Current Grid operation and future role of the Grid

Oxana Smirnova
Lund University / NeIC
CHEP12, New York, 22 May 2012

using input from various Grid organizations

Context: HEP

HEP communities kick-started Grids, and remain dominant users

Context: HEP

- Fact: Grid is an integral part of most HEP experiments
- Fused with generic research infrastructures

- HEP "vendor lock": Scientific Linux etc
- Operational homogeneity obscures technical deficiencies

- No ownership homogeneity
- Different administrative domains
  - National projects are often a priority

- Not everything called "Grid" actually uses Grid technologies
  - For purposes of this talk:

> Grid is a federation of heterogeneous conventional systems, enabled by fast networks and a middleware layer that provides single sign-on and delegation of access rights through common interfaces for basic services

- No guarantee it suits all other researchers

HEP Grid operation model: hierarchy

- Originally motivated by relatively slow networks: data processing was faster than data movement
  - Grid paradigm itself was motivated by fast networks
- Tiers are often coupled to national Grid projects

Tier1: top service levels

Tier2: lower service levels

- Technically, hierarchy is not needed
  - Mostly operational necessity: Regional Operating Centers (ROC)
  - Really fast networks are still too expensive

Real operations: complexity

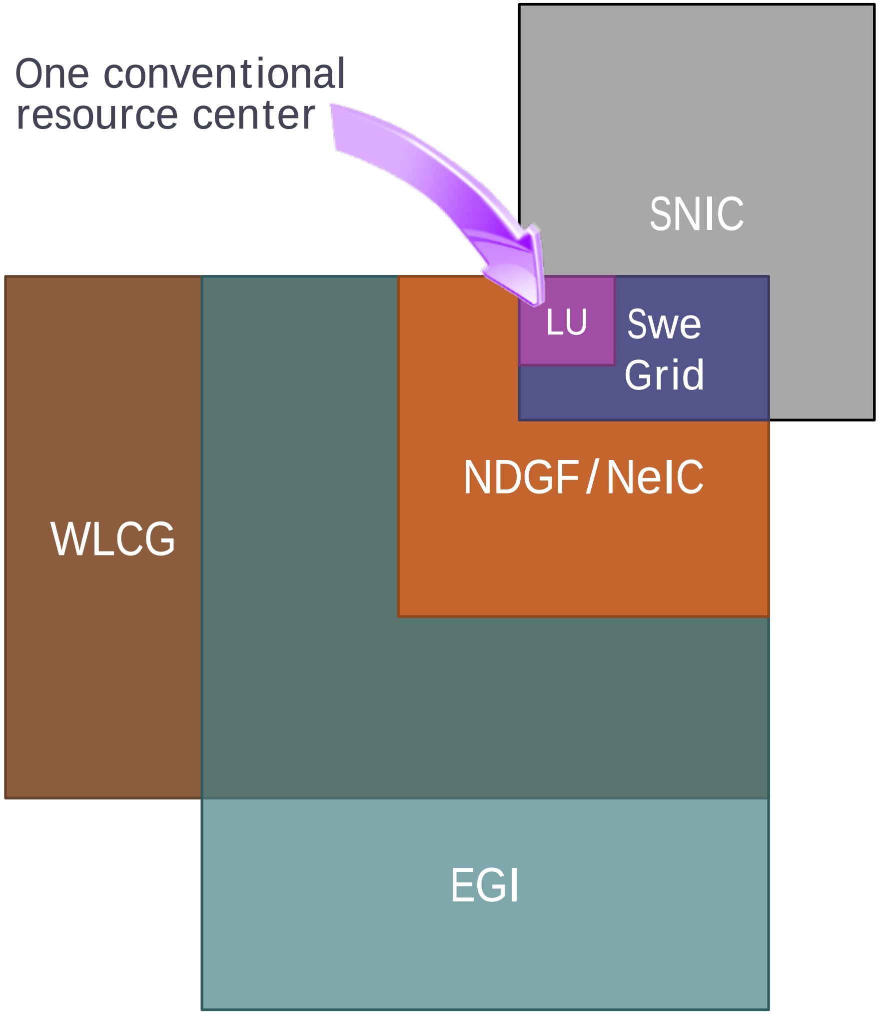

- An example: a Linux cluster in Lund University is a part of
  - National infrastructure (SNIC)
  - National Grid Infrastructure (SweGrid)
  - Regional infrastructure (NeIC)
  - Continental Grid infrastructure (EGI)
  - Worldwide Grid infrastructure (WLCG)
- Each infrastructure is not fully uniform
  - Different contributors
  - Different users/applications
  - Different middlewares
  - Different policies and priorities
  - Different funding cycles
  - And they overlap
- Each user expects top service levels
  - Hardware failures affect them all

A small Grid infrastructure: NDGF

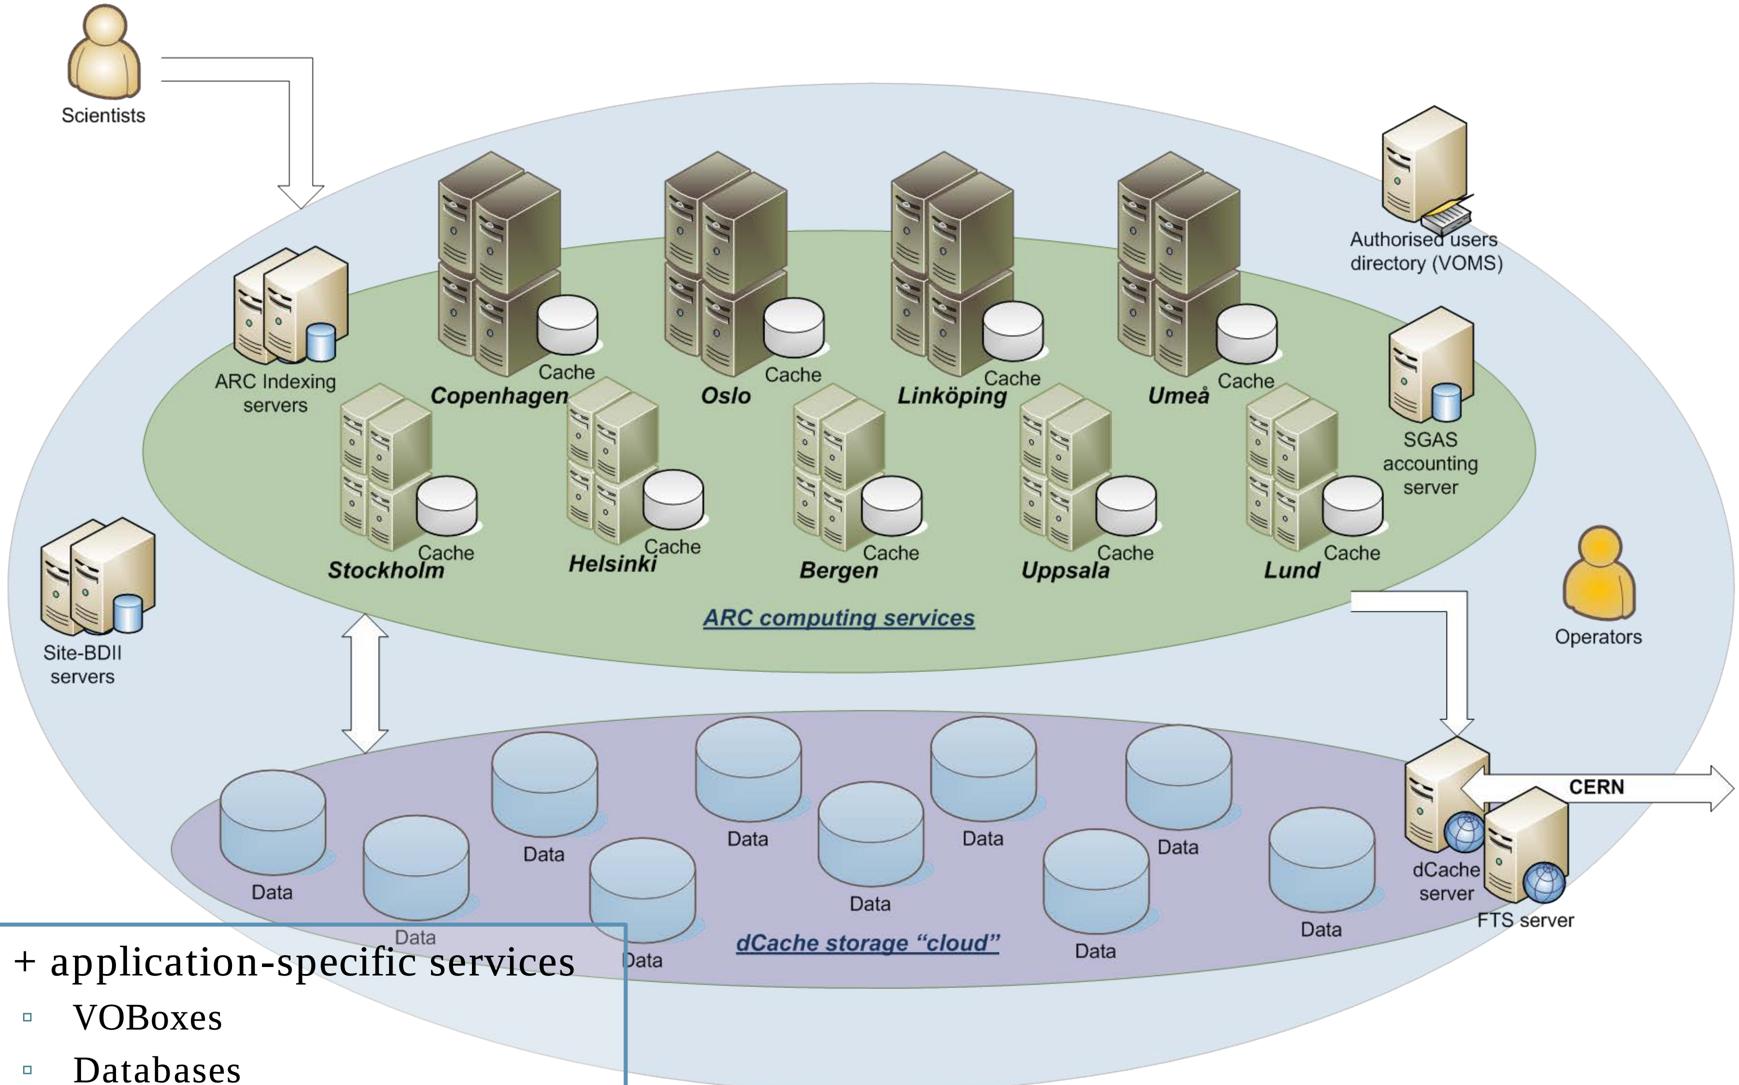

- + application-specific services
  - VOBoxes
  - Databases
  - etc

Larger picture: WLCG Grid Sites

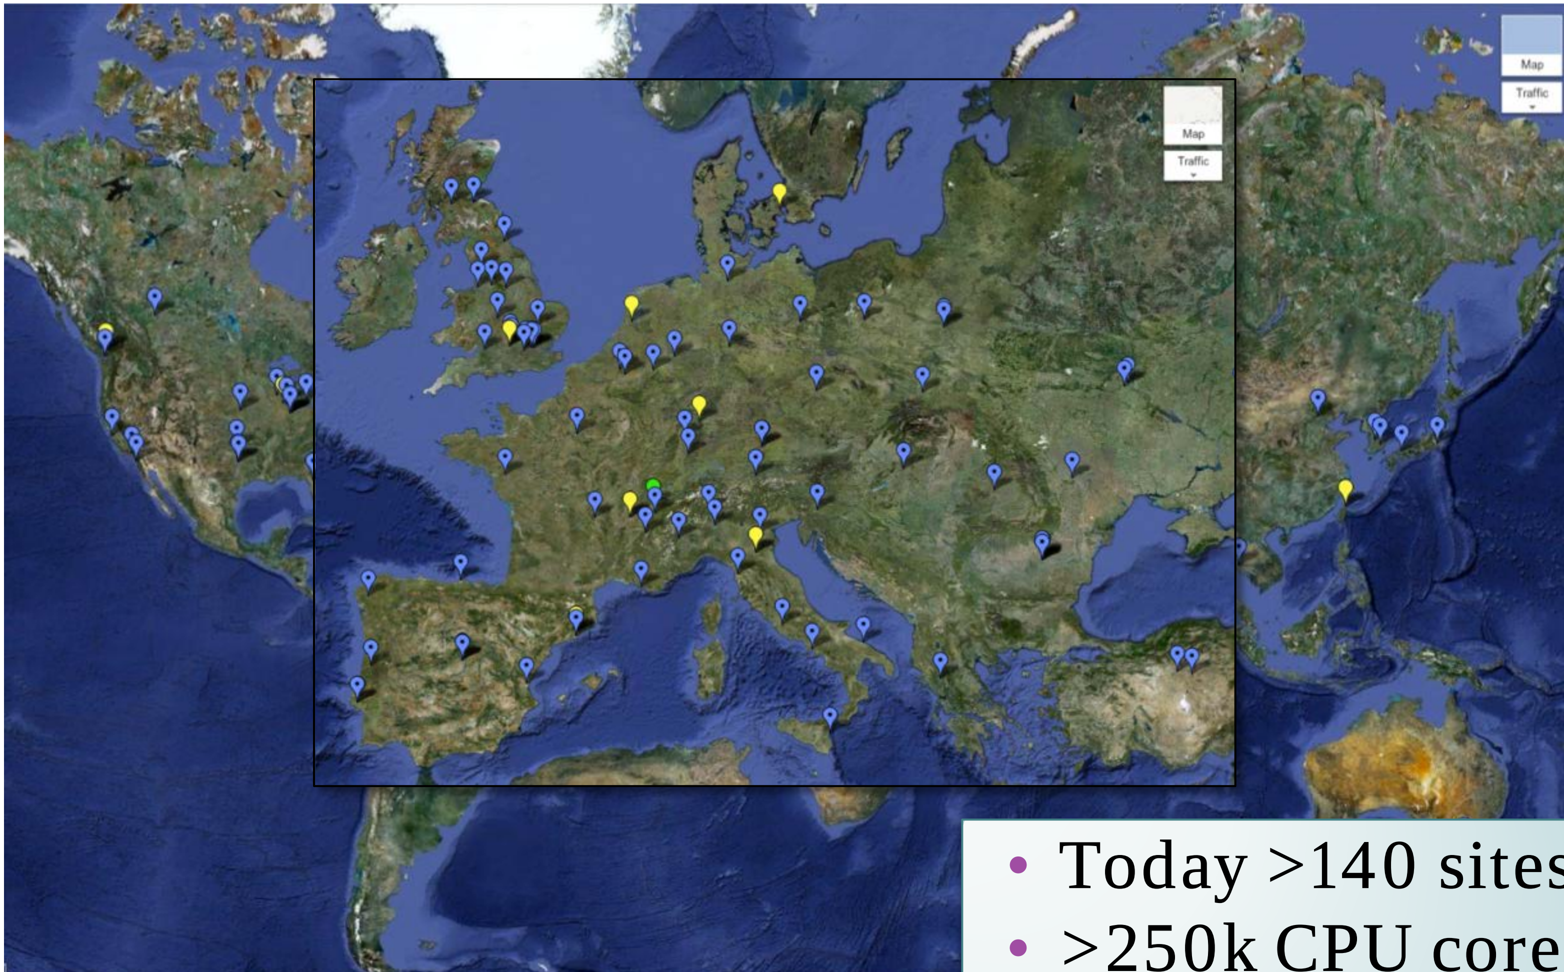

- Today >140 sites
- >250k CPU cores
- >150 PB disk

Tier 0 Tier 1 Tier 2

WLCG usage: continues to grow

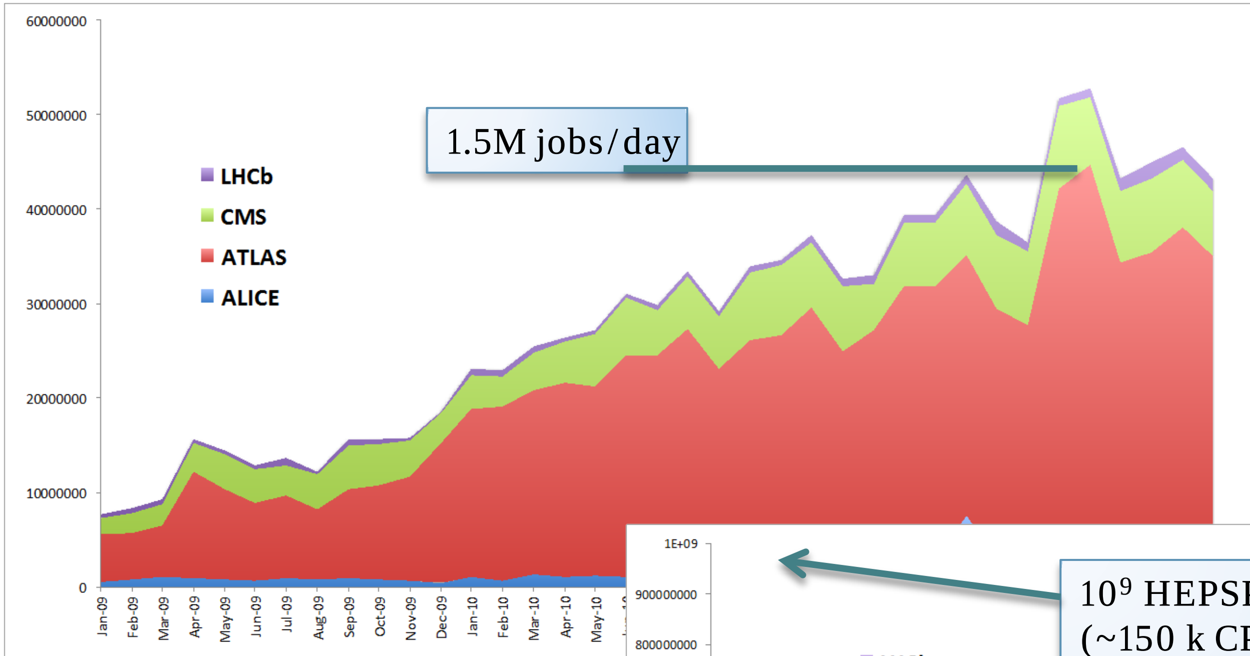

- WLCG usage pattern:
  - Continuous
  - Ever increasing load
  - Some spikes
  - Mirrored by contributing national Grids

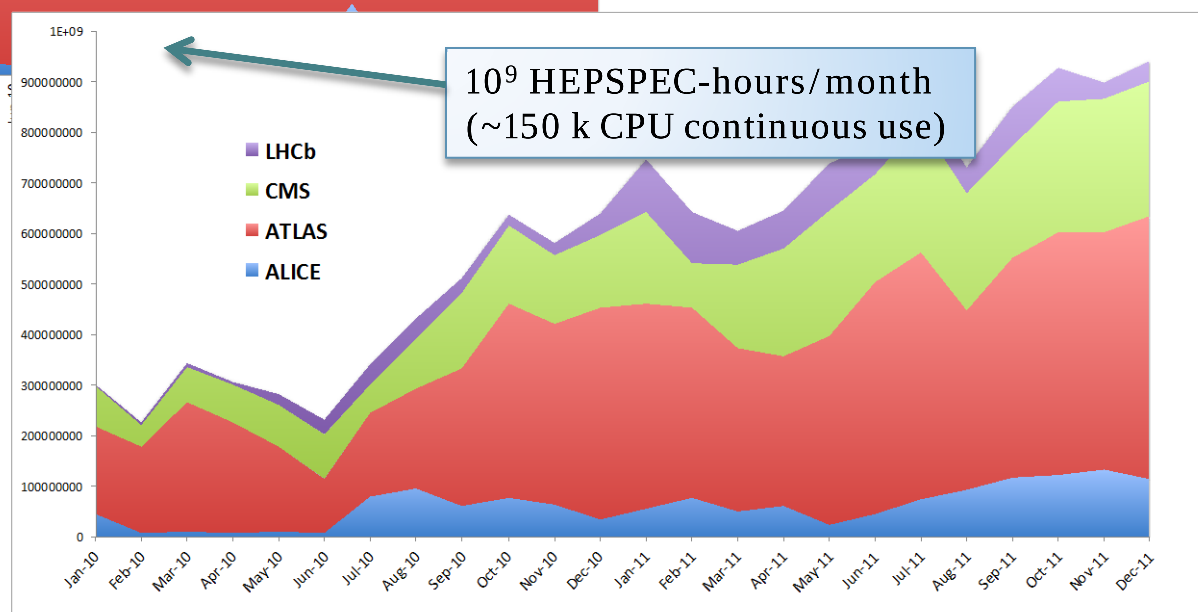

- Evolution:
  - Being discussed by *Technical Evolution Groups*
  - Operations is one of concerns

Multi-science Grids

Grid user communities are very many, but few are as large as HEP to afford own resources, so they use existing infrastructures

Beyond WLCG: more complexity

- Europe: EGI-rooted hierarchy
  - Centralized operations relying on ROCs and National Grid Initiatives (NGI)
  - Actively engaged in collaboration with Asian, African and Latin American Grid initiatives
- Asia – Pacific: collaboration
  - NGI-based model is not always feasible
  - Successful model: APGI and APROC, affiliated with EGI
- Latin America: collaboration
  - LGI rooted in NGIs and Equivalent Domestic Grid Structures (EDGS)
  - Strong cooperation with Europe
- USA: federations
  - Infrastructures largely defined by technologies
  - May contribute to each other, like OSG contributes to XSEDE
- Canada: a consortium
  - Includes WestGrid
- Africa: projects
  - Mediterranean – collaboration via EUMedGrid
  - South Africa: established National Grid
- Australia and New Zealand: collaboration
  - ARCS National Grid

EGI Resource Centres (April 2012)

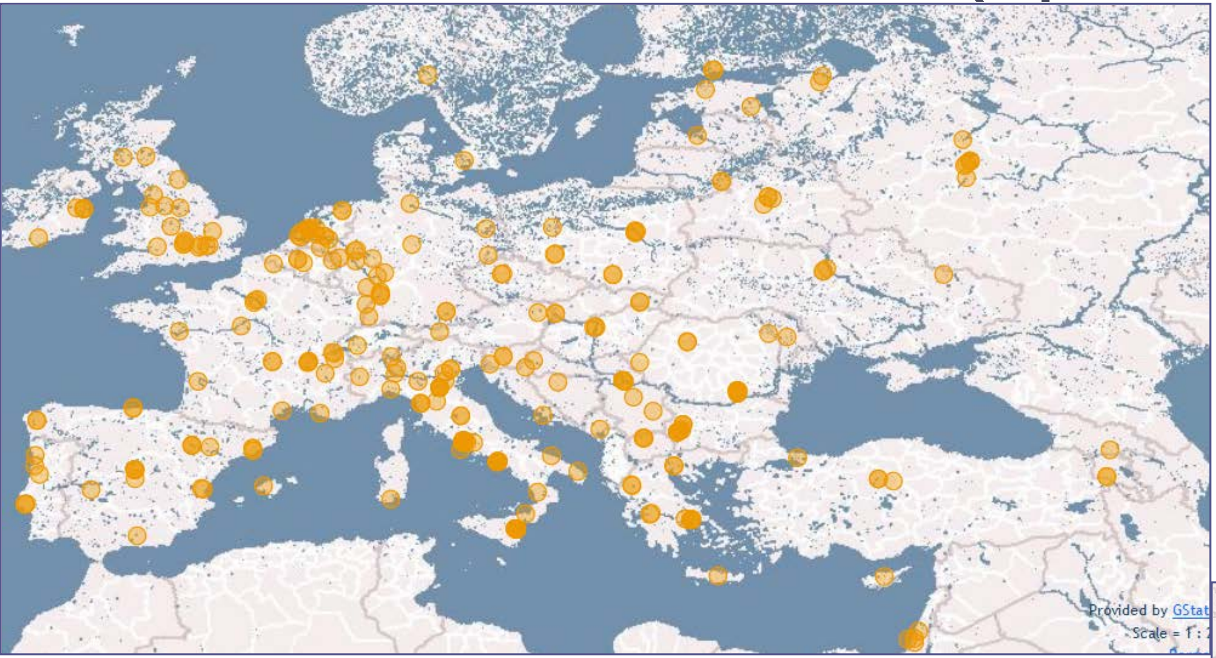

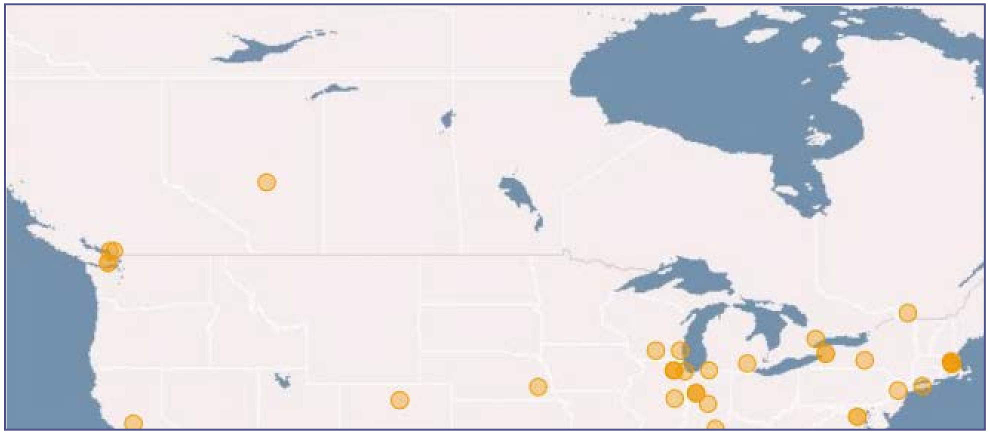

EGI Metrics (April 2012)

Metrics			Value (Yearly increase)
Installed Capacity	Logical CPUs	EGI-InSPIRE and EGI Council members	270,800 (+30.7%)
		Including integrated RPs	399,300
	HEP-SPEC 06	EGI-InSPIRE and EGI Council members and integrated RPs	2.96 Million (+49.5%)
	Storage	Disk (PB)	139 PB (+31.4%)
		Tape (PB)	134.3 PB (+50%)
Resource Centres	EGI-InSPIRE and EGI Council members		326
	Including integrated RPs		352
	Supporting MPI		90
Countries	EGI-InSPIRE and Council members		42
	Including integrated RPs		54
Performance	Monthly Availability/Reliability		94.50%/95.42%
Utilization	HEP-SPEC 06 Hours		10.5 Billion (+52.91%)
	Jobs		492.5 Million Jobs /year 1.35 Million Jobs/day (+46.42%)

OSG resources

- Resources accessible through the OSG are contributed by the community
  - Their autonomy is retained.
  - Resources can be distributed locally as a campus infrastructure
- >100 sites
- >70,000 cores accessible
- >30 research communities

XSEDE partners

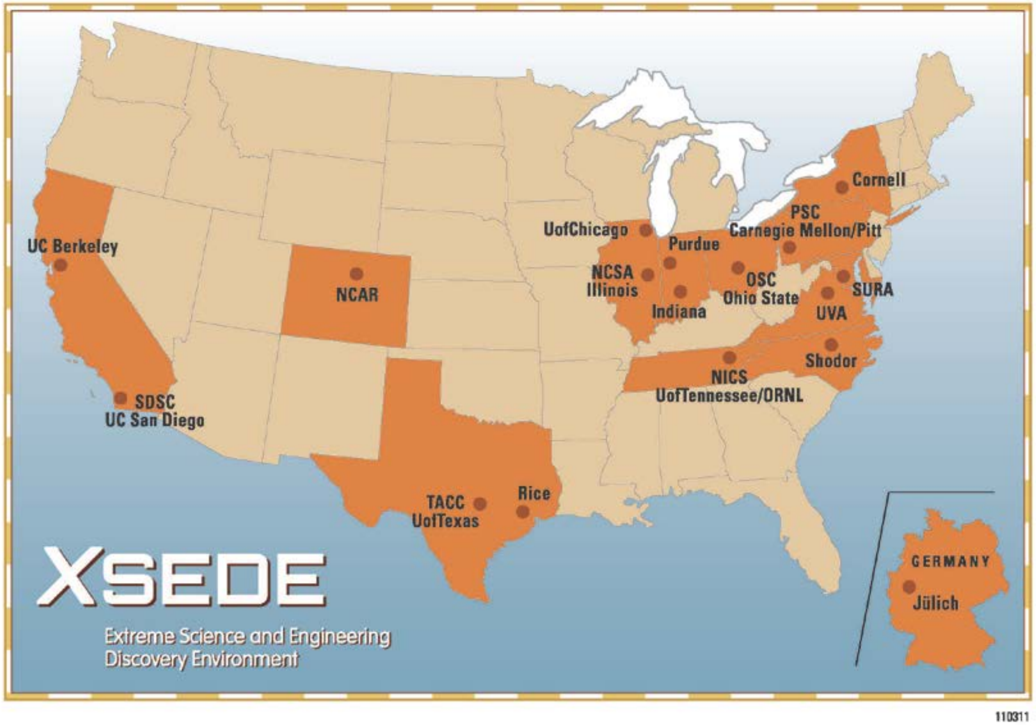

- 16 supercomputers and high-end visualization and data analysis resources
- HTC resources: OSG, Condor Pool
- >10000 scientists

Latin American Grid Sites

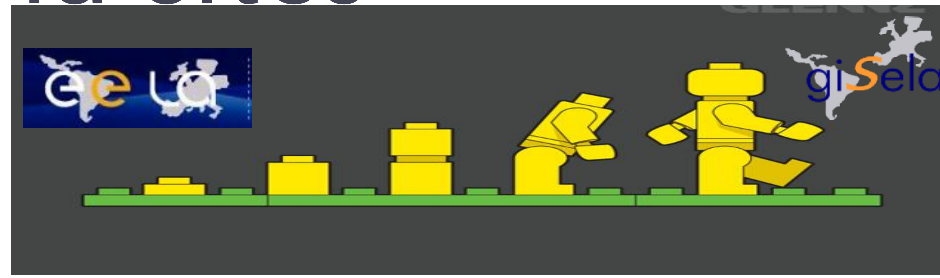

--31/08/2012

GISELA

--01/09/2010

--31/03/2010

EELA-2

--01/04/2008

--31/12/2007

EELA

--01/01/2006

- 18 Resources Centres provide skilful personnel to support and develop services in the region
- 435 trained users / 73 inst.
- 1794 cores platform
- Operation 24/7

GISELA: 34 production applications inherited from EELA / EELA-2 represent potential communities

GISELA e-Infrastructure inherited from EELA and EELA-2

Africa & Arabia Grid sites

Analysis

All Grid infrastructures face similar challenges

Mixed messages

- Current Grid operation is a mixture of:
  - cross-national application-specific domains (WLCG, EUMedGrid etc)
  - generic national domains (USA, Canada etc)
  - overall coordinating bodies (EGI)
- Different levels of expertise within each domain and between domains
- Different service levels
- Different middlewares and other technologies, especially between domains
- Convergence, even within a domain, is not always easy

- Resource provider perspective:
  - Centralized operations are good: real-time monitoring, downtime handling, upgrades, issue tracking etc
  - Overall performance figures matter
  - Like security

- User community perspective:
  - Every failed job or file transfer is a catastrophe
  - Prefer to seek help within the community
  - Hate security

  - Nobody likes buggy middleware
  - Nobody likes middleware updates

Grid operations challenges

- Middleware design problems are solved by operational means
  - Service keeps acting up – put a babysitter next to it
- Hardware failures and software bugs are unavoidable
  - Minor event on a local scale can become a major global issue
- Non-uniformity and non-predictability
  - Resource owners prefer standard services and predictable usage patterns
  - Users prefer their own ways
- Difficult to coordinate maintenance slots and users' activities
  - Technically, all sites can shut down simultaneously
- Mixed policies: no common authority
- Complex service levels and critical services definitions
  - What is critical for one VO is non-important for others
  - Different monitoring levels
  - Incident response – different priorities
  - Users keep watch 24/7
  - Operators work 8/5 and use different tools
- Reliable operations are rather expensive

Theory vs practice

Theory

- Grid by nature is a technology to federate resources; viable Grid organizations must be based on federation principles
  - different owners, different consumers, different technology
  - standardized service levels and interfaces
- Data storage and handling is secondary to computing

Practice

- Many Grid organizations are still heavily rooted in HEP
  - Technologies and policies are largely driven by WLCG needs – de facto standards
  - Global Grid standardization pace has slowed down, lacking support from user communities
- Data storage and handling is the source of most problems

If we want to build real Grids

Do

- Collaboration
  - Federated operations
- Common standardization and convergence effort
  - Common forum
  - Rare exceptions and minor modifications are allowed
- Common practice Open Source development and distribution
- Common "exchange" to share excess resources
  - Somewhat academic, as there is no excess yet
- Take good care of data
  - Storage and access
  - Transfer

Don't

- Single-rooted hierarchy
  - Globalized monitoring and support
- Selected core technologies
  - Preferred systems or architectures
- Proprietary code and distribution via e-mail
- Privileged VOs
  - Everybody can be a user
  - Everybody can be a resource provider, too
- Prioritize computing over data
  - We can not re-use storage like we re-use CPUs

The Cloud elephant in the room

- Research e-Infrastructure sustainability relies on ever increasing need for computing and storage capacities
  - Public funding
  - Grid business model: countries invest **as much as they can**, researchers use **as much as they need**
    - Works as long as we are much cheaper than commercial providers
    - Keeps national funding inside countries

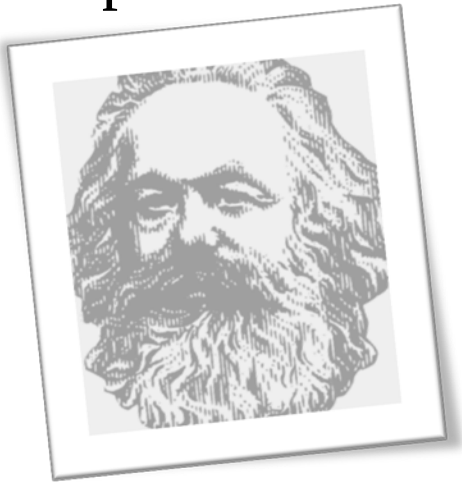

- Cloud technologies look appealing (Cycle Computing: 50000 cores á 5000 USD/hour)
  - If we trust commercial clouds with our data
  - If we can offer services comparable to commercial providers for lower costs
    - How?
    - Who?

Familiar patterns

Parallels with transportation services, though not exact, help putting the technology progress into context

Personal use – workstations

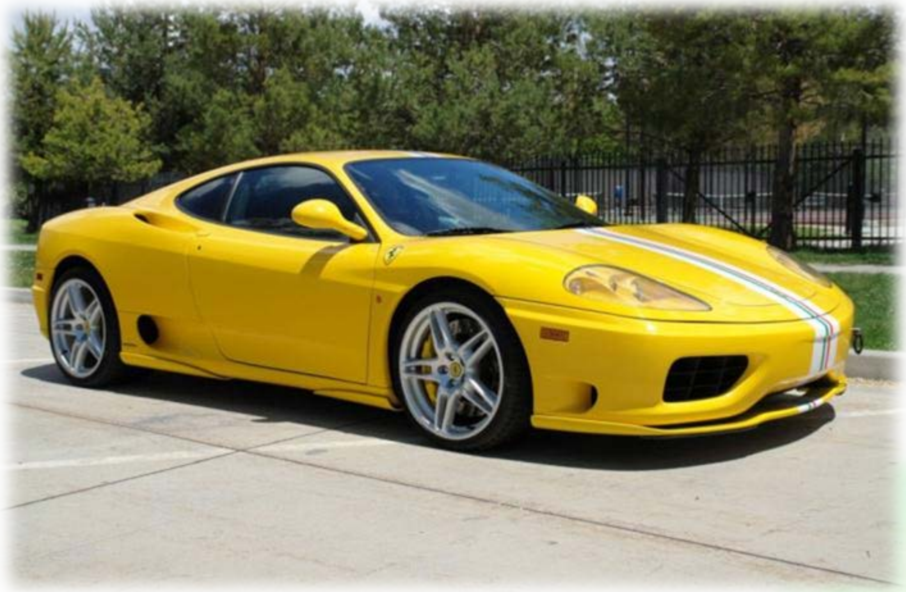

Everybody likes to have one or two

Customized shared service – Grids

There is always demand and supply

- Opinions 10 years ago: Grid will make local computing obsolete
- Reality: most scientific computing is still local
  - Successful distributed computing stories are so rare, they get press releases

Generic shared service – Clouds

Now exists for computing and data storage

- Opinions today: Clouds will make Grid obsolete
- Reality in future: Grid is likely to stay as long as diverse specialized facilities exist

Vision

Knowing past and present, we can try to predict future

Past, present and future

Past: Grid organizations fused with HEP research programs

Present: Grid organizations serve a variety of scientific communities

Future: Grid will mature into permanent national and international infrastructures

- We adopted the Grid, its future is our responsibility
  - Ignore standardization and common practices – Grid will become limited to HEP & Co
  - Serious investment in standard approaches – Grid will become useful for everybody

Further into future

All the mentioned technologies will be used but something else will come

- There will be use cases that will require different technologies
  - Currently, core infrastructure is the same for all existing distributed computing paradigms
- Truly new technologies will come with new core infrastructures
  - Different processing units, different network protocols, different media, different operating system principles etc

Conclusions

- Grid works (for ~1% of all scientists)
  - Allows LHC to achieve scientific results almost instantaneously
- Operations are complex and costly
  - Still immature middleware and faulty hardware
  - Highly customized very different application frameworks
  - Different resource ownership and service levels
- Clouds will not make Grids cheaper or redundant
  - But surely will add extra complexity
- Grid is here to stay
  - Scientific data will always be distributed
  - Global science is a collaborative effort, and so is Grid
- Standardization and convergence to common approaches is badly needed
  - Otherwise Grid efficiency will remain relevant only to few selected applications, like HEP
- Something totally different will certainly come