Managing Huge Data Sets with Grid Technology

Jon K. Nilsen
j.k.nilsen@fys.uio.no

Outline

- Large Hadron Collider
- Grid
- Grid Data Management
- Example
- Conclusion

Large Hadron Collider

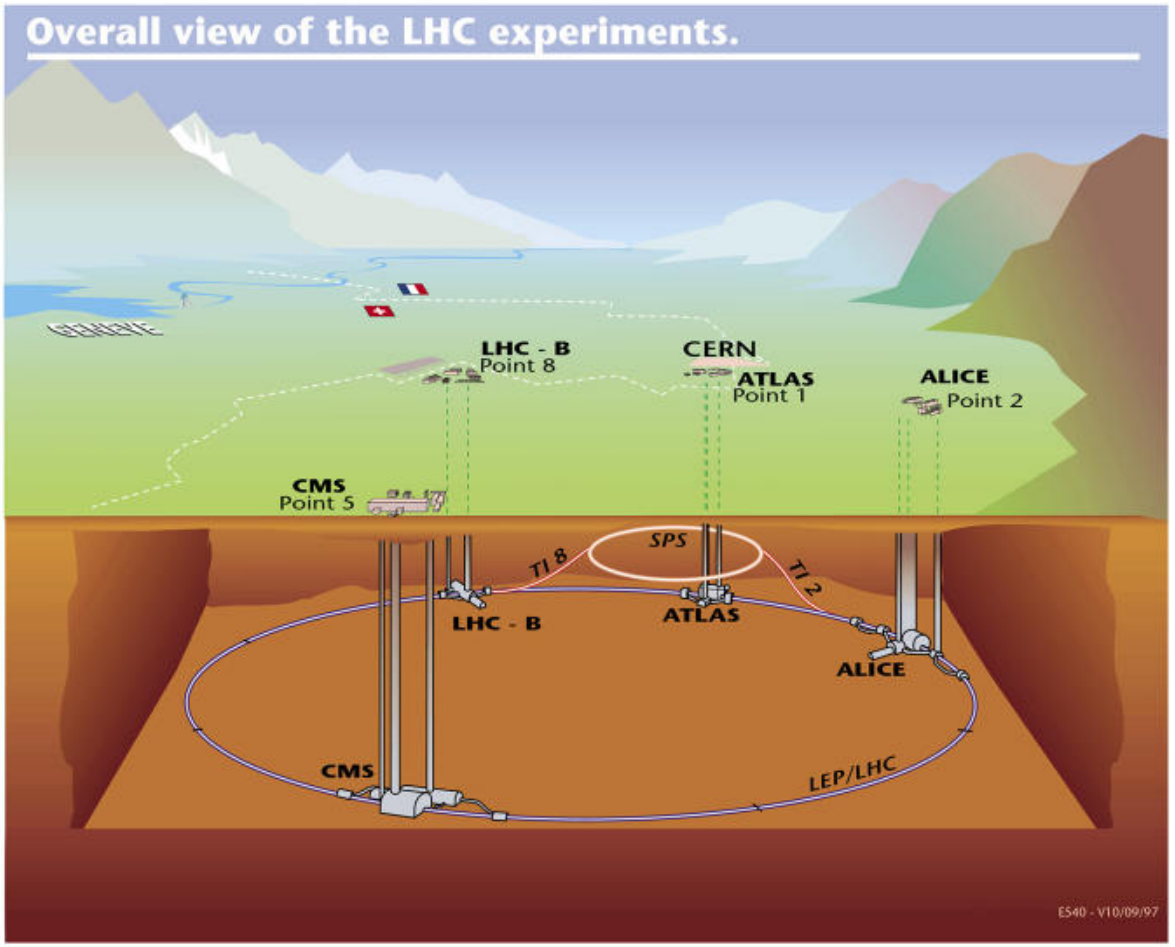

Atlas, LHC

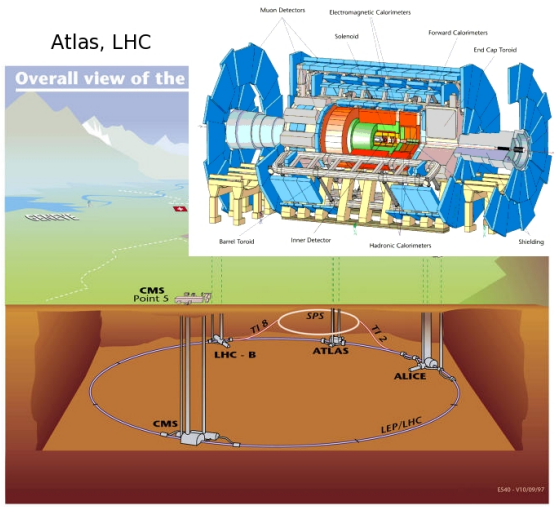

~320MB/s

Tier 0

~1GB/s

Tier 1

~20MB/s per T1

Tier 2

20% of real data

LHC

- LHC will produce ∼15 PB/year
- Everything needs to be safely stored and replicated
- Access must be provided for ∼5000 scientists in ∼500 research institutes and universities worldwide
- All data must be available over the estimated 15-year lifetime of LHC
- Needs to store (and replicate) more than 225 PB to disk and tape
- Too much to store on one site → need for Grid storage

Grid

- What's a Grid?

Grid

- What's a Grid?
- Checklist:
  - Coordinates resources that are not subject to centralized control
  - Use standard, open, general purpose protocols and interfaces
  - Deliver non-trivial qualities of service (the gain from combining resources must be significantly higher than just the sum of resources)

Grid

- What's a Grid?
- Checklist:
  - Coordinates resources that are not subject to centralized control
  - Use standard, open, general purpose protocols and interfaces
  - Deliver non-trivial qualities of service (the gain from combining resources must be significantly higher than just the sum of resources)
- Resource coordination through Virtual Organisations
- Standardization through e.g. Open Grid Forum (OGSA, OGSI and JSDL) and Web services
- Non-trivial qualities of service through blood, sweat and tears

Grid

- Grid layers

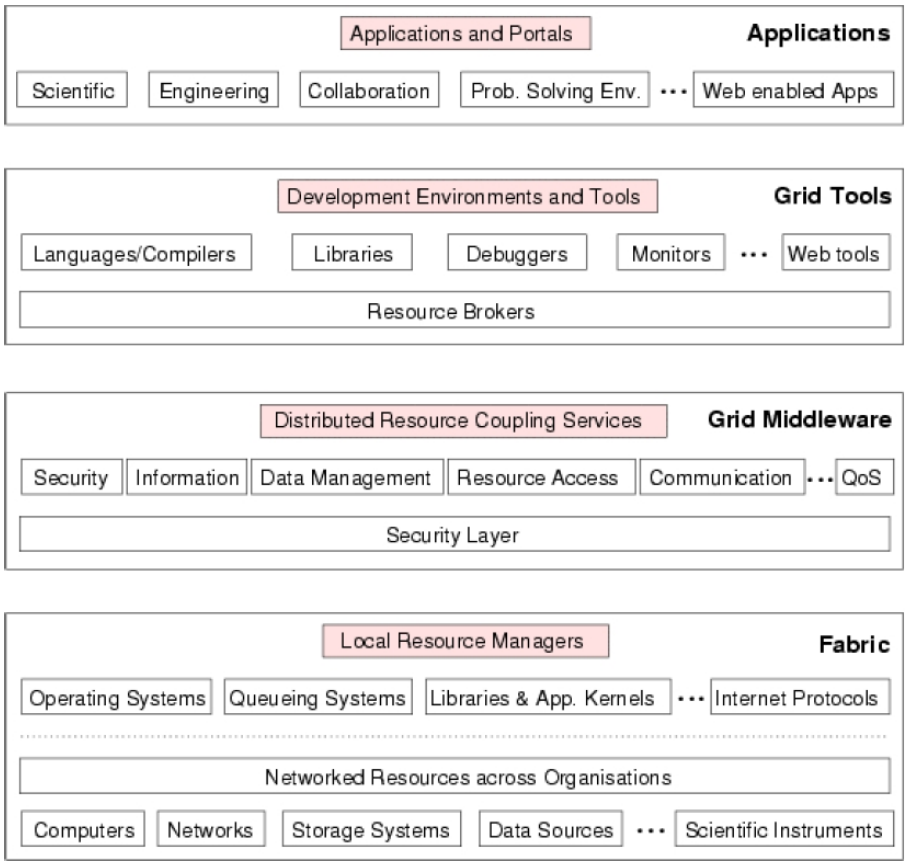

Grid

- Connecting fabrics

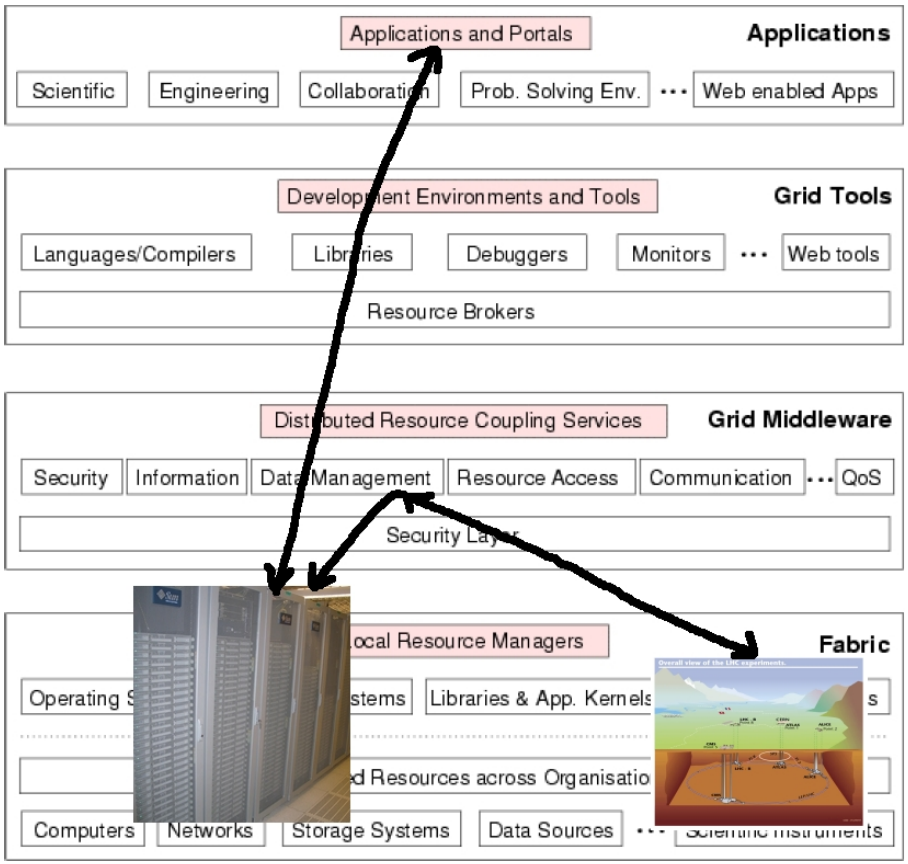

Grid Data Management

- Needed services:
  - Storage manager
  - Storage controllers
  - Storage elements
  - Indexing services
  - File transfer services
- Internal and external communications through web services

LHC Computing Grid

- Storage Resource Manager (storage manager)
- dCache, Castor or similar (storage controllers, storage elements)
- Globus Replica Location Service (indexing service)
- GridFTP2, gLite FTS (file transfer services)

Nordic tier 1

- Grid in the grid
- The 7 biggest Nordic compute centers, dTier-1s, form the NDGF Tier-1
- Resources are distributed
- Services are centralized
- Reduntant (all files replicated on more than one dTier-1)
- Load and hardware failiures distributed $\rightarrow$ high availability

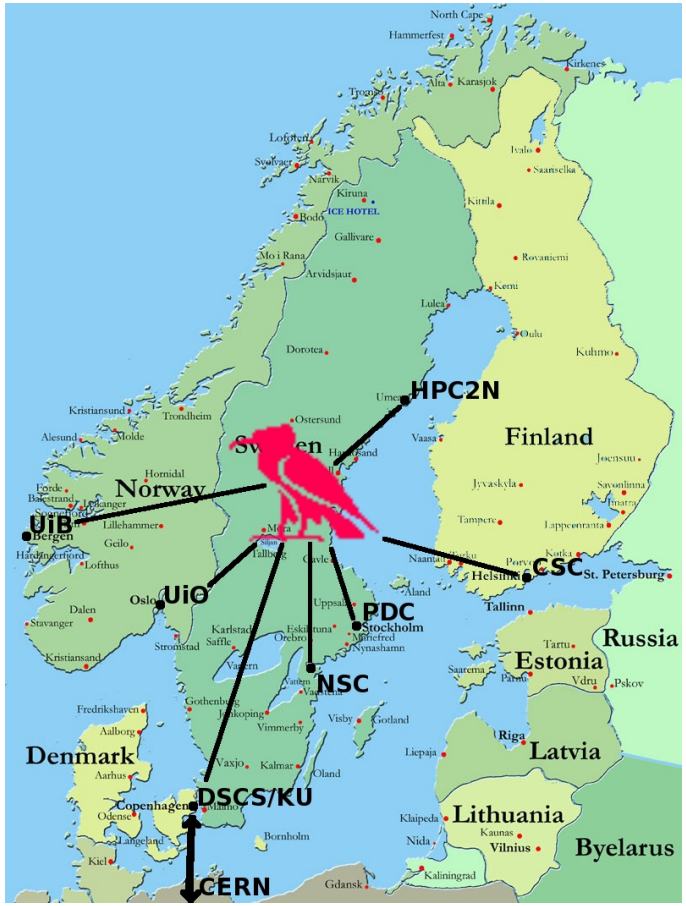

Nordic tier 1

- Uses dCache for controlling and storing data
- Uses SRM to manage data
- Register meta data to RLS
- Uses GridFTP2 for moving data internally

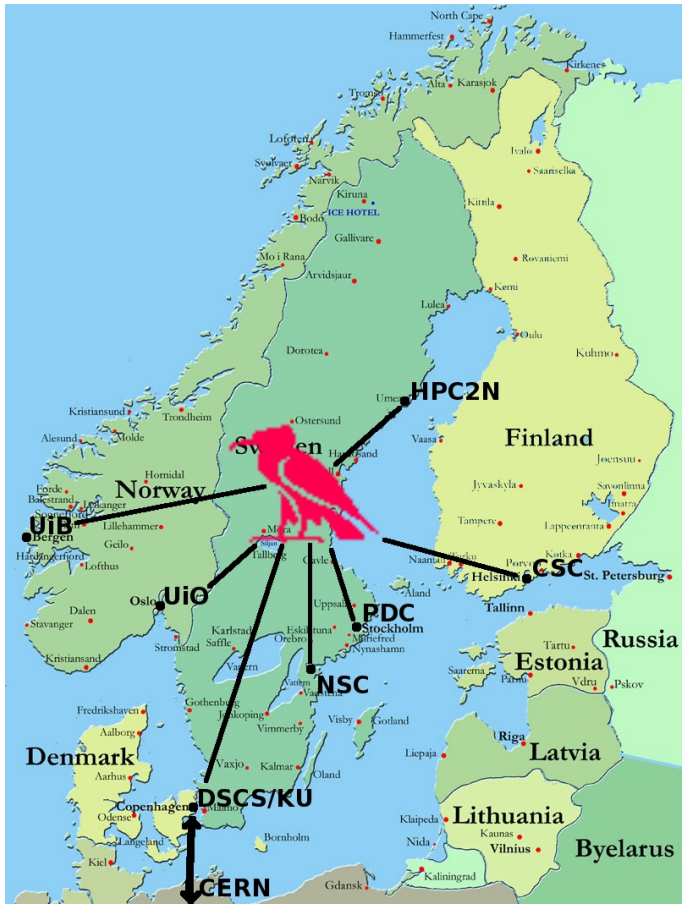

Nordic tier 1

- Connected with GEANT 10Gbit fiber to CERN
- Inter-connected with shared 10Gbit network from NORDUnet
- Dedicated 10Gbit LAN covering all dTier-1 centers from next year
- 200 MBytes/s sustained
- 100 TB stored, 184 TB capacity
- Will double both capacity and bandwidth within the end of the year

Conclusion

- There is a need for storing huge amounts of data
- Storing everything in one place is not feasible
- The Grid provides an infrastructure for handling these amounts of data
- The Grid is a nice way of connecting HPC resources
- The Grid must handle huge data sets properly to be useable
- Working solutions are emerging

Thank you!