



www.chameleoncloud.org

PRACTICAL REPRODUCIBILITY WITH CHAMELEON

Kate Keahey

keahey@anl.gov

University of Chicago, Argonne National Laboratory

July 21, 2021

CDAC Summer Lab Speaker Series



CHAMELEON IN A NUTSHELL

- ▶ We like to change: a testbed that adapts itself to your experimental needs
 - ▶ Deep reconfigurability (bare metal) and isolation
 - ▶ 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 distributed over 2 sites (UC, TACC) connected with 100G network
 - ▶ Diverse: ARMs, Atoms, FPGAs, GPUs, Corsica switches, etc.
- ▶ Cloud++: leveraging mainstream cloud technologies
 - ▶ Powered by OpenStack with bare metal reconfiguration (Ironic) + “special sauce”
 - ▶ Blazar contribution recognized as official OpenStack component
- ▶ We live to serve: open, production testbed for Computer Science Research
 - ▶ Started in 10/2014, available since 07/2015, renewed in 10/2017, and just now!
 - ▶ Currently 5,500+ users, 700+ projects, 100+ institutions, 300+ publications



BY THE NUMBERS

300+
Papers
published

45
Countries

700+
Projects

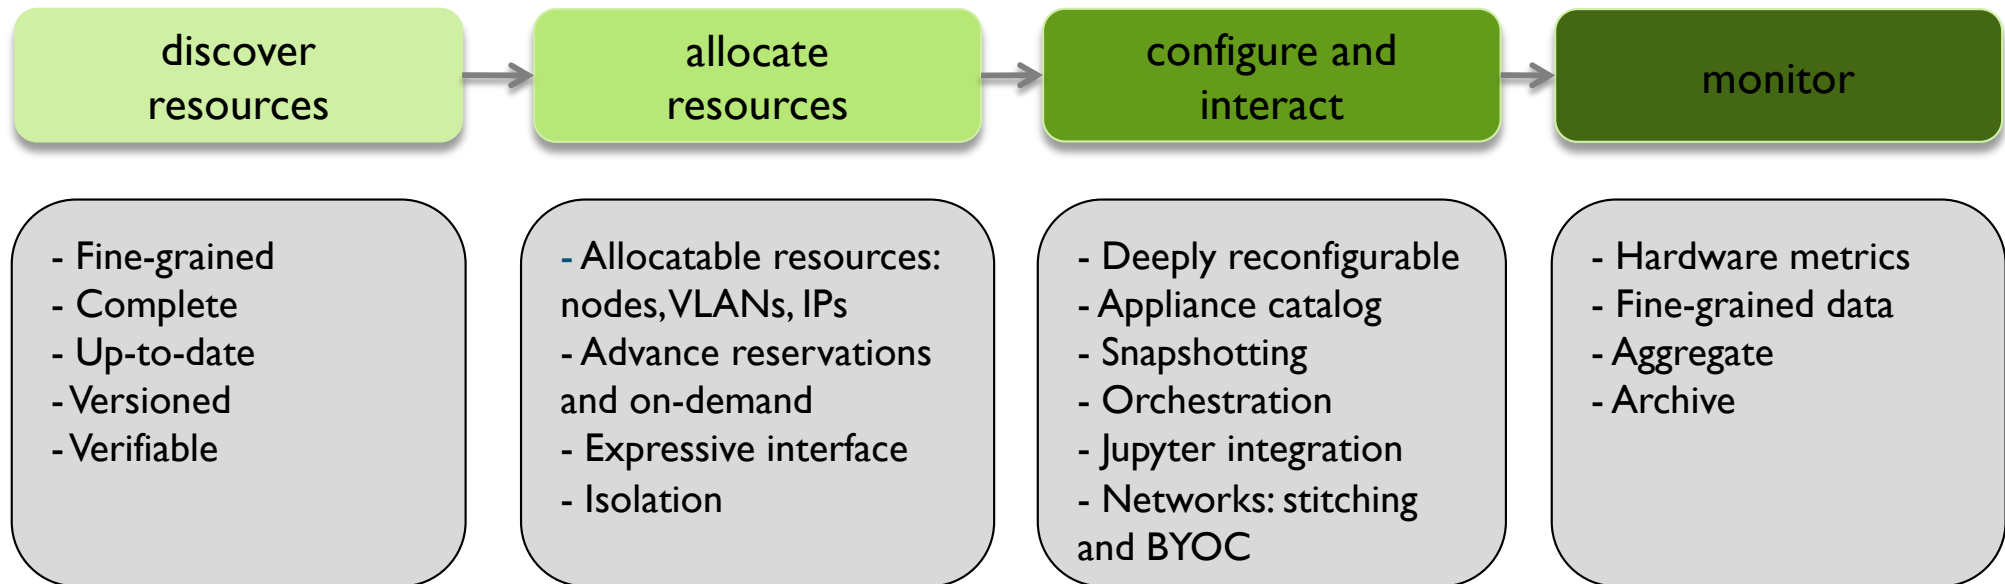
160+
Institutions

5,500+
Users

5+
Years Old

and 3+ more
years to grow!

CHI EXPERIMENTAL WORKFLOW



*Authentication via federated identity,
Interfaces via GUI, CLI and python/Jupyter*

REPRODUCIBILITY DILEMMA

Should I invest in making my experiments repeatable?

Should I invest in repeating this experiment?



Should I invest in more new research instead?

Is it faster to reinvent the wheel?

- ▶ Reproducibility has a cost: what is the right value proposition?
- ▶ Value of repeating experiments: variation and entertainment
- ▶ **Can experiments be as sharable/repeatable as papers are today?**
- ▶ **Holy grail:** reducing time to package (side-effect) and time to repeat

PRACTICAL REPRODUCIBILITY

- ▶ Public resource: sharing hardware via instruments held in common
- ▶ Clouds: sharing experimental environments
 - ▶ Disk images, orchestration templates, and other artifacts
- ▶ What is missing?
 - ▶ Telling the whole story: hardware + experimental container + experiment workflow + data analysis + story – literate programming
 - ▶ The easy button: it has to be easy to package, easy to repeat, easy to find, easy to get credit for, easy to reference, etc.
 - ▶ Nits and optimizations: declarative versus imperative, transactional versus transparent

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

REPRODUCIBILITY BUILDING BLOCKS

- ▶ Hardware: the baseline
 - ▶ >105 hardware versions over 5 years
 - ▶ Expressive allocation
- ▶ Clouds: images and orchestration
 - ▶ >130,000 images, >35,000 orchestration templates and counting
 - ▶ Portability and federation
- ▶ Packaging and repeating: integration with JupyterLab
- ▶ Share, find, publish and cite: Trovi and Zenodo



PACKAGING SHARABLE EXPERIMENTS



Literate Programming with Jupyter

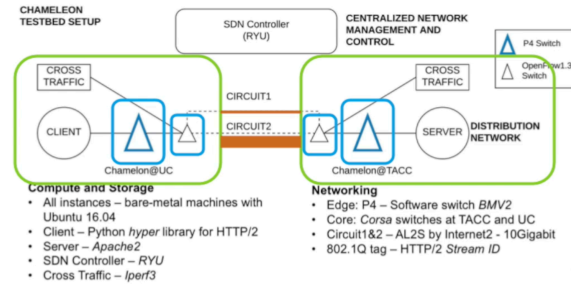
```

File Edit View Run Kernel Help
Trestle@redhat:~/jupyterhub -- LCN@demo2018@lynx
Python 3
2. Start P4 switches
[11]: from chi_ssh import Remote
import time
from toph.autonotebook import toph
p4_uc = Remote(ip='192.168.1.177')
p4_tc = Remote(ip='192.168.1.178')
for switch in toph([p4_tc, p4_uc]):
    switch.run('sudo p4ll -f 100%')
    time.sleep(2)
switch.run('screen -d -m sudo ipnet/switch -l 100m0.182 -l 100m0.282 -l 100m0.282 -l 100m0.282')
toph.write('Switch {} restarted.'.format(switch.host))
Switch 192.168.188.177 restarted. 100% 2/2 (0:00-0:00, 2.70s)
Switch 192.168.188.178 restarted.
Switch 192.168.188.177 restarted.
3. Start cross-traffic
[11]: from chi_ssh import Remote
import time
cross_uc = Remote(ip='192.168.188.187')
cross_tc = Remote(ip='192.168.188.187')
time.sleep(2)

```



*Experimental storytelling:
ideas/text, process/code, results*

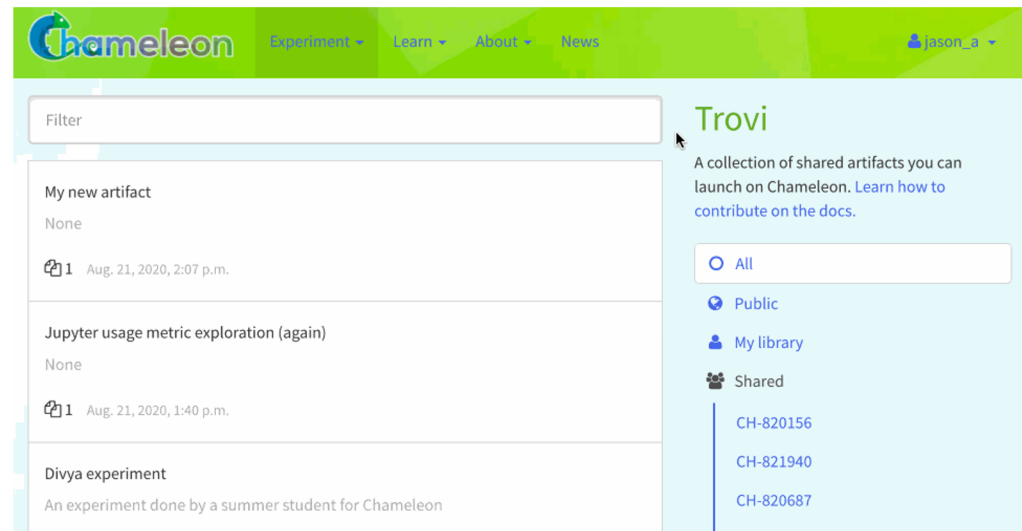



Complex Experimental containers

- ▶ Repeatability by default: Jupyter notebooks + Chameleon experimental containers
 - ▶ JupyterLab for our users: use jupyter.chameleoncloud.org with Chameleon credentials
 - ▶ Interface to the testbed in Python/bash + examples (see LCN'18: <https://vimeo.com/297210055>)
 - ▶ Especially for highly distributed experiments (CHI@Edge) notebook as terminal multiplexer

Paper: "A Case for Integrating Experimental Containers with Notebooks", CloudCom 2019

TROVI: CHAMELEON'S EXPERIMENT PORTAL



Create a new packaged experiment out of any directory of files in your Jupyter server. It is private to you unless shared. Supports sharing similar to Google Drive.


Any user with a Chameleon allocation can find and "replay" the packaged experiment.

SOME RESOURCES

- ▶ Trovi depositions
 - ▶ <https://www.chameleoncloud.org/experiment/share/?filter=tag%3Aexperiment>
 - ▶ Go to Chameleon Trovi and filter on “experiment” to find packaged experiments
 - ▶ Can be used to play with or teach
- ▶ Youtube channel, student screencasts:
 - ▶ AlexNet
 - ▶ DAWN Bench
 - ▶ LinnOS
 - ▶ DTRS experiment from tinyTailFlash
- ▶ SC20 poster: “Future-Proof Your Research: Designing for Replicability and Reproducibility”

SHARING EXPERIMENTS: PUBLICATION

Familiar research sharing ecosystem



Digital research sharing ecosystem



- ▶ Trove: a digital sharing platform
 - ▶ Make your experiments sharable within a community of your choice with one click
 - ▶ A library of reproduced experiments from foundational papers for research and education (see e.g., Brunkan et al., “Future-Proof Your Research”, SC20 poster)
- ▶ Integration with Zenodo: make your experimental artifacts citable via Digital Object Identifiers (DOIs) (export/import)
- ▶ Coming soon: the Chameleon daypass!

zenodo

PARTING THOUGHTS

- ▶ Time to package is important – but time to repeat is critical!
- ▶ Reproducible experiments rely on the existence of public resources
- ▶ Use-what-you-have approach: leveraging testbeds/clouds, existing digital artifacts, frameworks, patterns, etc. has the potential to lower the “price” of reproducibility and make it affordable
- ▶ We need to create a “marketplace” for repeating research
 - ▶ Repeatability and reproducibility can be thought of as the same thing at different “price points”
 - ▶ Recognition for published digital artifacts (software, data, experiments, etc.)
 - ▶ Starting early: education is an unappreciated tool for fostering reproducible research
- ▶ Coming soon: Chameleon daypass and repeatability hackathon!



We're here to change

www.chameleoncloud.org

keahey@anl.gov