Distributed Computing for Small Experiments

Daniela Bauer
Imperial College London

ICNFP 2017

Imperial College London

GridPP

Overview

- What is distributed computing ?
- The Grid – and what it can do for you.
- How does it work ?
- (Worked) Examples.

Distributed Computing

Local Machine

Distributed Computing

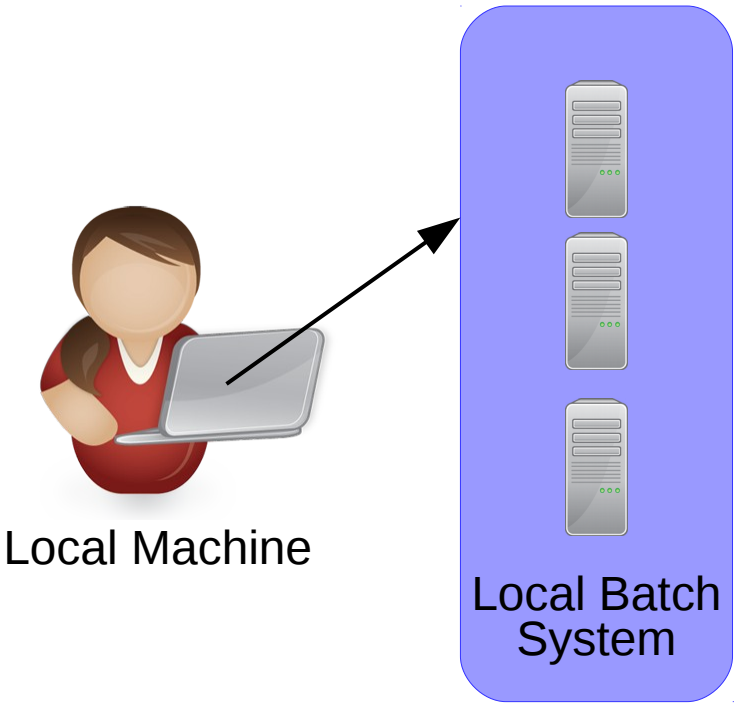

Distributed Computing

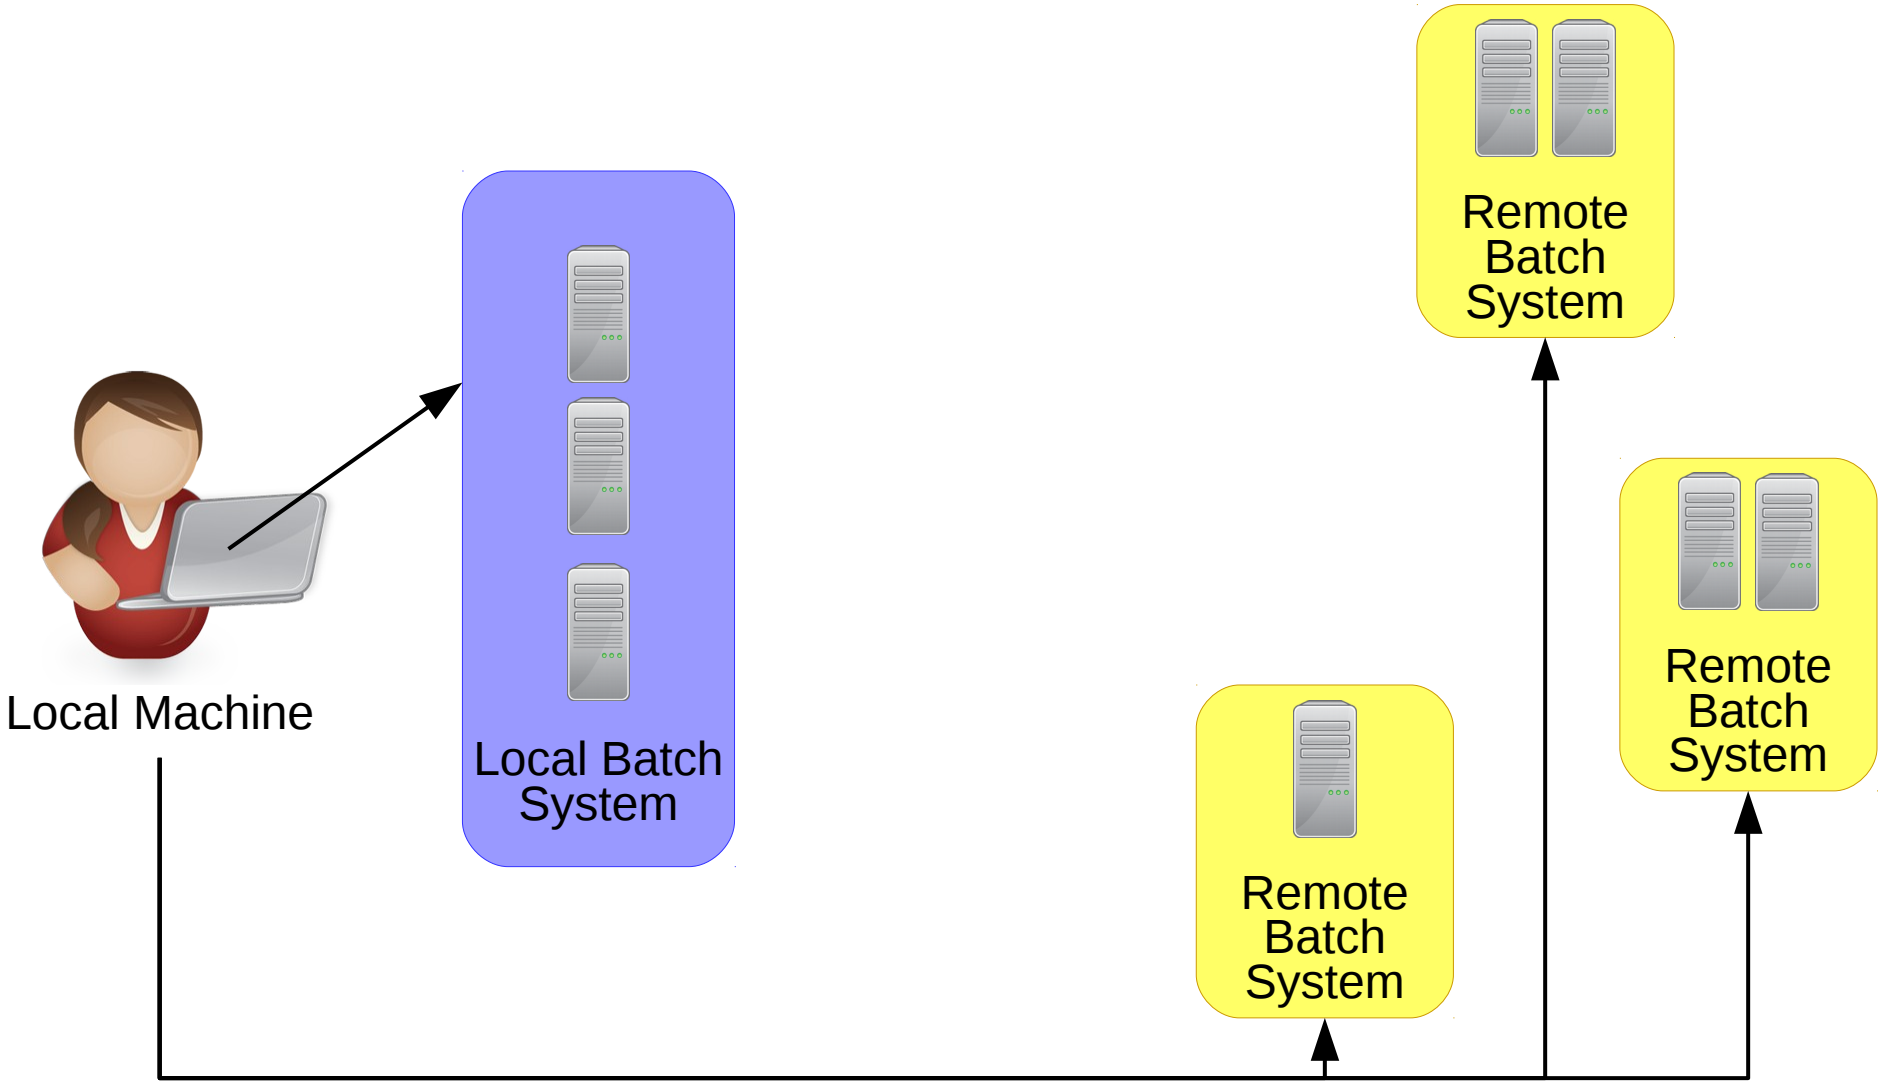

Distributed Computing

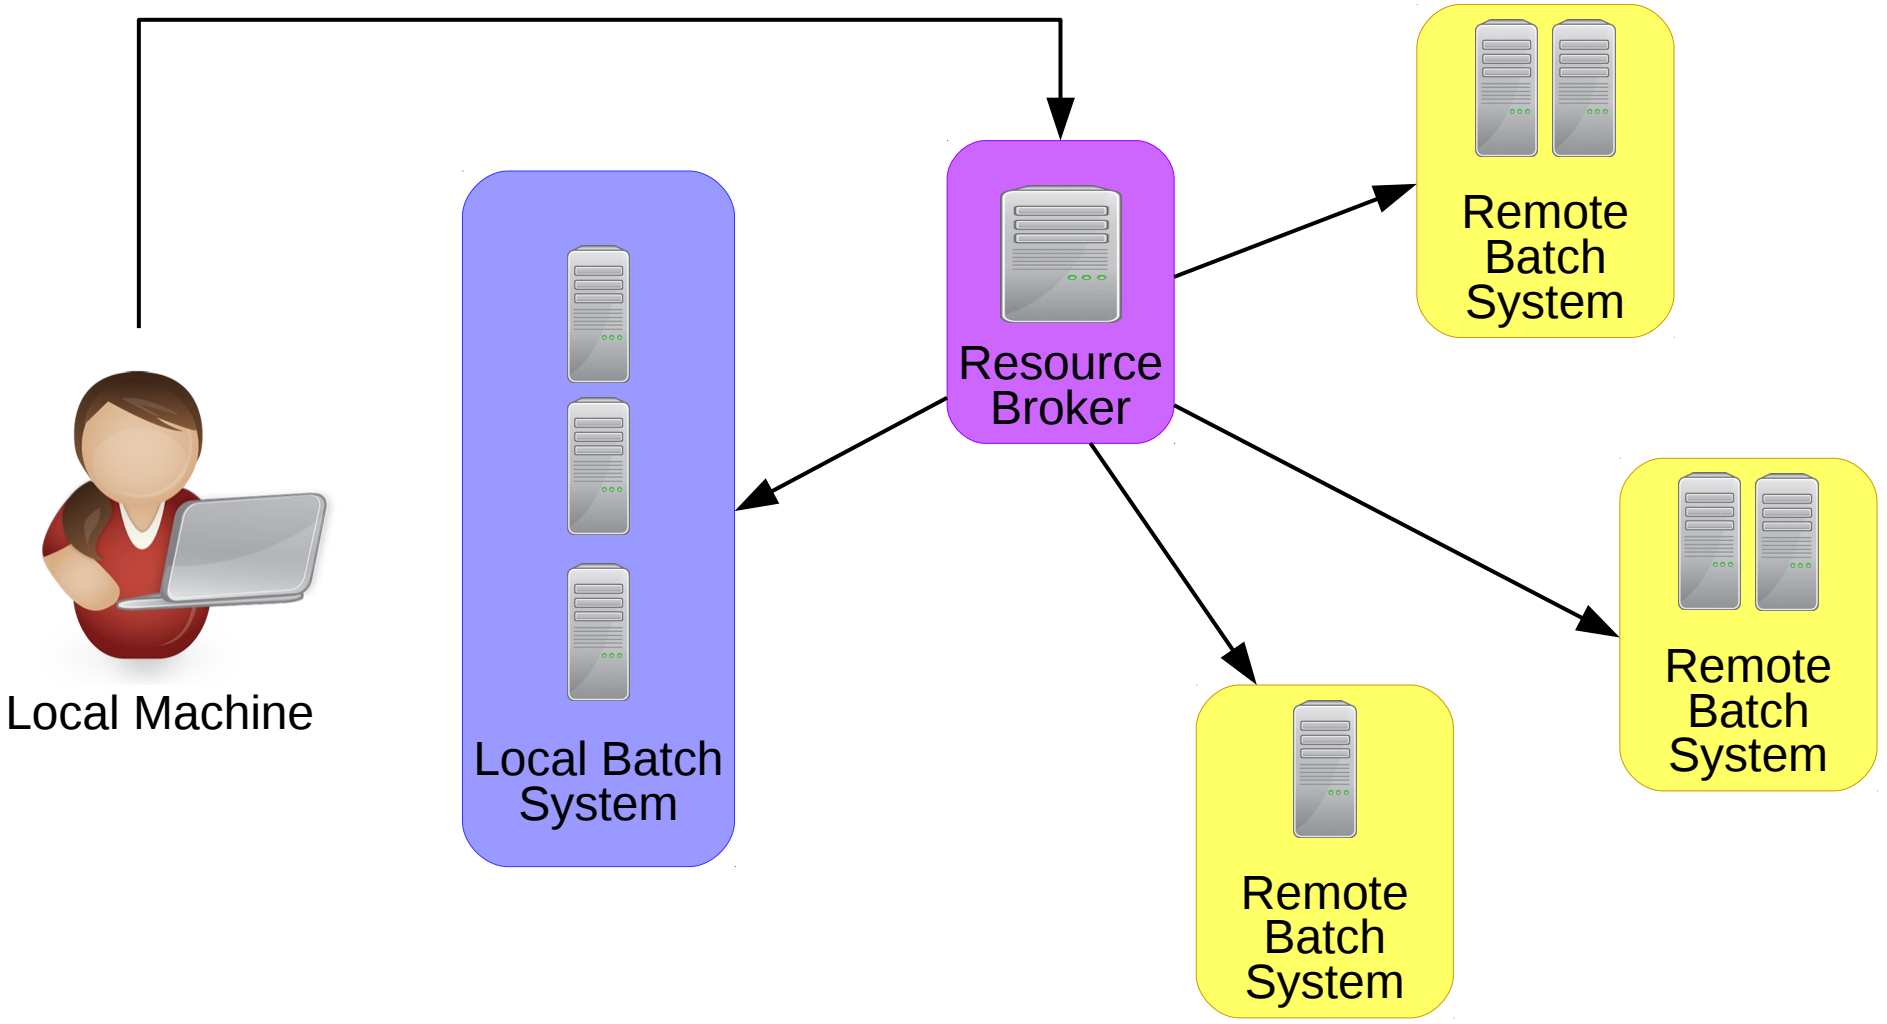

The Grid

- Originally developed as a tool to analyse data from the LHC: WLCG Webpage

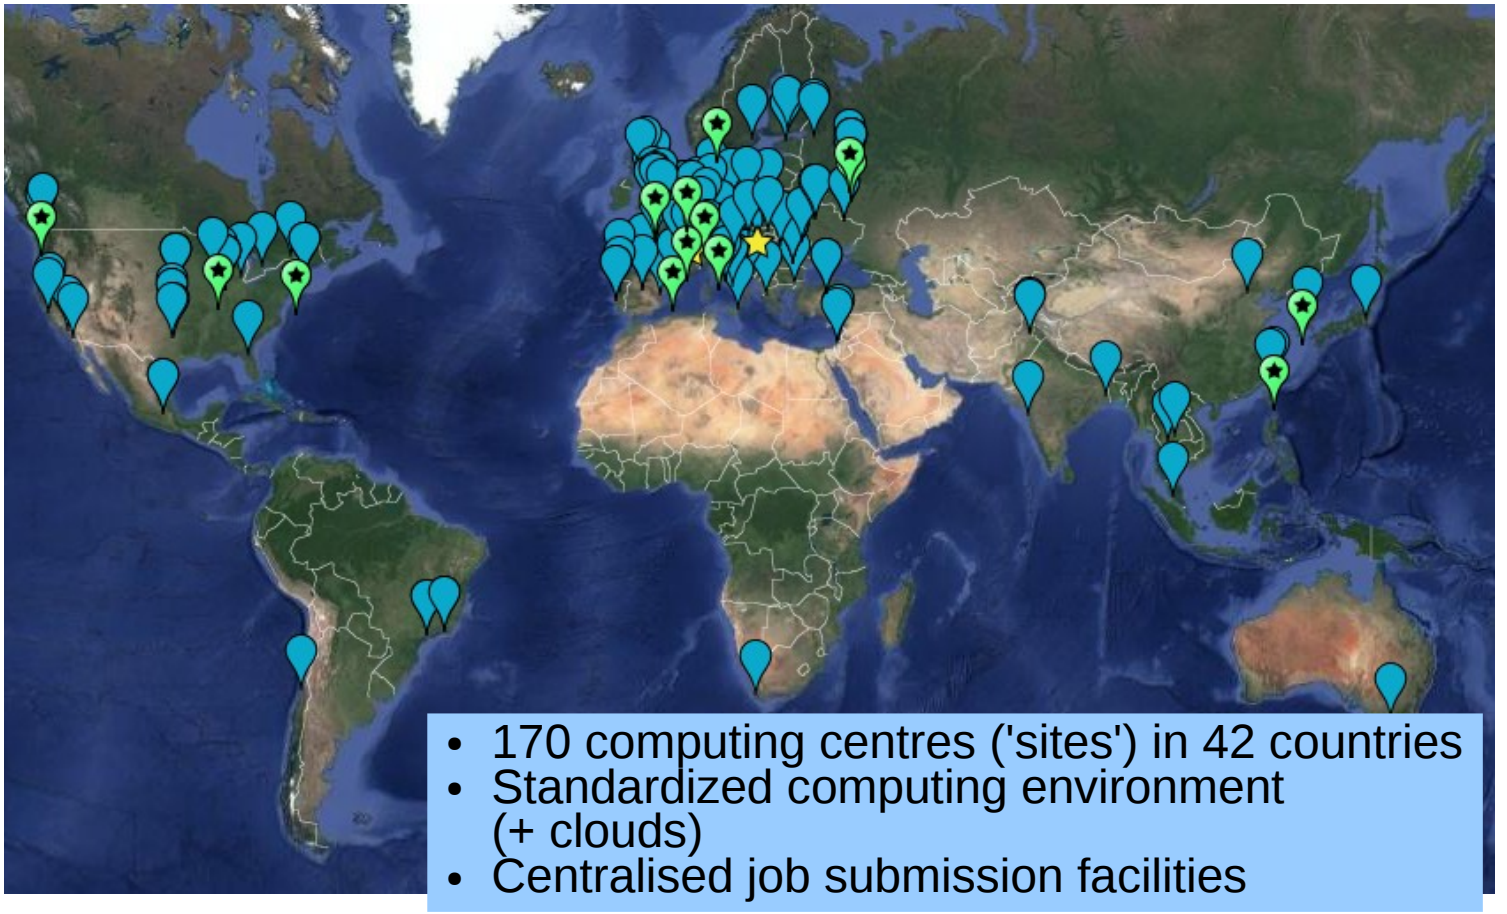

- 170 computing centres ('sites') in 42 countries
- Standardized computing environment (+ clouds)
- Centralised job submission facilities

Imperial College London

GridPP

GridPP

- A collaboration of 19 UK institutes providing Grid and Cloud based computing services to particle physics and other experiments.
- Hosts 69k job slots and 51 PB of storage.

Imperial College London

GridPP

How is this relevant to you ?

- The Grid is not just for the big LHC experiments.
- E.g. GridPP commits 10% of its resources specifically to non-LHC experiments.
- Similar initiatives exist in other countries.
- There is ample opportunity for opportunistic usage.
- (Almost) Free CPU !
- (Almost) Free Storage !
- But: No HPC, Licensing, not much GPU (yet).
- Where's the catch ?

  → Complex system, effort required.

Grid Terminology: What is a VO ?

- Wikipedia: In grid computing, a virtual organization (VO) refers to a dynamic set of individuals or institutions defined around a set of resource-sharing rules and conditions.
- A collaboration corresponding to an experiment:
  - CMS, Atlas, t2k.org, snoplus.snolab.ca, lz, vo.moedal.org, pheno
- Umbrella VOs:
  - For individual researchers whose projects are too small individually to justify their own VO:
    - WeNMR: Computational biology, gridpp: for researchers at UK universities, vo.france-grilles.fr: for researchers at French universities

What is a 'small' VO ?

- Small VOs/experiments: < 100 users
- These VOs have no dedicated computing support
- Users are typically familiar with the concept of a batch system
- Most (but not all) users are familiar with the concept of the grid
- GridPP has so far helped nine small VOs onto the Grid
- Three of these VOs were new to distributed computing

Imperial College London

GridPP

How do I access The Grid ?

- You need a certificate to prove who you are:

```
lx01:>  openssl x509 -in ~/.globus/usercert.pem -noout -text
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 46395 (0xb53b)
  Signature Algorithm: sha256WithRSAEncryption
    Issuer: C=UK, O=eScienceCA, OU=Authority, CN=UK e-Science CA 2B
    Validity
      Not Before: Nov  3 17:02:30 2016 GMT
      Not After : Dec  3 17:02:30 2017 GMT
    Subject: C=UK, O=eScience, OU=Imperial, L=Physics, CN=daniela bauer
    Subject Public Key Info:
      Public Key Algorithm: rsaEncryption
        Public-Key: (2048 bit)
        Modulus:
          00:b0:06:60:a8:3e:de:5a:0b:7e:1b:22:24:8b:cc:
```

Imperial College London

GridPP

Certificates

- Usually issued by a national certificate authority
  - e.g. France, Germany, UK
  - CERN

- With your certificate you need to join a VO (or make your own VO):
  - e.g. Joining the Moedal VO (needs certificate loaded in your browser):
    https://lcg-voms2.cern.ch:8443/voms/vo.moedal.org/register/start.action

Joining a VO

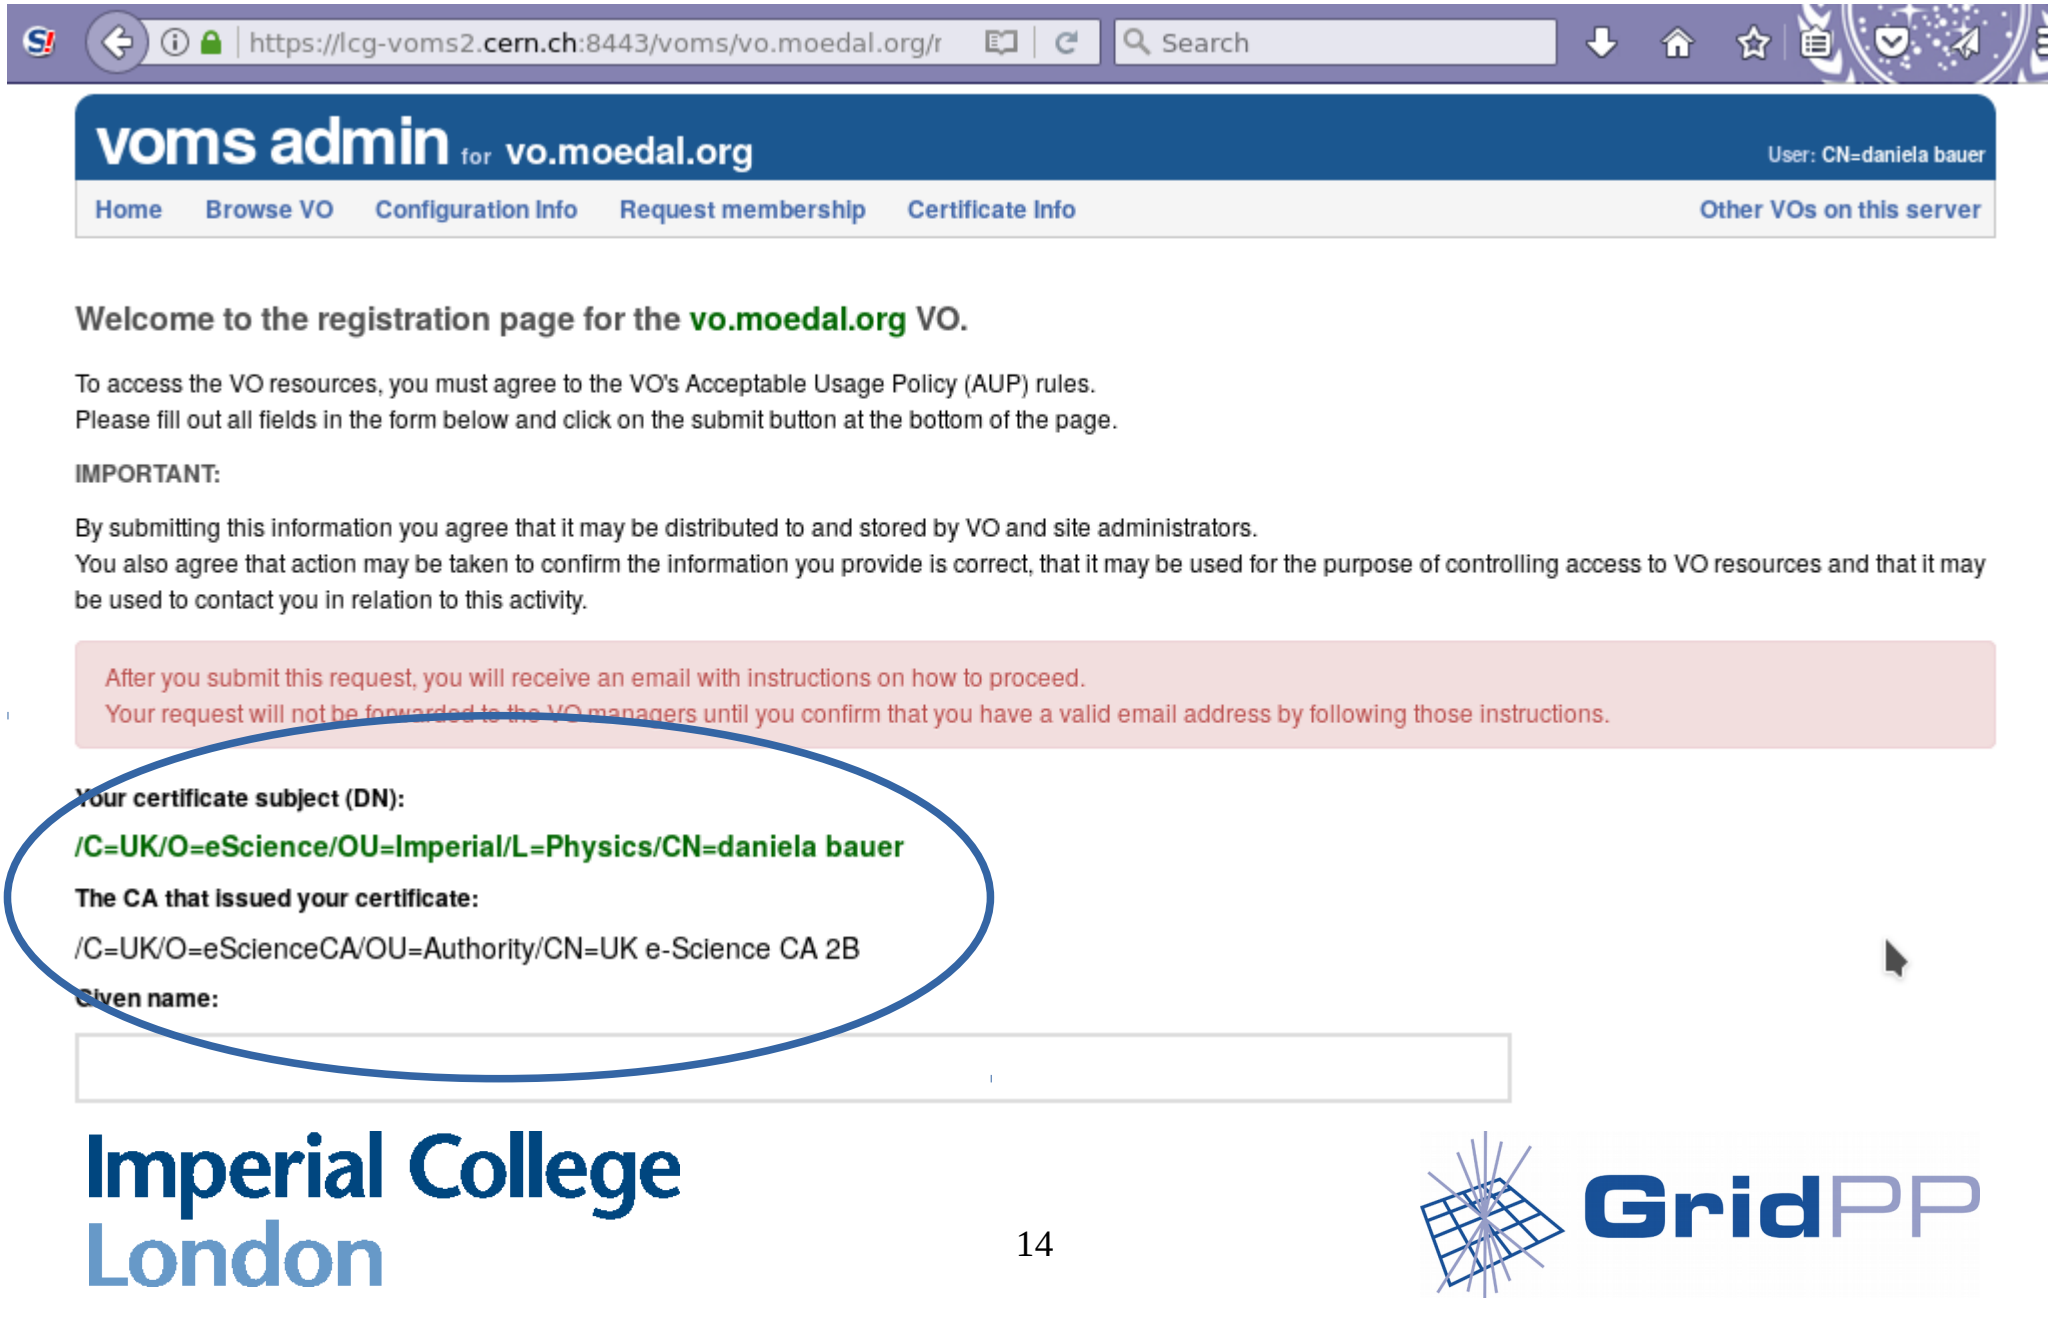
https://lcg-voms2.cern.ch:8443/voms/vo.moedal.org/r
Search
voms admin for vo.moedal.org
User: CN=daniela bauer
Home
Browse VO
Configuration Info
Request membership
Certificate Info
Other VOs on this server
Welcome to the registration page for the vo.moedal.org VO.
To access the VO resources, you must agree to the VO's Acceptable Usage Policy (AUP) rules.
Please fill out all fields in the form below and click on the submit button at the bottom of the page.
IMPORTANT:
By submitting this information you agree that it may be distributed to and stored by VO and site administrators.
You also agree that action may be taken to confirm the information you provide is correct, that it may be used for the purpose of controlling access to VO resources and that it may be used to contact you in relation to this activity.
After you submit this request, you will receive an email with instructions on how to proceed.
Your request will not be forwarded to the VO managers until you confirm that you have a valid email address by following those instructions.
Your certificate subject (DN):
/C=UK/O=eScience/OU=Imperial/L=Physics/CN=daniela bauer
The CA that issued your certificate:
/C=UK/O=eScienceCA/OU=Authority/CN=UK e-Science CA 2B
Given name:

Where to send your jobs: DIRAC

- The DIRAC framework for distributed computing was originally developed for LHCb.
- GridPP uses the DIRAC middleware to make their resources available to user communities.
- This includes Moedal :-)
- It provides a pilot based job submission framework and a data management system.
- There are currently at least 11 DIRAC instances in use worldwide.
- Homepage: diracgrid.org

DIRAC

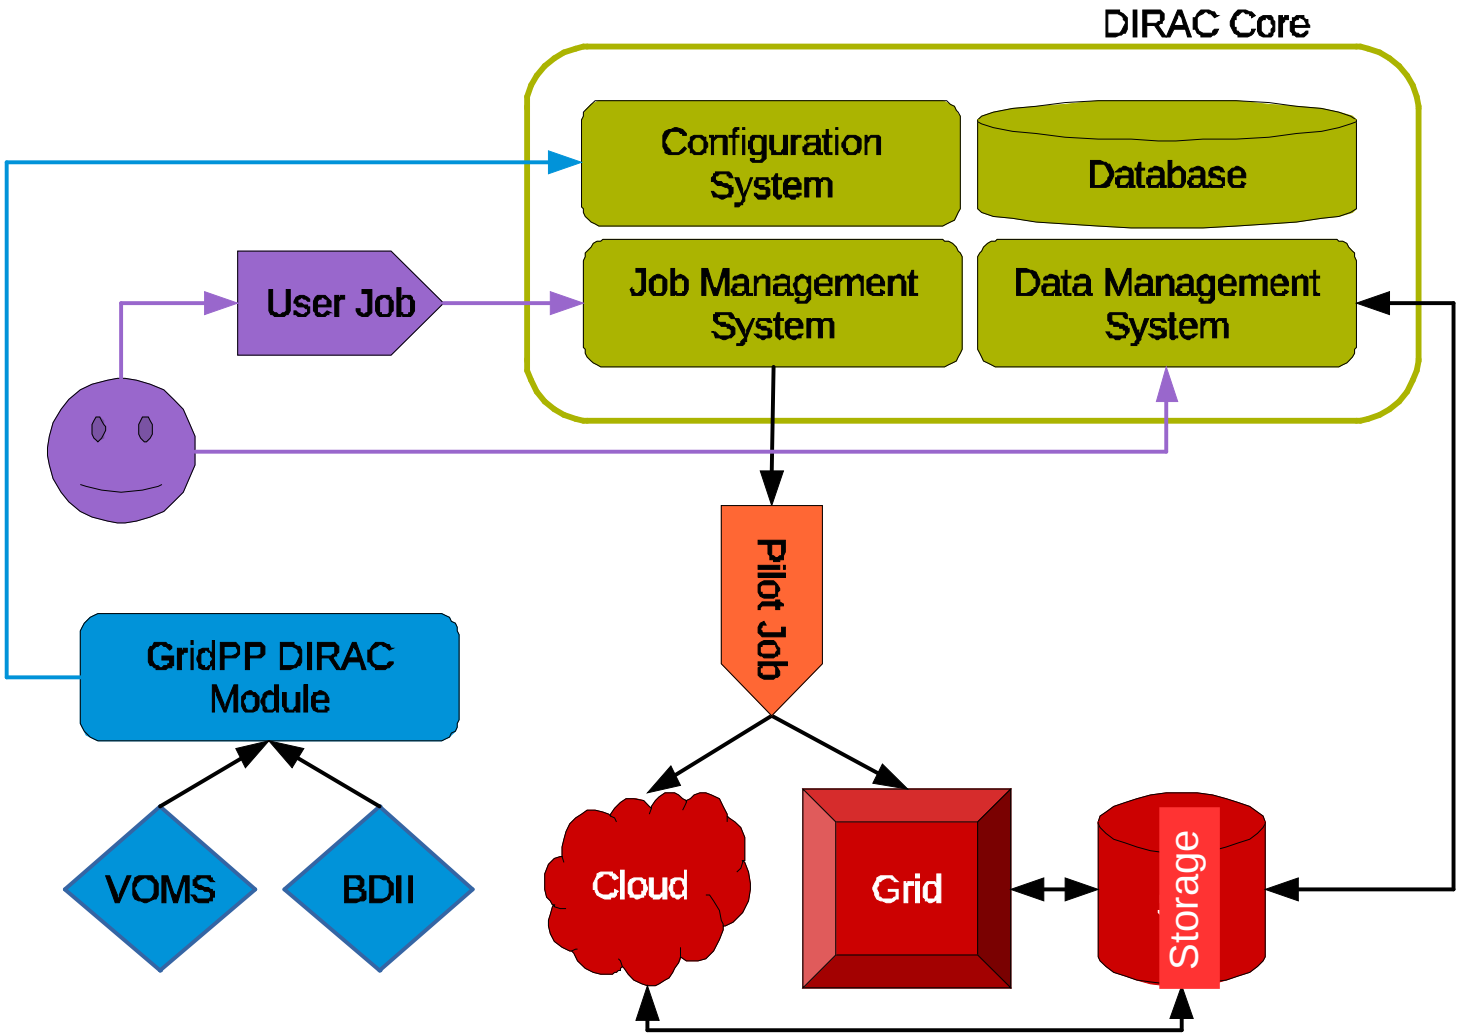
DIRAC Core
Configuration System
Database
User Job
Job Management System
Data Management System
Pilot Job
GridPP DIRAC Module
VOMS
BDII
Cloud
Grid
Storage

Accessing a DIRAC instance

- There are multiple ways to access a DIRAC instance:
  - via scripts (“DIRAC UI”) from any SL6 system.
  - via a python API.
  - using Ganga: https://ganga.web.cern.ch/ganga/
- (Usually) accessible to all members of a supported VOs
- If you are a member of the moedal VO and you would like to try out the system: https://www.gridpp.ac.uk/wiki/Quick_Guide_to_Dirac

Where to get help ?

- (moedal) access via DIRAC: gridpp-dirac-users@imperial.ac.uk
- your friendly sysadmin
- umbrella VO helplists
- GGUS

Job dashboard

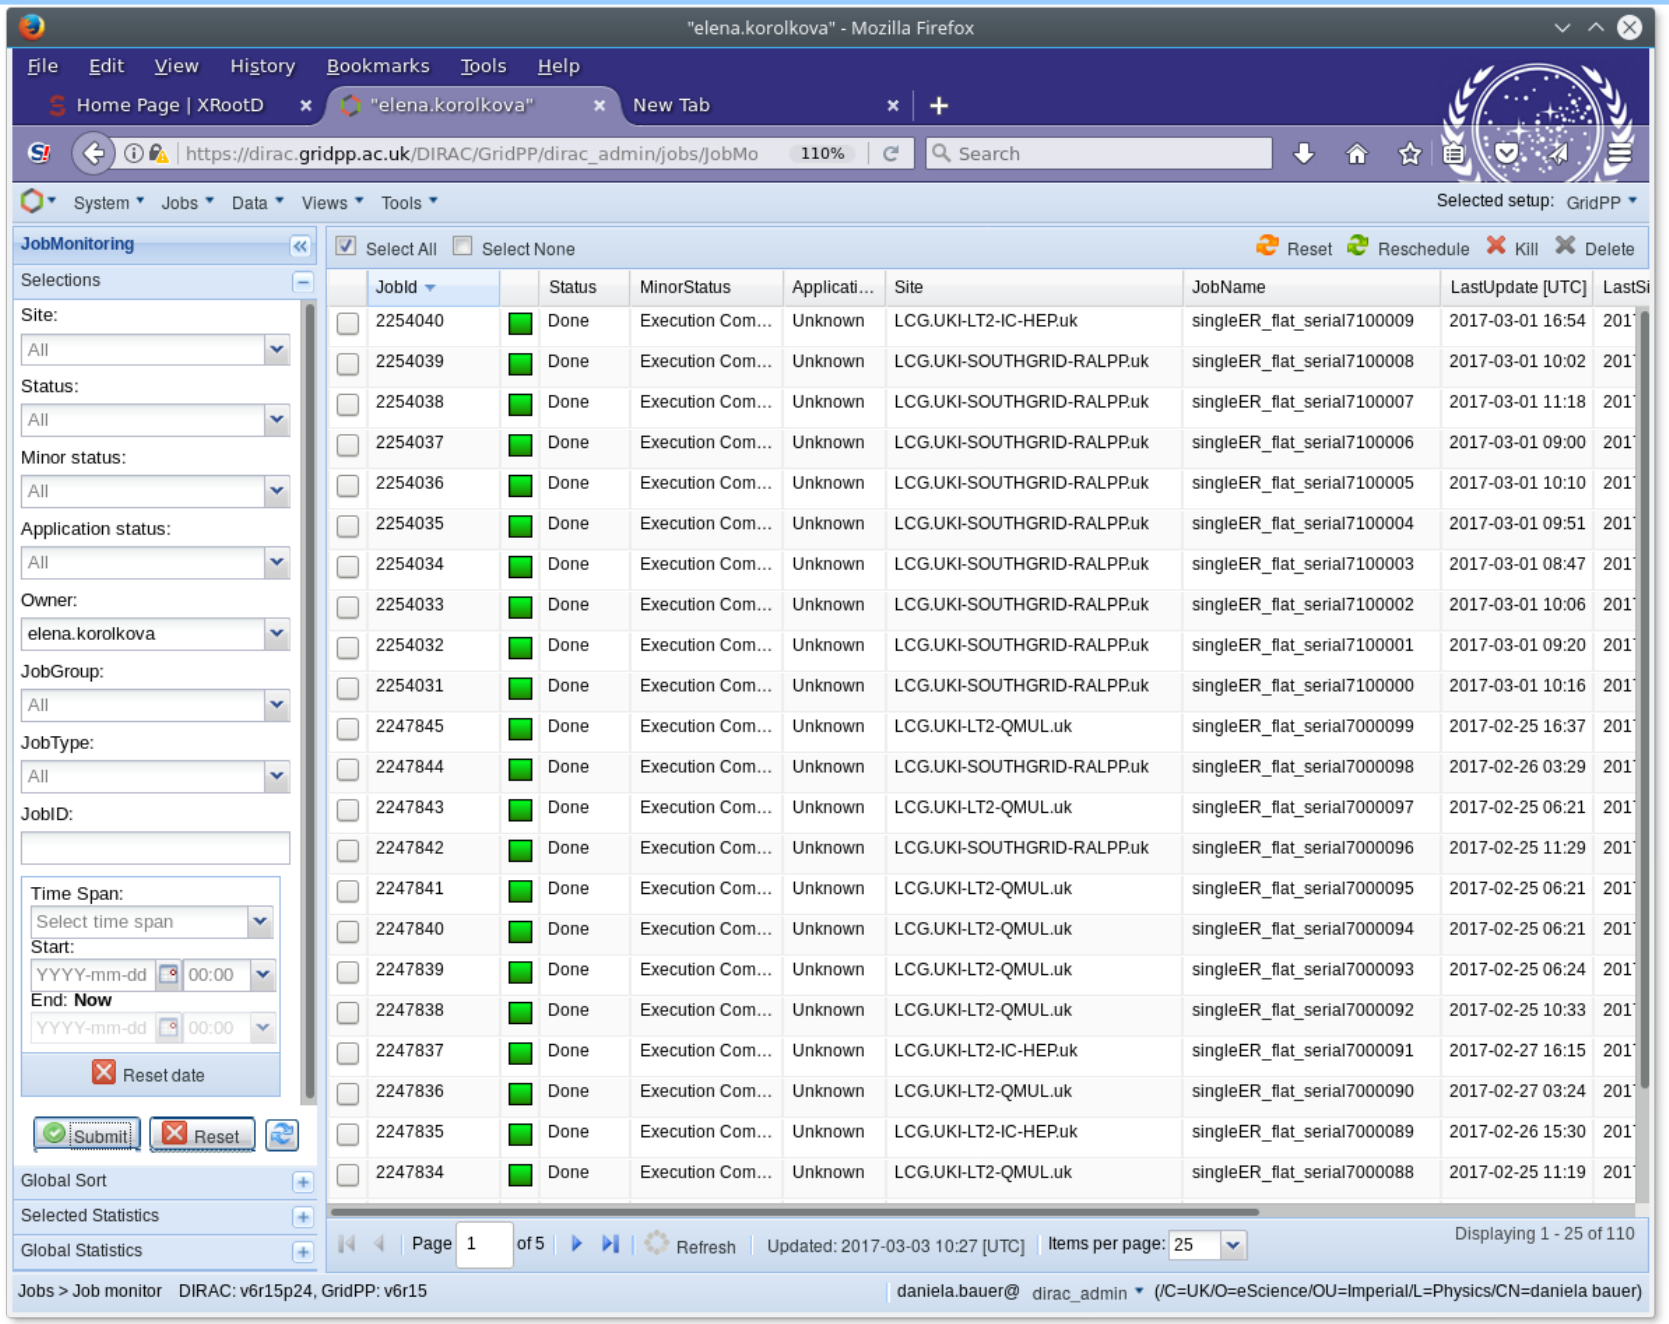

Software

- Software is usually distributed via cvmfs:
  - uploaded to central cvmfs server ('stratum 0')
  - cvmfs servers are typically hosted by WLCG T0/T1
  - software is downloaded to the sites as needed
  - caching at sites keeps popular software available for a certain amount of time
  - users access software in /cvmfs/moedal.cern.ch
- Software must be transferable
  - Anything not available in a standard SL6/SL7 environment must be packaged with your software
  - We can help, but we can't fix your code.

Data

- Small amounts of storage (~ 10 TB) are usually provided through an affiliated university.
- Larger amounts are a question between you and your funding agency.
- Data stored on grid enabled storage ('SE') is accessible from all grid sites **and** your desktop.
- A file catalogue keeps track of where your data is stored.
- The users only sees the Logical File Name (LFN): e.g.
  - /gridpp/user/dirac01.test/dirac01.testfile.txt

Real jobs from real users

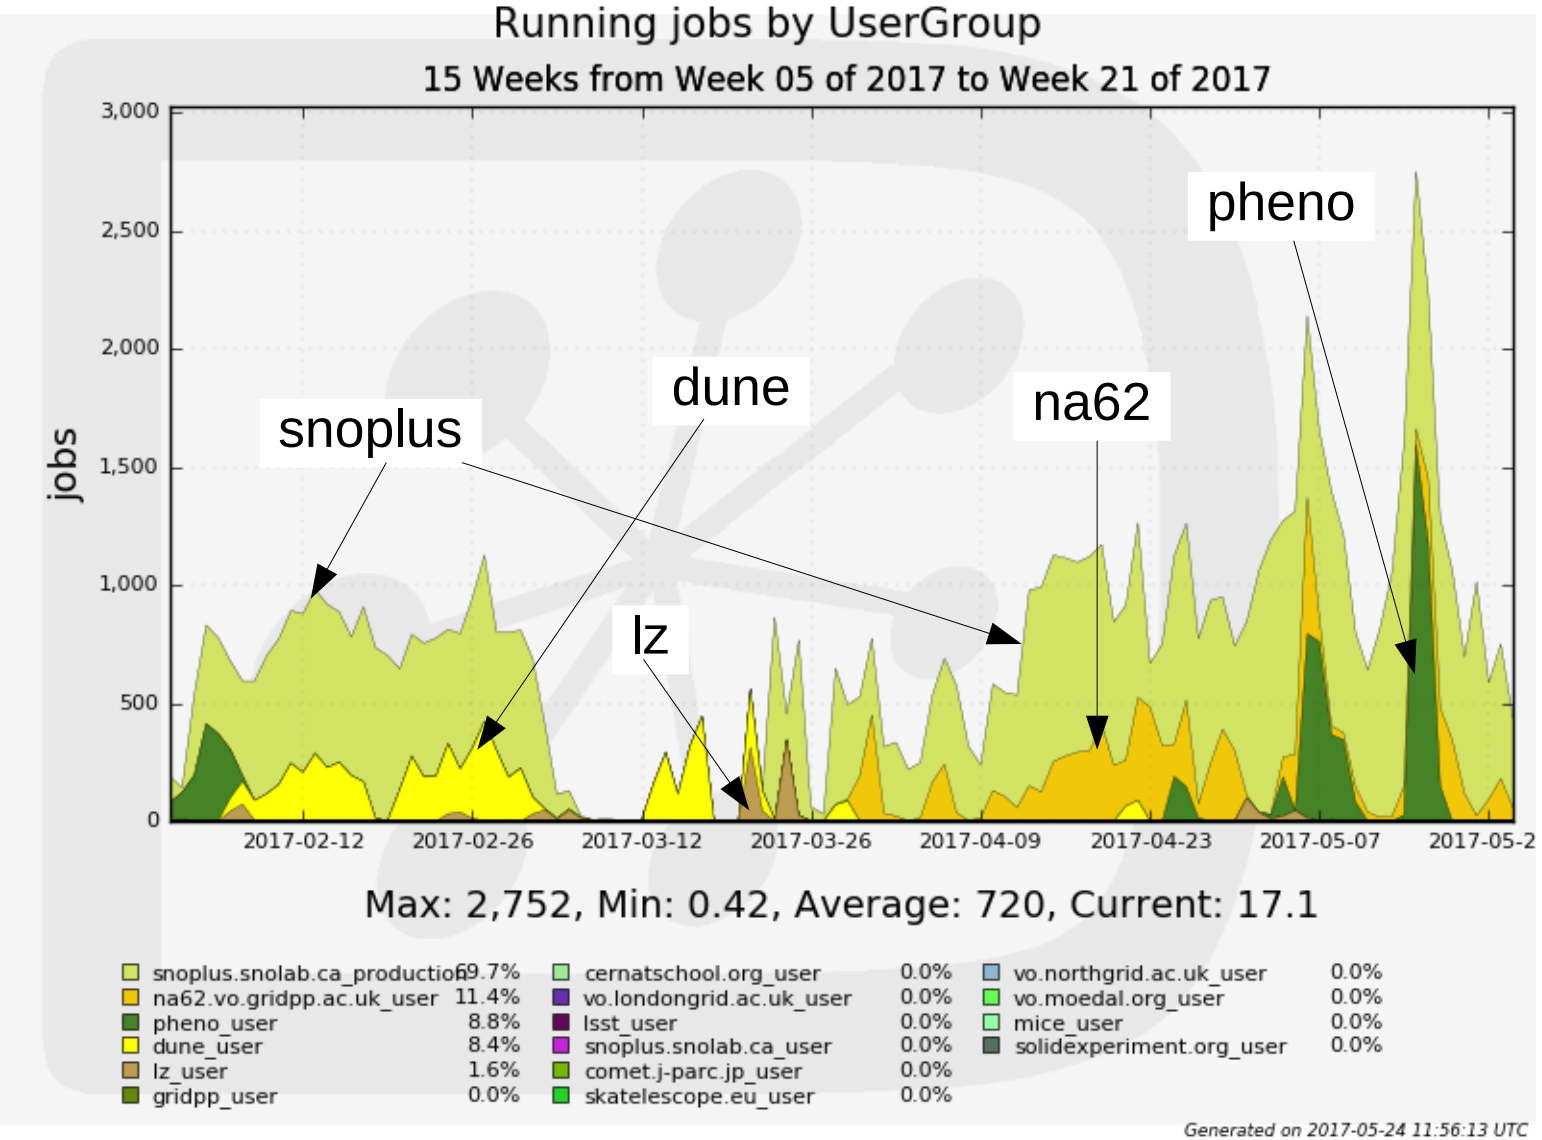

Imperial College London

GridPP

LSST

- Currently being build in Chile: http://www.lsst.org/
- Current UK contribution: Shape classification of galaxies on data taken by a predecessor (Dark Energy Survey)
- 100 million galaxies, data divided into 30000 files.
- Experiment specific software.
- Classification of a galaxy is an independent measurement, uses 10-20 s of CPU
- All data for a specific galaxy is contained in the same file
- **This looks a lot like particle physics.**

- No previous experience in grid computing.
- Settled on Ganga and DIRAC file catalogue for job submission and data access.
- 40 days FTE work required to setup and successfully complete large scale production workflow.

LuxZeplin

- Dark matter experiment based in South Dakota:
  - GEANT4 based detector simulations
  - Imperial College acts as their UK Data Centre
    - But data is processed exclusively on the grid
  - Started with experiment specific python scripts using the DIRAC API (physicists)
  - Now: Custom job submission system (software engineer)
  - Recently completed Mock Data Challenge:
    - 133 TB of MC Data comprising 732288 files

The power of opportunistic computing

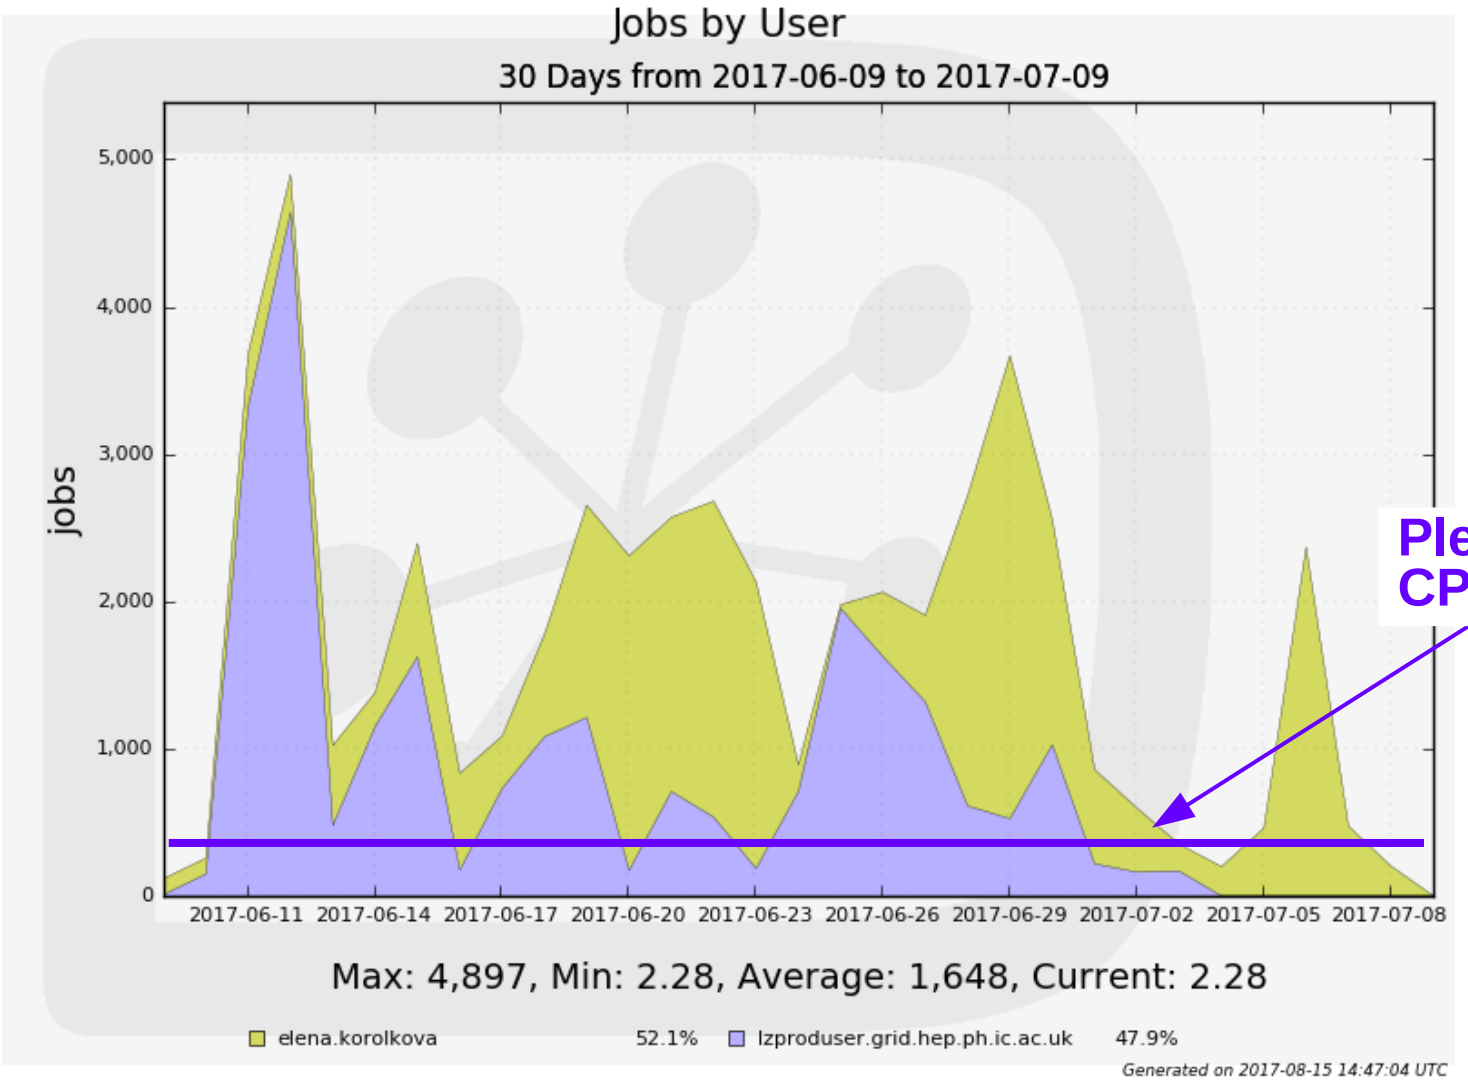

Imperial College London

GridPP

pheno/gridpp

- The **pheno VO** is used by the phenomenology group based at Durham/UK
- It uses Ganga/DIRAC to generate home grown MC
  - The NNLO QCD corrections to Z boson (...)

- The **gridpp VO** hosts very different projects:
  - GHOST: Geant Human Oncology Simulation Tool

    Using DIRAC directly, Geant4 based simluations
  - Systematic infrastructure testing (similar to the UK nagios tests, including network tests)

Imperial College London

GridPP

Summary

- Distributed computing has been successfully used by small collaborations in a production setting.
- If you have a project that needs some additional computing resources to succeed, distributed computing can help.
- If you want to outsource all your data processing needs, distributed computing can do that too (funding permitting).

Useful Links/Email

vo.france-grilles.fr:

- http://www.france-grilles.fr/presentation-en/users/
- http://www.france-grilles.fr/faq-en/

GridPP (and the 'gridpp' VO):

- https://www.gridpp.ac.uk/
- technical questions: daniela.bauer@imperial.ac.uk, jeremy.coles@cern.ch
- formal request of collaboration: David Colling (d.colling@imperial.ac.uk) and David Britton (david.britton@glasgow.ac.uk)

Moedal:

- https://lcg-voms2.cern.ch:8443/voms/vo.moedal.org/register/start.action