


BOINC + CernVM

Ben Segal / CERN

(describing the work of many people at CERN and elsewhere)

**Pre-GDB on Volunteer Computing
CERN, November 11, 2014**

“Real physics”: Challenge issued in 2006



Challenges (1) of “real LHC physics” with BOINC

- **It is not practical to deploy HEP applications on volunteer computers using conventional BOINC porting:**
 - Each application usually has a huge codebase (order of several GBytes) and requires a lot of 3rd party dependencies
 - Applications usually require specific platform/OS combinations to run
 - Applications are updated frequently

Challenges (2) of “real LHC physics” with BOINC

- Experiments would like BOINC volunteer resources to be transparently integrated into existing Grid infrastructures, including their own job schedulers, and/or as “Clouds”
- In contrast to Grid resources, volunteer computing resources are not “managed” (and thus cannot be “trusted”, particularly with Grid credentials)

Development Timeline

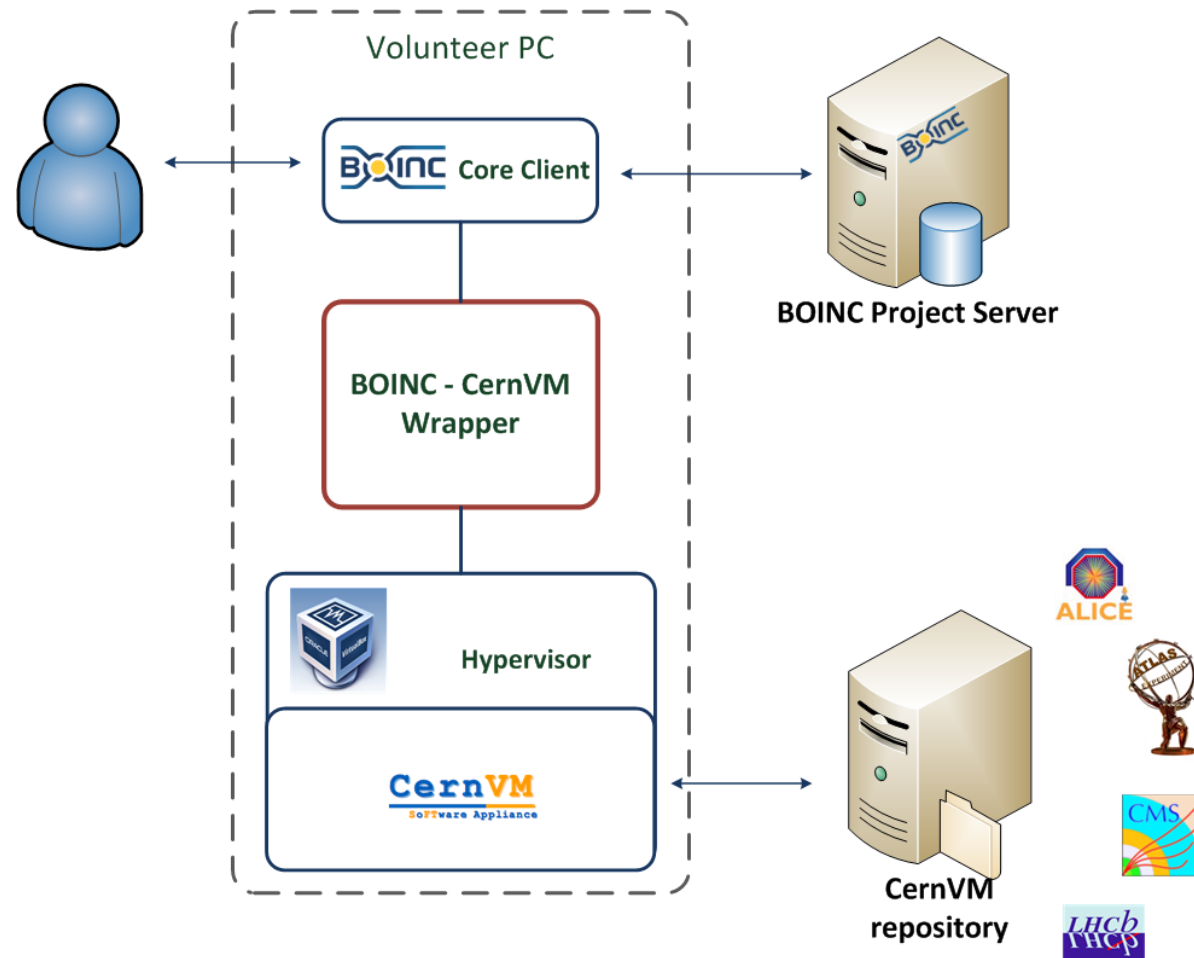
- 2006 – 2007 : VM's added to BOINC as proof of concept:
ATLAS code run under BOINC with VMWare but image too big and no scheduler interface.
- 2008 - 2009 : CernVM interfaced to BOINC, solving image size problem. CernVM VirtualBox BOINC wrapper developed
- 2009 – 2010 : CoPilot added, solving VM credential problem
- 2010 : ATLAS (PanDA), ALICE (ALIEN) and Theory jobs run
- 2011 : Test4Theory production begins

Publications at CHEP, BOINC Workshops, etc.

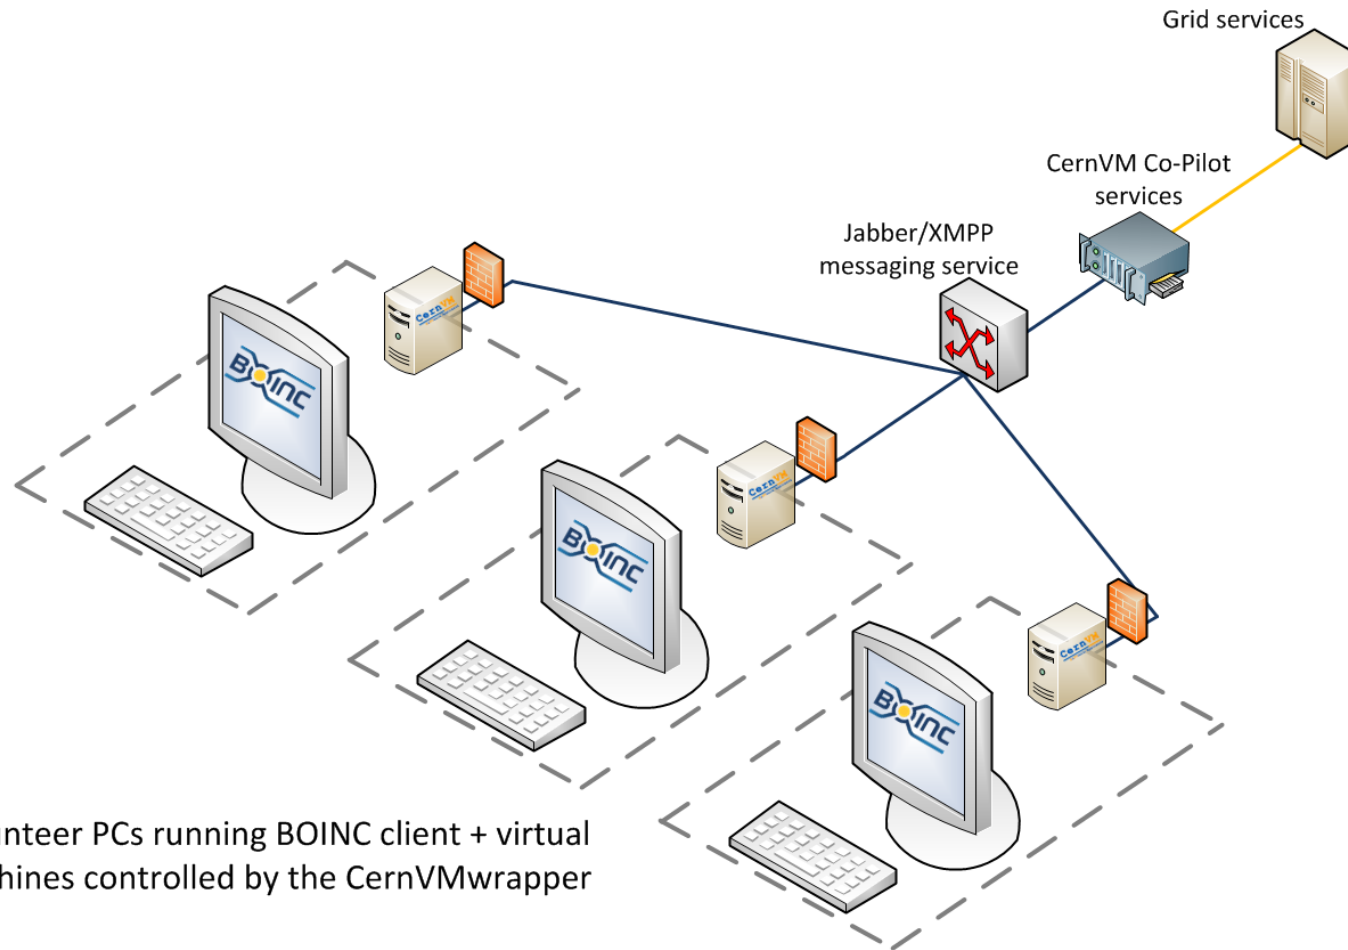
Development effort

- CERN effort : BS plus students, and PH-SFT (CernVM + CoPilot). **Volunteer spirit !**
- Funding from CCC and Shuttleworth Foundation paid for CernVM BOINC wrapper development (Daniel Lombrana)
- « Volunteer Cloud » approach, BOINC only used to start VM
- Web technologies used where possible (viewing VM, etc.)
- Always targeted LHC experiments, but first client was the CERN Theory Group for MCPlots project (« Test4Theory »).

BOINC – CernVM Architecture



BOINC + CernVM + Co-Pilot => Volunteer Cloud



Volunteer PCs running BOINC client + virtual machines controlled by the CernVMwrapper

What has Test4Theory achieved so far?

- About 11000 volunteers and 20000 PC's have contributed since 2011
- Currently about 2000 getting credit with about 3000 active PC's
- About 2500 simultaneously active VM's in a 24 hour period
- About 10000 successful Monte Carlo jobs run per day
- Each job is 100,000 events: so far over 1.5 trillion events run
- Graphic displays show the volunteers details of their CERN job activity, accumulated job statistics, plus offer online tuition
- In minimum-maintenance mode since 2 years

LHC@home 2.0

(Test4Theory project => vLHC@home)

- *We have built a “Volunteer Cloud”*
- *It is in production since 2011 for CERN’s Theory Group running Monte Carlo QCD event generation*
- *The main computing resource for the MCPlots project*
- *Recently reconfigured as “vLHC@home” to accommodate other CERN LHC experiments with similar requirements*