


50th EUGridPMA, AARC & EnCo meeting

Trust Coordination for Research Collaboration in the era of EOSC

*David Groep (Nikhef) et al.
September 2020
co-supported by SURF &
the GEANT4-3 project*

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’

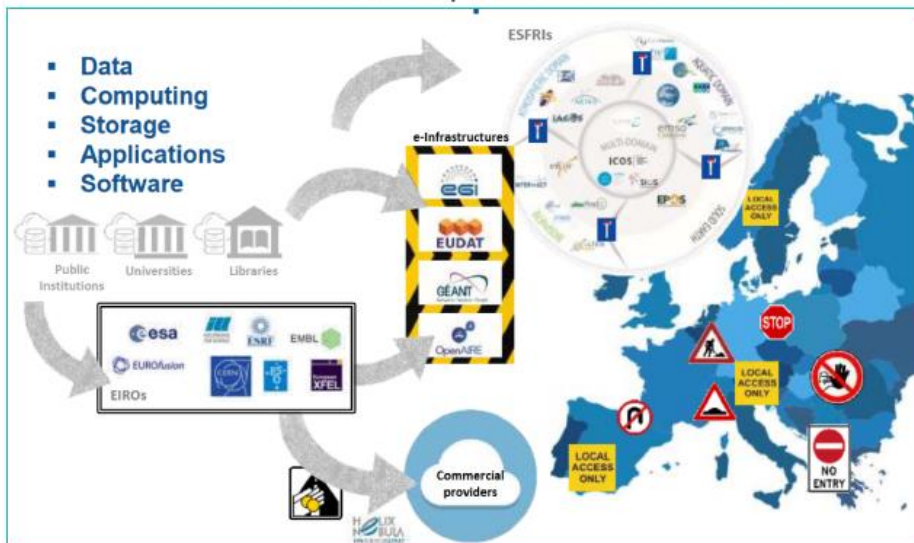


**whatever it is, it will be structuring
data-driven research in Europe in the 2020s**

Photo by Pop & Zebra on Unsplash

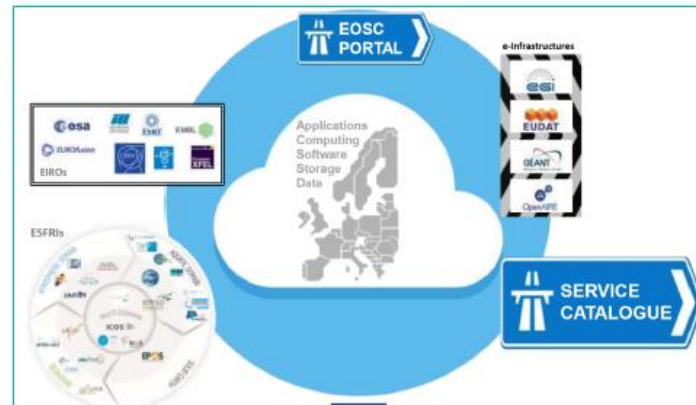
EOSC vision

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

EOSC – an ‘ecosystem of projects’ towards sustainability

Like almost any EU endeavour, EOSC is a process



1. **EOSC Pilot**, *the design study project*
2. **EOSC hub**, *towards a core based on infrastructures*
3. **EOSC {synergy, pillar, Nordic, ...}**, *expanding scope in domains and regions*
4. **EOSC Secretariat**, *modelling governance and moving towards the ‘EOSC ivzw’*
5. **EOSC Future ...** *bringing together the infrastructures and communities in a common portal architecture, and supporting the technical roadmapping efforts*

Image: Maria Teneva on unsplash

A driving force for both infrastructures and domains

Clustering of Infrastructures in Europe is amalgamating research either you're in a large infrastructure, or you're in the 'long tail' ...

"We are creating a European Open Science Cloud now. It is a trusted space for researchers to store their data and to access data from researchers from all other disciplines. We will create a pool of interlinked information, a 'web of research data'. Every researcher will be able to better use not only their own data, but also those of others. They will thus come to new insights, new findings and new solutions."



Ursula von der Leyen,
European Commission President
World Economic Forum in Davos,
January 2020

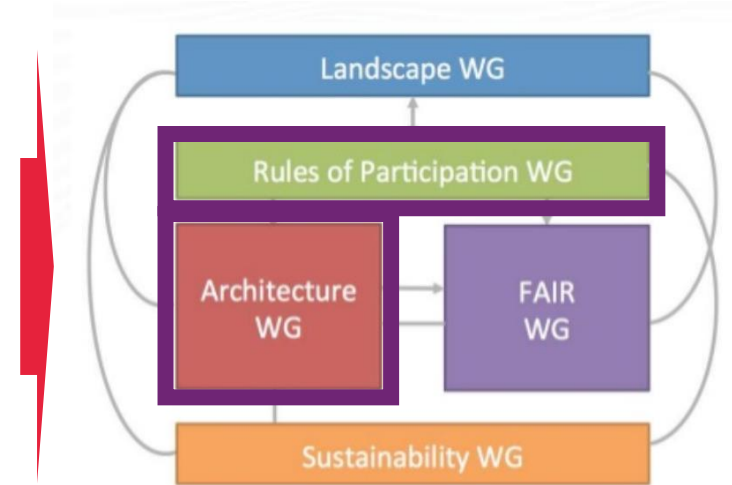
2016	I-2017	IV 2017	I-2018	II-2018	IV-2018	I-2019	II-2019	III-2019	IV-2019	I-2020	II-2020	III-2020	IV-2020	I-2021	II-2021	
GEANT																
EOSC Pilot																
eInfraCentral																
Freya																
EOSC Hub																
OpenAIRE-Advance																
RDA Europe 4.0																
PaNOSC																
EOSCsecretariat.eu																
FAIRsFAIR																
ENVRI-FAIR																
SSHOC																
EOSC-Life																
ESCAPE																
OCRE																
ARCHIVER																



sources: eosc-portal.eu, eoscsecretariat.eu

EOSC is much wider than what we have known before


Architecture	federated infrastructures and (data) resources
Data	FAIR data management and tools
Services	environment of user-oriented services
Access & Interface	access across disciplines, and compliance with open data
Rules	rules of participation for services, compliance with legal & trust needs
Governance	governance of the ecosystem and leadership in data-driven science



From: EOSC portal, by way of *The added value of EOSC for research in EOSC* Zoe Cournia (NI4OS-Europe)

Twinning the 'EOSC' to the e-infrastructure

'EOSC' could be seen as a twin sister (or brother) of the e-infrastructure organisations. One offering the compute and connectivity services and the other servicing the data and the interoperability.



13

source: Karel Luyben, IFFIS2019 conference, Stockholm November 2019
<https://www.slideshare.net/kbredaktion/european-open-science-cloud-205323223>

An ecosystem more than an 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
- Security & Operations
- Help Desk

EUROPEAN OPEN SCIENCE CLOUD

Open Consultation for the EOSC Strategic Research and Innovation Agenda

Have your say and let us shape the future of EOSC!

SAVE THE DATE: 20th July to 31st August

ACCESS EOSC SERVICES & RESOURCES


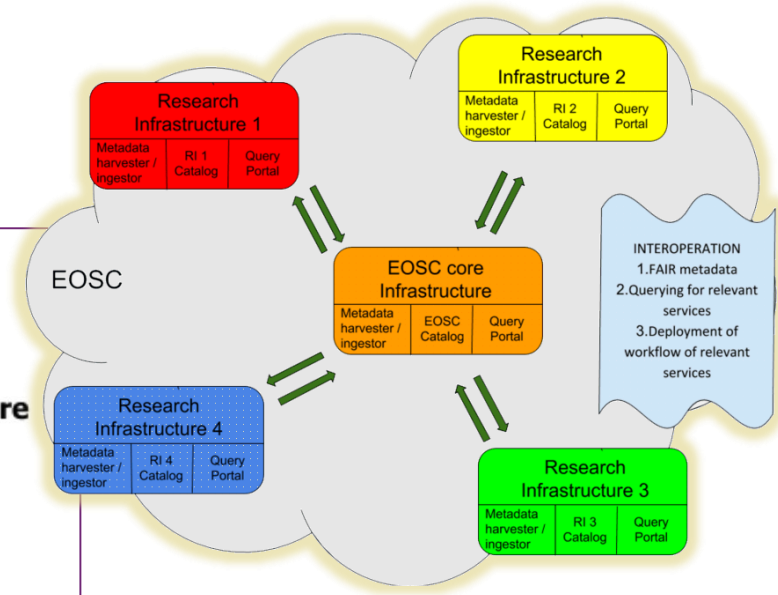
Icons representing various services: a computer monitor, a gear, a server rack, and a cloud.

With a structuring EOSC 'core'

Possible core functions for 'EOSC' in 2020+



- **Develop and govern federating core**
- **Manage compliance framework**
- **Manage trusted certification**
- **Manage 'EOSC' trademark(s)**
- **Implement PID policy guidelines**
- **Develop outreach to stakeholders**
- **Contribute to Horizon EU policy**
- **Monitor services and transactions**

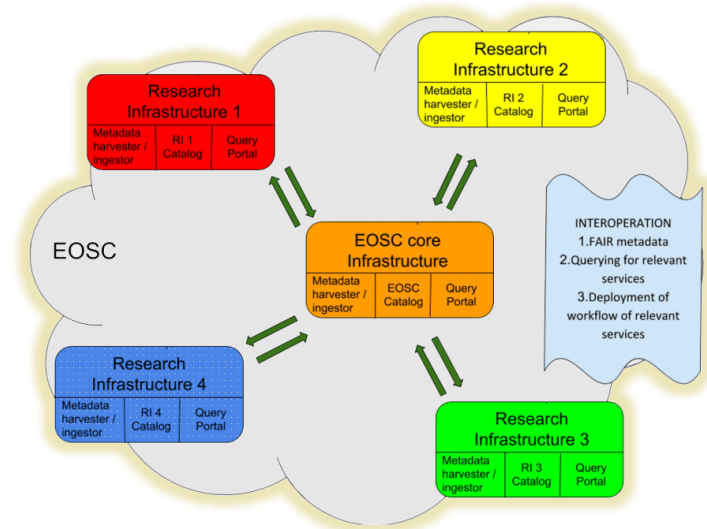


sources: EOSC Secretariat, Karel Luyben, EOSC-Future drafts, CNECT_RTD_Orientation_Skeleton_RIs_v14

Core services and 'the exchange'

What constitutes a 'core service'? A thin layer, with

- at least the portal itself
- finding & sharing of services, with a recommendation engine & messaging
- authentication and authorization, based on the 'AARC BPA'
- IT service management for the core
- operational security capabilities, trust policy, and security risk structuring



rest set via including criteria of the Sustainability and Architecture WGs

the *Architecture WG* and its taskforces will set the interoperability standards

A challenging landscape

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

An open ecosystem – rules of participation will favour no 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 – aim includes composability and
collective service design through an open AAI federation

Great (trust and security) expectations

‘the EOSC is a journey, and not at its final destination just yet’

Core

‘a distributed and participatory EOSC-Core in a collaborative way by reaching consensus on interoperability standards, APIs, and their implementation via best practices’

Exchange & Portal

‘research enabling services’

- national & regional
- institution & domain based
- including commercial providers

‘a catalogue ... covering the full research life cycle’

So what are the requirements on each? And the interdependencies?

Minimum Viable EOSC

today's world is agile, so focus is on the

'MVE – MINIMUM VIABLE EOSC'

including mechanisms to encourage adoption through policies

For trust & security, who should provide that capability?

- the infrastructures, or the service providers?
- a core team near the EOSC portal? Or (also) close to the AAI?

it will be a mix, but in all cases providers will play an important role

... and Sirtfi shows that is not completely unrealistic

MVE will emerge if political, technical and human/sociological conditions are met

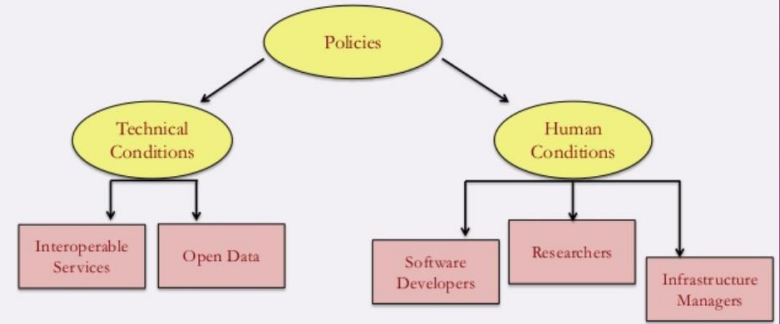


Photo: Patrick Perkins (Unsplash)

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

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



From *promoting and monitoring capabilities* to *managing 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

Making the EOSC a trusted place

Risk-centric self-assessment framework

- based on federated ISM guidance including WISE SCI

Baselining security policies & common assurance

- AARC, REFEDS, IGTF, PDK & practical implementation measures

fostering trust through a known skills programme

- so that your peers may have confidence in service provider abilities

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

Assessing risk ... in a peer-review framework

An information security risk assessment framework for EOSC services based on a federated evolution of the WISE SCI framework and a multi-tier maturity model, inclusive of data security and protection.

- risks 'play out' differently in different infrastructures
- more than just storage or compute, but also risks for (open!) data

Many risks are common, some need domain expertise to assess. Or are under regulated regime

WISE COMMUNITY

A Trust Framework for Security Collaboration among Infrastructures
SCI version 2.0, 31 May 2017

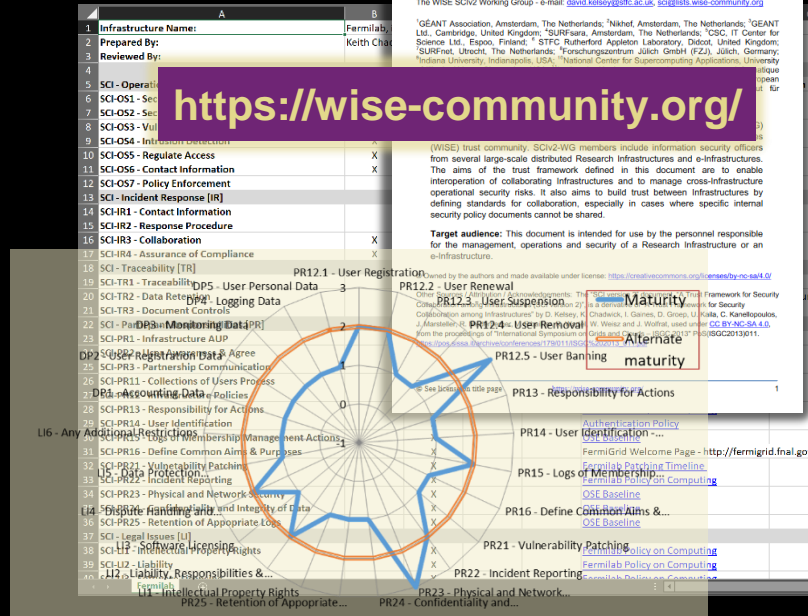
L Florio¹, S Gabriel², F Gagadis³, D Groep⁴, W de Jong¹, U Kaila⁵, D Kelsey⁶, A Moens⁷, I Neilson⁸, R Niederberger⁹, R Quisk¹⁰, W Raquet¹¹, V Rbailier¹², M Sallé¹³, A Scicchitano¹⁴, H Short¹⁵, A Stiegel¹⁶, U Stevanovic¹⁷, G Venkatesh¹⁸ and R Warter¹⁹

The WISE SCI Working Group - e-mail: david.kelsey@brf.ac.uk, sci@hds.wise-community.org

https://wise-community.org/

(WISE) trust community. SCIv2-WG members include information security officers from several large-scale distributed Research Infrastructures and e-Infrastructures. The aims of the trust framework defined in this document are to enable interoperability of collaborating infrastructures and to manage cross-infrastructure operational security risks. It also aims to build trust between infrastructures by defining standards for collaboration, especially in cases where specific internal security policy documents cannot be shared.

Target audience: This document is intended for use by the personnel responsible for the management, operations and security of a Research Infrastructure or an e-Infrastructure.



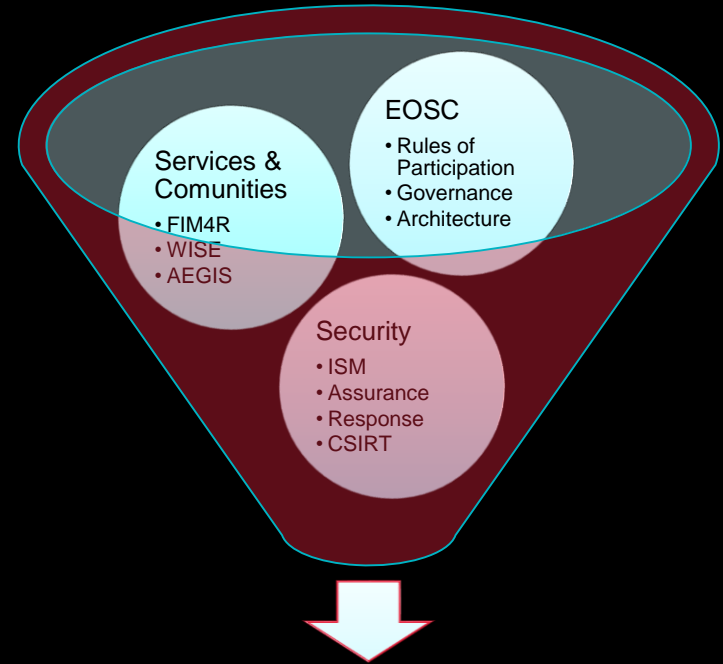
this spider diagram is fictional – idea by Urpo Kaila, CSC

Managing a policy baseline and assurance

A diverse set of requirements

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

and remain practical and manageable



security baseline, trust
and assurance profiles

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



Image credit: ZULTAX, <https://www.youtube.com/watch?v=NRznoYCJOHg>

Diversification is complex

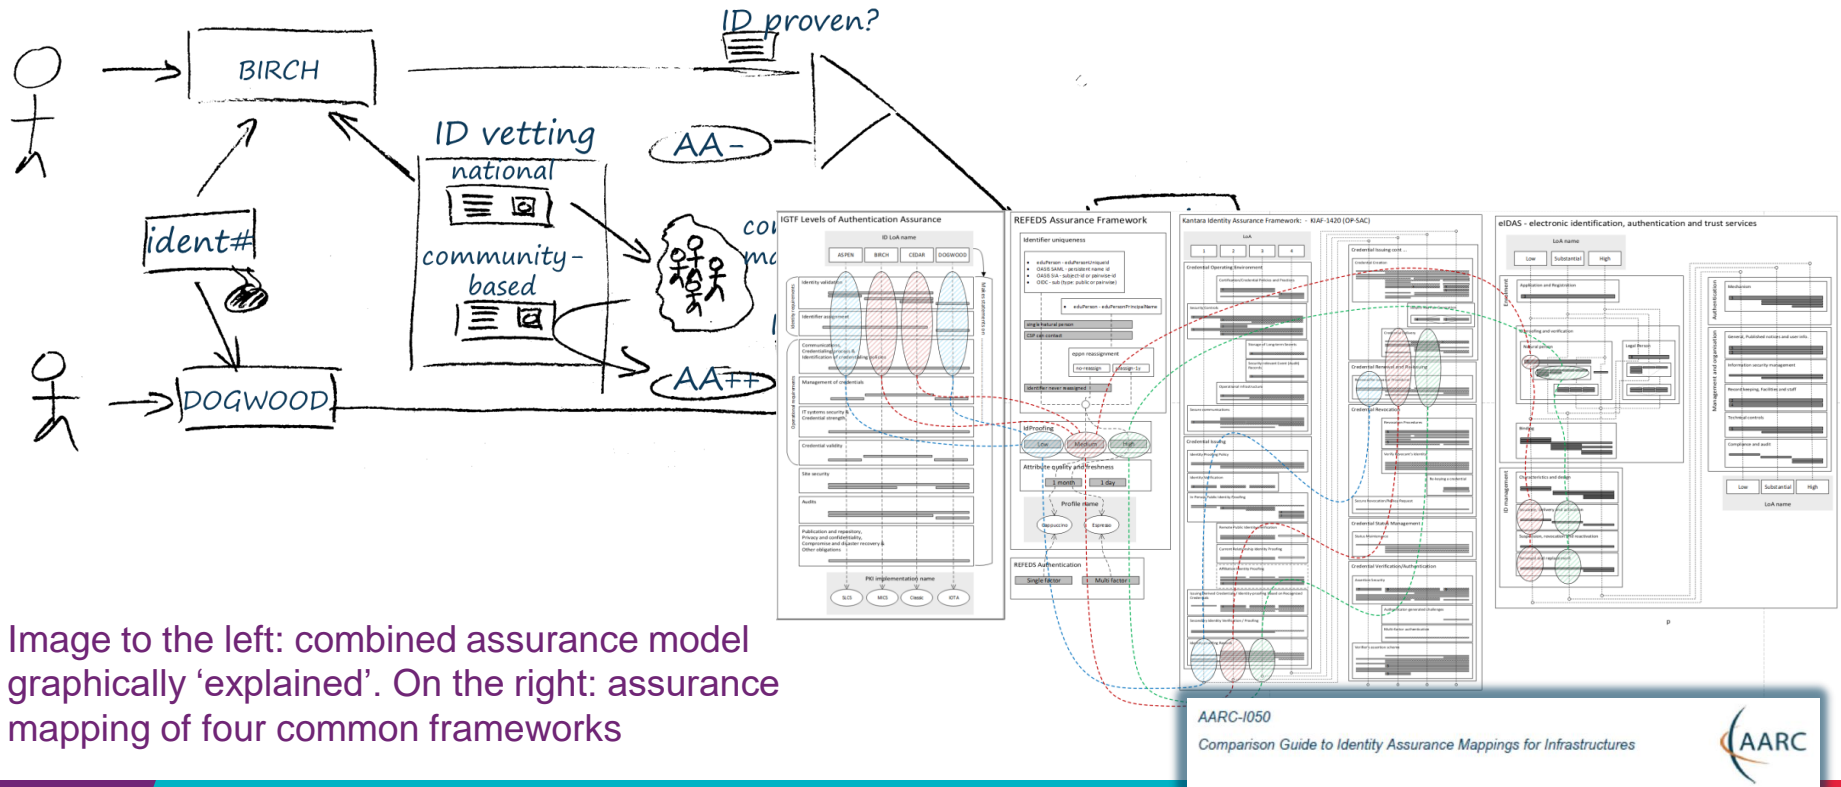
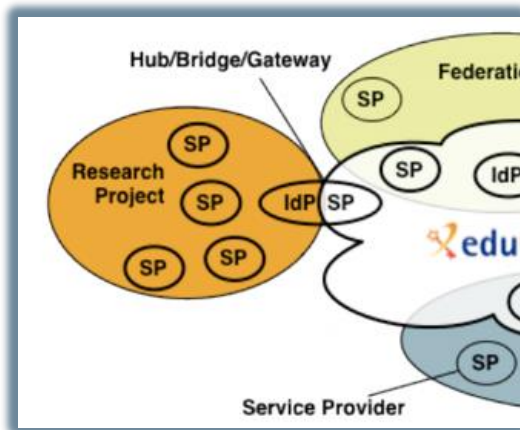


Image to the left: combined assurance model graphically 'explained'. On the right: assurance mapping of four common frameworks

Snctfi, maybe?

Scalable Negotiator for a Community Trust Framework in Federated Infrastructures



DERIVED FROM **SCIV2**:
FRAMEWORK ON
**SECURITY FOR
COLLABORATION IN
INFRASTRUCTURES VIA
WISE**

REFERENCE POLICIES
SUPPORTING **SNCTFI**
FULFILMENT IN THE POLICY
DEVELOPMENT KIT

Guidelines

The **AARC Guidelines** complement the **AARC Blueprint Architecture (BPA)** and the **policy best practices** recommended by the AARC project. The guidelines can apply to any topic that helps to advance Federated Identity Management for research and collaboration.

The AARC Guidelines help communities and infrastructures to implement and operate an AAI for research and collaboration more effectively and in an interoperable way.



Architecture Guidelines

Policy Guidelines

Targeted Guidelines

Upcoming Guidance

AARC-G014 Security Incident Response Trust Framework for Federated Identity

Sniff provides a mechanism to identify trusted, operationally secure eduGAIN participants and facilitate effective incident response collaboration

... more information ...

AARC-G015 Scalable Negotiator for a Community Trust Framework in Federated Infrastructures

The Sniff framework identifies operational and policy requirements to help establish trust between an infrastructure and identity providers either in an R&E Federation or in another infrastructure, in each case joined via a Service Provider to Identity Provider proxy

... more information ...

AARC-G021 Exchange of specific assurance information between

infrastructures and generic e-infrastructures compose an 'effective' assurance profile derived from resulting assurance assertion obtained between infrastructures so that it need not be re-computed by the infrastructure provider. This document describes the assurance profiles recommended to be used by the infrastructure

... more information ...

Architecture Guidelines

Policy Guidelines

Targeted Guidelines

Upcoming Guidance

AARC-G040 Preliminary Policy Recommendations for the LS AAI (application to R&S and CoCo)

The Life Sciences AAI Service (LS AAS), developed in joint collaboration with EDG, EUDAT and GÉANT, will result in a production-equivalent service to be operated for the Life Sciences community by the joint e-infrastructures. As the pilot enters its second phase the LS AAI has to declare compliance to R&S and CoCo towards the R&E federations. This document provides preliminary guidance for the operators of the pilot LS AAI

... more information ...

**aarc-
project.eu/guidelines**

***Snctfi covers both
service-centric and
some researcher-
centric policies***

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 template provides a step-by-step breakdown following a security incident.
Membership Management Policy	Infrastructure Management	Research Community (abides by)	This policy template defines the rules for their members, including their expiration.
Acceptable Use Policy	Infrastructure Management	Research Community	This is a placeholder for acceptable assurance profiles of user credentials.
Policy on the Processing of Personal Data	Infrastructure Management & Data Protection Contact	Research Community, Services (abide by)	This table can be used for identifying when a Protection Impact Assessment is required. This document defines the obligations on Infrastructure Participants processing personal data.
Service Operations Security Policy	Infrastructure Management	Services (abide by)	This can be used to define requirements for collected and processed data in the infrastructure. This policy defines the requirements for running a service within the Infrastructure.
Acceptable Use Policy	Infrastructure Management (for baseline) & Research Communities (for community specific restrictions)	Users (abide by)	This is a template for a policy that users must follow. Research Infrastructure augmented by the Research Communities.

Showing 1 to 9 of 9 entries

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
Acceptable Authentication Assurance	Infrastructure Management	Research Community, Services (abide by)	This is a placeholder for the Infrastructure to determine rules for the acceptable assurance profiles of user credentials.
Policy on the Processing of Personal Data	Infrastructure Management & Data Protection Contact	Research Community, Services (abide by)	This document defines the obligations on Infrastructure Participants when processing personal data.
Service Operations Security Policy	Infrastructure Management	Services (abide by)	This policy defines requirements 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.

[HTTPS://AARC-PROJECT.EU/POLICIES/POLICY-DEVELOPMENT-KIT/](https://AARC-PROJECT.EU/POLICIES/POLICY-DEVELOPMENT-KIT/)

But also the 'PDK' complexity will need to be managed ...

Questions to ask yourself when defining this policy:

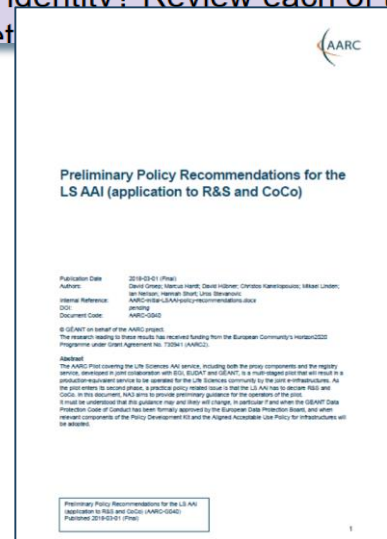
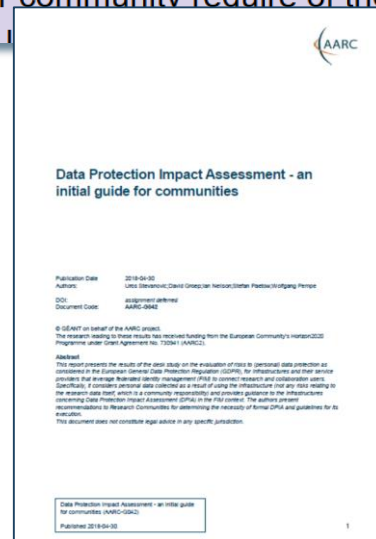
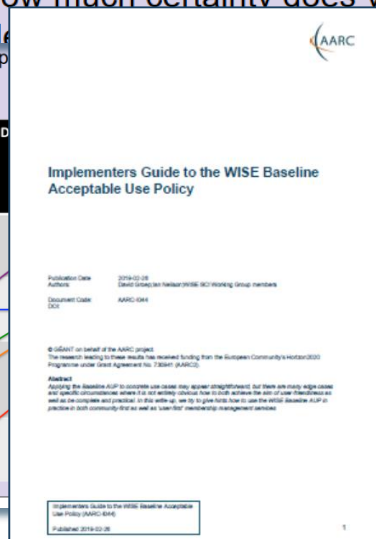
- Which identity providers are acceptable for your infrastructure? SAML Identity Federation IdPs? Social providers such as Google, Facebook etc?
- How much certainty does your community require of the identity? (e.g. authentication strength). How much certainty does your community require of the identity?
- How can you ensure that each user has the capability at their home organization to use your services, or a subset, of your services?

Questions to ask yourself when defining this policy:

- Which identity providers are acceptable for your infrastructure? SAML Identity Federation IdPs? Social providers such as Google, Facebook etc?
- How much certainty does your community require of the identity? Review each of the

The following chart can be used to help determine an appropriate level of assurance. Refer also to [AARC Guideline 21](#):

Should identifiers be unique, personal and traceable?	Should identifiers be unique across the infrastructure?	How fresh do attributes need to be?	What kind of ID Proofing is required?
Unspecified	Unspecified	Unspecified	Unspecified
Yes	Yes	1 month	Low (self asserted)
			Medium (e.g. postal credential delivery)
			High (e.g. face to face)



Start with baselining

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

Good Practice

implementation guidance

small number of assurance profiles (REFEDS, IGTF, eIDAS), AARC secure operations standards, AEGIS recommendations, CSIRT capability

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 ...

Technical guidance

e.g. expression of identity assurance


Rules of Participation

minimal set of capabilities –initially maybe just contact information and responsiveness


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 EOOSC portal providers and research communities
- desirable, but not yet likely, to have training a requirement for participation *but hard to realise for an EOOSC that does not wish barriers to entry* 😞



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



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'

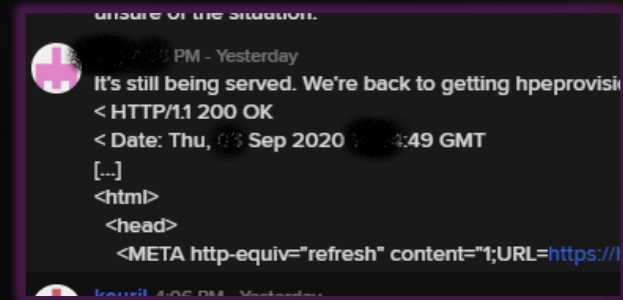


Actionable Response – coordination involving the Core

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
- it will have a tightly-knit team of experts looking after the security of the core
- who can work collaboratively with peer infrastructures and groups

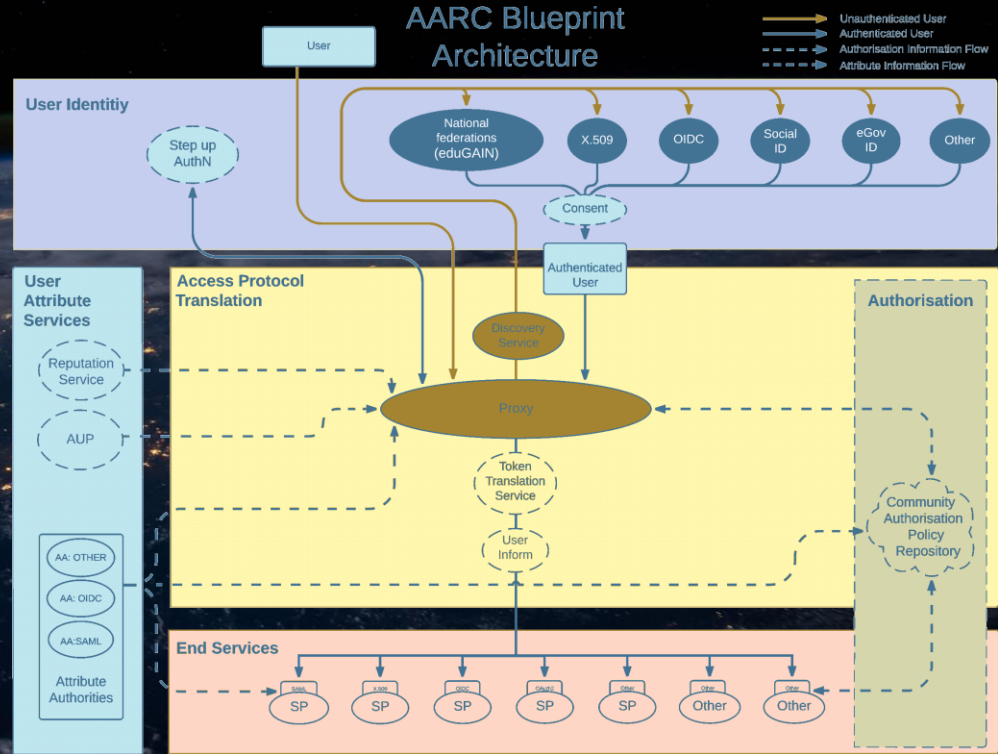
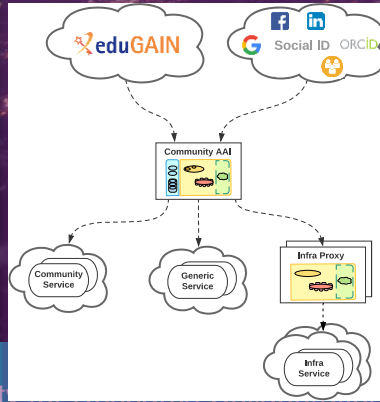


this team is essential to glue together the information during incidents – leveraging the trust built up before through engagement

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

... we really heard that one ...

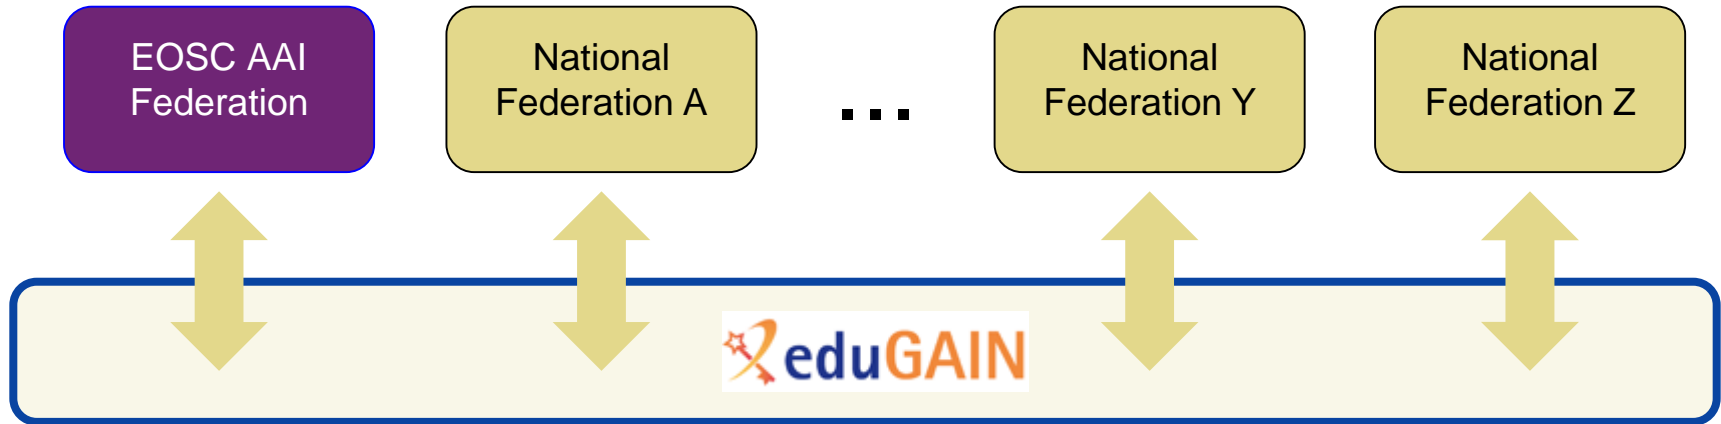
and although the AAI is a core service of the EOSC, it only does 'what it says on the tin'



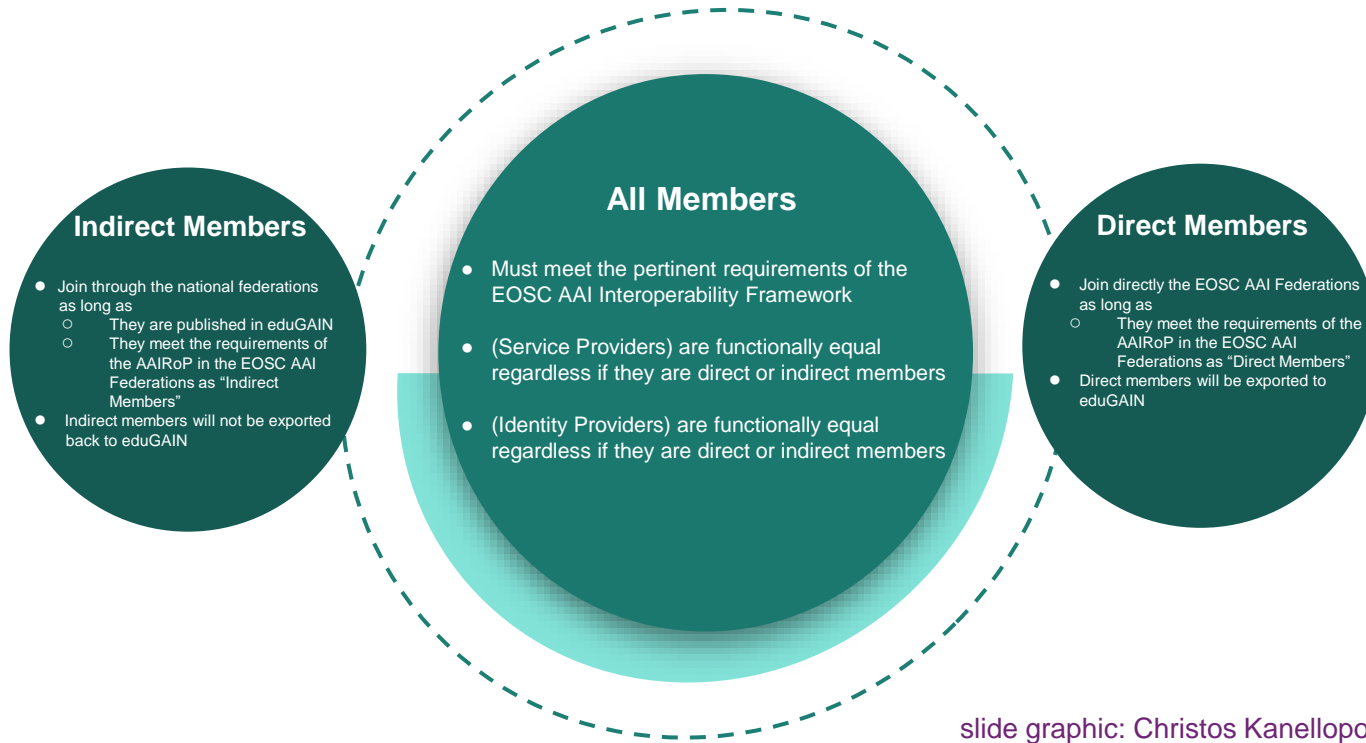
Linking the providers and users together - AAI

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



Linking into the EOSC federation



It has to be linked to eduGAIN, and both the EOSC and eduGAIN should mutually strengthen each other.

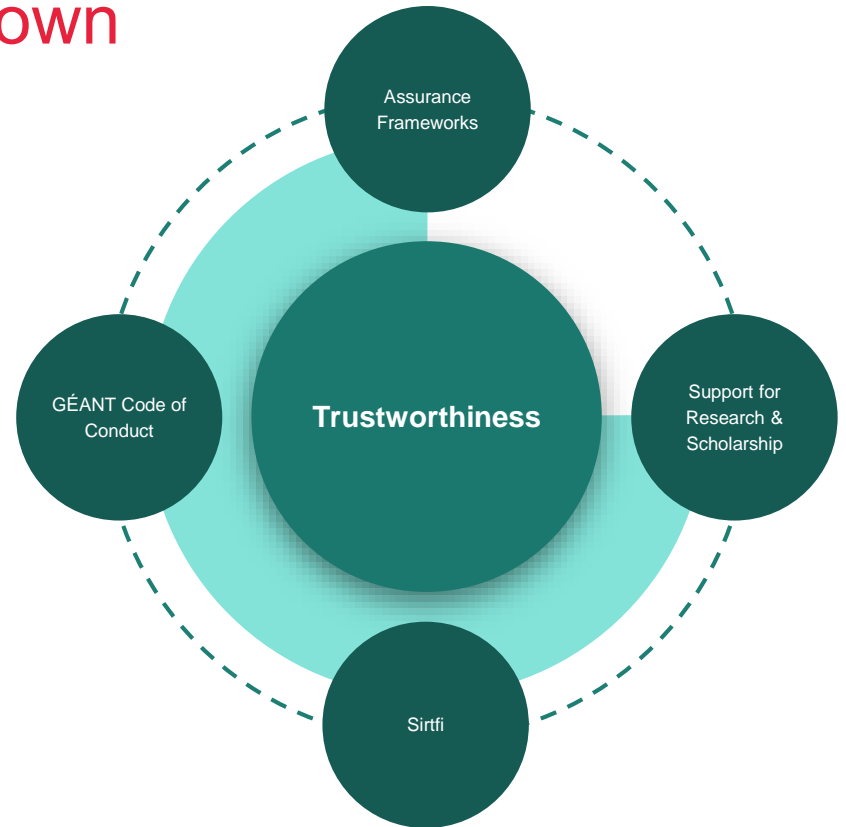
Given the broad reach of the EOSC, it may well contain new entities, both from the private sector and from international collaborations and research infrastructures

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

But now ... turtles all the way down

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

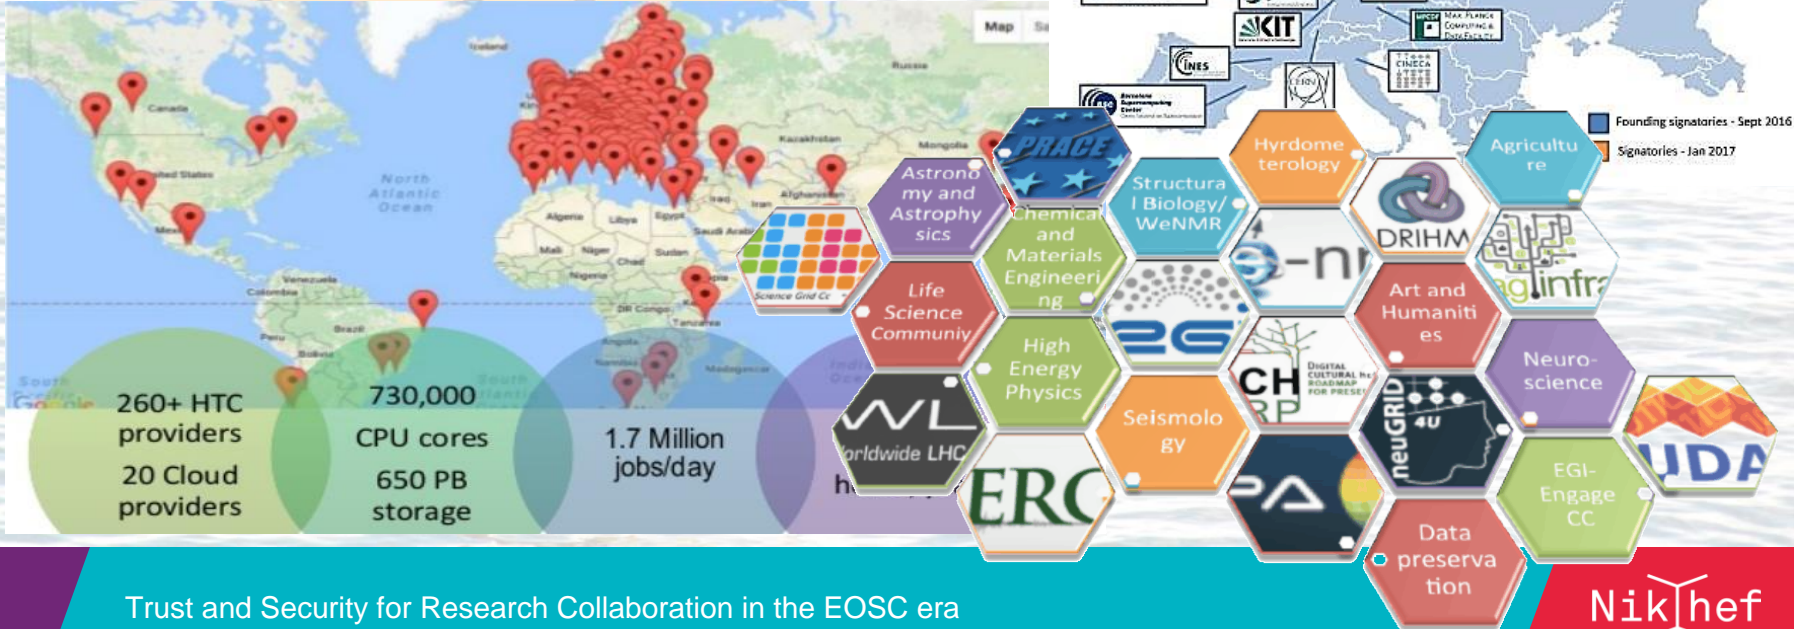
- inspired by eduGAIN constitution and other sources
- leveraging existing trust frameworks
- and not repeating earlier mistakes so implement a baseline at the start



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

So if not the AAI ... who then?

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

- an **infrastructure** – both providing and using services – can provide coordination amongst similar services, making that much easier
- security merit is that service providers in an infrastructure can **benefit from their commonalities** in response & security management
- and for the EOSC that a mature security capability can be structured at the infrastructure level in a scalable way across many service providers

and remember 'services' are very broad and includes data, publications, &c

Profiting from having a common set of services

egi CSIRT

Home Activities Materials Trainings Contacts News About

EGI CSIRT

Computer Security Incident Response Team

MISSION
EGI CSIRT coordinates operational security activities within the EGI Infrastructure to deliver a secure and stable infrastructure, giving scientists and researchers the protection and confidence they require to safely and effectively carry out their research.
[Find more about EGI CSIRT](#)

WHAT WE DO

- Prevention of security incidents (security monitoring, software vulnerability handling, risk assessment and mitigation)
- Incident response
- Security policy and procedures
- Security Trainings

[Find out more about our activities](#)

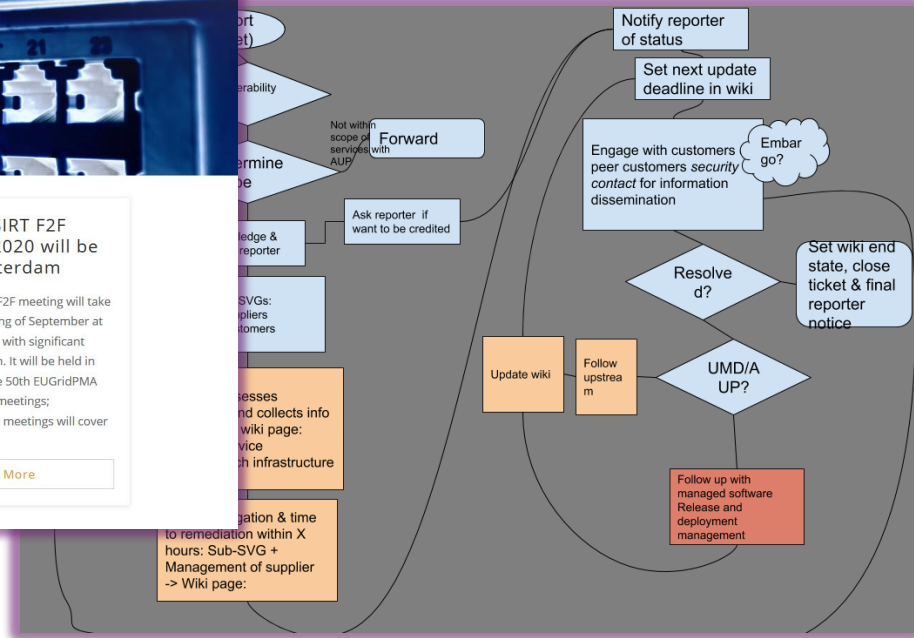
CONTACTS
[To report a security incident:](#)

TRAININGS
Keeping the EGI infrastructure secure requires an understanding of attack and defense

Third EGI-CSIRT F2F meeting in 2020 will be held in Amsterdam

The next EGI CSIRT F2F meeting will take place at the beginning of September at Nikhef, Amsterdam, with significant remote participation. It will be held in conjunction with the 50th EUGridPMA and EOSC-hub ISM meetings; discussions at these meetings will cover a range...

[Read More](#)



Complementarity within the service & infrastructures

- information security management maturity
looking after service integrity, responsive contacts, also for exercises, and monitoring for intrusions & vulnerabilities
- vulnerability assessment and management
pro-active security management in general
- incident response and resolution within the infrastructure and service
and smooth collaboration with the (EOSC) core team

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
- training and capability enhancement
- coherence of response, community readiness/collaboration, and information sharing
- resolution, forensics, resolution and remediation for core and stakeholders

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

Core in EOSC-Future



EGI

membership
contributions

VA
core

EUDAT

GEANT



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 a core team – incident response and forensics experts and coordinators – be busied consistently with service-specific response, where the portal would not be able to add to the trust 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



*based on the
white paper co-authored with
Jens Jensen, Dave Kelsey,
Daniel Kouřil, Maarten Kremers, and Hannah Short
and on discussions in the EOSC Future
Security Operations & Policy collaboration*

Nik|hef

David Groep

davidg@nikhef.nl

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



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

long read:

Groep, David L. *et al.*, Trust Coordination for Research Collaboration in the EOSC era

<http://doi.org/10.5281/zenodo.3674677>

Planning ahead!

Almost regardless of what happens next, we do need comprehensive security for the EOSC. If we don't act, or leave corners open, it will come back to haunt us.

- what the EOSC will be, is still being shaped
- yet connecting services, content, data will happen, and on a *much wider, more distributed, and multi-stakeholder* scale

We need to engage with the new and evolving stakeholders who will not know us – and likely not trust us – until we gain such trust outside our 'usual' zone

- education, awareness and training
- security exercises based on recognised trust frameworks + Rules of Participation
- ensure collaboration of everyone when time comes – we need the portal on board
- operational expertise, forensics, remediation, and demonstrable practical impact are key to success!

this work is co-supported by the Trust and Identity workpackage of the GEANT4 project - phase 3



THANK YOU

Read: *Trust Coordination for Research Collaboration in the EOSC era*
<http://doi.org/10.5281/zenodo.3674677>

davidg@nikhef.nl



NETWORKS · SERVICES · PEOPLE
WWW.GEANT.ORG



THIS WORK IS ALSO SUPPORTED BY A PROJECT THAT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NO. 856726 (GN4-3).