



www.chameleoncloud.org

CHAMELEON: AN INNOVATION PLATFORM FOR COMPUTER SCIENCE RESEARCH AND EDUCATION

Kate Keahey

Mathematics and CS Division, Argonne National Laboratory

CASE, University of Chicago

keahey@anl.gov

February 3rd, 2021

MAGIC meeting

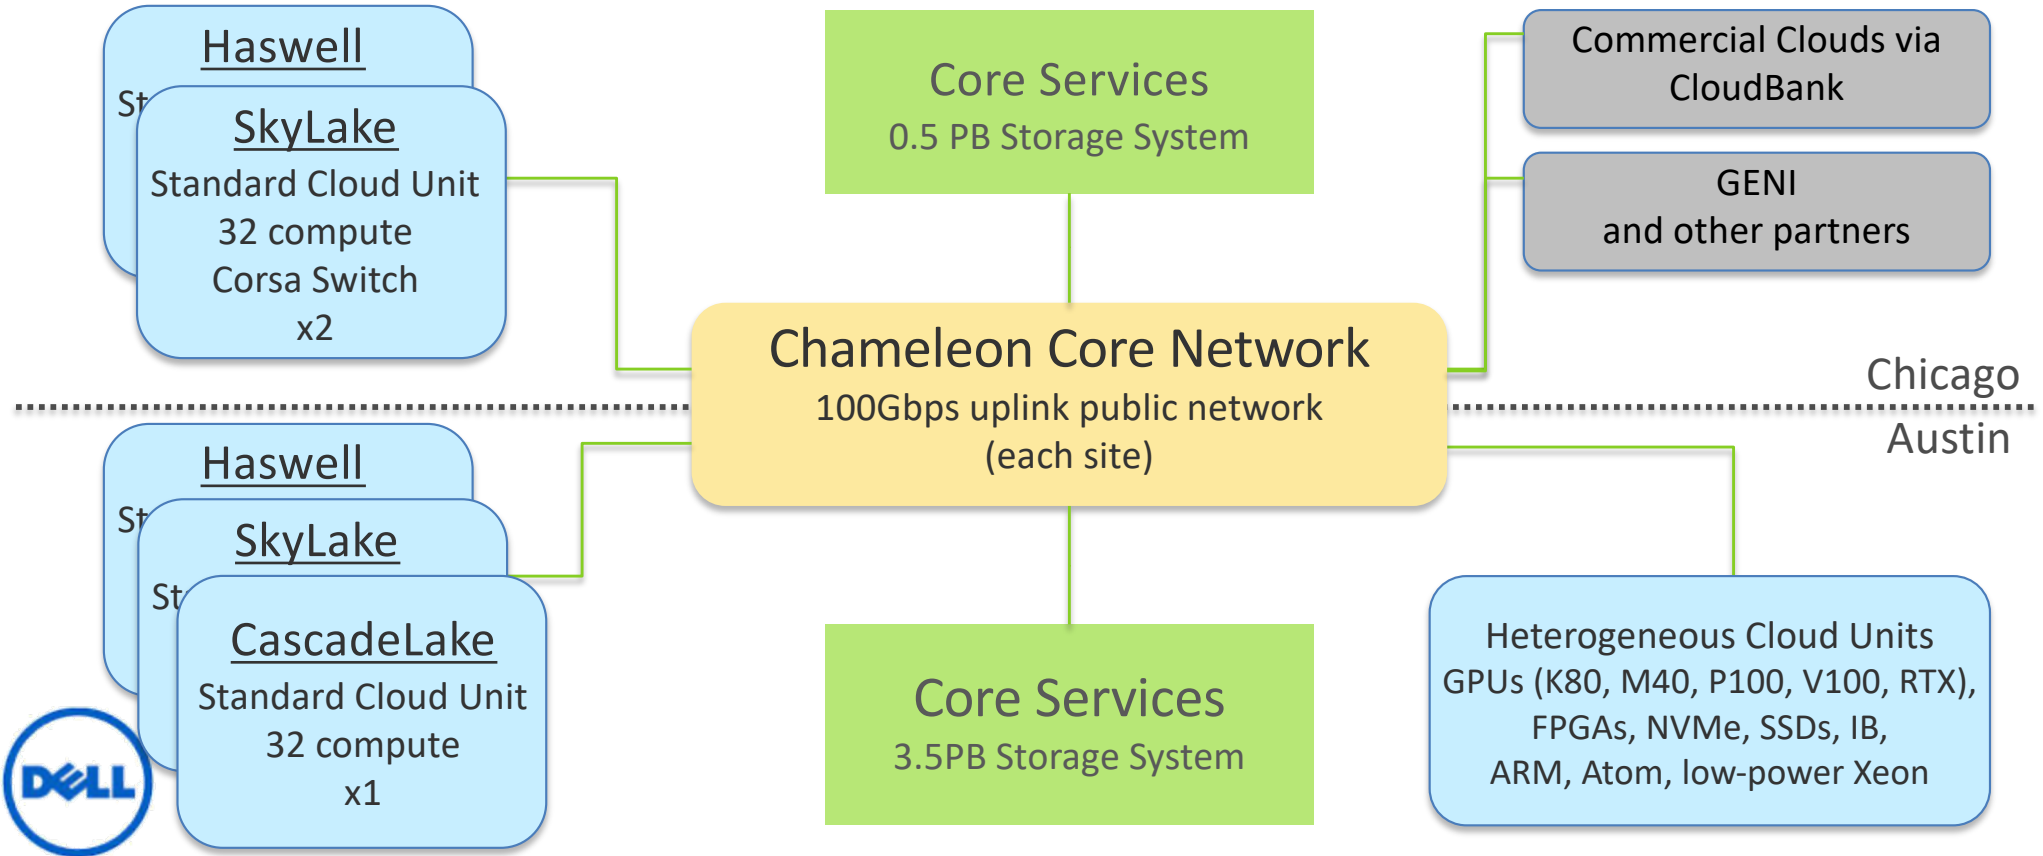


CHAMELEON IN A NUTSHELL

- ▶ 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, Corsas switches – and more to come!
- ▶ 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.
- ▶ 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, recently renewed till 10/2024!
 - ▶ Currently 5,000+ users, 600+ projects, 100+ institutions, 300+ publications



CHAMELEON HARDWARE




CHAMELEON HARDWARE (EXISTING)

- ▶ “Start with large-scale homogenous partition”
 - ▶ 12 Haswell Standard Cloud Units (48 node racks), each with 42 Dell R630 compute servers with dual-socket Intel Haswell processors (24 cores) and 128GB RAM and 4 Dell FX2 storage servers with 16 2TB drives each; Force10 s6000 OpenFlow-enabled switches 10Gb to hosts, 40Gb uplinks to Chameleon core network
 - ▶ 3 SkyLake Standard Cloud Units (32 node racks); Corsa (DP2400 & DP2200) switches, 100Gb uplinks to Chameleon core network
 - ▶ CascadeLake Standard Cloud Units (32 node rack) , 100Gb uplinks to Chameleon core network
 - ▶ Allocations can be an entire rack, multiple racks, nodes within a single rack or across racks (e.g., storage servers across racks forming a Hadoop cluster)
- ▶ Shared infrastructure
 - ▶ 3.6 + 0.5 PB global storage, 100Gb Internet connection between sites
- ▶ “Graft on heterogeneous features”
 - ▶ Infiniband with SR-IOV support, High-mem, NVMe, SSDs, P100 GPUs (total of 22 nodes), RTX GPUs (40 nodes), FPGAs (4 nodes)
 - ▶ ARM microservers (24) and Atom microservers (8), low-power Xeons (8)

COMING SOON TO A TESTBED NEAR YOU...


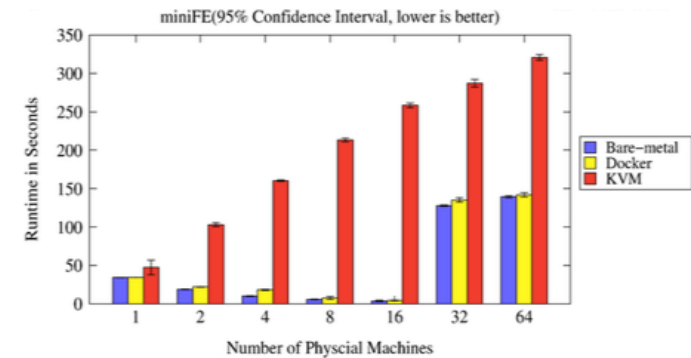
- ▶ New servers for old: upgrading Haswell to CascadeLake at TACC and UC
- ▶ More accelerators: V100s, AMD + AMD GPU, Xilinx
- ▶ Storage hierarchy options: Optane, a mix of enterprise NVMe and SSDs
- ▶ Composable hardware: IB rack with GPU/storage nodes (UC), LiQid (TACC)
- ▶ Networking: replicating FABRIC Hank design, P4 switches
- ▶ IoT devices & CHI@Edge
- ▶ Also a strategic reserve – we want to hear from you!

EXPERIMENTAL WORKFLOW



VIRTUALIZATION OR CONTAINERIZATION?

- ▶ Yuyu Zhou, University of Pittsburgh
- ▶ Research: lightweight virtualization
- ▶ Testbed requirements:
 - ▶ Bare metal reconfiguration, isolation, and serial console access
 - ▶ The ability to “save your work”
 - ▶ Support for large scale experiments
 - ▶ Up-to-date hardware



SCI5 Poster: “Comparison of Virtualization and Containerization Techniques for HPC”


CLASSIFYING CYBERSECURITY ATTACKS

- ▶ Jessie Walker & team, University of Arkansas at Pine Bluff (UAPB)
- ▶ Research: modeling and visualizing multi-stage intrusion attacks (MAS)
- ▶ Testbed requirements:
 - ▶ Easy to use OpenStack installation
 - ▶ A selection of pre-configured images
 - ▶ Access to the same infrastructure for multiple collaborators

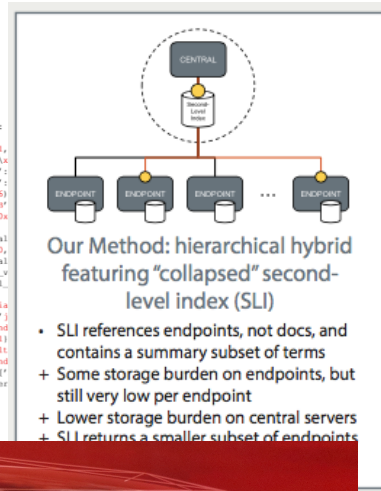


DATA SCIENCE RESEARCH

- ▶ ACM Student Research Competition semi-finalists:
 - ▶ Blue Keleher, University of Maryland
 - ▶ Emily Herron, Mercer University
- ▶ Searching and image extraction in research repositories
- ▶ Testbed requirements:
 - ▶ Access to distributed storage in various configurations
 - ▶ State of the art GPUs
 - ▶ Easy to use appliances and orchestration

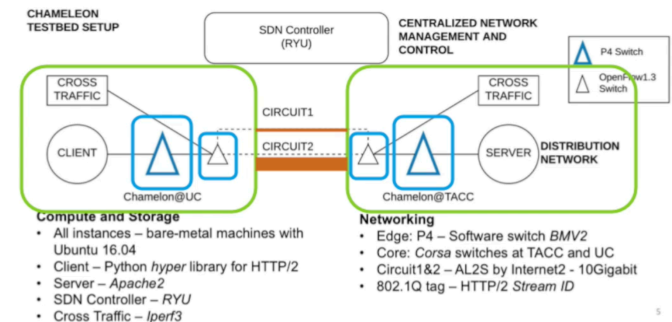


```
{  
  'header_info':  
    'file': '237',  
    'file_unit': '1',  
    'file_verification':  
    'file_density':  
    'dpi': (96, 96)  
    'image_mode': 'rgb'  
    'dimensions': '930x'  
  'color': {  
    'mean_pixel_val'  
    'extrema': (0,  
    'mode_pixel_val'  
    'median_pixel_v'  
    'std_dev_pixel_  
  'system': {  
    'path': '/media'  
    'extension': '}'  
    'file': 'img.img'  
    'size': 1158111  
    'image_text': ['hall'  
    'name_tags': ['Herron'  
    'svm_class_tags': ['  
    'mean_colors_cluster
```



ADAPTIVE BITRATE VIDEO STREAMING

- ▶ Divyashri Bhat, UMass Amherst
- ▶ Research: application header based traffic engineering using P4
- ▶ Testbed requirements:
 - ▶ Distributed testbed facility
 - ▶ BYOC – the ability to write an SDN controller specific to the experiment
 - ▶ Multiple connections between distributed sites
- ▶ <https://vimeo.com/297210055>



LCN'18: “Application-based QoS support with P4 and OpenFlow”

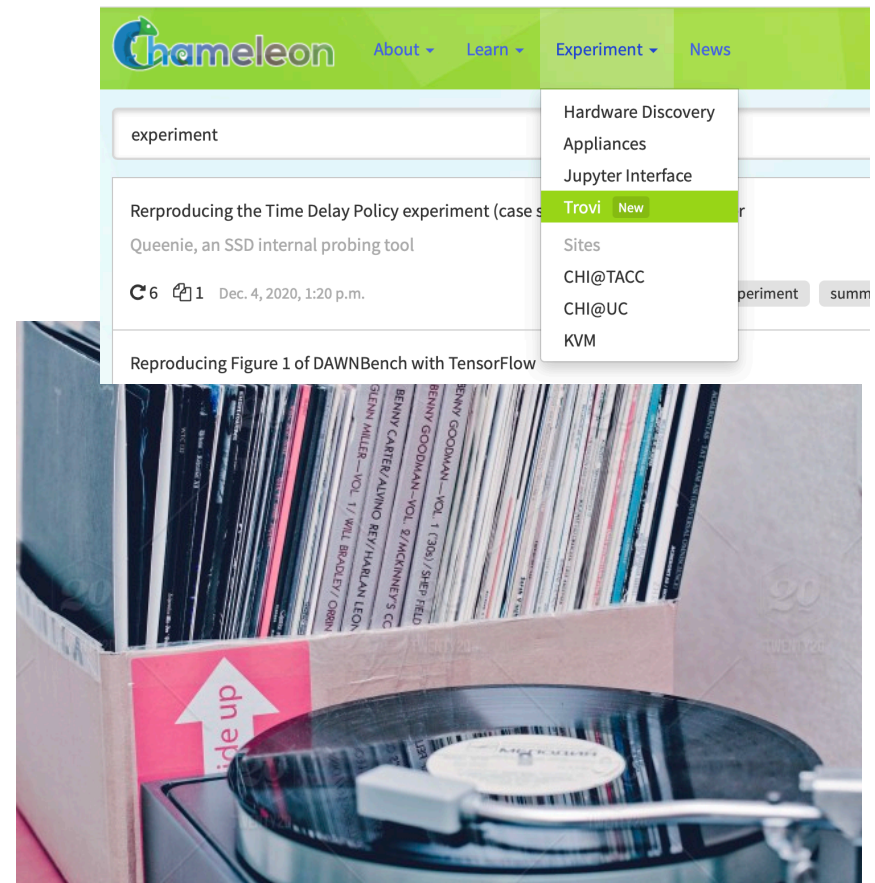
TOWARDS SHARING EXPERIMENTS

- ▶ Towards a world where experiments are as sharable as papers today
- ▶ Instruments held in common: sharing hardware
- ▶ 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


TROVI: SHARING EXPERIMENTS

- ▶ Testbed is a “player” for existing experiments/research
- ▶ Trovi: find existing experiments
 - ▶ Library including quickstarts (Python, OpenFlow, power management), and packaged of experiments from foundational papers
- ▶ Trovi: share experiments
 - ▶ Package your experiment, create a bundle, edit sharing settings, add keywords



PUBLISHING

Familiar research sharing ecosystem



Digital research sharing ecosystem



- ▶ Digital publishing with Zenodo: make your experimental artifacts citable via Digital Object Identifiers (DOIs)
- ▶ Integration with Zenodo
 - ▶ Export: make your research citable and discoverable
 - ▶ Import: access a wealth of digital research artifacts already published
- ▶ Towards making research findable: the digital sharing platform



PHASE 3 ADVERTISEMENT

- ▶ Already here: access via federated accounts!
- ▶ New hardware: traditional servers, new GPUs and FPGAs, storage upgrades (FLASH arrays), composable hardware (LiQid), networking (P4, integration with FABRIC), IoT devices -- and strategic reserve
- ▶ New capabilities: federation, Bring Your Own Device (BYOD) & CHI@Edge, networking (allocatable switches with P4), core capability improvements
- ▶ Infrastructure: CHI-in-a-Box, integration with production infrastructures
- ▶ Research sharing: better methods of experiment packaging, publishing, and discovery, digital experiment libraries, engagement

PARTING THOUGHTS

- ▶ Science does not stand still: laying the pavement as you walk on it!
- ▶ Chameleon is a shareable research instrument – but it is also a sharing platform
- ▶ The easy button: sharing experiments as naturally as we share papers
 - ▶ Clouds help us package experimental environments as a side-effect of using them and serve as “player” for such experiment
 - ▶ Literate programming is a convenient vehicle for “closing the gap”: packaging the whole experiment
 - ▶ Critical mass of research content: check out what you can use!
- ▶ Come, help us change: chameleoncloud.org



We're here to change – come and change with us!

www.chameleoncloud.org