# CSC computing environment

Kimmo Mattila

- **CSC computing environment**
  - **Sisu** supercomputer
  - **Taito** cluster
  - **Hippu** application server

  - Usage is free for researchers working in Finland (but you must register)
  - Possibility to work with terabyte level datasets
  - Plenty of bioinformatics tools available
  - Usage through linux command line

# Connecting Servers of CSC

- Terminal connections ( ssh, PuTTY, SUI)
  - Usage through typed commands
  - Graphics requires Xterm connection
- Scientist's User Interface
  - Usage through web interface
  - Mostly used for managing your account and files
  - No bioscience applications
- FreeNX virtual desktop
  - Requires local client installation
  - Norman terminal connection can be used
  - Enables using grapical interfaces and displaying images

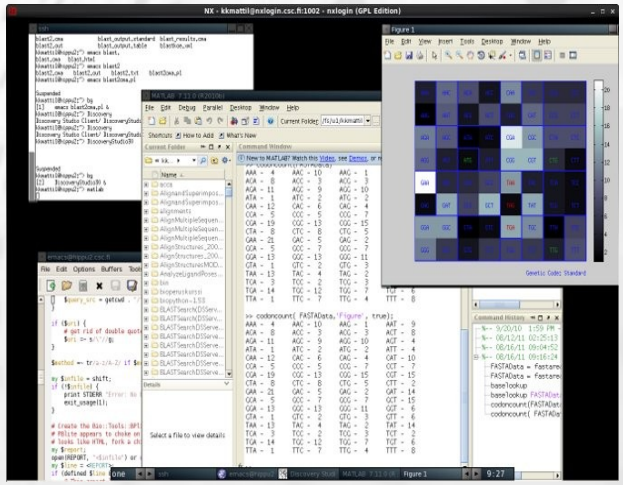

# Hippu

- 2x HP ProLiant DL58 G7 (Hippu3, Hippu4)
  - 4x 8-core Intel Xeon X7560/node
  - 64 cores total
  - 1 TB shared memory/node
- Meant for interactive jobs
  - job length not limited
  - no queue system installed
  - Plenty of bioinformatics tools installed
- Will be replaced during 2014
- Hippu user's guide:
  - http://www.csc.fi/english/pages/hippu_guide


# Sisu

- Cray XC30 Massively Parallel Processor (MPP) supercomputer
  - 1472 eight-core 2.6-GHz Intel Sandy bridge 64-bit processors
  - 11776 cores
  - 2 GB memory/core
  - Aires interconnects
- Meant for jobs that parallelize well
  - normally 64-512 cores/job (MPI)
  - can be increased for Grand Challenge projects
- Modest selection of bioinformatics tools
  - Molecular dynamics codes: gromacs, namd, Amber
- Sisu user's guide

http://research.csc.fi/sisu-user-guide

# Taito

- HP CP4000 BL ProLiant supercluster
  - Node: 2 x 8-core 2.6 GHz Intel Sandy Bridge 64-bit processors
  - 560 nodes with 64 GB memory (4 GB/core)
  - 16 nodes with 256 GB memory (16 GB/core)
  - 4 login nodes with 64 GB memory (4 GB/core)
  - Total of 9280 cores
- Meant for serial and mid-size parallel jobs
  - 1-256 cores/job (more posible after scalability tests)
- Wide selection of Bioinformatics tools
- Taito user's guide

http://research.csc.fi/taito-user-guide

# Taito cluster

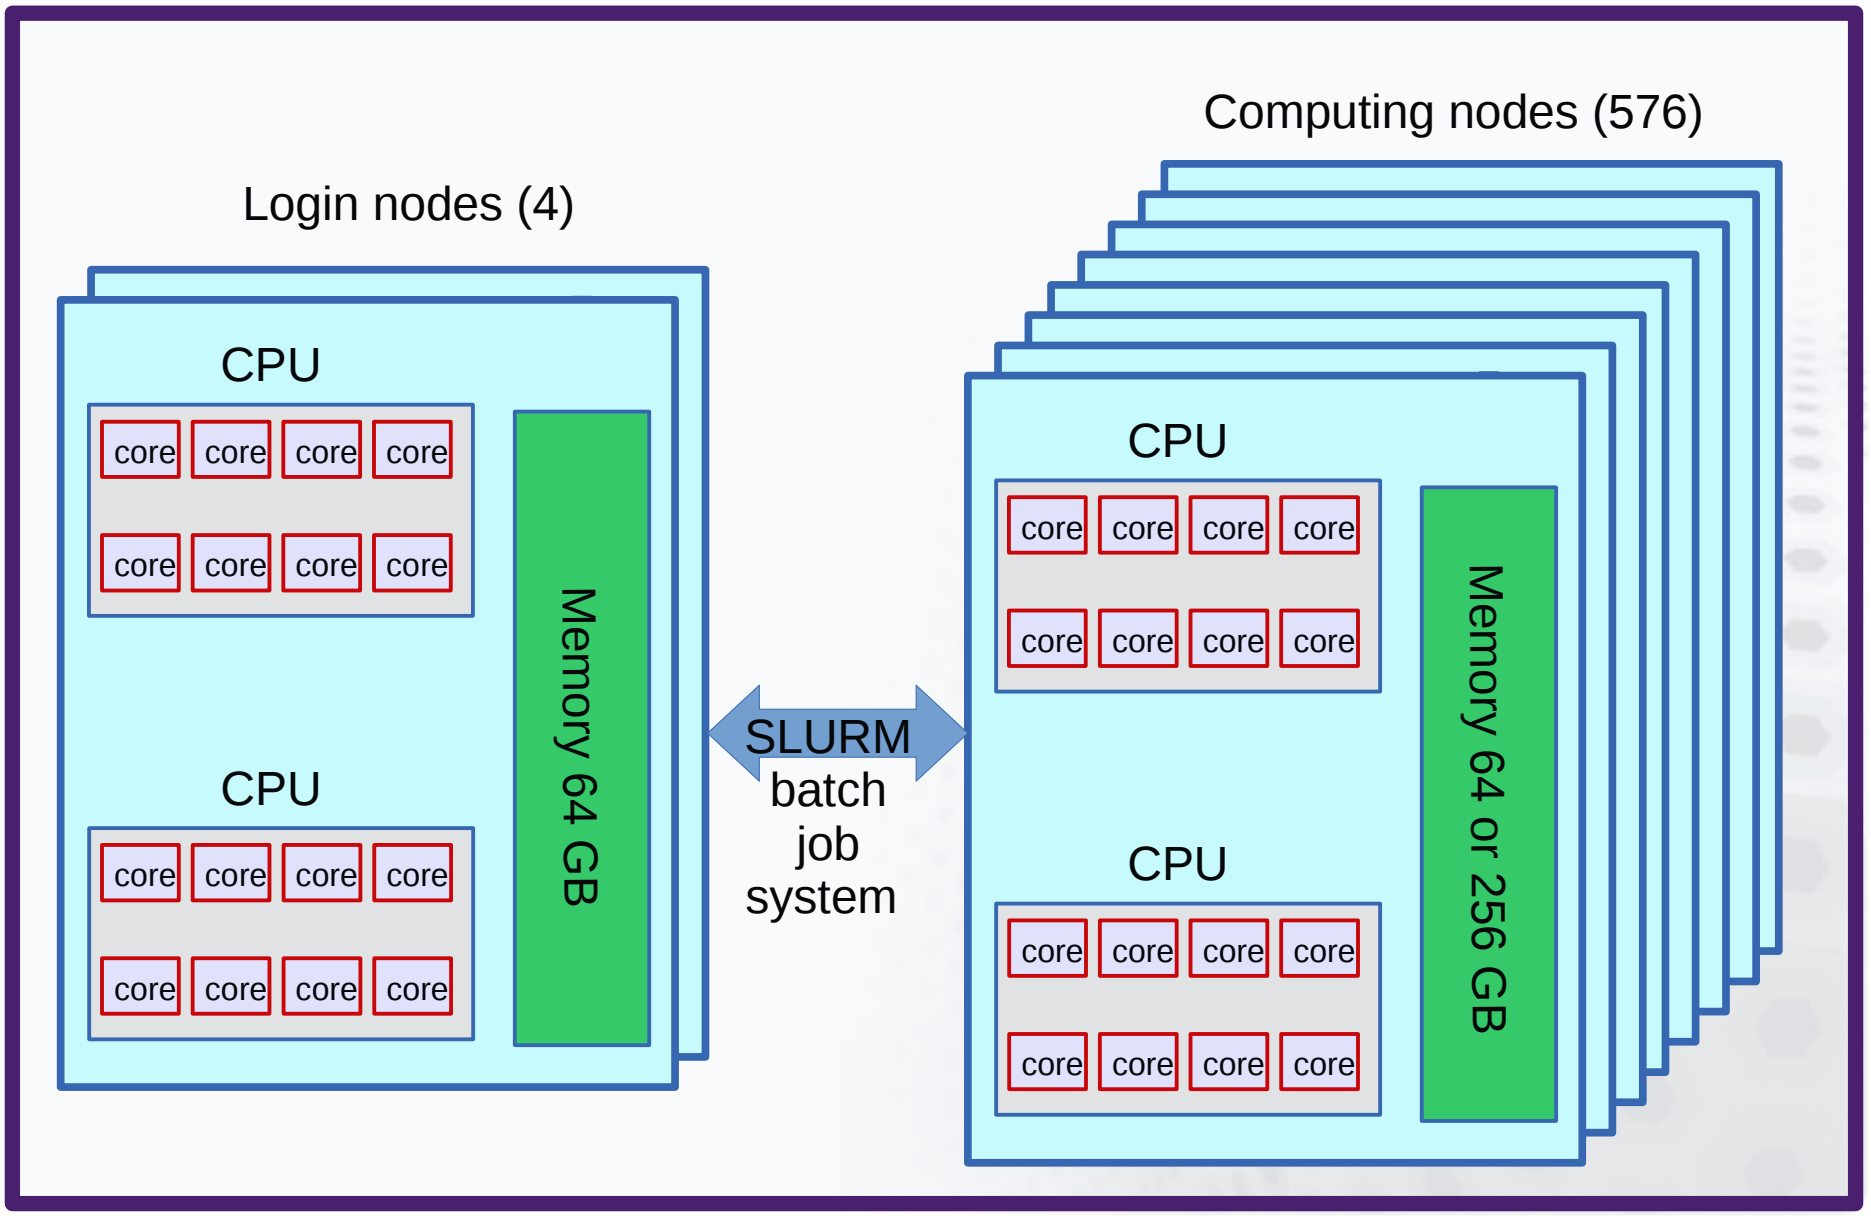

# Using Taito

- All “real computing” should be done through the batch job system. Login nodes are just for submitting jobs
- **module load biokit** sets up most of the bioinformatics tools
- Own software installations are possible ($USERAPPL)
- $WRKDIR for processing data
- HPC Archive and IDA for long term storage ab backup

# Using Taito

Default user specific directories in Sisu and Taito

| Directory or storage area | Intended use | Default quota/user | Storage time | Backup |
|---|---|---|---|---|
| $HOME | Initialization scripts, source codes, small data files. Not for running programs or research data. | 20 GB | Permanent | Yes |
| $USERAPPL | Users' own application software. | 20 GB | Permanent | Yes |
| $WRKDIR | Temporary data storage. | 5 TB | Until further notice. | No |
| $TMPDIR | Temporary users' files. | | 2 days | No |
| project | Common storage for project members. A project can consist of one or more user accounts. | On request. | Permanent | No |
| HPC Archive* | Long term storage. | 2 TB | Permanent | Yes |

# Batch jobs

- User has to specify necessary resources
  - In the batch job script or given as command line options for *sbatch* (or a combination of script and command line options)
- Resources need to be adequate for the job
  - Too small memory reservation will cause the job to use swap memory (very slow) or fail.
  - When the run time limit is exeeded, the job will be terminated whether finished or not
- Requested resources can affect the time the job spends in the queue
  - Especially core number and memory reservation
- Realistic resource request give best results
  - Not always easy to know beforehand
  - Usually best to try with smaller tasks first and check the used resources

# Batch jobs

Sample batch job file for TopHat

```
#!/bin/bash -l
#SBATCH -J TopHat
#SBATCH -o output_%j.txt
#SBATCH -e errors_%j.txt
#SBATCH -t 24:00:00
#SBATCH -n 1
#SBATCH --nodes=1
#SBATCH –cpus-per-task=4
#SBATCH –mem-per-cpu=4000
#

module load biokit

tophat -p $SLURM_CPUS_PER_TASK HS_bwt2_index reads1.fq reads2 -o tophat_results
```

Job is submitted with command:

```
sbatch batch_job_file.sh
```

# Batch jobs in Taito

| Queue | Number of cores | Maximum run time |
|---|---|---|
| serial | 16 ( one node) | 3 days |
| parallel | 448 (28 nodes) | 3 days |
| longrun | 16 (one node) | 7 days |
| test | 32 (two nodes) | 30 min |

# Using parallel computing Taito

- Embarrassingly parallel tasks:
  - Job can be split to numerous sub jobs
  - You can use *array jobs* and splitting utilizing tools like *pb blast, cluster_interproscan*, *trinity*, *miso.*
- Threads/ OpenMP based parallelization
  - Many bioinformatics tools use this approach. Bowtie2, BWA, Tophat, ....
  - All the parallel processes must see the same memory -> all processes must run within one node -> can utilize max 16 cores
  - Applications rarely benefit from more than 4-8 cores
- MPI parallelization.
  - Shared memory -> can utilize several nodes
  - Check scaling before launching big jobs

# IDA and HPC archive Storage services

# Moving data to and from CSC

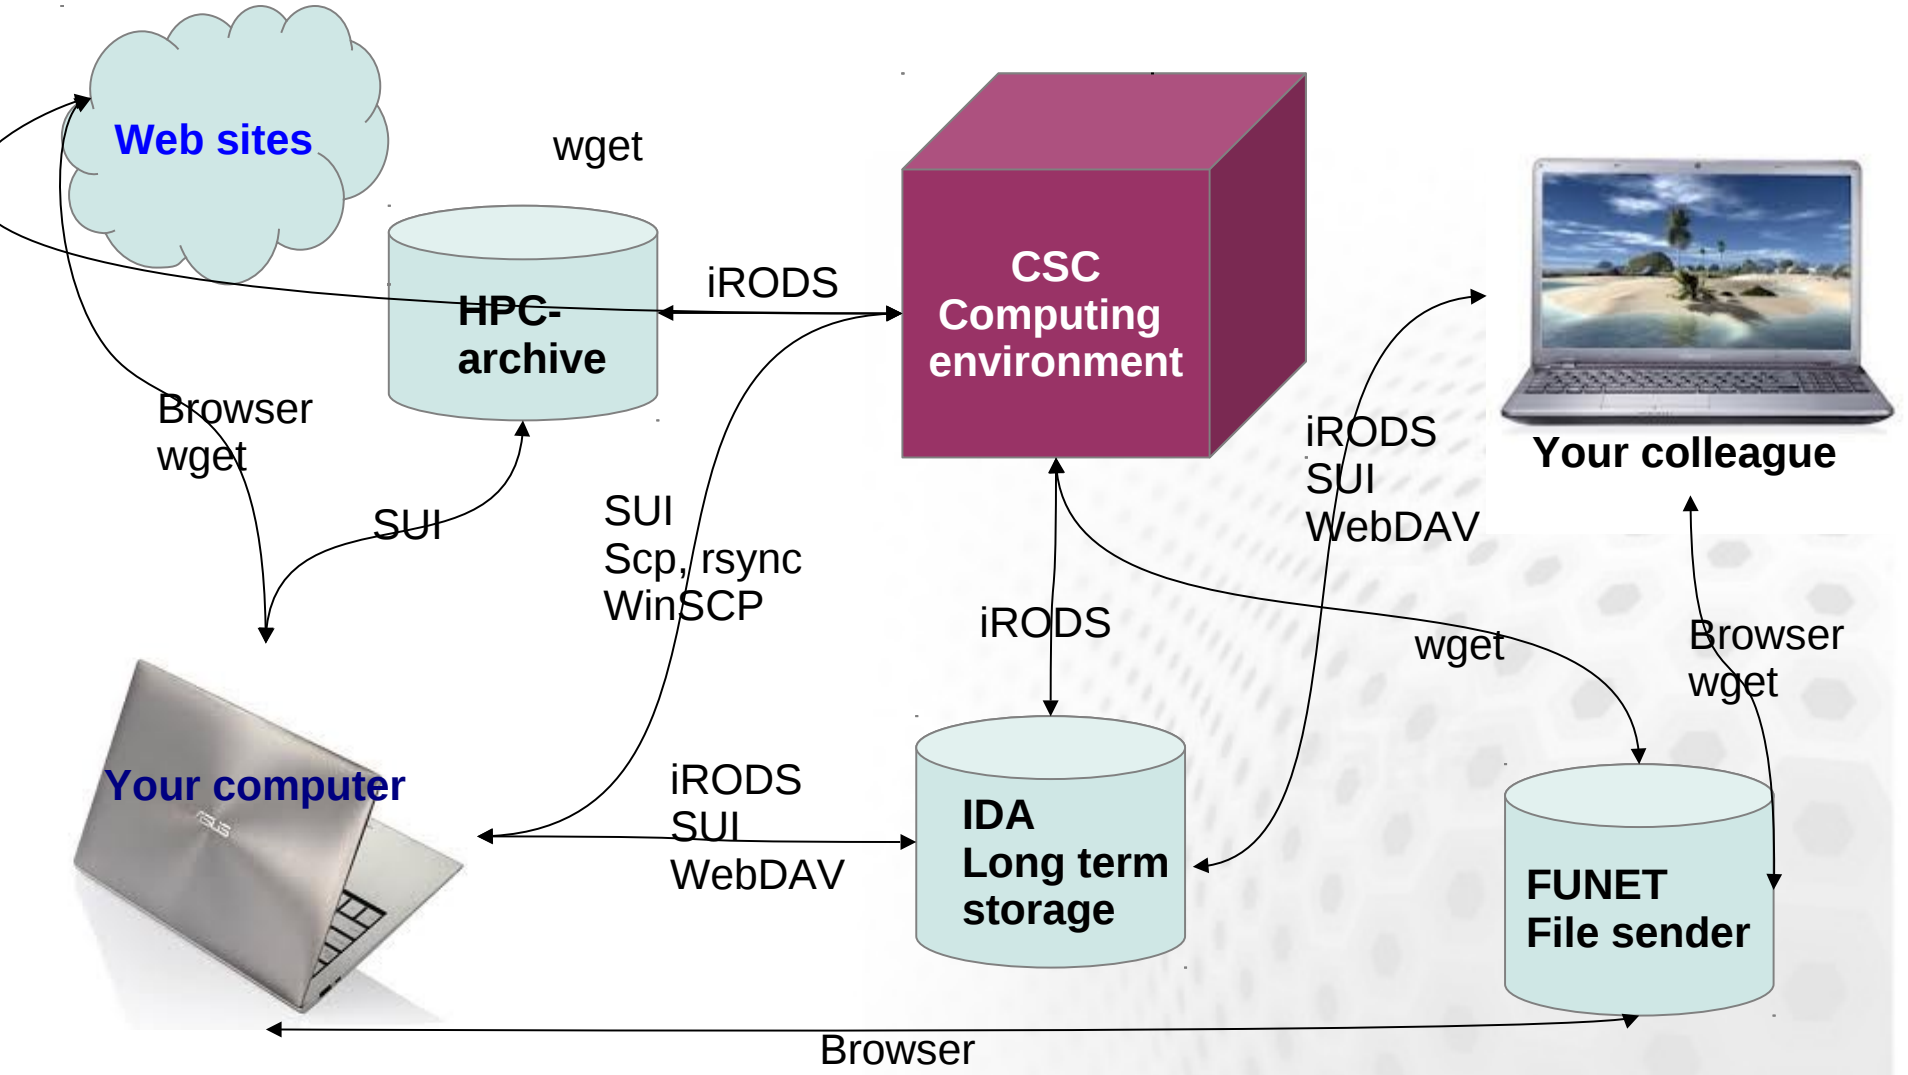

# IDA storage service

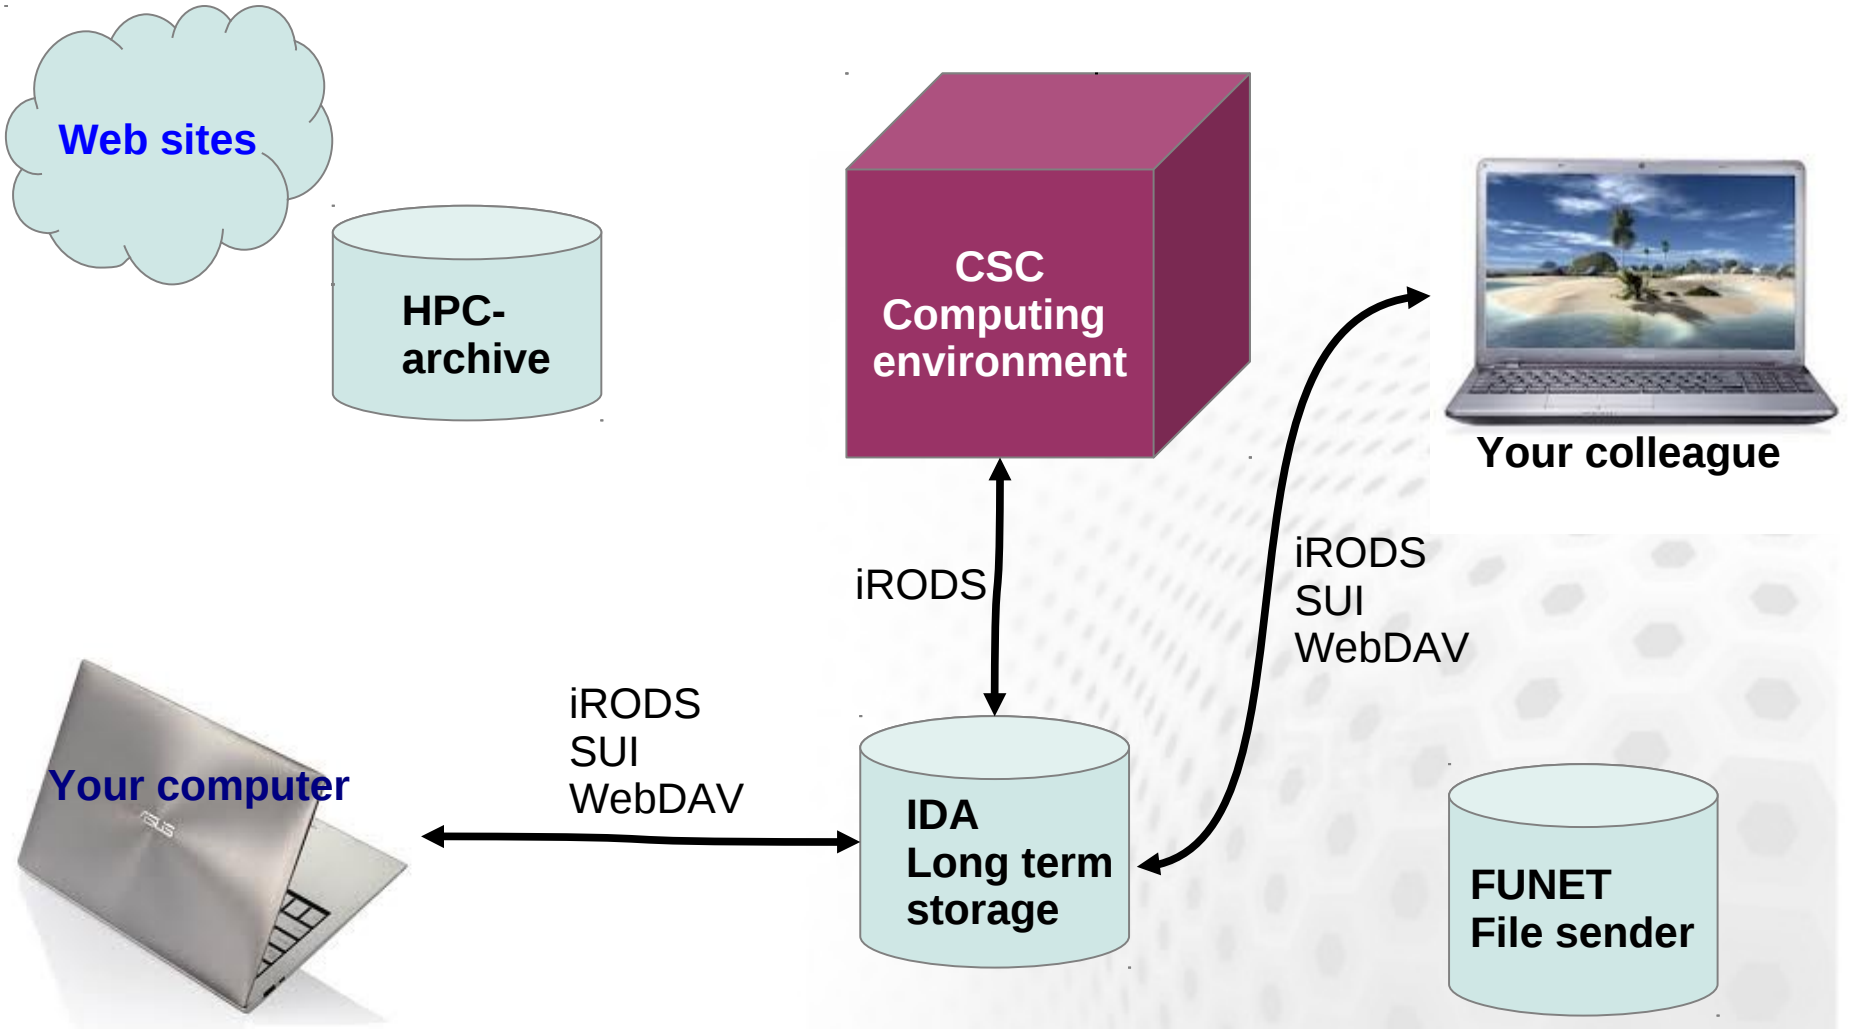

# HPC Archive and IDA

- **IDA**

- Storage service for research data
- quotas are grated by the Universities and Academy of Finland
- several different interfaces
- accessible through normal network connections
- part of the Tutkimuksen tietoaineistot (www.tdata.fi)

- **HPC Archive**
  - Intended for CSC users
  - 2TB / user
  - Replaces the $ARCHIVE
  - Only command line interface to the CSC servers

# IDA storage service

- iRODS based storage system for storing, archiving and sharing data
- The service was launched 2012
- Usage through personal accounts and projects
- Each project has a shared directory too
- Speed: about 10 GB/min at the servers of CSC
- CSC host's the service

Three interfaces:

- WWW interface in Scientists' User Interface
- network directory interface for Linux, Mac (and Windows XP)
- command line tools (i-commands installed at the servers of CSC)

# IDA interfaces at CSC

Some iRODS commands

- `iput` *`file`* — move file to IDA
- `iget` *`file`* — retrieve file from IDA
- `ils` — list the current IDA directory
- `icd` *`dir`* — change the IDA directory
- `irm` *`file`* — remove file from IDA
- `imv` *`file file`* — move file inside IDA
- `irsync` — synchronize the local copy with the copy in IDA
- `imkdir` — create a directory to IDA
- `iinit` — Initialize your IDA account

**IDA In Scientist's User Interface**

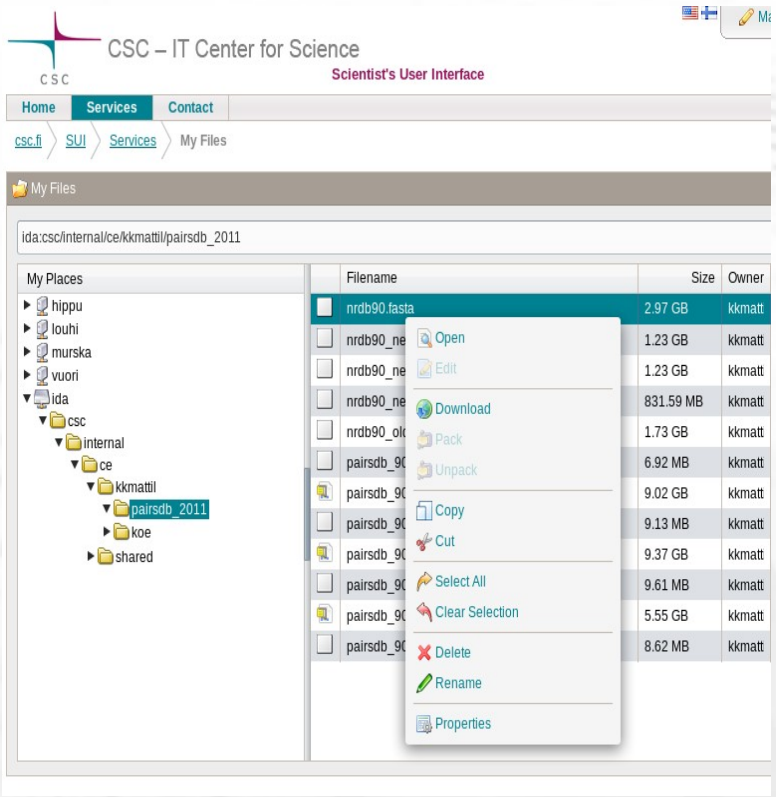

# Some brief generalizations:

- It’s usually faster to move one large file than many small ones
- On the other hand you should avoid too large files
  - it’s nicer to re-send one 10 GB chunk than the whole 100 GB file
- Connsider compression
- Create a hierarcical data structure to your archive.
- Data should be packaged for saving in Archive server.

# A National Computing Grid: FGI

Ulf Tigerstedt, Kimmo Mattila, Luis Alves

# What is a computing Grid?

- **Site:**
  - One or more computing clusters joined together via network interfaces
  - Can also host other computing “**services**”
  - Generally geographically located at one place
- **Computing GRID:**
  - A number of sites (usually) geographically distributed where the frontends are connected via network interfaces to the GRID network (aka. internet)
  - Software to direct users' jobs to a site is referred to as **Grid Middleware**

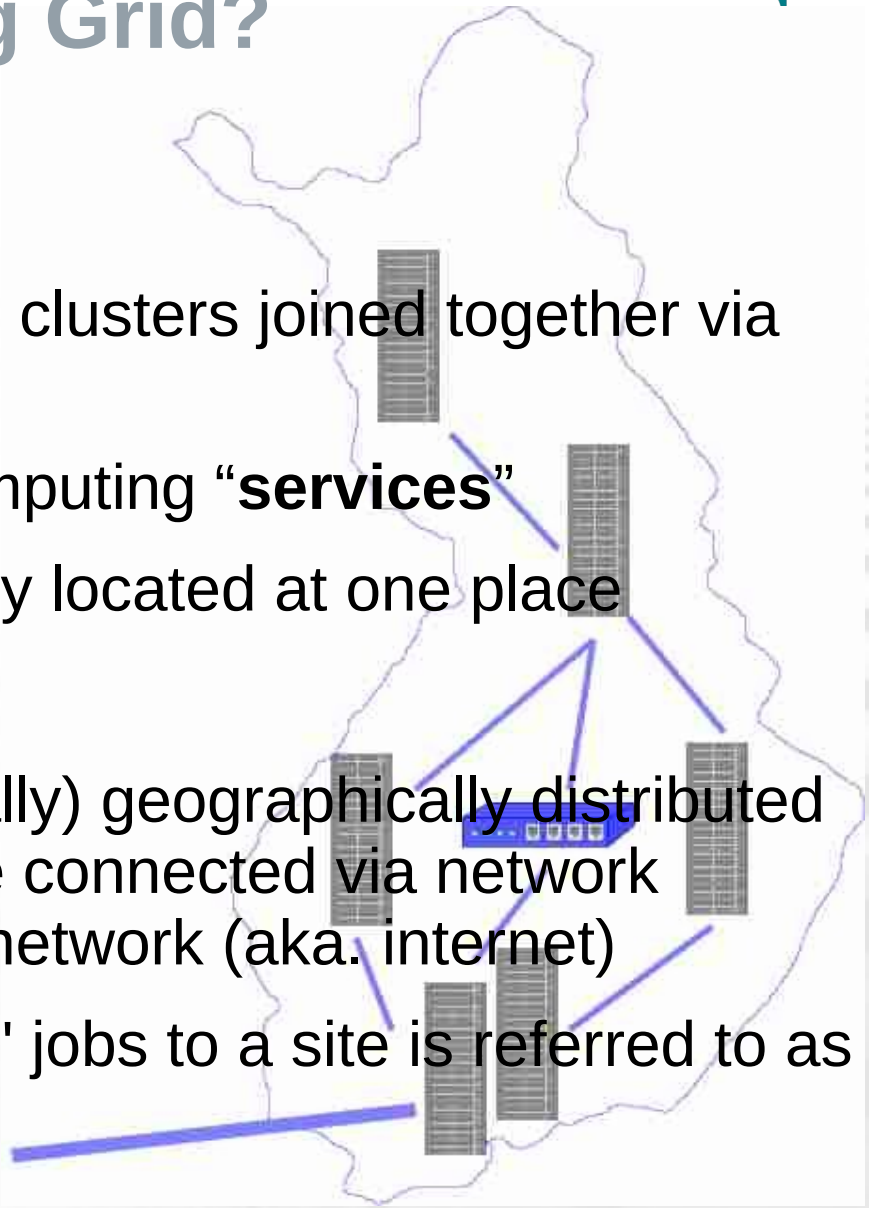

- **Virtual Organisation (VO):**
  - A community which has a common research or software requirement
    - Not geographically bound
    - Existing VOs can be easily joined
    - New VOs can be easily created

    (more on VOs further on)

# FGI

- FGI has been fully operational since spring 2012

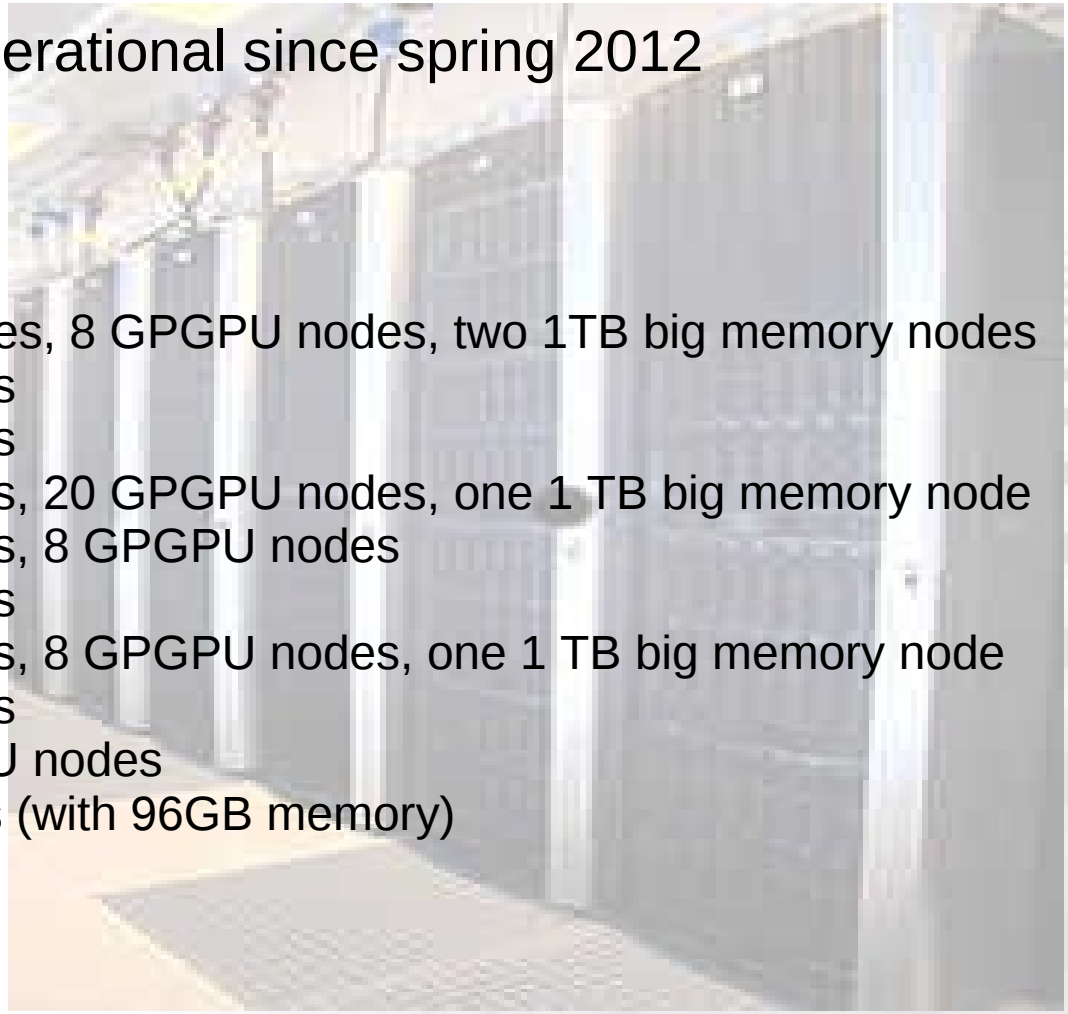

## FGI servers

| | |
|---|---|
| Aalto: | 112 nodes, 8 GPGPU nodes, two 1TB big memory nodes |
| Lappeenranta: | 16 nodes |
| Eastern Finland: | 64 nodes |
| Helsinki: | 49 nodes, 20 GPGPU nodes, one 1 TB big memory node |
| Jyväskylä: | 48 nodes, 8 GPGPU nodes |
| Oulu: | 30 nodes |
| Tampere (TUT): | 37 nodes, 8 GPGPU nodes, one 1 TB big memory node |
| Turku: | 20 nodes |
| Åbo Akademi: | 8 GPGPU nodes |
| CSC: | 24 nodes (with 96GB memory) |

# Normal clusters

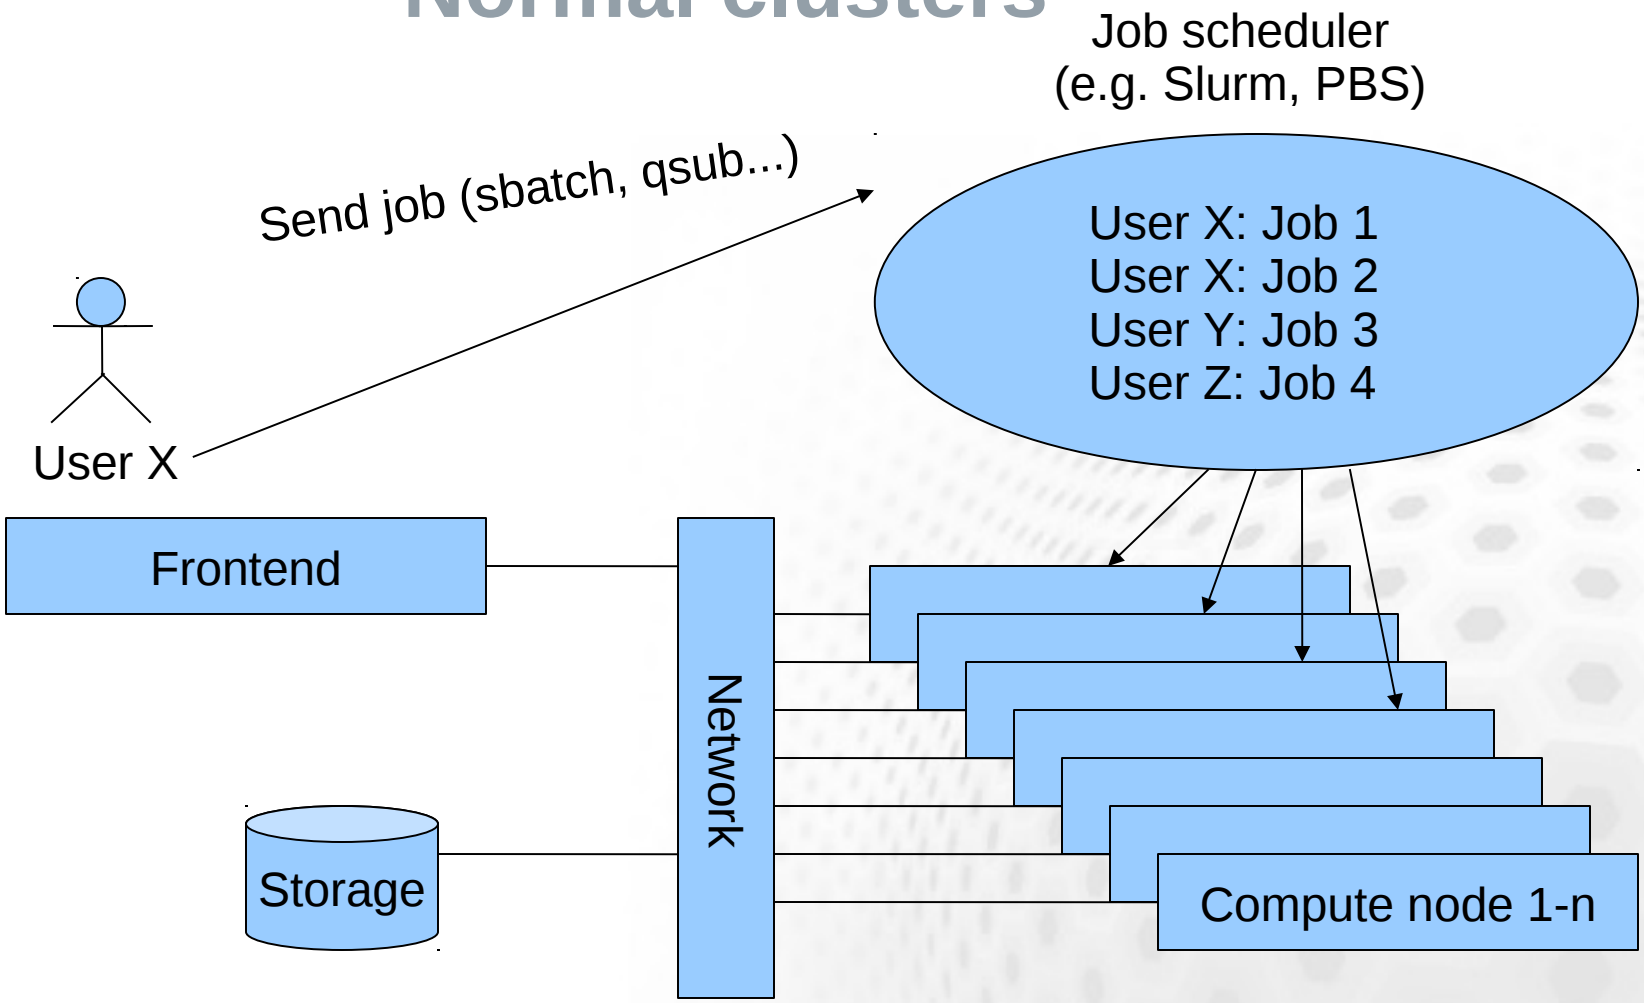

# Grids

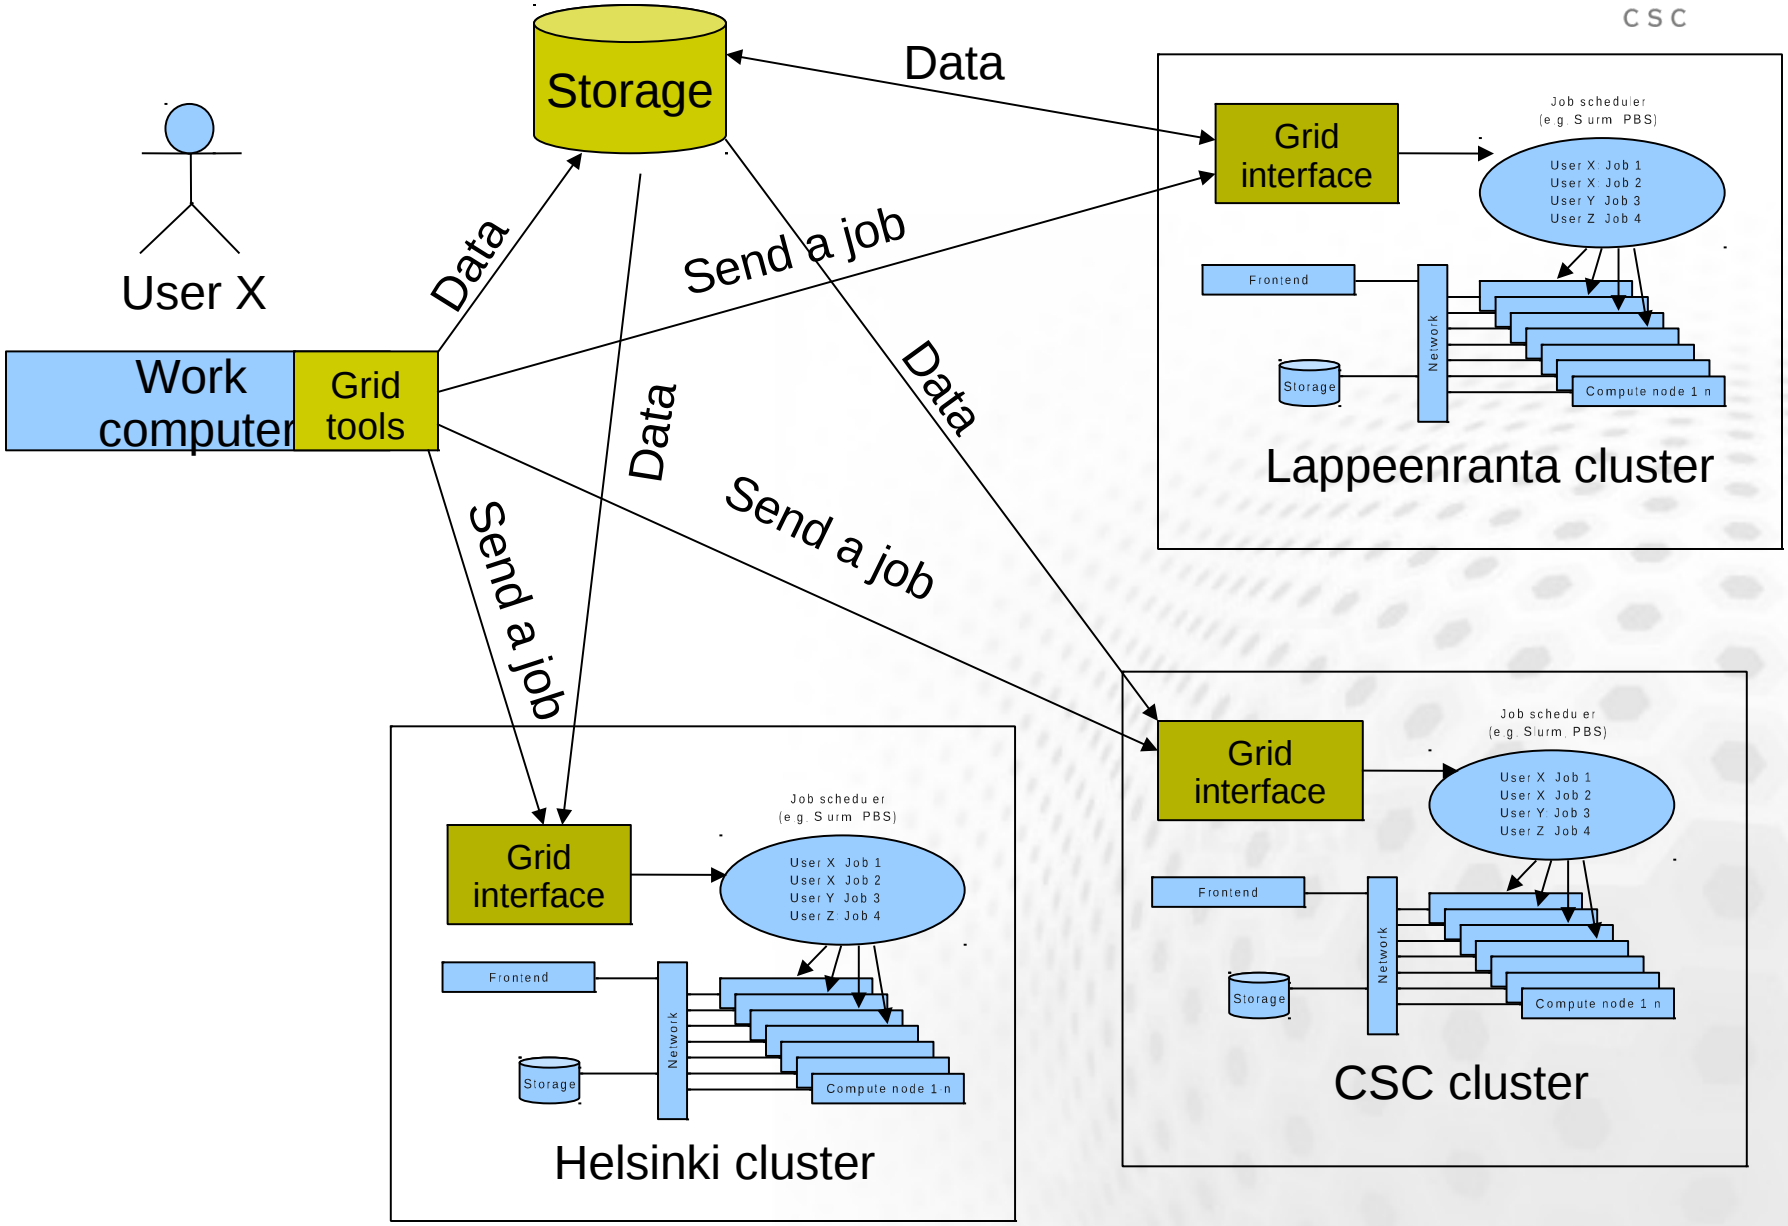
Storage
Data
User X
Data
Send a job
Work computer
Grid tools
Data
Data
Send a job
Send a job
Grid interface
Job scheduler (e.g. Slurm, PBS)
User X: Job 1
User X: Job 2
User Y: Job 3
User Z: Job 4
Frontend
Network
Storage
Compute node 1..n
Lappeenranta cluster
Grid interface
Job scheduler (e.g. Slurm, PBS)
User X: Job 1
User X: Job 2
User Y: Job 3
User Z: Job 4
Frontend
Network
Storage
Compute node 1..n
Helsinki cluster
Grid interface
Job scheduler (e.g. Slurm, PBS)
User X: Job 1
User X: Job 2
User Y: Job 3
User Z: Job 4
Frontend
Network
Storage
Compute node 1..n
CSC cluster

# Grids and ARC

- The ARC middleware is used in FGI
  - Server side
  - Client tools
- Also other grid middleware used in Europe
  - gLite
  - Unicore
  - condor

# User accounts in grids

- Grids do not use usernames and passwords
- Users authenticated using “certificates”
  - Easy to get a certificate
- FGI Users must join **fgi.csc.fi VO**
  - Resources give access to **VOs** not users
  - In Finland we use the  fgi.csc.fi VO
  - Simple to join (http://research.csc.fi/fgi-user-guide)

# What do you need?

- Certificate
- VO membership
- The ARC client tools
  - Installable on
    - most Linux versions
    - MAC OSX
  - ARC is available at CSC (Taito and Hippu )
    - Also available on your local cluster login node

# Starting with FGI

- FGI User guide: http://research.csc.fi/fgi-user-guide
- FGI user pages: http://confluence.csc.fi/display/FGI
  - Central place for all documentation and information about FGI
  - Getting started
  - Available software, and how to use it
- helpdesk@csc.fi
  - Problems? Requests?

# Software in FGI

- Some scientific software is pre-installed
  - Primarily open source software
- Runtime environment defines a software setup

(analogous to *environment modules* in clusters)

https://confluence.csc.fi/display/fgi/Runtime+Environments

- You can also run your own programs in FGI
- If you have suggestions, contact us

# Other ways to use the FGI

- Arcrunner: automatic job submission tool for large grid-job sets
- Automatic command line interfaces for: AutoDock, BLAST, BWA, InterProScan, SHRIMP and Exonerate
- Matlab Compiler Runtime
- batch script wizard on SUI!

# Scientist's User Interface: SUI

## http://sui.csc.fi

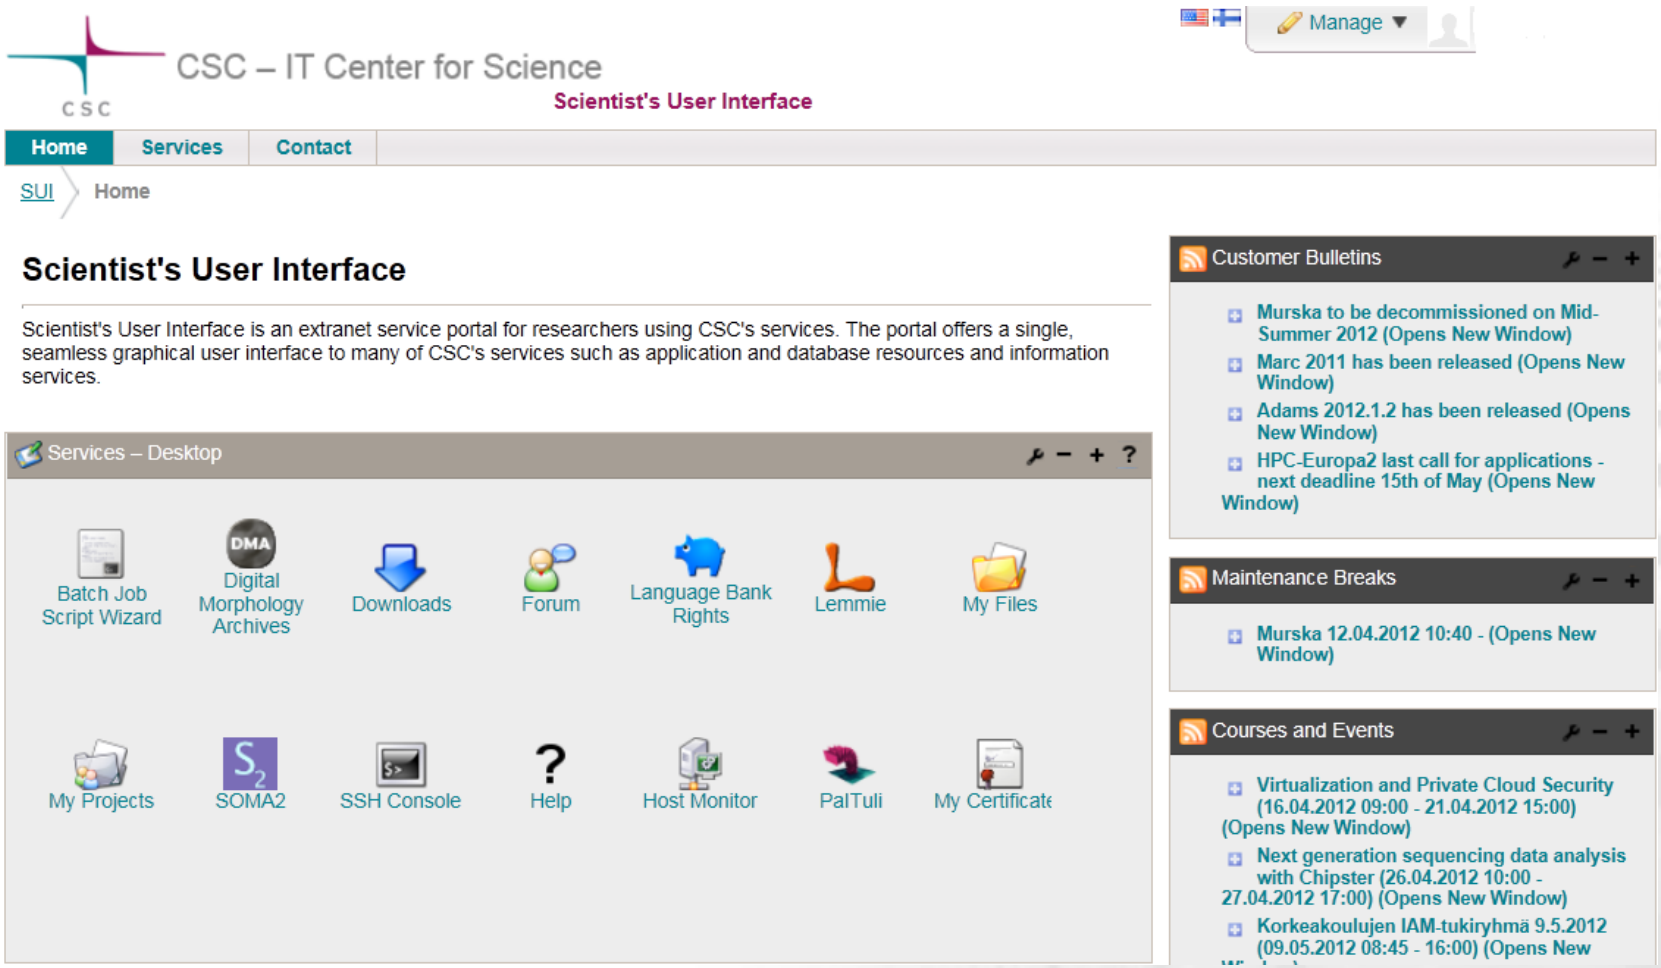


More info at: http://www.csc.fi/english/research/SUI_presentation/download

# ARC batch script tool now available!

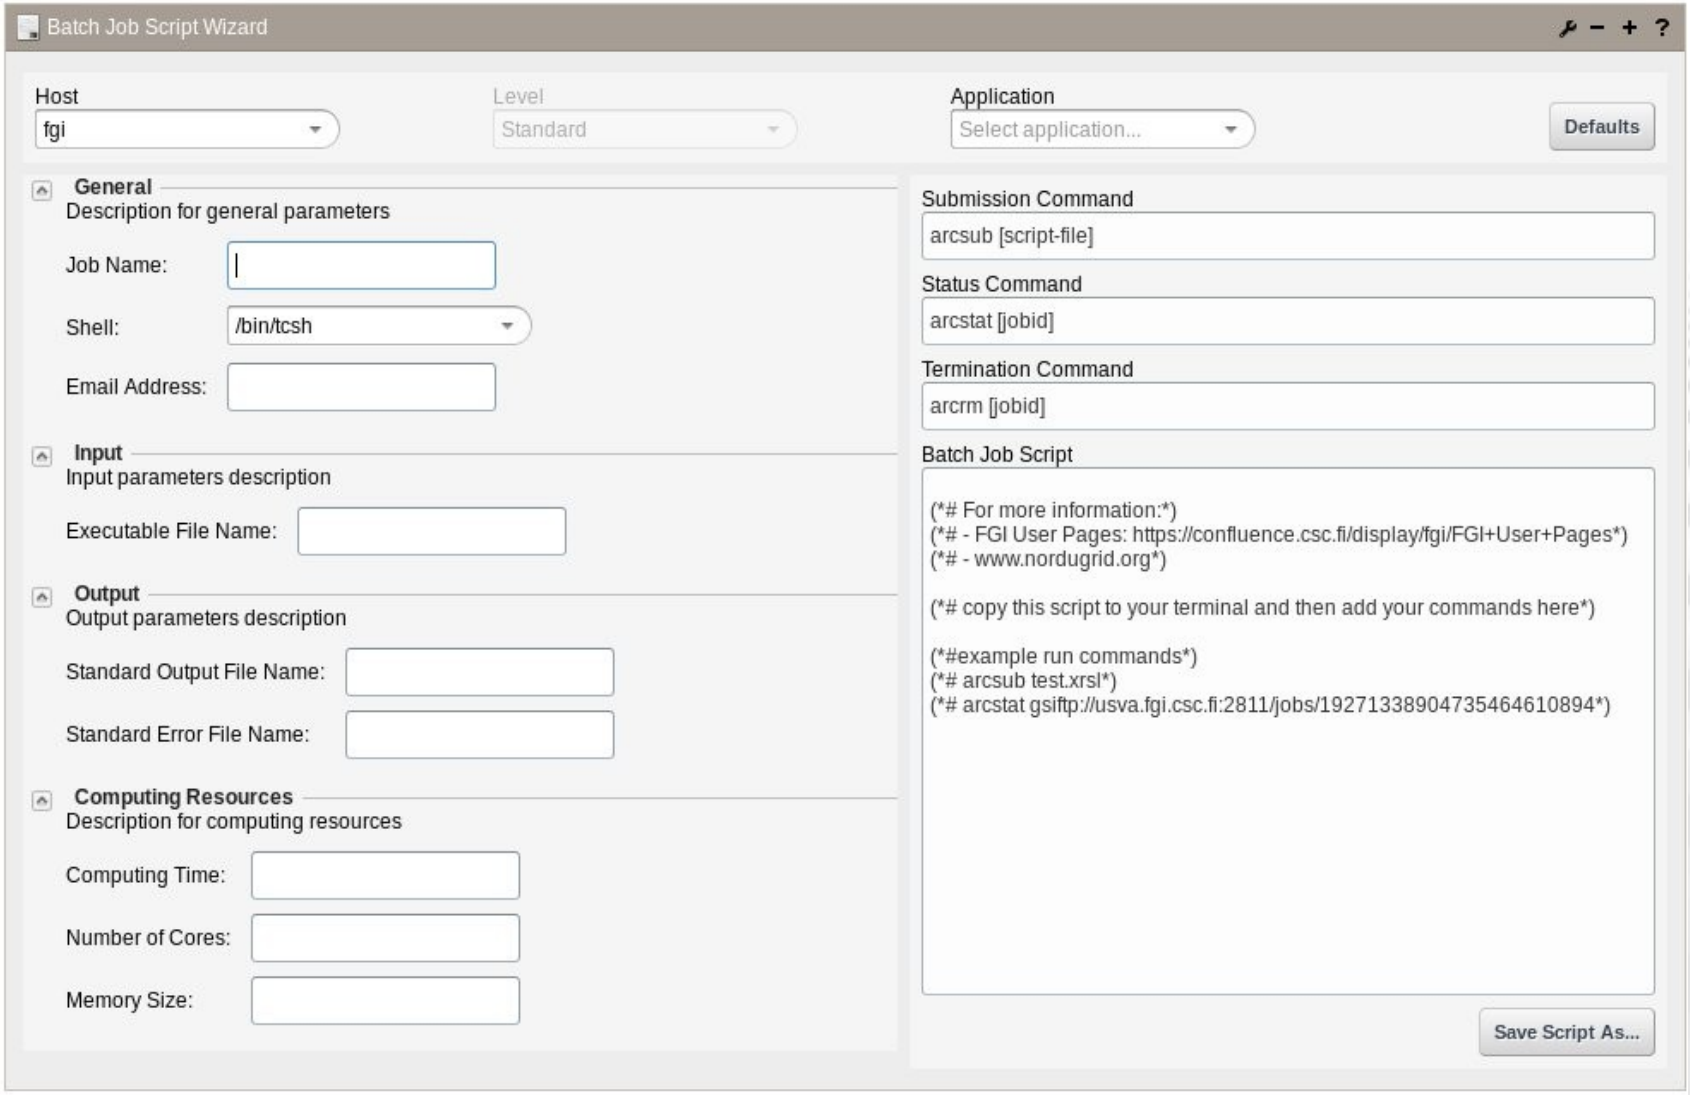

# ARC batch script tool now available!