

Exceptional service in the national interest



An Introduction to DevSim

Anthony M. Agelastos, Ryan P. Shaw, Joel O. Stevenson

Org. 9326: Scientific Applications & User Support



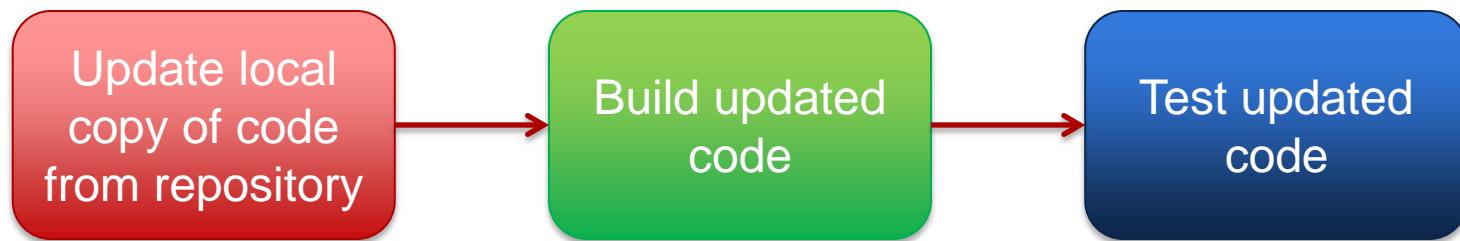
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2014-2516P

Executive Summary

- The objective is to understand and simulate representative developer workflows and assess their bi-directional impact with computing platforms.
- Analysis:
 - A typical SNL developer workflow includes 3 steps:
 - update a local codebase with updates from others (e.g., manually or with a software configuration management (SCM) tool like Git or CVS)
 - build the code and, if applicable, relevant documentation
 - test the code
 - Scripts and processes, collectively named DevSim (Development Simulator), were created to simulate the 3 steps listed above for ASC Sierra and ASC CTH.
- Results:
 - DevSim is currently used for building and testing by some developers.
 - DevSim was used to help CEE assess impact by changing file systems.

Developer Workflow

- SCENARIO:
 - A developer is ready to push their code update to the repository.
 - They will need to first pull in any changes others have been made since the local code base has been updated from the repository, build this updated code base which still contains the update, and then test it.
 - If this is successful, then their update can be pushed to the repository.



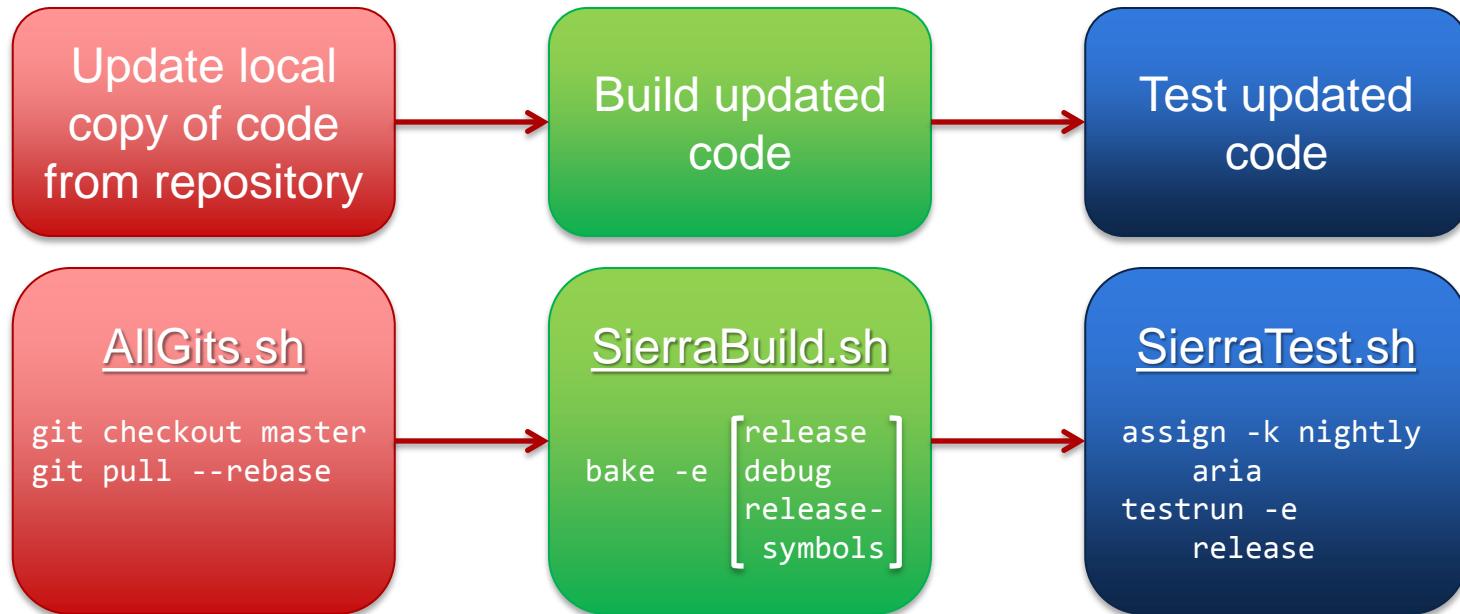
- GOALS:
 - Identify code(s) representative of typical scientific development at Sandia National Laboratories.
 - Create scripts and/or programs that encapsulate each piece of the above workflow for the identified code(s).

Chosen Simulation Codes

	ASC Sierra	ASC CTH
SCM	Git (x5)	CVS (x1)
Total Size	~200 GB	~5 GB
Build System	Boost BJam	GNU Make
Primary Languages (Excluding Third- Party Libraries)	C++	C, Fortran
No. of Primary Applications	6	1

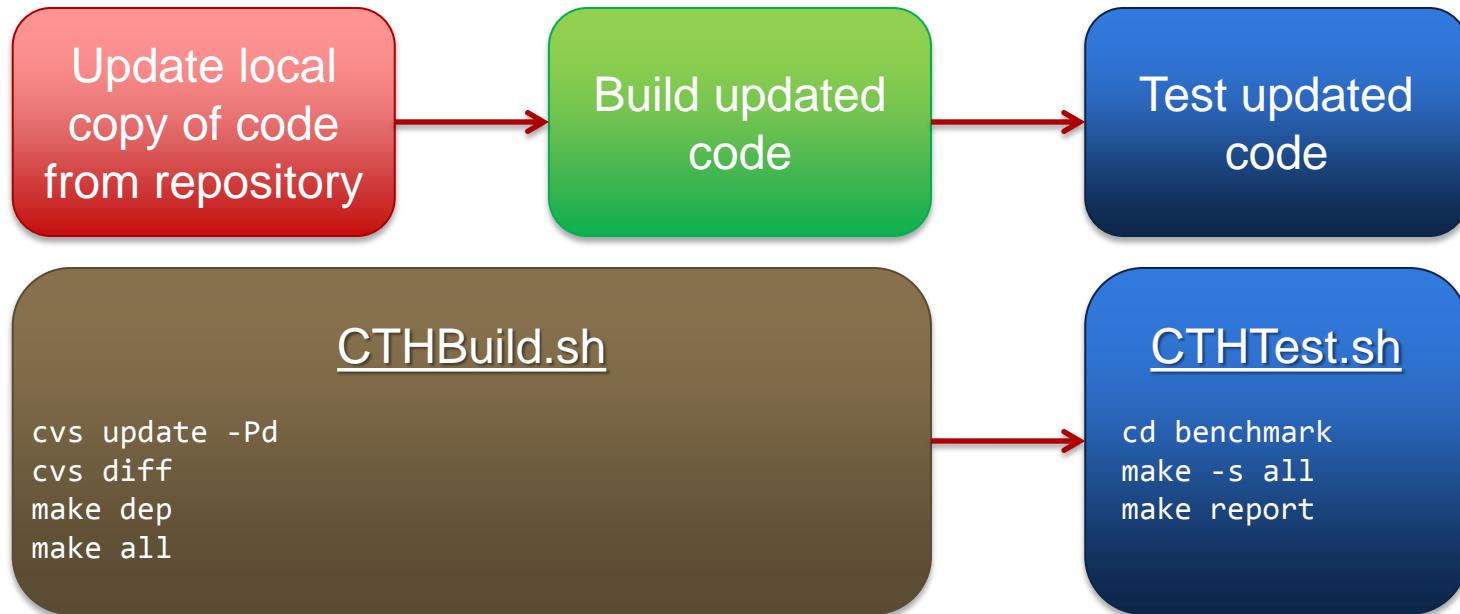
- ASC Sierra and ASC CTH were chosen because they encapsulate the typical languages, codebase footprint, and SCM usage of most Sandia National Laboratories' scientific software development teams.

ASC Sierra Scenario Simulator



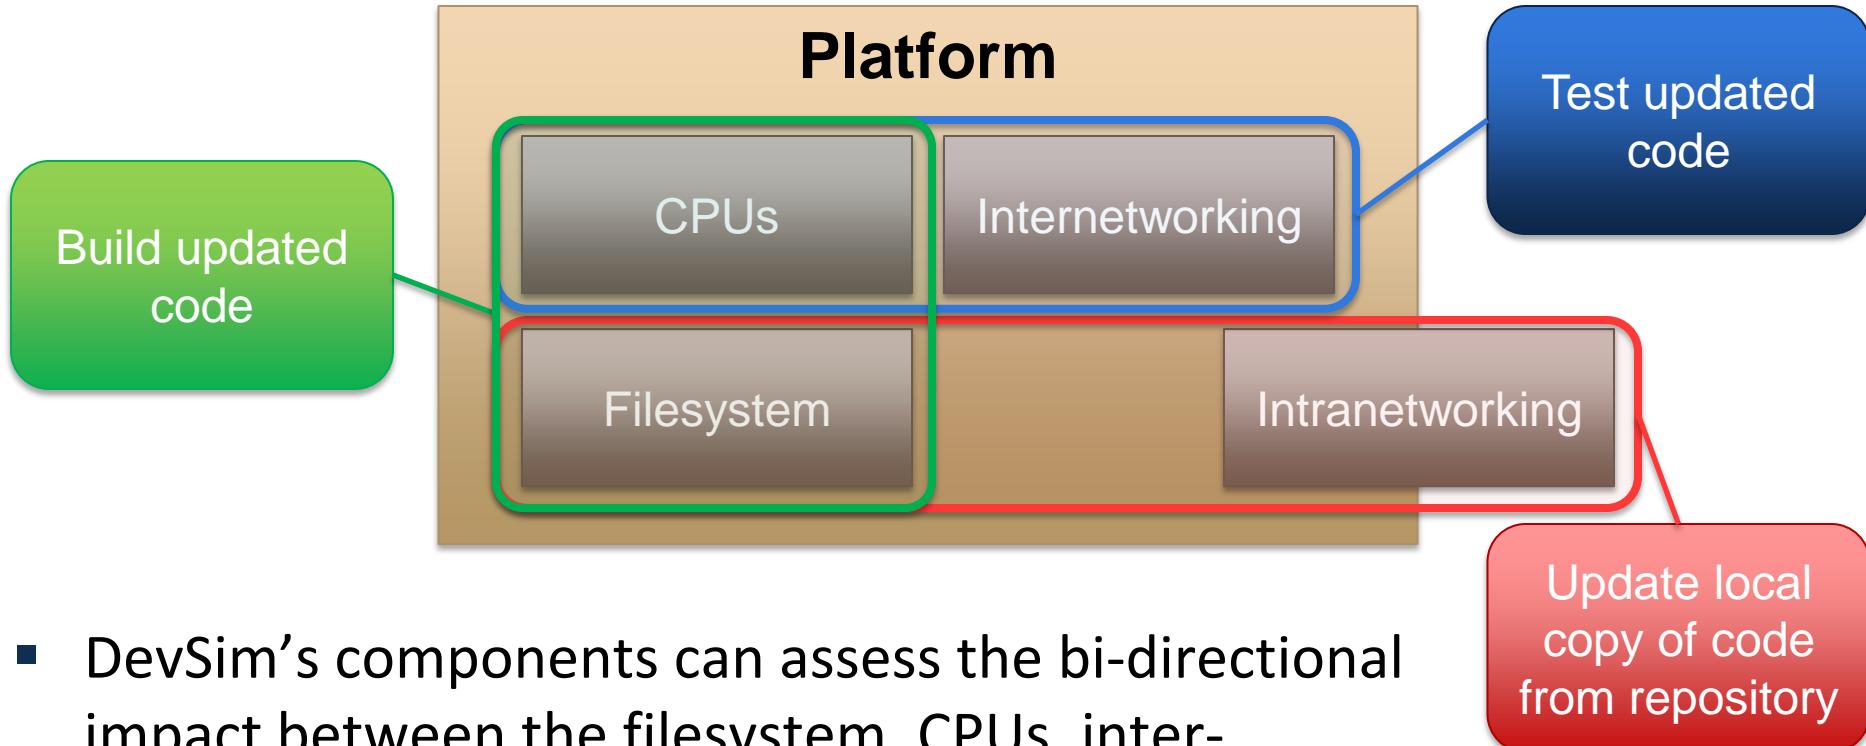
- Sierra has 5 Git repositories, so these are all updated at the same time with `AllGits.sh`. This update process is file system intensive.
- `SierraBuild.sh` currently builds in a threaded fashion (typically no more than 12 threads) and can use a lot of RAM.
- `SierraTest.sh` currently tests SIERRA Multimechanics Module: Aria with its nightly MPI tests.

ASC CTH Scenario Simulator



- `CTHBuild.sh` updates the source and performs the build. This build will serially perform the compile and link operations.
- `CTHTest.sh` will run all tests, which utilize serial and MPI CTH binaries, and generate the final report.

Typical Use Cases



- DevSim's components can assess the bi-directional impact between the filesystem, CPUs, inter-networking, and intranetworking of a platform and a typical developer workflow.
 - Changes to the platform can be tested to quantify impact.
 - DevSim can be used for platform planning and procurement activities.