NorduGrid

From World Wide Web to World Wide Grid - creating Nordic Testbed for Wide Area Computing and Data Handling

www.nordugrid.org

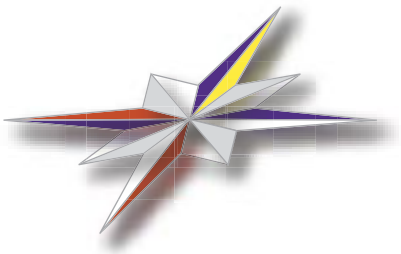

The aim of this brochure is to describe the NorduGrid research project and its achievements. NorduGrid research project is funded by NorduNet2 during years 2001-2002.

The project coordinator is prof. John Renner Hansen, NBI. The NorduGrid Steering Group consists of prof. P. Eerola (Lund), prof. T. Ekelöf (Uppsala), prof. J. Renner Hansen (NBI), docent S. Hellman (Stockholm), assoc. prof. F.Ould-Saada (Oslo), programme director J.Tuominiemi (HIP).

The NorduGrid Technical Working Group consists of Dr. M. Ellert (Uppsala), Dr. A. Konstantinov (Oslo), Dr. B. Kónya (Lund), Dr. O. Smirnova (Lund), Dr. A. Wäänänen (NBI).

For more information, contact us at contact@nordugrid.org or visit our homepage http://www.nordugrid.org

Copenhagen, October 2002.

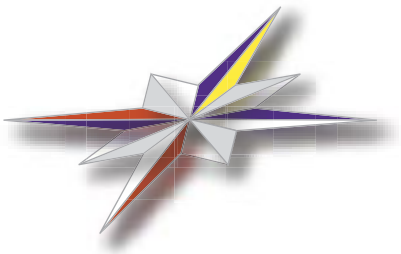

NorduGrid objectives and participants

The Nordic Testbed for Wide Area Computing And Data Handling project (NorduGrid) is a part of the Nordunet2 programme, aimed to develop networked applications with extensive usage of modern utilities and tools in the Nordic countries.

Objective

The aim is to establish an inter-Nordic test bed facility for implementation of wide area computing and data handling. The facility will provide the infrastructure for interdisciplinary feasibility studies of GRID-like computer structures. The project shall collect and document experience, as an input to the decision process on the future computer infrastructure strategy for sciences with distributed PByte storage requirements and processing power in the order of multi-teraflop.

What is Grid?

Grid is a technology to share and access seamlessly computing resources that are not subject to a centralized control. Grid will be the future infrastructure of computing and data management. The computing resources are connected together through a layer of software called the middleware, which uses standard, open, general purpose protocols and interfaces. This middleware forms the glue binding the resources into a virtual system.

In the same way as the World Wide Web gives us access to information, the World Wide Grid will give us access to computing capacity and data storage in the future.

Participants

DK Research Center COM
DK DIKU
DK Niels Bohr Institute

SE Lund University
SE Uppsala University
SE Stockholm University
SE Royal Institute of Technology

NO Oslo University
NO Bergen University

FI Helsinki Institute of Physics

Why do we need Grid in the Nordic countries?

Ever increasing amounts of data: an example

The European High Energy Physics community is in the process of construction and deployment of the Large Hadron Collider (LHC) - the world biggest accelerator, being built at the European Particle Physics Laboratory (CERN) in Geneva.

Challenges to be faced by physicists are unprecedented. Four experiments will be constructed at the LHC to observe events produced in proton-proton and heavy ion collisions. Data collected by these experiments will allow for exploration of new frontiers of the fundamental laws of nature, such as the Higgs mechanism with possible discovery of the Higgs boson, CP-violation in B-meson decays, supersymmetry, large extra dimensions, mini black-holes, and so on.

One of the greatest challenges of the LHC project will be the acquisition and analysis of the data. When, after a few years of operation, the accelerator will run at its design luminosity, each detector will observe bunch collisions at a rate of 4·107 per second. A set of filter algorithms, implemented in hardware and on state-of-art programmable processors, aims to reduce the event rate to less than 1000 events per second for final storage and analysis. The equivalent data volume is between 100 MByte/sec and 1 GByte/sec. Each experiment is expected to collect 1 PByte of raw data per year. The two LHC general purpose experiments, ATLAS and CMS, have each more than 150 participating institutes distributed all over the world. 2000 physicists per experiment contribute to the development of hardware and software and they expect to have almost instantaneous access to the data and to a set of up-to-date analysis tools.

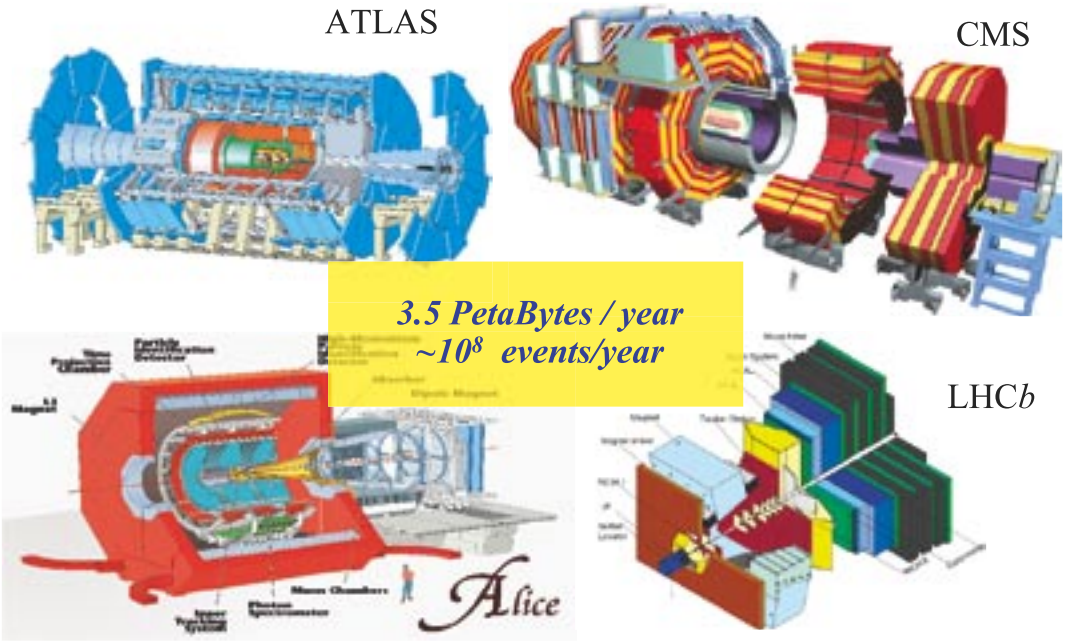

NorduGrid achievements

Project description

NorduGrid is a Grid research project which:
-develops openly available middleware (NorduGrid Toolkit),
-operates a production quality Grid Testbed,
-pursues basic research on Grid computing and surveys current Grid technologies,
-exposes the infrastructure to end-users in different scientific communities.

The NorduGrid Testbed consists of 10 sites in the Nordic countries, with about 200 CPUs and 4 TB disk storage capacity in total.

The architecture

The NorduGrid architecture consists of the following elements:
-Computing Element (clusters)
-Storage Element
-Replica Catalog
-Information System, including NorduGrid Information Model
-Grid Manager
-User Interface and Broker

The NorduGrid architecture is described in more detail in the documents available from the project's homepage at http://www.nordugrid.org/documents

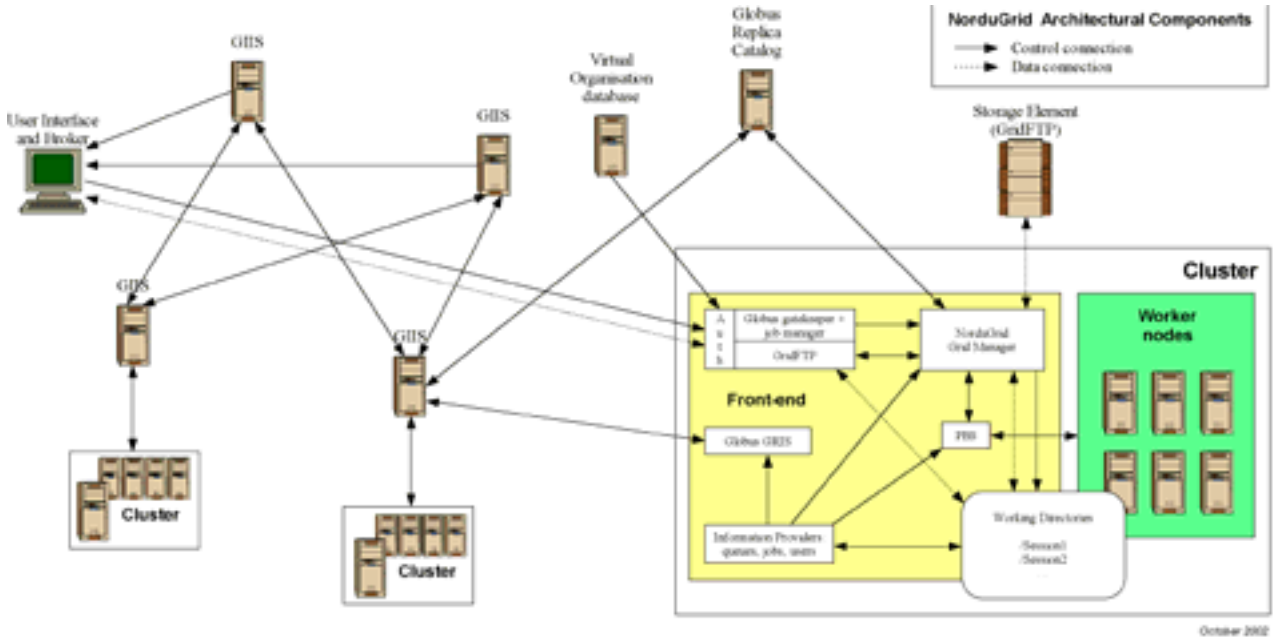

NorduGrid Toolkit

The NorduGrid toolkit is a Globus-based Grid middleware, developed to meet the basic requirements of distributed computing and data handling. The toolkit supports a broad variety of Linux platforms (Redhat, Mandrake, Slackware, Debian) and it is available freely from the NorduGrid site at www.nordugrid.org.

Its most important components are:
-the MDS-based Information System, providing reliable dynamic information about the status of the services, clusters, jobs and users,
-the Grid Manager, residing at a gatekeeper and performing job pre- and post-processing, as well as job control and submission to the local resources,
-the User Interface, responsible for the user's interaction with the Grid. The User Interface contains a built-in resource broker,
-extended resource specification language (XRSL) for formulating job requests,
-Load Monitor web interface to the NorduGrid Testbed.

The software package is easily installable on top of an existing Globus installation and comes with detailed documentation.

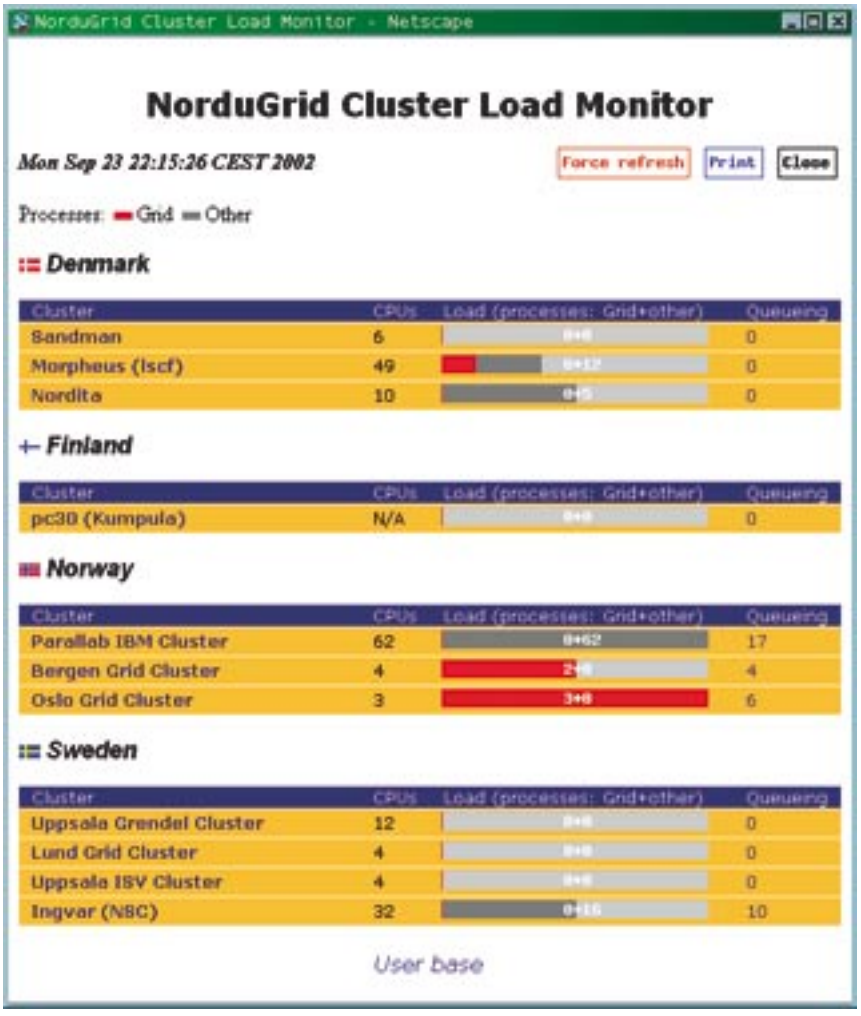

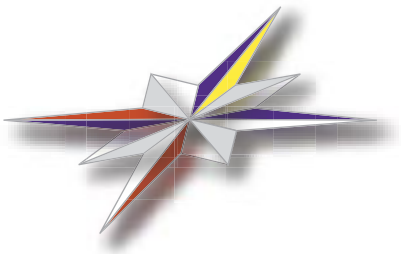

NorduGrid in production...

During summer 2002, the NorduGrid testbed and middleware were successfully deployed in a production environment. Nordic physicists participating in the ATLAS experiment at LHC used, for the first time, the NorduGrid toolkit to simulate a large amount of data with several computer clusters located in Denmark, Norway and Sweden.

The need for Simulation

In order to design and optimize the ATLAS detector, and prepare for the further data analysis, it is mandatory to simulate LHC events now. The procedure follows three main steps. Particles emerging from collisions are first generated using computer programs based on physics theories. A detailed simulation of the passage of the particles through the virtual detector is then performed. The resulting interactions are then converted into electronic signals similar to those produced by the real detector. In a final step, particle trajectories and energies are reconstructed before physics analysis takes over.

The ATLAS Data Challenges

How large data samples are needed to get ready for the first real LHC collisions? ATLAS opted for a series of Data Challenges of increasing size and complexity to prepare and test the software in stages. At the begining, the necessary software to simulate and analyse a small sample of events was collected and used at CERN. The goals of the next steps were to use better software and to collect a larger sample. For example, the simulation part of first such step corresponded to 71.000 CPU-days. Due to this huge amount of computing power, it was essential to distribute the task between 37 institutes from 18 countries participating in ATLAS.

Gridifying

It is natural that the Nordic ATLAS physicists would contribute to Data Challenges. What is unique is that they were the first ones who could entirely use Grid tools to complete their share of the task. The testbed included 8 Linux clusters across Denmark (Copenhagen), Norway (Bergen, Oslo) and Sweden (Lund, Uppsala, Linköping). Despite having different operating systems and hardware characteristics, the clusters performed as a single virtual computing center, having jobs distributed in an optimal way, and writing the output onto a dedicated storage area in Oslo. In total, 1300 jobs were successfully executed, distributed across the Nordic countries. The processing of 200 GB of input data produced 765 GB of output data, in 530 CPU-days. The average time to process a single event was 150 seconds. The ATLAS challenge was successfully met, with the failure rate being negligibly small, with all failures not attributed to the NorduGrid middleware. In the future, it is foreseen to continue running ATLAS Data Challenges on the NorduGrid, using upgraded middleware and increased computing resources. The NorduGrid tools and ATLAS software have now been installed at some Nordic computing centers (Linköping, Bergen), such that larger clusters are available for the next phases of the ATLAS Data Challenges. The second phase was designed to start in November 2002, with special emphasis on the use of Grid tools. NorduGrid will be present and the next challenge will be a distributed physics analysis across the Nordic countries, and hopefully worldwide.