DESERT HEAT
ASU NAU UA

NORTHERN ARIZONA UNIVERSITY

ASU ARIZONA STATE UNIVERSITY

THE UNIVERSITY OF ARIZONA

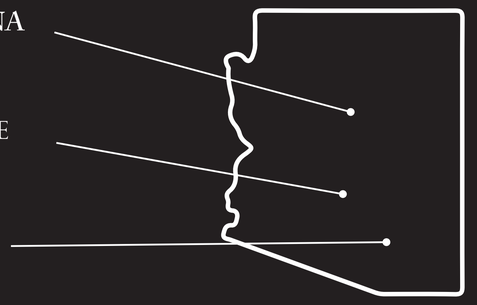

OUR SPONSORS

PARTNER

Hewlett Packard Enterprise

SC15 Austin, TX | hpc transforms.

MEET THE TEAM & SCHOOLS

We are the Arizona Tri-University Team, Desert Heat, consisting of students from Arizona State University, Northern Arizona University, and The University of Arizona. These schools are separated by many miles and are located in areas with distinct surroundings, ranging from beautiful desert landscapes to snowcapped mountains and ponderosa pine trees. Collectively, our three universities create a superb representation of the diversity found in the great state of Arizona.

Arizona State University (ASU) is a comprehensive public research university that is committed to excellence, access and impact. ASU measures itself by those it includes, not by those it excludes. ASU pursues research that contributes to the public good, and ASU assumes major responsibility for the economic, social and cultural vitality of the communities that surround it. Located within the metropolitan Phoenix area – a thriving urban sprawl in the middle of the Sonoran Desert, ASU is one of the largest universities in the United States, boasting over 83,000 enrolled students, of which 50,000 are on the Tempe campus alone. ASU has two students participating this year: Omri Mor and Derek Jones. Omri is a junior studying computer science, with a minor in mathematics. He currently works assisting the operations team at ASU Research Computing. Derek is in his third year at ASU, learning computer science and specializing in information assurance.

Northern Arizona University (NAU) is the smallest of the three Arizona universities, located in the mountains of Northern Arizona at 7,000 feet. NAU is represented by Mike Gilbert, Brandon Garling, and Jason Travis. Mike is a junior majoring in computer science and currently works in NAU's High Performance Computing department, assisting in cluster maintenance and user support. Brandon is a junior working towards his BS in computer science. He currently works on campus for ITS on the Systems Integration and Architecture team, helping maintain and create new and innovative tools to handle DNS, DHCP, LDAP, and CAS. This is his first time working with a high performance computing cluster. Jason is a senior studying computer science and working as a bioinformatics intern at Translational Genomics Research Institute (TGen North).

The University of Arizona (UA) is located in the heart of Southern Arizona. Established in 1885, The University of Arizona is the state's super land-grant university with two medical schools and revolutionizing the fields of space science, optics, biosciences, medicine, arts and humanities, business, and technology transfer all in pursuit of the mission of improving the lives of southern Arizonans. Although UA does not have current representation on the final team roster, the UA staff was actively involved in recruiting, mentoring and working with students to compete in the challenge. Due to various circumstances, the original members of the team from UA were unable to continue their participation in the cluster competition.

THE HARDWARE

Our cluster consists of seven nodes, each with 128GB of RAM and two Intel E5-2690 v3 12-core processors, for a total of 896GB of RAM and 168 cores capable of a theoretical 6.95 teraflops of computing power. The nodes use FDR Infiniband for fast communication and have two mirrored SSDs for rapid I/O. We are using Red Hat Enterprise Linux 6.6 for the operating system on each node as well as the Slurm workload manager to dispatch jobs.

We considered using accelerators, but decided against them after studying the various performance characteristics of the competition applications. We found that overall our Haswell processors provided better performance per watt than accelerators, due in part to the lack of accelerator support in some of the applications.

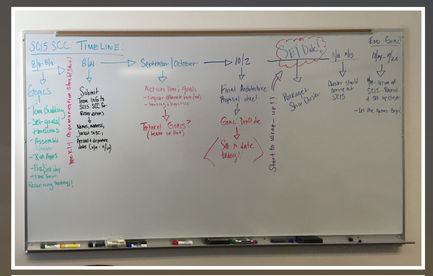

PREPARATION

Team Desert Heat is split geographically across hundreds of miles, making in person collaboration difficult. We tried to mitigate the lack of locality by using online meeting technologies.

One way we prepared for the competition was by meeting on Google Hangouts twice a week. At each meeting every member of the team was assigned responsibilities along with presenting the deliverables they had worked on each week. We used Slack and email distribution lists as our primary forms of communication outside of the biweekly meetings.

Another preparation technique involved dissecting the applications that may be utilized at the competition. For each application we discussed whether it was I/O or compute-bound and whether or not it could utilize accelerators. We then assigned team members to be responsible for certain applications. The individual team members worked on testing and optimizing their application and documented their progress. Once the applications were fully optimized and documented, each team member presented their findings to the team. We also attempted to cross-train where possible to ensure no single point of failure.

RUNNING & OPTIMZING APPS

We began with basic compiler configurations for each of the programs using both the Intel compiler suite, generously donated by Intel, and GCC. In all cases, we found the Intel compiler to be more efficient. Once we had the programs installed, we tested each with various sizes of sample data to ensure that they worked correctly. Next, we went through several iterations of making extra optimizations and retesting until we reached the optimum performance for each program.

When testing the applications on the cluster, we found that many of us were testing programs simultaneously, creating undesirable results and performance issues. We chose to use the Slurm workload manager to address this, which allowed us to maintain peak performance of the cluster at all times.

THE WINNING TEAM

Being new to the HPC community and with team members spread throughout the state, preparing for the competition did not come without its challenges. Examples of the challenges overcome included:

· Planning for turnover in the team, we recruited 12 potential student team members (four from each university): one primary team and one secondary team as a backup. Over the course of several months our team dwindled down to only five members due to various circumstances that affected participation.

· We had limited time with our competition cluster (one month total) as a result of a miscommunication with our first vendor and their cluster offering. The original cluster delivered was two generations old and wouldn't have been competitive. Our team's university leadership worked tirelessly to find another vendor partner to meet our technical requirements. Once our new vendor partners were identified, IT Partners and Hewlett Packard Enterprise, the new cluster was quickly sent into production. While the team waited for the competition cluster to be delivered, they worked on an academic cluster at ASU. Finally, on October 2nd we powered on our competition cluster and began final preparations for the competition.

· Due to unforeseen circumstances, a change in mentor leadership had to be made in the last two months leading up to the competition. Christopher Coffey, the HPC systems administrator at NAU, stepped into this role, sharing both his experience and the rare inspirational sense of wonder and discovery that make learning fun.

From these challenges and the many hours of hard work that went into compiling applications on our great cluster known as the Honey Badger, a winning team evolved. We all have our specific applications that we have mastered and know inside and out. We have spent countless hours making everything work as smoothly as possible. Our cluster is a well-oiled machine that will win this competition!

In addition to the mentorship provided by Chris and other members of university staff, we also worked with Logan Sankaran, an HPC Strategist at Hewlett Packard Enterprise. Logan graciously gave the team much of his time and expertise, sharing valuable information on the competition applications, as well as addressing any possible problems we ran into. The team would like to thank Logan for his time!