




## DB TEG discussion guide

WLCG workshop, 20 May 2012

Dave Dykstra  
[dwd@fnal.gov](mailto:dwd@fnal.gov)

Fermilab is operated by the Fermi Research Alliance, LLC  
under contract No. DE-AC02-07CH11359 with the United States Department of Energy




# Squid Monitoring

- Site squid monitoring should be WLCG responsibility – multi-experiment, multi-app
  - WLCG Frontier monitoring, only basic
- Should we set up a task force to figure out how best to integrate with current WLCG monitoring?
- Don't want to give up detailed performance info now plotted with MRTG
  - Done by SNMP polls every 5 minutes
  - How to store information about address of each squid? ATLAS now storing in AGIS, CMS ad-hoc but audited against CVS checkins of site configs



# Squid Discovery

- Need standard method for jobs to discover squids at site
  - Agreement on an environment variable would be simplest, but maybe should use information system
  - Should distinguish between squids for production or opportunistic use
  - Squid addresses related to monitoring but not necessarily the same due to private networks and round-robins



# Hadoop HBase

- Hadoop HBase is the most promising of the NoSQL database technologies
  - Most scalable, robust, plus has powerful MapReduce features
  - Also has SQL add-on interface
  - CERN IT is setting up a cluster, who knows the details of that?