Towards a Production Volunteer Computing Infrastructure for HEP

2015 J. Phys.: Conf. Ser. 664 022023

(http://iopscience.iop.org/1742-6596/664/2/022023)

Towards a Production Volunteer Computing Infrastructure for HEP

N Høimyr[1], M Marquina[1], T Asp[1], P Jones[1], A Gonzalez[1] and L Field[1]

[1] CERN, CH-1211 Geneva, Switzerland

E-mail: Nils.Hoimyr@cern.ch

Abstract. Following the successful inclusion of virtualisation to volunteer computing for theory simulations back in 2011, the use of volunteer computing with BOINC and CernVM has been extended to cover simulations for the LHC experiments ATLAS, CMS and LHCb. This paper describes the status of the BOINC volunteer computing platform at CERN used for LHC@home and how it has been designed to address a heterogeneous environment of different user communities with different computing infrastructure. The aim of the recent developments is to provide a volunteer computing platform that the experiments can build upon to exploit opportunistic resources. Furthermore, new developments of common solutions to span user authentication domains are explained.

1. Introduction

Volunteer computing allows projects to harvest computing power from opportunistic resources, such as idle office desktops, server farms, or contributions from volunteers among the general public who donate CPU time on their home PCs. In addition to accessing computing power that otherwise would not be available, volunteer computing also provides for outreach and publicity for scientific projects. In this paper, we describe the dominant technology for volunteer computing, how it has been adapted to support virtualisation, and the current status of volunteer computing projects at CERN. Finally, we explain the underlying service architecture and its components, and how it addresses the particular challenges of volunteer computing for the High Energy Physics (HEP) community.

2. BOINC

The *de-facto* standard software for volunteer computing projects is the Berkeley Open Infrastructure for Network Computing (BOINC) [1] that emerged in 2002 from Climateprediciton.net [2] and the experience of SETI@home [3] at the University of California, Berkeley. Since then, a number of BOINC-powered volunteer computing projects have been launched, including LHC@home [4] in 2004. The popularity of the BOINC software and the volunteer community has evolved considerably. Today active BOINC projects together harness about 7.5 Petaflops of computing power [5].

The BOINC system provides a client server application, with a Web server to handle scheduling, task and user management as well as forums to communicate with the users. BOINC clients are available for major operating systems ranging from Windows, MAC-OS and Linux to Android.

The BOINC server hosting a given project distributes an application and jobs to all volunteer clients that have attached to it. The project server receives results from the finished jobs, validates them and deals with issues of client downtime or unreliability. BOINC also provides message boards to allow communication between volunteers and the project administrators. One of the aspects behind the successful uptake of BOINC is that volunteers are awarded credit for their contributions. Although the credit scheme is purely virtual, for part of the community of volunteers competing for credit is a source of motivation for them to contribute CPU-power.

2.1. Virtualisation support in BOINC

Applications traditionally running under BOINC initially had to be compiled for each and every possible client operating system. Thanks to developments at CERN [6,7], this step can be avoided by distributing a Virtual Machine (VM) [8] to the volunteer computers via BOINC. This use of virtualisation under BOINC was pioneered by the Test4Theory LHC@home project during 2010-2011 [6, 7, 9].

The VM runs under a hypervisor such as Oracle's *VirtualBox* [10] that is installed on the local computer. As described by Harutyunian et al [7] and Høimyr et al [9], using the CernVM [6] image with BOINC would address porting issues for physics applications. Virtualisation technology was integrated with BOINC by using a VM wrapper initially prototyped at CERN, and later included as part of the official BOINC software distribution as the *vboxwrapper* [11]. The wrapper uncompresses the VM image after it has been downloaded to the client, launches it and controls it. For example, it will pause or resume the VM when the BOINC task has to be paused or resumed to enforce the computing preferences of BOINC.

Use of virtualisation with BOINC allows for running VMs with a HEP Linux flavour (e.g. Scientific Linux) on any volunteer computer host that can run a local hypervisor such as Virtual Box. This broadens the potential of volunteer computing for High Energy Physics applications, as there is no need to port applications to operating systems, nor to link the application with the BOINC system libraries. In addition to using computers belonging to volunteers, there is also scope for using this technology on other opportunistic computing resources, such as desktop grids and (temporarily) available clusters in HEP computer centres.

Other technologies than BOINC for volunteer computing are not addressed in this paper, but the evolution of software in this domain should be monitored as there may be other opportunities in the future to use simpler, fully Web-enabled technologies. (Notably in the domain of crowd-sourcing there has been an evolution of user-friendly Web applications with a lower threshold for volunteer participation than for BOINC, such as Zooniverse [12].) Such environments allow to run smaller applications directly in the browser, e.g. by converting the application into Java-script byte code.

3. LHC@home Applications

Volunteer computing has been used successfully at CERN since 2004 with the LHC@home project, and has provided additional computing power for CPU intensive applications with small data sets as well as an outreach channel for CERN activities. The following is a short description of the projects that are currently in production or development at CERN.

3.1. Sixtrack

Sixtrack [13] is a program for the simulation of charged particle trajectories in accelerator beams; it has been running under LHC@home since 2004. Some 120,000 users with more than 300,000 PCs have been active LHC@home volunteers since its launch. This has provided significant computing power for accelerator physics studies, for which there was no equivalent capacity available in the regular CERN compute clusters. The simulations carried out with Sixtrack under BOINC have provided important insights for the improvement to the LHC machine performance.[14].

The Sixtrack code is FORTRAN-based and was ported for use with BOINC to Windows and Linux by incorporating calls to the BOINC application programming interface (API) library and recompiling and re-linking the source code to produce executables for each client platform. Since 2004, the application code has undergone several updates to adapt to new BOINC versions as well as to improvements to Sixtrack itself [15, 16]. Recently the studies with Sixtrack have been extended to include simulations for the future HL-LHC. Volunteers contributing to Sixtrack have delivered a sustained processing capacity of more than 45 TeraFlops.

3.2. Virtual LHC@home (Formerly Test4Theory)

As a joint effort by the CernVM team and CERN-IT, the first BOINC project using virtual machines was launched in 2011: the LHC@home Test4Theory project. In this project BOINC is used to dispatch CernVM images that are pre-contextualized to run simulations based on Monte Carlo (MC) models of high energy collider physics using different event generators. The simulation software is distributed via the CernVMFS file system [17] that is mounted in the VM. The distribution of jobs to VMs and validation of results was achieved using the CernVM-CoPilot job management layer [7], thus decoupling the job handling from BOINC. With this approach the VMs distributed by BOINC to volunteers' computers appear similar to cloud resources to the application scientists [18] and the VM images can also potentially be launched in private or public cloud computing environments.

During three years of operation, the volunteers connected to the project have generated more than 1.7 trillion collision events, providing a large amount of simulated MC statistics that can be compared with experimentally measured results from detectors at CERN. On average the project has been providing more than 3 TFlops of sustained computing power.

With the vision to add more applications from the LHC experiments, the project was renamed to Virtual LHC@home in the summer of 2014, in order to reflect the broader scope of applications beyond theory simulations.

3.3. Beauty@home

The LHCb collaboration were the first to explore the use of volunteer computing among the LHC experiments with the *Beauty@home* BOINC test project from 2012. The approach is similar to the theory simulations. However, only invited volunteers, who are members of the LHCb collaboration, are accepted as an X.509 credential from an IGTF-accredited CA is required to communicate directly with LHCb's Dirac [19] distributed computing framework and hence can only be used by trusted parties. Efforts to address this issue are currently ongoing and once ready, the application can be added to the Virtual LHC@home project alongside with the theory simulations application, after which the general public be allowed to contribute.

3.4. ATLAS@home

Building upon the experience with virtualisation for BOINC at CERN with Test4Theory and at IHEP in Beijing for CAS@home, a pilot project for ATLAS was launched in 2014 and later brought under the BOINC service at CERN as part of the virtual LHC@home infrastructure [20]. Like the other applications, the ATLAS simulations are also using a CernVM image.

For job management ATLAS has included the ARC CE [21] so that jobs from the ATLAS PanDA [22] job management framework can be submitted directly to BOINC similarly to any ATLAS Grid site [20]. Communication between ARC CE and BOINC is done via a shared NFS job buffer on top of a Ceph [23] file system. The system is used to dispatch MC simulations of ATLAS events with Geant4 [24]. To-date approximately 5000 simultaneous jobs are running on computers from about 1300 users. With this level of contribution of about 2 TFlops, BOINC is already the largest ATLAS simulation site. \

Fig.1 illustrates running jobs from LHC@home, with 4000-7000 concurrent jobs of ATLAS and Theory simulations, and a peak up to 200 000 running jobs of Sixtrack, during their simulation campaign in spring. Sixtrack (in blue), dwarfs the other applications during such simulation batches.

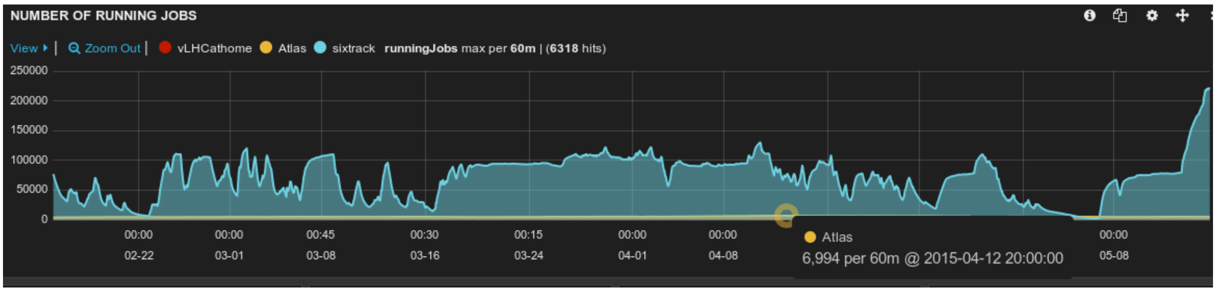

Figure 1. Running jobs on CERN Boinc projects over 3 months

.

3.5. CMS@home

A test project to validate the use of volunteer computing for the CMS collaboration was launched during the summer of 2014 [25]. The project has gradually gained experience and in this context, the concept of a data bridge was developed to span that authentication domains of volunteer and Grid computing. This concept and its concrete implantation is described later in Section 4.2.

4. BOINC Service Infrastructure at CERN

Since the initial deployment in 2011 [9], the BOINC service infrastructure at CERN has evolved to exploit the layered infrastructure offered by CERN's Agile Infrastructure, with Puppet-managed VMs on OpenStack. The project directories for the BOINC servers are hosted on a central Network File System service, with dedicated buffer volumes on Ceph to allow for high I/O rates. Furthermore, the MySQL [26] database back-end for the BOINC application and forums is now hosted using the Database-on-Demand service [27]. (Fig 2.) The database drivers of the BOINC server application had to be adopted to be more flexible to handle this set-up, and the modifications done to the BOINC code at CERN has been included in the upstream code release. This infrastructure allows for rapid creation of pre-configured test servers, thus enabling new application or experiment teams to evaluate BOINC. In addition to projects at CERN, this service infrastructure has also hosted a test project in collaboration with EPFL for the Human Brain Project.

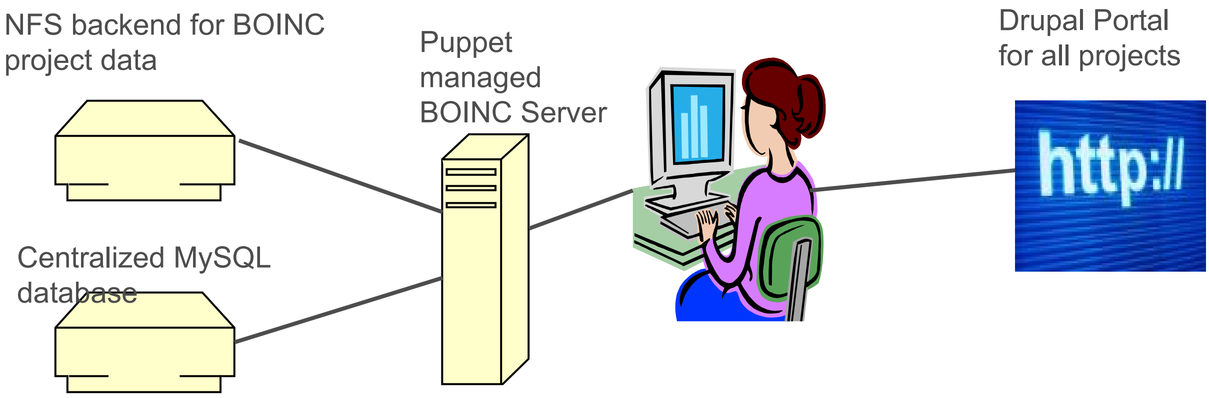

Figure 2. BOINC server infrastructure at CERN.

The current support model for BOINC is to host projects operated by dedicated teams, and to provide the underlying application and server support. Volunteers who wish to connect and contribute via BOINC can attach to one or several of the projects by selecting the corresponding project URL in the BOINC client. Once a VM application has been tested and validated, it can be added as an application to the Virtual LHC@home project. The project server can easily handle more applications from a wider user community as the communication of results of jobs between the VMs on volunteer computers and the job management frameworks of the applications is handled outside BOINC.

For scalability, the Sixtrack and ATLAS applications, which use BOINC directly to handle the application I/O, are currently hosted on separate BOINC project servers. Additionally, the job buffer for the application is stored on a dedicated Ceph-based NFS volume, which allows for scaling up the infrastructure with dedicated upload- and download-servers as part of a distributed BOINC server set-up. While this technology would also allow for hosting of all applications in a single BOINC project, there are also non-technical reasons of sociological and organisational character for our choice of implementation.

If we were to start the BOINC service at CERN from scratch today, all the applications from ATLAS, CMS, Beauty as well as Sixtrack and Theory would have been supported by the same BOINC project. However due to the historical evolution, outreach, discussions in forums and the work done in different teams and locations before a central BOINC service team was established, the applications cannot easily be merged. This diversity of applications and job management environments has been addressed by adding a custom LHC@home look and feel to all the BOINC projects to make them easily recognisable; a common front-end LHC@home Web portal has been established.

4.1. Outreach portal

A Web portal [4] with information about the science behind LHC@home applications as well as instructions and FAQs for users has been set up using the Drupal [28] content management system. The portal is shared among all LHC@home projects, and allows the different teams providing applications for LHC@home to easily create content. It acts as a user friendly entry point for volunteers who wish to contribute.

Projects that are actively advertising and promoting their research to the volunteer community tend to get more active volunteers and more donated computing capacity. Prime examples of this are SETI@home and Einstein@home, where the latter has rewarded top contributors among volunteers with printed diplomas for their contributions. Thus in order to get more contributions from the general public, we should be ready to strengthen the LHC@home brand, and also consider possible rewards in addition to the usual BOINC credit.

4.2. Integration with experiment frameworks

As the LHC experiment collaborations are using different systems for workload management, an important component to ensure wider use of volunteer computing is the link between the experiment job management frameworks and volunteer computing resources. Unlike the Grid and cloud resources that are hosted in data-centres, volunteer computing resources from the general public and other opportunistic resources cannot be trusted and therefore no X.509 credential from an IGTF accredited CA should be delegated.

To address this problem, the concept of the *DataBridge* [25] was developed and initially deployed for the CMS pilot project as well as Theory simulatons. As the resources are untrusted, any output files should be sand-boxed and validated before being assimilated for further use. The *DataBridge* does this by spaning the authentication domains so that output files can be written with the volunteer's credential that is provided by BOINC and read by as standard X.509 Grid credential. Similarly input files can be written to the DataBridge using an X.509 Grid credential and read using the volunteer's credential. This functionality has been provided by building upon the Dynafed component and extending this to support authentication via BOINC. In addition a simple message queue is also available to support the injections of job description if needed.

5. Conclusion and outlook

With virtualisation, volunteer computing can now support a wider range of physics applications and the BOINC service at CERN has evolved to become a more generic platform for volunteer computing. The service infrastructure has been designed to scale up by leveraging CERN's Agile Infrastructure. We anticipate that shortly the DataBridge will be deployed for several applications, and production can be ramped up on the Virtual LHC@home platform. When ready, the CMS and Beauty applications will be added to the virtual LHC@home project that is currently only providing the Theory application. With a view to attract more volunteers, further emphasis will be placed on outreach and consolidation of the Drupal outreach site for LHC@home.

6. Acknowledgements

Volunteer computing under LHC@home would not have been possible without the active participation from the developers and project teams who are providing the applications in this environment. We would like to thank all the teams for collaborating to make LHC@home a success. We would also like to thank the authors of the software packages and IT infrastructure teams who provide the foundation for the set-up of the BOINC and LHC@home project infrastructure.

Last but not least, we would like to express our gratitude to all volunteers that contribute computing power to LHC@home and other BOINC projects and in particular volunteers who help test beta applications and provide advice to other users in the forums.

References:

[1] Anderson, D *Grid Computing, 2004. Proceedings. Fifth IEEE/ACM International Workshop on , vol., no., pp. 4- 10*
[2] Climateprediction.net URL: http://www.climateprediction.net/
[3] Seti@home URL: http://setiathome.ssl.berkeley.edu/
[4] LHC@home portal at CERN. URL: http://lhcathome.cern.ch
[5] BOINC URL: http://boinc.berkeley.edu/
[6] Buncic P e a 2010 *Journal of Physics: Conference Series, 219, 042003*
[7] Harutyunyan A e a 2011 *Journal of Physics: Conference Series,* ***331*** *062013*
[8] Smith J E and Nair R 2005 *Virtual Machines S. Francisco: Morgan Kaufmann, Elsevier*
[9] Høimyr N e a 2012 *Journal of Physics: Conference Series, 396, 032057*
[10] Oracle VM Virtual Box URL: https://www.virtualbox.org/
[11] Vboxwrapper: URL: http://boinc.berkeley.edu/trac/wiki/VboxApps
[12] Zooninverse Citizen Science portal: *URL https://www.zooniverse.org/*
[13] Schmidt F 1994 CERN-SL-94-56
[14] Fartoukh S and Giovannozzi M 2012 *Nucl. Instrum. & Methods A 671*
[15] Giovannozzi M e a "LHC@home" 2012 *International Particle Accelerator Conference*.
[16] De Maria R e a "Recent developments and future plans for Sixtrack" 2013 *International Particle Accelerator Conference*.
[17] Blomer J e a *"Distributing LHC application software and conditions databases using the CernVM file system" 2011 J. Phys.: Conf. Ser. 331 042003*
[18] Segal B e a 2010 *Proceedings of the XIII International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT10), Jaipur*
[19] Stagni F e a 2012 *Journal of Physics: Conference Series, 396 032104*
[20] Cameron D e a ATLAS@Home: Harnessing Volunteer Computing for HEP (CHEP-2015)
[21] Filipčič A e a 2012 *Journal of Physics: Conference Series, 396 032039*
[22] Mauni T e a 2008 *Journal of Physics: Conference Series,* 119 062036
[23] Weil, S e a 2006 In *Proceedings of the 7th symposium on Operating systems design and implementation* (OSDI '06). USENIX Association, Berkeley, CA, USA, 307-320.
[24] Cosmo G e a 2014 *Journal of Physics: Conference Series 513 022005*
[25] Field L e a CMS@Home: Enabling Volunteer Computing Usage for CMS (CHEP-2015)
[26] MySQL URL: http://www.mysql.com/
[27] DB on Demand URL: http://cern.ch/dbondemand
[28] Drupal content management system: https://www.drupal.org/