CI/CS WORKSHOP

THE COMMUNITY TOGETHER

Panel: Ups and downs of cloud computing in open science

Ilya Baldin, CI CoE Pilot, RENCI/UNC Chapel Hill

Cloud Computing in Open Science

- The community is grappling with cloud adoption. It is a space of many tradeoffs
  - Costs, capabilities, staff training,
- Many practical questions
  - Architectures and best practices for transition
  - Security, monitoring, porting the software to ‘cloud native’
- Best of both worlds?
  - Hybrid infrastructures
  - Private vs. Public clouds

Interesting space for experiments pushing the limits

Open Science Grid About News Contact Docs

SDSC and IceCube Center Conduct GPU Cloudburst Experiment

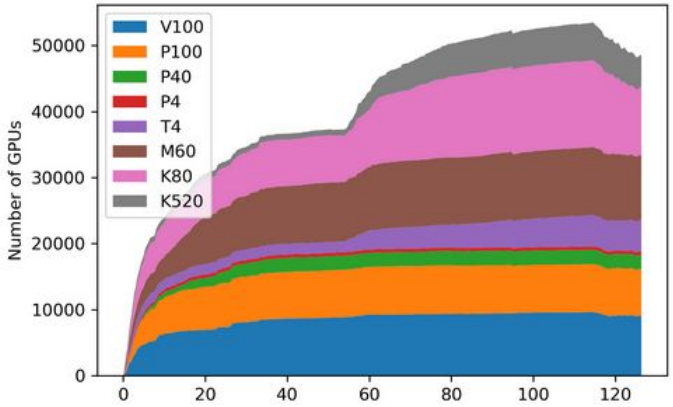

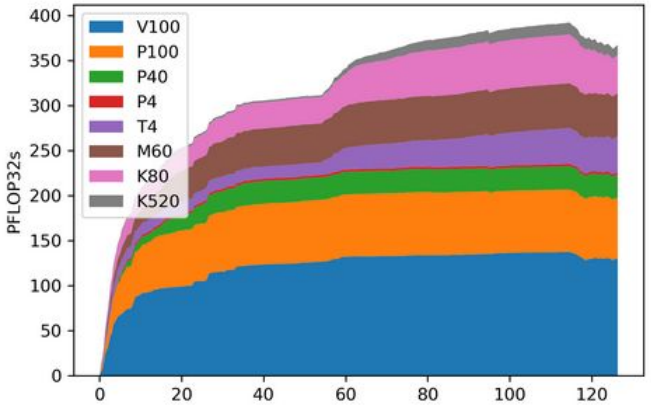

Today's panelists

- Benedikt Riedel, IceCube
- David Hancock, Jetstream/IU
- Mike Stanfield, RSOC/IU


CI/CS WORKSHOP
ResearchSoc
CI CoE
PILOT

Panelist questions

- Which parts of scientific CI are best suited for cloud transition?
- How best to think about hybrid infrastructures (on-prem + cloud) - benefits of keeping parts of CI on prem?
- Strategies for dealing with large scientific datasets and the cloud viz storage and transfer costs?
- Risk of vendor lock in and strategies for dealing with it
- Implications to staff expertise and training for cloud or hybrid science CI
- Accurately evaluating and projecting costs of cloud transition
- Best practices and pitfalls of adapting scientific application stacks to be ‘cloud-native’
- Monitoring cloud and hybrid infrastructures
- Security implications for running portions of scientific CI in the cloud, including identity management
- Building a productive mutually beneficial relationship with the cloud provider
- Going beyond the ‘Big Three’ (AWS, Azure, Google) - benefits of smaller/niche cloud providers
- Scientific computing at the edge - exploring the value of edge cloud providers to science

Thank you