

# CVMFS Introduction

Doug Benjamin

Duke University

# CVMFS in ATLAS

- CVMFS first came to ATLAS via the CERN VM project
- Next adopted by the ATLAS Tier 3 efforts
- Those at RAL saw the benefit and tested it at scale
- It has been adopted by other Tier 1/Tier 2 sites
  - RAL, QMUL, Wuppertal

# What you will hear in this session

- DB will present what we will make available via CVMFS
- Jakob Blomer (CVMFS lead developer) will teach us about CVMFS
- Alex will talk about the development for the nightlies
- Alessandro will talk about CVMFS and the grid
- Ian will talk about CVMFS deployment
- DB will present near term schedule and concluding remarks

# Current planned repository/directory structure for ATLAS CVMFS

VO\_ATLAS\_SW\_DIR=/cvmfs/atlas.cern.ch/repo/sw  
VO\_ATLAS\_CVMFS\_DIR="/cvmfs/atlas.cern.ch/repo"

Athena software will be installed in

\$VO\_ATLAS\_SW\_DIR/software/<version>

As of Today (6-April-11) for extra software we would have:

\$VO\_ATLAS\_SW\_DIR/cctools

\$VO\_ATLAS\_SW\_DIR/ddm

\$VO\_ATLAS\_SW\_DIR/atlas-gcc

\$VO\_ATLAS\_SW\_DIR/python (v2.5)

Note - other software can be added as needed.

\$VO\_ATLAS\_CVMFS\_DIR/ATLASLocalRootBase

\$VO\_ATLAS\_CVMFS\_DIR/conditions (for the conditions db files)

\$VO\_ATLAS\_SW\_DIR/database/DBRelease/<DB version> for DB Releases

\$VO\_ATLAS\_SW\_DIR/nightlies for the Nightly releases

# Actual directory structure & update frequency

directories would be:

/cvmfs/atlas.cern.ch/ATLASLocalRootBase

(**updated daily** - content mostly does not change, expect Grid CRLs)

/cvmfs/atlas.cern.ch/repo/conditions

(**updated several times per day**, new files created during data taking)

/cvmfs/atlas.cern.ch/repo/sw/database/DBRelease/<DB version>

(**updated periodically** - coupled to software releases)

/cvmfs/atlas.cern.ch/repo/sw/software/<version>

(**updated periodically** as new releases are created)

/cvmfs/atlas.cern.ch/sw/atlas-gcc

(**updated infrequently**)

/cvmfs/atlas.cern.ch/sw/cctools

(**updated infrequently**)

/cvmfs/atlas.cern.ch/repo/sw/ddm

(**updated infrequently**)

/cvmfs/atlas.cern.ch/sw/python

(**updated infrequently**)

/cvmfs/atlas.cern.ch/repo/sw/nightlies

(**updated nightly**)

/cvmfs/atlas.cern.ch/repo/sw/local/<Pool file Catalog>

(**updated several times per day** - see conditions directory)

Note- because of how cmt operates - we must install ATLAS software and ATLASLocalRootBase as/cvmfs/atlas.cern.ch/... This is because of the way pacman does the software installation. On the client machines the users will see the software located in /cvmfs/atlas.cern.ch/..

# Repositories

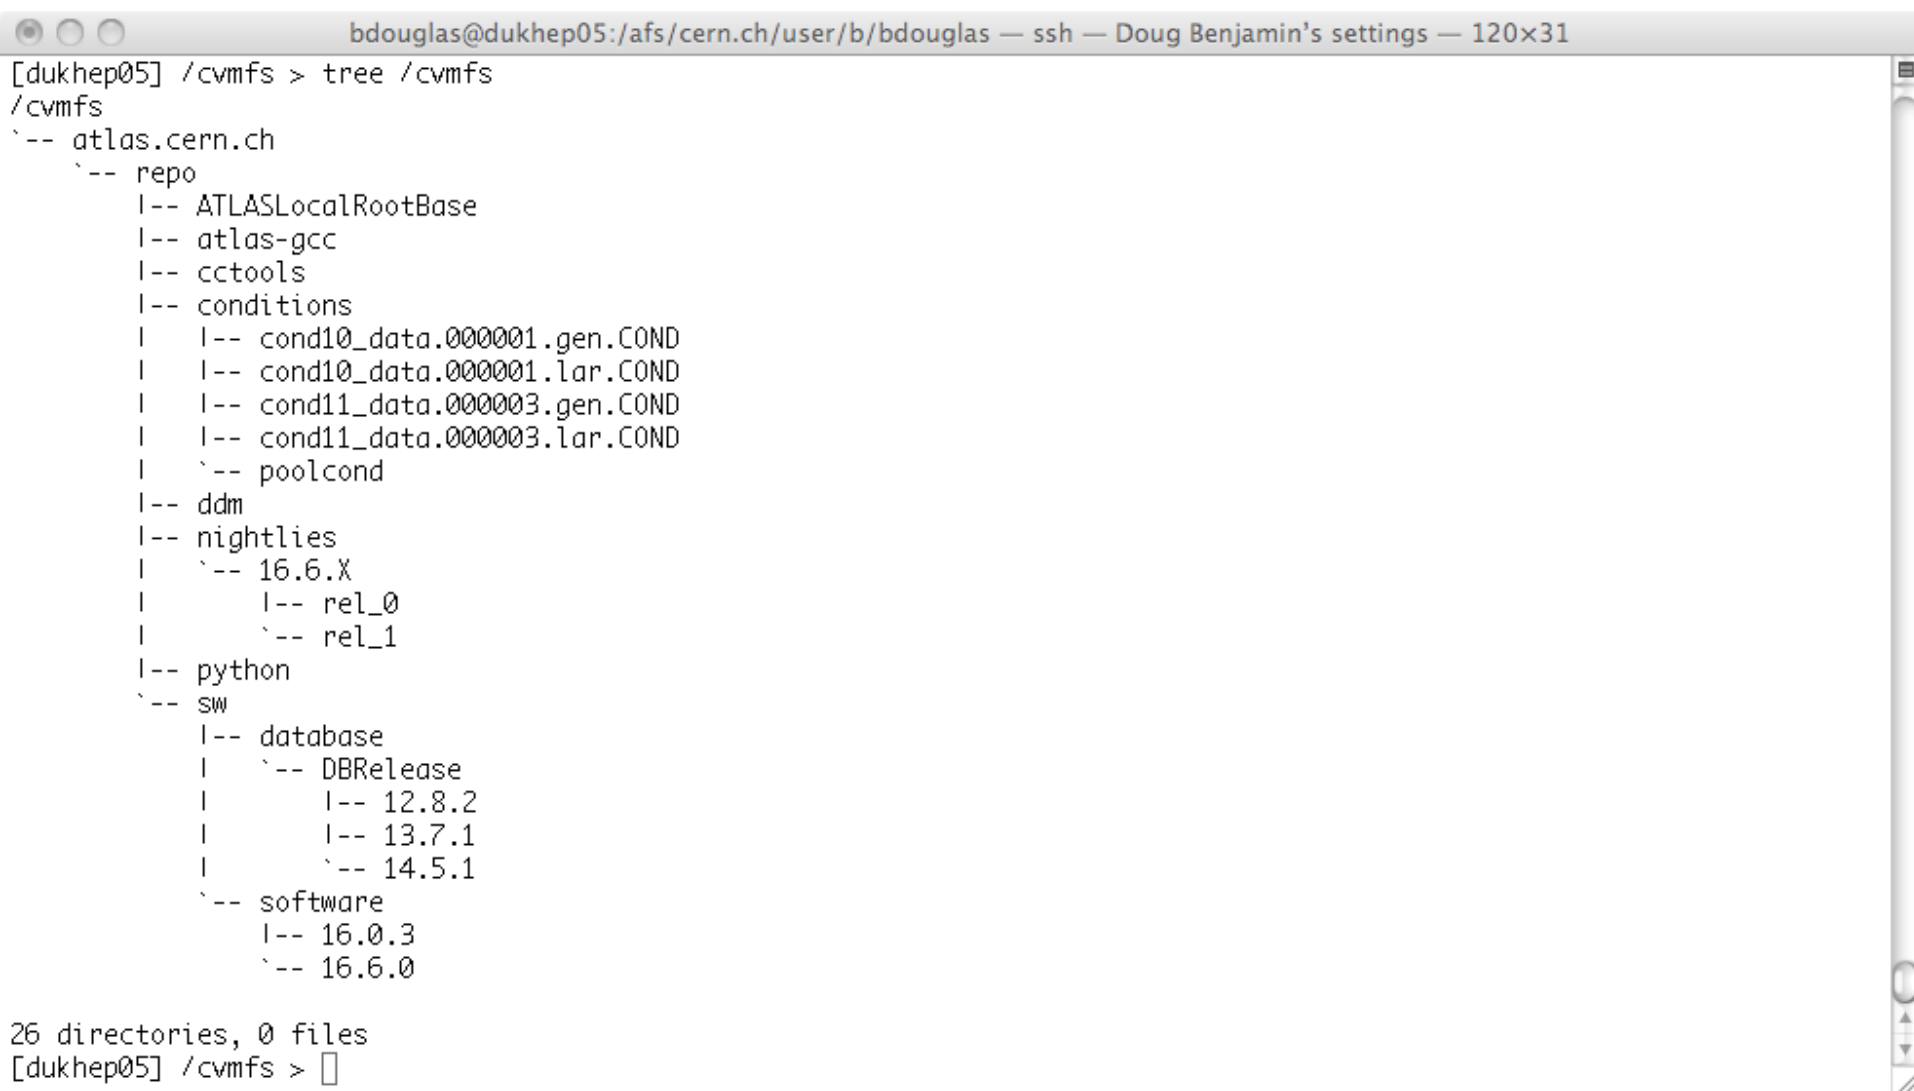
The main repository is:

`$VO_ATLAS_CVMFS_DIR`

The nested repositories are:

- `$VO_ATLAS_CVMFS_DIR/ATLASLocalRootBase`  
(`/cvmfs/atlas.cern.ch/ATLASLocalRootBase`)
- `$VO_ATLAS_CVMFS_DIR/conditions/<Conditions DB dataset>`  
(`/cvmfs/atlas.cern.ch/repo/conditions/<Conditions DB dataset>`)
- `$VO_ATLAS_SW_DIR/database/DBRelease/<version>`  
(`/cvmfs/atlas.cern.ch/repo/sw/database/DBRelease/<version>`)
- `$VO_ATLAS_SW_DIR/software/<version>`  
(`/cvmfs/atlas.cern.ch/repo/sw/software/<version>`)
- `$VO_ATLAS_SW_DIR/nightlies/<nightly branch>/<version>`  
(`/cvmfs/atlas.cern.ch/repo/sw/nightlies/<nightly branch>/<version>`)
  - Note each Nightly Branch is nested and within a branch each version (rel\_0, rel\_1... etc) is nested

# Tree view - some of the directories seen by users



```
bdouglas@dukhep05:/afs/cern.ch/user/b/bdouglas — ssh — Doug Benjamin's settings — 120x31
[dukhep05] /cvmfs > tree /cvmfs
/cvmfs
|-- atlas.cern.ch
    |-- repo
        |-- ATLASLocalRootBase
        |-- atlas-gcc
        |-- cctools
        |-- conditions
        |   |-- cond10_data.000001.gen.COND
        |   |-- cond10_data.000001.lar.COND
        |   |-- cond11_data.000003.gen.COND
        |   |-- cond11_data.000003.lar.COND
        |   |-- poolcond
        |-- ddm
        |-- nightlies
        |   |-- 16.6.X
        |       |-- rel_0
        |       |-- rel_1
        |-- python
        |-- sw
            |-- database
            |   |-- DBRelease
            |       |-- 12.8.2
            |       |-- 13.7.1
            |       |-- 14.5.1
            |-- software
                |-- 16.0.3
                |-- 16.6.0

26 directories, 0 files
[dukhep05] /cvmfs > 
```

NB – the total number of directories will be much larger

# Notes about negotiating this structure

- On Tier 3 sites - env variables are set and users do not need to know explicate paths
  - ATLASLocalRootBase is the magic
- Grid sites - pilots and env variables perform same function
- What about lxplus interactive nodes?
  - Since ATLASLocalRootBase is in cvmfs this can be part of the solution here