



CernVM - Virtual Machine for LHC Experiments

Predrag Buncic (CERN/PH-SFT)

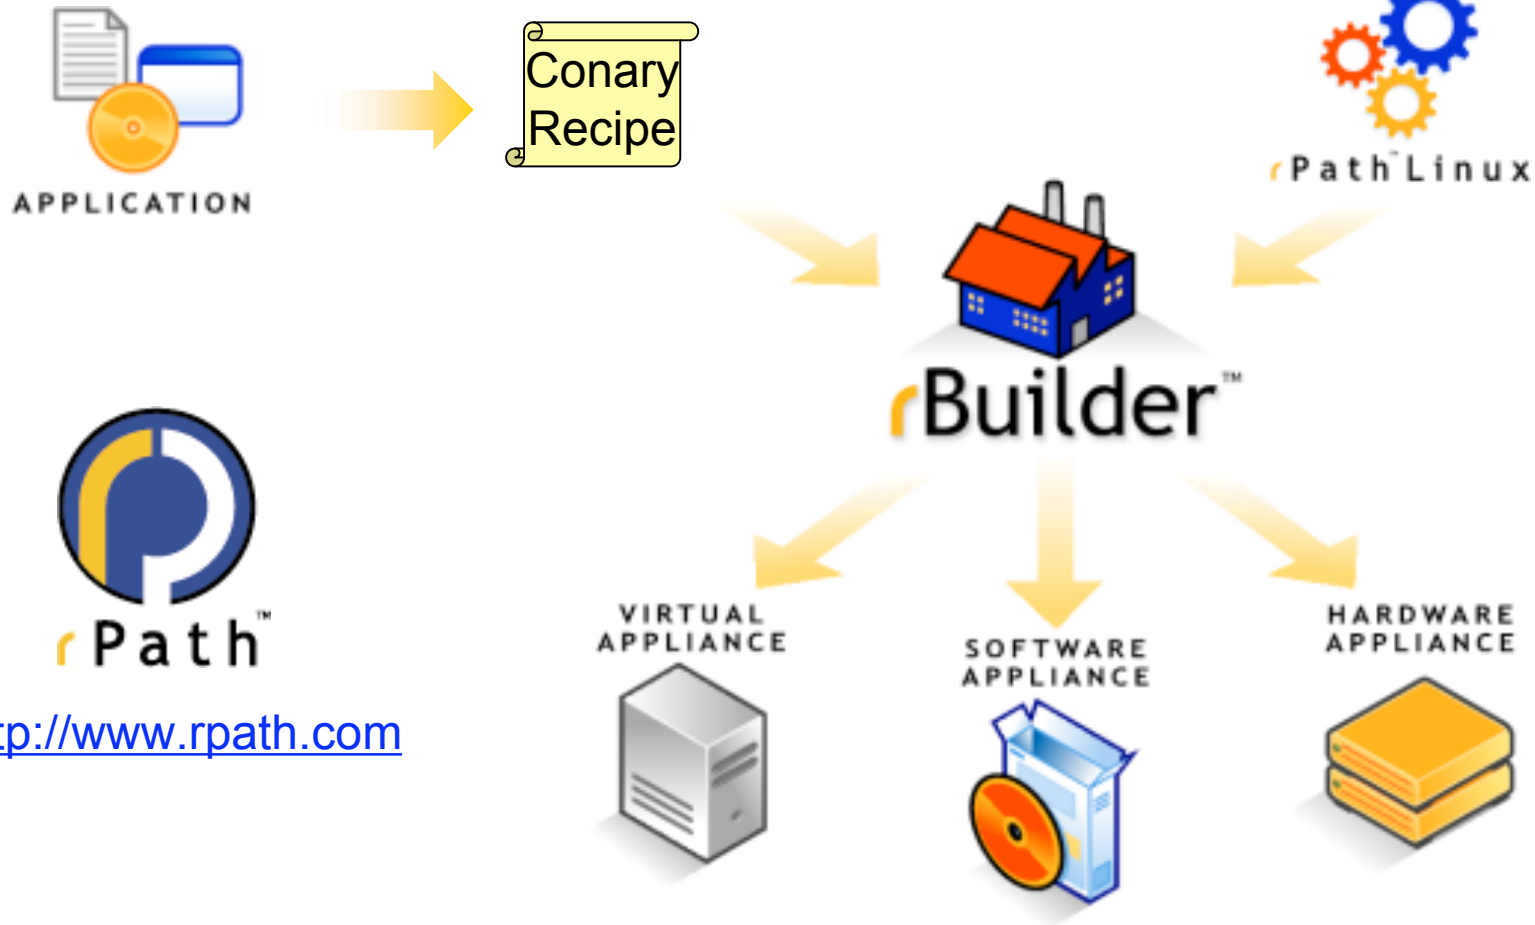
- Over the past couple of years, the industry is silently redefining the meaning of some familiar computing terms
 - Shift from glorious ideas of a large public infrastructure and common middleware towards end-to-end custom solutions and private corporate grids
- New buzzwords
 - Amazon Elastic Computing Cloud
 - Breakthrough in industry approach to computing
 - Everything is for rent (CPU, Storage, Network, Accounting)
 - Blue Cloud (IBM) is coming
 - Software as a Service (SaaS)
 - Google App Engine
 - Virtual Software Appliances and JeOS
- In all this cases, *virtualization* is emerging as a key enabling technology
- Trend supported by CPU manufacturers
 - Multiple cores
 - Hardware virtualization (Intel VT, AMD-V)

- Software @ LHC Experiment(s)
 - Millions of lines of code
 - Complicated software installation/update/configuration procedure, different from experiment to experiment
 - Only a tiny portion of it is really used at runtime in most cases
 - Often incompatible or lagging behind OS versions on desktop/laptop
- Multi core CPUs with hardware support for virtualization
 - Making laptop/desktop ever more powerful and underutilised
- Using virtualization and extra cores to get extra comfort
 - Zero effort to install, maintain and keep up to date the experiment software
 - Reduce the cost of software development by reducing the number of compiler-platform combinations
 - Decouple application lifecycle from evolution of system infrastructure

How do we want to do that?

- Build a “thin” Virtual Software Appliance for use by the LHC experiments
- This appliance should
 - provide a complete, portable and easy to configure user environment for developing and running LHC data analysis locally and on the Grid
 - be independent of physical software and hardware platforms (Linux, Windows, MacOS)
- This should minimize the number of platforms (compiler-OS combinations) on which experiment software needs to be supported and tested, thus reducing the overall cost of LHC software maintenance.
- All this is to be done
 - in collaboration with the LHC experiments and OpenLab
 - By reusing existing solutions where possible
- Resources
 - 7.5 FTE (staff) + 7.5 FTE (fellow) over 4 years

From Application to Appliance



<http://www.rpath.com>

Groups & Builds

Build types

Edit Group

group-cernvm-level version 0.2.0

Version: 0.2.0

Description:

Manage Components ▾

Trove	Version	Lock Version	Delete
group-cernvm	<u>Latest</u>	<input type="checkbox"/>	X
xterm	<u>Latest</u>	<input type="checkbox"/>	X
openssh	<u>Latest</u>	<input type="checkbox"/>	X
openssh-client	<u>Latest</u>	<input type="checkbox"/>	X
openssh-server	<u>Latest</u>	<input type="checkbox"/>	X
emacs	<u>Latest</u>	<input type="checkbox"/>	X
gcc	<u>Latest</u>	<input type="checkbox"/>	X
gcc-c++	<u>Latest</u>	<input type="checkbox"/>	X
gdb	<u>Latest</u>	<input type="checkbox"/>	X
cvs	<u>Latest</u>	<input type="checkbox"/>	X

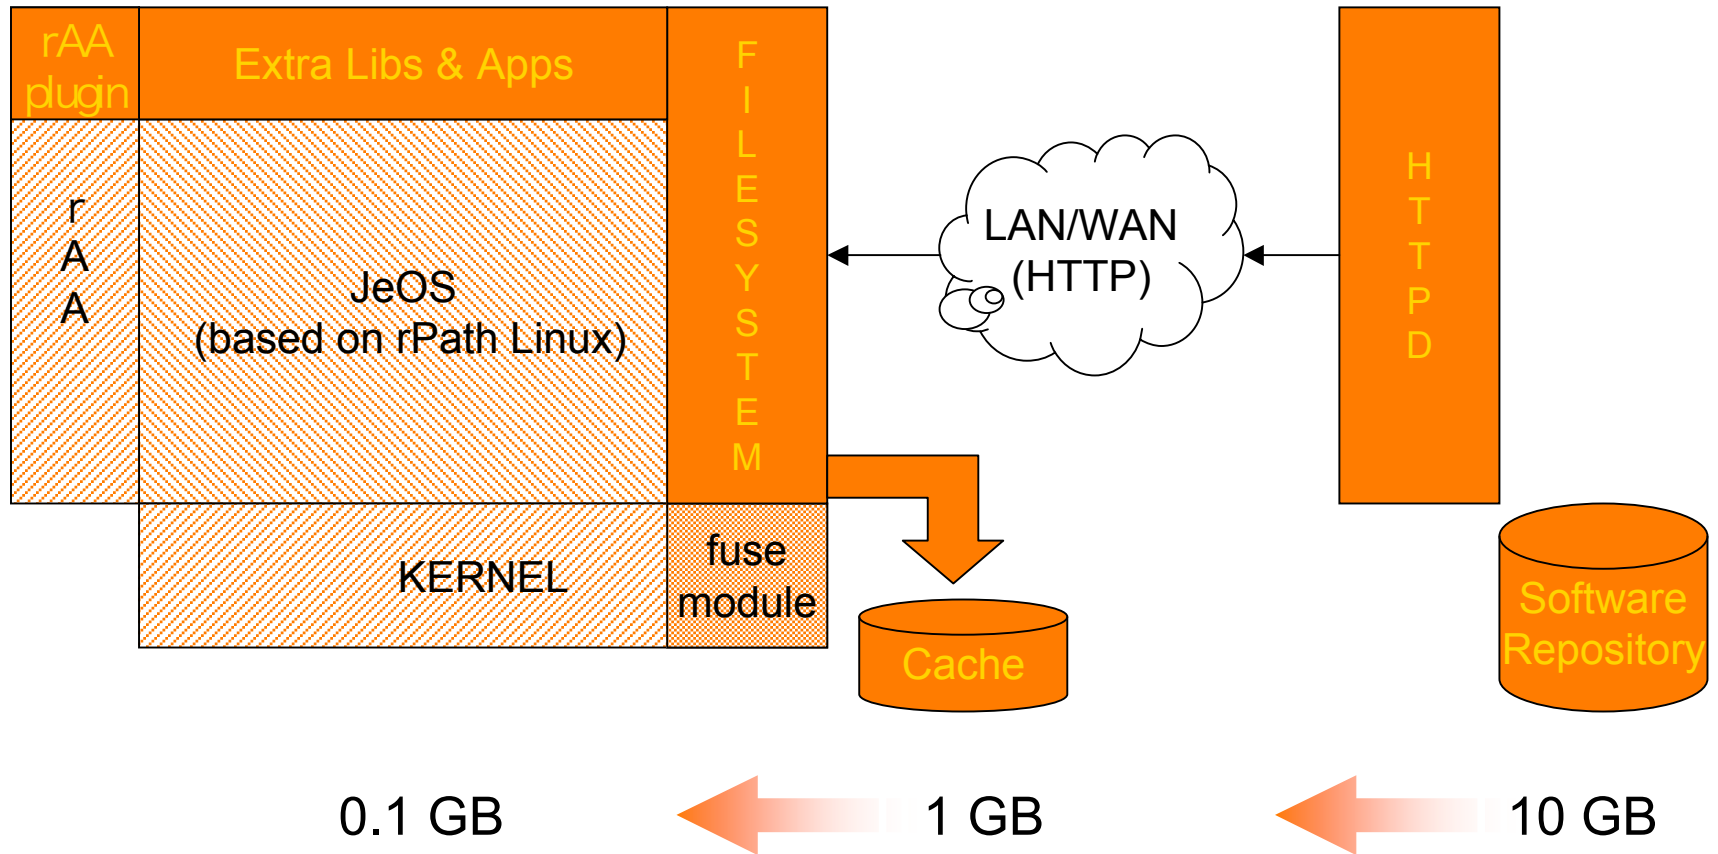
Hover your mouse over the trove version to see the fully-expanded Canary version.

Save and Cook | Save Changes Only | Delete This Group



- Installable CD/DVD
- Stub Image
- Raw Filesystem Image
- Netboot Image
- Compressed Tar File
- Demo CD/DVD (Live CD/DVD)
- Raw Hard Disk Image
- VMware® Virtual Appliance
- VMware® ESX Server Virtual Appliance
- Microsoft® VHD Virtual Appliance
- Xen Enterprise Virtual Appliance
- Virtual Iron Virtual Appliance
- Parallels Virtual Appliance
- Amazon Machine Image
- Update CD/DVD
- Appliance Installable ISO
- Sun VirtualBox Image

“Thin” Software Appliance



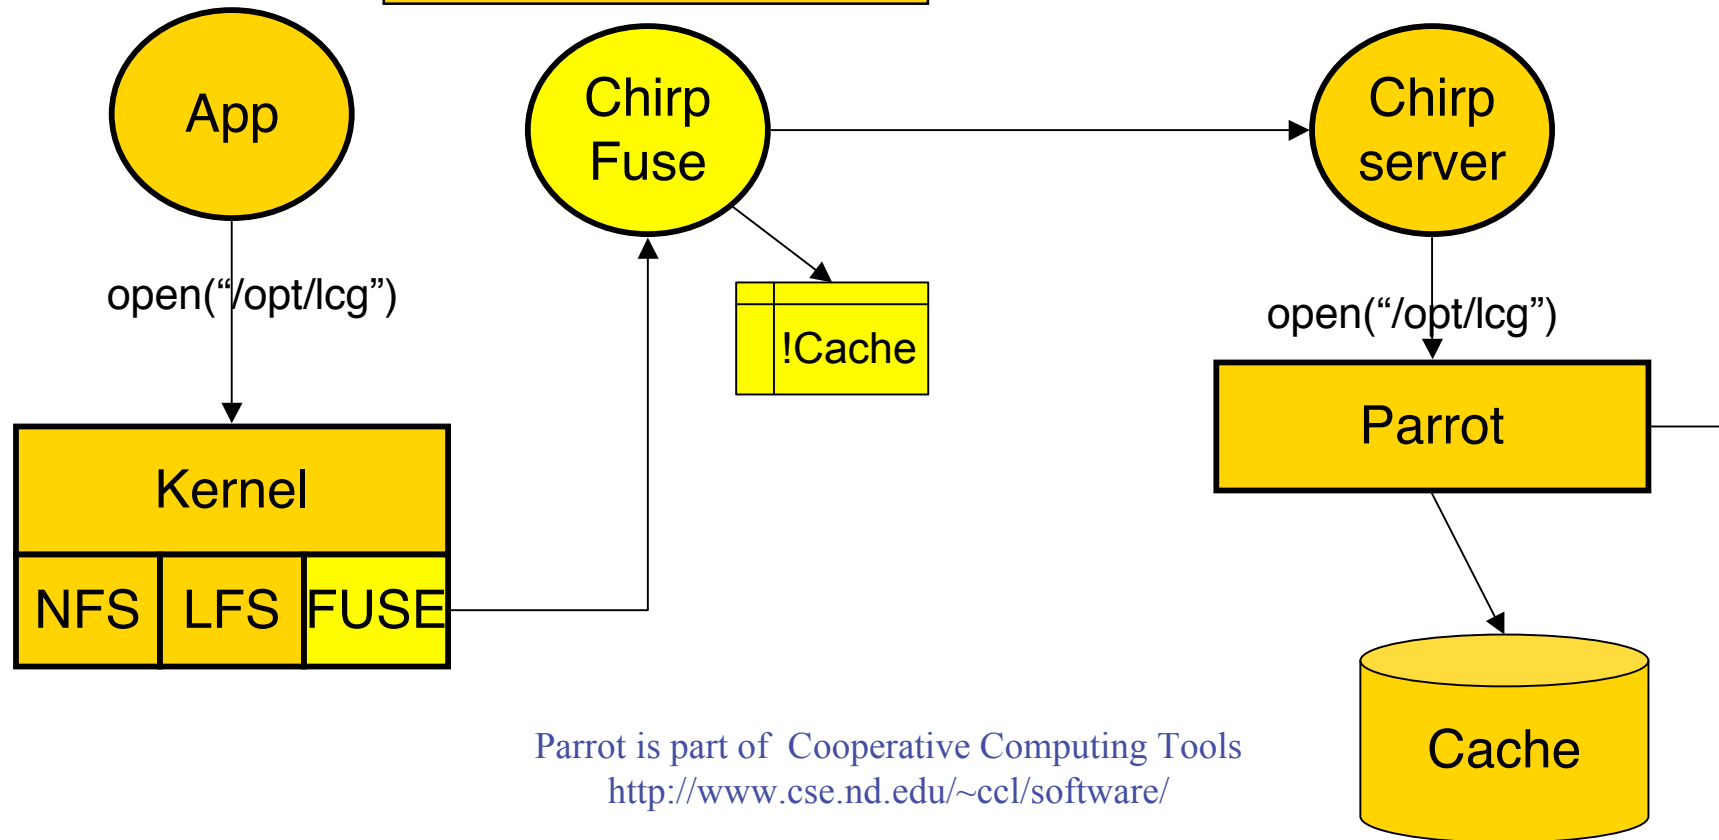
Reusing existing tools

On same host:

```
/opt/lcg
-> /chirp/localhost/opt/lcg
```

On File Server

```
/opt/lcg
-> /grow/host/opt/lcg
```

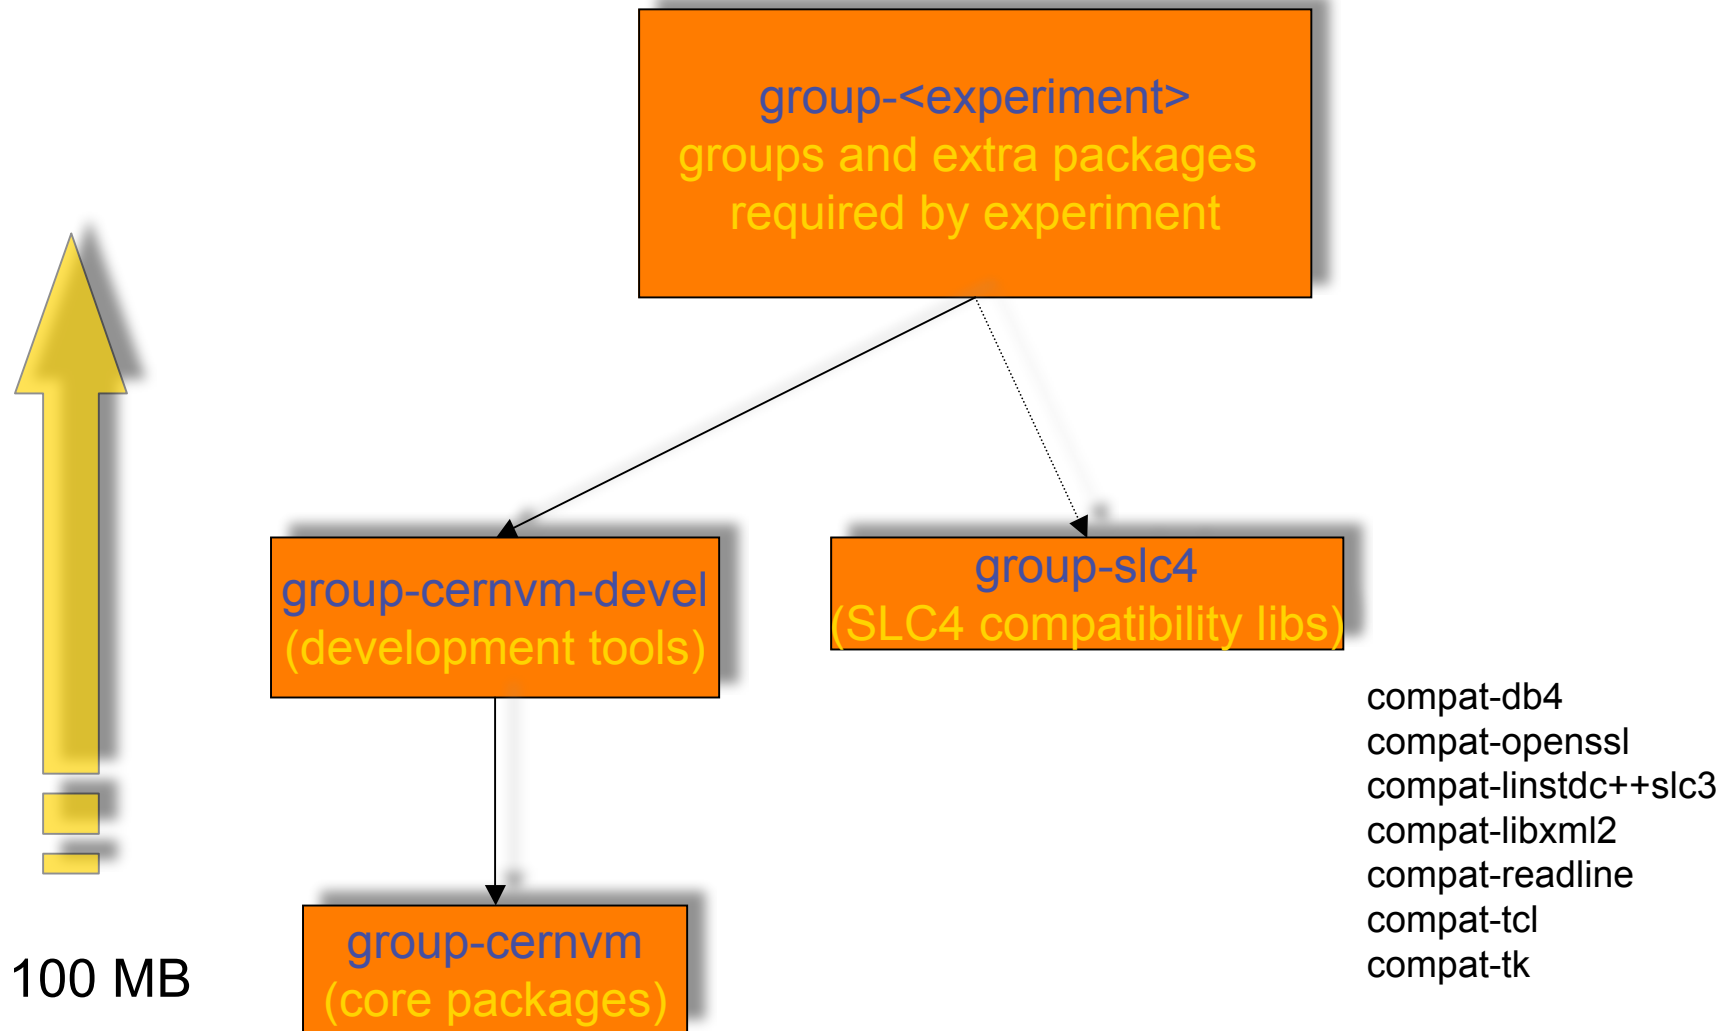


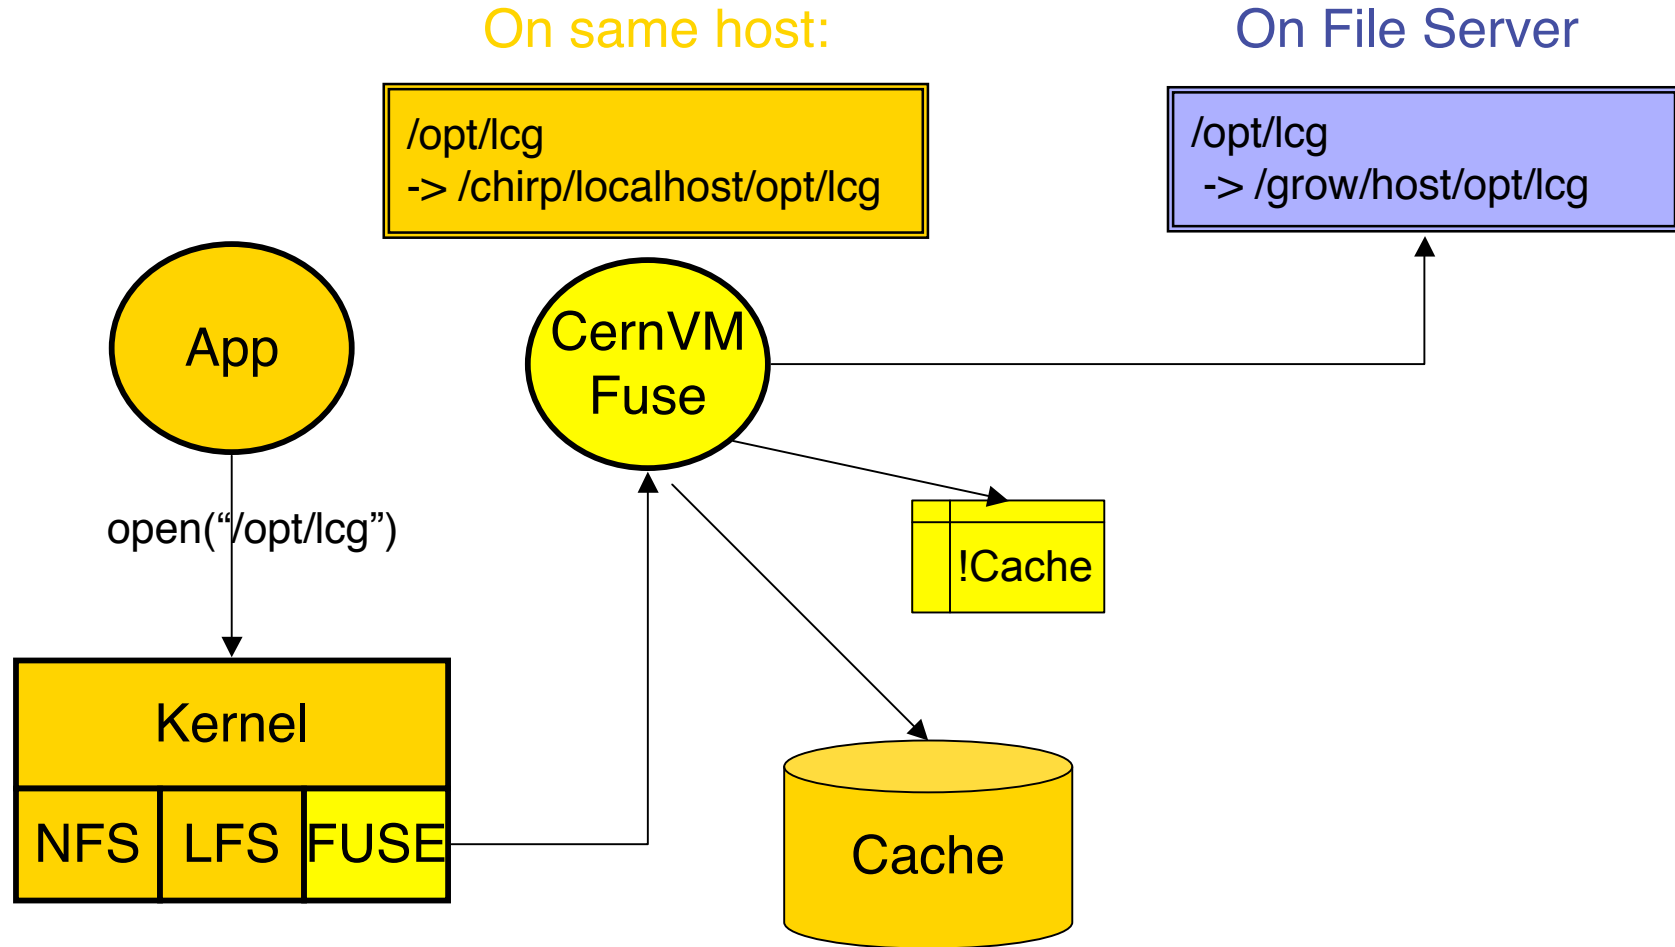
Parrot is part of Cooperative Computing Tools
<http://www.cse.nd.edu/~ccl/software/>

- Available for download from
 - <http://rbuilder.cern.ch/project/cernvm-devel/releases>
- Can be run on
 - Linux (Xen,VMware Player)
 - Windows(WMware Player)
 - Mac (Fusion, Parallels)
- Release Notes
<http://cernvm.web.cern.ch/cernvm/index.cgi?page=ReleaseNotes>
- HowTo
<http://cernvm.web.cern.ch/cernvm/?page=HowTo>
- Limited testing to verify that appliance can be configured and used with Alice, LHCb, ATLAS (and CMS) software frameworks
 - Thanks to Yushu Yao (ATLAS), Dimitrije Maletic (CMS) and Marco Clemencic (LHCb) for help in installing the experiment software and debugging problems

- Not an issue for ALICE
 - They provide a consistent set of dependencies down to gcc level
 - Framework directly depends only on ROOT
- No big problems with ATLAS
 - Requirements to allow flexible cache management (cernvm-cache)
- No feedback from CMS
 - Their approach is similar to ALICE, no big problems expected
- Number of requests from LHCb
 - xauth added to allow users to use ssh -X (or ssh -Y) to forward X11 sessions
 - vim, less and mc packages added to group-cernvm-devel
 - kbdconfig tool to select keyboard layout added to group-base
 - added SLC3 compatibility libraries for gcc 3.2.3
 - added SLC4 compatibility libraries
 - db4, readline, openssl, tcl, tk, libxml2

New package group structure





- Release 0.6 (30/06/08)
 - Bug fixes
 - Adding missing packages and features based on feedback from experiments
 - Addressing SLC4 binary compatibility issues (new package group structure)
- Release 0.7 (04/08/08)
 - Write access to software repository
 - Allowing experiments to publish their releases
 - Support for x86_64 platform
- Release 0.8 (01/09/08)
 - New (simplified and improved) file system
- Release 0.9 (29/09/08)
 - Bug fixes
- Release 1.0 (03/11/08)

Other interesting projects

- BOINC
 - Open-source software for volunteer computing and grid computing
 - CernVM is now extended to include BOINC client
 - Summer student project in OpenLab
 - based on LHC@HOME experience and CernVM image to run ATLAS simulation using BOINC infrastructure
 - <http://boinc.berkeley.edu/>
- Globus Virtual Workspace Service
 - An open source infrastructure for the deployment and management of virtual machines
 - Google Summer School (hosted at ANL) project to deploy a one-click, auto-configuring virtual Grid overlay for Alice/AliEn
 - <http://workspace.globus.org/>

Building the community...

- Mailing lists
 - cernvm-talk@cern.ch (Open list to discuss about design, user experience and related issues with the CernVM project)
 - cernvm.support@cern.ch (End-user support for the CernVM project)
- Savannah Portal
 - Please submit bugs and feature requests to Savannah at
 - <http://savannah.cern.ch/projects/cernvm>
- Web site(s)
 - <http://cernvm.cern.ch>
 - <http://rbuilder.cern.ch>