# Adapting for Adoption: Tailoring HPC Training for Resource-Constrained Communities

Bryan Johnston
University of Cape Town
Inter-University Institute for Data Intensive Astronomy
bryan@idia.ac.za

James Carson
TACC, USA
jcarson@tacc.utexas.edu

Charlie Dey
TACC, USA
charlie@tacc.utexas.edu

Jennifer M. Schopf
TACC, USA
jms@tacc.utexas.edu

Christopher J Whalen
RDCT, USA
cwhalen@researchdata.us

Grace Kebirungi
IDI, Uganda
gkebirungi@idi.co.ug

## ABSTRACT

Successful High-Performance Computing (HPC) adoption in resource constrained environments requires more than hardware deployment; it requires training models that are explicitly designed for local realities rather than transplanted from well-resourced contexts. This position paper argues that leadership-class HPC curricula must be re-engineered into pragmatic, context-aware training pathways if emerging HPC communities - particularly in Africa - are to achieve sustainable operational capability. Drawing on the adaptation of the Texas Advanced Computing Center's (TACC)[1] leadership-class training into an onboarding program for African HPC sites, we illustrate how global best practices can be reframed into actionable modules aligned with infrastructural, staffing, and operational constraints. We contend that such contextualisation is not optional but foundational for long-term HPC sustainability, and we outline principles for designing training that empowers institutions to deploy and operate HPC systems sustainably and effectively within resource-limited environments.

## KEYWORDS

HPC training, Resource-Constrained Environments, Leadership-Class Expertise, African HPC, Workforce Development

## 1 INTRODUCTION

High-Performance Computing (HPC) adoption in emerging or resource-constrained environments (RCEs) cannot rely on hardware redistribution alone. Sustainable operation requires training models that translate leadership-class best practices into locally applicable skills. Yet, most global HPC training materials are designed for mature, well-resourced centres and assume the presence of stable infrastructure, adequate staffing, and access to established communities of practice. When transplanted directly into RCEs, these models fail to address the operational realities that shape sustainability and effective implementation of advanced research computing.

The Texas Advanced Computing Center (TACC) has long extended the life and impact of flagship HPC systems through its Legacy Computing Program[2], which loans retired hardware to broaden access to advanced computing. The HPC Ecosystems Project in Africa complements this hardware pipeline by focusing on repurposing the hardware, training the workforce, capacity building, and long-term operational support for emerging HPC sites [5]. Together, these initiatives highlight a central challenge: access to hardware does not automatically translate into operational capability.

This paper argues that leadership-class HPC training must be re-engineered for RCEs through pragmatic, context-aware onboarding frameworks. Preparing African recipients of Stampede2[3] hardware required a substantial redesign of TACC's existing curriculum, originally developed for experienced administrators in mature US-centric environments. The adapted program addressed operational challenges such as limited staffing, intermittent connectivity, and isolated technical communities, while also providing practical guidance for deploying and operating the Xeon Phi Knights Landing (KNL)[4] architecture within modern software stacks [3].

### Position Statement

We contend that contextualised training is foundational - not optional - for meaningful HPC adoption in RCEs. By reframing leadership-class concepts into practical modules aligned with local constraints, we outline principles for designing training that supports both immediate onboarding and long-term institutional growth.

### Defining Leadership-Class Training

**Leadership-class**, in the context of this position paper, refers to training, operational practices, and architectural expertise developed by an experienced large-scale, production-grade HPC centre in a mature computing environment. This denotes a depth of expertise and operational maturity. Leadership-class training, therefore, encompasses advanced system design principles, policy frameworks, and operational strategies used by established HPC centres with long-standing experience, stable infrastructure, and well-developed user communities.

[1] https://tacc.utexas.edu/

[2] TACC - https://tacc.utexas.edu/partnerships/legacy-computing-program/
[3] Stampede2 (Top500.org) – https://www.top500.org/system/179045/
[4] Intel Product Finder - https://www.intel.com/content/www/us/en/ark.html

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, along with the DOI information below. To copy otherwise, republish, post on servers or redistribute to lists requires prior specific permission from the first author.

© 2026 Journal of Computational Science Education
https://doi.org/10.22369/issn.2153-4136/17/2/2

## 2 AN OVERVIEW OF TRAINING RESOURCES GAPS

Although extensive documentation and training materials exist globally for HPC system deployment[5], they overwhelmingly reflect the assumptions of established centres in the Global North. These resources presume reliable power supply, deep layers of staffing, mature user communities, and ready access to peer support - conditions that rarely exist in resource-constrained environments. As a result, the materials do not address the realities faced by emerging HPC sites, where both technical staff and users may be encountering HPC for the first time.

The OpenHPC Virtual Lab is an example of hands-on training that accounts for resource constraints [4]. It provides scalable, accessible system administrator training without requiring advanced infrastructure. However, its focus remains primarily technical and does not address higher-level concerns such as governance, policy design, or sustainability planning. Most global training materials share this limitation: they emphasise operations but overlook the broader institutional ecosystem required for long-term success.

TACC's leadership-class curriculum illustrates this gap. It was developed for administrators already embedded in established HPC environments, whereas implementers of legacy hardware in RCEs often operate in isolation, with limited staffing, limited expertise, and no established user community. Effective training for these sites must therefore extend beyond technical instruction to include leadership principles, policy frameworks, and strategies for cultivating a sustainable HPC culture.

To address these gaps, we combined TACC's leadership-class training expertise with more than a decade of on-the-ground experience from the HPC Ecosystems Project in deploying HPC systems to RCEs in Africa. This integration enabled the development of training materials that retain the strategic depth of leadership-class instruction while grounding it in the operational realities of emerging African HPC sites. The resulting framework, described in the next section, re-engineers leadership-class concepts into pragmatic, locally tuned onboarding for RCEs.

## 3 TRAINING DESIGN

Early lessons from the HPC Ecosystems Project revealed a structural challenge: distributing decommissioned hardware does not automatically create operational capability. This is particularly true for architectures such as Intel's deprecated Xeon Phi Knights Landing (KNL), whose memory modes, NUMA behaviour, and performance characteristics require nuanced operational knowledge that is no longer widely documented. Without targeted onboarding, sites receiving such systems face a steep learning curve. This can undermine the intended impact of hardware distribution.

TACC's long history of delivering leadership-class training - through programs such as the Summer Institute[6] and advanced operations workshops - provides a strong foundation of strategic and architectural expertise [1, 2]. However, these curricula assume conditions typical of established HPC ecosystems. These assumptions do not hold in most RCEs, where administrators may be operating an HPC system for the first time, and often in isolation.

To bridge this gap, we adapted TACC's leadership-class curriculum into a practical onboarding program tailored specifically for RCEs. The goal was not to simplify the material, but to translate high-level system design principles into usable guidance aligned with local constraints. This required re-engineering the training around three core principles:

- **Contextualisation over transplantation** — leadership-class concepts must be reframed for environments with intermittent power, constrained staffing, and limited access to vendor or community support.
- **Operational pragmatism over architectural idealism** — training must prioritise reliability, reproducibility, and sustainability within resource-constrained environments rather than cutting-edge optimisation.
- **Capacity building over instruction** — administrators must be empowered not only to operate the system, but to teach users, enforce policy, and cultivate a local HPC culture.

The resulting program retained the strategic depth of the original TACC 5-day workshop ("Designing and Administering Large-scale Systems")[7] while restructuring it around the needs of RCEs:

- **Original Curriculum:** An advanced 5-day workshop for leadership-class HPC operations.
- **Adaptation Goal:** Translate high-level design principles into practical, context-aware guidance.
- **Audience:** System administrators with Linux and networking experience but limited exposure to operational HPC environments.
- **Outcome:** Participants gained the ability to design, deploy, and operate Stampede2 KNL systems with production-grade reliability while embedding strategic operational policies to ensure the integration of sustainability and user-support principles into daily operations.

## 4 PRINCIPLES FOR CONTEXT-AWARE HPC TRAINING

The adaptations described above reveal a set of recurring design considerations that extend beyond this specific training program. These patterns form a generalisable framework for developing HPC training that is operationally realistic, contextually-aware, and aligned with the constraints of RCEs. The following principles distill the lessons learned from adapting leadership-class training for RCEs and provides a foundation for designing future training models that support sustainable adoption:

- Design for operational reality, not idealised assumptions.
- Embed policy into onboarding to reduce support burden.
- Prioritise reproducibility and documentation over optimisation.
- Translate architectural concepts into actionable heuristics.
- Build local capacity by training administrators to train users.
- Plan for sustainability from the outset, even when growth is uncertain.

[5] Example 1: sighpceducation.hosting.acm.org/resources/hpcresources/;
Example 2: sighpceducation.hosting.acm.org/resources/systems/

[6] Summer Institute - https://www.rdworldonline.com/a-new-paradigm-for-training-in-advanced-computing/

[7] Archived advert from 2020 - https://web.archive.org/web/20200925225118/https:/www.tacc.utexas.edu/education/institutes/large-scale-systems

**Table 1: Training Program Overview**

| Day | Content |
|---|---|
| Day 1 | **Foundations of scientific system architecture:** capability vs. capacity computing, workload bottlenecks, and understanding how policy functions as an architectural mechanism. |
| Day 2 | **User education, policy, and HPC user experience:** onboarding as policy enforcement, effective help-request practices, and reducing support load through structured user education. |
| Day 3 | **Slurm, KNL, and hardware-aware scheduling:** configuring NUMA, MCDRAM, and scheduler policies for real-world KNL workloads. |
| Day 4 | **Designing cluster policy:** fairshare, quotas, and policy-encoded scheduling rules to ensure equitable access and reduce long-term administrative overhead. |
| Days 5–7 | **Containers, data lifecycle, security, and sustainability:** reproducible science, data-lifecycle policies, security practices, and sustainable workflows to reduce administrator burnout. |
| Day 8 | **Capstone hackathon:** integrating architecture, policy, storage, and monitoring into a demonstrable HPC environment, reinforcing capacity building through hands-on synthesis. |

## 5 AGENDA OVERVIEW (ADAPTED FOR RCES)

The adapted curriculum preserved the strategic depth of the original leadership-class workshop while restructuring each day around the operational realities of RCEs. The agenda emphasised actionable skills, policy-driven operations, and reproducible practices. The agenda reflected in Table 1 operationalises the principles of contextualisation, pragmatism, and capacity building: demonstrating how leadership-class concepts can be reframed for emerging HPC communities.

## 6 KEY ADAPTATIONS

Adapting a leadership-class curriculum for resource-constrained environments required more than simplifying technical content; it demanded a fundamental reframing of what effective HPC training looks like outside mature ecosystems. The adapted program shifted emphasis from architectural design toward pragmatic operational strategies that reflect the realities of African HPC sites. User experience was treated as a system-design problem, embedding policy directly into onboarding so that administrators could reduce support load and enforce good practices from the outset. Advanced hardware topics (such as KNL memory modes, NUMA behaviour, and MCDRAM configuration) were distilled into actionable scheduling heuristics that complemented the KNL architectural theory. The training prioritised community documentation and reproducibility over cutting-edge optimisation, recognising that sustainability in RCEs depends on clarity, consistency, and shared operational knowledge. Finally, capacity building was embedded as a core objective: administrators were trained not only to operate the system but to teach users, cultivate local expertise, and seed a self-sustaining HPC culture.

### Summary of Key Adaptations

- Shifted emphasis from leadership-class design to practicalities for resource-constrained environments.
- Reframed user experience as a system design problem, embedding policy into onboarding.
- Simplified advanced hardware topics (KNL, NUMA, MCDRAM) into practical scheduling practices.
- Prioritised community documentation and reproducibility over optimisation.
- Embedded capacity-building goals: empowering admins to teach users, not just operate systems.

These adaptations reflect the principles outlined in Section 3 and demonstrate how leadership-class training can be effectively re-engineered to meet the realities of RCEs.

## 7 OUTCOMES

The adapted training model produced outcomes that reinforce the central argument of this position paper: context-aware training is essential for sustainable HPC adoption in resource-constrained environments. Participants reported increased confidence in deploying and managing Stampede2 systems, particularly in navigating KNL-specific operational challenges that are poorly documented in global training materials nowadays. By formalising troubleshooting paradigms and reframing user onboarding as a policy mechanism, the program reduced friction between administrators and users - a critically important factor in environments where staffing is limited and support burdens can quickly become overwhelming. Perhaps most significantly, the training seeded a community of practice rooted in the long-standing experience of the HPC Ecosystems Project, enabling sites to share operational knowledge, documentation, and lived experience despite limited vendor support. These preliminary outcomes suggest that leadership-class training, when re-engineered for local realities, has the potential to empower emerging HPC centres to operate legacy systems sustainably and contribute meaningfully to regional HPC growth.

### Summary of Outcomes

- Participants gained confidence in deploying and managing repurposed Stampede2 and GPU systems.
- Training reduced friction between users and administrators through structured troubleshooting paradigms.
- The adapted curriculum seeded a community of practice inspired by the HPC Ecosystems Project, ensuring sustainability despite limited vendor documentation.

These outcomes suggest that the principles guiding the adapted training model hold promise for similar contexts.

## 8 FUTURE WORK

Building on this preliminary wireframe, we intend to run a larger cohort to gather more definitive quantitative and qualitative feedback. This will allow us to assess whether the adapted training model demonstrates promise beyond its preliminary pilot, and to strengthen our understanding of the model's efficacy and further refine the curriculum based on evidence from a broader and more diverse participant group.

## 9 CONCLUSION

Adapting leadership-class HPC training for resource-constrained environments suggests that sustainable capacity building requires both technical and cultural transformation. By embedding policy into user education and reframing operations around pragmatic, context-aware practices, the adapted program enabled new adopters to operate legacy KNL systems with confidence and autonomy. More importantly, it highlighted the potential value in re-engineering global HPC training models to reflect the realities of emerging centres rather than the assumptions of mature ones. Future work will expand shared documentation, deepen regional collaboration, and evaluate the long-term impact of this training framework across multiple African HPC sites. We encourage the broader HPC community to consider adopting these principles as a framework for designing training models that meaningfully support emerging and resource-constrained centres.

## ACKNOWLEDGMENTS

The authors would like to thank the many individuals whose expertise, support, and operational efforts contributed to the success of the Building Machines for Science workshop.

We gratefully acknowledge Mitchell Collins, Joe Allen, Erik Ferlanti, Lena Duplechin Seymour, and Bob Garza for their contributions in teaching, mentoring participants, and supporting workshop activities throughout the program.

We also thank Nathaniel Mendoza and Nic Lewis for designing, configuring, and maintaining the training infrastructure used during the workshop.

This workshop was heavily inspired by the original Designing and Administering Large Scale Systems Institute. We would like to recognize Nick Thorne, Cyrus Proctor, and Je'aime Powell for their foundational work in developing the original institute model and curriculum that informed this training effort.

Finally, we would like to acknowledge Dan Stanzione for his continued leadership and support of workforce development, research computing education, and scientific cyberinfrastructure initiatives at Texas Advanced Computing Center.

Some editing of grammar and phrasing was performed using Microsoft CoPilot. NIH Grant G11TW013072.

## REFERENCES

[1] William J. Allen, Kelsey M. Beavers, Erik Ferlanti, Lorenzo Concia, Joshua Urrutia, Ernesto A. B. F. Lima, John M. Fonner, Felix Zuo, H. E. Duplechin Seymour, Ari B. Kahn, Joe Stubbs, Anagha Jamthe, Stephanie N. Baker, Tabish Khan, and James P. Carson. 2025. A Model for Teaching Machine Learning, Deep Learning, and Research Computing to Domain Scientists on HPC Resources. In *Proceedings of the SC '25 Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC Workshops '25)*. Association for Computing Machinery, New York, NY, USA, 401–408. https://doi.org/10.1145/3731599.3767380

[2] Erik Ferlanti, William J. Allen, Ernesto A. B. F. Lima, Yinzhi Wang, and John M. Fonner. 2023. Perspectives and Experiences Supporting Containers for Research Computing at the Texas Advanced Computing Center. In *Proceedings of the SC '23 Workshops of the International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W '23)*. Association for Computing Machinery, New York, NY, USA, 155–164. https://doi.org/10.1145/3624062.3624587

[3] Bryan Johnston, Charles Crosby, Quinn G. Reynolds, Jennifer M. Schopf, and Christopher J. Whalen. 2025. Sticking (with) the Landing A modern case for Knights Landing in Resource-Constrained Environments. *PEARC 2025 - Practice and Experience in Advanced Research Computing 2025: The Power of Collaboration* (7 2025). https://doi.org/10.1145/3708035.3736075;PAGE:STRING:ARTICLE/CHAPTER

[4] Bryan Johnston, Lara Timm, Eugene de Beste, and Mabatho Hashatsi. 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. *Journal of Computational Science Education* 15, 1 (March 2024), 64–71. https://doi.org/10.22369/issn.2153-4136/15/1/13 Presented at Supercomputing SC23.

[5] 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)*. ACM, 1–11. https://doi.org/10.1145/3626203.3670537