

Nik|hef




Maastricht University


David Groep

Samenwerken in onderzoek
een Nederlandse e-Infra aanpak

SURF Utrecht
11 juli 2022



An example - computing & data collaboration for the LHC

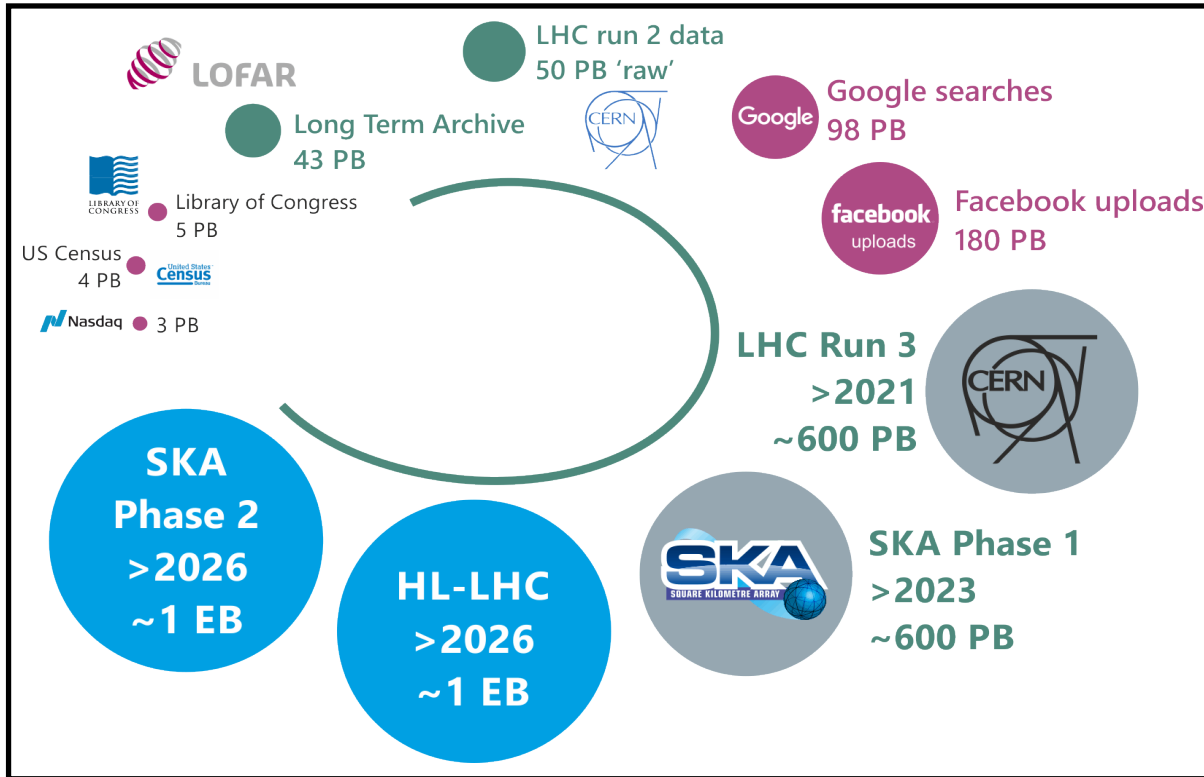


~ 1.4 million CPU cores
~ 1500 Petabyte
disk + archival


170+ institutes
40+ countries
13 'Tier-1 sites'
1 **NL-T1:**
SURF & Nikhef

e-Infrastructures
EGI
PRACE-RI
EuroHPC
OpenScienceGrid
XSEDE (ACCESS)

Data in the Dutch National e-Infrastructure



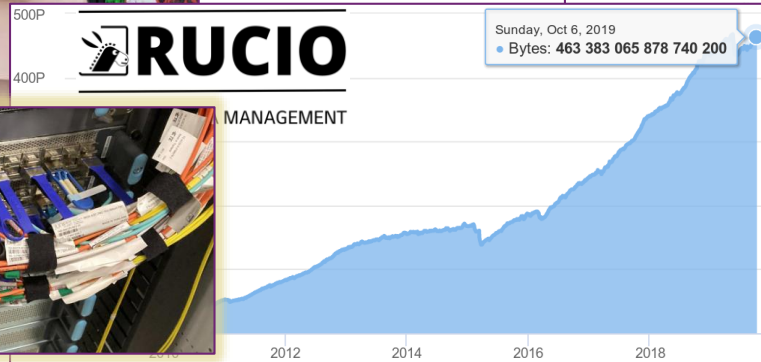
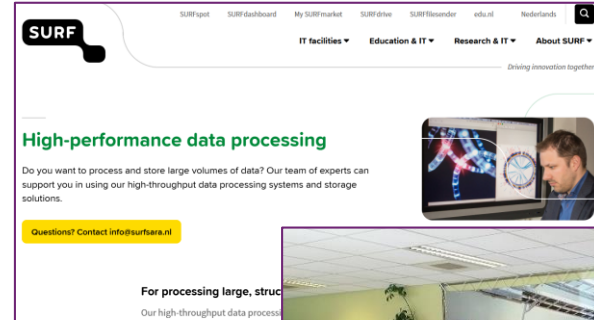
Fundamental
Sciences
E-infrastructure



approximate volumes for LoC, Google searches and FB uploads from ~ 2018. Source: WLCG publicity presentations
Imagery: SKA mid (South Africa) courtesy SKAO. Silicon tracker composite photo Nikhef, KM3Net DOM deployment module: Nikhef and NIOZ

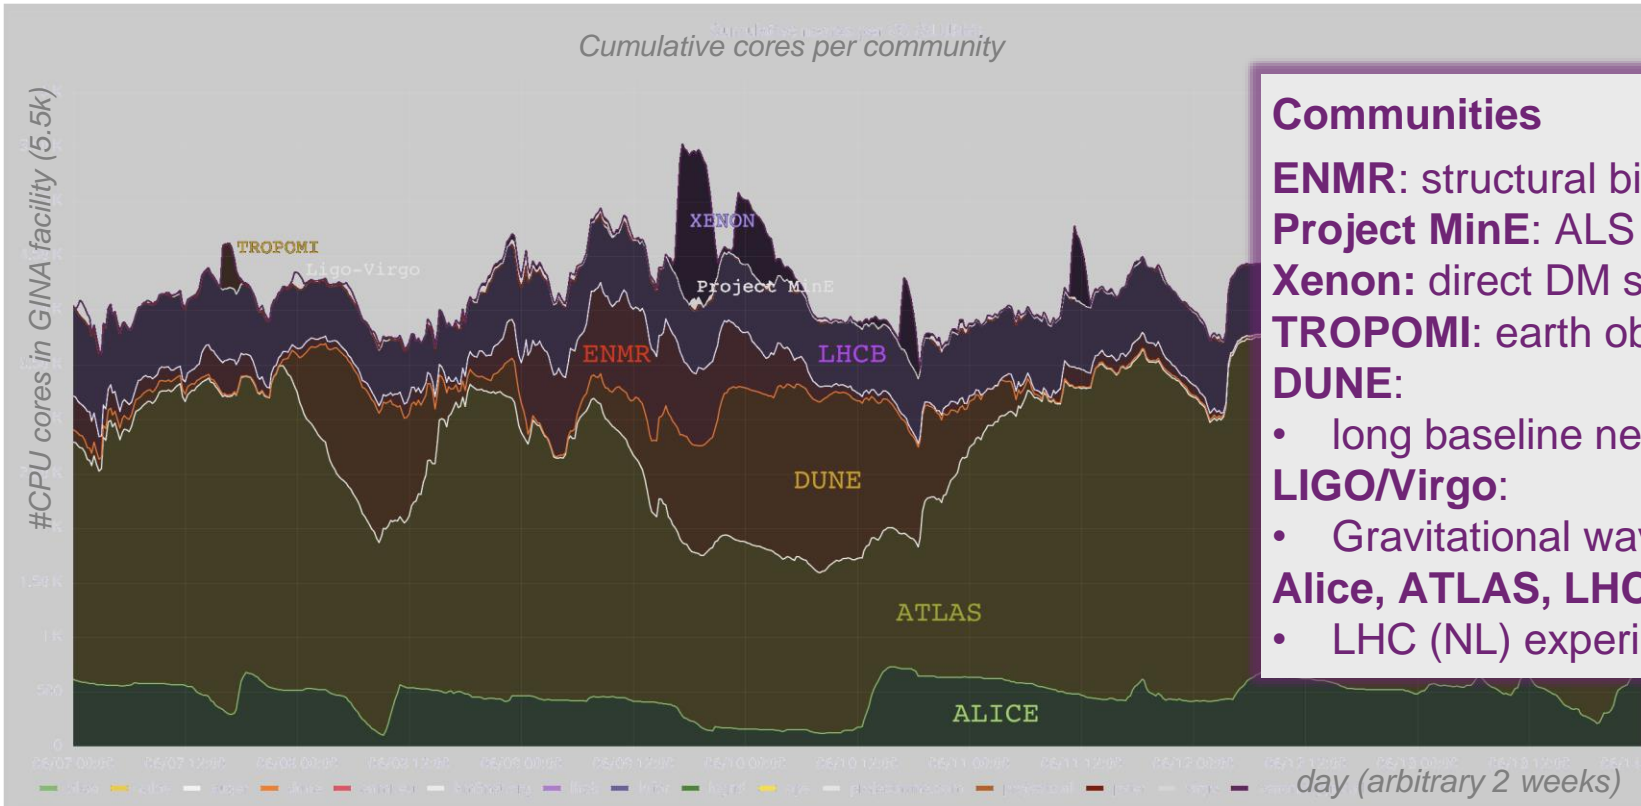


Infrastructure for research is an ecosystem: hardware, software, services, and ... people



Images: ATLAS Rucio volume, (from rucio.cern.ch); optical network: NDPF 'deel'; User meeting Stoomboot Office Hours (both Nikhef); Snellius opening visit; HPDC service page (both SURF)

Collaboration in the data processing coordinated by SURF



Communities

- ENMR:** structural biochemistry
- Project MinE:** ALS (health)
- Xenon:** direct DM searches
- TROPOMI:** earth observation
- DUNE:**
 - long baseline neutrinos
- LIGO/Virgo:**
 - Gravitational waves
- Alice, ATLAS, LHCb**
 - LHC (NL) experiments



Scalable HPC strategy: from local “T2” to European “T0”



Nikhef “Stoomboot”
Analysis Facility



...



SURF National Infrastructure
solid foundation *and* essential stepping stone



How to exploit
these unique systems?
*access, expertise, and ...
a long-term vision
on how research scales up*

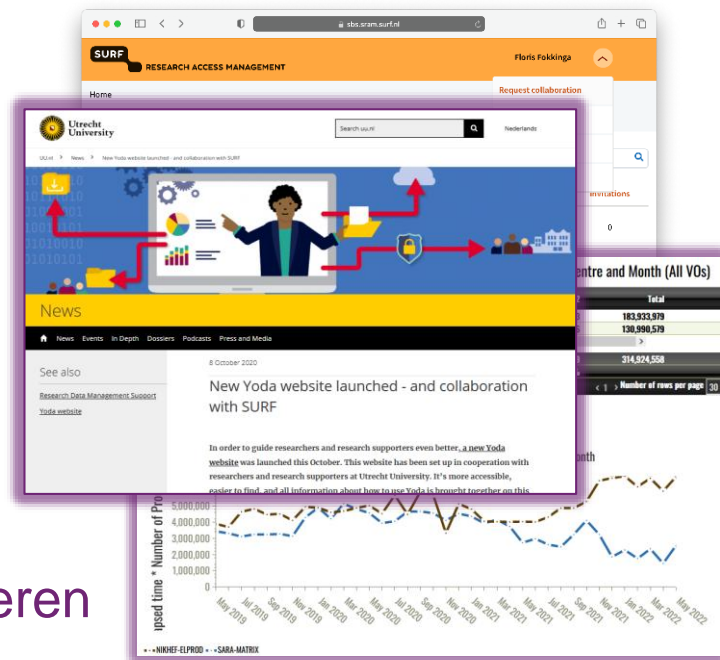
SURF als 'linking pin' tussen instellingen – en de wereld

Gezamenlijk optrekken werkt

- 'DNI' federative computing service
- Data repositories DANS & 4TU.RD
- HPC 'Tier-2' network (SURF & 8 instellingen)
- data management - zoals 'yoda' met UU
- 'digital competences': data levend houden

Maar dit werkt alleen als ook *instellingen* infrastructuur als 'asset' zien, en durven investeren

en dan is dit een effectieve oplossing die misschien eens *niet* aan vendor lock-in ten prooit valt zoals public cloud ... of zoals PURE & Mendelay, of Figshare (van die andere uitgever)



Sharing more than resources: from data to software & expertise

Beyond the 'F & A' of FAIR – *infrastructure for reproducible research with re-usable software*

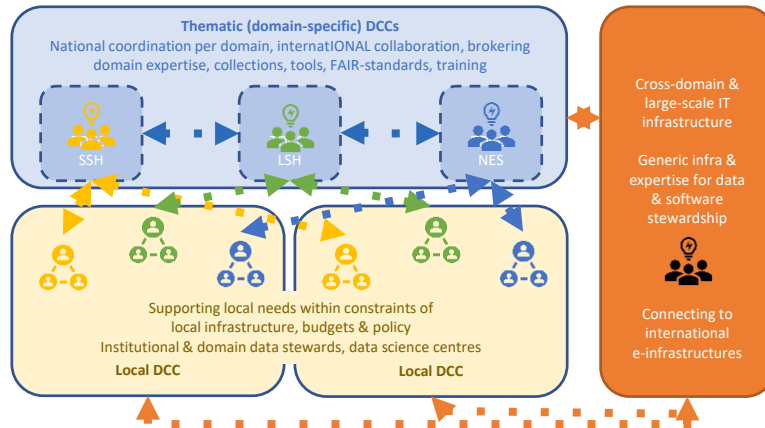


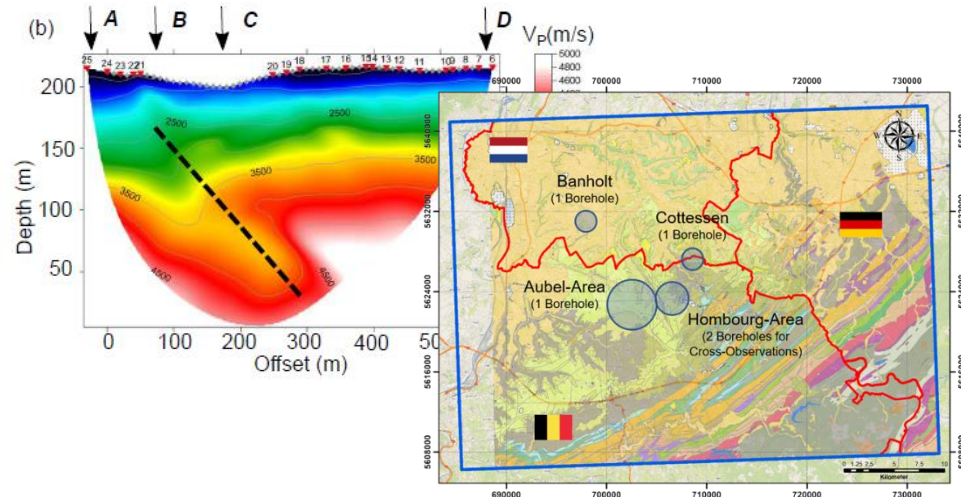


Image by: Ruben Kok, LSH TDCC and DTL

Software and infrastructure essential to bring 'dead' data to life!



Sharing expertise across domains – the ‘NES’ example

Case study: Einstein Telescope seismic studies in EUregio Meuse-Rhine in *E-TEST*



Data collected here is also useful for many others - outside of the ET planning ... that span many ‘local’ organisations (UM/Nikhef, KNMI, Liege, Aachen,...)


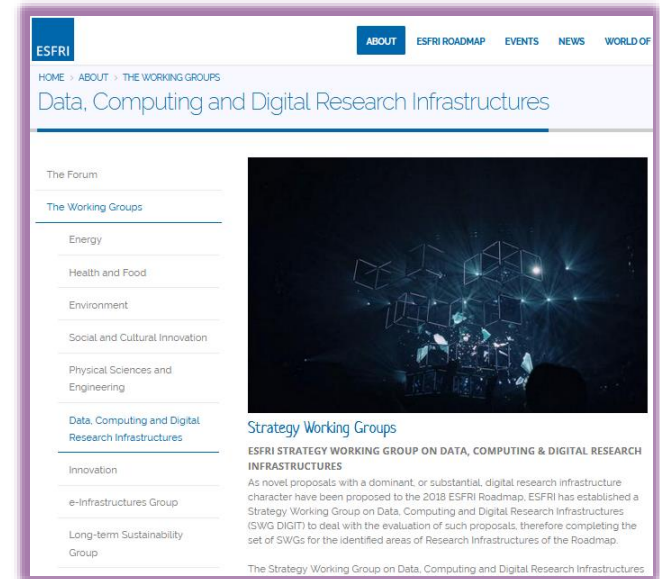
ET impression: Marco Kraan (Nikhef) from “Terziet drilling campaign” <https://www.nikhef.nl/wp-content/uploads/2019/10/Terziet-Drilling-Campaign-Final-NoC.pdf>
Seismic data: S Koley (VU and Nikhef) *Sensor networks to measure environmental noise at gravitational wave detector sites*, ISBN 978-94028-2054-6; map image: etest-emr.eu project site

ICT als onderzoeks *instrument*

'ICT infrastructuur voor onderzoek is geen office IT'

- research data is 'born digital' – research domein clusters steeds beter in bepalen volume en inschatten behoeftes
- 'digitale paragraaf' in ESFRI RIs is strategische resource
- open science, open data? groeiend beslag infrastructuur!
- research software idem: maintenance 'eet' in innovatie

Computing vraag blijft: open science is geen project
en 'de instellingen' gaan dit nu nooit alleen redden!



Infrastructure research for research infrastructures

Improving infrastructure



SURF DNI, Dutch data centres,
joint tenders, systems innovation

Accessing infrastructure



Workflow engines, access, Rucio
data lakes, eduGAIN, SRAM, ...

Improving use & efficiency

```
// Allocate on device memory
Matrix d_C;
d_C.width = d_C.stride = C.width; d_C.height = C.height;
size = C.width * C.height * sizeof(float);
cudaMalloc(&d_C.elements, size);

// Invoke kernel
dim3 dimBlock(BLOCK_SIZE, BLOCK_SIZE);
dim3 dimGrid(B.width / dimBlock.x, A.height / dimBlock.y);
MatMulKernel<<<dimGrid, dimBlock>>>(d_A, d_B, d_C);

// Read C from device memory
```

Pipelines, machine learning,
GPU's, Quantum Computing

Structure: fuse-infra.nl, images: Nikhef (nationalespeeltuin.nl), AARC Community, LHCB

Riding the Infrastructure Innovation Chain together

Computing Sciences Research

Operational Research (near-term, NextGen storage, 800G+ network, QC simulators)

Operational innovation (procurement, systems vendor co-engineering)



Research in the Network Systems (supers@surf.nl)

Computer Systems! Our modern and, in turn, on sustained remarkable technology area

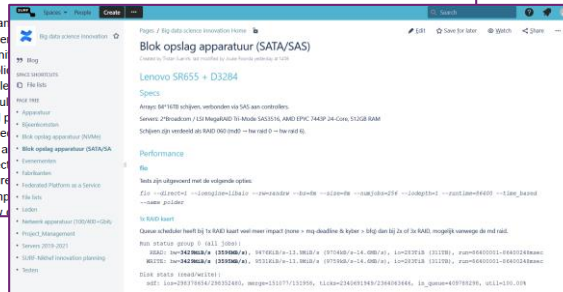
With an outstanding impact on society, this manifesto focuses on the Netherlands, where data centers and related ICT infrastructure enable over 3.3 million jobs and over 60% of the GDP, and novel services and products, and where every €1 invested in such systems generates €15 in added value. Sustained investments in capable networking infrastructure have made the Netherlands home to one of the largest data hubs in the world. Cloud adoption exceeds 90% for economic organizations, and 65% for government and public education organizations. Our overarching goal with this document is to highlight the grand societal, technological, and scientific opportunities and challenges in future computer systems and networking (the CompSys area), and to outline how to maintain the leading position the Netherlands has in this area.

Future-proof digitalization requires ICT research and development (R&D), now. The Dutch Government and societal stakeholders have identified an urgent need to expand economic and social activities by insuring and integrating ICT in their knowledge, expertise, and capabilities.

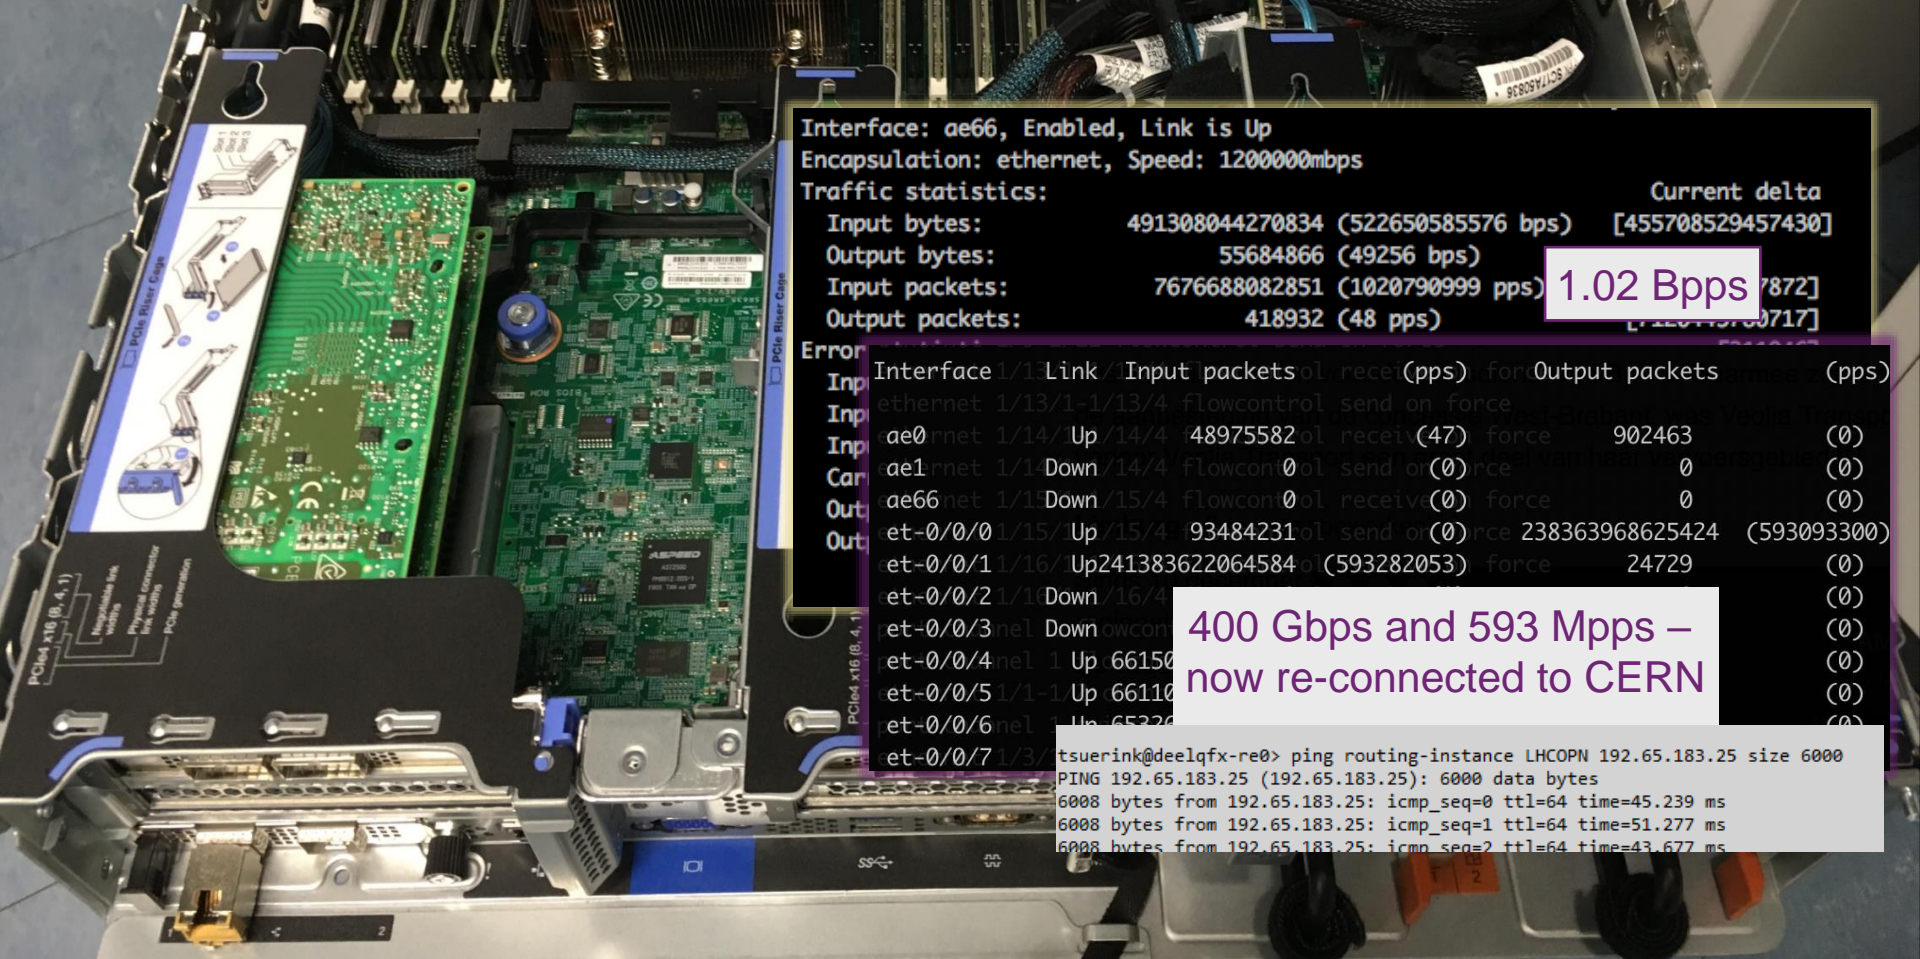
From Quark to Quantum ... with LHCb

Jacco de Vries: Maastricht University/Nikhef jdevries@nikhef.nl
 Daniel Campora: Maastricht University/CERN dcampora@cern.ch
 Kareljan Schoutens: UvA/QuSoft c.j.m.schoutens@uva.nl
 David Groep: Nikhef davidg@nikhef.nl
 Ariana Torres: SURF ariana.torres@surf.nl

Introduction



involving non-CS research domains



Interface: ae66, Enabled, Link is Up
Encapsulation: ethernet, Speed: 1200000mbps

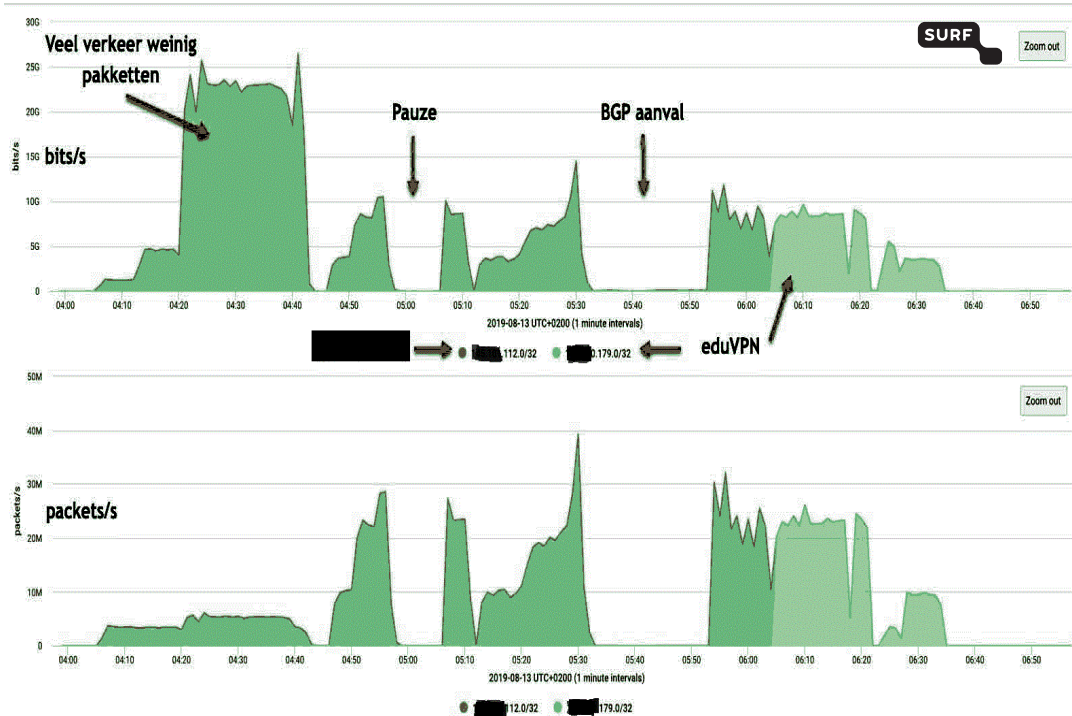
Traffic statistics: Current delta
Input bytes: 491308044270834 (522650585576 bps) [455708529457430]
Output bytes: 55684866 (49256 bps)
Input packets: 7676688082851 (1020790999 pps) **1.02 Bpps** [7872]
Output packets: 418932 (48 pps) [12079700717]

```
Error:
Interface 1/1 Link / Input packets ol receive (pps) for Output packets (pps)
Inp ethernet 1/13/1-1/13/4 Flowcontrol send on force
Inp ae0 rnet 1/14/Up /14/4 48975582 ol receive (47) force 902463 (0)
Inp ae1 rnet 1/1-Down /14/4 Flowcontrol send or (0) rce 0 (0)
Car ae66 net 1/1-Down /15/4 Flowcontrol receive (0) force 0 (0)
Out et-0/0/0 1/15/Up /15/4 93484231 ol send or (0) rce 238363968625424 (593093300)
Out et-0/0/1 1/16/Up 241383622064584 ol (593282053) force 24729 (0)
et-0/0/2 1/1-Down /16/4
et-0/0/3 nel Down owcon
et-0/0/4 nel 1 Up 66150
et-0/0/5 1/1-1-Up 66110
et-0/0/6 nel 1 Up 65320
et-0/0/7 1/3/ tsuerink@deeljfx-re0> ping routing-instance LHCOFN 192.65.183.25 size 6000
PING 192.65.183.25 (192.65.183.25): 6000 data bytes
6008 bytes from 192.65.183.25: icmp_seq=0 ttl=64 time=45.239 ms
6008 bytes from 192.65.183.25: icmp_seq=1 ttl=64 time=51.277 ms
6008 bytes from 192.65.183.25: icmp_seq=2 ttl=64 time=43.677 ms
```

400 Gbps and 593 Mpps –
now re-connected to CERN


Our science data looks akin to a DoS ... hence ...


shown: testing against SURFnet7
other targets cannot be disclosed publicly



evaluating resilience to cyberattack - in a cooperative way

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





SURF

Nikhef



David Groep

davidg@nikhef.nl

<https://www.nikhef.nl/~davidg/presentations/>

 <https://orcid.org/0000-0003-1026-6606>

this work co-funded by and contributing to the Dutch National e-Infrastructure coordinated by SURF

