
CERN

Update on Volunteer Computing

GDB 9-Sept 2015

Nils Høimyr, IT/PES on behalf the BOINC service team and LHC@home

Outline

- Why Volunteer Computing?
- LHC@home
  - Status of different projects and applications
- BOINC Service at CERN
- Summary
- Questions

Why volunteer computing?

- Free* resources
  - 100K hosts achievable for large projects
    - Actual core count is higher
    - Support for virtualisation - CernVM
- Community engagement
  - Outreach channel
  - Community participation and support

* Attracting and interacting with volunteers has a cost

The Vacuum Model

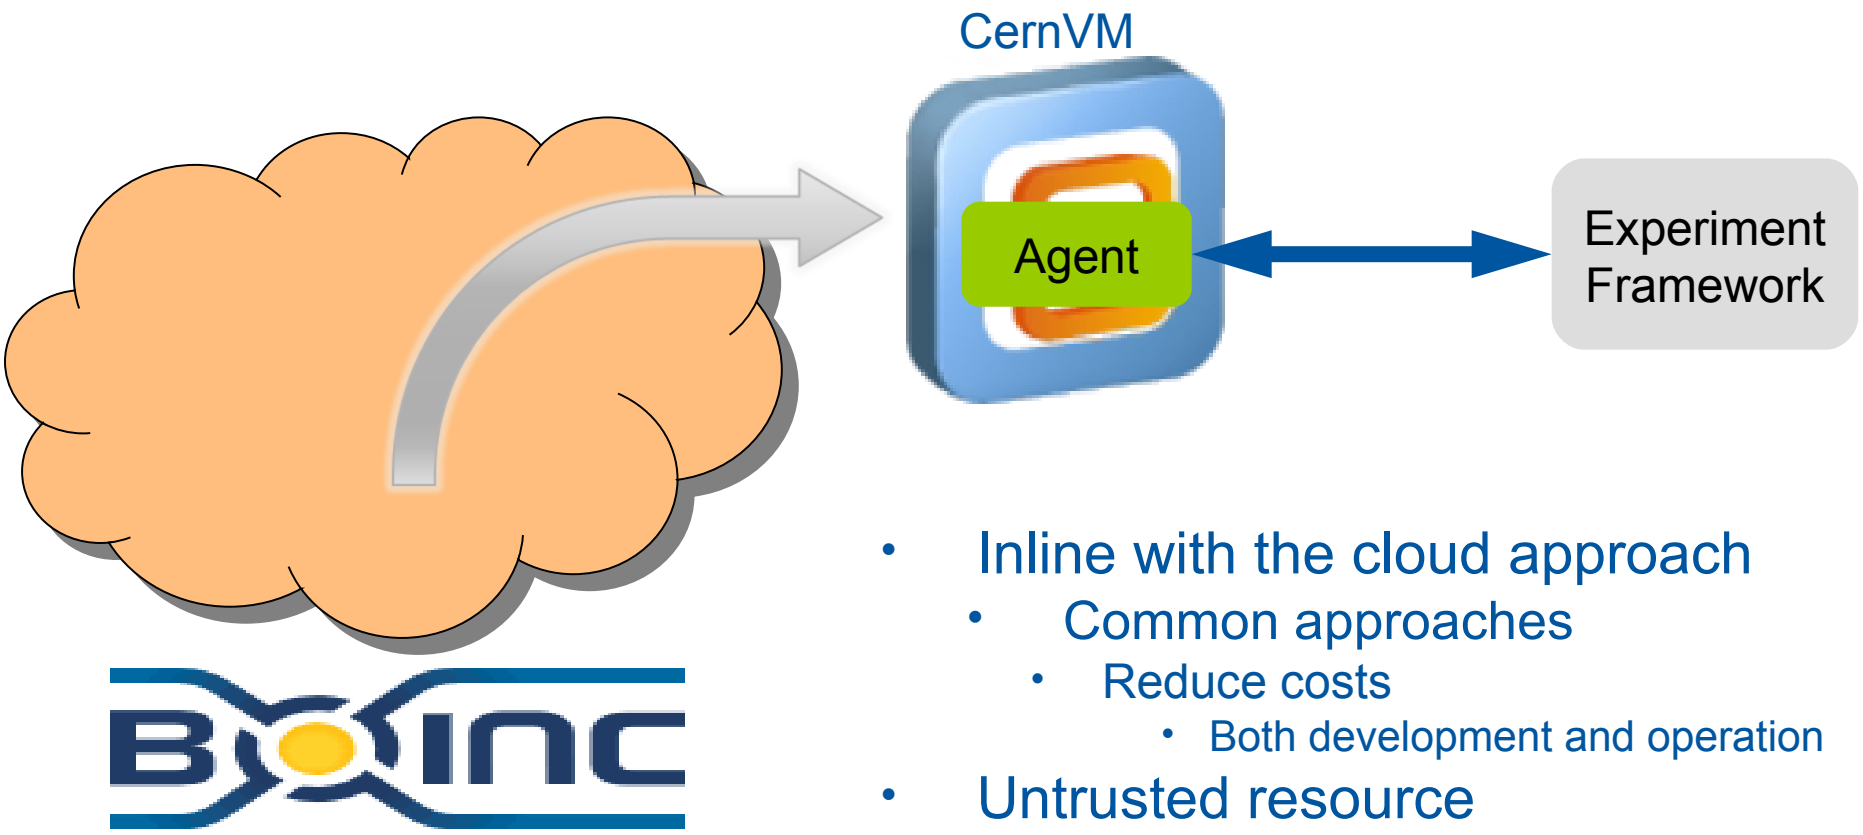

- Inline with the cloud approach
  - Common approaches
    - Reduce costs
      - Both development and operation
- Untrusted resource
  - Authentication
  - Validation

LHC@home
Volunteer computing for the LHC

Search this site

Search

HOME ABOUT PROJECTS JOIN US! HELP & FAQ CONTACT

Welcome

What is the Universe made of? How did it start? Physicists at CERN are seeking answers, using some of the world's most powerful particle accelerators.

LHC@Home is a volunteer computing platform where you donate idle time on your computer to help

News from the forums

- BOINC 7.6.9 release for Windows and Mac
  01 Sep 2015 / BOINC news

SixTrack (LHC machine)

- Original classic BOINC project for beam simulations
  - Calculates stability of proton orbits in the LHC accelerator
    - Simulates particle trajectories
- Based on experience from the Compact Physics Screensaver (CPSS)
  - Ran SixTrack on desktop computers at CERN
- Outreach project for CERN's 50th anniversary 2004
  - Also Year of Physics (Einstein Year) 2005
- Application written in FORTRAN
  - Runs on Linux, Mac and Windows platforms
- Renewed effort for LHC upgrade studies (HL-LHC)
  - 12K Active Users
  - 19K Active Hosts
  - 45 TeraFLOPS peak

Test4Theory

- Theoretical fitting of all past experimental data
  - Including from the LHC
  - Using Monte Carlo simulation based on Standard Model
- Launched 2011
  - In partnership with the Citizen Cyberscience Centre (CCC)
- Pioneered use of Virtualization with BOINC
- Uses recent developments from CERN's PH-SFT Group
  - CernVM
  - CernVMFS
  - CoPilot, later also Data Bridge from IT/SDC
    - CERN 60 year challenge with CernVM web api
- Wide range of potential (physics) applications
  - In 2014 changed name to **Virtual LHC@home**

Test4Theory Usage

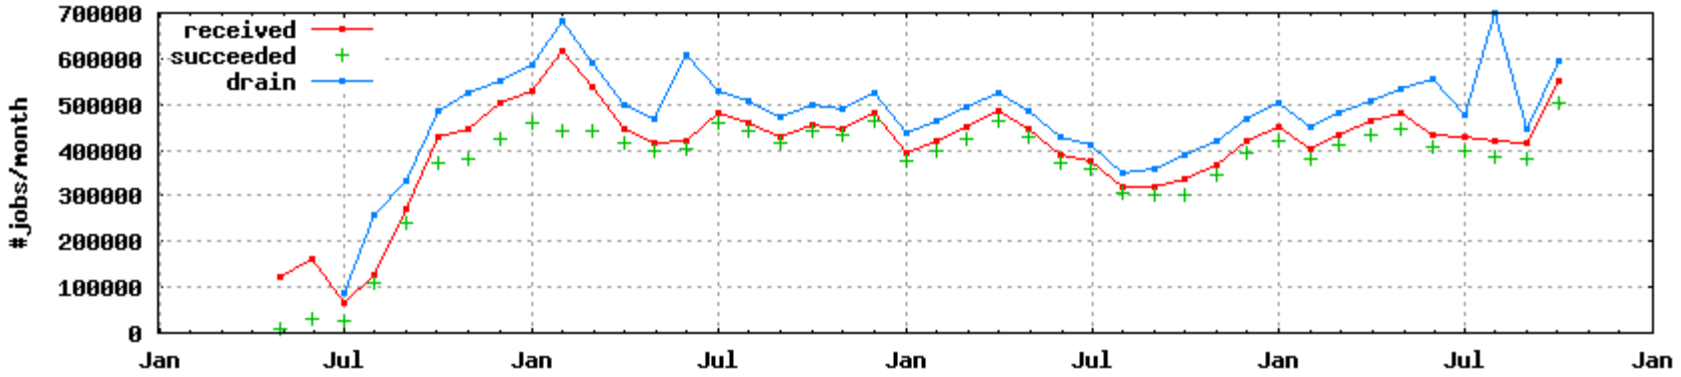

- More than 2 trillion events simulated since 2011
- Source: MC Plots (http://mcplots-dev.cern.ch/production.php)
- See also: http://cern.ch/go/9nRz

ATLAS@home

- Started as pilot beginning of 2014
  - Open to the public since more than a year
    - http://atlasathome.cern.ch
- Also using CernVM and virtualization
  - Classic BOINC model
- ARC CE used to interface with BOINC
  - PanDA for job management
- Supports simulations
  - Potentially other types of ATLAS workloads
- Job size and 64bit image limits to “hardcore” volunteers
  - Already significant CPU contribution
- Integrated with LHC@home environment
  - ARC-CE and BOINC sharing data via NFS

ATLAS@home Usage

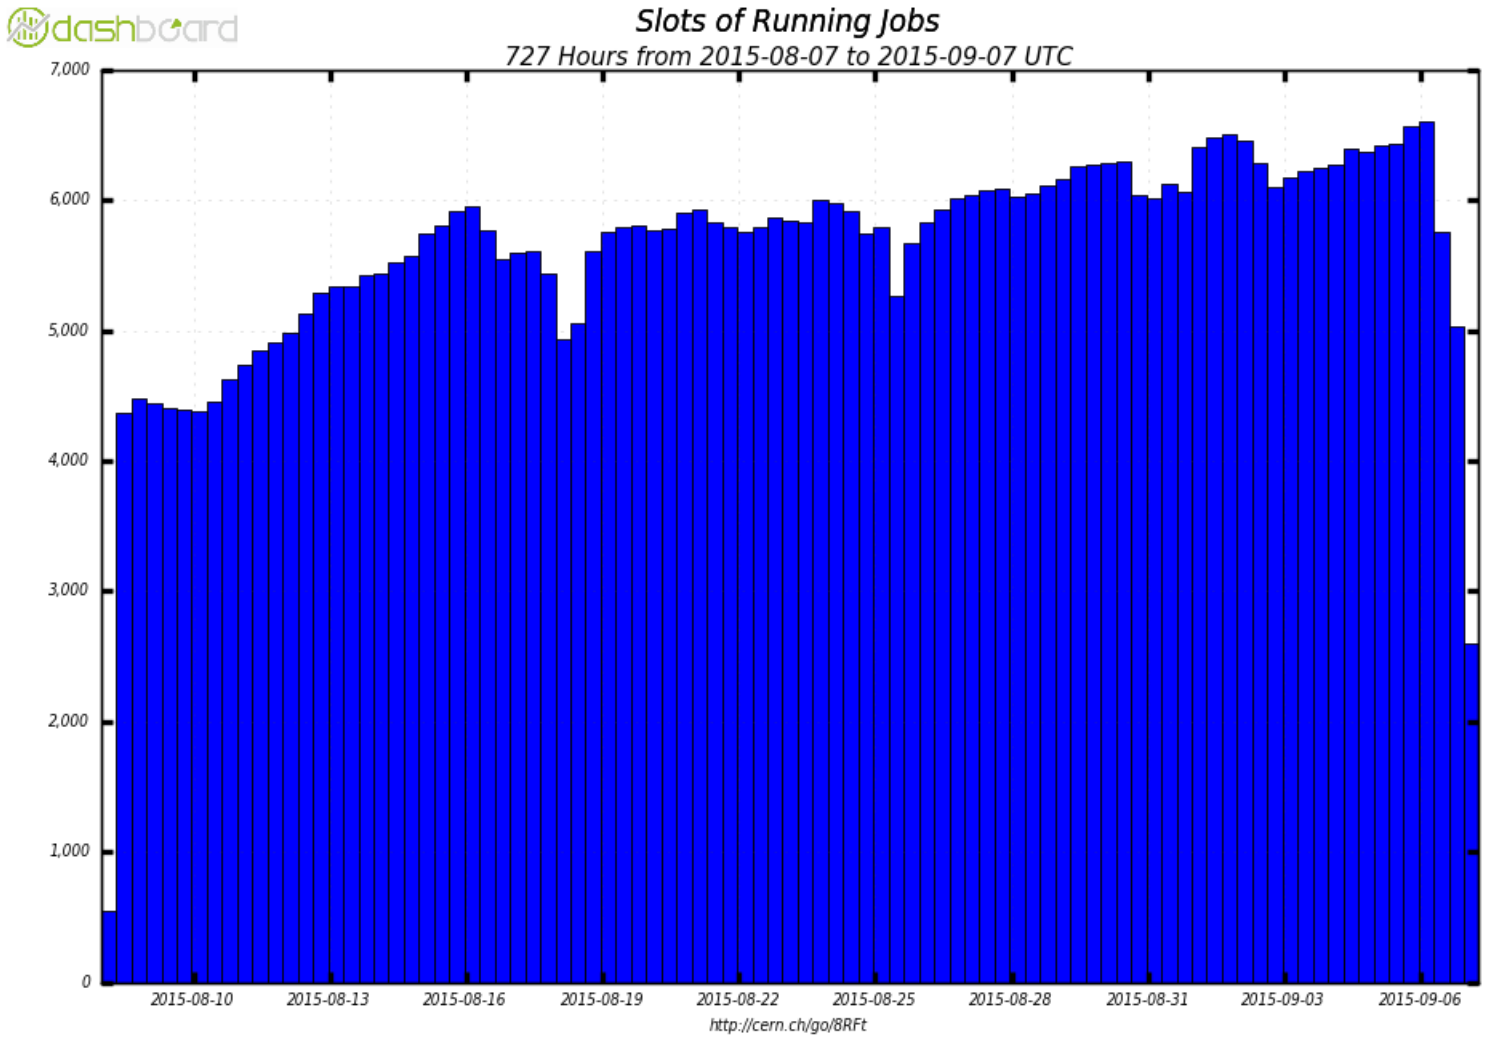

Beauty@home

- In pilot mode since 2012
  - Requires x509 credential in the client VM
    - Volunteers from within LHCb collaboration
    - Development of new proxy layer soon ready
      - Expect to be ready for the general public soon (2015)
- Communicates directly with DIRAC
- Vboxwrapper application
  - Using uCernVM

CMS@home

- Initial development summer 2014
  - Prototype service running
    - Rapidly gaining experience
    - Many improvements over the last months
      - Integration with CMS production environment on-going
      - Data Bridge component available for other projects
- To be added as application in vLHC@home
  - Once stable with production jobs

BOINC Service Monitor

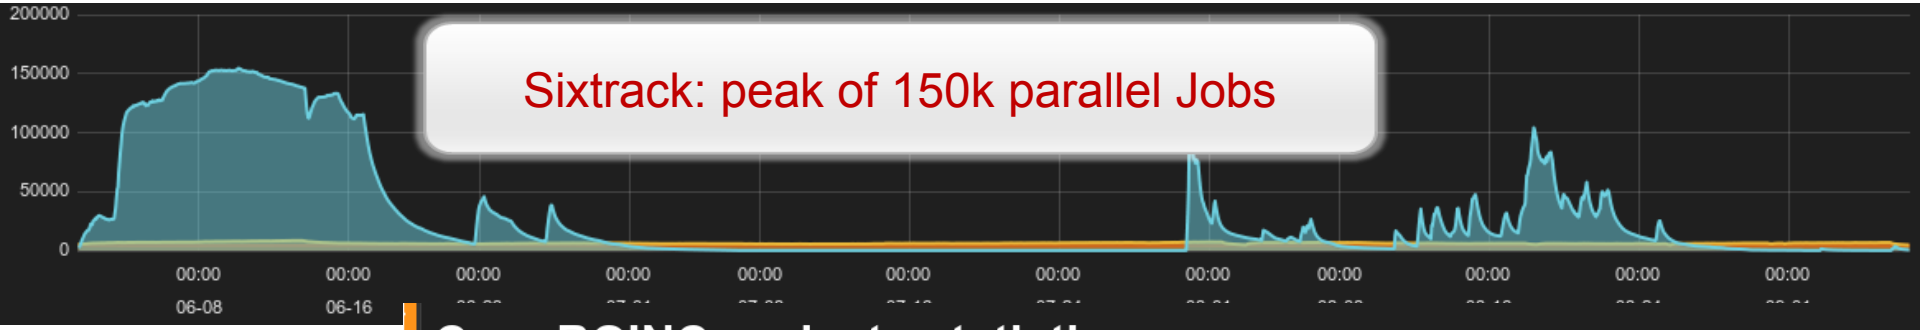

Sixtrack: peak of 150k parallel Jobs

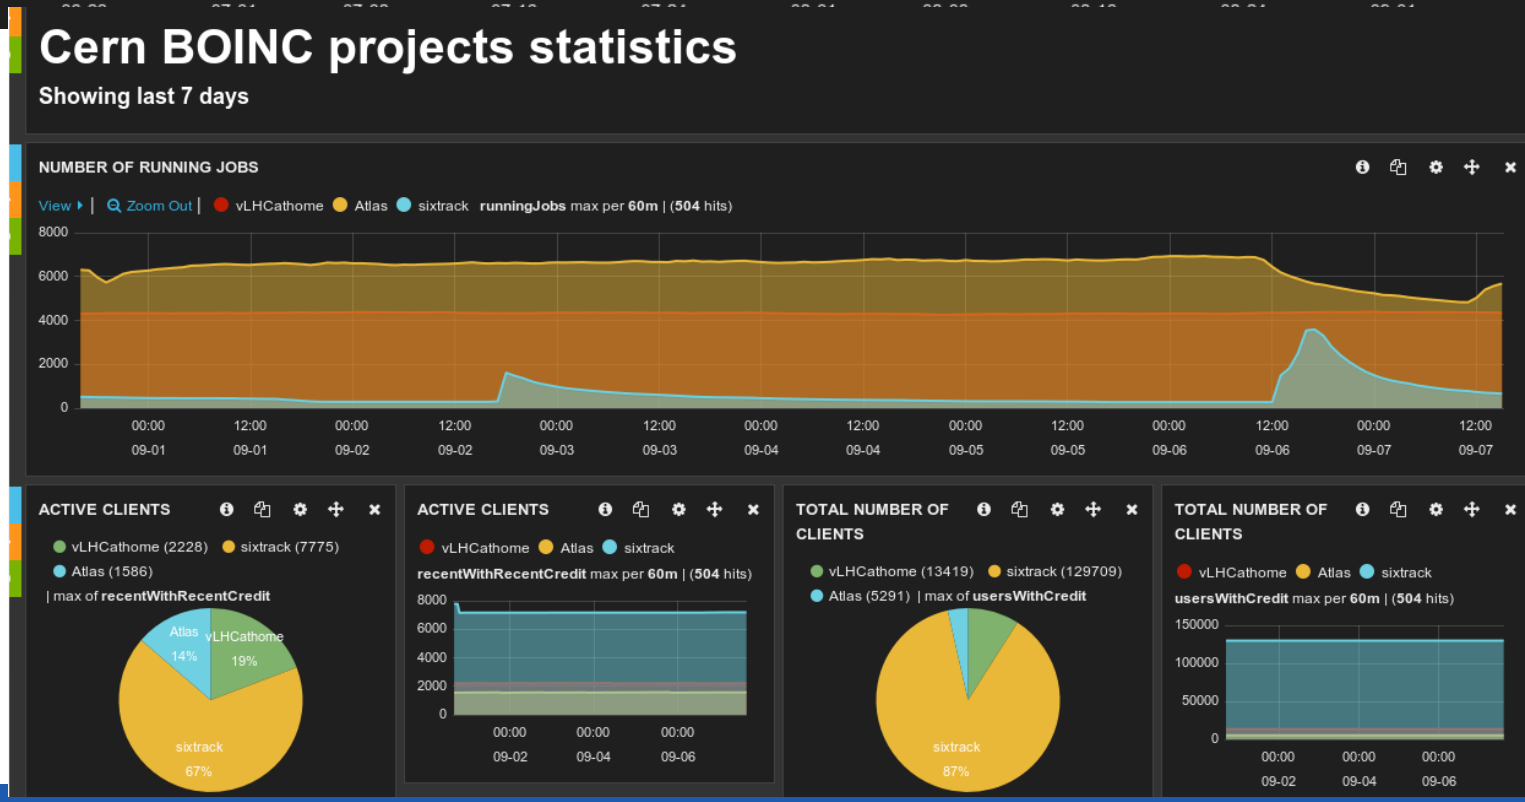

http://cern.ch/go/9nRz

BOINC at CERN – service layers

- Service Infrastructure
  - Drupal portal for lhcathome.cern.ch
  - OpenStack VMs used for the BOINC server cluster
    - Puppet managed machines.
    - Allows the quick creation of servers and even clients for testing purposes
  - MySQL DB back-end on DB on Demand service
    - BOINC server code modified and incorporated with upstream release
  - NFS back-end for project space and data buffers
    - Recently migrated from Netapp to Ceph-based NFS servers

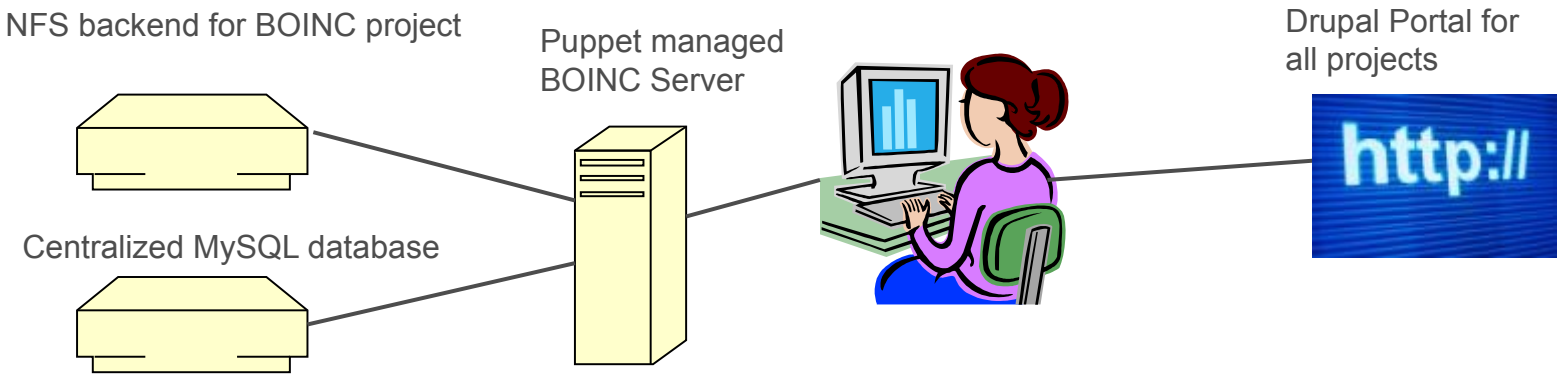

BOINC service evolution

- VM applications that report back to a local job management framework can be part of the Virtual LHC@home BOINC project
- Other projects (Sixtrack, ATLAS) currently hosted on separate servers to avoid I/O bottleneck
  - Distributed BOINC server setup tested for improved scalability
- Use Drupal portal as common entry point for all BOINC projects and applications
  - Improved experience for volunteers, plan to streamline registration
- Aim for standardisation on a volunteer cloud common job management solution
  - Data Bridge component available from IT/SDC

Summary

- 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 hosted as part of LHC@home
- The experiments are trying to exploit this opportunistic resource
  - Several production and pilot applications underway
- Requires investment
  - Initial integration
  - Attracting volunteers
  - Supporting volunteers via the forum
- Work towards a common platform
  - Share Development, Maintenance and Operations

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, Anton Karneyu et al
- Sixtrack: Eric McIntosh, Riccardo de Maria, Massimo Giovannozi, Igor Zacharov et al
- ATLAS: David Cameron, Andrej Filipic, Eric Lancon, Wenjing Wu, Efrat Tal Hod
- CMS: Laurence Field, Hendrik Borras, Daniele Spiga, Hassan Riahi
- LHCb: Federico Stagni, Cinzia Lucci, 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
  - More details in talks at the preGDB last fall.

www.cern.ch