Federated services: spanning countries and crossing borders

Building across the technology chain –
when networks and systems,
trust and identity combine

Maastricht University | Department of Advanced Computing Sciences

We live in a federated world

Collaboration: an inherently-cross-domain issue ...

Example from the LHC Computing infrastructure WLCG

170 sites
~50 countries & regions
~20000 users

just *how* many interactions ??

AuthN & AuthZ, **architecture and trust** should align with **collaboration structures**, and be **outward facing:** open, scalable, & multi-domain

people photo: a small part of the CMS collaboration in 2017, Credit: CMS-PHO-PUBLIC-2017-004-3; site map: WLCG sites from Maarten Litmaath (CERN) 2021

AARC BPA: The Blueprint Architecture

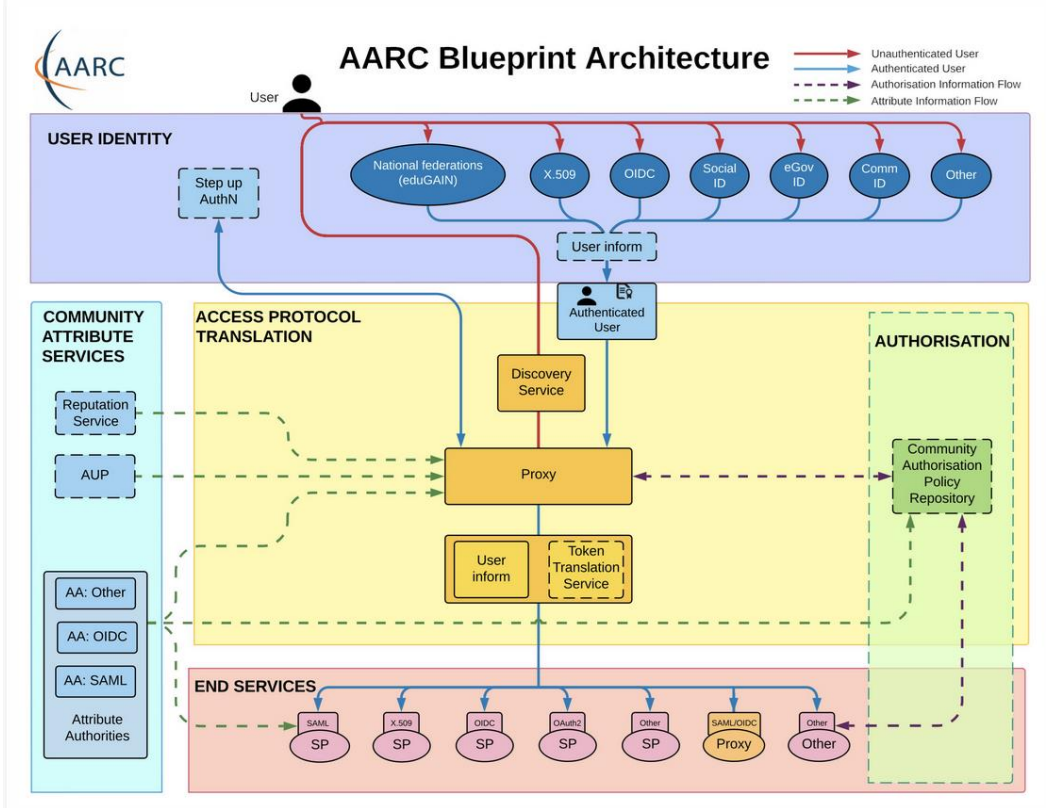

AARC BPA: proxies as first-class citizens of T&I space

- Access services using **identities from users' Home Organizations**, but **hide complexity** of multiple IdPs, federations, AA technologies
- O**ne persistent identity** across all the community's services through **account linking**
- **Access** services **based on role(s)** users have **in the collaboration**.
- For both **web** and **non-web** resources
- Integration of **guest identity solutions**
- **Support for stronger authentication assurance** mechanisms

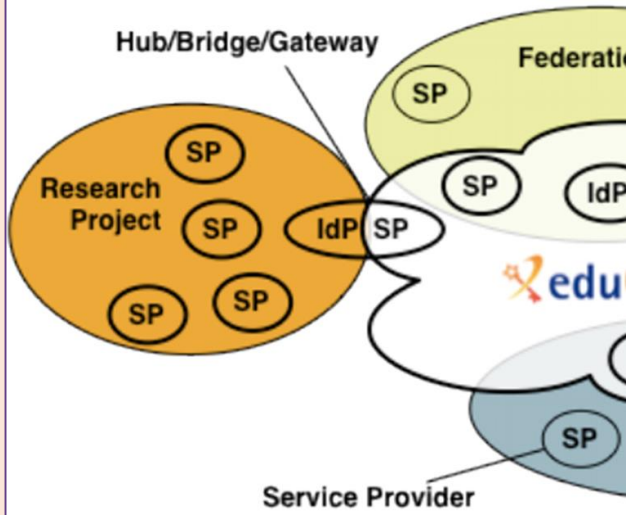

Graphics: Ann Harding and Lukas Hammerle (SWITCH)– from a long time ago now!

Authentication and Authorization for Research Collaboration – https://aarc-community.org/

Live applications of the AARC BPA

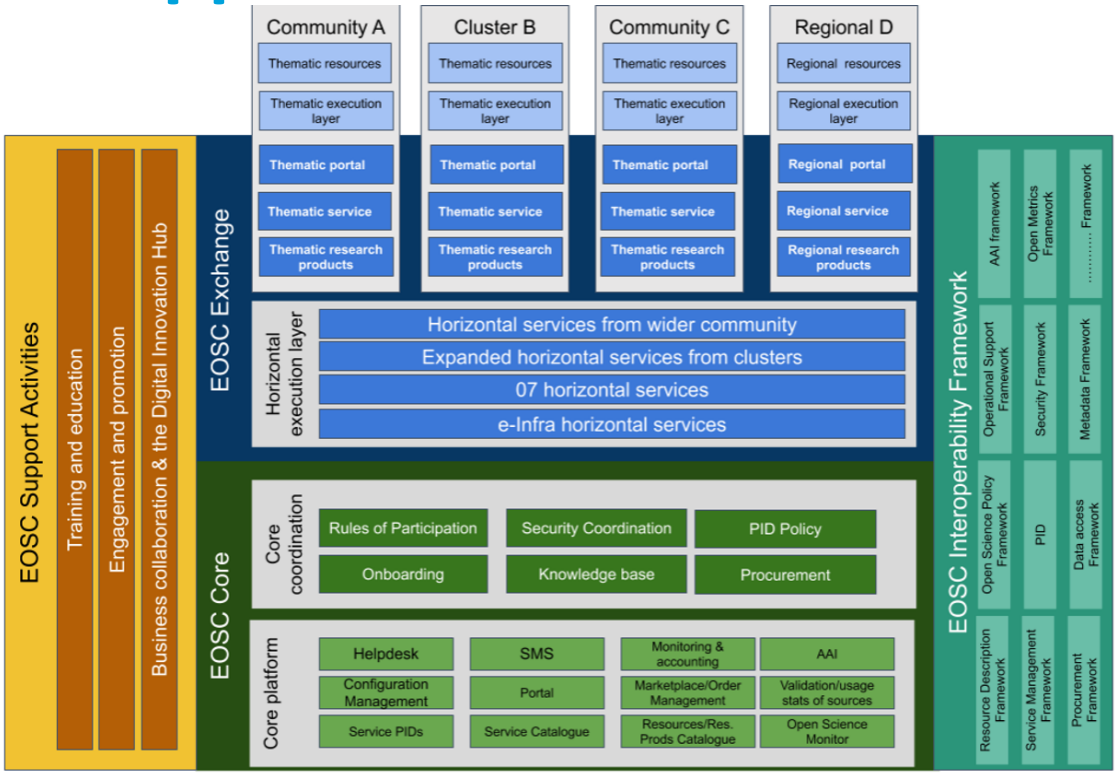

and many more systems and ‘data spaces’ besides EOSC: *e.g.* Copernicus EO data, GAIA-X, sectoral spaces, …

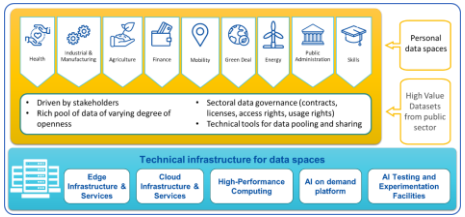

EOSC: https://eoscfuture.eu/wp-content/uploads/2022/04/EOSC-Core.pdf; data spaces image: https://digital-strategy.ec.europa.eu/en/library/building-data-economy-brochure

Where researchers find services & collaboration …

represented by logos: some of the (AARC BPA) Research Communities (top) providing federated access using the AAI proxy architecture.
At the ~ bottom: (global) e-Infrastructures, which all use the AARC BPA collaborative model

Distributed collaborative ICT instrumentation, *a more technical example*

Credential translation in the AARC BPA
… building RCauth.eu
Leveraging federation and collaboration for ubiquitous research credentials

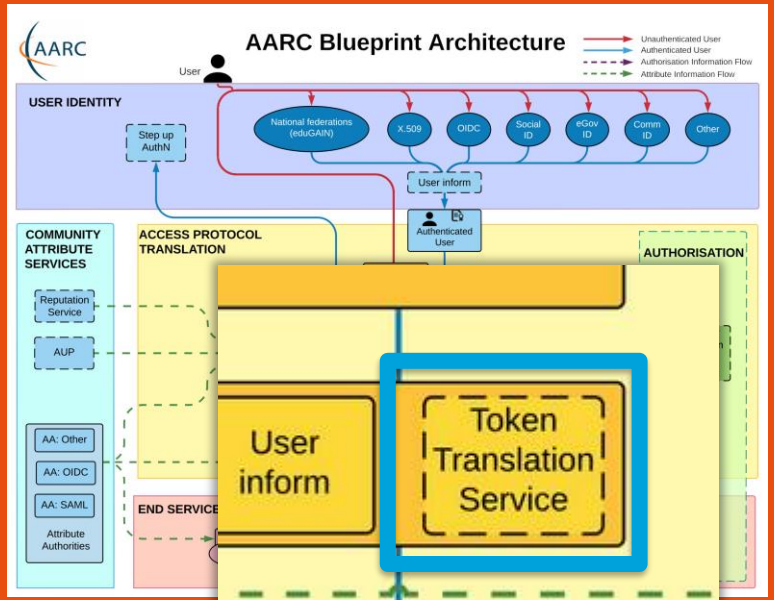

Bridges and Token Translation Services

TCS - for users that manage to grasp the idea

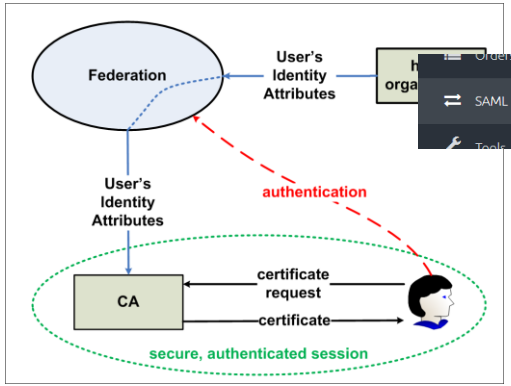

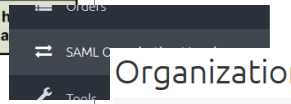

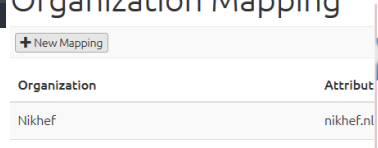

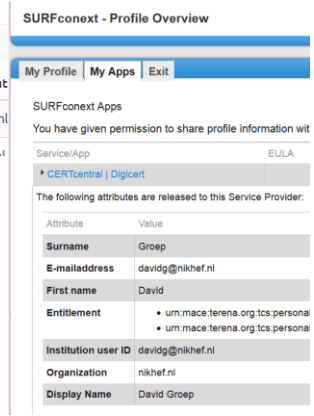

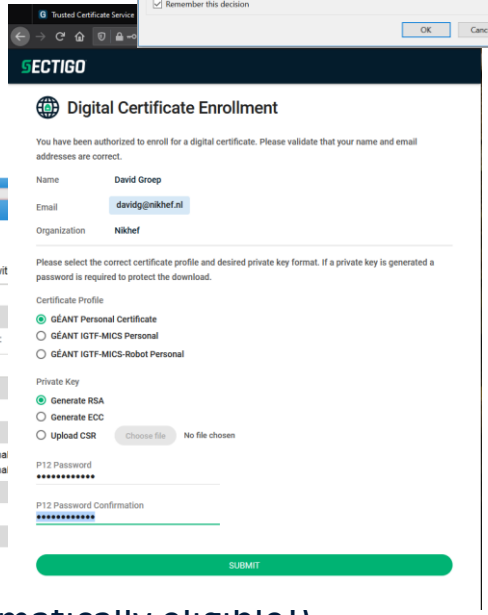

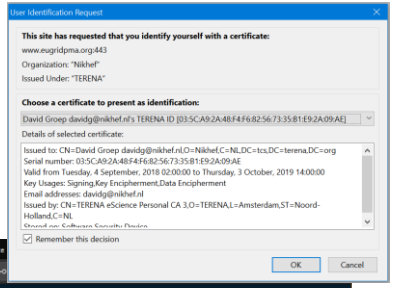

TCS is a SAML Service Provider (today by Sectigo)
to eduGAIN: where eligible authenticated users obtain client certificates for access to many research services
A globally recognized identity for all employees & students (they are automatically eligible!).

GEANT Trusted Certificate Service - https://ca.dutchgrid.nl/tcs/,
https://cert-manager.com/customer/surfnet/idp/clientgeant, https://www.geant.org/Services/Trust_identity_and_security/Pages/TCS.aspx

Ingredients for credential minting & token translation

eduGAIN (global R&E) *Entity Categories*	e-Infrastructure IGTF Authentication Profiles	Use of proxy bridging components
Curated grouping of entities 'REFEDS R&S' *this is a research service* 'DP CoCo' *abides by GDPR* 'Sirtfi' *cares for security response* slower adoption process adding identity assurance needs action at all 60+ Feds & 4k+ IdPs	Common baseline and profiles *co-defined by relying parties* user-centric ID harmonisation with unique global naming 'BIRCH' *real person with real name* 'DOGWOOD' *persistent linkable identifier* research-specific user base	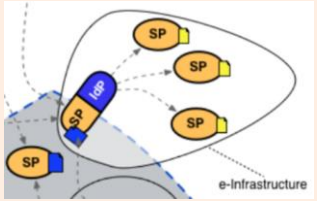 Identity and access 'proxy' harmonised eduGAIN IdPs *based on entity categories* *leverage Sirtfi and 'R&S'* *proxying is bi-directional* responsibility on the proxy operator

https://wiki.geant.org/display/AARC/Current+Status+of+SAML+Entity+Categories+Adoption - https://www.rcauth.eu/ - https://aarc-community.org/

Seamless in-line token translation services from ‘SAML’ to PKIX

user facing | hidden back-end

Community Science Portal

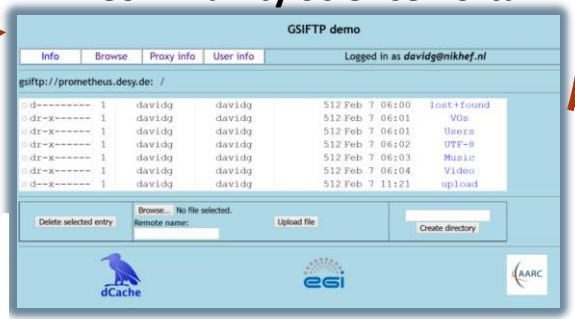

Infrastructure Master Portal Credential Store

IGTF accredited PKIX Authority

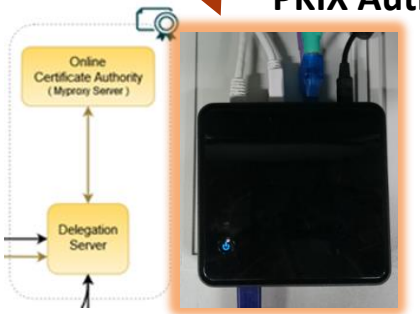

User Home Org
or Infrastructure IdP

see also https://rcdemo.nikhef.nl/

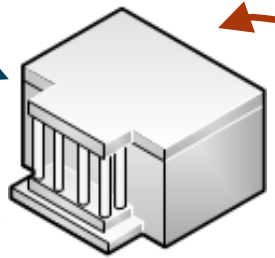

REFEDS R&S
Sirtfi Trust

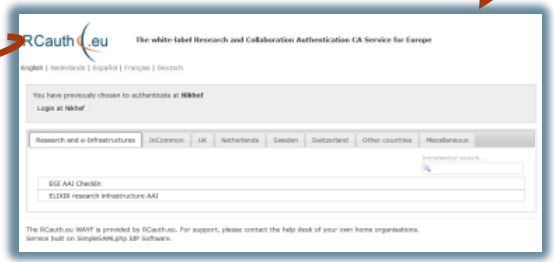

Policy Filtering WAYF to eduGAIN

Built on CILogon and MyProxy, see www.cilogon.org CILogon

Our Registration Authorities: the Federated IdPs

- Distributed RAs: the eligible IdPs
  - connected through a federation, primarily: the ensemble of IdPs in eduGAIN that meet the policy requirements of this CA
  - since authN and authZ are split, need is for non-reassigned identifier and point-in-time incident response
- eligible applicants are then all affiliated to an RA

Three eligibility models

- Direct relationship CA-IdP, with agreement declaration
- Rest of eduGAIN: – "Sirtfi" security incident response and OpSec capabilities plus
  – REFEDS "R&S section 6" non-reassigned identifiers & name ('personalized')
  are required, and tested via statement in 'meta-data' and by releasing the proper attributes
- within the Netherlands, SURFconext Annex IX* already ensures compliance for all IdPs
- "IdPs within eduGAIN are deemed to have entered materially into an agreement with the CA"

The ‘back side’ of a typical RCauth portal data flow

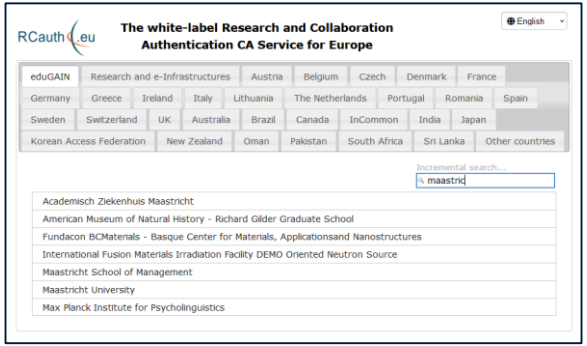

Parsed ID Token:

```
stdClass Object
(
    [typ] => JWT
    [kid] => E01796EA0367564935B0981731B9B116
    [alg] => RS256
)
stdClass Object
(
    [sub] => P70081609@unimaas.nl
    [idp] => http://login.maastrichtuniversity.nl/adfs/services/trust
    [eduPersonTargetedID] => http://login.maastrichtuniversity.nl/adfs
    [idp_display_name] => Maastricht University
    [cert_subject_dn] => CN=Groep\, David (DACS) KWwWAnhI4psmiGTw 1,O=
    [name] => Groep, David (DACS)
    [eduPersonPrincipalName] => P70081609@unimaas.nl
    [given_name] => David
    [family_name] => Groep
    [email] => david.groep@maastrichtuniversity.nl
    [iss] => https://aai.egi.eu/mp-oa2-server
```

Proxy information:

```
subject   : /DC=eu/DC=rcauth/DC=rcauth-clients/O=maastrichtuniversity.nl/CN=Groep, David (DACS) KWwWAnhI4psmiGTw 1/CN=208760481/CN=466908503
issuer    : /DC=eu/DC=rcauth/DC=rcauth-clients/O=maastrichtuniversity.nl/CN=Groep, David (DACS) KWwWAnhI4psmiGTw 1/CN=208760481
identity  : /DC=eu/DC=rcauth/DC=rcauth-clients/O=maastrichtuniversity.nl/CN=Groep, David (DACS) KWwWAnhI4psmiGTw 1/CN=208760481
type      : RFC compliant proxy
strength  : 2048 bits
path      : /tmp/x509up_uiI4TkF
```

With a single, yet fully compliant, ‘Heath Robinson’ CA

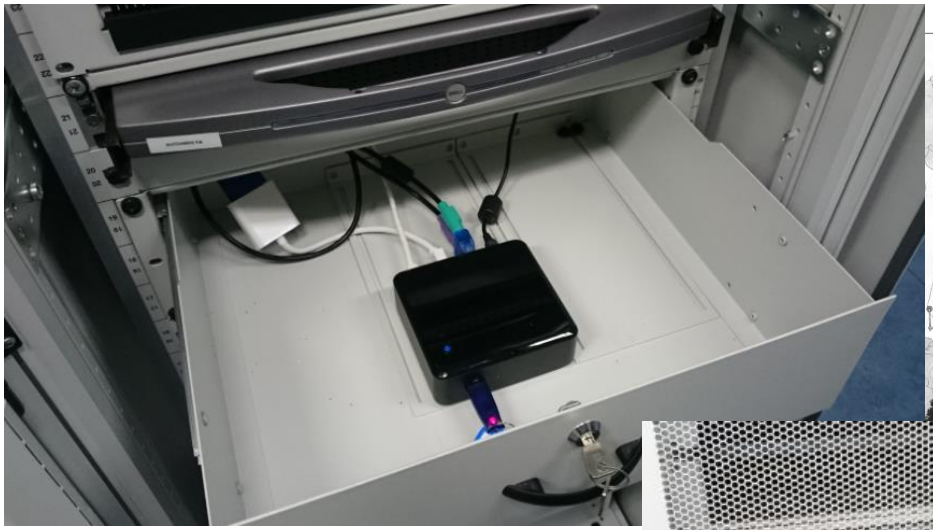

One ‘locally-highly-available’ RCauth at Nikhef Amsterdam

- Most ‘fault-prone’ components are
  - Intel NUC (single power supply)
  - HSM (can lock itself down, and the USB connection is prone to oxidation)
  - DS front-end servers (physical hardware, albeit with redundant disks and powersupplies)

Eliminated SPOFs first using ‘local HA’

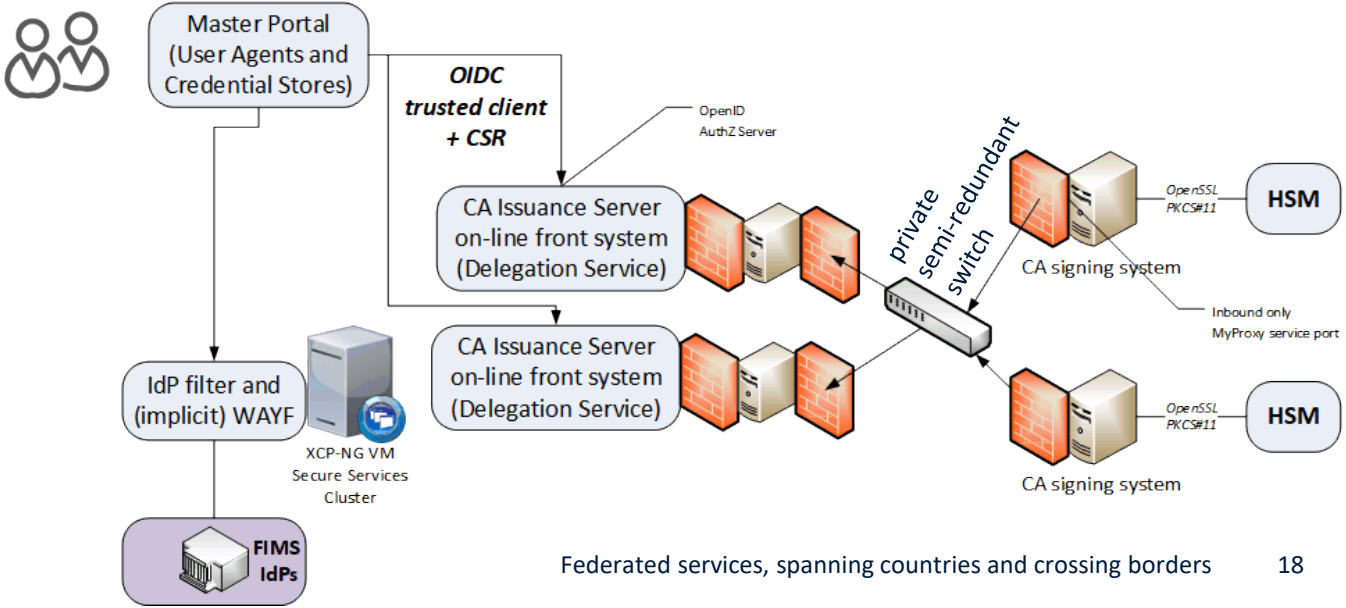

Distributing RCauth.eu across three cooperating sites

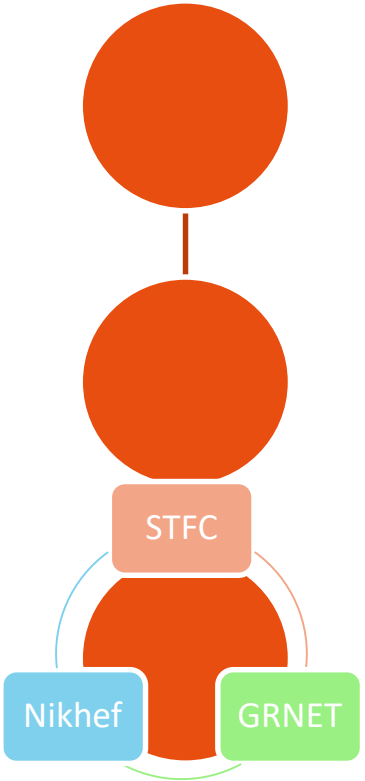

Governance Board

Representatives (and one alternate) from each Materially Contributing Stakeholder EGI.eu, EUDAT (ETFC), GÉANT, Nikhef (SURF)

RCauth PMA

Individuals drawn from the wide community […] experts in the field of identity management for research and collaboration, PKI technology and identity bridging

Ops Coordination Team

Operations people from each of the **hosting partners** with a (copy of) the RCauth.eu signing key, and those partners otherwise involved in OPS

… to a 3-fold, continuously-consistent, European setup

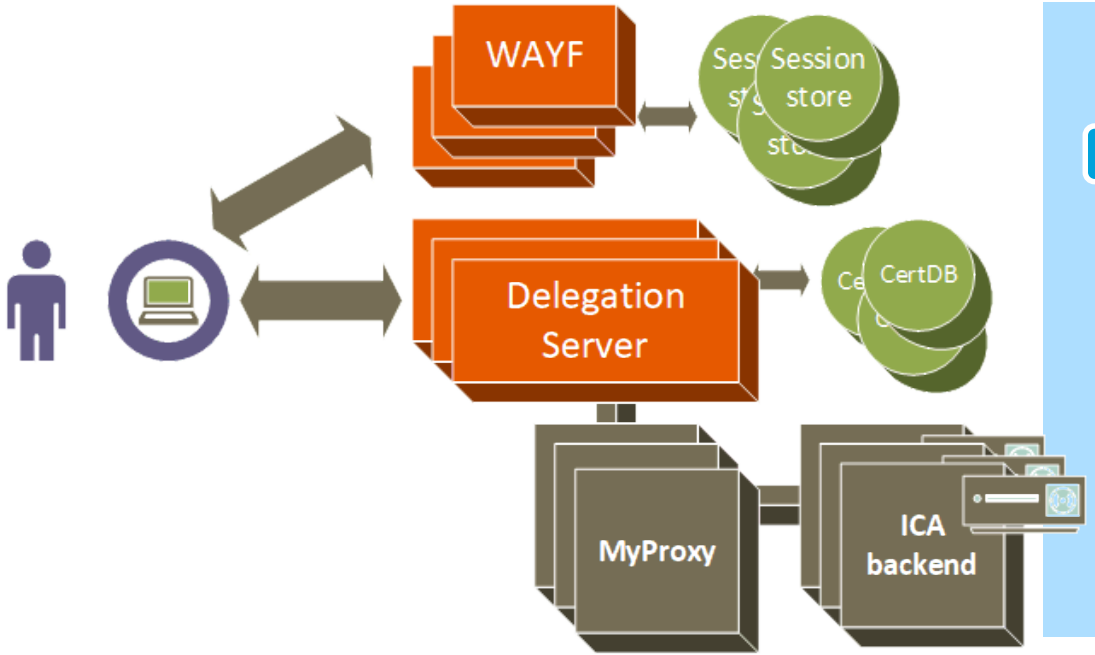

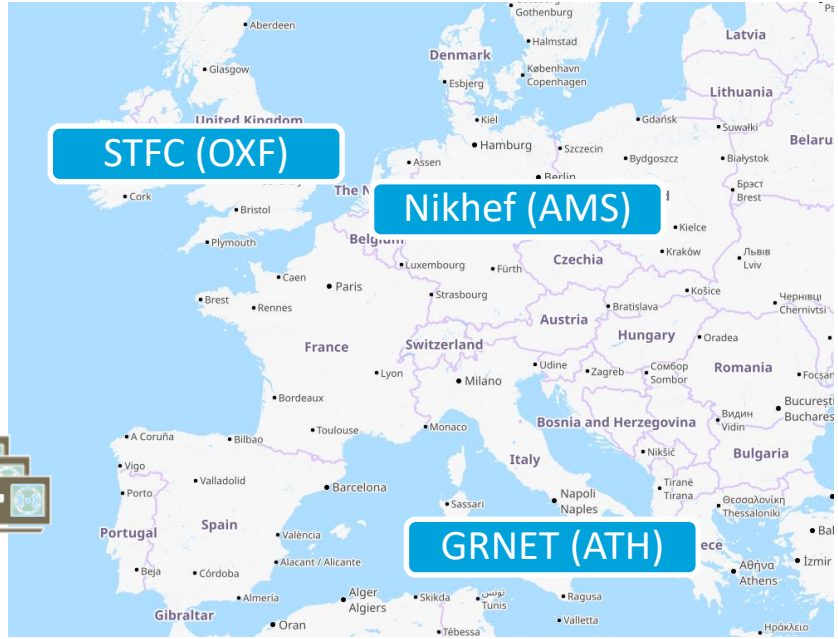

Since we do not like SPOFs …

- Distributed High Availability setup

- across the 3 sites
- design for minimal effort
- readily-available techniques
  - L3 VPN (OpenVPN) or L2 VPC
  - Linux HAProxy

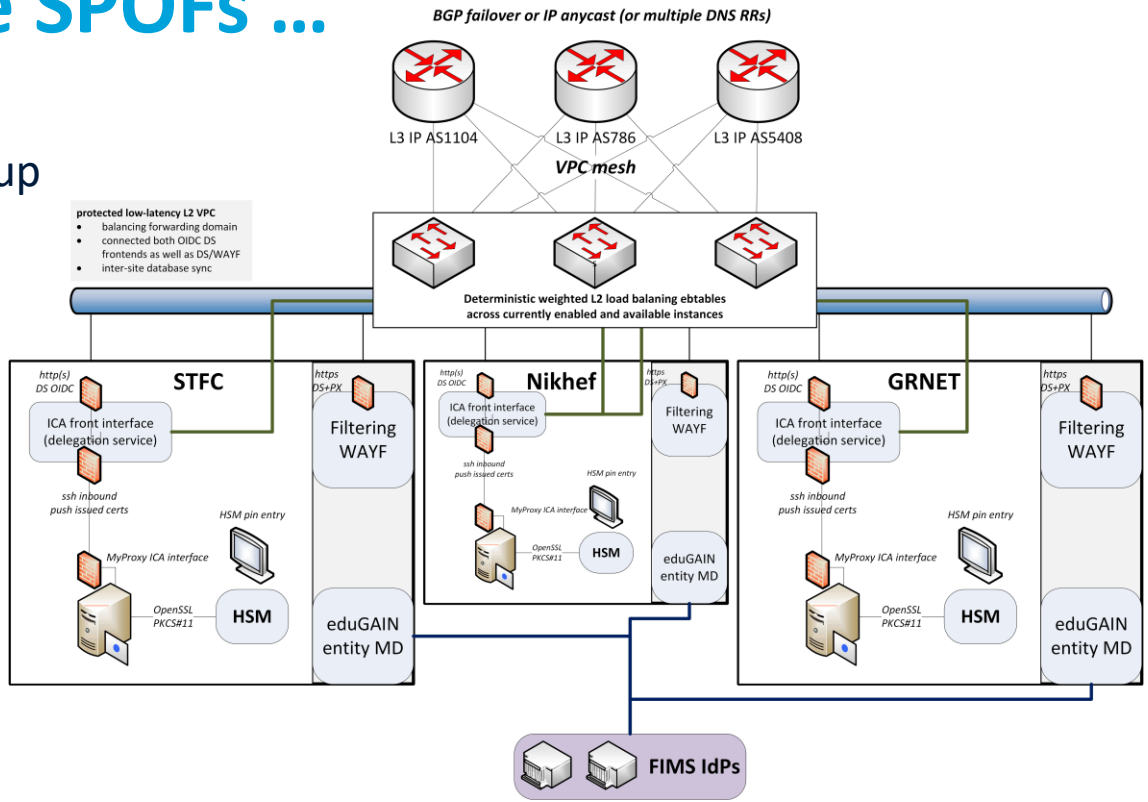

work supported by the EOSC Hub and EOSC Future Horizon Europe projects

A transparent multi-site setup is needed for the user

- User
- connects to HA proxy at {wayf,pilot-ica-g1}.rca
- HA proxy sends users to “closest” working ser
- primarily forward to its own DS when available

If a HA loses its backend DS, can still route to another DS over VPC/VPN backend

Straightforward proven solution is IP anycast

wherever the user is, the service is at

- **2a07:8504:01a0::1**
- or for legacy IP users at 145.116.216.1

selected imagery: Mischa Sallé, Jens Jensen, Nicolas Liampotis

Anycast: when the same place exists many times

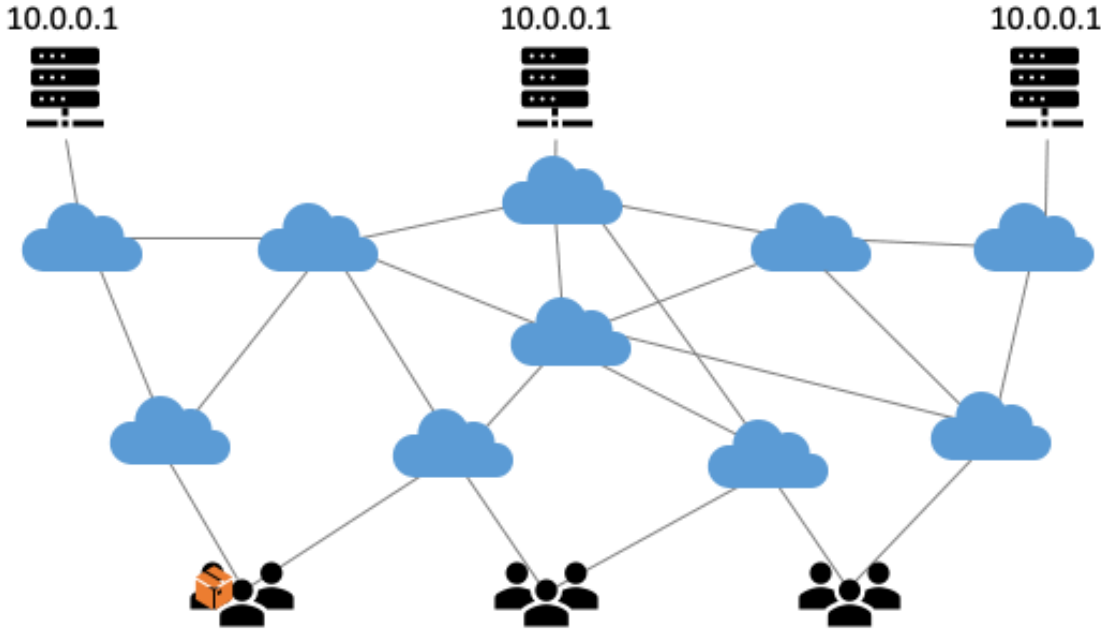

So we used

- 3 (for now: 2) sites
- one VM at each site exposing 2a07:8504:01a0::1
- smallest v6 subnet (/48)
- bird + a service probe
- each site's own ASN
- some IRR DB editing
- IPv4 is similar, with a /24

and some monitoring

routing image: SIDNlabs - https://www.sidnlabs.nl/en/news-and-blogs/the-bgp-tuner-intuitive-management-applied-to-dns-anycast-infrastructure

Getting 2a07:8504:1a0::/48 out there

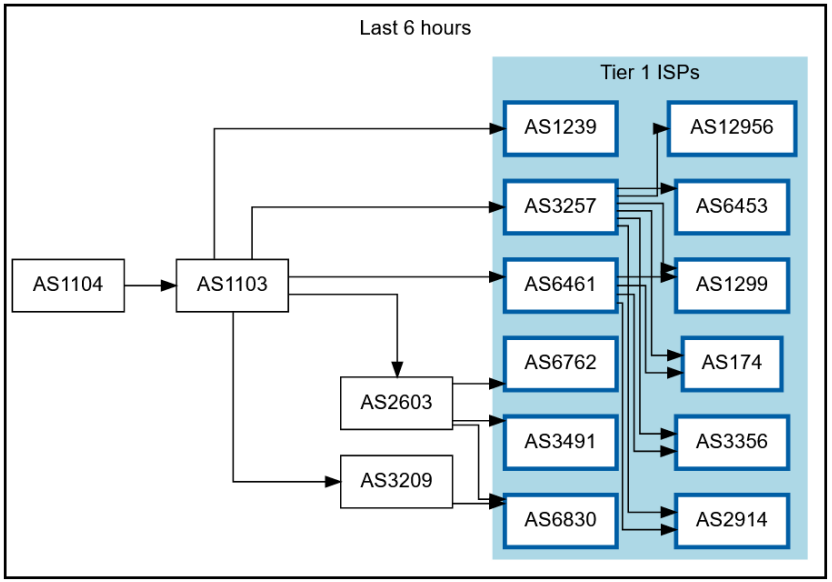

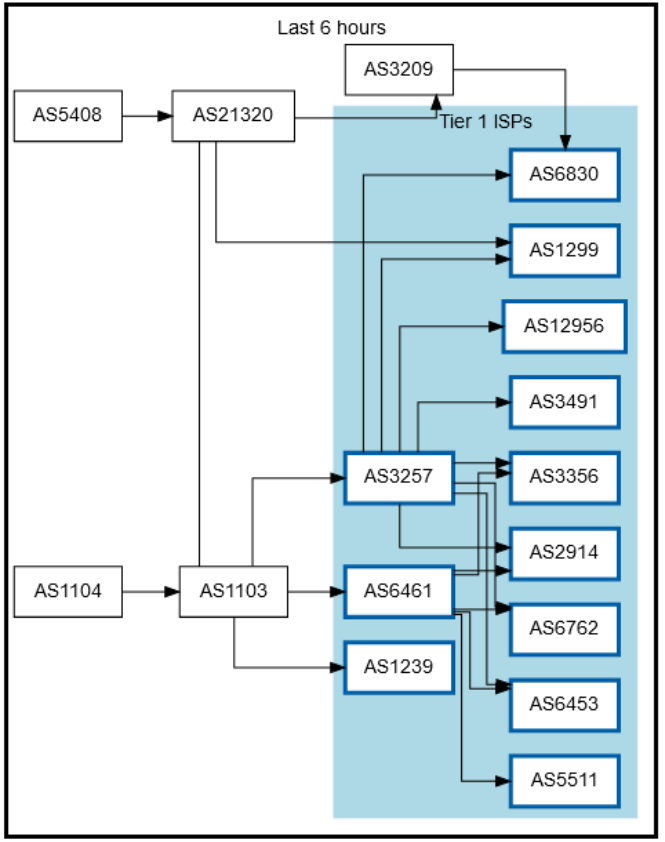

route maps: bgp.tools for 2a07:8504:1a0::/48 – IPv4 for 145.116.216.0/24 is similar – imagery from November 2022

‘Birds flock together’

```
# /etc/bird.conf route config
define ASN_OWN          = 65530;
define ASN_NEIGHBOUR    = 1104;
define ADDR_NEIGHBOUR4 = 194.171.98.94;
define ADDR_NEIGHBOUR6 = 2a07:8500:120:e011::1;
```

```
# /etc/bird.d/anycast-prefixes.conf
# …
# Generated 2023-02-05 14:49:36.047801
#  by anycast-healthchecker (pid=1299)
# 10.189.200.255/32 is a dummy IP Prefix. …
define ACAST_PS_ADVERTISE =
    [
        10.189.200.255/32,
        145.116.216.1/32
    ];
```

```
# /etc/anycast-healthchecker.d/haproxy.conf
[haproxy]
check_cmd      = /usr/local/sbin/check_haproxy.sh
check_interval = 20
check_timeout  = 5
check_fail     = 2
check_rise     = 2
check_disabled = false
on_disabled    = withdraw
ip_prefix      = 145.116.216.1/32

[haproxy6]
check_cmd      = /usr/local/sbin/check_haproxy.sh
check_interval = 20
check_timeout  = 5
check_fail     = 2
check_rise     = 2
check_disabled = false
on_disabled    = withdraw
ip_prefix      = 2a07:8504:1a0::1/128
```

Bird Internet Routing Daemon http://bird.network.cz/; for anycast-healthchecker: anycast-healthchecker-0.9.2.dev9-1.noarch Pavlos Parissis, Mischa Sallé

And you get reasonable load balancing in Europe for free

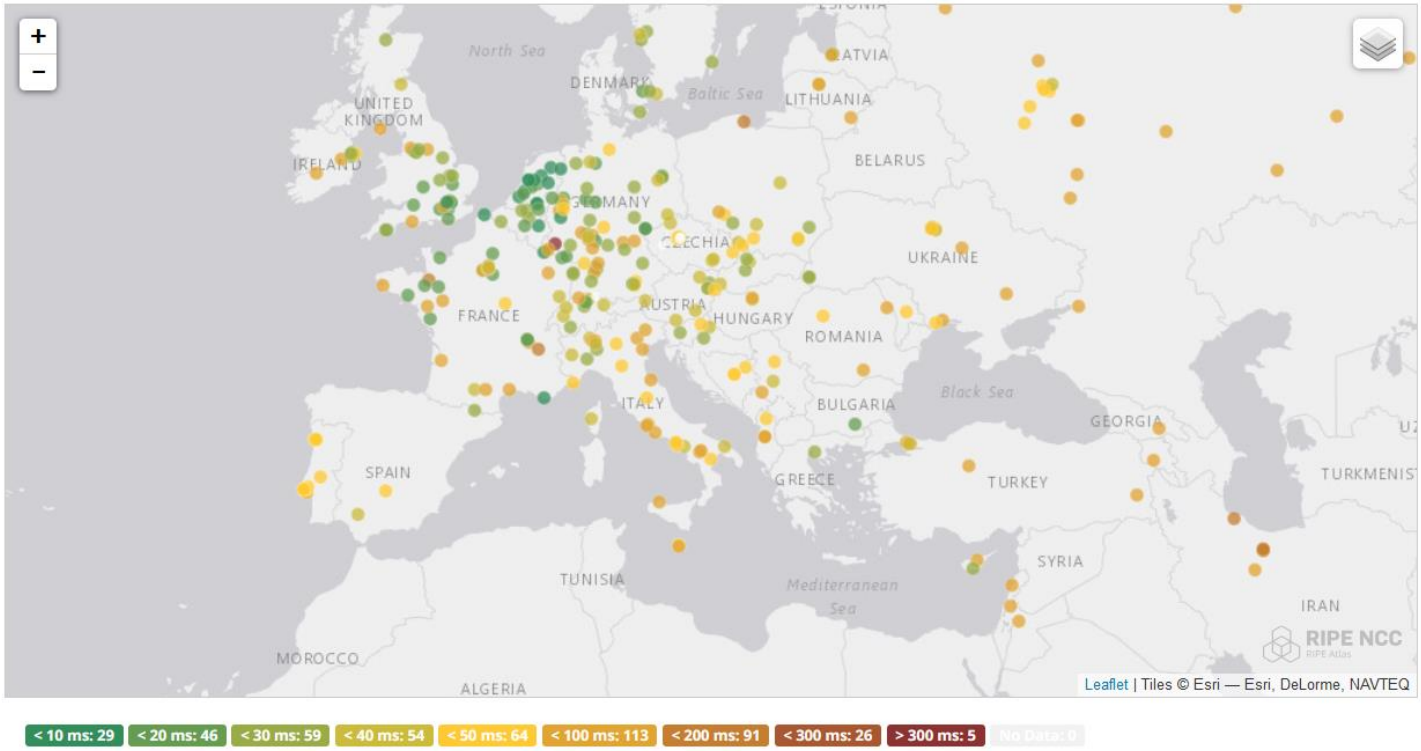

map: RIPE NCC RIPE Atlas - 500 probes, distributed across Europe (https://atlas.ripe.net/measurements/50949024/)

Shortest path, also when mixing with the default-free zone

```
[root@kwark ~]# traceroute -IA 145.116.216.1
traceroute to 145.116.216.1 (145.116.216.1), 30 hops max, 60 byte packets
 1  lo0-3.bras0.fi001.nl.freedomnet.nl
      (185.93.175.232) [AS206238]
 2  et-0-0-2-1001.core0.fi001.freedomnet.nl
      (185.93.175.223) [AS206238]
 3  as1104.frys-ix.net (185.1.203.182) [*]
 4  protnet-gw.nikhef.nl
      (194.171.98.94) [AS1104]
 5  gw-anyc-01.rcauth.eu
      (145.116.216.1) [AS786/AS5408/AS1104]
```

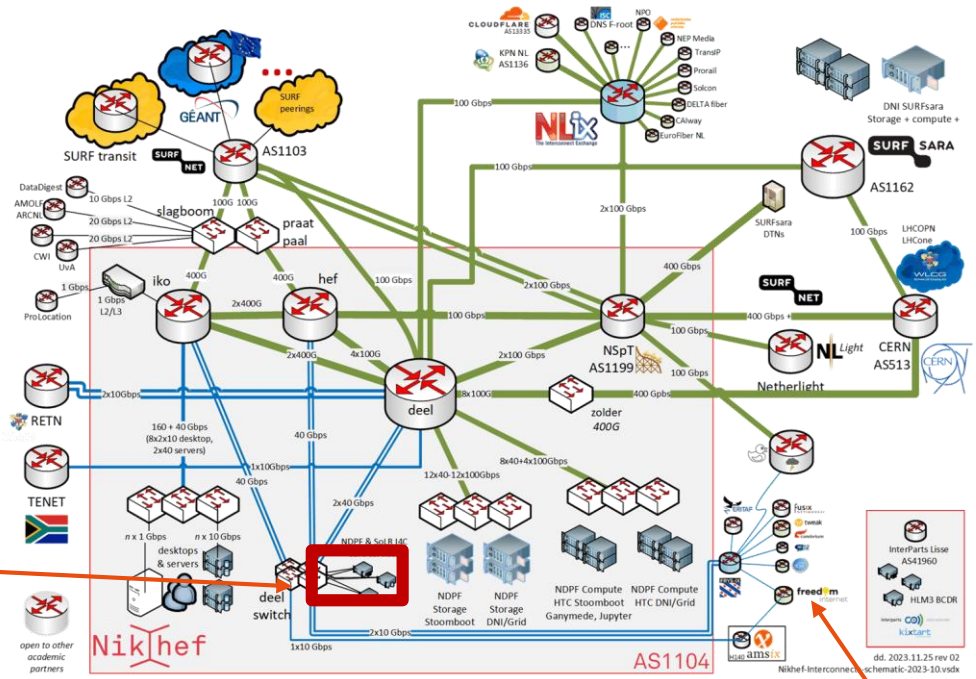

Route from home to RCauth.eu, from my home Freedom Internet ISP

So is there a common pattern?

- Infrastructure may be distributed, but that's nothing truly 'magic'
  - and every collaborating organization, university, and national lab is part of it and can do it!
- Move complexity and volume requirements to the edge
  - the edge scales horizontally and scaling from 2+ is much easier than from 1→ 2
- Any central (network) components should be passive and as stateless as possible
  - research (and computing education) infrastructure performance ought to just be 'a given'
  - any stateful device in the data path will block performant data transfers and reliability
  - although persistent storage obviously has to retain some state ☺
- Scaling collaboration infrastructure, trust & identity, and federation of expertise needed as much as we need scaling of our computing and networks

David Groep
david.groep@maastrichtuniversity.nl
https://www.nikhef.nl/~davidg/presentations/
https://orcid.org/0000-0003-1026-6606

Nikhef

Maastricht University | Department of Advanced Computing Sciences