

Singularity: Mobility of Compute with Containers

When: Thursday, May 19 from noon - 1pm

Where: 200C Warren Hall, 2195 Hearst St (see building access instructions [on parent page](#)).

Event format: The reading group is a brown bag lunch (bring your own) with a short <20 min talk followed by ~40 min group discussion.

Presenter: Gregory M. Kurtzer, Berkeley Research Computing Program / LBNL HPC Services

Facilitator: Patrick Schmitz, Research IT

Singularity is software that provides a way for researchers and others to package together the artifacts of their computational workflows - one or more Linux applications, with all their dependencies - and run them successfully in a variety of Linux environments. This makes it much easier to develop a research workflow on a laptop or a laboratory's server, then bundle it up and copy it: to a departmental cluster; to Savio, the campus's high-performance computing (HPC) cluster; to a national HPC resource; or to a cloud computing environment. When it is necessary to run existing computations faster, or to work with larger datasets, Singularity can help reduce the time and effort required to get existing computational environments up and running on a new and more powerful system.

Greg Kurtzer, Linux Cluster Architect for the BRC Program's Savio cluster and LBNL's HPC Services group, is the lead developer on the Singularity project, and will describe how it works and the kinds of use cases addressed by the software. We'll also discuss experiments currently underway to explore the feasibility and utility of deploying Singularity in the Berkeley Research Computing Program's Savio cluster, in support of campus computational research. Please bring your own use cases to the discussion to evaluate whether Singularity is a technology that will support your or your group's research.

===

Before the meeting, please review [Greg's slide deck \(PDF\)](#). Greg will give a brief presentation from the deck, but in order to leave ample time for discussion he will lean toward responding to questions rather than dive deeply into technical details covered in some of the slides.

Presenting: Greg Kurtzer

Attending:

Aaron Culich, Research IT
Aron Roberts, Research IT
Barbara Gilson, SAIT
Bill Allison, CTO / IST-API
Chris Hoffman, Research IT
Chris Paciorek, BRC / Statistics
Dav Clark, D-Lab
Greg Kurtzer, BRC / LBNL
Jason Huff, BRC / CGRL
John White, BRC / LBNL (via Zoom)
Kelly Rowland, BRC / Nuclear Engineering
Krishna Muriki, BRC / LBNL
Michael Jennings, BRC / LBNL
Olivia Habern, Research IT
Patrick Schmitz, Research IT
Paul, Student Intern, LBNL
Quinn Dombrowski, Research IT
Sophia Xia, Student Intern, BRC/LBNL
Steven Carrier, School of Education
Yong Qin, BRC / LBNL
Zashary, Student Intern, LBNL

Context: People on campus who need their own software environments in which to run their own code, but where their familiar/tuned environment differs from the HPC environment Savio has installed. BCE. "Mobility of Compute" -- that's the basic goal/concept: be able to make an application / workflow work in multiple, heterogeneous environment.

Greg's presentation; [see slides](#).

SDSC Comet: uses KVM.

Adaptive Computing (w/ MOAB/TORQUE) supports Docker integrated into scheduling mgmt
Docker Username space

...

Landscape moving quickly overall!

Map only namespaces that one needs to differentiate. Same user inside & outside the container. This allows workflow to run "through the container."

Greg: Make a Singularity image on a machine (or VM) where you have root. Then upload to Savio (or other HPC resource), and run *as yourself/user* -- not root. [Dav: Docker a good way to create that Singularity image, sounds like; Greg: Yes]

Patrick: similar approach to Shifter

Kelly: Shifter, we saw at a prior Reading Group, makes all shifter images available to all users. Singularity does not?

Greg: Correct.

Singularity images default to 1GB max; can be grown (by 0.5GB at a time by default). Image is a sparse file.

Patrick: Question about security -- making it 'easier' to run an old insecure OS or set of libraries into a modern, secure system. Is this a problem?

Greg: Worrysome. But because Singularity user can't escalate beyond her/his privileges in the containing OS, this worry is partly mitigated ... privilege escalation is the common mode of operation for malware / hacks.

Greg: Have requested audits of the Singularity code. It's open, please feel free to review.

PLS: Big, complex applications -- have they run through Singularity? E.g., climate change software, coded long ago?

Greg: Have not run complex applications, but no reason to believe that complexity would be an obstacle.

Yong: Can exploits be launched against vulnerable software (e.g., Apache) running in a container?

Greg: Yes. There are mitigations, but really running a web server (or other daemon) is a good use case for Docker, not Singularity.

Greg: will have version I've been demonstrating released early next month if all goes according to expectation. Hope to have a version on Savio soon after that. Actively seeking use cases. BCE ... (Patrick suggestions: the Linguistics' dept's "Phonetics Machine").