

Early Experience with CERN's Volunteer Cloud

Ben Segal / CERN

and:

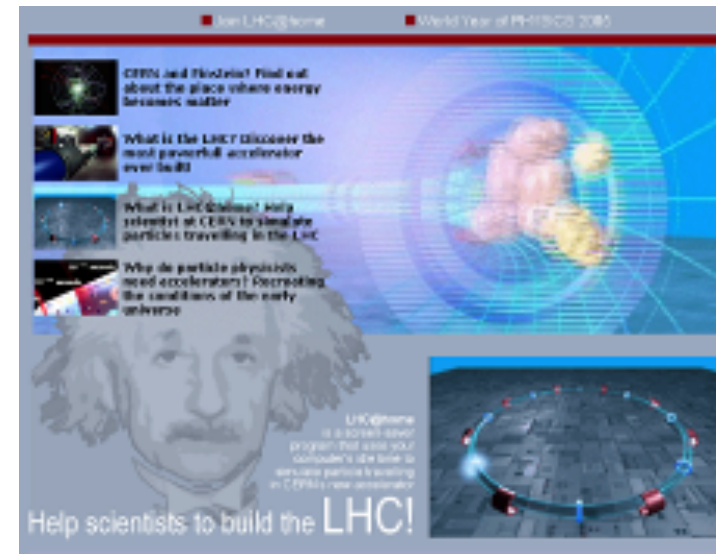
**Daniel Lombraña González / Citizen Cyberscience Centre
Predrag Buncic, Jakob Blomer, Artem Harutyunyan / CERN-PH
Peter Skands, Anton Karneyeu / CERN-TH-LPCC**

7th BOINC Workshop, Hanover

August 18, 2011

LHC@home : (2004 - ...)

- Calculates stability of proton orbits in CERN's new LHC accelerator
- System is nonlinear and unstable so numerically very sensitive. Hard to get identical results on all platforms
- About 40 000 users, 70 000 PC's... over 1500 CPU years of processing
- Objectives: extra CPU power and raising public awareness of CERN and the LHC - both successfully achieved.
- Started as an outreach project for CERN 50th Anniversary 2004; used for Year of Physics (Einstein Year) 2005; **still runs...**



- **USED CLASSIC BOINC PORTING**

BOINC & LHC physics code

Problems with “normal” BOINC used for LHC physics

- 1) A project’s application(s) must be ported to each volunteer platform: most clients run Windows, but CERN runs Scientific Linux and porting to Windows is impractical. Each experiment has its own environment.**
- 2) The physics code changes frequently so apps must be updated often.**
- 3) The projects’ work must be fed into the BOINC server for distribution, and results recovered. Job submission scripts must be developed for this, but CERN physics experiments won’t change their current setups.**
- 4) Job management is primitive in BOINC, whereas physicists want to know where their jobs are and be able to manage them.**

BOINC & Virtualization

The basic approach

- Solve client application porting problems using VM's
- Use "VMwrapper" to link VM's to BOINC core client & server
- Provide sufficient Host <-> Guest-VM communication/control
- No other BOINC changes needed "in principle" ..

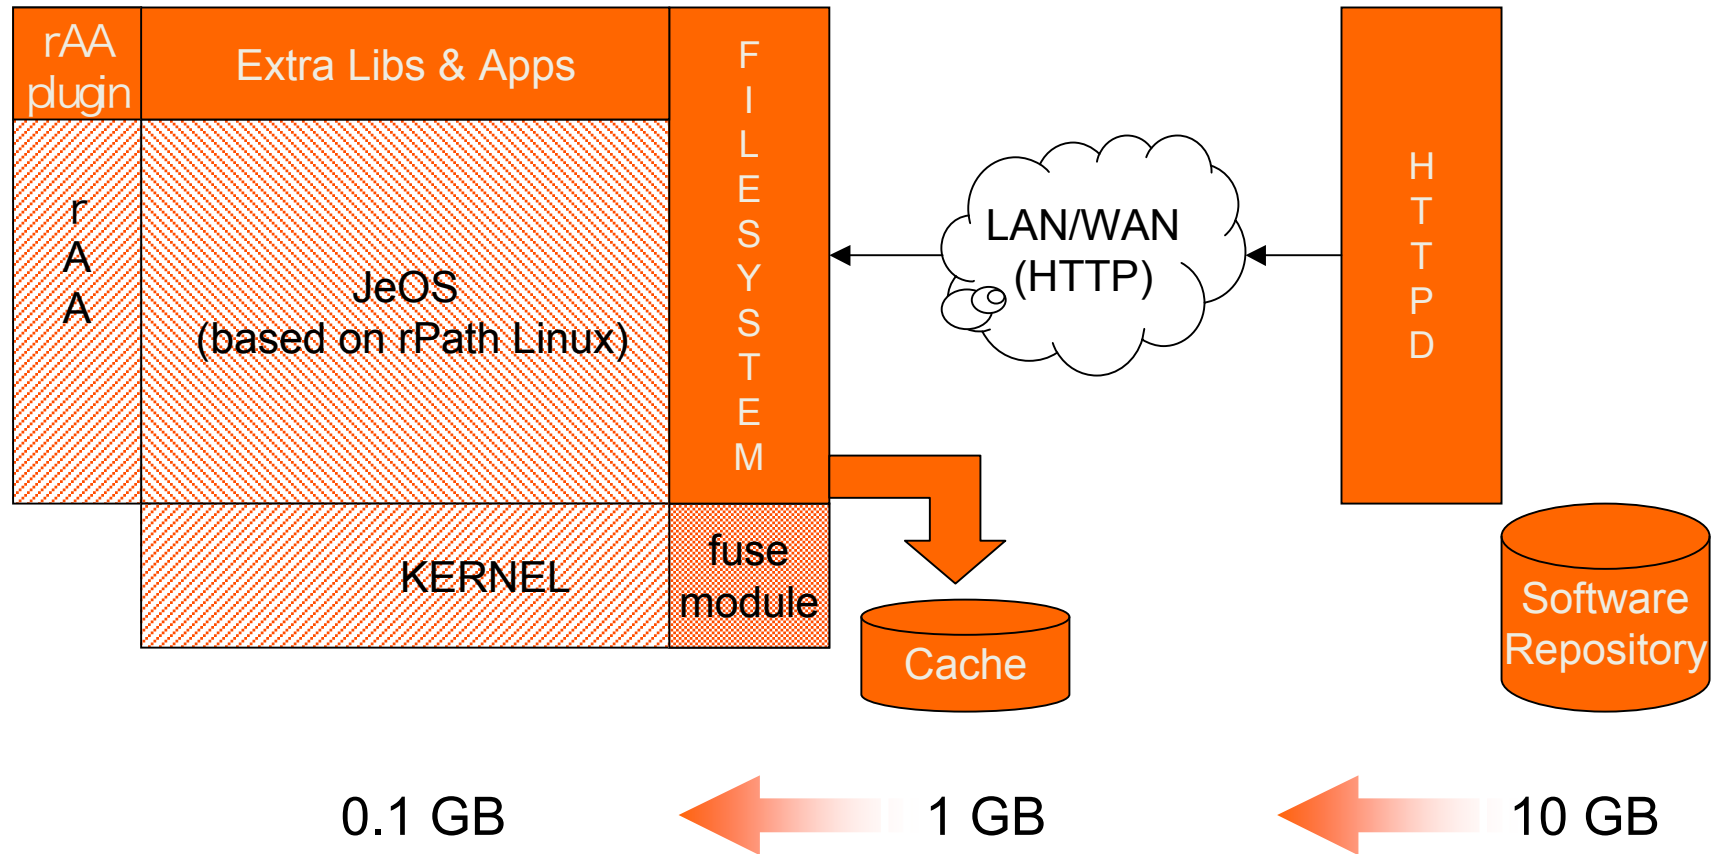
.. and, by using the **CernVM appliance** together with its **Co-Pilot adapter system** ..

.. solve the **image management problems** .. and ..
.. link to CERN's **physics job production interfaces**

LHC@home 2.0


- Started as a small development project in 2006 to explore use of virtual machines in the BOINC environment.
- **Originally** wanted to run BOINC applications “compatibly”: either in a Host CPU or in a Guest VM, keeping BOINC job management and configuration conventions unchanged (like GPU’s today ?).
- **Originally** based on an ambitious cross-platform Python Wrapper and VMController supporting full Host<->Guest VM communication, multiple hypervisors, etc.
- **Now based on a simple C++ BOINC VM_Wrapper to control the VM and hypervisor (currently VirtualBox but extensible to VMware etc).**
- **Uses the CernVM and CoPilot systems to:**
 - **AVOID BOINC PORTING OF ANY CERN APPLICATION**
 - **MINIMIZE VM IMAGE SIZE WITH DYNAMIC CODE UPDATES**
 - **CONNECT TO EXTERNAL JOB MANAGEMENT SYSTEMS**
- **OFFERS A CLOUD INTERFACE TO ALL APPLICATIONS**

CernVM Software Appliance




CernVM File System

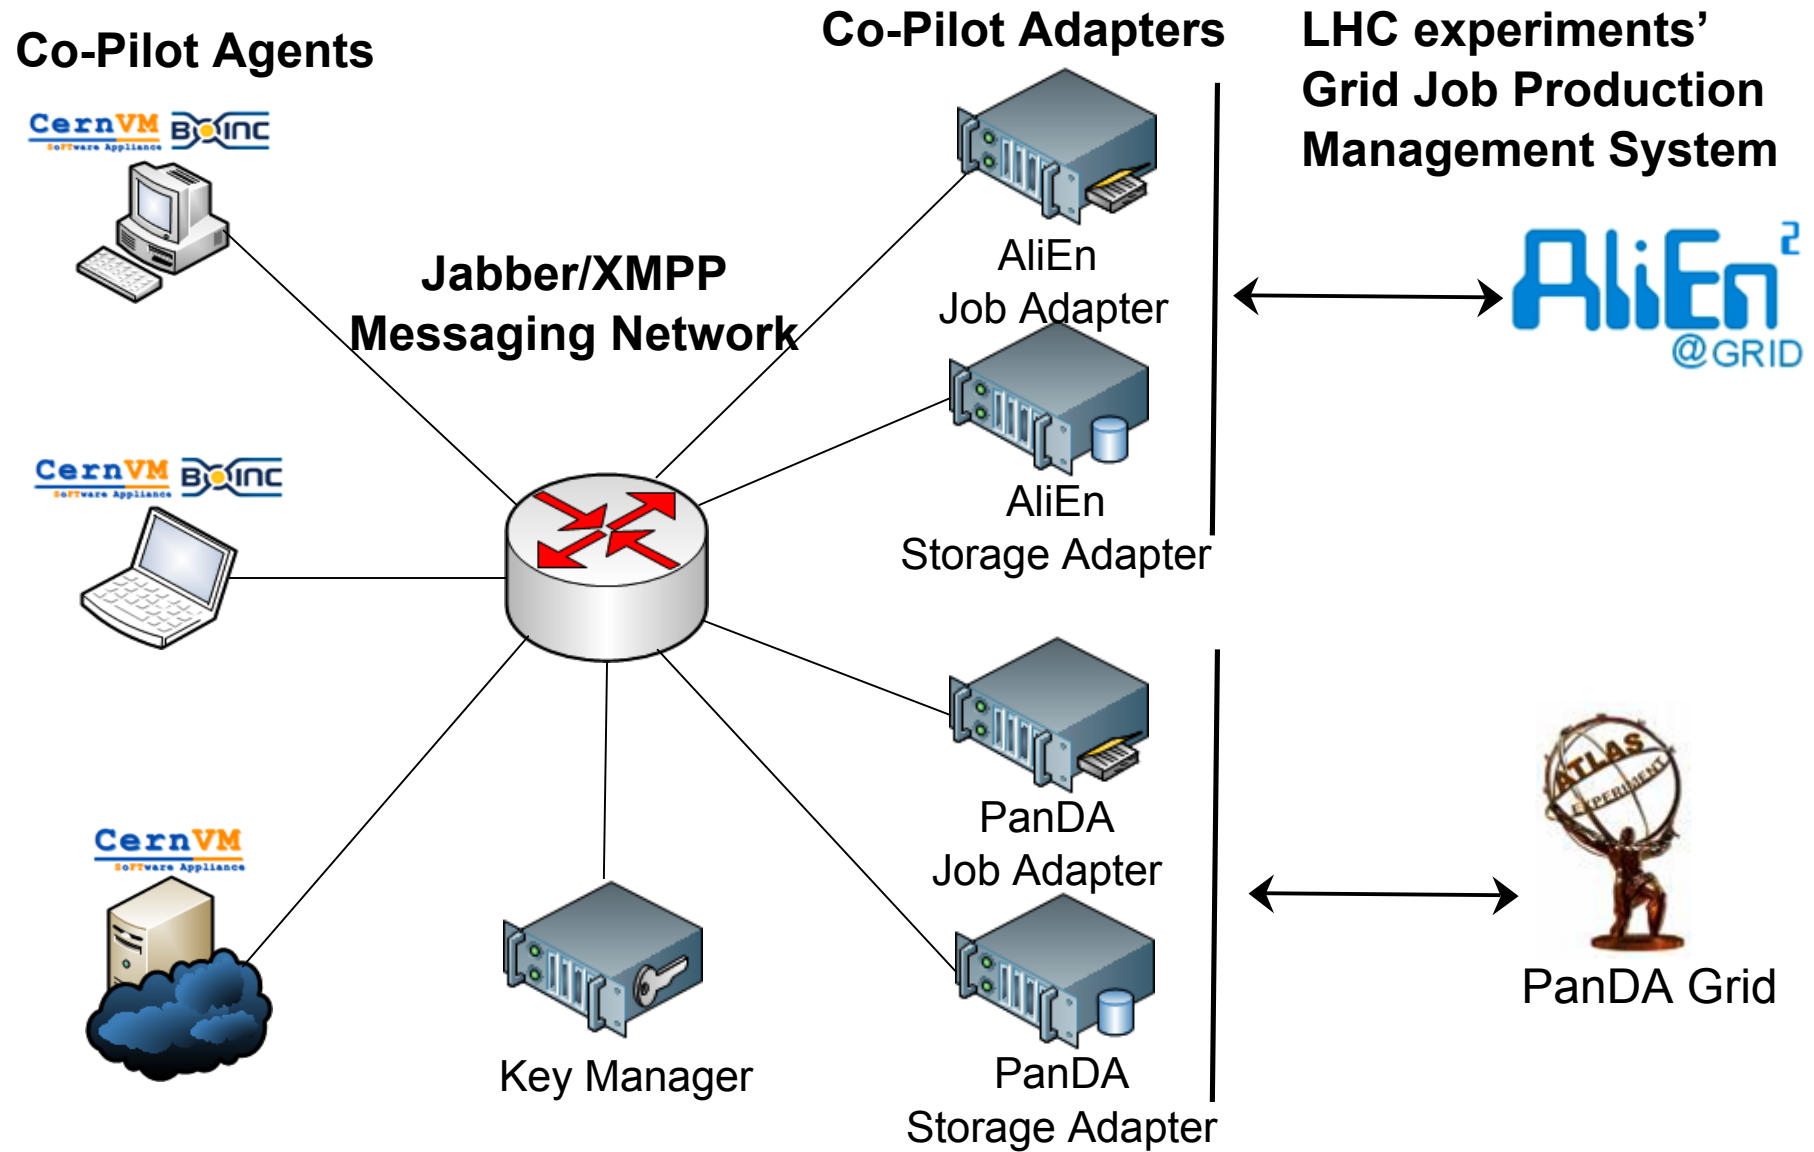
Software Distribution for Virtual Machines



- “Ready to run” binaries, i. e. `/mnt/cvmfs` mirrors destination of `make install`
- Read-only, public files



CernVM Co-Pilot architecture



The use of Jabber/XMPP allows to scale the system in case of high load by just adding new Adapter instances

“Volunteer Cloud”

- **Final Summary:**
 - Solved porting problem to all BOINC client platforms
 - Solved image size and dynamic code update problems
 - Solved job production interface problem
- All done without changing existing BOINC infrastructure (client or server side) ***** But see remarks later about this *****
- All done without changing physicists' code or procedures
- **We have built a “Volunteer Cloud” : now in Beta test**

So what have we achieved so far?

- **About 2000 active volunteers working and getting credit (with about the same number of active PC's)**
- **About 1000 simultaneously active VM's handled OK**
- **About 10000 average CERN jobs per day successfully run (already like TH Group's total computing power at CERN !)**
- **Some instabilities to fix in the job supply system to go further**

- **A lot of feedback and interest from the media and CERN**
- **CERN IT Department just assigned us some manpower (so far, our only CERN support was from PH Department)**
- **A big volunteer demand is waiting to join after we finish debugging the system**

So what have we learned?

- **Poor support of VirtualBox by BOINC gives us many user support problems**
 - Want auto-install of VirtualBox plus correct user privilege settings
(do we need to use the VirtualBox open source package?)
- **Granting credit conveniently and consistently is not yet solved by the standard BOINC methods**
 - Want “partial credits” (like CPDN) to avoid fixed length BOINC tasks
 - Want correct host-normalized scaling (CreditNew?)
 - Want each credit increment proportional to the integrated product:
(BOINC task Δ running time) * (host CPU used by the VirtualBox process)
- **We need collaboration from the BOINC project to solve these problems**

So what have we learned?

- **With our very small dev team (4 part-timers), and about 100 helpful volunteer testers, we debugged the system at Alpha level over a period of 6 months. This was quite easy to do.**
- **The recent transition to Beta level was tough as media buzz sent us 8000 newbies in 4 days !! It also stressed out our software but we survived. We are still under pressure.**
- **Several BOINC volunteer experts have helped us enormously, including the stressful move from Alpha to Beta test level. They appeared by magic and stayed through the toughest times. They are still working with us and we thank them enormously !**
- **The collaboration with the BOINC volunteer community is alive and well !!**

So what are our future plans?

- **Next step: about 10000 active volunteers and 15000 PC's**
- **Up to 10000 simultaneous VM's to be supported**
- **Accept jobs from LHC experiments as well as from the CERN Theory Group**

- **A simple graphic display to show a volunteer's CERN job activity (in final development)**

- **An interactive game to let volunteers "tune" the theory models and fit them to real measured data (in development)**

Thanks to everyone - join us!

- **Project link:**

<http://cern.ch/LHCathome/Physics>

- **Wrapper source code link:**

<http://boinc.berkeley.edu/trac/browser/trunk/boinc/samples/vboxwrapper/cernvm/cernvmwrapper.cpp>