From Femtoscale to Exascale: Computing at CERN

Andrzej Nowak, CERN openlab

The Register & Intel Live 2011 – London, November 22nd

CERN openlab

CERN

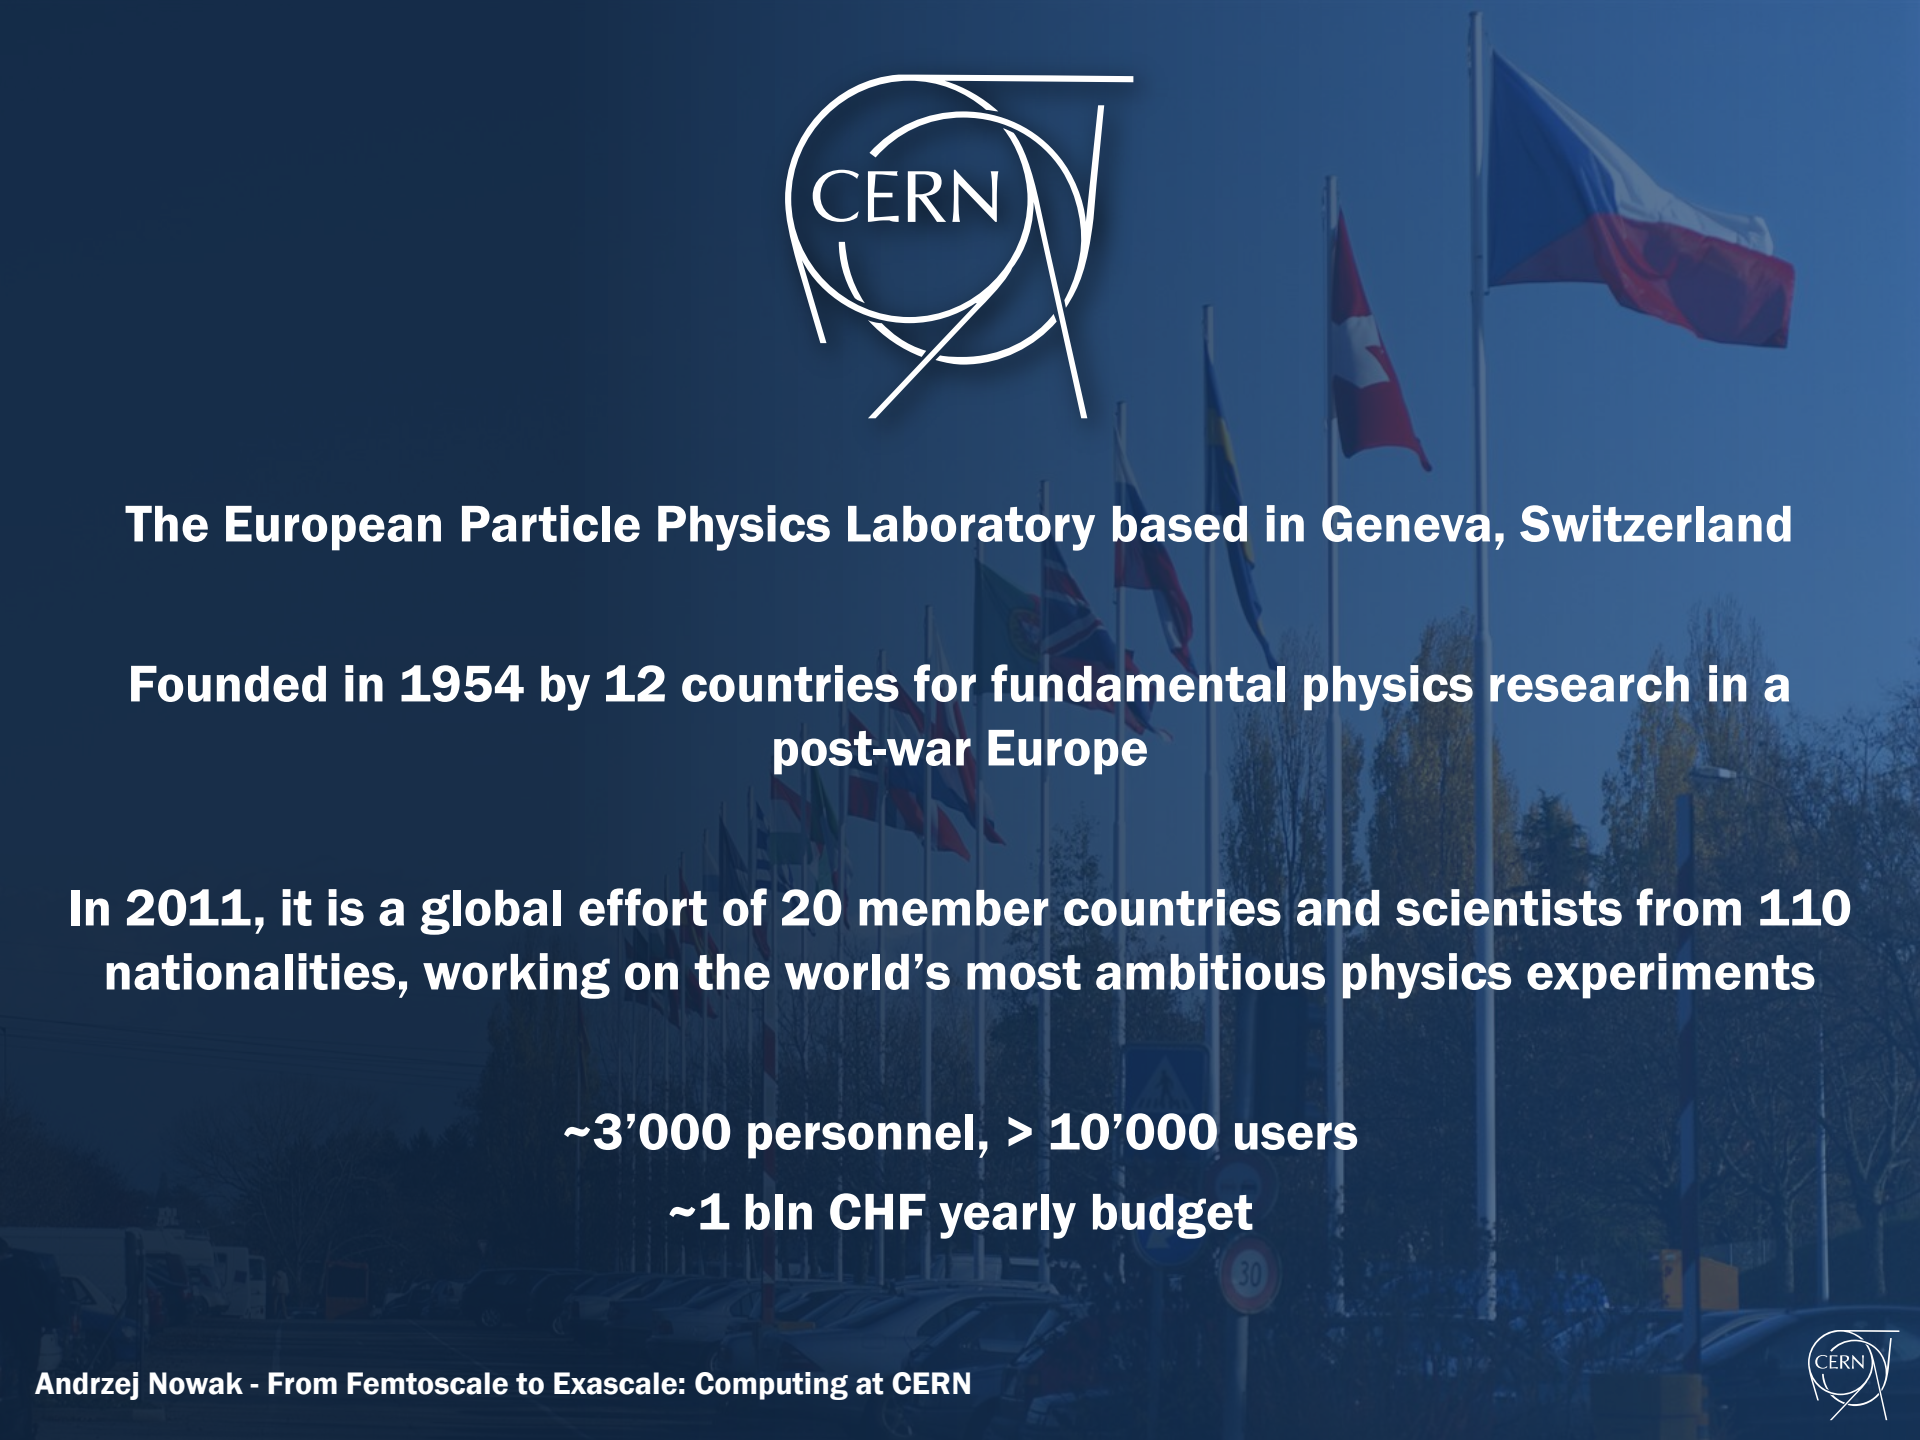

The European Particle Physics Laboratory based in Geneva, Switzerland

Founded in 1954 by 12 countries for fundamental physics research in a post-war Europe

In 2011, it is a global effort of 20 member countries and scientists from 110 nationalities, working on the world's most ambitious physics experiments

~3'000 personnel, > 10'000 users

~1 bln CHF yearly budget

The Large Hadron Collider

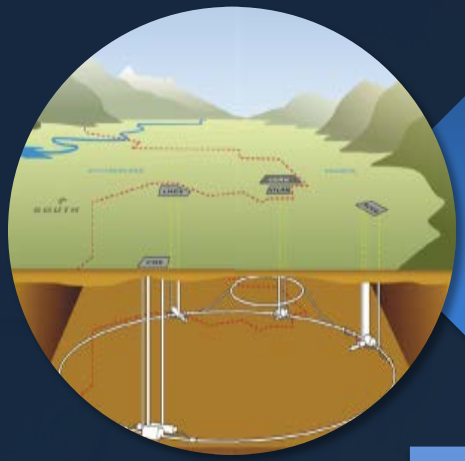

27 km underground superconducting ring – possibly the largest machine ever built by man

40 million collisions per second

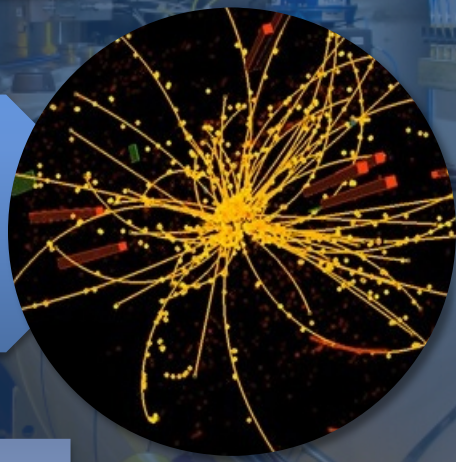

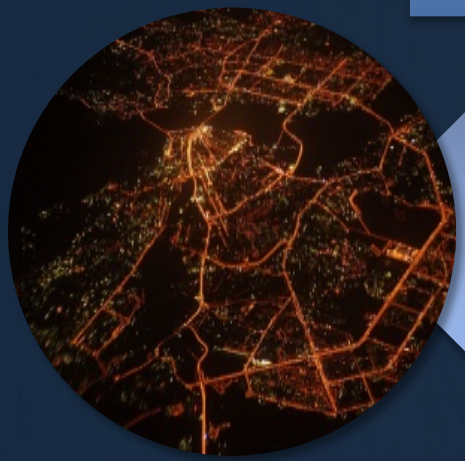

150-200 MW power consumption


Mont Blanc (4,808m)
Geneva (pop. 190'000)
Lake Geneva (310m deep)
LHCb
ATLAS
CERN Meyrin
CERN Prévessin
SPS 7 km
SUISSE
FRANCE
CMS
ALICE
LHC 27 km

CERN's accelerator complex

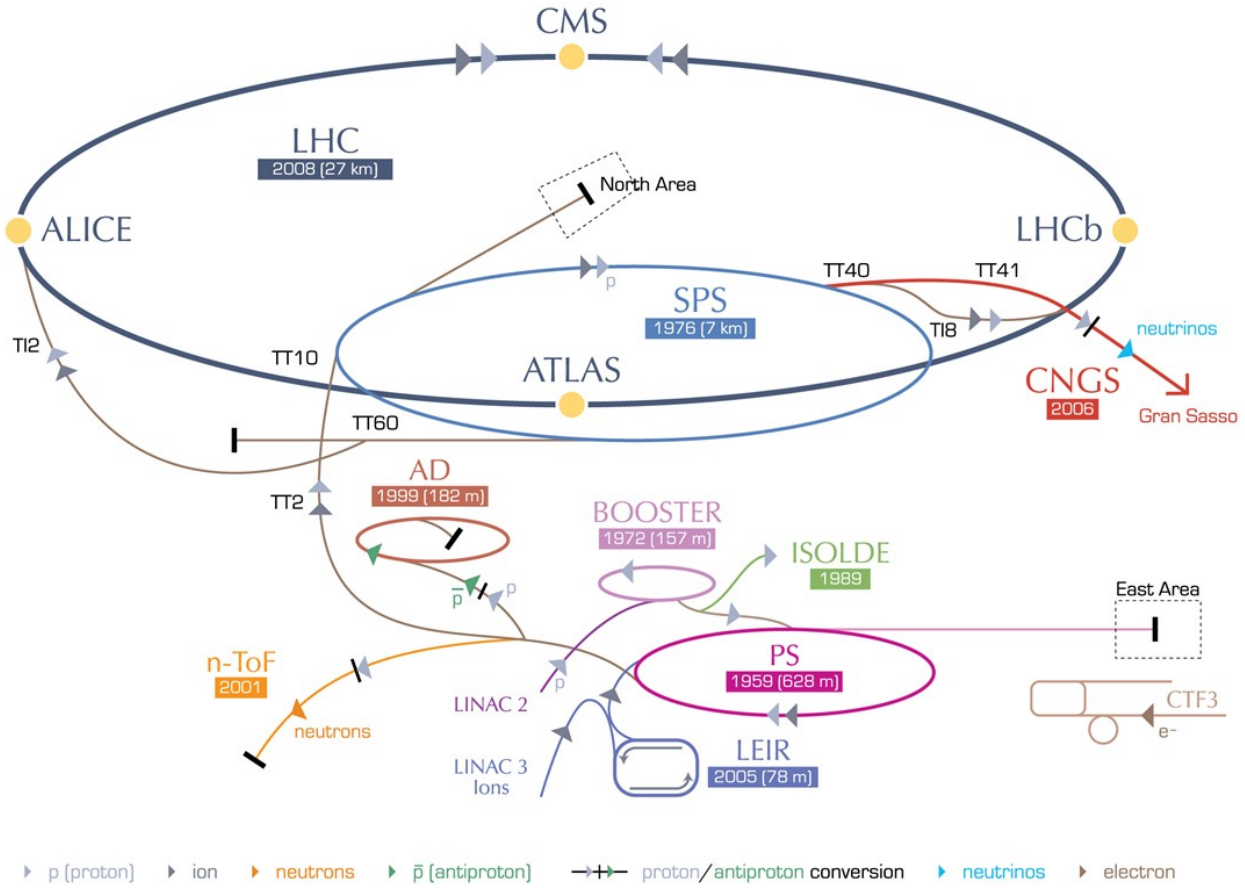

Intense data pressure creates strong demand for computing

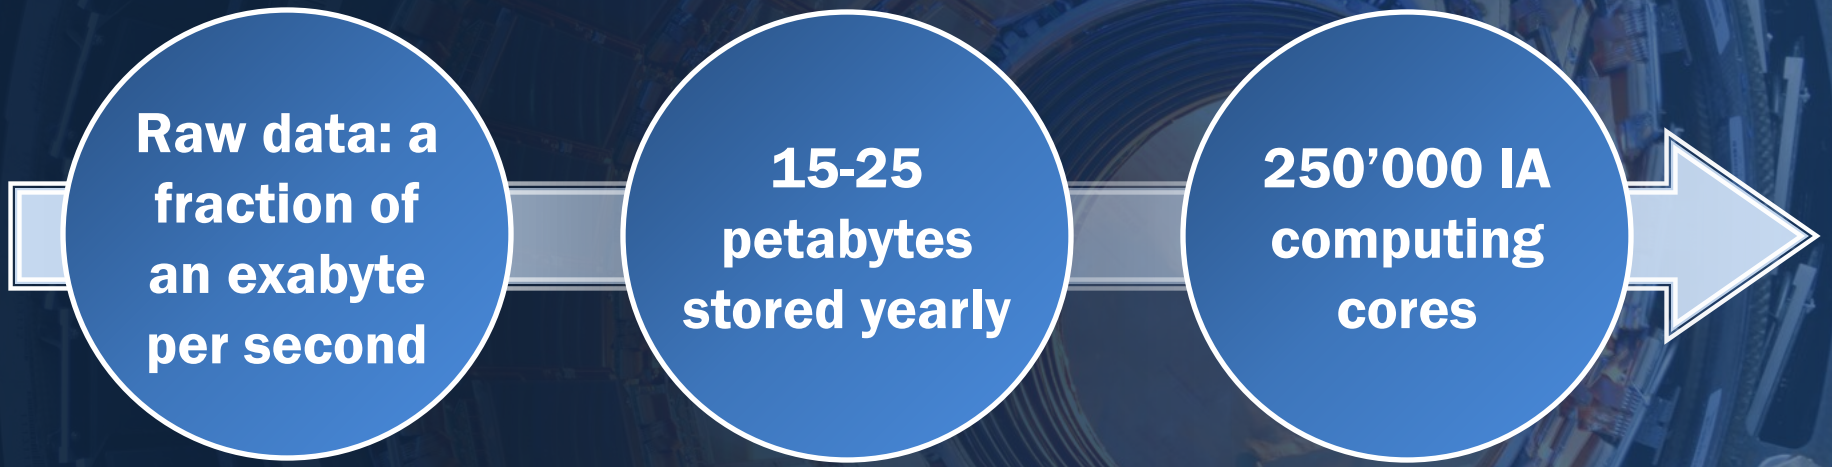

A rigorous selection process enables us to find that one interesting event in 10 trillion (10^{13})

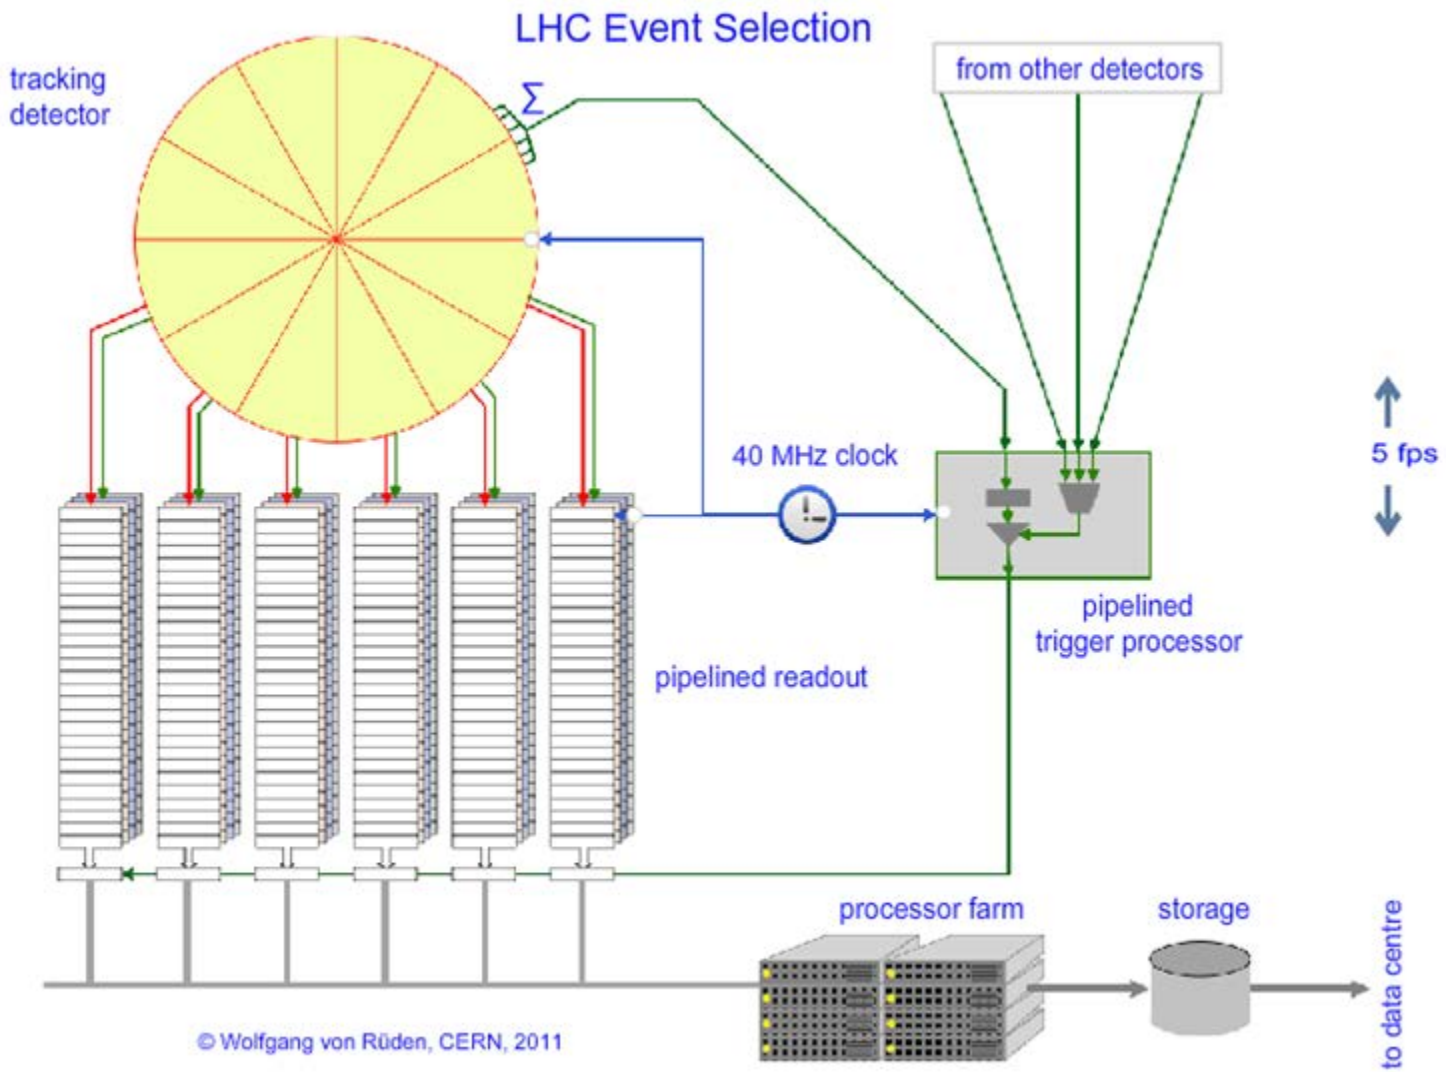

Triggering in underground computing farms

Data flow from the LHC detectors

Online triggering and filtering in detectors

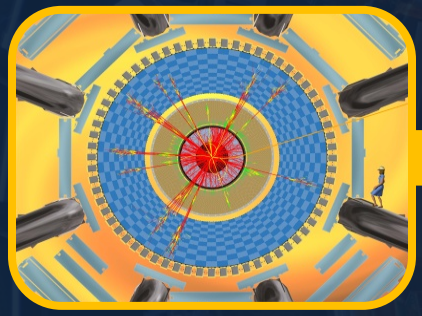

Event simulation

Reconstruction

Selection and reconstruction

Raw Data (100%)

Event reprocessing

Event summary data (10%)

Analysis

Analysis objects (1%)

Batch physics analysis

Processed data

Intel Processor features

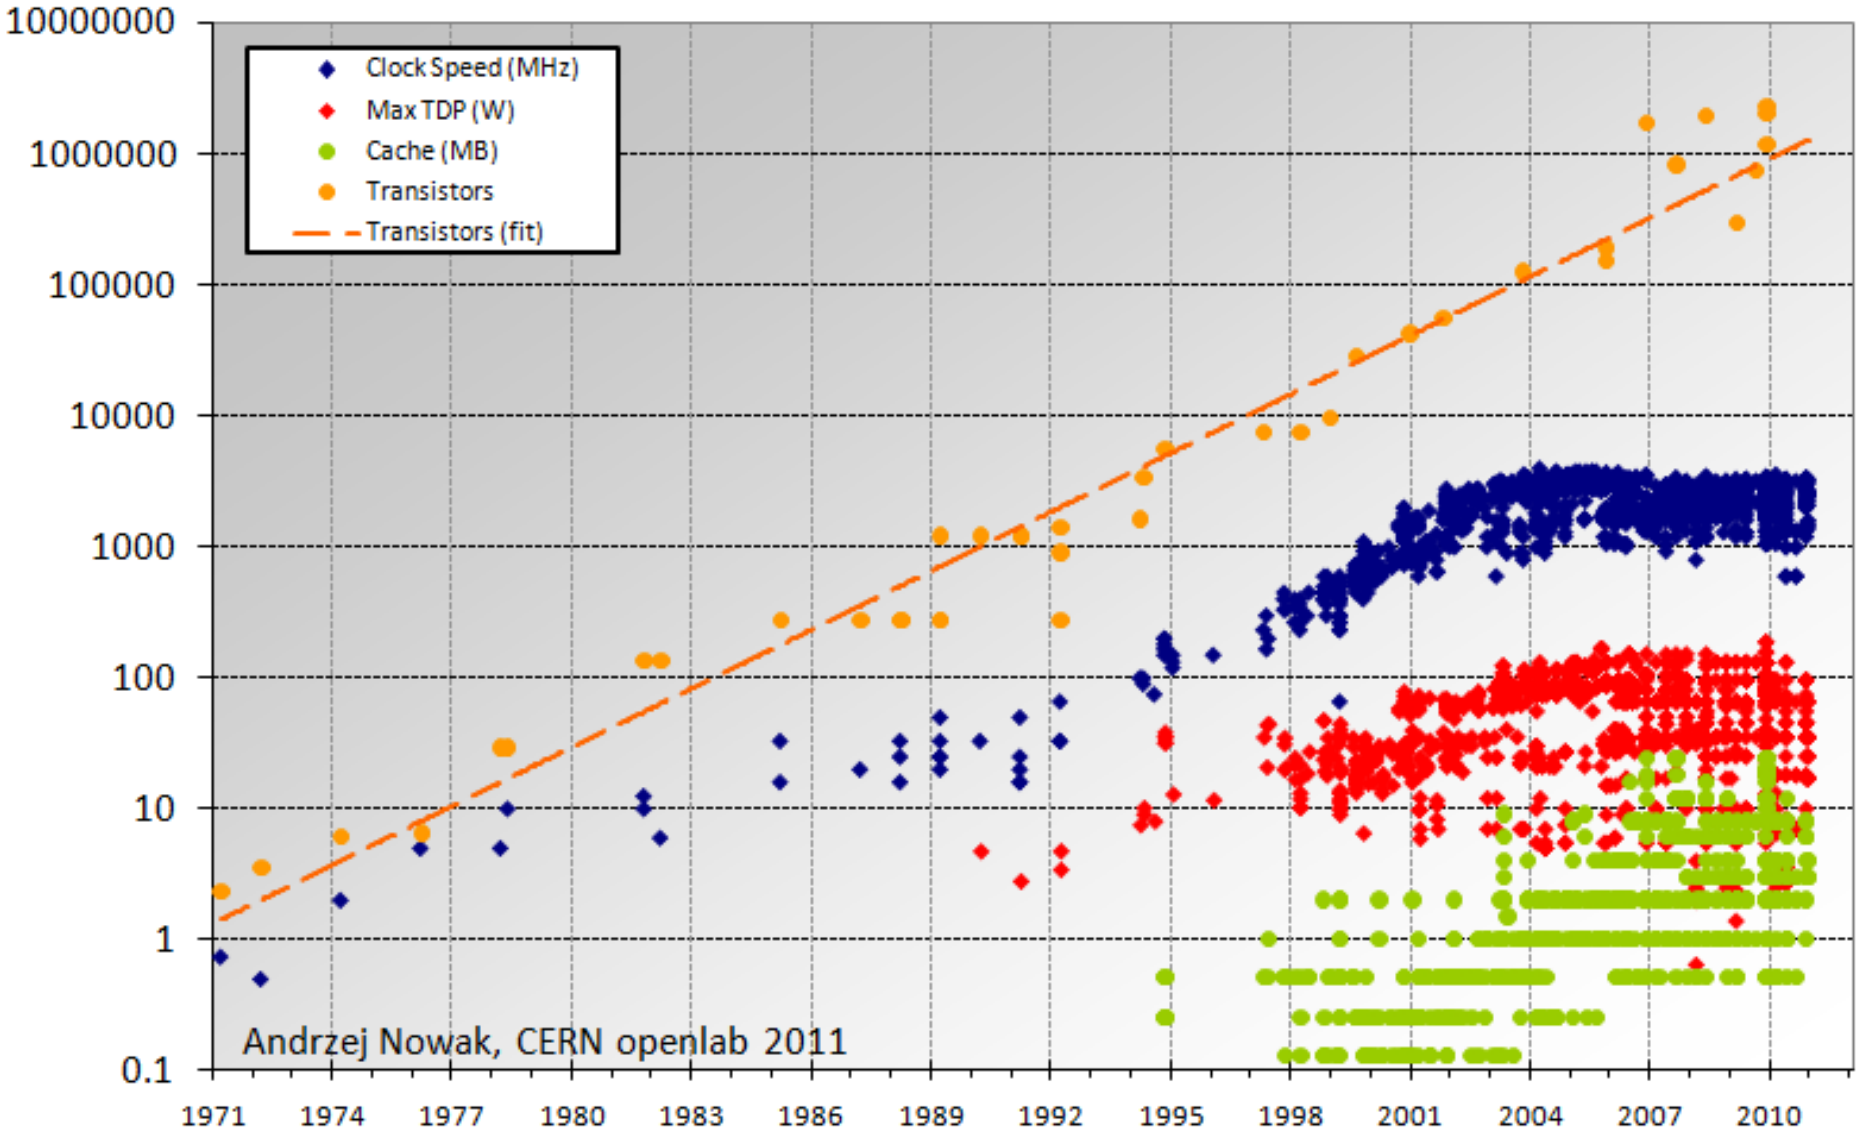

Digression: CERN software in general terms

Millions of lines of code

Key foundation: Linux + GCC

Compute scales with a combination of SPECFP and SPECINT

Independent, parallel events

Little or no HPC needed. Throughput is king.

Large aggregate requirements, but chaotic workload

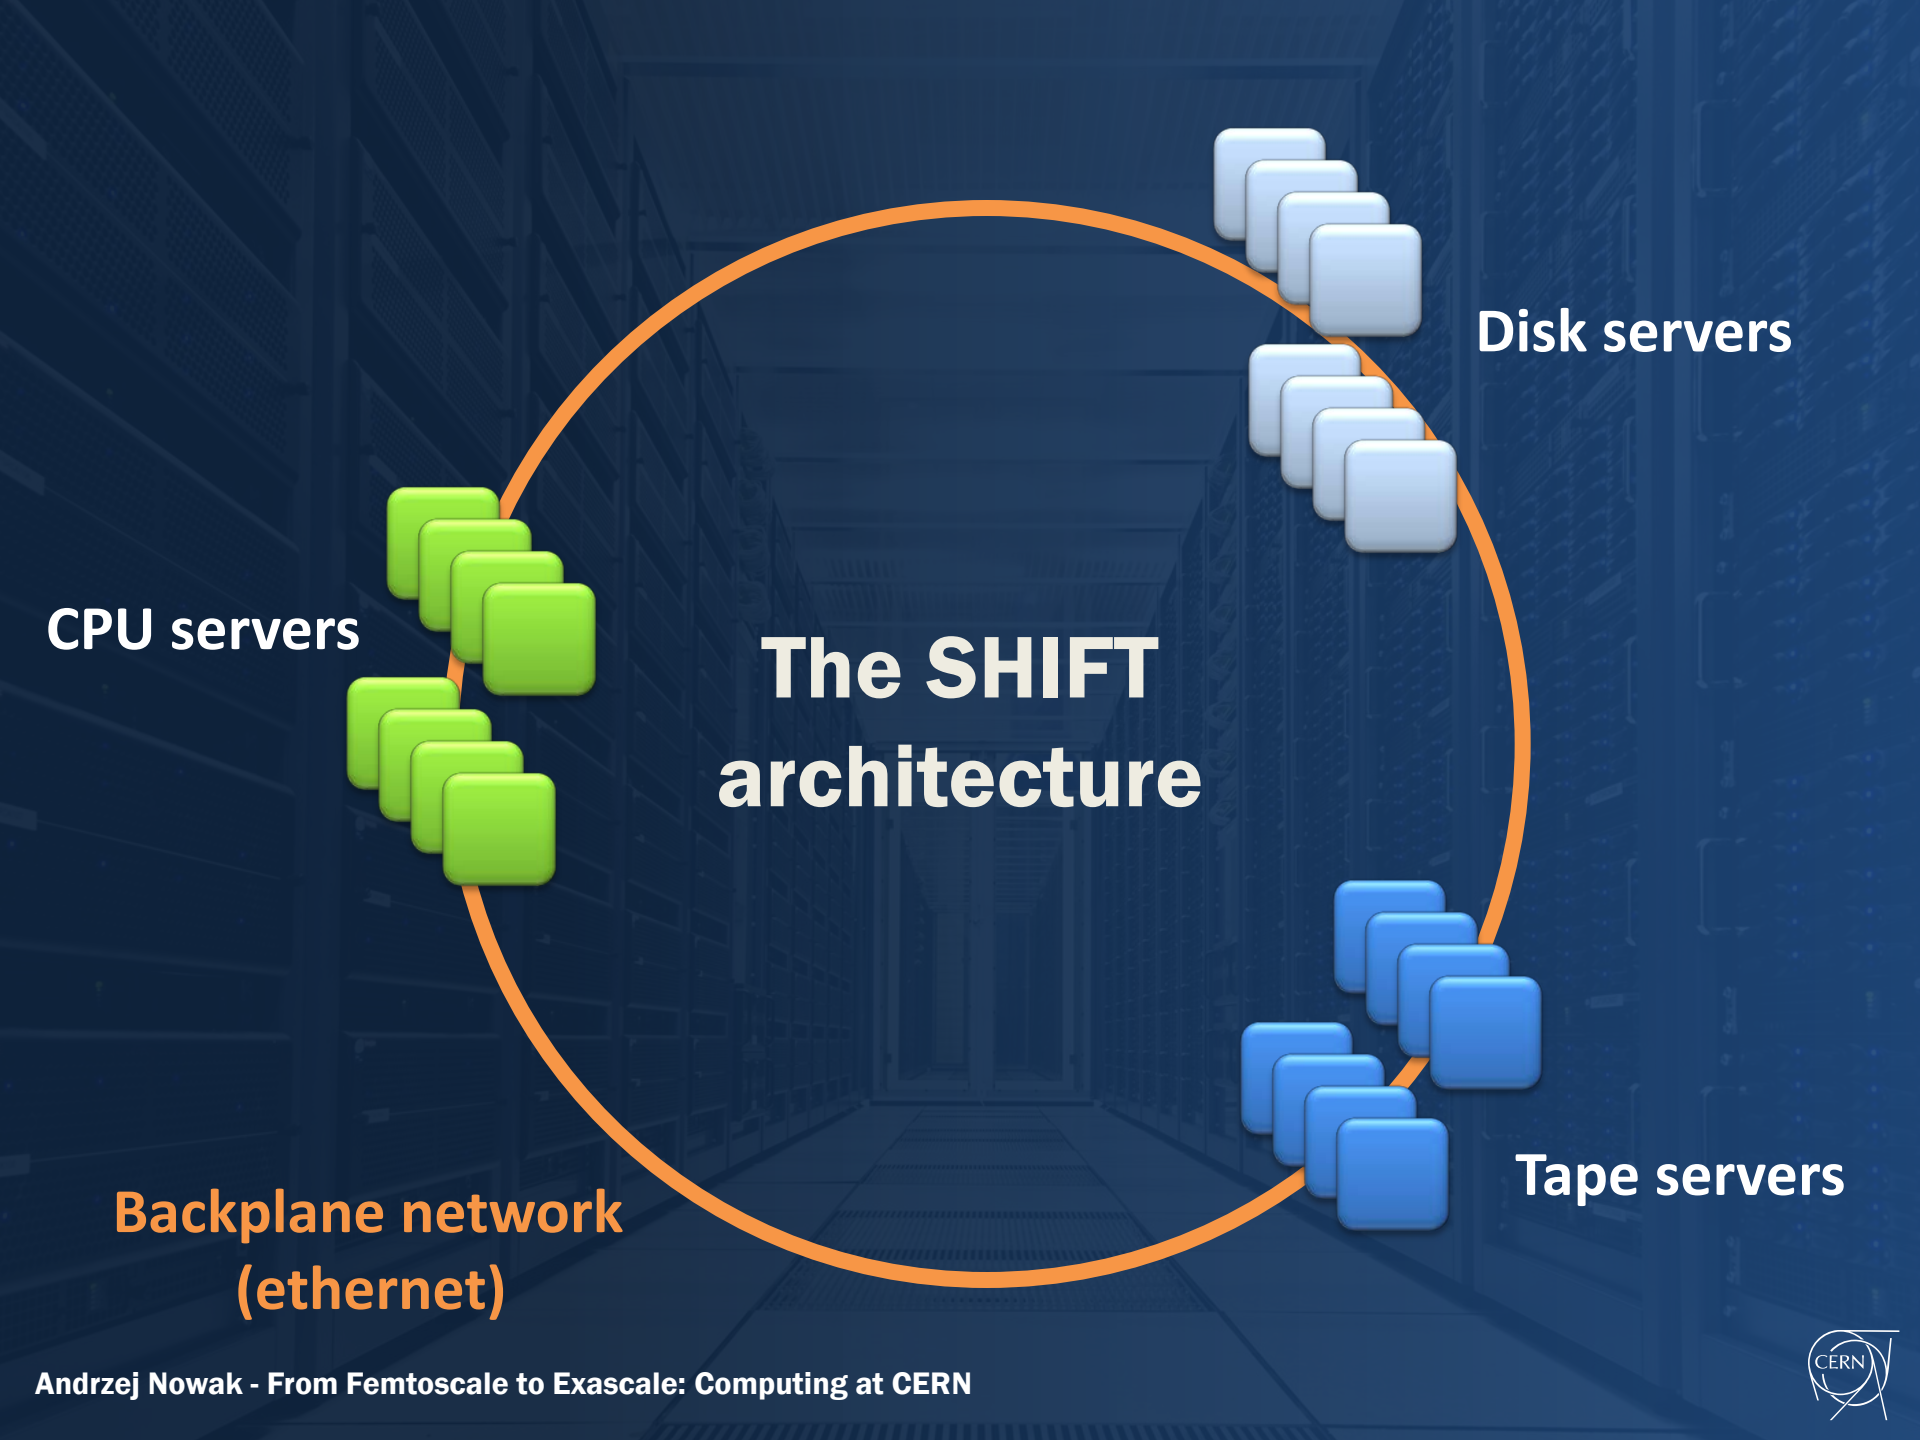
Disk servers
CPU servers
The SHIFT architecture
Tape servers
Backplane network (ethernet)

Central Tier-0 computing farm

Main characteristics

- Retrofitted mainframe building from the 70s
- 2 floors

Approximate figures

- 8’500 systems
- 64’000 cores
- 62 petabytes of disks
- 70 petabytes of tape
  - 80 if compressed
- 2.9 MW limit for power consumption (3.5 soon)
- 4.8 Tbit switching capacity

Commodity components

Exception: Electronics	• Own designs for online data acquisition in caverns • Custom networks, boards, thousands of FPGAs
Cpu server	• Periodic procurements (latest x86 generation) • 2GB memory / core
Disk server	• Simple processor • A couple dozen drives
Tape server and tape	• Tape servers linked to robots and libraries • Tape is our key means of reliable storage
Network	• Fiberoptics + standard copper ethernet • 10G for backbones, 1G for clients
Software	• Batch computing: own, Linux, GNU, Oracle • Client computing: Linux, MS Windows

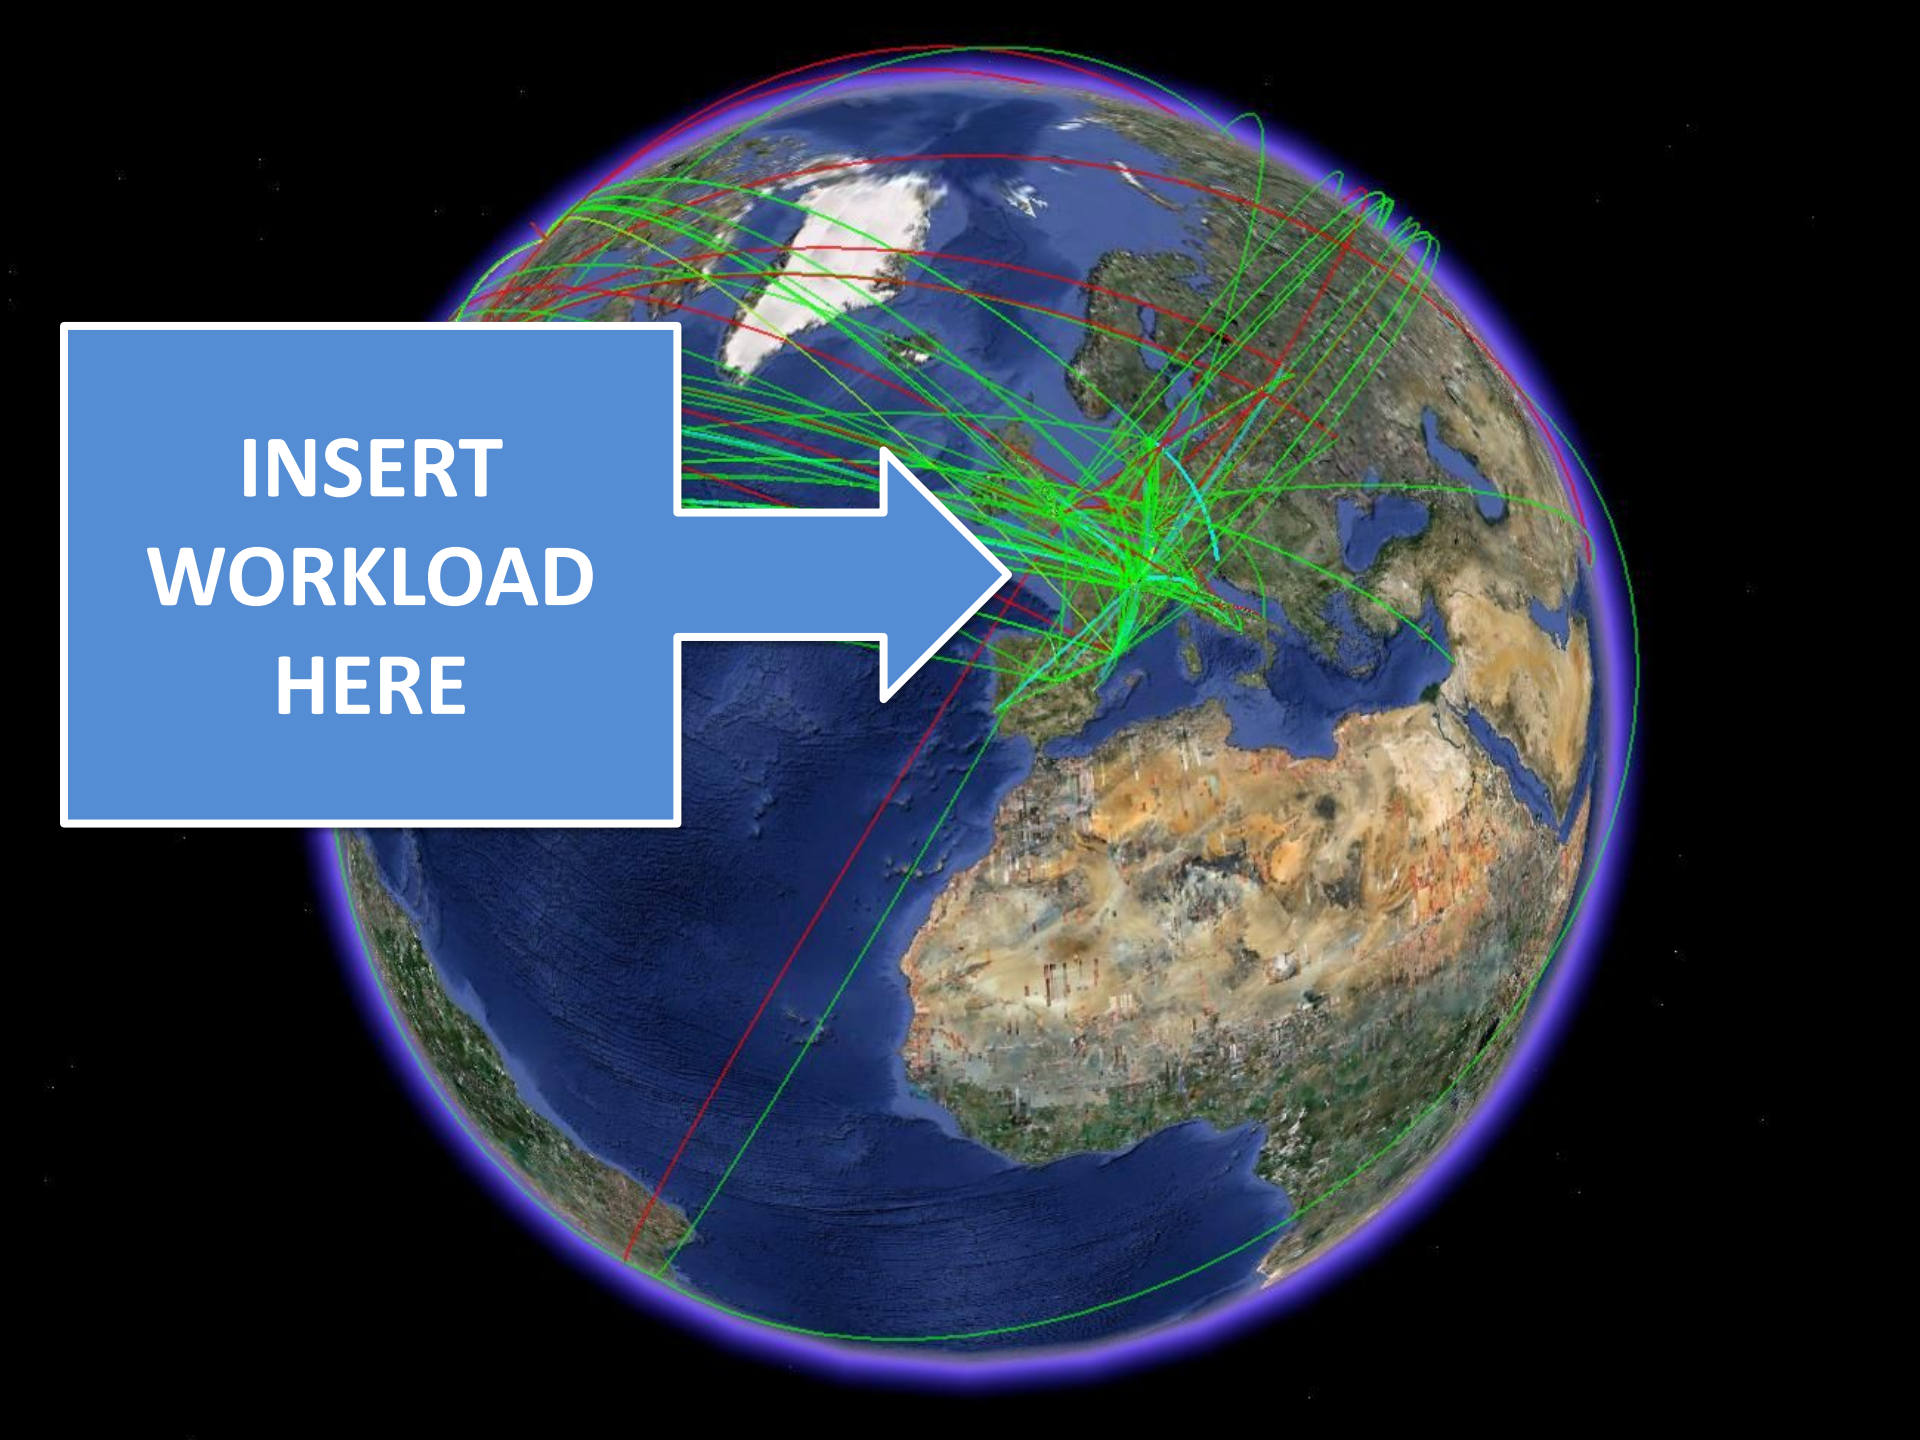
INSERT
WORKLOAD
HERE

The Worldwide LHC Computing Grid

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis

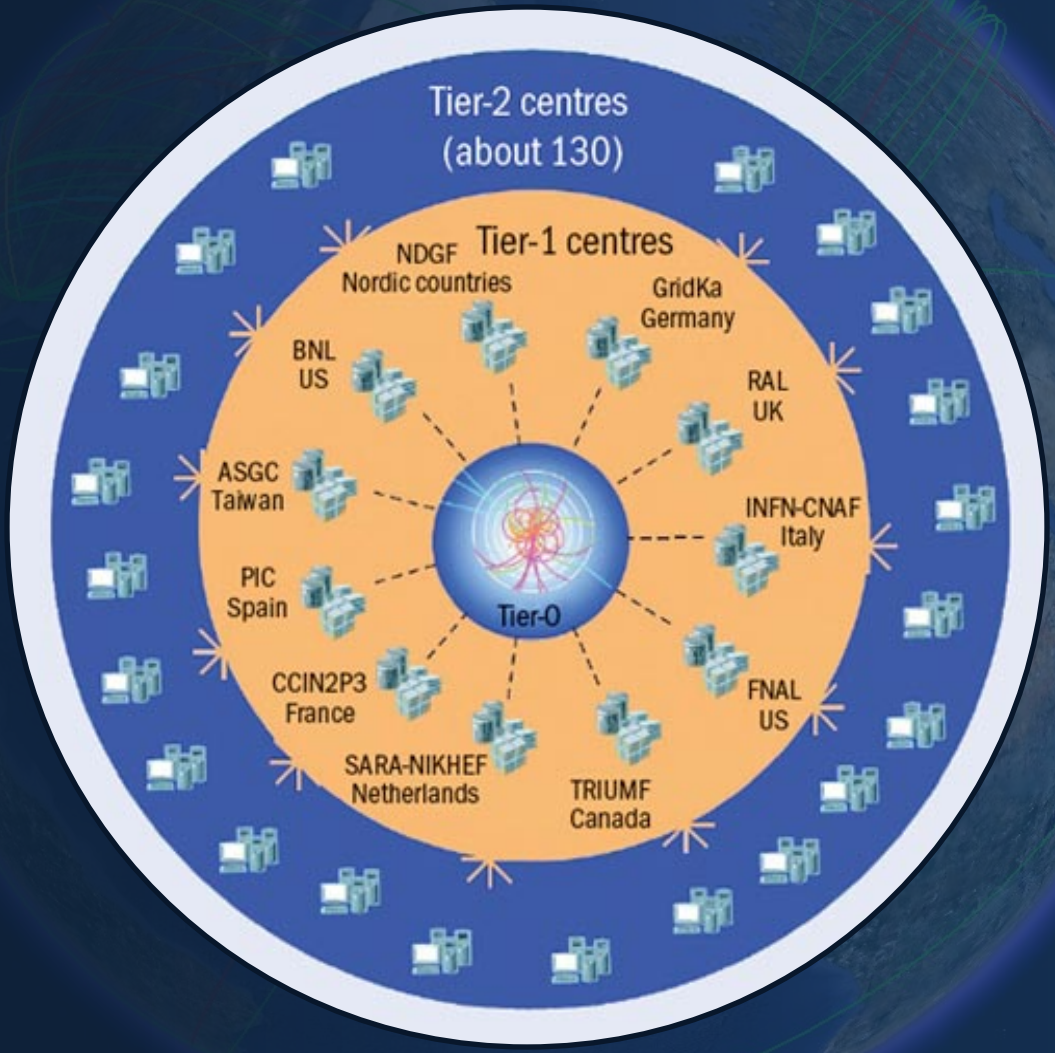

nearly 160 sites

~250'000 cores

173 PB of storage

> 1 million jobs/day

High throughput networking for the LCG

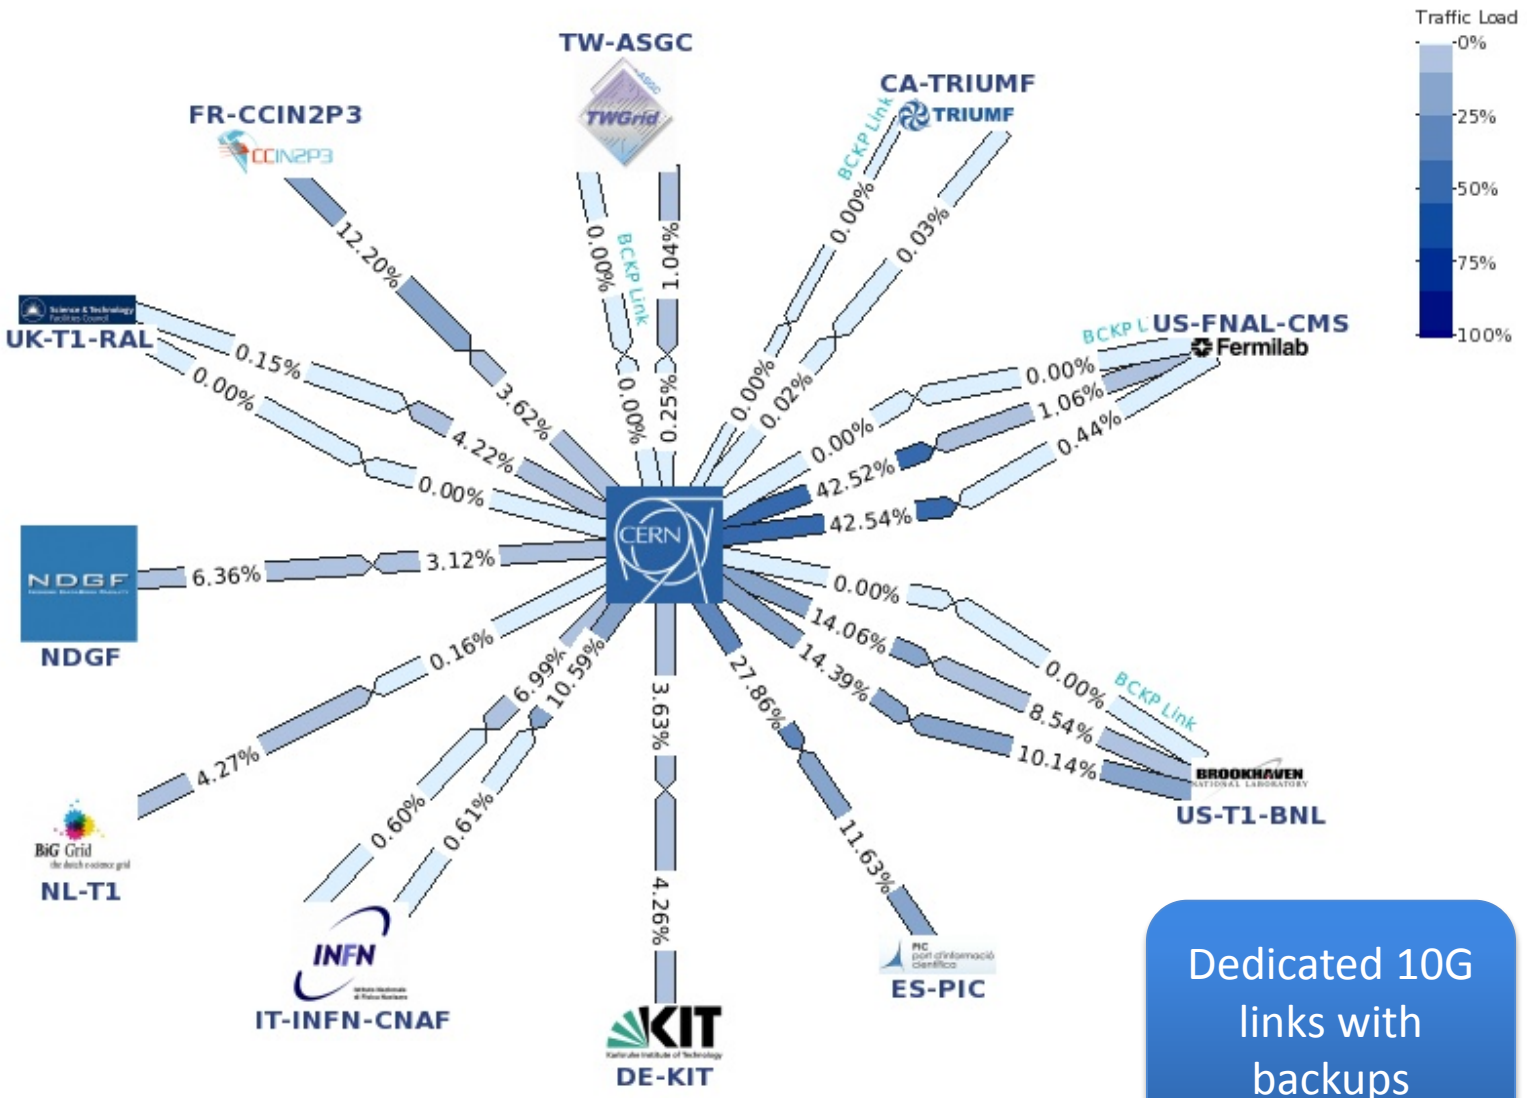

Dedicated 10G links with backups

CERN and the question of the Cloud

Data: persistence, access, jurisdiction

SLA and guarantees

Interoperability

A long track of innovation in computing

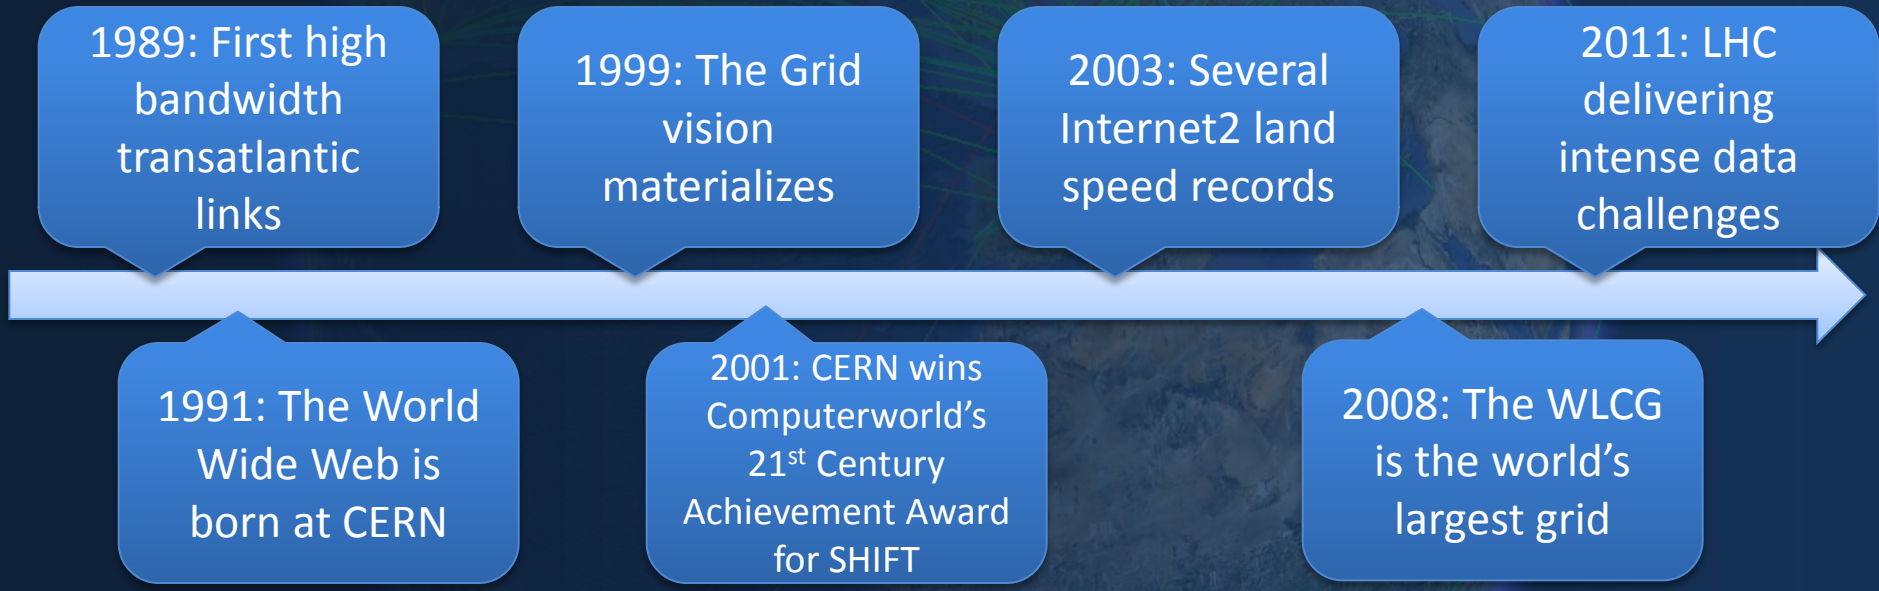

The CERN openlab

A unique research partnership of CERN and the industry
Objective: The advancement of cutting-edge computing solutions to be used by the worldwide LHC community

- Partners support manpower and equipment in dedicated competence centers
- openlab delivers published research and evaluations based on partners' solutions – in a very challenging setting
- Created robust hands-on training program in various computing topics, including international computing schools; Summer Student program
- Past involvement: Enterasys Networks, IBM, Voltaire, F-secure, Stonesoft, EDS; Future involvement: Huawei
- Entering phase IV: 2012-2014

http://cern.ch/openlab

CERN openlab

PARTNERS

hp invent intel ORACLE SIEMENS

Intel has been openlab's long-term partner – continued support since the inception over 10 years ago

Wide range of R&D activities touch on data centers, networking, performance optimization and many-core technologies

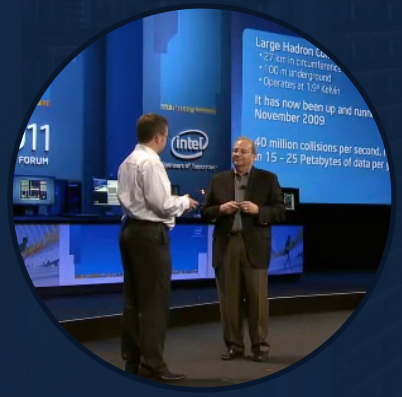

CERN openlab was one of the first Intel partners worldwide to start work and deliver results on the Intel MIC architecture

CERN openlab and Intel – a lasting partnership

Additional areas for openlab phase IV: Exascale, Cloud, Security, Embedded

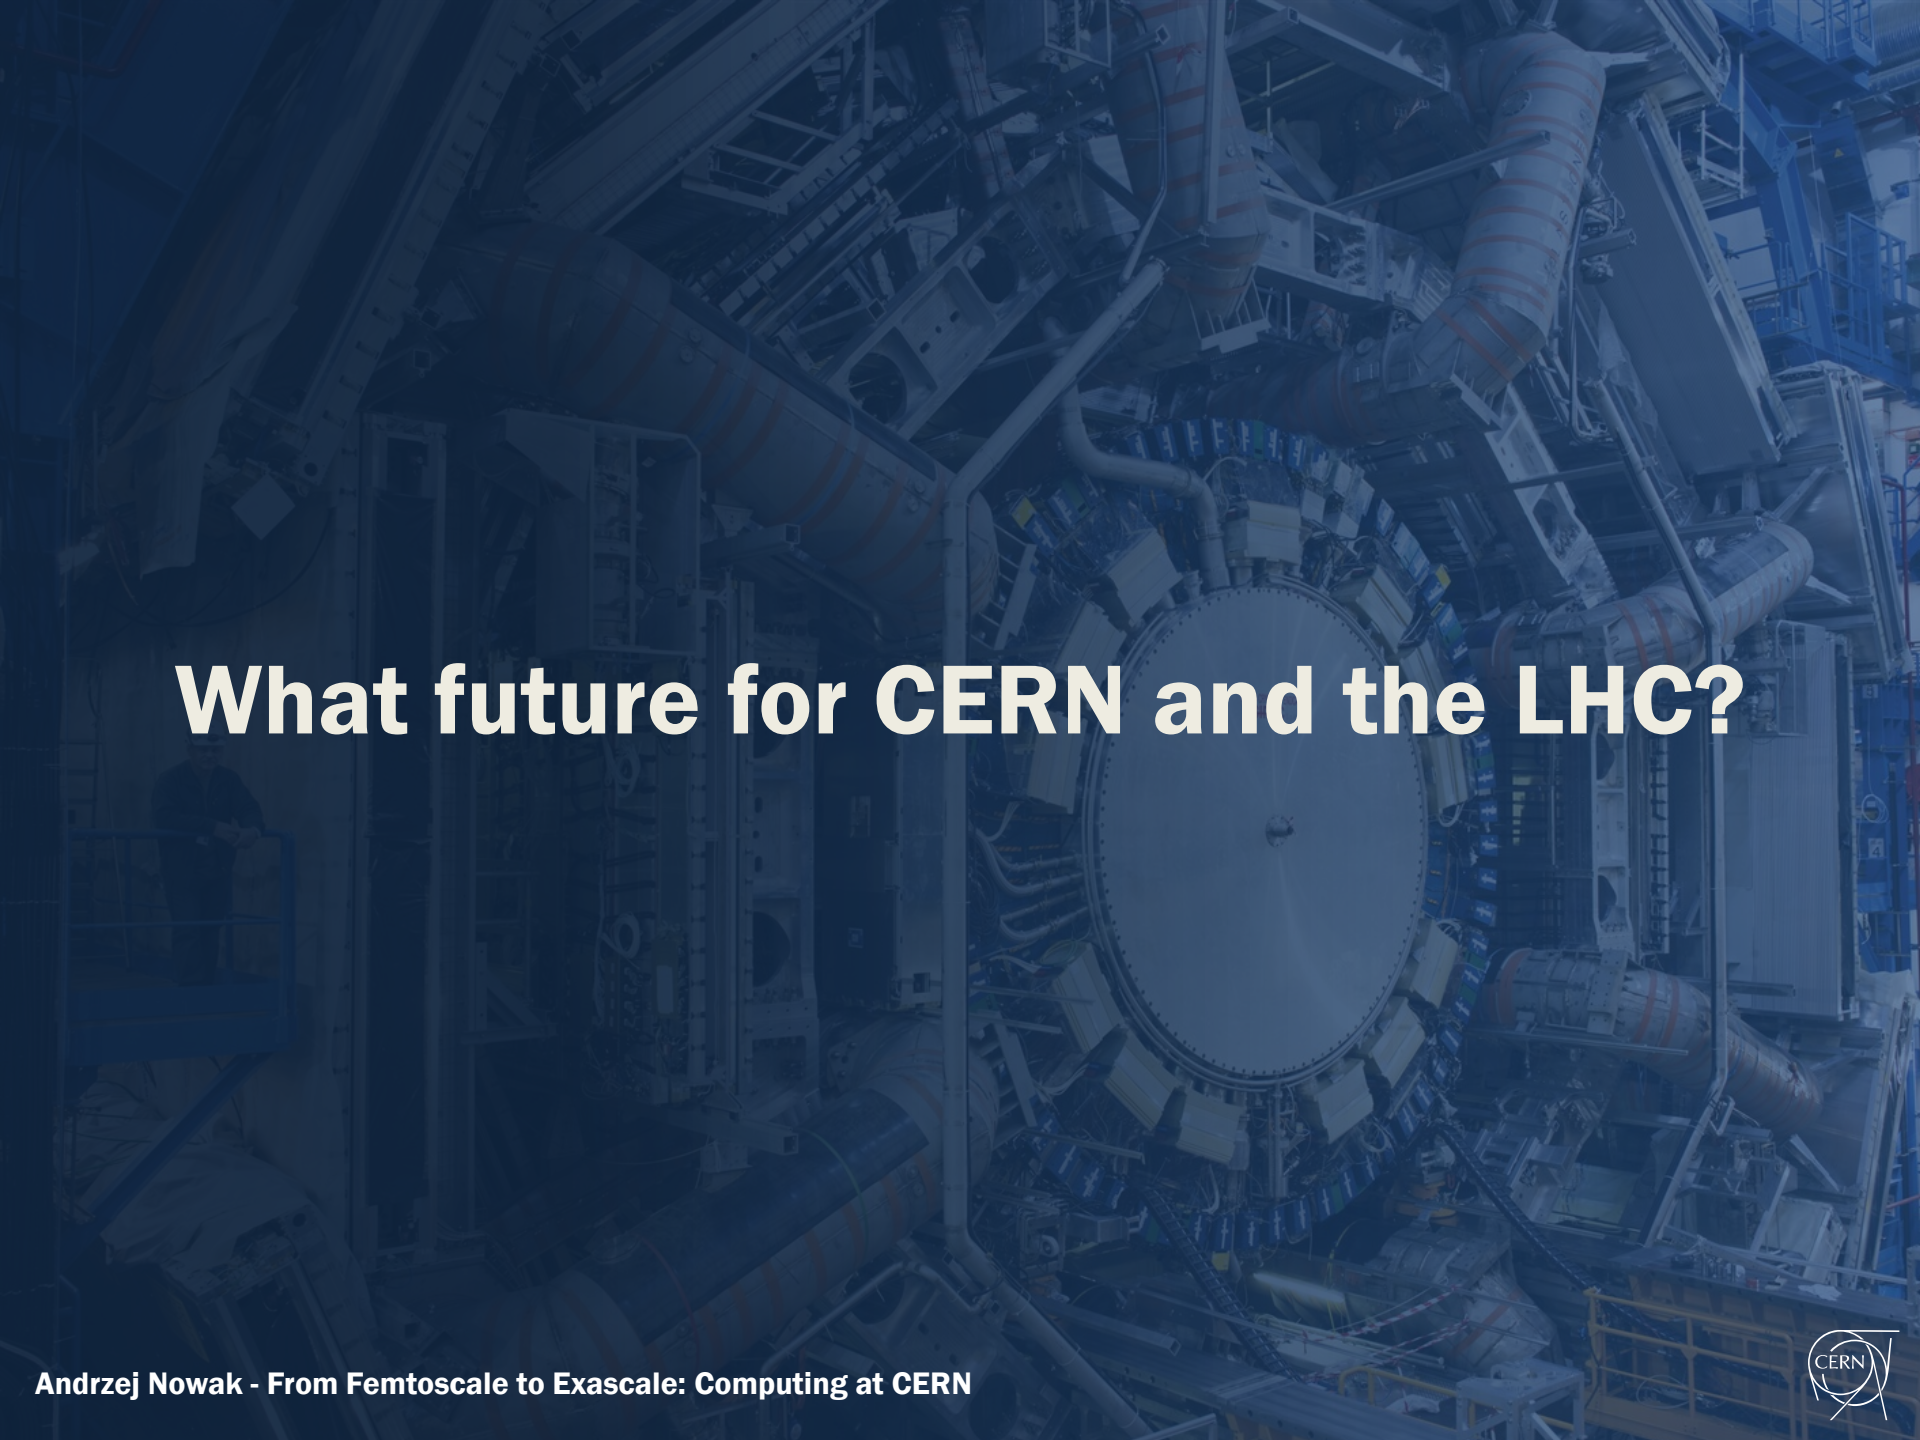

What future for CERN and the LHC?

Q & A