



NORDUGRID

*Grid Solution for Wide Area
Computing and Data Handling*



End-to-end Grid System

Thomas Frågåt

On behalf of the EPF Oslo Grid Team





Department of Physics, University of Oslo

Spåtind, January 3rd, 2008

- Grid
 - What is it?
 - Why do we need it?
- End-to-end Grid system
 - Development, deployment, and support
 - Usage


- EPF Oslo Grid Team
 - David Cameron, Thomas Frågåt, Aleksandr Konstantinov, Jon K. Nilsen, Farid Ould-Saada, Katarina Pajchel, Weizhong Qiang, Alex Read, Bjørn H. Samset, Adrian Taga

Expected need for computing power at LHC



Storage:
Sampling rate of raw data 0.1 – 1 GB/s
Total amount of data 10-15 PetaBytes/year

Computing power:
Need for 200.000 of the fastest computers of today



Challenges

- When LHC starts running, a tremendous increase in computing power and storage is needed
- Impossible to build a supercomputer covering all the needs at CERN
 - Too expensive to both build, run, maintain, and support for one site to do
 - No known scalable cooling system available for such a computer
 - Not enough electric power in the Geneva area

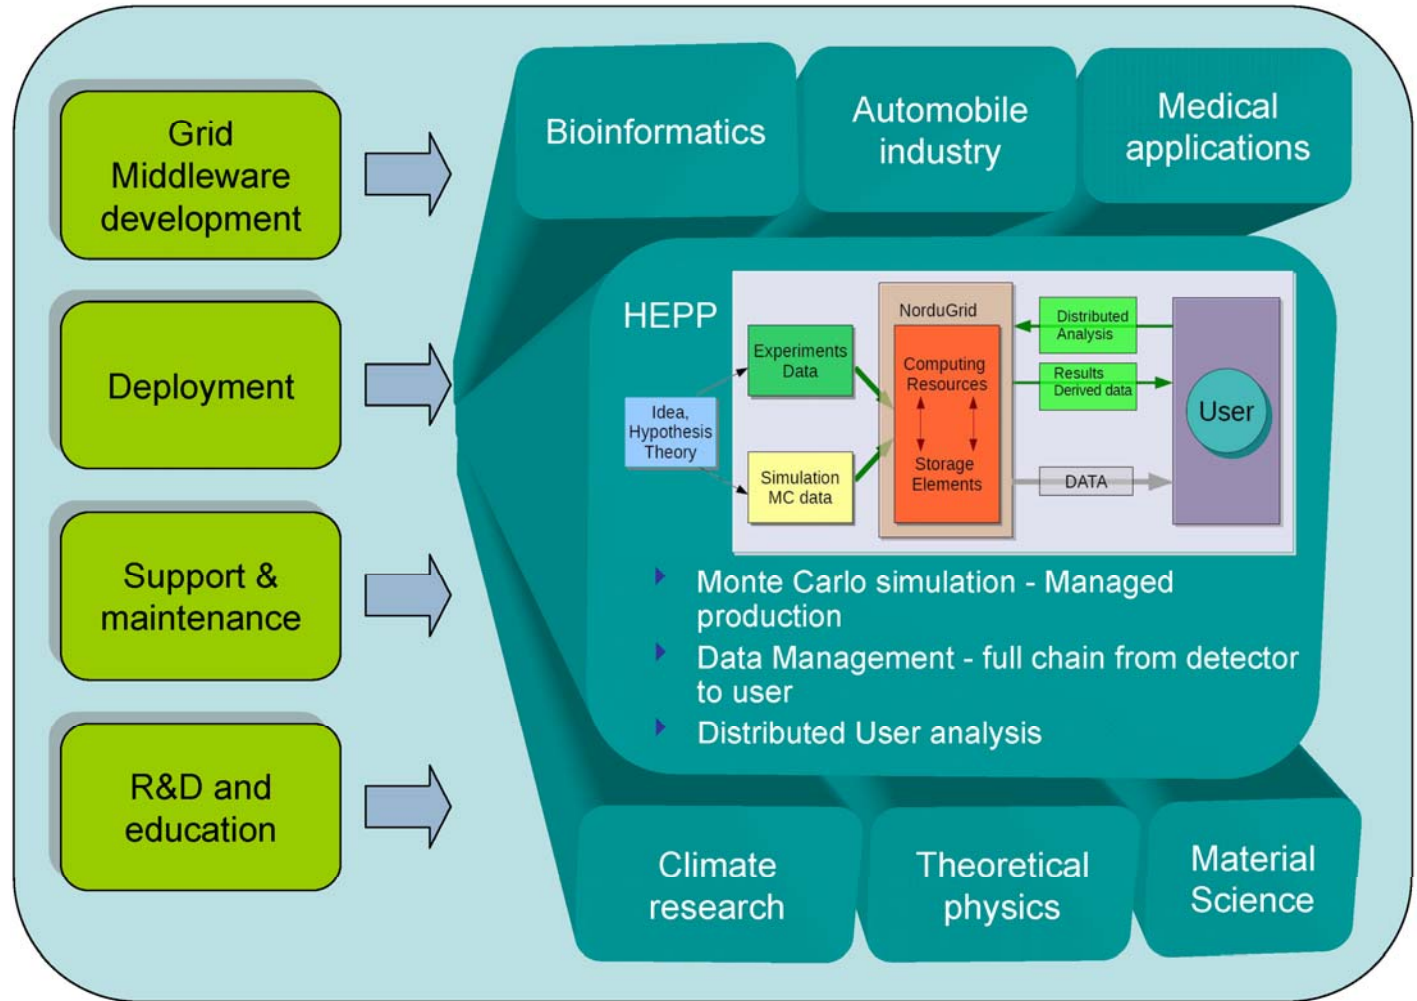
Solution: Grid

- The idea of connecting computers and clusters from all over the world into a Grid network
- A “glue”, *the middleware*, will connect the computers in a seamless manner
 - One interface to the resources: The user does not need to bother about what is happening behind the scenes
 - Creates a virtual supercomputer for the user
- Effective use of the distributed resources



A complete Grid environment for research

Synergy



Nordic Grid projects



- NorduGrid Collaboration (www.nordugrid.org)

- Develop, support, and coordinate inputs to the Advanced Resource Connector (ARC) middleware

- The KnowARC project (www.knowarc.eu)

- Create a novel, powerful Next Generation Grid middleware based on ARC and bridge the gaps between business and academia in Grid development



- NDGF (www.ndgf.org)

- NDGF is a production grid facility that leverages existing, national computational resources and grid infrastructures



- NGIn (www.nordugrid.org/ngin)

- Train new Grid experts, securing further technology development



- NGN (www.kbfi.ee/projects/NordicGrid/)

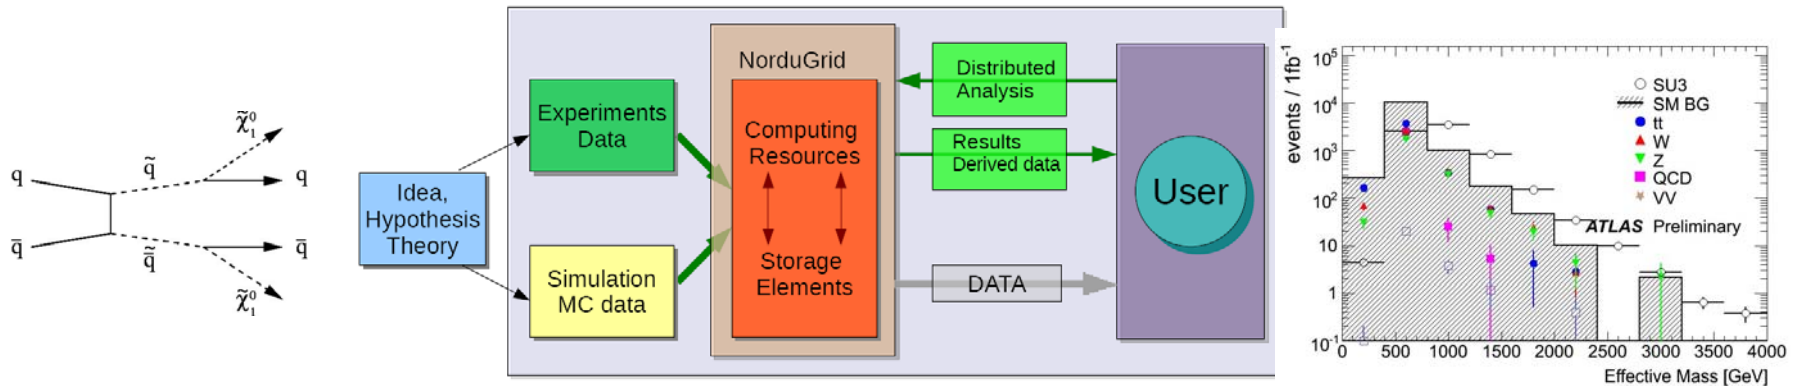
- R&D and education network including the Nordic and Baltic countries, and Russia

NorduGrid and the HEP community

- The Nordic HEP community was the initiator of the NorduGrid Collaboration
- The extreme HEP computing requirements were the main driving force of the development of the ARC middleware and Grid solutions
- The software, support, and infrastructure delivered by NorduGrid and related projects have become a necessity in the daily life of a particle physicist



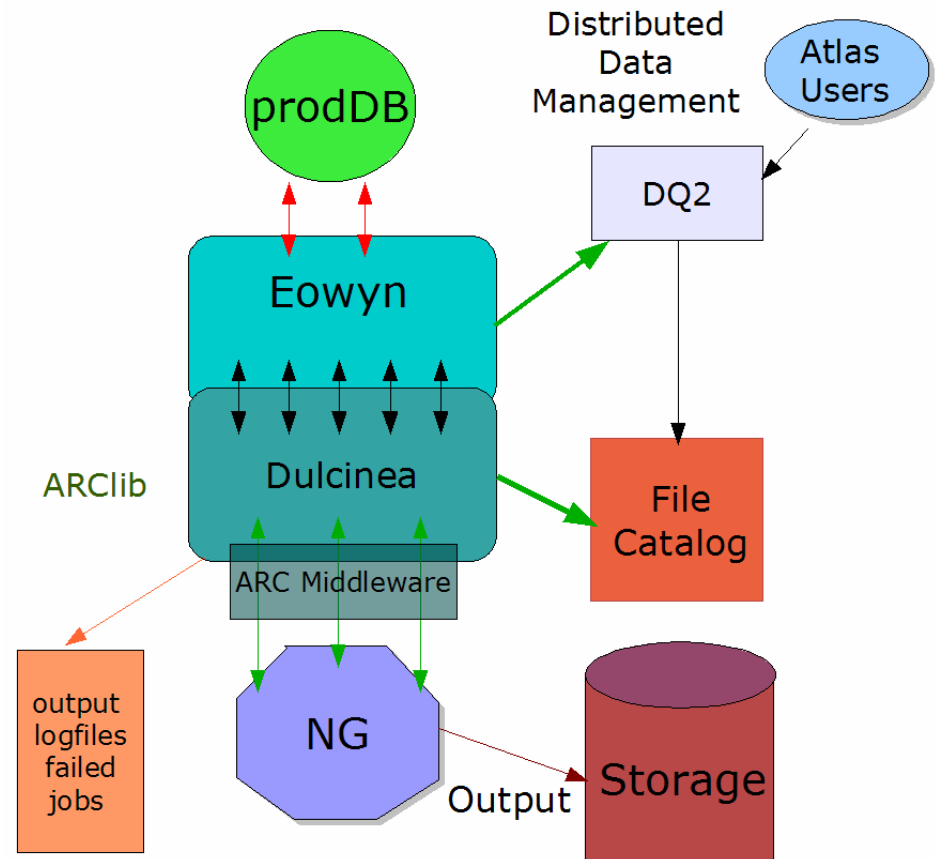
End-to-end Grid system for the Particle physicist



- Monte Carlo simulation and data processing
- Data management – full chain from detector to user
- Distributed user analysis

ATLAS production system

- Common *supervisor* (Eowyn) for all grids
 - Pulls jobs from the central production database (prodDB)
 - Updates job info in prodDB
 - Registers output files
- Dulcinea is the NorduGrid implementation of the *job executor* for required production routines
 - Prepares and submits jobs
 - Monitors the job status
 - Post-processes the job when finished, registers output, reports back to prodDB

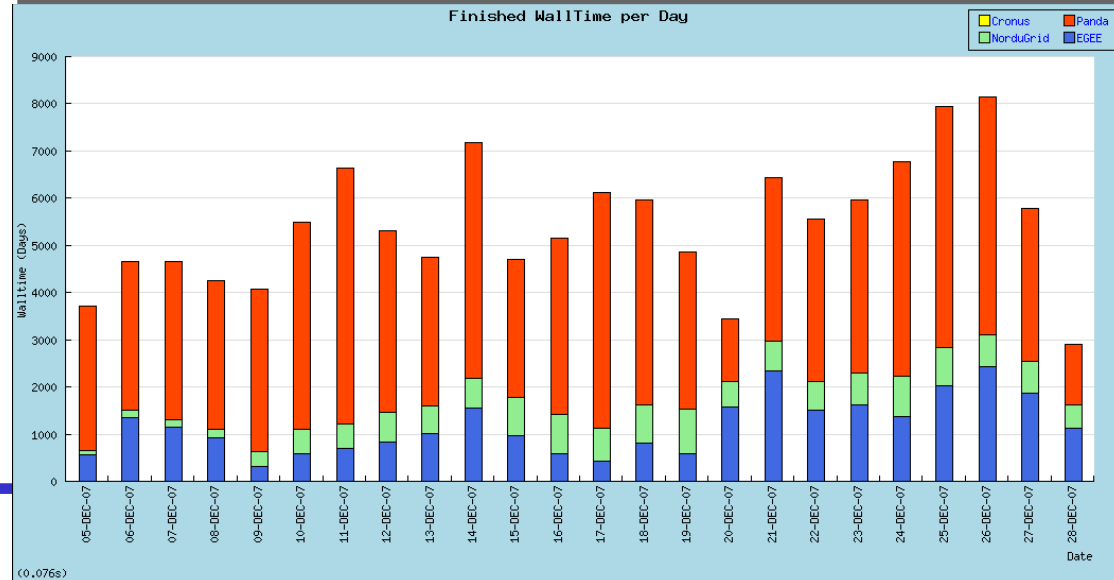



NorduGrid ATLAS production system: Achievements I

- Running since 2002
- The system is now robust and highly automated
 - Immediate submission of new tasks
 - Quick response and completion of high priority validation tasks
 - High rate of finished tasks
- Fully operational Tier1 with continuously increasing resources

Achievements II

- 2006:
 - 239 k jobs
 - average 150 CPU-days/day
- 2007 (status after 8 months):
 - 450 k jobs
 - Typical throughput > 600 CPU-days/day
 - 5-10% of Atlas
 - 300-500 GB a day on average during active periods of production

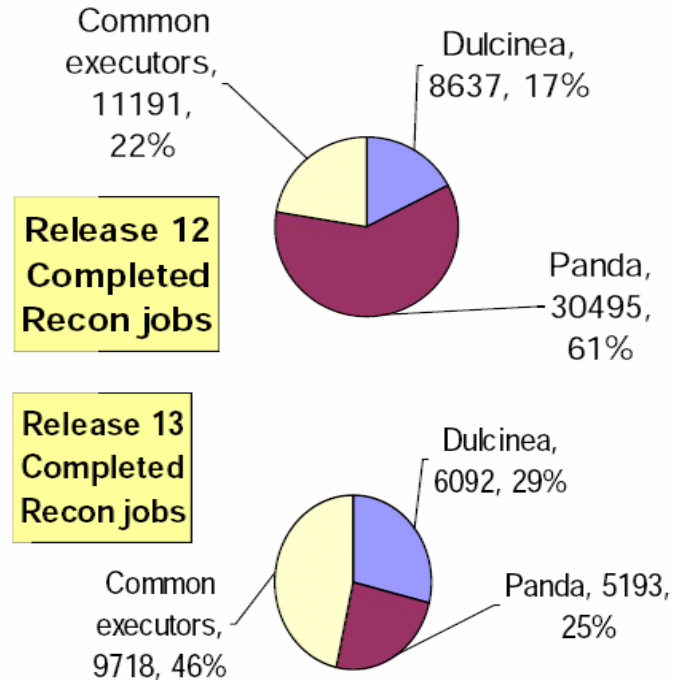


January 3rd, 2008

Achievements III


- The NorduGrid part of the ATLAS production system
 - Delivers 17–29% of the production full-chain (data useful for physics users)
 - Contributes with only 5–6% of the computing and storage resources

- In Atlas the NorduGrid resources are referred to as “one site”
 - proves that this ARC-based grid solution is a realization of the true Grid idea




Distributed Data Management (DDM)


- Operation and development in collaboration with NDGF
- Status: Fully operational distributed Tier1
- NDGF has become fully integrated in Atlas DDM
- MC data produced by NorduGrid are immediately available for the whole collaboration
- Currently 3 M Atlas files in the NorduGrid system
- Tier0 to Tier1 tests in July 07:
 - successful stable transfer at nominal rate 50 MB/s



- M4 cosmic ray data
 - The full Atlas data chain worked successfully
 - Nordic Tier1 participated and received all data
- The Nordic Tier1 participated also successfully in the M5 Cosmic run



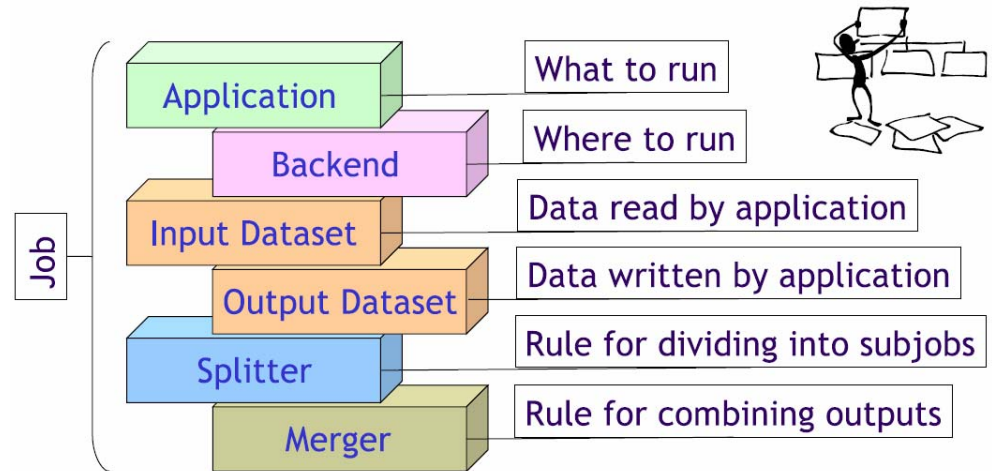
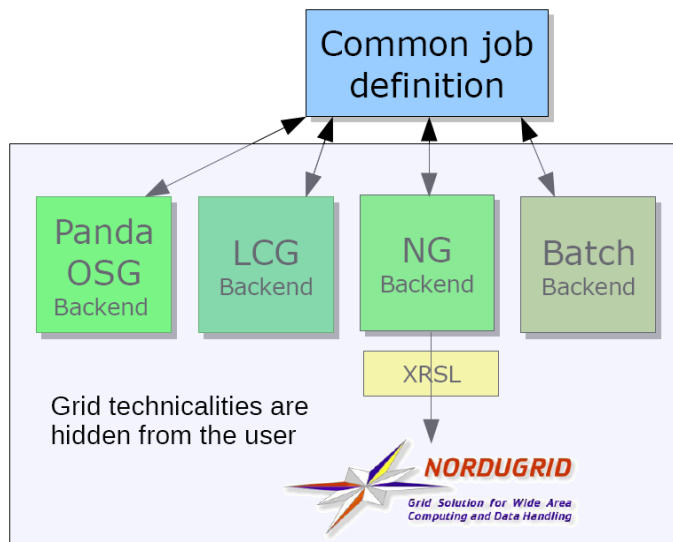
M4 Data transferred per day (GB) Aug. 23 - Sep. 8



M4 NDGF Throughput (MB/s) Aug. 28 -Sep. 4

Distributed user analysis

- Atlas Data Model: "Jobs go to data"
- Users require an uniform access method to all grid resources
- Ganga: user-friendly job definition and management tool
 - Allows an almost seamless switch between different computing resources
 - Built-in support for experiment-neutral applications
 - Built-in support for Atlas applications (based on Athena, DQ2)



Conclusion

- NorduGrid has developed a high quality production grid solution suitable for High Energy Physics
- It offers a fully operational system for
 - Monte Carlo and real data production
 - Distributed Data management
 - Distributed User Analysis
- The system performs very well
 - Rapid response to new tasks
 - High rate of finished jobs and tasks
 - Efficient utilization of the computing resources