Har•mon•y: A Harness Monitoring System for the OLCF Test Harness

Cameron Kuchta, Reuben D. Budiardja, Verónica G. Vergara Larrea
Oak Ridge National Laboratory, Oak Ridge, TN

INTRODUCTION

Harmony is a monitoring system developed for the Oak Ridge Leadership Computing Facility (OLCF) Acceptance Test (AT) Harness. Summit was used as the target system as it was undergoing acceptance during the duration of this work.

OBJECTIVES

- **Monitor** tests executed by the OLCF AT harness and alert staff when failures occur.
- **Record** results in filesystem into to a database (MariaDB)
- **Analyze and report** results via a database-backed web interface.

BACKGROUND & MOTIVATION

- **Summit:** The OLCF's newest supercomputer, currently #1 in the TOP500 list (Nov. 2018). The system has a theoretical peak performance of ~200 PF and has 27,658 GPUs.
- **Acceptance Testing (AT):** complex/large-scale systems require rigorous testing to ensure its various components are functioning correctly before the system goes into production.
- **OLCF Acceptance Testing Harness:**
  - Developed in-house and used for the past three #1 TOP500 systems deployed.
  - Simulates a realistic workload by automatically building, submitting, and checking various tests from a set of applications selected by the OLCF.

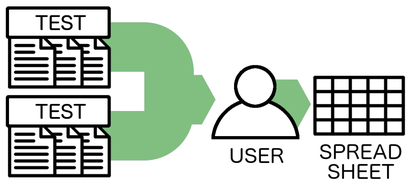

Fig. 1: *Failure classification method*

- **Harmony** provides solutions for the following center needs:
  - Gather all of the different results, spread in multiple directories across filesystem, in *query-able* format.
  - Check tests *manually* (Fig. 1) and record to spreadsheet → burdensome with large number of tests
  - A more automated way to categorize, sort, analyze, plot data (failures, success rate, number of jobs, runtime variability)

Harmony is Open Source

Source code and documentation: https://github.com/olcf/harmony

Part I: Monitoring Tests & Notification System

- Uses the LSF (job scheduler) API to query the status of jobs in the queue (Fig. 2):

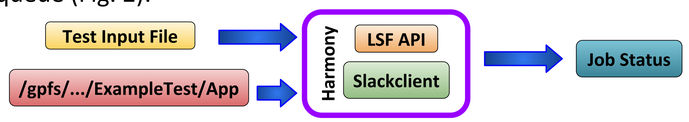

Fig. 2*. Harmony structure*

- Alerts test admin for any *missing test* not in the queue (Fig. 3).
  - Using Slack messages with Slack API.
  - Using email messages.
- Uses Slack Bot to communicate with machine's queuing system
  - Lets test admin check on jobs easily without having to log in

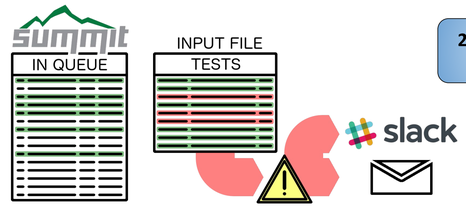

Fig. 3: *Monitoring Summit's LSF queue*

Part II: Recording Results

Harmony records data following these steps:

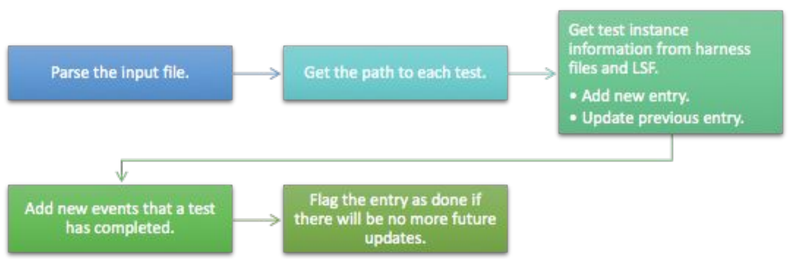

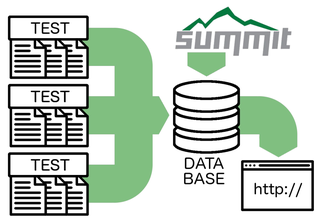

Fig. 4: *Harmony Recording and Reporting Workflow*

Part III: Reporting

- Web interface reports stored test information to users for analysis.
  - Django, HTML, CSS, and JavaScript.
  - Contains List pages and Detail pages (Fig. 5).
  - Follows the same hierarchy as the OLCF test harness as shown (Fig. 6).

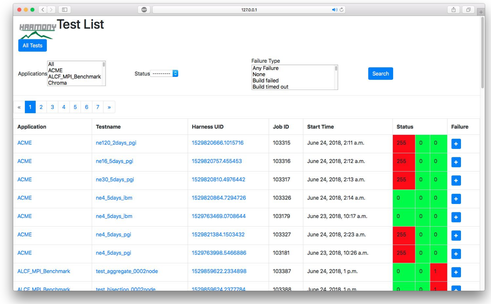

Fig. 5: *Test monitoring via website*

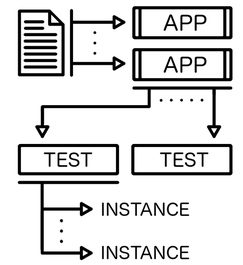

Fig. 6: *File system hierarchy*

- Pages inherit from generic templates for uniformity.
- AT requires that all test failures have an explanation. A pop-up window was added to the web interface to easily enter the type of failure and a description of what occurred.

CONCLUSIONS & FUTURE WORK

- **Harmony** is modular solution to that addresses the need to:
  - automate monitoring and alerting
  - automate recording
  - accurately reporting results from test harness runs on large-scale systems like Summit.
- **Harmony** complements the OLCF Test Harness to execute regression and acceptance tests on diverse OLCF systems.
- Future work includes:
  - Adding support for other job schedulers
  - Improving analysis capability

ACKNOWLEDGEMENTS

- We thank Don Maxwell, Jason Kincl, Arnold Tharrington, and Wayne Joubert for their contribution to this work.
- This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.
- This research was supported in part by an appointment to the Oak Ridge National Laboratory Oak Ridge Science Semester Program sponsored by the U.S. Department of Energy and administered by the Oak Ridge Institute for Science and Education.