

ATLAS pCache

Graeme Stewart
University of Glasgow

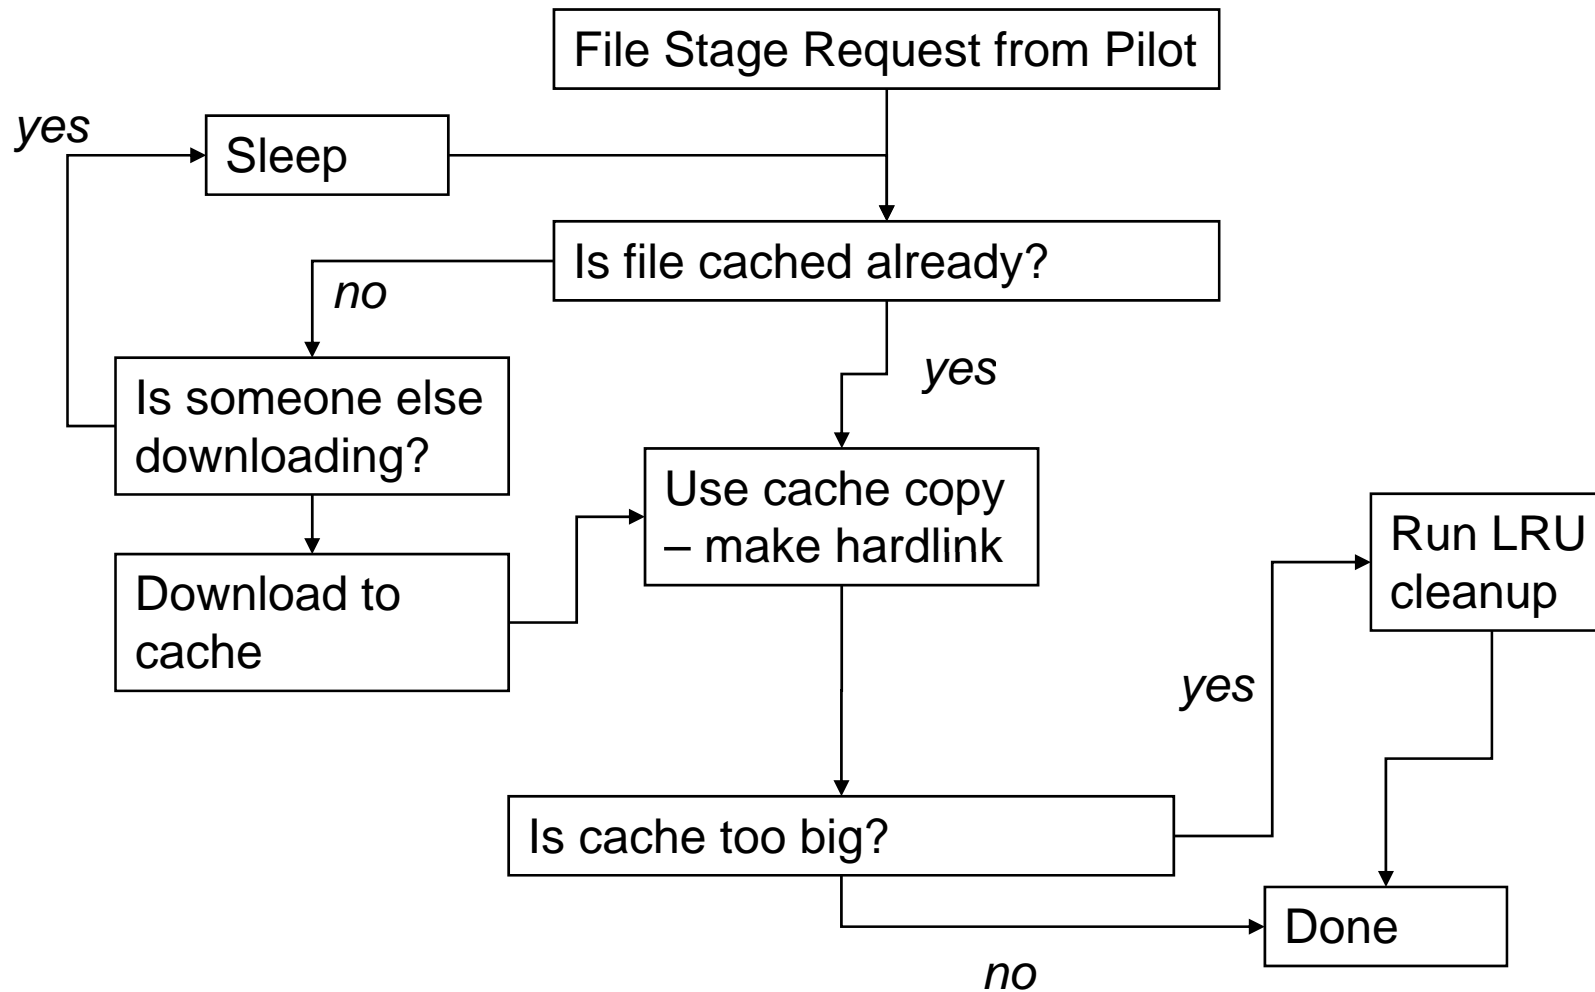
Reprocessing

- Reprocessing of combined cosmics runs over Christmas relied on a large SQLite file with conditions data
- This large CDRelease.tar.gz was downloaded to each WN and unpacked
- This file became very hot in the storage system and led to many job failures
 - Particular pain at RAL and NIKHEF

pCache

- We have a possible way to tackle hot files which might help sites
 - We can add a small cache wrapper to the production pilot which keeps a local cache on each WN
 - This can be as large or small as you like
 - You can safely delete the cache anytime
- (pCache was written by Charles Waldman and has been used at the U. of Chicago T2)

pCache Workflow



From the site point of view:

- You would need to tell us
 - Where to put the cache
 - How large it can be
 - This could be a %age fill on the worker node disk
 - Or an absolute value
 - Or both...
- You should not worry about deleting the cache if you have to
 - Though obviously if you have to often we should discuss the parameters above

Conclusion

- pCache should help overcome the hot files problem
 - It's nicely SE neutral and can be used on any site
 - It greatly eases the burden on your SE
 - With 8 core worker nodes the chances of cache hits are high
 - Of course it does leave files on the worker node
 - But these are useful files
 - And the system is well controlled
 - And what else is that WN 500GB disk doing anyway?
- 😊