Ten years of the HPC Ecosystems Project

Transforming HPC in Africa for the past decade

Bryan Johnston
CSIR, Centre for High Performance Computing,
Cape Town, South Africa
bjohnston@csir.co.za

Lara Timm
CSIR, Centre for High Performance Computing,
Cape Town, South Africa
ltimm@csir.co.za

David Macleod
CSIR, Centre for High Performance Computing,
Cape Town, South Africa
dmacleod@csir.co.za

John Poole
Clemson University, Center for Human Genetics,
Greenwood SC, USA
jopoole@clemson.edu

ABSTRACT

The Council for Scientific and Industrial Research (CSIR)'s National Integrated Cyberinfrastructure System (NICIS) plays a pivotal role in advancing two key initiatives that focus on developing cyberinfrastructure across the African continent: the Southern African Development Community (SADC) Cyberinfrastructure Framework, and the Square Kilometre Array (SKA) Partner Countries Big Data initiative. Within NICIS these initiatives are managed through the HPC Ecosystems Project, which has two primary objectives: distributing entry level High Performance Computing (HPC) systems by repurposing decommissioned tier-1 HPC systems, and cultivating a skilled HPC workforce across Africa. The first deployment of HPC systems under the project occurred in 2013, using repurposed hardware from the Texas Advanced Computing Center's decommissioned Ranger HPC system. These systems were allocated to bolster research capabilities at local research institutes in South Africa and within partner countries of the SKA project across Africa. A decade later, at the close of 2023, the HPC Ecosystems Project has deployed 35 HPC systems in 11 countries and delivered more than 30 formal HPC training workshops to over 700 participants, surpassing 21000 total participation hours. There is an active and growing virtual community exceeding 230 HPC practitioners globally. This paper provides a high-level overview of the first ten years of the project's lifespan; outlining the various approaches towards establishing sustainable cyberinfrastructure and HPC workforces in Africa. Included is a reflection on the challenges experienced, lessons learned, and progress made towards delivering cyberinfrastructure resources and HPC training to resource-constrained environments.

CCS CONCEPTS

• **Social and professional topics** → Professional topics; Management of computing and information systems; Project and people management; Computer and information systems training; Professional topics; Management of computing and information systems; Project and people management; Systems planning; User characteristics; Geographic characteristics.

KEYWORDS

Resource Constrained Environments, HPC Ecosystems Project, Africa, Workforce Development

ACM Reference Format:
Bryan Johnston, Lara Timm, David Macleod, and John Poole. 2024. Ten years of the HPC Ecosystems Project: Transforming HPC in Africa for the past decade. In *Practice and Experience in Advanced Research Computing (PEARC '24), July 21–25, 2024, Providence, RI, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3626203.3670537

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
PEARC '24, July 21–25, 2024, Providence, RI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0419-2/24/07
https://doi.org/10.1145/3626203.3670537

1 INTRODUCTION

The HPC Ecosystems Project ("the project") is an initiative of the South African Department of Science and Innovation[1], managed by the Centre for High Performance Computing (CHPC)[2], that facilitates readiness in advanced research computing in Africa [8]. At its inception in 2013, the project was mandated to prepare community partners for the upcoming exascale computing requirements of the world's largest radio telescope, the Square Kilometre Array (SKA). The larger of the two SKA telescopes, SKA-MID, will be hosted in Africa with the bulk of the dishes located in South Africa's Karoo [10]. Over the past decade, the project has expanded its regional scope to address the gaps in resolving the cyberinfrastructure strategic needs of countries in the region. At the end of 2023, the project included partner institutes across Africa that encompass the SKA partner countries[3], the Southern African Development Community (SADC)[4], and affiliate members from other regions.

The project involves both the deployment of HPC resources and the advancement of the HPC workforce to operate these resources. The past ten years have imparted many lessons towards successfully facilitating the deployment and sustainable operation of HPC resources in countries that lack infrastructure and human capital resources.

[1]www.dst.gov.za
[2]NICIS CHPC - https://www.chpc.ac.za/
[3]African Square Kilometre Array partner countries https://www.gov.za/news/media-statements/african-square-kilometre-array-ska-partner-countries-03-sep-2010
[4]Member States | SADC - https://www.sadc.int/member-states

Table 1: The total number and breakdown of HPC site deployments per Phase of the HPC Ecosystems Project

	TACC Ranger	Cambridge M1000e	CHPC C6100	TACC Stampede	Change	Total
Phase 1 – 2013 to 2016	5 (+5)	N/A	N/A	N/A	+5	5
Phase 2 – 2016 to 2018	11 (+6) *	2 (+2) *	2 (+2)	N/A	+10	15
Phase 3 – 2018 to 2020	11*	2*	2 (+0)	5 (+5)	+5	20
Phase 4 – 2020 to 2023	11*	2*	7 (+5)*	15 (+10)	+15	35

* All available resources have been allocated from this pipeline

While this paper focuses on our African community, we believe our story holds value for other groups with similar characteristics outside of Africa. To universalise the relevance of our journey, we classify our community partners as being from resource-constrained environments (RCEs), which we determine to be lower-income regions, and from the SIGHPC Chapter on Resource Constrained Environments[5] classification referring to under-represented regions in the HPC community.

The HPC Ecosystems Project has undergone several major operational revisions during its first ten years. A range of factors triggered these revisions, some of which include the changing HPC hardware pipeline, the expanding project community, and through the feedback and lessons learned during project execution. This paper describes these major revisions in four phases. Through a chronological account of the project, we aim to share our experience of learning and adaptation, to allow readers to relate to their own journeys, and to hopefully contribute towards best practices in facilitating the development of HPC resources and personnel in RCEs.

1.1 The Origin Story of the HPC Ecosystems Project

The HPC Ecosystems Project was conceived when the demand for HPC resources in Africa was starting to accelerate alongside the diminishing economic viability of large centres to sustain years-old flagship HPC systems. The accelerating demand in Africa was prompted primarily by the growing scientific computing needs of the African research community and the anticipated requirements of the SKA [1, 10, 15]. At the time, efforts had been made by others to engage in HPC capacity building activities and deliver "seed" cluster hardware in Africa, but these were relatively modest in scope and, importantly, were not focused on the SKA partner countries [1, 2].

The gradual and consistent expansion of the HPC Ecosystems target audience is testament to the sizable demand for HPC computing resources and the need to develop an HPC workforce in Africa. Table 1 demonstrates this growth by summarising the HPC system site deployments facilitated through the HPC Ecosystems Project up to the close of 2023.

While the target community has grown significantly, the key objectives of the project to deliver HPC resources and train a sustainable HPC workforce have remained relevant, affirming that there is still a prevalent digital divide in Africa.

[5]SIGHPC-RCE (acm.org) - https://sighpc-rce.acm.org/

2 PHASE 1 – ESTABLISHING THE PROJECT: 2013 TO 2016

In 2013 the Ranger HPC system[6], one of the National Science Foundation's flagship production systems and the first open science petaflop system, was decommissioned. Instead of the prospect of a still-usable system like Ranger becoming disposable e-waste, the Texas Advanced Computer Center (TACC) decided to break Ranger down into several independent racks and distribute them to a small selection of recipients. Africa received 40 racks: 16 to Tanzania; 4 to Botswana; and 20 shipped to the CHPC (South Africa). The CHPC intended to distribute the systems to South African research institutes and SKA partner sites [12, 16, 19]. In so doing, the Ranger Project – the predecessor to the HPC Ecosystems Project – was born.

At this initial stage, there was no existing site selection process. While SKA partner countries were obvious recipients of this hardware, South African recipients were not predetermined. The initial response from local research institutes in South Africa to the call for interest to adopt end-of-life hardware was substantial.

An HPC training course based on an in-house HPC software stack[7] was developed to establish which institutes were willing to allocate staff hours to participate in the training, and by extension, to participate in the Ranger Project. This led to a short list of institutes with which the Ranger Project could work closely to identify the suitability of their respective sites to host one of the entry level clusters. The areas of concern were: (1) available room in a suitable data centre environment; (2) ability to provide adequate power and cooling; and finally, (3) physical access space for the fully populated racks to be moved into the data centre. Sites already operating HPC services were not considered for selection.

By the close of this phase, the Ranger project had distributed its first five systems to partner sites (see Table 1 The total number of HPC Ecosystems site deployments per phase) and once recipient sites received their hardware, the software deployment was supported by project staff during an on-site engagement.

2.1 Phase 1 Technology Overview

The Ranger Project pipeline comprised of a single system: the Sun Blade x6420. Each rack consists of 48 nodes, 768 AMD Opteron 8356 (2.3GHz) cores, 1.5TB RAM, and QDR InfiniBand [5, 18].

The initial software stack was an internally developed integrated collection of open-source tools commonly used to configure and

[6]Ranger - TOP500 - https://www.top500.org/system/176183/

[7]ACE Lab HPC Ecosystems Software Stack - https://tinyurl.com/ACE-HPC-Stack-RangerProject/

administer HPC systems. The stack comprised of shell scripts and automation tools, namely: Puppet for configuration management, Cobbler for provisioning, and TORQUE for distributed resource management.

2.2 Phase 1 An Introduction to HPC Technologies

The initial training was performed at the CHPC in cohorts. Over the course of a week attendees were provided a series of introductory lectures relating to parallel computing and HPC technologies, supported by hands-on lab tutorials to install and manage a sample of the in-house software stack. Upon completion, participants were invited to register their interest in hosting an HPC system at their local institute.

2.3 Phase 1 Lessons Learned and Responses

Despite the intention of the selection process to identify earnest and plausible interest from suitable candidates, it was evident that some applications were the result of the Bandwagon Effect. Availability of personnel for the training did not consistently lead to active utilization of the HPC resource after delivery. Given that more hardware was in the project pipeline, it was mooted that a targeted approach to HPC recipient selection may be more successful in future rollouts.

At this stage, simply shipping hardware to qualifying sites was considered a success. The prevailing assumptions were that site personnel who had completed the initial HPC training would be adequately equipped to manage the HPC resource and could sustain their own personal development in HPC upskilling. Unfortunately, these assumptions did not hold true.

Many participants lacked institutional support and had no community engagement for sustainable development. Their isolation was exacerbated by the custom-built software stack lacking mainstream support. In some cases, technical complexities were insurmountable for sites lacking deeper knowledge of Linux systems. Since many of the sites had little Linux experience, the lack of readily available resources and community help prevented them from resolving problems.

Furthermore, the combination of inadequate support hardware (storage and management servers to support the donated HPC systems) and inexperience in effective public relations, coupled with low urgency from local leadership at the recipient sites, resulted in some HPC systems being under-utilised, or even remaining dormant. In many cases there was a lack of broader awareness of the presence of the HPC resource by potential local users which compounded the problem. It became clear that sustainable development in HPC centres across Africa would require further investment in support infrastructure and a stronger focus on forming and supporting HPC communities [3, 4, 11, 20].

3 PHASE 2 – ESTABLISHING THE HPC ECOSYSTEMS COMMUNITY: 2016 TO 2018

Phase two marked the addition of new HPC system configurations to the Ranger Project pipeline. With this divergence from only Ranger equipment, there came a need to refresh the *Ranger Project* moniker. The title *HPC Ecosystems Project* was adopted to remove reference to any specific configuration of hardware in the pipeline, while reflecting the focus on developing African HPC ecosystems through the project's interventions.

Accompanying the new HPC Ecosystems identity came the goal to entrench a deeper engagement from regional members towards the growth of a single HPC community for the region. A virtual community was established using Google Groups, which was used as a central communications channel, support platform and community knowledgebase.

To improve the success of HPC deployments, prospective recipients were required to first complete a site readiness evaluation complemented by a site assessment by the project team before hardware would be allocated to the site. This also aimed to ensure a faster turnaround from delivery to operationalisation of the HPC system.

Phase two saw the diversification of hardware resources in the project pipeline. This change offered research and skills development opportunities to the project, contributing to the RCE HPC deployment knowledgebase. The new technologies afforded recipients the opportunity to adopt a unique contributory role to the HPC community in Africa. Diversifying the hardware also necessitated tailored modifications to be made to the CHPC software stack and training materials.

To ensure that sites were better equipped to operationalise the donated HPC systems, management and storage servers were procured for each new recipient site. This added the benefit of reducing the software stack modifications required to accommodate diverse HPC support resources.

Given the relatively insignificant number of personnel to train, from only a few sites with hardware, the training demands were considered manageable at this stage for the small Ecosystems project team. Similarly to phase 1, training was performed in-person at each recipient site and aimed to render the nominated HPC technical team capable of operating the donated HPC systems.

Separately, the SADC Cyberinfrastructure (SADC/CI) Framework had been introduced through a regional governmental forum to address the understanding that continued skills development in HPC infrastructure and skills within the SADC region could not happen in isolation [14]. Given the overlap between the SADC/CI and SKA communities' common HPC goals, the SADC/CI was incorporated into the HPC Ecosystems Project.

During this period, ten more HPC systems were distributed, bringing the total deployed to fifteen (see Table 1 The total number of HPC Ecosystems site deployments per phase).

3.1 Phase 2 Technology Overview

Two additional HPC resources, as well as a standardised management and storage server, were incorporated into the site deployment pipeline during the second phase. The first HPC system - Cambridge's M1000e system – was a 16-node cluster with 192 Westmere 2.66GHz CPU cores, 768GB RAM, DDR Infiniband (20Gbps) and a 10U form factor. The second - the CHPC's own decommissioned Westmere system - was the first pipeline system to support modular deployment, allowing the project team to tailor an HPC system to more appropriate capacities for recipient sites. Each of the CHPC's C6100 chassis was a 4-blade system with 48 Westmere 2.93GHz

Table 2: Software Stack Deployment Status at Partner Sites (late phase 2)

	Number	Number Operational	Ratio
CHPC Software Stack	9	3	33% (22% *)
Alternative in-house Software Stack	5	3	60%
Not Deployed	1	0	0%
Overall	15	6	40% (33% *)

* At least one site had significant Software Stack bugs

CPU cores, 144GB RAM, QDR Infiniband and 2U form factor. Dell PowerEdge T400 series servers were procured as management and storage servers, with 18TB of raw storage to provide basic resource functionality for recipient sites.

With some minor adjustments, the in-house software stack developed in phase 1 could be deployed on the new C6100 and M1000e hardware.

3.2 Phase 2 Lessons Learned & Responses

Until this point the nomination of trainees was entrusted to leadership within recipient institutes, with Ecosystems staff sharing guidelines of appropriate prerequisites. The training provided to each site on their specific hardware was expected to swiftly produce HPC competent personnel able to administer their HPC systems. This, however, was found not to be the case. The challenges to rapidly produce personnel equipped with basic HPC skills could largely be attributed to the prevailing lack of understanding of the fundamental technical skills required to operate an HPC system. The diverse range of nominated technical HPC personnel – some with extensive scientific computing background (the ideal scenario) to others with no awareness of Linux or system administration experience (the most challenging scenario) – impacted the effectiveness of the training as well as the sustainable success of a site deployment. It was evident that individual sites were often not inclined to expose their knowledge gaps in the public forum (i.e. the Google Group), preventing them from obtaining community support to make progress with their HPC deployments and skills development.

It was found that using standardised HPC management servers greatly improved inter-site collaboration and community support, as well as reduced the administrative and technical complexity of preparing solutions and training material for the Ecosystems community. Additionally, the modularity of the CHPC's C6100 system provided more flexibility to distribute an appropriate node configuration to recipient sites, based on their infrastructure capacity and user demand.

As more sites operationalised their systems, the support and maintenance burden of the software stack grew significantly. To encourage self-sustainability and learning through self-guided discovery, sites were permitted (and encouraged) to modify and improve the CHPC software stack for their deployments, and to feed improvements back to the community. This approach quickly led to two prominent problems: (1) software stack configuration drift leading to disparate deployments, limiting the benefit of community support; and (2) sites discovering more design problems with the base software stack, which had originally been obscured by developer blindspots and the original design's specificity.

These effects can be seen in Table 2, which shows the low success rate of software stack implementations across the fifteen sites that had received hardware. Notably, sites that implemented their own in-house software stack showed a higher rate of success over the sites that implemented and received training on the CHPC Software Stack.

The lack of network driver support for the Ranger hardware compounded the software stack problems, prompting the University of Namibia (UNAM) to evaluate a relatively new community software stack that was in development at the time – OpenHPC [17]. Their initial exploration found that the OpenHPC stack provided native support for the Ranger network hardware and that there was already an active and established support community. If utilised, this would allow Ecosystems project community members to independently receive support for their HPC software stack.

The initial site-readiness evaluation was focused solely on technical matters with no concern for pertinent non-technical components of a successful HPC deployment. Subsequently, many sites that were confirmed to be technically able to operate the systems were frustrated by non-technical concerns, including: a lack of institutional support to provide further training opportunities for technical staff; restrictions on open access policies for interdepartmental users; general lack of awareness of the HPC resource; a dwindling or short-term operating budget. Appropriate changes were made to the site-readiness evaluation form to reflect these non-technical components.

Finally, international imports/exports of HPC hardware to partner countries were incurring significant delays. Personnel at recipient sites could wait months, even years, for the hardware to be in place and to be powered-on. Once the hardware landed at sites, it would remain unused until personnel could conduct their technical training. The project team felt that a more agile approach to HPC personnel readiness was necessary. Readiness could be improved if technical teams were prepared well before the sites received the hardware.

4 PHASE 3 – FOCUS ON GROWING THE COMMUNITY: 2018 TO 2020

Motivated by the investigations done by UNAM in 2017, the HPC Ecosystems community software stack was officially migrated from the in-house developed solution towards the community supported OpenHPC stack [18]. The philosophy of OpenHPC, which is to provide an HPC system software stack that is customisable and hardware agnostic, was a good fit for the Ecosystems project.

At this stage there were more institutes requesting hardware than the pipeline could support. To address equitable distribution of the remaining units in the HPC pipeline, a competitive bid process was undertaken where candidates submitted a comprehensive proposal for hardware, including commitment to a set of terms and conditions for the use of the hardware. To support the bid process, an updated site readiness evaluation form was introduced which included capturing information for key site stakeholders (such as personnel responsible for networking, cybersecurity, infrastructure, and facilities management) and more detailed data centre specifications and operational budgets. The readiness levels of potential technical personnel were evaluated in the revised forms.

Moving towards strengthening the community engagement, a community-wide in-person event was hosted at one of the South African partner institutes where members of all partner sites were invited to participate in a community conference.

To address the need for a more effective, direct, and synchronous communication platform for the community, a Slack workspace for HPC Ecosystems was established. The workspace complemented the Google Groups service, which changed into a general mailing list for alerts and announcements.

Five Stampede systems were distributed during phase 3, with twenty systems deployed in total to date (see Table 1 The total number of HPC Ecosystems site deployments per phase).

4.1 Phase 3 Technology Overview

TACC's Stampede HPC system was incorporated into the existing pipeline for distribution. Since Stampede was based on Dell's C8000 series chassis, it was the second modular HPC system available through the project. Each 4U chassis consisted of 64 Sandy Bridge 2.70GHz cores, 128G RAM, FDR InfiniBand, and Intel Xeon Phi Knights Corner accelerators.

The adoption of OpenHPC moved the HPC software stack development away from a small development team operating in isolation to a community-built tool with support resources provided by the global community.

4.2 Phase 3 Community Supported Training

During this phase, the training focus shifted towards deploying an OpenHPC-managed cluster, with much of the foundational training relating to parallel computing concepts being removed from the curricula entirely. The Ecosystems training material presented in the on-site training workshops was adapted from the official OpenHPC installation recipe[8]. The intention was to produce a curated version that addressed the steps necessary to deploy the HPC Ecosystems hardware resources. By removing parallel computing concepts from the training, the training model began shifting towards a blackbox HPC deployment model.

4.3 Phase 3 Lessons Learned & Responses

Training personnel to deploy an OpenHPC system proved successful in exposing administrators to the underlying backend of the HPC system and how to deploy it, but was not sufficient to prepare system administrators to manage an HPC system through daily use. The understanding gained through system deployment using the OpenHPC stack was, however, an improvement from the previous software configuration. It became clear that training personnel only upon hardware delivery led to delays in operationalising the HPC resource since it was dependant on technical personnel first learning how to deploy an HPC system. Given the lack of priority from site management, several sites had no urgency to complete the deployment of their system which proved to be a significant bottleneck after hardware was delivered.

As the training models were refined and improved during phase 3, the project became a victim of its own success. Upskilling technical personnel with scarce system administration and HPC skills meant they were more attractive hires to other companies. Some sites could not retain their HPC staff after they completed the training, and the sites reverted to their initial problem of not enough expertise to manage the HPC systems, leading some to require CHPC personnel to return to provide the training again. Consequently, the HPC Ecosystems team began developing a more sustainable training solution to offset the increased demands for in-person training.

As the project community continued to grow, it was apparent that sites who were not currently hardware recipients would invariably become eligible. The scope of training was thus expanded to include anyone in the region irrespective of whether they had an HPC system allocation or not. To provide a scalable and accessible training solution, it was determined that there would need to be an online and on-demand training platform for HPC practitioners.

5 PHASE 4 – ENABLING THE COMMUNITY: 2020 TO 2023

With the onset of COVID-19 and a lockdown on international travel in 2020, the transition to a community-focused approach for training and workforce development was a timely one for the HPC Ecosystems Project. Work commenced towards a scalable and accessible online training platform during phase 3. Phase 4 saw the team release an updated virtual training platform that enabled continuous HPC training, even when international travel was impossible [9].

Given the mitigating factors of COVID lockdown, an increase in the number of supported sites by 75% over phase 3, and the availability of a viable training solution; the project directed its focus towards achieving more with less. The CHPC project team remained small, with staff time divided on other projects. It became imperative to tap into the community virtually to compensate for the lack of physical engagements. The community resonated with the renewed vision of the project under the circumstances, and collectively embraced the virtual collaboration to ensure continued progress.

The bulk of deployments in phase 3 were determined competitively, so many of the earliest selected sites were more established institutes with appropriate resources to support the HPC donations. Accordingly, in phase 4 several of the remaining prospective sites were smaller institutes with fewer resources and less user demand. Taking advantage of the modular Stampede HPC systems, smaller

[8] Install_guide-CentOS7.1-1.0.pdf - http://openhpc.community/wp-content/uploads/Install_guide-CentOS7.1-1.0.pdf

sites were supplied with smaller HPC systems - typically half racks of Stampede (20 nodes) with a PowerEdge T400-series management and storage support server.

To accelerate system operationalisation and the onboarding of users to the HPC systems as early as possible, management servers were pre-configured (blackboxed) with the OpenHPC software stack. The anticipation was that the HPC systems could simply be powered on upon arrival and users could immediately make use of the resources. This blackbox philosophy marked a shift from enabling technical staff to drive the HPC adoption, to a user-centric approach where the users would drive the HPC adoption. From a technical training perspective, the move towards a 'blackbox' approach was supported by the questionable value in training technical personnel to deploy an HPC system that may never need to be re-installed, in contrast to focusing on training of personnel to operate and maintain a running HPC system.

All recipient sites were expected to have their technical teams complete an OpenHPC virtual workshop (available online in quarterly virtual classes, or available on-demand with community support available in the Ecosystems Slack) to ensure that the staff are familiar with the basic OpenHPC deployment principles. Additionally, sites were required to identify a pilot user to work alongside the allocated HPC staff. The pilot user must have some HPC user, and other technical experience, and must be available to work closely with the HPC technical personnel to operationalise the system.

Fifteen more HPCs were distributed during this period, bringing the total deployed to thirty-five systems (see Table 1 The total number of HPC Ecosystems site deployments per phase).

5.1 Phase 4 Technology Overview

The deployment of modular systems (CHPC's C6100 and TACC's Stampede) allowed flexibility in customising a system appropriate for a site's particular capabilities and user demand. All sites were still free to deploy a software stack of their choosing, but the HPC Ecosystems team would only support a specific flavour of OpenHPC (RockyLinux, Slurm, with stateless Warewulf).

5.2 Phase 4 Online and On-Demand Training

The dual factors of the COVID19 pandemic and the ever-increasing demand for training necessitated that teaching be transformed into an on-demand and accessible product that could be delivered at scale. A standalone virtual lab was made available with supporting training videos released on YouTube[9].

The first fully online training course was delivered in October 2020. Since then, the OpenHPC virtual lab has delivered training to 385 participants over 8 virtual workshops, with additional participants using the training lab independently outside of the curated quarterly workshops. An increase in community involvement accompanied the launch of the virtual labs, where community members frequently volunteer their time to assist with the training workshops and contribute to the ongoing maintenance and testing of the virtual lab.

During phase 4 the project's reach grew significantly, with the virtual training extending beyond HPC Ecosystems sites to more than 25 countries spanning five continents. The ability to reach such scale is enabled in no small part by leveraging the HPC Ecosystems community to support the training and continuous development. Primary support for sites is complemented by community members through Slack. With an ever-increasing number of sites, this has proven to be a more sustainable approach to facilitate advanced research computing readiness in the project's target regions, and beyond.

5.3 Phase 4 Lessons Learned & Responses

Although the virtual training workshops attract many attendees, a digital experience cannot replace an in-person engagement. Site visits are still considered to be the most effective way to guarantee dedicated time from the site's technical team to learn about their systems and to connect the HPC providers with their research community. Site visits reinforce the relevance and importance of the work that is being performed on the HPC systems. As is common with online learning environments, it is difficult to gauge the effectiveness of the online training workshops, particularly from afar [13]. A mechanism to verify course completion is in its early release stages – it is hoped that this information can be used to gauge the efficacy of the virtual training lab.

By providing online pre-training that is performed in small virtual communities, hardware recipients can gain confidence in Linux environments and a greater understanding of the software stack that they will be administrating. If a practical and sustainable approach can be found to still conduct on-site training, then a hybrid training model may extend the impact of the virtual HPC training with hands-on experiential learning [6].

We have found greater success where sites form an HPC steering committee, consisting of representatives from institutional leadership, HPC technical personnel, and a research pilot user. Once the system is considered ready by the steering committee, more users are incrementally onboarded, and after further validation and checks, the system is opened to the wider community.

6 REFLECTIONS / WHERE WE ARE

The task of building and sustaining an African HPC community is not easy. The HPC Ecosystems team is under no illusion that showcasing many deployed systems interrogates the quantity or quality of computational science being performed on any of these systems. While the metrics of the project's performance have focused on the size of the community and the number of HPC institutes, we are aware that there are many more meaningful measures to reflect upon, and that we must also resist the temptation to turn any measures into targets.

This paper is intended as a basis for a wider high-level overview of the project rather than a narrow deep dive into specific challenges and lessons learned over the past ten years. The myriad of lessons learned, and challenges faced, are too varied and diverse to be accounted for in a single paper. Likewise, many aspects of the project's journey have not been discussed in this overview, such as the support for community members' participation and representation at SC Supercomputing conferences since 2015, or the

[9]OpenHPC101 YouTube Training Series: https://www.youtube.com/playlist?list=PL2s6Yr_Iu_ke16_di1C3dowXHF-hRLmC9

STEM-Trek funded international workshops to sustain the face-to-face community engagements, or the University of Witwatersrand's modifications of their donated hardware for bespoke use cases.

Recognition of the community's achievements has been acknowledged through two HPCwire awards (2018 – Readers' choice - Workforce Diversity Leadership, and 2022 – Editors' Choice - Workforce Diversity & Inclusion Leadership) as well as several African Human Capital Development awards.

There are two recurring mantras within the HPC Ecosystems Project. The first is "if it's free, someone will take it" - unfortunately in some cases, this has shown itself to be true for all the wrong reasons. Even when a site is unable to support or use an HPC system, because it has been offered for free with training, there is a belief that accepting the resource and undergoing the training on offer will produce a successful HPC centre. It is evident from the first ten years that an appetite for a free HPC system, delivered with training, is not a recipe for success.

The second mantra is "HPC systems live fast but retire young" – we are aware that there will be a continuous pipeline for decommissioned HPC resources. One of the impending challenges is to determine how to effectively manage this growing pipeline in a sustainable manner to avoid the risk of, in effect, shifting e-waste disposal to Africa.

Finally, there is an urgent need to implement more effective auditing, reporting, and accountability mechanisms for sites, but this remains a difficult challenge, particularly since there is no reasonable recourse. The project attempts to guard against these risks through the implementation of diligent readiness evaluations to pre-empt any problems but it remains an ongoing challenge.

7 CONCLUSION

Given the relative novelty of HPC computing resources to RCEs in Africa, despite the provision of an operational HPC system, there is still a gulf in digital skills for computational scientists to use the systems. Incorporating user-friendly tools such as Open OnDemand to facilitate wider access to the HPC system is gaining attention within the project [7].

The response to the virtual HPC training lab has been overwhelmingly positive, and feedback has shown that there is more required to successfully deploy and operationalise HPC infrastructure than simply providing hardware and initial technical training. There is a notable demand to equip researchers to use an HPC resource and to address this, the HPC Ecosystems community is pursuing the enhancement of the virtual training lab by developing a downloadable, composable, flexible virtual cluster environment to enable future scalable on-demand training for HPC users.

The past ten years has shown us that there are three factors to consider towards the successful deployment and sustainable operation of an HPC system in an RCE, all of which are addressed by the establishment of an HPC steering committee:

1. Technical Personnel: staff require computational science awareness to appreciate the unique role that HPC systems play and to understand that it is simply not a case of more server resources for an enterprise IT datacentre. To ensure a sustainable workforce pipeline, the technical team must consist of two, ideally more, technical personnel.
2. Users: if there is no user demand, the system is a white elephant. Sites must ensure that there are users with an interest in the system and that this interest is generated through awareness of the system.
3. Management / Leadership support: personnel administering the HPC system require time and financial support to attend training, meet the wider community, and should be recognized for their niche skillset and how it is vital to facilitate the HPC resource. Staff with scarce in-demand skills are at risk of being poached by other organisations and must be incentivised to stay.

As with any community, the HPC Ecosystems community's strength lies with its members. While it may currently be too early to evaluate the impact of the new HPC centres in terms of scientific computing in the region, the regular contact with the community through training events and the project's digital platforms has allowed ample opportunity to engage. Through such engagements it is clear that there are numerous enthusiastic, talented, and knowledgeable HPC practitioners rising within the Ecosystems community. As stars continue to shine through, the community is in good hands and the future for HPC in Africa is bright.

ACKNOWLEDGMENTS

The HPC Ecosystems Project funders: South Africa's Department of Science and Innovation (DSI) and SADC/CI provided funding to support the implementation plan of SADC CI framework; DARA: Development in Africa with Radio Astronomy for SKA partner countries; SADC Member States. Furthermore, a project dedicated to the development of a global community cannot achieve this without the community – many people and organisations have played a pivotal role in the inception and ongoing development of the HPC Ecosystem Project's initiatives, including: Texas Advanced Computing Center, Cambridge University, the HPC Ecosystems Project team (past and present), Elizabeth Leake (STEM-Trek), Airlink, SC conference committees, all the HPC Ecosystems community trainers and teaching assistants and content contributors, the HPC Ecosystems Community at large, Eugene de Beste for continuing to support the project from afar, Anton Limbo for finding OpenHPC, Nick Thorne for starting our journey in 2013, and countless others for championing our community, supporting our vision, and investing their time.

REFERENCES

[1] Abiona, O., Onime, C., Cozzini, S. and Hailemariam, S. Capacity building for HPC Infrastructure setup in Africa: The ICTP experience.
[2] Amolo, G.O. 2018. The Growth of High-Performance Computing in Africa. *Computing in Science and Engineering*. 20, 3 (May 2018), 21–24. DOI:https://doi.org/10.1109/MCSE.2018.03221926.
[3] Apostolos, K. and Alivisos, S. 2010. Internet-mediated communities of practice (IMCoPs): A meta-analysis of critical elements. *Proceedings - 2nd International Conference on Intelligent Networking and Collaborative Systems, INCOS 2010*. (2010), 1–7. DOI:https://doi.org/10.1109/INCOS.2010.23.
[4] Cahill, K.J., Lathrop, S. and Gordon, S. 2017. ScienceDirect Building a Community of Practice to Prepare the HPC Workforce. 2017, (2017), 12–14. DOI:https://doi.org/10.1016/j.procs.2017.05.059.
[5] Erich, S., Jack, D., Horst, S. and Martin, M. 2008. *Top 500 list - Ranger – SunBlade x6420, Opteron Quad2Ghz,Infiniband*.
[6] Gordon, S.I., Demmel, J., Destefano, L. and Rivera, L. 2015. Extending access to HPC skills through a blended online course. *ACM International Conference Proceeding Series*. 2015-July, (Jul. 2015). DOI:https://doi.org/10.1145/2792745.2792760.

[7] Hudak, R., Bitterman, D.E., Carey, T., Johnson, P., Franz, D., Brady, E. and Diwan, S. 2018. Open OnDemand: A web-based client portal for HPC centers. 43, (2018), 1–43. DOI:https://doi.org/10.21105/joss.00622.

[8] Johnston, B. 2019. HPC Ecosystems Project: Facilitating advanced research computing in Africa. *ACM International Conference Proceeding Series*. (Jul. 2019). DOI:https://doi.org/10.1145/3332186.3333264.

[9] Johnston, B., Timm, L., de Beste, E. and Hashatsi, M. 2024. Let's Get Our Heads Out of the Clouds: A scalable and sustainable approach to HPC Training Labs for Resource Constrained Environments and anyone else stuck in the Clouds. *The Journal of Computational Science Education*. 15, 1 (Mar. 2024), 64–71. DOI:https://doi.org/10.22369/issn.2153-4136/15/1/13.

[10] Jongerius, R. 2016. *Exascale Computer System Design: The Square Kilometre Array*. Technische Universiteit Eindhoven.

[11] Katernyak, I., Loboda, V. and Kulya, M. 2018. eLearning within the Community of Practice for sustainable development. *Higher Education, Skills and Work-based Learning*. 8, 3 (Aug. 2018), 312–322. DOI:https://doi.org/10.1108/HESWBL-03-2018-0030/FULL/PDF.

[12] Lang, C. 2015. Cultivating HPC in South Africa. *HPCwire*.

[13] Michinov, N., Brunot, S., Le Bohec, O., Juhel, J. and Delaval, M. 2011. Procrastination, participation, and performance in online learning environments. (2011). DOI:https://doi.org/10.1016/j.compedu.2010.07.025.

[14] Motshegwa, T., Wright, C., Sithole, H., Ngolwe, C. and Morgan, A. 2018. Developing a Cyber-infrastructure for Enhancing Regional Collaboration on Education, Research, Science, Technology and Innovation. (2018).

[15] Nkundwe Mwasaga, M., Apiola, M., Suhonen, J. and Joy, M. Integrating High Performance Computing into a Tanzanian IT Engineering Curriculum.

[16] Ranger Supercomputer Begins New Life, Makes Global Journey to Africa - Research & Development World: https://www.rdworldonline.com/ranger-supercomputer-begins-new-life-makes-global-journey-to-africa/. Accessed: 2024-02-28.

[17] Schulz, K.W., Baird, C.R., Brayford, D., Georgiou, Y., Gregory, A., Kurtzer, M., Pittsburgh, D.S., Center, S., Sterling, T., Sundararajan, N., Eric, D. and Hensbergen, V. 2016. Cluster Computing with OpenHPC. (Nov. 2016). DOI:https://doi.org/10.1145/1235.

[18] Shapopi, J., Limbo, A. and Backes, M. 2023. Namibia's first high performance computer. *South African Computer Journal*. 35, 1 (2023), 4–18. DOI:https://doi.org/10.18489/SACJ.V35I1.1189.

[19] Trader, T. 2014. NSF Petaflopper Skips Retirement for New Life in Africa. *HPCwire*.

[20] Wenger, E. 2008. Communities of Practice: Learning as a Social System. (2008).