AARC Authentication and Authorisation for Research and Collaboration

# Trust Assurance and Policy for Federated Identity

## Requirements for Assurance and Policy along the AAI chain

**David Groep**

Policy and Best Practice Activity coordinator, AARC

*Nikhef Physics Data Processing programme*

Nikhef

Security Workshop ISGC 2016

March 2016

# Remember these days … ?

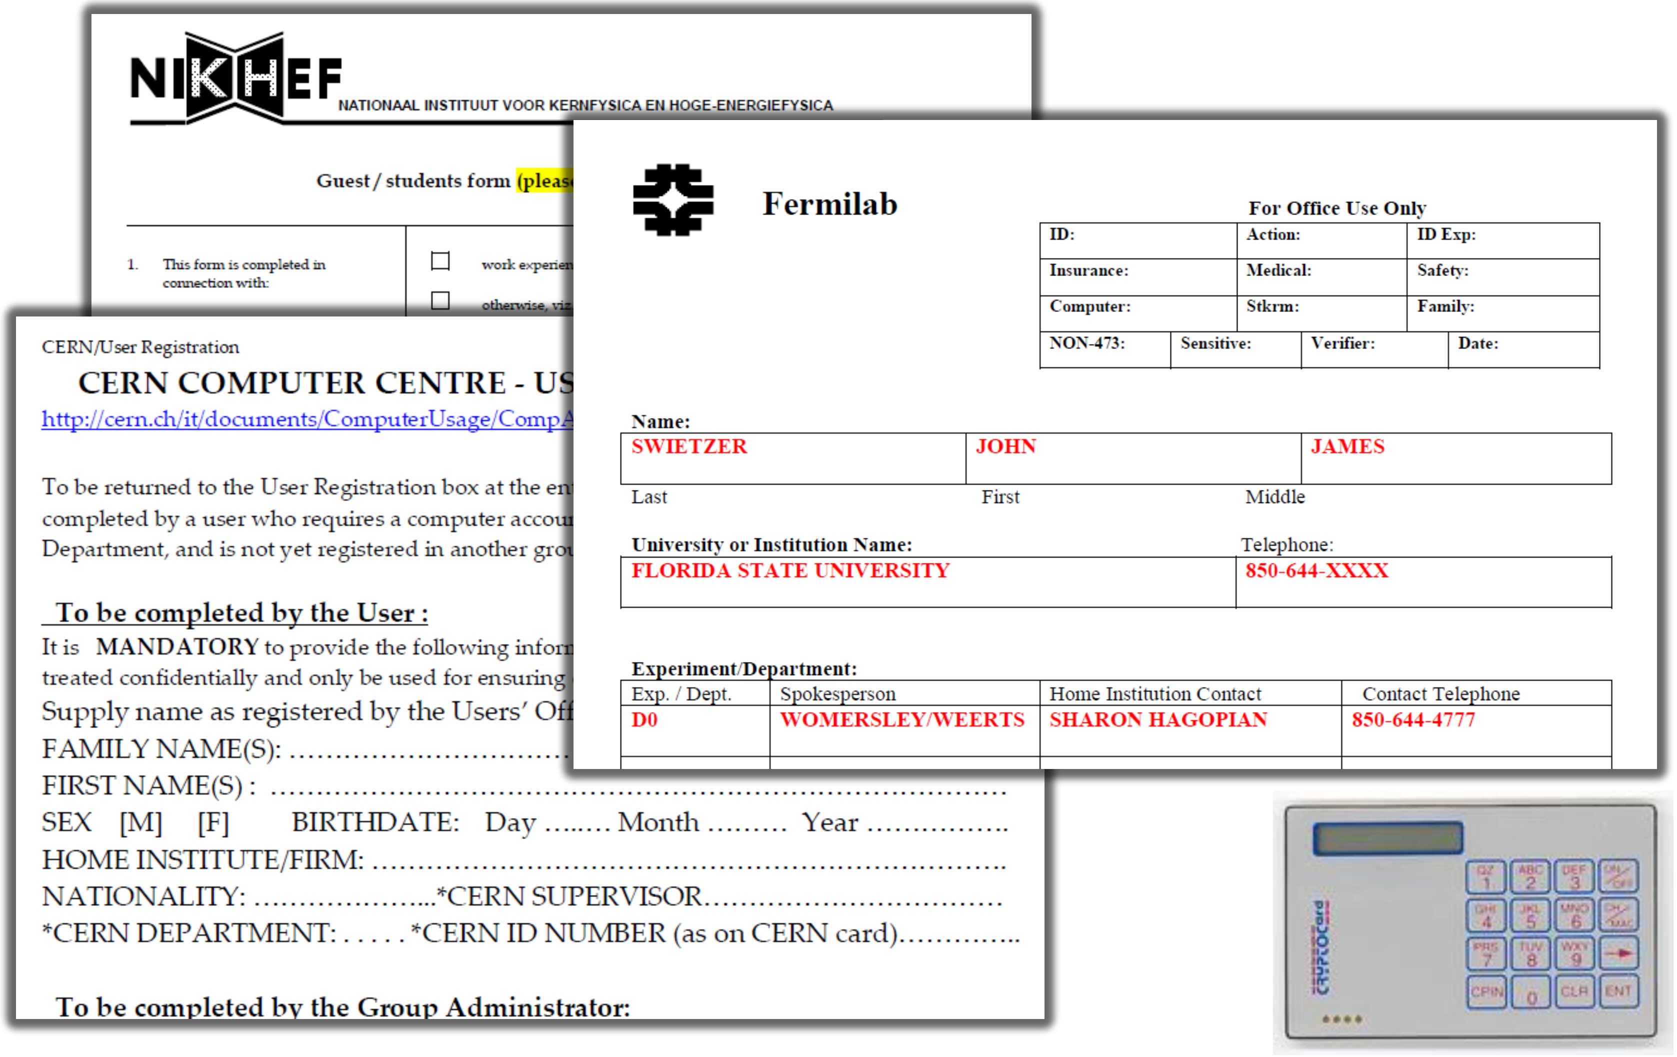

NIKHEF NATIONAAL INSTITUUT VOOR KERNFYSICA EN HOGE-ENERGIEFYSICA

Guest / students form (pleas

1. This form is completed in connection with: □ work experien □ otherwise, viz

CERN/User Registration

CERN COMPUTER CENTRE - US

http://cern.ch/it/documents/ComputerUsage/CompA

To be returned to the User Registration box at the en completed by a user who requires a computer accou Department, and is not yet registered in another gro

To be completed by the User :

It is MANDATORY to provide the following inforn treated confidentially and only be used for ensuring

Supply name as registered by the Users' Off

FAMILY NAME(S): ……………………………

FIRST NAME(S) : ……………………………………………………………………

SEX [M] [F] BIRTHDATE: Day …..… Month ……… Year ……………

HOME INSTITUTE/FIRM: ………………………………………………………………

NATIONALITY: ………………..*CERN SUPERVISOR……………………………

*CERN DEPARTMENT: . . . . . *CERN ID NUMBER (as on CERN card)…………..

To be completed by the Group Administrator:

Fermilab

For Office Use Only

| | | |
|---|---|---|
| ID: | Action: | ID Exp: |
| Insurance: | Medical: | Safety: |
| Computer: | Stkrm: | Family: |

| | | | |
|---|---|---|---|
| NON-473: | Sensitive: | Verifier: | Date: |

Name:

| | | |
|---|---|---|
| SWIETZER | JOHN | JAMES |
| Last | First | Middle |

| University or Institution Name: | Telephone: |
|---|---|
| FLORIDA STATE UNIVERSITY | 850-644-XXXX |

Experiment/Department:

| Exp. / Dept. | Spokesperson | Home Institution Contact | Contact Telephone |
|---|---|---|---|
| D0 | WOMERSLEY/WEERTS | SHARON HAGOPIAN | 850-644-4777 |

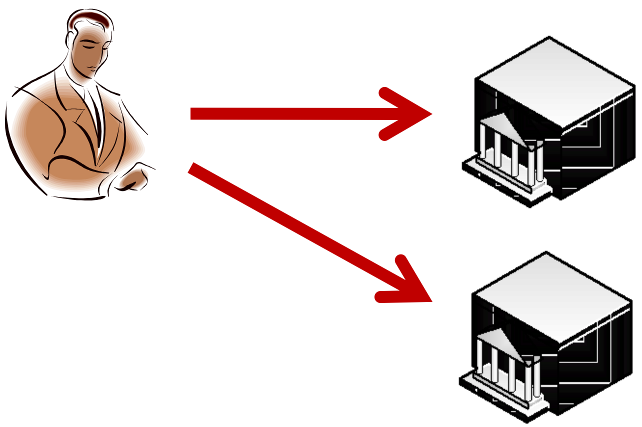

# And now we have this!

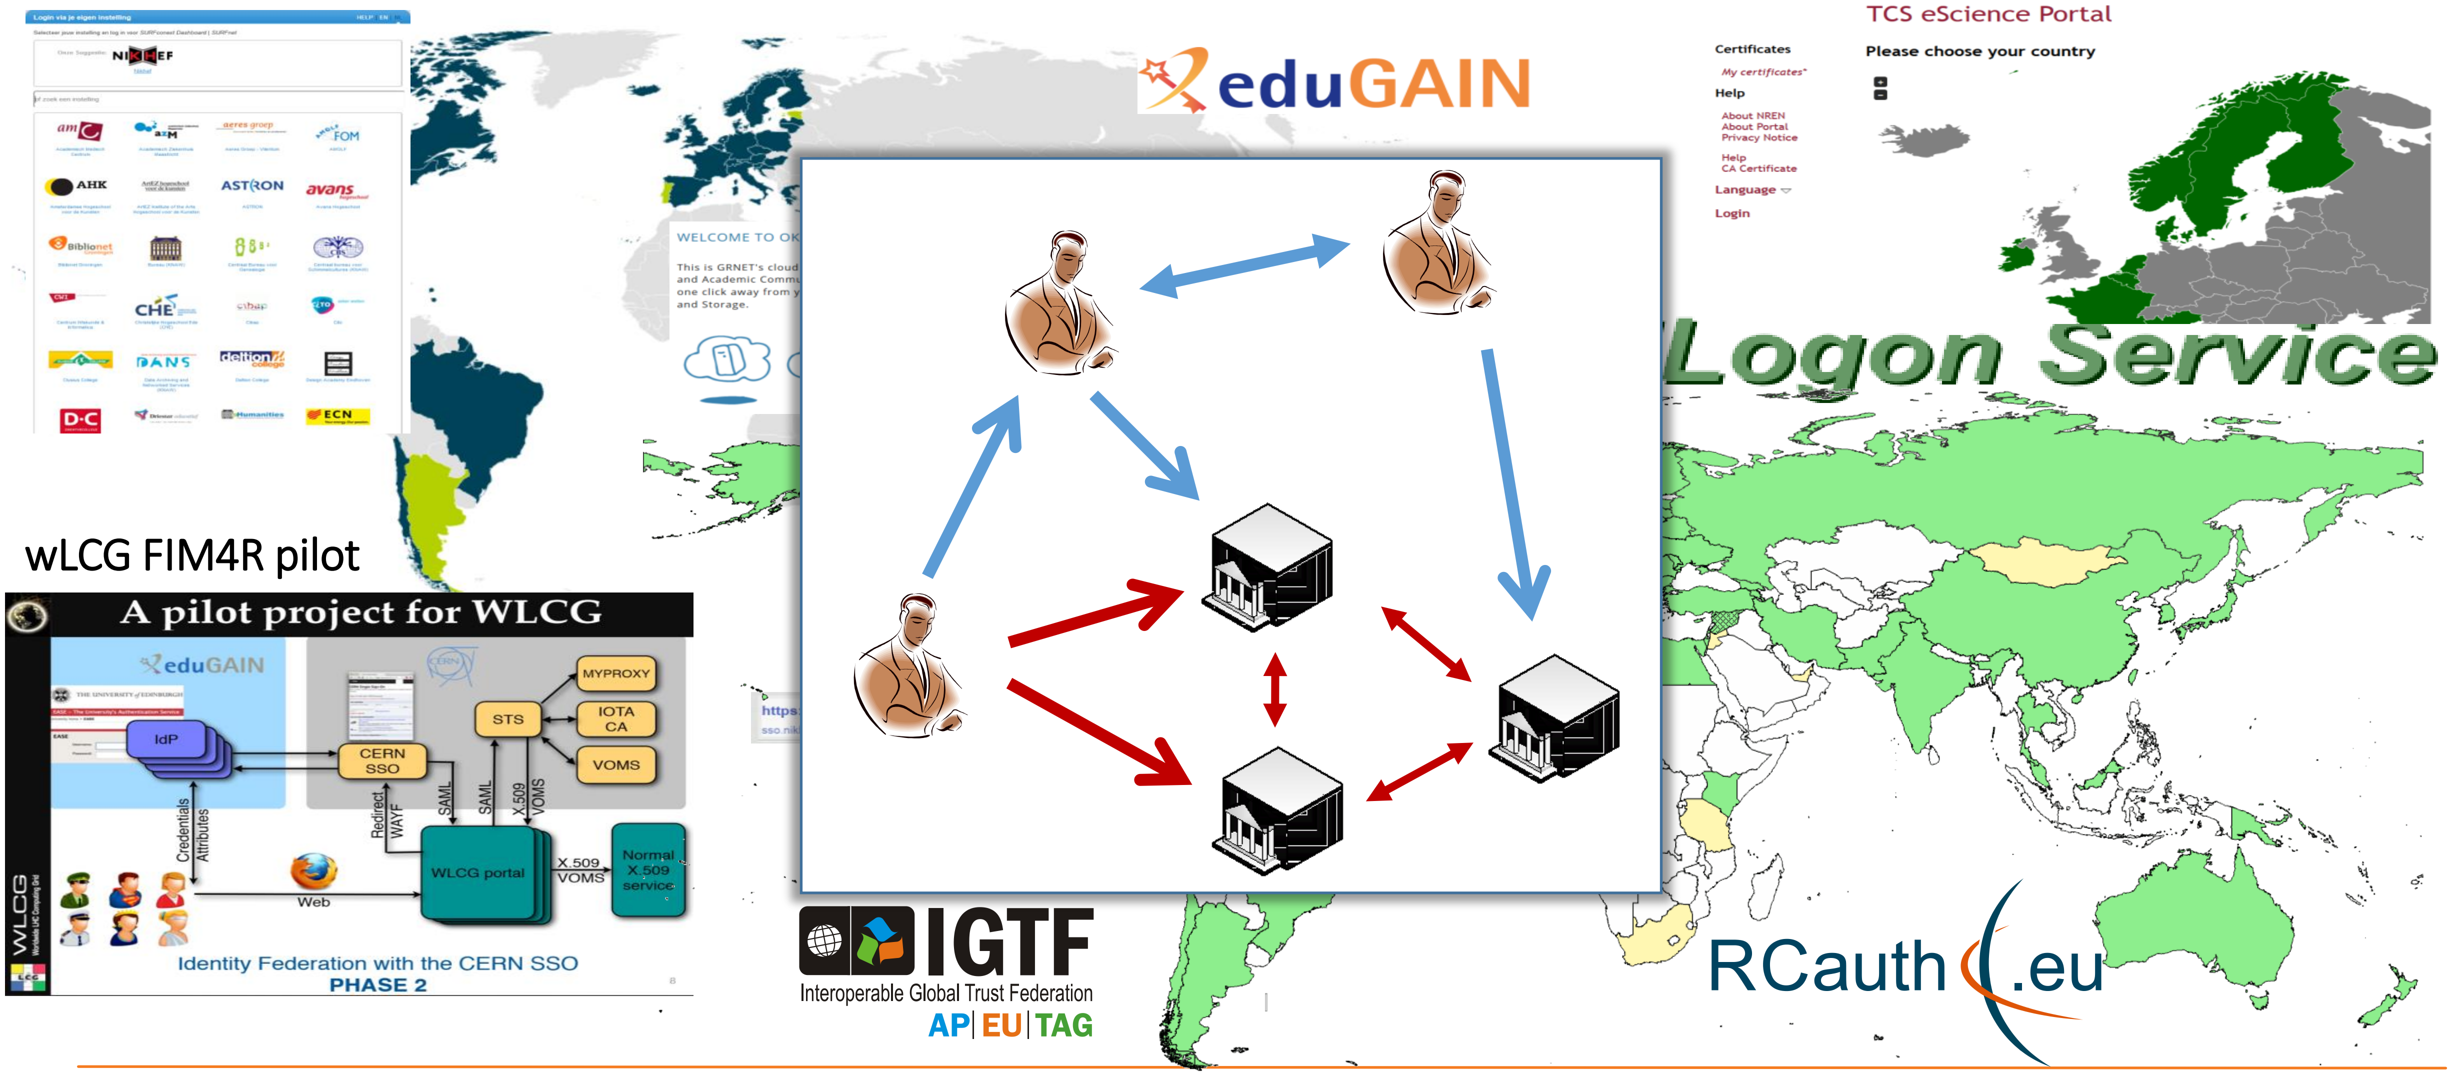

Login via je eigen instelling

Selecteer jouw instelling en log in voor *SURFconext Dashboard | SURFnet*

Onze Suggestie: Nikhef

of zoek een instelling

TCS eScience Portal

Certificates

*My certificates**

Help

About NREN
About Portal
Privacy Notice

Help
CA Certificate

Language

Login

Please choose your country

WELCOME TO OK

This is GRNET's cloud
and Academic Commu
one click away from y
and Storage.

wLCG FIM4R pilot

A pilot project for WLCG

Identity Federation with the CERN SSO
PHASE 2

background: *eduGAIN connected federations as of November 2014, and others – Brook Schofield, TERENA*

# Identifying responsibility – assessing risks

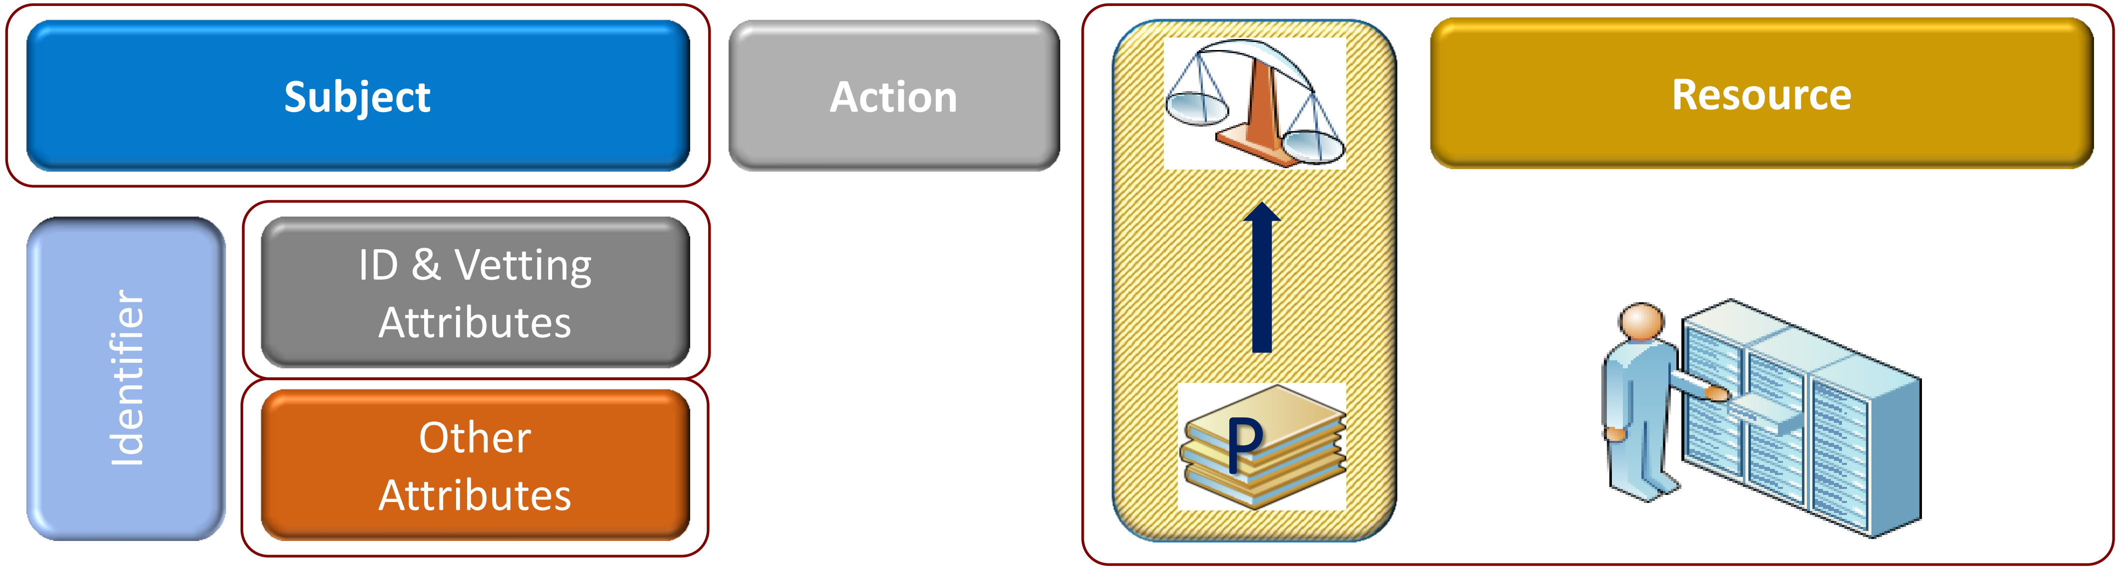

Independent administrative domains collaborate in **distinct roles**

# Chain of assurances …

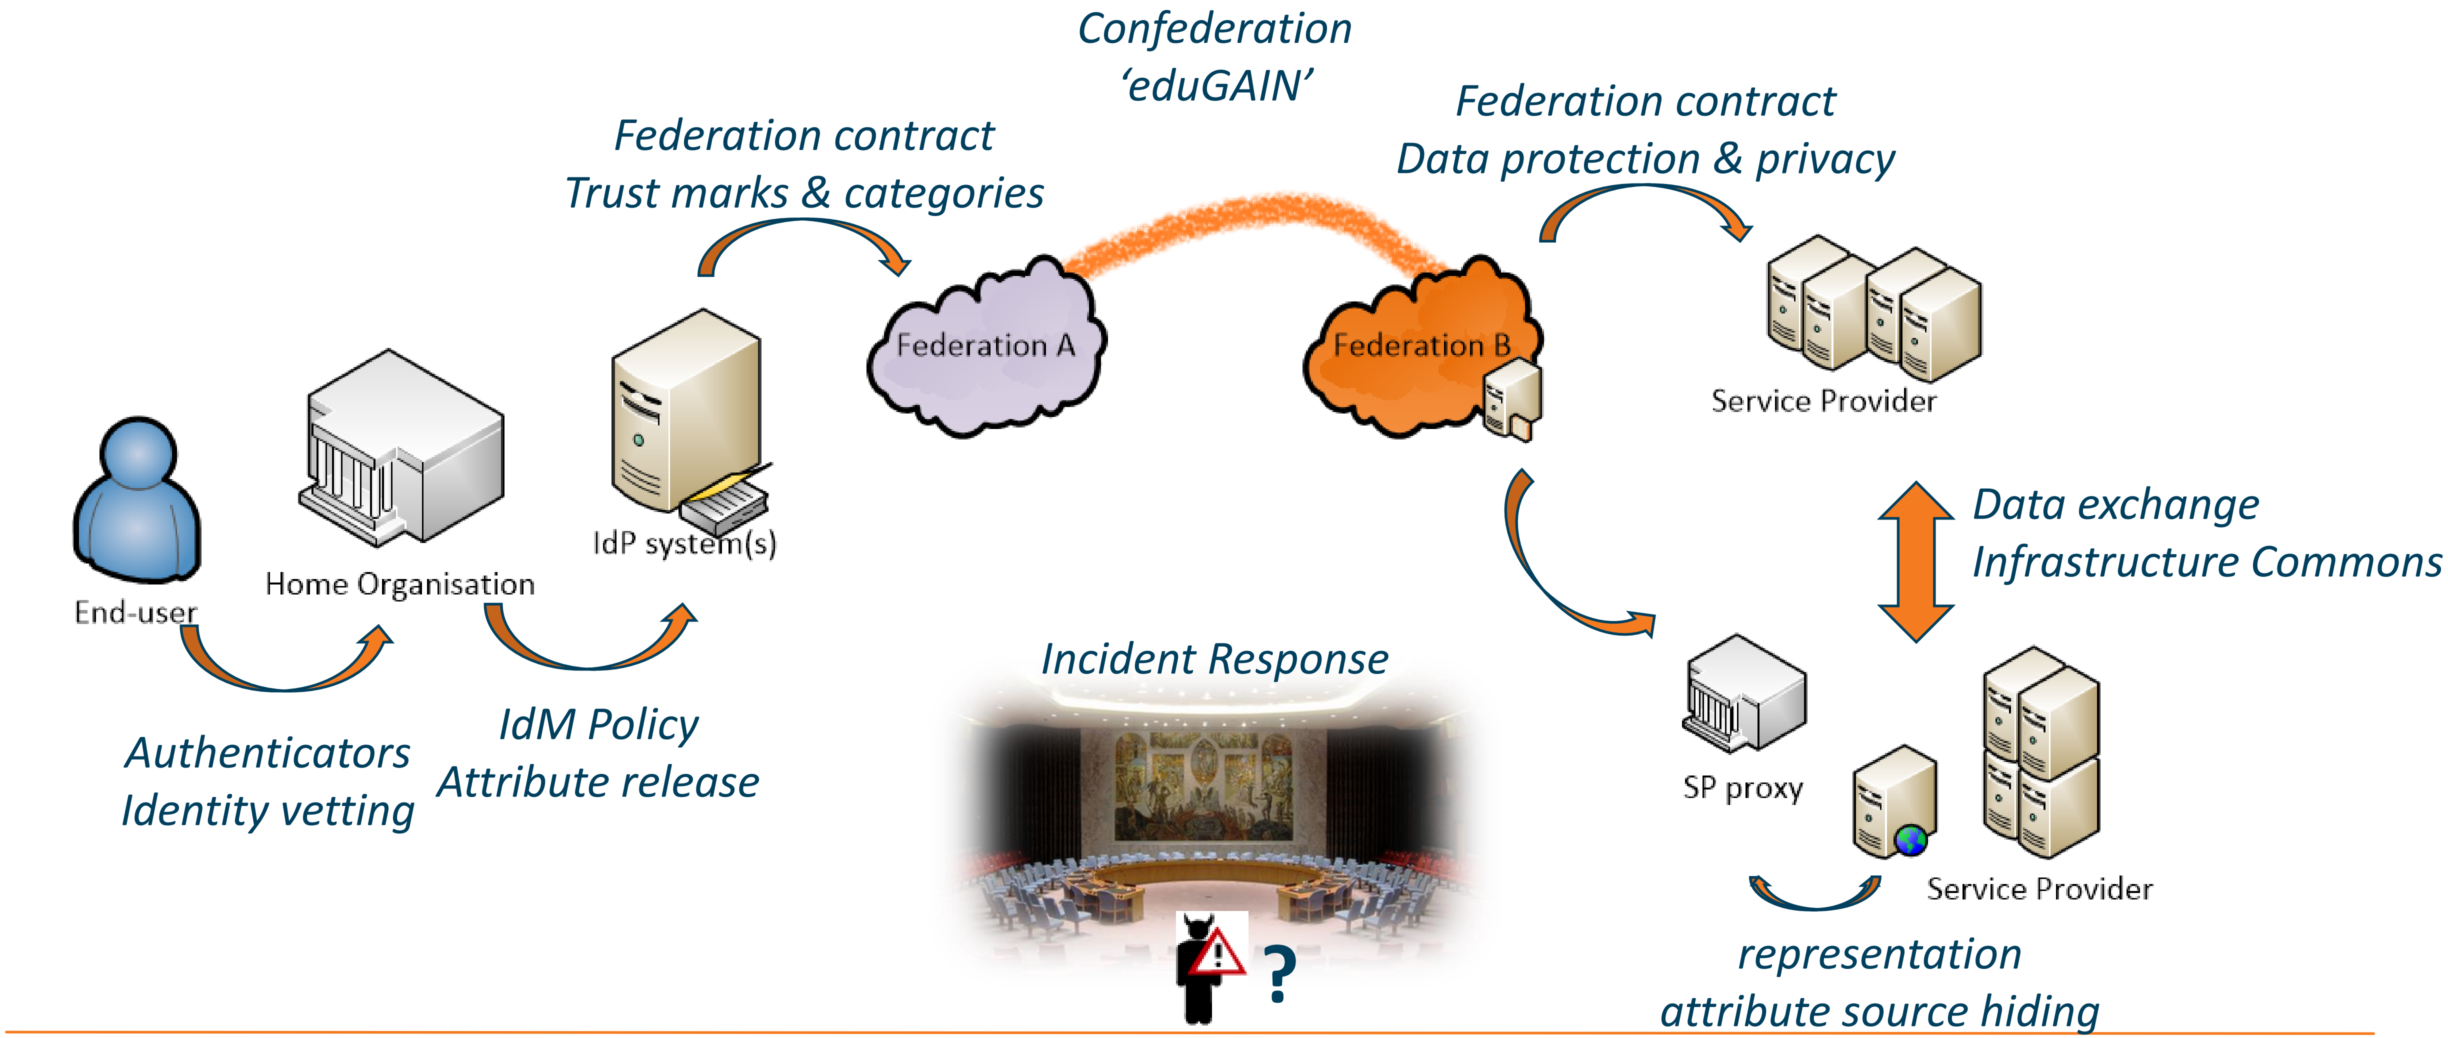

*Confederation ‘eduGAIN’*

*Federation contract*
*Trust marks & categories*

*Federation contract*
*Data protection & privacy*

*Data exchange*
*Infrastructure Commons*

*Authenticators*
*Identity vetting*

*IdM Policy*
*Attribute release*

*Incident Response*

*representation*
*attribute source hiding*

# Assurance as an element of risk assessment

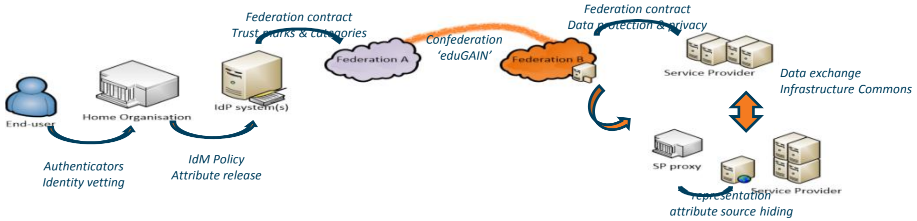

- Are the assertions provided by the IdP true?
  At the time of issuance? At a later time - i.e. can they revoke, and will they notify you?

- How do you trust the IdP? Does the registrar (federation) assert the right trust marks for it?
  Will the federation support you in contacting the IdP, and the IdP in finding the end-user?

- Is the federation meta-data complete and understandable?
  Can you make sweeping trust statement for the whole federation or
  do you have to assess each and every IdP?

- Do you see the IdP you're about to trust? Or do you just see a H&S federation masking it?
  Or is it a user community, hiding all of the world recursively with an 'IdP-SP proxy' setup?
  *When an IdP-SP proxy does attribute assertion, you cannot see if these are the original attributes … what is the assurance given by the VO/community?*

# Relying Parties as a Defining Element?

**Service providers (‘relying parties’) absorb almost all of the residual risk – as they host and manage resources under threat**

- Sources of ‘subject authority’ should **align with RP interests** to be useful
- RP must have policy controls to **compose sources of authority**
- RP must be **equipped with effective controls** to mitigate risks

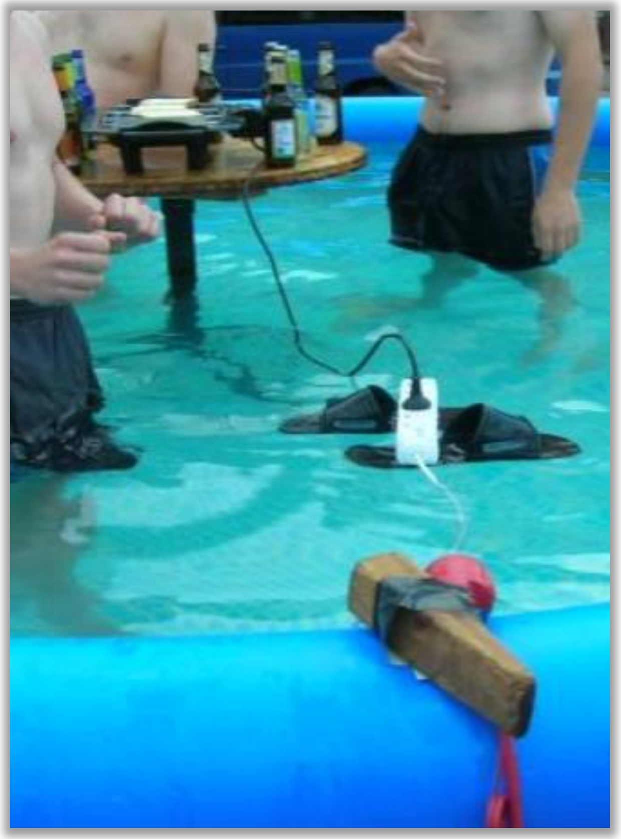

To *make* the risk assessment, you need to ‘predict’ the behaviour of others action chain:
Home organisation, its IdP, Federation Operators, Attribute Authorities, peer service providers

*source: NorthWood LAN party 7 - http://www.linuxno.de/*

# Expressing trust and assurance in different, complementary ways

## Formal and ITSM approaches

- ISO 27000 series compliance
- ISO 29115
  (NIST SP800-63, OMB M0404, OECD, Kantara)
- Assessment Frameworks
  (RFC 3647, ISO27002, SP800-53)
- External, periodic, independent audits
- Risk-based, multi-layer classifications

**But also**

- Costly, suitable mostly for large legal entities
- Depends often on regulatory pressure, so different in each country, even the mere specs

## Prevalent e-Infrastructure and RI approaches

- Standards-inspired frameworks
- Community standards (IGTF Assurance Profiles, GEANT Federation Template)
- Self-assessment and peer review
- Single or baselined classification of risk
- Personal trust groups

**But also**

- 'blind' trust (framework-only models, eduGAIN declaration)
- sometimes no trust definition at all
  (in a 'non-intervention' mode)

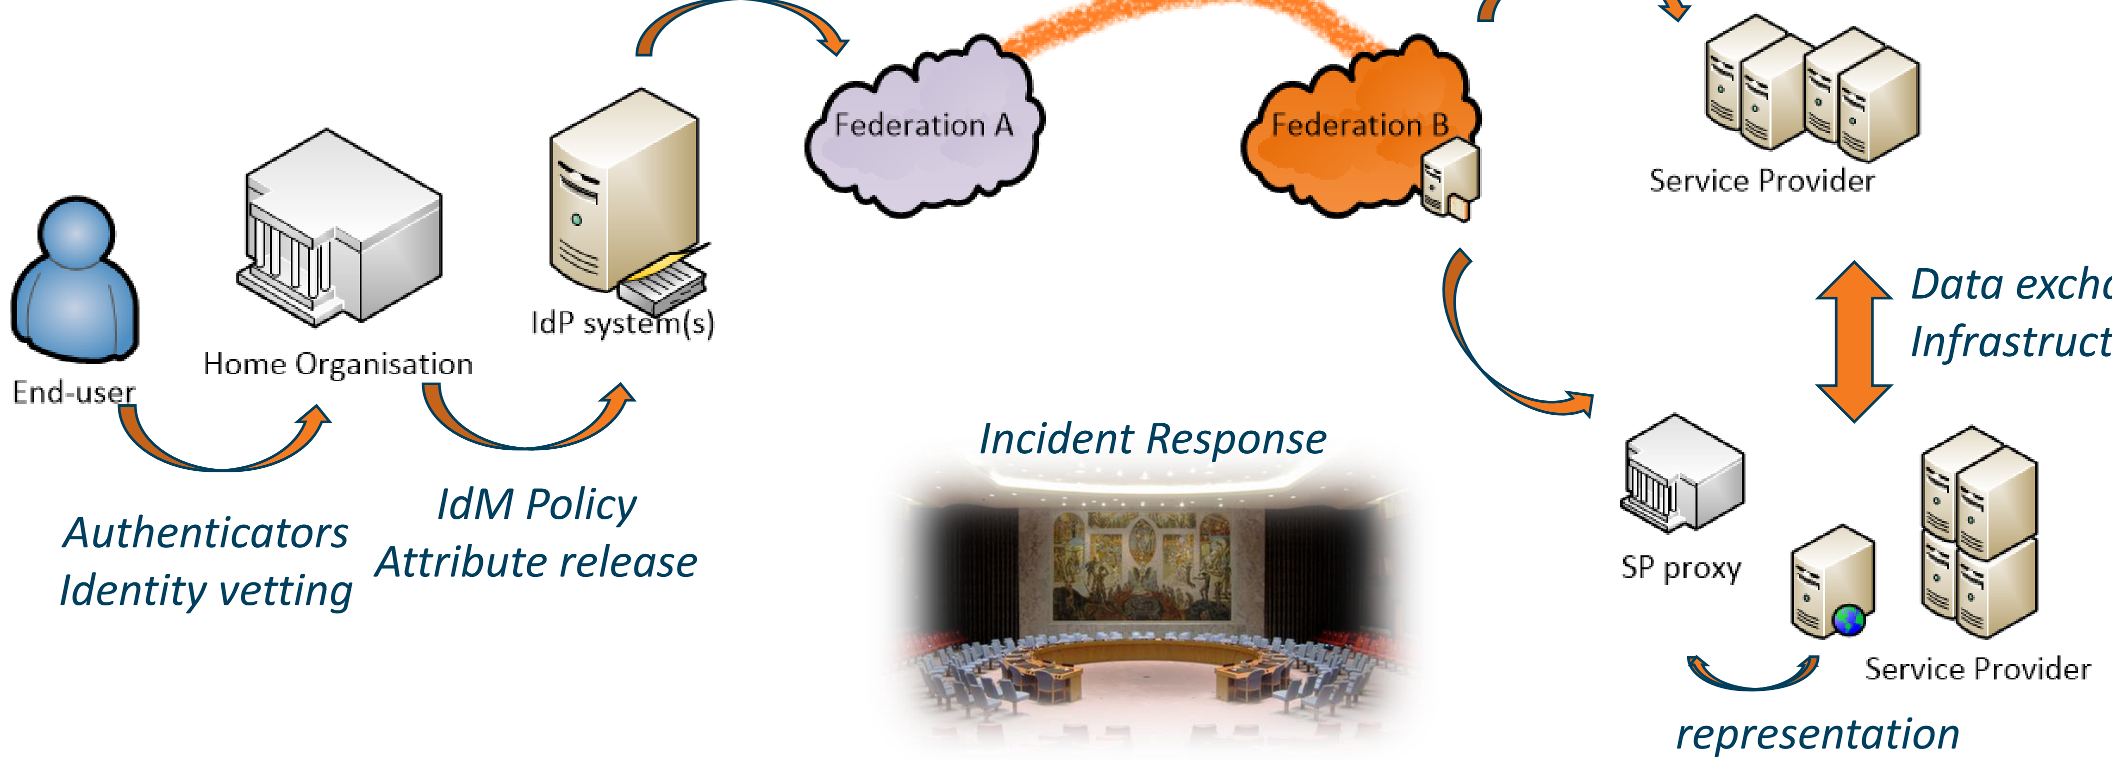

Confederation
'eduGAIN'

Federation contract
Trust marks & categories

Federation contract
Data protection & privacy

Data exchange
Infrastructure Commons

Incident Response

Authenticators
Identity vetting

IdM Policy
Attribute release

representation
attribute source hiding

# Authentication and identity vetting

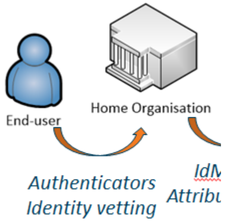

- A wide variety of assurance levels exists
  - Formal specifications: ISO29115, Kantara, NIST SP800-63 (OMB M-04-04) – usually four levels
  - Europe put in a new scheme for their “eIDAS” regulation, using three levels (including a “substantial”)
  - CI and e-Infrastructures around the IGTF now have ~ two levels (classic and ‘identifier-only’)
- Levels are usually a combination of authenticator strength (passwords, tokens, biometrics) and identity vetting assurance (face-to-face, two photo IDs, verified home address, naming)

## How to get a defined assurance level?

- Negotiate 1-on-1 with each IdP and get e.g. the eduPersonAssurance attribute set
- Rely on a ‘policy filter’ like the IGTF – only compliant and assessed IdPs make it through *and how that is done depends on the CA model, e.g. TCS uses specific entitlements, InCommon Silver uses eduPersonAssurance and institutional accreditation*
- For low assurance use cases, infer it from other IdP properties or ‘generate’ uniqueness *e.g. if an organisation has some incident response capability & send useful attributes …*

# A baseline LoA for federated identity

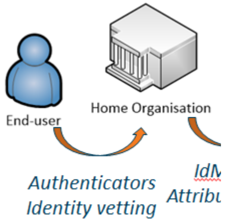

- At the moment, there is no *common* baseline assurance for identity. Individual federations may however have guidelines. E.g. at SURFnet:
  - Organisations adequately protect and configure systems and networks used for the IdP
  - Ensures that attributes of users are up-to-date and complete, and that the user has been identified on enrolment (but does not ask *how the user was identified in the first place!*)
  - Some attributes are mandatory (displayname, email, affiliation status, &c)
  - Access is restricted to a specific class of users only (so even if you express another affiliation, you still cannot authenticate third parties, even if you've identified them … and you cannot act as an IdP of last resort)
  - Organisations maybe asked to make their processes transparent (but that's not compulsory!)
  - They comply with SURFnet's privacy interpretation
- But this is – in a global context – already a very tightly controlled assurance level … there are much more open federations and IdPs out there

# A common baseline for assurance for e-Infrastructures

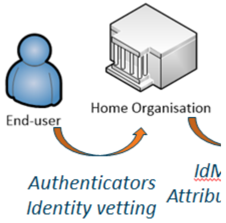

AARC MNA3.1 (Mikael Linden et al.):
**“Recommendation on minimal assurance level relevant for low-risk research use cases”**

*Based on depth-interviews with the major Research communities and e-Infra’s in Europe ...*

- Accounts belong to a known individual (i.e. no shared accounts)
- Persistent identifiers (i.e. are not re-assigned)
- Documented identity vetting (not necessarily F2F)
- Password authN (with some good practices)
- Departing user’s account closes/ePA changes promptly
- Self-assessment (supported with specific guidelines)

*A REFEDS task force will evolve these recommendations into globally implementable guidelines*

*https://wiki.geant.org/display/AARC/LoA+-+Level+of+Assurance*

# Expressing identity assurance

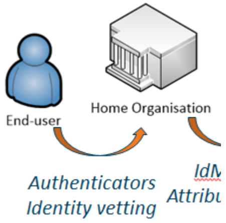

The CIs and research infrastructures are ‘happy’ with self-assessment & peer review while it works well is smallish (< ~150 people) communities, anything beyond this ‘Dunbar limit’ needs some documented trust process and automation support

## Self-assessment tool*

- Self-completion of maturity surveys – with optional peer review indicators
- Distribution of assessment status to federation participants, also via SAML meta-data
- Comparison of status between assessed participants (IdPs & SPs, against various policies)

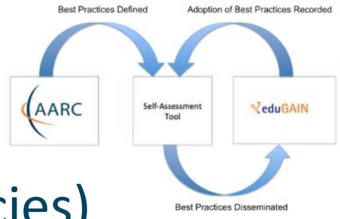
Best Practices Defined
Adoption of Best Practices Recorded
AARC
Self-Assessment Tool
eduGAIN
Best Practices Disseminated

## IETF registry for LoA

- Defined vocabulary for IdP LoA
- RFC 6711 (Leif Johansson) established https://www.iana.org/assignments/loa-profiles/
- One new AARC/REFEDS artefact (Sirtfi, https://refeds.org/sirtfi) submitted recently!

* https://docs.google.com/document/d/10kguCdxWn38z_EGRnrdjCI4GSeO44zFGeXWHGmzz27o

# Attribute release & IdM policy

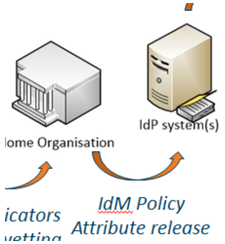

*So the organisation can authenticate and qualify users … and now what?*

- you have seen various models for federation: hub-&-spoke, full-mesh, and hybrids
- how can the recipient know about user identity, status, qualities? Via an attribute statement:

```
<samlp:Response ID="_2af69a..." Version="2.0" IssueInstant="2016-03-12T04:05:57Z"
              Destination="https://engine.surfconext.nl/authentication/sp/consume-assertion" ... >
    <saml:Issuer>https://sso.nikhef.nl/sso/saml2/idp/metadata.php</saml:Issuer>
    <ds:Signature xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
<saml:AttributeStatement>
  <saml:Attribute Name="eduPersonUniqueId“ NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:basic">
  <saml:AttributeValue xsi:type="xs:string">40ea621a0a7355cf4fb1ca8d4f22a53d@nikhef.nl</saml:AttributeValue>
```

But you may not get information: many considerations pose challenges for an IdP operator:

- some are assuming decisions on behalf of their users (unwillingness to release)
- unclear regulatory constraints scare IdPs operators:
  data protection in the EU + New Zealand, Argentina, Canada, Uruguay, &c ; FERPA in the US
- differing rules per country make for complex interfederation negotiations

Yet most ‘hub-&-spoke’ federations provide some *collective qualities* for their amalgamated IdPs

*https://addons.mozilla.org/en-US/firefox/addon/saml-tracer/*

# Attribute sets and release policies

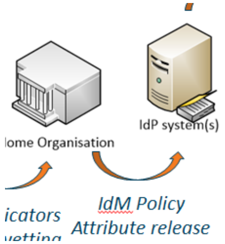

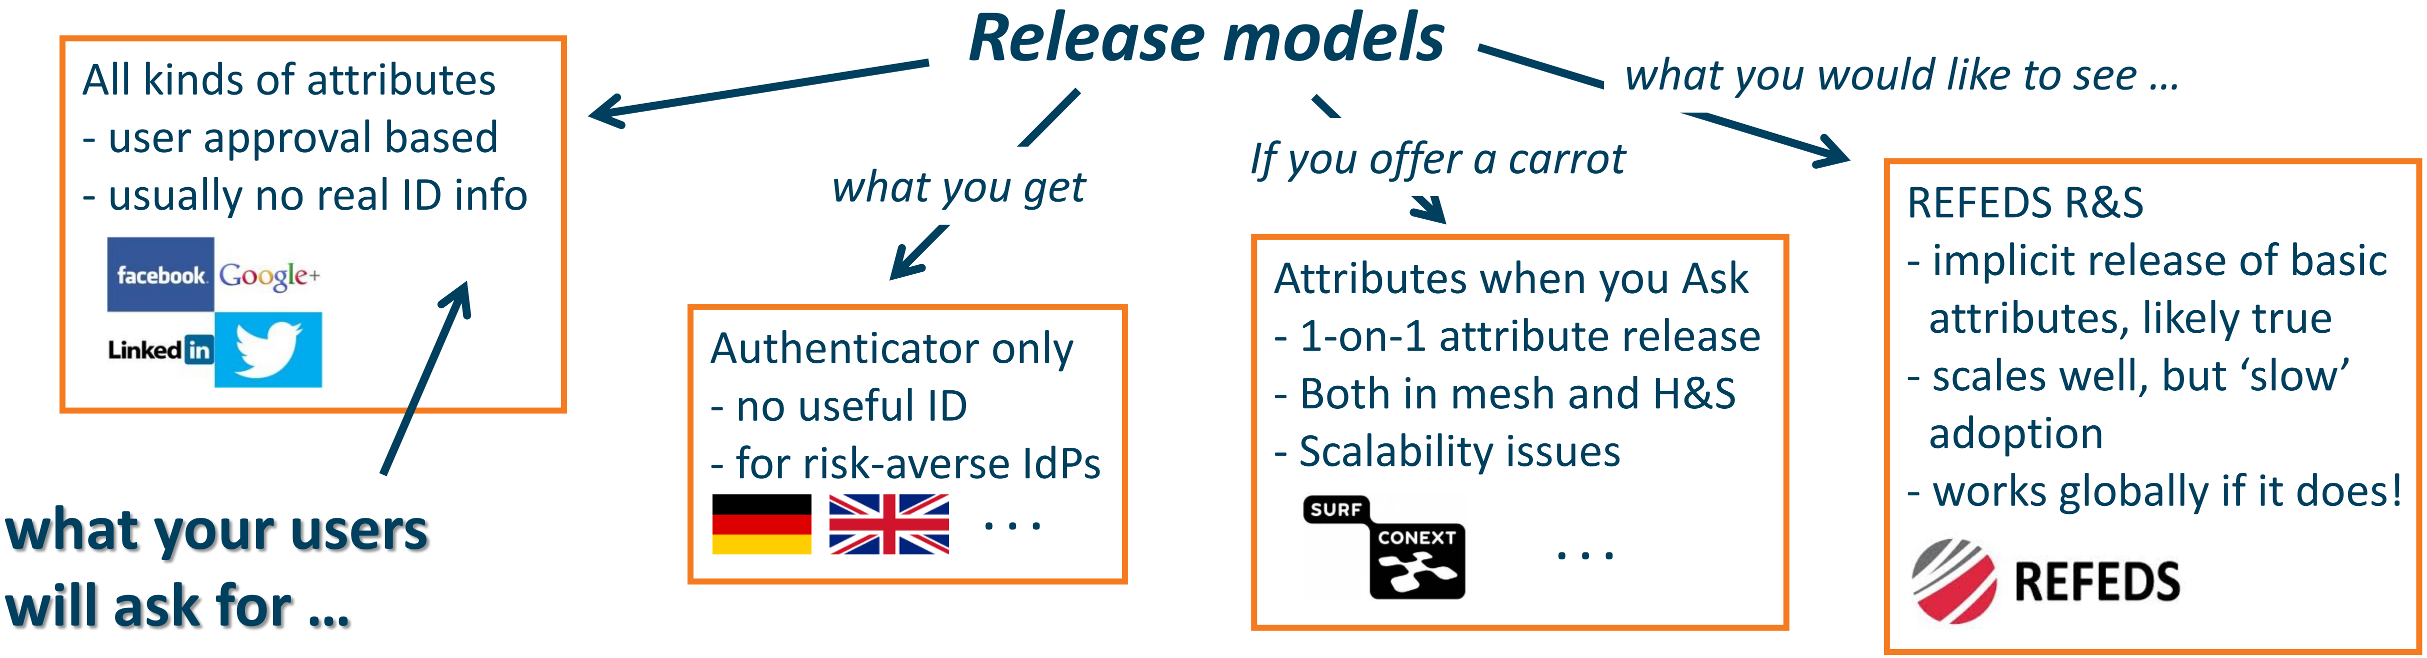

**Release models**

All kinds of attributes
- user approval based
- usually no real ID info

*what you get*

Authenticator only
- no useful ID
- for risk-averse IdPs

. . .

*If you offer a carrot*

Attributes when you Ask
- 1-on-1 attribute release
- Both in mesh and H&S
- Scalability issues

. . .

*what you would like to see …*

REFEDS R&S
- implicit release of basic attributes, likely true
- scales well, but ‘slow’ adoption
- works globally if it does!

**what your users will ask for …**

*“**R&S Service Providers MUST** resolve issues of non-compliance within a reasonable period of time from when they become aware of the issue. Failure to do so MUST result in revocation of the entity’s membership in the R&S category.” – but **nothing on IdPs** ☹*

**Research and Scholarship Attribute Set**

**displayName** *or* (**givenName** *and* **sn**)
**mail**
**eduPersonPrincipalName**
eduPersonTargetedID
eduPersonScopedAffiliation
**[https://refeds.org/category/research-and-scholarship]**

# Federation policies for IdPs – Entity Categories and trust marks

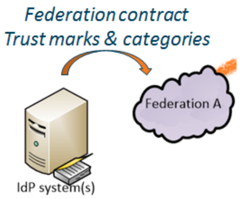

Federations operate according to very different assurance levels

- Rather strict: SURFnet, HAKA, SWAMID, …
- Very loose: ACOnet, InCommon, …
- Little common practice in older federations, new ones might follow Marina’s Best Practice*

**When assessing a federation itself for usefulness, look for these elements:**

- Technology Profiles (WebSSO, eduroam, …) – some technologies are transparent, others not
- Authentication LoA profiles – is there a common baseline in the federation
- Data Protection – can IdPs in that federation work with you and release attributes?
- Operating practices (service level, contacts, &c) – are they willing to talk to you, as an outsider?
- Governance and eligibility – what kind of entities can you expect in this federation?

** https://wiki.refeds.org/display/FBP/Policy+Templates*

# Interfederation – eduGAIN framework and interfederation

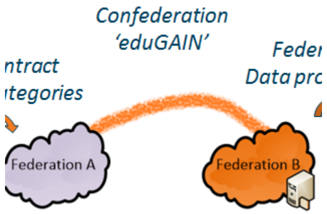

eduGAIN is a inter-federation scheme to connect federations to each other

- although the one global known entity,
  it is not *supposed* to be seen by anyone but federation operators
- yet an SP and IdP should be aware of it, since it conveys your assertions and impacts trust
- Limited policy around it: a top-level Declaration and a Constitution

**Governance is by federation operators, and at two levels**

- Public policy level Steering Group, and an Executive (actually: just the GEANT Project Executive)
- Operations Team (Brook Schofield)

**Besides this, there are informal groups where things actually happen …**
the REFEDS Federation Operators Group (FOG)

- wider scope than just eduGAIN, and more practical - it's the one reliable comms mechanism
- closed - with Fed operators usually concerned about PR issues like outages handled there

*services.geant.net/edugain/Resources/*

# Interfederation – eduGAIN policy-less operation

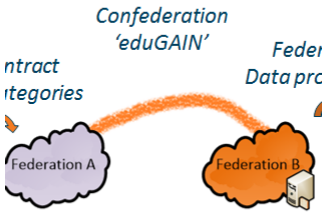

Technical implementation: a large meta-data feed at **http://mds.edugain.org/**

- You *could* filter on entities in there, and – as an independent service provider – you probably should
- But it’s not much encouraged: you need your own software to process this. But then deploying a few scripts is not too complex for a service provider (unless you’re hidden behind a classic hub-n-spoke federation operator), and e.g. *mod_mellon* has this as a default feature

## What to consider as an independent Service Provider

- Ability to suspend IdPs and registrars (=federations)
- Filter on Entity Categories (attributes in the meta-data)
- Subscribe to a central emergency suspension service? Of the self-assessment tool? But those are not there yet …

**Use the meta-data feed and MET to explore the world, via https://met.refeds.org/**

https://github.com/UNINETT/mod_auth_mellon - and see *MellonIdPIgnore*

# Federation policies: the impact on SPs

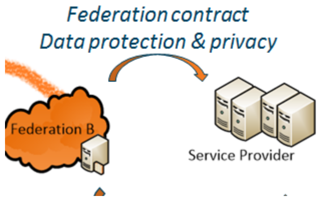

No way for general global SPs to connect directly to eduGAIN

Always has to go either via a federation, or via one-to-one agreements with IdPs

**But as an SP you can go federation shopping**

- Some require a fee for signup (e.g. InCommon)
- Some require you're either academic or have a contract in place (e.g. SURFconext)
- Some are open for everyone that can demonstrate a good reason (e.g. ACOnet)

**Same applies for policies**

- Some require strict adherence to e.g. a Data Protection Code of Conduct (e.g. HAKA)
- Some are more a conventional registrar and having good response and conditions is fine

There's quite a lack of 'market' for registrars in eduGAIN, but as a global SP you can pick your favourite ☺

# Data Protection and privacy

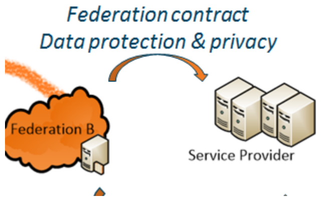

- Rapidly evolving global landscape related to privacy and data protection
- High-profile data leaks have raised public awareness, with Europe most pushy in this area *but other countries follow in order to do business …*

**You can declare adherence (in the EU) to a 'GEANT Data Protection Code of Conduct'**

- Uses mostly *informed consent* of the user as a basis for getting personal data from the IdP

**Plus**

- helps negotiating attribute release with IdPs, since they 'know' this model
- If you have a comprehensive policy framework, you may be compatible already

**Minus**

- mostly duplicates what's EU law anyway (so why state that again …)
- doesn't work outside of Europe – and the International CoCo can't progress for now

# Policy mapping and GEANT DPCoCo

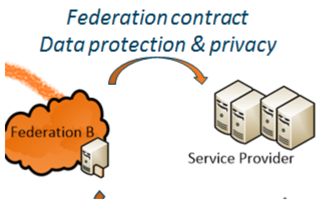

- Legal compliance: in SP jurisdiction
- Purpose limitation: collect only those data needed to enable access
- Data minimisation: also use the *least intrusive* attribute you can use for a purpose
- Deviating purposes: not permitted
- Data retention: delete when no longer needed (some federations say 6 months only! ☹)
- Third parties: no transfer, except if needed to give access, when 3rd party is also CoCo, or on consent
- Security measures: secure your service and protect data (see e.g. what the GDPR pushes for)
- Inform the end-users, via a privacy policy with certain rights.
  *But the right to delete data is not supposed to apply to security logs, luckily*
- Inform the IdP: via a machine-readable compliance statement (EC) and link to the Privacy Policy

Note that specific agreements take precedence, and it formally only works in the EU/EEA for now

https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home
http://geant.net/uri/dataprotection-code-of-conduct/

# Exchanging data between SPs

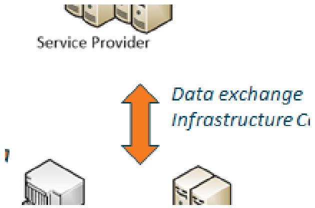

## Having agreed to DPCoCo or its principles, now what?

- To exchange use data (accounting) and logged information, you need to stay within purpose
- So in your policies and statement, add ‘security’ as a purpose of the data collection

> *Personal Data of End Users (hereinafter “Personal Data”) shall be Processed only for those administrative, operational, accounting, **monitoring and security purposes** that are necessary for the safe and reliable operation of Infrastructure services, without prejudice to the End Users’ rights under the relevant laws.*

*Make sure your peer service providers in the infrastructure sign up to the same policy suite*

“Personal Data may only be transferred to or otherwise shared with individuals or organisations where the recipient -

- has agreed to be bound by this Policy and the set of common Infrastructure policies, or
- is part of a recognised Computer Incident Response Team framework and as part of an incident investigation to prevent active or suspected misuse of Infrastructure services, or
- presents an appropriately enforced legal request.”

*Ian Neilson et al. based on the EU BCR model*

e.g. https://documents.egi.eu/document/669, and https://indico.egi.eu/indico/event/2923/contribution/2/material/0/

# IdP – SP proxies and credential translation services

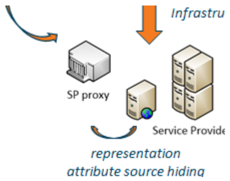

Community stepping in where IdP has left off

- adding collaboration and increasing assurance
- also *hides* back-end IdP assurance

Roles

- when in the attribute flow, becomes data processor
- otherwise it’s a conventional filter service (members) with a set of *membership policies* and *registration practices*

Communities come in various shapes & sized

- Highly structured, like WLCG
- Very loose small communities

*You need differentiates assurance at the SP end to filter!*

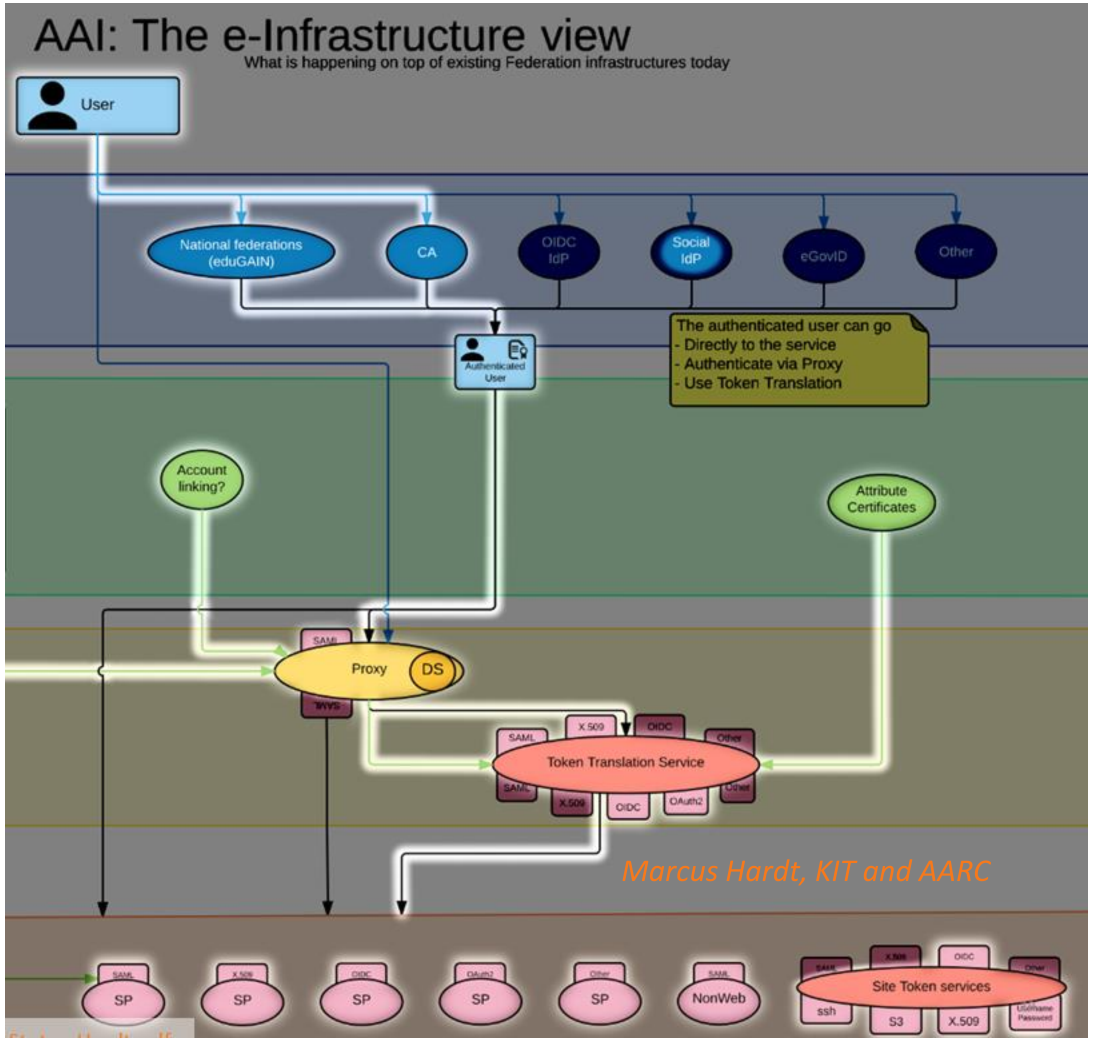
AAI: The e-Infrastructure view

What is happening on top of existing Federation infrastructures today

The authenticated user can go
- Directly to the service
- Authenticate via Proxy
- Use Token Translation

Marcus Hardt, KIT and AARC

https://aarc-project.eu/wp-content/uploads/2015/11/20151102-JRA1.2-Blueprint-Status-Hardt.pdf

# Improving scalability and avoiding silos

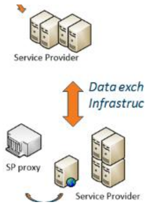

- When possible, ask common things as early as possible, and/or in a common way
- Otherwise you create a silo at the point of asking
- Common AUP is hard, comparable AUP works even across widely different infrastructures

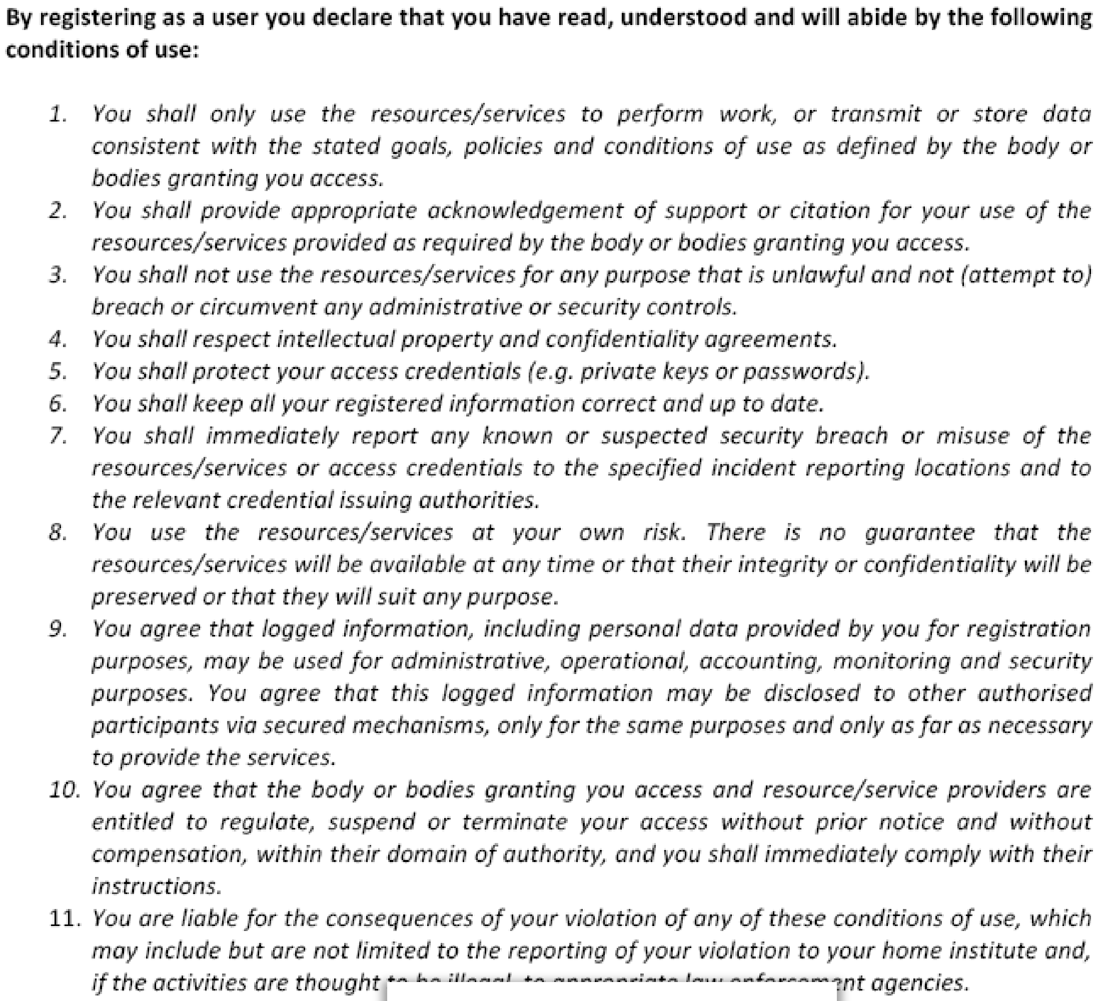

**By registering as a user you declare that you have read, understood and will abide by the following conditions of use:**

1. *You shall only use the resources/services to perform work, or transmit or store data consistent with the stated goals, policies and conditions of use as defined by the body or bodies granting you access.*
2. *You shall provide appropriate acknowledgement of support or citation for your use of the resources/services provided as required by the body or bodies granting you access.*
3. *You shall not use the resources/services for any purpose that is unlawful and not (attempt to) breach or circumvent any administrative or security controls.*
4. *You shall respect intellectual property and confidentiality agreements.*
5. *You shall protect your access credentials (e.g. private keys or passwords).*
6. *You shall keep all your registered information correct and up to date.*
7. *You shall immediately report any known or suspected security breach or misuse of the resources/services or access credentials to the specified incident reporting locations and to the relevant credential issuing authorities.*
8. *You use the resources/services at your own risk. There is no guarantee that the resources/services will be available at any time or that their integrity or confidentiality will be preserved or that they will suit any purpose.*
9. *You agree that logged information, including personal data provided by you for registration purposes, may be used for administrative, operational, accounting, monitoring and security purposes. You agree that this logged information may be disclosed to other authorised participants via secured mechanisms, only for the same purposes and only as far as necessary to provide the services.*
10. *You agree that the body or bodies granting you access and resource/service providers are entitled to regulate, suspend or terminate your access without prior notice and without compensation, within their domain of authority, and you shall immediately comply with their instructions.*
11. *You are liable for the consequences of your violation of any of these conditions of use, which may include but are not limited to the reporting of your violation to your home institute and, if the activities are thought to be illegal, to appropriate law enforcement agencies.*

**'Taipei Accord', 2005**

https://documents.egi.eu/document/2623

# Collaborating in incident response among the AAI chain

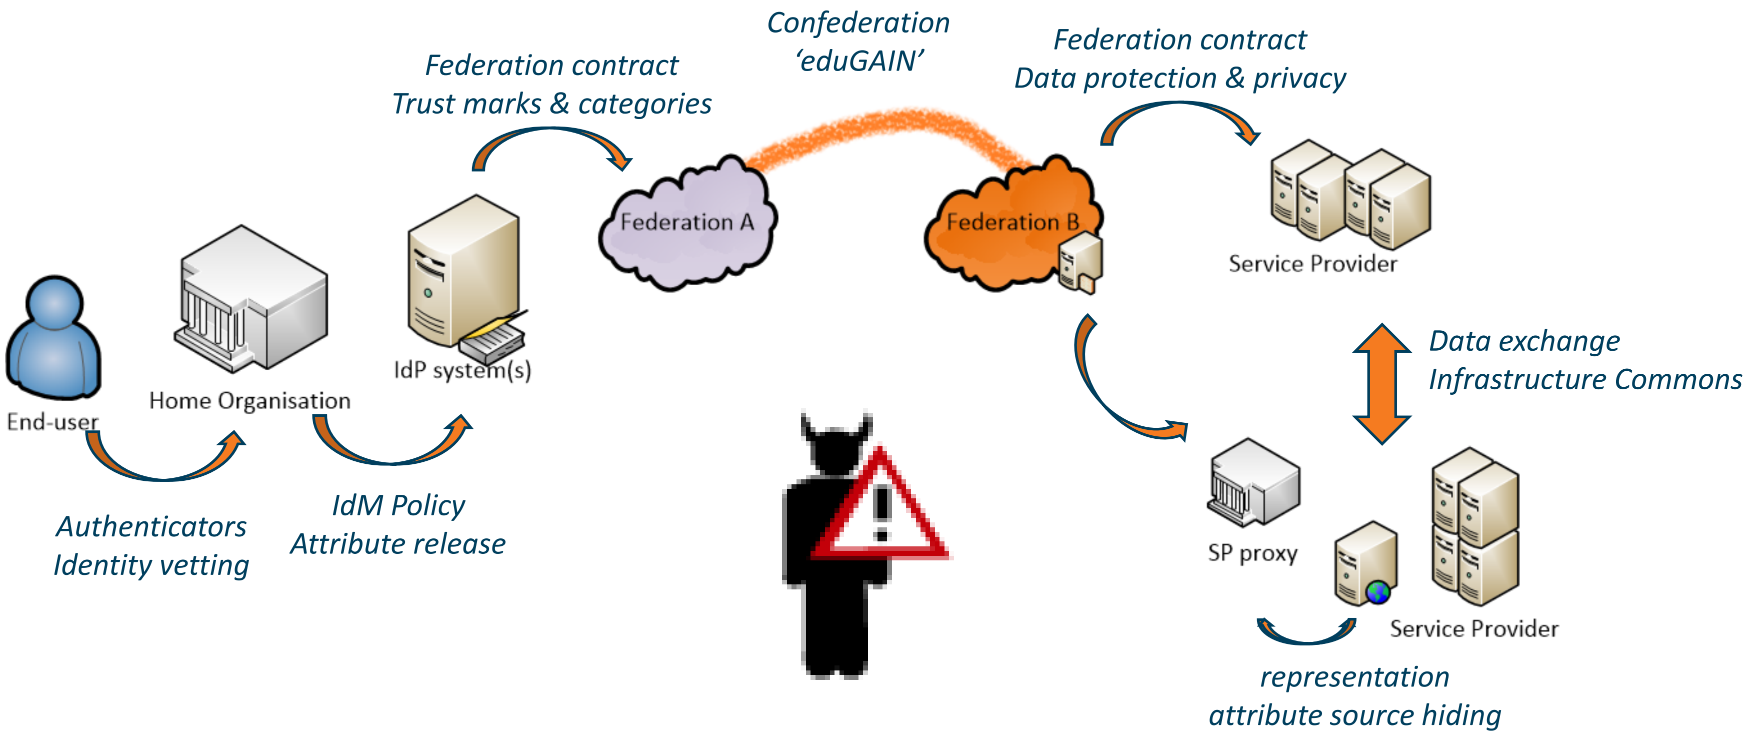

## Sirtfi – Security Incident Response Trust Framework for Federated Identity

# Sirtfi Framework for Incident Response Collaboration

Asserting Sirtfi compliance by an IdP or SP (e.g. via SAML meta-data) means concretely

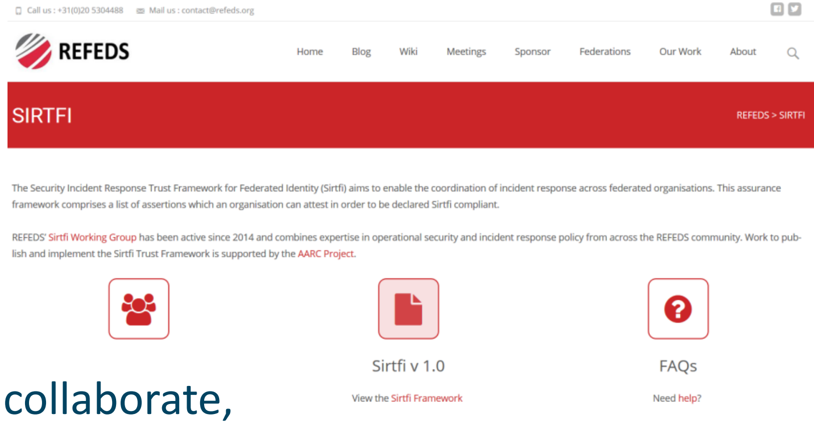

The Security Incident Response Trust Framework for Federated Identity (Sirtfi) aims to enable the coordination of incident response across federated organisations. This assurance framework comprises a list of assertions which an organisation can attest in order to be declared Sirtfi compliant.

REFEDS' Sirtfi Working Group has been active since 2014 and combines expertise in operational security and incident response policy from across the REFEDS community. Work to publish and implement the Sirtfi Trust Framework is supported by the AARC Project.

Sirtfi v 1.0

View the Sirtfi Framework

FAQs

Need help?

- Operational Security
  - Managed vulnerability patching and processes, some intrusion detection, account management, contact information available, CSIRT established
- Incident Response
  - Provide contact information, timely response to assistance requests, willing to collaborate, respect privacy and TLP classifications
- Traceability
  - Logs are timestamped, retained, and available; and there is a policy governing its retention
- Participant Responsibilities
  - There is an AUP for participants, users are made aware of it and agree to abide by it

*Some, not all, federation agreements actually cover some of this (e.g. an AUP in SURFconext)*

https://refeds.org/SIRTFI

# Sirtfi – the road ahead!

- Phase I: develop the SIRTFI Trust Framework specification, defining basic security incident response capabilities - organizations can self-assert compliance
  - Define meta-data schema for security contacts
  - Entity category definition and registration

- Phase II: means for member organisations in R&E federations to indicate their compliance
  - Training and communication
  - Implement contact and category definitions in production federations (and keep it up to date)

- Phase III: means for proactive notification of account compromise along the chain
  - Automated mechanisms for incident notification in the federation chain
  - Look e.g. at initiatives like Confyrm – incidents mostly spread based on commons user across SPs and IdPs
  - Tooling is needed, but maintaining privacy is always a challenge …

# Sirtfi in practice – Security Contact Meta-data in federations

- Security Contact Data in the meta-data specification
- Extends InCommon contactType (and will likely be renamed):

```
<ContactPerson contactType="other"
xmlns:icmd="http://id.incommon.org/metadata"
icmd:contactType=
"http://id.incommon.org/metadata/contactType/security">
```

Who to expect there?

- Somebody in the entity organisation security team who knows about TLP and confidentiality
- Promptly (within one business day) acknowledge receipt of the security incident report.
- As soon as circumstances allow, investigate incident reports regarding resources, services, or identities for which they are responsible.
- Respond to the incident reporter and any other impacted parties when the incident is resolved.

# Along the Chain

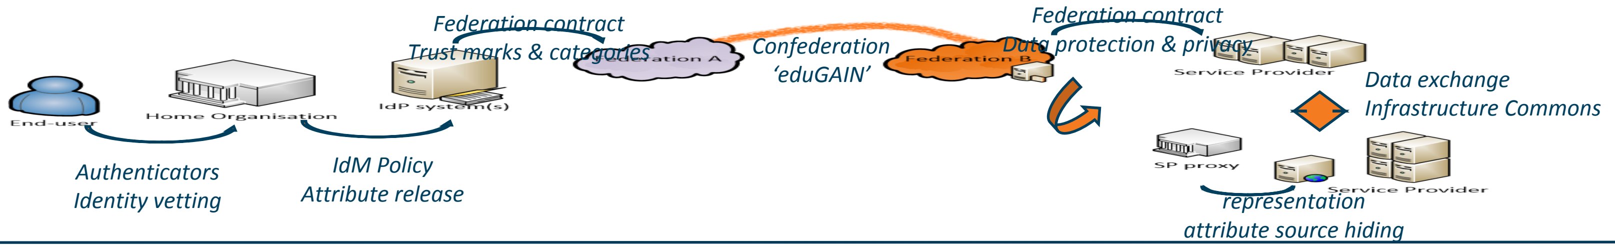

Federation contract
Trust marks & categories

Confederation
'eduGAIN'

Federation contract
Data protection & privacy

Data exchange
Infrastructure Commons

Authenticators
Identity vetting

IdM Policy
Attribute release

representation
attribute source hiding

- **Levels of Authentication Assurance**: common baseline being drafted. Usually you get reasonable entities, protected with a username-password, and data not older than a few years
- **IdP maturity level** is generally low (undocumented), with exceptions. Some federations include a minimum bar, but this will differ per country.
  At least REFEDS R&S gives you non-reassigned identifiers - but you're not sure which one.
- Assignment of **entity categories by IdP (SirTFi) and federation (REFEDS R&S)** enhances trust
- **eduGAIN** is great for finding meta-data and some contacts (e.g. in the MET tool), but
  no common policy apart from being "something academiccy" and a registrar name.
  Remember that it's nothing more – and that real discussions are in closed federation-only groups
- User attributes and data **sharing between service providers** may involve regulatory issues
- When there's a science gateway involved, VO, or **IdP-proxy, that's the place to get information**
- Security contact information and collaboration is **evolving through Sirtfi**

https://www.nikhef.nl/~davidg/presentations/ISGC-SecWS-2016-Assurance-and-Federation-AARC.ppdf

# Thank you

# Any Questions?

davidg@nikhef.nl

AARC

https://aarc-project.eu

The work leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 653965 (AARC).