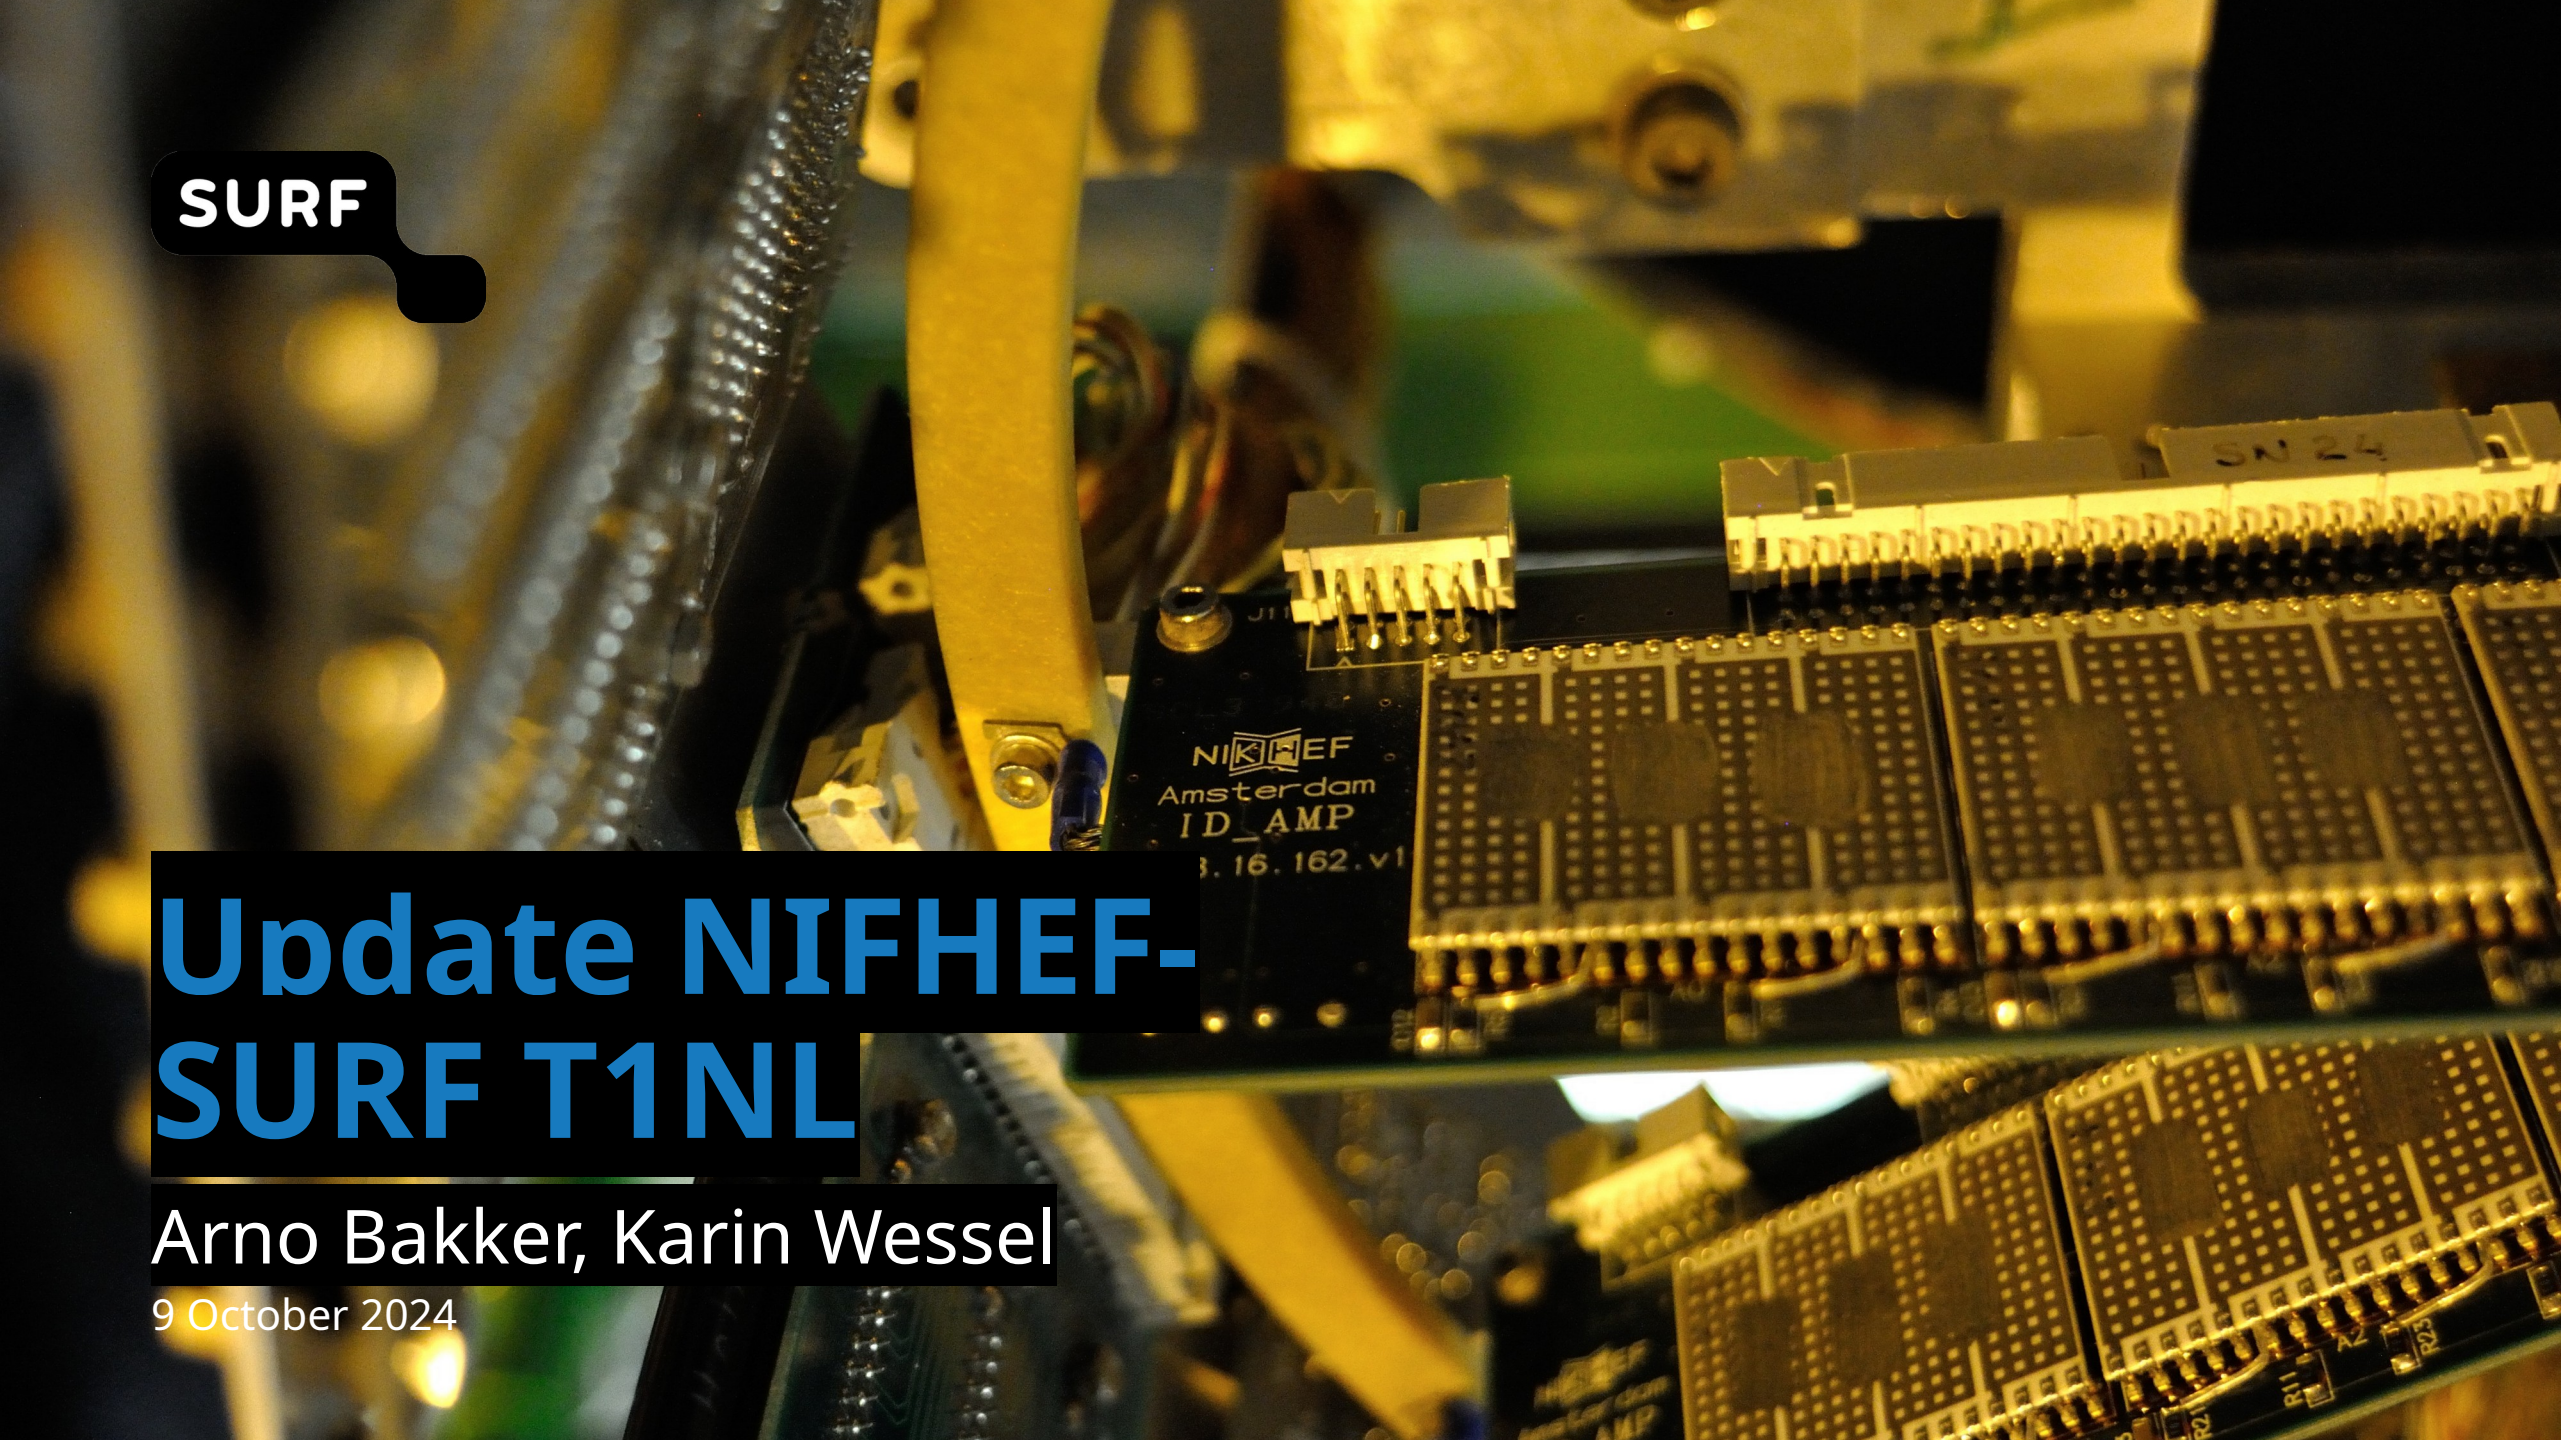

Update NIFHEF-SURF T1NL

Arno Bakker, Karin Wessel

9 October 2024

Who Am I?

Arno Bakker

Email: arno.bakker@surf.nl

Mobile: +31 6 18 24 23 97

New SURF Technical Product Manager

- NetherLight
- International Connectivity

Previously Lecturer in Security and Network Engineering at Univers of Amsterdam

- IETF RFC7574 ;o)

Who Am I?

Arno Bakker

Email: arno.bakker@surf.nl

Mobile: +31 6 18 24 23 97

New SURF Technical Product Manager

- NetherLight
- International Connectivity

Previously Lecturer in Security and Network Engineering at University of Amsterdam

- IETF RFC7574 ;o)

Overview

- About SURF = Dutch NREN
- NetherLight = Global Exchange Point in Amsterdam
- 800G Trials Amsterdam-Geneva
- Changes in Cross-Border Fibers
- Time-Frequency Transfer (TFT) in The Netherlands
- PerfSONAR Measurement Platform @ SURF

About SURF, the Dutch NREN

WHO: SURF is a member owned association focused on aligning its extensive tech services with the needs of its members – Research & Higher Education Institutes.

WHAT: Networking, Data Management, Processing & Storage, HPC, Security, Trust & Identity & Access Management

HOW: Collaboration, Innovation, and Sharing knowledge with our members & partners

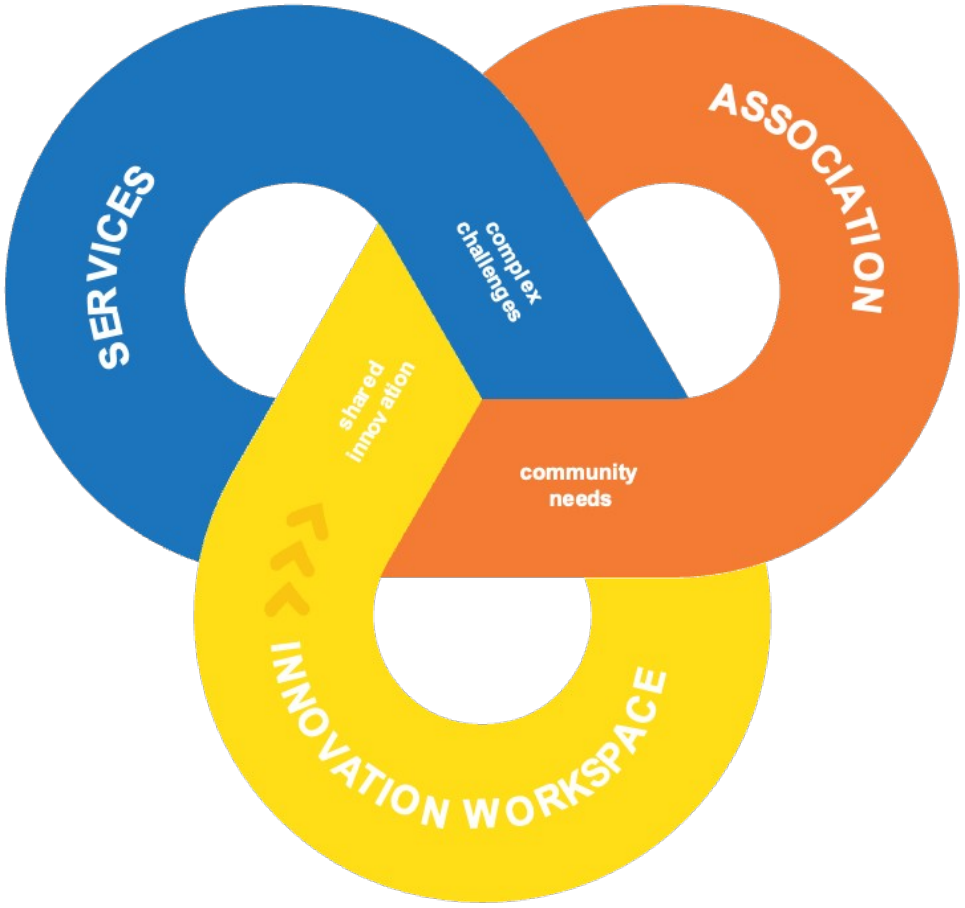

What is NetherLight

- **Kev GLOBAL EXCHANGE POINT** focused on connecting educational and research networks internationally
- Facilitates **high-speed connectivity for research purposes** up to 400 Gbit/s per interface
- Offers **interconnections** via Ethernet VPN (multipoint and point-to-point)
- **Access for our members to services** from connected parties, such as NREN services worldwide, but also services from various cloud providers
- **Current direct international hops**: Aachen, Brussels, Cairo, Geneva, Hamburg, London, Montreal, New Delhi, New York, Paris, Seeb, Seoul, Taipei City, Thuwal, Tokyo

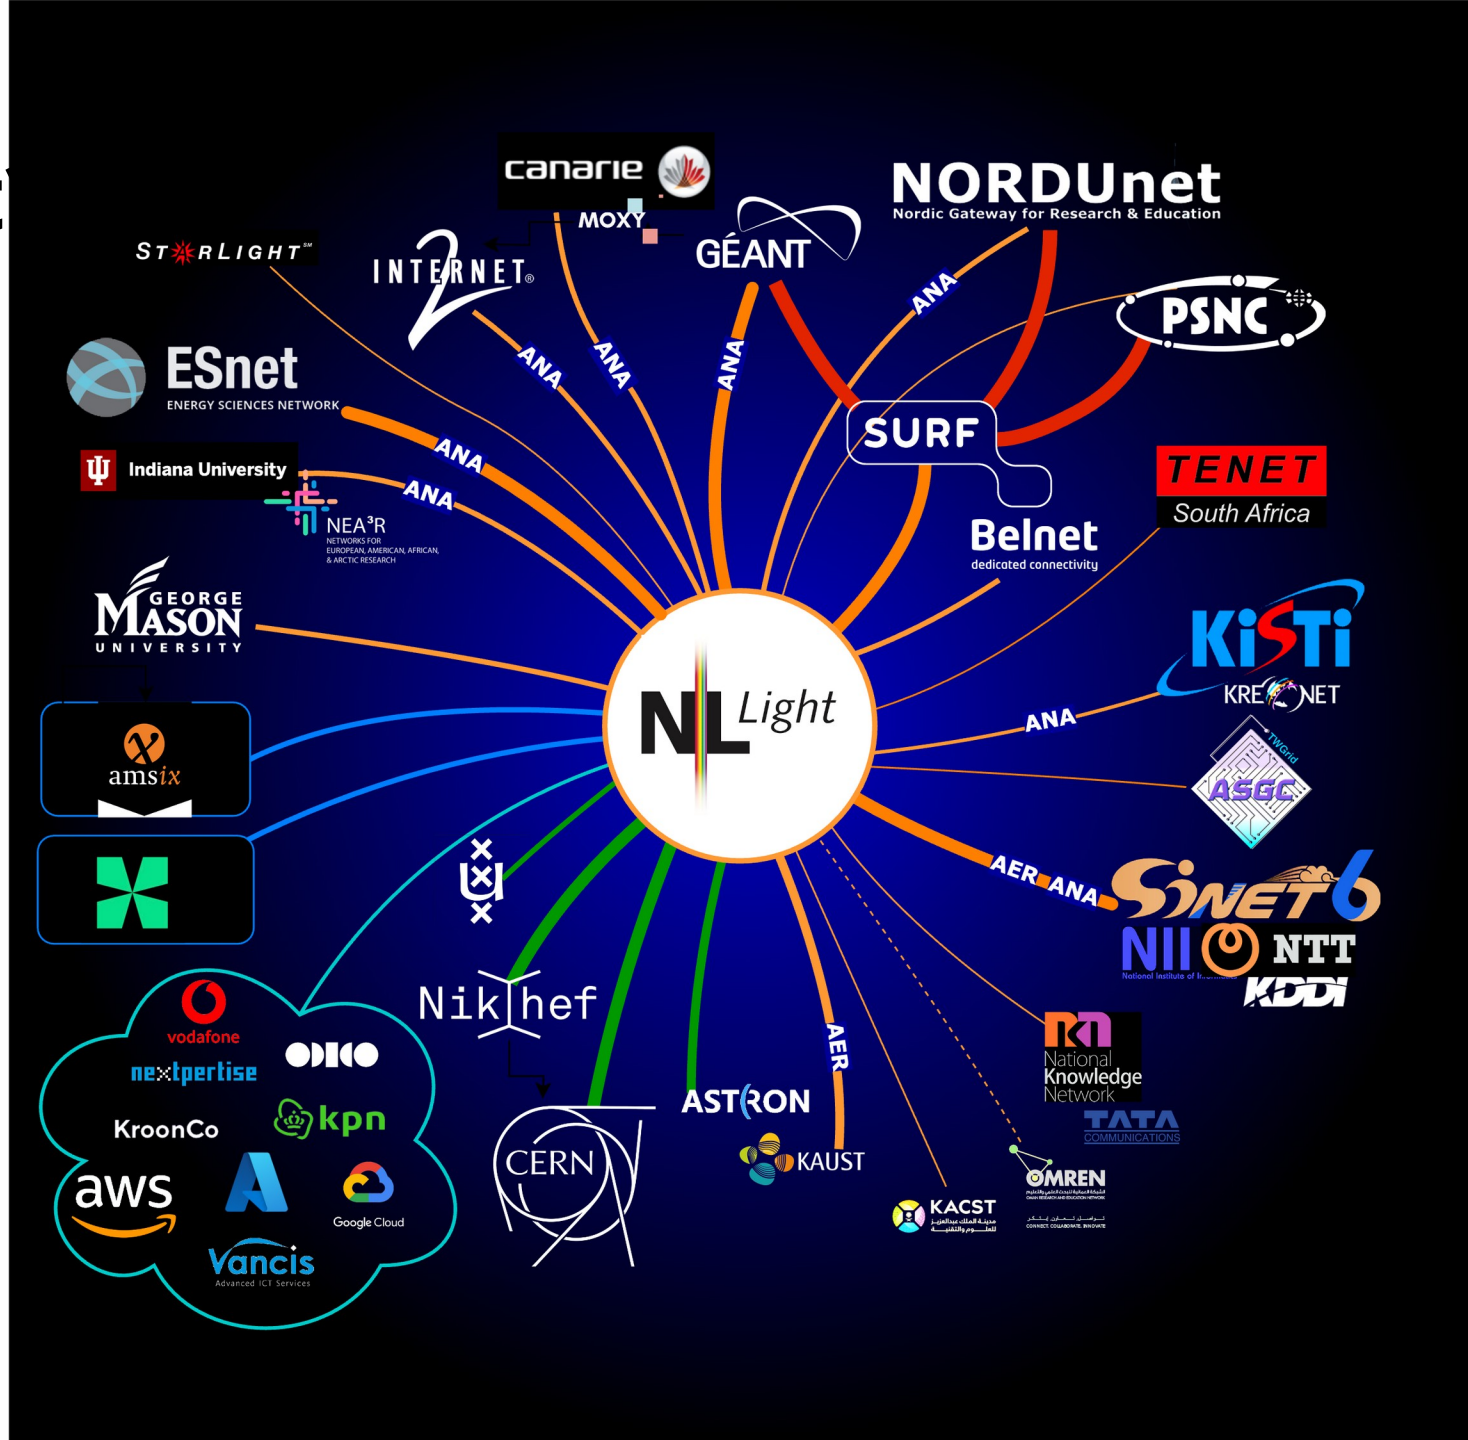

NetherLight

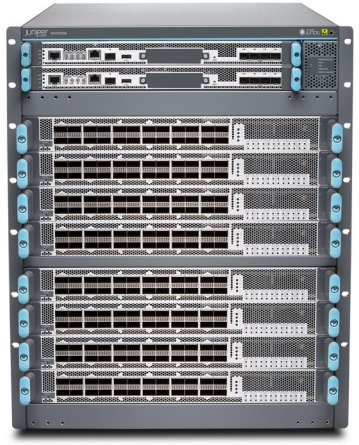

- **Two locations in Amsterdam:**
- **ASD001B** (Digital Realty AMS9, Science Park)
  - **Juniper MX10008**
    - Up to 9.6T connectivity available (100G and 400G)
    - Cf. AMS-IX recent peak at 13.2T
  - **Juniper MX480**
    - For < 100G connectors (1GE and 10GE)
    - Uplinks with n*100G to the MX10008
- **ASD002A** (Equinix AM7)
  - **Juniper MX480**
    - Mainly for cloud connectors that would like to provide redundancy
    - Facilitates up to 100GE connections (1G, 10G, 100G)

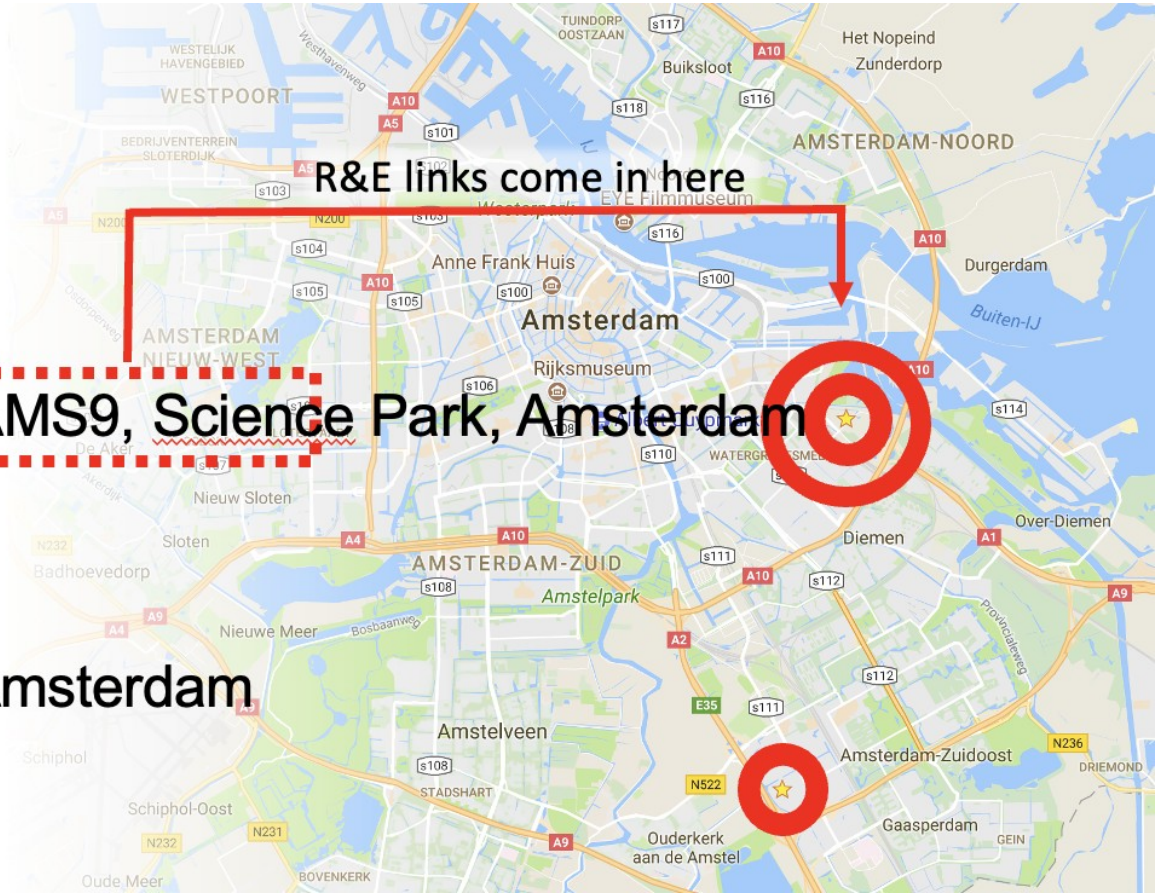

Cross Border Fiber system

Cross Border Fiber system (CBF)

- Production traffic from/to CERN
- Well-connected to the four Open Exchange Points in Europe
    - •NetherLight @ Amsterdam
    - •GÉANT Open @ Londen
    - •GÉANT Open @ Parijs
    - •CERNLight @ Genève
- Sharing capacity on the CBF with other partners that connect on the Open Exchange Points
- Perform cutting-edge trials (800G, 1+ Tbps)

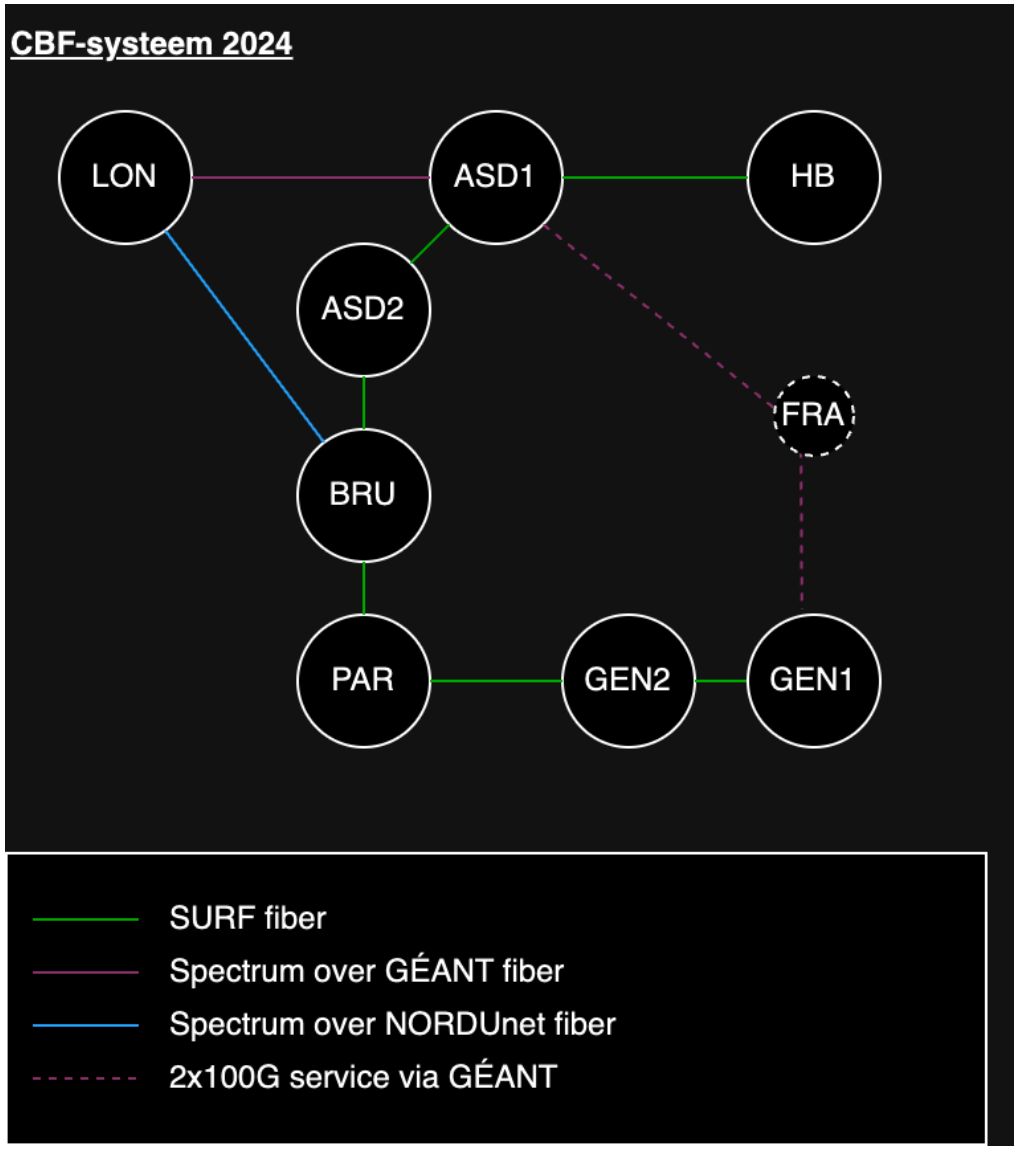

800G Trials

- Nokia (presented in Catania)
- Ribbon

800 G Trials Nokia

https://www.bbc.com/news/articles/clkkllxm4jzo

URF

800G Trials Nokia

- Challenging 800G over 1648 km terrestrial fiber
- Slides borrowed from:
- **Joachim Opdenakker & Edwin Verheul**
  Network Engineers @SURF

Why 800Gbps testing?

LHC

World’s largest and highest-energy particle collider. Located across the border of France and Switzerland.

Experiments

Alice, LHCb, ATLAS and CMS

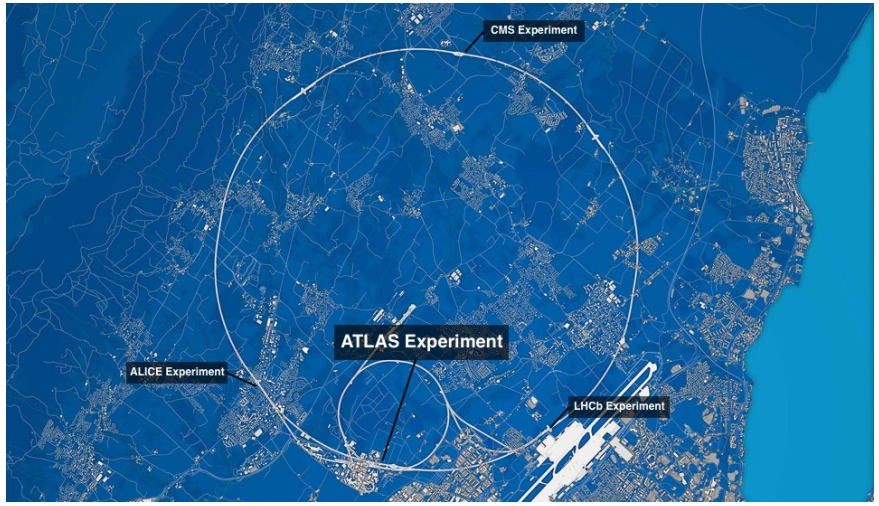

Experiments

ATLAS

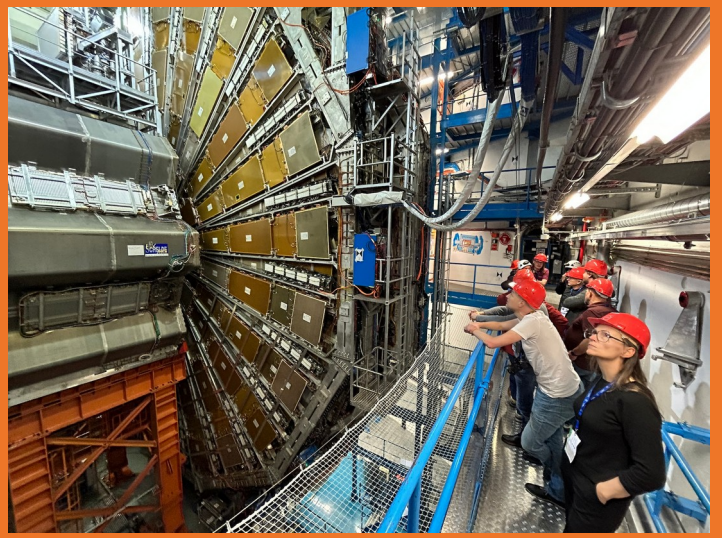

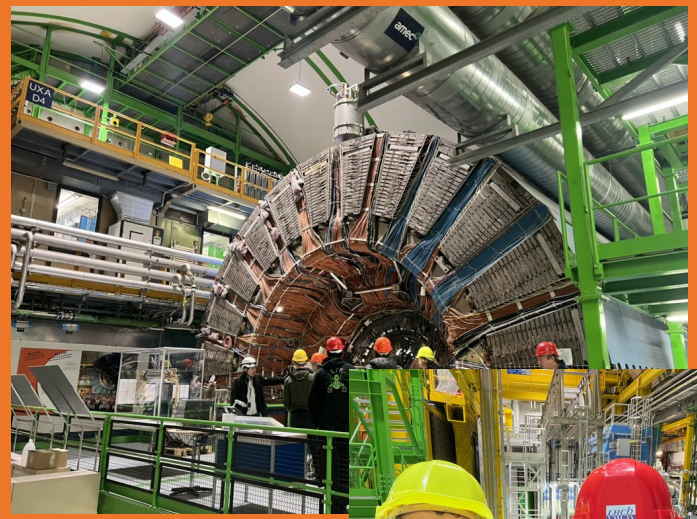

LHCb

Alice

ATLAS detector

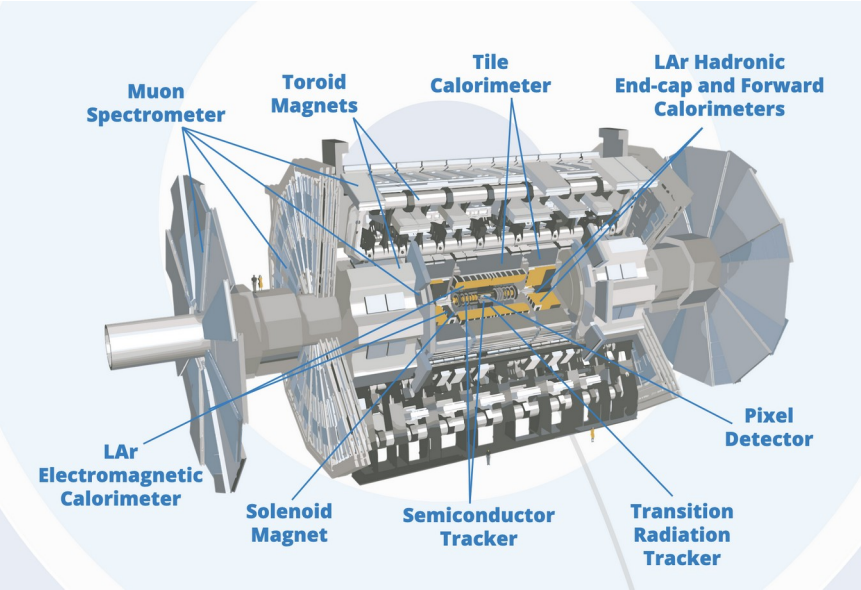

Event rates:

- At a beam luminosity of 10 34 cm-2 s-1, there will be about 20 collisions per buncl crossing.
- 40 million bunch crossings per second.
- Yields about 1 billion collisions per second.
- Level 1 trigger filters that down to about 75 000 events per second.
- Level 2 trigger reduces it to about 2 000 events per second.
- The Event Filter then selects for permanent storage abo second.

Trigger and Data Acquisition (TDAQ)

TDAQ has a 3 level Trigger system (reduction in three steps).
Total event reduction factor by the trigger system: 200 000.

- 1st level trigger: Hardware, level 1 is done using special-purpose processors.[1]
- 2nd level trigger: Software, large computing farms with ~ 500 dual pc processors.
- 3rd level trigger: Software, large computing farms with ~ 1700 dual pc processors.

The rates and reduction factors at 14 TeV are summarized as:

	Incoming event rate per second	Outgoing event rate per second	Reduction factor
Level 1	40 000 000	100 000	400
Level 2	100 000	3 000	30
Level 3	3 000	200	15

TDAQ records 320 Mbytes per second, which would fill more than 27 CDs per minute.

From the ATLAS fact sheet (pre HL-L

HL-LHC data to NL-T1

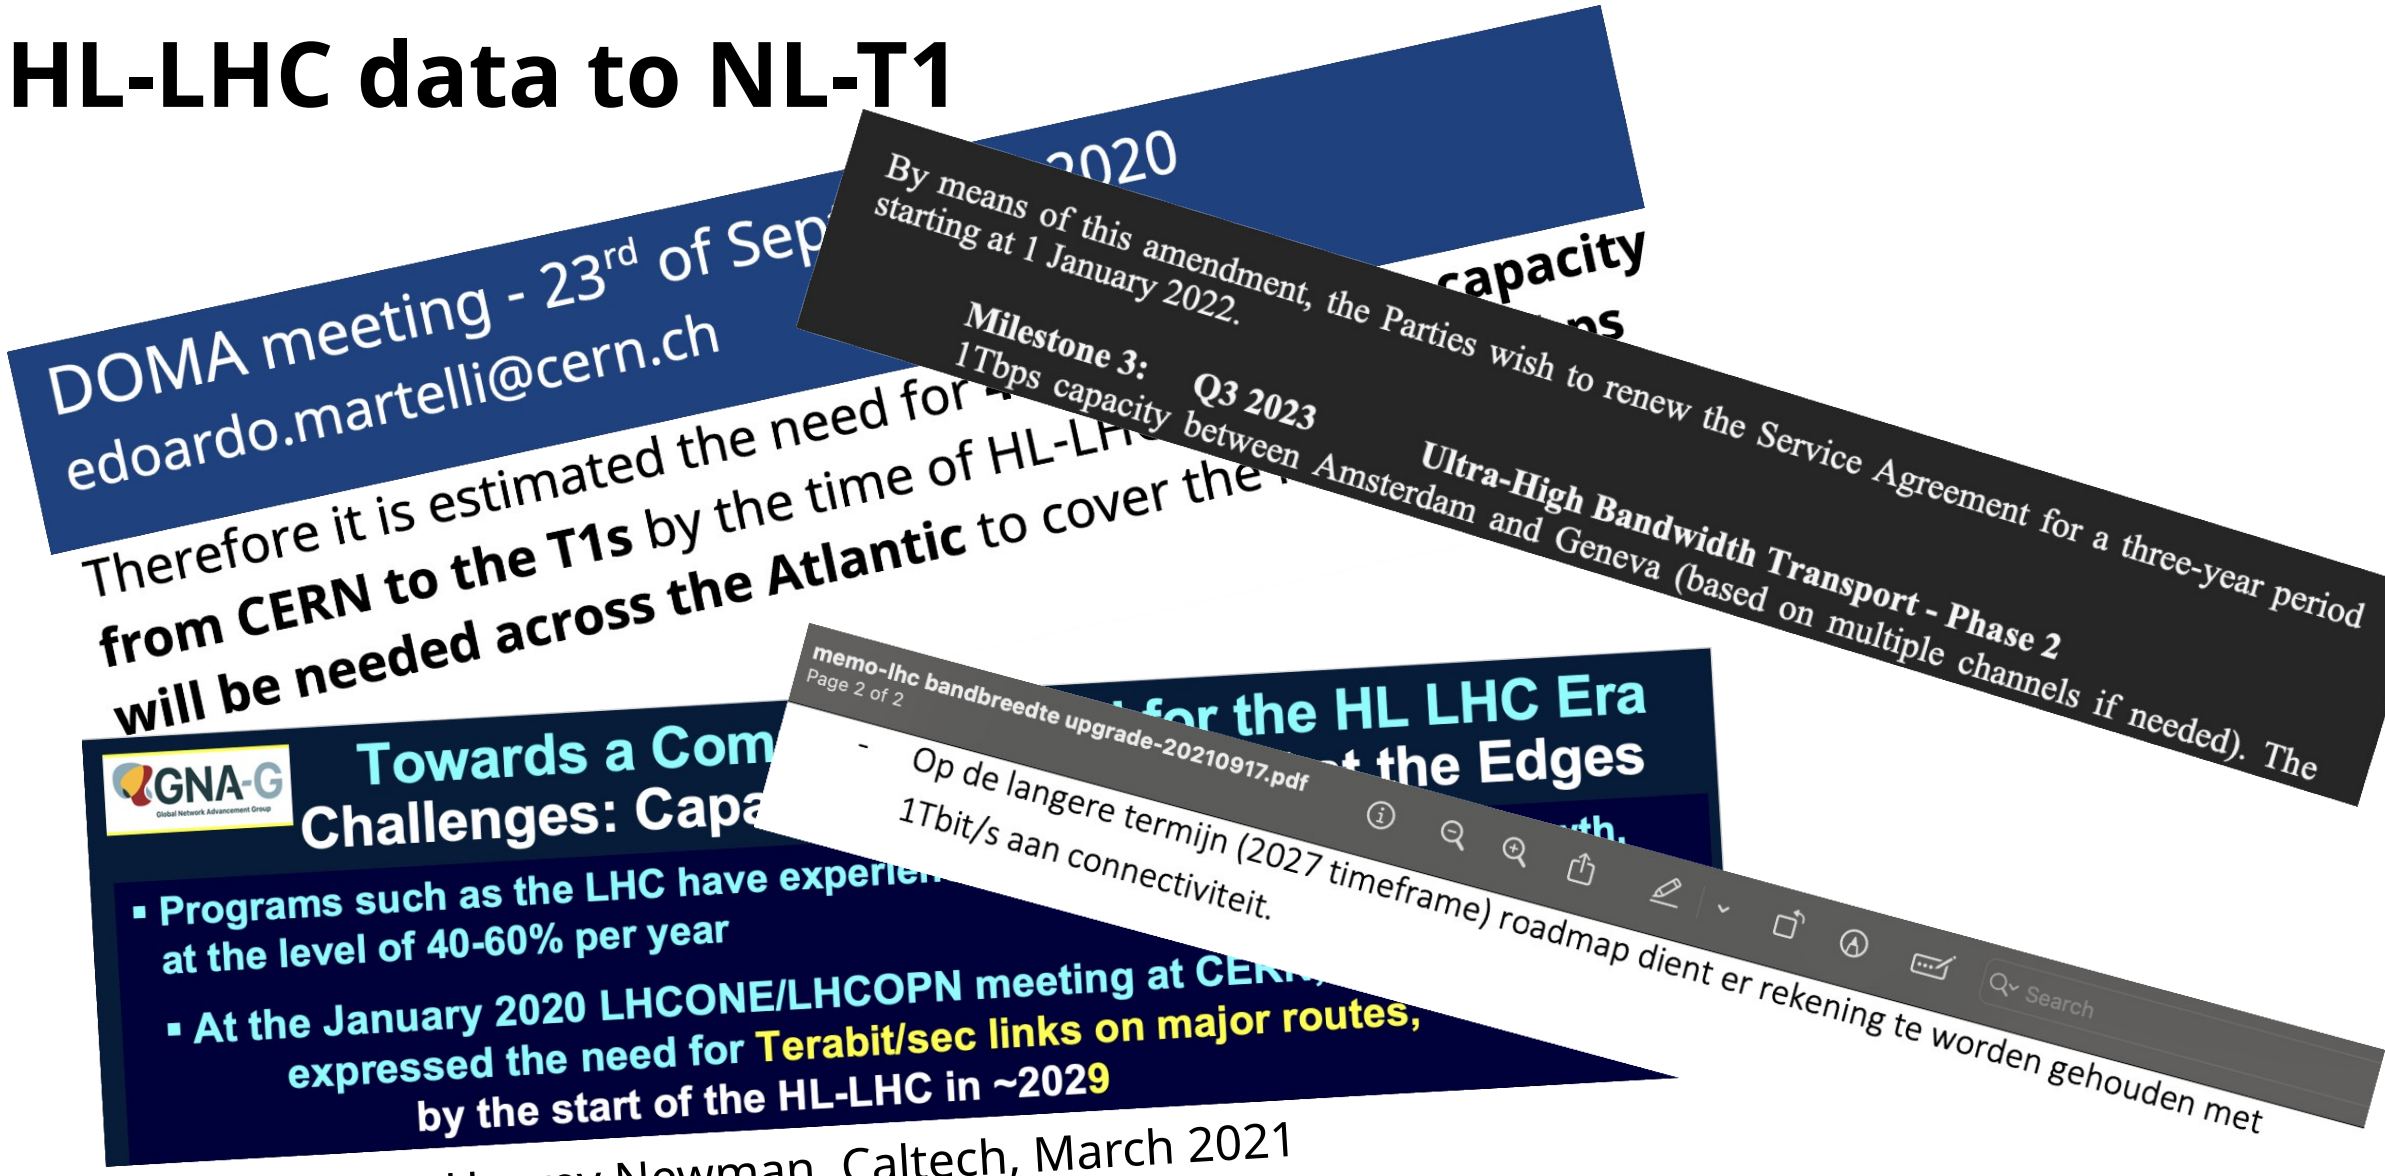

Harvey Newman, Caltech, March 2021

Data transport: tiered

Tiered storage locations need to be connected

CERN is T0 storage site

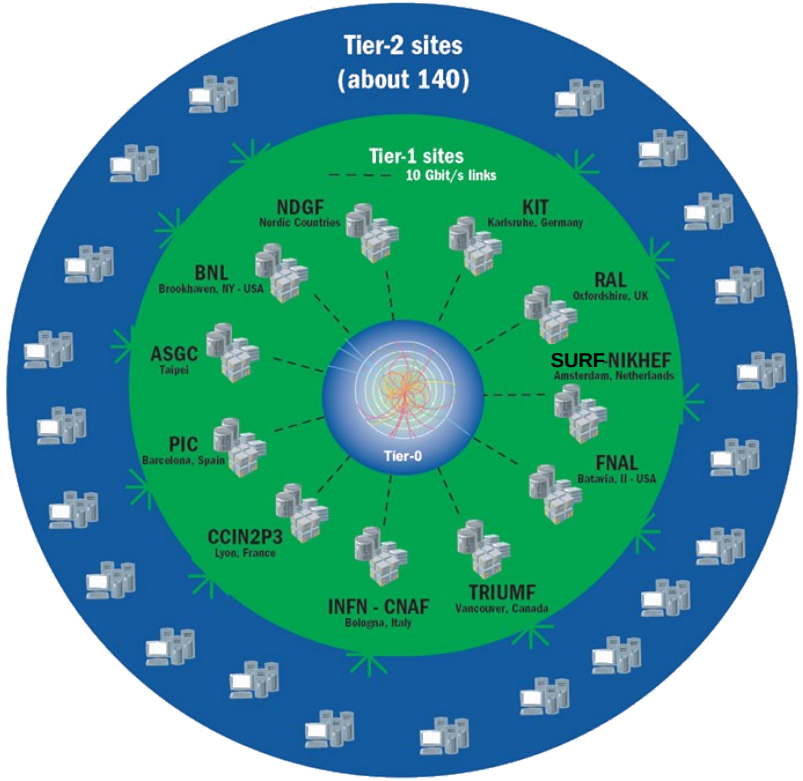

Data transport: T0 -> T1: LHCOPN

LHCOPN – Large Hadron Collider Optical Private Network

Private network between T0 and T1 storage sites.

Netherlands has a shared T1 storage facility

- SURF (formerly SURFsara)
- Nikhef (Dutch National Institute for Subatomic Physics)

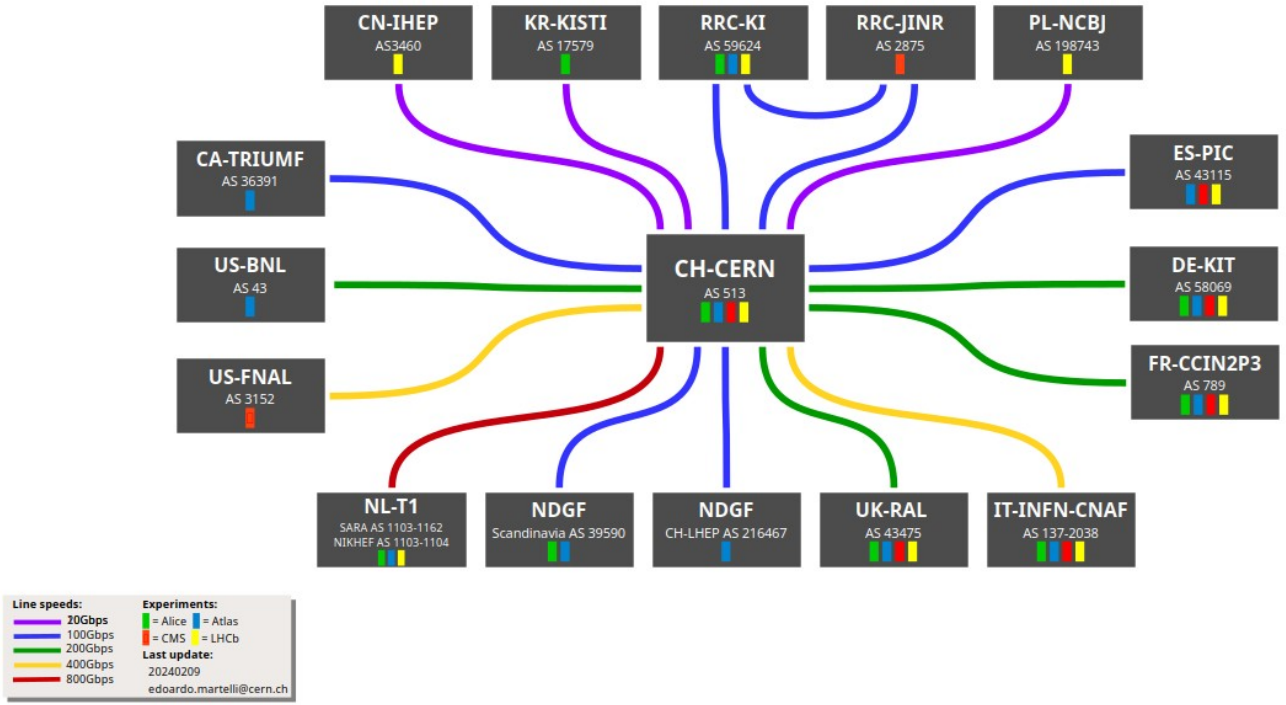

Data transportation: T1 -> T2: LHCONE

LHCONE – Large Hadron Collider Open Network Environment

Private network between T1 and T2 storage sites.

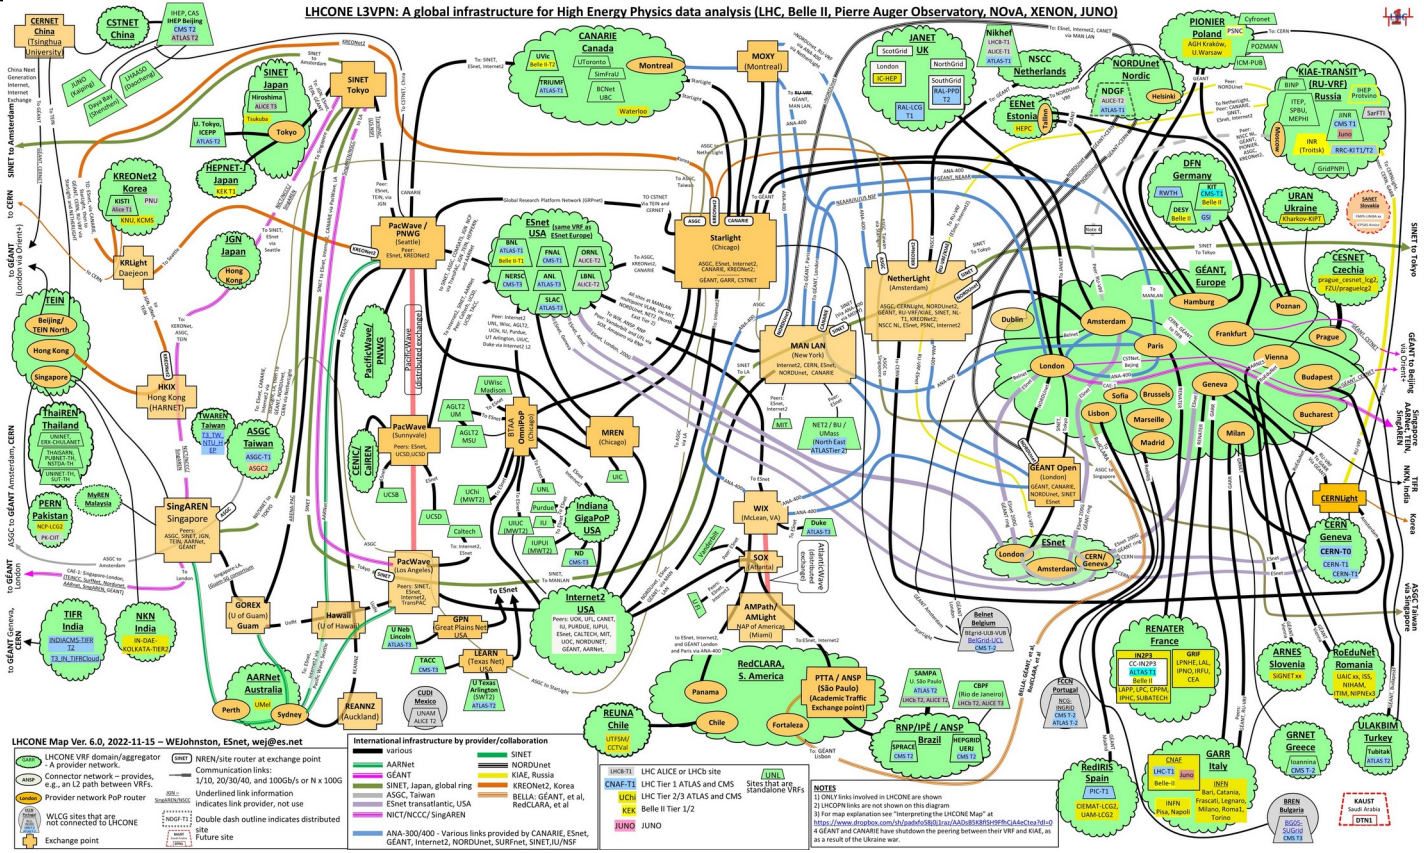

SURF’s line system on a map

Amsterdam – Geneva

- Total fiber distance (one-way) is 1648 km.

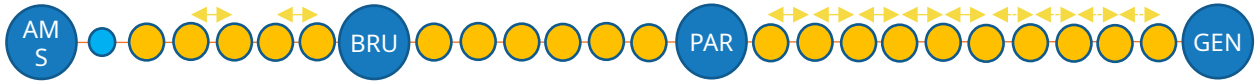

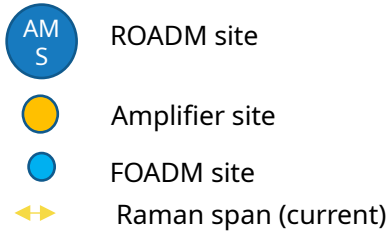

SURF’s line system between Amsterdam and Genev

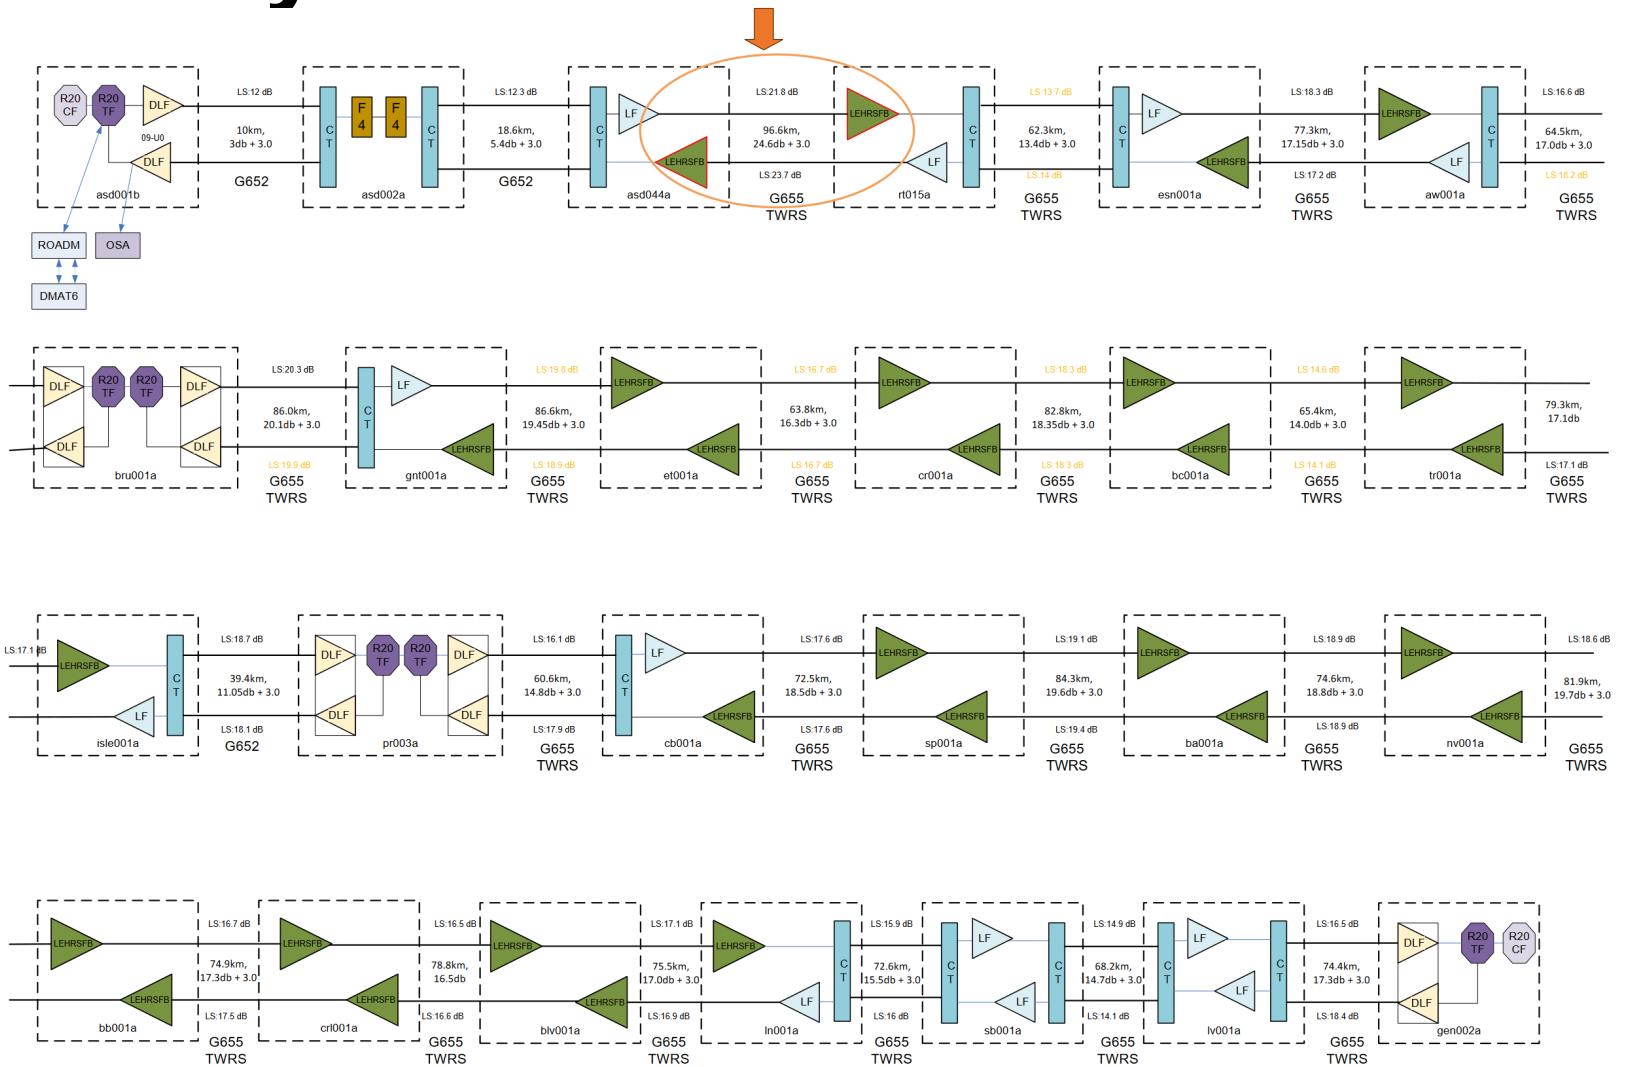

Preparing…

Q1/Q2 2023 - New software release for NMS and amplifiers

- Minor improvement of signal quality

Q3 2023 – RAMAN upgrades

- Replacing Raman for Raman with VOA and optimise OSC in scope, reducing tilt.

Q4 2023 – RAMAN installations

- Replacing EDFA type for Raman on some spans, improve link performance.

Week 8 / 2024 – Trial with Nokia transponder and IP equipment

SURF's LHCOPN Production topology (as of end 2023)

400GE capable transport system

- 1648km fiber trajectory long between Amsterdam & Geneva
- 400GE capable routers on each end
- 2x 200GE lambda between linesystem ends

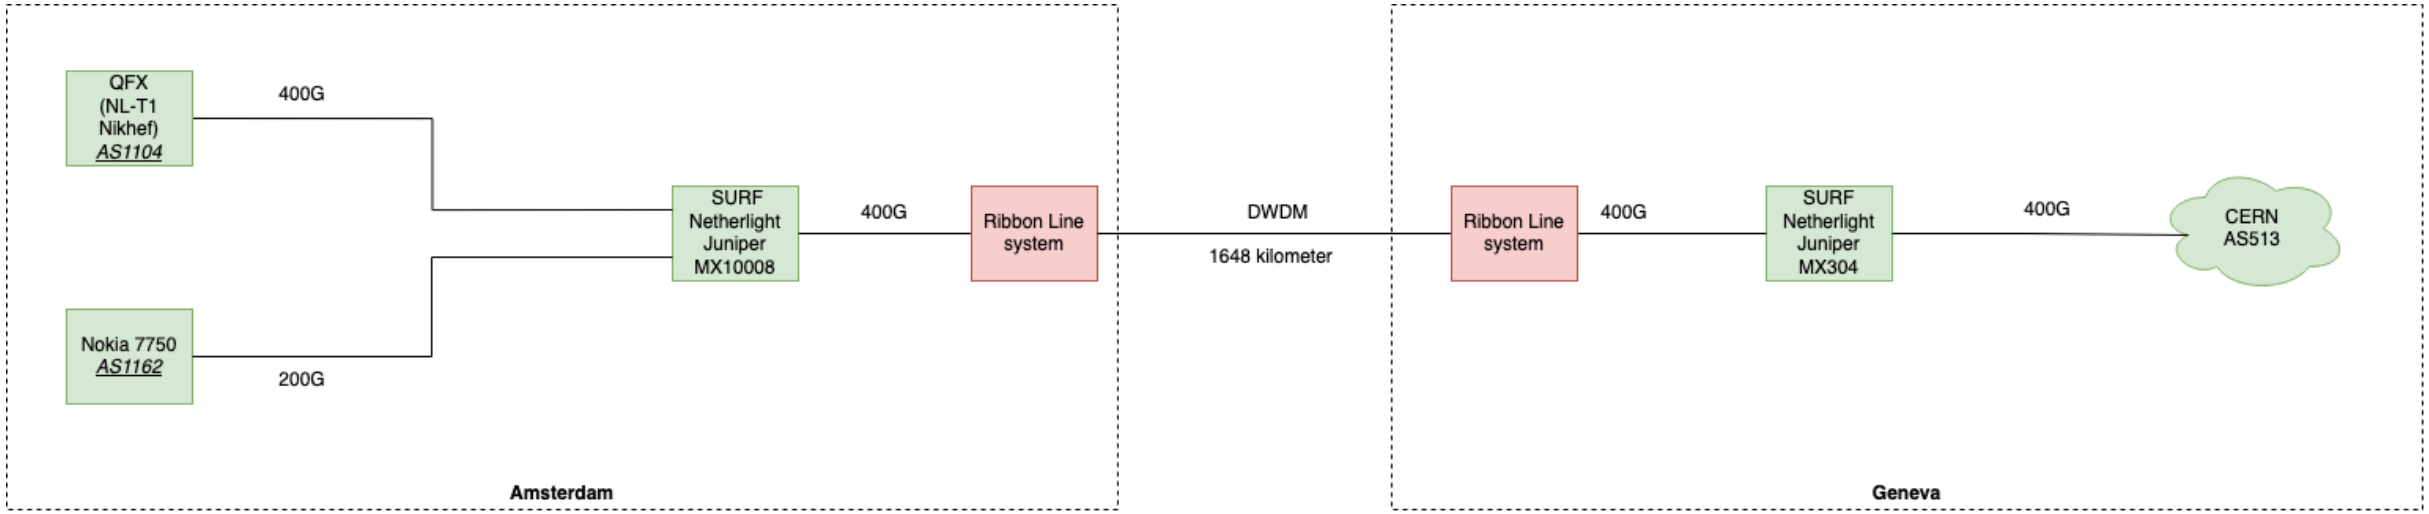

SURF's LHCOPN 1st Intermediate topology

400GE capable transport system

- Add in Nokia in between MX304 and CERN
- Moved service attach point from MX304 to Nokia

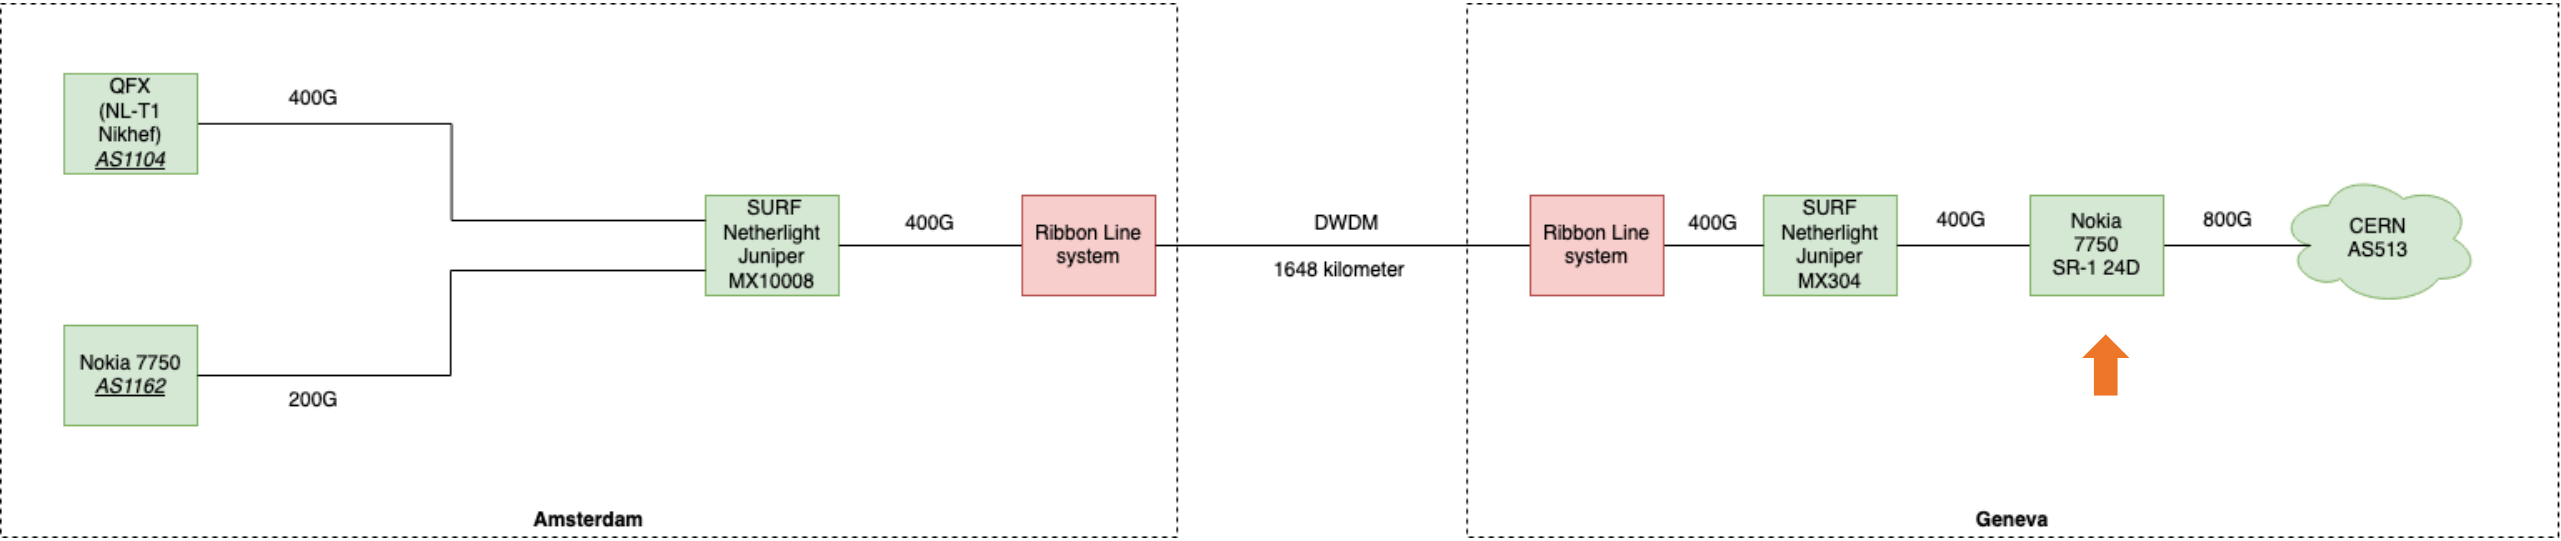

SURF's LHCOPN 2nd Intermediate topology

400GE capable transport system

- Add in Nokia between SURF and Nikhef in Amsterdam
- Add extra service attach point to Nokia

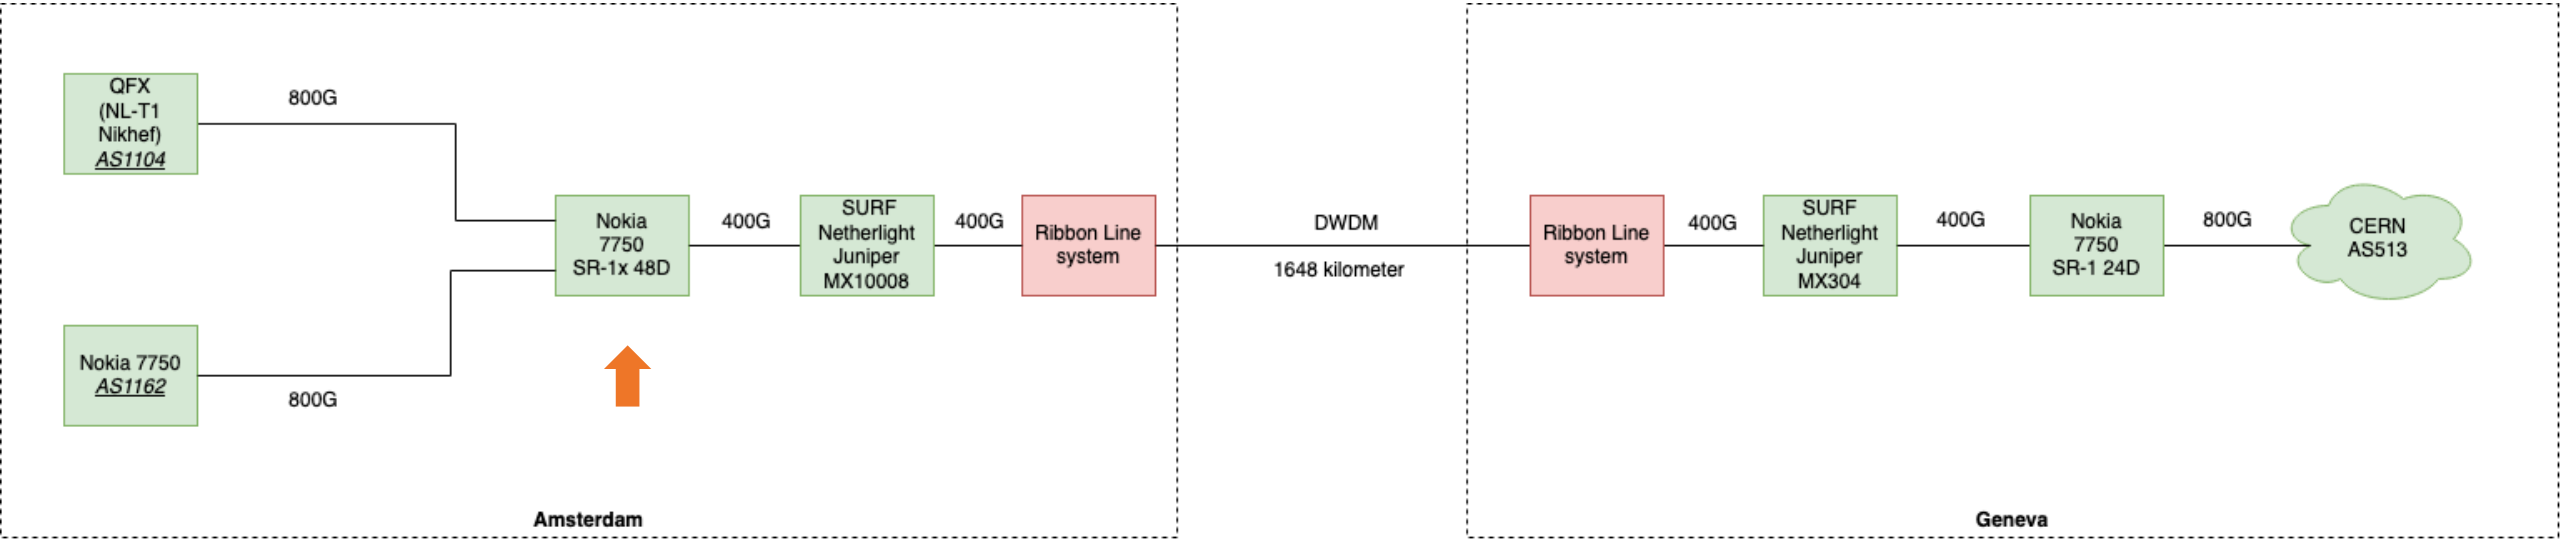

SURF's LHCOPN Final Test topology

800GE capable transport system

- 1648km fiber trajectory long between Amsterdam & Geneva
- 800GE capable routers on each end (stacking 400G towards sites)

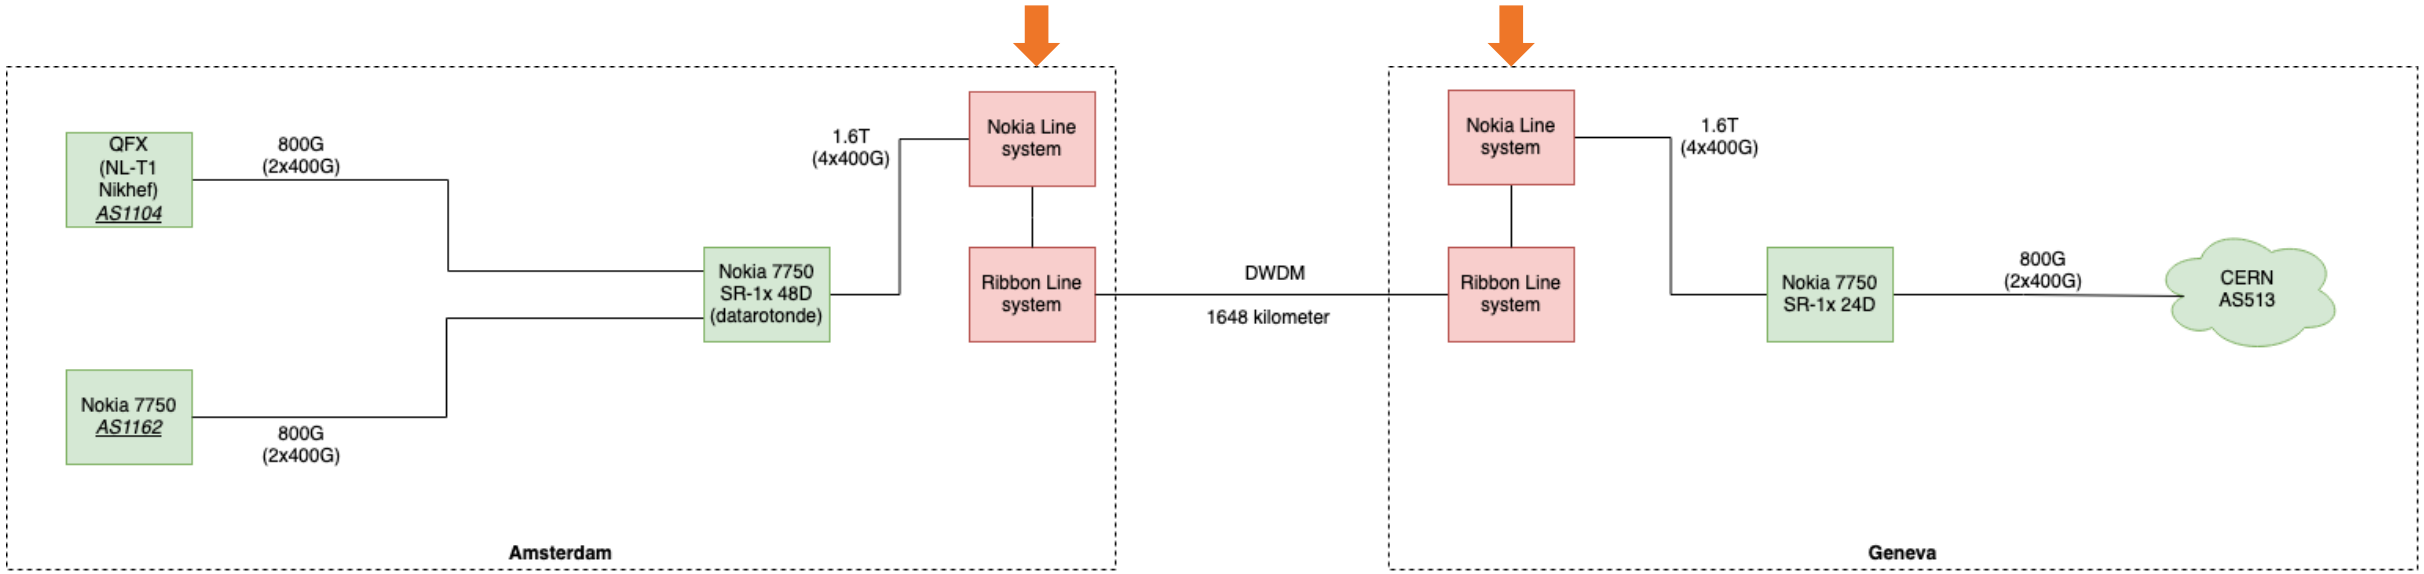

3 Experiments

We conducted 3 separate trials:

1. Optical
2. Packet generation from AMD server boxes
3. Transporting CERN ATLAS production data

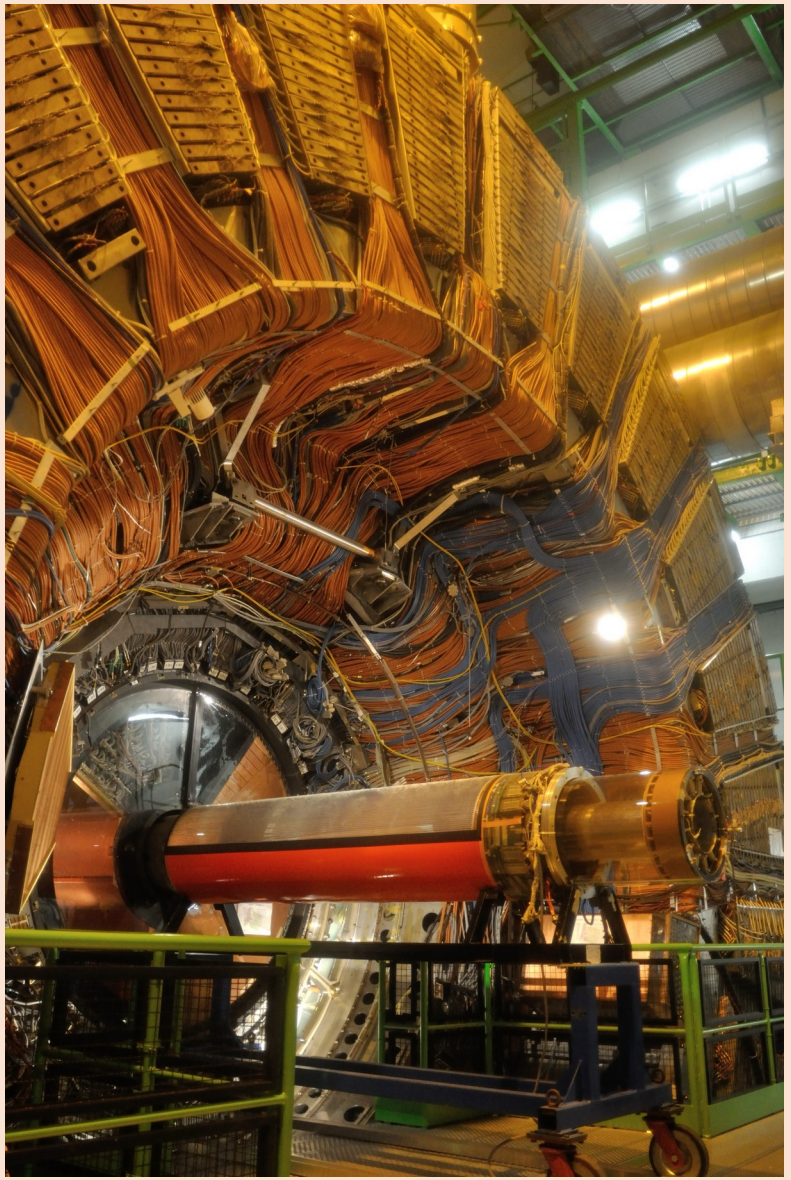

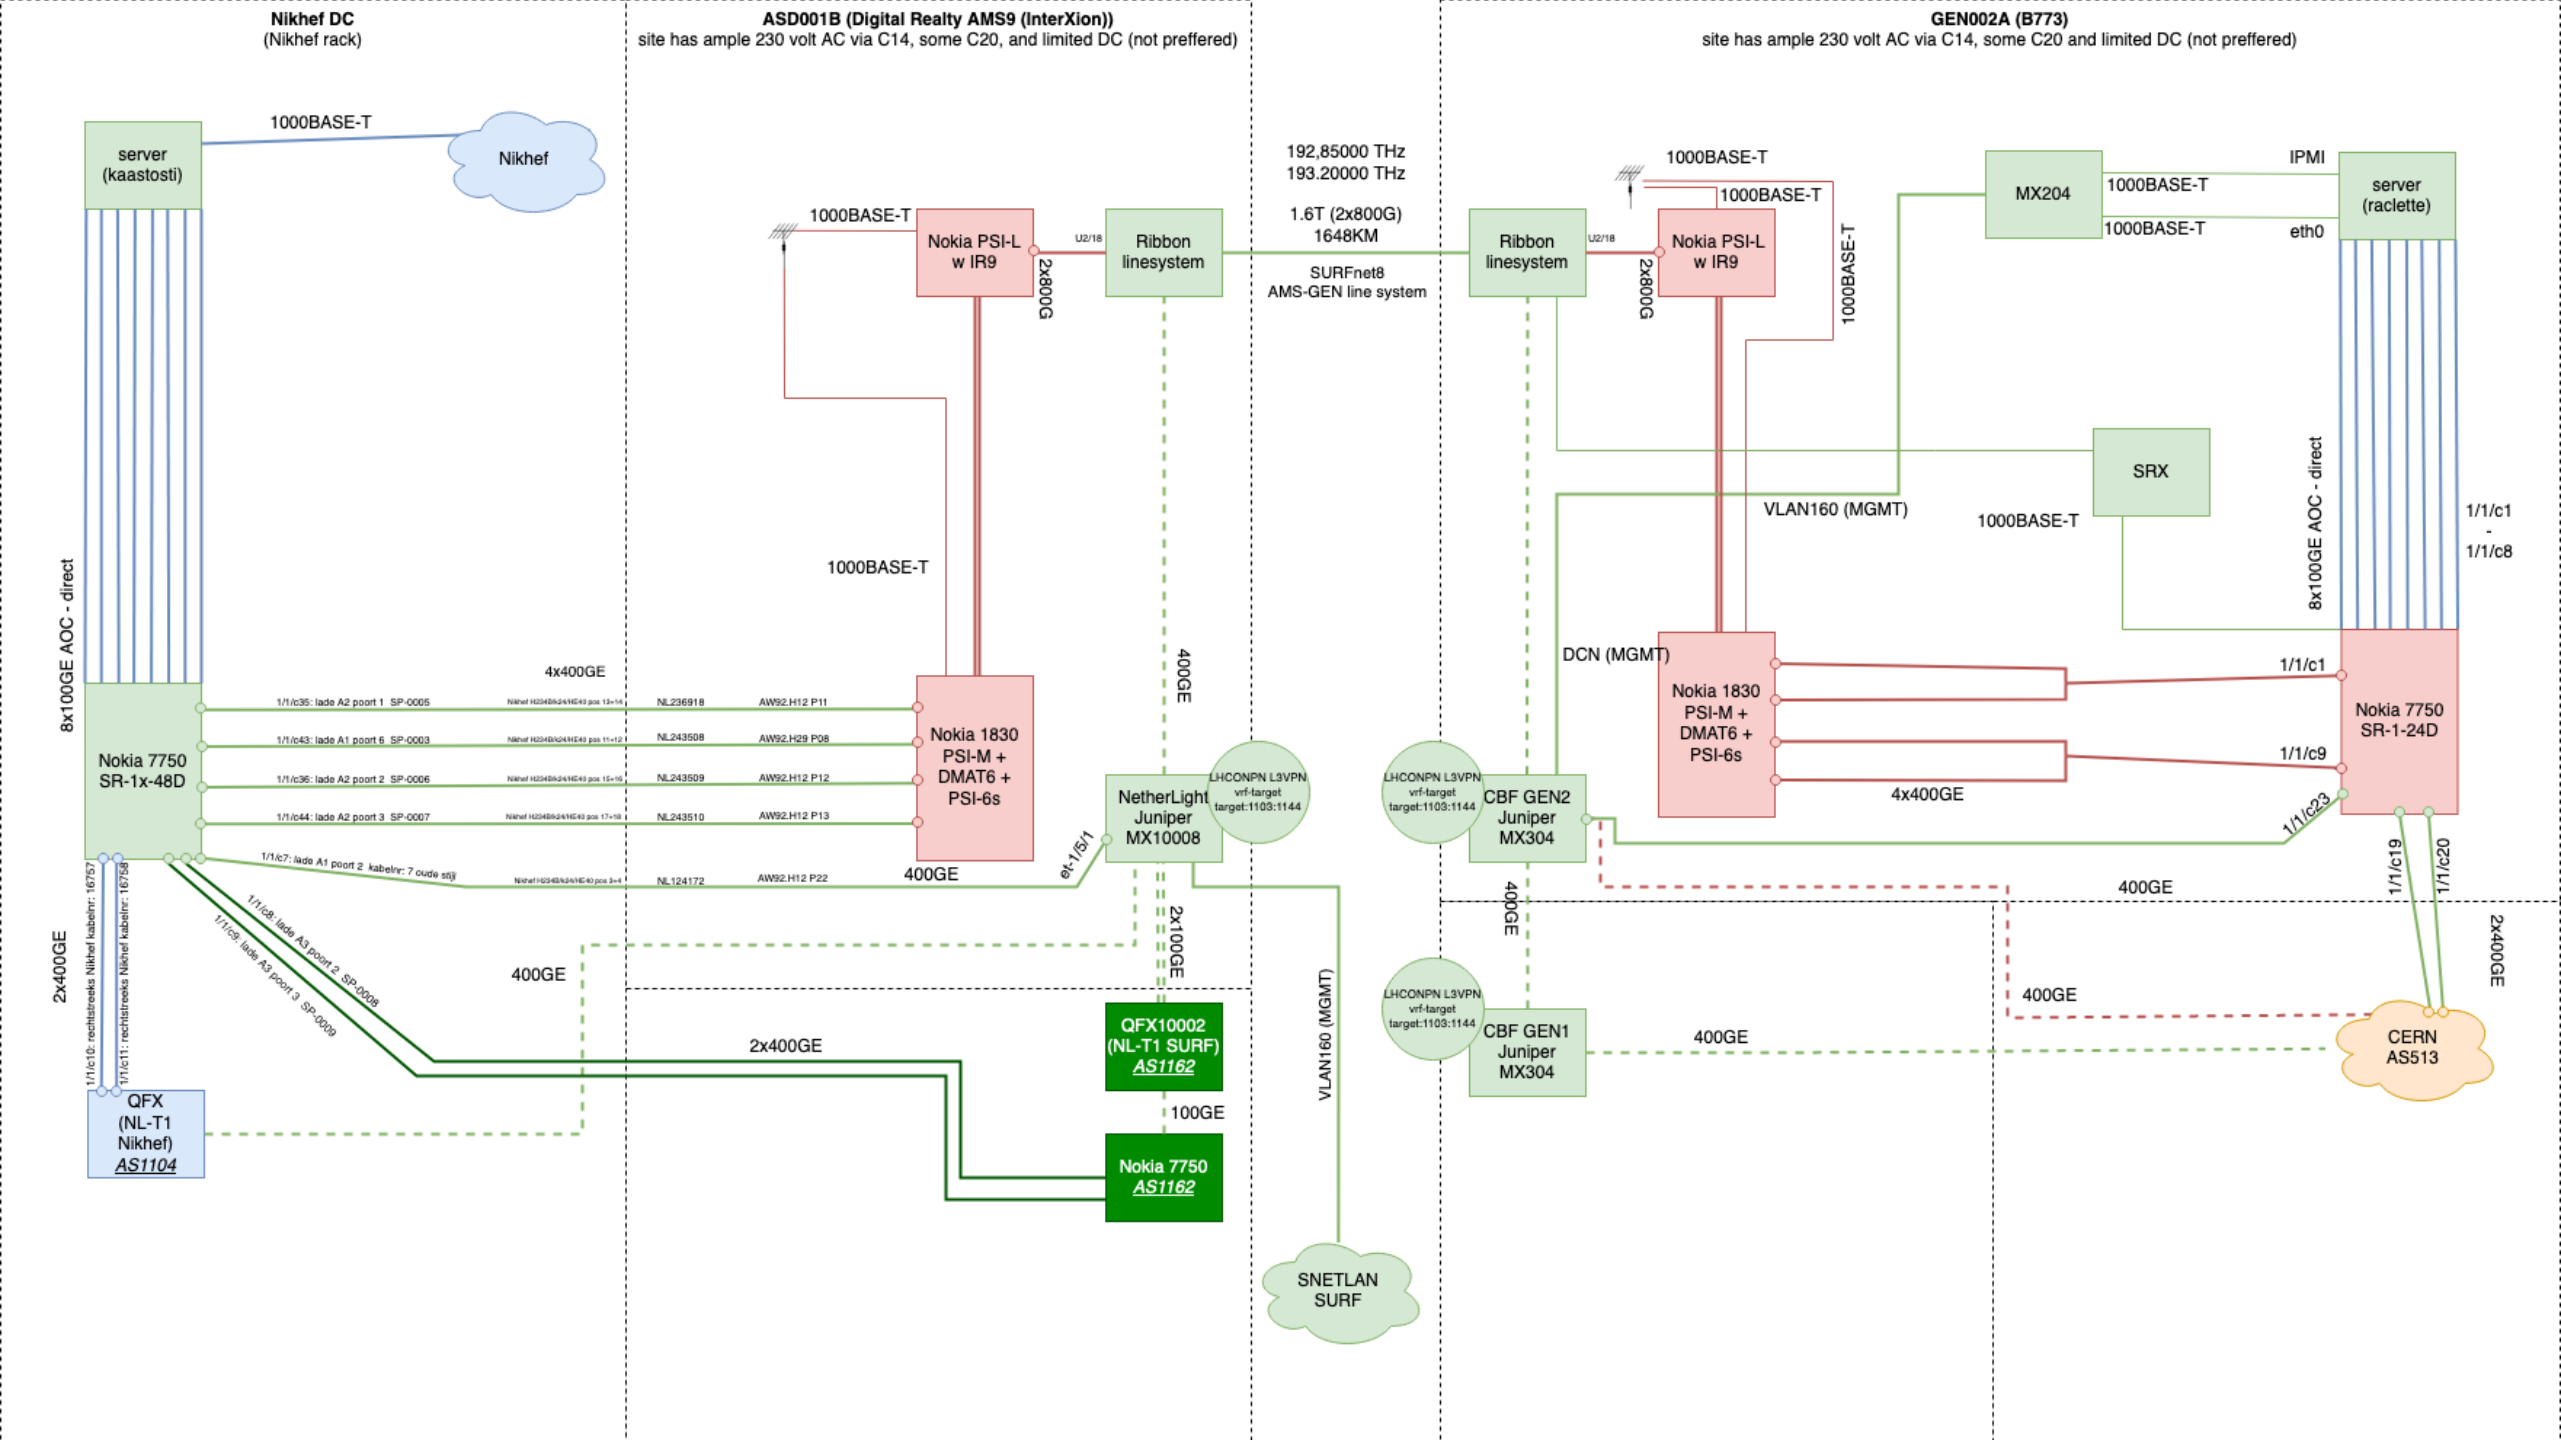
Nikhef DC
(Nikhef rack)
ASD001B (Digital Realty AMS9 (InterXion))
site has ample 230 volt AC via C14, some C20, and limited DC (not preffered)
GEN002A (B773)
site has ample 230 volt AC via C14, some C20 and limited DC (not preffered)
server (kaastosti)
1000BASE-T
Nikhef
8x100GE AOC - direct
Nokia 7750 SR-1x-48D
4x400GE
1/1/c35: lade A2 poort 1 SP-0005
1/1/c43: lade A1 poort 6 SP-0003
1/1/c36: lade A2 poort 2 SP-0006
1/1/c44: lade A2 poort 3 SP-0007
1/1/c7: lade A1 poort 2 kabelnr: 7 oude stijl
1/1/c8: lade A3 poort 2 SP-0008
1/1/c9: lade A3 poort 3 SP-0009
2x400GE
1/1/c10: rechtstreeks Nikhef kabelnr: 16757
1/1/c11: rechtstreeks Nikhef kabelnr: 16758
QFX (NL-T1 Nikhef) AS1104
400GE
1000BASE-T
Nokia PSI-L w IR9
U2/18
2x800G
Ribbon linesystem
192,85000 THz
193.20000 THz
1.6T (2x800G)
1648KM
SURFnet8
AMS-GEN line system
1000BASE-T
Nokia 1830 PSI-M + DMAT6 + PSI-6s
400GE
NetherLight Juniper MX10008
LHCONPN L3VPN vrf-target target:1103:1144
et-1/5/1
400GE
2x100GE
QFX10002 (NL-T1 SURF) AS1162
100GE
Nokia 7750 AS1162
VLAN160 (MGMT)
SNETLAN SURF
1000BASE-T
Nokia PSI-L w IR9
U2/18
2x800G
Ribbon linesystem
1000BASE-T
MX204
IPMI
1000BASE-T
1000BASE-T
eth0
server (raclette)
SRX
VLAN160 (MGMT)
1000BASE-T
8x100GE AOC - direct
1/1/c1 - 1/1/c8
DCN (MGMT)
Nokia 1830 PSI-M + DMAT6 + PSI-6s
1/1/c1
1/1/c9
4x400GE
Nokia 7750 SR-1-24D
1/1/c23
1/1/c19
1/1/c20
2x400GE
400GE
LHCONPN L3VPN vrf-target target:1103:1144
CBF GEN2 Juniper MX304
400GE
LHCONPN L3VPN vrf-target target:1103:1144
CBF GEN1 Juniper MX304
400GE
400GE
CERN AS513

Hardware used: optical transport

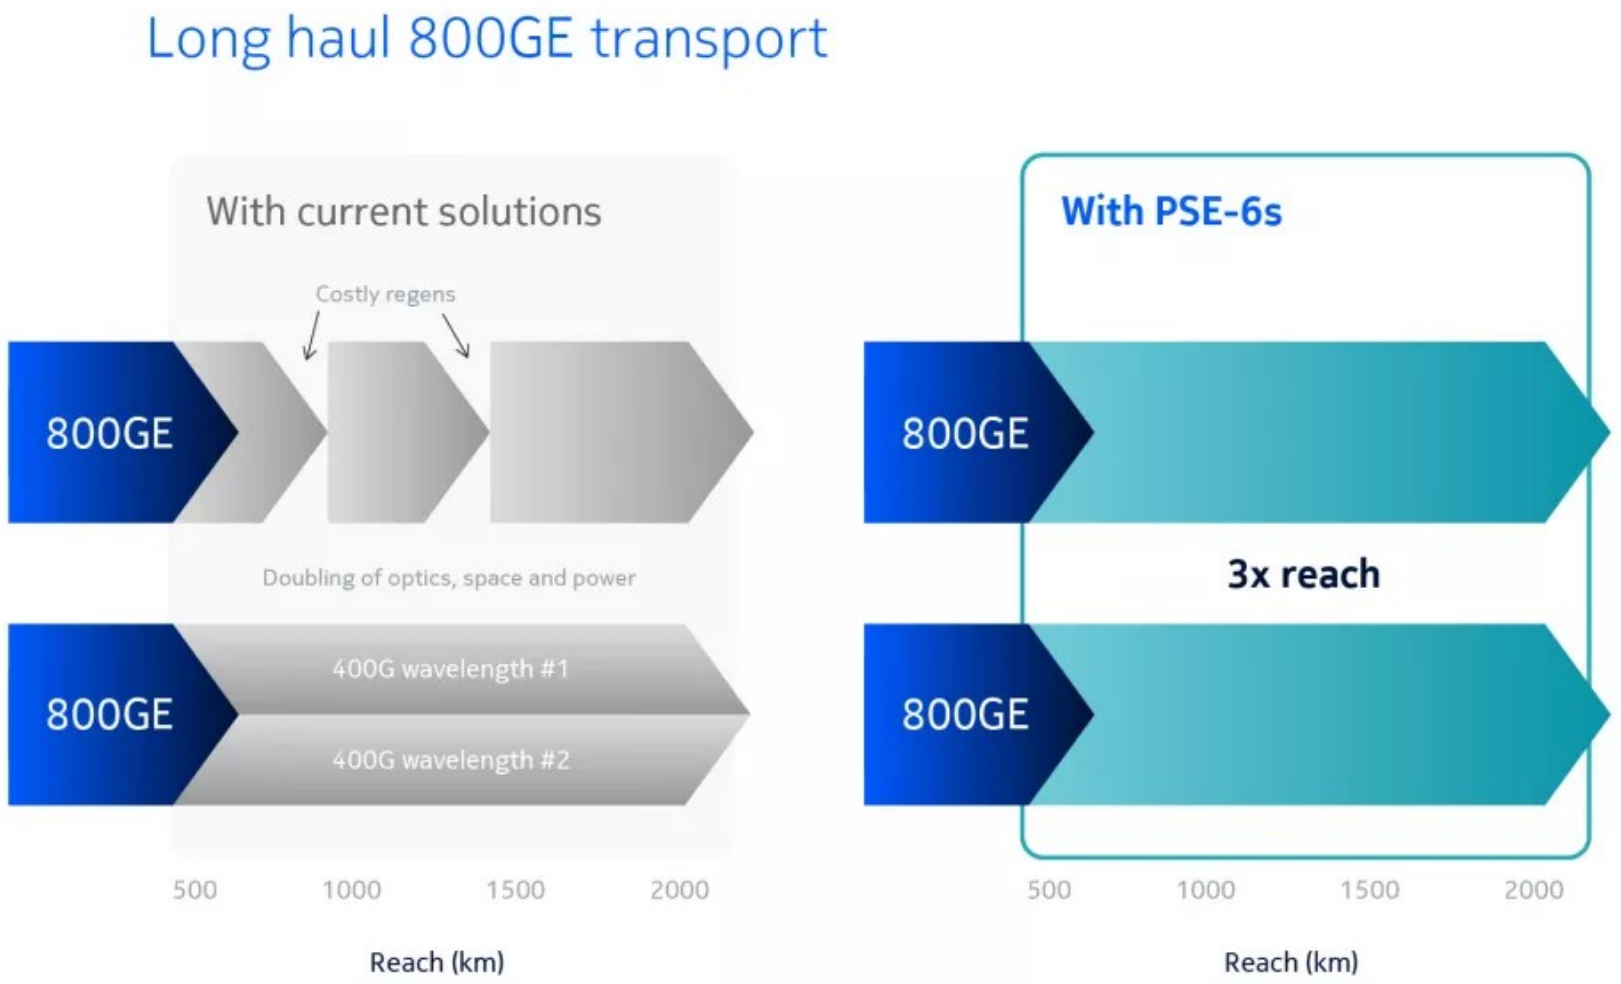

Hardware used: SURF IP

Nokia FP5 based SR-1 routers

Amsterdam

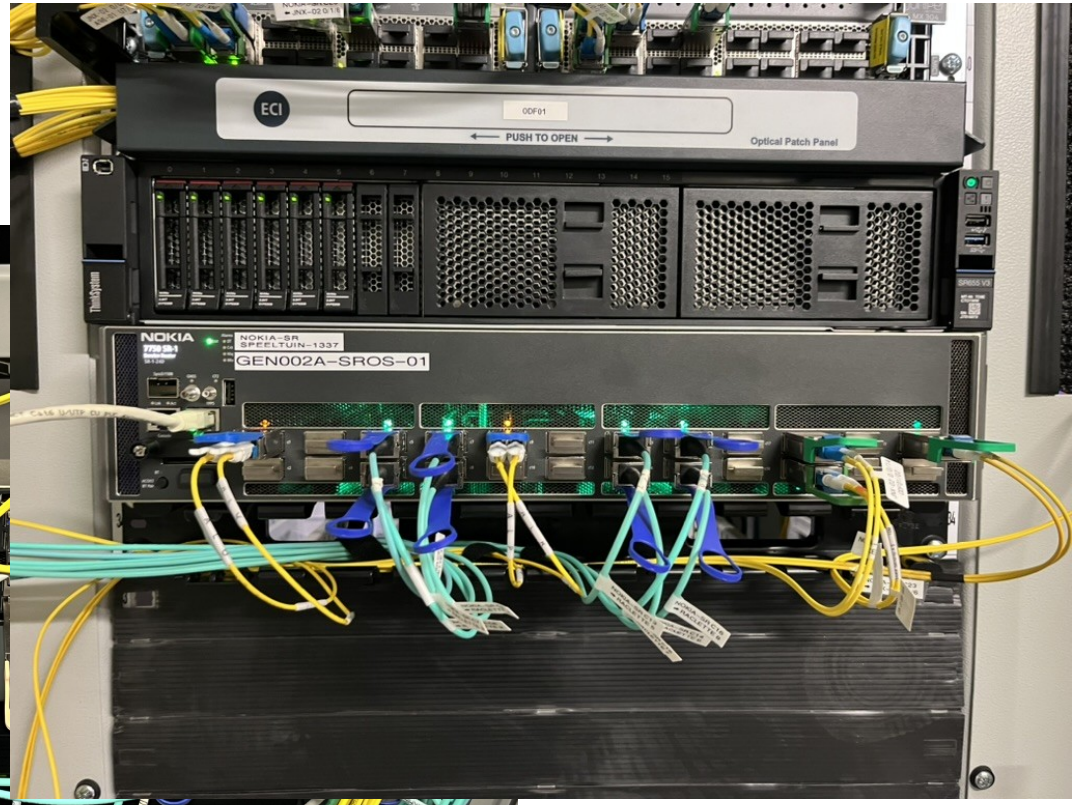

Geneva

SURF’s LHCOPN 800G test: Geneva / CERN side

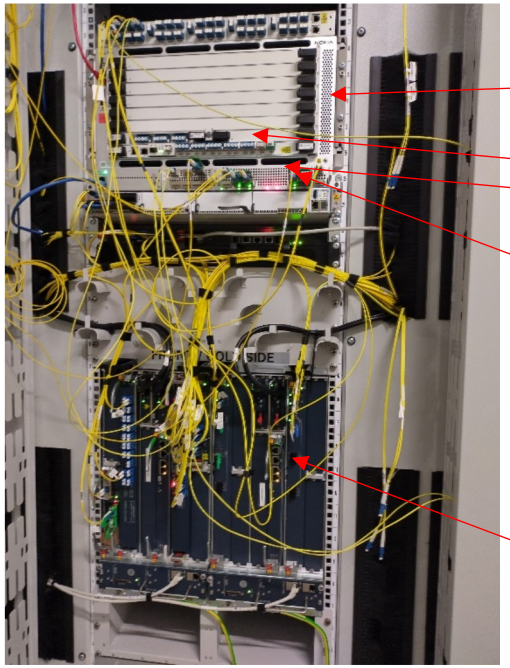

PSI-L

IR9-LP

PSIM
+DMAT6+
SFM6

Ribbon
production
chassis

EXFO OSA

SURF’s LHCOPN 800G test: Amsterdam side

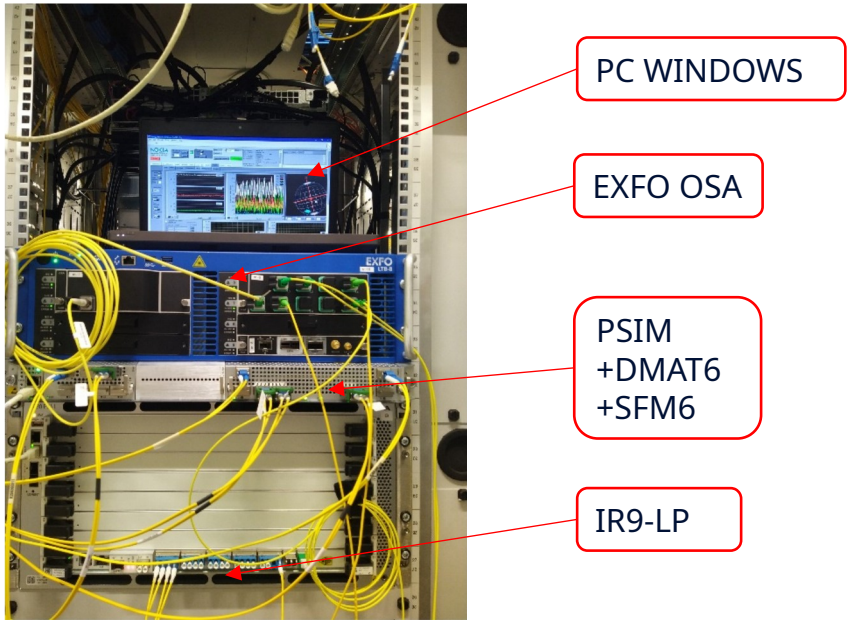

Total throughput Amsterdam <-> Geneva

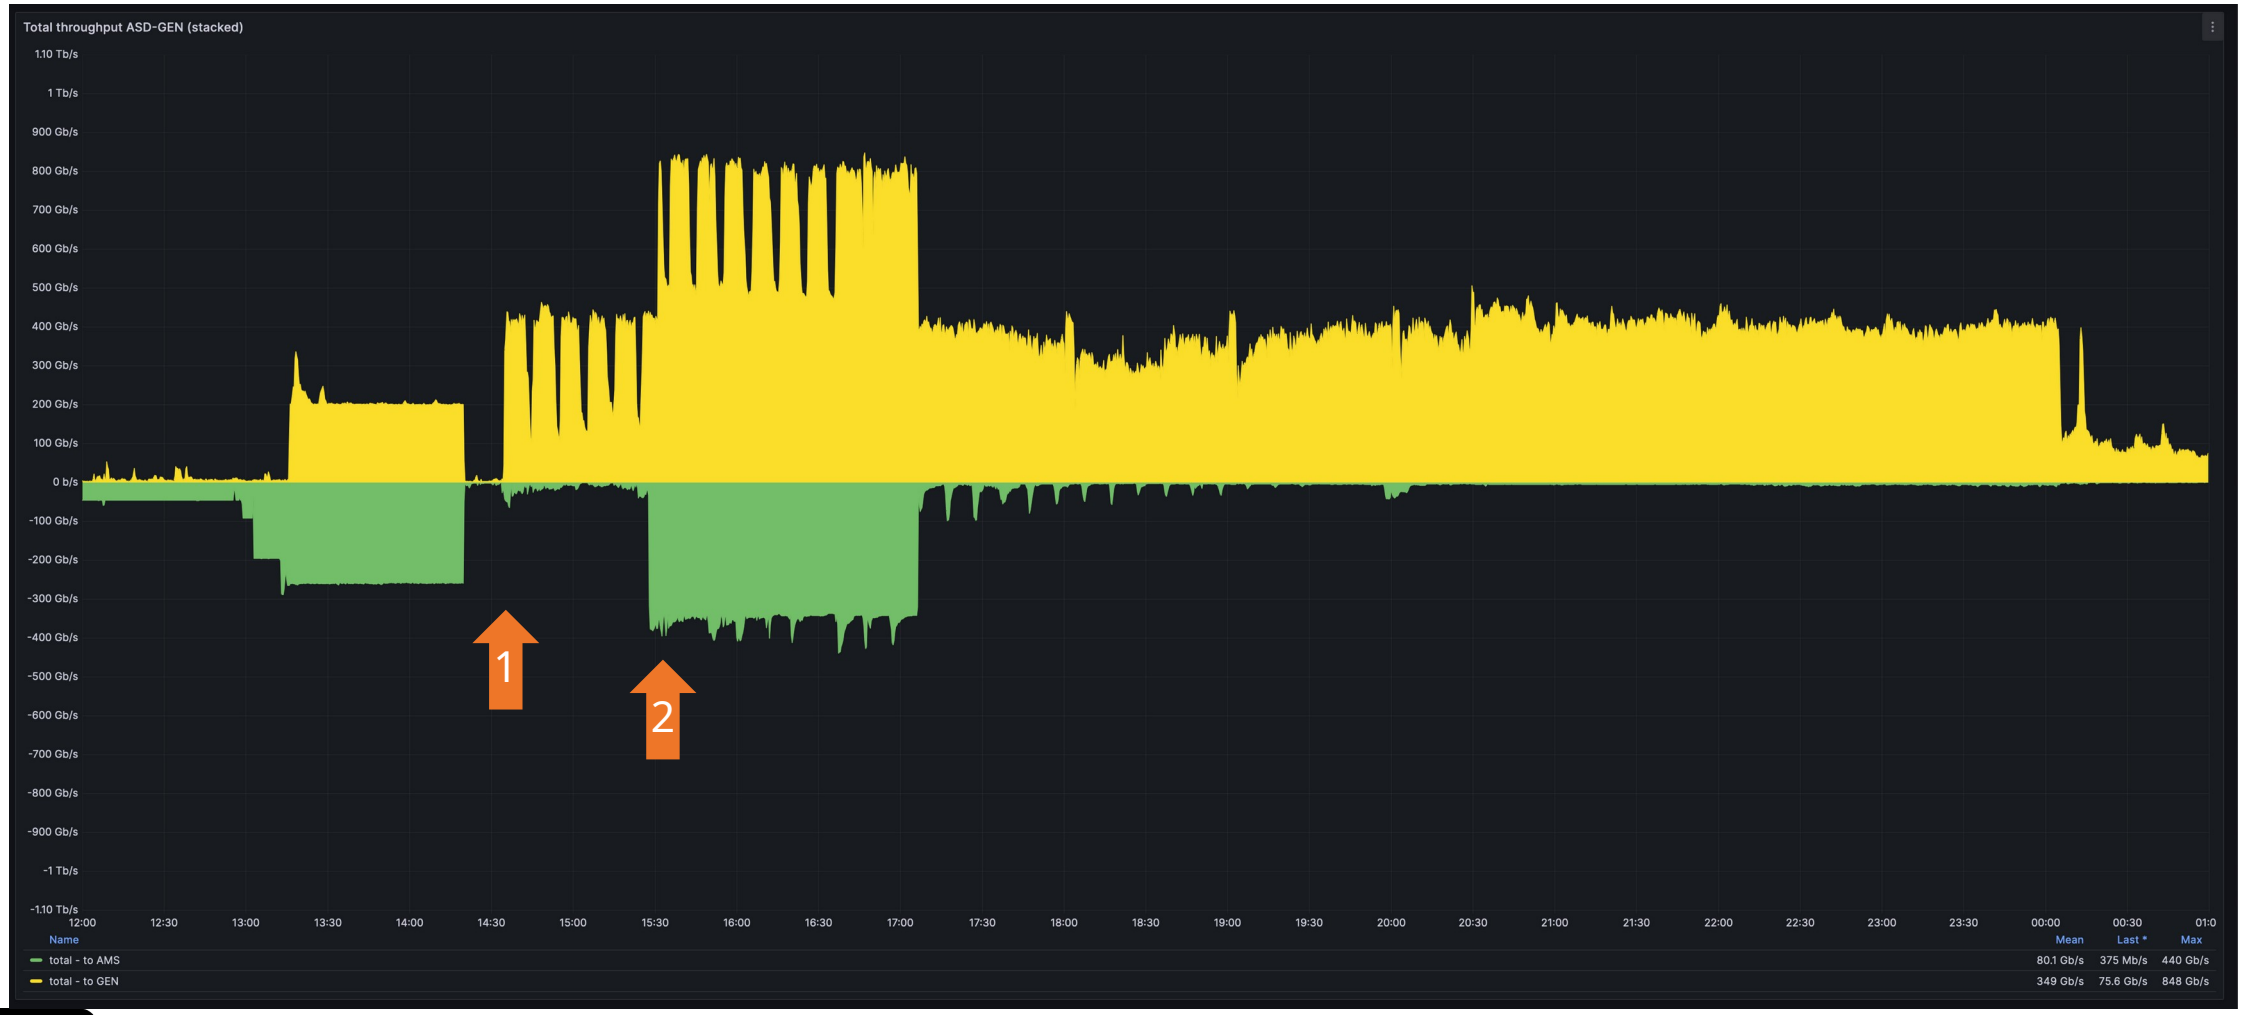

Last day of ATLAS production testing only: GEN->AMS traffic

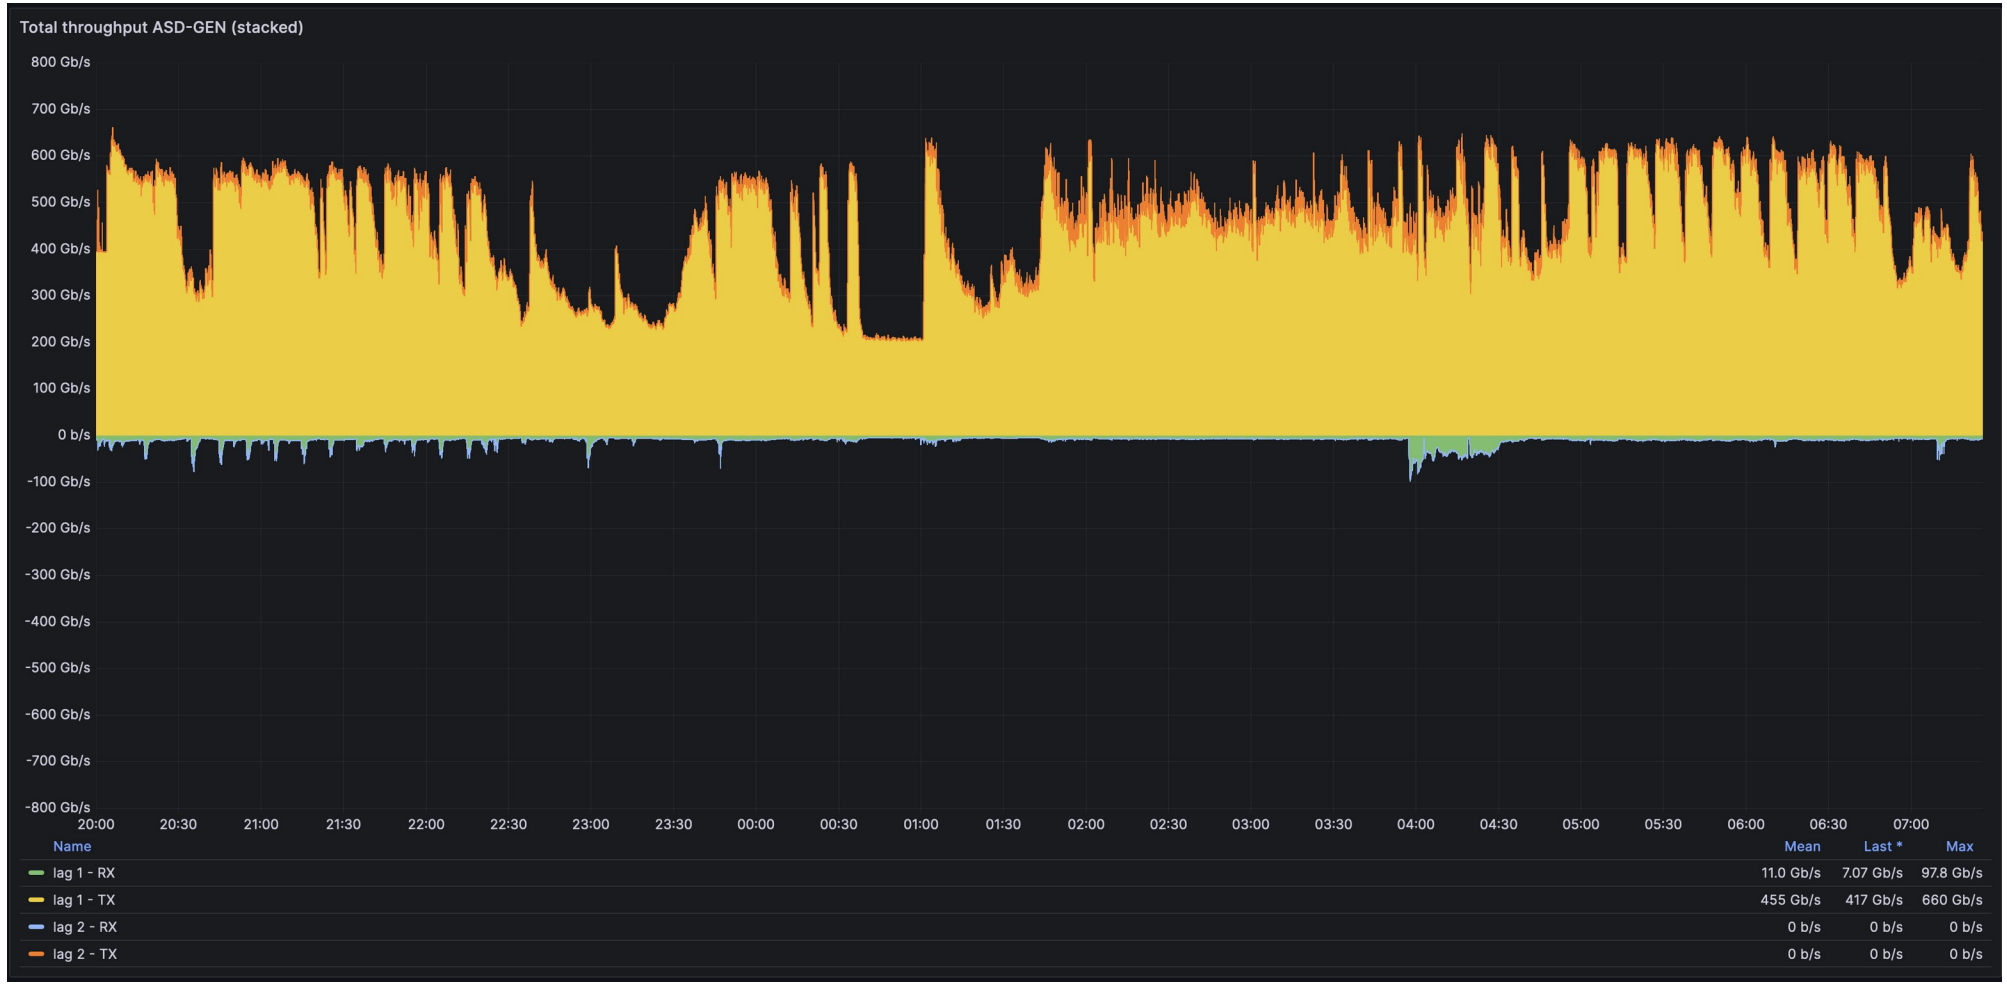

Last day ATLAS production testing: 2 storage sites in AMS

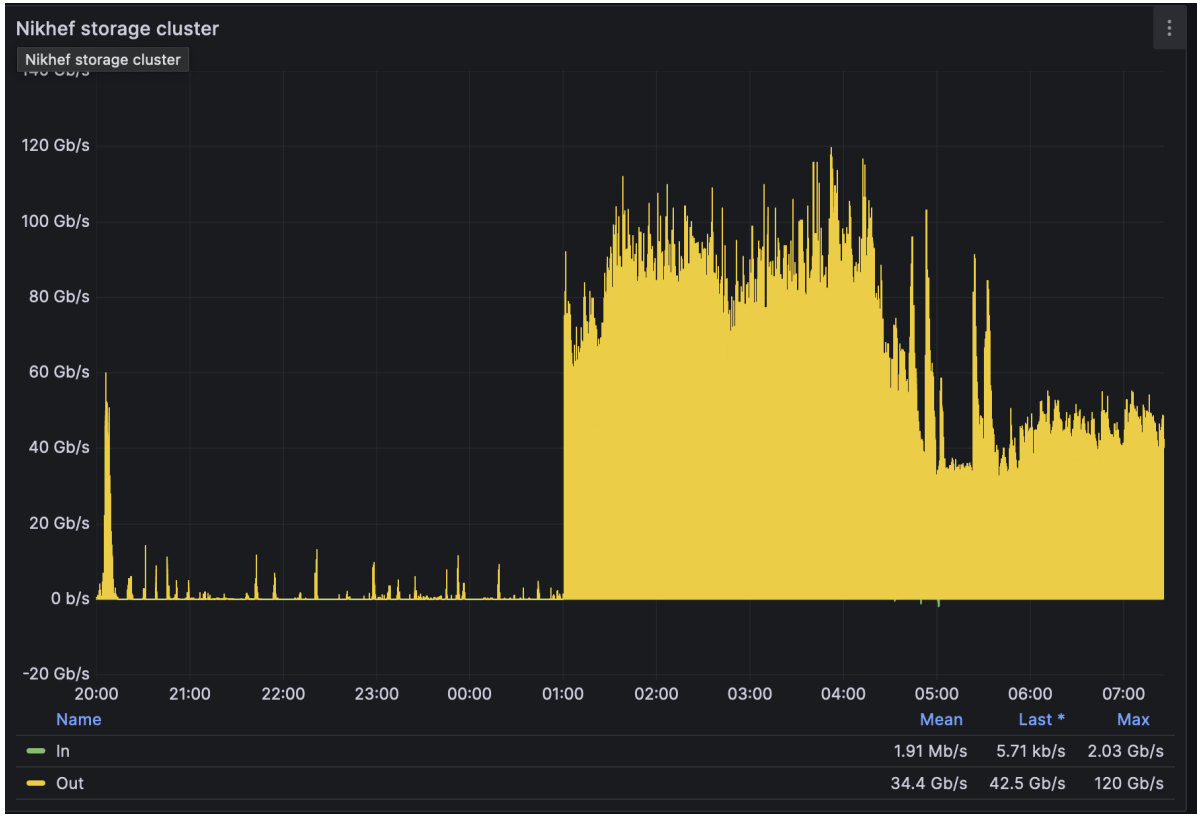

NIKHEF storage site

SURF storage site

dCache storage software throughput

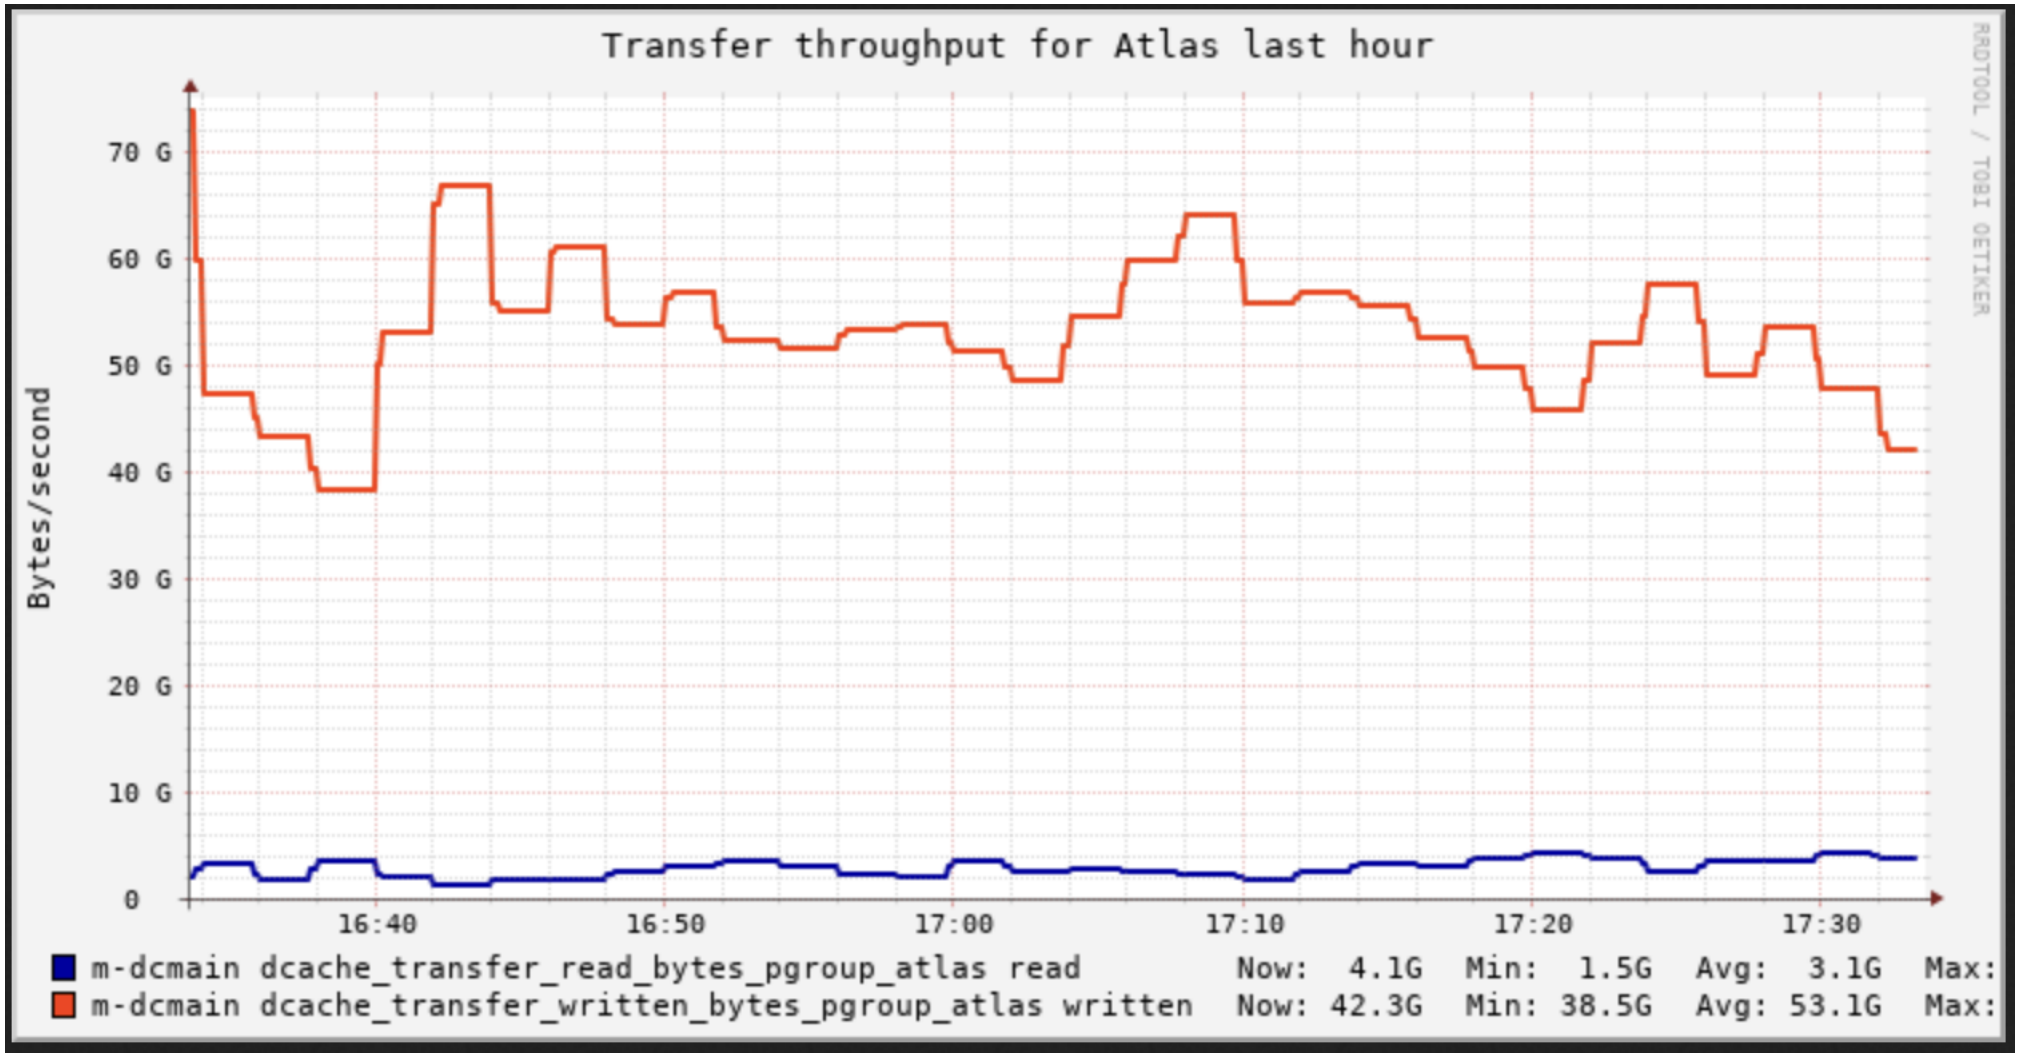

dCache power consumption

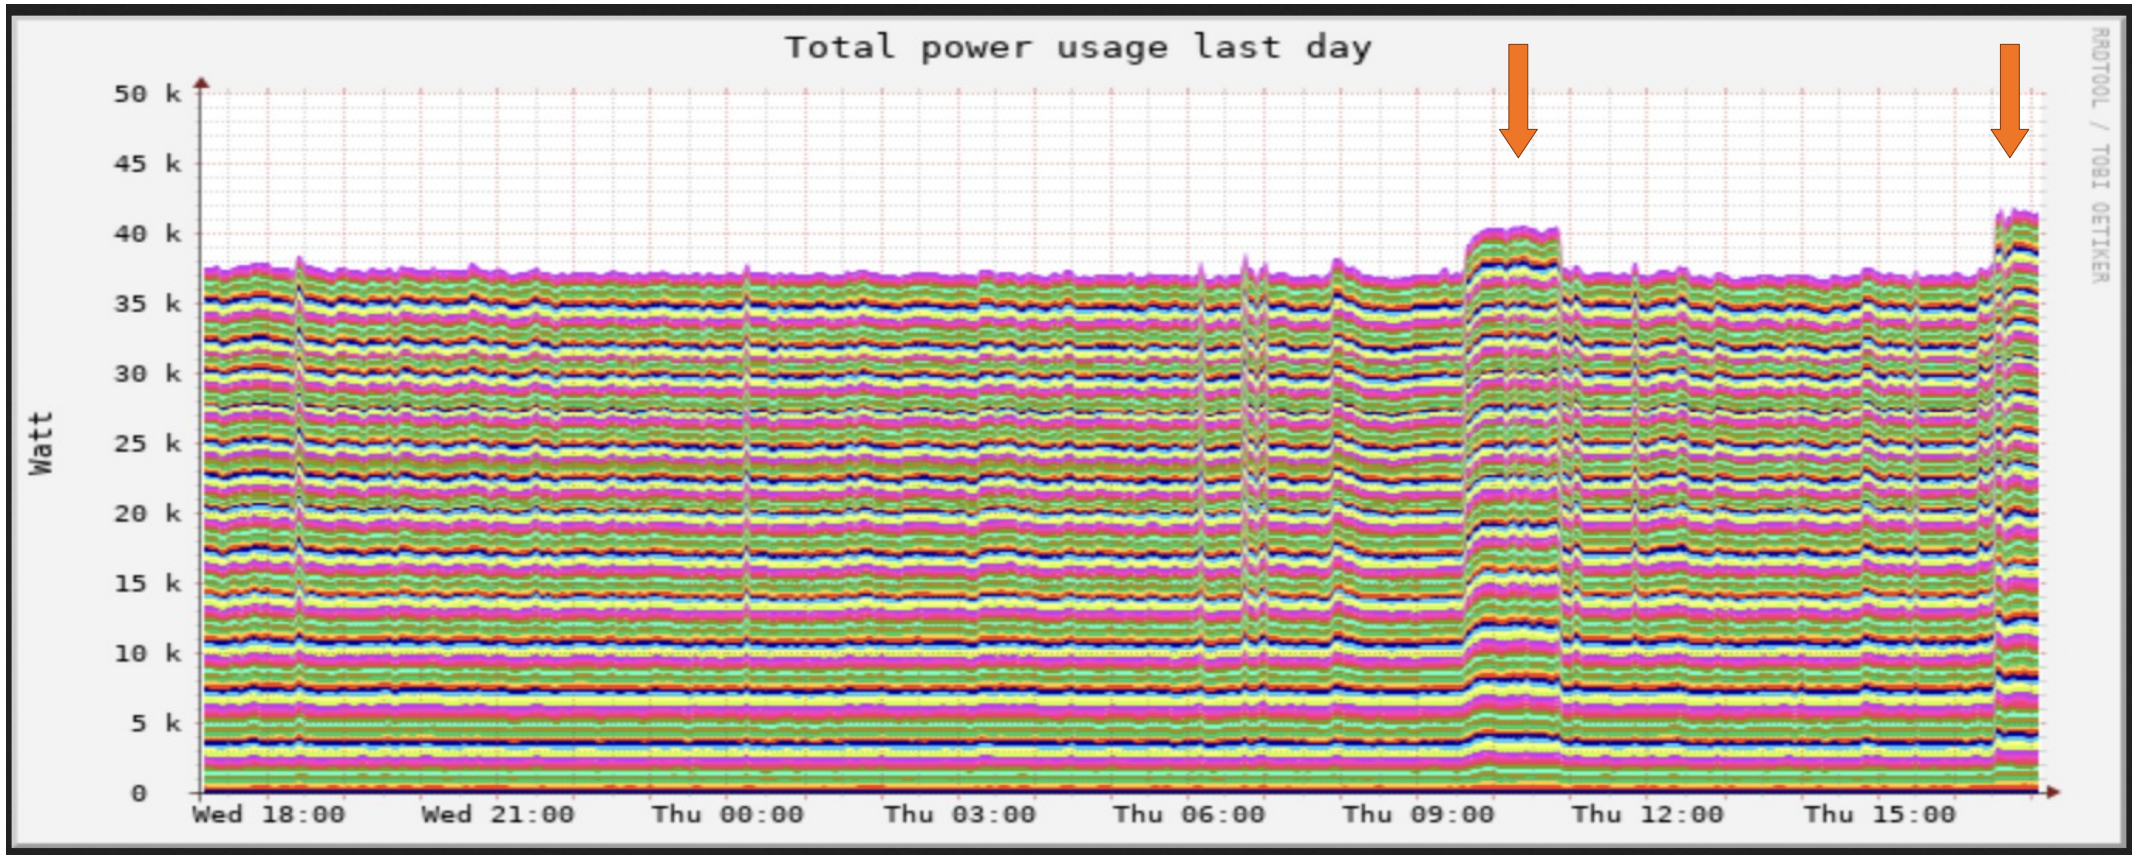

Spectrum analyzer: CERN

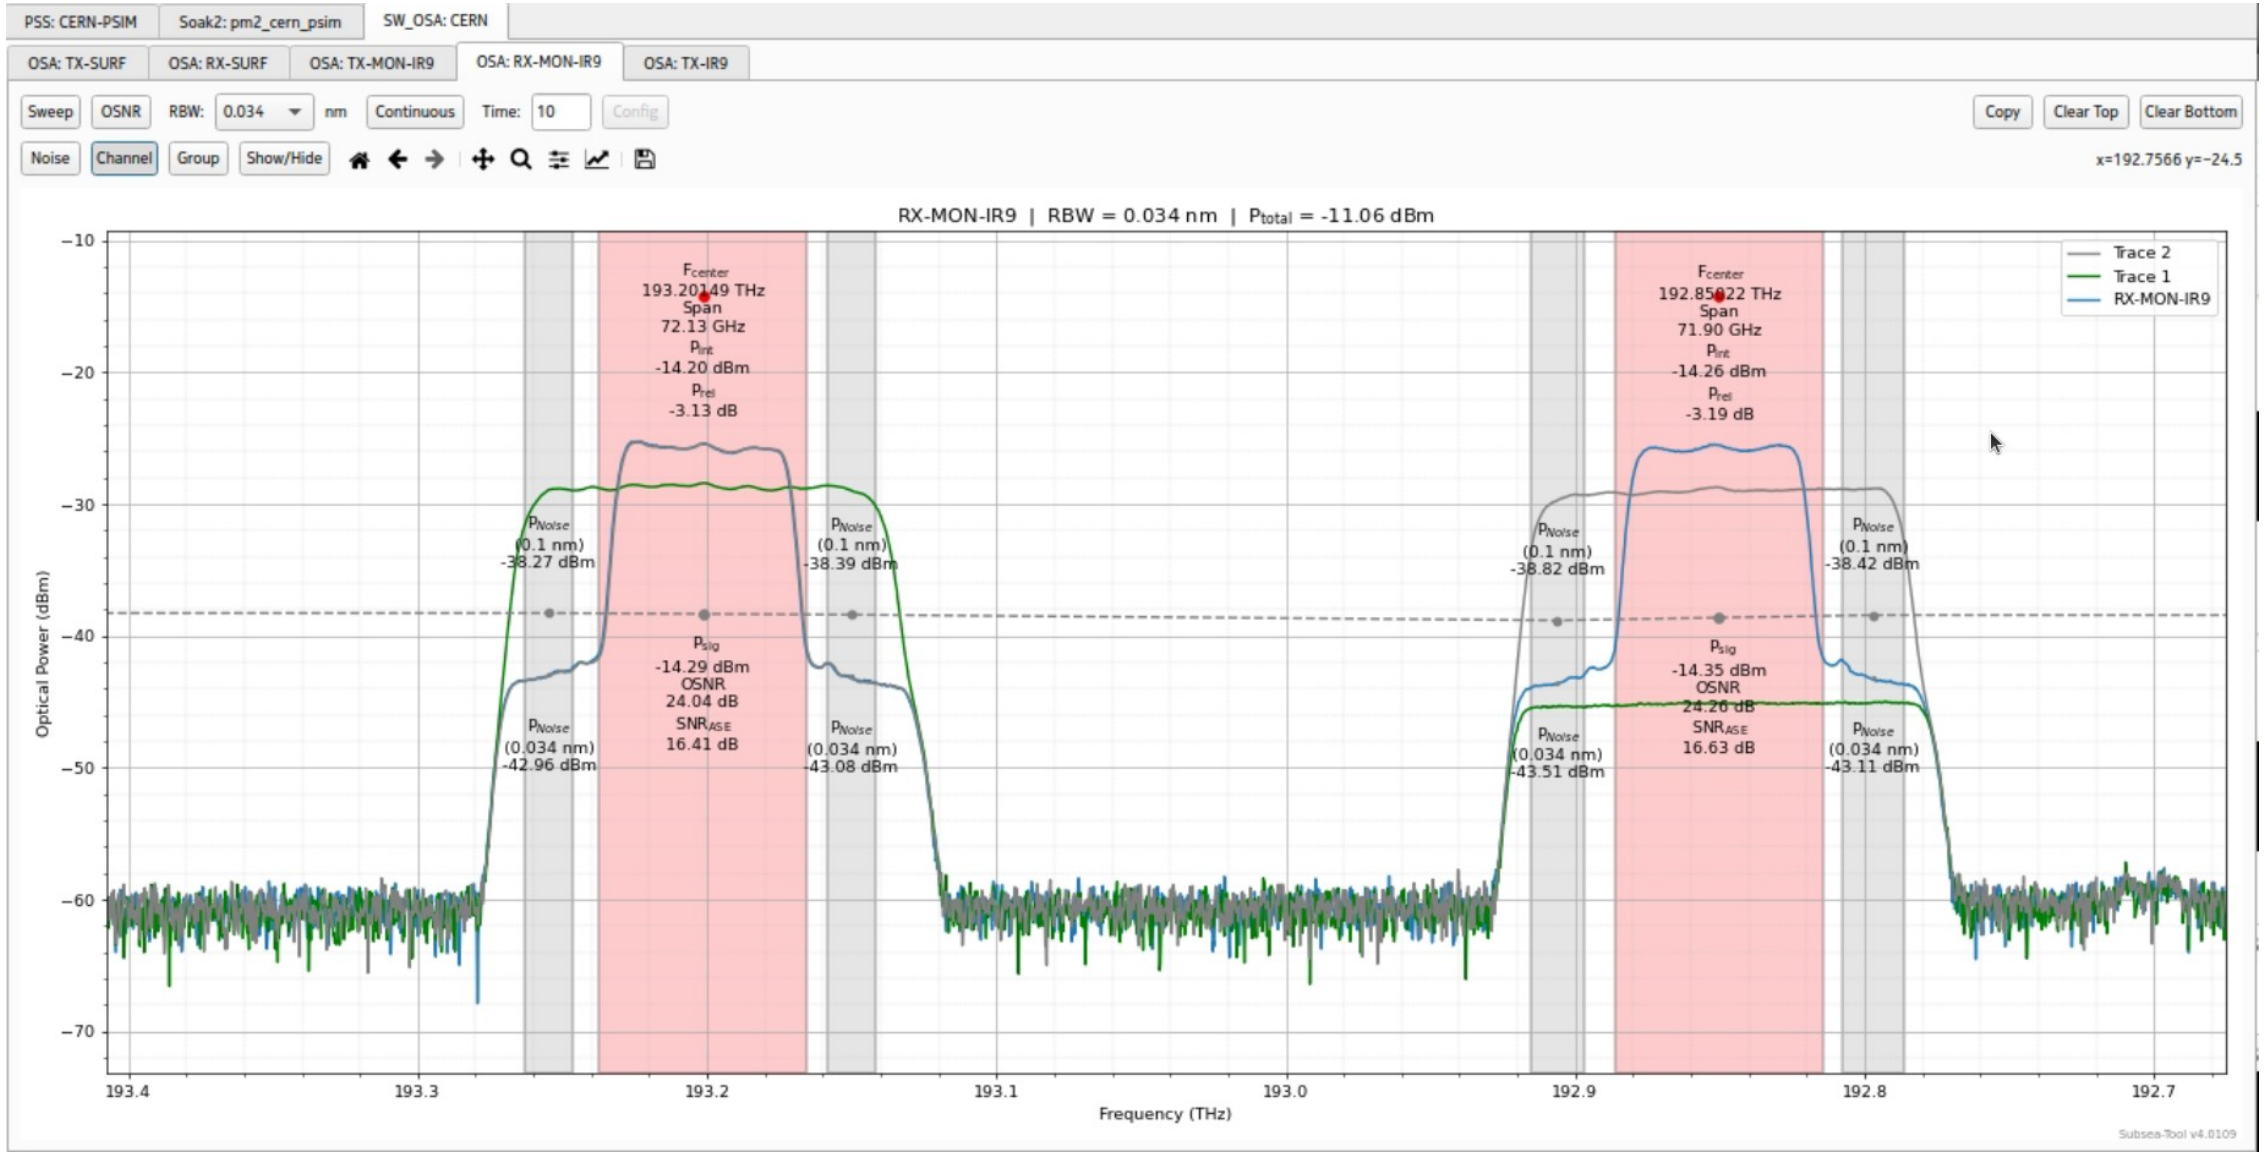

Q Margin AMSTERDAM

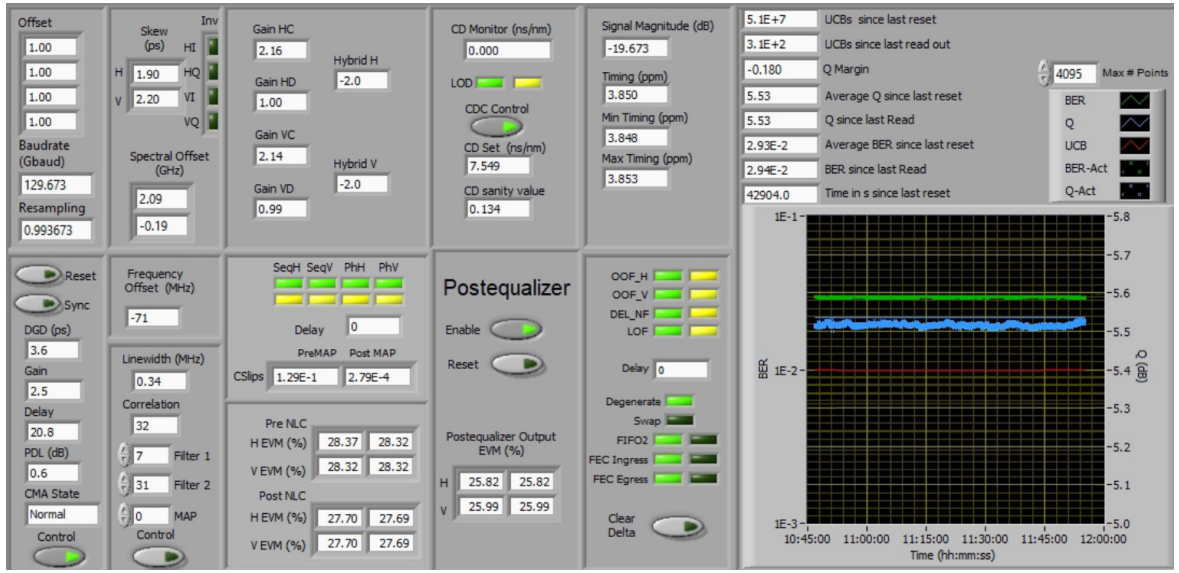

800G

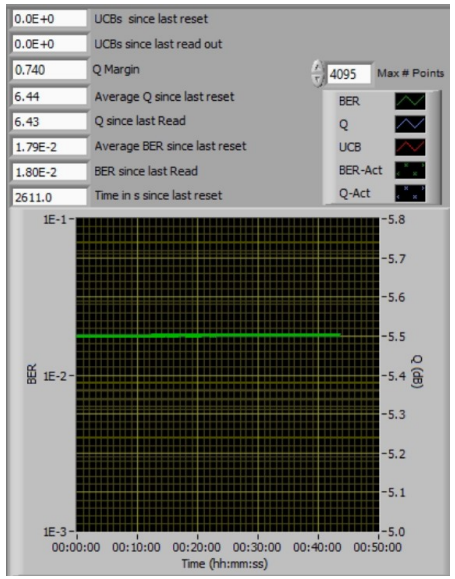

700G

Q Margin CERN

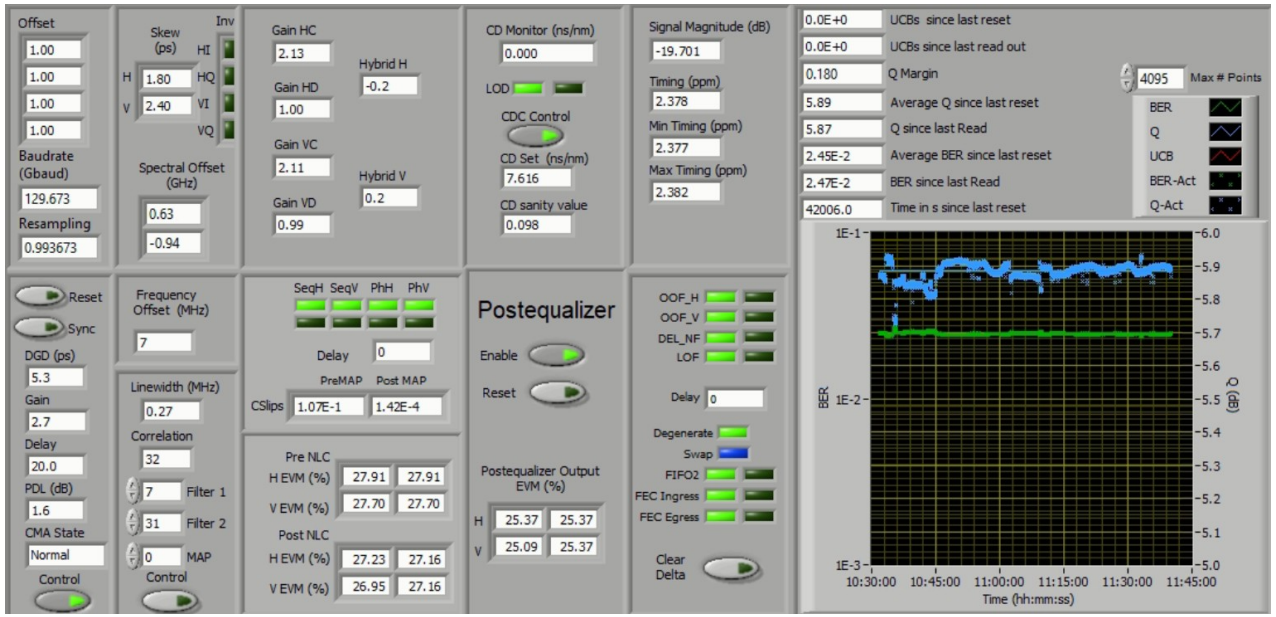

800G

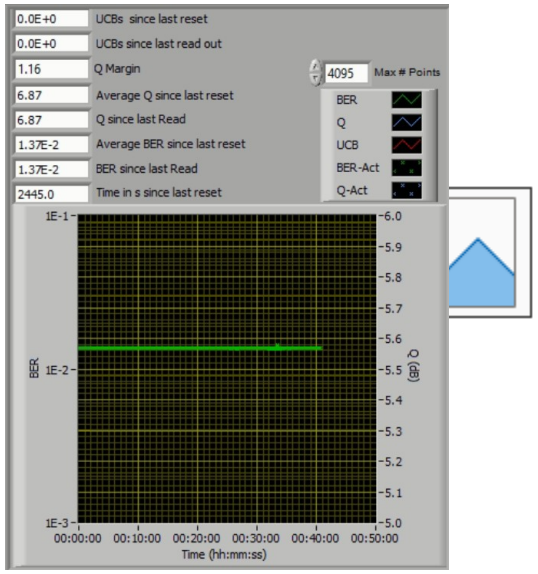

700G

Lessons learned

- Fiber AMS-GEN is an older type. Yet could reach 800G, but further optimization is required.
- Different vendors on the line system side and transponder side seem to work quite well.
- IP hardware used in trial worked as expected. (no 200G Ethernet available yet)
- L3VPN interop between Nokia and Juniper worked great
- Packet canons can easily reach 800G with 8x100G
- We are the fastest LHC T1 both on network and storage throughput!! Reached up to 661Gbit/s
- [661 was achieved with 2x 400G channels using DMAT-6 and SFM-6 to reach 800Gb/s Connectivity to NL-T1 storage]

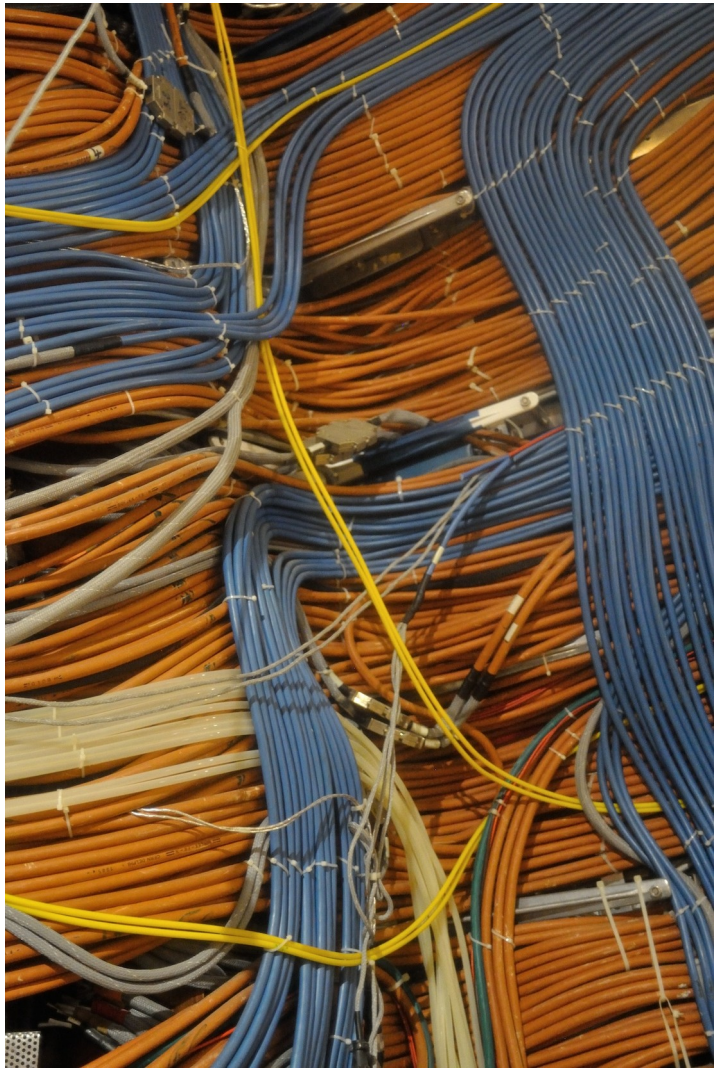

800G Trials Ribbon

Press Release

“successfully achieved 800G over a brownfield 1,650 km fiber optic link connecting research institutes including Nikhef with The Large Hadron Collider located on the CERN campus in Geneva.

The trial demonstrated a number of Ribbon’s advanced transport solutions:

- Apollo TM800_2, which uses industry-leading 5nm-140Gbaud transmission technology to deliver capacity-reach optimized 800G transport.

- Apollo Open Optical Line Systems, which include hybrid EDFA-Raman amplifiers that maximize the capacity of SURF’s brownfield G655 and G652 fiber, and have a proven ability to carry third-party vendor wavelengths.

- NPT 2400 metro router, which is interoperable with SURF’s network and delivered 2x400GbE uplinks running EVPN services on top of BGP to 8x100G ports on that network.”

- https://ribboncommunications.com/company/media-center/press-releases/ribbon-and-surf-achieve-interoperable-2x400gbe-transport-single-800g-

Cross-Border Fibers

- **Continental upgrades**
- Slides by Migiel de Vos

International Connectivity via Consortiums

SURF is closely involved with

- Advanced North Atlantic (ANA)
    - •100G service Amsterdam-Montreal (together with partners)
- Asia Europe Ring (AER)
    - •100G service London-Singapore (together with partners)

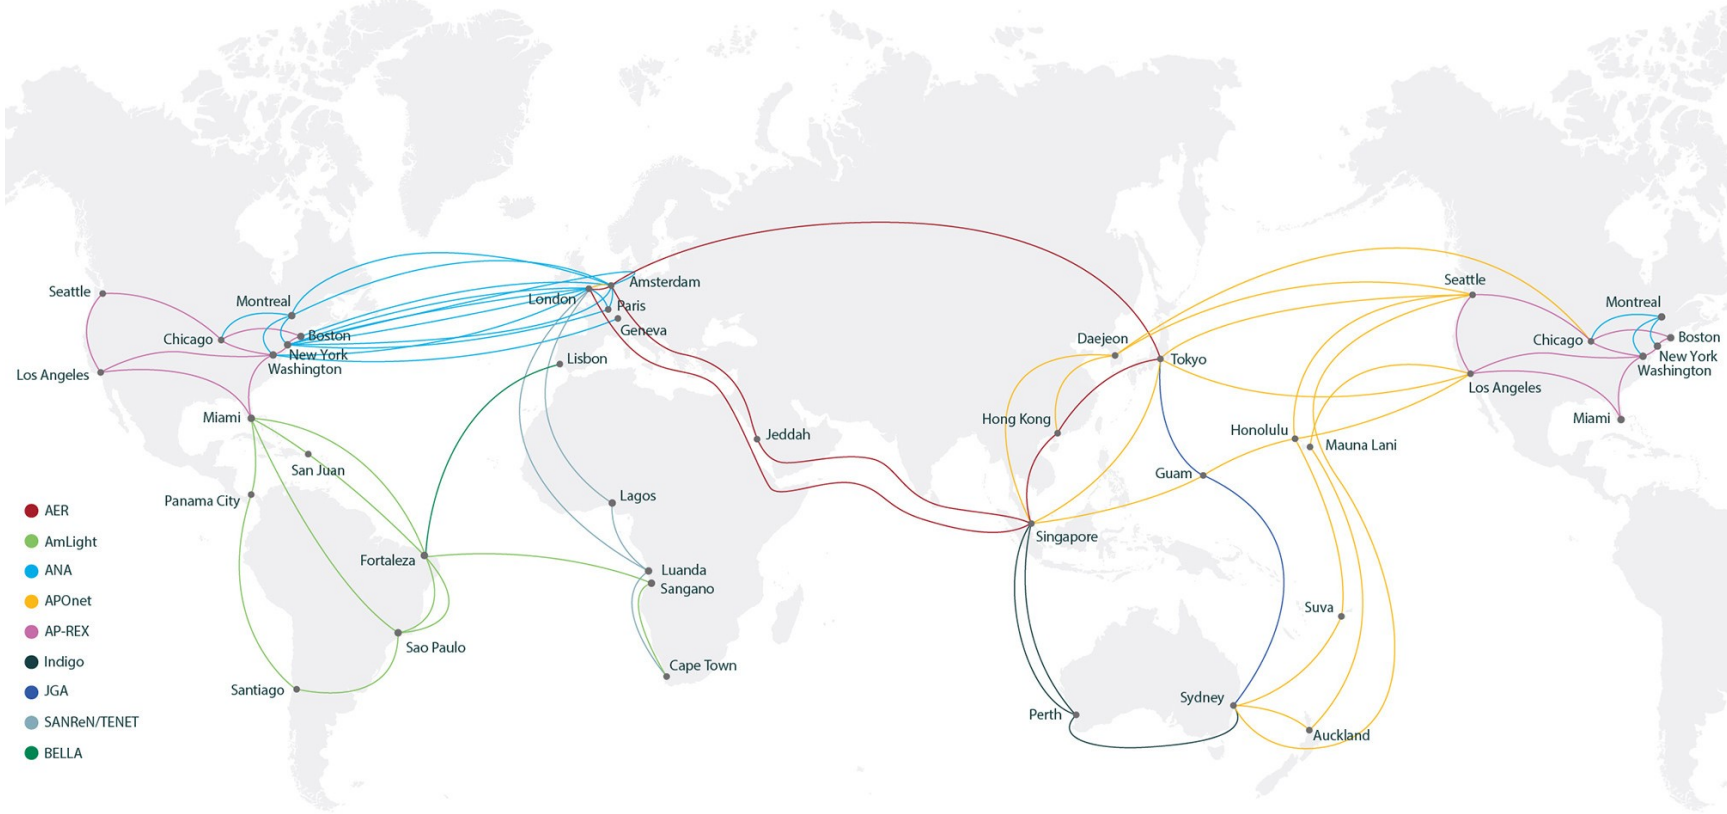

Cross Border Fiber system

Cross Border Fiber system (CBF)

- Production traffic from/to CERN
- Well-connected to the four Open Exchange Points in Europe
    - •NetherLight @ Amsterdam
    - •GÉANT Open @ Londen
    - •GÉANT Open @ Parijs
    - •CERNLight @ Genève
- Sharing capacity on the CBF with other partners that connect on the Open Exchange Points
- Perform cutting-edge trials (800G, 1+ Tbps)

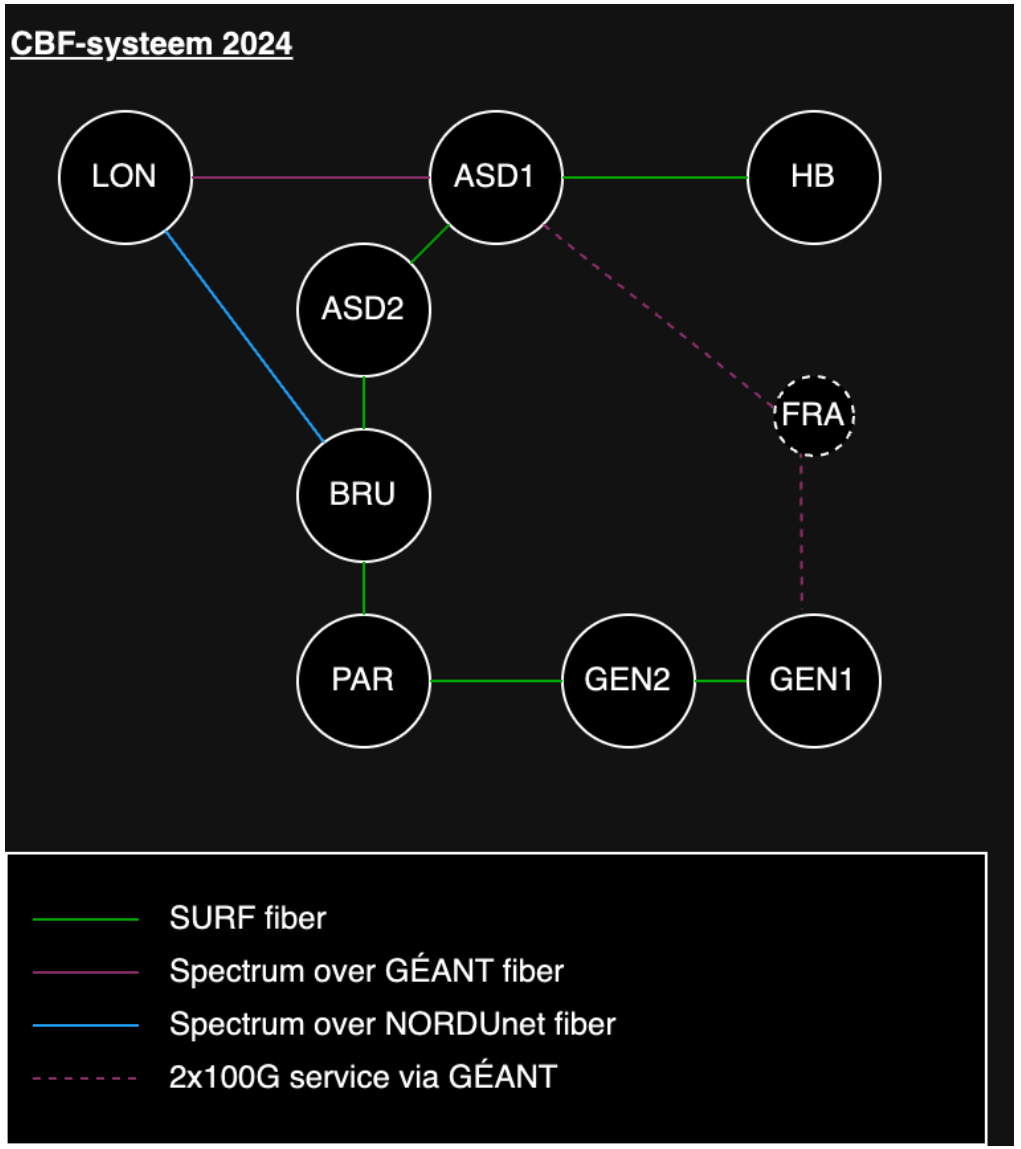

Cross Border Fiber system plans

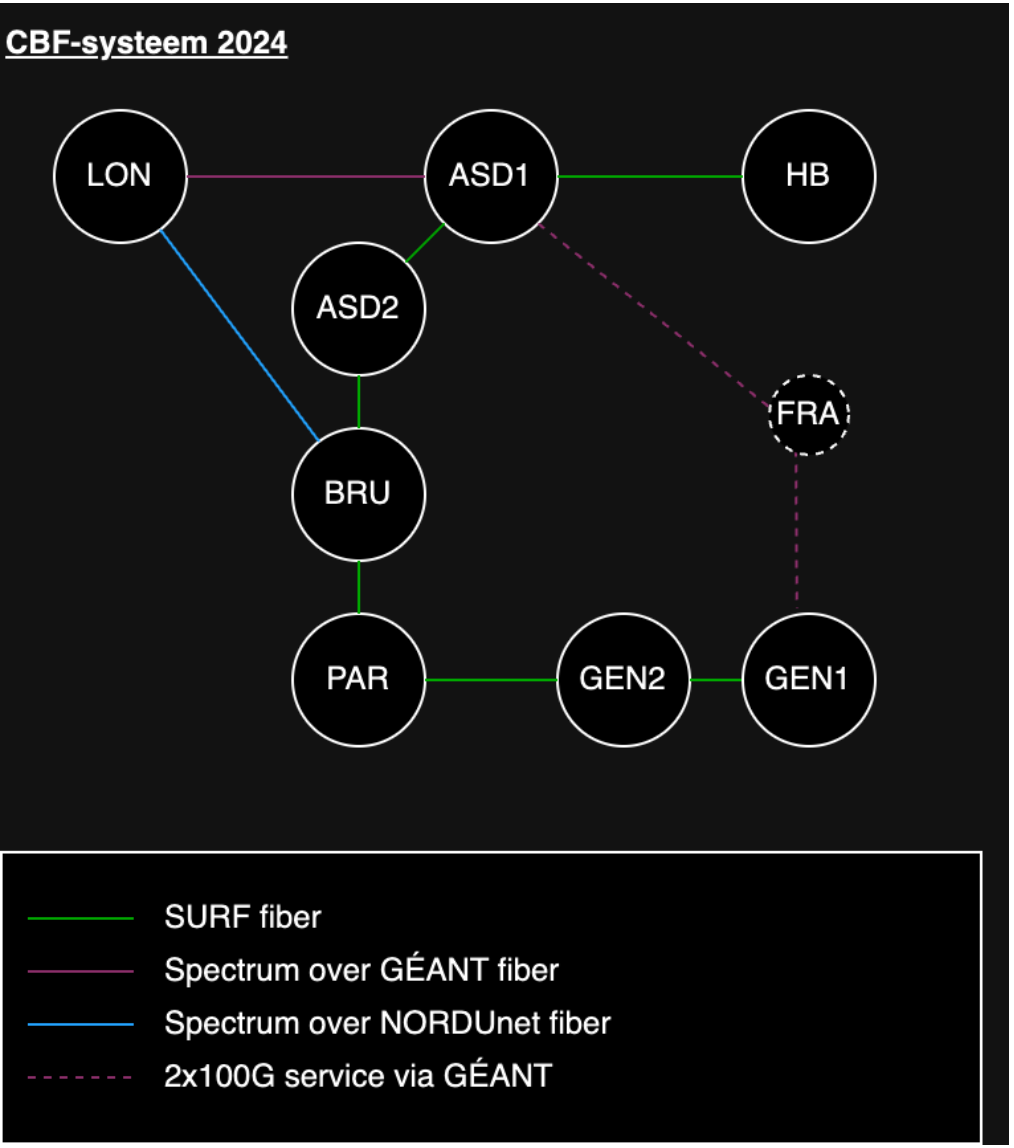

2024:

- 2x100G corelink between ASD1-GEN1 via FRA-route
- 400G ASD1-GEN2 (via PAR) with Ribbon TM800_2

2025+ (provisional):

- New fiber and line system ASD*-BRU-PAR-GEN2
- New fiber and line system between AAC* and FRA
- Spectrum via GÉANT on FRA-GEN1
- Break-out at Groningen / Aachen

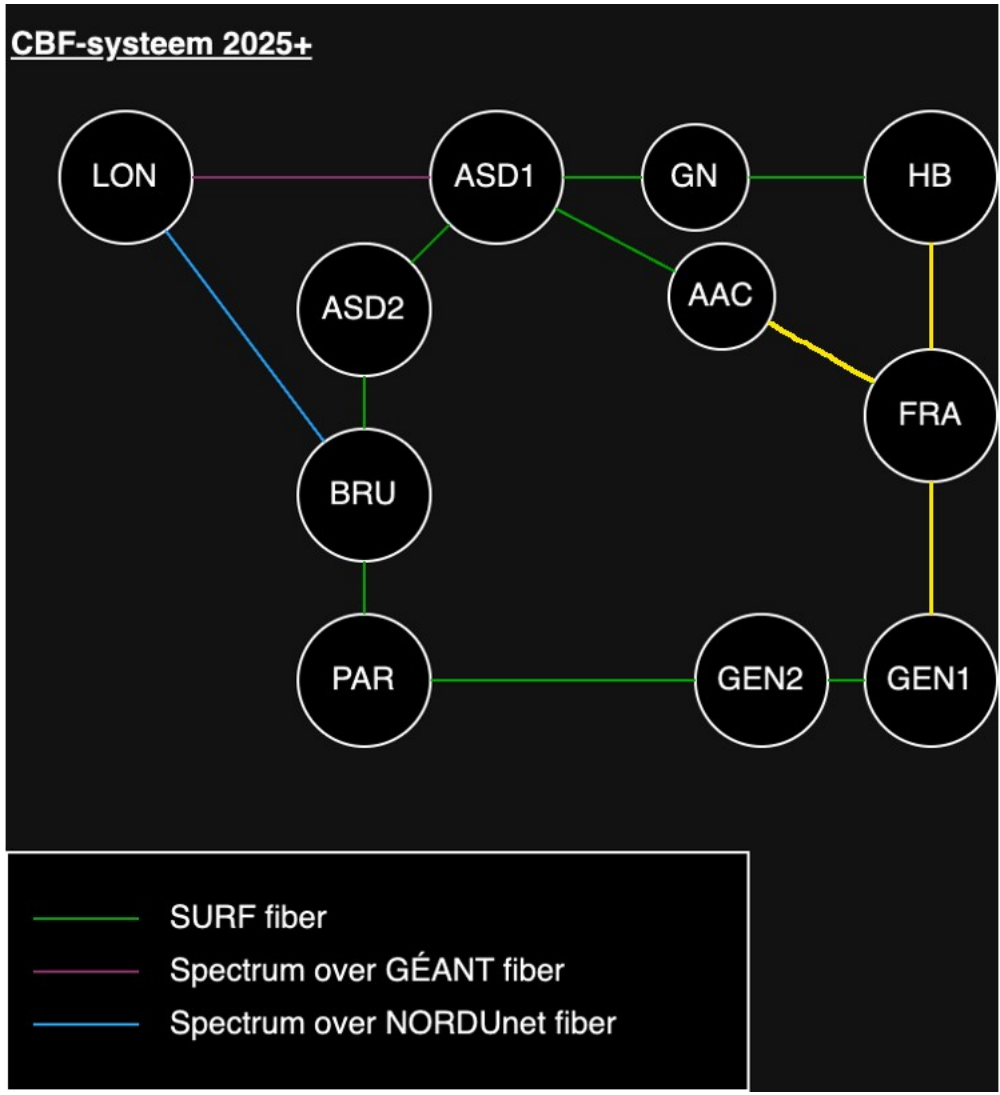

SURF TIME & FREQUENCY IN THE NETHERLANDS

Slides by Sander Klemann
Sander.Klemann@surf.nl

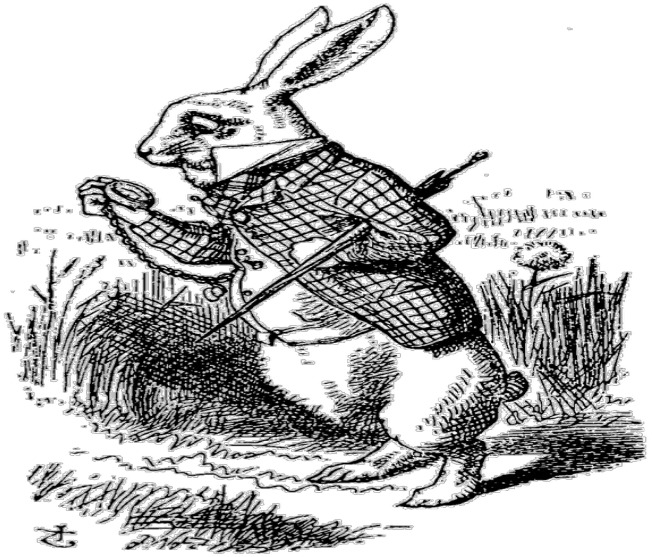

High-Level Design SURF Time & Frequency pilot

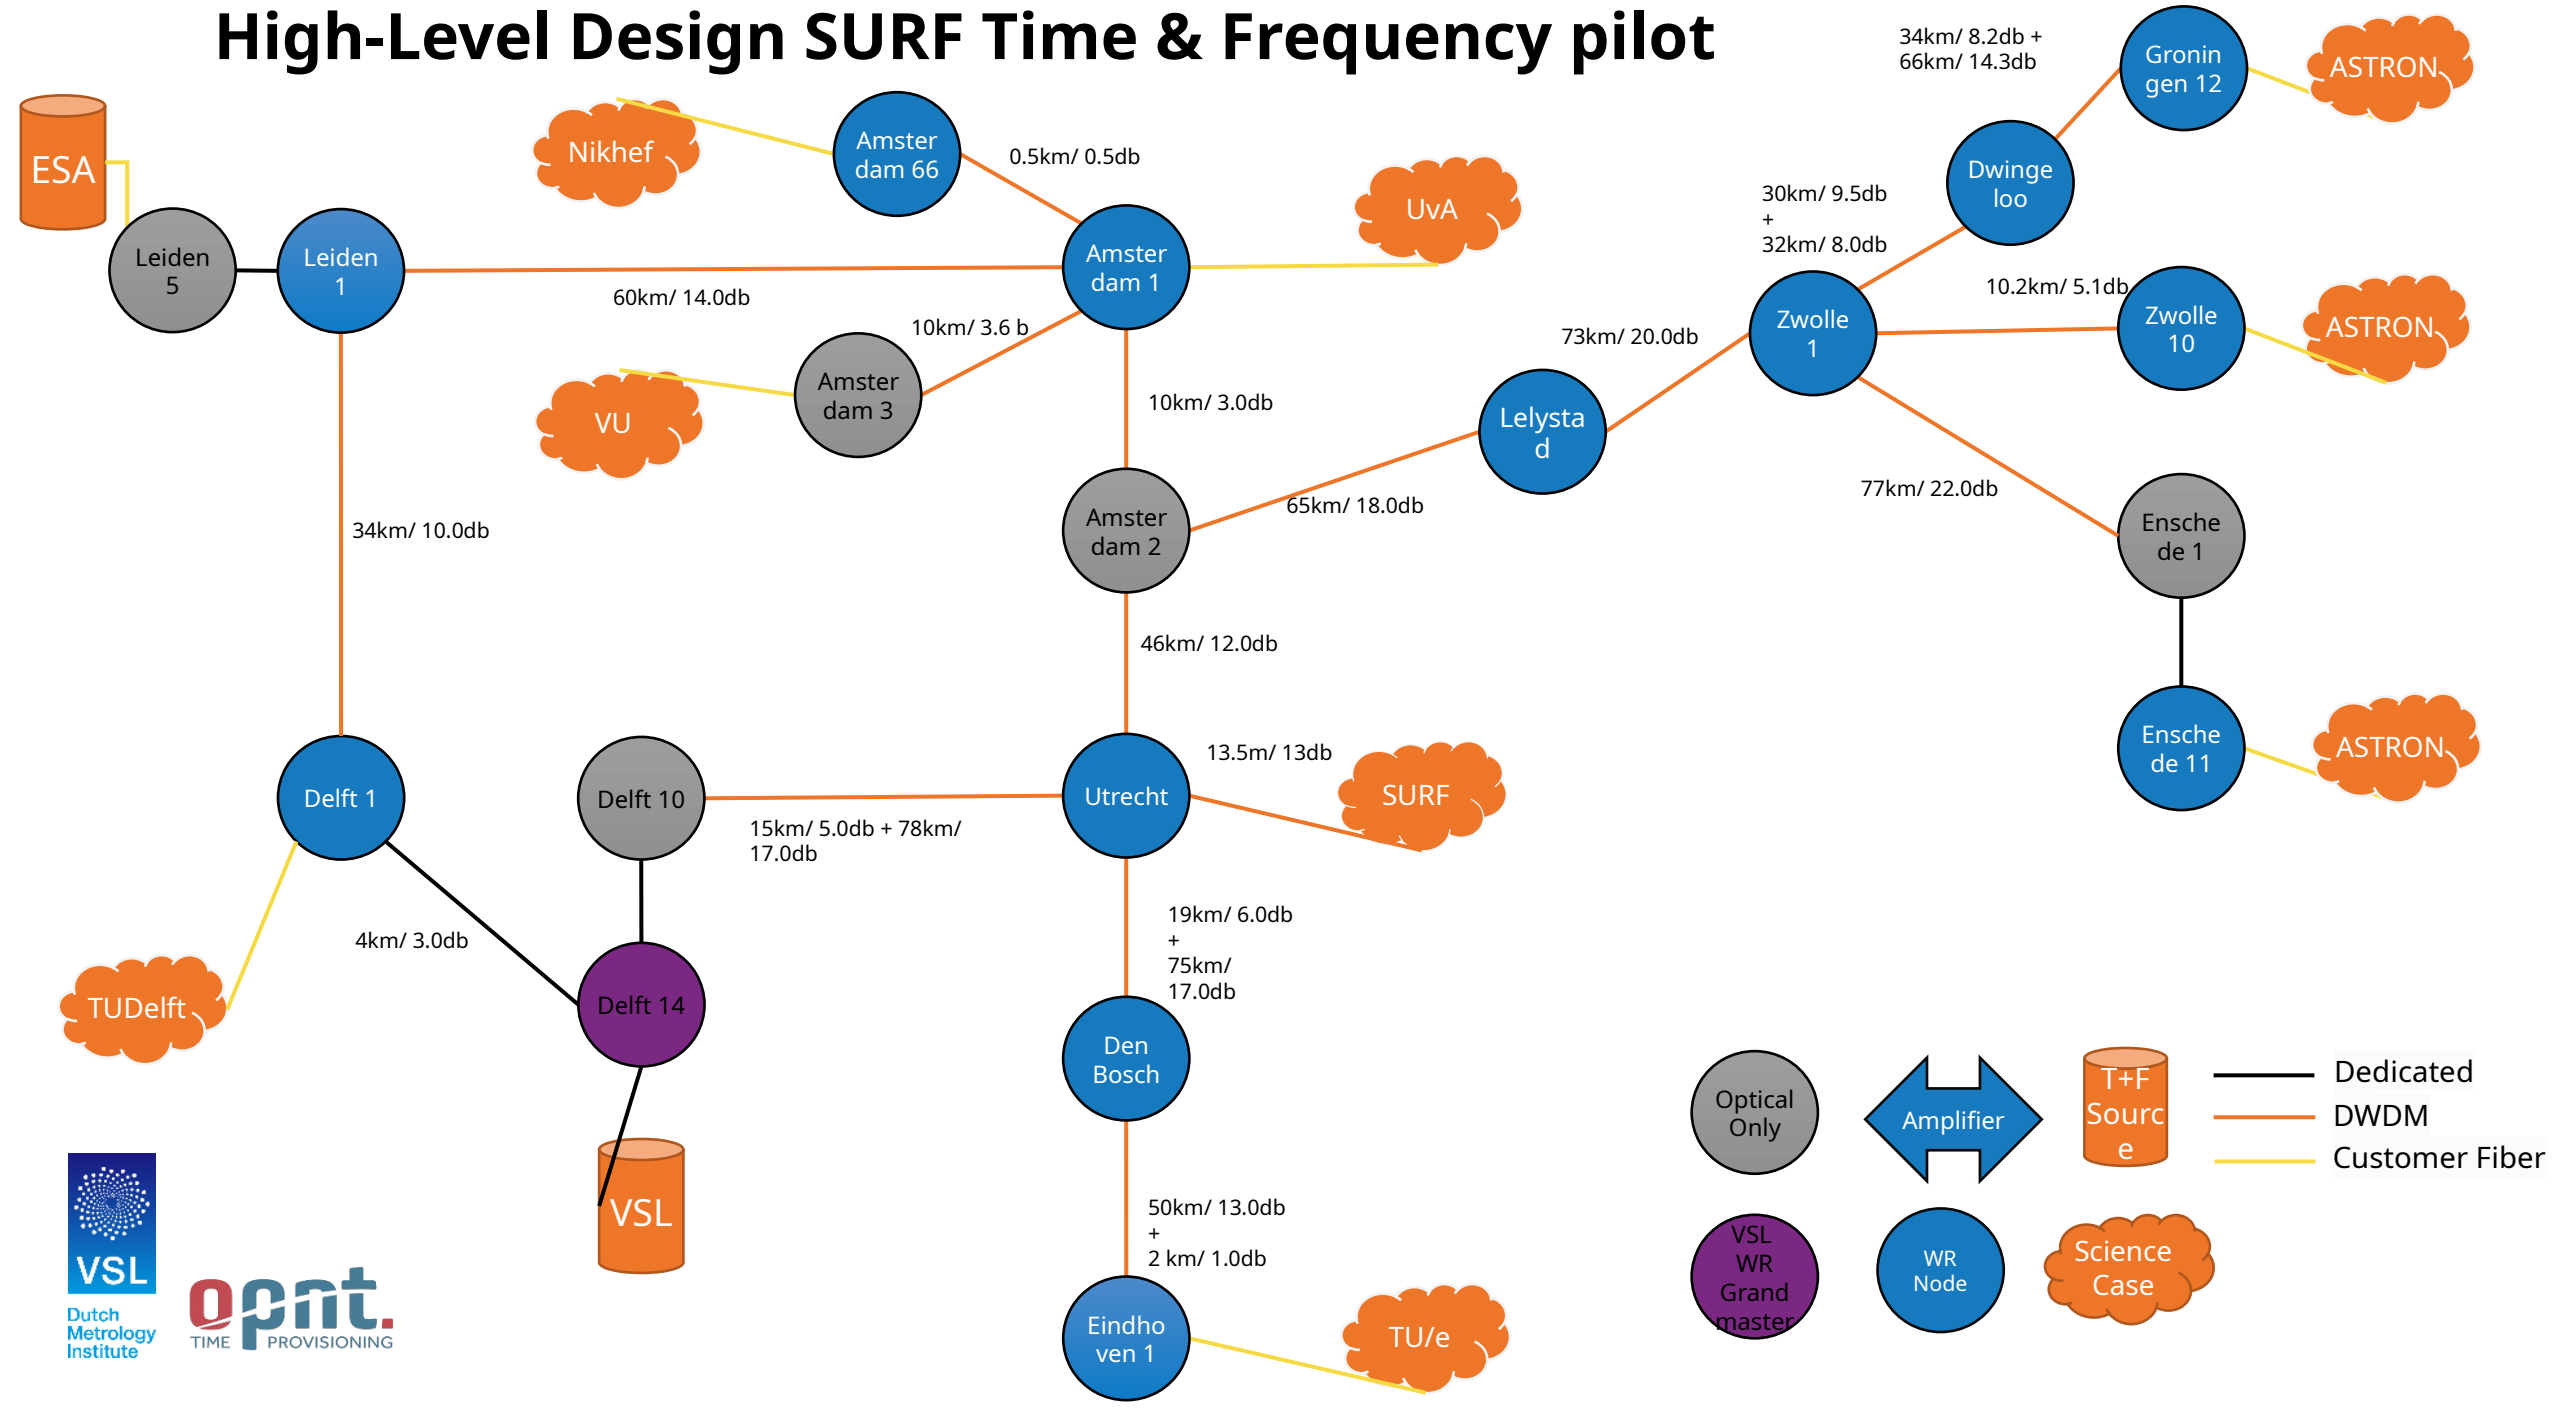

VSL Dutch Metrology Institute

opnt. TIME PROVISIONING

Current TFT Science Use Cases:

- TU/e: Quantum Computer Synchronisation
- VU: Time & Frequency Transfer Research.
- Nikhef: WR Protocol Development, TFT research
- UvA: Development of Quantum Clock, iqClock.
- ESA ESTEC: Time synchronisation and determining UTC-NL.
- TU Delft: SuperGPS -2 Research
- **ASTRON: LOFAR 2.0 Radio telescope Image Correction due to high precision clock synchronisation**
- SIDN: Time source for public NTP service Time.nl

SURF Time & Frequency
& White Rabbit in LOFAR 2.0


Onsala
Irbene
Birr
Dutch stations
Chilbolton
Norderstedt
Bałdy
Potsdam
Borówiec
Jülich
Effelsberg
Tautenburg
Łazy
Nançay
Unterweilenbach
Medicina


Onsala
Irbene
Birr
Dutch stations
Chilbolton
Norderstedt
Bałdy
Potsdam
Borówiec
Jülich
Effelsberg
Tautenburg
Łazy
Unterweilenbach
Medicina
6 arcsec


Onsala
Irbene
Birr
Dutch stations
Chilbolton
Norderstedt
Bałdy
Potsdam
Borówiec
Jülich
Effelsberg
Tautenburg
Łazy
çay
Unterweilenbach
Medicina
6 arcsec
Credit: L.K. Morabito; LOFAR Surveys KSP

Precision clock for International Stations

Calibrated image without correction for clock/ionosphere

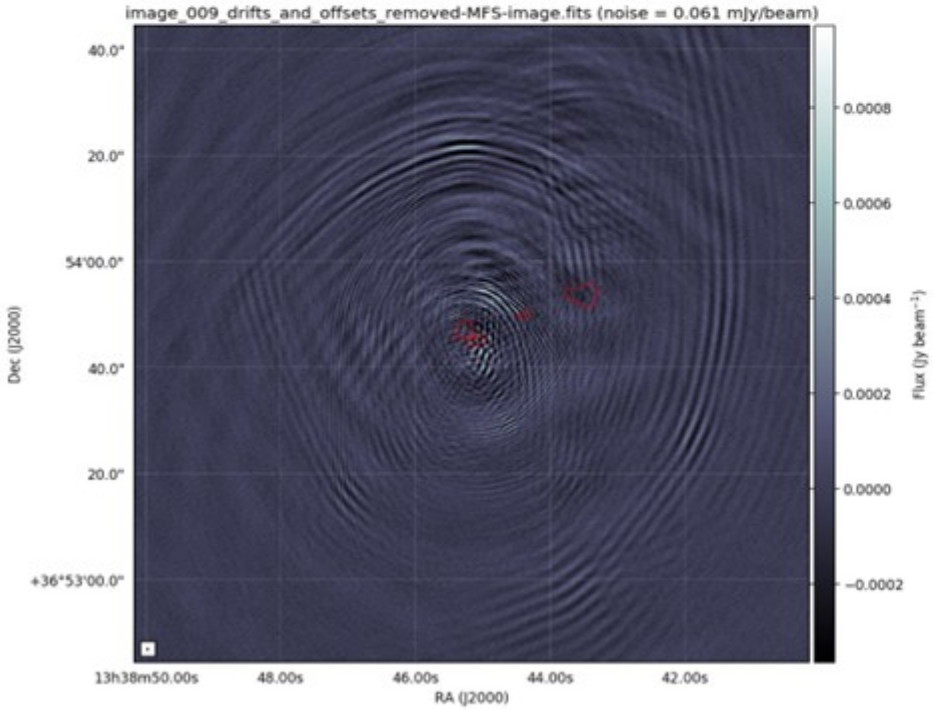

Calibrated image including correction for the clock offset and clock drifts

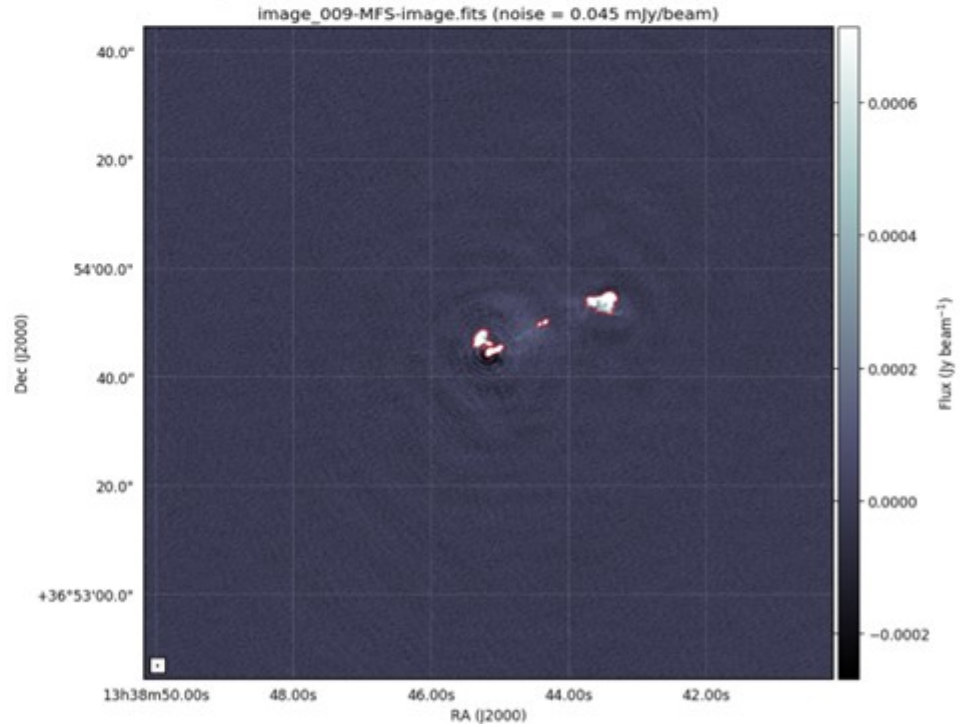

“Super-GPS” An urban navigation system with 10-centimeter accuracy

Colaboration between TUdelft, VU University, VSL and SURF

Use Mobile Telecommunication Network to distribute Atomic Clock time

More precise time is better accuracy

Publication in Nature in 2022

Koelemeij, J.C.J., Dun, H., Diouf, C.E.V. *et al.* A hybrid optical–wireless network for decimetre-level terrestrial positioning. *Nature* **611**, 473–478 (2022)

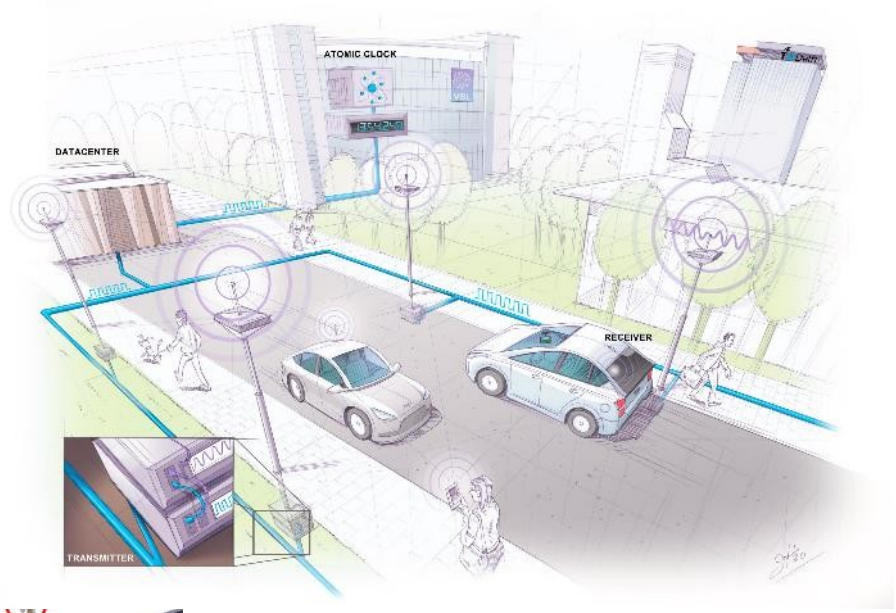

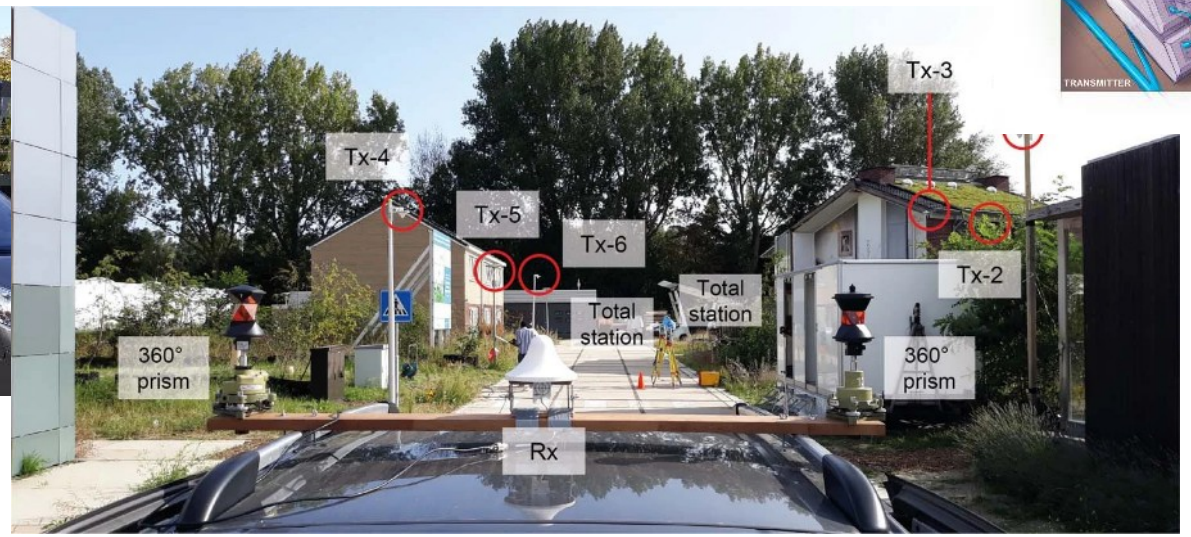

Example Use Cases in The Netherlands

TU Delft: Extra High Voltage Power Cable Research TU Delft
Real Time Digital Simulator laboratory (RTDS) Lab within a Faraday Cage, No access to GPS, needed for precise on off switching.

Eindhoven University of Technology: Hybrid Quantum Computer
Synchronisation of support devices.

University of Amsterdam: Stable Frequency for building next generation Optical Clocks
iqClock - Integrated Quantum Clock


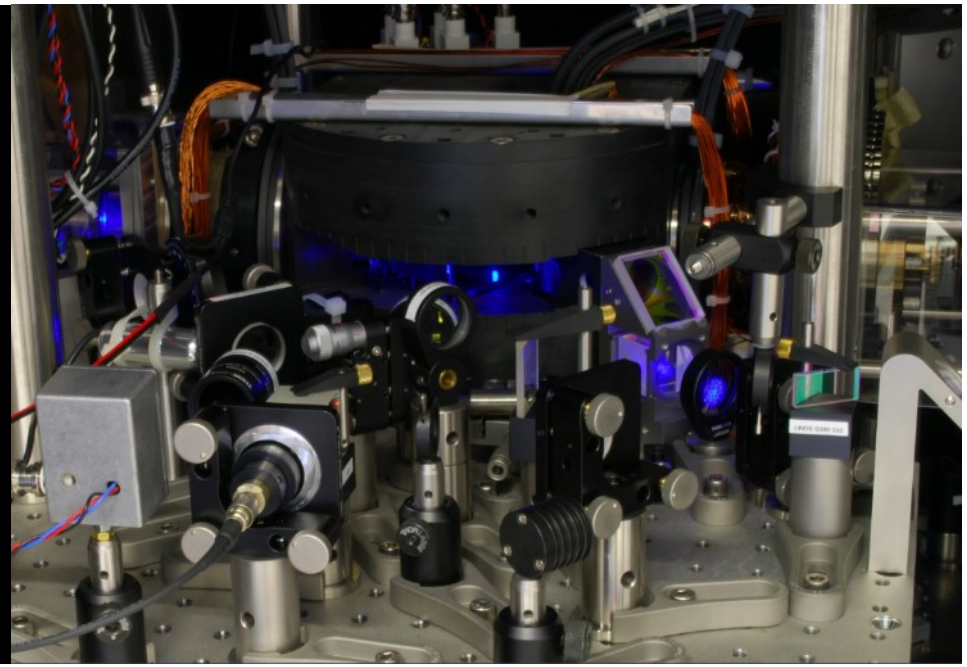

PerfSONAR Measurement Platform @ SURF

- Two nodes available:
- ps1.netherlight.net
- ps2.netherlight.net
- 100G
- Available for experimentation
- (currently reprovisioning ;o)

Questions?

Technical

Arno Bakker

Email: arno.bakker@surf.nl

Mobile: +31 6 18 24 23 97

Cooperation / Connecting / Spectrum sharing

Karin Wessel

Email: karin.wessel@surf.nl

Mobile: +31 6 39 66 94 14

LinkedIn