Computing challenges in modern particle physics

Oxana Smirnova
FYST17/FKF050
March 17, 2008, Lund

Short list of questions for physicists

- Can gravity be included in a theory with the other three interactions?
- Why do the particles have the masses we observe, and what is the origin of mass?
- How many space-time dimensions do we live in?
- Are the known elementary particles fundamental or do they possess structure?
- Why is the electrical charge of the electron equal and opposite to that on the proton?
- Why are there three generations of quarks and leptons?
- Why is there overwhelmingly more matter than anti-matter in the Universe?
- Are protons unstable?
- What is the nature of the dark matter and dark energy that pervade our Galaxy?
- Are there new states of matter at exceedingly high density and temperature?
- Do the neutrinos have mass, and if so why are they so light?

Large Hadron Collider to the rescue

Collisions at the LHC

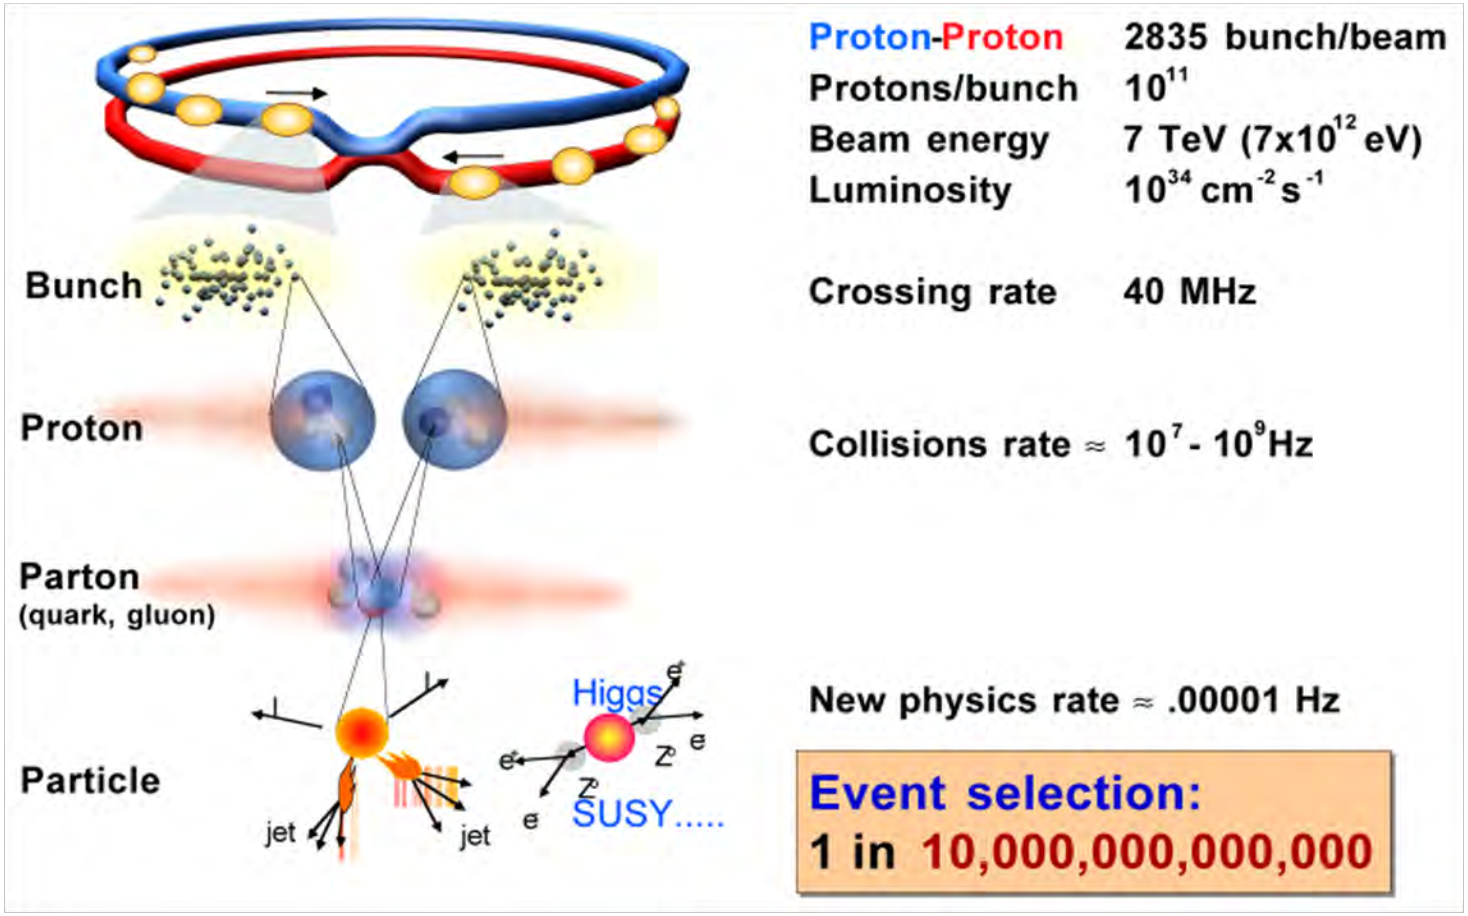

ATLAS detector at LHC

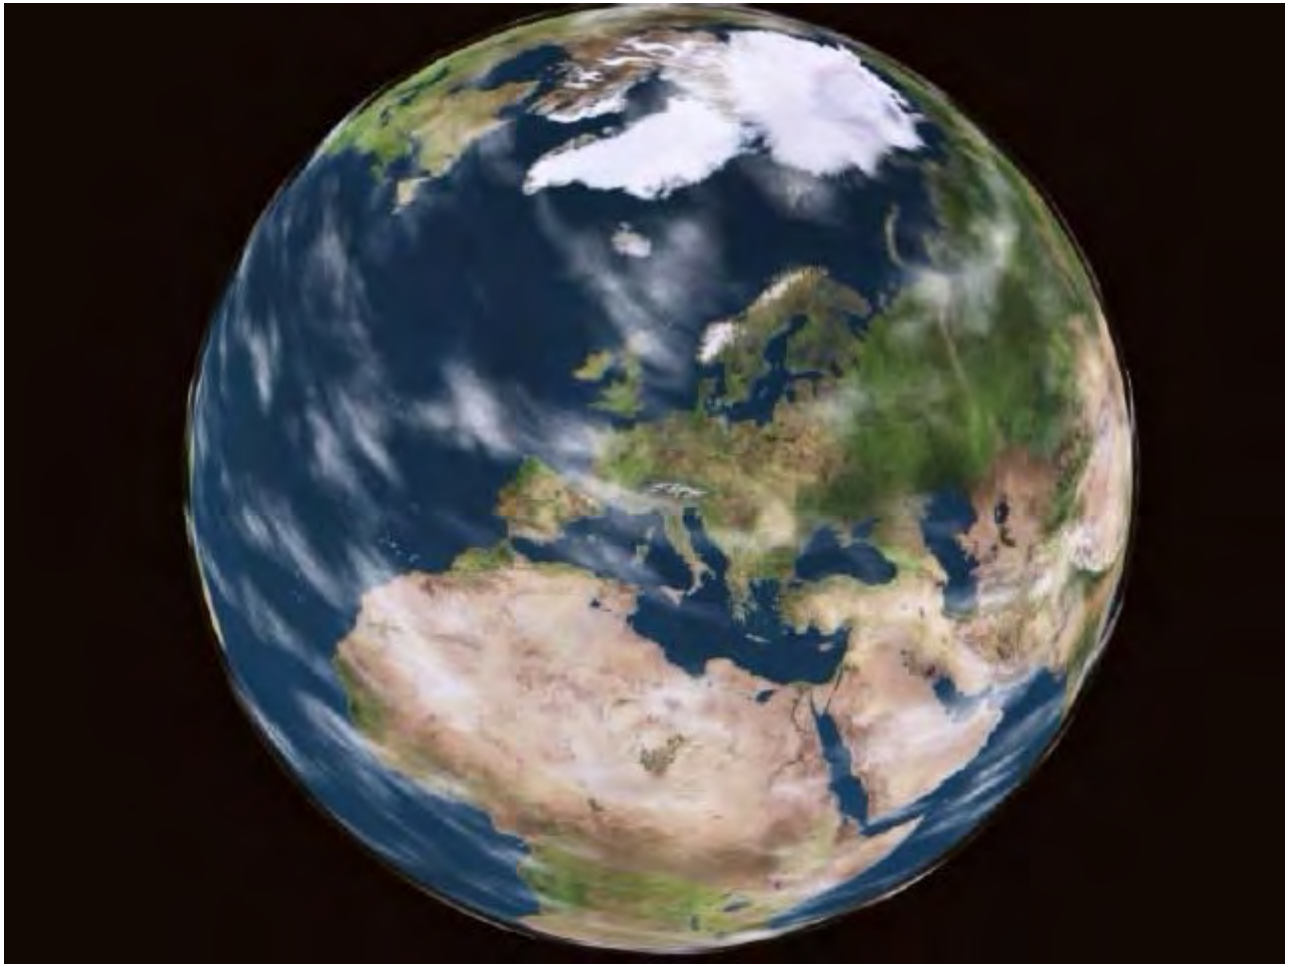

“Full chain” of HEP data processing

- Event generation (*Pythia*)
- Detector simulation (*Geant4*)
- Hit digitization
- Reconstruction
- Analysis data preparation
- Analysis, results (*ROOT*)

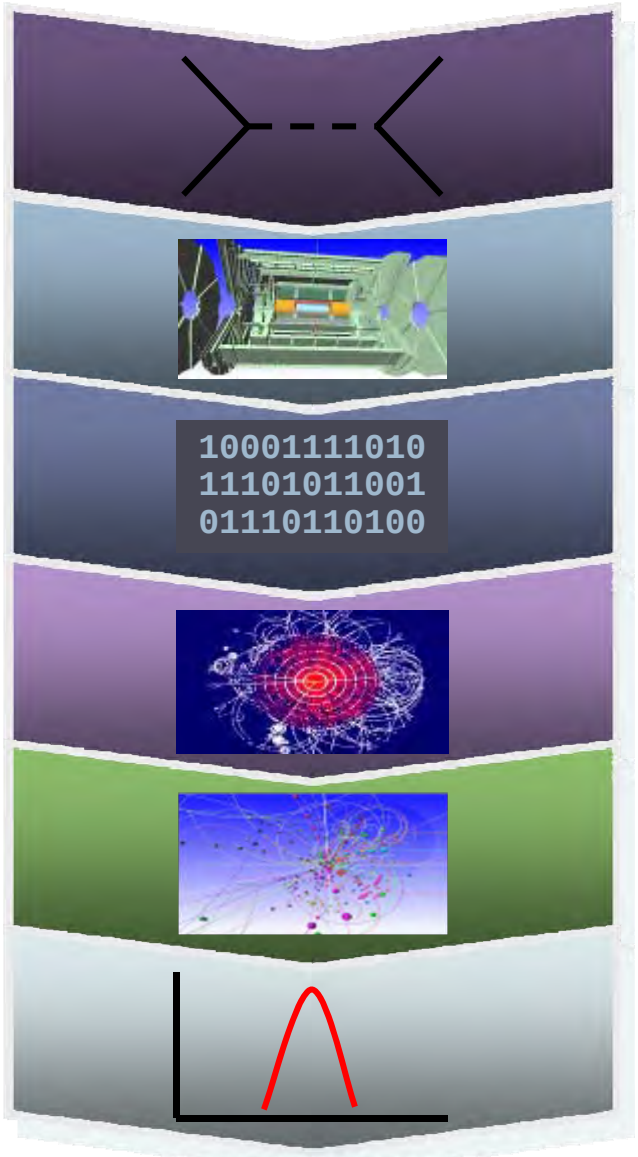

Slide adapted from Ch.Collins-Tooth and J.R.Catmore

Monte Carlo data production flow (10 Mevents)

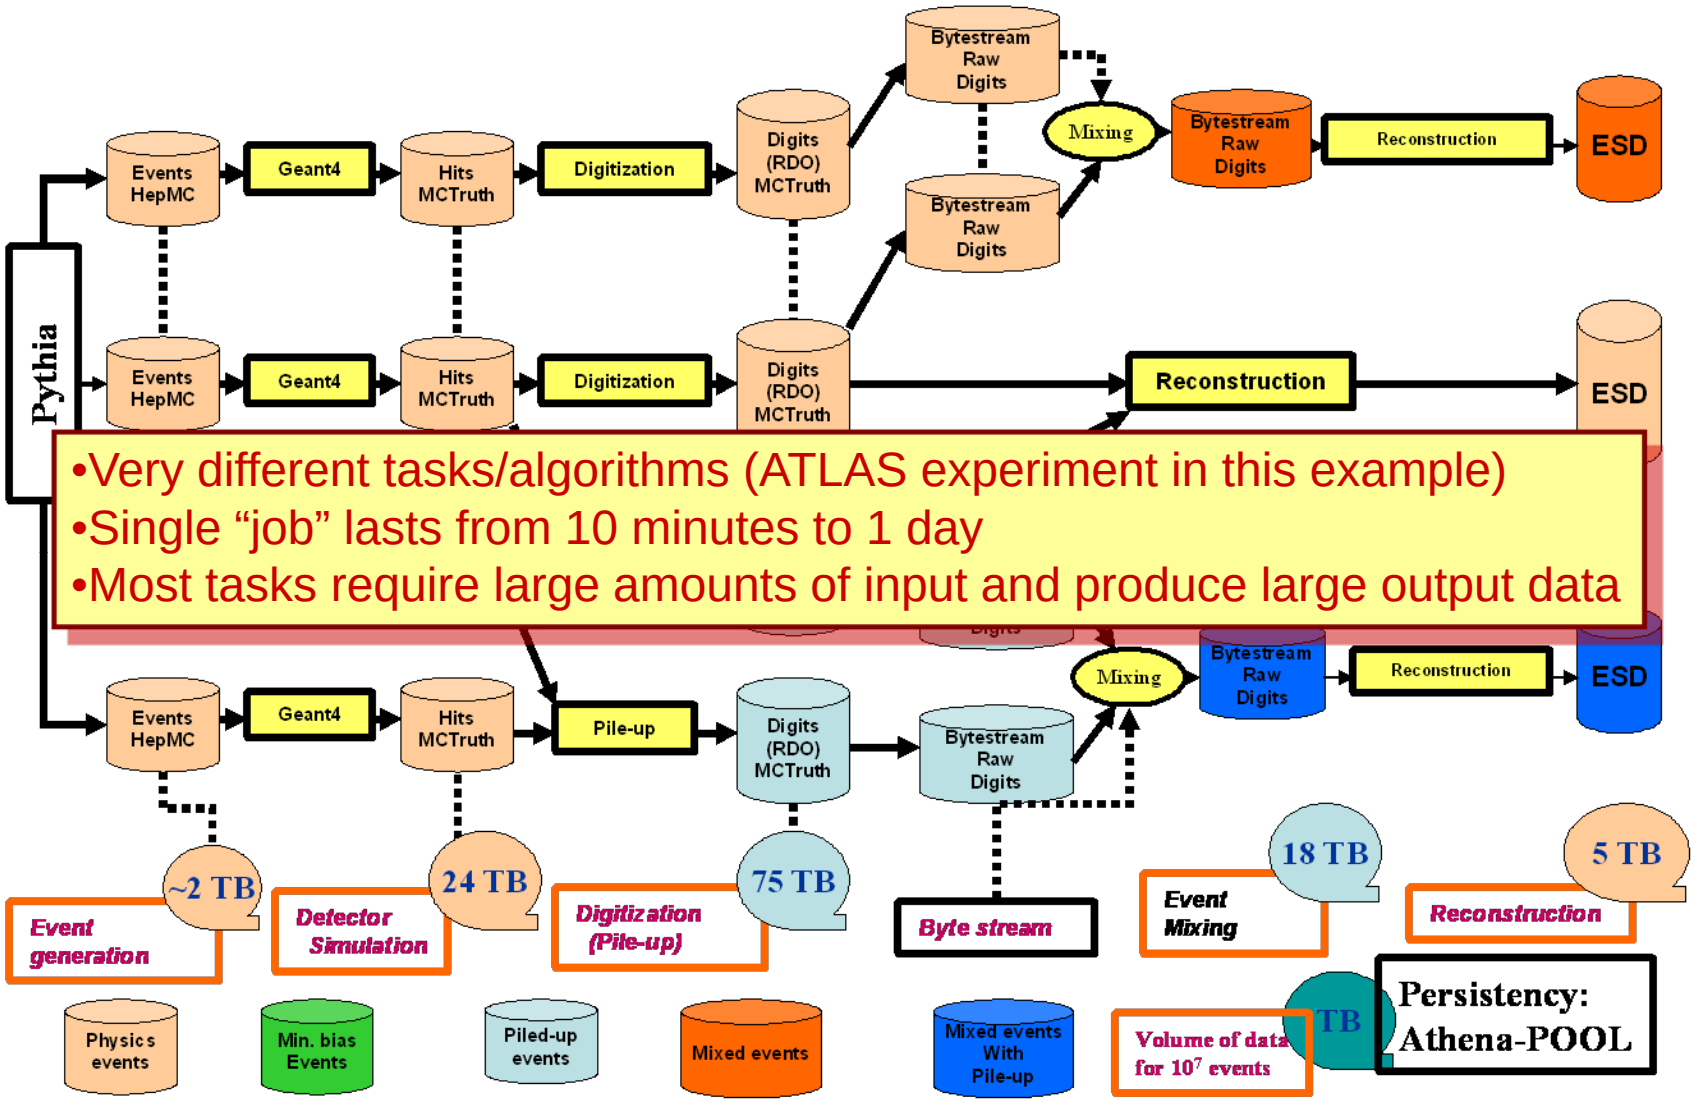

Software for HEP experiments

- Massive pieces of software
  - Written by very many different authors in different languages (C++, Java, Python, Fortran)
  - Dozens of external components
  - Occupy as much as ~10 GB of disk space each release
- Frequent releases
  - Every experiment can release as often as once a month during the preparation phase (which is *now* for LHC)
- Difficult to set up outside the lab
  - Experiments can not afford supporting different operating systems and different computer configurations
- For a small university group it is very difficult to manage different software sets and maintain hardware
  - ALICE, ATLAS, PHENIX etc – all in many versions
  - Solution: use external computing resources

HEP computing specifics

- Data-intensive tasks
  - Large datasets, large files
  - Lengthy processing times
  - Large memory consumption
  - High throughput is necessary
- Very distributed resources
  - Distributed computing resources of *modest* size
  - Produced and processed data are hence distributed, too
  - Issues of coordination, synchronization and authorization are outstanding
- HEP is by no means unique in its demands, but we are first, we are many, and we ***badly*** need it

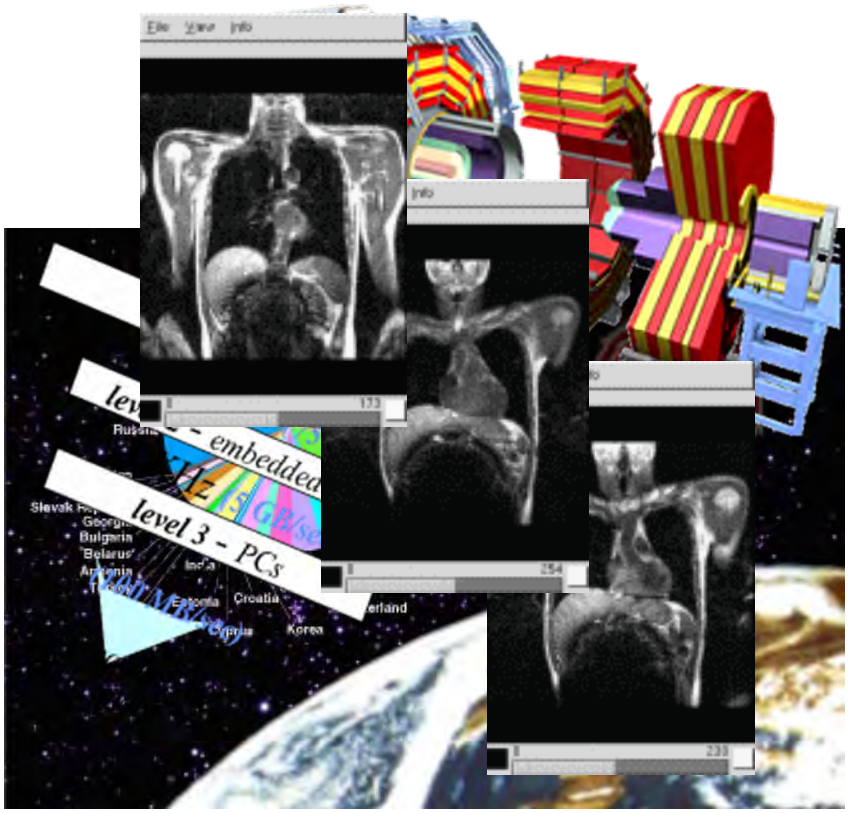

New technology: the Grid

- Proposed in the USA by Ian Foster and Carl Kesselman around 1997
- Refers to computing grids as analogy of power grids
  - Many producers
  - Competing providers
  - Simple for end-users
- Spelled "grid" or "Grid"
  - Except in French: "Grille de calcul"

Grid is a result of IT progress

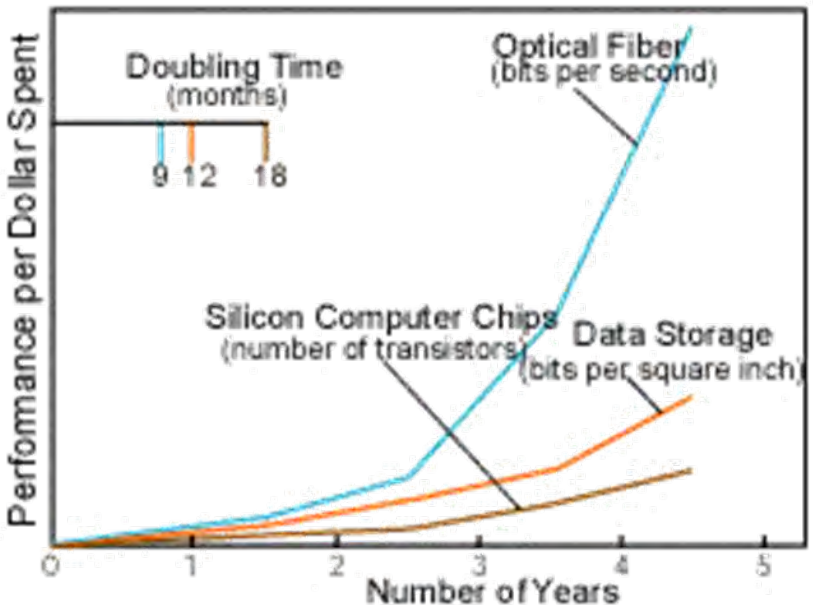

- Network vs. computer performance:
  - Computer speed doubles every **18** months
  - Network speed doubles every **9** months
- 1986 to 2000:
  - Computers: 500 times faster
  - Networks: 340000 times faster
- 2001 to 2010 (projected):
  - Computers: 60 times faster
  - Networks: 4000 times faster

Bottom line: CPUs are fast enough; wide area networks are very fast – gotta make use of it!

Slide adapted from the Globus Alliance

What can Grid do for us

- Distributed supercomputer, based on commodity PCs and fast WAN
- Access to the great variety of resources by a single pass – certificate
- A possibility to manage distributed data in a synchronous manner
- A new commodity

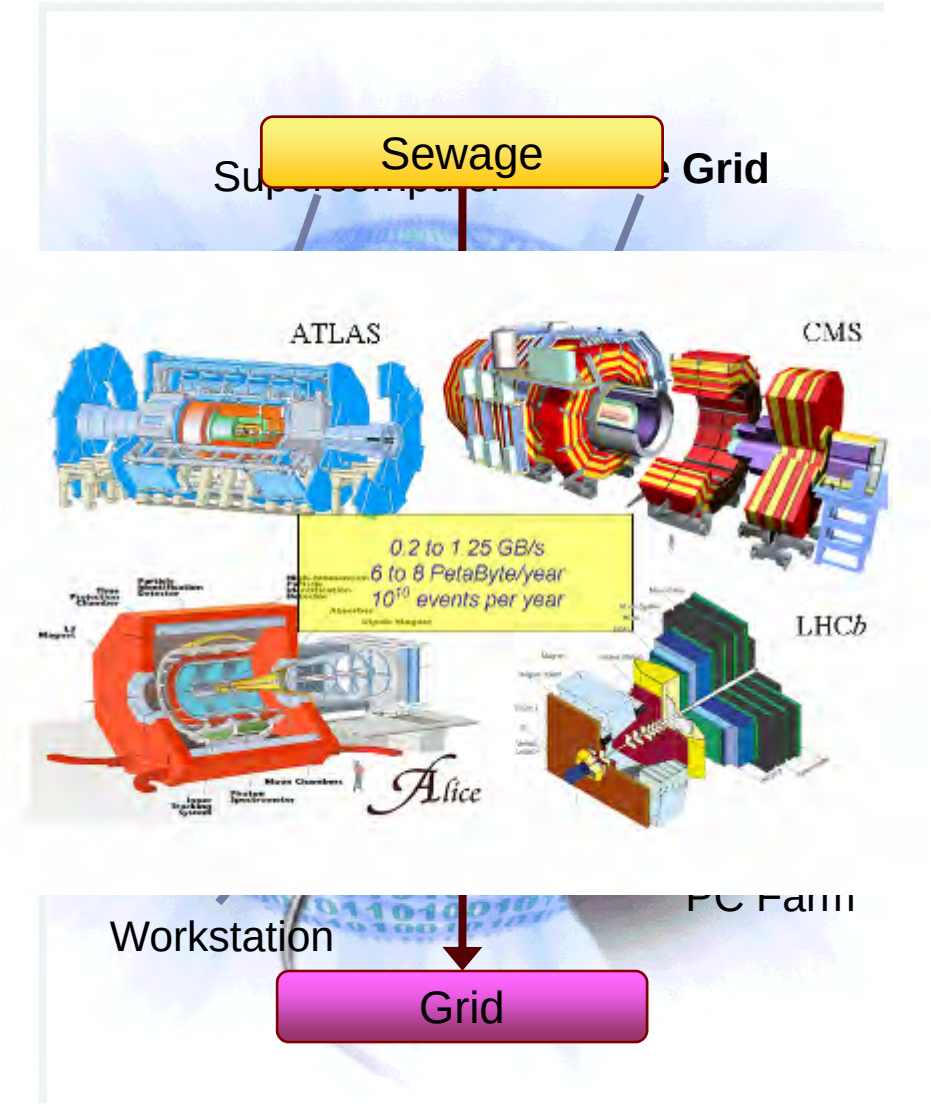

From distributed resources …

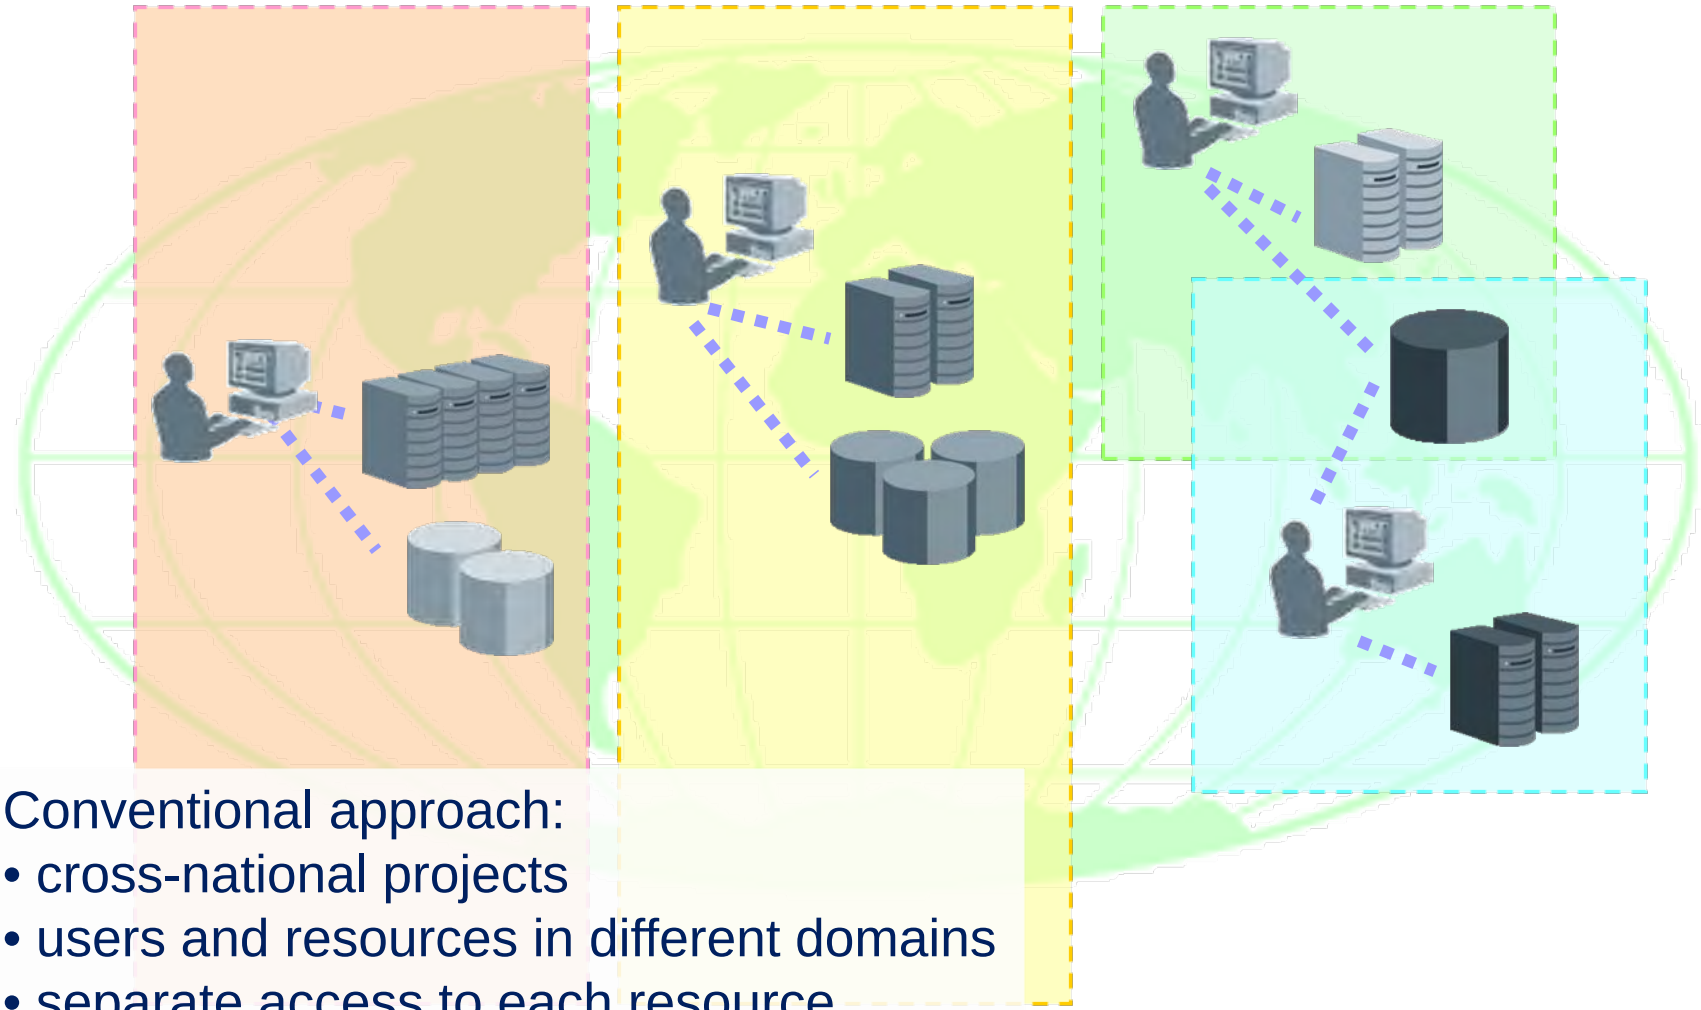

Conventional approach:
- cross-national projects
- users and resources in different domains
- separate access to each resource

… to World Wide Grid

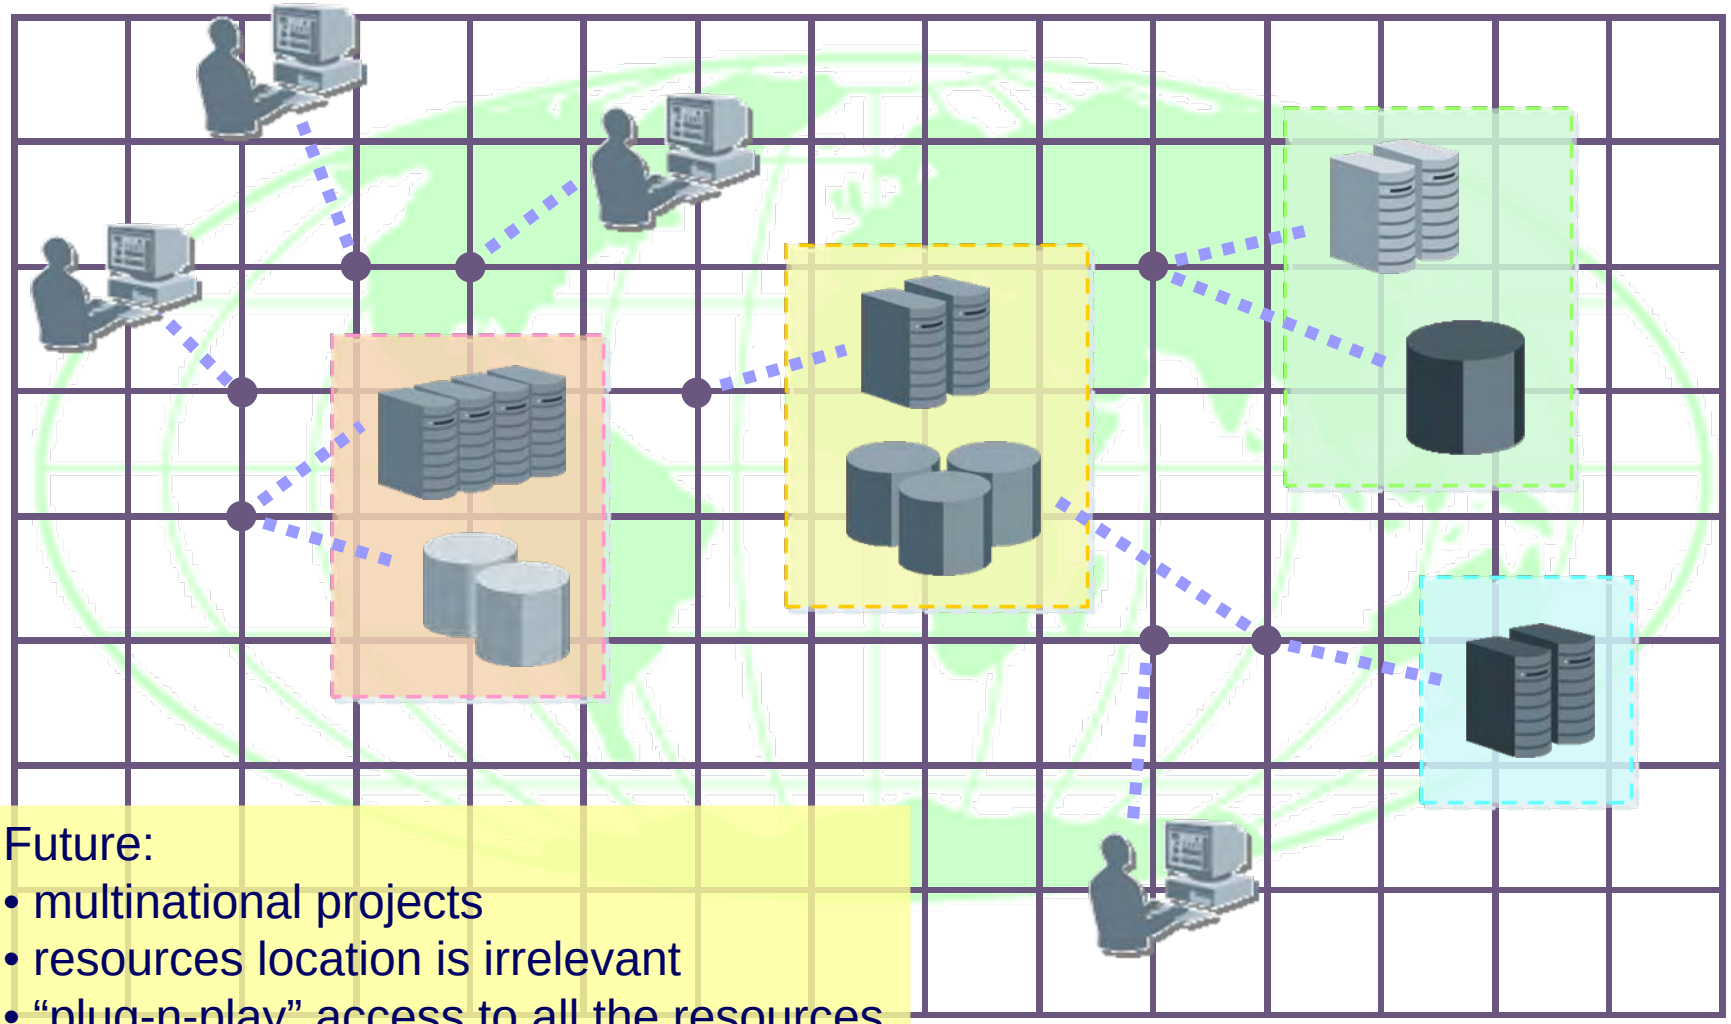

Future:
- multinational projects
- resources location is irrelevant
- “plug-n-play” access to all the resources

Global Grid Community

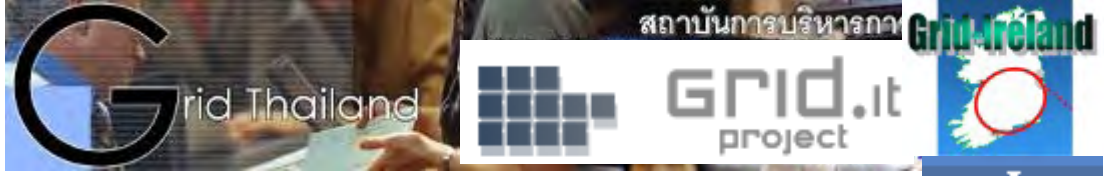

K*Grid
moreDream
KMI
Testbed
GFK
APEC APGrid
Access Grid
GNOC
Grid NOC

Slide originally byVicky White, FNAL

Grids in LHC experiments

- Almost all Monte Carlo and data processing today is done via Grid
- There are 20+ Grid flavors out there
  - Almost all are tailored for a specific application and/or specific hardware
- LHC experiments make use of only 3 Grid flavors:
  - gLite
  - ARC
  - OSG

- All experiments develop own higher-level Grid middleware layers
  - ALICE – AliEn
  - ATLAS – PANDA and DDM
  - LHCb – DIRAC
  - CMS – ProdAgent and PhEDEx

How did Grid appear in Lund

- Back in 2001, High Energy Physics Institutes from Scandinavia wanted to share their computing resources and jointly contribute to CERN/LHC computing
  - We needed Grid
  - The Grid hype just begun…
  - … and we created a NorduGrid project (Lund, Uppsala, Copenhagen, Oslo, Helsinki and many others)
- No production ready grid software (***middleware***) was available or seen on the horizon in fall 2001
- In February 2002, NorduGrid boldly decided to develop own Grid middleware
  - Was baptized ARC, for Advanced Resource Connector
  - Since May 2002 ARC is extensively used in ATLAS production and other scientific computing projects
- Now ARC is used to make a distributed computing center for High Energy Physics: the NDGF "Tier1"

Nordic DataGrid Facility

- Provides a unique distributed "Tier1" center for CERN
- Involves 7 largest Nordic academic HPC centers
- Plus a handful of University centers
- Connected to CERN directly with GEANT 10GBit fiber
- Inter-Nordic shared 10Gbit network from NORDUnet
- A dedicated 10Gbit LAN covering all the sites to be deployed in 2008

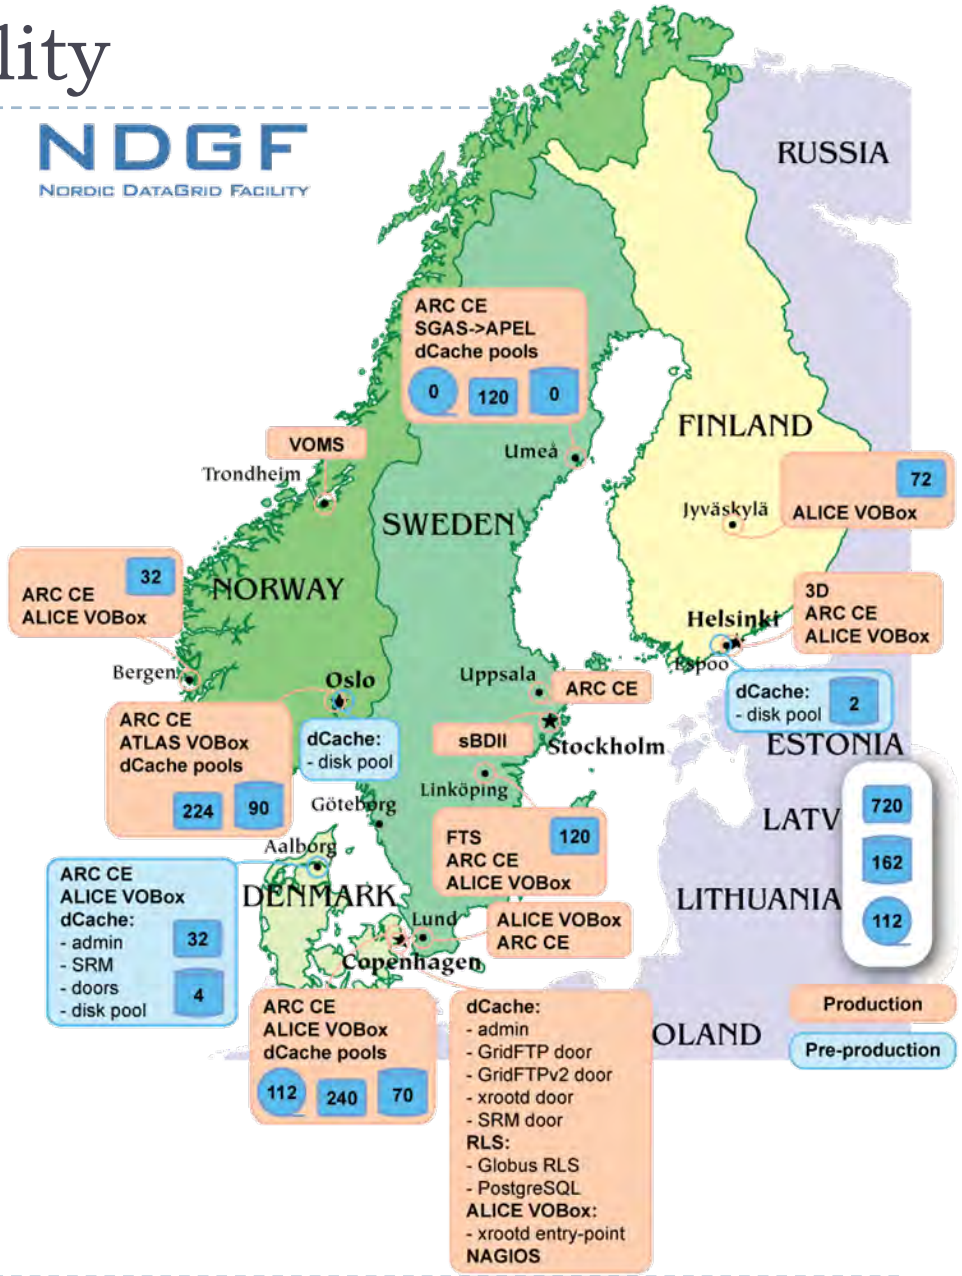

What is an LHC "Tier1" center

- WLCG: Worldwide LHC Computing Grid
  - A CERN project aiming to provide HEP computing infrastructure
  - Tiered structure: Tier0 at CERN, a dozen of regional Tier1s, many local Tier2s etc
- WLCG Tier1 is primarily a set of services:
  - 24/7 on-call support system
  - Infrastructure: network, power, cooling, safety etc
  - Authorization, specific software for entire multinational VOs
  - Job submission interface
  - Data indexing service
  - Storage resource management interface
  - File transfer services between Tiers
  - Experiment-specific interfaces ("VOBoxes")
  - Database service
  - Other: information system, monitoring, logging etc

ATLAS Multi-Grid Infrastructure

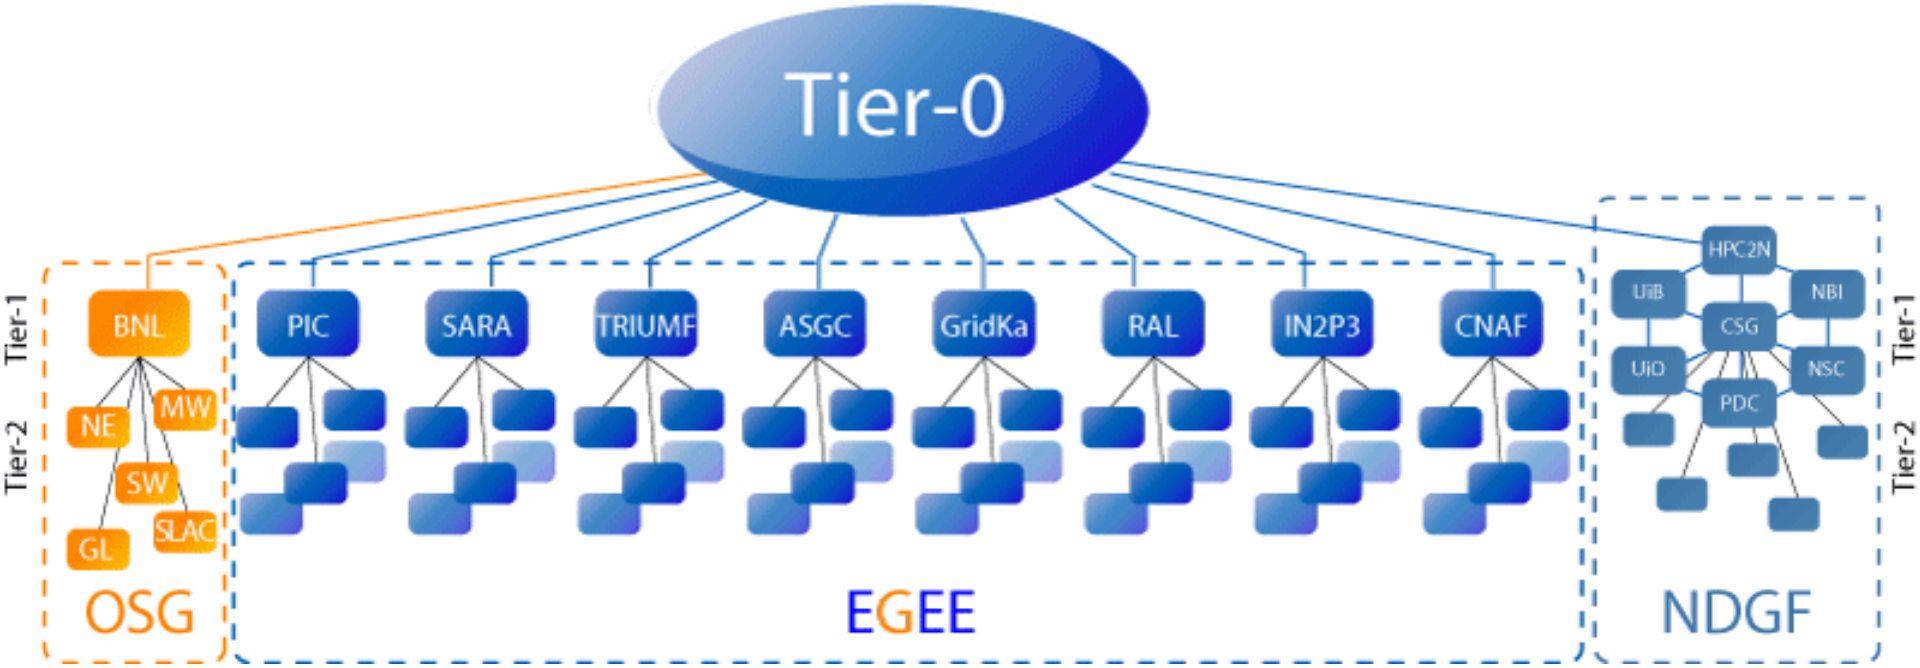

Graphics from a slide by A.Vaniachine

How to build a Grid: ARC in a nutshell

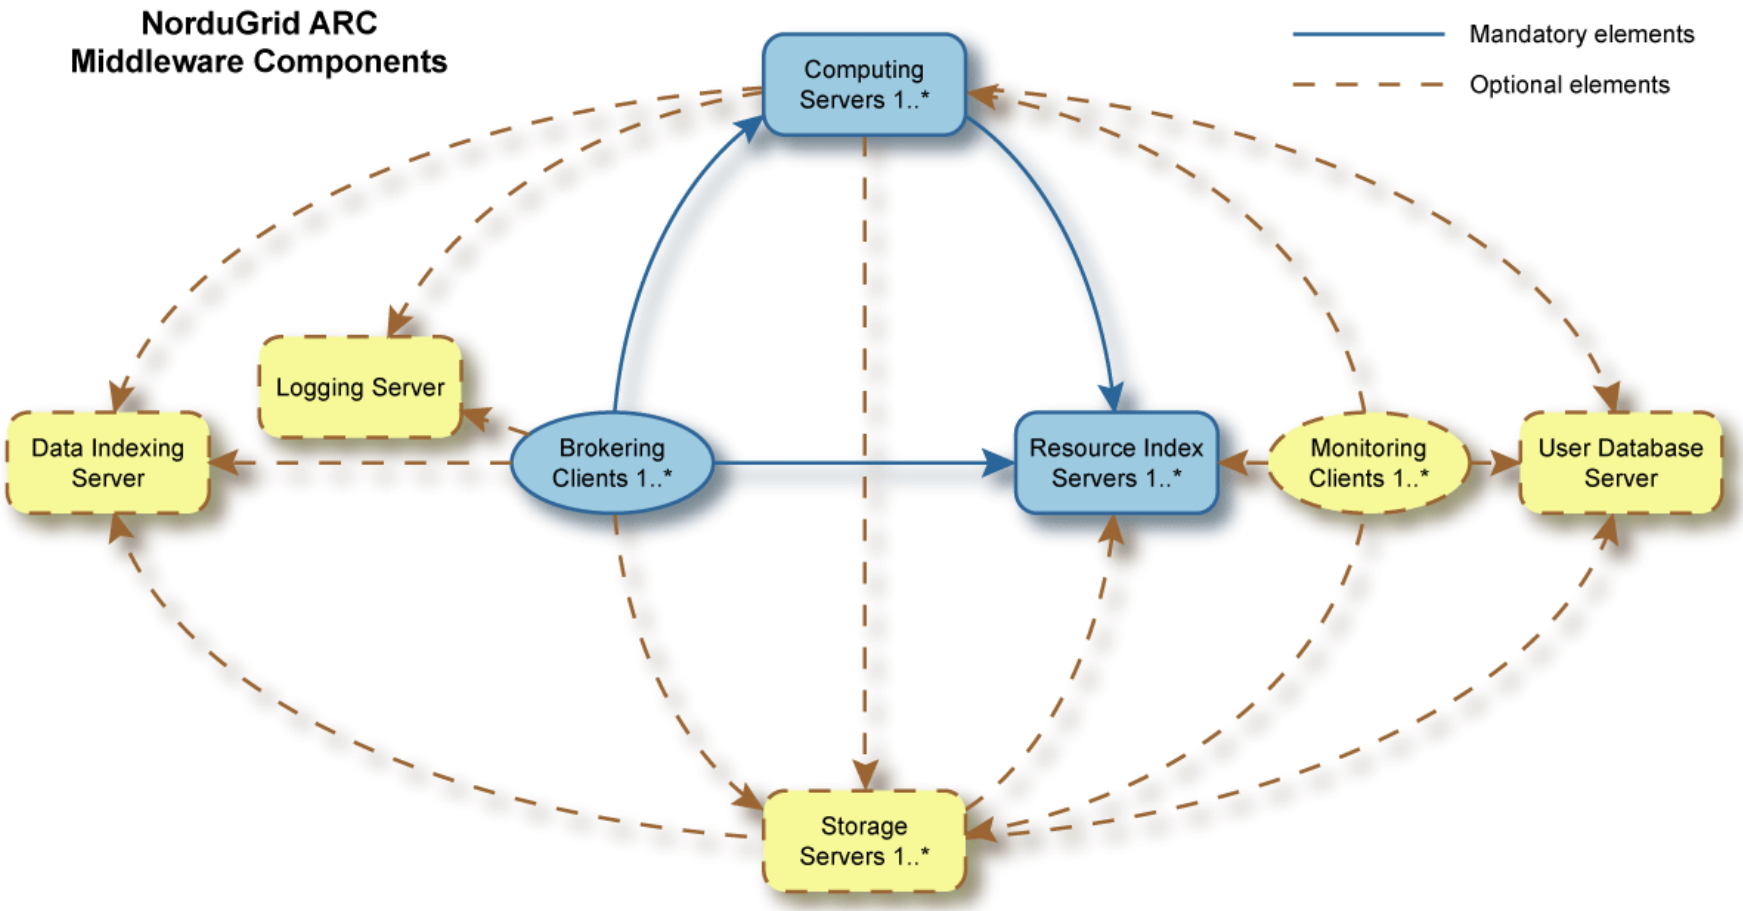

Goal: no single point of failure

ARC and ATLAS

- Via NorduGrid/ARC, Nordic countries contribute to ATLAS Data Challenges since 2002
  - Resources donated to ATLAS by national Grid projects, enthusiastic owners
    - Highly heterogeneous (OS: Fedora Core N, Red Hat, Debian, Gentoo,...; LRMS: PBS/Torque, Condor, SGE…)
    - No common policies enforced
    - Loosely coupled
  - Currently, ca 10% of ATLAS production tasks
    - Only 2 persons in charge of the production
    - Highest resource usage efficiency, reliability
  - Accumulated ~40TB of ATLAS data in ~50 locations
    - Includes e.g. Ljubljana – still, indistinguishable for jobs and outside users

ATLAS Monte Carlo production with NorduGrid/ARC

Jobs/day

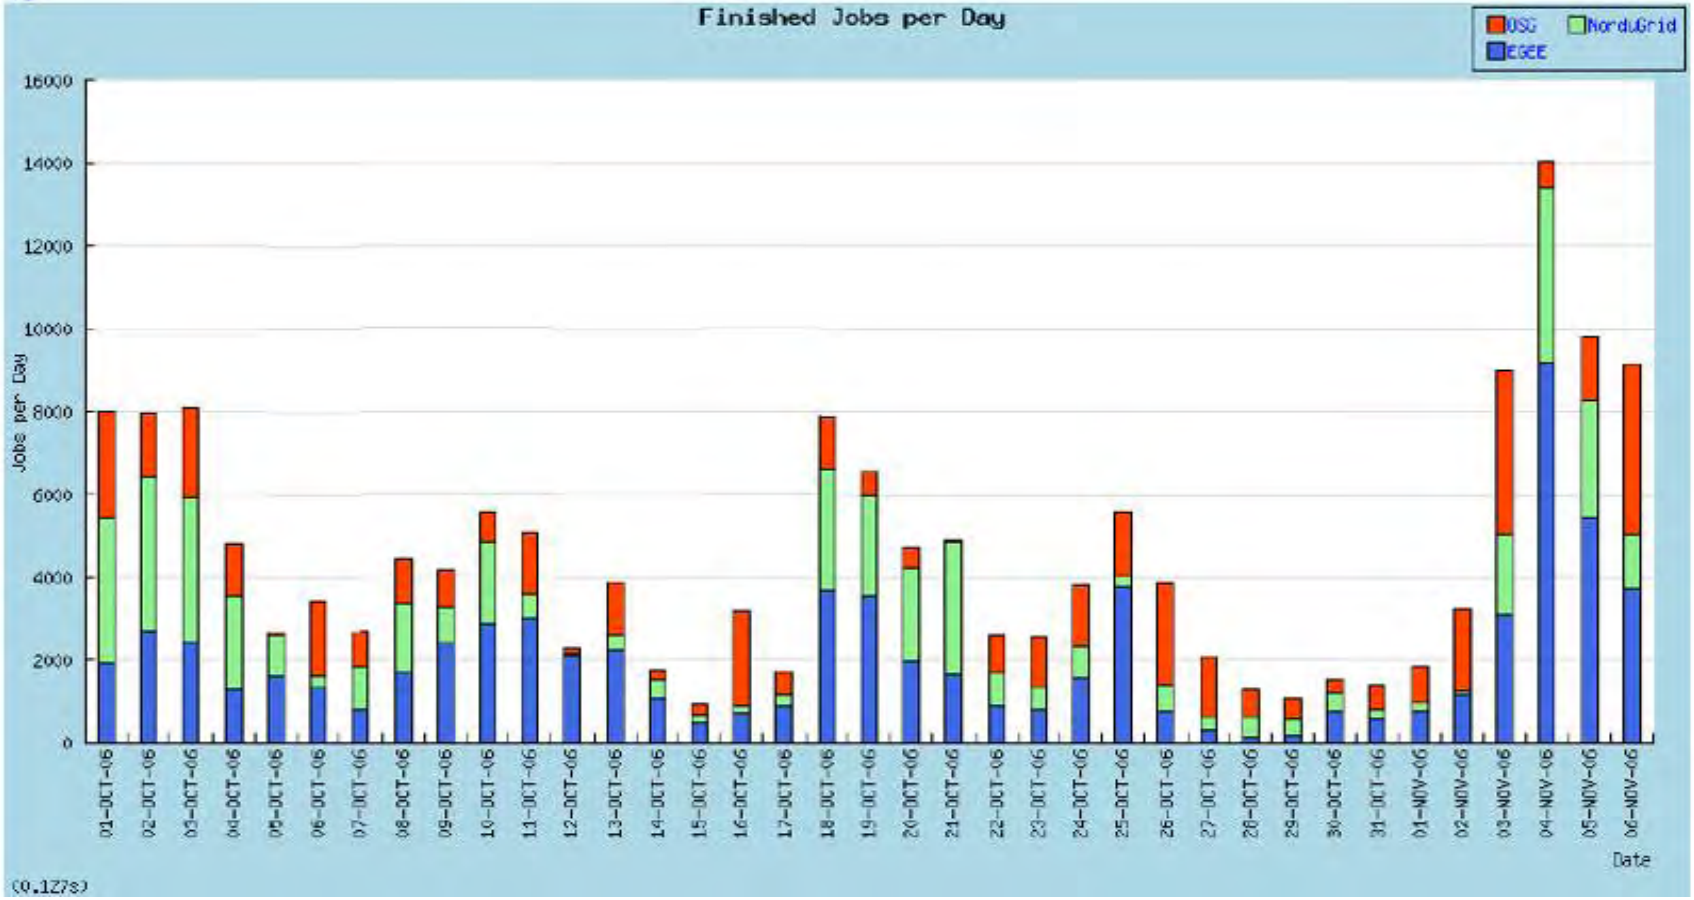

Where do I start on my way to Grid?

- Ask local sysadmin to install a Grid client
  - ARC standalone client is the easiest: you can install it yourself
- Apply for a "passport": the Grid certificate
  - Every country has a Certificate Authority
- Ask a knowledgeable person which Grid is adopted by your collaboration/group
- Apply for a "visa": become an appropriate Virtual Organization (VO) member
- Read the manual

Conclusion

- Particle physics community is the major consumer of Grid technologies
  - Every HEP researcher sooner or later will have to learn Grid basics
  - HEP community invests massive efforts into Grid development
    - If Grid won't help, it is unclear what would be the "backup solution"
    - The data will eventually be processed, the question is – how soon and how accurate
- Many other sciences are on-looking
  - Bioinformatics and radioastronomy appear to be the next in line
    - Huge data volumes, trivially parallel processing, distributed user base
- Grid could be the next big thing introduced by particle physicists after the World Wide Web