07-05-2017

Deliverable DSA1.4: Pilots to improve access to R&E-relevant resources

Deliverable: DSA1.4

Contractual Date:	31-03-2017
Actual Date:	07-05-2017
Grant Agreement No.:	653965
Work Package:	SA1
Task Item:	Task 3
Lead Partner:	PSNC
Document Code:	DSA1.4
Authors:	M. Jankowski (PSNC), P. van Dijk (SURFnet), P. Birkinshaw (DAASI), M. Brzeźniak (PSNC), N. van Dijk (SURFnet), P. Gietz (DAASI), M. Hardt (KIT), D. Hübner (DAASI), N. Liampotis (GRNET), B. Oshrin (SURFnet), M. Prochazka (CESNET), M. Sallé (NIKHEF), P. Solagna (EGI), U. Stevanovic (KIT)

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

Abstract

This document provides an overview of the pilots realised within Service Activity 1 Pilots (SA1), Task 3 Pilots to improve access to R&E-relevant resources and (commercial) services (T3) during the term of the AARC project. Fourteen pilots took place, in three main areas: AAI mechanisms to access non-web resources, bridging e-infrastructures and access to (commercial) cloud services.

Table of Contents

Table of Tables

Executive Summary

This document provides an overview of the pilots realised within Service Activity 1 Pilots (SA1), Task 3 Pilots to improve access to research-and-education-relevant resources and (commercial) services (T3), during the term of the AARC project.

The work within this Task focused on three main areas, corresponding to the Task objectives defined in the Technical Annex:

- Authentication and authorisation infrastructure mechanisms to access non-web resources.
- Bridging e-infrastructures.
- Access to (commercial) cloud services.

Crossing all three areas was the objective to provide access to other services that are not yet accessible via the federated framework.

As a result of the Task's activities, 14 pilots were undertaken. Each of them addresses at least one of these objective-related areas. The selection of the components to be piloted was driven by the findings of Joint Research Activity 1 Architectures (JRA1). JRA1 identified requirements, existing gaps and issues, as well as promising tools and technologies to be used by e-infrastructures and communities.

The Task ran the pilots in close collaboration with the AARC communities, and obtained useful insights regarding the deployabilty and usability of the tested tools and technologies. Where needed, documentation was improved and, to facilitate adoption by potential communities, deployment scripts and dissemination materials were created. As a further significant benefit, since most partners in SA1 also participated in the drafting of the blueprint architecture (BPA), the piloting and mapping of components has helped to verify and improve the BPA.

Some of the tested solutions have already been promoted to production-level services; others are considered mature enough to be deployed in a production setting in the near future. In certain areas, particularly the integration of e-infrastructures, work will continue in AARC2.

1 Introduction

Service Activity 1 Pilots (SA1) Task 3 aims to improve access to relevant (non-web) resources across institutional, country and e-infrastructure/community boundaries. The main improvement is to leverage existing authentication and authorisation infrastructures (AAIs) that provide verified institutional user credentials and (external) authorisation attributes instead of using locally managed users. While many successful implementations exist already for web-based services, the technology for non-web scenarios (e.g. secure shell (SSH) or X.509 access) is still immature. As a result, the use of institutional credentials for research infrastructures and research services is still limited. The focus of this Task, therefore, has been on suitable approaches and services for token translation to bridge the web and non-web worlds.

The work was divided into three main areas, corresponding to the Task objectives defined in the Technical Annex [Technical_Annex]:

- AAI mechanisms to access non-web resources.
- Bridging e-infrastructures.
- Access to (commercial) cloud services.

Crossing all three areas was the objective to provide access to other services that are not yet accessible via the federated framework.

1.1 Access to Non-Web Resources

The following pilots focused on access to non-web resources:

- CILogon-like.
- ELIXIR TTS with CILogon.
- LDAP Facade.
- SSH key management using COmanage and OpenConext.
- SSH key provisioning for VMs using WaTTS.
- RCauth certificate delivery using WaTTS.

While the initial and main focus in this area was on translating from web to non-web technologies, the Task discovered use cases in which a reverse translation is required as well. With services such as CILogon/RCauth, users have been able to use the credentials from their home organisations to access services that require certificate-based authentication. In addition to the pilots focused on access to non-web services, the Task performed one to show that users holding Interoperable Global Trust Federation (IGTF) certificates were able to access federated services which are available via eduGAIN.

As a preparation to the pilot work on non-web access, the Task undertook a brief comparison of methods. This desk research included CILogon and LDAP Facade (both being tested in AARC) as well as Unity (tested in a slightly different context within the EUDAT-PRACE pilot) and Moonshot (not piloted in AARC so far). The comparison is available at the public AARC wiki [Non-Web-Comp].

1.2 Bridging e-Infrastructures

The Task undertook three cross-infrastructure integration pilots. The main aim of these pilots was to investigate and pilot approaches to making e-infrastructures accessible for users from other e-infrastructures, so that compute and data facilities become mutually exchangeable. Both web-based and non-web-based AAI scenarios were considered:

- EUDAT-PRACE integration (non-web-based scenario).
- EUDAT-EGI integration (web-based scenario).
- EUDAT-EGI integration (non-web-based scenario).

1.3 Access to Cloud Services

A third category of pilots aimed to demonstrate that third-party (often commercially hosted) services can be accessed leveraging AAIs based on Security Assertion Markup Language (SAML). These pilots included:

- Federated access to ORCID (this work was later extended to leverage ORCID as an attribute authority as well).
- Federated access to LibreOffice/Nextcloud.
- Federated access to Seafile.

1.4 In this Document

The following sections present summaries, in table form, of the pilots in each area.

Each section has a short introduction to the work area, outlining the objective, a general explanation of the challenges and pilot highlights.

Each table contains basic information (including focus, approach, components piloted, gain) and links to further resources, such as demos, screenshots, documentation, source code, etc. More in-depth information on these pilots is available on the AARC public wiki [AARCWikiPilotResults].

2 Pilot Summaries

The following tables summarise the main information about the pilots in each area, and provide links to more detailed documentation on the AARC project wiki.

2.1 Access to Non-Web Resources

2.1.1 Introduction

We have tested several solutions to improve access to non-web resources based on AAIs that provide verified institutional user credentials. While many successful implementations exist already for web portals, the technology to use institutional accounts for non-web scenarios (e.g. secure shell (SSH) or X.509 access) is still immature. In this paragraph, we present a number of promising pilot setups that enable users to use an institutional account to access non-web based resources. In addition, we performed one pilot that shows that users with an IGTF certificate can access SAML based web-resources (The IGTF to eduGAIN bridge).

2.1.1.1 *RCauth.eu*

One effort, resulting from the CILogon-like pilot (see Table 2.1 below), deserves particular mention. In order to allow easier integration of the pilots (and, later on, the production services) with the distributed relying party service providers in the European and global e-infrastructures, the Research and Collaboration Authentication (RCauth.eu) service was implemented and deployed [RCauth]. The service acts as an online certification authority (CA) and is based on CILogon software [CILogon], with a few adaptations to conform to European privacy regulations. It is currently being operated by FOM-NIKHEF and the Dutch national e-infrastructure coordinated by SURF. The service operates with a high trust assurance, in compliance with all policies, including International Global Trust Federation (IGTF), under the Identifier-Only Trust Assurance (IOTA) profile. The service is accessible for users who can successfully authenticate via eduGAIN and identity providers (IdPs) that meet the policy requirements. As such, RCauth.eu has been included as a central component in several other AARC pilots (e.g. ELIXIR TTS and WaTTS certificate delivery).

2.1.2 Summary Tables

Pilot	CILogon-like
Focus	Pilot a combination of solutions (workarounds) to enable access to non-web, X.509-based resources that are common in the GRID world with SAML-based credentials in an end-user-friendly way

Pilot	CILogon-like
Approach/AARC identified solution	CILogon is the central component in this pilot, which provides X.509 certificates using OpenID Connect and a SAML-based federated identity. End users interact with web portals (science gateways), which are represented in the pilot in the form of a demonstrator "VO Portal". Additional components, such as a separate "Master Portal", have been introduced to hide complexity for the users and science gateways and to provide sufficient scalability. At the same time, interfacing with the Virtual Organisation Membership Service (VOMS) as an attribute provider is being investigated and piloted.
Components piloted	• CILogon software (including OA4MP, Shibboleth, MyProxy, Simple CA) • VO Portal + Master Portal • VOMS
Gain for end users/ administrators	• Providing X.509-based access capabilities to the end user without the need for them to maintain or understand PKI • No need to help end users with difficult-to-manage certificate credentials • A single entry point can provide access to a wide range of resources, both web and non-web based
Demo/video	• A clear explanation of this effort is provided as a blog post on the AARC project website: Digital certificates behind the scenes • Two working demos have also been created, one showing the API and flow and one showing a PoC storage portal, leveraging the integration with VOMS
Detailed technical description	FOM-NIKHEF AARC pilot wiki RCauth.eu presentation
Documentation of components	FOM-NIKHEF AARC pilot wiki
Software source(s)	https://github.com/rcauth-eu
Lead	FOM-NIKHEF
Community partners	ELIXIR, EGI, CESNET, CSC
Status	After successful pilots with the ELIXIR and EGI community, a white-label CA pilot service has been created: RCauth.eu. In a next step, command line access to proxy credentials using SSH keys will be piloted.

Table 2.1: Pilot: CILogon-like

Pilot	ELIXIR TTS with CILogon
Focus	• ELIXIR AAI login to a cloud VM • ELIXIR AAI login to a data transfer portal

Pilot	ELIXIR TTS with CILogon
Approach/AARC identified solution	Use X.509 certificates and credential translation from ELIXIR AAI's web-based protocols (currently SAML, in the future also OpenID Connect (OIDC)) [ELIXIR].
Components piloted	• CA issuing the certificates (operated by FOM-NIKHEF) • Master Portal (operated by CESNET) • MyProxy (operated by CESNET)
Gain for end users/ administrators	• Users can use services requiring X.509 authentication without needing to hold one • Administrators can build portals that can operate on behalf of the user
Demo/video	Screenshots
Detailed technical description	• Master Portal, MyProxy and RCauth.eu description and installation are described at FOM-NIKHEF AARC pilot wiki • Client side is described at ELIXIR AAI GitHub
Documentation of components	RCauth.eu pilot wiki Perun Client
Software source(s)	Perun Client
Lead	CESNET
Community partners	ELIXIR, FOM-NIKHEF
Status	Pilot is working. For registered users it can be tested on https://wuotan.ics.muni.cz

Table 2.2: Pilot: ELIXIR TTS with CILogon

Pilot	LDAP Facade
Focus	The pilot aims to provide access to non-web resources (e.g. SFTP, SSH console) to non-grid users using existing AAIs, without the need to obtain user certificates
Approach/AARC identified solution	LDAP Facade is a single software component, which needs to be installed at the SP. It makes use of the local accounts prepared during the registration step. The software is able to replace the traditional LDAP server, as it provides the same interface. As a result, LDAP Facade can be used as a local user manager, as well as an authentication and authorisation component, without any modification to servers (not only is core code unchanged, even implementing a specialised plugin is not necessary). On the other side (collaboration with external IdP), the deployment is the same as for any other SAML-based SP.

Pilot	LDAP Facade
Components piloted	LDAP Facade
Gain for end users/ administrators	• No need to install additional software on user side • Limited use of user groups (depends on available attributes) • Federated access to non-web services • No need to manually maintain local accounts on each resource • No modification to server software • Up-to-date identity information (requires ECP or AQ SAML profiles support by the IdP. Workarounds have been proposed but are not yet implemented)
Demo/video	Demo (see flow description and instructions here)
Detailed technical description	AARC wiki
Documentation of components	LDAP Facade wiki at KIT
Software source(s)	Git at KIT
Lead	PSNC
Community partners	PSNC, KIT
Status	Finalised. The LDAP Facade portal is working. It is possible to log in there using eduGAIN credentials and register to a test service. The user is then able to log in to the resource using a local password. Limitation: the user's credentials are not checked against his/her IdP while logging in to the resource.

Table 2.3: Pilot: LDAP Facade

Pilot	SSH key management using COmanage and OpenConext
Focus	A common requirement for research organisations is to provide access via SSH to command line tools hosted on Unix-based systems. Provisioning and deprovisioning the accounts required for access, and leveraging the researcher's existing credentials to authenticate the access, are common challenges in providing this access.
Approach/AARC identified solution	The solution leverages COmanage [COmanage] to enrol a researcher to a collaboration, collect the researcher's SSH public key, and create a Unix account for the researcher
Components piloted	• COmanage Registry • SP proxy (OpenConext) • LDAP server

Pilot	SSH key management using COmanage and OpenConext
Gain for end users/ administrators	• No need to manually set up a shell account and manage SSH keys • Binding SSH key authentication to existing credentials
Demo/video	Video
Detailed technical description	• COmanage Registry, used to manage participant registration in the collaboration • SP proxy and OpenConext, used to manage connectivity to federated identity and other authentication services • LDAP server, provisioned by COmanage and used by the Unix server as a source of account information
Documentation of components	COmanage OpenConext OpenLDAP
Software source(s)	COmanage OpenConext OpenLDAP
Lead	SURFnet
Community partners	Tested with AARC partners
Status	Works as a reliable pilot. Requires some minor enhancements to be generally suitable for production use.

Table 2.4: Pilot: SSH key management using COmanage and OpenConext

Pilot	SSH key provisioning for VMs using WaTTS
Focus	General: provision of desired final credentials at selected services while authenticating at home institution. Specifically piloted: provisioning SSH keys across EGI virtual machines while authenticating via EGI SP-IdP proxy.
Approach/AARC identified solution	WaTTS Token Translation Service [WaTTS] is used to provide desired final credentials (e.g. SSH keys, X.509, S3 access tokens, etc.). WaTTS accepts many different implementations of OIDC providers. The solution can be used at any community/infrastructure/service where there is a need to "bridge" between different technologies, and can be run as a standalone "plug-and-play" solution. The only requirement is that the authentication through WaTTS is done with OIDC.
Components piloted	WaTTS
Gain for end users/ administrators	Provision SSH keys, e.g. in a remote VM
Demo/video	WaTTS service
Detailed technical description	AARC wiki
Documentation of components	WaTTS source WaTTS GitBook
Software source(s)	WaTTS source
Lead	KIT
Community partners	GRNET
Status	Works as a reliable pilot. The Task is seeking use cases to promote it to production.

Table 2.5: Pilot: SSH key provisioning for VMs using WaTTS

Pilot	RCauth certificate delivery using WaTTS
Focus	General: provision of desired final credentials at selected services while authenticating at home institution. Specifically piloted: provisioning X.509 certificates of the RCauth.eu CA.
Approach/AARC identified solution	WaTTS Token Translation Service [WaTTS] is used to provide desired final credentials (e.g. SSH keys, X.509, S3 access tokens, etc.). WaTTS accepts many different implementations of OIDC providers. The solution can be used at any community/infrastructure/service where there is a need to "bridge" between different technologies, and can be run as a standalone "plug-and-play" solution. The only requirement is that the authentication through WaTTS is done with OIDC.

Pilot	RCauth certificate delivery using WaTTS
Components piloted	WaTTS
Gain for end users/ administrators	Obtain certificate in VO Portal or command line
Demo/video	WaTTS service
Detailed technical description	AARC wiki
Documentation of components	WaTTS source WaTTS GitBook
Software source(s)	WaTTS source
Lead	KIT
Community partners	GRNET, FOM-NIKHEF
Status	Works as a reliable pilot. The Task is seeking use cases to promote it to production.

Table 2.6: Pilot: RCauth certificate delivery using WaTTS

Pilot	IGTF X.509 certificates to eduGAIN proxy
Focus	The main goal of this pilot is to allow end users holding valid certificates issued by an IGTF Classic or MICS CA to access services that are available via eduGAIN
Approach/AARC identified solution	The IGTF to eduGAIN bridge is a SAML identity provider that supports authentication via X.509v3 certificates. The information that is available in the user's certificate is extracted during the authentication process and is released to the service provider in the form of SAML attribute assertions.
Components piloted	SAML 2.0 identity provider (SimpleSAMLphp)
Gain for end users/ administrators	Enables end users to access services: • Using the X.509 certificates that they already have • Not forcing them to register/use home organisation accounts • Achieving a high level of assurance (LoA) Enables service operators/administrators to: • Follow the paradigm shift to federated access • Not lose/alienate their user base • Maintain LoA requirements

Pilot	IGTF X.509 certificates to eduGAIN proxy
Demo/video	• Staging instance of IGTF to eduGAIN proxy https://edugain-proxy-pilot.igtf.net/simplesaml • Production instance of IGTF to eduGAIN proxy https://edugain-proxy.igtf.net/simplesaml
Detailed technical description	The IGTF to eduGAIN bridge is a SAML identity provider based on SimpleSAMLphp that supports authentication via X.509v3 certificates using the authX509toSAML module. It is configured to accept only certificates issued by certification authorities accredited under the IGTF Classic or MICS profiles. The IGTF to eduGAIN bridge IdP does not store any user information. The information, which is released in the form of SAML assertions, is extracted from the user's certificate during the authentication process. Read more at AARC wiki.
Documentation of components	AARC wiki SimpleSAMLphp documentation authX509toSAML module documentation
Software source(s)	SimpleSAMLphp GitHub authX509toSAML module GitHub
Lead	GRNET, FOM-NIKHEF
Community partners	EGI
Status	The pilot has been promoted to production and is now available through eduGAIN as an identity provider supporting the REFEDS Research and Scholarship (R&S) entity category and Sirtfi.

Table 2.7: Pilot: IGTF X.509 certificates to eduGAIN proxy

2.2 Bridging e-Infrastructures

2.2.1 Introduction

The main objective for the cross-infrastructure pilots was to propose and verify solutions that could effectively bridge the usage of user credentials between e-infrastructures, as well as interconnect identity providers and service providers from so far separated infrastructures. The Task conducted three pilots in this category:

- EUDAT-PRACE.
- EUDAT-EGI for cross-infrastructure access to web-browser-based resources.
- EUDAT-EGI for cross-infrastructure access to non-web-browser-based resources.

The challenges included, on the one hand, the different technologies and approaches used by infrastructures, and on the other hand, the different policies and attributes that apply. As a result of the pilots, further harmonisation efforts have been initiated.

2.2.2 Summary Tables

Pilot	EUDAT-PRACE
Focus	Achieve interoperability between the two infrastructures regarding authentication and authorisation, in particular using X.509 certificates to access non-web resources on both sides. Examine how Unity IdM technology may be used for this kind of task [Unity].
Approach/AARC identified solution	Use Unity IdM HTTP API to synchronise user accounts and group information between three parties: PRACE LDAP, EUDAT B2ACCESS (based on Unity IdM) and EUDAT B2STAGE services. Map different credentials (including different types, e.g. X.509, SAML) of single user to single account across many resources.
Components piloted	• Unity IdM (B2ACCESS-like configuration) • B2STAGE • iRODS • LDAP server
Gain for end users/ administrators	• Automated account provisioning on EUDAT infrastructure for PRACE users respecting local security policies • Access to multiple resources by single user using different credentials
Demo/video	• Test B2ACCESS user console: https://b2access.eudat.psnc.pl:2443/home/home • Test B2STAGE/iRODS service, available for gridFTP connections at gsiftp://eptest.eudat.psnc.pl" • Screenshots and description of the steps: EUDAT-PRACE pilot on the AARC wiki
Detailed technical description	Both the PRACE nodes and the EUDAT B2STAGE service use X.509 certificates for authentication and authorisation. The authentication is provided by X.509 PKI out of the box if both infrastructures accept CA signing of the user's certificate (at the moment, any IGTF-approved CA). The proper authorisation requires synchronisation of accounts and groups between the infrastructures. User management in PRACE is based on LDAP, with manual account provisioning by the LDAP administrator. Accounts for PRACE users are then automatically provided in B2ACCESS (EUDAT user management system) by a script run periodically. B2ACCESS is able to accept different types of credentials and link multiple credentials to a single identity. B2ACCESS is also a token translation service that provides different types of final credentials required by EUDAT services. In particular, the piloted solution implemented a script that

Pilot	EUDAT-PRACE
	automatically and on user login provisions the user account and verifies user's group membership on B2STAGE service.
Documentation of components	Unity IDM documentation B2STAGE GridFTP GitHub
Software source(s)	Unity IDM GitHub B2STAGE GridFTP GitHub
Lead	PSNC
Community partners	EUDAT, PRACE
Status	The basic version (provisioning accounts) is working. Group handling is in progress. The software is being evaluated by EUDAT to be used in production.

Table 2.8: Pilot: EUDAT-PRACE

Pilot	EUDAT-EGI for cross-infrastructure access to web-browser-based resources
Focus	The main goal of this pilot is to allow end users to transparently access EGI [EGI] and EUDAT [EUDAT] web-browser-based resources
Approach/AARC identified solution	Use EGI CheckIn as an IdP to EUDAT B2ACCESS SP proxy and EUDAT B2ACCESS as an IdP to EGI CheckIn
Components piloted	• EUDAT B2ACCESS IdP/SP proxy service (Unity IdM) • EGI CheckIn IdP/SP proxy service (SimpleSAMLphp)
Gain for end users/ administrators	End users are able to access EGI and EUDAT web-browser-based resources using the same credentials
Demo/video	• Staging instance of B2ACCESS service https://unity.eudat-aai.fz-juelich.de:8443 • Staging instance of EGI CheckIn service https://aai-dev.egi.eu/proxy • Staging instance of EGI Applications Database https://appdb-dev.marie.hellasgrid.gr/
Detailed technical description	Following an exchange of SAML metadata, EGI CheckIn acts as an IdP to EUDAT B2ACCESS SP proxy while, at the same time, EUDAT B2ACCESS acts as an IdP to EGI CheckIn
Documentation of components	SimpleSAMLphp documentation Unity IdM documentation
Software source(s)	SimpleSAMLphp GitHub Unity IdM GitHub
Lead	EGI, GRNET

Pilot	EUDAT-EGI for cross-infrastructure access to web-browser-based resources
Community partners	EGI, EUDAT
Status	Completed

Table 2.9: Pilot: EUDAT-EGI for cross-infrastructure access to web-browser-based resources

Pilot	EUDAT-EGI for cross-infrastructure access to non-web-browser-based resources
Focus	The main goal of this pilot is to allow end users to transparently access EGI [EGI] and EUDAT [EUDAT] non-web-browser-based resources using a common online CA (RCAuth.eu)
Approach/AARC identified solution	Use a common online CA (RCauth.eu)
Components piloted	• EUDAT B2ACCESS IdP/SP proxy service (Unity IdM) • EGI CheckIn IdP/SP proxy service (SimpleSAMLphp) • EGI Master Portal • EGI VO Portal • RCauth.eu online CA / delegation service
Gain for end users/ administrators	End users are able to access EGI and EUDAT non-web-browser-based resources
Demo/video	• Staging instance of B2ACCESS service https://unity.eudat-aai.fz-juelich.de:8443 • Staging instance of EGI CheckIn service https://aai-dev.egi.eu/proxy • Staging instance of EGI Master Portal https://masterportal-pilot.aai.egi.eu/mp-oa2-server/ • Staging instance of EGI Demo VO Portal https://masterportal-pilot.aai.egi.eu/vo-portal • Staging instance of RCauth.eu online CA / delegation service https://ca-pilot.aai.egi.eu/
Detailed technical description	Following an exchange of SAML metadata, EGI CheckIn acts as an IdP to EUDAT B2ACCESS SP proxy while, at the same time, EUDAT B2ACCESS acts as an IdP to EGI CheckIn. EGI and EUDAT operate dedicated Master Portal instances, both of which are registered as OIDC clients to RCauth.eu online CA / delegation service. Moreover, EGI operates a VO Portal that is registered as an OIDC client to the EGI Master Portal.
Documentation of components	SimpleSAMLphp documentation Unity IdM documentation

Pilot	EUDAT-EGI for cross-infrastructure access to non-web-browser-based resources
Software source(s)	SimpleSAMLphp GitHub Unity IdM GitHub Master Portal GitHub VO Portal GitHub RCauth.eu Online CA / delegation server GitHub
Lead	EGI, GRNET, FOM-NIKHEF
Community partners	EGI, EUDAT
Status	Completed

Table 2.10: Pilot: EUDAT-EGI for cross-infrastructure access to non-web-browser-based resources

2.3 Access to Cloud Services

2.3.1 Introduction

AARC activities regarding collaboration with commercial providers were partly based on the results of the GÉANT GN3plus project Support to Clouds Service Activity. AARC evaluated the results and content of the cloud providers survey and the cloud services catalogue developed by GÉANT. This collaboration between AARC and GÉANT worked in both directions. In particular, the GN4 project organised and implemented the Infrastructure as a Service (IaaS) tender, in which the community made the explicit mandatory requirement for SAML v2 support in commercial IaaS and related services. AARC actively contributed to raising awareness related to federated AAI in cloud systems and applications, both on the user-community side as on the industry side.

Another step towards solving federated AAI adoption issues at service providers is to set up a trusted proxy service, run on the community side, taking care of authentication, authorisation and privacy. Commercial cloud services then have to deal with only one single point of contact, while at the same time there is more freedom in terms of available integration technologies (e.g. OIDC in addition to SAML). This results in a lower threshold for integrating with research and education AAIs.

As a result of these efforts, industry partners are improving support for AAI in their products. For instance, in response to explicit user requirements expressed by several community parties, Seafile, a data sharing and syncing solution (see Table 2.11 below) developed and included support for SAML-based authentication to both web and non-web clients (including desktop and mobile clients). Similar developments happened in other areas, e.g. LibreOffice/Nextcloud and ORCID. Importantly, industry itself started to consider support for SSO as a selling point for their product.

The Task conducted three pilots to show the feasibility of connecting to federated AAIs and to prove that a proxy approach lowers technical thresholds for service providers:

- Federated access to ORCID (Table 2.12), which provides persistent identifiers for researchers. This pilot was later extended (in Task 2) to leverage ORCID as an attribute authority as well (see Table 2.13).
- Federated access to LibreOffice/Nextcloud, an attempt to implement a service for online, collaborative versions of documents (similar to Google Docs) with federated access, that can be run in a private cloud.
- Enable federated access to Seafile, a data sharing and syncing solution.

2.3.2 Summary Tables

Pilot	Seafile with SAML federation
Focus	Seafile [Seafile] is a cloud storage system with file encryption, group sharing, synchronisation, etc. available both as a free and open source Community edition and as a priced Pro edition. Seafile supports Shibboleth to enable SSO. The Pro edition supports many storage backends, including commercial object stores (through its S3 plugin). This pilot aimed to test federated access to the Seafile service using a community WAYF service [WAYF] as a proxy to multiple SAML IdPs.
Approach/AARC identified solution	Seafile software is designed to work with a single IdP. Therefore the missing element to enable federated access to a Seafile service is the discovery service. The approach was to configure the existing WAYF service as SAML IdP for Seafile Shibboleth authentication.
Components piloted	• Seafile server • Seafile clients: web based and Java • Shibboleth SP • SWITCH WAYF
Gain for end users/ administrators	Seafile Pro supports a number of storage backends including Amazon S3 [S3], OpenStack Swift [Swift] and Ceph [Ceph]. These storage interfaces are used by many commercial cloud providers (in particular S3, which is the most popular and recognised). With Seafile it is possible to build a service with federated access on one side and data stored by commercial providers on the other side. In this way, the service may provide a wide range of research communities with access to scalable commercial resources, while at the same time preserving data privacy, thanks to user-side encryption. The Java-based Seafile desktop client supports non-web access to the services and can perform local vs. remote storage synchronisation. The Collabora Online (LibreOffice) cloud suite integration supports collaborative work on online documents via Seafile (feature available in Seafile Pro edition) in a way that is similar to Google Docs.
Demo/video	Running service: https://box.pionier.net.pl/

Pilot	Seafile with SAML federation
Detailed technical description	A customised version of the Seafile service was configured to use the PIONIER.id WAYF service as an IdP. The solution follows WAYF Shibboleth 2.x flow. See AARC wiki page for the pilot details.
Documentation of components	Seafile manual on Shibboleth authentication Shibboleth SP WAYF
Software source(s)	Seafile GitHub
Lead	PSNC
Community partners	PIONIER.id federation
Status	This pilot demonstrated an approach to using Seafile as a suitable solution in a federated identity context. Following this effort, the Seafile service (Community edition) is now officially available for PIONIER.id users. PSNC recently purchased a Pro licence and the Pro edition is scheduled to be deployed in Q2 2017 in the production instance of the Seafile service.

Table 2.11: Pilot: Seafile with SAML federation

Pilot	Federated access to ORCID
Focus	Work with ORCID to make the ORCID service, which provides persistent identifiers for research and education workflows, available through eduGAIN
Approach/AARC identified solution	ORCID is acting as a SAML service provider
Components piloted	• ORCID • SimpleSAML ORCID • SimpleSAMLphp • Shibboleth
Gain for end users/ administrators	Federated access to ORCID greatly simplifies getting access to ORCID IDs
Demo/video	ORCID production service: https://orcid.org/signin
Detailed technical description	The core of the setup is an ORCID account linking services, which contains an ORCID SP and an ORCID AA component. The ORCID SP allows end users to link their SAML-based home institution account to their ORCID. This is done by logging in twice: once at the home institution, and once at ORCID. The combination of the ePPN attribute value and the ORCID is then stored in a database. An ORCID AA component allows SAML-based SPs to query the attribute authority.

Pilot	Federated access to ORCID
Documentation of components	ORCID ORCID API SimpleSAMLORCID module Shibboleth AA setup (example) SimpleSAMLPHP attribute aggregator
Software source(s)	See above
Lead	SURFnet
Community partners	Initially Dutch and Italian Research community later tested among AARC partners
Status	Completed. In production use.

Table 2.12: Pilot: Federated access to ORCID

Pilot	ORCID as an attribute provider
Focus	Work with ORCID identifiers provided by the ORCID service, available to services in a SAML-based profile
Approach/AARC identified solution	ORCID is acting as an attribute authority. A proxy is used to aggregate ORCID identifiers on behalf of services.
Components piloted	• ORCID • OpenConext • SimpleSAMLphp • Shibboleth
Gain for end users/ administrators	Providing ORCID identifiers as part of the SAML-based authentication greatly simplifies working with ORCID identifiers for SAML-based services
Demo/video	Screenshots: https://wiki.surfnet.nl/display/ORCIDAA/Technical+Setup

Pilot	ORCID as an attribute provider
Detailed technical description	The core of the setup is an ORCID account linking services, which contains an ORCID SP and an ORCID AA component. The ORCID SP allows end users to link their SAML-based home institution account to their ORCID. This is done by logging in twice: once at the home institution, and once at ORCID. The combination of the ePPN attribute value and the ORCID is then stored in a database. An ORCID AA component allows SAML-based SPs to query the attribute authority. In this scenario, as SURFnet is operating a hub-and-spoke federation, the Task also investigated how the hub could be used to provide the ORCID attribute to SPs connected to the hub as part of the regular authentication flow. For this, an attribute query client was used in the hub (OpenConext) to query the ORCID AA component just before passing on an authentication to a SP. More details are available on SURFnet wiki.
Documentation of components	ORCID ORCID API SimpleSAMLORCID module Shibboleth AA setup (example) SimpleSAMLPHP attribute aggregator
Software source(s)	See above
Lead	SURFnet
Community partners	Tested among AARC partners
Status	Completed.

Table 2.13: Pilot: ORCID as an attribute provider

Pilot	LibreOffice/Nextcloud
Focus	The original intention of this pilot was to integrate ownCloud [ownCloud] as a backend storage service for the Collabora Online office suite (by providing a SAML/OAuth authentication proxy) as an example of bridged web and non-web services. However, the product developers recently started to cooperate to provide built-in integration (with Collabora as a *backend* to ownCloud) using Microsoft's WOPI protocol's own version of OAuth. These developments made the original pilot redundant. The pilot continued as a demonstration of integrating WOPI applications with SAML-based services, and as a SAML-aware, open source alternative to Google Docs. A second instance of the pilot explored attribute aggregation by complementing proxied SAML identities with attributes from an LDAP directory.

Pilot	LibreOffice/Nextcloud
Approach/AARC identified solution	Nextcloud [Nextcloud] (a fork of ownCloud) was used for the file management component instead of ownCloud, as it has a built-in and free SAML authentication option. The free but limited “CODE” edition of Collabora Online [CollaboraOnline] (bundled as a Docker image) was used for the office suite component. Three different demonstration services were built to explore various configurations.
Components piloted	• Nextcloud • Nextcloud SSO plugin • CODE Collabora Online Docker image • Shibboleth SP • OpenLDAP
Gain for end users/ administrators	Collaborative work on documents that often contain sensitive information is common practice in a lot of research communities. It is therefore useful for research infrastructures to have functionality similar to Google Docs provided by more trusted service providers, such as university computing centres, or research infrastructures like DARIAH.
Demo/video	Pilot demonstration information
Detailed technical description	Information on the implementation
Documentation of components	Nextcloud Nextcloud SAML CODE
Software source(s)	CODE Nextcloud
Lead	DAASI
Community partners	DARIAH, Collabora
Status	The CODE edition of Collabora is limited to 10 concurrent users, so existing pilots are not suitable for production.

Table 2.14: Pilot: LibreOffice/Nextcloud

3 Conclusions

The practical result of Task 3 is a total number of 14 pilots, described in this document. Each pilot addresses at least one of the Task's objectives, as defined in the Technical Annex. It is worth highlighting that some pilots have been promoted to production-level services, namely: the IGTF to eduGAIN proxy (Table 2.7), Federated access to ORCID (for eduGAIN users) (Table 2.12) and Seafile with SAML federation (for PIONIER.id users) (Table 2.11). In addition, with RCauth.eu, a white-label certification authority service has been established for research. Many piloted solutions are considered highly relevant for the community and, in terms of maturity, are close to production-level quality. Some pilots, such as the cross-infrastructure pilots (Section 2.2), are currently undergoing further evaluation. For other piloted solutions, such as those using WaTTS (Table 2.5 and Table 2.6), the Task is exploring whether they can be applied in a production setting, and to this end is seeking real-life use cases. Of the tested technologies in this Task, LDAP Facade seems to be too restrictive for particular use cases as it requires, for example, IdPs with SAML ECP profiles, which currently are not very common in Europe.

The Task observed that there is still a gap regarding technology to access SSH resources on physical machines. The solutions investigated either lack maturity, or require either usage of X.509 certificates, which does not seem appropriate in many use cases, or additional components that are not available yet (e.g. ECP support on the IdP side – see [DJRA1.1] for more details). Moonshot technology appears to offer the desired functionality. The piloting and evaluation of Moonshot [Moonshot] is planned for the AARC2 project.

The integration of e-infrastructures needs further discussion, testing and maturing. The Task has been able to take some first integration steps but further adjustment and harmonisation is needed to be able to move to production-level integrations. Further elaboration of cross-infrastructure AAI-related topics will take place in AARC2.

Overall, the pilot work undertaken by Task 3 is regarded as very useful, as it provided the opportunity to test the integration of various solutions side by side, communities were able to test-drive these solutions, relevant guidelines were defined, and possibilities for the further development of AAI solutions were identified. As a further significant benefit, observations and feedback gathered from this pilot work provided useful input for drafting and refining the blueprint architecture (see [MJRA1.4], [DJRA1.2], [Blueprint-Guidelines]).

References

[AARCWikiPilotResults] https://wiki.geant.org/display/AARC/Pilot+results+and+demos

[Blueprint-Guidelines] *AARC Blueprint Architecture and Implementation Guidelines* [in progress – link not yet available]

[Ceph] http://ceph.com/

[CILogon-Pilot-Blog] AARC blog post: "Digital certificates behind the scenes: the AARC CILogon pilot https://aarc-project.eu/digital-certificates-behind-the-scenes-the-aarc-cilogon-pilot/

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

[CollaboraOnline] https://www.collaboraoffice.com/collabora-online/

[COmanage] https://spaces.internet2.edu/display/COmanage/Home

[DJRA1.1] *Deliverable DJRA1.1: Analysis of user community and service provider requirements* https://aarc-project.eu/wp-content/uploads/2015/10/AARC-DJRA1.1.pdf

[DJRA1.2] *Deliverable DJRA1.2: Blueprint architectures* [in progress – link not available]

[EGI] https://www.egi.eu/

[ELIXIR] https://www.elixir-europe.org/services/compute/aai

[EUDAT] https://www.eudat.eu/

[iRODS] https://irods.org/

[LDAP-Facade-Wiki] http://wiki.data.kit.edu/index.php/LDAP-Facade

[MJRA1.4] *Milestone MJRA1.4: First Draft of the Blueprint Architecture* https://aarc-project.eu/wp-content/uploads/2016/08/MJRA1.4-First-Draft-of-the-Blueprint-Architecture.pdf

[Moonshot] https://www.jisc.ac.uk/rd/projects/moonshot

[Nextcloud] https://nextcloud.com/

[NIKHEF-AARC-Wiki] https://wiki.nikhef.nl/grid/AARC_Pilot

[Non-Web-Comp] https://wiki.geant.org/x/WgCIAw

[OpenConext] https://openconext.org/

[OpenLDAP] http://www.openldap.org/

[ORCID] https://orcid.org/

[ownCloud] https://owncloud.org/

[Perun] https://perun.cesnet.cz

[PRACE] http://www.prace-ri.eu/

[RCauth] https://www.rcauth.eu/

[S3] https://aws.amazon.com/s3/

[Seafile] https://www.seafile.com/

[Shibboleth] https://shibboleth.net/

[SimpleSAMLphp] https://simplesamlphp.org/

[Swift] https://wiki.openstack.org/wiki/Swift

[Technical_Annex]	https://aarc-project.eu/wp-content/uploads/2015/04/technical_annexB_chap1_3_v1_0-FINAL.pdf
[Unity]	http://www.unity-idm.eu
[WaTTS]	https://www.gitbook.com/book/indigo-dc/token-translation-service/details
[WAYF]	https://www.switch.ch/aai/support/tools/wayf/

Glossary

AA	Attribute Authority
AAI	Authentication and Authorisation Infrastructure
AARC	Authentication and Authorisation for Research and Collaboration
API	Application Programming Interface
AQ	Attribute Query
BPA	Blueprint architecture
CA	Certification Authority
Ceph	A unified, scalable object, block, and file storage system
CI	Cyber Infrastructure
CILogon	CILogon enables users to authenticate with their home organisation and obtain a certificate for secure access to cyber infrastructure
ECP	Enhanced Client or Proxy
eduGAIN	International interfederation service interconnecting research and education identity federations
EGI	European Grid Infrastructure
EPC	End Point Criterion
ePPN	eduPersonPrincipalName Shibboleth attribute
Git	A free and open source distributed version control system
HTTP	Hypertext Transfer Protocol
IaaS	Infrastructure as a Service
IdM	Identity Management
IdP	Identity Provider in the context of SSO scenarios, such as supported by Shibboleth
IGTF	Interoperable Global Trust Federation – a body to establish common policies and guidelines that help establish interoperable, global trust relations between providers of e-infrastructures and cyber infrastructures, identity providers, and other qualified relying parties
IOTA	Identifier-Only Trust Assurance
iRODS	Integrated Rule-Oriented Data System
JRA1	Joint Research Activity 1 Architectures
KIT	Karlsruhe Institute of Technology
LDAP	Lightweight Directory Access Protocol
LoA	Level of Assurance – degree of certainty that the user has presented a credential that refers to that user's identity
MICS	Member-Integrated Credential Service
NREN	National Research and Education Network
OA4MP	OAuth for MyProxy provides an OAuth-compliant REST web interface to the MyProxy service for providing user certificates to science gateways
OAuth	OAuth is an open standard for authorisation
OIDC	OpenID Connect
ORCID	A not-for-profit organisation that provides a unique identifier for individuals to use with their name as they engage in research, scholarship, and innovation activities across disciplines, borders and time

Perun	A wide system providing user management and user-connected services to various types of facilities in various infrastructure sizes
PKI	Public Key Infrastructure
PoC	Proof of Concept
PRACE	Partnership for Advanced Computing in Europe
R&E	Research and Education
R&S	Research and Scholarship
RCauth.eu	The white-label Research and Collaboration Authentication CA Service for Europe
REFEDS	Research and Education Federations
REST	Representational State Transfer
S3	Amazon Simple Storage Service
SAML	Security Assertion Markup Language is an XML-based, open-standard data format for exchanging authentication and authorisation data between parties, in particular, between an identity provider and a service provider
SA1	Service Activity 1 Pilots
Seafile	Open source file sync and share software
SFTP	SSH File Transfer Protocol
Sirtfi	Security Incident Response Trust Framework for Federated Identity
SP	Service Provider in the context of SSO scenarios, such as supported by Shibboleth
SSH	Secure Shell
SSO	Single Sign-On
Switft	OpenStack Object Storage
T3	SA1 Task 3 Pilots to improve access to research and education relevant resources and (commercial) services
TTS	Token Translation Service. RCauth.eu is a Token Translation Service that translates SAML to X509
VM	Virtual Machine
VO	Virtual Organisation – a dynamic set of individuals or institutions defined around a set of resource-sharing rules. Resource sharing is, necessarily, highly controlled, with resource providers and consumers defining clearly and carefully exactly what is shared, who is allowed to share, and the conditions under which sharing occurs
VOMS	Virtual Organisation Membership Service
WOPI	Web Application Open Platform Interface
WaTTS	A plugin-based Token Translation Service developed by KIT in the context of the INDIGO Data Cloud project
WAYF	Where Are You From
X.509	Standard for a public key infrastructure to manage digital certificates
XML	Extensible Markup Language