

# NIKHEF OVERVIEW

DENNIS VAN DOK

HEPiX autumn 2023 Workshop, University of Victoria,  
Tuesday 2023-10-17

# ABOUT NIKHEF

- long term partner in the HEPiX community
- traditionally in accelerator physics, but also:
- neutrinos (KM3NeT)
- Dark matter (XENON)
- Cosmics (Auger)
- Gravitational waves (LIGO-VIRGO-KAGRA)
- also theory, eEDM

# ORGANISATION

- ~243 scientific staff (96 permanent, 147 PhD candidates and Postdocs)
- ~122 technical and support staff

(source: <https://www.nikhef.nl/en/facts-figures-2022/personnel-2022/>)

Nikhef is a partnership between the Institutes  
Organisation of NWO and six universities:

Nik|hef

Radboud University



rijksuniversiteit  
 groningen



UNIVERSITY OF AMSTERDAM



Maastricht University



Instituten



Utrecht University

VU



VRIJE  
UNIVERSITEIT  
AMSTERDAM

# DATA CENTRE HOUSING

The early involvement with the internet puts Nikhef right on one of the oldest exchanges in the world.

Currently being extended due to popular demand.

See [talk by Floris Bieshaar](#) later this week.

# HIGH THROUGHPUT COMPUTING

- Running a modest facility 'at scale'.
- Limited to 400 kW total power because of air cooling capacity<sup>\*</sup>

<sup>\*</sup> But see [this slide on computing innovations](#)



# THE DUTCH NATIONAL INFRASTRUCTURE

- Together with SURF (SARA), we provide compute power and storage to projects that have a dutch partner.
- NL-T1 Jointly with SURF
- SURF provides funding for personnel and equipment.




# CLUSTERS

- Still running Torque, but thinking very hard now about a transition to HTCondor.
- Under investigation: running all payloads in containers to support more platforms
- All our current systems (< 5y old) have single socket AMD CPUs.
- Three ARC CEs on the front-end

# STORAGE

- Happily running dCache for several years now.
- Latest addition was a 'bargain' 6PB in 6 Lenovo SR655/D3284 84x16TB (1.3 PB raw, 1.0 PB net)
- 3 installed in 2022, 3 more in May 2023.



Drawer 1 (Top)

Drawer 1 status LEDs

Enclosure status LEDs

3x Drive rows in each drawer (14x 3.5-inch SAS hot-swap drives in each row)

2x Hot-swap sideplanes (Left and Right) on each drawer (Right sideplanes shown)

Drawer 2 status LEDs

Drawer 2 (Bottom)

# CONFIGURATION MANAGEMENT

- Transition to Saltstack for stateful management since 2016.
- See [presentation by Andrew Pickford](#) on Thursday

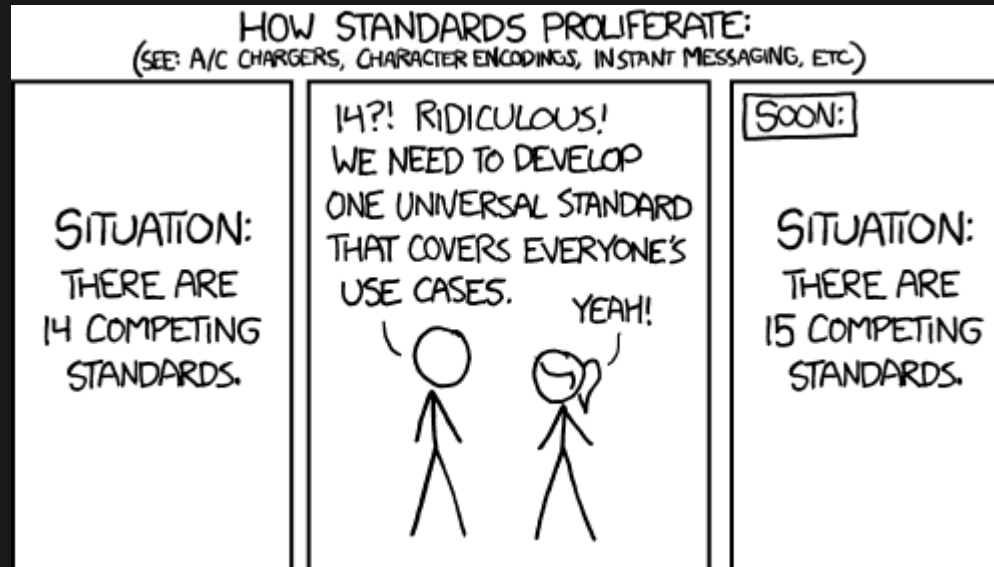
# TOKEN TRANSITIONS

There are too many incompatible token profiles doing the rounds today.

- Scitokens profile
- WLCG profile
- EGI/AARC profile

Everybody agrees it would be a good idea to introduce:

*The GUT (Grand Unified Token) profile*



(source: <https://xkcd.com/927/>)

# PROVIDING NEUTRAL TERRAIN

- We don't have a special interest in pushing any token technology.
- But we do like to talk to everyone.
- The main goal is interoperability.
- So we also advocate for a mapping plug-in standard.



Leading this effort is Mischa Sallé.

# DDOSSING OUR GOVERNMENT (WITH PERMISSION)



# JOINT EXERCISE BETWEEN GOVERNMENT, IT PROVIDERS, ACADEMIA

- High bandwidth enables use for 'other' purposes
- Integrated exercise between government departments, IT providers and academia (the red team)
- Actually DDOSsing the live systems (during night time to have minimal disruption)
- Regularly repeated to improve resilience over time
- ...but attacks get more sophisticated, too



# INNOVATION

Several new(ish) developments were tested to see how they suit our needs.

# STORAGE

- Tested Fungible (no longer on offer) with NVMe over TCP
- see <https://www.nikhef.nl/en/news/22705/>
- Managed 6.55 MIOPS.

## STORAGE II

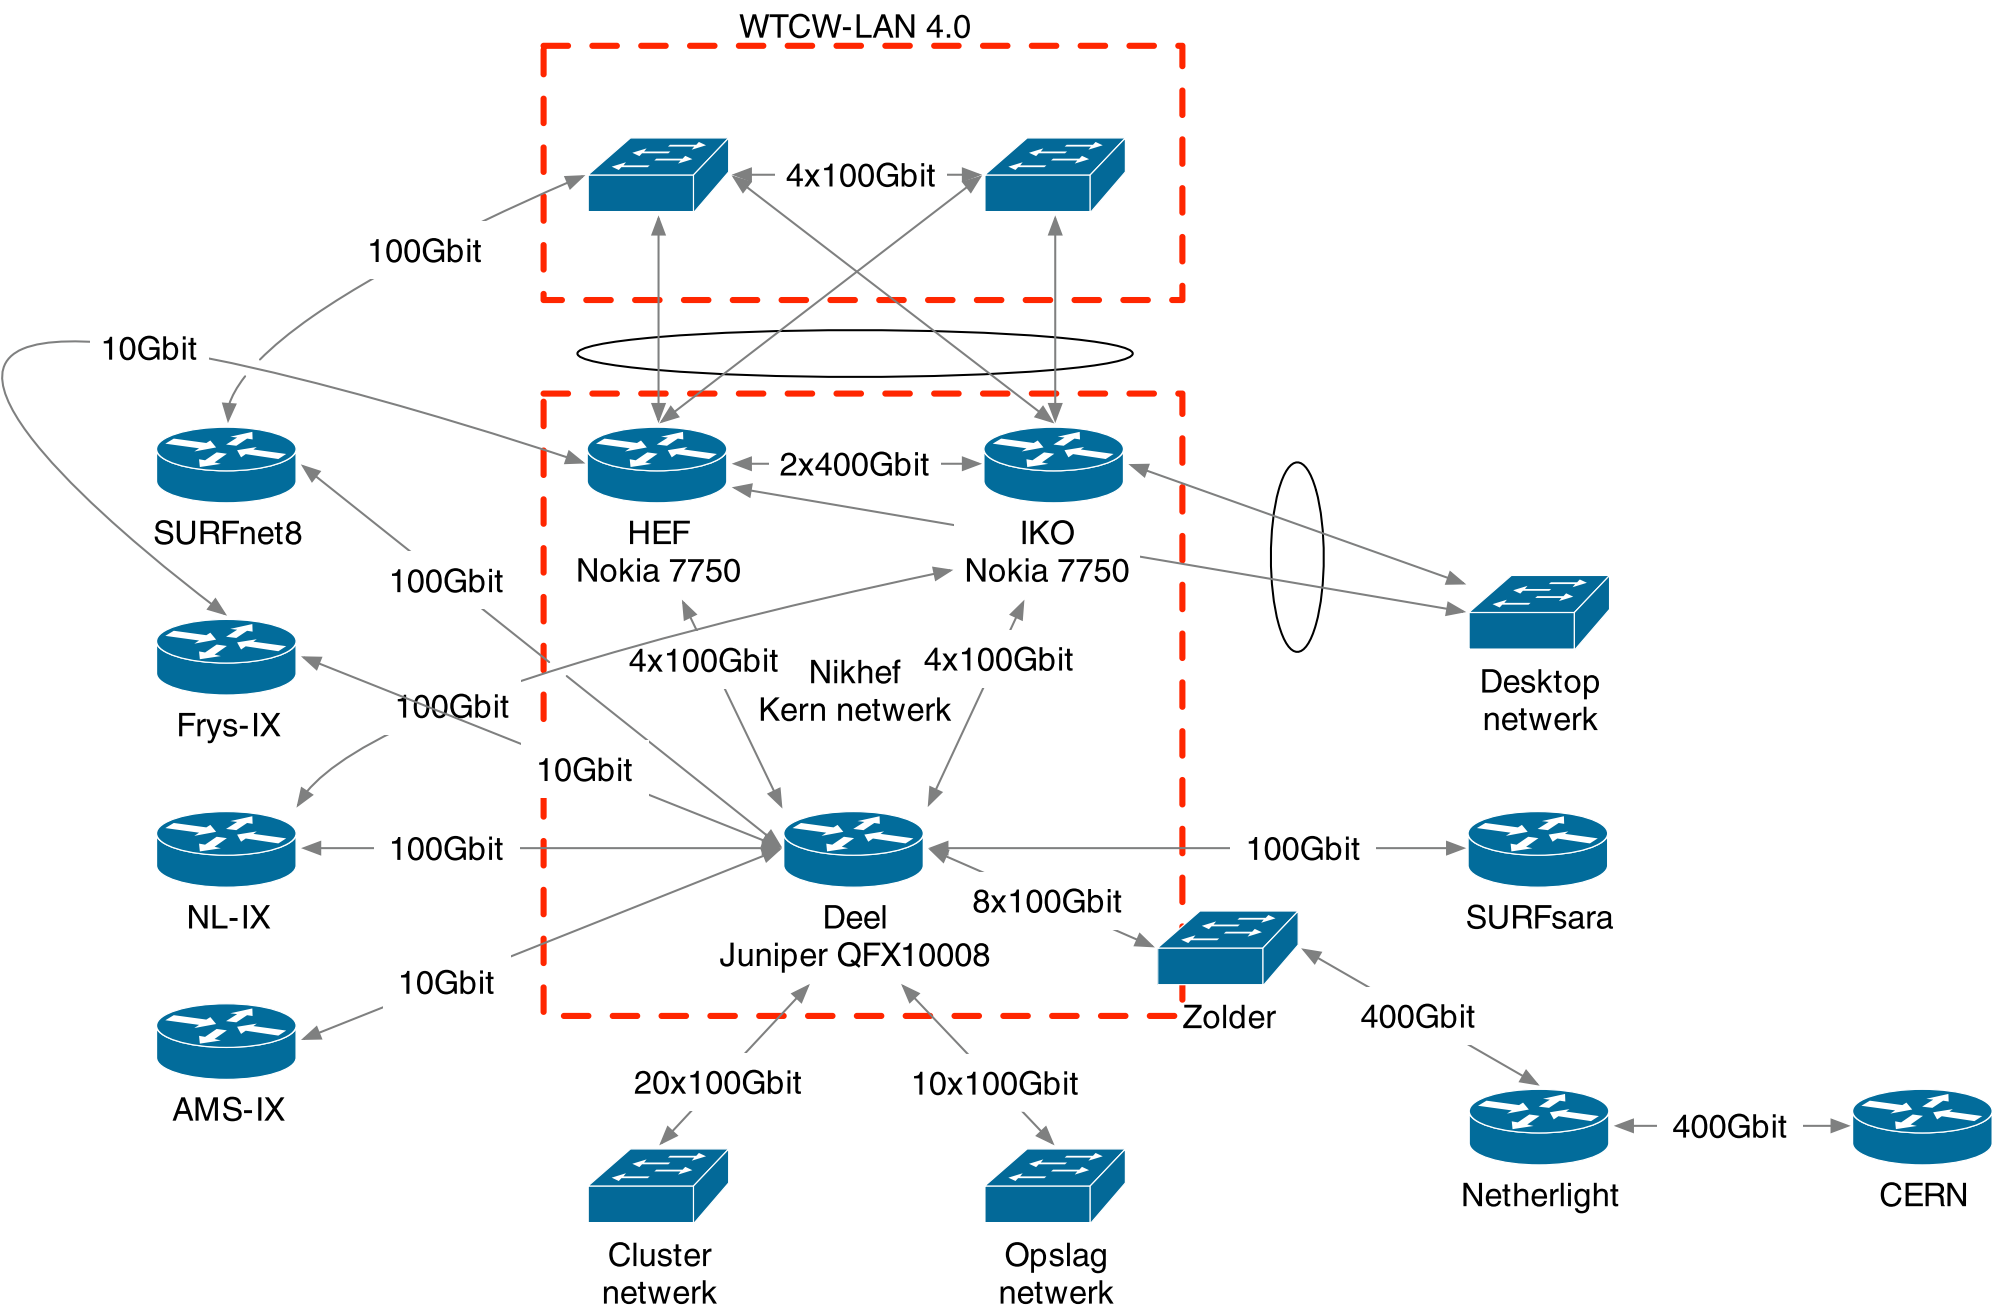
- WEKA distributed file system
- extensively tested
- interesting maybe for large shared scratch space across worker nodes
- No support for IPv6?

# STORAGE III

- SSD backed data transfer node
- System not capable of 400 Gbit/s yet
- reached about 30 GB/s

# NETWORKING

- New core router setup, dual Nokia 7750
- 800 Gbit/s ready
- More in [talk by Bart van der Wal](#)
- Plans for 800 Gbit/s (coherent) link to CERN.
- Generally testing routing/switching platforms





# COMPUTING

- AMD Genoa test (96 cores, 192 threads)  
Seems to deliver same performance per thread as AMD Rome delivers per core.
- AMD Bergamo (128 cores) is coming.

# END OF GENERIC COMPUTING?

- Trade-offs between number of cores and memory bandwidth
- Adding more cores won't increase the total memory bandwidth.
- Genoa-X has a (much) larger L3 cache, which could be interesting for applications that need high throughput with small data.

Conclusion: type of system really begins to depend on the type of application we are aiming for.

# MORE COMPUTING

- AMD GPU MI210 now on par with NVIDIA offerings

Generally: we're having good relationships with hardware vendors and are open for collaboration on testing!

# AND MORE COMPUTING

In the new lab (post-renovation) we're going to dive into water cooling.

- Currently too many competing systems and standards
- Understanding what is out there and what we could buy is key
- Avoiding vendor lock-in
- This would lift the 400kW air-cooling limit

*To be presented at a future HEPiX!*

# RENOVATION

- Oldest building on the campus. Needed a facelift.



# CHALLENGES

- Many challenges:
- This is not your ordinary office building
- Keeping the lights on for computing facilities 24/7
- Building a new extension to our housing facility.

See [talk by Floris](#) on Friday.





An architectural rendering of a modern building interior. The space is characterized by a high ceiling with a grid of white beams and large glass windows. The floor is made of light-colored wood. In the foreground, three silhouettes of people are standing on the wooden floor. In the middle ground, a large area is set up with many white chairs, suggesting a meeting or workshop. In the background, there are more people and a large glass wall that looks out onto a landscape. The overall atmosphere is bright and open.

*Come visit Nikhef in our new building  
some time! Maybe for the HTCondor  
Europe Workshop in 2024.*

# THANKS FOR YOUR ATTENTION

With thanks to my colleagues who provided this material:

Mischa Sallé, Tristan Suerink, David Groep, Erik Kooistra

