

SC19 Birds of a Feather

HPC System Testing: Procedures, Acceptance, Regression Testing, and Automation

Verónica G. Melesse Vergara
Oak Ridge Leadership Computing Facility

Bilel Hadri
KAUST

BOF Goals

- Share acceptance test procedures
- Learn about tools used successfully for testing
- Compile list of resources, tools, and tests.
- Start a working group

Live Survey

<https://tinyurl.com/system-test-bof>

Room: SC19HPC

Lightning Talk Presenters

- KAUST: Bilel Hadri
- NCSA: Brett Bode & Celso Mendes
- NERSC: Tina Declerck
- OLCF: Verónica G. Melesse Vergara

Thank you!

<https://olcf.github.io/system-test-wg/events/sc19bof.html>

Comment or send PR:

<https://github.com/olcf/system-test-wg>

OLCF Acceptance Testing

Verónica G. Vergara Larrea

Reuben Budiardja

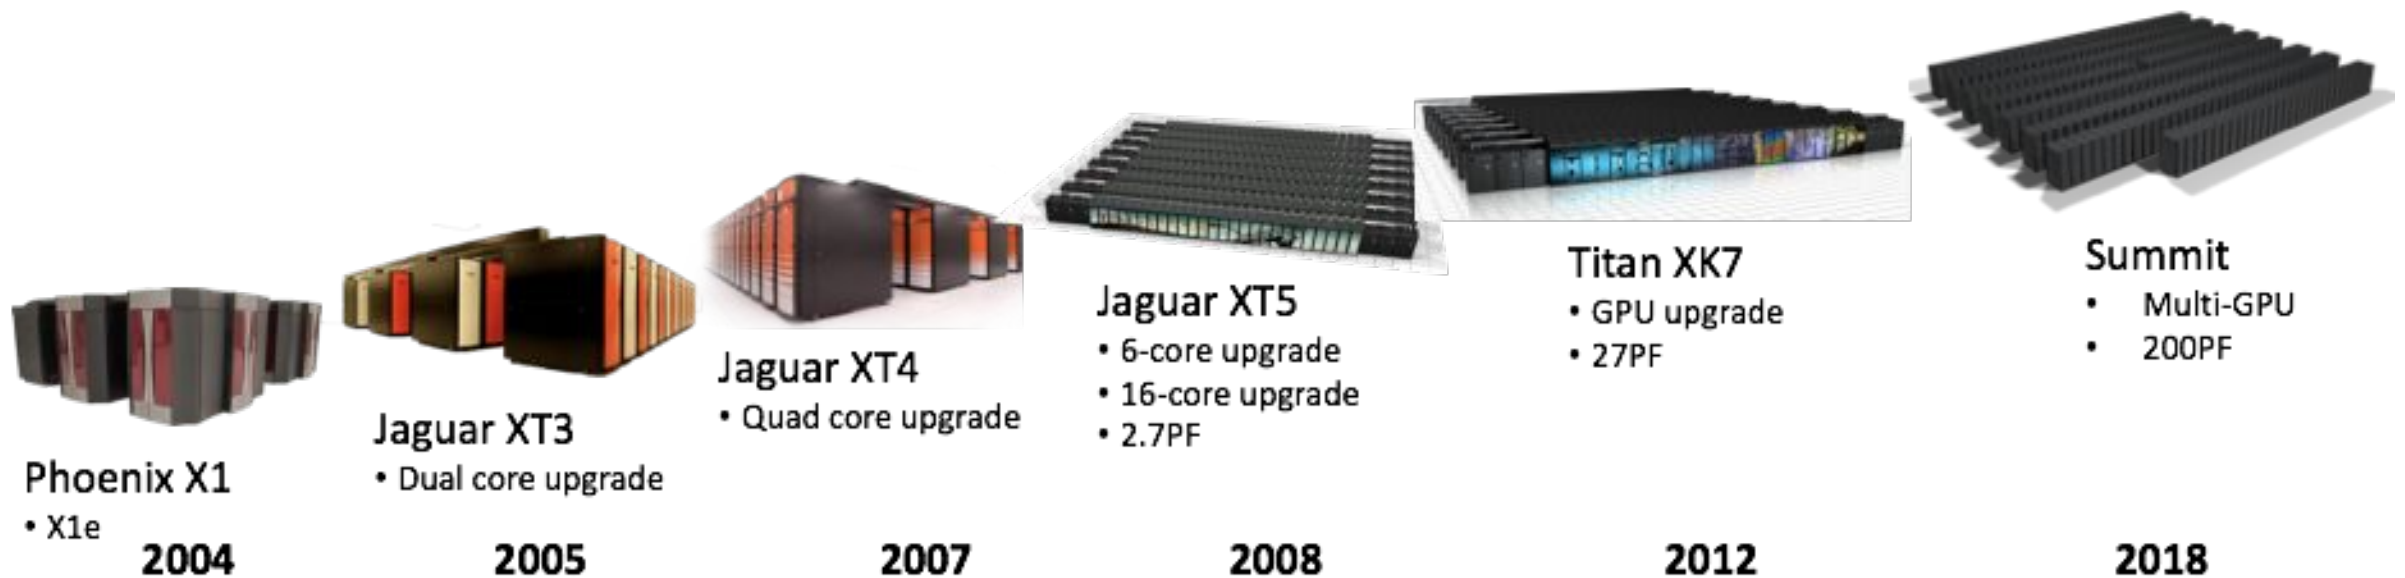


Outline

- What is the OLCF?
- Summit
- Acceptance Testing
- Post-acceptance Testing

What is the Oak Ridge Leadership Computing Facility?

- Deploy and operate computational and data resources required to tackle global challenges.
- Offer leadership-class computing resources to researchers who have many of the largest computing problems in science.
- Partnering has been essential to delivering science in a rapidly changing computational environment.



Summit

Compute System

- 256 compute racks
- 4,608 compute nodes
- Mellanox EDR IB fabric
- 200 PFLOPS – FP64
- ~11 MW, 70F cooling water
- 10.2 PB Total Memory



Compute Rack

- 18 Compute Servers
- Warm water (70°F direct-cooled components)
- RDHX for air-cooled components




39.7 TB Memory/rack
55 KW max power/rack

Compute Node

- 2 x POWER9
- 6 x NVIDIA GV100
- NVMe-compatible PCIe 1.6 TB SSD

IBM POWER9

- 22 Cores
- 4 Threads/core
- NVLink



NVIDIA GV100

- 7 TF
- 16 GB @ 0.9 TB/s
- NVLink



25 GB/s EDR IB- (2 ports)
512 GB DRAM- (DDR4)
96 GB HBM- (3D Stacked)
Coherent Shared Memory

GPFS File System

250 PB storage
2.5 TB/s read, 2.5 TB/s write
(*2.5 TB/s sequential and 2.2 TB/s random I/O)



Acceptance Testing

HW

- Hardware Acceptance Test: Complete hardware diagnostics.

FT

- Functionality Test: Demonstrate that basic hardware and software functionality meet essential requirements.

PT

- Performance Test: Demonstrate that the system hardware and software meet performance and scalability requirements of the suite of applications defined in the Agreement.

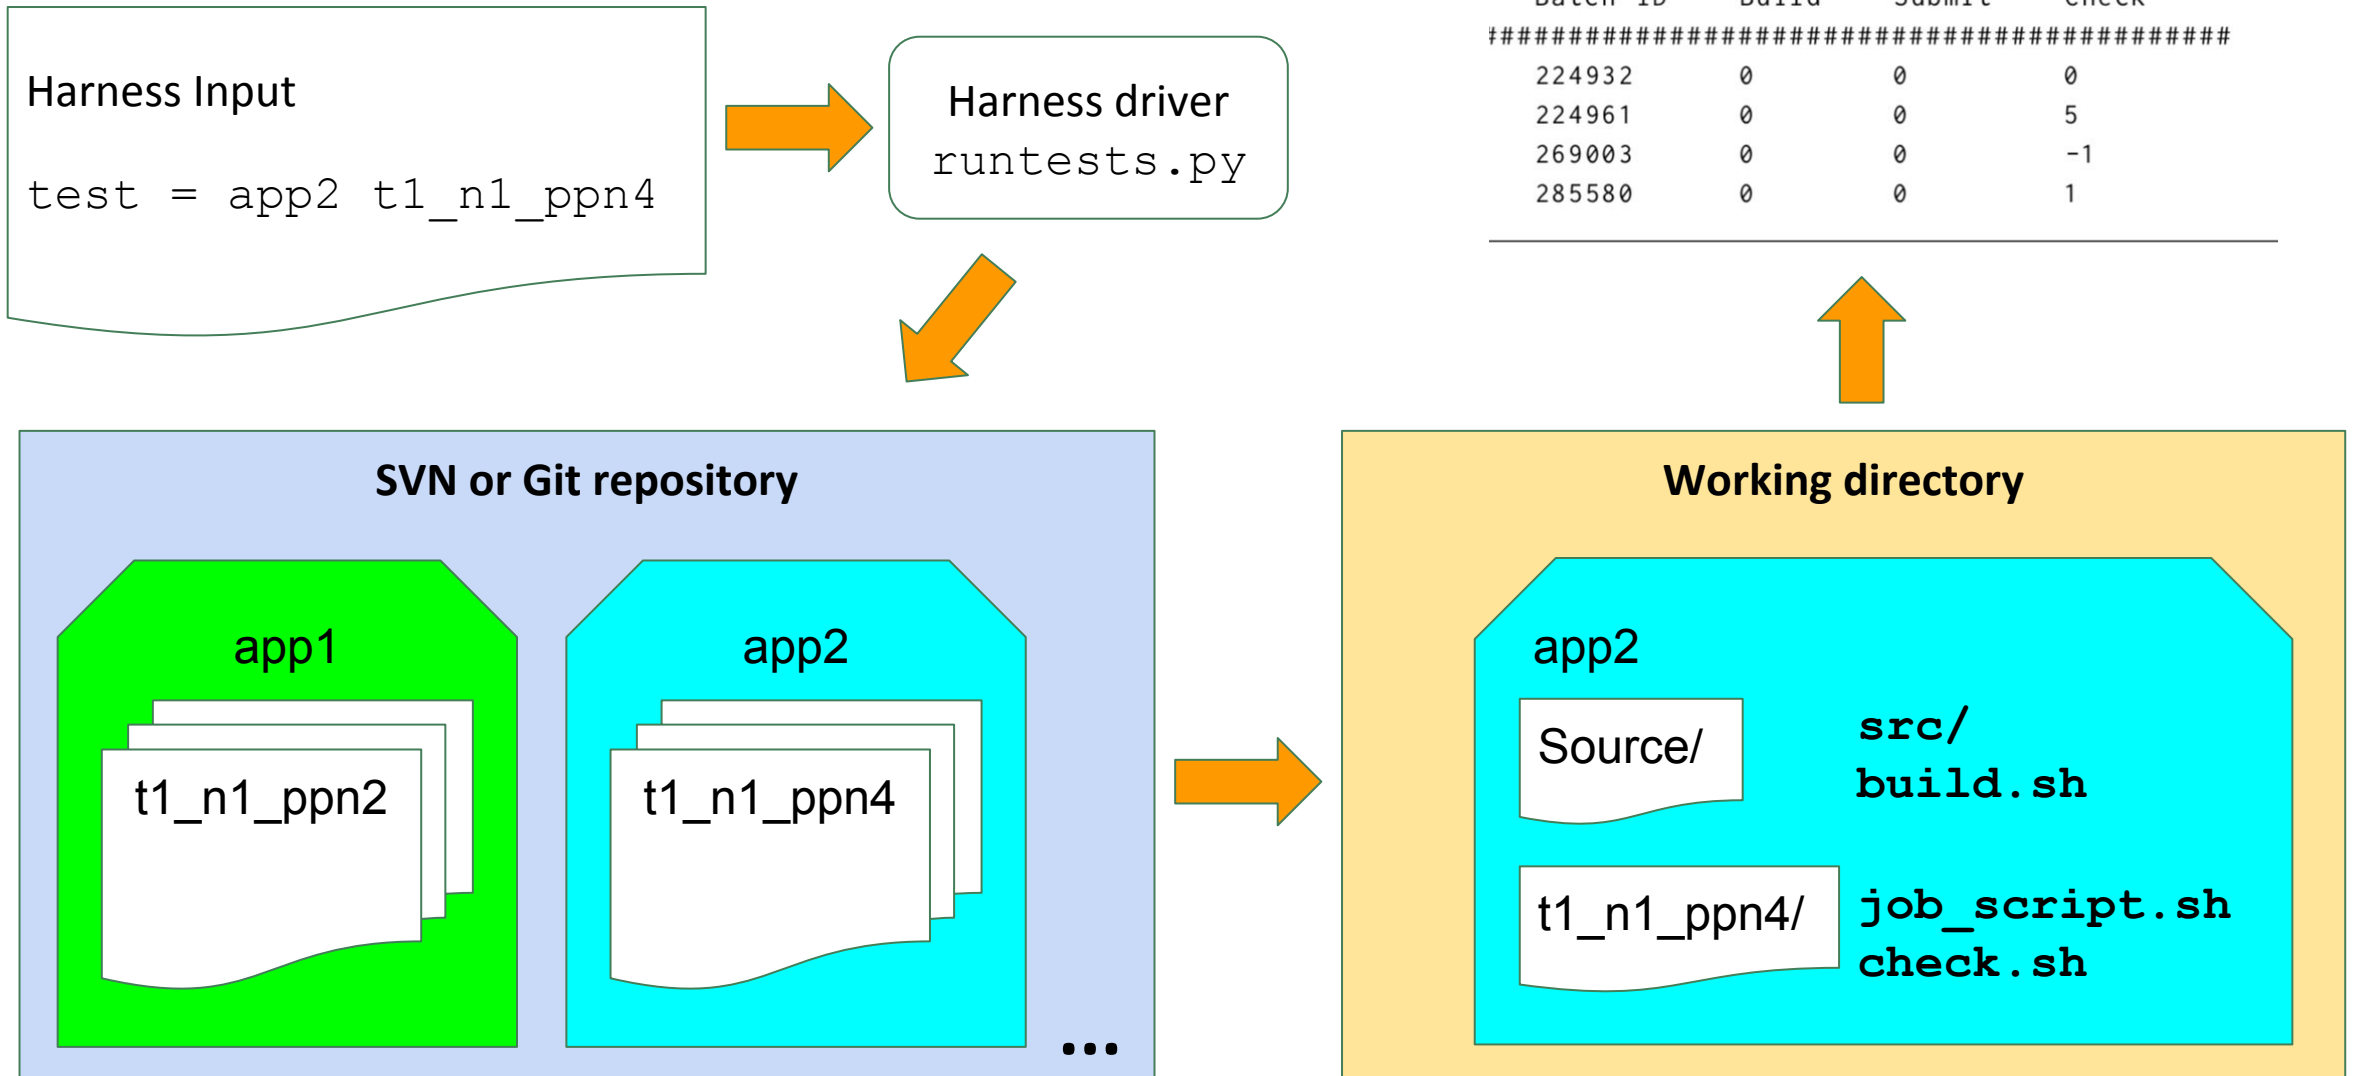
ST

- Stability Test: Demonstrate stability across a mix of simulated code development activity and production simulations.

Acceptance Tests (cont'd)

- Benchmarks
 - contractual, feature specific
- Real-world Applications
 - past workloads and expected workloads
- Workflows
- Tools!
 - Profilers (nvprof)
 - Test in single- and multi-host modes on applications
 - Test traces, profiles, analysis metrics for CUDA programs
 - Debuggers (ARM DDT)
 - Ensure it can run at-scale (20% full system) reliable and start within 5 minutes
 - Used in *offline* (non-interactive) mode
 - Breakpoints set on application, output captured and validated with script

OLCF Test Harness



Post-acceptance Testing

- A shortened version of the acceptance test is prepared to validate new versions of the HPC SW stack.
- Expand tests to include applications from the new allocation cycle.
- Multi-stage testing:
 - Starts at the smaller scale on the test & development system
 - Schedule a 8-12 hour testshot on Summit
- Delicate balance: downtime vs. risk of introducing issues into production
- Automate regression testing:

