

Discussion material

- Some thoughts on how to expand volunteer computing for HEP
- Extending the base of volunteers
- Desktop grids or other ad-hoc controlled resources
- Proposed volunteer computing strategy
- Areas for more work and collaboration

Some figures

- Number of desktops
 - 2010: 1 Billion Internet-connected PCs, 55% privately owned
 - If 100M people participate:
 - 100 PetaFlops, 1 Exabyte storage (Source: D. Anderson)
- Comparisons
 - GPUGRID: 1.3 petaFLOPS
 - Uses GPU's, about 5000 recent hosts
 - LHC@home classic: 41 TeraFLOPS
 - 18 hours, 156,314 cores, \$33k, theoretical maximum of 1.21 petaFLOPS on Amazon
 - 240 teraFLOPS for \$1279/h on Amazon
 - E.g. Estimated Amazon cost of current Sixtrack load \$215/h

How to get more volunteers?

- **Volunteers do not come by themselves**
 - Invest more in outreach and media
 - Coordinated effort with CERN outreach and press office
 - Further develop software (BOINC or alternative portals) to make running volunteer clouds more attractive
 - *Requires more resources than currently allocated*
- **Cost/benefit analysis**
 - Additional volunteers v.s. cost of additional manpower
 - Cost of investment in volunteer computing vs more servers in the data centre
 - What fraction of HEP computing can benefit?

Desktop grids

- Use spare capacity at CERN and HEP labs
 - Low hanging fruit
- Currently use of BOINC on desktops is on a voluntary basis
- Other sites operate large desktop grids with mandatory BOINC client installations
 - Automate deployment (e.g. CMF for NICE)
 - Standard run-time user (credit not a concern)
 - Environmental policy: Switch off all idle PCs that do not run BOINC
- Example: University of Westminster in the UK (talk of Dario Ferrer)

Desktop grids - links

- International Desktop Grid Federation: <http://desktopgridfederation.org/>
- Crowd Computing: <http://crowdcomputing.eu/>

Power consumption aspects

- Cost of opportunistic capacity versus extra servers in the data centre?
 - Would be interesting to measure power consumption of a set of jobs on a typical desktop and on a recent server
- Others have also looked at the subject:
 - Low energy BOINC: <http://low-energy-boinc.cslparis.fr/>
 - Talks by Peter Hanappe at the 2 last BOINC workshops
 - BOINC heat and energy considerations (user documentation)

Volunteer vs other cloud?

- The Volunteer Computing model is good for desktops and some small Tier-2's
- However, other volunteer resources are uncontrolled and untrusted, so access to Grid storage is problematic
- Tier 2 centers with local storage may be better off using a cloud framework like VAC, running VMs with credentials to take advantage of Grid storage elements
- XtremWeb is another alternative

A set of recommendations for which system for which resource may be useful in the LCG context

Volunteer computing strategy

- Volunteer Computing should be seen as an extension of the cloud model of the experiments
 - BOINC used as a vehicle to distribute CernVM images
 - This is also a good outreach channel, currently up to each experiment to exploit
- Desktop grids and some volunteer resources come with minimal additional cost
- Let's use these resources before turning to commercial cloud providers

VC – areas to work on

- We have seen promising developments around Volunteer Computing
- Areas for more work within HEP
 - CernVM image / CernVMfs, pre-cached images etc
 - Job piloting and handling standardisation (data bridge etc)
 - Learn from active projects and share!