

# XSEDE

Extreme Science and Engineering  
Discovery Environment

# Programmable Cyberinfrastructure e

## Introduction to building Clusters in the Cloud

PEARC 18

7/22/2018

Eric Coulter

# Speakers

- Eric Coulter
  - Indiana University, XCRI Engineer
- Jeremy Fischer
  - Indiana University, Senior Technical Adviser, Jetstream
- Suresh Marru
  - Indiana University, Deputy Director, Science Gateways Research Center

# What is XCRI?

## (XSEDE Cyberinfrastructure Resource Integration)

- Very similar to Campus Bridging!
- XCRI provides software toolkits to ease use of local resources, and facilitate easy transitions between local and XSEDE resources
- We also do site visits and remote consultation!
- Continually looking for feedback from XSEDE users, Campus Champions, and service providers to keep our offerings up-to-date with current needs

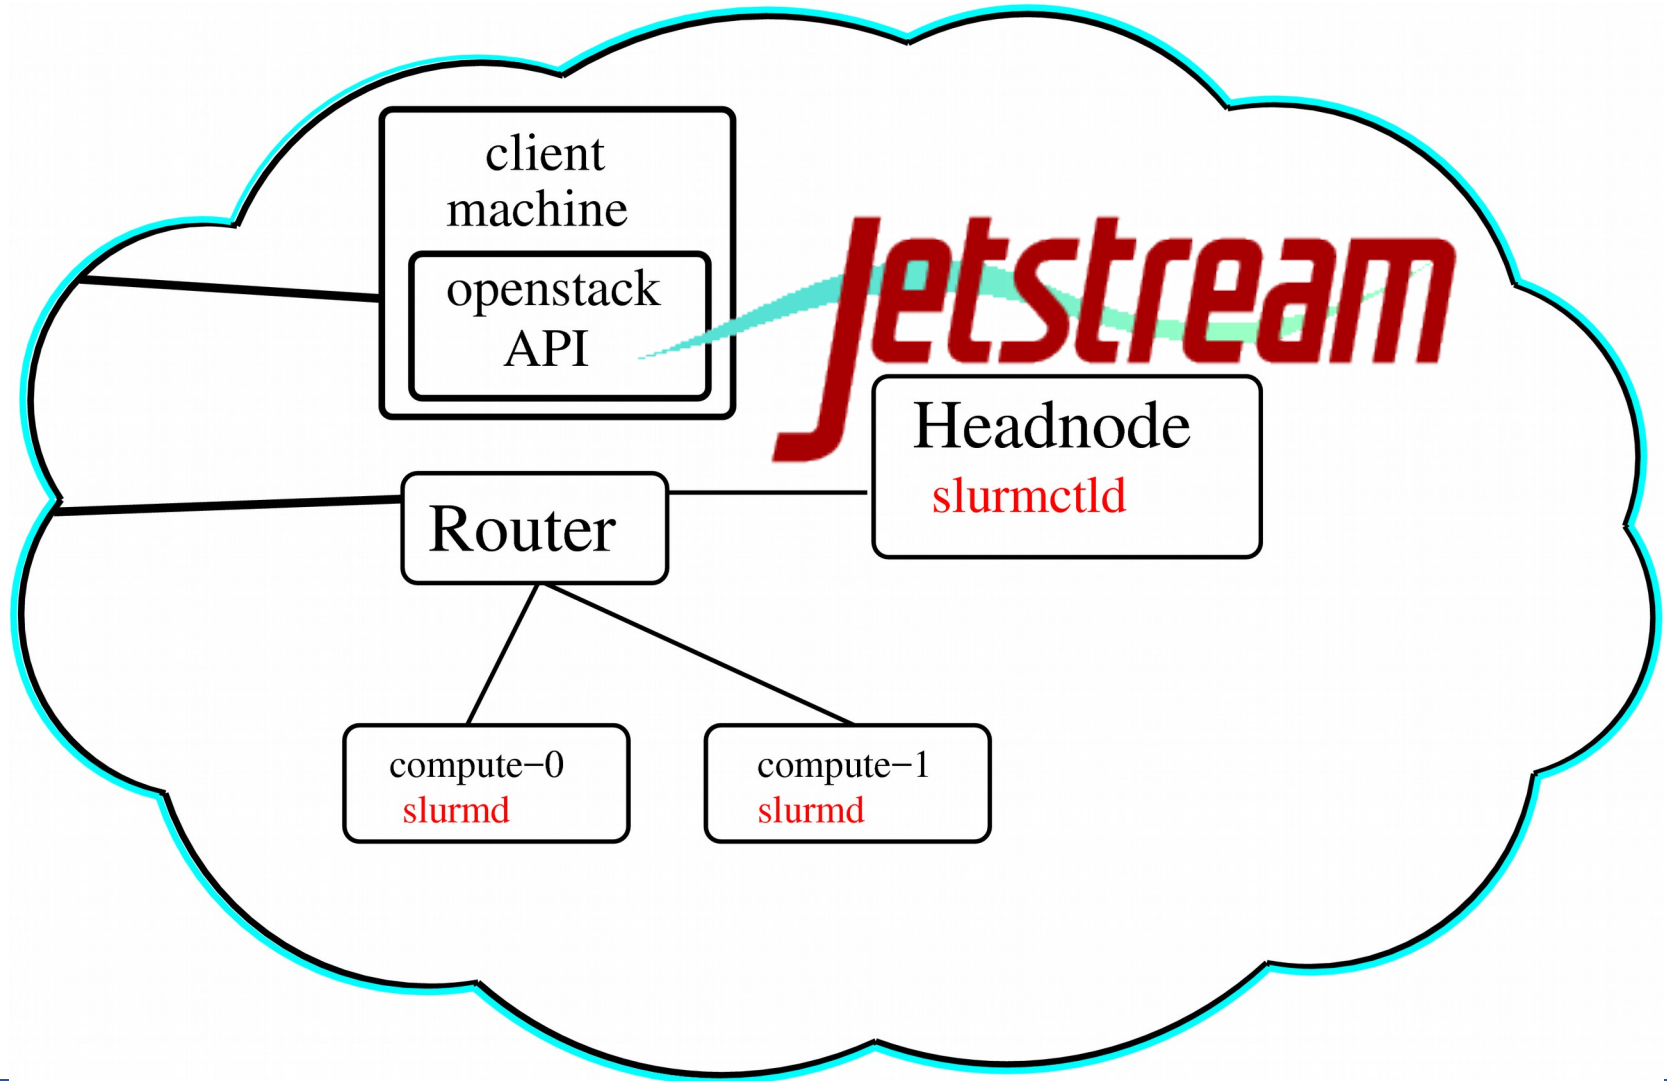
# XCRI Toolkits

- XCBC
  - Build a cluster based on OpenHPC
- XNIT
  - Get open-source scientific software
- Globus Connect Server configuration management
  - Easily set up a local globus connect server using Ansible
- Jetstream Virtual Clusters
  - Build a cluster in Jetstream, with Openstack and Ansible
- Cluster Monitoring toolkit
  - Easily set up Ganglia and XDMod for cluster load and usage statistics

# Jetstream Virtual Clusters

- Inspired by a need for more resources on Science Gateways
- Science Gateways allow users to submit jobs through a web interface, to a variety of resources – local, XSEDE, or cloud.
- The Airavata middleware developed by the SGRC (Science Gateways Research Center at Indiana University) makes these easy to use
- This model of virtual cluster was specifically developed for the SEAGrid project, available and easily configurable for anyone else.

# What we're going to build:



# Pieces of the whole:

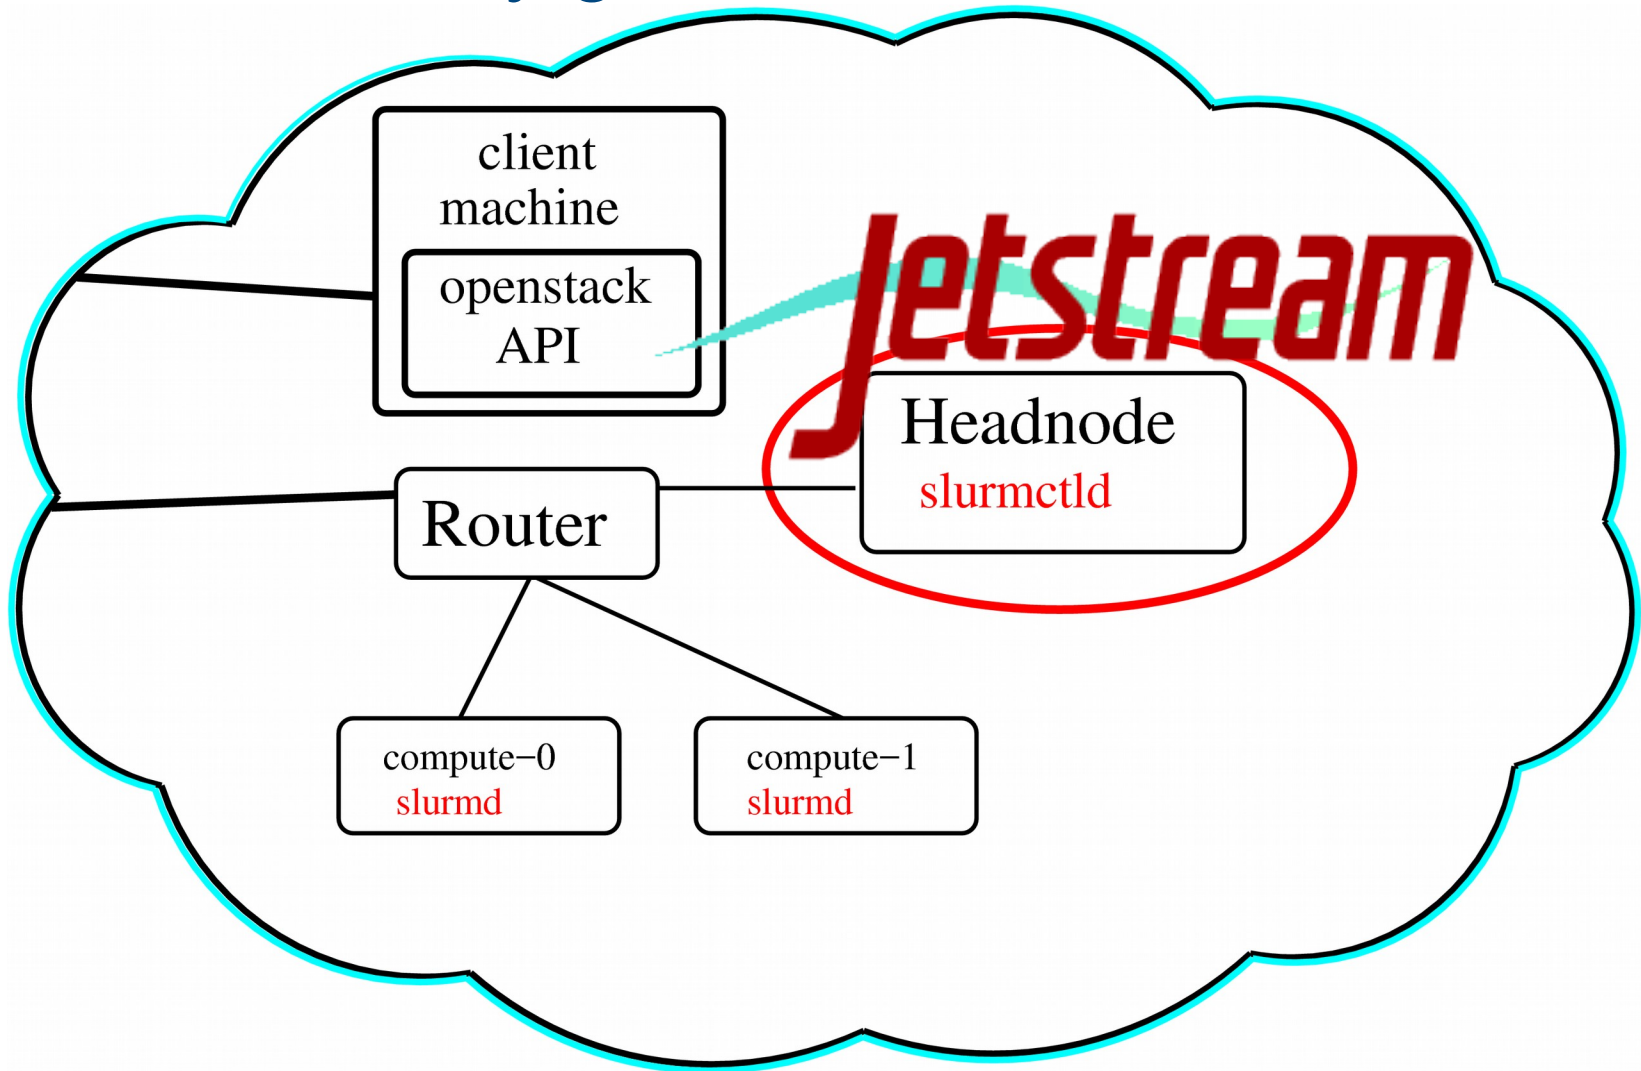
- OpenHPC Project
  - Slurm – for managing compute resources and scheduling
  - Spack – for building software
  - Lmod – for module environment
- Ansible – for compute node configuration
- Openstack client - for elasticity

# Let's get started!

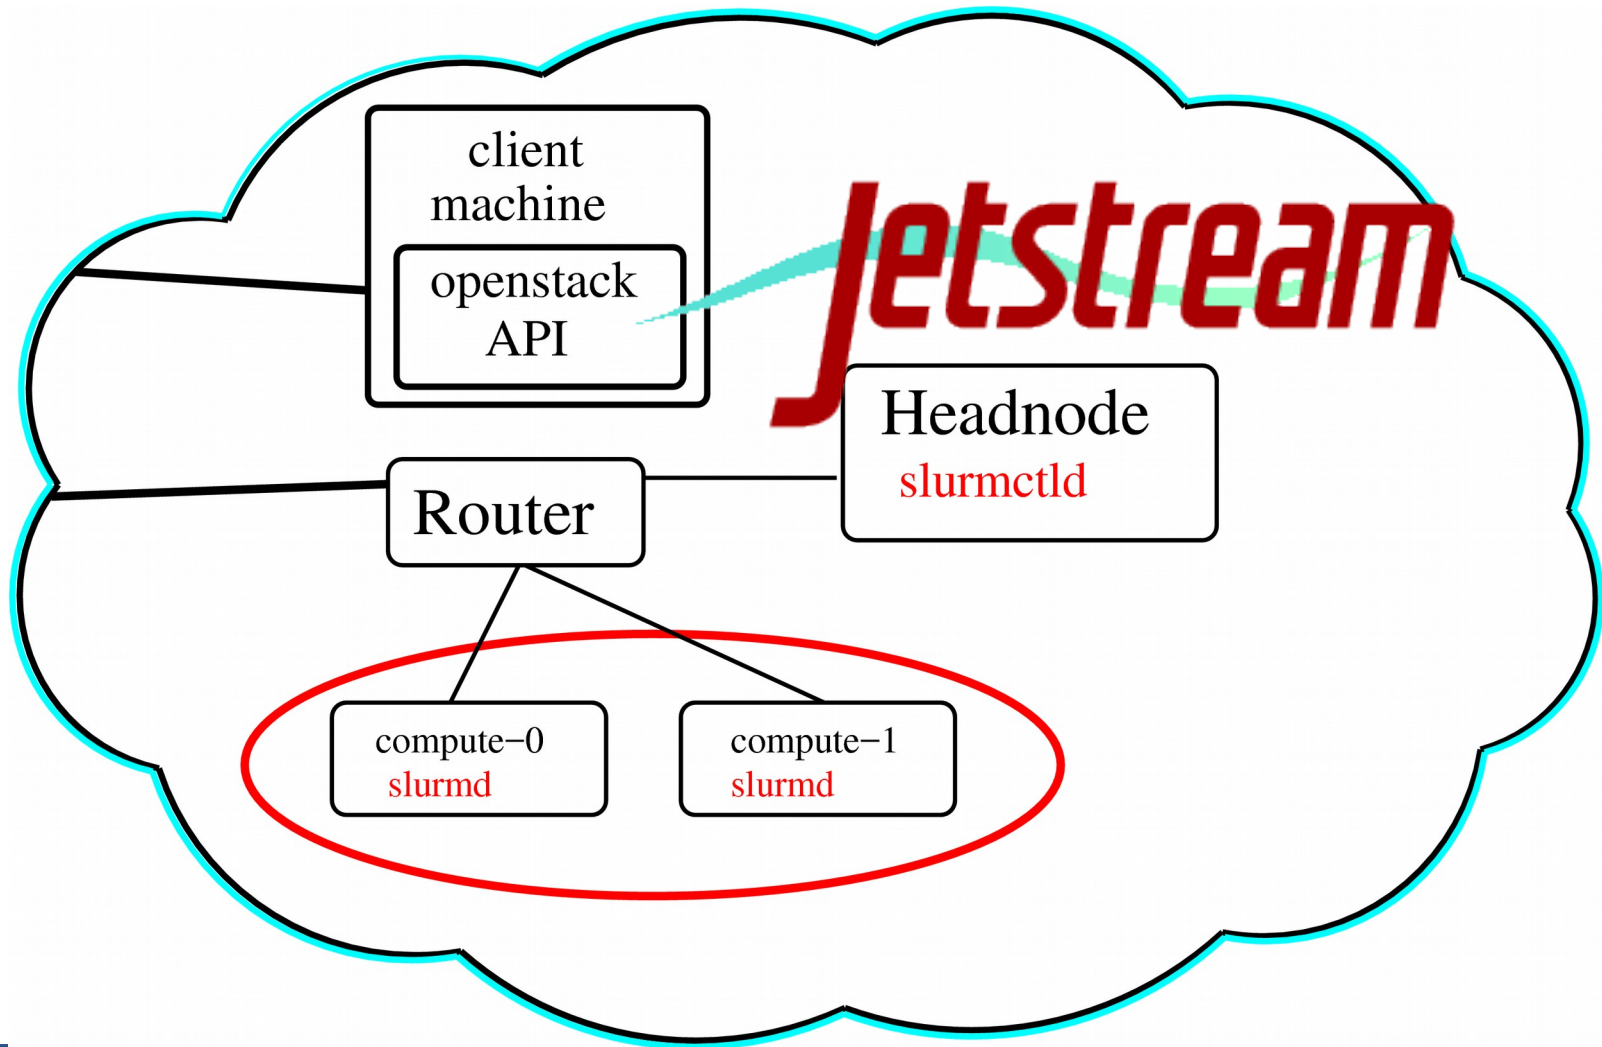
<https://goo.gl/FmoHZ5>



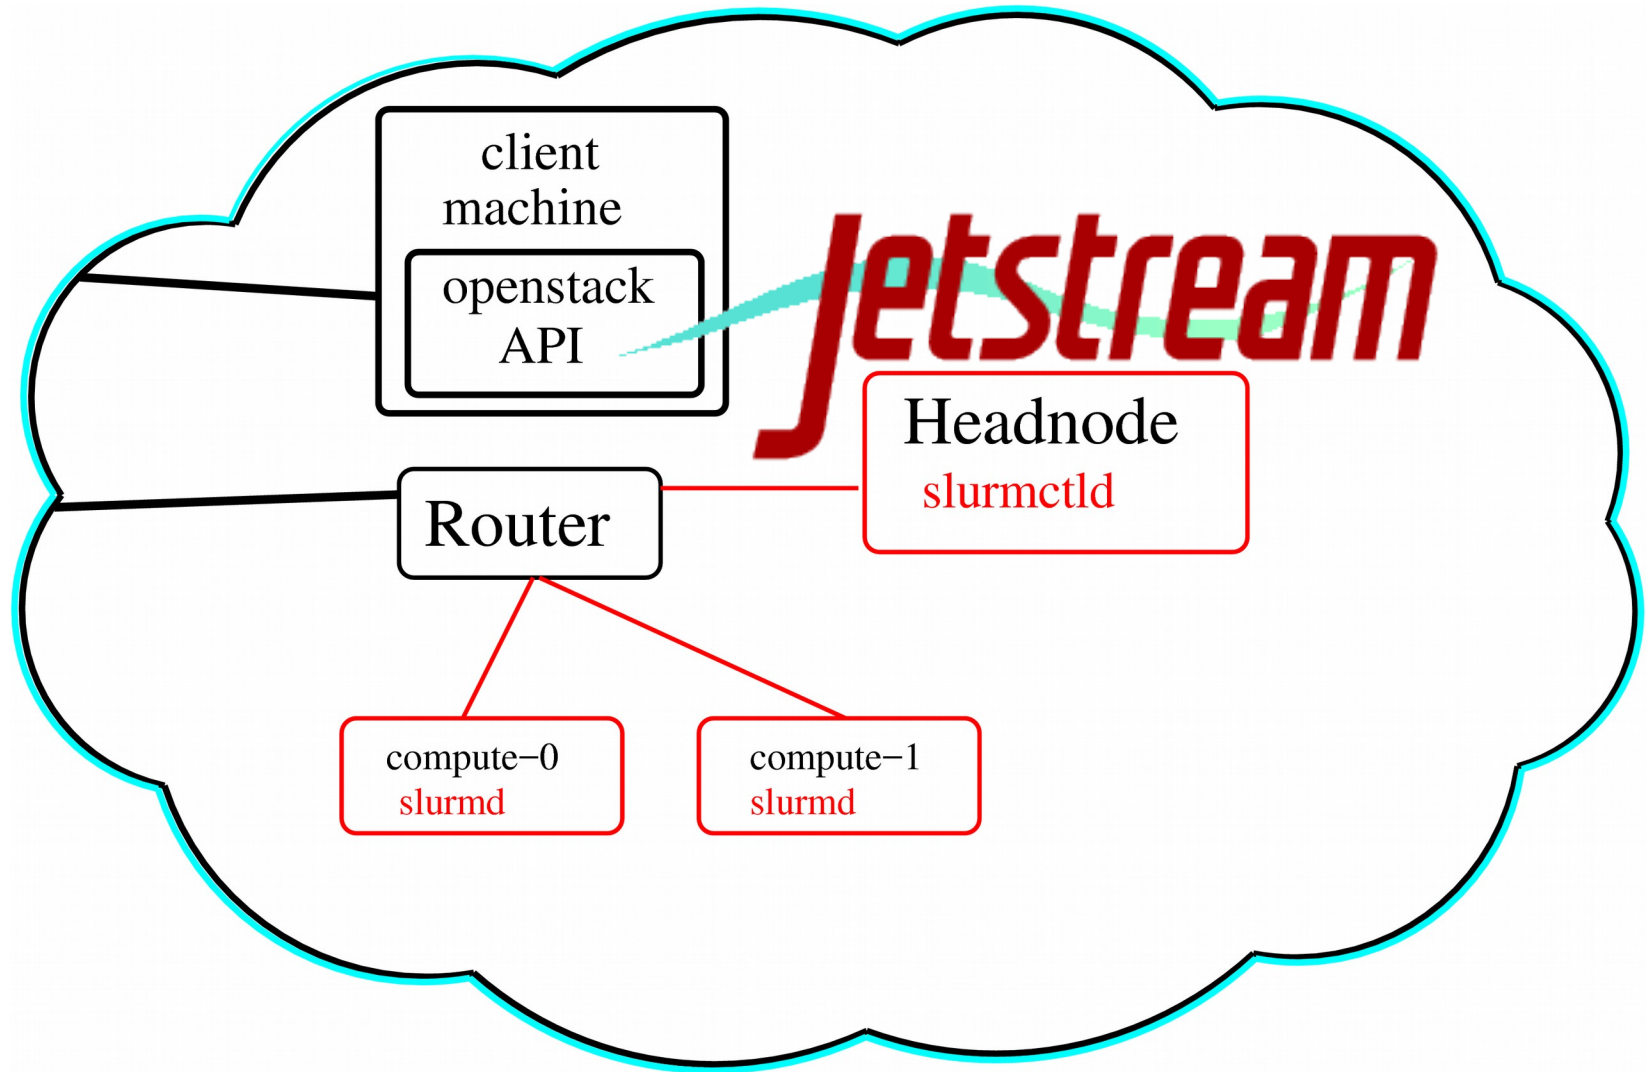
You've already got a headnode:



Now, create the compute nodes:



Now, configure the scheduler (etc.)!



# Where can I get help?

- **Wiki / Documentation:** <http://wiki.jetstream-cloud.org>
- User guides: <https://portal.xsede.org/user-guides>
- XSEDE KB: <https://portal.xsede.org/knowledge-base>
- Email: [help@xsede.org](mailto:help@xsede.org)
- Campus Champions: <https://www.xsede.org/campus-champions>
- Training Videos / Virtual Workshops (TBD)

# Spack configuration for OpenHPC:

- Spack installs in
  - `/opt/ohpc/admin/spack/0.11.2/`
  - Configuration files in:
    - `/opt/ohpc/admin/spack/0.11.2/etc/defaults/`
    - `/root/.spack/linux/`
  - Should be run as privileged user
  - Need to make sure software will install in a public directory:
    - Edit `install_tree` and `tcl modules path` in `/opt/ohpc/admin/spack/0.11.2/etc/defaults/config.yaml`
    - Provide OpenHPC compiler Module names in `/root/.spack/linux/compilers.yaml`