



# Introduction to CVMFS

A way to distribute HEP software on cloud

Tian Yan

(IHEP Computing Center, [yant@ihep.ac.cn](mailto:yant@ihep.ac.cn))

BESIIIICGEM Cloud Computing Summer School


2015/9/9, IHEP, Beijing



# Outline



- ❖ Overview
- ❖ How does it work?
- ❖ How to install, configure, and debug it?



Part I

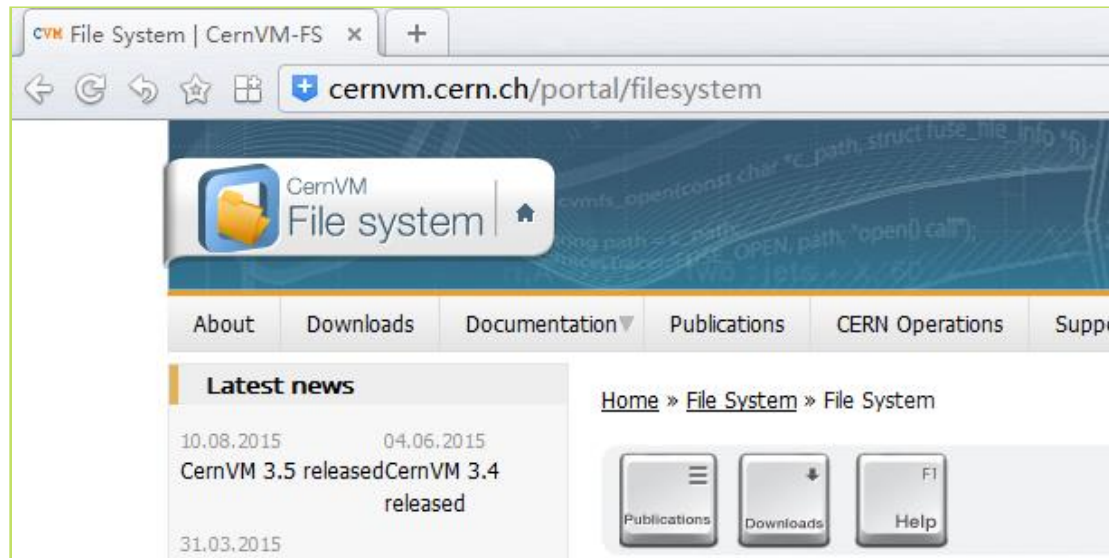
# OVERVIEW



# What is CVMFS?



- ❖ CVMFS = CERN VM File System
- ❖ POSIX read-only network file system based on HTTP
- ❖ designed and optimized for HEP software distribution
- ❖ official site: <http://cernvm.cern.ch/portal/filesystem>





# Why we use it?



- ❖ Widely used in HEP community
  - developed in CERN and widely used by WLCG and other HEPs
- ❖ Fast
  - based on HTTP, allow exploitation of various web caches (e.g. squid, commercial content delivery networks)
  - aggressively cached to reduce latency
  - transfer data and metadata on demand
- ❖ Flexible
  - filesystem versioning and hotpatching file-by-file
  - software in VM/Docker image need further packaged



# Why we use it?



## ❖ Scalable

- one stratum 0 (release manager machine), many stratum 1 mirrors
- automatic mirror server selection based on geographic proximity
- use web/local cache to reduce direct access to server

## ❖ Reliable

- it verifies data integrity by cryptographic hashes

## ❖ Easy to install, configure and maintain

- only outgoing HTTP, firewall friendly
- can be installed by one shell script (less than one page)
- it's stable, so it rarely need maintain effort



Part II

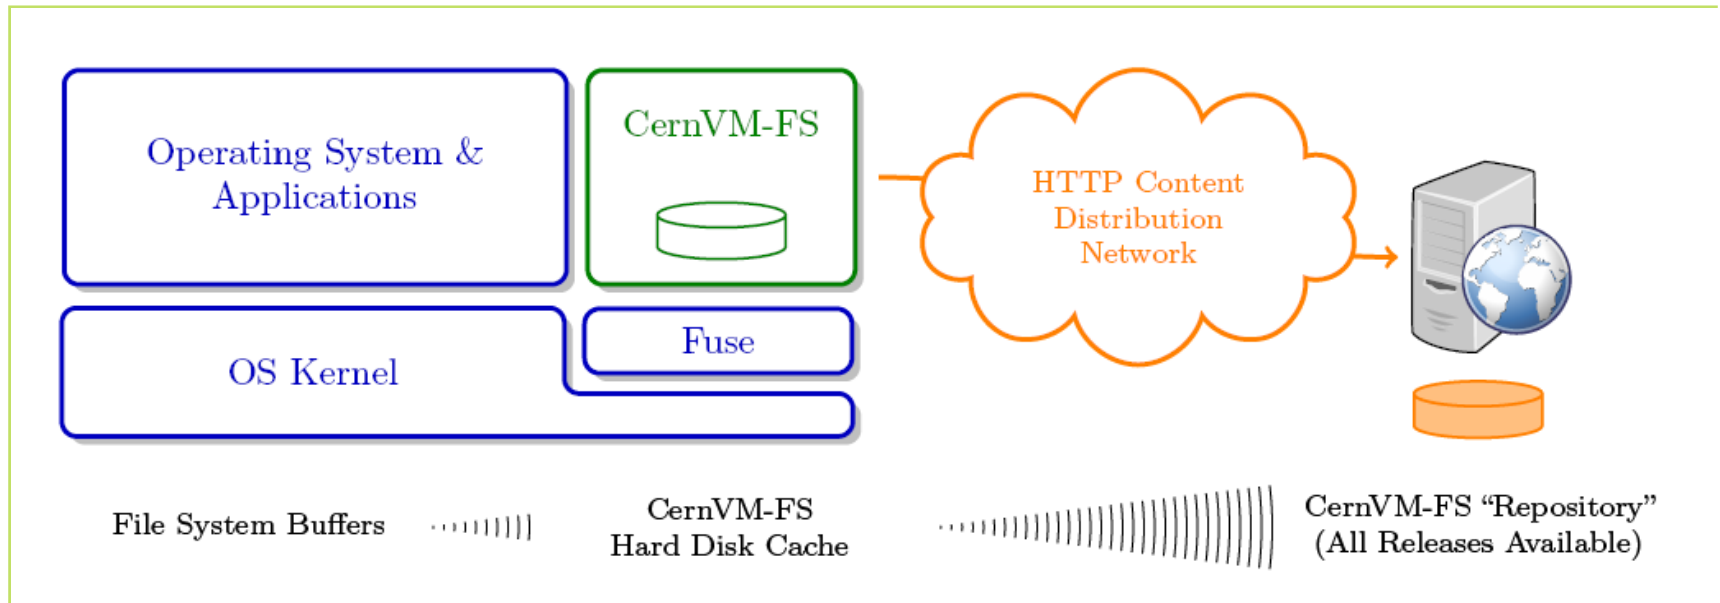
# HOW DOES IT WORK?



# A glance



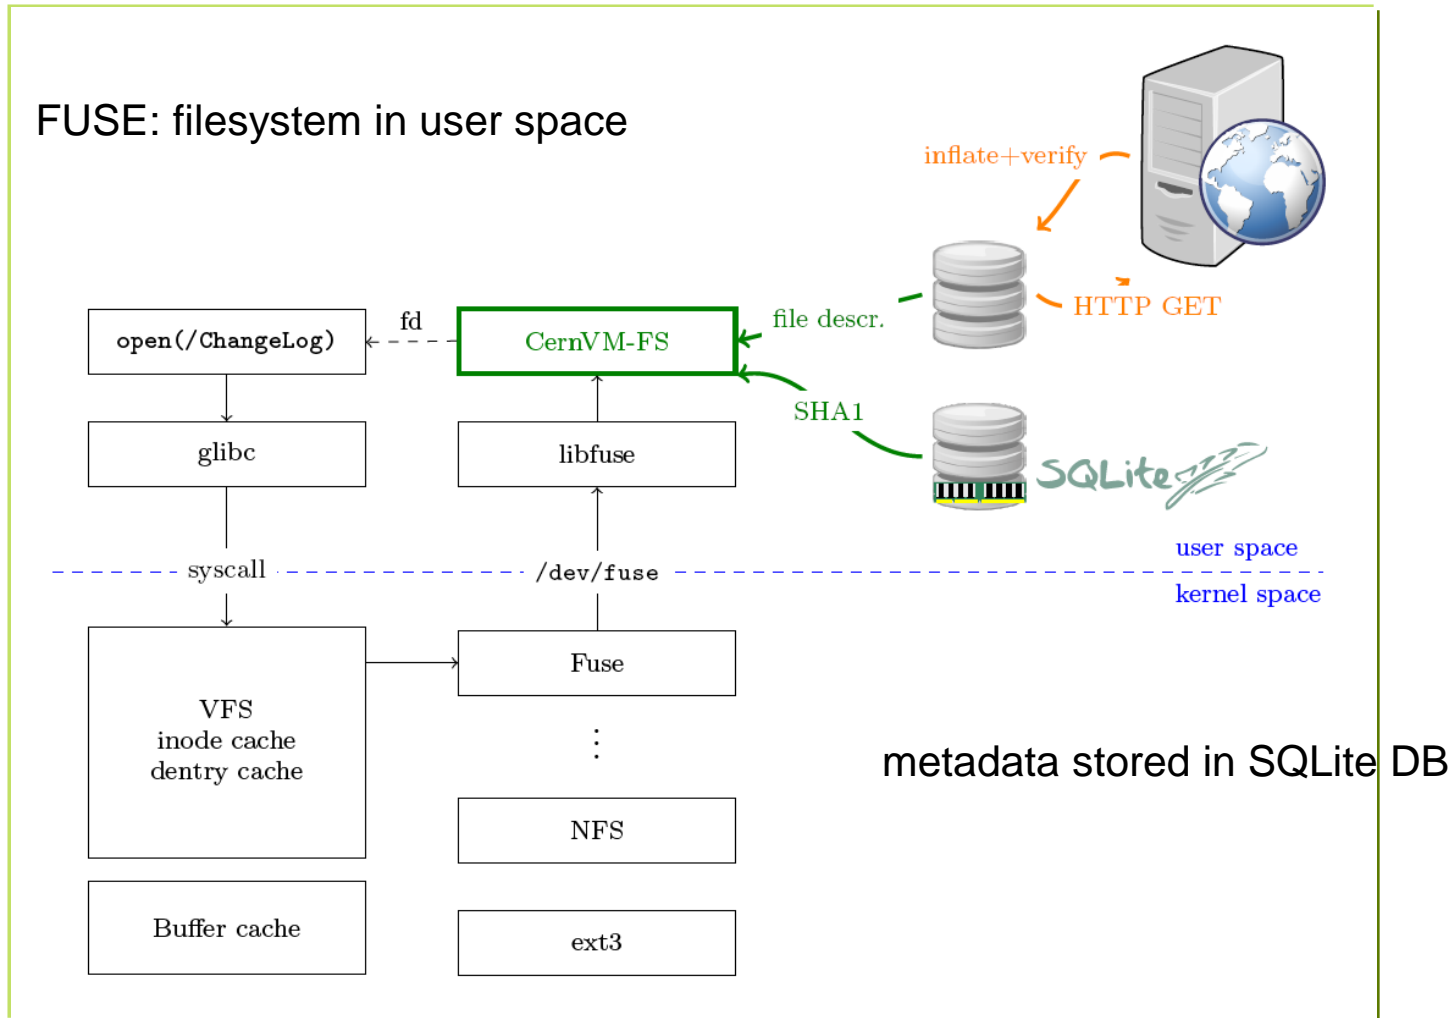
- ❖ left side is client; right side is server with http interface
- ❖ virtual filesystem based on FUSE in user space
- ❖ load data only on access





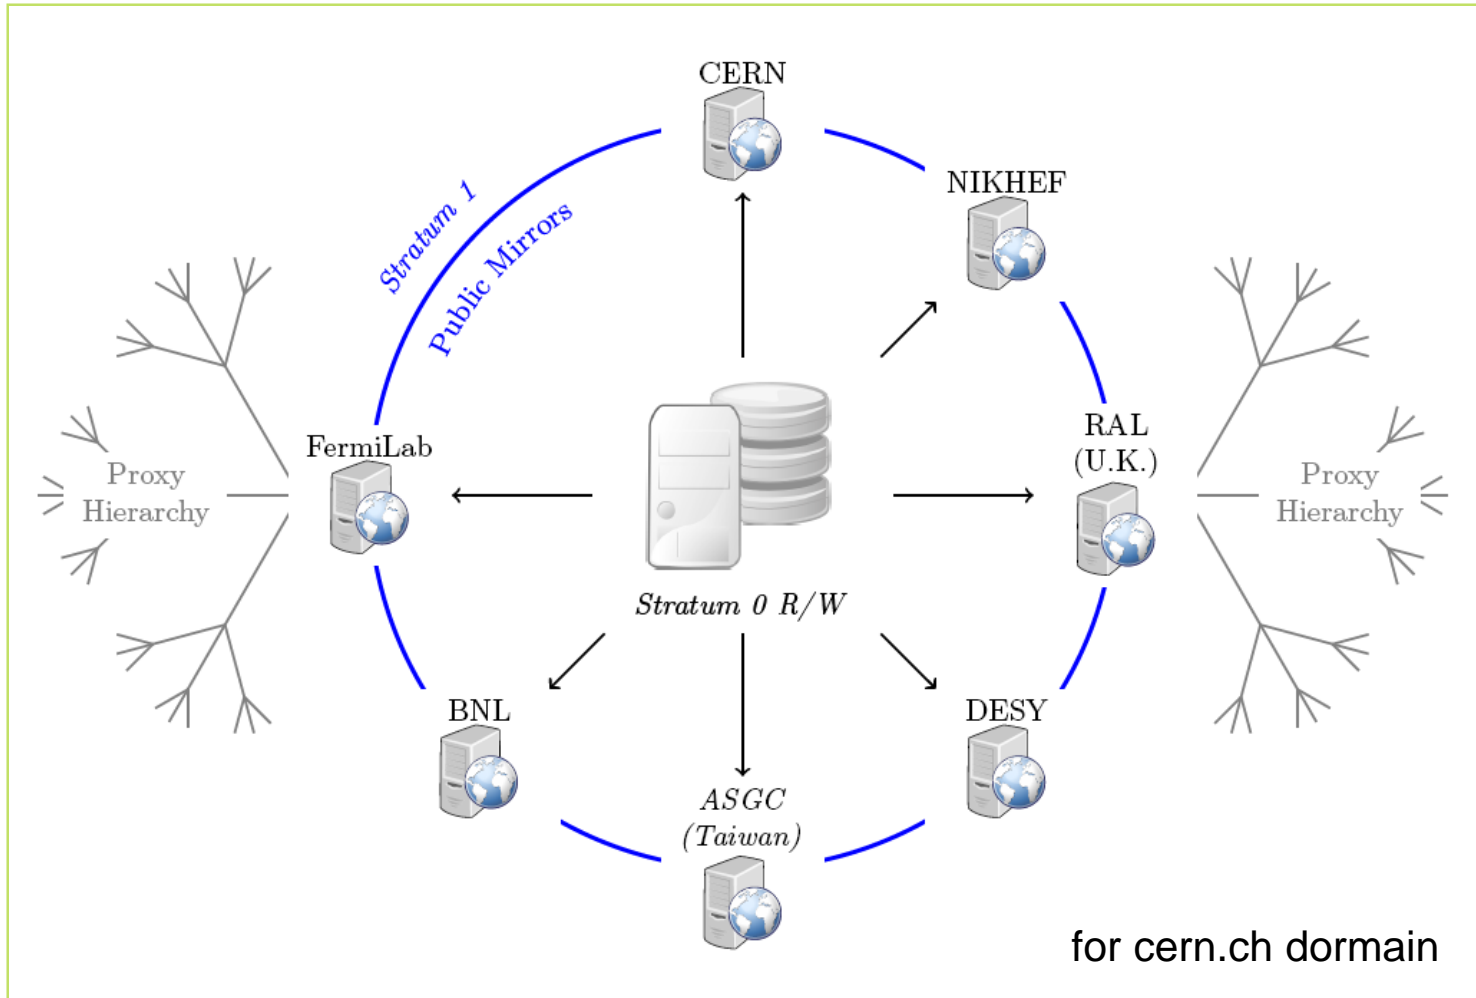


# Opening a file on CVMFS





# CVMFS server stratum

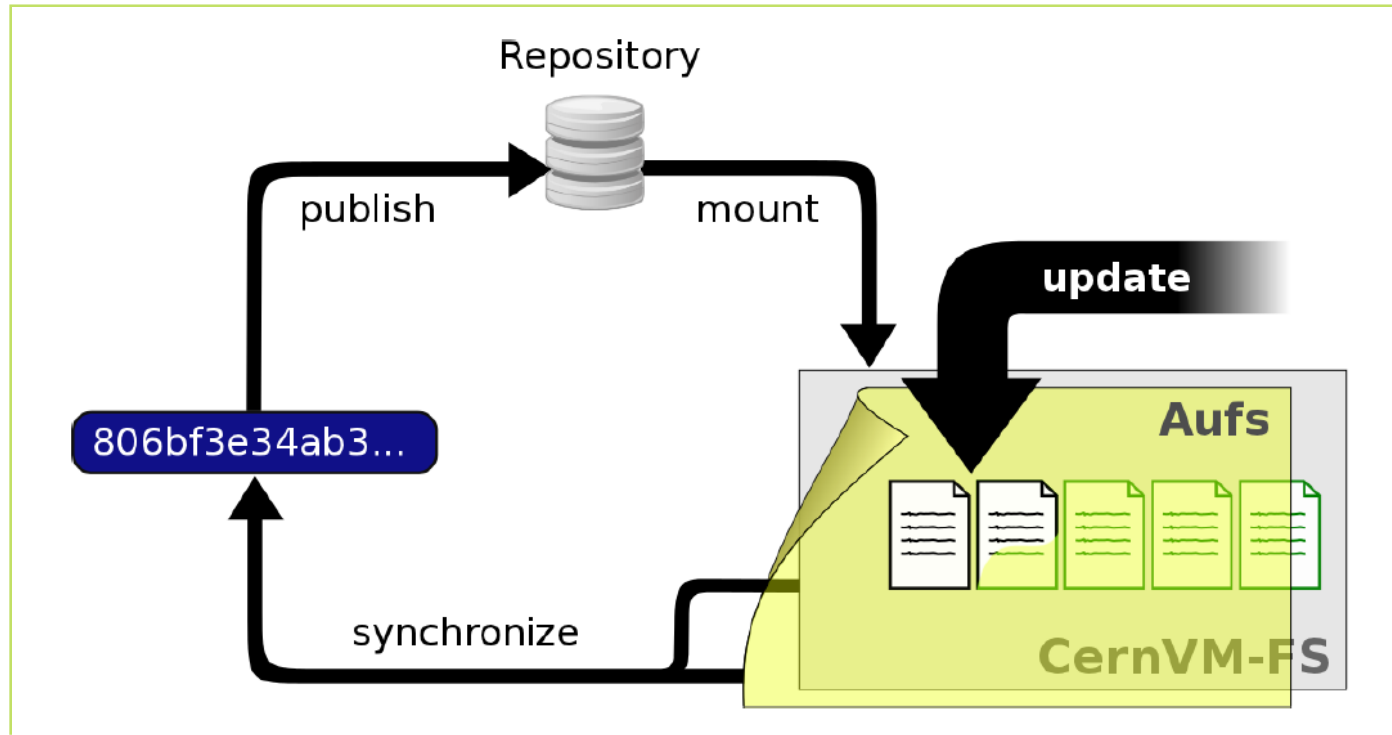




# CVMFS repository



- ❖ CVMFS repository is a form of content-addressable storage
- ❖ publish: create file catalog, compress data, calculate hash





# metadata



---

Field	Type
Path MD5	128 Bit Integer
Parent Path MD5	128 Bit Integer
Hardlinks	Integer
SHA1 Content Hash	160 Bit Integer
Size	Integer
Mode	Integer
Last Modified	Timestamp
Flags	Integer
Name	String
Symlink	String
uid	Integer
gid	Integer

---

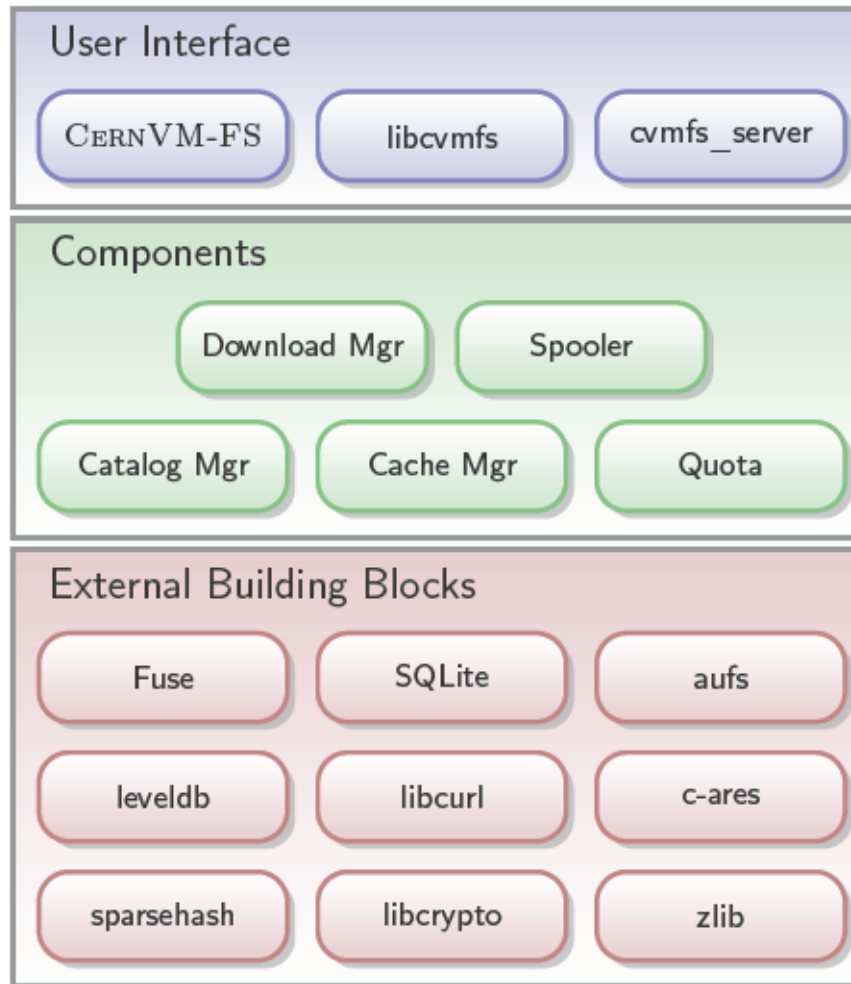
---

Flags	Meaning
1	Directory
2	Transition point to a nested catalog
33	Root directory of a nested catalog
3	Regular file
4	Symbolic link

---



# CVMFS building blocks





Part III

# HOW TO INSTALL, CONFIGURE AND DEBUG IT?



# Install stratum 0 server



## ❖ Install AUFS-enabled kernel

```
# wget https://ecsft.cern.ch/dist/cvmfs/cvmfs-release/cvmfs-release-2-5.noarch.rpm
```

```
# yum install cvmfs-release*.rpm
```

```
# yum -disablerepo="*" -enablerepo="cernvm-kernel" install kernel
```

```
# yum -enablerepo="cernvm-kernel" install aufs2-util
```

```
# reboot
```

## ❖ Install cvmfs packages

```
# yum install cvmfs cvmfs-server
```



# Creating a repository



- ❖ create a repository (bes3.ihep.ac.cn)
  - # cvmfs\_server mkfs bes3.ihep.ac.cn
- ❖ update the repository
  - # cvmfs-server transaction bes3.ihep.ac.cn
    - add/edit/delete files/directories in /cvmfs/bes3.ihep.ac.cn
  - # cvmfs\_server publish bes3.ihep.ac.cn
- ❖ if you want to clear all changes and start over again:
  - # cvmfs\_server abort bes3.ihep.ac.cn





# Paths in CVMFS server



- ❖ Some notable file paths in cvmfs server
- ❖ `/srv/cvmfs`, `/srv/cvmfs/<fqrn>`, `/var/spool/cvmfs` can be symlinked to another location before creating the repository

Path	Description
<code>/cvmfs</code>	repository mount point, read-only AUFS mountpoints
<code>/srv/cvmfs</code>	central repo. storage location
<code>/srv/cvmfs/&lt;fqrn&gt;</code>	storage location of a specific repo.
<code>/var/spool/cvmfs</code>	internal states of the repo.
<code>/etc/cvmfs</code>	configuration files and keychains



# Install client



- ❖ Install cvmfs packages:

```
# yum install cvmfs cvmfs-config-default
```

- ❖ create a file `/etc/cvmfs/default.local` with content:

```
CVMFS_REPOSITORIES=bes3.ihep.ac.cn
```

```
CVMFS_HTTP_PROXY=DIRECT
```

```
CVMFS_CACHE_BASE=/path/to/your/cache/dir
```

```
CVMFS_QUOTA_LIMIT=10240
```

- ❖ create file `/etc/cvmfs/config.d/bes3.ihep.ac.cn` with content:

```
CVMFS_SERVER_URL=http://your.server/cvmfs/bes3.ihep.ac.cn
```

```
CVMFS_PUBLIC_KEY=/etc/cvmfs/keys/bes3.ihep.ac.cn.pub
```



# Install client



- ❖ copy the key in CVMFS server to the same dir in client:

```
/etc/cvmfs/keys/bes3.ihep.ac.cn.pub
```

- ❖ setup

```
# setenforce 0
```

```
# cvmfs_config setup
```

```
# service autofs restart
```

- ❖ check if it can be mounted

```
# cvmfs_config probe
```



# Debugging



- ❖ Check if misconfiguration exist

```
# cvmfs_config chksetup
```

- ❖ show all the configuration parameters

```
# cvmfs_config showconfig bes3.ihep.ac.cn
```

- ❖ to exclude autofs/automount as a source of problem, try to mount repo. manually

```
# mkdir /mnt/cvmfs_test
```

```
# mount -t cvmfs bes3.ihep.ac.cn /mnt/cvmfs_test
```

- ❖ to exclude SELinux as a source of problem, disable it by

```
# setenforce 0
```



# Summary



- ❖ CVMFS is a good way to distributed HEP software in grid, cloud, and cluster sites over WAN
- ❖ It's a POSIX read-only network filesystem based on FUSE, HTTP, AUFS etc.
- ❖ You can try setting up a server and client on VMs.

Thank you!