# Mapping Campus Research Computing to the Cloud

High Performance Research Computing
DIVISION OF RESEARCH

Michael Lau, Stuti Trivedi, Tri Pham, Emily Murphy and Dhruva Chakravorty
High Performance Research Computing, Texas A&M University, College Station, Texas

CC* CYBERTEAM SWEETER

## Introduction

Emerging research workflows are likely to benefit from a hybrid computing environment, that offers seamless integration between on-campus and off-campus cloud resources. Public and Federal clouds offer researchers access to computing platforms that may not be available on their campus. The large number of cloud offerings makes cost management, and workflow transitions to appropriate platforms challenging. While much has been accomplished in improving access, matters are complicated by the ever-expanding range of available resources on the cloud, dynamic billing scenarios, changing computing locations, and institutional security and data-management policies. Successfully mapping workflows from on-campus resources to the cloud, and leveraging the available cost structures to find economical cost models are critical steps to enabling researcher access to this vast resource.

## Research Question

- What are the economics and resource configuration models best suited for open-science computing researchers in the Cloud?

## Methods

Acknowledging the diverse campus and cloud offerings, the Terra Cluster at Texas A&M HPRC with 320 nodes and 5PB of general storage was used as a benchmark. Amazon Web Services (AWS) was the cloud provider of choice for this study.

### Resource Mapping:

- An exhaustive resource analysis of scientific jobs of Terra for an academic year was performed to identify researcher usage patterns
- All AWS offerings were studied using publicly available information since March 2020 and was last updated on Feb 26, 2021
- Resource mapping was dictated by storage and networking capabilities (EFA) associated with the AWS instances
- Intel Xeon CPUs on Terra were mapped to c5n.18xlarge (CPU, Intel Xeon) instances
- NVIDIA K89 GPUs on Terra was mapped to p3dn.24xlarge (GPU, NVIDIA V100) instances based on the ratio of maximum double precision FLOPs, the common model of usage in scientific computing

### Cost Management:

- Terra jobs were priced with the cheapest instance/s where the instance satisfied the required memory, cores, and GPU requirements
- Calculated the equivalent cost to run each job from Terra in 2020 on the AWS cloud using market prices for that date
- The "***direct equivalence***" method calculated the cost of allocating an equivalent "always-on" cluster on AWS that would mimic Terra.
- A "***flexible orchestration***" method was also developed to allow for pricing efficiencies. Here, jobs were categorized into single node CPU, single node GPU, multi-node CPU, and multi-node GPU jobs. Multinode jobs were priced against the best fit EFA enabled CPU and GPU (p2/p3) instances.

### Storage Costs:

Storage was priced in relation to Terra's 5 PB storage solution against AWS Lustre 200 MB/s/TiB baseline and AWS standard S3. Importantly egress and S3 read and write operations are not included in this model.

## Results

### Research Computing Workload Analysis

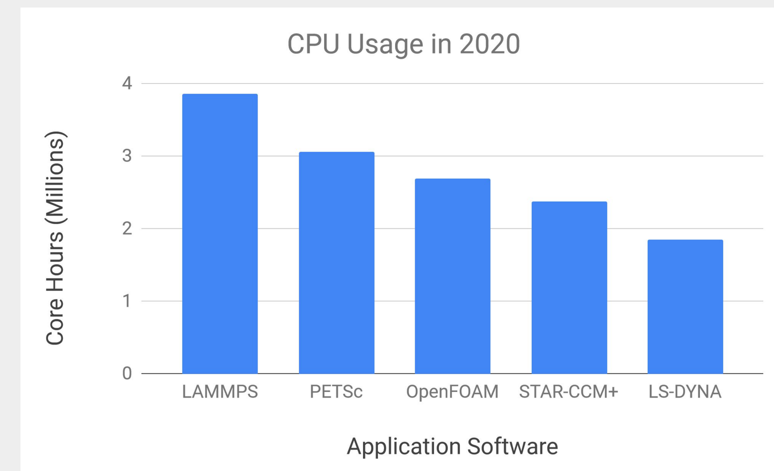


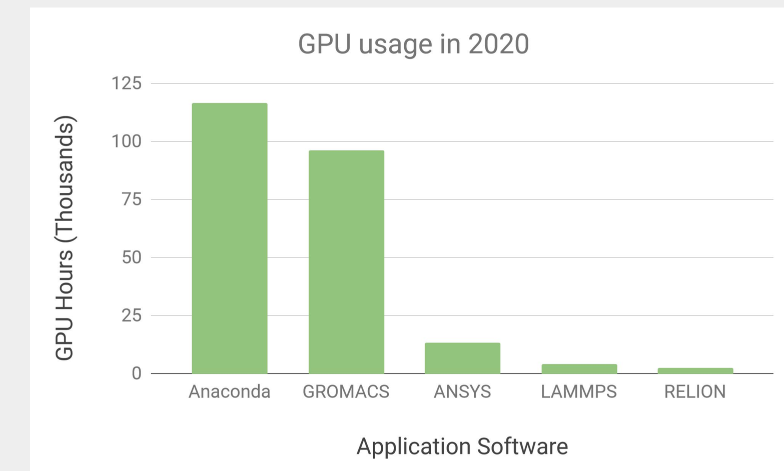


- Software module usage was collected by cross-referencing Slurm resource allocation logs with LMOD "userload" logs. Job resources are associated with each module loaded in the job
- LAMMPS, OpenFOAM, and other applications have high CPU utilization on Terra. Gromacs, Ansys, and Anaconda (for AI/ML) represent the applications with the greatest GPU usage

### Computing in the Cloud

Annual costs were calculated using publicly available AWS pricing for "on-demand" and "reserved" instances using the "***direct equivalence***" method. This method assumes ready and high availability of resources and can not be combined with "spot pricing" on AWS that can offer up to 90% discounts but only support fault-tolerant workloads.

| AWS Pricing Tier | On-Demand | Reserved EC2 |
|---|---|---|
| GPU Compute | $820,251 | $253,339 |
| CPU Compute | $7,935,719 | $2,508,487 |
| Additional Discounts | $0 | $138,091 (5%) |
| **Final Compute Cost** | **$8,755,970** | **$2,623,735** |

### Storage Costs

| Storage Solution | Annual Cost |
|---|---|
| AWS Lustre 200 MB/s/TiB baseline | $17,400,000 |
| AWS S3 Standard | $1,260,000 |

*Storage costs do not include I/O and data egress charges. Researcher may choose to pay for additional throughput for Lustre.

### Efficiencies Through Flexible Orchestration

While the "***flexible orchestration***" with resource mapping found a total price of compute for all jobs ran in calendar year 2020 to be **$6,341,819.**

## Research Computing Bazaar Software

The Research Computing Bazaar was developed to automate the cost calculation of on-campus workloads using flex orchestration.

- Taking an API-driven approach, it automatically determines the price for each job on a campus cluster with actual cloud pricing data
- This is a flask/python approach built on a PostgreSQL database populated with cloud pricing and on-prem cluster workloads
- Containerization, load-balancing, and scale are design elements
- Cost data is stored and visualized through Apache Superset a visualization platform. The dashboard provides reports on the costs of running on-premise workloads on the cloud providers
- The dashboard offers daily, and cumulative costs for real dates

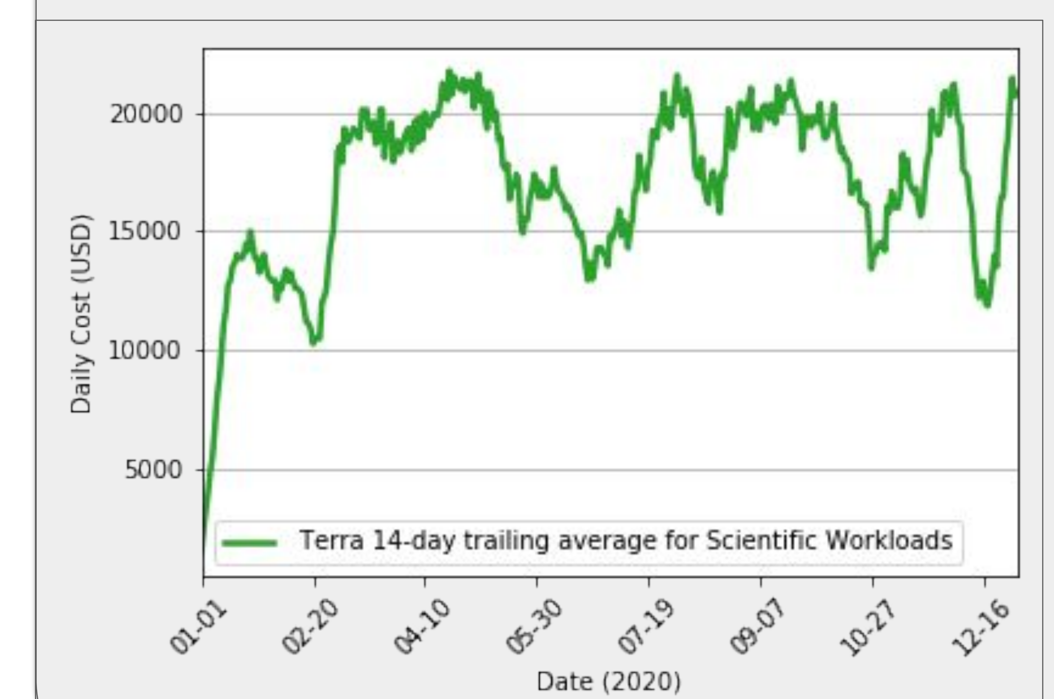


14-day trailing average cloud cost for 2020. Running costs are higher during periods of greater academic activity and costs decline around holidays, offering possible benefits from dynamic pricing models.

## Conclusions

- Federal and public clouds offer several computing resources and offer unique advantages to local systems such as near instantaneous procurement time.
- Resource mapping offers researchers access to new architectures but is a continuing process requiring almost daily updates.
- The flexible orchestration approach matches the needs of the scientific application to the best resource requires more work up-front, but is the more cost-effective strategy to cloud adoption. The lift-and-shift approach that recreates a campus cluster in the cloud offers an easier transition path for campus computing but is more expensive.

## Future Work

★ The Flexible Orchestration model will be extended to include predictive pricing, and support for AWS spot-pricing type ephemeral instances
★ Preliminary work with AI/ML models show promise in predicting cloud-pricing and resource matching
★ Costs associated with storage need to include I/O operations, campus direct-connect fees and egress fees could increase the reported prices

## References

1. Alan Chalker, Curtis W. Hillegas, Alan Sill, Sharon Broude Geva, and Craig A. Stewart 2020. DOI:https://doi.org/10.1145/3311790.3396642
2. Posey, Brandon, "Dynamic HPC Clusters within Amazon Web Services" (2016). *All Theses.* https://tigerprints.clemson.edu/all_theses/2392
3. Git repository: https://github.com/tamu.edu/HPRC/cloud-bazaar

## Acknowledgements

We gratefully acknowledged support from the National Science Foundation award number 1925764 for the CC* Cyberteam SWEETER, and staff at Texas A&M High Performance Research Computing, AWS, and TAMU IT.

NSF