



[www.chameleoncloud.org](http://www.chameleoncloud.org)

## PRACTICAL REPRODUCIBILITY WITH CHAMELEON

**Kate Keahey**

Mathematics and CS Division, Argonne National Laboratory

CASE, University of Chicago

*keahey@anl.gov*


**October 14, 2021**

**6th Annual CROSS Research Symposium**



# PRACTICAL REPRODUCIBILITY

- ▶ Can experiments be as sharable as papers are today?
- ▶ Could it be as easy to provide conditions for reviewers to repeat experiments or data analysis in a paper as it is to organize a PC meeting?
- ▶ Can I simply integrate somebody's model into my research instead of reinventing the wheel?
- ▶ Can I have so much fun playing with somebody's experiment that discover a new result?
- ▶ Can I develop exercises for my class based on most recent research results?



*The existence of powerful open testbeds is a fundamental requirement for practical reproducibility*

# CHAMELEON IN A NUTSHELL

- ▶ We like to change: a testbed that adapts itself to your experimental needs
  - ▶ Deep reconfigurability (bare metal) and isolation – but also a small KVM cloud
  - ▶ power on/off, reboot, custom kernel, serial console access, etc.
- ▶ Balance: large-scale versus diverse hardware
  - ▶ Large-scale: ~large homogenous partition (~15,000 cores), ~6 PB of storage originally distributed over 2 sites (**UC, TACC**) connected with 100G network
  - ▶ Diverse: ARMs, Atoms, FPGAs, GPUs, Corsa switches, etc.
  - ▶ **CHI-in-a-Box** sites at Northwestern, NCAR, IIT, and other places
- ▶ Cloud++: CHameleon Infrastructure (CHI) via mainstream cloud tech
  - ▶ Powered by OpenStack with bare metal reconfiguration (Ironic) + “special sauce”
  - ▶ Blazar contribution recognized as official OpenStack component
- ▶ Reproducibility, repeatability and sharing
  - ▶ Packaging (via Jupyter), sharing, discovering, and publishing experiments



## OPEN TESTBED – BY THE NUMBERS

400+

Papers  
published

45

Countries

750+

Projects

Over

6,000

Users

160+

Institutions

6+

Years Old

and 3 more  
years to grow!

# TESTBED AS SHARING PLATFORM

- ▶ **Instruments held in common** are a reproducibility imperative
  - ▶ Hardware and hardware versions: >105 versions over 5 years
  - ▶ Expressive allocation
- ▶ **Sharing via **cloud pattern****
  - ▶ Disk images, orchestration templates, and other artifacts
  - ▶ Chameleon >130,000 images, >35,000 orchestration templates and counting
- ▶ Testbed as “player” for environments



*Paper: “The Silver Lining”, IEEE Internet Computing 2020*

# WHAT IS MISSING?

- ▶ Packaging: complete, imperative, non-transactional, integrated (literate programming)

- ▶ Get access for reproducibility

- ▶ Discover/find experiments through various channels



- ▶ Package experiment in a way that is cost-effective but also user-friendly

- ▶ Give access for reproducibility

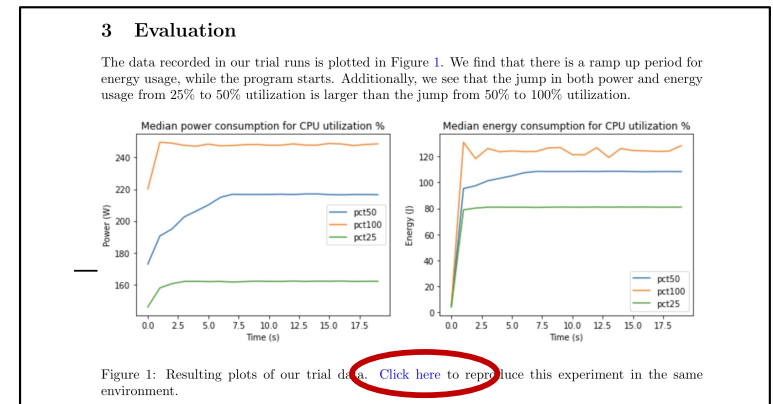
- ▶ Share work in progress; publish and advertise completed work





# TESTBED ACCESS WITH CHAMELEON DAYPASS

- ▶ Authors create a subproject with multiple short-term leases that are long enough to reproduce the experiment
- ▶ Readers click through data of a published experiment, request a daypass, and reproduce either the experiment or data analysis



Chameleon About Learn Experiment Blog Log in

Artifacts / Getting Started with Chameleon: Power management experiment

## Getting Started with Chameleon: Power management experiment

This notebook is a short example of how to use Chameleon notebooks to run a simple experiment, and analyze the data, using the python-chi interface.

Estimated duration: 1 hour  
Support contact: help@chameleoncloud.org

2 Sept. 29, 2021, 12:44 p.m. example experiment

Authors  
Jason Anderson (University of Chicago)  
Mark Powers (University of Chicago)

Request day pass

If you do not have an active Chameleon allocation, or would prefer to not use your allocation, you can request a temporary one from the PI of the project this artifact belongs to.


Versions

Version 2	Sept. 29, 2021, 12:43 p.m.
Version 1	Sept. 29, 2021, 12:37 p.m.



# SHARING AND FINDING EXPERIMENTS

*Familiar research sharing ecosystem*



*Digital research sharing ecosystem*



- ▶ Digital publishing with Zenodo: make your experiments citable via Digital Object Identifiers (DOIs)
- ▶ Trovi: sharing work in progress
  - ▶ BINs to collect all the artifacts, fine-grained sharing, versioning
  - ▶ Portal to browse, filter, and find interesting experiments
  - ▶ Integrated with Jupyter/Chameleon, Swift, Zenodo, and github (in progress)



# PARTING THOUGHTS

- ▶ Testbeds as a reproducibility fundamental
  - ▶ Public resource: instrument held in common
  - ▶ Instruments and methods enabling practical reproducibility
  - ▶ Packaging, access, and sharing
- ▶ Practical reproducibility == making reproducibility affordable
  - ▶ Time to package is important but time to repeat is critical!
  - ▶ Reproducibility marketplace
- ▶ Potential
  - ▶ Integration of research and teaching
  - ▶ Integration of Computer Science research in emergent applications

# *Think Big!*

*(with the help of a small reptile)*



[www.chameleoncloud.org](http://www.chameleoncloud.org)