Nikhef – a Journey in Physics and Data Processing

David Groep
Nikhef
PDP - Advanced Computing for Research

Nationaal instituut voor subatomaire fysica

Verleggen van de grenzen van onze kennis

- **Accelerator-based particle physics**
  Experiments studying interactions in particle collision processes at particle accelerators, in particular at CERN;
- **Astroparticle physics**
  Experiments studying interactions of particles and radiation emanating from the Universe.

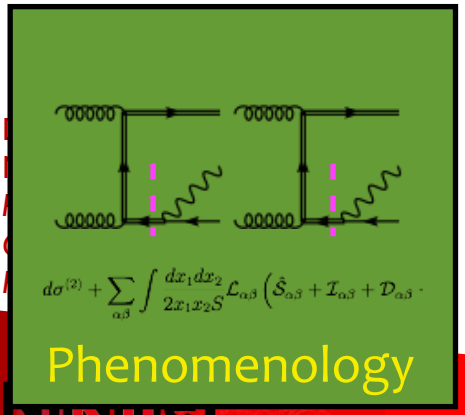

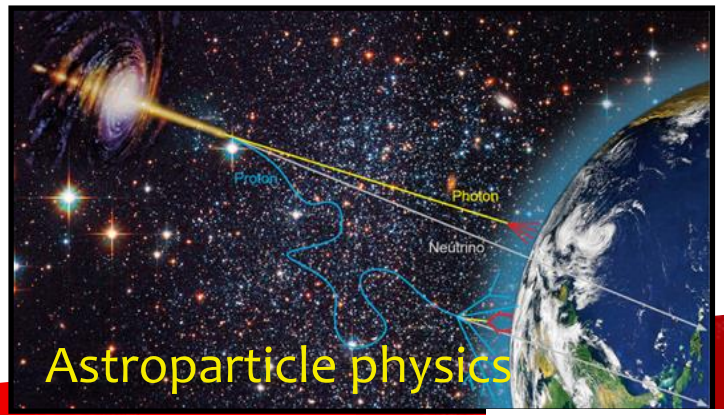

Large Hadron Collider

CMS

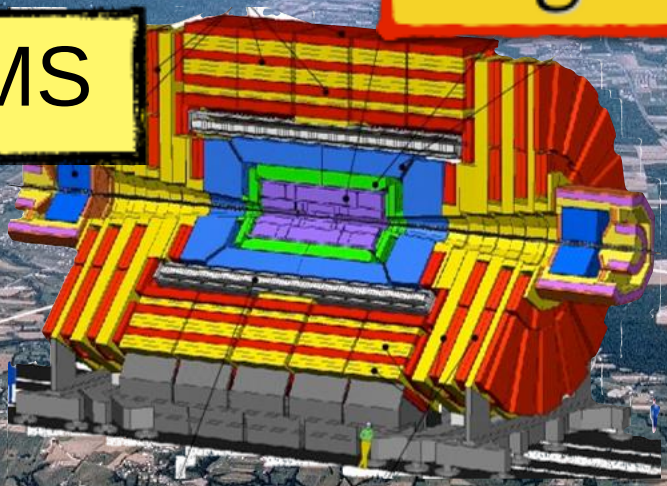

Nikhef

LHCb

ALICE

ATLAS

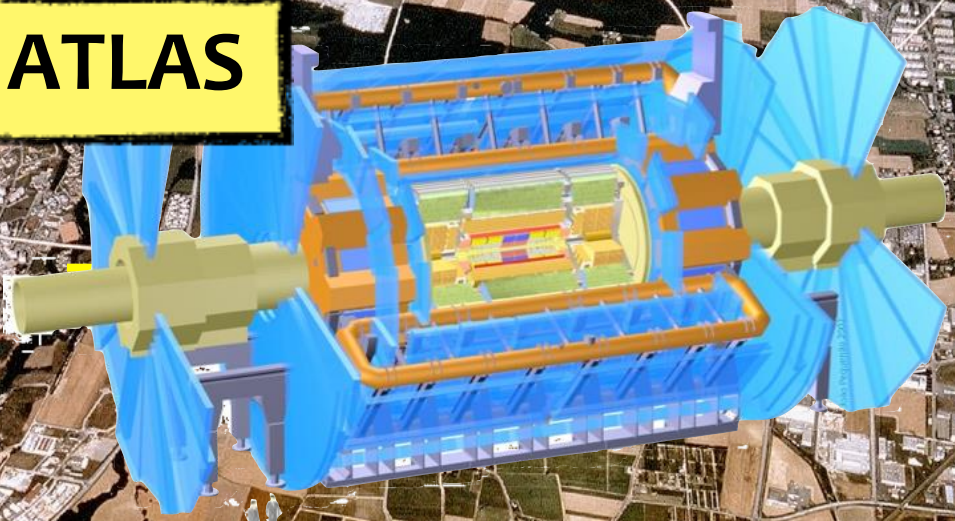

Imagery: CERN

Nikhefs neutrino-detector: KM3NeT

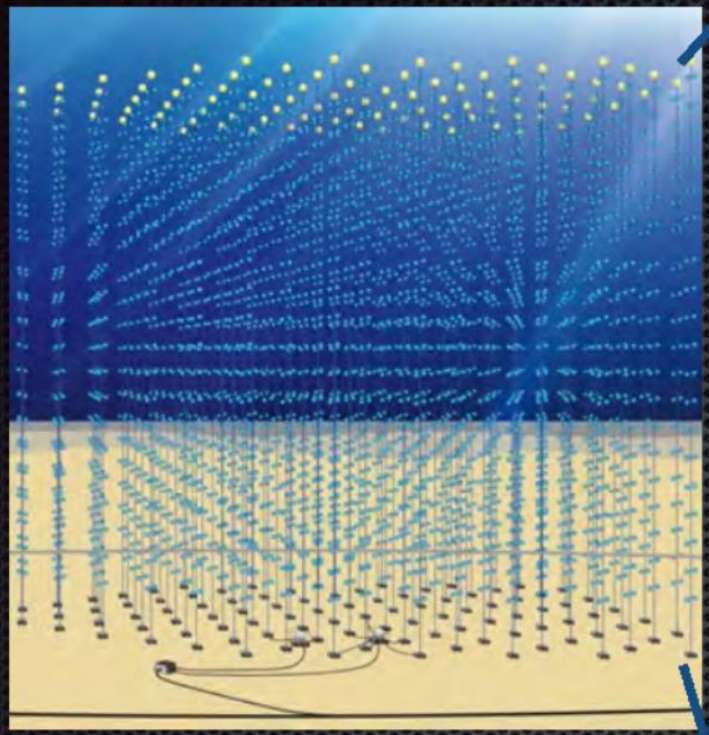

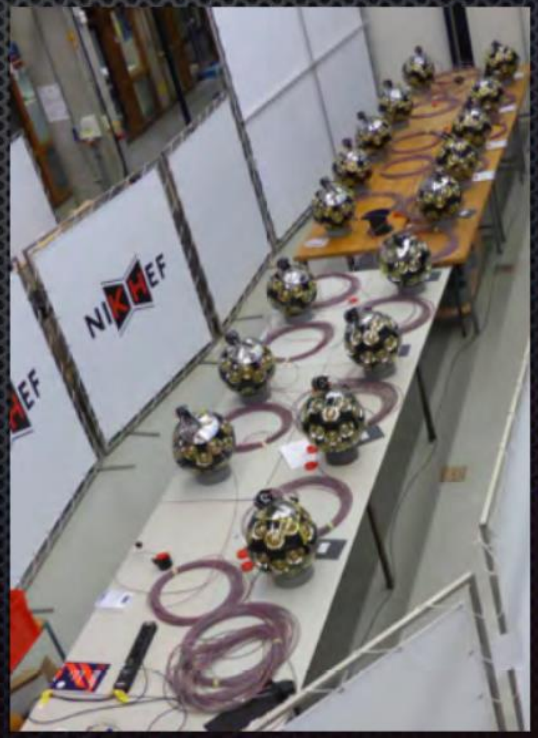

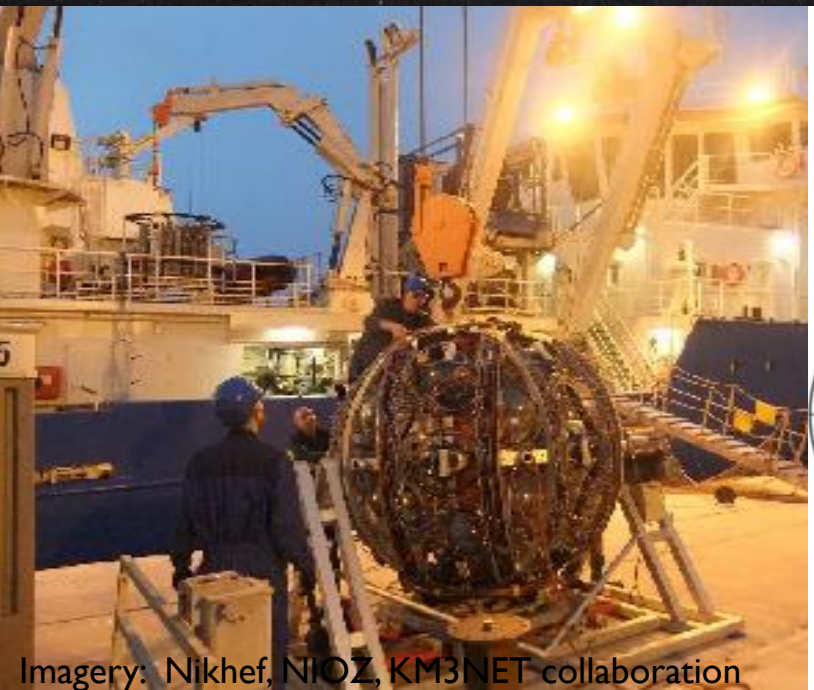

De Melkweg

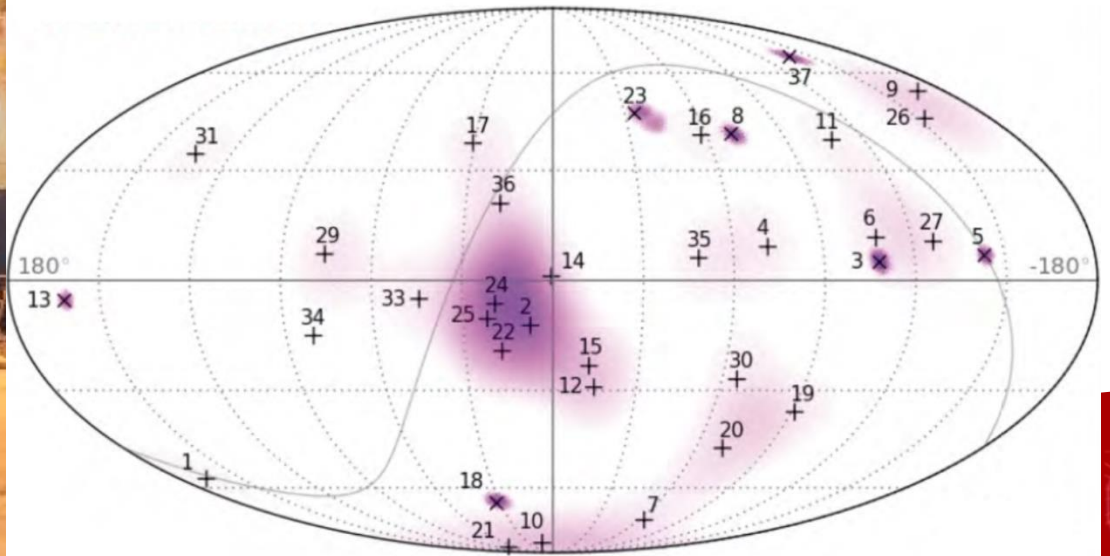

Imagery: Nikhef, NIOZ, KM3NET collaboration

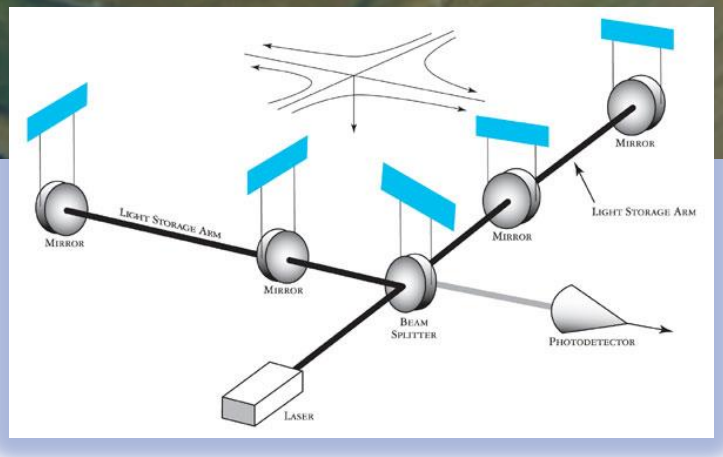

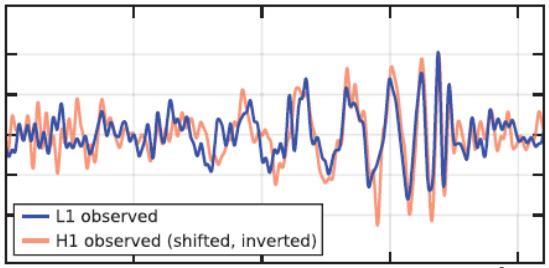

LSC LIGO Scientific Collaboration

VIRGO

Imagery: gw-astronomy collaborations, LSC

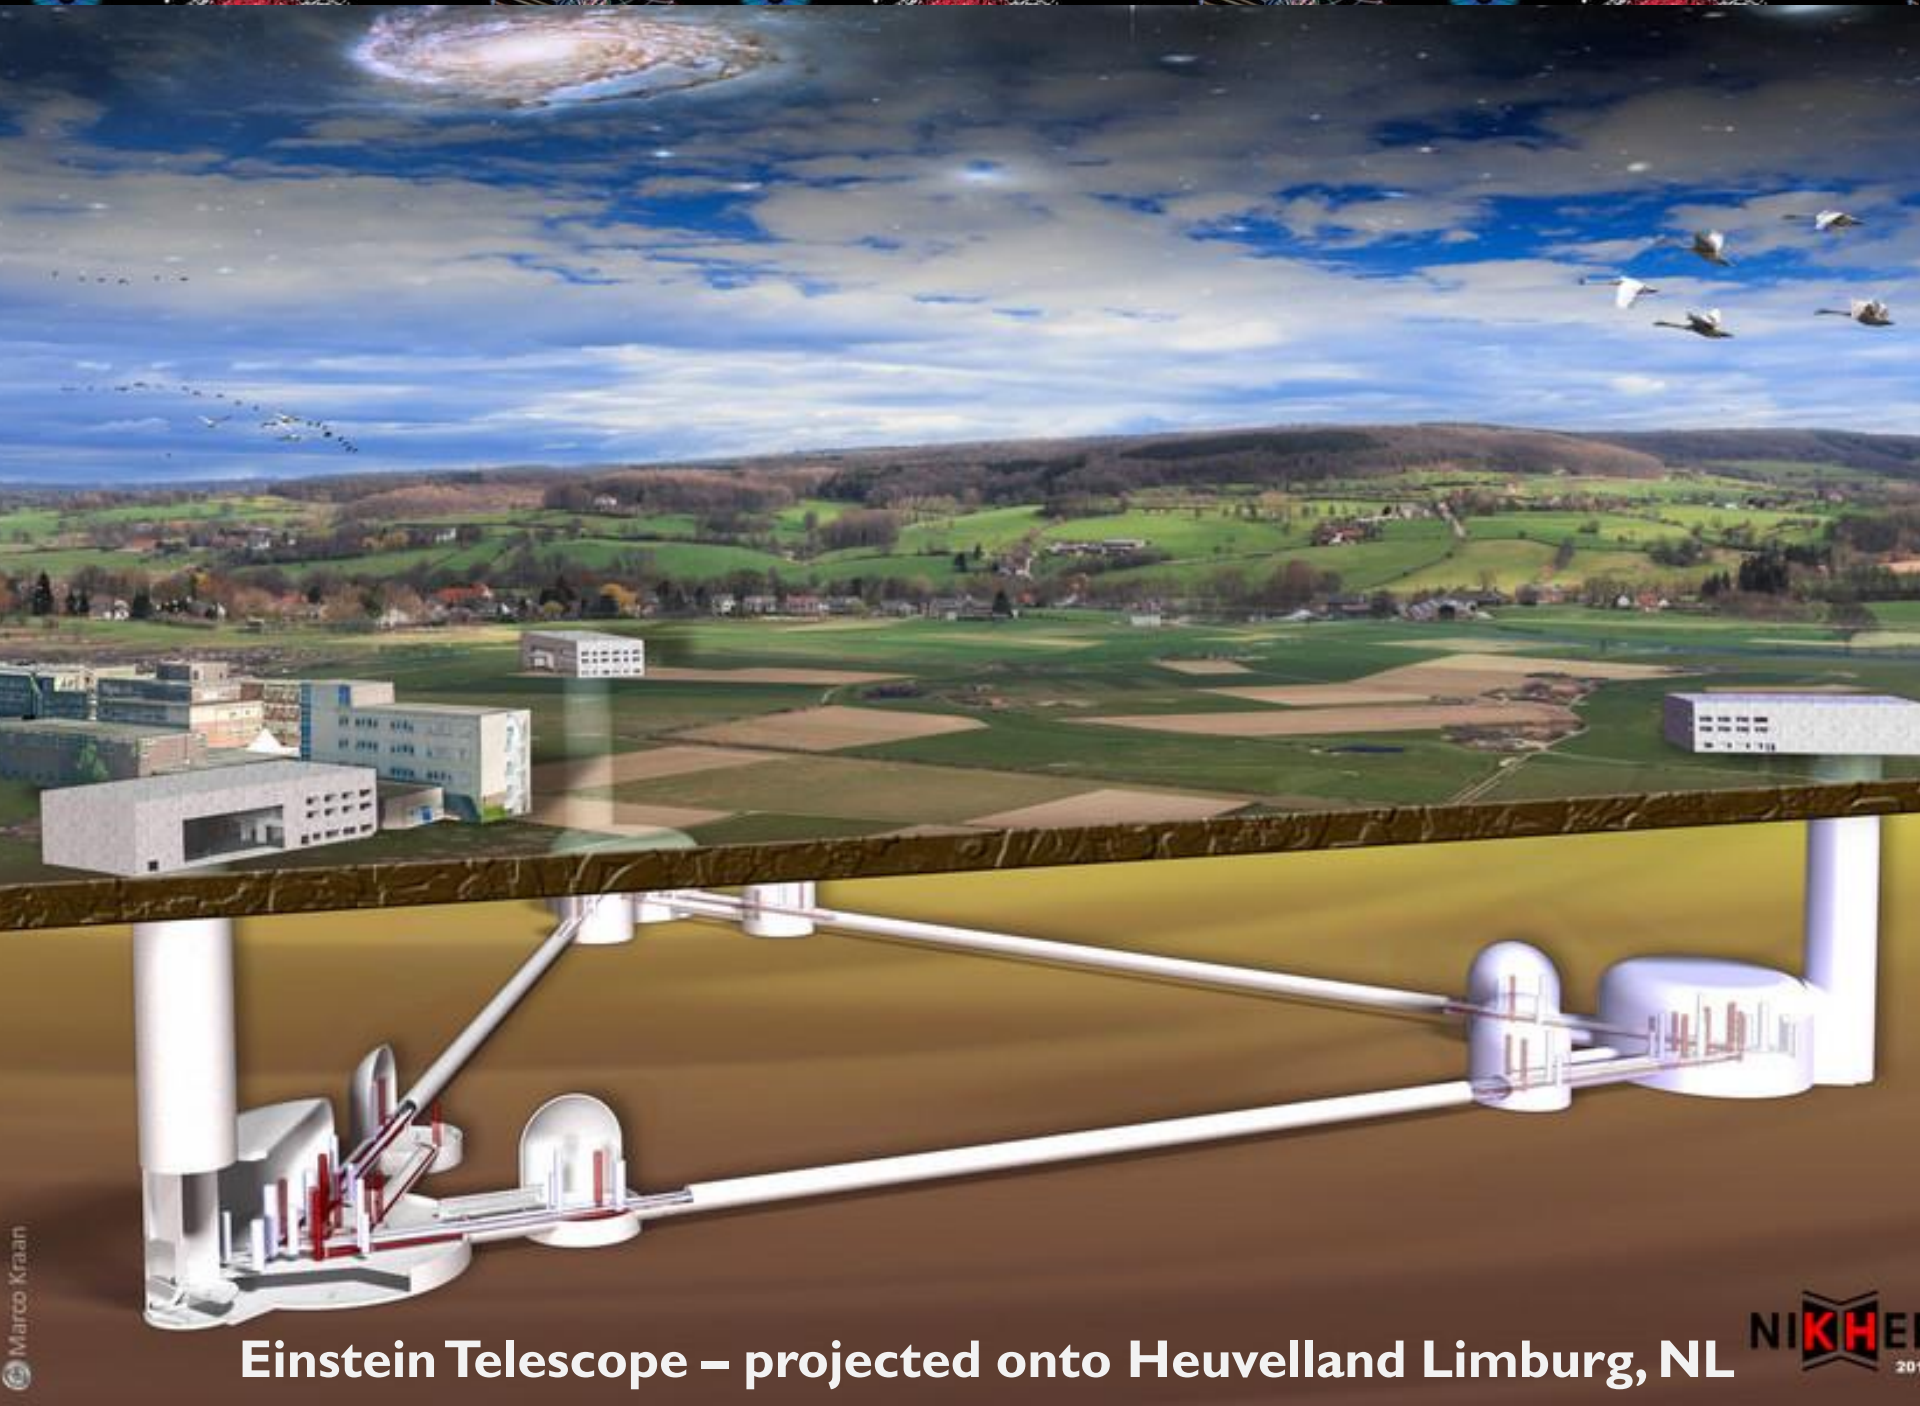

Einstein Telescope – projected onto Heuvelland Limburg, NL

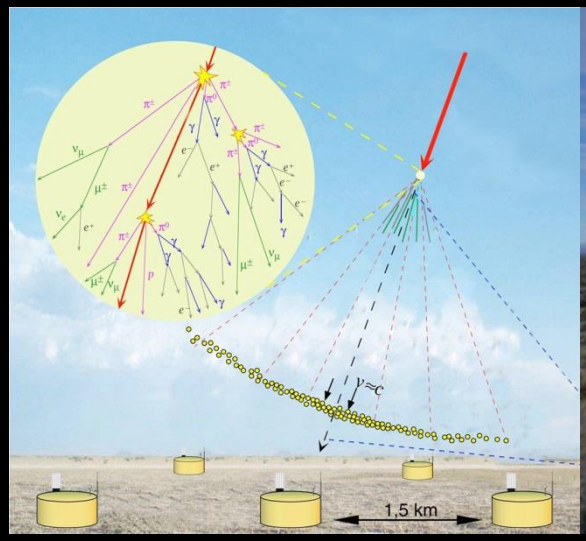

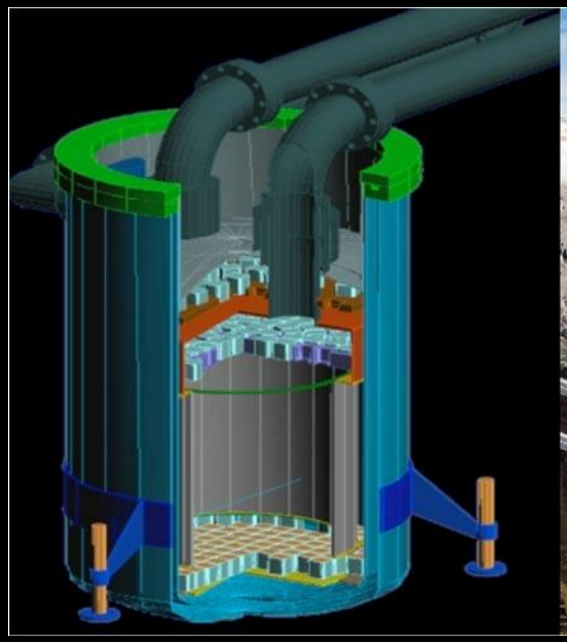

Image sources: LNGS/INFN, Xenon collaboration; Pierre Auger collaboration; Nikhef

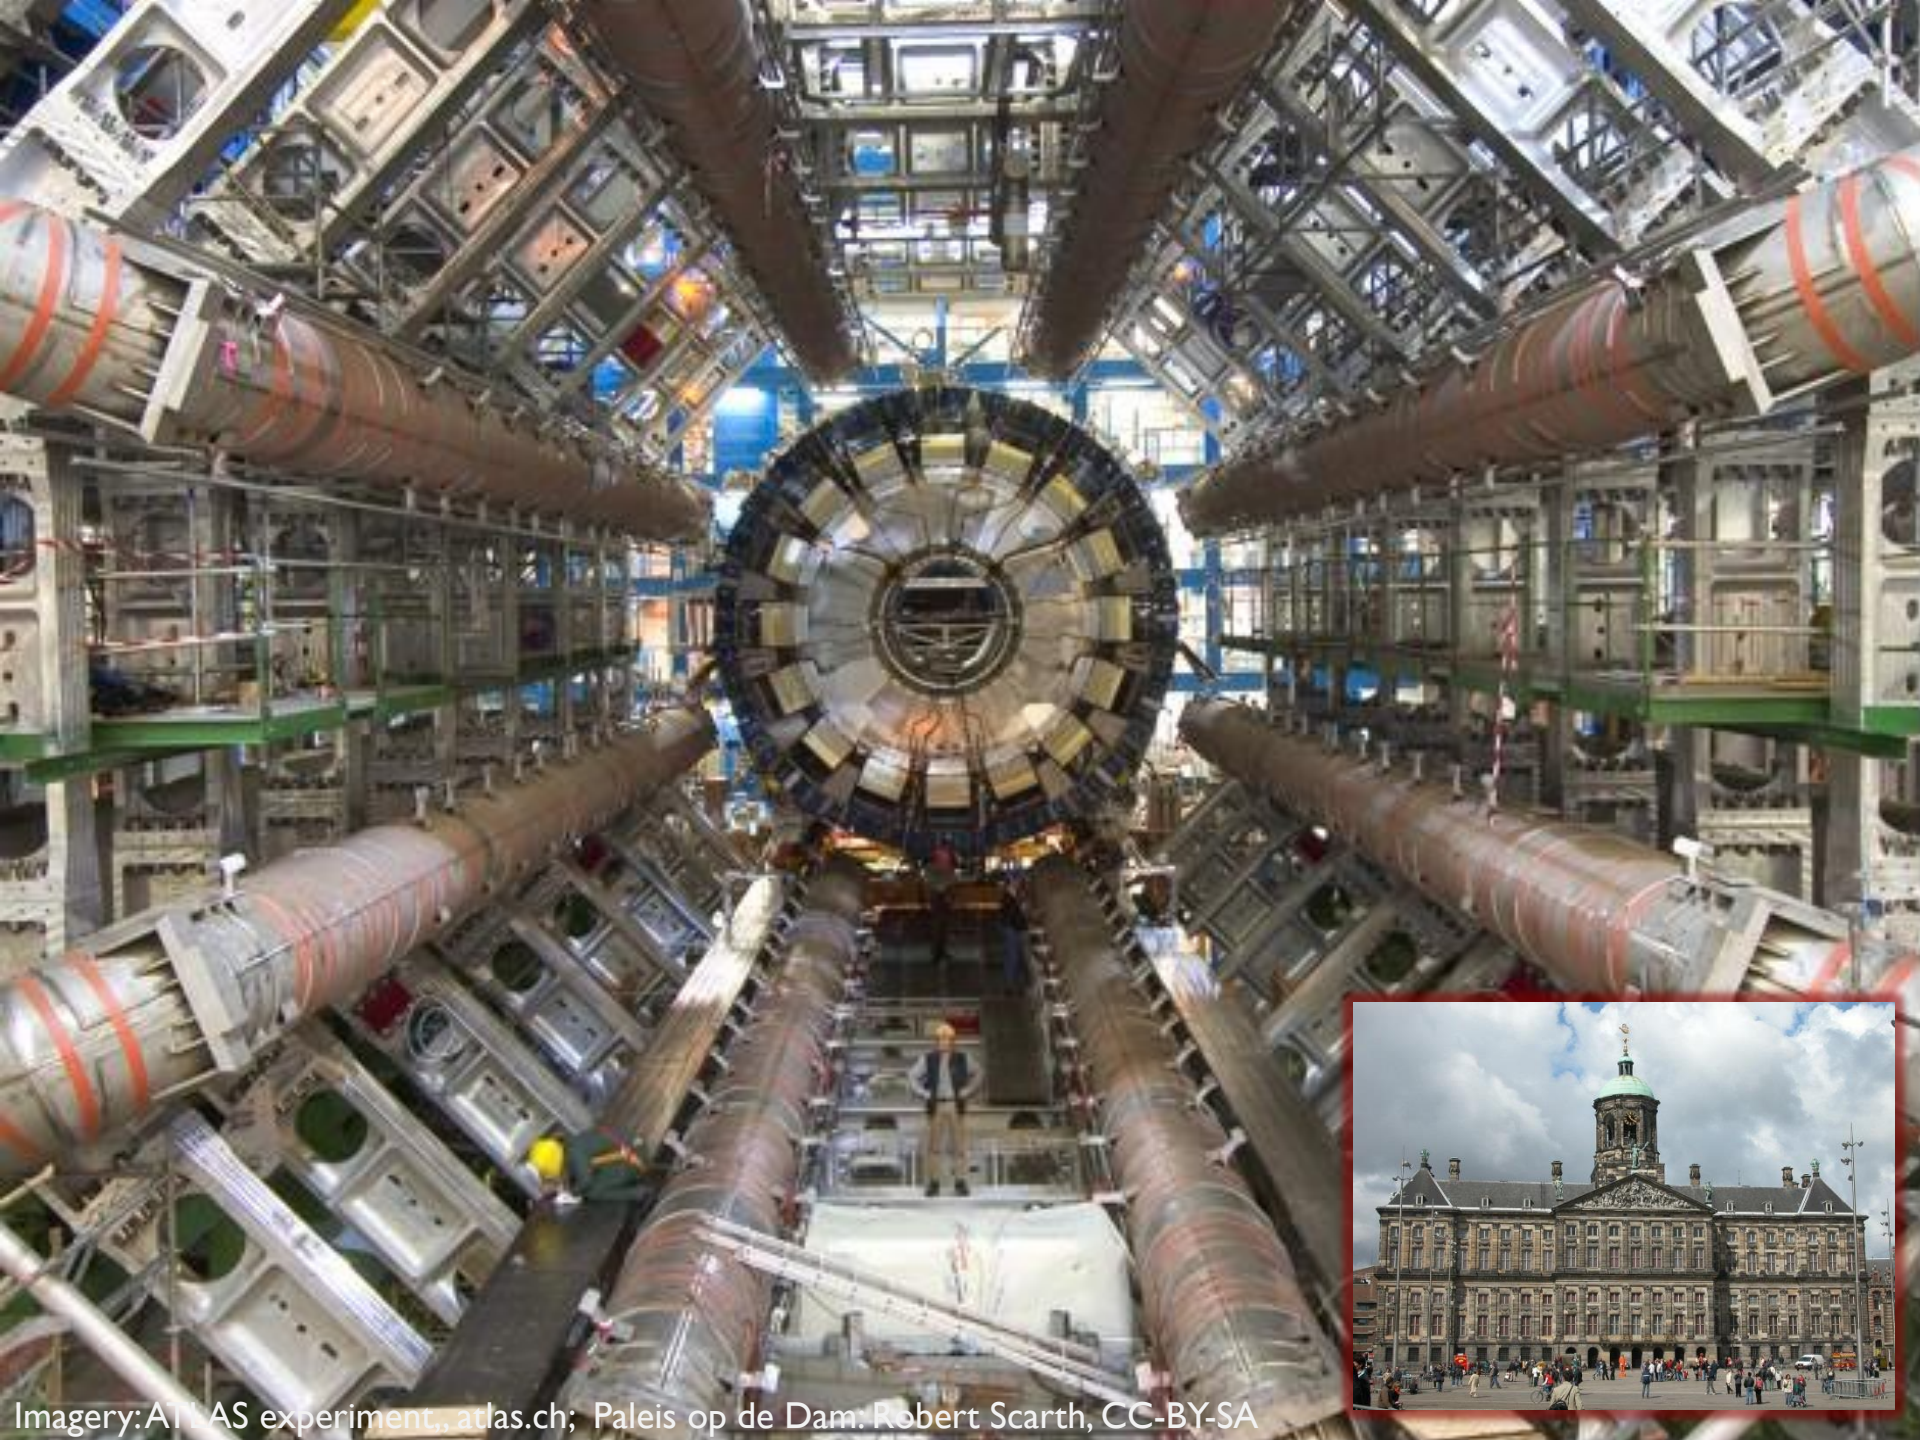

Imagery: ATLAS experiment,, atlas.ch; Paleis op de Dam: Robert Scarth, CC-BY-SA

Deeltjes botsingen

Deeltjes botsingen

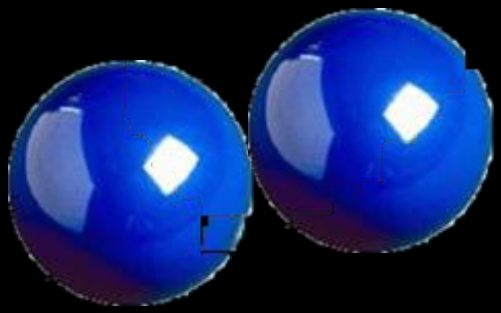

Deeltjes botsingen

$$E = mc^2$$

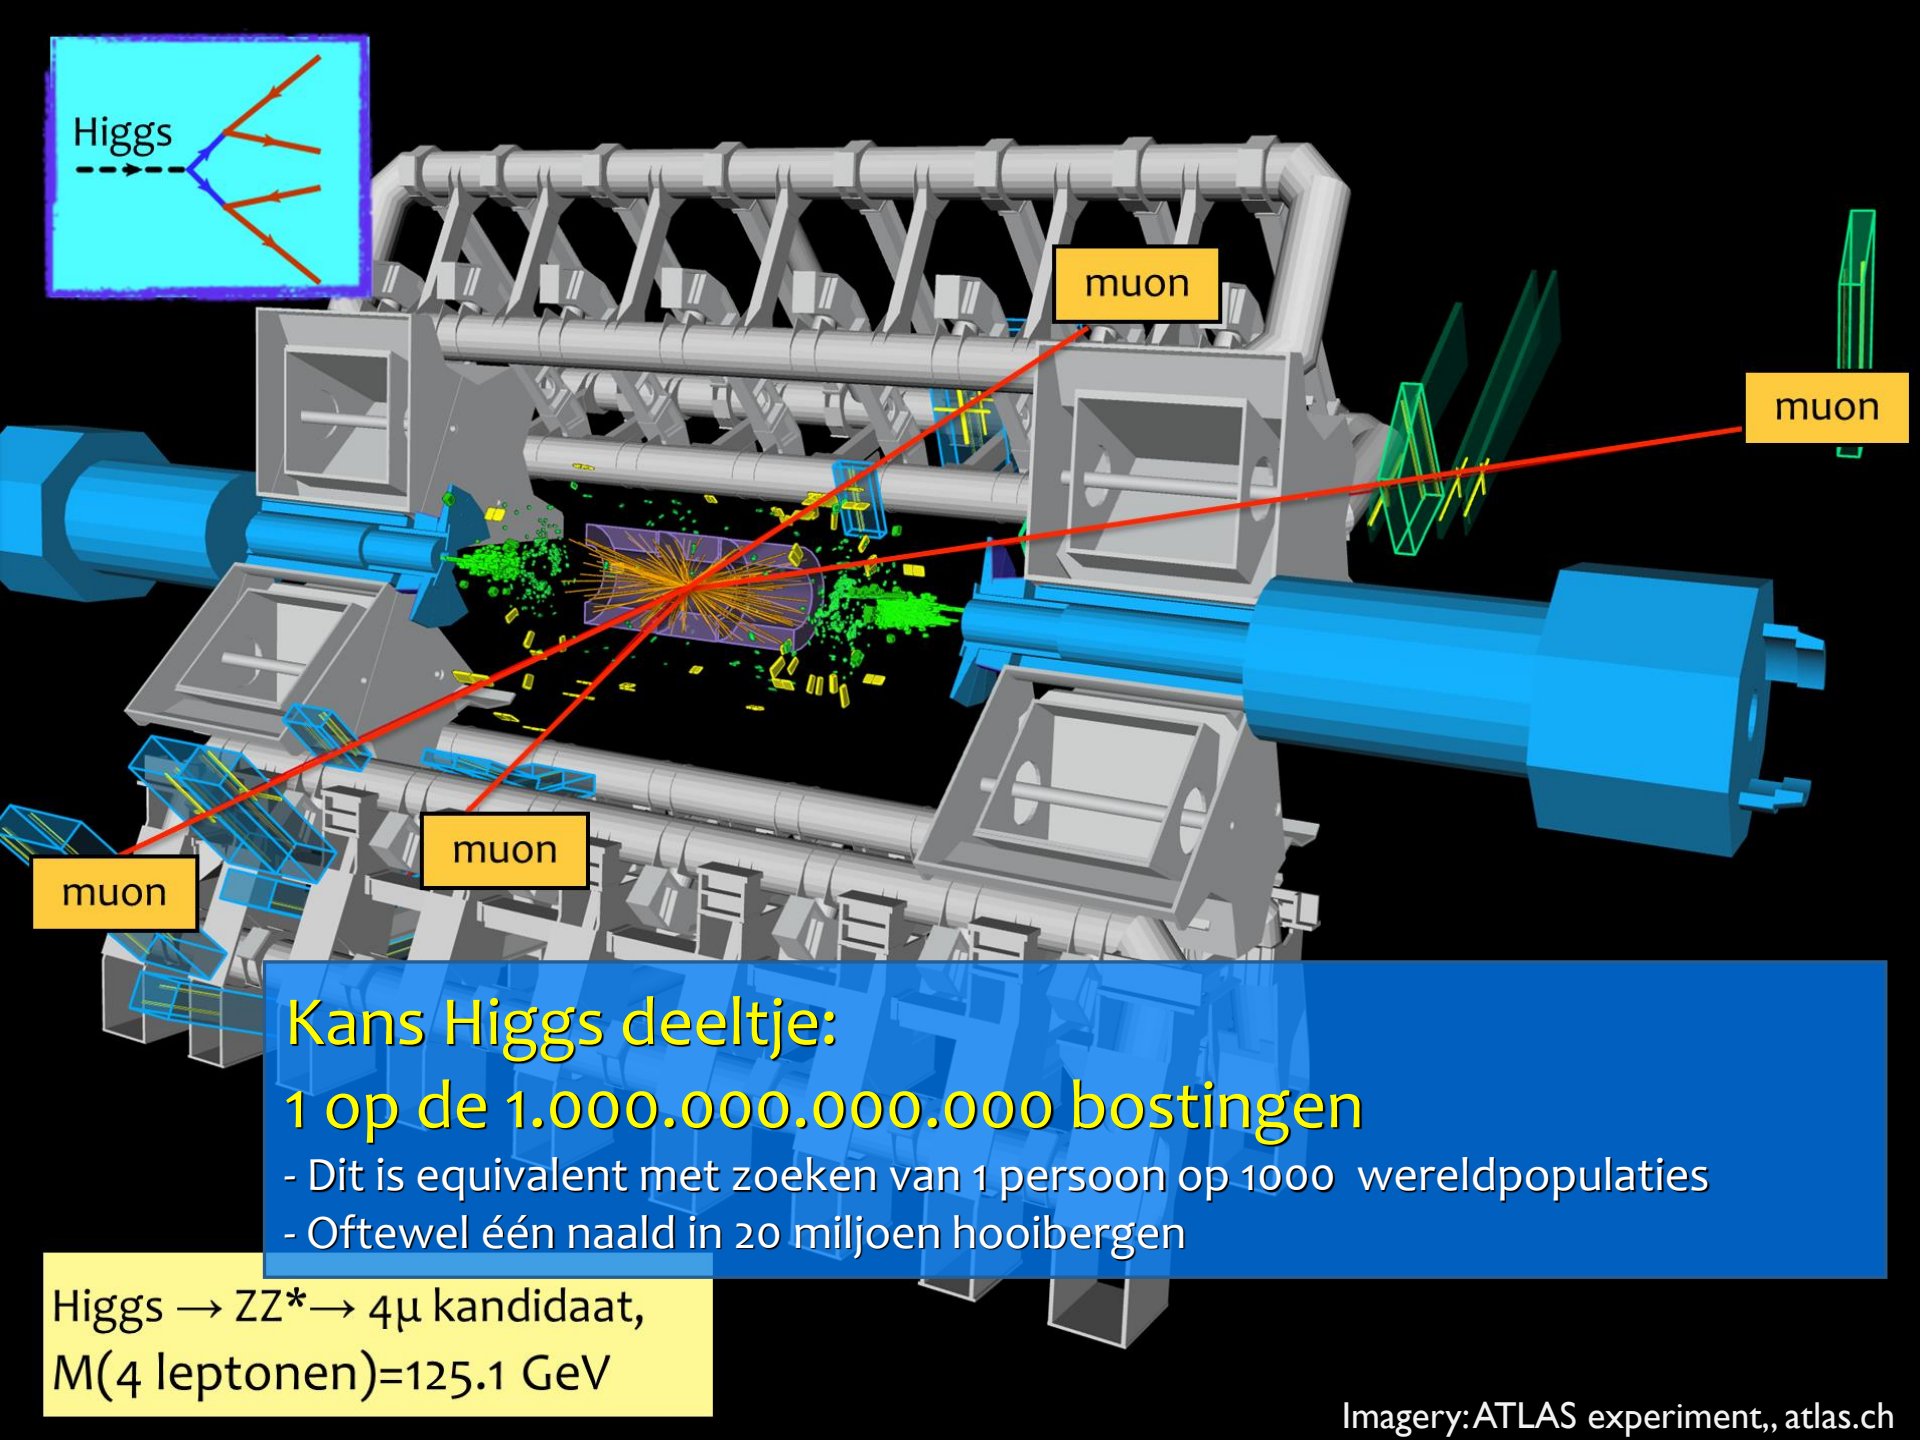

Kans Higgs deeltje:
1 op de 1.000.000.000.000 bostingen

- Dit is equivalent met zoeken van 1 persoon op 1000 wereldpopulaties
- Oftewel één naald in 20 miljoen hooibergen

Higgs → ZZ*→ 4μ kandidaat,
M(4 leptonen)=125.1 GeV

Imagery: ATLAS experiment,, atlas.ch

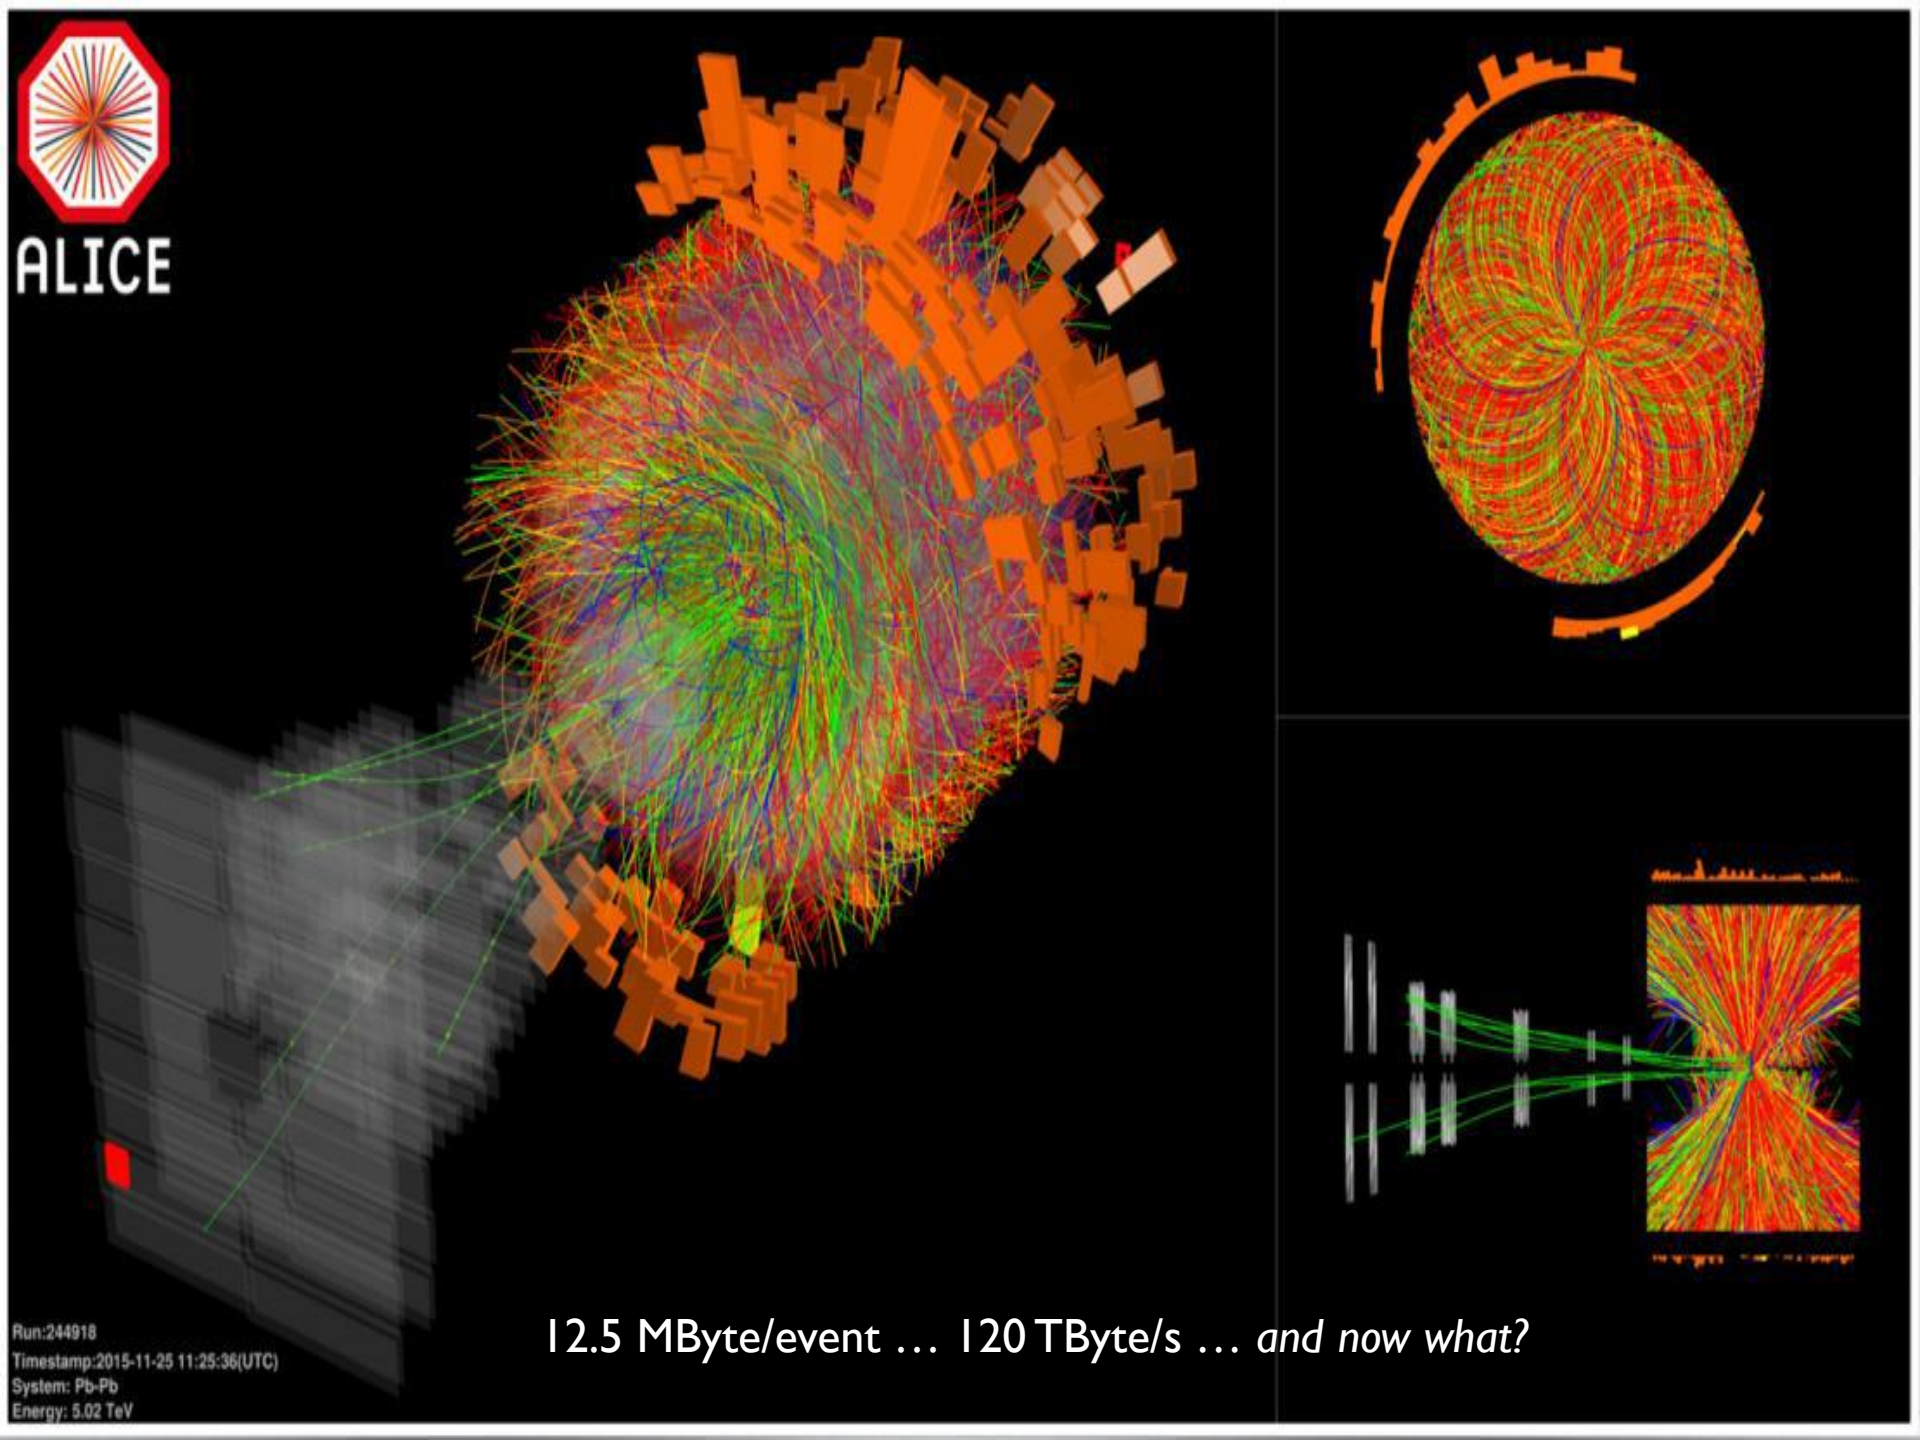

12.5 MByte/event … 120 TByte/s … *and now what?*

Detector to doctor ...

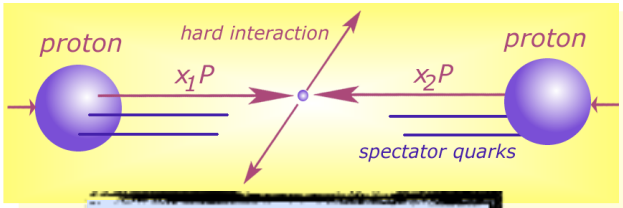

40 miljoen / seconde

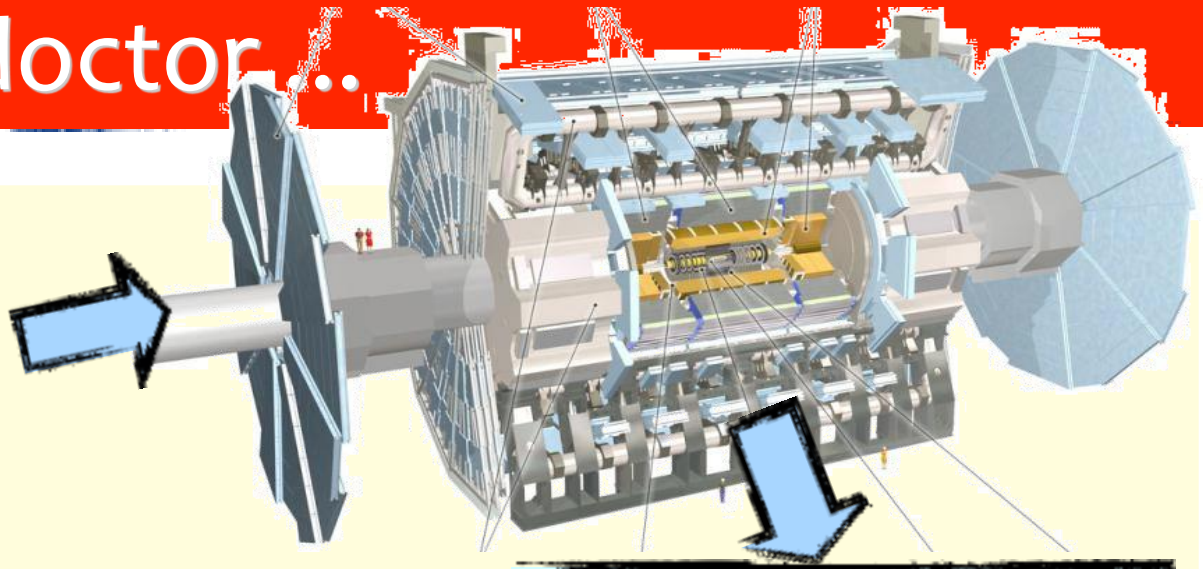

Trigger systeem selecteert 600 Hz
~ 1 GB/s data

Analyse van botsingen door promovendi

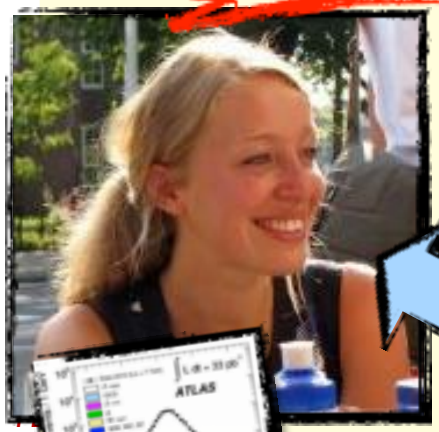

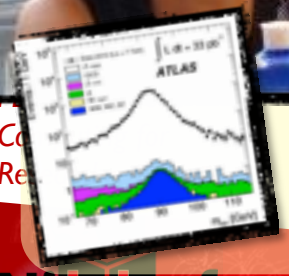

and processing

Data distributie met GRID computers

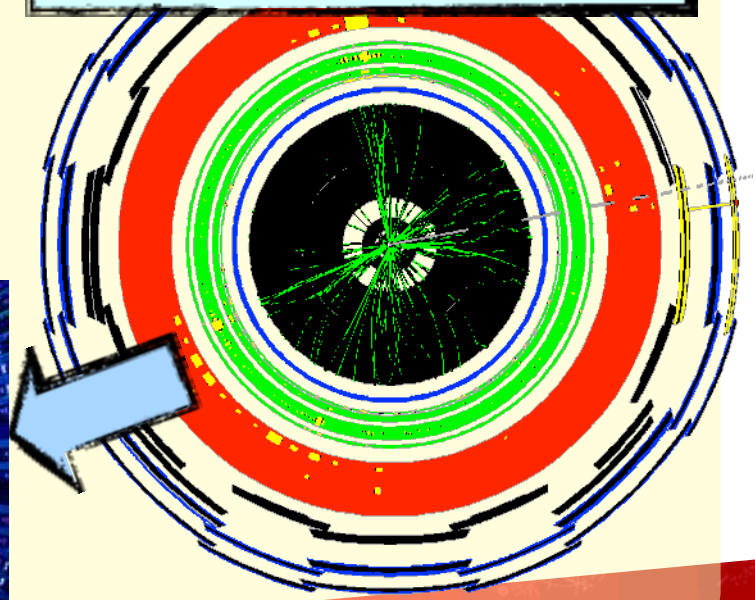

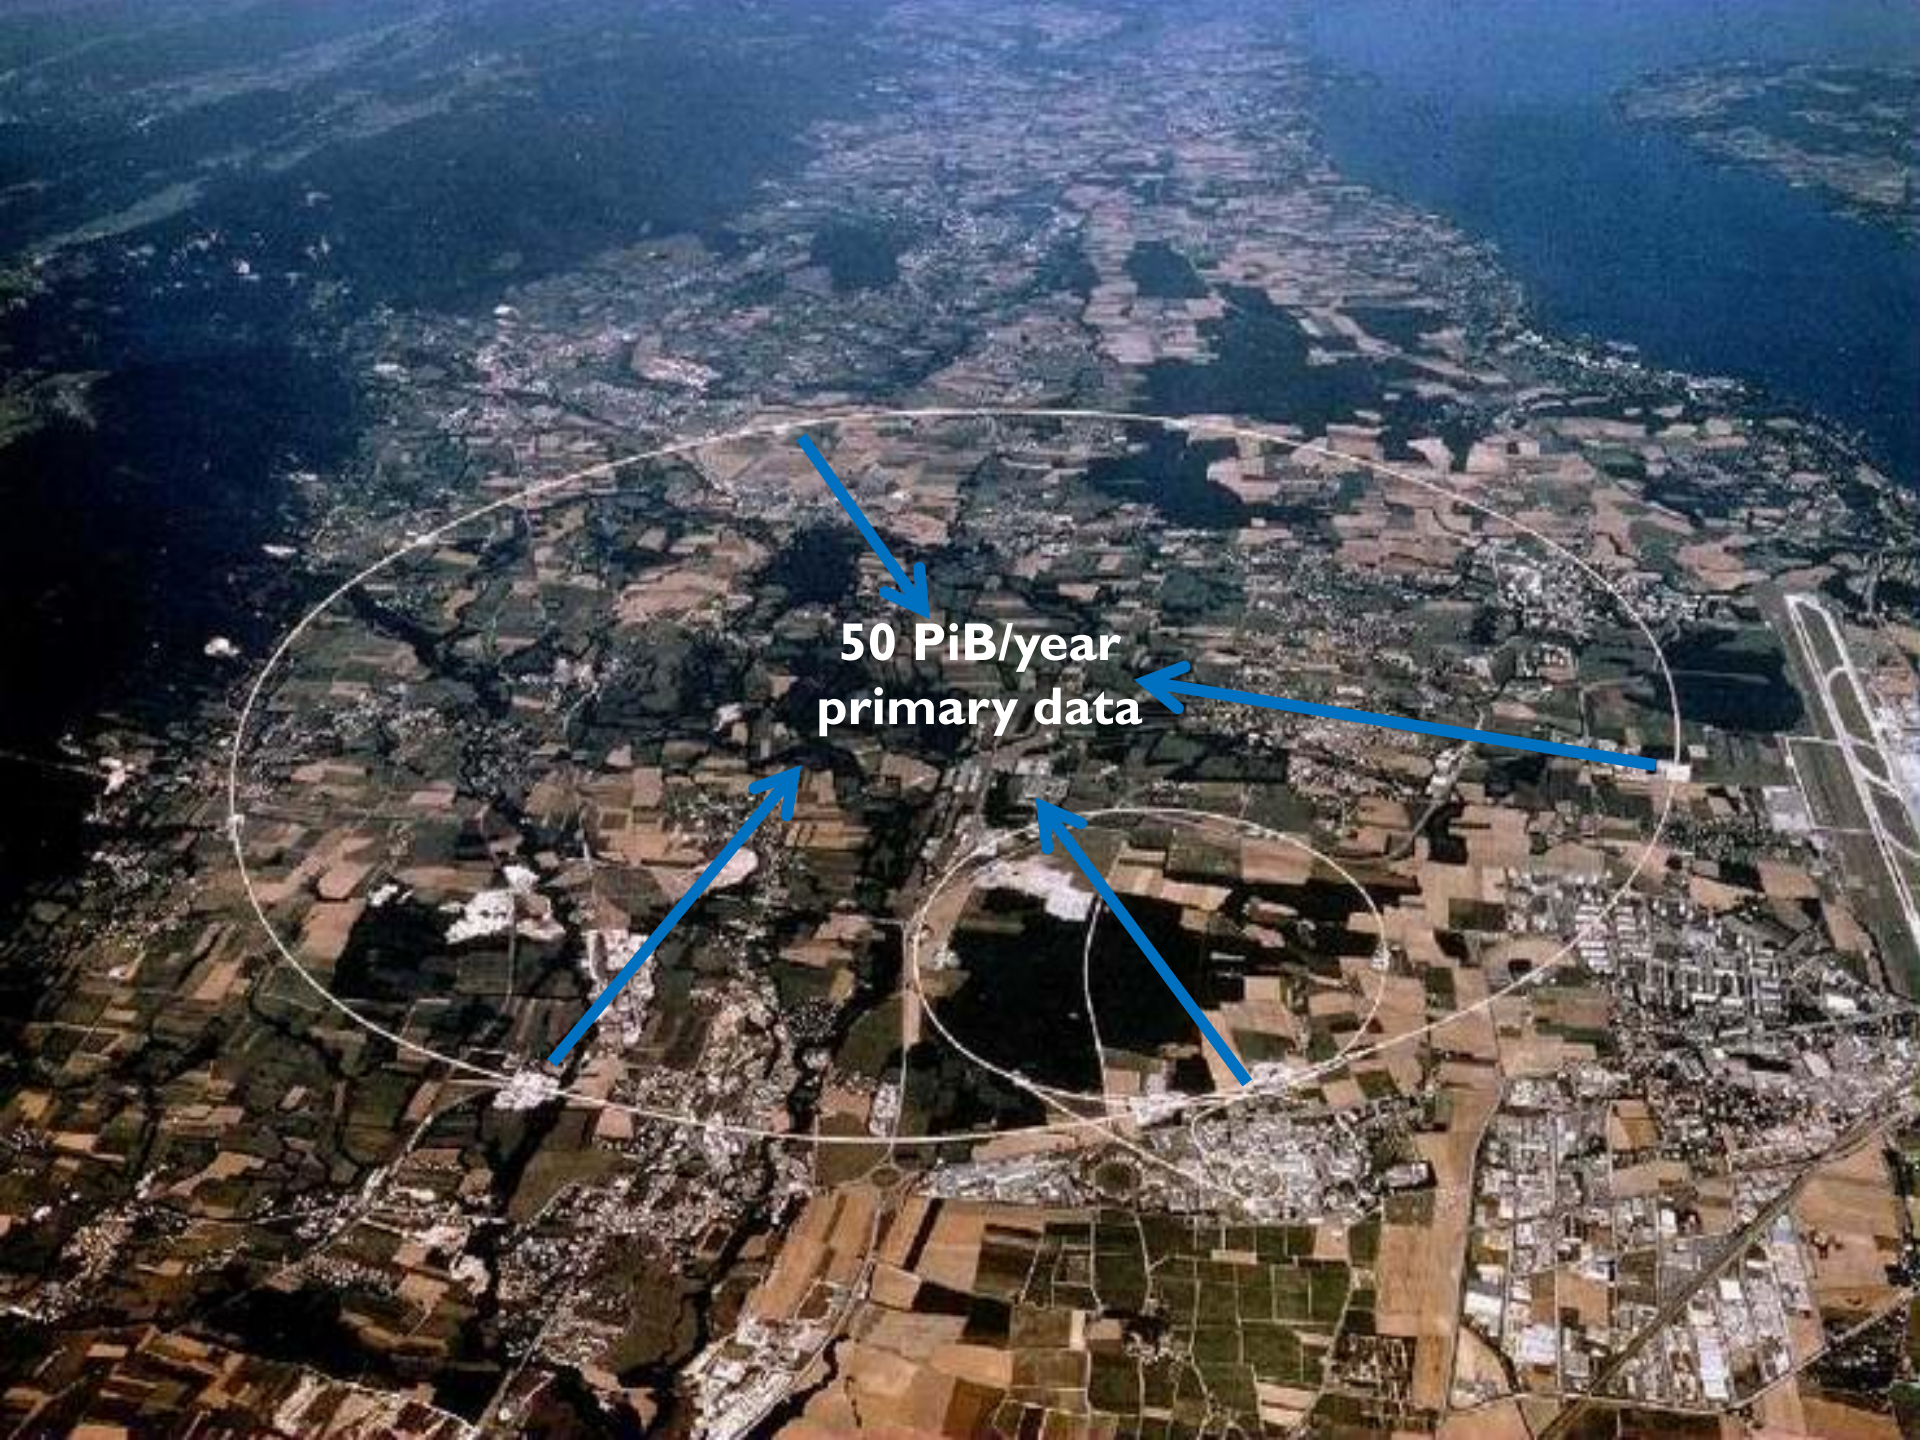
50 PiB/year
primary data

Image source: joint CERN (wLCG) and EGI

Organisations participating in the global collaboration of e-Infrastructures

Even just for wLCG, supporting the CERN LHC programme

More than 200 independent institutes with end-users
More than 50 countries & regions
More than 300 service centres
One independent 'policy-bridge' identity service
Handful regional 'service coordination organisations'
500 000 CPU cores, 200+PByte storage

PRACE

Open Science Grid

EGI

XSEDE
Extreme Science and Engineering Discovery Environment

From SC04, CCRC08, STEP09, .. to today

Global transfer rates now > 40 GB/s – acquisition:10 PB/mo (~x2 derived data)

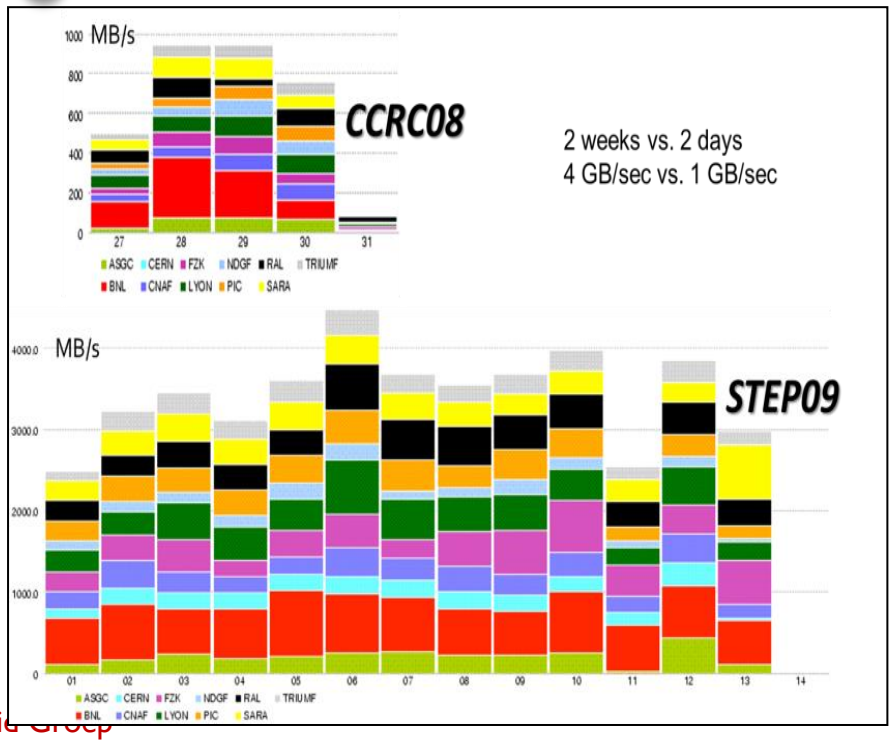

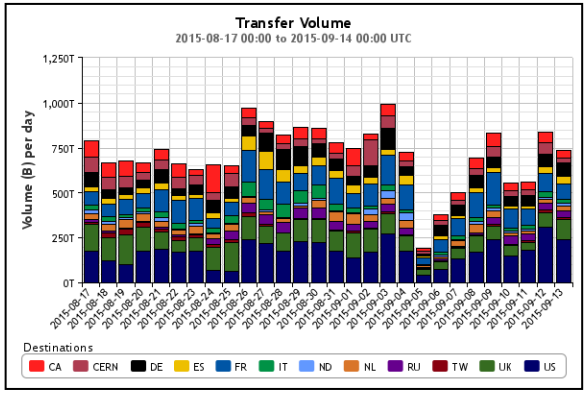

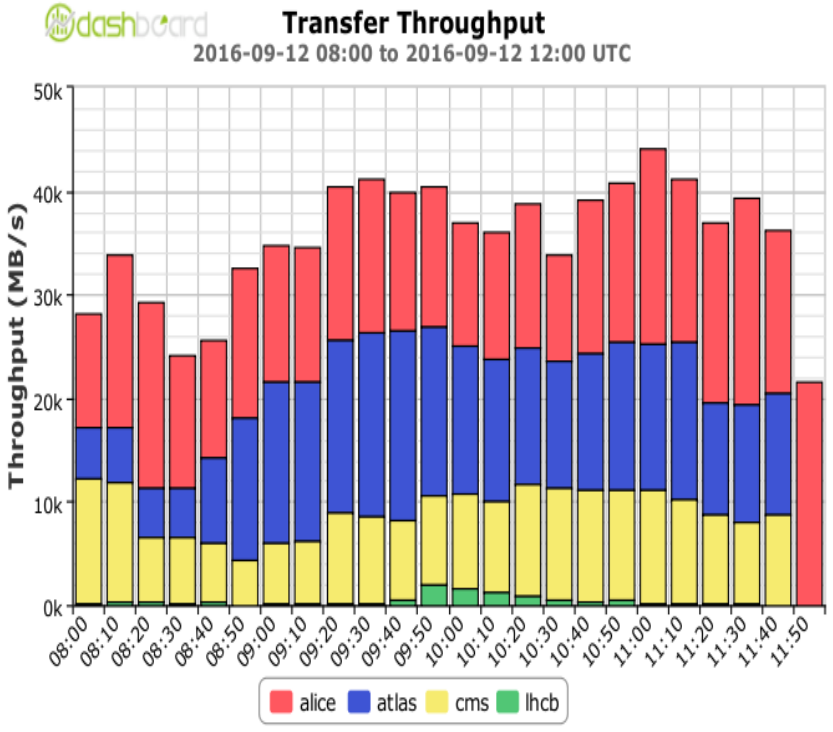

David Groep
Nikhef
PDP - Advanced Computing for Research

Atlas: ~50 TByte/day raw data to tape; 1000 TByte/day processed data transfers

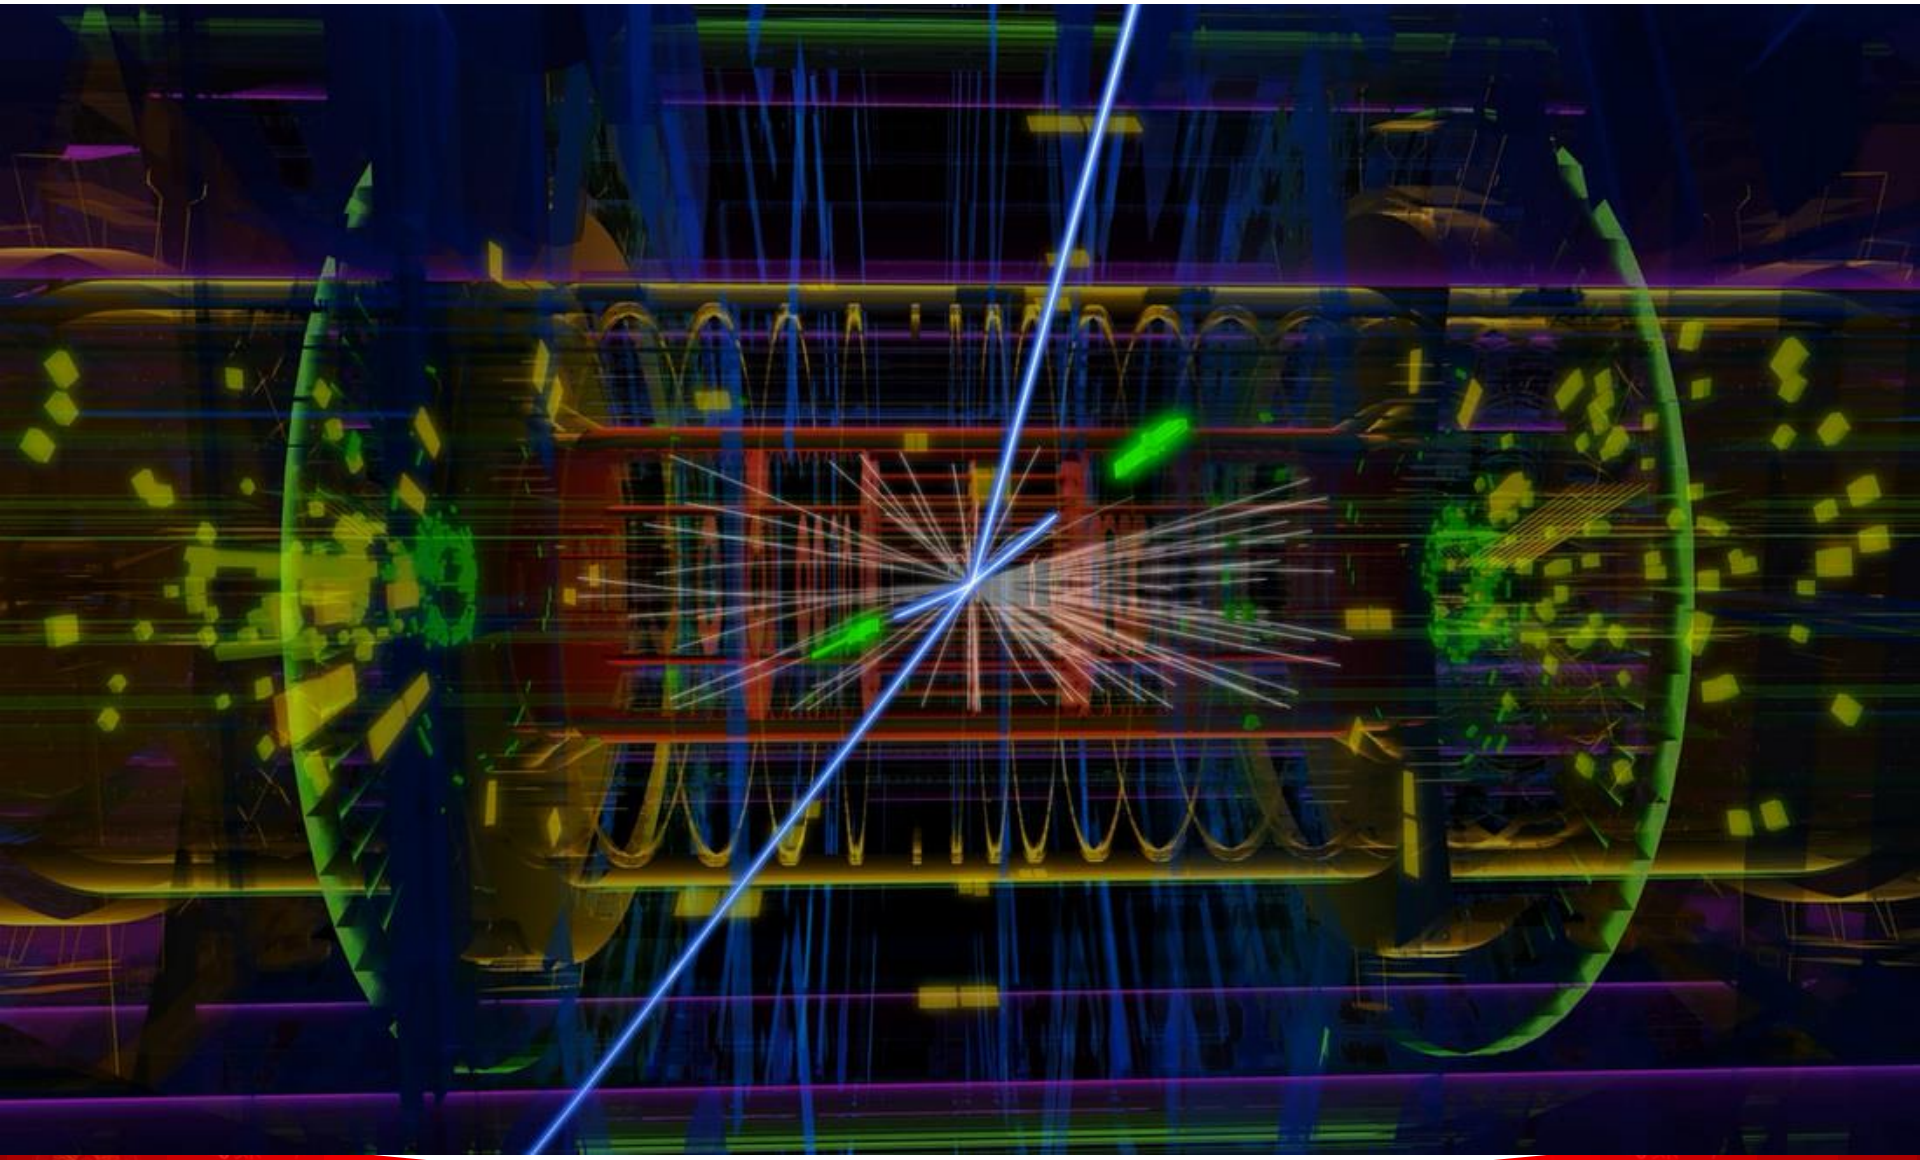

Nikhef

Image source: CERN, Atlas

… and tomorrow ?!

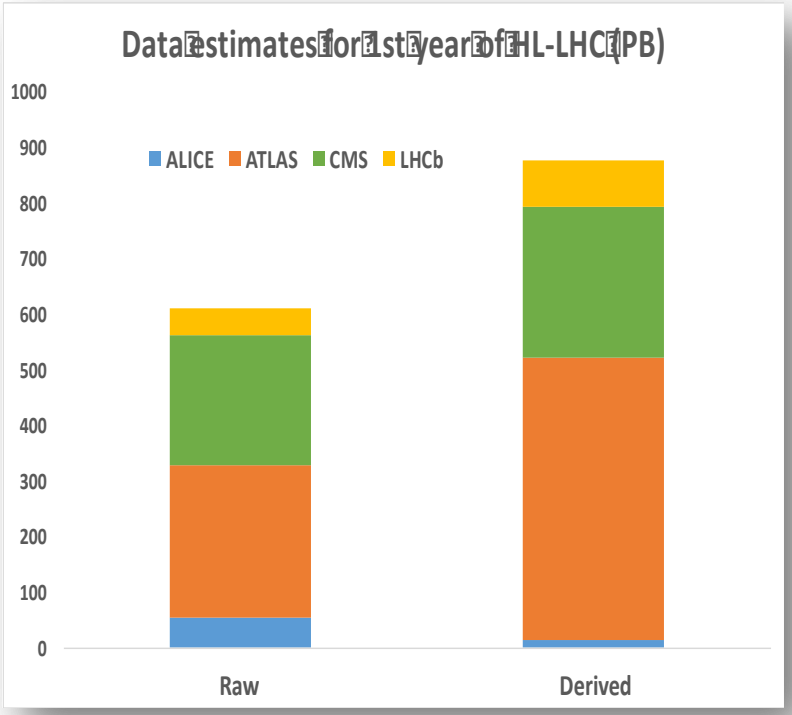

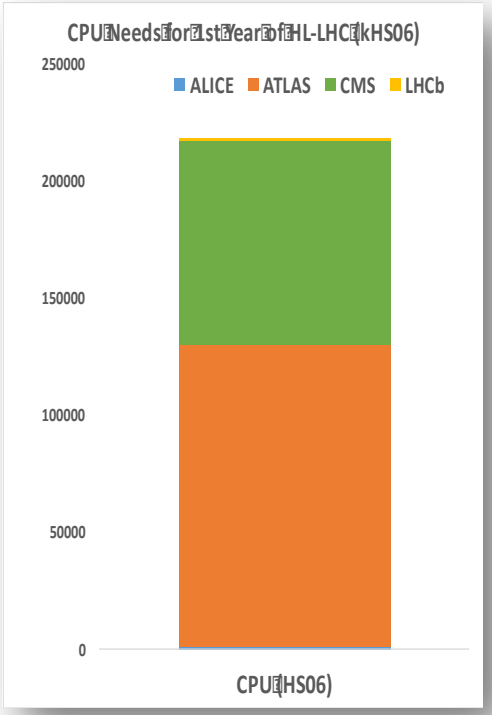

Data:

- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU:

- x60 from 2016

Technology at ~20%/year will bring x6-10 in 10-11 years

David Groep
Nikhef
PDP - Advanced Computing for Research

Infrastructure for research: balancing network, CPU, and disk

- CPU and disk both expensive, yet idling CPUs are 'even costlier'
- architecture and performance matching averts any single bottleneck
- but requires knowledge of application (data flow) behaviour data pre-placement (local access), mesh data federation (WAN access)

This is why e.g. your USB drive does not cut it – and neither does your 'home NAS box'

… however much I like my home system using just 15 Watt idle and offering 16TB for just € 915 …

Building the infrastructure for the LHC data

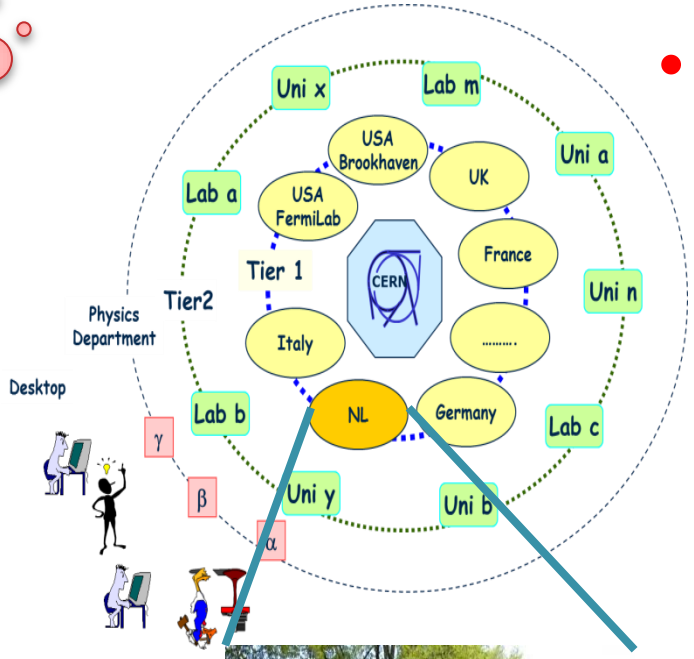

- From hierarchical data distribution to a full mesh and dynamic data placement

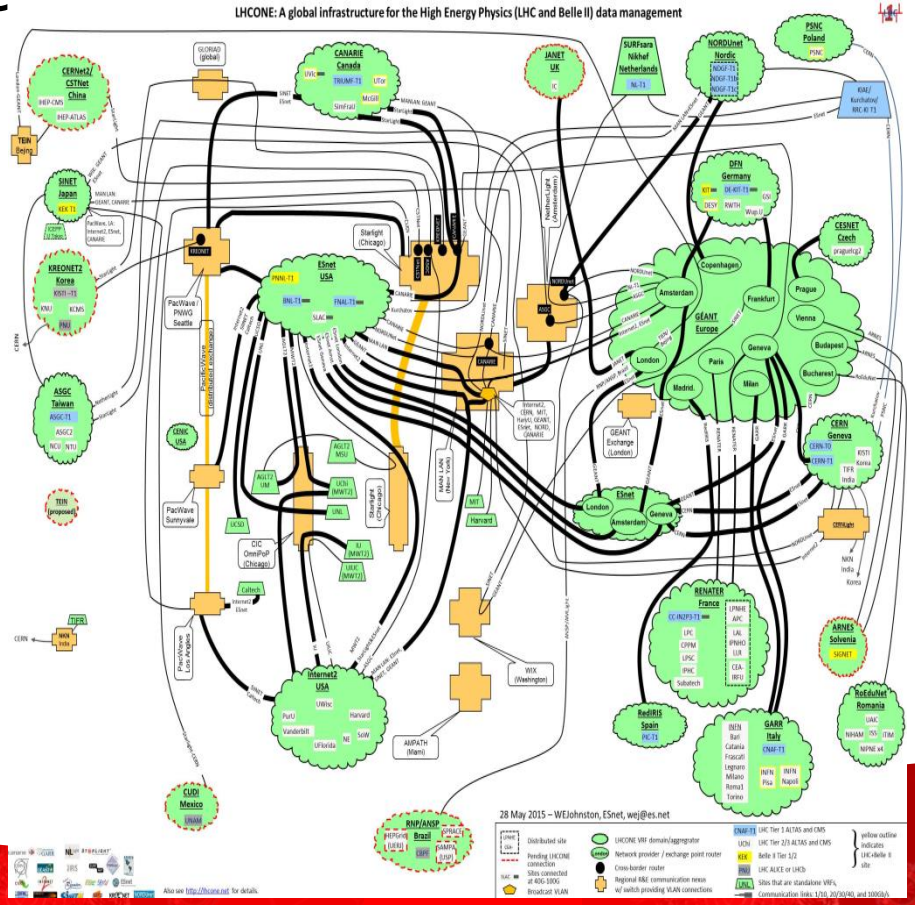

David Groep
Nikhef
PDP - Advanced Computing for Research

LHCOne graphic: lhcone.net

Connecting Science through Lambdas

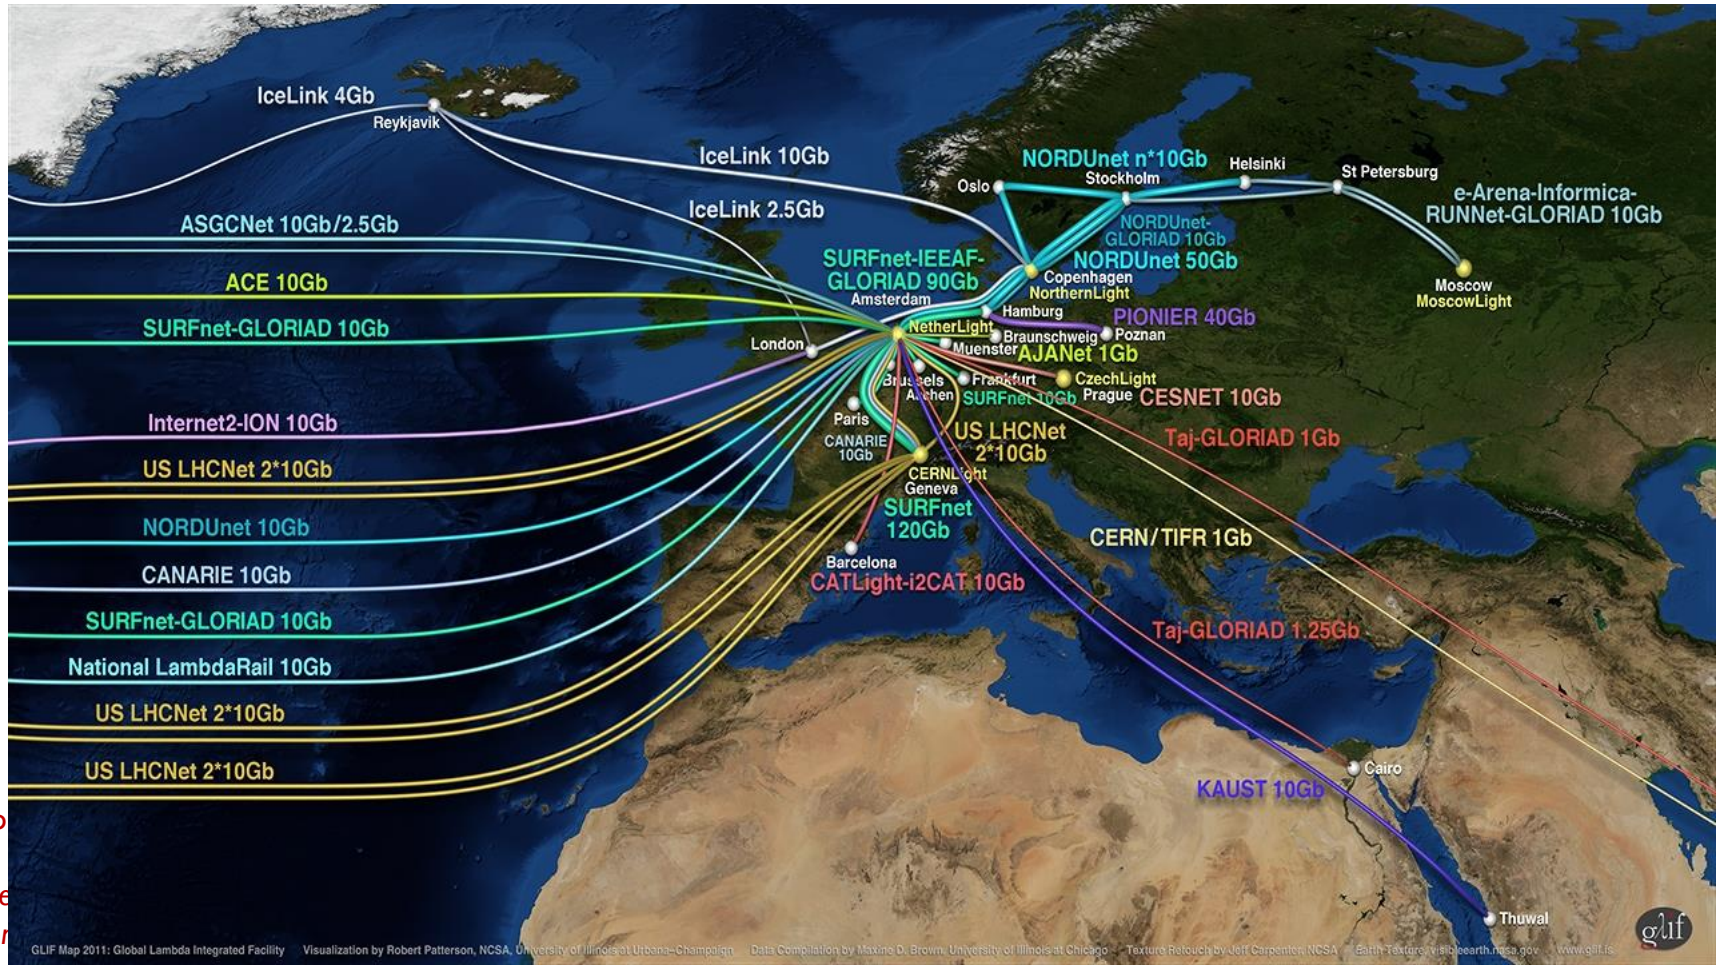

David Groep
Nikhef
PDP - Advance
Computing for
Research

Nikhef

Network built around application data flow

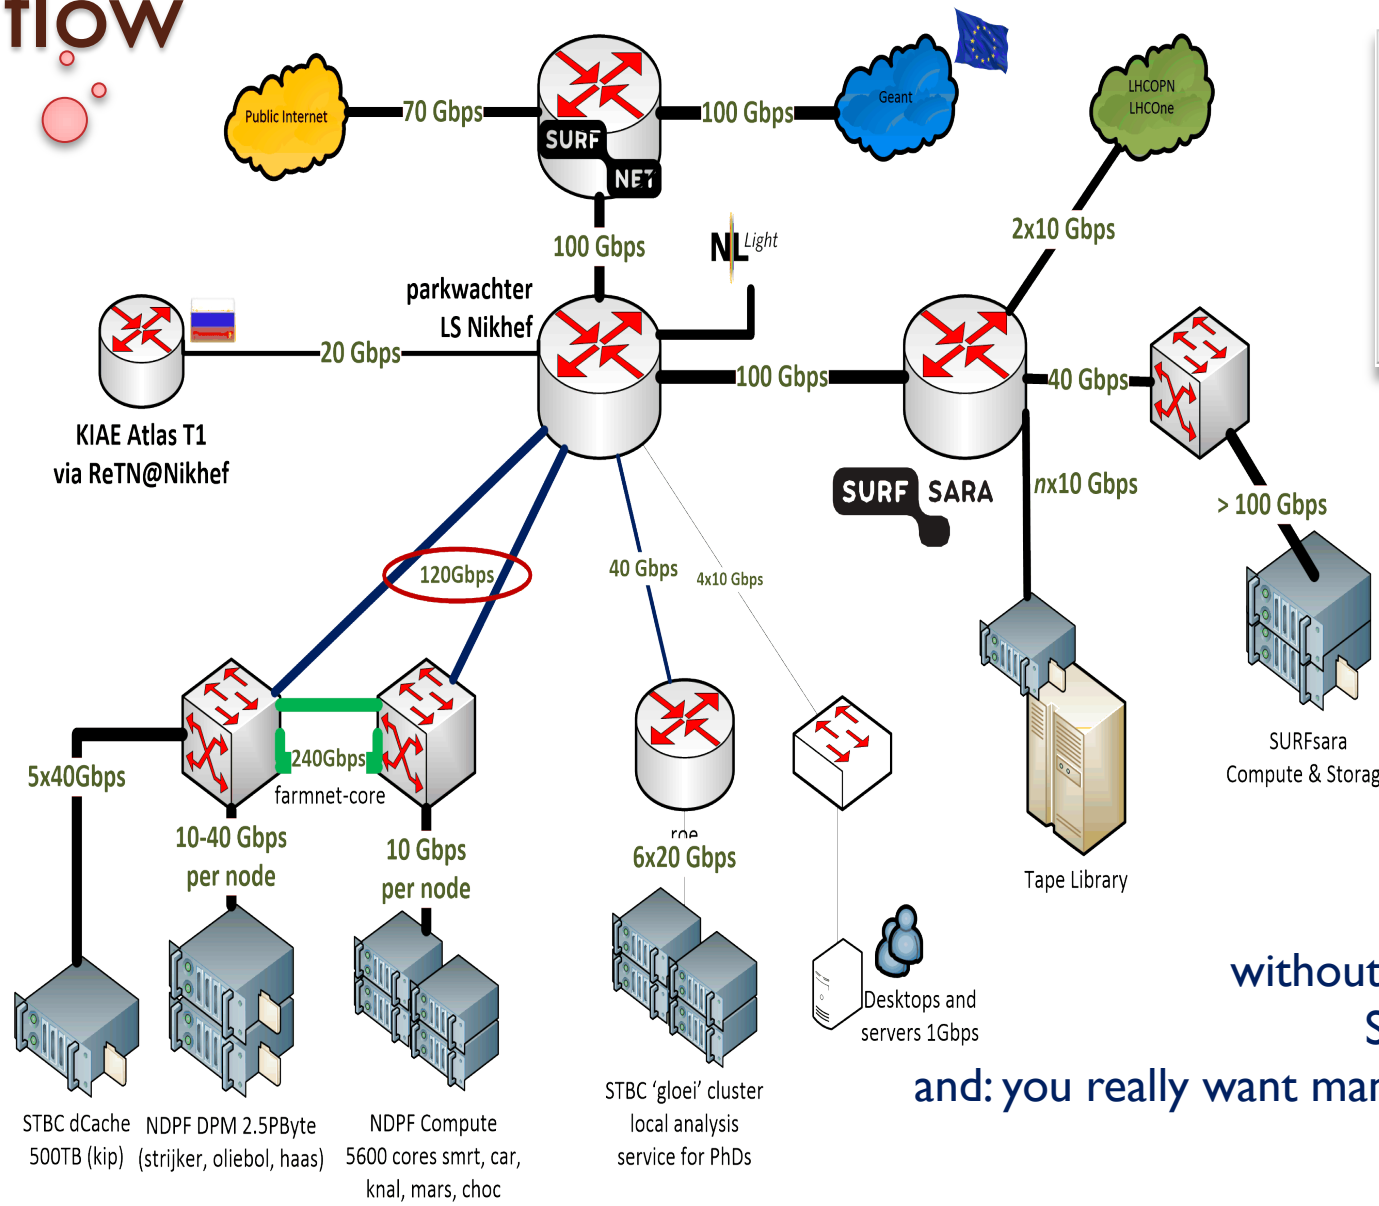

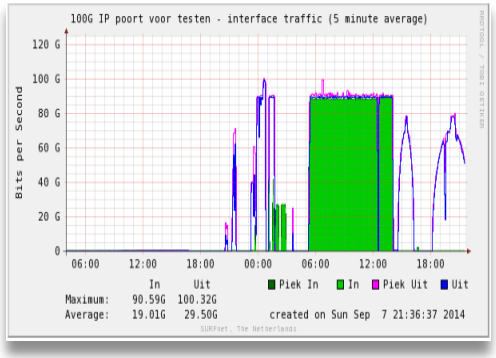

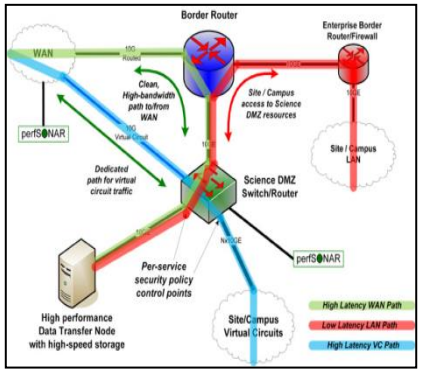

***Need** to work together!*

without our SURFsara peering, SURFnet gets flooded ☺

and: you really want many of your own peerings

100Gbit

Nikhef → SURFnet → RUG-CIT||UvA

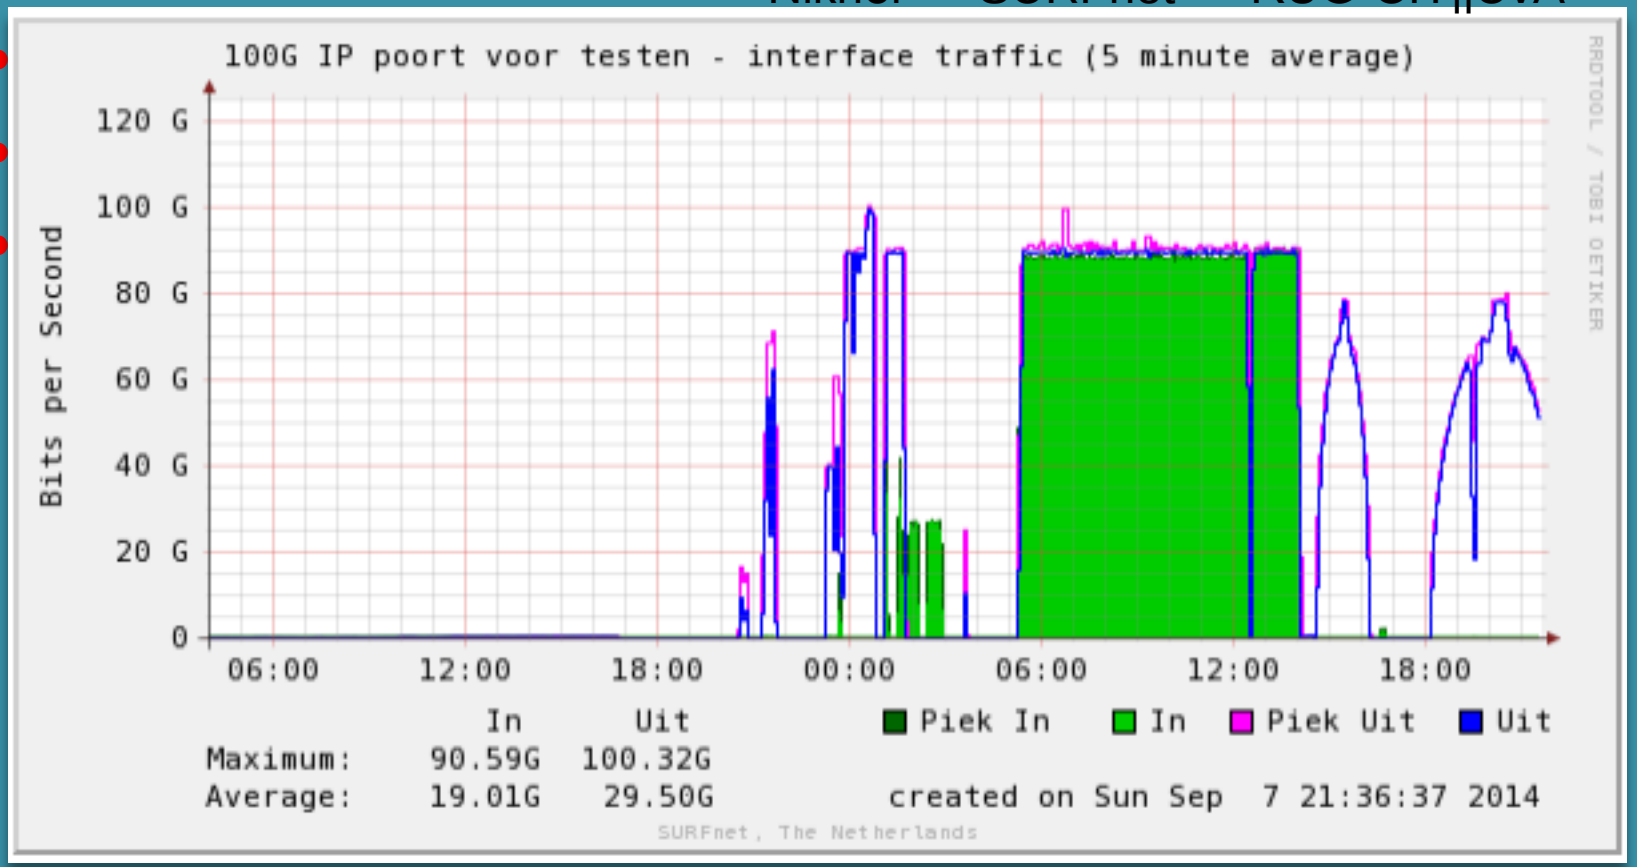

T Suerink
Nikhef
Amsterdam
PDP & Grid

slides courtesy Tristan Suerink

26

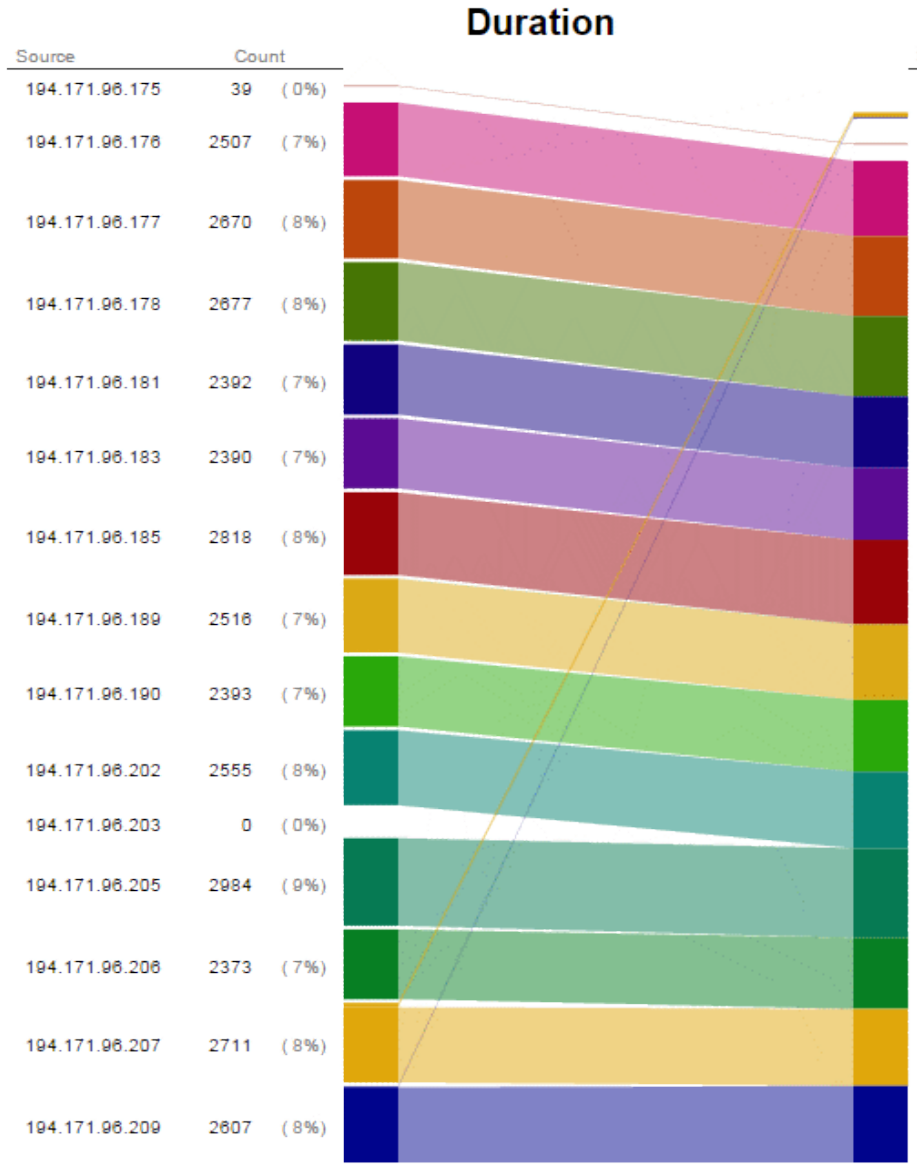

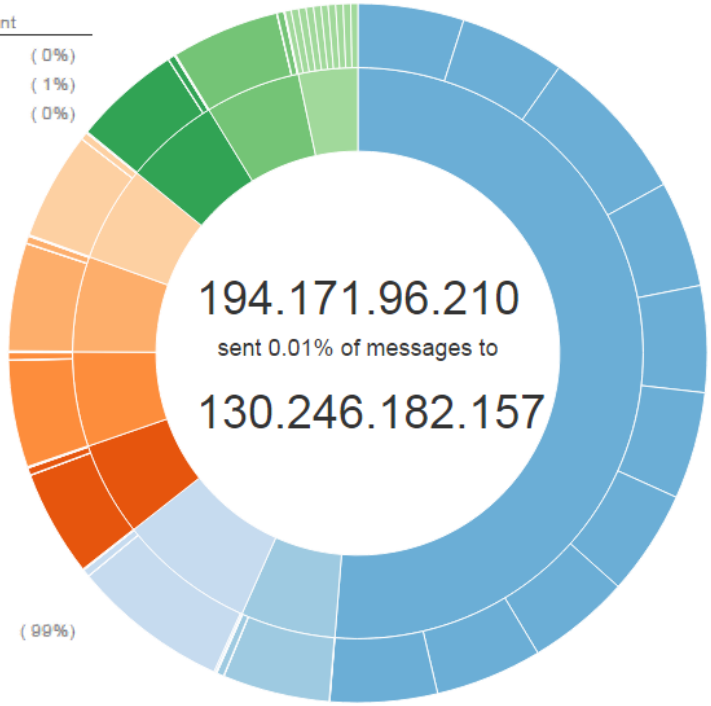

Data flows: a user at UC Irvine reading a data set from Nikhef, with some background from CERN and KIT Karlsruhe

Graphics courtesy Jouke Roorda and Olivier Verbeek

Nikhef

Getting more bytes through?

- Power vs x64: more PCI lanes & higher clock should give more throughput – *if all the bits fit together*
- Only way to find out is … by trying it!
  joint experiment with Nikhef and SURFsara
  on comparing IO throughput between x86 & P8

HGST: 480 TByte gross capacity/4RU

yet more is needed

- RAID card are now a performance bottleneck
- JBOD changes CPU-disk ratio
- closer integration of networking to get >100Gbps

Matching systems architecture

- Most applications using x86 today, and probably will for a long time
- alternatives (GPGPU or Power) not quite viable … although for 'dedicated farms' FPGAs help, and KNH works better (we need the memory)

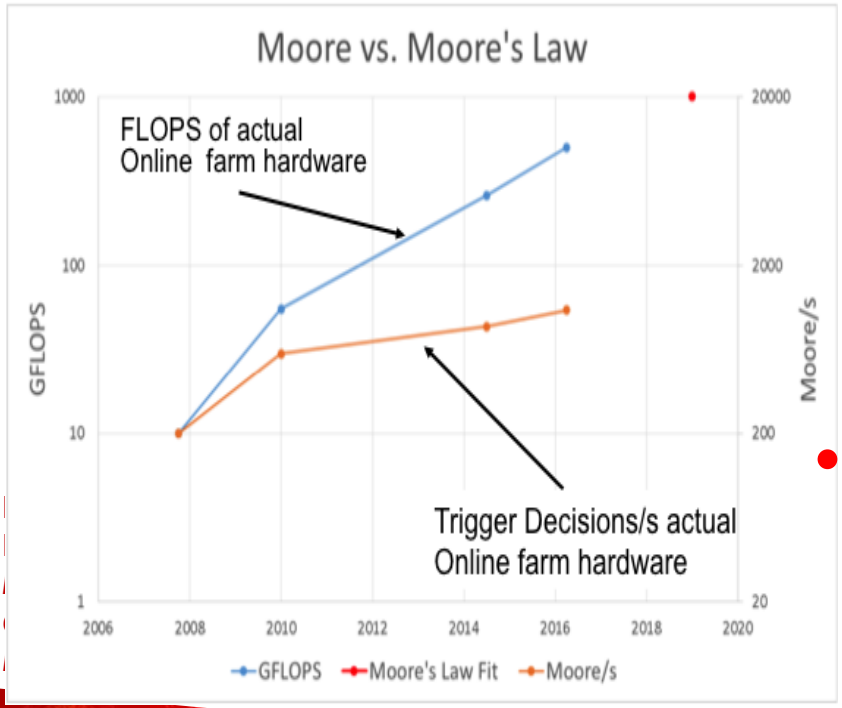

sales volume of different architectures

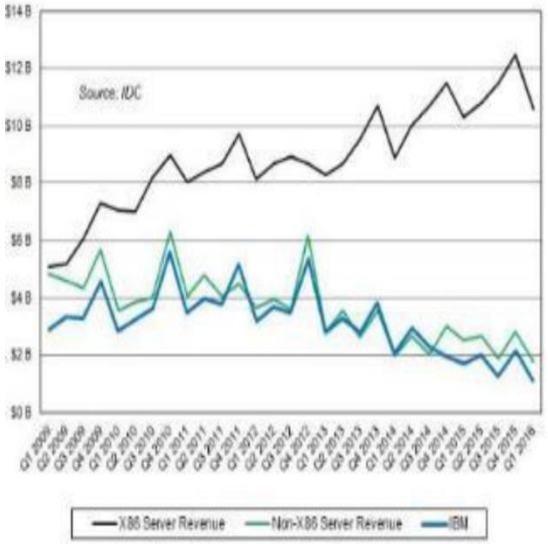

- Yet change must be: most gain to be had from SIMD vectorization and improved memory access patterns

Improvements at the application layer

- ‘traditional’ (1990’s) style HEP applications were ‘lean’, and fail to scale even in pipelining
- let alone vector instructions or multicore

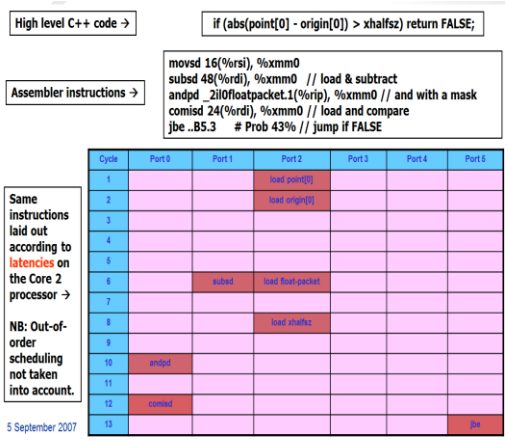

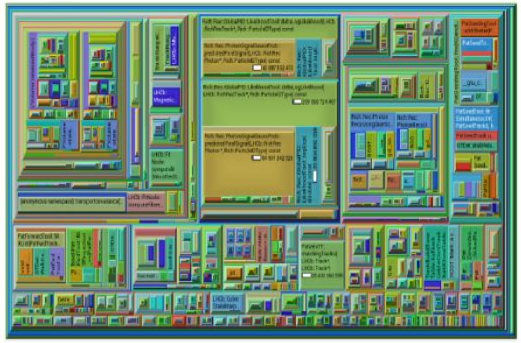

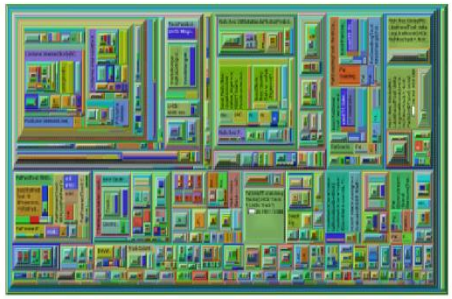

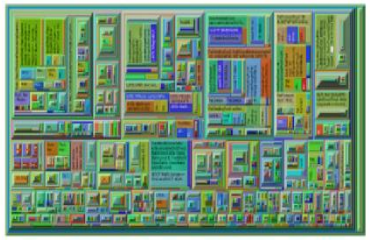

2012 v45r1 v48r1 v48r1 (2015 reco)

review of algorithms gave overall +34% in LHCb – memory layout still to be done …

To use current processor generations, you need better – machine-aware! – code

2007 Core 2 efficientcy: Sverre Jarp, CHEP2007
Performance data: M Schiller (CERN) et al. for LHCb, - profiling progress (update) 2015;; work by: Gerhard Raven, Nikhef

Systems architecture and your application

- Many things you only find in production …
  - When you're 'embarrassingly parallel' with a memory challenge *why not try 'priming' of memory for the first few events and then fork?*

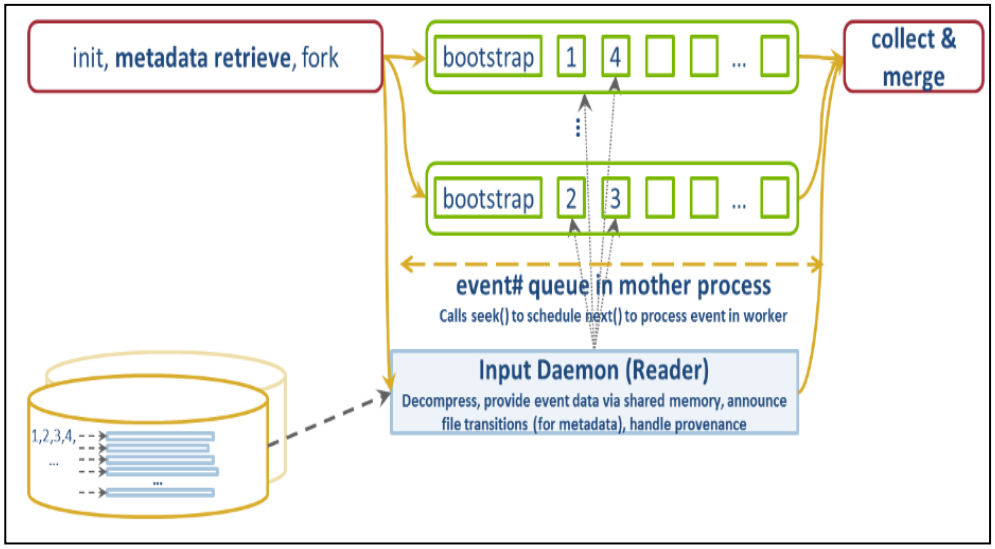

- Towards single-socket systems: cache coherence limits performance – there's a penalty to pay for massive multi-socket-big-memory hosts!

Graphic: P van Gemmeren et al. ATLAS collaboration, ATL-SOFT-PROC-2015-020

Systems for Research @ Nikhef

David Groep
Nikhef
PDP - Advanced Computing for Research

Nikhef

Statistics

BiG Grid
the dutch e-science grid

Dutch National e-Infrastructure coordinated by SURF

"BiG Grid" HTC and storage platform services

- 3 core operational sites: SURFsara, Nikhef, RUG-CIT
- 25+ PiB tape, 10+ PiB disk, 12000+ CPU cores

@Nikhef

~ 5500 cores and 3.5 PiB

focus on large/many-core systems

> 45 install flavours (service types)

and a bunch of one-off systems

Shared infrastructure, efficient infrastructure!

- >98% utilisation, >90% efficiency

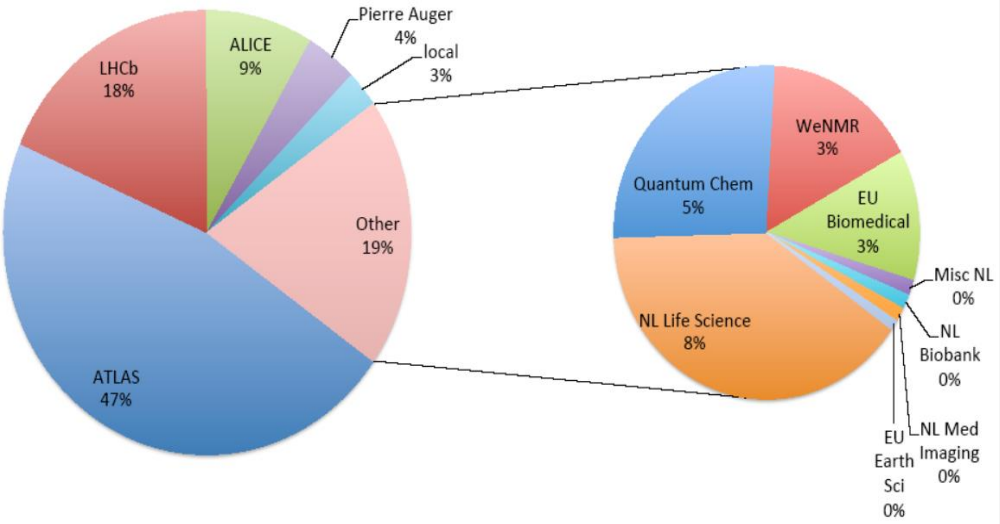

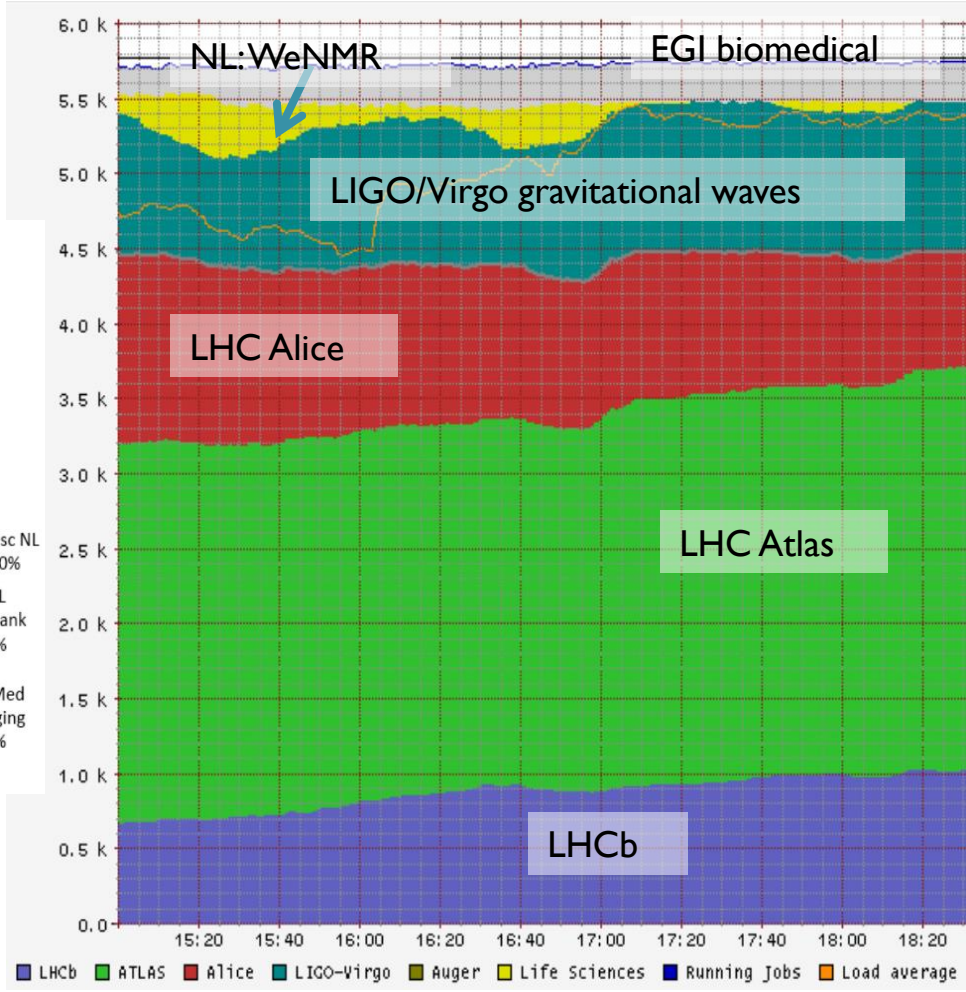

David Groep
Nikhef
PDP - Advanced Computing for Research

Right:: NIKHEF-ELPROD facility, Friday, Dec 9th, 2016
Left: annual usage distribution 2013-2014

Federation of high-throughput services

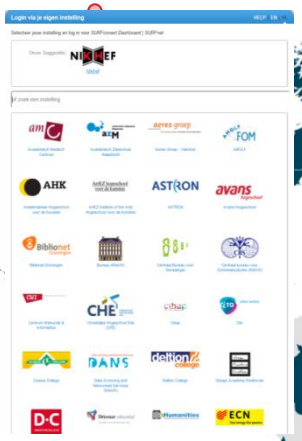

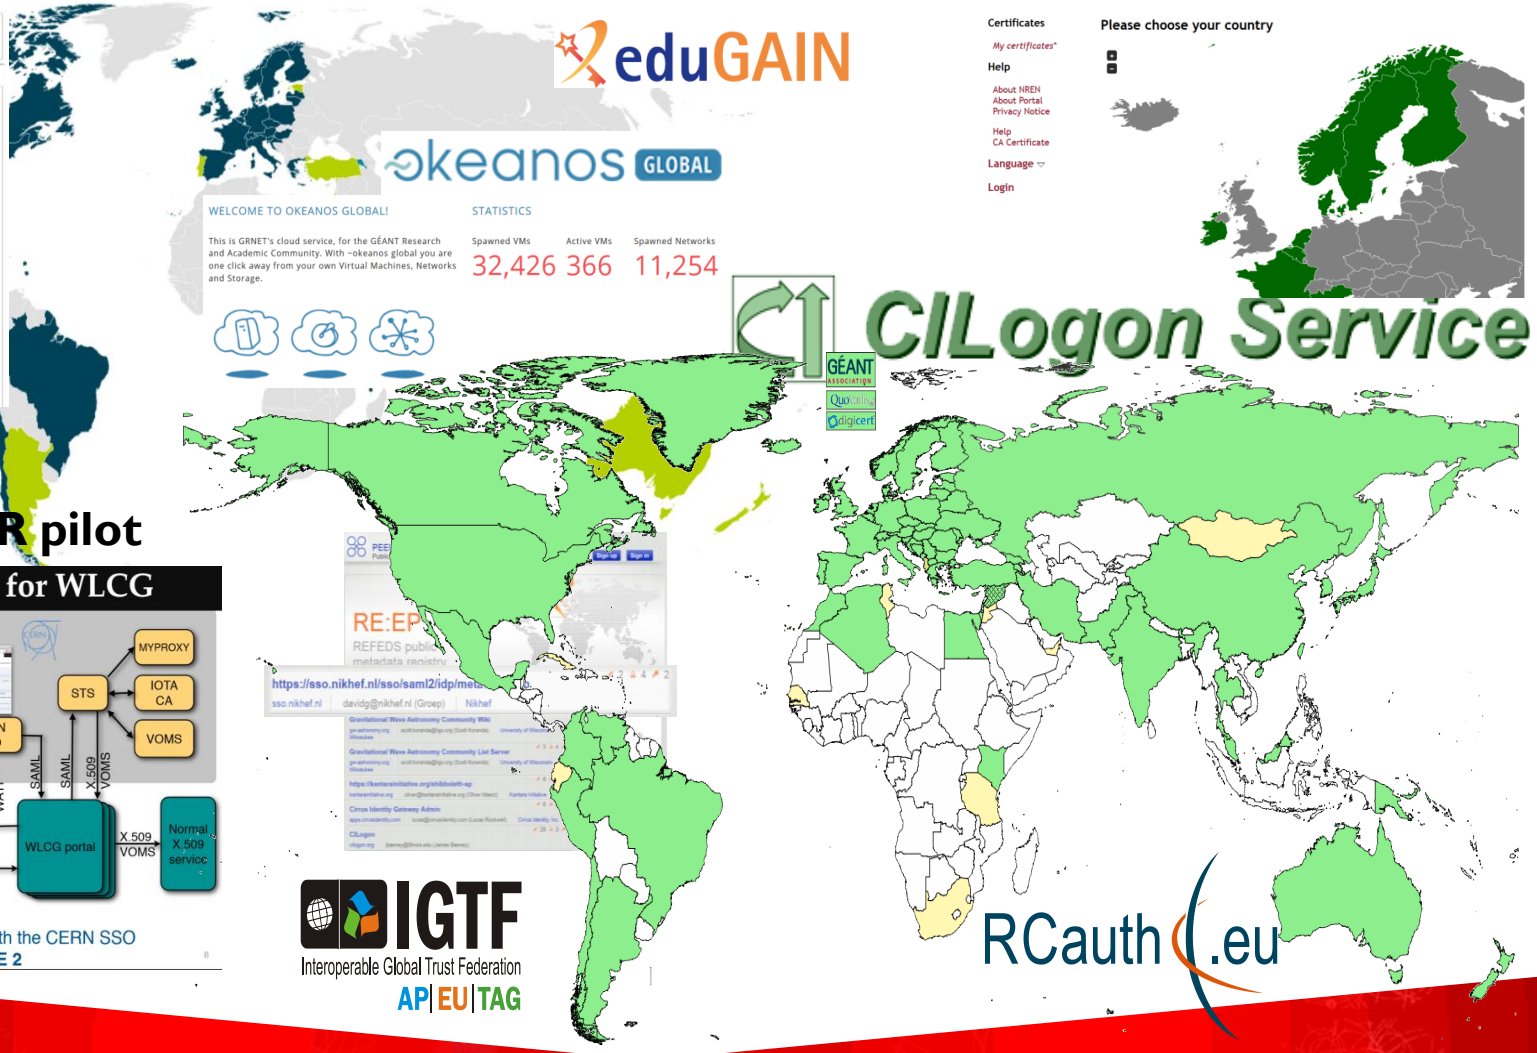

wLCG FIM4R pilot

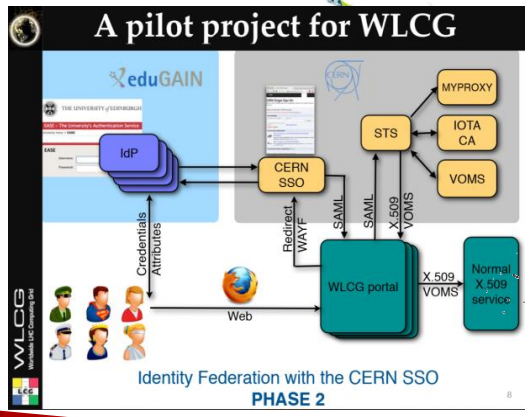

Nikhef

‘cloud’ is a means, not an end-all solution

/cvmfs/softdrive.nl

Nikhef SURF SARA

E. Tejedor et al., CERN, SWAN Service for Web-based Analtsis, CHEP 2016

- Docker: single thin image, not managed by the user!
- CVMFS: configurable environment via “views”
- CERNBox: custom user environment

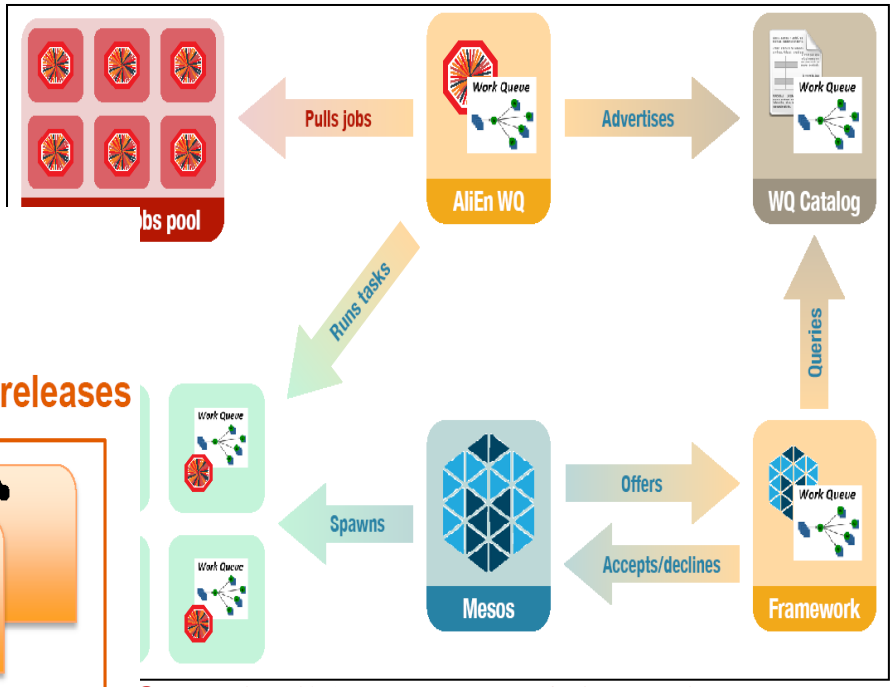

@cern.ch - CHEP 2016 - Experiences with the ALICE Mesos infrastructure

docker + C

CernVM File system

LCG releases

CERN software

+ CERNBox User software

Collaborative advantage:
joint effort of infrastructure and users

http://doc.grid.surfsara.nl/en/latest/Pages/Advanced/grid_software.html#softdrive

For (informed) fun & testing – some random one-off systems …

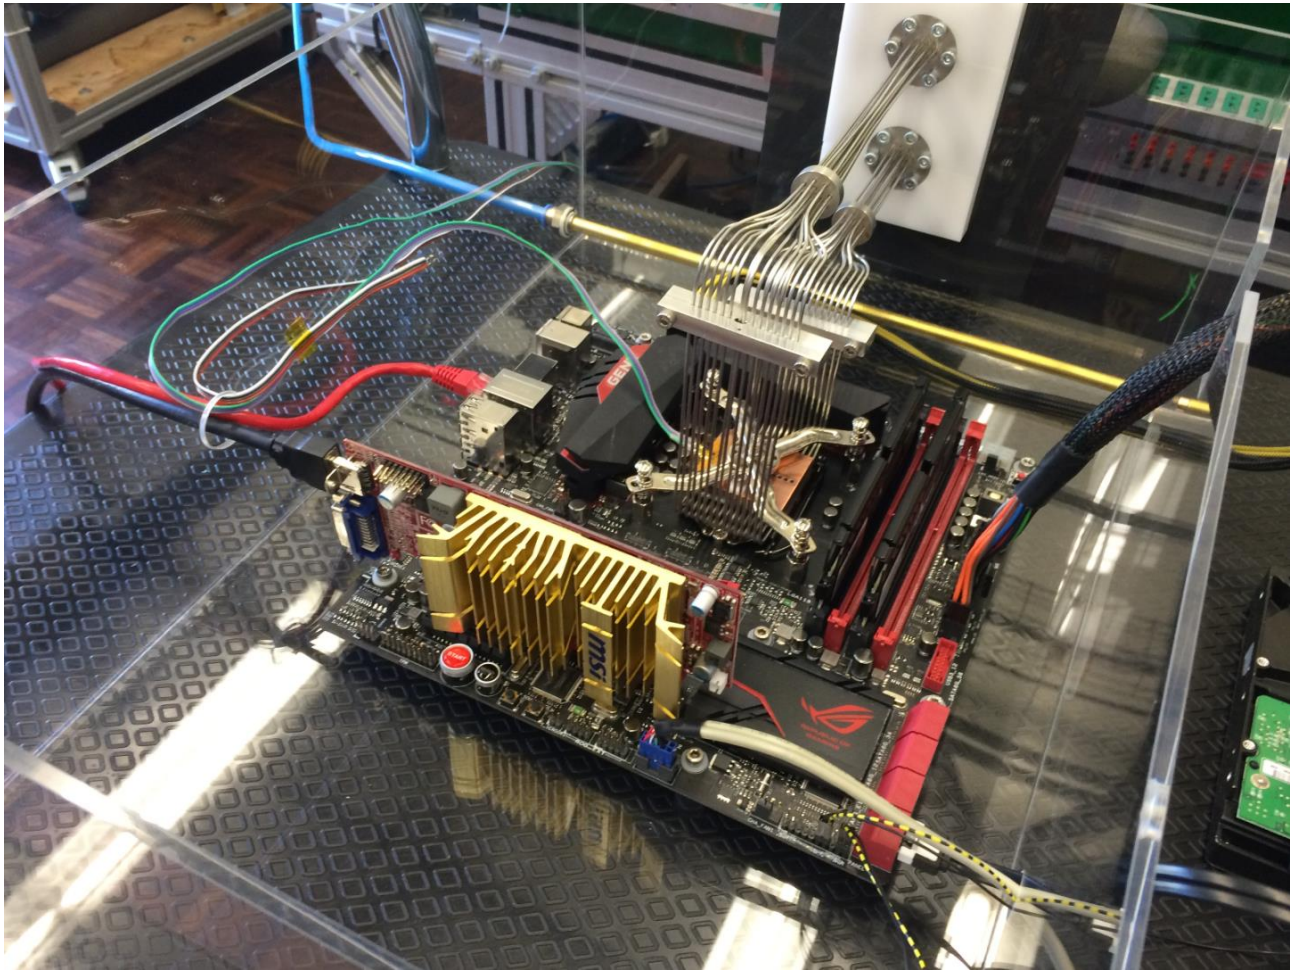

David Groep
Nikhef
PDP - Advanced Computing for Research

CO2-cooled Intel CPUs @6.2GHz

For (informed) fun & testing – some random one-off systems …

David Groep
Nikhef
PDP - Advanced Computing for Research

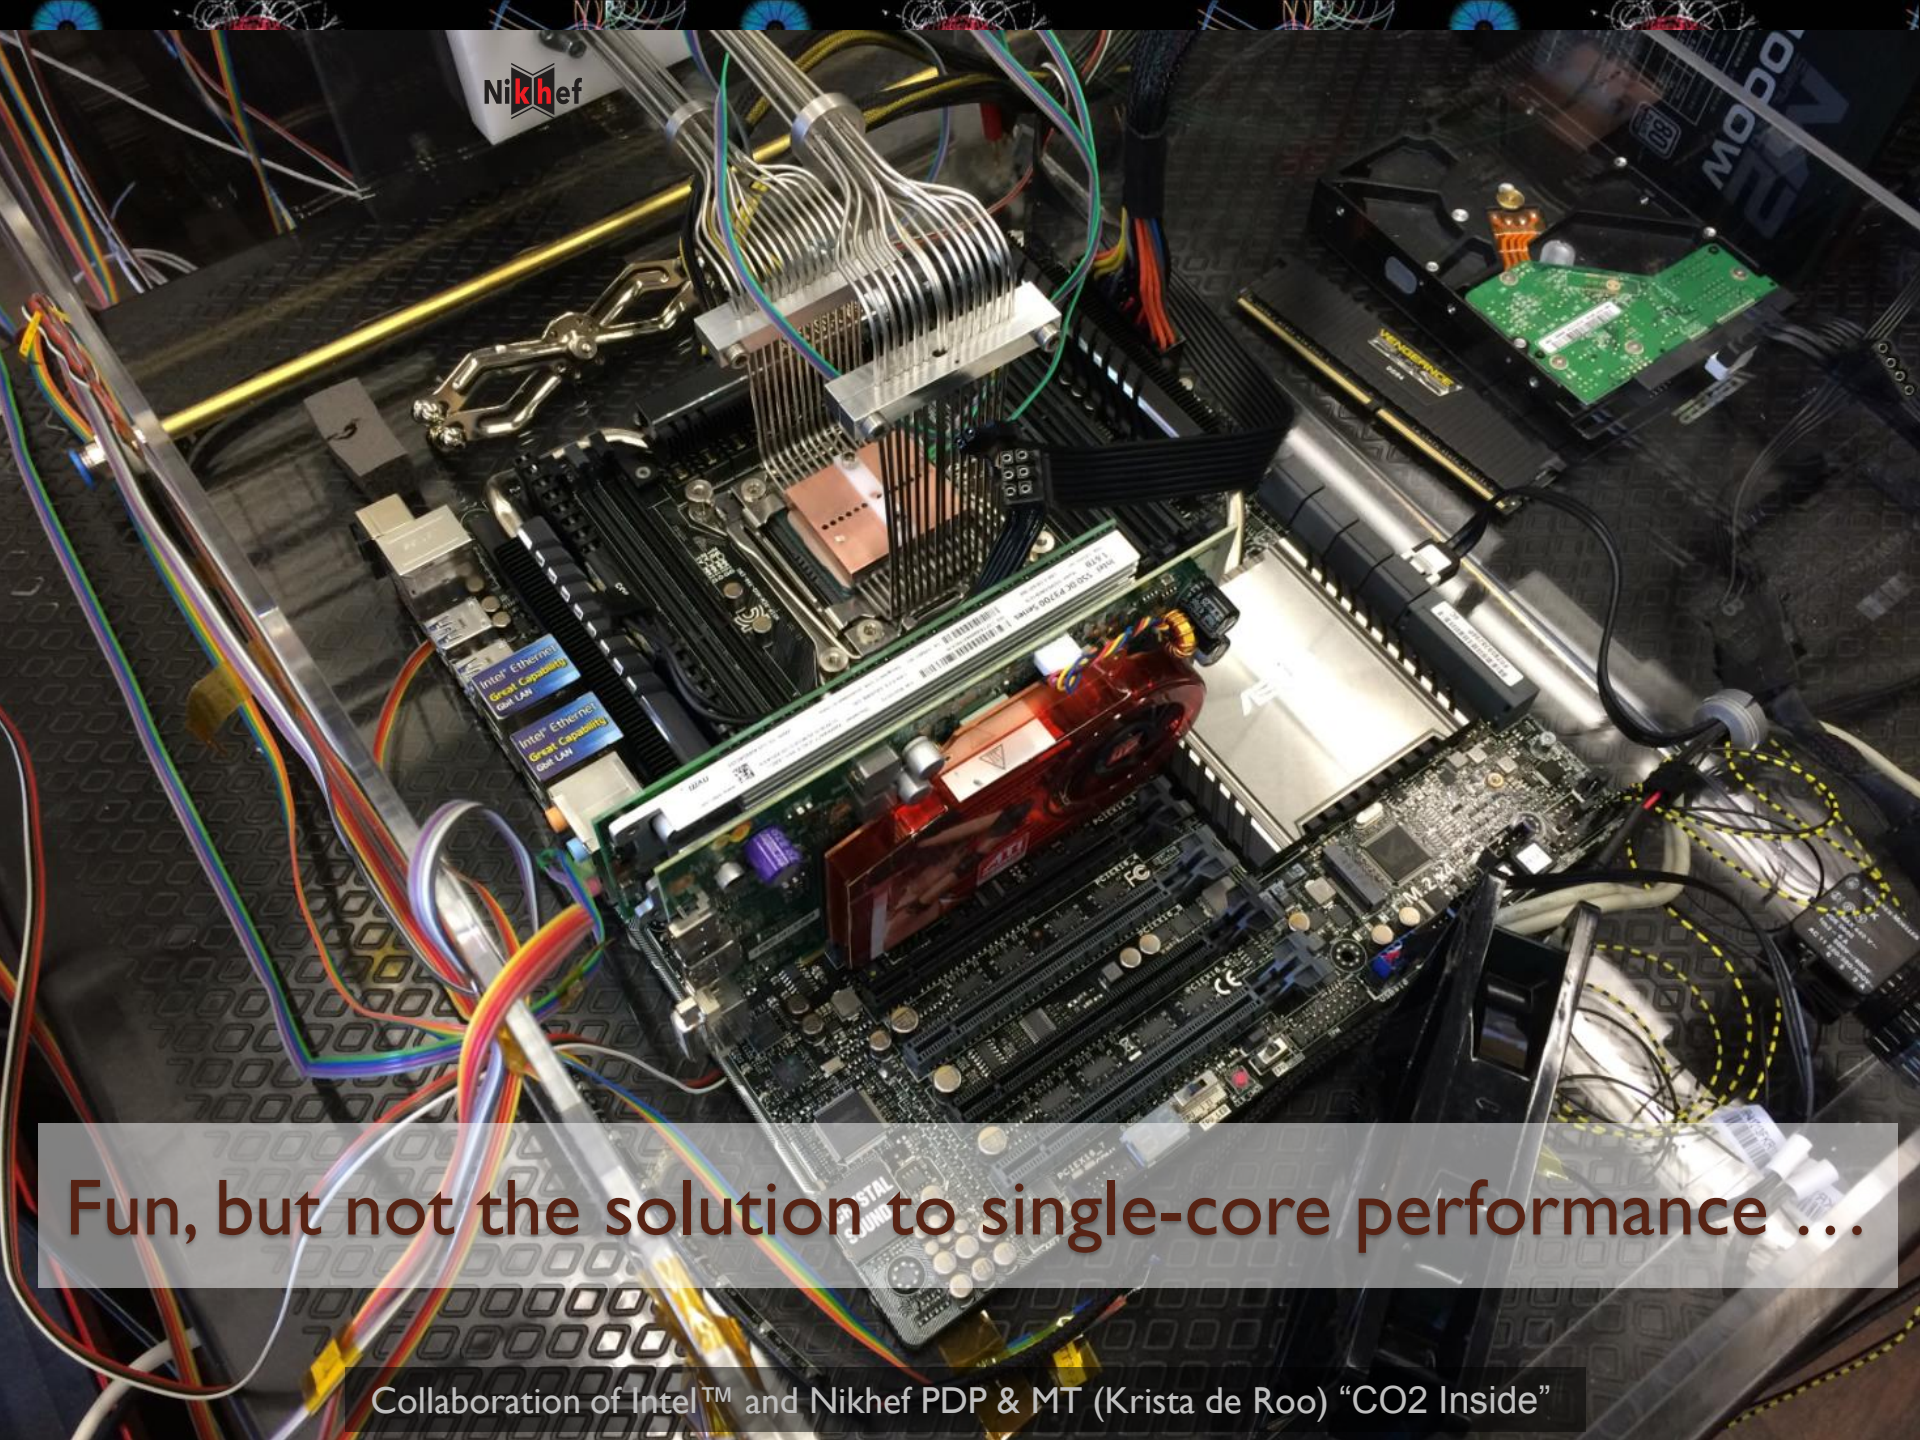

Fun, but not the solution to single-core performance …

Collaboration of Intel™ and Nikhef PDP & MT (Krista de Roo) “CO2 Inside”