

BOINC

Volunteer Computing at CERN

Pre-GDB 11/11-2014

Nils Høimyr, IT/PES on behalf the BOINC service team

Outline

- Why Volunteer Computing?
- Infrastructure / Middleware – BOINC
- BOINC system architecture
- BOINC Compute Power
- Virtualisation with BOINC
- BOINC Service at CERN
- Summary
- Questions

Why Volunteer Computing?

Target	Deployment	Benefit
Volunteers	Uncoordinated, opportunistic	<ul style="list-style-type: none">• Get additional, “free” compute cycles• Engage with communities outside HEP: outreach and publicity for HEP and science
Institute desktops	Coordinated, opportunistic	<ul style="list-style-type: none">• Get additional, “free” compute cycles
Small to midsize server farms	Coordinated, pledged	<ul style="list-style-type: none">• Easier to deploy than complete Grid middleware


Infrastructure / Middleware

- Most commonly used middleware: BOINC
 - Other choices: XtremWeb (Oleg's presentation), HTCondor, ...
 - Other initiatives based on virtualisation and clouds
- LHC@home uses BOINC
 - Focus on that for the rest of this presentation

BOINC

“Berkeley Open Infrastructure for Network Computing”

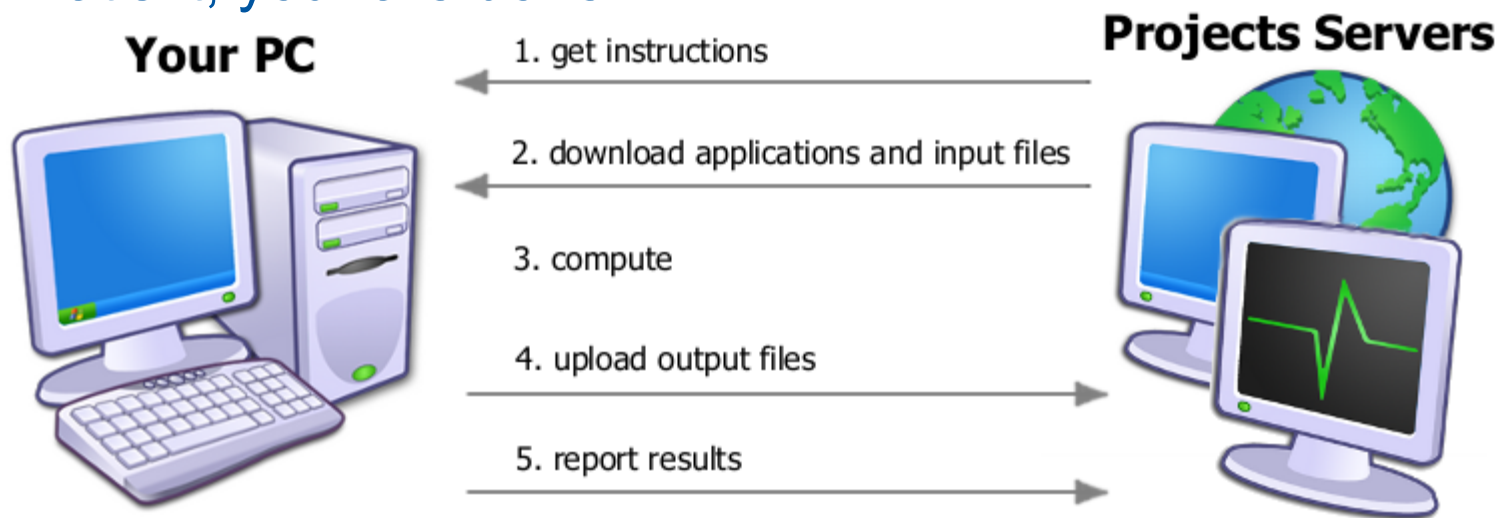
- Software platform for distributed computing using volunteered computer resources
- Client – server architecture
- Free and open source
- Used for
 - SETI@home
 - Climateprediction.net
 - Einstein@home
 - LHC@home
 - ...



<http://boinc.berkeley.edu>

Volunteer Perspective of BOINC

- Download and run BOINC software
- Choose a project
- Enter an email address and password in the BOINC Manager
 - You can also make a silent connection with a key from the BOINC client
- That's it, you are done!



BOINC client - SLC6

- 3 steps and your machine is crunching:

```
#yum install boinc-client-7.2.33-3.git1994cc8.el6.x86_64
```

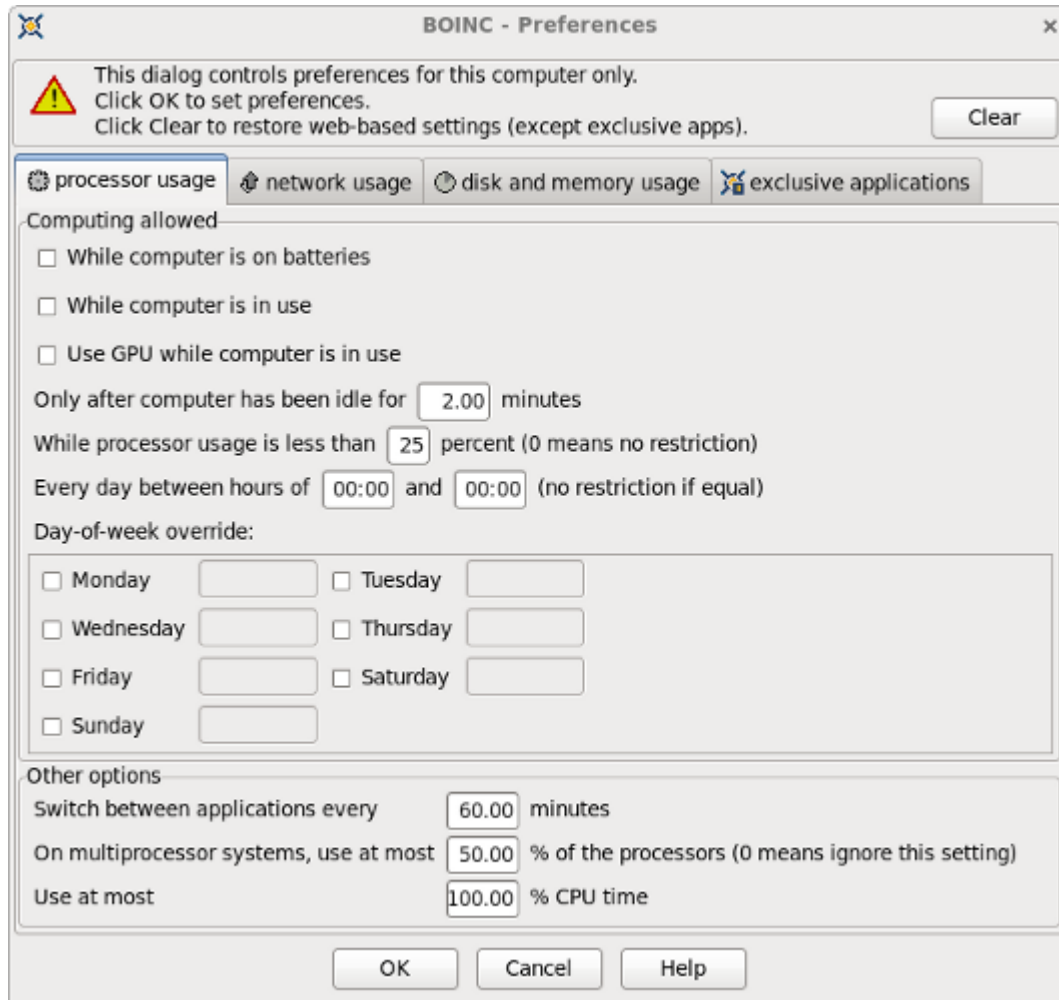
```
#yum install VirtualBox-4.3-4.3.12_93733_el6-1.x86_64.rpm
```

```
bash$ boinc_client --attach_project
```

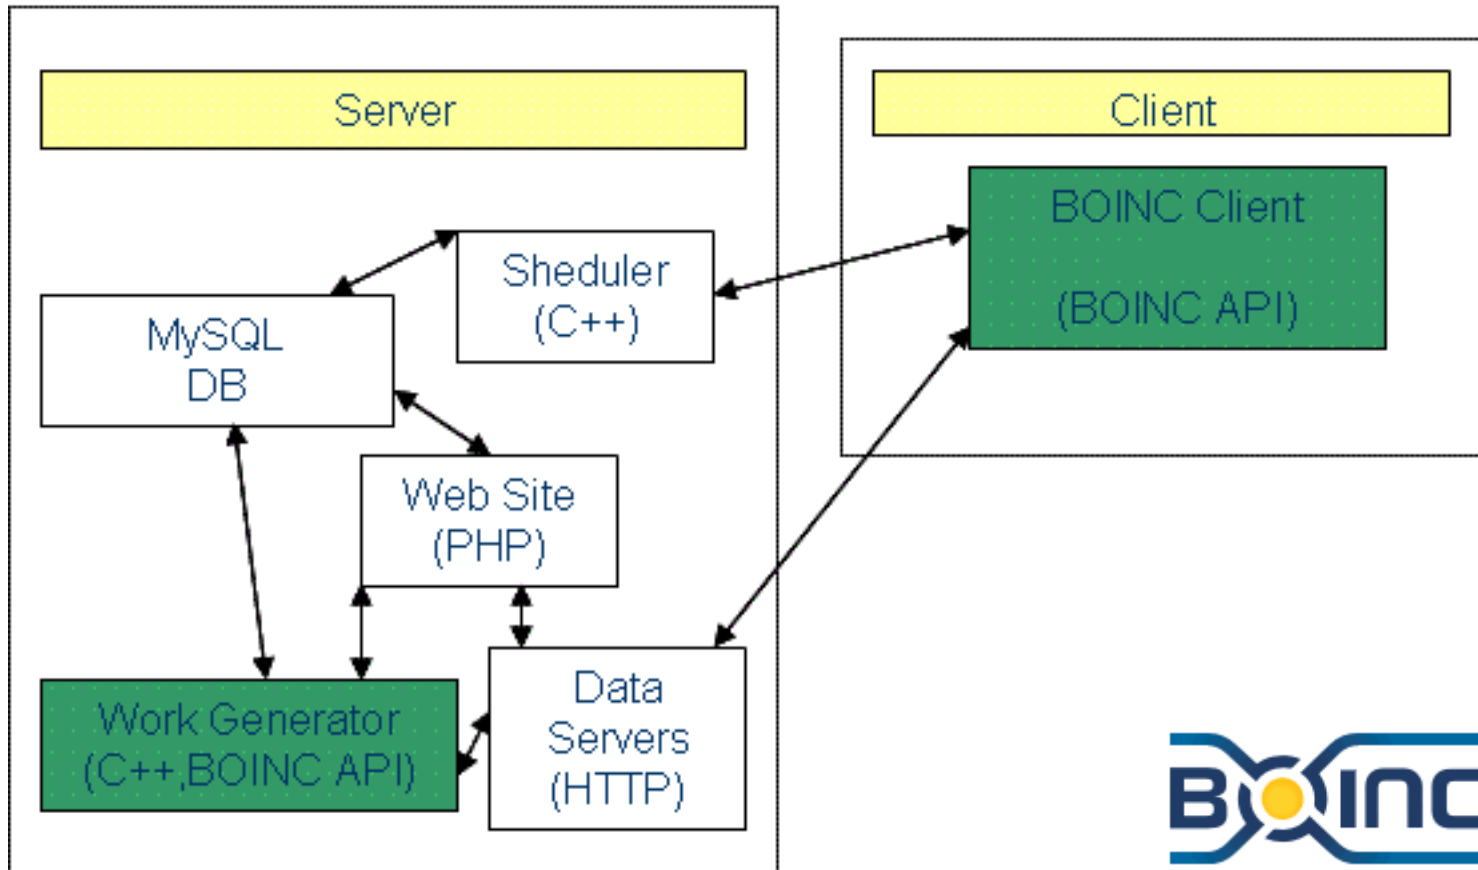
```
http://lhathome2.cern.ch/vLHCathome a25c5d7c0887ef8faf5f9e46d3ca54fc
```

- Additionally “boincmgr” can be used on desktops for interactive use (attach to projects, suspend tasks, settings etc.)

BOINC client – designed not to disturb the user



BOINC system architecture - 1



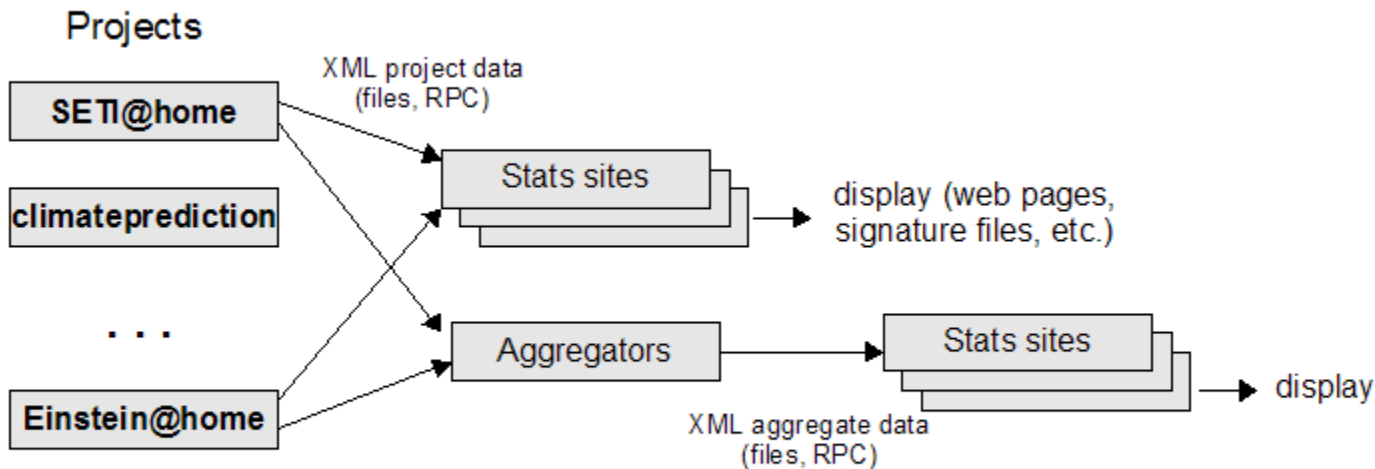
BOINC system architecture - 2

- Server:
 - MySQL db back-end
 - Web interface (PHP and CGI)
 - Includes forums for interaction with volunteers
 - Scheduler (C++)
 - Daemons: Work generator, assimilator, validator,
 - API (Python, C++, Fortran)
- Client:
 - Manager GUI
 - Core client, running as a daemon
 - XML files with project and job settings



BOINC credit and statistics

- A key concept used to motivate volunteers with BOINC is **credit** for contributed CPU
 - The credit is based on CPU benchmarks on the host
 - Credit is collected from multiple BOINC projects



BOINC Compute Power

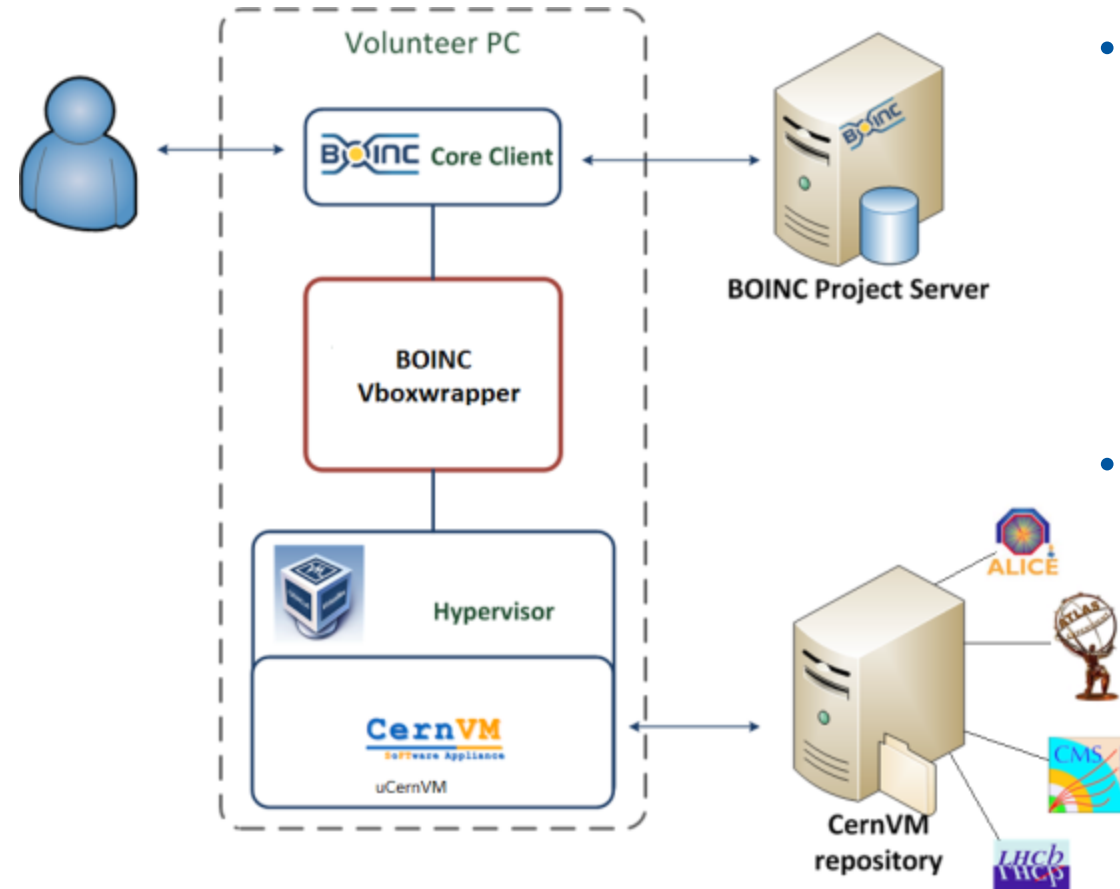
Project	Average power
SETI@home	653 TFlops
Einstein@home	637 TFlops
World Community Grid	421 TFlops
LHC@home -classic	31 TFlops
Virtual LHC@home	2.6 TFlops

According to BOINCstats.com 10.11.2014

BOINC and Virtualisation - 1

- Pioneered at CERN by Test4Theory and the CernVM team in PH/SFT 2010-2011
- Later brought into BOINC mainstream code as “Vboxwrapper”
 - Ref: <http://boinc.berkeley.edu/trac/wiki/VboxApps>
 - BOINC developers very helpful with improvements
- BOINC projects currently deploying Virtualisation:
 - RNAword, Climateprediction.net, CAS@home
 - CERN (Theory, Atlas, CMS, LHCb)

BOINC and Virtualisation - 2



- BOINC distributes VMs to client machines along with a wrapper application
- The BOINC client installation for Windows now includes Virtual Box

Classic BOINC vs. Virtualisation

Classic BOINC	Virtualisation
<ul style="list-style-type: none">• Applications are native binaries• Unknown environment• Multitude of OS• Application building/testing and result verification is very labour-intensive	<ul style="list-style-type: none">• Applications run in a VM• Typical hypervisor; VirtualBox (installed with BOINC on some OS)• Application to be built for one environment only
<ul style="list-style-type: none">• BOINC takes care of job management• Local application framework must be integrated with BOINC	<ul style="list-style-type: none">• BOINC takes care of distributing VM image• External job manager possible

Applications, Deployment, Relationships

Volunteers	Institute desktops, server farms
<p>Suitable applications:</p> <ul style="list-style-type: none">• CPU-centric with relatively small I/O• watch memory footprint	<p>Suitable applications:</p> <ul style="list-style-type: none">• Watch memory footprint
<ul style="list-style-type: none">• Runs sandboxed in unknown environment	<ul style="list-style-type: none">• Known and controlled environment• Clients start as service, all running as the same user
<ul style="list-style-type: none">• Publicity and marketing essential• Outreach, user forums, user credits	<ul style="list-style-type: none">• Mostly irrelevant

BOINC in CERN/IT


- BOINC server and application support for LHC@home 2004-2006
 - Activity stopped, exile at Queen Mary
- Renewed effort from 2011
 - Virtualisation support, launch of LHC@home 2.0 Test4Theory
 - LHC@home brought back to CERN
 - New team in IT/PES (0.5 FTE + Tech Student now)
 - Move from project support to more generic BOINC service

BOINC Service at CERN

- BOINC server cluster
 - LHC@home servers
 - Sixtrack, Theory, ATLAS
 - Test servers (LHCb, project with EPFL, dev environments)
- BOINC server application support
 - Configuration, monitoring
 - MySQL database server back-end
 - BOINC server application configuration and updates
- *Handled by the user project teams:*
 - *Porting of applications to BOINC*
 - *Application specific job management framework*
 - *Communication with users about scientific projects*
 - *Content of forums and portal*

BOINC at CERN – recent developments

- Service Consolidation
 - Drupal portal for lhathome.cern.ch
 - Started to use OpenStack for VM servers for BOINC server
 - Puppet managed machines.
 - Allows the quick creation of servers and even clients for testing purposes
 - Use of centralised MySQL service
 - BOINC server code modified by Tomi Asp (Graduate student at CERN for 1 year)
 - Use of centralised NFS disk space
 - Take advantage of the central services features and support



Planned service improvements

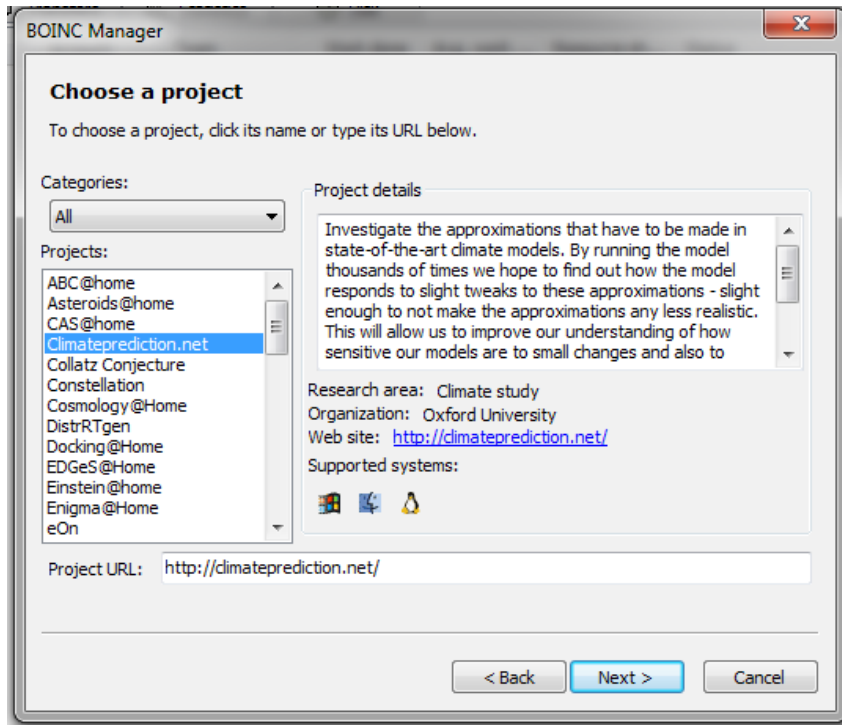
- Scalability: Scale out with separate servers for download/upload etc
 - Ceph buffer for data cache
- Tools to automate deployment of applications and upgrades
- Upgrade [Drupal portal](#) to updated environment, adding elements used by other large BOINC sites
- Security and usability enhancements
 - Https (traditionally not used by BOINC projects)
 - Alternative login methods (e.g. Open ID)

LHC@home legacy and diversity

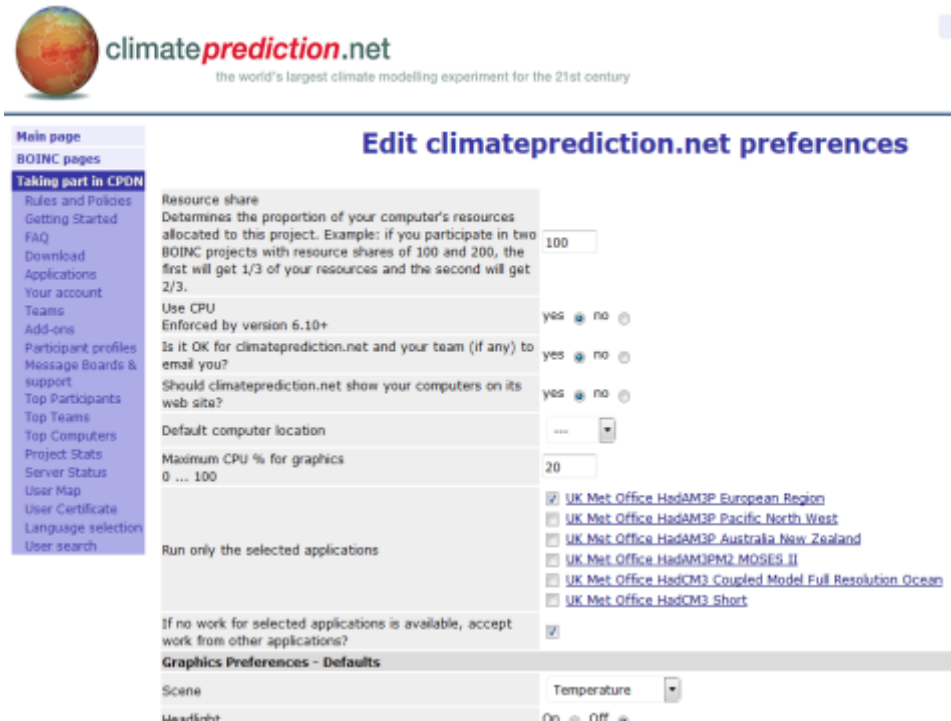
- Sixtrack uses BOINC in the classic way
- Other projects use CernVM and the Vboxwrapper application
- Common denominator for VM projects:
 - cernvm image
 - getting apps from Cernvmfs file system
 - e.g. BOINC is used to extend the local cloud to volunteers
- However different job managers are used
 - (Co-Pilot, Panda+ArcCE, Crab, Dirac...)

How to organise many BOINC projects?

Individual projects



One project with many apps



Current BOINC service approach

- Add VM applications that report back to a local job management framework to the [Virtual LHC@home](#) BOINC project
- Other projects (Sixtrack, ATLAS) currently hosted on separate servers to avoid I/O bottleneck
 - A distributed server setup with separate servers for upload/download would allow for a single project
- Use [Drupal portal](#) as common entry point for all BOINC projects and applications
- Aim for standardisation on a volunteer cloud common job management solution in the future

Conclusions

- Volunteer computing offers a lightweight way to distribute jobs
- BOINC is the de-facto standard middleware for volunteer computing
- Thanks to virtualization support, BOINC is now suitable for a wider range of HEP applications
- Applications running under CernVM and getting data from CernVMFS can be hosted as part of LHC@home
- The size of the application data sets is a bottle neck
- Outreach and communication is essential to get contributions from the general public
- Desktops and other opportunistic local resources offer capacity that can be exploited

Acknowledgements

- BOINC service: Pete Jones, Tomi Asp, Alvaro Gonzalez
- Also Miguel Marquina, Helge Meinhard, Manuel Guijarro, Ignacio Reguero
- Test4Theory: Ben Segal, Peter Skands, Jakob Blumer, Ioannis Charalampidis, Artem Harutyunyan, Predrag Buncic, Daniel Lombrana Gonzalez, Francois Grey et al
- Sixtrack: Eric McIntosh, Riccardo de Maria, Massimo Giovannozi, Igor Zacharov et al
- ATLAS: David Cameron, Andrej Filipic, Eric Lancon, Wenjing Wu
- CMS: Laurence Field, Hendrik Borrás, Daniele Spiga, Hassan Riahi
- LHCb: Federico Stagni, Joao Medeiros et al
- BOINC: David Anderson, Rom Walton
- All our IT colleagues offering a layered service, DB on Demand, Openstack, Puppet, AFS, NFS filers, Linux, network... :-)

References

- <http://cern.ch/lhcathome>
- <http://boinc.berkeley.edu>
- Contact the [BOINC service team](#) at CERN