

CI/CS WORKSHOP

THE COMMUNITY TOGETHER



ResearchSOC



CI CoE
PILOT


IceCube Computing in the Cloud

Benedikt Riedel

Global Computing Coordinator for IceCube Neutrino Observatory

UW-Madison

CI/CS WORKSHOP


 ResearchSOC

 CI CoE PILOT

Why am I here?

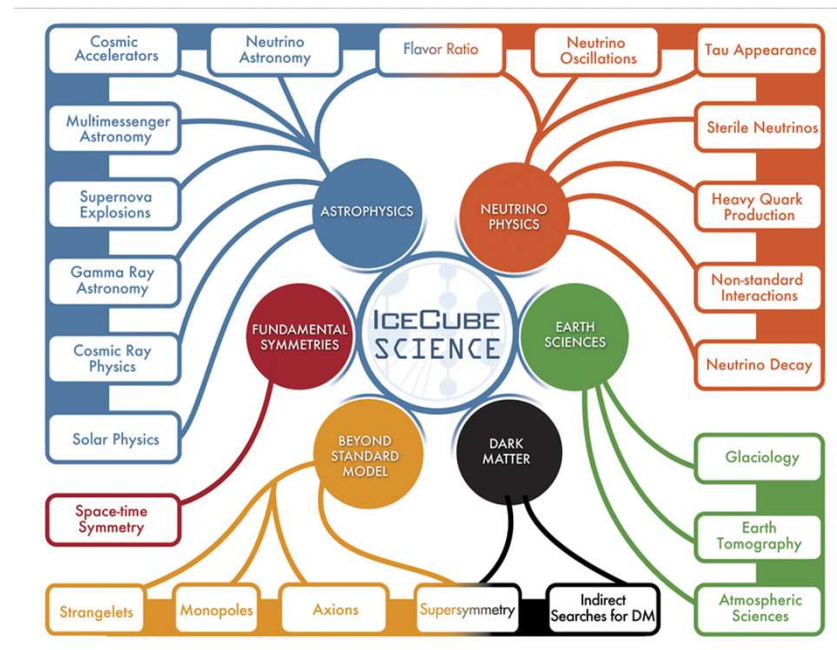
- Part of Phase 1 of [Exploring Clouds for the Acceleration of Science](#) (E-CAS)
- GPU Cloud Burst Experiment with collaborators from UCSD/SDSC - 51000+ GPUs across 3 cloud providers

IceCube



CI/CS WORKSHOP

- Neutrino Detector deployed in the Antarctic ice shelf
- 5160 detectors buried in the ice shelf
- Broad Science reach
- Classical Particle Physics Computing



CI CoE PILOT

E-CAS


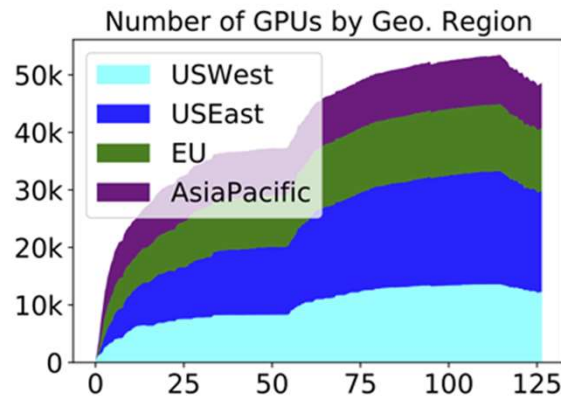

- Scaling
 - Multi-Messenger Astrophysics - IceCube wants to notify community about events as soon as possible
 - Workload scales with the no. of cores applied - Where to get cores “on-demand”? - Cloud
- Novel Compute Architectures
 - FPGAs, TPUs, ARM, etc. - Testing “new” architectures
 - Will it break? - Scaling to test for use on HPC

What did we learn?

- Cloud is great for on-demand scaling - Need 40,000 cores now!
- Cloud has good return on investment (ROI) the right application - CPU heavy, little data movement
- Funding (overhead) and university bureaucracy are the big hurdle - 1.5 years to sign contracts, etc.

GPU Cloud Burst

- Two Experiments (so far):
 - Go Big! - Aggregate 51500 GPUs across 3 cloud providers/28 cloud regions, \$50-150k
 - Go Efficient - Aggregate ~15000 most cost-efficient GPUs (T4, V100, P40) across 3 cloud providers/28 cloud regions, 1 ExaFLOP-hour, \$60k



What did we learn?

- IceCube's workflow can scale - Thanks to HTCondor
- Cloud is expensive for the wrong application - Replacing all of IceCube resources would cost O(\$50-100M) per year
- Data Movement can blow the budget - Networking cost > compute cost
- Social Engineering Needed
 - University has to be willing to do the leg-work in establishing partnerships
 - Need people inside the cloud vendors to advocate - Fish in a big pond

Cloud Extends Beyond


- “Cloud” not just AWS, Azure, GCP, etc.
- Business Services - Email, project management, human resources, etc.
- Data Center Services - Monitoring for CI, control planes for CI, etc.

Where do we see the cloud?

- ROI is a big question
 - Great - E-mail, Code Repository
 - Depends on scale/application
 - Jira, Trello, etc. - Expensive for large groups
 - Computing - CPU-heavy and little data
 - Replacing our CI with the cloud - No, simply too expensive
- Social Engineering needed at the funding agency level, university, and cloud providers - CloudBank, but for all

Questions?

CI/CS WORKSHOP


 ResearchSOC

 CI CoE PILOT

Thank you!

E-mail: briedel@icecube.wisc.edu

CI/CS WORKSHOP

 ResearchSoc

 CI CoE
PILOT