

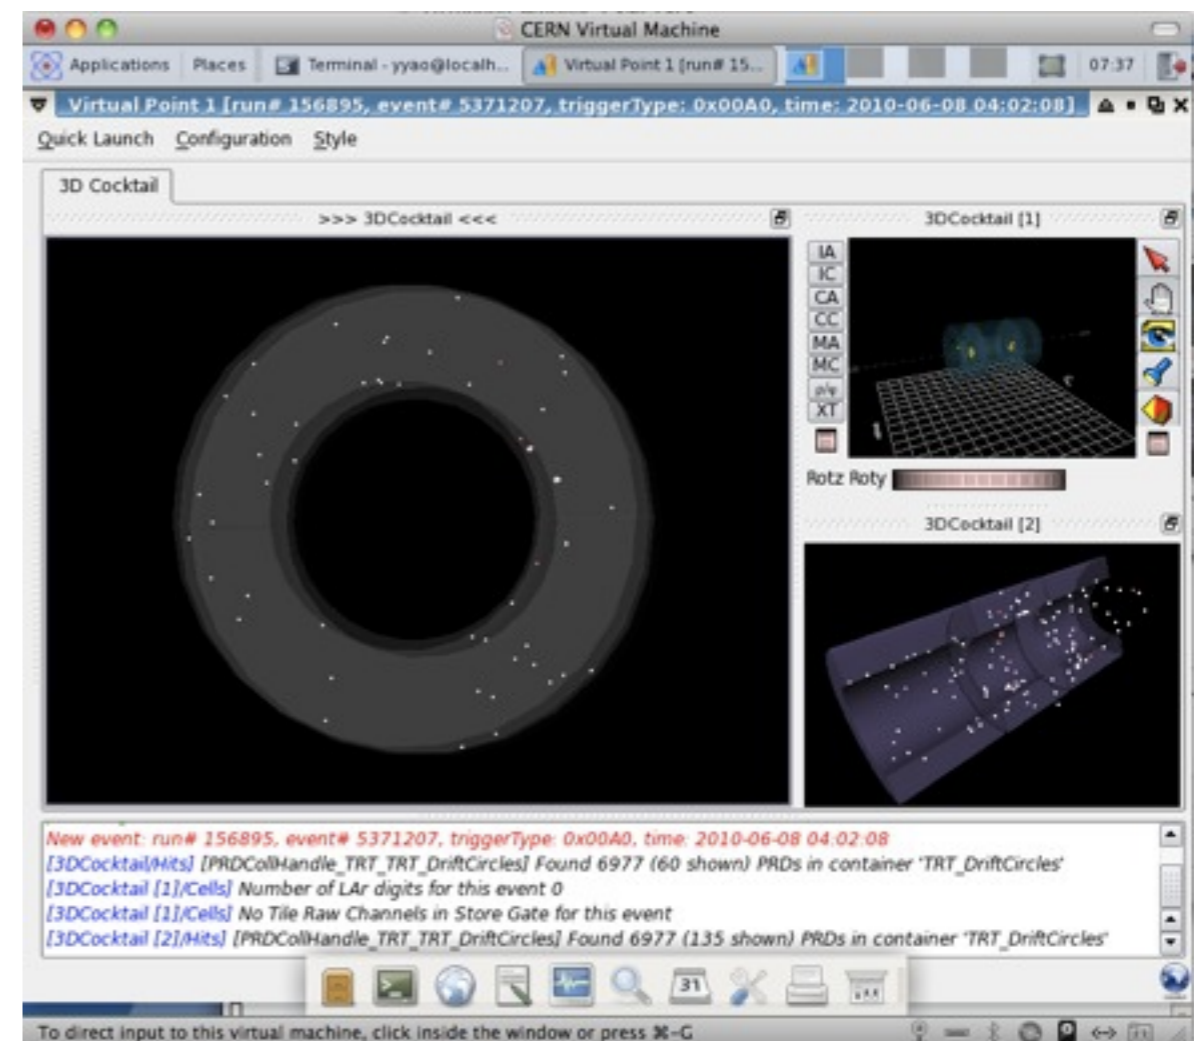
# Cern VM-FS (CVMFS) in ATLAS Tier3 - Painless ATLAS Software Distribution

Yushu Yao  
June 8, 2010

# Briefing

- CernVM and CernVM-FS (old name CVMFS)
- Why CernVM-FS
- How CernVM-FS works
- What CernVM-FS Provides for Tier3 Users
- How to use CernVM-FS

- The simplest yet fully functional Tier3-workstation.
- All ATLAS SW, Grid Job Submission
- I-click VPI Live
- Tutorial: <https://twiki.cern.ch/twiki/bin/view/Atlas/CernVMTutorialHead>




# Getting ATLAS software, the existing solution


- Before CernVM-FS:
  - Each tier3 site will have a manager (a graduate student) who install all the releases via kit. Which works but:
- ATLAS has too **MANY BIG** Releases
  - **10 GB per releases:**
    - waste of disk space/network bandwidth since not all are needed.
  - **One major releases every month, several patch releases every week:**
    - The graduate student who's managing the SW says: come on, I need time to do my thesis!

# CernVM-FS

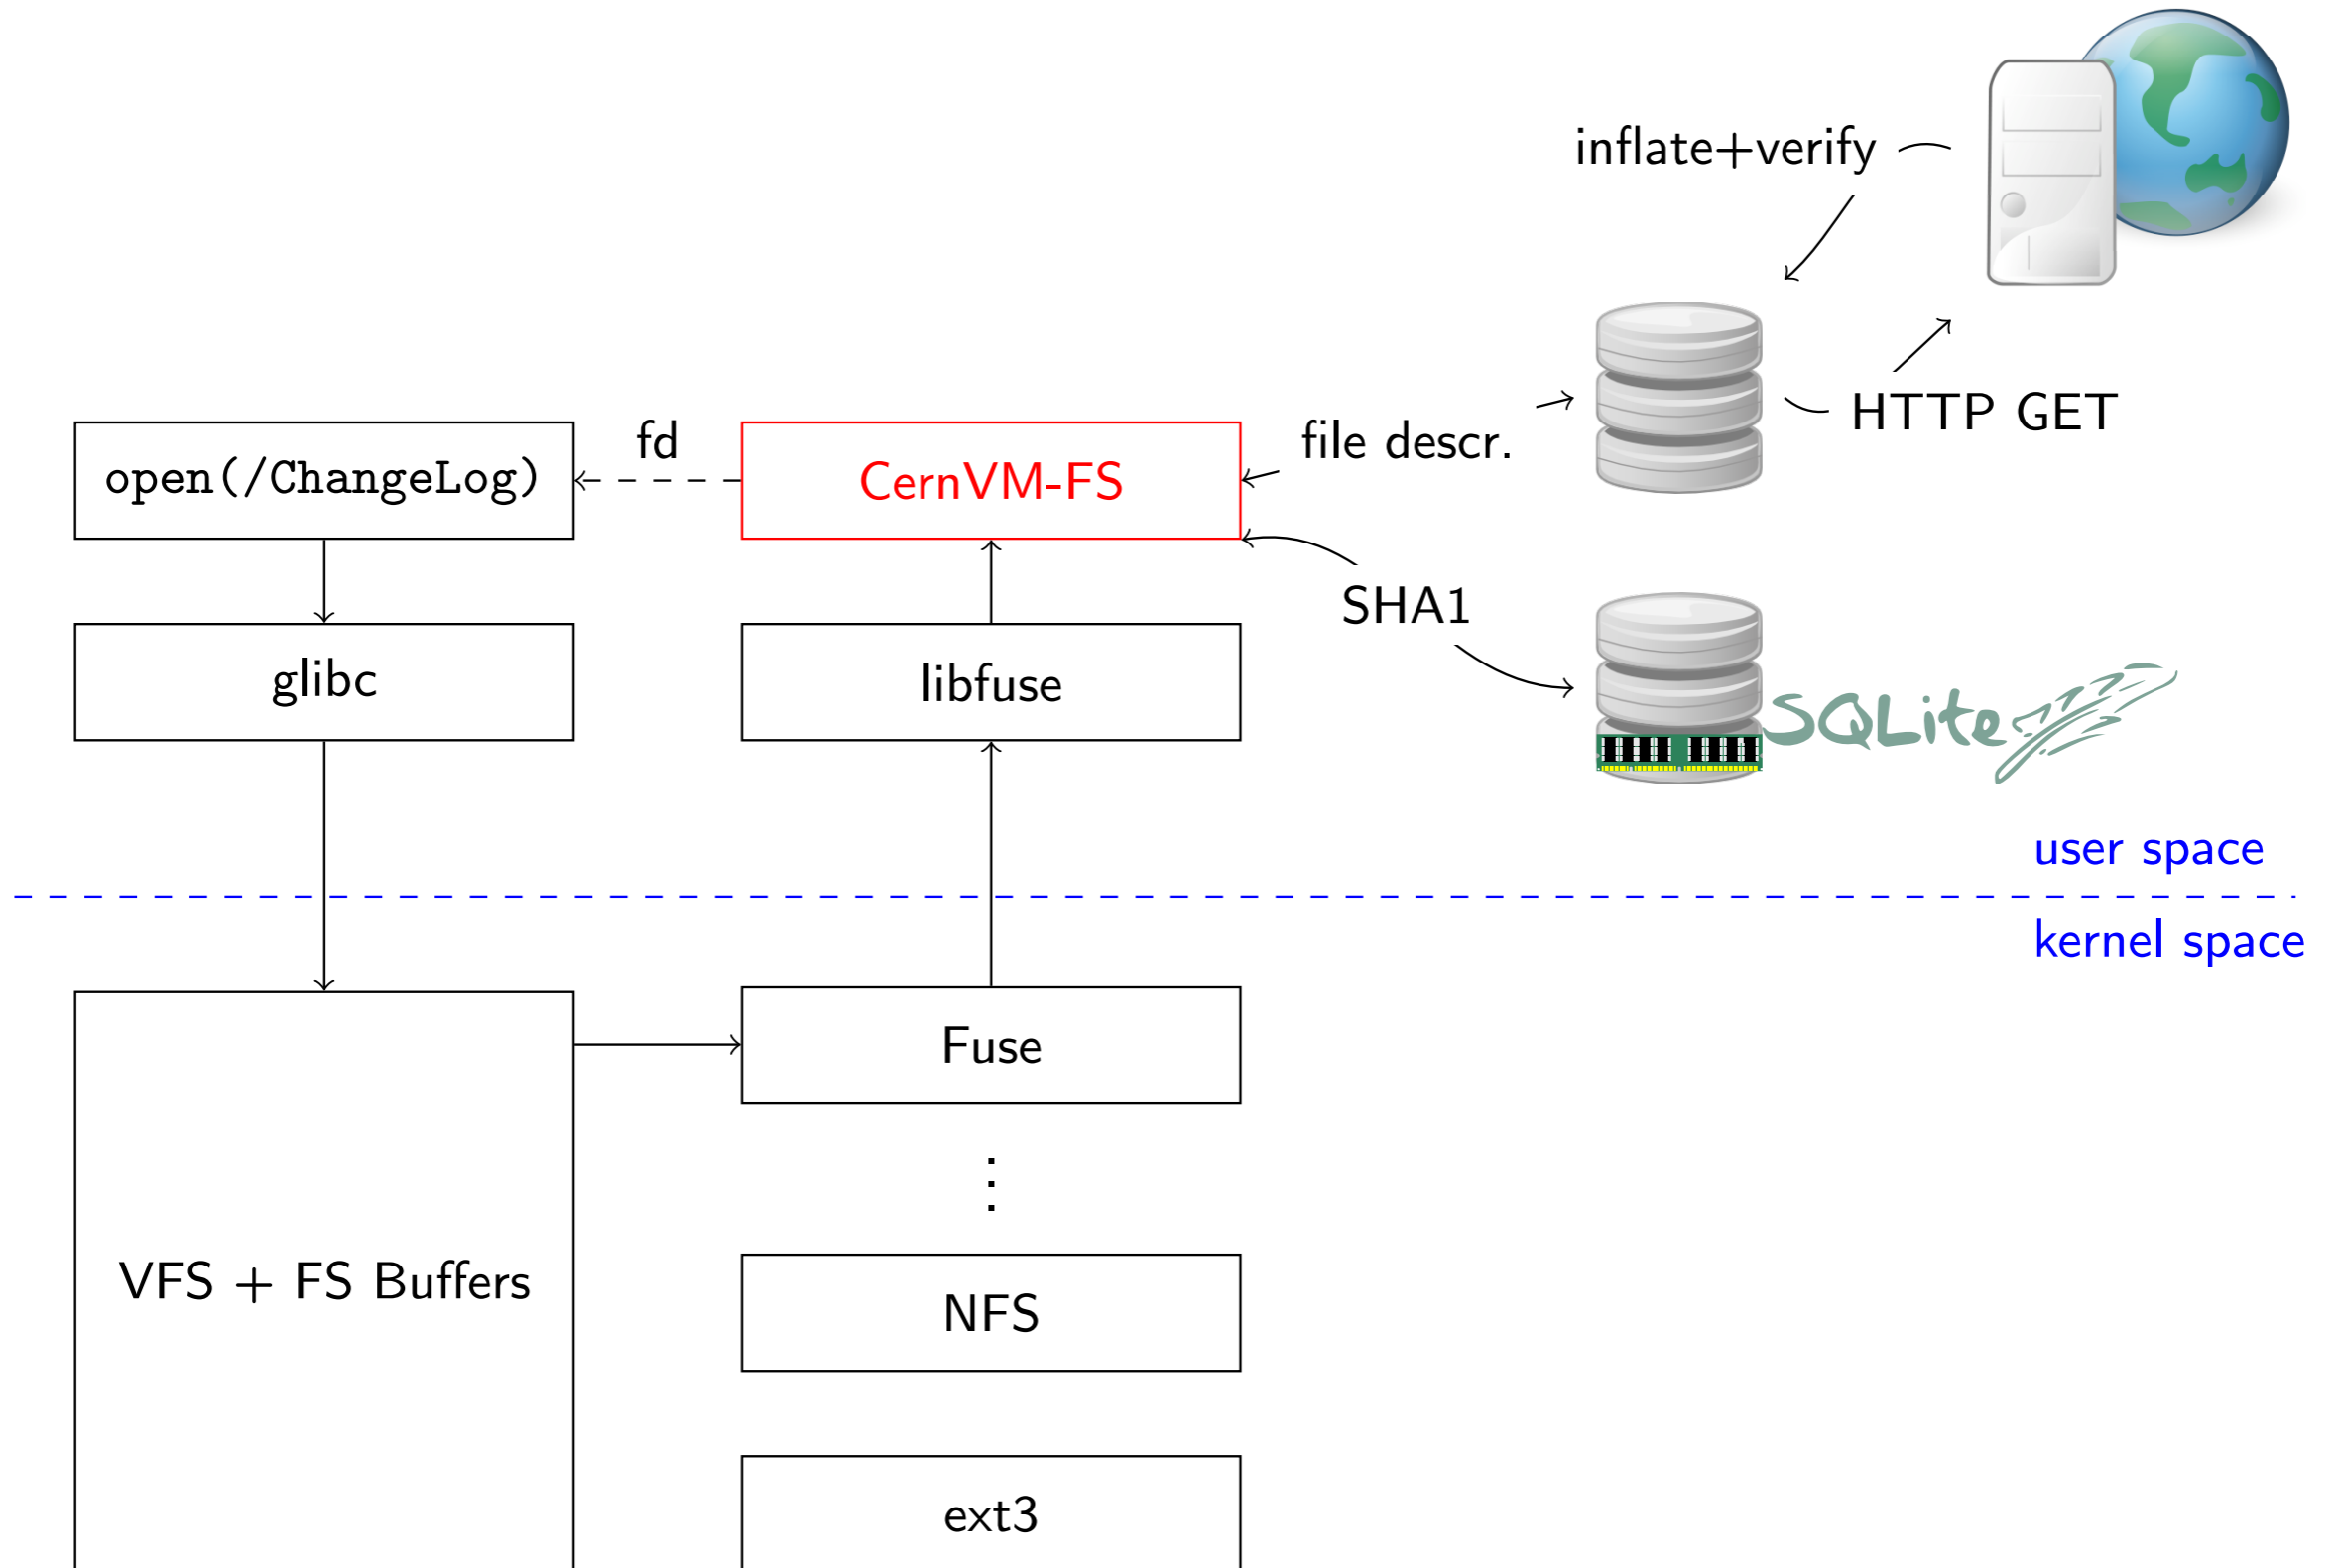
CernVM-FS is a filesystem to distribute ATLAS software/CondDB



1. ATLAS SW is put on central server and centrally managed
  2. CernVM-FS present the file catalog to Users, it looks like a normal Linux FS
  3. CernVM-FS will obtain via HTTP a file only if it is opened
  4. If a file is opened a 2nd time, it present the locally cached copy
  5. HTTP Proxy servers are used to speed up web access.
- For the files you never touch, they will never reach your machine.
    - only **800MB** of files are downloaded to run full reco (compare to 10GB releases)



# How CVMFS Works



# Highlights of CernVM-FS

- Requires only **outgoing** HTTP(S) connection, i. e. works with practically every Internet connection, even you are behind multi-layer firewalls.
- **Verifies** file integrity on download by SHA1 checksum
- Automatic failover for chain of forward/reverse proxy servers
- Possibility to pre-load cache
- Offline mode
- Multi-Mount
- Trace file system operations
- Nested catalogs
- Catalogs can be signed by X.509 certificate
- **Catalogs are stored together with a time to live, which allows for automatic updates**

# Performance of CernVM-FS

- Close to local disk performance once cached.
- Way better than NFS mounted.
- Ideal for batch clusters (where access pattern are similar)

Rik Yoshida (ANL, March 2010)

Dell R710: 8 cores (16 hyperthreaded)

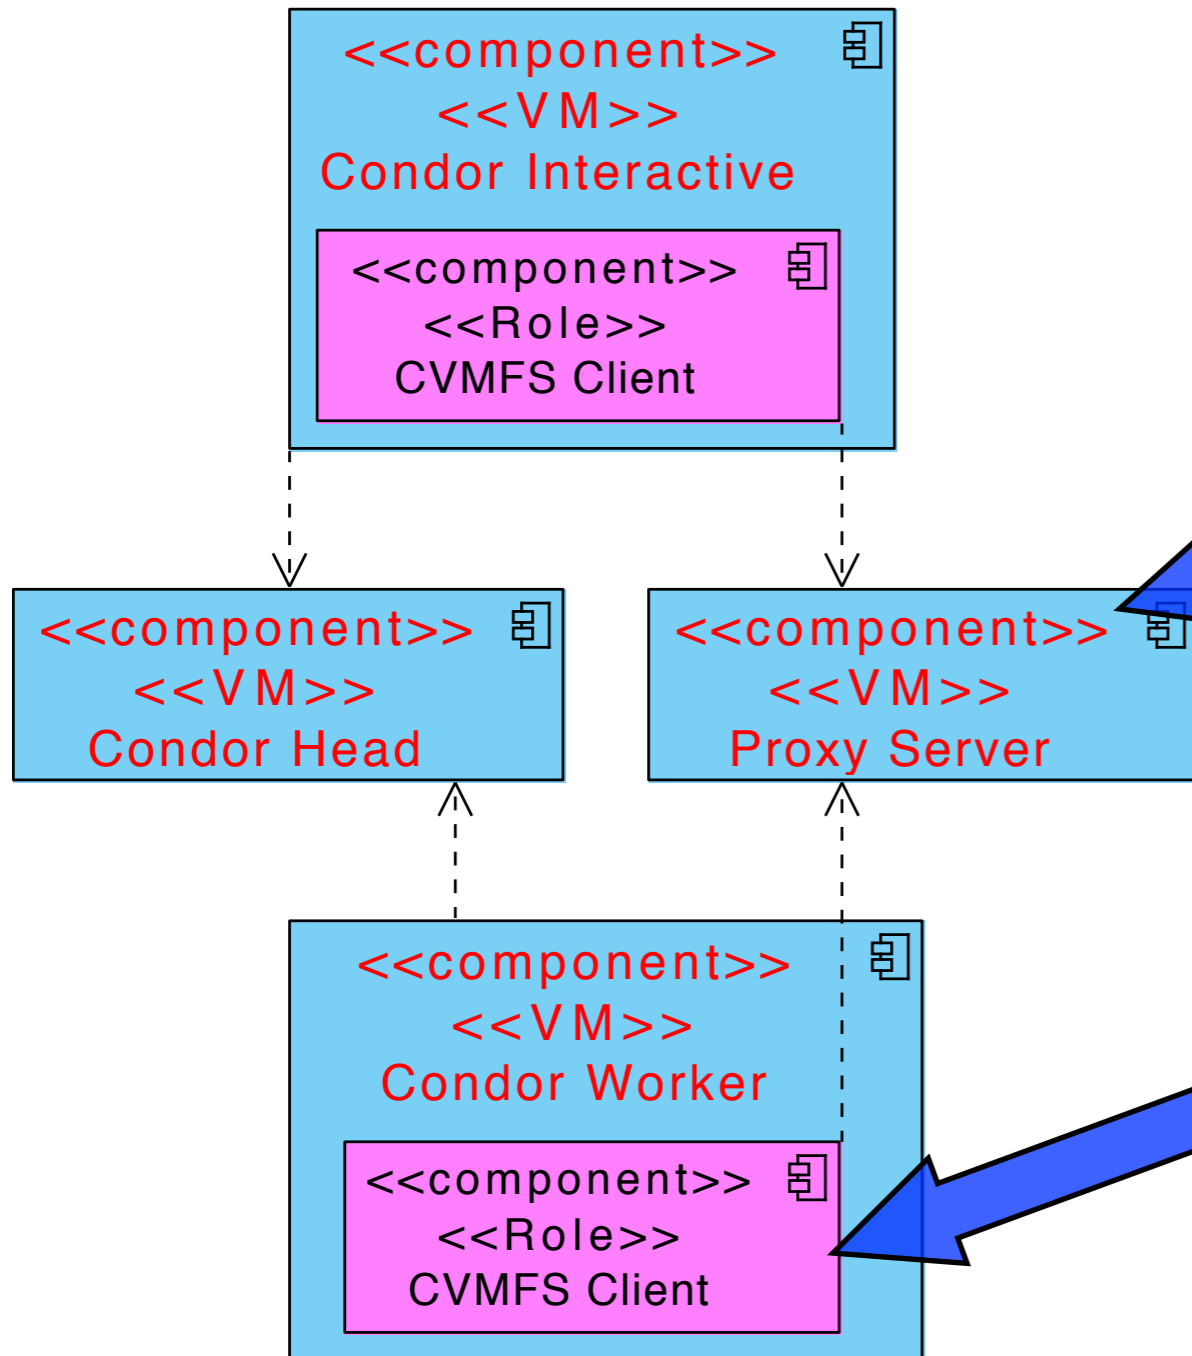
No. Simultaneous Condor jobs:	1	4	8	14
NFS4	7 min	15 min	60 min	
CVMFS2	7 min		8 min	11 min



# CernVM-FS repositories for ATLAS

- We provide three repositories:
  - **ATLAS software releases (Ready RightNow)**
    - All major releases and patches
    - Adding this to the official ATLAS release building procedure, making sure everything is available
    - Fully Compatible with manageTier3SW (Rik mentioned)
  - **ATLAS conditions database**
    - Dedicated server at CERN. Synchronized and published daily.
  - **ATLAS software nightly builds**
    - Dedicated server at CERN, updated for each

# Setting Up CernVM-FS in Your Tier3 Cluster



- Proxy Server
  - Need 100GB disk, 2GB+ Memory
  - We provide (or will provide):
    - Pre-configured RPMs
    - Ready-to-run Virtual Appliance
    - Puppet Definitions

Also work for Frontier CondDB access (i.e. one stone two birds)

- CernVM-FS Client
  - Need 50GB disk
  - Install RPM
  - Customize
  - Start Service


- Preferred: a good network connection between Proxy Server and CernVM-FS Clients

# Maintaining CernVM-FS

- NO Action needed when new SW releases or CondDB be come available. CernVM-FS will find figure this out by itself (normally within 1 day)

# Now what we have to ensure you get Atlas SW

- CernVM-FS
- locally manage Tier3SW installed releases
- Manual Kit installations



# Summary

- CernVM-FS provides ATLAS SW and Conditions DB
- Centrally Managed Repository (No need to install new releases anymore)
- Easy Installation / Near Zero Maintenance
- **READY FOR PRODUCTION**
- Help Provided

# Backup

# Secure CernVM-FS

