

Nikhef




Maastricht University

Physics Data Processing  
and Computing Technologies

Fancy a (peta)?b(y|i)te ...

David Groep  
Nikhef Jamboree 2023

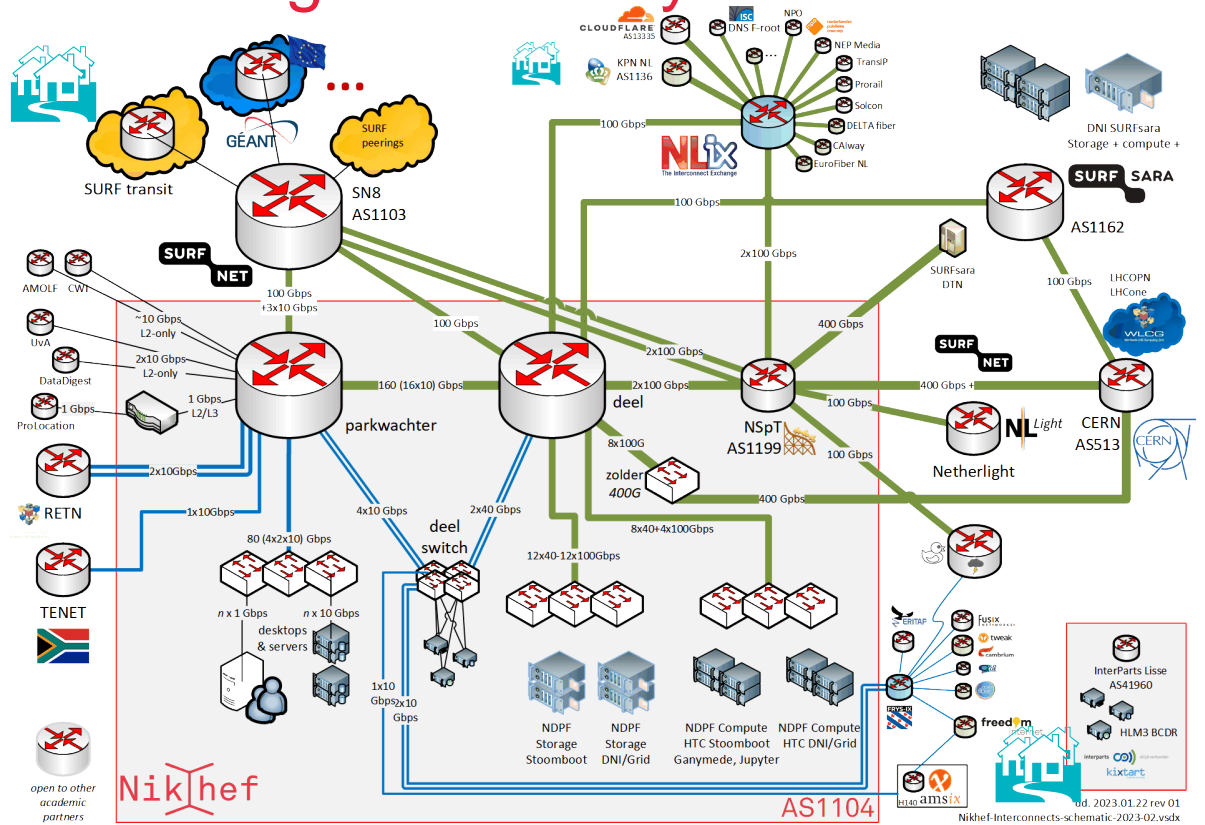


# Computing can make an image with many tracks as well 😊

Nikhef AS1104  
internal network  
and private peerings

Time for some track finding!  
Curious about your way home?  
From home, use


`traceroute -A myip.nikhef.nl`



# The three pillars of Nikhef Physics Data Processing


## Algorithmic design patterns and software

- designing software for (GPU) accelerators, new algorithms, high-performance processors
- software design patterns for workflow & data orchestration




## Infrastructure for trusted collaboration

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



## Infrastructure, network & systems co-design R&D

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



# Towards heterogeneous computing at scale

- Allen for LHCb's GPU-based HLT1 commissioned
- Roadmap for more *hybrid* computing
  - alternative architectures:  
non-x86/ARM, GPU+FPGA hybrid dies

SURF




## For the long term: Quantum Computing algorithms exploration

- in collaboration with our experiments (notably LHCb and GW), QuSoft, SURF, IBM, ...
- personal expectation: 'production' use far away (>2035?),  
but work on algorithms, even if ultimately not QC, very interesting anyway

Image: LHCb's *Allen* team: Daniel Campora (Nikhef & UM), Roel Aaij (Nikhef), Dorothea vom Bruch (LPNHE) (source: LPNHE)



# Innovation on infrastructure



## 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.



**CUSTOMER**  
Nikhef

**INDUSTRY**  
Subatomic Physics

**CHALLENGES**  
Increasing data throughput with higher I/O and memory bandwidth

**SOLUTION**  
Deploy AMD EPYC™ 7002 and 7001 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 volumes

**AMD TECHNOLOGY AT A GLANCE**  
AMD EPYC™ 7002 processors with 32 cores  
AMD EPYC™ 7001 processors with 24 cores  
AMD Radeon Instinct MI50 GPUs

**TECHNOLOGY PARTNER**  
Lenovo

**AMD + NIKHEF CASE STUDY**

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 education for pioneering capabilities for subatomic physics research is Nikhef, the Dutch National Institute concentrating on this area. Nikhef has provided computing that has helped with the discovery of gravitational waves in 2016, the Higgs boson, and the fundamental physics in lepton, including confirmation that many of the heavy elements in the universe are produced in neutron star mergers.

"The institute performs many research to learn more about the nature of the universe and the building blocks of matter," explains Hoel 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 Smeek, IT Architect at Nikhef. The more computing power that the institute can have at its disposal, 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.

**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 these 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 OPERA, 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.

"We were able to be the workable number one in five public projects like *Assess@Home* and *Worldwide Community Grid* with the AMD EPYC cluster," Hoel Aaij, Scientific Staff Member, Nikhef.

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 Smeek. "I think there are few experimental physics papers that do not end with 'we need more data.' And in the 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 less 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 open 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.

## FUNGIBLE

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

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



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

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.

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

### Forse ddos-aanvallen en nerdgrapjes tijdens nachtelijke oefening overheid

Door Rutger Otto

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

202 reacties

Het begon in 2018. Een bijzondere samenwerking tussen overheden om ervoor te zorgen dat de dienstverlening niet wordt onderbroken.

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

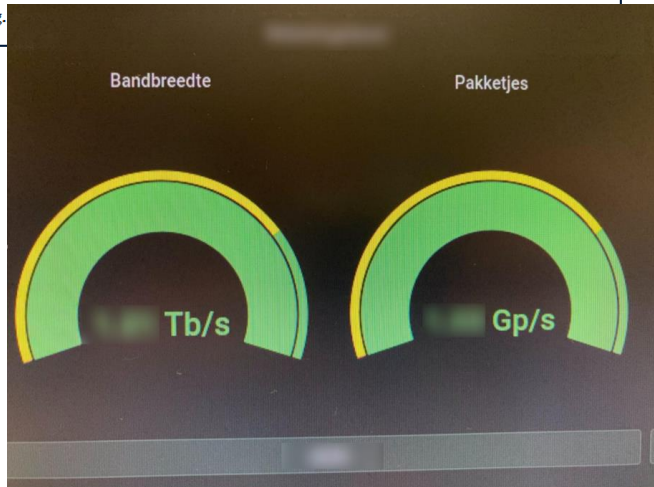
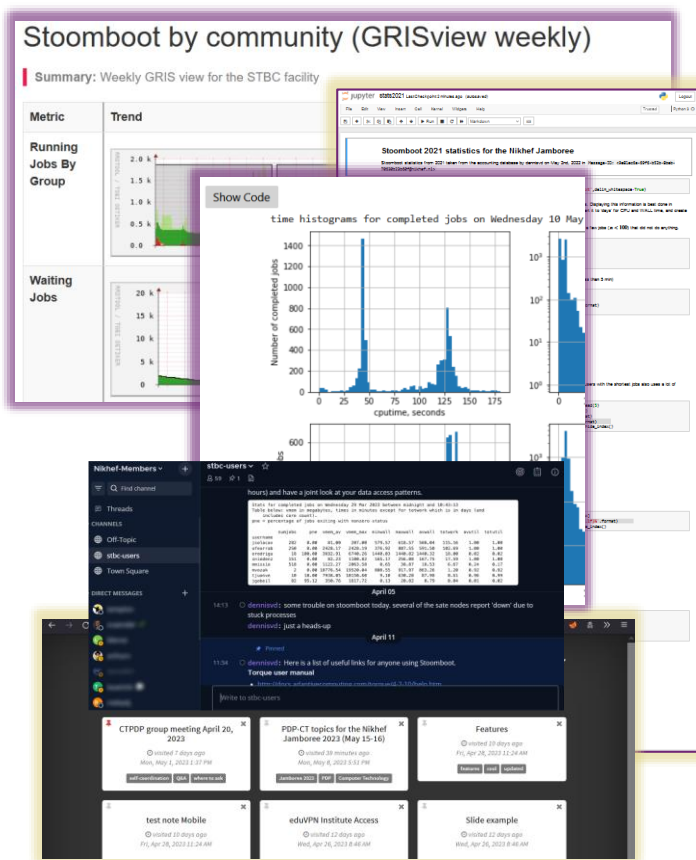


Image sources: belastingdienst.nl, rws.nl, nu.nl, werkentegennederland.nl

# Some services you use today ...

- **Stoomboot:** ~2000 cores and ~3 PByte dCache
- **Callysto:** JupyterHub with \$HOME and SSO login
- **eduVPN:** securely access Callysto and your home
- Mattermost's '**STBC-users**' channel to talk & ask
- **eVA, SURFdrive, and FileSender** to collaborate
- **Experimental services:** ShareMD, Commute, ...



But do read <https://www.nikhef.nl/pdp/doc/experimental-services> before using experimental services ...

Stoomboot statistics: <https://www.nikhef.nl/pdp/doc/stats/stbc-grisview-week>, [https://www.nikhef.nl/pdp/stats/stbc/intern/stbc\\_summ\\_plots](https://www.nikhef.nl/pdp/stats/stbc/intern/stbc_summ_plots)



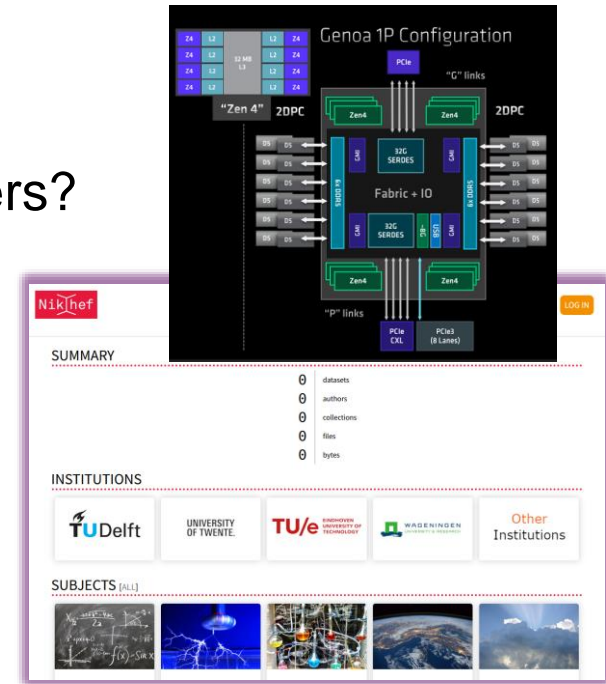
# ... towards ~ tomorrow

But why not get results faster on our next-gen clusters?

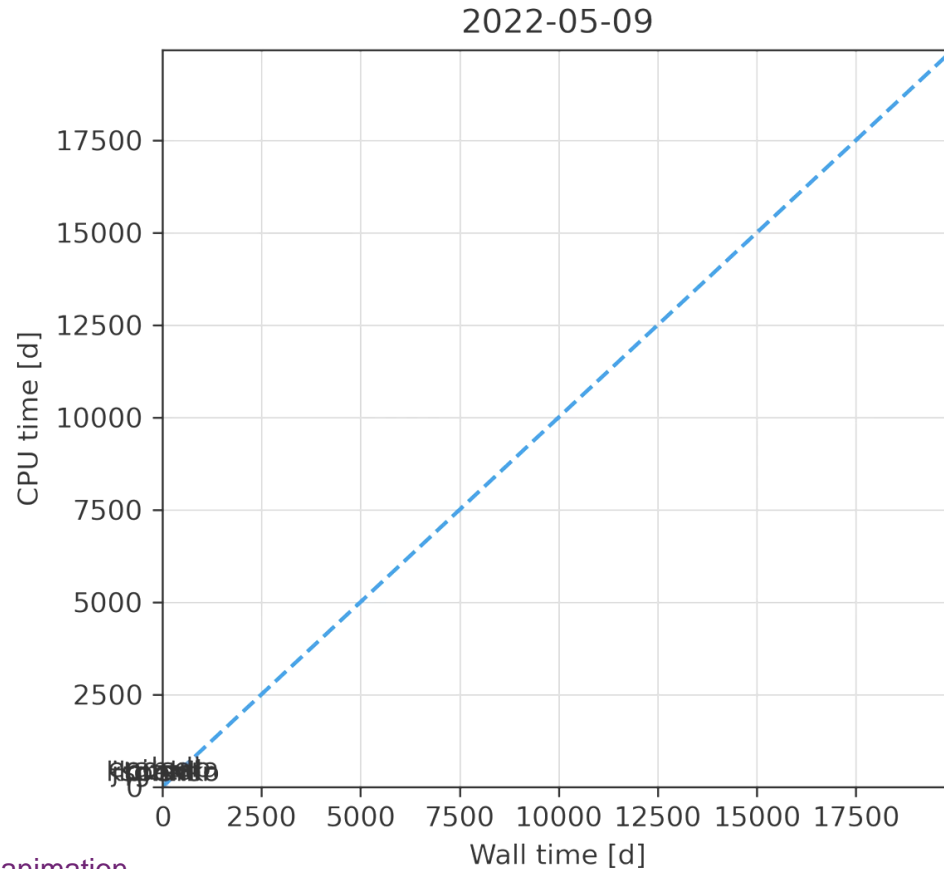
- new clusters likely by the end of 2023 ...
- today's Genoa already +30% HEPscore performance *although ... why is nobody using stbc-iexp?*

And containerise your work in the future!

- better access to **GPUs on stoomboot**
- run on **newer hardware** *that will use Rocky 8 or 9*
- prepare for **data analysis preservation** and good research data management *when we link Stoomboot dCache to our institutional Research Data Management*



And this year,  
the stoomboot  
winner is ...



Animation: Dennis van Dok,  
source: <https://gitlab.nikhef.nl/pdp/stoomboot-animation>

# Your new knowledge base starts here, on [kb.nikhef.nl/ct](https://kb.nikhef.nl/ct)

The screenshot shows a web browser window with the URL [https://kb.nikhef.nl/ct/Interactive\\_GPU\\_nodes.html](https://kb.nikhef.nl/ct/Interactive_GPU_nodes.html). The page title is "Interactive GPU nodes". The left sidebar contains a navigation menu with items like "Home", "Accounts", "Calendar & Mail", "Printing", "Laptops & Hardware", "Network (WiFi/VPN)", "Telephone", and "Meeting Rooms". The main content area has a heading "Interactive GPU nodes" and a sub-heading "Instructions". The text under "Instructions" reads: "Registering to indico.nikhef.nl. You **do not have** to register yourself to the Indico service. Login with your single sign-on Nikhef credentials (account name and password). Usually there is no need to login in case you only want to have a look at the agenda. You have to login in case you want to manage events, upload categories and are managed by top level".

This screenshot shows the "Instructions" page for connecting to a Nikhef desktop. The page title is "Instructions". The main text says: "If you want to connect to your Nikhef desktop after installing eduVPN, follow [these instructions](#)." Below this text is a section titled "Windows" which contains a sub-section "Instructions" with the text: "Visit the eduVPN website to download the software:". To the right of this text is a sub-section "Images" with the heading "1. Download eduVPN" and a list of links: "voor Windows", "voor Android", "voor macOS", and "voor iOS". There is also an image of the eduVPN logo. The left sidebar is visible, showing the "Network (WiFi/VPN)" section expanded to "Connecting to Networks", where "eduVPN" is highlighted. The right sidebar shows a "Table of contents" with items like "Introduction", "Description", "Secure Internet", "Privacy considerations", "Institute Access", "Privacy considerations", and "Useful links".

# Fancy cake?

## Meet at the Computing Office hours – ‘ask us anything’

- 1300-1500, first Thursday of every month
- SP110 ‘Spectrum Noord’
- slice of cake with every question

## Join the Nikhef Computing Course for a structured learning approach

- <https://www.nikhef.nl/pdp/computing-course/>
- materials at <https://indico.nikhef.nl/e/computingcourse2022>



*Next up*

“A tour of Italian CPUs, and our speedy road south”  
(Erik Kooistra)



David Groep

davidg@nikhef.nl

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

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



Maastricht University

Nikhef

