

Robinhood Trial

Ruth Klundt SNL
Monitoring Workshop
8/12/2015



*Exceptional
service
in the
national
interest*



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Overview

- Robinhood service collects info on lustre file systems 2 ways:
 - Scan, via the client mount
 - Subscription to changelog facility on lustre MDS
- Capabilities based on config options
 - Generate alerts and run arbitrary scripts based on events
 - Schedule reports
 - Define and schedule purging
 - Define file classes and custom behaviors for those classes
 - File Migration

Initial Results

- Original test setup:
 - 2.5.1 Robinhood
 - 2.1.x lustre servers and client
 - Mysql database
- Full scan of FS1 with ~130M files took ~4 days. FS2 with ~7M files took ~2 hours. This is due to both lustre client limitations and DB insertion rates.
- Using the robinhood database query tool, generating a list of file over X days old took 2-4 days. This is all DB access.

Tuning mysql DB

- Useful links for tuning
 - <https://raw.githubusercontent.com/major/MySQLTuner-perl/master/mysqltuner.pl>
 - <http://www.mysqlperformanceblog.com/2008/11/21/how-to-calculate-a-good-innodb-log-file-size>
- Managed to decrease scan time by half for our site, (some of these are now recommended on the robinhood site):
 - `innodb_file_per_table=1`
 - `innodb_buffer_pool_size=40G` # 50-90% memory
 - `innodb_flush_log_at_trx_commit=0`
 - `innodb_flush_method=O_DIRECT`
 - `innodb_log_file_size=512M`
 - `innodb_log_buffer_size=64M`
 - `innodb_thread_concurrency=24` # 2*number cores
 - `innodb_max_dirty_pages_pct=25`

Problems and Future Plans

- DB tuning did not help speed up report generation
 - Suspicion is that our namespaces are far deeper than what CEA may see – report tool must traverse those paths within the mysql DB in order to generate the report
 - Faster reports were the main motivation for trying this tool
- A variety of old bugs which have since been resolved caused us to take the service down and de-register the changelog users
- Lustre servers for the larger file systems are now running lustre 2.5.x, and performance improvements have been reported in Robinhood so we will give it another try – Real Soon Now