ESRFUP – WP11
Data Management Requirements

Outline

1. Examples of data intensive science
2. The ESRF in numbers
3. Today's data management @ ESRF
4. Infrastructure overview
5. ESRF Upgrade and data management
6. GRID + data management
7. International collaborations
8. Conclusion

1. Examples of Data Intensive Science

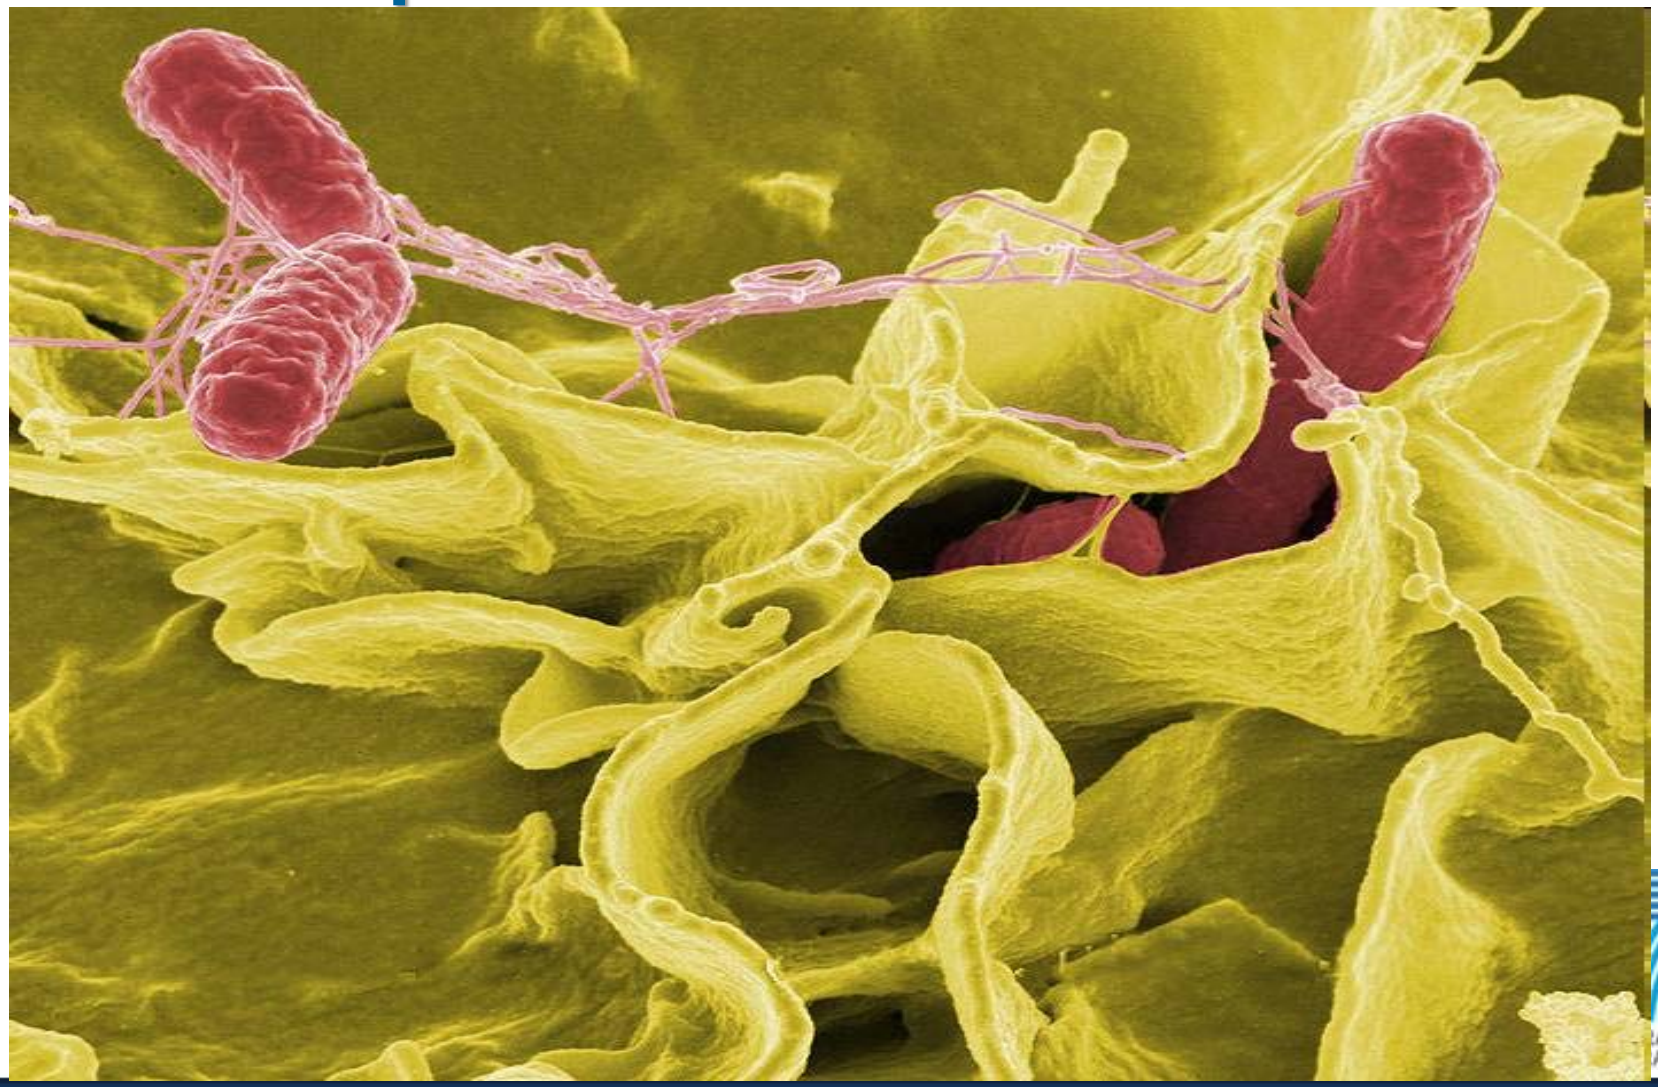

Structural Biology of Deinococcus Radiodurans

- a non-pathogenic bacterium extremely resistant to ionising radiation
- able to withstand up to 15,000 grays of ionising radiation
- *D. radiodurans* is capable of surviving and repairing hundreds of double strand breaks
- in *E. coli* five double strand breaks are invariably fatal
- *D. radiodurans* was amongst the first organisms for which the complete genome sequence was made available

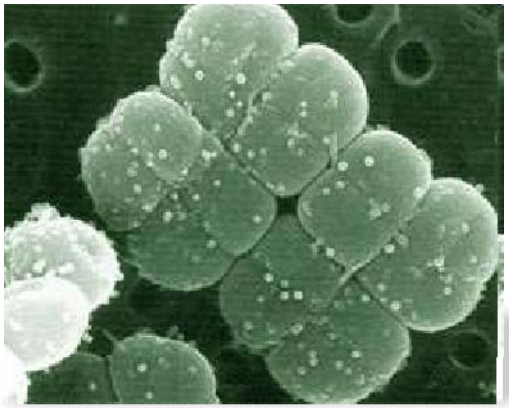

Deinococcus Radiodurans

- The ESRF MX-group has selected protein targets for structural studies
- Targets included proteins involved in DNA-damage repair, desiccation resistance, and oxidative stress response
- 24 structures have already been deposited in the PDB
- **Many more will need to be investigated to understand the repair mechanisms.**

- **MX projects range from 10 to 500 GB, with an average of 100 GB**

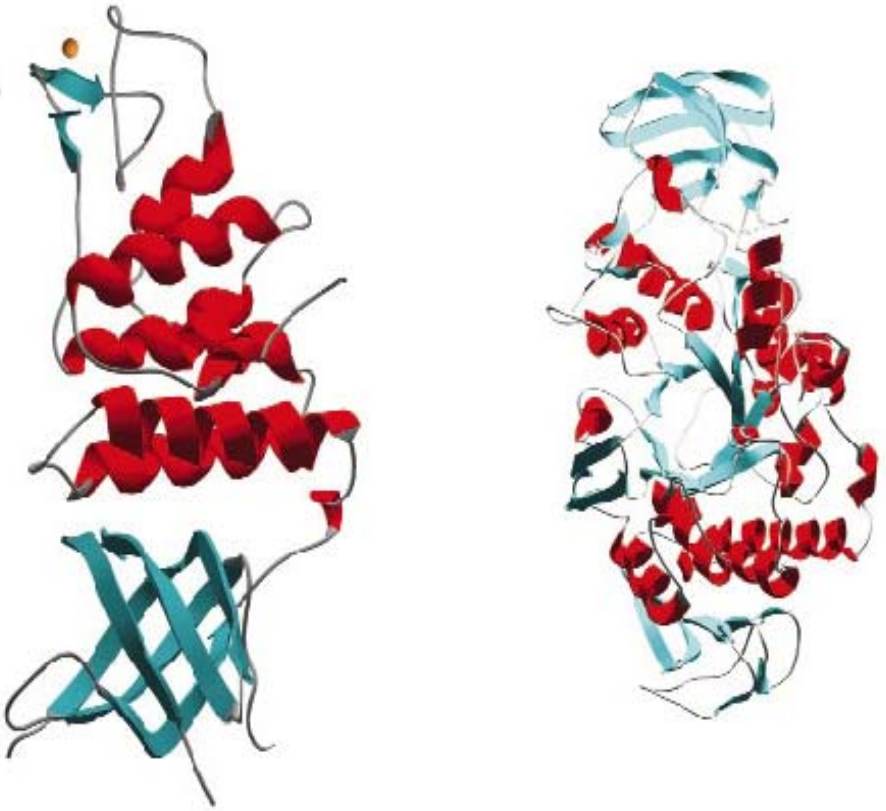

Examples of protein structures solved from *Deinococcus radiodurans*.
Credits: D. Hall, I. Leiros, H-K. Leiros, E. Micossi, E. Gordon, S. Macedo, U. Kapp, C. Jamin, J. Timmins, S. McSweeney; ESRF

Propagation Phase Contrast Radiography of opaque amber

Virtual 3D extraction of organisms in opaque amber

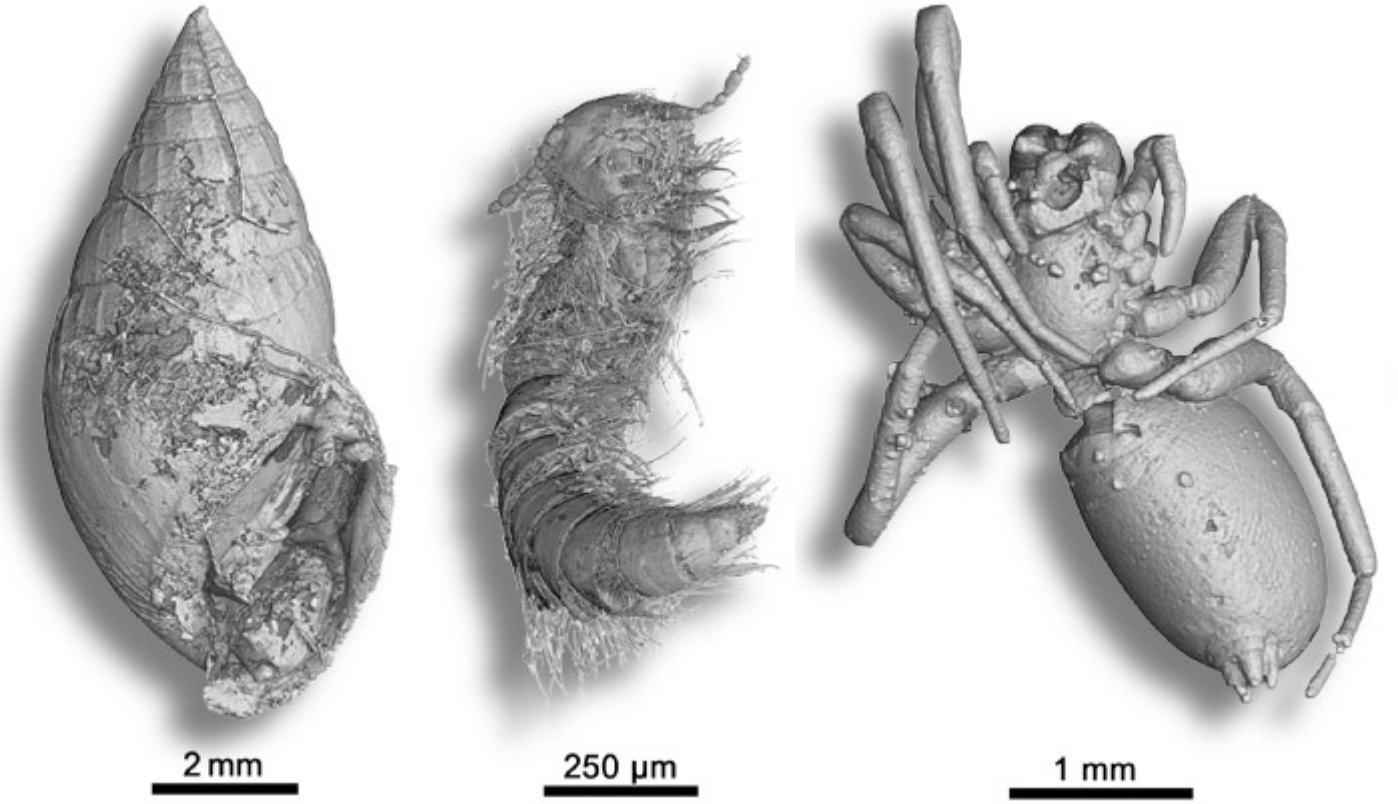

Gastropod Ellobiidea

Myriapod Polyxenidae

Arachnid

Credits: M. Lak, P. Tafforeau, D. Néraudeau (ESRF Grenoble and UMR CNRS 6118 Rennes).

SEVENTH FRAMEWORK PROGRAMME

Virtual 3D extraction of organisms in opaque amber

Conifer branch (Glenrosa)

Isopod crustacean Ligia

Insect hymenopteran Falciformicidae

Credits: M. Lak, P. Tafforeau, D. Néraudeau (ESRF Grenoble and UMR CNRS 6118 Rennes).

SEVENTH FRAMEWORK PROGRAMME

Toumai – our oldest ancestor?

- Micro tomography phase contrast studies of Toumai (Sahelanthropus tchadensis) reveal the 3D structure of the skull and allows remodeling of missing fragments.
- 1 scan = volume of 160 x 160 x 200 mm, pixel size of 45 µm = ~ $1*10^{11}$ Bytes = ~100 GB
- The raw data set is in the order of 400 GB
- The processed data represents more than 4 TB.
- Processing has taken several 1000 h CPU time on a 2.2 GHz AMD Opteron system

- **Tomography projects range from 200 GB to 10 TB with an average of 800 GB**

SEVENTH FRAMEWORK PROGRAMME

2. The ESRF in numbers

The ESRF in numbers

- 42 Beamlines (31 ESRF, 11 CRGs)
- 6 200 user visits in 2007
- ~7 500 research entities registered in our database
- 2 000 proposals, 900 accepted in 2007
- 15 300 eight hour shifts for experiments in 2007
- >1 800 peer reviewed publications/year

- Two disciplines generate mountains of data
  - Life Science
  - Imaging
- An average of 2 TB/day of data in 2008
- >80 000 000 files in 2008

SEVENTH FRAMEWORK PROGRAMME

Evolution of the ESRF data production

- 12 years → data volume x 400.
  In 2007: 300TB ~$1*10^8$ files,
  In 2008: ~400TB ~$8*10^7$ files
- We managed to provide the necessary network bandwidth, disk space, backup capacity at constant budget
- However, we have doubled the data centre infrastructure (m^2, kW, cooling)
- Today: ~250 kW, ~300 m^2

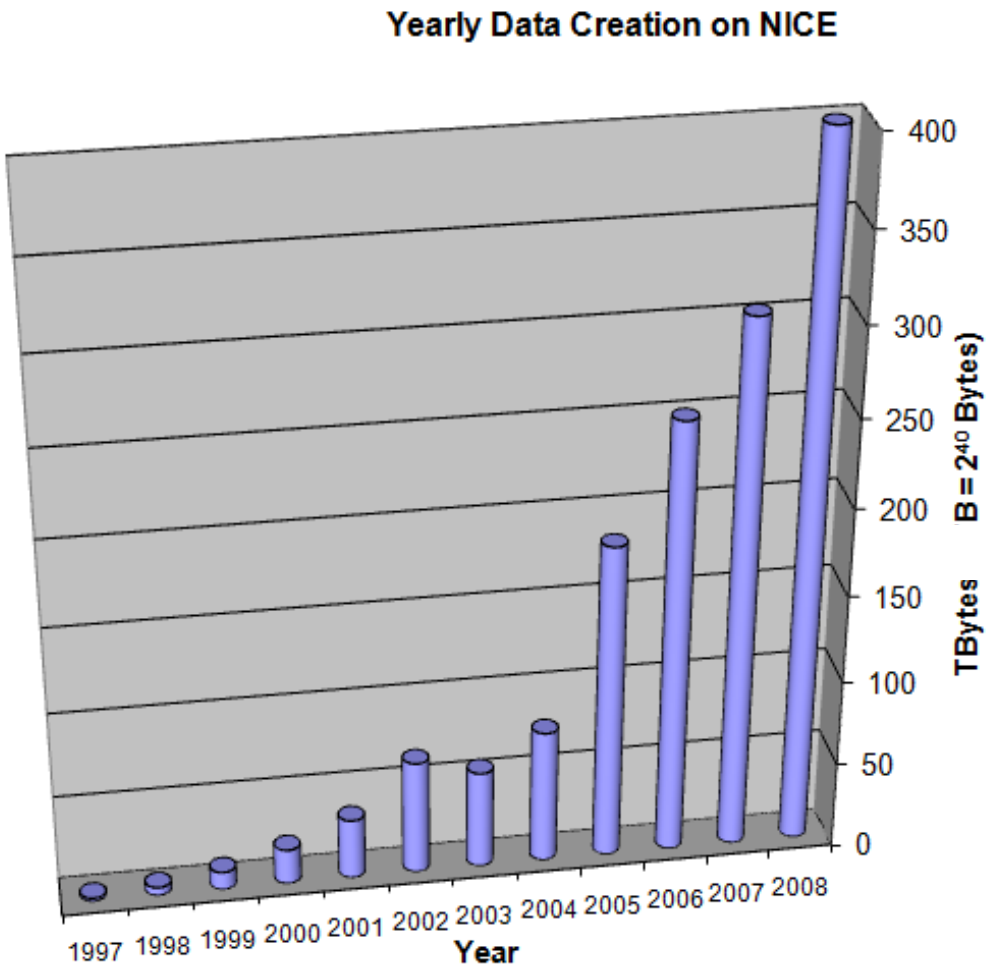

Data Storage – disk space allocation

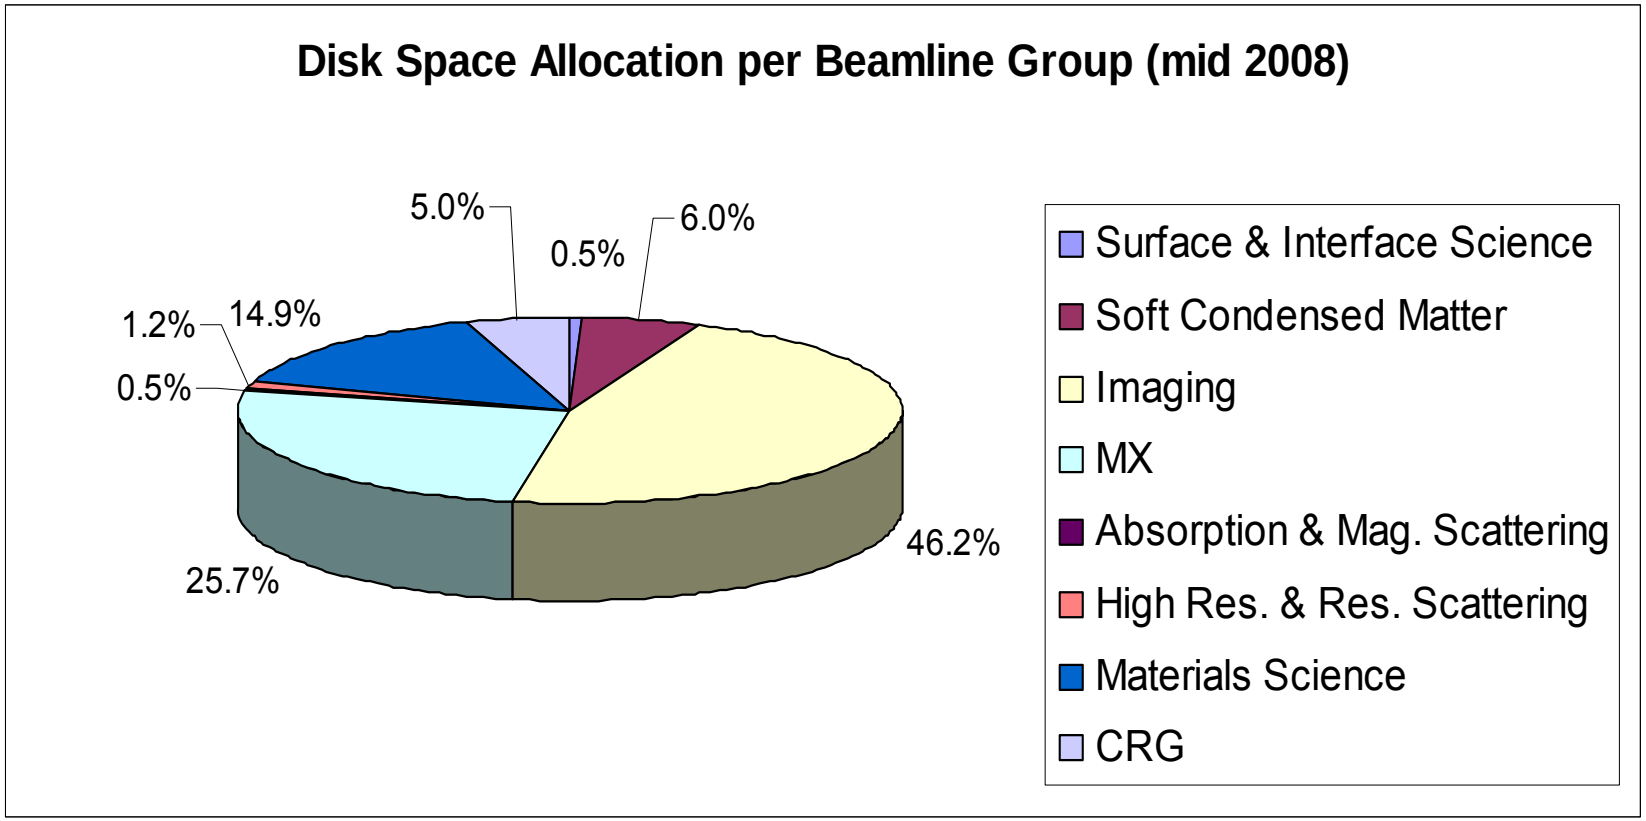

SEVENTH FRAMEWORK PROGRAMME

3. Data management @ ESRF

ESRF Storage Policy for Scientific Data

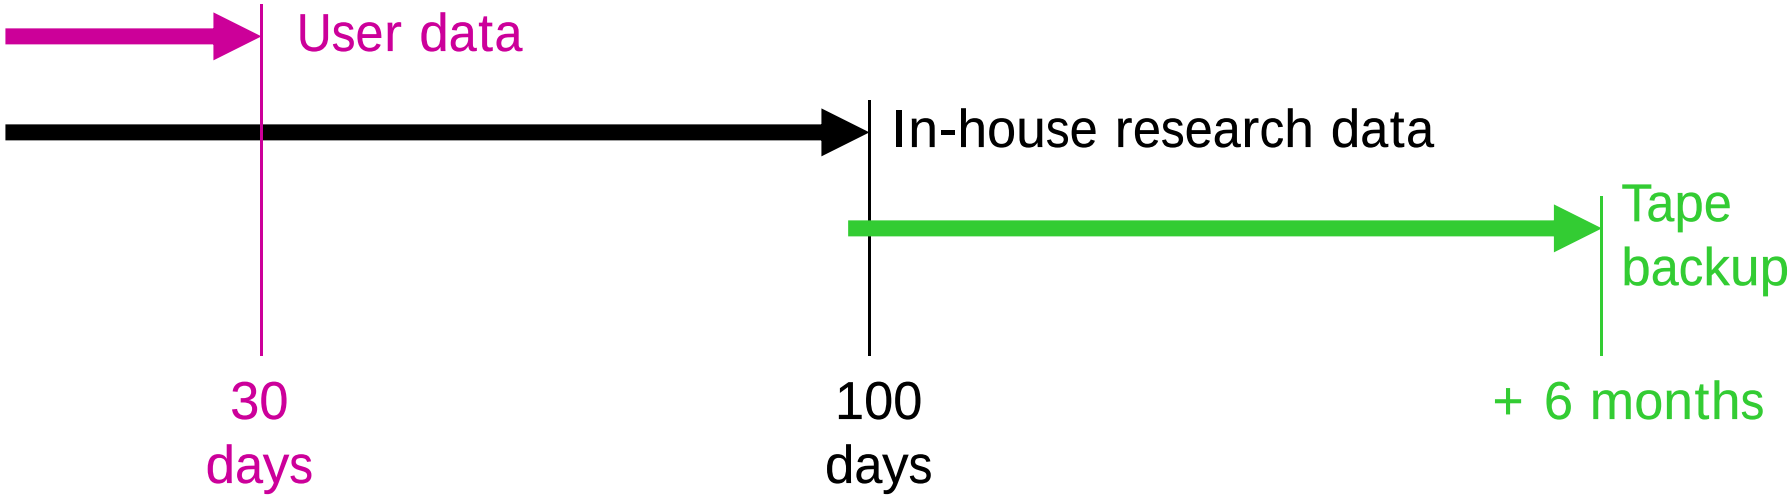

- After 6-9 months data has disappeared from our data centres!
- Archiving of the data is currently left to the users!
- This policy is now under debate.

ESRF Storage Policy for Scientific Data - 2009

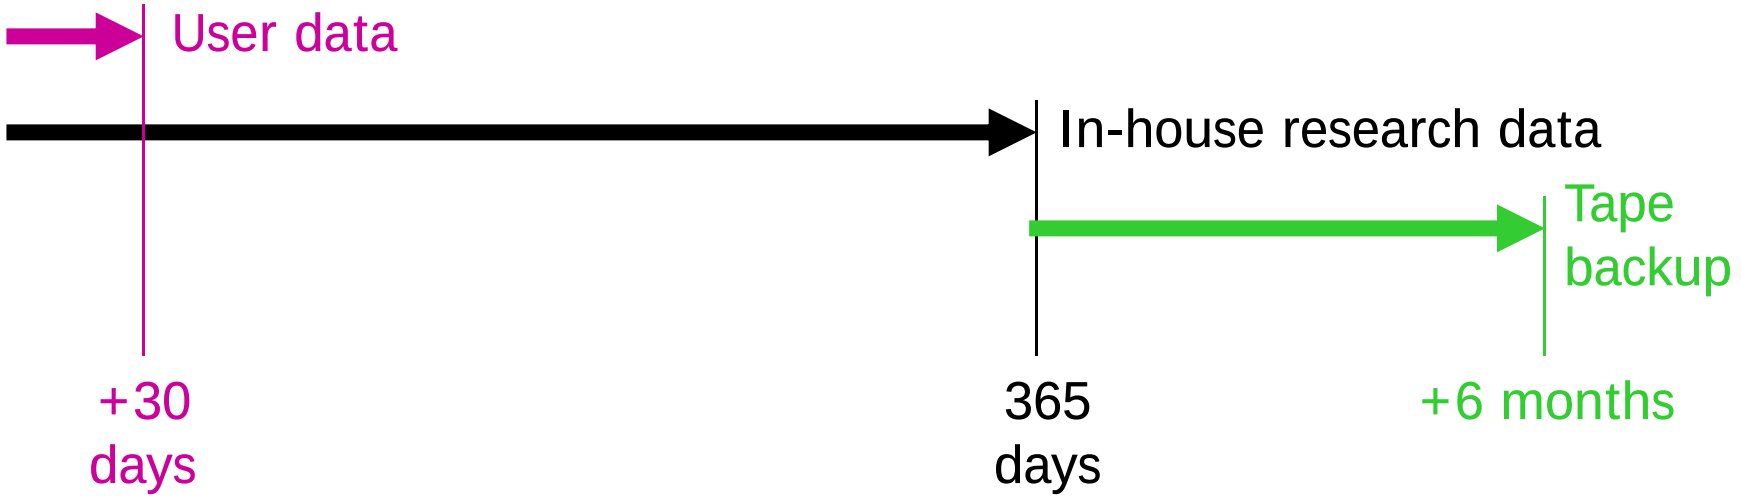

- User data: erased 30 days after departure
- In-house research data: 12 months
- Backup's: + 6 months

What is our current situation?

Many photon laboratories apply similar data policies

Consequence:

- Data remains “private” property of scientists
- No easy network access to data
- No guarantee that data can be found/read after a few years
- Some scientists are left with an un-manageable problem
- A dramatic loss of data!
- Data analysis can become very difficult

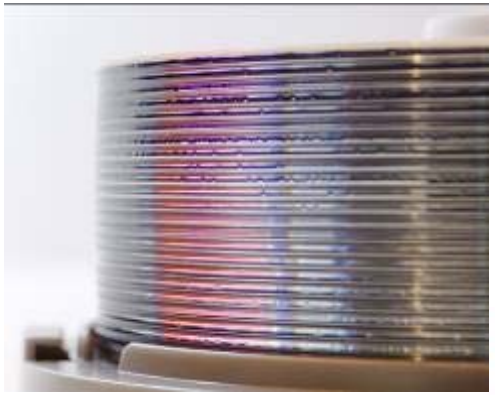

- This situation has let to the **FP7 EDNP proposal**
- and to the **VEDAC proposal**

SEVENTH FRAMEWORK PROGRAMME

4. Infrastructure overview

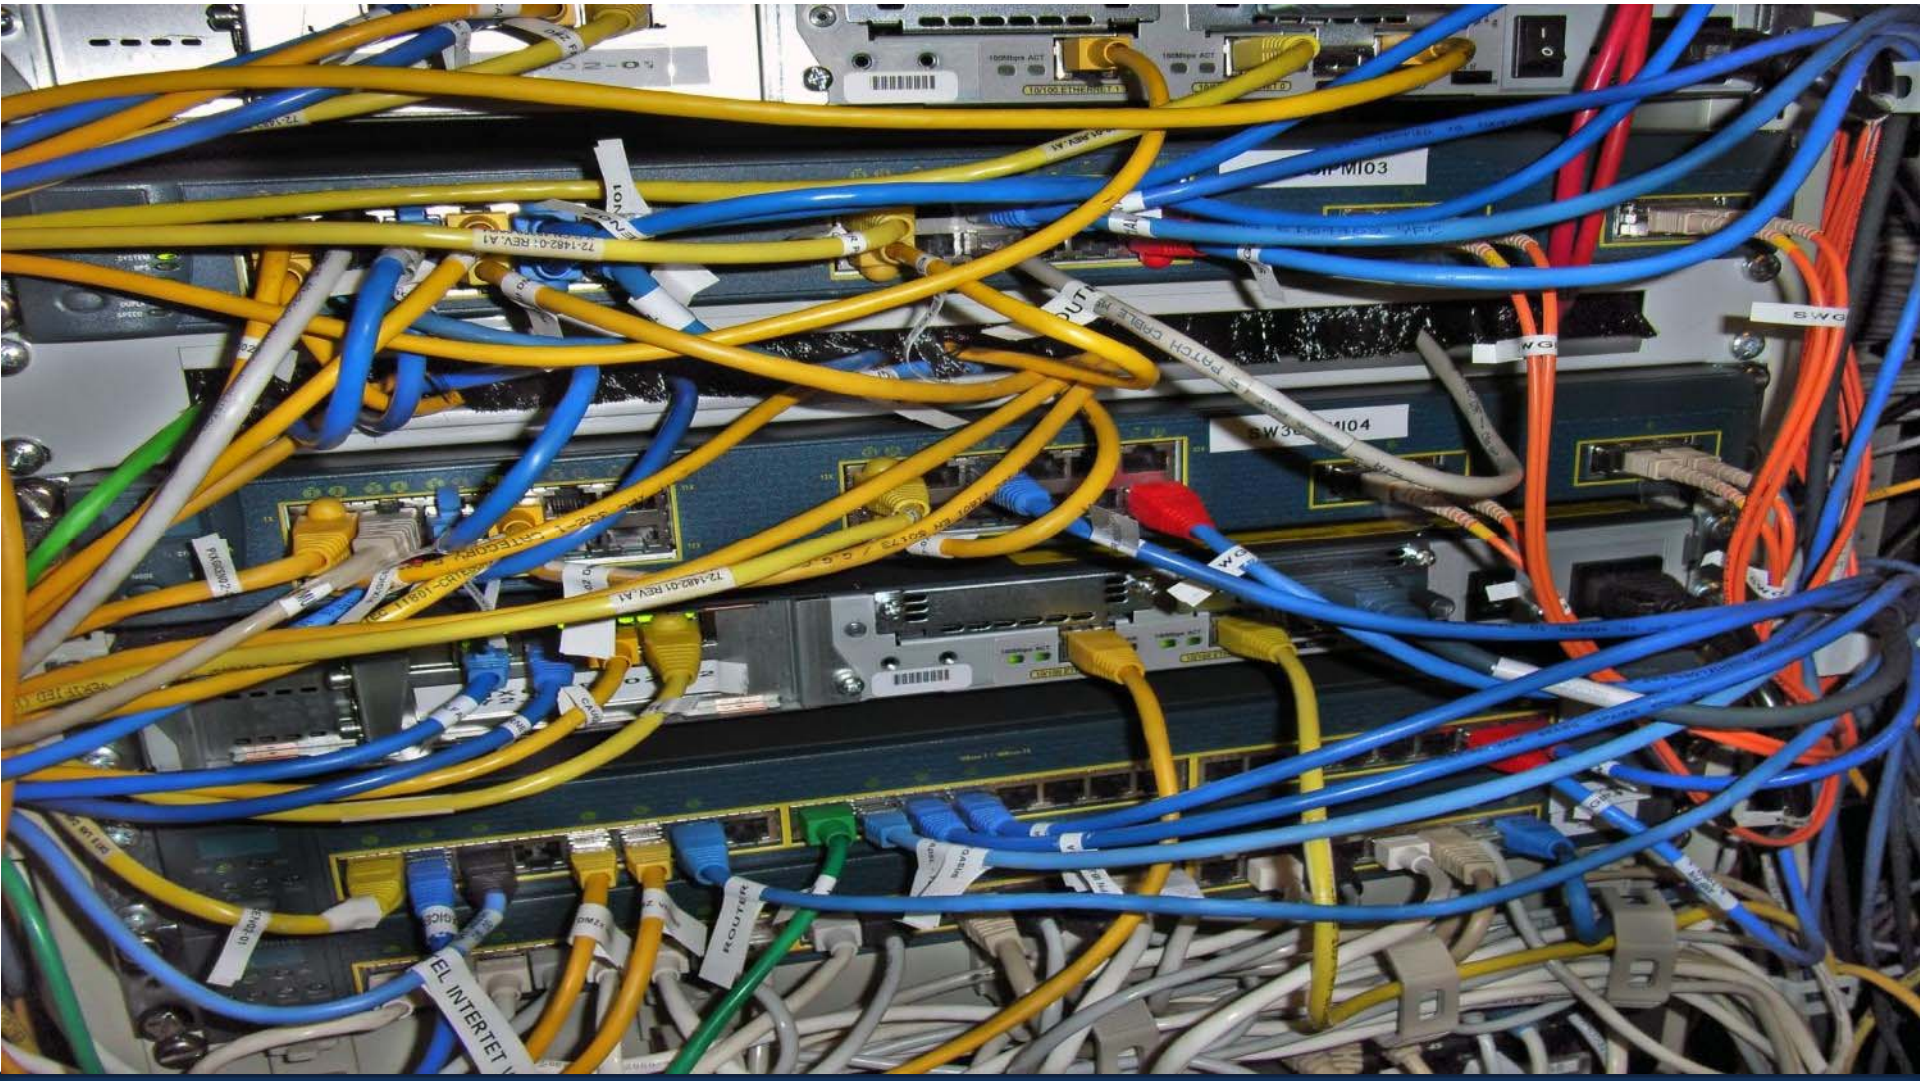

Data Network

- The LAN is the spinal cord of the laboratory
- Reliability of all components is fundamental

ESRF operates with a class B address:

- 160.103.a.b; a=subnet, b=host

Network entirely based on Extreme Network switches:

- BlackDiamond10k and BlackDiamond8k
- Multiple aggregated 10 Gbps backbone links
- Edge devices on the beamlines are Extreme Summit 400-48T switches

SEVENTH FRAMEWORK PROGRAMME

The backbone network

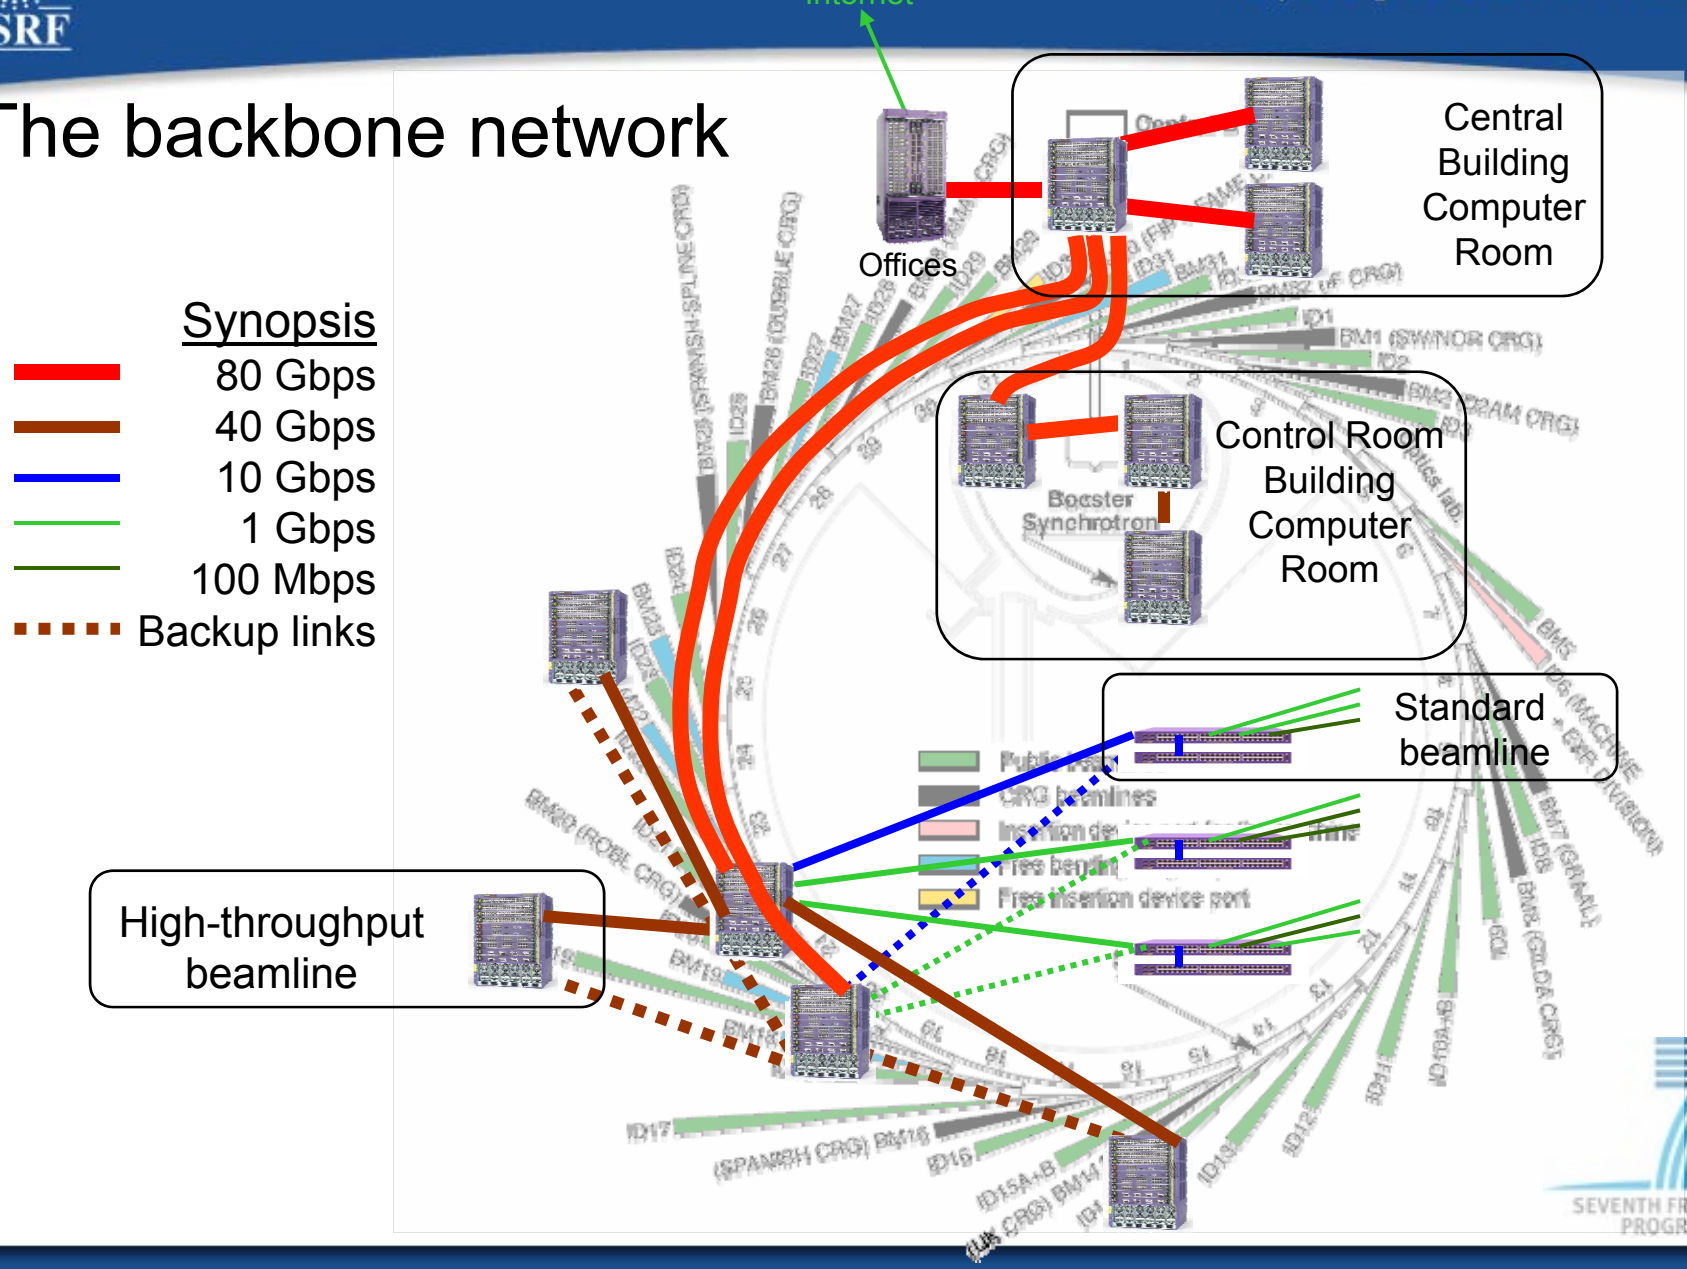

ESRF – connected to the world via RENATER

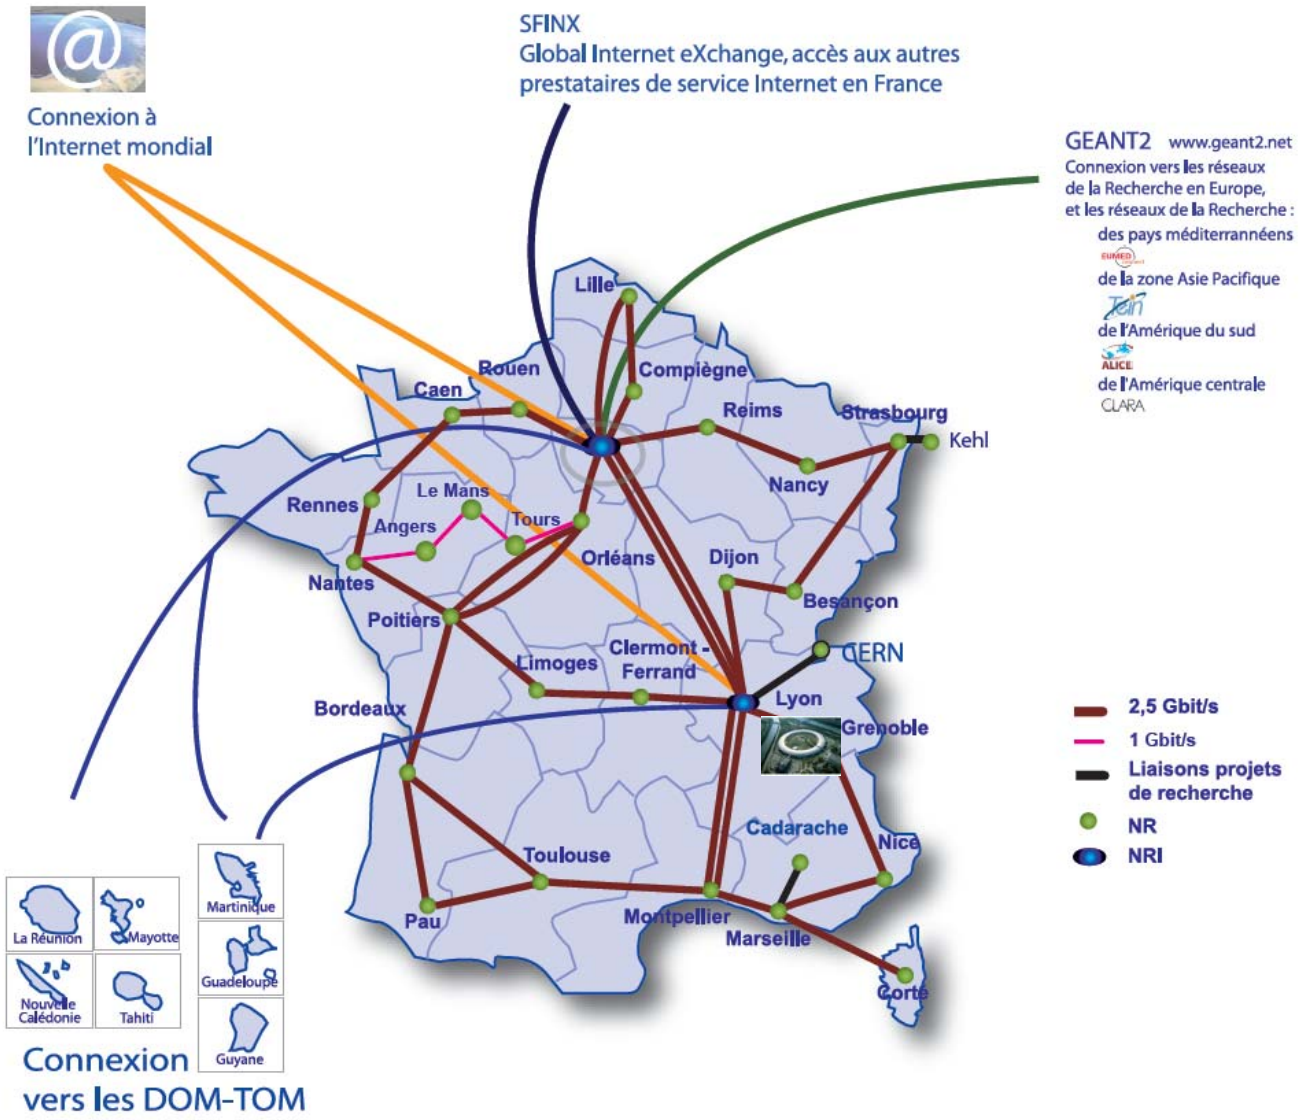

SEVENTH FRAMEWORK PROGRAMME

GEANT2 network infrastructure

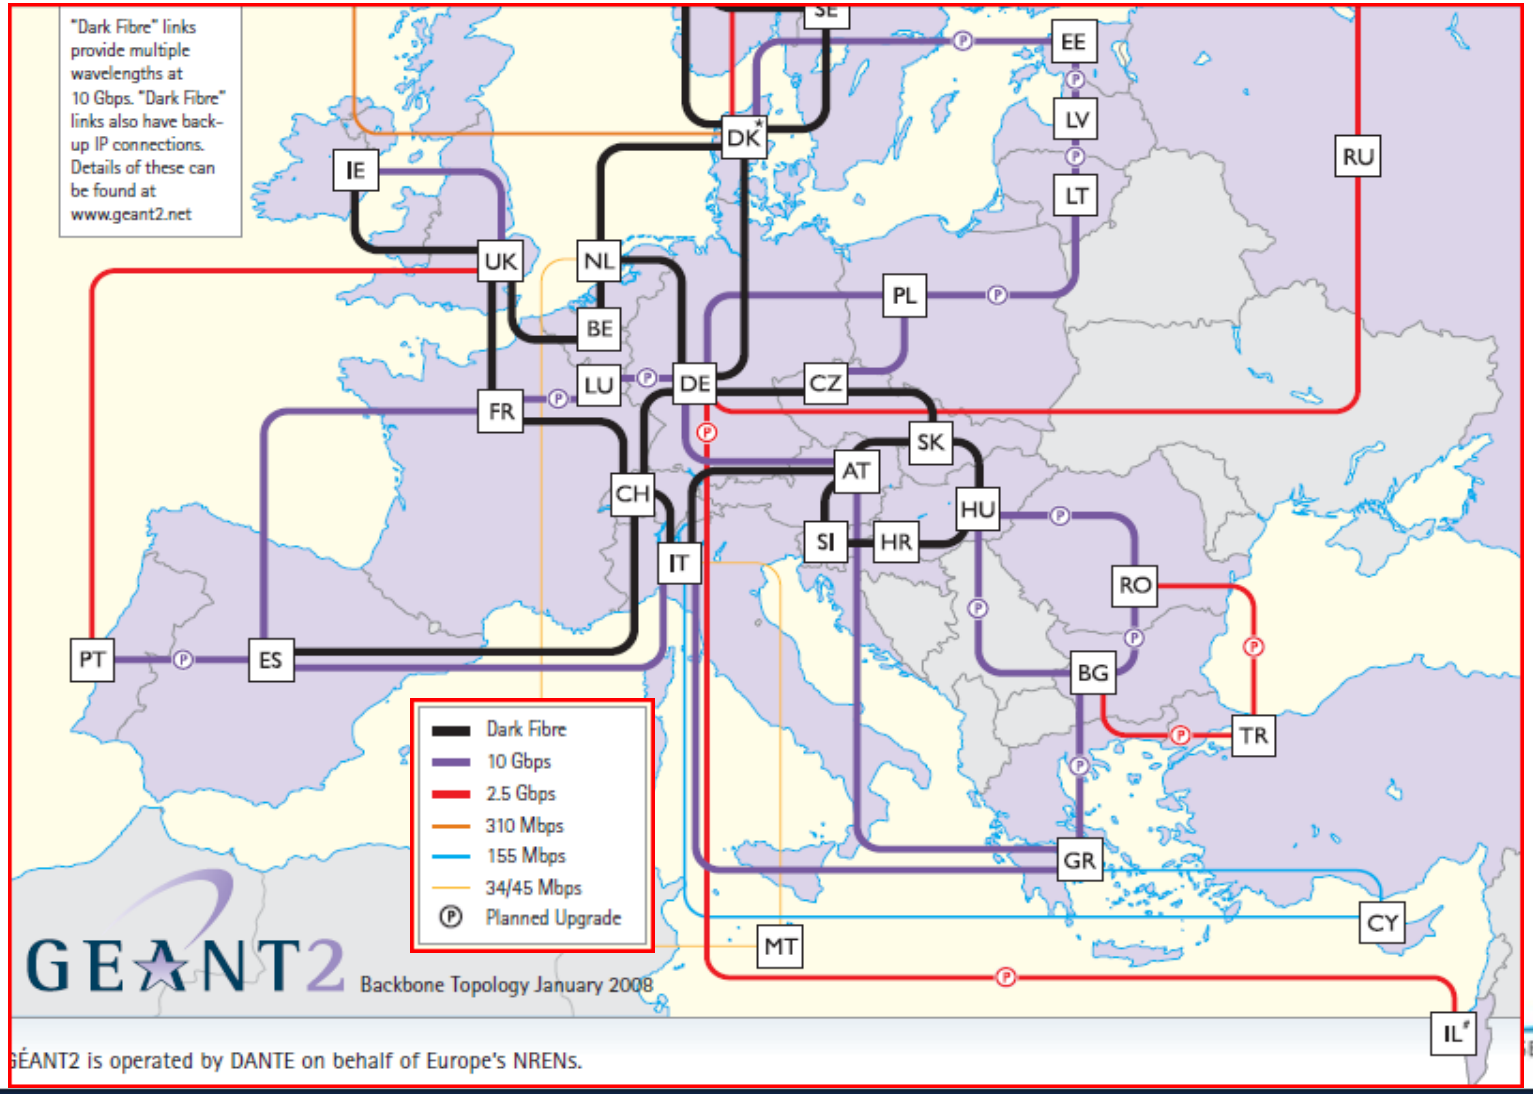

SEVENTH FRAMEWORK PROGRAMME

Central Data Storage: NAS systems

4 x Network Appliance FC940 cluster
- Eight controllers, 1GE network attachments
- 70 TB

2 x Network Appliance FAS3050 cluster
- Eight controllers, 1GE network attachments
- 140 TB

2 x Network Appliance FAS6070 cluster
- Eight controllers, 10GE network attachments
- 400 TB

600 TB total disk storage capacity, currently NFS only

75 + 250 MB/s single process NFS read/write performance

~700 MB/s cluster access for data processing

SEVENTH FRAMEWORK PROGRAMME

Data Storage: Tape Backup

4 x StorageTek L700, 700 slots,
    4-10 LTO-2 and LTO-3 tape drives

1 x StorageTek L8500, 3 500 slots,
    18 tape drives LTO-3
    12 tape drives LTO-4
The L8500 can be extended to >10 000 slots

Software: Quadratec Time Navigator
Backup of very large file systems is difficult

SEVENTH FRAMEWORK PROGRAMME

5. The ESRF UP and data management

The ESRF Upgrade Programme

A 290 M€, 10 year programme to keep the ESRF at the forefront of photon science

- Nano-Science
- Biology + Soft matter
- X-ray imaging
- Pump-probe, time resolved
- Extreme conditions

- A vast building programme
- Extended beamlines
- An enhanced X-ray source
- Enhanced instrumentation
- Powerful computing

- The first phase of the Upgrade Programme starts in 2009 and extends over 7 years
- **170 M€ will be invested**

SEVENTH FRAMEWORK PROGRAMME

Its all about data

Why will the data avalanche become worse?

1. Automation
2. Remote access
3. New methods
4. Faster detectors, higher resolution

Screening & Automation for Macromolecular Crystallography

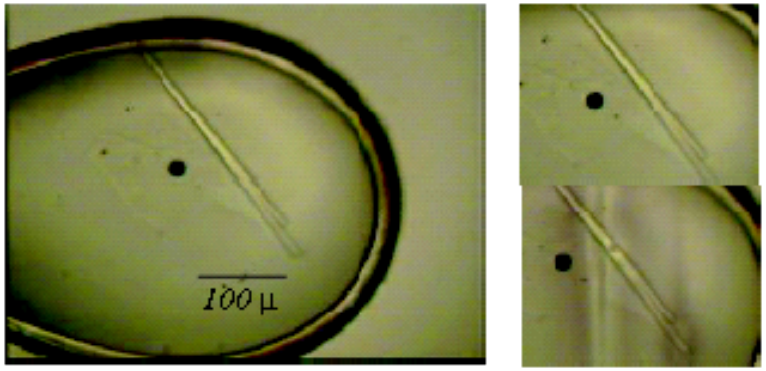

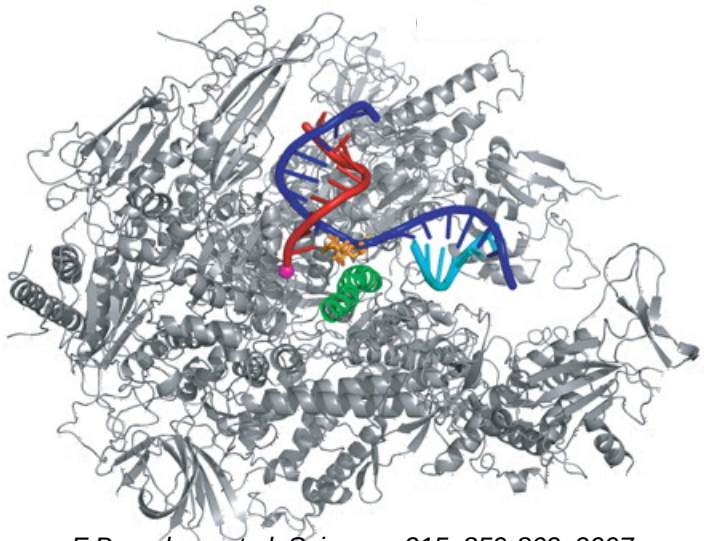

F Brueckner et al. Science , 315, 859-862, 2007

- Available crystals become smaller
- Interesting complexes become larger
- 2007 figures (>150.000 samples tested) expected to increase by more than a factor of 10

- Massif automation and nano beams will substantially increase data volumes

SEVENTH FRAMEWORK PROGRAMME

New methods: nanobeams & raster scans

Micrometer

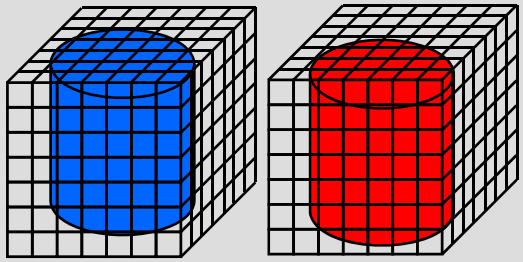

3D, 2 modalities

Low sampling

Nanometer

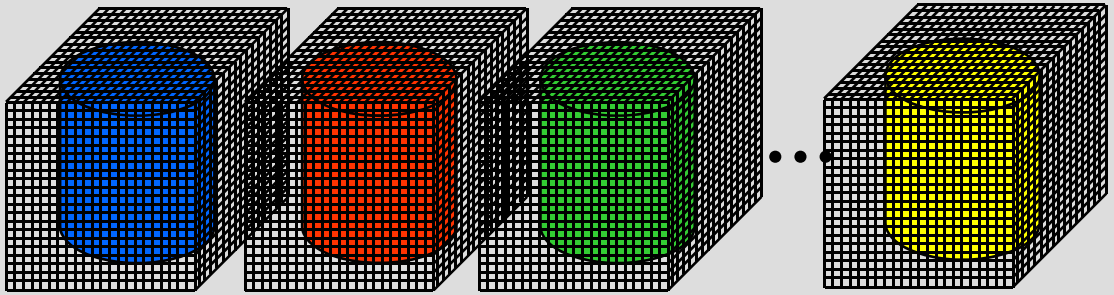

3D, Many modalities + in-situ

High sampling

"Few" small files

Many big files

KiloBytes | **MegaBytes** | **GigaBytes** | **TeraBytes** | **PetaBytes**

Store
Process
Reduce
On-line Analyses

Transfer
Off-line Analyses
Long term storage

Credits: Pierre Bleuet, ESRF

SEVENTH FRAMEWORK PROGRAMME

New Detectors – More Data

Quantuum ADSC Q315

- 6k x 6k pixels = 72 MB/image;
- ½ frame/s = **36 MB/s**

Frelon2k

- 16 frames/s = **66 MB/s**

Dalsa Pantera TF1M60

- 60 frames/s = **150 MB/s**

Sarnoff CAM512

- 512 x 512, 400 frames/s, **200MB/s**
- 1k x 1k = **500MB/s**

Pixel Detectors

- Very high frame rates, up to 6M pixels, **100 - 1000 MB/s**

SEVENTH FRAMEWORK PROGRAMME

How to stay abreast of the data avalanche?

Yearly Data Creation on NICE

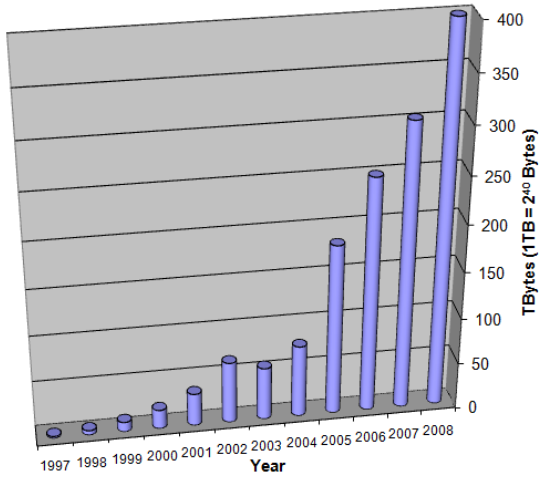

Today the ESRF produces 2 TB/day

How much data will the ESRF produce in 5-10 years?

If the trend of the last 10 years simply continues, i.e. a doubling of the rate every 18 months:

- In 5 years ⇨ **~ 15 TB/day**
- In 10 years ⇨ **~ 120 TB/day, i.e. 25 PB/y**

The ESRF Upgrade Programme is likely to accelerate this trend if and when pixel detectors come into the game

CERN LHC = 15 PB/year once fully operational)

SEVENTH FRAMEWORK PROGRAMME

Computing & the ESRF Upgrade

Computing is (and will remain) on the critical path for science at ESRF.
We have to make significant progress on two fronts:

Lower latency ⇨ diminish the time to measure, store, analyse data

- On-line data analysis
- Network, Storage, and Cluster upgrades
- Grid

Add functionality ⇨ new ways to measure, store, analyse data

- On-line and off-line data analysis
- Scientific Management Information System (SMIS), data mining
- Grid, remote access
- Data curation

SEVENTH FRAMEWORK PROGRAMME

6. Grid and data management

Grid & Data Management

...or why do we look into Grid technology?

Data Access

Today access to data is granted via generic accounts, e.g. ls2252

- Issues: Computer security, password management, mutual data protection, etc.

→Individual authentication needed

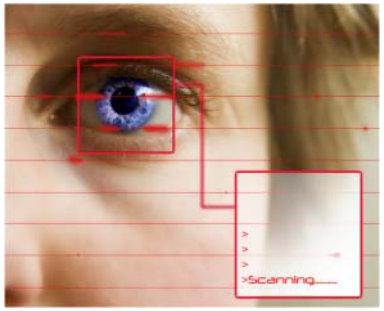

Today remote access to data is complicated

- Issues: Firewalls, encryption, accounts, file locations, etc.

→ Access to data must be easy and secure

SEVENTH FRAMEWORK PROGRAMME

Grid & Data Management (contd.)

…or why do we look into Grid technology?

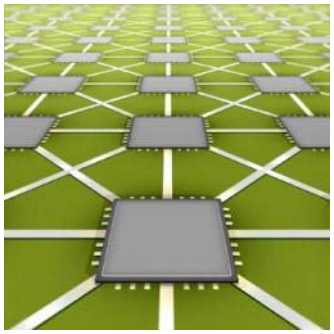

Data transport

Some experiments at ESRF may generate too much data … to carry away

- Detector: 250MB/s … USB disk: 10MB/s …Internet: ?

→ Leave some of the data at ESRF?

Data Storage

Some users may not have the resources to store (nor to analyse) the data

→ Use Grid storage and compute elements

SEVENTH FRAMEWORK PROGRAMME

Grid & Data Management (contd.)

...or why do we look into Grid technology?

Data from cross facility, cross discipline experiments

Merging data sets from several photon labs

Merging data sets from neutron and photon labs

- Issues: data formats, metadata description, data access

→ Data catalogues, easy data access

Data Preservation

Depositing data securely, curation of data

→ A user-friendly interface to data/metadata, browsing of data

SEVENTH FRAMEWORK PROGRAMME

7. International Collaborations

On-going International Collaborations

- **ESFRI Project ESRFUP**
  - SMIS and GRID work packages
- **European Data infrastructure for Neutrons and Photons (EDNP)**
  - Common data format, metadata definition, sustainable data storage, data curation and access, Grid
  - FP7 call in progress
- **Virtual European Data Analysis Centre (VEDAC)** for photon science
  - fast detector read-out, online data analysis
  - Project within the SR round-table
- **TANGO** control system collaboration
  - SOLEIL, ELETTRA, ALBA, DESY, ESRF

SEVENTH FRAMEWORK PROGRAMME

FP7 proposal EDNP (Infra 2008-1.2.2)

European **D**ata Infrastructure for **N**eutron and **P**hoton Sources

Sustainable data storage and data management tools to:

- Harmonise data policies in laboratories
- Manipulate large data sets
- Standardise data formats and annotation of data
- Allow transparent and secure remote access to data
- Establish sustainable and compatible distributed data catalogues
- Allow long term preservation of data
- Provide tools/interfaces for curating data
- Provide compatible open source data analysis software

~30 000 Scientists, >300 Instruments, ~4000 experiments/y

Current Status: a decision is expected before mid January 2009

8. Conclusion

The Grid may find its place in photon science for:

- **Authenticating users,**
- **Storing, accessing, browsing data remotely,**
- **Analysing large data sets,**
- **Curating data,**
- **Working together,**

if we can make it

user friendly!

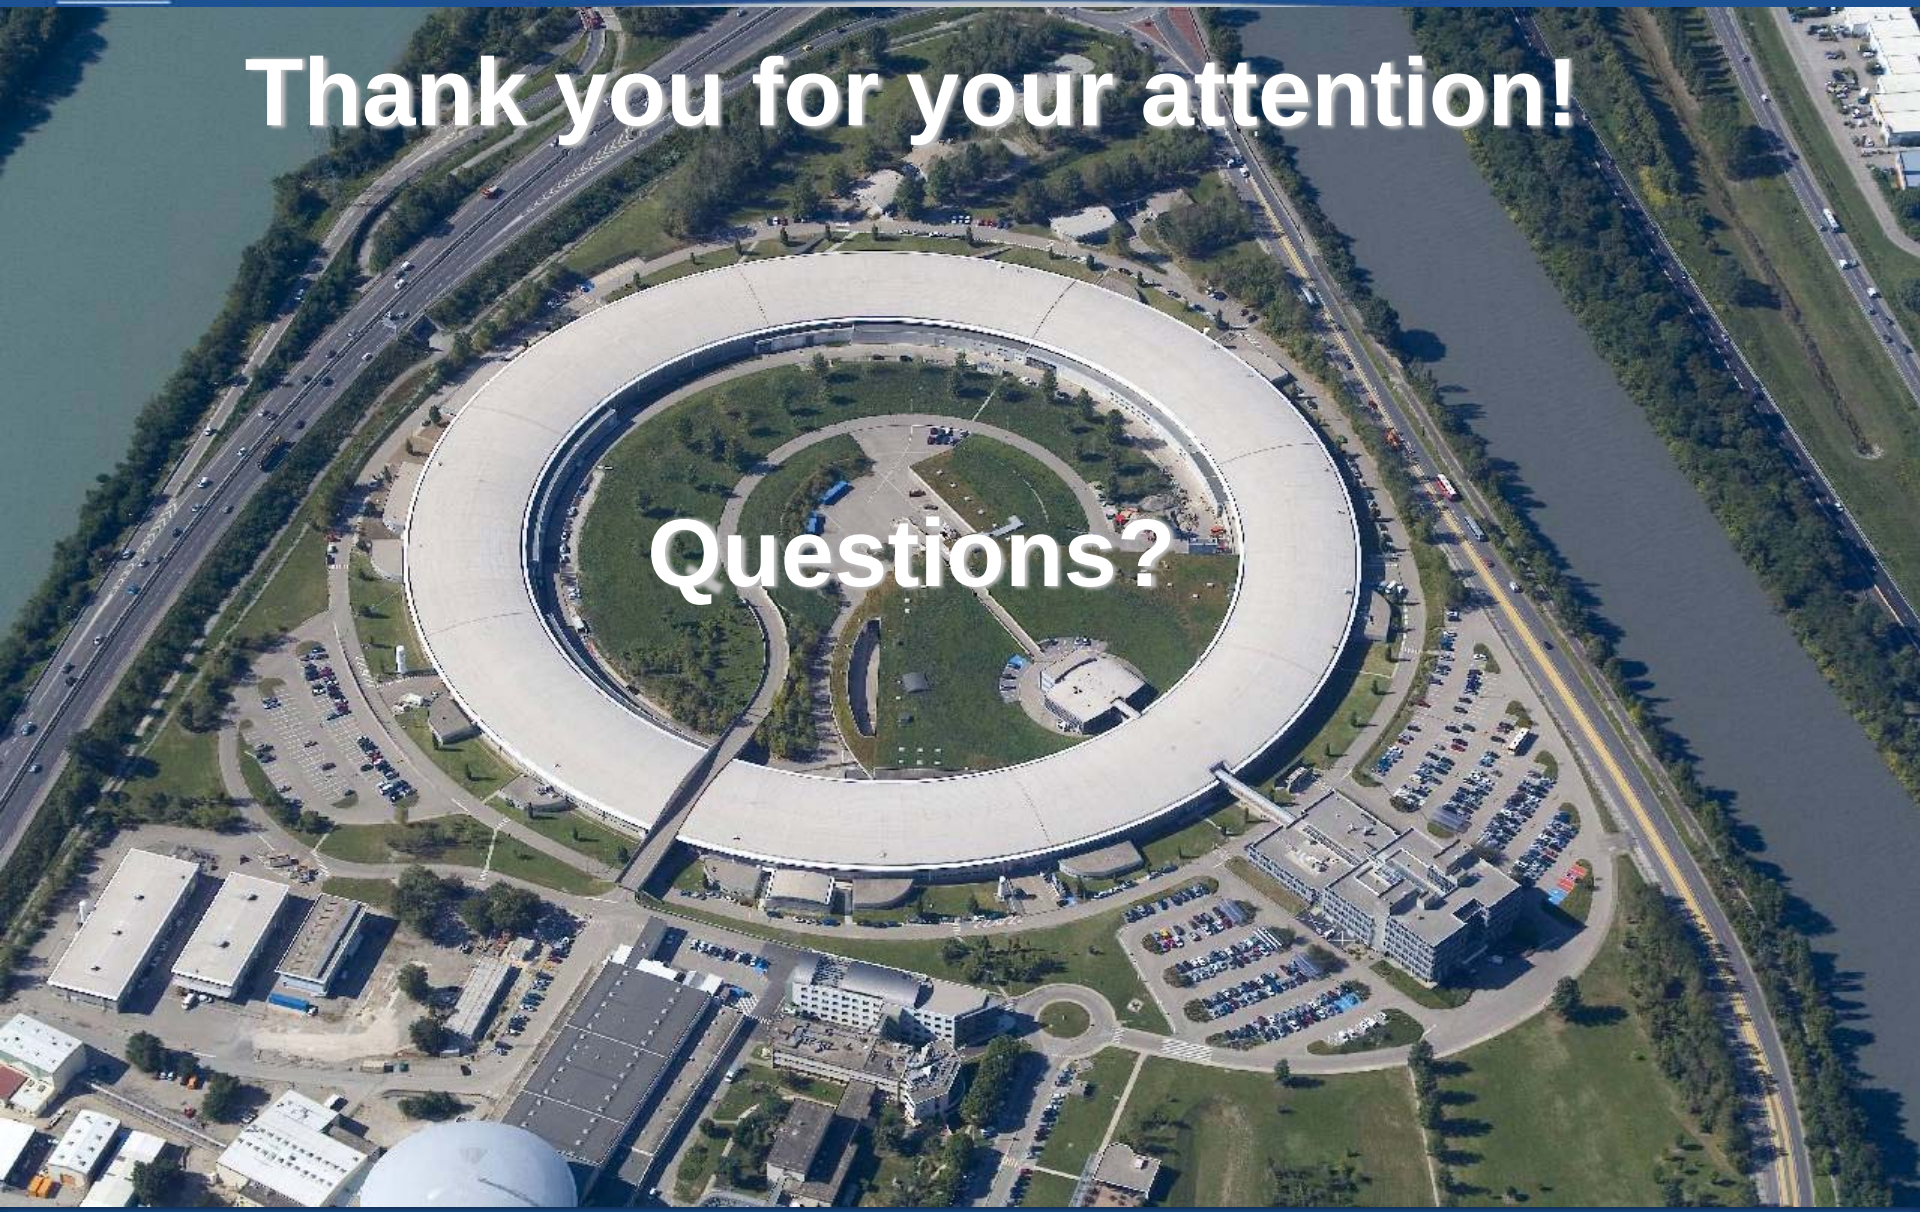
Thank you for your attention!
Questions?