M-grid system architecture, collaboration and grid use

Arto Teräs <arto.teras@csc.fi>

Sec-tv meeting

CSC, Espoo Dec 14, 2005

Contents

- **Introduction to the Finnish Material Sciences Grid (M-grid)**
- **System architecture**
- **Shared system administration**
- **Grid use and resource sharing**

The Material Sciences Grid (M-grid)

- **Goal: Throughput computing capacity mainly for the needs of physics and chemistry researchers**
- **Joint project between seven Finnish universities, Helsinki Institute of Physics and CSC**
  - Partners mainly laboratories and departments, not university IT centers
- **Jointly funded by the Academy of Finland and the participating universities**
  - Funding application Nov 2003, deployment Oct 2004
- **First large initiative to put Grid middleware into production use in Finland**
- **Platform: Linux based PC clusters**

CSC

Grid environment

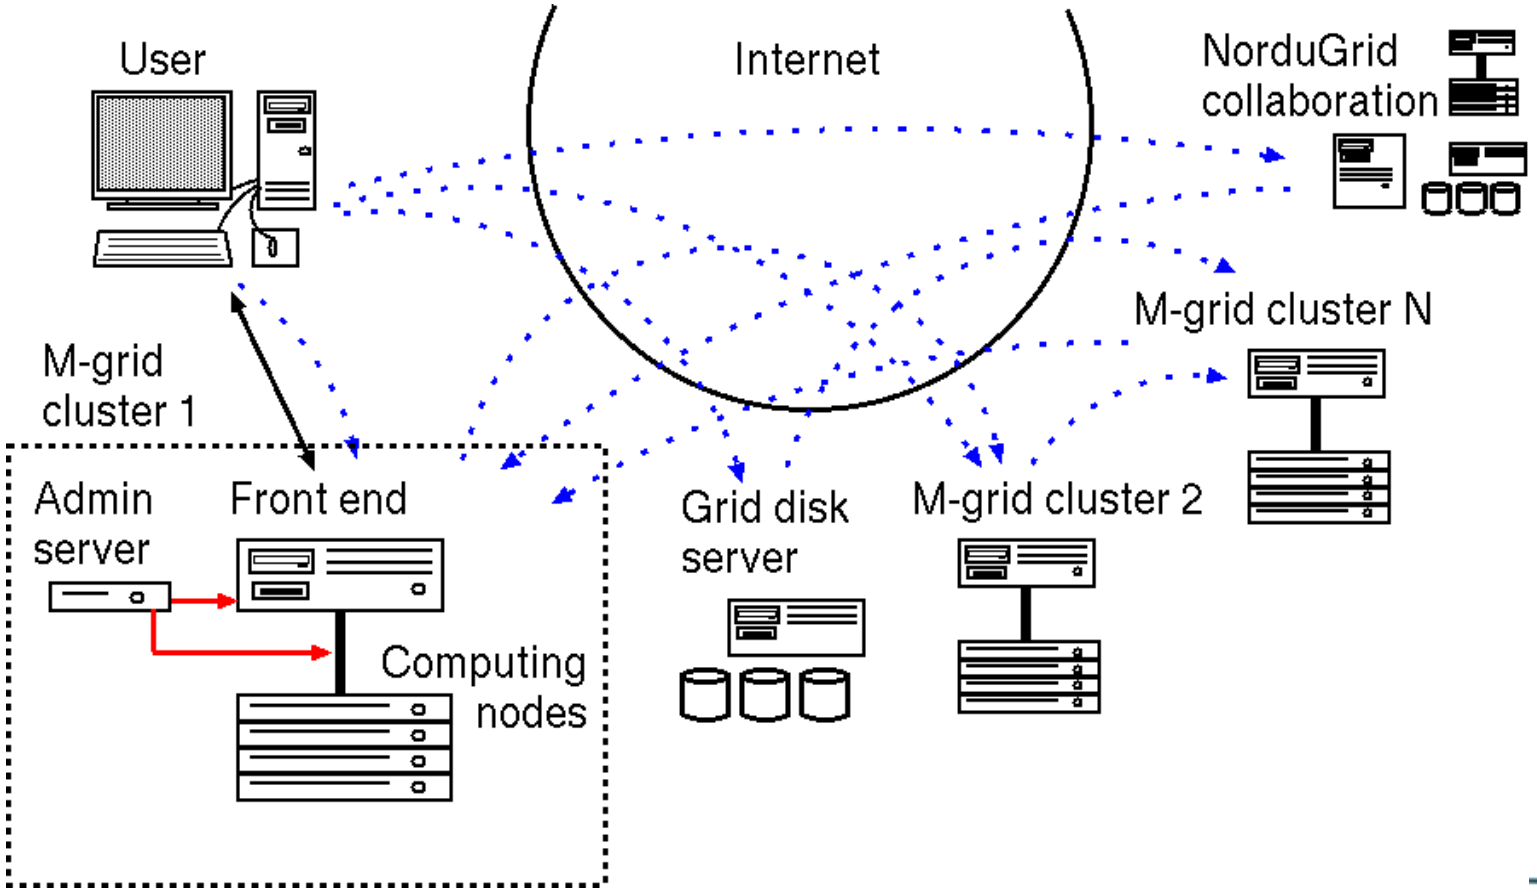

Hardware and CPU distribution

- **Ten clusters of varying size**
  - Dual AMD Opteron computing nodes (HP DL145): 1.8-2.2 GHz, 2-8 GB RAM, 80-320 GB local disk
  - Front end (HP DL585): 1-2 TB shared disk
  - Network 2 x Gbit Ethernet + remote administration network
- **Total 778 (CSC) + 434 (universities) CPUs in the computing nodes, theoretical total computing power 5 TFlop/s.**

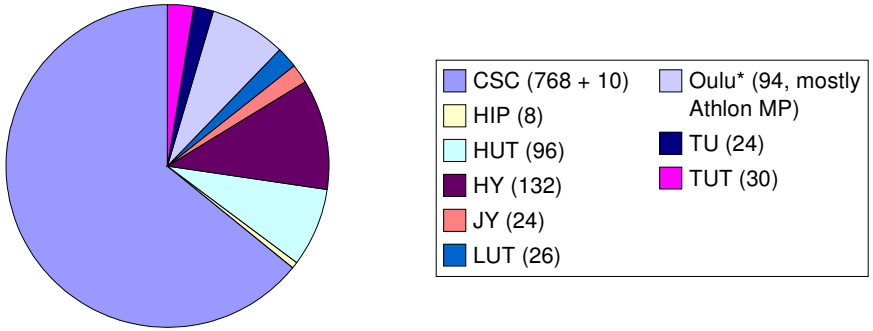

One M-grid cluster

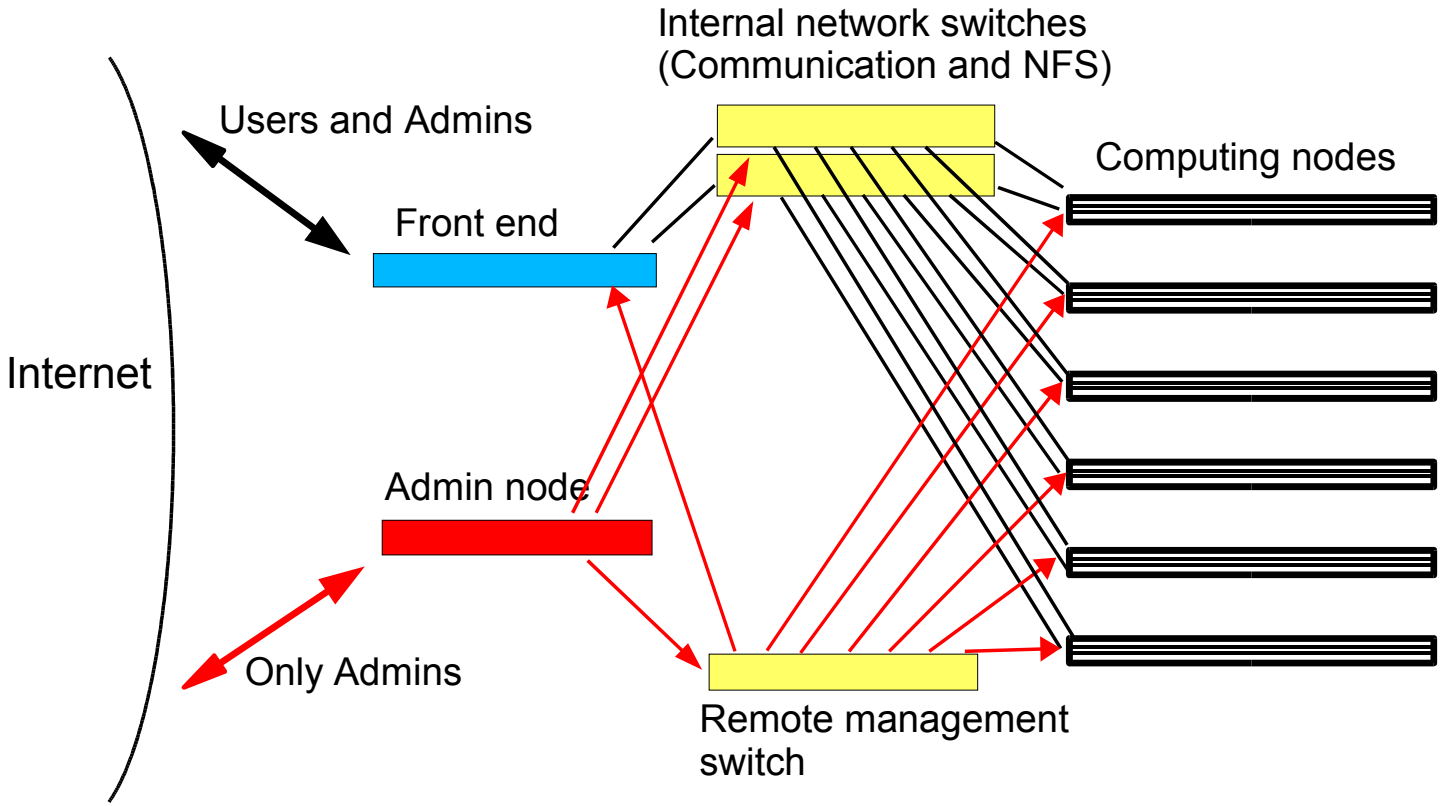

Operating system and Grid middleware

- **NPACI Rocks Cluster Distribution**
  - Cluster oriented Linux distribution, main developer San Diego Supercomputing Center, U.S.A.
  - Based on Red Hat Enterprise Linux, but not a Red Hat product
  - http://www.rocksclusters.org

- **N1 Grid Engine batch queue system**
  - Local resource management in each cluster

- **NorduGrid ARC Grid middleware**
  - Enables shared use of the systems, the middleware selects a free resource automatically
  - http://www.nordugrid.org

System administration in M-grid

- **Tasks divided between CSC and site administrators**
- **CSC administrators:**
  - Maintain (remotely) the operating system, batch queue system, Grid middleware and certain libraries for all sites except Oulu
  - Separate small test cluster for testing new software releases
- **Site administrators**
  - Local applications and libraries, system monitoring, user support
- **Regular meetings of administrators every two months, common mailing list**

Installing updates

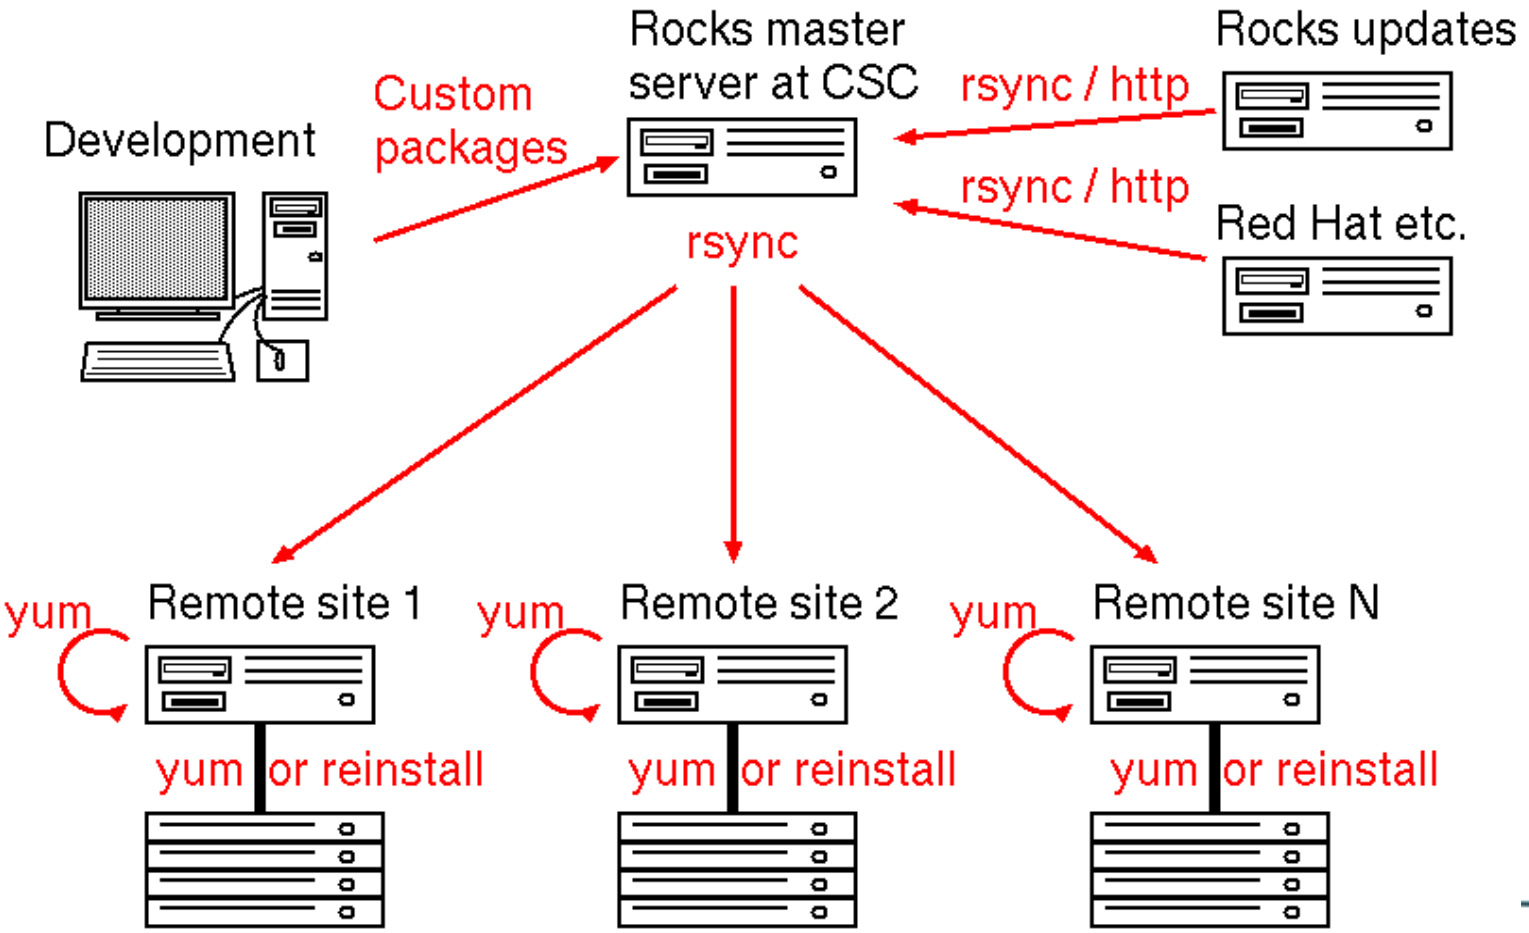

Goals of Shared System Administration

- **Centrally administered foundation while maintaining local control**
  - A new paradigm -- traditionally in Finland academic HPC resources have been centralized at CSC
- **Easier for universities than setting up their own cluster from scratch**
  - However, needs a significant amount of work both from CSC and the local sysadmins
- **Take advantage of the local sysadmin expertise**
  - Site administrators know the software of their own group best => faster and better user support

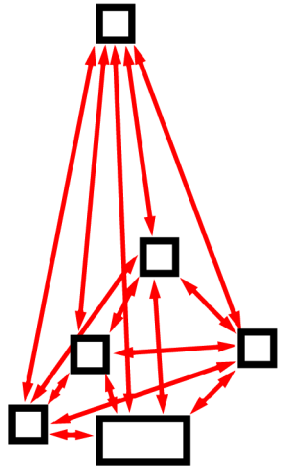

36 pairs for collaboration!

Positive experiences

- **Site administrators have found CSC support valuable**
  - On the other hand local control (root access) enables quick fixes and is important psychologically
- **Site administrators have picked up tasks which benefit everyone — CSC has not done everything**
- **Collaboration has strengthened relationships between groups also in their research**
- **Systems are close to the user**
  - Easier to talk to the own group sysadmin, less support requests to CSC
- **Most site administrators are also users => direct usability feedback to CSC**

Negative experiences

- **Wiki based FAQ hasn't become popular, questions and answers are buried on the mailing list**
  - The Wiki model can also be a success: e.g. Wikipedia
- **Some users found support poor**
  - Varying experiences: on some sites users are very happy
  - There were gaps in the user documentation mainly due to lack of human resources — documentation can be written in a distributed group but compiling it needs central coordination
- **Communication with university IT centers not always optimal**
  - IT centers not informed well enough — CSC assumed too much that the local groups will discuss all necessary topics with them

Grid use and resource sharing

- **Policy: Jobs can be submitted both to the local queue and through the grid interface**
  - Queue priority: local jobs 80%, grid jobs 20%
- **Goal is to minimize waste of resources: empty nodes are always available for use (dynamical sharing)**

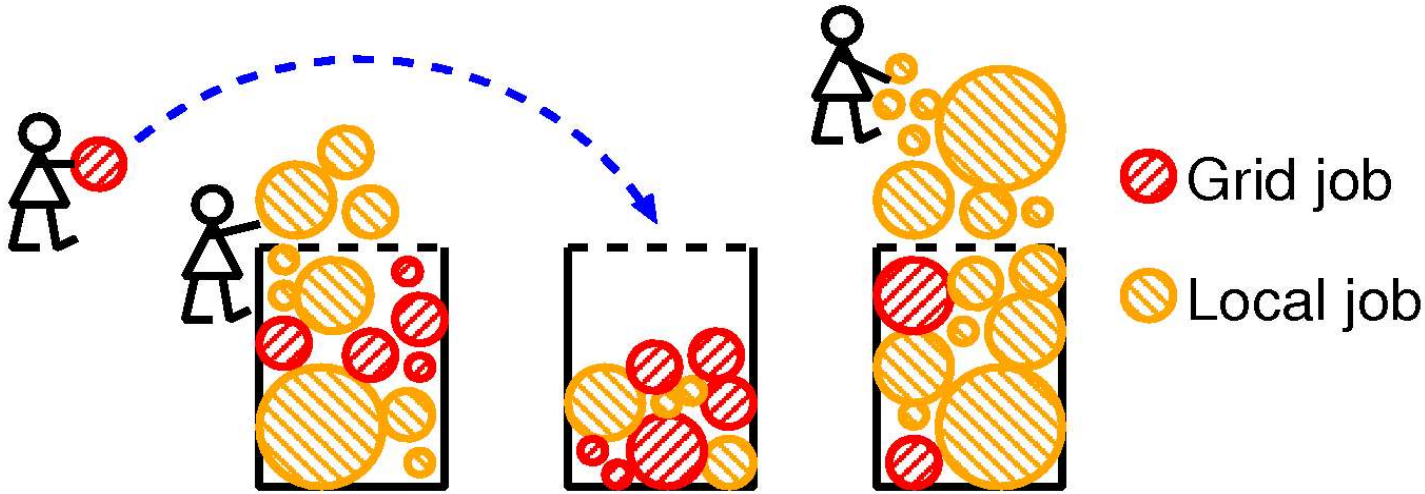

Role of grid middleware

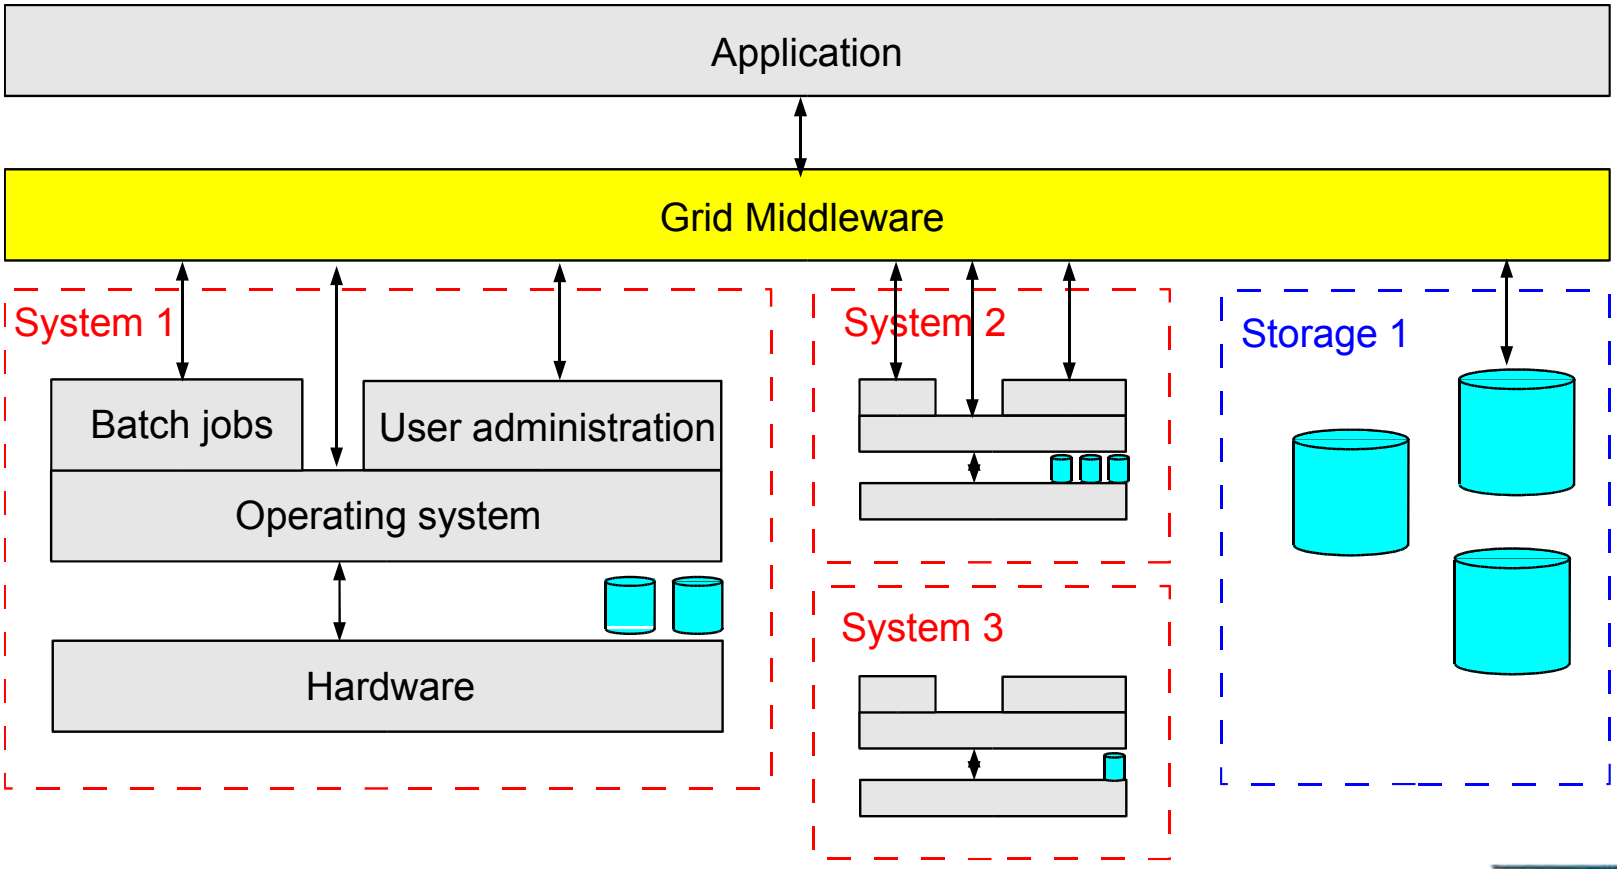

Authentication and authorization

- **Standard username/password (or ssh key) for local use**
- **Grid authentication based on personal X.509 certificates**
  - Using NorduGrid Certificate Authority, CSC acting as a RA
  - Certificates stored as files in the home directories of the users
- **Currently all M-grid users who request a certificate are authorized in all M-grid clusters**
  - No direct logins via ssh, but ability to run arbitrary binaries in the grid jobs
  - Each group decides their policy of giving user accounts in their cluster
- **Other user groups from NorduGrid are not currently authorized in any of the clusters, but may be in the future**

Grid experiences

- **Grid use started August 2005**
  - Installation was delayed due to other tasks and a few technical problems
  - Environment still in development
- **Grid environment must be better than the existing one, otherwise nobody will use it!**
  - Long queue in the local cluster and empty resources on the Grid may be a good enough incentive
- **Currently only a few Grid users, time will show how well the Grid environment will be adopted**
- **Collaboration model has been successful: Grid projects always have other aspects than just the technology**

Grid collaboration and security

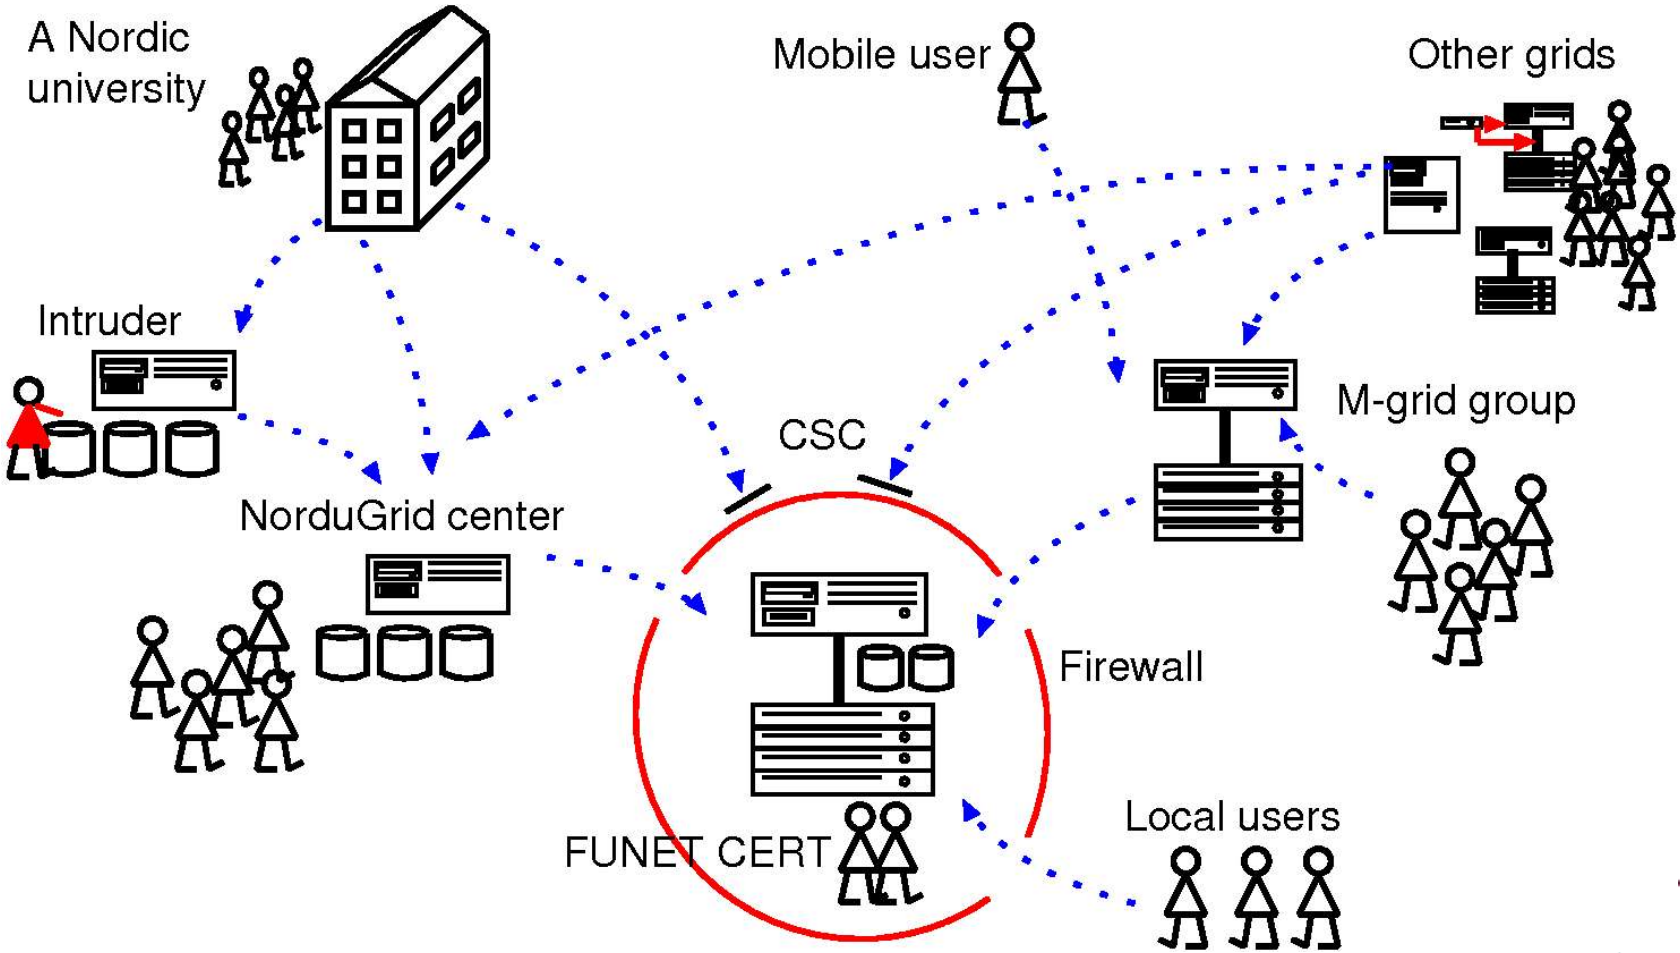

More information

- **M-gridin homepage: http://www.csc.fi/proj/mgrid/**
- **Rocks homepage: http://www.rocksclusters.org**
- **NorduGrid homepage: http://www.nordugrid.org**
- **Contact people:**
  - Arto Teräs <arto.teras@csc.fi>
  - Kai Nordlund <kai.nordlund@helsinki.fi>
  - Olli-Pekka Lehto <oplehto@csc.fi> (Rocks)
  - Urpo Kaila <urpo.kaila@csc.fi> (security)
- **Thank you! Questions?**