Grids — Computing and Collaboration

Arto Teräs <ajt@iki.fi>

CSC, the Finnish IT center for science

University of Pune, India, March 12th 2007

Think Without Limits

- **"We believe computer science is poised to become as fundamental to biology as mathematics has become to physics."**
- **"Big challenges for future computing systems have elegant analogies and solutions in biology, such as the development and evolution of complex systems, resilience and fault tolerance, and adaptation and learning."**

— Towards 2020 Science, Microsoft Research, 2006
http://research.microsoft.com/towards2020science/

CSC

Contents

- **Introduction to grid computing**
- **The Finnish Material Sciences Grid (M-grid)**
- **Grid use and experiences in M-grid**
- **Data centric grids**
- **Identity federation**
- **Summary**

What is a Grid?

- **A grid is the combination of networked resources and the corresponding grid middleware, which provides grid services for the user.**
  (definition from Dieter Kranzlmüller, deputy director, EGEE project)
- **Uniform and secure access to a geographically distributed heterogeneous system**
- **Both the set of users and connected resources vary dynamically**
- **Grids go across multiple organizations!**

A Computing Grid

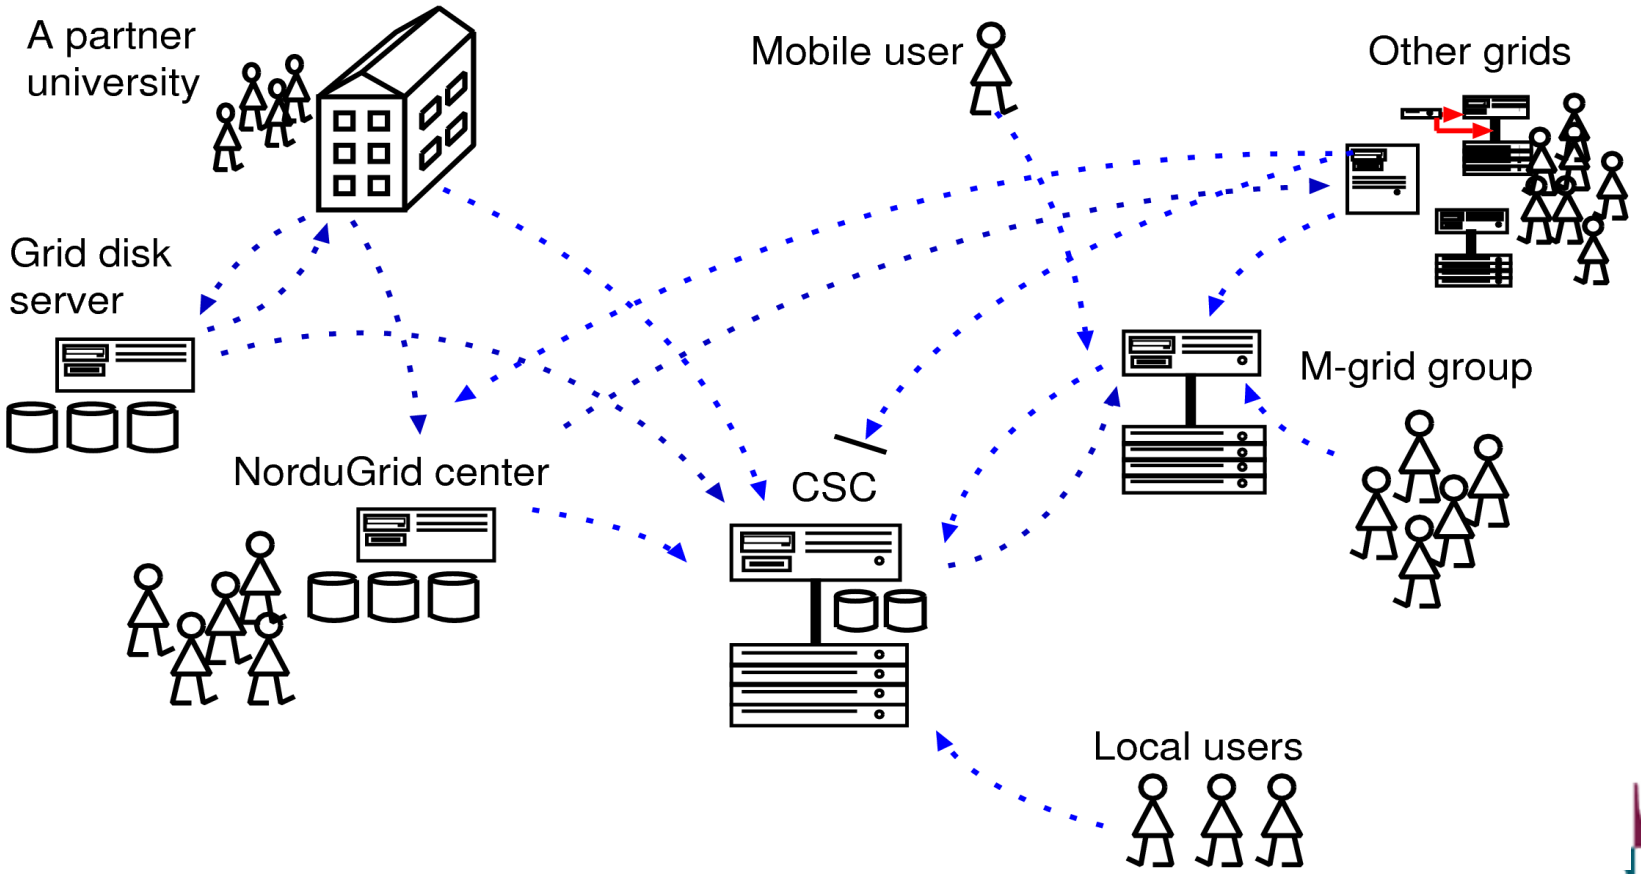

Common Misconceptions

- **Grid increases resources**
  - Popular comparison with the World Wide Web is misleading:
    - One web server may serve a thousand users, but one grid user wants to use a thousand servers...
  - Effective use of resources can bring some savings, but new services and easy access much more important
- **Grid magically binds software together**
  - Vision: Computing power as electricity from the plug
    Reality: still quite far from it
  - If data formats or APIs are incompatible grid doesn't help
  - Possibility to monitor job execution is important — trying to make a black box easily results in a black hole

A Few Major Grid Projects

- **EGEE (http://www.eu-egee.org)**
  - The largest in the world, started in Europe but now expanding worldwide
- **DEISA (http://www.deisa.org)**
  - Main focus on connecting large scale supercomputers in Europe, collaboration with centers outside Europe
- **TeraGrid (http://www.teragrid.org)**
  - U.S. national grid, large sites and very high speed networks between them
- **NAREGI (http://www.naregi.org)**
  - Japanese grid, started later but making a large investment in software development

CSC

Grid Software (Middleware)

- **There are many different software packages, called grid middleware, which can be used to build a grid**
  - The Globus Toolkit
  - gLite
  - NorduGrid ARC
  - Unicore
  - Shibboleth
  - ...

- **In this talk I try to present generic concepts instead of comparing different middleware**

Role of Grid Middleware

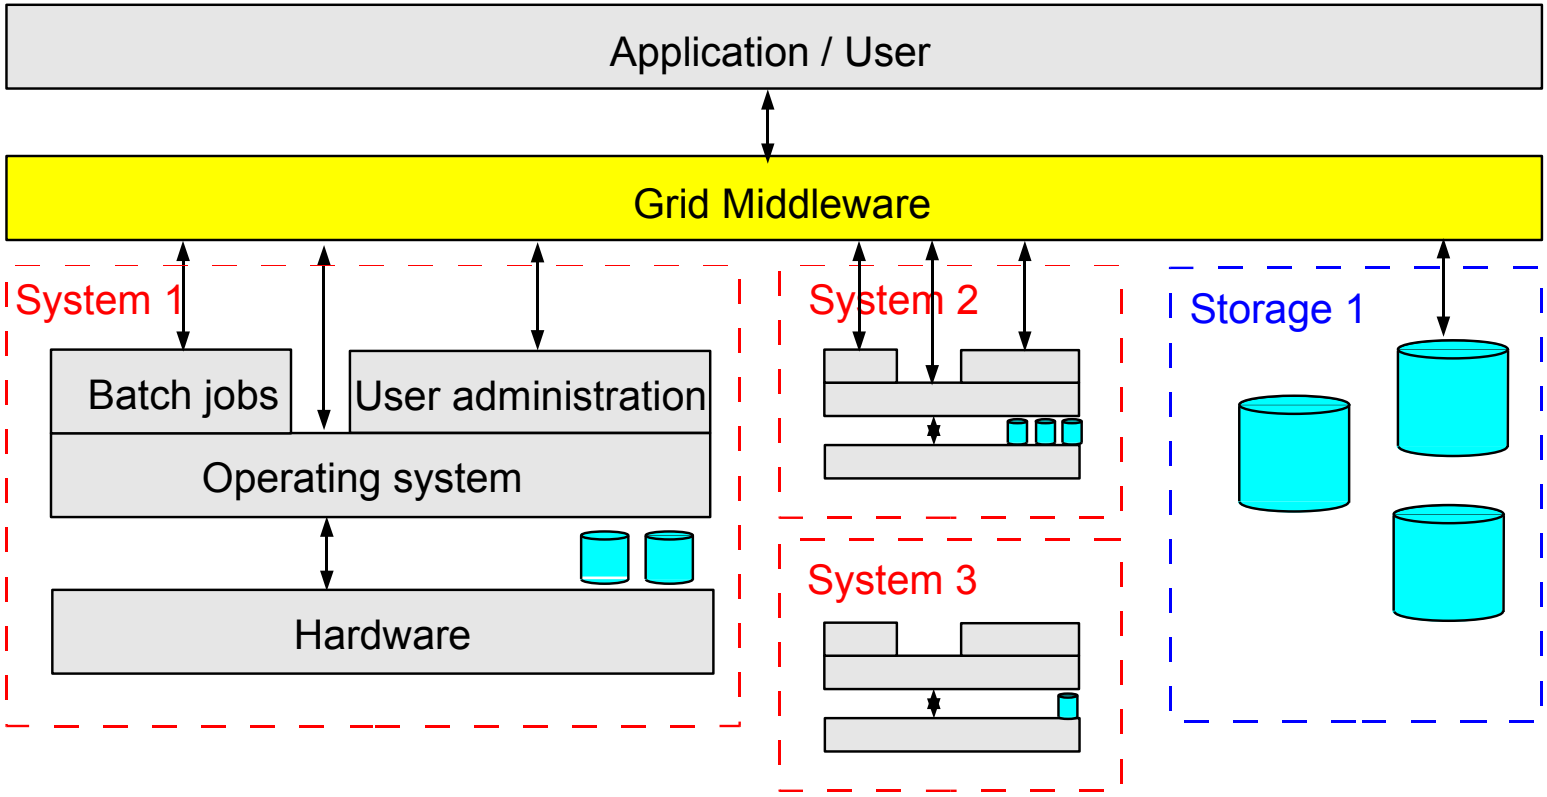

Steps to Start Using a Grid

1) **Parallelize your problem**
   - Make the subtasks as independent as possible
2) **Apply for access to a grid environment**
   - Scientists can typically get access to national grids
   - International grid projects are typically either focused on a certain scientific field, consist of a number of collaborating institutes (need to work in one of them to get access) or try to find scientifically important applications (anyone can apply)
3) **Download client software and follow instructions**
   - Usually specific to each grid

Does One Need to Change Existing Applications?

- **Three different approaches:**

1) Using the application as is: grid middleware will move the executable and the data to the target system
   - Library dependencies often need to be resolved by linking statically or packing them to go with the application
2) Installing the application on the target system and using it via the grid interface
   - Batch processing type applications normally work without changes, interactive applications may be more difficult
3) Modifying the application to fully exploit a distributed environment
   - Initially more work but easiest for users!

Authentication: Local vs Grid

- **Local resources**
  - User name and password
  - “Login” authenticates and usually also authorises to use local resources
- **Grid environments**
  - Authentication usually based on X.509 certificates granted by a third trusted party, Certificate Authority (CA)
  - Each user has his/her own personal certificate
  - Authentication is separate from authorization => having a valid certificate does not automatically give access to resources

Certificate Based Authentication

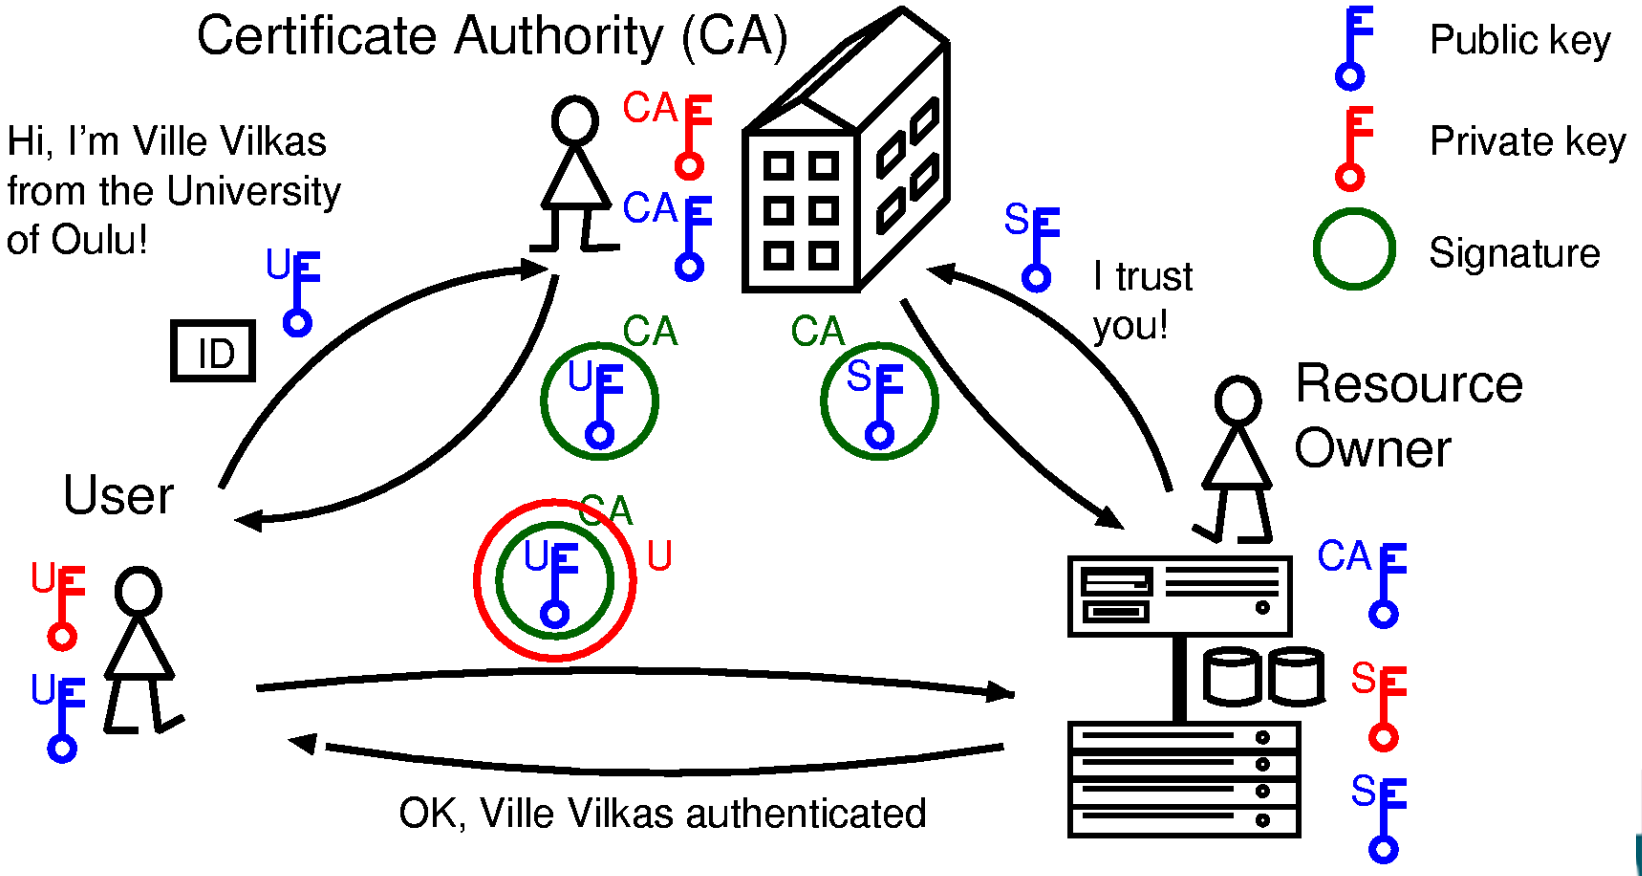

Authorization in Grids

- **Users are grouped into Virtual Organizations (VO)**
  - Based on common research area, nationality, funding agency or project
  - The same user can belong to several VOs
- **Resource providers grant access to VOs**
  - Scales better than managing individual users at every resource
  - Implies trust towards the organization managing the VO
  - E.g. in the Finnish M-grid the users of each site form one VO, and we can combine all to a larger "M-grid VO" for negotiating resource sharing with external parties

Grid Security

- **Objective: Convenient but secure access to different kinds of resources**
  - Grid account is a pass to computers beyond organizational domains!
  - Great power => great damage
- **Security aspects have been considered in the technology**
  - Strong authentication and encryption: no plain-text passwords
  - Identity usually tied to a certificate: revocation blocks access in the whole grid
- **Implementation details may vary from site to site**
  - Be careful if your data is sensitive!

Resource Sharing

- **Many different models are in use**
  - Anarchy: for example local resources in laboratories — relies on solidarity and personal relations
  - Centralized allocation within an organization, organization level agreements
  - Giving away free cycles while local jobs have higher priority: a model used in several NorduGrid clusters
- **Challenges in resource allocation and sharing**
  - User friendliness
  - Maximal resource utilization rate
  - Technical implementation, lack of standards

Grid Standardization

- **Standardization efforts led by Open Grid Forum, OGF (formerly GGF)**
  - Three big yearly meetings, open for everyone
  - Participants from both academia and industry
- **Existing standards used when possible, but many new grid-specific standards still emerging and not ready**

  => Grid middleware packages often incompatible with each other

Example: The Finnish 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, access through grid interface since Jul 2005
- **First large initiative to put grid middleware into production use in Finland**
- **Platform: Linux based PC clusters**

CSC

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.**

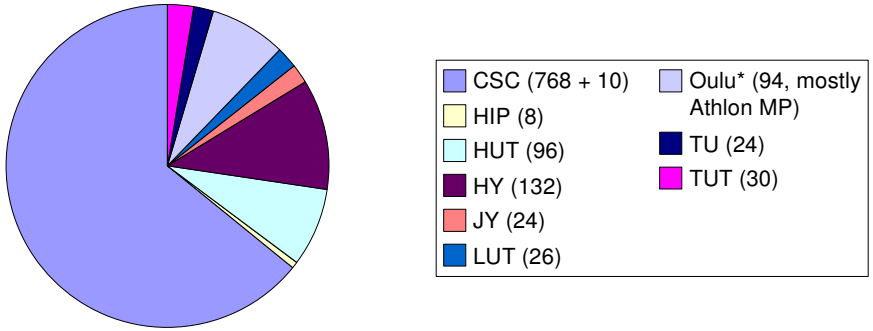

Operating System and Grid Middleware

- **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

Resource Sharing in M-grid

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

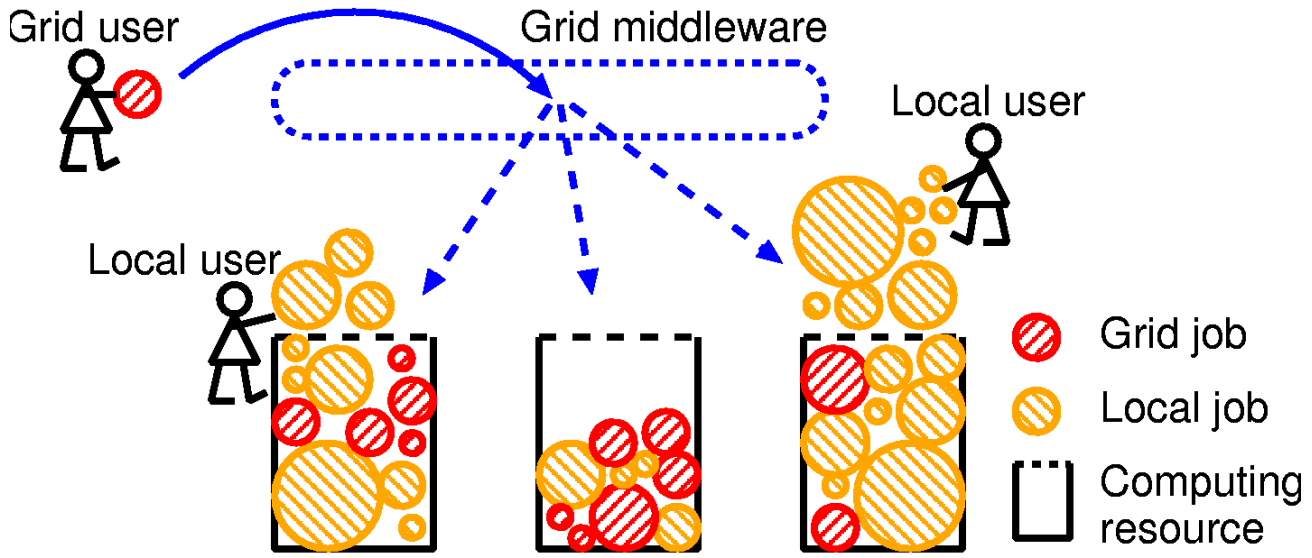

Grid Use July 2005 - May 2006

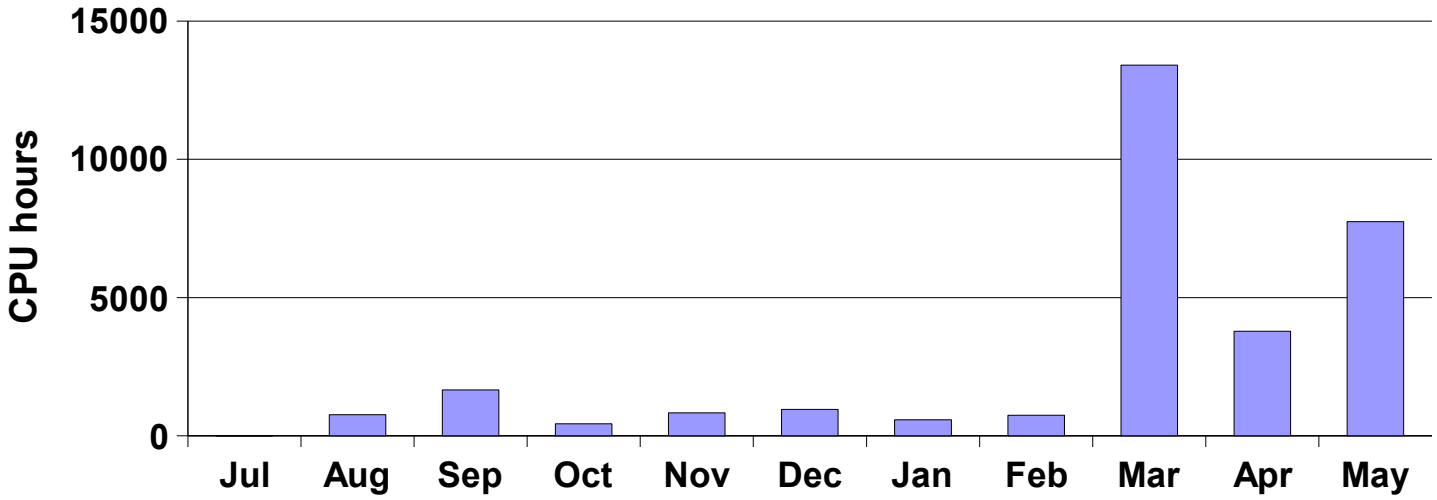

- **Total 31 000 CPU hours of jobs submitted through the grid interface: only 1.5% of the total use of the clusters**
- **Towards the end of the period grid use increased as some initial problems were solved**

M-grid Experiences

- **Only a handful of active grid users**
  - About one out of five users continued after initial tests
  - Most users still happily submitting jobs directly to the local batch queue instead of using the grid interface
- **Grid environment must be better than the existing one, otherwise nobody will use it!**
  - Users should not need to put extra effort on adapting their applications to the grid
  - Long queue in the local cluster and empty resources on the grid is a good enough incentive to some users
- **Collaboration model in system administration has been successful: grid projects always have other aspects than just the technology**

Obstacles to Grid Adoption

- **Need to request a certificate**
  - Not a very big problem: we had many more certificates requested than actual active users
- **Different job description syntax compared to the local batch queue system**
- **Higher failure rate and less determined execution times require job management tools**
  - Most of our active users were competent in scripting and developed their own job management aids
- **File management: shared disk is only visible within the local cluster (typical in many grid environments, some use distributed filesystems)**
- **Too few applications adapted for the grid**

CSC

Summary of M-grid Experiences

- **The M-grid systems are in heavy use locally, most users are satisfied**
- **Collaboration model in system administration has been successful**
- **Only a handful of active grid users, but they have been able to take advantage of otherwise unused capacity**
- **M-grid experiences will be used when planning new grid deployments in Finland**

Data Centric Grids

- **Computing power requirements may be minimal but data maintained at several locations**
- **Grid technologies may be useful for providing authenticated but uniform access to distributed datasets**
- **Example: EMBRACE project integrates databases and software tools in bioinformatics from different centers around the world**

Solving the Multiple Passwords Problem

- **Multiple usernames and passwords should always be avoided!**
  - Both inconvenient for users and a security risk
  - Many grids rely on personal certificates to provide single sign on capability, but identity federation is a user-friendly alternative (currently available for web based services but still in development for computing grids)
- **Example: Haka identity federation in the Finnish higher education**
  (following slides from Mikael Lindén and Arto Tuomi, CSC)

Haka Federation, the Identity Federation of Finnish Higher Education

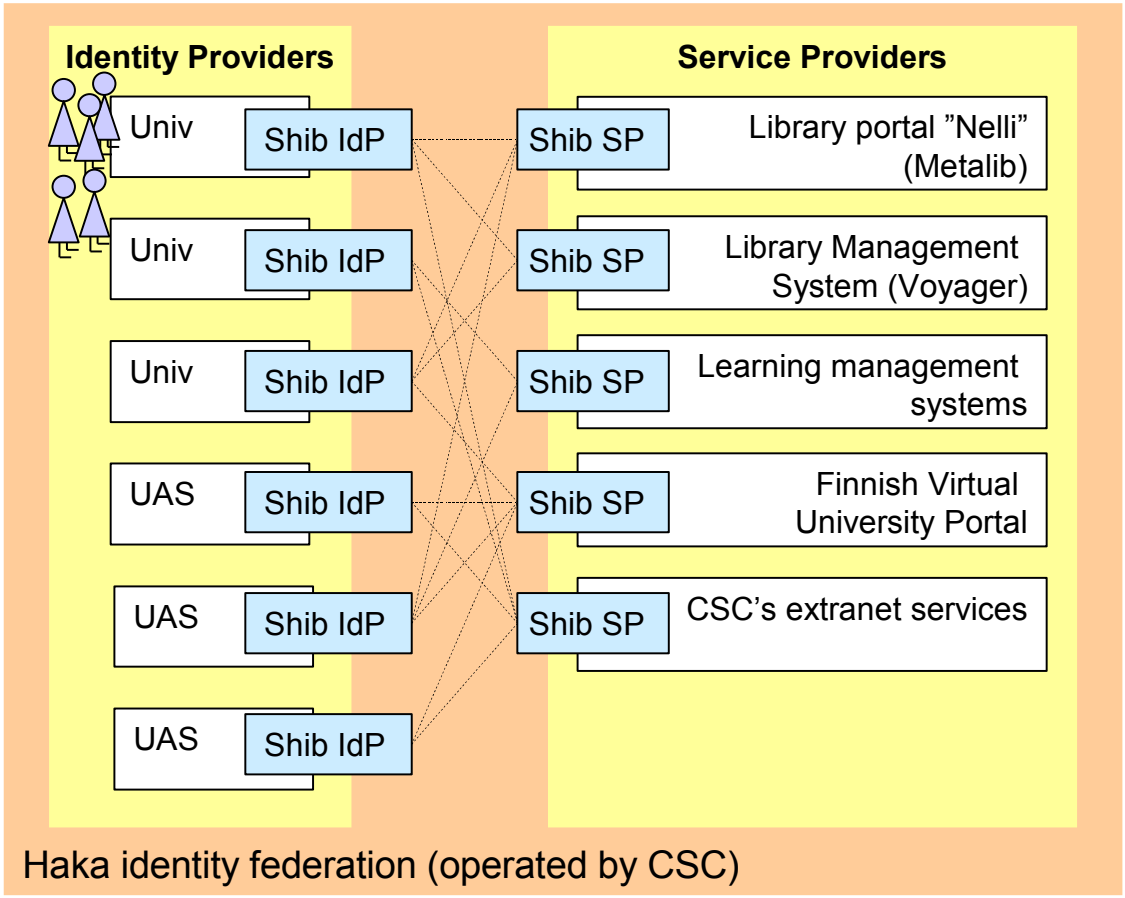

- User's home institution (Identity Provider, IdP) maintains user's identity and attributes (name, contact info, role, major etc)
- Home institution authenticates the user (e.g. by password)
- Home institution releases attributes to the Service Provider (on user consent)
- Based on the attributes, service provider decides what kind of service the user gets

Technology in Haka Federation

- **Protocol: Shibboleth ver 1.2/1.3**
  - protocol specified by Internet2 (US), based on standards (SAML 1.1, SOAP, XML)
  - free open source implementation by Internet2 available since 2003
  - version 2.0 compatible with Liberty Alliance
  - available for Apache and Microsoft IIS
- **Schema: funetEduPerson**
  - specification of attributes (name, mail, affiliation, study subject…)
  - based on eduPerson of Internet2
- **Certificates: Sonera CA (a commercial Finnish CA)**
  - only server certificates needed

Shibboleth in Action

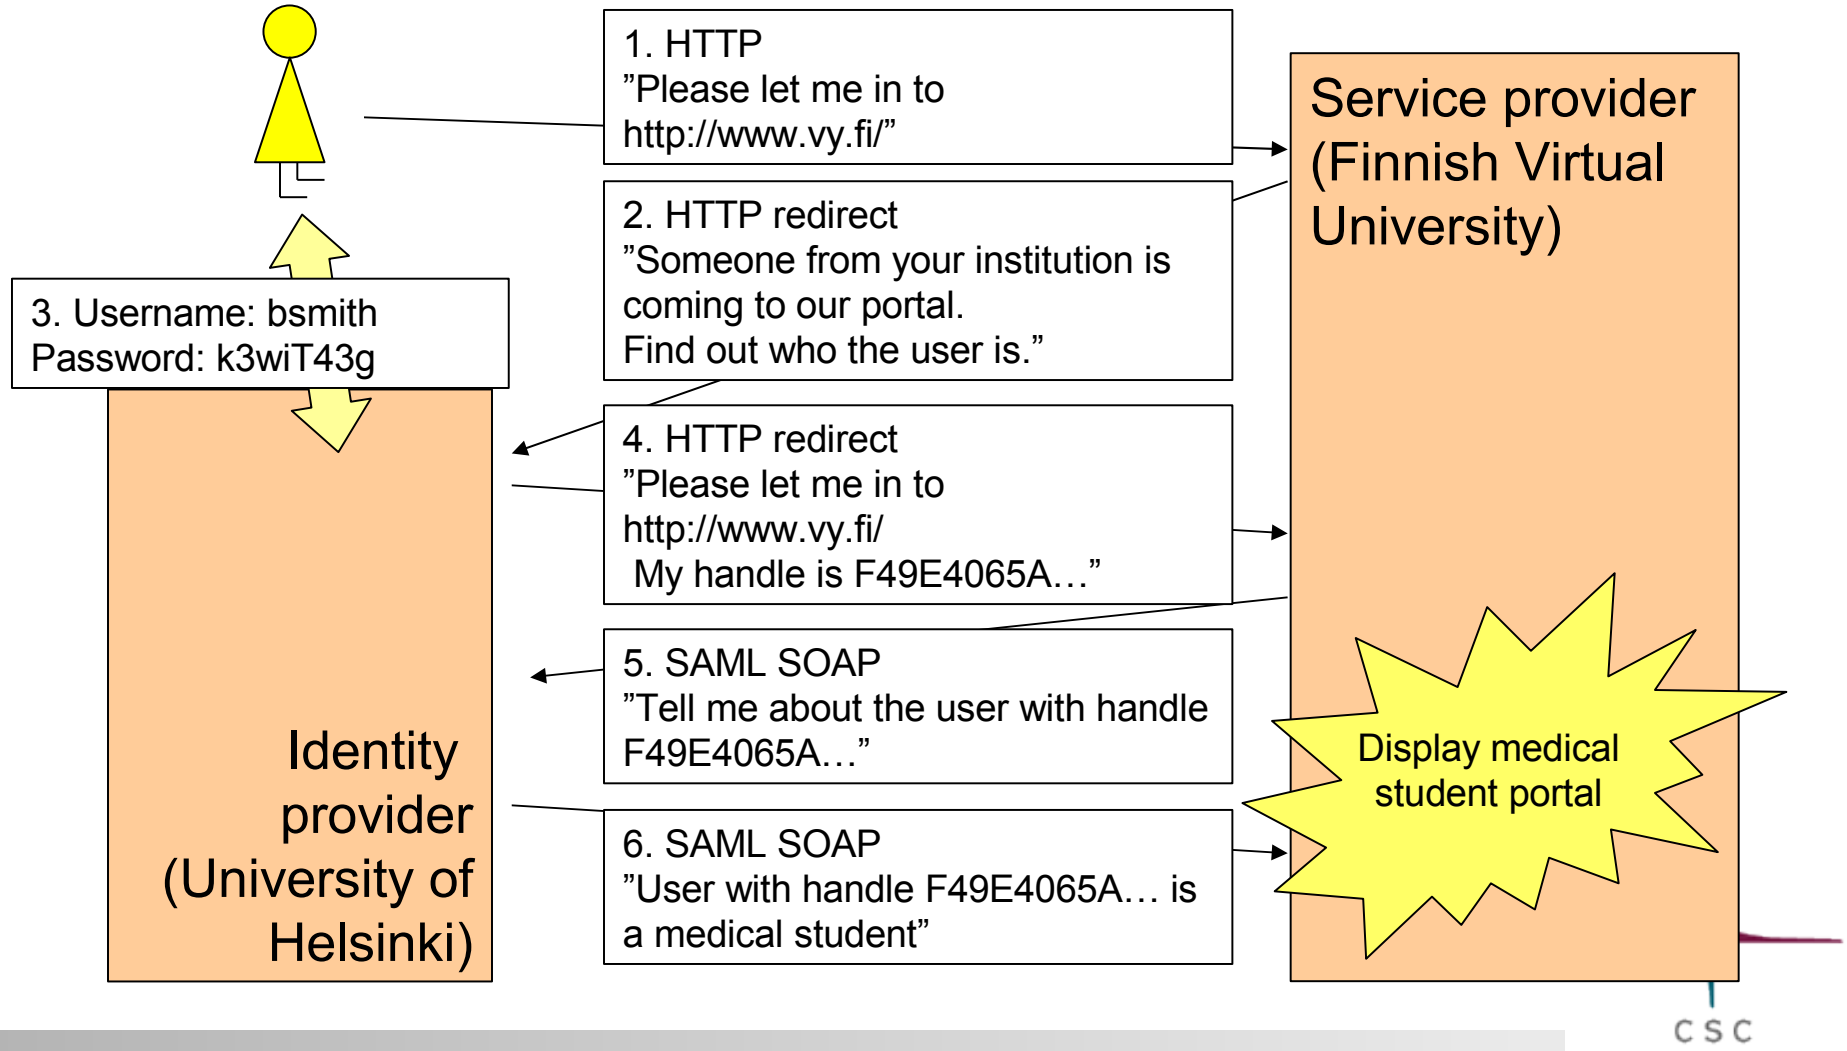

Status of Haka Identity Federation

- **Pilot federation operational 12/2003**
- **Production level federation operational 8/2005**
- **Current members: 13 universities (out of 20), about half of polytechnics**
  - Big universities; coverage 90% of end users in universities
- **Identity Providers (IdP) and Service Providers (SP)**
  - 11 IdPs
  - 13 SPs
  - ~160 000 logins to services in February 2007

Service Categories in Haka

1. **Library services**
   - The library management system (Voyager), the library portal (Metalib), the digital content repository (Encompass, work in progress)
   - The content providers (work in progress)
2. **eLearning services**
   - Learning management systems (Moodle, A&O, Optima)
   - Electronic application form for becoming a visiting student in another Finnish university (www.joopas.fi)
3. **Nationally provided services**
   - CSC's extranet services to researchers
   - Research funding application form (work in progress)
4. **ASP services in the administration of an institution**
   - Circulation of travel expence reports & incoming invoices (work in progress)
   - HR software/Employee self-service (work-in-progress)

Collaboration — Key to Success

- **Choosing right technology is important but even more important is good collaboration between partners**
- **Grid projects often require changes to current practices inside each organization**
  - Operational practices, user management etc. need to be compatible between participants
  - Partners must be able to trust each other
- **Need to involve people from many different levels in the discussions**
  - Users
  - System administrators, including network security staff
  - Management

Summary

- **A grid is the combination of networked resources and the corresponding grid middleware, which provides grid services for the user**
- **Problems which can be parallelized into independent subtasks are the most suitable for grid environments**
- **Grid environment must be better than the existing one, otherwise nobody will use it**
- **Standards still emerging and middleware packages incompatible with each other but grids are already used to do real science**
- **Good collaboration between partners is crucial for grid projects to succeed**

More Information

- **Grid Today: http://www.gridtoday.com/**
- **International Science Grid This Week: http://www.isgtw.org/**
- **Open Grid Forum: http://www.ogf.org**
- **M-grid homepage: http://extras.csc.fi/mgrid/**
- **Haka homepage: http://www.csc.fi/hallinto/haka**
- **Thank you! Questions?**