials at the same time that can be used directly for access to large number of RP services – any service that any of its users has access to. However, it does not hold much personal data, and – through revocation of the user credentials from the Delegation Service – incident containment and recovery is possible.

The service must be run in a secured environment, in accordance with specific security guidelines and secure best practice[2]. It has to establish bilateral trust with VO portals (so that it will not delegate user credentials to rogue portals through OIDC), and with at least one Delegation Service (to allow it to authenticate users to it).

- The **Delegation Service** provides the actual token translation between the federated SAML user ID, the PKI certificate, and the OpenID Connect authentication by the user (via the Master Portal). Since this component will create omnidirectional PKI credentials for the user that may be long-lived (typically between 2 weeks and 13 months), and can request such credentials for arbitrary names, it is a highly sensitive component. In the PKI domain, it acts as a certification authority (CA) trusted third party and its credentials will be (and have to be) accepted as authoritative by all resources and service providers in the infrastructure. The **Online CA** is annexed to the Delegation Service but separate in order to ensure compliance with the minimum security requirements and IGTF PKI technology guidelines.

  In order to be acceptable as a trusted source of identity to the relying party RP services, it needs to be run in a highly-secure way, have very well documented policies that are compliant to infrastructure requirements, be externally reviewed and re-assessed periodically using e.g. a peer-review process, and avails over specific hardware security modules (HSMs) that protect tampering with the system signing data. Generally this is done via accreditation to the Interoperable Global Trust Federation (IGTF) for services in the major generic e-Infrastructures like EGI and PRACE in Europe, XSEDE and Open Science Grid in the USA, and most national e-Infrastructures in Europe.

- The **Federated Identity Service** is the established set of national federations that are interlinked through eduGAIN, although specific 'identity providers of last resort' can be added to the scheme as long as such are interoperable with the SAML2Int profile used by the federations and they comply with the basic assurance needs of the Delegation Service: the non-reassignment of 'user names'.

2.1 Where Are You From usability considerations

The (combination of) systems also needs to authenticate the end-user (having come in through a VO and Master Portal) at the federated identity provider (IdP), typically one of the many institutions participating in a federation or in eduGAIN. This is most imperative for the Delegation Service, since it connects to the existing federation space that already contains thousands of IdPs. To do that, it must present the user with a choice of IdPs, or be provided up-front with the (single) proper IdP – which may however be different for each user.

However, confronting the user with multiple dialogues where a selection of authentication sources is presented in confusing, and it is unlikely that the user will recognise any other entity besides the (research) community offering the service (the VO portal), and his or her home institution (the Federated IdP 'in the purple box'). The other elements in the chain should – if at all possible – not offer a 'where-are-you-from' choice, but have a many-to-one link (i.e. many VO portals talk a single Master Portal, but each VO portal talks only to one).

[2] AARC is planning to develop and evolve such guidelines, based on information security management practices and the IGTF Recommendations for the operation of a Credential Store (CS) by a trusted Operator (https://www.eugridpma.org/guidelines/trustedstores/)

3 Deployment models

The elements of the service each could be deployed in several ways: independently by each research group, by a research infrastructure, by a generic e-Infrastructure, by a consortium, or under contract by a (commercial) third party. For each of the elements the model choices are discussed.

3.1 VO portals

The VO portals are the most visible part of the system for end-users, and experience in the research area shows clearly that scientific workflows are highly specialised and their computing support almost universally tailored specifically to each domain and set of workflows. Rapid development cycles, on-demand instantiation and short 'product' life cycles are characteristic of many VO portals. Larger analytics platforms exist, but there is no single dominant one even within a single research community. It is realistic to assume that most VO portals will be widely distributed, and many will be run on an ad-hoc basis.

Integration of the VO portals with a system based on the pilot will have to be light-weight and be based on well-supported public domain libraries that do not incur a support burden on the (set of) operators of a CILogon-like TTS service. The choice of OpenIDConnect – a technique widely used for many open-source, community and commercial AAIs – places the maintenance of this software outside the responsibility of the CIlogon-like TTS service operator, as long as open standards are adhered to.

3.2 Master Portal and Credential Store

The Master Portal is a security-sensitive component, whose availability and reliability directly affects the VO portals connected to it. As such, it should earn the trust of the operators of the VO portals: they can usually connect their portal to just a single Master Portal. The current state of technology does not easily allow for many-to-many connections. Yet they have a choice: set up their own master portal, or join up with an existing provider.

Setting up a Master Portal in a way that is secure and sufficiently trustworthy for the infrastructure service operators is complex and may be costly. It requires secure service hosting, specific trained personnel, and a security incident response capability. It is likely that VO portal operators will look towards either their home organisation or to the larger community of research users in their field (e.g. at the ESFRI level) for a cheaper or more readily available alternative. Especially if the service is already present and trusted by the infrastructure relying parties, it is likely that new VO portals will prefer to associate with their 'natural home'.

Operating and maintaining a Master Portal will have an ongoing cost element: software and systems maintenance, self-assessment and compliance testing with respect to the operating guidelines set by the relying party service providers, and on-boarding and delisting of VO portals.

Several placement options remain and could be considered:

- at the institutions (universities and research centres): although the end-users and the ad-hoc VO portal operators will find this an attractive option, the user base within each institution is relatively small. There are not that many researchers engaged in collaborative research compared to, say, the number of students in a university (an issue that has been an impediment to the uptake of federated ID for research in general). Also, as soon as the use of a portal expands beyond a single institution, there is no longer a 'natural home', a sustainable funding model, and organisational issues become burdensome.

- at the generic e-Infrastructures (EGI, EUDAT, GÉANT, PRACE, …): larger e-Infrastructures with a mission to deliver or coordinate services have the expertise to run secured services in a trusted way. Being organisationally close to the relying party service providers (resource centres, HPC centres, data repository services) makes them a naturally trusted partner to operate credential stores. Yet – as generic infrastructure brokers – they are rather remote from the researchers that are setting up VO portals and run the scientific workflows. It may be more complex to form relationships with a large distributed body of portal operators who may not know the 'way to' the generic e-Infras. Where existing relationships exist, and also for generic VO portals that are operated by, on behalf of, or in close cooperation with the generic e-Infrastructures, they will be a natural focal point to establish a Master Portal.

  Where possible, Master Portals should also be run under a highly-available service level, since many VO portals (and research workflows) will depend on a single instance. Availability models could include both a service offered by a single organisation in a distributed way, or by distributing coordinated copies of the service across multiple organisations that collaborate.

  The operators of the VO portals will have a loose association with the generic e-Infrastructures, and each VO portal operator will be a relatively small entity. A likely result is a limited ability to fund the Master Portal by means of charging the VO portal operators.

- at the Research Infrastructures (ELIXIR, WLCG, …): researchers and research supporters are likely to feel affinity with the pan-European Research Infrastructure (RIs) in their domain, and may have elements of their ICT usage also integrated with or source from Research Infrastructures and (global) research communities. Experience and 'marketing' of VO portal technology is likely to spread through domain conferences and meetings, and it is likely that expertise in setting up VO portals will follow the same communication. This hints at the structured pan-European RIs being a place to offer a Master Portal: these RIs are 'close' (community-wise) to the VO portal operators and researchers, and as a result they are 'easily found' as a service by the community. The RIs are also well organised and have the ICT expertise and ability to either establish a properly secured production service themselves, or to procure one to their specification and desired availability level in the market – or from one of the generic e-Infrastructures. The expectations on availability and service level are also likely to be comparable within a single research domain, and RIs can jointly opt to procure (or operate) a single instance – which could be specifically branded to each RI and community.

  The VO portal operators will have a closer association with the RIs, and it may be within the remit of the RIs to support and sustain a service for VO portals that support their research domain. The VO portal operators remain small entities, and are unlikely to be a direct source of income for the Master Portal operators.

- In a hybrid model, where an e-Infrastructure or consortium runs a Master Portal in a multi-tenant mode on behalf of one or more Research Infrastructures or large communities. The technical implementation of a Master Portal cannot actually distinguish between being a 'multi-tenant' infrastructure provided under contract to a research infrastructure, or directly to specific VO portals. Logically it does make a significant difference in the sustainability model, since the on-boarding and user support for the specific VO portals in this model remains with the Research Infrastructures, and only the technical and security operations are managed by the generic e-Infrastructure(s). Funding in this model will likely be via the now-explicit RI-to-e-Infrastructure relationship.

An alternative cost recuperation model in any placement model is to charge a fee per transaction. This pay-per-use is common in commercial authentication systems, where the service provider 'outsources' identification of users to a third party, and then pays per authentication transaction[3]. Since the relying party service providers are not directly connected to the Master Portal (they authenticate the end-user mediated via the VO portal), the Master Portal cannot charge the service directly. It is unlikely that charging the user will be effective when there are alternative routes that may be more complex to use, but gratis. The remaining options include contracts with (consortia of) relying party service providers (e-Infrastructures), or VO portal consortia (Research Infrastructures and user communities). Either option is effectively equivalent to those entities outsourcing the Master Portal operation.

In summary, a mixed model of Master Portals at both the generic e-Infrastructures and at the Research Infrastructures could emerge, with some RIs running their own Master Portal instances compliant with the relying party service provider needs, while other outsource this function to a cluster consortium or a (consortium of) generic e-Infrastructures.

3.3 VO membership services linked to the Master Portal

Not detailed in the graphic are the VO membership services (VOMS) that are run by or on behalf of user communities and research collaborations. These membership database services are sensitive resources (in terms of exposure to security incidents) and should be run according to specific best practice guidelines[4].

Today these are run either by the major lead organisation in a few large user communities (WLCG), or run on behalf of a number of smaller communities by the generic e-Infrastructures (EGI VOMS, and the PRACE User Directory) and national infrastructures (NGIs).

It is not foreseen that this model will change for the VOMS and LDAP directory services. The sustainability model for in-line attribute authorities for federated authorization is out of scope of this study.

3.4 Delegation Service and On-line CA

The Delegation Service is a highly sensitive component that has both an initial set-up cost and a continued maintenance cost. The initial costs consist of the secure hardware setup (certified HSM hardware modules,

[3] This is a model used by e.g. ProtectNetwork for their IdP service. In countries where banks or mobile telephone operators offer general-purpose authentication services this is also on a pay-per-use basis.
[4] Such as the Attribute Authority Operations guidelines, see https://www.igtf.net/guidelines/aaops/

multiple physical systems, protected locked racks, physical security controls, data centre protection and environmental controls), policy development (the human effort expended in documenting the policy and how it is implemented, specific naming, identifier assignment requests, administrative and liability considerations, data protection policy and compliance, compliance with R&S and Sirtfi specifications), systems installation (software, configuration), and accreditation (self-audit, peer review process, IGTF participation, enrolment in a national federation and in eduGAIN). There are also ongoing costs including software and systems maintenance, annual audits, and maintaining community standing with the relevant trust groups.

The usability of the entire system depends highly on the availability of the Delegation Service for the widest possible user base. The preferred option is for the Delegation Service to support every possible IdP in any federation that is able to pass the minimum assurance baseline: non-reassigned identifiers, point-in-time traceability of identity, and willing to participate in resolving information security incidents.

Partial participation of IdPs and federations is also a consideration when there would be multiple Delegation Services, and the user – in the authentication WAYF selection – would accidentally pick a Delegation Service that is not connected to his or her home institute: by then the user will have 'no way back' and has to try again. This could be perceived by the user as a frustrating experience.

The Delegation Service is exposed to the end-user and the VO portal operators only because the authentication flow passes through it, and it (likely) presents a WAYF selection and data release consent dialog box. It is 'hidden' behind the Master Portal in the authentication flow, and has to establish relations with just those Master Portals and with the federation service (eduGAIN). Yet it has to be acknowledged that the trust relationship between the Delegation Server (and CA) and the Master Portals is a strong one: a credential compromise even on the Master Portals will have an impact (both in terms of effort expenditure as well as in reputational damages) on the Delegation Server. As such agreement-negotiation and a level of oversight and auditing (and the effort needed to maintain these) must be factored into the cost and deployment models.

There are several possible deployment models:

- at the institution of the end-user or VO portal operator: this is a very unlikely scenario, since the cost and administrative overhead of establishing a Delegation Service and On-line CA is out of proportion to the benefit (which would be 'independence'), and the calendar time needed to establish a Delegation Service (in the order of at least several months) is prohibitively long. There are also technical and operational reasons why relying party service providers would not be willing to support a large number of trust anchors (currently in the order of 100 trust anchors globally, with 1-2 changes per month).

- one per generic e-Infrastructure: the e-Infrastructures such as EGI are technically capable of operating a Delegation Service. Their close association with the relying party service providers will ease the process of acceptance of such on-line CA. Yet until a single common entity emerges, having each generic e-Infrastructure run an independent and separate Delegation Service will create non-interoperable silos: for reasons of trust the same user will be identified differently depending on the Delegation Service used, likely impairing cross-infrastructure workflows.

  Recuperating the set-up and maintenance cost of the Delegation Service for the e-Infrastructures could be based on service fees within each e-Infra, as long as the Master Portal operators feel bound to a single e-Infra. Alternatively, a market may emerge where different e-Infrastructures offer Delegation Services in a competitive way (based on price, service level, or in a package deal with other services).

There is potential for emergent 'vendor lock-in' since many of the e-Infrastructures have organisational ties with relying party service providers and may (unwittingly) prefer their native Delegation Service.

- one per Research Infrastructure (or user community): the RIs do have the critical size to operate production ICT services, and could develop the technical and policy capabilities to run Delegation Servers. However, there is a large number of RIs and communities (approx. 20-50 in Europe) and – like in the one-per-e-Infra case – each Delegation Service creates an 'identity silo' across which trust will not extend. This will not be an issue for the one-per-RI model as long as researchers stay within a single research domain for all their work and during their entire research career. Based on experience we consider such 'domain loyalty' to be unlikely.

  Recuperating the set-up and maintenance cost will be harder, since the user base is smaller than in the e-Infrastructure case. A more direct or imminent need for the service may however balance that consideration: the need to have such a service to enable specific research use cases may make the establishment of a (costly) service still a proper business decision. The costs would be sustained in a way similar to the Master Portal.

- a single Delegation Service for Europe: if the (user-wise) best option is to have a Delegation Service that accepts as many identities from all of eduGAIN, then – in addition to the existing but slightly differently targeted US CILogon instance – a single instance of the service in Europe could suffice. Running a (logically) single instance would address the 'WAYF' issue for end-users (everyone Master Portal would redirect through the single instance), and there is a common security incident control point that can deal with revocation of credentials, traceability, and the storage of personal data.

  Recuperating the cost could be easier, since the user base (number of Master Portals) is broader and most of the costs are per instance, not per-authentication. The total cost of the European Research Area is lower if there is a single instance. Yet it may suffer from clear (moral) ownership: the entity operating it should have a solid recuperation model (or be rather altruistic). It should also take care of branding and 'perceived image' in the community. For better or worse, many RIs, e-Infras, and projects perceive a need for branding and for identifying services as 'their own'. A single service branded by a perceived competitor will not be deemed acceptable. Therefore, any single European service should either be brand-less, brandable on-demand by the VO or Master Portal(s), or white-label.

- a single Delegation Service or CILogon for the world: one of the options considered was to use the existing CILogon service operated by NCSA for any user world-wide. Although attractive, the need to transfer attributes and personal data from the EU to the USA would very likely impact the willingness of institutions in Europe to release attributes and participate in the scheme. It would also – in its current form – not address all use cases for the VO portals. Instead, a very close collaboration has been established between AARC and the CILogon team through mutual support letters to ensure aligned development of solutions for global research. Conversely, it is unlikely that USA-based relying party services and science gateways would opt to use exclusively a European-based service (if only for latency reasons).

There is also an obvious risk for a single European or global Delegation Service: a single common service will have to balance the needs of all users, and thus will have to compromise in case of conflicting requirements. Requirements on minimal baseline assurance level and name uniqueness are the issues that could potentially lead to divergent services. Having a differentiated assurance classification, expressed between the Delegation

Service and the Master Portal (and potentially propagated to the VO portals), *together with the software capability to make decisions based on that assurance level statement*, will be necessary.

Given the usability, technical and cost considerations, it would be advisable to limit the number of Delegation Services (and the on-line CAs) as much as possible. The study proposes to agree on a single instance, acceptable to all RIs and e-Infrastructures. This leaves open the various sustainability models that could underpin such a single service.

The fee-per-transaction model – discussed for the Master Portal – would have similar issues for the Delegation Service. It is unlikely that the end-user or researcher will pay at point of use when gratis alternatives are available, and the other cost distribution models are logically equivalent to outsourcing by the e-Infrastructures or RIs.

The benefits on a single instance and the cost saving of sharing the service between all RIs and e-Infrastructures can be achieved in two ways:

- joint procurement of the service by the RIs and e-Infrastructures from a trusted third party. The competitive procurement of trust services has a precedent in the European R&E community in the form of the Server Certificate Service and Trusted Certificate Service since 2005. In a series of procurements different vendors have provided this service under a brand and with an interface specified by the community. However, with a specialised custom service (contrary to web site server certificates, the Delegation Service is a rather dedicated service that is not common in the industry) the risk of vendor lock-in is higher, and the potential list of competitors smaller.

- cooperative operation of the service by several e-Infrastructures: in a closely coordinated way, and with an independent policy management authority that is responsible for the service characteristics and policy compliance, the e-Infrastructures and their partners may establish a distributed redundant service. Each operating centre in itself may operate at a different service level, as long as the collective service meets the needs of the RIs, communities, and all VO portals. The expertise to establish and maintain the service is not a commodity, but is present in several organisations within the European R&E landscape.

  This model is reminiscent of the distributed set-up of CILogon in the USA, centrally coordinated but with its technical equipment distributed between a primary and alternate site. This model also protects against site and geographical failures, and can be implemented technically given the proper tools and hardware. The initial cost of such a distributed setup will be higher than a single instance, but may be off-set by lower staffing and operational costs.

The ultimate choice depends on the desired service level of the Research Infrastructures and user communities, and on the managerial and organisational choices made in the European community of researchers.

4 Summary

In this study we have surveyed the various sustainability models that could accompany a production deployment of the AARC CILogon-like TTS pilot system. This system and its close derivatives are being piloted in major RIs and e-Infrastructures, and can bridge the gap that currently exists between institutional user identities, R&E identity federations, and the non-web compute, cloud, and data services and brokering that use client certificates, such as those in EGI, PRACE, EUDAT, and cloud-based data transfer services of e.g. Globus Online.

A plausible deployment model would see large numbers of VO portals for domain researchers that are developed and sustained locally (inside an organisation or user group). There would be a limited number of Master Portals, the sensitive credential repositories and gateways that will manage the access keys of the users to the infrastructure. These will likely be run by RIs and e-Infrastructures, and there will be several of these in the system. Joint operation or delegation by RIs to e-Infrastructures has the potential for cost saving.

There should preferably be a single instance of a pan-European Delegation Service that acts a as secure and trusted gateway between the largest possible number of federations and institutions, and all of the e-Infrastructure and RI relying party services. This may be a jointly procured single service, or a distributed service offered collectively by a consortium of e-Infrastructures.

Pay-for-use by the researcher is unlikely to offer a sustainable income source as gratis (albeit more complex) alternatives for the end-users exist.

5 The RCauth.eu pilot service

As part of the AARC SA1 Pilot activity, a set of Master Portals is being deployed with Research Infrastructures (initially: ELIXIR) and e-Infrastructures (initially: EGI). Some of these are variants of the AARC pilot setup that have been adapted within the respective infrastructure (in particular the EGI instance has been modified by the EGI Engage project to suit specific EGI Competence Centre needs). To demonstrate end-to-end-feasibility of the pilot, a single Delegation Service and on-line IOTA CA "RCauth.eu" is being established by Nikhef for AARC SA1. The RCauth.eu CA will allow the demonstrators to integrate in an easier way with the distributed relying party service providers in the European and global e-Infrastructures and Cyberinfrastructures.

The RCauth.eu service is specifically scoped:

- as a 'white label' service so that it can be integrated by existing projects without creating a branding conflict. The name – inspired by Research and Collaboration – was chosen to be generic and on purpose does not include AARC – which is a transient project.

- it is operated by Nikhef and the Dutch National e-Infrastructure coordinated by SURF, with a high trust assurance and in compliance with all policies, and including the accreditation of the Delegation Service and on-line CA with the IGTF under the IOTA assurance profile. The service is established and accredited to demonstrate feasibility and to identify any weak points (also in the policy negotiation mechanism) that would be blocking issues in a production service.

- Nikhef and the Dutch National e-Infrastructure will offer this service to any user that can successfully authenticate to eduGAIN and meet the policy requirements, and will integrate with Master Portals from established Research Infrastructures and e-Infrastructures in Europe – in the expectation that these will be a relatively small number and cooperate with AARC.

- The service is run on a best-effort basis, for an indeterminate time period, and we plan to operate it as long as reasonably needed. We have the expectation that the service elements be taken up by qualified other parties, maybe based on the models in this study. Yet we cannot give guarantees that it will run uninterrupted, run for any given number of years, or is suitable for any specific community. Operations, maintenance and ongoing accreditation status are funded exclusively and voluntarily by Nikhef and the Dutch National e-Infrastructure coordinated by SURF.

The differentiated assurance issue raised in section 3.4 has been taken into account in setting up the RCauth.eu Pilot ICA G1 (Delegation Service), by aligning with the lowest assurance profile level that has been declared permissible by the major e-Infrastructures and Research Infrastructures that participate in the IGTF. This "Identifier-Only Trust Assurance" profile [5] (codenamed DOGWOOD [6]) only needs a non-reassigned identifier and cooperation of the IdP in security incident response. It has been tested with a large number of research use cases in the US by the CILogon Basic service, where it interconnects users of the InCommon

[5] https://www.igtf.net/ap/iota
[6] https://www.igtf.net/ap/loa

federation with a large number of XSEDE science gateways. More fine-grained differentiated assurance levels may be implemented in a future service, but it should be noted that substantial assurance is already available in many European countries through the GEANT Trusted Certificate Service – which bridges between federated identity plus eduGAIN and PKIX for a wide range of research, e-Infrastructure, and general purpose trust use cases.

Nikhef strongly encourages the e-Infrastructures, Research Infrastructures, and user consortia to find a persistent and sustainable solution within a reasonable amount of time. We note that the AARC project will end around mid 2017.

References

[CILogon] Jim Basney et al.; *CILogon, a project of the Cybersecurity Directorate at the National Center for Supercomputing Applications, University of Illinois with support from the US National Science Foundation and others*; http://www.cilogon.org/

[CILPPW] Mischa Sallé, Tamas Balogh, David Groep; *AARC CILogon Pre-Pilot Work*; https://wiki.nikhef.nl/grid/CILogon_Pre-Pilot_Work