Nikhef

Maastricht University

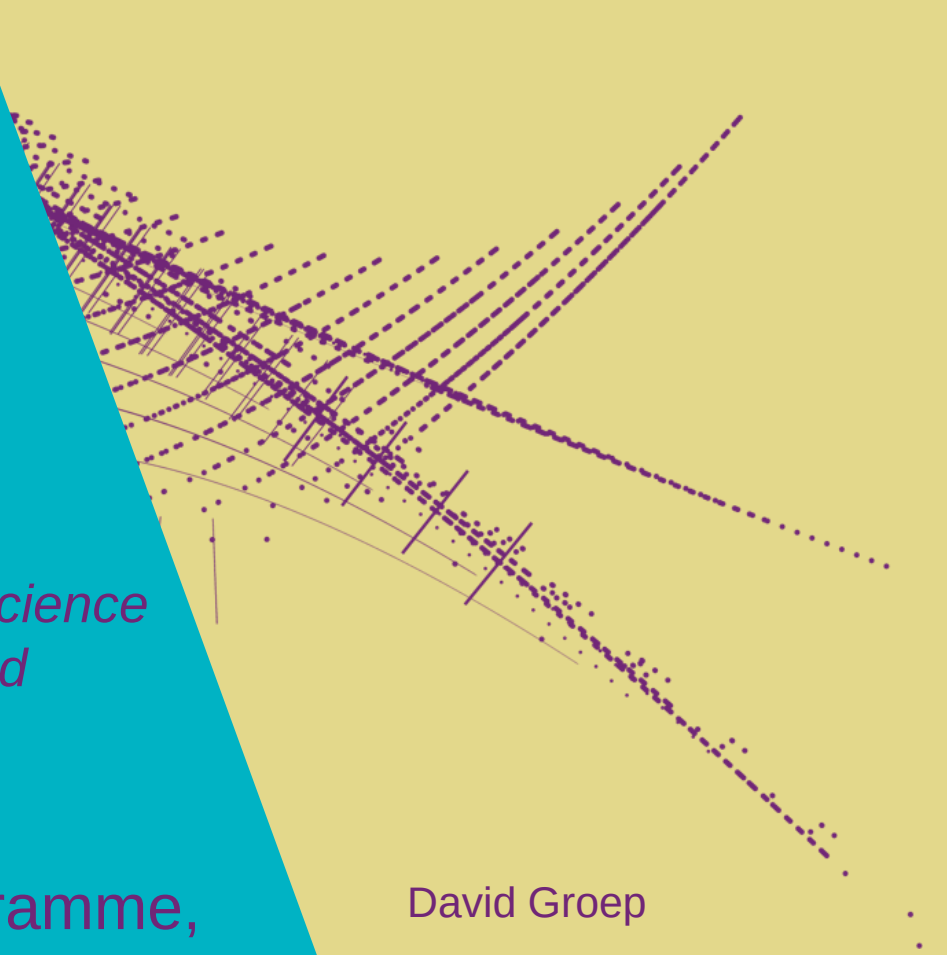

Federated computing ecosystems for science
Connecting the Netherlands to the world

Nikhef, its science and ICT programme, and the NikhefHousing data centre

David Groep

2025

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

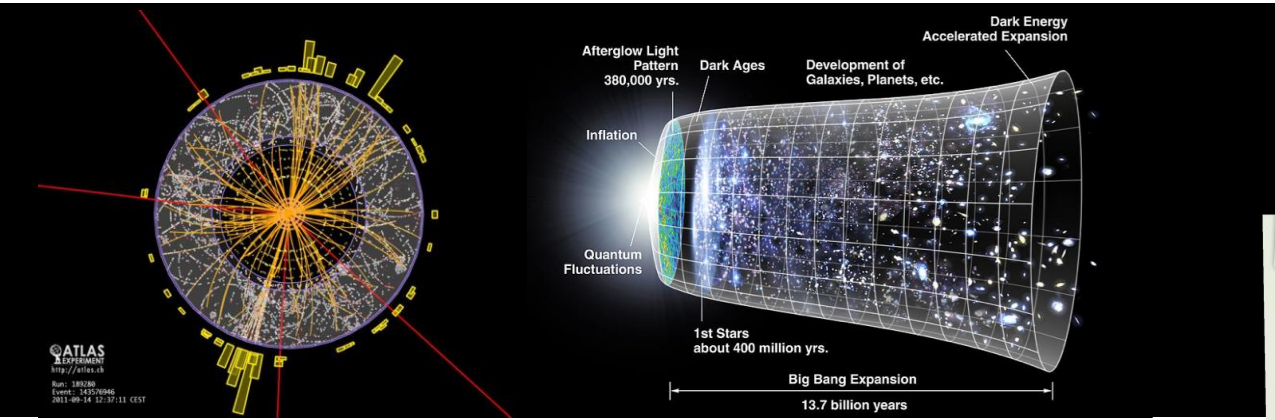

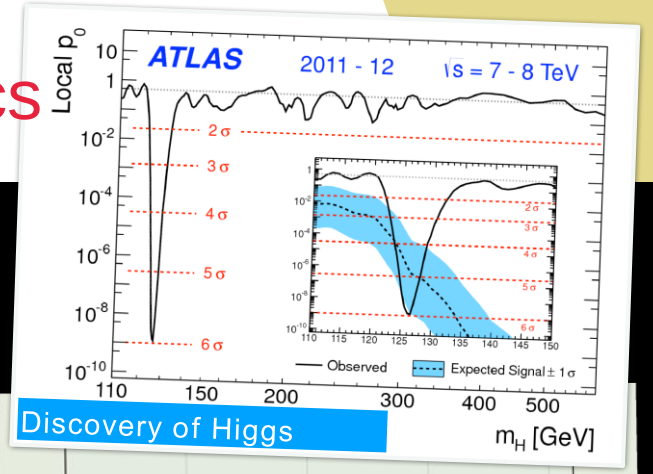

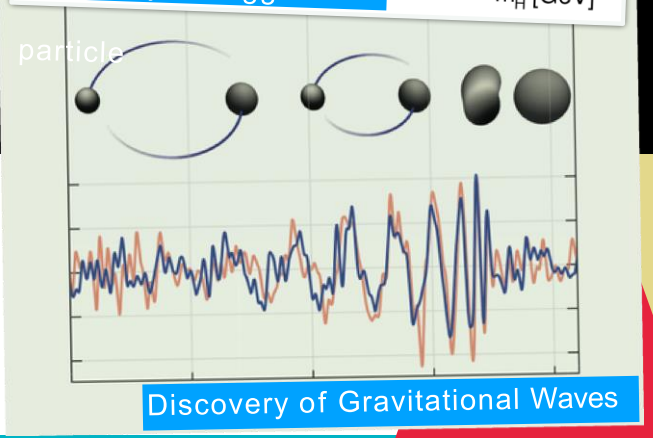

We probe our world, made of particles and fields,

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

Nikhef, its science and ICT programme, and the NikhefHousing data center

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

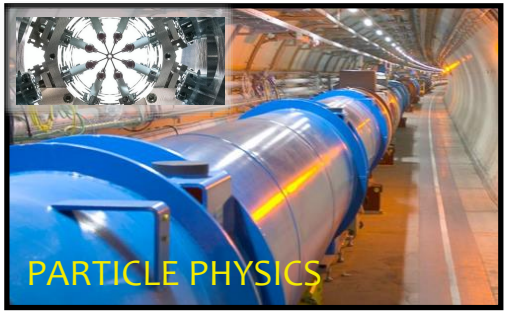

- Atlas
- LHCb
- Alice
- eEDM

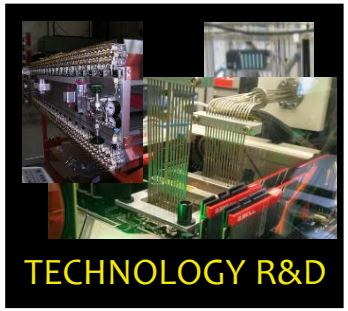

- Detector R&D
- Physics Data Processing

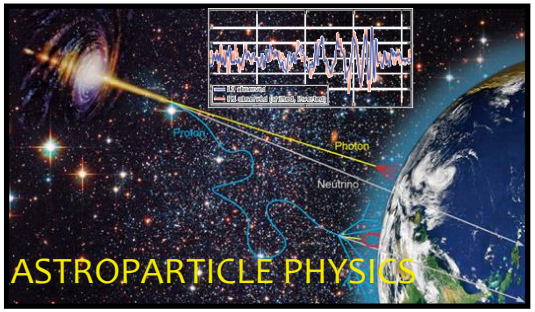

- Neutrinos
- Gravitational waves
- Cosmic Rays
- Dark Matter

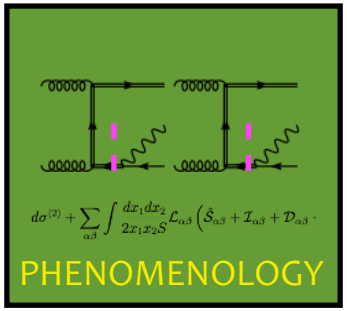

- Theoretical Physics

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

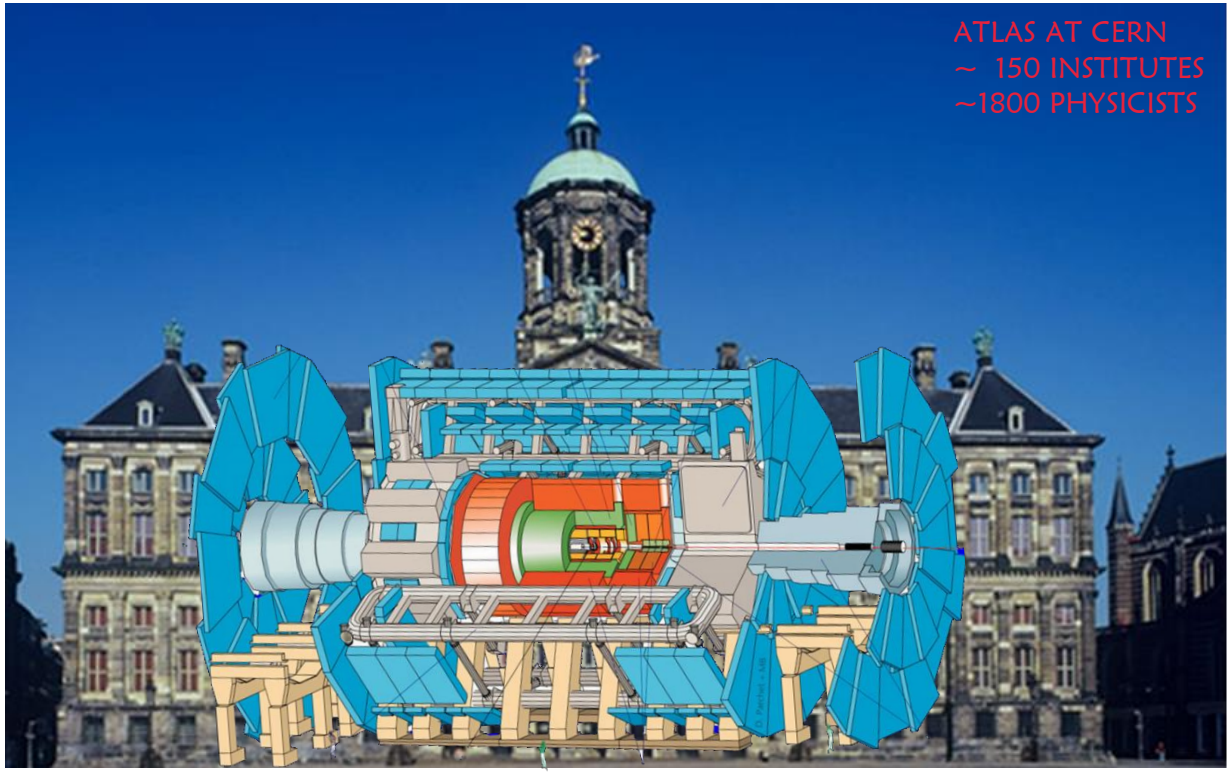

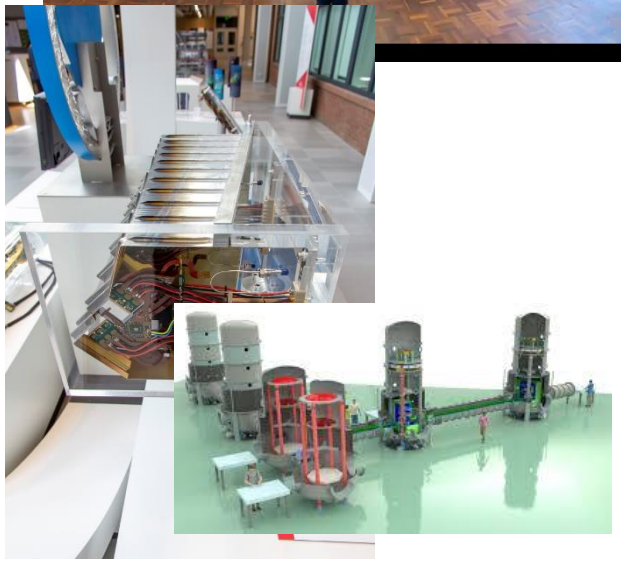

Slide materials from: Stan Bentvelsen, 2016, ATLAS collaboration, CERN; LHCb VELO and RF box at Nikhef and ET Pathfinder visualisation: Marco Kraan, Nikhef

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

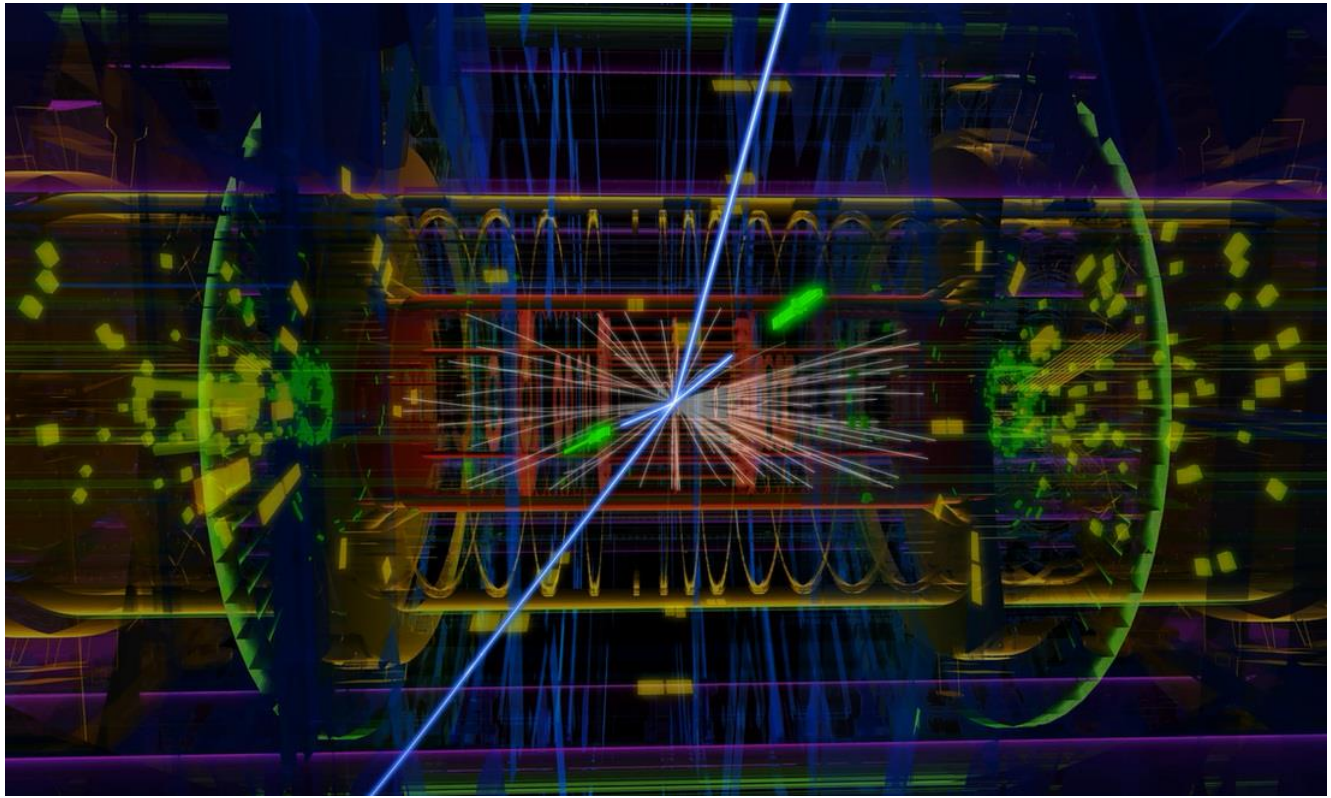

Image source: CERN, Atlas collaboration; ATLAS approx. 140M active channels, mean event size ROD 1.6 MByte

And decelerate for ultra-high-precision measurements

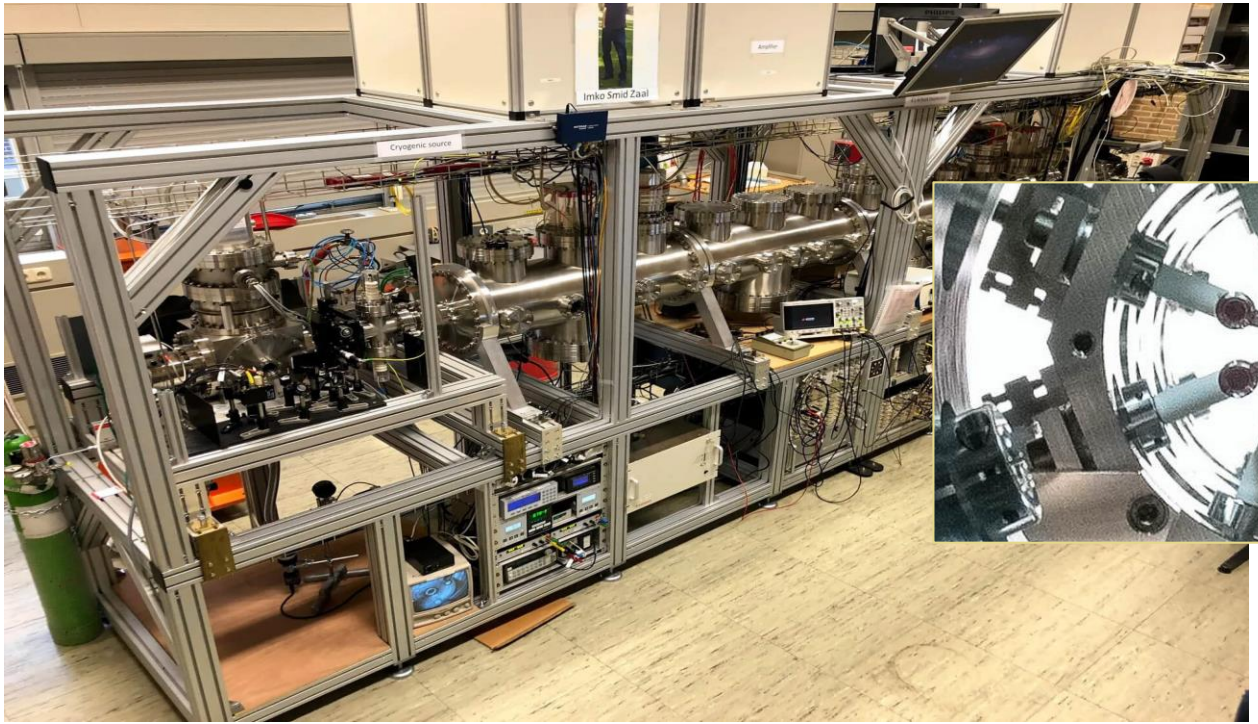

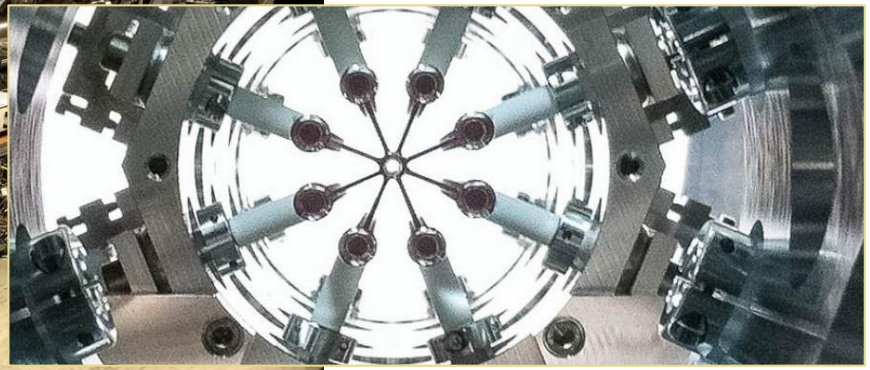

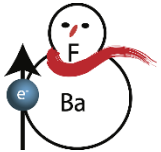

https://www.eedm.nl/ Nikhef and Rijksuniversiteit Groningen, images: Steven Hoekstra

Particles at Extreme Energies

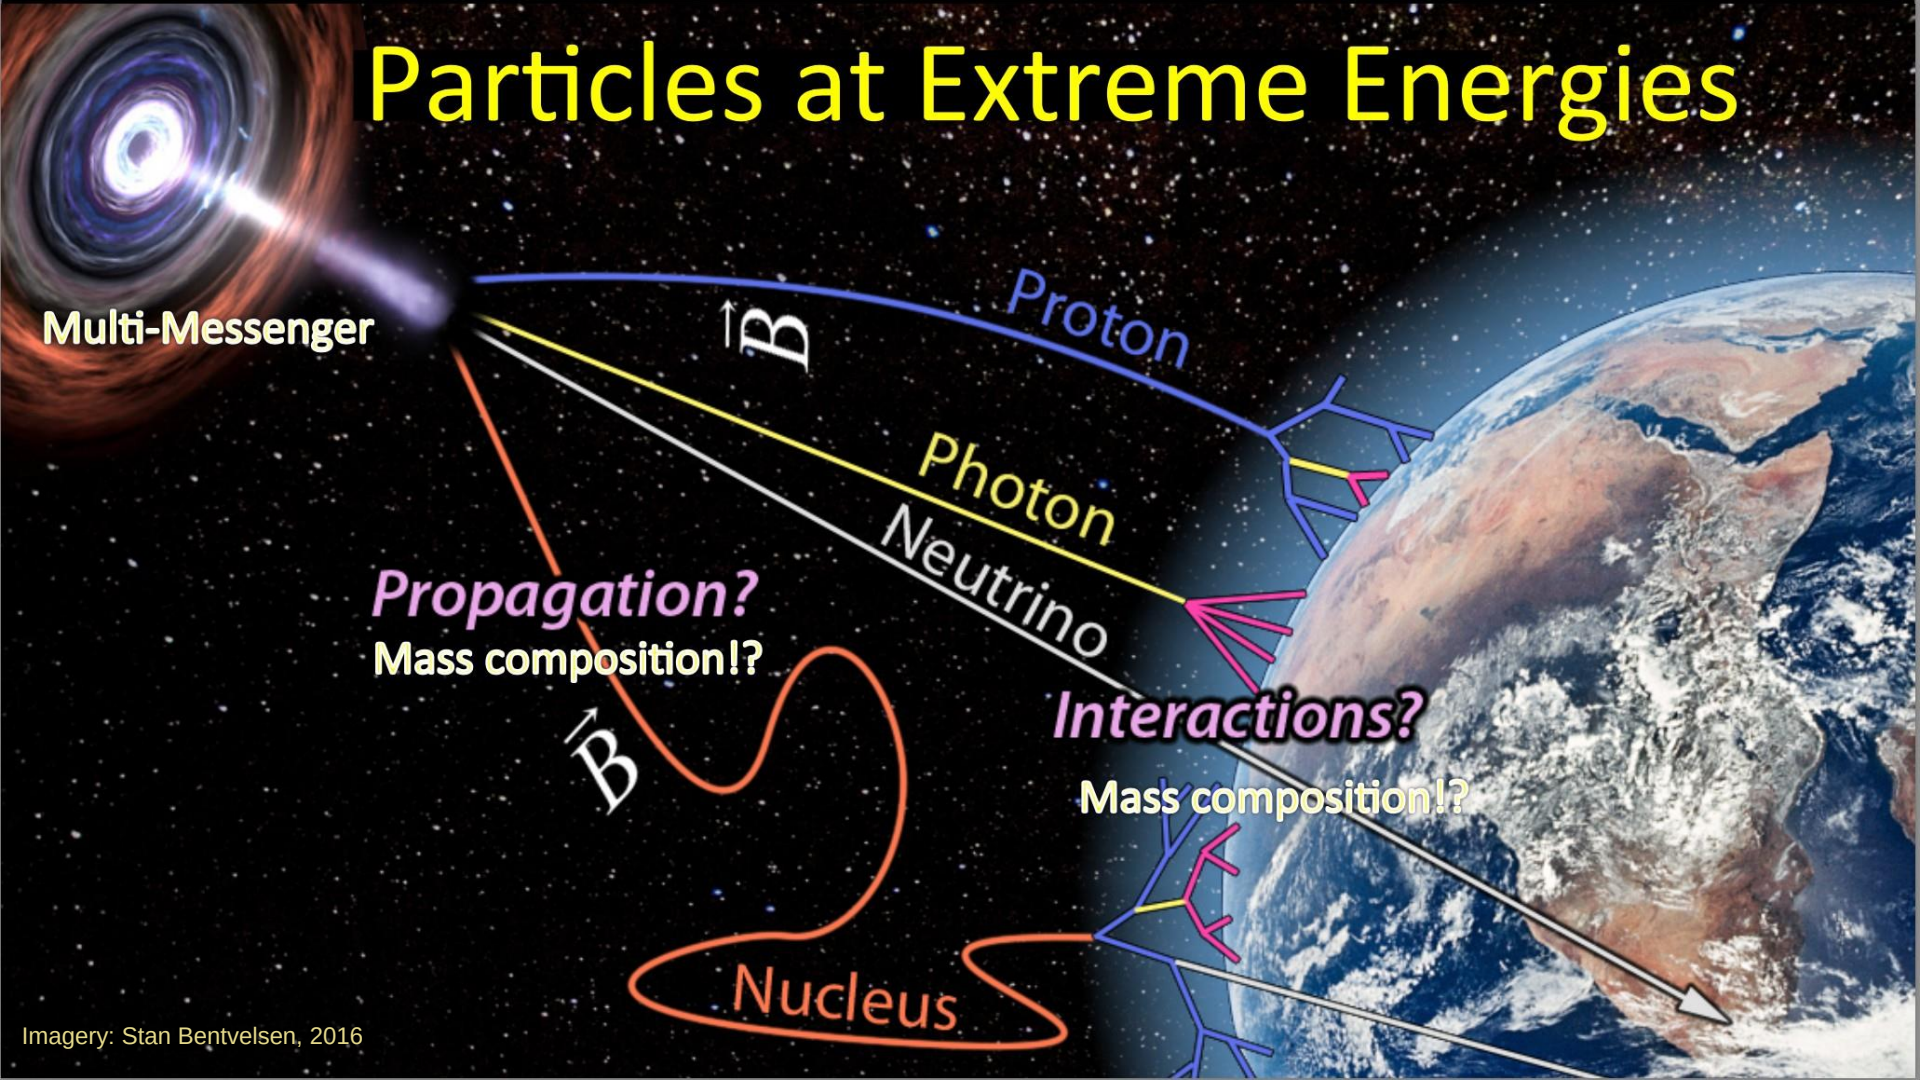

Imagery: Stan Bentvelsen, 2016

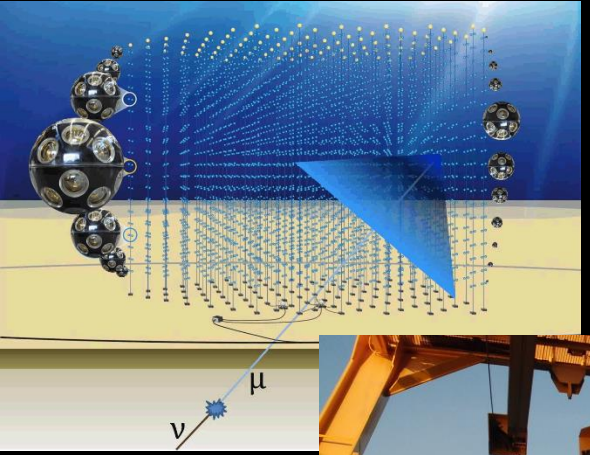

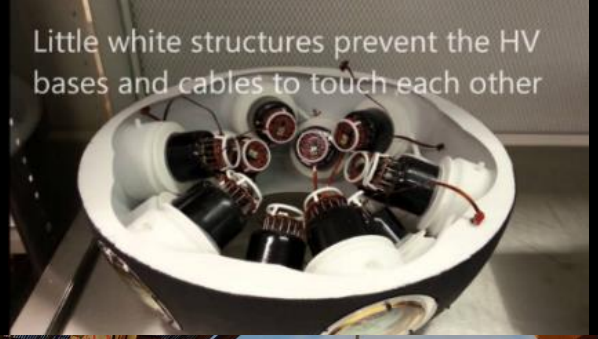

Image sources:
Nikhef, NIOZ,
KM3NET collaboration,
NRC Handelsblad

www.et-emr.eu

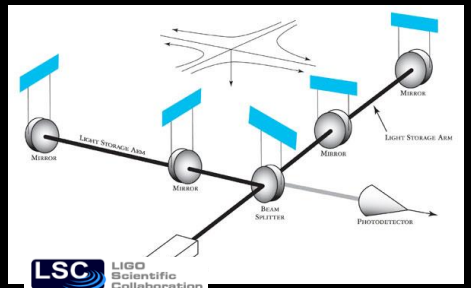

Einstein Telescope projected in the **Euregio Maas Rijn**, images: Marco Kraan;
ET Pathfinder at UM, Maastricht, NL

GW150914 event:
gw-astronomy collaborations, LIGO

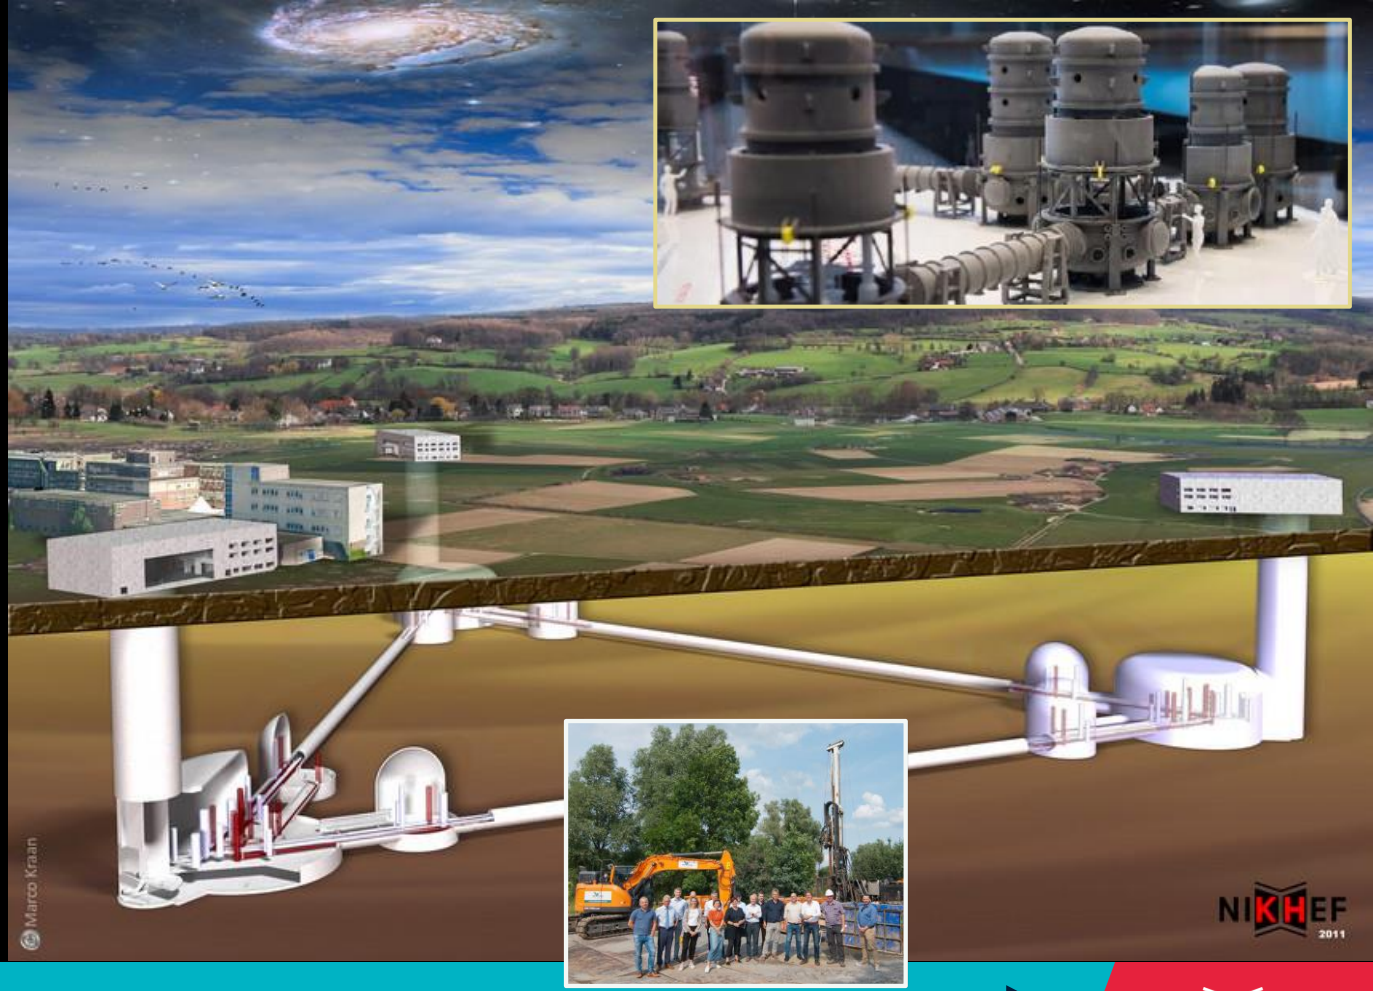

Enabling Research Programmes – next gen infrastructure

Detector R&D

Theoretical Physics

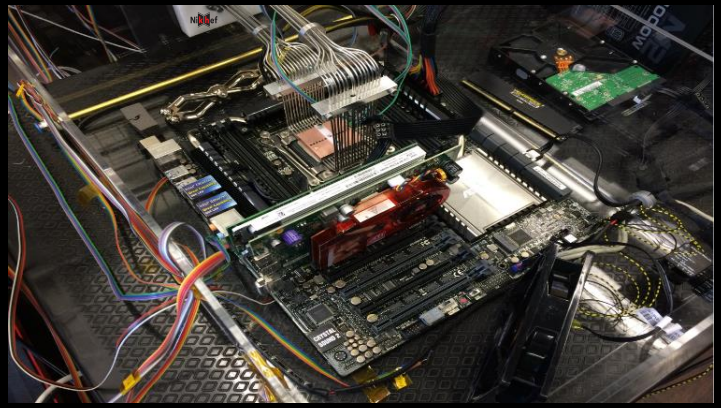

Physics Data Processing

Infrastructure research for research infrastructures at Nikhef

Algorithmic design patterns and software

- scientific software (GPU) acceleration, new algorithms, high-performance processors
- software design patterns for workflow & data orchestration, and (energy) efficiency

Infrastructure, network & systems co-design R&D

- building ‘research IT facilities’
- co-design & development
- big data science innovation
- research *on* IT infrastructure

Infrastructure for trusted collaboration

- trust and identity for enabling communities
- managing complexity of collaboration mechanisms
- securing the infrastructure of our open science cloud

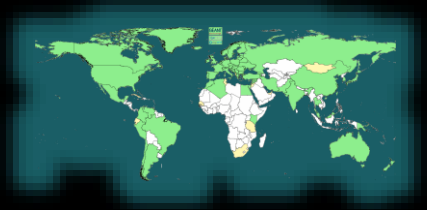

Algorithms and detectors go hand in hand

Reducing complexity is both in hardware and algorithms

If we ‘do nothing’, the HL-LHC intensity will result in hundreds of overlapping ‘images’ that are impossible to disentangle:

- improve timing resolution to 10-50ps (more ‘frames’ per collision)
- accelerated algorithms and dedicated GPU kernels, also ‘off-line’

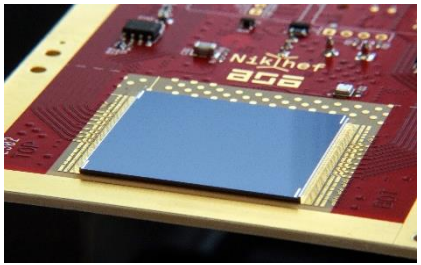

From: FASTER (LHCb images) and R.Geertsema LHCP2022; right: TimePIX4 (also used at e.g. M4i)

Data Processing infrastructure – the Dutch SURF example

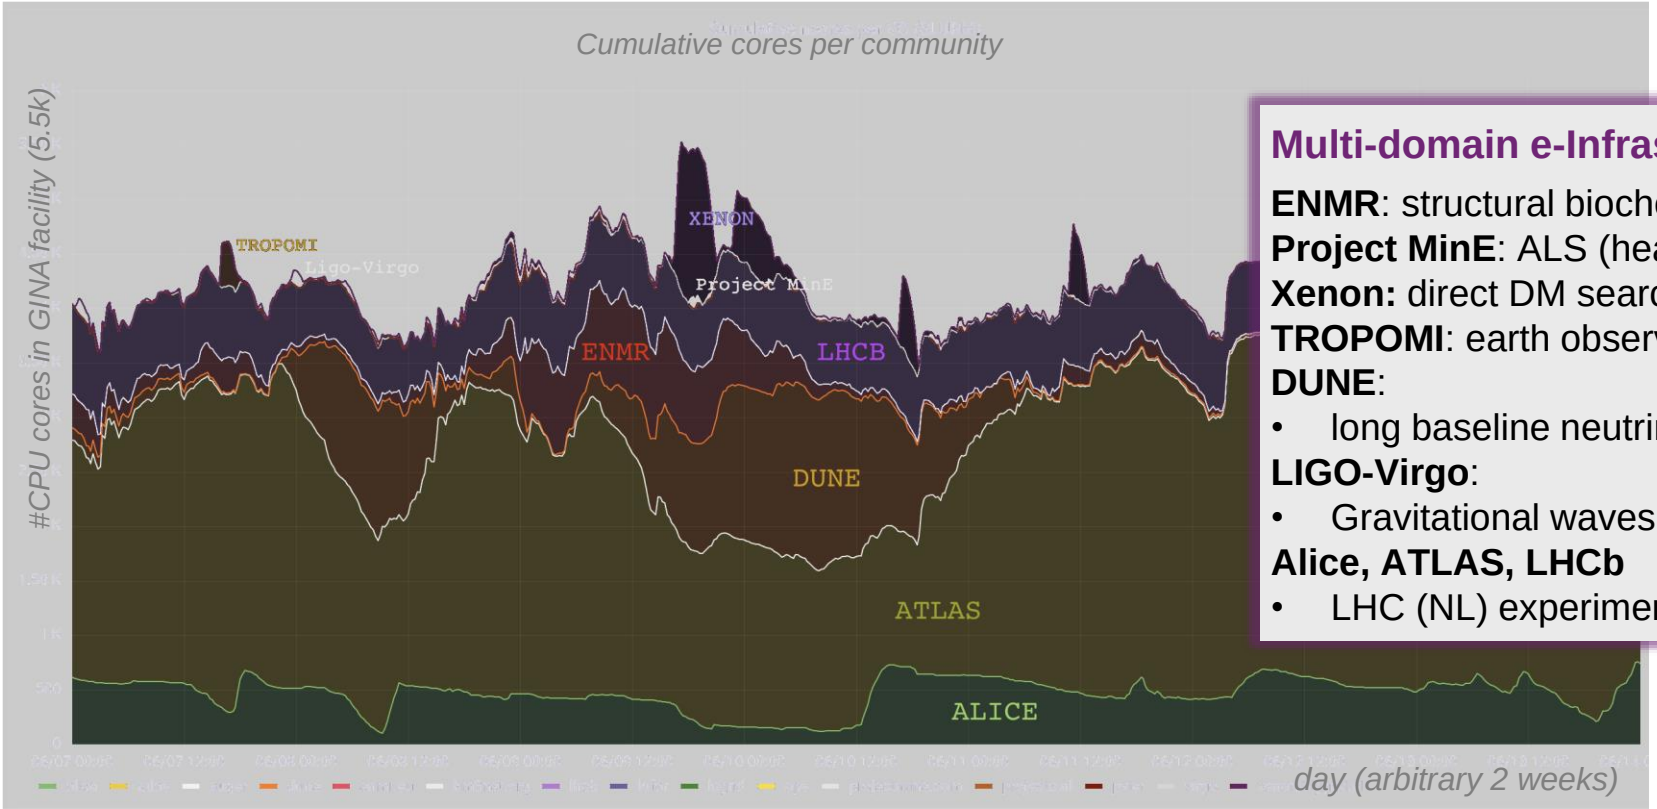

Multi-domain e-Infrastructure

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

Riding the infrastructure innovation chain … together!

Computing Sciences Research

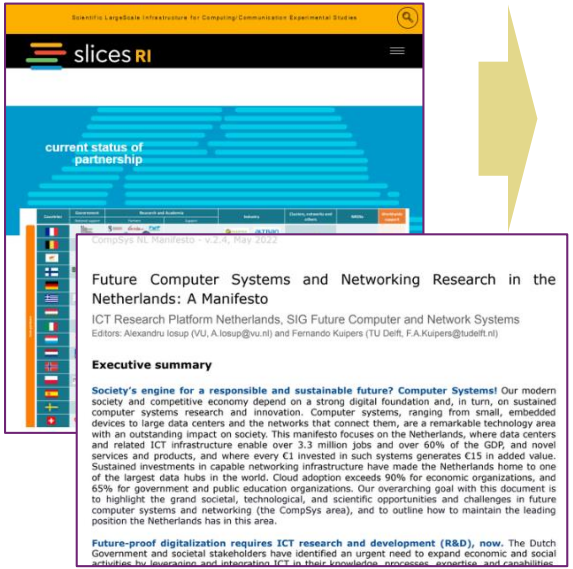

Operational Research (near-term, NextGen storage, 800G+ network, QC & simulators)

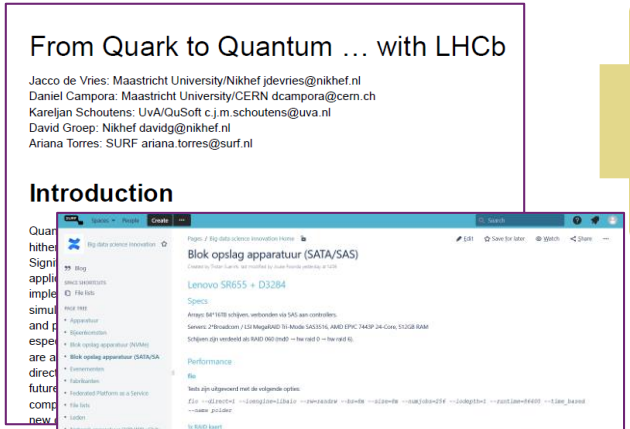

involving non-CS research domains

Operational innovation (procurement, systems vendor co-engineering)

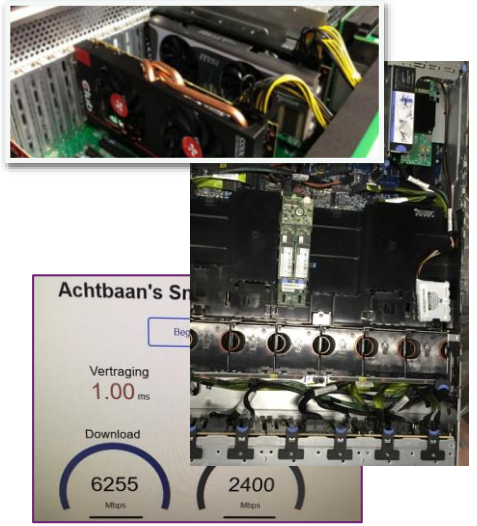

Images: SLICES-RI, CompSysNL, UM & SURF QC, SURF SOIL BDRI, CentraleSpeeltuin.nl @Nikhef

SURF Experimental Technologies Platform

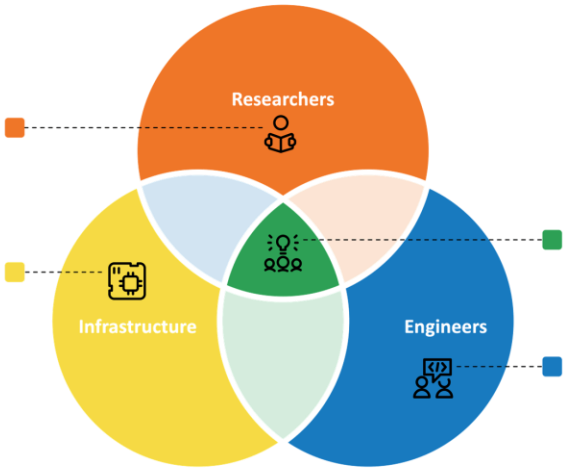

SURF–ETP
Open and collaborative environment to foster the assessment of cutting-edge technologies and methodologies.

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

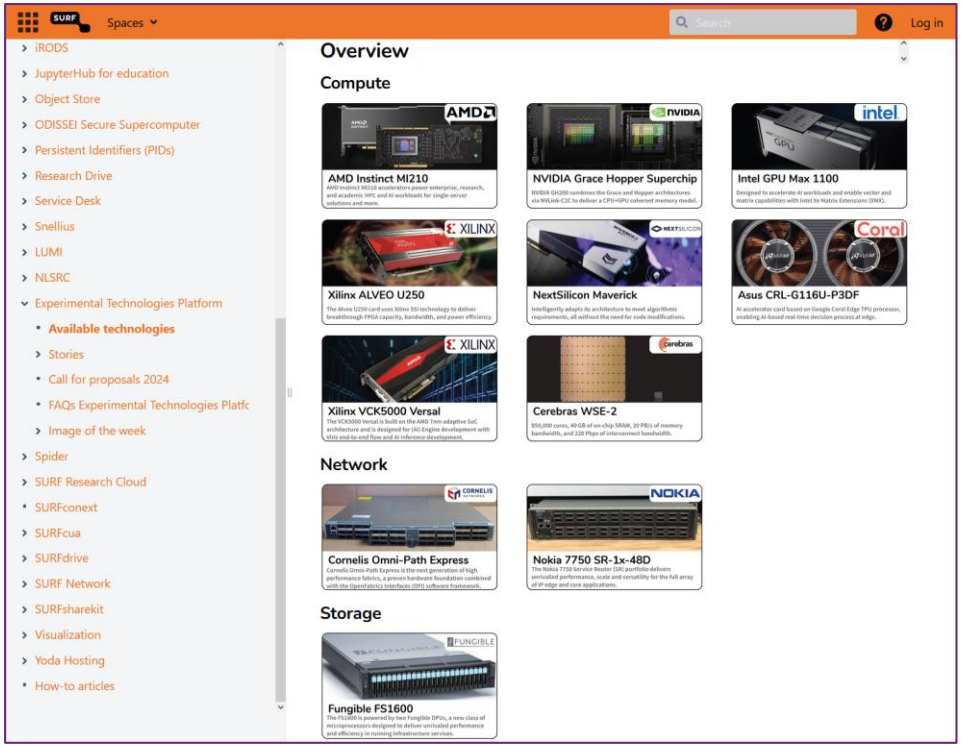

https://servicedesk.surf.nl/wiki/display/WIKI/Experimental+Technologies+Platform and https://www.surf.nl/en/etp - contact Raymond Oonk at SURF for more info

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

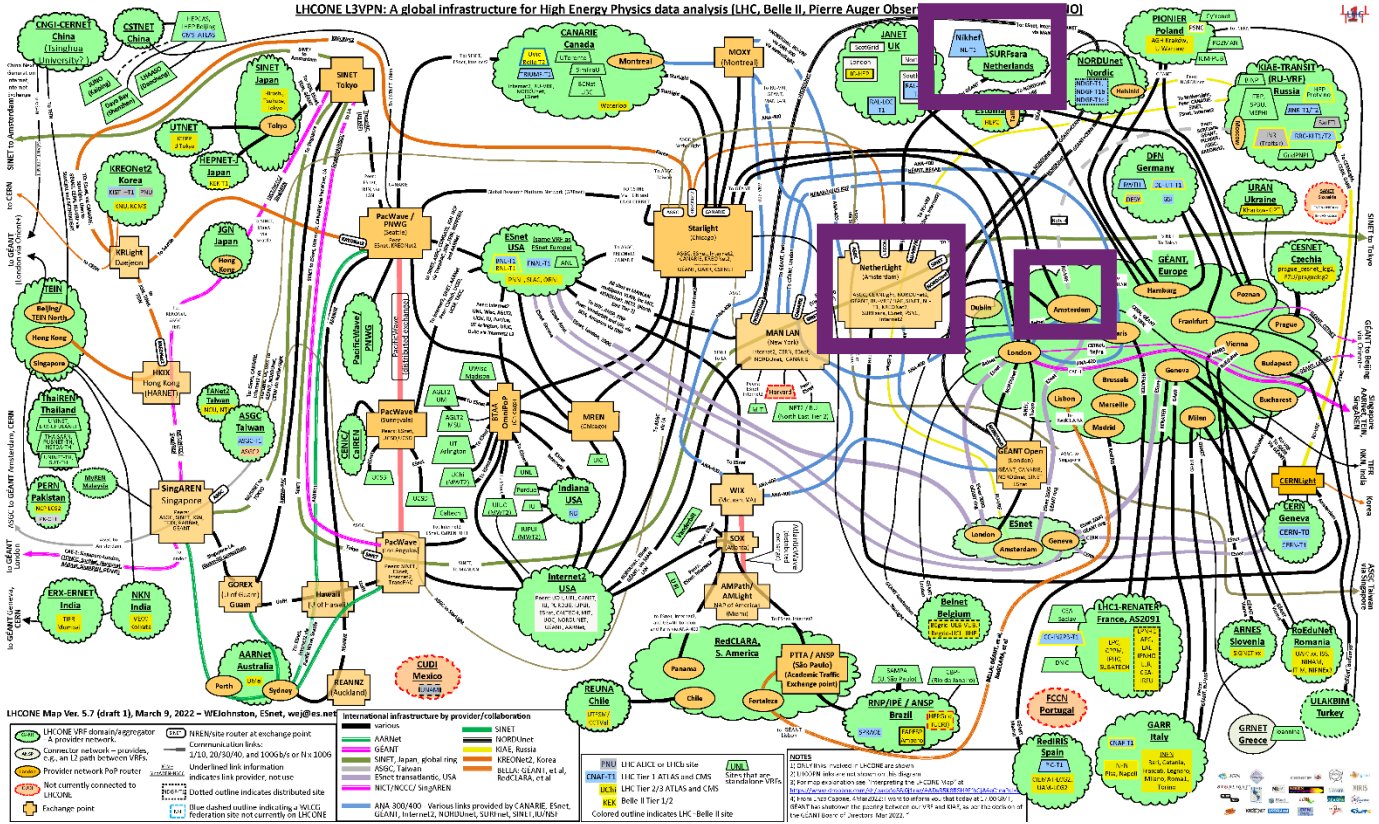

Nikhef, its science and ICT programme, and the NikhefHousing data center

Graphic: Bill Johnston, ESnet, for LHCone

And since speed *does* matter ..

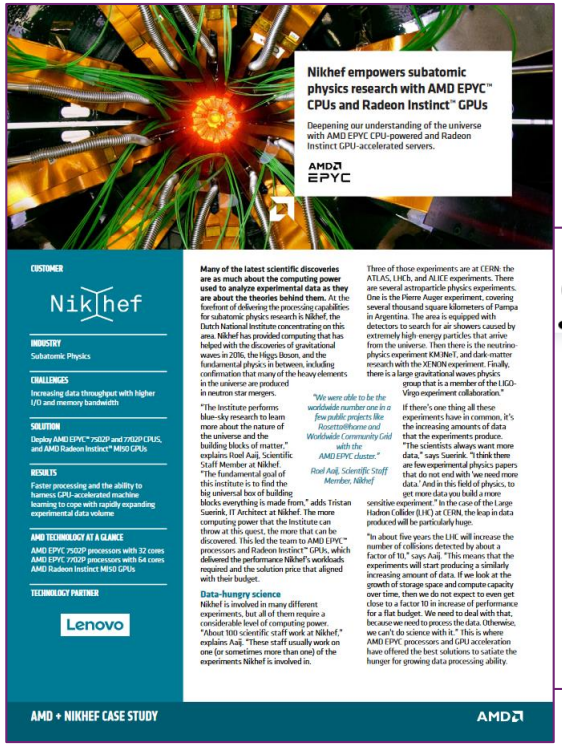

Nikhef empowers subatomic physics research with AMD EPYC™ CPUs and Radeon Instinct™ GPUs

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

AMD EPYC

CUSTOMER

Nikhef

INDUSTRY

Subatomic Physics

CHALLENGES

Increasing data throughput with higher I/O and memory bandwidth

SOLUTION

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

RESULTS

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

AMD TECHNOLOGY AT A GLANCE

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

TECHNOLOGY PARTNER

Lenovo

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

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

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

Roel Aaij, Scientific Staff Member, Nikhef

Data-hungry science

Nikhef is involved in many different experiments, but all of them require a considerable level of computing power. "About 100 scientific staff work at Nikhef," explains Aaij. "These staff usually work on one (or sometimes more than one) of the experiments Nikhef is involved in.

Three of those experiments are at CERN: the ATLAS, LHCb, and ALICE experiments. There are several astroparticle physics experiments. One is the Pierre Auger experiment, covering several thousand square kilometers of Pampa in Argentina. The area is equipped with detectors to search for air showers caused by extremely high-energy particles that arrive from the universe. Then there is the neutrino-physics experiment KM3NeT, and dark-matter research with the XENON experiment. Finally, there is a large gravitational waves physics group that is a member of the LIGO-Virgo experiment collaboration."

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

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

AMD + NIKHEF CASE STUDY

AMD

FUNGIBLE

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

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

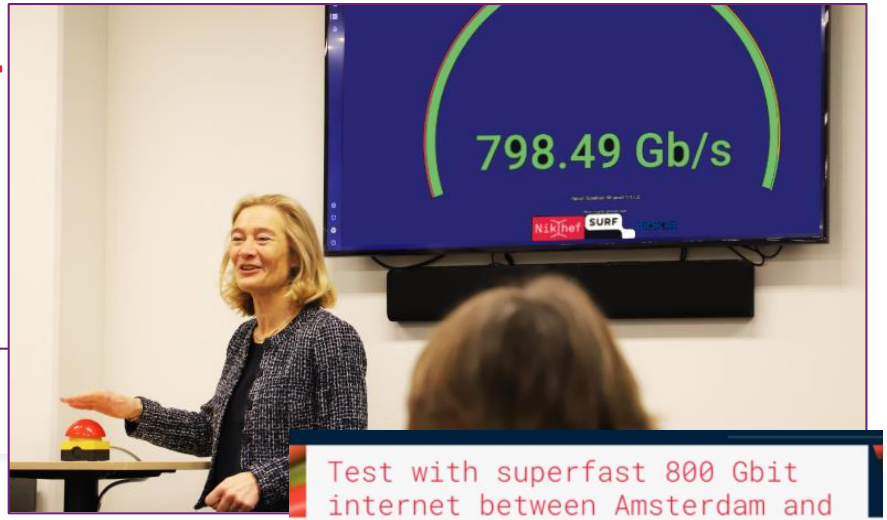

Test with superfast 800 Gbit internet between Amsterdam and CERN successful

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

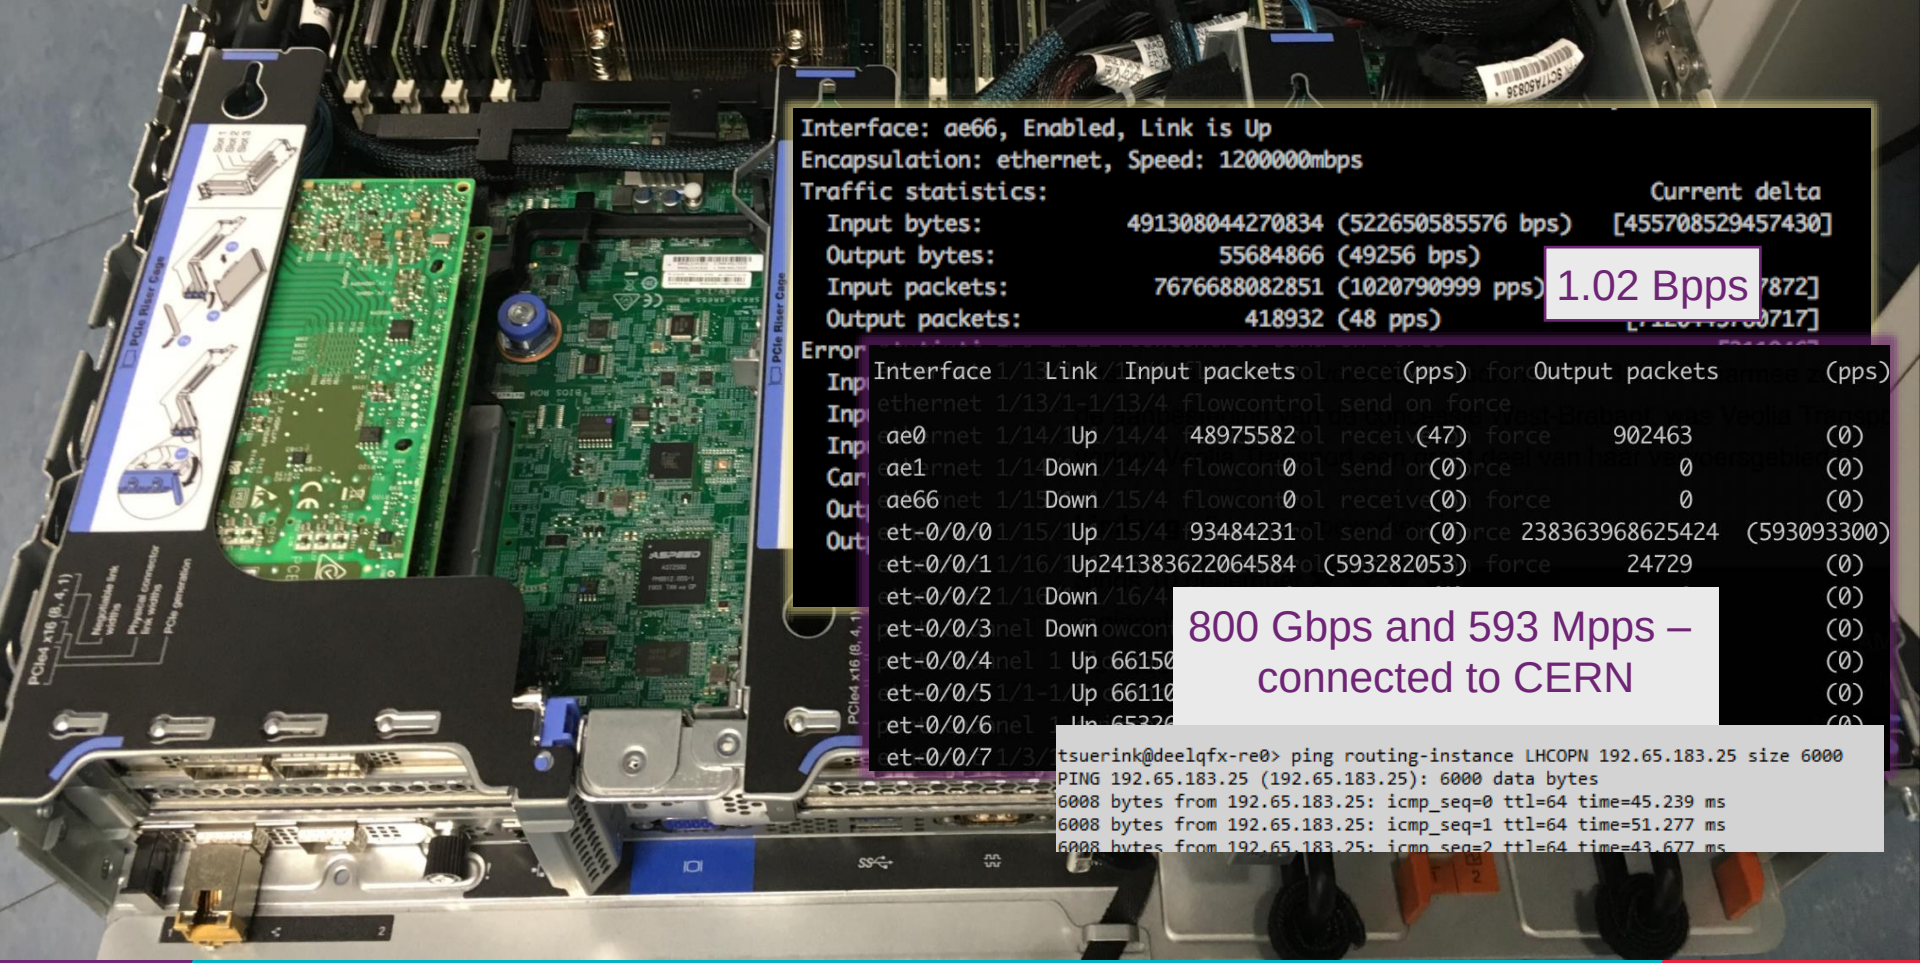
Interface: ae66, Enabled, Link is Up
Encapsulation: ethernet, Speed: 1200000mbps
Traffic statistics:
Current delta
Input bytes: 491308044270834 (522650585576 bps) [455708529457430]
Output bytes: 55684866 (49256 bps)
Input packets: 7676688082851 (1020790999 pps)
Output packets: 418932 (48 pps)
1.02 Bpps
Interface Link Input packets (pps) Output packets (pps)
ae0 Up 48975582 (47) 902463 (0)
ae1 Down 0 (0) 0 (0)
ae66 Down 0 (0) 0 (0)
et-0/0/0 Up 93484231 (0) 238363968625424 (593093300)
et-0/0/1 Up241383622064584 (593282053) 24729 (0)
et-0/0/2 Down (0)
et-0/0/3 Down (0)
et-0/0/4 Up (0)
et-0/0/5 Up (0)
et-0/0/6
et-0/0/7
800 Gbps and 593 Mpps – connected to CERN
tsuerink@deelqfx-re0> ping routing-instance LHCOPN 192.65.183.25 size 6000
PING 192.65.183.25 (192.65.183.25): 6000 data bytes
6008 bytes from 192.65.183.25: icmp_seq=0 ttl=64 time=45.239 ms
6008 bytes from 192.65.183.25: icmp_seq=1 ttl=64 time=51.277 ms

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

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

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,
vapenen tegen
alitie. Onlangs

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

Betastingsdienst

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.

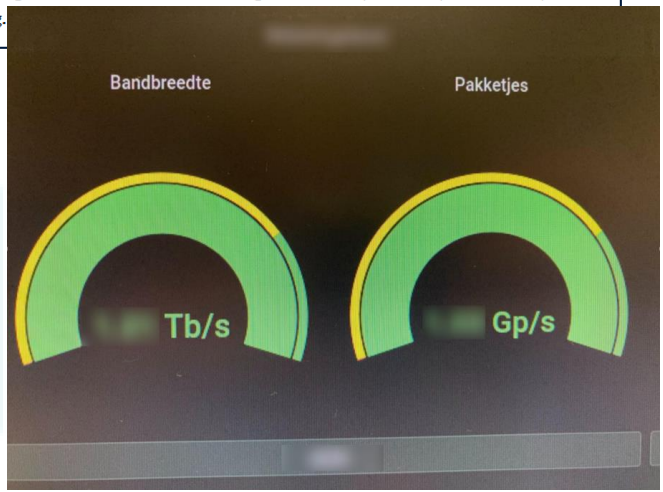

Image sources: belastingdienst.nl, rws.nl, nu.nl, werkentegennederland.nl

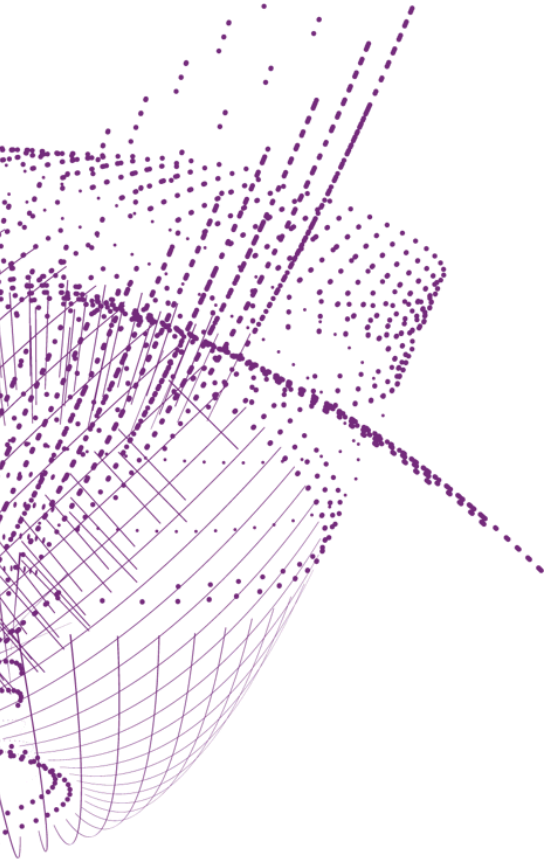

A collaborating ‘ecosystem’
for science and innovation

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 …

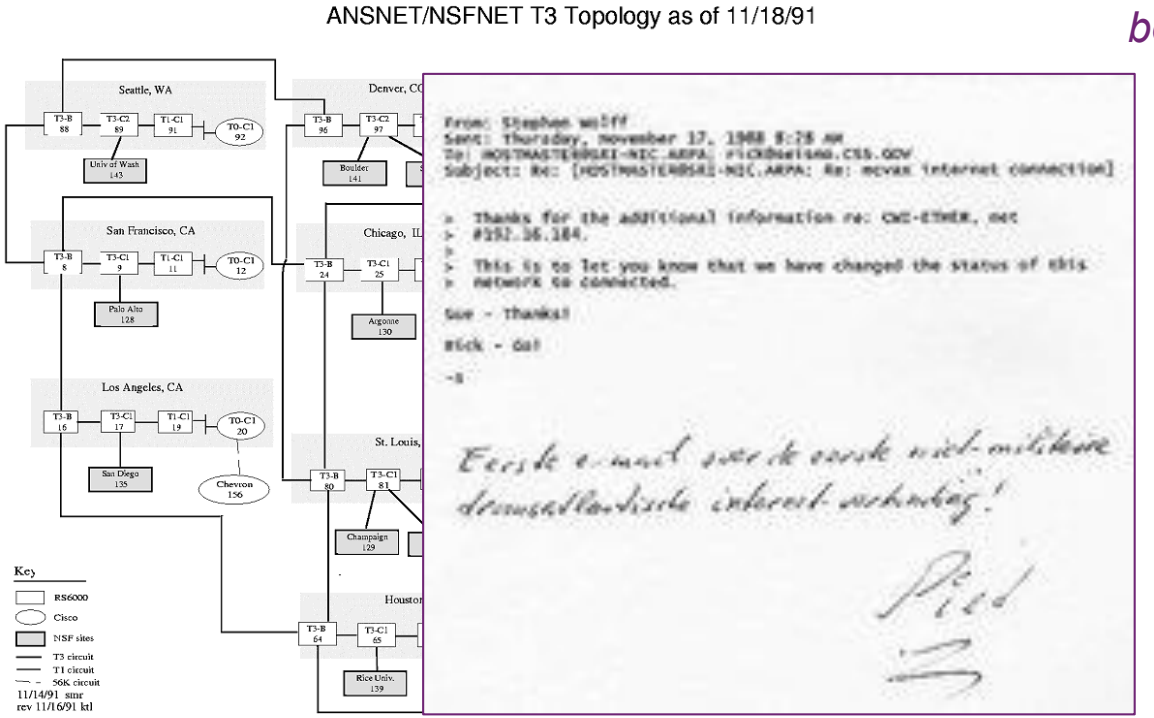

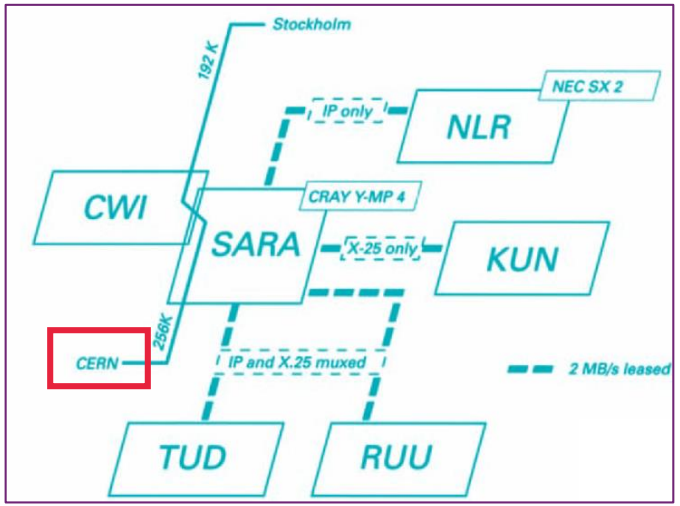

See https://personalpages.manchester.ac.uk/staff/m.dodge/cybergeography/atlas/historical.html for more historic maps ; right-hand image: SURFnet2, 1990
first email to MCVAX at CWI from https://www.cwi.nl/en/news/cwi-celebrates-25-years-of-open-internet-in-europe-in-november/ (Piet Beertema, CWI, 1988)

‘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

1990’s: Computers got smaller and there was space left…

What happens if you welcome two other networks on 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!

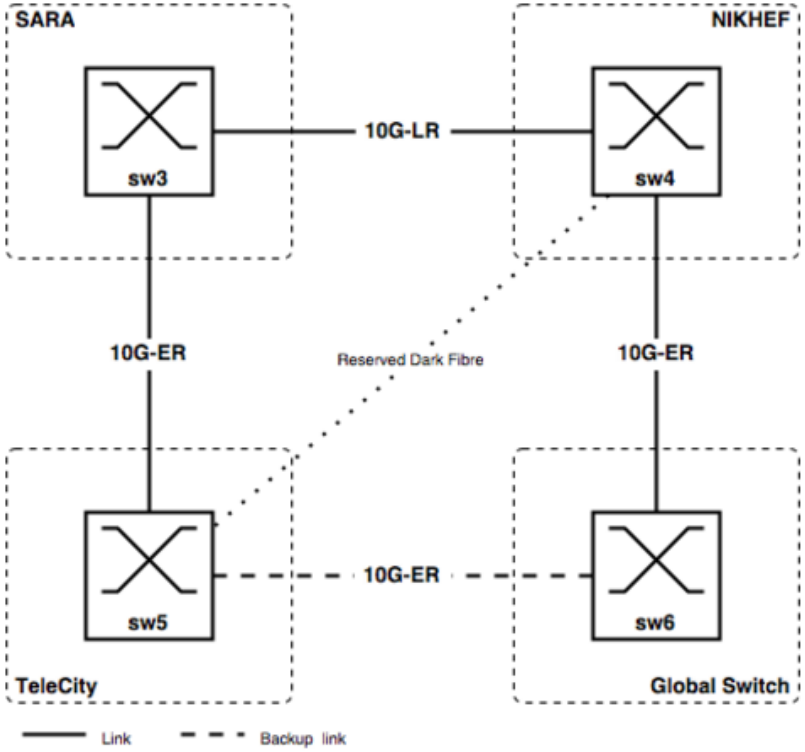

AMS-IX topology, 2002

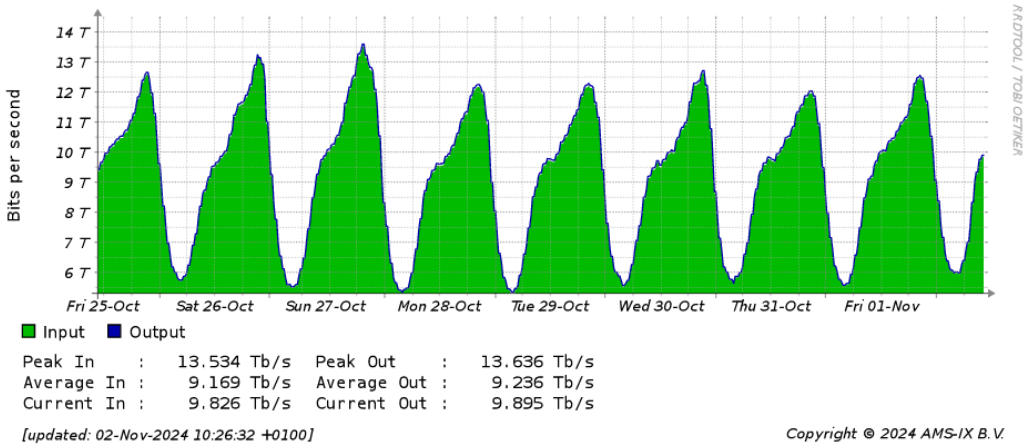

AMS-IX traffic Nov 2, 2024; https://stats.ams-ix.net/ https://www.ams-ix.net/ams/colocations

Just one of the many autonomous systems ...

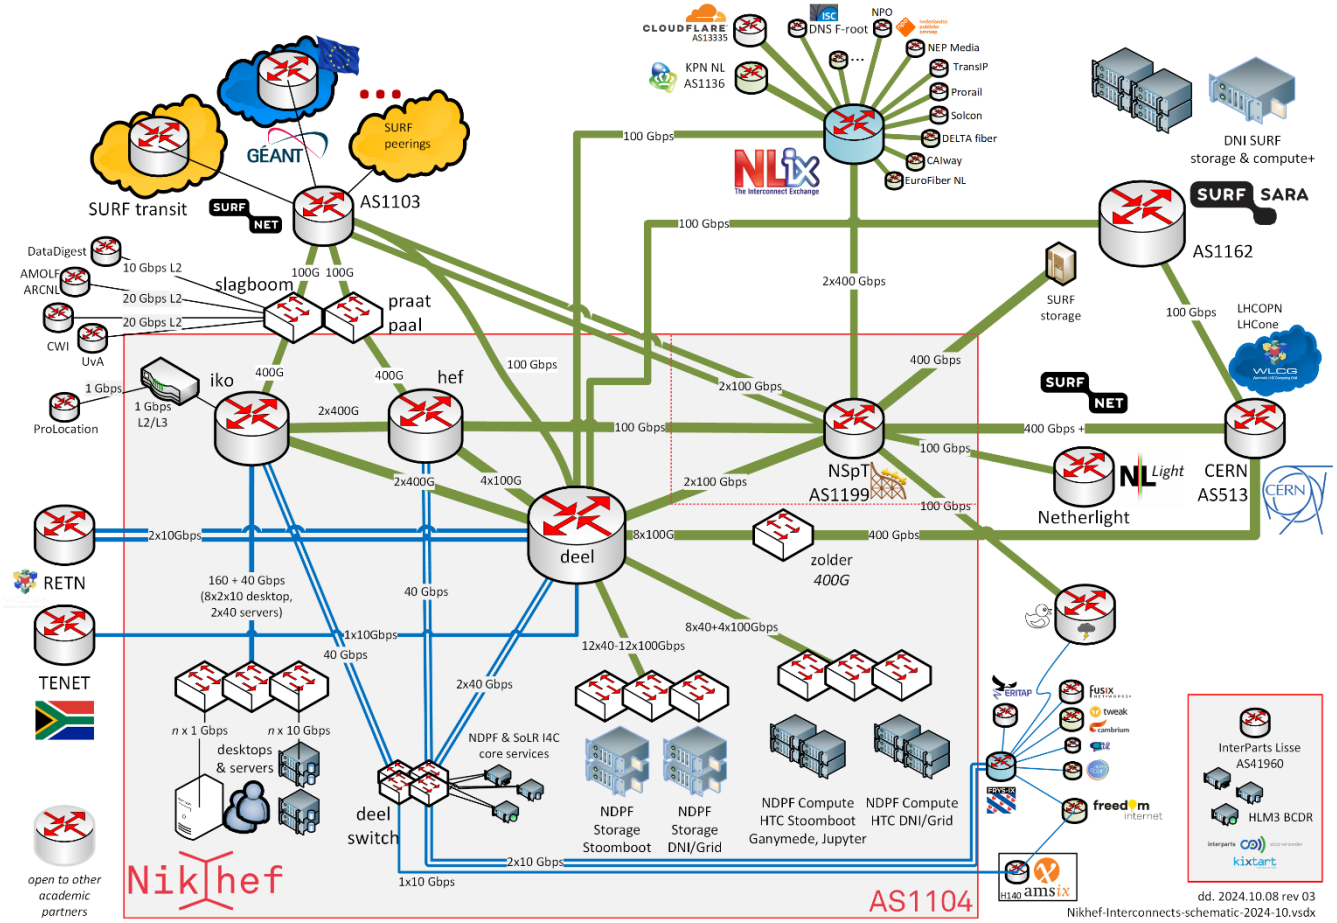

AS1104

Connect each other … and scientific data and instruments

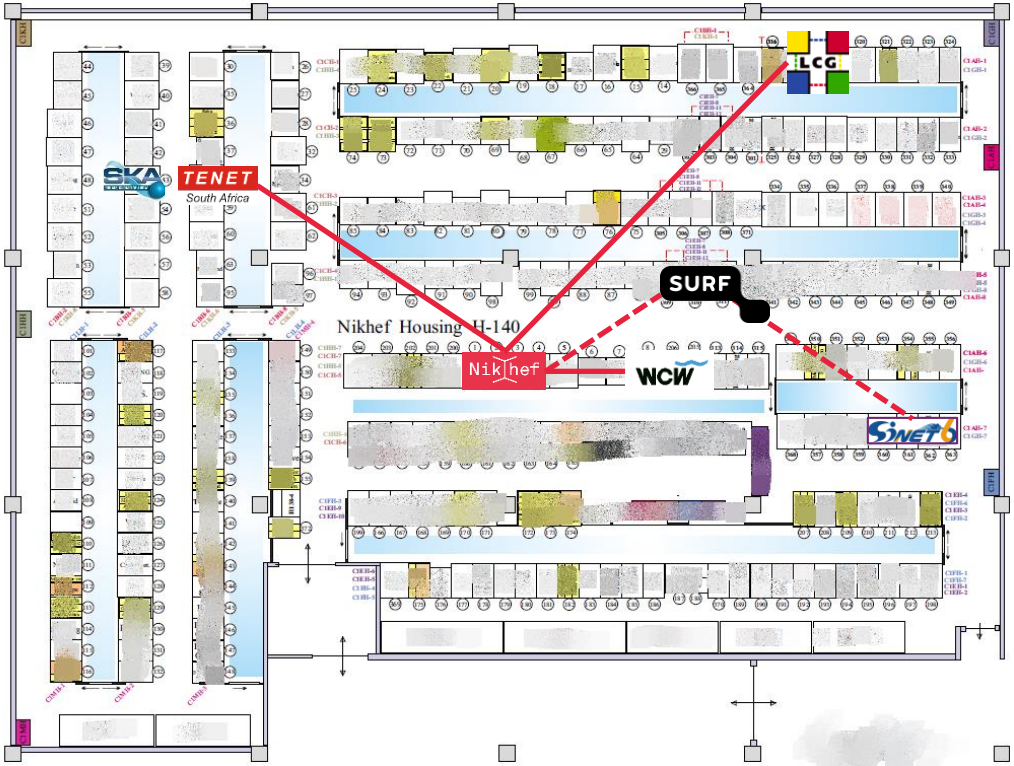

dashed lines: traffic is routed via ancillary facilities of SURF and GEANT first (currently: AMS9 at WCW) SKA Meerkat traffic via TENET

What happens inside a data centre …

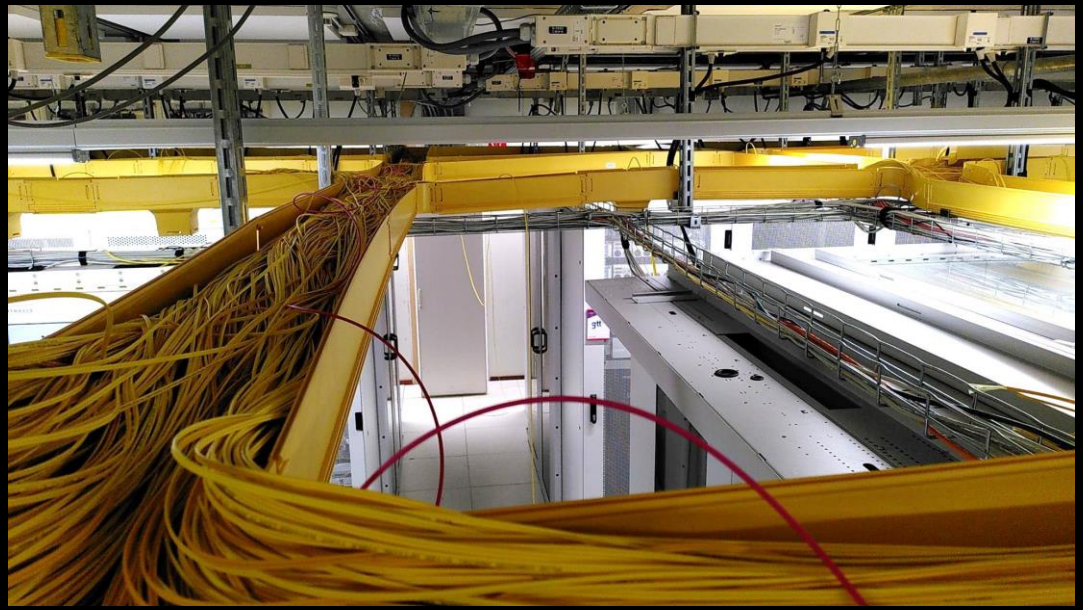

‘Connectivity’ housing and ‘hosting’ are different things:

- 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 ~500 racks
- many different connectivity parties
- **396** 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

Connected networks: see https://www.peeringdb.com/fac/18

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

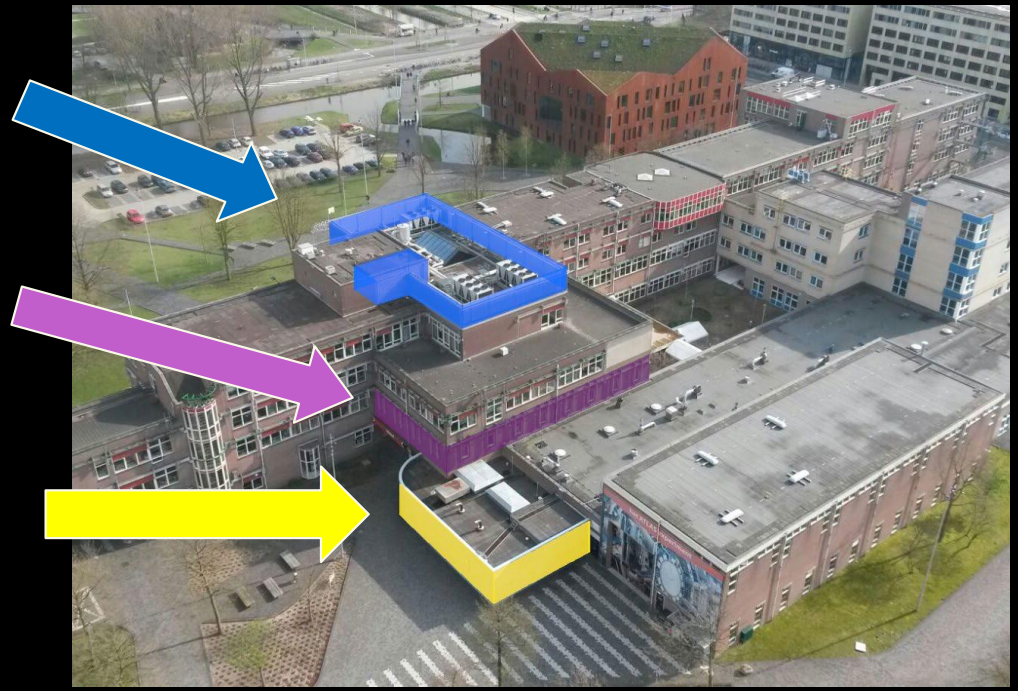

Image: Floris Bieshaar, Nikhef

Power in … and power out …

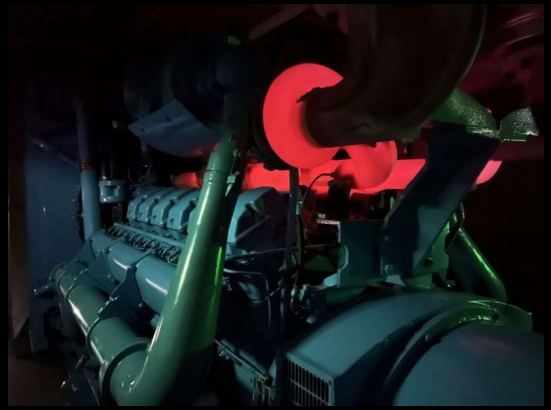

Three generators

- 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 …

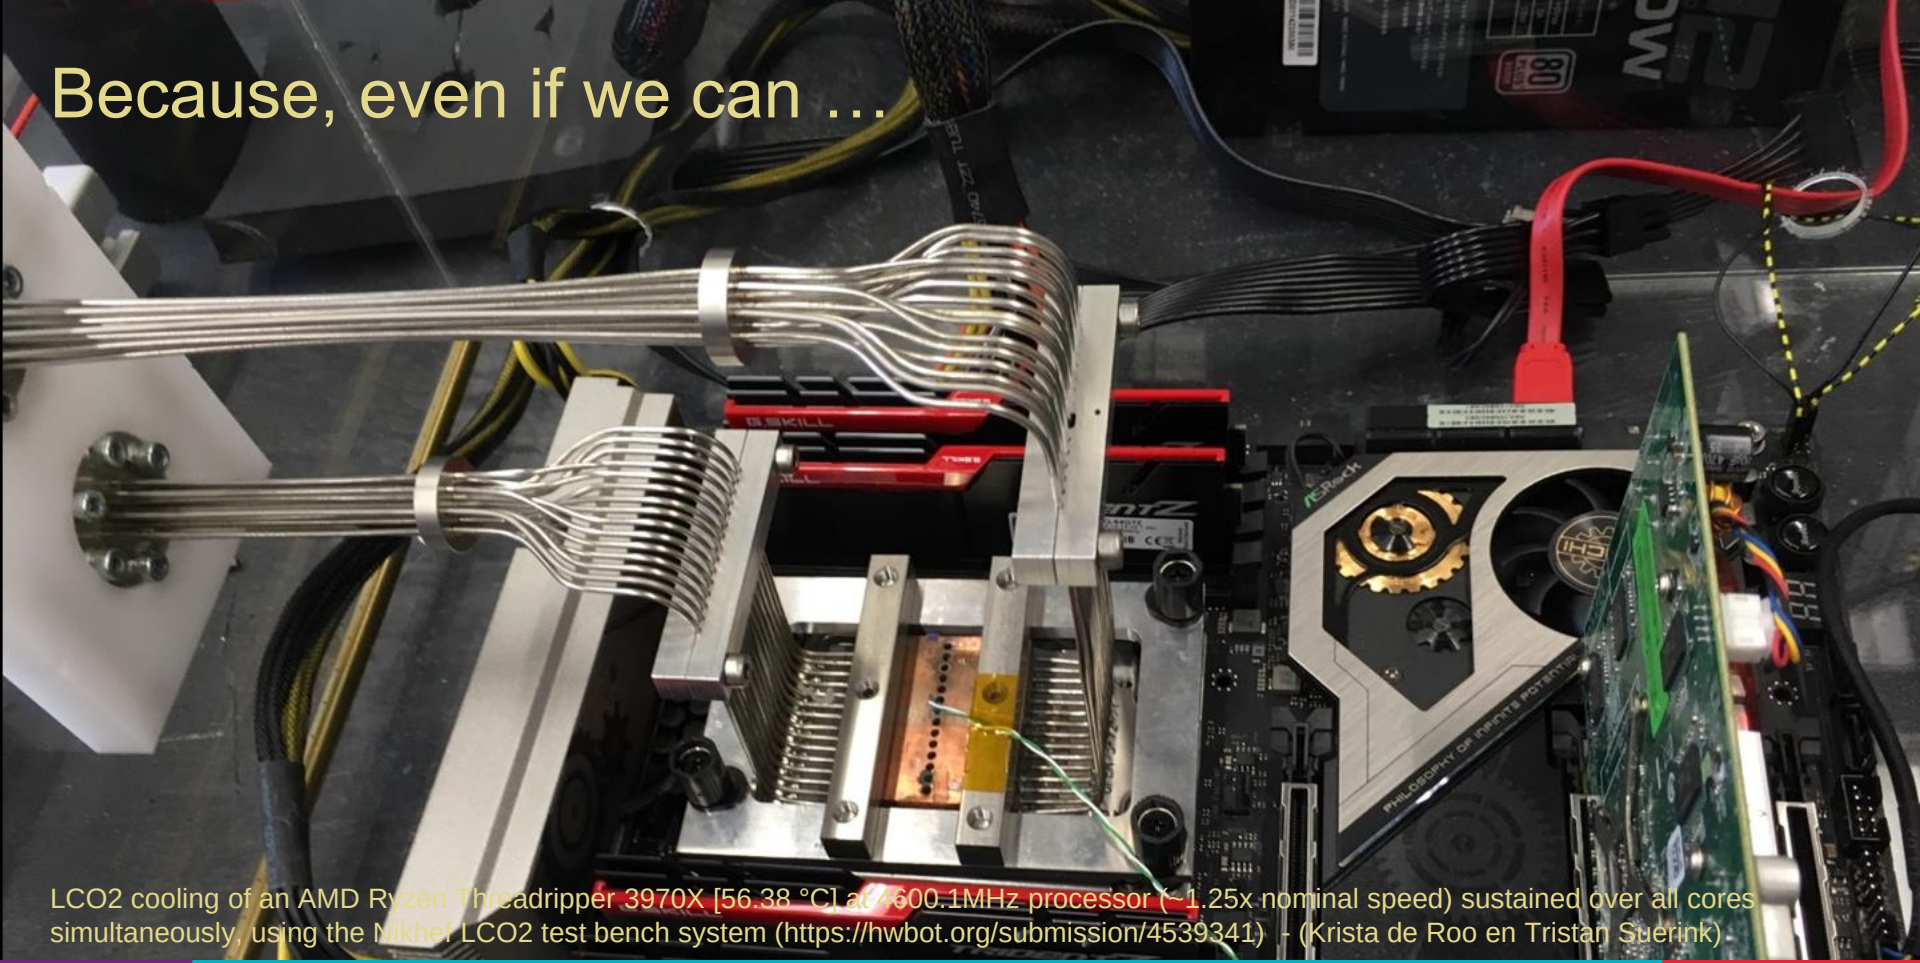

LCO2 cooling of an AMD Ryzen Threadripper 3970X [56.38 °C] at 4600.1MHz processor (~1.25x nominal speed) sustained over all cores simultaneously, using the Nikhef LCO2 test bench system (https://hwbot.org/submission/4539341) - (Krista de Roo en Tristan Suerink)

… it is not always the most scalable solution!

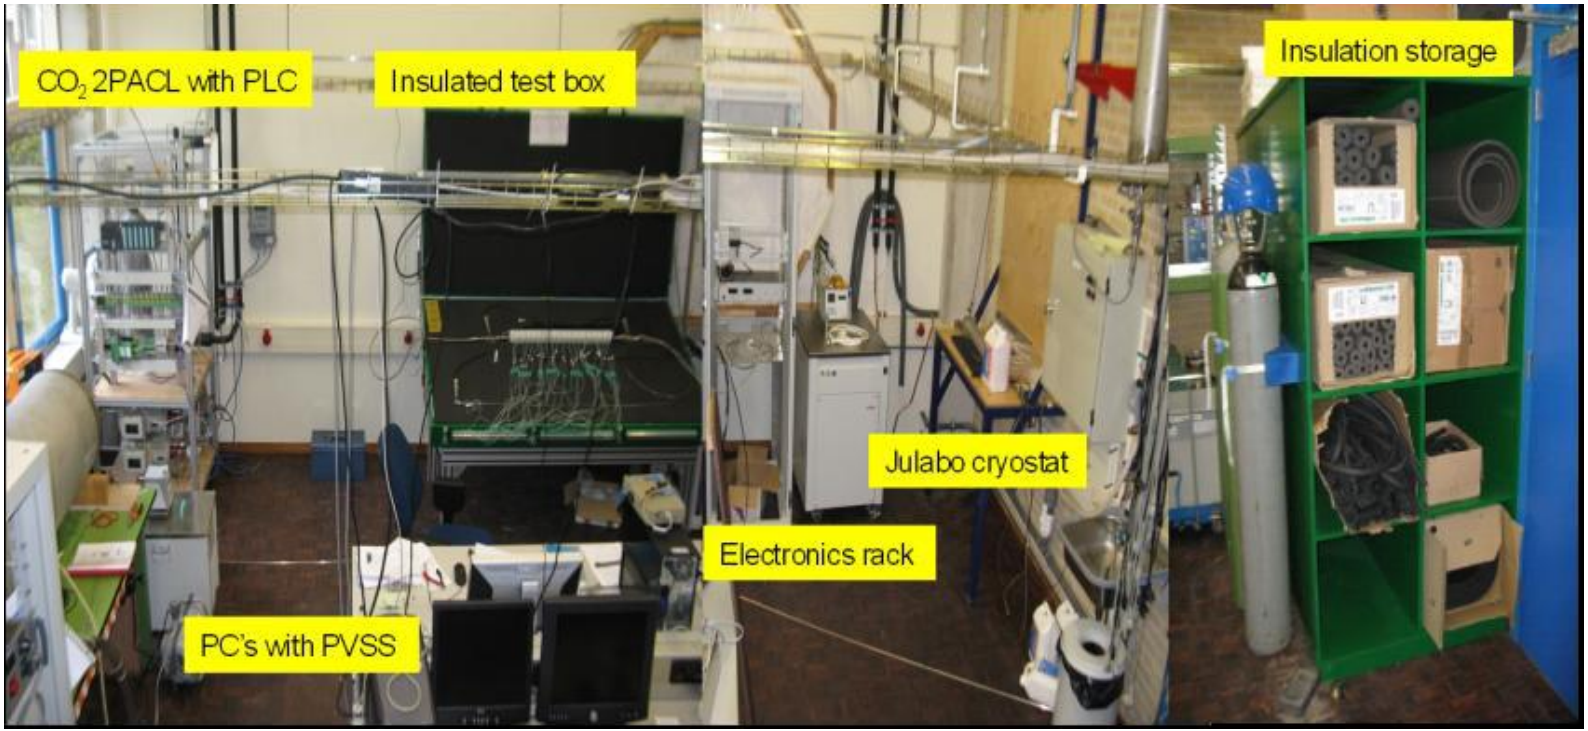

Nikhef 2PA LCO2 cooling setup. Image from Bart Verlaat, Auke-Pieter Colijn *CO2 Cooling Developments for HEP Detectors* https://doi.org/10.22323/1.095.0031

Thanks, and enjoy Nikhef

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

Maastricht University

Nikhef

Tour

hereafter:

NikhefHousing ‘the backside of the internet’
with Anton Mors, Floris Bieshaar, & me