Ref paper: dx.doi.org/10.1145/2834996.2834998

version 2015-05
as of 13 Nov 2015

TALES FROM THE TRENCHES: CAN USER SUPPORT TOOLS MAKE A DIFFERENCE?

Robert McLay[1]

Doug James[1], Si Liu[1], Todd Evans[1], Bill Barth[1], Antia Lamas-Linares[1], Reuben Budiardja[2], Mark Fahey[3]

[1]Texas Advanced Computing Center (TACC)
[2]National Institute for Computational Sciences (NICS)
[3]Argonne National Laboratory (ANL)

HUST 2015
20 Nov 2015
Austin, TX

REFLECTS OUR COLLECTIVE EXPERIENCE SUPPORTING STAMPEDE USERS

Stampede: TACC flagship and XSEDE workhorse

- Delivering 2 million core-hours a day to thousands of active users from hundreds of institutions
- 6 million jobs and 2 billion core-hours in its first 33 months of production service
- Both high impact and representative of larger community's needs and experience

One of 6,400 compute nodes on the 9+ petaflop Stampede system at TACC

OVERVIEW

Support Tools

Case Studies

Ad Hoc Observations

SUPPORT TOOLS

XALT (Fahey, McLay) – Job-Level Usage Data

- Job-Level: build-time and run-time data on essentially every job
- Usage: what software and libraries do researchers actually use?
- Of particular value to decision-makers and support staff
- Separate HUST 2015 talk

Download: github.com/Fahey-McLay/xalt

```
| intel/13.0.2.146                          | 3520152 |
| mvapich2/1.9a2(intel/13.0.079)            | 1170548 |
| intel/14.0.1.106                          |  611698 |
| intel/15.0.2                              |  543938 |
...
| impi/4.1.3.049(intel/13.0.079)            |  167063 |
| gc/4.9.1                                  |  133638 |
| impi/4.1.0.030(intel/13.0.079)            |  103845 |
| mvapich2/2.1(intel/15.0.2)                |   98403 |
| phdf5/1.8.9(intel/13.0.079:mvapich2/1.9a2)|   65302 |
| cuda/5.5                                  |   52962 |
| mvapich2/2.1(gcc/4.9.1)                   |   48891 |
| fftw3/3.3.2(intel/13.0.079:mvapich2/1.9a2)|   48802 |
| boost/1.55.0(intel/13.0.079)              |   45351 |
| boost/1.55.0(gcc/4.7.1)                   |   41794 |
| opencv/2.4.6.1(intel/13.0.079)            |   39791 |
| qt/4.8.4                                  |   33427 |
```

TACC Stats (Barth, Evans, Hammond) – Job-Level Performance Data

- Job-Level: performance snapshots on essentially every job
- Automatic filters generate daily reports on jobs that may require attention
- Web access to integrated dashboard view of any job

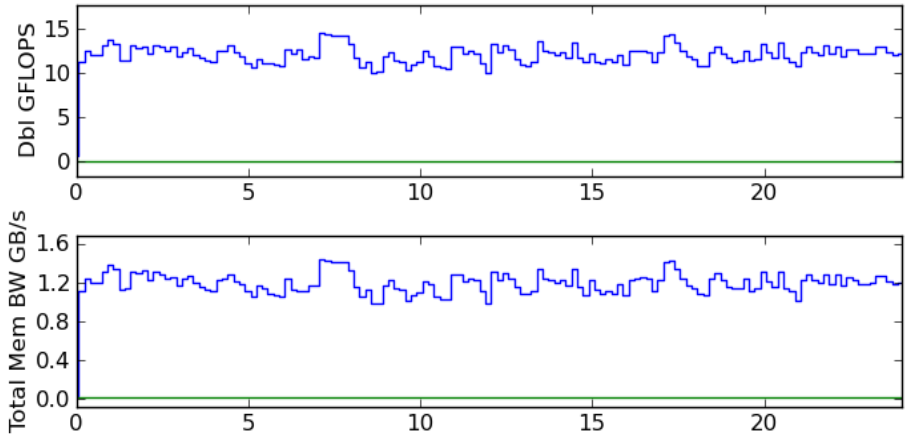

Download: github.com/TACC/tacc_stats

Lmod (McLay) -- User Ownership of the Environment

- Control: load, manage, use the software I need to do my research
- Repeatability: robust mechanisms to replicate software environment
- Visibility: know what's available under what conditions
- Protection: detect, prevent, eliminate common mistakes

```
TACC: Starting up job 5577472
TACC: Setting up parallel environment for MVAPICH2+mpispawn.
*******************************************************
WARNING: Your MPI Environment is : mvapich2/2.0b
Your executable was built with: impi/5.0.2
*******************************************************
*******************************************************
WARNING: Your Compiler Environment is : intel/14.0.1.106
Your executable was built with: intel/15.0.2
*******************************************************
TACC: Starting parallel tasks...
Fatal error in MPI_Init: Other MPI error, error stack:
MPIR_Init_thread(784).................:
MPID_Init(1323).......................: channel initialization failed
MPIDI_CH3_Init(141)...................:
dapl_rc_setup_all_connections_20(1386): generic failure with errno = 872609295
getConnInfoKVS(849)...................: PMI_KVS_Get failed
```

Download: sourceforge.net/projects/lmod

Sanity Tool (Liu, McLay, Wilson) – Account-Level Diagnostics

- Push-button assessment of account configuration
- Saves time and serves as general health screening: repeatable checks for common problems that are often notoriously hard to diagnose
- Evolves to reflect lessons learned

Download: sourceforge.net/projects/sanitytool

```
login1 62$ sanitycheck
Sanitytool Version: 1.1

     1: Check SSH permissions:
      Failed
      Error: group permission on $HOME will cause RSA to fail!
      Error: other permission on $HOME will cause RSA to fail!

      Make sure you have a .ssh directory under your $HOME directory.
      You can use the following commands to set the proper permissions:
      $ chmod 700 $HOME #(750 and 755 are also acceptable)
      $ chmod 700 $HOME/.ssh
      $ chmod 600 $HOME/.ssh/authorized_keys
      $ chmod 600 $HOME/.ssh/id_rsa
      $ chmod 644 $HOME/.ssh/id_rsa.pub
     2: Check SSH Key:
      Passed
     3: Check environment variables (e.g. HOME, WORK) and file system access:
      Passed
     4: Check user's queue accessibility (Stampede Only):
      Passed
     5: Check allocation balance:
      Passed
     6: Check quota for $HOME and $WORK:
      Passed
     7: Check module environment:
      Passed
     8: Check compilers:
      Passed
     9: Check scheduler commands:
      Passed

-------------------------------
1/9 tests failed
-------------------------------
login1 63$
```

Tips Database (McLay) – Spontaneous Hints

- A single random "tip of the day" appears on each user's console at the end of the login sequence
- Tips include command-line basics (Bash, Linux, editors, modules, etc.), best practices (change management, backups, job submission, etc.), system-specific info (e.g. Xeon Phi basics)

```
Tip 197 (See "module help tacc_tips" for
          features or how to disable)

  If you don't need 48 hrs for your job, don't
  request that much: the job scheduler will have
  an easier time finding a slot for the 1 hour
  you really need.

        ...or...

Tip 4 (See "module help tacc_tips" for
         features or how to disable)

  Use "!!" to repeat the most recent command.
  Use "!mpi" to repeat the most recent command
  that started with "mpi" and "!?mpi?" for the
  most recent one that contained "mpi".
```

Email TACC for access to source code

Synergy through Integration

- Optional Interfaces
  - XALT supplies meta-data for TACC Stats
  - XALT prototype supplies usage data to XSEDE XDMoD dashboards
  - Lmod supplies module data to XALT
- Related Functionality
  - Lmod produces usage data (as in XALT)
  - XALT will soon check build-vs-run consistency (as in Lmod)
- Cross-Pollination
  - XALT syslog implementation informing Lmod migration
  - Tips database highlights other tools

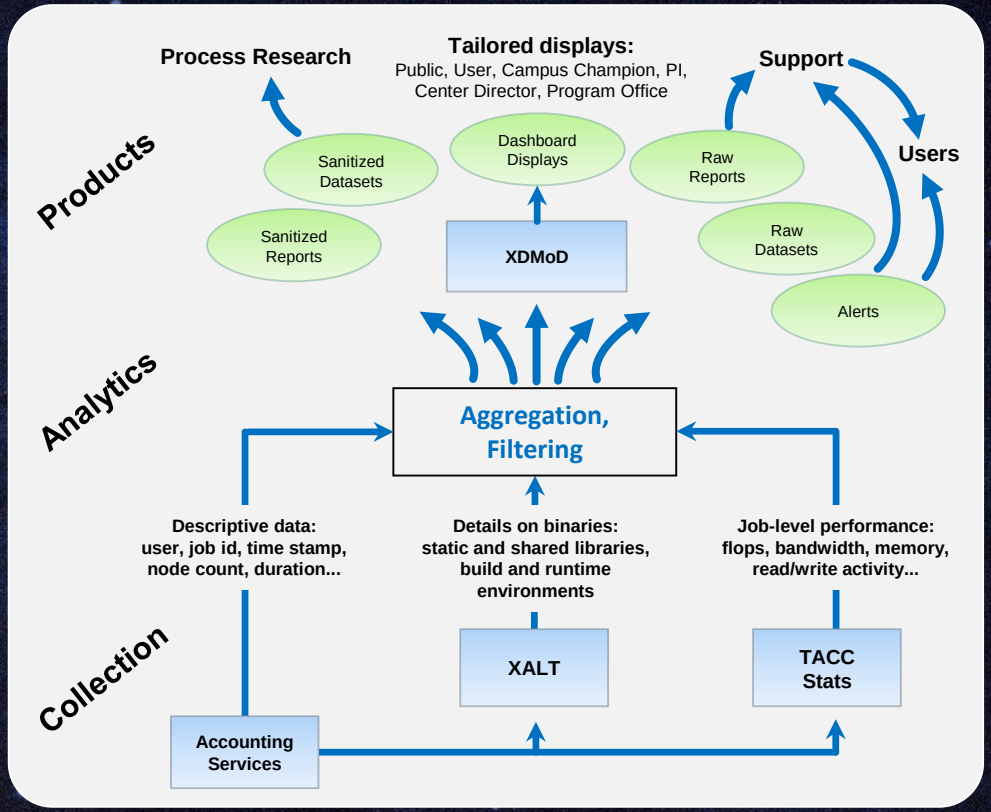

CASE STUDIES

Queue Wait Times: Targeted Outreach

- Dramatic increase in demand during 2015
  - Normal queue wait time went from 2-3 hrs to 24+ hrs
  - Largemem queue wait time reached 4+ days
- Scheduler tuning helped, but usage data was revealing
  - High percentage of jobs specified substantially more time than actually needed
  - XALT suggested opportunities to move some largemem users to normal queues by reducing tasks/node
- Interventions: steady, targeted outreach (email and tickets), tips database
- Results: demand's still up, but queue times are down from their peaks (normal queue down 50%; largemem 75%)

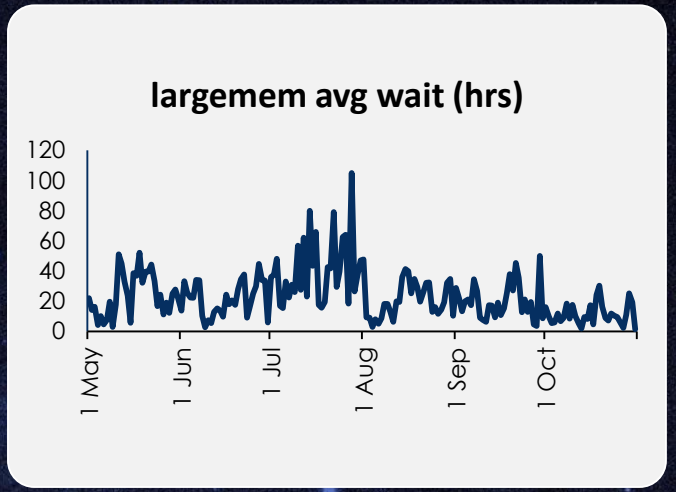

NumberOfJobs	user	queue	TotalSUs
29	userA	largemem	17522
3	userB	largemem	4100
8	userC	largemem	2158
4	userD	largemem	1507
19	userE	largemem	995
6	userF	largemem	716
2	userG	largemem	232
2	userH	largemem	98
4	userI	largemem	23
1	userJ	largemem	1

Compiler Migration: Managing Change

- Original plan: remove four very old compilers and associated software stacks; install new one
- Usage data: three of four had light usage; oldest had 2000+ users
- Interventions: change in plan, targeted outreach, nag messages

```
login1 69$ module load intel/13.0.079

------------------------------------------------------------
There are messages associated with the following module(s):
------------------------------------------------------------

intel/13.0.079:
   This module is deprecated. It is marked for removal in
   Dec 2015. Please transition to a newer Intel compiler
   (intel/15.0.2 recommended) before this occurs.

------------------------------------------------------------
```

Help Desk: Recurring Issues

- Recurring ticket: jobs fail to launch (mpi_spawn errors)
- ssh-related root causes difficult to diagnose unless you know what to look for
- Intervention: automation (job submission, account diagnostics)
- Recurring errors are the motivation behind both Sanity Tool and the build-vs-runtime consistency checks in Lmod/XALT.

```
login1 62$ sanitycheck
Sanitytool Version: 1.1

     1: Check SSH permissions:
     Failed
     Error: group permission on $HOME will cause RSA to fail!
     Error: other permission on $HOME will cause RSA to fail!

     Make sure you have a .ssh directory under your $HOME directory.
     You can use the following commands to set the proper permissions:
     $ chmod 700 $HOME #(750 and 755 are also acceptable)
     $ chmod 700 $HOME/.ssh
     $ chmod 600 $HOME/.ssh/authorized_keys
     $ chmod 600 $HOME/.ssh/id_rsa
     $ chmod 644 $HOME/.ssh/id_rsa.pub
     2: Check SSH Key:
     Passed
```

Job-Level Performance Data: Open-Ended Discovery

- Collect first and see what we learn
- Certain recurring patterns jumped off the page
- Led to auto-detection filters and daily reports: common problems, need for intervention
- Also learned to turn to job-level and user-level integrated dashboards before first response to a ticket

Sanitized (user names hidden)

Data-Driven Design: Benchmarking

- Analyzed 55,000 jobs with XALT and TACC Stats looking for general characteristics that could affect design decisions for next-generation systems
- Top codes cluster naturally in ways that amount to performance signatures
- Selected four representative codes as cornerstone of real-world benchmark suite; will begin to validate using Lonestar 5
- Issues remain: e.g. licenses that make it difficult to publish performance results

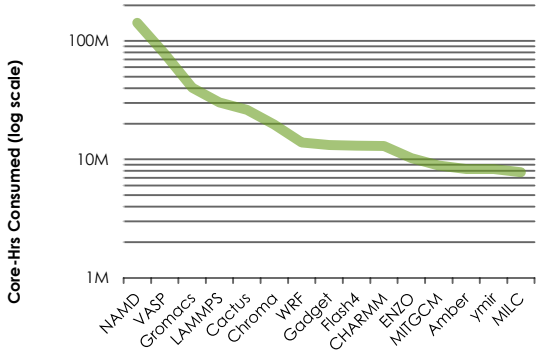

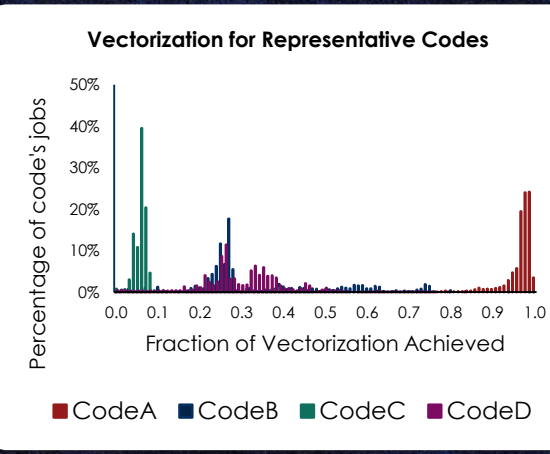

AD HOC OBSERVATIONS

OUR WORLDVIEW

- Labor hours are more precious than core-hours
- Support needs are not proportional to core-hours
- Formal training may not move the needle
- There's too much actionable data
- Urgent need: put tools and data in the hands of users
- Integration is power
- Targeted outreach gets users' attention

Ref paper: dx.doi.org/10.1145/2834996.2834998

SUMMARY

- User support tools can and do make a difference
- Support at scale requires...
  - Targeting small allocations, inexperienced users
  - Putting data and tools in the hands of users
  - Putting a priority on labor costs
- We gratefully acknowledge NSF support through grants ACI 10-53575 (XSEDE), ACI-1134872(Stampede), 1339708 (XALT), and ACI-1203560 (TACC Stats)

Ref paper: dx.doi.org/10.1145/2834996.2834998

Robert McLay
mclay@tacc.utexas.edu

For more information:
www.tacc.utexas.edu

Ref paper: dx.doi.org/10.1145/2834996.2834998

LICENSE