


Nikhef SAC April 2023

Physics Data Processing

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


David Groep



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, network & systems co-design R&D

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



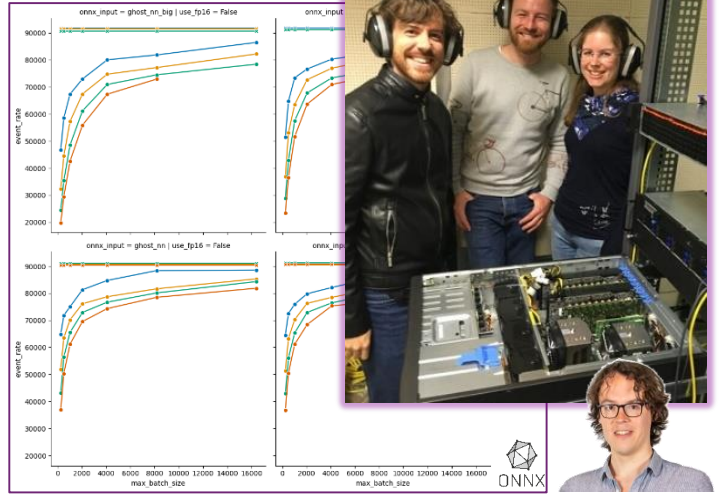
Infrastructure for trusted collaboration

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



Efficient computing and ‘accelerated results’

- FASTER: computing for HL-LHC & ‘4D’ reco
- LHCb’s Allen full-GPU trigger for HLT1
 - now adding HLT2 and CPU NN implementations
- NLeSC GPU acceleration in LHCb and parallel inference with ONNX+tensor ML libraries
- R&D roadmap for hybrid computing
 - link to infrastructure innovation & engineering with vendors
 - alternative architectures: non-x86, watercooling, GPU+FPGA, hybrid dies
 - ‘scaling and validation’, collaboration with computer science (SLICES-RI) & ML algorithms (@UM + RU)



For the long term: Quantum Computing algorithms exploration



- in collaboration with our experiments (notably LHCb and GW), QuSoft (CWI, UvA, *et al.*), SURF, and IBM (Zürich) via Maastricht
- 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). Graphs: Allen inference event rate vs batch size (NLeSC) <https://github.com/LHC-NLeSC/run-allen-run>

Infrastructure for Research

High-through compute (HTC) + HT storage

- National e-Infrastructure coordinated by SURF
- NL-T1, IGWN, KM3NeT, Xenon, DUNE + WeNMR, MinE, ...
- ‘Stoomboot’ local analysis facility + IGWN cluster & ‘submit node’ ~ 12 000 cores, 13 PByte storage installed, with pretty competitive TCO!

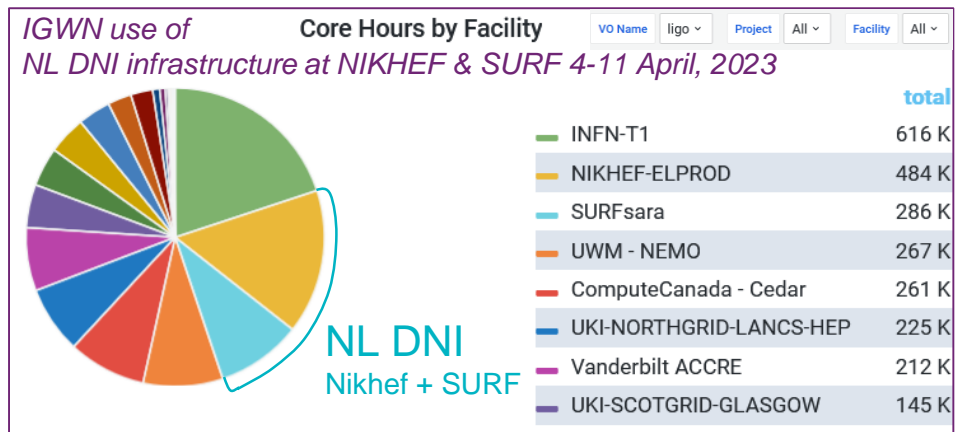


Occupancy: NDPF DNI processing facility. Faster processors (AMD Rome) allows processing in fewer cores using less power

Common solutions are essential for our 'national' facility

We promote alignment of general purpose e-Infrastructure and use by experiments (LHC, GW, KM3NeT, DUNE, Xenon, ...)

- *common solutions* since bespoke systems for each experiment do not scale for Nikhef (NL)
- needed for both efficient resource sharing and importantly: for our DevOps people



Continuation of our long-term strategy - from EU DataGrid onwards:


- drives collaborative efforts we work with for identity management and common protocols: ACCESS-CI and CILogon (US) – which are key players for LIGO, DUNE, and US-ATLAS
- common processing framework development (e.g. together with KM3NeT in its INFRADEV project)

Data on Tue April 11th 2023 from the OSG accounting for the LIGO VO for past week (with also SURFsara fully in production)

<https://gracc.opensciencegrid.org/d/9u1-Q3vVz/cpu-payload-jobs?orgId=1&var-ReportableVOName=ligo&var-Project=All&var-Facility=All&var-Probe=All&var-interval=1d&from=1680566400000&to=1681257600000>

Innovation *on* infrastructure

- Network – systems integration
- Storage throughput & DPUs
- Systems design and tuning



- early engineering engagement with vendors to build us suitable systems
- co-design of our national HPC systems ('Snellius')
- data-intensive compute with DPUs. Or on-NIC FPGAs?
- next-gen networks: >800Gbps (and 400G to CERN)




Contact Us

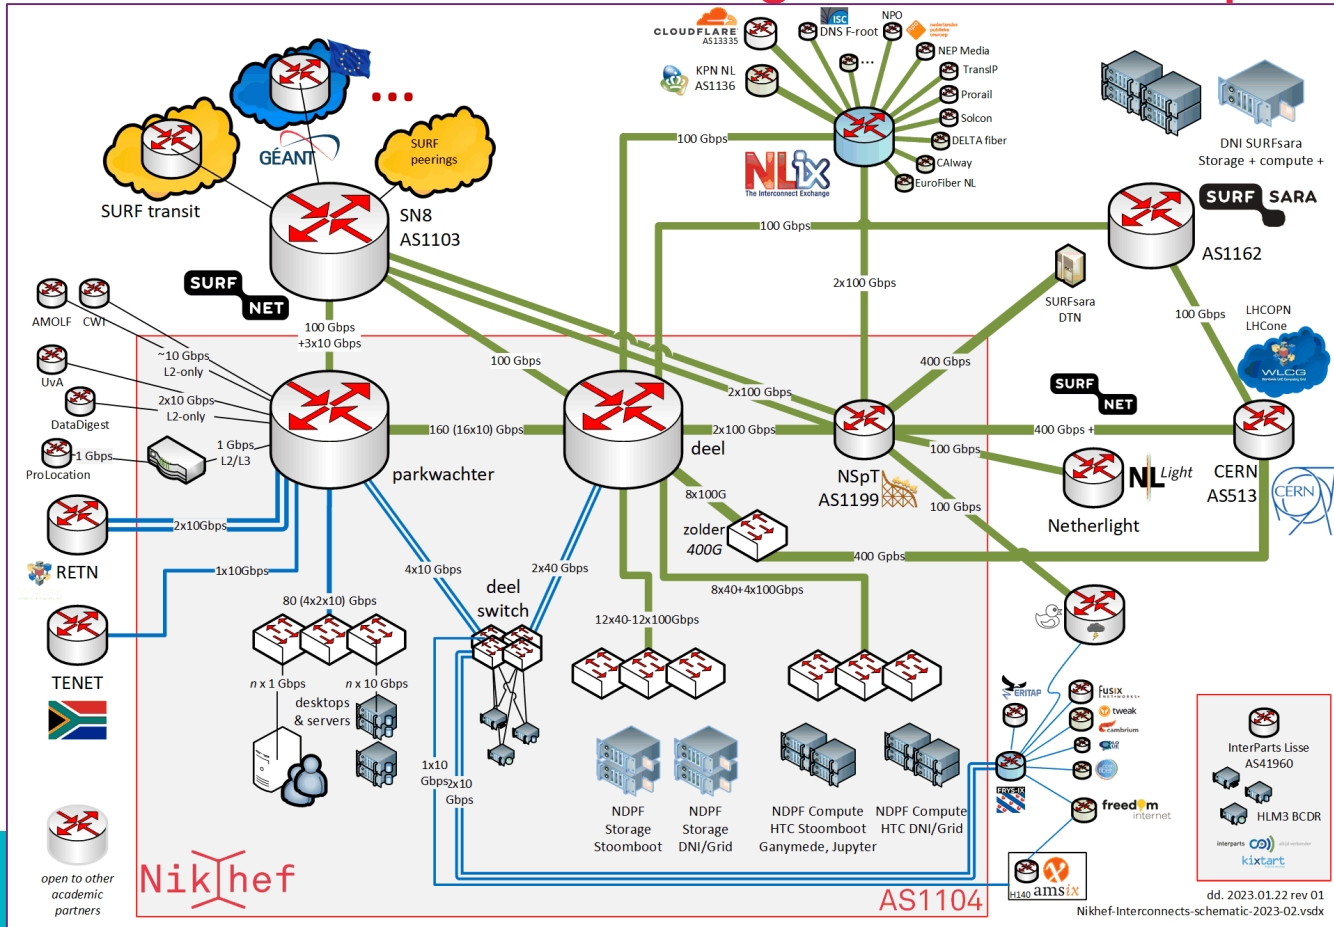
NIKHEF, SURF AND FUNGIBLE SET NEW BENCHMARK FOR THE WORLD'S FASTEST STORAGE PERFORMANCE

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>

Next Gen networks – connecting our data and people



Our science data flows are somebody else's DDoS attack



Het begon in 2018. Een bijzondere samenwerking tussen overheden.

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

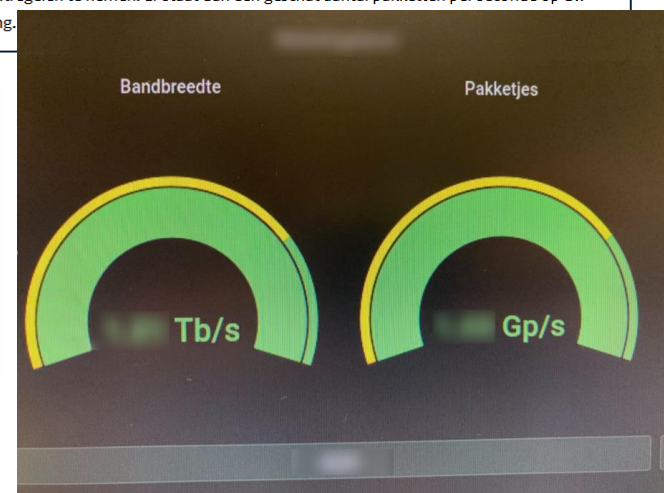

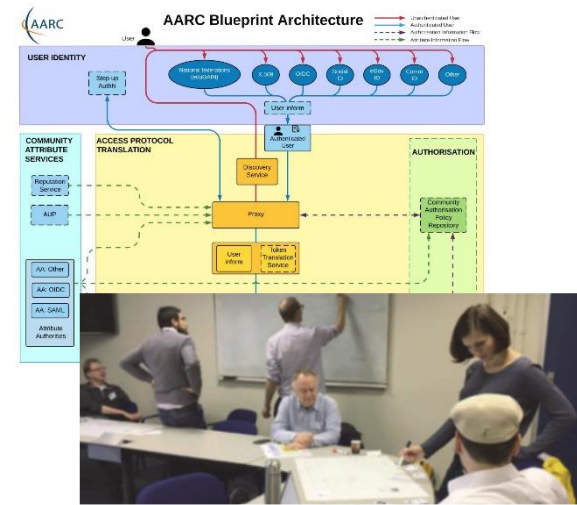


Image sources: belastingdienst.nl, rws.nl, nu.nl, werkentegennederland.nl

Infrastructure for Collaboration

Target high-impact specialized areas, bridging policy & technical architecture for 'AAI' & 'OpSec'

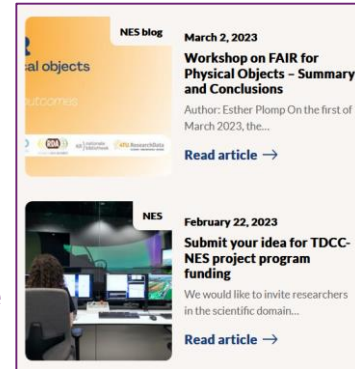
- authentication & authorization for research collaboration
 - **AARC project & community**, GEANT GN5, REFEDS & eduGAIN, EOSCFuture, TCS & RCauth.eu, ...
and planned: EOSC Core, AARC-TREE, EOSC-Q, ...
 - **policy for interoperability** for data protection, seamless service access, single-click acceptable use
 - continuous **technical evolution** driving IGWN, WLCG in line with the AARC BPA and global RIs
- embedding of data processing needs of our experiments in the EOSC landscape
 - **EOSC** Interoperability Framework
 - EOSC-A (AAI-TF), **EGI**, **GEANT community**



Collaboration: Research Data Management beyond 'FA'

FAIR for **live data and large volumes**, shifting gears towards the "I" and "R"

- not *that* many disciplines with really **voluminous data**
 - so nationally join forces with those who do: ASTRON (SRCnet), KNMI (earth observation, seismology)
- work with those who care about **software** to bring data to life:
 - NLeSC, 4TU.RD/TU Delft, CWI, and with those who ensure the *infrastructure*: SURF
- for our own analyses and the local (R&D) experiments we work towards *continuous deposition* of **re-usable** data and software
 - NWO Thematic Data Competence Centre for the Natural and Engineering Sciences
 - develop Djehuty RDM repository software link to 'Stoomboot' analysis cluster storage
 - research data management is 'just' part of a good science workflow



PDP strategic directions and their supporting projects

New initiatives and projects

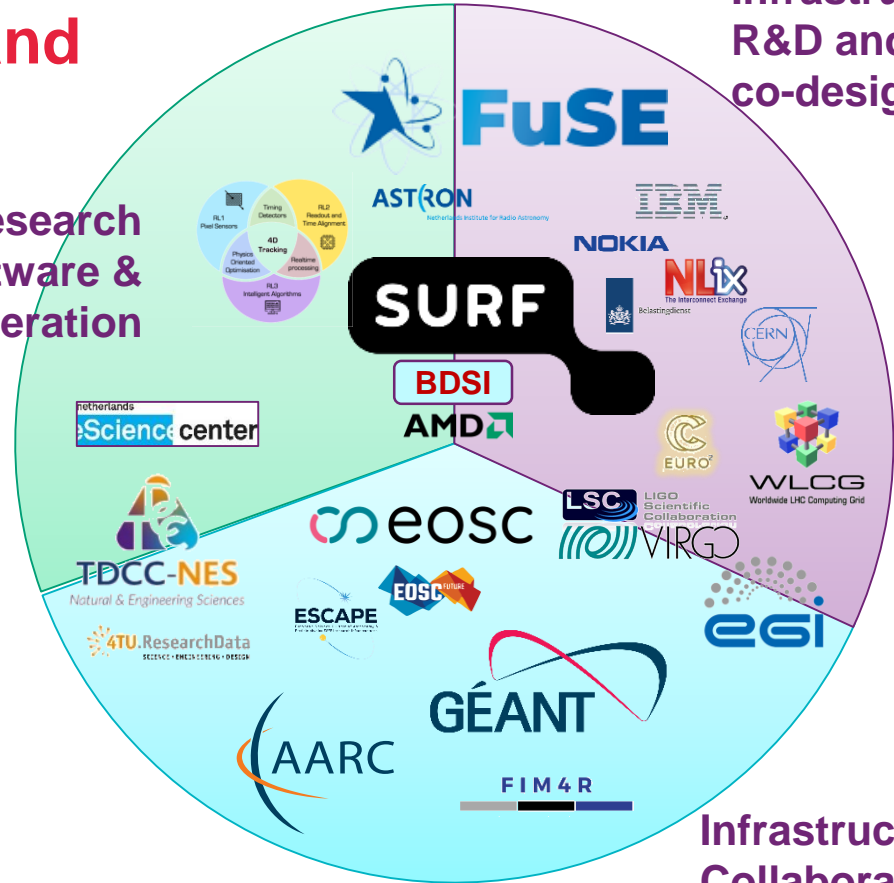
- strengthen the strategic areas
- ensure continuity of research and infrastructure

planned project pathways include GN5-*, AARC-TREE, EOSC-Q, EOSC Core, LHC4D Roadmap, ...

Public partner R&D engagement

AMD, Nokia, NVidia/MLNX, IBM, NL-ix, ...
Dutch national government


Research Software & Acceleration



Infrastructure
R&D and system
co-design

Infrastructure for
Collaboration

Because we can ... does not mean it's the scalable way 😊



LCO2 cooling of an AMD Ryzen Threadripper 3970X [56.38 °C] at 4600.1MHz processor (~1.5x nominal speed) sustained, using the Nikhef LCO2 test bench system (<https://hwbot.org/submission/4539341>) - (Krista de Roo en Tristan Suerink)




Nikhef

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>



Policy development in ICT and our collaborative values


Our Research Infra (and Open Science) needs a *collaborative* ‘values framework’

- frequently threatened by increasingly ‘corporate’ approach to ICT services
- continuous remedial action needed at many levels: from European Commission and EP, down to even our ‘own’ centralized national institutes organization ...

Continued vigilance on IT infra is part of PDP programme
to keep our research ICT infrastructure open, e.g. through:

- reviewed & improved NIS2 directive, with our Opinion taken up by the EP;
- watch carefully and continuously NWO’s push to ‘corporate IT’ and its impact on academic integrity & freedom;
- promote trustworthy federated access via AARC (globally), SRAM (our SURF national scheme), promote trust & identity with our university partners
- build scalable security solutions rather than pay corporate ‘ISO tick-box’ providers

luckily many times in collaboration with GEANT, JISC, STFC, EGI, and our peer institutes in NL



<https://doi.org/10.5281/zenodo.4629136>