

# LHC@home – Past, Present, Future

## Volunteer Computing at CERN

Helge Meinhard, Nils Høimyr / CERN  
for the CERN BOINC service team



# Outline

- Why volunteer computing
- Infrastructure: BOINC
- BOINC and virtualisation
- LHC@home applications
- BOINC service at CERN-IT



# Why Volunteer Computing?

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



# Infrastructure / Middleware


- Most commonly used middleware: BOINC
  - Other choices: XtremWeb, 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
  - ...

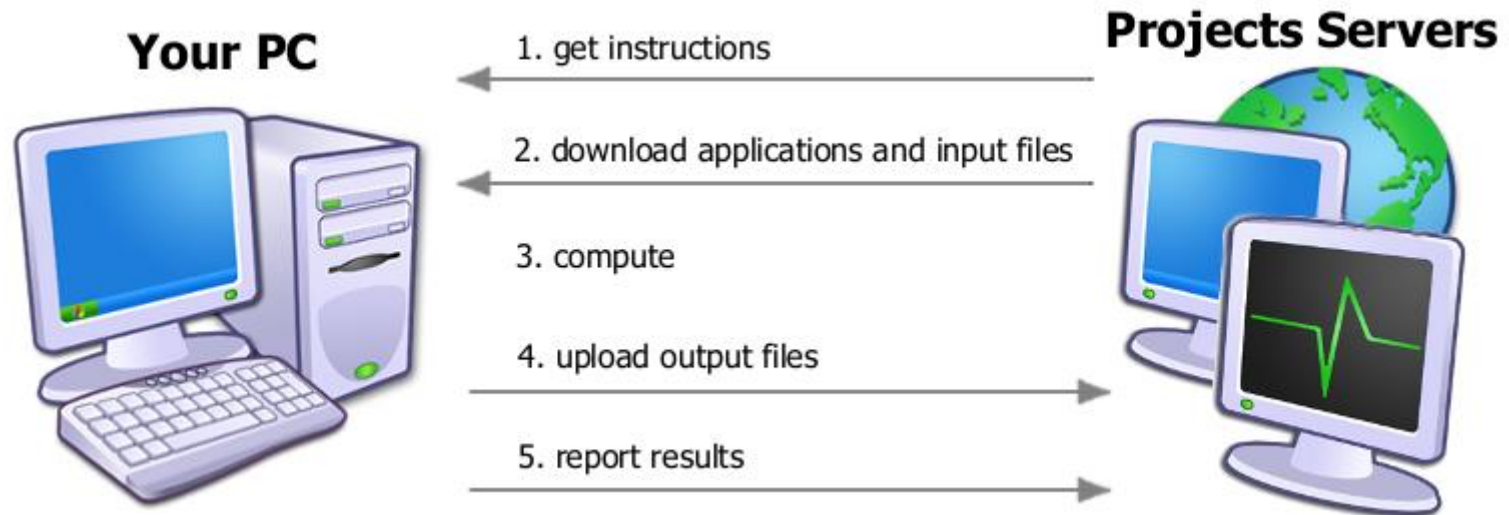


<http://boinc.berkeley.edu>



# 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, your are done!



# BOINC Compute Power

Project	Average power	<i>Equivalent dual Haswell</i>
<u>SETI@home</u>	734 TFlops	4'000
<u>Einstein@home</u>	558 TFlops	3'100
<u>World Community Grid</u>	408 TFlops	2'250
<u>LHC@home</u>	12 TFlops	65
<u>LHC@home 2.0</u>	2.4 TFlops	13

(BOINCstats.com 14-Oct-2014)

*Equivalent Haswell is coarse estimate,  
input by Manfred and Michele*



# Applications, Deployment, Relationships

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



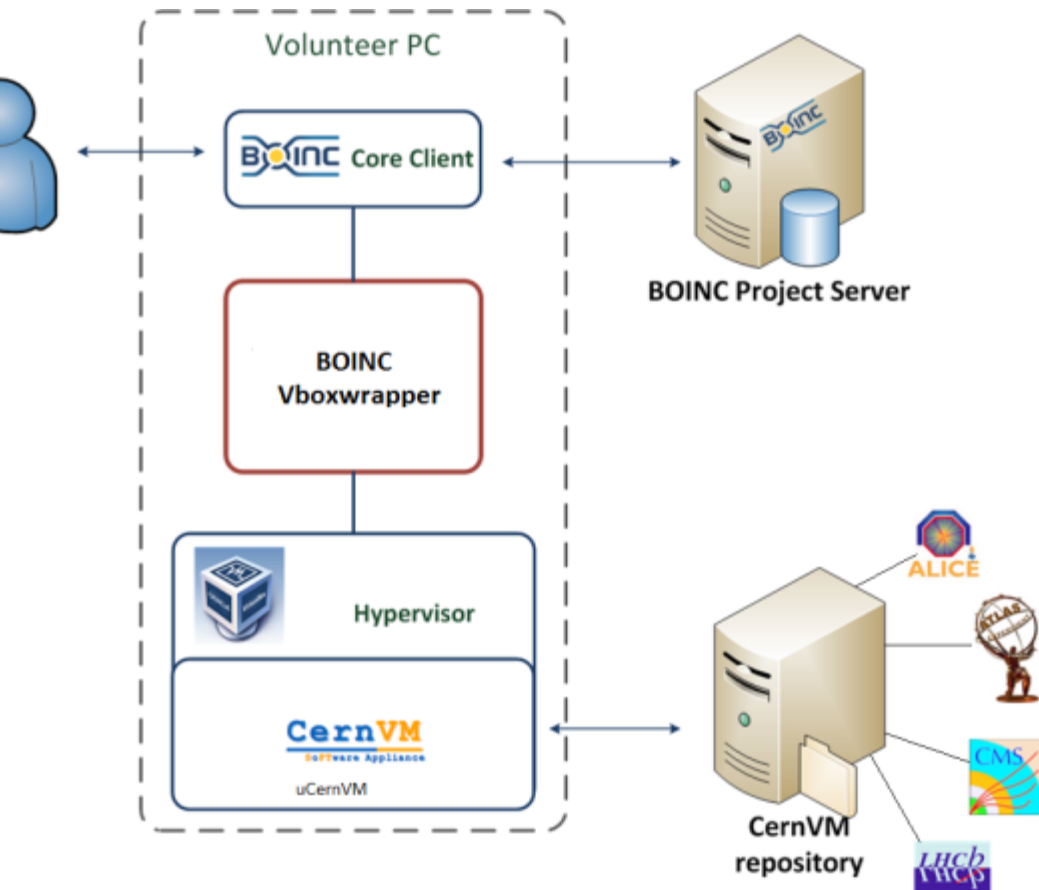


# Classic BOINC vs. Virtualisation

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



# BOINC and Virtualisation



- BOINC distributes VMs to client machines
- Vboxwrapper now part of official BOINC distribution

# LHC@home Applications: Sixtrack

- Started as outreach project for CERN's 50th anniversary 2004, used for Year of Physics (Einstein Year) 2005
- Calculates stability of proton orbits in the LHC accelerator
- Written in FORTRAN, simulates particle trajectories
- Uses the classic BOINC approach
- Client runs on Linux, Mac and Windows platforms
- Renewed effort for LHC upgrade studies (HL-LHC)
- Total 118'000 volunteers, about 20'000 active recently
- Compute power: Peak 40 TFlops, average 12 TFlops



# LHC@home Applications: Test4Theory

- Launched 2011 in partnership with the Citizen Cybercience Centre – CCC
- Theoretical fitting of all past experimental data (including LHC) using Monte Carlo simulation based on Standard Model
- One trillion events simulated by volunteers since 2011
- Portal changed name in 2014 to **Virtual LHC@home**
  - View to host more applications
- Uses a virtual machine on the volunteer computers
- User installs the BOINC client and VirtualBox
- BOINC client downloads a VM-wrapper “vboxwrapper”, that gets the image (CernVM) and the job
- Job reads data from CernVMFS
- External job management: CoPilot
- CernVM, CernVMFS, CoPilot: developed by CERN PH-SFT
- Wide range of potential (physics) applications




# LHC@home Applications: ATLAS

- Started as pilot within ATLAS early this year, now public
  - <http://atlasathome.cern.ch>
  - ARC CE integrated with BOINC, using PanDA for job management
- Supports simulations and potentially other types of ATLAS jobs
  - Job size and 64bit image limits to “hardcore” volunteers
  - Already significant CPU contribution (ref PanDA monitor)
- Integration with LHC@home environment in progress
  - BOINC server migration to IT BOINC service
  - ARC-CE and BOINC sharing data via NFS



# Running jobs

168 Hours from 2014-09-17 to 2014-09-24 UTC



- |                           |                  |                           |                         |                        |
|---------------------------|------------------|---------------------------|-------------------------|------------------------|
| ■ BNL-ATLAS               | ■ RAL-LCG2       | ■ CERN-P1                 | ■ CERN-PROD             | ■ INFN-T1              |
| ■ DESY-HH                 | ■ MWT2           | ■ TRIUMF-LCG2             | ■ AGLT2                 | ■ NDGF-T1              |
| ■ FZK-LCG2                | ■ BOINC          | ■ BU ATLAS TIER2          | ■ UKI-NORTHGRID-MAN-HEP | ■ UKI-SCOTGRID-GLASGOW |
| ■ UKI-LT2-QMUL            | ■ SIGNET         | ■ INFN-NAPOLI-ATLAS       | ■ PRAGUE LCG2           | ■ UKI-SOUTHGRID-RALPP  |
| ■ WT2                     | ■ IN2P3-CC       | ■ LRZ-LMU                 | ■ TAIWAN-LCG2           | ■ UNIBE-LHEP           |
| ■ CERN-T0                 | ■ UKI-LT2-RHUL   | ■ SARA-MATRIX             | ■ WUPPERTALPROD         | ■ SWT2_CPB             |
| ■ UKI-NORTHGRID-LANCS-HEP | ■ UKI-LT2-BRUNEL | ■ GOEGRID                 | ■ UKI-SOUTHGRID-OX-HEP  | ■ TOKYO-LCG2           |
| ■ OU OCHEP SWT2           | ■ NIKHEF-ELPROD  | ■ CSCS-LCG2               | ■ IN2P3-LAPP            | ■ SFU-LCG2             |
| ■ UNF-FREIBURG            | ■ HU_ATLAS_TIER2 | ■ CA-VICTORIA-WESTGRID-T2 | ■ CA-MCGILL-CLUMEQ-T2   | ... plus 82 more       |

Maximum: 133,280 , Minimum: 37,990 , Average: 91,338 , Current: 86,320



# LHC@home Applications: Other

- LHCb: Beauty
  - Since 2012, still in test phase
  - Volunteers from within LHCb collaboration
  - Job management: DIRAC
  - Vboxwrapper application with CernVM image
- CMS
  - Started work in summer 2014
  - Very early time, prototype running
  - Job management: Crab linked to messaging service
  - To be added as beta application on Virtual LHC@home once stable
- Adopt a Neuron
  - Pilot project with EPFL for Human Brain project



# CERN-IT BOINC Service (1)

- BOINC server cluster
  - LHC@home servers
    - Sixtrack, Theory, ATLAS
    - Test servers (LHCb, Adopt a Neuron, 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*





# CERN-IT BOINC Service (2)

- Recent evolutions
  - Drupal portal for [lhathome.cern.ch](http://lhathome.cern.ch)
  - OpenStack puppet-managed virtual machines as BOINC servers
    - Allows for quick creation of servers and even clients for testing purposes
  - Use of centralised MySQL service (DB on Demand)
    - BOINC server code modified by Tomi Asp (Graduate student at CERN for 1 year)
  - Shared NFS space for project hosting
    - Take advantage of CERN NFS filer service (part of CERN IT's layered approach)



# Conclusion

- Volunteer computing offers a lightweight way to distribute jobs
- Thanks to virtualization support, BOINC is now suitable for a wider range of HEP applications
- Outreach and communication is essential to get contributions from the general public
- 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
- Stay tuned for more information:
  - Pre GDB on Volunteer Computing 11<sup>th</sup> of November
  - Update at a future HEPiX meeting:-)



# References

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

