

Volunteer Computing at CERN

BOINC workshop Sep 2014, Budapest
Tomi Asp & Pete Jones, on behalf the LHC@Home team

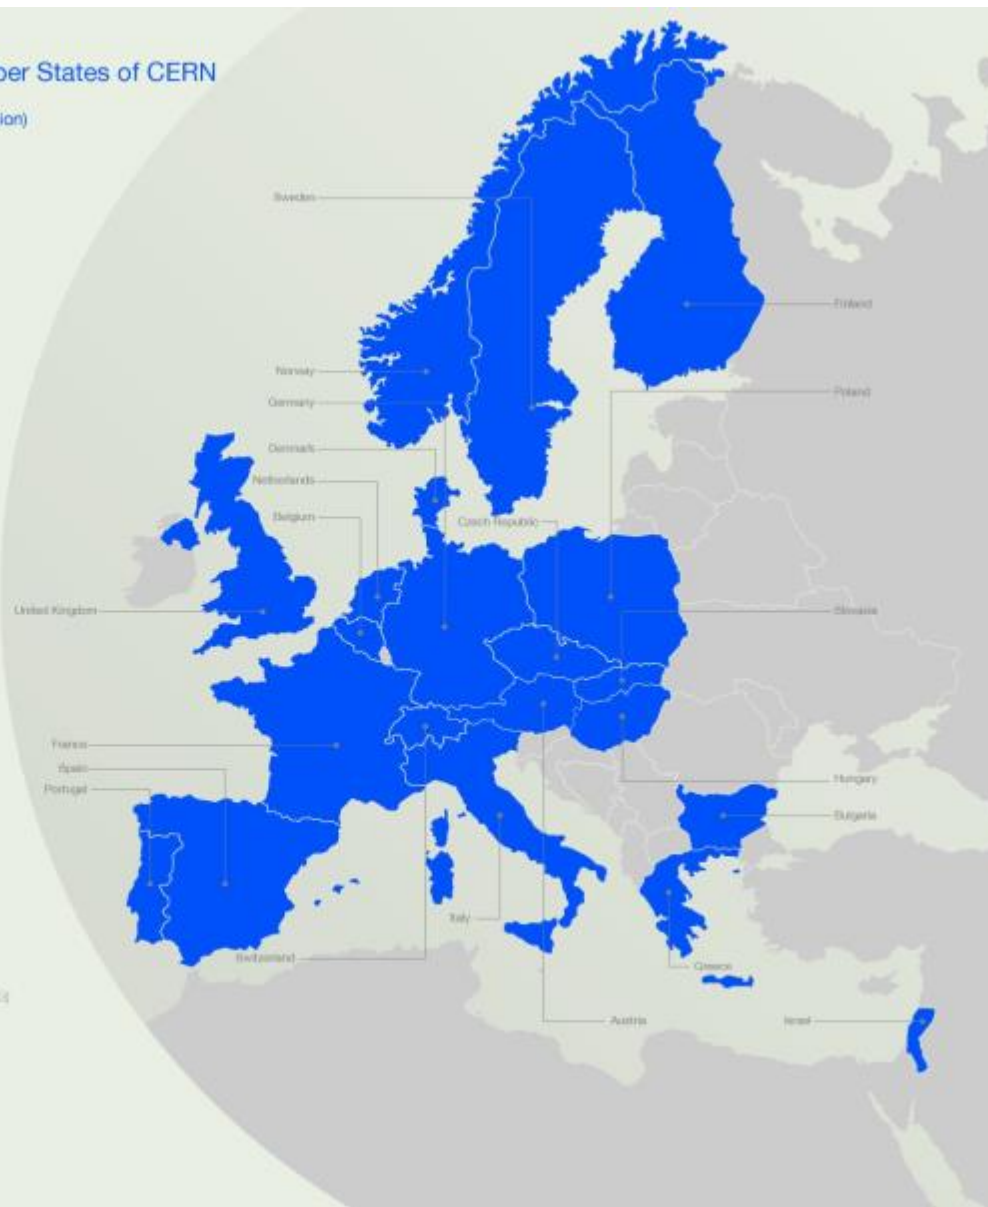
Agenda

- Overview
- Status of the LHC@Home projects
- Additional BOINC projects
- Service consolidation
- Planning for the future
- Questions

The twenty one Member States of CERN

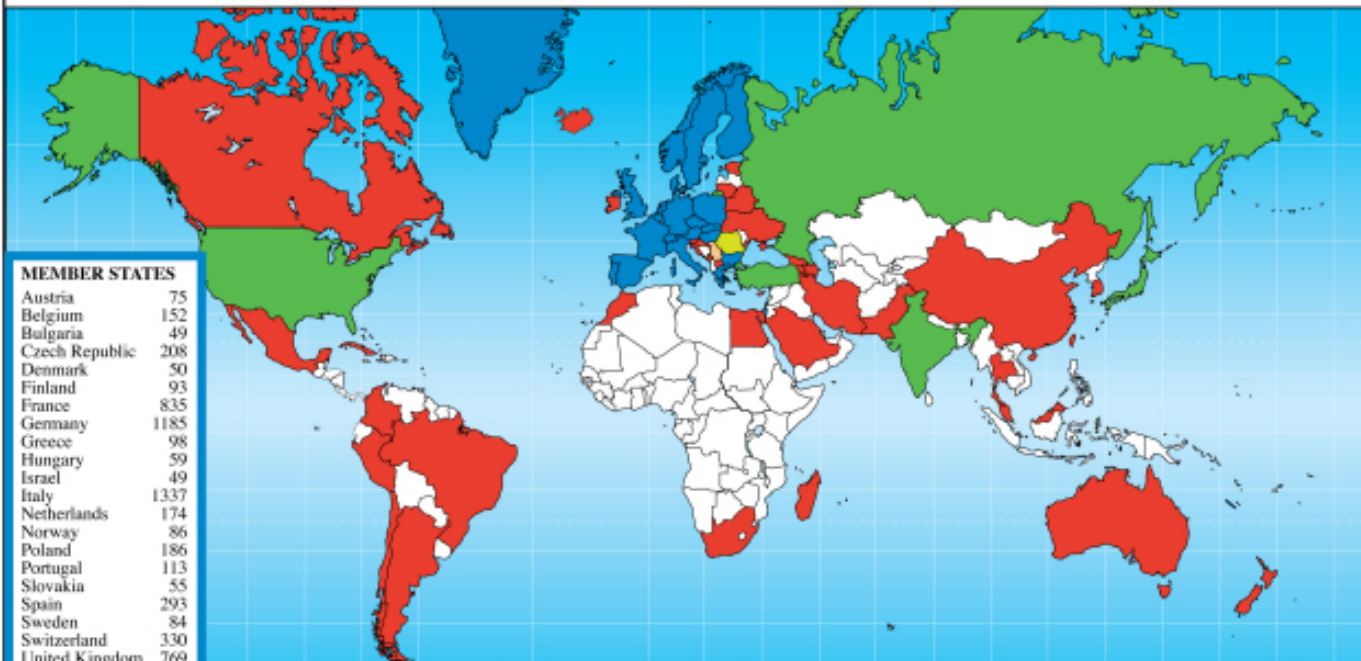
Member States (Dates of accession)

-  Austria (1958)
-  Belgium (1953)
-  Bulgaria (2000)
-  Czech Republic (2009)
-  Denmark (1962)
-  Finland (1992)
-  France (1953)
-  Germany (1953)
-  Greece (1962)
-  Hungary (2002)
-  Israel (2014)
-  Italy (1953)
-  Netherlands (1953)
-  Norway (1962)
-  Poland (1987)
-  Portugal (1986)
-  Slovakia (1993)
-  Spain (1/1981, 12/1985, 1/1986)
-  Sweden (1969)
-  Switzerland (1959)
-  United Kingdom (1963)



The European Organization for Nuclear Research

Distribution of All CERN Users by Location of Institute on 14 January 2014



MEMBER STATES

Austria	75
Belgium	152
Bulgaria	49
Czech Republic	208
Denmark	50
Finland	93
France	835
Germany	1185
Greece	98
Hungary	59
Israel	49
Italy	1337
Netherlands	174
Norway	86
Poland	186
Portugal	113
Slovakia	55
Spain	293
Sweden	84
Switzerland	330
United Kingdom	769

6280

OBSERVERS

India	153
Japan	217
Russia	890
Turkey	110
USA	1724

3094

CANDIDATE FOR ACCESSION

Romania	86
---------	----

ASSOCIATE MEMBER IN THE PRE-STAGE TO MEMBERSHIP

Serbia	30
--------	----


OTHERS

Argentina	13	China	122	Iran	20	Pakistan	18
Armenia	16	China (Taipei)	71	Ireland	5	Peru	2
Australia	39	Colombia	10	Korea	105	Saudi Arabia	3
Azerbaijan	2	Croatia	23	Lithuania	13	Slovenia	25
Belarus	24	Cuba	3	Madagascar	3	South Africa	32
Brazil	116	Cyprus	13	Malaysia	8	Thailand	8
Canada	147	Egypt	18	Mexico	46	T.F.Y.R.O.M.	1
Chile	8	Estonia	17	Montenegro	1	Ukraine	24
		Georgia	11	Morocco	6		
		Iceland	4	New Zealand	5		

982

CERN is the home of the LHC

- The CERN laboratory sits astride the Franco-Swiss border near Geneva.
- The Large Hadron Collider (LHC) is the world's largest and most powerful particle collider
- LHC data analysis requires a computing power equivalent to ~100,000 of today's fastest PC processors



~30 Petabytes of data is
annually generated by the
LHC

- CERN can only provide ~20% of the computing capacity
- The rest is running on the World LHC Computing Grid
- Volunteer computing could provide additional capacity for tasks like simulation

BOINC at CERN - LHC@HOME

- **Sixtrack**

- Since 29-Sep 2004 to celebrate CERN's 50th birthday
- Simulates particles accelerating through the LHC to find their orbit stability.
- In one work-unit, 60 particles are simulated
- 300,000 clients have signed up.
- Peak computing power up to 40 Tflops
- Classic BOINC architecture – Mac, Linux and Windows Clients
- BOINC URL: <http://lhathomeclassic.cern.ch/sixtrack/>

- **Test4Theory**

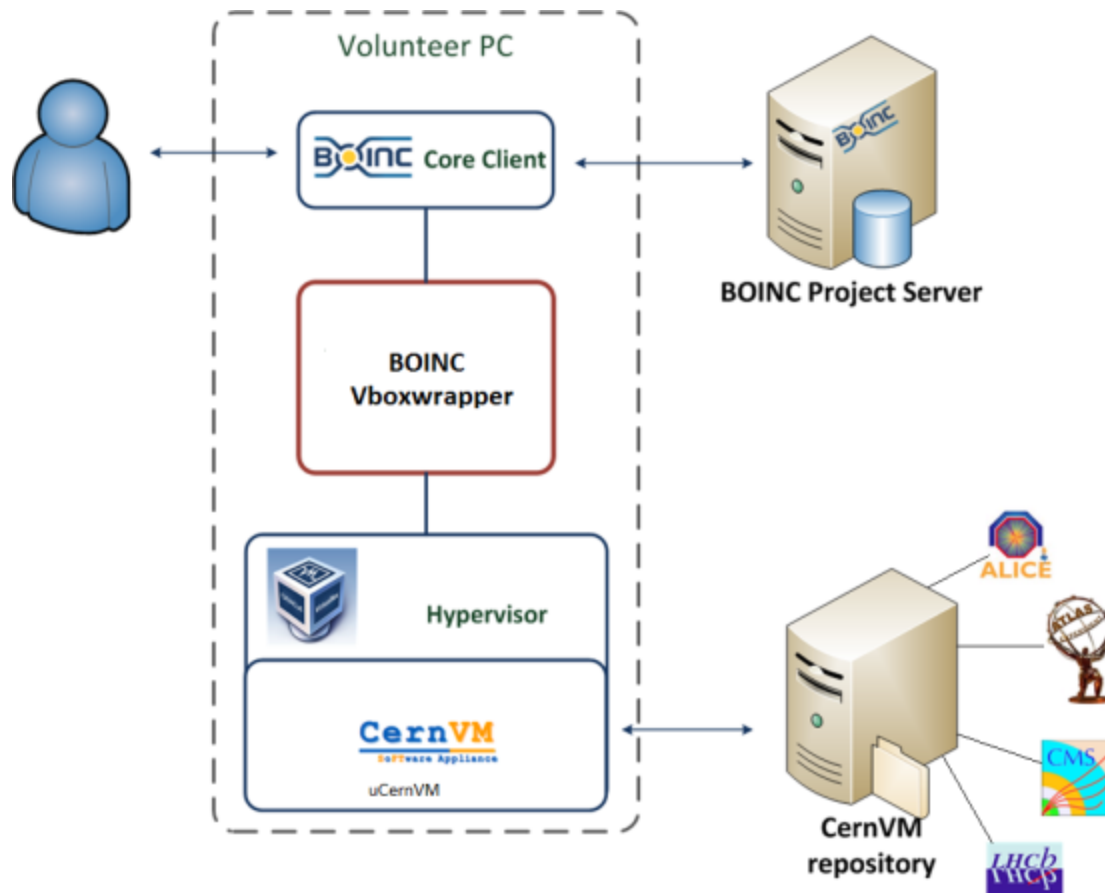
- Since 2011 and does research about the elusive Higgs particle
- Monte Carlo simulations of physics events at the LHC.
- 30,000 volunteers
- Computing power up to 2.5 Tflops
- Modified architecture to traditional BOINC - uses the CERN developed CernVM and Vboxwrapper application
- Changed name in 2014 to **Virtual LHC@home** - with the view to host more apps than Theory simulations
- BOINC URL: <http://lhathome2.cern.ch/vLHCathome/>

- **Both in present form since 2011**

LHC@Home – other projects in progress

- **Atlas@Home**
 - Lots of progress during 2014.
 - Running Atlas simulation
 - Over 5000 registered users
 - See talk later given by Wenjing Wu: Atlas@home
 - BOINC URL <http://atlasathome.cern.ch/>
- **LHCb: Beauty**
 - Since 2012 and still in test phase.
 - The LHCb experiment collaboration's cloud and grid computing management solution
 - Vboxwrapper application, with CernVM image
 - Planned to be an app on vLHC@home
- **CMS**
 - Started work on this during summer 2014
 - In very early stages, prototype running
 - To be added as a beta-application on vLHC@home once stable
- **Adopt a Neuron**
 - Pilot project with EPFL for a Human Brain project
 - For tasks in neuroscience for a volunteer computing context similar to CERN's Test4Theory
 - Slowly progressing

BOINC and Virtualisation - 1




BOINC and Virtualisation - 2

- Vboxwrapper is a success
 - Thanks to Rom Walton for improvements and bug fixes
 - What about a forum for VM based projects and developments?
- Our aim is to keep the image and job size small
 - uCernVM is a small core image
 - An environment for developing and running LHC data analysis independent of OS
 - Pull applications from CernVMfs
 - A software distribution mechanism to deliver experiment software efficiently
 - Multiple tasks/cores:
 - BOINC supports multiple BOINC tasks, e.g. multiple vboxwrapper apps
 - What about multiple tasks per VM?

BOINC at CERN – recent developments

Service Consolidation

- Take advantage of CERN IT services features and support
 - OpenStack - open source cloud computing platform
 - Puppet managed VMs
 - Use of centralised Database service
 - BOINC server code modified by Tomi Asp (Graduate student at CERN for 1 year)
 - Centralised NFS disk space
 - Backup services
 - Drupal portal for lhathome.cern.ch

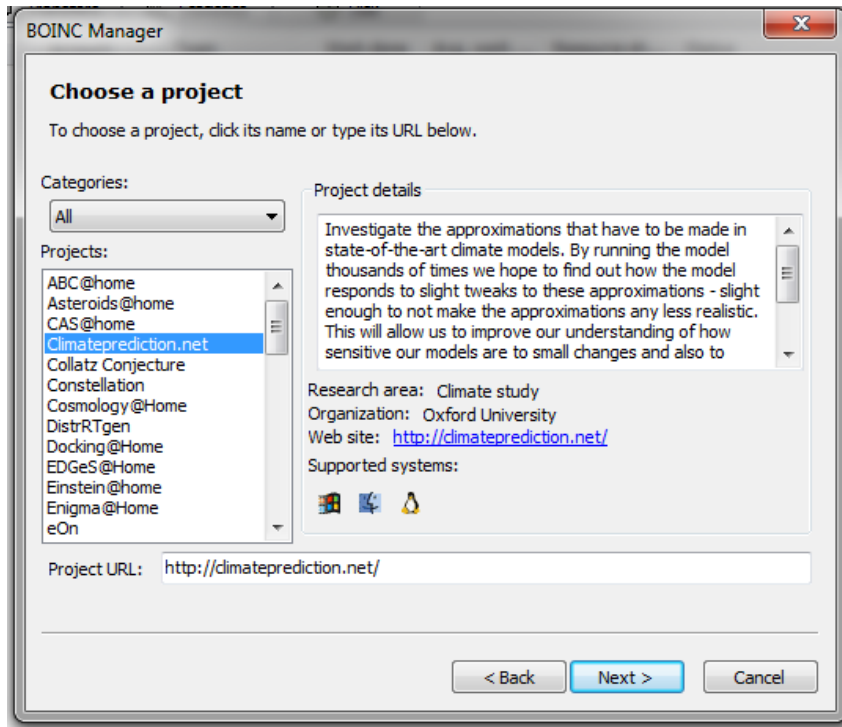


Different project approaches

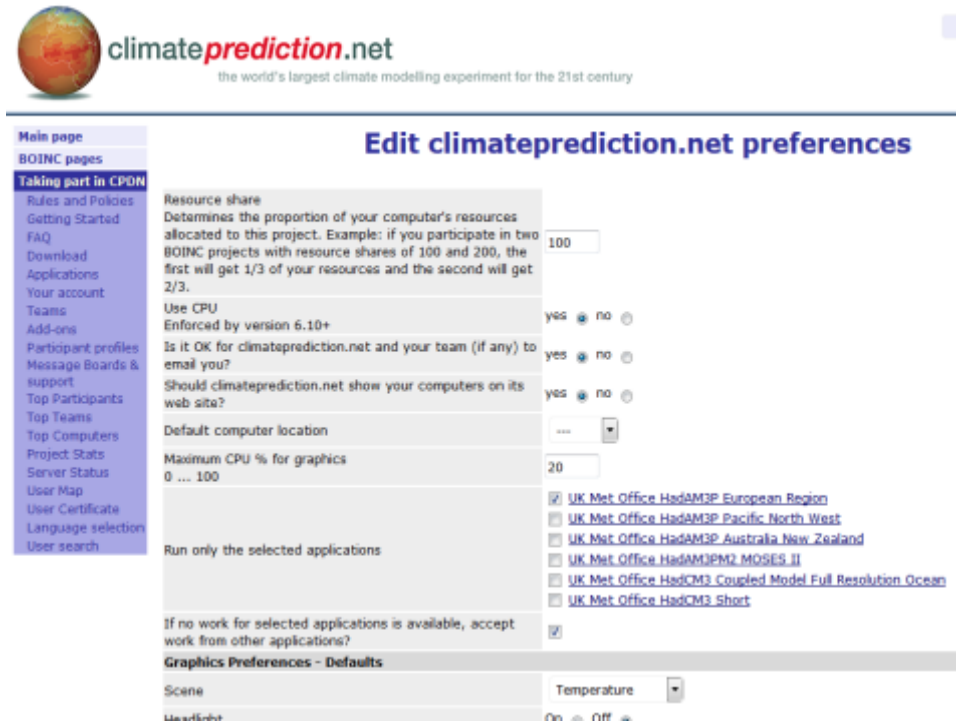
- We have to consider the needs of the multiple experiments that are starting to use BOINC
- Sixtrack uses BOINC in the classic way.
- Other projects use the Vboxwrapper application
- Common denominator for VM projects:
 - cernvm image
 - getting apps from Cernvmfs file system.
- However different job managers are used
 - (Co-Pilot, Panda+ArcCE, Crab, Dirac...)

How to organise many BOINC projects?

Individual projects



One project with many apps



We want to centralize the service and at the same time respect the needs of multiple experiments.

One project with several apps

- Pros
 - One URL for the volunteer
 - One user database
 - One set of forums
 - Probably less maintenance (server and BOINC)
- Cons
 - Merging the database (e.g. users=yes, forums=no)
 - Managing different stakeholder requirements
 - The effect of many apps on the server performance
 - Handling app credit (however this seems to be fixed)

Individual BOINC projects

- Pros
 - Experiment independence
 - Easier to manage heterogeneous projects
 - Avoid merging issues
- Cons
 - Common URL – is this possible?
 - Will mean separate user registration
 - Separate forums (possible duplications)

Summary

- Increased interest in Volunteer Computing at CERN
- We provide and a centralized BOINC service
 - Integrating existing IT services
 - Accommodate the needs of all current and future projects
- Various configuration options analysed
 - All have their risks
 - Separate projects are less risky to handle
 - But the app feature seems a sensible solution
- Our setup must be able to scale
- It must be coherent for volunteers
- We welcome feedback, comments and questions

