

CMSSW on CvmFS

Bockjoo Kim

University of Florida

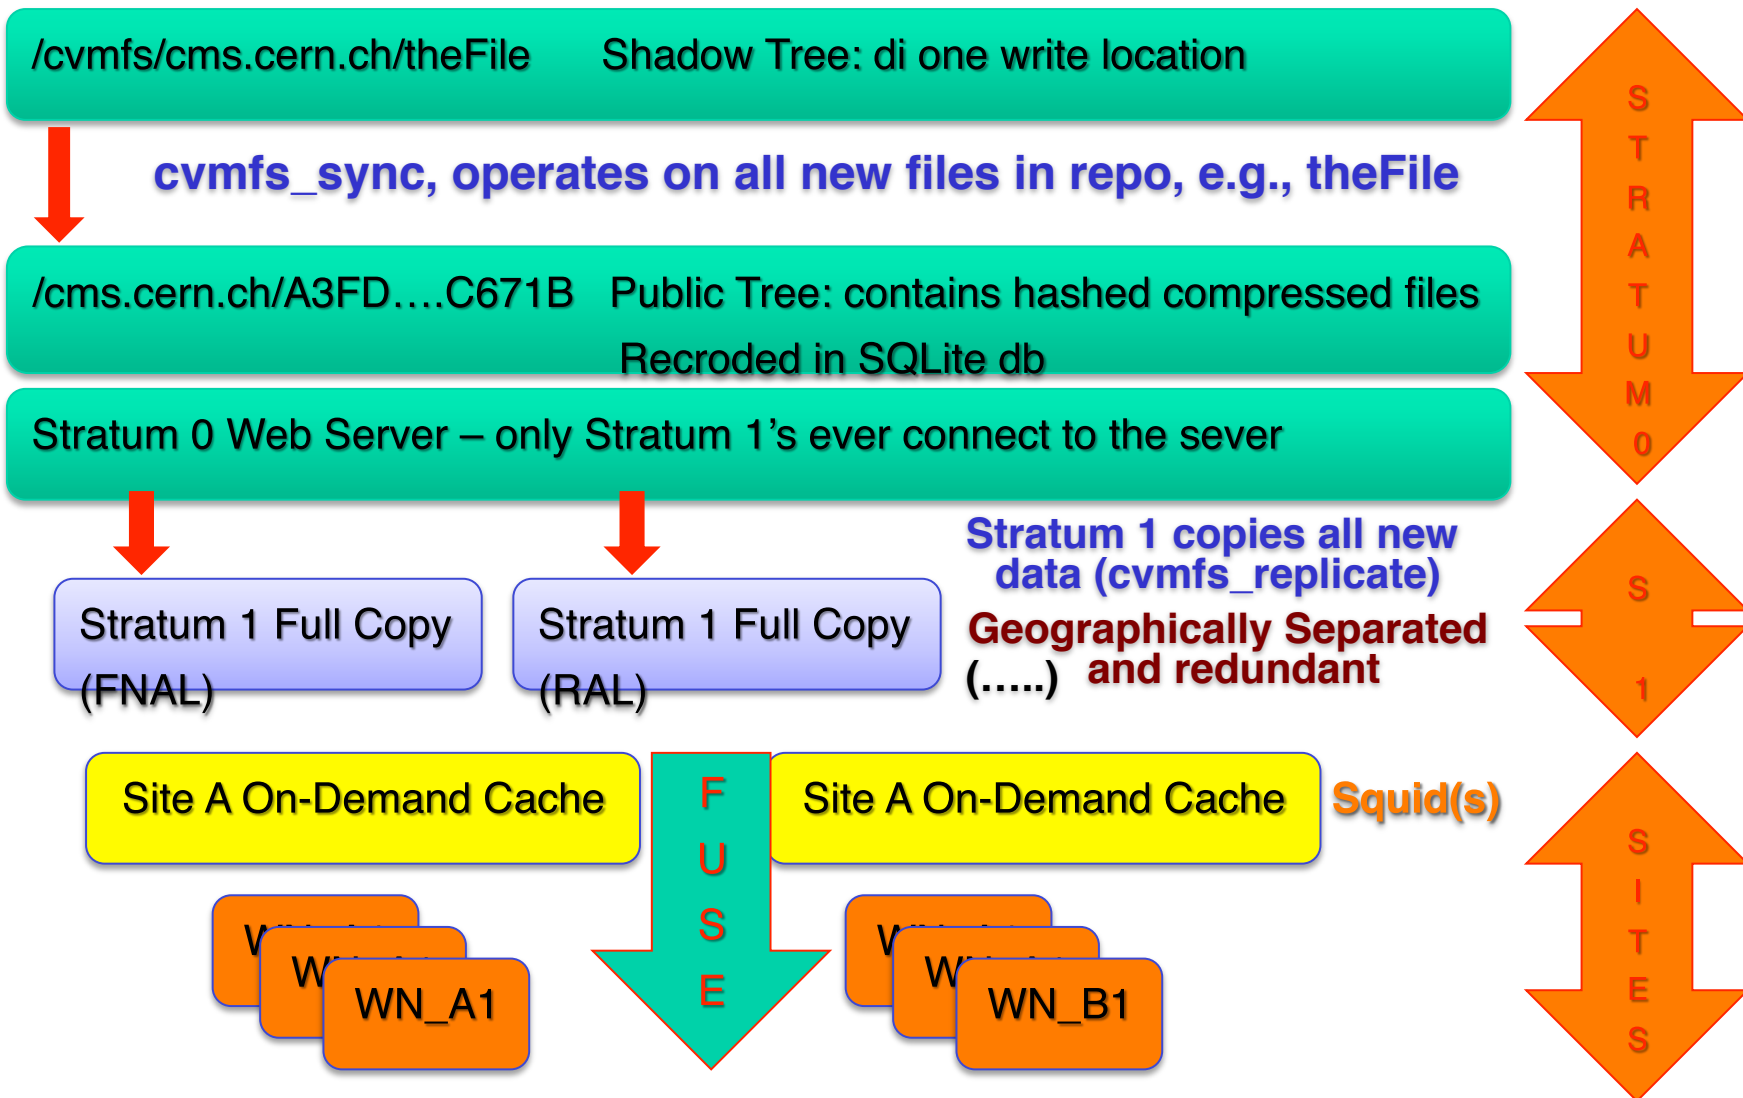
CMS Tier-2/3 Workshop

March 11, 2013

What is CvmFS?

- CvmFS: CERN Virtual Machine File System
- Developed for the CERN VM
 - Standalone software - VM not needed
 - A Network File System designed with SW dist. In mind: WORM
 - Addresses scalability and reliability
- CvmFS Design
 - HTTP – SQUID
 - fuse
 - Local cache
 - Content Addressable Storage vs Location Based Addressing
 - Deduplication (CAS) : Useful for env with fraction of files change
 - File integrity check on every file.

CvmFS File Publication/Delivery Chart



Why CvmFS?

- WLCG recommends it
- CMS intends to move to use CVMFS as the main GRID site software distribution mechanism
- Can utilize the existing squid(frontier)/fuse(hadoop) infrastructure
- Every site can use an identical CMSSW
- You can avoid publishing it on the BDII
- Cron based auto installation still available

CvmFS Client Installation (I)

- **RPMs:** <http://cernvm.cern.ch/portal/>
 - RHEL5/6 RPMs and Debian package
- **Installation Instruction Pages**
 - **OSG Instruction Twiki:**
 - <https://twiki.grid.iu.edu/bin/view/Documentation/Release3/InstallCvmfs>
 - **Glite/egee/emi sites:**
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/CernVMFS4cms>
 - **CMS OSG :**
<https://twiki.cern.ch/twiki/bin/view/CMS/CMSSoftDeployOSG>

CvmFS Client Installation (II)

- **Pre-requisite: fuse, autofs, account 'cvmfs', proper epel and yum repository setup**
- **Three different ways to install CvmFS client**
 1. **Installation with yum**
yum install cvmfs
 2. **Installation with rpm**
rpm -Uvh fuse*
rpm -Uvh cvmfs*
 3. **Installation with my script**
wget http://oo.ihepa.ufl.edu:8080/cmssoft/install_cvmfs_c.sh
sed -i "s|workdir=/root/RPMS/CVMFS|workdir=\$(pwd)|g" install_cvmfs_c.sh
./install_cvmfs_c.sh --help

CvmFS Client Configuration (I)

```
user_allow_other >> /etc/fuse.conf
```

fuse

```
/cvmfs /etc/auto.cvmfs >> /etc/auto.master
```

auto mount

```
CVMFS_REPOSITORIES=cms.cern.ch >> /etc/cvmfs/  
default.local
```

```
CVMFS_HTTP_PROXY=\"$squids\" >> /etc/cvmfs/default.local
```

```
export CMS_LOCAL_SITE=T2_US_Purdue >> /etc/cvmfs/  
default.local
```

```
CVMFS_QUOTA_LIMIT=10000 >> /etc/cvmfs/default.local  
(10GB)
```

```
CVMFS_.....
```



**CvmFS
Client
C
O
N
F
I
G**

CvmFS Client (2.1.6) Configuration (II) as an NFS Server

`user_allow_other >> /etc/fuse.conf`

fuse

~~`/cvmfs /etc/auto.cvmfs >> /etc/auto.master`~~

~~**auto mount**~~

Configure CvmFS client same as Configuration (I) for
CVMFS_REPOSITORIES, CVMFS_HTTP_PROXY, CMS_LOCAL_SITE, CVMFS_QUOTA_LIMIT

CvmFS Client

CvmFS Additional Configuration for NFS:

`CVMFS_NFS_SOURCE=yes >> /etc/cvmfs/default.local`
`CVMFS_MEMCACHE_SIZE=256 >> /etc/cvmfs/default.local`

Configure /etc/exports as an NFS Server, e.g.:

```
/cvmfs/cms.cern.ch 10.13.16.0/21(ro,sync,no_root_squash,no_subtree_check,fsid=101)
10.13.24.0/22(ro,sync,no_root_squash,no_subtree_check,fsid=101)
10.13.28.0/22(ro,sync,no_root_squash,no_subtree_check,fsid=101)
10.13.54.0/24(ro,sync,no_root_squash,no_subtree_check,fsid=101)
128.227.253.64/27(ro,sync,no_root_squash,no_subtree_check,fsid=101)
128.227.221.0/24(ro,sync,no_root_squash,no_subtree_check,fsid=101)
```

**NFS
Server
(RHEL6)**

Mount CvmFS on the NFS Clients (Worker Nodes), e.g.;

```
mount -t nfs -o nfsvers=3,noatime,ac,actimeo=60 ce:/cvmfs/cms.cern.ch /cvmfs/
cms.cern.ch
```

**NFS Client
RHEL5/RHEL6**

Configuration Files

- **/etc/cvmfs/default.conf** : global configuration
- **/etc/cvmfs/default.local** ← **Overrides default.conf**
- **/etc/cvmfs/domain.d/cern.ch.conf** : repo default
- **/etc/cvmfs/domain.d/cern.ch.local** : overrides above
- **/etc/cvmfs/config.d/cms.cern.ch.conf**: CMS specific
- **/etc/cvmfs/config.d/cms.cern.ch.local** : overrides above

Typical CvmFS Client Configuration

```
CVMFS_REPOSITORY_NAME=cms.cern.ch
CVMFS_USER=cvmfs # from /etc/cvmfs/default.conf
CVMFS_NFILES=65536 # from /etc/cvmfs/default.conf
CVMFS_CACHE_BASE=/var/lib/cvmfs # from /etc/cvmfs/default.conf
CVMFS_CACHE_DIR=/var/lib/cvmfs/shared # from
CVMFS_MOUNT_DIR=/cvmfs # from /etc/cvmfs/default.conf
CVMFS_QUOTA_LIMIT=10000 # from /etc/cvmfs/default.local
CVMFS_SERVER_URL=http://cvmfs1.fnal.gov:8000/opt/cms;http://cvmfs2.fnal.gov:8000/opt/cms;http://cvmfs-stratum-
one.cern.ch:8000/opt/cms # from /etc/cvmfs/domain.d/cern.ch.local
CVMFS_DEBUGLOG=
CVMFS_HTTP_PROXY=http://squid.ihepa.ufl.edu:3128|http://squid2.ihepa.ufl.edu:3128 # from /etc/cvmfs/default.local
CERNVM_GRID_UI_VERSION=
CVMFS_SYSLOG_LEVEL=
CVMFS_SYSLOG_FACILITY=
CVMFS_TRACEFILE=
CVMFS_TIMEOUT=5 # from /etc/cvmfs/default.conf
CVMFS_TIMEOUT_DIRECT=10 # from /etc/cvmfs/default.conf
CVMFS_PUBLIC_KEY=/etc/cvmfs/keys/cern.ch.pub:/etc/cvmfs/keys/cern-it1.cern.ch.pub:/etc/cvmfs/keys/cern-
it3.cern.ch.pub
CVMFS_MAX_TTL=
CVMFS_RELOAD_SOCKETS=/var/run/cvmfs # from /etc/cvmfs/default.conf
CVMFS_DEFAULT_DOMAIN=cern.ch # from /etc/cvmfs/default.conf
CVMFS_MEMCACHE_SIZE=
CVMFS_KCACHE_TIMEOUT=
CVMFS_ROOT_HASH=
CVMFS_REPOSITORIES=cms.cern.ch # from /etc/cvmfs/default.local
CVMFS_PROXY_RESET_AFTER=300 # from /etc/cvmfs/default.conf
CVMFS_MAX_RETRIES=1 # from /etc/cvmfs/default.conf
CVMFS_BACKOFF_INIT=2 # from /etc/cvmfs/default.conf
CVMFS_BACKOFF_MAX=10 # from /etc/cvmfs/default.conf
CVMFS_IGNORE_SIGNATURE=no # from /etc/cvmfs/default.conf
CVMFS_STRICT_MOUNT=no # from /etc/cvmfs/default.conf
CVMFS_SHARED_CACHE=yes # from /etc/cvmfs/default.conf
CVMFS_NFS_SOURCE=no # from /etc/cvmfs/default.conf
CVMFS_CHECK_PERMISSIONS=yes # from /etc/cvmfs/default.conf
CVMFS_AUTO_UPDATE=
```

Initial Debugging

- **Dump CvmFS client configuration**
 - cvmfs_config showconfig
 - source /etc/cvmfs/config.sh ; env | grep CVMFS_
- **Enable a log file**
 - CVMFS_DEBUGLOG=/tmp/cvmfs.log
- **Mount CvmFS outside the auto mounter**
 - mkdir /mnt/cvmfs
 - mount -t cvmfs /cvmfs/cms.cern.ch /mnt/cvmfs
- **Check /var/log/messages**
 - tail -f /var/log/messages

More CvmFS Command Tools

- **Configuration related tools**

 - `cvmfs_config`

 - `source /etc/cvmfs/config.sh` ← this provides various bash functions

- **Talking to the CvmFS Client to make it do something**

 - `cvmfs_talk`

- **CvmFS cache management command**

 - `cvmfs_fsck`

Migrating to CvmFS on OSG (I)

- **Turn \$OSG_APP/cmssoft/cms into CVMFS:**
 - Stop installation cron if any
 - `cd $OSG_APP/cmssoft ; ln -s /cvmfs/cms.cern.ch cms`
- **Configure SITECONF**
 - Choose `$OSG_APP/cmssoft/cms/SITECONF` to soft-link to, e.g., `/etc/cvmfs/SITECONF` (visible from all machines)
 - Add `export CMS_LOCAL_SITE=/etc/cvmfs/SITECONF` to `/etc/cvmfs/default.local`
 - SITECONF on CvmFS is updated every hour. Site can also rely on the SITECONF on CvmFS

Migrating to CvmFS on OSG (II)

- **Publish Installed CMSSWs for CRAB, etc**
 - See Step3 of “Five Steps to CVMFS migration” in <https://twiki.cern.ch/twiki/bin/view/CMS/CMSSoftDeployOSG>
- **Validation**
 - Watch swinst test status in the SUM page
 - More proactively, submit CRAB jobs or ask the prod team to run a few test jobs

Operations: Detecting Problems

- **CvmFS site problem with prod jobs in CMSSW**
 - **Cache or Automount**
 - **Common mistakes:**
 - CVMFS_QUOTA_LIMIT mismatch with disk
 - **Cache full**
- **A nice script to detect CvmFS cache issue (Wisconsin)**
`wget http://www.hep.wisc.edu/~dan/check_cvmfs`
`chmod a+x check_cvmfs`
`./check_cvmfs cms.cern.ch`
- **Also cvmfs_talk and cvmfs_fsck can be used**
`wget http://oo.ihepa.ufl.edu:8080/cmssoft/cmssoft_check_cvmfs.sh`
`chmod a+x cmssoft_check_cvmfs.sh`
`./cmssoft_check_cvmfs.sh`

CvmFS Sites

- T2_US_Nebraska
 - First CVMFS migrated site on OSG
- T2_US_Purdue
- T2_US_Wisconsin
- T2_US_Caltech (3/6/2013)
- T2_US_Vanderbilt
- T3_US_Baylor
 - Due to the issue with the glusterfs, they decided to migrate to the CVMFS
- CMS site count with CMSSW on CvmFS
 - 5 T1, 27 T2 (5 sites from OSG) and 4 T3 (one OSG site)

Contacts

- **CMSSW Central Installation**
 - christoph.wissing@desy.de (CMS contact)
 - Marko.Petek@cern.ch (main central installation)
 - bockjoo@phys.ufl.edu (US/Brazil/OSG support)
- **European EMI Sites**
 - Wolf.Behrenhoff@cern.ch
 - Stijn.Blyweert@cern.ch
 - michael.maes@cern.ch
 - christoph.wissing@desy.de
- **OSG Sites**
 - bockjoo@phys.ufl.edu
- **CvmFS Technical Questions:**
 - cvmfs-talk@cern.ch

Documentation and Support

- Main Twiki
 - <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CompOpsCVMFS>
- CVMFS Installation Twiki
 - CMS-OSG page
<https://twiki.cern.ch/twiki/bin/view/CMSPublic/CernVMFS4cms>
 - Wisconsin Server/Client Installation Instruction
www.hep.wisc.edu/cms/comp/ops/cvmfs.html
- CVMFS Migration on OSG
 - <https://twiki.cern.ch/twiki/bin/viewauth/CMS/CMSSoftDeployOSG>
- Savannah ticket (probably the best way to get support)
 - cmscompinfrasup-cvmfs

Summary

- CvmFS is designed for VO Software distribution in mind. WLCG recommends it and CMS might/will require it
- CvmFS client can be installed in two different modes. All CvmFS client or NFS exported CvmFS client. Installing/migrating aren't difficult (< 1 day should be enough)
- Main issues are cache size/corruption and automount. Sites might want to monitor these two popular issues proactively.
- So far, 36 sites have CMSSW on CvmFS