



WP9 - Workshop Summary

<http://cern.ch/cernvm>

Predrag Buncic (CERN/PH-SFT)

- Summary of Talks in Virtualization Session
- To Do List
- Building the community
- Conclusions

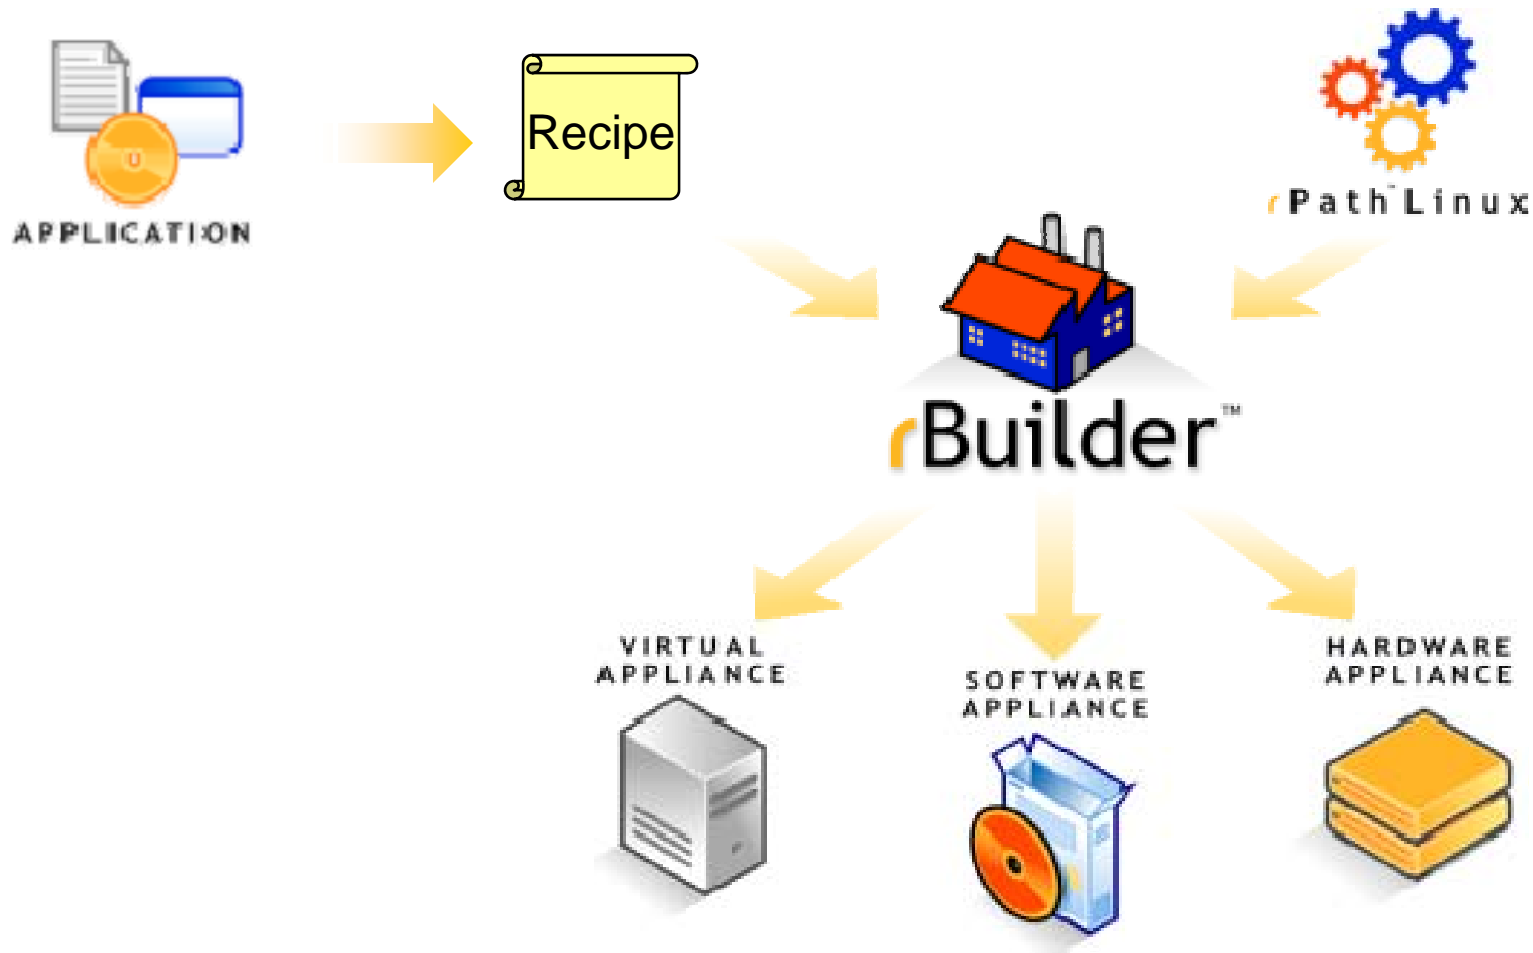
Summary of Talks

- Intel (H Cornelius)
 - Multi and many core CPUs are not future, they are reality in Intel labs
- VMware (Richard Garsthagen)
 - Virtualization technology is much more than simple server virtualization
 - High Availability, Debugging, Vmotion tools, Power saving...
 - This extras do not come for free
- rPath (Tim Gerla)
 - rBuilder makes it really easy to build your own VM images
 - Independent of virtualization technology
 - rPath Appliance Platform makes it simple to upgrade, manage components and services in the appliance
- IgorFS (Thomas Fuhrman, TU Munich)
 - Using P2P technology to build fast and scalable file system for software distribution


CernVM: Initial Project Plan

1. Evaluation of the available virtualization technologies
 - Understand and validate technologies by checking their performance in HEP environment
2. Evaluation of the tools to build and manage Virtual Appliances
3. Collect User Requirements from experiments
 - Confront them with available technologies
 - Suggest an optimal choice for a given use case
4. Development and deployment of a read-only distributed Network File System for software distribution
 - Essential to keep the basic appliance small in size (<100MB)
 - Validate performance, scalability and usability of such approach
5. Provide prototypes for at least two LHC experiments
 - Assist experiments in adapting their software practices to this platform
6. Setup a service and support infrastructure

CernVM: Built using rBuilder



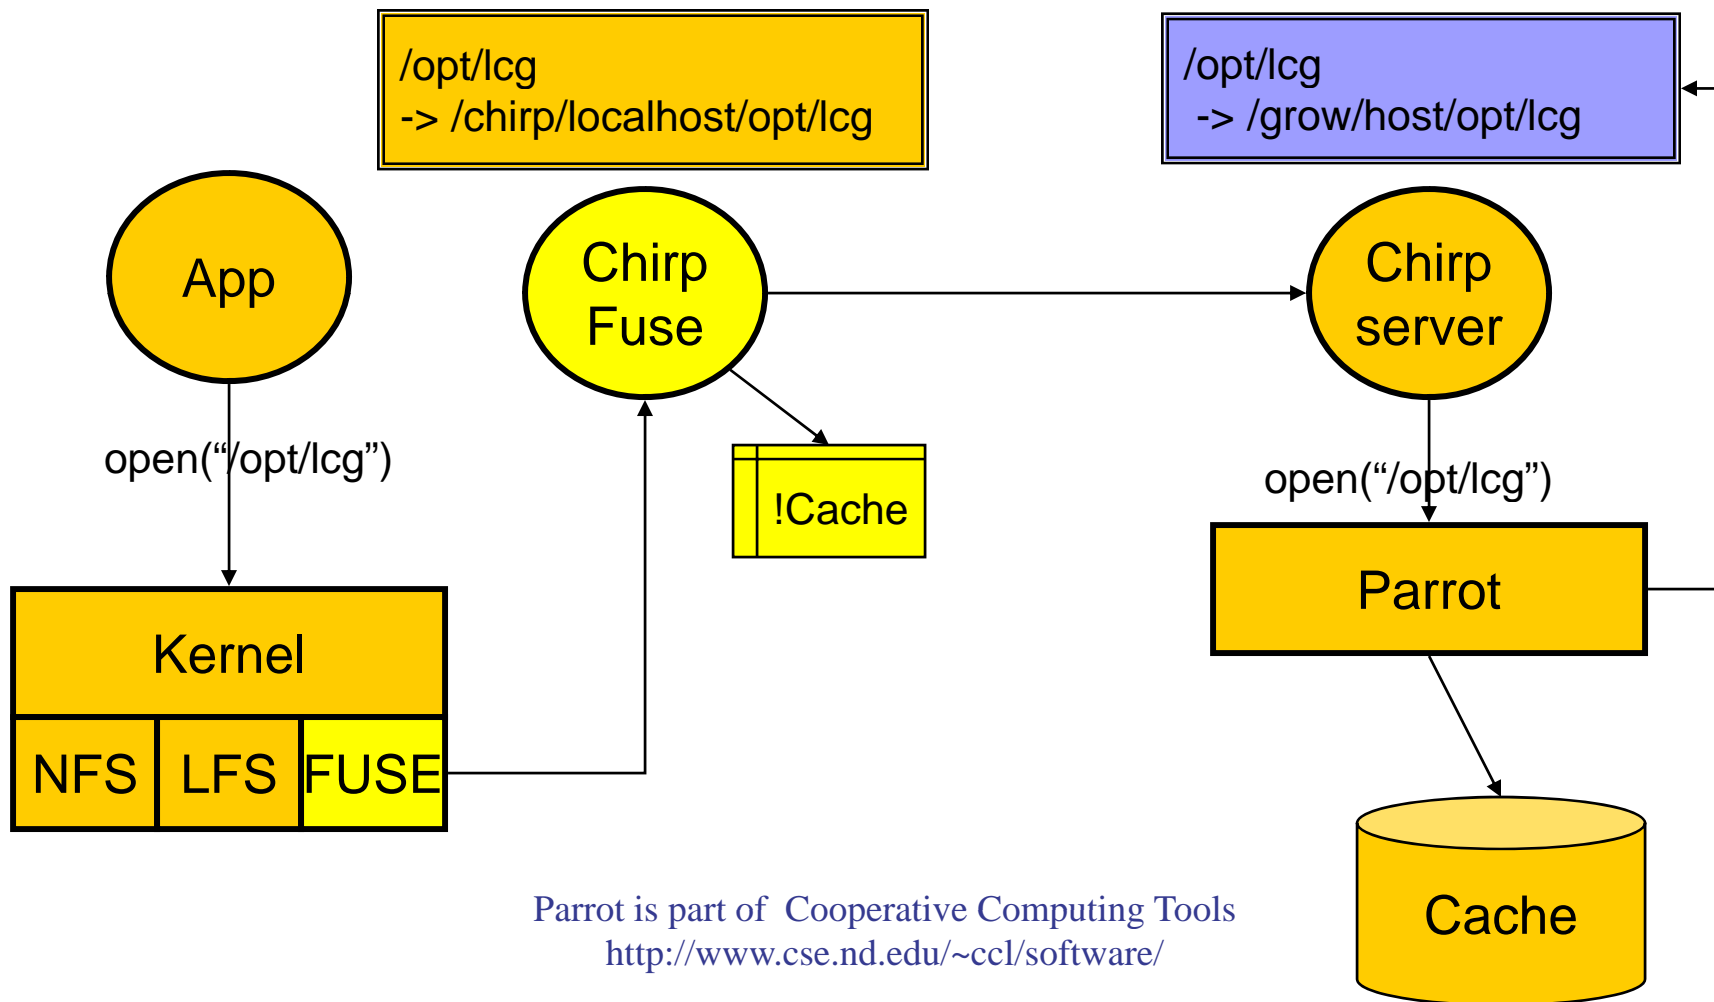
“Thin” Software Appliance



Reusing existing tools

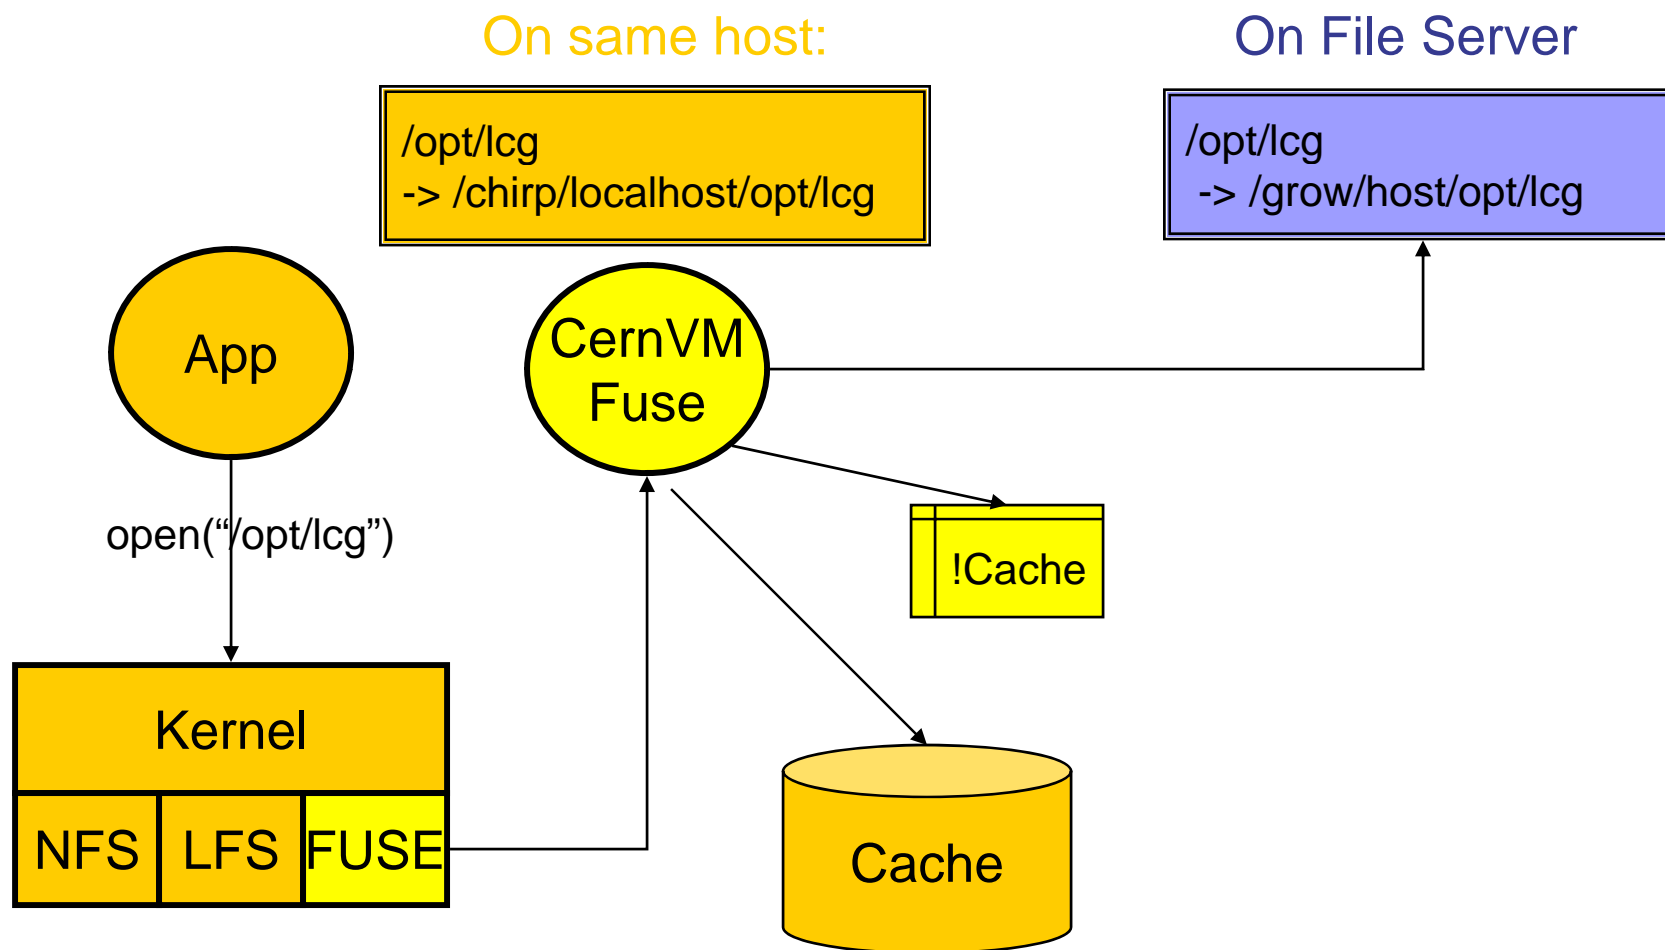
On same host:

On File Server



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

Foreseen development



Alpha release 0.3

QuickTime™ and a
decompressor
are needed to see this picture.

- 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)
- Minimal HowTO at
- <http://cernvm.web.cern.ch/cernvm/?page=HowTo>
- Limited testing to verify that appliance can be configured and used with Alice and LHCb software frameworks
- ATLAS is next in line
- Waiting for input from CMS

Feedback from Experiments

- ALICE (Stefan Boettger):
 - HLT cluster
 - Currently 100, will grow to 900 cores
 - Running grid jobs on VMware VM when not busy with HLT
 - Concern: the price for VMware server product line
 - Question: Can afford CERN site license?
- ATLAS (Yushu Yao):
 - Performance comparison of QEMU and VMWare
 - Conclusion that VMWare offers acceptable performance for casual user
 - Suggestion to distribute Atlas software as virtual disk that users can attach to the VM
- ATLAS (Amir Farbin):
 - Very welcome summary of user requirements based on two years of practical experience
 - Outlines some of the problems that we have tackled in CernVM prototype
 - Stresses concern about I/O performance and optimal way to move things in and out of VM

- SFT (Jakob Blomer):
 - Benchmarking HEP applications in VM
 - Work done as summer student last year
 - Applications tested: Geant4, ROOT, AliRoot
 - Results point to very good performance of Xen under Linux
 - Surprisingly good performance of VMware and Parallels under MacOs
- IT (Andreas Unterkircher):
 - Using VMs in test environment
 - Deploying grid components for testing
 - VNODE tool
 - Xen supported by IT (Quattor profile for Xen exist for SLC4)
- SFT (Carlos Aguado Sanchez):
 - Looking for Open Source tools to deploy VMs on a small physical clusters
 - Enomalism2 chosen

- Analyse Atlas requirements
 - Formalize in URD
- Continue development of read-only file system for software distribution
 - Compatible with both Atlas presentations
 - Simplify and make it faster
 - Add extra management interfaces
 - Explicitly fetch subdirectory
 - Pin the directory content in cache
 - Clear the cache
- Continue work towards first release
 - Add support for ATLAS, CMS
 - Functional test of exp frameworks
 - Performance test
- Add capability to act as Grid User interface
 - Verify that Grid Data Management works
 - Know problem with rfio
- Documentation

Building the community

- Regular meetings (2 weeks)
- Mailing list
- Web site
 - <http://cern.ch/cernvm>

- The work on the project has started according to initial planning and is progressing well
 - Planning based on initial discussions with experiments trying to identify interest for collaboration
 - Update plan with items in To Do list and suggestions from the floor
- Contacts have been established with goal to carry out further research
- The evaluation of performance indicates that it is possible to use virtualization technologies in HEP context
- rBilder instance installed at CERN and seem to live up to the promises
- Work on setting up scalable service infrastructure is in progress
- Prototype of file system for s/w distribution is implemented
- Alpha release is available and is ready for 2 experiments
 - Work on Atlas will start next
- Expecting first release that could be used by the real users in September