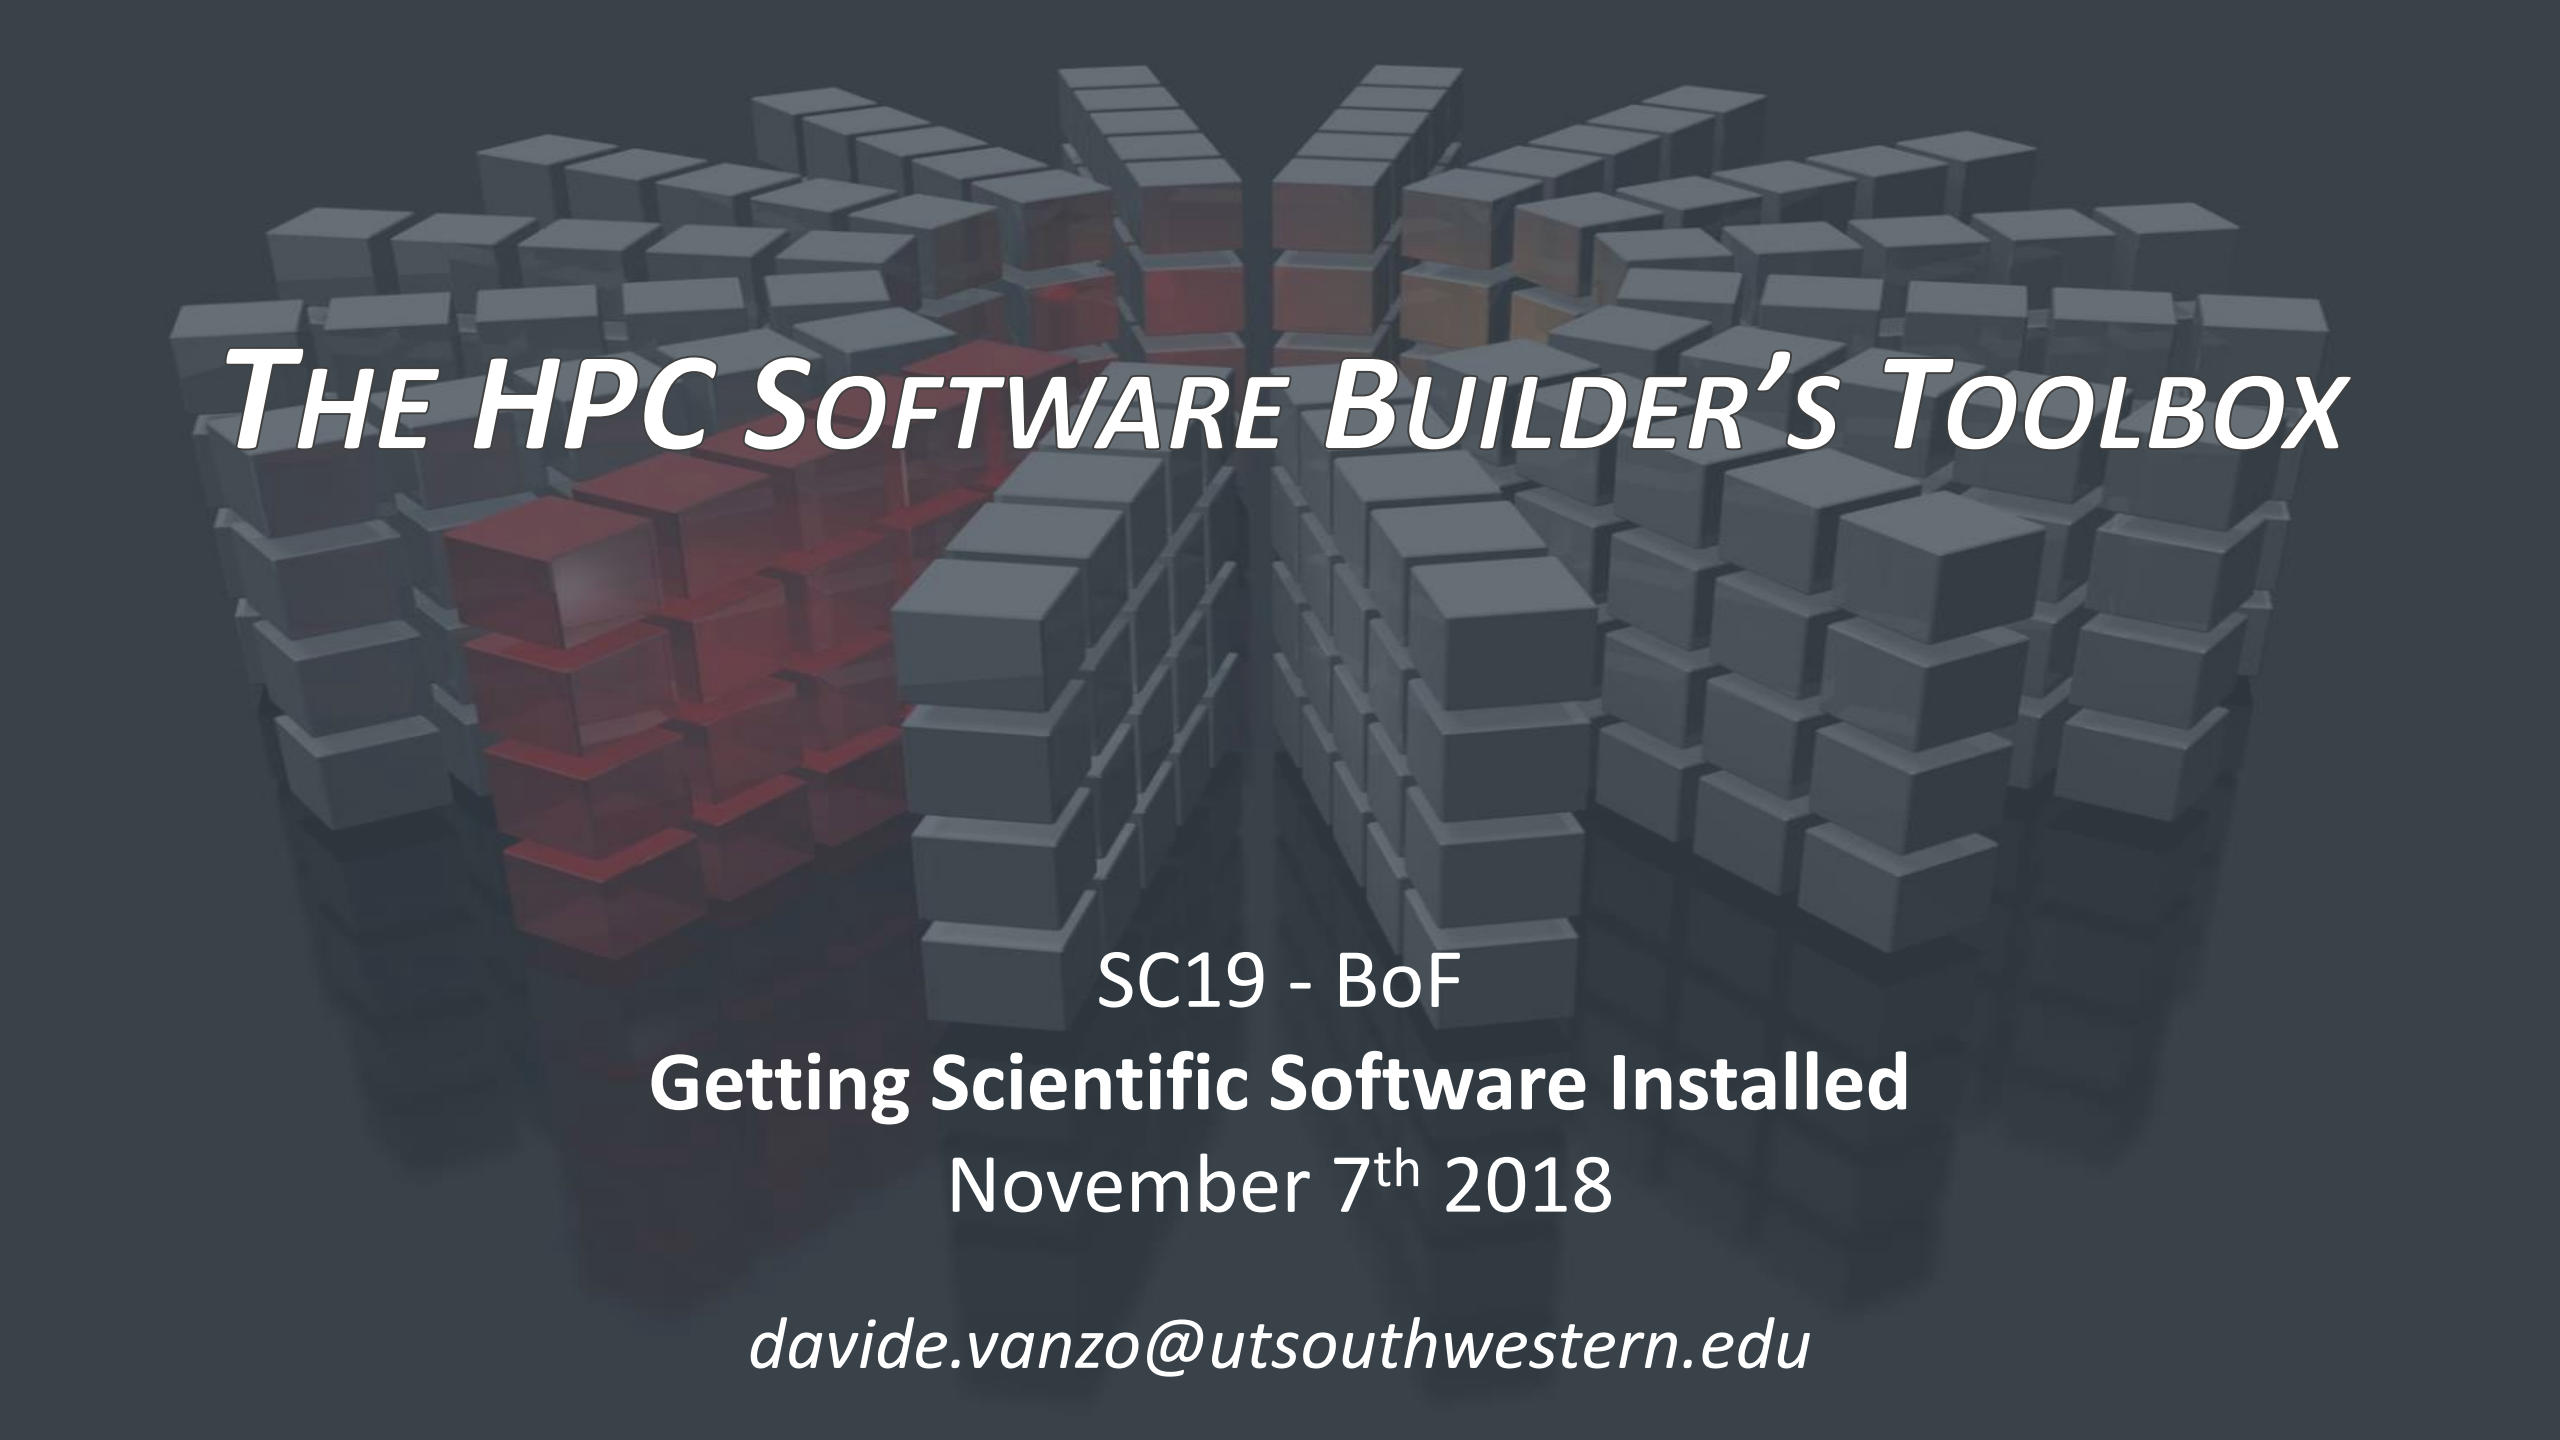

THE HPC SOFTWARE BUILDER'S TOOLBOX

SC19 - BoF

Getting Scientific Software Installed

November 7th 2018

davide.vanzo@utsouthwestern.edu

GETTING SCIENTIFIC SOFTWARE INSTALLED

Managing a scientific software stack is a significant problem for all HPC sites around the world.

- Built from source for maximum performance
- Provide reliability and reproducibility
- Tedious, time-consuming, frustrating process
- Significant burden for users support teams
- Very little collaboration across HPC centers (until recently)

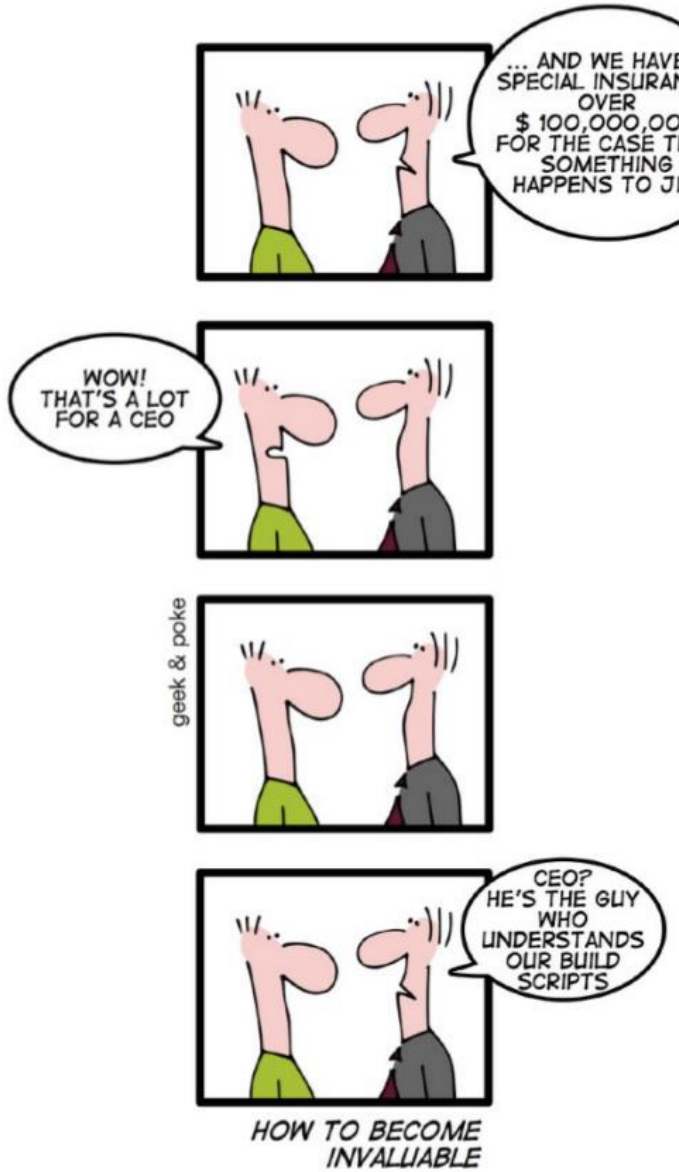

http://geekandpoke.typepad.com/geekandpoke/2010/05/how-to-become-invaluable.html

EASYBUILD

Framework to build and install scientific software on HPC clusters.

http://easybuilders.github.io/easybuild

- Implemented in **Python 3 (and 2)**. FOSS (**GPLv2**)
- Supports **Linux** and **Cray** systems
- Architectures: **x86_64**, **POWER**, **ARM**
- Designed for **HPC user support teams** and **end users**
- Builds from **source** with attention to **performance**
- Build based on **“recipe” files**
- **2,000+** supported software
- Highly community driven

EASYBUILD

100+ SITES

UNIVERSITY OF BIRMINGHAM

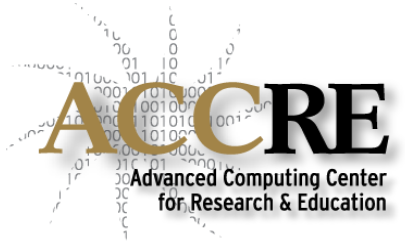

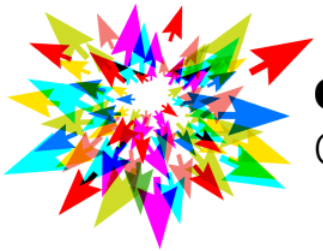

compute canada

CSCS

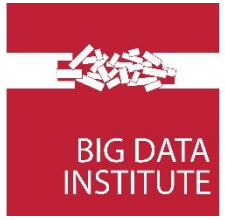

NUS
National University
of Singapore

NEWS

- Major version 4 released Sep 2019. Supports both Python 2 and 3.
- No more required python packags. Only Python standard libraries.
- Software installation directories names independent of module naming scheme.
- New software-specific easyblock for OpenMPI
- Use SYSTEM toolchain rather than deprecated dummy toolchain

5TH EASYBUILD USER MEETING

The EasyBuild User Meeting is an open and highly interactive event that provides a great opportunity to meet fellow EasyBuild enthusiasts, discuss related topics and learn about new aspects of the tool.

Jan 29th - 31st 2020 *(week before FOSDEM'20)*
Parc Tecnològic - Barcelona, Spain

https://github.com/easybuilders/easybuild/wiki/5th-EasyBuild-User-Meeting

Organized in collaboration with:

HPC Now!

GHENT UNIVERSITY

SC19 BoF – Getting Scientific Software Installed

Spack

Flexible package manager for HPC software

https://spack.io

- Implemented in **Python 2** and **3**. OSS (**Apache-2.0**, **MIT**)
- Supports **Linux**, **MacOS** and **Cray**
- Architectures: **x86_64**, **POWER**, **ARM**
- Designed for **user support teams**, **developers**, **end users**
- Builds from **source** with attention to **performance**
- Advanced **software dependency graph** resolution
  - Extensive flexibility in dependency choice
- **RPATH** linking
- 2,900+ supported software packages
- Highly community driven

SPACK @ SC19

https://spack.io/spack-at-sc19

Thurs., November 21

- **12:15pm - 1:15pm**, in 503-504
  The second **Spack Community BOF** at SC, will feature a brief presentation by core developers on the latest Spack release and roadmap directions, followed by an interactive survey.

- **2:30pm - 3:30pm**, at DOE Booth 925
  Ask the Spack developers anything at the DOE Booth. Once again, core developers will be available to discuss roadmap directions, issues, collaborations, or anything else Spack-related.

Fri., November 22

- **9:45am - 10:00am**, in 405-406-407
  Carson Woods, Matt Curry, and Anthony Skjellum will be **presenting Implementing a Common HPC Environment in a Multi-User Spack Instance** at the HPC System Professionals Workshop. (See also PR #11871).

Spack

Sites


News

- Support for virtual environments (for HPC codes, Python, R)
- Support for spack.yaml / spack.lock files embedded in git repositories
  - Clone a project, cd into it, run `spack install` to install all dependencies
  - Use spack.lock file to reproduce the entire build
- Build infrastructure for binary packages now deployed in AWS -- public binaries coming soon
- Finer-grained control over compiler flags
- Relicensed entire project from LGPL to Apache-2.0/MIT

Conda, Guix, Nix

https://conda.io

Package, dependency and environment management for "any language".

OS: **Linux, MacOS, Windows**

Platforms: **x86_64**

Implementation: **Python 2/3, YAML**

Target: **End users**

- Binary packages installation
- Anaconda cloud + channels
- YAML package recipes
- 3,500+ supported software

https://nixos.org/nix

The purely functional package manager

OS: **Linux, MacOS, Unix**

Platforms: **x86_64**, **AArch64**

Implementation: **C++**

Target: **System administrators**

- Binary packages installation or builds from source if not available
- Strong focus on reproducibility
- Nix DSL package recipes
- 13,000+ supported software (scientific software only minority)

https://gnu.org/software/guix

The GNU package manager

OS: **Linux**

Platforms: **x86_64**, **AArch64**

Implementation: **Scheme, C++**

Target: **System administrators**

- Binary packages installation or builds from source if not available
- Strong focus on reproducibility
- GNU Guile package recipes
- 6,500+ supported software (scientific software only minority)

SINGULARITY

Container solution for HPC environments.

https://singularity.lbl.gov

- **BSD** licensed
- Only **Linux** currently supported.
- Containerization ensures **full application portability**
- Good support for **MPI**, **IB**, **accelerators**, etc.
- Containers executed in **user space**
  No privilege escalation allowed
- **Version 3.5**
  - Encryption applied at rest, transit and during execution.
  - Support for AMD GPUs

HPC Container Maker

Generates container specification files based on a “recipe”.

- Open source project (**Apache-2.0**) by Nvidia in **Python 2** and **3**
- Supports **Singularity** and **Docker** (full container abstraction)
- Higher level of abstraction through modular building blocks
- Generated containers are fully portable
- Performance optimization may impact portability
- Support for Singularity multi-stage builds

https://github.com/NVIDIA/hpc-container-maker

OpenHPC

Collection of open source software for deploying and managing HPC clusters.

- Supports **Linux** on **x86_64** and **ARM** platforms
- Pre-built packages for easy deployment

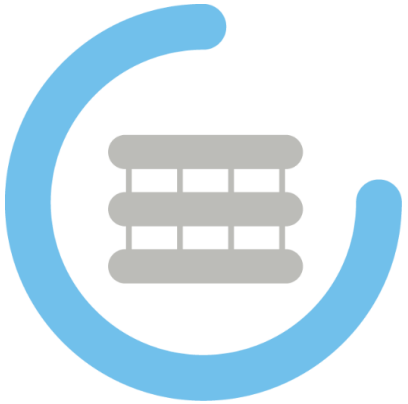

https://openhpc.community

Software mgmt.	EasyBuild, Spack, Lmod
Compilers	GCC, Intel (BYOL)
Containerization	Singularity, Charliecloud
Resource mgmt.	Slurm, PBS Professional
Provisioning	Warewulf, xCAT
Administration	ClusterShell
MPI	OpenMPI, MPICH, MVAPICH2
Numerical libs	OpenBLAS, ScaLAPACK, FFTW, GSL, Metis, Trilinos, PETSc, Mumps and more
⋮	⋮