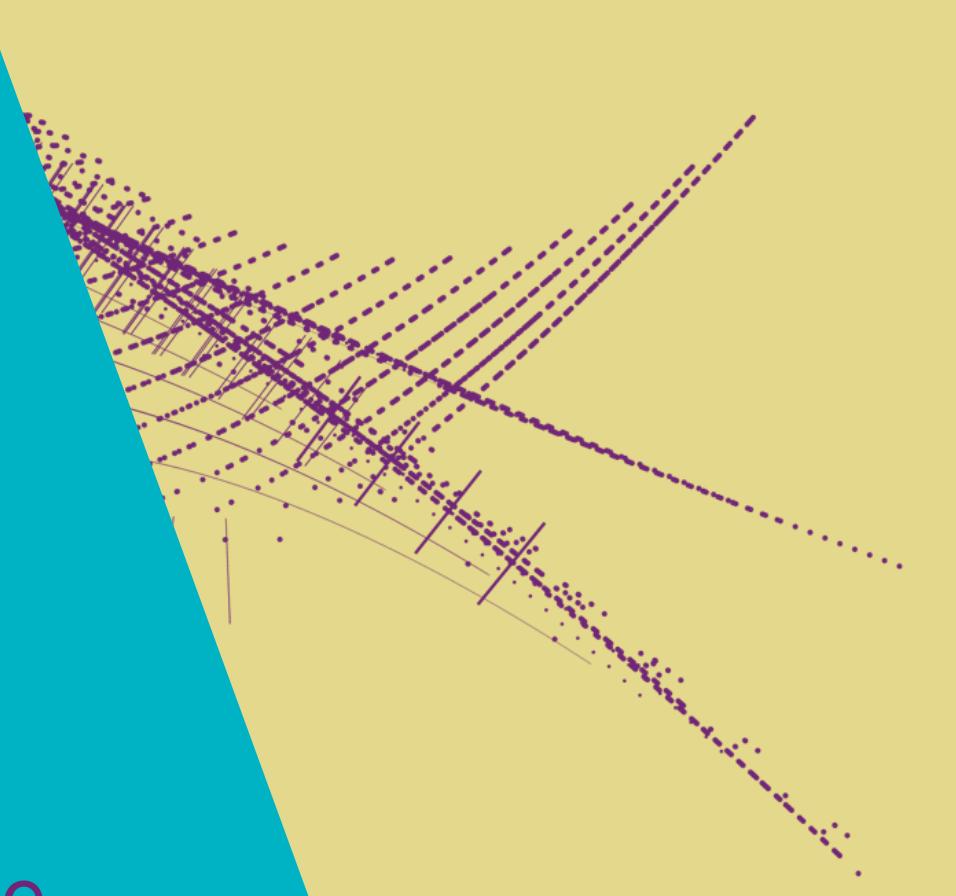


RWTH colloquium

Infrastructure:
for the small and the large
build, trust, collaborate, govern



David Groep,
February 2026

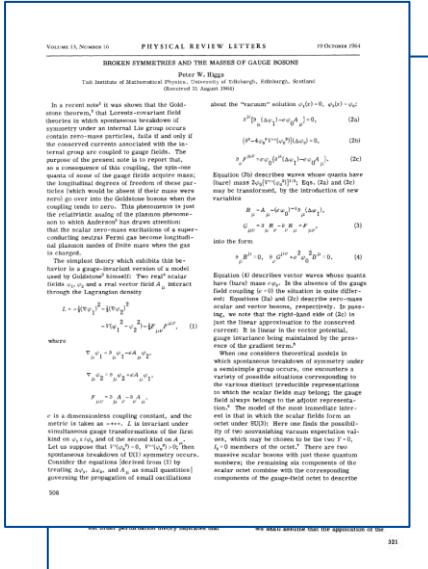


INFRASTRUCTURE: FOR THE SMALL AND THE LARGE

Peter Higgs and Francois Englert at the 2013 Nobel prize press conference, Stockholm. Photo: Bengt Nyman, <https://www.flickr.com/photos/97469566@N00>

Exploding data? the Large Hadron Collider at CERN

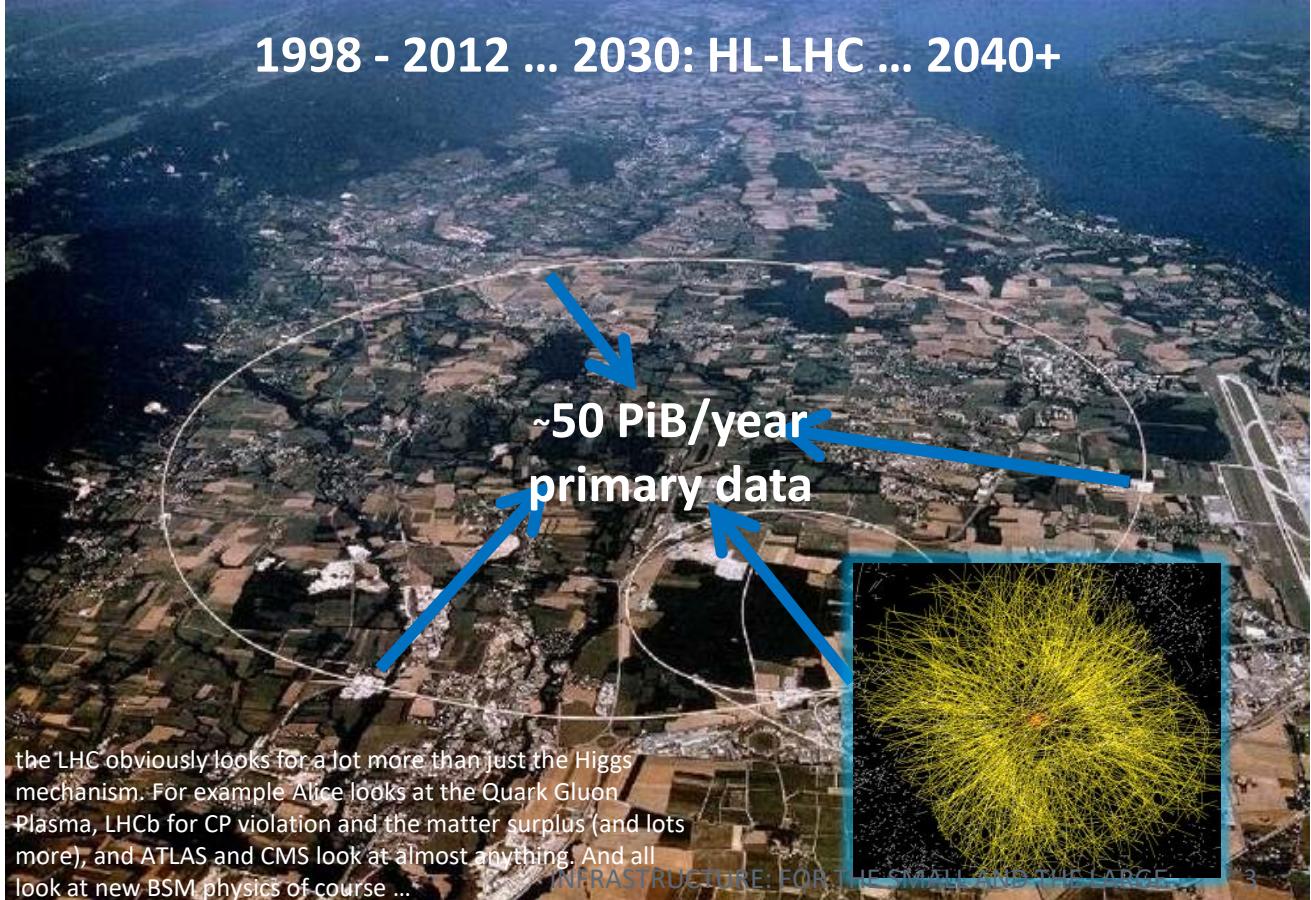
1964



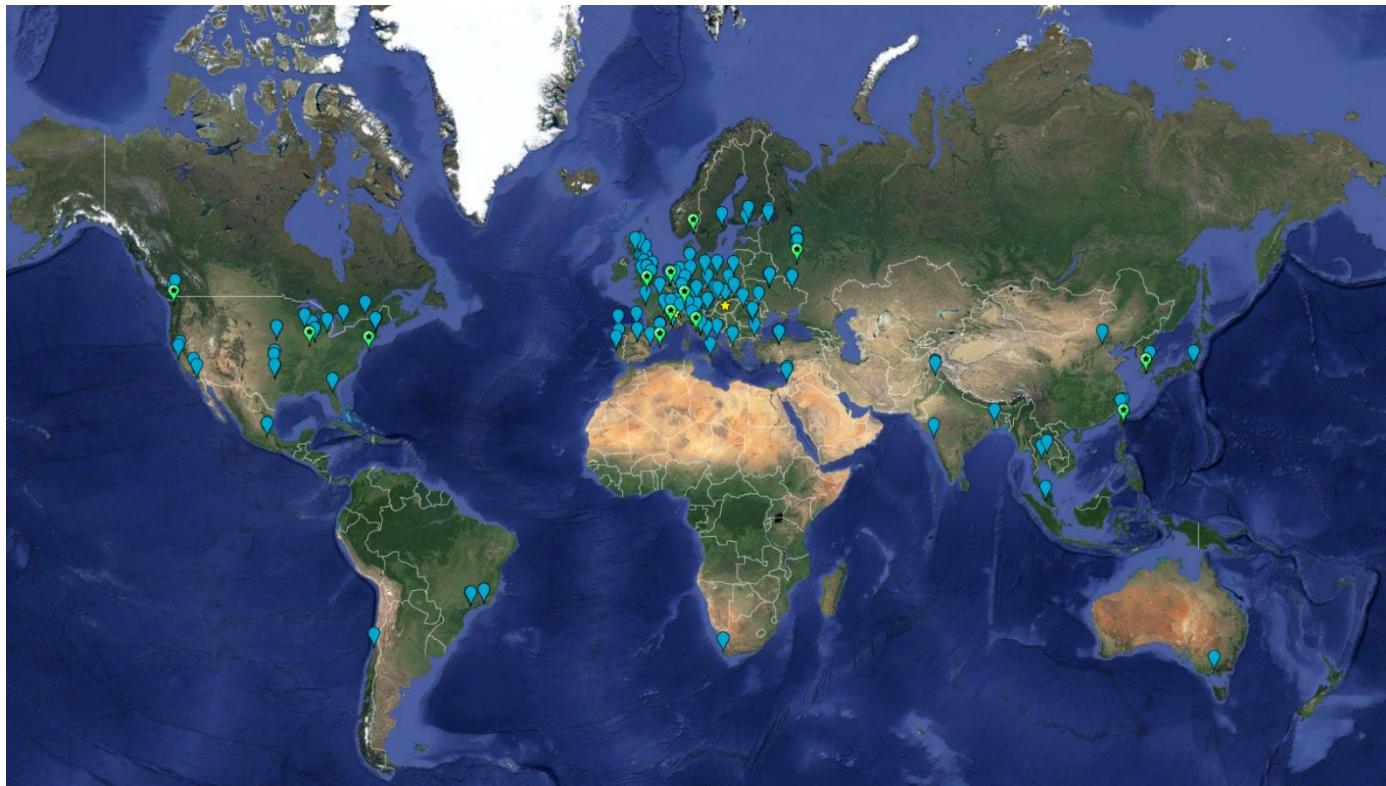
P. Higgs, Phys. Rev. Lett. 13, 508:

16823 characters, 165 kByte PDF

1998 - 2012 ... 2030: HL-LHC ... 2040+



Networked complexity: the worldwide LHC Computing



~ 1.6 million CPU cores
~ 2000 Petabyte disk + archival

160+ institutes
40+ countries
13 'Tier-1 sites'

NL-T1:
SURF & Nikhef

largely based on generic e-Infrastructures
EGI
EuroHPC
NEIC
OpenScienceGrid
ACCESS-CI

Earth background: Google Earth; Data and compute animation: STFC RAL for WLCG and EGI.eu; Data: <https://home.cern/science/computing/grid>

For the LHC Computing Grid: wlcg.web.cern.ch, for EGI: www.egi.eu; NSF ACCESS-CI: <https://access-ci.org/>, for the NL-T1 and FuSE: fuse-infra.nl, <https://www.surf.nl/en/research-it-infrastructure>: for the small and the large

Volume and computational complexity



ATLAS RAW single event
ROD File
1.60 MB

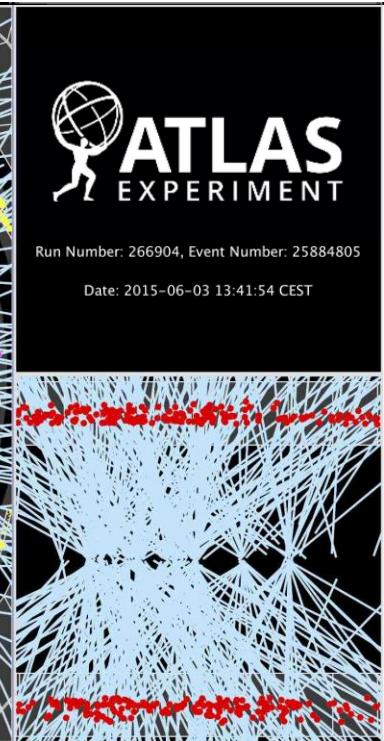
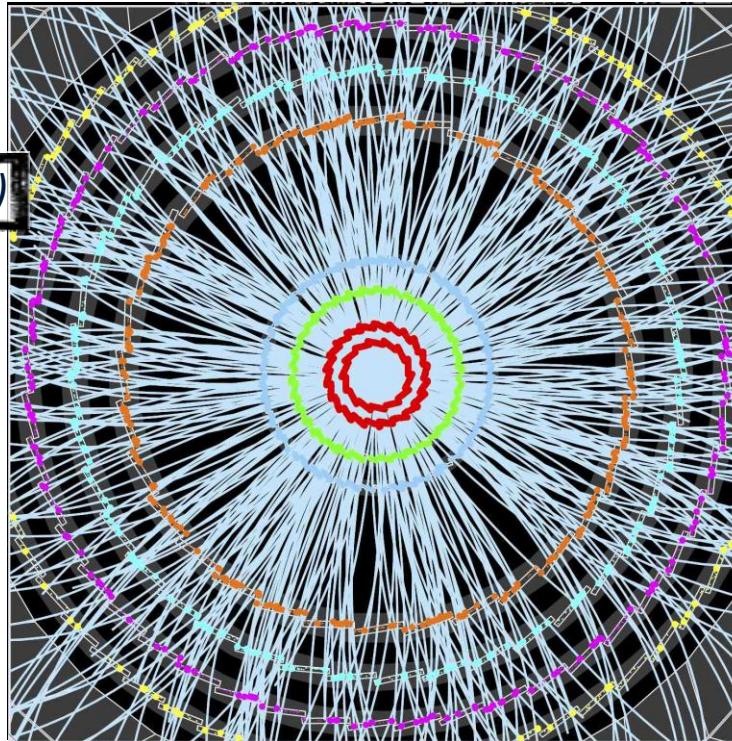
~60 TByte/s (compressed)

Trigger system selects
600 Hz ~ 1 GB/s data

~ 10 seconds compute for
a single event at ATLAS
with 'jets'
containing ~30 collisions

~10k researchers

CERN and ~170 institutes

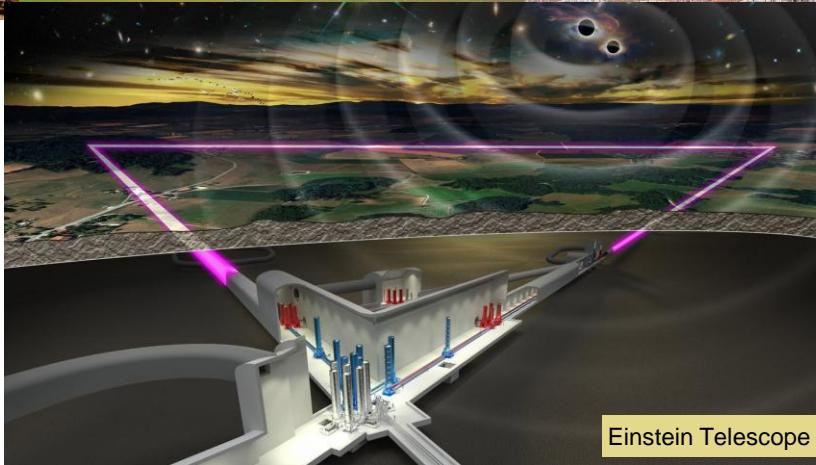


Display of a proton-proton collision event recorded by ATLAS on 3 June 2015, with the first LHC stable beams at a collision energy of 13 TeV;
Event processing time: v19.0.1.1 as per Jovan Mitrevski and 2015 J. Phys.: Conf. Ser. 664 072034 (CHEP2015)

Scaling computing infrastructure – a common need



Gravitational Waves



Einstein Telescope

SKA-Low (impression, Australia)



EOSC-WeNMR portals
@Bonvinlab

Announcing the EOSC-WeNMR Portals for Integrative Modeling of Biomolecular Interactions @ Bonvinlab, December 17, 2020

The user will find their favorite in 400,000+ entries on which features you would like to see implemented over their development with you.

Research

Research within the Computational Structure Biology group focuses on the development of novel methods and computational tools for biomolecular structure determination and analysis. Our research interests are centered on the use of high-throughput experimental data, structural bioinformatics and machine learning to determine atomic-resolution structures of proteins, nucleic acids and their complexes. We are also interested in developing a framework for the analysis of ensemble structures and in the use of machine learning to predict the structure of proteins and nucleic acids.

Services

ADDOCK

ADDODC is a magnetic platform for protein docking. It is a high-resolution docking platform that is able to dock a large variety of rigid and flexible molecules, from nucleic acids to proteins, and is able to predict the binding affinity of the complex.

proAiGy

proAiGy is a new predictor for protein-protein interaction prediction.

DISCOVER

DISCOVER identifies the structural and functional features of proteins and nucleic acids.

POWER

POWER automatically the atomic models of protein density maps.

Spot

SPOT is a new predictor for protein-protein interaction prediction.

WHISCY

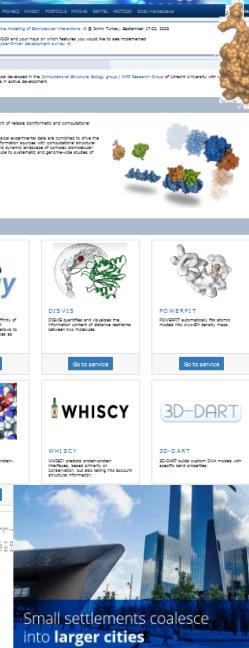
WHISCY predicts protein-protein interaction interfaces, local contacts, and local contacts.

3D-DART

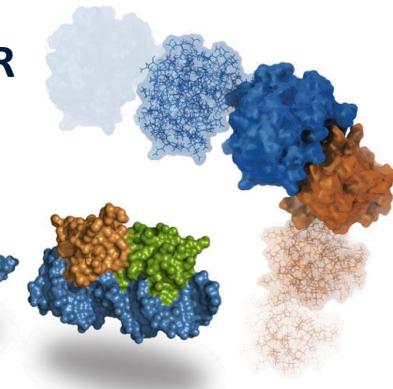
3D-DART predicts protein-DNA models with specific local properties.

proABC 2

proABC 2 is a new predictor for protein-protein interaction prediction.



WeNMR



Institutions for Collective Action
www.collective-action.info

HOME ABOUT ICAS RESEARCH ON ICAS COMMON RESOURCES EVENTS ABOUT US

COMPLETED PROJECTS

MODELLING INSTITUTIONAL DYNAMICS IN HISTORICAL COMMUNITIES (MIDI)

The Modelling institutional dynamics in historical communities (MIDI) project adopted an agent-based model approach to contribute empirically-grounded knowledge and insights on mechanisms driving the dynamics of institutional change. The starting point is a dataset of European communities management institutions compiled by the Research Team Institutions for Collective Action in the Common Rules-project.

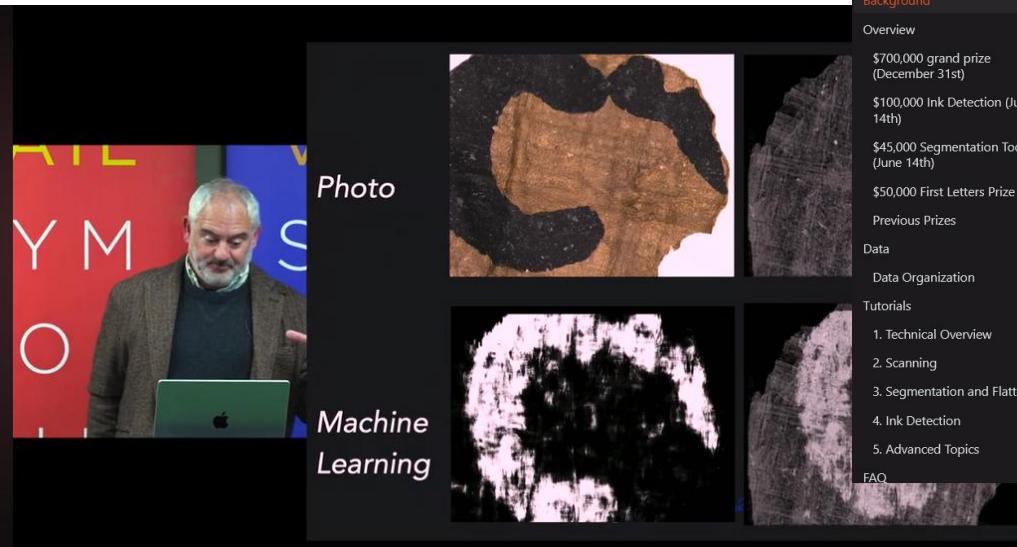
The goal of the project was to expand our theoretical/analytic competencies to model the dataset on the basis of a systematic institutional analysis framework, to fully exploit its potential through the use of data mining and evolutionary analysis techniques, and to embed the resulting knowledge in an agent-based model capturing the essential dynamics of institutional change.

The project involves researchers from various research disciplines from Linnaeus University (Sweden), Delft University of Technology (The Netherlands), and Tine De Moor and René van Beursem from the Institutions for Collective Action Research Team (The Netherlands). The project is funded by the Research Foundation Flanders (FWO) and runs from 1 January 2010 until the end of 2012.

Sources: Einstein Telescope: <https://et-emr.eu/>; CERN <https://wlcg.web.cern.ch/>; HADDOCK, WeNMR, @Bonvinlab <https://wenmr.science.uu.nl/>; Virgo, Pisa, IT; SKAO: the SKA-Low observatory, Australia <https://www.skatelescope.org/> - OpenMOLE simulation on EGI - https://cdn.egi.eu/app/uploads/2022/04/EGI_Use_Cases.pdf; agent-based modelling of ICAs: <https://collective-action.info/research-on-icas/> Molood Dehkordi (TUDelft), Tine de Moor (EUR RSM)

Collaborative computing changing fields you may not expect

Brent Seales' work on En-Gedi and Herculaneum scrolls with virtual unrolling and machine learning



A screenshot of a website for the 'Vesuvius Challenge'. The URL is https://scrollprize.org. The main content area features a large image of a volcanic eruption with the text: 'Resurrect an ancient library from the ashes of a volcano. Win \$1,000,000.' Below this, a paragraph reads: 'The Vesuvius Challenge is a machine learning and computer vision competition to read the Herculaneum Papyri.' On the left, a sidebar lists competition details: '\$700,000 grand prize (December 31st)', '\$100,000 Ink Detection (June 14th)', '\$45,000 Segmentation Tooling (June 14th)', '\$50,000 First Letters Prize', 'Previous Prizes', 'Data Organization', 'Tutorials' (with links to '1. Technical Overview', '2. Scanning', '3. Segmentation and Flattening', '4. Ink Detection', '5. Advanced Topics'), and 'FAQ'. The University of Kentucky logo is at the bottom.

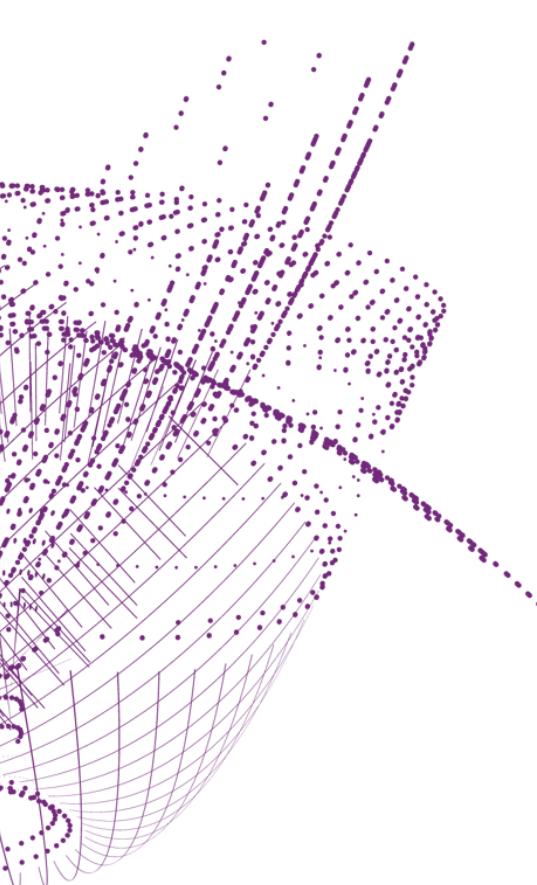
Photograph Herculaneum scrolls: The Digital Restoration Initiative/PA; capture Brent Seales from youtu.be/TUmWqsFrJpk; ML challenge: scrollprize.org

This is a tour of {a,one} large-scale IT landscape

'Exploring the e-infrastructure with use cases from data intensive research'

- **building** a compute, storage, data, and network facility
for high-throughput computing at the LHC scale
- a global **collaborative** infrastructure
with trust and identity, in a secure way
- **sustaining the ecosystem** we have built
*the Research Infrastructure Commons, the GORC, and
how research principles can guide digitalisation*

... to make ICT a research instrument rather than 'just a tool'



Building the facilities

Single CPU scaling stopped around 2004

- limitation is power, not circuit size
 - and clock frequency is most 'power-hungry'
 - still some packages now @ TDP of 400W
- multiple cores on the same die helps:
 - AMD EPYC Genoa (Zen 4) has 96 cores/die
 - Intel Granite Rapids, Nvidia GraceHopper, ...
 - but e.g. Intel Cascade Lake AP was less useful
- CPU design-level performance gains left
 - predictive and out-of-order execution
 - on-die parallelism (multi-core)
 - pre-fetching and multi-tier caching
 - execution unit sharing ('SMT')

but at increased risk for security/integrity

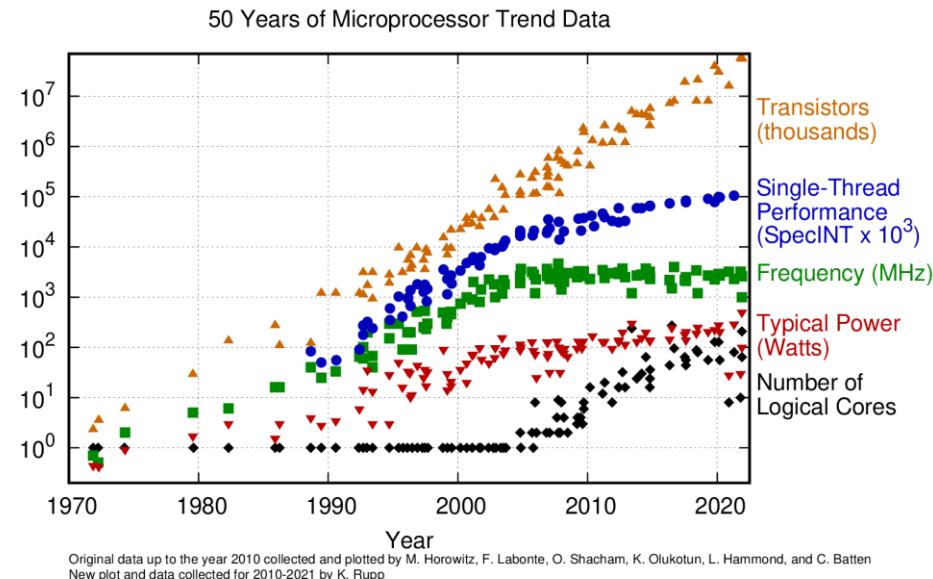
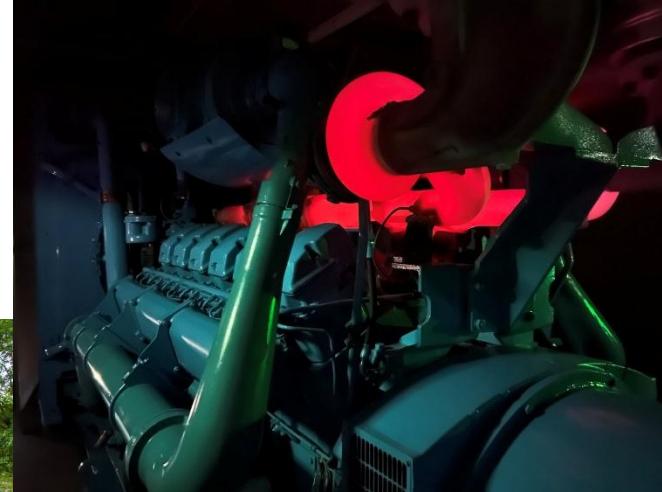


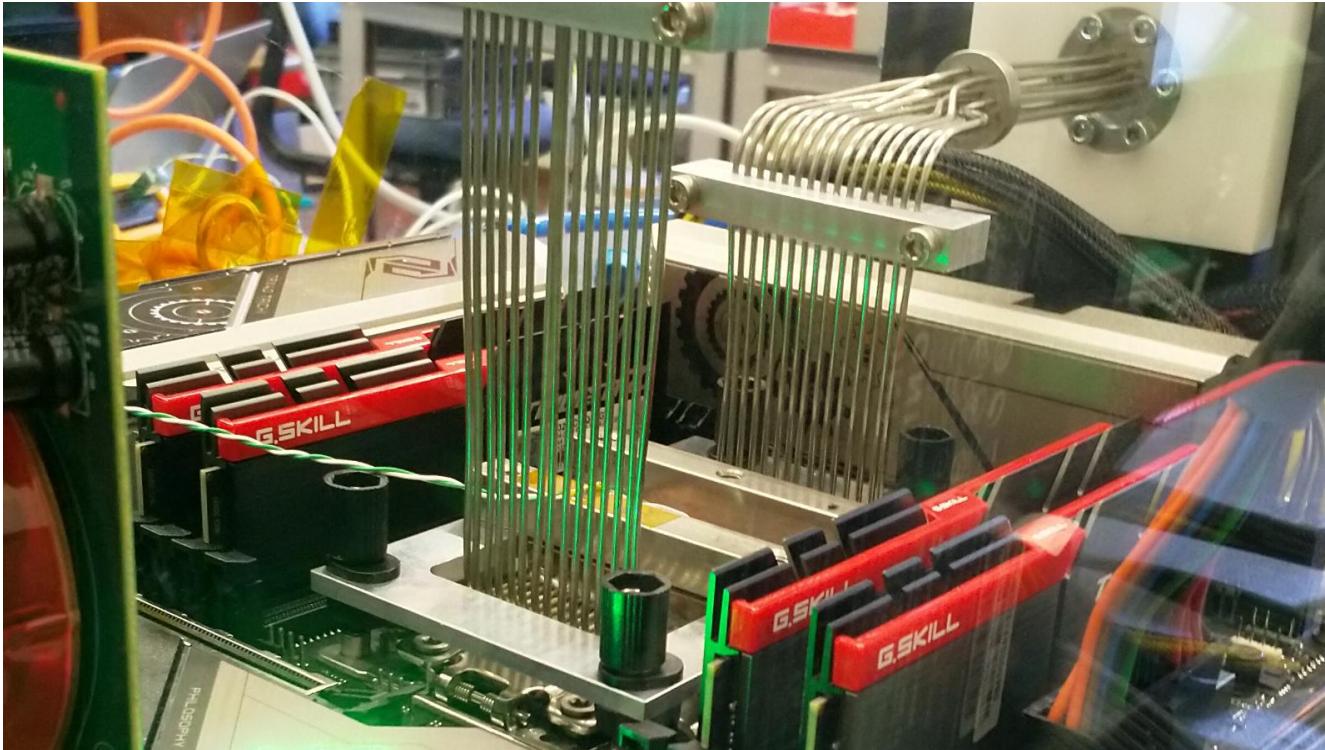
Image: K Rupp, <https://github.com/karlrupp/microprocessor-trend-data>

'I got the power ...'



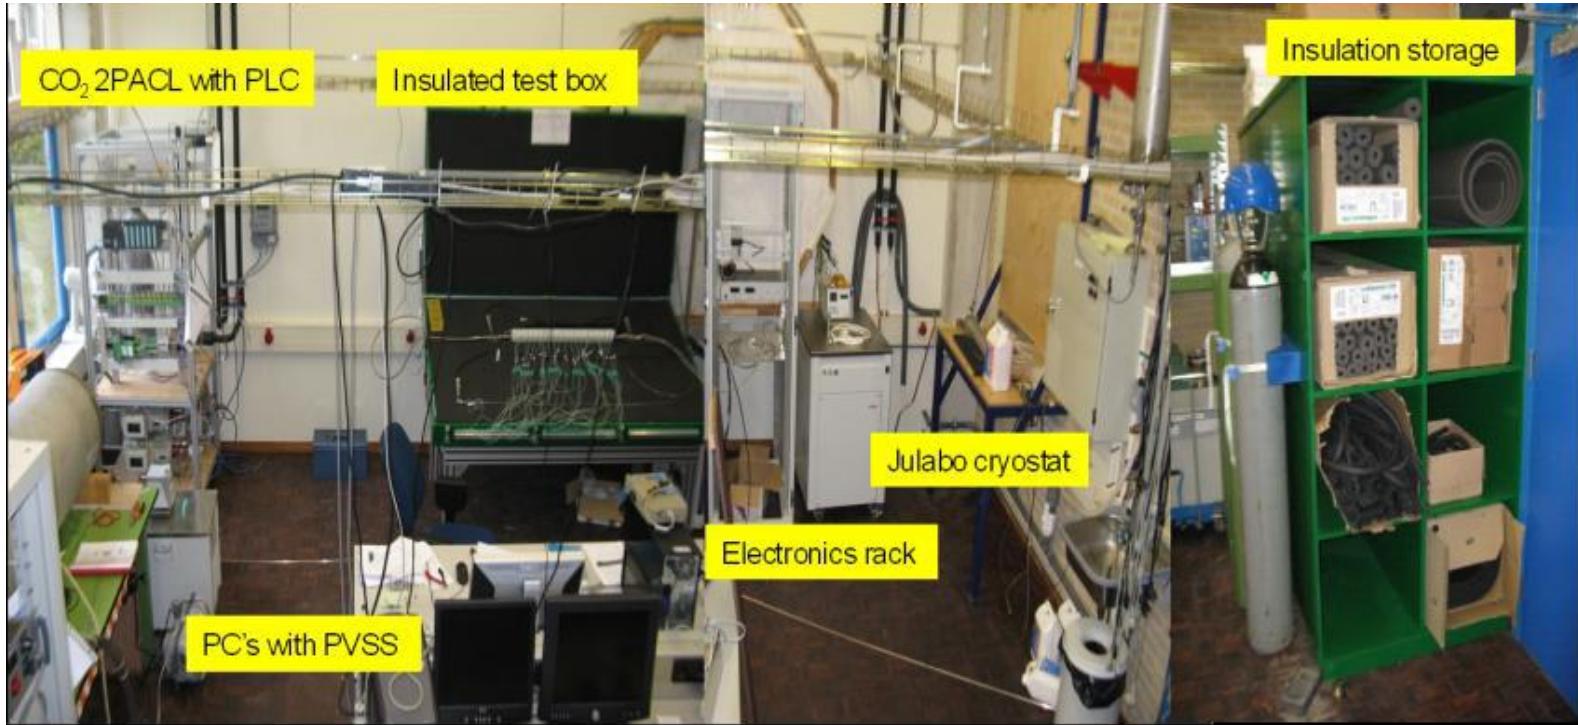
Images: Anton Mors, David Groep, Nikhef

Fix the thing that didn't scale well, CPU frequency??



LCO2 cooling of an AMD Ryzen Threadripper 3970X [56.38 °C] at 4600.1MHz processor (~1.25x nominal speed) sustained over all cores simultaneously, using the Nikhef LCO2 test bench system (<https://hwbot.org/submission/4539341>) - (Krista de Roo en Tristan Suerink)

... since you then need this around it ...



Nikhef 2PA LCO₂ cooling setup. Image from Bart Verlaat, Auke-Pieter Colijn *CO₂ Cooling Developments for HEP Detectors* <https://doi.org/10.22323/1.095.0031>

So we scaled up *inside* one system

Multiple cores and SMT on a single die

- ‘trivial’ step-up is to do multiple sockets in one system
2-socket, sometimes 4 socket on a motherboard
- appears as a single shared memory system, but requires *cache coherency* between CPU cores and sockets ...
which is useful for tightly coupled parallel applications
but not needed for ‘trivially parallel’ high throughput needs
- depending on architecture
cache coherency may limit single-thread performance
(although AMD did better here than Intel *lakes)

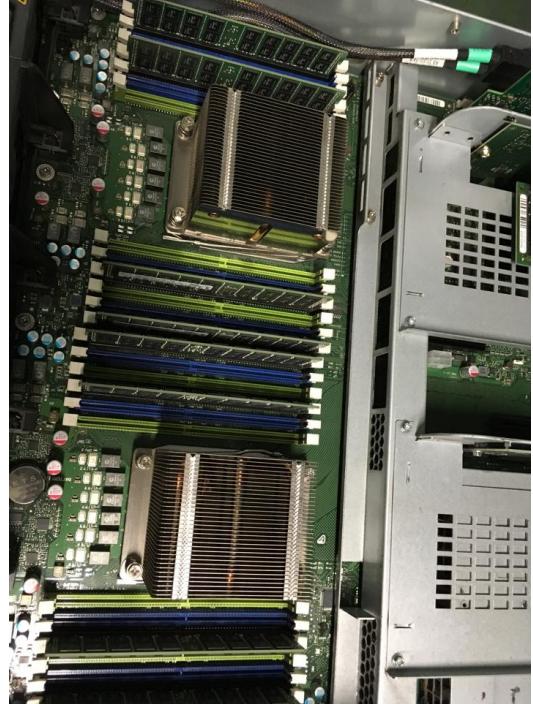


Image: dual-socket Fujitsu system at the Xenon experiment site, 2019. source: Tristan Suerink, Nikhef

The advantages of LHC-like data

- ‘tightly coupled’ HPC and (computational) cluster computing:
 - modelling for weather/climate, fluid dynamics, but also e.g. QC-simulation
- HTC and data-intensive processing for horizontal scaling:
 - lots of data, as in High Energy Physics (HEP), *omics and protein docking, ...
 - conveniently parallel,
but (intensive) local I/O requirements on memory and scratch storage
- portals and many web applications: ‘horizontal’ scaling,
but for RI use cases often backed by HPC or HTC resources ...
 - science gateways and portals (like <https://wenmr.science.uu.nl/>)
 - interactive notebooks and analysis environments

HPC: High Performance Computing; HTC: High Throughput Computing

CPU design changes may fit application, or not

Example: AMD EPYC effective for LHC-like workloads

- Naples → Rome added shared memory die
- links all cores directly to memory

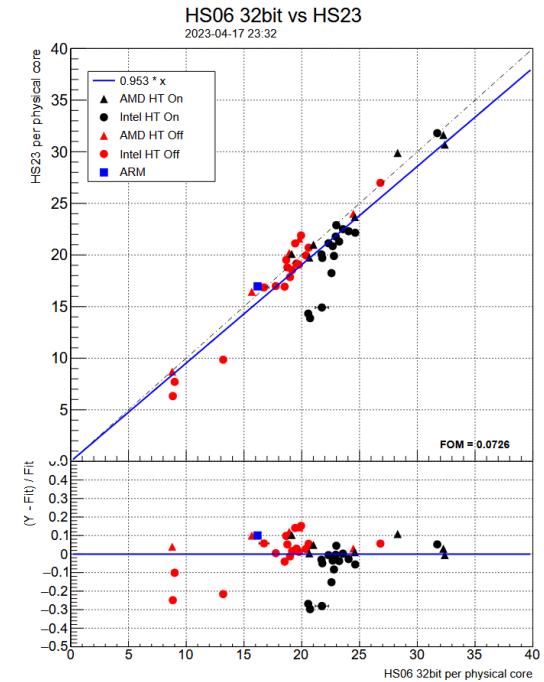


Image source: AMD, retrieved from <https://m.hexus.net/tech/news/cpu/135479-amd-shares-details-zen-3-zen-4-architectures/>

HEPscore: <https://doi.org/10.48550/arXiv.2306.08118> ; AMD-EPYC architecture benefits memory-intensive HEPscore23 over HS06 ('memory subset of SPECint06')

CPU design changes may fit application, or not

But the
Rome-Milan improvement ...?

- shared L3 cache
benefits tightly coupled HPC,
but not HTC, limited by
'off-die memory'

*Which is also why single-socket
systems outclass dual-socket
(also on TCO)*

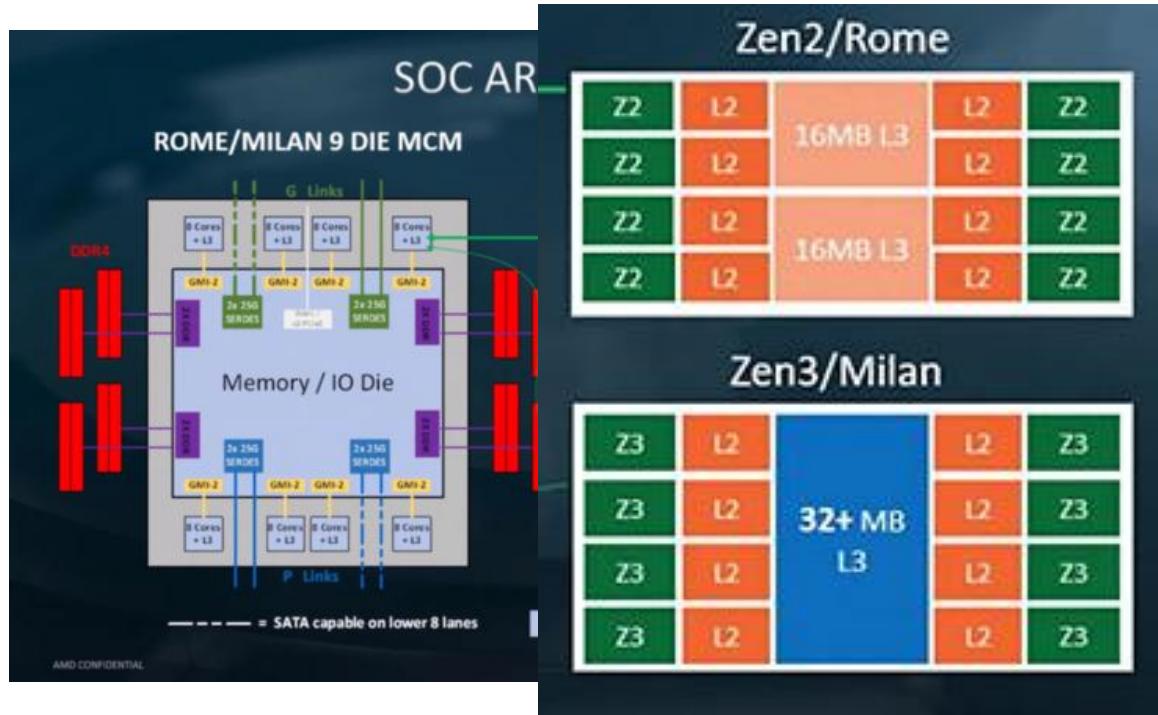


Image source: AMD, retrieved from <https://m.hexus.net/tech/news/cpu/135479-amd-shares-details-zen-3-zen-4-architectures/>

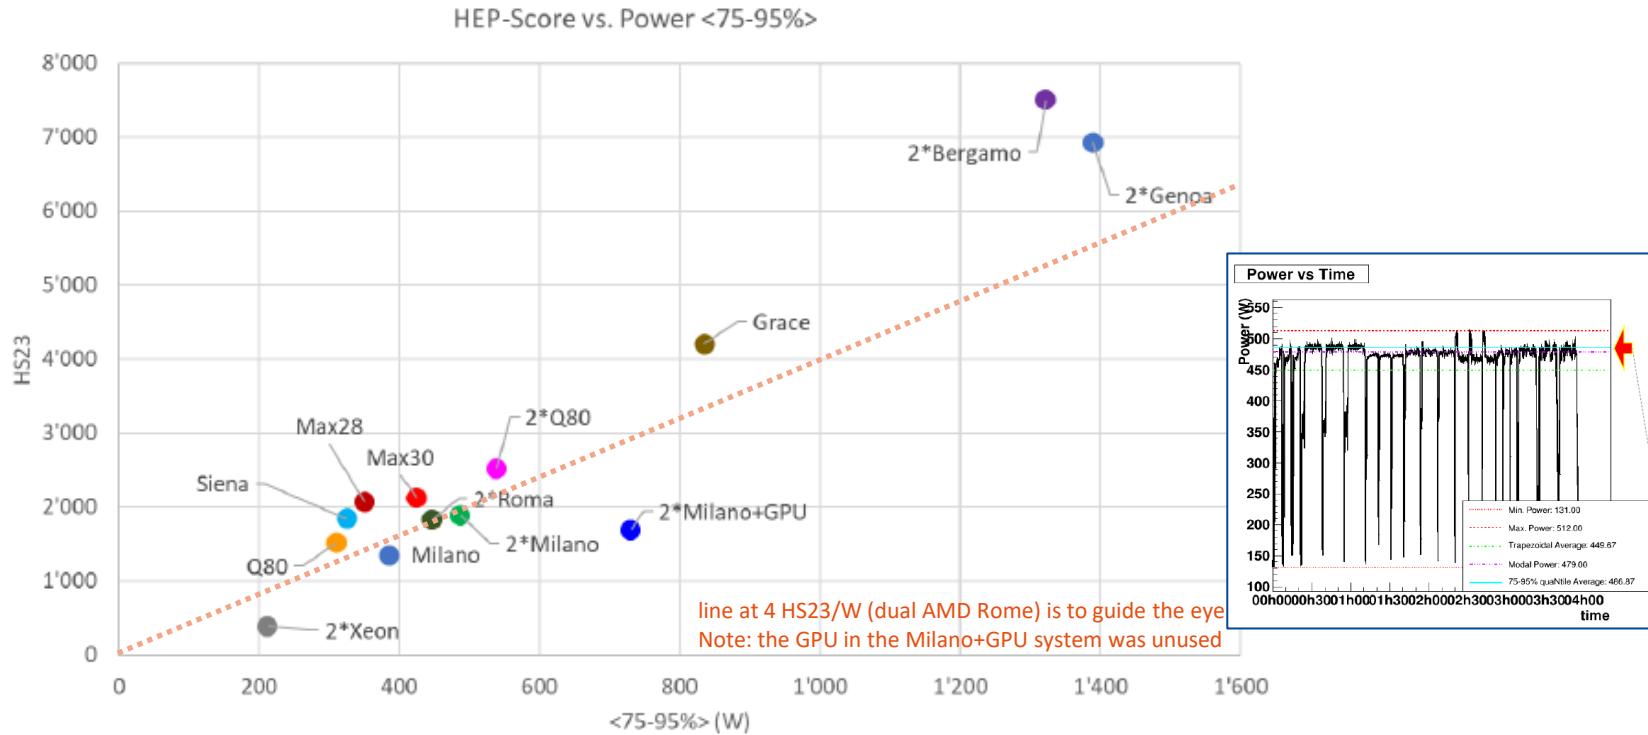
... and indeed we see it in the HEPscore benchmarks

Generation	HEPscore/core	Clock(Ghz)	HEPscore/Ghz	W	Cores	HEPscore/W
naples	18.192	2.5	7.28	180	32	3.23
rome	27.171	2.6	10.45	280	64	6.21
milanX	26.171	2	13.09	280	64	5.98
Genoa	35.551	2.45	14.53	280	64	8.13
Genoa	29.724	2.40	12.385	360	96	7.926

The HEPscore/W is the most relevant number for an ‘always full’ system
and for TCO due to energy price, at least until memory prices exploded in October ’25

Performance data from Emily Kooistra *A Tour of Italian CPUs and the fastest road South*, Nikhef Jamboree 2023.
HEPscore: Giordano *et al.* EPJ Web of Conferences 295, 07024 (2024) <https://doi.org/10.1051/epjconf/202429507024>

The energy bottleneck: architecture 'figure of merit'



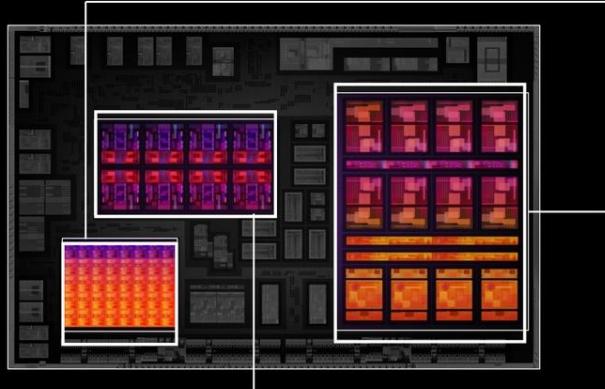
Data and graphs: Emanuele Simili, Glasgow University, at CHEP2024 (<https://indico.cern.ch/event/1338689/contributions/6011562/>)

HEPSPEC23 benchmark: <https://gitlab.cern.ch/hep-benchmarks/hep-benchmark-suite> ('memory-intensive' high throughput processing application benchmark)

Hybrid SOCs and heterogeneous architectures

3rd Gen AMD Ryzen™ AI

Next-gen AI PC experiences demand the best of NPU, CPU, and GPU architectures



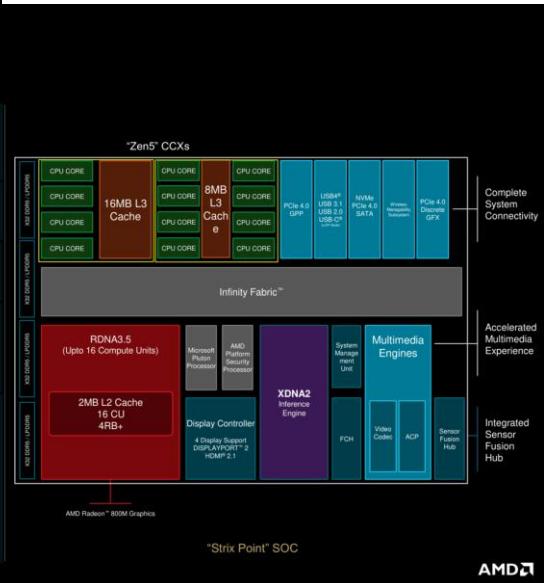
(note these are laptop/desktop SKUs, not servers)

NPUs, GPUs, APUs ...

AMD “Strix Point” SOC

- CPU**
 - 4C8T Zen5 – 1MB L2/core, 16MB L3 CCX
 - 8C16T Zen5c – 1MB L2/core, 8MB L3 CCX
 - Datapath – 32B/cycle port each
- GPU**
 - 8 WGP (16 CU) RDNA 3.5
 - Datapath – 4 x 32B/cycle ports
- NPU**
 - 4 x 8 Array XDNA 2 Inference Engine
 - Datapath – 32B/cycle
- Accelerators / uControllers**
 - Video Encode/Decode
 - Audio Co-processor
 - Display Controller
 - System Management, Security, Wireless Manageability
- IO**
 - 128b LPDDR5/DDR5 (7500/5600 MT/s)
 - 16L PCIe Gen4
 - 4 Simultaneous display streams
 - 8 USB ports
 - 2 USB v1
 - 1 USB Type-C
 - 2 USB 2 Gen2
 - 3 USB2
 - I2c, SPI/eSPI, GPIO

16 NEXT GENERATION ZEN 5™ CORE



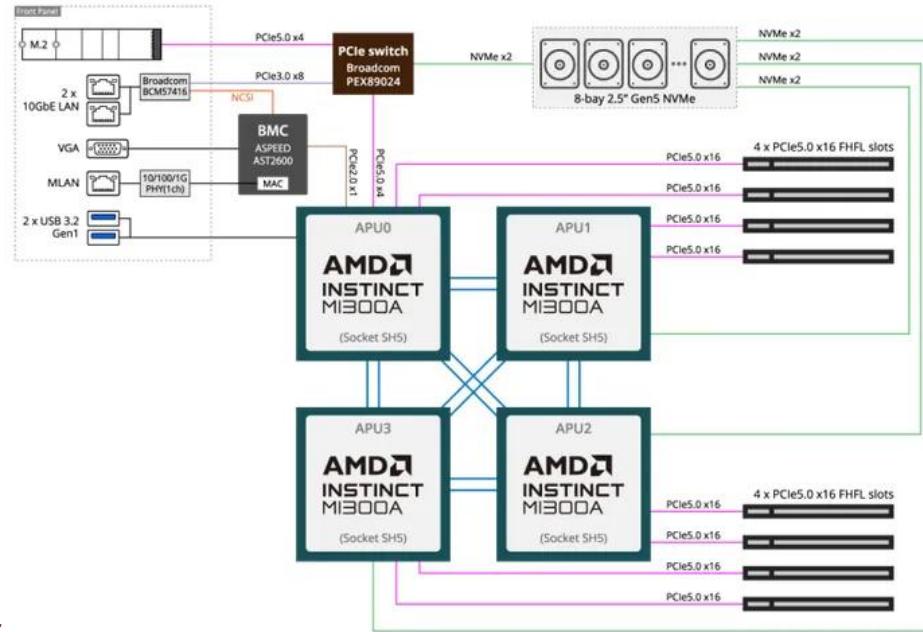
Images: AMD Ryzen 9 HX 370 AI, Strix SOC – compare also Intel Lunar Lake architecture

ML ‘big physics models’ are changing that, but at a cost ...

Current models tend to be very large, training barely fitting in an H100, and inference also needs 48-96 GB

- conventional GPUs for training are outgrowing budgets very fast
- validate hybrid/APU architectures?

Challenge of course is software porting if doing more than just pyTorch – but that is more of a worry when writing kernels for real-time applications like in the HLT



Systems block diagram: GigaBYTE G383-R80-AAP1, Nikhef SIF “Bordercollie” ML training system

but there is also a serious issue with sockets ...

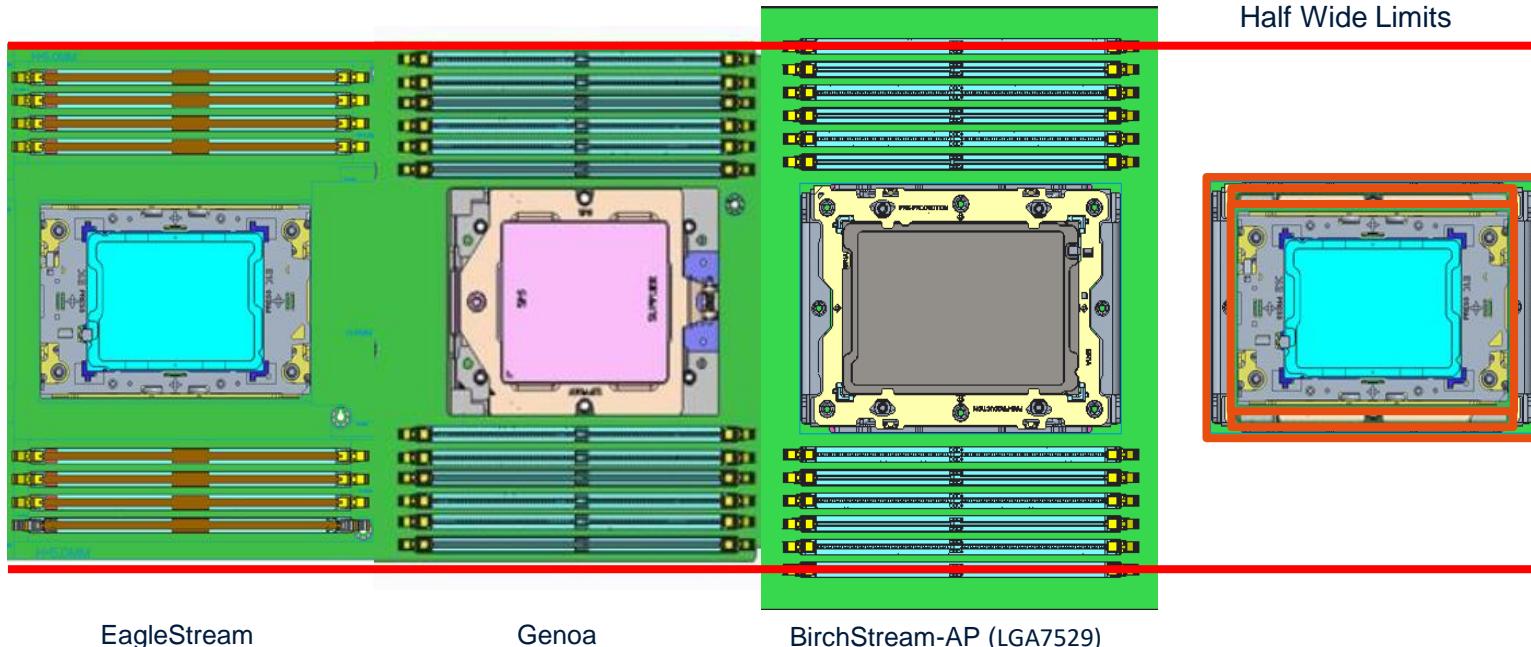


Image thanks go to Rick Koopman – Lenovo at the HTCondor Workshop 2024 <https://indico.cern.ch/event/1386170/>

So if large-scale IT does not quite fit ... ahum ...



SuperMicro (branded as 'Lambda Blade')
4U chassis, supporting 10 consumer-grade GPUs ...
... with a bump

Image source: <https://lambdalabs.com/products/blade>

And it's hot in there ...

- Heat capacity of liquid is much larger than air
- by now (almost) standard for HPC systems
- immersive systems
look cool, but are 'a bit
hard' on maintenance



Strongly depends on systems engineering:
when water inlet temperature can be >40 degC, you have almost always free cooling

Image source dual-board system: Lenovo, ThinkSystem SD650
immersive cooling image <https://hypertec.com/blog/sustainable-emerging-tech-liquid-immersion-cooling/>, PIC T1 centre, Barcelona, ES

Scaling up – beyond one lone system



Typical compute farm @Nikhef for ‘milking’ computer clusters

Continuous design challenge

- **balanced features** for node throughput
CPU, storage, memory bandwidth
& latency, NIC & network speed

For example for WLCG:

- **single-socket** multicore systems are fine,
today typically 64-128 cores per system
- **network**: 2x25/2x100Gbps (matching #cores)
- **memory**: say ~ 8 GiB/core
- **local disk**: 8-16 TB NVME (~100GB/core)
- + space (physical + power) to add **GPUs**

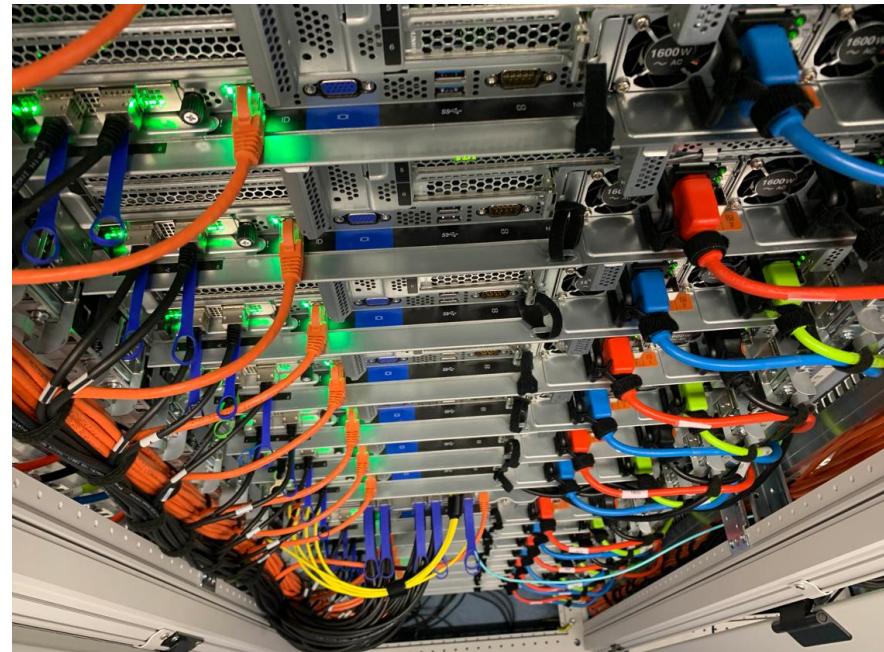
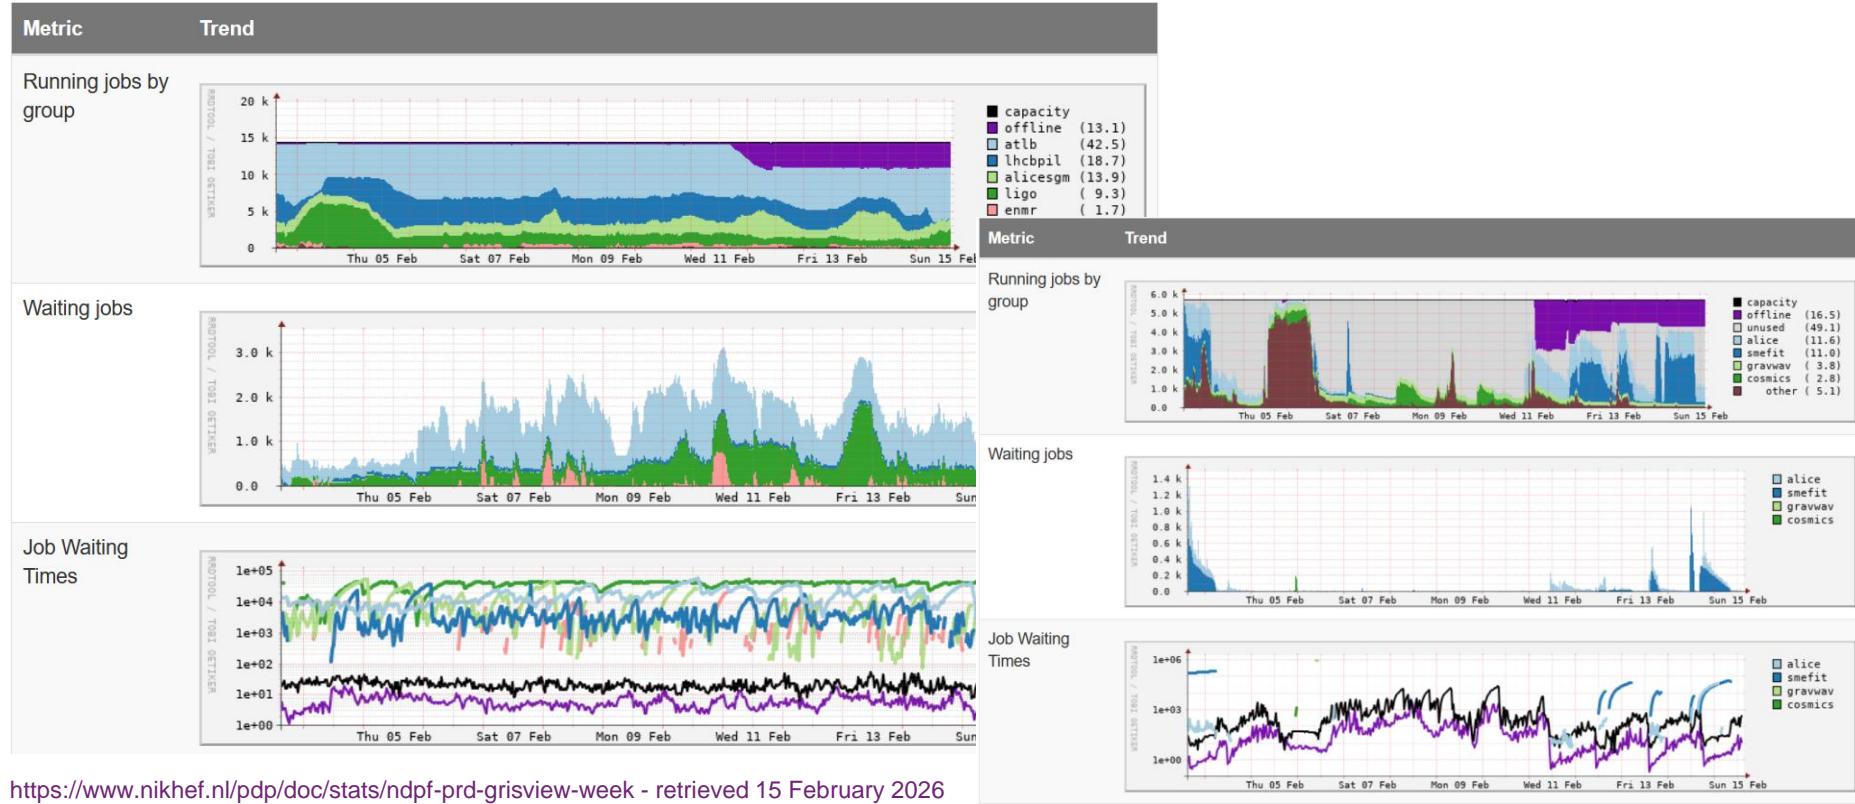


Image: Cluster ‘Lotenfeest’ at the Nikhef NPDF, acquired March 2020. Lenovo SR655 with AMD EPYC 7702P 64-Core single-socket. Some with 4 L40s Nvidia GPUs

To fill or not to fill, that's the question ...



<https://www.nikhef.nl/pdp/doc/stats/ndpf-prd-grisview-week> - retrieved 15 February 2026

Occupancy: balance efficient use of resources and happy users

For organized ‘production’ computing (planned months in advance in WLCG)

- ***predictable scheduling*** is more important (steady flow of results)
- **maximizing efficiency**: resource cost is the limiting factor in (physics) results
- co-scheduling with data (pre-placement) is required
- community-authorization based access to data sources only

For ‘local’ users, e.g. students whose progress tomorrow depends on results *today*

- ***response time*** is more important than efficiency
- fast turn-around/short waiting times by heterogeneous (‘competing’) user base
- data access must be parallelism-ready, but is ‘always’ local on-site
- with local credentials and sharing with desktop and Jupyter environments

so offering two distinct classes of services is (in this case) intentional

Standard interfaces for compute and data?

hourglass model ‘kind-of’ worked for IP and web with http as common standard

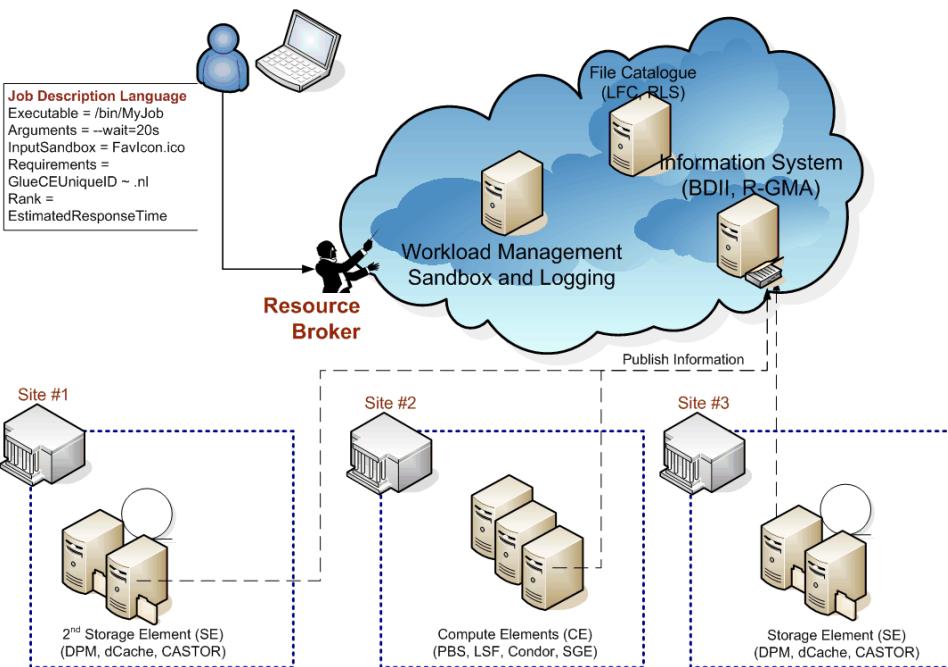
- a very simple stateless interface

protocols for higher-level services never quite reached this level of global interop

- requirements too complex and stateful
- use cases were usually scoped

slowly changing now but only for similarly simple things, like on-line object storage

Is distributed computing too bespoke ...?



Interoperable cloud? Compare OGF's OCCI WG GFD.221 (<https://www.ogf.org/documents/GFD.221.pdf>) with e.g. Amazon S3 API or the OwnCloud CS3 interfaces

DIRAC: spanning heterogeneous resource models

Add a scheduling layer!

‘any (IT) problem can be solved by adding an extra level of indirection’**

DIRAC is just one example

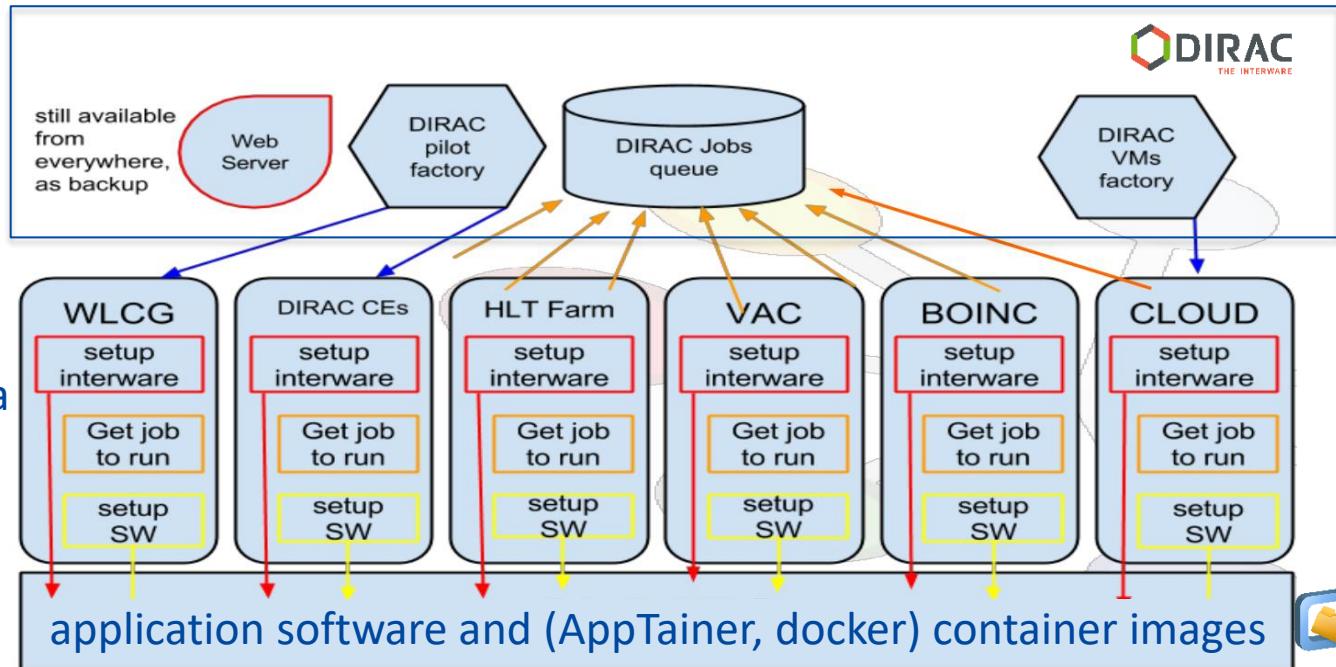
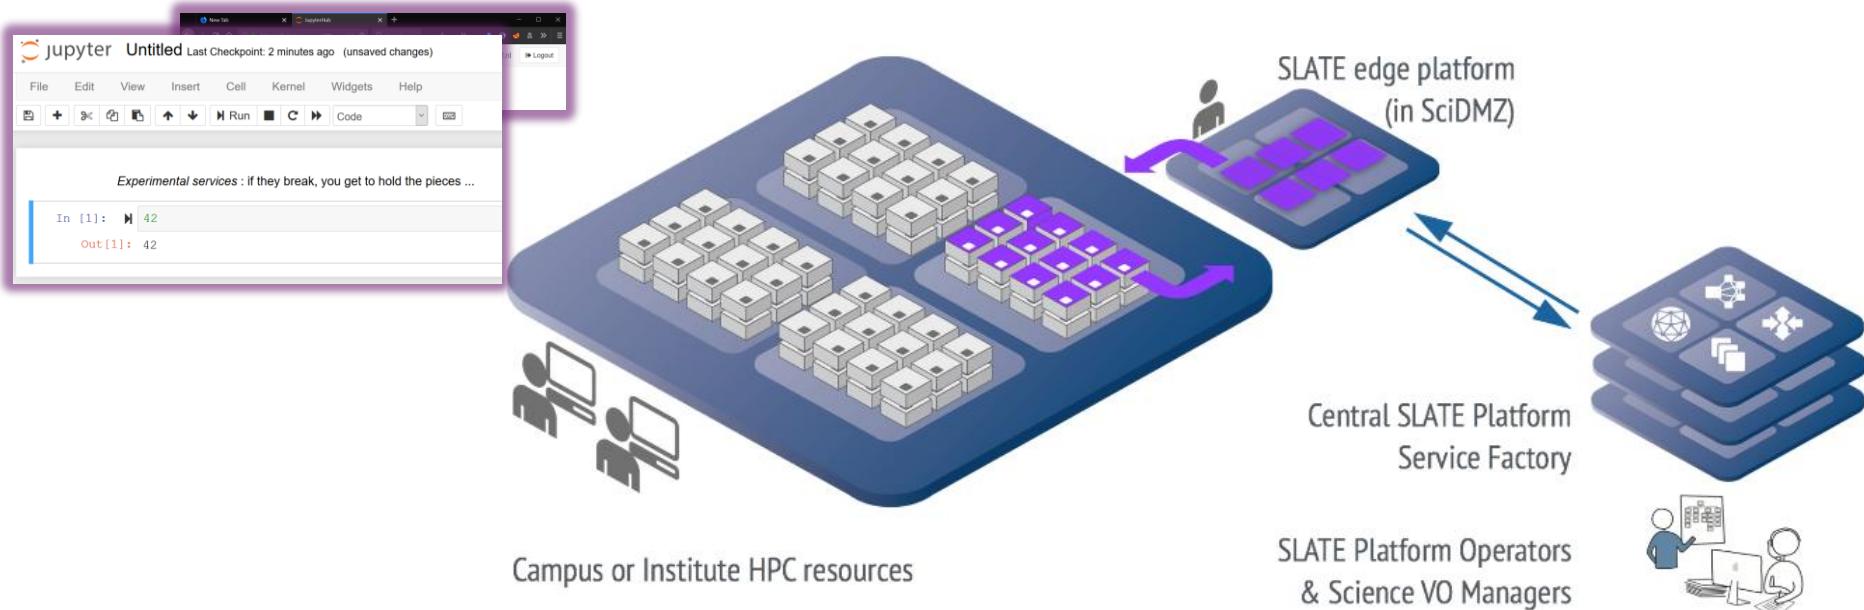


Image: DIRAC project, A. Tsaregorodtsev *et al.* CPPM Marseille, from <https://dirac.readthedocs.io/> ; CVMFS (CERN VM File System) is a common software distribution platform using distributed signed data objects in a cached hierarchy using CDN techniques, see <https://cernvm.cern.ch/fs/> * thanks to Miron Livny



An overlay network of containers

Nobody wants a cloud per-se ... what folk want is a solution ...

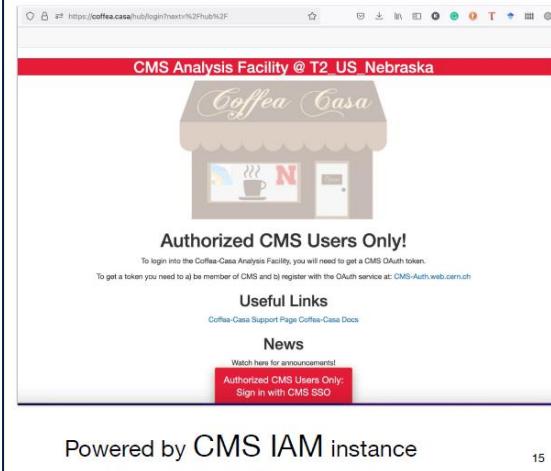


‘alien containers’ HPC integration - container computing, using curated application images

Image sources: NPDF JupyterHub service “Callysto”; SLATE: Service Layer At The Edge – Rob Gartner (UChicago), Shawn KcMee (UMich) *et al.* – slateci.io

Containerised workloads: between ‘PaaS’ and ‘SaaS’

CMS Coffea-Casa Analysis Facility: <https://coffea.casa>



CernVM
File system

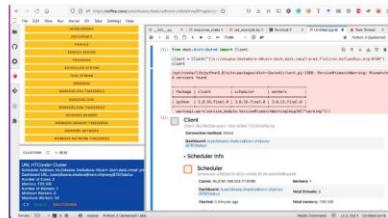


locks: easy integration with scalable computing



task-management computational
tion (based on the manager-worker

in HPC clusters, running a variety of
including SLURM, LSF, SGE and
“*task-jobqueue*”
us to create a user-level
system via queuing up in the



Dask can be used inside Jupyter or you can simply
launch it through Jupyter and connect directly from
your laptop

See also EESSI: the European
Environment for Scientific
Software Installations ...



Images: Oksana Shadura et al (UNebraska Lincoln), Brian Bockelman (Morgridge Institute) at CHEP2023 <https://indico.jlab.org/event/459/contributions/11610/>
EESSI software distribution (<https://www.eessi.io/>) is CVMFS + Modules

Storage for high-throughput processing

Basic storage properties are well known

- throughput
- IOPS – I/O Operations per Second
- seek-time (latency)

but not many **file systems** support *concurrent parallel access* by many clients

- both data **and** (file system or index) meta-data must be scalably distributed
- typically sacrifice either instant consistency, or (POSIX) semantics, (or scalability) in a distributed storage system

Common commercial solutions: GPFS, ... but also NetApp, HDS, Dell-EMC, have their own

Common open source: BeeGFS, gluster, dCache, CephFS, Lustre, ...

... likely do not use a file system if object storage does the job, but then you need a catalogue/database

'Interesting' distribution: client-side-managed GlusterFS

- scalable through independence of both clients and servers
- design is stateless: file system meta-data kept in each server's file system
- data itself can be replicated and protected, but ... inconsistencies in metadata linger around the corner in case of client failures (e.g. batch system worker nodes)

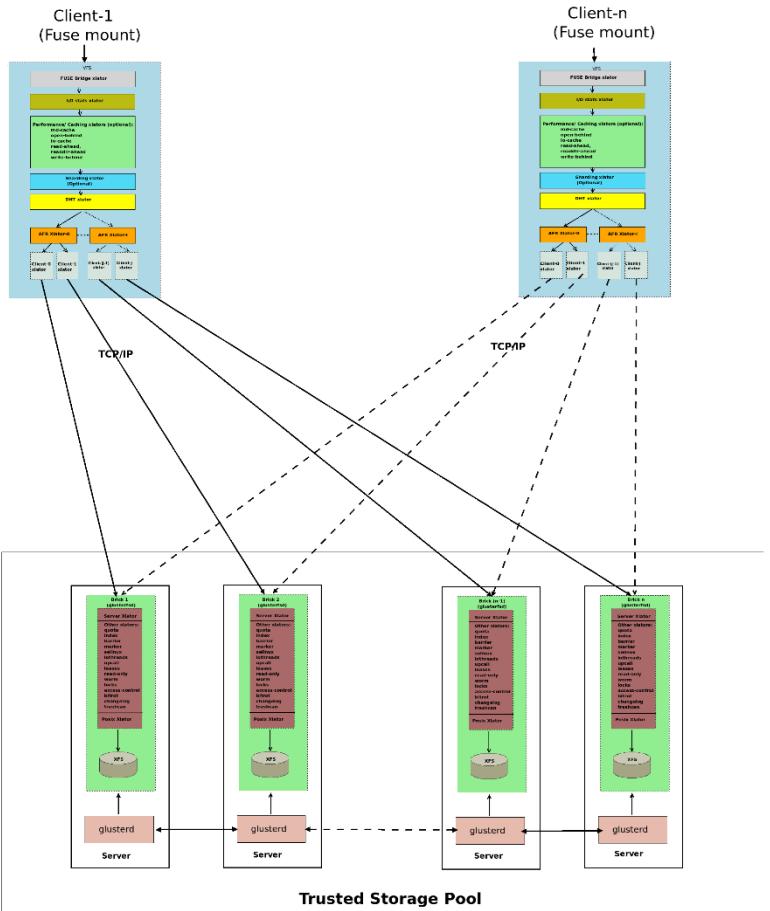
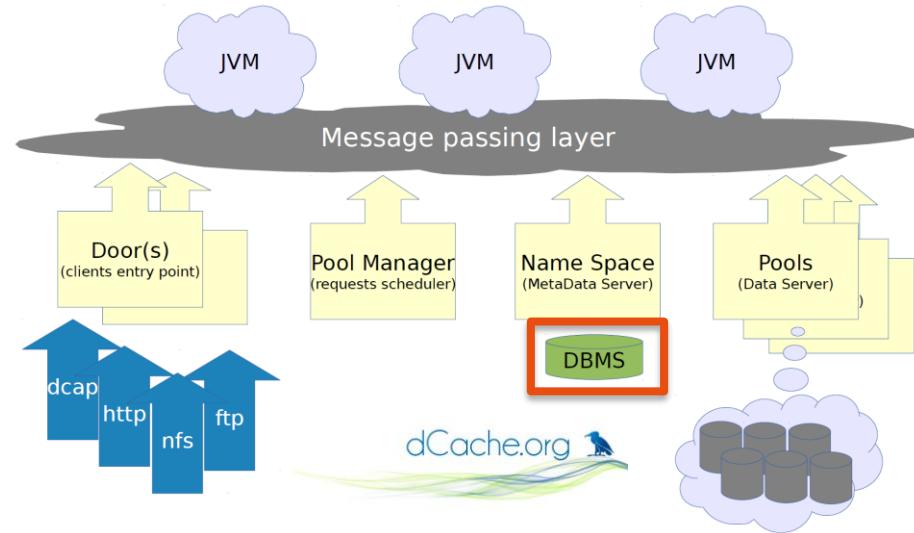
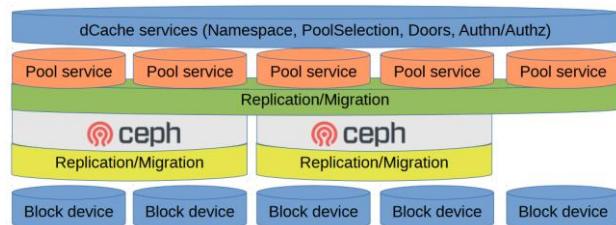


Image source Gluster community: <https://docs.gluster.org/en/main/Quick-Start-Guide/Architecture/>

Example: server-coherent distribution – dCache

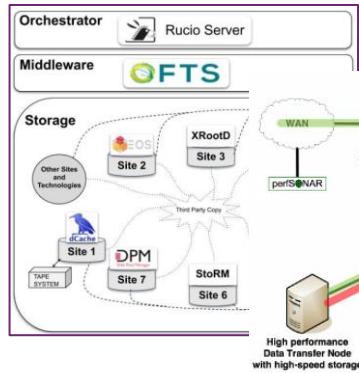
- separate client entry points, storage access scheduling, filesystem meta-data (namespaces), and storage
- message layer for eventual consistency
- redirect-based access
 - doors and pools usually on all nodes
 - now also feature of standard NFSv4.1



Images: Tigran Mkrtchyan (DESY, dCache.org), *dCache on steroids - delegated storage solutions*, ISGC 2016, <https://dcache.org/manuals/publications.shtml>

Structure of application data placement impacts storage (hardware) systems design

pre-staging all data locally allows for **latency hiding**,
posix-style access with `lseek(2)`, and a fast, local, '\$TMPDIR'
e.g. why there are *Data Transfer Nodes (DTNs)* in the 'Science DMZ' concept



but, nowadays, pre-staging started coming at a cost, when using **SSDs** as local 'scratch' area ... because of their hardware characteristic 'endurance'

Photo HGST nVME from: Dmitry Nosachev on Wikimedia Commons CC-BY-SA; Image Science DMZ and Data Transfer Nodes: ESnet fasterdata.es.net

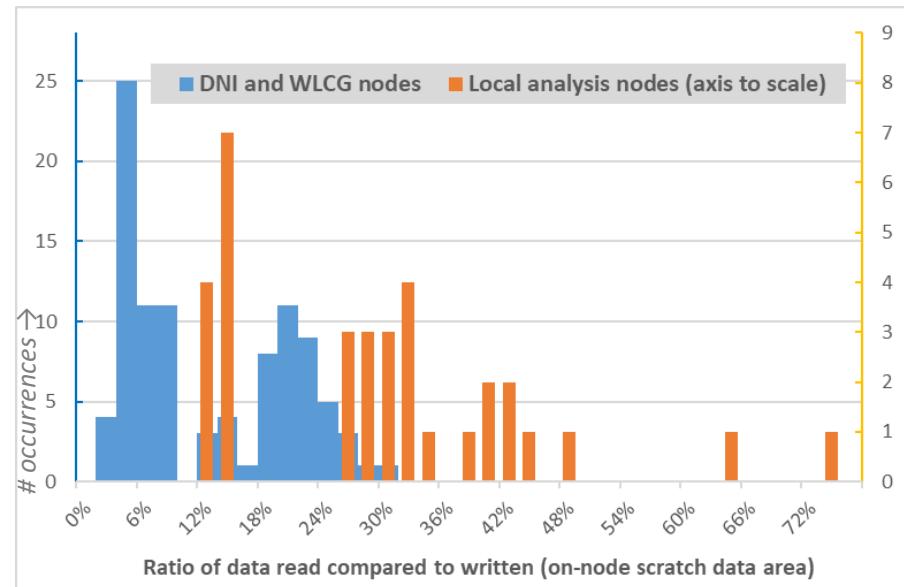
Especially with *WORN* storage: Write Once Read Never

Frequency distribution of **read-back vs. write** volume, observed on local scratch for NDPF execution nodes for ***outside ('grid') access (blue) vs local access (orange)***

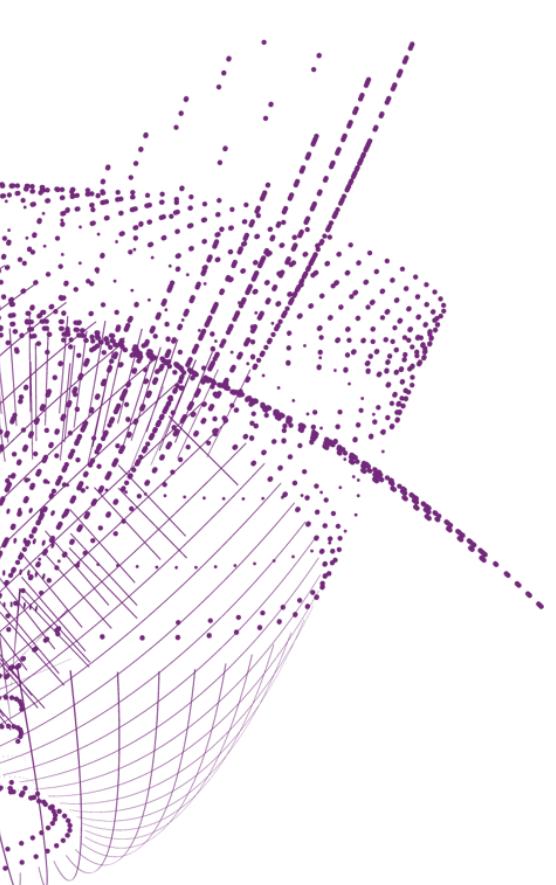
Access pattern is rather different. But why?

- external users pre-stage, because it is built into data management frameworks (like DIRAC, Athena),
- 'local' users stream output data (dCache with NFSv4) and use \$TMPDIR mainly for merging partial results

Different types of workload (here analysis vs processing) determine the choice of systems hardware



Data: NDPF execution nodes, based on SSD SMART data, integrated over total device lifetime; plot shows number of local analysis nodes scaled to DNI-WLCG count; collected using smartctl on 2020-10-28 – in total 97 'DNI' and 34 'STBC' SSDs were used in the analysis



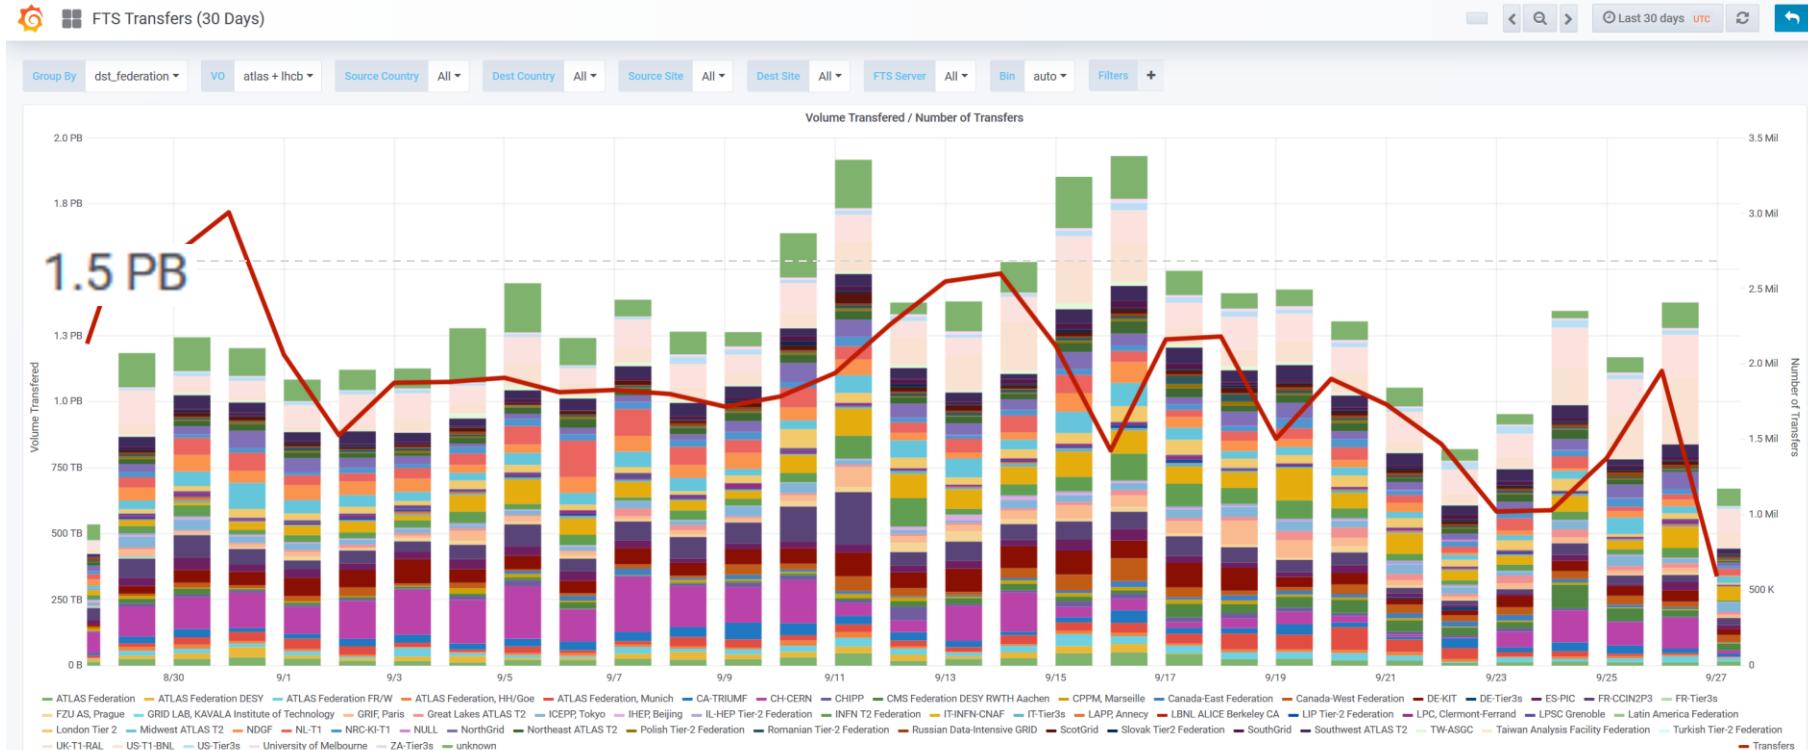
Beyond the single site

It's all about data
... globally interconnected



<https://wlcg.web.cern.ch/>

High throughput computing includes data and networks



source: <https://monit-grafana.cern.ch/d/000000420/fts-transfers-30-day> ; data: November 2020 ; CERN FTS instance WLCG: daily transfer volume ATLAS+LHCb

Can hardly be said better than Eli Dart did at TNC23

The Value Of Routine Performance

- It's important to get to where high performance is normal
- No magic, no arcana, things just normally work – for petabytes of data
- DOE HPC facilities now easily shuffle around hundreds of terabytes
 - Some people have smaller data sets too
 - But the point is that it's normal and routine
- What follows is one specific example, chosen because of some specific features

Network is more than just what it says on the tin

More network bandwidth does not mean your *data* gets there faster

- memory requirements (since TCP needs a capability to re-transmit)
- tcp 'slow start'
- congestion control algorithms

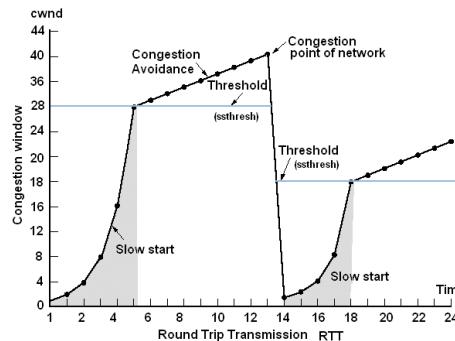
TCP throughput calculator

Theoretical network limit

rough estimation: rate $< (\text{MSS}/\text{RTT}) * (\text{C}/\sqrt{\text{Loss}})$ [C=1] (based on the Mathis et.al. formula)
network limit (MSS 9000 byte, RTT: 150.0 ms, Loss: 2.304×10^{-11} ($2 \times 10^{-9}\%$)) : **100000.00 Mbit/sec.**

Bandwidth-delay Product and buffer size

BDP (100000 Mbit/sec, 150.0 ms) = **1875.00 MByte**
required tcp buffer to reach 100000 Mbps with RTT of 150.0 ms \geq **1831054.7 KByte**
maximum throughput with a TCP window of 1831054 KByte and RTT of 150.0 ms \leq **100000.00 Mbit/sec.**



Useful sources: https://www.switch.ch/network/tools/tcp_throughput/, <https://fasterdata.es.net/>

tcp slow-start graphic from Abed et al, *Improvement of TCP Congestion Window over LTE- Advanced Networks* IJoARiC&CE 2012

The cat video that destroyed it all ...

- TCP protocol sensitive to packet loss
 - 3 lost packets is enough to trigger this
- different congestion avoidance algorithms exists (~ 20 by now)
- loss severely impacts links w/large 'bandwidth-delay-product' (BDP)
- NL: ~ 3 ms, US East: 150ms

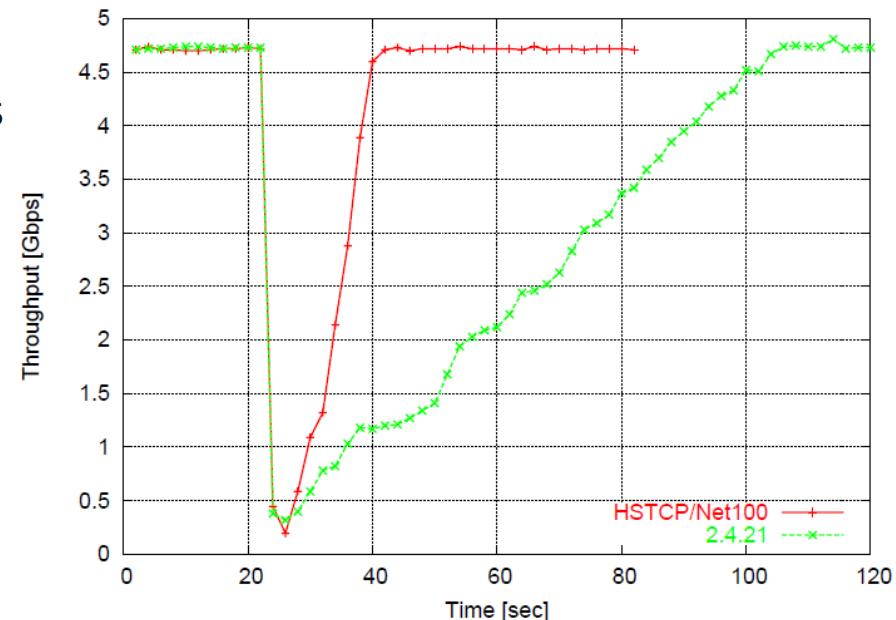
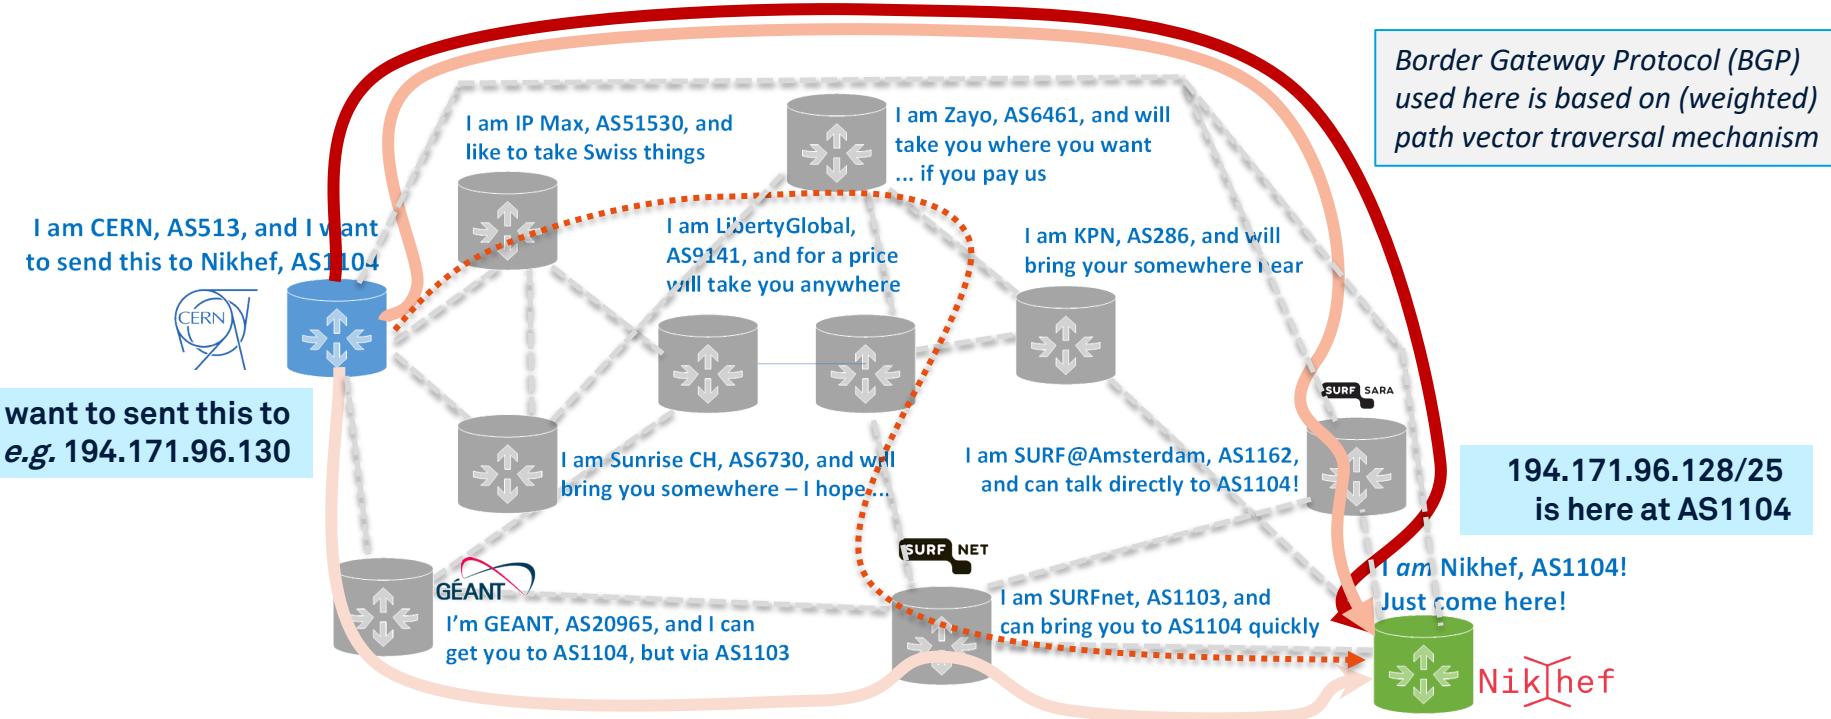


Figure 10: HSTCP versus stock TCP recovery time

source: Catalin Meirosu et al. *Native 10 Gigabit Ethernet experiments over long distances* in FGCS, doi:10.1016/j.future.2004.10.003 – aka. ATL-D-TN-0001

Fast track - and getting rid of cat videos ...



grey-dash lines for illustration only: may not correspond to actual peerings or transit agreements; red lines: the three existing LHCOPN and R&E fall-back routes; yellow: public internet fall-back (least preferred option)

LHCOPN – distributing raw data

LHCOPN

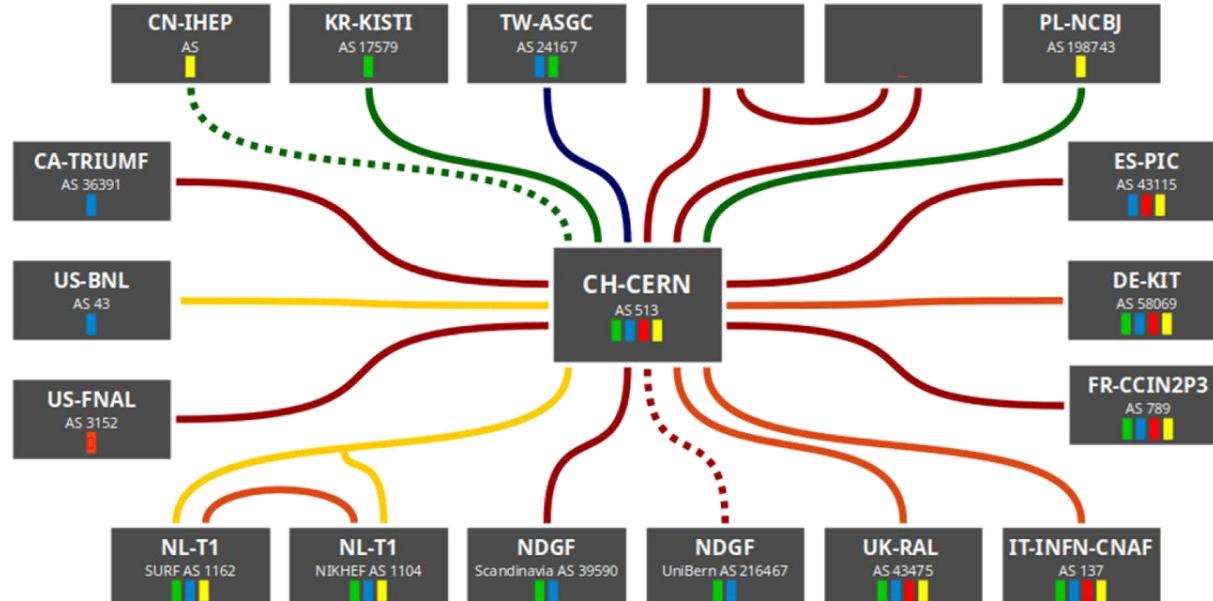
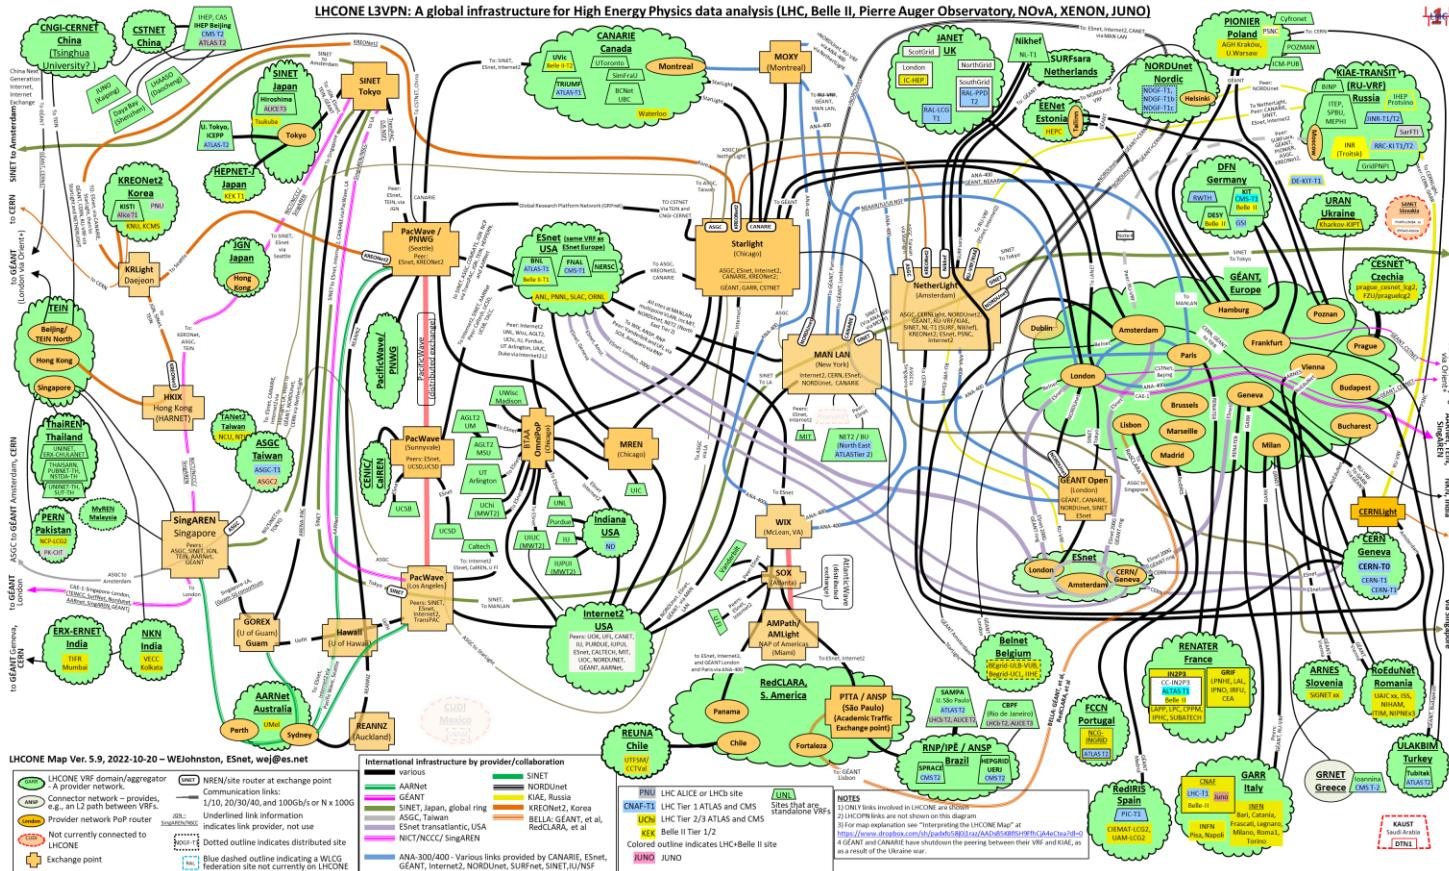


Image source: Edoardo Martelli, CERN, <https://lhcopn.web.cern.ch/>

LHConne



LHCOne (“LHC Open Network Environment”) – visualization by Bill Johnston, ESnet version: October 2022 – updated with new AS1104 links

‘ScienceDMZ’

Predictable performance
and data access for research

‘where research services,
data, and researchers meet’

- latency hiding through caching
- **security zoning/segmentation**
protects specific data sets
- **outside any enterprise perimeter**

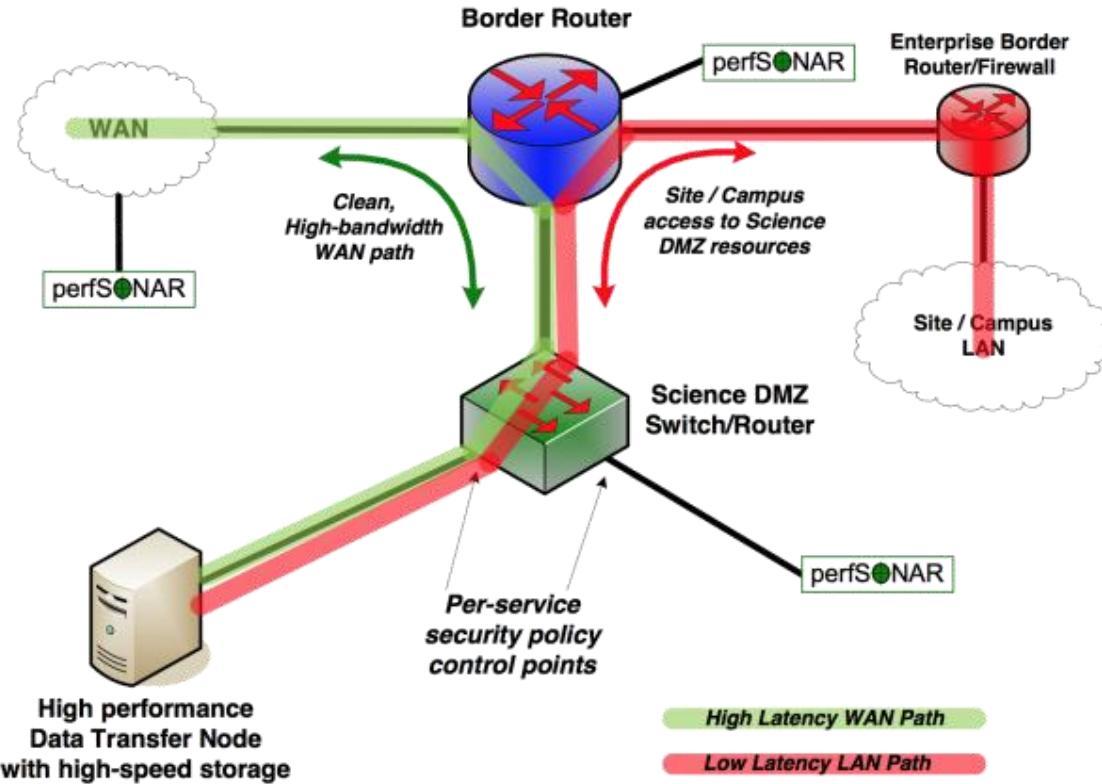
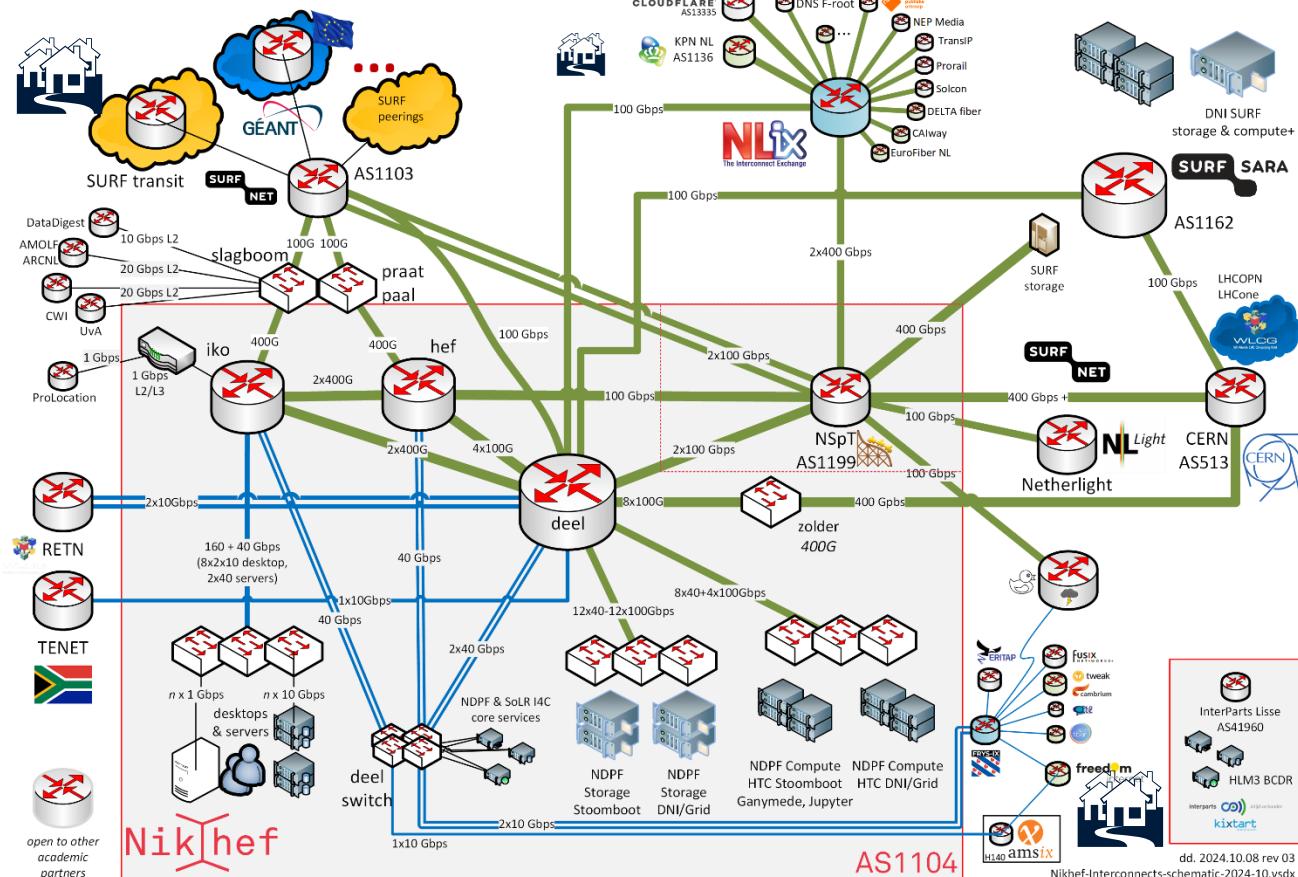


Image and ‘ScienceDMZ’ concept promulgated by ESnet (see fasterdata.es.net)

Just one random autonomous system: AS1104



AS1104
state as of Oct 2024



Nikhef

Exercising the network – sensor data and events

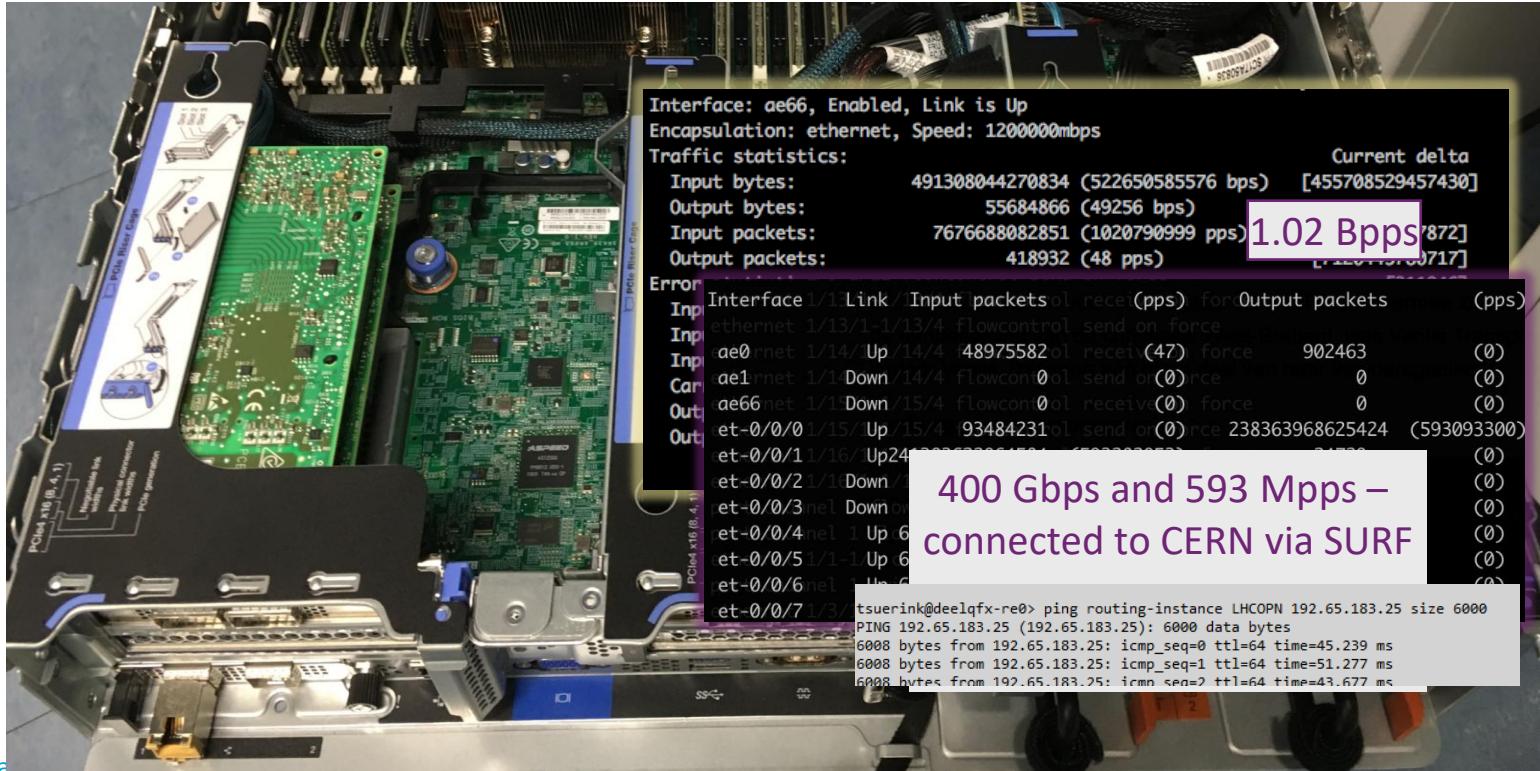


Image: <https://www.balawandurkhan.com/>, Photo: SURFnet

For example for HL-LHC, or SKA, more is needed > 2028 ...

- ‘Typical’ network is now mixed 400G-100G
- Push experiments to 800Gbps in metro area, and a local (AMS) loop has been demonstrated
- next: 800 → 1600G AMS-GVA ☺



Web screenshot: btg.org,

Images Nokia 7750-SR1x in Nikhef AMS H234b: Tristan Suerink



[Home](#) | [BTG](#) | [BTG Services](#) | [INTUG](#) | [Innovatielab](#) | [Activiteiten](#) | [Lobby & Opinie](#) | [Publicaties](#)

Minister Adriaansens opent testomgeving voor volgende generatie netwerktechnologieën

In Amsterdam is door minister Micky Adriaansens van Economische Zaken en Klimaatierotonde is een testomgeving waar SURF en Nikhef gaan experimenteren met nieuwe beschik over een internetsnelheid van 800 Gbit/s, wat meer dan 1000 keer sneller gemiddeld huishouden in Nederland. De innovatiertotonde stelt Nederlandse te doen naar de volgende generatie netwerktechnologieën.

en onderzoek naar bandbreedte op het internet groeit. Onderzoekers willen steeds meer over de landsgrenzen heen met elkaar delen. De bandbreedte van het netwerk speelt grote hoeveelheden data snel te kunnen verwerken, en de verwachting dat 8000bit/s. De innovatiertotonde maakt het mogelijk om te experimenteren met nieuwe

Scaling data access: ‘system-aware design’ at application layer

Reading data ‘scattered’ in a file - simply using POSIX-like IO - when done over the network severely exposes latency

and TCP slow-start makes that even worse

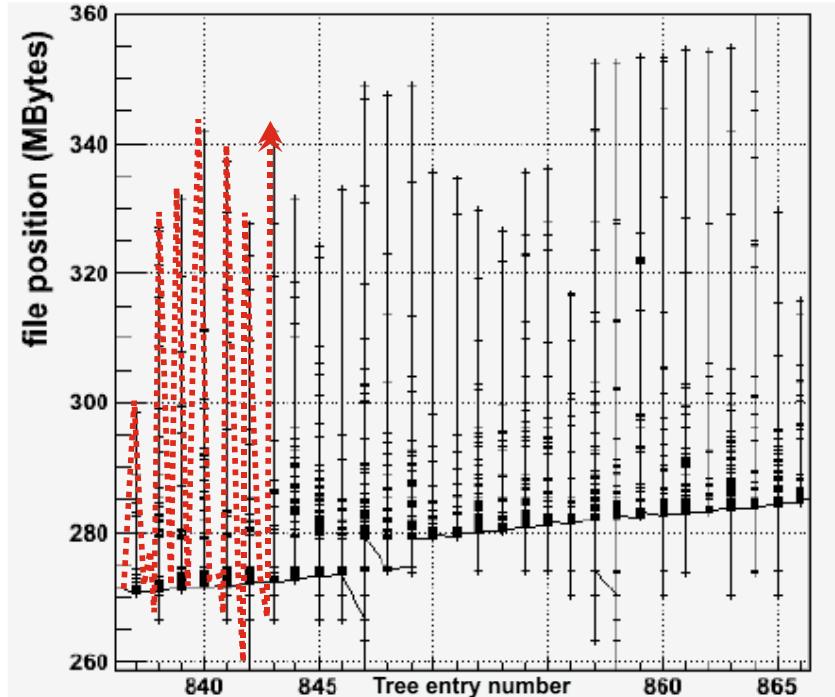
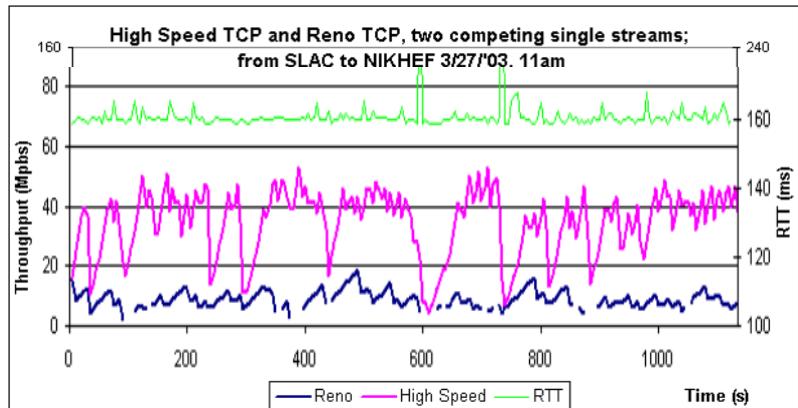


Image of TCP slow-start and packet loss impact (in Mpps): Antony Antony et al., Nikhef, for DataTAG, 2003(!)

Right: base graphic: Philippe Canal “Root I/O: the fast and the furious”, CHEP2010 Access pattern reflects Root versions < 5.28, before Ttree caching and ‘baskets’

And some traffic is triggered by researchers scaling up 'accidentally' from a laptop to a cluster without too much thought

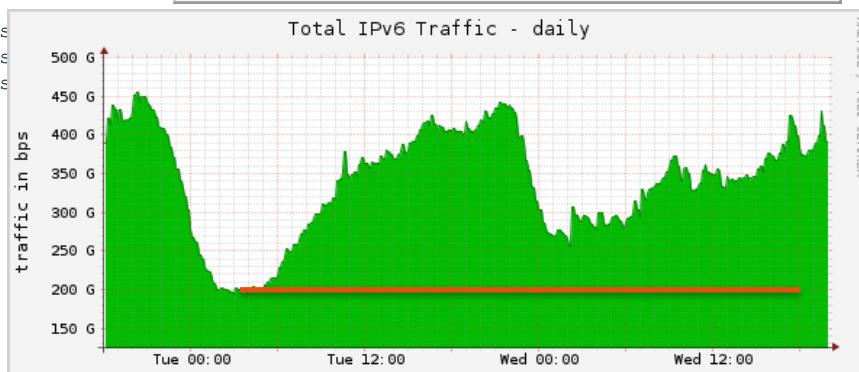
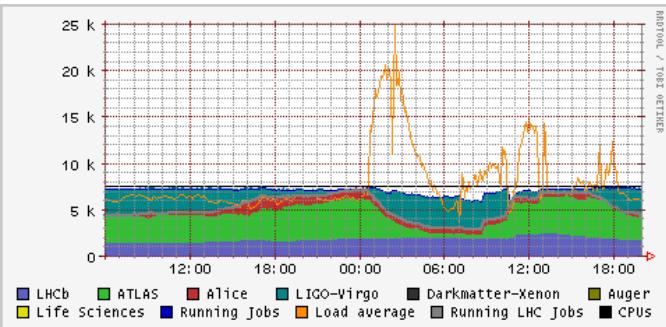
A researcher doing mass creation of containers, rebuilding their python 'virtual env' for each job, running on >> 4000 cores

```
[root@wn-pep-002 ~]# top
top - 09:40:47 up 71 days, 12:17, 2 users, load average: 110.38, 101.43, 106.3
Tasks: 700 total, 7 running, 666 sleeping, 0 stopped, 27 zombie
%Cpu(s): 17.0 us, 2.0 sy, 0.0 ni, 81.0 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 39462902+total, 23514457+free, 10406320 used, 14907812+buff/cache
KiB Swap: 67108860 total, 66841340 free, 267520 used. 37964784+avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
82661	ligo000	20	0	5618756	396356	924	R	360.0	0.1	5:14.43	mksquashfs
72615	ligo000	20	0	5626336	248516	816	R	90.0	0.1	5:44.11	mksquashfs
83257	ligo000	20	0	5611608	219300	852	S	90.0	0.1	1:17.66	mksquashfs
...											



Pulling the python packages at line rate and downloading public python repositories ultimately will trigger Cloudflare and flood SURFnet



June 28th, 2023, data from Nikhef NDFP stats & cricket (top), SURFnet asd001b-jnx-01 to asd001b-jnx-04 (left), AMS-IX SFlow <https://stats.ams-ix.net/sflow/index.html> (bottom)

Research data traffic looks like ... a DDoS to others 😊



Belastingdienst

Home Menu Zoeken

Home > Actueel > ICT en informatievoorziening > De systemen testen dankzij een unieke samenwerking

Lees voor

De systemen testen dankzij een unieke samenwerking

Dinsdag 14 maart 2023 | Het laatste nieuws het eerst op NU.nl



Forse ddos-aanvallen en nerdgrapjes tijdens nachtelijke oefening overheid

Door Rutger Otto

12 feb 2023 om 05:02
Update: een maand geleden

202 reacties

Het begon in 2018. Een bijzondere samenwerking tussen overheden, internetproviders- en exchanges, academische instanties, non-profitorganisaties, de politie, de krijgsmacht, de krijgsmacht en de krijgsmacht. **Een goed begin** De voorbereidingen van de avond beginnen ver. Elke organisatie bepaalt welke systemen ze wil uitgevoerd worden. Het 'red team' is verantwoordelijk voor de verdediging. Eén van de partijen die aanwezig waren, Nikhef, geeft aan dat zij dit belangrijk vinden, omdat de motivatie.

Nikhef is het Nationaal instituut voor subatomicaire en nucleaire wetenschappen. Ze zijn verantwoordelijk voor een gigantische bandbreedte, wat noodzakelijk is voor de samenwerking.



Betastingsdienst

Home

Home > Aanslagen > Ik heb een DDoS aanslag ontvangen - wat nu?

Ik heb een DDoS aanslag op mijn netwerk ontvangen - wat nu?

U ontvangt een DDoS aanslag op uw netwerk, bijvoorbeeld omdat u vergeten bent werkende tegenmaatregelen te nemen. Er staat dan een geschat aantal pakketten per seconde op uw monitoring.

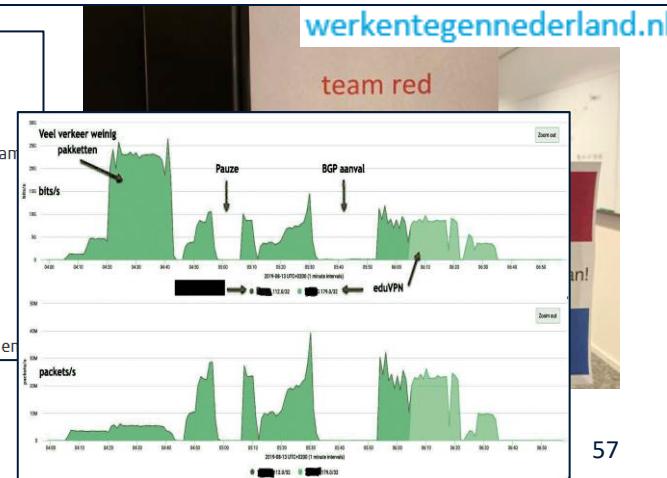
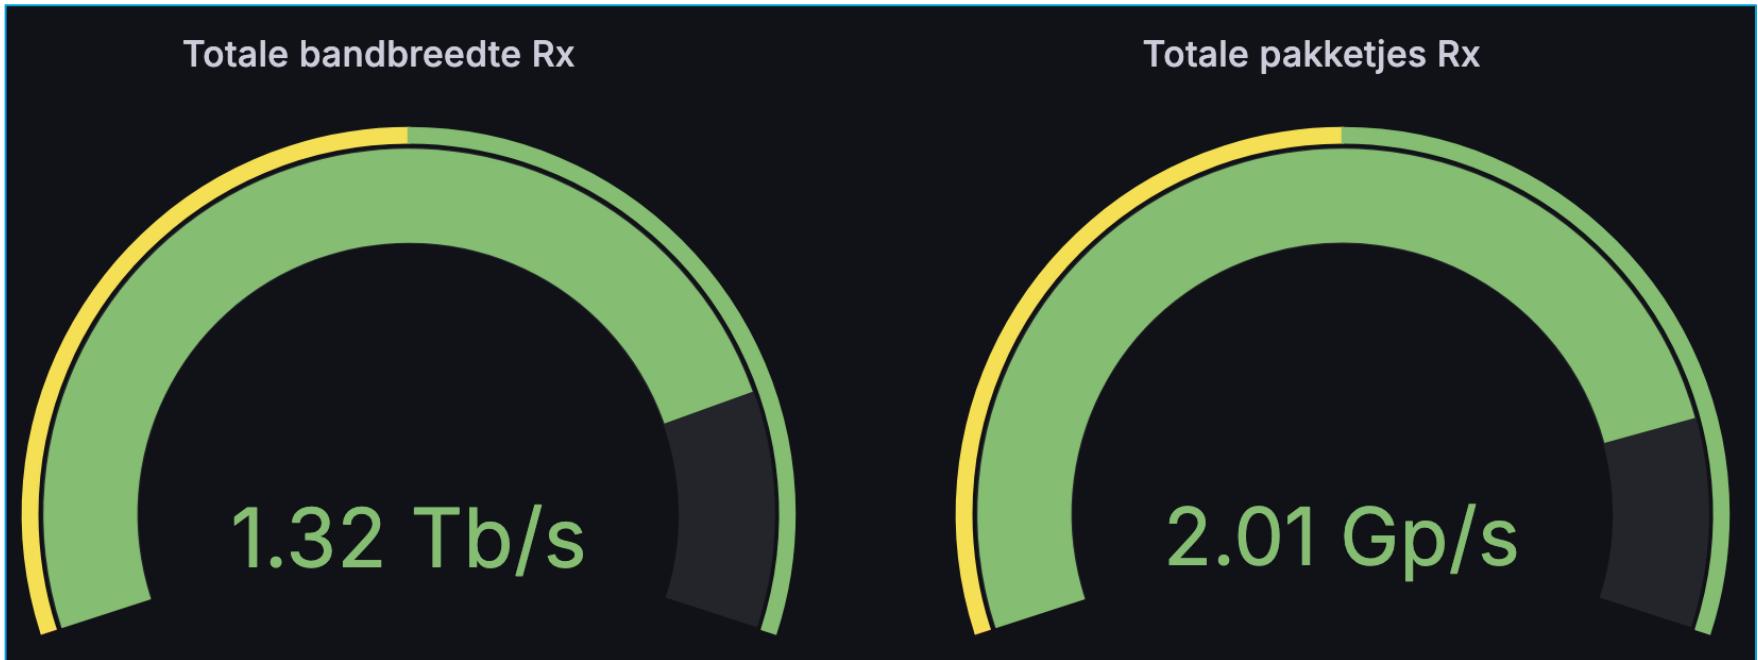


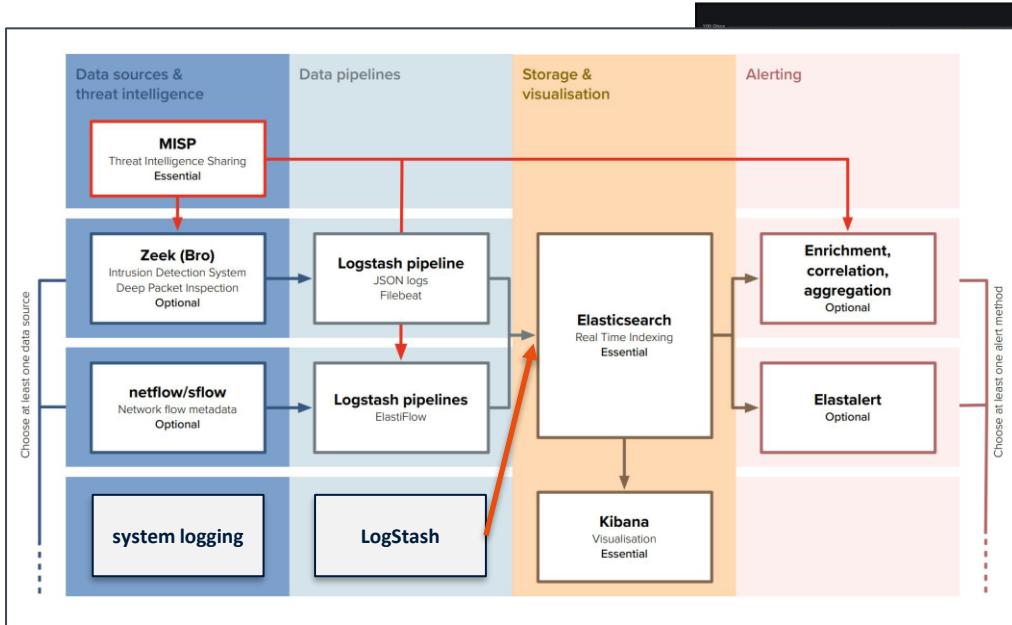
Image sources: belastingdienst.nl, rws.nl, nu.nl

with packets being more destructive than bandwidth ...



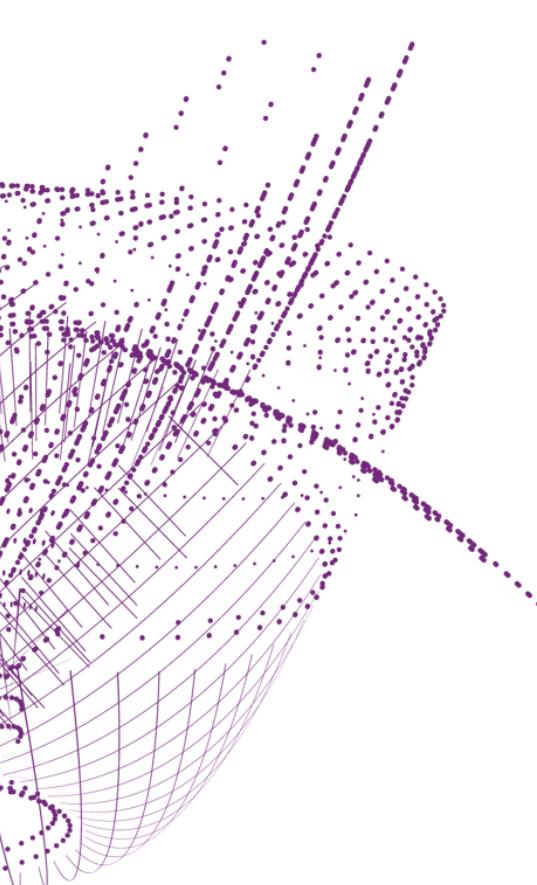
https://wiki.nikhef.nl/grid/2Bpps_Machine - in preparation for the 2025 Resilience Exercise -

And 'open' does not mean 'insecure' – the WLCG SOC model



650 GByte/day ingest; 100Gbps+ monitoring through optical taps and mirroring; MISP intel from CERN, SURF, and private intel sources

Nikhef SOC design/management by Jouke Roorda. WLCG SOC WG model: Liviu Valsan (CERN) and David Crooks (STFC RAL)



In Infrastructure We Trust

Infrastructure: for the small and the large



Collaborations: from small ...



Nikhef user room H1.37 – terminal stations in the early 1990's – image source: Nikhef

... to large collaborations (and shown here is a subset ...)



a small part of the CMS collaboration in 2017, photo credit CERN on behalf the CMS collaboration, CMS-PHO-PUBLIC-2017-004-3

How many interactions? And just how many logins?



Worldwide LHC Computing Grid (~ 2024)
~ 1.4 million **cpu** cores
~ 1500 Petabyte disk + archival

170+ institutes
42+ countries
13 ‘Tier-1 sites’
some multi-community:
NL-T1 @ SURF & Nikhef

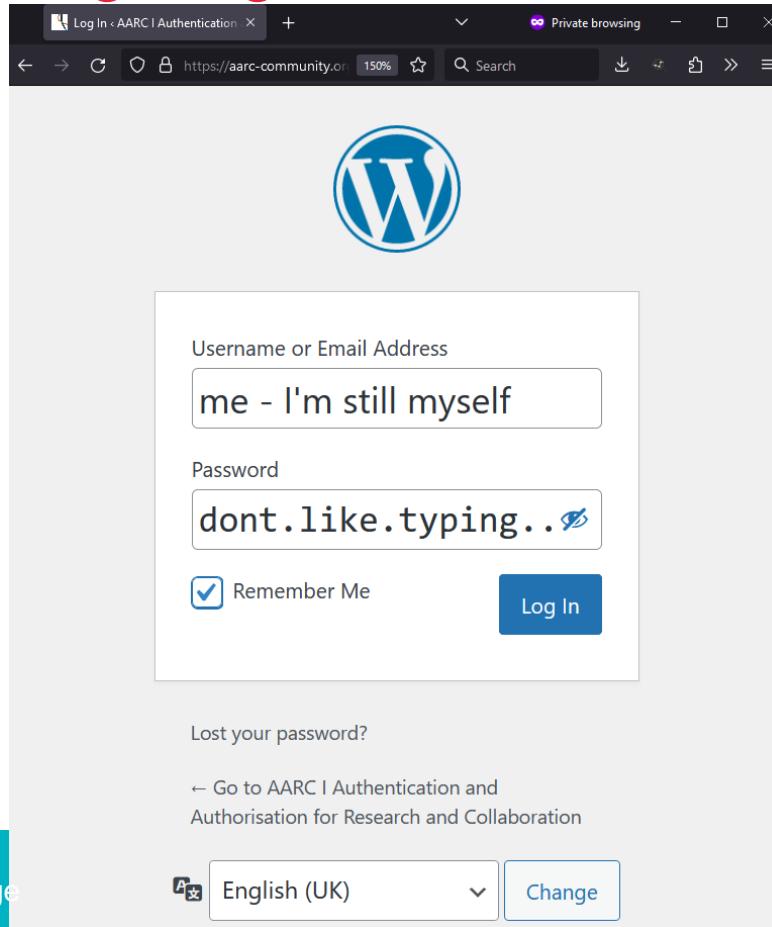
Earth background: Google Earth; Data and compute animation: STFC RAL for WLCG and EGI.eu; Data: <https://home.cern/science/computing/grid> ; LHC Computing Grid: wlcg.web.cern.ch, EGI: www.egi.eu; ACCESS CI: <https://access-ci.org/>, NL-T1 and FuSE: fuse-infra.nl, <https://www.surf.nl/en/research-it>

When you are asked to login again ... 12 000 x 170+ times?

Authentication

demonstrating 'you are you'

- ***authenticator***
'you' remains same 'you'
- ***vetted identity***
'you' can be pseudonymous
'you' can be a vetted person



Self-asserted or 'pseudonymous' often not enough



NATIONAAL INSTITUUT VOOR KERNFYSICA EN HOGE-ENERGIEFYSICA

Guest / students form (please tick)

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

CERN/User Registration

CERN COMPUTER CENTRE - US

<http://cern.ch/documents/ComputerUsage/CompA.html>

To be returned to the User Registration box at the end of the experiment. It is completed by a user who requires a computer account in the CERN Computer Department, and is not yet registered in another group.

To be completed by the User:

It is MANDATORY to provide the following information in a manner that is treated confidentially and only be used for ensuring the security of the experiment. Supply name as registered by the Users' Office.

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:	Stkrn:	Family:	
NON-473:	Sensitive:	Verifier:	Date:



Scaling credentials: per service per user

Many start with *credentials* dedicated to each service where you need access

- In a multi-organizational system becomes

$$\mathcal{O}(n_{\text{services}}) * \mathcal{O}(n_{\text{users}})$$

- usually creates a strong link to authorization:

*different accounts for different roles,
multiplying the number of credentials per user*

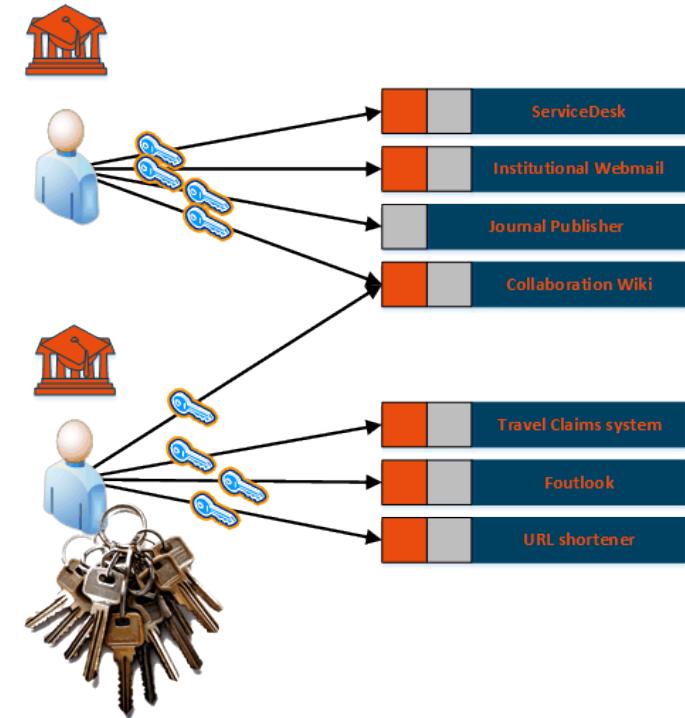


Image inspired by AARC NA2 training module “Authentication and Authorisation 101” – keychain image created by generative AI

bilateral 'SSO': a single service, or a single identity source

#credentials required?

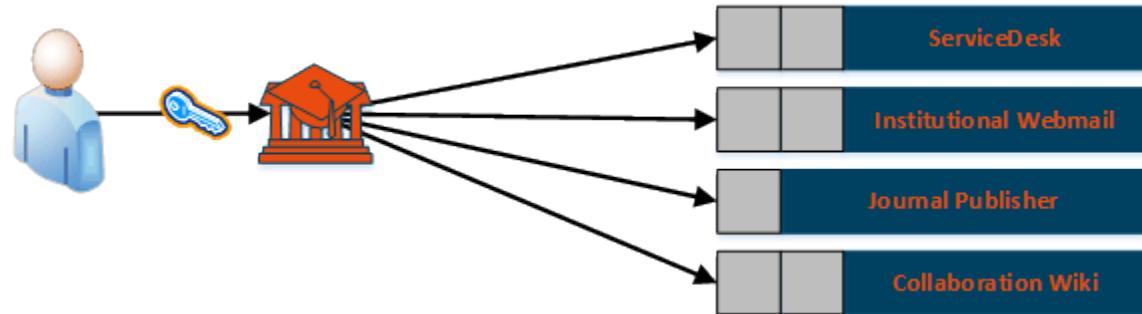
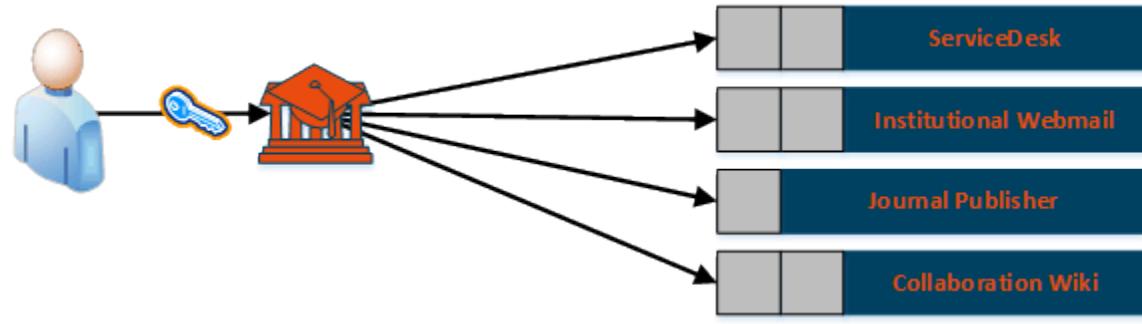
from previously

$$\mathcal{O}(n_{\text{services}}) * \mathcal{O}(n_{\text{users}})$$

to

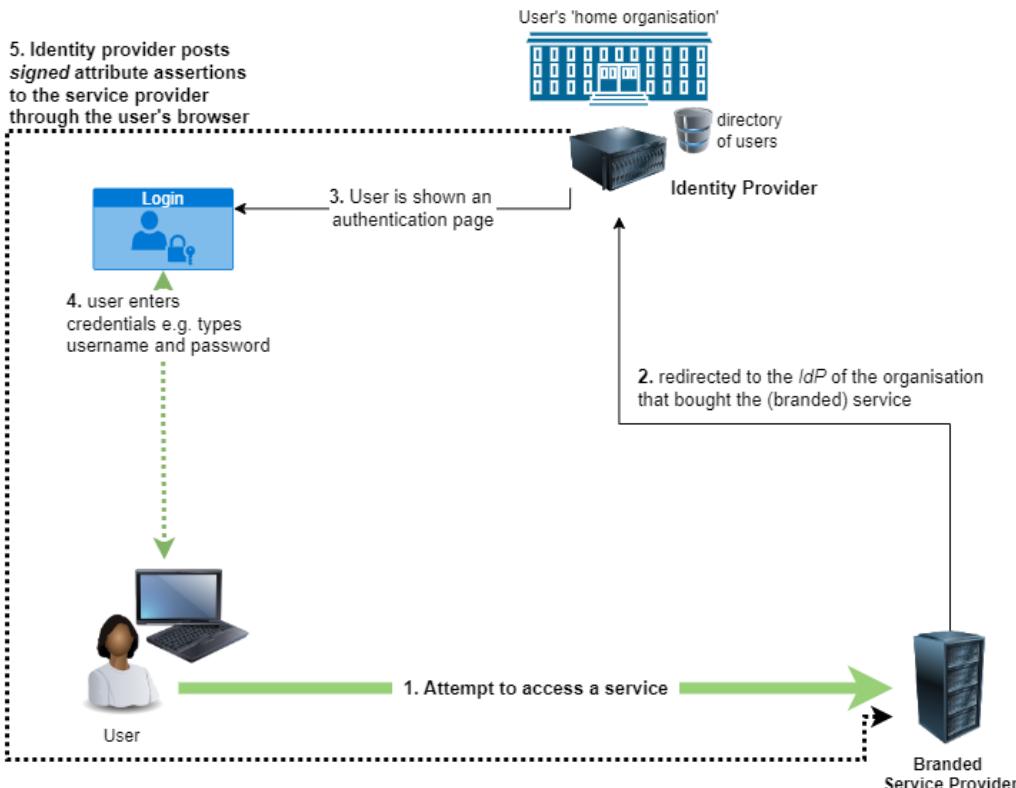
$$\mathcal{O}(n_{\text{users}}) + \mathcal{O}(n_{\text{services}} * n_{\text{home-orgs}})$$

in first order at least



Single sign-on – why your browser keeps loading things

5. Identity provider posts *signed* attribute assertions to the service provider through the user's browser



Extension: (SAML-tracer) - SAML-tracer — Mozilla Firefox

X Clear II Pause Autoscroll Filter resources Colorize Export Import

GET https://commute.nikhef.nl/
GET https://commute.nikhef.nl/favicon.ico
GET https://commute.nikhef.nl/commute/?auth=nikhef-sso
GET https://sso.nikhef.nl/sso/saml2/idp/SSOService.php?SAMLRequest=fVJL SAML
GET https://sso.nikhef.nl/sso/module.php/nikhef/loginuserpass.php?AuthState=_9d4f7
GET https://sso.nikhef.nl/sso/module.php/consent/getconsent?StatId=_9d4f753ffc12c
GET https://sso.nikhef.nl/sso/resources/icons/favicon.ico
GET https://sso.nikhef.nl/sso/module.php/consent/getconsent?saveconsent=1&StatId=
POST https://commute.nikhef.nl/simplesaml/module.php/saml/sp/saml2-ac.php SAML
GET https://commute.nikhef.nl/commute/?auth=nikhef-sso
GET https://commute.nikhef.nl/favicon.ico

Glossary

'SAML' is the "Security Assertion Mark-up Language" an XML blob with information, usually digitally signed

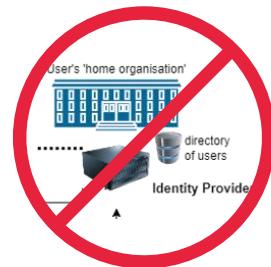
HTTP Parameters SAML Summary

Version="2.0"
IssueInstant="2025-02-28T11:49:04Z"
>
<saml:Issuer>https://sso.nikhef.nl/sso/saml2/idp/metadata.php</saml:Issuer>

SAML-tracer plugin by Tim van Dijken (SSC-ICT) et al.
<https://github.com/simplesamlphp/SAML-tracer>

User-centric identity: 'I take my passport anywhere by myself'

Your 'home organisation' does not have to be in the loop ...

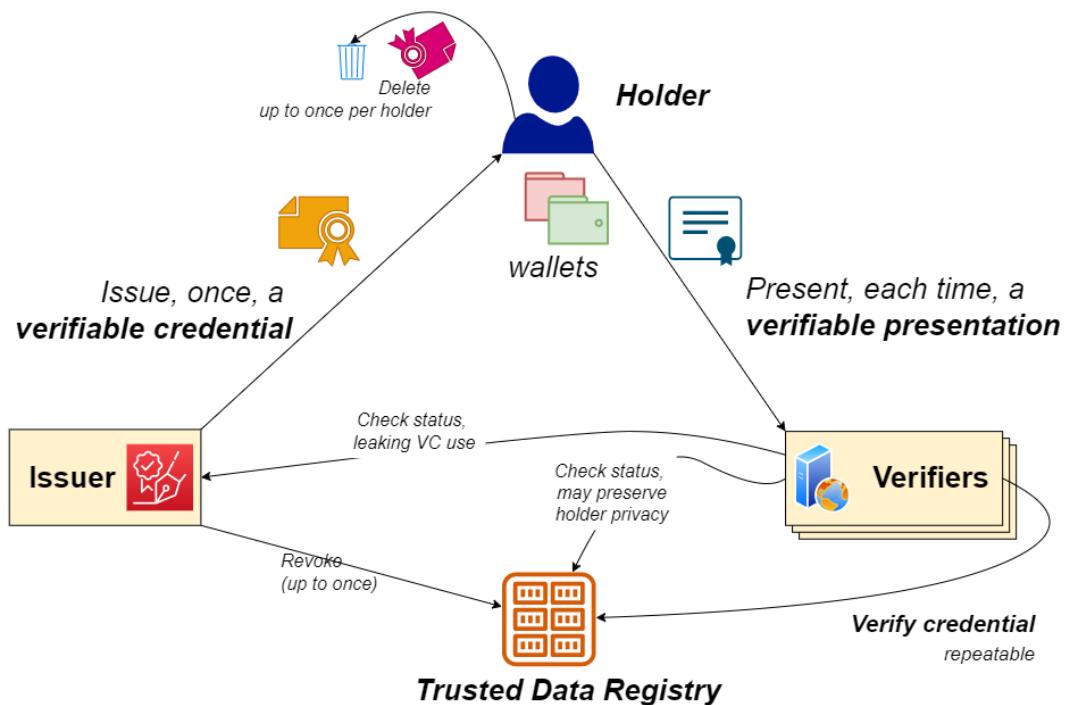
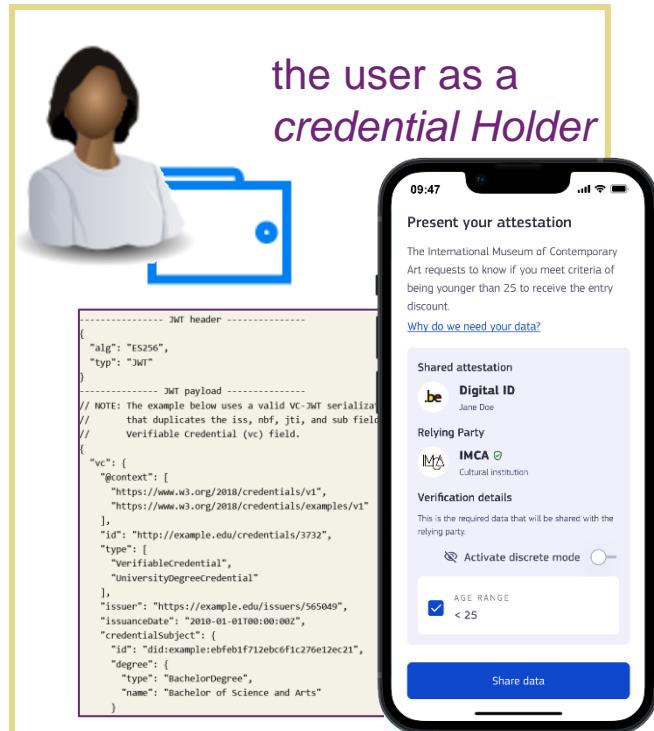


*user-centric trust: you yourself hold a credential from a trusted third party and can use it *without having to ask 'home' each time*:*

- Public Key Infrastructure client certificates ("X.509")
- Verifiable credentials in wallets
- *and who remembers CardSpace?*

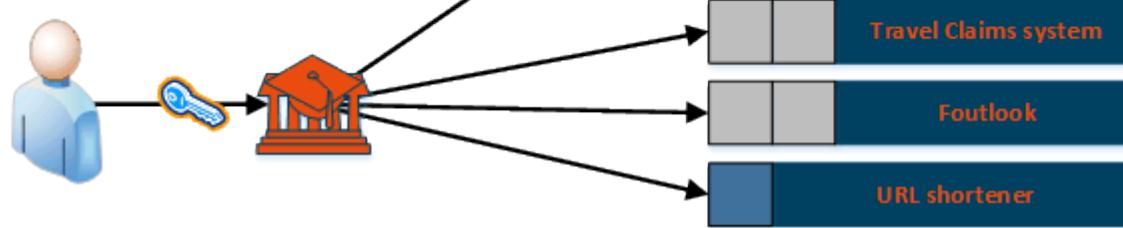
Passport image: cropped from original by Jon Tyson on Unsplash <https://unsplash.com/photos/Hid-yhommOg>

Identity wallets, held by the user, are another



Flow diagram inspired by: Lifecycle Details (5.1), Verifiable Credentials Data Model v1.1, W3C Recommendation 03 March 2022, <https://www.w3.org/TR/vc-data-model/>
EU eID Wallet from https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en
Annimage: European Commission, at <https://ec.europa.eu/digital-building-blocks/sites/display/EUDIGITALIDENTITYWALLET/Security+and+Privacy>

Can we scale better with an ‘federated’ Authentication and Authorisation Infrastructure (‘AAI’)



Where are 'you' in the federated space – discovery!

The diagram illustrates the cross-institutional service flow from HARICA to various institutions. It shows a sequence of five screens:

- HARICA Login Page:** Shows a 'Login' form with fields for 'Email address' and 'Password', and a red box around the 'Academic Login' button.
- Choose Your Institution Page:** Shows a list of recent institutions: Nikhef, CERN, and Maastricht University. The 'Academic Login' button is also present here.
- Nikhef Single Sign-On Page:** Shows the 'Enter your Nikhef username and password' form.
- CERN Single Sign-On Page:** Shows the 'CERN Single Sign-On' form with fields for 'Username' and 'Password'.
- Maastricht University Single Sign-On Page:** Shows the 'Maastricht University' form with fields for 'Username' and 'Password'.

Arrows indicate the flow from the HARICA login to the discovery page, then to each institution's SSO page. A purple box highlights the 'Edit' button on the discovery page.

An example cross-institutional service by HARICA, the GEANT TCS G5 provider, presenting a SeamlessAccess.org discovery page

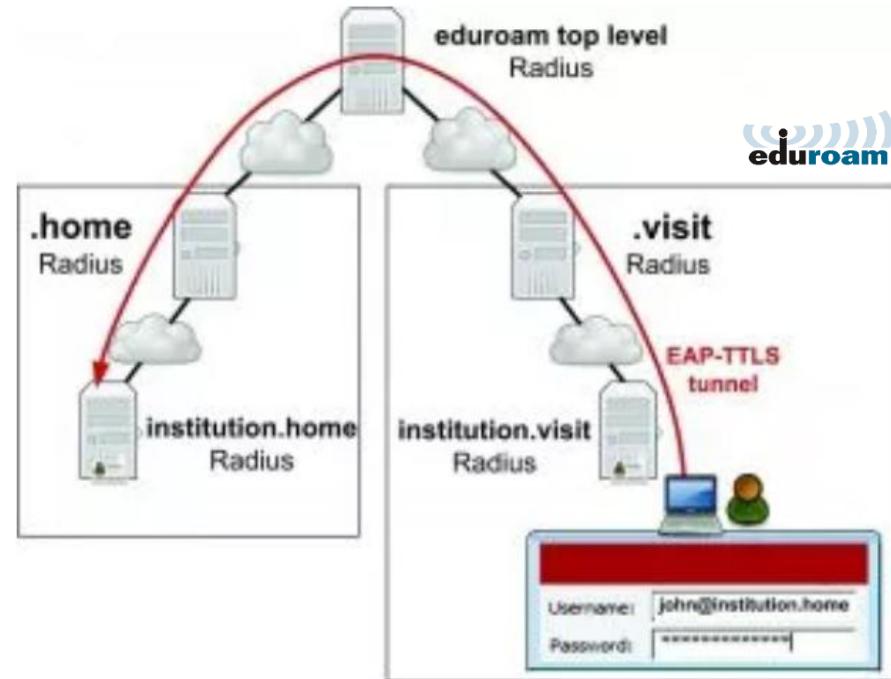
The R&E federation that was there first ...

service-specific trust
between organisations

hierarchical server path, based on
a network-specific secure exchange

sending your credentials back
to *only* your home institution

found via <anon@domain.name>



eduroam image from <https://eduroam.org/how/>, GEANT ; RADIUS: RC2865 <https://www.rfc-editor.org/rfc/rfc2865>; see also freeradius.org

IGTF: a policy-bridged global federation for research computing

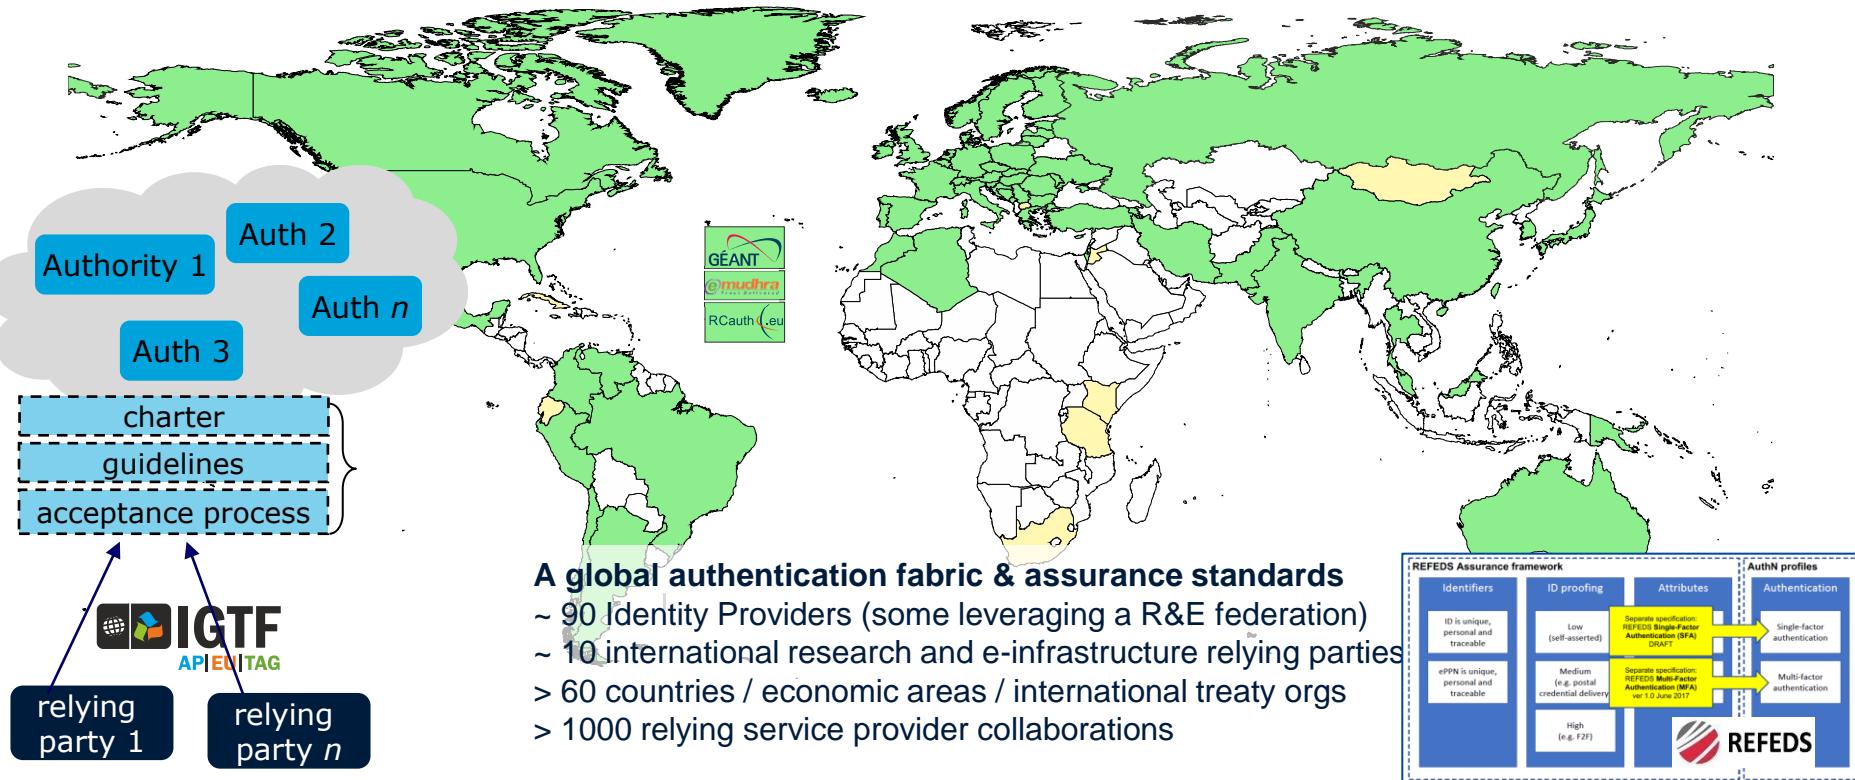
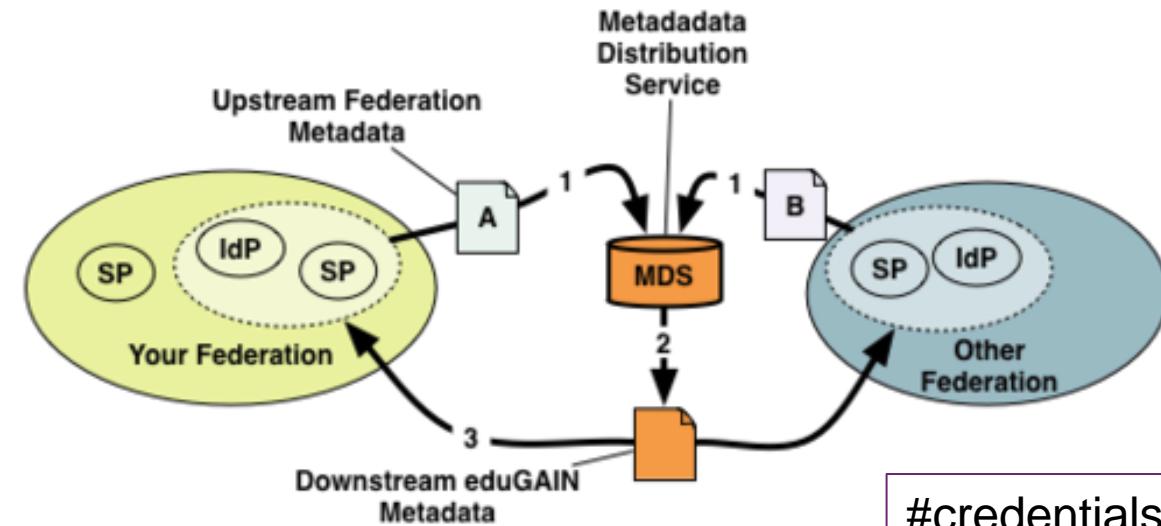


Image: Interoperable Global Trust Federation IGTF, <https://igtf.net/>; REFEDS Assurance Framework RAF: <http://refeds.org/assurance>, <https://refeds.org/profile/mfa>

We live in a federated world!



Meta-data and trust in IdP-SP 'multi-lateral' federations

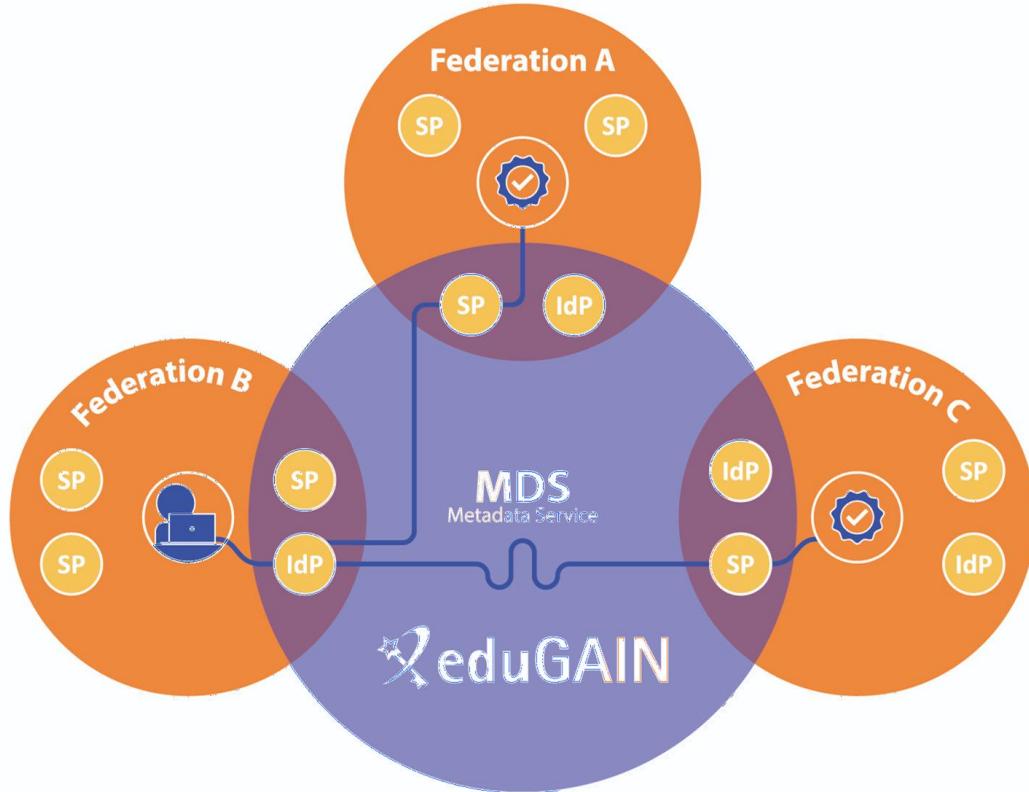


Listing of all entities (1972), 42 federations		
Entity ID	Entity details	ECCS
IdP 29 Mayis University	Entity ID: https://kimlik.29mayis.edu.tr/simpleسامي/saml2/idp/metadata.php Entity categories: Data Protection Code of Conduct v1; REFEDS Research and Scholarship Registrar: YETKİM Org: İstanbul 29 Mayıs University	ECCS
IdP A*STAR - Agency for Science, Technology and Research	Entity ID: https://a-star.singaren.net.sg/simpleسامي/saml2/idp/metadata.php Entity categories: REFEDS Research and Scholarship Registrar: Singapore Access Federation - SGAF Org: A*STAR - Agency for Science, Technology and Research	ECCS
IdP AAF Virtual Home	Entity ID: https://vbaaf.edu.auiidp/shibboleth Entity categories: REFEDS Research and Scholarship SIRTFI asserted Registrar: AAF Org: AAF Virtual Home	ECCS
IdP AARNet	Entity ID: https://shibboleth.aarnet.edu.au/idp/shibboleth Entity categories: REFEDS Research and Scholarship SIRTFI asserted Registrar: AAF Org: Australian Academic and Research Network (AARNET)	ECCS
IdP ACCESS CI	Entity ID: https://access-ci.org/idp Entity categories: REFEDS Research and Scholarship; http://id.incommon.org/category/registered-by-incommon SIRTFI asserted Registrar: InCommon Org: National Center for Supercomputing Applications	ECCS
IdP ACOnet staff	Entity ID: https://idp.aco.net/idp/shibboleth Entity categories: Data Protection Code of Conduct v1; REFEDS Research and Scholarship Registrar: ACOnet Identity Federation Org: ACOnet staff	ECCS
IdP AIRcentre	Entity ID: https://ids.aircentre.org/idp/shibboleth	ECCS

#credentials required?

from $\mathcal{O}(n_{\text{users}}) + \mathcal{O}(n_{\text{services}} * n_{\text{home-orgs}})$
to $\sim \mathcal{O}(n_{\text{users}}) + \mathcal{O}(n_{\text{home-orgs}}) + \mathcal{O}(n_{\text{services}})$

MDS meta-data flow: <https://wiki.geant.org/display/eduGAIN/Metadata+Flow+in+eduGAIN>
eduGAIN meta-data <https://mds.edugain.org/edugain-v2.xml> ; table excerpt from <https://technical.edugain.org/entities> showing only R&S IdPs, i.e. those supporting research ...



eduGAIN image: Davide Vaghetti, GARR for GN*-*

78

Identity Federations

5100+

Identity Providers

3600+

Service Providers

We progressed a lot since 2003 with identity federation

NIKHEF NATIONAAL INSTITUUT VOOR KERNFYSICA EN HOGE-ENERGIEFYSICA

Guest/ students form (please with a copy of your identity card)

CERN/User Registration
Date: 01.03.2004

CERN COMPUTER CENTRE - USER REGISTRATION FORM
<http://cern.ch/it/documents/ComputerUsage/CompAccountRegistrationForm-English.pdf>

To be returned to the User Registration box at the entrance of Building 513, after being completed by a user who requires a computer account in a Central Service provided by IT Department, and is not yet registered in another group or system or has already signed it before.

Fer

To be completed by the User:
It is **MANDATORY** to provide the following information (except those with an *). It will be treated confidentially and only be used for ensuring correct identification.
Supply name as registered by the Users' Office or HR Division.

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):

Name: SWIETZER
Last:
University or Institute: FLORIDA STATE

To be completed by the Group Administrator:

Experiment/Department:

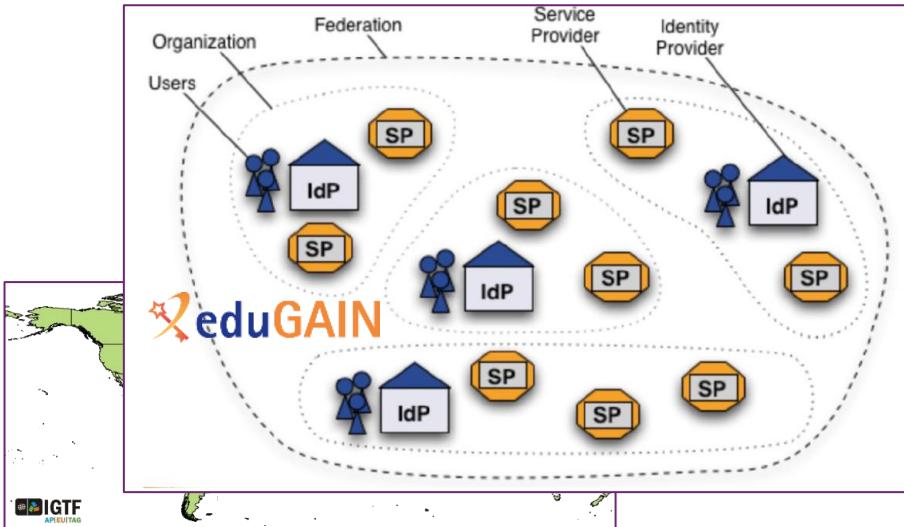
Exp / Dept.	Spokesperson	Home Institution Contact	Contact Telephone
D0	WOMERSLEY/WEERTS	SHARON HAGOPIAN	850-644-4777

CPNS

4 5 6 7 8 9 0 1 2 3 0 0 0

PRN 7 TUE 8 WED 9 0 0 0

CPNS 0 CLR ENT



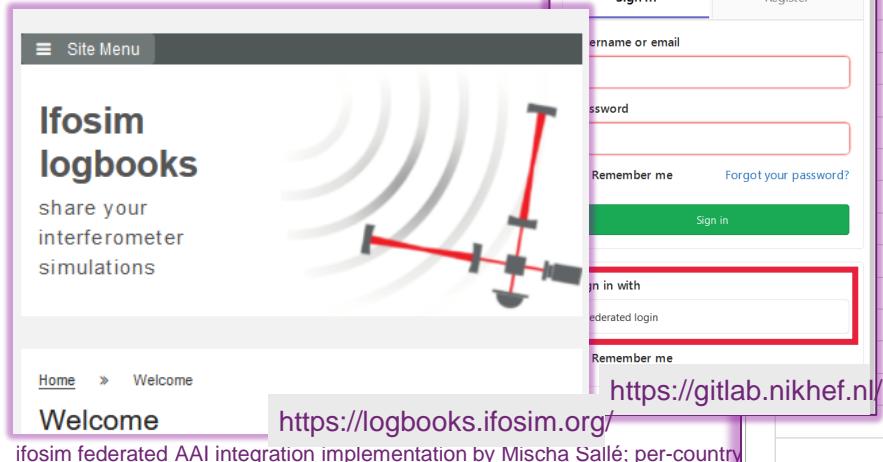
For eduGAIN federation the IdPs provide **authentication** from the home organisation, for the user-centric PKI IGTf trust fabric, the CAs do. Then **Service providers** perform **authorization**, ... maybe using attributes provided by the IdP. But do they get them??

Right-hand image: Shibboleth IdP federation, Lukas Hammerle, SWITCH (CH), user-centric PKI credentials: Interoperable Global Trust Federation, <https://igtf.net/>

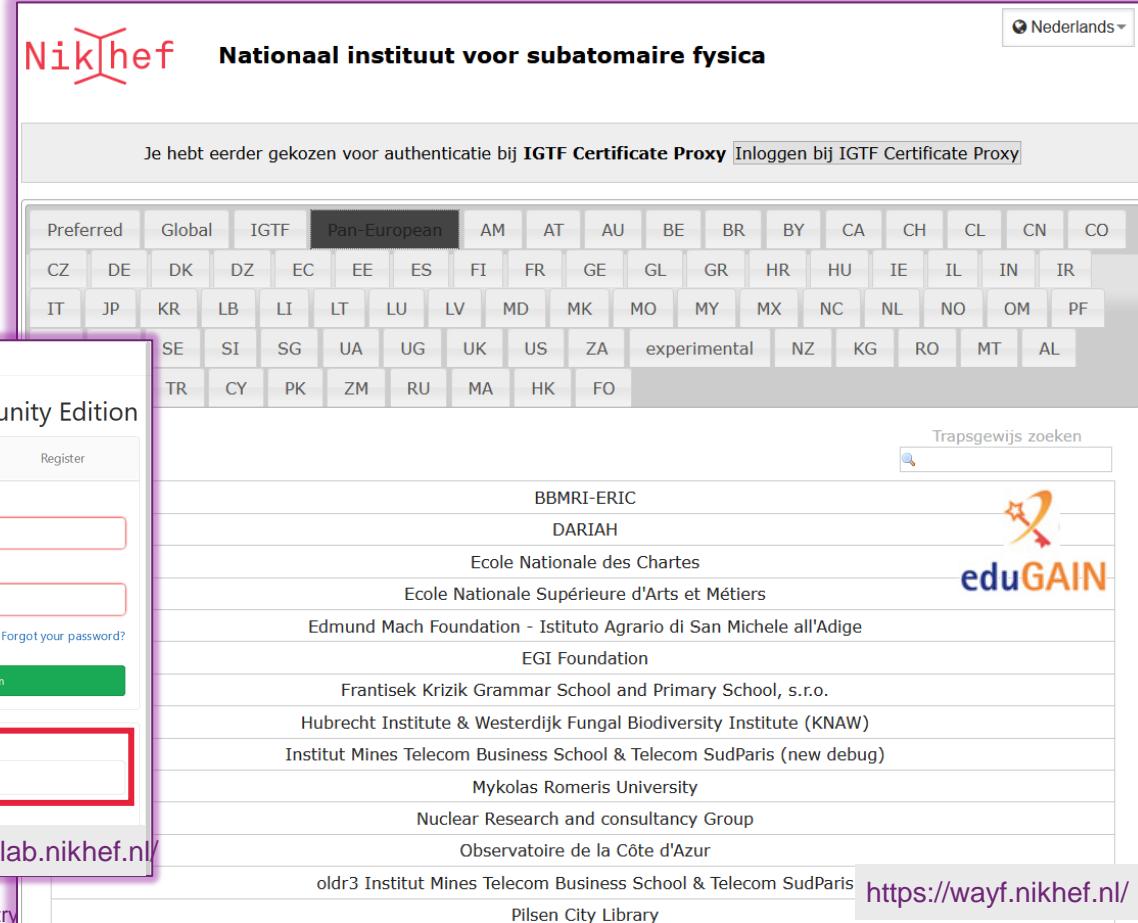
Federated Success!

Login to GW's ifosim.org, to gitlab, or ... via the service proxy

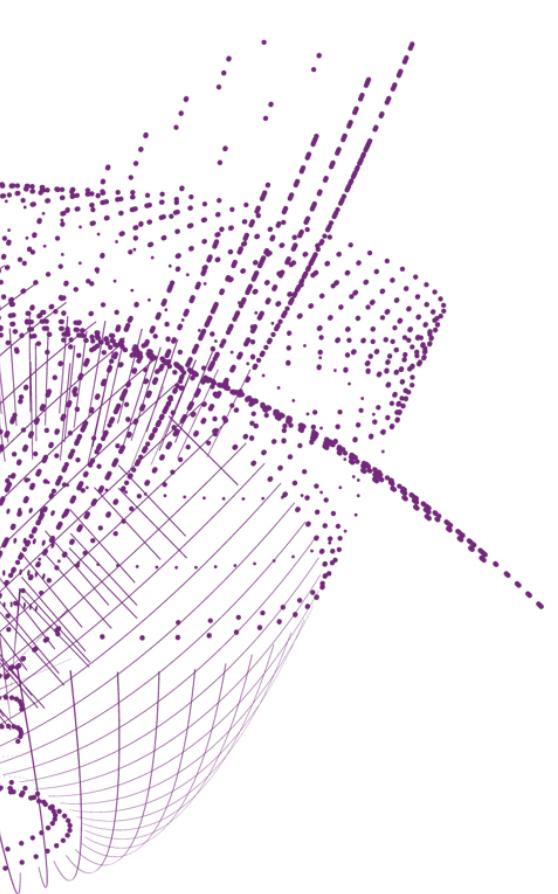
*with any eduGAIN IdP
for user authentication*



The screenshot shows the Ifosim logbooks website. The top navigation bar includes a 'Site Menu' button. The main content area is titled 'Ifosim logbooks' and describes 'share your interferometer simulations'. Below this is a 'Welcome' message and a URL 'https://logbooks.ifosim.org/'. At the bottom, it says 'ifosim federated AAI integration implementation by Mischa Sallé; per-country'. A red box highlights the 'Federated login' button on the login form, which is part of a larger 'GitLab Community Edition' interface.



The screenshot shows the Nikhef service proxy login interface. At the top, the Nikhef logo and the text 'Nationale instituut voor subatomaire fysica' are displayed. A button for 'Nederlands' is in the top right. Below this is a message: 'Je hebt eerder gekozen voor authenticatie bij **IGTF Certificate Proxy** [Inloggen bij IGTF Certificate Proxy](#)'. A grid of buttons for different authentication providers is shown, with 'Pan-European' selected. The grid includes: Preferred, Global, IGTF, Pan-European, AM, AT, AU, BE, BR, BY, CA, CH, CL, CN, CO, followed by a row of country codes: CZ, DE, DK, DZ, EC, EE, ES, FI, FR, GE, GL, GR, HR, HU, IE, IL, IN, IR, IT, JP, KR, LB, LI, LT, LU, LV, MD, MK, MO, MY, MX, NC, NL, NO, OM, PF, SE, SI, SG, UA, UG, UK, US, ZA, experimental, NZ, KG, RO, MT, AL, and TR, CY, PK, ZM, RU, MA, HK, FO. Below the grid is a search bar for 'Trapsgewijs zoeken'. The right side of the page lists various federated partners, including BBMRI-ERIC, DARIOH, Ecole Nationale des Chartes, Ecole Nationale Supérieure d'Arts et Métiers, Edmund Mach Foundation - Istituto Agrario di San Michele all'Adige, EGI Foundation, Frantisek Krizik Grammar School and Primary School, s.r.o., Hubrecht Institute & Westerdijk Fungal Biodiversity Institute (KNAW), Institut Mines Telecom Business School & Telecom SudParis (new debug), Mykolas Romeris University, Nuclear Research and consultancy Group, Observatoire de la Côte d'Azur, oldr3 Institut Mines Telecom Business School & Telecom SudParis, Pilzen City Library, and the eduGAIN logo.



Federation works quite well for Authentication, but ... not (federated) Authorisation – the important element for collaborative (research); with different complementary sources of authority, and decision power at the RP and its coordinating (e)-infrastructure

The Forgotten A in AAI

Authorization – what you are allowed to do

soon needs specifying **access rights** to resources, based on an access **policy**

- might be implicit or ad-hoc
- be in formal policy language like XACML (*example: Argus PDP*)
- or be service-specific
example: Linux sssd config

```
resource "http://cern.ch/authz/ce1" {
  action "http://cern.ch/authz/actions/ce-submit" {
    rule permit {
      vo="atlas"
      pilot-job="true"
    }
    rule deny {
      pilot-job="true"
    }
  }
}
```

simplified Argus policy language – can map directly to XACML



```
ldap_access_order = filter,authorized_service
ldap_access_filter = (|(memberOf=cn=gridSrvAdministrators,ou=DirectoryGroups,dc=farmnet,dc=nikhef,dc=nl)(memberOf=cn=gridMWSecurityGroup,ou=DirectoryGroups,dc=farmnet,dc=nikhef,dc=nl)(memberOf=cn=nDPFPrivilegedUsers,ou=DirectoryGroups,dc=farmnet,dc=nikhef,dc=nl))
```

Policy example: Argus system, <https://argus-documentation.readthedocs.io/en/stable/misc/examples.html>; service-specific: sssd.conf ldap auth_provider

Authorization policy subjects

AuthZ policies need subject attributes ('claims')

- **bound to an verifiable identity statement**
 - e.g. visa are strongly linked to a specific entity, and asserted by a trusted party (by the service)
- **be a bearer token**
 - scoped to a relying party, a service, or an action
- **self-asserted**
 - quite useless unless backed by verifiable evidence, like in self-sovereign identity schemes

Transport mechanisms (see also RFC2903)

- pushed alongside the service access,
- pulled from the source as needed, or
- pushed by the attribute source as an agent

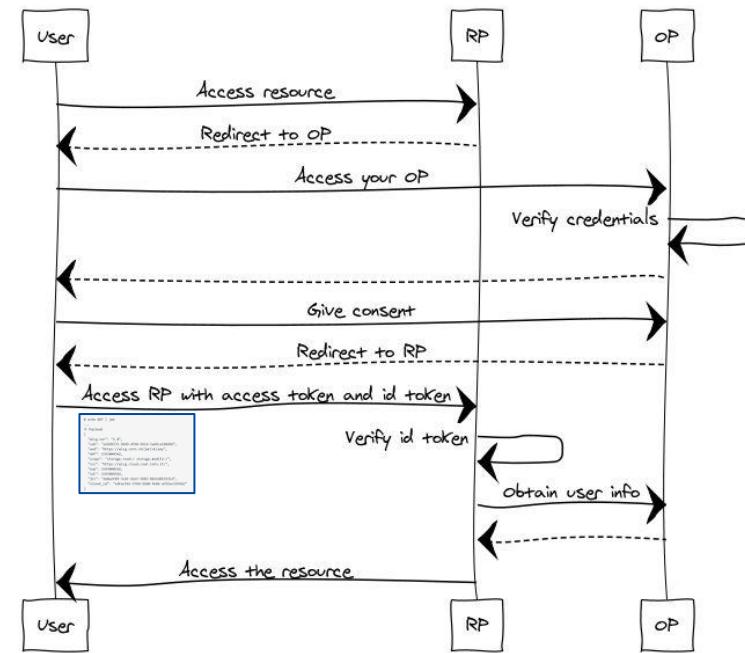
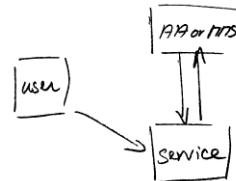
USA visa image source: <https://2009-2017.state.gov/m/ds/rls/rpt/79785.htm> ; RATP bearer token, issued for the Paris public transport system



OpenID Connect and OAuth2: the ‘modern’ way

- Quite well-known
(used by lots modern ‘non-enterprise’ SSO)
- shows signs of its initial design objective:
one source of identity (Openid Provider, ‘OP’), and many services (Relaying Parties, ‘RP’)

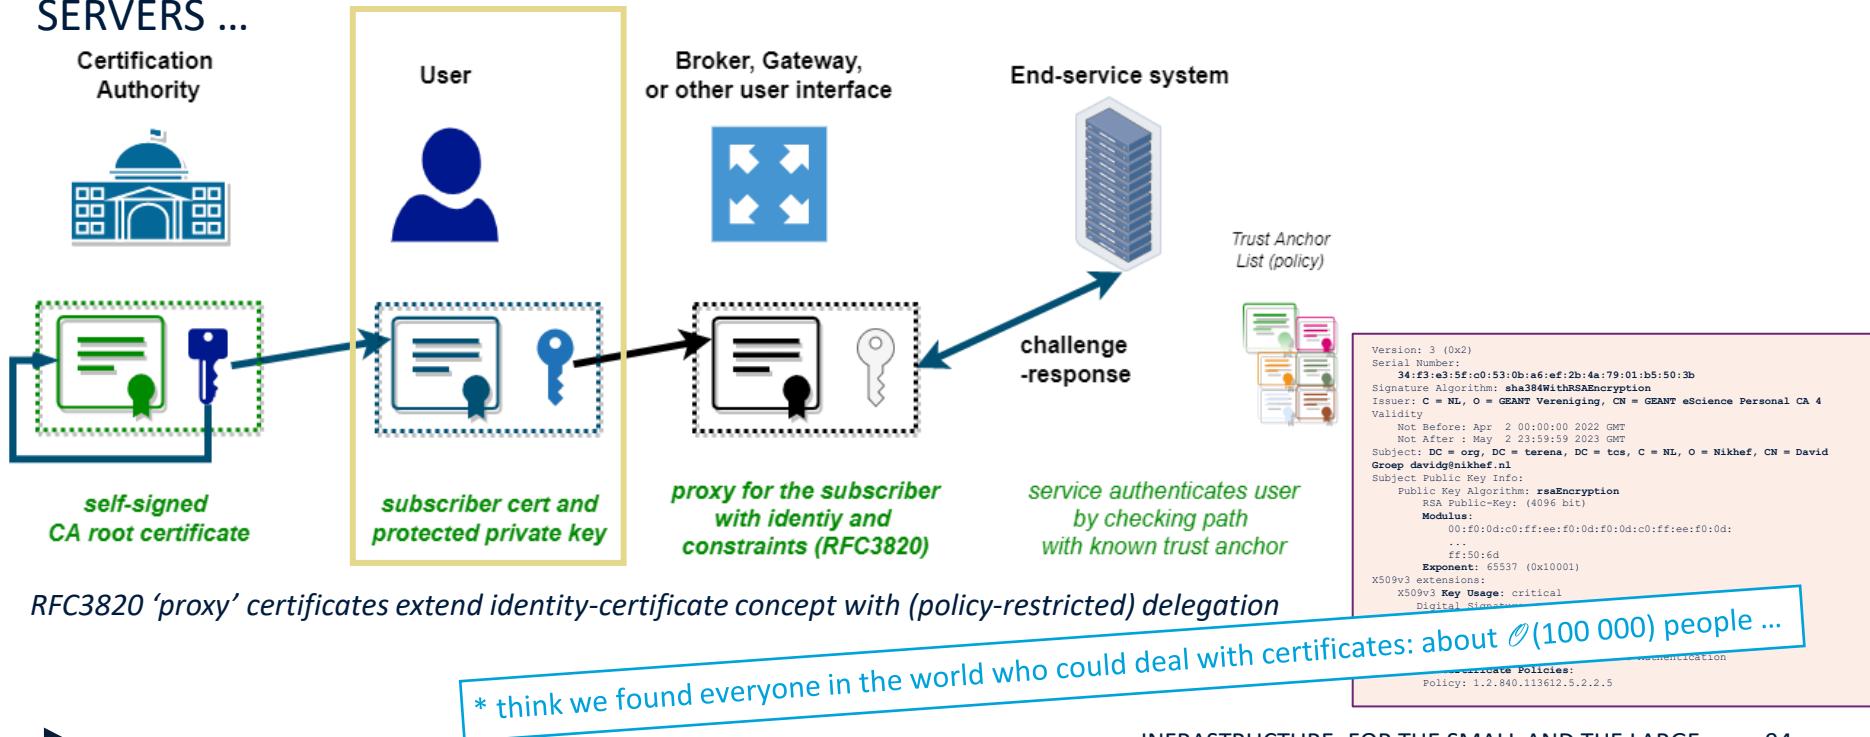
Show OpenID Connect Client	
Name	hekel.nikhef.nl
Description	Hekel using mod_auth_openidc
Client id.	_f6bfe81892e680e4ecfc3b41ecf1a15d141c0d106b
Client secret	██
Auth. source	saml2
Redirect URI	https://hekel.nikhef.nl/rp/redirect_uri
Scopes	openid profile email assurance
	Return Reset secret



Shown is the ‘implicit flow’, other flows possible. Image source: AARC NA2 training on AAI 101
See <https://openid.net/> for protocols and standardization work

PKI client certificates – user* *client held* credentials

YOU HAVE SEEN *HTTPS*, BUT THE SAME PKI CERTIFICATES CAN BE USED FOR CLIENTS, NOT SERVERS ...



Different tech, also an AAA push concept: X.509 and a trust PKI

Version: 3 (0x2)

Serial Number:

34:f3:e3:5f:c0:53:0b:a6:ef:2b:4a:79:01:b5:50:3b

Signature Algorithm: **sha384WithRSAEncryption**

Issuer: **C = NL, O = GEANT Vereniging, CN = GEANT eScience Personal CA 4**

Validity

Not Before: Apr 2 00:00:00 2022 GMT

Not After : May 2 23:59:59 2023 GMT

Subject: **DC = org, DC = terena, DC = tcs, C = NL, O = Nikhef, CN = David Groep davidg@nikhef.nl**

Subject Public Key Info:

Public Key Algorithm: **rsaEncryption**

RSA Public-Key: (4096 bit)

Modulus:

00:f0:0d:c0:ff:ee:f0:0d:f0:0d:c0:ff:ee:f0:0d:

...

ff:50:6d

Exponent: 65537 (0x10001)

X509v3 extensions:

X509v3 **Key Usage:** critical

Digital Signature, Key Encipherment

X509v3 Basic Constraints: critical

CA:FALSE

X509v3 **Extended Key Usage:**

E-mail Protection, TLS Web Client Authentication

X509v3 **Certificate Policies:**

Policy: 1.2.840.113612.5.2.2.5

You should be able to get an 'IGTF-DOGWOOD' assurance certificate from RCauth.eu.

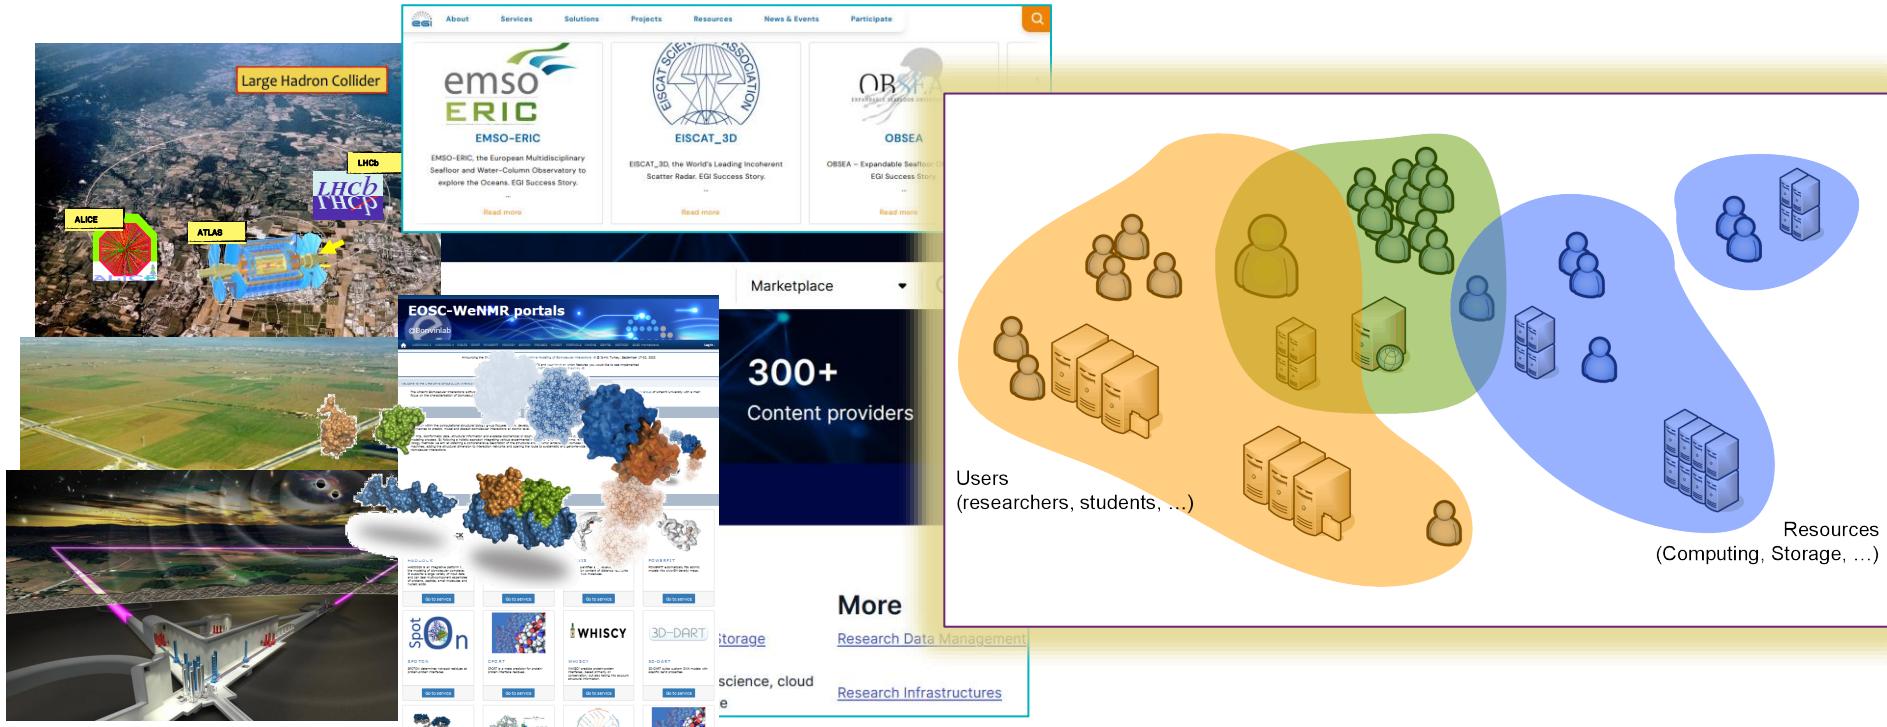
Go to <https://rcdemo.nikhef.nl/> and select the 'Basic demo' and use 'run non-VOMS' to get and view your short-lived certificate

are back-channel interactions

run non-VOMS demo



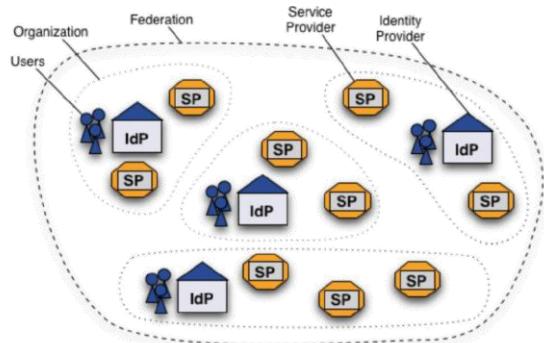
Science infrastructures using our R&E 'federated access'



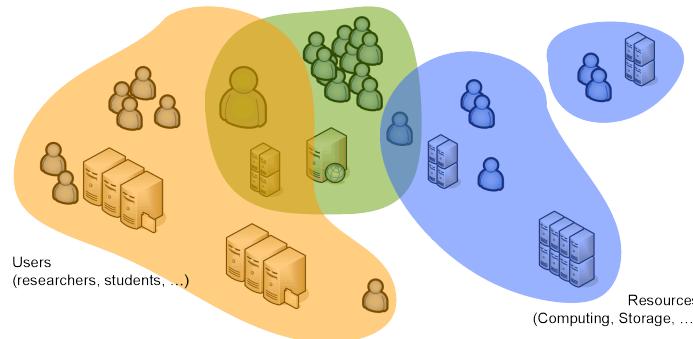
Images: CERN <https://wlcg.web.cern.ch/>; HADDOCK, WeNMR, @Bonvinlab <https://wenmr.science.uu.nl/>; Virgo, Pisa, IT; artist impression Einstein Telescope EMR region; EOSC portal in 2023, EGI catalogue <https://www.egi.eu/>

They look similar, yet they are not ...

In the **Identity federation** picture,
the source of authority is the
home organisation via its IdP



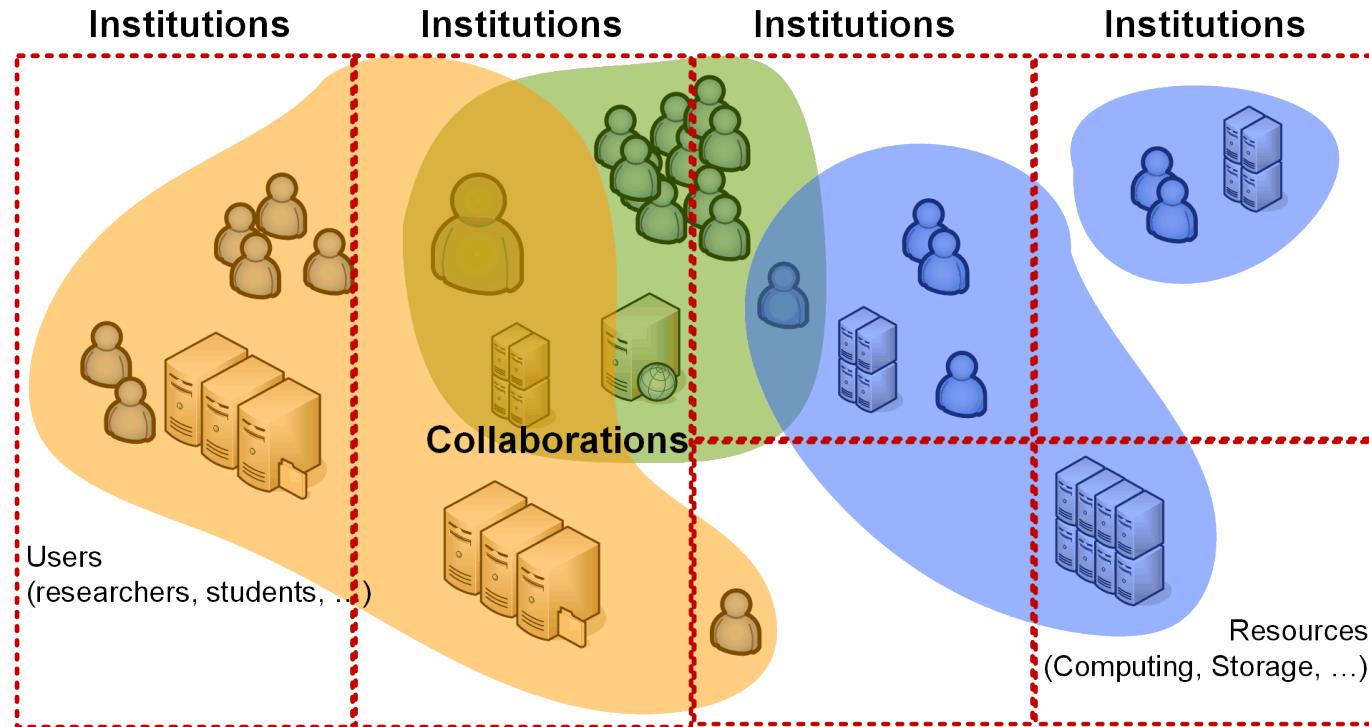
In the **Community** picture,
the source of authority
is *the community itself*



the AuthN-AuthZ separation is fundamental
to the Federated (R&E) AAI, global IGTF PKI, VOMS, 'AARC BPA' AAI architecture ...

Right-hand image: Shibboleth IdP federation, Lukas Hammerle, SWITCH (CH)

Since collaborations and institutions slice in different ways



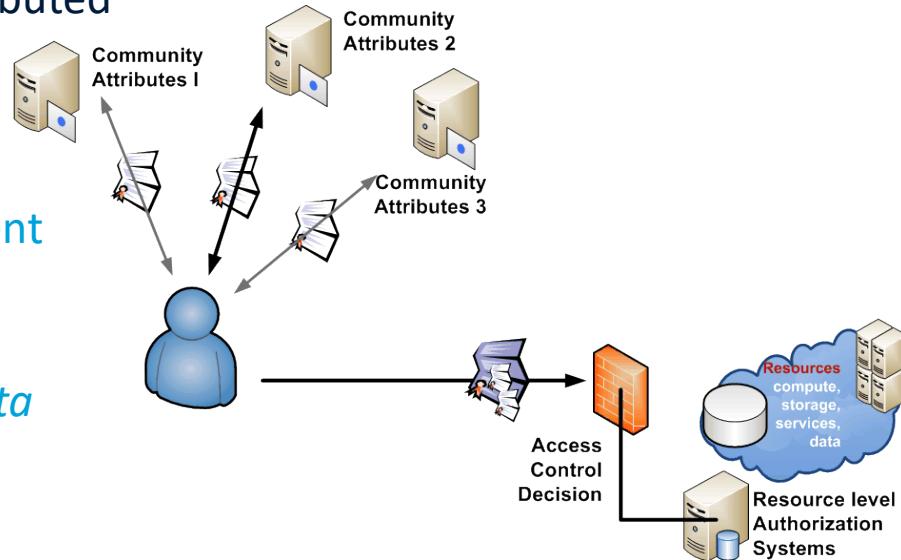
Multiple sources of authority: the community

- authorization assertion providers (attribute authorities) use the identifier(s) from authentication in their membership services
- *source of authority* for attributes is distributed

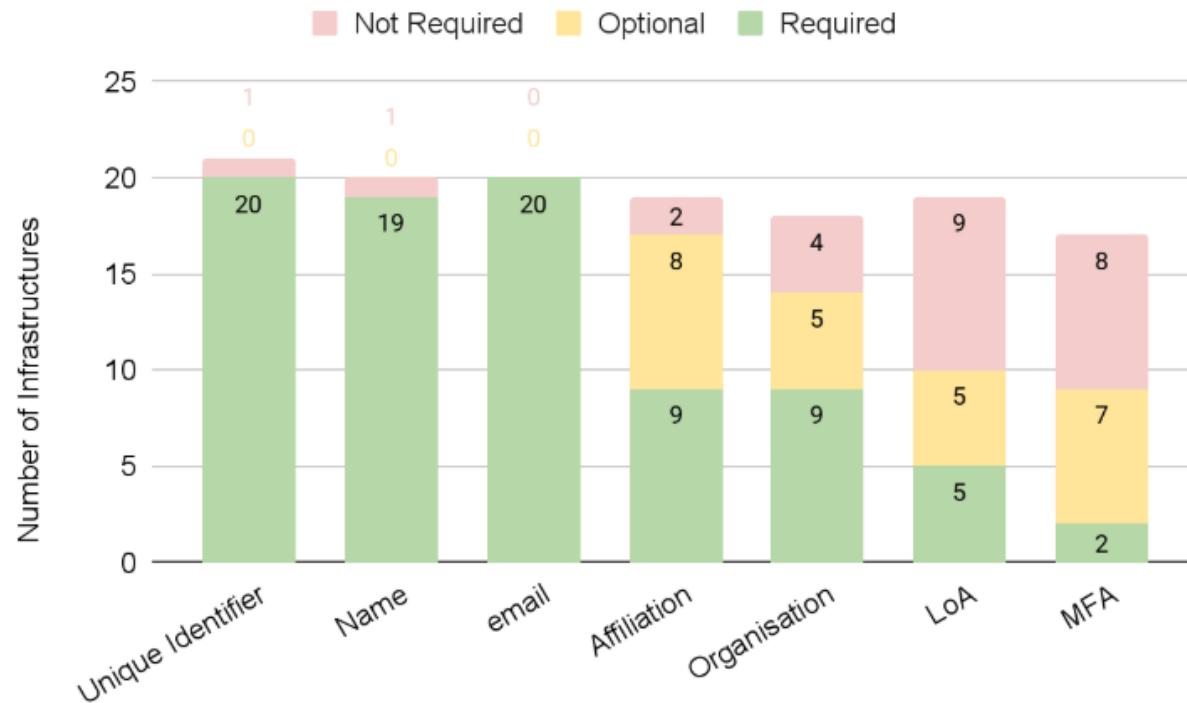
for example:

- community membership from an experiment
- affiliation status from home organisation

*may be jointly needed to access sensitive data
that is subject to medical-ethical clearance*



Research Infrastructures: what they *actually* need from ‘home’



Glossary

Affiliation: what *type* of entity are you (student, faculty, alumnus, ...)

LoA: level of authentication assurance (like passport identity vetting and ‘freshness’ of data)

MFA: multi-factor authentication (password, 6-digit code, SMS, fingerprint)

Source: Marina Adomeit, Janos Mohasci, *et al.* AARC TREE Use-case collection and analysis (D3.2), 2025 (under review)

The one infra that did ‘not need a unique identifier’ actually stated: “<our infra> assingns own identifier upon registration” – so the unique identifier is *still* there!

For starters: sharing good user identifiers is non-trivial 😞



REFEDS Spaces ▾

Entity-Categories

Pages / Entity-Categories Home

Research and Scholarship

Created by Nicole Harris, last modified on Apr 30, 2020

...

For IdP Operators

What attributes should be released by an R&S IdP?

The Research & Scholarship specification defines a bundles of attributes that Identity Providers are encouraged to release to R&S services:

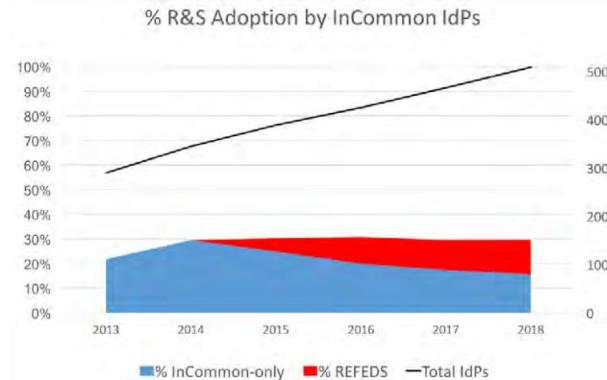
- personal identifiers: email address, person name, eduPersonPrincipalName
- pseudonymous identifier: eduPersonTargetedID
- affiliation: eduPersonScopedAffiliation

Category support is defined as follows:

An Identity Provider indicates support for the R&S Category by exhibiting the R&S entity attribute in its metadata. Such an Identity Provider MUST, for a significant subset of its user population, release all required attributes in

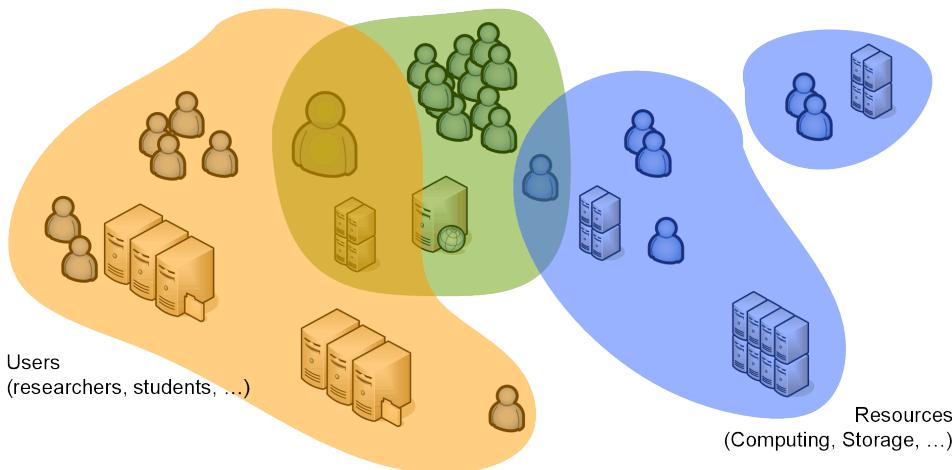
of 6019 identity providers
in 77 federations,
only 1994 support R&S or Personalised access **33%**

~ constant since 2018 😞

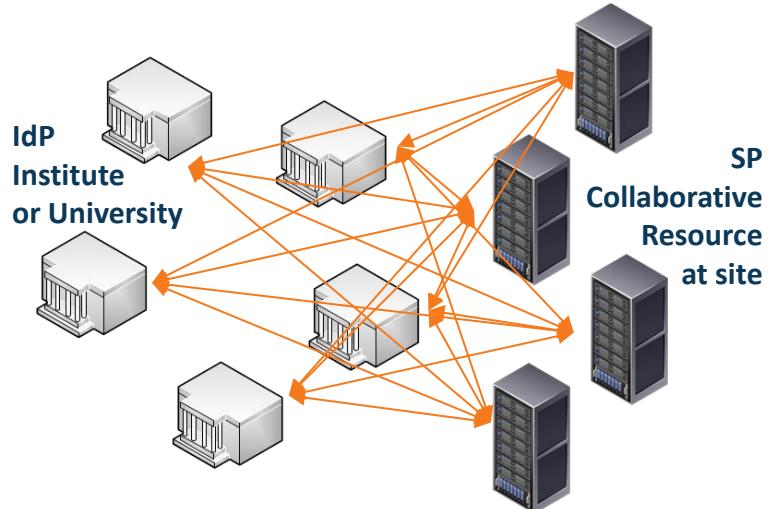


Graph: InCommon: Attributes-WG-Recommendations-May2018.pdf; Entity Category stats as per 2025-03-03, from <https://technical.edugain.org/entities>

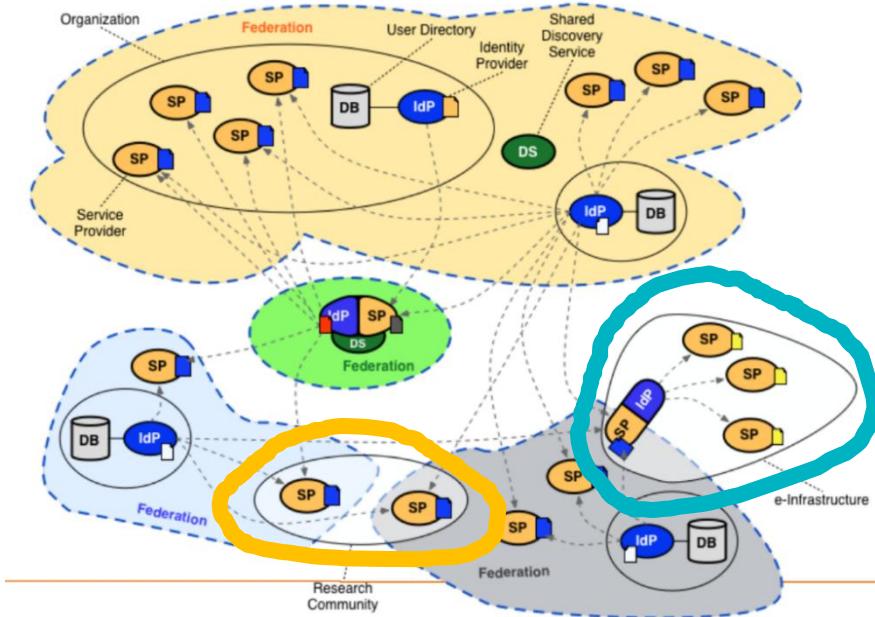
A fundamental scaling issue remained unique to research



for identity and user data
 $'n \times m'$ agreements remain(ed)

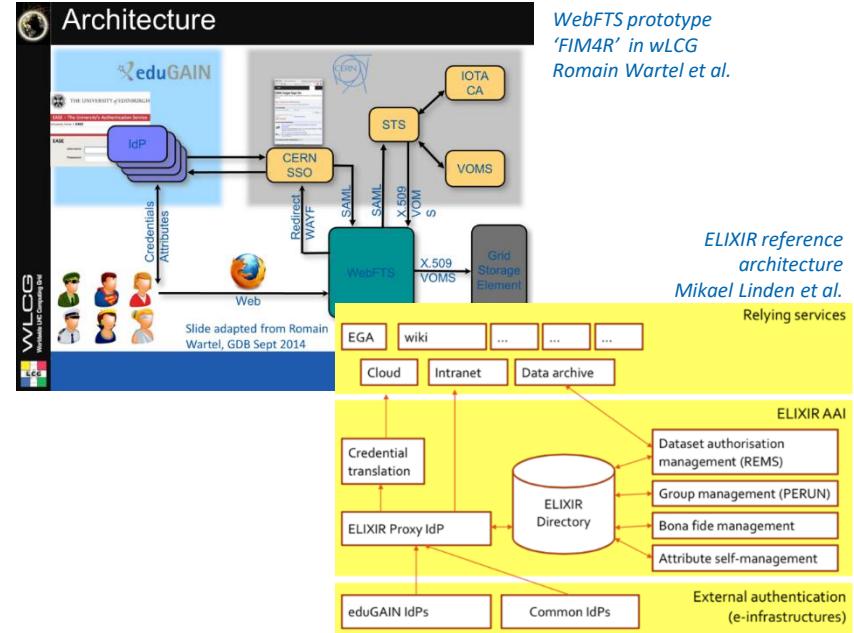


Managing complexity: distributed diverse identity sources



they were composed of many services
each of which had to manage federation complexity

Community images: Romain Wartel, CERN; Mikael Linden, CSC; Federation image (R): Lukas Hammerle, SWITCH

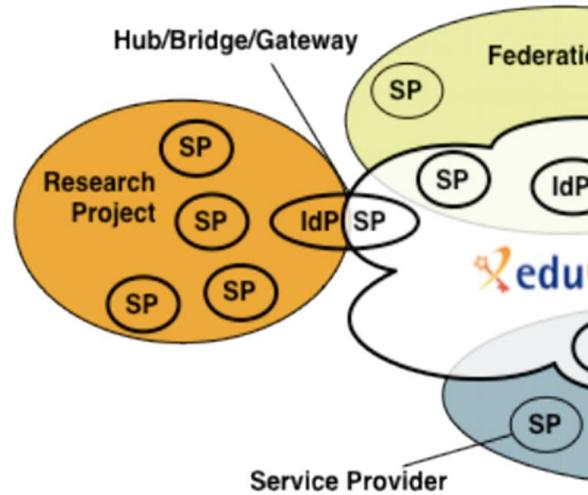


but most communities had started to invent
their own 'proxy' model to abstract complexity

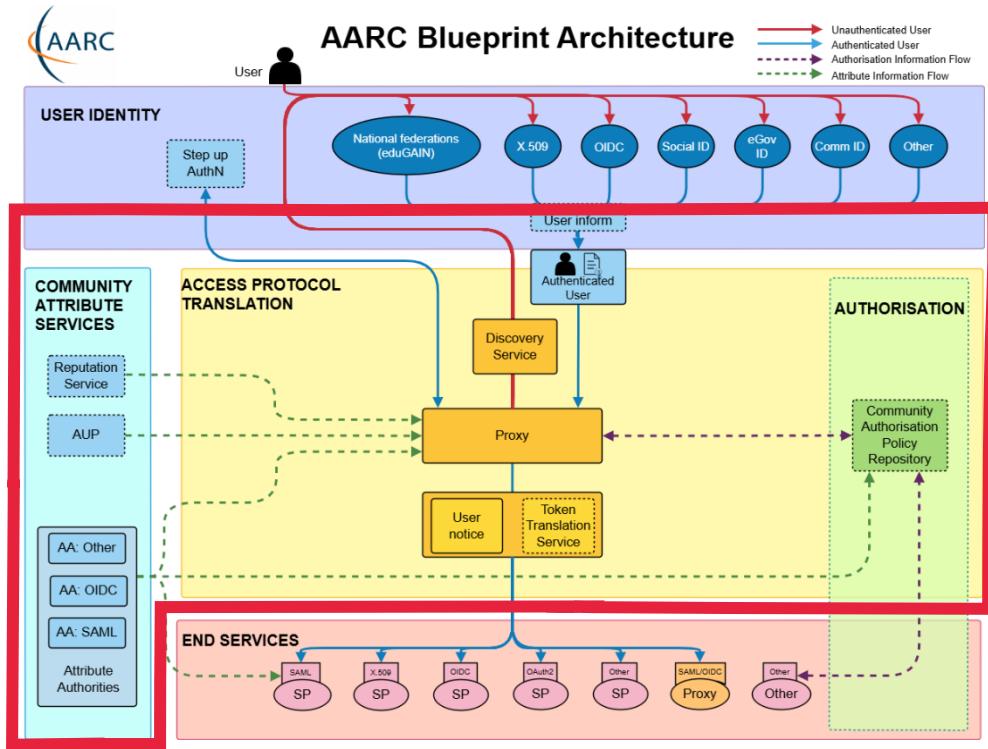
The IdP-SP bridge

- Access services using **identities from users' Home Organizations**,
- but **hide complexity** of multiple IdPs, federations, and different technologies for authentication and authorisation
- **One 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**

often known as proxy!



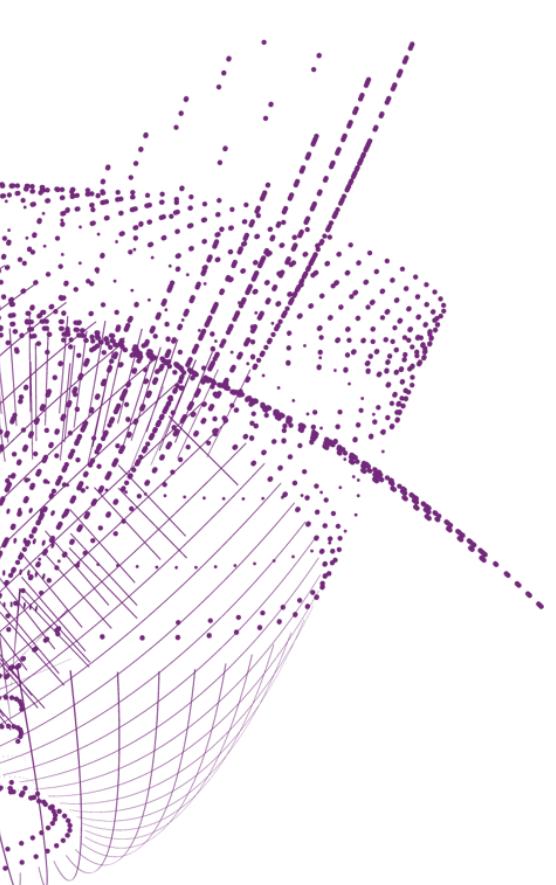
AARC Blueprint – making the bridge a first-class citizen



Manage users and access rights
with interoperable **building blocks**
for '**AAI infrastructure**' architects

that are

- technology-agnostic
- have multiple implementations
- come with policy templates & good practice guides

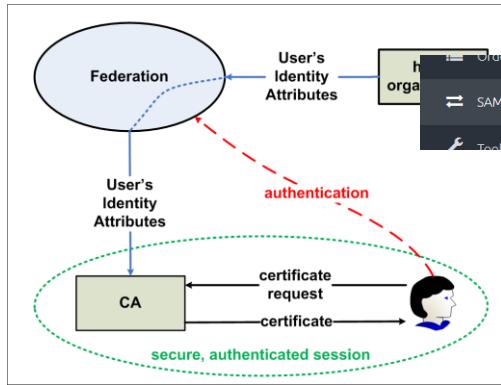


TCS and RCauth

Token translation

Bridges and Token Translation Services

TCS - for users that manage to grasp the idea



Organization Mapping

SURFcontext - Profile Overview

Organization	Attributes
Nikhef	nikhef.nl
EOM instituut AMOLF	amolf.nl

SURFcontext Apps
You have given permission to share profile information.

Service/App	Value
CERTcentral Digicert	urn:mace:terena.org:tcs

The following attributes are released to this Service Provider

Attribute	Value
Surname	Groep
E-mailaddress	davidg@nikhef.nl
First name	David
Entitlement	urn:mace:terena.org:tcs
Institution user ID	davidg@nikhef.nl
Organization	nikhef.nl
Display Name	David Groep

User Identification Request

This site has requested that you identify yourself with a certificate:
www.sugridma.org:443
Organization: Nikhef
Issued Under: TERENA

Choose a certificate to present as identification:
David Groep,davidg@nikhef.nl's TERENA ID [01:5CA9:2A4BFAF6:6256:73:35:81:E9:2A:09:AE]

Details of selected certificate:
Issuing CA: GÉANT Personal Certificate
Subject: 01:5CA9:2A4BFAF6:6256:73:35:81:E9:2A:09:AE
Valid from Tuesday, 4 September, 2019 02:00:00 to Thursday, 3 October, 2019 14:00:00
Key Usage: Signing, Key Encipherment, Data Encipherment
Email address(es): davidg@nikhef.nl
Issued by: TERENA eScience Personal CA 3.0=TERENA,=Amsterdam,ST=Noord-Holland,CN=...
Check for updates, Continue

OK Cancel

Trusted Certificate Service

SECTIGO

Digital Certificate Enrollment

You have been authorized to enroll for a digital certificate. Please validate that your name and email addresses are correct.

Name: David Groep
Email: davidg@nikhef.nl
Organization: Nikhef

Please select the correct certificate profile and desired private key format. If a private key is generated a password is required to protect the download.

Certificate Profile:
 GÉANT Personal Certificate
 GÉANT IGTF-MICS Personal
 GÉANT IGTF-MICS-Robot Personal

Private Key:
 Generate RSA
 Generate ECC
 Upload CSR No file chosen

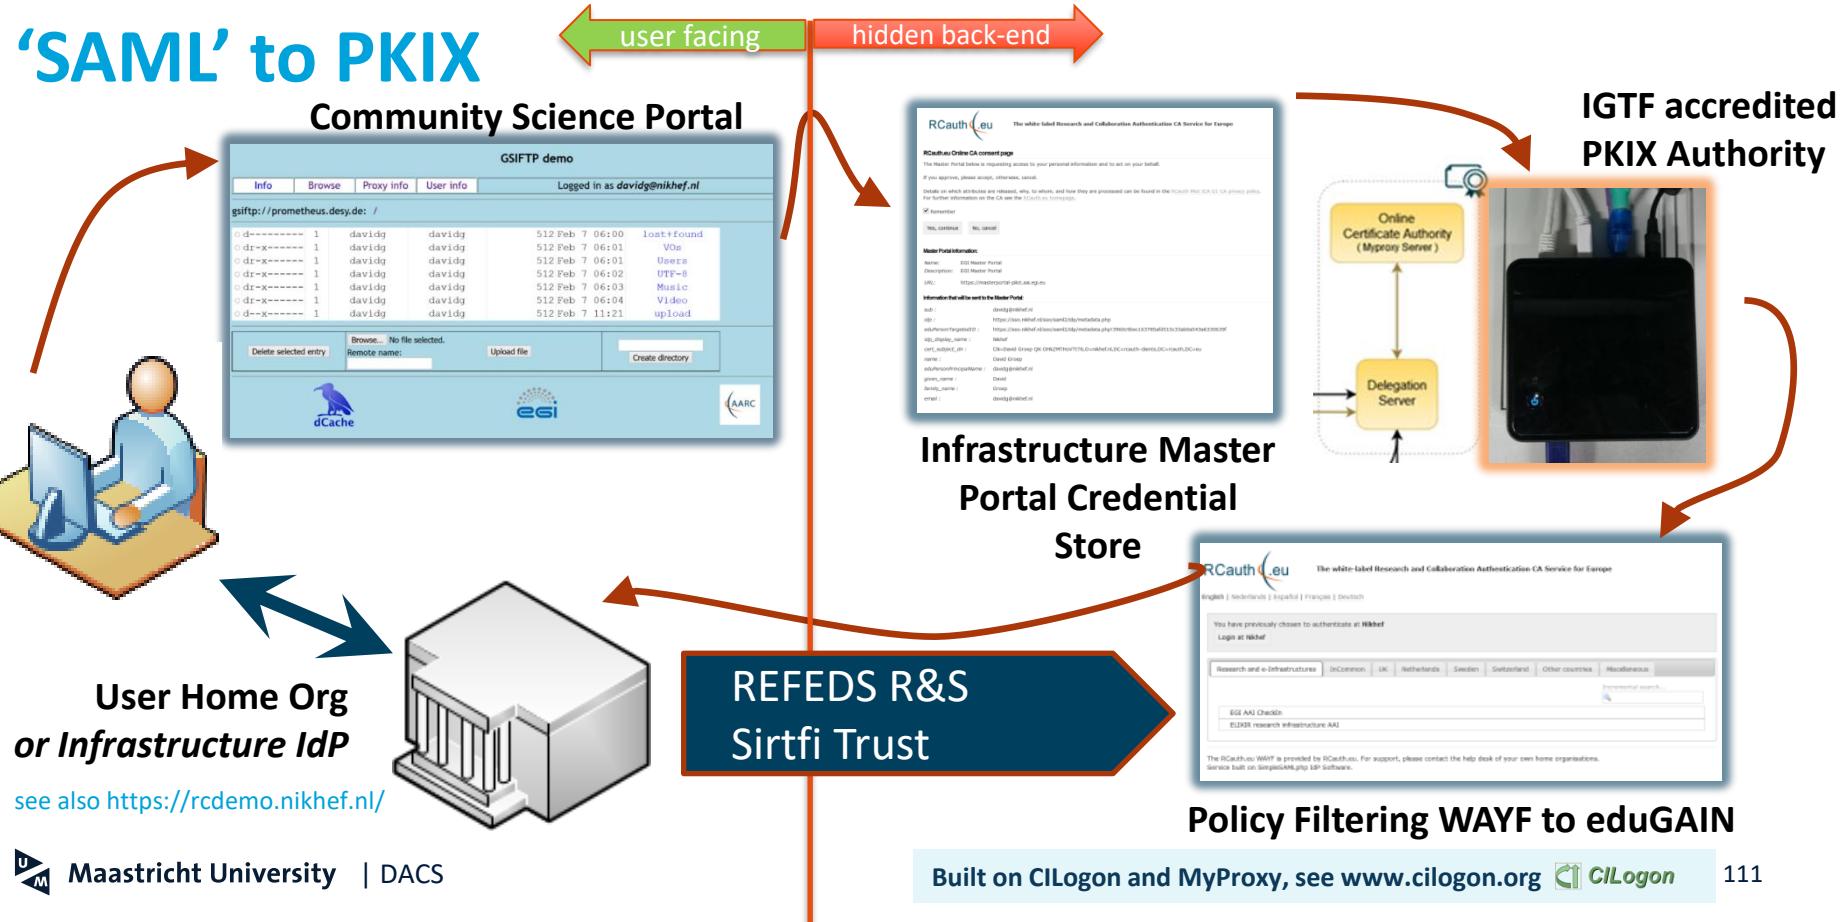
P12 Password: P12 Password Confirmation:

SUBMIT

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

Seamless in-line token translation services from 'SAML' to PKIX



Unique certificated from FIM via eduPerson and REFEDS R&S

Sources of naming and uniqueness, that work *today*

- **eduPersonPrincipalName** – scoped point-in-time unique identifier, which could be, but usually is not, privacy preserving: “davidg@nikhef.nl”, “P70081609@maastrichtuniversity.nl”
- **eduPersonTargetedID** – scoped transient non-reassigned identifier, like <urn:geant:nikhef.nl:nikidm:idp:sso!27c8d63ed42c84af2875e2984>
- **subject-id** - a scoped persistent non-reassigned identifier, which should be privacy-preserving: <44f7751265a6e8b228f9@nikhef.nl>

Plus the (domain-name based) schacHomeOrganisation and a ‘**representation of the real name**’

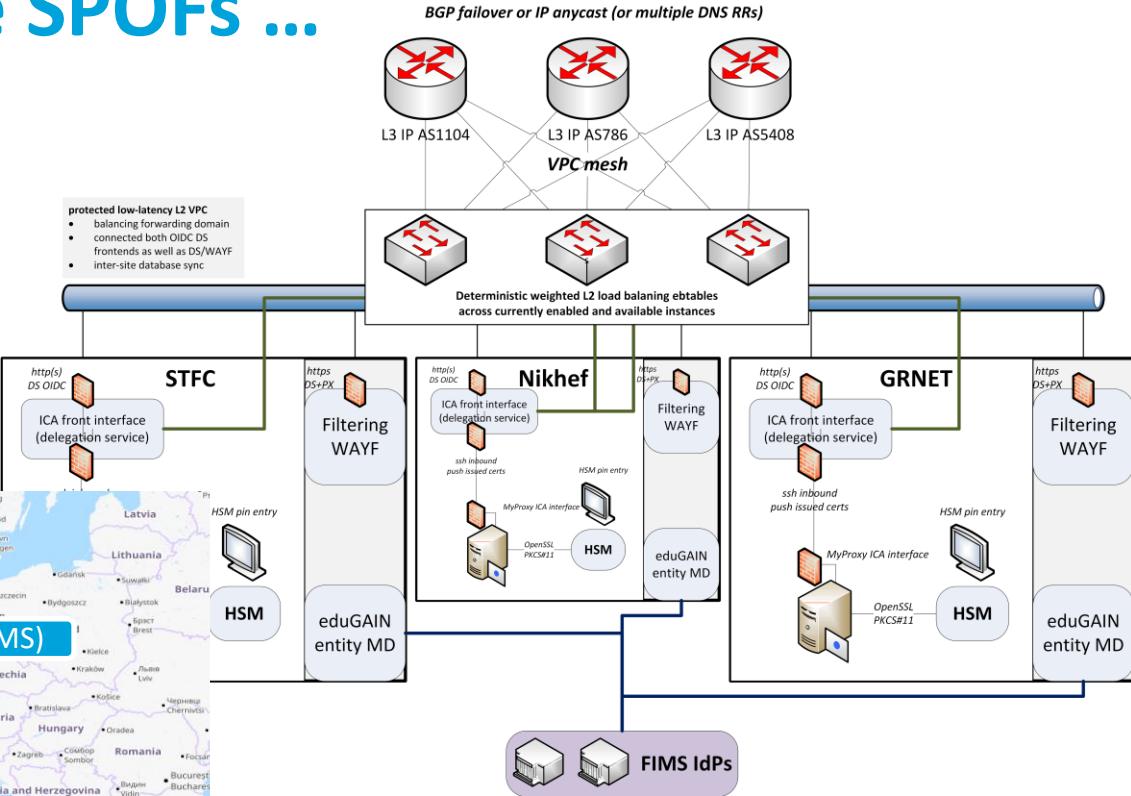
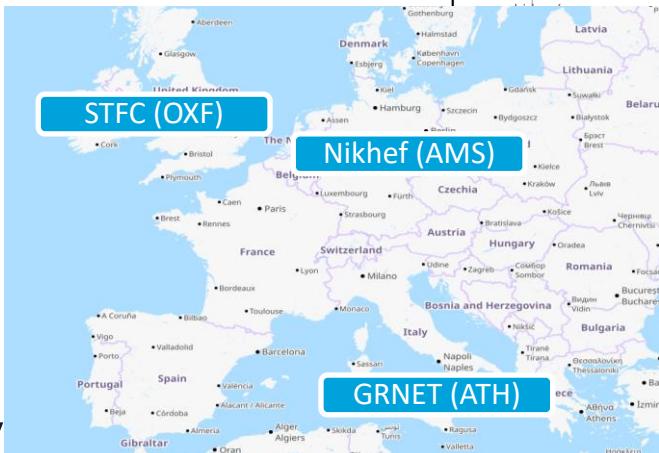
/DC=eu/DC=rcauth/DC=rcauth-clients/O=orgdisplayname/CN=commonName +uniqueness

uniqueness will be added to commonName via hashing of *ePPN*, *ePTID*, *subject-id*, so that an enquiry via the issuer allows unique identification of the vetted entity”

Since we do not like SPOFs ...

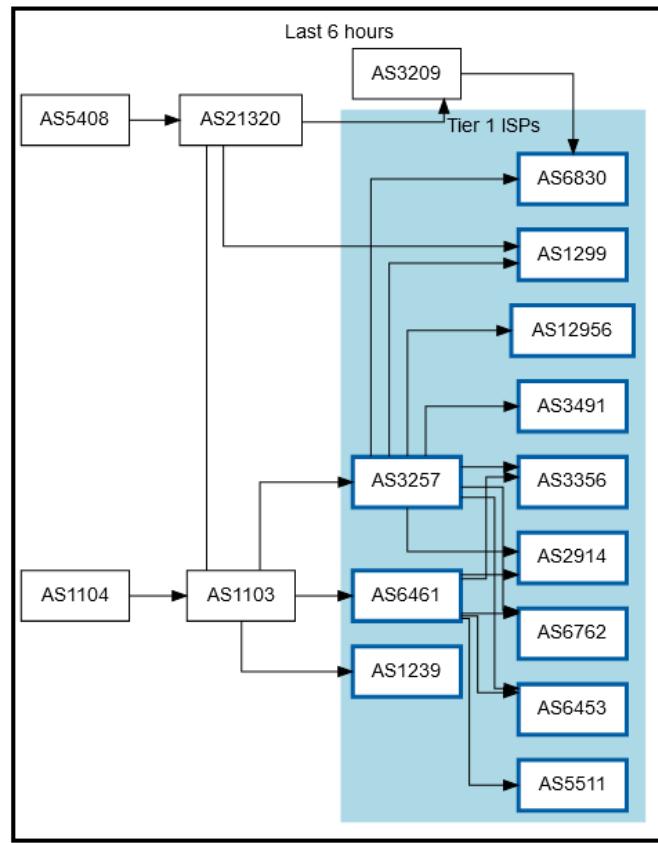
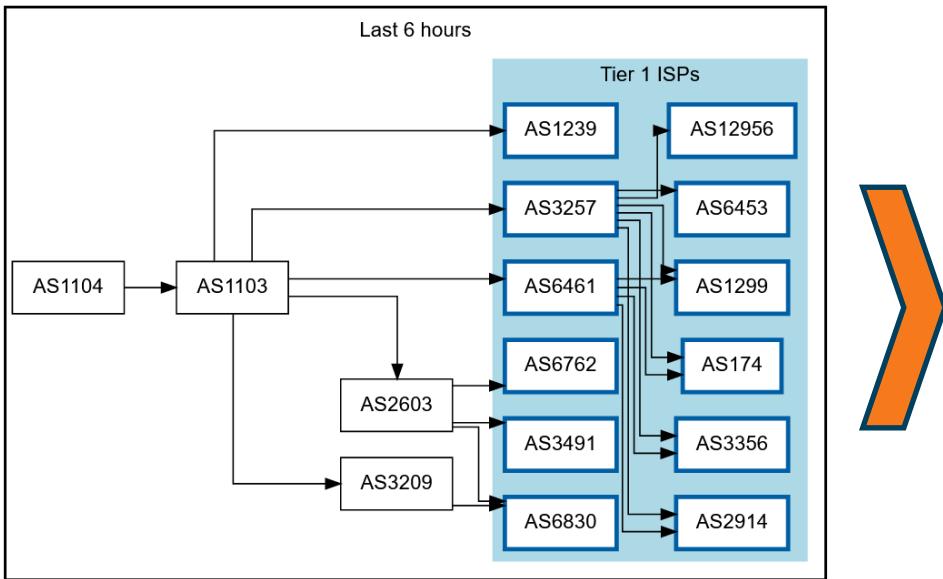
Distributed High Availability setup

- 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

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



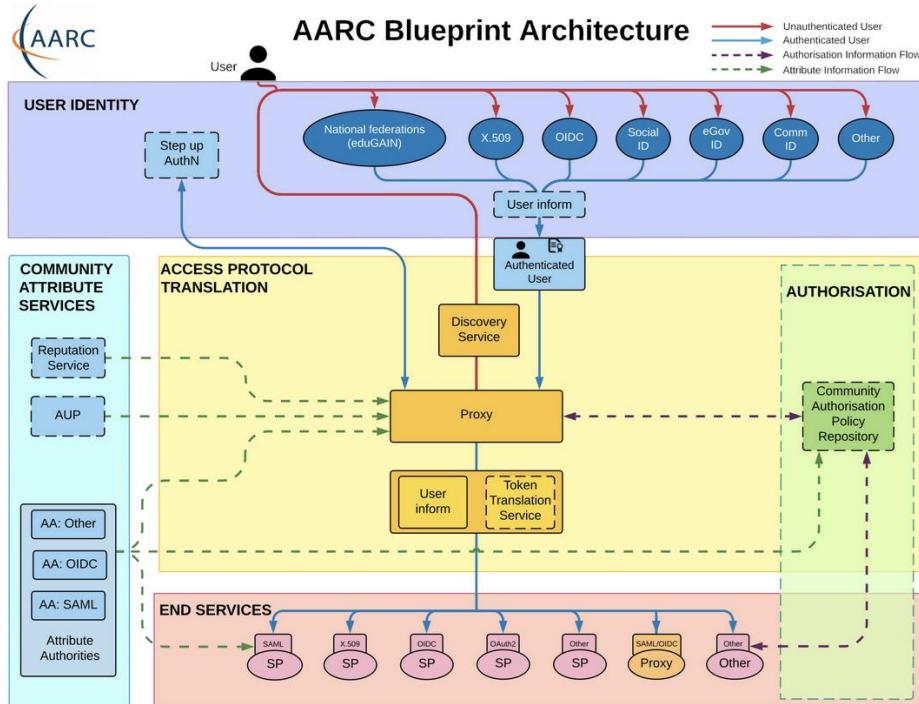
route maps: bgp.tools for 2a07:8504:1a0::/48 – IPv4 for 145.116.216.0/24 is similar – imagery from November 2022

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/>)

More than just nice colours



<https://aarc-community.org/guidelines/>

User Identity:

- How should I integrate Social Media Identity Providers? [AARC-G008](#)
- How should users link accounts, and how does that affect Assurance? [AARC-G009](#)
- How should services indicate that they would like users to authenticate with multifactor authentication, and how should my proxy forward that information? [AARC-G029](#)

Assurance:

- How should assurance information of external identities be calculated? [AARC-G031](#)
- What can I say about assurance of identities from social media accounts? [AARC-G041](#)
- How is assurance impacted by account linking? [AARC-G009](#)
- How should assurance information be shared with other infrastructures? [AARC-G021](#)
- Which Assurance Profiles should I use, there are so many! [AARC-I050](#)

Community Attribute Services:

- How should attributes from multiple sources be aggregated? [AARC-G003](#)
- How should I express the home institute of a user? [AARC-G025](#)
- How should I express the identifier of a user? [AARC-G026](#)
- What are the best practices for running my Attribute Authorities securely? [AARC-G071](#)
- Which Acceptable Use Policy should I use to facilitate interoperability? [AARC-I044](#)
- How should I infer the affiliation of a user? [AARC-Q057](#)

Access Protocol Translation:

- Which best practices should I follow for my Token Translation Services? [AARC-G004](#)
- How should I translate from Identity Federation information to X.509 certificates? [AARC-G010](#)

Authorisation:

- How should I manage authorisation information from multiple sources? [AARC-G006](#)
- How should group and role information be expressed to facilitate interoperability? [AARC-G002](#)
- How should resource capabilities be expressed? [AARC-G027](#)

End Services:

- My service needs to act on behalf of the user – how should I handle credential delegation and impersonation? [AARC-G005](#)
- My services are not web based, how can I use identities from the proxy? [AARC-G007](#)
- How should Services hint which idP they would like users to use? [AARC-G049](#)
- Which Security practices should I follow? [AARC-G014](#)

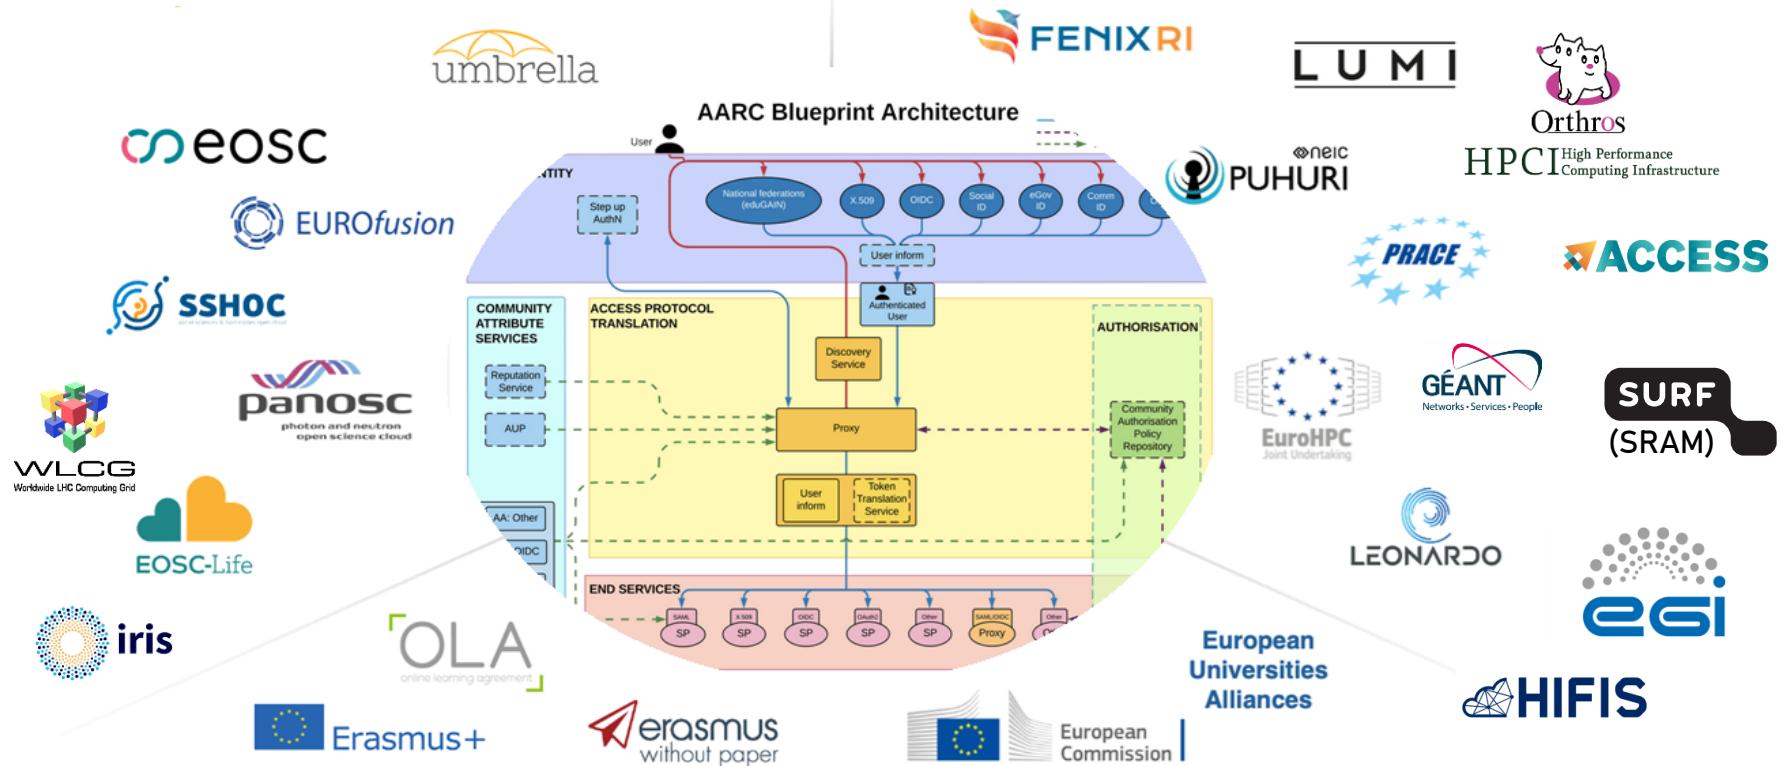
What next? Are you looking for a kick start with your policies? Take a look at the [Policy Development Toolkit](#) which provides a set of templates.

Service Operations Security Policy	Infrastructure Management	Services (abide by)	policy + service (for service specific policies)	Infrastructure and its personnel, both service in the infrastructure, as well as the infrastructure itself, should complete the template.
Acceptable Use Policy	Infrastructure Management (for baseline) & Research Communities (for community specific restrictions)	Users (abide by)	This policy defines requirements for running a service within the Infrastructure.	This is a template for the acceptable use policy that users must accept. It is augmented by the Research Communities.
				Google Doc

PDK

Showing 1 to 9 of 9 entries

There is plenty of AARC deployments ...



Example: SURF Research Cloud Secure Supercomputing

Welcome to your SURF research cloud dashboard

Quick actions

- Create new workspace
- Create new storage

Start here to create everything you need for a workspace.

Start here to create everything you need for storage.

Workspaces Storage IP addresses (advanced) Networks (advanced)

Workspaces

- Encrypted Mounting State: running
- SANE linux data owner portal State: running
- Centos VM port 1194 State: running

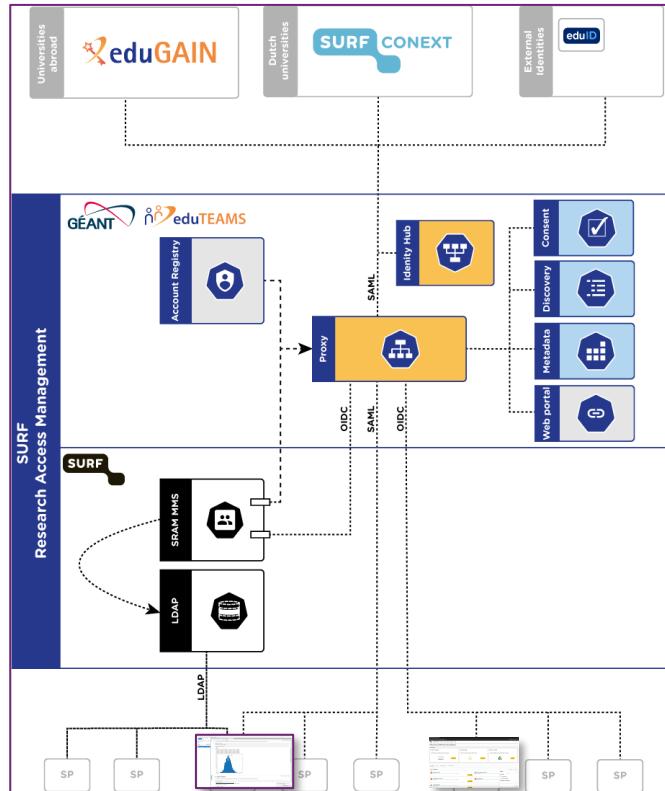
SANE Tinker Webinar Project

Filters

- Show only
- My workspaces
- Deleted workspaces

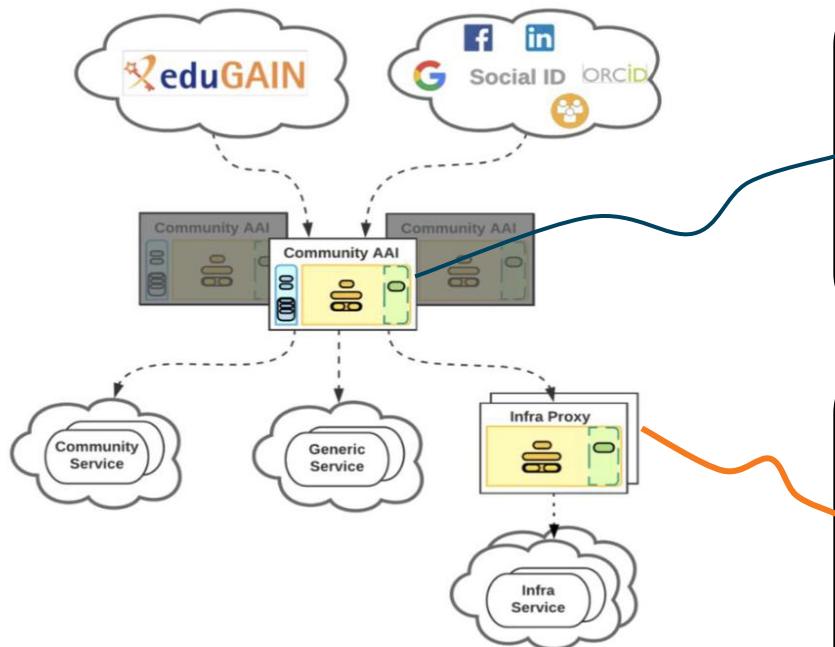
SANE data server

Collaborative organisations



SURF SRAM architecture, Raoul Teeuwen et al. from
<https://servicedesk.surf.nl/wiki/display/IAM/Dienstbeschrijving+SURF+Research+Access+Management>
SURF Research Cloud capture: from Introduction to SANE (Secure ANalysis Environment)
webinar February 2024, by Martin Brandt et al., SURF
<https://www.surf.nl/themas/onderzoeksinfrastructuur/sane-veilige-omgeving-voor-analyse-van-gevoelige-data>

... but one proxy is not enough in a research cloud



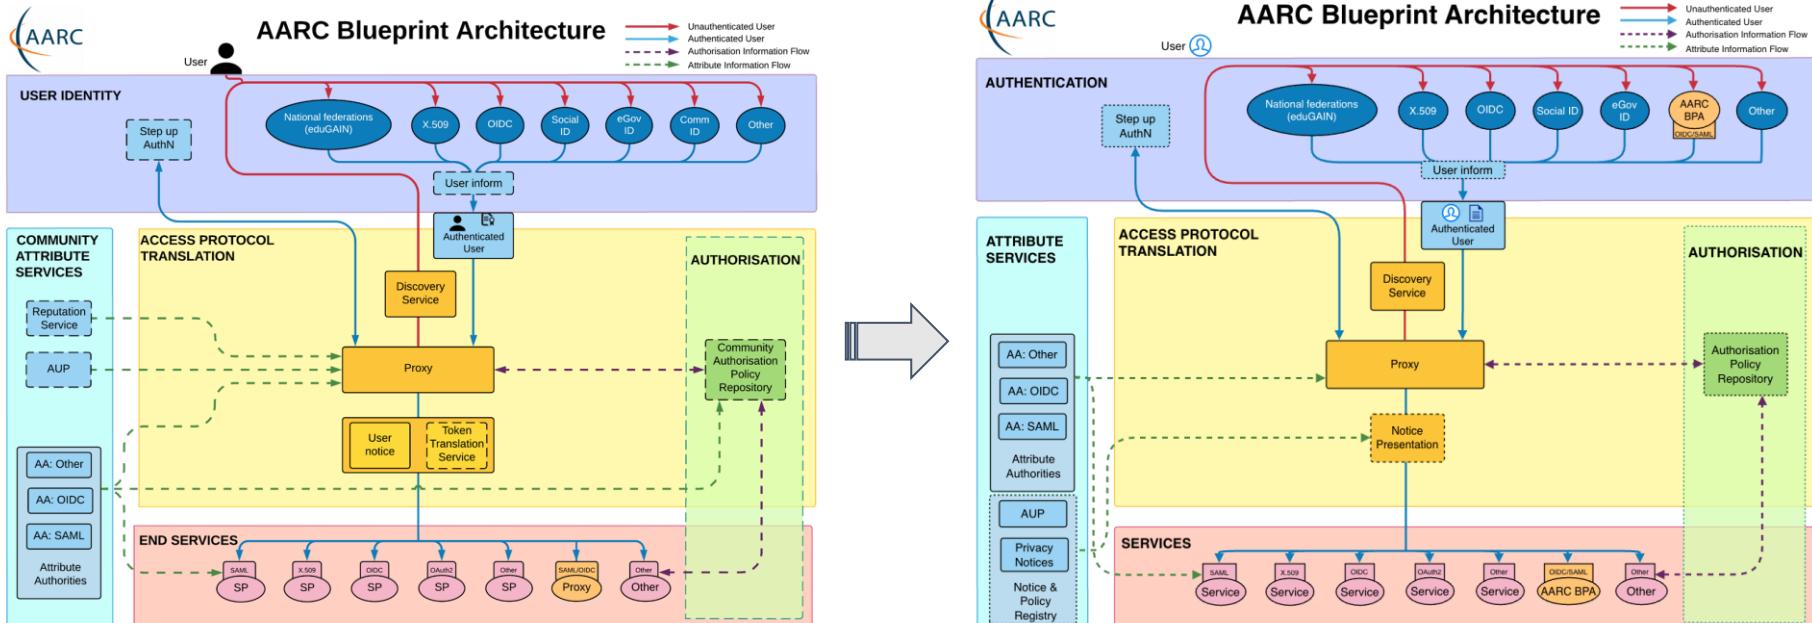
Community AAI

streamline researchers' access to services, both those provided by their own infrastructure as well as the services provided by shared infrastructures from other communities.

Infrastructure Proxy

enables Infrastructures with large number of resources, to provide them through a single integration point, where the Infrastructure can maintain centrally all the relevant Policies and business logic for making available resources to multiple communities

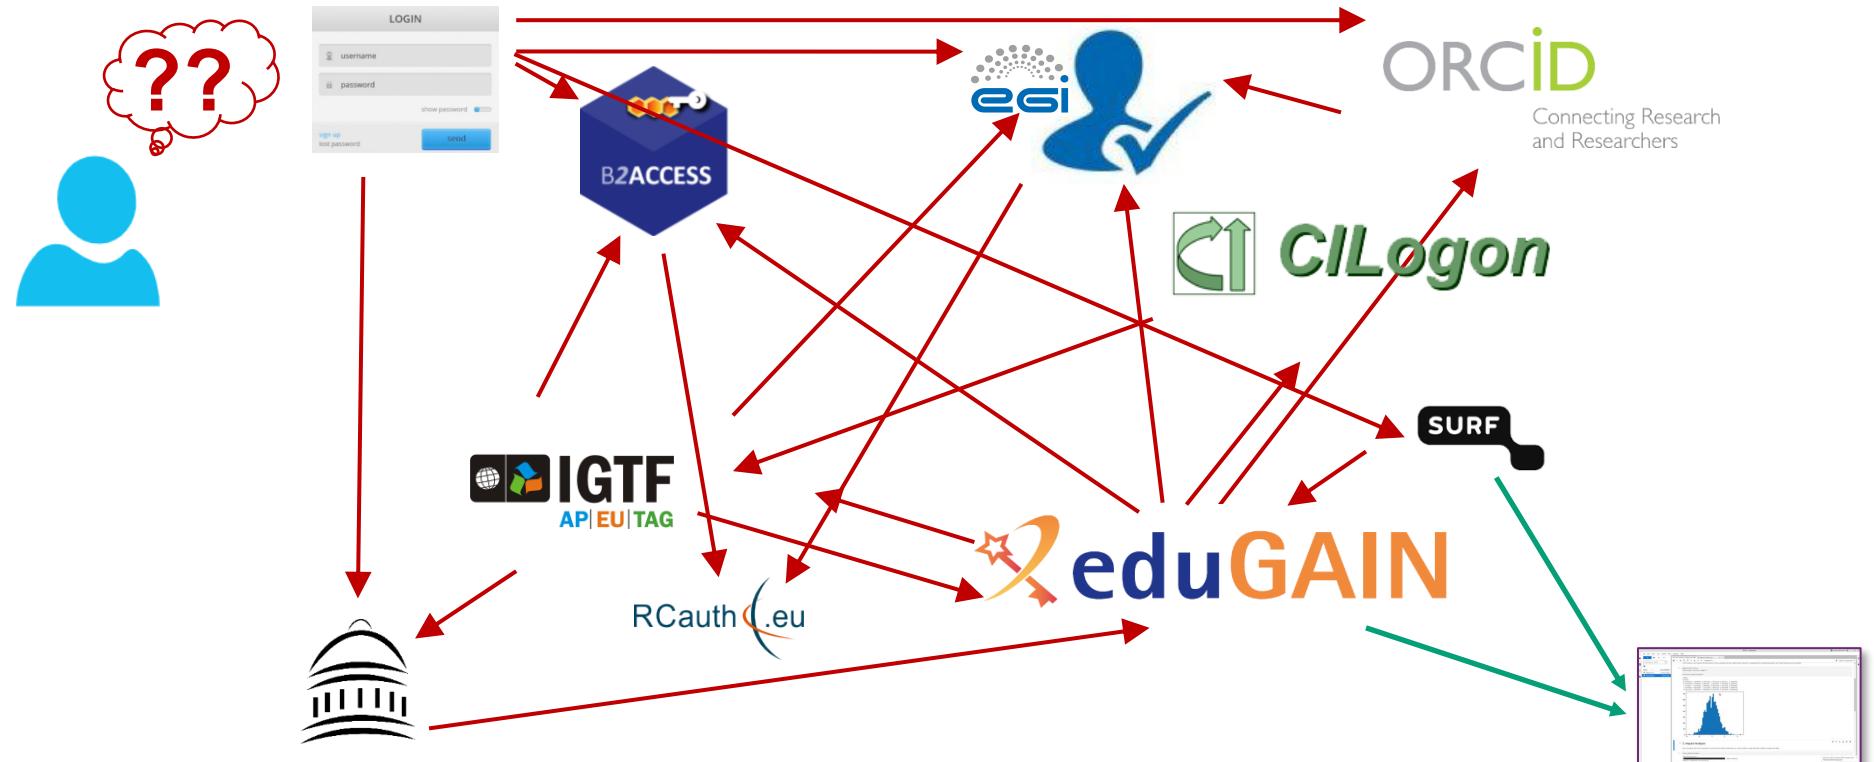
AARC Blueprint Architecture 2025: Component Layers - What has changed since AARC-BPA-2019? Can you spot the differences?



AARC-BPA-2019

AARC-BPA-2025

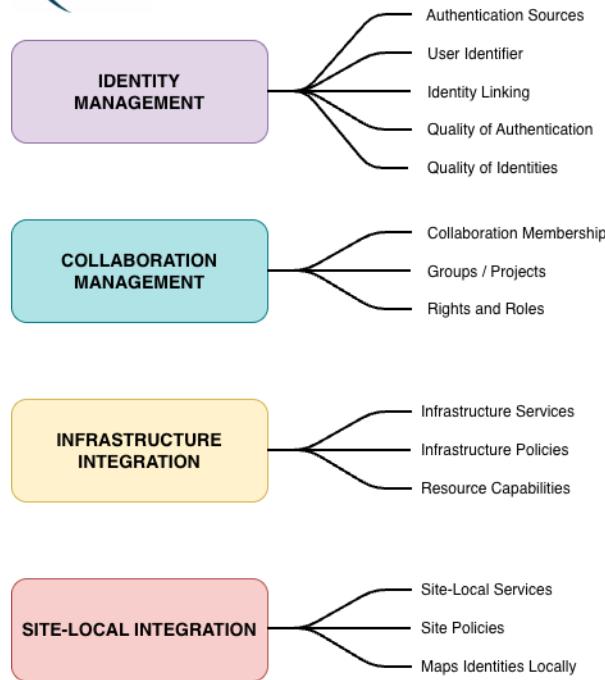
Identity spaghetti: 1-loop, 2-loop and higher order diagrams



AARC Blueprint Architecture 2025: Functional Capabilities



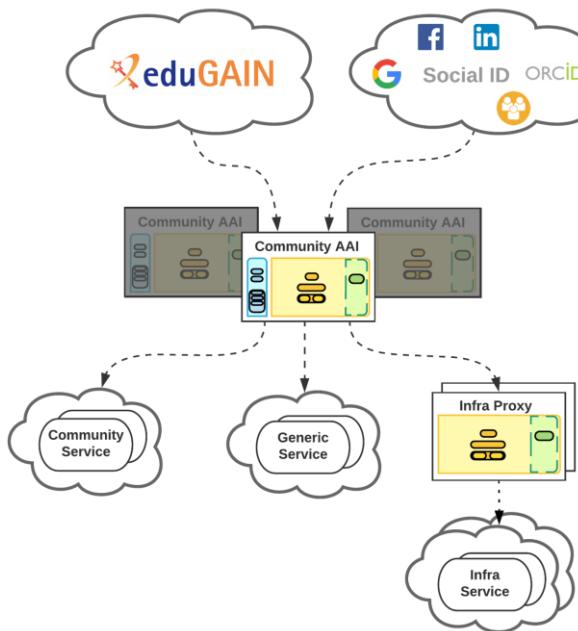
AARC Blueprint Architecture



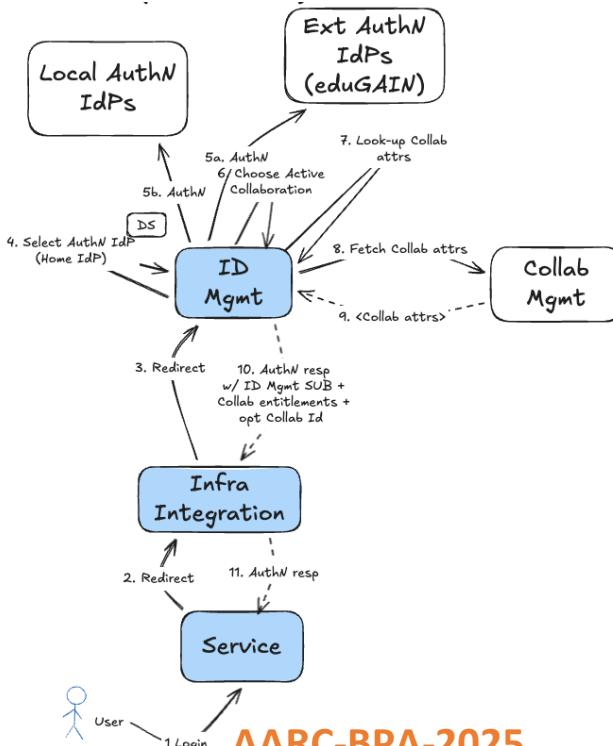
What has changed since AARC-BPA-2019?

- **Added Identity Management capability:**
 - Groups identity-related functions such as unique identifier assignment, identity assurance, authentication assurance, and identity linking
- **Community AAI → Collaboration Management**
- **Added Site-local Integration capability:**
 - Enables integration of federated users into local environments

A new way to access resources with AARC-BPA-2025

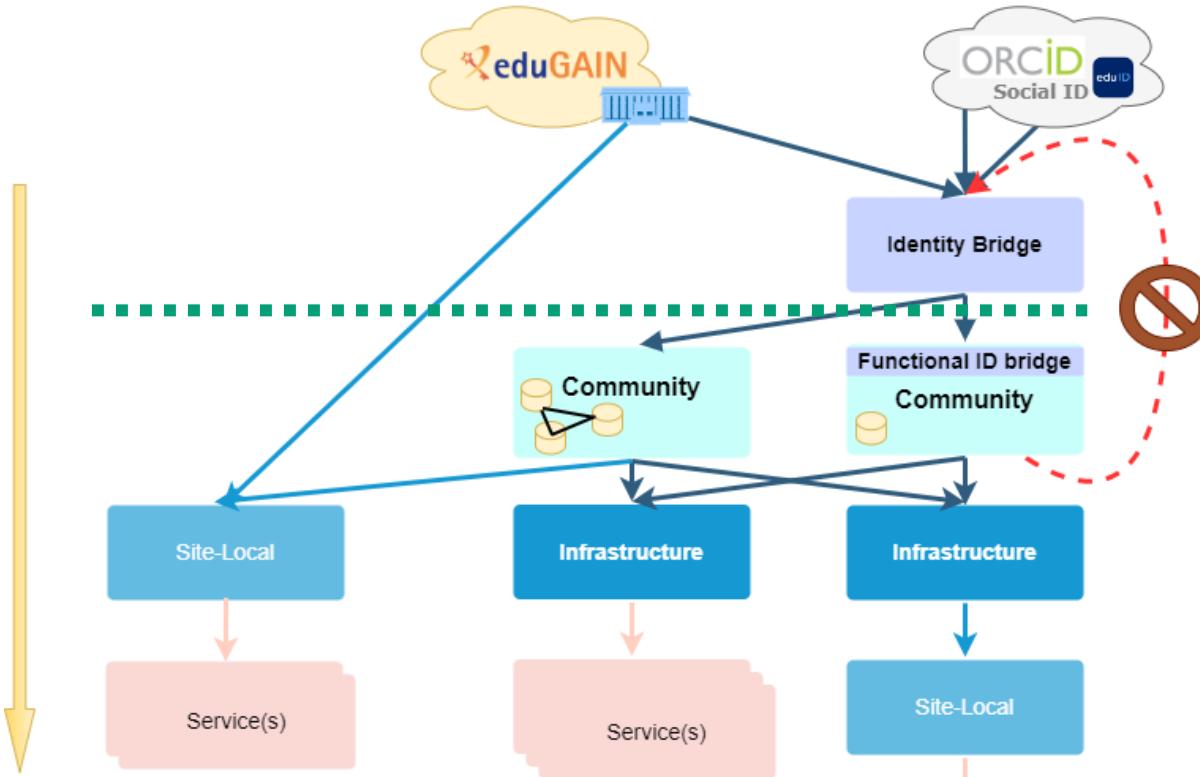


AARC-BPA-2019 aka “Community-first”



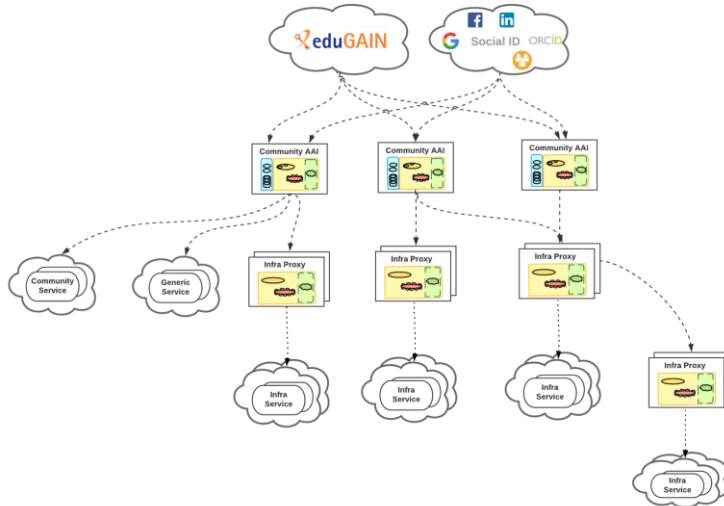
AARC-BPA-2025

Not all that is possible is allowed in the AARC BPA

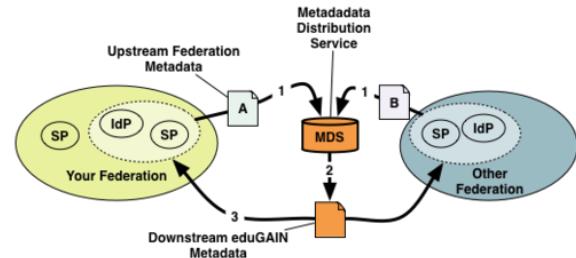


We have seen many arrows before ... it needs federation!

Identity, community, infrastructure proxies and services form a **federation of proxies**



- bilateral registration
but then you have a scalability issue again
- meta-data distribution of trust paths
 - **OpenID Federation**
 - **SAML** meta-data
- discovery and identity provider hinting



European Open Science Cloud federation (2023 edition)

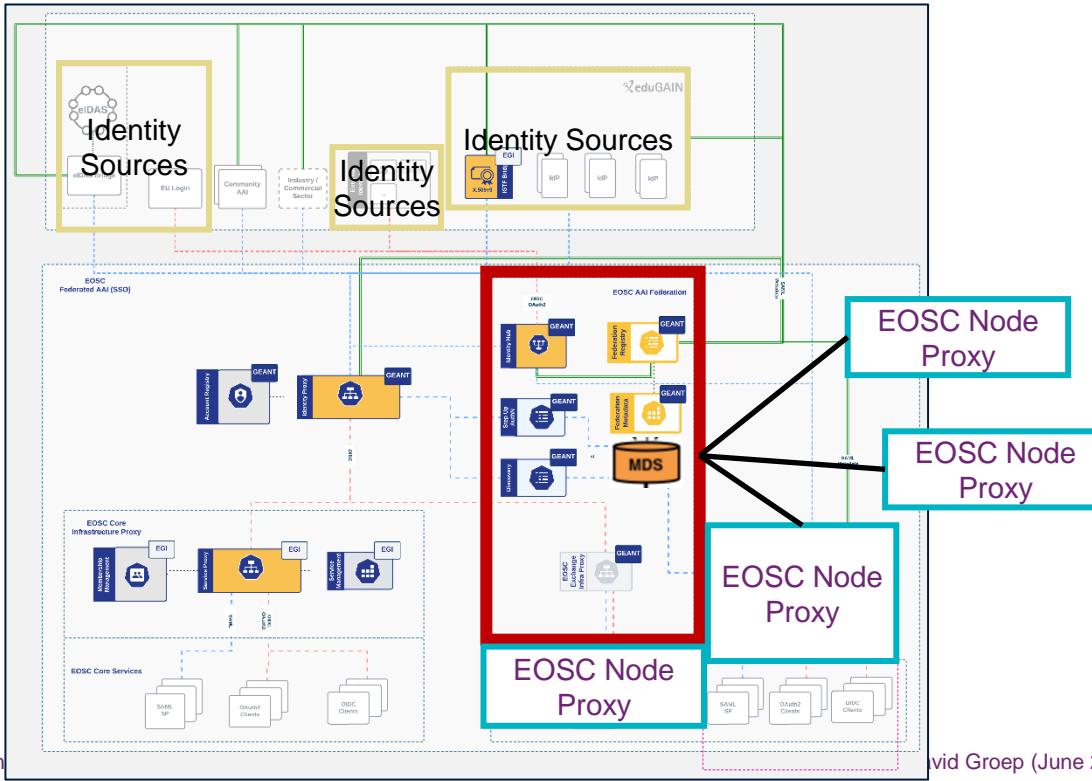


Image: EOSC AAI for the EOSC Core and Exchange

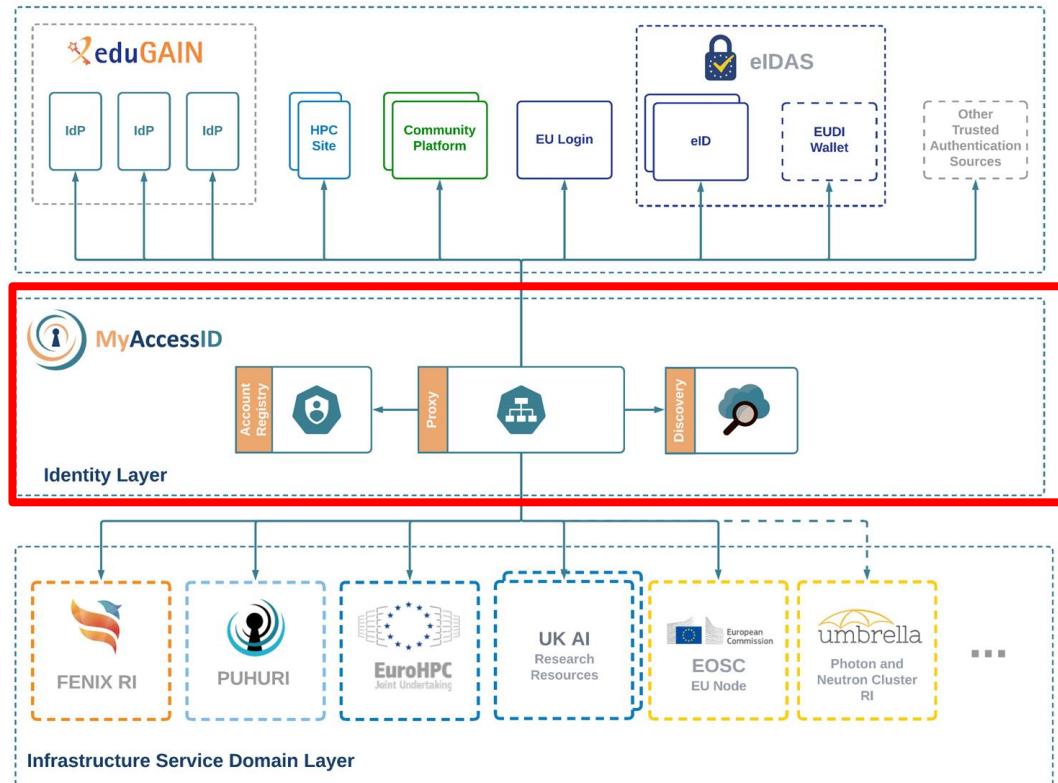
avid Groep (June 2023)

AARC BPA Deployment example: MyAccessID and EOSC AAI

MyAccessID: A common Identity Layer for Science



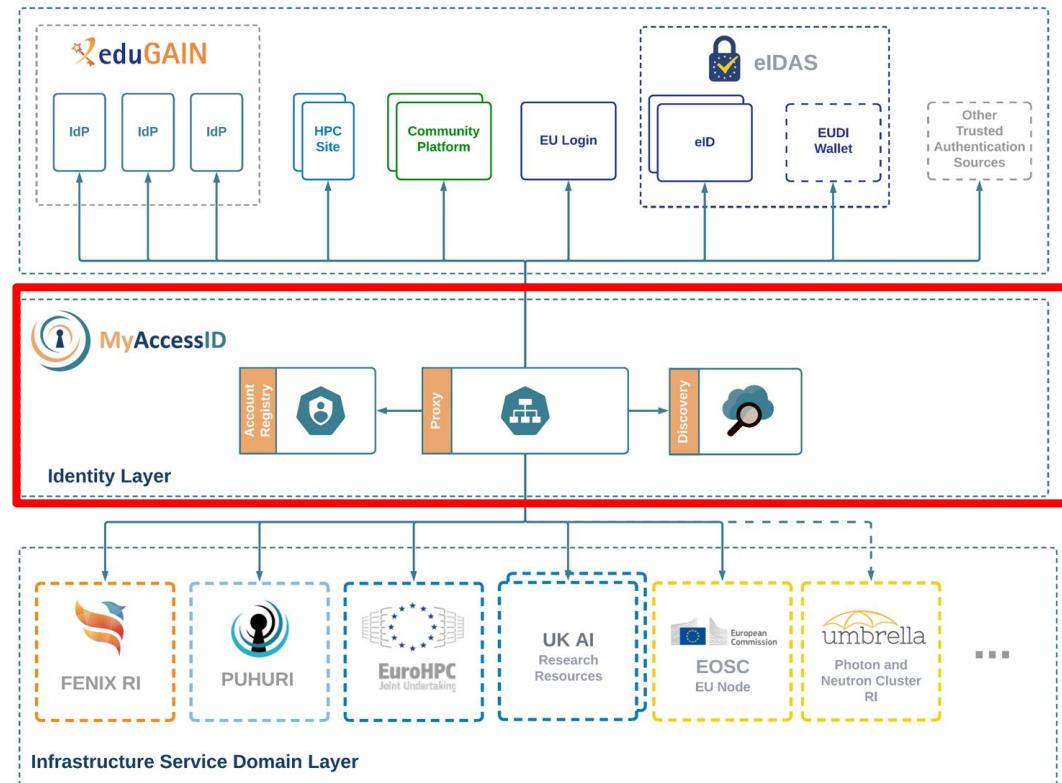
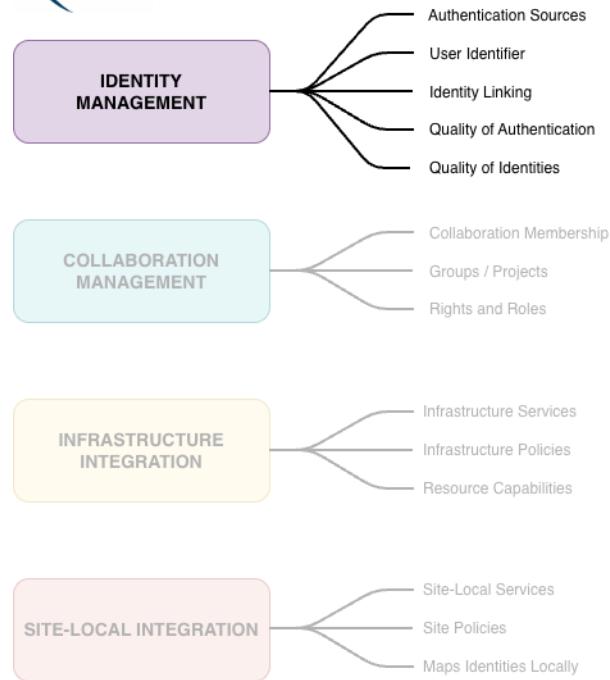
- HPC Datacenters are in the process of transforming to **Infrastructure Service Providers** with a **diverse Service Portfolio**
- These services become available in different administrative and policy domains, which we call **Infrastructure Service Domains**
- **A common Authentication and Authorization Infrastructure** enables uniform accessibility to scientists and engineers at European scale

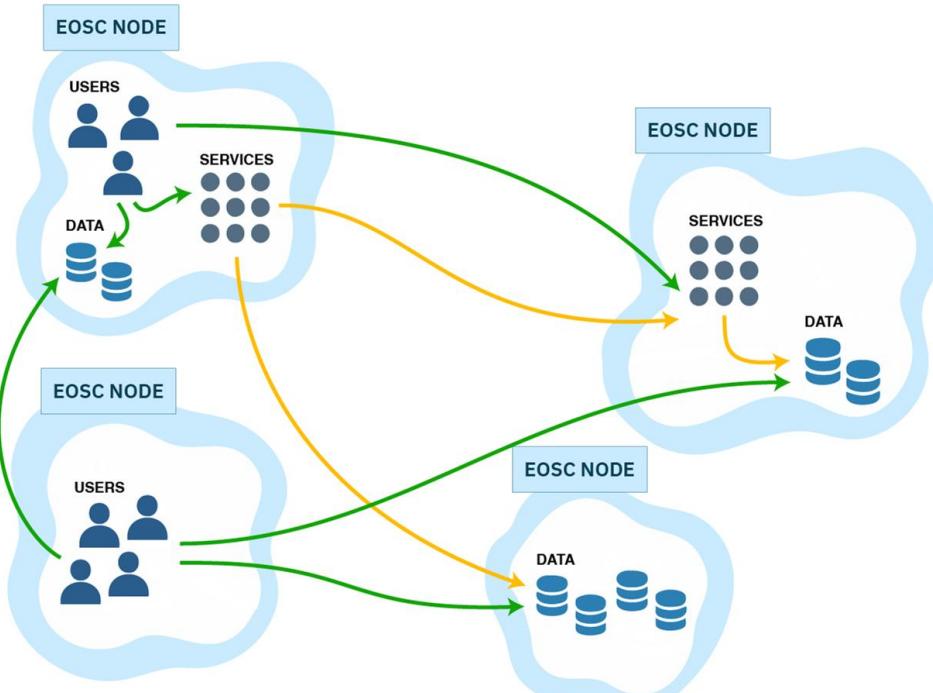


MyAccessID: A common Identity Layer for Science



AARC Blueprint Architecture



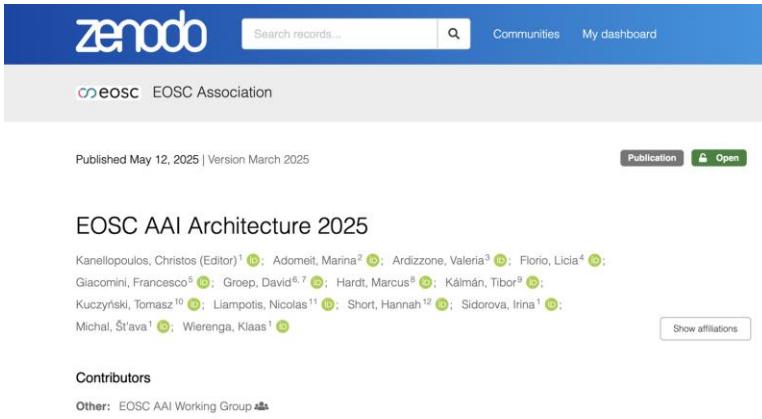


→ interactive login of users
→ service to service

Initial EOSC AAI Use Cases

- Single Sign On Across Nodes
- Cross Node Workflows

* See Licia's presentation from the FIM4R session



The screenshot shows the Zenodo document page for the EOSC AAI Architecture 2025 report. The page has a blue header with the Zenodo logo, a search bar, and links for 'Communities' and 'My dashboard'. Below the header, it says 'eosc EOSC Association'. The main content area has a blue header bar with the title 'EOSC AAI Architecture 2025'. Below this, the document details are listed: 'Published May 12, 2025 | Version March 2025', 'Publication', and a green 'Open' button. The document title is 'EOSC AAI Architecture 2025'. The contributors listed are: Kanellopoulos, Christos (Editor)¹; Adomeit, Marina²; Ardizzone, Valeria³; Florio, Licia⁴; Giacomini, Francesco⁵; Groep, David^{6,7}; Hardt, Marcus⁸; Kálmán, Tibor⁹; Kuczyński, Tomasz¹⁰; Liampotis, Nicolas¹¹; Short, Hannah¹²; Sidorova, Irina¹; Michal, Št'ava¹; Wierenga, Klaas¹. There is a 'Show affiliations' button. Below the document details, there is a 'Contributors' section and a note about the EOSC AAI Working Group.

This document presents recommendations for the initial implementation of the EOSC AAI Federation, offering background on prior work and summarising recent advancements, including updates to the AARC Blueprint Architecture.

AAI implementers who wish to go directly to the technical requirements may refer to the "Implementation" section, while those interested in the rationale behind the architectural choices are encouraged to also read the "Background Information" section.

The overarching goal of the EOSC AAI Federation is to eventually support a full-mesh, dynamic topology without introducing a centralised component into the European AAI ecosystem. However, current technological constraints — particularly those associated with OpenID federation — limit the feasibility of such a model.

The work required at the architecture level will certainly extend beyond 2025, while efforts at the tooling and policy levels have yet to begin. This gap has been recognised in the EOSC AAI WG and there has been a clear decision that although the work towards the desired final architecture should continue without any delays, we need to provide practical solutions that can support the needs of today.

To be more specific, the high priority requirements recognised are the needs for enabling SSO across the first wave of EOSC Nodes that will be forming the EOSC Federation and executing workflows that utilise resources across multiple Nodes.

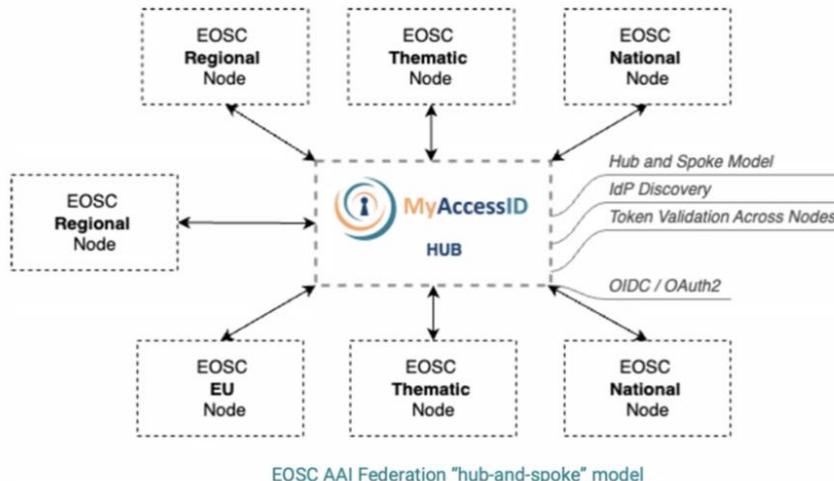
The design for this first implementation is guided by three core principles:

<https://doi.org/10.5281/zenodo.15388270>

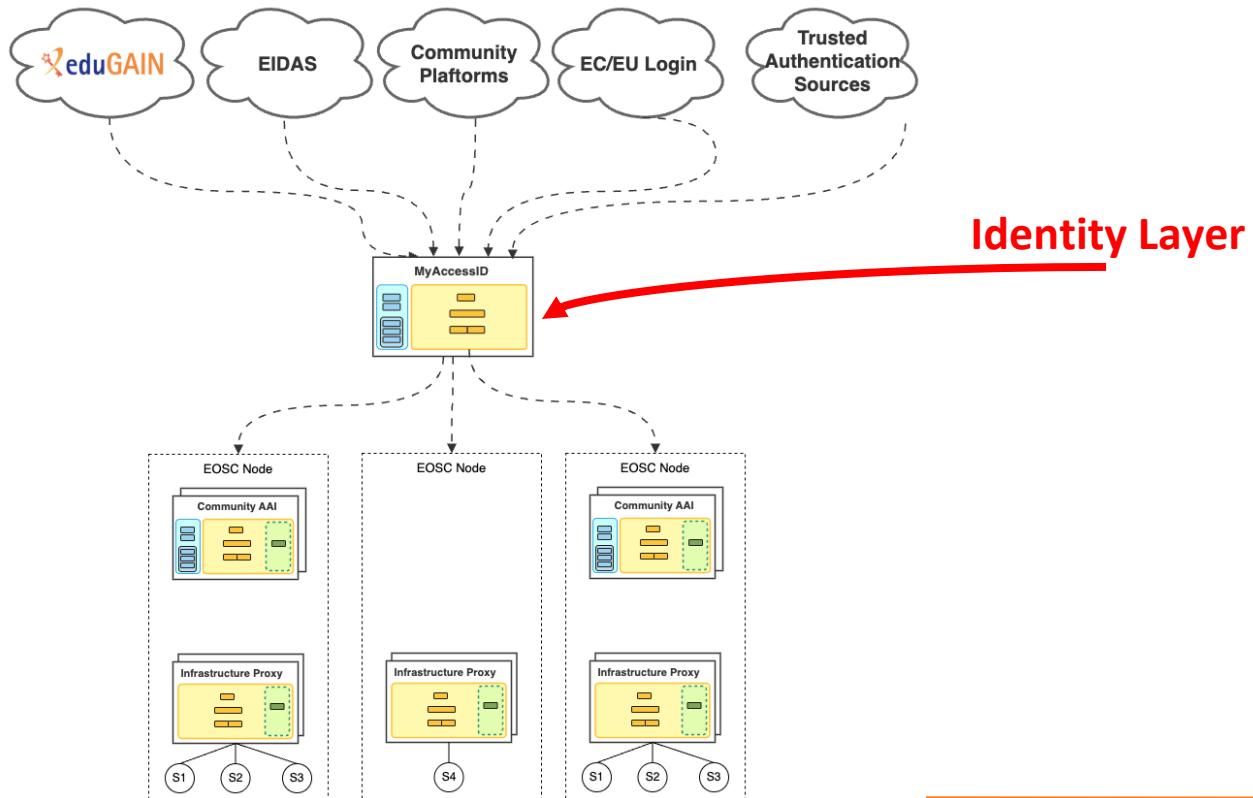


DOI 10.5281/zenodo.1538827
EOSC AAI Architecture 2025
May 12, 2025

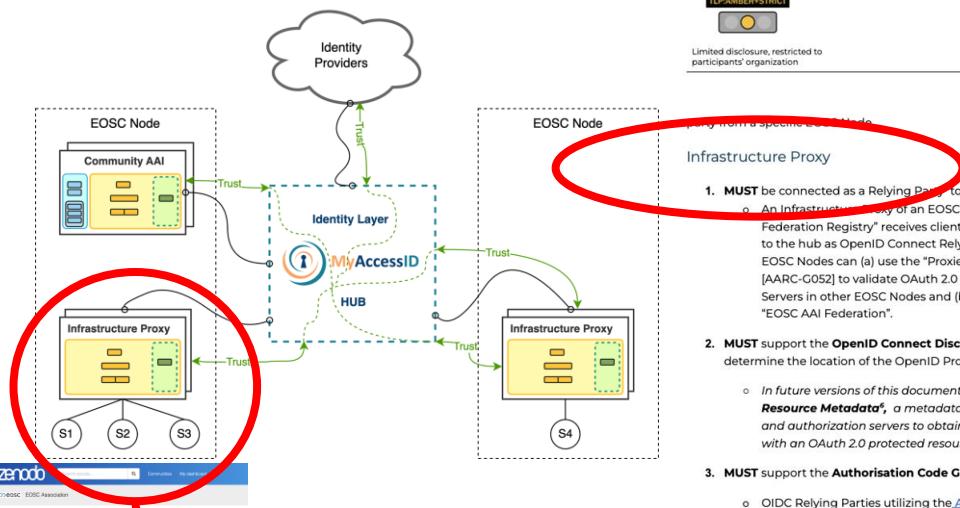
- The document presents **recommendations for the initial implementation of the EOSC AAI Federation**, offering background on prior work and summarising recent advancements, including updates to the AARC Blueprint Architecture.
- It is intended as a **practical guide for candidate EOSC Nodes**, outlining the steps necessary to connect with the EOSC AAI Federation. In the EOSC model, Nodes act as the primary integration points for services as it is described in the EOSC Federation Handbook, services are onboarded to individual Nodes rather than directly to the Federation.
- The overarching goal of the EOSC AAI Federation is to eventually support a **full-mesh, dynamic topology** without introducing a centralised component into the European AAI ecosystem.



EOSC AAI Federation - The Identity Layer



Infrastructure Proxy



4. MUST make use of the nonce Parameter

- o OIDC Relying Parties **MAY** use the nonce parameter (as specified in [\[OpenID_Connect_Core\]](#)[[OIDC-Core]]) alongside the corresponding nonce claim in the ID Token.
- o This further mitigates replay attacks and ensures the integrity of the authentication process.

5. MUST support Token Introspection

- **Compliance with OAuth 2.0 Token Introspection [RFC7662]**
Authorisation Servers in the "EOSC AAI Federation" **MUST** implement the OAuth 2.0 Token Introspection endpoint to verify the validity and active state of tokens they have issued.

- **Scope and Policy Enforcement**

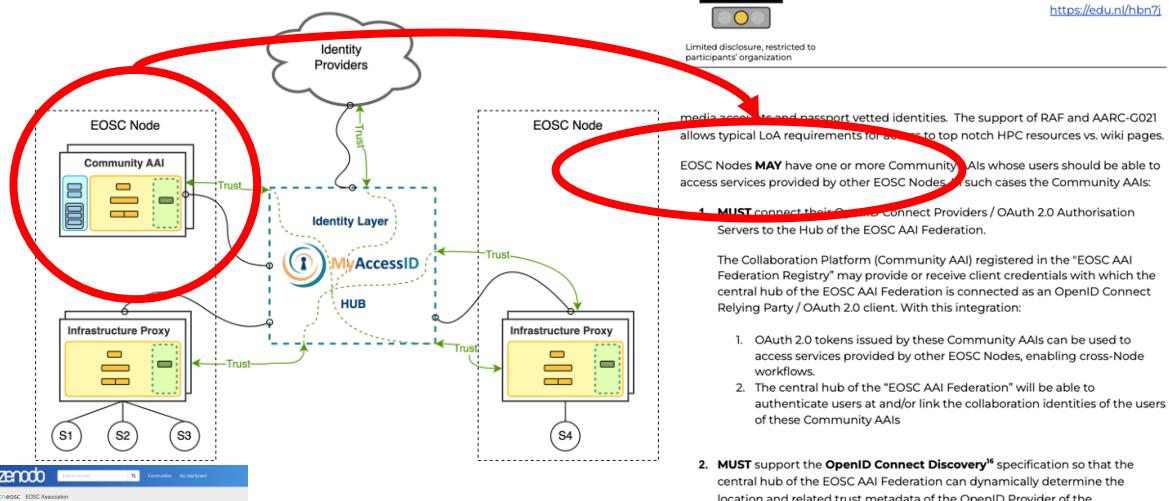
By performing token introspection, the EOSC Node's services can retrieve authorised scopes and user claims from the token response. This information **SHOULD** be used to enforce fine-grained access control policies and ensure the requestor's permissions match the resource's requirements.

8 Reduced Exposure of Access Tokens

EOSC Nodes **SHOULD** minimize exposing raw access tokens to various service components. Instead, they **SHOULD** rely on a dedicated component (e.g., the Infrastructure Proxy) to handle introspection, thus limiting security risks.

8. Token Revocation and Freshness

Regular introspection checks help detect revoked or expired tokens before granting access to protected resources. The EOSC Node **SHOULD** define a suitable caching or re-validation strategy (e.g., time-based) to balance performance with security needs.



zenodo

EOSC AAI Architecture 2025

Published May 12, 2025 | Version 1.0 | DOI: [zenodo.1520000](https://doi.org/10.5281/zenodo.1520000)

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Open Access

Attribution

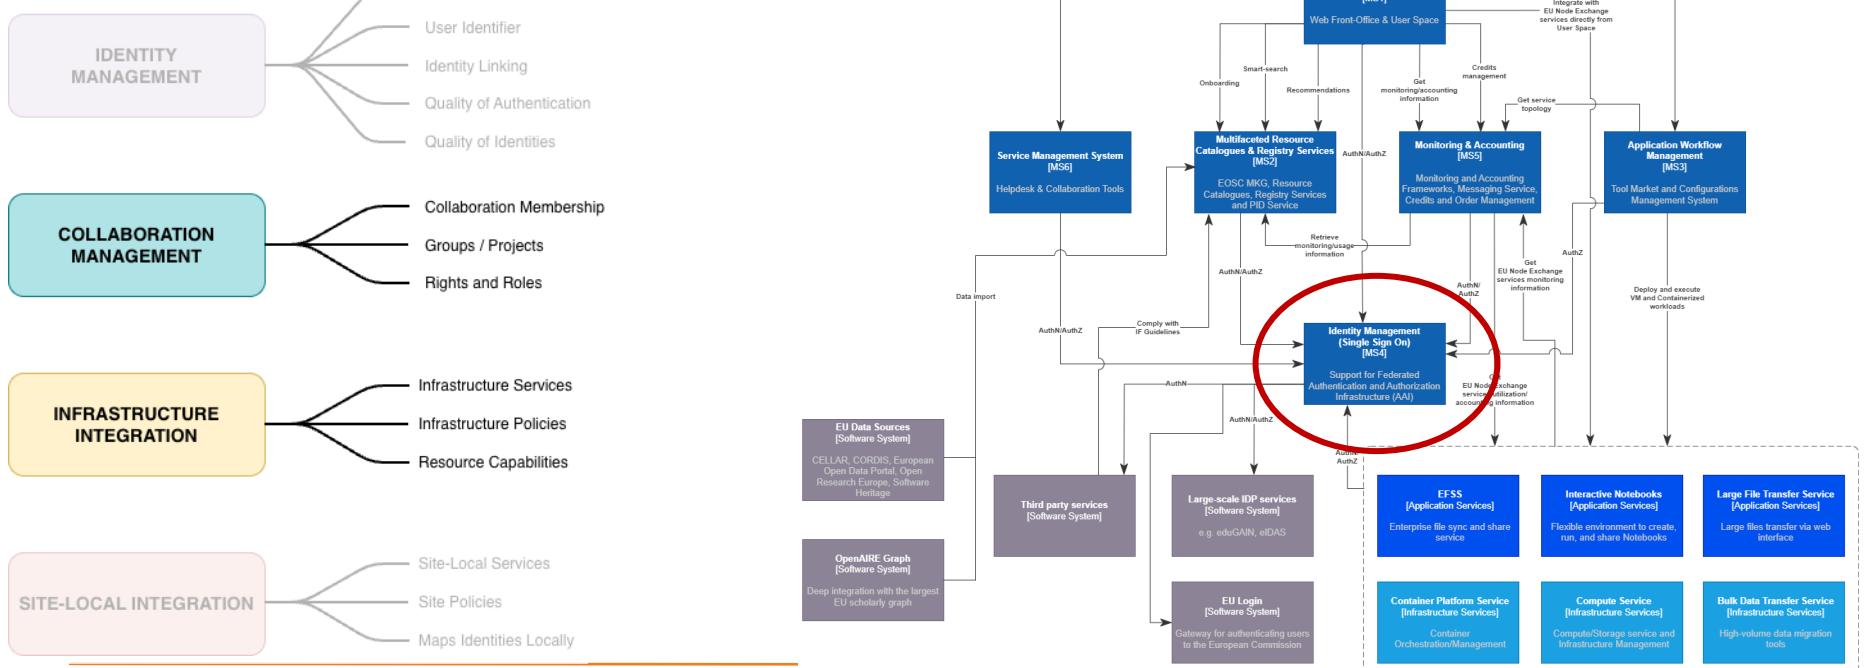
Non-Commercial

Share Alike

No Derivatives

CC-BY-NC-ND 4.0 International

Blueprint Architecture



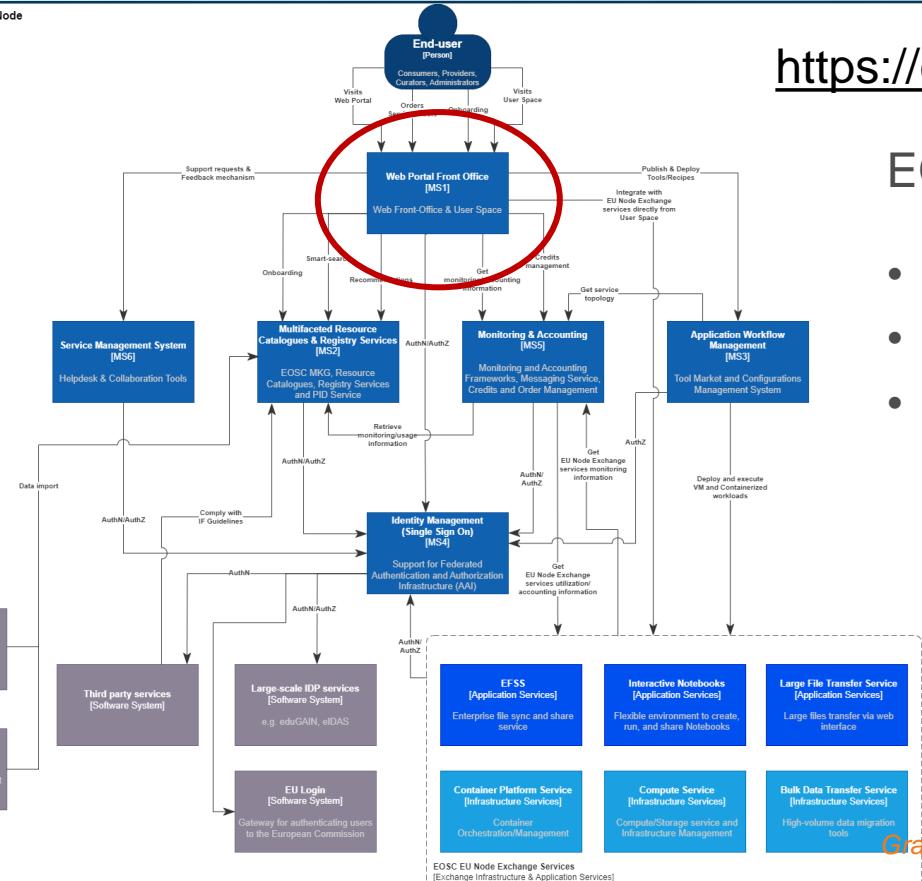
Graphics: Christos Kanellopoulos and the AARC Community

[System Landscape] EU Node

Conceptual view

Legend

Person
Lot-1 Managed Service
Lot-2 Managed Service
Lot-3 Managed Service
External Software System

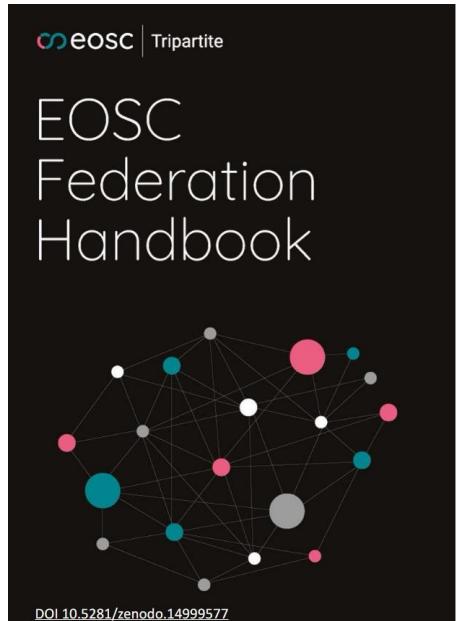


<https://open-science-cloud.ec.europa.eu/>

EOSC EU Node Web Portal

- Users register
- Users requests and manage projects
- Integrates with the EOSC EU Node AAI (SCIM)

Joining the EOSC Federation



Excerpt from the slides by Bob Jones (EOSC-A) for the November 2025 EOSC Symposium

Consensus on Technical Requirements

Mandatory Technical requirements

• Federated AAI

Adhere to the EOSC Node Federated AAI requirements as defined in EOSC AAI Architecture 2025 (March 2025) <https://doi.org/10.5281/zenodo.15388270>

• Federated catalogues

Register their Service catalogues and Research Product catalogues in the EOSC EU Node Resource Catalogue (September 2025) <https://doi.org/10.5281/zenodo.15516020>

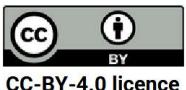
Recommended

- Application Workflow Management
- Service Monitoring
- Service and Research Product Accounting
- Order Management
- Helpdesk
- Management System

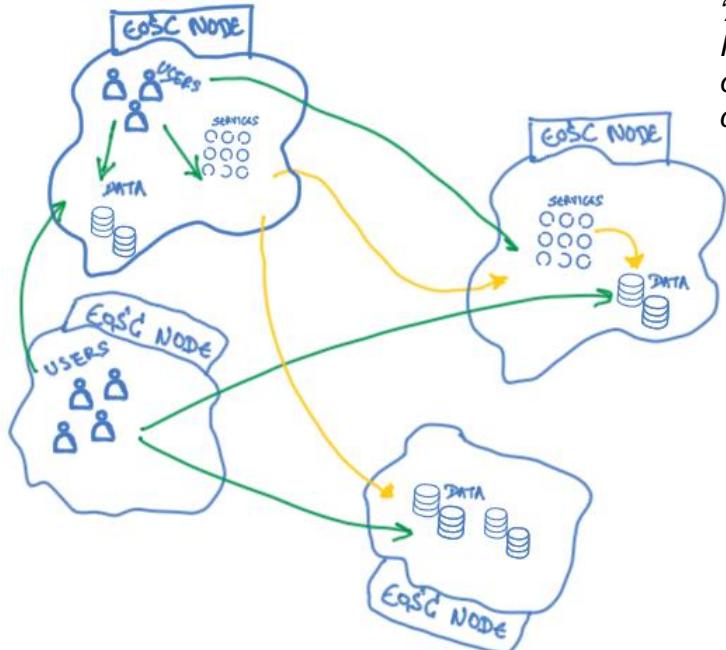
DOI [10.5281/zenodo.14999577](https://doi.org/10.5281/zenodo.14999577)

EOSC Federation Handbook

March 27, 2025

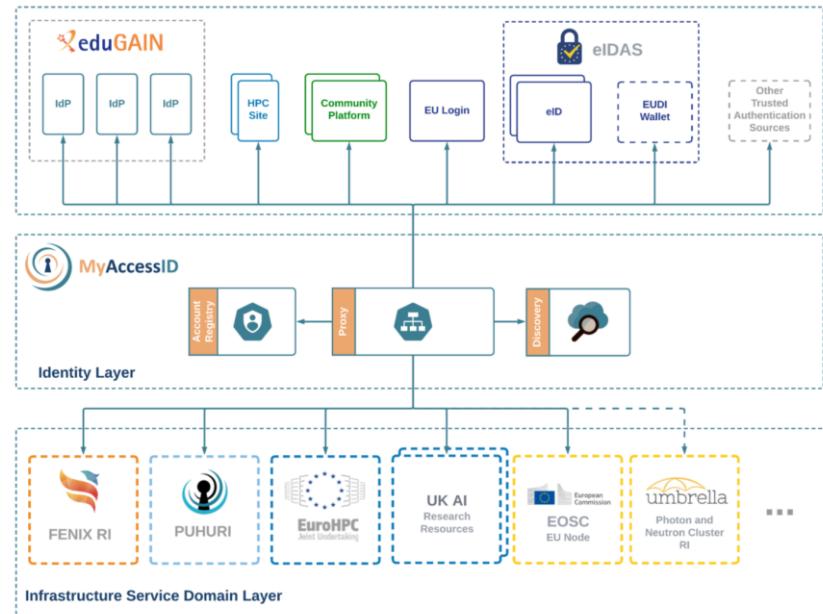


EOSC Federation structure and the evolution of its AAI

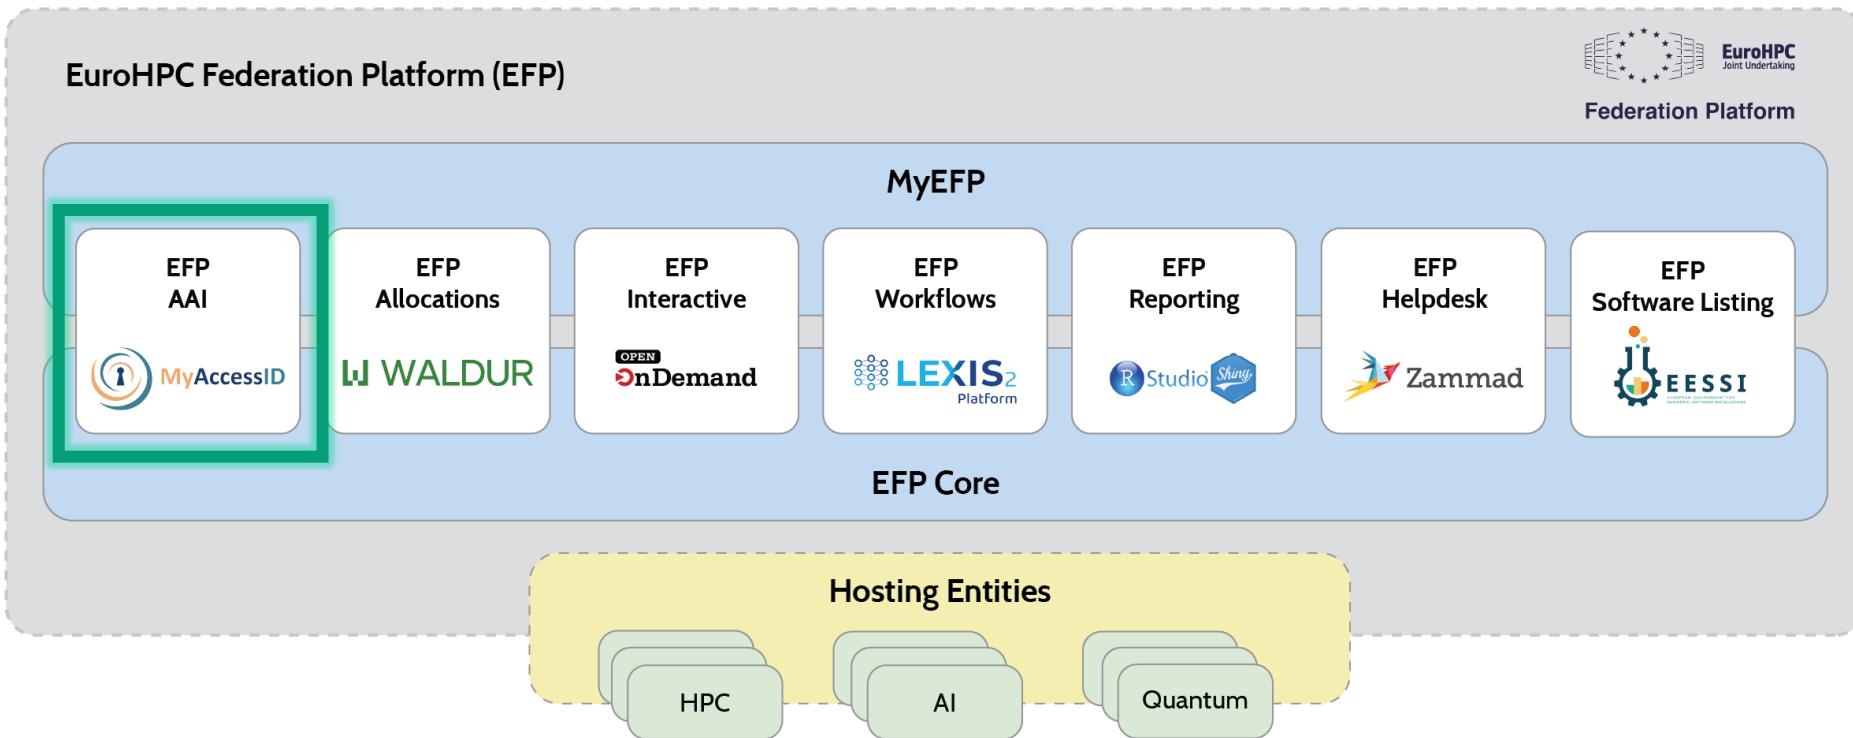


EOSC AAI Architecture 2025, <https://doi.org/10.5281/zenodo.15388269>
(EOSC AAI Working Group, 2025)

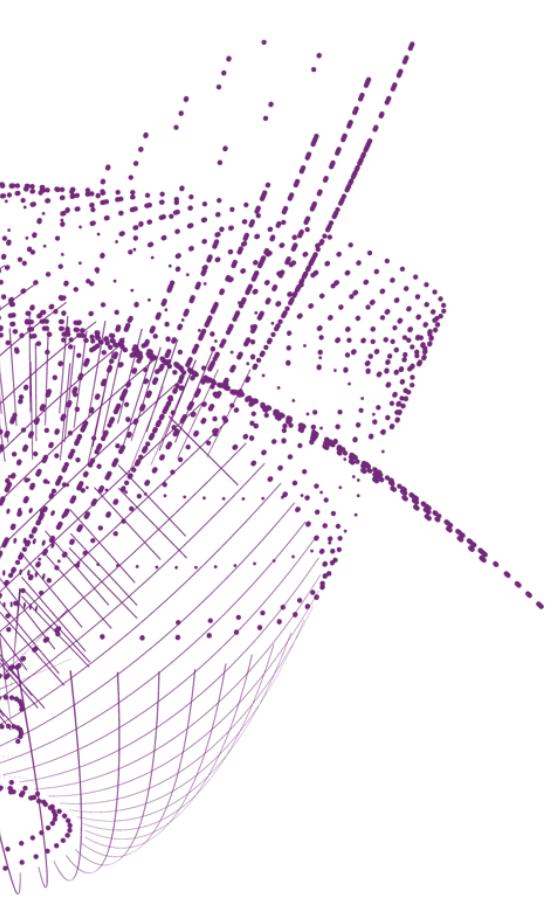
'Adopting the “hub-and-spoke” model in the initial phase of the EOSC AAI Federation is a practical step forward, and it is implementable today, while the design and development work for the OpenID Federation and “full-mesh” topologies continues in the background in the AARC Architecture WG and the EOSC AAI WG.'



And the EOSC is not alone in adopting this structure

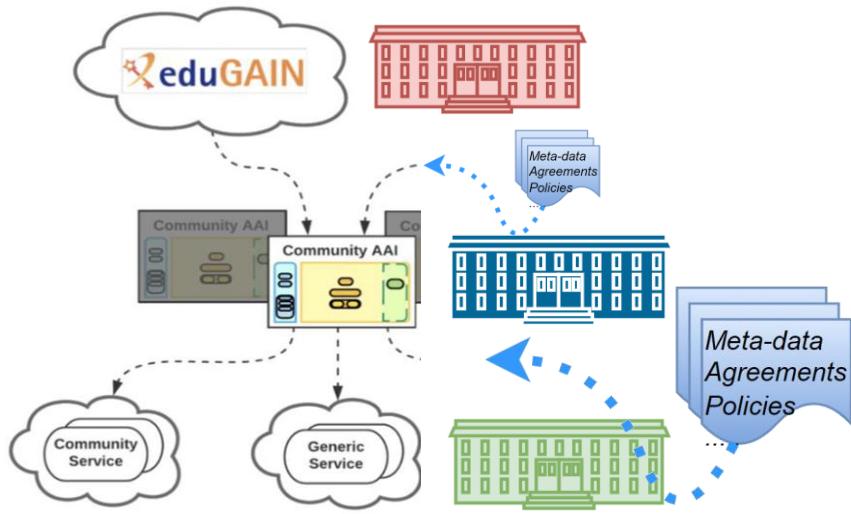


EuroHPC JU Federation Platform, see e.g. <https://my-eurohpc.eu/> (image retrieved from <https://my-eurohpc.eu/> February 2026)



Trust and the AARC Policy Development Kit

Now we need to 'decorate' the arrows with trust

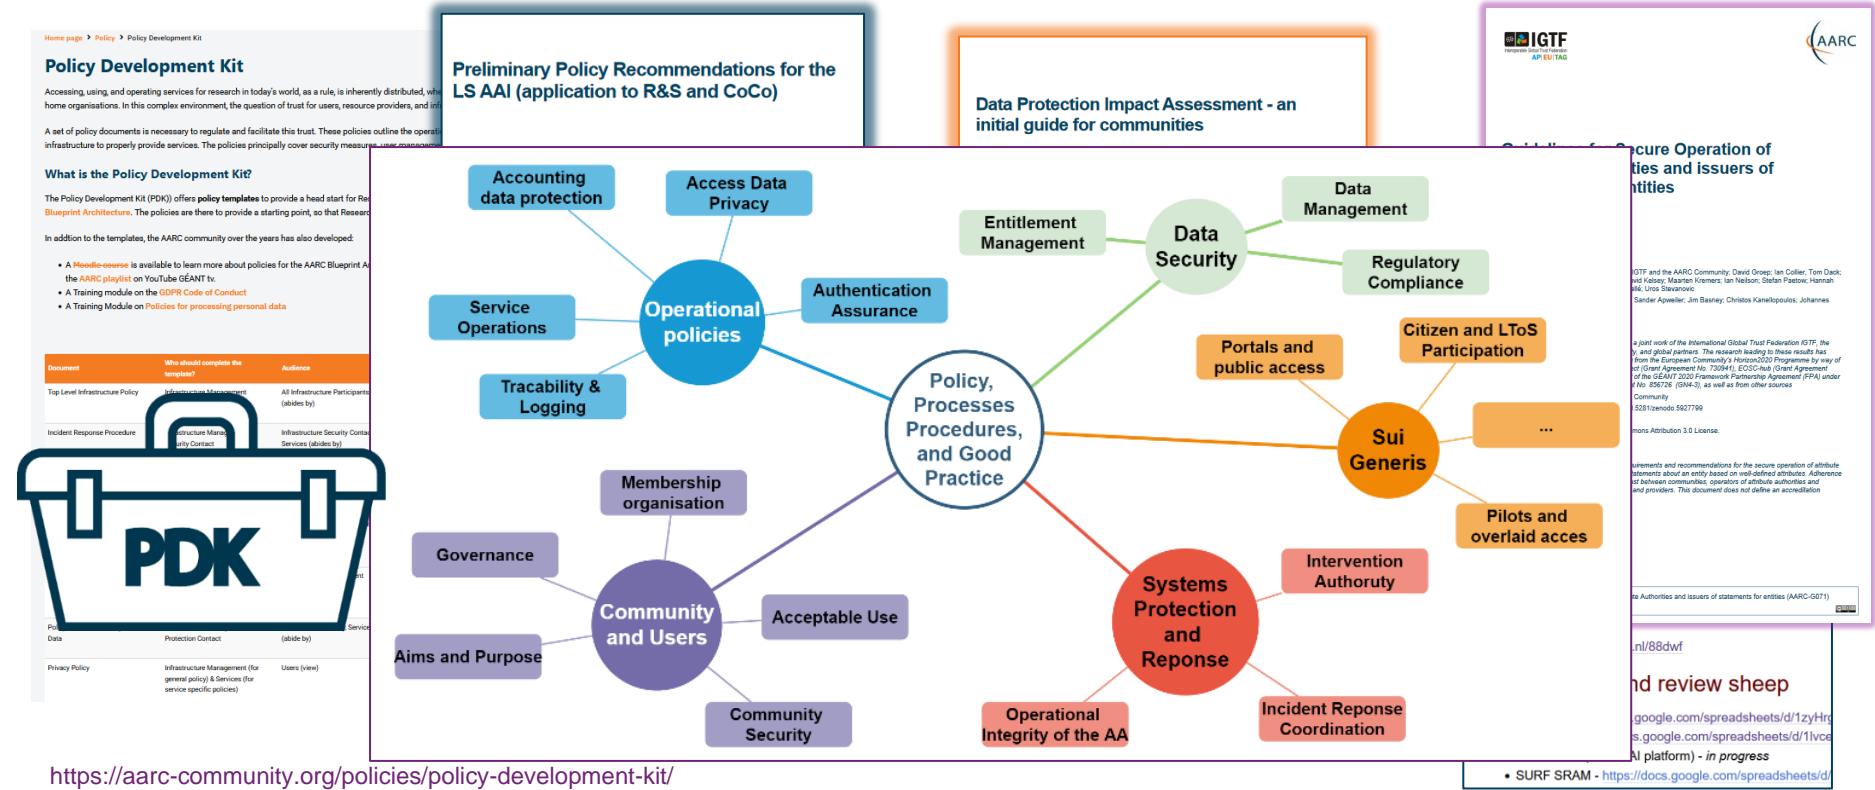


Each side of each arrow has *independent* parties

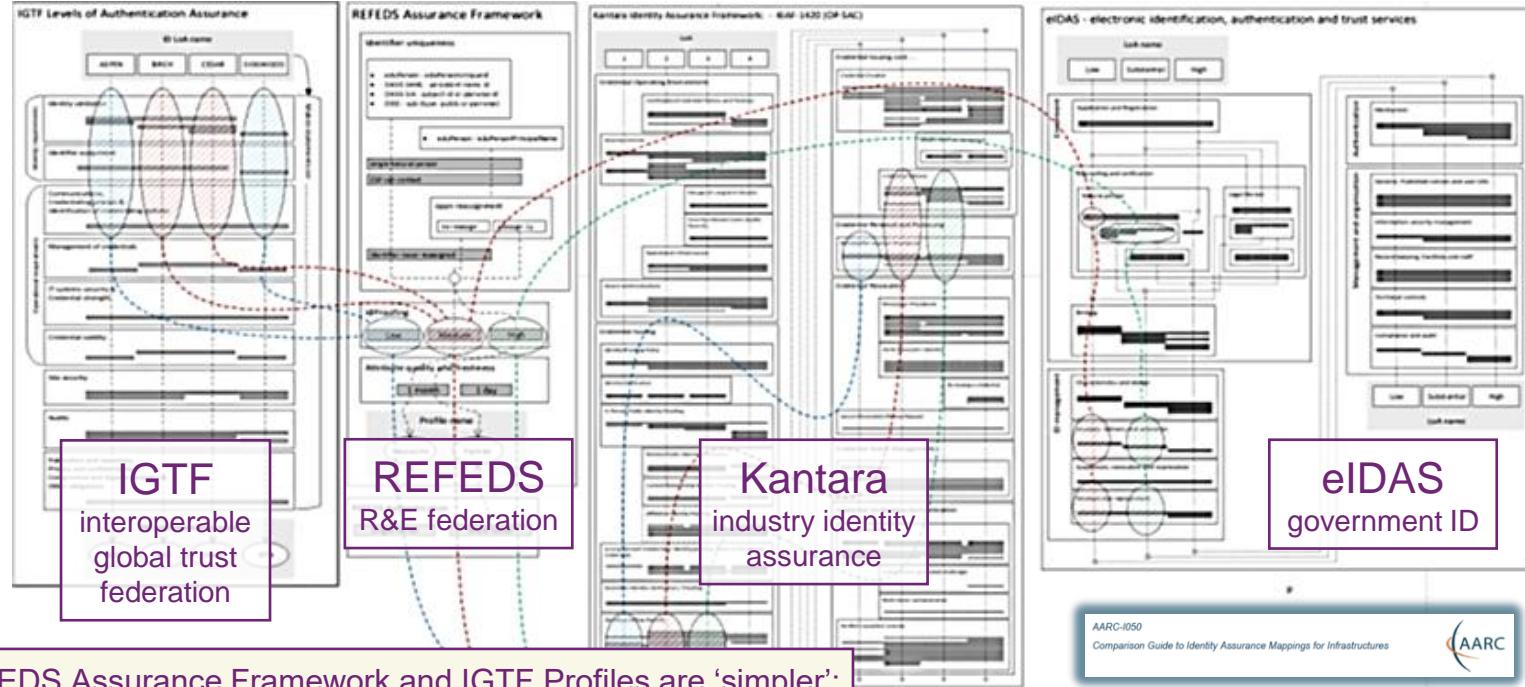
- we allow *them* to do part of the work we would otherwise do
- to make it easier and faster for users to perform their research
- but **we relinquish some control** beyond our organisation, our own policies, our own jurisdiction

Why would we trust them to do that?

Structuring trust ‘between boxes and arrows’ is complex!



And even a simple 'Who are you?' is not always easy ...



REFEDS Assurance Framework and IGTF Profiles are 'simpler':
academia is a higher-trust environment,
leveraging self-assessed peer review

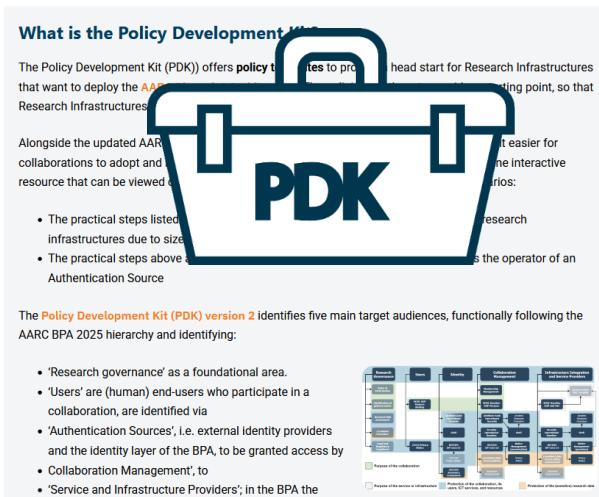
Source: <https://aarc-community.org/guidelines/aarc-i050>, Ian Neilson et al.

Developing the Trust framework, guidelines and best practice for BPA proxies and interaction with research services



minimise the number of divergent policies

empower identity providers, service providers, user communities to rely on interoperable policies



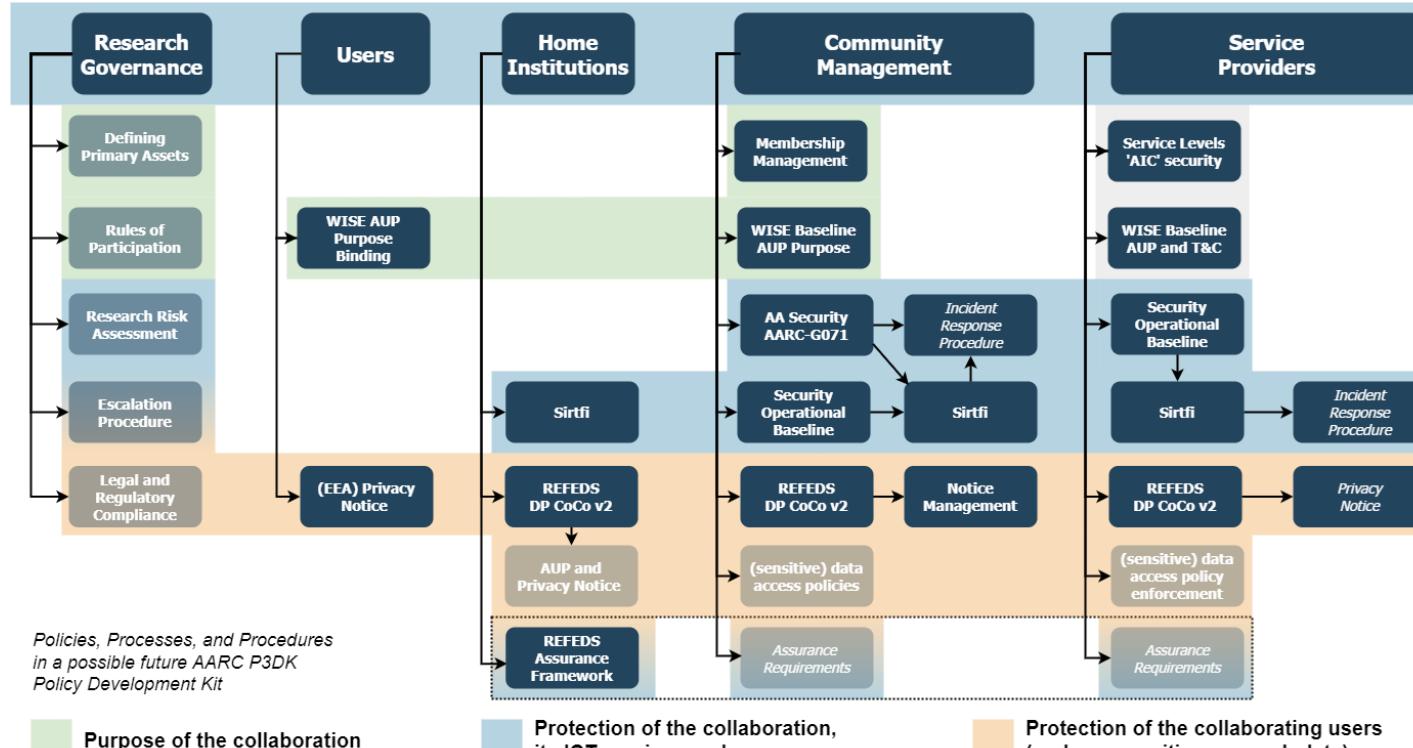
From the AARC2 infrastructure-oriented Policy Development Kit to

a simpler and deployment-oriented Policy, Process, and Procedure Development Kit version 2

- comprehensive review of existing policy suite to reduce complexity
- input from national research infrastructures and EOSC nodes, but not *only* in Europe but e.g. also Australia
- leverage the works we co-created with REFEDS and EOSC

<https://aarc-community.org/policies/policy-development-kit/>

Building the trust framework: development of the new full PDK structure



Today specialised AAI platform providers have established themselves

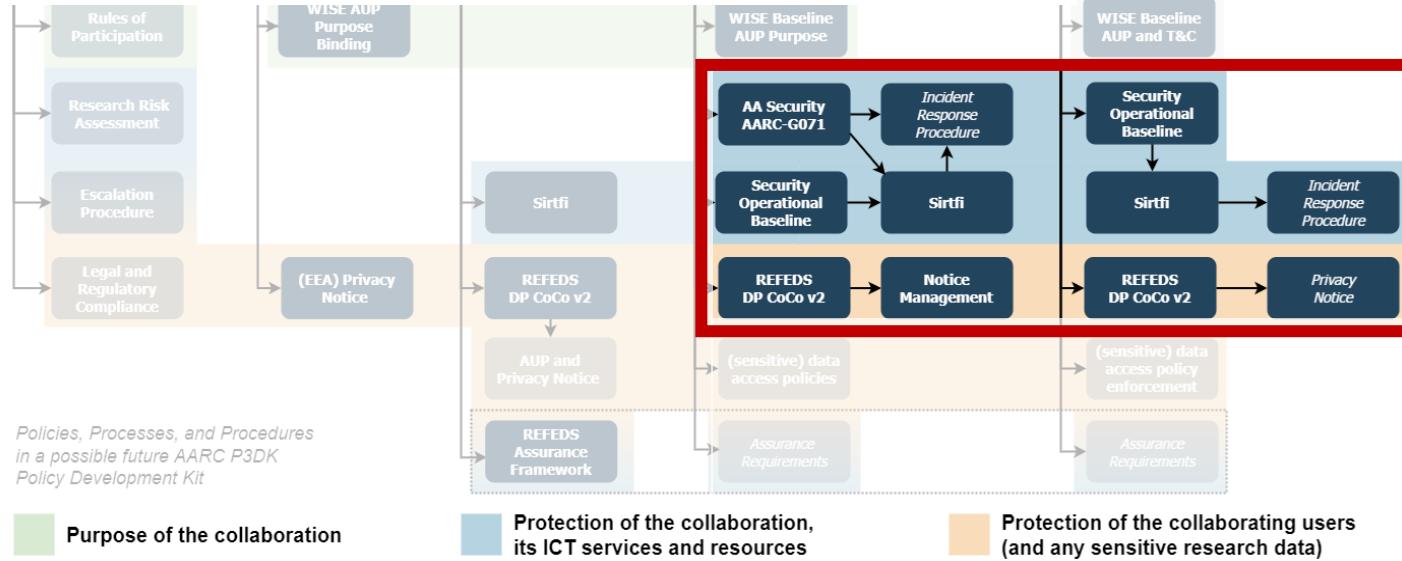
- Previous PDK policies targeted primarily at *infrastructure AAI*s and at *operators* of the few multi-community AAI
- BPA2025 identifies platform layers, and AAI platform *operators* serving many collaborations and infrastructures with a common layer are a key player today
- A 'trusted proxy operator' can now be either self-hosted or used 'as a service'

**This has changed the policy landscape:
the more complex policy implementations can now be 'sourced' from trusted providers**

AAI infrastructure providers for communities: a new 'Snctfi' trust mark

review and enhance effectiveness of Snctfi 'revamped'

*the set of guidelines that describe
a (self-) assessable baseline for the proxy operator
a set of service providers behind an AARC BPA Proxy*



Collaboration: foundational guidance

Practical steps to getting started with Policies for a Research Collaboration

Policy may appear a daunting or overly complex task if you start on your research collaboration journey, but with eight simple steps you can quickly navigate the policy space and avoid the most common pitfalls. Expand each step to learn the why and how of starting with your trusted collaboration quickly and smoothly:

- Define a unique name for your collaboration, preferably from the domain name system (DNS)
- Identify a governance body to make policy decisions
- Define the purpose of your collaboration - this will be used for your AUP
- Think about your crown jewels, risks, any regulations and legal things, privacy - and what to do if things go wrong ...
- Define the purpose of your collaboration - this will be used for your AUP
- **Why?** As you connect services and infrastructures to your collaboration via the AAI, these will have their 'acceptable' (and unacceptable) use defined. They provide services based on what you, as a collaboration, are planning to do, pay for, or because of shared goals and ambitions. Your users should be acting as part of your community, so also they need clarify what the collaboration is for. To prevent each and every infrastructure and service provider asking the users to comply with their acceptable use - and having to remember on your behalf what the collaboration's goal in life is - the common WISE Baseline AUP can do that in one go. But for that the purpose of use needs to be clear. Only you (as in: the collaboration) can provide that clarity
- **Recommendation:** be clear and concise in how to word your purpose. A one-line sentence is needed to be inserted verbatim into the WISE Baseline AUP that you should show to users enrolling in your collaboration (or that your AAI service provider will show on your behalf when new users join). This is not the place to write a grant proposal ...
- **Applicable guidance:** WISE AUP, AARC-I044 (AUP implementation guide), AARC-G083 (notice management), Governance - primary assets, Governance - risk assessment
- Think about your crown jewels, risks, any regulations and legal things, privacy - and what to do if things go wrong ...
- Define or adopt as-is the basic set of six policy documents for collaboration - and seek endorsement by your governance body
- **Why?** This basic set of 6 documents helps get a sufficient set of collaboration guidelines quickly - you can always adapt them later
- **Recommendation:** these are the documents you surely need - or you need to ask from your AAI provider:
 - Membership Management
 - Acceptable Use and Terms and Conditions
 - Privacy Notice
 - Attribute Authority operational security (AAOPC)

Pages / ... / Membership Management

Template for a Community Membership Management Policy

Created by David Groep, last updated on Jan 14, 2026 • 4 minute read

Template for a Membership Management Policy

Membership Management Policy for <Collaboration name>

This policy is effective from <insert date>.

The current collaboration manager can be found at <insert link>.

INTRODUCTION

This policy establishes practices that are adopted by <collaboration X> in the management of its members. Accurate management of a collaboration's members and their authorisation attributes is fundamental to ensuring secure access control. Trust between <collaboration X>, underlying infrastructure and partner collaborations may be established by rigorous application of this policy.

COLLABORATION MANAGER

<Collaboration X> defines a Collaboration Manager role and assigns this role to two or more individuals. The Collaboration Manager is responsible for meeting the requirements identified in this policy. This responsibility may be devolved to designated personnel in the Collaboration or in the Infrastructure, and their trusted agents (such as Institute Representatives or Resource Centre Managers).

MEMBERSHIP LIFE CYCLE REQUIREMENTS

More importantly:

AARC Guidelines series as a pathways to policy sustainability and impact



AARC-I082 Trust framework for proxies and Snctfi research services landscape analysis and structure

AARC-G083 Guidance for Notice Management by Proxies reducing user frustration by streamlining

AARC-G084 Security Operational Baseline trusted and secure infrastructure and incident response

AARC-I085 eID Assurance Model Assessment investigates capabilities for leveraging national eID

AARC-I086 Membership Management Policy Development at light-weight and infrastructure-level

AARC-PDK Policy Development Kit an interactive resource for jumpstarting collaboration

Cross-cutting guidelines

AARC-G080 Blueprint Architecture 2025 as the conceptual foundation

AARC-G081 Recommendations for Token Lifetimes balancing usage patterns and security

Adoption stimuli through the Policy Development Kit version 2 for

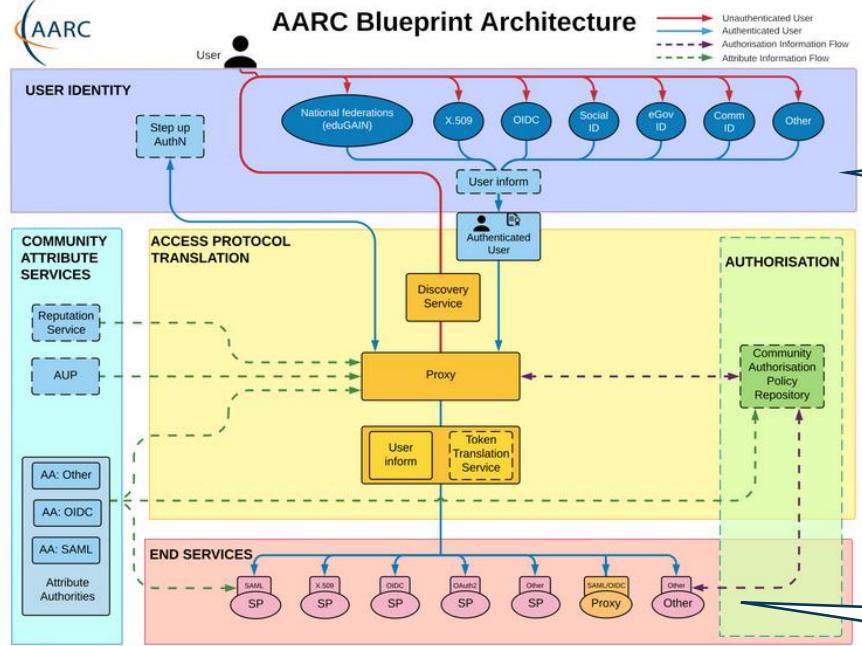
AARC-G071 'Attribute Authority and Proxy Operations', **AARC-I044 'Baseline AUP implementation'**

AARC-I051 and SIRTFI federated incident response, REFEDS DPCoCo v2, **AARC-G042 'DPIA'** for research collaborations, REFEDS Assurance Framework

<https://aarc-community.org/guidelines/>; PDK: <https://aarc-community.org/policy/policy-development-kit/>

Practices we already have, practices we need to harmonise

AARC-G071



Authentication/identity sources

NIST SP800-63

FIPS140

ISO 27001

IGTF AP Profiles

REFEDS MFA

REFEDS Assurance Framework

*so ... what about standards for the
Community Attribute Authority (AA)
or for operation of the Proxy?*

Service provider operations

ISO27k

NIS2

ITSRM2

How to establish secure operation for your (AARC BPA) proxy?

AARC-G071



The Challenge

- How to securely operate proxies, attribute authorities and issuers of statements for entities?

Guideline

- [AARC-G071 Guidelines for Secure Operation of Attribute Authorities](#)

Summary

- Operational security processes and procedures
- Requirements on traceability, auditability, and logging
- Requirements on the secure operation
- Requirements on securing the interactions



Guidelines for Secure Operation of Attribute Authorities and issuers of statements for entities

Publication Date

2022-04-11

Authors:

Members of the IGTF and the AARC Community; David Groep; Ian Collier, Tom Dack; Jens Jensen; David Kelsey; Maarten Kremers; Ian Neilson; Stefan Paetow; Hannah Short; Mischa Sallé; Uros Stevanovic

With feedback from

Marina Adomeit; Sander Apweiler; Jim Basney; Christos Kanellopoulos; Johannes Reetz

AARC Document Code: **AARC-G071**

Supported by:

This guideline is a joint work of the International Global Trust Federation IGTF, the AARC community, and global partners. The research leading to these results has received funding from the European Community's Horizon2020 Programme by way of the AARC2 project (Grant Agreement No. 730941), EOSC-hub (Grant Agreement 777536), as part of the GÉANT 2020 Framework Partnership Agreement (FPA) under Grant Agreement No. 856726 (GNA-3), as well as from other sources

Publishing Organisations: IGTF and AARC Community

DOI:

<https://doi.org/10.5281/zenodo.5927799>

Deployment guidance, self-assessment, and peer feed-back

AARC-G071



4.2. Attribute Management and Attribute Release

AMR-1

The Community must define and document the semantics, lifecycle, data protection, and release policy of attributes stored or asserted by the AA.

The community should follow the guidance from relevant policy documents. In particular, the Policy Development Kit has recommendations on Community Membership Management. It is recommended to use standardised attributes where possible, e.g. from eduPerson [EPSC] or SCHAC [SCHAC], and their semantics must be respected.

If Communities make modifications to the attribute set, their semantics, or release policies, it is recommended that they inform both their relying parties as well as the AA Operator thereof, since the AA operator may have implemented checks for schema consistency. The Community is ultimately responsible for the values and semantics of the attributes.

AMR-2

The AA Operator must implement the community definitions as defined and documented, for all the AAs it operates.

By implementing these requirements, the AA operator will support the chain of trust between Community and the RPs. An AA Operator must only host those communities for which it can implement the requirements.

AMR-3

It is recommended that the AA Operator provide a capability for the community to

Assessments and review sheet

- WLCG - <https://docs.google.com/spreadsheets/>
- UK-IRIS - <https://docs.google.com/spreadsheets/>
- eduTEAMS (Core AAI platform) - *in progress*
- SURF SRAM - <https://docs.google.com/spreadsheets/>
- NFDI - Academic ID - <https://docs.google.com/spreadsheets/>
- NFDI - didmos - <https://docs.google.com/spreadsheets/>
- NFDI - Reg App - <https://docs.google.com/spreadsheets/>
- NFDI - Unity - <https://docs.google.com/spreadsheets/>

<http://wiki.eugridpma.org/Main/AAOperationsGuidelines>



Proxy Operations: Information Security and Security Operational Baseline

AARC-G084



'address information security for disciplines and infrastructures - some of which process sensitive data'

Service Security Policy from AARC PDK v1
was successful but diverged in several directions:

- national implementations and specialisations
- included in EOSC Interoperability Framework as 'Security Operational Baseline'

The new PDK in AARC TREE converges on a common Baseline - with guidance and FAQ

- Included in the EOSC AAI WG Federation 2025



AARC-G084
Security Operational Baseline for Proxies and Services

3. Security Baseline

To adhere to the Security Operational Baseline, you must:

1. comply with the SIRTFF1 security incident response framework for structured and coordinated incident response
2. ensure that your Users agree to an Acceptable Use Policy (AUP) or Terms of Use, and that there is a means to contact each User.
3. promptly Inform Users and other affected parties if action is taken to protect their Service, or the Infrastructure, by controlling access to their Service, and do so only for administrative, operational, or security purposes.
4. honor the confidentiality and integrity of information gained as a result of your Service's participation in the Infrastructure.
5. respect the legal and contractual rights of Users and others with regard to the personal data processed, and only use access personal data for administrative, operational, accounting, monitoring or security purposes.
6. retain systematically deleted information (logs) in order to allow the reconstruction of a compromise and completion of actions as part of a security incident (the who, what, where, when, and to whom), for a minimum period of 180 days, to be used during the investigation of a security incident.
7. follow, as a minimum, generally accepted IT security best practices and governance, such as pro-actively applying secure configurations and security updates, and taking appropriate action in response to security vulnerabilities and threats, and agree to participate in drills and simulations to test Infrastructure resilience as a whole.
8. operate services and infrastructure in a manner which is not detrimental to the security of the Infrastructure nor to any of its Participants or Users.
9. collaborate in a timely fashion with others, specifically those with which there is a direct trust relationship, in the reporting and resolution of security events or incidents related to their participation in the Infrastructure and those affecting the Infrastructure as a whole.
10. honour obligations on security collaboration and log retention (clauses 1, 6, and 9 above) for the period of 180 days after the Service is retired from the Infrastructure, including the retention of logs when in virtual environments are decommissioned.
11. not hold Users or other Infrastructure participants liable for any loss or damage incurred as a result of the delivery or use of the Service in the Infrastructure, except to the extent of the law or any licence or service level agreement.
12. maintain an agreement with representatives for individual service components and suppliers that ensures that engagement of such parties does not result in violation of this Security Baseline.

Publication Date: 2025-03-29
Authors: David L. Grose (ed.), Al Moers, Daniel Kouli, Baptiste I. Kelsey, Jan Nelson, Linda Cornwall, Matt Viljoen, Pieter P. Nederbrugge, Romain Warot, Sven Gabril, and Uro P. Dzakovic

Document Code: AARC-G084
(to be assigned)

DOI:

Framework: Policy Development Kit v2

© by the authors and the AARC Community, 2003-2025

Abstract

The Security Baseline provides a reference set of minimum expectations and requirements for those offering services to users, communities, and other participants in a distributed system. It specifies the minimum security requirements for the provision of services, and the responsibilities of the participants in the system. It aims to trust between all Participants in the Infrastructure to enable reliable and secure operations.

Security Operational Baseline for Proxies and Services (AARC-G084)

Published 2025-03-28

The 12 points of AARC-G084

1. comply with the SIRTFI security incident response framework for structured and coordinated incident response [ref to SIRTFI]
2. ensure that your Users agree to an Acceptable Use Policy (AUP) or Terms of Use, and that there is a means to contact each User.
3. promptly inform Users and other affected parties if action is taken to protect their Service, or the Infrastructure, by controlling access to their Service, and do so only for administrative, operational or security purposes.
4. honour the confidentiality requirements of information gained as a result of your Service's participation in the Infrastructure.
5. respect the legal and contractual rights of Users and others with regard to the personal data processed, and only use access personal data for administrative, operational, accounting, monitoring or security purposes.
6. retain system generated information (logs) in order to allow the reconstruction of a coherent and complete view of activity as part of a security incident (the 'who, what, where, when', and 'to whom'), for a minimum period of 180 days, to be used during the investigation of a security incident.
7. follow, as a minimum, generally accepted IT security best practices and governance, such as pro-actively applying secure configurations and security updates, and taking appropriate action in relation to security vulnerability notifications, and agree to participate in drills or simulation exercises to test Infrastructure resilience as a whole.
8. operate services and infrastructure in a manner which is not detrimental to the security of the Infrastructure nor to any of its Participants or Users.
9. collaborate in a timely fashion with others, specifically those with which there is a direct trust relationship, in the reporting and resolution of security events or incidents related to their participation in the infrastructure and those affecting the infrastructure as a whole.
10. honour the obligations on security collaboration and log retention (clauses 1, 6, and 9 above) for the period of 180 days after their Service is retired from the Infrastructure, including the retention of logs when physical or virtual environments are decommissioned.
11. not hold Users or other Infrastructure participants liable for any loss or damage incurred as a result of the delivery or use of the Service in the Infrastructure, except to the extent specified by law or any licence or service level agreement.
12. maintain an agreement with representatives for individual service components and suppliers that ensures that engagement of such parties does not result in violation of this Security Baseline

FAQ and implementation guidance

<https://wiki.geant.org/spaces/AARC/pages/1049624759/view>

Pages / ... / AARC-G084 Security Operational Baseline

Security Operational Baseline FAQ and Recommendations

Created by David Groep on May 24, 2025 • 11 minute read

The Security Operational Baseline (AARC-G084) sets minimum expectations and puts requirements on the behaviour of those offering services to and on communities connected to a federated infrastructure, when interacting with the infrastructure peers and services. Worded in an intention concise manner, the 12 key requirements may give rise to additional questions, or in general can benefit from concrete examples and guidance. "FAQ" document, each of the key baseline items is put in context with additional examples, best practices, and generally helpful ideas.

- Can you elaborate on what is meant by item 9 and its incident response requirements?
- What are 'IT security best practices' in item 7?
- What does "honour the confidentiality requirements of information" in item 4 mean?
- What are "the legal and contractual rights of Users and others with regard to their personal data processed as part of service delivery" in item 5?
- "Retain system generated information (logs)" in item 6 sounds rather open-ended. What do I need to do? And why?
- "Aggregated centrally wherever possible, and protected from unauthorised access or modification" in item 6, how and why?
- Log aggregation in the layered and composite infrastructure
- What about the 'reconstruction of a coherent and complete view of activity' when you have a a 'layered technology stack' mentioned in item 6?
- What are "Named persons"?

Can you elaborate on what is meant by item 9 and its incident response requirements?

Item 3 talks about security incident response. In an interwoven environment it is vital that data about incidents is shared and communicated to detect, analyse, contain and eradicate malicious actors while preserving the necessary evidence for analysis and post-processing. For most infrastructures, there is a dedicated team of incident response specialists to aid with this task. This team can also communicate between different service providers affected by

[Home page](#) > [Guidelines](#) > AARC-G084

March 20, 2025

AARC-G084 Security Operational Baseline

The Security Baseline provides a reference set of minimum expectations and requirements of the behaviour of those offering services to users, communities, and other participants in a distributed proxy ecosystem, and of those providing access to services or assembling service components. It aims to establish a sufficient level of trust between all Participants in the Infrastructure to enable reliable and secure Infrastructure operation.

Document URL: <https://wiki.geant.org/download/attachments/999948380/AARC-G084-Security-Operational-Baseline-POKv2.pdf>

Development information: <https://wiki.geant.org/spaces/AARC/pages/999948380/AARC-G084+Security+Operational+Baseline>

Status: pending approval by AEGIS

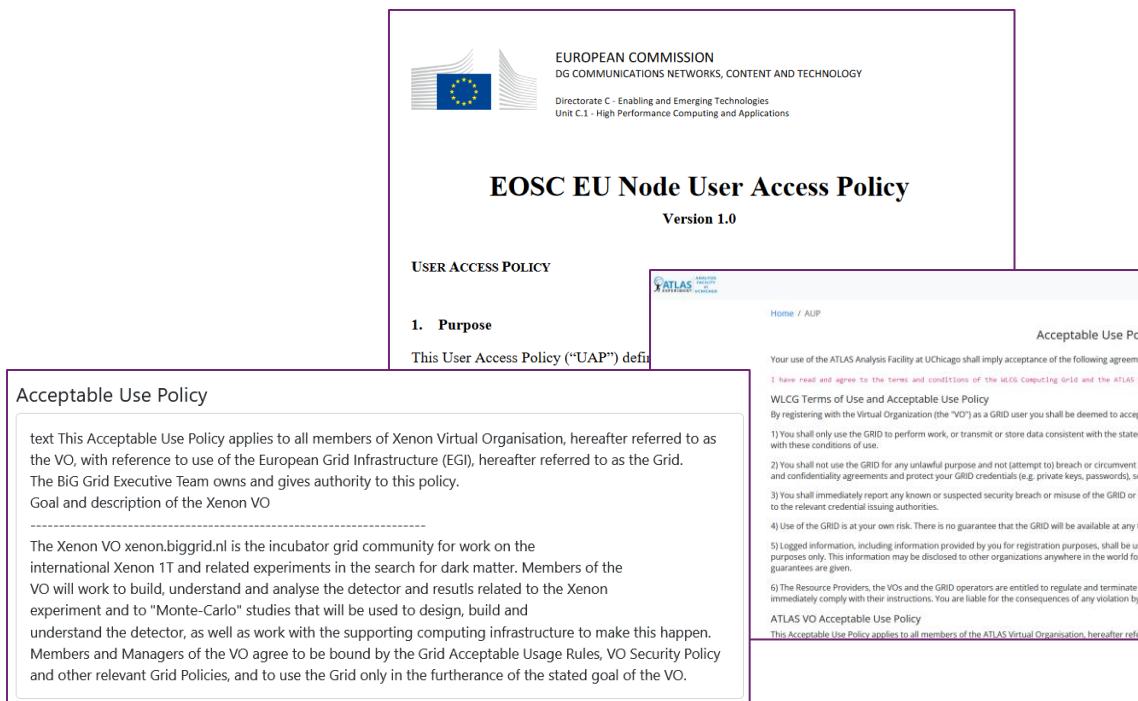
DOI: [10.5281/zenodo.17349890](https://doi.org/10.5281/zenodo.17349890)

Errata: none

Supersedes:

Supporting documentation, implementation suggestions and background information is available in the [Security Operational Baseline FAQ and Recommendations](#).

Helping community and users: how much clicking through?



Acceptable Use Policy

This Acceptable Use Policy applies to all members of Xenon Virtual Organisation, hereafter referred to as the VO, with reference to use of the European Grid Infrastructure (EGI), hereafter referred to as the Grid. The BiG Grid Executive Team owns and gives authority to this policy.

Goal and description of the Xenon VO

The Xenon VO xenon.biggrid.nl is the incubator grid community for work on the international Xenon 1T and related experiments in the search for dark matter. Members of the VO will work to build, understand and analyse the detector and results related to the Xenon experiment and to "Monte-Carlo" studies that will be used to design, build and understand the detector, as well as work with the supporting computing infrastructure to make this happen. Members and Managers of the VO agree to be bound by the Grid Acceptable Usage Rules, VO Security Policy and other relevant Grid Policies, and to use the Grid only in the furtherance of the stated goal of the VO.

EGI Configuration Database Acceptable Use Policy and Conditions of Use (AUP)

This Acceptable Use Policy and Conditions of Use ("AUP") defines the rules and conditions that govern your access to and use (including transmission, processing, and storage of data) of the resources and services ("Services") as granted by the EGI Federation, and the Virtual Organisation to which you belong, for the purpose of meeting the goals of EGI, namely to deliver advanced computing services to support researchers, multinational projects and research infrastructures, and the goals of your Virtual Organisation or Research Community.

1. You shall only use the Services in a manner consistent with the purposes and limitations described above; you shall show consideration towards other users including by not causing harm to the Services; you have an obligation to collaborate in the resolution of issues arising from your use of the Services.

2. You shall only use the Services for lawful purposes and not breach, attempt to breach, nor circumvent administrative or security controls.

3. You shall respect intellectual property and confidentiality agreements.

4. You shall protect your access credentials (e.g. passwords, private keys or multi-factor tokens); no intentional sharing is permitted.

5. You shall keep your registered information correct and up to date.

6. You shall promptly report known or suspected security breaches, credential compromise, or misuse to the security contact stated below; and report any compromised credentials to the relevant issuing authorities.

7. Reliance on the Services shall only be to the extent specified by any applicable service level agreements listed below. Use without such agreements is at your own risk.

8. Your personal data will be processed in accordance with the privacy statements referenced below.

9. Your use of the Services may be restricted or suspended, for administrative, operational, or security reasons, without prior notice and without compensation.

10. If you violate these rules, you may be liable for the consequences, which may include your account being suspended and a report being made to your home organisation or to law enforcement.

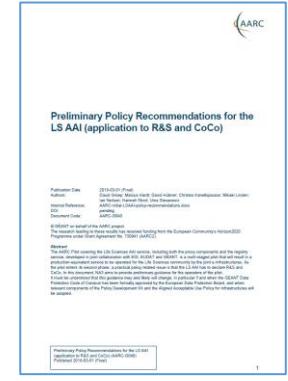
The administrative contact for this AUP is: operations@egi.eu
The security contact for this AUP is: abuse@egi.eu
The privacy notice is located at <https://gocdb-preprod.egi.eu/privacy.html>.

[Return to GOCDB homepage.](#)

Proxies have their 'experience challenges': AUPs, T&Cs, Privacy notices, ...

For large 'multi-tenant' proxies

- some subset users in some communities use a set of services – how to I present their Terms and Conditions, and their privacy policies, so that the users
 - only see the T&Cs and notices for services they will access
 - this does not need to be manually configured for each community
 - is automatically updated when services join



beyond bespoke guidance

as well as for **community and dedicated proxies**

- when new (sensitive) services join, who actually needs to see the new T&Cs?
- can we communicate acceptance of T&Cs to services even if 'we' are small and 'they' are large?

What is an acceptable user experience in clicking through agreements?
 What is most effective in exploiting the WISE Baseline AUP? What do you need?

With Fewer Clicks to More Resources!

Good common practice: the WISE Baseline AUP



Acceptable Use Policy and Conditions of Use

This Acceptable Use Policy and Conditions of Use ("AUP") defines the rules and conditions that govern your access to and use (including transmission, processing, and storage of data) of the resources and services ("Services") as granted by {community, agency, or infrastructure name} for the purpose of {describe the stated goals and policies governing the intended use}.

<To further define and limit what constitutes acceptable use, the community, agency, or infrastructure may optionally add additional information, rules or conditions, or references thereto, here or at the placeholder below. These additions must not conflict with the clauses 1-10 below, whose wording and numbering must not be changed. >

1. You shall only use the Services in a manner consistent with the purposes and limitations described above; you shall show consideration towards other users including by not causing harm to the Services; you have an obligation to collaborate in the resolution of issues arising from your use of the Services.
2. You shall only use the Services for lawful purposes and not breach, attempt to breach, nor circumvent administrative or security controls.
3. You shall respect intellectual property and confidentiality agreements.
4. You shall protect your access credentials (e.g. passwords, private keys or multi-factor tokens); no intentional sharing is permitted.
5. You shall keep your registered information correct and up to date.
6. You shall promptly report known or suspected security breaches, credential compromise, or misuse to the security contact stated below; and report any compromised credentials to the relevant issuing authorities.
7. Reliance on the Services shall only be to the extent specified by any applicable service level agreements listed below. Use without such agreements is at your own risk.
8. Your personal data will be processed in accordance with the privacy statements referenced below.
9. Your use of the Services may be restricted or suspended, for administrative, operational, or security reasons, without prior notice and without compensation.
10. If you violate these rules, you may be liable for the consequences, which may include your account being suspended and a report being made to your home organisation or to law enforcement.

<Insert additional numbered clauses here>

The administrative contact for this AUP is:

{email address for the community, agency, or infrastructure name}

The security contact for this AUP is:

{email address for the community, agency, or infrastructure security contact}

The privacy statements (e.g. Privacy Notices) are located at: {URL}

Applicable service level agreements are located at: <URLs>



Purpose binding

ensure use is as intended for access grant



Terms and Conditions

research data access conditions, permits, grant conditions



WISE Baseline AUP

common 10 commandments that allow seamless cross-sectoral user movement



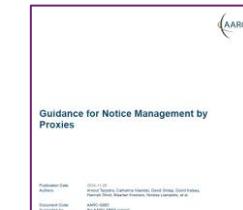
Service level agreements

promises and recourse



Privacy notice references

for access personal data policies



<https://wise-community.org/wise-baseline-aup/>

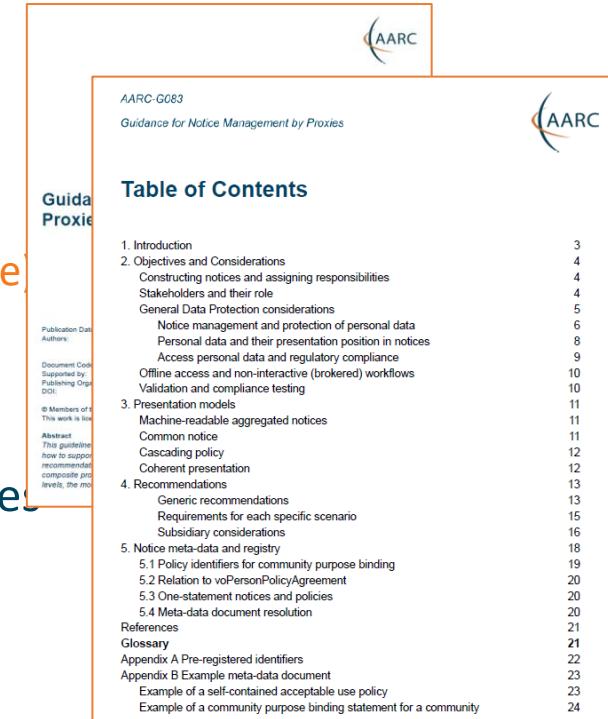
New AARC guidance on Notice Management by Proxies

Four presentation models In order of preference

1. machine-readable aggregated notice
2. common notice (single common authority domain)
3. cascading notices (assume responsibility for underlings)
4. coherent presentation: you show what you need (but not more)

Recommend WISE Baseline AUP plus model to
construct notices and communicate acceptance
 based on the AARC ID-community-infra hierarchy of proxies

- sufficient to build you a comprehensive WISE Baseline AUP
- and a set of privacy notices (for those GDPR encumbered)
- plus a namespace inspired by RFC6711's LoA registry

	AARC-G083 <i>Guidance for Notice Management by Proxies</i>																																																												
Table of Contents																																																													
Guidance for Notice Management by Proxies <i>Publication Date: [REDACTED]</i> <i>Author(s): [REDACTED]</i> <i>Document Code: [REDACTED]</i> <i>Supported by Publishing Organization DOI: [REDACTED]</i> <i>© Members of the AARC [REDACTED]</i> <i>Abstract: This document provides how to support recommendations for composite proxy levels, the model for how to support notices and how to support notices and policies.</i>	<table border="0"> <tr> <td>1. Introduction</td> <td>3</td> </tr> <tr> <td>2. Objectives and Considerations</td> <td>4</td> </tr> <tr> <td> Constructing notices and assigning responsibilities</td> <td>4</td> </tr> <tr> <td> Stakeholders and their role</td> <td>4</td> </tr> <tr> <td> General Data Protection considerations</td> <td>5</td> </tr> <tr> <td> Notice management and protection of personal data</td> <td>6</td> </tr> <tr> <td> Personal data and their presentation position in notices</td> <td>8</td> </tr> <tr> <td> Access personal data and regulatory compliance</td> <td>9</td> </tr> <tr> <td> Offline access and non-interactive (brokered) workflows</td> <td>10</td> </tr> <tr> <td> Validation and compliance testing</td> <td>10</td> </tr> <tr> <td>3. Presentation models</td> <td>11</td> </tr> <tr> <td> Machine-readable aggregated notices</td> <td>11</td> </tr> <tr> <td> Common notice</td> <td>11</td> </tr> <tr> <td> Cascading policy</td> <td>12</td> </tr> <tr> <td> Coherent presentation</td> <td>12</td> </tr> <tr> <td>4. Recommendations</td> <td>13</td> </tr> <tr> <td> Generic recommendations</td> <td>13</td> </tr> <tr> <td> Requirements for each specific scenario</td> <td>15</td> </tr> <tr> <td> Subsidiary considerations</td> <td>16</td> </tr> <tr> <td>5. Notice meta-data and registry</td> <td>18</td> </tr> <tr> <td> 5.1 Policy identifiers for community purpose binding</td> <td>19</td> </tr> <tr> <td> 5.2 Relation to voPersonPolicy/Agreement</td> <td>20</td> </tr> <tr> <td> 5.3 One-statement notices and policies</td> <td>20</td> </tr> <tr> <td> 5.4 Meta-data document resolution</td> <td>20</td> </tr> <tr> <td> References</td> <td>21</td> </tr> <tr> <td> Glossary</td> <td>21</td> </tr> <tr> <td> Appendix A Pre-registered identifiers</td> <td>22</td> </tr> <tr> <td> Appendix B Example meta-data document</td> <td>23</td> </tr> <tr> <td> Example of a self-contained acceptable use policy</td> <td>23</td> </tr> <tr> <td> Example of a community purpose binding statement for a community</td> <td>24</td> </tr> </table>	1. Introduction	3	2. Objectives and Considerations	4	Constructing notices and assigning responsibilities	4	Stakeholders and their role	4	General Data Protection considerations	5	Notice management and protection of personal data	6	Personal data and their presentation position in notices	8	Access personal data and regulatory compliance	9	Offline access and non-interactive (brokered) workflows	10	Validation and compliance testing	10	3. Presentation models	11	Machine-readable aggregated notices	11	Common notice	11	Cascading policy	12	Coherent presentation	12	4. Recommendations	13	Generic recommendations	13	Requirements for each specific scenario	15	Subsidiary considerations	16	5. Notice meta-data and registry	18	5.1 Policy identifiers for community purpose binding	19	5.2 Relation to voPersonPolicy/Agreement	20	5.3 One-statement notices and policies	20	5.4 Meta-data document resolution	20	References	21	Glossary	21	Appendix A Pre-registered identifiers	22	Appendix B Example meta-data document	23	Example of a self-contained acceptable use policy	23	Example of a community purpose binding statement for a community	24
1. Introduction	3																																																												
2. Objectives and Considerations	4																																																												
Constructing notices and assigning responsibilities	4																																																												
Stakeholders and their role	4																																																												
General Data Protection considerations	5																																																												
Notice management and protection of personal data	6																																																												
Personal data and their presentation position in notices	8																																																												
Access personal data and regulatory compliance	9																																																												
Offline access and non-interactive (brokered) workflows	10																																																												
Validation and compliance testing	10																																																												
3. Presentation models	11																																																												
Machine-readable aggregated notices	11																																																												
Common notice	11																																																												
Cascading policy	12																																																												
Coherent presentation	12																																																												
4. Recommendations	13																																																												
Generic recommendations	13																																																												
Requirements for each specific scenario	15																																																												
Subsidiary considerations	16																																																												
5. Notice meta-data and registry	18																																																												
5.1 Policy identifiers for community purpose binding	19																																																												
5.2 Relation to voPersonPolicy/Agreement	20																																																												
5.3 One-statement notices and policies	20																																																												
5.4 Meta-data document resolution	20																																																												
References	21																																																												
Glossary	21																																																												
Appendix A Pre-registered identifiers	22																																																												
Appendix B Example meta-data document	23																																																												
Example of a self-contained acceptable use policy	23																																																												
Example of a community purpose binding statement for a community	24																																																												

Notice presentation (PoC example implementation from the Validator)

Notice Presentation Component

<https://another-community.org>

Url: <https://another-community.org>

Description: A research community beyond suspicion.

Augments:

no description

<https://some-community.org>

Url: <https://some-community.org>

Description: A community somewhere researching for the betterment of mankind (hopefully)

Augments:

detector construction and experiment analysis for the search of dark matter using Xenon detectors

urn:doi:10.60953/68611c23-ccc7-4199-96fe-74a7e6021815

urn:idk:123456

Url: <https://the-community.org>

Description: no description

Augments:

Deze Gebruiksvoorwaarden betreffen het gebruik van netwerk en computers bij Nikhef. Iedere gebruiker van deze middelen of diensten wordt geacht op hoogte te zijn van deze voorwaarden en deze na te leven.

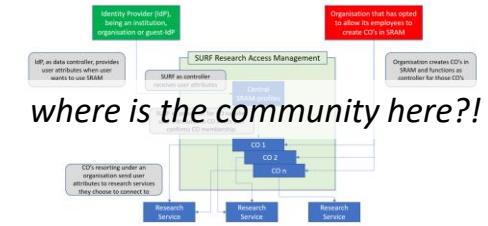
Agree **Reject**

Helping out the community – a simpler policy toolkit for communities

provide a revised policy development kit for mid-sized communities using the research infrastructures

Requirement from the AAI operators in FIM4R and BPA operators:

“small to mid-sized communities do not have the resources to maintain a bespoke community management policy”



But both communities and operators of membership management services are today unclear about trust assurance level of members: current templates in toolkit too complex and prescriptive

- develop ‘minimum viable community management’ for most small and mid-sized use cases
- give template and implementation guidance (FAQ) on community lifecycle management
- leverage complement of PDK practices that communities can ‘source’ from trusted providers

I086: Simplified Community Management policy - down to five items!

Each Community must

- Have a **unique name** (we recommend use DNS domain names)
- Require **members to accept an AUP** that defines the community goals and does not conflict with the Infrastructure AUP. It is recommended for the AUP to include the WISE Baseline AUP and follow the (AARC G083) notice management scheme
- Inform members about how their **personal information is processed**, follow local legal and regulatory requirements (e.g. by means of a Privacy Notice)
- Ensure its **members and their authorizations are valid** and enforced (e.g. who is an administrator and who is in which group)
- Be prepared for, and collaborate in, **security incident response**. You should be able to trace and take action on user accounts, and be prepared to participate in resilience exercises. Ensure that your provider can and will participate in incident response and meets security requirements including *Sirtfi* by providing contacts and sufficient logging.

PDK 2.0 Lightweight Community Security Policy

INTRODUCTION

Access to Infrastructure resources is commonly granted to members of a Community. To help protect those resources from damage or misuse, a Community has responsibilities in the manner it manages its membership and the way it behaves towards the Infrastructure. This policy aims to establish a sufficient level of trust to enable reliable and secure Infrastructure operation.

Guidance on this implementation is available in the [References and Notes](#) section, which may be updated from time to time, and does not form part of the effective policy.

DEFINITIONS

Entities identified by a leading capital letter in this document are defined in the Infrastructure Security Policy.

SCOPE

This policy applies to each Community whose members make use of the Infrastructure.

POLICY

Each Community must

1. Have a unique name -> recommend use DNS
2. Require members to accept an AUP that defines the community goals and does not conflict with the Infrastructure AUP. It is recommended for the AUP to include the WISE Baseline AUP and follow the (AARC G083) notice management scheme
3. Inform members about how their personal information is processed, follow local legal and regulatory requirements (e.g. by means of a Privacy Notice)
4. Ensure its members and their authorizations are valid and enforced (e.g. who is an administrator and who is in which group)
5. Be prepared for, and collaborate in, security incident response. You should be able to trace and take action on user accounts, and be prepared to participate in resilience exercises. Ensure that your provider can and will participate in incident response and meets security requirements including *Sirtfi* by providing contacts and sufficient logging.

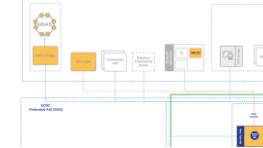
More diverse sources of researcher identity & assurance with eID wallets

Most reliable (and most 'available') source of assurance could be government identity!

- Step-up can now readily be done 'at home' by users through their national eID schemes
- eID wallets could solve the blockage by home IdPs to release assurance

... but their applicability to research and education use cases remains limited:

- eIDAS 1.0 suffers from inconsistent national uptake, asymmetrical cross-border connectivity, and protocol incompatibilities
- eIDAS 2 *at this point in time*, has incomplete roll-out, national implementations vary widely, and support for non-governmental use cases remains immature
- non-European users in Europe and international linking are not addressed at all today



Verifiable Credentials and digital wallets offer a complementary path forward, but lack of ecosystem maturity, lack of common standards, and adoption are (too) far in the future ...

Build trusted baseline expectations to increase reach of RI proxies

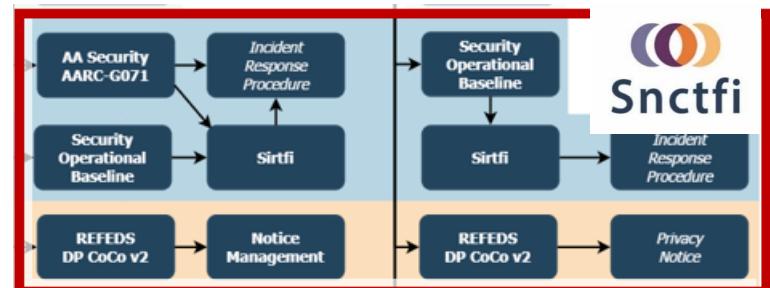
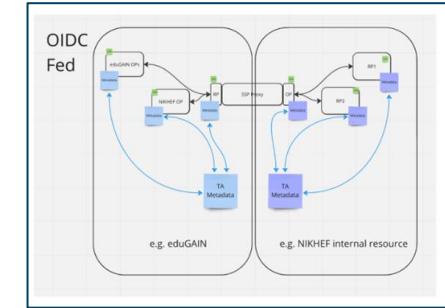
... with R&E identity providers ... and with new sources of information



Service and data providers need *unique identifier and affiliation*, with name and email, and 'fresh' assurance from home IdPs, but:

- proxies have met with scepticism by IdPs:
lack of even basic personalised and R&S attribute release
- how do these trust qualities 'traverse' proxies?
- how do operators rely on adherence to guidelines by their 'downstream' providers?

Position of the proxy makes trust bidirectional, and ***platform operators are facilitating this trust today***



Bringing it together: the Policy Development Kit

Practical steps to getting started with Policies for a Research Collaboration

Policy may appear a daunting or overly complex task if you start on your research. But with the Policy Development Kit, you can quickly navigate the policy space and avoid the most common pitfalls. Starting with your trusted collaboration quickly and smoothly:

- › Define a unique name for your collaboration, preferably from the domain name
- › Identify a governance body to make policy decisions
- › Define the purpose of your collaboration - this will be used for your AUP
- › Think about your crown jewels, risks, any regulations and legal things, privacy
- › Define or adopt as-is the basic set of six policy documents for collaboration
- › Review the AEGIS endorsed guidelines required for AARC compliance and ensure they are adopted
- › Ensure that the policies are presented to and accepted by the relevant audience
- › Publish your documents and responsible parties at a suitable location

Identify a governance body to make policy decisions

Define the purpose of your collaboration - this will be used for your AUP

Why? As you connect services and infrastructures to your collaboration via the AAI, these will have their 'acceptable' (and unacceptable) use defined. They provide services based on what you, as a collaboration, are planning to do, pay for, or because of shared goals and ambitions. Your users should be acting as part of your community, so also they need clarity as to what the collaboration is for. To prevent each and every infrastructure and service provider asking the users to comply with their acceptable use - and having to remember on your behalf what the collaboration's goal in life is - the common WISE Baseline AUP can do that in one go. But for that the purpose of use needs to be clear. Only you (as in: the collaboration) can provide that clarity

Recommendation: be clear and concise in how to word your purpose. A one-line sentence is needed to be inserted verbatim into the WISE Baseline AUP that you should show to users enrolling in your collaboration (or that your AAI service provider will show on your behalf when new users join). This is not the place to write a grant proposal ...

Applicable guidance: WISE AUP, AARC-I044 (AUP implementation guide), AARC-G083 (notice management), Governance - primary assets, Governance - risk assessment

Think about your crown jewels, risks, any regulations and legal things, privacy - and what to do if things go wrong ...

Define or adopt as-is the basic set of six policy documents for collaboration - and seek endorsement by your governance body

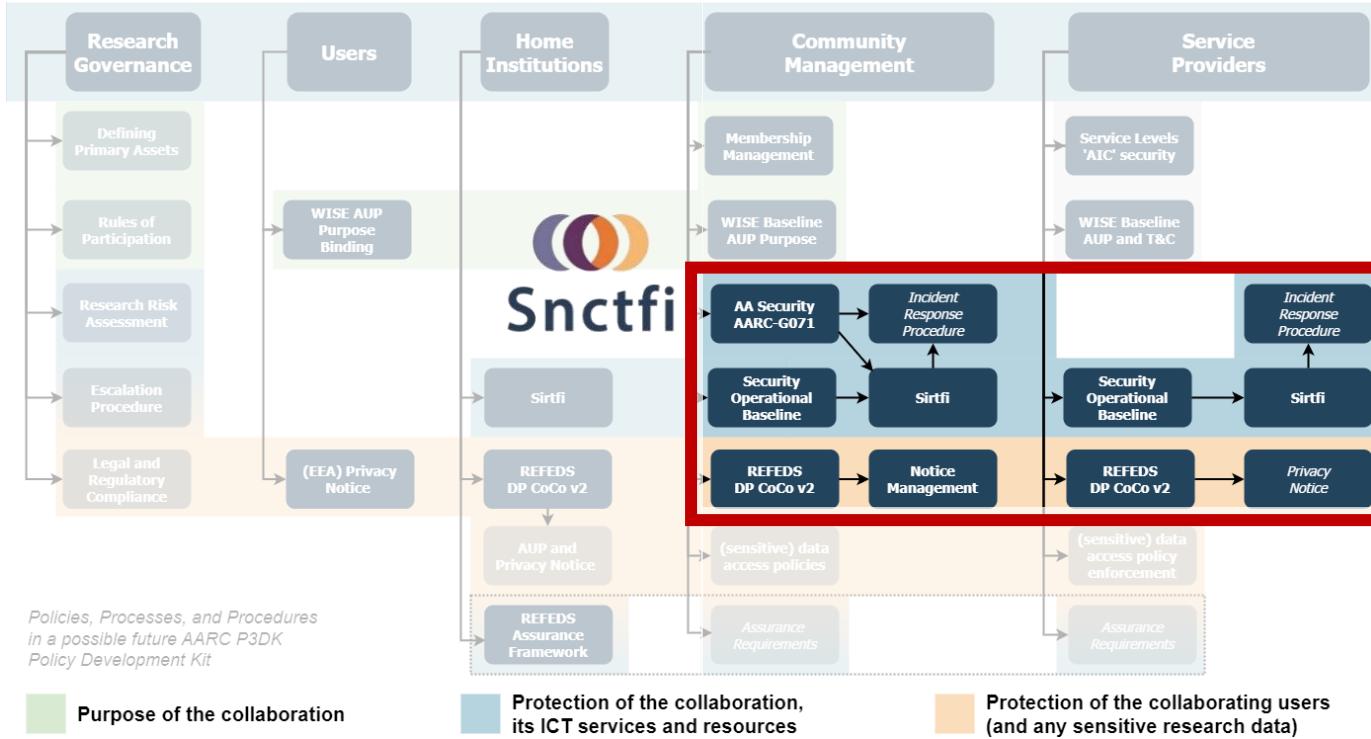
Why? This basic set of 6 documents helps get a sufficient set of collaboration guidelines quickly - you can always adapt them later

Recommendation: these are the documents you surely need - or you need to ask from your AAI provider:

- Membership Management
- Acceptable Use and Terms and Conditions
- Privacy Notice
- Attribute Authority operational security (AAOPS)

<https://aarc-community.org/policies/policy-development-kit/>

Providers manage complexity for research communities



communities sourcing
'well-operated'
community platforms



and a few more ...

through their scale gets
federations to trust our
AARC 'middle boxes'

Towards the AARC Compendium



AARC

AARC ▾ Architecture Policy Guidelines AARC TREE

[Home page](#) > AARC TREE Compendium Workshop – a summary



AARC TREE Compendium Workshop – a summary

[September 29, 2025](#)

On 17 September 2025, the AARC-TREE project held a dedicated workshop to gather community input on the upcoming release of the Compendium of Best Practices and Recommendations. The event, hosted at CERN and co-located with the OSCARS Composability Workshop, focused on refining the Compendium's structure, scope, and usability for diverse audiences.

<https://aarc-community.org/aarc-tree-compendium-workshop-a-summary/>

[Home page](#) > First Draft of the AARC Compendium Released for Community Feedback

AARC COMPENDIUM

Released



Please contribute your comments directly to the [AARC Compendium](#) on GitHub.

First Draft of the AARC Compendium Released for Community Feedback

[November 15, 2025](#)

AARC TREE has released the first draft of the AARC Compendium, an introductory guide to implementing federated identity management for research infrastructures and their communities.

Based on the [AARC Blueprint Architecture \(AARC BPA\)](#), the Compendium provides a practical overview of how to design and operate Authentication and Authorisation Infrastructures (AAIs). It includes a glossary of key terms, an FAQ section, and guidance on topics such as implementation scenarios, technical and policy requirements, security, and data protection.

<https://aarc-community.org/first-draft-of-the-aarc-compendium-released-for-community-feedback/>

AARC Compendium - AARC - X +

wiki.geant.org/spaces/AARC/pages/1278607380/AARC+Compendium

Spaces Blogs

Search Log in

WP2: Trustworthy Harmonisation and Interoperability

- WP3: Use Cases Collection and Analysis
- WP4: Adoption and Validation
- WP5: Compendium & Recommendations
 - (Draft) Compendium
 - AARC Compendium**
 - 1. What is a Research Collaboration AAI?
 - 2. What is the AARC Blueprint Architecture?
 - 3. Benefits and Value Proposition
 - 4. AARC Guidelines and Compliance
 - 5. How to implement an AARC Compliant AA
 - 6. Available Software and Services
 - 7. Case Studies
 - 8. Summary & Recommendations
 - 9. Glossary
 - 10. FAQ
 - Compendium Diagram
 - Ongoing Notes WP5- WP6: Bootstrap Communication and Exploitation

Pages / ... / **WP5: Compendium & Recommendations**

AARC Compendium

Created by [Hannah Short](#), last updated by [Sally Chambers](#) on [Jan 28, 2026](#) • 1 minute read

Introduction

The Authentication and Authorisation for Research and Collaboration (AARC) Compendium is intended as an introductory guide to implementing federated identity management for research collaborations, based on the AARC Blueprint Architecture (AARC BPA). This guide provides an introduction to Authentication and Authorisation Infrastructure (AAI). A glossary of key terms and their definitions is provided as well as a list of Frequently Asked Questions (FAQs).

The compendium covers a number of different topics, including: what is the AARC Blueprint and why has it been developed, how to implement an AAI service outlining a number of implementation scenarios. An overview of the landscape of existing AAI solutions is provided including commonly used software and services as well as hosted services.

Specific topics such as technical requirements, security, data protection and policy related issues are covered, including how to build the necessary bridges between legal, policy and technology.

Audiences

This guide is written with several audiences in mind. Wherever relevant, information is presented in multiple ways to best suit these audiences.

- Research Community Management **AUDIENCE: RESEARCH COMMUNITY MANAGEMENT**
- AAI Implementors and Operators **AUDIENCE: AAI IMPLEMENTORS AND OPERATORS**
- Funding Agencies **AUDIENCE: FUNDING AGENCIES**
- All **AUDIENCE: ALL**

Space tools

Back to top

Compendium Outreach Campaign: Engagement through Use Cases



 **SCIENCE-CLUSTERS.eu**
Research Infrastructures for Open Science

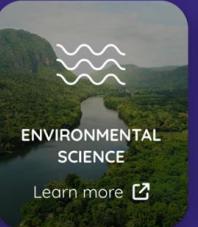
[About](#) [Joint Strategy](#) [Publications](#)

Research Infrastructures and Communities

The science clusters have grown out of five collaborative projects funded by the European Union in 2019 to link ESFRI and other world-class Research Infrastructures (RIs) to the European Open Science Cloud (EOSC). The services developed by the clusters and other outcomes of the projects are cornerstones of the emerging EOSC fabric and support both disciplinary communities and multidisciplinary initiatives with harmonised models for access to data, tools, workflows and training. Each cluster unites multiple RIs in their specific scientific domain.



ASTRONOMY AND PARTICLE PHYSICS
[Learn more](#)



ENVIRONMENTAL SCIENCE
[Learn more](#)



LIFE SCIENCE
[Learn more](#)



PHOTON AND NEUTRON SCIENCE
[Learn more](#)



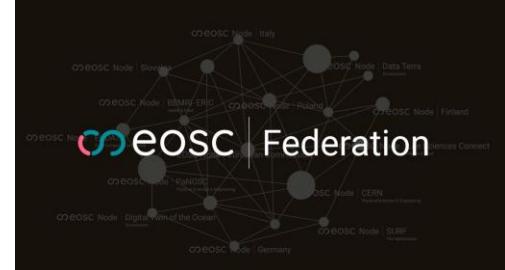
SOCIAL SCIENCES AND HUMANITIES
[Learn more](#)

<https://science-clusters.eu/>

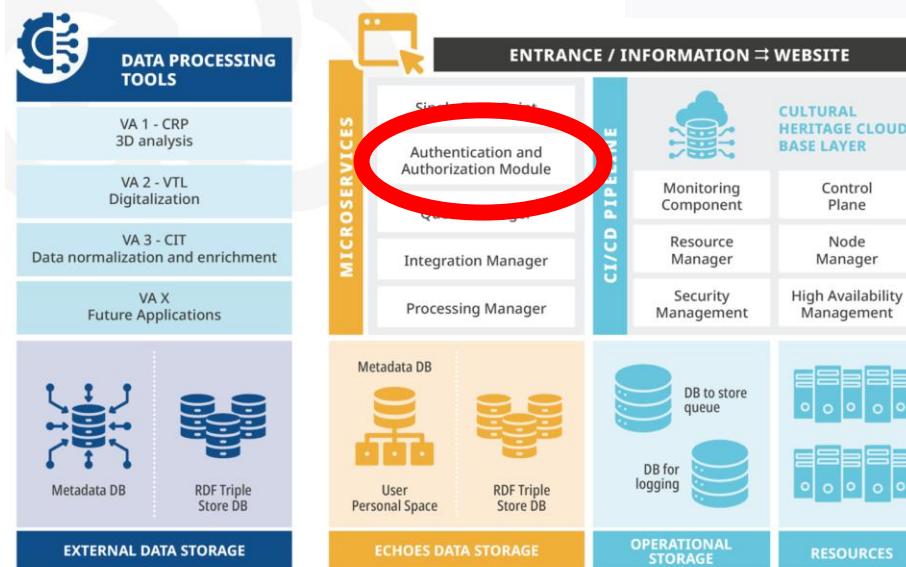
- Highlight Compendium within various communities, e.g. compendium as a resource in the [SSH Open Marketplace](#)



<https://oscars-project.eu/>



Compendium Outreach Campaign: Engagement through Use Cases



<https://www.echoes-eccch.eu/infrastructure/>



<https://www.echoes-eccch.eu/>

ECHOES Integration Task Force (EITF)

Kotzinos, D., Chambers, S., Barbot, L., Durco, M., & Dalla Torre, G. (2025). *ECHOES Integration Strategy for datasets, tools and workflows with potential for reuse in ECHOES*
<https://doi.org/10.5281/zenodo.17751335>

Sister Projects ▾

AUTOMATA

TEXTaiLES

HERITALISE

COLOURS

EXCALIBUR

MusicSphere

PlaceMUS XR

StratiGraph

UNICHE

ARXIVE

KINETIKA

INFINITY

Recommendations: Key messages

- There is benefit to investing in common AAI solutions, both for funding agencies and research collaborations themselves
- AAI is complex and our community generally recommends using a hosted/managed solution rather than starting from scratch. This will aid future interoperability as the landscape is highly dynamic.
- A “consulting” service where Research Collaborations can seek advice would be highly useful
- Many thanks to **all** contributors. Particular mention to the **Australian Access Federation (AAF)**.

Visit

<https://wiki.geant.org/spaces/AARC/pages/1278607380/AARC+Compendium>

The AARC Blueprint – a very digestible architecture ... so



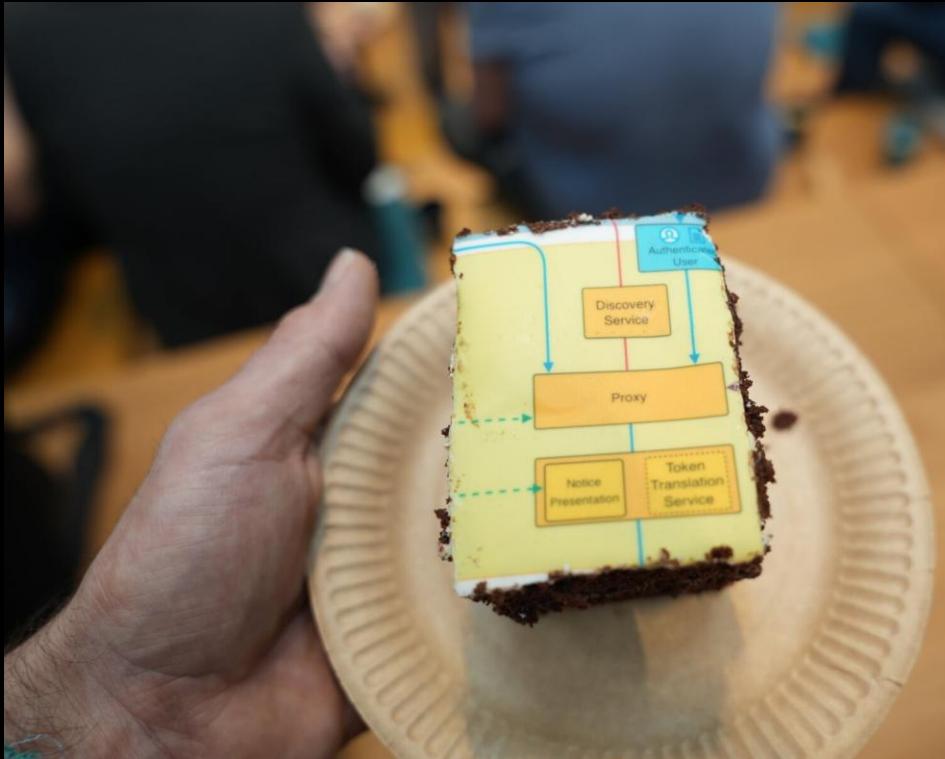
Photo credit: Marcus Hardt

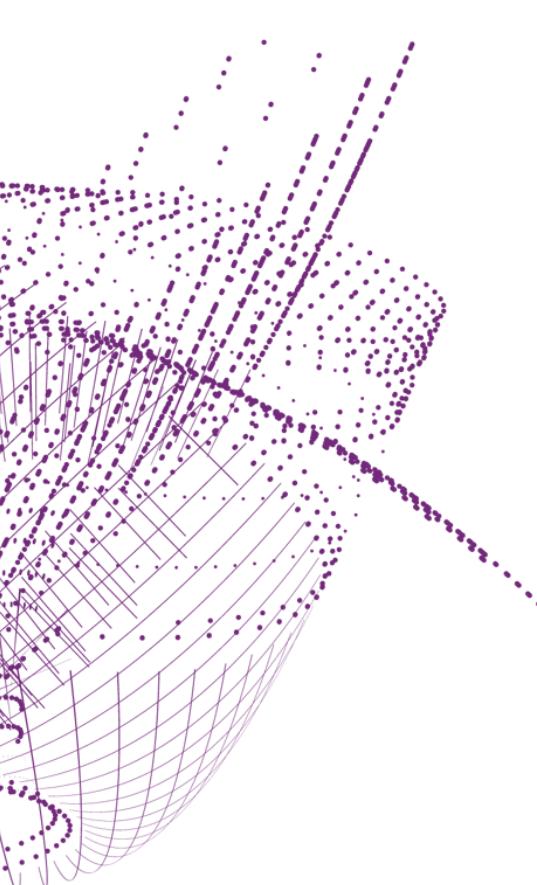
The AARC Blueprint – take a piece and feed collaboration!



Photo credit: Marcus Hardt

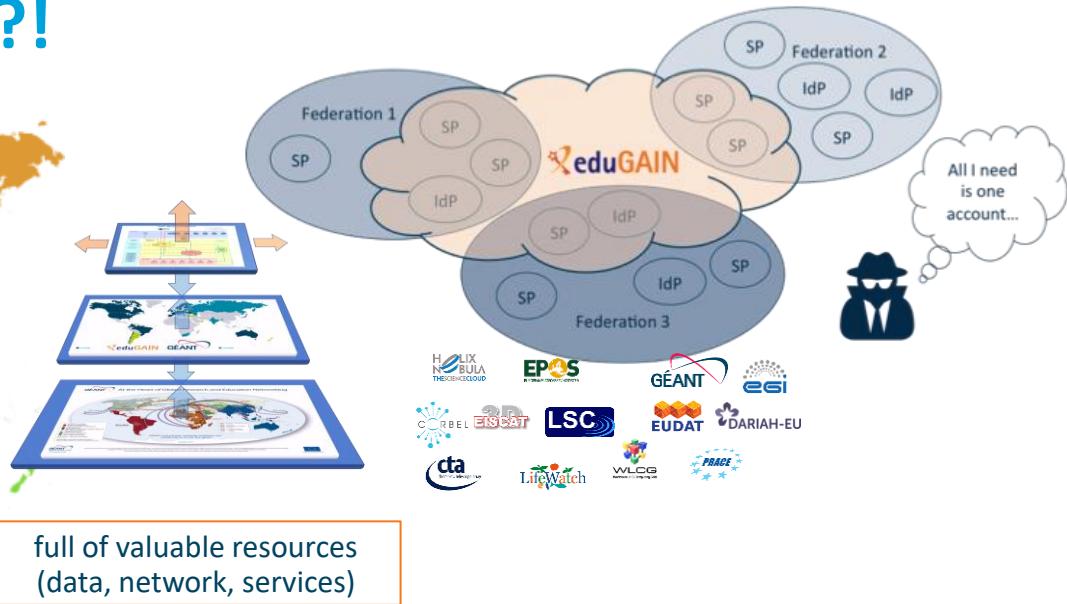
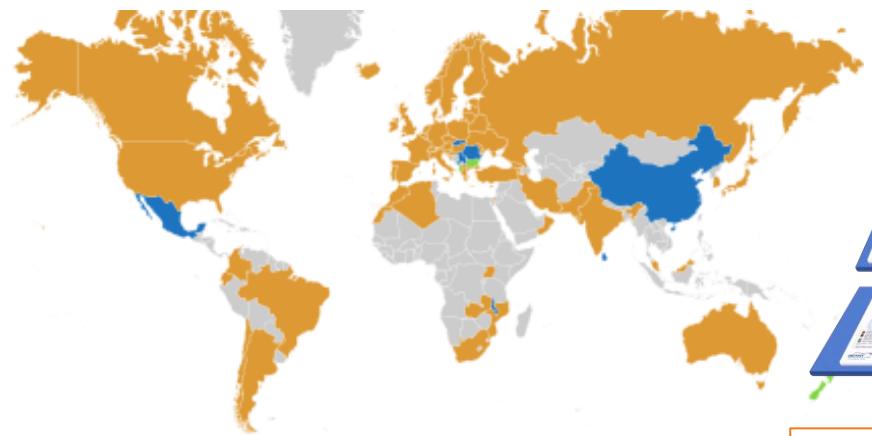
And Marcus Hardt then ate the proxy





Securing our federated world

Now *what* have we built?!



We have federation and single sign-on ...
... but can we share security information when needed?
... timely and confidentially, protecting everyone's reputation?

left: eduGAIN interederation extent in 2020; logos on the right from the European e-Infrastructures and ESFRIs; center graphic: AARC collaboration

‘Sirtfi’ – what makes federated security different?



SIRTFI

Security Incident Response Trust
Framework for Federated Identity

Organisations probably do ‘something reasonable’ for their own security ... but may not realise the implications for others

Sirtfi targets coordinated **response in a federated context**:

1. Enable **communication** and coordination in managing federated security incidents
2. Relevant **event data** is available to help collaborating incident responders.
3. **Security protections are applied** to federated transactions

Define capabilities for security incident response an IdP or SP
organisation can self-asserts in federation meta-data

<https://refeds.org/sirtfi>

The image shows a screenshot of the SIRTFI Version 2 document. The cover page includes the Refeds logo, document version (2.0, dated 28 July 2022, page 1/10), and the title 'A Security Incident Response Trust Framework for Federated Identity (Sirtfi) Version 2'. The table of contents on the right lists four main sections: Operational Security, Incident Response, Tracability, and User Rules & Conditions. A note at the bottom left of the document page states: 'This document is intended for use by the personnel responsible for operational security of federated entities such as Identity Providers, Service Providers and Attribute Authorities, and by Federation Operators who may facilitate its adoption by their member organisations.'

- [IR3] Notify security contacts of entities participating in Sirtfi when a security incident investigation suggests that those entities are involved in the incident. Notification should also follow the security procedures of any federations to which your organisation belongs.

Sirtfi – Security Incident Response Trust framework for Federated Identity



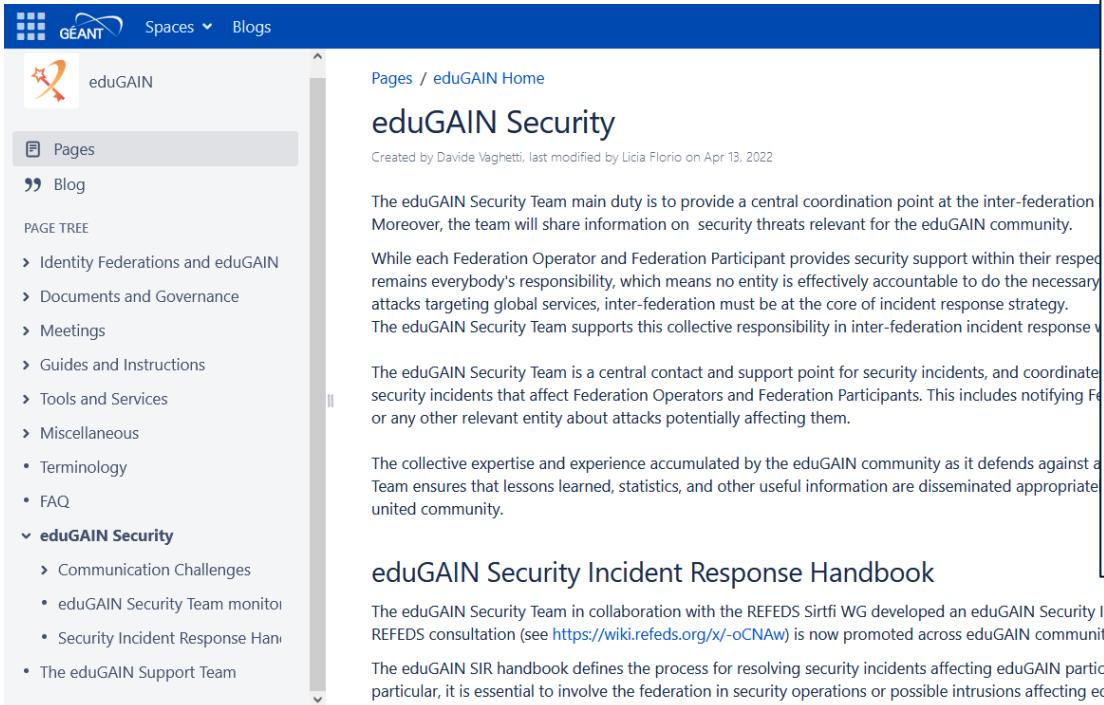
A means by which to enable a **coordinated response to a security incident in a federated context** that does not depend on a centralised authority or governance structure to assign roles and responsibilities for doing so.

Defines a set of capabilities and roles associated with security incident response that an IdP or SP **organisation self-asserts**. The Sirtfi trust framework posits that organisations asserting conformance with these will coordinate their response to security incidents.

Derived from the first four elements of the SCI Framework:

- **Operational Security:** patch and vulnerability management; IDS and threat mitigation; service ownership management; user suspension and termination; CSIRT capability
- **Incident Response:** CSIRT contact in meta-data; timely response; collaborate in IR; defined processes; privacy respect; TLP information sharing
- **Traceability:** timestamped accurate logs are available; log retention process in place
- **Participant Responsibilities:** users agree to an AUP; awareness and acceptance of the AUP

The eduGAIN Security Handbook



The screenshot shows the eduGAIN website interface. At the top, there is a navigation bar with the GÉANT logo, 'Spaces', and 'Blogs' options. Below the navigation bar, the 'eduGAIN' logo is displayed. The main content area has a left sidebar with links to 'Pages', 'Blog', 'PAGE TREE', and various sections under 'Identity Federations and eduGAIN', 'Documents and Governance', 'Meetings', 'Guides and Instructions', 'Tools and Services', 'Miscellaneous', 'Terminology', 'FAQ', and 'eduGAIN Security' (which is expanded to show 'Communication Challenges', 'eduGAIN Security Team monitoring', 'Security Incident Response Handbooks', and 'The eduGAIN Support Team'). The main content page is titled 'eduGAIN Security' and was last modified on April 13, 2022. It contains text about the eduGAIN Security Team's role in providing a central coordination point and sharing security threats. It also discusses the team's support for Federation Operators and Participants, and its role in incident response and community coordination. A section titled 'eduGAIN Security Incident Response Handbook' is also present, mentioning its development in collaboration with the REFEDS Sirtfi WG and its purpose in defining a process for resolving security incidents.

Pages / eduGAIN Home

eduGAIN Security

Created by Davide Vaghetti, last modified by Licia Florio on Apr 13, 2022

The eduGAIN Security Team main duty is to provide a central coordination point at the inter-federation. Moreover, the team will share information on security threats relevant for the eduGAIN community.

While each Federation Operator and Federation Participant provides security support within their respective communities, the security of the entire inter-federation remains everybody's responsibility, which means no entity is effectively accountable to do the necessary actions to defend against attacks targeting global services. inter-federation must be at the core of incident response strategy.

The eduGAIN Security Team supports this collective responsibility in inter-federation incident response via the following mechanisms:

- The eduGAIN Security Team is a central contact and support point for security incidents, and coordinates responses to security incidents that affect Federation Operators and Federation Participants. This includes notifying Federation Operators and Participants, and any other relevant entity about attacks potentially affecting them.
- The collective expertise and experience accumulated by the eduGAIN community as it defends against attacks is used to improve the security of the entire inter-federation. The eduGAIN Security Team ensures that lessons learned, statistics, and other useful information are disseminated appropriate to the entire inter-federation and the wider eduGAIN community.

eduGAIN Security Incident Response Handbook

The eduGAIN Security Team in collaboration with the REFEDS Sirtfi WG developed an eduGAIN Security Incident Response (SIR) Handbook, which after REFEDS consultation (see <https://wiki.refeds.org/x/oCNAw>) is now promoted across the eduGAIN community for adoption.

The eduGAIN SIR handbook defines the process for resolving security incidents affecting eduGAIN participants involving all key stakeholders. In particular, it is essential to involve the federation in security operations or possible intrusions affecting eduGAIN entities.

<https://edugain.org/edugain-security/references/> eduGAIN Security activities supported by the GN4-3 and GN5-1 Trust and Identity activities

eduGAIN Security Incident Response Handbook

Preface

Chapter 1. Understanding Your Role and Responsibilities

Introduction	2
Roles	2
Scope	3
Responsibilities	3
Federation Participants	4
Federation Operators	4
eduGAIN Security Team	4

Chapter 2. Security Incident Response Procedures

Federation Participants	5
Federation Operators	6
eduGAIN Security Team	7

Preface

As with products of any REFEDS Working Group, in this instance the SIRTFI Working Group,

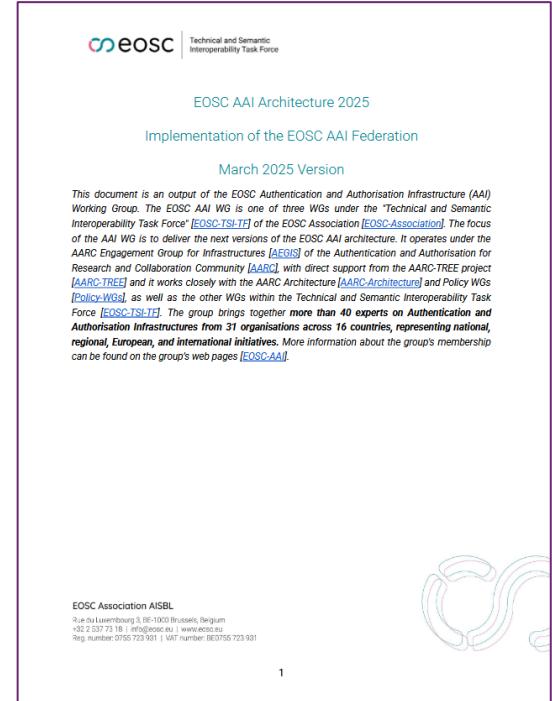
And security is a mandatory corner stone also of the EOSC

11. MUST support common security procedures

Security procedures define procedures and duties to allow an organised incident response. Distributed systems, in particular when spanning multiple organisational domains and countries, need a common approach for security related matters.

The following are well established in distributed infrastructures, and therefore mandatory for being supported:

- Security Incident Response Trust Framework for Federated Identity Sirtfi [\[REFEDS-SIRTIFI\]](#)
- Security Operational Baseline [\[AARC-G084\]](#) to enable secure infrastructure operation
- **Data Protection** for access to personal data: Compliance with the REFEDS Code of Conduct version 2 [\[REFEDS-DPCoCo\]](#) or other GDPR-aligned code of conduct.Collaboration Platform (Community AAI)

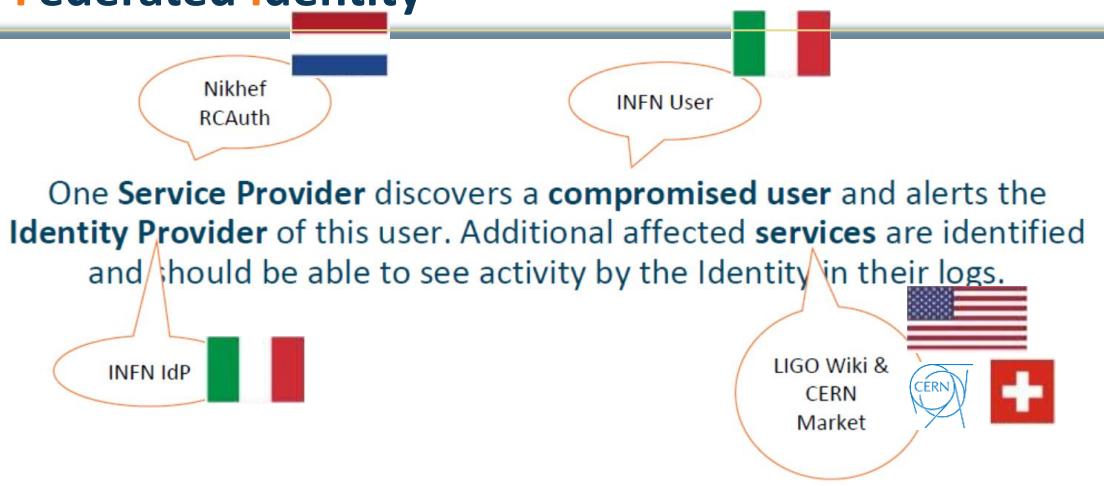


The image shows the EOSC AAI Architecture 2025 logo at the top left, featuring a blue hexagon with a white 'e' and the text 'eosc' and 'Technical and Semantic Interoperability Task Force'. The main document cover is a white page with a purple border. It features the EOSC logo, the title 'EOSC AAI Architecture 2025', the subtitle 'Implementation of the EOSC AAI Federation', and the date 'March 2025 Version'. Below this is a detailed description of the document's purpose and the EOSC AAI Working Group's role in delivering the next versions of the EOSC AAI architecture. The text also mentions the AARC Engagement Group for Infrastructures, the AARC-TREE project, and the AARC-Architectural and Policy WGs. It highlights the group's focus on authentication and authorisation infrastructures across 16 countries, representing national, regional, European, and international initiatives. The document is an output of the EOSC Association AISBL, located in Brussels, Belgium, with a registration number BE0755 723 931 and a VAT number BE0755 723 931. The cover also features a stylized blue and green 'C' shape graphic on the right side.

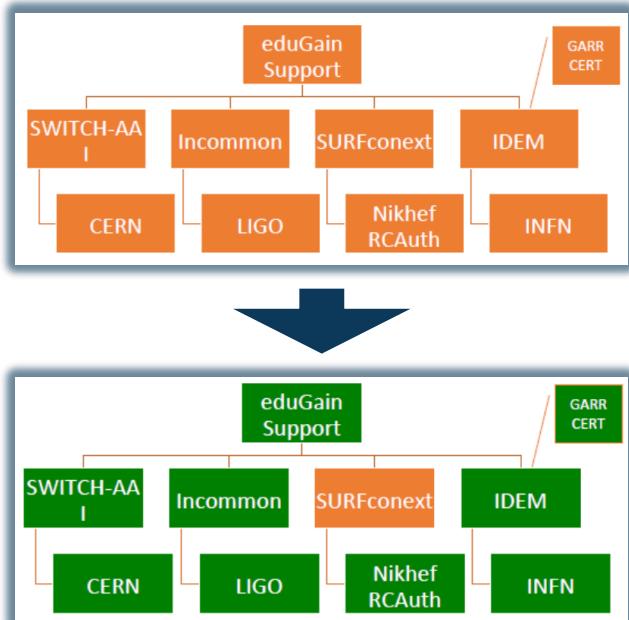
A federated community security challenge

Can we coordinate our collective R&E response?
'challenges' based on the *Sirtfi* contact model

Security Incident Response Trust Framework for Federated Identity



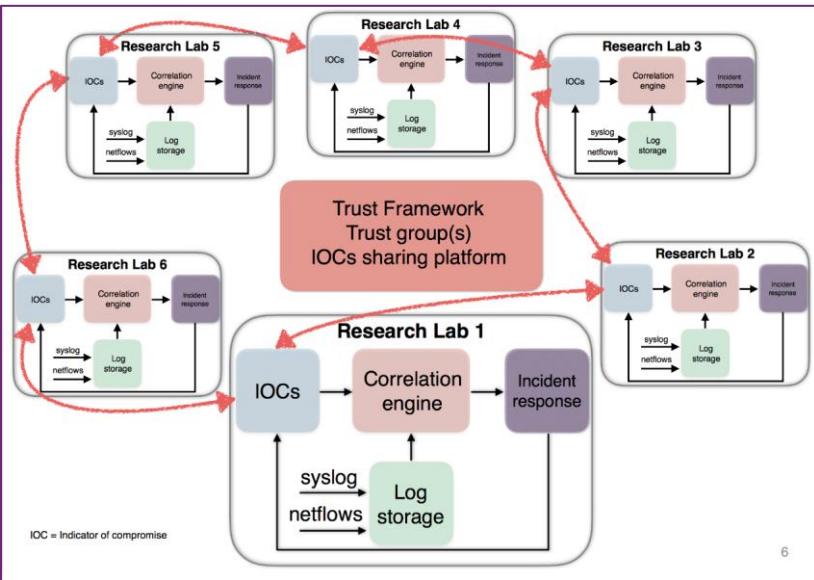
Report-outs see <https://wiki.geant.org/display/AARC/Sirtfi+Communications+Challenges%2C+AARC2-TNA3.1>



parties involved in response challenge

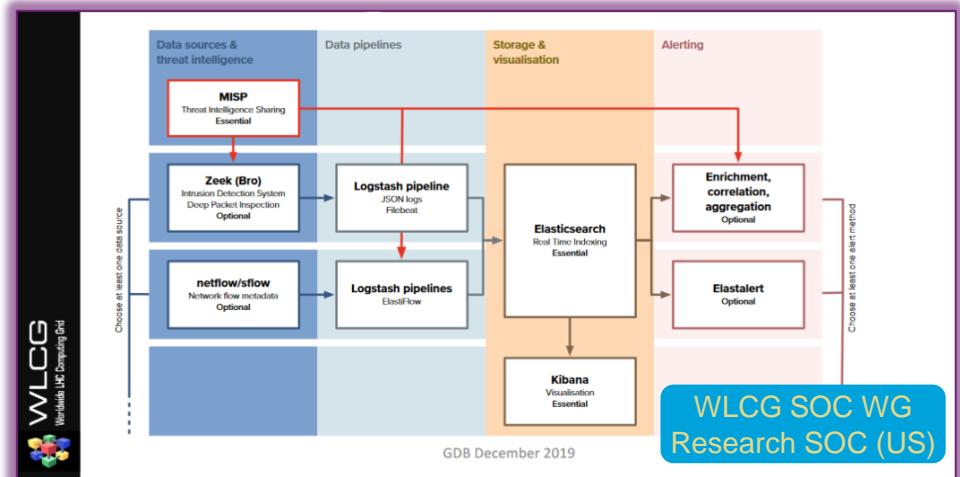
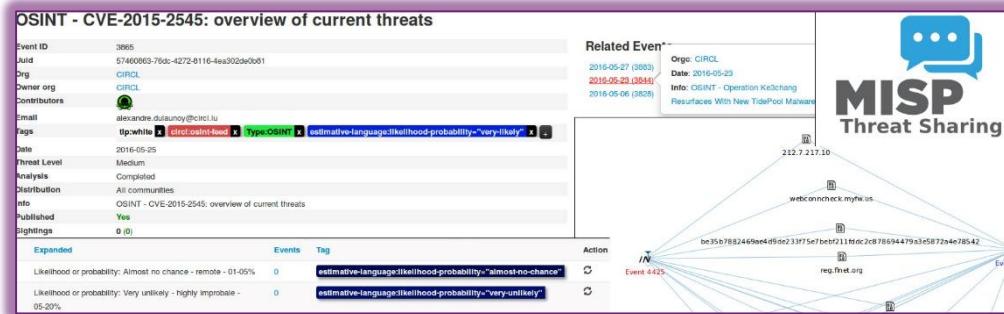
Sharing threat intel – working with our community

MITRE | ATT&CK®

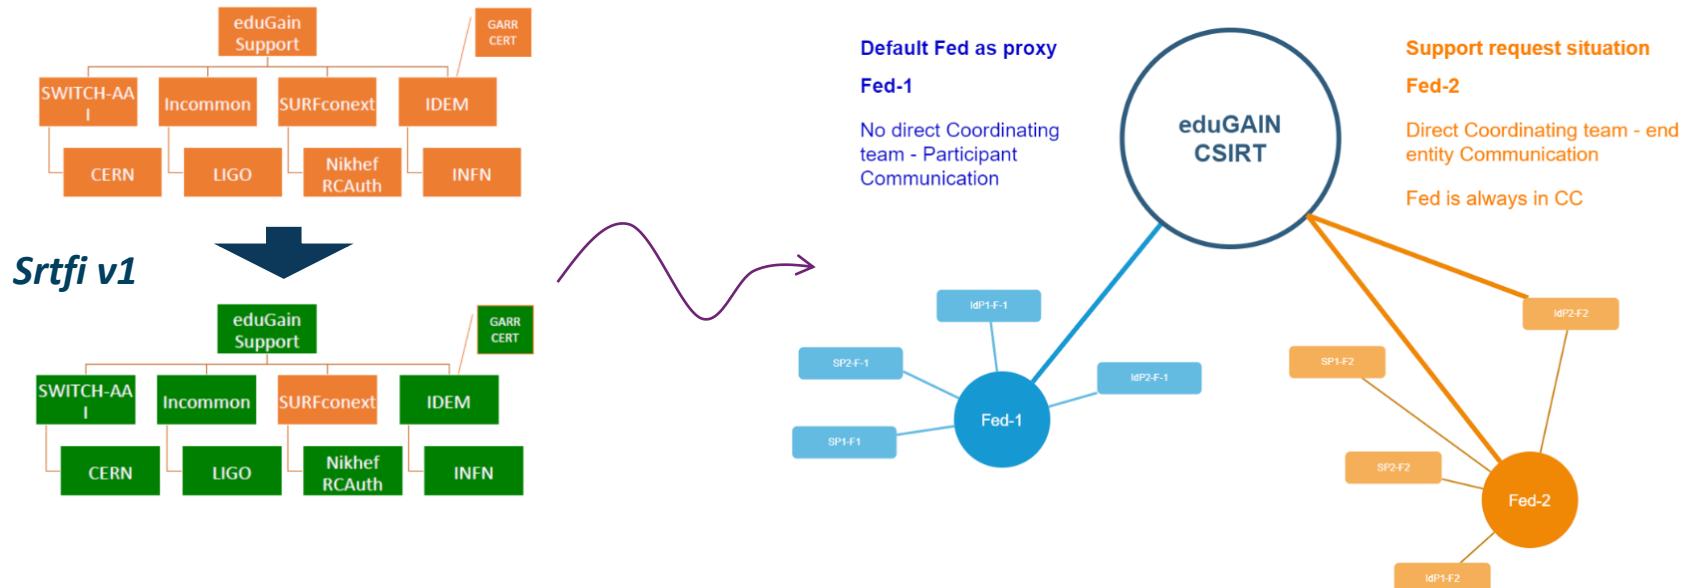


The logo for the American Association for Respiratory Care (AARC) is located in the bottom right corner. It consists of the letters "AARC" in a bold, black, sans-serif font, with a stylized orange and blue swoosh graphic to the left of the "A".

AARC I-051 Guide to federated incident response
<https://aarc-community.org/guidelines/aarc-i051/>



Response across IdP-SP Proxies: the limits of Sirtfi version 1

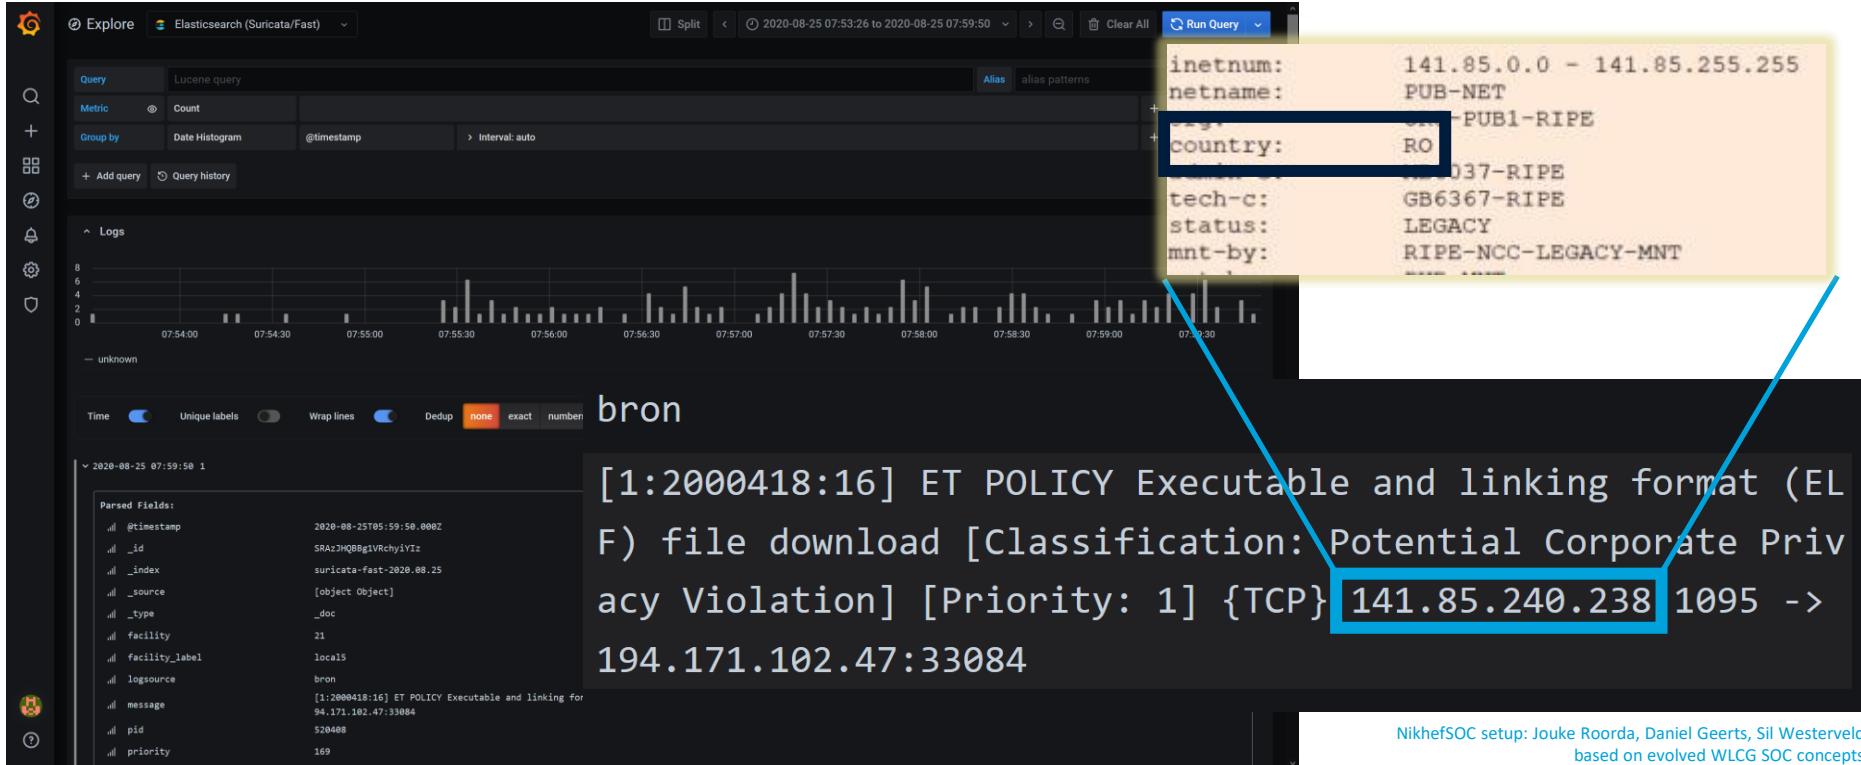


joint work with GN5 EnCo
and eduGAIN CSIRT



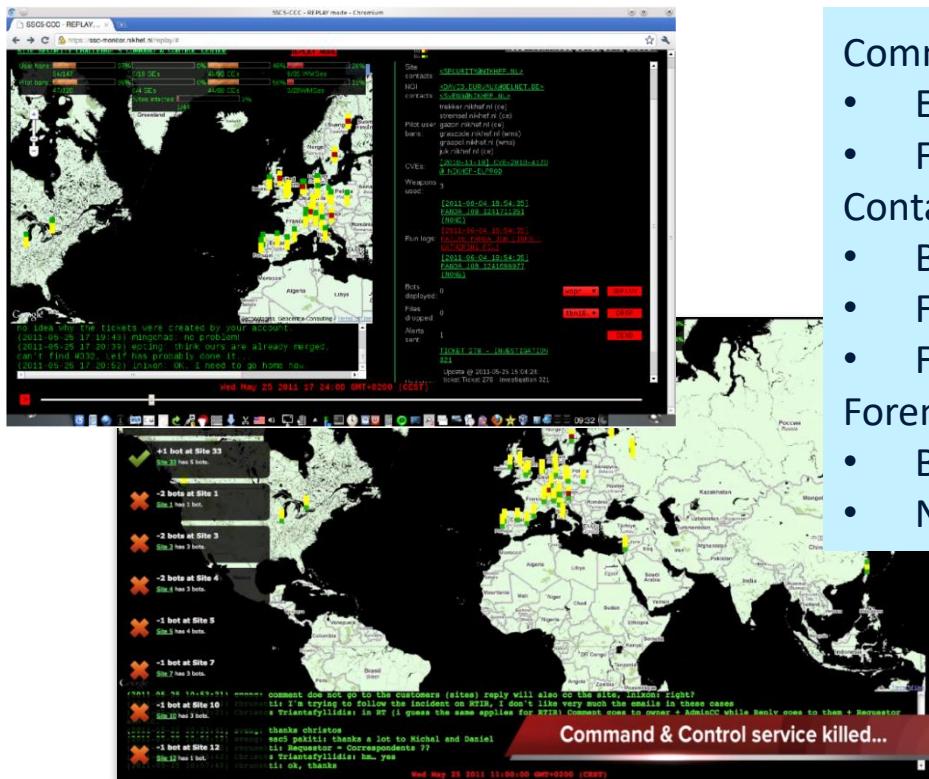
A single site sees only so much ...

many ‘false warnings’ when industry-standard (e.g. Suricata) rules are used. You need R&E specific ones!



NikhefSOC setup: Jouke Roorda, Daniel Geerts, Sil Westerveld
based on evolved WLCG SOC concepts

A question of *when*, not *if* – hence we run security challenges



Communication:

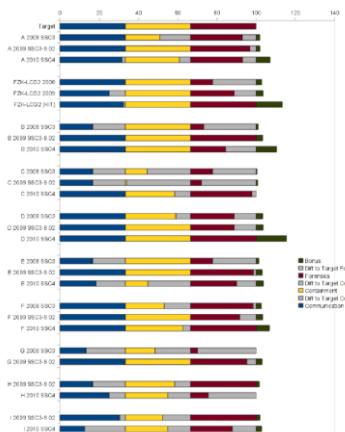
- Endpoints valid?
- Form/Content OK ?

Containment

- Ban "malicious" users
- Find/Stop malicious processes
- Find submission IP

Forensics

- Basic Forensics on binary
- Network traffic



Nikhef

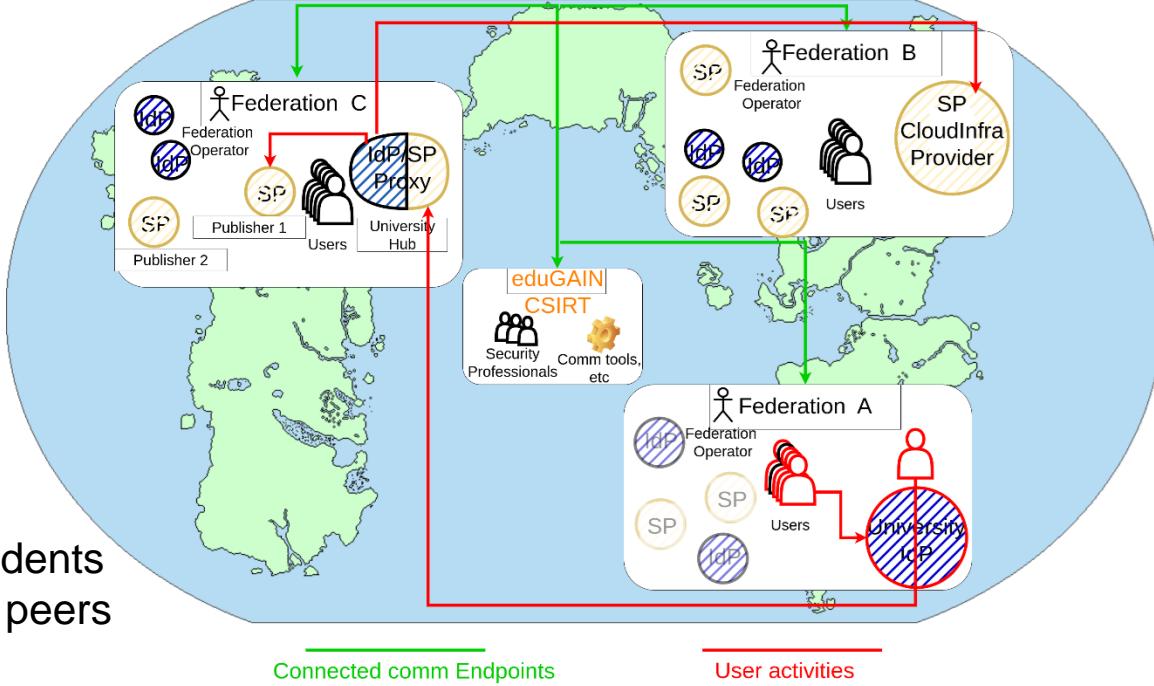
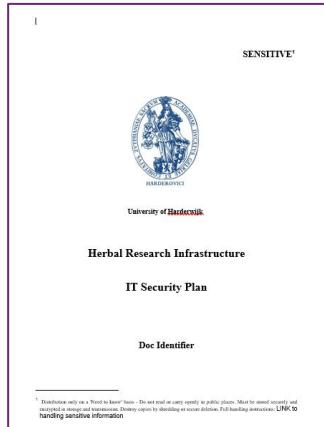
Nikhef CSIRT Traceability Challenge

Introduction

Deze Traceability Challenge bestaat uit drie onderdelen, in (naar verwachting) oplopende moeilijkheidsgraad. Iedere challenge begint met een externe 'trigger' – aan het eind van dit document staan de hints en de opdracht (of in ieder geval: de 'gewenste') oplossing.

Vast nietzien

Federation security table-top exercises

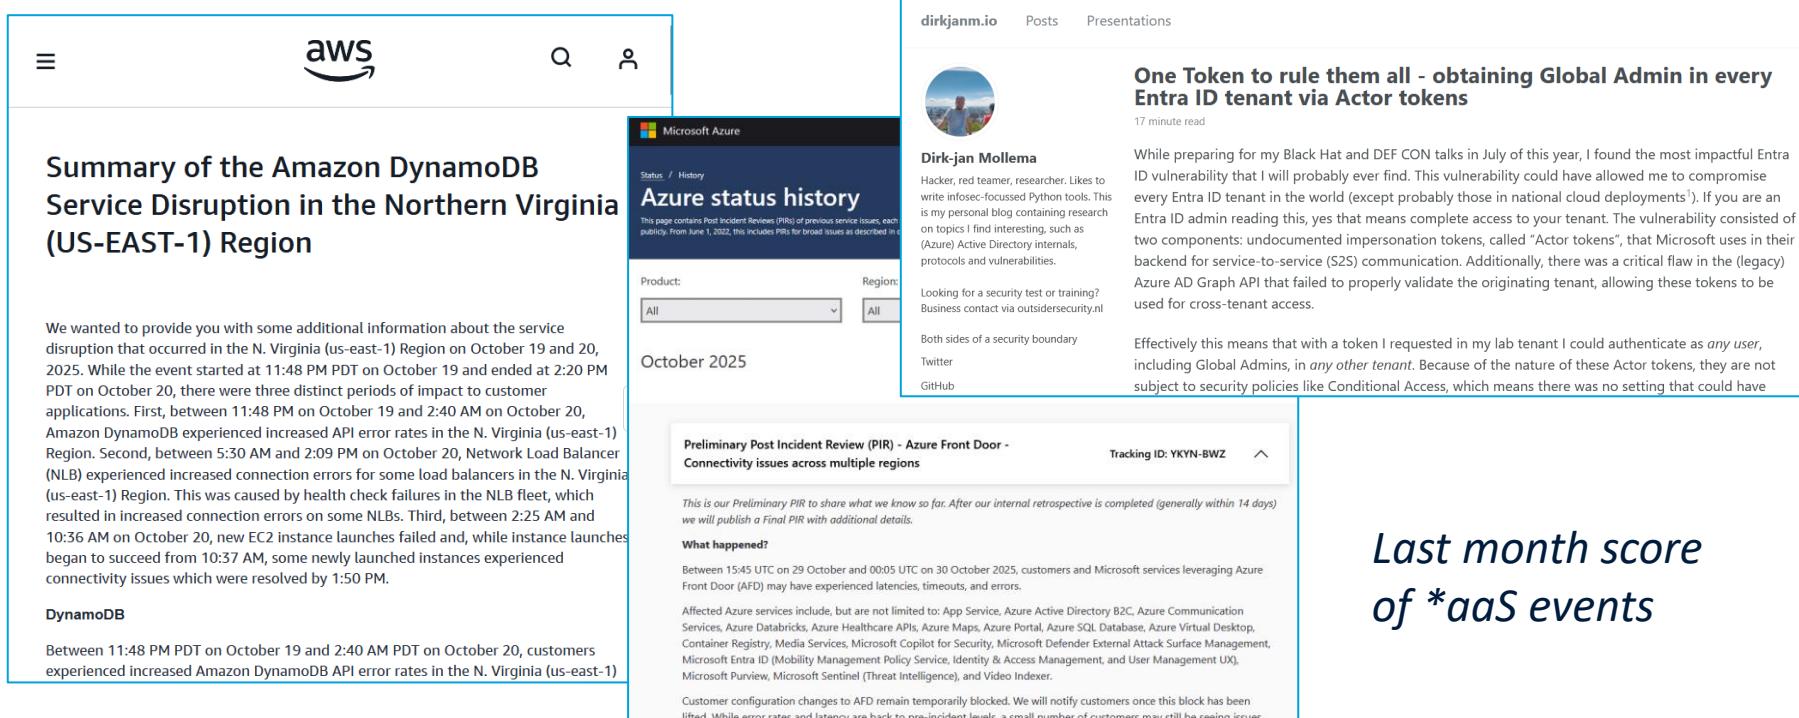


What your role play brings you 😊

- real time pressure to contain incidents
- true gratitude for protecting your peers
- collective recovery
- exploring some gruelling conflicts of interest!

eduGAIN TTX – role play scenario from the ISGC Security Workshop 2024, 2025

On leaky abstractions and circular dependencies



Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-EAST-1) Region

We wanted to provide you with some additional information about the service disruption that occurred in the N. Virginia (us-east-1) Region on October 19 and 20, 2025. While the event started at 11:48 PM PDT on October 19 and ended at 2:20 PM PDT on October 20, there were three distinct periods of impact to customer applications. First, between 11:48 PM on October 19 and 2:40 AM on October 20, Amazon DynamoDB experienced increased API error rates in the N. Virginia (us-east-1) Region. Second, between 5:30 AM and 2:09 PM on October 20, Network Load Balancer (NLB) experienced increased connection errors for some load balancers in the N. Virginia (us-east-1) Region. This was caused by health check failures in the NLB fleet, which resulted in increased connection errors on some NLBs. Third, between 2:25 AM and 10:36 AM on October 20, new EC2 instance launches failed and, while instance launches began to succeed from 10:37 AM, some newly launched instances experienced connectivity issues which were resolved by 1:50 PM.

DynamoDB

Between 11:48 PM PDT on October 19 and 2:40 AM PDT on October 20, customers experienced increased Amazon DynamoDB API error rates in the N. Virginia (us-east-1) Region.

Azure status history

Status / History

October 2025

Preliminary Post Incident Review (PIR) - Azure Front Door - Connectivity issues across multiple regions

Tracking ID: YKYN-BWZ

What happened?

Between 15:45 UTC on 29 October and 00:05 UTC on 30 October 2025, customers and Microsoft services leveraging Azure Front Door (AFD) may have experienced latencies, timeouts, and errors.

Affected Azure services include, but are not limited to: App Service, Azure Active Directory B2C, Azure Communication Services, Azure Databricks, Azure Healthcare APIs, Azure Maps, Azure Portal, Azure SQL Database, Azure Virtual Desktop, Container Registry, Media Services, Microsoft Copilot for Security, Microsoft Defender External Attack Surface Management, Microsoft Entra ID (Mobility Management Policy Service, Identity & Access Management, and User Management UX), Microsoft Purview, Microsoft Sentinel (Threat Intelligence), and Video Indexer.

Customer configuration changes to AFD remain temporarily blocked. We will notify customers once this block has been lifted. While error rates and latency are back to pre-incident levels, a small number of customers may still be seeing issues

One Token to rule them all - obtaining Global Admin in every Entra ID tenant via Actor tokens

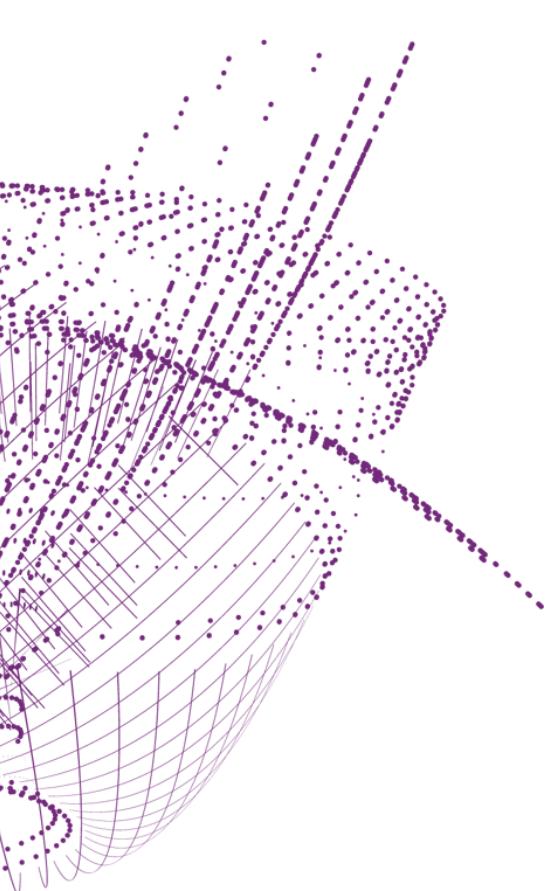
17 minute read

While preparing for my Black Hat and DEF CON talks in July of this year, I found the most impactful Entra ID vulnerability that I will probably ever find. This vulnerability could have allowed me to compromise every Entra ID tenant in the world (except probably those in national cloud deployments'). If you are an Entra ID admin reading this, yes that means complete access to your tenant. The vulnerability consisted of two components: undocumented impersonation tokens, called "Actor tokens", that Microsoft uses in their backend for service-to-service (S2S) communication. Additionally, there was a critical flaw in the (legacy) Azure AD Graph API that failed to properly validate the originating tenant, allowing these tokens to be used for cross-tenant access.

Effectively this means that with a token I requested in my lab tenant I could authenticate as *any user*, including Global Admins, in *any other tenant*. Because of the nature of these Actor tokens, they are not subject to security policies like Conditional Access, which means there was no setting that could have

*Last month score of *aaS events*

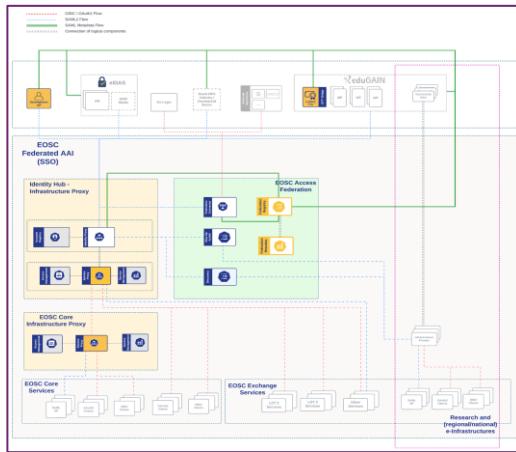
<https://aws.amazon.com/message/101925/>, <https://azure.status.microsoft.com/en-us/status/history/>, <https://dirkjanm.io/obtaining-global-admin-in-every-entra-id-tenant-with-actor-tokens/>



Collaboration for research

Enabling research: using the ‘EOSC’ with federated login

AARC compliant federation of ‘national’ and ‘thematic’ nodes in the European Open Science Cloud
linked with other ‘data spaces’ and infrastructures



 eosc | Federation



The organisations invited to join the March 2025 kick-off workshop for the build-up phase of the EOSC Federation. All of the organisations are among the membership of the EOSC Association.

<https://eosc.eu/eosc-about/building-the-eosc-federation/contributing-to-the-build-up-phase-of-the-eosc-federation/>; See also <https://wiki.geant.org/display/AARC/EOSC+AAI>

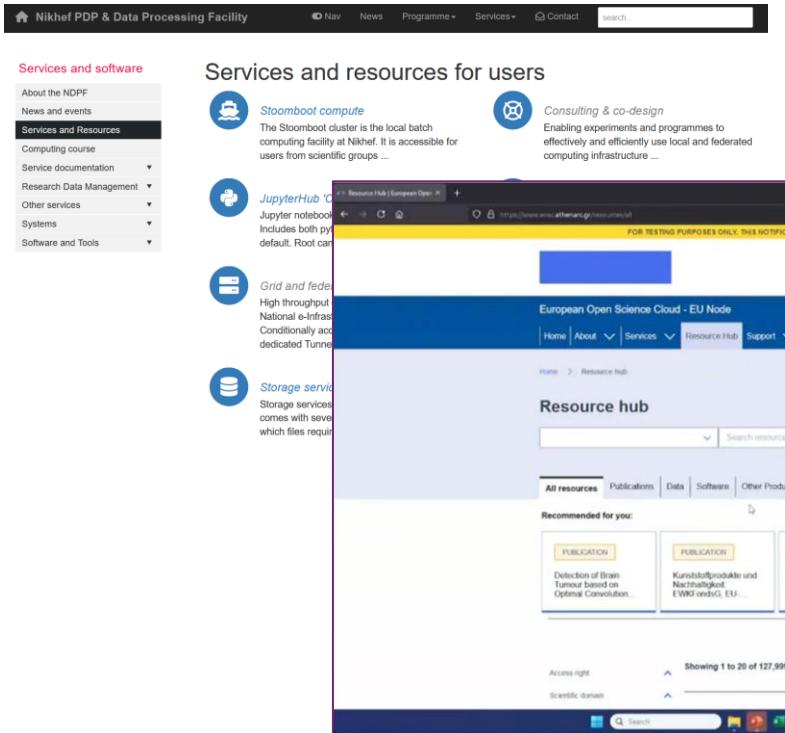
Where do researchers find services & collaboration



represented by logos: some of the (AARC BPA) Research Communities (top) providing federated access using the AAI proxy architecture.

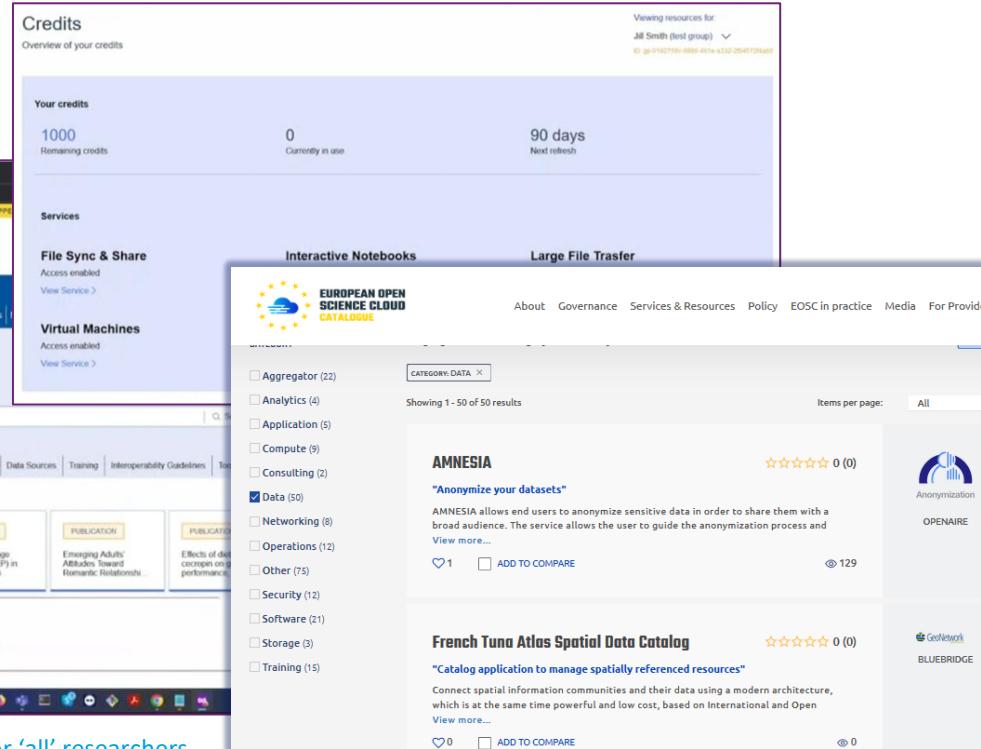
At the ~ bottom: (global) e-Infrastructures, which all use the AARC BPA collaborative model

Service portfolios – what do you offer, and to whom



The screenshot shows the Nikhef PDP & Data Processing Facility website. The left sidebar is titled 'Services and software' and includes links for 'About the NPDF', 'News and events', 'Services and Resources' (which is highlighted in a dark box), 'Computing course', 'Service documentation', 'Research Data Management', 'Other services', 'Systems', and 'Software and Tools'. The main content area is titled 'Services and resources for users' and lists four services with icons: 'Stoomboot compute' (batch computing facility), 'Consulting & co-design' (enabling experiments and programmes), 'JupyterHub 'C'' (Jupyter notebook), and 'Grid and fede' (High throughput National e-Infra). Below these is a screenshot of a browser showing the 'European Open Science Cloud - EU Node' Resource Hub. The hub displays a search bar, a 'Resource hub' menu, and a grid of recommended resources. At the bottom, it shows 'Showing 1 to 20 of 127,999,850 resources'.

Catalogues from Nikhef, European Open Science Cloud EU Node (free VMs for 'all' researchers, subject to <https://open-science-cloud.ec.europa.eu/system/files?file=2024-10/EOSC-EU-Node-User-Access-Policy-v1.0.pdf>)



The screenshot shows the European Open Science Cloud Catalogue. At the top, it displays 'Credits' (1000 remaining, 0 currently in use, 90 days until refresh). Below this are sections for 'Services' (File Sync & Share, Virtual Machines), 'Interactive Notebooks', and 'Large File Transfer'. The main area shows a list of services categorized under 'EUROPEAN OPEN SCIENCE CLOUD CATALOGUE'. Categories include Aggregator (22), Analytics (4), Application (5), Compute (9), Consulting (2), Data (50) (which is checked), Networking (8), Operations (12), Other (75), Security (12), Software (21), Storage (3), and Training (15). Below this is a section for 'AMNESIA' (Anonymize your datasets) with a star rating of 0 (0) and a link to 'View more...'. Another section for 'French Tuna Atlas Spatial Data Catalog' is also shown with a star rating of 0 (0) and a link to 'View more...'. The bottom right corner shows the 'GeoNetwork BLUEBRIDGE' logo.

‘Services await us’ in global research & e-infrastructures

both in *thematic* and in
horizontal e-Infrastructures

The screenshot shows the WorkflowHub interface. On the left, a sidebar lists 'Workflows' with filters for 'Created At' (Any time), 'Updated At' (Any time), and 'Tool' (MultiQC, SAMtools, fastp, BWA, Bowtie 2, Cutadapt). The main area shows a list of 350 workflows, with one highlighted: 'Galaxy Protein conformational ensembles generation' (Status: Ready). On the right, a login screen for 'LIFE SCIENCE RI' is displayed, showing 'Login with' options for 'maastricht' and 'Search the extended list of Identity Providers' (Maastricht University, Academisch Ziekenhuis Maastricht, maastricht nl, Maastricht School of Management, maastricht nl). Below the login is a note: 'Hiding complexity and transparent access to' and a link to 'Further info: <https://wiki.escape/>'.

ELIXIR RI and Life Sciences AAI (left),
ESCAPE Data Lake by Ricardo Di Maria (CERN)
CS3MESH4EOSC – Science Mesh and Services
<https://cs3mesh4eosc.eu/science-mesh>

The diagram illustrates the ESCAPE Data Lake architecture, showing the Orchestrator, Rucio Server, and various middleware components (GFTS, Storage, Caching Solution, and XRootD). It also shows the Science Mesh architecture, where users access data via SURF DRIVE, Cubbit, cesnet, CERNBox, and ailleron, which are connected to the Science Mesh Data Services. The Science Mesh Data Services are shown as a central hub for Data Science Environments, including JupyterLab, QuantStack Voilà, CodiMD, and VyourTechnology.

how to leverage all this effectively and achieve what we want?
Given our strategy strives for an attractive research climate
“Met hoogwaardige onderzoeksfaciliteiten stellen we hen in staat om excellent onderzoek te doen” – which includes ICT

Collaborative services are distributed and federated

Collaborative services are spread across the research community

- logbooks with federated login from LIGO and IGWN for the ET pathfinder
- analysis notebooks and control software in open to the collaboration via eduGAIN
- our aforementioned RCauth.eu

need mix of local expertise and resources, national systems, research infra services, and European (global) resources

'every partner contributes to the trip to Stockholm'

The screenshot shows a web browser window with the Nikhef logo and the text "National Institute for Subatomic Physics". A message at the top says "You have previously chosen to authenticate at Maastricht University". Below this is a grid of buttons for various institutions, with "Preferred" at the top left. The grid includes: BE, BG, BR, BY, CA, CH, CL, CN, CO, CY, CZ, DE, DK, DZ, EC, EE, ES, FI, FO, FR, GE, GI, GL, GR, HK, HR, HU, IE, IL, IN, IR, IS, IT, JP, KE, KG, KR, LB, LT, LK, LT, LU, LV, MA, MD, MK, MO, MT, MX, MY, NC, NG, NL, NO, NZ, OM, PF, PK, PL, PT, RO, RS, SA, SE, SG, SI, SK, SO, TH, TR, UA, UG, UK, US, ZA, ZM, Experimental, and Outsourced. Below the grid is an "Incremental search..." input field. A list of projects follows:

- T Kevin Heijhoff / tpx4_configs
- P Jordy Degens / Plotting
- D XAMS / DAOC
- L Uwe Kraemer / labEnvironmentLogger

Each project entry includes a star rating, a file count, a download count, and an "Updated" timestamp. At the bottom right, the text "INFRASTRUCTURE: FOR THE SMALL AND THE LARGE" is visible.

Infrastructure for research is an ecosystem: hardware, software, services, and ... people

SURF

RUCIO

MANAGEMENT

Sunday, Oct 6, 2019
• Bytes: 463 383 065 878 740 200

2010 2012 2014 2016 2018

SURF

High-performance data processing

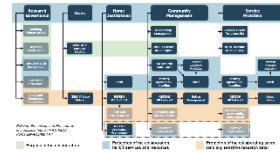
Do you want to process and store large volumes of data? Our team of experts can support you in using our high-throughput data processing systems and storage solutions.

Questions? Contact info@surfara.nl

For processing large, structured data

Our high-throughput data processing systems support structured data, such as event records, log files, and databases, and unstructured data, such as images, videos, and documents.

Images: ATLAS Rucio volume, (from rucio.cern.ch); optical network: NDPF 'deel'; User meeting Stoomboot Office Hours (both Nikhef); Snellius opening visit; HPDC service page (both SURF)



... but what about that first pillar
of the Policy Development Kit?

Collaboration & governance

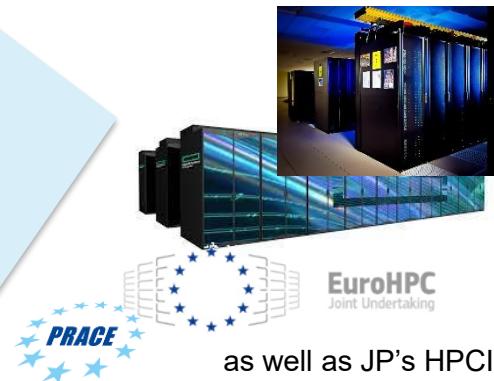
Make and treat computing as the research instrument it is today – institutionally and globally



Institutional:
Nikhef “Stoomboot”
Analysis Facility



National Infrastructure
SURF Snellius HPC



as well as JP's HPCI,
US's AccessCI, &c of course!

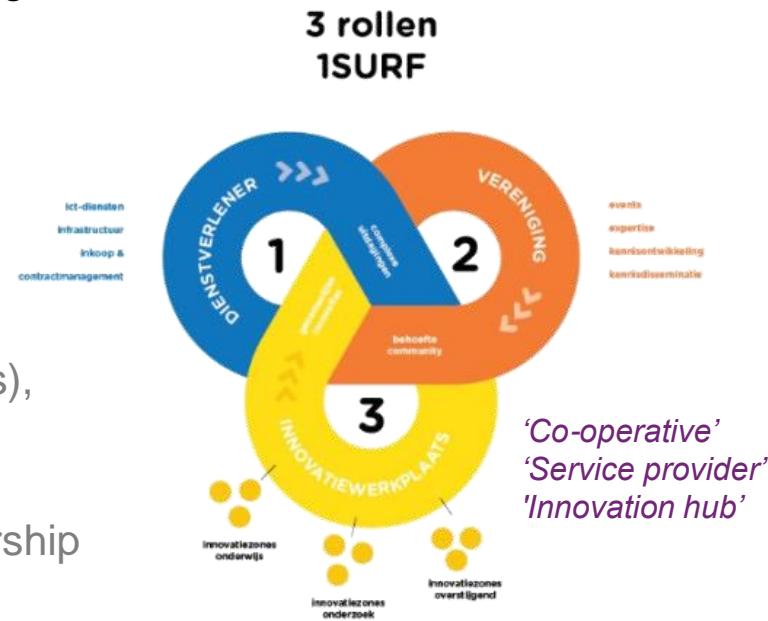
There are today as much part of science as detectors are to physics and: users should move seamlessly between tiers

Photos: Nikhef NDPF, DelftBlue/TUDelft, SURF Data Repository, Snellius, SURF @ DigitalRealty; EuroHPC images: EuroHPC, LUMI Consortium, Jules Verne consortium

SURF, our 'collaborative organisation for IT in edu and research'

Besides 'commodity' services (network, software licensing, joint procurement, &c) as a service provider, SURF adds

- **Research IT facilities**
HPC 'Snellius', HPC Cloud, Grid, Data Archive, collaborative analysis platform, ResearchCloud, ...
- **research support and open science**
national coordination research data management (LCRDM), digital competence centre initiatives (DCCs), European Open Science Cloud, expertise networks
- **Federated access and collaboration**
SURFconext, projects, international liaison & membership
- **Innovation of the IT knowledge basis**
at SURF and at co-creating member organisations



Integral Approach to ICT Digital Competences for Research

- need for a federated networked scheme for data, computing (and expertise) remains as relevant today as it was in 2017
- ‘local’ digital competence centres in their role as *“node in a federated network for data, computing en expertise”* did not get attention for *infrastructure* that was intended
- expertise bundling and development of “Tier-2” facilities in national landscape set as institutional responsibility, (‘strengthening research support’) with some central funding
- but using national funding also means: be open to national collaboration, and ensure facilities (expertise, but also datasets, computing, storage, networks) are actually accessible in a FAIR and federated way, open to researchers from outside – based on e.g. federated SRAM, MyAccessID, or IGTF



Integrale aanpak voor digitalisering
in de wetenschap

Uitvoeringsplan investeringen
digitale onderzoeksinfrastructuur



A comprehensive approach to research digitalization

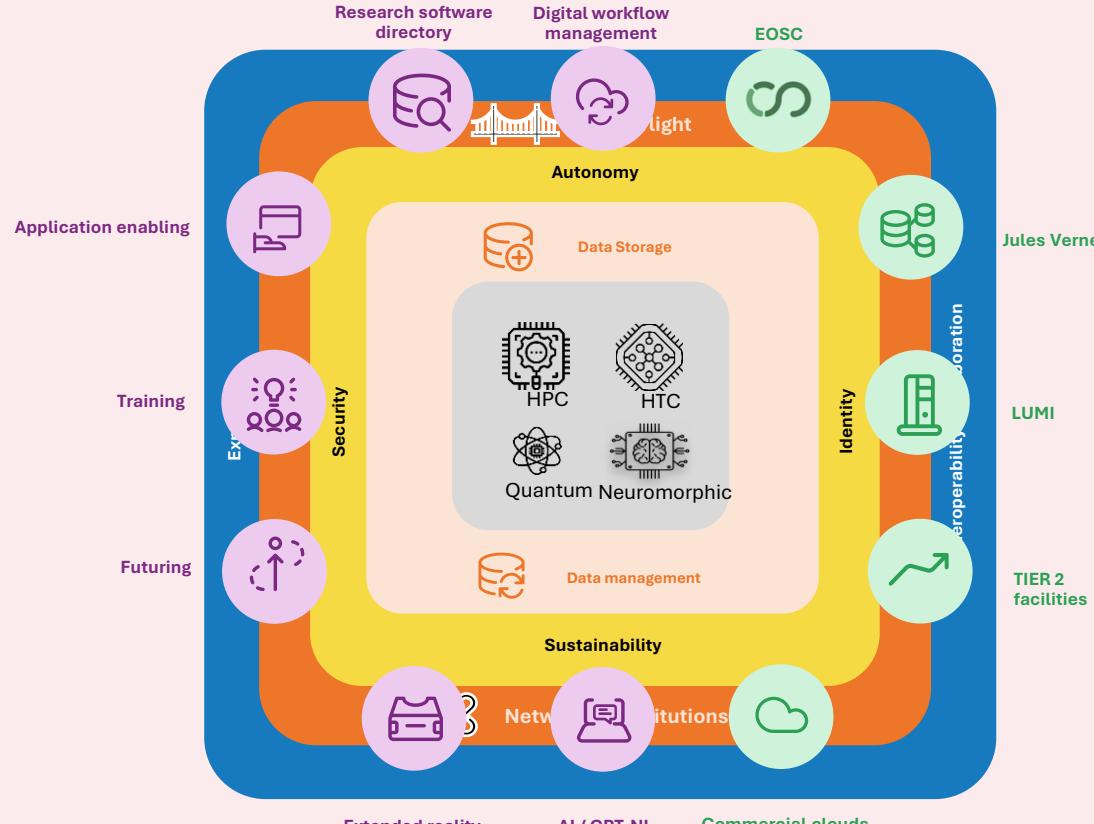
Specific execution plans at SURF for research support,
funded by the NWO 'Apers' means

- A. Rekenfaciliteiten: Aanschaf **nieuwe supercomputer**
- B. Rekenfaciliteiten: Toekomstige **investeringen in HPC**
& jaarlijkse gebruikersbijdrage reserveringen HPC
- C. Rekenfaciliteiten: Investeringen in **overige rekenfaciliteiten en opslag hardware**
- D. Rekenfaciliteiten: **Vernieuwing van de kennisbasis**
- E. Rekenfaciliteiten: **Expertise en ondersteuning** van rekenfaciliteiten en datacentra

- F. Digitalisering: Impulsfinanciering **lokale DCC's**
- G. Digitalisering: Stimulering **thematische DCC's**
- H. Digitalisering: Ondersteuning DCC's door SURF
- I. Digitalisering: Investeringen in **eScience**



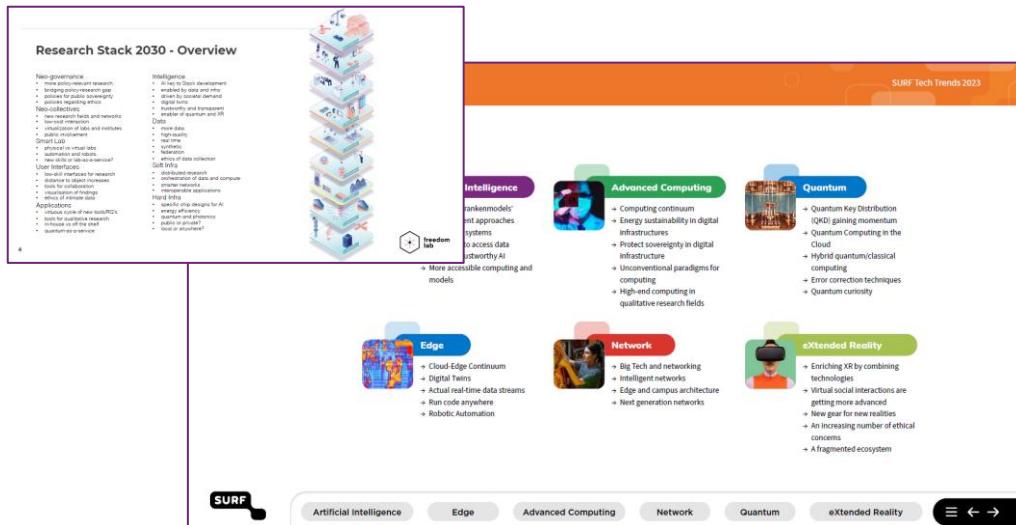
SURF Research Infrastructure Ecosystem



Acknowledgements: SURF, Jet de Ranitz, Magchiel Bijsterbosch

Joint innovation and the Dutch HPC Tier2 landscape

Representative expert group on knowledge basis innovation drives joint innovation with institutes on tech trends:
AI, Advanced Computing, Quantum, Edge, Network, XR



Integrale aanpak voor digitalisering in de wetenschap
Uitvoeringsplan investeringen digitale onderzoeksinfrastructuur



Strengthening Tier-2 infrastructure

- support interaction between SURF and institutes & universities
- workshops on Federated HPC & Data
- strengthen organisational Research IT support

Federating hitherto independent centres is non-trivial

Even with national funding available, getting a coherent national plan for Tier-2 facilities proved challenging

- Local build-up of computing expertise
- Technical alignment (ESSII/CVMFS, SRAM, OpenOnDemand, ssh, ...)
- Scaling: how to reach researchers?
- Divergent business models between institutions *and* between institutions and national/EU resources
- ‘Competition’ from ‘free’ national (NWO) compute grants for ‘pilot/small requests’

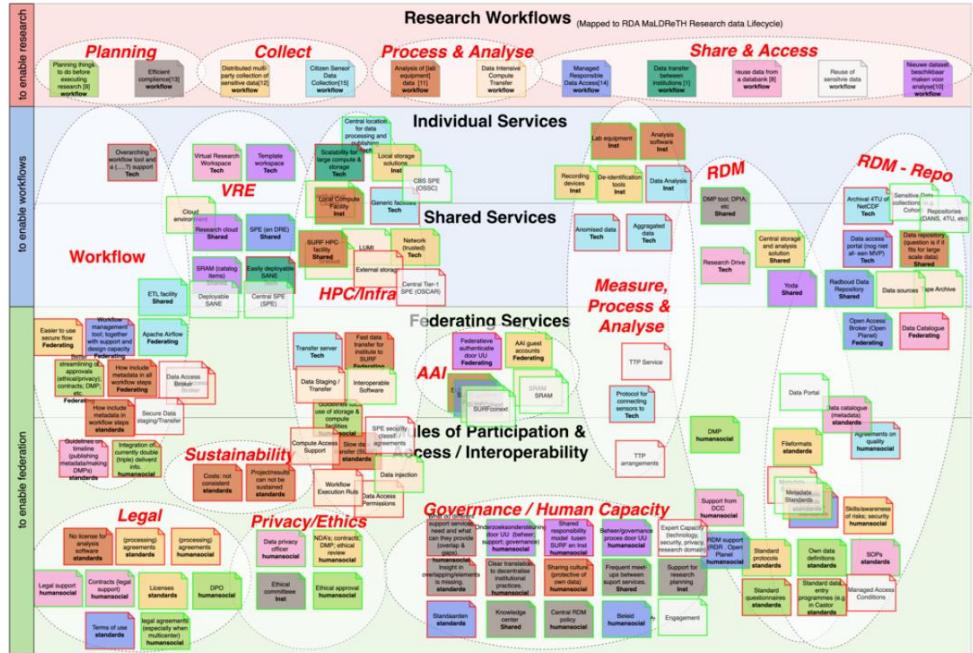


Figure 6: Output of workshops, for illustration purposes

SURF and NWO DCC-T2 proposal 1.0

Limitations to the membership model ?

Coherence between funding, research mission, and *representation* is challenging

- *Representation in the co-operative* is from mostly enterprise-IT focussed organisational units – that then drive key decision processes on the whole co-operative
- *Influence* on services by researchers, research funding, and funders is hard to anchor in the governance (no representation = limited influence)
- *Divergent directions* even within the co-operative
example: innovation action on digital sovereignty has no effect on enterprise-member-enforced discontinuation of existing autonomous services – regardless of strategy and impact 😔

Collaborating on ICT infrastructure collaboration

Innovatiezone: *Gemeenschappelijke Digitale Soevereiniteit*

- Bewustwording en kennis over alternatieven, creëren van draagvlak en bereidheid om dit serieus te ontwikkelen;
- Sturing op aanpassing van de Sourcing strategieën (van alle leden), acceptatie van Open Source;
- Investering in AI op basis van alternatieve ontwikkelingen;
- Mandaat om hier binnen SURF al pilots op te ondernemen; en
- Breed communiceren over beschikbare alternatieven en geleerde lessen.

Innovatiezone **Gemeenschappelijke Digitale Soevereiniteit**

Naar meer regie over de digitale omgeving op basis van publieke waarden

Bijeenkomst (virtueel): 13 september 2024, 12:30 - 14:00

SURF

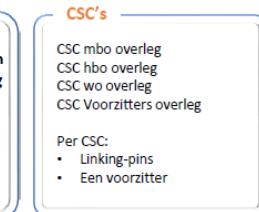
Agendapunt 1.

Governance van de Innovatiezone

De dagelijkse operationele aansturing als ook de facilitering van de innovatiezone zal door SURF worden verzorgd. Net als andere zones is iedere innovatiezone een complex vraagstuk waarvoor we samenwerken aan een oplossing. Dat doen we als SURF-organisatie met leden én als leden onderling.

Voor de invulling van de governance van de innovatiezone zijn er meer opties mogelijk en liggen er kansen om met behulp van bestaande gremia hier een passende vorm bij te vinden.

Er zijn meerdere bestaande gremia die een rol zouden kunnen spelen als het gaat om aansturing, mandaat, voortgang, uitdragen of advies, denk hierbij aan:



Even in NL - traditionally a big-tech-lax country ...

‘Digital autonomy is the starting point for government.

We choose a European digital infrastructure [...]

so that more Dutch and European SMEs can participate.

[...] based on security-by-design, zero-trust, sovereignty, open source, and supply-chain security.

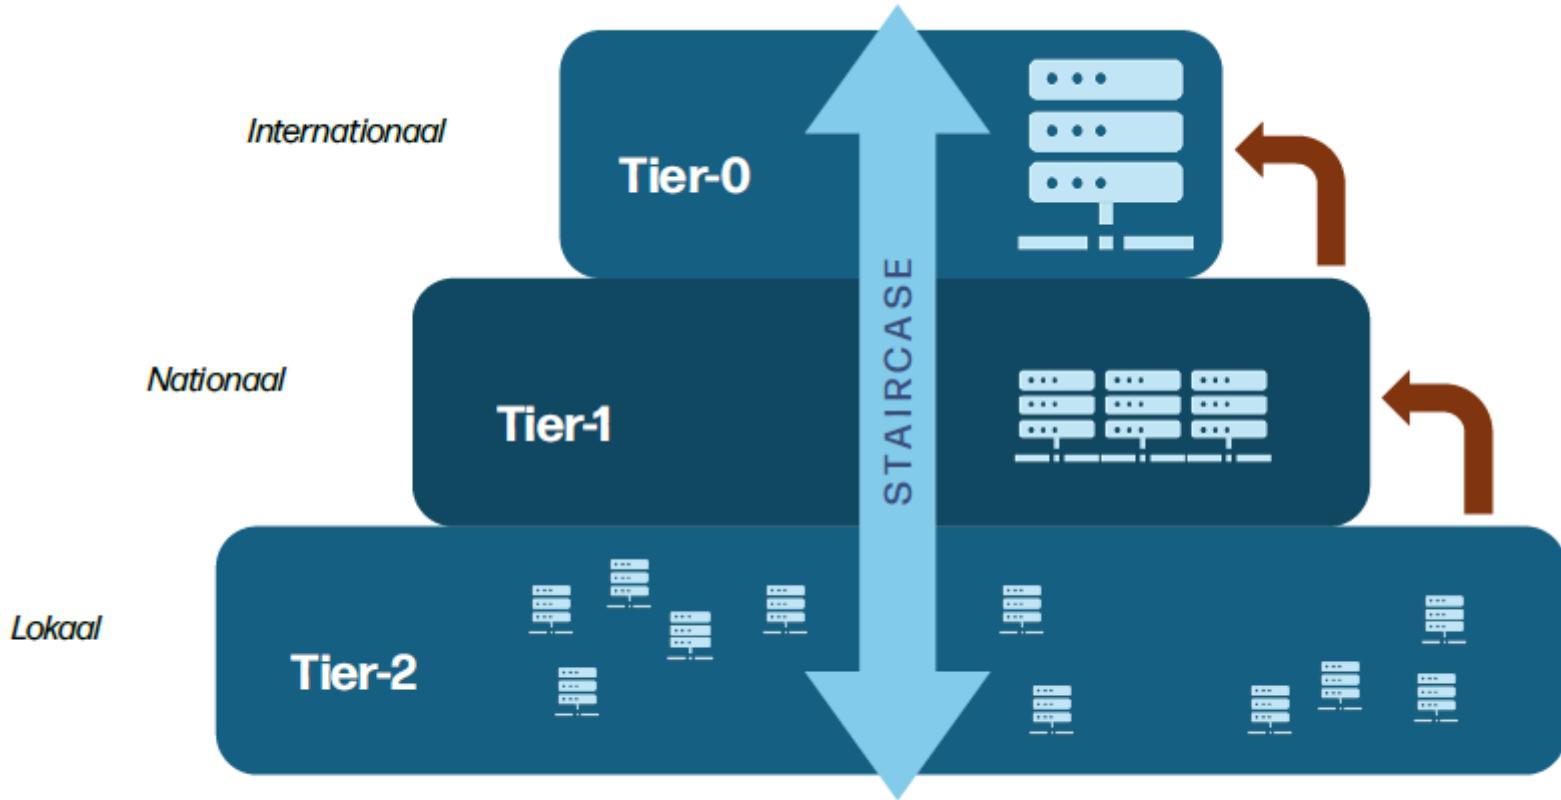
Government authorities will use their purchasing power to enforce safe standard and will put in place minimum requirements for security for central government.’

‘Digitale autonomie moet het uitgangspunt zijn voor de overheid. We kiezen voor een Europese digitale infrastructuur, bouwen strategische afhankelijkheden in cloud, data en cruciale systemen doelgericht af, en we splitsen grote projecten op zodat meer Nederlandse en Europese mkb’ers kunnen meedoen.

Digitale inkoop en aanbestedingen worden gestandaardiseerd en gecentraliseerd, gestuurd op security-by-design, zero-trust, soevereiniteit, open source en ketenveiligheid. De overheid benut haar marktmacht om veilige standaarden af te dwingen en stelt rijksbrede minimumvereisten op voor security.’

<https://www.kabinetsonderhoud2025.nl/documenten/2026/01/30/aan-de-slag---coalitieakkoord-2026-2030>;
Akhan, Groep, Ritzen and Rounding: <https://doi.org/10.53330/ZMFF2486>, Rapport Wennink <https://www.rapportwennink.nl/>





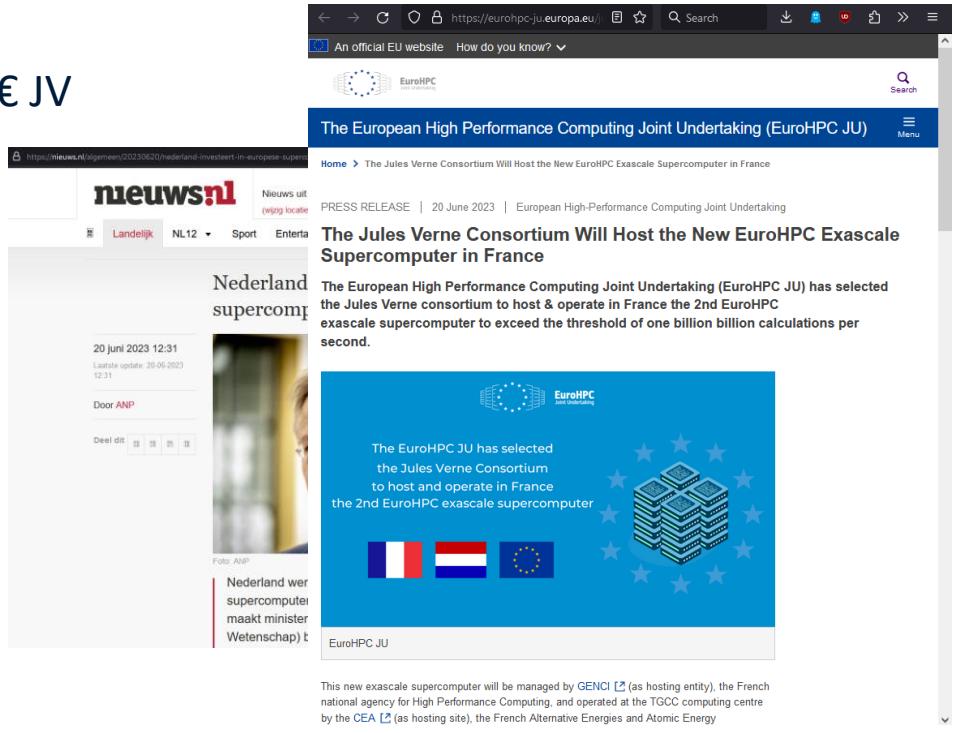
EuroHPC targets large-scale *compute* (and some data)

Dutch direct investments: 2M€ LUMI, 8M€ JV
+ access through ‘Europe’ and the JU

But: it’s not the ‘one single solution’ ...

e.g. EuroHPC has overly many controls,
it being subject to more export controls

- harder to use for research
(like for DestinE portals) that need to run services or use service accounts
- tension with open and citizen science



The image shows a split-screen comparison. On the left is a screenshot of a news article from Nieuws.nl. The headline reads 'Nederland supercompu'. On the right is a screenshot of the EuroHPC website. The main headline on the website is 'The Jules Verne Consortium Will Host the New EuroHPC Exascale Supercomputer in France'. Below the headline is a sub-headline: 'The European High Performance Computing Joint Undertaking (EuroHPC JU) has selected the Jules Verne consortium to host & operate in France the 2nd EuroHPC exascale supercomputer to exceed the threshold of one billion billion calculations per second.' The website features the EuroHPC logo and the European Union flag.

Images: <https://nieuws.nl/algemeen/20230620/nederland-investeert-in-europese-supercomputer/>, https://eurohpc-ju.europa.eu/jules-verne-consortium-will-host-new-eurohpc-exascale-supercomputer-france-2023-06-20_en. EuroHPC comments, see also Thomas Geenen, ECMWF & DestinE (at EGI2023)

Dealing with the digitisation ‘explosion’

Reken er maar (niet meer) op:
de digitale infrastructuren voor onderzoek
2027-2035



NWO en SURF

16 november 2025

<https://www.rijksoverheid.nl/documenten/rapporten/2025/11/18/reken-er-maar-niet-meer-op-de-digitale-infrastructuren-voor-onderzoek-2027-2035>

On ESFRIs, GWIs, the European Open Science Cloud

eoSC EOSC Collaborations The Association Activities EOSC Focus

Search

eoSC Building the EOSC Federation

NWO GROOTSCHALIGE WETENSCHAPPELIJKE INFRASTRUCTUUR

Home Permanente Commissie Nieuws GWI Internationaal Roadmap NL Contact

eoSC Tripartite

Building the EOSC Federation

The vision for EOSC is to put in place a system in Europe to find and access data and services for research and innovation. This is to help researchers store, share, process, analyse and reuse FAIR research outputs within and across disciplines and borders.

The deployment of a network between data repositories and services will be instrumental for Open Science to progress in Europe. For this, the EOSC Federation of nodes is being created.

ESFRI

ESFRI LANDSCAPE ANALYSIS 2014 EXECUTIVE SUMMARY Section 1 THEMATIC AREAS Section 2 CROSS-DOMAIN Annex AUTHORS

Section 1 THEMATIC AREAS DATA, COMPUTING AND DIGITAL RESEARCH INFRASTRUCTURES DOWNLOAD SECTION 1 DOWNLOAD CHAPTER

ESFRI Report Access to Research Infrastructures and Charter on Access to RIs

DOI 10.5281/zenodo.10555986

KAART LIJST

Grootschalige Wetenschappelijke Infrastructuur

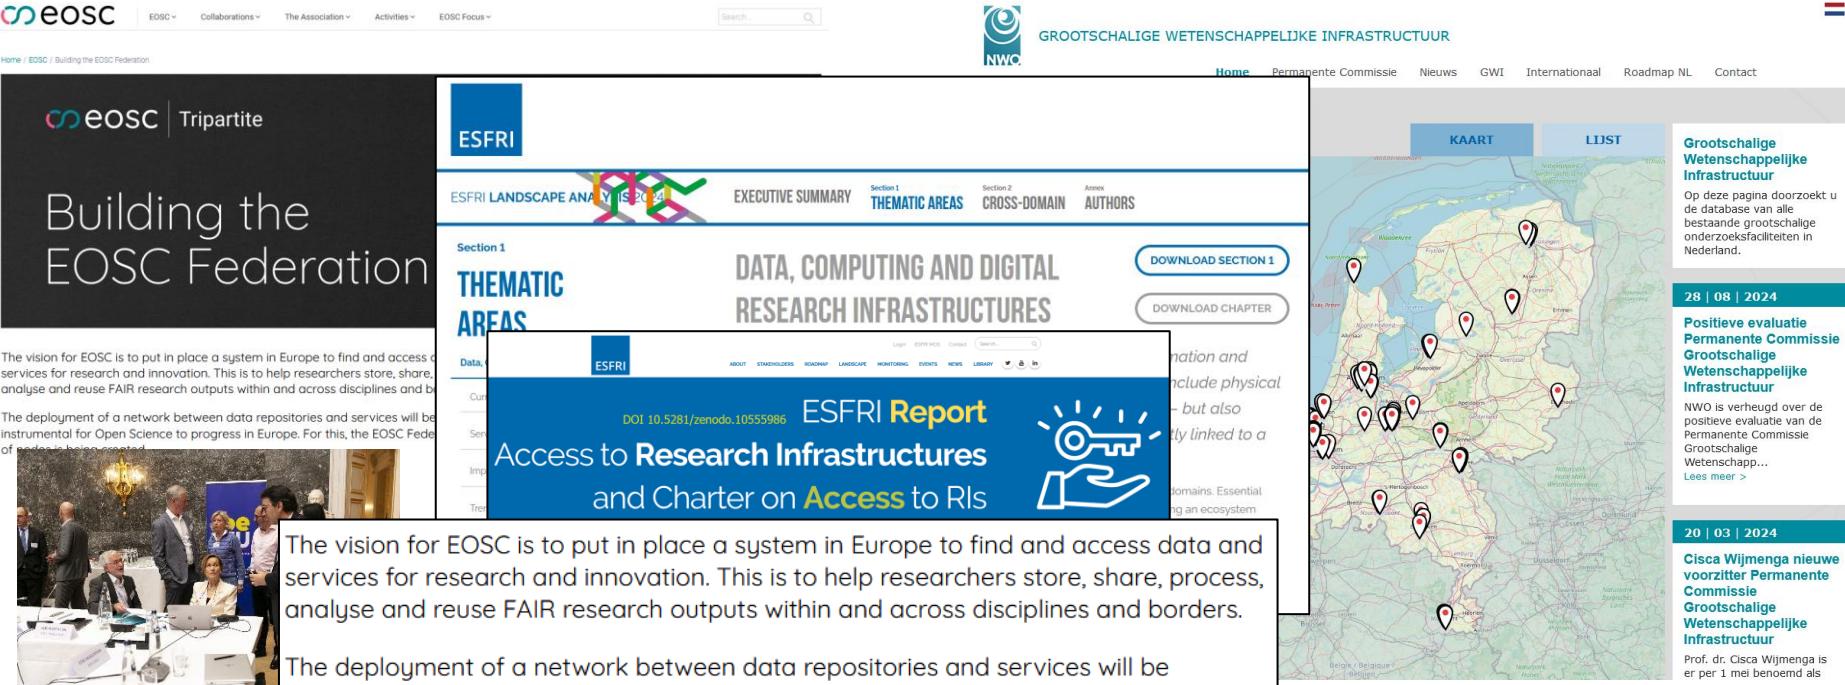
Op deze pagina doet u de database van alle bestaande grootschalige onderzoeksinfrastructuren in Nederland.

28 | 08 | 2024 Positieve evaluatie Permanente Commissie Grootschalige Wetenschappelijke Infrastructuur

NWO is verheugd over de positieve evaluatie van de Permanente Commissie Grootschalige Wetenschappelijke Infrastructuur. Lees meer >

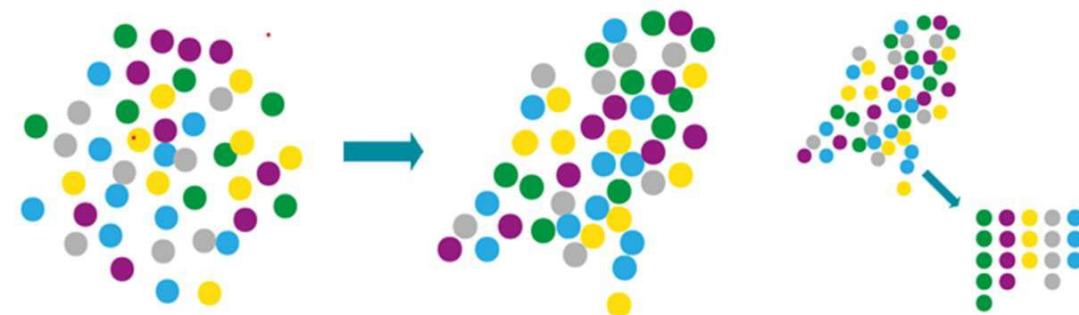
20 | 03 | 2024 Cisca Wijmenga nieuwe voorzitter Permanente Commissie Grootschalige Wetenschappelijke Infrastructuur

Prof. dr. Cisca Wijmenga is er per 1 mei benoemd als



See e.g. <https://www.onderzoeksfaciliteiten.nl/>; <https://www.esfri.eu/>; <https://landscape2024.esfri.eu/>; <https://eosc.eu/building-the-eosc-federation/>

The Dutch Research Infrastructure Landscape Process



Nationale Roadmap Grootchalige Wetenschappelijke Infrastructuur



Landschapsinventarisatie
Continu open voor indiening

GWI – Landschap
elke 5 jaar wordt een volledige update gepubliceerd

Van Landschap naar Roadmap Groepen
Selectie door PC-GWI

Veldproces
Afstemmen, verbinden, prioriteiten stellen door het veld

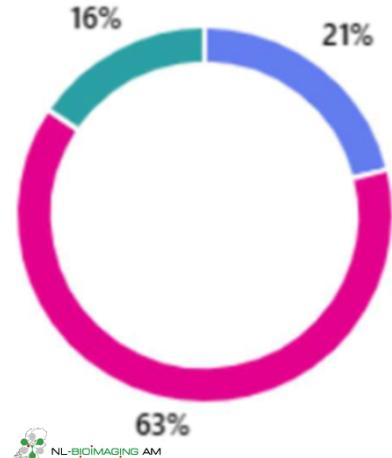
Nationale Roadmap GWI
elke 5 jaar gepubliceerd

<https://www.nwo.nl/en/large-scale-research-infrastructure>

Digitisation is everywhere

80%+ of all LSRIIs in the Netherlands have significant data or digital requirements

- Data/digital is its primary function (e.g. genomic database, AI resource) 38
- It is not its primary function but it has a significant requirement for data/digital infrastructure (e.g. it...) 114
- No 28



With LSRIIs 'practicing what we preach', e.g.:



ARISE (biodiversity), **ODISSEI** (social sciences), **FuSE** (HEP physics & astronomy),
hDMT (human Disease Model Technologies), **NL-BioImaging** (distributed multi-model microscopy)
(... and apologies to those forgotten)

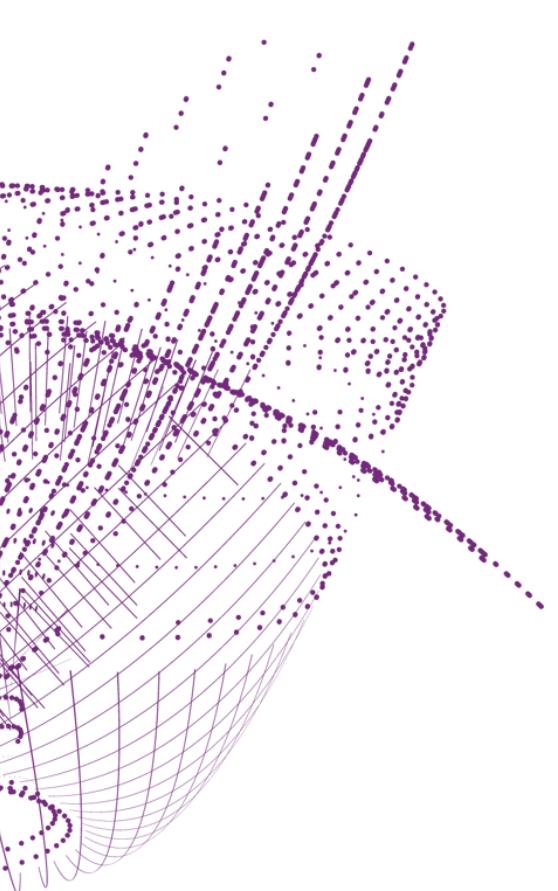
Data from NWO GWI Landscape Analysis, 2025 (Katrien Uytterhoeven *et al.*, NWO); logos from the respective GWIs

And more LSRI's relying on *much more data* than before

GWI	Opslagbehoefte wereldwijd (geschat voor einde van de periode)
CERN / HL-LHC	~800 PB/jaar
SKA	~700 PB/jaar
LOFAR	~6 PB/jaar
ESO / (ALMA WSU, VLT, ELT)	~9 PB/jaar
Copernicus/ESA	~50 PB/jaar
EMBL / ELIXIR	~10 PB/jaar
ITER	>10 PB/jaar

Tabel 5.1.2: Geschatte opslagbehoefte per jaar voor een aantal GWI's waaraan Nederland deelneemt.

From the report 'Reken er maar (niet meer) op', November 2025, <https://open.overheid.nl/documenten/be82689f-08e7-40de-8898-181c119e080a/file>



Research Infrastructure Commons

Infrastructure: for the small and the large



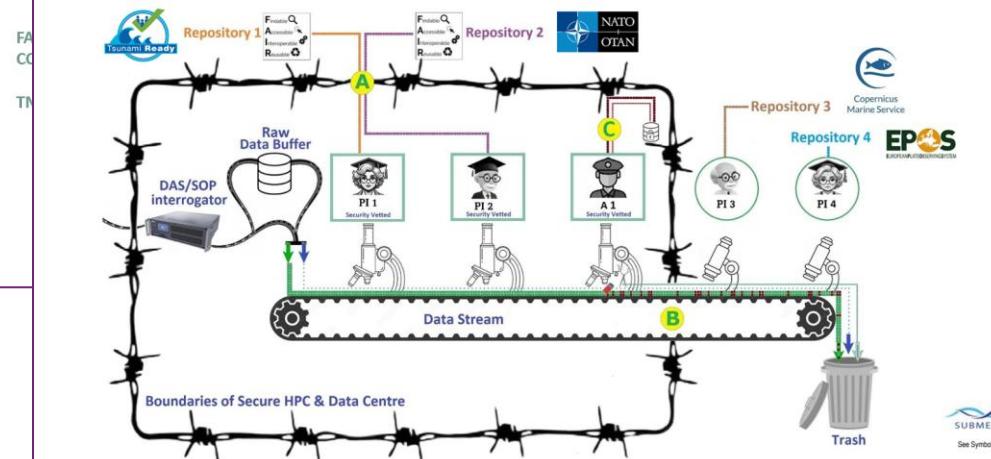
Nikhef

Example of common challenges and foundations: OSSC and .. fiber-optic sensing?

tnc25

Fibre Optic Sensing

A FOS Trusted Research Environment (TRE)



DeIC

Get data that was previously off-limits

The Challenge

Valuable datasets are often inaccessible due to privacy restrictions.

Rich employment records, family studies, and granular demographics could transform your research—but require a secure environment to work with them.

The SANE Solution

Access sensitive data safely in a secure environment.

Use familiar tools like RStudio and Jupyter Notebooks, while SANE ensures all privacy requirements are met.

Why It Works

Built on the Five Safes framework and independent security testing.

SANE lets you start analyzing data in just 30 minutes—no months of approvals needed.

From research idea to secure analysis in four steps

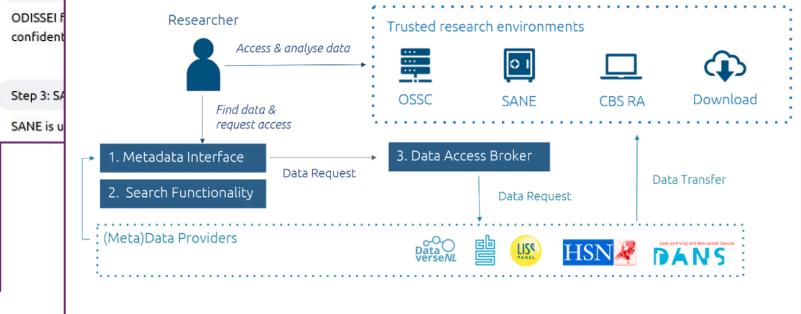
Step 1: Submit Your Data Access Request

Submit data access request describing your research project.

Step 2: Access & Analyse

Step 3: Share & Publish

Step 4: Reuse & Reproduce



FOS slides from DeIC, Rene Belsø, at TNC25 (<https://indico.geant.org/event/5/contributions/420/>), and ODISSEI Secure Supercomputing (OSSC) by SURF & ODISSEI GWI consortium (<https://odissei-data.nl/facility/secure-analysis-environment-sane/>), <https://odissei-data.nl/facility/odissei-portal/>

Identifying elements of the Commons in LSRIs

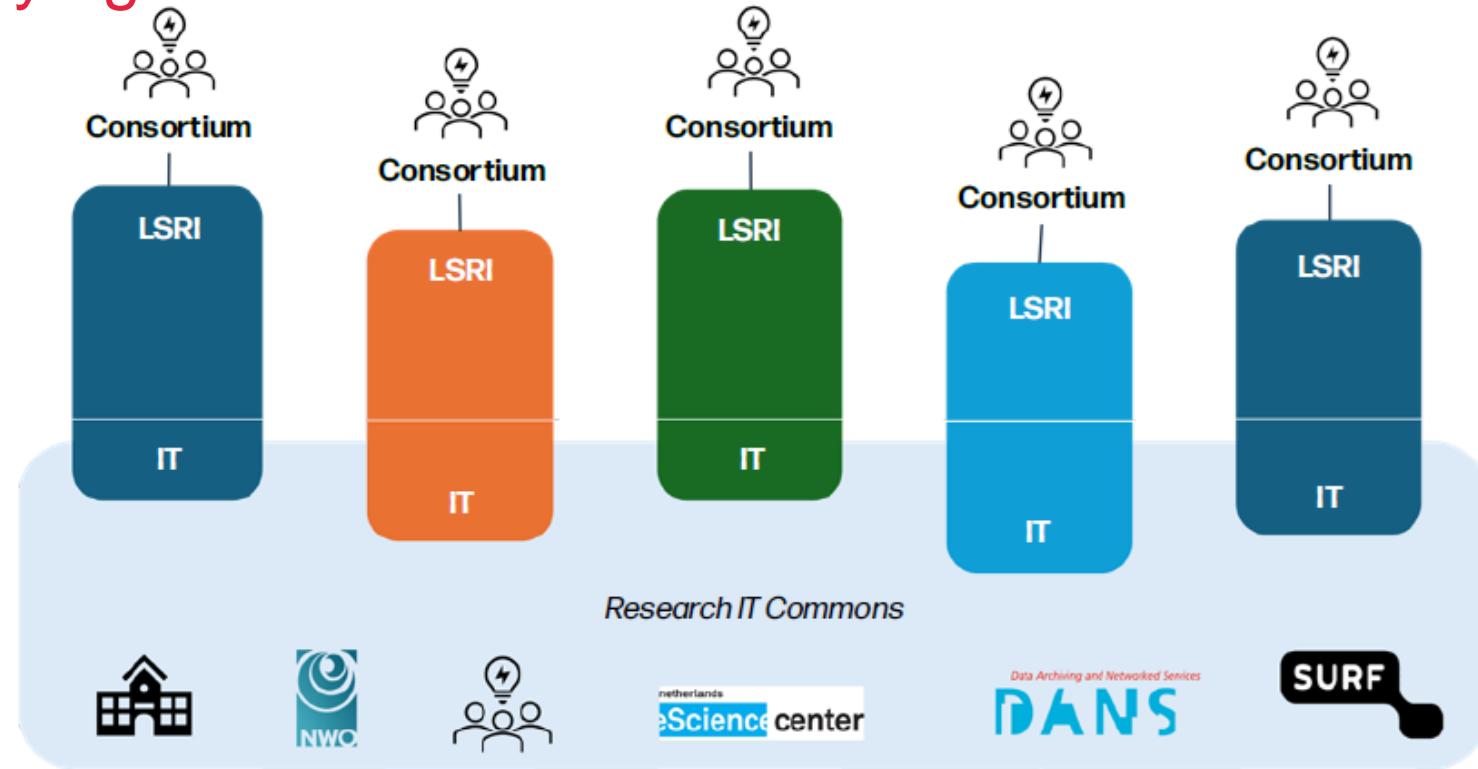
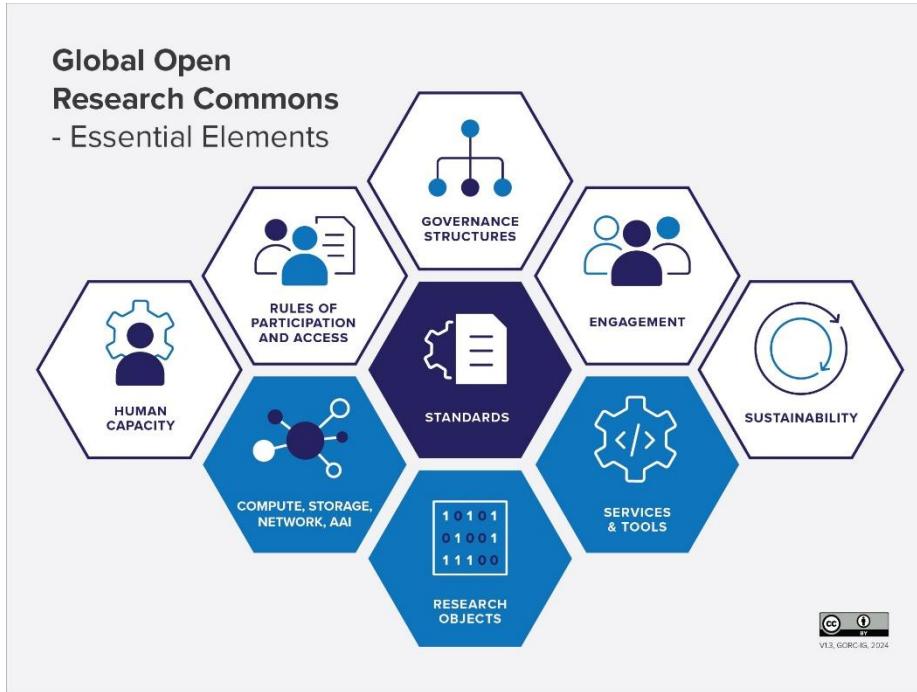


Image source: Machgiel Bijsterbos, SURF

‘Managing the commons’



GORC IG: Typology and Definitions, <https://doi.org/10.15497/RDA00087> - for the original book 'Governing the Commons' see of course Elanor Ostrom ☺

Horizontal platforms – GWI-DIGIT ... and ESFRI-DIGIT



Ontdek de GWI faciliteiten:

- Overzicht 0
- Op de kaart 0

NWO
Grootschalige
Wetenschappelijke
Infrastructuur

Digit-platform

<https://www.esfri.eu/working-groups/data-computing-and-digital-research-infrastructures>

ESFRI

ABOUT STAKEHOLDERS ROADMAP LANDSCAPE MONITORING EVENTS NEWS LIBRARY

HOME > ABOUT > THE WORKING GROUPS

Data, Computing and Digital Research Infrastructures

The Forum

The Working Groups

- Data, Computing and Digital Research Infrastructures
- Energy
- Environment
- Health and Food
- Physical Sciences and Engineering
- Social Sciences and Humanities
- Monitoring Committee
- ESFRI-EOSC Coordination Task Force
- Implementation
- Innovation
- e-Infrastructures Group
- Long-term Sustainability Group
- Neutron Landscape Group

The People



Strategy Working Group & Task Forces

ESFRI STRATEGY WORKING GROUP ON DATA, COMPUTING & DIGITAL RESEARCH INFRASTRUCTURES

As novel proposals with a dominant, or substantial, digital research infrastructure character have been proposed to the 2018 ESFRI Roadmap, ESFRI has established a Strategy Working Group on Data, Computing and Digital Research Infrastructures (SWG DIGIT) to deal with the evaluation of such proposals, therefore completing the set of SWGs for the identified areas of Research Infrastructures of the Roadmap.

The Strategy Working Group on Data, Computing and Digital Research Infrastructures is chaired by Dr Isabel Campos Plasencia. Currently, the group consists of fifteen experts with a complementary expertise that covers the three thematic sub-fields. All experts were nominated by the ESFRI delegations and approved by the forum. The European Commission representatives in the DIGIT WG are actively involved in all activities.

How to make ICT infrastructure into our ‘research instrument’ ?

All these use cases seem diverse, but still result in **common infrastructure capabilities**

- Interactive analysis, collaboration and ‘research service bursting’ platform
 - DSRI is there now to fill this space –can evolve to the ‘interactive gateway’ for all users
- HTC/HPC computing facilities at reasonable ‘T2’ scale, based on application co-design
 - solves short-turnaround needs at limited scale, is the place for growing expertise for scale out to national (SURF) and international (EuroHPC, EGI, EOSC, ...) level
- High-throughput data storage and sharing services
 - targeting data processing compute integration and effective fast access to FAIR data
- Open network for collaborative & data intensive sciences
 - ‘ye shall not have stateful devices in thy data path’ – ScienceDMZ or better
 - is *essential* prerequisite for open science, EOSC, and collaborative (& citizen science) services
- Tools for digital research collaboration
 - sustainable research software, collaborative spaces with *global* partners, SRAM, eduGAIN & EOSC federated access, ubiquitous access to *external R&S* services

Infrastructure is more than just the tools or technology

The '*Uitvoeringsplan*' ('commissie Apers', 2019) deliberately identified digital competences to be broad and include not only data, but also software **and a federated expertise network** at the 'local' digital competence centres (LDCCs):

- "Knooppunt in een gefedereerd netwerk voor data, computing en expertise"
- "Belangrijk is dat de aangesloten lokale infrastructuren middels het gefedereerde systeem geïntegreerd moeten kunnen worden in de European Open Science Cloud (EOSC), die in ontwikkeling is."

This means we require expertise and alignment, also for governance and policy, with the goals for federated Open Science which our nationally initiatives are funding

<https://zoek.officielebekendmakingen.nl/kst-29338-189>

Redefining the concept of service management

IT Service Management (ISO20000, ITIL, FitSM) promote a *specific definition* of service portfolio management:

“Product/service portfolio The product/service portfolio is the complete set of products and/or services *that are managed by the organization*, and it represents the organization’s commitments and investments across all its customers and market spaces. ...”

[ITIL v4, chapter 5]

Concept of ‘services’ in collaborative research is entirely different:

- coming from existing collaborations and infrastructures, many or most services *already exists and used extensively* by research and collaborative administration;
- they are an *essential part of collaborative research*: they should be embraced;
- whether a ‘service’ is operated by a third party, our outside (local) ITSM control is *immaterial to the value of the service*

Research Infrastructures both users and providers ... and 'we' are as well!

The screenshot shows two main web pages. On the left, the WLCG (Worldwide LHC Computing Grid) website has a sidebar with links for 'Who can use', 'Security', 'Certificates', 'Software', 'Tools', and 'Monitoring/Visualisation'. The main content area is titled 'Who can use WLCG' and lists 'Schools and Individuals', 'Site admins and grid users', and 'Set up new sites or new Federations'. On the right, the IGWN (International Gravitational-Wave Network) Public Alerts User Guide page has a sidebar with links for 'IGWN | Public Alerts User Guide', 'Getting Started Checklist', 'Observing Capabilities', 'Data Analysis', 'Alert Contents', 'Sample Code', 'Additional Resources', 'Early-Warning Alerts', 'Change Log', and 'Glossary'. The main content area is titled 'LIGO/Virgo/KAGRA Public Alerts User Guide' and features images of a lightcurve from Fermi/GBM, a gravitational-wave time-frequency map, and a map of the sky showing the locations of LIGO/Virgo, Fermi/GBM, and INTEGRAL, with a specific event marked.



All European Social Survey

All European Social Survey (ESS) data and documentation is now only accessible through the ESS Data Portal.

The European Social Survey (ESS) is a pan-European research infrastructure providing freely accessible data for academics, policymakers, civil society and the wider public.

<https://rtd.igwn.org/>; <https://www.europeansocialsurvey.org/>; <https://wlcg.web.cern.ch/>; <https://www.slices-ri.eu/consortium-netherlands/>

The screenshot shows the SLICES-NL (Scientific Large Scale Infrastructure for Computing/Communication Experimental Studies) website. A yellow banner at the top reads 'Scientific Large Scale Infrastructure for Computing/Communication Experimental Studies'. Below it, a blue banner features the text 'Netherlands'. The main content area describes the SLICES-NL project, stating: 'The SLICES-NL will support experimentation on variety of technologies related to data-centric Complex Cyber Infrastructure (CCI), Big Data Infrastructure and technologies (BDIT), and future Internet. This includes programmable network infrastructure, power and energy optimisation in distributed cloud to edge computation and data repository sharing to support innovation in Open Science and Industry 4.0. The SLICES-NL will support architecture research for future Internet and data centric and user centric infrastructures, including architecture aspect of the future RI Platform as a Service (PRIaaS), multi-cloud

And it needs everyone to work together

To scale trust in research infrastructures,
we need to keep challenging ourselves ...

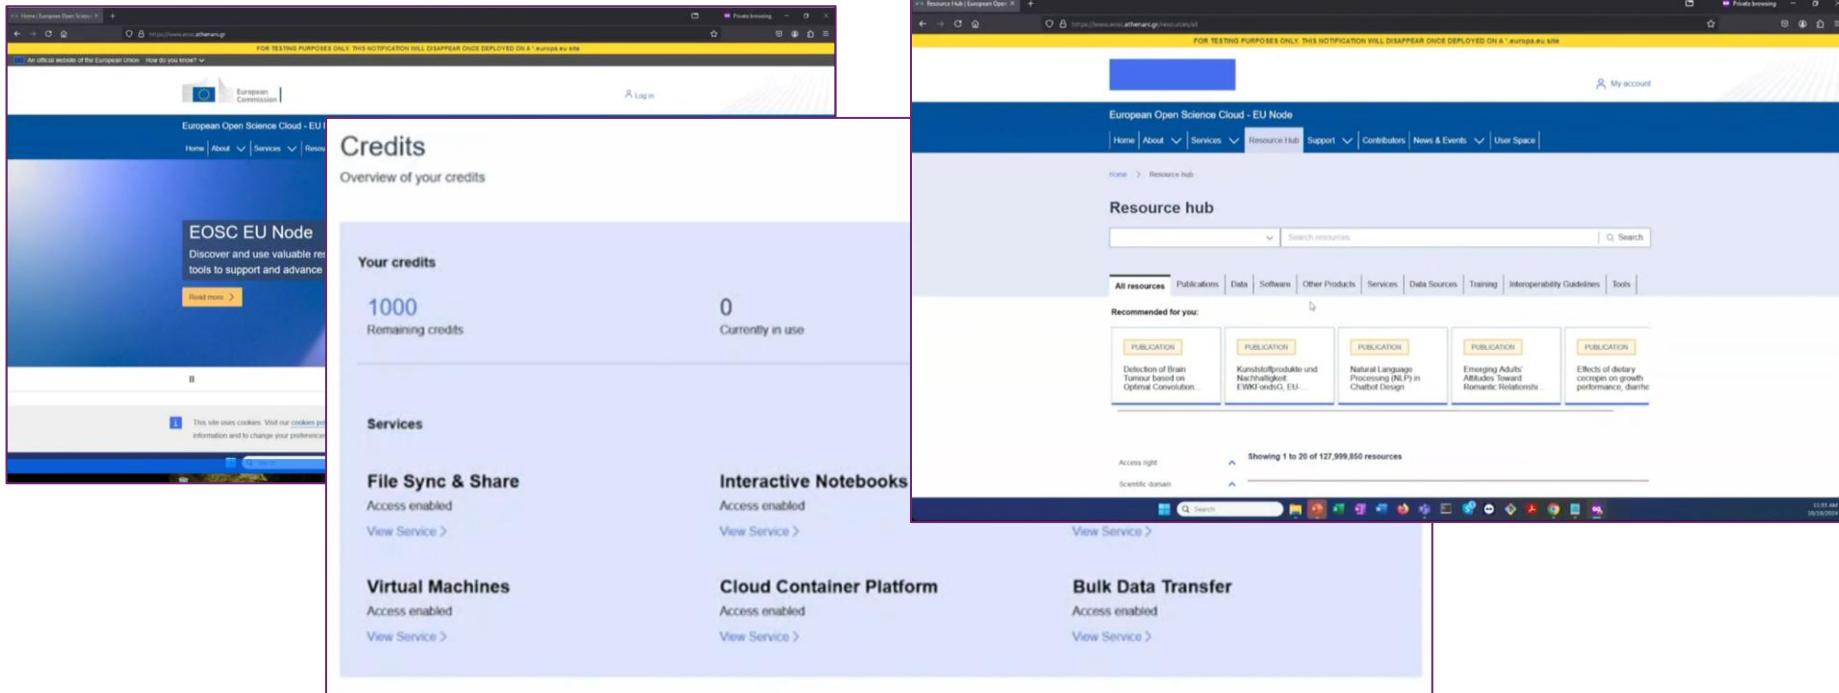
- *for eduGAIN: do we choose more trustworthiness and target baseline assurance, or more inclusiveness, but maybe less trust?*
- *for your university IT department: prioritize the primary mission of education and research, as both are now globally connected*
 - ‘we can use existing services from outside’
 - ‘we can contribute in collaborations in education and research’
 - ‘we teach our students to understand, study, and work with interconnected services and systems that are globally connected’

... rather than get stuck in an enterprise egg-shell approach?
- *do our networks support a perimeter ‘fit for collaboration’?*

The screenshot shows the GEANT Entities Database Explorer interface. At the top, there are navigation links: Members, Joining, Tools, Policy framework, and Operations. Below that is a sub-navigation bar: Tools > Entities Database Explorer. The main area has two filter dropdowns: 'Entity Type' (set to 'All') and 'Entity category filter' (set to 'No filter'). A search bar contains the text 'Slovenia'. A dropdown menu for 'Entity category filter' is open, showing a list of categories starting with 'No filter', 'Any category', and 'Data Protection Code of Conduct v1'. The search results list 'Slovenia' as the first result. Below the search results is a 'start typing below to narrow the list' input field. The bottom part of the screenshot shows a 'Error - Access denied' message from Maastricht University. The message says: 'Message from your institution: Please contact servicedesk-icts@maastrichtuniversity.nl if you need access to this service.' It also includes a note: 'You cannot use edu.nl because your institution limits access to it (the "Service Provider") with an authorization policy. Please contact the service desk of your institution if you think you should be allowed access to edu.nl.' At the bottom of the message are links for 'SURFconext Wiki', 'Service desk', and 'servicedesk-icts@maastrichtuniversity.nl'.

Images: <https://technical.edugain.org/entities>, Maastricht University blocking access to ... a privacy-friendly URL shortener ☺,

And service use is coming from the EOSC 'whether you want it or not' ...



The image displays two screenshots of the European Open Science Cloud (EOSC) EU Node interface. The left screenshot shows the 'Credits' page, which provides an overview of your credits. It displays '1000 Remaining credits' and '0 Currently in use'. Below this, there are sections for 'Services' (File Sync & Share, Virtual Machines, Interactive Notebooks) and 'Cloud Container Platform'. The right screenshot shows the 'Resource hub' page, which displays a search bar and a grid of recommended publications and services. The publications listed are: 'Detection of Brain Tumour based on Optimal Convolution', 'Kunststoffprodukte und Nachhaltigkeit', 'Natural Language Processing (NLP) in Chatbot Design', 'Emerging Adults' Attitudes toward Romantic Relationships', and 'Effects of dietary changes on growth performance in mice'.

<https://webcast.ec.europa.eu/eu-node-technical-launch-event-24-10-10>

And education labs are much like ad-hoc research collaboration



just slightly more organised than research ... ?

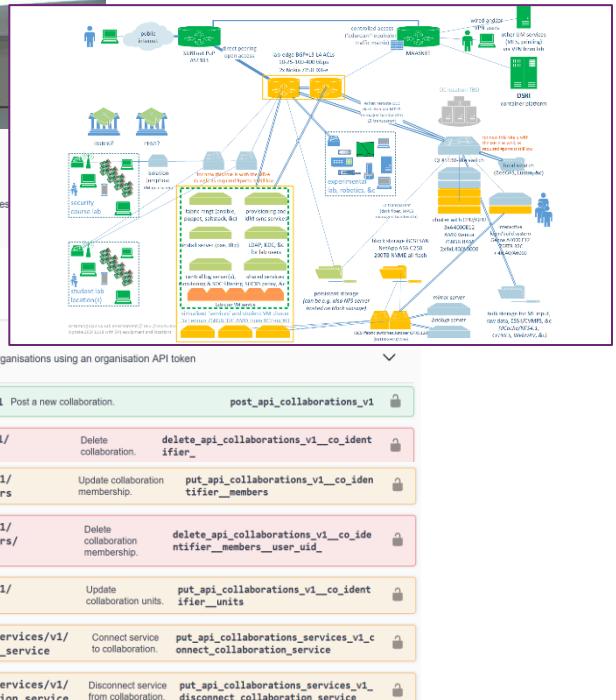
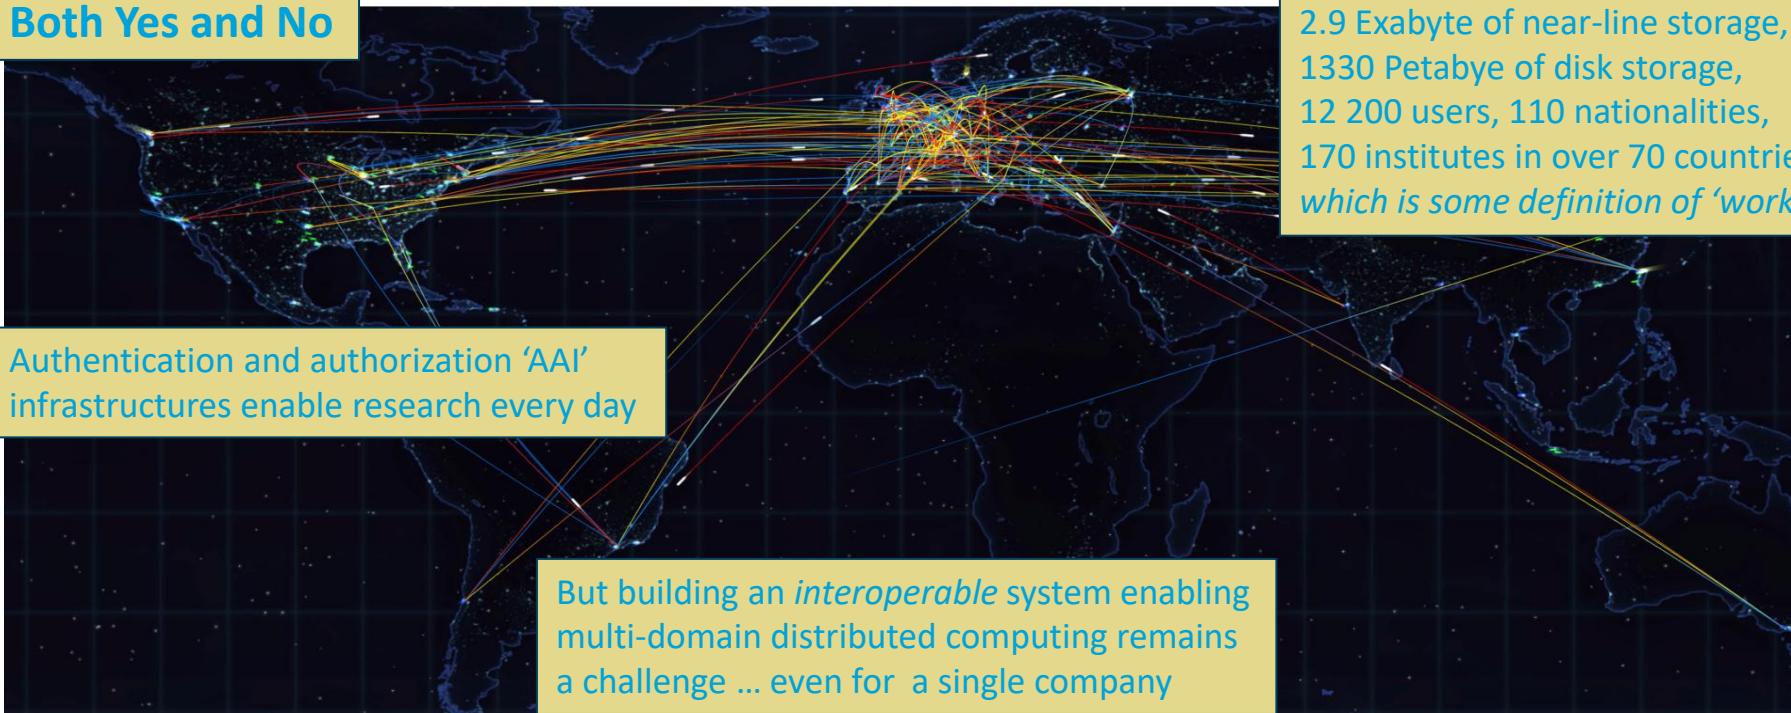


Photo by sunrise University on Unsplash; network diagram: FSE CSLab, Maastricht University; SRAM API: <https://sram.surf.nl/apidocs/>

So did we solve our - research - infrastructure challenge?

Both Yes and No

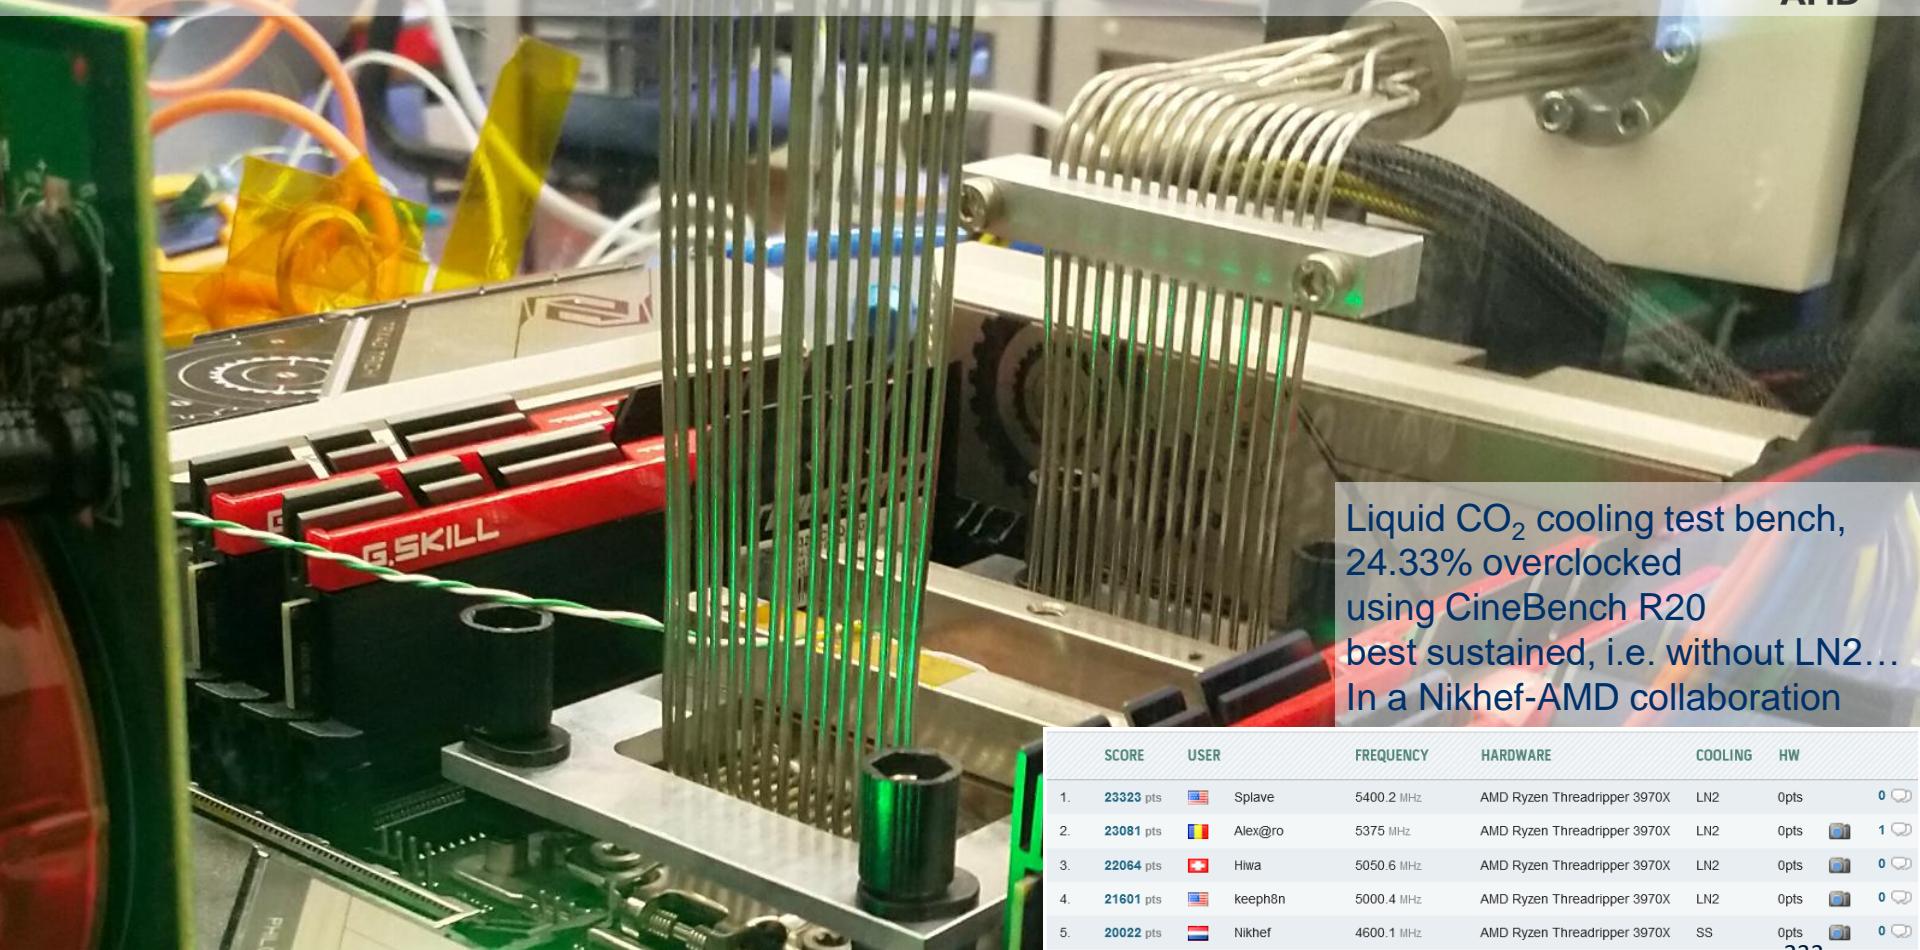


Even WLCG alone today comprises over 900 000 CPU cores, 2.9 Exabyte of near-line storage, 1330 Petabyte of disk storage, 12 200 users, 110 nationalities, 170 institutes in over 70 countries which is some definition of 'works'

Looking for the common pattern ...

- It's all about *balance*
 - systems are like congested highways: no use solving just *one* bottleneck
 - and the bottlenecks may be inside the system as well as in interconnects, trust, interoperability, and governance
- Scaling *collaboration and trust* federation is as complex as scaling systems
 - composing services across administrative domains is ubiquitous
 - but beyond a certain size, $\mathcal{O}(100)$, you will also find need for policy and review

And you may move problems around, but it's hard to actually *solve* them!



Liquid CO₂ cooling test bench,
24.33% overclocked
using CineBench R20
best sustained, i.e. without LN2...
In a Nikhef-AMD collaboration

SCORE	USER	FREQUENCY	HW	COOLING	HW
1. 23323 pts	Splave	5400.2 MHz	AMD Ryzen Threadripper 3970X	LN2	0 pts
2. 23081 pts	Alex@ro	5375 MHz	AMD Ryzen Threadripper 3970X	LN2	0 pts
3. 22064 pts	Hiwa	5050.6 MHz	AMD Ryzen Threadripper 3970X	LN2	0 pts
4. 21601 pts	keeph8n	5000.4 MHz	AMD Ryzen Threadripper 3970X	LN2	0 pts
5. 20022 pts	Nikhef	4600.1 MHz	AMD Ryzen Threadripper 3970X	SS	0 pts

Discussion time ... !

All views expressed herein are my own, and do not necessarily reflect the position of any co-funding organisations, the European Commission, SURF, NWO, or collaborators mentioned below. Yet works presented here are part of many collaborations. Thanks especially to – in random order – Mischa Sallé, Tristan Suerink, Dennis van Dok, Mary Hester, Andrew Pickford, Jeff Temperton, Roel Aaij, Emily Kooistra, Jouke Roorda, Sven Gabriel, Lennie de Roo, Arjen van Rijn, Krista de Roo (Nikhef), Nicolas Liampotis, Kyriakos Glinis (GRNET), Liam Atherton, Jens Jensen, Dave Kelsey, David Crooks (STFC RAL), Hannah Short, Liviu Valsan (CERN), Uros Stevanovic, Marcus Hardt (KIT), Maarten Kremers (SURF), Licia Florio (NORDUNET), Christos Kanellopoulos, Klaas Wieringa (GEANT), Tom Barton (Internet2, UChicago), Tiziana Ferrari, Matt Viljoen, Baptise Grenier (EGI.eu), and the EGI, GEANT, PRACE, WLCG, REFEDS, WISE, IGTf, AARC & FIM4R communities!

David Groep

david.groep@maastrichtuniversity.nl

davidg@nikhef.nl Nikhef

<https://www.nikhef.nl/~davidg/presentations/>

 <https://orcid.org/0000-0003-1026-6606>



for all content not explicitly otherwise attributed in the slides
All trademarks, logos and brand names are the property of their respective owners. All company, product and service names used herein are for identification purposes only. Use of these names, trademarks and brands does not imply endorsement.



Elements of this work have been co-supported by projects that have received funding from the European Union's Horizon research and innovation programmes under grants GN4-3, EOSC Future, AARC2, AARC-TREE, GN5-1, and GN5-2.



Elements of this work have been co-supported by SURF, the collaborative organisation for IT in Dutch education and research.





This work has also been co-supported by projects that have received funding from the European Union's Horizon research and innovation programmes under AARC-TREE and GN5-2.

SURF

This work is supported by SURF under the Innovation Programme, part of the Execution Plan Digital Infrastructures for Research and other sources

time for (more) discussion ...



Maastricht University

Nikhef

David Groep
davidg@nikhef.nl

<https://www.nikhef.nl/~davidg/presentations/>
 <https://orcid.org/0000-0003-1026-6606>



Nulla folia post hoc sunt

Thanks for watching!

*“En daarmee, geachte luisteraars, laat ik u over aan
de verpozing die uw babbelklant u gemeenlijk pleegt te bieden.”*

