cloud cloUT

Kelly Conlon, Juan Paez, Brock Davis, Nicholas Aufiero, Evan Chang, Jack Gaither

TEAM INTRODUCTION

- Team composed of students from the University of Texas at Austin and the University of Arkansas
- Mentored by staff at the Texas Advanced Computing Center
- Brought together diverse set of individuals with unique skills and ideas; focused on diversity of backgrounds and interests

Kelly Conlon: Computational Engineering Student at UT Austin, focused on scaling HPCG

Evan Chang: Electrical Engineering Student at UT Austin, worked on scaling IO500

Brock Davis: Studying Computer Science at UT Austin, worked on Reproducibility Challenge

Nicholas Aufiero: Computational Engineering & Plan II at UT Austin, focused on optimizing CESM

Juan Paez: 4th-year Electrical Engineering Student, worked on tuning Gromacs

Jack Gaither: Studying Computer Science at the University of Arkansas Fayetteville, focused on tuning linpack

SYSTEM DESIGN AND MANAGEMENT

- Head node on Azure Cyclecloud for cluster management
- CentOS 7 operating system with Slurm scheduler selected due to high performance and previous experience
- Multiple node arrays to allow us to select instances dynamically during competition
- BeeGFS parallel file system to share data between VM's, selected due to high throughput and ease of deployment
- Both Intel and GCC compilers will be used on an as-needed basis by each application; applications will be compiled ahead of time with training budget
- Build dependencies and applications prior to competition and keep in persistent storage

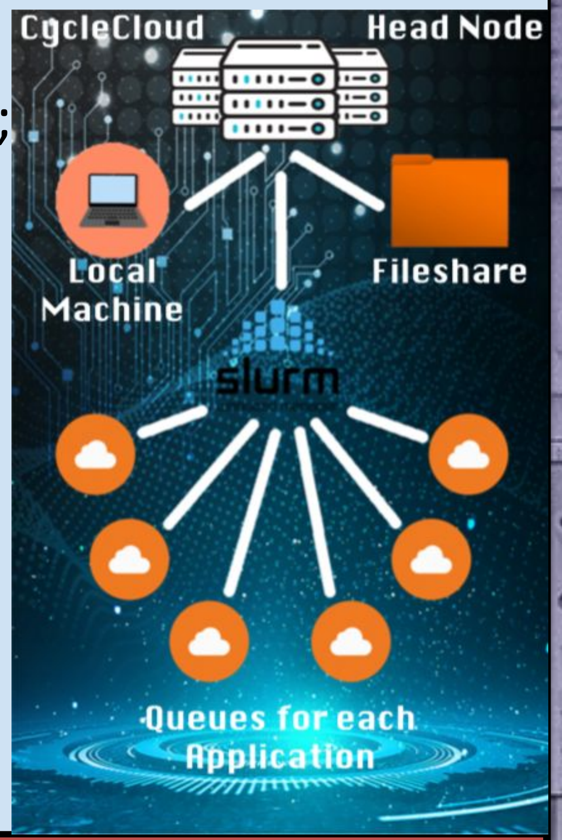

APPLICATION STRATEGIES

HPCG
- Depending on the size of the problem, we will utilize the appropriate number of nodes, probably in the range of 8-12 nodes
- We will start with the D-series, as they offer low latency, and high-speed local storage at low cost pricing

Linpack:
- Utilize the most cost efficient compute nodes available to Azure to scale HPL on a single node to maximize local compute usage due to high RAM and CPU availability

APPLICATION STRATEGIES, CONT'D

IO 500:
- The ability to toggle different IO500 tests allows us to systematically improve our overall score in a piecemeal manner
- Toggling variables and system specific information will optimize individual tests at a time

CESM:
- Allocate more processors to expensive components, and use the timing document to load balance from there
- Use the restart run capability to split runs into smaller time units and leverage spot instances
- H-series instances will be best due to high memory demand

Gromacs:
- Optimal hardware setup varies from problem to problem due to varying bottlenecks (long range vs. short range interactions) and problem size
- Start with 4-16 low-end N-series instances (e.g. NC6), add/remove nodes during tuning
- Eventually invest in high-end N-series to generate final results
- May use high-end non-GPU nodes (e.g. F-series) for problems without PME

MemXCT:
- Analyze performance on a variety of CPU-based and GPU-based nodes (e.g. D-series and N-series)
- Compare results with paper. Runs are relatively short, allowing to save funds for other areas.

Mystery Application:
- Select instances based on problem type (e.g. ND series for deep learning) and scale out on the spot

PREPARATION

- Develop initial familiarity with HPC and applications through TACC supercomputers and public resources

The Frontera Supercomputer at TACC -- #8 in Top 500

- Use cloud credits for testing and strategy development with Azure-specific instances
- Each team member becomes a domain expert for a single application, learning about its scalability properties, optimal parameters, compilation strategies, etc.
- Everyone responsible for understanding job management and instance deployment, Azure CycleCloud interface and basic tuning and scaling techniques for any application (CUDA stack, MPI, OpenMP)

WHY WE WILL WIN

- Actively communicating, meeting, and improving our approach since February.
- Diverse background skills specifically tailored to each team pairing,
- Leveraging our access to TACC resources and mentors