



## US/Pan-African HPC Workshop: On Common Ground @ SC16

Salt Lake City, Utah-US, Nov. 11-19, 2016

A team of international high performance computing (HPC) trainers and facilitators is planning an SC16 co-located training exercise for campus champions who support education and research applications for advanced cyberinfrastructure in rural, resource-constrained regions of the U.S. and sub-Saharan Africa.

**Call for participation released August 10, 2016; applications due Sept. 2, 2016**

Systems administrators, educators, researchers, cybersecurity professionals, tech project managers, and network engineers **who also serve as campus champions** in rural, resource-constrained regions are welcome to apply for travel support (including conference registration). National Science Foundation (NSF) eXtreme Science and Engineering Discovery Environment (XSEDE) Campus Champions from EPSCoR states and territories, and their African counterparts from Southern African Development Community (SADC) member states are especially encouraged. Non-EPSCoR XSEDE Campus Champions who perform significant EPSCoR or Minority Serving Institution (MSI) HPC education and outreach are also welcome to apply.

Final selections will be made by STEM-Trek's advisory board and volunteer sub-committee (with representation from the steering committee). **Awardees will be notified of their acceptance by September 15, 2016.** A variety of factors will be considered, including delegates' ability to contribute expertise and train others in one or more data- and computationally-intensive fields. Special consideration will be given to candidates from resource-constrained regions, and demographics that are under-represented in the use of HPC.

### Planned activities

Common challenges will be addressed with a series of technical workshops and discussions, including basic HPC skills training, open-source computation and data tools, and coaching in the art of advocacy so delegates are prepared to foster administrative and legislative support for locally-hosted HPC.

Social networks and affinity pairings will be formed, and participants will be encouraged to collaborate on posters and white papers that address common challenges. All will be encouraged to submit posters and papers to future CHPC National Meetings, XSEDE'XX, and SCXX conferences. Members of the planning committee and/or STEM-Trek's proposal adviser will offer continued oversight following the conference, and the activity's social network will be maintained and expanded by STEM-Trek.

## Introduction and Background

The XSEDE Campus Champions program has more than 200 members. Each of the 50 states, plus the Puerto Rico territory, is represented by one or more delegates. Champions help campus researchers, faculty, staff, and students find computational and data resources that accelerate the process of discovery and learning. They recommend the best solution for each challenge, and instruct others in the use of locally-hosted or nationally-distributed systems and services.

The NSF Experimental Program to Stimulate Competitive Research (EPSCoR) allows organizations in 25 states or territories to compete for NSF EPSCoR program opportunities. EPSCoR eligibility is determined by the level of NSF research funding received over the past three years. Consequently, many EPSCoR Campus Champions work in resource-constrained environments. Many can't afford to refresh computers on an industry-recommended cycle, and lack the time and money necessary to train, or support new and emerging "long-tail" research communities.

The Southern African Development Community (SADC) HPC Forum delegates, and others who manage HPC resources in sub-Saharan Africa, share many of the same responsibilities and challenges with their EPSCoR counterparts. There are always urgent and competing priorities for limited funds on their campuses. Therefore, all would benefit from learning how to build and present a stronger case for campus research computing which also fosters support for a greater national investment in CI.

Each US/Pan-African state is unique, with industrial landscapes that reflect their region's environmental factors, geology, minerals, and local economy. College campuses prepare the regional workforce, and the research accomplishments of each campus highlight its resources and body of expertise. Every state's research community makes a unique contribution to their nation's ability to compete in the global marketplace.

Many US/Pan-African states are sparsely populated, and a larger percentage earns a living from agriculture. The US has always been a leader in agricultural research and development, and sub-Saharan Africa is important to the global agricultural R&D community for a variety of reasons. There are more than 100 unique native plant species on the continent, and an emerging community of agricultural scholars who are eager to learn from US experts. The human capital framework represented by rural champions would benefit from mastering the computational and data tools employed by the agricultural research community, but they often lack access to resources, or the opportunity to train.

Global food security is a grand challenge that experts agree will be solved through a greater investment in agricultural research and development (including widely-distributed, geospatially-explicit, data-intensive research). Climate change and drought are causing new biotic stressors that affect the way agricultural diseases and pests behave, and the footprint of production is shifting. Consequently, tropical regions are of increasing importance, including Puerto Rico (US-EPSCoR) and Mauritius (SADC). There are similarities between the sisters that warrant research collaborations, including growing and economic conditions. For example, both import a large amount of their food while once productive land is no longer cultivated.

Because densely-populated, urban communities can more easily support subscription-based commodity networks and energy infrastructure, rural communities often face power and network shortcomings. Their universities are usually the last to be connected to high-speed research networks, and often have last-mile challenges if the schools can't afford to install the appropriate

switches or routers needed to engage. In these settings, light-at-the-end-point, computational, training and data solutions are useful.

Many rural regions are arid, and therefore struggle to power and cool HPC systems. Data center managers in the American Southwest could share lessons learned with counterparts from desert nations, like Botswana and Namibia, for example. They are currently suffering from the driest season in 30 years and their policy-makers must choose between greater bandwidth, or funding alternative ways to deliver fresh water for human consumption and agricultural use. Their champions would benefit from learning how to advocate on behalf of a greater capital investment in computational resources and services, despite competition for assets.

Due to the factors described above, it's likely these champions will support the technical resources and applications needed by the following research communities:

- Renewable Energy, Natural Resources, and Environment
- Food Safety, Nutrition, and Health
- Plant Health and Production, and Plant Products
- Animal Health and Production, and Animal Products
- Agricultural Systems and Technology
- Agricultural Economics and Rural Communities

### **History of the SADC HPC Forum**

The SADC HPC Forum formed in 2013 when the University of Texas donated their decommissioned, NSF-sponsored *Ranger* system to the South African Center for High Performance Computing (CHPC). CHPC divided 25 Ranger racks into ten individual clusters that were installed in universities in South Africa, Botswana, and Lesotho. Tanzania, Mauritius, Namibia, Zambia, and Zimbabwe also have HPC. In 2016, similar hardware was donated by the University of Cambridge, UK and it was installed in Madagascar. The CHPC continues to lead SADC training efforts and a dozen or so US and European HPC industry experts volunteer to advise as the SADC shared CI project continues to gain traction.

The Square Kilometre Array (SKA) project—the most ambitious and well-funded technology project ever developed—is being installed in the “radio quiet” Karoo region of South Africa. SKA is contributing computational and data science training (and resources) to some SADC HPC Forum participant schools to help prepare the indigenous workforce pipeline that will be needed for SKA's 50-year lifespan. Additionally, they are investing in the national research and education networks that will support SADC's shared CI.

Forum delegates have trained as a cohort since 2013, and it has been a successful exercise in science diplomacy. Among them are network engineers, sysadmins, educators, computational, and domain scientists. While the SADC region has multiple language and other cultural disparities, as they trained together with a common goal, this cross-functional team has coalesced, despite the differences. They are creating a procedural framework that will continue to inform policy-makers and newcomers. Their synergy will likely have a lasting impact that will translate effectively in other arenas of governance throughout the peaceful and diplomatic SADC region.

## Broader Impacts

The *broader impacts* of this effort are clear and far-reaching. When the Ranger system was donated to SADC, it had already helped many thousands of US researchers advance discovery and innovation during its production lifetime, before being supplanted by the more powerful NSF-supported Stampede system. The donation of the hardware allowed students and researchers at more than a dozen African universities to not only benefit from access to computational resources, but through the additional training and knowledge transfer provided by the NSF-supported fall 2015 training exercise at TACC, and ongoing efforts to train the SADC HPC Forum cohort as they master the skills necessary to deploy and operate cluster computers of their own.

In March, 2016 STEM-Trek Nonprofit conducted a survey to determine how many universities would be interested in adopting decommissioned HPC systems. Seventy favorable responses (including 35 US-EPSCoR, and 12 African schools) confirmed there is a demand for upcycled systems, and the investment necessary to refurbish them for light research and HPC workforce development purposes. Since the SADC HPC Forum has learned from their Ranger and UK system refurbishing experience, it would be useful to foster ongoing collaborations with US stakeholders who are likely to adopt decommissioned hardware in the future. Additionally, it's useful when newly-configured systems that attach to peered networks follow NSF XSEDE and OSG community cluster standards. In theory, that means more cycles for grid-enabled science around the world!

The merit of this SC16 exercise rests in the exchange of information and best practice. Open source software and operations procedures developed by programs developed under NSF funding will be exchanged with US and African champions. They will learn useful skills necessary to support local hardware, and engage with federated, cloud-based resources and services sponsored by XSEDE. US and pan-African participants will learn from each other, and their relationships will support additional research collaborations among their respective communities in the future.

All workshop materials will be shared with delegates and hosting organizations so they are prepared to teach other researchers, educators and systems administrators in the future. At the close of the workshop, a wrap-up meeting will be held to hear suggestions for future proposals. Notes from this meeting will be distributed, and may be suitable for a BoF or paper submission for future Supercomputing Conferences (or ACM Special Interest Group for HPC in Resource Constrained Regions). Affinity groups (established on the first Saturday workshop) will be encouraged to collaborate on posters and papers that will be submitted to XSEDE17, and/or SC17 and/or CHPC17.

This activity will strengthen the support toolkit used by US/Pan-African campus champions who build and manage cluster computing resources, and help scientists understand how they might use HPC to power discoveries that change the world.

The SC16 activity will precede the annual CHPC National Meeting held December 5-9, 2016 in East London, South Africa. STEM-Trek, XSEDE Campus Champions and the CHPC will co-market news about their collective activities beginning at the XSEDE conference in July, 2016. BoFs will be held at the CHPC meeting in December, and subsequent SC and XSEDE meetings.

## **Need for the activity**

Over the past 50 years, high performance computing (HPC) has supported tremendous advances in all areas of science. Yet, CI managers and campus champions from under-served regions of the U.S. and Africa often lack access to HPC training, which limits their ability to support others, position their centers for growth, contribute to scientific advances, or prepare their region's workforce. Consequently, the majority of research advances have been made by researchers at universities or national laboratories in urban-serving regions (with more robust training programs).

## **Benefits to US/Pan-African Delegates**

With increased exposure to resources and communities of interest that foster broad, international collaborations, this training exercise will pave the way for US and African discoveries in climate science, agriculture, water use, transportation, and telecommunications, among others – all of which would benefit African and US interests. By working with the African universities to make HPC more accessible to their communities, the effort will produce results in both computer-focused education, and science education.

In addition to 5-35 participants who will be supported to attend SC16 and the co-located training activity, XSEDE Campus Champions and members of the “Challenges of Managing Small HPC Centers” group could be invited to participate in some of the planned activities where they will meet the delegates (as funds allow). It is estimated that more than 200 would benefit directly, and thousands indirectly (through future training that is conducted by participants). Tier III center managers will be invited to attend one evening group activity (pending sponsorship).

## **Meeting Organization**

The workshop will be organized by an international facilitation team led by Elizabeth Leake (STEM-Trek Nonprofit), with trainers from the US and Africa. Planning for this activity was coordinated with Happy Sithole (Director of the South African Centre for High Performance Computing); SADC Secretariat Anneline Morgan; SADC HPC Forum Chair Tshiamo Motshegwa (U-Botswana); and XSEDE Campus Champion Coordinator Kay Hunt (Purdue).

## **Organizing Committee:**

Dana Brunson, Oklahoma State University (XSEDE, Campus Champion Liaison);  
Elizabeth Leake, STEM-Trek Nonprofit (Chair and Facilitator);  
Henry Neeman, University of Oklahoma (XSEDE, Training Coordinator);

## **With assistance from:**

Bryan Johnston, South African Centre for HPC (SADC Trainer, Cape Town);  
Nick Thorne, Texas Advanced Computing Center (U-Texas, formerly SADC/CHPC);  
Alana Romanella, Virginia Tech (XSEDE/STEM-Trek);  
And additional volunteer reviewers (*thank you!*).

## Eligibility and Application Process

Systems administrators, educators, researchers, cybersecurity professionals, tech project managers, and network engineers who also serve as campus champions are welcome to apply. As a campus champion (staff or graduate student), applicants must work in a support role helping students, faculty and staff leverage locally-hosted, or remotely-accessed advanced CI for education and research (at least 50 percent of the time).

Applicants must be at least 21 years of age. A CV or resume, photograph of your face and brief biography (200 words) should be included with a cover letter, adviser/supervisor recommendation letter, and copy of an identification card or passport. US applicants may provide a photocopy of their driver's license, or university ID. African applicants must provide a photocopy of their passport photo page. International travelers are responsible for securing their own visas (where applicable), and should have at least three blank pages in their passport.

We expect the process to be competitive. Therefore, the cover letter should describe why the applicant should be supported. *Include specific and unique reasons.*

The adviser/supervisor letter of recommendation should declare their support and describe how the university (or host institution) will benefit from the activity, that the release time from work is approved, and how the individual intends to share what they have learned with others. STEM-Trek beneficiaries are encouraged to pay-it-forward by helping those who are less fortunate. Each delegate will be paired with a pay-it-forward mentor that will help them follow through. For example, staff might conduct a tech meetup, workshop or seminar, or teachers could share lessons learned with their students when they return to the classroom. Then, their experience can be documented in the form of a blog post, or white paper.

### Applications must include:

1. Cover letter (one page).
2. Photo ID copy (or passport photo page).
3. Photograph/Head shot (that clearly shows the individual's face). *By providing your photo, you acknowledge that STEM-Trek and its collaborators may use it for publicity purposes on the public web site, and any communications collateral developed for the activity. Photos will be taken during the activity will also be used for promotional purposes unless you officially opt out (see number 7).*
4. Letter of recommendation (adviser/supervisor)
5. CV/Resume (2-pages maximum).
6. 200-word biography explaining your experience and interests (as it relates to HPC; in addition to what is already on your CV/Resume).
7. If you do NOT authorize use of your photo, name or bio for publicity purposes, you must still provide a full-facial photograph with your application, identification, and reasons why you wish to be exempt in your cover letter. For example, people who perform cybersecurity forensics or auditing 50 percent of the time might opt for anonymity's sake. In that case, disclose the conflict in your cover letter, and ask your adviser/supervisor to support the business case for anonymity in their letter.

The application, with .jpg and .doc (MSWord) attachments, must be received by [info@stem-trek.org](mailto:info@stem-trek.org) no later than Sept. 2, 2016. Notification letters will be sent to awardees by September 15, 2016.

