

Sierra System Status

LC User Meeting

2019-03-28

John Gyllenhaal

 Lawrence Livermore
National Laboratory



LLNL-PRES-770834

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

We want you on SIERRA ASAP! (With some caveats)

- Preparing to support the ATCC-7 campaign on SIERRA
 - Campaign started this week for other AT systems (Sequoia, Trinity)
- Focus on code developers and ATCC-7 participants
 - Please request accounts and start building/testing if haven't already
 - LLNL users can use the Inlexec bank to prepare
 - Help us find problems early, especially blockers for your team
- ATCC-7 scheduled to start when SIERRA goes GA
 - We are driving towards ~April 17th however 2 week delay possible
 - Potential roadblock discovered yesterday testing latest software
 - If need new software drop to fix, may push GA back to May 1st
 - Sierra will not be perfect when it goes GA

SIERRA awesome but not perfect!

Known issues as of 3/28/19

- Draining nodes can cause next job launches to fail
 - “Fixed” version delivered yesterday, unexpected 50 second delay
 - So might be May before we can get usable fix installed (GA blocker)
- Regression tests, 100s of bsubs, can cause nodes to drain
 - Several WCI codes cause scheduler daemon to wedge in kernel
 - Epilogue drains node, triggers issue above, killing unrelated jobs
 - Still trying to generate non-export controlled reproducer for IBM
 - Running multiple jsruns in one bsub minimizes impact but not perfect
- Ssh keys with passphrases not supported (long term)
 - If have ssh keys with passphrases, recommend moving aside
 - bsub and lalloc now automatically create ssh keys for you if needed
 - If use to access git, will need to submit new public key to git

SIERRA awesome but not perfect!

Known issues as of 3/28/19 continued

- We cannot yet detect bad GPUs as well as your codes can!
 - We continue adding checks as we find quick tests/diagnostics
 - The same GPU failing in different runs makes good suspects
 - Please tell us the node name, GPU id, and times/dates of failures
 - Allows us to get more details out of node logs around timestamp
 - We are happy to drain small numbers of suspect nodes
- SIERRA does not deal well with bad nodes and error states
 - Hard! Slow improvement in tolerance over the lifetime of SIERRA
 - We continue to improve our scanning for detecting bad nodes
- Mixing XL and CLANG OpenMP GPU code link issues
 - Working with IBM to get superset runtime to enable mixed linking
 - No one is known to be currently blocked on this, caused by new features
- XL partial C++14 support disabled by default
 - Have debug flag workaround but working on general solution

SIERRA awesome but not perfect!

Known issues as of 3/28/19 cont 3

- MPI executables with read-bit disabled breaks MPI
 - Several teams mark their public apps as execute only
 - Linux then disables key functionality MPI uses extensively (CMA)
 - To prevent loophole for reading executable anyway
 - We have asked IBM for option to avoid CMA in MPI
 - May lower MPI performance (but running is infinite improvement)
 - Don't have timeline yet (but IBM just agreed to do it 3/28)

New features coming soon to SIERRA

- CUDA 10.1 support expected around May 1st
 - Testing on test system butte started yesterday, 3/27
 - Huge IBM software drop due ~April 19th, includes CUDA 10.1 support
 - Cannot roll out new XL fixes until CUDA 10.1 supported in kernel
- Beta gcc 8.2 support expected mid-May on CORAL systems
 - Once we have it working, we will ask codes to test it out
 - Targeting default switchover in mid Summer if things go well
- mpibind/lrun support for SMT 2, 3, and 4 expected in April
 - Edgar Leon has implementation that we are starting to test
 - Automatically scales up to 2, 3, or 4 threads per task based on OMP_NUM_THREADS setting (i.e., 10, 20, 30 or 40 for 10 cores)
- SPINDLE support for not hammering NFS during big runs (soon)
 - We have hand-crafted workarounds if doing full system runs now

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.



**Lawrence Livermore
National Laboratory**