Compute roadmap and initiatives

Jorik van Kemenade

31 January, 2025

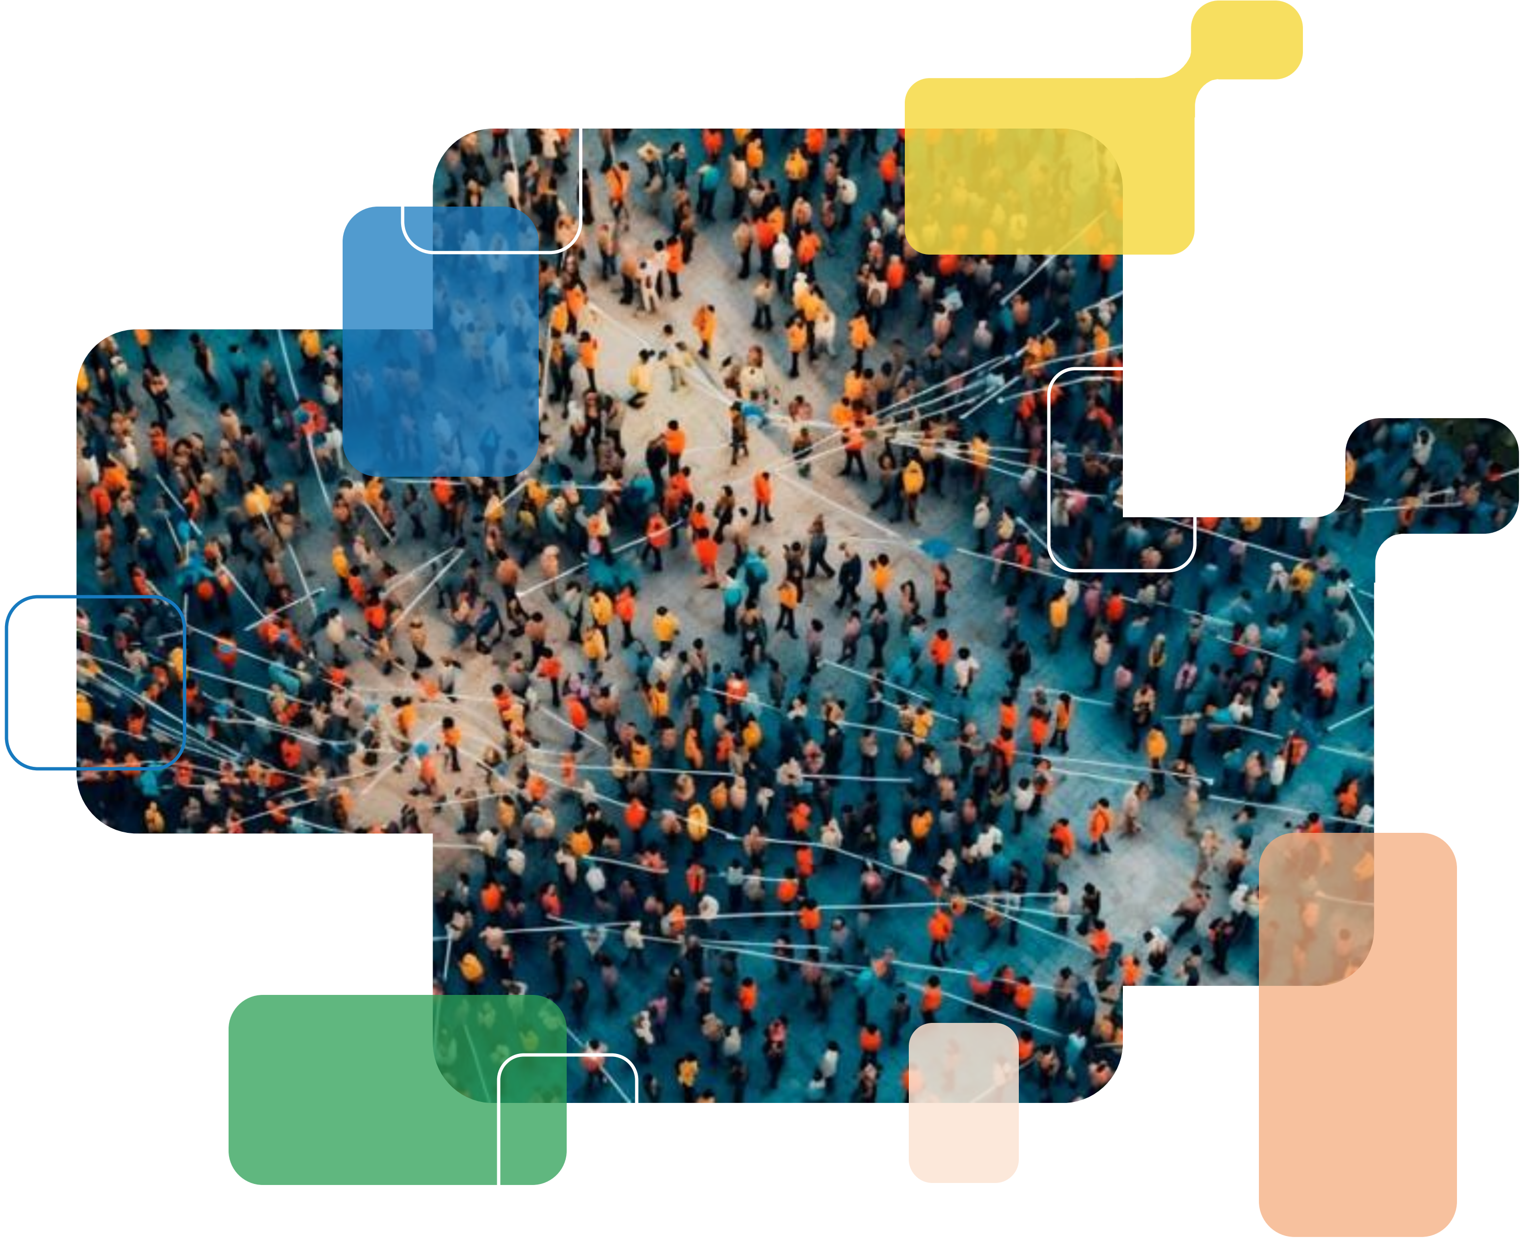

The compute domain roadmap

High Performance Computing

2025 · 2026 · 2027

Ease of use	Openondemand (UI)	SRAM SSO	Common rates & access routes	Expanding AI support and services (e.g. AI Hub)	
Capacity	Tier-2 Connection & Access	Alice Recoque participation	Hybrid Quantum computer	AI Factory	New National Supercomputer
Efficiency	Capacity management with ML	Energy-efficient computing	Integration of HPC / HTC / Cloud infra		
Sensitive data	Developing OSSC(AR) and connecting to SANE				

Data Processing

2025 · 2026 · 2027

Ease of use	Improved interoperability of other services	Common rates & access routes	Community based tooling
Capacity	Preparation for LOFAR 2.0 and SKA	Supporting CERN growth (LH-LHC)	
Efficiency	Batch computing alternatives	Integration of HPC / HTC / Cloud infra	
Sensitive data	Trusted clusters-as-a-service		

Cloud computing

2025 · 2026 · 2027

Ease of use	Improve SRC UI	Common rates & access routes	Starting National EOSC Node
Capacity	Access and interoperability with public clouds	Expand SRC components catalog	Container-based services
Efficiency	Data center redundancy	Integration of HPC / HTC / Cloud infra	
Sensitive data	Developing SANE	Developing data access broker	

SURF

The compute domain roadmap - Capacity Building

High Performance Computing

2025 — 2026 — 2027

Ease of use	Openondemand (UI)	SRAM SSO	Common rates & access routes	Expanding AI support and services (e.g. AI Hub)	
Capacity	Tier-2 Connection & Access	Alice Recoque participation	Hybrid Quantum computer	AI Factory	New National Supercomputer
Efficiency	Capacity management with ML	Energy-efficient computing	Integration of HPC / HTC / Cloud infra		
Sensitive data	Developing OSSC(AR) and connecting to SANE				

Data Processing

2025 — 2026 — 2027

Ease of use	Improved interoperability of other services	Common rates & access routes	Community based tooling
Capacity	Preparation for LOFAR 2.0 and SKA	Supporting CERN growth (LH-LHC)	
Efficiency	Batch computing alternatives	Integration of HPC / HTC / Cloud infra	
Sensitive data	Trusted clusters-as-a-service		

Cloud computing

2025 — 2026 — 2027

Ease of use	Improve SRC UI	Common rates & access routes	Starting National EOSC Node
Capacity	Access and interoperability with public clouds	Expand SRC components catalog	Container-based services
Efficiency	Data center redundancy	Integration of HPC / HTC / Cloud infra	
Sensitive data	Developing SANE	Developing data access broker	

SURF

The compute domain roadmap – HPC and HTC

High Performance Computing

	2025	2026	2027
Ease of use	Openondemand (UI); SRAM SSO	Common rates & access routes	Expanding AI support and services (e.g. AI Hub)
Capacity	Tier-2 Connection & Access; Alice Recoque participation	Hybrid Quantum computer	AI Factory; New National Supercomputer
Efficiency	Capacity management with ML	Energy-efficient computing	Integration of HPC / HTC / Cloud infra
Sensitive data	Developing OSSC(AR) and connecting to SANE		

Data Processing

	2025	2026	2027
Ease of use	Improved interoperability of other services	Common rates & access routes	Community based tooling
Capacity	Preparation for LOFAR 2.0 and SKA	Supporting CERN growth (LH-LHC)	
Efficiency	Batch computing alternatives		Integration of HPC / HTC / Cloud infra
Sensitive data	Trusted clusters-as-a-service		

Cloud computing

	2025	2026	2027
Ease of use	Improve SRC UI	Common rates & access routes	Starting National EOSC Node
Capacity	Access and interoperability with public clouds	Expand SRC components catalog	Container-based services
Efficiency	Data center redundancy		Integration of HPC / HTC / Cloud infra
Sensitive data	Developing SANE; Developing data access broker		

Running HEP Workloads on Dutch HPC (2019)

Functional Test

- LHCB jobs run on Cartesius
- Support effort needed to install and maintain CVMFS
- No outside connectivity, so data had to be staged to shared filesystem
- Can fully load nodes
- Memory usage not an issue with multi-processing

Going forward

- Tested only O(few) jobs running simultaneously
- Further investigation required for running at scale
- Investigate automated staging (input and output) and job management

Running HEP Workloads on Dutch HPC (2019)

Functional Test

- LHCB jobs run on Cartesius
- Support effort needed to install and ma
- No outside connectivity, so data had to
- Can fully load nodes
- Memory usage not an issue with multi-p

Going forward

- Tested only O(few) jobs running simulta
- Further investigation required for runnin
- Investigate automated staging (input and o

Part of two iniatives to test HEP workflows on HPC

- Promissing first results
- Both iniatives got stuck because of resources
- Expected challenges on more data intensive jobs
- Some lessons learned used for Snellius
- Investigating options to test again on Snellius

SURF–ETP

Researchers
Bring domain knowledge and pose challenges and questions to be solved.

Infrastructure
Evolves continuously, adapting to the needs of future research problems.

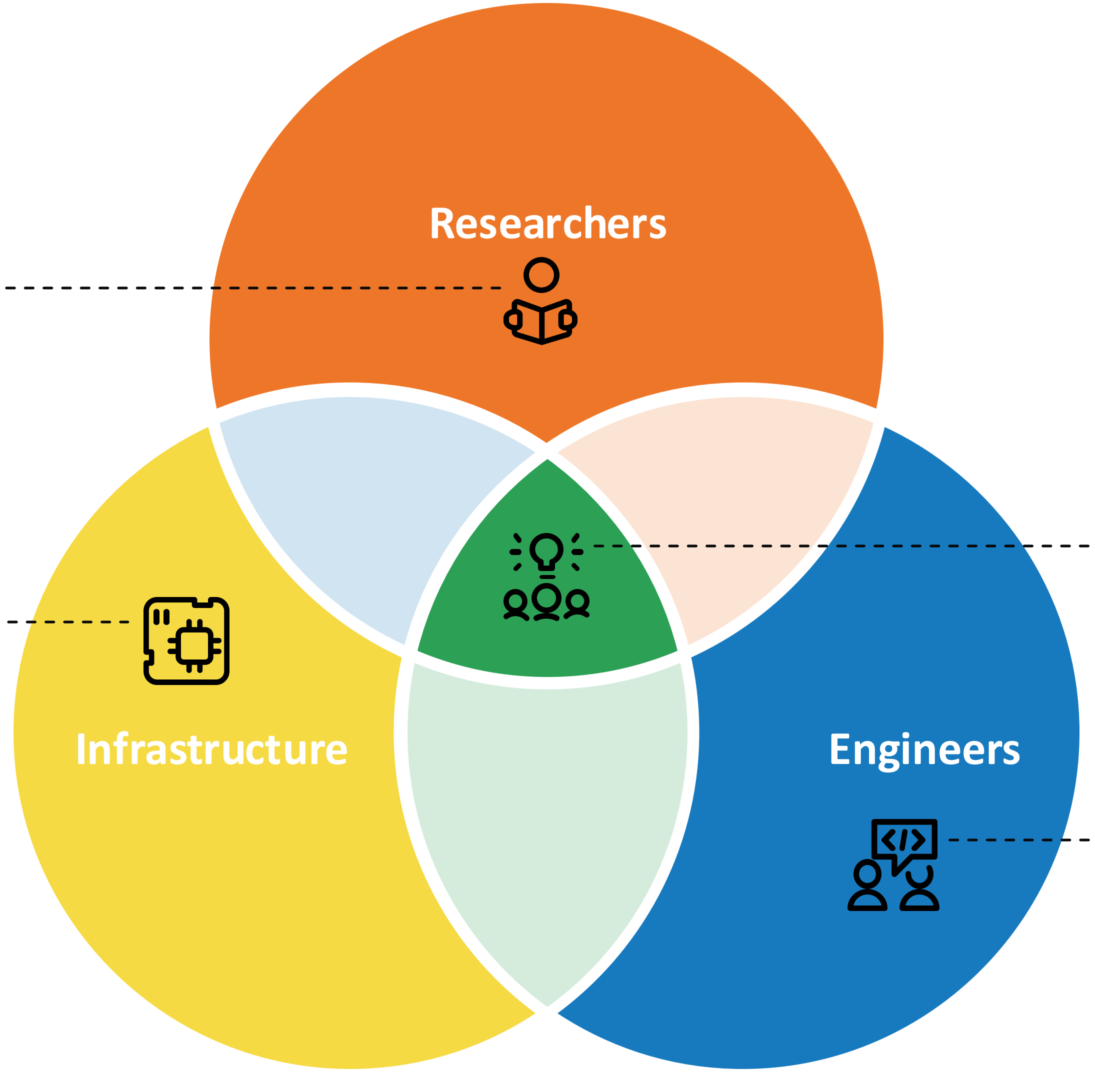

SURF–ETP
Open and collaborative environment to foster the assessment of cutting-edge technologies and methodologies.

Engineers
Leverage their technical skills and motivation to surf the state-of-the-art.

SURF - ETP

AMD Instinct MI210

AMD Instinct MI210 accelerators power enterprise, research, and academic HPC and AI workloads for single-server solutions and more.

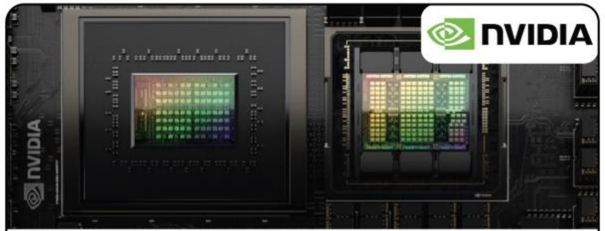

NVIDIA Grace Hopper Superchip

NVIDIA GH200 combines the Grace and Hopper architectures via NVLink-C2C to deliver a CPU+GPU coherent memory model.

Intel GPU Max 1100

Designed to accelerate AI workloads and enable vector and matrix capabilities with Intel Xe Matrix Extensions (XMX).

Xilinx ALVEO U250

The Alveo U250 card uses Xilinx SSI technology to deliver breakthrough FPGA capacity, bandwidth, and power efficiency.

NextSilicon Maverick

Intelligently adapts its architecture to meet algorithmic requirements, all without the need for code modifications.

Asus CRL-G116U-P3DF

AI accelerator card based on Google Coral Edge TPU processor, enabling AI-based real-time decision process at edge.

Xilinx VCK5000 Versal

The VCK5000 Versal is built on the AMD 7nm adaptive SoC architecture and is designed for (AI) Engine development with Vitis end-to-end flow and AI Inference development.

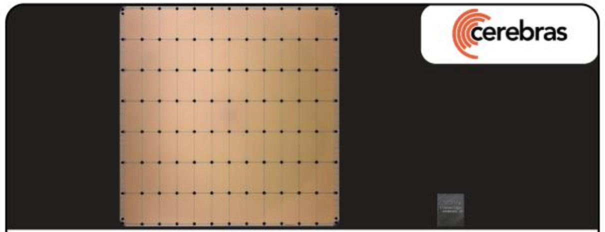

Cerebras WSE-2

850,000 cores, 40 GB of on-chip SRAM, 20 PB/s of memory bandwidth, and 220 Pbps of interconnect bandwidth.

SURF

Data pipelines for data intensive science

SURF participants	Partners
Duncan, Mohsen, Xavier	Cerebras

- Cerebras Systems' CS-2 WSE is a purpose-built system designed to accelerate scientific computing and AI workflows. As such, it delivers the wall-clock compute performance of tens to hundreds of GPUs on specific workloads. To comprehend such unconventional hardware, SURF partnered with Cerebras to conduct a technology assessment.
  - ETP contribution
- Investigated the Cerebras documentation and SDK.
- Executed a set of AI workflows.
- Partnered with a cutting-edge hardware vendor, Cerebras, with the ultimate goal of facilitating access to Cerebras systems for Dutch researchers in the future.

SURF

SURF's CBF and international connectivity

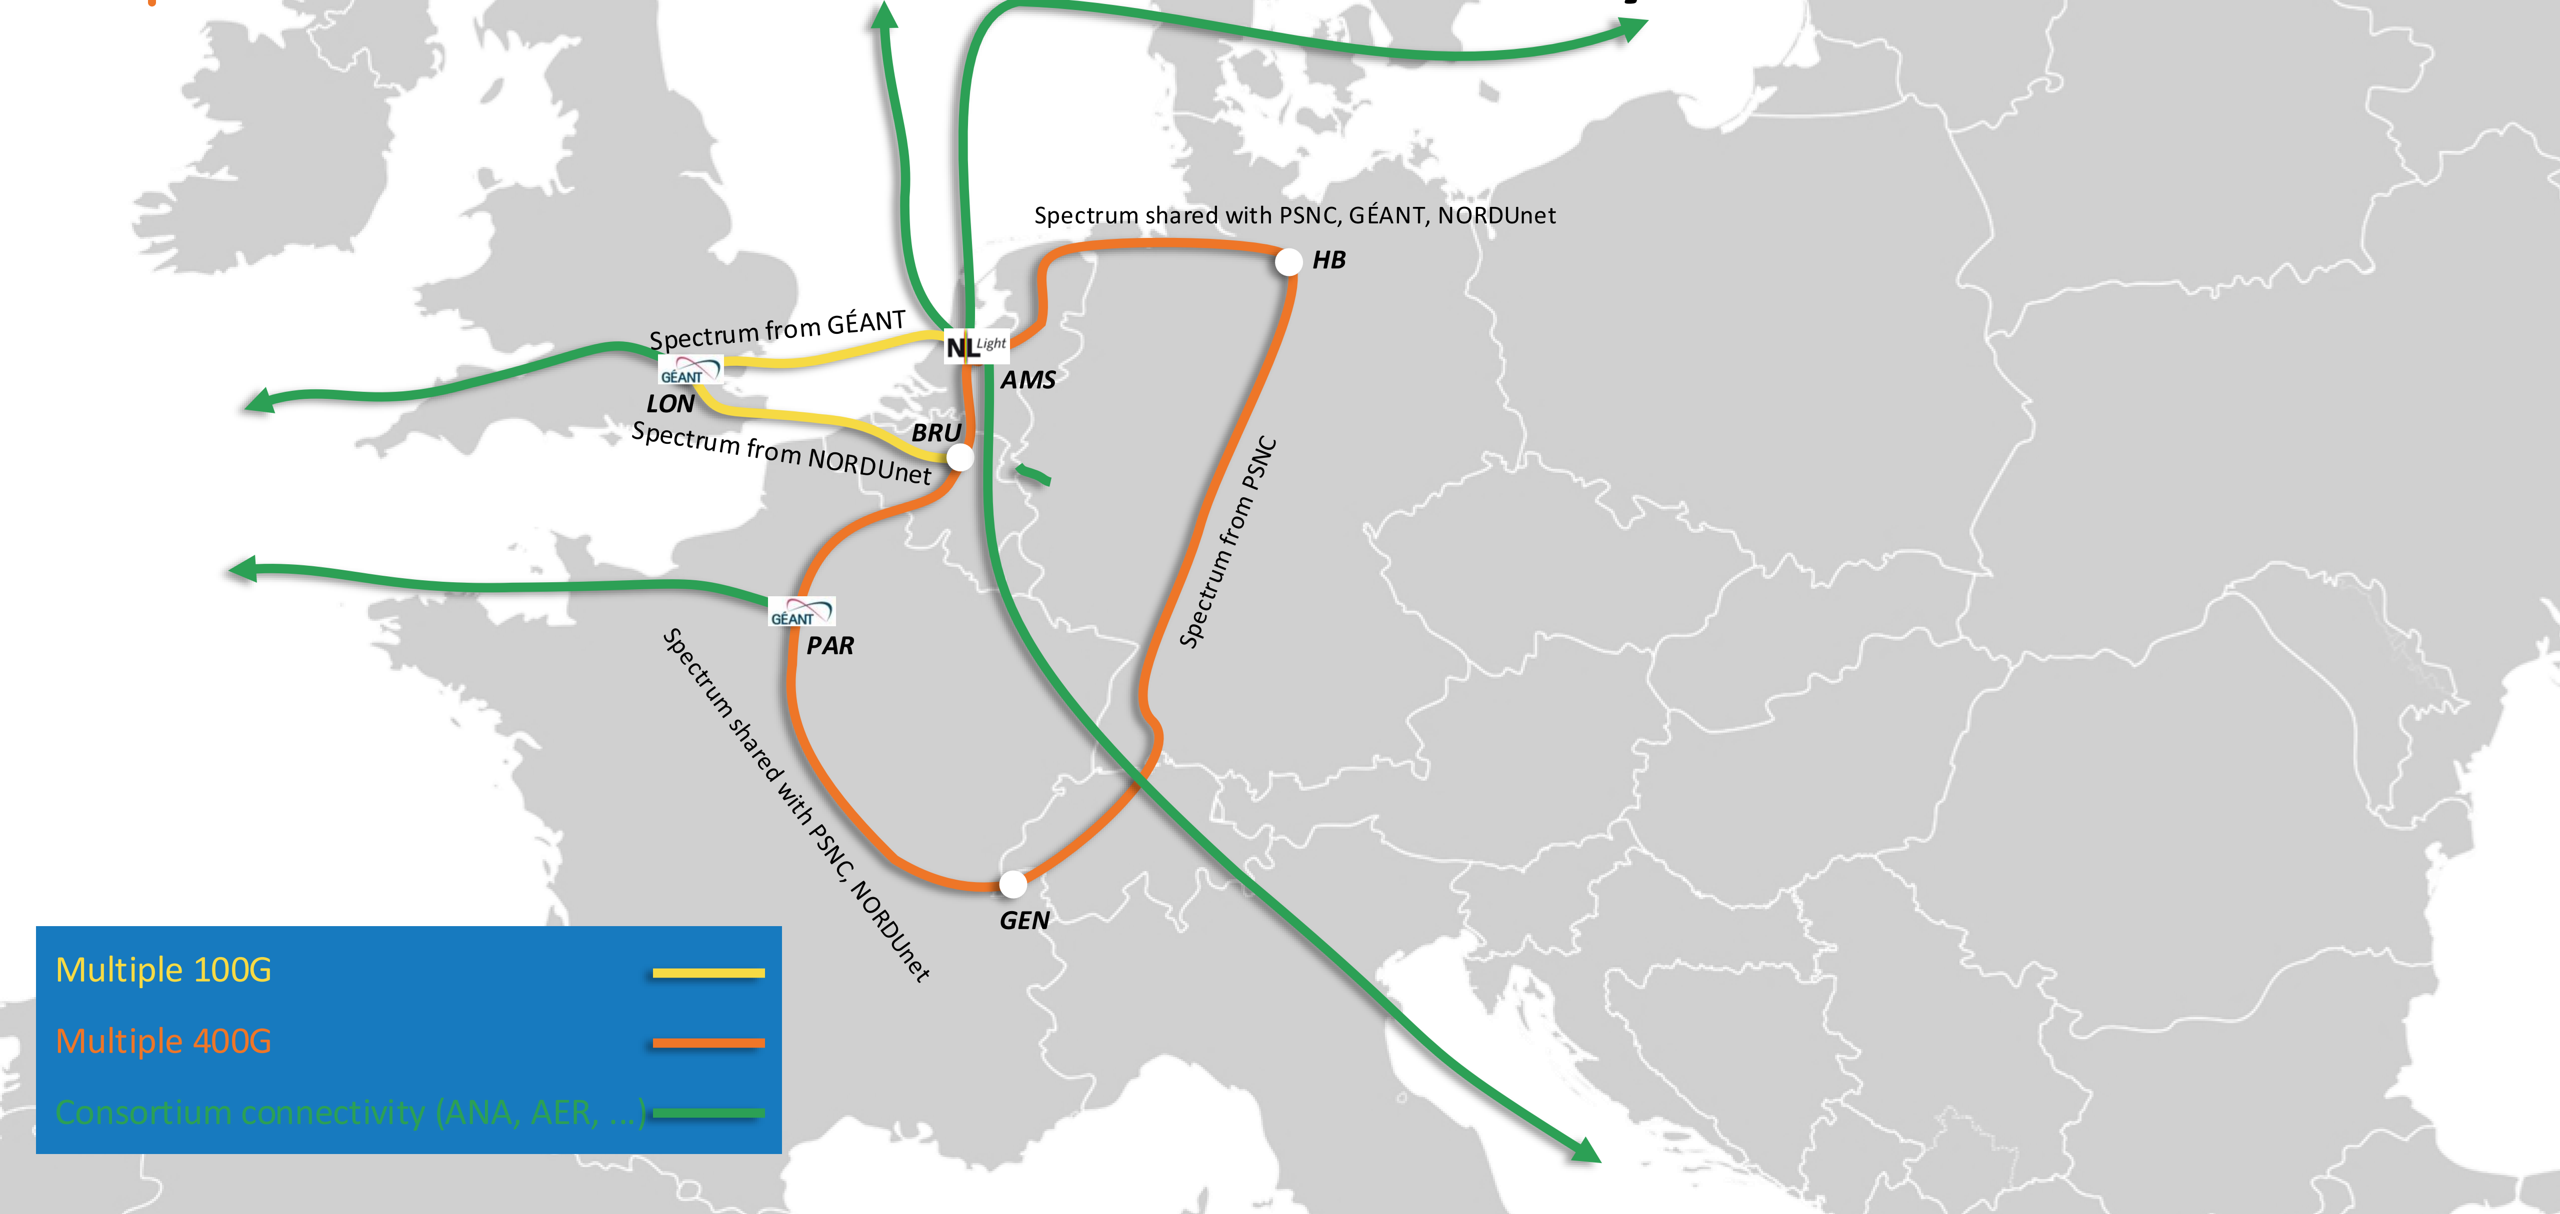

Network innovation - 800G Trial to Geneva

What did we build? What did we test?

- Optical: 800G single channel (150Ghz) on 1648km fiber trajectory with older fiber type (G.655)
- IP: 800GE capable routers on each end (multiple 400G's towards end points)
- Data Storage: End-to-end ATLAS data transfer, storage to storage, between NL-T1 and CERN

Why is this important?

Our goal is to get ready got the high luminosity LHC data explosion and the need for terabit networking.

For this trial we wanted to figure out whether we could get the most out of new transponders on current infrastructure (linesystem and fiber), and also combine this with real data transfers to test scalable storage systems. We want to make sure that we will be prepared when the data increase happens and for this, we need to have the knowledge and the infrastructure in place.

What did we achieve?

- Data processing: writing to disk was impressive – peaks were reached up to 661Gbps, with an average of 540 Gbps. Test was more difficult than when using packet generators, with a lot more parties involved and storage at both ends and their hardware and software, RUCIO, FTS, dCache, et cetera.

- Network: with some optimizations on the line system, tuning, and firmware upgrades from Nokia on the transponders we were able to reach 800G single channel from Amsterdam to Geneva. We were able to transfer up to 800Gbit/s with combination of real traffic and synthetical traffic.

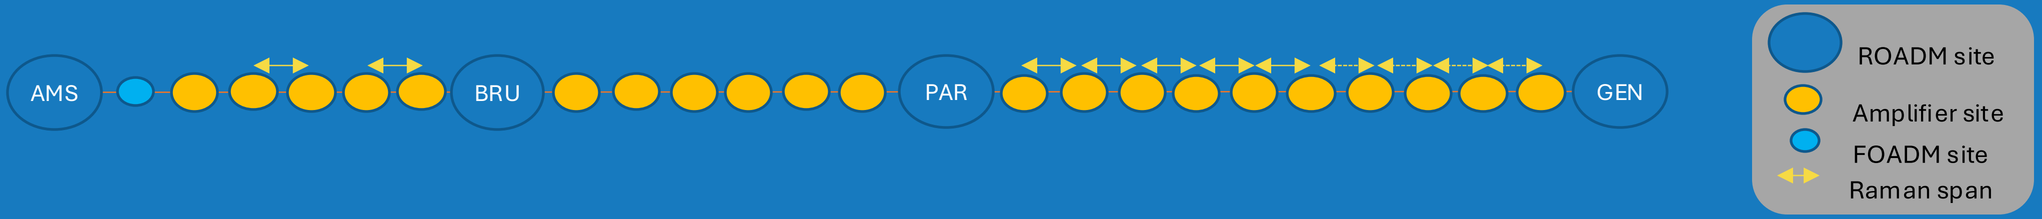

dCache storage throughput during the 800G trial

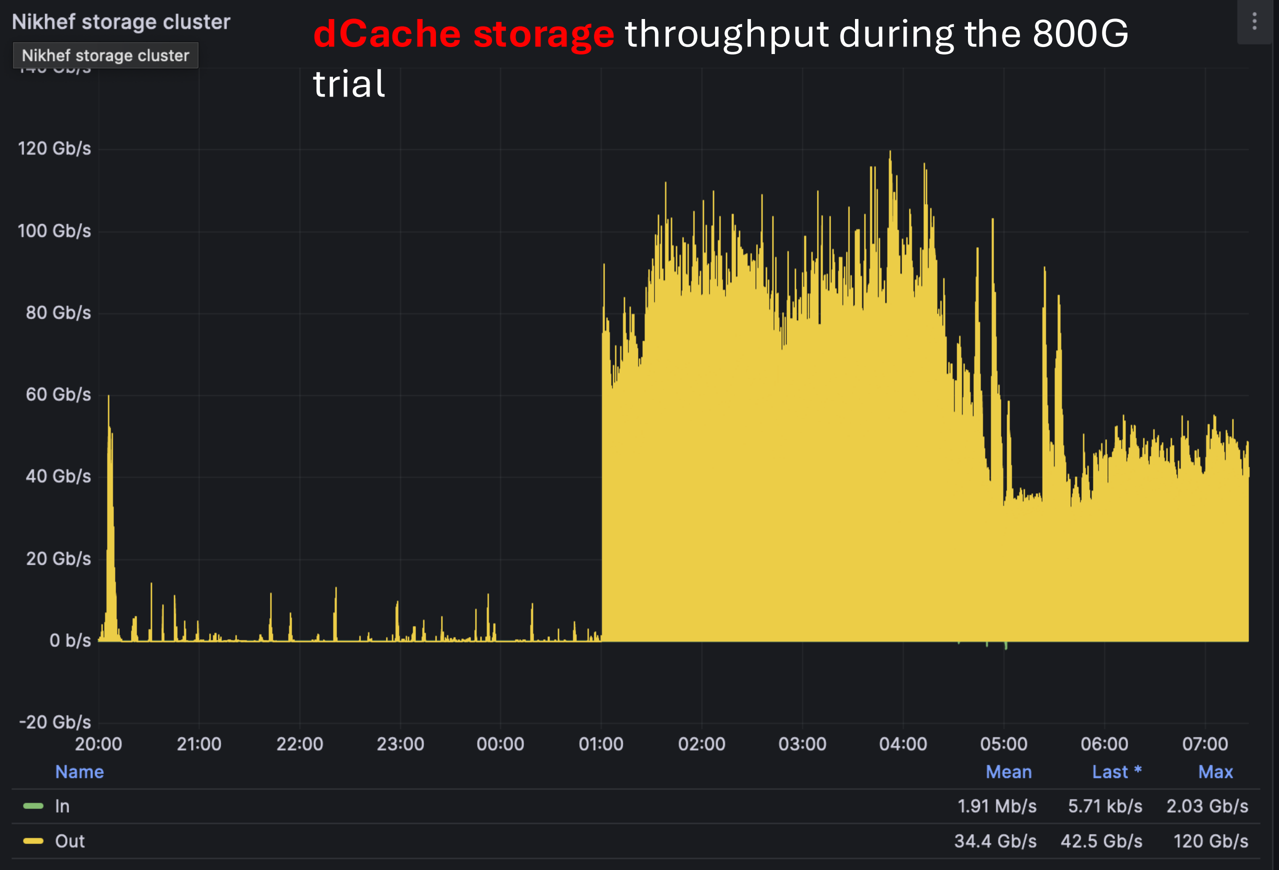

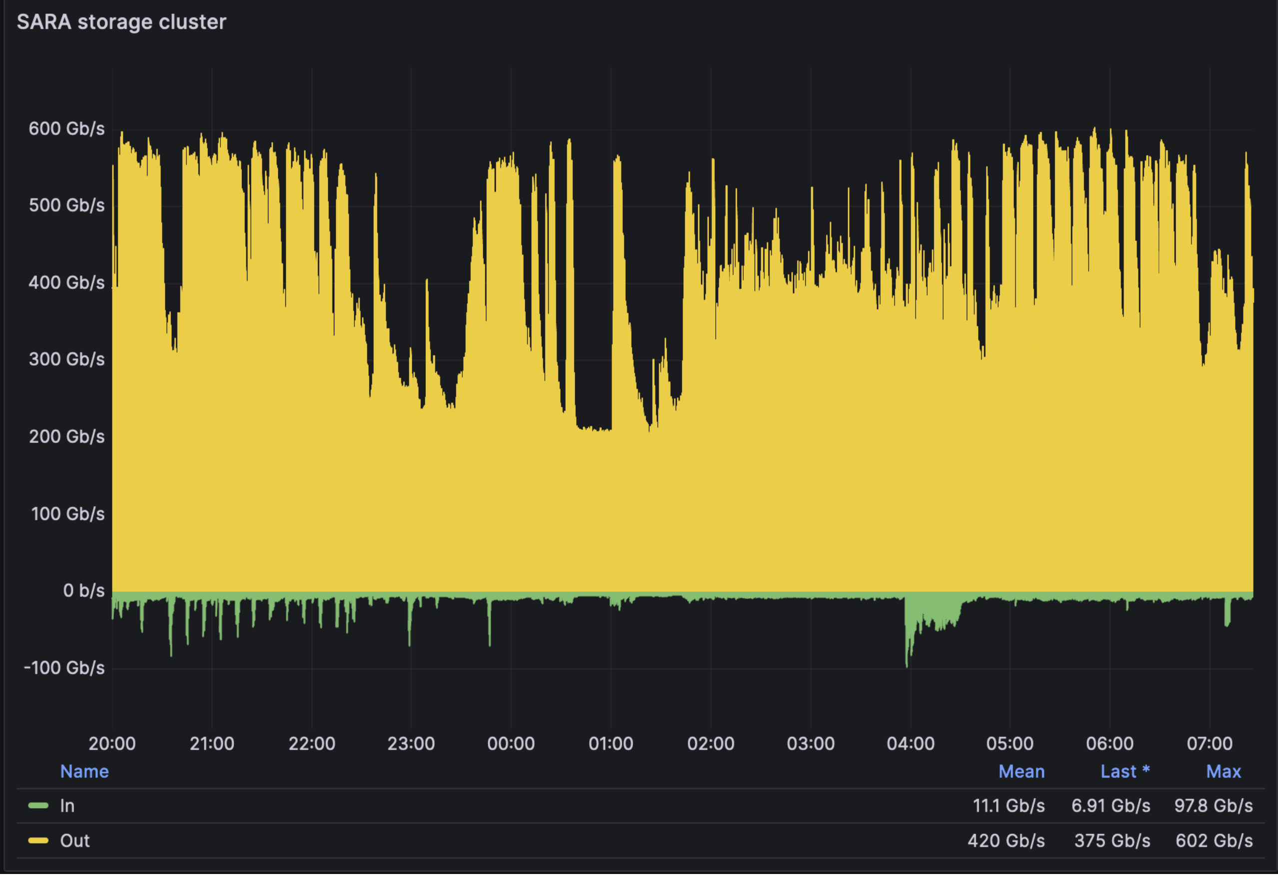

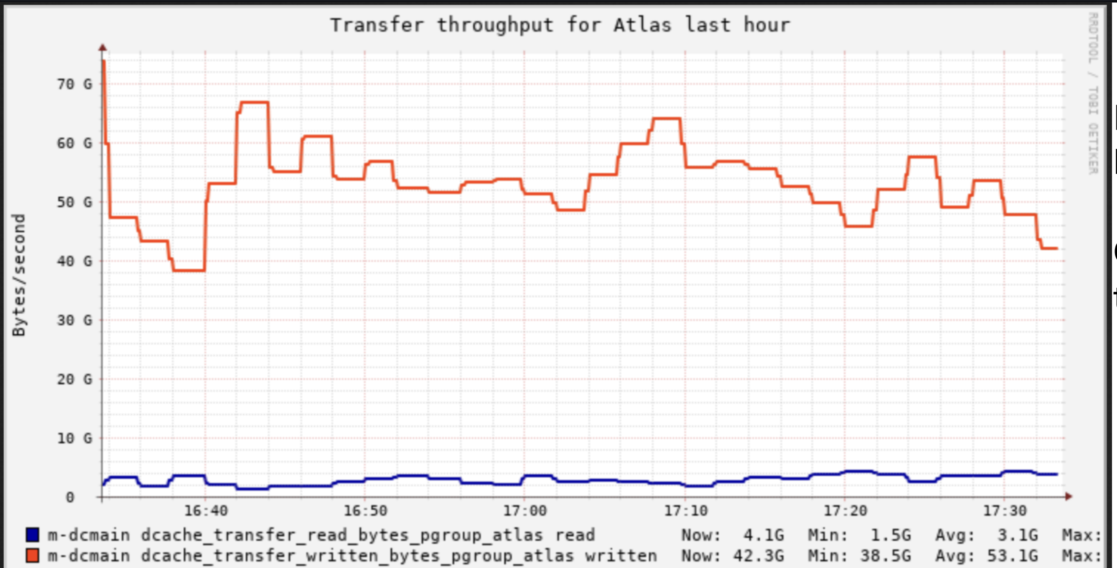

During the trial, om the 800Gbit link, we have reached up to 661 Gbit/s in storage and an hourly average of 530 Gbt.

On dCache: in 2023, 206 PB of data went through dCache, incoming, outgoing, from tape to tape, all of it, 171 PB of which came from CERN.

Network innovation – Kajaani trial

- Create a light path between Amsterdam and Kajaani
  - Distance of 2400 km+
  - With SURF, NORDUNET, SUNET and FUNET
- Test bandwith with different vendors
  - 400 Gbps+, Ribbon, Nokia
- Test use-cases between CERN – Snellus – LUMI
- **Initial trial period Feb – May 2025**

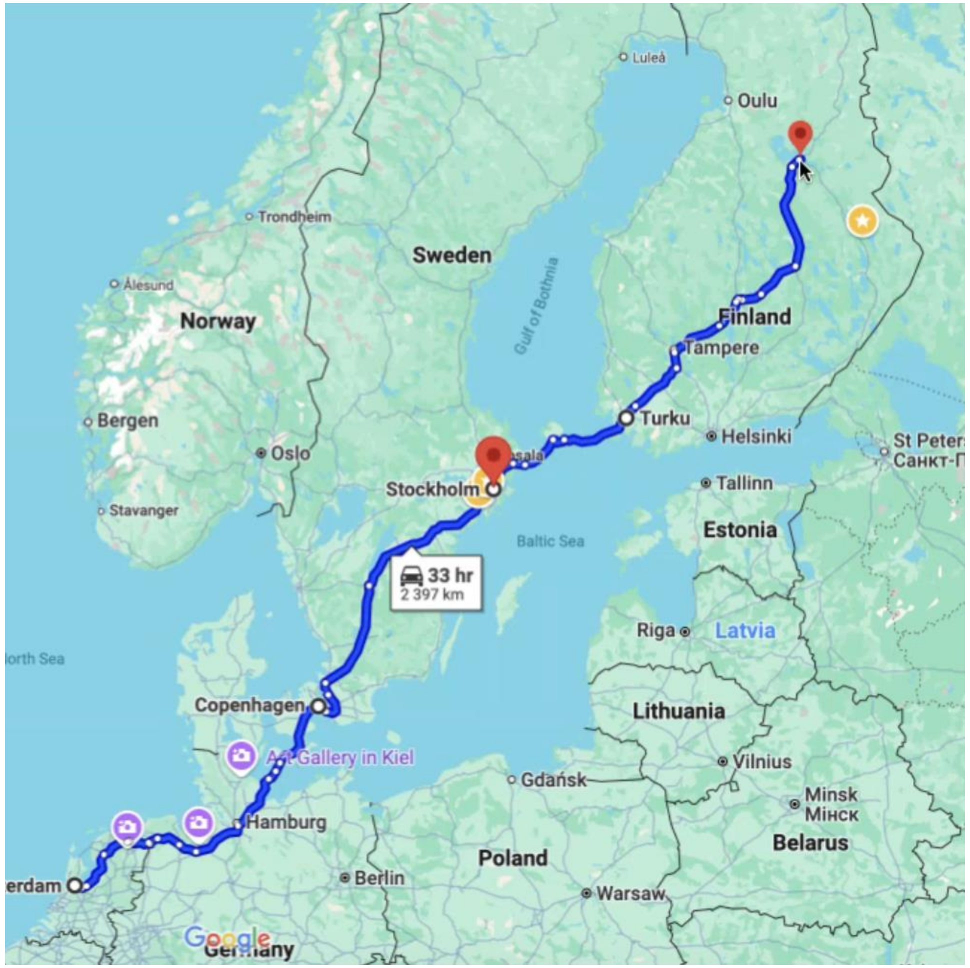

SURF

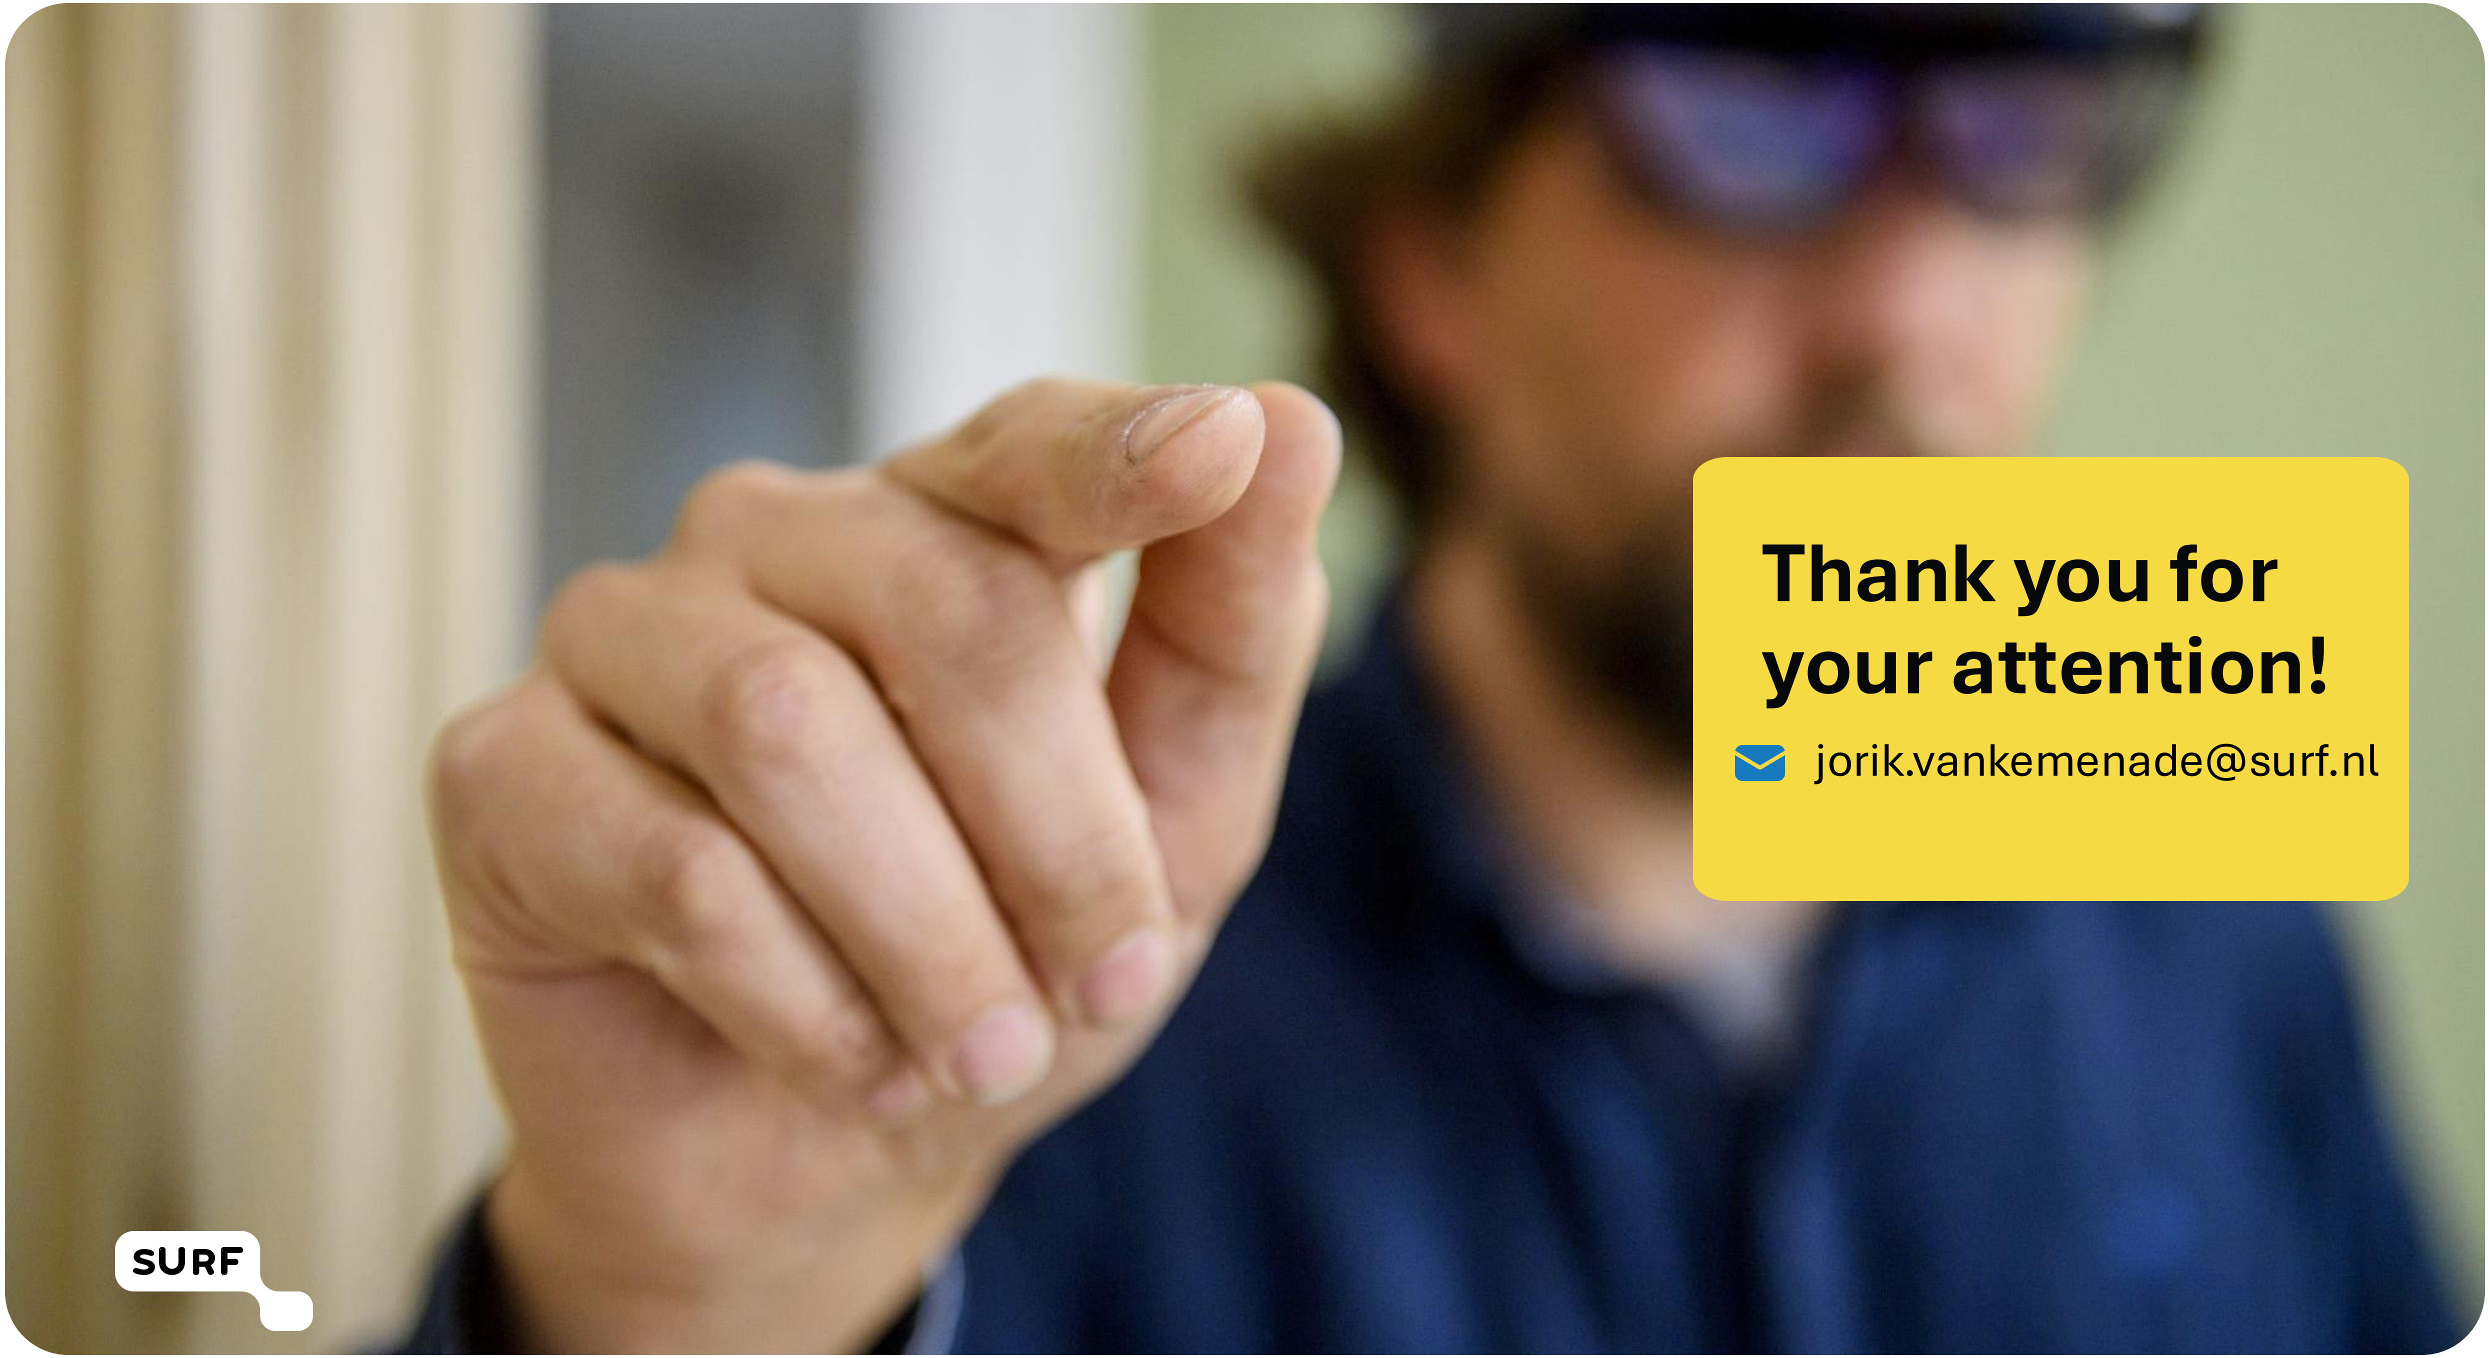

Thank you for your attention!

jorik.vankemenade@surf.nl

SURF