CERN IT Department

The (B)right future of software installation

S. Bagnasco, L. Betev, F. Carminati, F. Furano,
C. Grigoras, A. Grigoras, P. Mendez Lorenzo,
A. Peters, P. Saiz

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

WHAT IS THE BEST METHOD TO INSTALL SOFTWARE?

- What to install
- Where/when to do it
- How should it be done

What to install

- Two types of packages:
  - Experiment software:
    ROOT (An Object-Oriented Data Analysis Framework), GEANT3, AliRoot
  - GRID interface: AliEn 2 @GRID
- Requirements
  - Self-contained
  - Platform dependent
  - User space
  - Size does matter

Build & Test systems

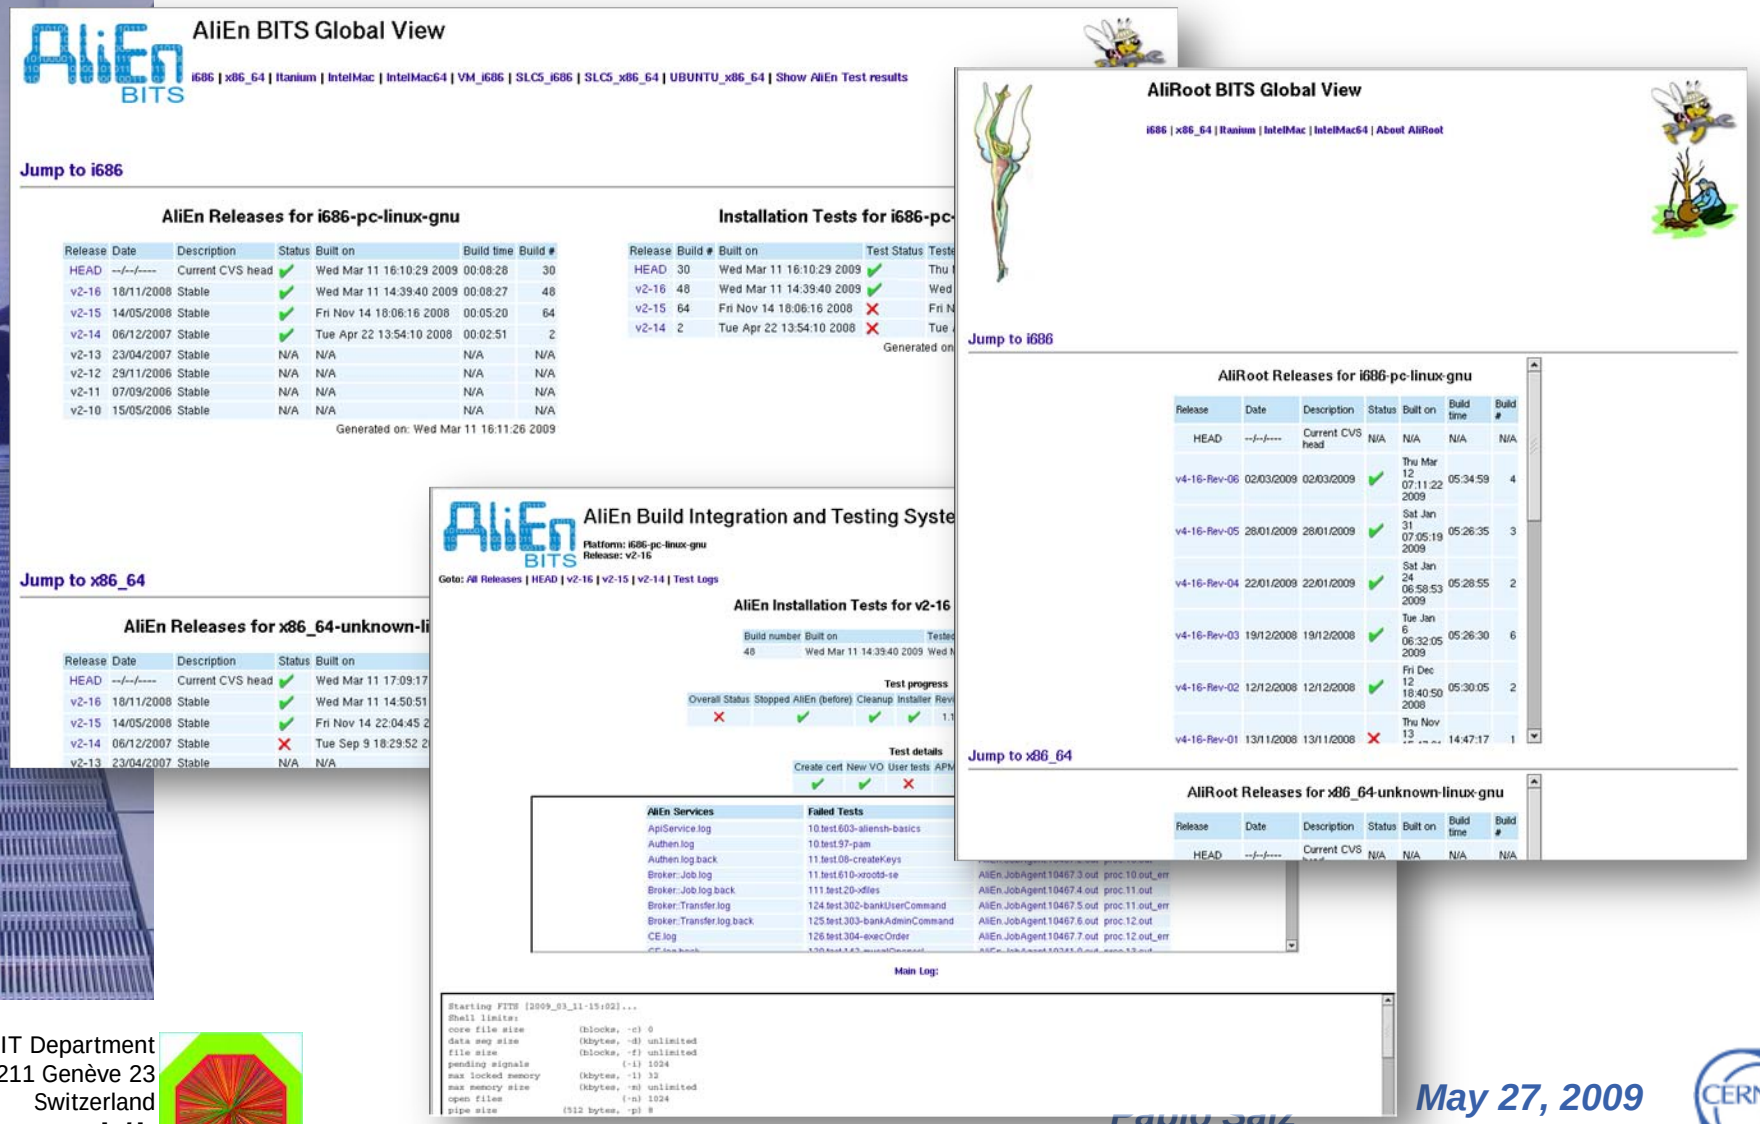

More than 200 individual packages

CERN IT Department

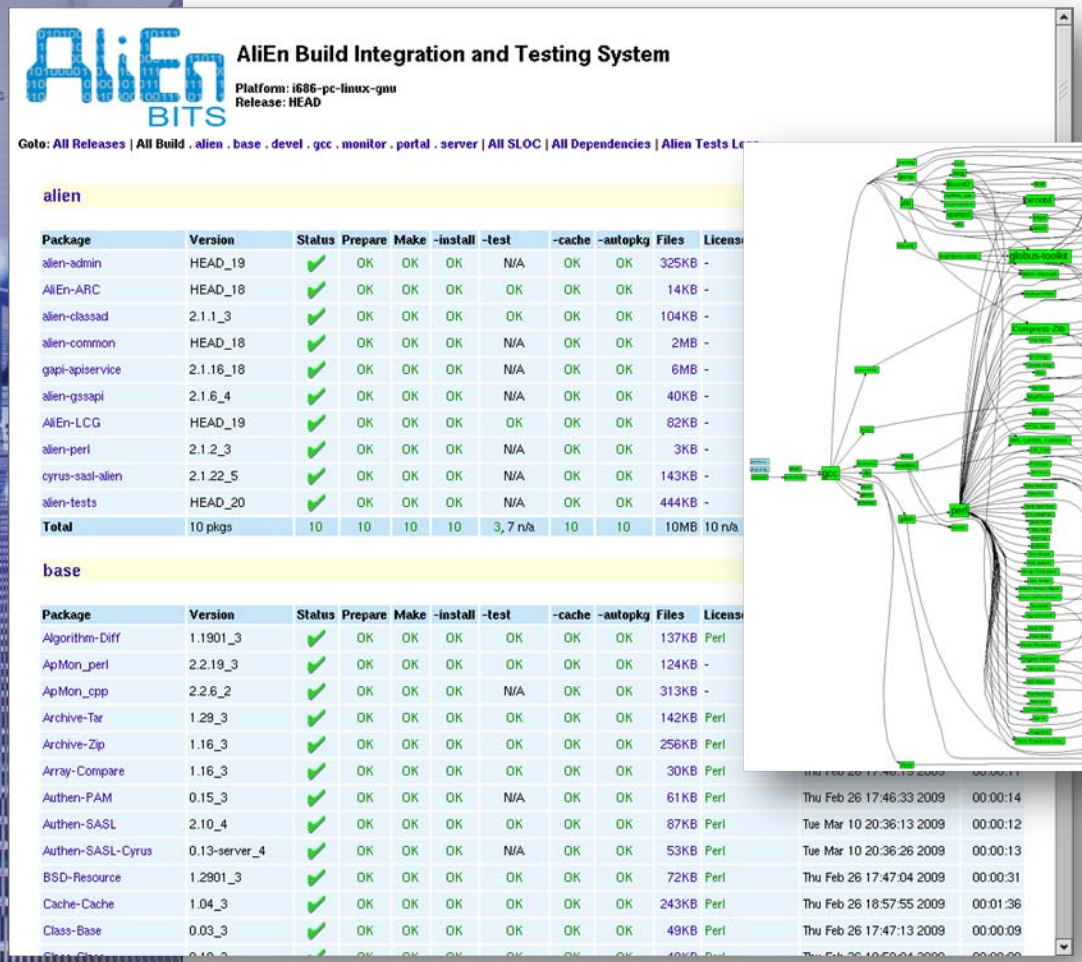

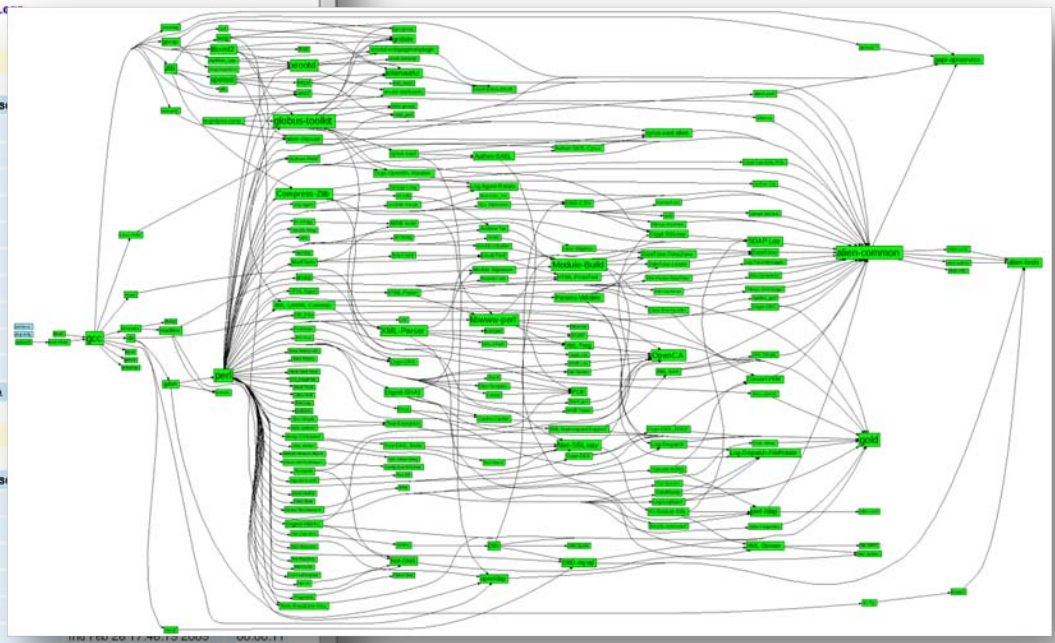

Easy to add new packages, move to new platform, get files from different sources, apply patches, define dependencies ...

Packaging & size

- Combine all the required grid packages into distributions
  - Full installation: 155 MB, mysql, ldap, perl, java...
  - VO-box: 122 MB, monitor, perl, interfaces,
  - User: 55 MB, API client, gsoap, xrootd
  - Worker node: 34 MB, min perl, openssl, xrootd
- Experiment software:
  - AliRoot: 160 MB
  - ROOT: 60 MB
  - GEANT3: 25MB

300 MB to run jobs

Where and when to do it

CERN IT Department

- Current scenario
  - One installation (per platform) per site
  - On a shared area
    - Might be slow
    - Security risk
  - AliEn installation:
    - New sites: running a script
    - Updating a site: triggered by a VO admin
  - Experiment software:
    - Install on demand according to the jobs
    - Before submitting JobAgents

Do we need it?

Can we do something better?

- Automatic installation on every worker node

- Automatic
- Self-contained
- User space
- Small software (300 MB)
- Job Agents (can run more that one small job)

How to do it

- Landing on a virgin worker node
- Sending a small script that:
  - Finds a scratch area in a local disk
  - Installs latest version of AliEn
    - Simple requirements: wget
  - Starts a Job Agent
    - JobAgent will install required experiment software
  - After finishing, removes everything
    - Or leaves it for the next agent... (still to be evaluated)

But how to get the files…

Let's go mainstream!!

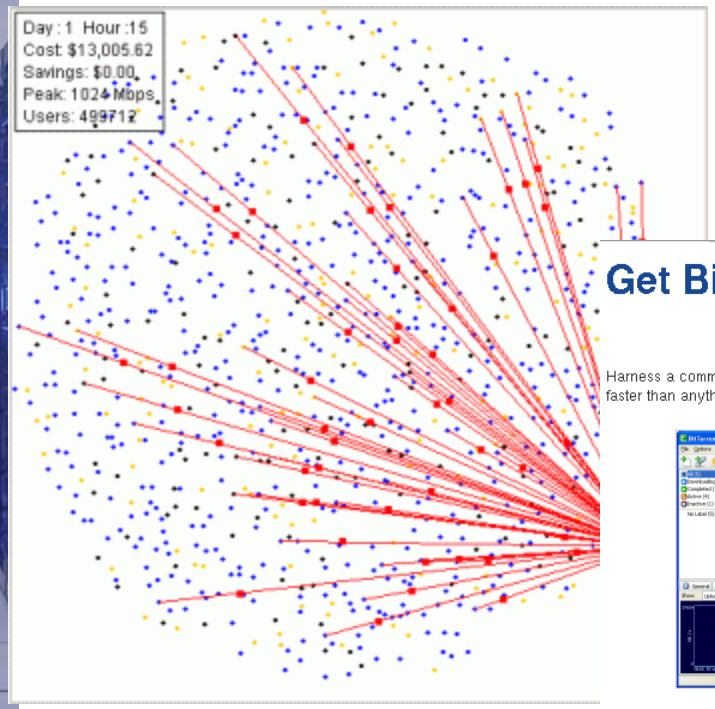

More than 150 million users!!

Get BitTorrent

Harness a community of over 150 million users to deliver files to your PC faster than anything else.

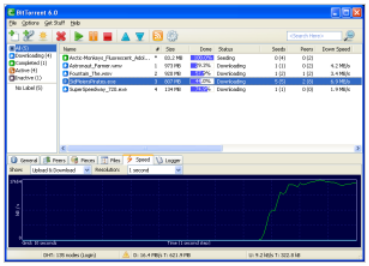

The new BitTorrent 6 for Windows brings together BitTorrent's proven expertise in networking protocols with µTorrent's efficient implementation and compelling UI to create a better BitTorrent client. For questions about the BitTorrent client, take a look at the FAQ in our support center, or visit the client forums.

Advertisement

Get BitTorrent

Download other versions:

- BitTorrent for Windows
- BitTorrent for MacOS X
- Linux, Source Code, and Older Versions

If you encounter any problems, please report them in the Client Forums.

BitTorrent 6 Features

- Lightweight client
- Local peer discovery
- Configurable bandwidth scheduler
- Global and per-torrent speed limiting
- RSS Downloader
- Always Spyware-Free

http://bittorrent.com

Torrent technology

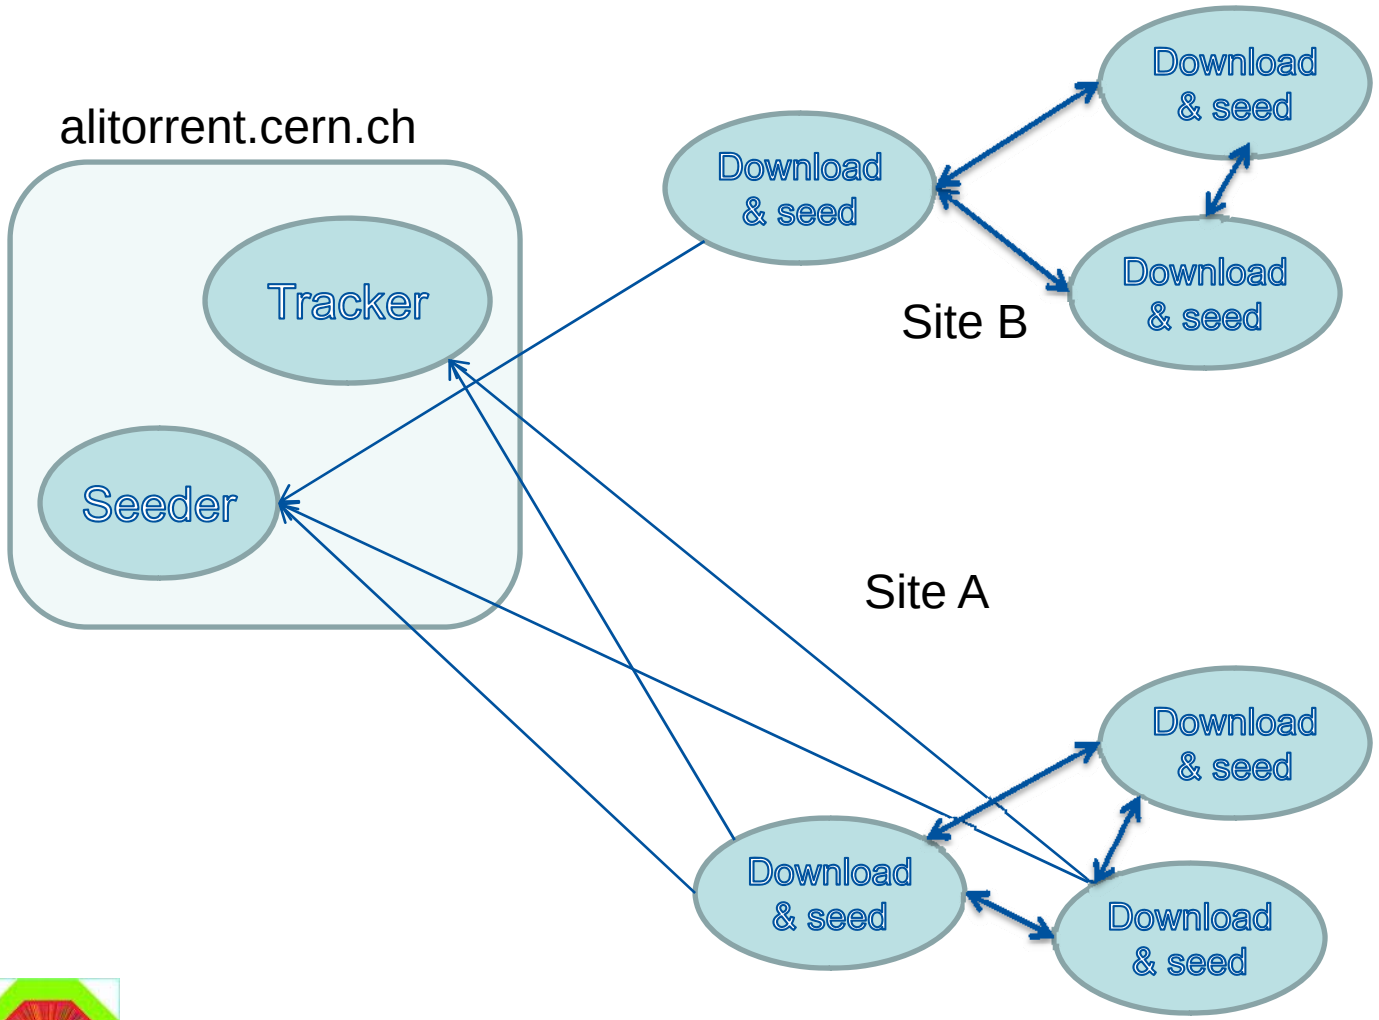

Transfer files

- Torrent files created from the build system
- One seeder at CERN
  - Standard tracker and seeder.
- Get torrent client from ALICE web server
  - Aria2c
- Download the files and install them
- Seed the files while the job runs

Load on the server

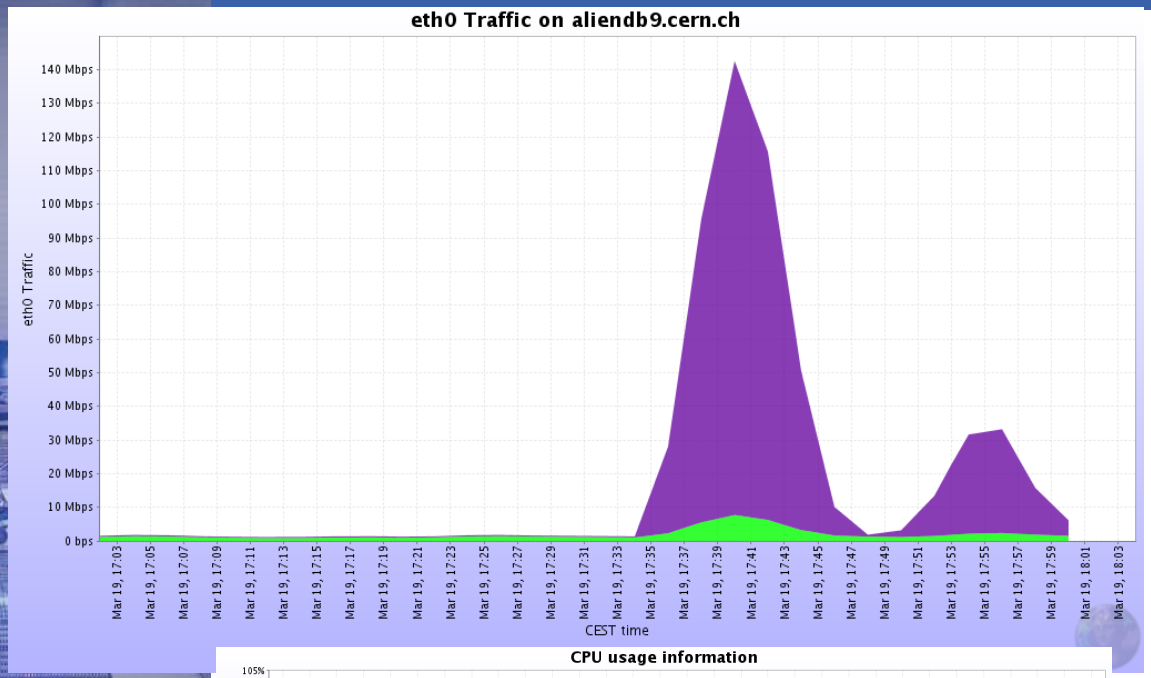

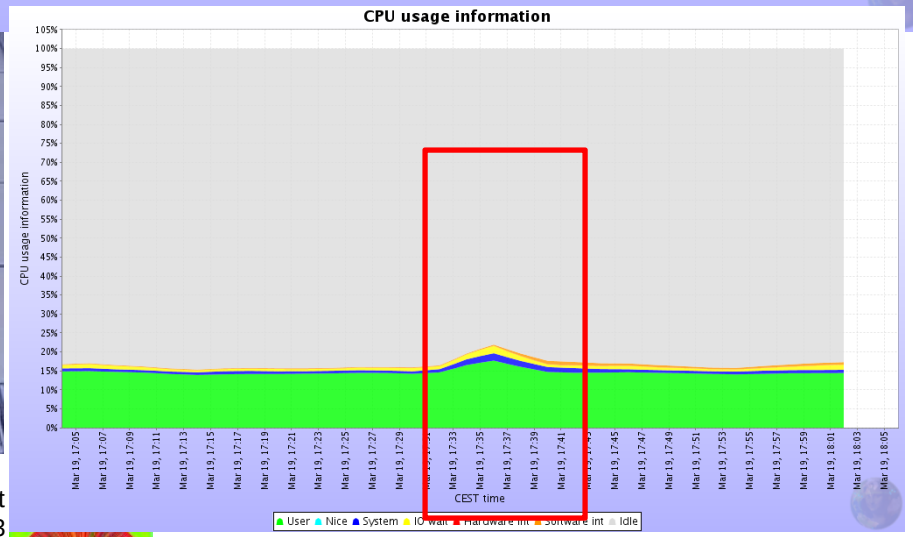

600 clients at CERN:
600x300MB=180GB

Downloaded from the server: 7GB

Fetched from other wn: 173GB

Negligible CPU load

Network load limited by conf

Network load does not increase with more clients

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Security

- Working in collaboration with CERN Security team
- Peer-peer is allowed for professional usage
- Torrent files have checksums to detect corrupted files/wrong files
- Only VO admin can register files in tracker
- Signing the packages

Security concerns

- What if a wn modifies a file?
  - Torrent technology detects and prevents it

    - Each chunk is individually signed
    - The whole file can be also signed with a certificate
- Can any user transfer files?
  - NO. Only VO admin can register files in the tracker
- Do site admins have to open ports on wn?
  - Only within the site

Next steps

- Involve more sites
  - Do we need more seeders?
    - One main seeder should be enough
    - Otherwise, add seeder to (some) voboxes, or to other machines
- Larger scale
  - We have verified up to O(1000) jobs
- Reuse downloaded files on the same wn

Software installation

- Build and test system for multiple platforms
  - Small, self-contained packages
- GRID software updates triggered by VO admin
- Automatic installation on every worker node
- Using torrent technology to transfer files
  - Already tested by millions of users
  - Reduce inter-site transfers
  - No extra requirements on worker nodes
- No need for a shared area
- Same algorithm could be used by other VOs

Great! How to use it

- To install alien:
  - wget http://alien.cern.ch/alien-torrent-installer
- Any file in the catalogue registered with torrent://…
  - This can be used for packages
  - Or for any other data…

Example


psaiz@pcegee02:~ - Shell No. 5 - Konsole
Session Edit View Bookmarks Settings Help
[aliendb06c.cern.ch:3307] /alice/cern.ch/user/p/psaiz/packages/ROOT/v5-22-00p/ > whereis Linux-x86 64
May 26 18:46:38 info The file packages/ROOT/v5-22-00p/source is in
SE => no_se pfn =>torrent://alitorrent.cern.ch/torrents/root_v5-22-00.Linux-x86_64-2.6-gnu-3.4.6.tar.gz.t
orrent
[aliendb06c.cern.ch:3307] /alice/cern.ch/user/p/psaiz/packages/ROOT/v5-22-00p/ > get Linux-x86_64
May 26 18:47:30 info The file packages/ROOT/v5-22-00p/Linux-x86_64 is in
SE => no_se pfn =>torrent://alitorrent.cern.ch/torrents/root_v5-22-00.Linux-x86_64-2.6-gnu-3.4.6.tar.gz.t
orrent
May 26 18:47:30 info Getting a security envelope..
May 26 18:47:30 info In getfile, with no_se
May 26 18:47:30 info Getting the file torrent://alitorrent.cern.ch/torrents/root_v5-22-00.Linux-x86_64-2.6-gnu-3.4.6.tar
.gz.torrent
May 26 18:47:37 info Getting a torrent file
May 26 18:47:37 info Getting the file http://alitorrent.cern.ch/torrents/root_v5-22-0
.torrent, using the IP 137.138.41.23
2009-05-26 18:48:15.136477 NOTICE - Your share ratio was 0.0, uploaded/downloaded=0B/10
Download Results:
gid|stat|avg speed |path/URI
1| OK| 1.7MiB/s|[MEMORY]root_v5-22-00.Linux-x86_64-2.6-gnu-3.4.6.tar.gz.torrent
2| OK| 3.7MiB/s|/tmp/psaiz/cache/root_v5-22-00.Linux-x86_64-2.6-gnu-3.4.6.tar.gz.t
2-00.Linux-x86_64-2.6-gnu-3.4.6
Status Legend:
(OK):download completed.
and 1
May 26 18:48:15 info Here we should check the signature
May 26 18:48:15 info And move the file
May 26 18:48:15 info Finally, start the seeder
May 26 18:48:15 info Returning
May 26 18:48:16 info Everything worked and got /tmp/psaiz/cache/root_v5-22-00.Linux-x86_64-2.6-gnu-3.4.6.tar.gz.torrent.
45161243356457
[aliendb06c.cern.ch:3307] /alice/cern.ch/user/p/psaiz/packages/ROOT/v5-22-00p/ >
Shell
Shell No. 2
Shell No. 3
Shell No. 4
Shell No. 5
Shell No. 6
Shell No. 7
Shell No. 8
Shell No. 9
Please, try it now!! File downloaded using torrent

Please, try it out...

- Login to your vobox
  - ssh ...
- Make sure that you have a valid proxy
  - alien proxy-info
- Connect to alien:
  - alien
- Go to my directory:
  - cd /alien/cern.ch/user/p/psaiz/packages/ROOT
- Check where the file is:
  - whereis v5-22-00p/Linux-x86_64
- Get the file:
  - Get v5-22-00p/Linux-x86_64

French questionnaire

CERN IT Department

Alice T1/T2 Tutorial
BitTorrent Usage
Questions & Concerns from the IN2P3 federation

- Disclaimer : this is a summary from exchanges in the LCG-FR discussion list for technical issues. I did my best to accurately report the questions. Answers are what I understood from Pablo and Costin's replies. JM.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Pablo Saiz

May 27, 2009

Jean-Michel questionnaire (I)

- Cache on every worker node?
  - For the time being, no cache. Everything deleted after the job finishes.
- If a worker node runs 8 jobs simultaneously…
  - … there will be 8 different installations
- It seems stupid to install the software 8 times if the machine runs 8 Alice jobs.
  - If this is a bottleneck, we will reuse the files. At the moment, we want to stress the download
- Will the cleaning be done if a job crashes ?
  - Yes, it will. Cleaning done automatically by CREAM, LCG-CE ...
- Compiling on every worker node?
  - No! Distributing binary code

Jean-Michel questionnaire (II)

- Could one user overwrite software of another user?
  - No! Each job runs inside an LCG sandbox
- If a file is corrupted, it will be very difficult to find where it came from
  - Torrent does corruption detecion
- How will local clients prefer using local seeder instead of all going to the CERN seeder ?
  - Torrent protocol. It favours closer seeders with high connectivity
- If the central machine is hacked, all wn are exposed
  - Yes, but:
    - Machine well under control
    - Are there other options? SE hacked? Build systems hacked? CVS server hacked?

Jean-Michel questionnaire (III)

- Single Point of Failure: CERN tracker
  - No. We could register alternative PFN
- The network traffic between worker nodes will increase and in an uncontrolled manner.
  - The network traffic from your nfs will decrease
- Do we have to open ports on worker nodes?
  - No!
- It would be nice to have a multi-vo solution…
  - This setup can be used by other VOs: convince them!!