

CVMFS inode problem

Alastair Dewhurst, Rod Walker

The Problem?

- If the CVMFS inode counter goes above the 32 bit limit, 32 bit applications not using 64 bit file system calls will break (e.g., most ATLAS software).
- This is most likely to affect ATLAS as we have a very large repository which updates every few hours.
- Even then, it still requires continuously mounted cvmfs for multiple months (~ 80 days)
- As the number of slots of on machines grows, chances are CVMFS will stay mounted.



Diagnoses & Fix

- If the inode limit is reached on a machine it will start failing jobs. Will fail on whichever 32 bit application runs first (e.g., random errors).
- Shifters can use Panda link to spot flakey nodes. e.g,

http://pandamon.cern.ch/wnlist?site=RAL-LCG2_SL6&jobStatus=wn&days=0.50

- Sites can check each WN by doing:

```
# cvmfs_talk -i atlas.cern.ch internal affairs
```

```
Inode Generation:
```

```
init-catalog-revision: 3157 current-catalog-revision: 3870 incarnation: 2 inode generation: 4364500012
```

- Sites can reset the counter by doing:
`cvmfs_config umount.`
- This requires that no ATLAS jobs are running.



Conclusion

- Shifters should be aware of problem.
- A lot of sites upgraded to SL6 in September/October, if they haven't rebooted machines since, they could be in for an unpleasant Christmas present.
- Sites should take opportunity to make sure repository is unmounted and inode counter reset while ATLAS are not submitting large numbers of jobs.
- Jakob Bloomer plans to improve inode usage in 2.1.16 so that it would take a much longer period of time to hit the limit.

