

Update on CERN infrastructure services

Nils Høimyr, on behalf of CERN IT/PES/IS

- Update on existing services
 - TWiki service status
 - Update on Version Control Services (SVN, CVS)
 - Engineering Service infrastructure (on linux)
- 2 new services:
 - Issue Tracking Service (JIRA)
 - BOINC volunteer computing project support

- **TWiki** offered as a collaborative Wiki solution since 2004 and widely used by LHC experiments and CERN users.
 - 124k topics (web pages)
 - 2 M monthly topic views
 - 55 k monthly topic updates
- TWiki 4.3.2 in production, with many plugins and extensions for CERN E-groups authorization.
- Upgrade to TWiki 5.1.1 planned for May.
 - Many **usability enhancements** and application platform changes.
 - Undergoing final tests and validation now.
- **Service evolution plans:**
 - Further integration with CERN SSO and web infrastructure.
 - Keeping an eye on alternatives, such as Foswiki.

- Central SVN service
 - 1014 hosted projects
- Legacy CVS service
 - 14 projects remaining, the most active is CMS Offline.
 - Plan to stop the CVS service during LS1 in 2013.
- Focus on consolidation of SVN
 - Reliability & performance
 - Running SVN 1.6, upgrade to 1.7.x planned
 - Back-end storage optimization (AFS parameters, NFS alternative)
 - Trac db moving from Sqlite to MySQL
 - Integration with Issue Tracking Service
 - Pilot Git setup for the autumn
 - Focus on security and scalability

<http://cern.ch/svn>

- Linux licence servers (Flexlm cluster, Mathlm, Nag..) all migrated to virtual infrastructure.
 - Some issues with Flexlm triads and the proprietary licence server extension for Ansys.
 - Flexnet monitoring application suite from vendor Flexera not very robust, yet helpful as a tool to manage licences.
- Engineering batch resource (part of LXBATCH)
 - 40 nodes with 10Gb Ethernet interconnects.
 - 20 first nodes with “Net Effect” low latency cards supporting the OFED stack (Infiniband protocol). 20 nodes added later.
 - Used for MPI applications; CFD (Ansys/Fluent, OpenFoam), Beam Dynamics (ORBIT), Field Calculations (GdfidL) etc.
 - AFS used for input & results, local disk for task i/o where possible.

- Issue tracking service based on JIRA from Atlassian
 - Requests from users in LHC experiments for a JIRA service
 - LCG Savannah portal old and needs a replacement
 - Trac too close to SVN and insufficient to meet all requirements
 - Many groups with private JIRA installations at CERN
- Central hosted JIRA Service, integration with :
 - SSO, eGroups, CERN Services
 - Version Control Systems (SVN, Git(?))
- Service starting up
 - Pilot projects running successfully
 - Offer an integrated service for new projects
 - Migrate existing JIRA installations to the central service
 - <http://cern.ch/its>

- JIRA is a commercial software from Atlassian, used by :
Apache, facebook, eBay, NASA, twitter, ...



- Issue tracking tool, highly customizable (plugins, workflows, ...), friendly UI
- Other Atlassian tools of interest :
 - GreenHopper, Bamboo, Confluence, Crucible, Fisheye, ...

- Supported by Atlassian :
 - https load balancing + Tomcat6 + Oracle
- Virtual Machines on SLC6
- Architecture :
 - 1 VM = 1 instance
 - Central instance for common needs (most projects)
 - Separate instances for special needs (no administration by ITS)
- Quatorrized
- Out-of-the-box SVN and Git plugin

- CERN SSO :
 - Shibboleth (mod_shib for Apache) protecting login page (anonymous access)
 - Custom Authenticator in Jira (java class that “filter” login requests)
- CERN eGroups
 - Jira Gadget : mapping eGroups to Jira Groups
- Service-now : basic link
 - Custom plugin

CERN eGroups mapper

eGroup

Select Jira group(s)

jira-users
 pf-developers
 root-developers
 simulation-developers
 spi-developers

[Map groups](#)

List of mapped groups

eGroup	Jira Group	Unmap selected
it-service-its-admin	jira-administrators	<input type="checkbox"/>


IMPORTANT

Mapping one eGroup to one (or more) Jira Groups means that this groups will be synchronized on login.

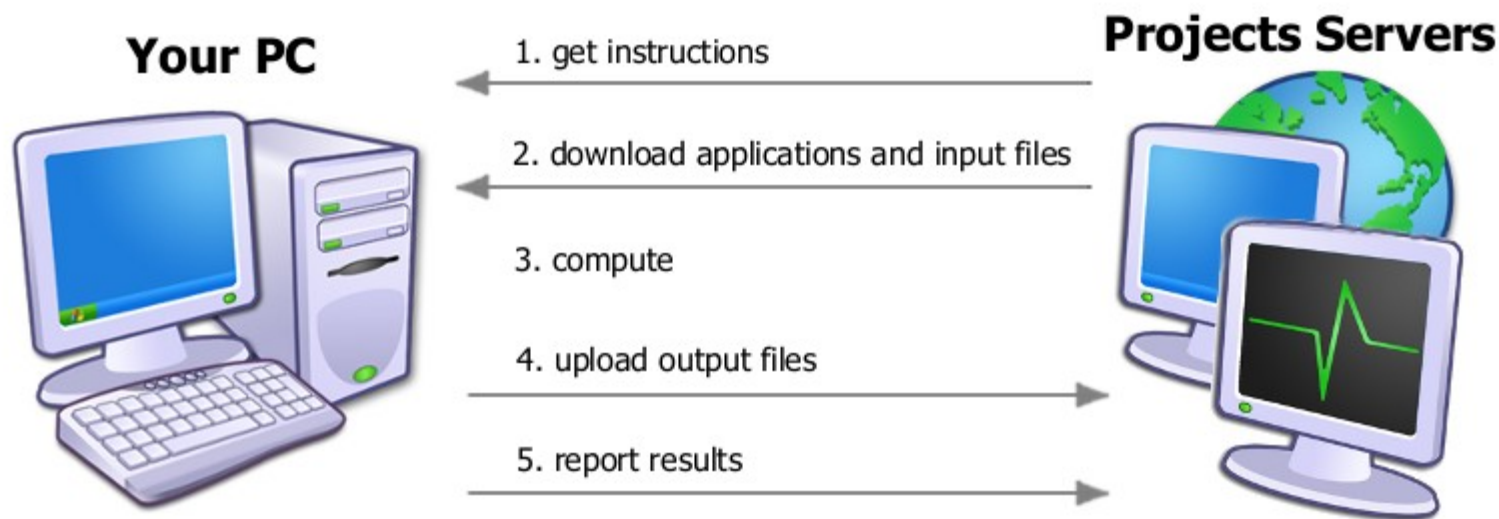
That's to say, if a user is member of one of the eGroup that has been mapped, he will be added to all the mapped Jira Groups.

If somehow, he is member of the Jira Group (added manually) but not of the mapped eGroup, he'll be removed from the Jira Group.

- “**Berkeley Open Infrastructure for Network Computing**”
- **Software platform for distributed computing using volunteered computer resources**
- <http://boinc.berkeley.edu>
- **Uses a volunteer PC's unused CPU cycles to analyse scientific data**
- **Client-server architecture**
- **Free and Open-source**
- **Example projects**
 - SETI@Home
 - Climateprediction.net
 - Einstein@Home



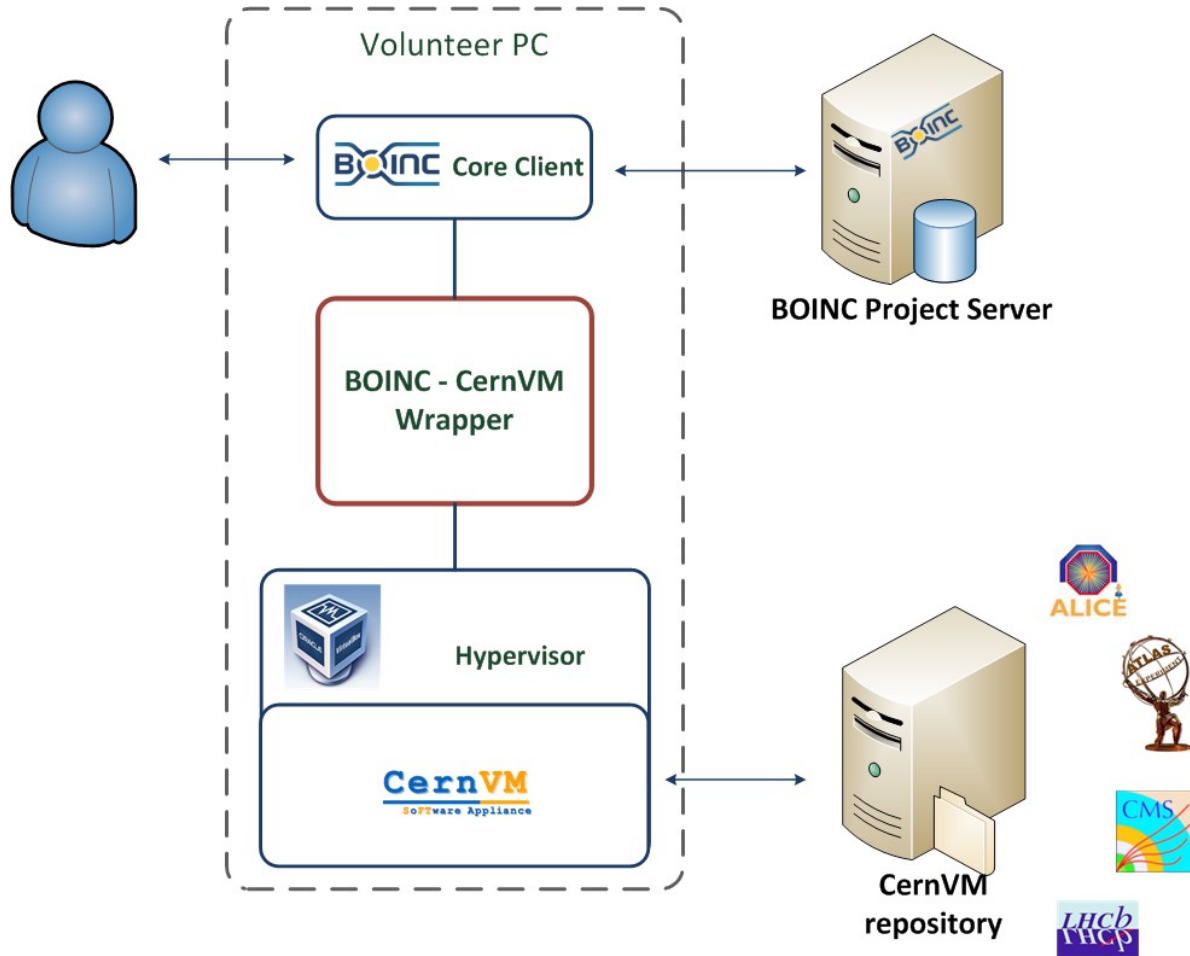
- Download and run BOINC software
- Choose a project
- Enter an email address and password in the BOINC Manager.
- That's it, your done!



- Calculates stability of proton orbits in the LHC accelerator.
 - SixTrack Fortran program, simulating particle trajectories.
- About 60 000 users, 100 000 PC's... over 3000 CPU years of processing.
- Client runs on Linux and Windows platforms
 - Classic BOINC architecture
- Started as an outreach project for CERN 50th Anniversary 2004; used for Year of Physics (Einstein Year) 2005.
- Objectives: extra CPU power and raising public awareness of CERN and the LHC - both successfully achieved.
- Ran (intermittently) since, with server in QMUL since 2007.
- Server migrated back to CERN in September 2011

- Project done by Ben Segal, a number of short term students and the CernVM team of PH/SFT
- Application case from the Theory group
 - Theoretical fitting of LHC data using Monte Carlo simulation based on Standard model
- Uses a virtual machine on the volunteer computers
 - User installs Virtual Box
 - User installs the BOINC client
 - Attach to project and ready to go with application on CernVM
 - The BOINC client downloads a VM-wrapper, that gets the image and the job to run on the VM
 - Potentially wide range of physics applications

BOINC – CernVM Architecture



- BOINC server cluster
 - VM infrastructure for each server
- BOINC server application
 - Configuration, monitoring
 - MySQL database server back-end
 - Portal with user forums etc
 - In progress:
 - Hosting of Co-Pilot, CernVmFS etc
- *Out of scope: BOINC user applications*
 - *Porting of applications to BOINC*
 - *Communication with users about scientific projects*

- Many thanks to my colleagues:
 - David Asbury (VCS)
 - Gautam Botrel (JIRA)
 - Owen Dyckhoff (student, TWiki)
 - Alvaro Gonzalez (ITS, Lic, VCS)
 - Pete Jones (TWiki, BOINC)
 - Georgios Koloventzos (VCS)
 - The rest of the PES group
 - PH & BE colleagues for LHC@home



Questions?

Nils.Hoimyr (at) cern.ch