Nikhef

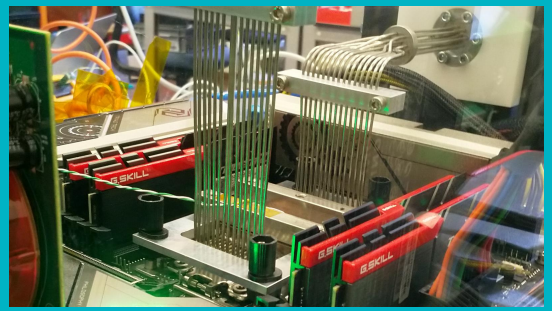

Nikhef SAC June 2025

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, algorithms & high-performance processors
- design patterns for workflow & data orchestration

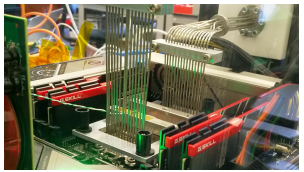

Infrastructure, network & systems co-design R&D

- 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

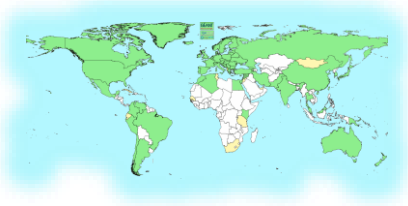

Ambitions and execution

- Accelerate time to results with **efficient use of computing**
- exploit **new systems design and integration:** different architectures, innovative vendors, faster networks
- ensure our users **access to the broadest range of services** with ease

Meanwhile

- ensure our computing capacity is sustainable
  we fund it to a large extent from generic, non-subatomic-physics sources
- continue push for generic access of services
  important to keep individual RIs like WLCG in line with (access to) generic infrastructures
  divergence puts our NL-T1 funding model (2.7M€ per year for WLCG alone) in jeopardy
- move more of our experiments to generic computing infrastructure so we can contribute
  more effectively, cheaper and faster, like we invested in IGWN and KM3NeT

Some computing, engineering and trust highlights

- significantly increased
  use of GPUs
  not only with Allen
  but also in our analysis facility

From single GPU demo to **500 node LHCb HLT1** makes tracking info directly available for classification

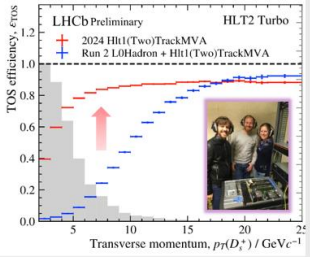

Now we get much higher efficiency (blue → red points) for *e.g.* D(s)→KKπ charm physics

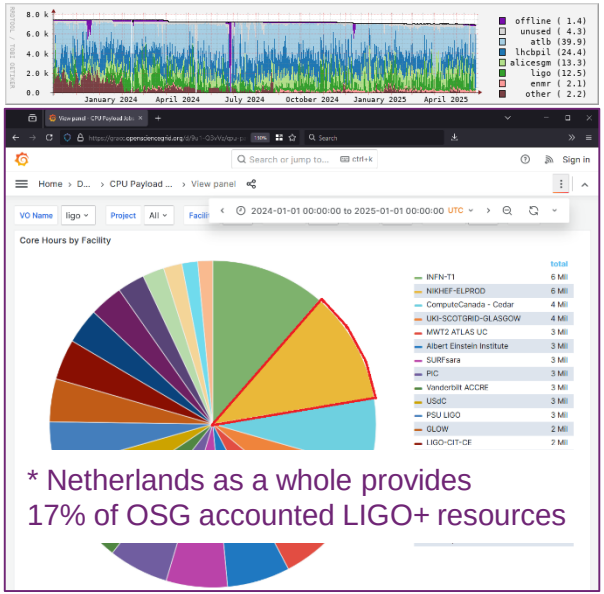

- **94.3% occupancy** of our NL-federated Tier-1 facility
- Nikhef **compute for IWGN now >11%** of global capacity*
- **800 Gbps coherent** long-distance AMS-GVA network
- Security and risk management **for EOSC EU Node**
- new **AARC architecture and trust framework**
  with tokens for more than just IGWN and WLCG

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 2024: Nikhef 11.1% + SURF 6.1%
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=1704067200000&to=1735689600000&viewPanel=9

ATLAS T0 export repeat

- T0 export rates are the most important use case and were not achieved
- The rates weren't achieved because they were queued behind production
  - T2 traffic is non negligible in ATLAS (42% dst, 25% src)
- Tests were repeated injecting one site at the time
  - Rates improved for the majority of sites
- Some differences:
  - SARA was testing 800 Gb/s after DC24; was injected with much larger rates
  - RAL wanted to test writing directly to tape in the second test; other limitations were identified
  - NDGF resolved the dcache bug that was affecting them

Site	T0 Export	DC24 best rates on day 1,2	% of expected rates	T0-T1 one T1at the time	% of expected rates
BNL-ATLAS	60	31.5	53%	61.3	102%
FZK-LCG2	32	26.4	83%	42.2	132%
IN2P3-CC	38	43	113%	50.9	134%
INFN-T1	23	19.3	84%		
NDGF-T1	15	13.8	92%		
SARA-MATRIX	15	12.2	81%		
pic	11	12.3	112%		
RAL-LCG2	38	15	39%		
TRIUMF-LCG2	25	23.9	96%	27.2	109%
T1 summary	257	197.4	77%	562.7	219%
T1 summary -SARA	242	185.2	77%	288.6	119%

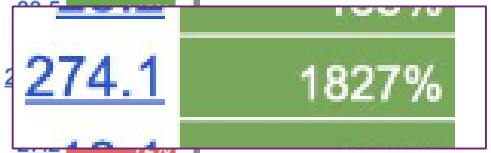

T1 summary	257	197.4	77%	562.7	219%
T1 summary -SARA	242	185.2	77%	288.6	119%

LHCONE/LHCOPN meeting, April 2024

25

800Gbps in practice

From: Katy Ellis (STFC RAL): DC24 report, LHCOPN meeting, Catania, IT, April 2024 https://indico.cern.ch/event/1349135/#6-dc24-report

A compact programme backed by computing technology

People: 2.1 FTE (spread over 3 staff)

- *our one postdoc got a permanent position (LHCb) at Maastricht University since July 1st*

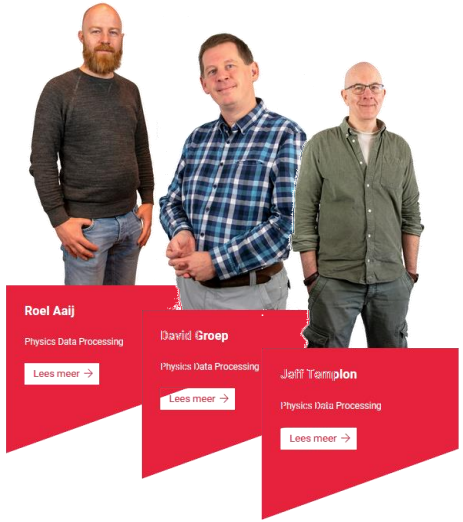

But: rely can on ~ 6.5 FTE from Nikhef Computing Technology

~ 3 FTE (nominally) infrastructure operations
for both our Tier-1 (co-funded by SURF) and Tier-3
while also providing researcher engagement & support

~ 1.5 FTE infrastructure for collaboration and cross-domain security
(funded by GEANT, EGI, and other EC contracts)

~ 2 FTE innovation, accelerated processing and next-gen networks
(mostly funded by SURF), doubling with support for advanced infra operations

Activities on accelerated computation, ML (and QC) spread across all programmes, with PDP participation, to ensure proper sustained embedding of their results in the programmes

Pathways via projects and collaborations

New initiatives and projects

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

future project pathways planned:
FASTTRACK, AI/ML Upgrade, …

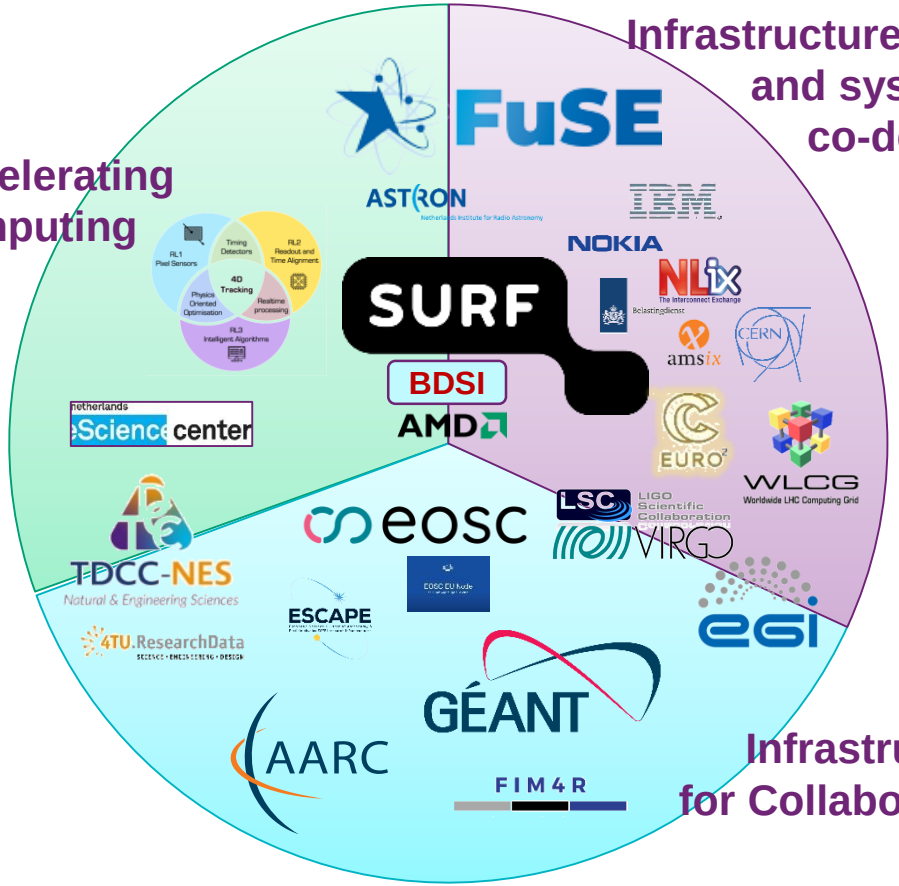

Public partner R&D engagement
AMD, Nokia, NVidia, NL-ix, AmsIX…
Dutch national government

Next to operations and innovation, we push ahead!

Drive change in analysis and re-processing with accelerators to become ‘the default’:

- we have **GPU accelerator expertise** with Allen for LHCb’s … on to a FASTER FASTTRACK
- there is already large embedded **effort on AI & ML** across the Nikhef portfolio
- as a T1 have **fast access to data**, and have a GPU-heavy national HPC system at SURF
- strong and coherent CompSys **research in CS**, and the **AI coalition across Dutch academia**

Push the high-throughput system integration boundaries

- **dedicated ‘lab’ data centre** for Tbps+ networks, energy & cooling, novel architectures
- excellent links with SURF, public sector, and vendor ecosystem

Make collaboration easier across domains

- keep our field connected to **AAI in R&E space**, evolving our federated access
- building on the foundations we brought and bring to eduGAIN/GEANT, EOSC, EGI, SURF, and the **trust and identity federations**
- executable papers, research workflows, **‘beyond’-Open-Science workflows**, Open Scholarly Infrastructure

Coherence and alignment – within and outside

Data processing across all of Nikhef

- Per-programme algorithmic collaboration
  from compiler flags to restructuring algorithms
- Computing Office Hours
  from printing to vectorisation, all in one go

… and across the national and global landscape

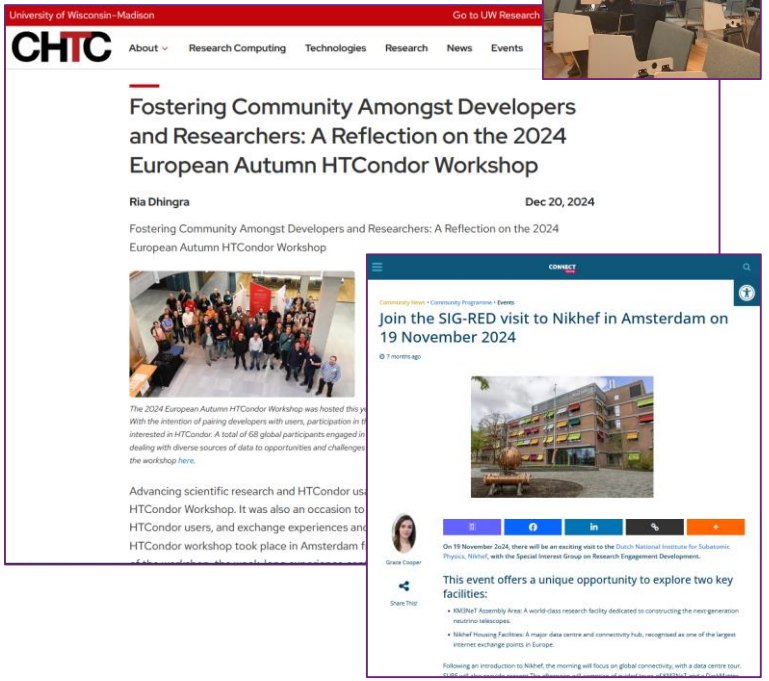

accelerating "time to science" through computing and collaboration

Nikhef

Challenges and opportunities

Computing operations and funding for capacity are a continuous challenge

- energy and personnel costs are dominating factors
- ‘AI bubble’ actually works against us for hardware pricing
- Dutch funding model with emphasis on short-term competition assumes operations are ‘just’ funded from institutional basis
- reliance on generic funding means we *must* remain aligned!

But we can also help and make an impact at the same time

- large users like us improve efficiency of whole ecosystem
- acceleration advances and ML help us exploit new systems, and allow us to do more with the resources we do get
- innovation programme aids on procuring national systems that fit also our future needs

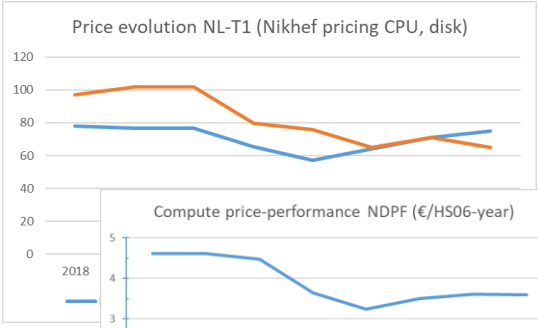

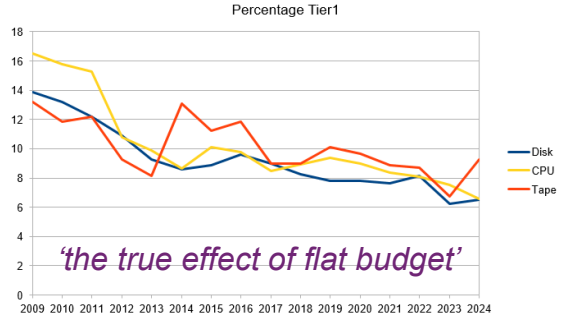

Nikhef Data Processing Facility TCO data. Percentage (WLCG) Tier-1 data: SURF, 2024

Continuity for capacity and an ‘AI/ML’ upgrade

Opportunities can actually be found in funding landscape

- Large-Scale Research Infrastructure upgrade call specifically mentions as eligible examples
  ‘joint e-Infrastructure’ and ‘access to and development of the latest AI and data processing techniques’
- we can build on BIGLHC (2013) and KM3NeT 2.0 (2019)
- GPU heavy extensions both for AI/ML and accelerated compute
- ensuring all types of data are available for training (and validation)

Examples of upgrade investments:
Upgrade investment can enable (beyond) state-of-the-art scientific research, increase the LSRI user base and bring about collaboration between existing LSRIs, resulting in wider impact and economies of scale. Contributions to greening of an LSRI can be part of the investment. For example, by integrating greening into the upgrade investment or by partially focusing it on greening. Examples of upgrade investments are investments in:

- development and/or procurement/construction of the latest state-of-the-art equipment for 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 LSRIs, connecting new user groups and/or setting up a joint e-infrastructure;
- Modifications to existing LSRI needed to reduce its carbon footprint.

which can complement other funding opportunities that are available to Nikhef for AI/ML

Physics Data Processing

towards results unconstrained by computing

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

Nikhef

Pushing computing boundaries, some examples

From single GPU demo to **500 node LHCb HLT1** makes tracking info directly available for classification

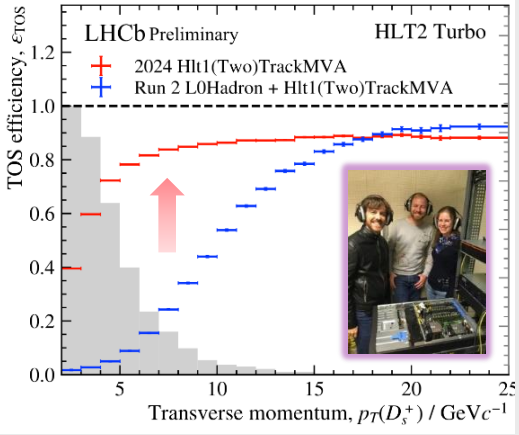

Now we get much higher efficiency (blue → red points) for *e.g.* D(s)→KKπ charm physics

Providing Security and Risk Management for the **European Open Science Cloud** "EU Node" core services (procured EC service, delivered with EGI)

*and next-gen trust & identity for LSRIs with **AARC TREE** and **eduGAIN** security*

AARC eduGAIN

Core Hours by Facility

LIGO (+ Virgo) global distributed compute all of 2024

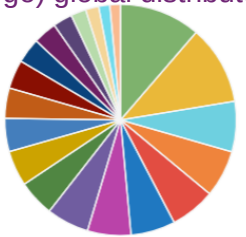

	total
INFN-T1	6 Mil
NIKHEF-ELPROD	6 Mil
ComputeCanada - Cedar	4 Mil
UKI-SCOTGRID-GLASGOW	4 Mil
MWT2 ATLAS UC	3 Mil
Albert Einstein Institute	3 Mil
SURFsara	3 Mil
PIC	3 Mil

Nikhef Data Processing Facility in our federated Dutch e-Infrastructure: WLCG Dutch Tier 1, Auger, KM3NeT gravitational waves … but also *e.g.* WeNMR structural biochemistry

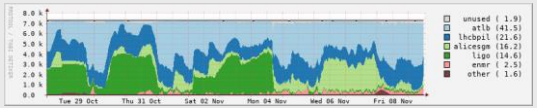

Image sources: EOSC EU Node https://open-science-cloud.ec.europa.eu/; IGWN utilisation from OSG GRACIA https://gracc.opensciencegrid.org/ and https://edu.nl/tr79b; LHCb HLT1: Roel Aaij (Nikhef), Daniel Campora (Nikhef & UM), Dorothea vom Bruch (LPNHE), photo: LPNHE; Run3 to Run2 comparison: https://lbfence.cern.ch/alcm/public/figure/details/3837; EOSC AAI Architecture: publications office of the European Union; AARC blueprint: https://aarc-community.org/; DNI: https://www.nikhef.nl/pdp/doc/stats/

Jumping over our computing boundaries

Efficient and accelerated computing

- **accelerators**, such as GPU- and hybrid-computing
  – *extending its scope to last trigger stage, to off-line processing and to analysis*
- more **machine learning,** for tracking and physics event classification
- **efficient architectures**: ARM systems (with our SURF Experimental Platform)
- exploration of **quantum computing**, for future gains in algorithmic complexity

Networking, both physical link and collaborative networks

- network bandwidths ***over* 800 Gbps to CERN**, and to more destinations
- evolving and **aligning global collaboration** technologies (unified tokens): authentication & authorization for WLCG, Gravitational Waves, SKA, DUNE, … and the European Open Science Cloud

Open Science at scale, with data for today's and tomorrow's science

- new forms of service delivery with **Analysis Facilities** and 'data lakes'
- **practical Open Science** for data-intensive systems (with 4TU.RD/TUDelft)
  – with a new Institutional Data Repository using Djuhuty open source software, **linked to our local analysis facility**, adding re-usable science software

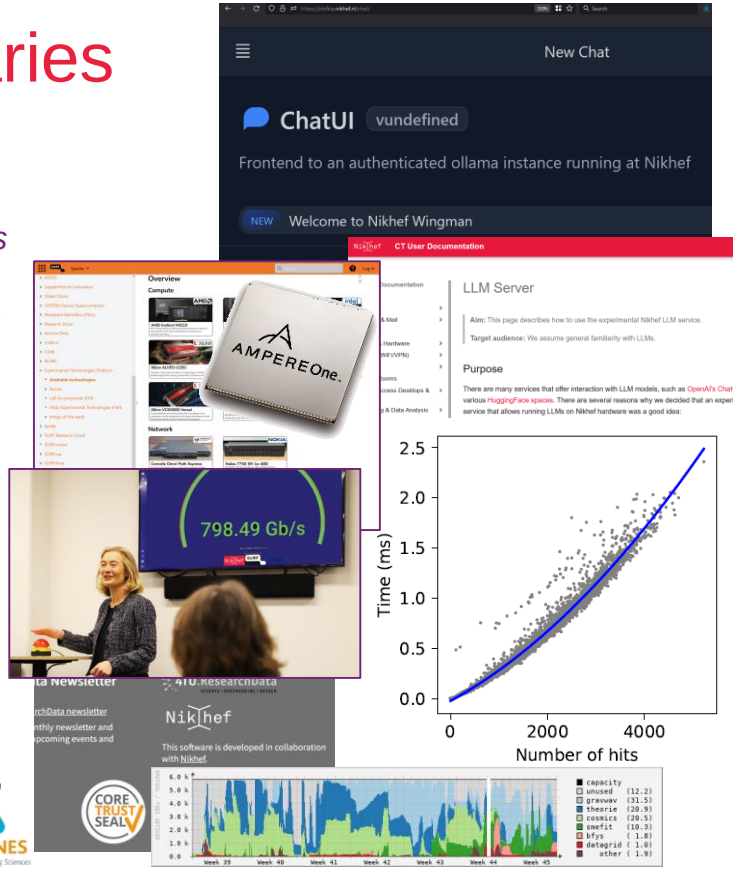

LLM service: https://plofkip.nikhef.nl/chat, services: kb.nikhef.nl/ct, AmpereOne SURF ETP ARM experimental platform pauperbak.nikhef.nl; scatter plot shows compute time per event estimated for ATLAS at the HL-LHC with the current processing model (from: FASTER ENW-XL); Djehuty open source Data repository: archive.nikhef.nl, data.4tu.nl, Catharina Vaendel (Nikhef) and Roel Jansen et al. (TUDelft).

Our science data flows … are somebody else's DDoS attack

Bandwidth for the HL-LHC, WLCG, and the trigger data streams can also 'exercise' our Dutch national critical infrastructure

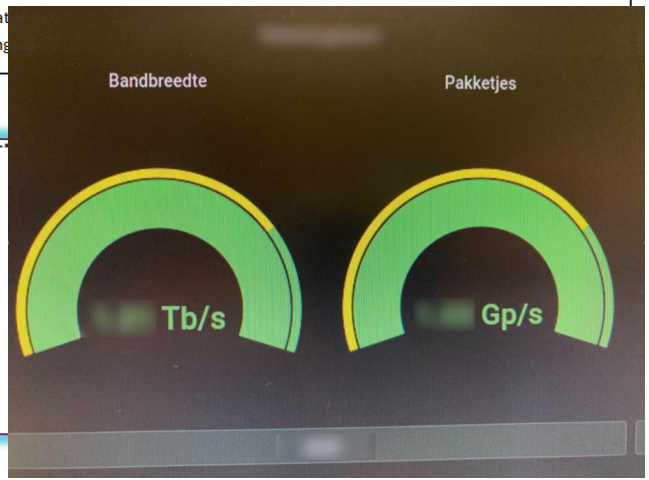

Image sources: belastingdienst.nl, rws.nl, werkentegennederland.nl, https://www.nomoreddos.org/deelnemers/

Infrastructure for Research

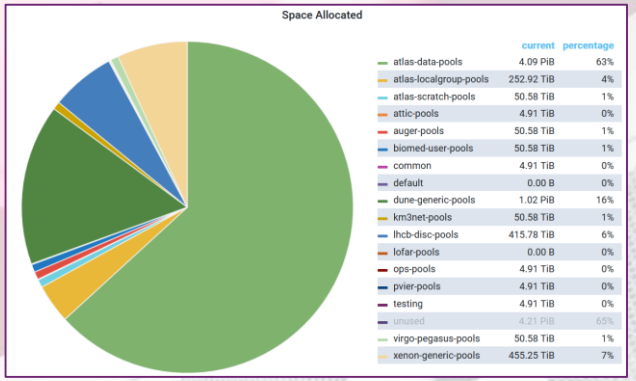

High-through compute (HTC) + high throughput storage

- provided by the national e-infrastructure
- coordinated by SURF
- WLCG NL-T1, IGWN, KM3NeT, Xenon, DUNE
  + other science domains WeNMR, Tropomi, project-MinE, …
- ‘Stoomboot’ local analysis facility

~ 11 000 cores, 16 PByte storage installed, with pretty competitive TCO

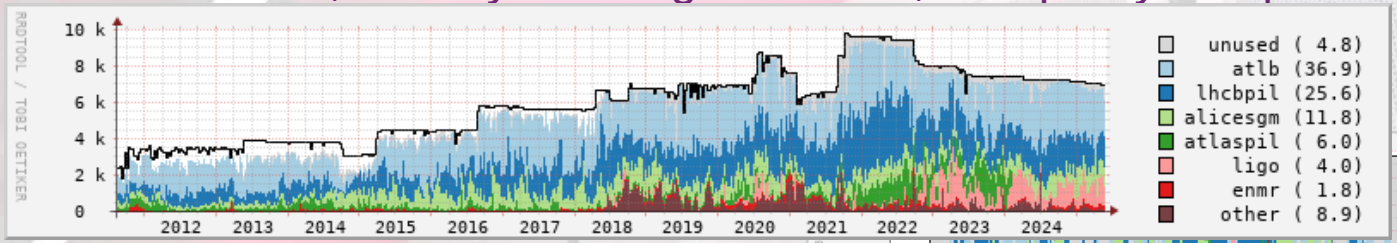

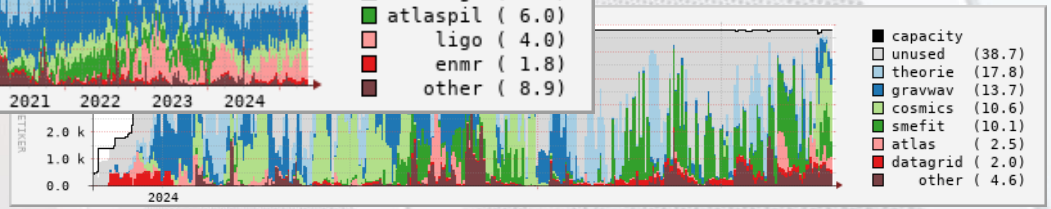

Occupancy: NDPF DNI processing facility. Faster processors (AMD Rome and Bergamo) allows processing in fewer cores and using less power
Storage data: https://steker.nikhef.nl:3000/d/c8fA5P0Wk/dcache-grid?orgId=1&refresh=5m&viewPanel=37 as of May 25, 2025

800G link Nikhef & SURF to CERN

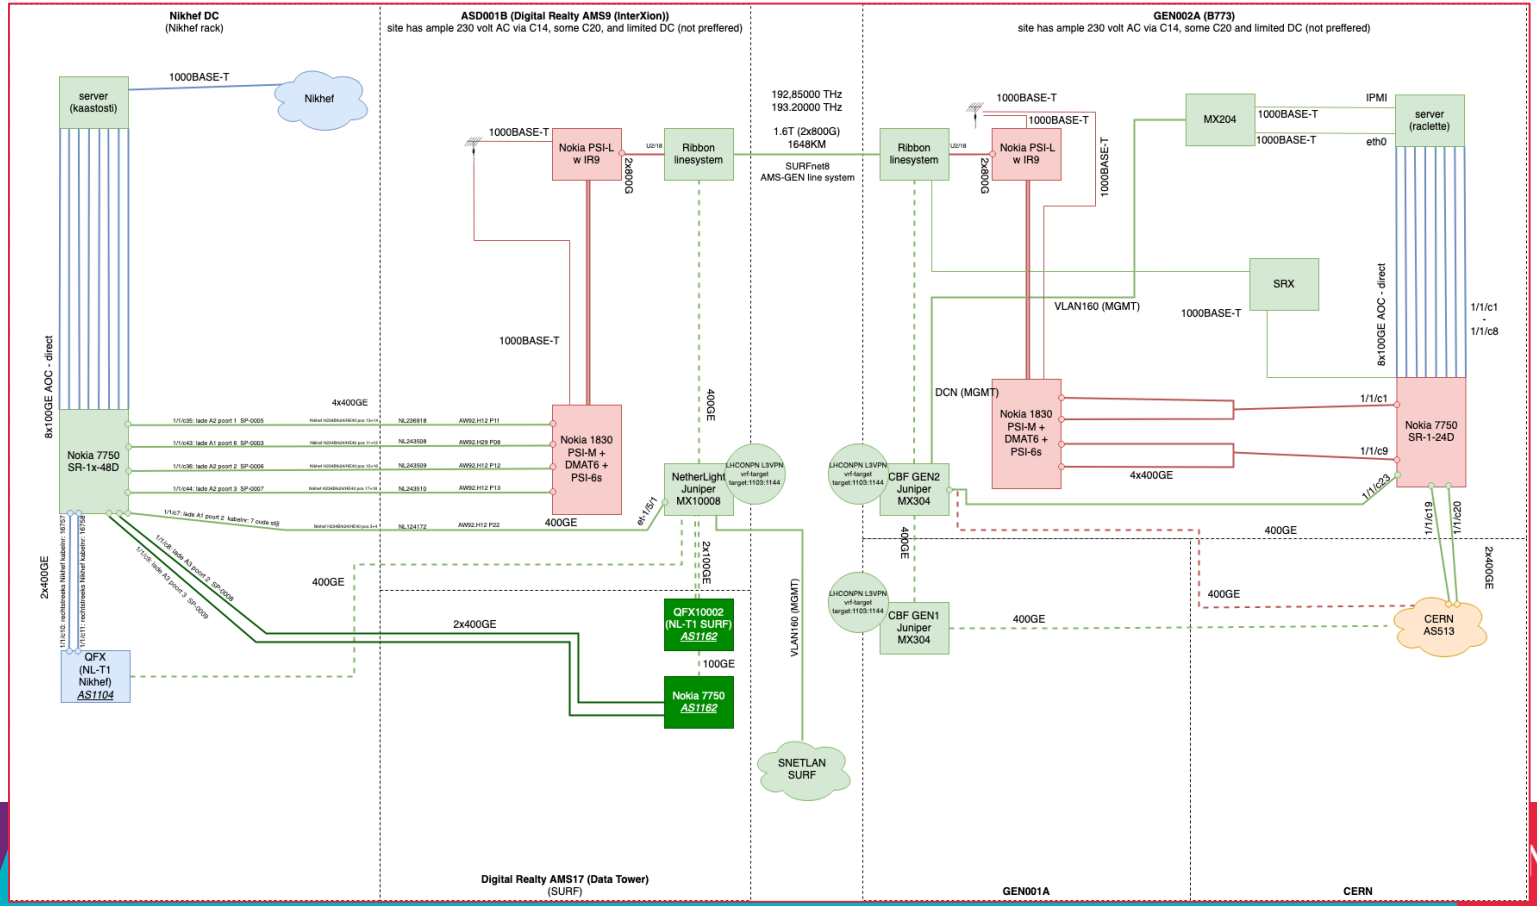

Next Gen networks – connecting our data and people

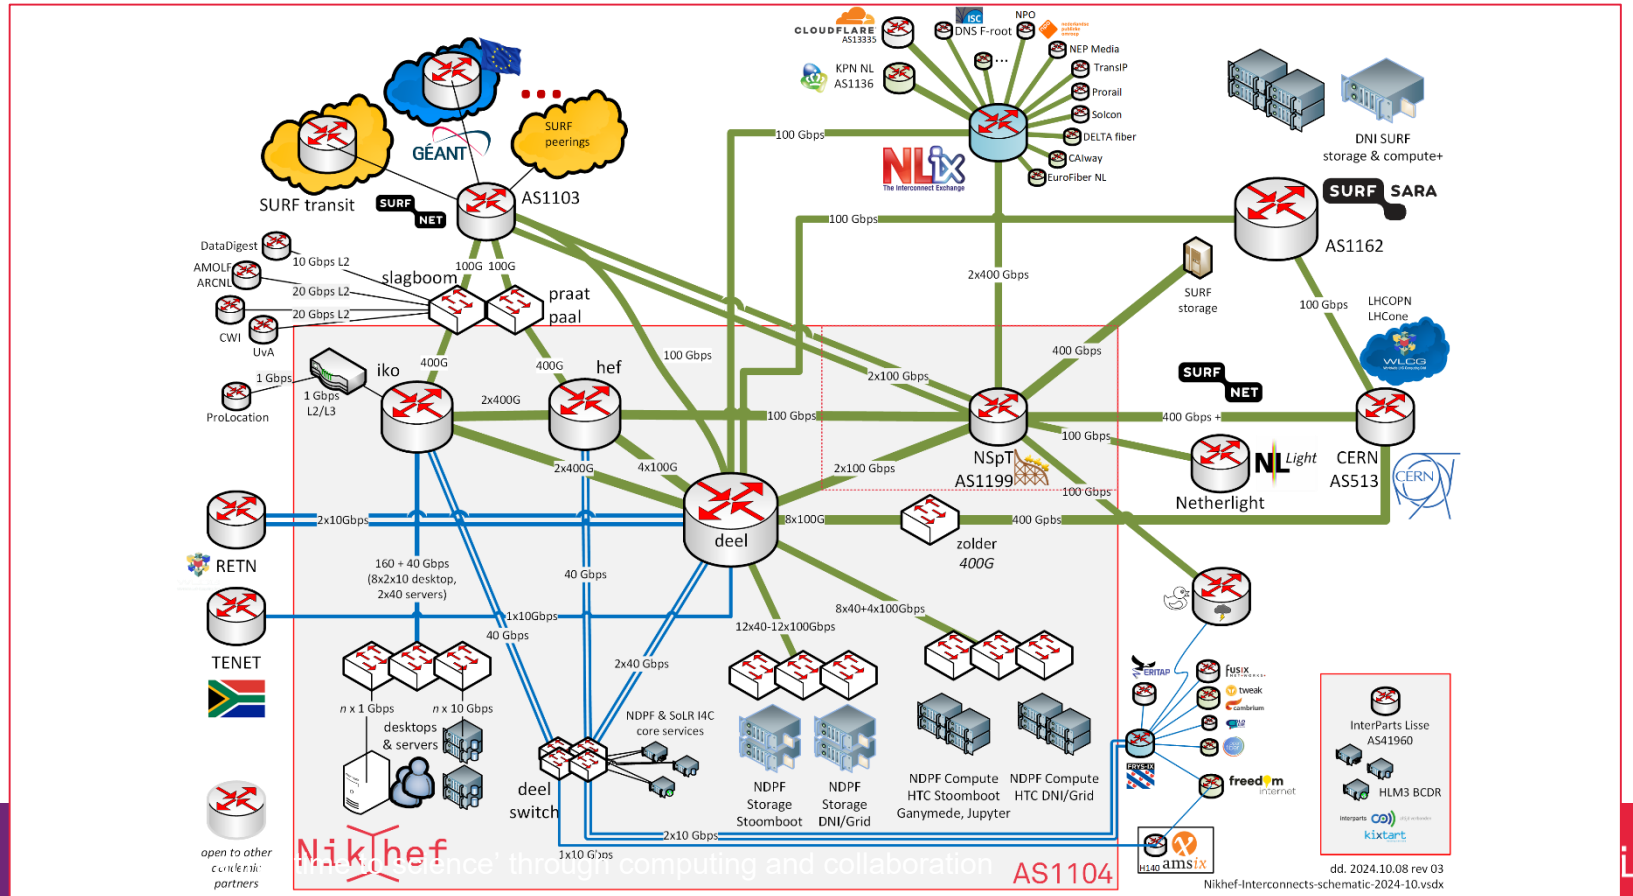

Infrastructure for Collaboration

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

- authentication & authorization for research collaboration
  - **AARC project & community**, GEANT GN5, REFEDS & eduGAIN, TCS & RCauth.eu, AARC TREE, EOSC
    policy for interoperability for data protection, seamless service access, single-click acceptable use
  - continuous **technical evolution** in line with the AARC BPA and global RIs
- embedding of data processing infrastructure of our programmes in the EOSC landscape

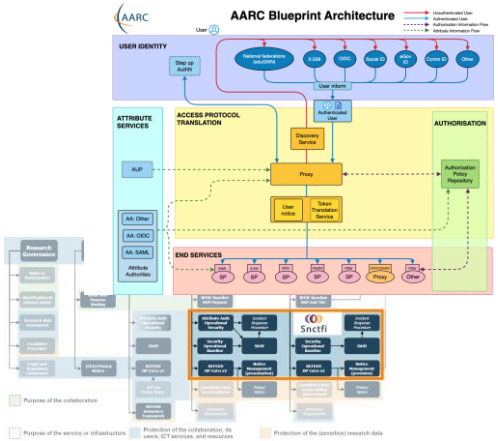

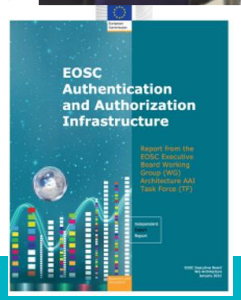

Collaboration: Research Data Management beyond ‘FA’

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/TUDelft, 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

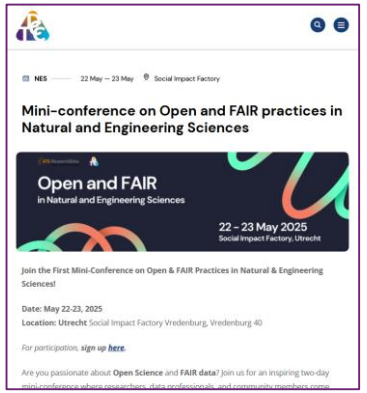

Foundational Principles for Digitalisation at Nikhef

Still needs explanation – despite digital sovereignty discussions. For both research and ‘enterprise’ ICT.
And ‘if they are many, they are not principles’, so we set only four:

1. Institutional strategy and mission directs ICT decision making
   ICT decisions are assessed based on the Nikhef strategic themes
2. Collaboration as a core value
   Nikhef stands for the whole of the Dutch community in (astro)particle physics and its European and global collaborations
3. Shared public values and responsible technology
   Nikhef employs, develops, and shapes technologies that preserve autonomy, justice and humanity, that builds on our academic sovereignty and integrity
4. Digitalisation reflects the continuity in our research programmes
   With research horizons measured in decades, ICT reflect this continuity in its choice of infrastructure, services, and information management, and in its human expertise

Policy development in ICT and our collaborative values

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

- frequently threated 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

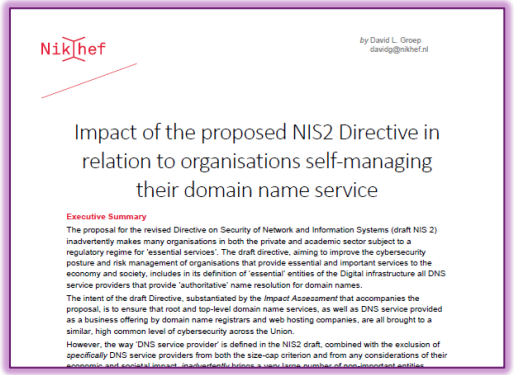
Nikhef

by David L. Groep
davidg@nikhef.nl

Impact of the proposed NIS2 Directive in relation to organisations self-managing their domain name service

Executive Summary

The proposal for the revised Directive on Security of Network and Information Systems (draft NIS 2) inadvertently makes many organisations in both the private and academic sector subject to a regulatory regime for 'essential services'. The draft directive, aiming to improve the cybersecurity posture and risk management of organisations that provide essential and important services to the economy and society, includes in its definition of 'essential' entities of the Digital infrastructure all DNS service providers that provide 'authoritative' name resolution for domain names.

The intent of the draft Directive, substantiated by the *Impact Assessment* that accompanies the proposal, is to ensure that root and top-level domain name services, as well as DNS service provided as a business offering by domain name registrars and web hosting companies, are all brought to a similar, high common level of cybersecurity across the Union.

However, the way 'DNS service provider' is defined in the NIS2 draft, combined with the exclusion of *specifically* DNS service providers from both the size-cap criterion and from any considerations of their

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

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

Where do we plan to go from here?

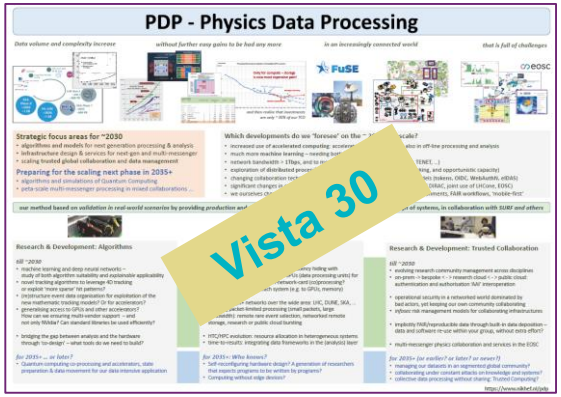

Strategic focus areas for ~2030

- algorithm design patterns for processing & analysis
- infrastructure design for multi-messenger
- trusted global collaboration and data management

on this ‘near term’ timescale

- accelerators (such as GPUs and hybrid computing) – also in off-line processing and analysis
- *much* more machine learning – needing both learning and inference
- network bandwidths > 1Tbps, and to more destinations (SURF, CERN, ESnet, TENET, …)
- exploration of distributed processing (integrating trigger, processing, networking, and opportunistic capacity)
- changing collaboration technologies: industry augmenting our bespoke models
- significant changes in service delivery (‘data lakes’, multi-user workflow systems, science networks, EOSC)

and in due course, we ourselves and as a community will change

executable papers, ‘notebooks’ for all Research Environments, ‘beyond’-Open-Science workflows, ‘mobile-first’ research computing, *and who knows what’s more to come?*

Because we can … does not mean it’s the scalable way ☺

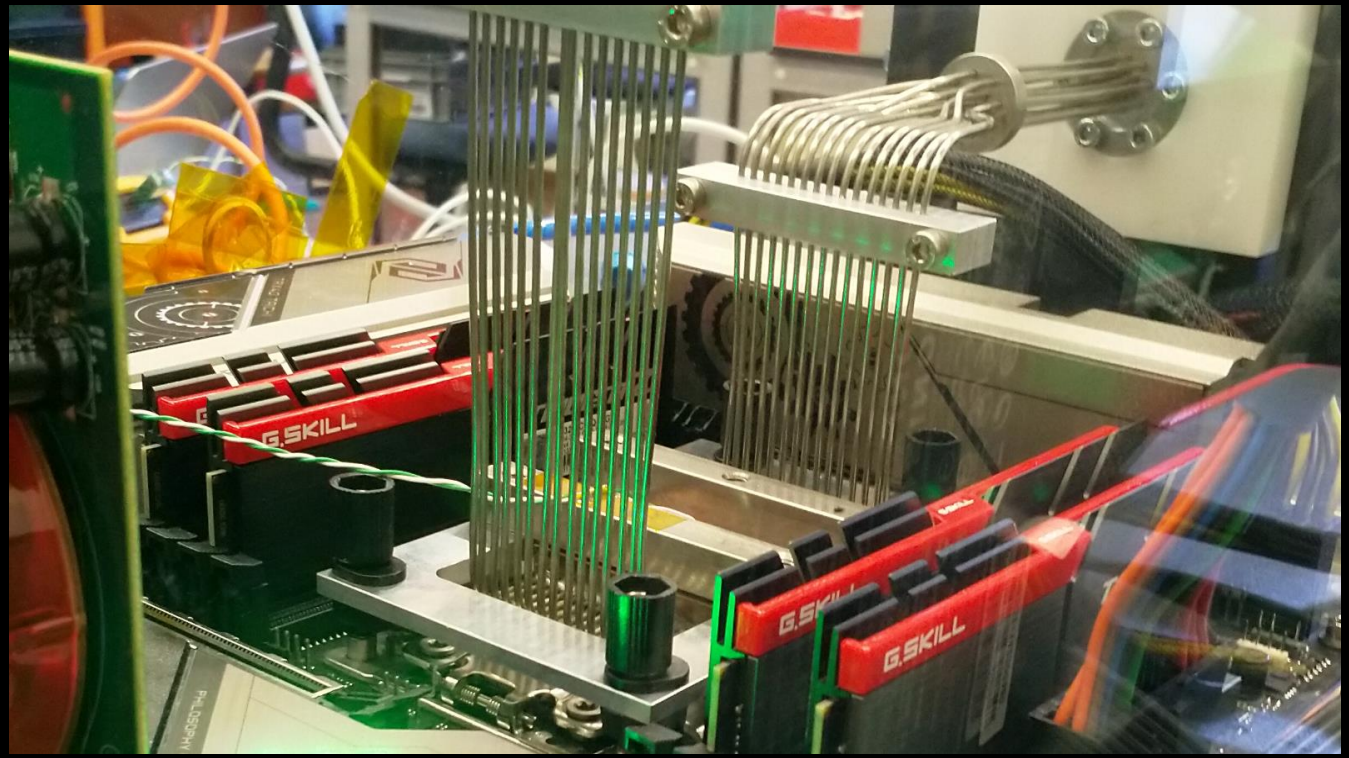

LCO2 cooling of an AMD Ryzen Threadripper 3970X [56.38 °C] at 4600.1MHz processor (~1.5x nominal speed) sustained, using the Nikhef LCO2 test bench system (https://hwbot.org/submission/4539341) - (Krista de Roo en Tristan Suerink)