Nikhef

Maastricht University

Federated computing ecosystems for science

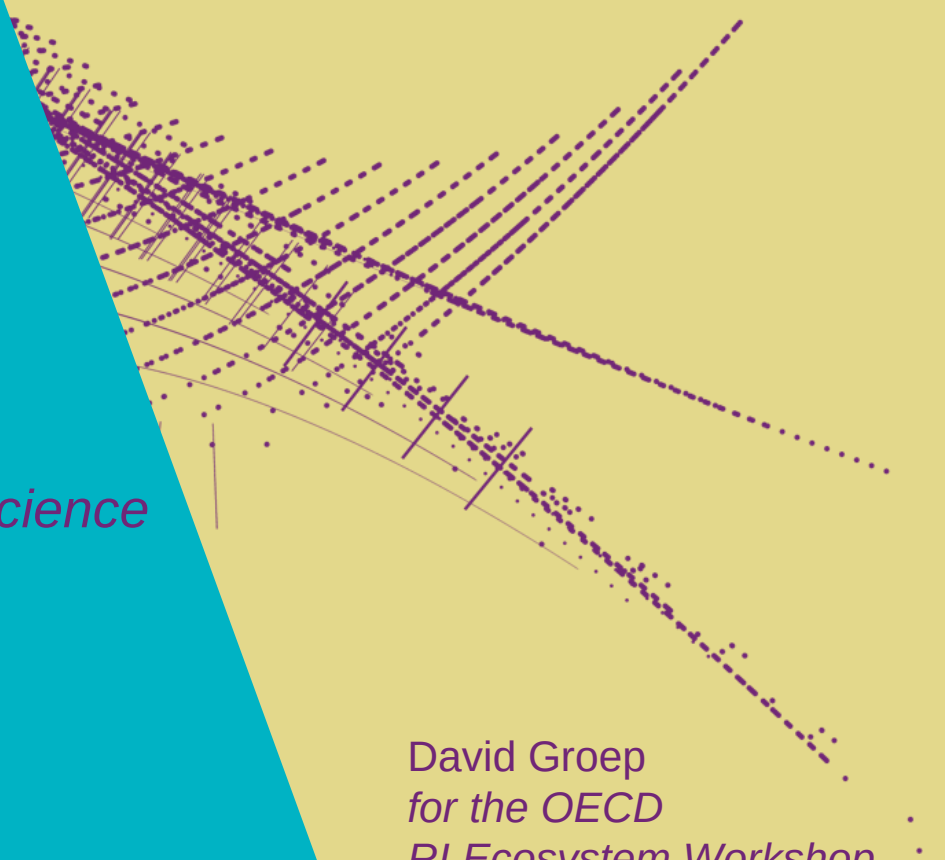

Nikhef, its Science Programme, and the Dutch National e-Infrastructure

David Groep
for the OECD
RI Ecosystem Workshop
November 4, 2024

Welcome to Nikhef the Dutch National Institute for Sub-atomic Physics

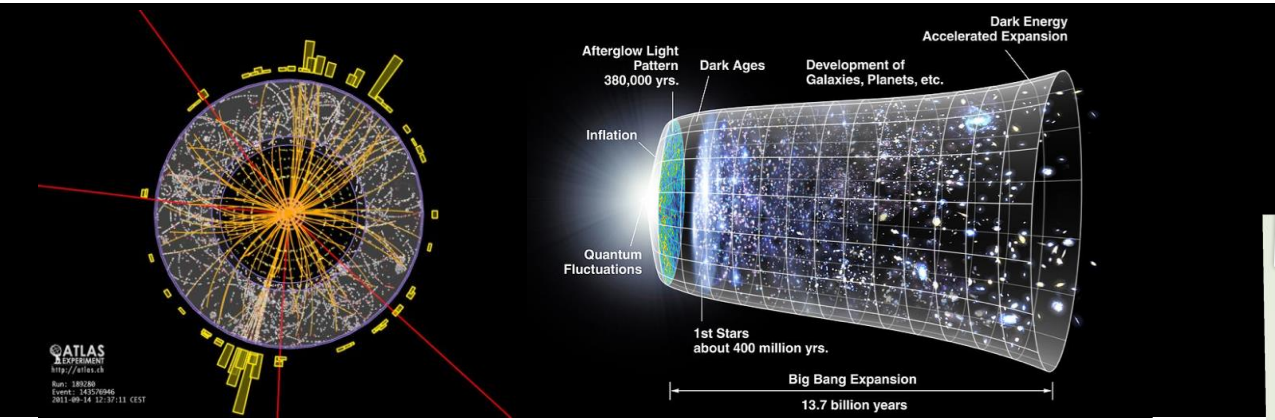

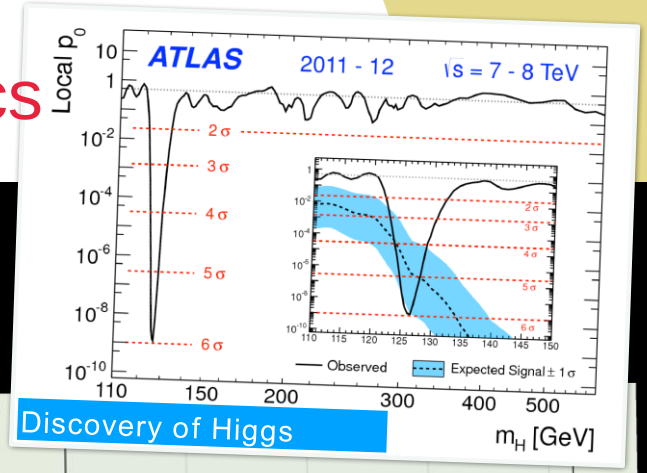

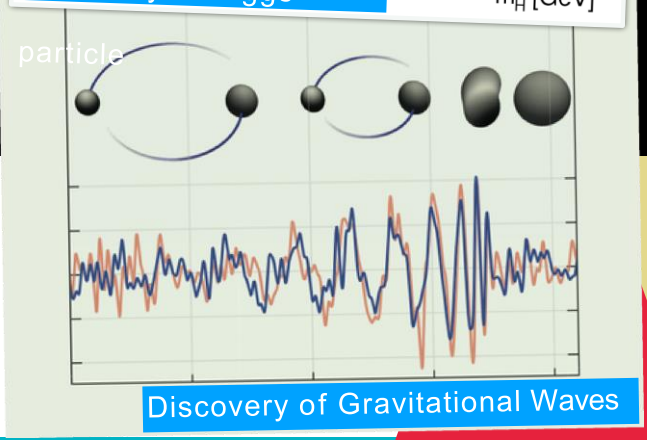

We probe our world, made of particles and fields,

- with collider physics, primarily at CERN
- astroparticle physics:
  particles, radiation, and ripples coming from the universe

images: atlas.cern.ch, CERN, LIGO-Virgo Collaboration

The Nikhef Partnership

In itself a collaborative ecosystem

- university partners co-lead (most) research programmes
- aligned with a joint national strategy

Permanent Staff	96
PhD candidates	125
Postdocs	43
Technical/engineer	88
Support	33

Nikhef Scientific Programmes

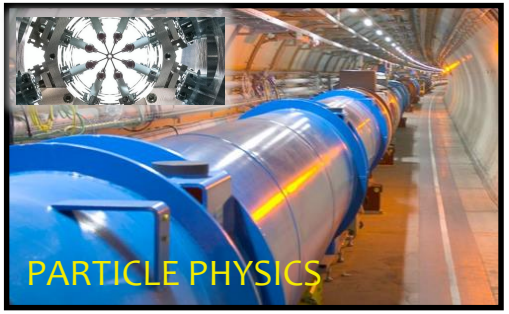

- Atlas
- LHCb
- Alice
- eEDM

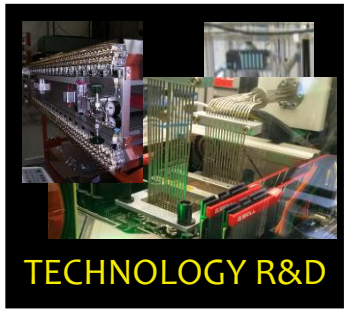

- Detector R&D
- Physics Data Processing

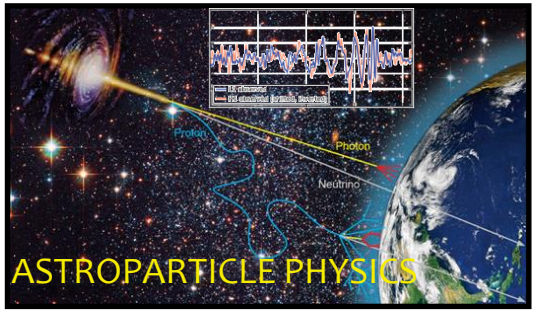

- Neutrinos
- Gravitational waves
- Cosmic Rays
- Dark Matter

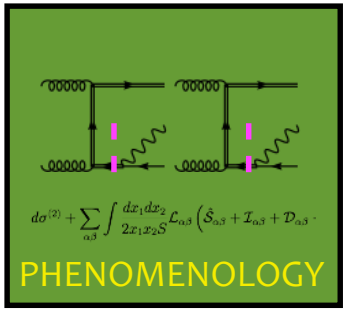

- Theoretical Physics

Technical Engineering and technology transfer

`The research at Nikhef relies on the development of innovative technologies. The knowledge and technology transfer to third parties, i.e., industry, civil society and general public, is an integral part of Nikhef's mission.'

pictures from Nikhef's 'Dimensions' magazine and Computing Office Hours; https://www.nikhef.nl/en/nikhef-mission/

Infrastructure-intensive science

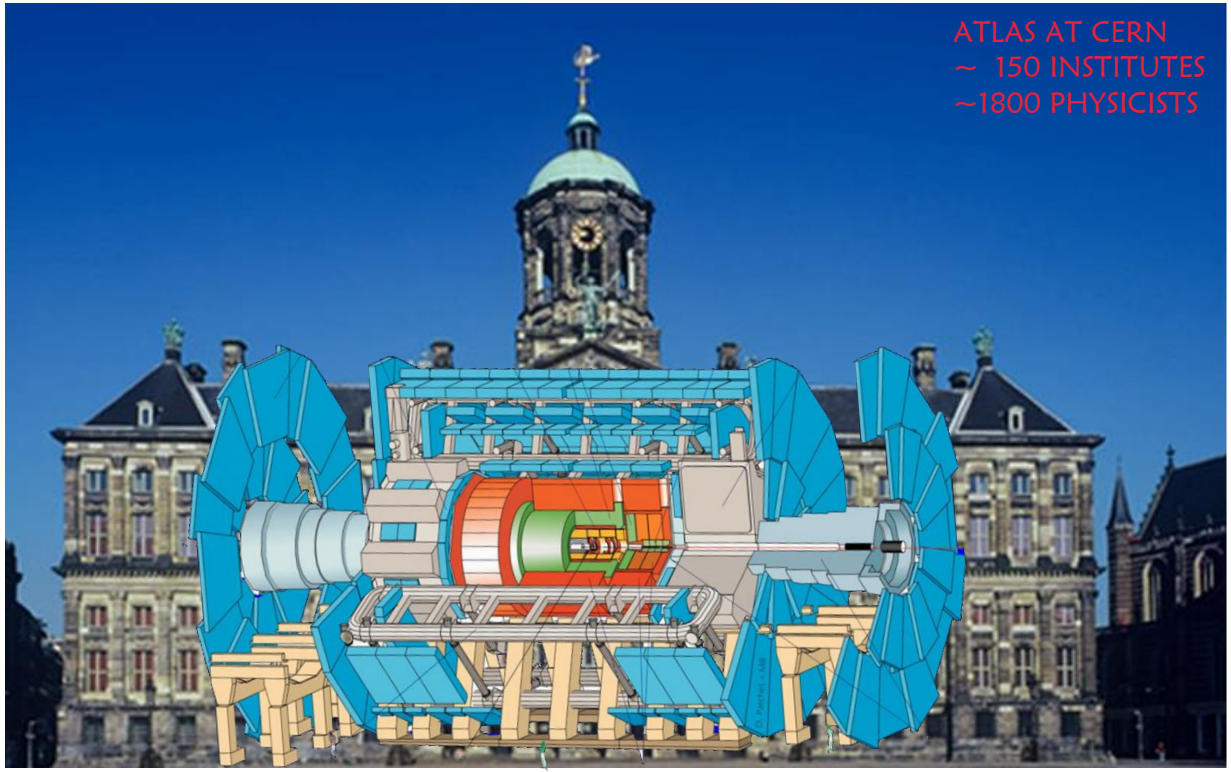

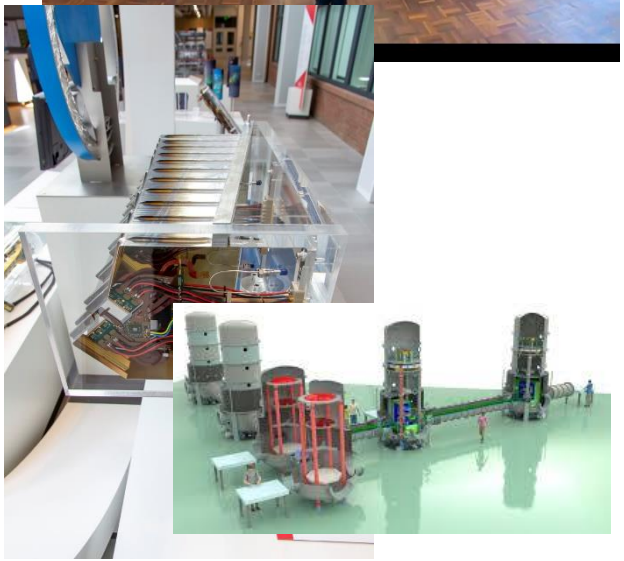

Slide materials from: Stan Bentvelsen, 2016, ATLAS collaboration, CERN; LHCb VELO and RF box at Nikhef and ET Pathfinder visualisation: Marco Kraan, Nikhef

Nikhef at CERN – LHCb, Atlas, Alice

CMS

Large Hadron Collider

LHCB

ALICE

ATLAS

Imagery: CERN, European Organisation for Nuclear Research, and Marco Kraan, Nikhef

Atlas: up to 40 Tbyte/day of fresh raw data

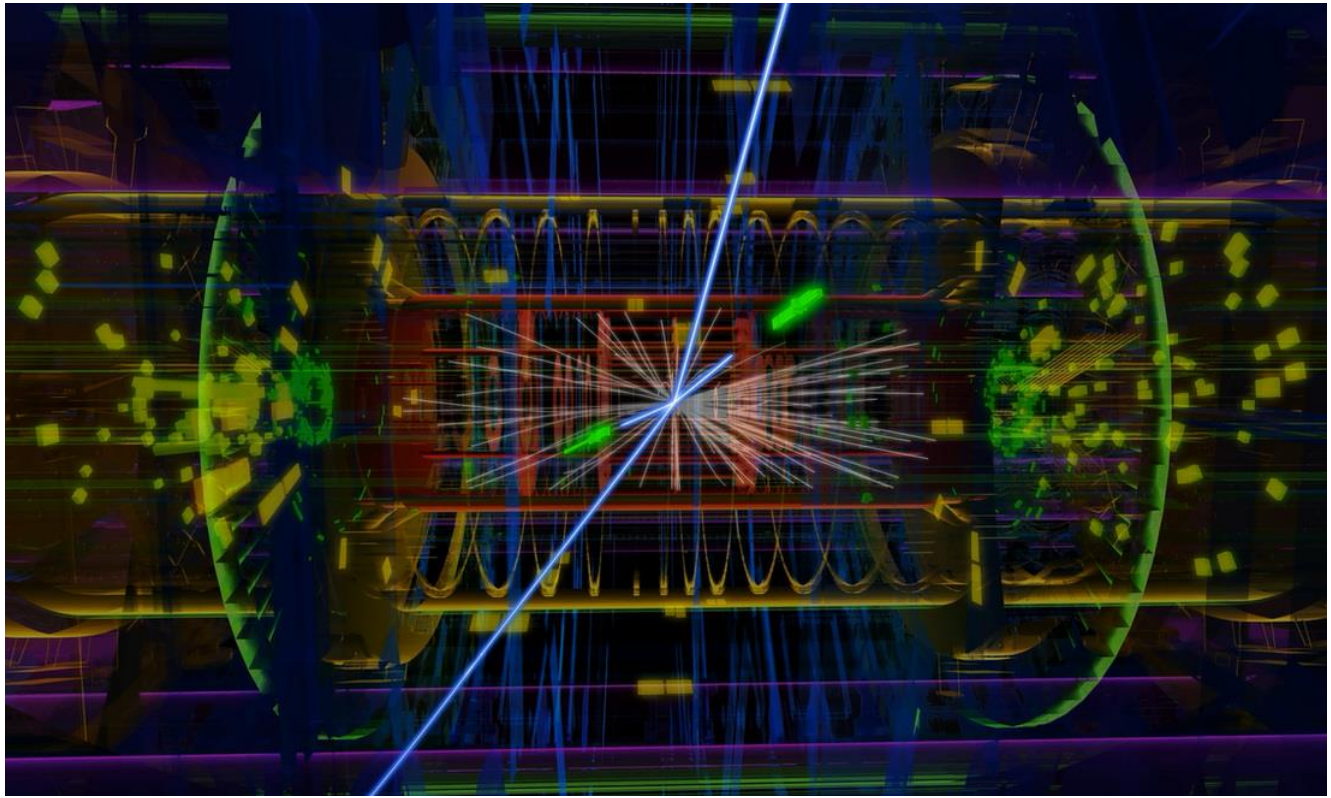

Image source: CERN, Atlas collaboration; ATLAS approx. 140M active channels, mean event size ROD 1.6 MByte

And decelerate for ultra-high-precision measurements

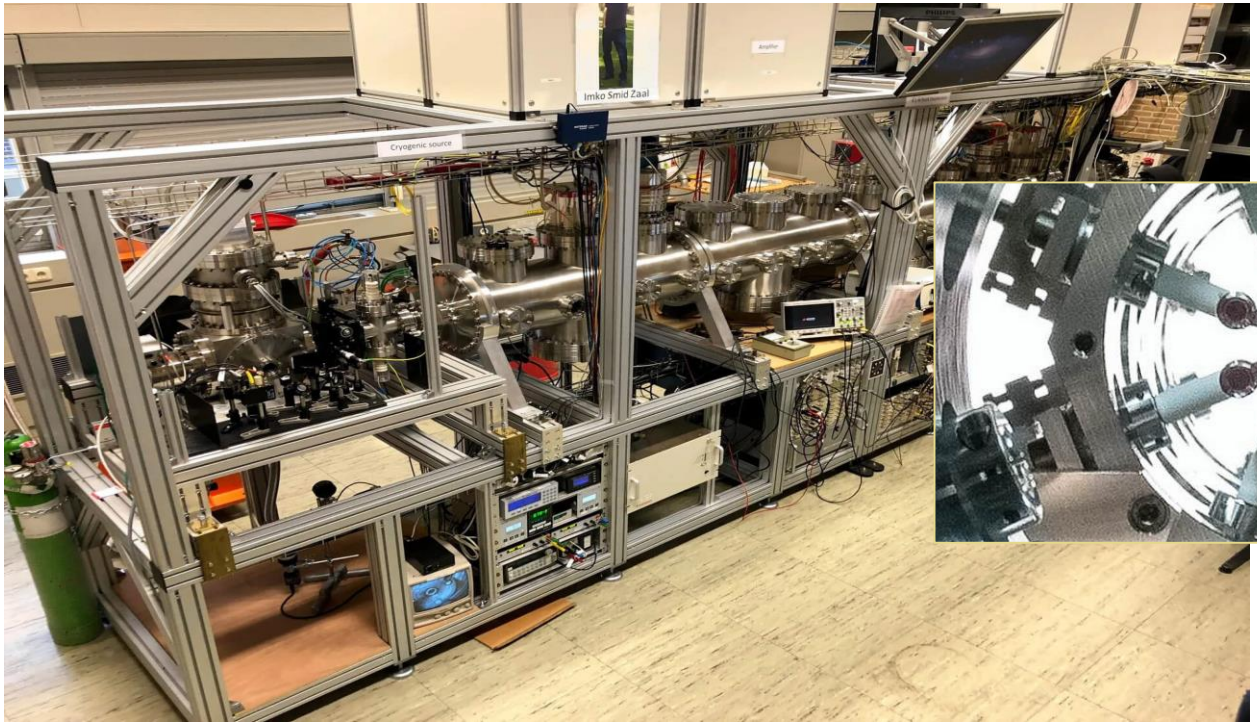

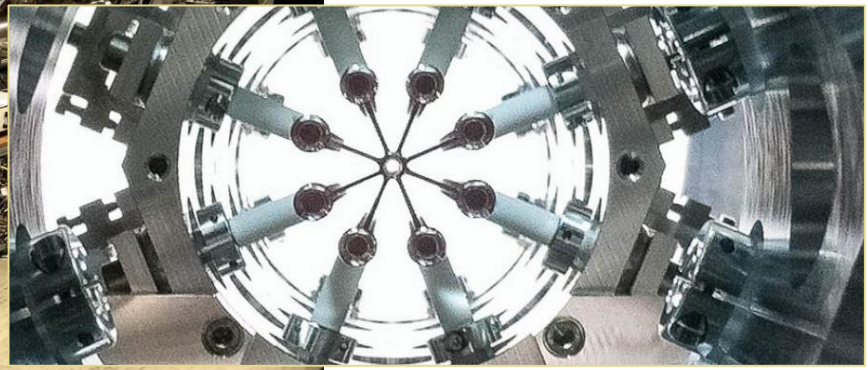

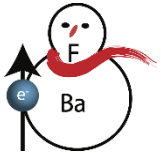

https://www.eedm.nl/ Nikhef and Rijksuniversiteit Groningen, images: Steven Hoekstra

Particles at Extreme Energies

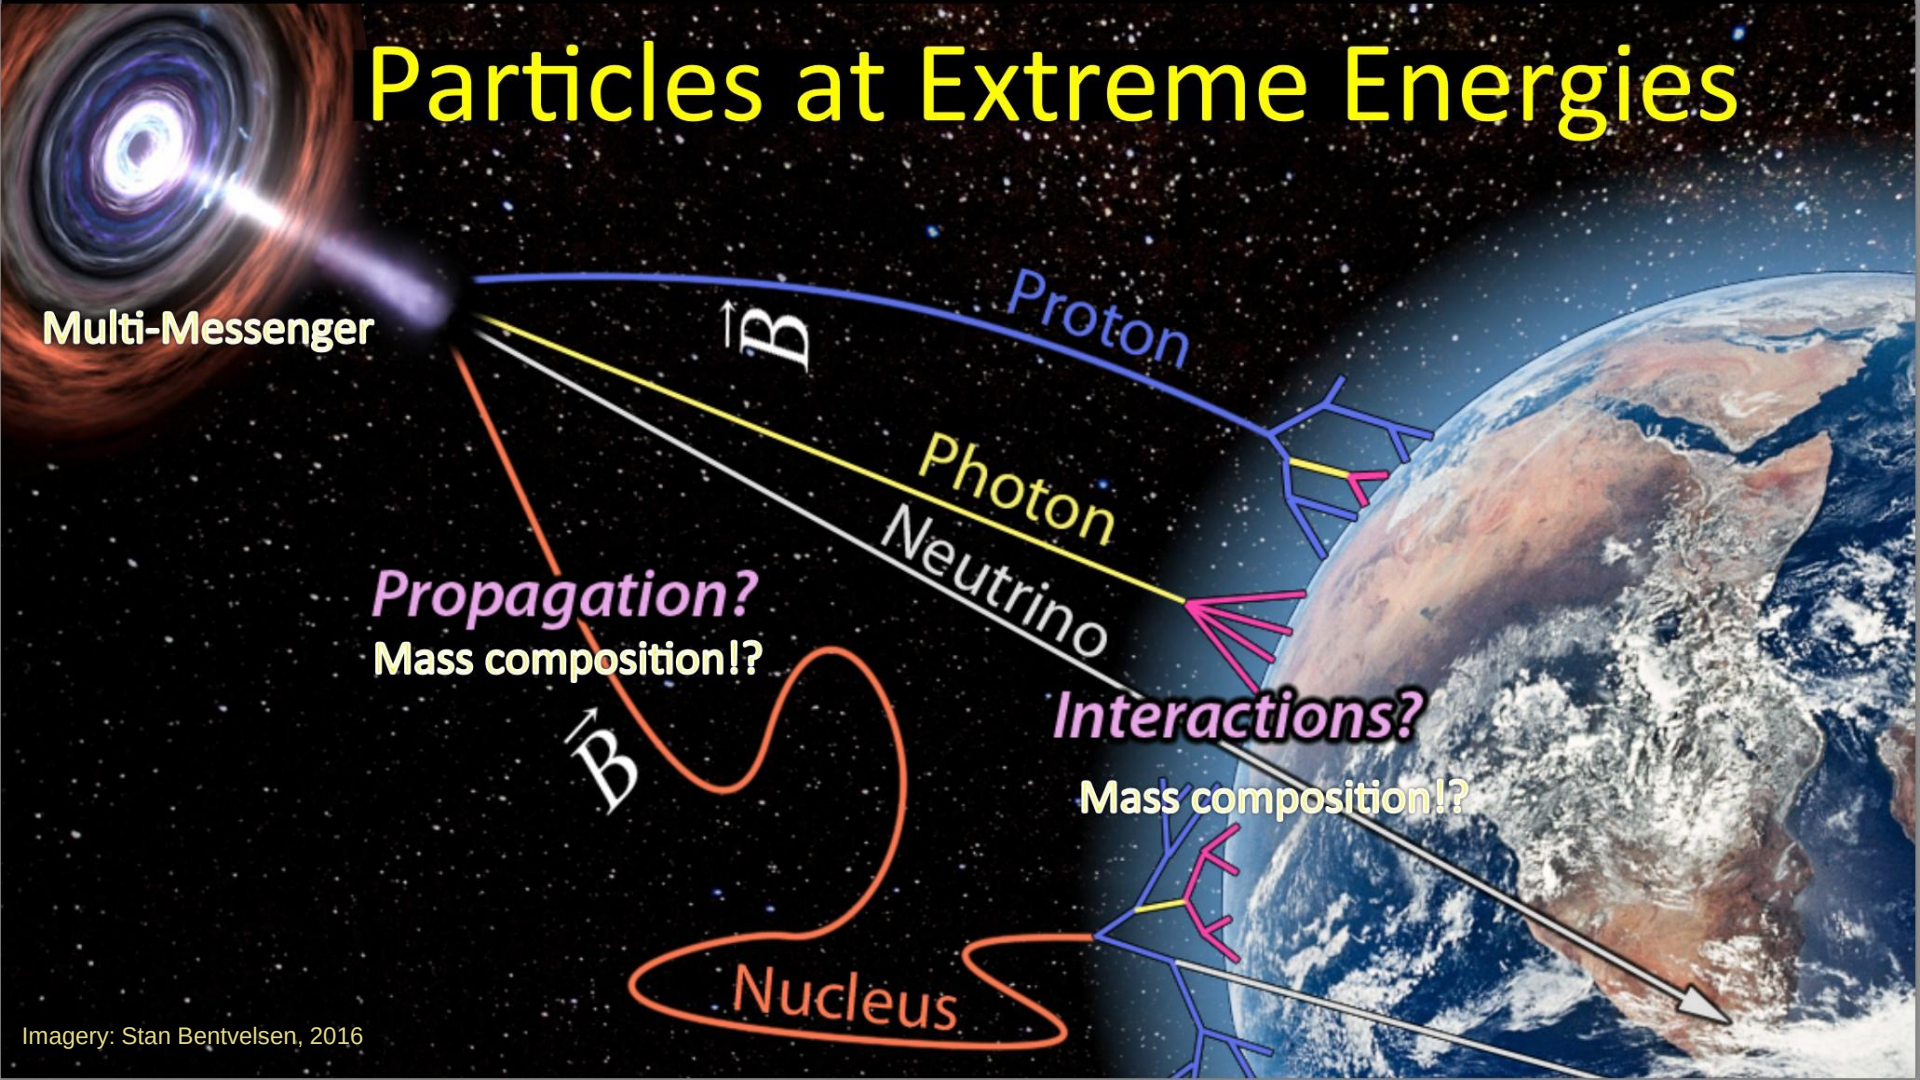

Imagery: Stan Bentvelsen, 2016

www.et-emr.eu

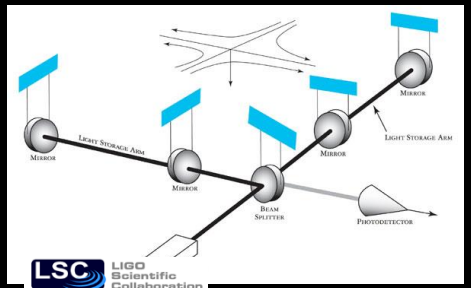

Einstein Telescope projected in the **Euregio Maas Rijn**, images: Marco Kraan; ET Pathfinder at UM, Maastricht, NL

GW150914 event: gw-astronomy collaborations, LIGO

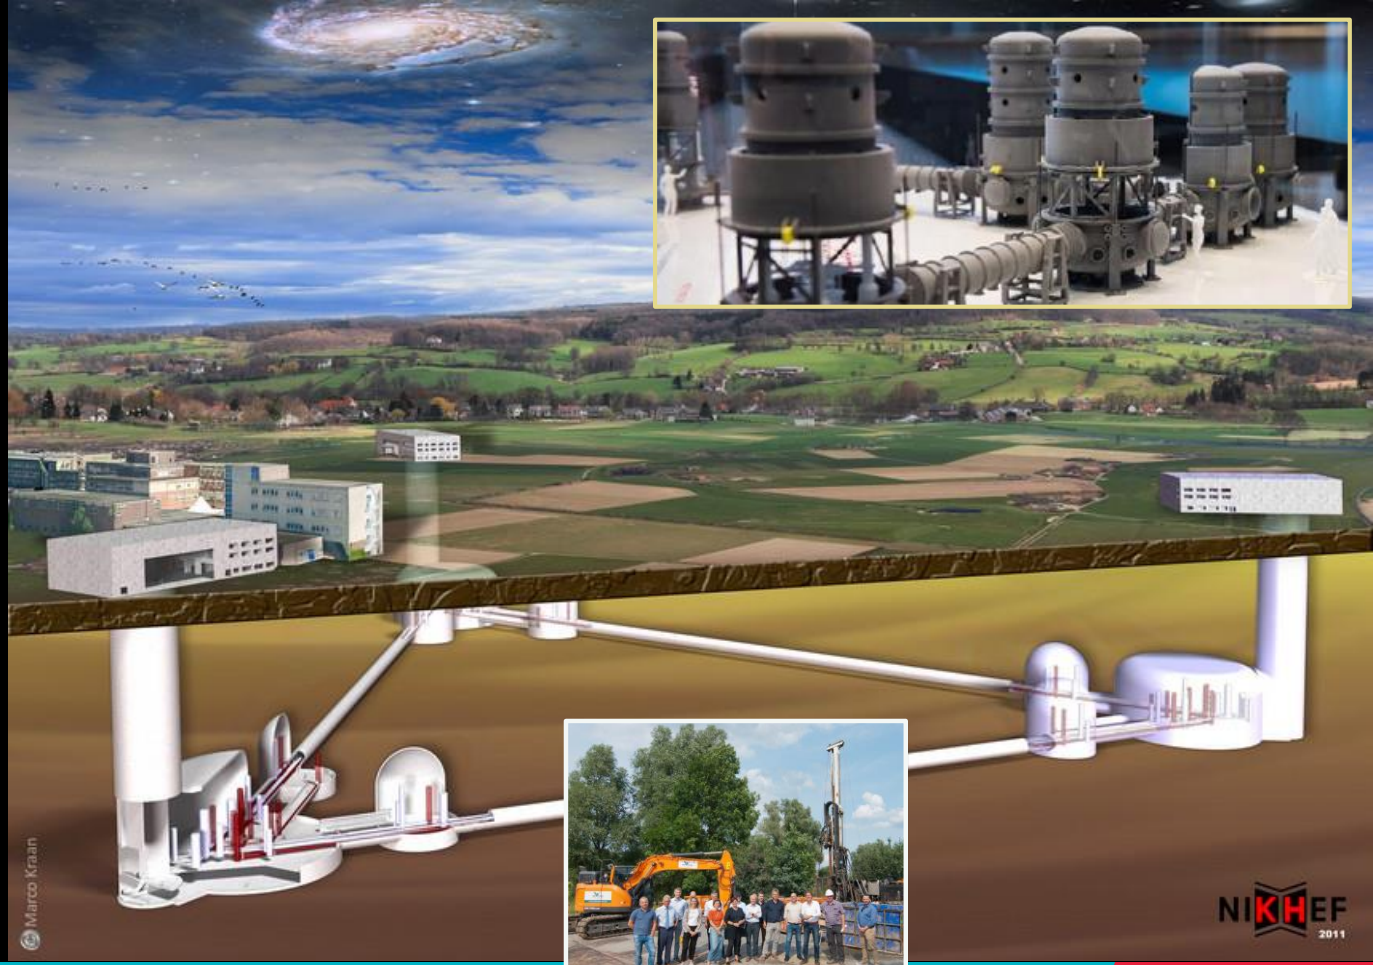

The astro-particle physics programmes Nikhef

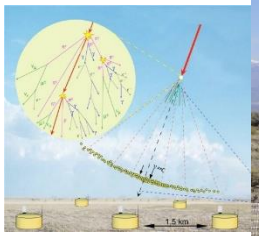

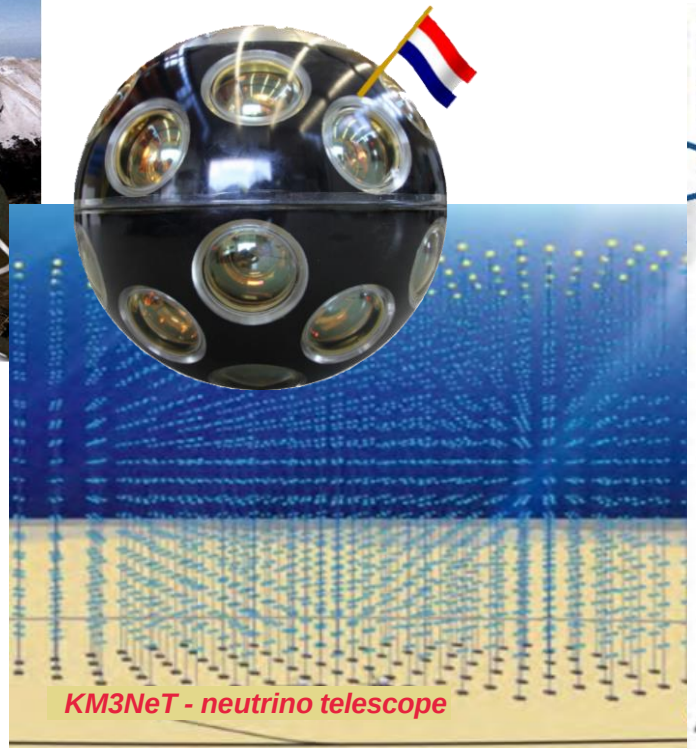

Slide: Stan Bentvelsen, Nikhef SEP 2023, KM3NeT collaboration, Virgo Collaboration, Xenon-nT collaboration, Pierre-Auger collaboration

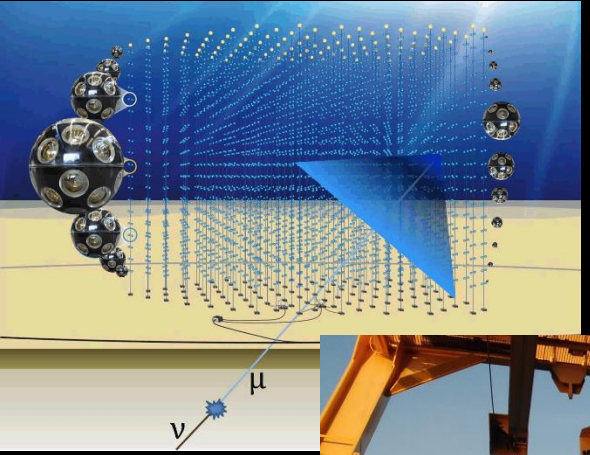

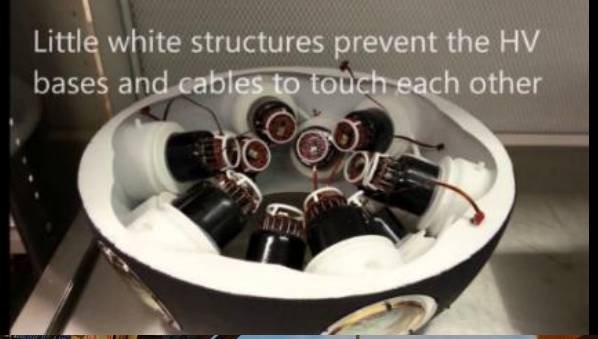

Image sources:
Nikhef, NIOZ,
KM3NET collaboration,
NRC Handelsblad

Enabling Research Programmes – next gen infrastructure

Detector R&D

Theoretical Physics

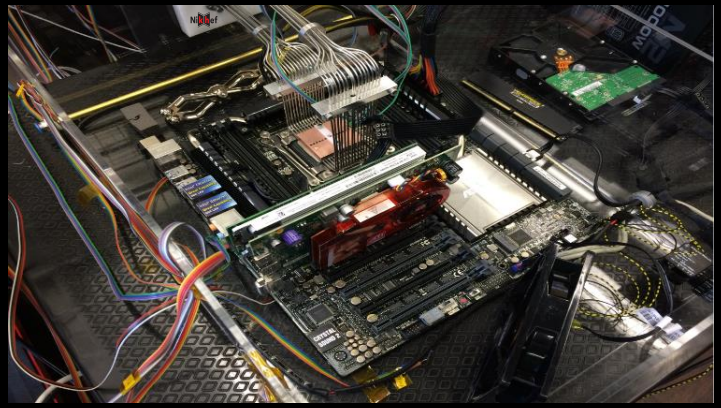

Physics Data Processing

LHC Computing – a global infrastructure

> 2 million CPU cores
> 2000 Petabyte
disk + archival
170+ data centres
42+ countries
13 ‘Tier-1 sites’
NL-T1 jointly by
SURF & Nikhef

built on many ‘generic’ e-Infrastructures
EGI, GEANT, EuroHPC
OpenScienceGrid
NSF ACCESS-CI, ESnet,
KEK, CalculCanada, …

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

Data in the Dutch National e-Infrastructure

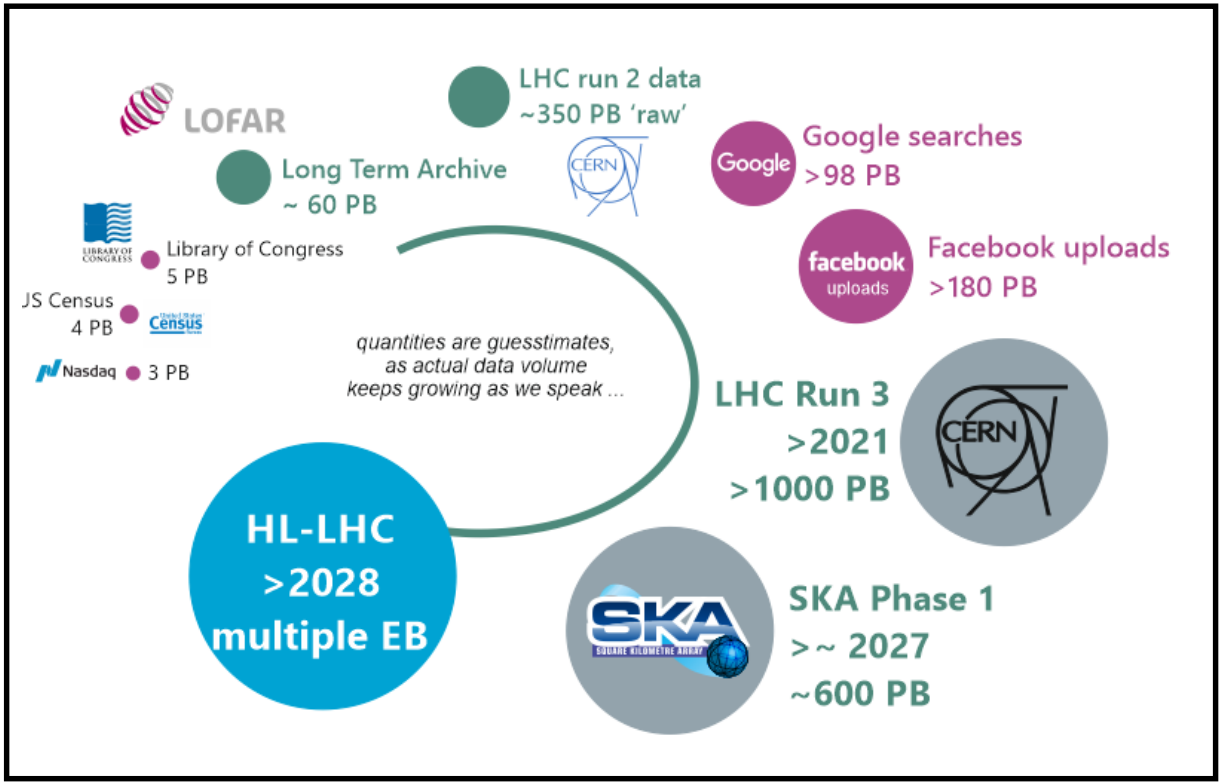

approximate volumes for LoC, Google searches and FB uploads from ~ 2018. Source: WLCG publicity presentations
Imagery: SKA mid (South Africa) courtesy SKAO. Silicon tracker composite photo Nikhef, KM3NeT DOM deployment module: Nikhef and NIOZ

FuSE

Nikhef

Fundamental
Sciences
E-infrastructure

ASTRON
Netherlands Institute for Radio Astronomy

SURF

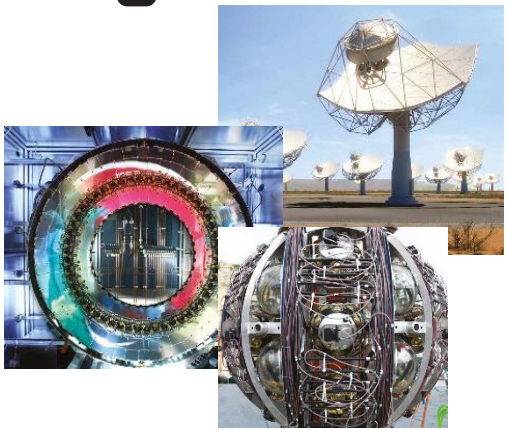

Data Processing infrastructure – the Dutch SURF example

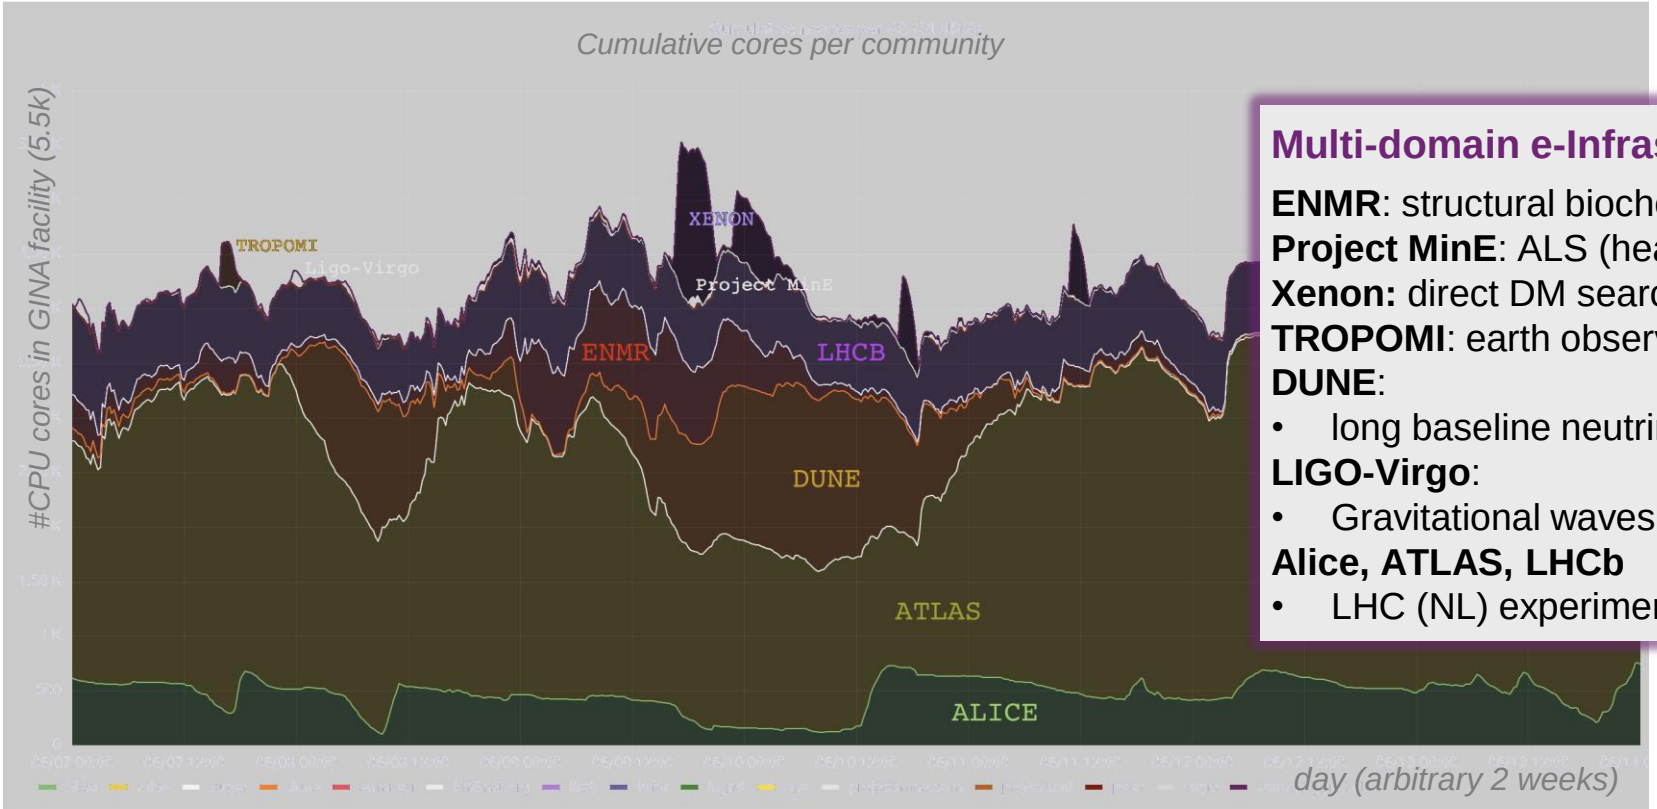

Multi-domain e-Infrastructure

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

Graphic: GINA DNI compute service coordinated by SURF

Coordinated e-Infrastructure: the SURF cooperative

SURF nationally aligns and supports computing for members (RPOs and education) and research council

- **networking** (the physical ones) with global links
- **HPC and HTC** facilities, data storage
- **federation services**: trust, identity, and security

Coordinates joint infrastructure, initiatives, and projects

- nationally with members and the LSRIs
- in Europe with GEANT, EGI, EUDAT, EuroHPC and ESFRIs,
- globally with collaborations like WLCG, SKA, IGWN, Internet2
- ‘digital competences’: FAIR data … and beyond
- technology innovation, digital autonomy, policy, and ‘futuring’

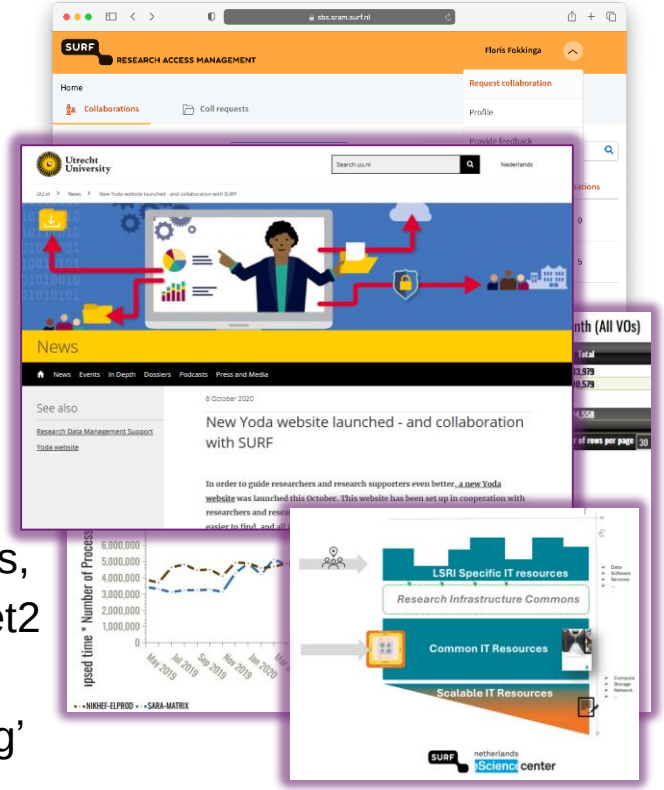

Images: EGI accounting (https://accounting.egi.eu/) for NL, all federated communities, 2020-2022; Yoda: Utrecht University and SURF, https://www.uu.nl/en/news/new-yoda-website-launched-and-collaboration-with-surf; SRAM: SURF Wiki; ; RI Commons and Common IT Resources: Magchiel Bijsterbosch, SURF

Scalable HPC: from local “Tier-2” to European “Tier-0”

Nikhef “Stoomboot”
Analysis Facility

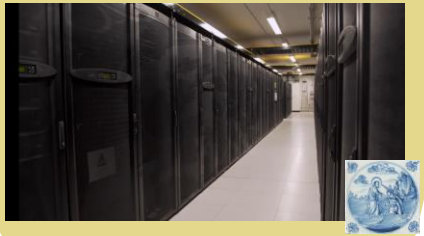

SURF National Infrastructure
foundation *as well as* stepping stone

EuroHPC Joint Undertaking

eosc

PRACE

You need **local expertise**
to enable exploitation
of European and
global resources

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

Sharing more than resources: data, software, research pipelines, expertise

Thematic Digital Competence Centres - beyond data stewards and 'dead' data for open science

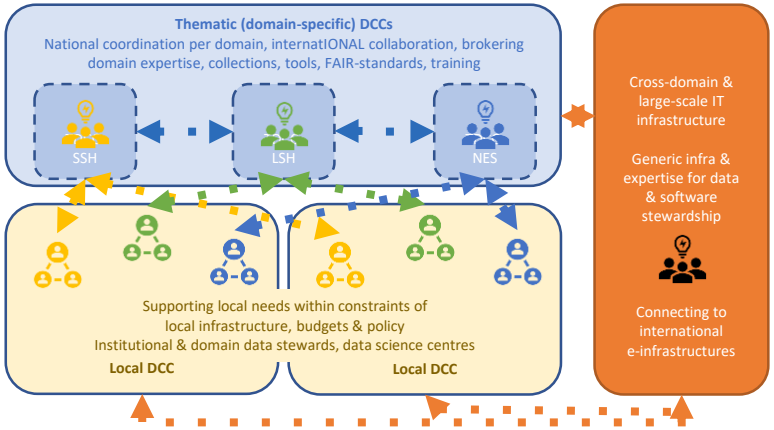

Image by: Ruben Kok, LSH TDCC and DTL

Software and infrastructure essential to bring 'dead' data to life!

Sharing expertise: thematic digital competence centres

an example from the Natural and Engineering Sciences domain

Case study: Einstein Telescope seismic studies in EUregio Meuse-Rhine in the *E-TEST* project

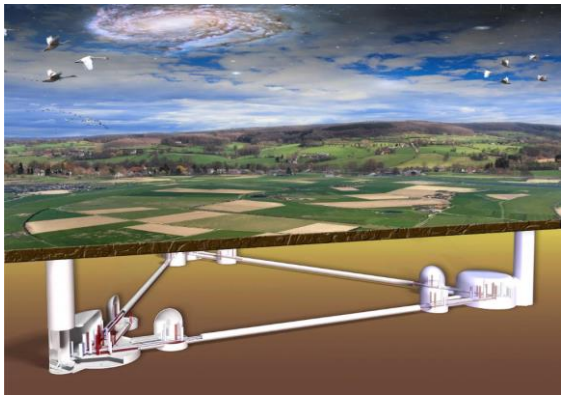

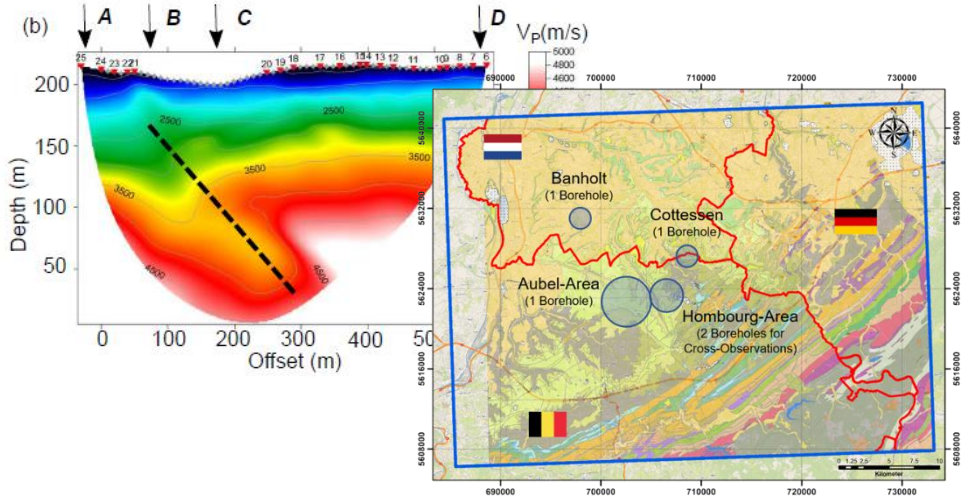

Data collected here is also useful for many other domains outside of the ET planning …

ET impression: Marco Kraan (Nikhef) from “Terziet drilling campaign” https://www.nikhef.nl/wp-content/uploads/2019/10/Terziet-Drilling-Campaign-Final-NoC.pdf
Seismic data: S Koley (VU and Nikhef) *Sensor networks to measure environmental noise at gravitational wave detector sites*, ISBN 978-94028-2054-6; map image: etest-emr.eu project site

Looking at ICT as a research instrument, like our detectors

‘ICT infrastructure for research – distinct from the office and enterprise service’

- research data today is ‘born digital’ and has joint challenges, hence: shared e-Infrastructure!

But our ‘open data’ increasingly stresses *capacity*

- data volume and bandwidth keeps growing: there is no sustainable funding model yet
- same for research software: continued maintenance eats into innovation

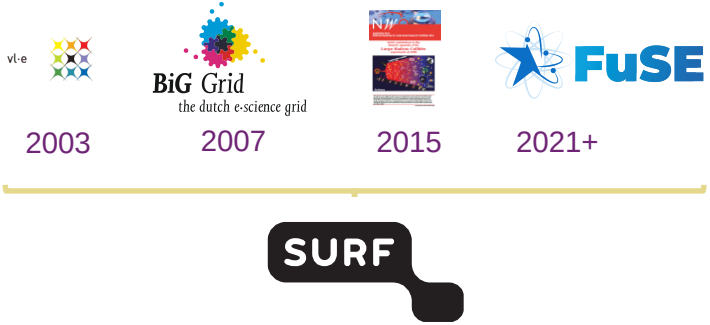

But project-based mechanisms are not solution for the sustainability issues like data & compute: a better programmatic approach needed, including long-term financial stability in the ecosystem

Riding the infrastructure innovation chain together

Computing Sciences Research

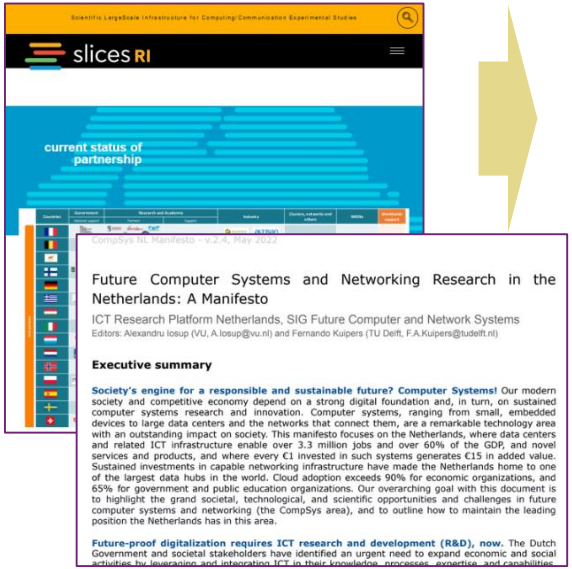

Operational Research (near-term, NextGen storage, 800G+ network, QC & simulators)

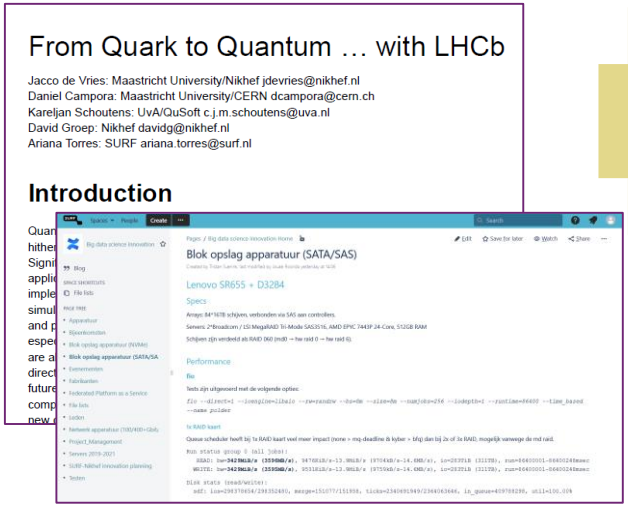

involving non-CS research domains

Operational innovation (procurement, systems vendor co-engineering)

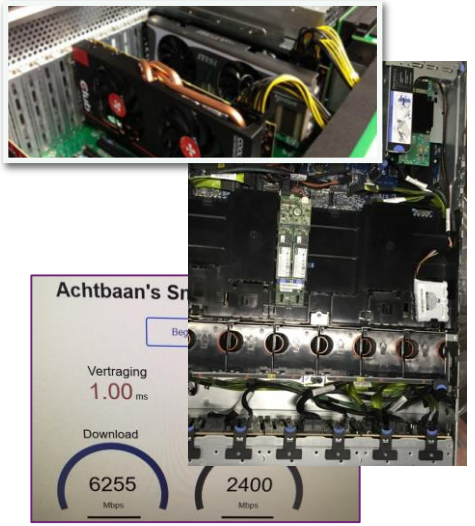

Images: SLICES-RI, CompSysNL, UM & SURF QC, SURF SOIL BDRI, NationaleSpeeltuin.nl @Nikhef

SURF Experimental Technologies Platform

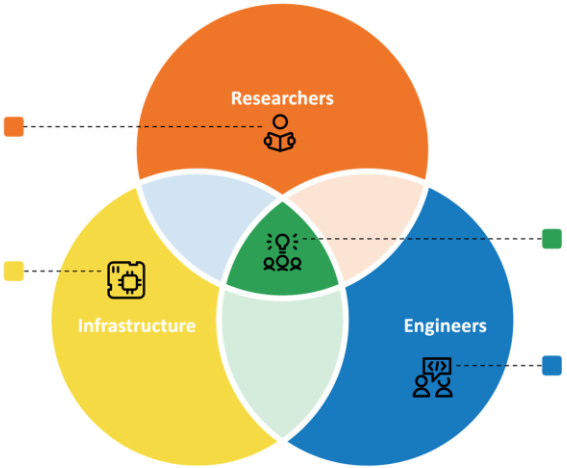

SURF–ETP
Open and collaborative environment to foster the assessment of cutting-edge technologies and methodologies.

Engineers
Leverage their technical skills and motivation to surf the state-of-the-art.

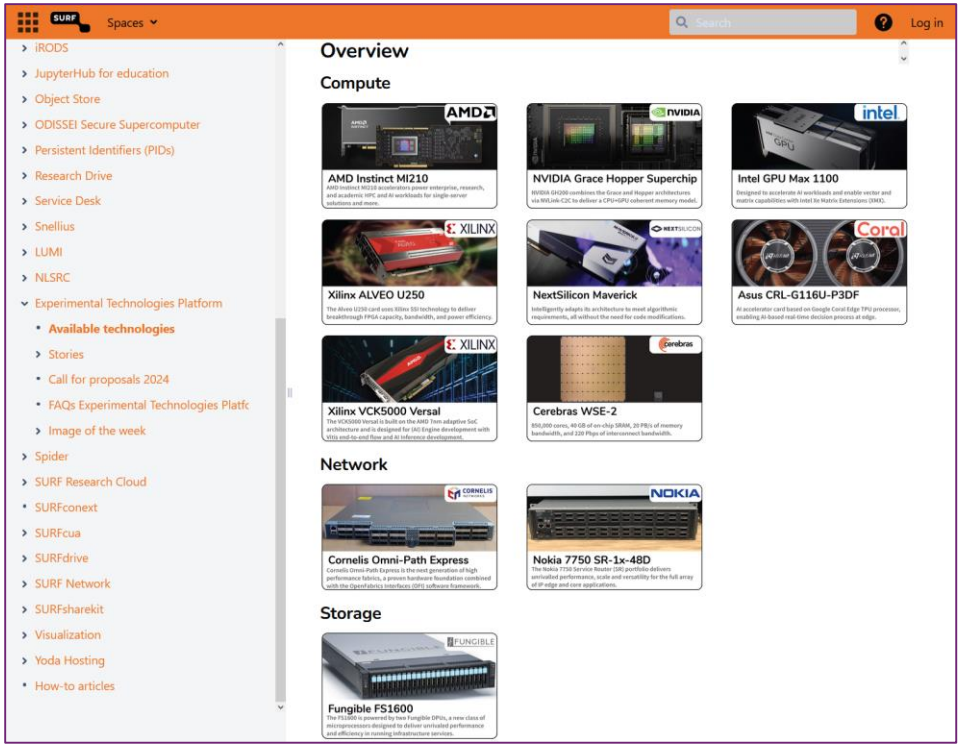

https://servicedesk.surf.nl/wiki/display/WIKI/Experimental+Technologies+Platform and https://www.surf.nl/en/etp - contact Raymond Oonk at SURF for more info

It’s all about networking – of all kinds, and globally

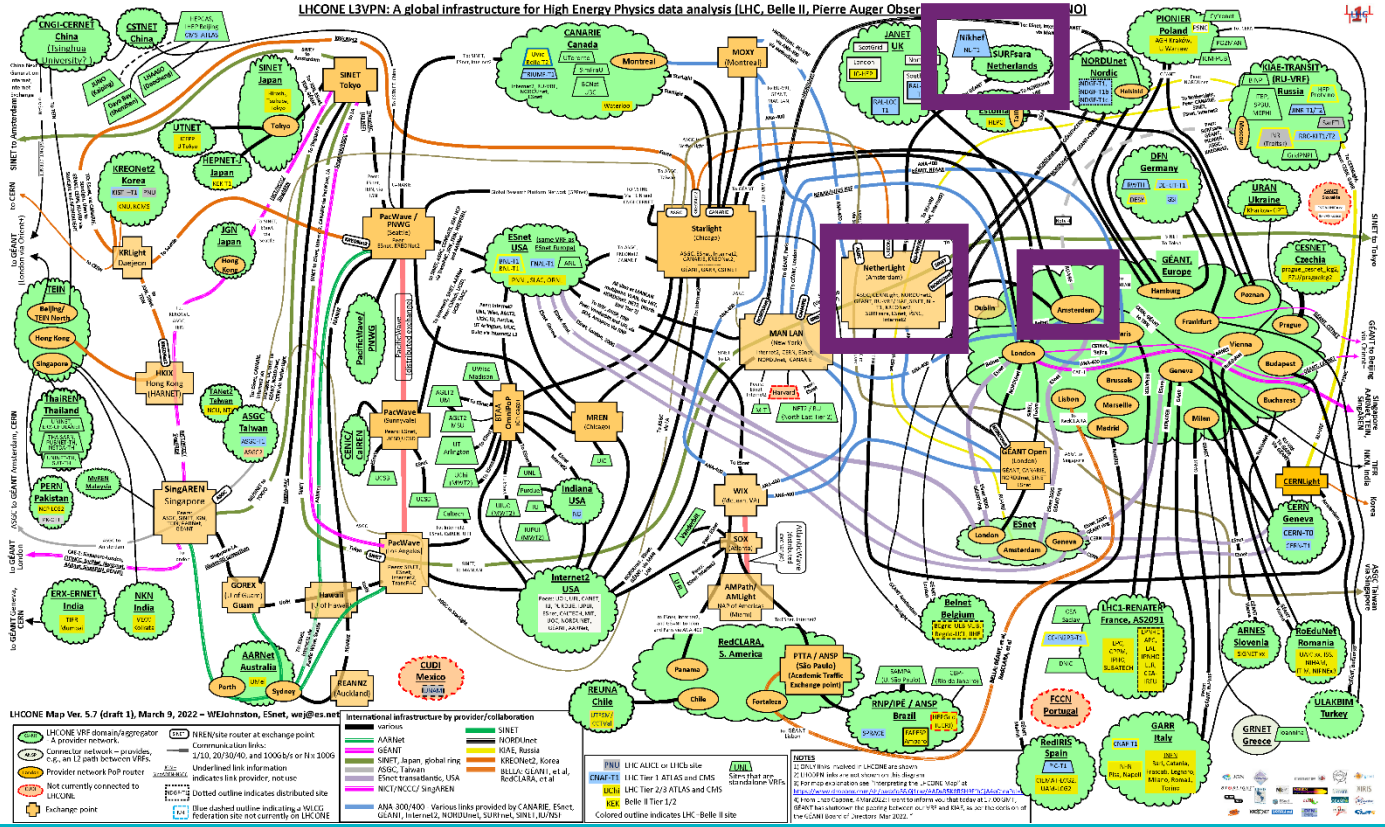

Graphic: Bill Johnston, ESnet, for LHCone

And since speed does matter ..

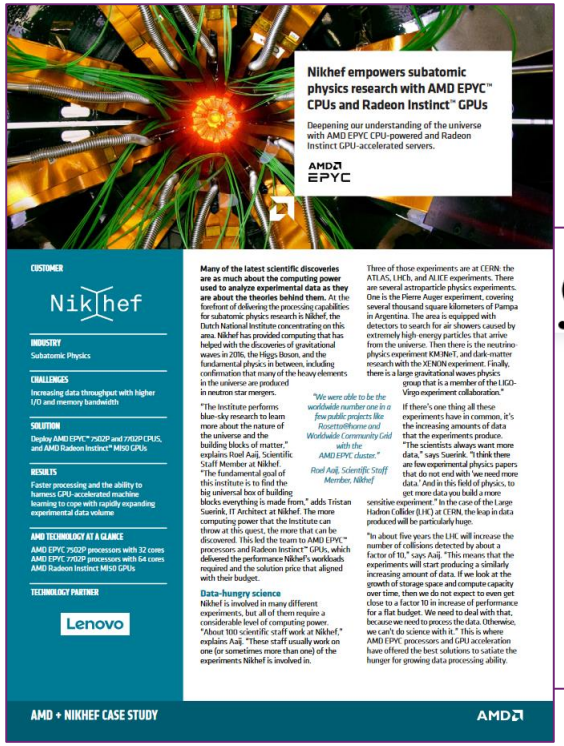

FUNGIBLE

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

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

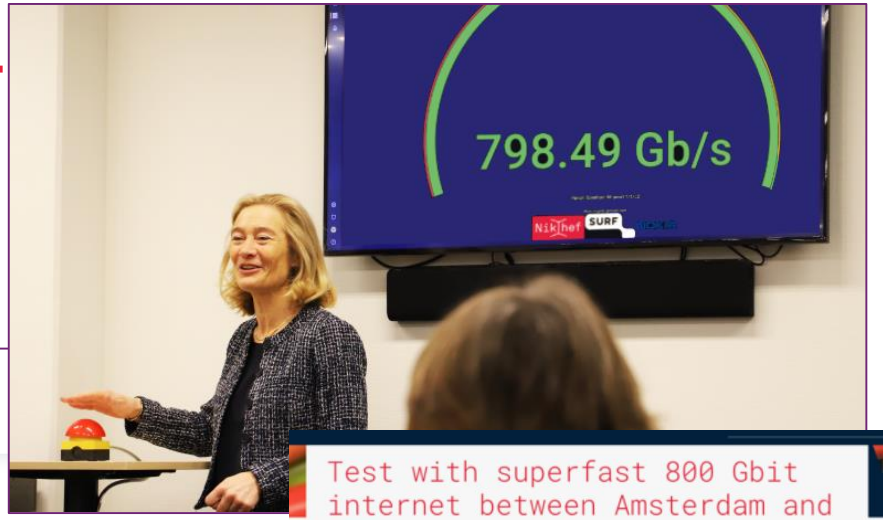

Test with superfast 800 Gbit internet between Amsterdam and CERN successful

15 April 2024

Nokia and SURF have successfully tested an 800 Gbit/s data connection between Nikhef in Amsterdam and CERN in Geneva. Such a connection is needed to transmit data from the upcoming high-luminosity LHC accelerator.

The test used existing fiber-optic connections through Belgium and France toward Geneva in Switzerland over a total distance of 1,648 kilometers. An 800 Gbit/s connection is about a thousand times faster than the Internet connection in an average household.

Nokia's latest photonic technology, the sixth-generation super-coherent Photonic Service Engine (SPE-6s), was deployed in the tests, along with 16QUM-shaped modulation. The results of the tests will be announced in more detail next week at a Nokia expert conference in Athens.

Data hub

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

Our science data flows are somebody else's DDoS attack

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-profitorganisaties en banken. Nadat duidelijk w… simpele DD… n platgelegd, H… vapenen tegen… alitie. Onlangs…

Op dez…
Een goe…
Examen…
Wat gebeurde er die nacht?
Vragen over dit artikel?
Terug naar boven

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

Door Rutger Otto

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

202 reacties

Het 'red team' is verantwoordelijk voor de aanvallen, het 'blue team' voor de verdediging. Een van de partijen die aan de avond meedoet is Nikhef. Tristan, IT architect bij Nikhef, geeft aan "dat zij dit

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.

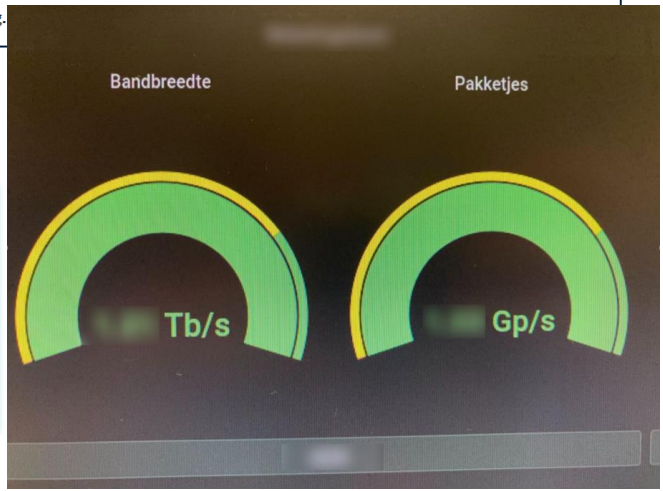

Image sources: belastingdienst.nl, rws.nl, nu.nl, werkentegennederland.nl

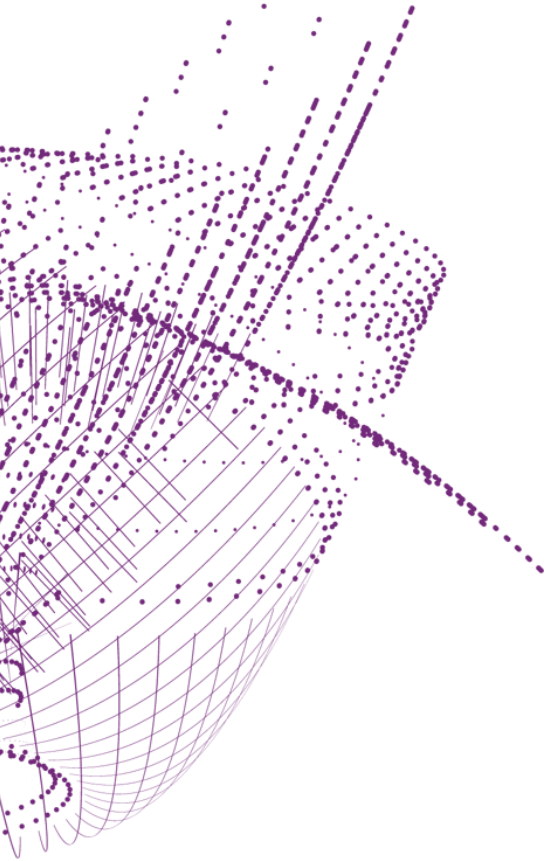

A collaborating ‘ecosystem’
for science and innovation

How did we get here?

Once upon a time … a green fields approach?

Watergraafsmeer and volkstuinencomplex Frankendael, 1974, starting the construction of the WCW. Image: Beeldbank Amsterdam, gemeente Amsterdam

The Nikhef data centre – at the end of the 1980s

Nikhef room H1.37 – terminal stations on the raised data floor of the computer room (H1.40, behind the glass-panel walls)

Gould, Sun, and DEC systems, taking several racks each

- 500 m2 floor area
- Raised floor: +60cm
- walls are ‘movable’ to accommodate expansion

Collaborative Research Infrastructures: all about networking!

be it human or computer networks …

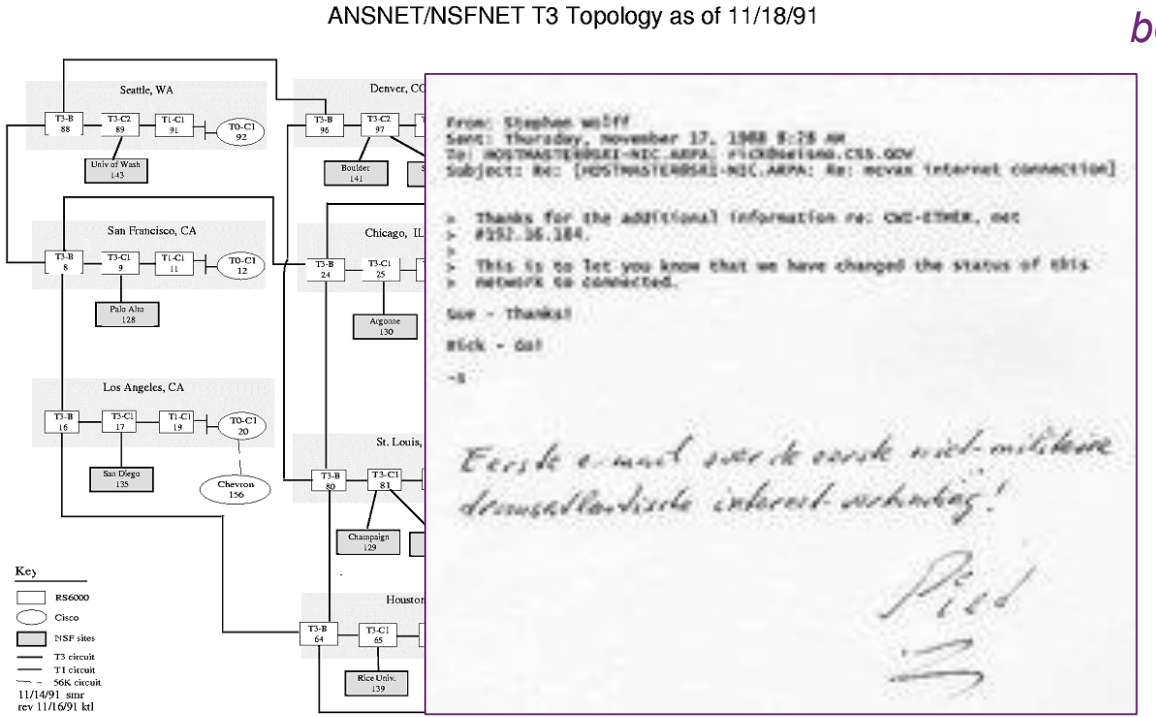

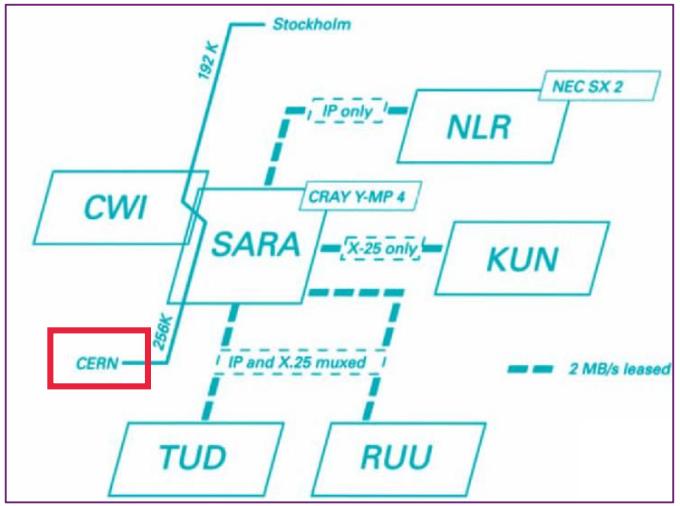

See https://personalpages.manchester.ac.uk/staff/m.dodge/cybergeography/atlas/historical.html for more historic maps ; right-hand image: SURFnet2, 1990
first email to MCVAX at CWI from https://www.cwi.nl/en/news/cwi-celebrates-25-years-of-open-internet-in-europe-in-november/ (Piet Beertema, CWI, 1988)

‘IBR-LAN’ at Nikhef – connecting local and global networks

International Backbone Router Local Area Network “IBR-LAN” at Nikhef, room H1.40 as seen in 1996. Right: H1.39 with nikhefh.nikhef.nl racks and early DAS-2 system

What happens if you welcome two networks onto your floor

Image: AmsIX at Nikhef H140 in 2007 – foto Beeldbank Amsterdam https://archief.amsterdam/beeldbank/detail/a95bc475-8fcc-d0d1-37f9-b077ba3729db/media/

A growing internet!

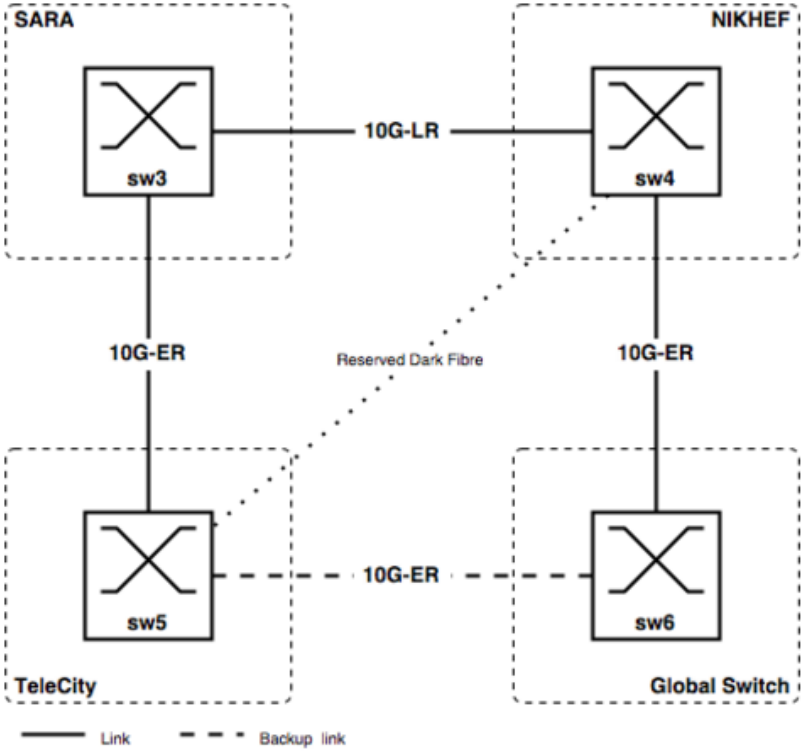

AMS-IX topology, 2002

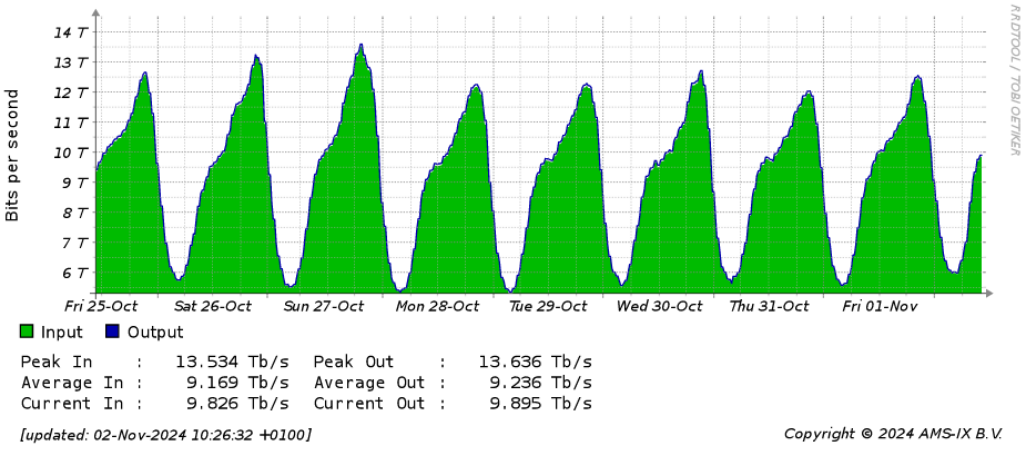

AMS-IX traffic Nov 2, 2024; https://stats.ams-ix.net/ https://www.ams-ix.net/ams/colocations

Connect each other … and scientific data and instruments

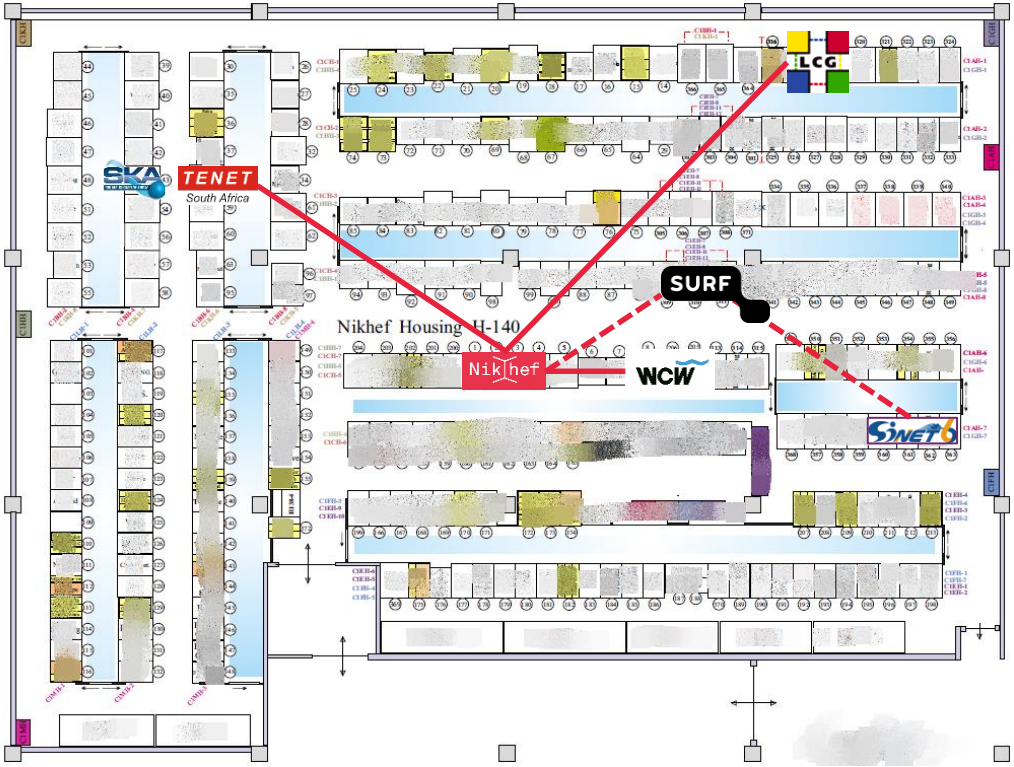

dashed lines: traffic is routed via ancillary facilities of SURF and GEANT first (currently: AMS9 at WCW) SKA Meerkat traffic via TENET

What happens inside a data centre …

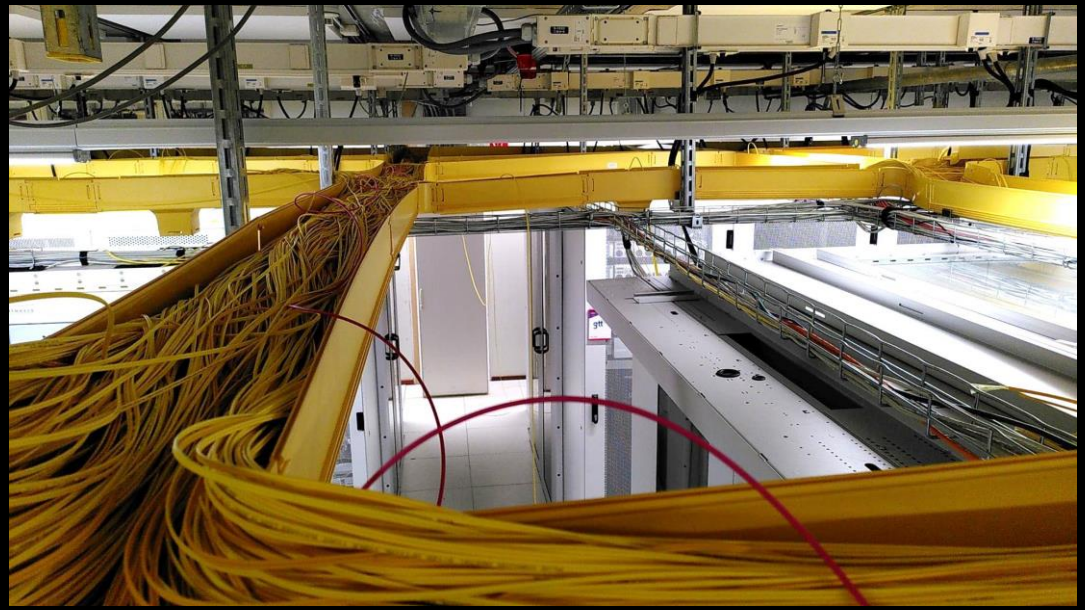

‘Connectivity’ housing and ‘hosting’ are different things:

- NikhefHousing (H140) has connectivity parties only, and does not host any content
- what you see on the 1st floor tour is *network* equipment: shipping data, but not keeping anything

2nd floor has our science data centre

And no single connectivity data centre is a single point of failure:
Internet protocols are engineered to re-route traffic

Today’s data centre at Nikhef

‘**NikhefHousing**’ data centre
- from the first 2 racks in a corner
- to now > ~400 racks
- many different connectivity parties
- **376** networks present in PeeringDB
- connectivity-focus, not hosting

Nikhef ‘science’ data centre H234b
- 47 racks and ~350 kW
- hosts Nikhef, CERN, gravitational waves, and SURF *research* data
- strengthens connectivity at, and uses NikhefHousing

Data centre installation management, ever growing

- active/free cooling chillers installed on the roof in 2009
- data floor: ~400 racks
- evolving hot-isle/cold-isle configuration
- electricity generator sets 2003, 2009, 2021
- aquifer thermal energy storage (ATES) system installed 2010

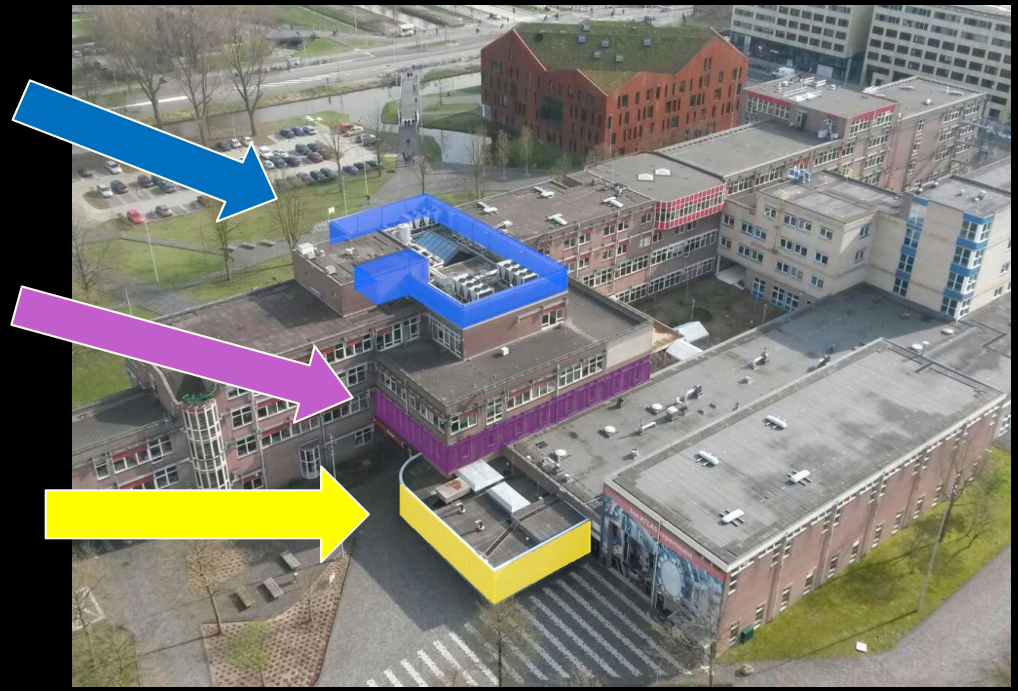

Image: Floris Bieshaar, Nikhef

Power in … and power out …

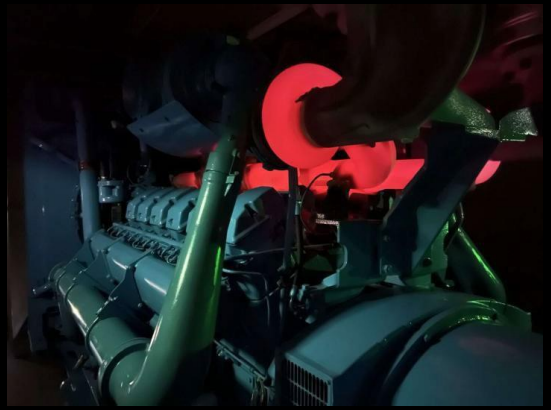

Three generators

- A-Feed 1250 kVA (pictured under load while testing)
- B-Feed 1700 kVA
- C-Feed 1250 KVA added with the current expansion

Separate redundant UPS for each

Heat re-use: aquifer thermal energy storage

re-use heat to warm our building (pretty warm) AND feed more heat to student housing opposite
nominal 'PUE' ~ 1.21

Generator image source: Floris Bieshaar. MacGillevrylaan sketch: Science Park Amsterdam

Because, even if we can …

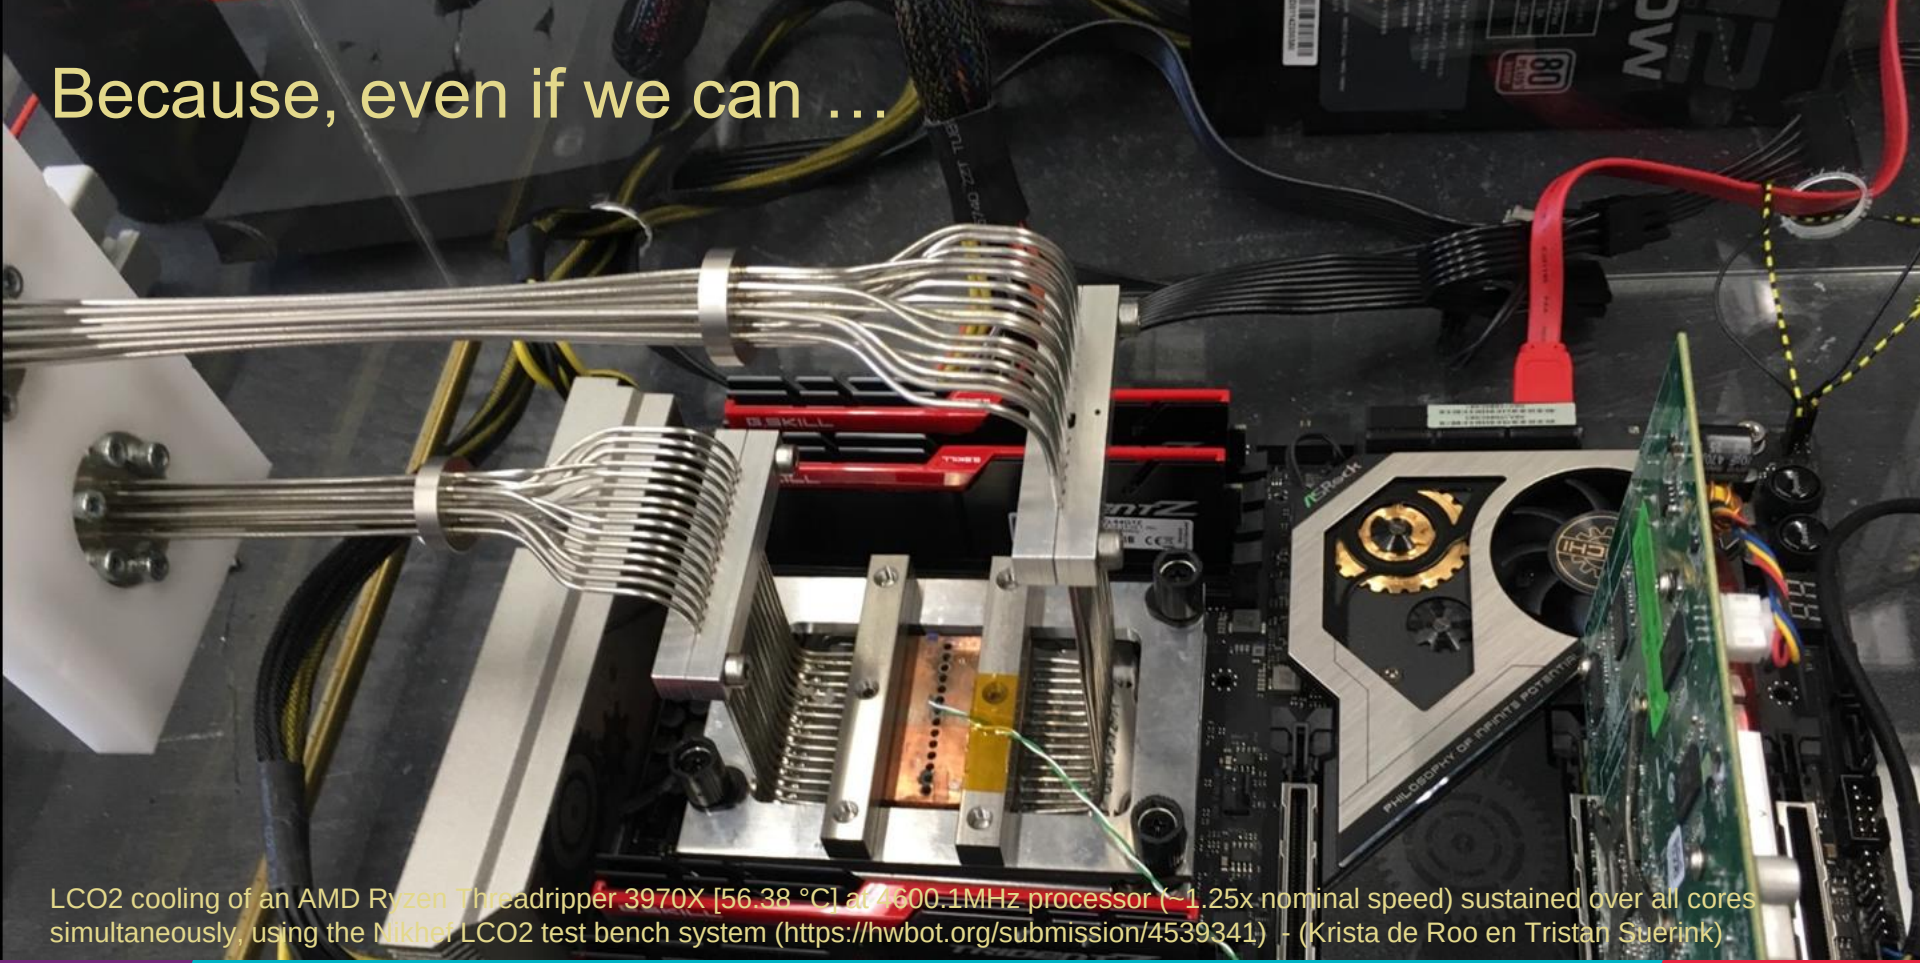

LCO2 cooling of an AMD Ryzen Threadripper 3970X [56.38 °C] at 4600.1MHz processor (~1.25x nominal speed) sustained over all cores simultaneously, using the Nikhef LCO2 test bench system (https://hwbot.org/submission/4539341) - (Krista de Roo en Tristan Suerink)

… it is not always the most scalable solution!

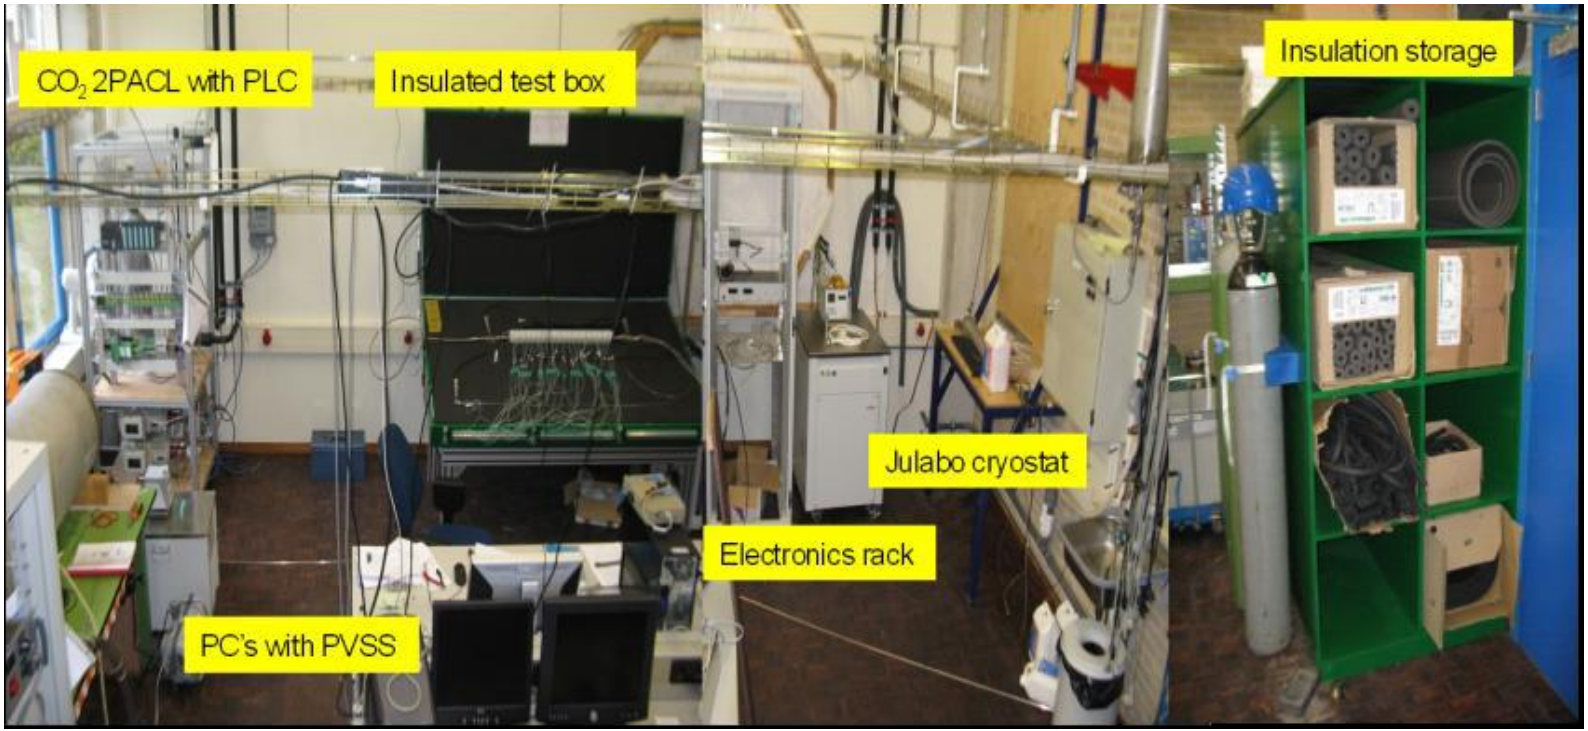

Nikhef 2PA LCO2 cooling setup. Image from Bart Verlaat, Auke-Pieter Colijn *CO2 Cooling Developments for HEP Detectors* https://doi.org/10.22323/1.095.0031

Thanks, and enjoy Nikhef and the Amsterdam Science Park

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

Maastricht University

Nikhef

Tours

1. Data Centre: assemble here on Nikhef ground-floor

2. KM3NeT BOL Assembly: Ronald Bruijn, from the NikhefHousing entrance

Return here for the
Introduction to the local ecosystem of the Amsterdam Science Park

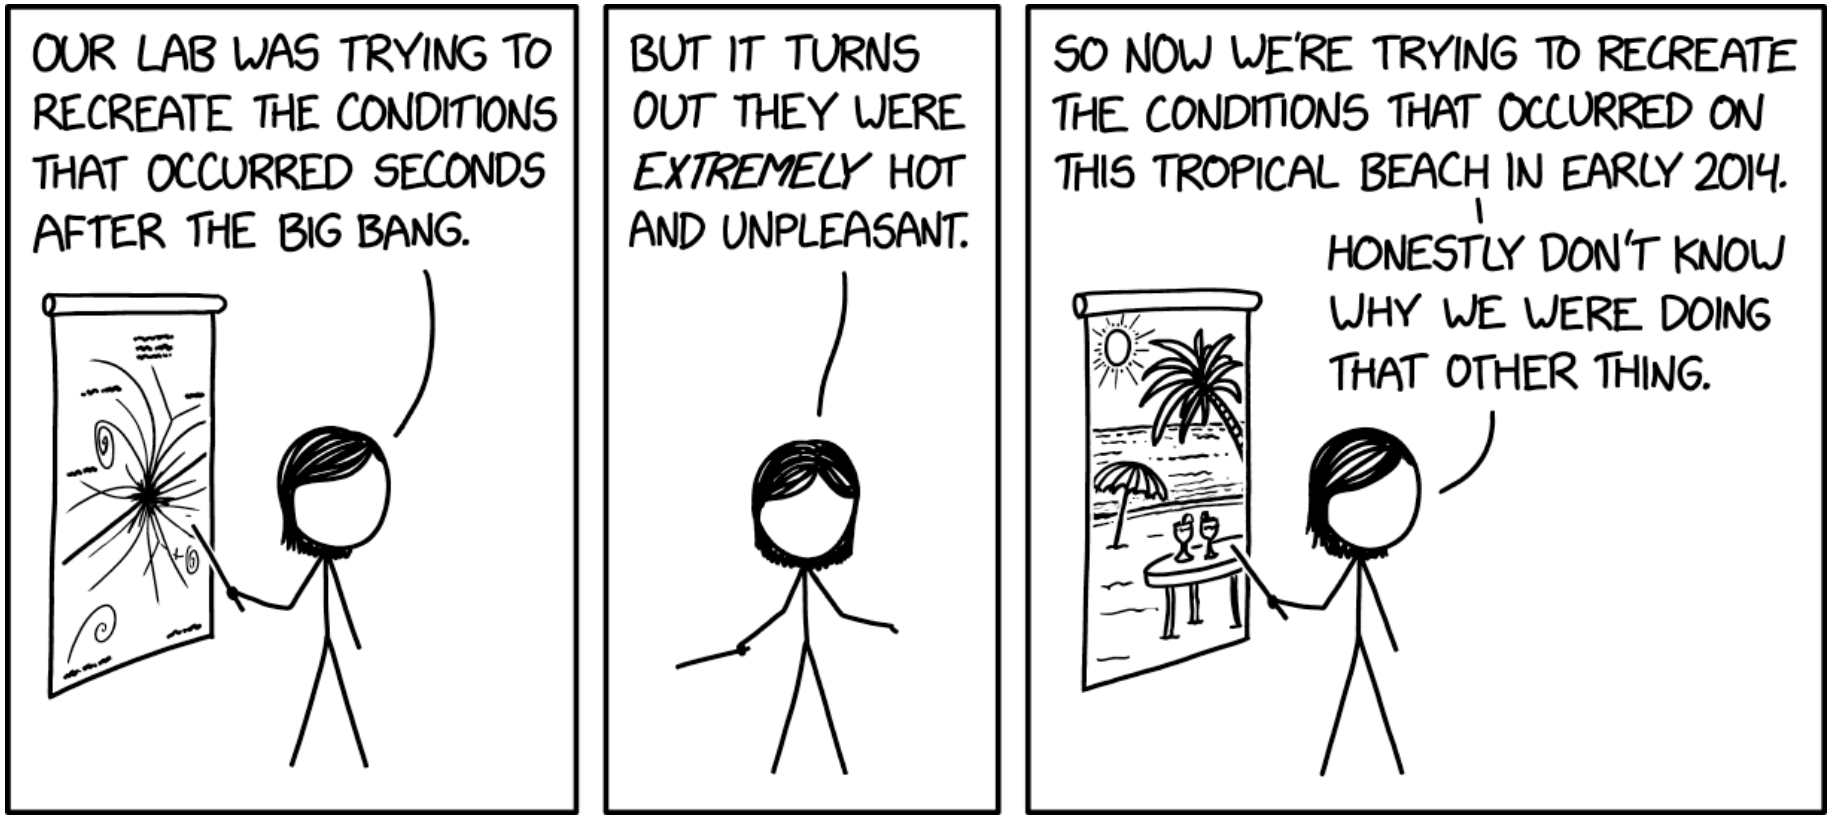

https://xkcd.com/2511/