

CVMFS AT TIER2S

Sarah Williams
Indiana University

Motivation

2

- Integration Program goals
 - to test CVMFS operations at the scale of a Tier 2
 - eventually deploy at all the Tier 2 sites
- Expected Advantages
 - Scales better than other solutions
 - No more installer jobs
 - Easier to deploy new software
 - New releases instantly available

Worker Node Setup

3

- ❑ `/etc/init.d/autofs start`
- ❑ `cd /etc/yum.repos.d/ wget`
<http://cvmrepo.web.cern.ch/cvmrepo/yum/cernvm.repo>
- ❑ `cd /etc/pki/rpm-gpg/`
- ❑ `wget` <http://cvmrepo.web.cern.ch/cvmrepo/yum/RPM-GPG-KEY-CernVM>
- ❑ `yum -y install fuse cvmfs cvmfs-init-scripts`
- ❑ **Create** `/etc/cvmfs/local.d/default.conf` with contents:

```
CVMFS_REPOSITORIES=atlas
CVMFS_CACHE_DIR=/var/cache/cvmfs2
CVMFS_HTTP_PROXY="http://iut2-  
grid1.iu.edu:3128;http://cernvm.lbl.gov:3128"
```
- ❑ `/etc/init.d/cvmfs start`

Panda SchedConfig

4

- 'appdir' : '/opt/atlas_app/atlas_rel'
- 'environ' :
**'VO_ATLAS_SW_DIR=/opt/atlas_app/atlas_rel
APP=/osg/app TMP=/scratch DATA=/osg/data'**

Testing on MWT2_IU

5

- Most MWT2 job slots have been integrated into the MWT2_UC queue, to be renamed MWT2
- The leftover in MWT2_IU are used for the CVMFS test
- 140 job slots, 2 Dual Core AMD Opteron(tm) Processor 285, 8Gb RAM, 100Gb hard drive

Results from MWT2_IU

6

- Production running successfully
- After 24 hours of production, cache size is 2-4Gb. CVMFS cache can be limited with CVMFS_QUOTA_LIMIT in the config, or by using the cvmfs-talk utility
- Number of files in cache is 3-8K
- Cache on squid server is 3.2Gb, shared with Frontier
- After 72 hours of production, cache size is 3-5GB, 6-10K files in cache, and the cache on the squid server is 3.3Gb

Testing on Analy_MWT2_X

7

- Analy_MWT2_X is a test queue which uses the MWT2 integrated cluster
- 3811 Job slots on various hardware
- Storage for this test used the Xrootd global redirector testbed to access MWT2 storage
- Since this is an analysis queue, we use Hammercloud jobs for testing

Results from Analy_MWT2_X

8

- 1152 of 1152 jobs completed successfully
- Cache size is 0.5-2GB
- Number of files in cache is 2K
- Cache on squid server is 3.3Gb, shared with Frontier
- Efficiency was in line with previous tests using NFS for releases

Future Plans

9

- Asoka De Silva, Doug Benjamin, Alex Undrus and Alessandro De Salvo are working on a new version of CVMFS. Plans are to demonstrate for testing next month.
- Mount point will change to `/cvmfs/atlas.cern.ch/repo`
- Will include the Athena kits (including nightlies), DBReleases and conditions pool files, and ATLASLocalRootBase for Tier3
- Will have a defined setup process for Tier1/2/3
- Replica servers are in testing at BNL and RAL so that the CERN server is not a single point of failure

Questions?

10