Nikhef

Making hardware, software, and people collaborate

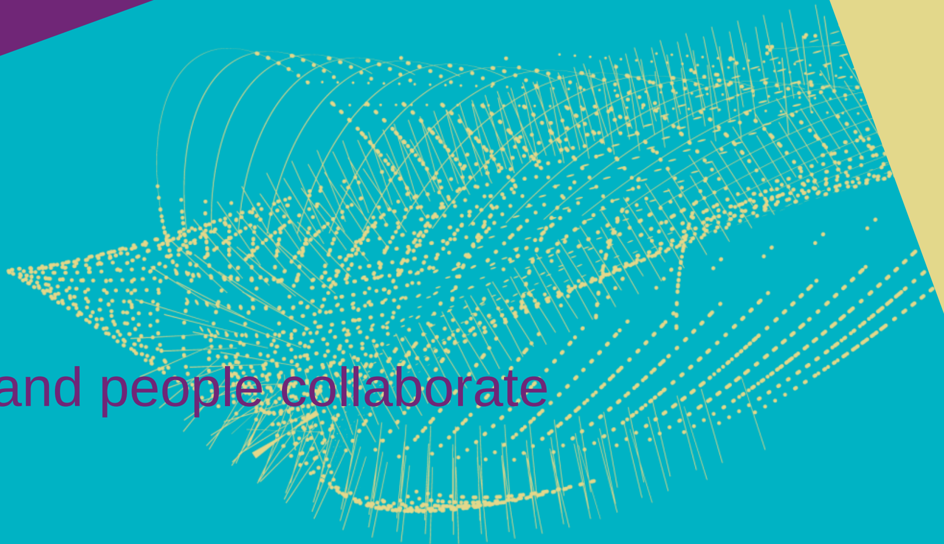

Building scalable e-Infrastructures for research

David Groep, Nikhef

November 10, 2020

Physics-DKE colloquium

Scaling & the 'data deluge' … with data comes processing

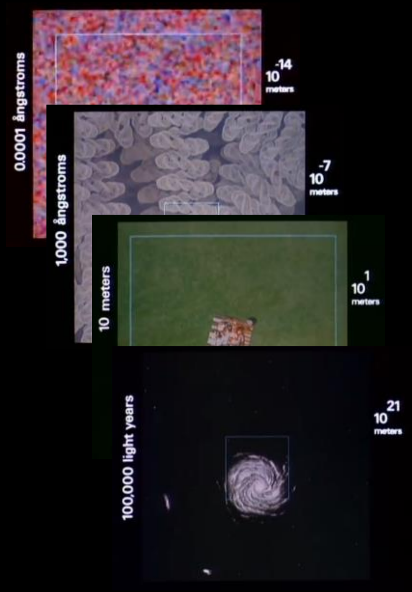

imagery: *Powers of Ten, Eames Office LLC (available at www.eamesoffice.com)*

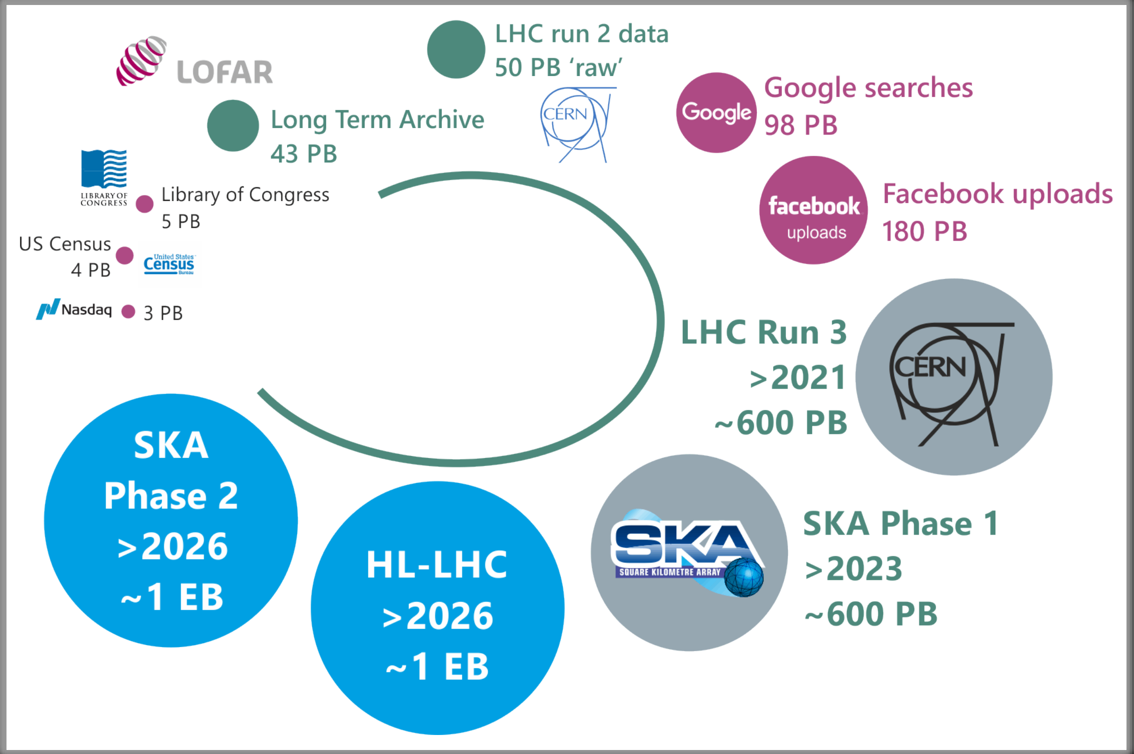

Data from various sources, for public entities: data ca. 2018, indicative, within ~ factor 2
LHC volumes: LCG Resource Scrutiny Group & CERN; 2020
SKA and LOFAR volumes: ASTRON/Michiel van Haarlem, 2020

Data is distributed – but collaboratively interpreted

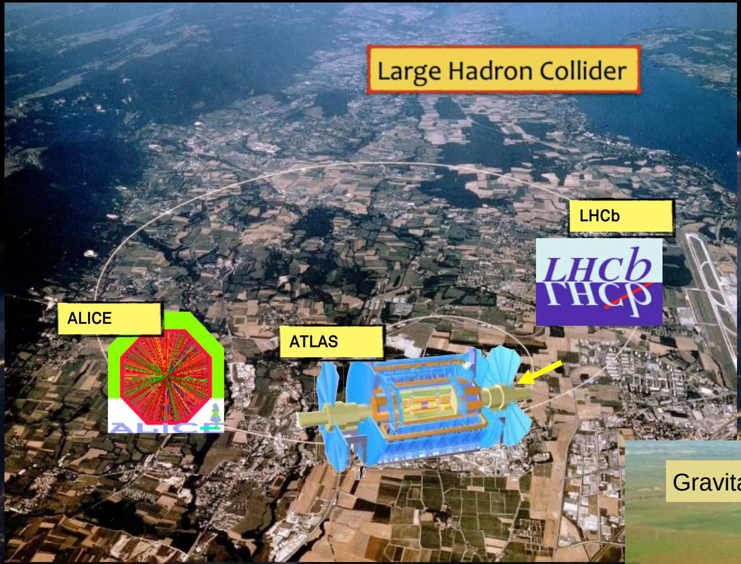

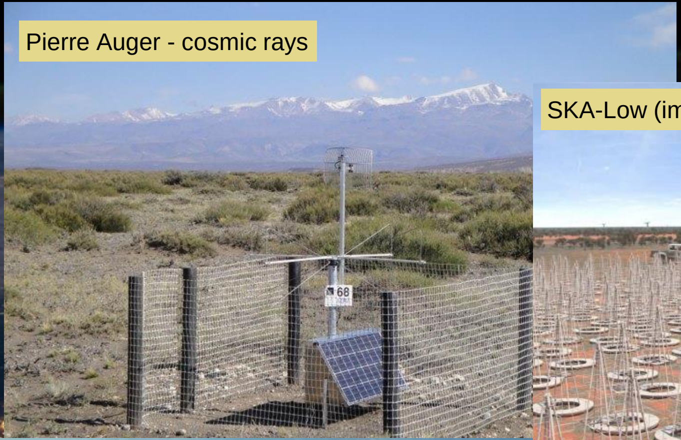

>170 institutes in >42 countries and economic regions

- *Computing* ~ 1,000,000 cores
- *On-line disks* > 310 PB
- *Archival* > 390 PB

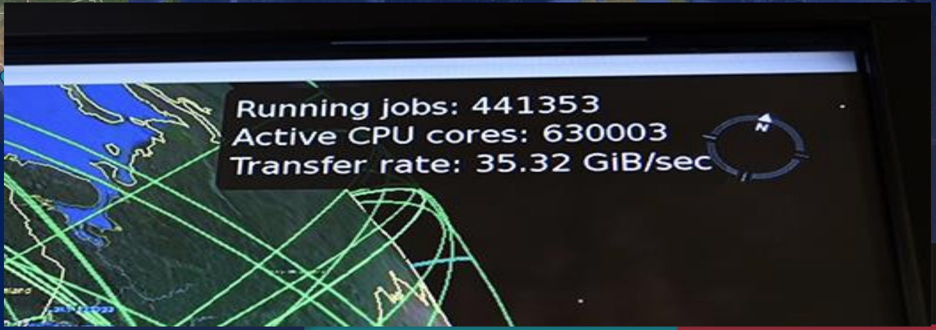

The High Luminosity Challenge

- The effort to reduce the gap between the estimations[...] resources is[...]
  - Changes in[...]
  - Improvem[...]
  - GPU comp[...]
  - Market m[...]
- Still, we are[...]

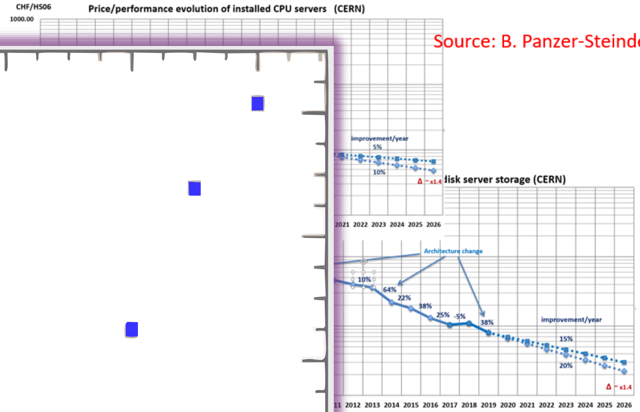

Source: B. Panzer-Steindel

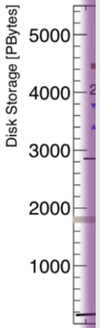

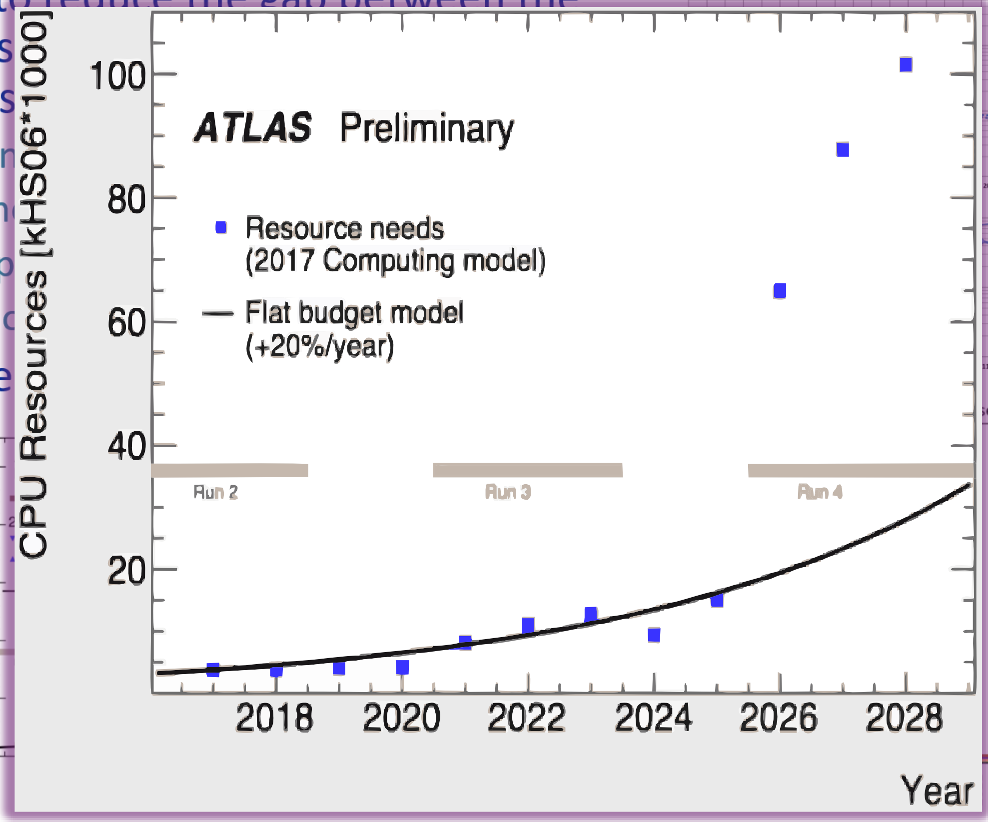

Source: D. Costanzo

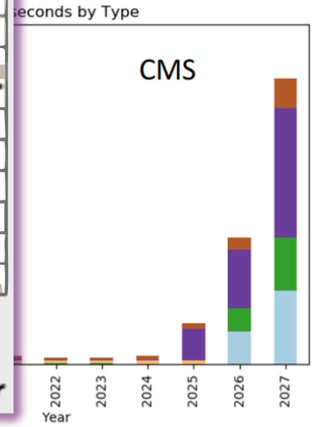

Source: D. Lange

source: Andrea Sciaba et al. for the WLCG Resource Evolution WG, CHEP2019

WLCG
Worldwide LHC Computing Grid

Scaling across all science domains

Dutch-only (!)
compute requirements for LHC, KM3NeT, and SKA, until 2030

Even then, it will be sufficient if and only if

- GPU and parallelism are fully exploited
- throughput per core continues to increase
- data access patterns will match system design

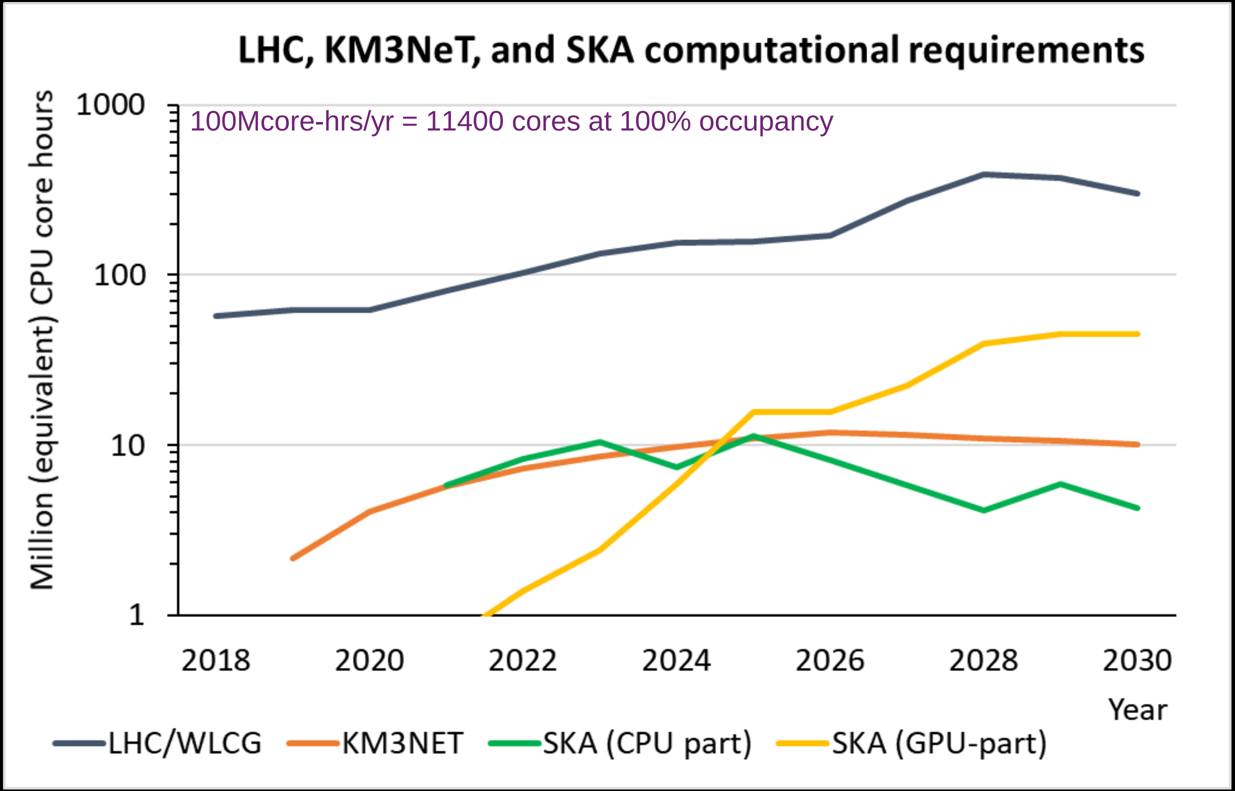

Source: FuSE, the Fundamental Sciences E-infrastructure, 2019
https://fuse-infra.nl/

Infrastructure: dealing with data processing at scale

1: matching algorithms and systems design

- designing for high-performance processors
- rethinking design patterns for work & data orchestration

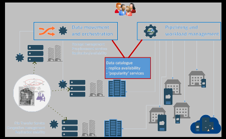

people - systems

2: collating compute, storage, and networks

- building 'facilities'
- peering and global networks
- stressing networks
- research 'cloud' services

systems - systems

3: accessing services, collaboratively & securely

- community building in a multi-national federation
- global trust and identity
- securing the infrastructure of an open science cloud
- our National e-Infrastructure

people - people

"It's just hardware"

performance goes up, doesn't it?

- processor performance
- memory bandwidth and on-die caching
- accelerators (GPU, &c)

- wider and faster PCI bus throughput
- faster storage and global interconnects

… but …

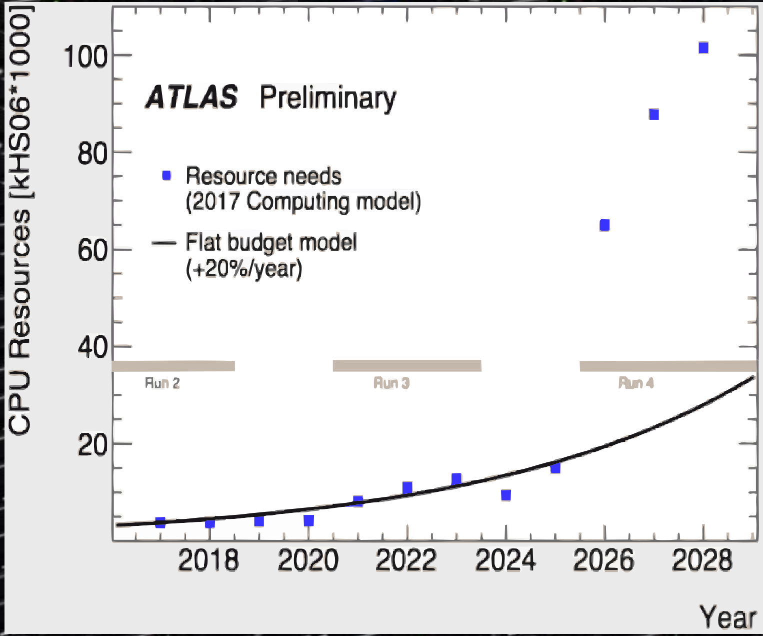

Atlas image source: Simone Campana for WLCG, 135th LHCC Meeting Open Session, September 2018

What we knew was coming …

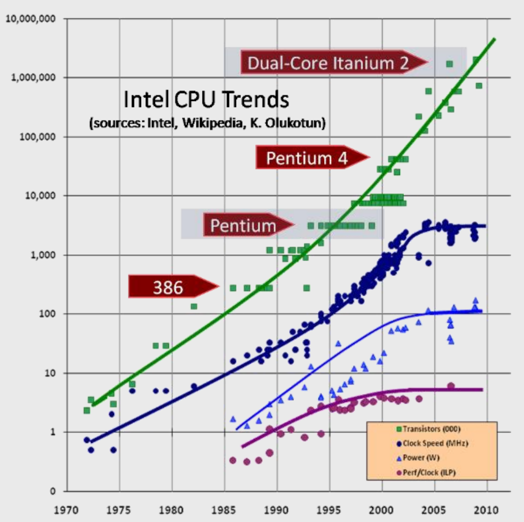

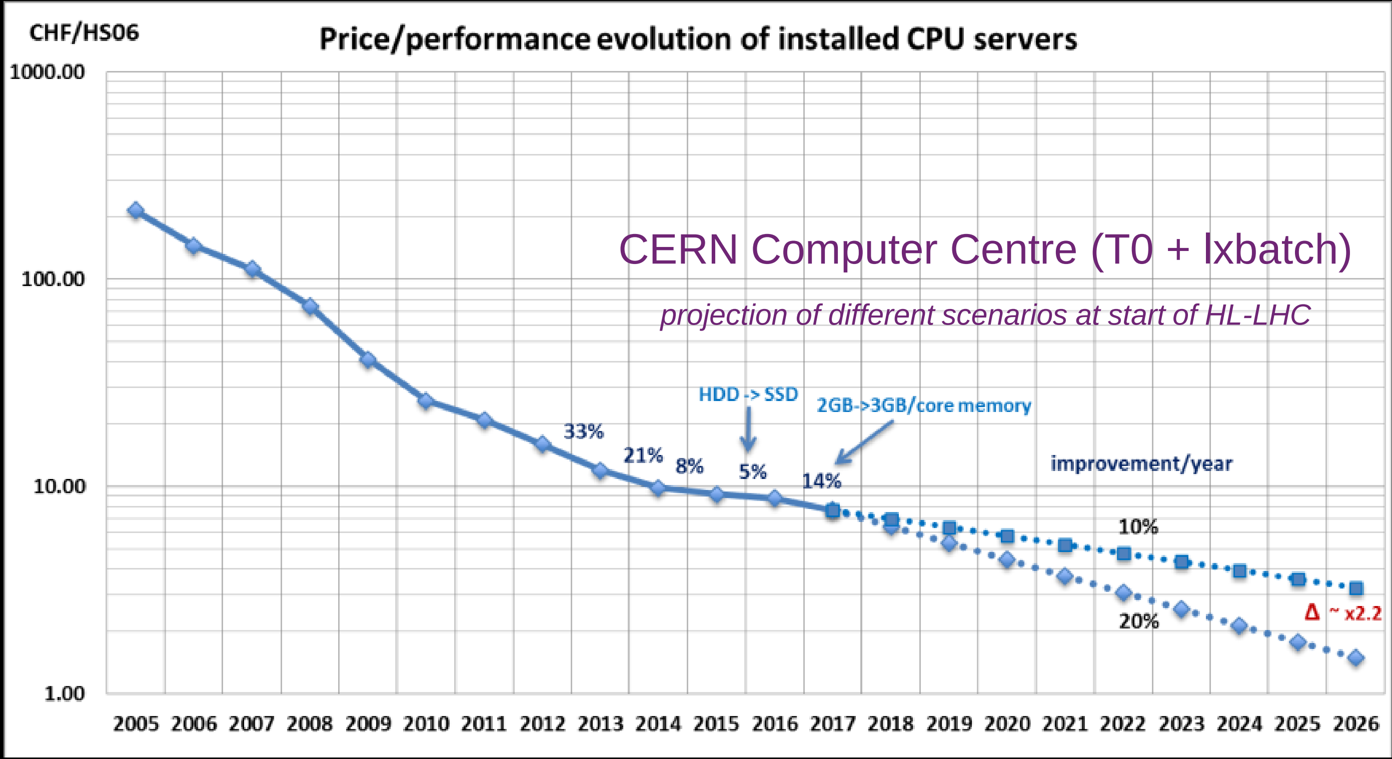

Helge Meinhard, Bernd Panzer-Steindel, Technology Evolution, https://indico.cern.ch/event/555063/contributions/2285842/

Figure left: Herb Sutter, Dr.Dobbs Journal 2004, updated 2009, see http://www.gotw.ca/publications/concurrency-ddj.htm

Yet exploiting even these improvements need people …

… *and implementations that take ‘hardware’ into account*

Application performance (“HEPSPEC06”) diverging from system capability (“SpecINT”)

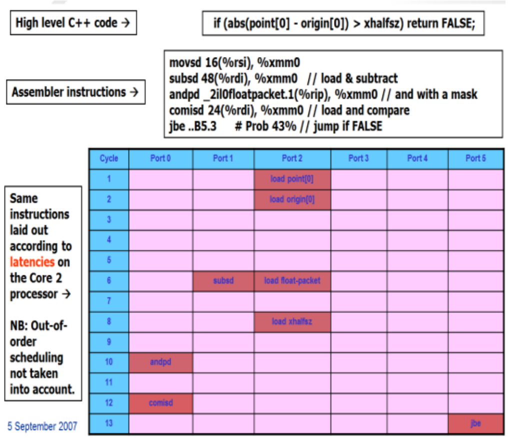

High level C++ code →

```
if (abs(point[0] - origin[0]) > xhalfsz) return FALSE;
```

Assembler instructions →

```
movsd 16(%rsi), %xmm0
subsd 48(%rdi), %xmm0   // load & subtract
andpd _2il0floatpacket.1(%rip), %xmm0 // and with a mask
comisd 24(%rdi), %xmm0 // load and compare
jbe ..B5.3     # Prob 43% // jump if FALSE
```

Same instructions laid out according to latencies on the Core 2 processor →

NB: Out-of-order scheduling not taken into account.

Cycle	Port 0	Port 1	Port 2	Port 3	Port 4	Port 5
1			load point[0]			
2			load origin[0]			
3						
4						
5						
6		subsd	load float-packet			
7						
8			load xhalfsz			
9						
10	andpd					
11						
12	comisd					
13						jbe

5 September 2007

*2007 Core 2 efficiency: Sverre Jarp, CHEP **2007** (!)*

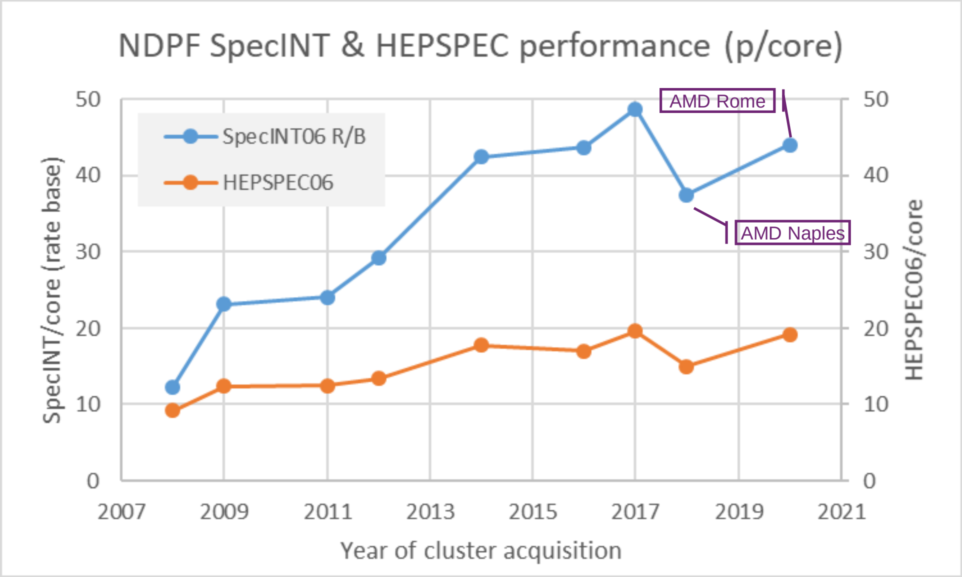

SPEC benchmark: spec.org, «Rate Base» (R/B) measures throughput under full load of all cores

Graph: measured HS06 and registered SpecINT06 Rate (base) performance per core, with SMT disabled, for the Nikhef Data Processing Facility NDPF (HTC compute)
line is not mononotically increasing because of other design choices (power efficiency) and price-performance optimisation chosen

And of course depends on hardware & CPU architecture

AMD “Naples to Rome” – boost in application (“HS06”) performance due to new memory (I/O) architecture and direct access to all memory banks

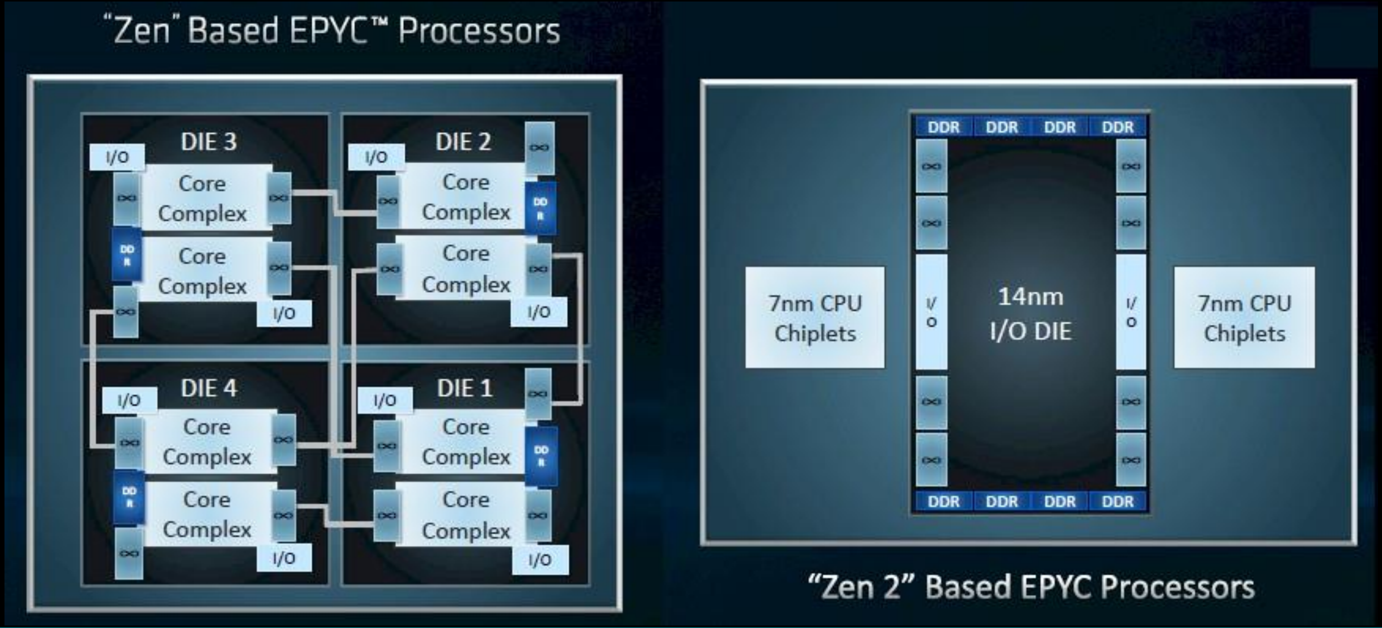

Image by AMD
retrieved from https://www.nextplatform.com/2018/11/06/amds-long-road-from-naples-to-milan-centers-on-rome/

And why some changes will not impact performance at all

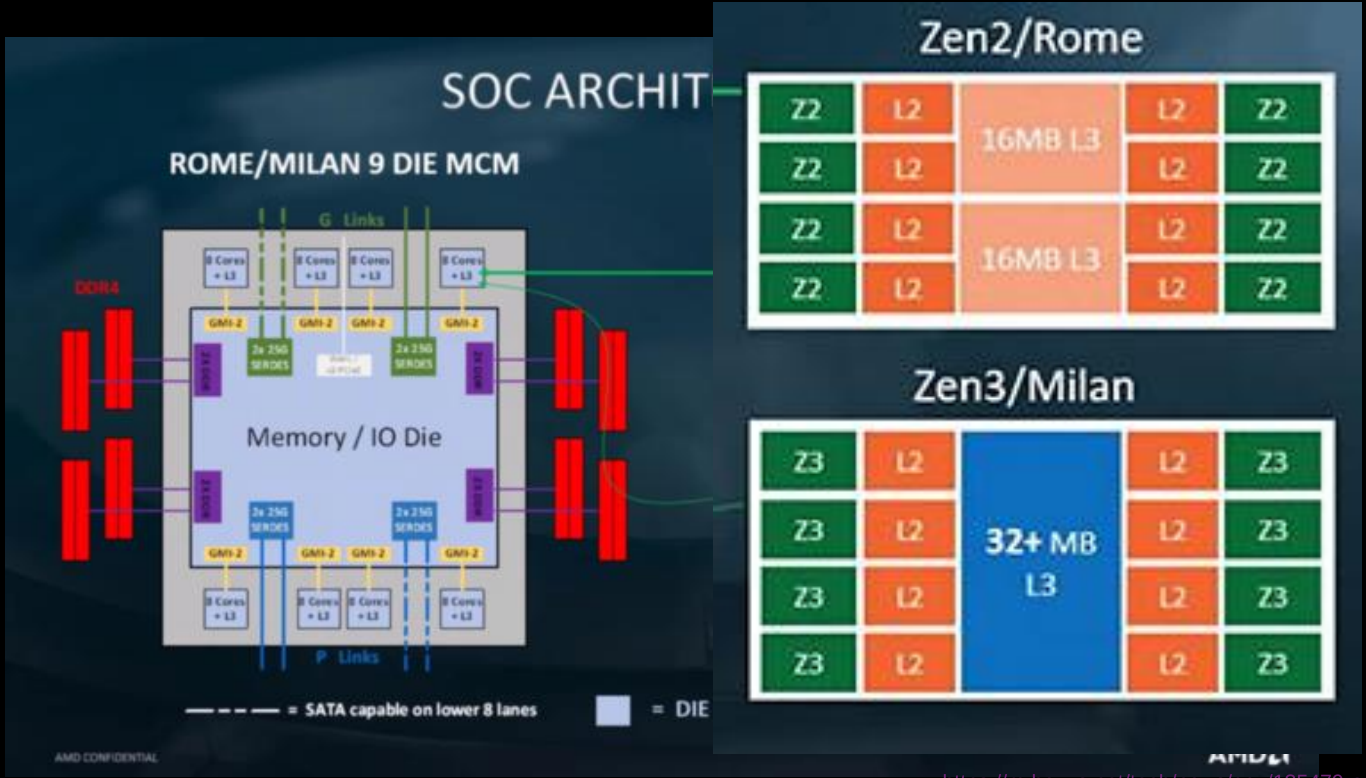

Image source: AMD, retrieved from https://m.hexus.net/tech/news/cpu/135479-amd-shares-details-zen-3-zen-4-architectures/

Nikhef

Bigger is Better - if you keep it together

Common element: moving data is ‘expensive’, so
‘keep on computing as long as you can, and don’t move data around’

- e.g. AMD (and for others: single-socket systems), are better since there are *no (useless) cache coherency delays and improved direct memory access*
- similarly, keep your GPU busy … as data comes from (slow) RAM

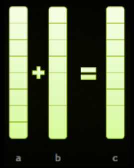

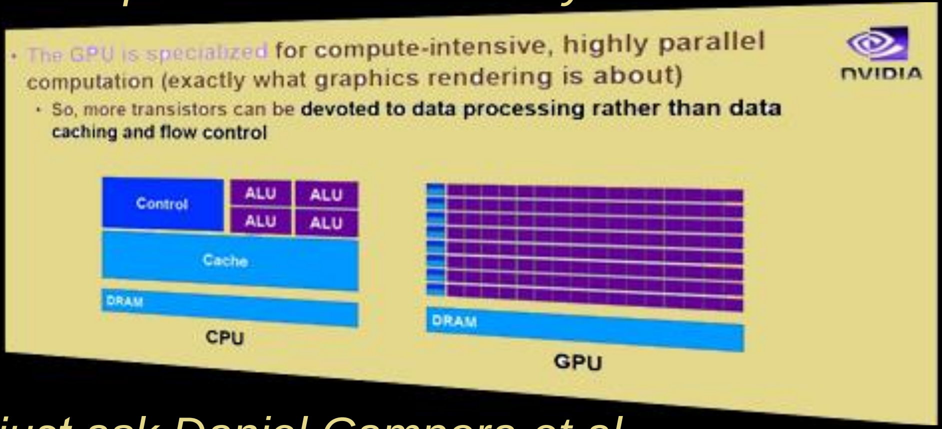

*Getting to be a quite specialised field – use **frameworks** to implement key code*

… or just ask Daniel Campora et al.

Image sources: NVidia ‘Massively Parallel Computing with CUDA’

and if it doesn’t quite fit …

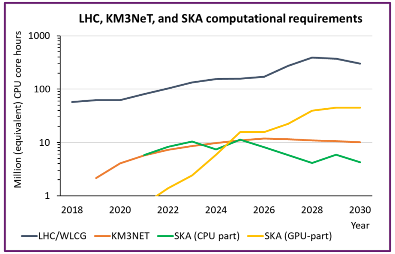

SuperMicro (branded as ‘Lambda Blade’)
4U chassis, supporting 10 consumer-grade GPUs …
… with a bump

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

Beyond the single box

Luckily, many things in this world are *conveniently parallel*

- HEP events & simulation 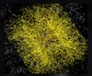
- ligand matching 
- structural biochemistry 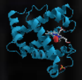
- …

```
korf.nikhef.nl:
                                                              Req'd   Req'd       Elap
Job ID                  Username    Queue    NDS   TSK   Memory  Time    S   Time
----------------------- ----------- -------- ----- ------ ------ --------- - ---------
33134895.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  37:46:21   wn-choc-023
33134901.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  40:04:09   wn-smrt-128
33134908.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  37:14:29   wn-choc-030
33134917.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  14:23:42   wn-smrt-072
33135197.korf.nikhef.n  atlb019     atlasmc      1      4  16040 208:00:00 R 183:02:04   wn-mars-018+
wn-mars-018+wn-mars-018+wn-mars-018
33135883.korf.nikhef.n  atlb019     atlasmc      1      4  16040 208:00:00 R 166:44:22   wn-mars-018+
wn-mars-018+wn-mars-018+wn-mars-018
33142633.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  37:30:47   wn-mars-043
33149106.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  10:23:30   wn-car-027
33149132.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:57 R  32:36:49   wn-mars-057
33149220.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:5                 4
33151669.korf.nikhef.n  lhcbpi08    lhcb         1      1  5120m  41:59:5                 9
33152704.korf.nikhef.n  atlb019     atlasmc      1      4  16040 208:00:0                 8+
wn-mars-018+wn-mars-018+wn-mars-018
```

challenge is not in the parallelism itself,
but in **global compute with data**
just like difference between SI06 and HS06 showed data as the driving factor

Image: Nikhef D0 farm in 2001

Conveniently parallel: a global infrastructure for research

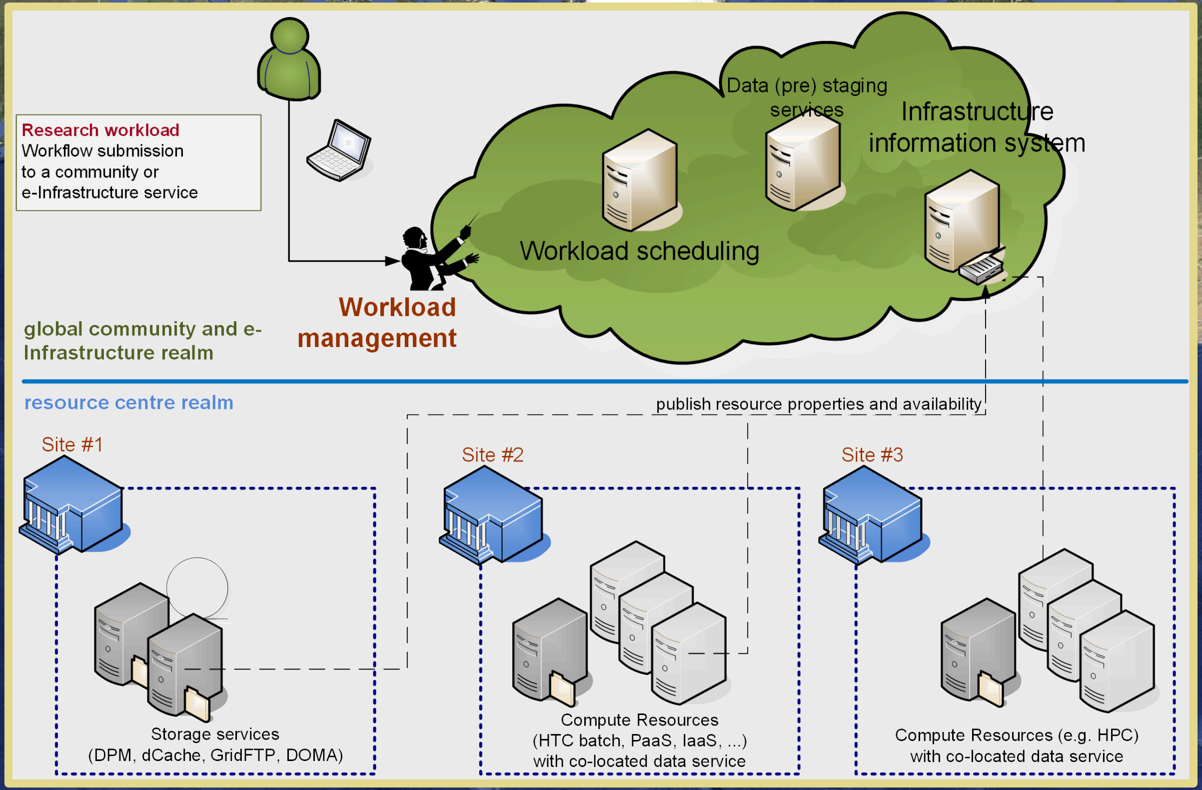

shared multi-community infrastructure

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

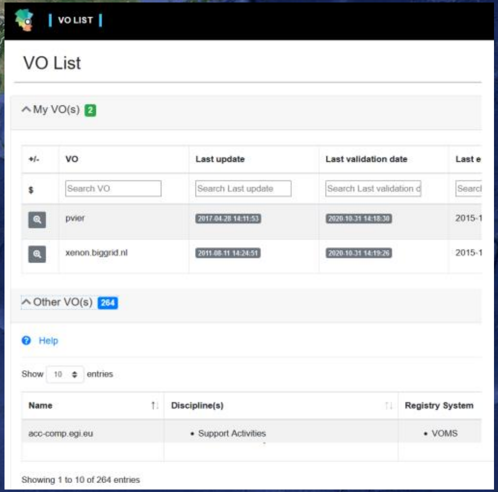

Right-hand graphic: EGI operations portal, https://operations-portal.egi.eu/vo/

Nikhef

Nikhef Data Processing Facility – multi-community service

'just one of these sites'

NDPF HTC platform

- member of a federated service with SURFsara, Nikhef, RUG-CIT
- high-throughput storage at SURFsara and Nikhef,
- long-term storage at SURFsara,
- interconnected by SURFnet, and authentication by TCS and IGTF

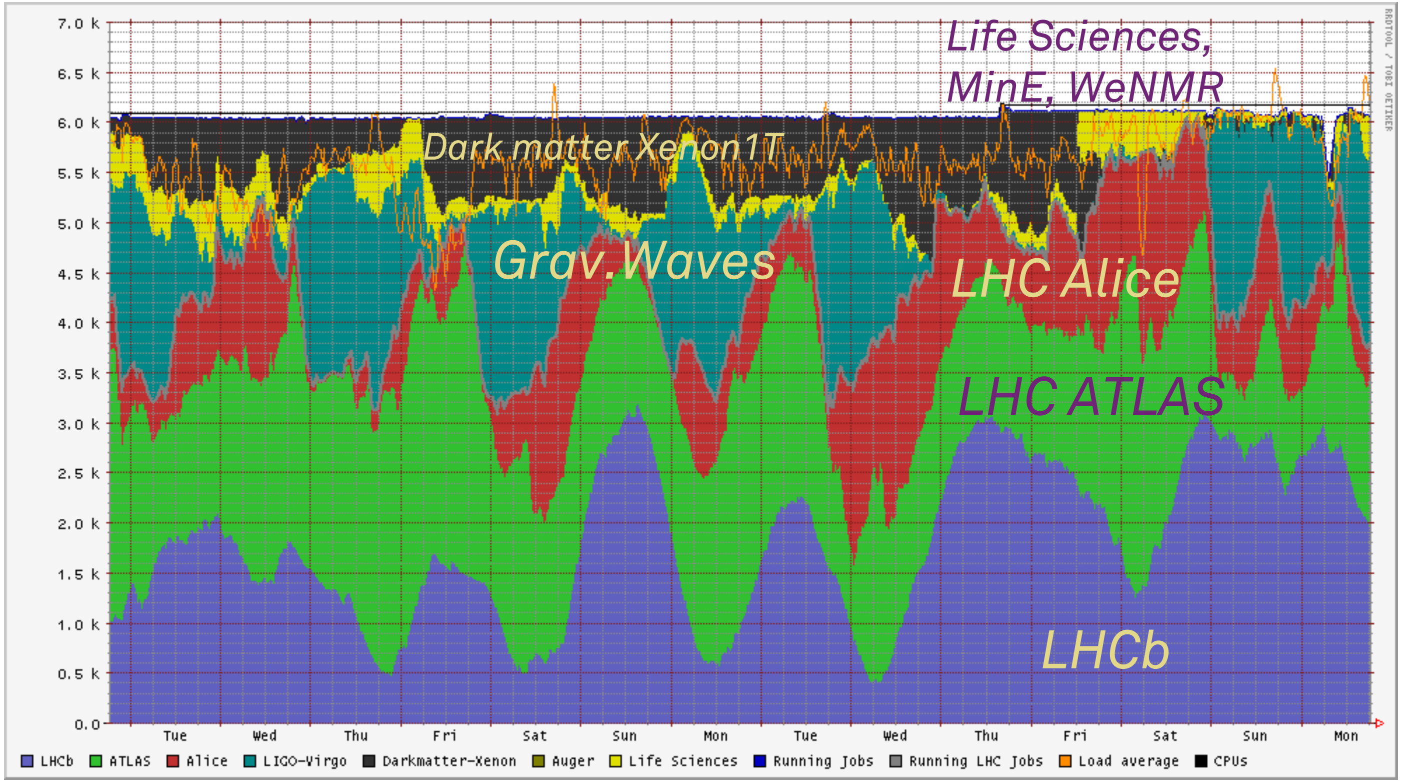

NDPF voview short 1 October 2018

Nikhef

Getting the data to process … in the right place

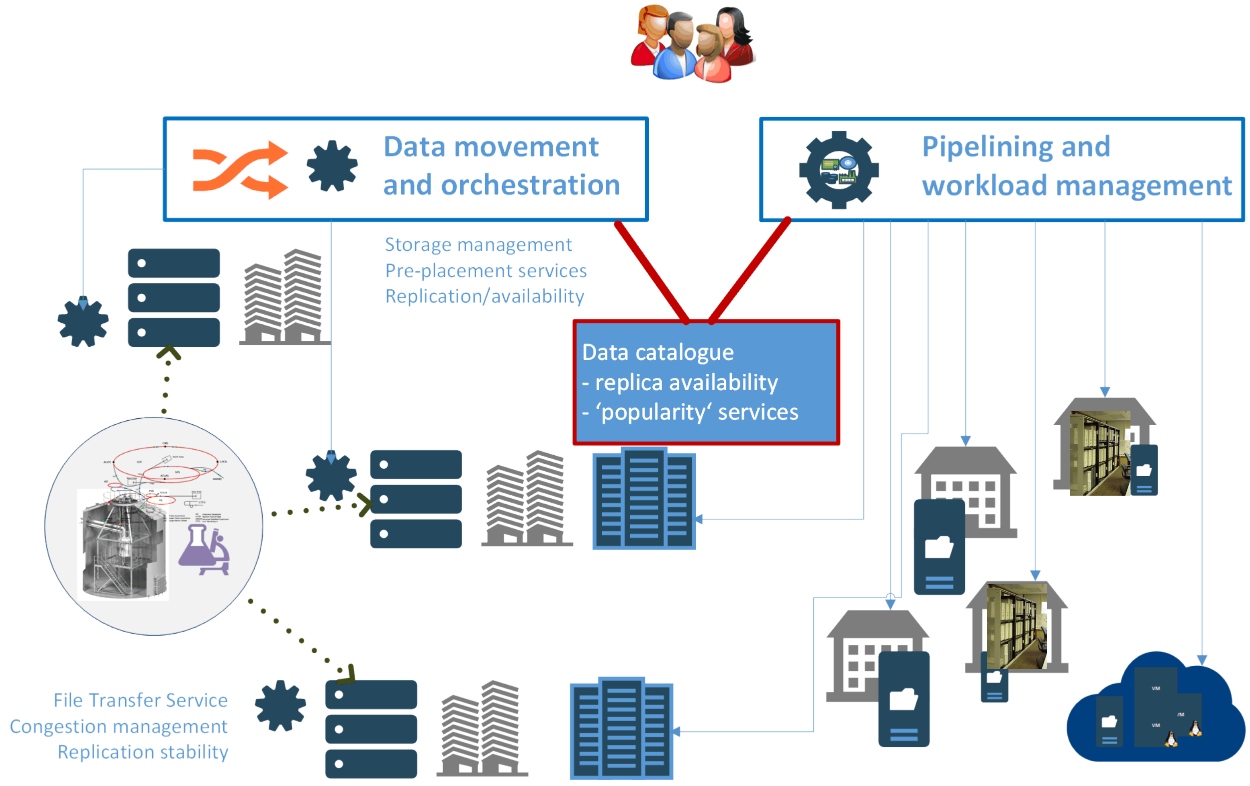

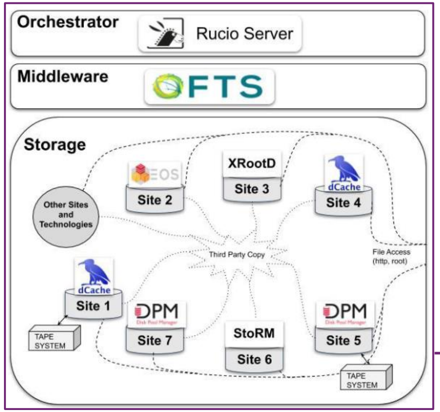

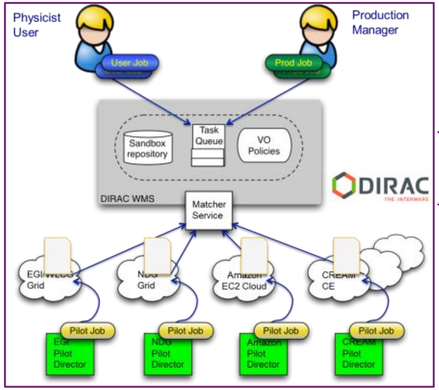

ESCAPE reference components: Rucio, FTS, and DIRAC

Rucio FTS image: “The ESCAPE Datalake” Aris Fkiaras, CERN, February 26th, 2020 ESCAPE WP2 eDIOS https://indico.in2p3.fr/event/20203/sessions/12778/#20200226
DIRAC image from: A. Tsaregorodtsev, CPPM IN2P3, EGI Webinar, 7 June 2016 https://indico.egi.eu/event/2978

Nikhef storage infrastructure

‘making disks collaborate’

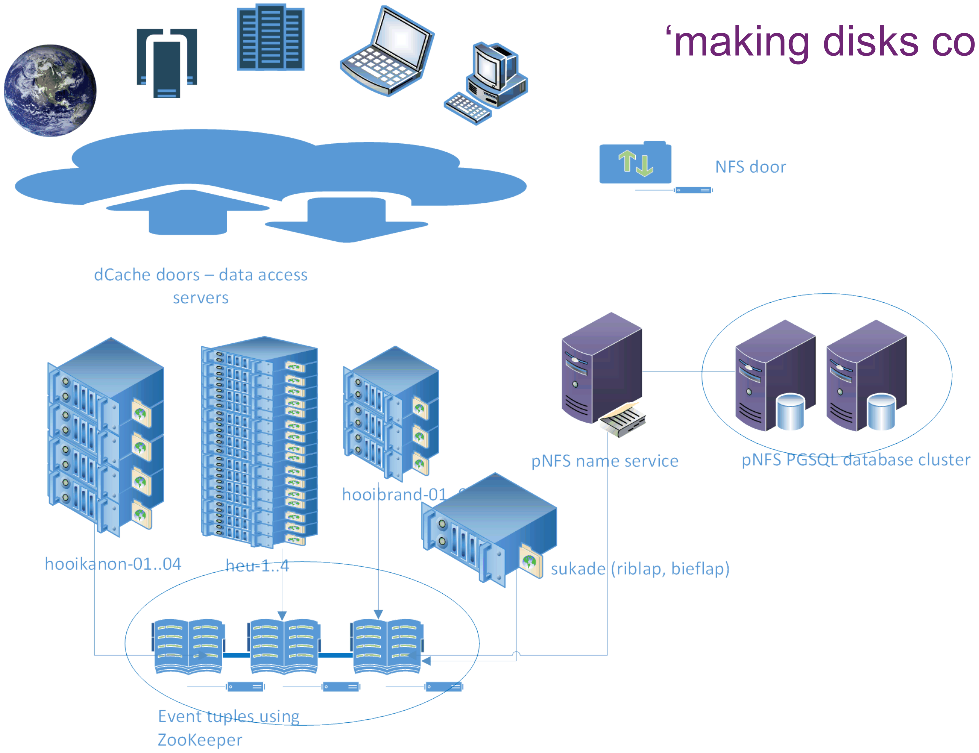

Storage as infrastructure – even a few PBytes requires some organization

- management challenge is **#files**, not capacity
- cost challenge is **throughput**, not capacity

for the management software, see https://dcache.org
alternatives exist, like DPM https://lcgdm.web.cern.ch/dpm

Data distribution for WLCG

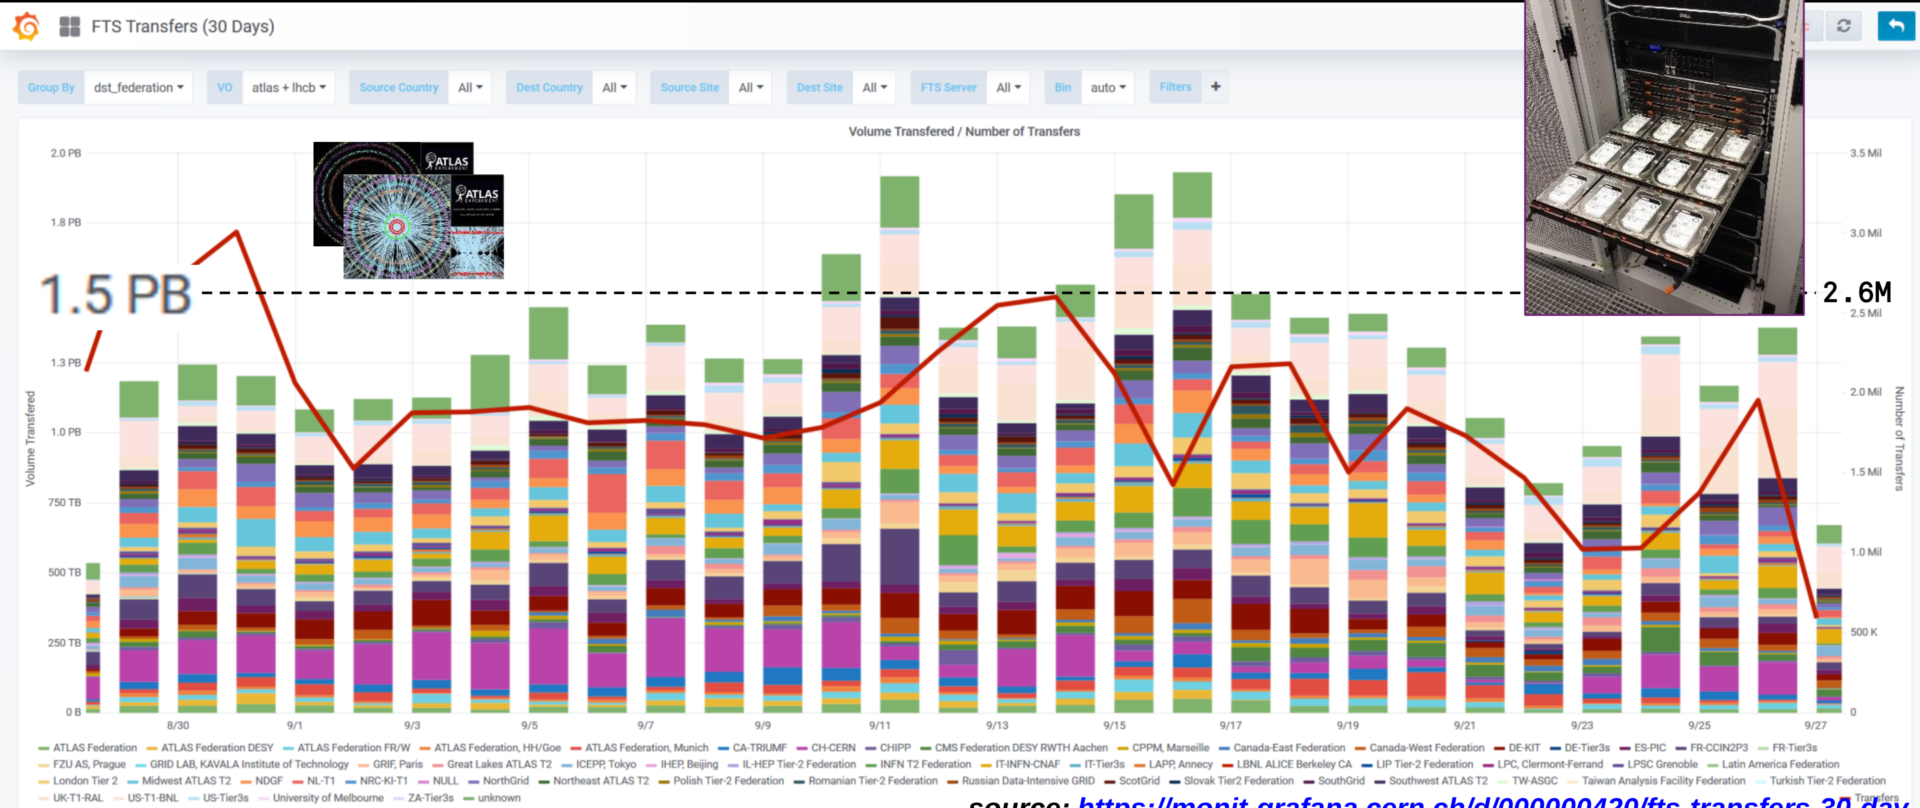

1.5 PB

2.6M

source: https://monit-grafana.cern.ch/d/000000420/fts-transfers-30-day

Structuring of frameworks impacts systems design

pre-staging all data locally supports
latency hiding, posix-style access with lseek(2), ‘$TMPDIR’
e.g. why there are Data Transfer Nodes (DTNs) in the ‘Science DMZ’ concept

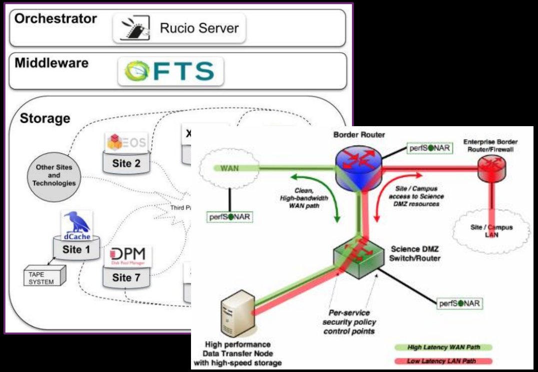

but, recently, pre-staging starts coming at a cost, when using SSDs as local data ‘scratch’ area … because of their unique element: ‘endurance’

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

WORN storage – Write Once Read Never

Frequency distribution observed on the NDPF execution nodes for outside ('grid') access (blue) and local access (orange)

Access pattern is rather different. But why?

- external users pre-stage, because that is built into the frameworks (like DIRAC, Athena), where local users can use streaming access ('dCache NFSv4')
  yet there are changes in pre-stage streaming behaviour over time
- different types of workload:
  ntuple-data analysis vs (re)processing
- …

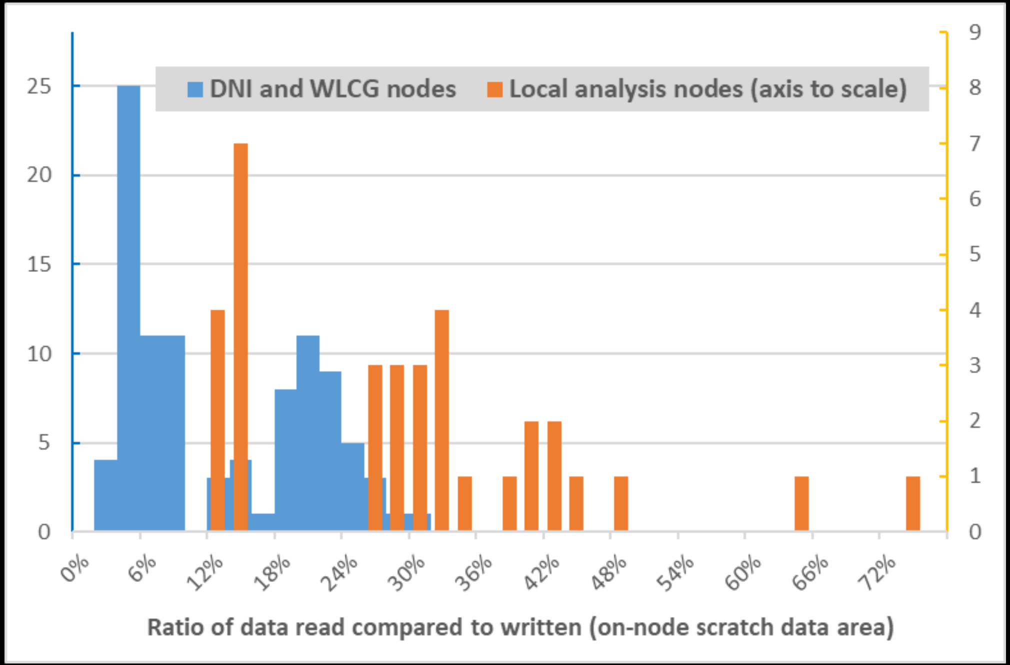

Data: NDPF execution nodes, based on SSD SMART data, integrated over total device lifetime
plot shows number of local analysis nodes scaled to DNI-WLCG count; collected using smartctl on 2020-10-28

Data comes from somewhere, and has to go somewhere …

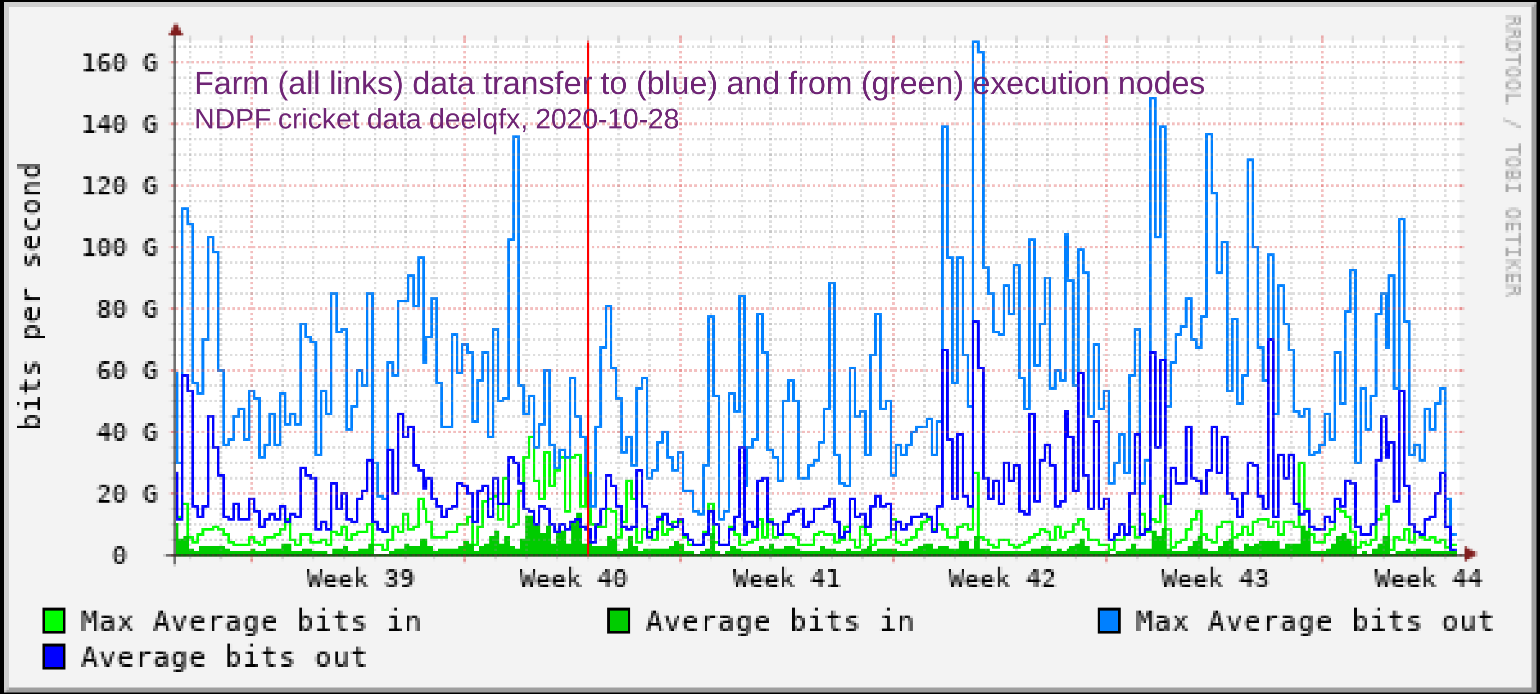

LHC Open Network Environment

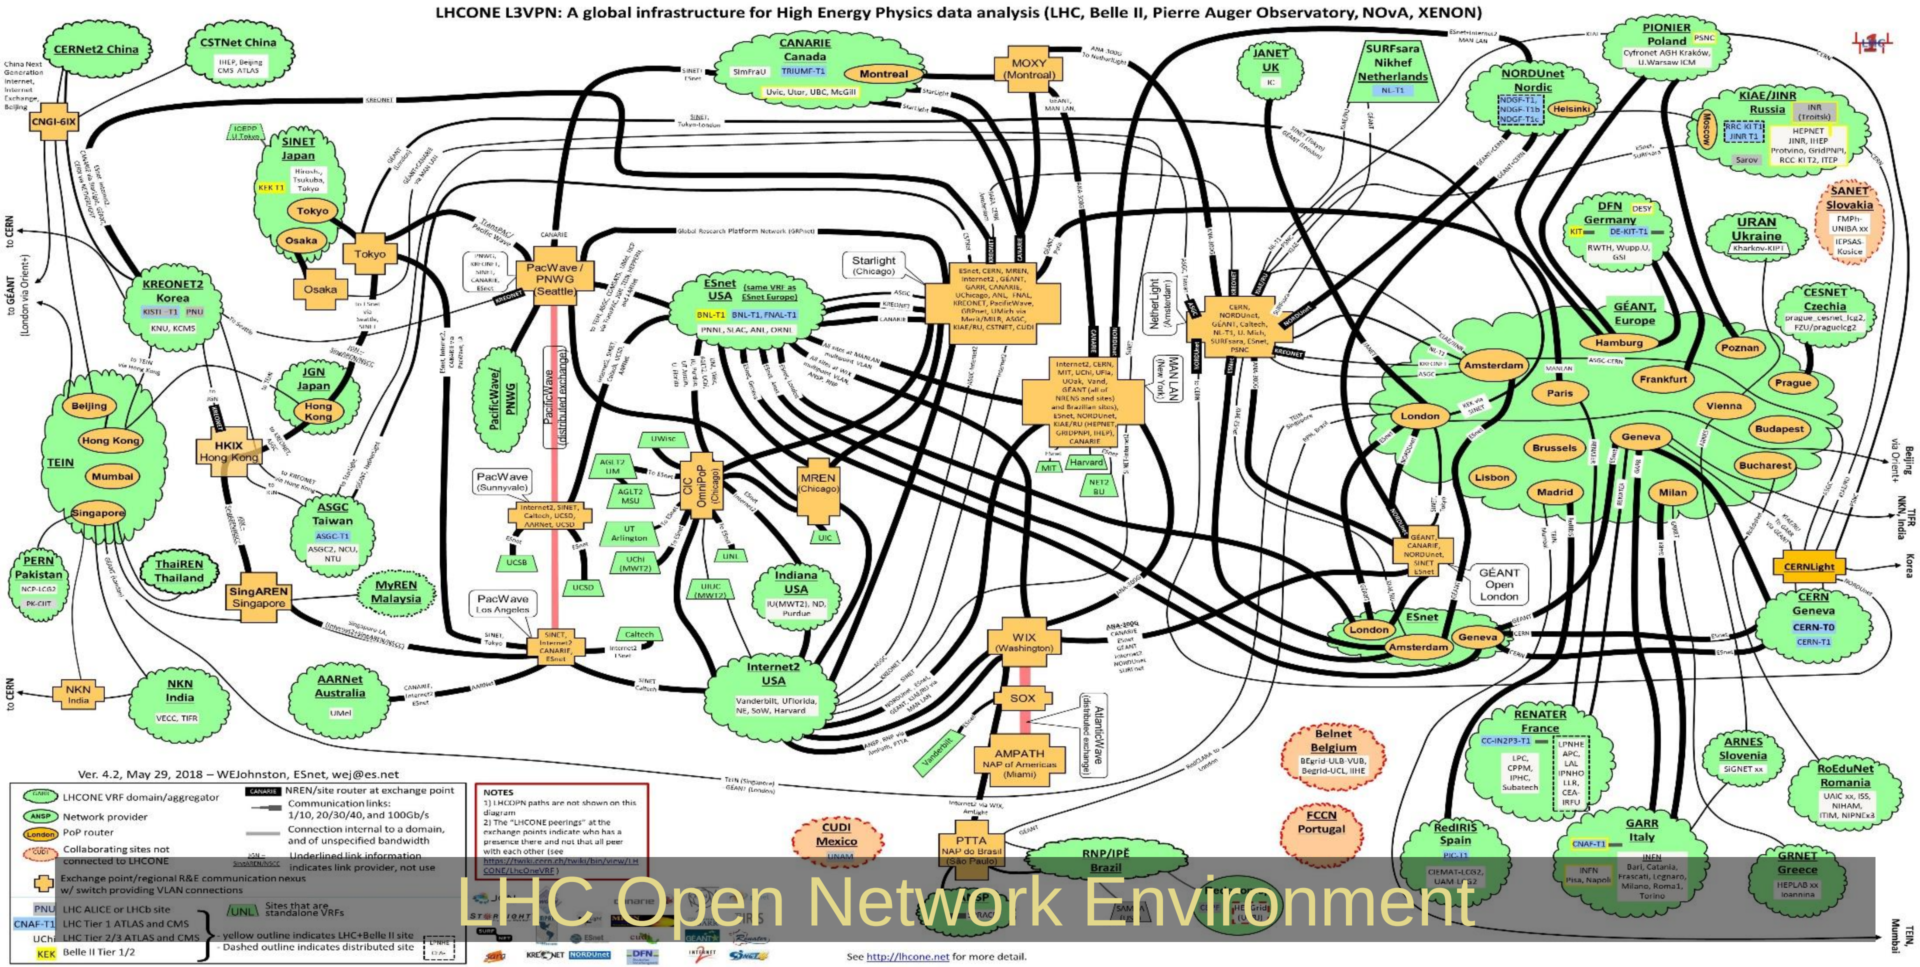

Nikhef

Evolving the physical network view

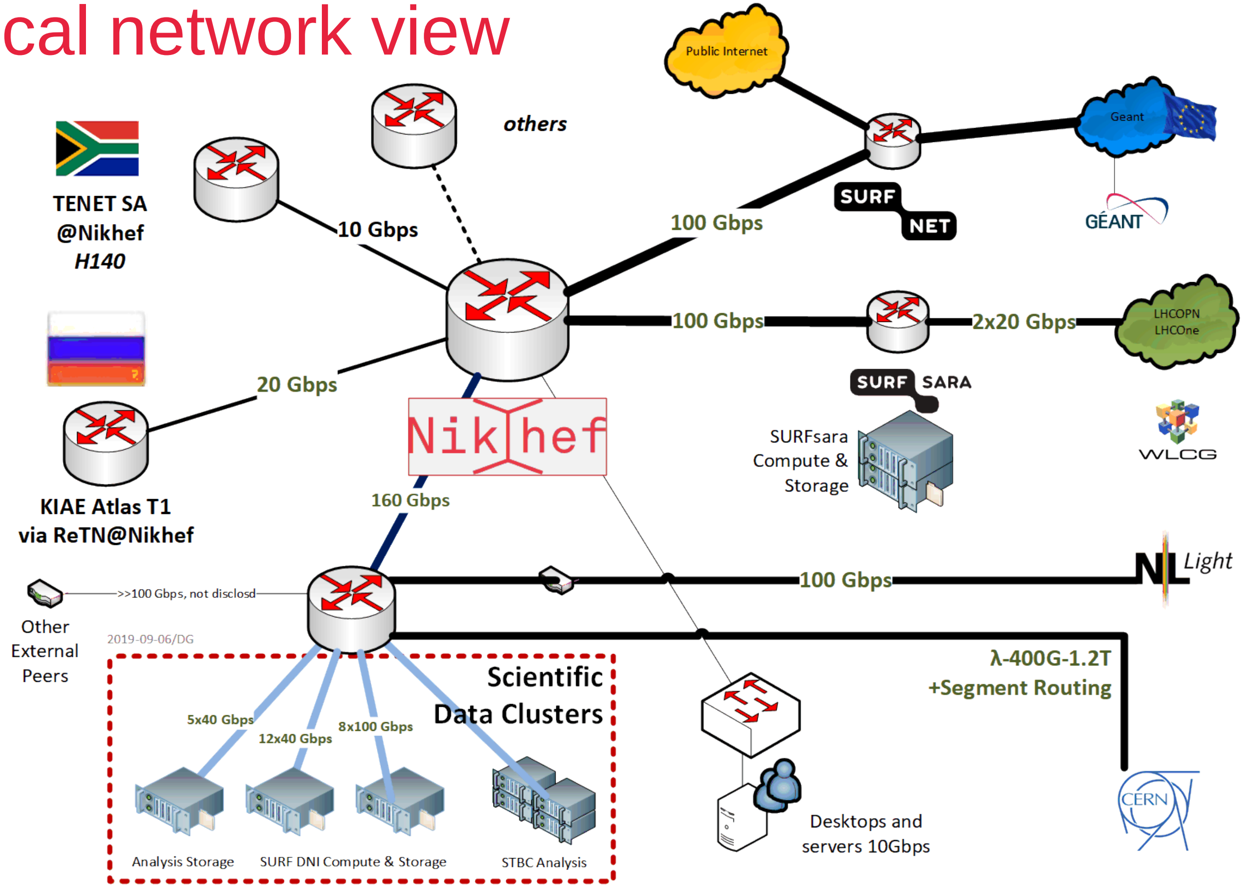

Left: IBR-LAN (1996) in H1.40;
Right: Nikhef peering visualisation (medio 2020)

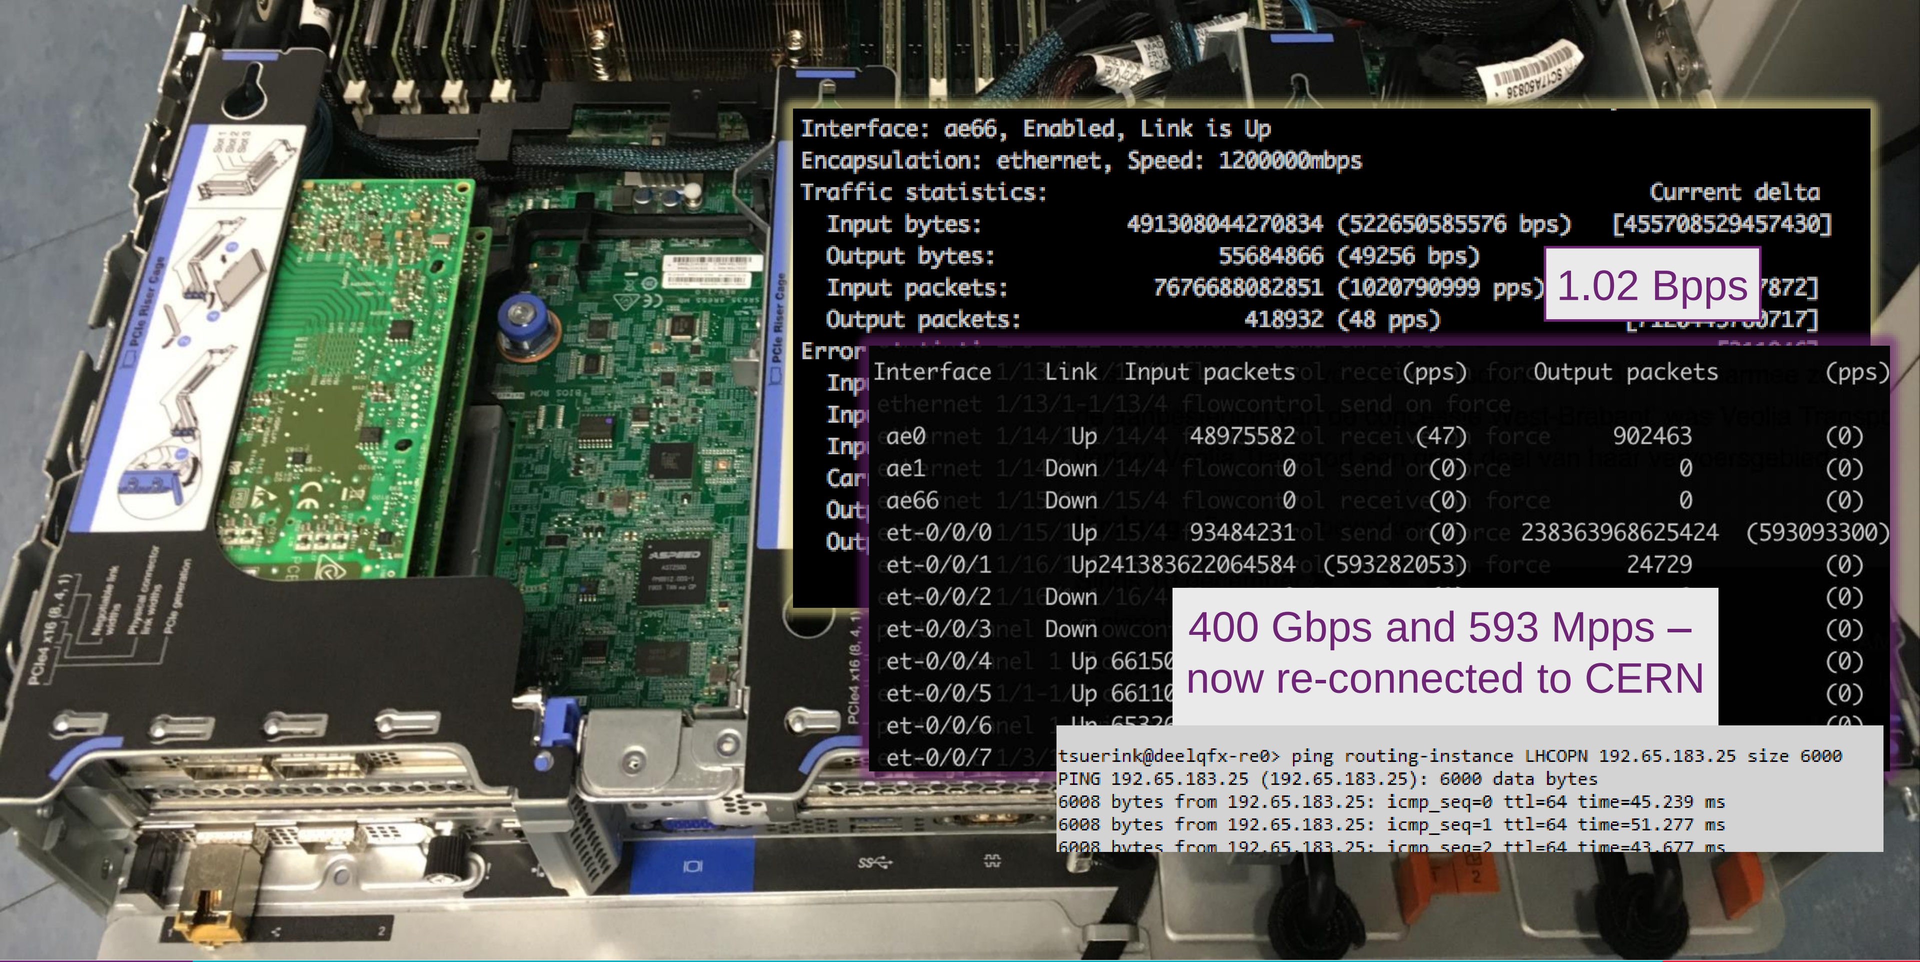

1.02 Bpps

400 Gbps and 593 Mpps – now re-connected to CERN

Image: Tristan Suerink

Our science data looks akin to a DoS

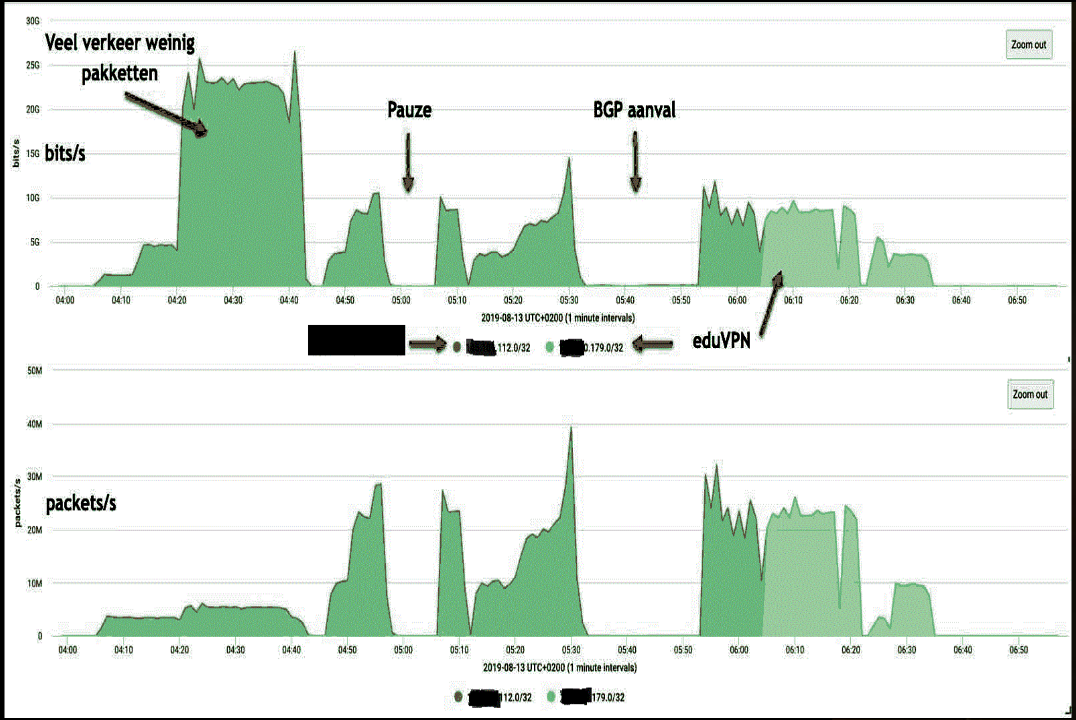

evaluating resilience to cyberattack – *in a cooperative way*

Segmentation: a network of 'private domain' clouds within

open-core research network model implements the enclave structure

protects against overload by *no stateful components in the network path*

and allows open research federated cloud using eVPN overlays

although you'll always have some (reputational) risks even if you advertise the block as 'customer network devices'

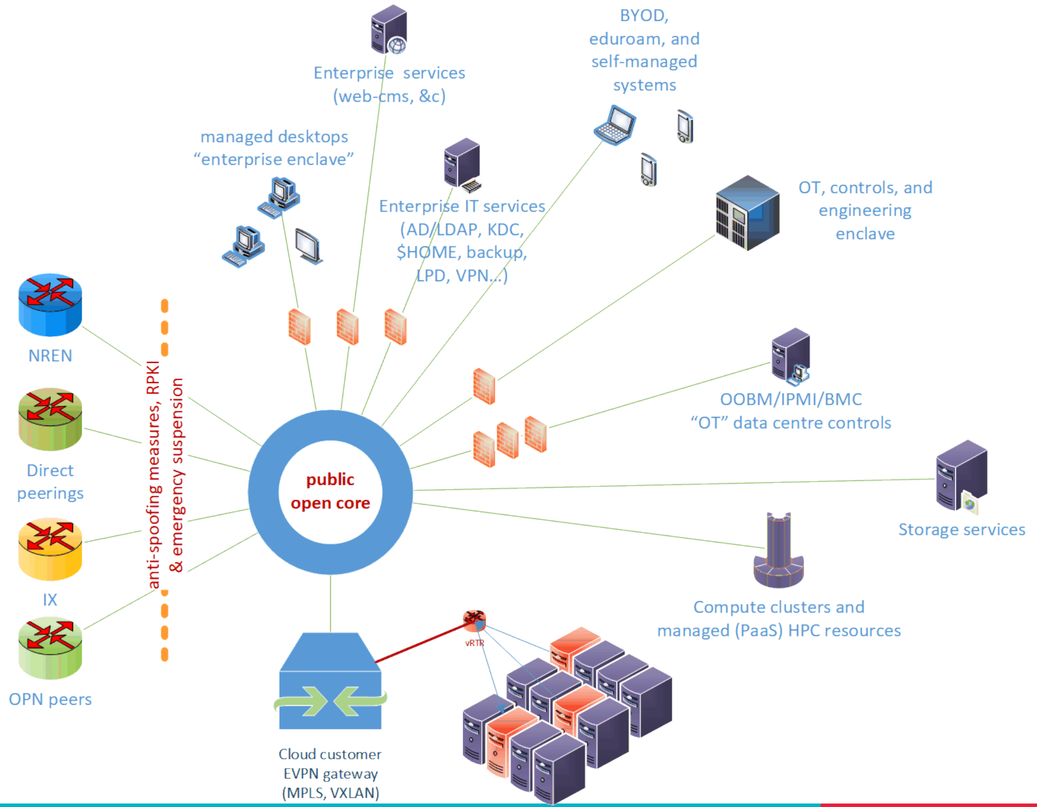

Nikhef

Nikhef cloud – targeting high-throughput use cases

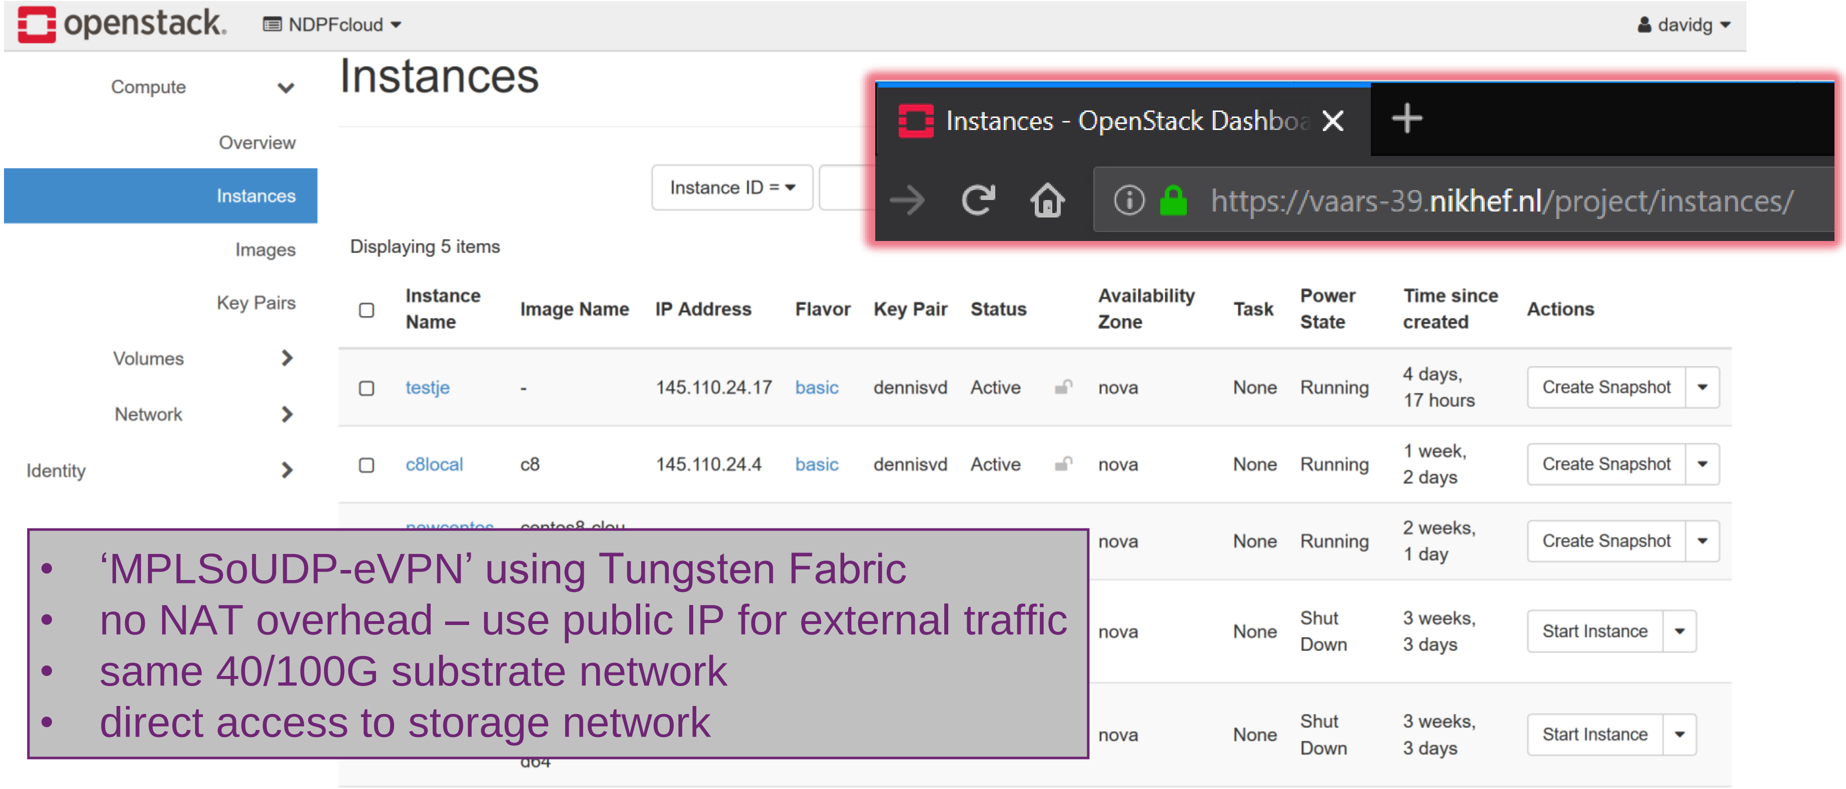

- ‘MPLSoUDP-eVPN’ using Tungsten Fabric
- no NAT overhead – use public IP for external traffic
- same 40/100G substrate network
- direct access to storage network

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

Nobody wants a cloud … you want a solution!

research community overlays and 'virtual clusters'

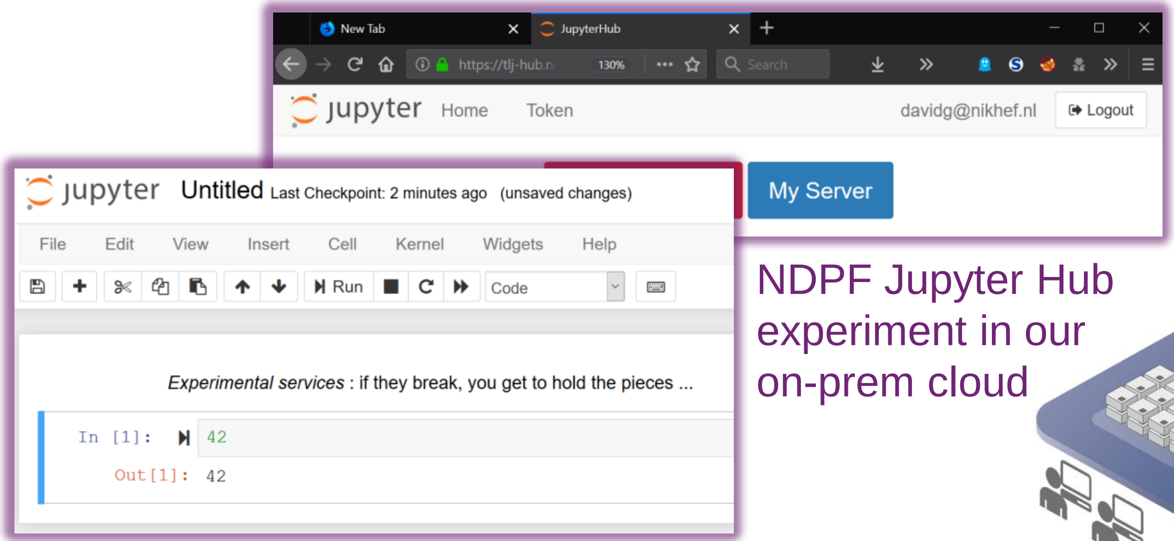

NDPF Jupyter Hub experiment in our on-prem cloud

a federated infrastructure

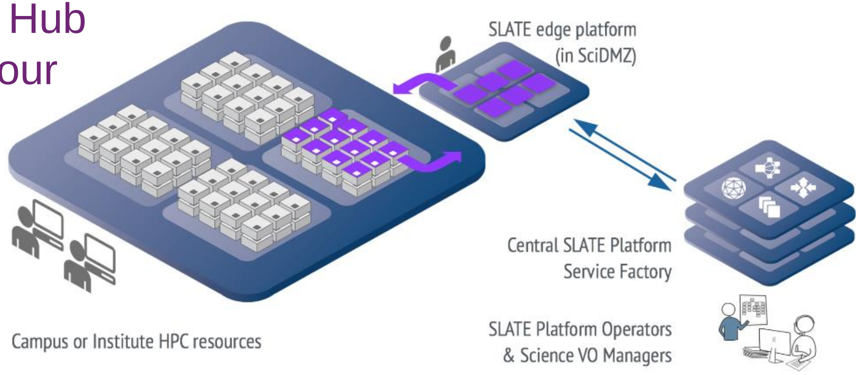

at scale: container computing, yet with curated application images – slateci.io

Cross-organisation infrastructure
we need an 'ecosystem' more than a cloud

on-prem cloud, or research cloud, is oft better

- very cost-efficient if utilised at capacity
- effective as it can provide more than 'IaaS'
- can leverage our own R&E federated access

and not all 'cloud' is what you think it is …

EOSC – the European Open Science Cloud

more an ecosystem (or 'web of data') than a 'cloud'

Photo by Pop & Zebra on Unsplash

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

An ecosystem built on federated infrastructures

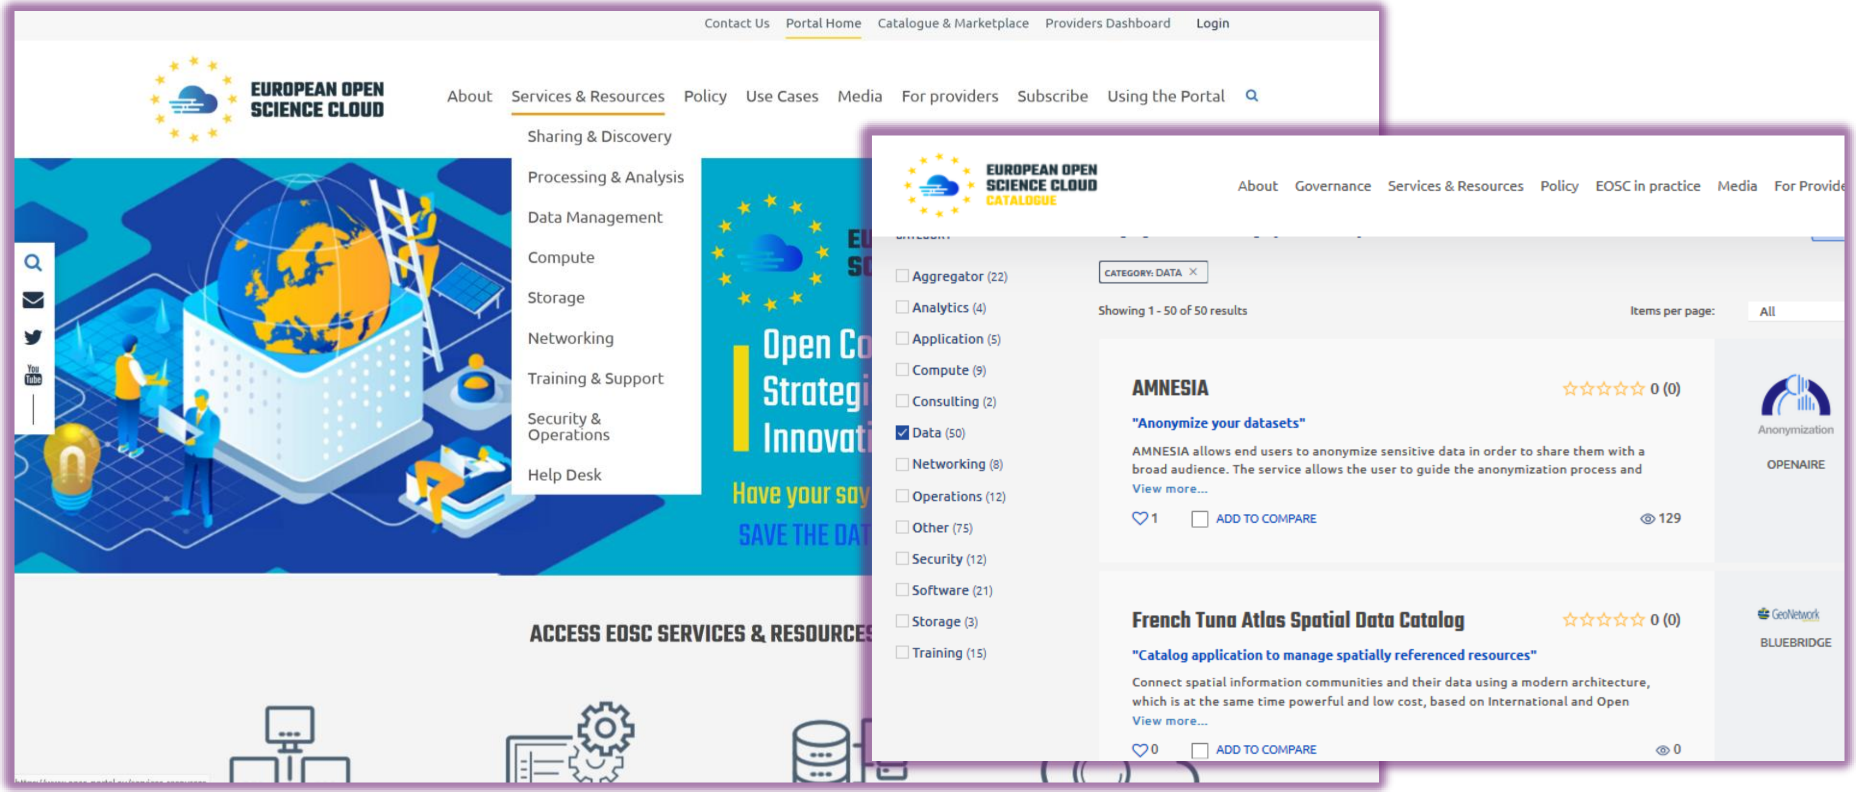

EOSC Portal (https://www.eosc-portal.eu/) – as built by EOSChub

Whence we came: the long road to federated access

From disparate systems in ~2000

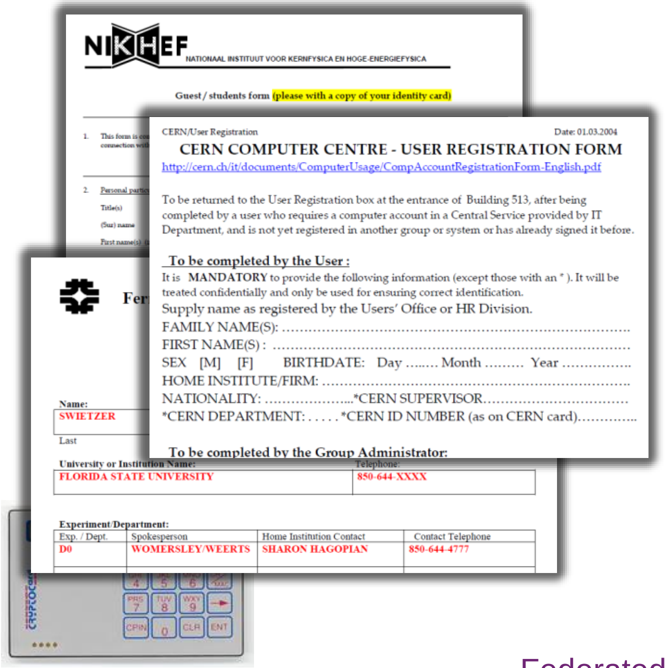

separated authentication and authorisation, splitting *identity sources, community membership,* and *services*

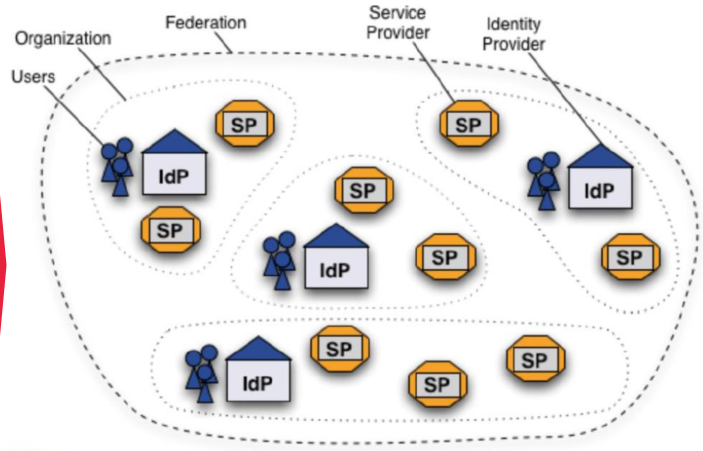

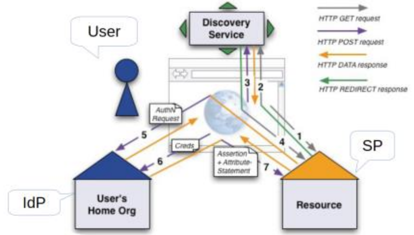

SAML2.0 auth flow

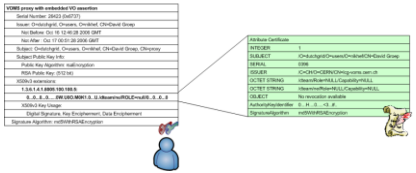

Federated (R&E) AAI, the global IGTF PKI, VOMS, and
'AARC BPA' AAI architecture all have this as fundamental property

VO Management System
VOMS attributes certificates

Shibboleth IdP image: SWITCH (CH)

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

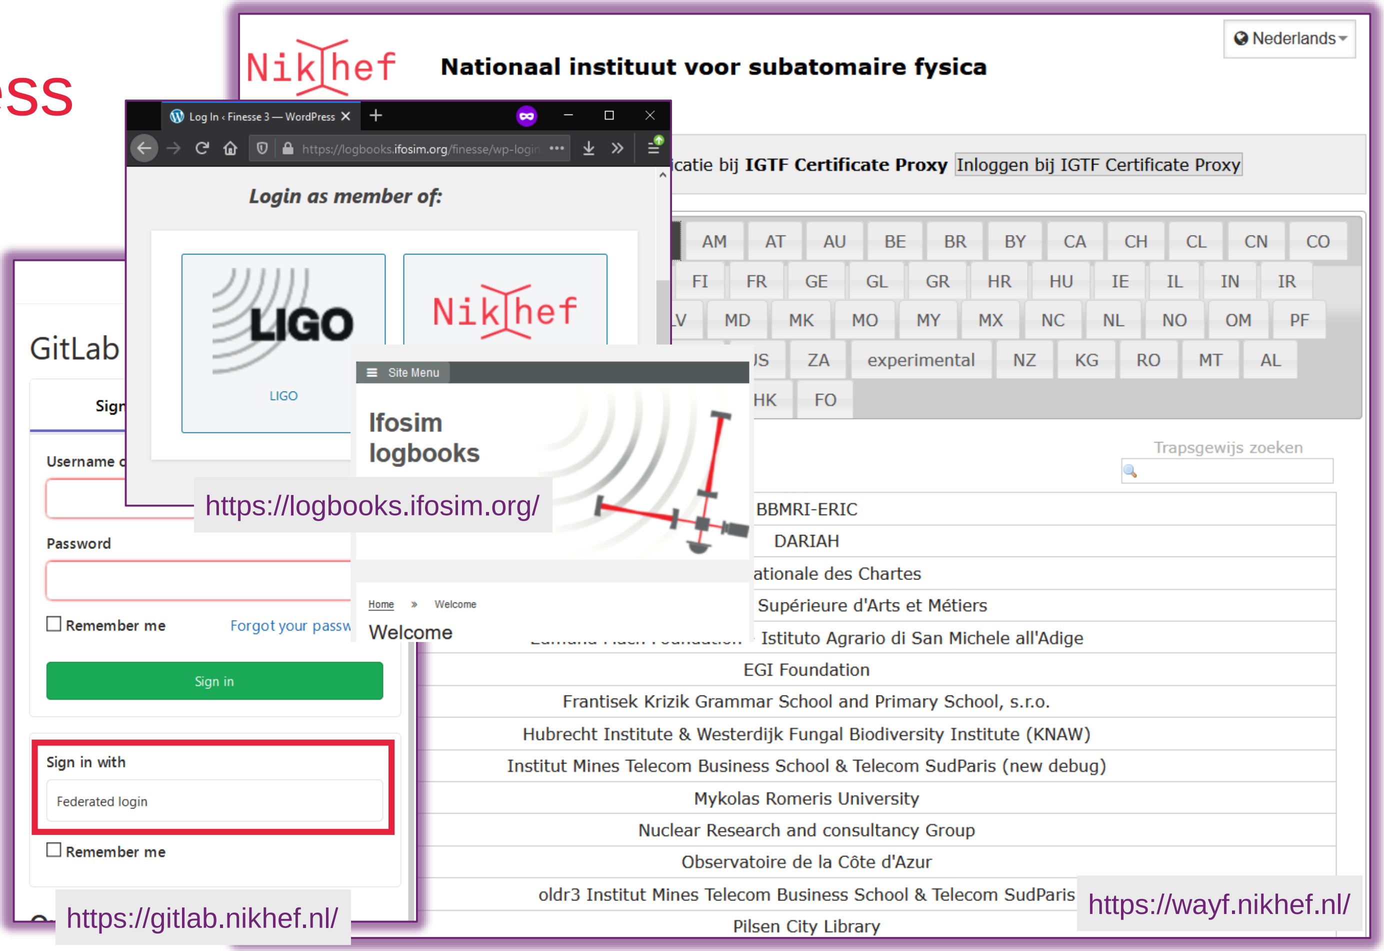

eduGAIN

ifosim federated AAI integration implementation
performed by Mischa Sallé

Challenges of scaling federation

Beyond 'enterprise' services, it becomes challenging!

Collaborations - by design - have their services distributed *and*

- not that many collaborations are a legal entity
- or not 'legally authoritative' for constituent services
- or run into risk-averse, or slow, 'home organisations'

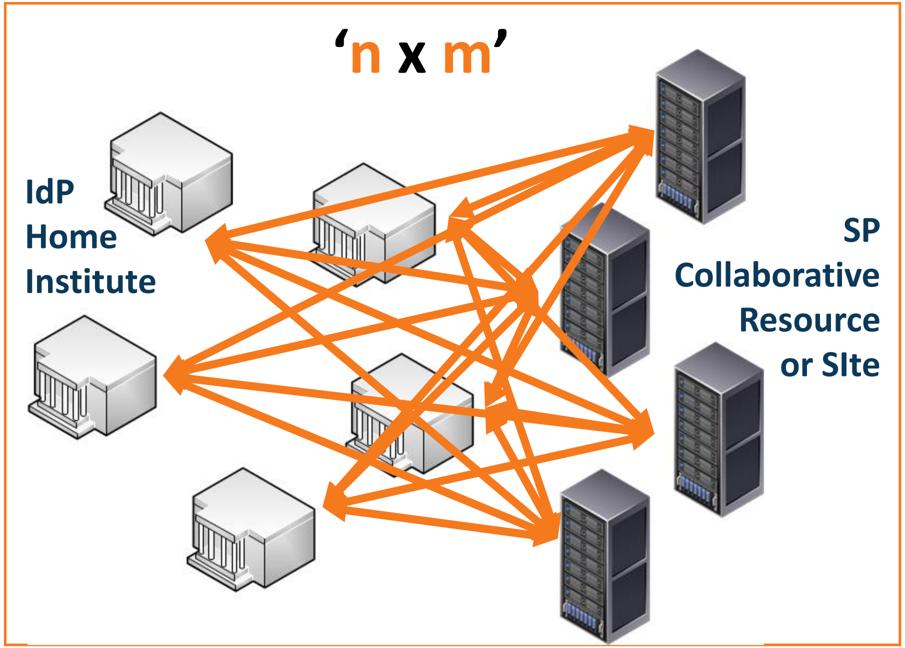

Scaling community and institutional trust

eduGAIN (global R&E) *Entity Categories*

Curated grouping of entities
‘REFEDS R&S’
this is a research service
‘DP CoCo’
abides by GDPR
‘Sirtfi’
cares for security response

slower adoption process
adding identity assurance needs
action at all 60+ Feds & 4k+ IdPs

e-Infrastructure IGTF Authentication Profiles

Common baseline and profiles
co-defined by relying parties

user-centric ID harmonisation with unique global naming
‘BIRCH’
real person with real name
‘DOGWOOD’
persistent linkable identifier

research-specific user base

Use of proxy bridging components

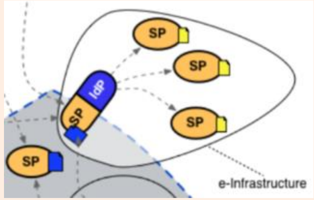

Identity and access ‘proxy’ harmonised eduGAIN IdPs

based on entity categories
leverage Sirtfi and ‘R&S’
proxying is bi-directional

responsibility on the proxy operator

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

Research-friendly federation: REFEDS R&S … or SRAM

REFEDS Spaces Search Log in

Entity-Categories

Pages

Blog

CHILD PAGES

Entity-Categories Home

Research and Scholarship

Research and Scholarship FAQ

Pages / Entity-Categories Home

Research and Scholarship

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

For IdP Operators

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

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

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

Category support is defined as follows:

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

‘a science collaboration zone’

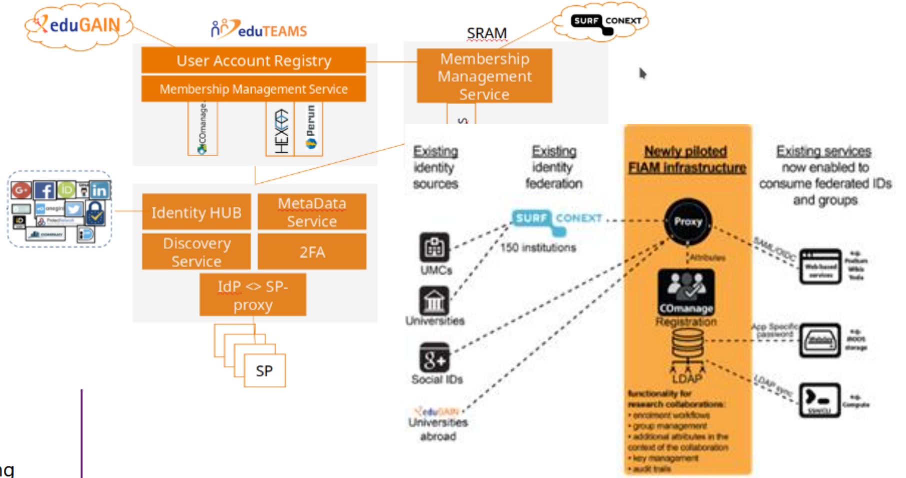

https://wiki.surfnet.nl/display/SRAM/

https://refeds.org/SIRTFI
https://refeds.org/assurance
https://refeds.org/category/research-and-scholarship

Bridges and Token Translation Services

GEANT Trusted Certificate Service

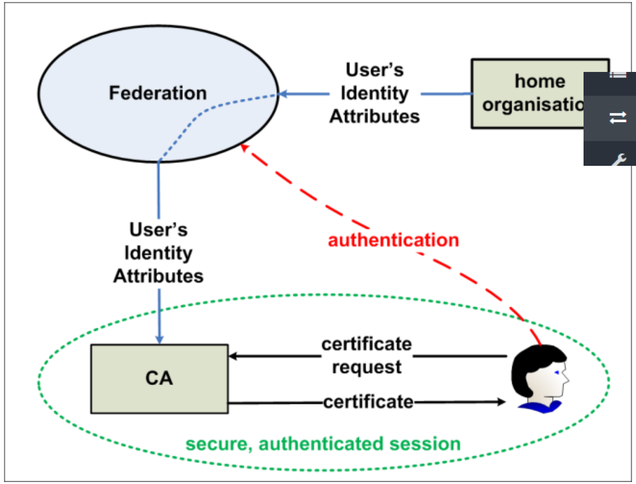

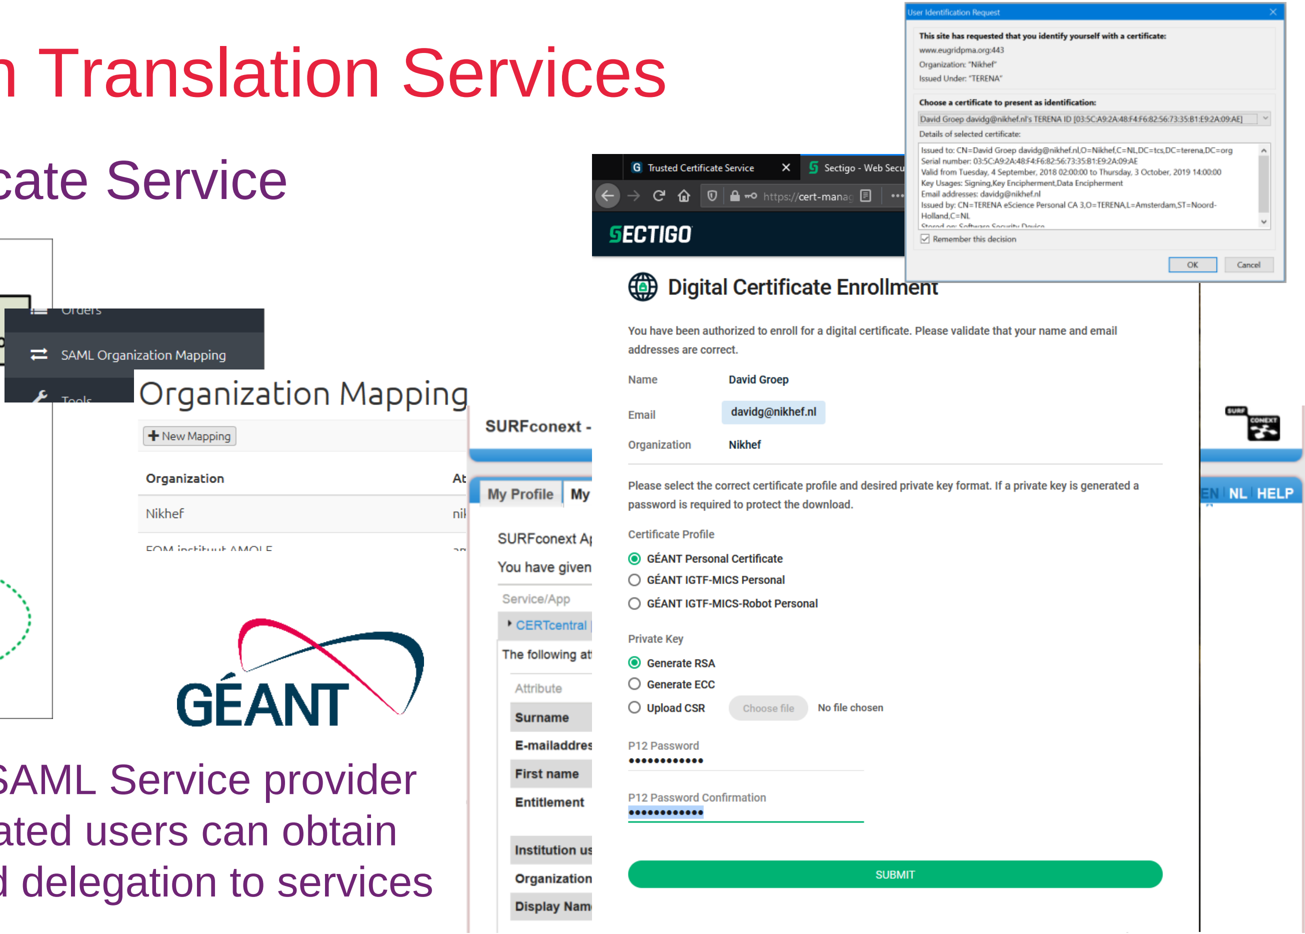

TCS (today: Sectigo) acts as SAML Service provider to eduGAIN: eligible authenticated users can obtain client certificate for access and delegation to services

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

Interoperable Global Trust Federation IGTF

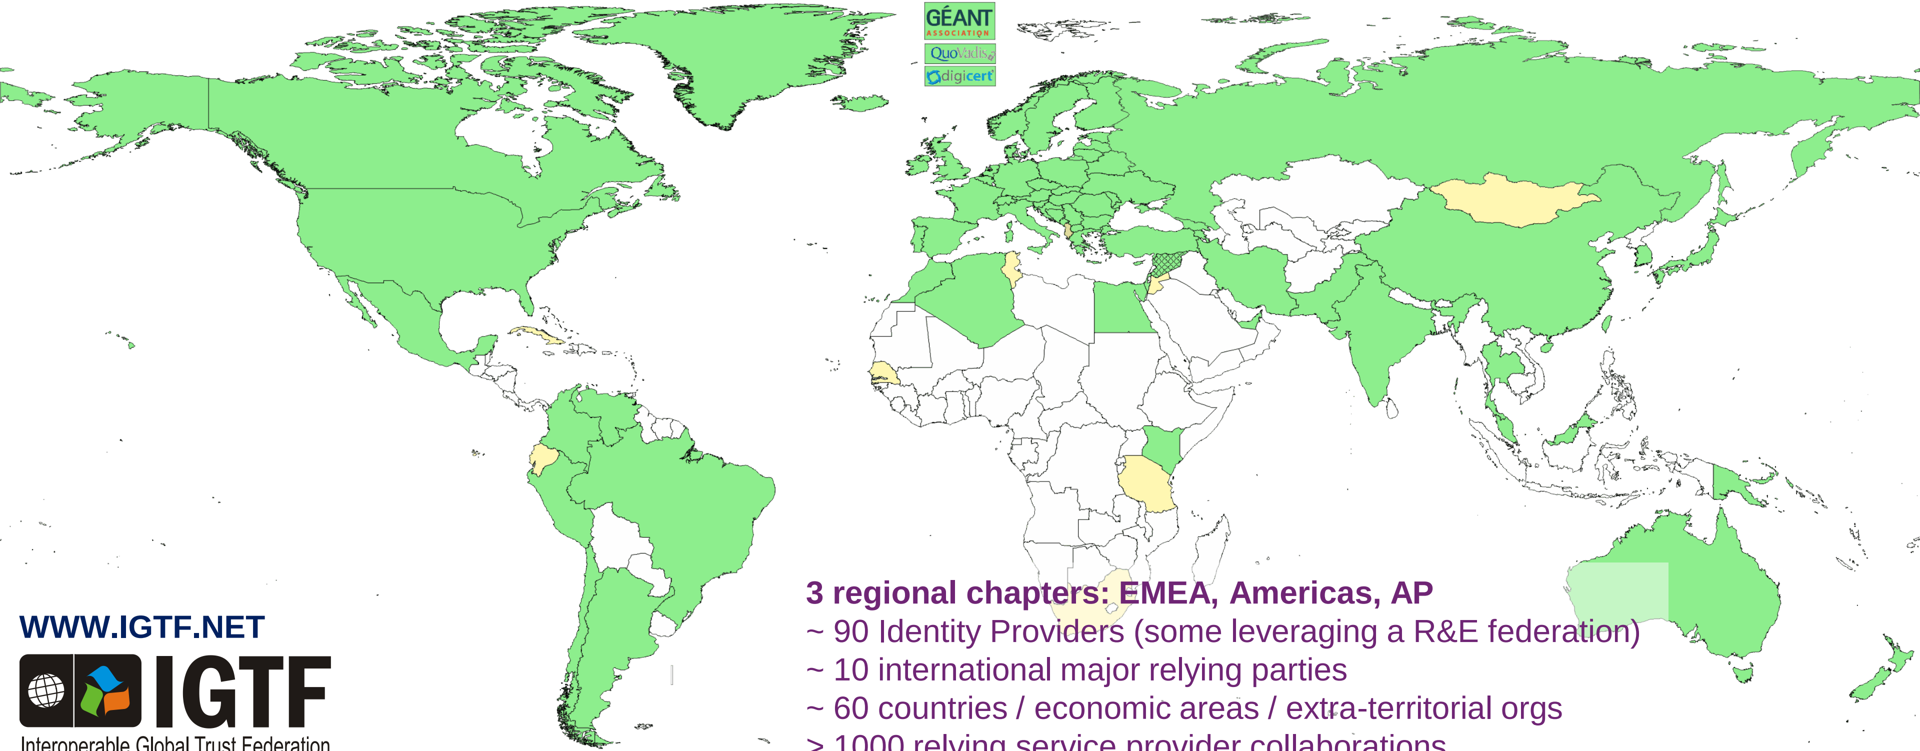

3 regional chapters: EMEA, Americas, AP
~ 90 Identity Providers (some leveraging a R&E federation)
~ 10 international major relying parties
~ 60 countries / economic areas / extra-territorial orgs
> 1000 relying service provider collaborations

WWW.IGTF.NET

IGTF
Interoperable Global Trust Federation
AP | EU | TAG

Managing complexities of distributed identity sources

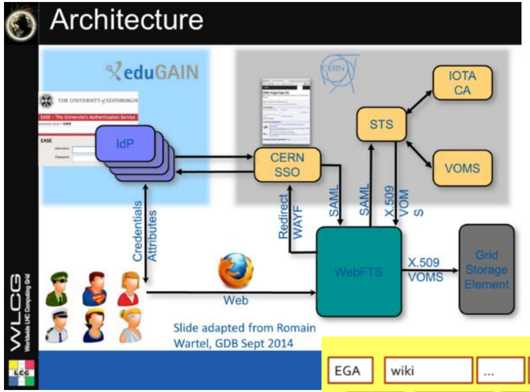

WebFTS prototype
'FIM4R' in wLCG
Romain Wartel et al.

ELIXIR reference architecture
Mikael Linden et al.

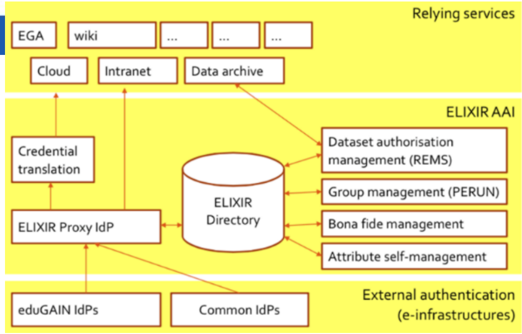

communities had either invented
their own 'proxy' model to abstract complexity

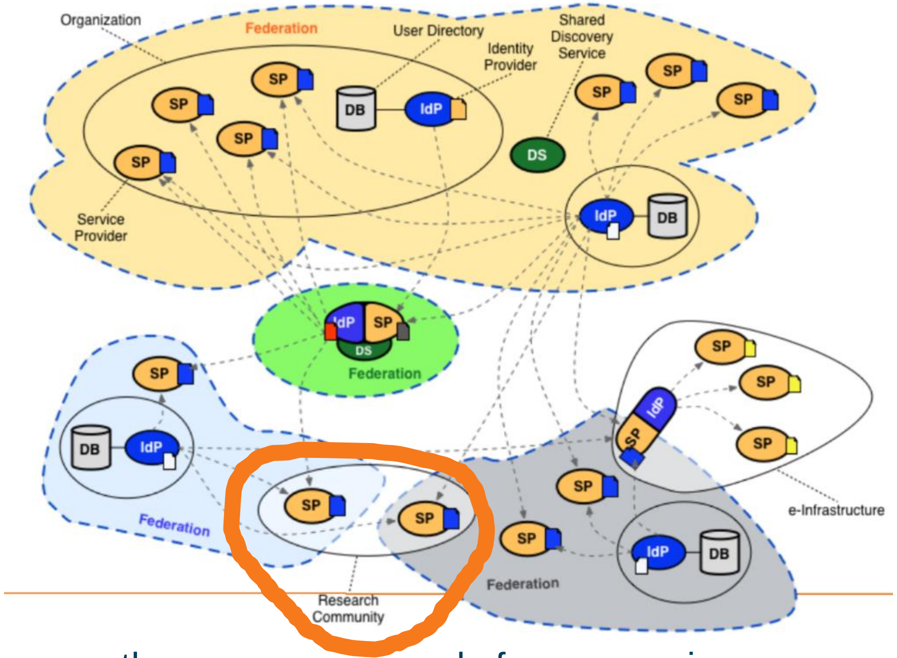

or they were composed of many services
each of which had to manage federation complexity

Community images: Romain Wartel, CERN; Mikael Linden, CSC; Lukas Hammerle, SWITCH

Nikhef

'Community First' AARC Blueprint Architecture: the Proxy

… user and group ID *same* across services
… minimize discovery 'wayf' & info screens

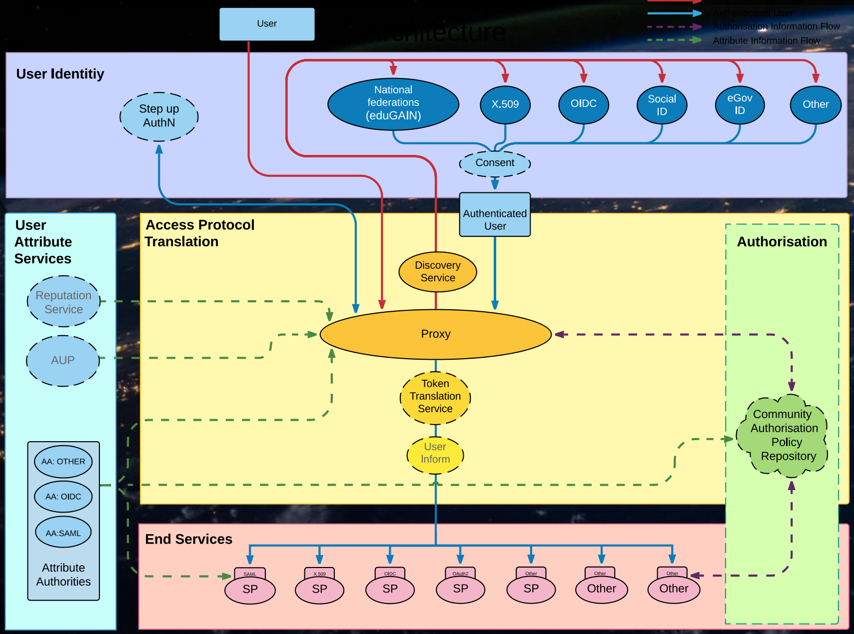

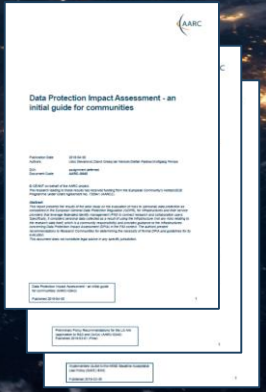

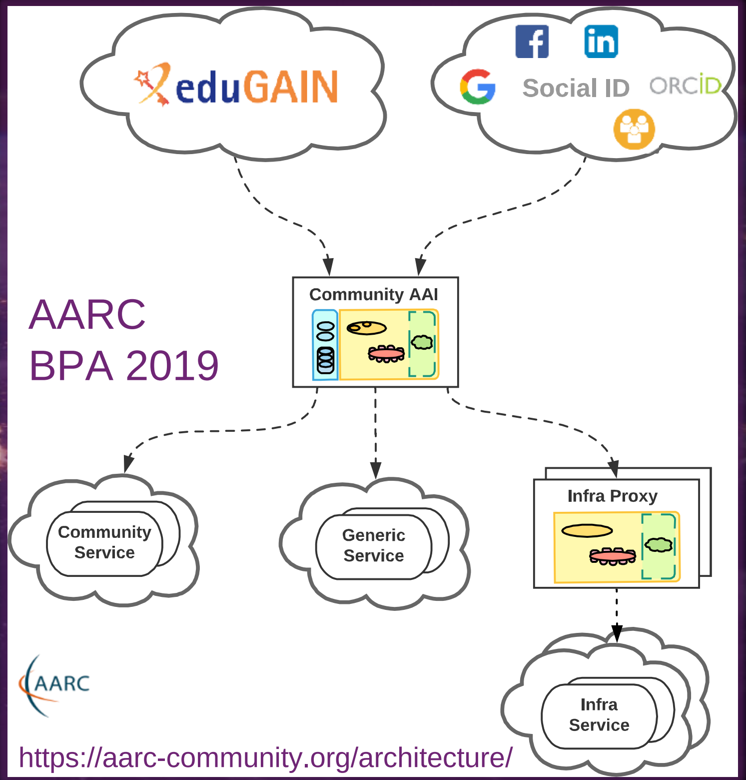

Interconnecting communities and infrastructures

GÉANT

elixir

LSC LIGO Scientific Collaboration

NIH National Institutes of Health
Turning Discovery Into Health

EUDAT

egi

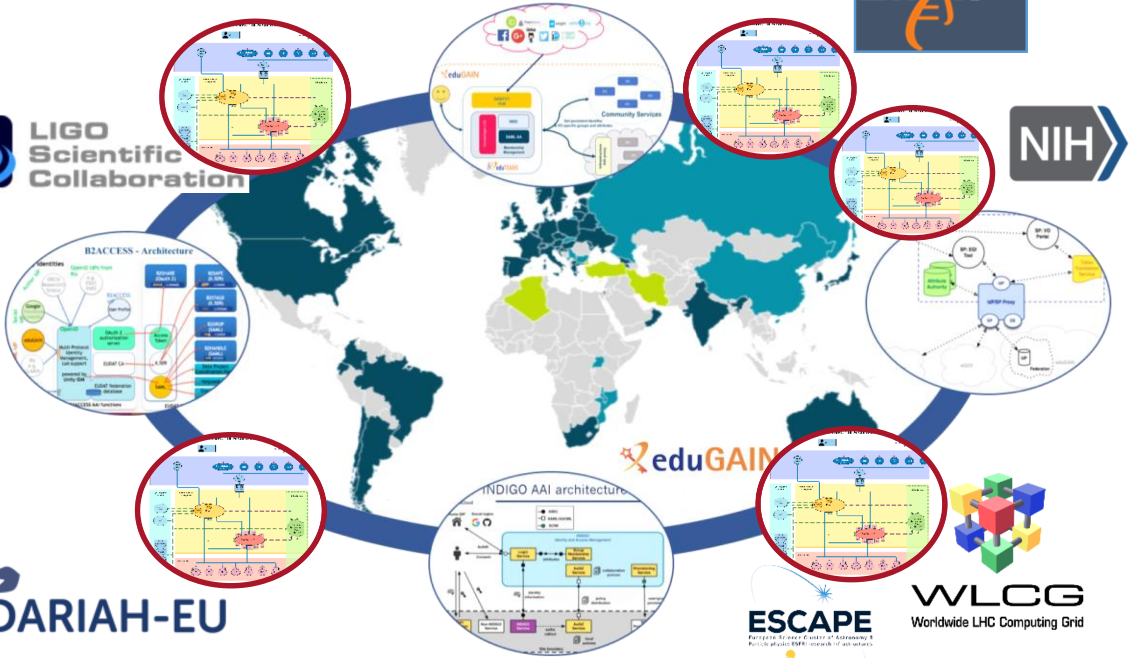

eduGAIN

DARIAH-EU

ESCAPE

WLCG
Worldwide LHC Computing Grid

INDIGO - DataCloud
Better Software for Better Science

https://aarc-community.org/about/aegis/

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*

turn EOSC entities into a federation itself, linked to eduGAIN, preventing 'user loops' & meeting common (security) baseline

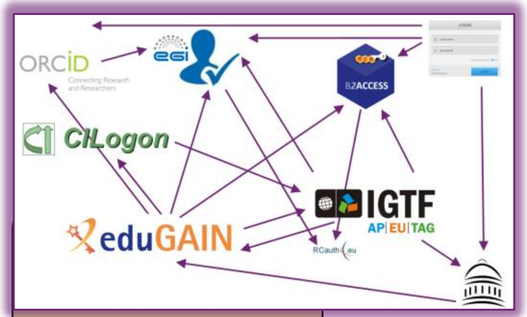

EOSC AAI Federation

National Federation A

...

National Federation Y

eduGAIN

Now *what* have we built?!

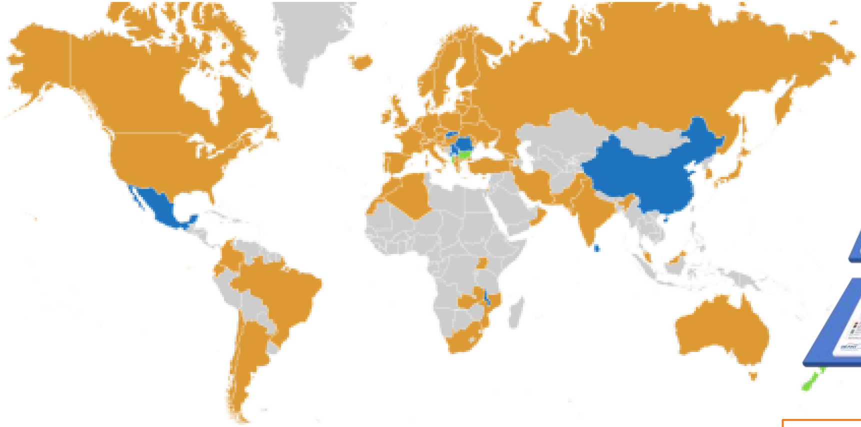

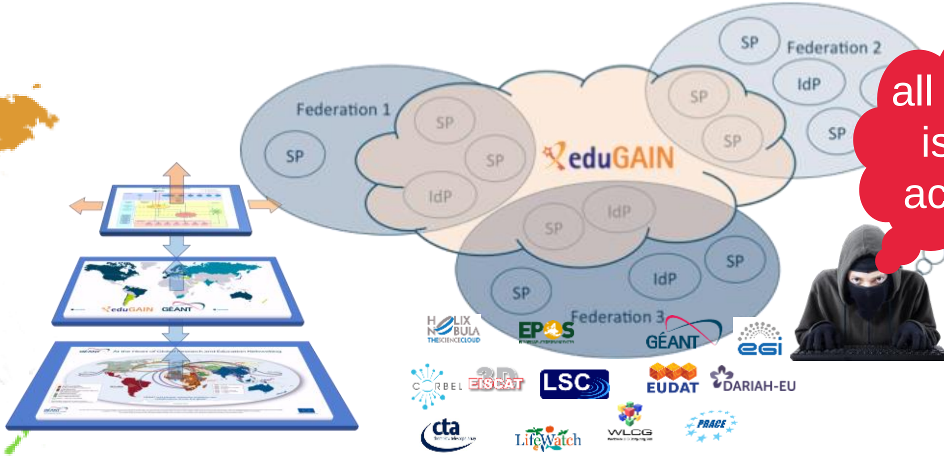

full of valuable resources
(data, network, services)

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
Logos right: e-Infrastructures and ESFRIs
Center graphic: AARC collaboration

Assessing risk … in a collaborative infra

InfoSec **risk assessment framework**
for (EOSC) services based on WISE SCI


National Institute for Subatomic Physics

Nikhef

CONFIDENTIAL

Nikhef NDPF IT standard technical and organisational measures

Introduction and scope

The Nikhef NDPF IT standard technical and organisational measures (TOM) for IT services apply to those services and systems ("in-scope Services") that are part of the Nikhef Data Processing Facility (NDPF), designated to provide services to the Target Audience, and are not otherwise covered by more specific TOM documents.

The Target Audience for in-scope services is defined to be Nikhef employees and registered guests, authorised visitors, users of the NDPF Services in the scope of the Dutch National e-Infrastructure coordinated by SURF, and partners in federated infrastructures in which Nikhef participates and that have been designated as an in-scope target audience by explicit designation or by default.

In-scope Services are those services that are managed and operated by or provided under the authority of the Computer Technology (CT) group and/or the Physics Data Processing (PDP) programme that are in addition intended to be used by the Target Audience ("managed services").

Contents

e.g. ISO27001 can help structure
or identify gaps in your knowledge,
but ISO27002 should not be blindly applied
without *considering the federated interactions*

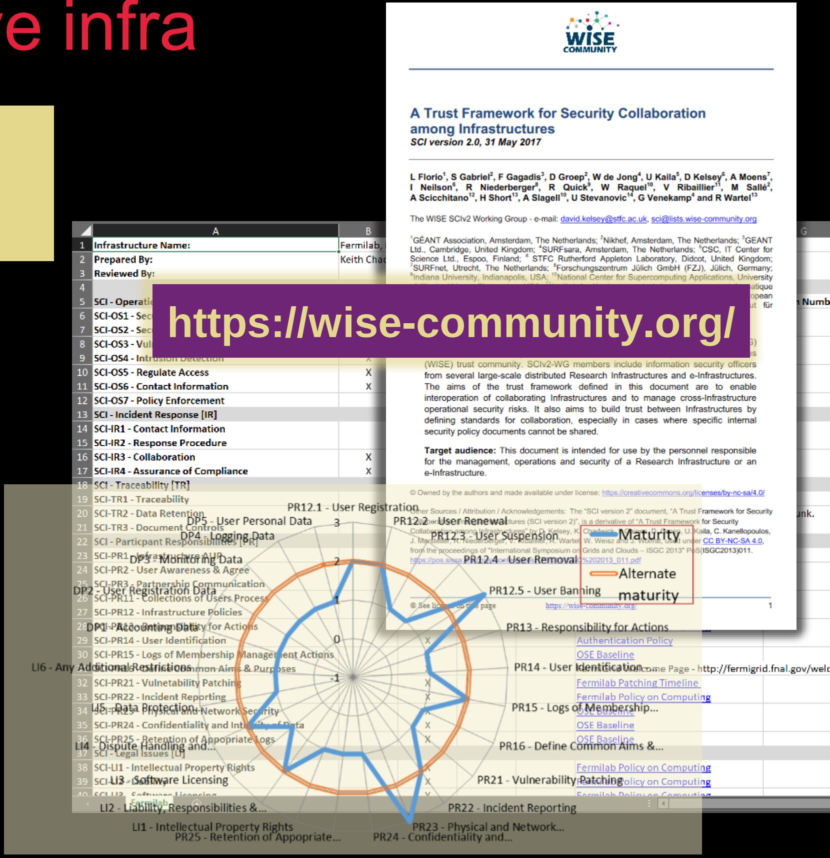

https://wise-community.org/

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

Containment & segmentation

matching the 'open core' research network
community data & systems 1st class citizens

BC/DR
Haarlem

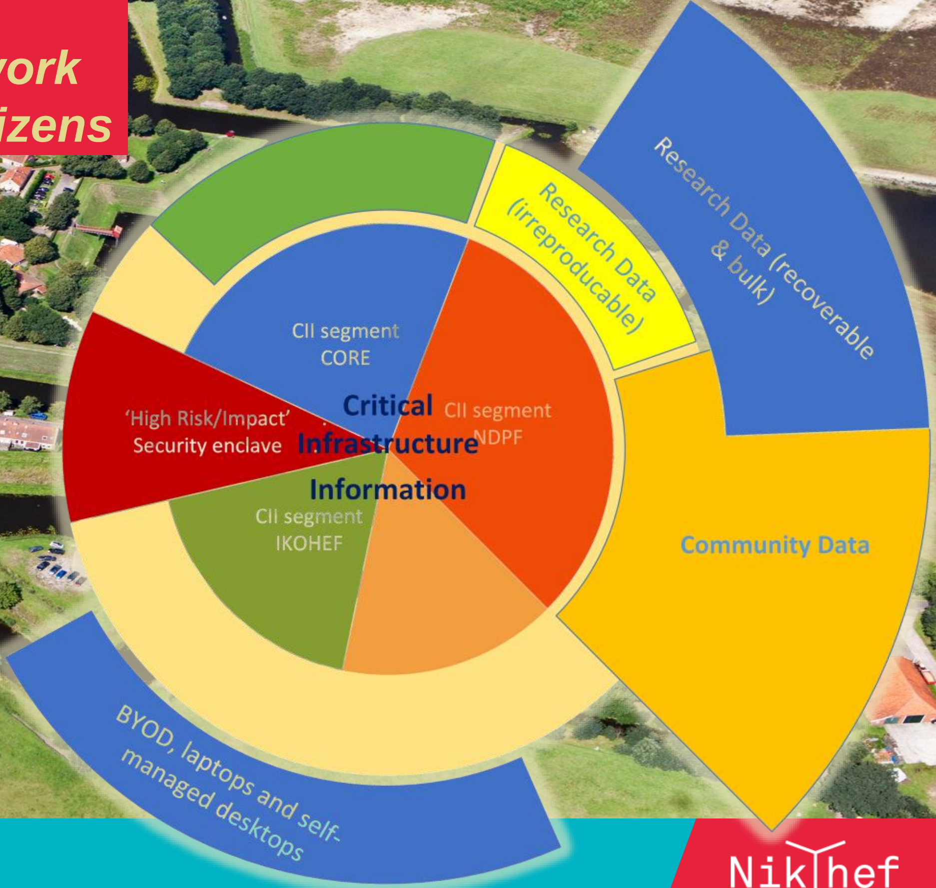

impression Nikhef network-level segmentation

beeld: stichting vesting Bourtange

A question of *when*, not *if*

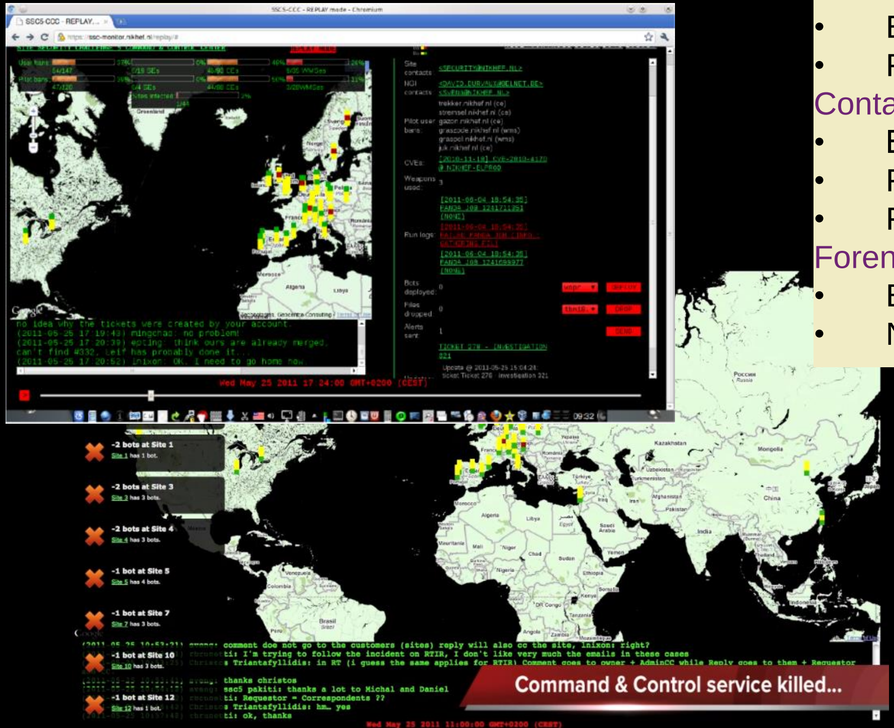

Communication:
- Endpoints valid?
- Form/Content OK ?

Containment
- Ban "malicious
- Find/Stop mali
- Find submissio

Forensics
- Basic Forensic
- Network traffic

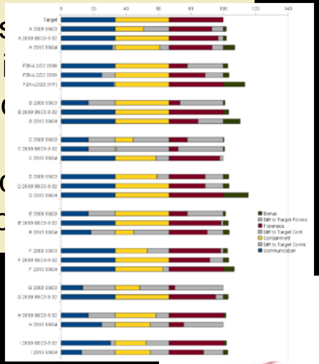

Nikhef

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

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

Security Incident Response Trust Framework for Federated Identity

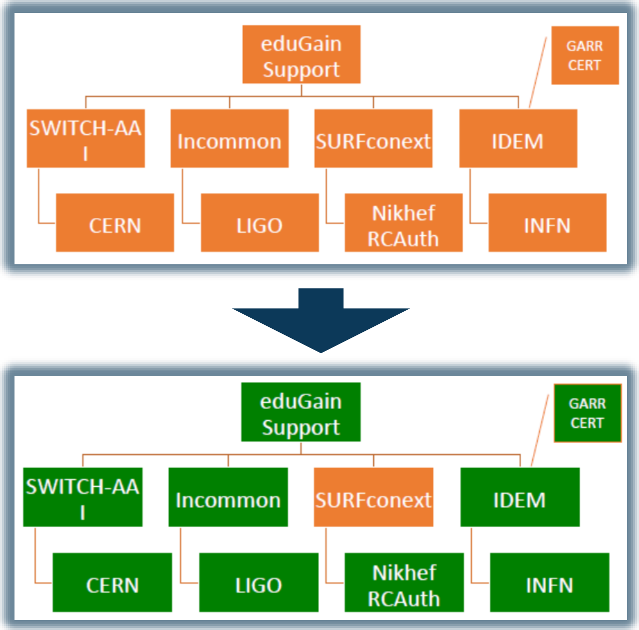

PARTIES INVOLVED IN RESPONSE CHALLENGE

REPORT-OUTS SEE **HTTPS://WIKI.GEANT.ORG/DISPLAY/AARC/SIRTFI+COMMUNICATIONS+CHALLENGES%2C+AARC2-TNA3.1**

Sharing threat intel – working with our community

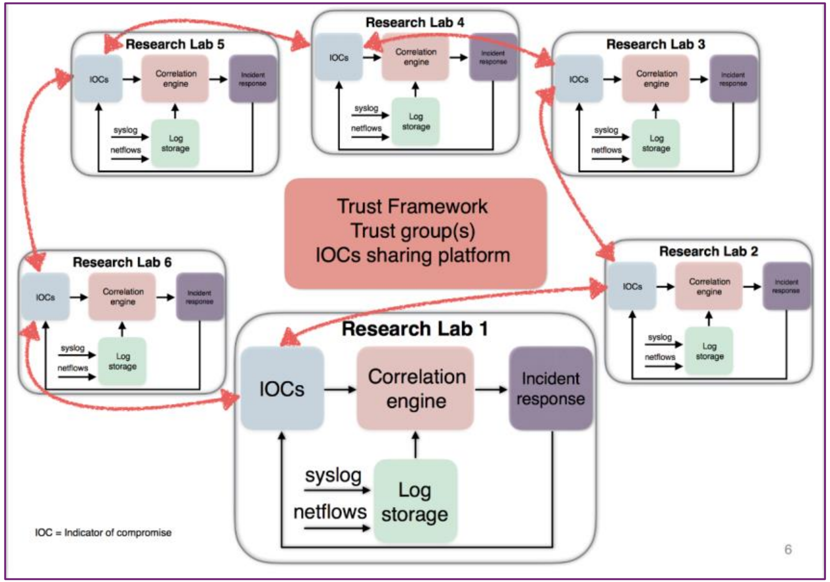

AARC I-051 Guide to federated incident response
https://aarc-community.org/guidelines/aarc-i051/

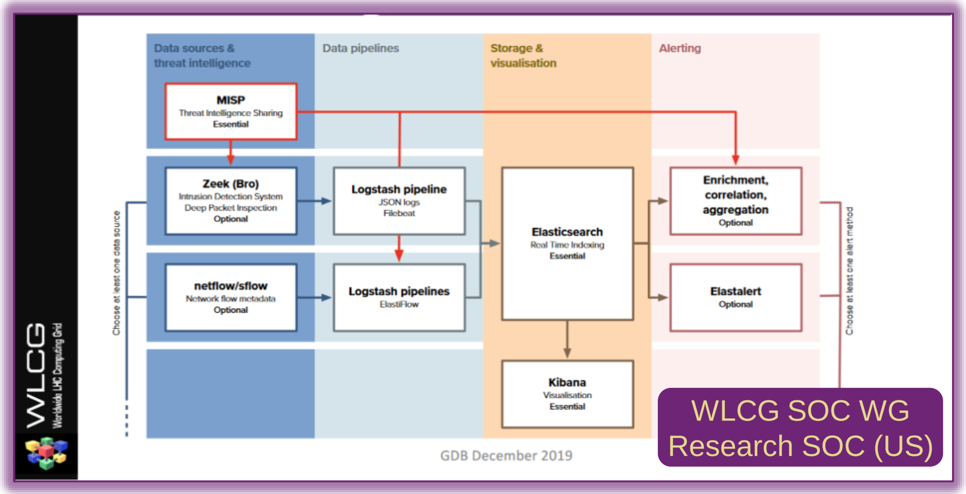

Nikhef SOC – NDPF traffic analysis

many 'false warnings': needs tuning

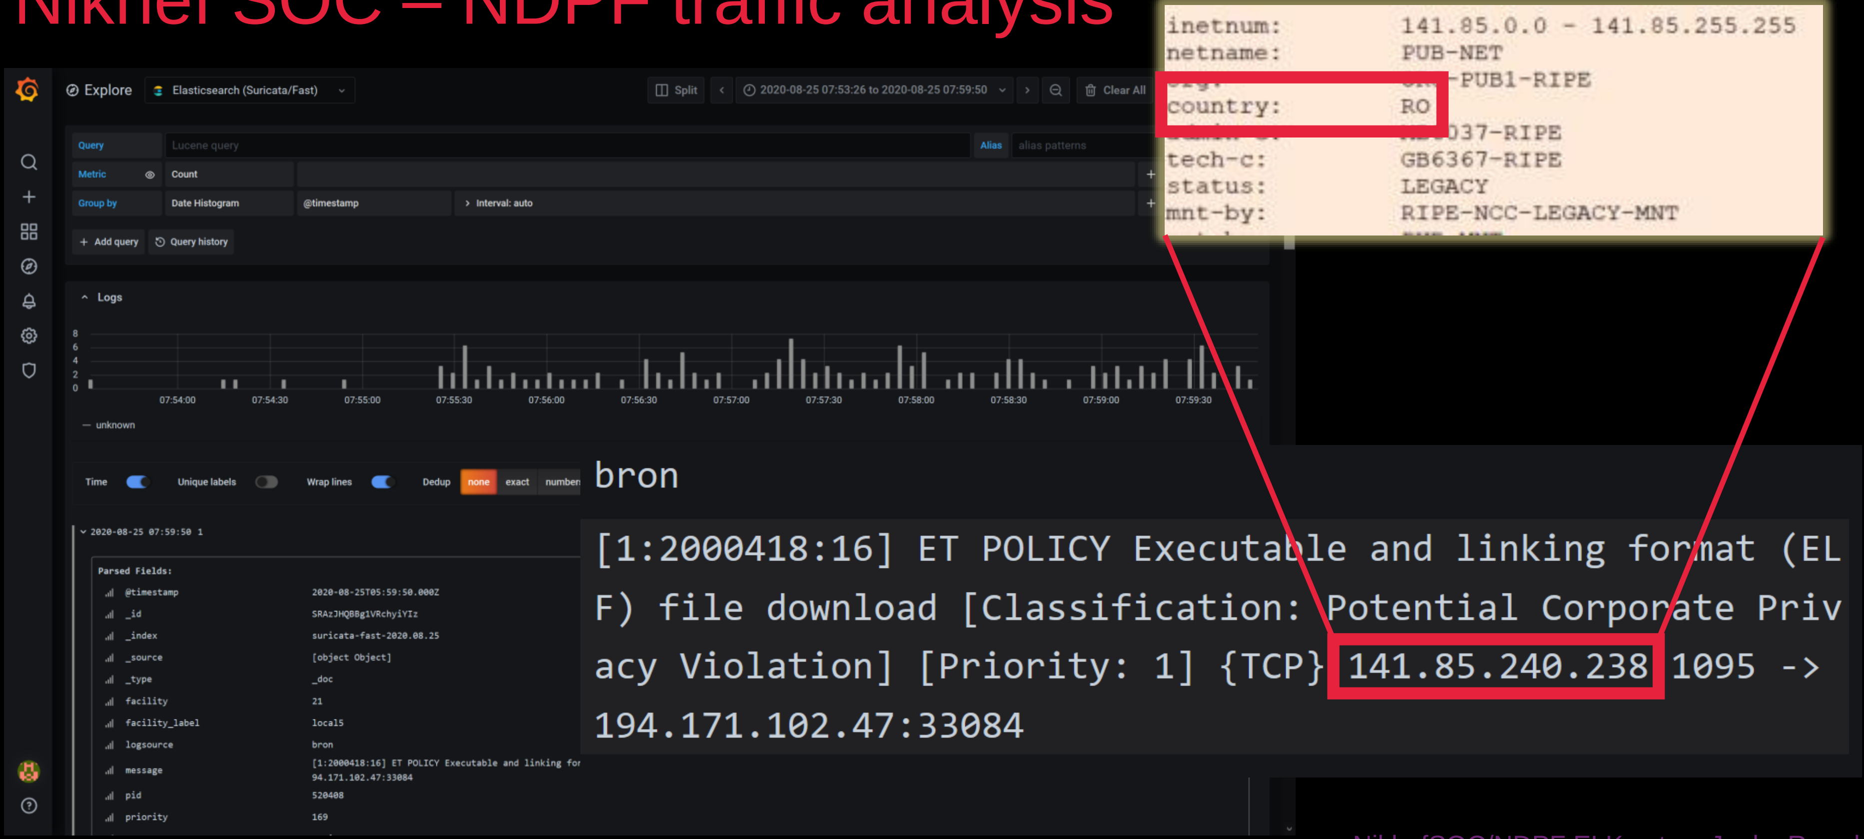

NikhefSOC/NDPF ELK setup: Jouke Roorda

e-Infrastructures: EGI, EUDAT, GEANT, PRACE, … and DNI!

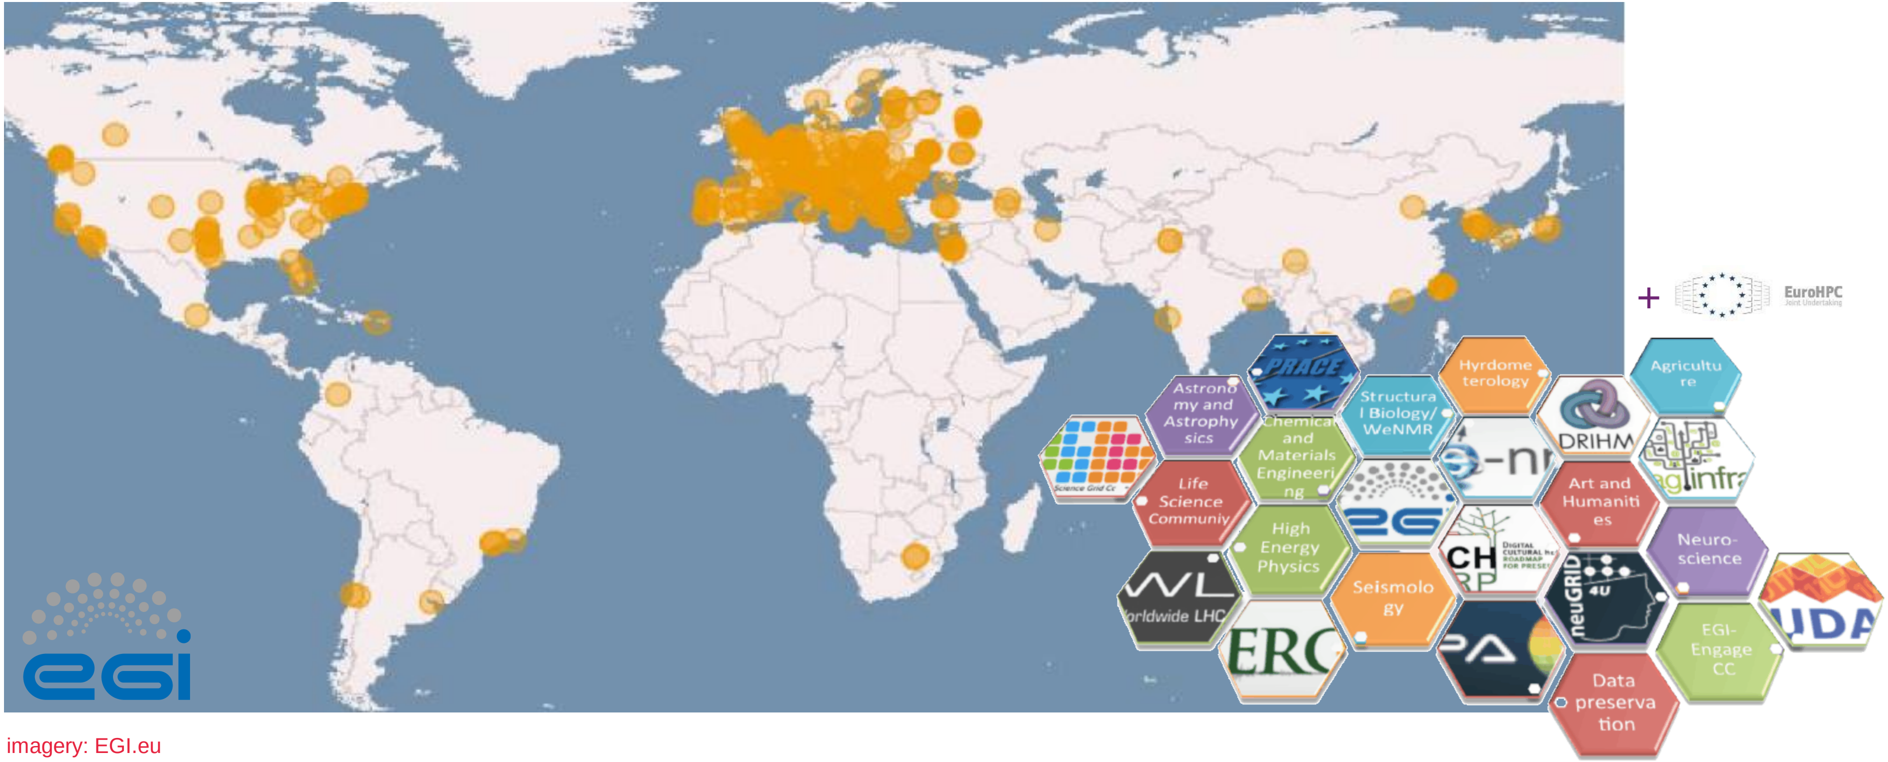

imagery: EGI.eu

‘DNI coordinated by SURF’

Coordinated **D**utch **N**ational e-**I**nfrastructure

- Single application portal (at SURF and NWO)
- Resources allocated at most-suitable partners
- Federated management and common innovation

BiG Grid
the dutch e-science grid

SURF

Nikhef

university of groningen

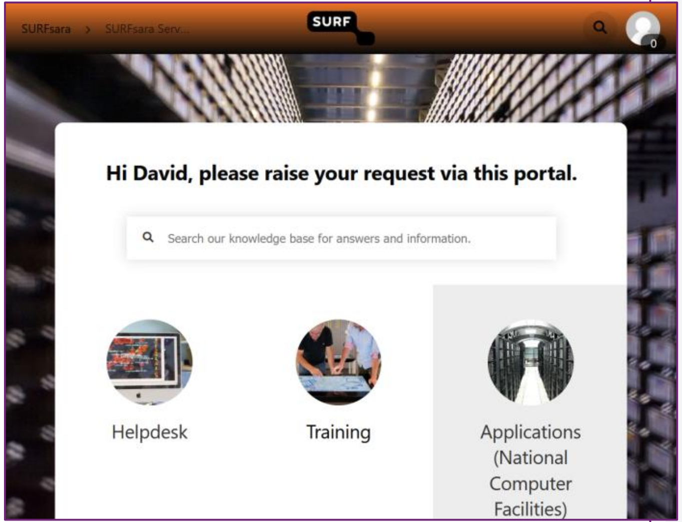

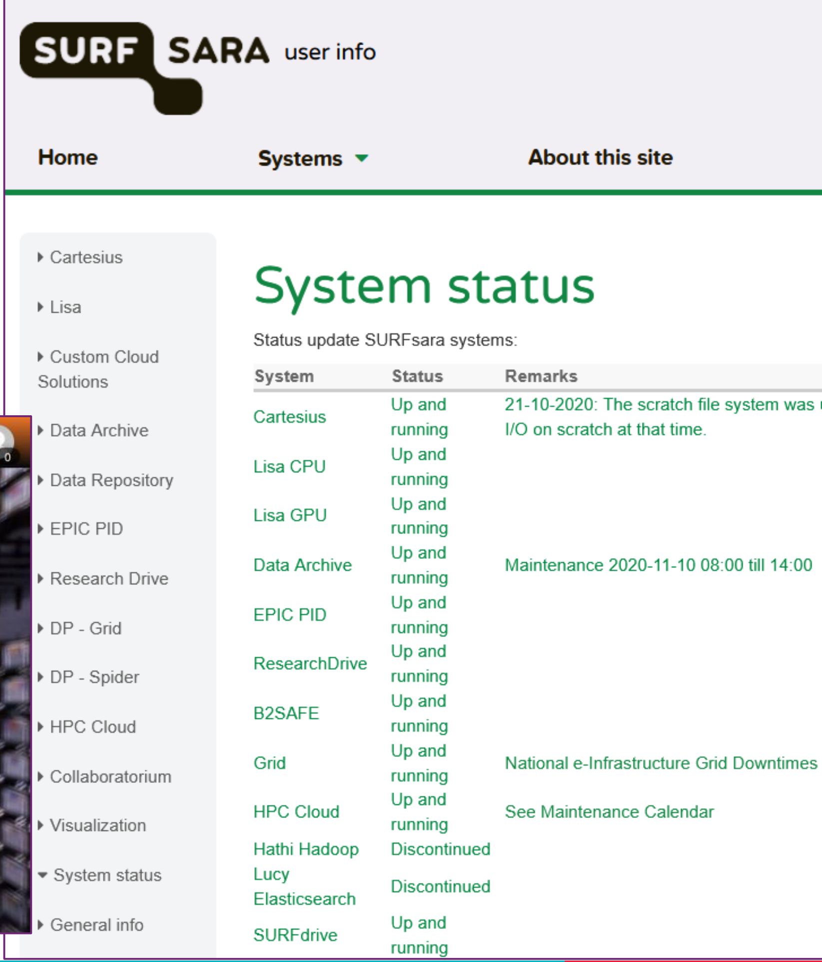

https://servicedesk.surfsara.nl/jira/plugins/servlet/desk/portal/1

Connecting resources – people – organisations – data

Programma

Rekentijd Nationale Computersystemen

Voorlezen Mail
Print Delen

‹ Naar de lijst van programma's

› **Rekentijd nationale computersystemen**
Achtergrond
Organisatie
Projecten
Contact

Geavanceerde (super)computersystemen worden gebruikt voor technisch-wetenschappelijk onderzoek waarbij grote rekenproblemen moeten worden opgelost. Door tijd beschikbaar te stellen via het programma 'Rekentijd nationale computersystemen' maakt NWO geavanceerde nationale computerfaciliteiten beschikbaar voor wetenschappelijk onderzoek. Hierdoor is hoogwaardig en competitief onderzoek mogelijk in Nederland.

Geavanceerde computersystemen worden bijvoorbeeld gebruikt voor berekeningen voor de weerverwachting, waarbij grote hoeveelheden data van satellieten en weerstations verwerkt moeten worden. Andere voorbeelden zijn de verwerking van grote hoeveelheden data in de radioastronomie, klimaatonderzoek, onderzoek naar grote historische archieven en tekstcorpora, of genoom- en eiwitanalyse.

Nieuws

8 oktober 2020
› NWO honoreert drie aanvragen voor Rekentijd op de Nationale Computersystemen

9 juli 2020
› NWO honoreert 10 aanvragen voor Rekentijd op de Nationale Computersystemen

20 mei 2020
› NWO honoreert 24 aanvragen voor Rekentijd op de Nationale Computersystemen

› **Al het nieuws voor Rekentijd nationale computersystemen**

Kalender

31 Rekentijd Nationale

JE MAINTIENDRAI

Staatscourant van het Koninkrijk der Nederlanden

Datum publicatie	Organisatie	Jaargang en nummer	Rubriek
07-06-2018 09:00	Nederlandse Organisatie voor Wetenschappelijk Onderzoek	Staatscourant 2018, 31287	Overig

The Dutch National e-Infrastructure

In this call, all applicants are asked to indicate the project's e-Infrastructure needs, in terms of compute hours, data storage capacity, lightpath connectivity, or otherwise. A 'use-or-explain' policy will be applied, meaning that

– projects *without* e-Infrastructure needs are asked to give a brief explanation;
– projects with clear e-Infrastructure needs are expected to select the hardware resources and services as part of the Dutch National e-Infrastructure as first option, and to indicate the expected extent of use;
– projects with clear e-Infrastructure needs that aim to use international (e.g. PRACE, XSEDE, etcetera) or commercial (e.g. web, cloud, etcetera) hardware and services instead are required to give a brief explanation.

The use of the Dutch National e-Infrastructure is not a requirement, nor is it a formal review criterion. However, in all cases in which the Dutch National e-Infrastructure is not used, a justification should be provided.

In this call, the Dutch National e-Infrastructure is defined as follows:

The definition distinguishes between hardware resources and services available to all researchers in the Netherlands (Category I), and those made available to a selected subset (Category II). The Category I e-Infrastructure, outlined below, is formed by the hardware resources and services provided and maintained by SURFsara, SURFnet, DANS, and – in part – also by Nikhef and RUG-CIT.

accessible to a selected group of researchers following thematic or geographic criteria. Examples of

https://www.nwo.nl/onderzoek-en-resultaten/programmas/Rekentijd+nationale+computersystemen

Balanced infrastructure - based on our joint science cases

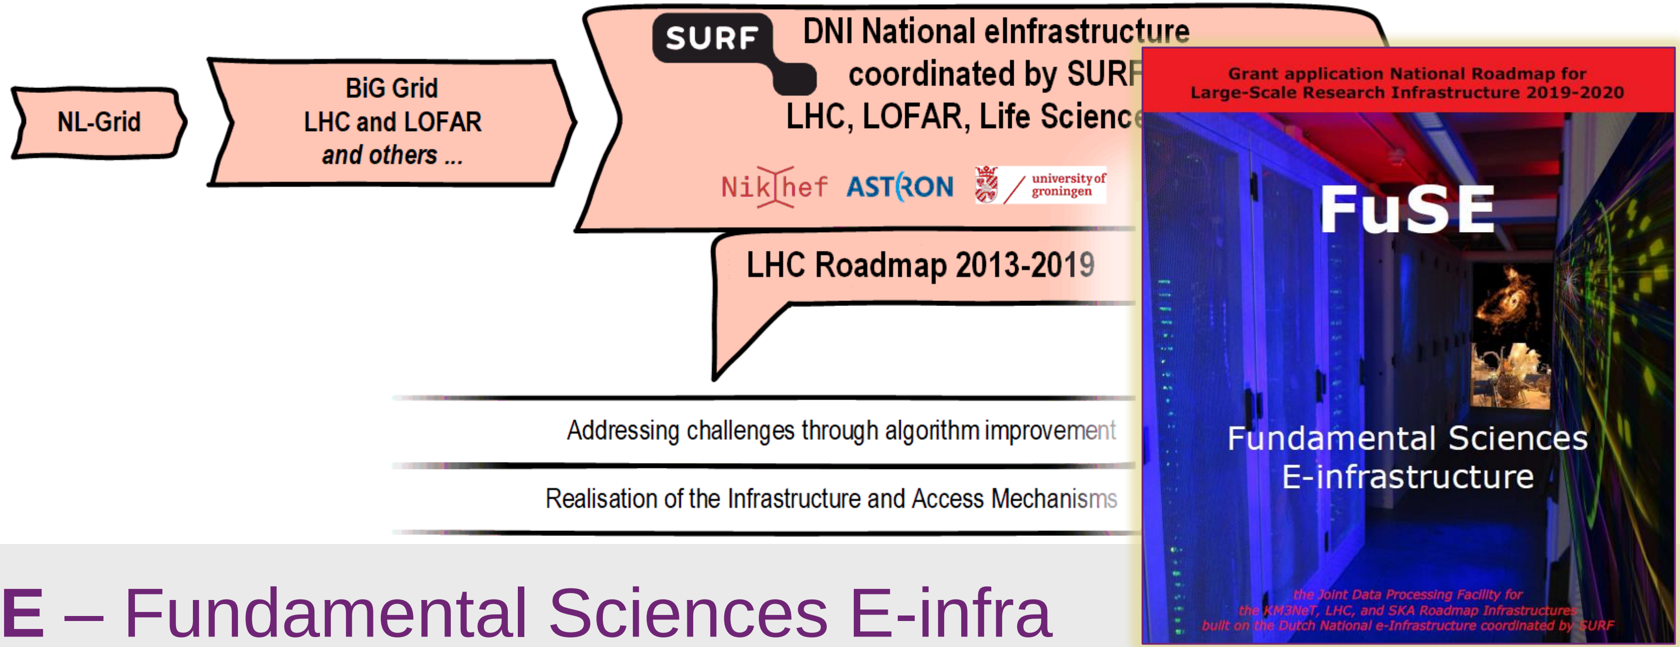

FuSE – Fundamental Sciences E-infra

an integrated infrastructure
for algorithms, hardware, networking, and collaboration

https://www.fuse-infra.nl/

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

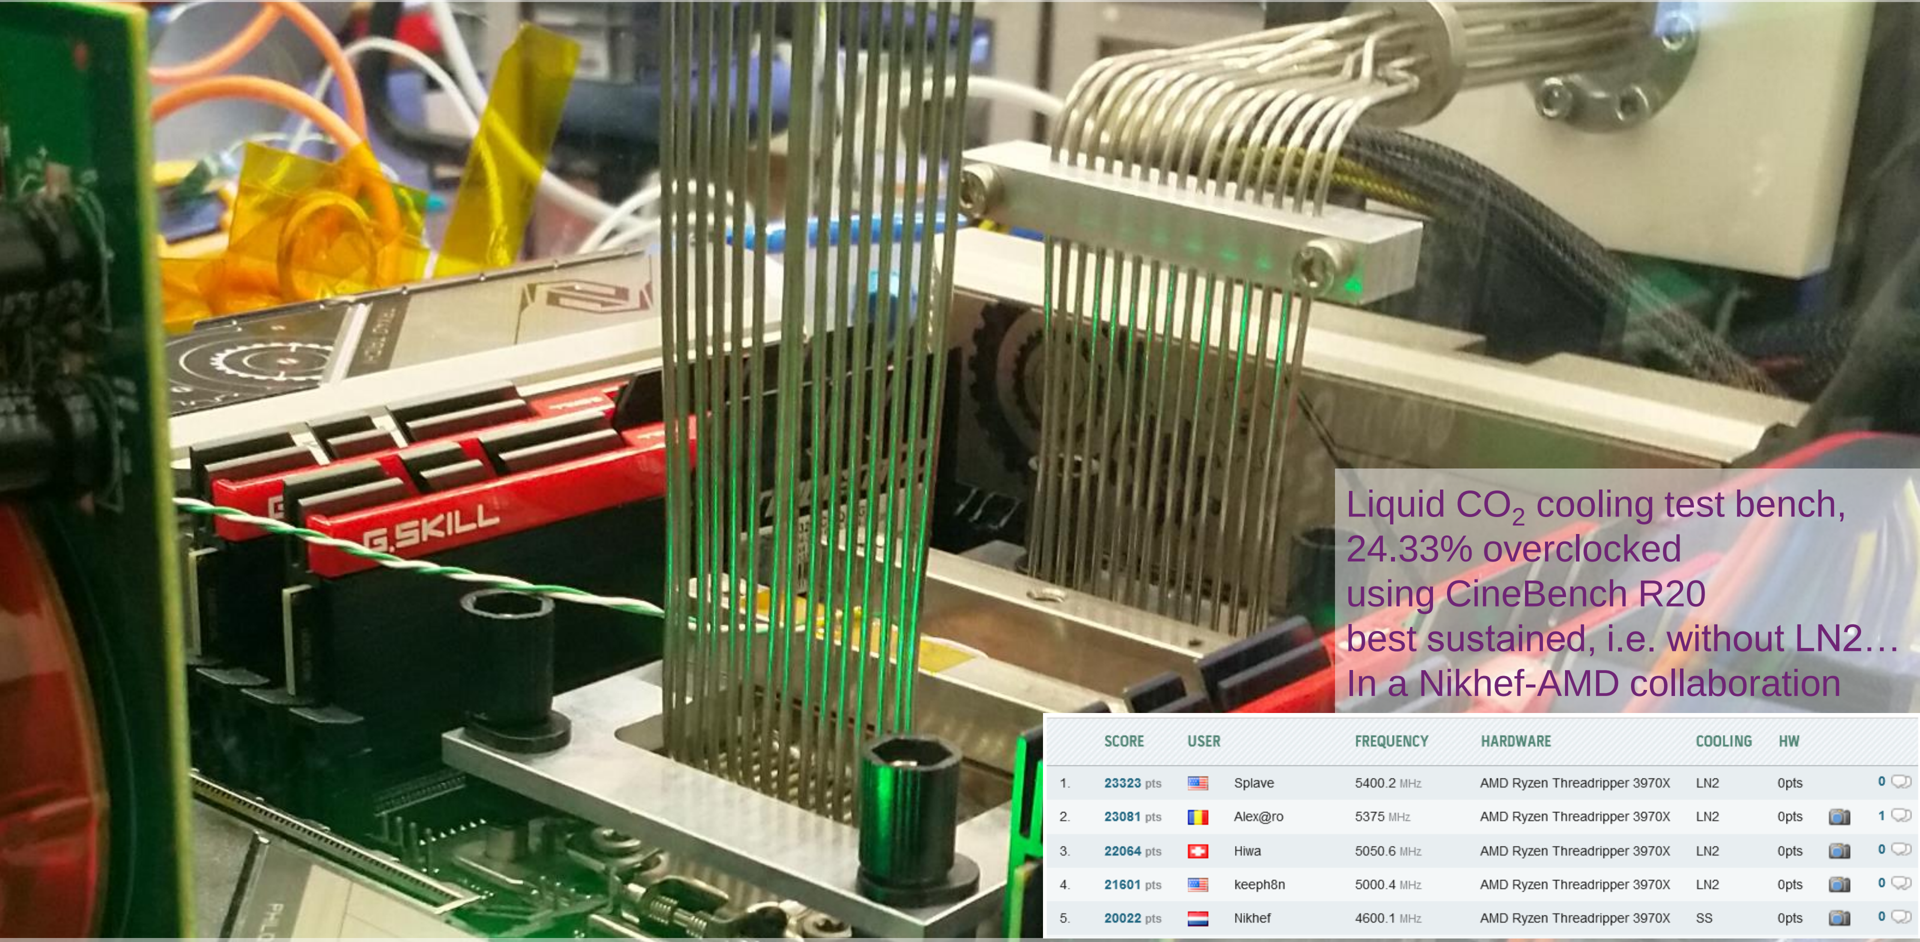

Liquid CO_2 cooling test bench,
24.33% overclocked
using CineBench R20
best sustained, i.e. without LN2…
In a Nikhef-AMD collaboration

	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

Let It All Collaborate!

Nikhef

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