ICT as the research instrument for our collaborative world

exploring the federated infrastructures for data, computing, networks, trust & identity

- *Research Infrastructures and computing needs*
- *'More than one': building computing and network fabrics*
- *Beyond commodity – innovation for enabling next gen research*
- *Infrastructure for Collaboration: trust and identity*
- *Research Overlays and the EOSC*
- *Infrastructure is everywhere: the RCauth example*
- *Much in common: horizontal ICT infra and digital competences*
- *Everyone should join in: expertise and essential ICT instrumentation*

David Groep
DACS and Nikhef

Nikhef

CI-Office | Verdiepingssessie
July 2023

Maastricht University | Department of Advanced Computing Sciences

ALFR. NOBEL
NAT. MDCCC XXXIII OB. MDCCC XCVI
PETER W. HIGGS
KUNGL. VETENSKAPS AKADEMIEN
FRANÇOIS ENGLERT
KUNGL. VETENSKAPS AKADEMIEN

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

A ‘big science’ facility: the Large Hadron Collider at CERN

1964

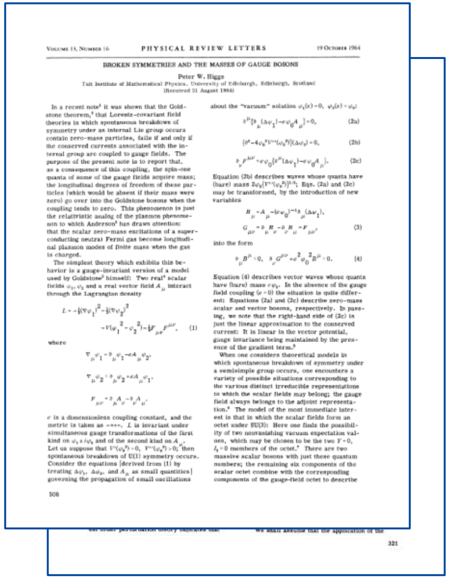

Volume 13, Number 16 PHYSICAL REVIEW LETTERS 19 October 1964

BROKEN SYMMETRIES AND THE MASSES OF GAUGE BOSONS

Peter W. Higgs

Tait Institute of Mathematical Physics, University of Edinburgh, Edinburgh, Scotland

(Received 31 August 1964)

In a recent note[1] it was shown that the Goldstone theorem,[2] that Lorentz-covariant field theories in which spontaneous breakdown of symmetry under an internal Lie group occurs contain zero-mass particles, fails if and only if the conserved currents associated with the internal group are coupled to gauge fields. The purpose of the present note is to report that, as a consequence of this coupling, the spin-one quanta of some of the gauge fields acquire mass; the longitudinal degrees of freedom of these particles (which would be absent if their mass were zero) go over into the Goldstone bosons when the coupling tends to zero. This phenomenon is just the relativistic analog of the plasmon phenomenon to which Anderson[3] has drawn attention: that the scalar zero-mass excitations of a superconducting neutral Fermi gas become longitudinal plasmon modes of finite mass when the gas is charged.

The simplest theory which exhibits this behavior is a gauge-invariant version of a model used by Goldstone[2] himself: Two real[4] scalar fields φ_1, φ_2 and a real vector field A_μ interact through the Lagrangian density

$$L = -\tfrac{1}{2}(\nabla\varphi_1)^2 - \tfrac{1}{2}(\nabla\varphi_2)^2 - V(\varphi_1{}^2 + \varphi_2{}^2) - \tfrac{1}{4}F_{\mu\nu}F^{\mu\nu}, \qquad (1)$$

where

$$\nabla_\mu\varphi_1 = \partial_\mu\varphi_1 - eA_\mu\varphi_2,$$

$$\nabla_\mu\varphi_2 = \partial_\mu\varphi_2 + eA_\mu\varphi_1,$$

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu,$$

e is a dimensionless coupling constant, and the metric is taken as $-+++$. L is invariant under simultaneous gauge transformations of the first kind on $\varphi_1 \pm i\varphi_2$ and of the second kind on A_μ. Let us suppose that $V'(\varphi_0{}^2) = 0$, $V''(\varphi_0{}^2) > 0$; then spontaneous breakdown of U(1) symmetry occurs. Consider the equations [derived from (1) by treating $\Delta\varphi_1$, $\Delta\varphi_2$, and A_μ as small quantities] governing the propagation of small oscillations about the "vacuum" solution $\varphi_1(x) = 0$, $\varphi_2(x) = \varphi_0$:

$$\partial^\mu\{\partial_\mu(\Delta\varphi_1) - e\varphi_0 A_\mu\} = 0, \qquad (2a)$$

$$\{\partial^2 - 4\varphi_0{}^2 V''(\varphi_0{}^2)\}(\Delta\varphi_2) = 0, \qquad (2b)$$

$$\partial_\nu F^{\mu\nu} = e\varphi_0\{\partial^\mu(\Delta\varphi_1) - e\varphi_0 A_\mu\}. \qquad (2c)$$

Equation (2b) describes waves whose quanta have (bare) mass $2\varphi_0\{V''(\varphi_0{}^2)\}^{1/2}$; Eqs. (2a) and (2c) may be transformed, by the introduction of new variables

$$B_\mu = A_\mu - (e\varphi_0)^{-1}\partial_\mu(\Delta\varphi_1),$$

$$G_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu = F_{\mu\nu}, \qquad (3)$$

into the form

$$\partial_\mu B^\mu = 0, \quad \partial_\nu G^{\mu\nu} + e^2\varphi_0{}^2 B^\mu = 0. \qquad (4)$$

Equation (4) describes vector waves whose quanta have (bare) mass $e\varphi_0$. In the absence of the gauge field coupling ($e = 0$) the situation is quite different: Equations (2a) and (2c) describe zero-mass scalar and vector bosons, respectively. In passing, we note that the right-hand side of (2c) is just the linear approximation to the conserved current: It is linear in the vector potential, gauge invariance being maintained by the presence of the gradient term.[5]

When one considers theoretical models in which spontaneous breakdown of symmetry under a semisimple group occurs, one encounters a variety of possible situations corresponding to the various distinct irreducible representations to which the scalar fields may belong; the gauge field always belongs to the adjoint representation.[6] The model of the most immediate interest is that in which the scalar fields form an octet under SU(3): Here one finds the possibility of two nonvanishing vacuum expectation values, which may be chosen to be the two $Y = 0$, $I_3 = 0$ members of the octet.[7] There are two massive scalar bosons with just these quantum numbers; the remaining six components of the scalar octet combine with the corresponding components of the gauge-field octet to describe

508

est order perturbation theory indicates that

we shall assume that the application of the

321

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

16823 characters, 165 kByte PDF

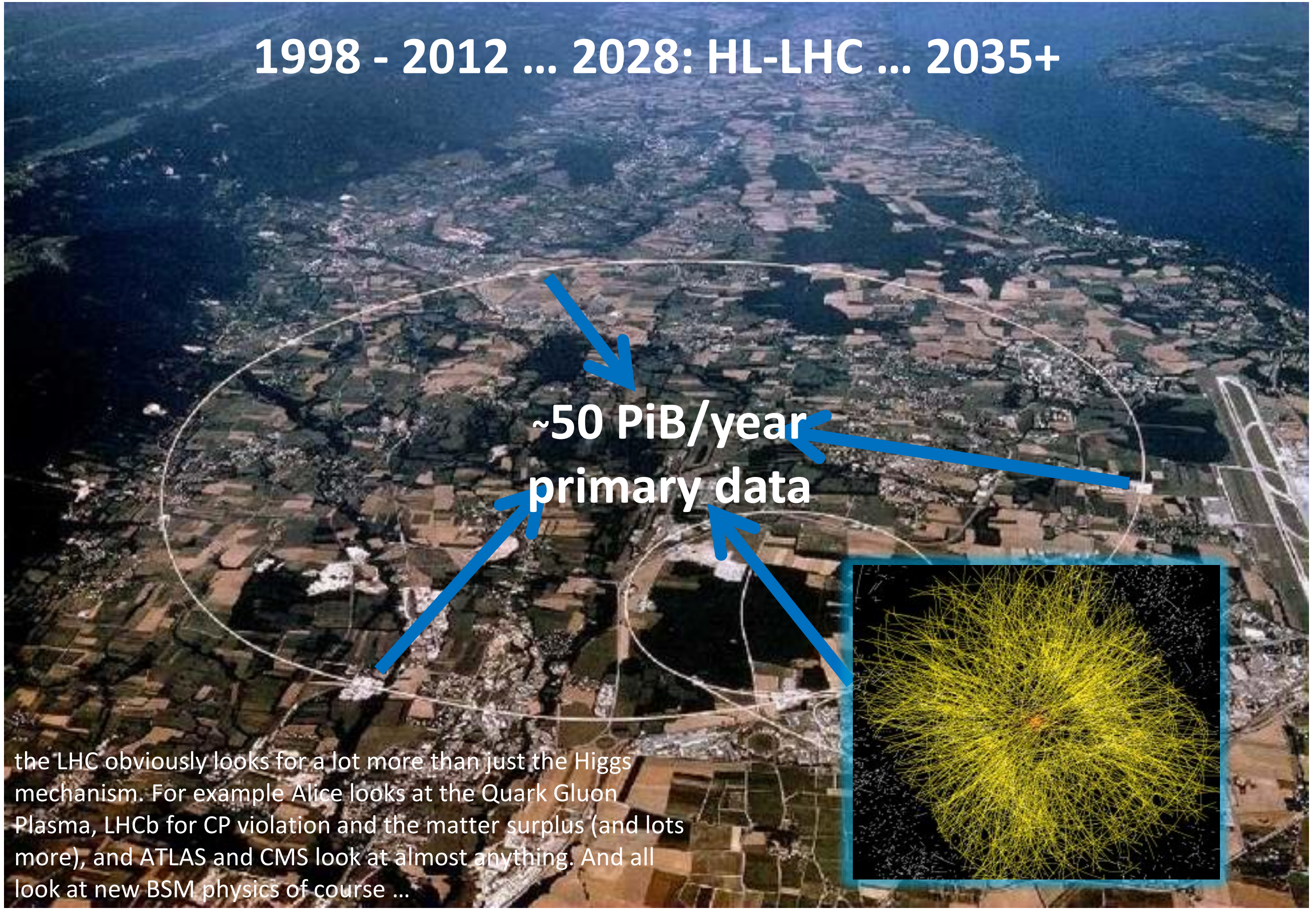

1998 - 2012 … 2028: HL-LHC … 2035+

~50 PiB/year primary data

the LHC obviously looks for a lot more than just the Higgs mechanism. For example Alice looks at the Quark Gluon Plasma, LHCb for CP violation and the matter surplus (and lots more), and ATLAS and CMS look at almost anything. And all look at new BSM physics of course …

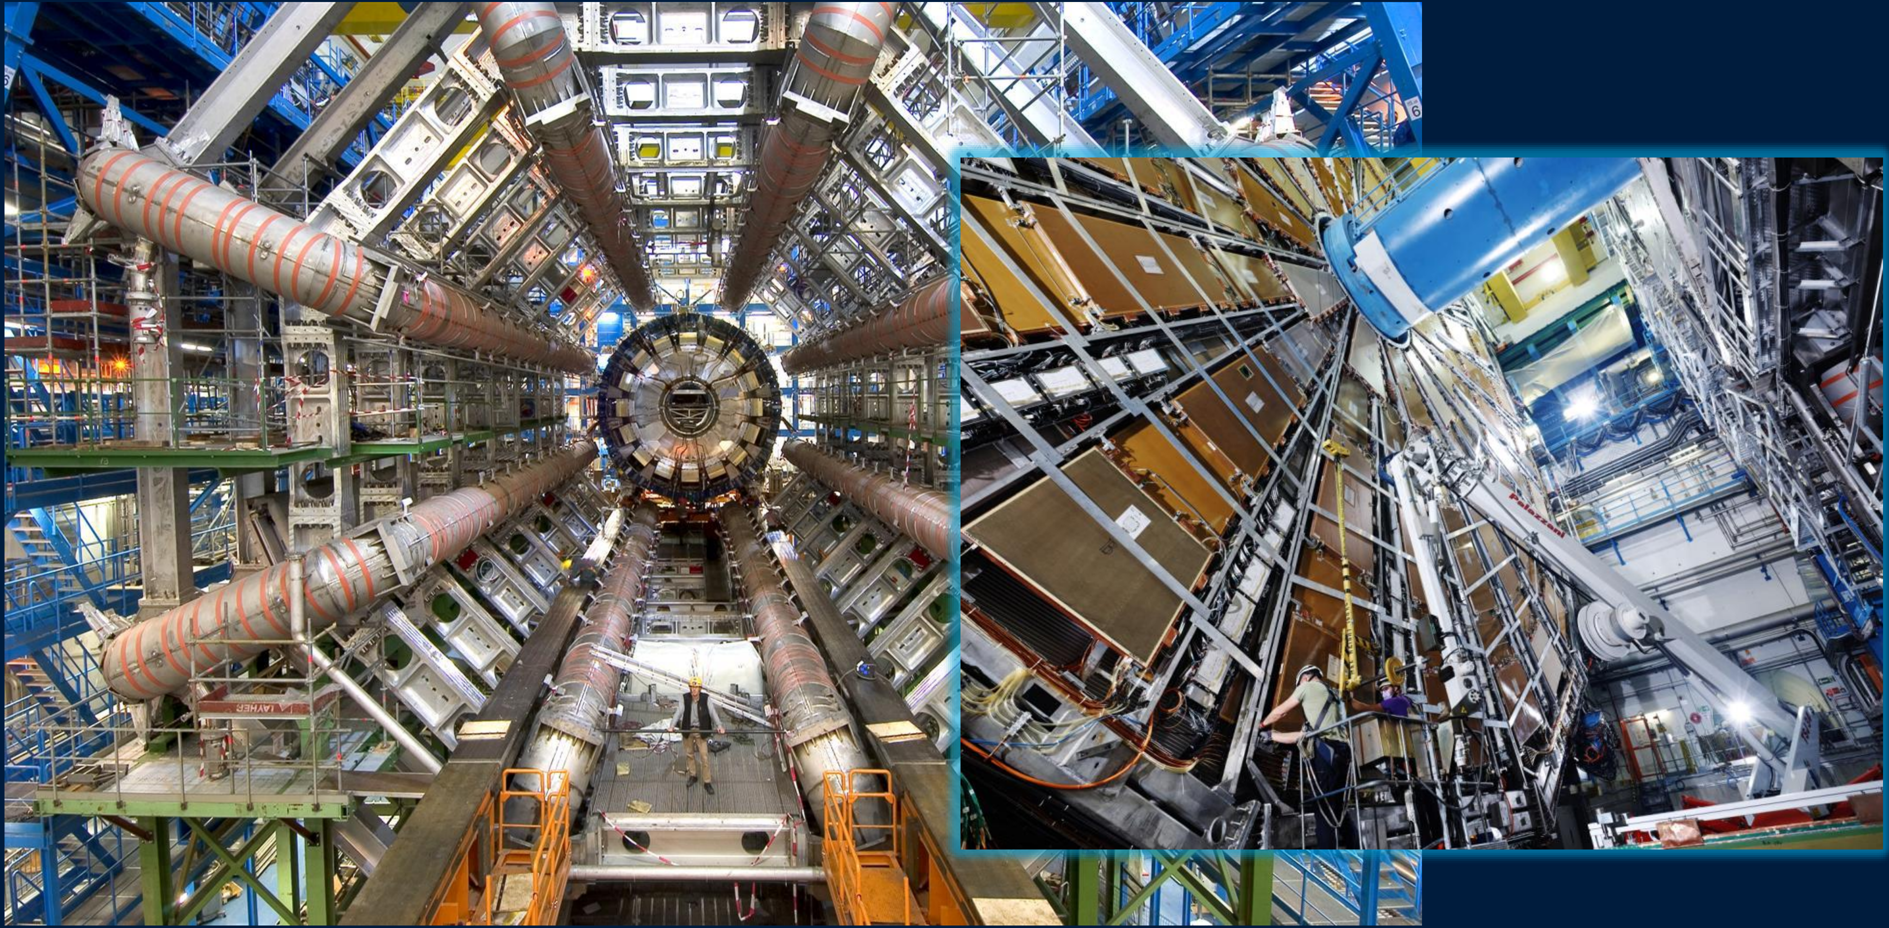

Images: ATLAS detector in the cavern at CERN. Source: CERN

Detector to doctor workflow (LHC example)

ATLAS: 1 of 4 at the LHC

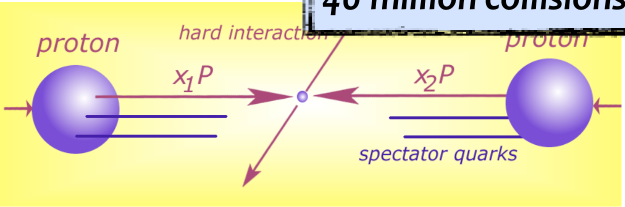

40 million collisions / second

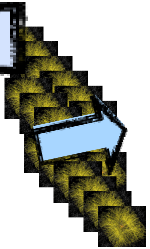

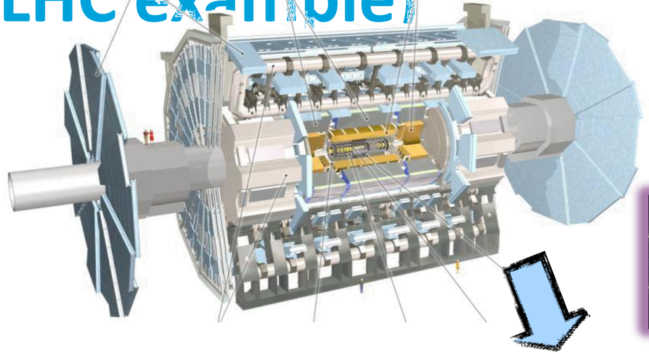

Trigger system selects 600 Hz ~ 1 GB/s data

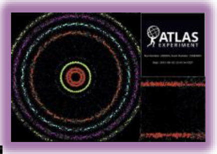

Classify particles in collision and their physics properties:
- *electrons*
- *muons*
- *jets consisting of hadrons*

…

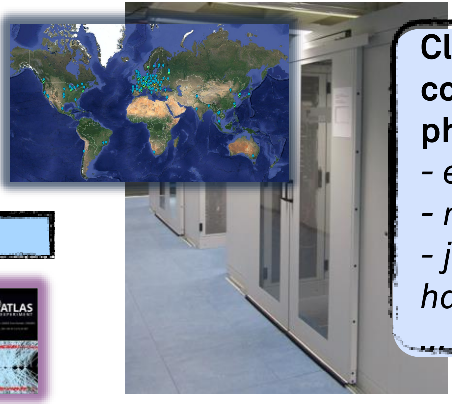

1 'T0' and 170 institutes

Physics analysis by (PhD) students, in papers & analysis notes

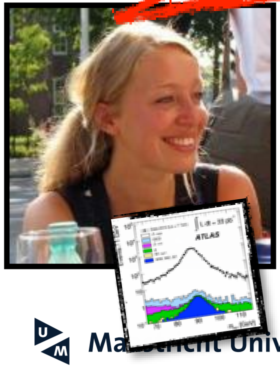

LHC: ~10k researchers

diagram adapted from Frank Linde; images: ATLAS collaboration, Nikhef …
and thanks to Rosemarie

Computing on lots of data – 40Mevents/sec

~ 10 seconds to compute a single event at ATLAS for ‘jets’ containing ~30 collisions

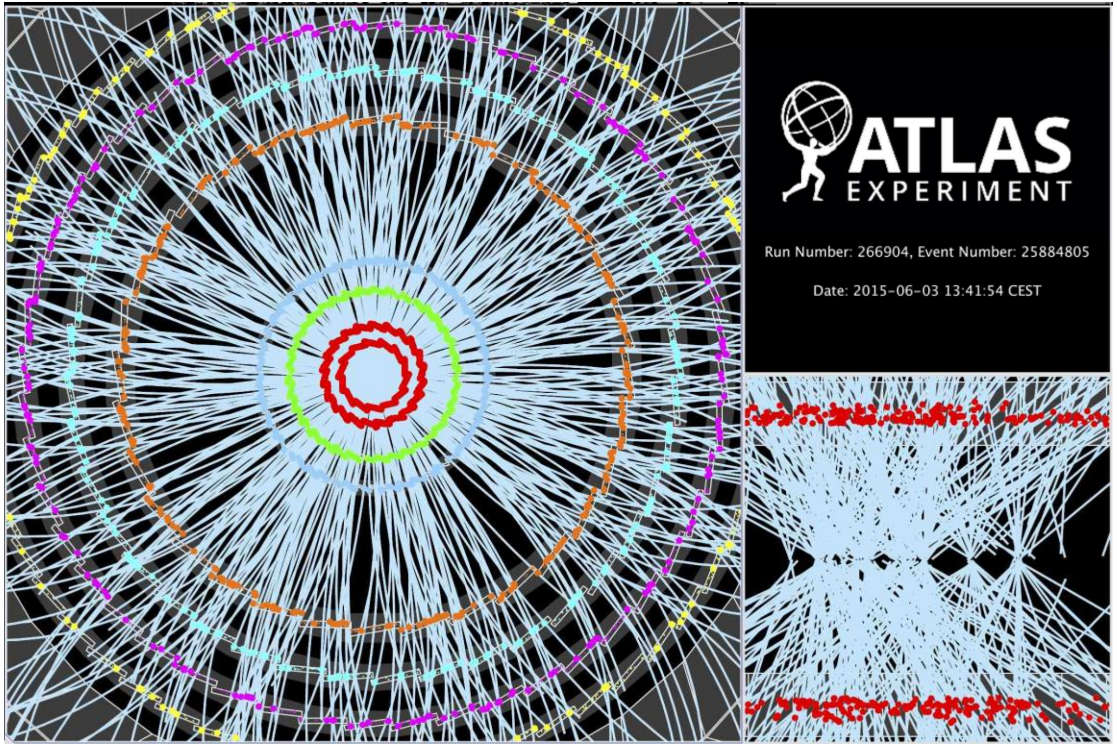
ATLAS
EXPERIMENT
Run Number: 266904, Event Number: 25884805
Date: 2015-06-03 13:41:54 CEST

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)

Processing at scale for data intensive science

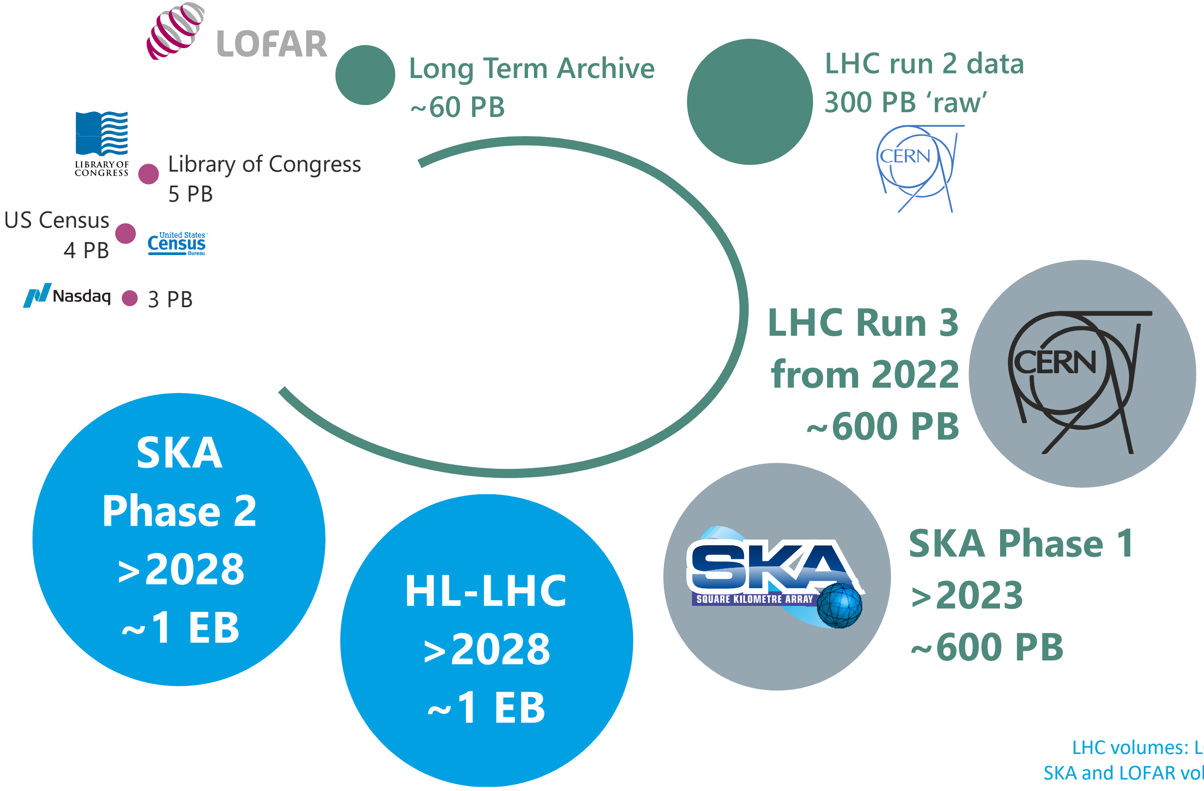

LOFAR

Long Term Archive
~60 PB

LHC run 2 data
300 PB 'raw'

Library of Congress
5 PB

US Census
4 PB

Nasdaq 3 PB

LHC Run 3
from 2022
~600 PB

SKA
Phase 2
>2028
~1 EB

HL-LHC
>2028
~1 EB

SKA Phase 1
>2023
~600 PB

Data from various sources, for public entities: data ca. 2018, indicative, within ~ factor 2
LHC volumes: LCG Resource Scrutiny Group & CERN; 2020
SKA and LOFAR volumes: ASTRON/Michiel van Haarlem, 2020

Scaling computing infra: volume not the only thing that matters

EOSC-WeNMR portals

@Bonvinlab

HADDOCK2.2 HADDOCK2.4 DISVIS CPORT POWERFIT PRODIGY SPOTON PROABC2 WHISCY PDBTOOLS FANDAS GENTBL ARCTIC3D EOSC marketplace

Announcing the EMBO Practical Course on Integrative Modelling of Biomolecular Interactions @ Izmir, Turkey, September 17-22, 2023

We would like your feedback on HADDOCK and your input on which features you would like to see implemented
User-Driven development survey

WELCOME TO THE UTRECHT BIOMOLECULAR INTERACTION WEB PORTAL

The Utrecht Biomolecular Interactions software portal provides access to software tools developed in the Computational Structural Biology group / NMR Research Group of Utrecht University with a focus on the characterization of biomolecular interactions. Please note that this site is in active development.

Research

Research within the computational structural biology group focuses on the development of reliable bioinformatic and computational approaches to predict, model and dissect biomolecular interactions at atomic level.

For this, bioinformatic data, structural information and available biochemical or biophysical experimental data are combined to drive the modelling process. By following a holistic approach integrating various experimental information sources with computational structural biology methods we aim at obtaining a comprehensive description of the structural and dynamic landscape of complex biomolecular machines, adding the structural dimension to interaction networks and opening the route to systematic and genome-wide studies of biomolecular interactions.

Services

HADDOCK

HADDOCK is an integrative platform for the modeling of biomolecular complexes. It supports a large variety of input data and can deal multi-component assemblies of proteins, peptides, small molecules and nucleic acids.

PRODIGY

PRODIGY predicts the binding affinity of protein-protein and protein-small molecules complexes and also allows to classify crystallographic interfaces as biological or not.

DISVIS

DISVIS quantifies and visualizes the information content of distance restraints between two molecules.

Go to service

POWERFIT

POWERFIT automatically fits atomic models into cryo-EM density maps.

Go to service

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 COMMONS (MIDI)

The Modelling institutional dynamics in historical commons (MIDI) project adopted an interdisciplinary perspective to contribute empirically-grounded and systematic knowledge of the mechanisms driving the process of institutional change. Its starting point is a dataset of European commons-management institutions composed by the Research Team Institutions for Collective Action in the Common Rules-project.

The goal of the project was to exploit our interdisciplinary competences to recode 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 (Netherlands), and Tine De Moor and René van Weeren from the *Institutions for Collective Action Research Team* (then at Utrecht University , The Netherlands). The project is funded by the Riksbankens Jubileumsfond and ran from 1 January 2018 until the end of 2019.

Sources: 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

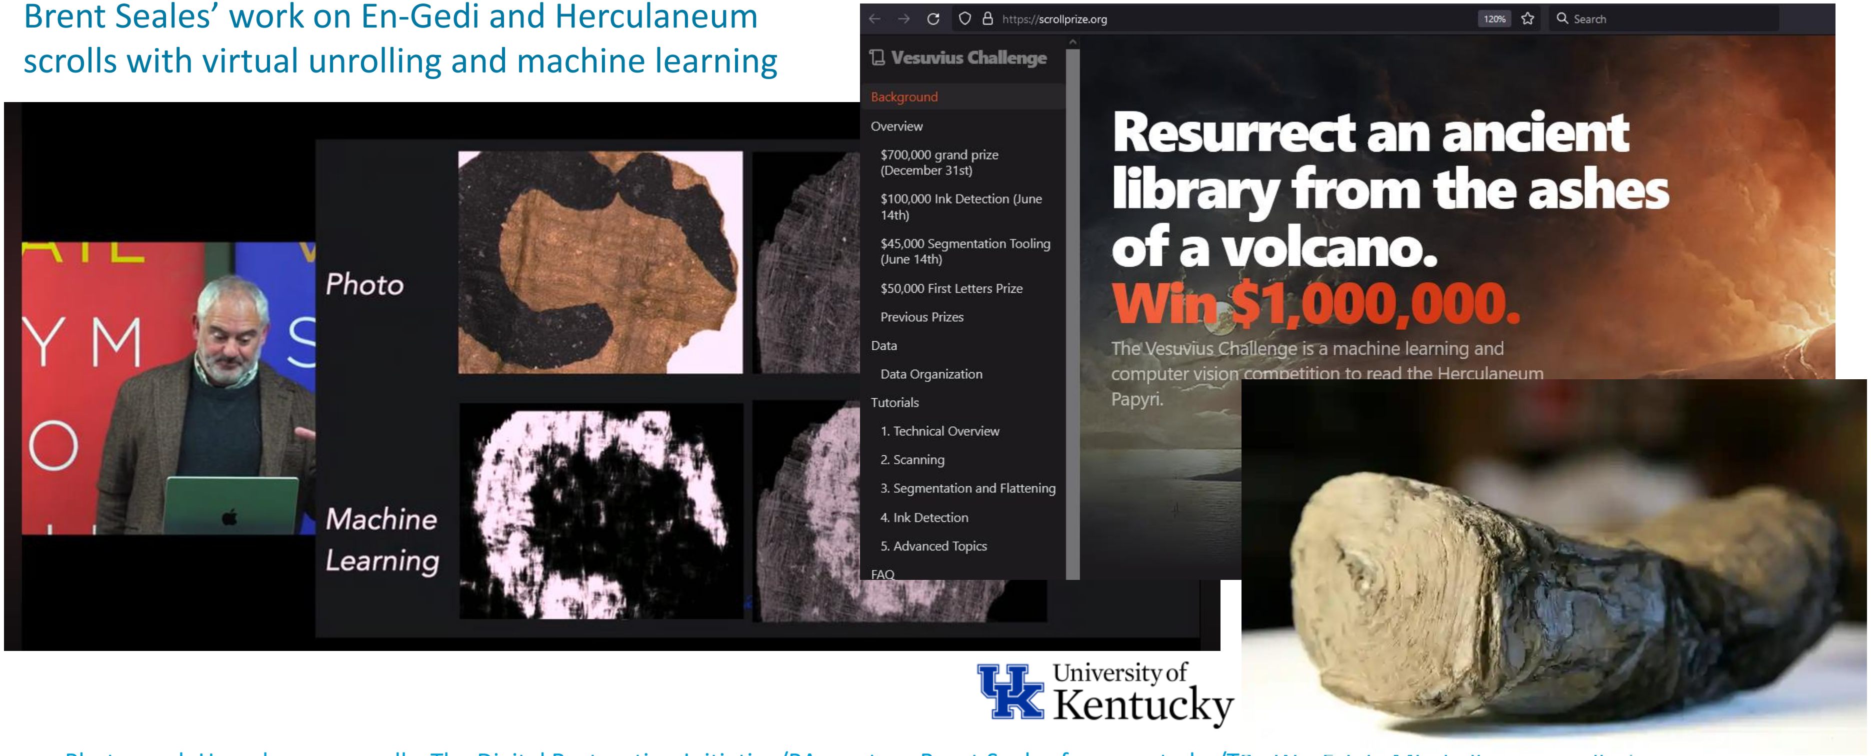

https://scrollprize.org

Vesuvius Challenge

- Background
- Overview
  - $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
  - Data Organization
- Tutorials
  - 1. Technical Overview
  - 2. Scanning
  - 3. Segmentation and Flattening
  - 4. Ink Detection
  - 5. Advanced Topics
- FAQ

Resurrect an ancient library from the ashes of a volcano.

Win $1,000,000.

The Vesuvius Challenge is a machine learning and computer vision competition to read the Herculaneum Papyri.

Photograph Herculaneum scrolls: The Digital Restoration Initiative/PA; capture Brent Seales from youtu.be/T0mWqsFrJpk; ML challenge: scrollprize.org

Computing is instrumentation just like a detector …

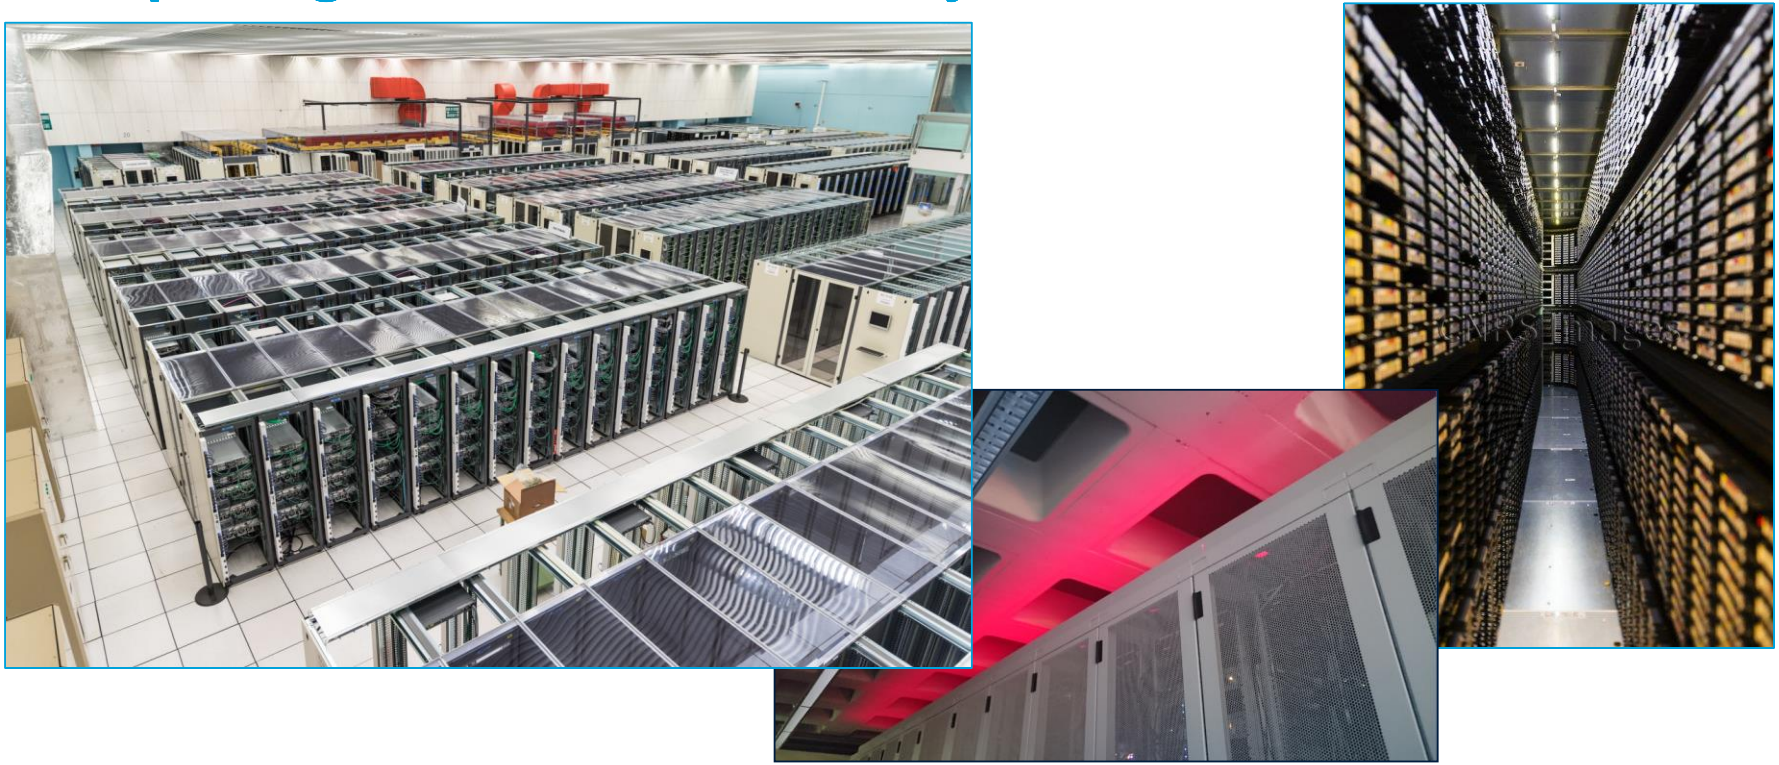

CERN Computing Centre B513, image: CERN, https://cds.cern.ch/record/2127440; tape library image CC-IN2P3 with LHC and LSST data; cabinets: Nikhef H234b

Infrastructures for research, built on computing services

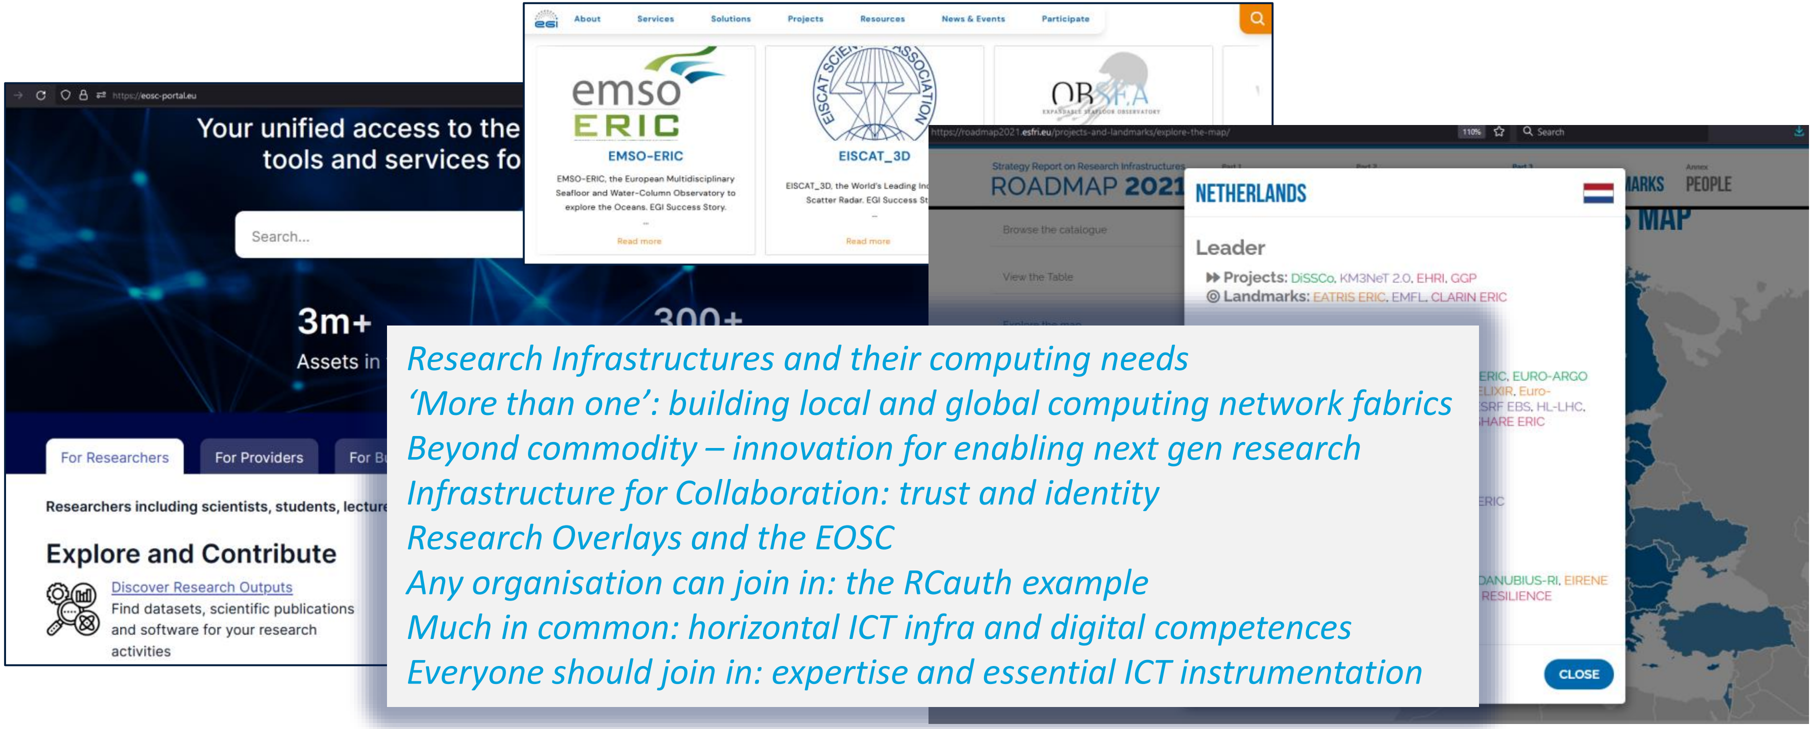

https://eosc-portal.eu

Your unified access to the tools and services fo

Search...

3m+ Assets in

300+

For Researchers | For Providers | For B

Researchers including scientists, students, lecture

Explore and Contribute

Discover Research Outputs
Find datasets, scientific publications and software for your research activities

About | Services | Solutions | Projects | Resources | News & Events | Participate

EMSO-ERIC
EMSO-ERIC, the European Multidisciplinary Seafloor and Water-Column Observatory to explore the Oceans. EGI Success Story.
Read more

EISCAT_3D
EISCAT_3D, the World's Leading Inc Scatter Radar. EGI Success St
Read more

https://roadmap2021.esfri.eu/projects-and-landmarks/explore-the-map/

Strategy Report on Research Infrastructures
ROADMAP 2021

Browse the catalogue
View the Table

NETHERLANDS

Leader

Projects: DiSSCo, KM3NeT 2.0, EHRI, GGP
Landmarks: EATRIS ERIC, EMFL, CLARIN ERIC

ERIC, EURO-ARGO
ELIXIR, Euro-
ESRF EBS, HL-LHC,
HARE ERIC

ERIC

DANUBIUS-RI, EIRENE
RESILIENCE

CLOSE

Research Infrastructures and their computing needs
‘More than one’: building local and global computing network fabrics
Beyond commodity – innovation for enabling next gen research
Infrastructure for Collaboration: trust and identity
Research Overlays and the EOSC
Any organisation can join in: the RCauth example
Much in common: horizontal ICT infra and digital competences
Everyone should join in: expertise and essential ICT instrumentation

Service catalogues from the EOSC Portal (eoc-portal.eu), EGI (egi.eu), and ESFRI (esfri.eu) Roadmap projects and Landmarks with Dutch involvement (2021 Roadmap)

Enabled by Computing, ‘more than one’

Facilities for the global infrastructure ecosystem
Collaborative workflows and services across multiple organisations
Networking, federated access and the ScienceDMZ concept

There is NO CLOUD, just other people’s computers

Image source: Free Software Foundation Europe - https://fsfe.org/

Local computational resources as a starting point

Many HTC applications like WLCG, SKA, or WeNMR are ‘conveniently parallel’

- **balanced features for node throughput**
  (CPU, storage, memory bandwidth, network)

- **single-socket** multicore systems are fine, typical: 64-128 cores per system
- **network**: 2x25Gbps
  (+ ‘out of band’ network for IPMI or Redfish)
- **memory**: 8 GiB/core (different from HPC)
- **local storage**: 4TB NVME PCIe Gen4 x4

+ space (physical + power) to add **GPU**

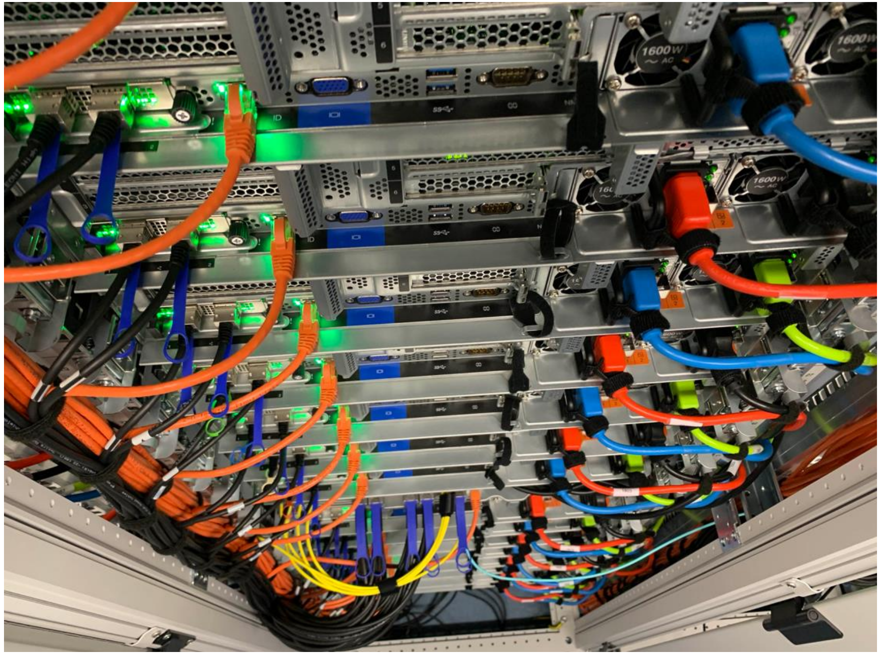

Image: Cluster ‘Lotenfeest’ at the Nikhef NDPF, acquired March 2020. Lenovo SR655 with AMD EPYC 7702P 64-Core single-socket

NDPF ‘WLCG and Dutch National Infra’ cluster

Running jobs:

period: March 2021 .. October 2022

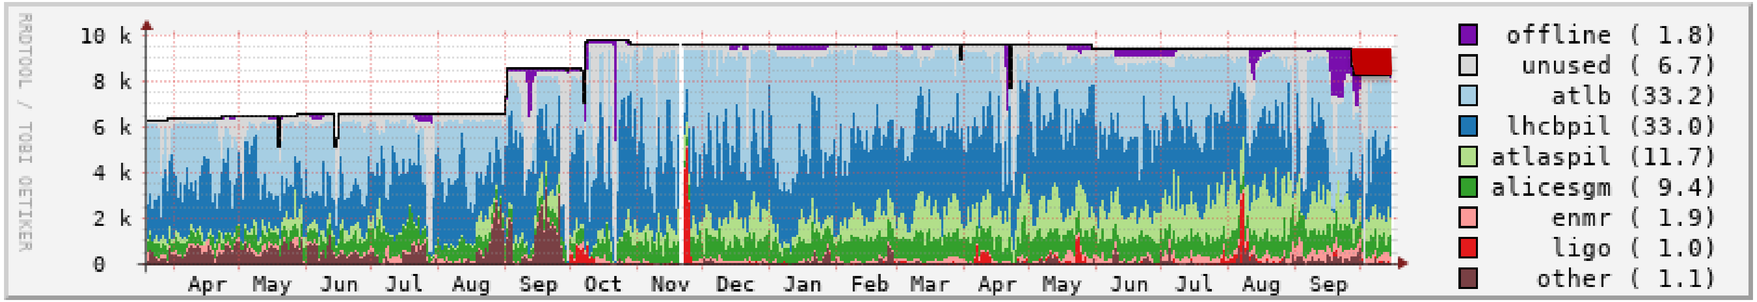

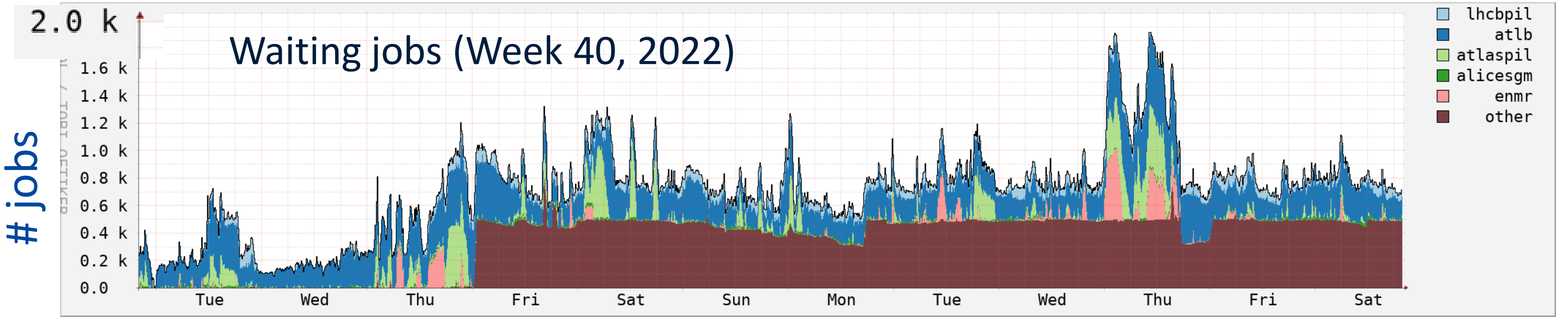

capacity move on Sept 27: nodes moved to LIGO-VIRGO specific cluster; Source: NDPF Statistics overview, https://www.nikhef.nl/pdp/doc/stats/
‘other’ waiting jobs are almost all for the Auger experiment - GRISview images: Jeff Templon for NDPF and STBC

WLCG NL-T1 and the Dutch National Infrastructure

HPC-HTC convergence with Snellius-LISA-GINA

Joint SURF & Nikhef collective service – part of EGI, WLCG and FuSE

hosts WLCG, but also LOFAR radio telescope data, and ~100 other projects

59 PByte near-line storage (tape), 42.5 PByte on-line (disk), 27.6 k cores (cpu)

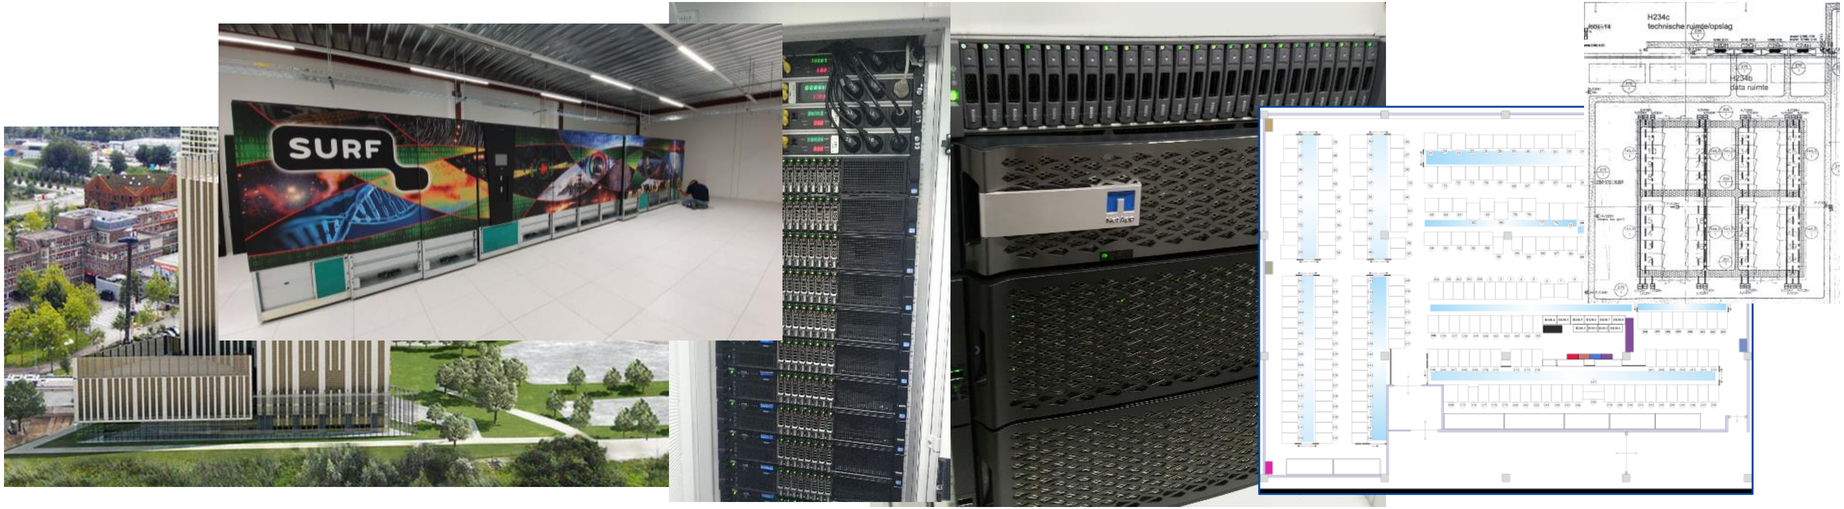

DNI and NL-T1 capacity from 2023 DNI NWO, LOFAR, and WLCG; see https://www.surf.nl/onderzoek-ict/toegang-tot-rekendiensten-aanvragen ; fuse-infra.nl
SURF tape total: ~80 PByte by end 2022; image library at Schiphol Rijk from Sara Ramezani; NikhefHousing: https://www.nikhef.nl/housing/datacenter/floorplan/

Dutch National e-Infrastructure: High Throughput GINA

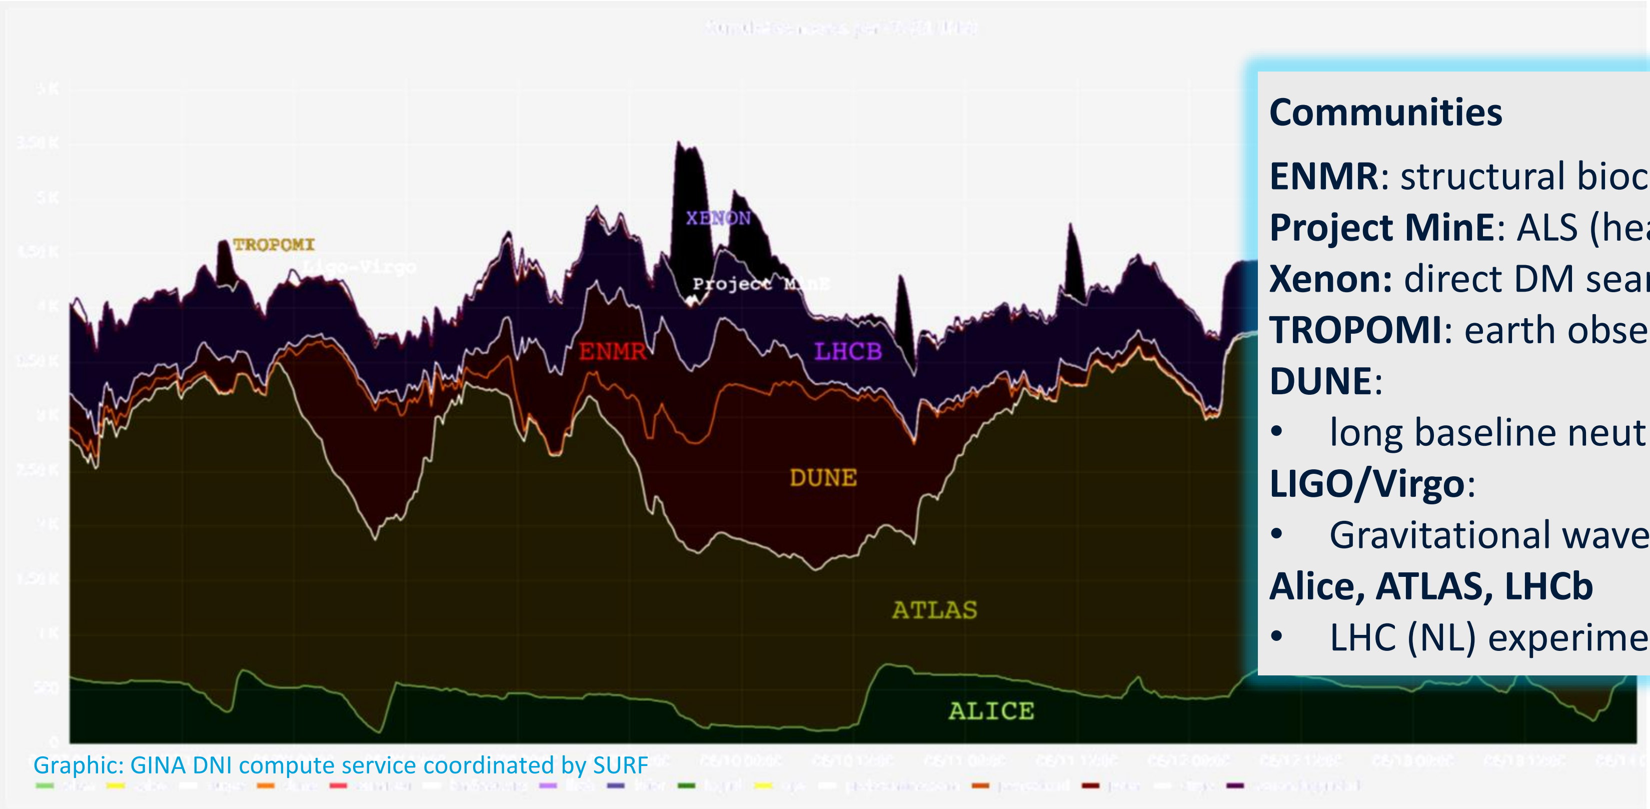

Graphic: GINA DNI compute service coordinated by SURF

Communities

ENMR: structural biochemistry
Project MinE: ALS (health)
Xenon: direct DM searches
TROPOMI: earth observation
DUNE:

- long baseline neutrinos

LIGO/Virgo:

- Gravitational waves

Alice, ATLAS, LHCb

- LHC (NL) experiments

More than one: the worldwide LHC Computing Grid

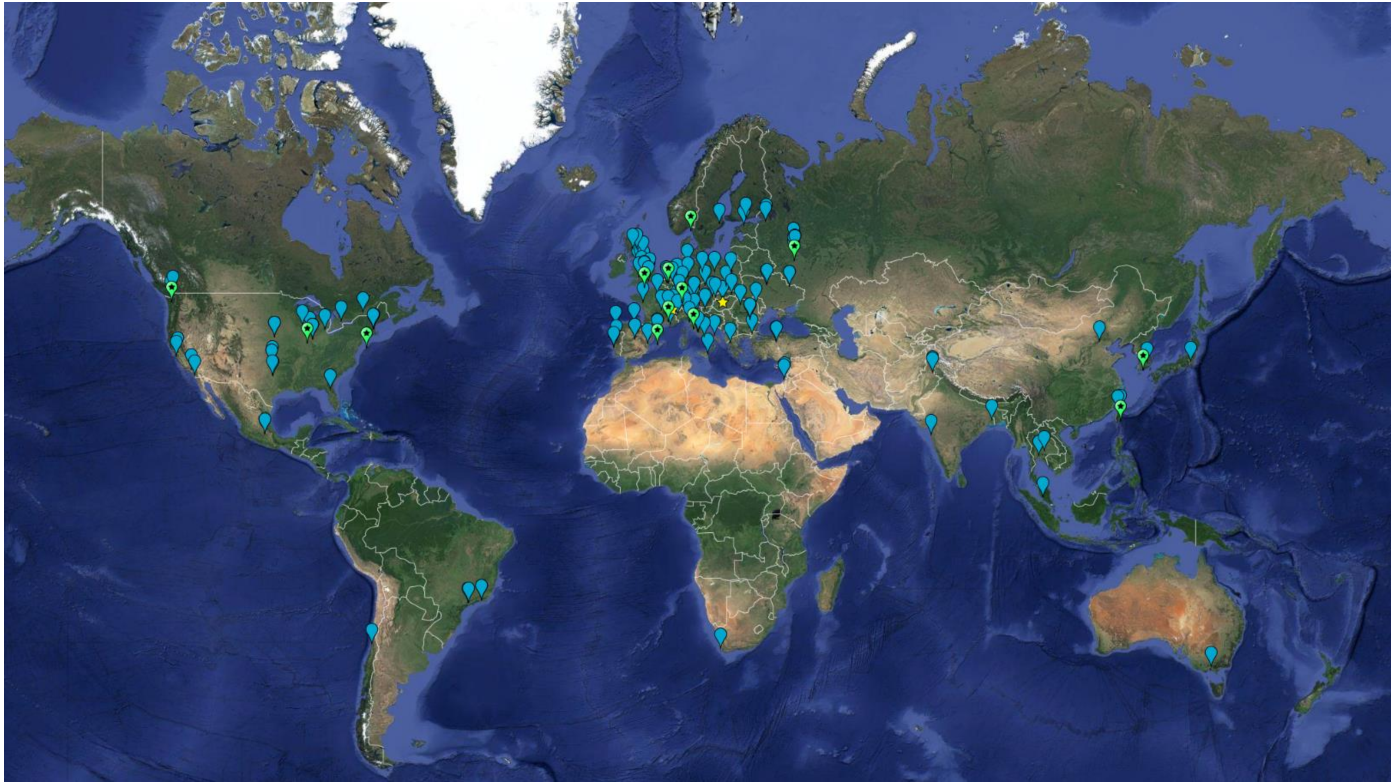

~ 1.4 million CPU cores
~ 1500 Petabyte
disk + archival

170+ institutes
40+ countries
13 ‘Tier-1 sites’
NL-T1:
SURF & Nikhef

built on e-Infrastructures
EGI
PRACE-RI
EuroHPC
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; ACCESS (XSEDE): https://access-ci.org/, for the NL-T1 and FuSE: fuse-infra.nl, https://www.surf.nl/en/research-it

Conveniently parallel: a global infrastructure for research

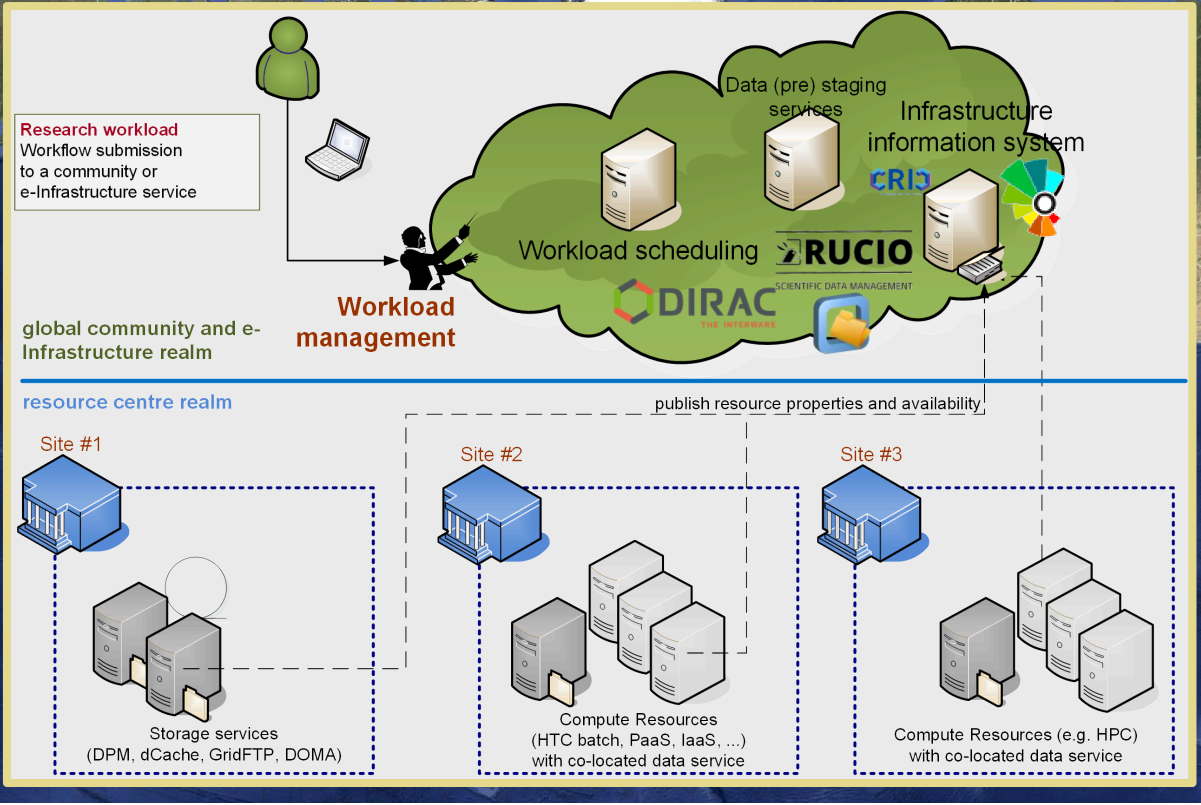
Research workload
Workflow submission
to a community or
e-Infrastructure service
global community and e-Infrastructure realm
Workload management
Workload scheduling
Data (pre) staging services
Infrastructure information system
CRIC
RUCIO
SCIENTIFIC DATA MANAGEMENT
DIRAC
THE INTERWARE
resource centre realm
publish resource properties and availability
Site #1
Site #2
Site #3
Storage services
(DPM, dCache, GridFTP, DOMA)
Compute Resources
(HTC batch, PaaS, IaaS, ...)
with co-located data service
Compute Resources (e.g. HPC)
with co-located data service

shared multi-community infrastructure

Already EGI e-infra has >250 communities just doing HTC

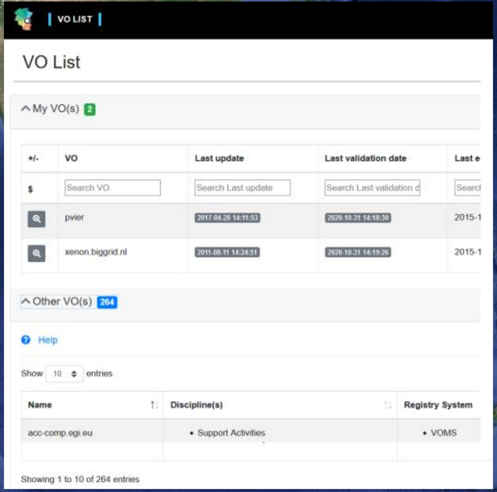

VO List

^ My VO(s) 2

+/-	VO	Last update	Last validation date	Last e
$	Search VO	Search Last update	Search Last validation d	Searcl
	pvier	2017-04-28 14:11:53	2020-10-31 14:18:30	2015-1
	xenon.biggrid.nl	2011-08-11 14:24:51	2020-10-31 14:19:26	2015-1

^ Other VO(s) 264

Help

Show 10 entries

Name	Discipline(s)	Registry System
acc-comp.egi.eu	• Support Activities	• VOMS

Showing 1 to 10 of 264 entries

Right-hand graphic: EGI operations portal, https://operations-portal.egi.eu/vo/ - project logos in workflow image for illustration only, other services exist

Global distribution of computing and data placement

WLCG and EGI Advanced Computing for Research

LHCOPN – traffic levels for T0T1 data transfer

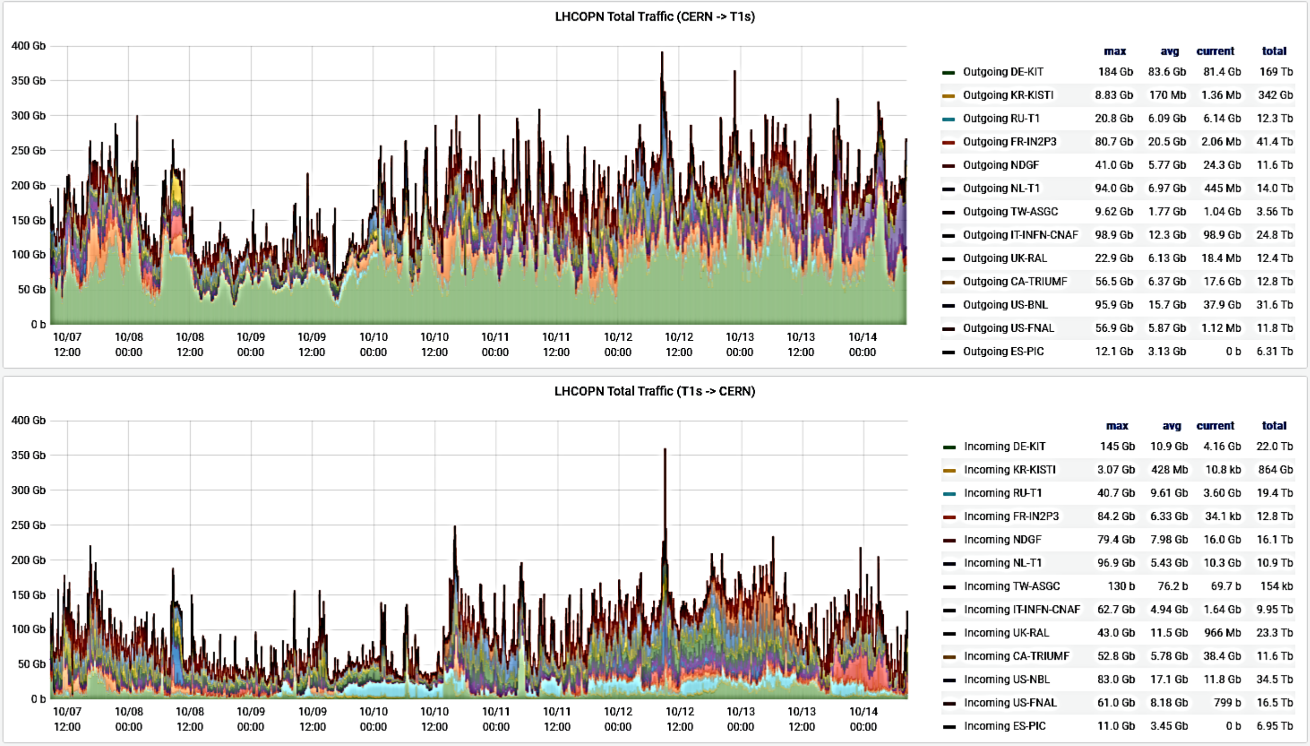

LHCOPN Total Traffic (CERN -> T1s)

	max	avg	current	total
▬ Outgoing DE-KIT	184 Gb	83.6 Gb	81.4 Gb	169 Tb
▬ Outgoing KR-KISTI	8.83 Gb	170 Mb	1.36 Mb	342 Gb
▬ Outgoing RU-T1	20.8 Gb	6.09 Gb	6.14 Gb	12.3 Tb
▬ Outgoing FR-IN2P3	80.7 Gb	20.5 Gb	2.06 Mb	41.4 Tb
▬ Outgoing NDGF	41.0 Gb	5.77 Gb	24.3 Gb	11.6 Tb
▬ Outgoing NL-T1	94.0 Gb	6.97 Gb	445 Mb	14.0 Tb
▬ Outgoing TW-ASGC	9.62 Gb	1.77 Gb	1.04 Gb	3.56 Tb
▬ Outgoing IT-INFN-CNAF	98.9 Gb	12.3 Gb	98.9 Gb	24.8 Tb
▬ Outgoing UK-RAL	22.9 Gb	6.13 Gb	18.4 Mb	12.4 Tb
▬ Outgoing CA-TRIUMF	56.5 Gb	6.37 Gb	17.6 Gb	12.8 Tb
▬ Outgoing US-BNL	95.9 Gb	15.7 Gb	37.9 Gb	31.6 Tb
▬ Outgoing US-FNAL	56.9 Gb	5.87 Gb	1.12 Mb	11.8 Tb
▬ Outgoing ES-PIC	12.1 Gb	3.13 Gb	0 b	6.31 Tb

LHCOPN Total Traffic (T1s -> CERN)

	max	avg	current	total
▬ Incoming DE-KIT	145 Gb	10.9 Gb	4.16 Gb	22.0 Tb
▬ Incoming KR-KISTI	3.07 Gb	428 Mb	10.8 kb	864 Gb
▬ Incoming RU-T1	40.7 Gb	9.61 Gb	3.60 Gb	19.4 Tb
▬ Incoming FR-IN2P3	84.2 Gb	6.33 Gb	34.1 kb	12.8 Tb
▬ Incoming NDGF	79.4 Gb	7.98 Gb	16.0 Gb	16.1 Tb
▬ Incoming NL-T1	96.9 Gb	5.43 Gb	10.3 Gb	10.9 Tb
▬ Incoming TW-ASGC	130 b	76.2 b	69.7 b	154 kb
▬ Incoming IT-INFN-CNAF	62.7 Gb	4.94 Gb	1.64 Gb	9.95 Tb
▬ Incoming UK-RAL	43.0 Gb	11.5 Gb	966 Mb	23.3 Tb
▬ Incoming CA-TRIUMF	52.8 Gb	5.78 Gb	38.4 Gb	11.6 Tb
▬ Incoming US-NBL	83.0 Gb	17.1 Gb	11.8 Gb	34.5 Tb
▬ Incoming US-FNAL	61.0 Gb	8.18 Gb	799 b	16.5 Tb
▬ Incoming ES-PIC	11.0 Gb	3.45 Gb	0 b	6.95 Tb

CERN OpenMonIT LHCOPN, period Oct 7 .. Oct 14 2022, from https://monit-grafana-open.cern.ch/d/HreVOyc7z/all-lhcopn-traffic

LHCOPN

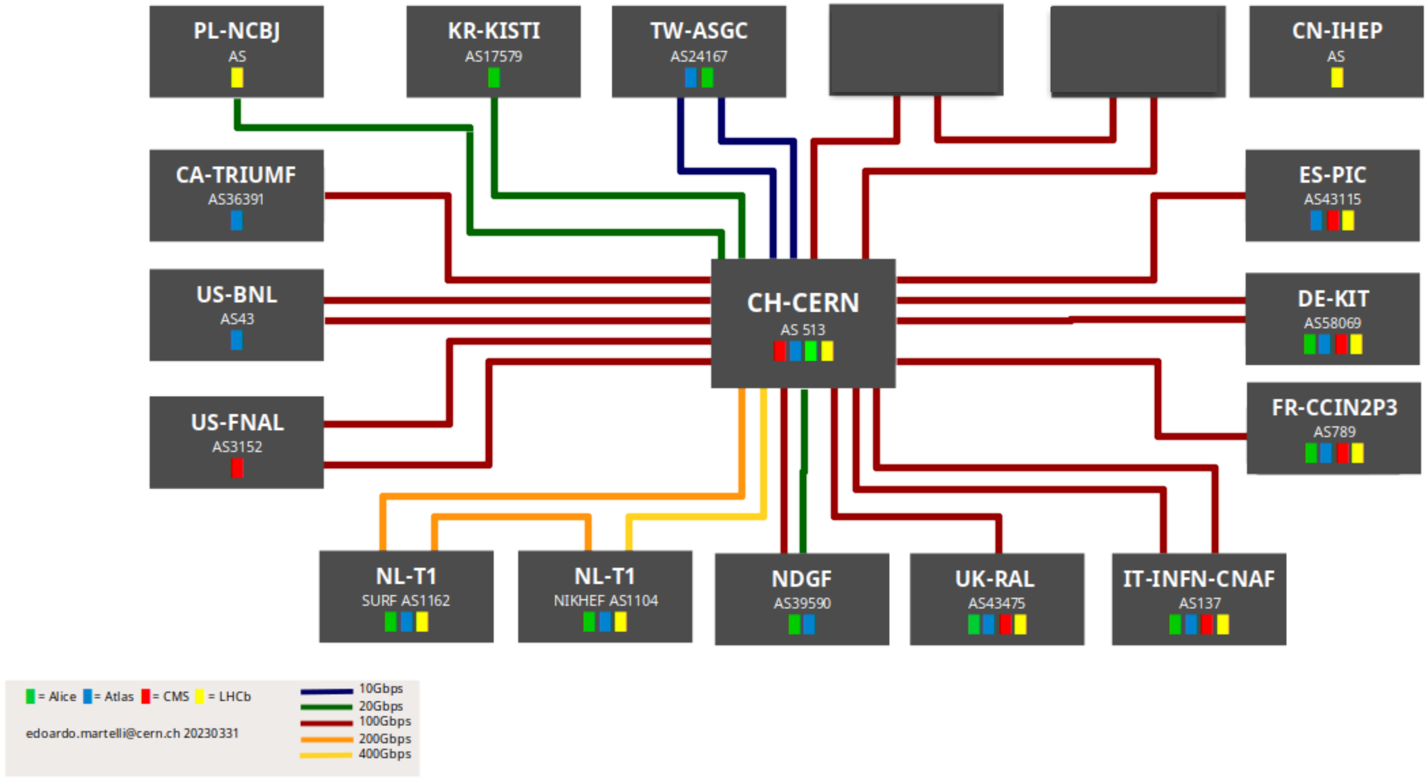

PL-NCBJ
AS
🟨

KR-KISTI
AS17579
🟩

TW-ASGC
AS24167
🟦🟩

CN-IHEP
AS
🟨

CA-TRIUMF
AS36391
🟦

ES-PIC
AS43115
🟦🟥🟨

US-BNL
AS43
🟦

CH-CERN
AS 513
🟥🟦🟩🟨

DE-KIT
AS58069
🟩🟦🟥🟨

US-FNAL
AS3152
🟥

FR-CCIN2P3
AS789
🟩🟦🟥🟨

NL-T1
SURF AS1162
🟩🟦🟨

NL-T1
NIKHEF AS1104
🟩🟦🟨

NDGF
AS39590
🟩🟦

UK-RAL
AS43475
🟩🟦🟥🟨

IT-INFN-CNAF
AS137
🟩🟦🟥🟨

🟩 = Alice 🟦 = Atlas 🟥 = CMS 🟨 = LHCb

edoardo.martelli@cern.ch 20230331

10Gbps
20Gbps
100Gbps
200Gbps
400Gbps

Edoardo Martelli, CERN (https://twiki.cern.ch/twiki/bin/view/LHCOPN/OverallNetworkMaps)

LHCone

WLCG

T1-T1 & T1-T2-T3

+ collaborations

DUNE, Belle-II

Pierre-Auger

NOvA, XENON

JUNO

Quite elementary expectations

IPv6,

jumboframes,

symmetric routing

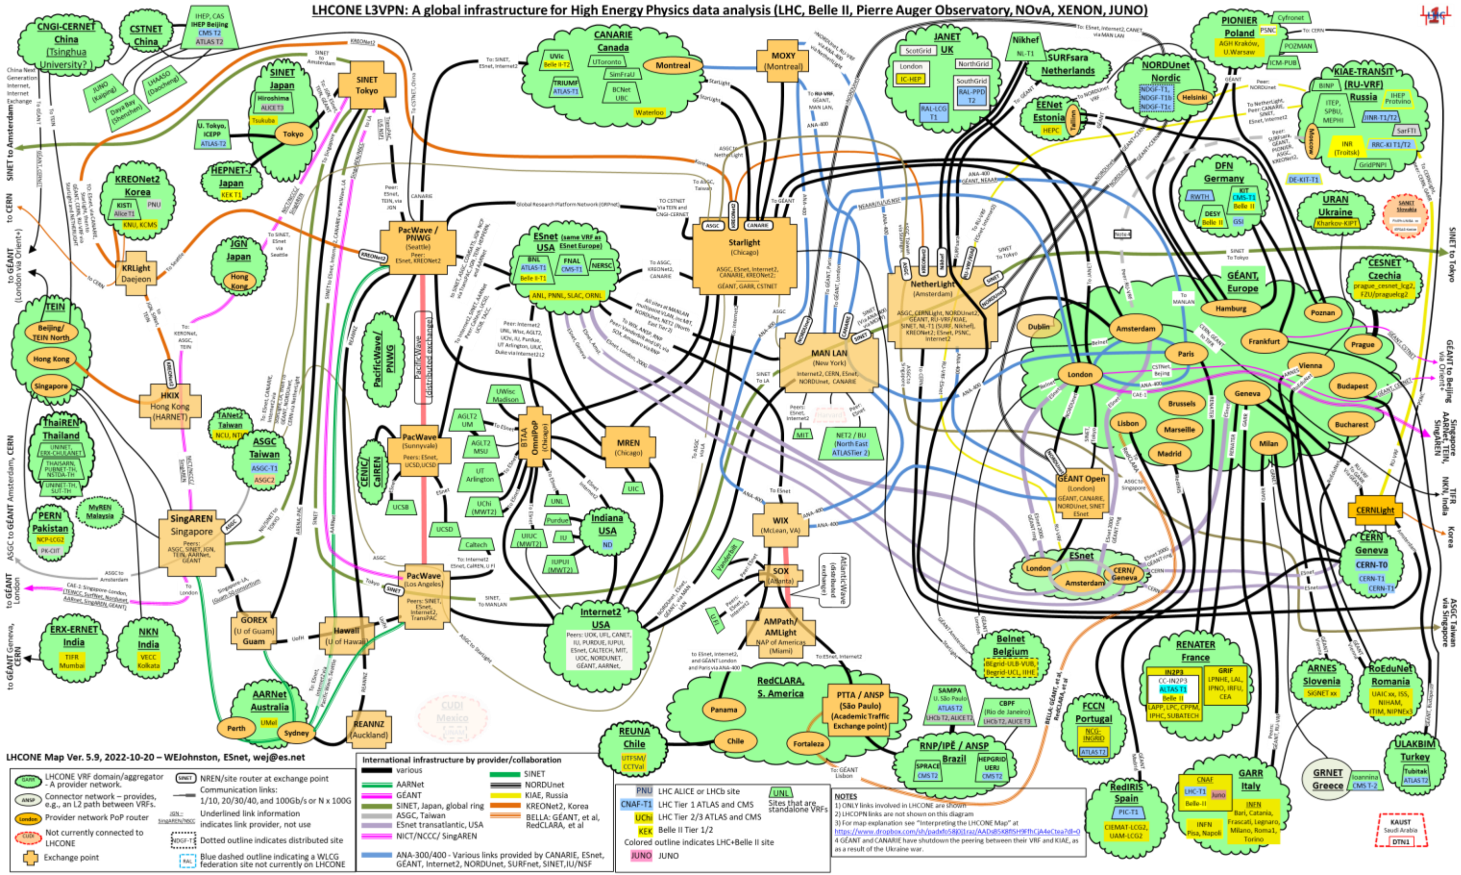

LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger Observatory, NOvA, XENON, JUNO)

LHCONE Map Ver. 5.9, 2022-10-20 – WEJohnston, ESnet, wej@es.net

NOTES
1) ONLY links involved in LHCONE are shown
2) LHCOPN links are not shown on this diagram
3) For map explanation see "Interpreting the LHCONE Map" at https://www.dropbox.com/sh/padxfo58j0j1raz/AADsB5K8flSH9FfhCjA4eCtea?dl=0
4 GÉANT and CANARIE have shutdown the peering between their VRF and KIAE, as as a result of the Ukraine war.

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

‘ScienceDMZ’

Predicable 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**

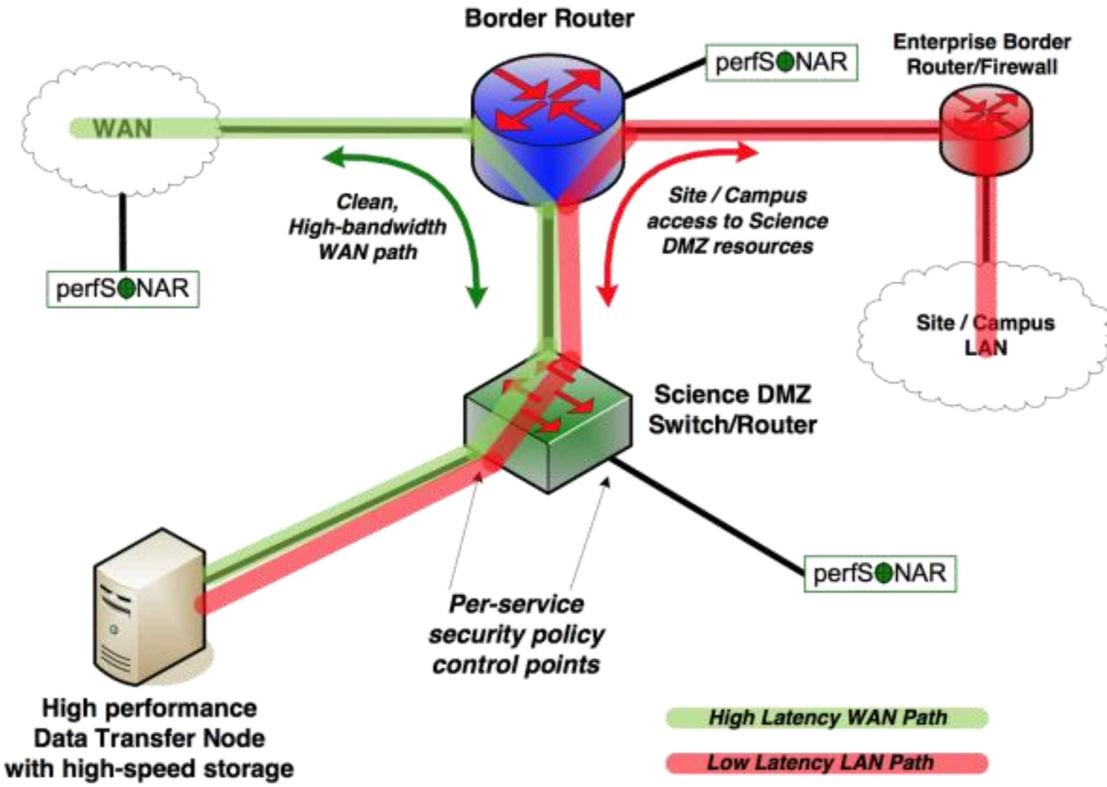
Border Router
perfSONAR
Enterprise Border Router/Firewall
WAN
perfSONAR
Clean, High-bandwidth WAN path
Site / Campus access to Science DMZ resources
Site / Campus LAN
Science DMZ Switch/Router
perfSONAR
Per-service security policy control points
High performance Data Transfer Node with high-speed storage
High Latency WAN Path
Low Latency LAN Path

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

Can hardly be said better than Eli Dart did at TNC23

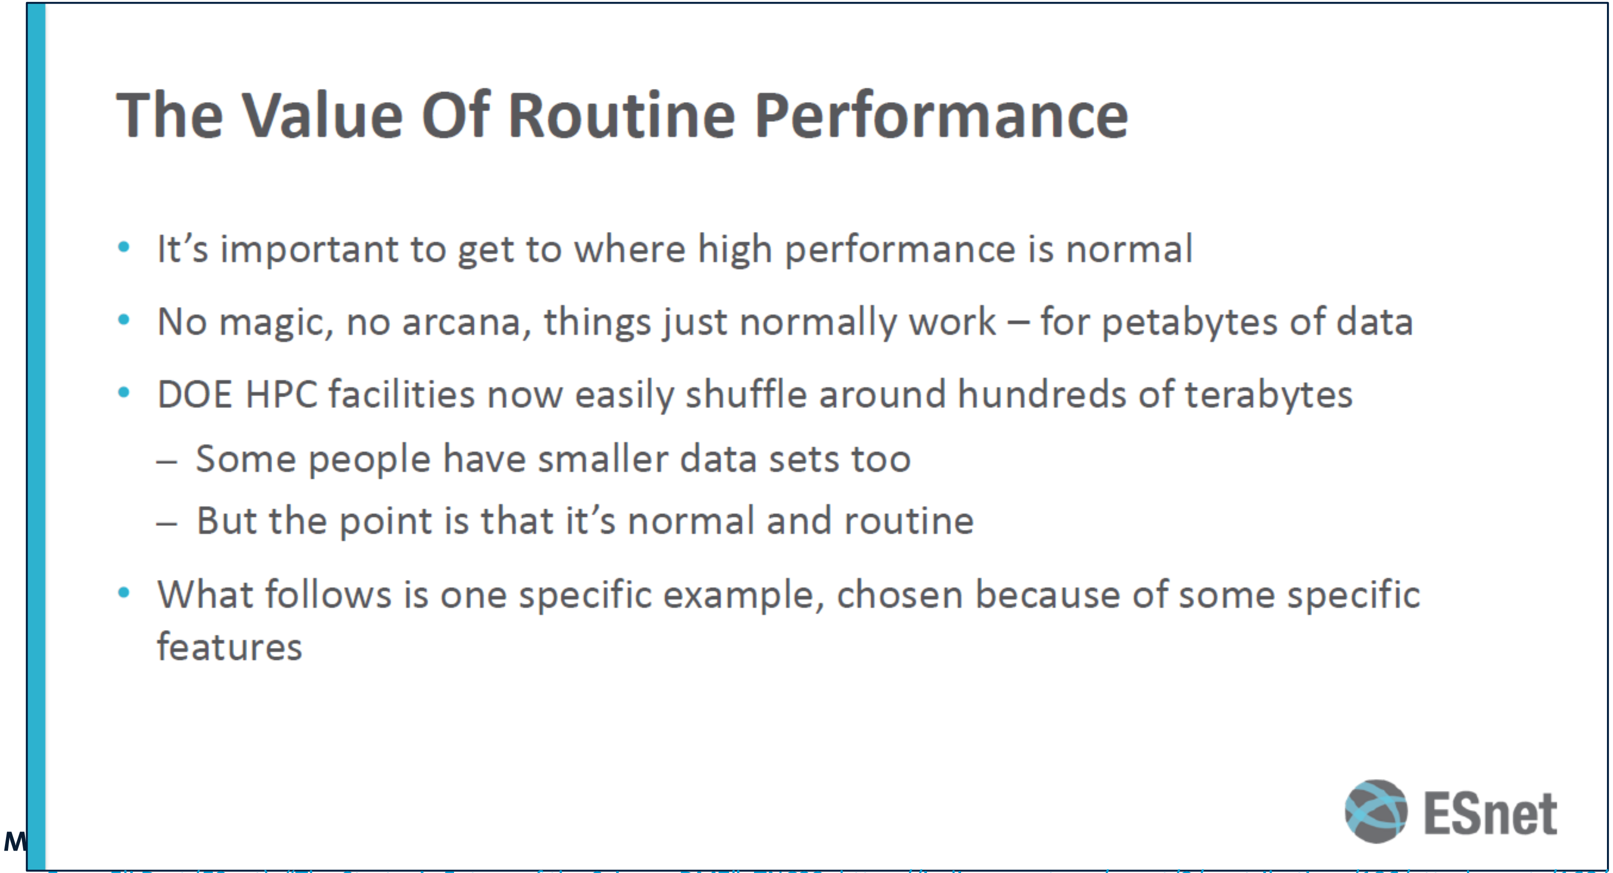

The Value Of Routine Performance

- 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

From Eli Dart (ESnet), “The Strategic Future of the Science DMZ”, TNC23, https://indico.geant.org/event/2/contributions/186/attachments/168/

The network is there to connect – ‘AS1104’ as an example

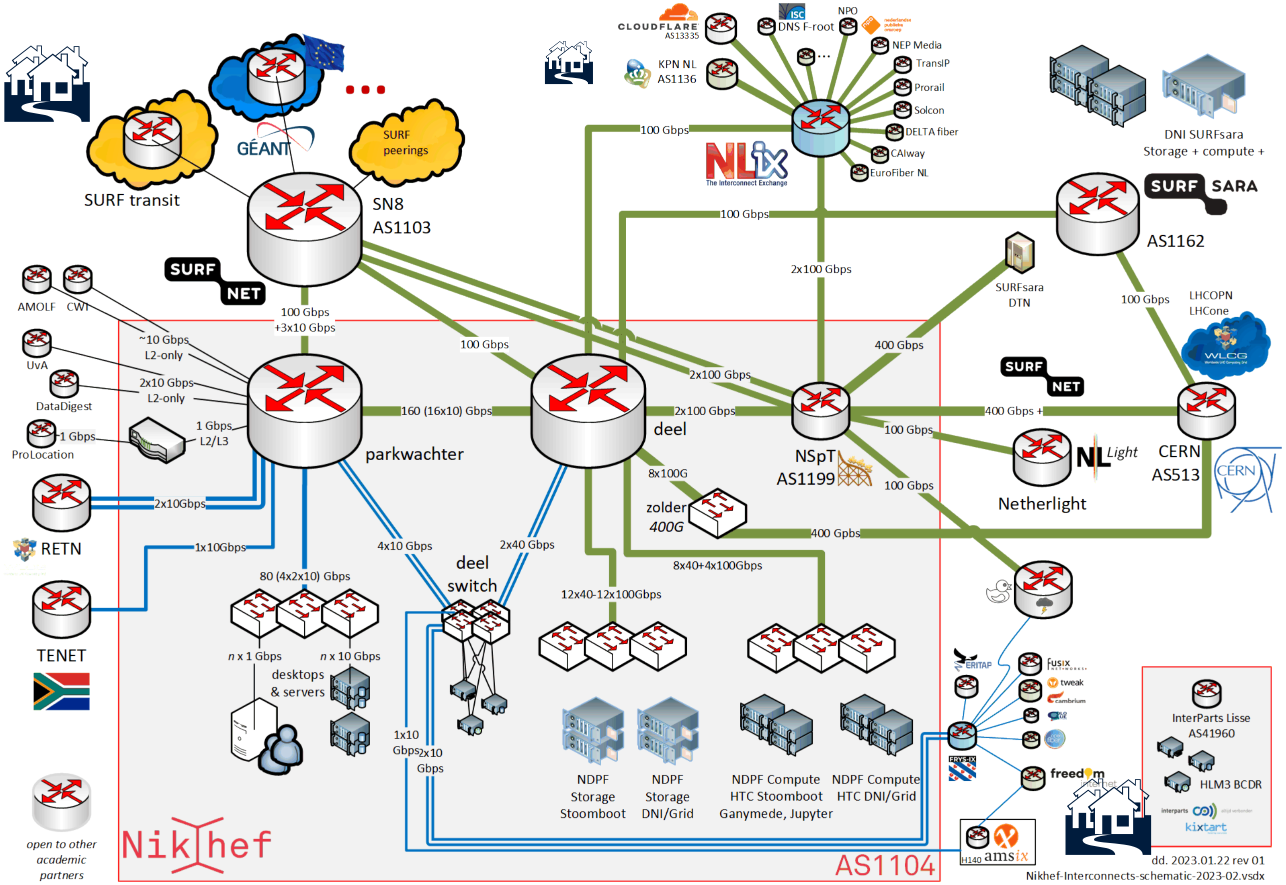

Make a guess … the Nikhef Institute (AMS) total network admin effort including desktop, wifi, servers, cloud, peering, and procurement needs _ FTE? – mind: doing networking right is not overly complex …

And for a research mission …

… you want
a **science network**
with a ‘back-office enclave’

‘open-core’ research network model implements enclave structure *and* protects against overload by having *no stateful components in the network path*

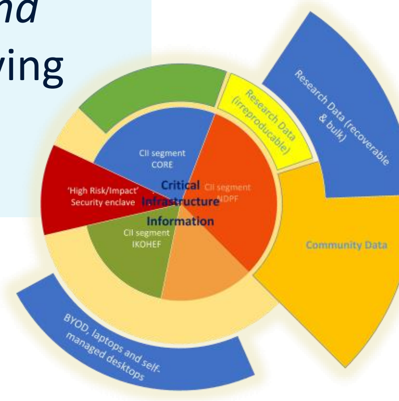

Research Data (irreproducible)
Research Data (recoverable & bulk)
CII segment CORE
'High Risk/Impact' Security enclave
Critical Infrastructure Information
CII segment NDPF
CII segment IKOHEF
Community Data
BYOD, laptops and self-managed desktops

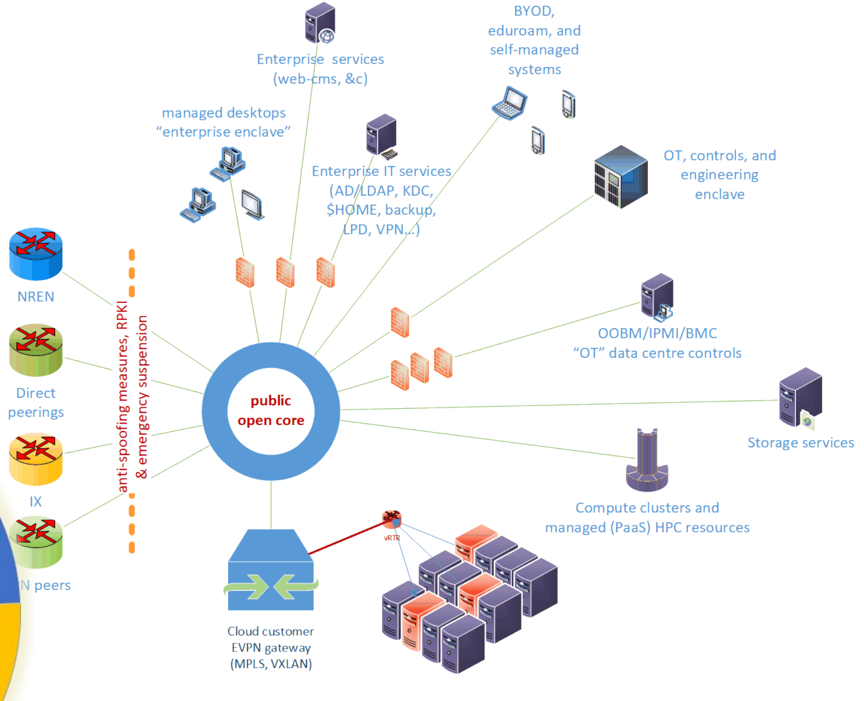

Enterprise services (web-cms, &c)
BYOD, eduroam, and self-managed systems
managed desktops “enterprise enclave”
Enterprise IT services (AD/LDAP, KDC, $HOME, backup, LPD, VPN...)
OT, controls, and engineering enclave
NREN
Direct peerings
IX
N peers
anti-spoofing measures, RPKI & emergency suspension
public open core
OOBM/IPMI/BMC “OT” data centre controls
Storage services
Compute clusters and managed (PaaS) HPC resources
vRTR
Cloud customer EVPN gateway (MPLS, VXLAN)

And sometimes 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
...
```

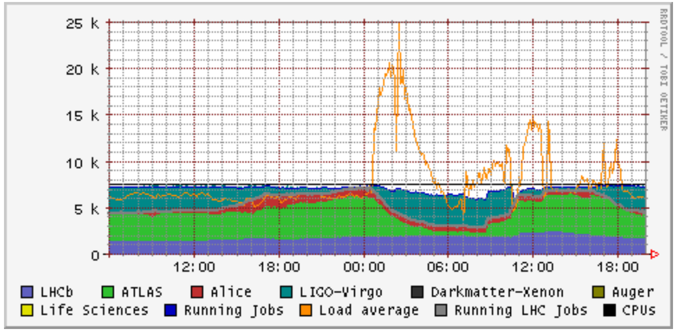

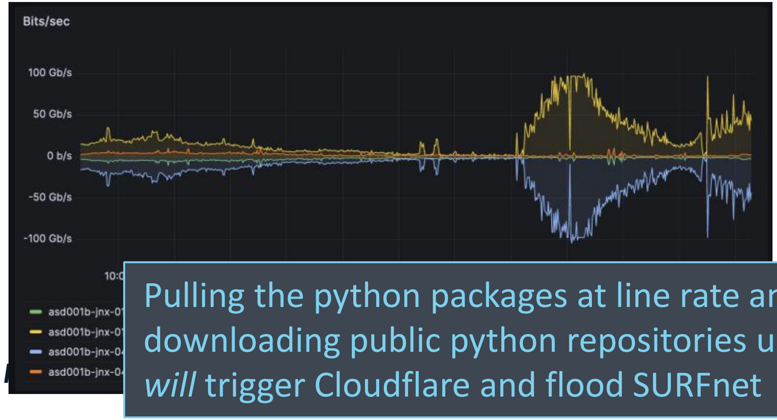

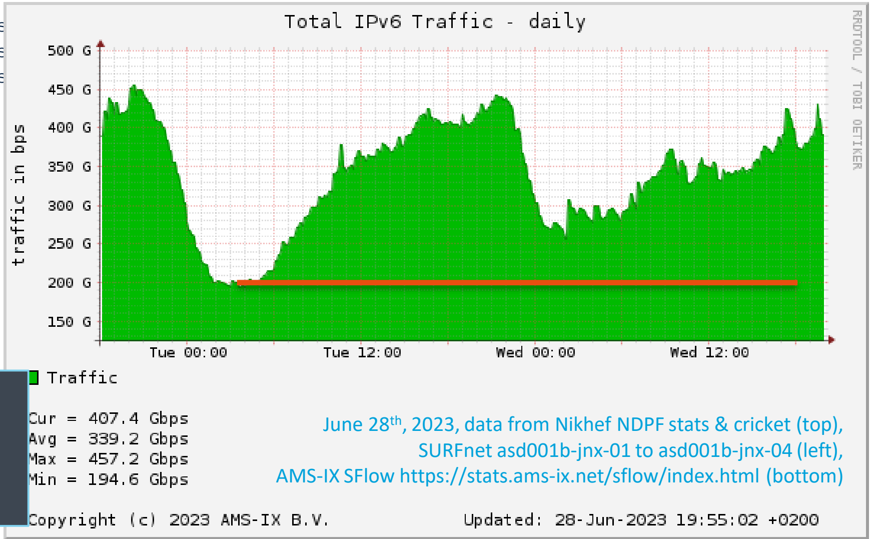

Pulling the python packages at line rate and downloading public python repositories ultimately *will* trigger Cloudflare and flood SURFnet

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

Science DMZ where ‘zero trust’ labelling comes in

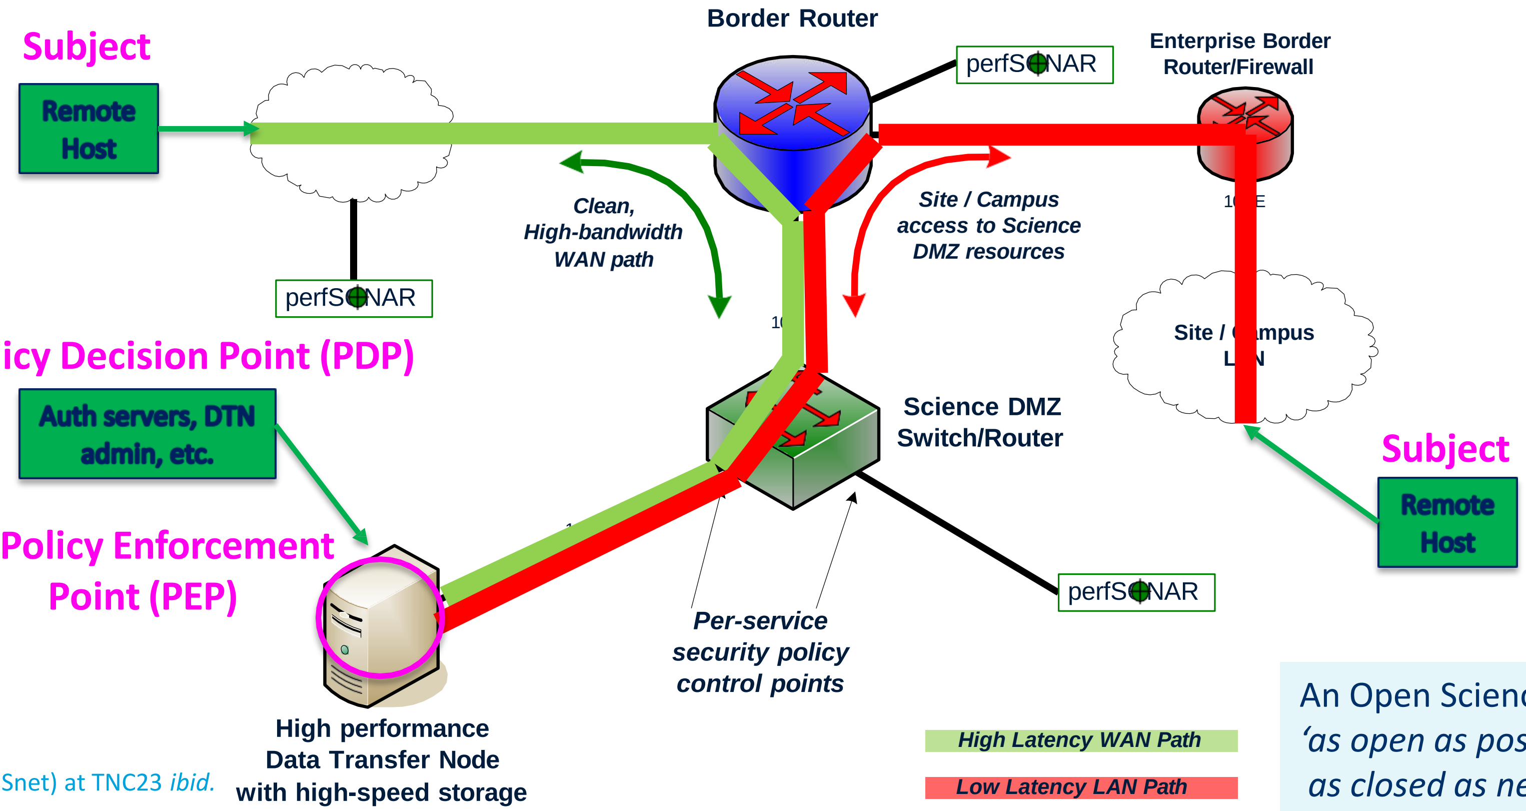

Subject

Remote Host

Border Router

perfSONAR

Enterprise Border Router/Firewall

Clean, High-bandwidth WAN path

Site / Campus access to Science DMZ resources

perfSONAR

Policy Decision Point (PDP)

Auth servers, DTN admin, etc.

Site / Campus LAN

Science DMZ Switch/Router

Subject

Remote Host

Policy Enforcement Point (PEP)

perfSONAR

Per-service security policy control points

High performance Data Transfer Node with high-speed storage

🟩 *High Latency WAN Path*

🟥 *Low Latency LAN Path*

From Eli Dart (ESnet) at TNC23 *ibid.*

An Open Science/FAIR net:
‘as open as possible,
as closed as necessary’

And ‘open’ does not mean ‘insecure’ – the WLCG SOC model

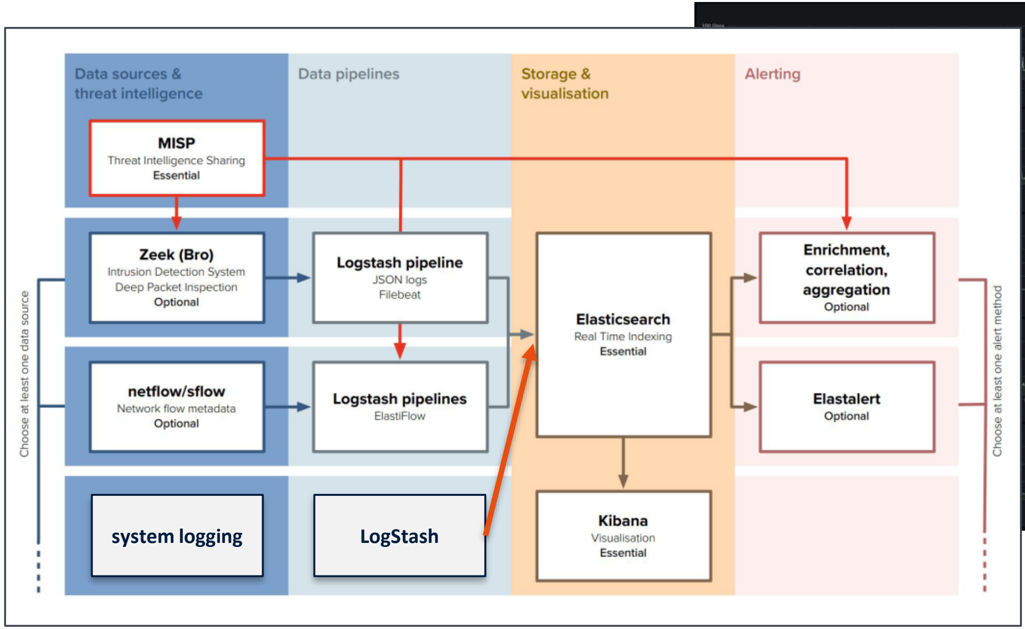

Data sources & threat intelligence

Data pipelines

Storage & visualisation

Alerting

MISP
Threat Intelligence Sharing
Essential

Zeek (Bro)
Intrusion Detection System
Deep Packet Inspection
Optional

netflow/sflow
Network flow metadata
Optional

system logging

Logstash pipeline
JSON logs
Filebeat

Logstash pipelines
ElastiFlow

LogStash

Elasticsearch
Real Time Indexing
Essential

Kibana
Visualisation
Essential

Enrichment, correlation, aggregation
Optional

Elastalert
Optional

Choose at least one data source

Choose at least one alert method

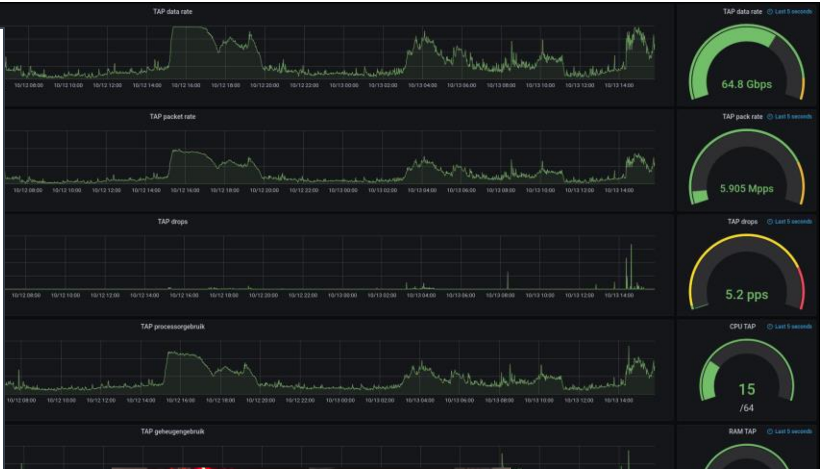

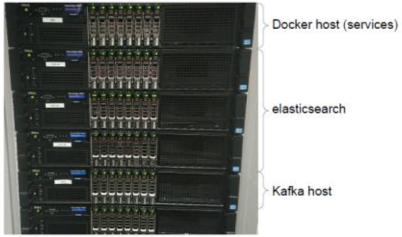

650 GByte/day ingest; 100Gbps+ monitoring through optical taps and mirroring; MISP intel from CERN, SURF, and private intel sources

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

Beyond today's commodity

From CS Research and SLICES-RI to Infrastructure Innovation

SURF Big Data Science Innovation, SURFNet 9, Snellius evolution

Data networks for the HL-HLC, SKA, and beyond

National resilience testing and innovation partnerships

Exercising the network – for sensor data or ‘rare’ HEP events

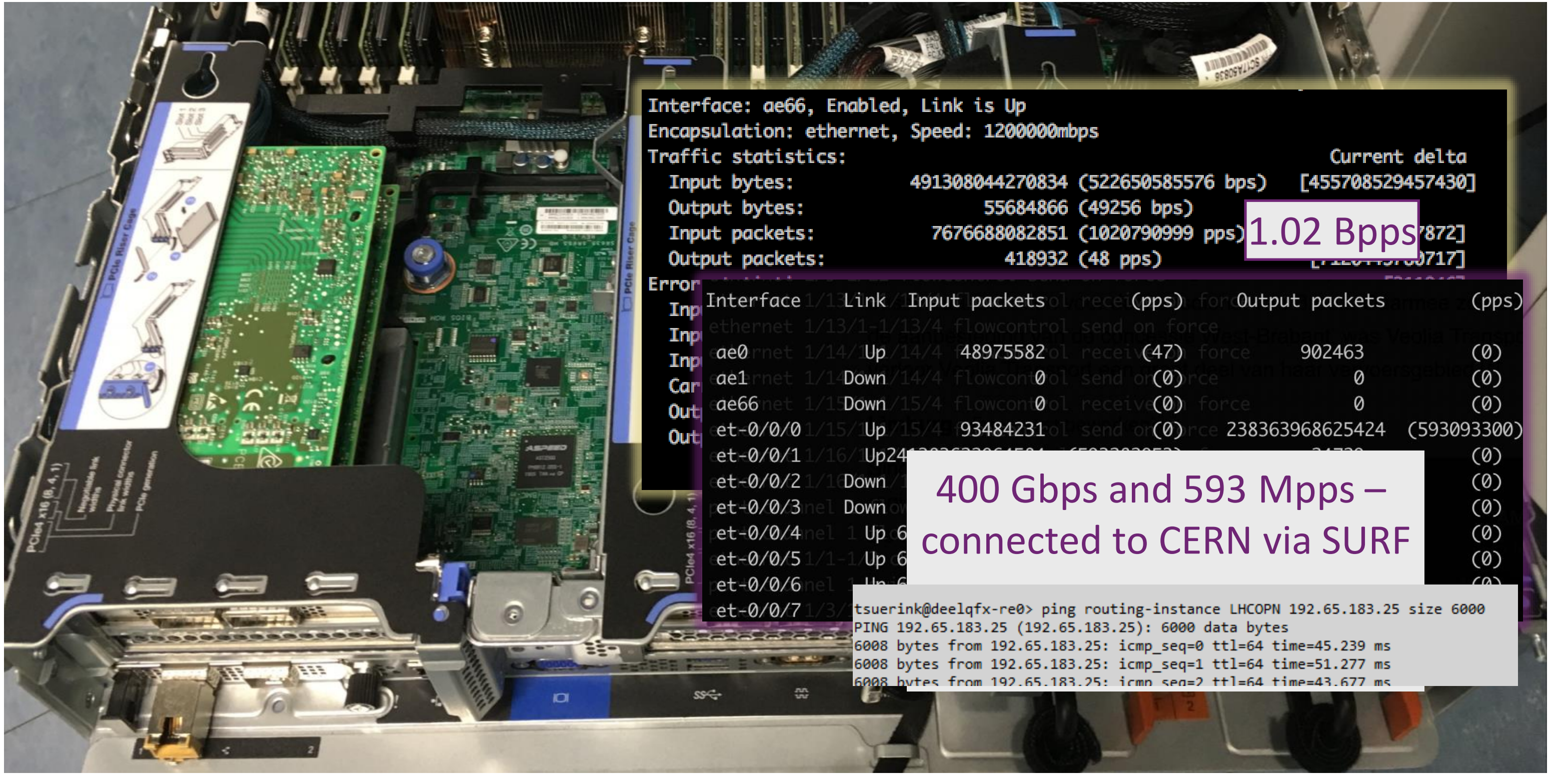

```
Interface: ae66, Enabled, Link is Up
Encapsulation: ethernet, Speed: 1200000mbps
Traffic statistics:                                       Current delta
  Input bytes:          491308044270834 (522650585576 bps)  [455708529457430]
  Output bytes:                55684866 (49256 bps)
  Input packets:          7676688082851 (1020790999 pps)               872]
  Output packets:                418932 (48 pps)                       717]
Error
  Inp
  Inp
  Inp
  Car
  Out
  Out
```

1.02 Bpps

```
Interface   Link  Input packets       (pps)   Output packets       (pps)
 ae0        Up        48975582        (47)        902463            (0)
 ae1        Down             0         (0)             0            (0)
 ae66       Down             0         (0)             0            (0)
 et-0/0/0   Up        93484231         (0)   23836396862542 4 (593093300)
 et-0/0/1   Up24                                                    (0)
 et-0/0/2   Down                                                    (0)
 et-0/0/3   Down                                                    (0)
 et-0/0/4   Up 6                                                    (0)
 et-0/0/5   Up 6                                                    (0)
 et-0/0/6
 et-0/0/7
```

400 Gbps and 593 Mpps – connected to CERN via SURF

```
tsuerink@deelqfx-re0> ping routing-instance LHCOPN 192.65.183.25 size 6000
PING 192.65.183.25 (192.65.183.25): 6000 data bytes
6008 bytes from 192.65.183.25: icmp_seq=0 ttl=64 time=45.239 ms
6008 bytes from 192.65.183.25: icmp_seq=1 ttl=64 time=51.277 ms
6008 bytes from 192.65.183.25: icmp_seq=2 ttl=64 time=43.677 ms
```

Image: ballenbak.nikhef.nl, Tristan Suerink

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: 400 → 800G AMS-GVA ☺

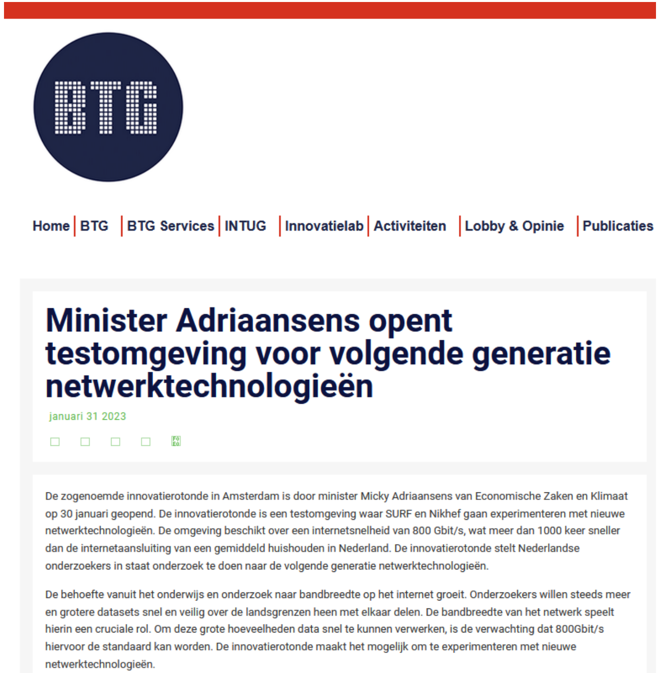
BTG
Home | BTG | BTG Services | INTUG | Innovatielab | Activiteiten | Lobby & Opinie | Publicaties

Minister Adriaansens opent testomgeving voor volgende generatie netwerktechnologieën

januari 31 2023

□ □ □ □

De zogenoemde innovatierotonde in Amsterdam is door minister Micky Adriaansens van Economische Zaken en Klimaat op 30 januari geopend. De innovatierotonde is een testomgeving waar SURF en Nikhef gaan experimenteren met nieuwe netwerktechnologieën. De omgeving beschikt over een internetsnelheid van 800 Gbit/s, wat meer dan 1000 keer sneller dan de internetaansluiting van een gemiddeld huishouden in Nederland. De innovatierotonde stelt Nederlandse onderzoekers in staat onderzoek te doen naar de volgende generatie netwerktechnologieën.

De behoefte vanuit het onderwijs en onderzoek naar bandbreedte op het internet groeit. Onderzoekers willen steeds meer en grotere datasets snel en veilig over de landsgrenzen heen met elkaar delen. De bandbreedte van het netwerk speelt hierin een cruciale rol. Om deze grote hoeveelheden data snel te kunnen verwerken, is de verwachting dat 800Gbit/s hiervoor de standaard kan worden. De innovatierotonde maakt het mogelijk om te experimenteren met nieuwe netwerktechnologieën.

Web screenshot: btg.org,
Images Nokia 7750-SR1x in Nikhef AMS H234b: Tristan Suerink

Innovation on infrastructure

Nikhef empowers subatomic physics research with AMD EPYC™ CPUs and Radeon Instinct™ GPUs

Deepening our understanding of the universe with AMD EPYC CPU-powered and Radeon Instinct GPU-accelerated servers.

AMD EPYC

CUSTOMER

Nikhef

INDUSTRY

Subatomic Physics

CHALLENGES

Increasing data throughput with higher I/O and memory bandwidth

SOLUTION

Deploy AMD EPYC™ 7502P and 7702P CPUS, and AMD Radeon Instinct™ MI50 GPUs

RESULTS

Faster processing and the ability to harness GPU-accelerated machine learning to cope with rapidly expanding experimental data volume

AMD TECHNOLOGY AT A GLANCE

AMD EPYC 7502P processors with 32 cores
AMD EPYC 7702P processors with 64 cores
AMD Radeon Instinct MI50 GPUs

TECHNOLOGY PARTNER

Lenovo

Many of the latest scientific discoveries are as much about the computing power used to analyze experimental data as they are about the theories behind them. At the forefront of delivering the processing capabilities for subatomic physics research is Nikhef, the Dutch National Institute concentrating on this area. Nikhef has provided computing that has helped with the discoveries of gravitational waves in 2016, the Higgs Boson, and the fundamental physics in between, including confirmation that many of the heavy elements in the universe are produced in neutron star mergers.

"We were able to be the worldwide number one in a few public projects like Rosetta@home and Worldwide Community Grid with the AMD EPYC cluster."

Roel Aaij, Scientific Staff Member, Nikhef

"The Institute performs blue-sky research to learn more about the nature of the universe and the building blocks of matter," explains Roel Aaij, Scientific Staff Member at Nikhef. "The fundamental goal of this institute is to find the big universal box of building blocks everything is made from," adds Tristan Suerink, IT Architect at Nikhef. The more computing power that the Institute can throw at this quest, the more that can be discovered. This led the team to AMD EPYC™ processors and Radeon Instinct™ GPUs, which delivered the performance Nikhef's workloads required and the solution price that aligned with their budget.

Data-hungry science

Nikhef is involved in many different experiments, but all of them require a considerable level of computing power. "About 100 scientific staff work at Nikhef," explains Aaij. "These staff usually work on one (or sometimes more than one) of the experiments Nikhef is involved in. Three of those experiments are at CERN: the ATLAS, LHCb, and ALICE experiments. There are several astroparticle physics experiments. One is the Pierre Auger experiment, covering several thousand square kilometers of Pampa in Argentina. The area is equipped with detectors to search for air showers caused by extremely high-energy particles that arrive from the universe. Then there is the neutrino-physics experiment KM3NeT, and dark-matter research with the XENON experiment. Finally, there is a large gravitational waves physics group that is a member of the LIGO-Virgo experiment collaboration."

If there's one thing all these experiments have in common, it's the increasing amounts of data that the experiments produce. "The scientists always want more data," says Suerink. "I think there are few experimental physics papers that do not end with 'we need more data.' And in this field of physics, to get more data you build a more sensitive experiment." In the case of the Large Hadron Collider (LHC) at CERN, the leap in data produced will be particularly huge.

"In about five years the LHC will increase the number of collisions detected by about a factor of 10," says Aaij. "This means that the experiments will start producing a similarly increasing amount of data. If we look at the growth of storage space and compute capacity over time, then we do not expect to even get close to a factor 10 in increase of performance for a flat budget. We need to deal with that, because we need to process the data. Otherwise, we can't do science with it." This is where AMD EPYC processors and GPU acceleration have offered the best solutions to satiate the hunger for growing data processing ability.

AMD + NIKHEF CASE STUDY

AMD

FUNGIBLE

NIKHEF, SURF AND FUNGIBLE SET NEW BENCHMARK FOR THE WORLD'S FASTEST STORAGE PERFORMANCE

Companies Double Current Performance Record, Setting the New Bar at 6.55 Million Read IOPS

798.49 Gb/s
Nikhef
SURF
BLUEFIELD HACKATHON
One SmartNIC. Field of Opportunities.
3rd Place
IN RECOGNITION of Your Exceptional Accomplishments
BlueField

Image: Minister of Economic Affairs M. Adriaansens launched the Innovation Hub with Nikhef, SURF, Nokia and NL-ix, January 2023. Composite image from https://www.surf.nl/nieuws/minister-adriaansens-lanceert-testomgeving-voor-supersnelle-netwerktechnologie

Research data traffic looks like … a DDoS to others ☺

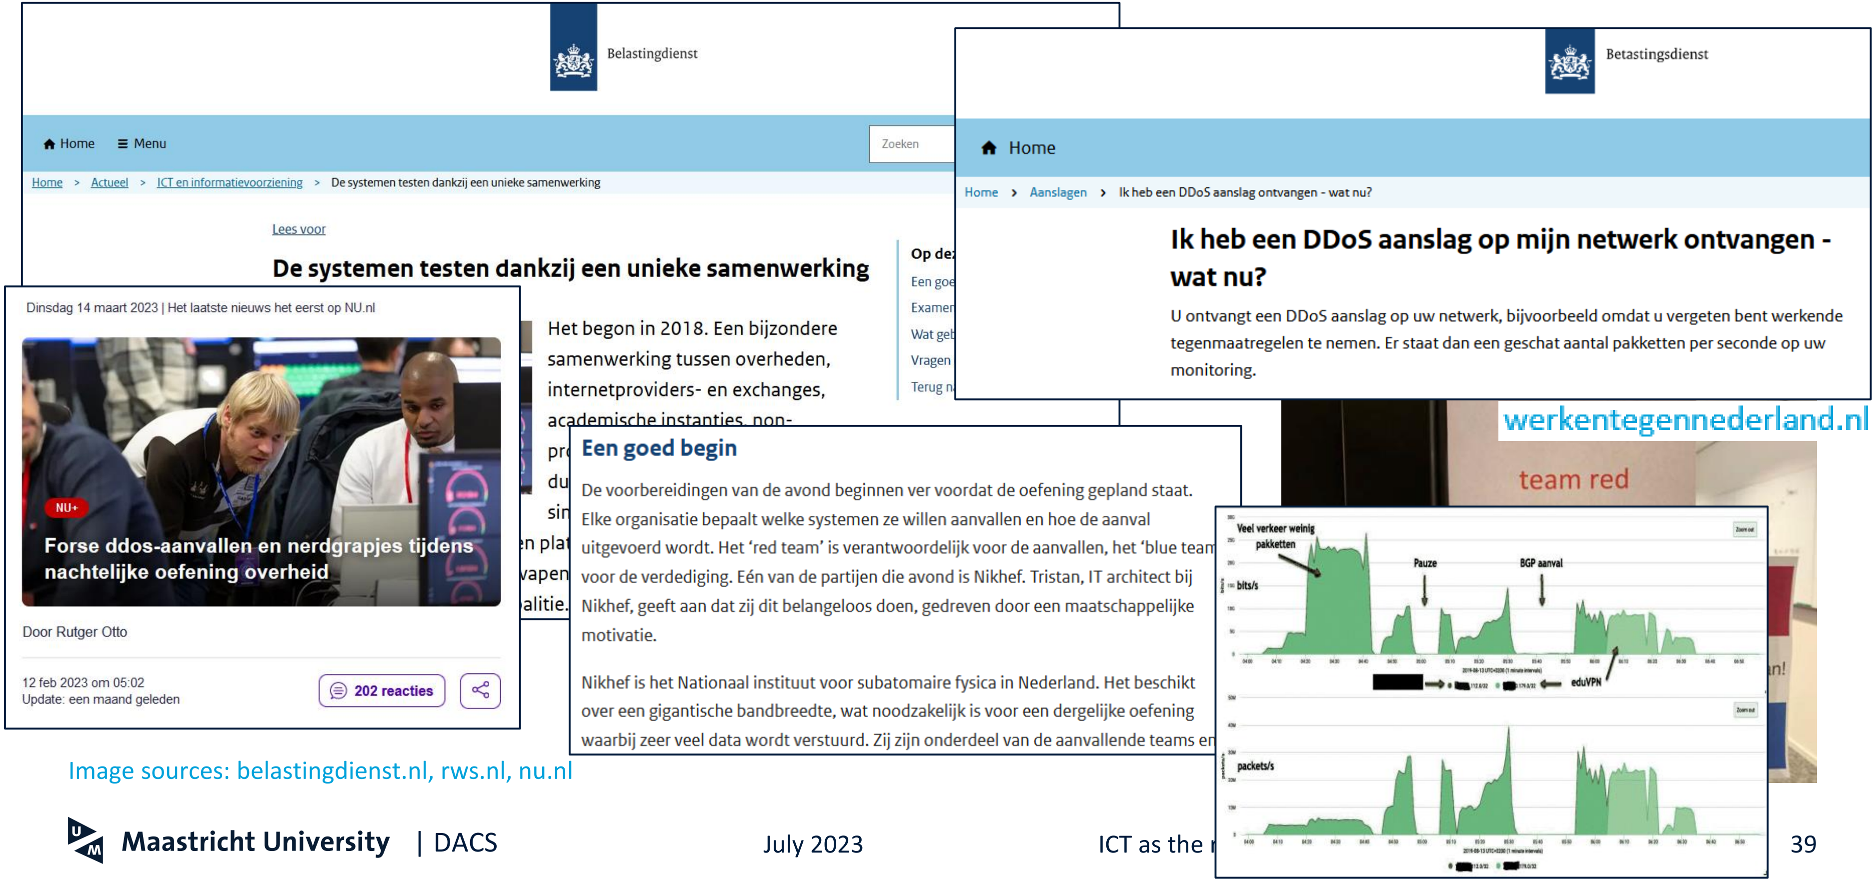

Belastingdienst

⌂ Home ≡ Menu Zoeken

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

Lees voor

De systemen testen dankzij een unieke samenwerking

Het begon in 2018. Een bijzondere samenwerking tussen overheden, internetproviders- en exchanges, academische instanties, non-
pr
du
sir
en plat
vapen
alitie.

Op de
Een goe
Examen
Wat geb
Vragen
Terug n

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

NU+

Forse ddos-aanvallen en nerdgrapjes tijdens nachtelijke oefening overheid

Door Rutger Otto

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

202 reacties

Belastingdienst

⌂ 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.

Een goed begin

De voorbereidingen van de avond beginnen ver voordat de oefening gepland staat. Elke organisatie bepaalt welke systemen ze willen aanvallen en hoe de aanval uitgevoerd wordt. Het 'red team' is verantwoordelijk voor de aanvallen, het 'blue team voor de verdediging. Eén van de partijen die avond is Nikhef. Tristan, IT architect bij Nikhef, geeft aan dat zij dit belangeloos doen, gedreven door een maatschappelijke motivatie.

Nikhef is het Nationaal instituut voor subatomaire fysica in Nederland. Het beschikt over een gigantische bandbreedte, wat noodzakelijk is voor een dergelijke oefening waarbij zeer veel data wordt verstuurd. Zij zijn onderdeel van de aanvallende teams en

werkentegennederland.nl

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

Trust & Identity

Safe access for open data processing

More than one user, *from*
more than one organizational domain, *in*
more than one country!

WLCG: when we met a global trust scaling issue

170 sites
~60 countries & regions
~20000 users
just *how* many interactions ??

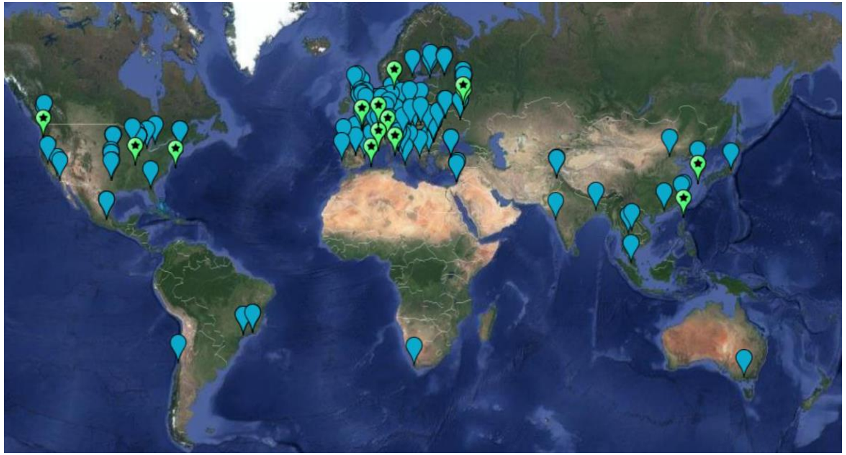

people photo: a small part of the CMS collaboration in 2017, Credit: CMS-PHO-PUBLIC-2017-004-3; site map: WLCG sites from Maarten Litmaath (CERN) 2021

We live in a federated world!

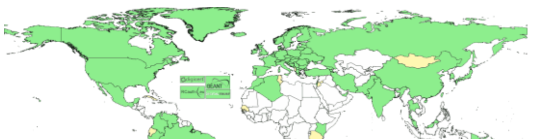

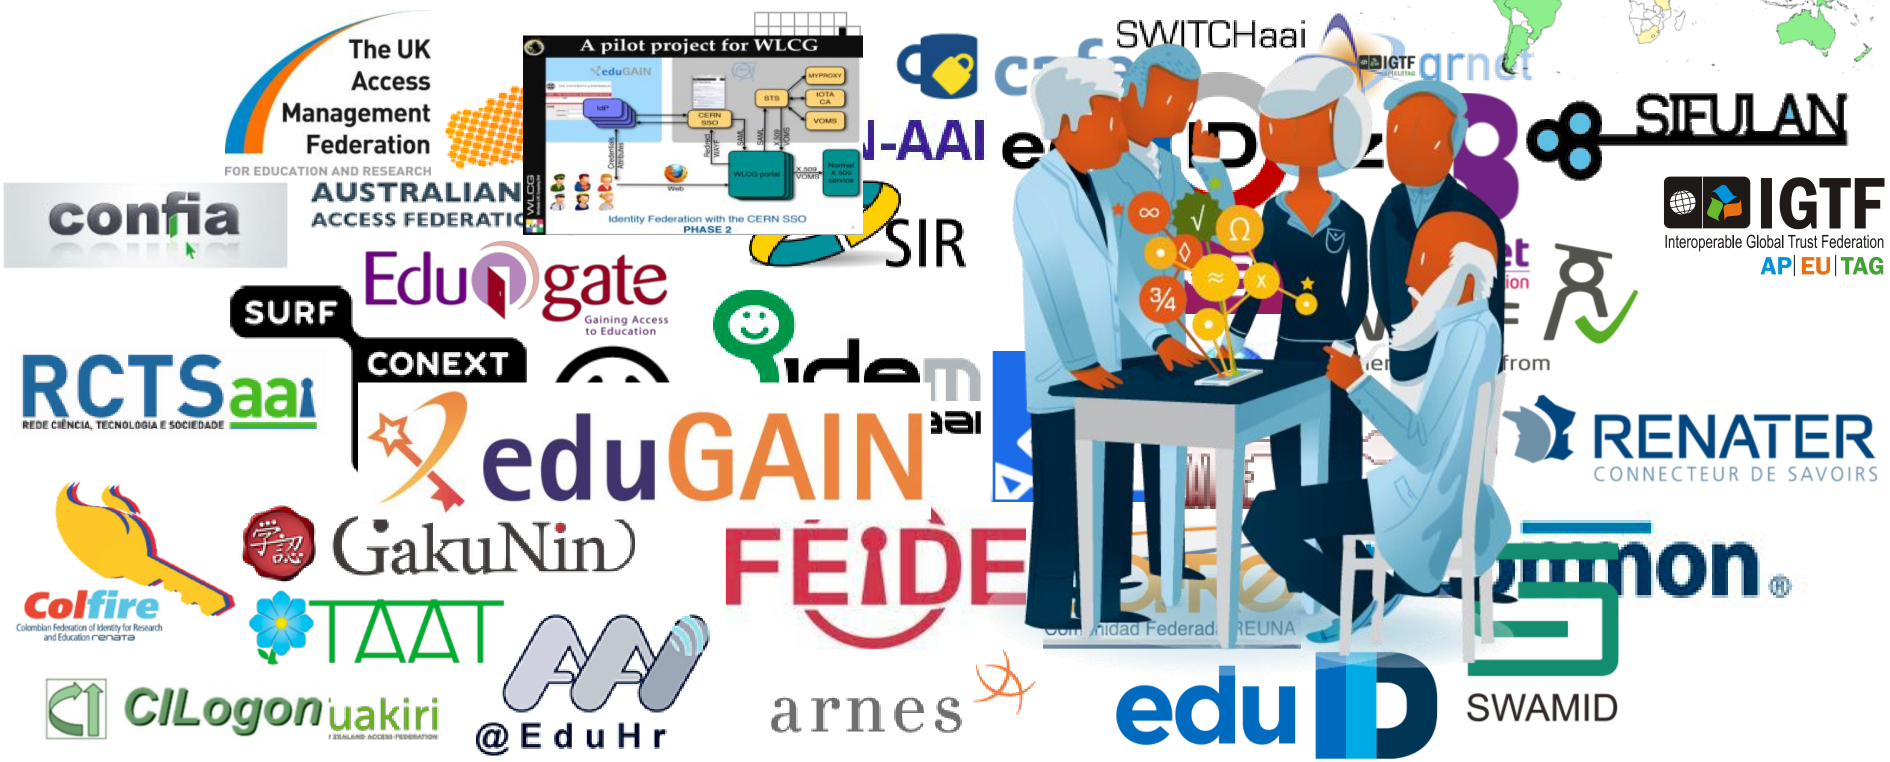

slide inspiration: Licia Florio, NORDUNET

eduGAIN

Implementation of eduGAIN Future WG recommendations

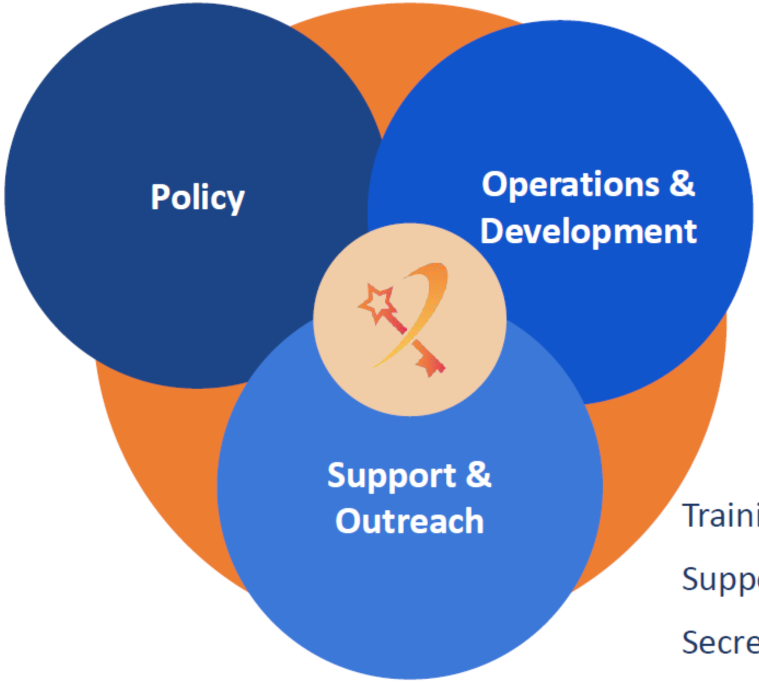

Policy

Operations & Development

Support & Outreach

Ongoing operations

New eduGAIN-OT

Evolution and duplication of core infrastructure

Training, Webinars , T&I town hall

Support and CSIRT

Secretariat, Business development

78

Identity Federations

5100+

Identity Providers

3600+

Service Providers

slide by: Maarten Kremers, SURF, for the GEANT 5-1 project

Federated Access

Login via the Nikhef service proxy to gitlab, ifosim.org, …

“Where are you from”

discovery screen showing entities from the eduGAIN global interfederation

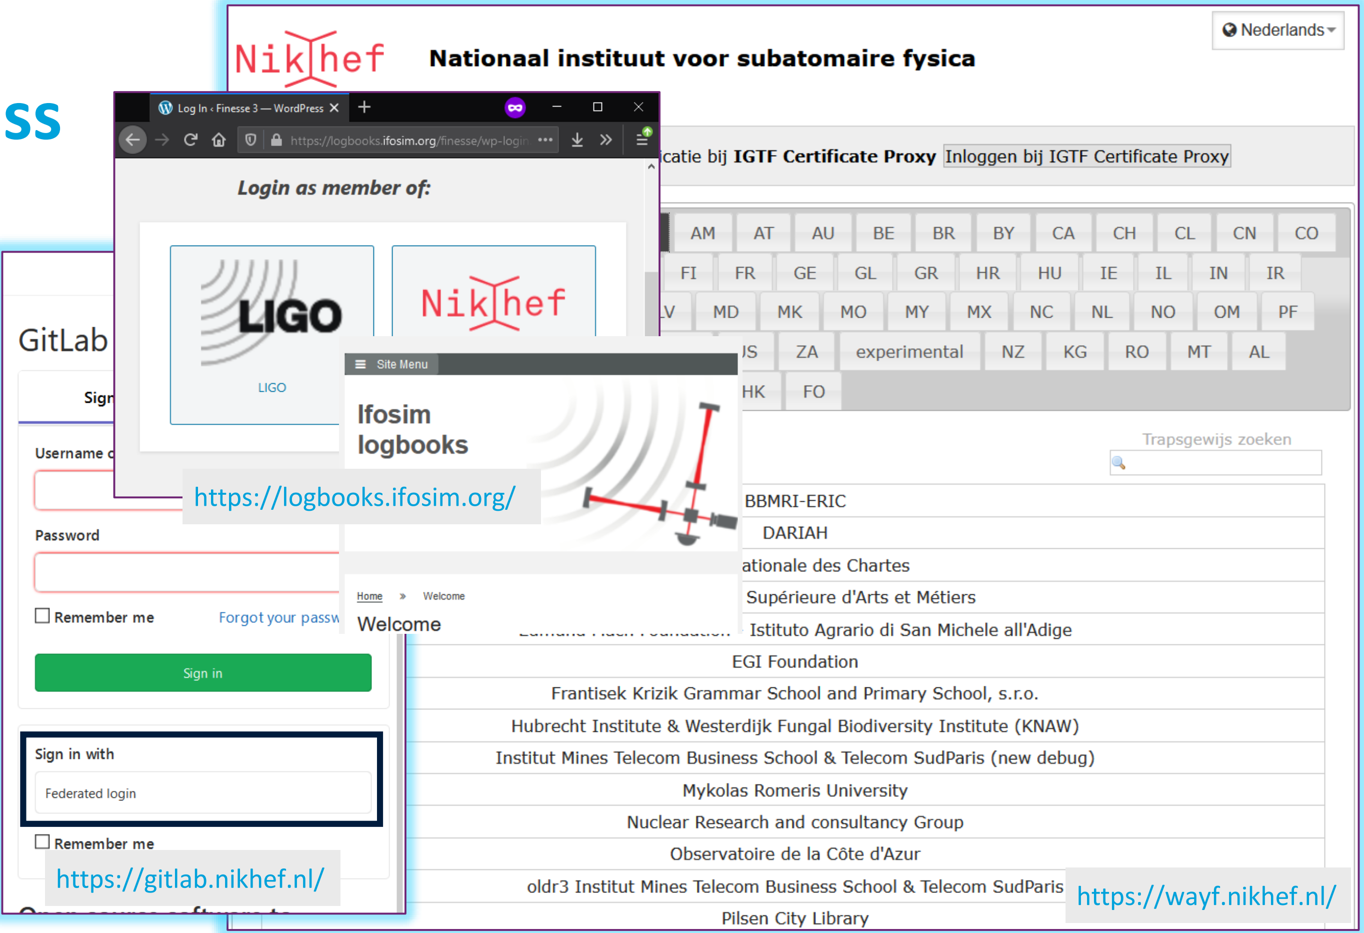

Nederlands

Nationaal instituut voor subatomaire fysica

…catie bij **IGTF Certificate Proxy** Inloggen bij IGTF Certificate Proxy

AM AT AU BE BR BY CA CH CL CN CO
FI FR GE GL GR HR HU IE IL IN IR
V MD MK MO MY MX NC NL NO OM PF
JS ZA experimental NZ KG RO MT AL
HK FO

Trapsgewijs zoeken

BBMRI-ERIC
DARIAH
ationale des Chartes
Supérieure d'Arts et Métiers
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
Pilsen City Library

https://wayf.nikhef.nl/

Log In ‹ Finesse 3 — WordPress

https://logbooks.ifosim.org/finesse/wp-login

Login as member of:

LIGO

Nikhef

Site Menu

Ifosim logbooks

Home » Welcome

Welcome

https://logbooks.ifosim.org/

GitLab

Sign

Username o

Password

□ Remember me   Forgot your passw

Sign in

Sign in with

Federated login

□ Remember me

https://gitlab.nikhef.nl/

ifosim federated AAI integration implementation by Mischa Sallé; per-country WAYF selection is a bespoke Nikhef WAYF feature

IGTF: a policy-bridged global federation for research computing

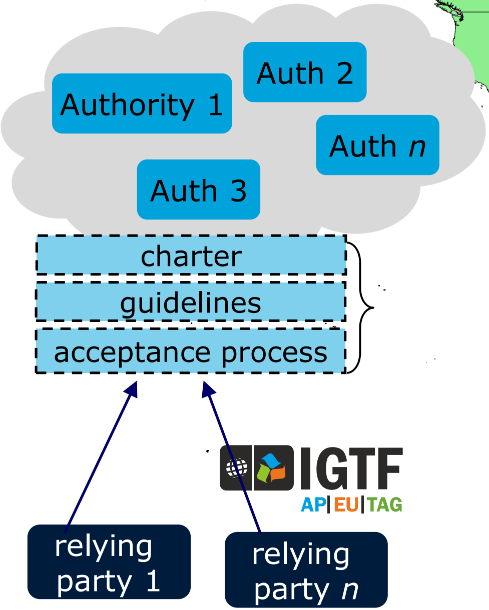

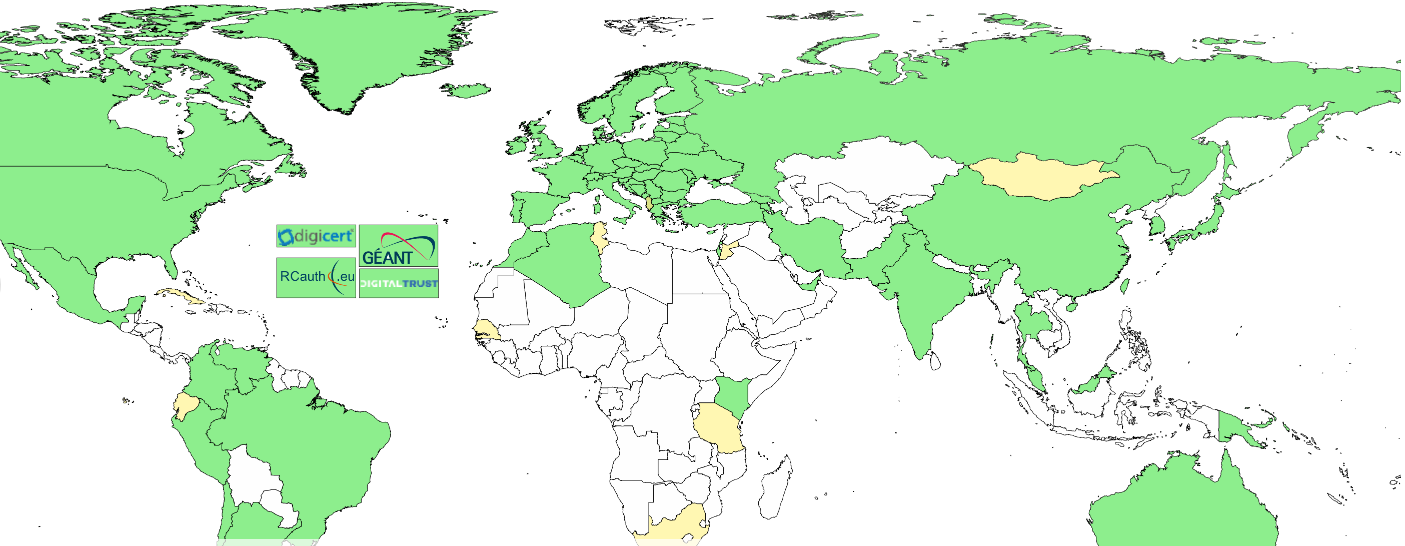

A global authentication fabric & assurance standards

~ 90 Identity Providers (some leveraging a R&E federation)

~ 10 international research and e-infrastructure relying parties

> 60 countries / economic areas / international treaty orgs

> 1000 relying service provider collaborations

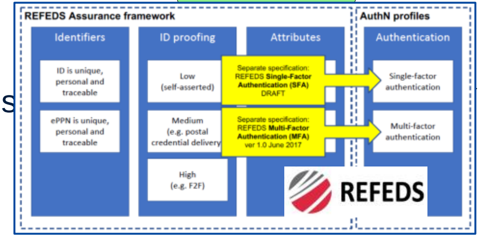
REFEDS Assurance framework
AuthN profiles
Identifiers
ID is unique, personal and traceable
ePPN is unique, personal and traceable
ID proofing
Low (self-asserted)
Medium (e.g. postal credential delivery)
High (e.g. F2F)
Attributes
Separate specification: REFEDS Single-Factor Authentication (SFA) DRAFT
Separate specification: REFEDS Multi-Factor Authentication (MFA) ver 1.0 June 2017
Authentication
Single-factor authentication
Multi-factor authentication
REFEDS

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

Separating source of authenticator, identity, and access

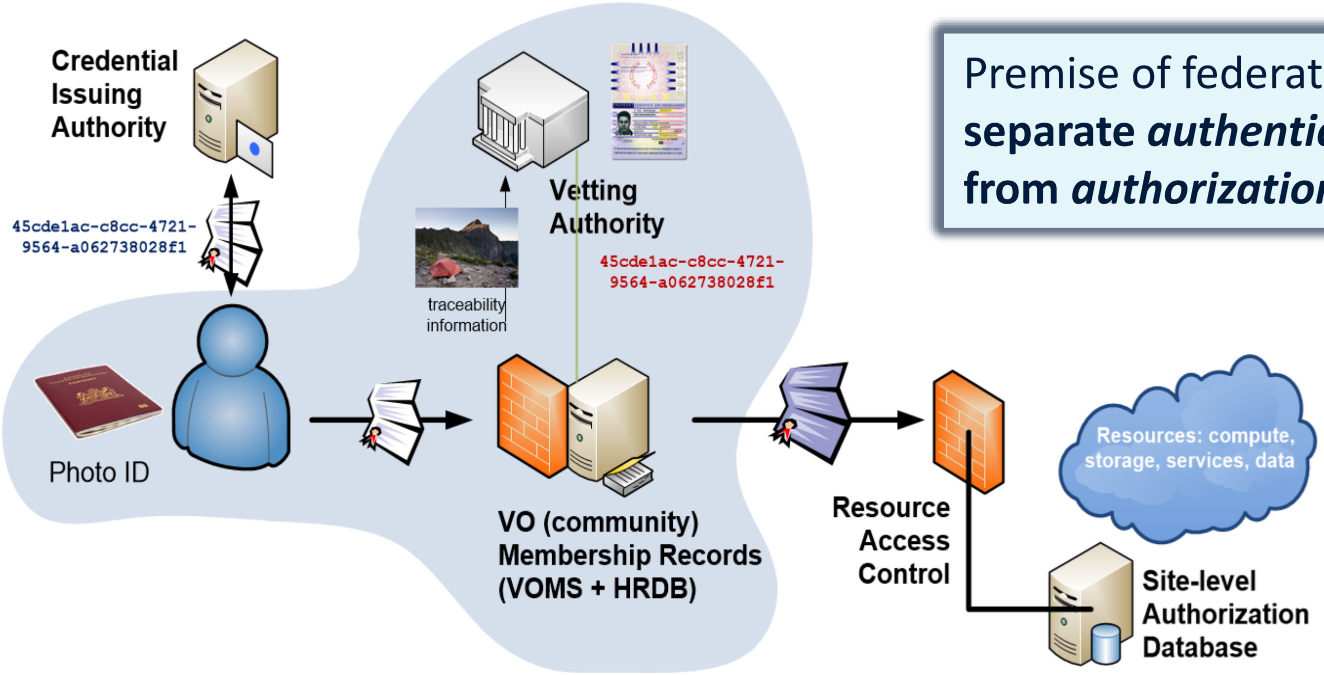

Premise of federated access **separate *authentication* from *authorization***

Research Infrastructures and access models based on eduGAIN

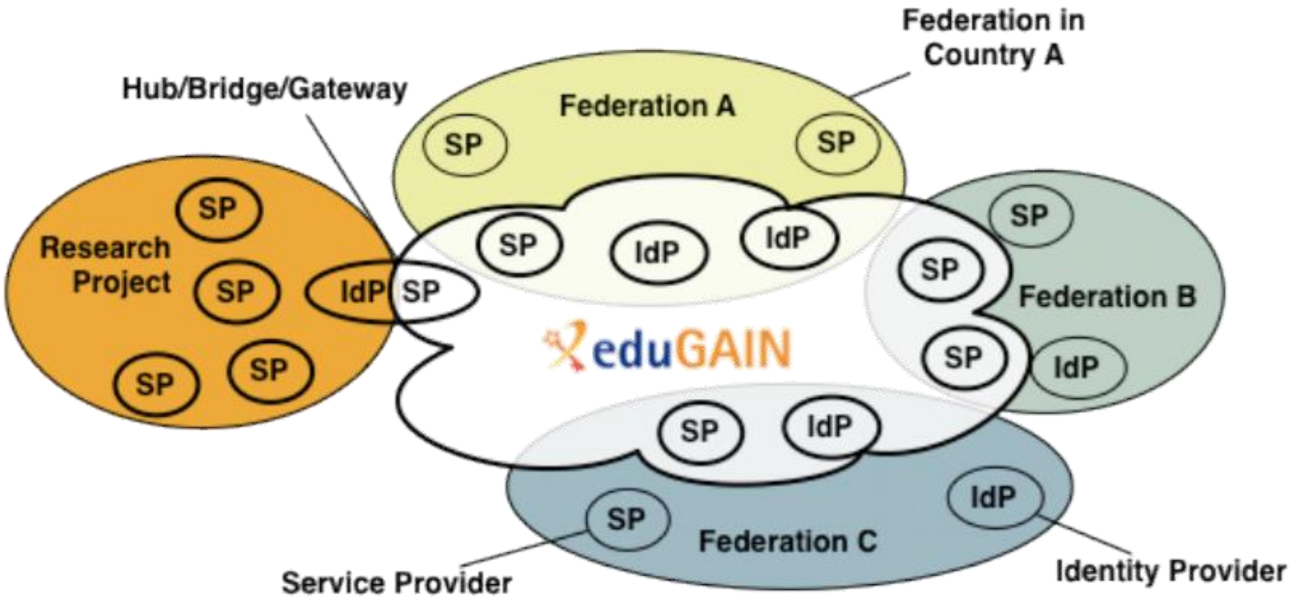

Federation with SP Proxy image by: SWITCH (CH)

Source of authority for access to research SPs defined by the research project (ERIC, ESFRI), not home organization IdP which only has affiliation

So ‘just eduGAIN’ is not enough for research collaboration

AARC

- Access services using **identities from their Home Organizations.**
- **Access** services **based on role(s)** users have **in the collaboration**. This info is not known to IdPs – or eduGAIN.
- Secure integration of **guest identity solutions** and **support for stronger authentication identity assurance** mechanisms.
- Requirement for **one persistent identity** across all the community’s services when needed and **account linking**.
- **Web** and **non-web** resources
- **Hide complexity** of multiple IdPs/feds/At Auth/ technologies.

Authentication and Authorization for Research Collaboration – AARC (Licia Florio *et al.*) – https://aarc-community.org/

Federated access for research collaboration – AARC

Authentication and Authorization architecture for Research Collaboration

Defines a model and building blocks to address researcher needs exploiting group membership for authorization

eduGAIN and the Identity Federations

Foundational federated access in R&E
Allows researchers to use ONE digital identity to access MANY services and resources available in eduGAIN

Network connectivity

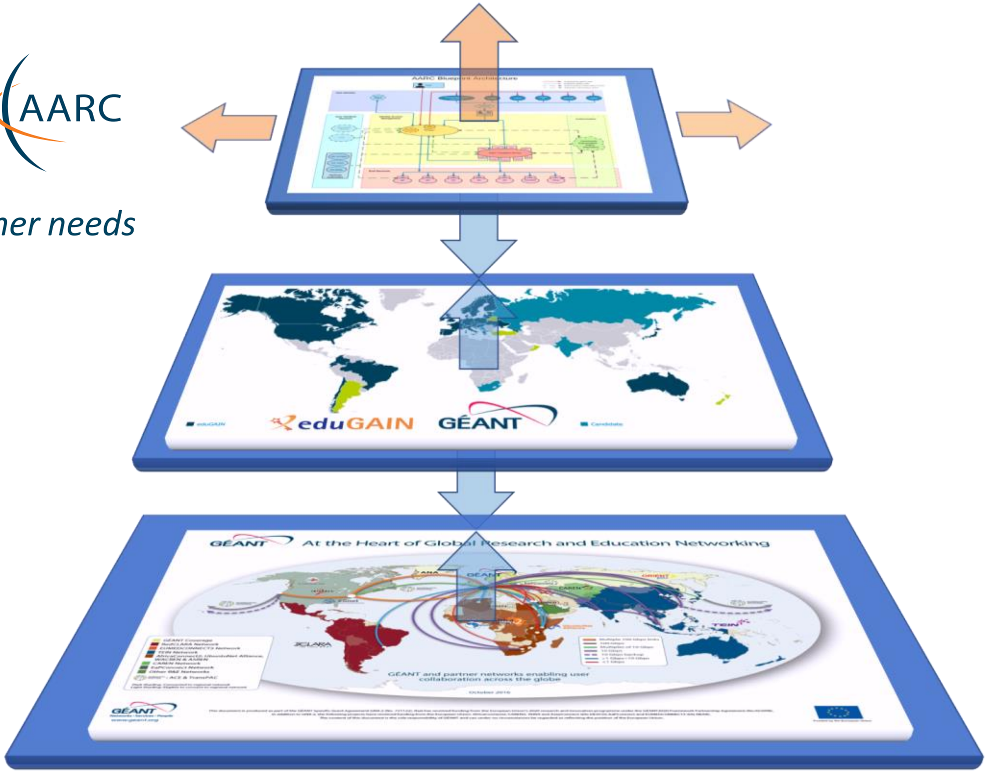
AARC Blueprint Architecture
eduGAIN
GÉANT
GÉANT At the Heart of Global Research and Education Networking
GÉANT and partner networks enabling user collaboration across the globe
October 2016

Trust flows from the research community

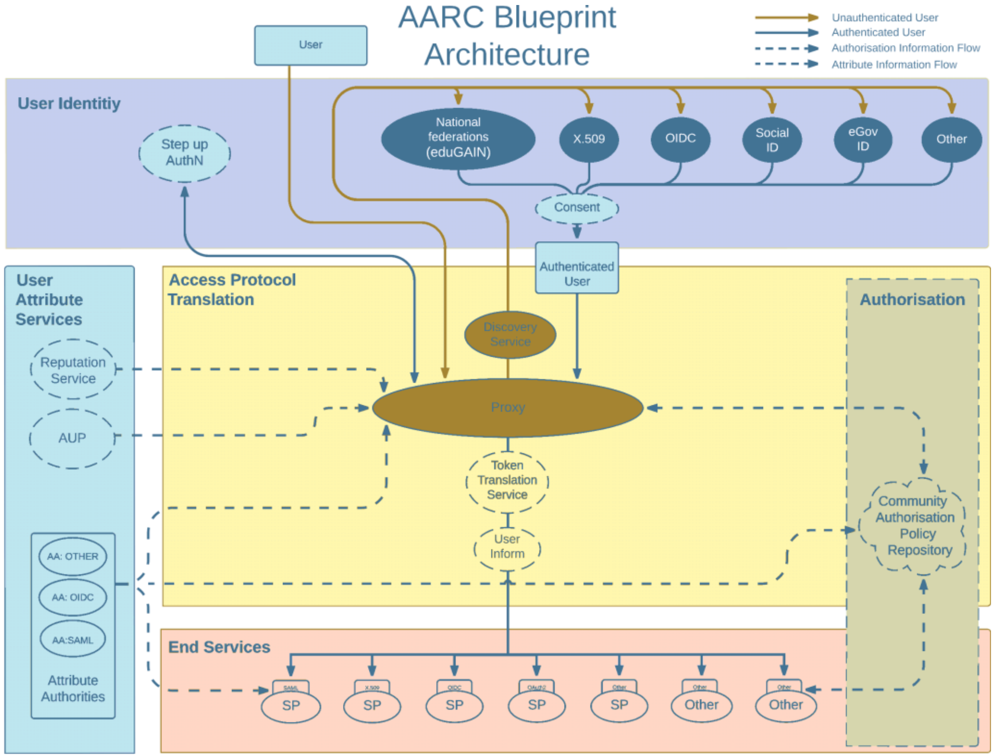

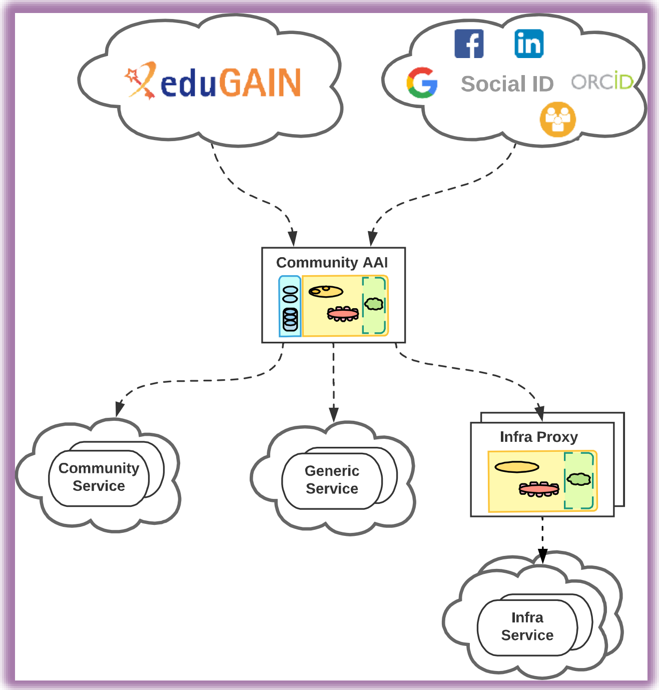

AARC Blueprint Architecture (2019) AARC-G045 https://aarc-community.org/guidelines/aarc-g045/; stacked proxies: EOSC AAI Architecture
EOSC Authentication and Authorization Infrastructure (AAI), ISBN 978-92-76-28113-9, http://doi.org/10.2777/8702

Composite AAIs – proxies beyond ‘just’ the EOSC

Proxy model supports harmonizing IdPs beyond research

- **eduID**-style identifiers
  - ‘life-long learning’ identifiers
  - independent student identifier (the ESI) for mobility & Erasmus-without-papers
  - eduGAIN-alignment foreseen: eduid.nl, Swiss eduID, …
- **eIDAS** and government eID (e.g. DigID)
  - identity assurance step-up
- **ORCID** provides this service for research in general
  - since it persists, also very useful to allow researchers consistent access independent of home org ☺

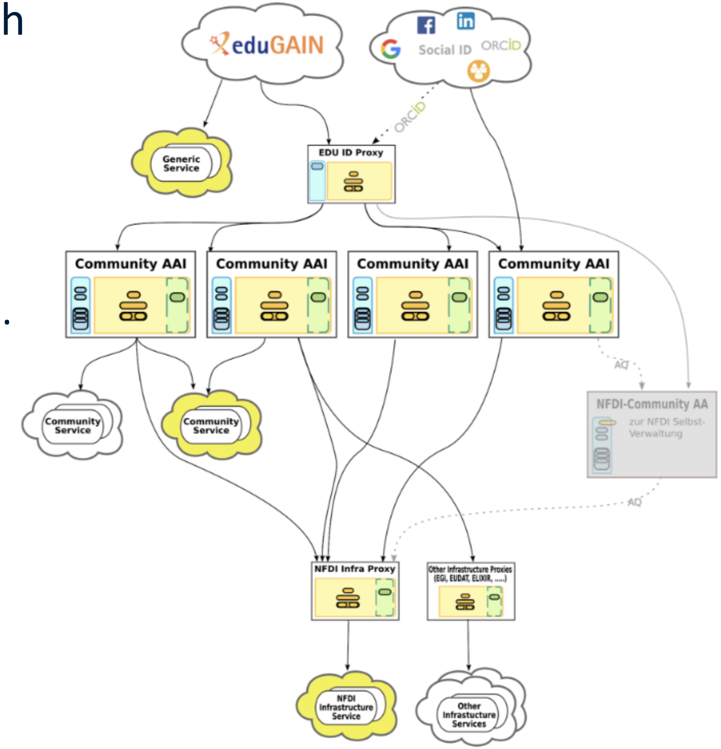

Composite AAI image source: Christos Kanellopoulos (GEANT), Marcus Hardt (KIT)

EOSC AAI Federation

<span>Identity assurance</span> brings the true value: authenticators are aplenty, and 'MFA' far less interesting than vetted identities. But HEI home IdPs seem reluctant to provide it …

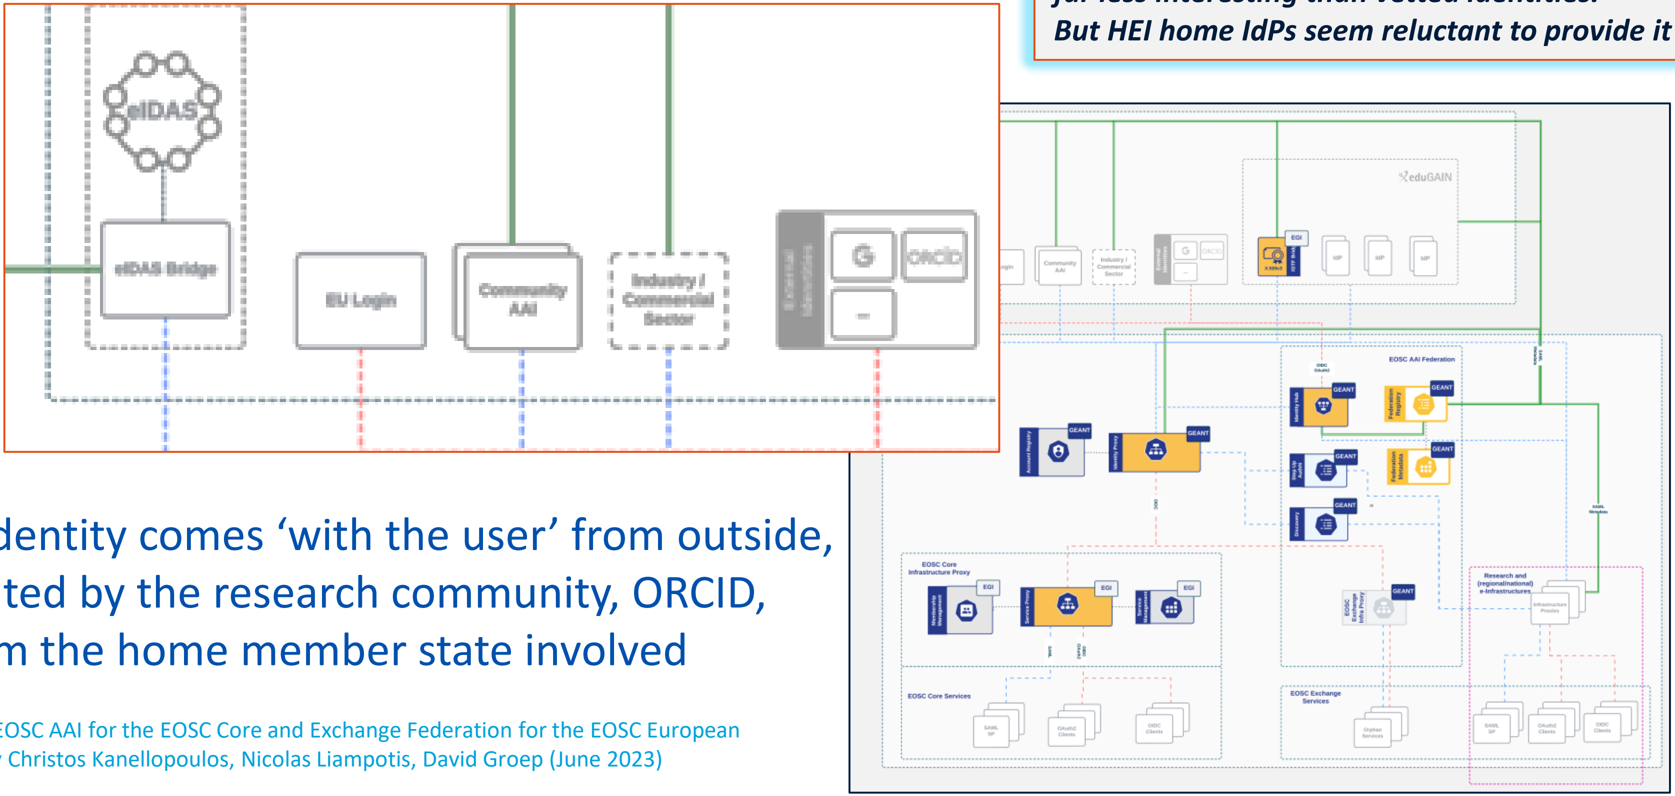

user identity comes 'with the user' from outside, mediated by the research community, ORCID, or from the home member state involved

Image: EOSC AAI for the EOSC Core and Exchange Federation for the EOSC European Node by Christos Kanellopoulos, Nicolas Liampotis, David Groep (June 2023)

Same blocks underlie e.g. the Fenix and Puhuri HPC ecosystem

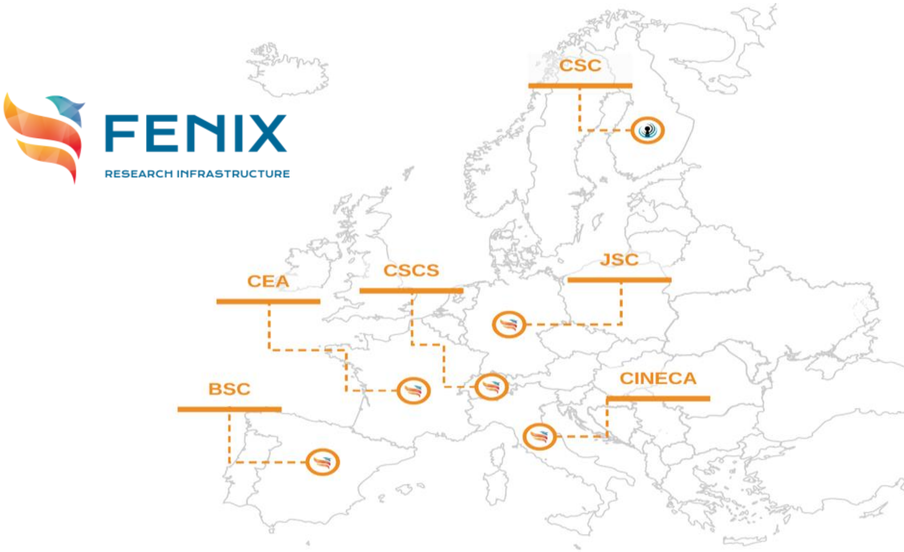

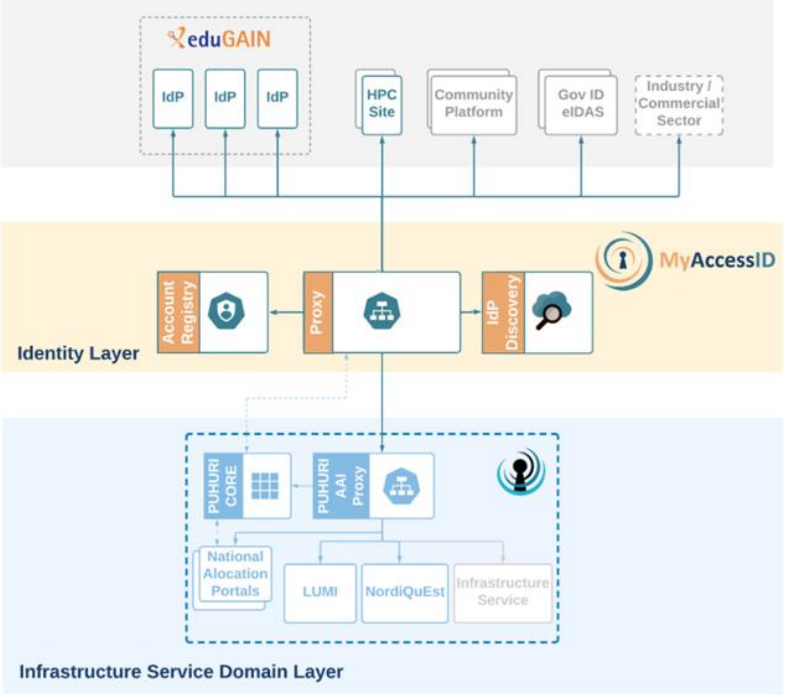

Fenix image via Christos Kanellopoulos, diagram via Anders Sjöström (NeIC, Puhuri) at the TNC23 workshop

And the blocks are the basis for education & Erasmus+

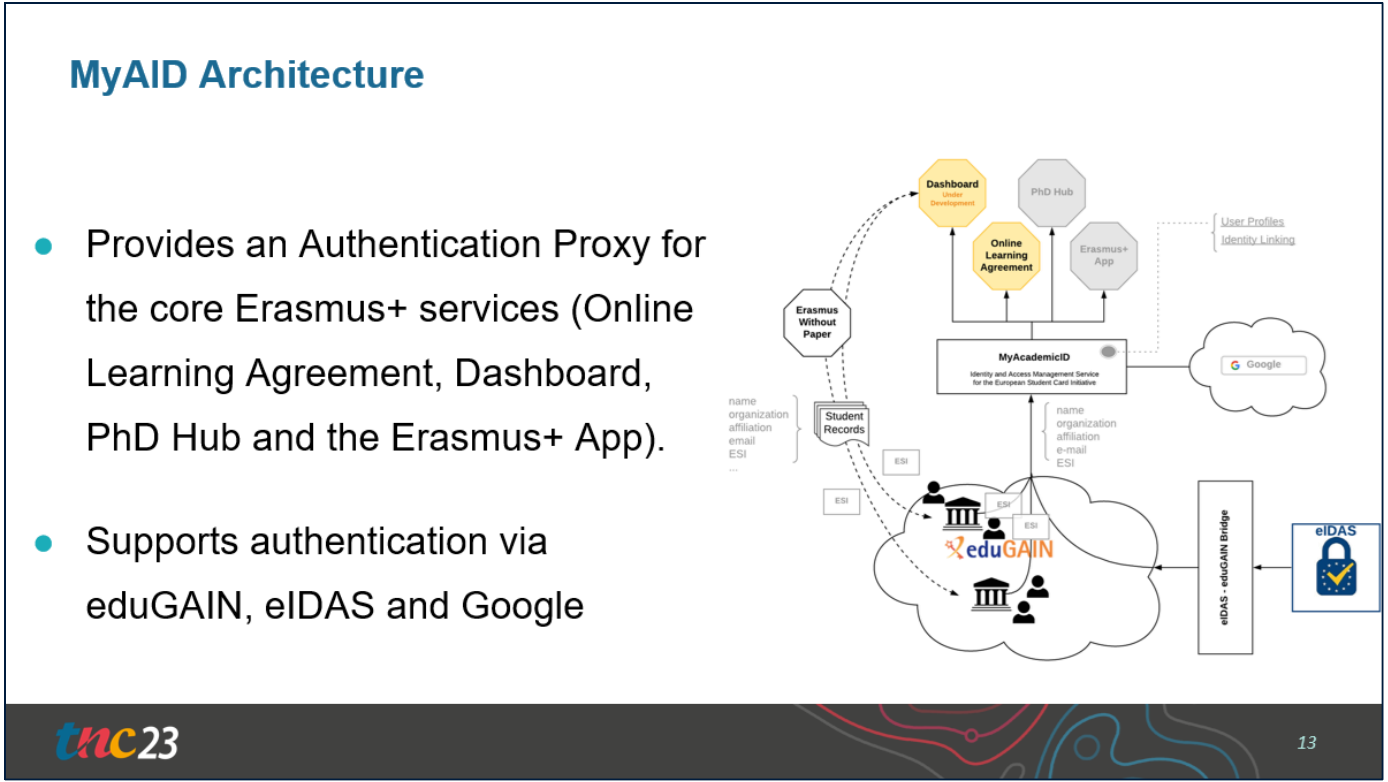

MyAID Architecture

- Provides an Authentication Proxy for the core Erasmus+ services (Online Learning Agreement, Dashboard, PhD Hub and the Erasmus+ App).
- Supports authentication via eduGAIN, eIDAS and Google

Christis Kanellopoulos (GÉANT) for the Erasmus+/Erasmus Without Papers programme

What value does our university ID bring in a life-long learning environment? Time to think less institution-centric?

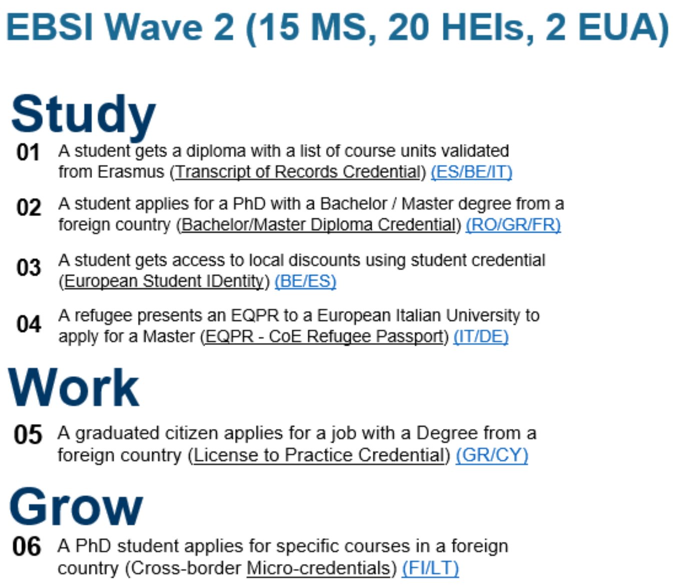

EBSI Wave 2 (15 MS, 20 HEIs, 2 EUA)

Study

01 A student gets a diploma with a list of course units validated from Erasmus (Transcript of Records Credential) (ES/BE/IT)

02 A student applies for a PhD with a Bachelor / Master degree from a foreign country (Bachelor/Master Diploma Credential) (RO/GR/FR)

03 A student gets access to local discounts using student credential (European Student IDentity) (BE/ES)

04 A refugee presents an EQPR to a European Italian University to apply for a Master (EQPR - CoE Refugee Passport) (IT/DE)

Work

05 A graduated citizen applies for a job with a Degree from a foreign country (License to Practice Credential) (GR/CY)

Grow

06 A PhD student applies for specific courses in a foreign country (Cross-border Micro-credentials) (FI/LT)

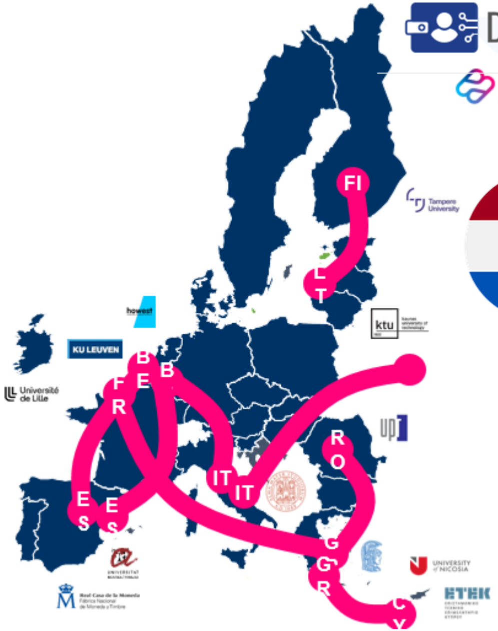

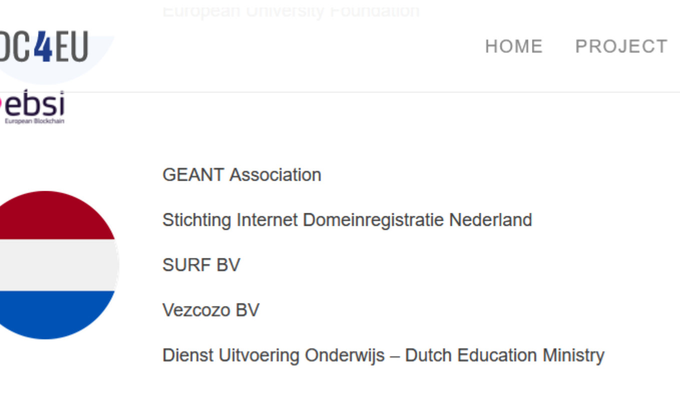

European University Foundation

HOME PROJECT

DC4EU
ebsi European Blockchain

GEANT Association

Stichting Internet Domeinregistratie Nederland

SURF BV

Vezcozo BV

Dienst Uitvoering Onderwijs – Dutch Education Ministry

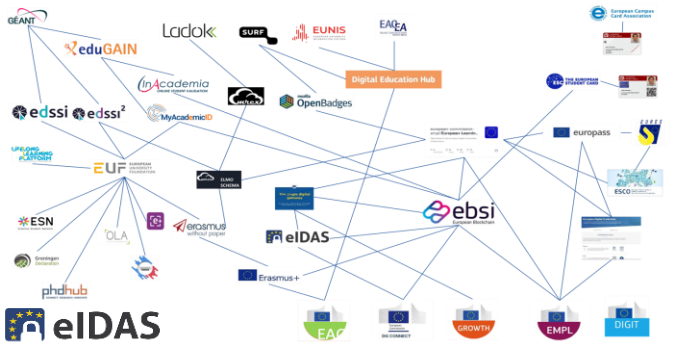

Images from Lluís Ariño, for the DC4EU project. See e.g. https://www.dc4eu.eu/consortium/#netherlands

Open science & infrastructure ecosystem enabled by Federation

Common infrastructure for many communities
ESFRI Clusters and the European Open Science Cloud EOSC

A global infrastructure of EGI, OSG and WLCG, …

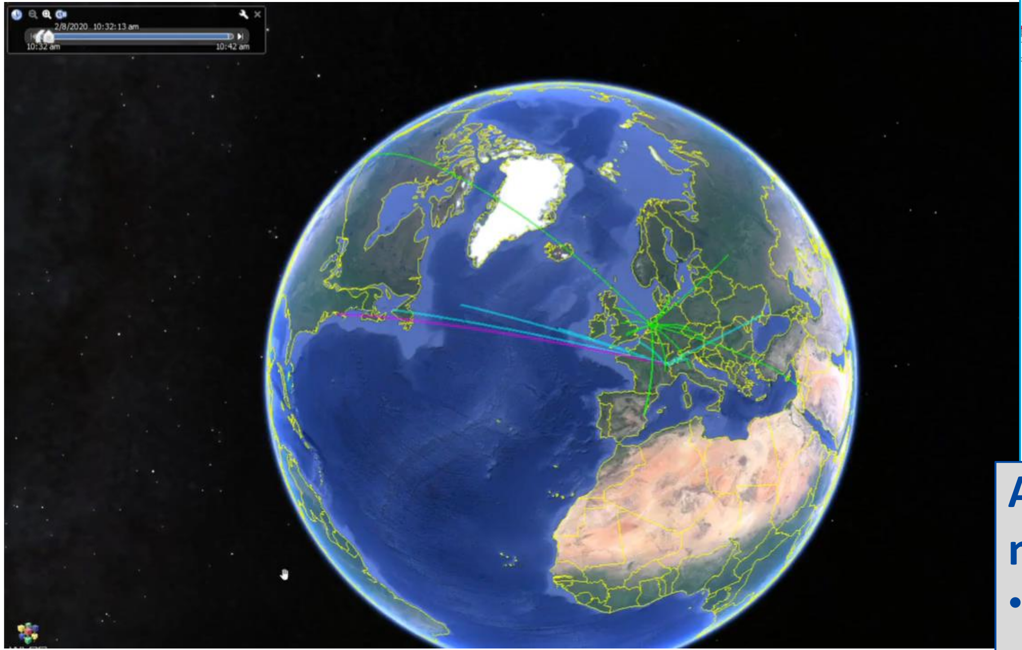

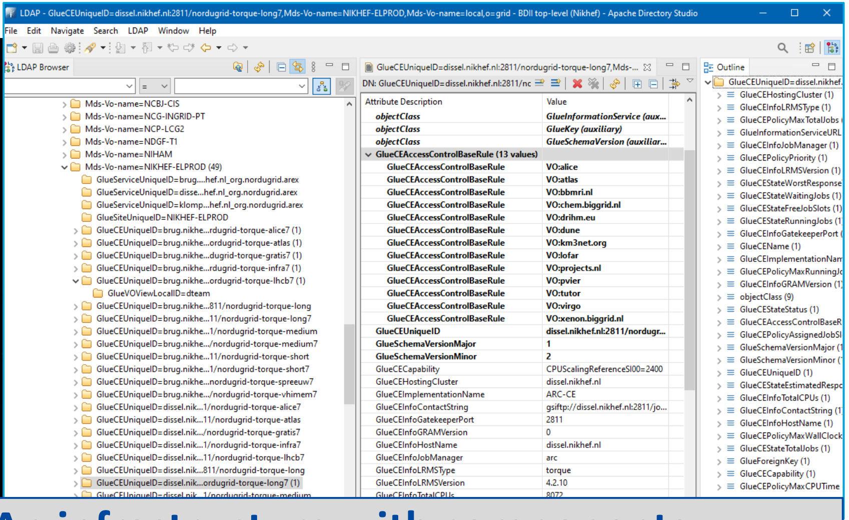

LDAP - GlueCEUniqueID=dissel.nikhef.nl:2811/nordugrid-torque-long7,Mds-Vo-name=NIKHEF-ELPROD,Mds-Vo-name=local,o=grid - BDII top-level (Nikhef) - Apache Directory Studio

```
> Mds-Vo-name=NCBJ-CIS
> Mds-Vo-name=NCG-INGRID-PT
> Mds-Vo-name=NCP-LCG2
> Mds-Vo-name=NDGF-T1
> Mds-Vo-name=NIHAM
v Mds-Vo-name=NIKHEF-ELPROD (49)
    GlueServiceUniqueID=brug....hef.nl_org.nordugrid.arex
    GlueServiceUniqueID=disse...hef.nl_org.nordugrid.arex
    GlueServiceUniqueID=klomp...hef.nl_org.nordugrid.arex
    GlueSiteUniqueID=NIKHEF-ELPROD
  > GlueCEUniqueID=brug.nikhe...rdugrid-torque-alice7 (1)
  > GlueCEUniqueID=brug.nikhe...ordugrid-torque-atlas (1)
  > GlueCEUniqueID=brug.nikhe...dugrid-torque-gratis7 (1)
  > GlueCEUniqueID=brug.nikhe...rdugrid-torque-infra7 (1)
  v GlueCEUniqueID=brug.nikhe...ordugrid-torque-lhcb7 (1)
      GlueVOViewLocalID=dteam
  > GlueCEUniqueID=brug.nikhe...811/nordugrid-torque-long
  > GlueCEUniqueID=brug.nikhe...11/nordugrid-torque-long7
  > GlueCEUniqueID=brug.nikhe...1/nordugrid-torque-medium
  > GlueCEUniqueID=brug.nikhe.../nordugrid-torque-medium7
  > GlueCEUniqueID=brug.nikhe...11/nordugrid-torque-short
  > GlueCEUniqueID=brug.nikhe...1/nordugrid-torque-short7
  > GlueCEUniqueID=brug.nikhe...nordugrid-torque-spreeuw7
  > GlueCEUniqueID=brug.nikhe.../nordugrid-torque-vhimem7
  > GlueCEUniqueID=dissel.nik...1/nordugrid-torque-alice7
  > GlueCEUniqueID=dissel.nik...11/nordugrid-torque-atlas
  > GlueCEUniqueID=dissel.nik.../nordugrid-torque-gratis7
  > GlueCEUniqueID=dissel.nik...1/nordugrid-torque-infra7
  > GlueCEUniqueID=dissel.nik...11/nordugrid-torque-lhcb7
  > GlueCEUniqueID=dissel.nik...811/nordugrid-torque-long
  > GlueCEUniqueID=dissel.nik...ordugrid-torque-long7 (1)
  > GlueCEUniqueID=dissel.nik...1/nordugrid-torque-medium
```

GlueCEUniqueID=dissel.nikhef.nl:2811/nordugrid-torque-long7,Mds-...

DN: GlueCEUniqueID=dissel.nikhef.nl:2811/nc

Attribute Description	Value
objectClass	*GlueInformationService (aux...*
objectClass	*GlueKey (auxiliary)*
objectClass	*GlueSchemaVersion (auxiliar...*
GlueCEAccessControlBaseRule (13 values)	
GlueCEAccessControlBaseRule	**VO:alice**
GlueCEAccessControlBaseRule	**VO:atlas**
GlueCEAccessControlBaseRule	**VO:bbmri.nl**
GlueCEAccessControlBaseRule	**VO:chem.biggrid.nl**
GlueCEAccessControlBaseRule	**VO:drihm.eu**
GlueCEAccessControlBaseRule	**VO:dune**
GlueCEAccessControlBaseRule	**VO:km3net.org**
GlueCEAccessControlBaseRule	**VO:lofar**
GlueCEAccessControlBaseRule	**VO:projects.nl**
GlueCEAccessControlBaseRule	**VO:pvier**
GlueCEAccessControlBaseRule	**VO:tutor**
GlueCEAccessControlBaseRule	**VO:virgo**
GlueCEAccessControlBaseRule	**VO:xenon.biggrid.nl**
GlueCEUniqueID	**dissel.nikhef.nl:2811/nordugr...**
GlueSchemaVersionMajor	**1**
GlueSchemaVersionMinor	**2**
GlueCECapability	CPUScalingReferenceSI00=2400
GlueCEHostingCluster	dissel.nikhef.nl
GlueCEImplementationName	ARC-CE
GlueCEInfoContactString	gsiftp://dissel.nikhef.nl:2811/jo...
GlueCEInfoGatekeeperPort	2811
GlueCEInfoGRAMVersion	0
GlueCEInfoHostName	dissel.nikhef.nl
GlueCEInfoJobManager	arc
GlueCEInfoLRMSType	torque
GlueCEInfoLRMSVersion	4.2.10
GlueCEInfoTotalCPUs	8072

```
v GlueCEUniqueID=dissel.nikhef.
  > GlueCEHostingCluster (1)
  > GlueCEInfoLRMSType (1)
  > GlueCEPolicyMaxTotalJobs (
  > GlueInformationServiceURL
  > GlueCEInfoJobManager (1)
  > GlueCEPolicyPriority (1)
  > GlueCEInfoLRMSVersion (1)
  > GlueCEStateWorstResponse
  > GlueCEStateWaitingJobs (1)
  > GlueCEStateFreeJobSlots (1)
  > GlueCEStateRunningJobs (1
  > GlueCEInfoGatekeeperPort (
  > GlueCEName (1)
  > GlueCEImplementationNam
  > GlueCEPolicyMaxRunningJo
  > GlueCEInfoGRAMVersion (1
  > objectClass (9)
  > GlueCEStateStatus (1)
  > GlueCEAccessControlBaseR
  > GlueCEPolicyAssignedJobSl
  > GlueSchemaVersionMajor (1
  > GlueSchemaVersionMinor (
  > GlueCEUniqueID (1)
  > GlueCEStateEstimatedRespo
  > GlueCEInfoTotalCPUs (1)
  > GlueCEInfoContactString (1
  > GlueCEInfoHostName (1)
  > GlueCEPolicyMaxWallClock
  > GlueCEStateTotalJobs (1)
  > GlueForeignKey (1)
  > GlueCECapability (1)
  > GlueCEPolicyMaxCPUTime
```

An infrastructure with components matched to application needs

- systems architecture, compute (clusters), networking, storage, and application structure
- in a cost-efficient, and energy-efficient, way

BerkeleyDB Information System for EGI, from top-level BDII at ldap://bdii03.nikhef.nl:2170/o=grid; Earth visualization: https://dashb-earth.cern.ch/, Google Earth

Job distribution overlay and pilot jobs in WLCG

Building a cross-site 'overlay' batch system

- work around local bespoke interfaces and specific semantics

provides the single **community-level interface**

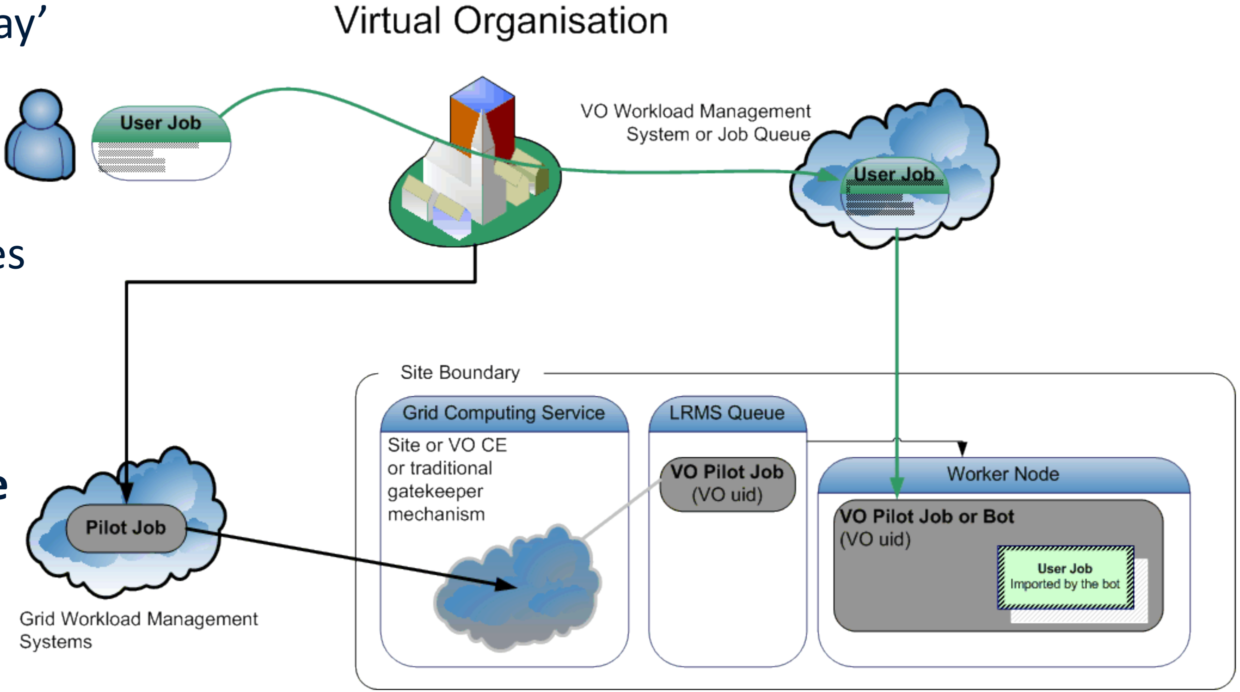
Virtual Organisation
User Job
VO Workload Management System or Job Queue
User Job
Site Boundary
Grid Computing Service
Site or VO CE or traditional gatekeeper mechanism
LRMS Queue
VO Pilot Job (VO uid)
Worker Node
VO Pilot Job or Bot (VO uid)
User Job
Imported by the bot
Pilot Job
Grid Workload Management Systems

Site Access Control with pilot jobs: gLExec, http://doi.org/10.1088/1742-6596/119/6/062032; GlideinWMS: https://glideinwms.fnal.gov/ based on Condor; also: PANDA

SLATE – structuring the research cloud overlay
Nobody wants a cloud per-se … what we want is a solution …

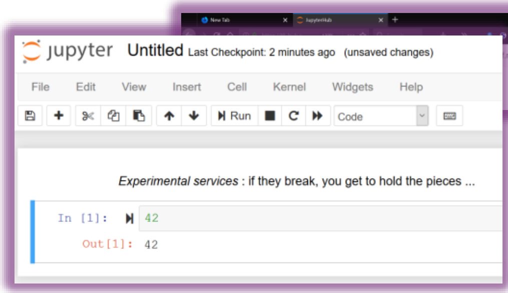

jupyter
Untitled Last Checkpoint: 2 minutes ago (unsaved changes)
File Edit View Insert Cell Kernel Widgets Help
Run
Code
Logout

Experimental services : if they break, you get to hold the pieces ...

```
In [1]:  42
Out[1]:  42
```

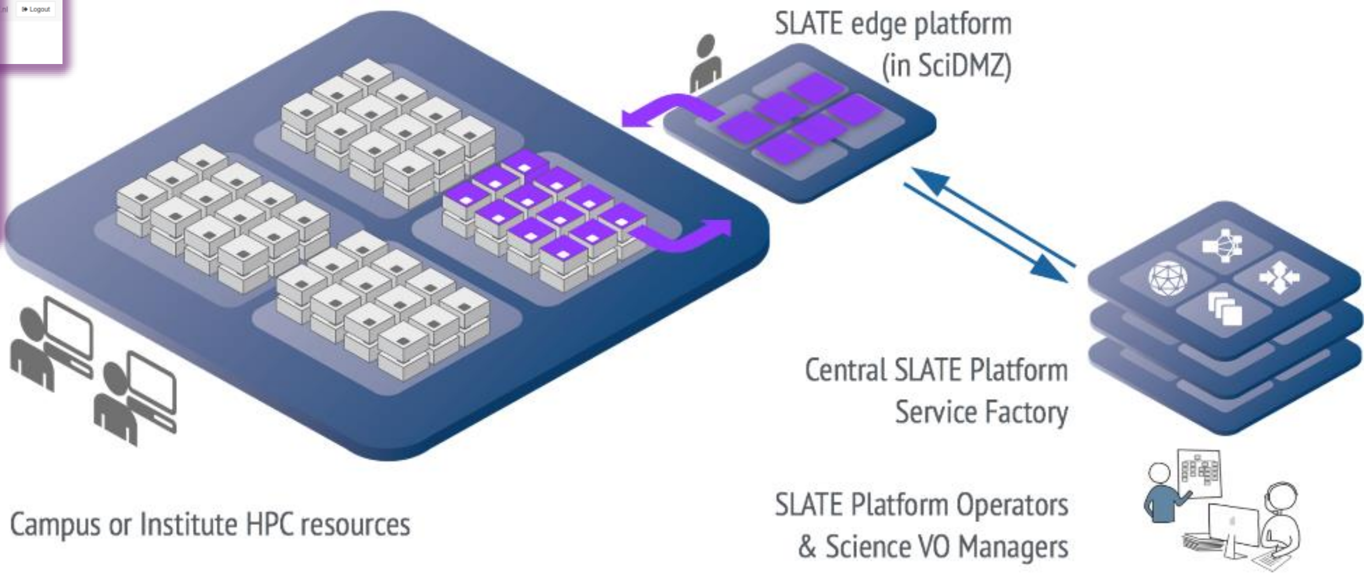

'alien containers' HPC integration - container computing, using curated application images

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

Beyond just technology: Analysis Facilities & Coffea Casa

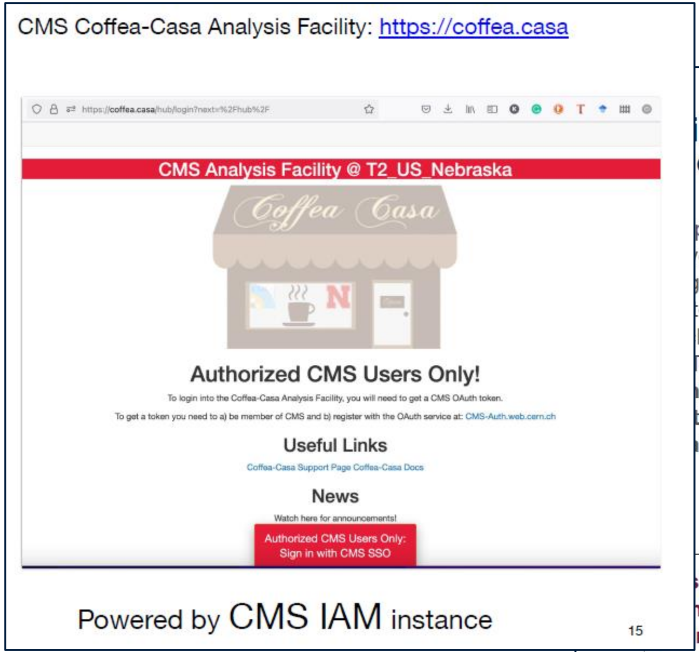

CMS Coffea-Casa Analysis Facility: https://coffea.casa

https://coffea.casa/hub/login?next=%2Fhub%2F
CMS Analysis Facility @ T2_US_Nebraska
Coffea Casa
N

Authorized CMS Users Only!

To login into the Coffea-Casa Analysis Facility, you will need to get a CMS OAuth token.

To get a token you need to a) be member of CMS and b) register with the OAuth service at: CMS-Auth.web.cern.ch

Useful Links

Coffea-Casa Support Page Coffea-Casa Docs

News

Watch here for announcements!

Authorized CMS Users Only:
Sign in with CMS SSO

Powered by CMS IAM instance

15

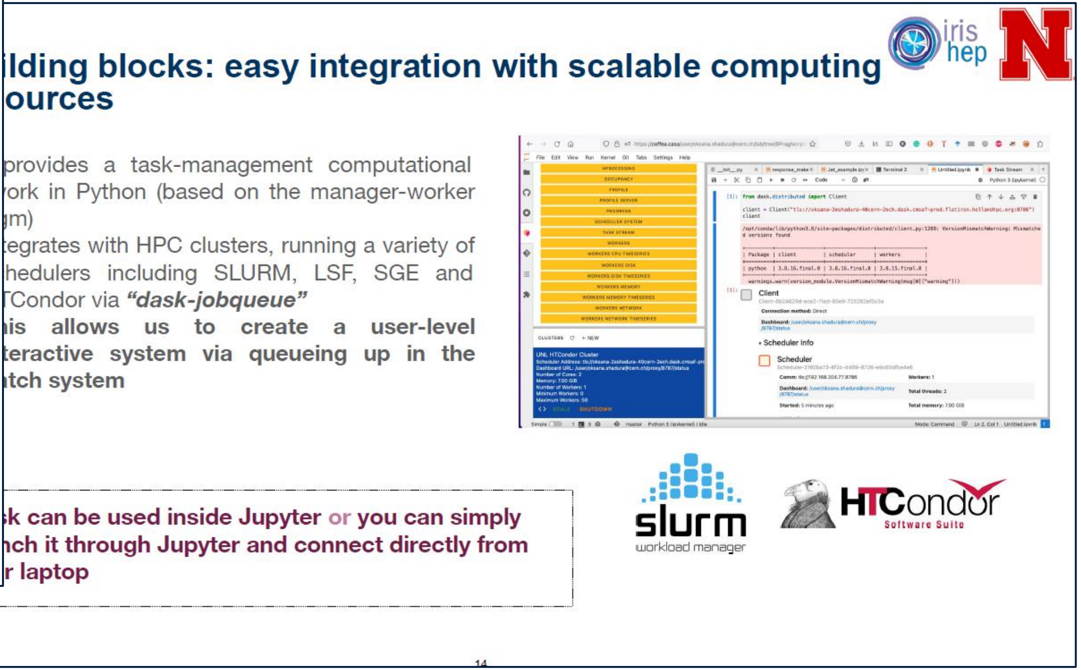

ilding blocks: easy integration with scalable computing
ources

provides a task-management computational
ork in Python (based on the manager-worker
m)
egrates with HPC clusters, running a variety of
hedulers including SLURM, LSF, SGE and
Condor via ***"dask-jobqueue"***
is allows us to create a user-level
teractive system via queueing up in the
tch system

UNL HTCondor Cluster
Scheduler Address: tls://oksana-2eshadura-40cern-2ech.dask.cmsaf-pr
Dashboard URL: /user/oksana.shadura@cern.ch/proxy/8787/status
Number of Cores: 2
Memory: 7.00 GiB
Number of Workers: 1
Minimum Workers: 0
Maximum Workers: 50

```
[1]: from dask.distributed import Client

     client = Client("tls://oksana-2eshadura-40cern-2ech.dask.cmsaf-prod.flatiron.hollandhpc.org:8786")
     client

/opt/conda/lib/python3.8/site-packages/distributed/client.py:1288: VersionMismatchWarning: Mismatche
d versions found

+---------+------------------+------------------+------------------+
| Package | client           | scheduler        | workers          |
+---------+------------------+------------------+------------------+
| python  | 3.8.16.final.0   | 3.8.16.final.0   | 3.8.15.final.0   |
+---------+------------------+------------------+------------------+
  warnings.warn(version_module.VersionMismatchWarning(msg[0]["warning"]))
```

[1]: Client

Client-8b24629d-ece2-11ed-80e9-725282ef5c3a

Connection method: Direct

Dashboard: /user/oksana.shadura@cern.ch/proxy/8787/status

- Scheduler Info

Scheduler

Scheduler-2160ba73-8f2c-4489-8726-e9c60dfbe4e6

Comm: tls://192.168.204.77:8786	Workers: 1
Dashboard: /user/oksana.shadura@cern.ch/proxy/8787/status	Total threads: 2
Started: 5 minutes ago	Total memory: 7.00 GiB

k can be used inside Jupyter or you can simply
nch it through Jupyter and connect directly from
r laptop

14

Images: Oksana Shadura et al (UNebraska Lincoln), Brian Bockelman (Morgridge Institute) at CHEP2023 https://indico.jlab.org/event/459/contributions/11610/

Community federated access to analysis

Analysis facility characteristics

- shared collaborative analysis
- data and compute access across all partners
- design for equitable access to global collaborations

The 'ESCAPE' ESFRIs are not the only ones

- AARC BPA design in EOSC & ESFRI clusters
- Netherlands: SRAM
- globally: CILogon, HPCI (JP)

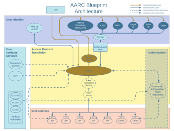

... *now extending AARC proxy model to **meshed collaborations***

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

AARC

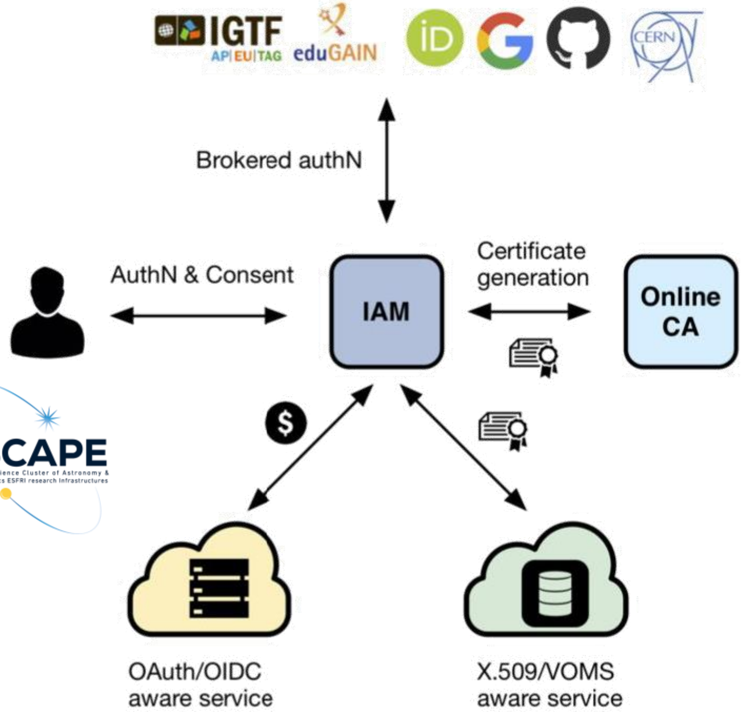

Indigo IAM structure diagram: Andrea Ceccanti et al. (INFN CNAF) ESCAPE IAM: https://projectescape.eu/ , Online CA: AARC RCauth CA, https://rcauth.eu/

EOSC: an ecosystem more than just services infrastructure

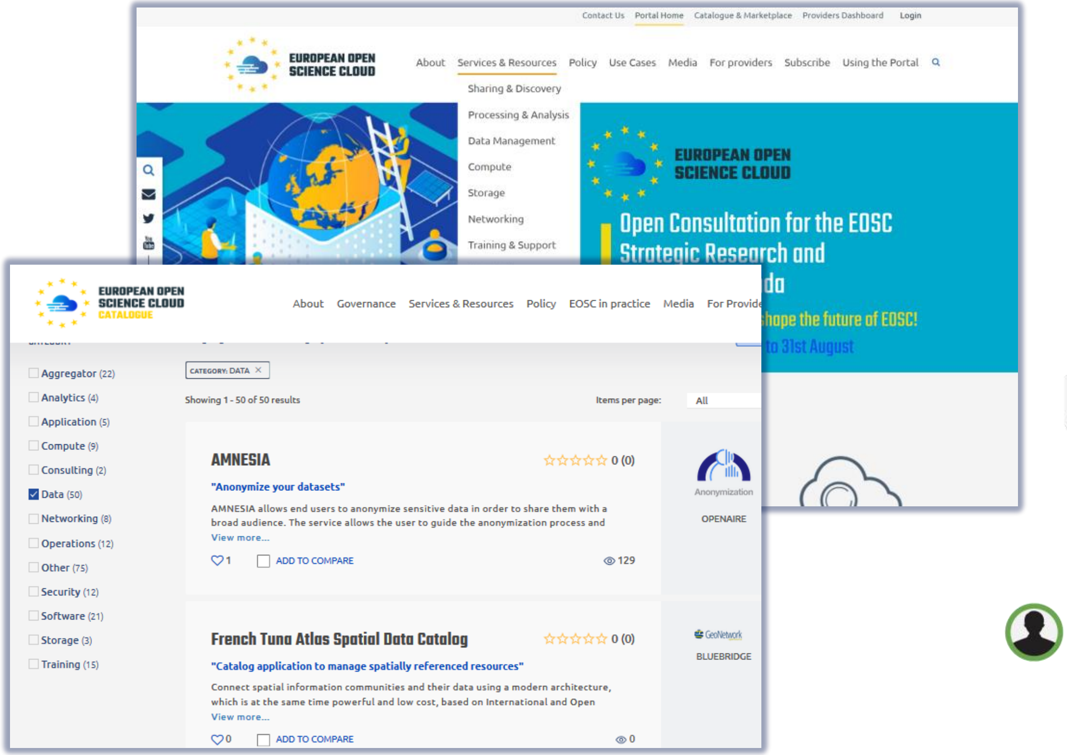

Contact Us | Portal Home | Catalogue & Marketplace | Providers Dashboard | Login

EUROPEAN OPEN SCIENCE CLOUD

About | Services & Resources | Policy | Use Cases | Media | For providers | Subscribe | Using the Portal

- Sharing & Discovery
- Processing & Analysis
- Data Management
- Compute
- Storage
- Networking
- Training & Support

EUROPEAN OPEN SCIENCE CLOUD

Open Consultation for the EOSC Strategic Research and ...da

...hape the future of EOSC!

...to 31st August

EUROPEAN OPEN SCIENCE CLOUD CATALOGUE

About | Governance | Services & Resources | Policy | EOSC in practice | Media | For Provide...

- □ Aggregator (22)
- □ Analytics (4)
- □ Application (5)
- □ Compute (9)
- □ Consulting (2)
- ☑ Data (50)
- □ Networking (8)
- □ Operations (12)
- □ Other (75)
- □ Security (12)
- □ Software (21)
- □ Storage (3)
- □ Training (15)

CATEGORY: DATA ×

Showing 1 - 50 of 50 results

Items per page: All

AMNESIA ☆☆☆☆☆ 0 (0)

"Anonymize your datasets"

AMNESIA allows end users to anonymize sensitive data in order to share them with a broad audience. The service allows the user to guide the anonymization process and

View more...

♡1 □ ADD TO COMPARE 129

Anonymization

OPENAIRE

French Tuna Atlas Spatial Data Catalog ☆☆☆☆☆ 0 (0)

"Catalog application to manage spatially referenced resources"

Connect spatial information communities and their data using a modern architecture, which is at the same time powerful and low cost, based on International and Open

View more...

♡0 □ ADD TO COMPARE 0

GeoNetwork

BLUEBRIDGE

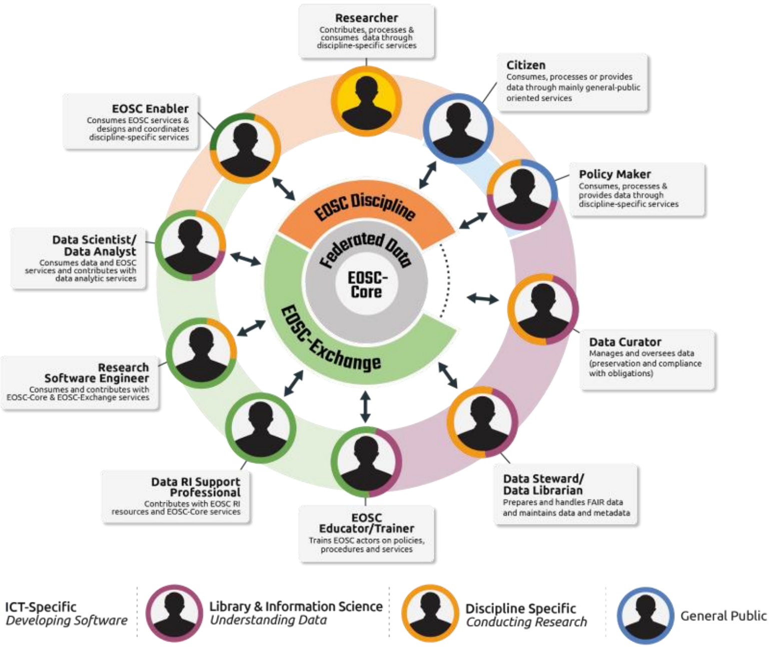
Researcher
Contributes, processes & consumes data through discipline-specific services
Citizen
Consumes, processes or provides data through mainly general-public oriented services
EOSC Enabler
Consumes EOSC services & designs and coordinates discipline-specific services
Policy Maker
Consumes, processes & provides data through discipline-specific services
Data Scientist/ Data Analyst
Consumes data and EOSC services and contributes with data analytic services
EOSC Discipline
Federated Data
EOSC-Core
EOSC-Exchange
Data Curator
Manages and oversees data (preservation and compliance with obligations)
Research Software Engineer
Consumes and contributes with EOSC-Core & EOSC-Exchange services
Data RI Support Professional
Contributes with EOSC RI resources and EOSC-Core services
EOSC Educator/Trainer
Trains EOSC actors on policies, procedures and services
Data Steward/ Data Librarian
Prepares and handles FAIR data and maintains data and metadata
ICT-Specific
Developing Software
Library & Information Science
Understanding Data
Discipline Specific
Conducting Research
General Public

Circle diagram from Ignacio Blanquer's ISGC 2022 keynote, Digital Skills for FAIR and open science: doi.org/10.2777/59065; EOSC Portal (https://www.eosc-portal.eu/) by EOSChub

The EOSC ecosystem – core and an ‘exchange’

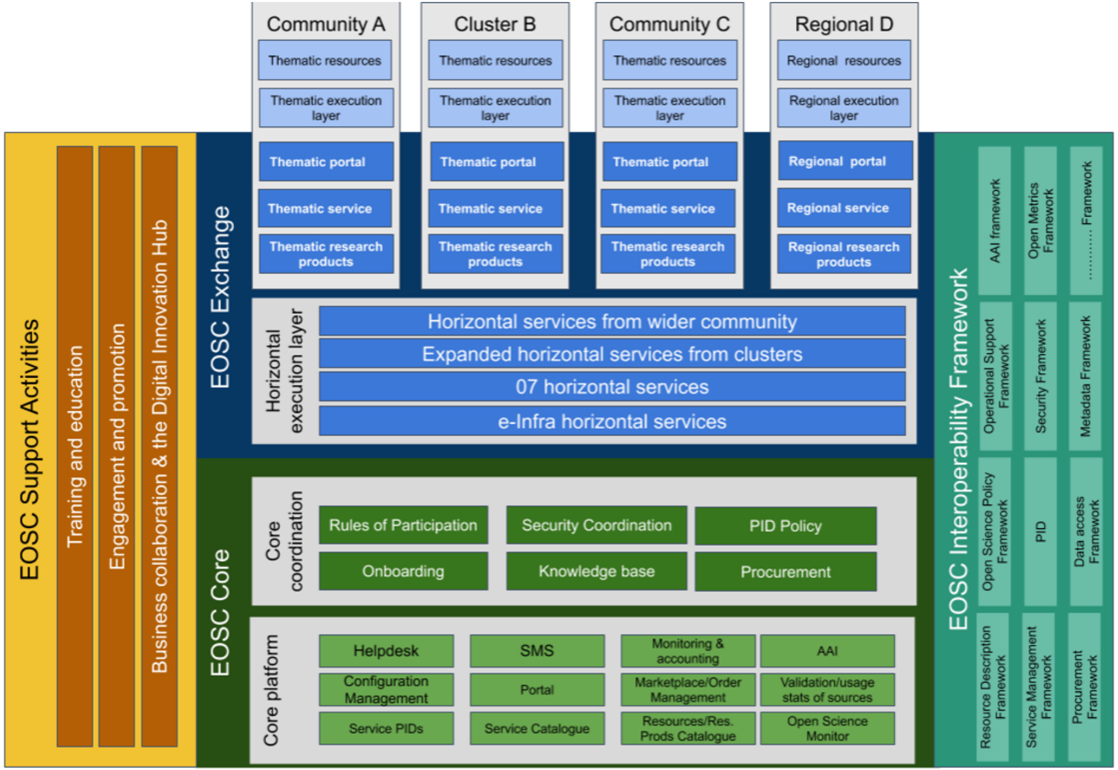

EOSC Future

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

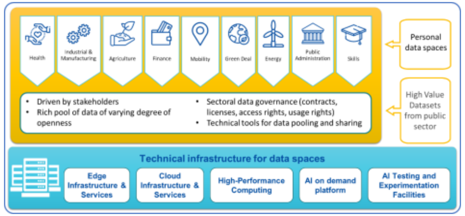

Health
Industrial & Manufacturing
Agriculture
Finance
Mobility
Green Deal
Energy
Public Administration
Skills
Personal data spaces
High Value Datasets from public sector
Technical infrastructure for data spaces
Edge Infrastructure & Services
Cloud Infrastructure & Services
High-Performance Computing
AI on demand platform
AI Testing and Experimentation Facilities

- Driven by stakeholders
- Rich pool of data of varying degree of openness
- Sectoral data governance (contracts, licenses, access rights, usage rights)
- Technical tools for data pooling and sharing

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

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

both in *thematic* and in *horizontal* e-Infrastructures

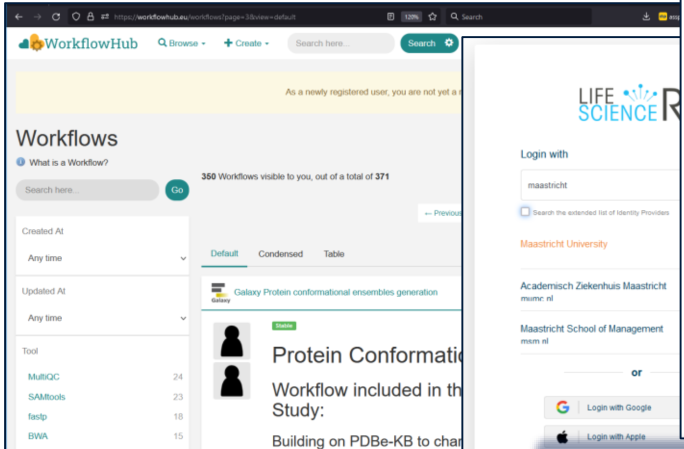

WorkflowHub

Browse

Create

Search here...

Search

As a newly registered user, you are not yet a

Workflows

What is a Workflow?

350 Workflows visible to you, out of a total of 371

Created At: Any time

Updated At: Any time

Tool	
MultiQC	24
SAMtools	23
fastp	18
BWA	15
Bowtie 2	11
Cutadapt	10

Default | Condensed | Table

Galaxy Protein conformational ensembles generation

Stable

Protein Conformati

Workflow included in th Study:

Building on PDBe-KB to char native proteins

This tutorial aims to illustrate the process of generating pr

LIFE SCIENCE R

Login with

maastricht

□ Search the extended list of Identity Providers

Maastricht University

Academisch Ziekenhuis Maastricht
mumc.nl

Maastricht School of Management
msm.nl

or

Login with Google

Login with Apple

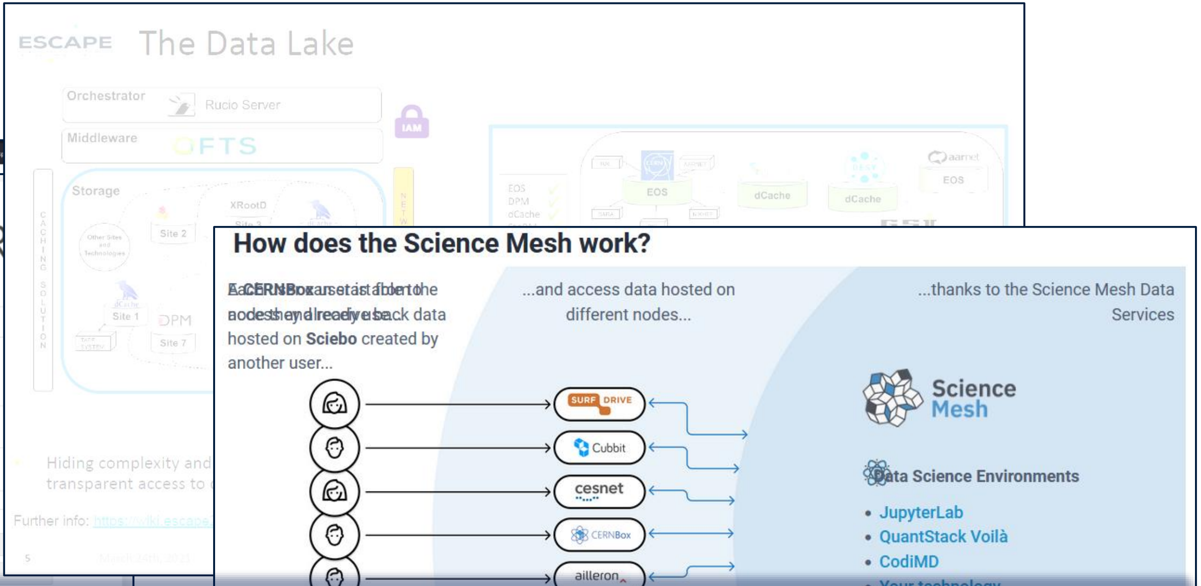

ESCAPE The Data Lake

Orchestrator: Rucio Server

Middleware: FTS

Storage

Hiding complexity and transparent access to

Further info: https://wiki.escape

How does the Science Mesh work?

EaCERNBoxarsetastabtetothe aodesthey already useack data hosted on Sciebo created by another user...

...and access data hosted on different nodes...

...thanks to the Science Mesh Data Services

Data Science Environments

- JupyterLab
- QuantStack Voilà
- CodiMD
- Your technology
- OnlyOffice

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

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!

Distributed collaborative ICT instrumentation, *a more technical example*

Credential translation in the AARC BPA

… building RCauth.eu

Leveraging federation and collaboration

for ubiquitous research credentials

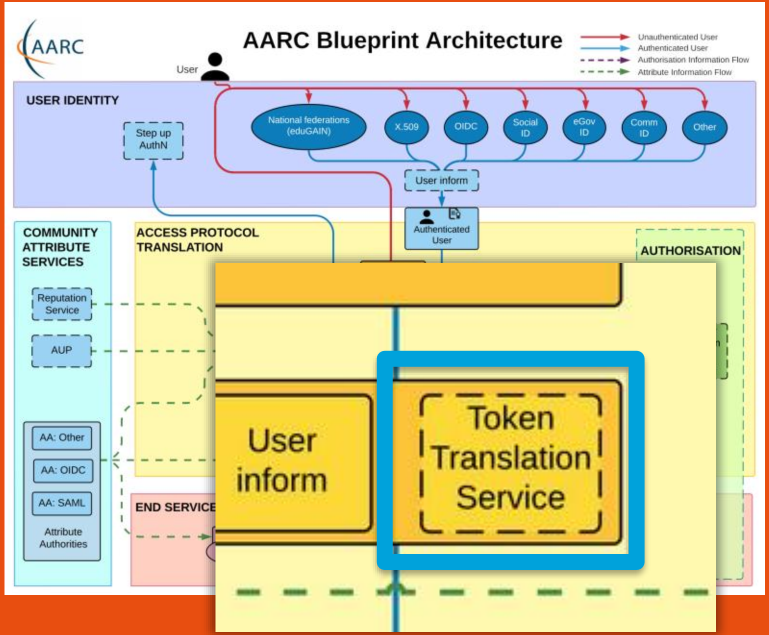

Bridges and Token Translation Services

TCS - for users that manage to grasp the idea

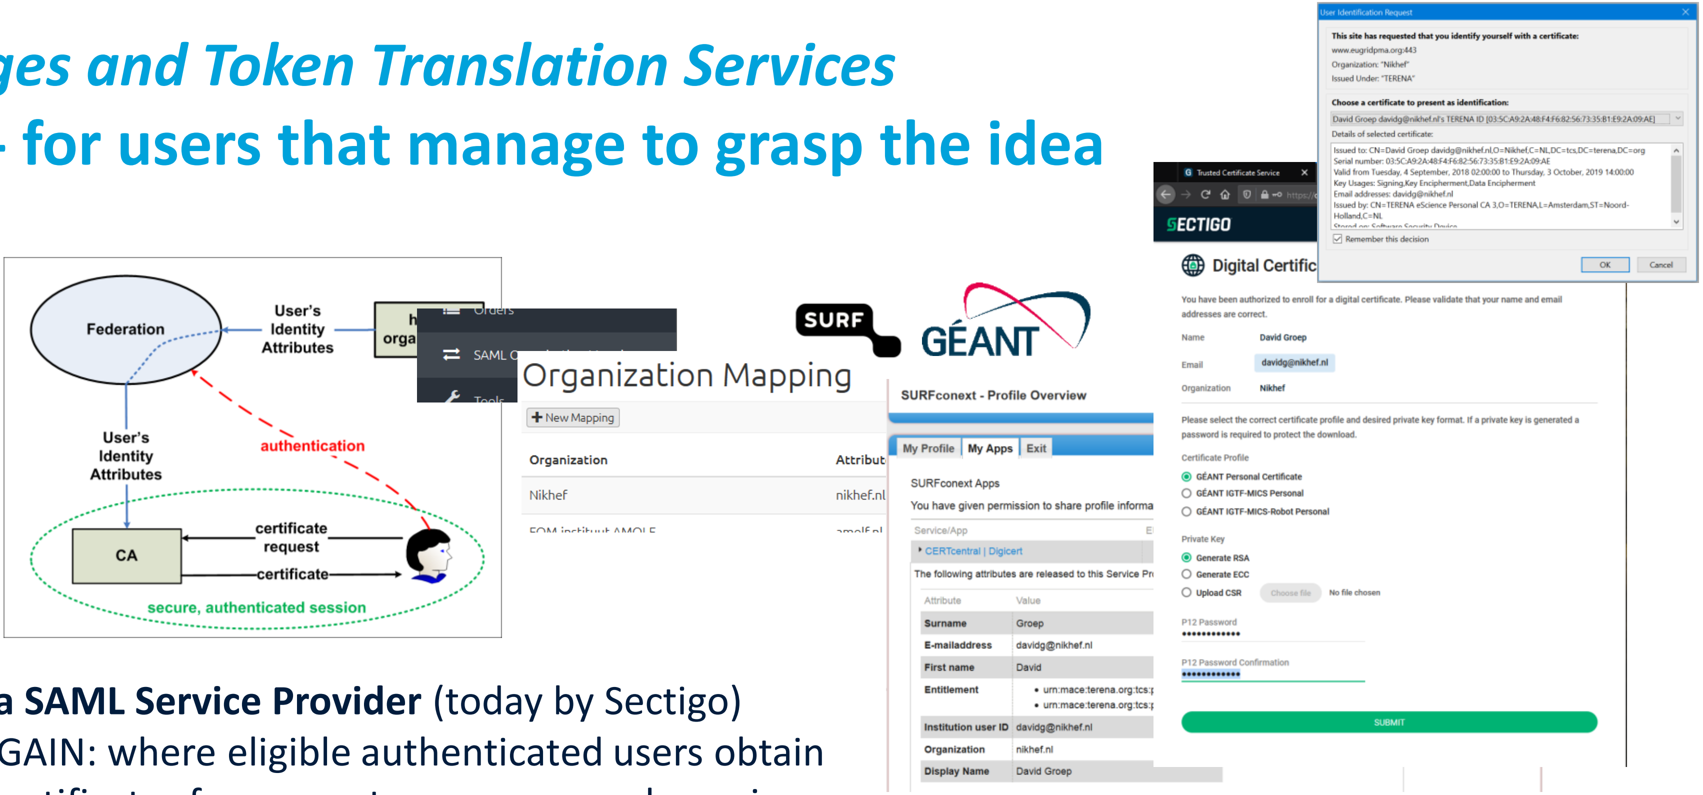
Federation
User's Identity Attributes
h
orga
User's Identity Attributes
authentication
CA
certificate request
certificate
secure, authenticated session
Orders
SAML C
Tools
SURF
GÉANT
Trusted Certificate Service
SECTIGO

Organization Mapping

+ New Mapping

Organization	Attribut
Nikhef	nikhef.nl
FOM instituut AMOLF	amolf.nl

SURFconext - Profile Overview

My Profile | My Apps | Exit

SURFconext Apps

You have given permission to share profile informa

Service/App

‣ CERTcentral | Digicert

The following attributes are released to this Service Pr

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

Digital Certific

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 Choose file No file chosen

P12 Password ••••••••••••

P12 Password Confirmation ••••••••••••

SUBMIT

User Identification Request

This site has requested that you identify yourself with a certificate:

www.eugridpma.org:443

Organization: "Nikhef"

Issued Under: "TERENA"

Choose a certificate to present as identification:

David Groep davidg@nikhef.nl's TERENA ID [03:5C:A9:2A:48:F4:F6:82:56:73:35:B1:E9:2A:09:AE]

Details of selected certificate:

Issued to: CN=David Groep davidg@nikhef.nl,O=Nikhef,C=NL,DC=tcs,DC=terena,DC=org
Serial number: 03:5C:A9:2A:48:F4:F6:82:56:73:35:B1:E9:2A:09:AE
Valid from Tuesday, 4 September, 2018 02:00:00 to Thursday, 3 October, 2019 14:00:00
Key Usages: Signing,Key Encipherment,Data Encipherment
Email addresses: davidg@nikhef.nl
Issued by: CN=TERENA eScience Personal CA 3,O=TERENA,L=Amsterdam,ST=Noord-Holland,C=NL
Stored on: Software Security Device

☑ Remember this decision

OK Cancel

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

user facing | hidden back-end

Community Science Portal

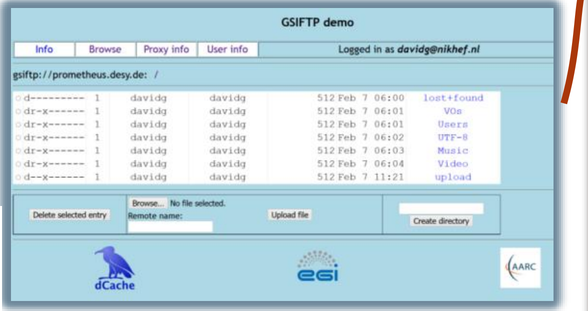
GSIFTP demo

Logged in as ***davidg@nikhef.nl***

gsiftp://prometheus.desy.de: /

```
○ d---------  1   davidg   davidg   512 Feb  7 06:00  lost+found
○ dr-x------  1   davidg   davidg   512 Feb  7 06:01  VOs
○ dr-x------  1   davidg   davidg   512 Feb  7 06:01  Users
○ dr-x------  1   davidg   davidg   512 Feb  7 06:02  UTF-8
○ dr-x------  1   davidg   davidg   512 Feb  7 06:03  Music
○ dr-x------  1   davidg   davidg   512 Feb  7 06:04  Video
○ d--x------  1   davidg   davidg   512 Feb  7 11:21  upload
```

No file selected.

Remote name:

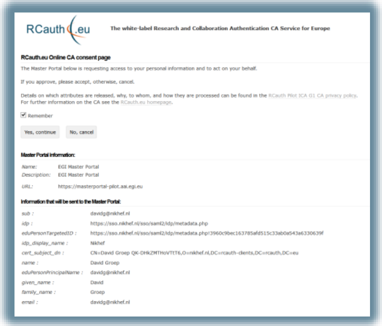
RCauth.eu — The white-label Research and Collaboration Authentication CA Service for Europe

RCauth.eu Online CA consent page

The Master Portal below is requesting access to your personal information and to act on your behalf.

If you approve, please accept, otherwise, cancel.

Details on which attributes are released, why, to whom, and how they are processed can be found in the RCauth Pilot ICA G1 CA privacy policy. For further information on the CA see the RCauth.eu homepage.

☑ Remember

Master Portal Information:

Name:	EGI Master Portal
Description:	EGI Master Portal
URL:	https://masterportal-pilot.aai.egi.eu

Information that will be sent to the Master Portal:

sub :	davidg@nikhef.nl
idp :	https://sso.nikhef.nl/sso/saml2/idp/metadata.php
eduPersonTargetedID :	https://sso.nikhef.nl/sso/saml2/idp/metadata.php!3960c9bec163785afd515c33ab0a543a6330639f
idp_display_name :	Nikhef
cert_subject_dn :	CN=David Groep QK-DHkZMTHoVTtT6,O=nikhef.nl,DC=rcauth-clients,DC=rcauth,DC=eu
name :	David Groep
eduPersonPrincipalName :	davidg@nikhef.nl
given_name :	David
family_name :	Groep
email :	davidg@nikhef.nl

IGTF accredited PKIX Authority

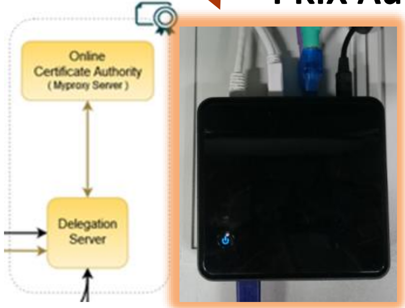

Infrastructure Master Portal Credential Store

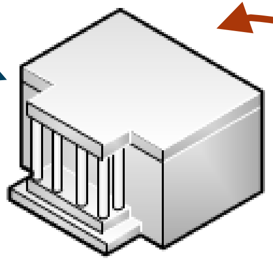

User Home Org *or Infrastructure IdP*

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

REFEDS R&S
Sirtfi Trust

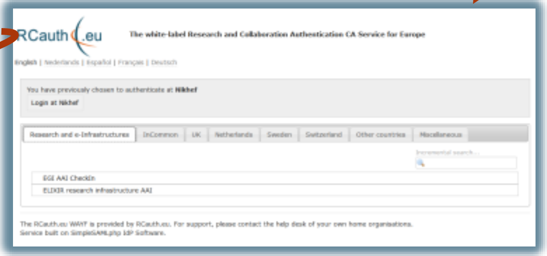
RCauth.eu — The white-label Research and Collaboration Authentication CA Service for Europe

English | Nederlands | Español | Français | Deutsch

You have previously chosen to authenticate at **Nikhef**

EGI AAI CheckIn

ELIXIR research infrastructure AAI

The RCauth.eu WAYF is provided by RCauth.eu. For support, please contact the help desk of your own home organisations.
Service built on SimpleSAMLphp IdP Software.

Policy Filtering WAYF to eduGAIN

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

Our Registration Authorities: the Federated IdPs

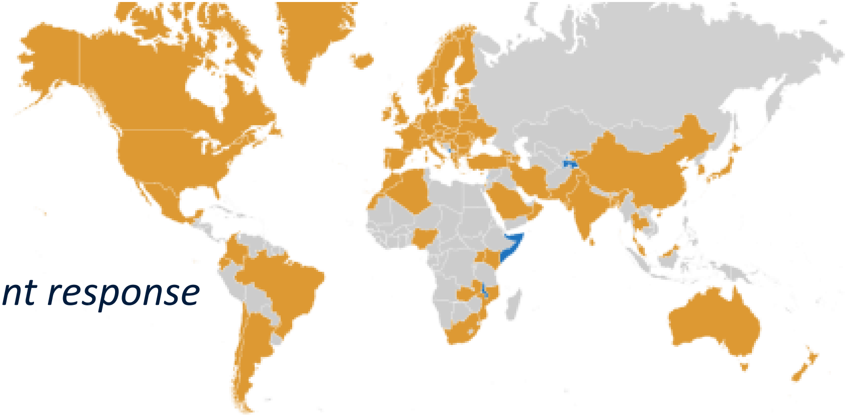

Distributed RAs: the *eligible IdPs*

- connected through a federation, primarily: the ensemble of IdPs in eduGAIN that meet the policy requirements of this CA
- since authN and authZ are split, need is for *non-reassigned identifier* and *point-in-time incident response*

eligible *applicants* are then all affiliated to an RA

Three eligibility models

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

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

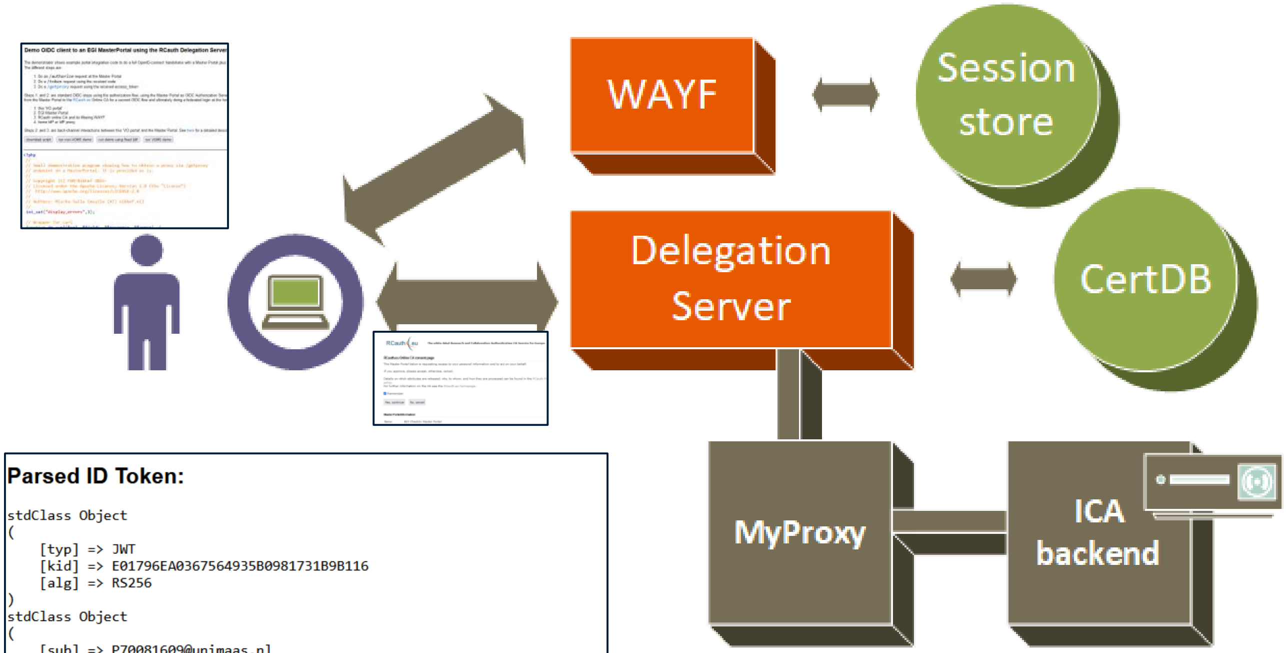
Demo OIDC client to an EGI MasterPortal using the RCauth Delegation Server
WAYF
Session store
Delegation Server
CertDB
MyProxy
ICA backend
RCauth.eu

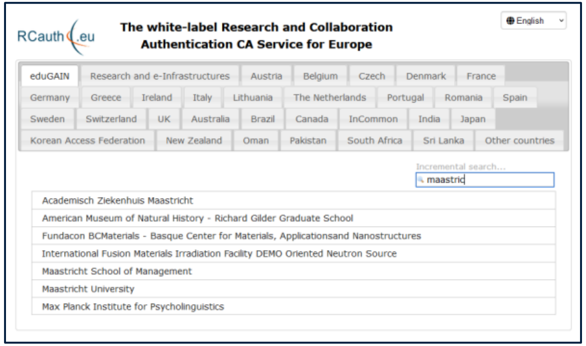

The white-label Research and Collaboration Authentication CA Service for Europe

- Academisch Ziekenhuis Maastricht
- American Museum of Natural History - Richard Gilder Graduate School
- Fundacon BCMaterials - Basque Center for Materials, Applicationsand Nanostructures
- International Fusion Materials Irradiation Facility DEMO Oriented Neutron Source
- Maastricht School of Management
- Maastricht University
- Max Planck Institute for Psycholinguistics

Parsed ID Token:

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

Proxy information:

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

With a single, yet fully compliant, 'Heath Robinson' CA

The locally-highly-available RCauth at Nikhef Amsterdam

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

Eliminated SPOFs first using ‘local HA’

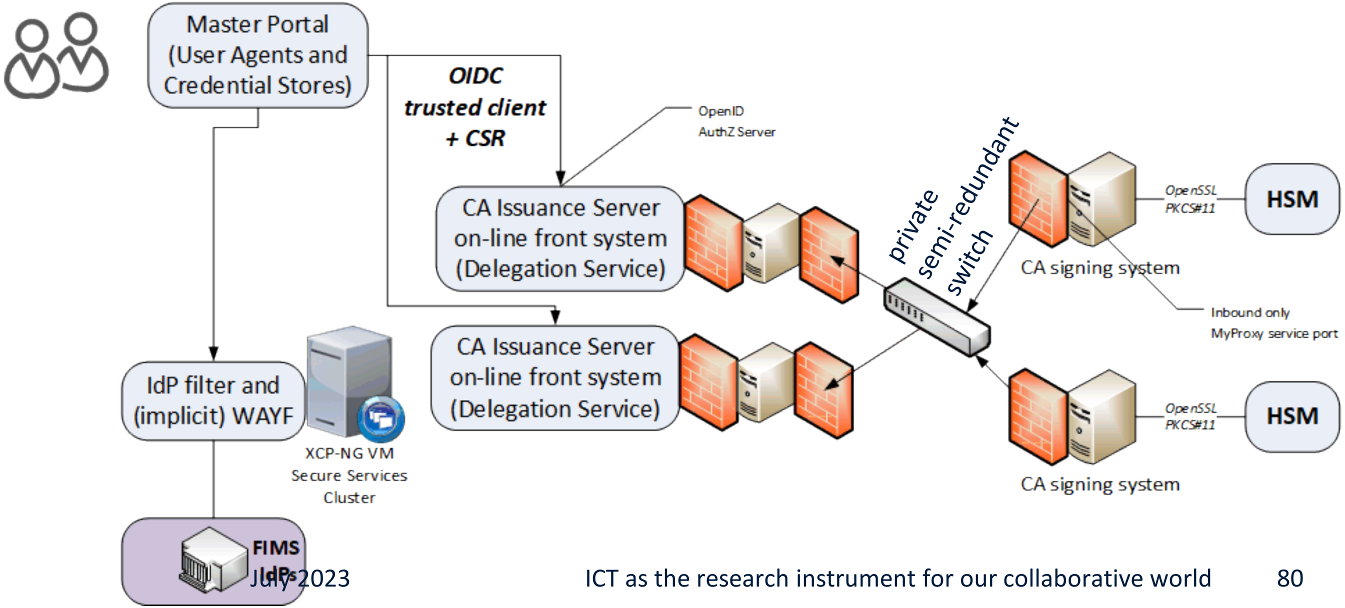

Master Portal (User Agents and Credential Stores)
OIDC trusted client + CSR
OpenID AuthZ Server
CA Issuance Server on-line front system (Delegation Service)
CA Issuance Server on-line front system (Delegation Service)
private semi-redundant switch
OpenSSL PKCS#11
HSM
CA signing system
Inbound only MyProxy service port
OpenSSL PKCS#11
HSM
CA signing system
IdP filter and (implicit) WAYF
XCP-NG VM Secure Services Cluster
FIMS IdPs

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

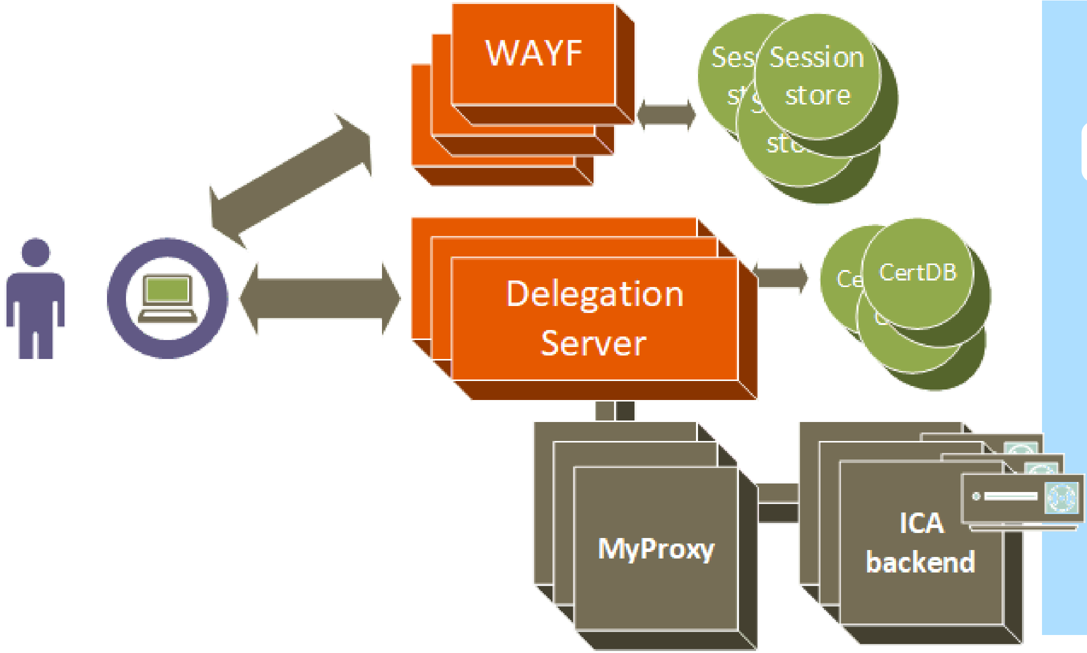

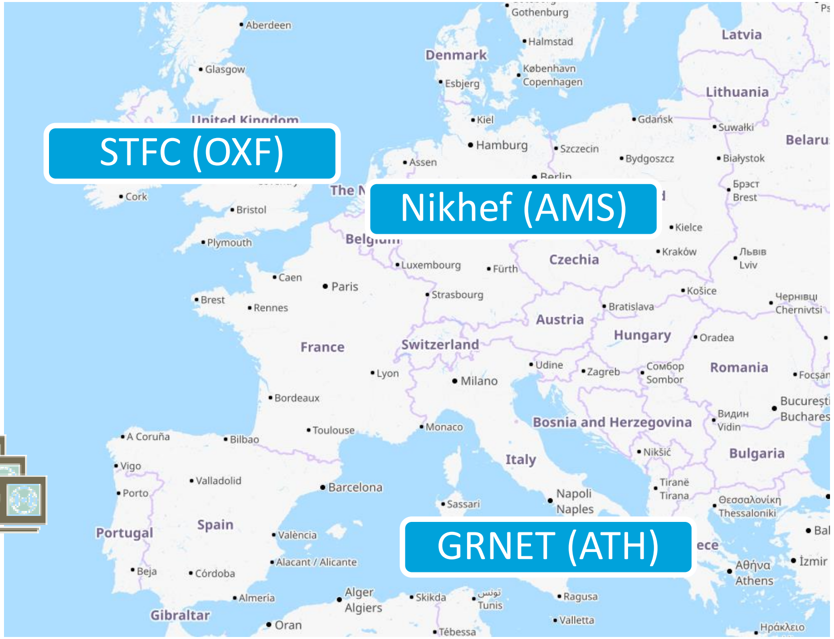

Since we do not like SPOFs …

Distributed High Availability setup

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

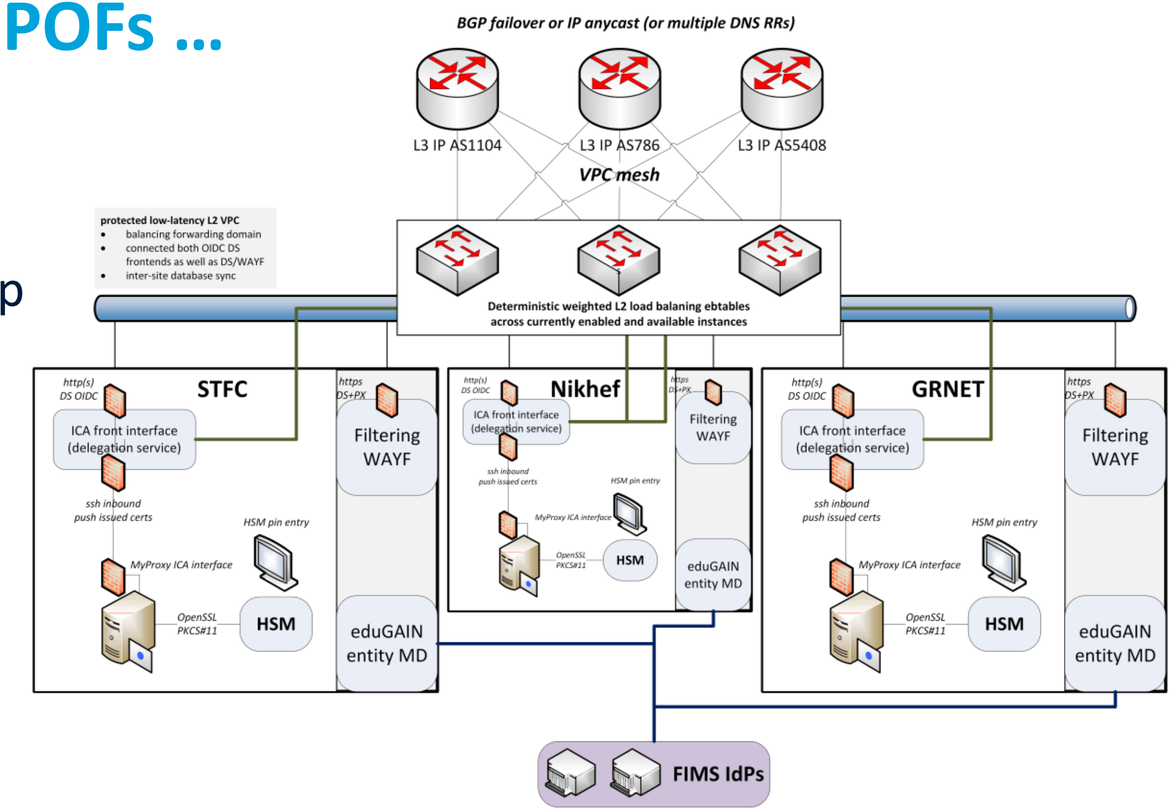

BGP failover or IP anycast (or multiple DNS RRs)
L3 IP AS1104
L3 IP AS786
L3 IP AS5408
VPC mesh
protected low-latency L2 VPC
balancing forwarding domain
connected both OIDC DS frontends as well as DS/WAYF
inter-site database sync
Deterministic weighted L2 load balaning ebtables across currently enabled and available instances
STFC
http(s) DS OIDC
ICA front interface (delegation service)
ssh inbound push issued certs
HSM pin entry
MyProxy ICA interface
OpenSSL PKCS#11
HSM
https DS+PX
Filtering WAYF
eduGAIN entity MD
Nikhef
http(s) DS OIDC
ICA front interface (delegation service)
ssh inbound push issued certs
HSM pin entry
MyProxy ICA interface
OpenSSL PKCS#11
HSM
https DS+PX
Filtering WAYF
eduGAIN entity MD
GRNET
http(s) DS OIDC
ICA front interface (delegation service)
ssh inbound push issued certs
HSM pin entry
MyProxy ICA interface
OpenSSL PKCS#11
HSM
https DS+PX
Filtering WAYF
eduGAIN entity MD
FIMS IdPs

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

A *transparent* multi-site setup is needed for the user

User

- connects to HA proxy at **{wayf,pilot-ica-g1}.rcauth.eu**
- HA proxy sends users to **“closest”** working service
- primarily **forward to its own DS** when available

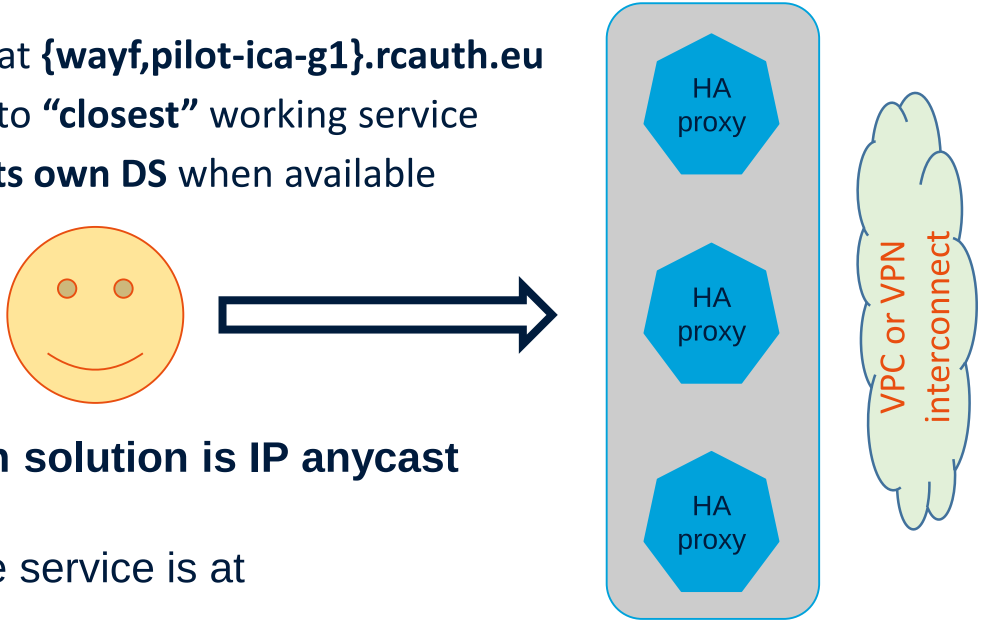

If a HA loses its backend DS, can still route to another DS over VPC/VPN backend

Straightforward proven solution is IP anycast

wherever the user is, the service is at

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

selected imagery: Mischa Sallé, Jens Jensen, Nicolas Liampotis

Anycast: when the same place exists many times

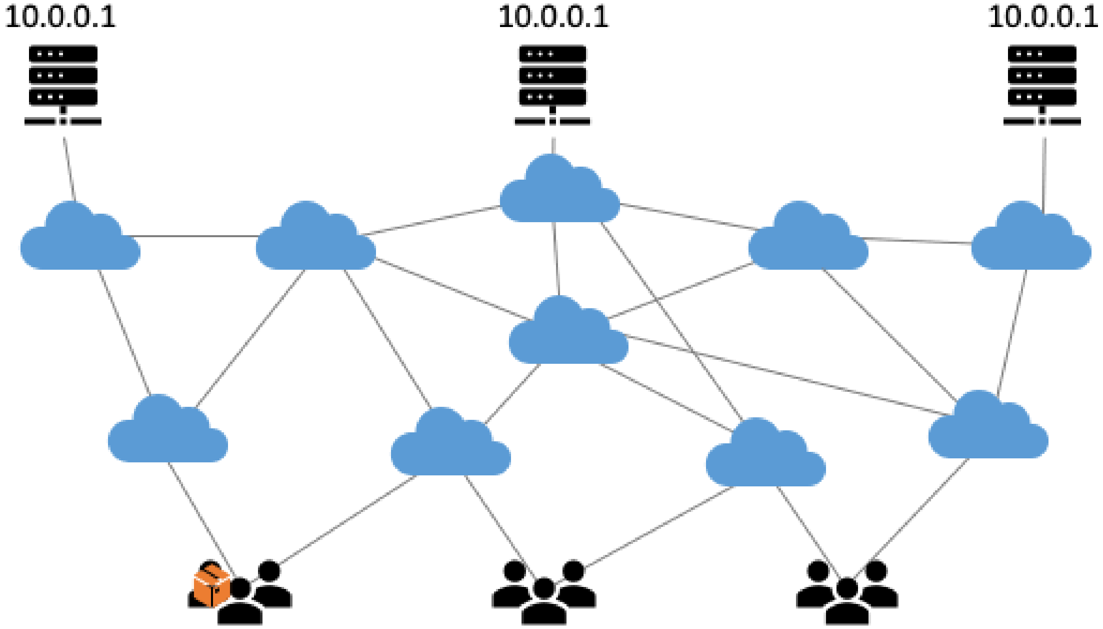

So we used

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

and some monitoring

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

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

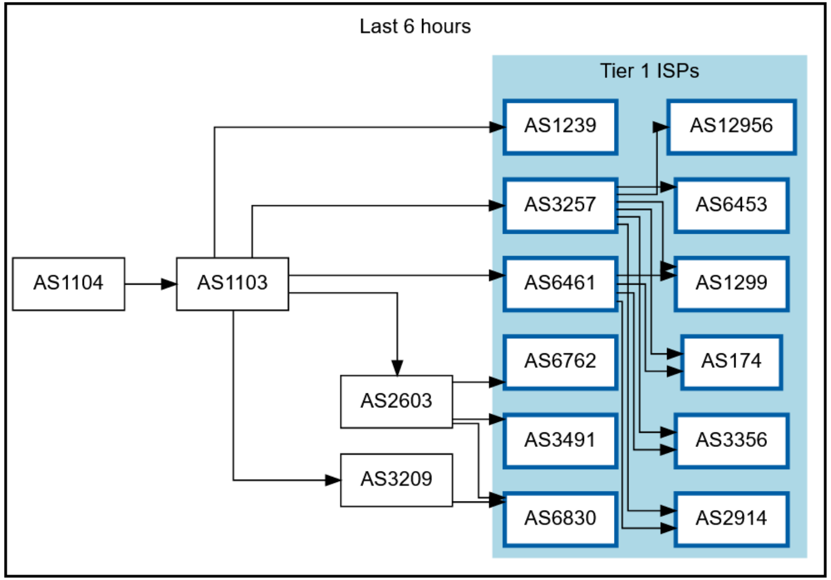

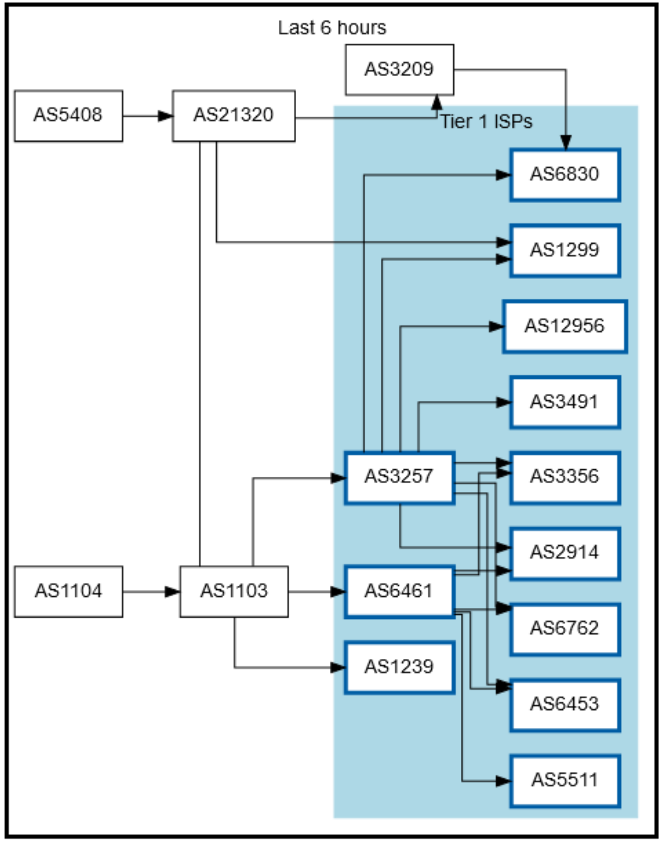

route maps: bgp.tools for 2a07:8504:1a0::/48 – IPv4 for 145.116.216.0/24 is similar – imagery from November 2022

And you get reasonable load balancing in Europe for free

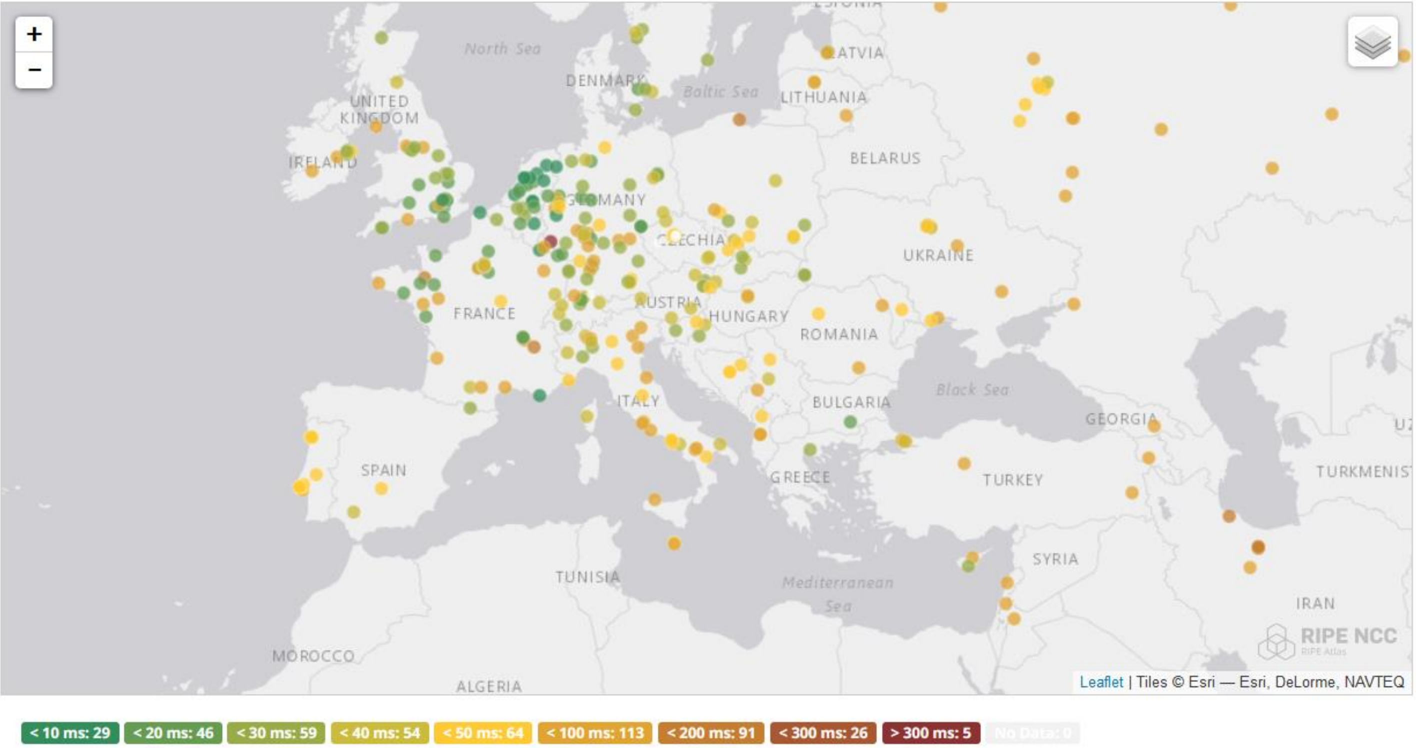

map: RIPE NCC RIPE Atlas - 500 probes, distributed across Europe (https://atlas.ripe.net/measurements/50949024/)

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

```
[root@kwark ~]# traceroute -IA 145.116.216.1
traceroute to 145.116.216.1 (145.116.216.1), 30 hops max, 60 byte packets
 1  cmbr.connected.by.freedominter.net
      (185.93.175.234) [AS206238]
 2  connected.by.freedom.nl
      (185.93.175.240) [AS206238]
 3  et-0-0-0-1002.core1.fi001.nl.freedomnet.nl
      (185.93.175.208) [AS206238]
 4  as1104.frys-ix.net (185.1.203.66) [*]
 5  parkwachter.nikhef.nl
      (192.16.186.141) [AS1104]
 6  gw-anyc-01.rcauth.eu
      (145.116.216.1) [AS786/AS5408/AS1104]
```

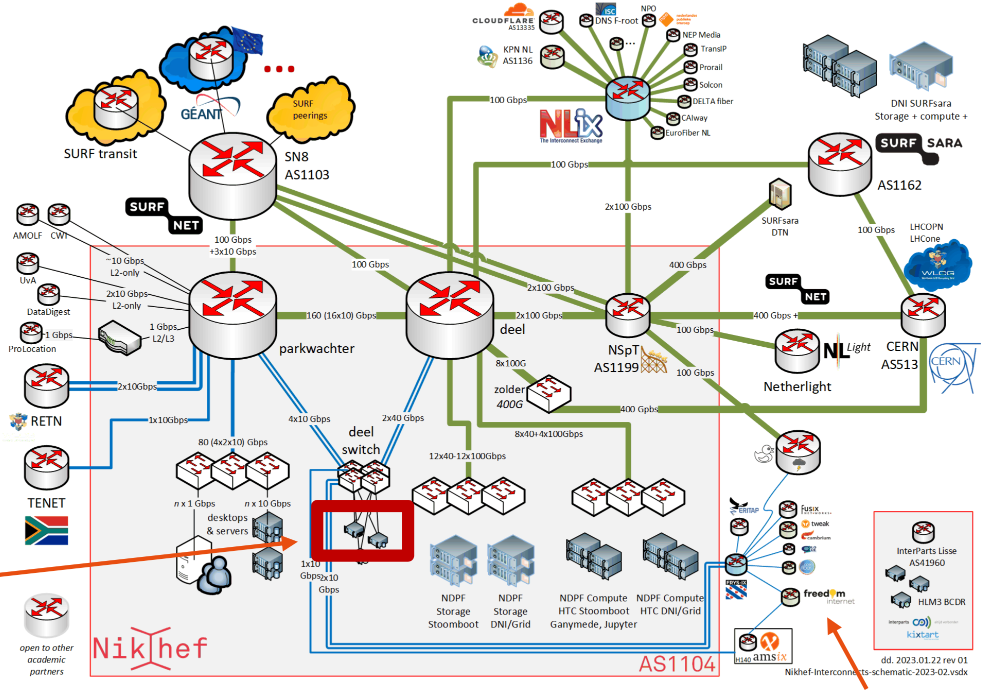

rcauth.eu HA proxy

me, at home

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

So can we discern a common pattern?

- Infrastructure is distributed, but that’s nothing truly ‘magic’
  - and *every* collaborating organization, university, and national lab is part of it and can do it

- Move complexity and volume requirements to the edge
  - the edge scales horizontally and scaling from 2+ is much easier than from 1→ 2

- Any central (network) components should be passive and as stateless as possible
  - research (and computing education) infrastructure performance ought to just be ‘a given’
  - any stateful device in the data path will block performant data transfers and reliability
  - although persistent storage obviously has to retain some state ☺

- Scaling *collaboration infrastructure, trust & identity,* and ***federation of expertise*** needed as much as we need scaling of our computing and networks

Diverse use cases, common vision

Supporting our mission on collaboration, Open Science,
and internationalization through scalable e-Infrastructures

ICT infrastructure landscape in the Netherlands

Using ICT as research instrumentation

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

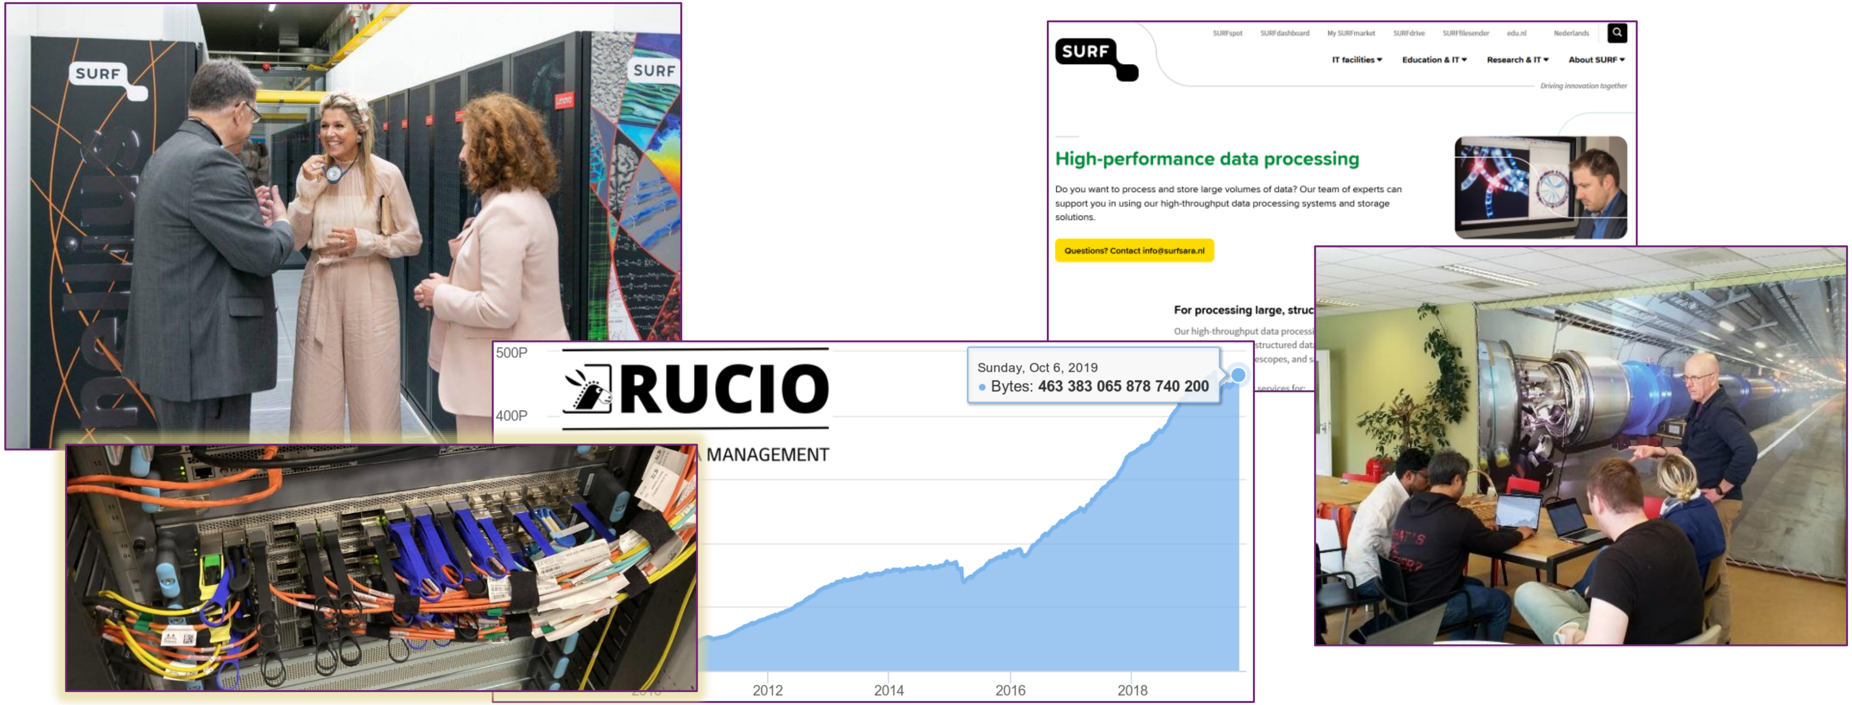

SURFspot SURFdashboard My SURFmarket SURFdrive SURFfilesender edu.nl Nederlands

IT facilities ▾ Education & IT ▾ Research & IT ▾ About SURF ▾

Driving innovation together

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@surfsara.nl

For processing large, struc

Our high-throughput data processi
structured dat
escopes, and s
services for:

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)

For example our HPC strategy: from local "T2" to European "T0"

Nikhef "Stoomboot"
Analysis Facility

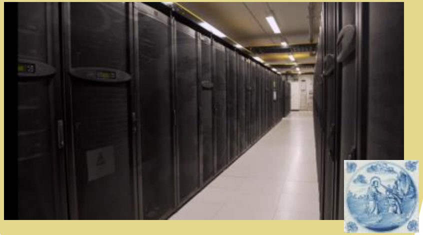

...

SURF National Infrastructure

PRACE

EuroHPC
Joint Undertaking

eosc

How to exploit
these unique systems?
access, expertise, and …
a long-term vision
on how research scales up

Photos: Nikhef NDPF, [illegible]t, SURF Data Repository, Snellius, SURF @ DigitalRealty
Steamship: Michael on Unsplash (https://unsplash.com/photos/944sDSMQ778), Tile: Nationaal Museum van Wereldculturen

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

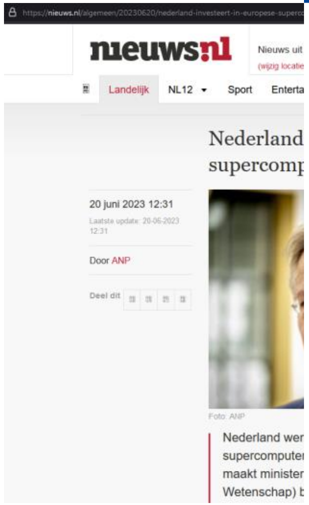

https://nieuws.nl/algemeen/20230620/nederland-investeert-in-europese-superco

nieuws.nl

Nieuws uit
(wijzig locatie

Nederland
supercomp

20 juni 2023 12:31
Laatste update: 20-06-2023 12:31

Door ANP

Deel dit

Foto: ANP

Nederland wer
supercompute
maakt minister
Wetenschap) b

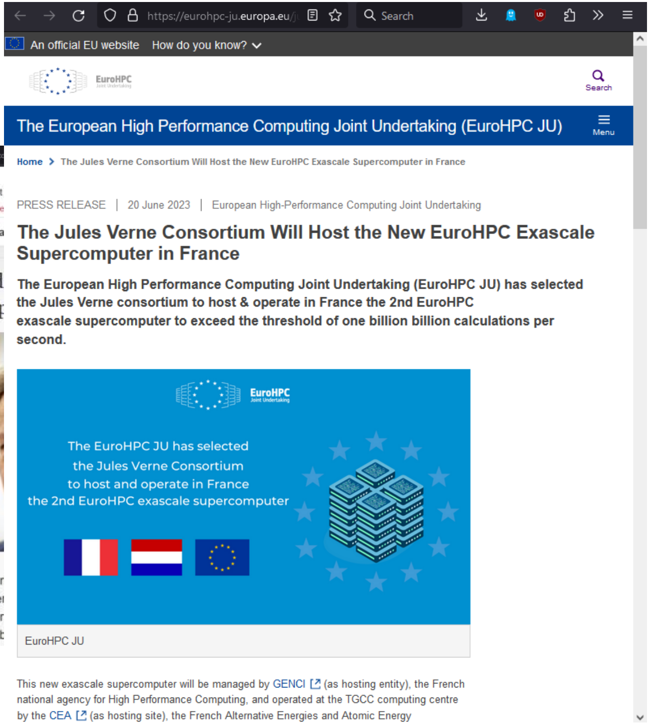

https://eurohpc-ju.europa.eu/

The European High Performance Computing Joint Undertaking (EuroHPC JU)

Home > The Jules Verne Consortium Will Host the New EuroHPC Exascale Supercomputer in France

PRESS RELEASE | 20 June 2023 | European High-Performance Computing Joint Undertaking

The Jules Verne Consortium Will Host the New EuroHPC Exascale Supercomputer in France

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 EuroHPC JU has selected the Jules Verne Consortium to host and operate in France the 2nd EuroHPC exascale supercomputer

EuroHPC JU

This new exascale supercomputer will be managed by GENCI (as hosting entity), the French national agency for High Performance Computing, and operated at the TGCC computing centre by the CEA (as hosting site), the French Alternative Energies and Atomic Energy

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)

Collaborative services are distributed and federated

Collaborative services are
spread across the research community

- logbooks with federated login from LIGO (LVK collaboration) for ET pathfinder and IGWN
- 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

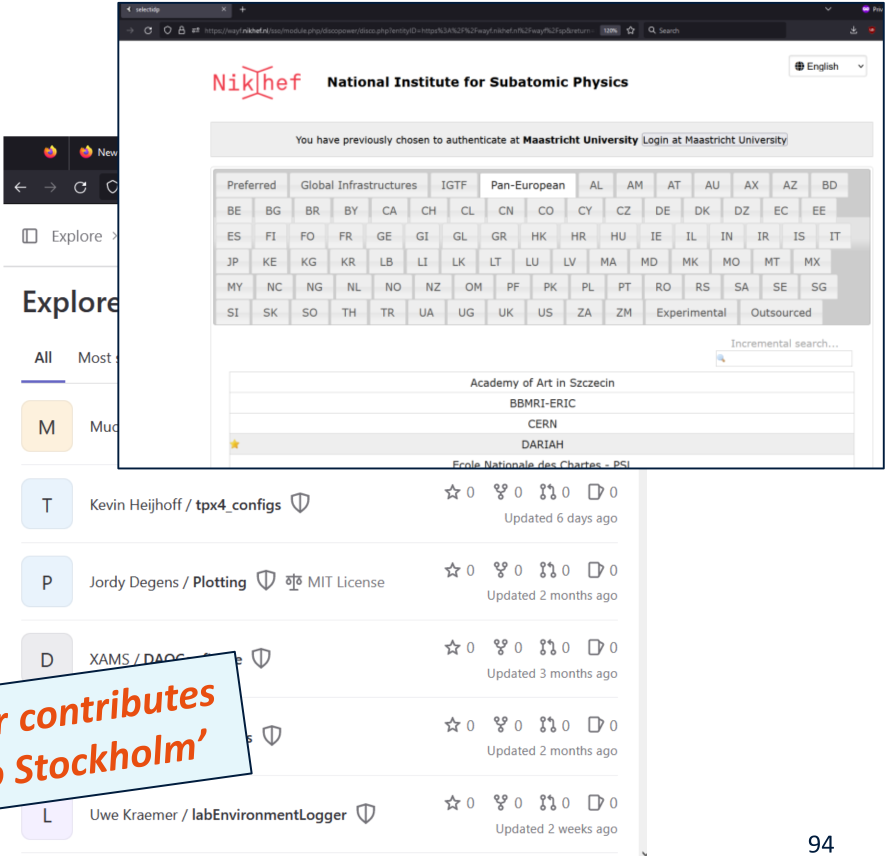
selectidp
https://wayf.nikhef.nl/sso/module.php/discopower/disco.php?entityID=https%3A%2F%2Fwayf.nikhef.nl%2Fwayf%2Fsp&return=
120%
Search
English
Nikhef
Preferred
Global Infrastructures
IGTF
Pan-European
AL AM AT AU AX AZ BD
BE BG BR BY CA CH CL CN CO CY CZ DE DK DZ EC EE
ES FI FO FR GE GI GL GR HK HR HU IE IL IN IR IS IT
JP KE KG KR LB LI 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 Outsourced
Incremental search...
New
Explore
All
Most

National Institute for Subatomic Physics

You have previously chosen to authenticate at **Maastricht University** Login at Maastricht University

Academy of Art in Szczecin
BBMRI-ERIC
CERN
DARIAH
Ecole Nationale des Chartes - PSL

Explore

M Mud

T Kevin Heijhoff / **tpx4_configs** ☆ 0 0 0 0 Updated 6 days ago

P Jordy Degens / **Plotting** MIT License ☆ 0 0 0 0 Updated 2 months ago

D XAMS / DAQC... ☆ 0 0 0 0 Updated 3 months ago

☆ 0 0 0 0 Updated 2 months ago

L Uwe Kraemer / **labEnvironmentLogger** ☆ 0 0 0 0 Updated 2 weeks ago

'every partner contributes to the trip to Stockholm'

So: ICT Digital Competences for research

Integrale aanpak voor digitalisering in de wetenschap

Uitvoeringsplan investeringen digitale onderzoeksinfrastructuur

NWO

- need for a federated networked scheme for data, computing (and expertise) remains as relevant today as it was in 2017
- LDCC role as “*knooppunt in een gefedereerd netwerk voor data, computing en expertise*” has not received much attention in terms of infrastructure
- expertise bundling and development of “Tier-2” facilities in national landscape is institutional responsibility, strengthening research support
- using national funding also means: be open to national collaboration, and ensuring the 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. SRAM, eduid.nl, and MyAcademicID

Collaboration 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

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 beyond just UM
  - sustainable research software, collaborative spaces with *global* partners,
    SRAM, eduGAIN & EOSC federated access, ubiquitous access to *external R&S* services

… since some things are fun, but not quite *that* scalable …

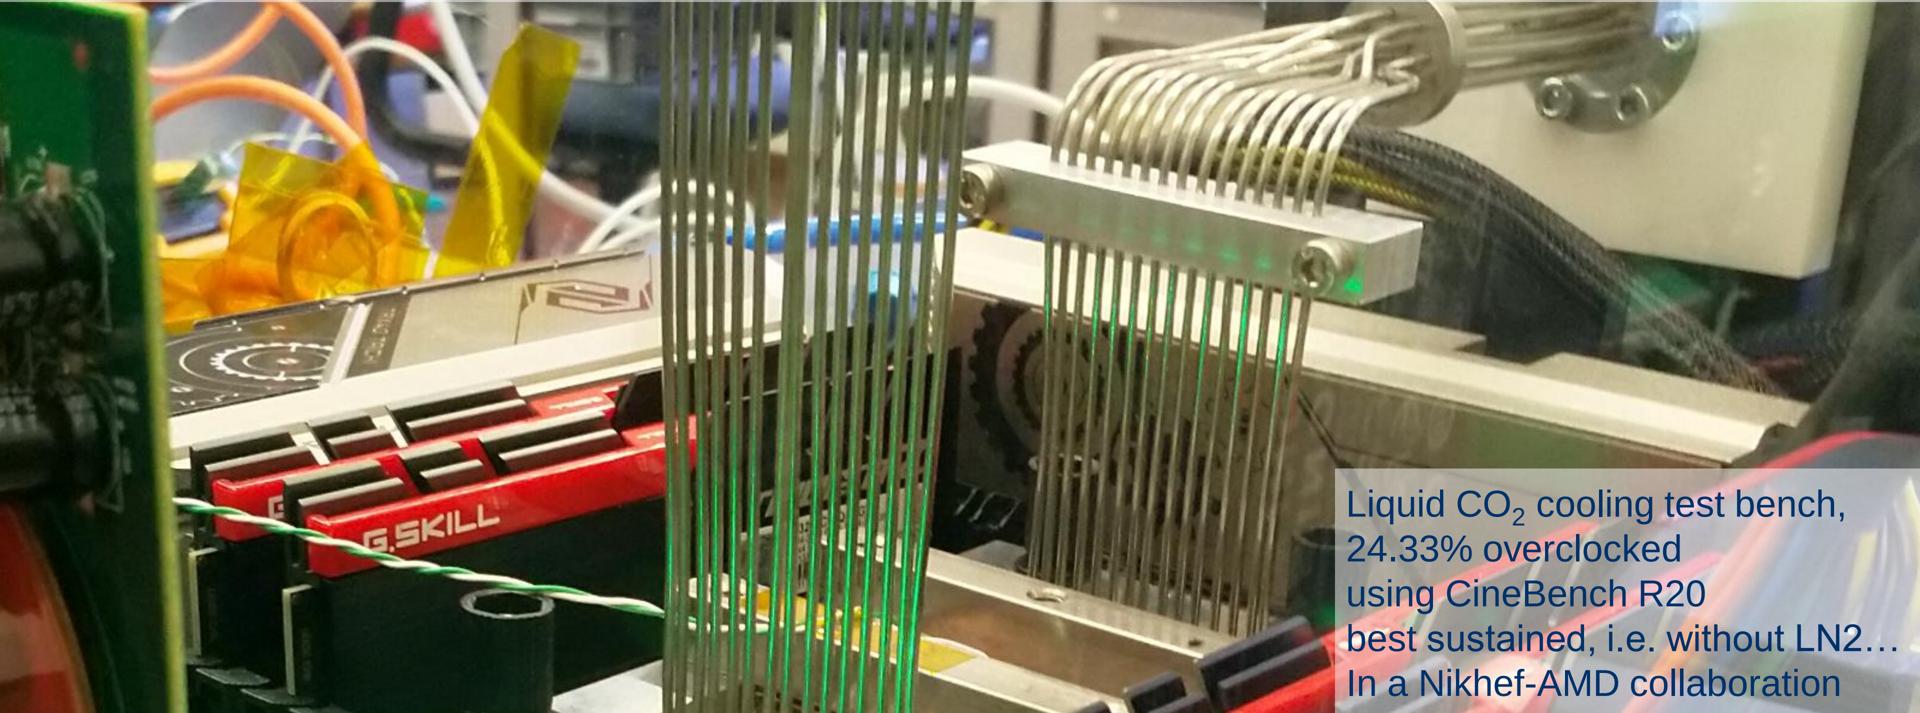

Liquid CO_2 cooling test bench,
24.33% overclocked
using CineBench R20
best sustained, i.e. without LN2…
In a Nikhef-AMD collaboration

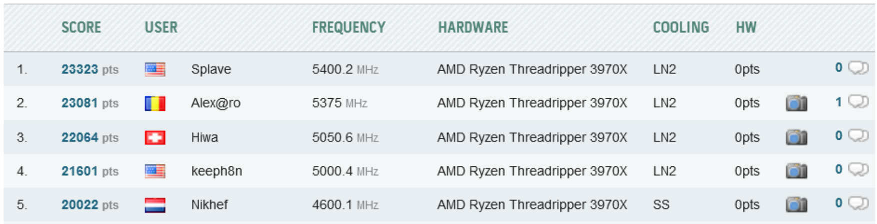

	SCORE		USER	FREQUENCY	HARDWARE	COOLING	HW		
1.	23323 pts	🇺🇸	Splave	5400.2 MHz	AMD Ryzen Threadripper 3970X	LN2	0pts		0
2.	23081 pts	🇷🇴	Alex@ro	5375 MHz	AMD Ryzen Threadripper 3970X	LN2	0pts	📷	1
3.	22064 pts	🇨🇭	Hiwa	5050.6 MHz	AMD Ryzen Threadripper 3970X	LN2	0pts	📷	0
4.	21601 pts	🇺🇸	keeph8n	5000.4 MHz	AMD Ryzen Threadripper 3970X	LN2	0pts	📷	0
5.	20022 pts	🇳🇱	Nikhef	4600.1 MHz	AMD Ryzen Threadripper 3970X	SS	0pts	📷	0

T Suerink, K de Roo: https://hwbot.org/submission/4539341_nikhef_cinebench___r20_with_benchmate_ryzen_threadripper_3970x_20022_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 the 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 Templon, Roel Aaij, Erik Kooistra, Jouke Roorda, Sven Gabriel, Lennie de Roo,, Arjen van Rijn, Krista de Roo (Nikhef), Nicolas Liampotis, Kyriakos Glinis (GRNET), Jens Jensen, Dave Kelsey, David Crooks (STFC RAL), Hannah Short, Liviu Valsan (CERN), Uros Stevanovic and Marcus Hardt (KIT), Maarten Kremers (SURF), Licia Florio (GEANT and 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, and FIM4R communities!

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

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 Grant Agreement No. 856726 (GN4-3), 101017536 (EOSC Future), 730941 (AARC2), and 101100680 (GN5-1).

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

Federation and security

Collaborative security

Sirtfi

Testing resilience and Sirtfi v2

eduGAIN Security and CSIRT

Now *what* have we built?!

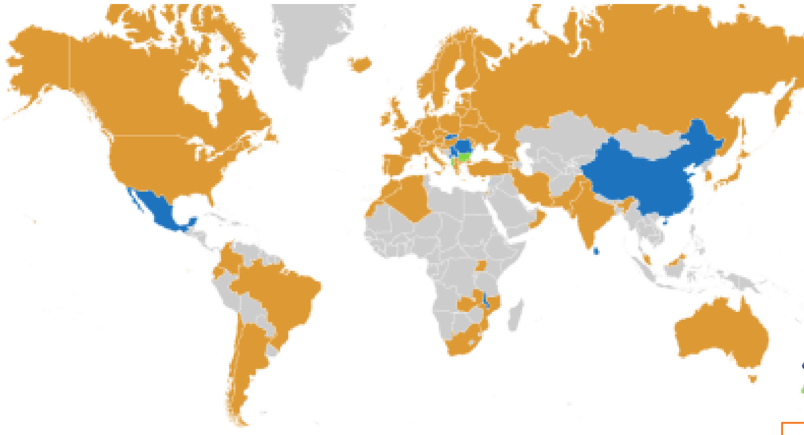

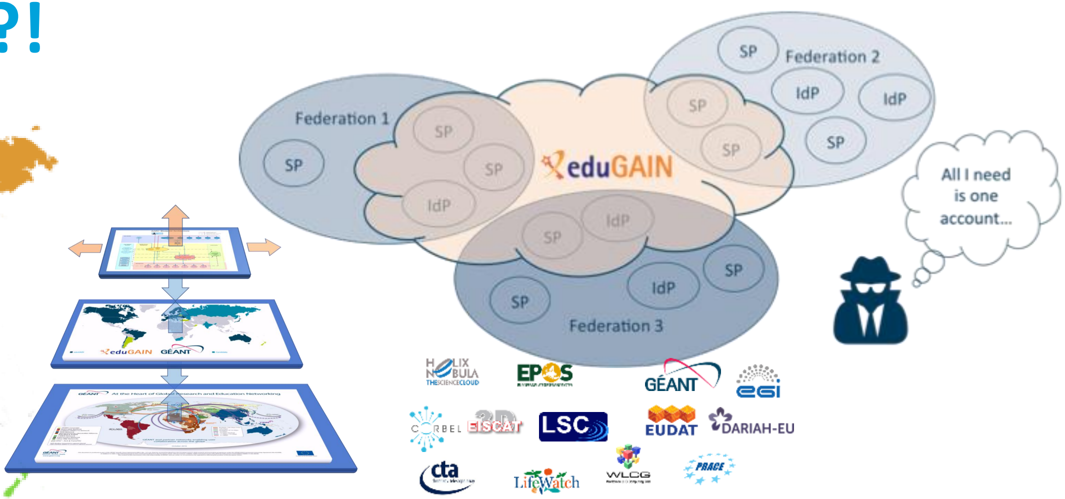

full of valuable resources
(data, network, services)

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

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

Sirtfi – Security Incident Response Trust framework for Federated Identity

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

https://refeds.org/SIRTFI

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

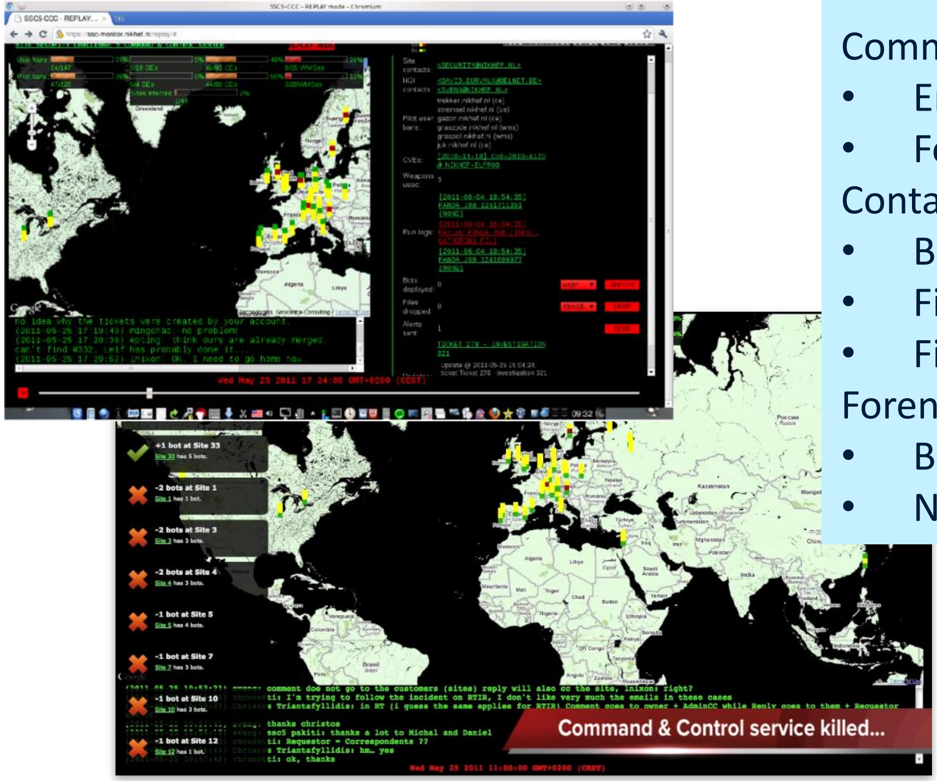

```
no idea why the tickets were created by your account.
(2011-05-25 17:19:43) mingchao: no problem!
(2011-05-25 17:20:39) epting: think ours are already merged.
can't find #332, Leif has probably done it...
(2011-05-25 17:20:52) inixon: OK. I need to go home now.
Wed May 25 2011 17:24:00 GMT+0200 (CEST)
```

```
(2011-05-25 10:53:21) avong: comment doe not go to the customers (sites) reply will also cc the site, inixon: right?
ti: I'm trying to follow the incident on RTIR, I don't like very much the emails in these cases
s Triantafyllidis: in RT (i guess the same applies for RTIR) Comment goes to owner + AdminCC while Reply goes to them + Requestor
avong: thanks christos
avong: ssc5 pakiti: thanks a lot to Michal and Daniel
ti: Requestor = Correspondents ??
s Triantafyllidis: hm… yes
ti: ok, thanks
Wed May 25 2011 11:00:00 GMT+0200 (CEST)
```

Communication:
- Endpoints valid?
- Form/Content OK ?

Containment
- Ban ”malicious” users
- Find/Stop malicious processes
- Find submission IP

Forensics
- Basic Forensics on binary
- Network traffic

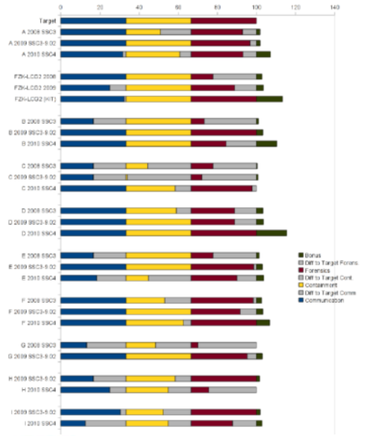

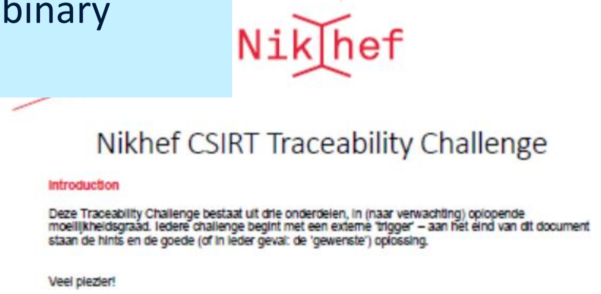

Nikhef

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 goede (of in ieder geval: de 'gewenste') oplossing.

Veel plezier!

A federated community security challenge

SIRTFI
Security Incident Response Trust Framework for Federated Identity

Can we coordinate our collective R&E response?

‘challenges’ based on the *Sirtfi* contact model

Security Incident Response Trust Framework for Federated Identity

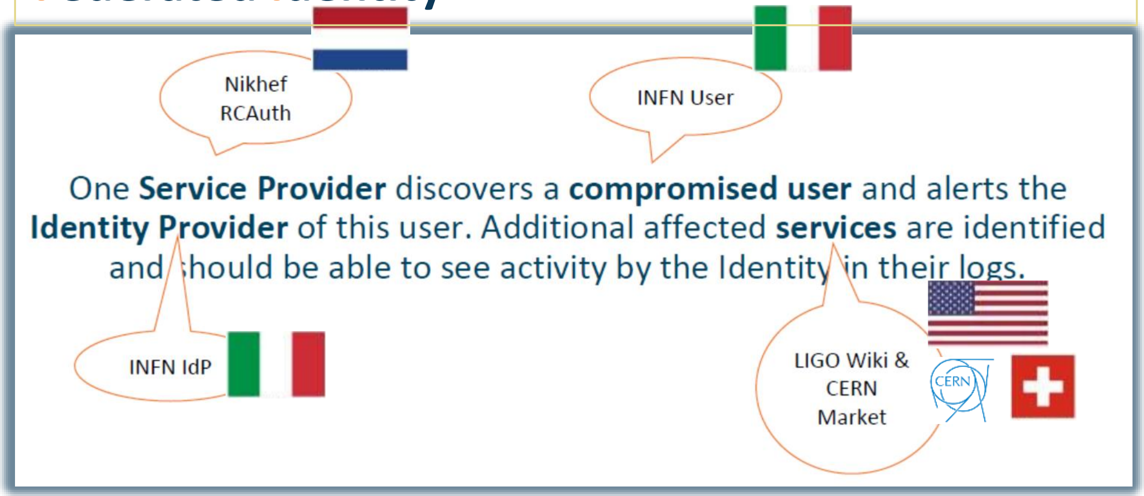

One **Service Provider** discovers a **compromised user** and alerts the **Identity Provider** of this user. Additional affected **services** are identified and should be able to see activity by the Identity in their logs.

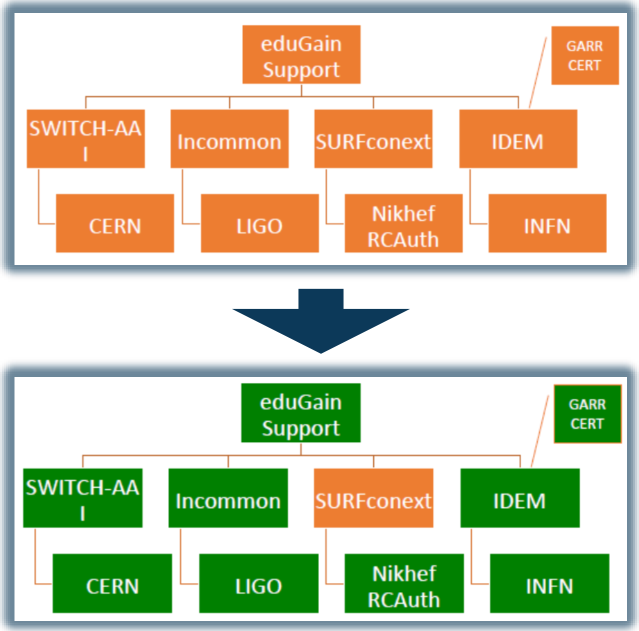

parties involved in response challenge

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

Sharing threat intel – working with our community

MITRE | ATT&CK®

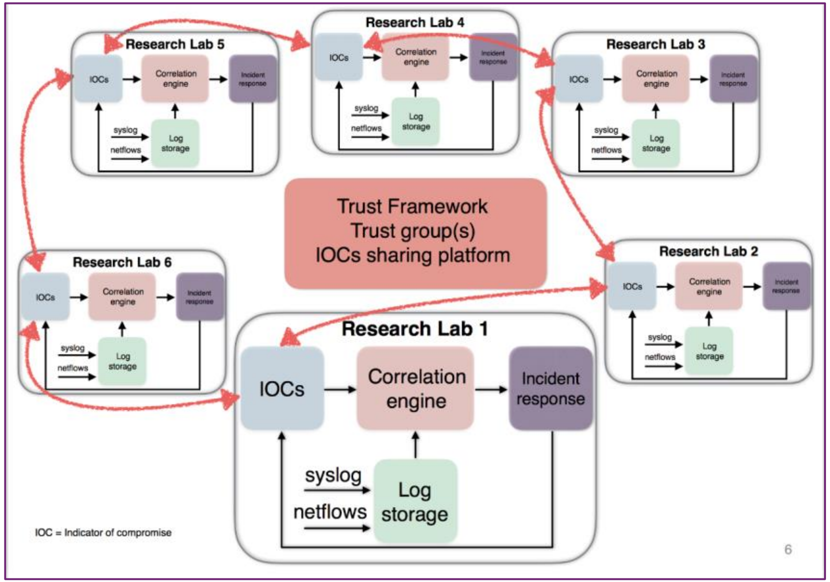

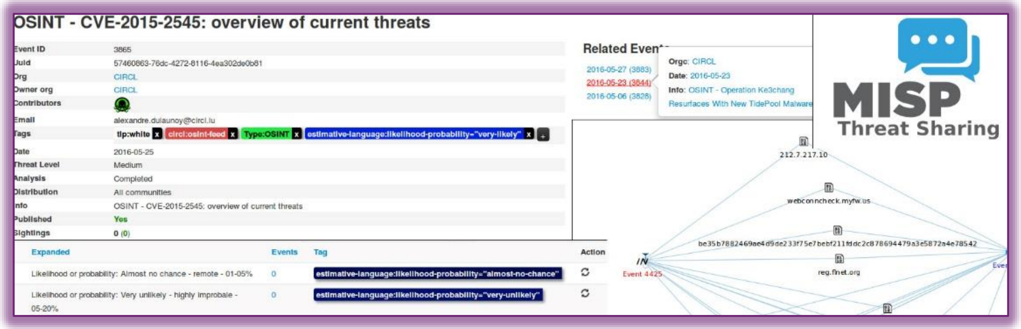

OSINT - CVE-2015-2545: overview of current threats

Event ID	3865
Uuid	57460863-76dc-4272-8116-4ea302de0b81
Org	CIRCL
Owner org	CIRCL
Contributors	
Email	alexandre.dulaunoy@circl.lu
Tags	tlp:white x circl:osint-feed x Type:OSINT x estimative-language:likelihood-probability="very-likely" x +
Date	2016-05-25
Threat Level	Medium
Analysis	Completed
Distribution	All communities
Info	OSINT - CVE-2015-2545: overview of current threats
Published	Yes
Sightings	0 (0)

Related Events

2016-05-27 (3883)
2016-05-23 (3844)
2016-05-06 (3828)

Orgc: CIRCL
Date: 2016-05-23
Info: OSINT - Operation Ke3chang Resurfaces With New TidePool Malware

Expanded	Events	Tag	Action
Likelihood or probability: Almost no chance - remote - 01-05%	0	estimative-language:likelihood-probability="almost-no-chance"	
Likelihood or probability: Very unlikely - highly improbale - 05-20%	0	estimative-language:likelihood-probability="very-unlikely"	

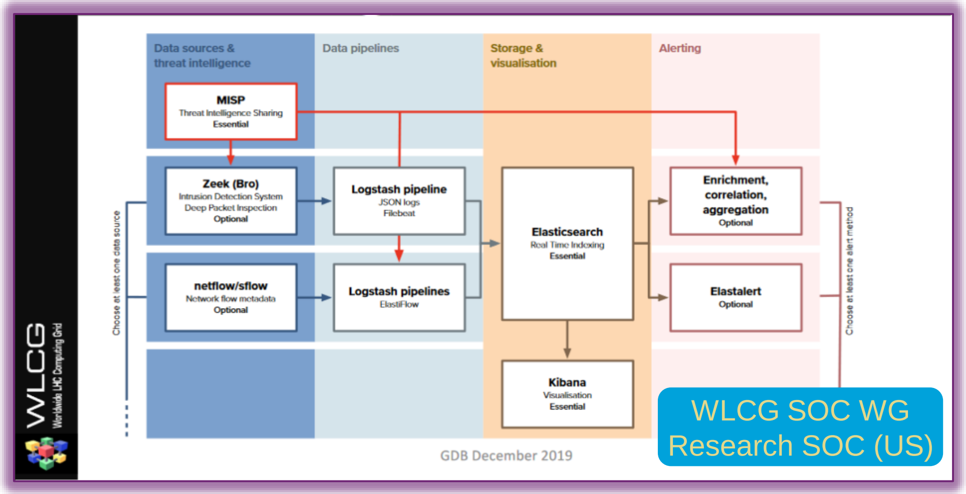

WLCG SOC WG
Research SOC (US)

AARC

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

Nikhef SOC – NDPF traffic analysis

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

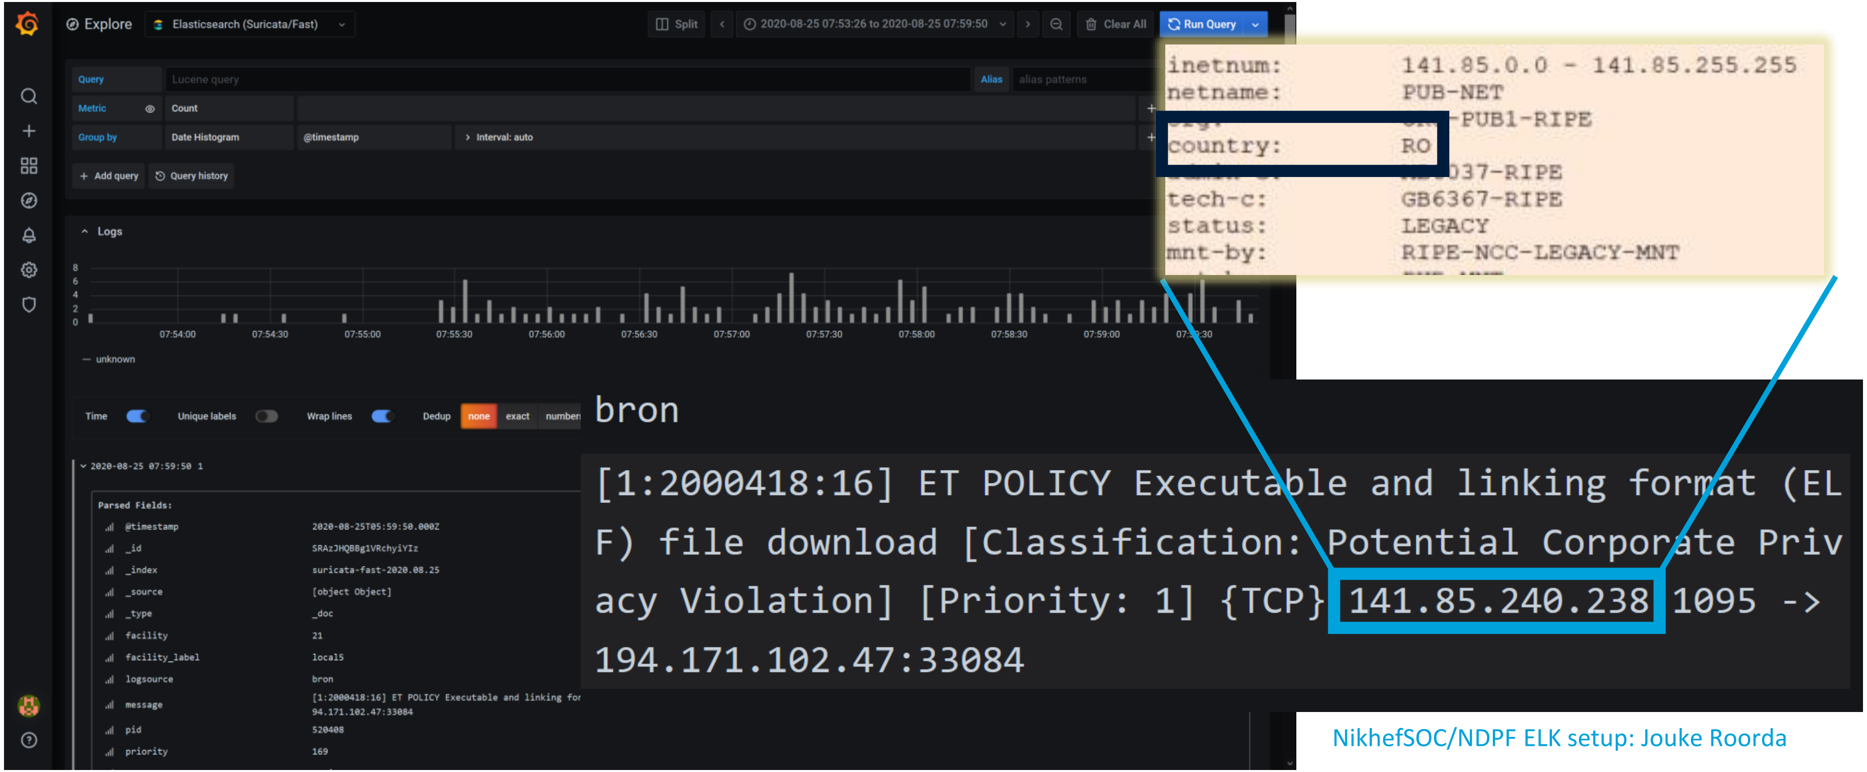

```
inetnum:        141.85.0.0 - 141.85.255.255
netname:        PUB-NET
                  -PUB1-RIPE
country:        RO
                   037-RIPE
tech-c:         GB6367-RIPE
status:         LEGACY
mnt-by:         RIPE-NCC-LEGACY-MNT
```

```
2020-08-25 07:59:50 1

Parsed Fields:
  @timestamp        2020-08-25T05:59:50.000Z
  _id               SRAzJHQBBg1VRchyiYIz
  _index            suricata-fast-2020.08.25
  _source           [object Object]
  _type             _doc
  facility          21
  facility_label    local5
  logsource         bron
  message           [1:2000418:16] ET POLICY Executable and linking for
                    94.171.102.47:33084
  pid               520408
  priority          169
```

```
bron

[1:2000418:16] ET POLICY Executable and linking format (EL
F) file download [Classification: Potential Corporate Priv
acy Violation] [Priority: 1] {TCP} 141.85.240.238:1095 ->
194.171.102.47:33084
```

NikhefSOC/NDPF ELK setup: Jouke Roorda

The eduGAIN Security Handbook

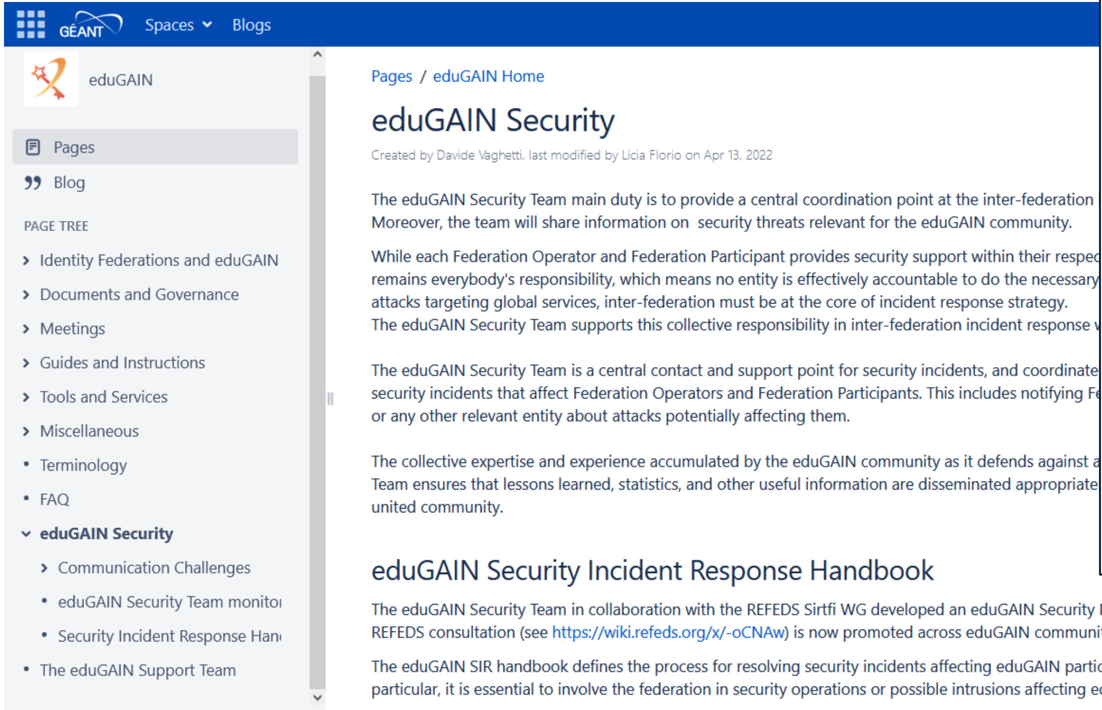

Pages / eduGAIN Home

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 respec
remains everybody's responsibility, which means no entity is effectively accountable to do the necessary
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 v

The eduGAIN Security Team is a central contact and support point for security incidents, and coordinate
security incidents that affect Federation Operators and Federation Participants. This includes notifying Fe
or any other relevant entity about attacks potentially affecting them.

The collective expertise and experience accumulated by the eduGAIN community as it defends against a
Team ensures that lessons learned, statistics, and other useful information are disseminated appropriate
united 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 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.

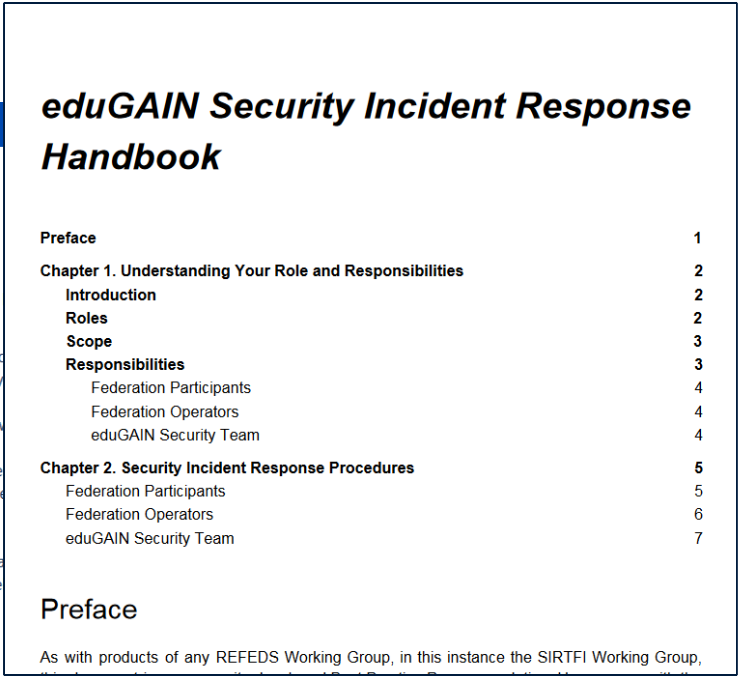

eduGAIN Security Incident Response Handbook

Preface

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

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

Nulla folia post hoc sunt

Thanks for watching!

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