Nikhef

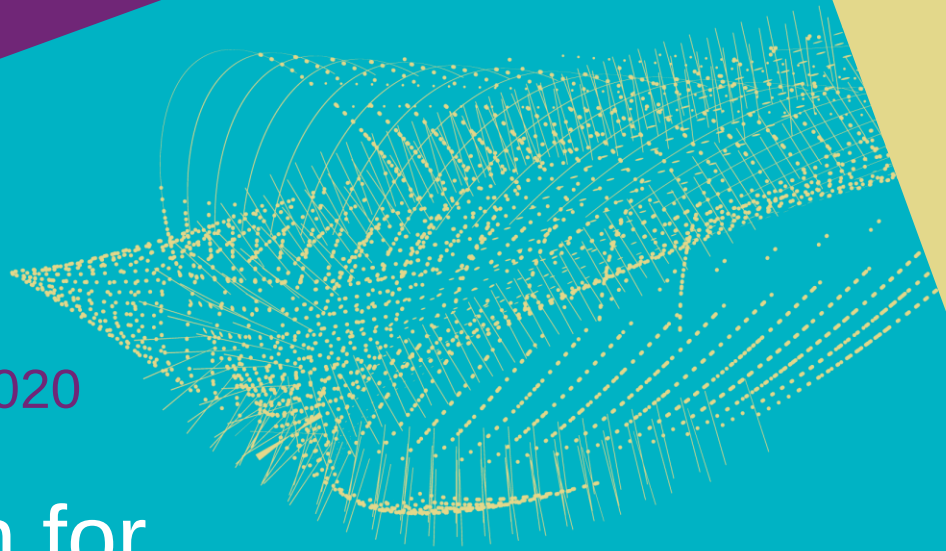

NSF Cybersecurity Summit 2020

Trust Coordination for Research Collaboration in the era of EOSC

David Groep, et al.
for the EOSC Trust & Security Operation collaboration
September 2020

EOSC? The “European Open Science Cloud”

- a ‘commons’ for research data aiming to combine all disciplines across all (European) countries
- an ongoing process, with both means and methods still very much evolving
- ‘a portal’, ‘a marketplace’, ‘a web of FAIR data’
- ‘an infrastructure’ … or its ‘data twin’

whatever it is, it will be structuring data-driven research in Europe in the 2020s

Photo by Pop & Zebra on Unsplash

EOSC vision

Current model of European data infrastructures

Source: EOSC Strategic Implementation Roadmap 2018-2020, May 2018, European Commission

From fragmentation and uneven access to information to a federated model, where access to data would be universal, building on a strong legacy

Future EOSC model: federation of data infrastructures

Credits: Ognjen Prnjat– Project Coordinator- NI4OS-Europe, EOSC Symposium, 26.11.2019, Budapest 14

sources: https://www.eoscsecretariat.eu/eosc-symposium-programme

An ecosystem more than an infrastructure

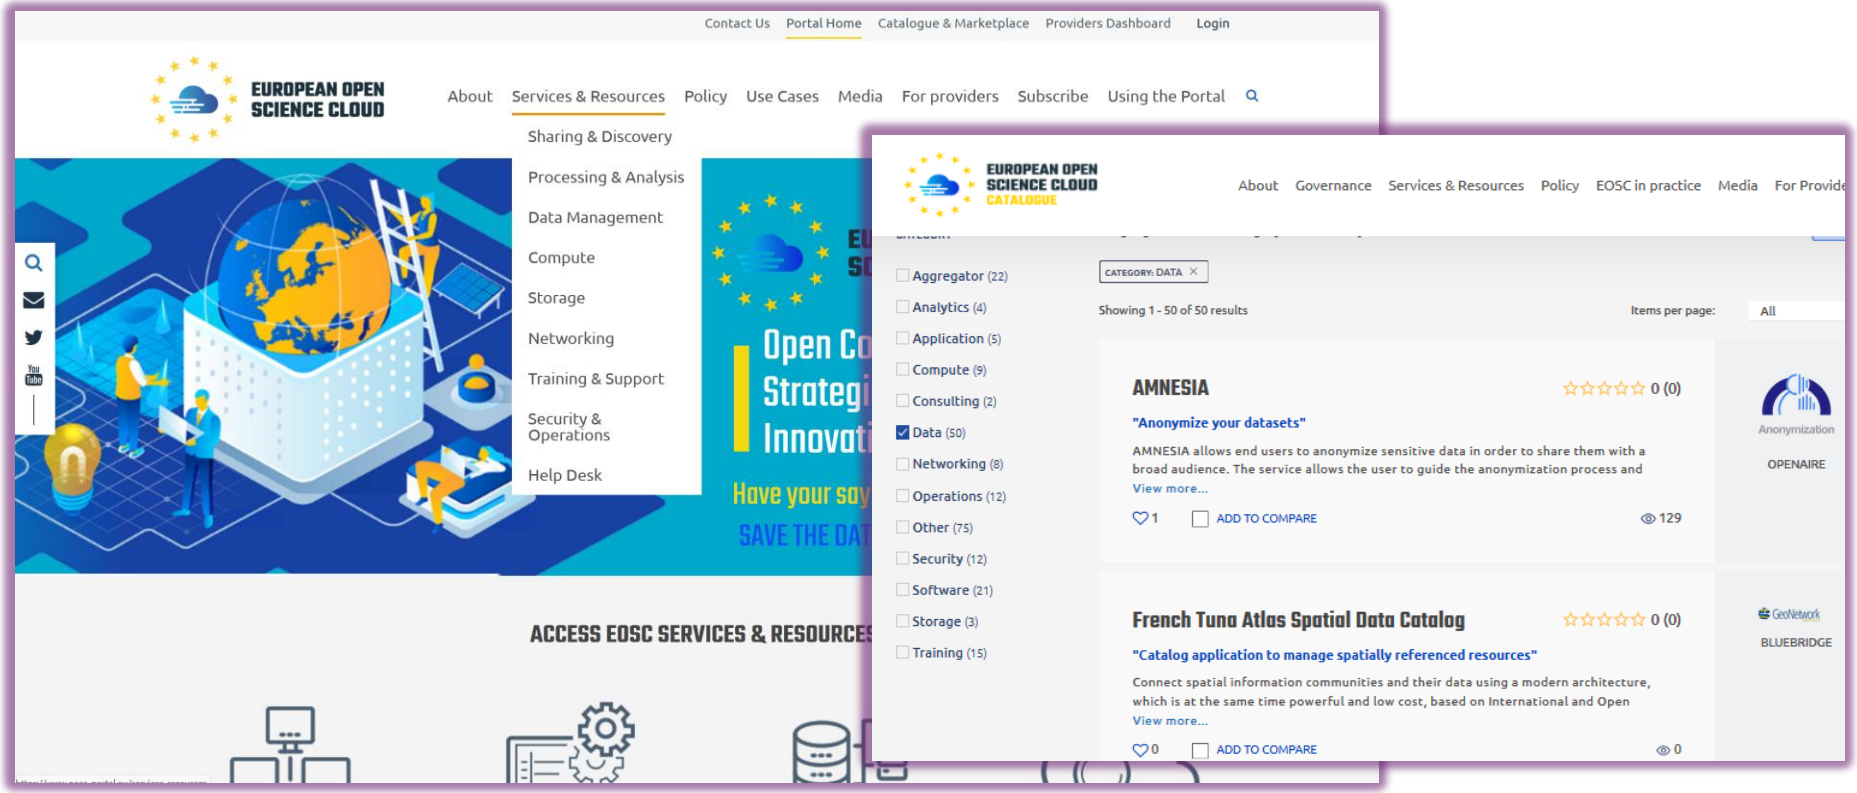

EOSC Portal (https://www.eosc-portal.eu/) – as built by EOSChub

A challenging landscape

Entities of all kinds – diversity in the EOSC range
from *data sets* to *storage* to *computing* to *publications & digital objects*

An open ecosystem – rules of participation will favour low barrier to entry regarding operational maturity, service management quality, &c

A diverse ecosystem – providers will come from e-Infrastructures, from member states, from research infrastructures, and private sector

An *interdependent* ecosystem – aiming for composability and collective service design through an open, core AAI federation

Core services and ‘the exchange’

What constitutes a ‘core service’? A thin layer, with

- at least the service catalogue (portal) itself
- governance, landscaping, and policy
  persistent identifiers, certifications, trademarking

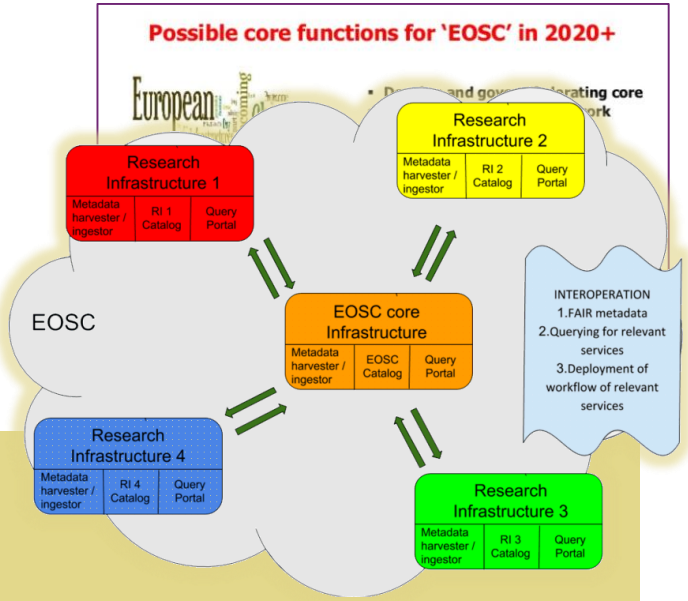

- **AAI federation** - authentication and authorization
  based on the ‘AARC BPA’
- IT service management for the (core) services
- **operational security capabilities, trust policy, and security risk structuring**

Sustainability and Architecture WGs set criteria for inclusion of additional services
Architecture WG and its taskforces set interoperability standards

and for the ‘BPA’ AARC Blueprint Architecture? See https://aarc-community.org/architecture/

image sources: EOSC Secretariat, Karel Luyben, EOSC-Future drafts, v14

Minimum Viable … EOSC

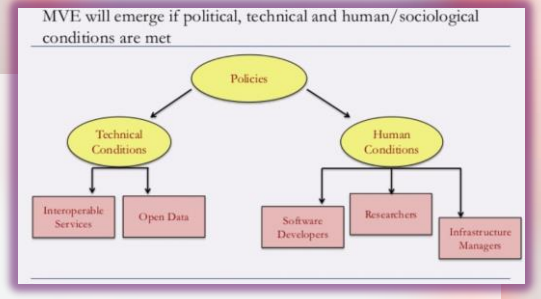

Great Expectations … but what about requirements?

‘MVE – MINIMUM VIABLE EOSC’

includes some *Rules of Participation* to aid security & trust

Core

- ‘distributed and participatory’
- ‘collaborative consensus’
- ‘interoperability standards, […] and implementation via best practices’

Exchange & Portal

- ‘research-enabling services’
- ‘national, regional, institutional, domain based, … and commercial’
- ‘catalogue …[for] research life cycle’

- it will be a mix, and in any case service providers will need to contribute
- *Sirtfi shows that is not completely unrealistic*

Sirtfi – security incident response trust framework for federated identity – see refeds.org/sirtfi

graphic: Prompting an EOSC in Practice – Isabel Campos, CSIC & EOSC HLEG

Photo: Patrick Perkins (Unsplash)

Back to Basics: the few tenets for the EOSC ecosystem security

From *promoting and monitoring capabilities* **to** *managing core risk*

A service provider should

- **do no harm** to interests & assets of users
- **not expose *other*** service providers in the EOSC ecosystem to enlarged risk as a result of *their* participation in EOSC
- **be transparent** about its infosec maturity and risk to its customers and suppliers

this will mean *some minimum requirements* in the Rules of Participation

Photo Hippokrates tomb: Melania Stubos, CC-BY-SA-3.0 http://himetop.wikidot.com/hippocrates-funeral-monument

Making the EOSC a trusted place

Risk-centric self-assessment framework
- based on federated InfoSec guidance including WISE SCI

Baselining security policies & common assurance
- AARC, REFEDS, IGTF, PDK & practical implementation measures

An incident coordination hub and a trust posture
- spanning providers and core, based on experience & exercises

Actionable operational response to incidents
- EOSC core expertise to support resolution of cross-provider issues

Fostering trust through a known skills programme
- so that your peers may have confidence in service provider abilities

WISE SCI: wise-community.org/sci
AARC&c: aarc-community.org, refeds.org, igtf.net
PDK: aarc-community.org/policies/policy-development-kit

Assessing risk … in a peer-review framework

InfoSec **risk assessment framework**
for EOSC services based on
a federated evolution of WISE SCI and
a multi-tier maturity model,
also addressing data security and protection

- risks ‘play out’ differently
  in different infrastructures
- more than storage or compute, but also
  risks for (open) data and for reputation

Many risks are generic, some need context and expertise to assess. Or are under regulated regime

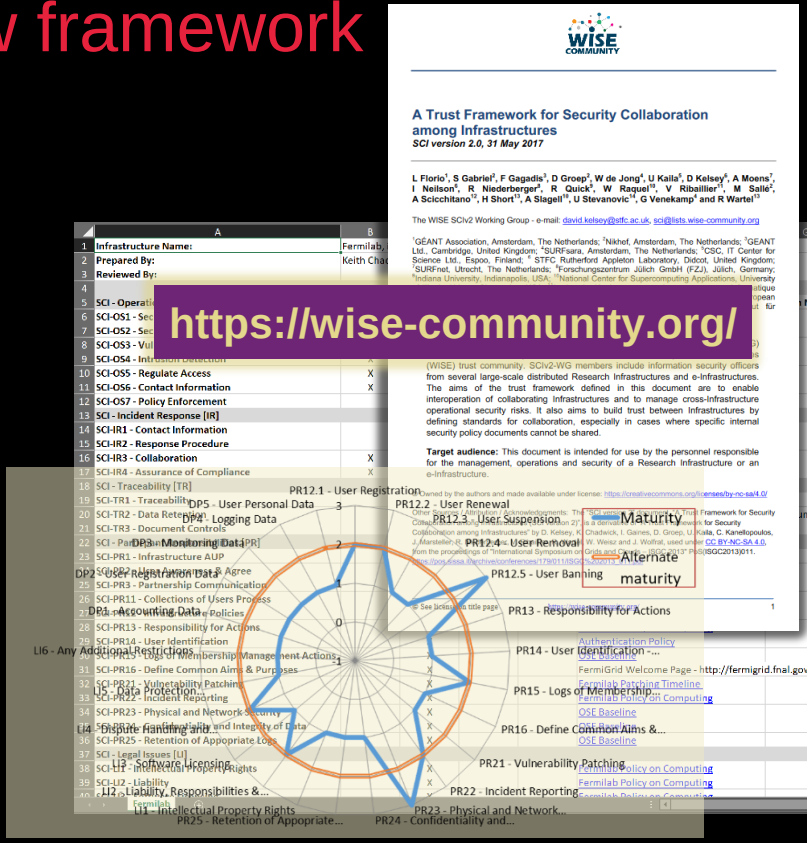

this spider diagram is fictional – idea by Urpo Kaila, CSC

Shared understanding of a baseline?

Closely coordinated infrastructures – e.g. WLCG, EGI –
started with a single common policy set and assurance level

- service providers and users ‘understand’ its meaning and compliance
  - and the understanding is shared

Move towards differentiated models
adds flexibility, but also complexity!

- different means to achieve same goal
- varying means to achieve
  different goals with diverse risk

Diversification is complex

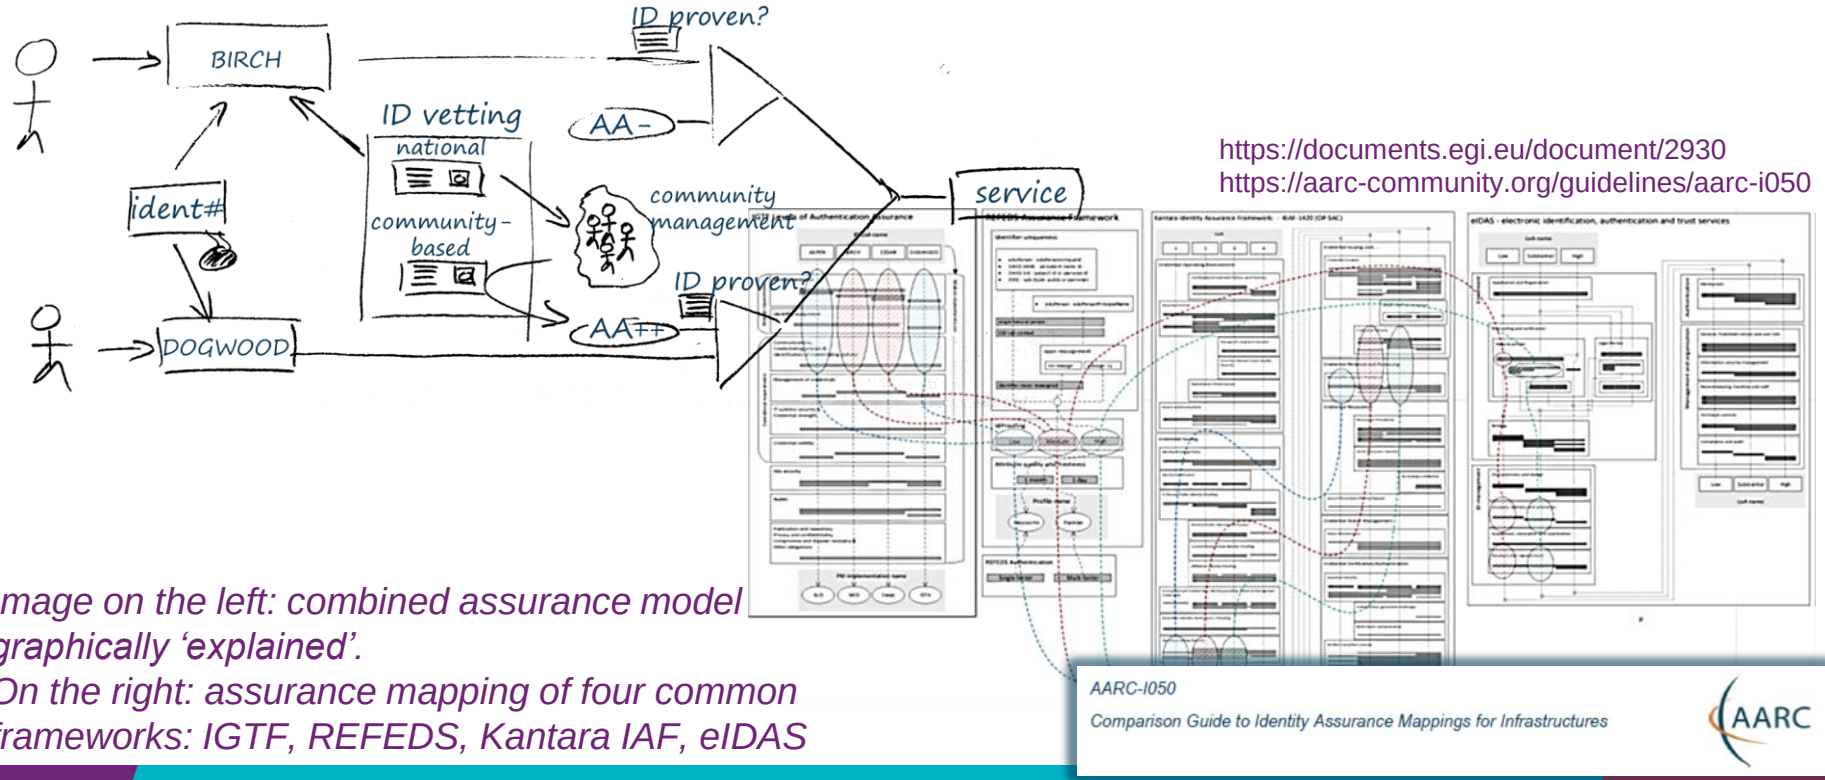

https://documents.egi.eu/document/2930
https://aarc-community.org/guidelines/aarc-i050

Image on the left: combined assurance model graphically 'explained'.
On the right: assurance mapping of four common frameworks: IGTF, REFEDS, Kantara IAF, eIDAS

Managing an EOSC policy baseline and assurance

A diverse set of requirements

- EOSC mechanisms & working groups
- Community and e-Infrastructure requirements
- Operational security need for response, containnment, and resolution

and remain practical and manageable

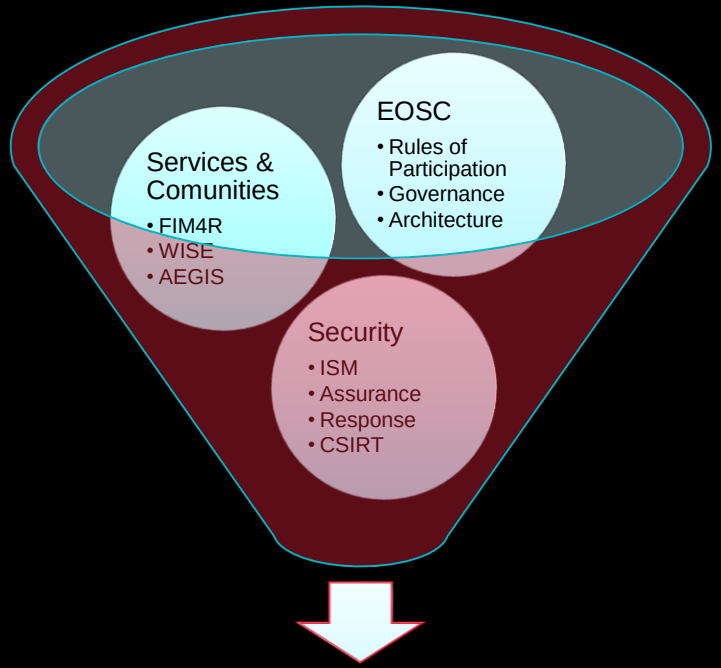

security baseline, trust and assurance profiles

Start with baselining

baselining has been very effective with Sirtfi, for R&S, and for InCommon …

Trust marks or seals
for specific service levels, access classes, types of data, regulatory domains, &c

SCI-based policy mapping
leverage common templates like the WISE Acceptable Use Policy, or membership management …

Good Practice
policy implementation guidance
small number of assurance profiles (REFEDS, IGTF, eIDAS), AARC secure operations standards, AEGIS recommendations, CSIRT capability

Technical guidance
e.g. expression of identity assurance

Rules of Participation
minimal set of capabilities – initially maybe just contact information, responsiveness, confidentiality

THE POLICY DEVELOPMENT KIT

Document	Who should complete the template?	Audience	Description
Top Level Infrastructure Policy	Infrastructure Management	All Infrastructure Participants (abides by)	This policy template defines the roles of actors in the Research Infrastructure and binds the policy set together
Incident Response Procedure	Infrastructure Management & Security Contact	Infrastructure Security Contact, Services (abides by)	This is a placeholder for the Infrastructure to determine rules for the
Membership Management Policy	Infrastructure Management	Research Community (abides by)	This
Acceptable Authentication Assurance	Infrastructure Management	Research Community, Services (abide by)	This
Policy on the Processing of Personal Data	Infrastructure Management & Data Protection Contact	Research Community, Services (abide by)	This
Service Operations Security Policy	Infrastructure Management	Services (abide by)	This for running a service within the Infrastructure.
Risk Assessment	Infrastructure Management, Services & Security Contact	Infrastructure Management (completes)	This table can be used as a starting point for identifying whether a full Data Protection Impact Assessment is required.
Acceptable Use Policy	Infrastructure Management (for baseline) & Research Communities (for community specific restrictions)	Users (abide by)	This

Showing 1 to 9 of 9 entries

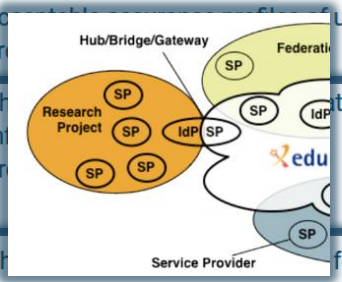


Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (Sctfi)

Licia Florio (GÉANT), David Groep (Nikhef), Christos Kanellopoulos (GÉANT), David Kelsey (STFC), Mikael Linden (CSC), Ian Neilson (STFC), Stefan Paetow (Jisc), Wolfgang Pempe (DFN), Vincent Ribaillier (IDRIS-CNRS), Mischa Salle (Nikhef), Hannah Short (CERN), Uros Stevanovic (KIT) and Gerben Venekamp (SURFsara)

Sctfi

https://aarc-project.eu/policies/policy-development-kit/
https://aarc-community.org/policies/sctfi/

graphic IdP-SP bridge: Lukas Hammerle and Ann Harding, SWITCH

Establishing the trust basis for response

Collaboration frameworks, processes, exercises – the basis of trust
since not everything can be done on personal trust and 'blind faith'

sources: GEANT CLAW, https://connect.geant.org/2020/02/19/claw-2020-save-the-date
Sirtfi: Hannah Short et al. https://wiki.geant.org/pages/viewpage.action?pageId=123766092

Do I know that you know what to know about what?

Training - and ability to exercise - intelligence sharing framework and best practices, but *also* collective technical and forensic expertise!

- build up expertise to desired maturity – esp. across EOSC portal providers and research communities
- desirable, but not yet likely, to have training a requirement for participation
  that is hard for an EOSC
  that does not wish barriers to entry ☹

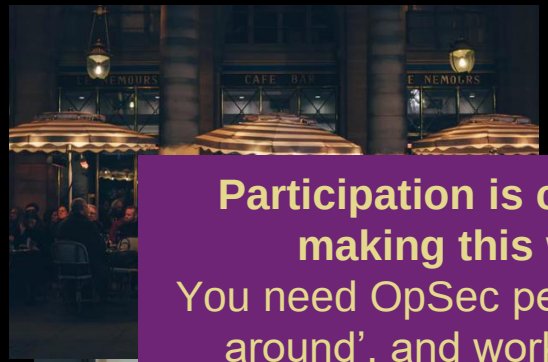

Participation is critical to making this work
You need OpSec people to 'get around', and work globally

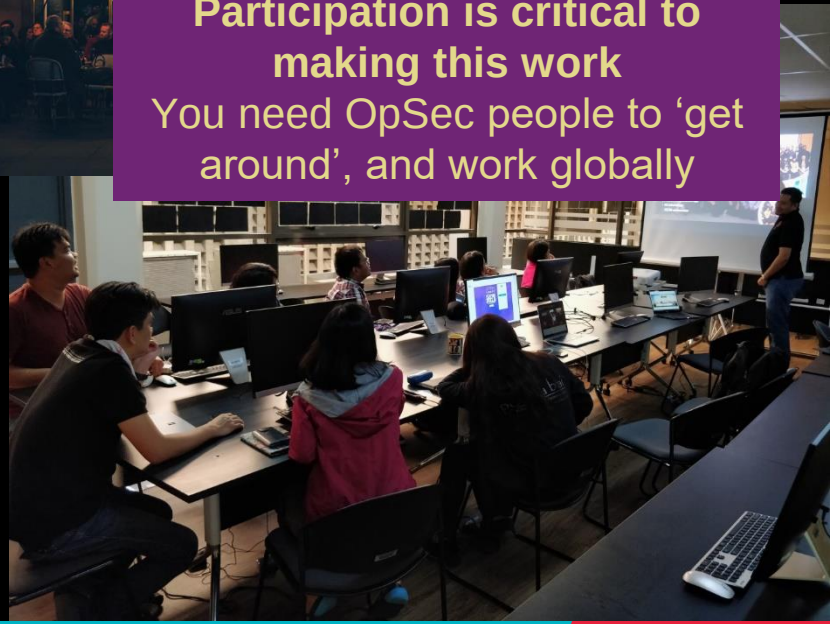

image credits: TRANSITS-I

Actionable Response – coordination involving the Core

We *know* we cannot address all needs, but we can make progress

‘in the end, the same people do the same work, together, and regardless of the project of funding label’

- EOSC core will itself be a significant hub
- tightly-knit team of experts
  looking after the security of the core
- who can work collaboratively
  with peer infrastructures and groups

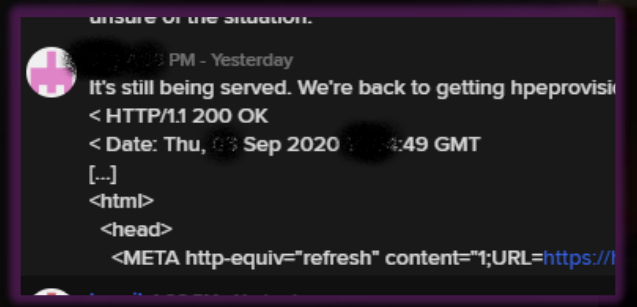

this team is essential to glue together the information during incidents – leveraging the trust built up before through engagement

But isn't 'AAI' going to solve all that 'as a service'?

… we really heard that one …

and although the AAI is a core service of EOSC …

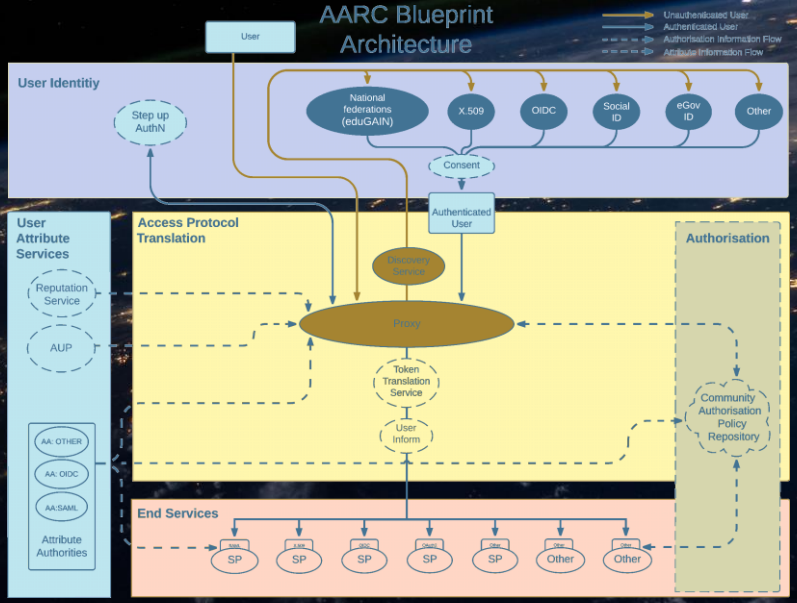

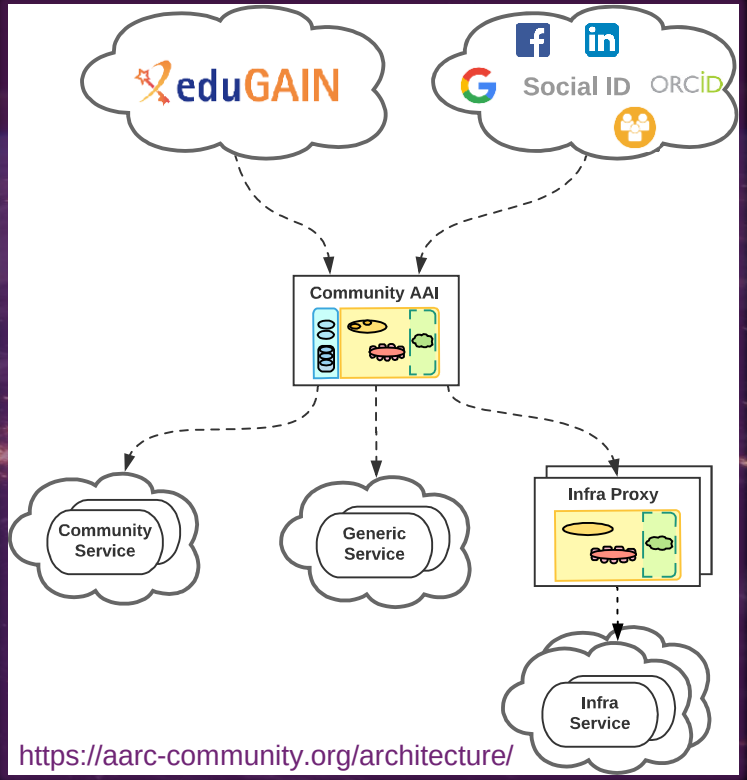

Linking the providers and users together - AAI

AARC BPA's 'community-first' model does not cover all EOSC cases, e.g. *infrastructures acting as providers **and** suppliers **and** as attribute authority*

You need to turn the EOSC entities into a federation in itself, with carefully forged links to eduGAIN to prevent 'user loop' inconsistencies

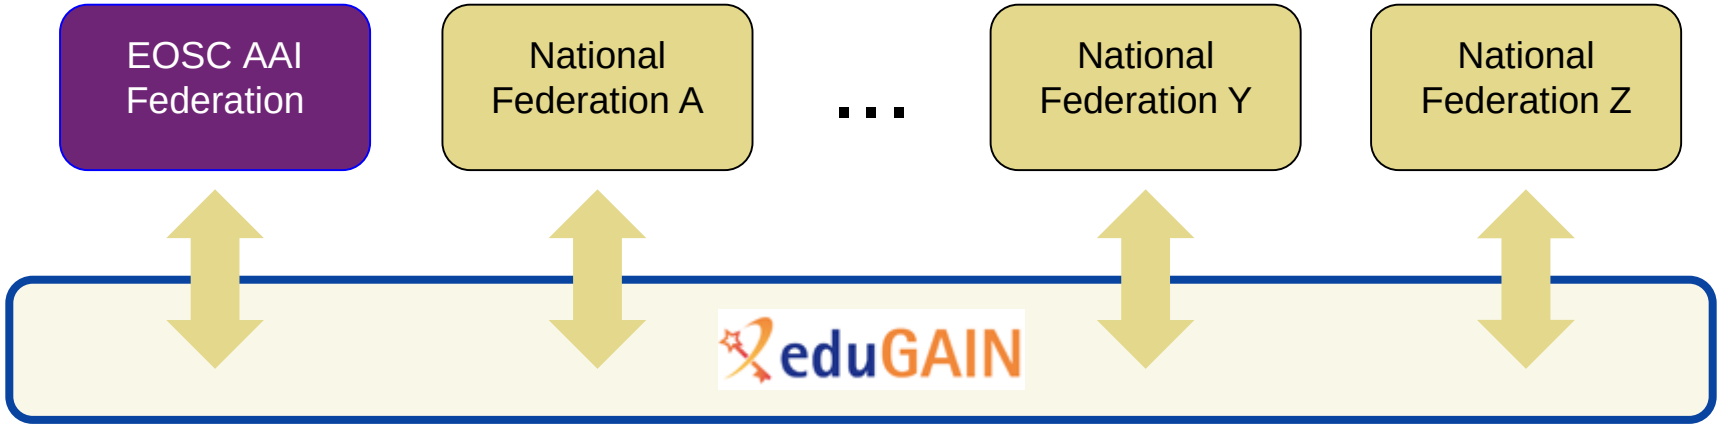

But now … turtles all the way down

… now that new 'EOSC' federation needs policies and a base line

- inspired by eduGAIN constitution and other sources
- leveraging existing trust frameworks
- and not repeating earlier mistakes so implement a baseline at the start

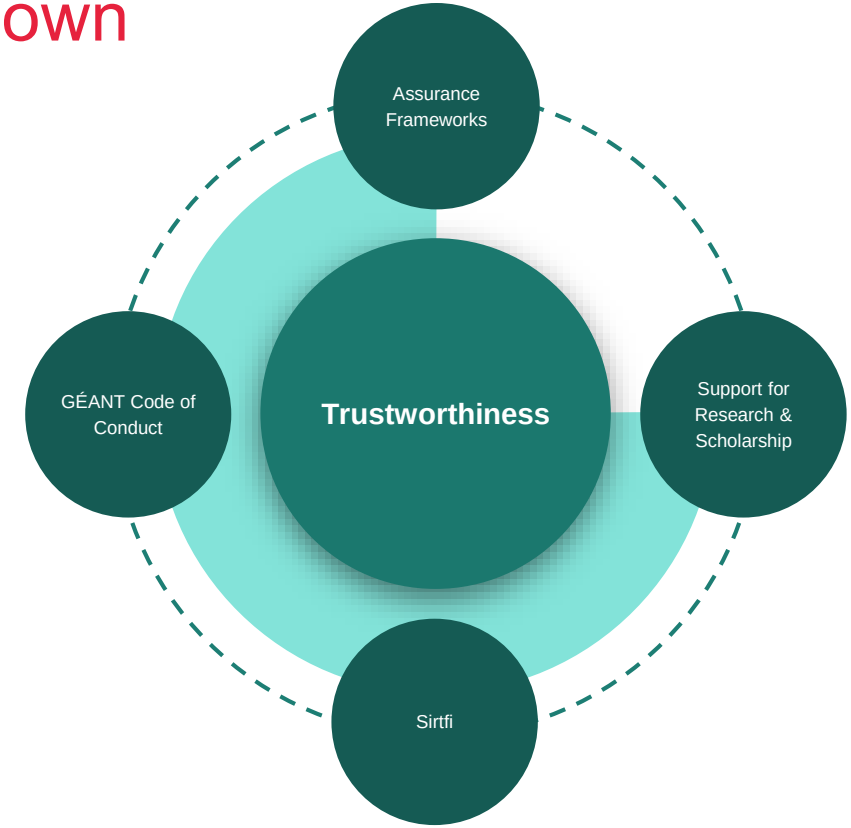

slide graphic: Christos Kanellopoulos, with NicolasL, DavideV, and DavidG

Must EOSC-level mechanisms solve everyone's issue?

do we face
an unbounded challenge?

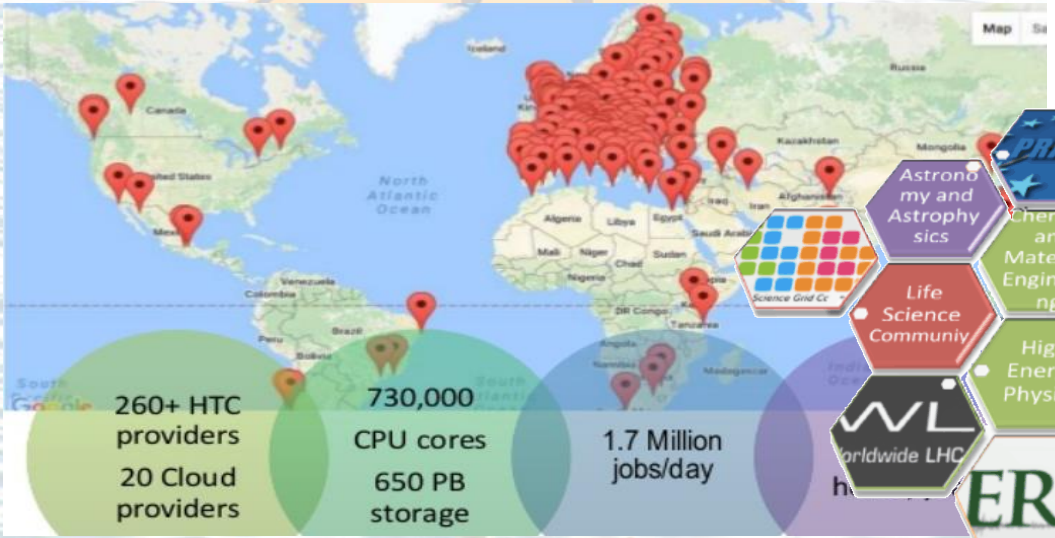

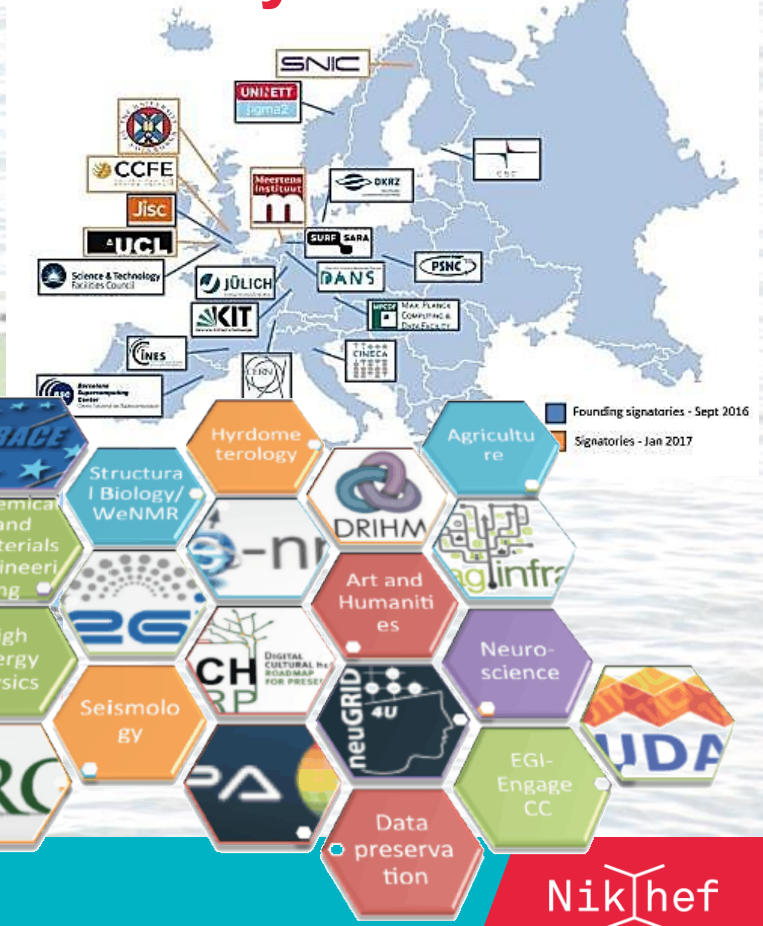

What we expect in the infrastructures and services

Service providers should be at, or grow towards, a mature security stance

and an **infrastructure** provides coordination amongst ‘similar’ things

- providers in an infrastructure can **benefit from their commonalities** *in response and security verification, and vulnerability management*
- a mature EOSC security capability can be structured with infrastructure in **a scalable way** across many service providers

While ‘services’ generally are very broad, including data, publications, &c

Infrastructures: profiting from having a shared services set

common vulnerabilities,
or common risk environment

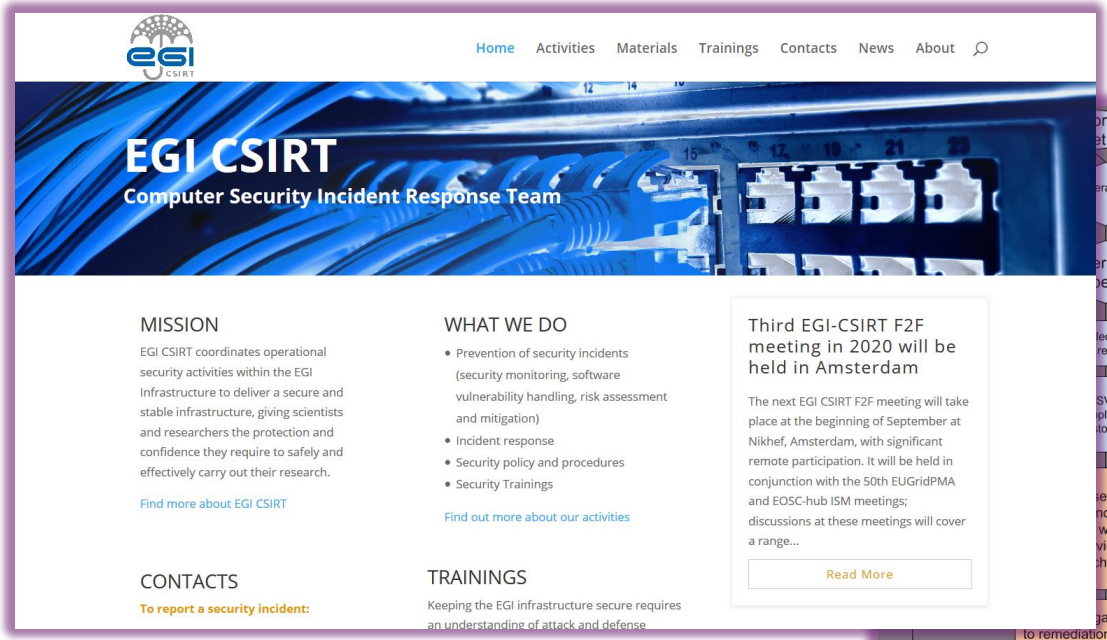

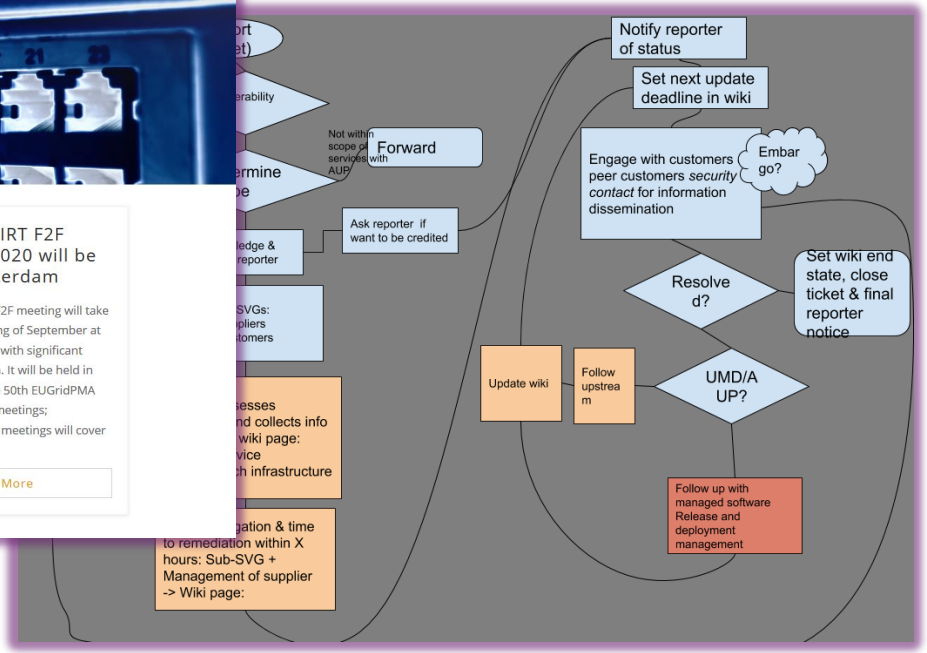

commonality in user base
and access patterns – and testing

image sources: csirt.egi.eu and EGI SVG

Thus even generic capabilities will be widely distributed

EOSC 'Portal' and ecosystem

security for a loosely coupled ecosystem

- risk management for collective services
- security baselining and trust marking
- coherence of response, community readiness/collaboration, and information sharing
- resolution, forensics, resolution and remediation for core and stakeholders
- training and capability enhancement

Core in EOSC-Future

(e-)Infrastructures, services, content

- service security & integrity, responsiveness, compliance monitoring
- vulnerability management and pro-active security management
- incident response and resolution within the infrastructure or service

EGI | EUDAT | GEANT | OpenAIRE | *ServiceX* | ..

See also *Trust Coordination for Research Collaboration in the EOSC era,* February 2020, https://g.nikhef.nl/eosc-sec-wp; https://doi.org/10.5281/zenodo.3674676

Common questions – open answers

Will the core team drown?

the incident response and forensics experts busied consistently with service-specific response, and the 'portal' not able to help through of its participating providers?

Or can we do better?

- a baseline policy bringing enough trust to keep an EOSC-like ecosystem secure?
- will service providers act collectively in the common interest?
- will diverse policy and assurance establish a common reputation for services?
- will provider self-assessment and mitigation of key risks, be seen as 'good value'?

And … do the users care?

- and: *care enough* to make trust and security worth the cost for service providers?

Photo by Yash Prajapati on Unsplash

long read: *Trust Coordination for Research Collaboration in the EOSC era* http://doi.org/10.5281/zenodo.3674677

so: do we stand a chance?

based on the white paper by
David Groep, Jens Jensen, Dave Kelsey,
Daniel Kouřil, Maarten Kremers, and Hannah Short
and on discussions in the EOSC Future
Security Operations & Policy collaboration
with, in addition, Urpo Kaila, Alf Moens, and Vincent Brillault

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

Parts of this work are based on results that received support from projects co-funded from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3), grant agreement No. 777536 (EOSC-hub), and under Grant Agreement No. 730941 (AARC2).

Nikhef