




Science and
Technology
Facilities Council

Authentication and Authorisation for the WLCG

An introduction to the key concepts, and
how they fit within the WLCG Space



Who am I?

I currently wear multiple Authentication and Authorisation hats...

- Service Manager for the Identity and Access Management service for the U.K. IRIS Collaboration
- Current Authorisation Working Group Chair for the WLCG
- Scrum master for the SKA agile team working on Authentication and Authorisation Infrastructure (AAI)

Working to ensure that all these communities (and others!) can interoperate

Questions? Feel free to contact at:

thomas.dack@stfc.ac.uk



Agenda

1 Introduction to AuthN & AuthZ

...and why does it matter anyway

2 AuthN & AuthZ For the WLCG

3 AuthN & AuthZ: Certificates & VOMS

4 AuthN & AuthZ: Tokens

5 Towards Tokens for WLCG



Take Aways

- What is the difference between Authentication and Authorisation
- An understanding of why this is important for research
- A basic understanding of the existing Certificate and VOMs infrastructure
- A more in-depth understanding of the OAuth and OIDC protocols
- A brief overview of the planned token transition for WLCG



Science and
Technology
Facilities Council

Introduction to Authentication & Authorisation

...and why does it matter anyway

Authentication and Authorisation

Or, Not letting everyone in

- Letting everyone access everything is often *a bad idea*
- Though this is not always true – the level of access control required involves considering the risks
- In the contest of research, we probably don't want anyone and everyone online being able to get in...



Authentication and Authorisation

Controlling access depends on:

- Verifying a user's access to an account or identity – **authentication (AuthN)**
- Knowing that a user is allowed to do what they want to do – **authorisation (AuthZ)**

These processes are usually combined, as key processes in a community's
Authentication and Authorisation Infrastructure - AAI

SIDE NOTE:

- **AuthoriZation** - 🇺🇸 & 🇨🇦
- **AuthoriSation** - 🇬🇧 & the rest of the English-speaking world

I will have typed **AuthoriSe**, except in occurrences where it is a **named term or attribute**

... also I use the Americanised shorthand, **AuthZ**

Authentication and Authorisation - Then

Traditionally a user would need to Authenticate separately to every new service or account they wanted to use

- Lots of accounts, scattered through every site a user has ever signed up for
- This results in *many* username and password pairs, which in turn leads to bad security practices
 - Credentials are reused for “simplicity”, or simply forgotten and lost

Authentication and Authorisation - Now

In more modern applications, the process is evolving

- Increased importance of account credibility
 - Processes and systems in place to show that an associated identity is verified
- Increased adoption of single sign-on mechanisms, using a unified identity
 - Log-in with Social IDs, such as Facebook, Google, ORCID
- Can grant authorised access based on an identity's attributes
 - Home institute, email address, etc

What makes AAI important for research?

A couple of key concepts underpin the use of identity in research:

- Confidentiality
- Traceability
- Suspension
- Attribution

What makes this important for research?

A couple of key concepts underpin this:

- Confidentiality
- Traceability
- Suspension
- Attribution



Confidentiality

- Whilst final research outputs are public, maintaining confidentiality before this point is important
- Particularly important for fields handling personal and medical data

What makes this important for research?

A couple of key concepts underpin this:

- Confidentiality
- Traceability
- Suspension
- Attribution




Traceability

- If something goes wrong – accidentally or maliciously – service owners need to be able to trace where this happened
- Knowing which account caused an issue is important for both support or suspension

What makes this important for research?

A couple of key concepts underpin this:

- Confidentiality
- Traceability
- Suspension
- Attribution




Suspension

- Processes to suspend an individual user account in case of compromise or malicious activity
- Avoids downtime caused by stopping an entire service or resource, by isolating the problematic identity

What makes this important for research?

A couple of key concepts underpin this:

- Confidentiality
- Traceability
- Suspension
- Attribution



Tumblr – [The Internet by Nedroid](#) / January 30th, 2013

Attribution

- A single central identity provides a mechanism through which research can be attributed
- ORCID is an example of this – life-long identifiers for researchers, which can be attached to publications, grant requests, etc
- Identity changes can present problems – name changes from marriage, gender transition, etc



Science and
Technology
Facilities Council

AuthN & AuthZ For the WLCG



AuthN & AuthZ for the WLCG

Providing global access to computing resources – not easy!

- Global user community, with many members
- Distributed single infrastructure
- Not guaranteed that users know each other
- Not guaranteed that users will ever meet

Need a system for provisioning access, to be trusted across the grid

- ***Not guaranteed that users know each other***

Who can you trust to know the users, in order to **authenticate** them?

Many options, including:

- The Infrastructure
- The Experimental Group or Research Community
- The Home Organisations
- A trusted Third Party

In most cases, a user's **Home Organisation** may have the most current information – especially if their access is a function of their affiliation.

- ***Not guaranteed that users will ever meet***

But! A user's **Experimental Group or Research Community** may be better placed to tell you...

- Which group the user belongs to
- What roles (permissions) they should have
 - *what are they **authorised** to do*
 - *Are they a user, an admin, a super-user, etc?*
- The status of a user's policy acceptance – e.g. whether they have accepted the latest acceptable usage policy

AuthN & AuthZ for the WLCG

Need to bring these concepts together in order to control access to the grid

We'll look at an overview of how this is done for the WLCG, looking at both the current system and the intended migration infrastructure:

Current:	Future:
Certificates & VOMS	Tokens & IAM



Science and
Technology
Facilities Council

AuthN & AuthZ: Certificates & VOMS

X.509 – a Recap

A Certificate is...

- A digital identity, representing an entity
 - could be a service/website, a machine, or a human individual
- Signed by a Certificate Authority (CA)
 - Self-signed certificates do exist but are not useful for authentication purposes!
 - Signed by taking a hash of the certificate, which is then encrypted with the CA's private key
- Long lived – typically a grid certificate will last a year
- User keeps an accompanying private key and password

VOMS - Virtual Organization Membership Service

VOMS is the certificate-based AAI for the WLCG

- A central attribute authority, and central repository for VO user information
- Enables sorting of users into group hierarchies, storing roles and other attributes
- This is used to issue trusted attribute certificates which can then be used in the Grid environment for authorisation purposes.

VOMS - Virtual Organization Membership Service

The typical VOMS user flow is as follows

- The user registers at the CERN User Office and with their experiment secretariat, and the user is entered into the CERN HR database.
- The user gets a certificate from any eligible Certificate Authority
- The user registers with VOMS Admin for their VO and presents their certificate to authenticate. Their email address is matched with their record in the CERN HR database to confirm that their identity has been verified. The VO manager approves the registration. From that time on, the HR ID is used as the unique ID in the VOMS DB, allowing the e-mail address to be changed independently at either end.
- VOMS attributes, such as groups and roles, are maintained by VO Managers.
- The user can submit jobs to the grid or access data on the grid by saving their User Certificate locally, running voms-proxy-init to generate the VOMS proxy, and running the relevant grid command; their proxy will get delegated to grid services as needed.

Proxy Certificates?

A Proxy Certificate is generated and signed using a user's certificate

- As it is signed by the user, the proxy certificate acts as a proof of identity – can be used for **authentication**

A VOMS Proxy adds extra, VO specific, information to the proxy – such as roles

- This adds **authorisation** information to the proxy

Proxy certificates are short lived, and can be used to in turn sign further proxies – therefore allowing for delegated access on behalf of the user

voms admin for VO: test.vo Current user: CN=test0

[Home](#) [Browse VO](#) [Configuration Info](#) [Request membership](#) [Certificate Info](#) [Other VOs on this server](#)

Your certificate information

Subject	/C=IT/O=IGI/CN=test0
Issuer	/C=IT/O=IGI/CN=Test CA
Serial number	9
Not valid after	Sep 24, 2022 17:39:34 (in 6 years, 341 days)

Your certificate:

- is **NOT** linked to any membership in this VO. This means you are **NOT** recognized as a VO member, and cannot get VOMS credentials using `voms-proxy-init` for this VO out of this certificate.

[Click here to register as a new member](#)



Science and
Technology
Facilities Council

AuthN & AuthZ: Tokens



...Tokens?

There are a few components in a token flow. Key to this, a token is...

- a JSON Web Token
 - “JWTs are an open, industry standard RFC 7519 method for representing claims securely between two parties.” <https://jwt.io/introduction/>
 - Tokens are encoded strings of data, issued by an issuer with which the client (receiver) has a trusted relationship.
- In this context, JWTs are used to communicate authentication and authorization information using the **OAuth 2.0** and **OIDC** protocols

OAuth 2.0	OIDC
Open Authorization	Open ID Connect


OAuth 2.0 – the Access Token




- An open standard for access delegation, and most often used for allowing users to grant a website access to their information on another website, without giving them their password
 - Examples include signing into a third-party website using your Google, GitHub, or Orcid account
- An access token is provided to the third party, from whichever service is acting as the Authorization Server, which it can then use to retrieve a protected resource

Oauth 2.0 – Terminology

- **Protected Resource**
 - The identity or data which is to be shared
 - *eg: your email address, or the ability to post to your Twitter*
- **Resource Owner**
 - The user who owns the **Protected Resource**
 - *eg: you!*
- **Client**
 - The application that wants access on behalf of the **Resource Owner**
 - *eg: the website you want to register with using your Google identity*
- **Authorisation Server**
 - The application which knows the **Resource Owner**, and where the **Resource Owner** already has an account
 - *eg: The Google, Twitters, Orcids of the world*
- **Resource Server**
 - Where the **Protected Resource** lives, and what the **Client** wants to use
 - *eg: the API from which the **Client** can access the **Protected Resource***



Oauth 2.0 – Terminology



- **Protected Resource**

- The identity or data which is to be shared
- *eg: your email address*

- **Resource Owner**

- The user who owns the resource
- *eg: you!*

- **Client**

- The application that requests access to the resource
- *eg: the website you are using*

- **Authorisation Service**

- The application which has an account with the Resource Server
- *eg: The Google, Twitter, Facebook*

- **Resource Server**

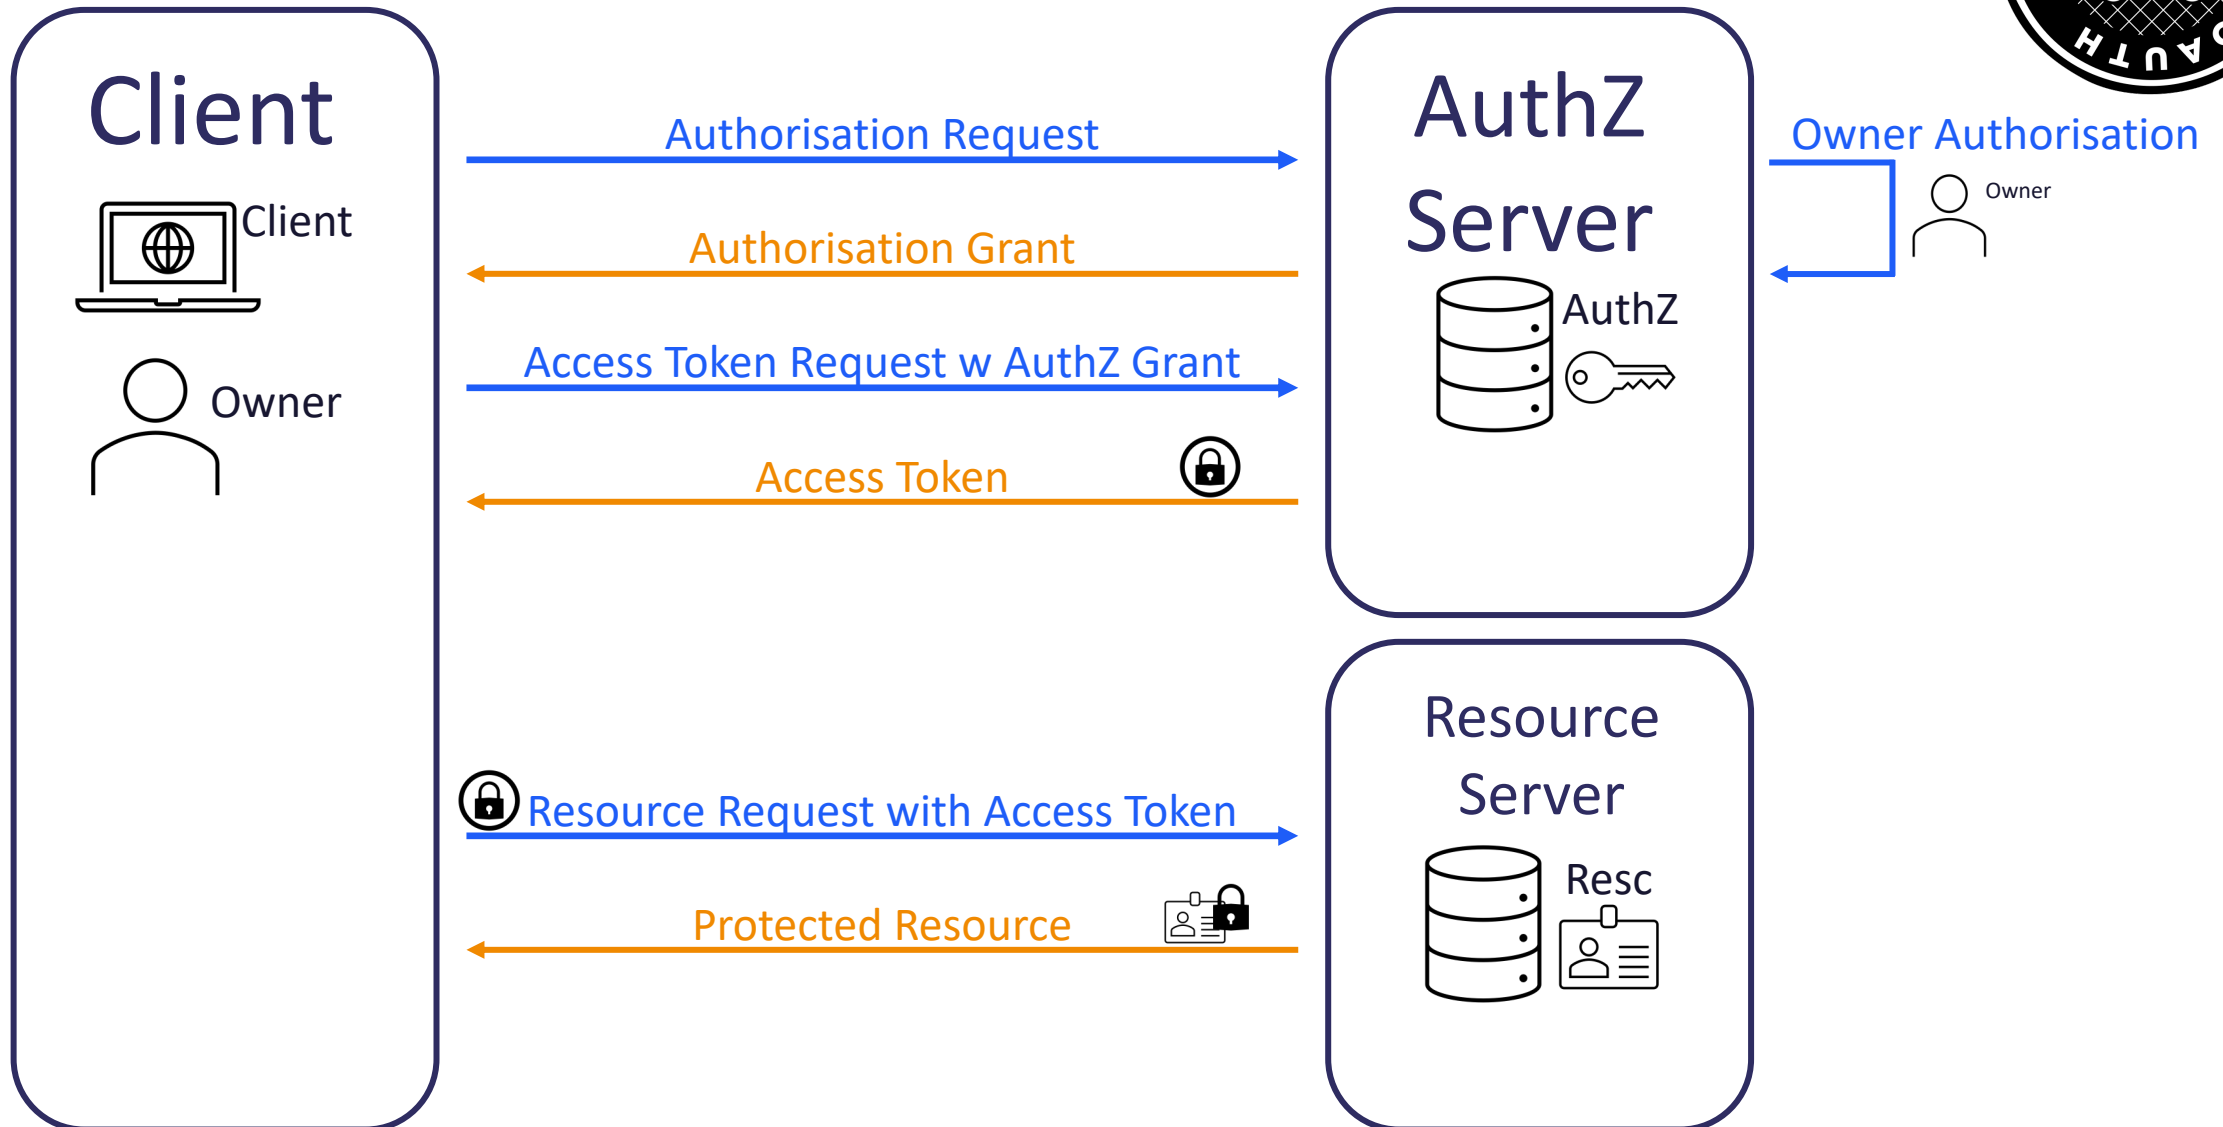
- Where the Protected Resource is stored
- *eg: the API from which the resource is accessed*

A quick aside...

OAuth 2.0 defines two types of clients:
Confidential and Public Clients

- Confidential applications can hold credentials with which they use to authenticate themselves to the AuthZ Server in a secure way. They require a trusted backend server to store the secret(s).
 - *eg – a web application with a secure backend*
- Public clients **cannot** hold credentials securely.
 - *eg – a native desktop or mobile application, or a JavaScript-based client-side web application (single-page app)*

OAuth 2.0 – How things work



OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants
 - Authorisation Code
 - PKCE
 - Client Credentials
 - Device Code
 - Refresh Token

Key Authorisation Grant Concepts

- Redirect URL – a URL at the client which the AuthZ server will deliver an issued token to
- ClientID – the “username” of the client
- ClientSecret – the “password” of the client

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants
 - Authorisation Code
 - PKCE
 - Client Credentials
 - Device Code
 - Refresh Token



Authorisation Code

- Used by both confidential and public clients
- An authorisation code is exchanged for an access token
- When the user returns to the client via the Redirect URL, the Authorisation Code is extracted from the URL and used to obtain the Access Token

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants

```
https://example.com/oauth/auth?  
response_type=code  
&scope=EXAMPLE_REQUESTED_SCOPES  
&client_id=EXAMPLE_CLIENTID  
&redirect_uri=EXAMPLE_REDIRECT_URI
```

- Authorisation Code Grant
- PKCE
- Client Credentials Grant
- Device Code Grant
- Refresh Token Grant

Public clients

- All authorisation code is exchanged for an access token
- When the user returns to the client via the Redirect URL, the Authorisation Code is extracted from the URL and used to obtain the Access Token

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants

```
• curl --header "Authorization: Basic EXAMPLE_SECRET"  
• --data "grant_type=authorization_code&code=EXAMPLE_CODE"  
• --request POST https://example.com/oauth/token
```

- Device Code
- Refresh Token

- Used by both confidential and public clients
- An authorisation code is exchanged for an access token
- When the user returns to the client via the Redirect URL, the Authorisation Code is extracted from the URL and used to obtain the Access Token

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants
 - Authorisation Code
 - PKCE →
 - Client Credentials
 - Device Code
 - Refresh Token

PKCE: Proof Key for Code Exchange

- An extension to Authorisation Code, which aims to prevent Cross Site Request Forgery (CSRF) and authorisation code injection attacks
- Not a replacement for a ClientSecret, and therefore does not enable treating a Public Client as Confidential

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants
 - Authorisation Code
 - PKCE
 - Client Credentials
 - Device Code
 - Refresh Token



Client Credentials

- Must **ONLY** be used by Confidential Clients
- The Client authenticates itself directly with the Authorization Server, for example using its ClientID and ClientSecret pair or with a public/private key pair
- As Client Authentication is used as the Authorisation grant, no further AuthZ is required
- Often used for a client to obtain information about itself

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through

```
curl --request POST \  
  --url 'https://EXAMPLE/oauth/token' \  
  --header 'content-type: application/x-www-form-urlencoded' \  
  --data grant_type=client_credentials \  
  --data client_id=EXAMPLE_CLIENT_ID \  
  --data client_secret=EXAMPLE_CLIENT_SECRET \  
  --data audience=EXAMPLE_AUDIENCE
```

ClientSecret pair or with an public/private key pair


- As Client Authentication is used as the Authorisation grant, no further AuthZ is required
- Often used for a client to obtain information about itself

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants

- Authorisation Code
- PKCE
- Client Credentials
- Device Code
- Refresh Token



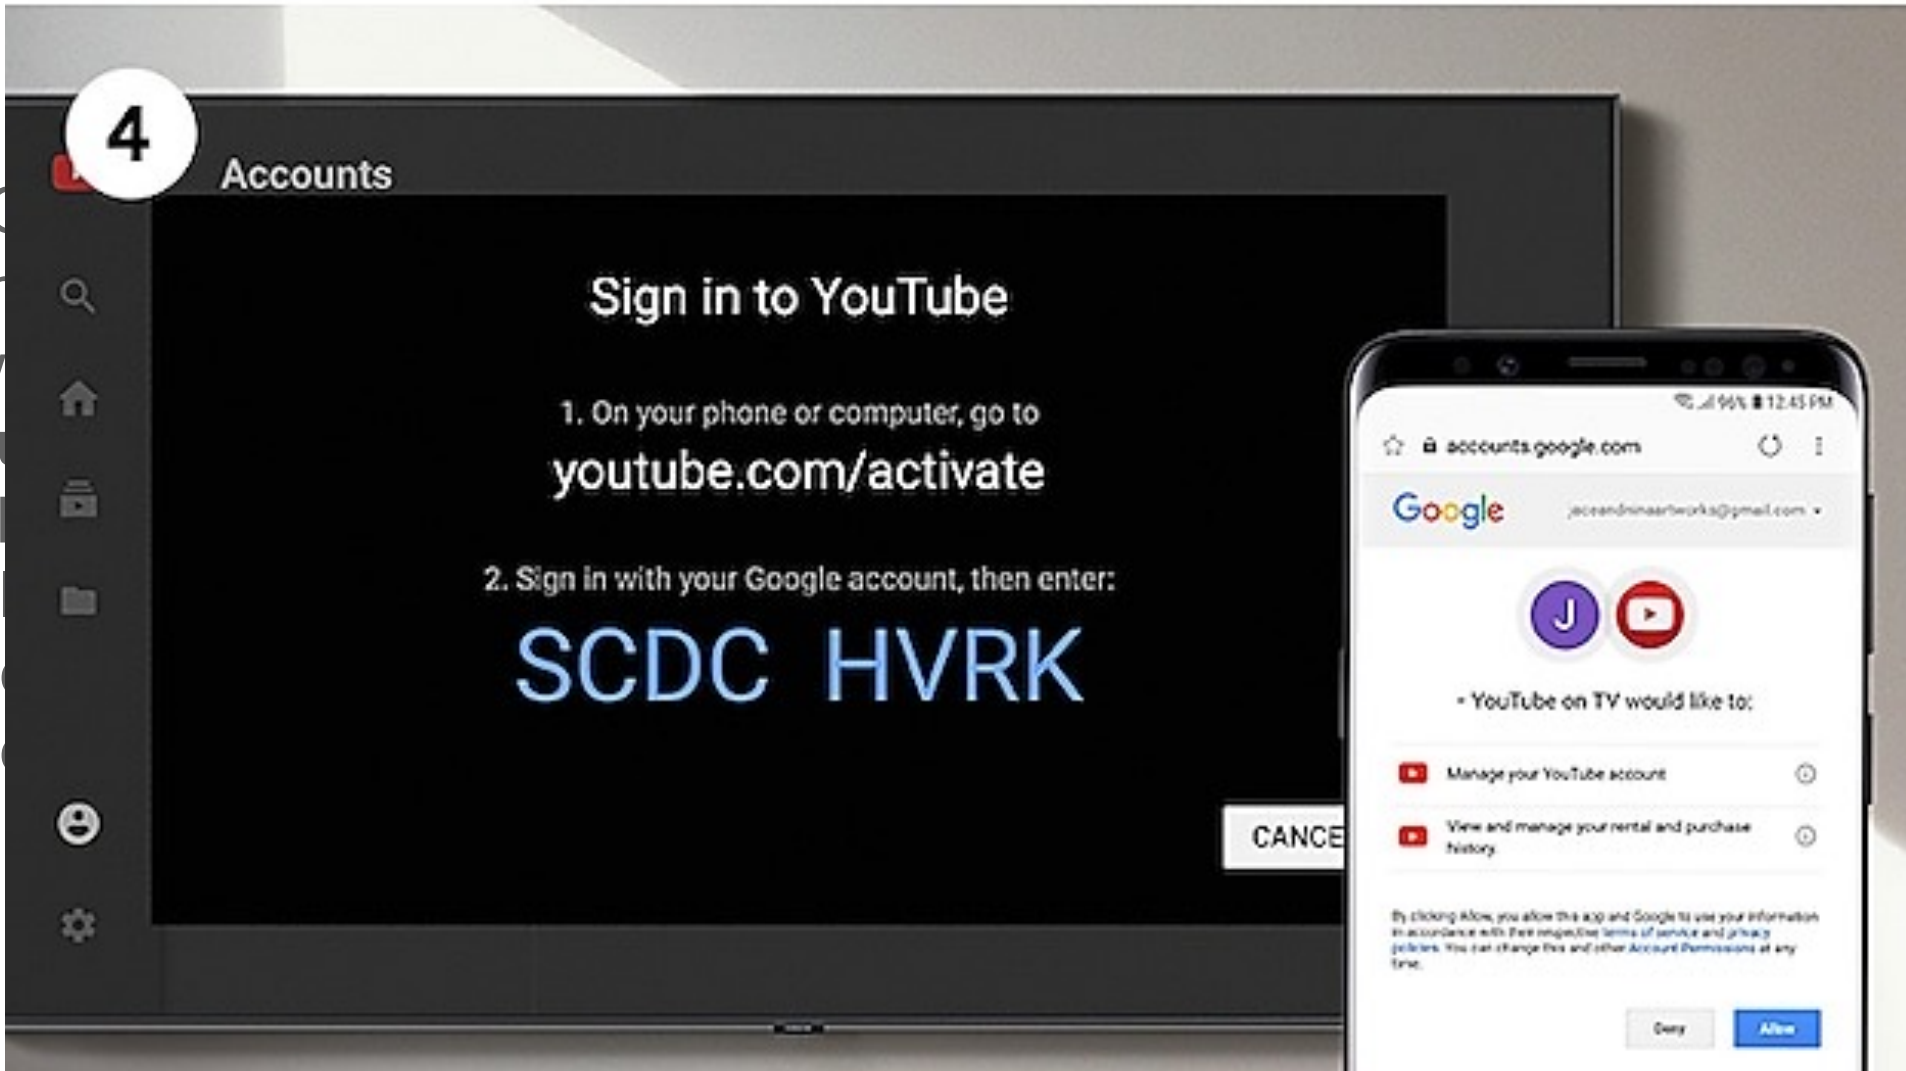
Device Code

- This flow is intended for use by browserless or input constrained clients – such as command line
- The user is directed to visit a URL at the AuthZ Server in a separate browser, along with a user code to identify the device
- The original device continuously polls the AuthZ server with the generated device code until the user completes the interaction, the code expires, or another error occurs

OAuth 2.0 – Authorisation Grants



- The O...
- which
- know
- A
- P
- C
- D
- R



out
ver in a
ify the
erver with
es the
interaction, the code expires, or another error occurs

OAuth 2.0 – Authorisation Grants



- The OAuth 2.0 framework specifies several different methods through which a **Client** can verify itself to the **Authorisation Server** - these are known as Authorisation Grants
 - Authorisation Code
 - PKCE
 - Client Credentials
 - Device Code
 - Refresh Token

Refresh Token

- This is a mechanism for allowing a client to get a new Access Token after an initial one has expired, without further user interaction
- Not a replacement for a ClientSecret, and therefore does not enable treating a Public Client as Confidential

OAuth 2.0 – Authorisation Grants



To get a Refresh token, the client needs to request the scope:

```
scope=offline_access&
```

```
curl --request POST \  
  --url 'https://{yourDomain}/oauth/token' \  
  --header 'authorization: Basic {yourApplicationCredentials}' \  
  --header 'content-type: application/x-www-form-urlencoded' \  
  --data grant_type=refresh_token \  
  --data 'client_id={yourClientId}' \  
  --data 'refresh_token={yourRefreshToken}'
```

therefore does not enable treating a Public Client as Confidential

OAuth 2.0 – the Access Token



- OAuth is an **Authorisation** flow – does not have a method for user **Authentication**
- Enter...

OpenID Connect: OIDC

OIDC: The ID Token



- **OpenID Connect (OIDC)** is an identity layer on top of the base OAuth 2.0 protocol
- OIDC provides the ability for app and web developers to authenticate users without the need to directly store credential sets
- OIDC enables the Authorisation Server to act as an **Identity Provider (IdP)**. An OAuth 2.0 Authorisation Server implementing OIDC is also referred to as an **OpenID Provider (OPs)**.
- A client which uses an OP for authentication is a **Relying Party (RPs)**.
- OIDC identifies a set of personal attributes that can be exchanged between Identity Providers and the apps that use them, and includes an approval step so that users can consent (or deny) the sharing of this information

OIDC: The ID Token



- In OAuth flows involving a user, the user will authenticate with the Authorisation Server before providing consent for the server to release information to the client. OpenID Connect utilises this process to authenticate to the client
- The Token Endpoint will provide two separate tokens: a standard OAuth Access Token, and the OIDC ID Token.
- The ID Token will contain information about the user pertaining to what the client has requested
- When the client receives the ID Token, it may then read off requested claims

OIDC: Requesting Info with Scopes



- When the client makes its token request, it must use **Scopes** to specify which privileges are being requested in the token
- In OAuth 2.0, Scopes correspond to what resources are available when a protected resource is accessed
- In OIDC, Scopes correspond to the specific sets of information to be made available as Claim Values
- OIDC defines a set of standard Scopes, but does allow additional Scope values to be defined and used

```
https://example.com/oauth/auth?  
response_type=code  
&scope=openid%20profile%20email  
&client_id=EXAMPLE_CLIENTID  
&redirect_uri=EXAMPLE_REDIRECT_URI
```

OIDC: Default Scopes

as defined by OIDC Core



Claim	Summary
openid	REQUIRED. Informs the Authorization Server that the Client is making an OpenID Connect request. If the <code>openid</code> scope value is not present, the behaviour is entirely unspecified.
profile	OPTIONAL. This scope value requests access to the End-User's default profile Claims, which are: <code>name</code> , <code>family_name</code> , <code>given_name</code> , <code>middle_name</code> , <code>nickname</code> , <code>preferred_username</code> , <code>profile</code> , <code>picture</code> , <code>website</code> , <code>gender</code> , <code>birthdate</code> , <code>zoneinfo</code> , <code>locale</code> , and <code>updated_at</code> .
email	OPTIONAL. This scope value requests access to the <code>email</code> and <code>email_verified</code> Claims.
address	OPTIONAL. This scope value requests access to the <code>address</code> Claim
phone	OPTIONAL. This scope value requests access to the <code>phone_number</code> and <code>phone_number_verified</code> Claims.
offline_access	OPTIONAL. This scope value requests that an OAuth 2.0 Refresh Token be issued that can be used to obtain an Access Token that grants access to the End-User's UserInfo Endpoint even when the End-User is not present (not logged in).

OIDC: Default Token Required Claims

as defined by OIDC Core



```
{ "iss":  
  "https://server.example.com",  
  "sub": "24400320",  
  "aud": "s6BhdRkqt3",  
  "exp": 1311281970,  
  "iat": 1311280970,  
  "auth_time": 1311280969"  
}
```

Claim	Summary
iss	REQUIRED: Issuer identifier for who issued the response. This is a case sensitive URL
sub	REQUIRED: Subject identifier. A locally unique identifier which must never be reassigned, this identifies the user within the issuer. A case sensitive string \leq 255 ASCII characters in length
aud	REQUIRED: Audiences that this token is intended for, using the Client ID as its identifying value.
exp	REQUIRED: Expiration time for the token, after which it MUST NOT be used for processing.
iat	REQUIRED: Time at which the JWT was issued
auth_time	Time when end-user authentication occurred. If a max_age request is made or auth_time is requested as an essential claim, this is REQUIRED – otherwise it is OPTIONAL

OIDC: Standard Identity Claims

as defined by OIDC Core



Claim	Summary
<code>name</code>	The end-user's full name, including all name parts. <code>given_name</code> , <code>family_name</code> , <code>middle_name</code> and <code>nickname</code> may also be used.
<code>preferred_username</code>	A shorthand name by which the user wishes to be referred to by at the RP. This MAY be any valid JSON string – including special characters such as <code>@</code> , <code>/</code> , or whitespace. The RP MUST NOT rely on this value being unique.
<code>email</code>	End-User's preferred e-mail address. Its value MUST conform to the RFC 5322 <code>addr-spec</code> syntax. The RP MUST NOT rely upon this value being unique.
<code>phone_number</code>	End-user's preferred telephone number.
<code>address</code>	End-User's preferred postal address.
<code>profile</code>	URL of the End-User's profile page. The contents of this Web page SHOULD be about the End-User.
... And others	<code>picture</code> , <code>website</code> , <code>email_verified</code> , <code>gender</code> , <code>birthdate</code> , <code>zoneinfo</code> , <code>locale</code> , <code>phone_number_verified</code> , <code>updated_at</code>

OIDC: Additional Claims



- The OpenID Connect Core specification only defines the previous small set of claims as standard
- However, OP's MAY provide additional claims about the End-User
- A typical example includes `groups`, a list of groups the user is registered with at the OP
- Any additional Claim will need to have a corresponding Scope, to allow it to be requested
 - Alternatively, extra claims could be included within the `profile` scope – this is how the OP to be used within the exercises provides its `groups` claim

OIDC: Putting this together...



The token is three Base64-URL strings separated by dots that can be easily passed in HTML and HTTP environments

Encoded <small>PASTE A TOKEN HERE</small>	Decoded <small>EDIT THE PAYLOAD AND SECRET</small>						
<pre>eyJraWQiOiJyc2ExIiwiaWF0IjoiUjMyNTYifQ.eyJzdWIiOiI0MGY4NzNhOS1hY2E4LTQxYmYtYmFkNS05OGQ1MWU3NjUyNDiLCjpc3MiOiJodHRwczp1wvaXJpcy1pYW0uc3RmYy5hYy51ayIsIm5hbWUiOiJUZXR0IFVzZXIiLCJncm91cHMlOlsiaUNTQ1wvdGVzdGluZyIsIm1DU0MiLCJyYWwtdGl1c2E1c2VybmFtZSI6IiRlc3Rvc2VyIiwib3JnYW5pc2F0aW9uX25hbWUiOiJlJlUkLTIEIjBTSIsImV4cCI6MTY3ODANNTU5MCwiaWF0IjoxNjc0MDkxOTkwLCJqdGkiOiJjYTc0ODVmNi0yYTA5LTQ1NjU0YThjOC04YTA1ZmI3NDlhYWMiLCJjbGllbnRfaWQiOiIyMWE2YzI2NC1mNGVmLTRlZDEtOGQwM0xNzQ4YzhhNDM2NDIiLCJlbWVpbCI6InQta1kQGHvdG1haWwud28udWcifQ.E6e5yuF9x7BTxM9bbGGx-jwk4VFy8wnkKesWMDWk1QjwI-oLHYrfCWfoZdJ7BU2ox3MxCMBYPBZa-vwpl4bD8vTYp9mz15z263JyuLDBhFVdxVNr1siZxQaQDcEPmnxzw703qDzmdFLAG9bZ79Bkn1MWiGdZ3erzSacEdEPBSEvSw6A18RqRksAfK8DV7_6d9kNcrr19NSn01NzX1cv6yYtxwli79iQuFv63UvRtJ_4F9-27jvIW_dwNgdw7ivhqEmThcMazcSPhn68EudLM3ytGTldlppOWKeD1cMHexeYFNcfvdyVgQSm1EZVXN7-FXUqZmK-zjEmgU7oEHWDQ</pre>	<table border="1"><thead><tr><th>HEADER: ALGORITHM & TOKEN TYPE</th></tr></thead><tbody><tr><td><pre>{ "kid": "rsa1", "alg": "RS256"}</pre></td></tr><tr><th>PAYLOAD: DATA</th></tr><tr><td><pre>{ "sub": "40f873a9-aca8-41bf-bad5-98d51e765241", "iss": "https://iris-iam.stfc.ac.uk", "name": "Test User", "groups": ["iCSC/testing", "iCSC", "ral-tier1", "localAccounts"], "preferred_username": "TestUser", "organisation_name": "IRIS IAM", "exp": 1678095590, "iat": 1678091990, "jti": "ca7485f6-2a09-4565-a8c8-8a05fb749aac", "client_id": "21a6c264-f4ef-4ed1-8d00-1748c8a43642", "email": "t-j-d@hotmail.co.uk"}</pre></td></tr><tr><th>VERIFY SIGNATURE</th></tr><tr><td><pre>RSASHA256(base64UrlEncode(header) + "." + base64UrlEncode(payload),</pre></td></tr></tbody></table>	HEADER: ALGORITHM & TOKEN TYPE	<pre>{ "kid": "rsa1", "alg": "RS256"}</pre>	PAYLOAD: DATA	<pre>{ "sub": "40f873a9-aca8-41bf-bad5-98d51e765241", "iss": "https://iris-iam.stfc.ac.uk", "name": "Test User", "groups": ["iCSC/testing", "iCSC", "ral-tier1", "localAccounts"], "preferred_username": "TestUser", "organisation_name": "IRIS IAM", "exp": 1678095590, "iat": 1678091990, "jti": "ca7485f6-2a09-4565-a8c8-8a05fb749aac", "client_id": "21a6c264-f4ef-4ed1-8d00-1748c8a43642", "email": "t-j-d@hotmail.co.uk"}</pre>	VERIFY SIGNATURE	<pre>RSASHA256(base64UrlEncode(header) + "." + base64UrlEncode(payload),</pre>
HEADER: ALGORITHM & TOKEN TYPE							
<pre>{ "kid": "rsa1", "alg": "RS256"}</pre>							
PAYLOAD: DATA							
<pre>{ "sub": "40f873a9-aca8-41bf-bad5-98d51e765241", "iss": "https://iris-iam.stfc.ac.uk", "name": "Test User", "groups": ["iCSC/testing", "iCSC", "ral-tier1", "localAccounts"], "preferred_username": "TestUser", "organisation_name": "IRIS IAM", "exp": 1678095590, "iat": 1678091990, "jti": "ca7485f6-2a09-4565-a8c8-8a05fb749aac", "client_id": "21a6c264-f4ef-4ed1-8d00-1748c8a43642", "email": "t-j-d@hotmail.co.uk"}</pre>							
VERIFY SIGNATURE							
<pre>RSASHA256(base64UrlEncode(header) + "." + base64UrlEncode(payload),</pre>							

Handy token decoding taken from: <https://jwt.io/>

Check it out for more info on JWTs or when using them yourself!



OIDC: Putting this together...



Encoded PASTE A TOKEN HERE

```
eyJraWQiOiJyc2ExIiwiaWF0IjoiUjMyNTYifQ.
```

Decoded EDIT THE PAYLOAD AND SECRET

HEADER: ALGORITHM & TOKEN TYPE

```
{  
  "kid": "rsa1",  
  "alg": "RS256"  
}
```

```
base64UrlEncode(header) + "." +  
base64UrlEncode(payload),
```

eyJraWQiOiJyc2ExIiwiaWF0IjoiUjMyNTYifQ.

```
{  
  "kid": "rsa1",  
  "alg": "RS256"  
}
```



Science and
Technology
Facilities

OIDC: Putting this together...



```
.E6e5yuF9x7BTxM9bbGGx-
```

```
jwk4VFy8wnkKesWMDWk1QjwI-  
oLHYrfCWfoZdJ7BU2ox3MxCMBByPBZa-  
vwp14bD8vTYp9mz15z263JyuLLDBhFVdxVNr1si  
ZxQaQDcEPmnxzwY703qDzmdFLAG9bZ79Bkn1MWi  
GdZ3erzSacEdEPBSEvSw6A18RqRksAfK8DV7_6d  
9kNcrr19NSn01NzX1cv6yYtxwli79iQuFv63UvR  
tJ_4F9-27jvIW_dwNgdw7ivhqEmThcMazcSPhn6  
8EudLM3ytGTldlpp0WKeD1cMHexeYFNcfvdyVgQ  
Sm1EZVXN7-FXUqZmK-zjEmgU7oEHWDQ
```

JWK signing key published at a URL at the OP,
such as:

<https://wlcg.cloud.cnaf.infn.it/jwk>

The specific URL is detailed in the OP
metadata, found at:

```
.well-known/openid-configuration
```

For Example:

<https://wlcg.cloud.cnaf.infn.it/.well-known/openid-configuration>

```
keys:  
  0:  
    kty: "RSA"  
    e: "AQAB"  
    kid: "rsa1"  
    n: "hN-YNx7NE0lK3v5f5xFMr-M1dIkVt8KaAMVTKKl08C_LctAWCzDGXvbfjS0LJ5kSyyyazv6_5YBq6Yeiyhdxr0-mQip0SBFEhp9bVY2mYQ0BauEi-  
z4czyyV9Ey4UppsJPJIXYUIGa8_mIr1Fw_4FLFDw2VrUJuoJbB8a-  
pMGhxcKSM8nDLjmw4WJ1ienJsurIttIsaGB0n1Kshydj8EnV6CBqWuzYdCsqxjIDCvXxzYzBtpTVHwoGsDXJR3K8MbMdm2xVY_Lts7Tg7A6InbitAr95APnvgr3G-  
WE1M9rw8R6j0lEPaUIvYEEqnFcmJub67w7lQ0I1exiBUgAPu61w"
```

OIDC: UserInfo Endpoint



- The ID token is not the only way for a RP to get information about a user – it can also make a request to the Ops `UserInfo` endpoint
- A UserInfo Request is made by the client using either a HTTP `GET` or HTTP `POST`. The Access Token previously obtained from an OIDC Authentication request **MUST** be included as a Bearer Token
- The OIDC Core recommendation is that the request uses the HTTP `GET` method, with the Access Token sent using the `Authorization` header field

A non-normative example
of a UserInfo Request:

```
GET /userinfo HTTP/1.1
Host: server.example.com
Authorization: Bearer SlAV32hkKG
```




Science and
Technology
Facilities Council

Towards Tokens for WLCG



Transition to Tokens - Motivations

- OAuth and OIDC protocols have been adopted by a wide range of software and systems, thanks to their prevalence in industry – in particular within the social identity space
 - This prevalence enables developers to utilise developed libraries, and facilitates integration and interoperability
- Certificates, whilst established within WLCG, are not familiar outside this context - and are often viewed as unintuitive. Conversely, Token-based flows are becoming increasingly commonplace, and can lead to a better user experience as a result

Token Authentication and Authorisation Infrastructure

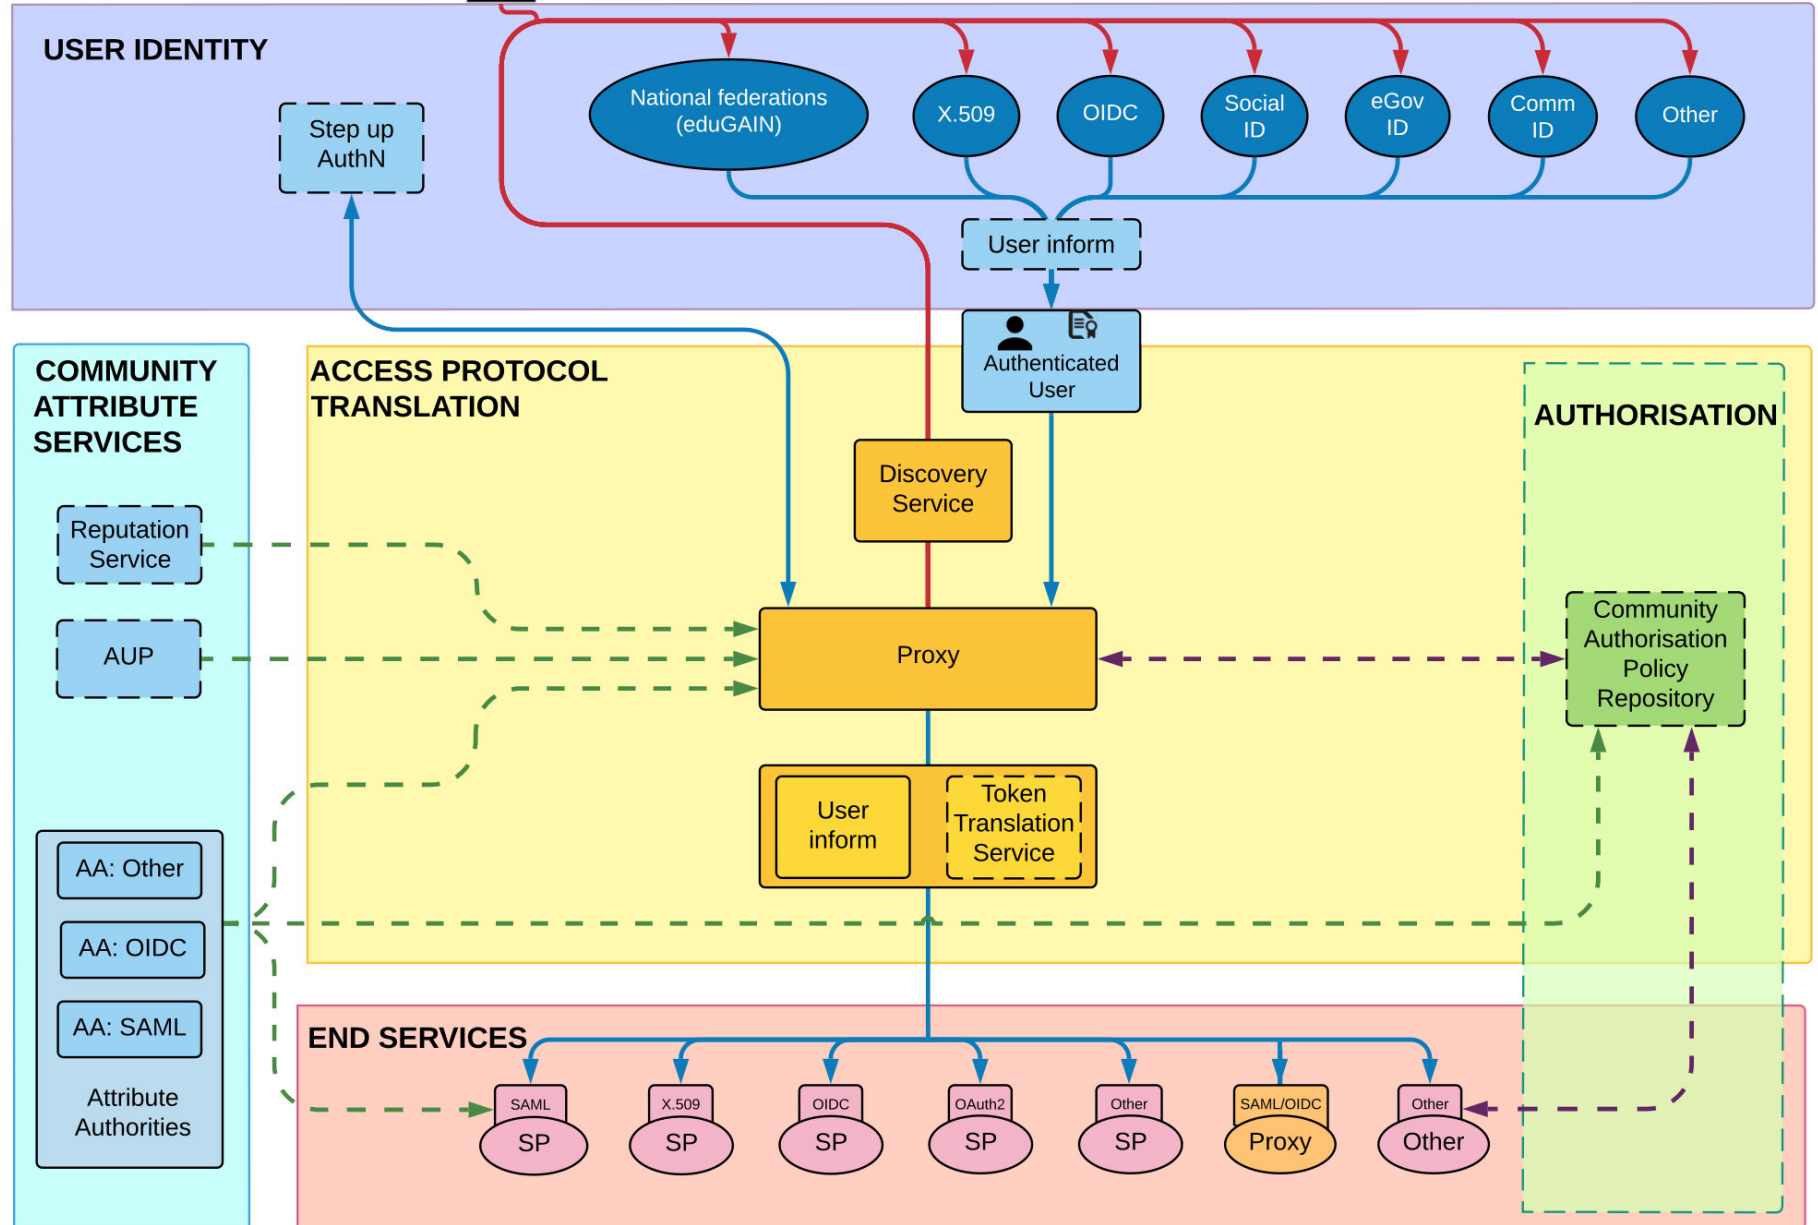
- In the planned WLCG infrastructure, there will be an OP per VO, which users may access using their CERN Account
- This model unifies the IdP and Research community, with both being represented by the Token Issuer
- The Infrastructure design has been informed by the **AARC Blueprint architecture** - a set of software building blocks that can be used to implement federated access management solutions for international research collaborations

The AARC Blueprint

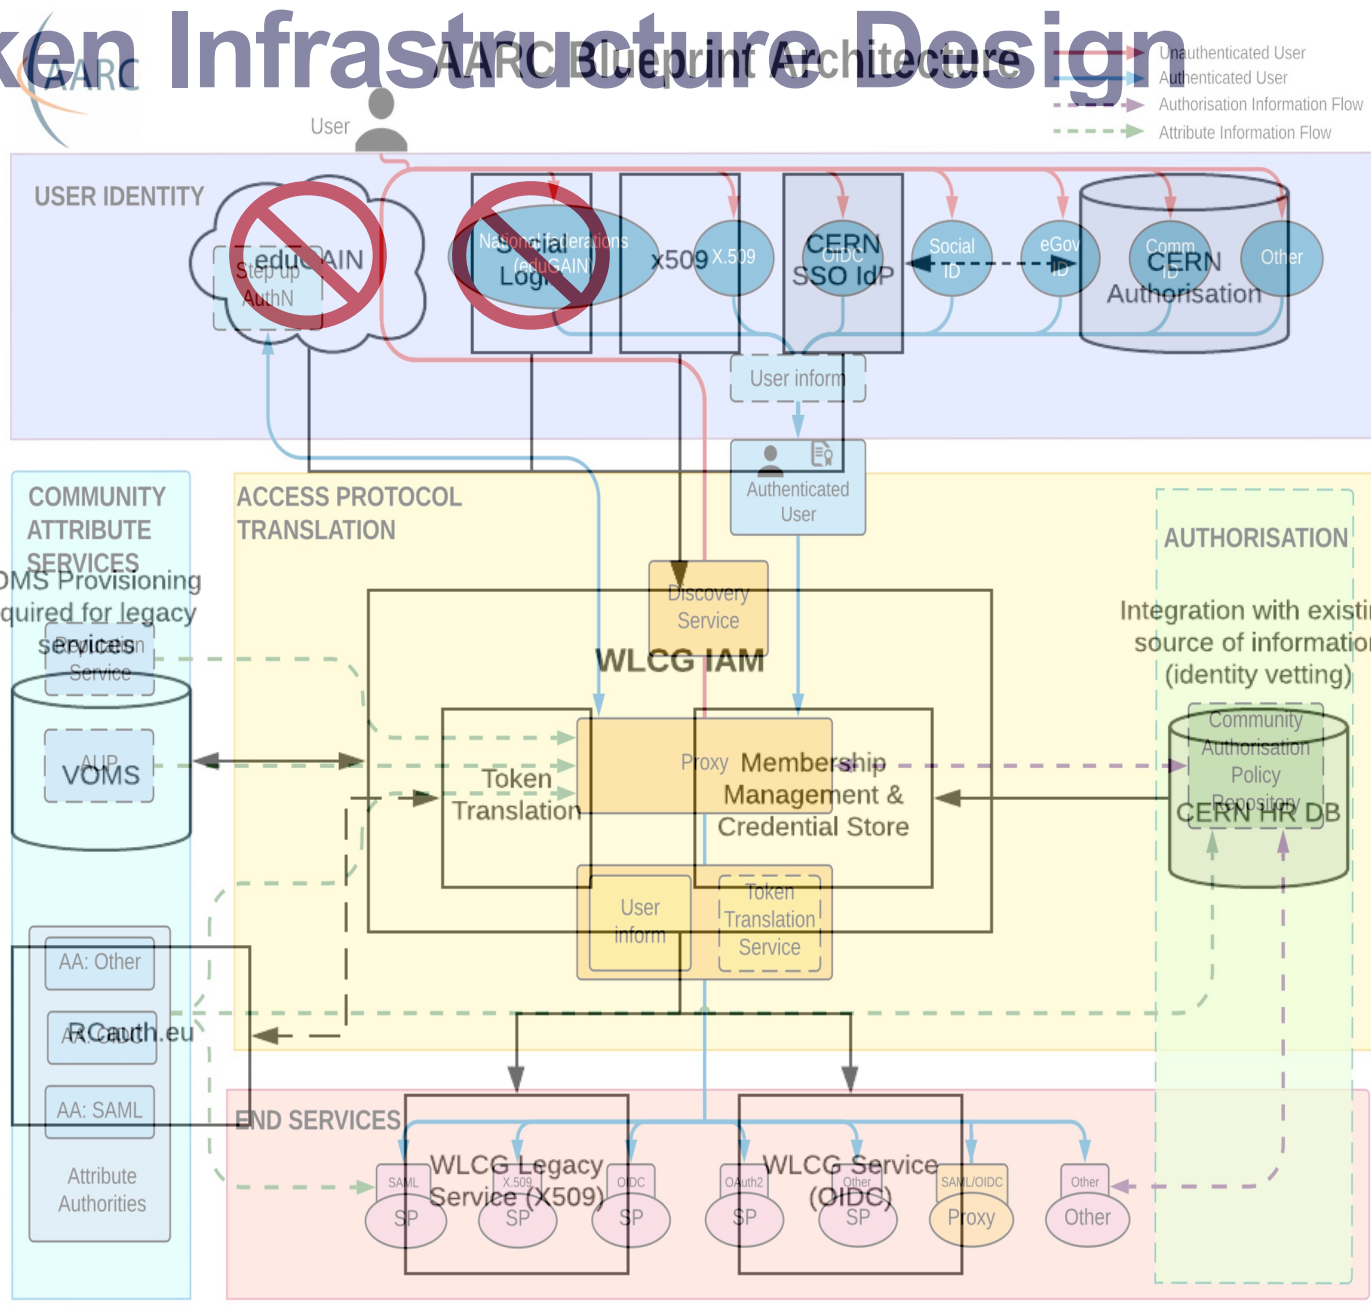


AARC Blueprint Architecture



- Unauthenticated User
- Authenticated User
- Authorisation Information Flow
- Attribute Information Flow



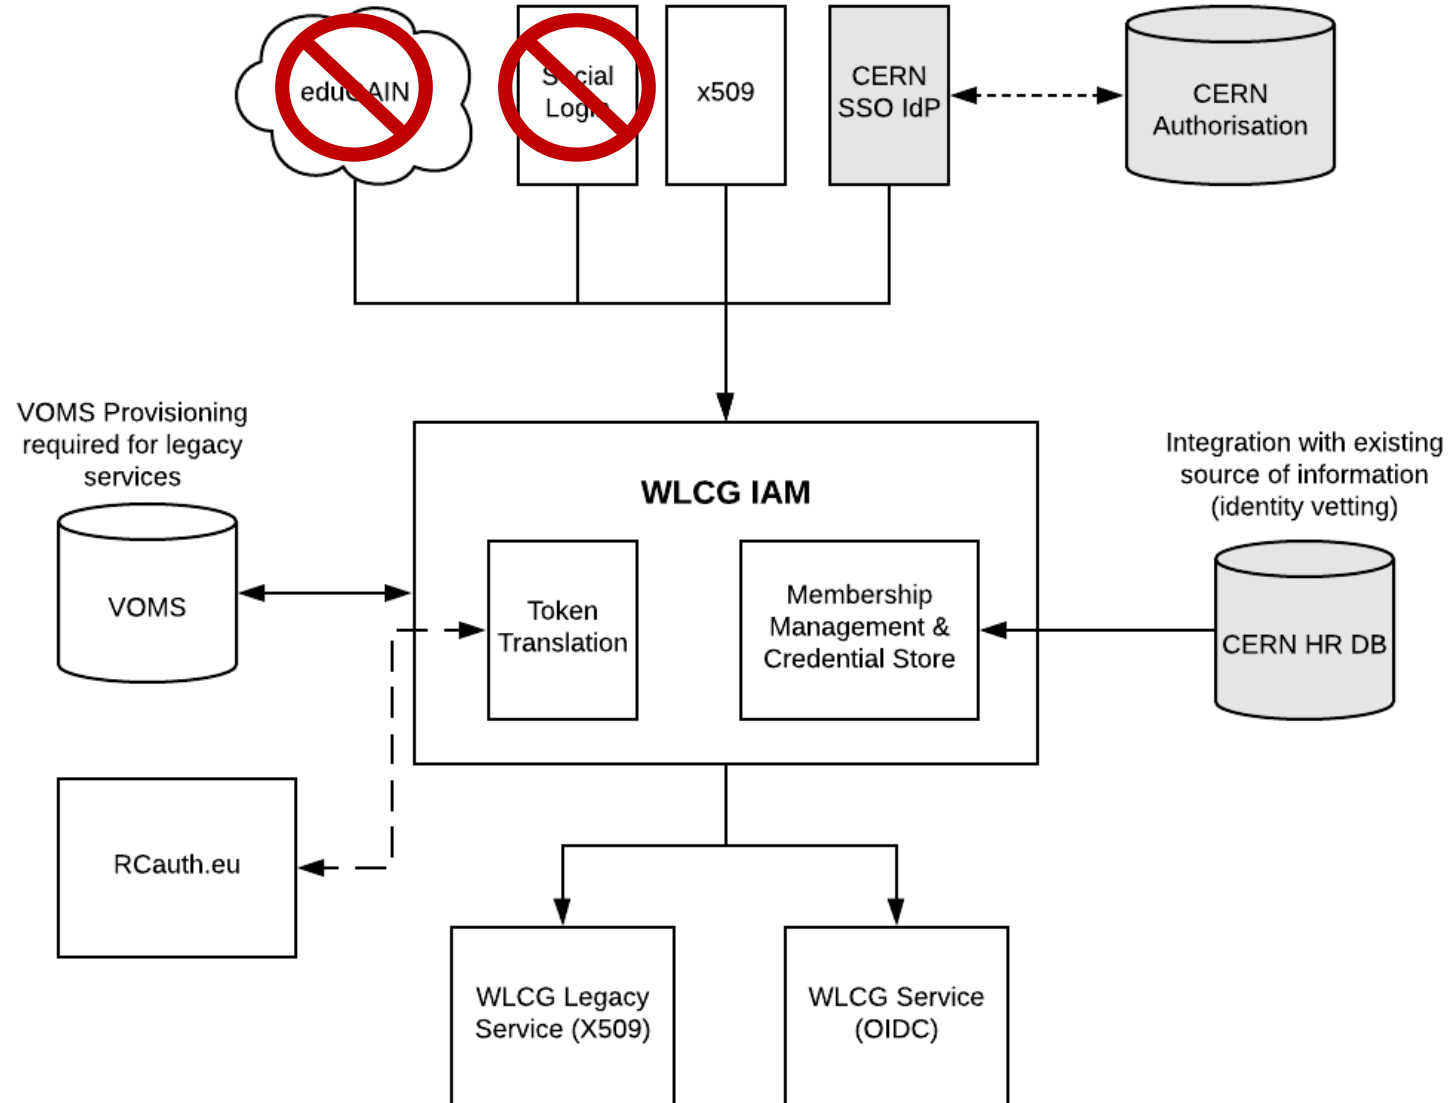
WLCG Token Infrastructure Design



COMMUNITY ATTRIBUTE SERVICES
VOMS Provisioning required for legacy services



WLCG Token Infrastructure Design

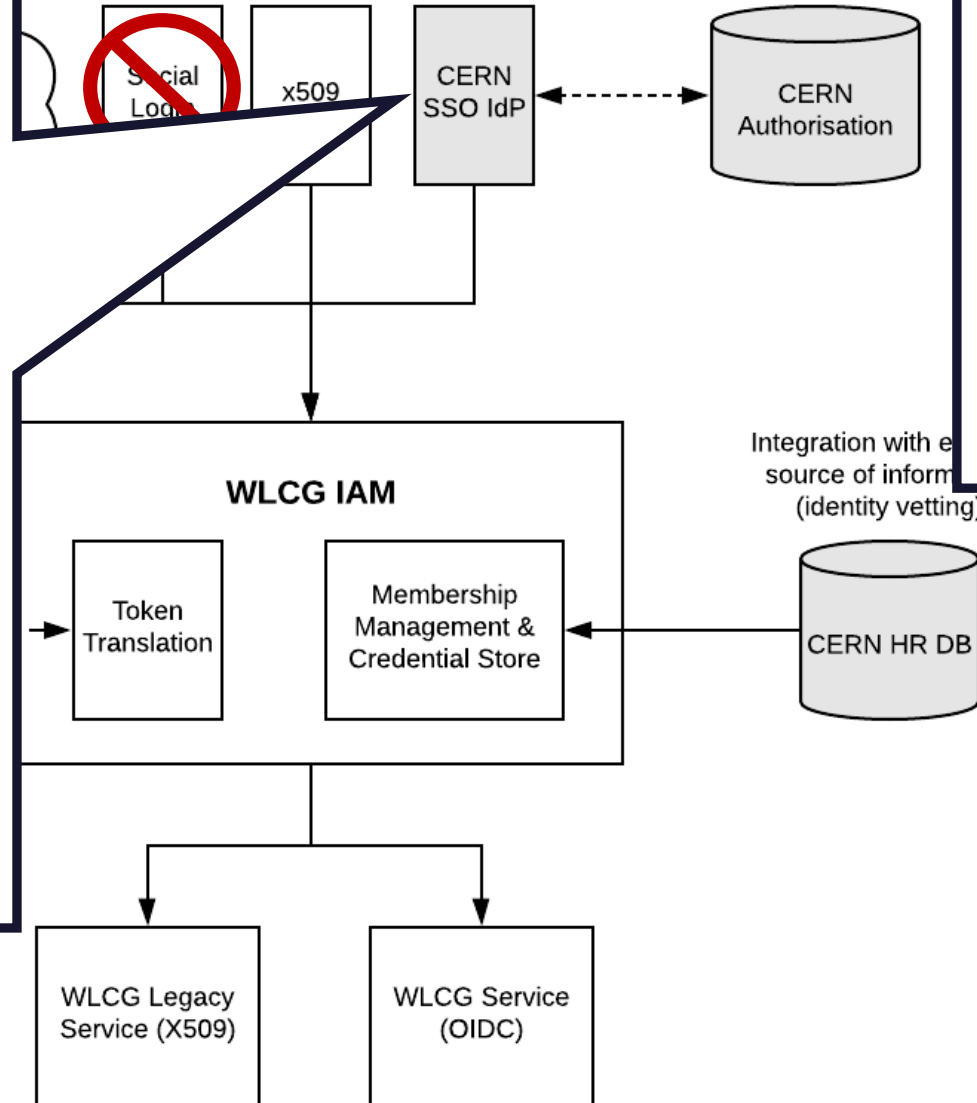


WLCG Token Infrastructure Design

CERN SSO releases:

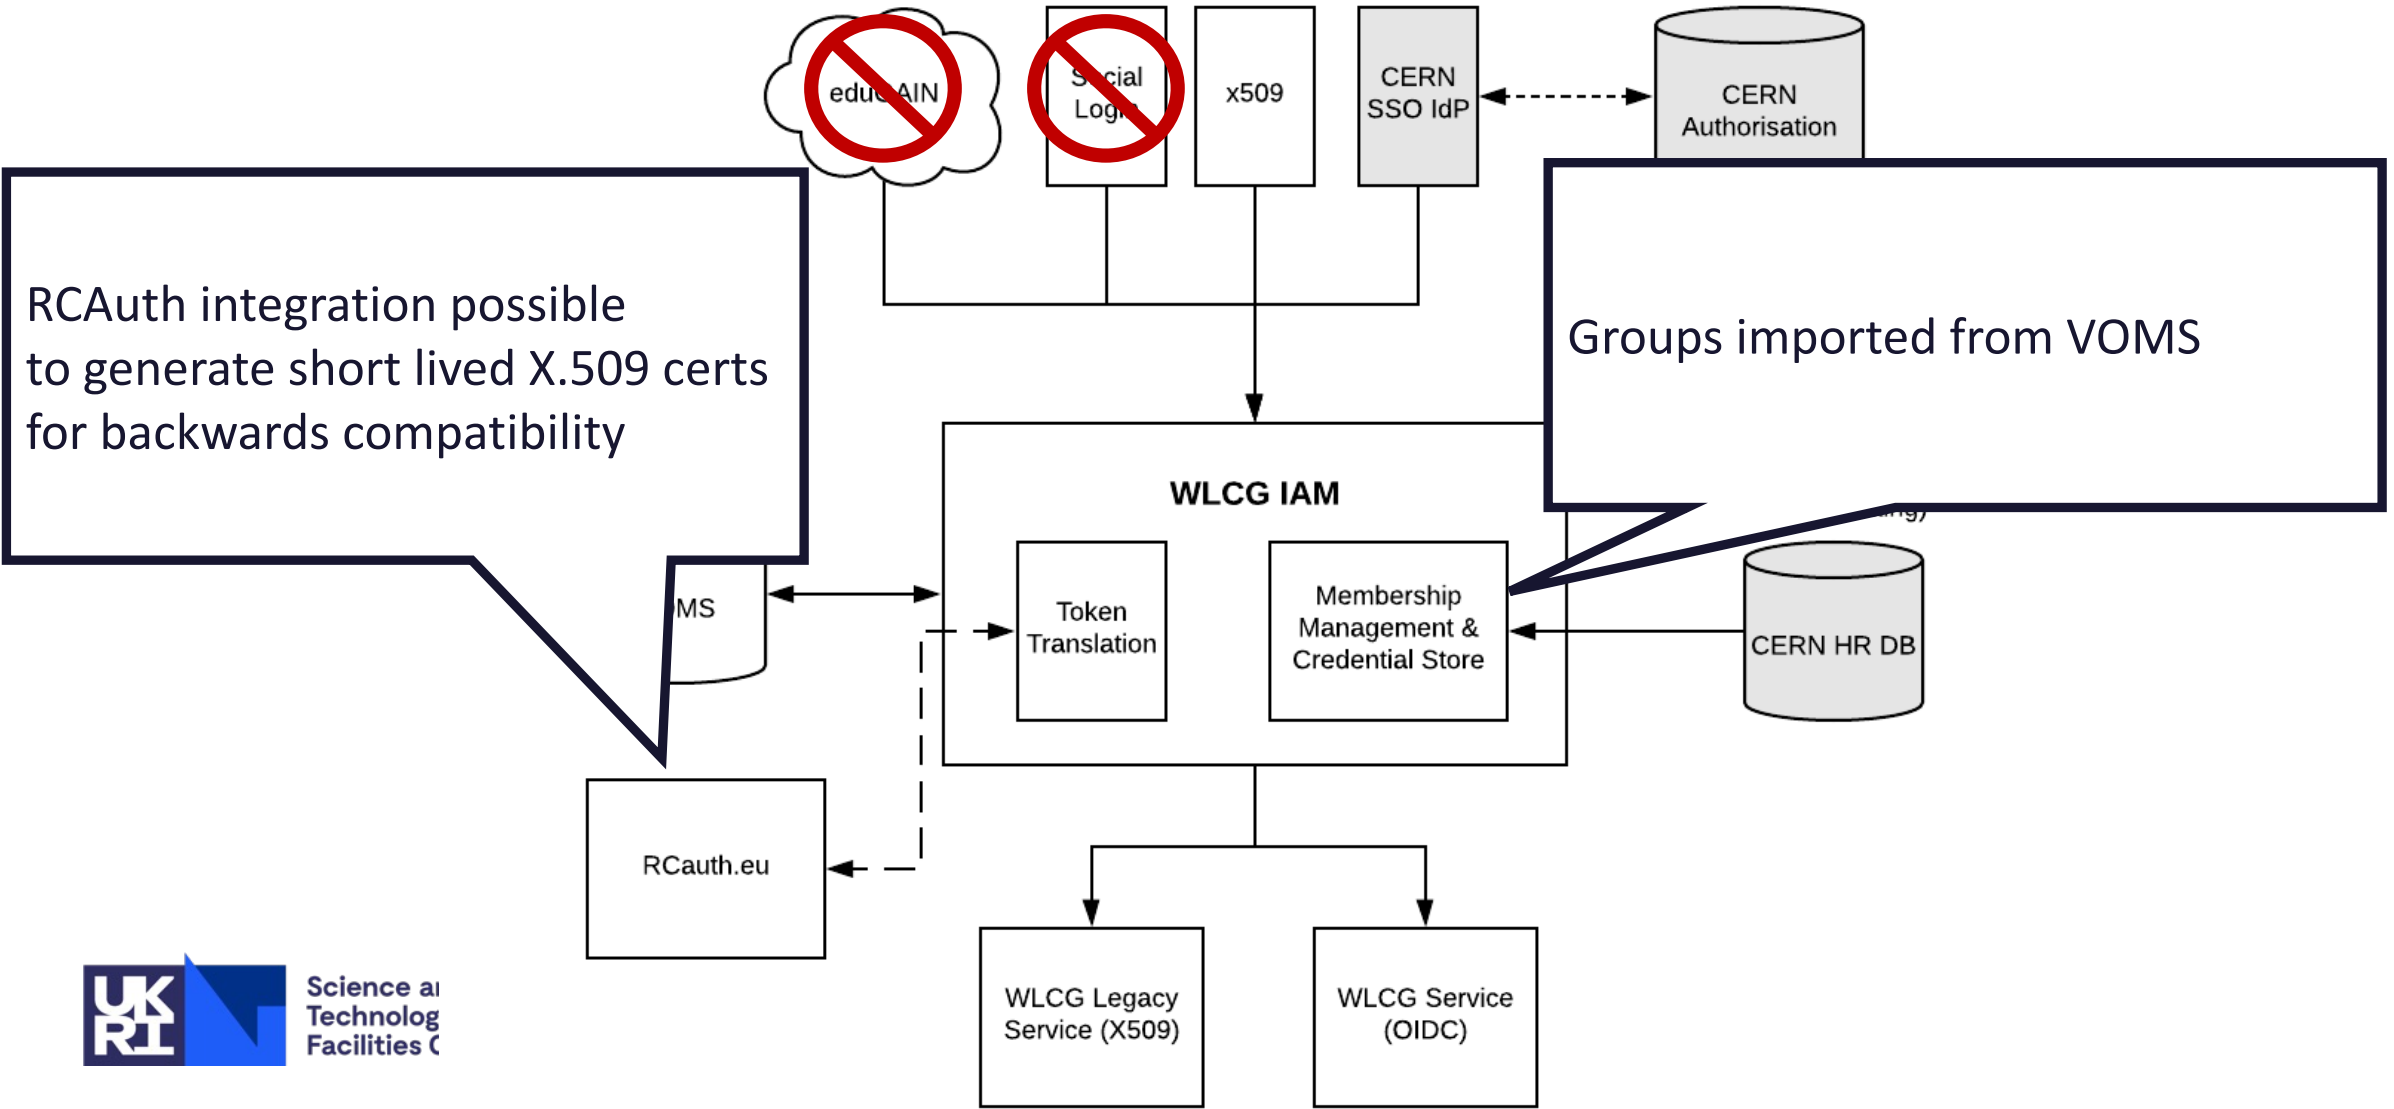
- Name,
- Email,
- CERN Person ID (indicates HR has performed ID check),
- CERN Kerberos Principal
- ...

Currently all researchers have CERN accounts but aim is to work towards removing this need in future



CERN Person ID is checked against CERN HR DB. Affiliation with Virtual Organisation (experiment) is verified, as well as end dates. If the check is OK, the membership is approved.

WLCG Token Infrastructure Design



WLCG Token Schema

- In order to serve authorisation information a VO, the WLCG token schema defines extra claims – `wlwg.groups` for Group-based authorisation and `scopes` for Capability-based authorisation
- `wlwg.groups` semantics are equivalent to existing VOMS groups, and will be initially imported directly from VOMS
 - Eg: `/atlas/production`
- `scopes` is used to provide capability to a specific token, rather than permanent authorisation to a user
 - Format `$AUTHZ:$PATH` where `$PATH` is mandatory (may be `'/'` for `*`)
 - Eg: `storage.read:/atlas`
- For more details, you can see the published schema:
<https://zenodo.org/record/3460258#.Y-YqUxPMLVs>



Science and
Technology
Facilities Council

Questions?



Science and
Technology
Facilities Council

Thank you



Science and Technology Facilities Council



@STFC_matters



Science and Technology Facilities Council