

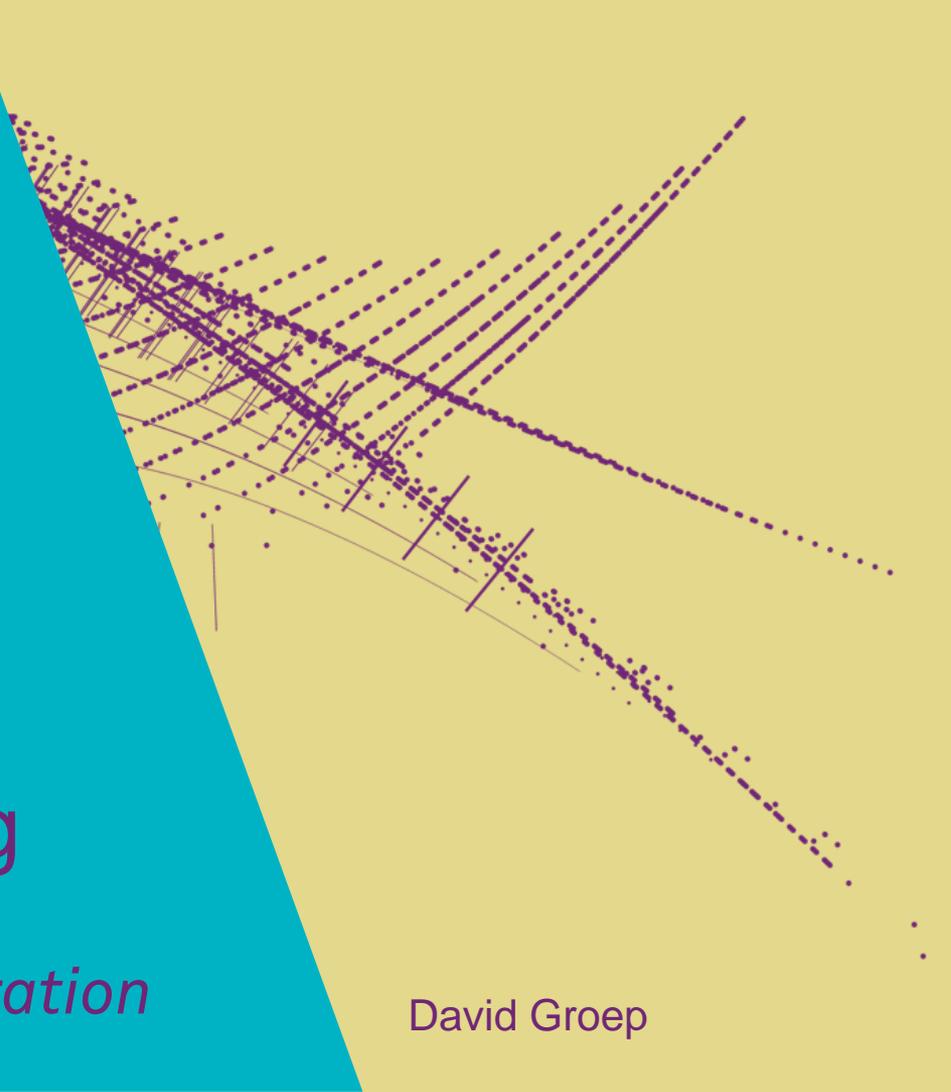


Nikhef SAC February 2026

Physics Data Processing

*accelerating 'time to science'
through computing and collaboration*

David Groep



Fourty years of data driven computing at Nikhef (almost)

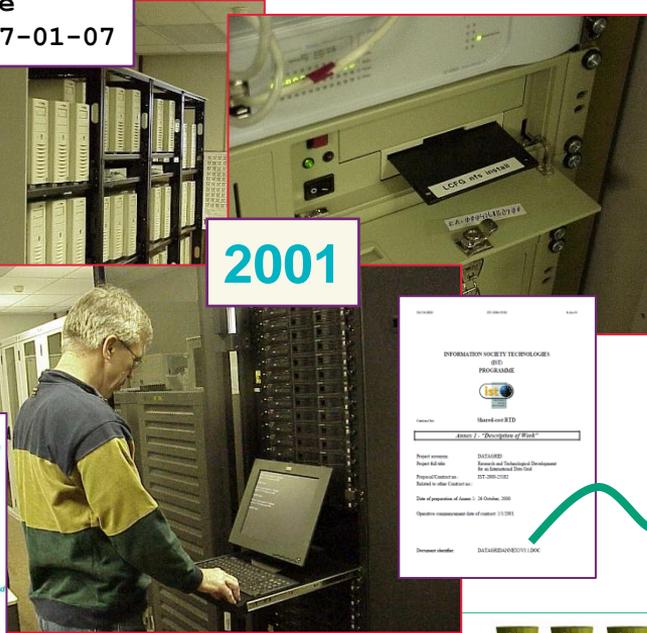
```
$ whois nikhef.nl
Domain name: nikhef.nl
Status:      active
Creation Date: 1987-01-07
```

1996

1990



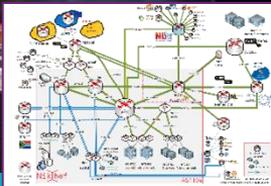
2001



plus WTCW Virtual Laboratory, DAS-2, Virtual Lab for eScience, BiGGrid, since 1999



2026



'Means and purpose' of processing

- **Accelerate** time to results with **efficient use** of computing through algorithmic design, novel software methods, and optimisation of data driven workflows and models
- exploit **new systems design and integration** for efficiency and scalability with heterogeneous architectures, innovative vendors, faster networks
- ensure physics **access to the broadest range of science services** with ease

while operationally

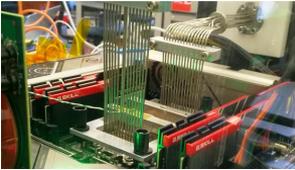
- ensure our computing capacity is sustainable ... *as we fund it to a large extent from non-HEP/APP sources, and have a commitment to open science and collaboration*
- continue push for generic access of services ... *important to keep individual RIs aligned*
- move more of our experiments to generic computing infrastructure ... *so we can contribute compute more effectively, cheaper & better to our physics programmes*



The pillars of Physics Data Processing

Algorithmic design patterns and software

- GPU & accelerator software, algorithms, heterogeneous processors optimisation
- algorithmic design R&D



Infrastructure, network & federated systems R&D

- building 'research IT facilities' & federated access for RIs
- data science at volume scale
- research *on* IT infrastructure



Infrastructure for scalable ML in HEP & APP

Inference optimisation for data throughput, HTC/HEP and HPC/AI-factory convergence, hybrid workflows
“scaling data processing from local to national to AI factories”

Infrastructure for trusted collaboration

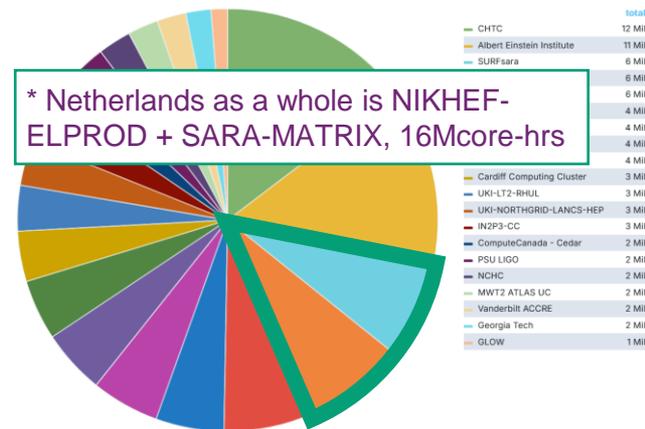
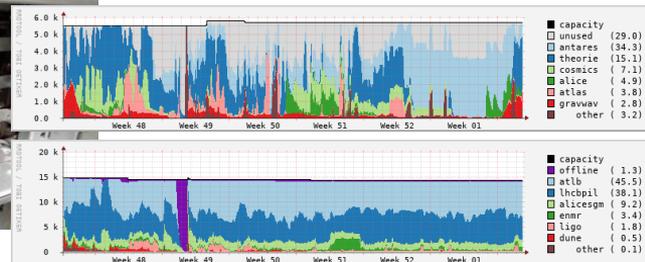
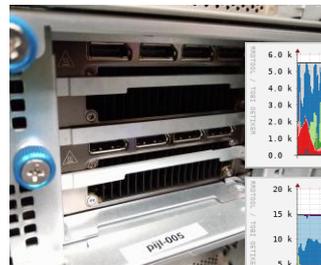
- community trust and identity
- managing complexity of collaboration mechanisms
- securing the infrastructure of our open science cloud
- scalability of policy negotiation



← T&I, Federation, Research Infra Commons

Some computing, engineering and trust highlights

- seeing significantly increased **use of GPUs** in our analysis facility for AI/ML
- consolidated migration to HTCondor for both local analysis and our federated facility – improved **98% occupancy** of our NL-federated Tier-1
- ‘NL-T1’ **compute for IGWN now >15%** of global capacity*
- faster networks & more packets with ‘interesting’ effects
- security and risk management **for EOSC EU Node**
- new **AARC architecture & trust framework** delivered tokens, OpenID Federation, and review of EU ID Wallets



Data: occupancy data for Nikhef NDPF NL-T1 and DNI compute (<https://www.nikhef.nl/pdp/doc/stats/ndpf-prd-grisview-year>). IGWN data: OSG Gracia portal, whole of 2025: Nikhef 7.5% + SURF 7.5% <https://gracc.opensciencegrid.org/d/9u1-Q3vVz/cpu-payload-jobs?orgId=1&var-ReportableVONName=ligo&var-Project=All&var-Facility=All&var-Probe=All&var-interval=1d&from=1735689600000&to=1765756800000&viewPanel=9>

Pathways via projects and collaborations

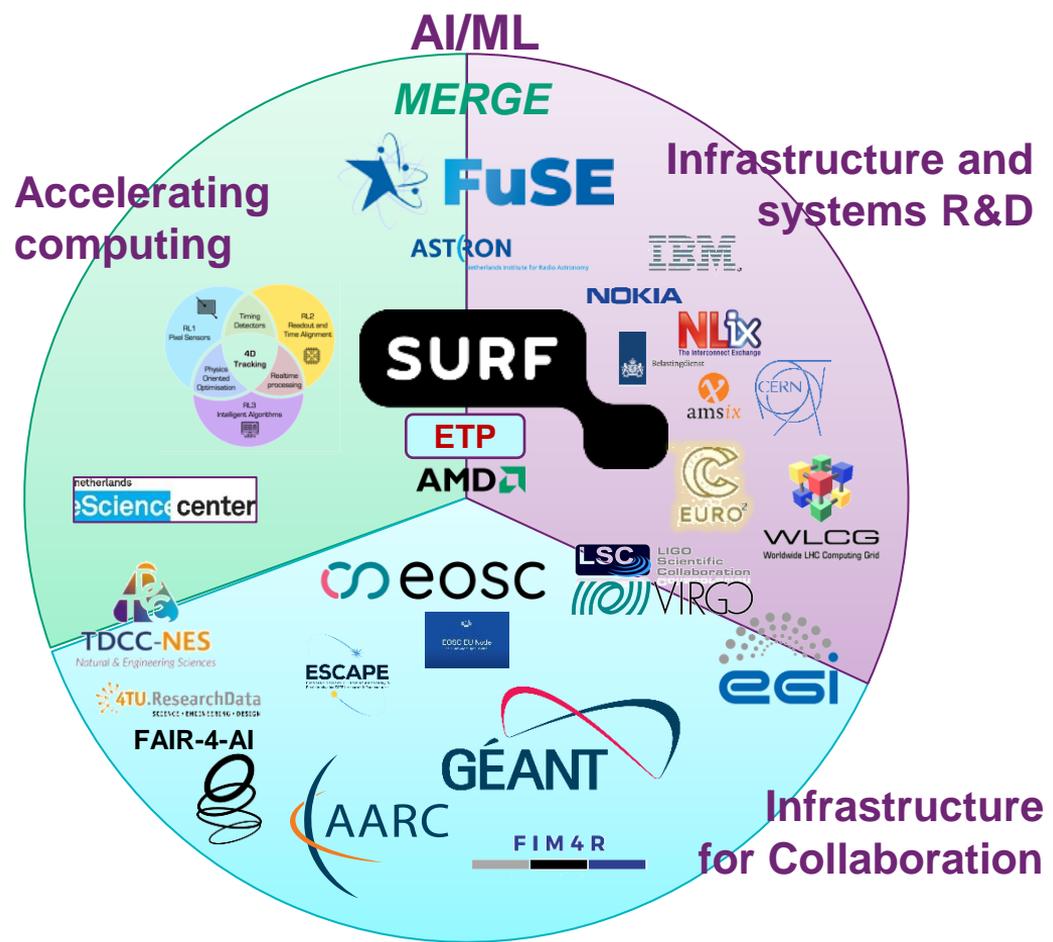
New initiatives and projects

- *strengthen* the strategic areas
- ensure *continuity* of research and infrastructure

future project pathways planned: MERGE, SHAKE, ...

Public partner R&D engagement

AMD, Nokia/Cerebras, NL-ix, AmsIX...
Dutch national government



From the SAC's last year observations

Report of the Nikhef Scientific Advisory Committee Meeting
June 5-6, 2025
E Elsen (chair), S Farrington, P Hernandez, M Kowalski, C Salgado, D Shoemaker and G A Stewart
<https://indico.nikhef.nl/event/6637/>

Introduction
The Nikhef Scientific Advisory Committee (SAC) held its meeting on June 5 and 6, 2025. The committee expressed gratitude to the management and the Nikhef community for the exceptional

• • • •

PHYSICS DATA PROCESSING
The PDP group has an ambitious remit, ranging from developing and supporting the use of modern algorithms on modern hardware, to the facility deployment at Nikhef (together with Nikhef Computing Technology), and essential collaborative infrastructure for scientific computing in international projects. The team has established good relationships with the physics teams locally and at the university partners and has a focus on training and help-desk support for software issues. As software is clearly critical (and frequently undersupported) across the institute's physics activities, we find the team size to be quite small for this role. Growing links to research software engineering teams at the university partners may be a means to broaden the support base for software and help address important issues of sustainability. Due consideration of software

4 / 7

lifecycles should be given and new software technologies pursued. In particular, growing use of AI/ML is being observed and this means that the PDP group should also play an important role in both resource provisioning and skills training for new ML-based software techniques. Resource delivery for international projects benefits hugely from stressing common interfaces and solutions, which we commend. Keeping in-house resources reduces risks and maintains vital expertise.

Looking back and forwards

- ‘Keeping in-house resources...’
- ‘... benefits hugely from stressing common interfaces & solutions’
- ‘... growing use of AI/ML is being observed’

‘... an ambitious remit ...’

Report of the Nikhef Scientific Advisory Committee Meeting June 5-6, 2025

‘Keeping in-house resources...’

We need our operational capability to be relevant* in the ecosystem

- we *are* the ‘staircase model’ put into practice: scaling from ‘Stoomboot’ to our federated ‘DNI’ national and global resources – *and influencing these by example*
Naturally aligned with the Research Infrastructure Commons
- we provide the ‘rapid-response’ capability of the innovation platforms and the (national) ‘Experimental Technologies Platform’ (ETP) for SURF

Outsourcing (or even off-site co-location) of facilities would jeopardize that role in innovation – and be detrimental to our talent case ...

But we (cannot) do not do this alone

- MERGE proposal explicitly includes the Commons as mechanism to import and export expertise on AI/ML and infrastructure usage



* aligns well with the Dutch Wennink report on national relevance and innovation

‘... benefits hugely from stressing common interfaces & solutions’

Designing for exploitation of resources for *scheduled processing* **and** *chaotic analysis*:
a combination of research & development + operations & support

In acquisition and real-time acceleration for 4D

- FASTTRACK accelerated R&D facility (WP5.1),
and algorithmic work *in the experiments themselves* for WP5.2
- Jointly with SURF ETP innovation collaboration

In making HPC resources (and AI factories) a better fit for data intensive HEP/APP

- with SURF: workflows on Snellius preparing for ‘Lorentz’ tender – leveraging ATLAS & Alice
- EuroHPC Federation and the Tier-1→Tier-0 transition:
systems design, AAI, Interfaces to Architecture Testbeds
- and also just the basics: helping promote the NWO ‘Rekentijd’ granting scheme

In developing a smooth transition ‘staircase’ path for both software and resulting data sets

- NWO Thematic Centre (TDCC) NES project on JupyterHub alignment of Stoomboot & Snellius
- OpenScienceNL project on reproducible software at scale (with 4TU, TUDelft, NLeSC)

Innovation: ‘... growing use of AI/ML is being observed’ ...

New AI/ML R&D group brings activities together in *both* the experiments *and* for PDP

- links strong physics programmes on AI/ML with the ‘data, physics, workflows, and infrastructure research’ role of PDP
- R&D on infrastructure extensions for AI/ML and accelerated compute generally

We will be evolving

- the physics staff and PoC hardware will be available starting mid-2026
- transformative facility for inference (MERGE, under funder’s review)

While other aspects remain as relevant in the future

- systems R&D and our role as a ‘Tier-1’ facility for HEP/APP
- in a complex multi-stakeholder world, where mixing long and short term objectives is non-trivial
- accelerating time to results from 4D fast timing, from close to the detectors to the final plots 😊



Examples of upgrade investments:

Upgrade investment can enable (beyond) state-of-the-art equipment to be used and bring about collaboration between scientists at a larger scale. Contributions to greening of an LSRI can be made by greening into the upgrade investment or by greening existing LSRI. Investments are investments in:

- development and/or procurement, construction or the latest state-of-the-art equipment of the LSRI;
- expansion and enrichment of data and sample collections;
- access to and/or development of the latest AI and data processing techniques;
- linking or integrating existing complementary LSRI, connecting new user groups and/or setting up a joint e-infrastructure;
- Modifications to existing LSRI needed to reduce its carbon footprint.

Balancing training & support with innovation R&D

‘seeing significantly increased use of GPUs in our analysis facility for AI/ML’

Pathway to practical impact and *efficient* use is via training and tutoring ... with infinite demand but *always* limited capacity

- computing courses need investment from two sides
- getting scalable AI/ML workflows:
 - locally and on ‘scale-up’ systems
 - from Stoomboot → Snellius → LUMI* → our federated HTC / Tier-1 in 2028+
- plus the research data management, FAIR ‘digital research process’ training, ‘regular’ computing, software engineering, knowledge safety, computer security ...

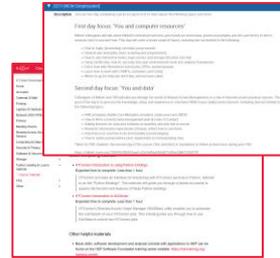


*LUMI: EuroHPC pre-exascale system in which NL has a share, as we do also in the Jules Verne consortium

Collaborative training & support for our physics programs

We lack a ‘sweet spot’ in trainers: it is beyond the helpdesk expertise, but our next layer is level-3+ engineering researcher support ... and although it’s really nice to be able to give trainings and lectures, and see the effect of your R&D efforts, time is always limited.

- ‘computing courses need investment from all sides’
 - Allowing time to partake in the trainings: long term benefit vs short term time investment
 - Help making courses scalable: a trainee-to-trainer programme for PhD & postdocs?
- Course development as a collaborative exercise
 - Use - and co-develop where needed - *contextualised* trainings
 - Carpentries, NWO-I DCC, TDCC NES, HSF Training hub is there, but need ‘local flavour’ to be attractive
- Giving the courses requires people, and is resource intensive
 - emerging AI group at Nikhef helps raise awareness and importance: more interest in getting trained
 - would giving formal recognition (e.g. in terms of certificates) for trainers help after their time @Nikhef?
We can recognise their contribution internally (and do!), but people tend to move 😊
 - Recognition & Rewards (SFDORA) relatively strong in NL and Europe, but not (yet) much elsewhere ...

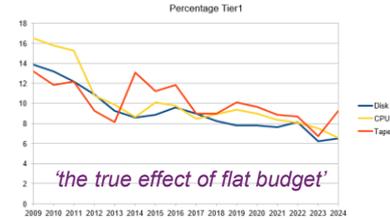


https://kb.nikhef.nl/ct/Course_Tutorials.html

Weaknesses, threats ... and mitigation scenarios

Weaknesses

- Continuity as 'cross-domain Tier-1 hub' based on project, not institutional, funding
 - fiche 'digital infrastructure', NWO-SURF report 'Reken er maar (niet meer) op' ... MERGE is only a stopgap measure for this, as, again, it is an upgrade *project*



Threats

- Market developments in semicon ☹
 - phased and deferred acquisition, market research, and vendor engagement/co-innovation to get better conditions is the best we can do – but likely not enough for 2026-2027
- Divergence of HEP/HTC and EuroHPC/AI factories
 - co-design of new Tier-1 systems and ensure 'Lorentz' as national facility, collaborate to instigate change in EuroHPC (SURF, EuroCC, SHAKE)
- Balancing operations, support and research as algorithms (and AI) become more important
 - engage physics groups in software engineering (as per the SAC2025 report), more 'project' software engineering and computing/AI/ML training capacity in CT-PDP
- Risk of divergence WLCG from most 'others', even in our own domain, remains significant
 - invest in control points such as HPC CoEs and in AAI, policy, and 'OpSec'

And some resource-centric ones ... no realistic mitigation?

Weaknesses

- We provide a lot of research services, of all kinds, but disinvesting is not usually an option
 - we possible we can join up with others (SURF, OSG, other GWIs in the Commons, maybe NWO-I), but these in turn look back to us ...

Threats

- Lack of people on the CT-(PDP) side in the longer term: with FASTTRACK, MERGE, and AI/ML the demand keeps growing.
 - the requested 'matching' engineering resources, esp. FASTTRACK, need to become available in CT
- Too many SPOFs, esp. in collaboration infrastructure, 'open science', and policy activities
 - 'dakpan Jeff' needs to take care of anchoring our 'horizontal' knowledge as well
 - makes for a very challenging job indeed!

Opportunities and strength ... (no mitigation scenarios 😊)

Strength

- Impact in our chosen focus areas: Accelerated computing, Infrastructure, Trust & Identity,
→ we are seen and proactively asked to partake in projects, and seen as trusted and active partner

Opportunities

- Acceleration advances and ML help us exploit new systems
→ alignment with EuroHPC (and AI factories) for HEP/APP use cases as the EuroHPC Federation forms
- Design of next gen national Tier-1 ('Lorentz' ?) facility and (to some extent) AI factory
→ improve suitability for HEP use cases through design choices
- Cross-programme AI/ML activity raises profile of computing in our programmes
→ may get us the audience awareness for the importance of computing and algorithms
- Heterogeneous and energy-efficient computing
→ more 'physics output per unit of .*'
- National Plan for Research Digitalisation ('fiche')
→ if granted, we can leverage far broader ecosystem at the institute and partner universities!

Physics Data Processing *towards results unconstrained by computing*



Nikhef

David Groep
davidg@nikhef.nl

<https://www.nikhef.nl/~davidg/presentations/>
 <https://orcid.org/0000-0003-1026-6606>

