SCIENCE FOR AFRICA FOUNDATION

MAPPING SCIENTIFIC COMPUTING TO POWER RESEARCH, BIG DATA, AND AI IN AFRICA

Evidence and Pathways for a Federated Ecosystem

Authors:

Paul Korir[2], Leah Mwangi[1], Kobus Herbst[3], and Evelyn Gitau[1]

[1] Science for Africa Foundation, Nairobi, Kenya

[2] SCI2PRO Africa Ventures Limited

[3] Africa Health Research Institute (AHRI)

EXECUTIVE SUMMARY

Africa's research institutions are engaging with scientific computing, but remain underprepared for the data- and AI-intensive science now reshaping global research. A pan-African mapping exercise conducted between March and April 2026 received 118 institutional submissions from 28 countries, of which 67 were retained after quality review. The analysis found an overall continental readiness of 40.1% (1.20 on a 0–3 maturity scale), with uneven capability across countries, regions, and institution types.

The constraint is not hardware scarcity alone, but ecosystem weakness across compute, artificial intelligence (AI) and graphics processing unit (GPU) access, storage, software, connectivity, energy, procurement, governance, funding, and skills. Convening participants reinforced this view, warning that AI ambition will not translate into durable scientific capability without coordinated investment in the infrastructure and workforce required to support it.

Scientific computing is now foundational research infrastructure and a matter of scientific sovereignty. Without it, African institutions remain dependent on external platforms, proprietary clouds, and foreign-controlled data and analytics.

FEDERATED AFRICAN SCIENTIFIC COMPUTING SERVICE ECOSYSTEM, BUILT AROUND FOUR PRIORITIES:

1. **Ecosystem-oriented infrastructure**: modular, interoperable, tiered services from local training nodes to regional centres and continental coordination.
2. **Demand and workforce development**: activated users, supported workflows, and skilled technical teams, measured by use rather than installed capacity.
3. **Strategic regional investment**: pooled procurement, energy-readiness, lifecycle funding, and National Research and Education Networks as core partners.
4. **Evolved data governance**: FAIR-aligned systems, sovereignty protections, secure access models, and investment in locally relevant African datasets.

The central ask is a shift in posture: scientific computing should no longer be treated as a cost centre or specialist technical service, but as strategic infrastructure for scientific sovereignty, digital transformation, and economic competitiveness. It should be funded, governed, and operated as a coordinated, regionally shared ecosystem.

The call to action is clear: African governments, research institutions, funders, National Research and Education Networks, and technical communities should now move from fragmented pilots to a coordinated continental roadmap for scientific computing. Implemented deliberately, this agenda can translate the ambitions of Agenda 2063 and STISA-2034 into usable research infrastructure, enabling African institutions to shape the datasets, models, discoveries, and innovations that matter for the continent.

INTRODUCTION

Scientific computing has become the third pillar of scientific discovery, standing alongside theory and experimentation as a foundation of modern research. As artificial intelligence reshapes science, Africa's ability to generate knowledge, solve complex challenges and drive innovation will increasingly depend on building integrated, sovereign and sustainable scientific computing ecosystems. The future of discovery will not be determined by scientific ambition alone, but by the computational infrastructure, expertise and collaboration that make transformative research possible

Scientific computing is now a third pillar of science, standing alongside theoretical and empirical inquiry as a foundational mode of knowledge production. Across disciplines, scientific discovery increasingly depends on the capacity to model complex systems, simulate physical and biological processes, analyse large datasets, and reproduce computational workflows. This shift is especially evident in fields such as climate science, genomics, epidemiology, engineering, materials science, geospatial analysis, and artificial intelligence, where progress is constrained not only by scientific imagination but also by access to advanced computational capability. Scientific computing has therefore moved from being a specialised technical support function to becoming core research infrastructure.

Scientific computing may be defined as the application of computational methods, software systems, data infrastructure, and advanced hardware to the solution of scientific problems. While often associated with high-performance computing, it is broader than HPC alone. It spans the full research lifecycle: data acquisition, storage, curation, analysis, modelling, simulation, visualisation, publication, preservation, and reuse. It also depends on research software engineering, high-speed networks, scalable storage, secure platforms, skilled technical personnel, and operational models capable of sustaining long-term scientific use. In this sense, scientific computing is closely aligned with the FAIR principles, since data and workflows must be findable, accessible, interoperable, and reusable if computational research is to be reproducible, cumulative, and widely useful.

The emergence of multimodal generative artificial intelligence has further accelerated the strategic importance of scientific computing. Systems capable of processing and generating text, images, audio, video, code, and scientific data are reshaping expectations about discovery, automation, modelling, and decision support. In research contexts, these systems can support literature synthesis, experimental design, simulation surrogates, imaging analysis, drug discovery, climate analytics, and language technologies. Yet their promise rests on substantial computational foundations: GPU-accelerated systems, high-throughput storage, robust data pipelines, reliable networking, energy infrastructure, and specialised human expertise. There is therefore a significant risk that enthusiasm around AI and digital transformation will outpace the infrastructure required to sustain meaningful impact, producing fragmented pilots rather than durable scientific capability.

The implications extend beyond universities and research institutes. Scientific computing is now part of national infrastructure, with direct relevance to health systems, agriculture, education, climate resilience, industrial development, public administration, and digital sovereignty. African countries cannot rely indefinitely on external platforms, proprietary clouds, or foreign-controlled data and model ecosystems for capabilities that are increasingly strategic to national development. Sovereign scientific computing capacity is therefore essential for protecting sensitive data, supporting African research priorities, developing local and low-resource language technologies, and ensuring that digital transformation is grounded in African institutional control. This aligns with the African Union's Agenda 2063 and the Science, Technology and Innovation Strategy for Africa 2034, both of which position science, technology, innovation, and digital capability as central to inclusive growth, industrialisation, and continental transformation.

Important prior work has already laid foundations for African scientific computing. The HPC Ecosystems Project has helped characterise institutional needs, capacity gaps, and opportunities for coordinated development. Large installations such as Sudan's HPC and grid computing efforts, South Africa's Centre for High Performance Computing, and Morocco's advanced computing investments demonstrate that significant capability exists on the continent. National Research and Education Networks are also emerging as important centres for connectivity, identity, collaboration, and potentially shared computing services. ICTP[1] has contributed through training, capacity building, and support for computational science communities. The African Open Science Platform adds a complementary policy and coordination layer, particularly around open access, research outputs, data stewardship, and the broader open-science agenda, although its infrastructure ambitions remain less mature than its convening and policy functions.

Despite these efforts, Africa's scientific computing landscape remains insufficient in scale, distribution, integration, and operational maturity. Existing capacity is unevenly distributed, often concentrated in a small number of countries or institutions, and frequently constrained by limited funding, fragmented governance, inadequate technical staffing, weak service models, and unreliable supporting infrastructure. This paper argues that scientific computing in Africa must be treated as a coherent ecosystem rather than as isolated hardware procurement. It examines the current state of institutional capability, identifies persistent gaps across infrastructure, software, connectivity, governance, energy, and human capital, and outlines a path toward a more coordinated, sovereign, and sustainable scientific computing agenda for the continent.

1 International Centre for Theoretical Physics

METHODS AND DATA

DATA COLLECTION

118
Institutional submissions received

28
African countries represented

67
Institutional profiles included in the final analysis

50%
Approximate continental representation

This paper draws on institutional capability data collected through a pan-African scientific computing infrastructure mapping exercise conducted between March and April 2026. The exercise aimed to assess scientific computing readiness across African research institutions, with particular attention to AI-driven and data-intensive workloads.

A total of 118 submissions were received from designated institutional representatives across 28 African countries, representing approximately half of the continent. After data cleaning and quality review, 67 submissions were retained for primary analysis as sufficiently complete and internally consistent. The responses are self-reported institutional profiles and are therefore used for strategic synthesis and ecosystem-level interpretation, rather than for verified institutional ranking.

Respondents included public and private universities, research institutes, national laboratories, and a smaller number of non-profit and private-sector organisations. While a handful of countries contributed the largest number of responses, no single country dominated the dataset. The sample therefore offers meaningful geographic and institutional diversity while not claiming full continental coverage.

The survey used a maturity-based framework covering multiple dimensions of scientific computing capability. Each scored item used four ordered categories mapped onto a 0–3 scale. Section scores were calculated using arithmetic means, and composite readiness scores were produced using weighted sections. CPU-based compute infrastructure and AI/GPU capability received the highest weights, followed by software systems and research software engineering. Data infrastructure, connectivity, and energy were weighted moderately, while governance and economic sustainability received lower but still material weights. A "no-compute gate" prevented institutions reporting neither HPC infrastructure nor GPU capability from scoring above Tier 1. The resulting continental readiness average was 40.1%, equivalent to 1.20 on the 0–3 scale.

To address uneven country response counts, country-level averages were adjusted using empirical Bayes shrinkage with k=5 pseudo-observations. This partially shrank small-sample country means toward the continental mean while allowing larger samples to retain greater influence from observed scores. Responses were solicited through Science for Africa Foundation mailing lists, associated institutional networks, professional networks, and LinkedIn.

VIRTUAL CONVENING

Following the survey, all 67 participants who had expressed interest in continued engagement were invited to a focused virtual convening to interpret survey findings and refine future directions. Sixteen participants registered, and nine joined the live session. Participants received pre-read materials consisting of an early draft of the paper, participation guidelines, and curated questions with each participant required to prepare responses to the two questions most relevant to their expertise or institutional context.

The session followed a round-robin format with two structured response rounds and an open discussion segment. Despite the small group, the discussion generated technically grounded insights. Participants endorsed the ecosystem framing and called for a shift from generic training toward workforce development. Regional infrastructure was viewed as a practical response to procurement delays, energy constraints, uneven expertise, and limited institutional funding, but participants stressed that it must be balanced with data sovereignty, secure access, equitable allocation, and clear ownership arrangements. Usability, coordinated implementation, staged technical training, and FAIR-aligned data governance were also identified as central requirements for a sustainable African scientific computing ecosystem.

A more detailed account of the methods and data collected are provided in the adjoining *Supplementary Materials*.

THE STATE OF SCIENTIFIC COMPUTING IN AFRICA

RESULT OVERVIEW

The survey results and convening discussion reveal a continent that is actively engaging with scientific computing, but whose overall capability remains uneven and immature relative to emerging global demands. Most participating institutions fall within the *Nascent* or *Emerging* tiers of scientific computing readiness (Fig. 1 and Fig. 2), indicating limited access to integrated computing environments capable of supporting large-scale scientific workloads. A smaller but meaningful group of institutions occupy the *Developing* tier, demonstrating established operational capability in selected areas. Only a very small number of institutions meet the criteria for *Mature* scientific computing environments.

Across the dataset, the weakest areas consistently include GPU and AI accelerator infrastructure, CPU-based compute capacity, software systems and practices, and connectivity (Fig. 3). This is also reflected in the top-10 priority areas for upgrades (Fig. 9). This pattern is important because it suggests that the challenge facing African scientific computing is not merely one of isolated hardware scarcity, but rather one of ecosystem maturity.

The majority of participants indicated that they are working in areas related to health and/or climate likely reflecting the preponderance of such fields in the target audience (Fig. 4). One striking observation though is the low level of utilisation from industrial applications such as materials science and engineering-related disciplines.

READINESS DISTRIBUTION ACROSS THE CONTINENT

The readiness distribution observed in the survey reveals substantial unevenness across the continent (Fig. 5 and Fig. 6). Higher levels of scientific computing capability are more visible in parts of North and Southern Africa, while comparatively weaker provision is observed across portions of the West–East corridor. These disparities extend beyond hardware infrastructure alone and are also reflected in software maturity, technical expertise, operational coordination, and integration into scientific workflows (Fig. 7). However, there is a strong interest by most participants to engage in federated African scientific computing (Fig. 8).

CONSTRAINTS ARE SYSTEMIC RATHER THAN ISOLATED

Institutions frequently face simultaneous limitations across multiple layers of the ecosystem (Fig. 7). The challenge is therefore not simply the absence of compute infrastructure, but a broader misalignment between infrastructure, software capability, governance, operational sustainability, and technical expertise.

ACCESS AND USABILITY AS MAJOR BARRIERS

Existing compute environments are often fragmented across institutions and projects (Fig. 3). Access may be localised to small research groups or specific departments. Allocation models are frequently informal or uneven. In many cases, weak software environments and limited operational support reduce the practical usability of otherwise functional infrastructure (Fig. 7).

MULTIPLE FACTORS CONTRIBUTE TO WEAK GPU AND AI CAPACITY

Institutions report significant procurement barriers including import restrictions, high acquisition costs, and administrative delays (Fig. 3). Even where hardware is available, operational expertise remains limited. Many institutions lack personnel experienced in managing GPU clusters, distributed training systems, or AI software stacks such as CUDA-enabled environments and orchestration frameworks.

The survey also suggests that demand for AI infrastructure is not yet consistently translated into operational planning. In many cases, institutions recognise the strategic importance of AI capability but lack clear pathways for translating this into sustainable infrastructure programmes.

Importantly, AI acceleration cannot be treated independently of broader scientific computing infrastructure. GPU systems themselves depend upon substantial CPU compute, memory, storage, networking, and operational capability. In practice, strengthening foundational scientific computing capacity substantially advances AI readiness simultaneously.

SOFTWARE WEAKNESS AS A MAJOR CONSTRAINT

One of the strongest signals emerging from the data concerns software systems and operational maturity (Fig. 7). Across many institutions, scientific computing environments appear to rely heavily on ad hoc workflows, inconsistent software practices, and limited research software engineering support.

These findings reinforce the argument that the marginal value of hardware investments diminishes sharply in the absence of strong software capability. Investments in research software engineering, reproducible workflows, usability, and operational tooling should therefore be considered foundational components of scientific computing strategy rather than secondary concerns.

CONNECTIVITY AS A SILENT BOTTLENECK

One of the more surprising findings emerging from the survey concerns connectivity (Fig. 5). Given the considerable investments made by National Research and Education Networks (NRENs) and regional connectivity initiatives across Africa, stronger connectivity scores might have been expected.

The results suggest that the principal issue may not necessarily lie in primary backbone provision, but rather in last-mile delivery within and across research institutions. In practice, institutions may possess nominal access to high-capacity connectivity while still facing operational barriers related to internal network architecture, reliability, firewall restrictions, bandwidth allocation, or access policies.

This finding is particularly significant because connectivity is foundational to any federated or modular scientific computing strategy. Weak last-mile connectivity directly undermines distributed computing, shared infrastructure access, cloud integration, and collaborative data-intensive research.

PROPOSED APPROACHES TO DEVELOP SCIENTIFIC COMPUTING PRACTICE IN AFRICA

Advancing scientific computing in Africa requires a shift from isolated infrastructure acquisition to a federated service ecosystem. The challenge is not simply to procure more computing hardware, but to integrate compute, data, software, connectivity, energy, procurement, funding, governance, and workforce capacity into usable services that support sustained scientific work. The proposed approach therefore rests on four linked priorities: building modular and interoperable scientific computing ecosystems; leading with demand and workforce development; investing strategically in regional capacity while treating procurement, energy, and funding as sustainability gates; and establishing data governance models that enable responsible openness without sacrificing sovereignty. This builds on the paper's broader framing of scientific computing as an integrated environment spanning hardware, data systems, software practices, networking, energy infrastructure, operational models, and human expertise.

AN ECOSYSTEM-ORIENTED APPROACH

10

Essential Components: Compute, storage, connectivity, software, operational tools, user support, data governance, procurement, energy and sustainable financing.

7

Governance Foundations: Allocation, service levels, cost recovery, cybersecurity, institutional responsibilities, sovereignty and equitable access.

5

Open Standards: Cloud-native tools, containerisation, standardised data, FAIR stewardship and research software practices.

If the current state of scientific computing in Africa reflects an ecosystem failure, then the appropriate response cannot be limited to hardware, training, or policy in isolation. Scientific computing becomes useful only when its constituent elements are designed to function together. Compute infrastructure must be linked to storage, connectivity, reproducible software environments, operational tooling, user support, data governance, procurement systems, energy provision, sustainable financing, and skilled personnel. The central design principle should therefore be ecosystem integration rather than infrastructure accumulation.

This requires a shift from procurement-led investment to service-oriented design. Investments should focus on deployable scientific computing units that bundle the basic conditions required for real use: compute, storage, networking, software environments, monitoring, documentation, onboarding, and user support. Such systems should not be treated as dormant assets awaiting expert users, but as usable environments through which researchers can learn, experiment, scale, and collaborate. Usability should therefore be a formal outcome of scientific computing investment.

The ecosystem should also be tiered and modular. Not every country or institution requires the same level of infrastructure, and uniform deployment would be inefficient. A more realistic model would include local access and training nodes, micro-scale workstations, departmental or institutional clusters, national service nodes, regional high-capacity centres, and a continental coordination layer. The purpose of this architecture is to guarantee minimum viable participation: every institution should have a clear access pathway, some local technical capacity, user onboarding mechanisms, and a route into regional or continental compute and data services, even where advanced infrastructure is hosted elsewhere.

Federation should remain a central design principle, but it must be treated as an operational commitment rather than a slogan. Modular and interoperable systems can reduce duplication, support cross-border collaboration, and allow institutions to participate at different levels of maturity. However, federation requires governance: allocation rules, service-level expectations, cost-recovery mechanisms, cybersecurity standards, institutional responsibilities, sovereignty protections, and equitable access arrangements. Without these conditions, federation risks becoming a rhetorical label for disconnected systems.

Interoperability and standards alignment are therefore essential. African scientific computing environments should rely, wherever feasible, on open standards, cloud-native tooling, containerisation, standardised data formats, FAIR-aligned data stewardship, and widely adopted research software practices. These choices reduce long-term dependency on proprietary systems, lower barriers to federation, and make it easier for African infrastructure to interoperate with global research systems while retaining local control.

Finally, the ecosystem must include mechanisms to prevent institutional and disciplinary silos. Participating institutions should map existing compute, storage, software, data, and specialist capacity; expose usable assets through service catalogues; and establish cross-disciplinary scientific computing governance structures. These structures should include researchers, IT units, data governance specialists, procurement and finance personnel, and institutional leadership. The aim is a coordinated system in which existing and new assets are made visible, governed, and usable by real communities of researchers.

LEADING WITH DEMAND AND WORKFORCE DEVELOPMENT

Scientific computing capability cannot be sustained where scientific demand is weak, poorly articulated, or disconnected from operational use. Infrastructure should not be deployed on the assumption that demand will automatically emerge once equipment is installed. Demand must be cultivated through disciplinary engagement, clear use cases, accessible services, and deliberate expansion of computational practice across the research ecosystem.

A demand-led approach begins with awareness but must move beyond general sensitisation. Many researchers, institutional leaders, and policy actors remain unfamiliar not only with what advanced computing can do, but also with the expertise, workflows, recurrent costs, and institutional commitments required to use it effectively. A federated scientific computing system should therefore include a discovery and engagement layer: service catalogues, exemplar workflows, decision-maker briefings, domain-specific demonstrations, researcher onboarding, and clear pathways from scientific problem to computational workflow. These mechanisms are necessary to convert latent interest into activated demand.

Success should not be measured primarily by installed capacity. A programme that reports only machines, cores, GPUs, or storage volumes risks reproducing the familiar failure mode of infrastructure without use. More meaningful indicators include activated users, repeat users, trained users, successful workflows, supported research outputs, institutional demand generated, and progression from introductory exposure to routine computational practice. Demand is both an input into the system and an outcome to be monitored.

This makes workforce development an operational layer of the infrastructure itself. The relevant need is broader than training or human capital in the narrow sense. Scientific computing requires a workforce architecture that includes system administrators, research software engineers, user-support scientists, trainers, data stewards, AI and machine-learning operators, cybersecurity personnel, workflow specialists, institutional champions, and scientific leaders able to connect computational capability to disciplinary problems. It also requires communities of practice that support peer learning, retention, and cross-institutional collaboration.

Deployment should therefore be staged through operational readiness gates. Before infrastructure is installed, each site should demonstrate who will administer the system, who will support users, how software environments will be maintained, how recurrent costs will be met, and where escalation support will come from. A federated model should include remote administration, regional expert pools, secondments, and mentorship arrangements so that local skills gaps do not immediately render systems unusable.

Training is most effective when paired with real use. Preconfigured or "black-box" deployments can be useful where they allow pilot users to begin running scientific workloads quickly while local technical staff are gradually inducted into administration, troubleshooting, and workflow support. Such models are not substitutes for expertise; they are mechanisms for sequencing learning. A practical progression is to deploy a minimally functional and well-supported system, onboard pilot users, train administrators against real problems, provide remote mentorship, and progressively transfer operational control to local teams.

Retention must also be designed into the system. Data science, research software engineering, and advanced computing expertise are scarce and difficult for academic institutions to retain, especially where industry offers more competitive compensation. A sustainable programme should include funded positions, fellowships, career ladders, formal recognition, institutional embedding, and opportunities to work on meaningful African scientific problems. Smaller countries and institutions should not be expected to develop every role family independently; regional training programmes, shared specialist teams, and cross-country communities of practice will be necessary.

STRATEGIC REGIONAL INVESTMENT AND OPERATIONAL SUSTAINABILITY

The survey results indicate significant disparities in scientific computing readiness across the continent, including comparatively lower levels of provision across parts of the West–East corridor. Future investment should therefore be deliberately structured to reduce regional asymmetries in access, capability, and support. This does not require identical infrastructure in every country. It requires strategically distributed nodes designed around regional accessibility, interoperability, training, and coordinated service delivery. Regional investment is important not only for equity, but also for resilience, utilisation, procurement efficiency, and broader scientific participation.

Africa already possesses assets on which such a system can build. National Research and Education Networks have established substantial connectivity infrastructure, with approximately 38 NRENs operating across the continent and coordinated through regional organisations including WACREN[1], UbuntuNet Alliance, and ASREN[2]. These institutions should be treated as core infrastructure partners for authentication, connectivity, routing, access policy, and federated service delivery. South Africa's National Integrated CyberInfrastructure System also offers a useful model, linking data infrastructure through DIRISA[3], connectivity through SANReN[4] and TENET[5], and compute capability through CHPC. The broader lesson is that scientific computing should be organised as a coordinated national and regional ecosystem rather than as a set of isolated institutional projects.

Sustainability must be treated as a first-class requirement. Scientific computing infrastructure has to be economically, operationally, and environmentally realistic from the outset. Commodity-based cluster approaches, Beowulf-style architectures, federated grid models, open-source software ecosystems, containerisation, and standards-aligned tooling remain important because they can reduce costs, strengthen local technical control, and limit dependence on proprietary systems. The strategic opportunity is not simply to acquire more infrastructure, but to connect, operationalise, and govern existing and new capabilities as a durable scientific computing system.

1 West and Central African Research and Education Network
2 Asian States Research and Education Network
3 Data Intensive Research Initiative of South Africa
4 South African National Research Network
5 Tertiary Education and Research Network of South Africa

PROCUREMENT

Procurement is a practical bottleneck that can neutralise even well-designed funding. Advanced computing acquisition is often slowed by tender requirements, vendor guarantees, import procedures, warranty constraints, and rapidly changing technical specifications. The consequence is not merely administrative delay. Failed procurement can lead to expired funding, obsolete quotations, abandoned deployments, and continued dependence on external platforms or overseas analysis.

A continental strategy should therefore include procurement architecture. This should include pre-qualified vendor lists, reusable technical specifications, template tender documents, framework agreements, regional demand aggregation, import and logistics support, warranty guidance, and procurement advisory capacity. Without such mechanisms, each institution repeats the same fragile process independently, often with limited technical support and weak bargaining power.

Procurement should also be linked to regional infrastructure design. Where acquisition barriers are severe, shared regional facilities and infrastructure-as-a-service models can provide access without requiring every institution to procure and operate high-end systems independently. However, regional sharing requires a shift away from narrow ownership cultures. Institutional assets intended to serve wider communities must be governed through clear access policies, allocation procedures, service catalogues, lifecycle costing, and cost-recovery arrangements.

ENERGY

Energy is one of the clearest tests of whether scientific computing infrastructure is real or merely nominal. Advanced computing systems require reliable power, cooling, backup arrangements, facilities management, and recurrent operational funding. In many African settings, grid instability, high energy costs, hot climates, and weak facilities support can turn otherwise valuable infrastructure into an underused or unusable asset. The risk of white-elephant infrastructure is therefore not caused only by lack of skills or procurement delays; it can also arise from unfunded energy obligations.

Energy readiness should be a deployment gate. Site assessment should include grid reliability, available electrical capacity, cooling capacity, UPS and generator arrangements, maintenance support, facilities staffing, and responsibility for recurring utility costs. Sites that cannot support high-density computing should not be excluded from the ecosystem; they should participate through appropriately scaled models, including low-power local systems, training nodes, edge devices, cloud credits, regional compute access, or compute-to-data arrangements.

A realistic architecture should segment infrastructure by power profile. Low-power local infrastructure can support training, prototyping, data preparation, and user onboarding. Medium-scale national systems can support routine research computing. High-density GPU and HPC systems should be concentrated in professionally managed facilities with adequate power, cooling, monitoring, security, and operational staffing. Energy-efficient facility design, workload scheduling, efficient cooling, renewable integration, and hybrid energy systems should be explored, but not treated simplistically: high-density computing still requires stable baseload planning and professional facilities management.

FUNDING

Scientific computing must be funded as a lifecycle commitment rather than a one-time capital purchase. The financing model must cover hardware, procurement, importation, installation, power, cooling, staffing, software, maintenance, cybersecurity, user support, training, governance, and refresh cycles. Many institutions can secure occasional project funds or donor grants, but few can independently sustain the recurrent costs of advanced computing. A credible model must therefore be pooled, recurrent, cross-institutional, sovereignty-aware, and tied to explicit operational responsibilities.

Weak institutional buy-in is itself a funding barrier. Governments, universities, research institutes, and finance units need to understand scientific computing as national research infrastructure, economic infrastructure, and sovereignty infrastructure. The investment case should be connected to public priorities such as health systems, climate resilience, agriculture, education, epidemics, genomics, AI governance, industrial development, and national data protection. Africa is already paying for inadequate infrastructure through outsourced analysis, delayed research, weaker data sovereignty, dependence on external platforms, and lost opportunities to build local technical capability.

Funding should also reward shared use. Project-specific equipment purchases and isolated deployments can create capacity that serves a narrow group while remaining unavailable to the wider scientific community. Major infrastructure proposals should specify who else can use the system, under what access policy, with what user support, how operating costs will be covered, and how the infrastructure will connect to national or regional services. Regional hosting should be funded explicitly, because institutions that host shared compute carry disproportionate recurrent costs. A plausible financial architecture would combine national subscriptions, public investment, donor capital funding, institutional co-funding, subsidised access for priority research, and paid tiers for heavier or commercial use.

EVOLVED DATA GOVERNANCE

A federated scientific computing ecosystem requires data governance that enables collaboration without sacrificing sovereignty, ethics, trust, or reproducibility. The central question is not only where African research data will be stored or processed, but how such data can be made usable, shareable, protected, locally relevant, ethically governed, and scientifically reproducible. Data infrastructure should therefore be treated as a first-order pillar of scientific computing, not merely as storage capacity.

The governing posture should be responsible openness. African scientific computing systems should support cross-border collaboration, reproducibility, and equitable partnerships while preserving sovereignty over sensitive local datasets. This is especially important for health and genomic data, biometrics, indigenous and community knowledge, education data, national datasets, commercially sensitive research, and data linked to vulnerable populations. Responsible openness rejects both unrestricted data release and defensive data closure. It requires clear rules for what can be shared, with whom, under what conditions, through which technical environments, and with what accountability.

The FAIR principles provide a necessary foundation, but they must be implemented operationally. Data infrastructure should include metadata standards, persistent identifiers, catalogues, interoperable formats, documentation, versioning, provenance records, dataset citation practices, access procedures, and reproducibility tooling. A federated African scientific computing ecosystem should therefore include common metadata templates, a federated data catalogue, data-management-plan requirements, controlled-access workflows, reproducible analysis environments, and mechanisms for tracing how datasets are used. FAIR does not mean that every dataset is publicly downloadable; it means that data should be findable, accessible under appropriate conditions, interoperable, and reusable where ethically, legally, and scientifically justified.

Data governance also requires workforce capacity. Institutions need data stewards, data managers, research software engineers, security personnel, ethics and governance specialists, legal support, and user-support teams able to translate rules into practice. These roles are essential for classifying data, preparing documentation, implementing access controls, supporting reproducible workflows, managing data-sharing requests, and advising researchers on legal, ethical, and technical requirements.

Regional infrastructure should be designed around data classification. Low-risk or openly reusable datasets may be hosted or analysed through regional platforms. Sensitive datasets may require national data nodes, trusted research environments (also known as secure data environments), federated analytics, compute-to-data arrangements, anonymisation, audit trails, approved regional facilities, and legally binding data-sharing agreements. In this model, data are local where necessary, regional where appropriate, and open where justified.

The system should also invest in locally relevant African datasets. FAIR governance should not only make existing data reusable; it should help produce better data resources that reflect African scientific, social, environmental, linguistic, and health realities. Priority areas include climate and hydrological data, genomics and public-health datasets, agriculture and food-systems data, geospatial data, education data, African language datasets, and domain-specific reference resources. These datasets should be curated, documented, versioned, cited, and maintained as shared research assets. Without such investment, African researchers will remain dependent on external datasets, models, and benchmarks that often poorly represent local conditions.

Evolved data governance is therefore not an administrative overlay. It is the trust layer through which scientific computing enables collaboration, reproducibility, and innovation while protecting sovereignty. The goal is to treat data as a governed research asset: findable, accessible under appropriate conditions, interoperable, reusable, locally relevant, legally protected, and embedded in trusted institutional arrangements.

IMPLICATIONS FOR GENERATIVE AI DEVELOPMENT IN AFRICA

Generative AI strengthens the case for a federated scientific computing service ecosystem. It depends on the same foundations already identified for scientific computing: compute, storage, connectivity, software environments, energy, procurement, lifecycle financing, data governance, and skilled personnel. The difference is that AI raises the threshold for meaningful participation. Weaknesses that previously constrained scientific computing now directly limit AI development, adaptation, deployment, and oversight.

AI as an Ecosystem Stress Test: Generative AI should not be treated as a separate infrastructure agenda. GPU access alone is insufficient. AI workloads require reliable CPU infrastructure, scalable storage, stable networks, secure data environments, reproducible software stacks, monitoring, user support, and trained operators. Investments labelled as AI capacity building will therefore remain fragile if they bypass broader scientific computing foundations.

Federated Services, Not Isolated Assets: Not every institution can or should host high-density AI infrastructure. Large GPU systems are expensive, power-intensive, operationally demanding, and fast-depreciating. A tiered federated model is more realistic: local nodes for training and prototyping; institutional or national nodes for applied ML, model evaluation, and fine-tuning; regional centres for heavier workloads, secure data enclaves, and specialist support; and a continental layer for standards, procurement templates, governance, catalogues, and benchmarking.

Demand and Workforce Development: Africa's near-term opportunity is likely to lie less in frontier-scale model training and more in adapting, evaluating, deploying, and governing models for local needs, including language, health, agriculture, climate, education, public administration, and scientific research. This requires demand-led development rooted in real domain problems. It also requires a broader workforce: system administrators, ML operations specialists, research software engineers, data stewards, cybersecurity personnel, model evaluators, domain experts, trainers, and ethics and legal specialists.

Sustainability and Data Governance: AI infrastructure should be placed where it can be reliably powered, cooled, secured, staffed, connected, maintained, and refreshed. Procurement, energy reliability, and lifecycle financing are deployment gates, not secondary details. Equally, indigenous AI development depends on governed African datasets that are curated, documented, legally usable, and representative of local realities. The appropriate posture is responsible openness: data should be findable and reusable where possible, protected where necessary, and accessed through secure, trusted arrangements.

Strategic Implication: Generative AI should be positioned as an accelerator of the federated scientific computing agenda, not as a standalone technology programme. Its promise cannot be realised through isolated pilots or hardware purchases alone. It requires coordinated investment in usable services, shared infrastructure, workforce development, governed datasets, operational sustainability, and regional coordination.

CONCLUSION

Scientific computing should be understood not as a cost centre, but as a strategic engine for scientific, technological, and economic development across Africa. The findings of this paper show that the continent's scientific computing challenge is not reducible to hardware scarcity alone. It is an ecosystem problem involving compute, data, software, connectivity, energy, governance, procurement, funding, and workforce capacity. Addressing this challenge therefore requires coordinated investment in scientific computing as foundational infrastructure for research, innovation, and digital sovereignty.

At the scientific level, advanced computational capability enables African researchers to participate more fully in frontier fields such as genomics, climate science, artificial intelligence, materials science, epidemiological modelling, digital health, precision agriculture, and engineering simulation. It strengthens reproducibility, accelerates discovery, and expands the range of scientific questions that can be addressed from within African institutions. Without such capability, African research risks remaining dependent on external platforms, collaborators, datasets, and analytical infrastructure, limiting both scientific autonomy and the continent's ability to respond to locally specific challenges.

At the economic level, scientific computing creates value beyond the academy. Investments in computational infrastructure can support digital industries, AI ecosystems, biotechnology, engineering services, climate services, advanced manufacturing, and startup formation. They also cultivate the technical workforce required for wider digital transformation, including research software engineers, data stewards, system administrators, AI practitioners, and domain specialists able to translate computational capability into usable scientific and industrial outputs. Scientific computing is therefore not only research infrastructure; it is economic infrastructure.

The strategic importance of scientific computing is particularly clear where African institutions seek to transform local data and local knowledge into globally relevant innovation. Climate-informed health systems, AI for African and low-resource languages, genomic surveillance, computational drug discovery, geospatial intelligence, and data-driven agriculture all depend on the ability to process, govern, and analyse data at scale. These capabilities must be developed in ways that preserve sovereignty, protect sensitive data, and support responsible openness, while still enabling collaboration and interoperability across institutions and borders.

The long-term vision, therefore, is not merely technological modernisation. It is the establishment of a continent-wide capability layer that supports scientific sovereignty, regional collaboration, economic competitiveness, and sustained innovation. Achieving this will require moving beyond fragmented initiatives toward a coordinated ecosystem of usable infrastructure, skilled workforce, trusted data governance, sustainable financing, and shared regional capacity. If developed with these principles in mind, scientific computing can become one of the central foundations for Africa's scientific and economic future.

Acknowledgements

We are grateful to all those who took part in completing the survey form on behalf of their institutions. In particular, we are grateful to the following nine (9) participants who joined the virtual convening to share their views:

1. Melaku Takele Abate, Oromia State University, Jimma Campus (Jimma College of Teachers' Education)
2. Mosadoluwa Fasasi, Nubian Research
3. Bryan Johnston, NICIS/CHPC HPC Ecosystems Project
4. Mabvuto Luhanga, MEIRU
5. Mwiza Jolly Mhango, Malawi Liverpool Wellcome Research Programme
6. Sicelo Nkambule, Royal Science and Technology Park
7. Joseph Oroja Osumeje, Ahmadu Bello University
8. Rajib Rusthum P, Poovath Aqua Solutions LLP
9. Eric Tutu Tchao, DIPPER Lab KNUST

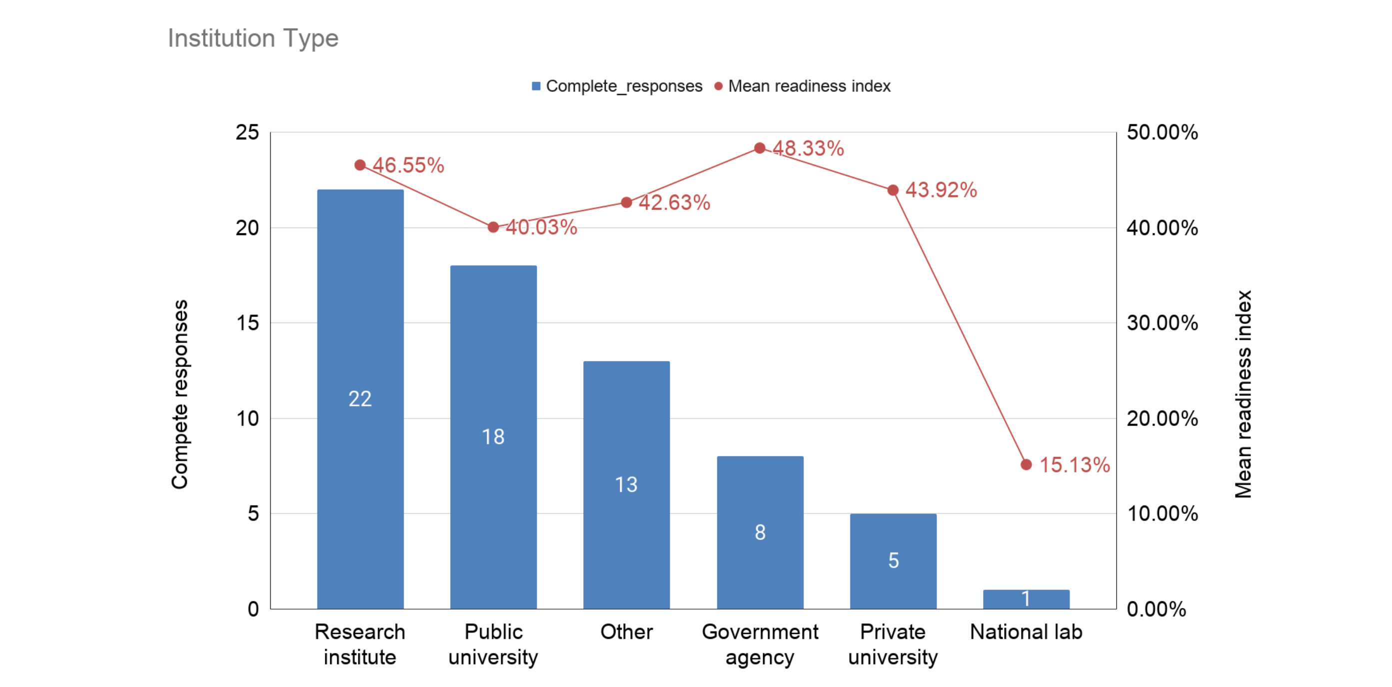

Figure 1: Number of responses and raw mean readiness by institution type.

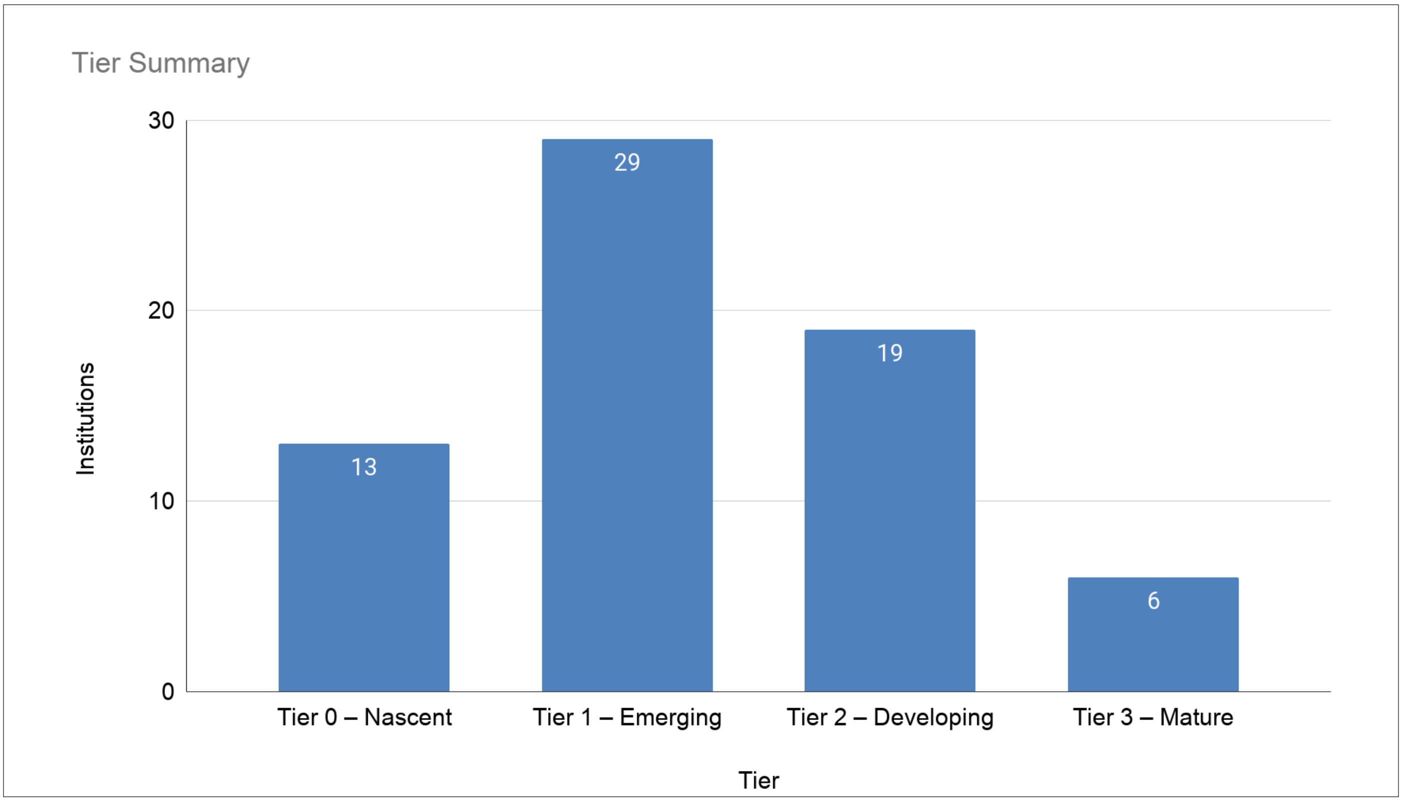

Figure 2: Summary of the number of institutions falling into the four tiers (nascent, emerging, developing and mature).

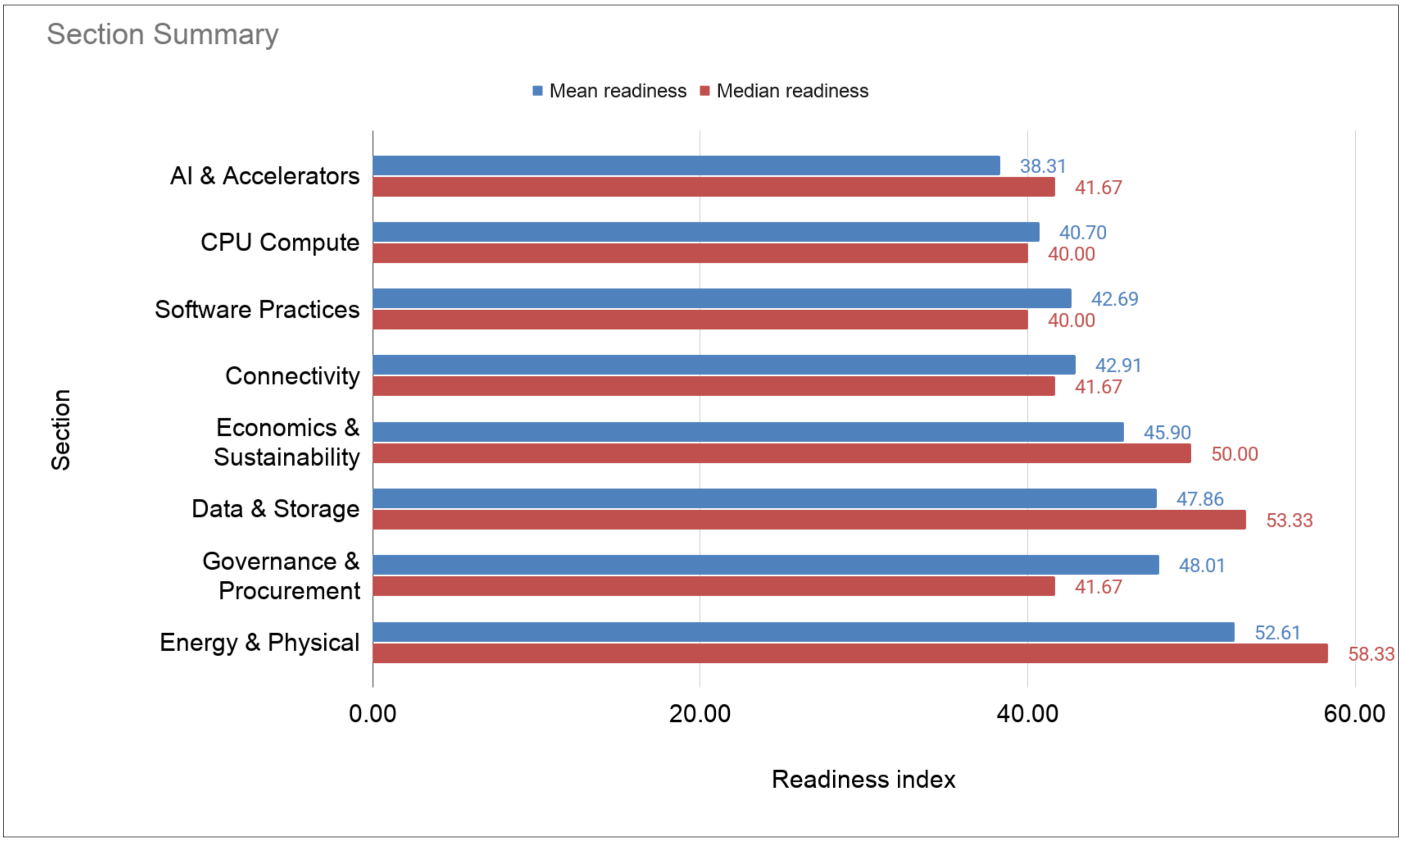

Figure 3: Mean and median raw readiness scores by survey section. Refer to the supplementary material for details of each section in the survey.

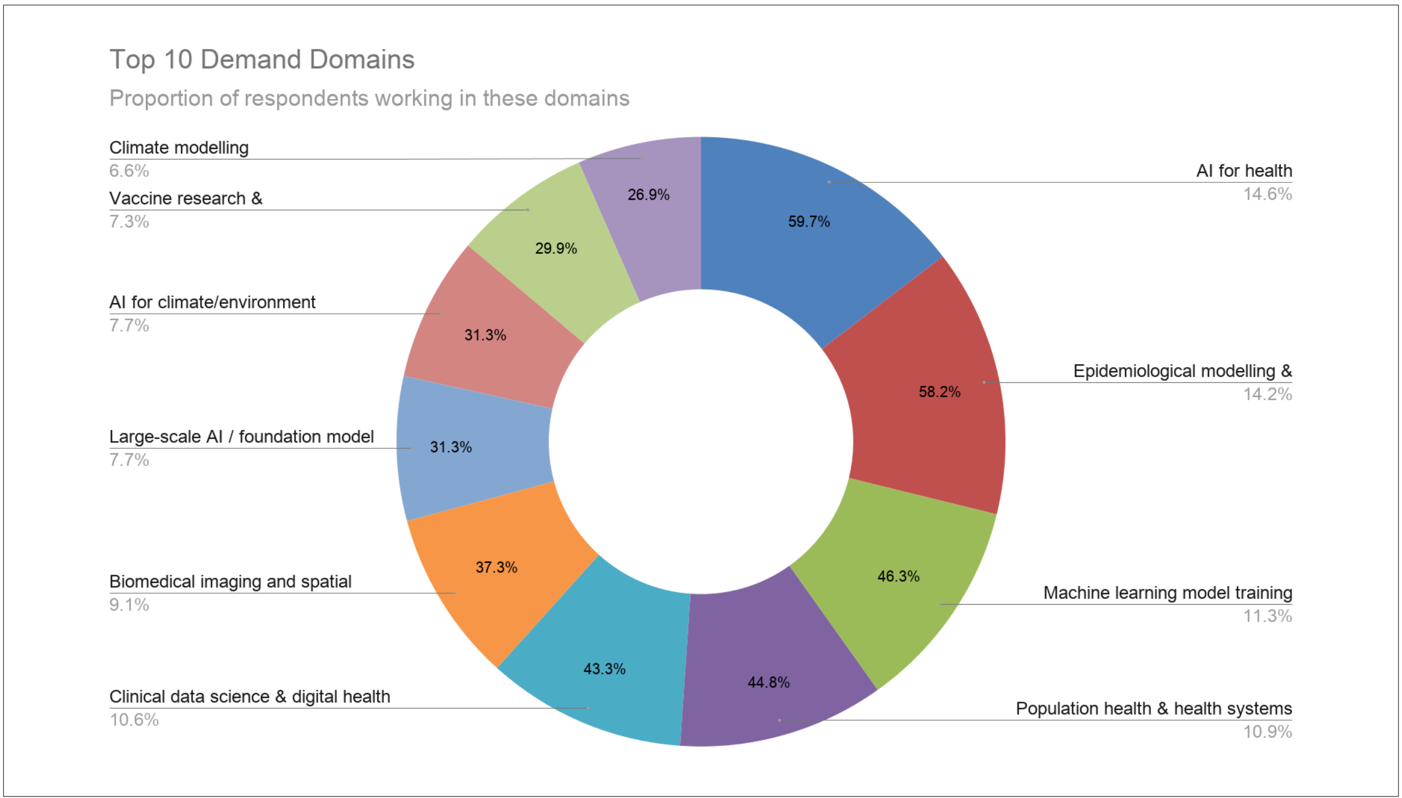

Figure 4: Top 10 domains showing demand for scientific computing. Domains are ordered clockwise from the 12-o'clock position showing a preponderance of health-related applications. However, this could be a bias of the audience polled.

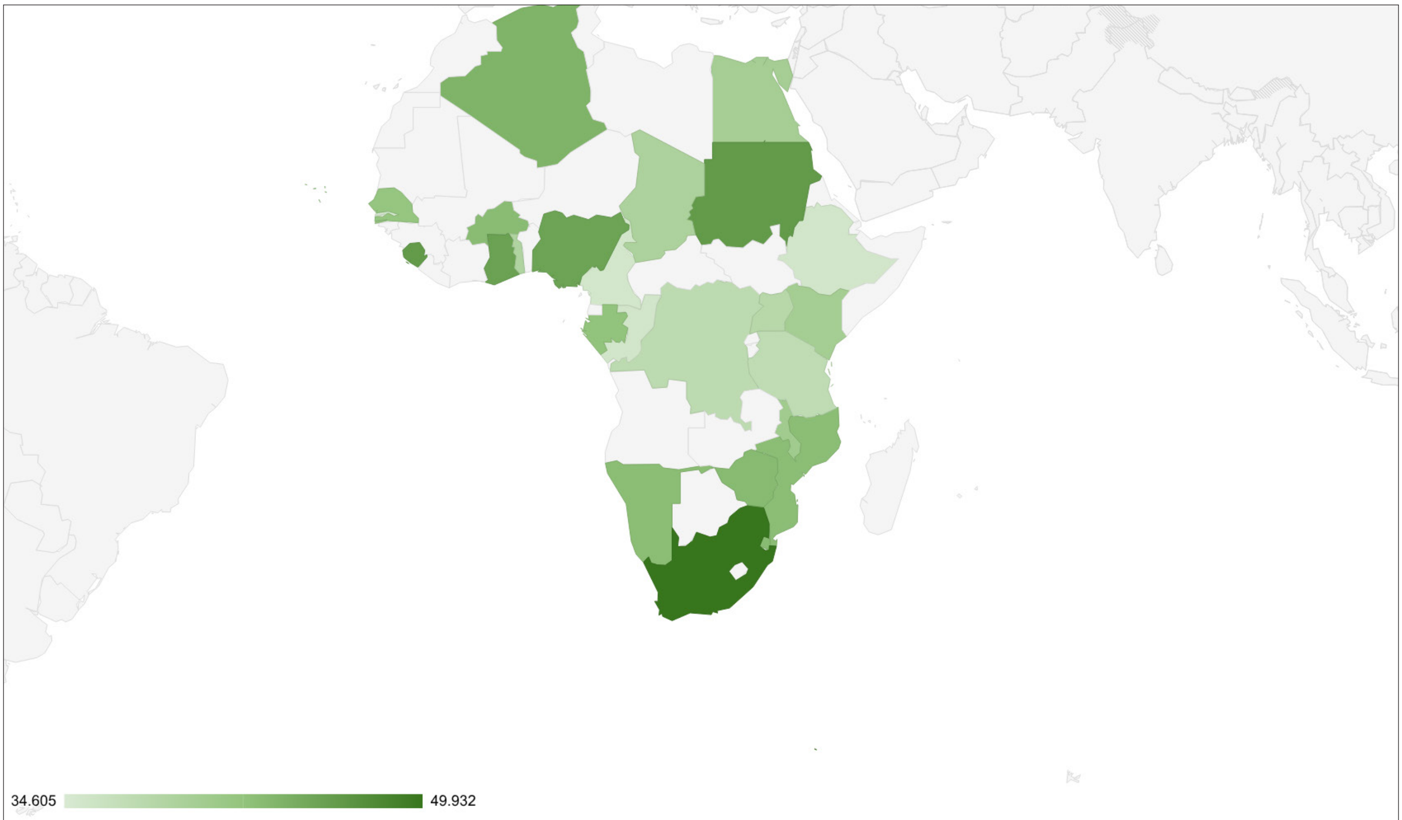

Figure 5: Readiness distribution across Africa. Readiness values are shrunk using empirical Bayes with k=5 for comparison. The readiness range in the key goes from 34.605% to 49.932%.

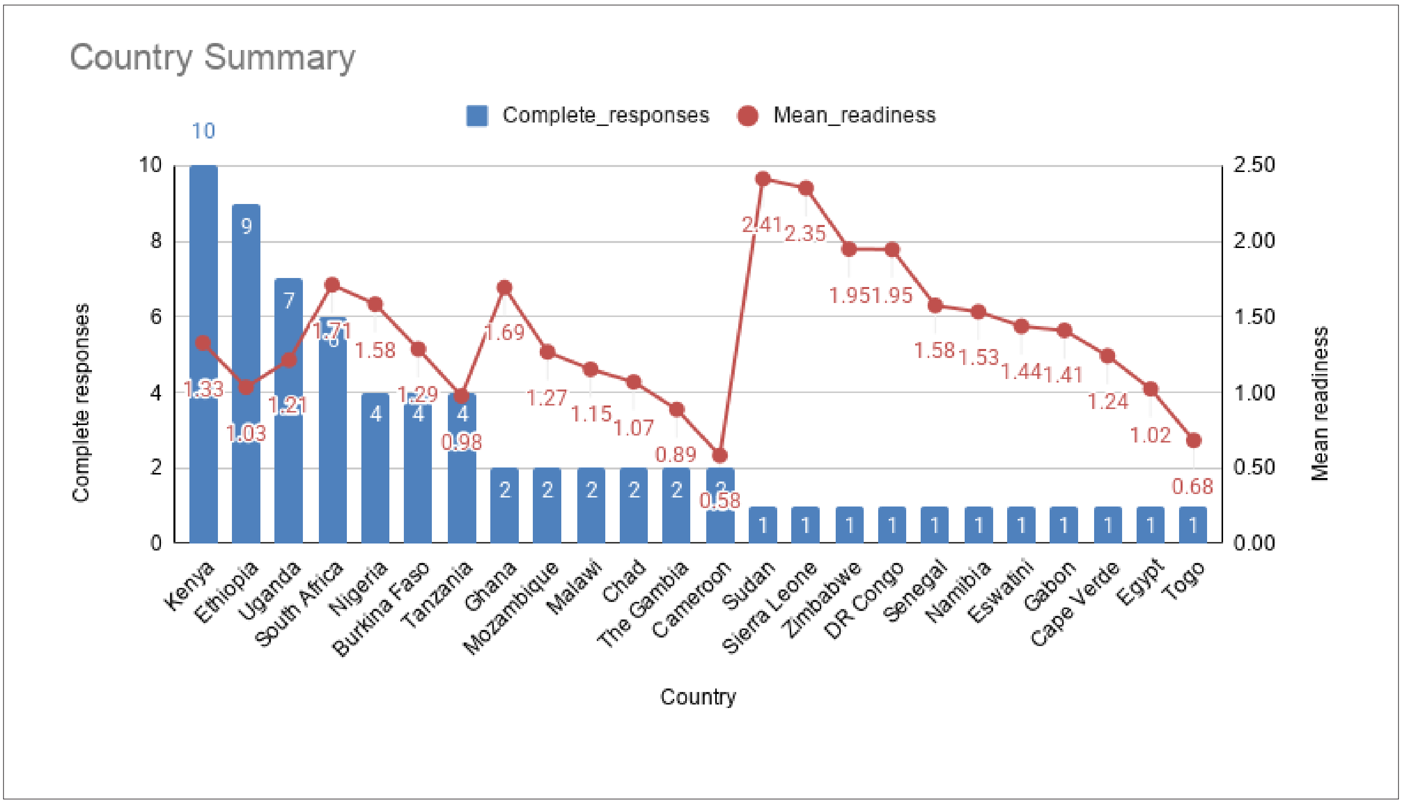

Figure 6: Number of responses and shrunk readiness scores for individual countries.

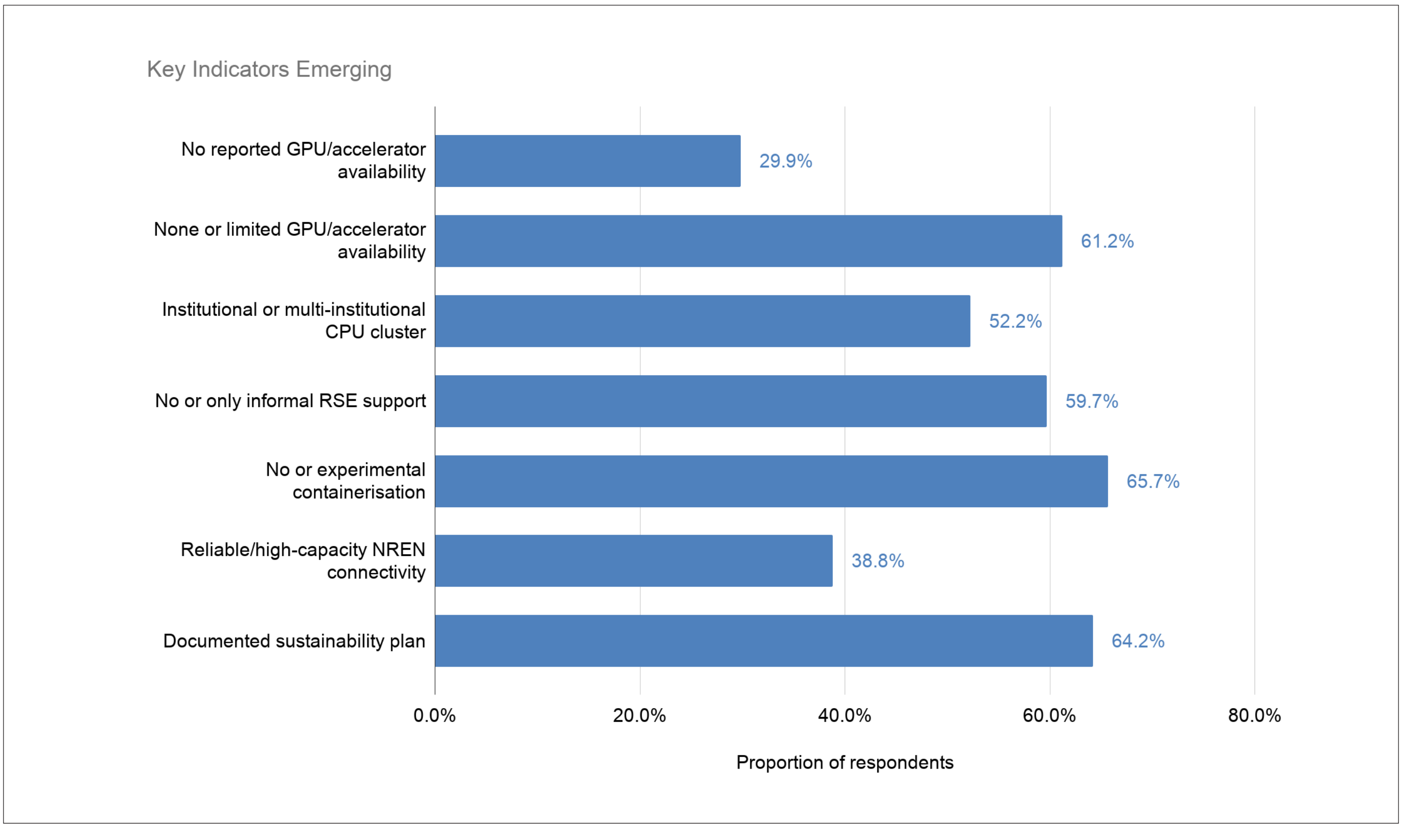

Figure 7: Indicators demonstrating main gaps illustrating systemic factors rather than dominance by a single factor.

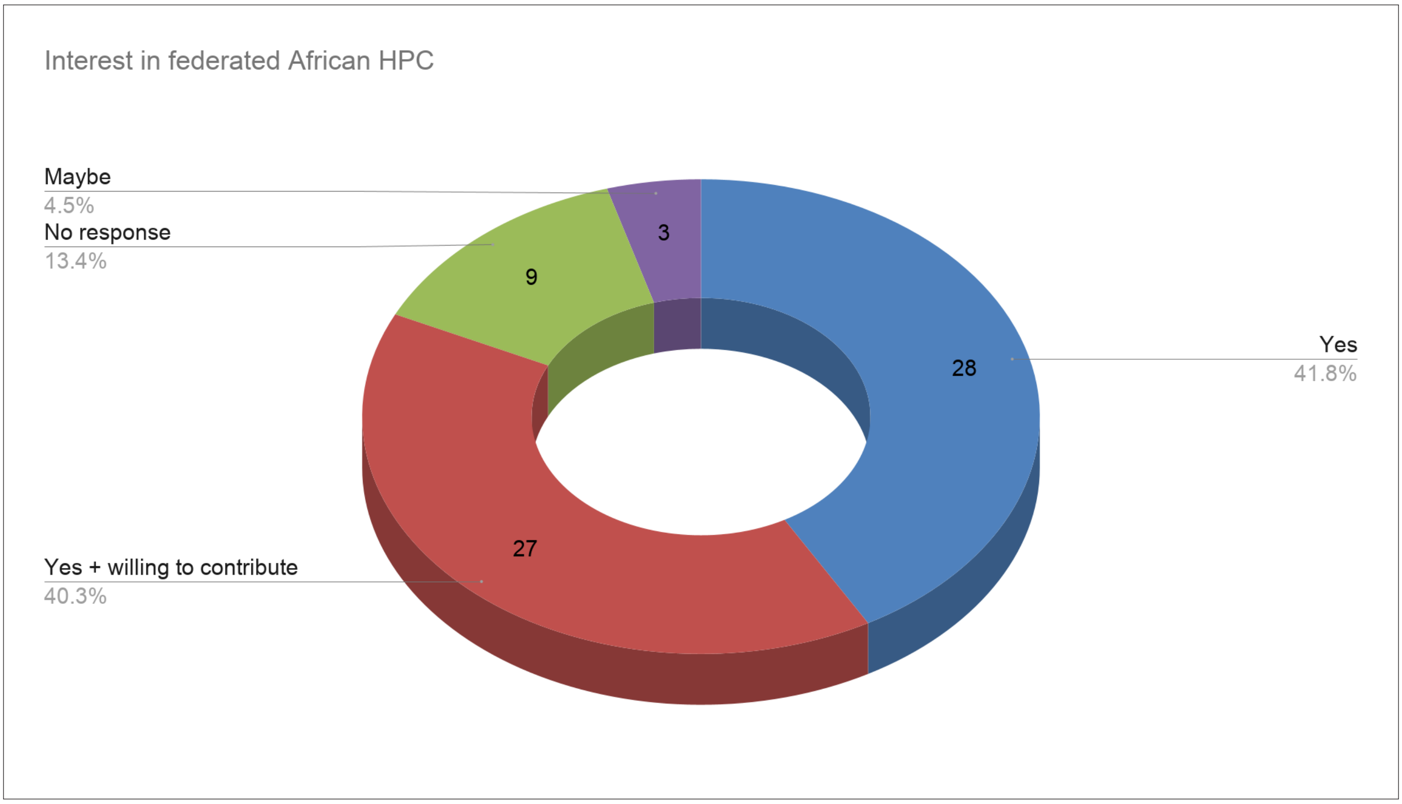

Figure 8: Interest in federated African HPC infrastructure. 83.1% of respondents indicated willingness to participate in federated HPC.

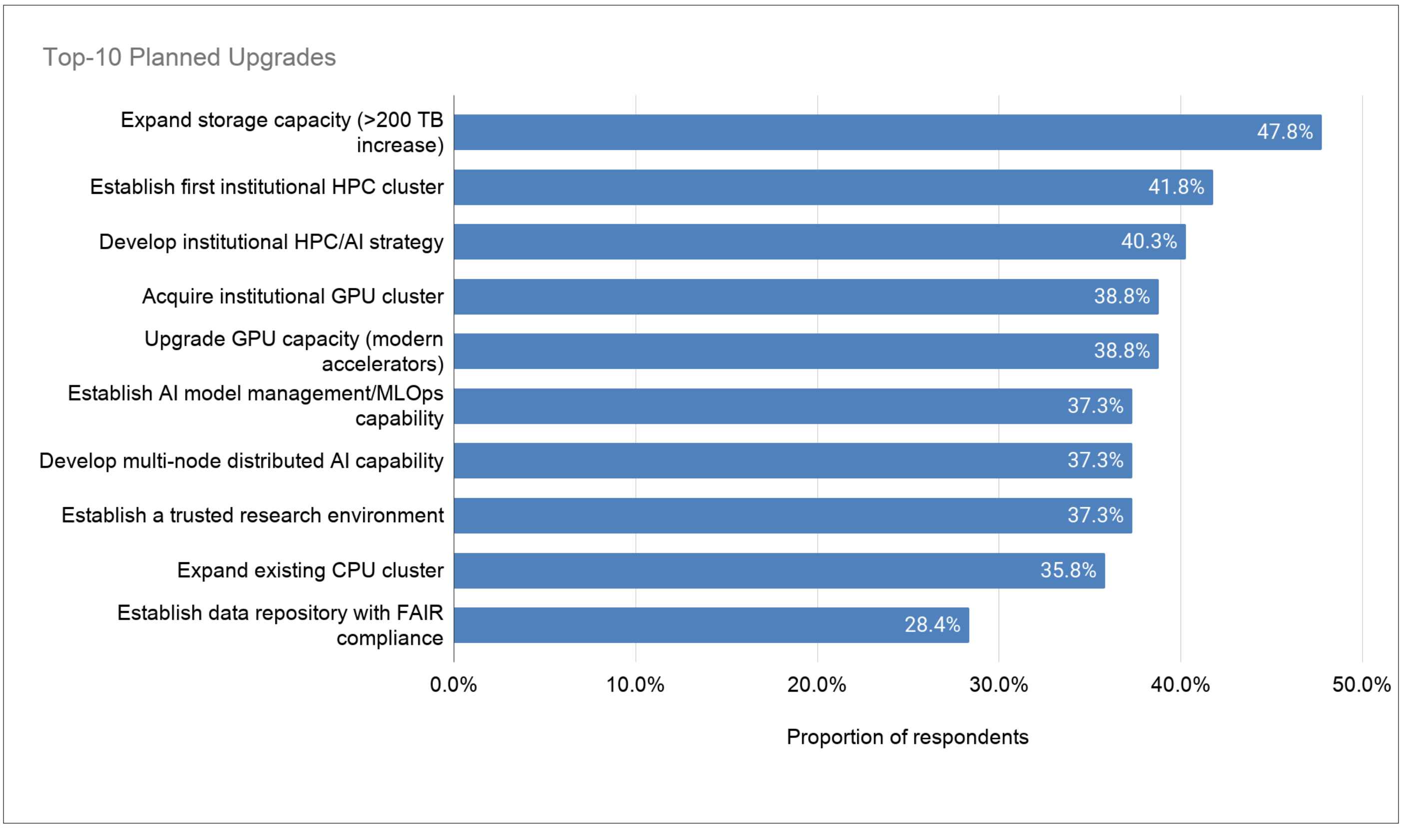

Figure 9: Top-10 upgrades planned over the next 3-5 years. Data storage topping the list is an indication of the existence of nascent compute for existing needs.

SUPPLEMENTARY MATERIALS

METHODS AND DATA

Data Collection

This paper draws upon institutional capability data collected through a pan-African scientific computing infrastructure mapping exercise conducted between March and April 2026. The survey was conducted in English, French, Portuguese and Arabic with translations generated through Google Translate. The exercise sought to develop a structured overview of the current state of scientific computing capability across African research institutions, with particular emphasis on emerging AI-driven and data-intensive scientific workloads.

A total of 118 submissions were received from designated institutional representatives spanning 28 African countries, representing approximately 50% of the continent. Following data cleaning and quality review, 67 submissions were identified as sufficiently complete and internally consistent for inclusion in the primary analysis. Responses consisted of self-reported institutional capability profiles and are therefore intended to support strategic synthesis and ecosystem-level interpretation rather than the production of verified institutional rankings.

The responding institutions included a diverse mix of public and private universities, research institutes, national laboratories, and a smaller number of non-profit and private-sector organisations. While responses from Kenya, Ethiopia, Uganda, and South Africa were numerically strongest, no single country overwhelmingly dominated the dataset. We therefore believe that the resulting dataset provides sufficient geographic and institutional diversity to support a credible pan-African analysis, even though the survey does not claim complete continental coverage.

The survey instrument was designed around a maturity-based framework intended to capture scientific computing capability across multiple dimensions of the ecosystem. Each scored question consisted of four ordered response categories mapped onto a monotonic maturity scale ranging from 0 to 3. Question responses were subsequently aggregated into section-level scores using simple arithmetic means. Composite readiness scores were then computed using weighted sections reflecting the relative strategic importance of different capability areas. CPU-based compute infrastructure and AI/GPU capability were weighted most heavily (weight = 2), followed by software systems and research software engineering practices (weight = 1.5). Data infrastructure, connectivity, and energy were weighted moderately (weight = 1), while governance and economic sustainability were assigned lower but still significant weights (weight = 0.75). To avoid overestimating institutional readiness, a "no-compute gate" was introduced such that institutions reporting neither HPC infrastructure nor GPU capability could not score above Tier 1 regardless of performance in other categories. From the final scores we can also compute a readiness index as a percentage with 0% indicating total unreadiness and 100% indicating total readiness. The overall continental average came to 40.1% (1.20/3.0).

Because the number of responses per country varied substantially, raw arithmetic means were adjusted using an empirical Bayes shrinkage approach to reduce distortion arising from very small sample sizes. Under this approach, country-level averages with few observations were partially "shrunk" toward the continental mean, while countries with larger numbers of responses were influenced more strongly by their observed scores. A shrinkage parameter of k=5 pseudo-observations was used, meaning that countries with fewer than approximately five responses were treated more conservatively in order to improve stability and comparability of estimates. This approach does not eliminate underlying differences between countries, but provides a more robust estimate of ecosystem readiness under conditions of uneven response coverage.

$$\hat{\mu}_c = \frac{n_c \mu_c + k\mu_0}{n_c + k}$$

Where n_c is the number of responses from country c, μ_c is the raw country mean, μ_0 is the continental mean and k=5 is the shrinkage parameter ("pseudo-observations").

Responses were initially solicited through Science for Africa Foundation mailing lists and associated institutional networks. Additional participation was encouraged through public calls distributed via professional networks and LinkedIn.

118

Institutional submissions received from across Africa.

28

African countries represented, covering approximately 50% of the continent.

67

High-quality institutional responses included after rigorous data cleaning and quality assurance.

40.1%

Average continental scientific computing readiness, highlighting significant opportunity for investment and capacity strengthening.

7

Core capability domains assessed, spanning compute infrastructure, AI capacity, software systems, data infrastructure, connectivity, energy, governance and sustainability.

VIRTUAL CONVENING

Following the completion of the data collection phase, all 67 institutional participants who had expressed interest in continued engagement were invited to register for a focused virtual convening intended to synthesise insights emerging from the survey and refine possible future directions for scientific computing development in Africa. A total of 16 participants registered for the convening, of whom nine (9) ultimately participated in the live session.

The convening was intentionally designed as a high-signal and highly structured discussion rather than a broad open forum. Participants received pre-read materials in advance, including an early draft of the position paper, a synthesis of preliminary survey findings, and a set of participation guidelines outlining the structure and expectations of the session. A curated set of questions categorised by topic was also circulated ahead of time to allow participants to prepare concise and considered responses.

Each topic category was structured into four types of questions: *experiential, conceptual, prospective*, and *feedback*-oriented. Participants were asked to prepare responses to the two questions that most strongly resonated with their expertise or institutional context. To maximise participation and maintain focus, written responses were capped at 200 words and presented verbally during the session.

The convening itself was conducted in a round-robin format consisting of two structured response sessions followed by an open discussion segment focused on reactions, clarifications, and synthesis. This structure was designed to ensure that all participants had equitable opportunity to contribute while preserving sufficient time for collective discussion and interrogation of emerging themes. Overall, the session lasted approximately 1 hour and 40 minutes.

Despite the relatively small number of participants, the convening produced rich and technically grounded discussion.

- Participants broadly endorsed the paper's ecosystem framing, emphasizing that Africa's scientific computing constraints are interconnected across compute, data, software, connectivity, procurement, power, governance, funding, and skills, rather than being reducible to hardware scarcity alone.
- They called for greater granularity in the skills agenda, shifting from a narrow focus on "training" or "human capital" toward workforce development, including system administrators, research software engineers, data stewards, AI/ML operators, trainers, user-support personnel, institutional leaders, and communities of practice.
- Regionality emerged as a practical response to several structural barriers, especially long procurement cycles, high energy and cooling requirements, uneven technical expertise, and limited institutional funding. Participants nevertheless stressed that regional infrastructure must be balanced against data sovereignty, secure access, equitable allocation, and clear ownership arrangements.
- There was strong emphasis on the need for coordinated implementation across funders, governments, universities, research institutions, NRENs, and technical communities. Participants argued that sustainability depends on aligning financing, procurement, energy provision, staffing, maintenance, governance, and user support, rather than treating these as separate challenges.
- Participants also highlighted that usability should be an explicit design goal. A hybrid implementation model was proposed: deploy functional "black-box" or preconfigured systems to enable immediate use, while simultaneously providing rigorous, staged technical training so local personnel progressively gain operational control and deeper expertise.
- Data governance was identified as a core part of the ecosystem. Participants emphasized FAIR-aligned data management, responsible openness, controlled sharing of sensitive data, and data-sharing agreements as mechanisms for enabling collaboration while preserving sovereignty, trust, and ethical use.

SCIENTIFIC COMPUTING IN AFRICA: CONVENING QUESTIONS

Please use the following form to prepare your responses for the convening. You will be required to create your own copy as illustrated in the screenshot below.

Participants are also encouraged to review the accompanying document, *Guidelines for Participation*, prior to the session.

To ensure that all participants have an opportunity to contribute during the round-robin discussions, kindly limit each written response to **150–200 words**. Responses should be concise, focused, and grounded in practical experience, strategic insight, or institutional perspective.

Please prepare responses to **at least one question from each thematic category**. During the convening, each participant will be given an opportunity to present their **top two responses** as part of the structured round-robin session with the possibility of a third if time allows.

Participants are encouraged to:

- Focus on clear and actionable insights rather than broad general commentary
- Draw where possible from institutional experience, operational challenges, or strategic lessons
- Engage constructively with both opportunities and constraints emerging from the survey findings and proposed approaches
- Consider both immediate operational realities and longer-term continental implications

The convening is intended to support focused synthesis and informed discussion toward the development of a coordinated scientific computing roadmap for Africa. Contributions that are technically grounded, solution-oriented, and concise will greatly strengthen the quality of the discussion and the resulting position paper.

CATEGORY	QUESTIONS
Computer Infrastructure	**Experiential** 1. In your institutional context, what is currently the single greatest obstacle preventing effective use of scientific computing infrastructure? 2. Several institutions reported limited access to GPU-enabled infrastructure. How does this limitation practically affect research workflows in your environment? **Conceptual** 3. To what extent should scientific computing infrastructure in Africa be approached as a shared regional capability rather than an institutional asset? **Prospective** 4. Looking ahead 5–10 years, which scientific workloads do you believe will place the greatest pressure on African compute infrastructure? **Feedback** 5. The pre-read proposes modular and federated scientific computing infrastructure deployed at different times
Your response *150-200 words to be read*	

SCIENTIFIC COMPUTING IN AFRICA: CONVENING QUESTIONS

CATEGORY	QUESTIONS
Data Infrastructure and Governance	**Experiential** 1. Have data governance, sovereignty, or cross-border data-sharing requirements constrained scientific collaboration or compute usage in your institution? Please elaborate briefly. **Conceptual** 2. What balance should Africa strike between open scientific collaboration and sovereign control of sensitive datasets? **Prospective** 3. Which areas of African science are most likely to generate data volumes that current institutional infrastructure will struggle to support? **Feedback** 4. The proposed approaches emphasise standards-aligned and interoperable infrastructure. What practical barriers exist today to achieving interoperability across institutions?
Your response *150-200 words to be read*	

CATEGORY	QUESTIONS
Connectivity and Federation	**Experiential** 1. The survey suggests that connectivity remains a significant bottleneck despite substantial NREN investment. Does this align with your experience? If so, where do the practical bottlenecks occur? **Conceptual** 2. What role should NRENs play in the future scientific computing ecosystem beyond connectivity provision? **Prospective** 3. What capabilities would become possible if reliable high-throughput connectivity between African research institutions became routine? **Feedback** 4. The pre-read proposes federated access to distributed infrastructure. What operational or governance challenges would need to be solved to make this workable?
Your response *150-200 words to be read*	

CATEGORY	QUESTIONS
Software Systems and Research Software Engineering	**Experiential** 1. How do limitations in software engineering support or operational tooling affect reproducibility, scalability, or collaboration in your environment? **Conceptual** 2. Should research software engineering be treated as core scientific infrastructure in the same way as compute and storage? **Prospective** 3. How do you expect the rise of AI-assisted software development and scientific workflows to change the demands placed on African scientific computing ecosystems? **Feedback** 4. The proposed approaches place strong emphasis on usability and software maturity. In your view, are these appropriately prioritised relative to hardware investments?
Your response *150-200 words to be read*	

SCIENTIFIC COMPUTING IN AFRICA: CONVENING QUESTIONS

CATEGORY	QUESTIONS
Human Capital and Skills	**Experiential** 1. Which scientific computing skills are currently most difficult to recruit, retain, or develop within your institution? **Conceptual** 2. What is the most effective model for developing sustainable scientific computing expertise in Africa: specialised centres, distributed institutional training, regional programmes, or some combination? **Prospective** 3. What capabilities should an African scientific computing workforce possess within the next decade to remain globally competitive? **Feedback** 4. The phased investment model proposed in the pre-read initially prioritises human capital and software capacity over large-scale hardware deployment. Do you agree with this sequencing? Why or why not?
Your response *150-200 words to be read*	

CATEGORY	QUESTIONS
Governance, Sustainability, and Strategy	**Experiential** 1. Which governance or procurement barriers most significantly delay or complicate scientific computing development in your context? **Conceptual** 2. What would a sustainable African scientific computing ecosystem realistically look like in operational and financial terms? **Prospective** 3. If substantial investment became available today, what types of investments would generate the greatest long-term impact over the next decade? **Feedback** 4. What important risks, blind spots, or unintended consequences do you believe are insufficiently addressed in the proposed approaches?
Your response *150-200 words to be read*	

CATEGORY	QUESTIONS
Closing / Synthesis	1. What is the single most important action that should be prioritised over the next 3–5 years to advance scientific computing in Africa? 2. What existing strengths or emerging opportunities give you the greatest optimism about the future of scientific computing on the continent?
Your response *150-200 words to be read*	

GUIDELINES FOR PARTICIPATION

Pan-African Scientific Computing Convening
14 May 2026 | Virtual Session (Zoom)

To ensure a focused, productive, and respectful discussion, we kindly ask all participants to observe the following guidelines:

Before the Session

- Participants are requested to review the pre-read materials listed below and prepare their responses to the guiding questions in advance.
  - » Towards Accessible and Sustainable Scientific Computing in Africa
  - » Scientific Computing in Africa: Convening Questions
- Please join the meeting **at least 5 minutes before the scheduled start time** so that we may begin promptly.
- Kindly ensure that your display name reflects your full name and institution.
- We encourage participants to join from a quiet environment with a stable internet connection where possible.

During the Session

- Please keep your microphone muted unless invited to speak.
- The session will follow a structured round-robin format. Participants will be given a brief, time-bound opportunity to present their prepared responses.
- To ensure equitable participation across institutions and regions, we kindly ask participants to keep interventions concise and focused on the guiding questions.
- Immediate reactions and discussion segments will also be time-limited to maintain momentum and allow broad participation.
- We encourage contributions that are constructive, evidence-informed, and solution-oriented.

Recording and Documentation

- The session will be recorded for purposes of synthesis and preparation of the position paper.
- Notes and selected proceedings from the discussion may be shared with participants following the convening.
- Please avoid sharing confidential or institutionally sensitive information that is not intended for broader circulation.

General Conduct

- Participants are encouraged to engage respectfully with differing perspectives and experiences.
- The convening is intended as a collaborative and technically grounded discussion focused on strengthening scientific computing capability across Africa.
- Where possible, contributions should focus on practical insights, operational experience, and strategic recommendations.

We thank all participants for their time, expertise, and commitment to advancing scientific computing on the continent.

SURVEY QUESTIONS BY SECTION

Section A: Institutional Profile

1. Institution name – free text
2. Country – dropdown (see List of countries at end)
3. Institution type – single select
   - Public university
   - Private university
   - Research institute
   - National lab
   - Government agency
   - NREN
   - International research institution
   - Private research lab
   - Other: _____________
4. Designated submitting authority – name (free text)
5. Designated submitting authority - designation (free text)
6. Designated submitting authority - department (free text)
7. Designated institutional address for clarification or coordination related to this submission (email address) *Declaration: "by submitting this information, you confirm that you consent to being contacted regarding this institutional capability profile and related strategic initiatives. Contact details will be used solely for coordination and clarification purposes and will not be shared publicly."*
8. Primary research domains supported – multi-select
   - Health/Biomedical
   - Genomics
   - Climate
   - Agriculture
   - Engineering
   - Physics
   - Chemistry
   - AI/ML
   - Social science
   - Other: _____________

Section B: Data Infrastructure & Storage Capacity

B1. Total institutional storage capacity available for research (online + nearline, excluding personal laptops)

A. <50 TB
B. 50–200 TB
C. 200 TB–1 PB
D. >1 PB

B2. Primary storage architecture

A. Mostly individual/departmental ad hoc storage (external drives, unmanaged servers)
B. Central shared storage but limited redundancy/management
C. Managed institutional storage (access control, quotas, monitoring)
D. Tiered architecture (hot/warm/cold) with defined lifecycle + automation

B3. Ability to support large-scale datasets (>100 TB) for active projects

A. Not possible
B. Possible for one-off projects with significant manual effort
C. Supported for multiple projects with established process
D. Supported routinely, including high-throughput ingestion + sharing

B4. Data governance framework status

A. None / unclear
B. Informal guidelines only
C. Institutional policy (data classification, access, retention)
D. Auditable compliance (formal governance + regular review; aligns with national/international requirements)

B5. Backup & disaster recovery

A. No reliable backups
B. Backups exist but manual/irregular, restore not tested
C. Automated backups + periodic restore tests
D. Disaster recovery plan (cross-site replication or equivalent) + tested procedures

Section C: Compute Infrastructure (CPU-Based HPC)

C1. Presence of HPC/HTC compute environment (CPU)

A. None (desktop/lab servers only)
B. Small departmental cluster or shared server(s)
C. Institutional cluster with managed access
D. Multi-institutional or national node / institutional cluster with clear expansion pathway

C2. Total CPU capacity available for research (best estimate)

A. <256 cores
B. 256–2,000 cores
C. 2,000–10,000 cores
D. >10,000 cores

C3. Job scheduling / resource management

A. None (manual SSH / first-come-first-served)
B. Basic queueing (scripts/manual allocation)
C. Full scheduler (e.g., Slurm/PBS/HTCondor)
D. Scheduler + accounting + quotas + monitoring + documented policies

C4. Utilisation / access pressure (choose the closest)

A. Compute is rarely used / no queue pressure
B. Moderate use; wait times usually <1 day
C. High use; wait times commonly 1–3 days
D. Capacity constrained; wait times often >3 days / frequent rejections due to limits

C5. Expansion capability (next 24 months)

A. No realistic path to expand (funding/space/power blocking)
B. Expansion possible but unfunded / uncertain
C. Expansion planned with partial funding and identified site requirements
D. Expansion budgeted and scheduled (procurement path and operations plan in place)

Section D: AI & Accelerator Infrastructure

D1. GPU/accelerator availability for research (institutional or managed access)

A. None
B. Limited (e.g., 1–7 GPUs total or consumer-grade only)
C. Moderate (e.g., 8–63 GPUs and/or at least some data-centre class GPUs)
D. Large-scale (≥64 GPUs and/or multi-node GPU cluster)

D2. GPU access model

A. Ad hoc (personal machines / unmanaged)
B. Shared but informal (first-come-first-served, unclear priorities)
C. Managed access (allocation policy, scheduler/booking, support)
D. Managed access + fair-share/quotas + project allocation + usage reporting

D3. Support for distributed GPU training

A. Not supported
B. Possible only via manual setup by expert users
C. Supported for multi-GPU single-node and some multi-node capability
D. Routinely supported multi-node distributed training with documented environment/modules

D4. AI infrastructure roadmap (next 3–5 years)

A. None
B. Under discussion; no costed plan
C. Costed plan exists; funding partially identified
D. Active roadmap with approved budget and implementation plan

Section E: Software Systems & Operational Practices

E1. Version control standardisation (Git, etc.)
A. No standard; code mostly shared via email/USB
B. Individual Git use; no institutional standard
C. Team-based Git workflows (GitHub/GitLab) used in key projects
D. Institution-standard Git + access control + contribution guidelines

E2. Containerisation (Docker/Apptainer/Podman)
A. Not used
B. Used by a few individuals experimentally
C. Used regularly in projects; documented images for key workflows
D. Standard practice across HPC/AI environments; institutional registry and guidance

E3. Workflow/pipeline management (Snakemake/Nextflow/etc.)
A. Mostly manual scripts/notebooks
B. Some reusable scripts; limited provenance
C. Versioned pipelines used in multiple projects
D. Reproducible workflows with provenance, CI checks, and shared templates

E4. CI/CD, testing, reproducibility standards
A. No testing or reproducibility standards
B. Occasional tests; project-dependent
C. Standard testing + environment management in key repos
D. Institution-wide standards for testing, releases, and reproducibility

E5. Research Software Engineering support (RSE)
A. None
B. Informal support (students/volunteers)
C. Dedicated RSE(s) supporting multiple projects
D. RSE team with defined mandate, career pathway, and service model

Section F: Connectivity & Network Integration

F1. Effective institutional internet bandwidth available to research workloads
A. <100 Mbps
B. 100 Mbps–1 Gbps
C. 1–10 Gbps
D. >10 Gbps

F2. NREN connectivity
A. Not connected / no relationship
B. Connected indirectly / limited service
C. Connected with reliable service for research
D. High-capacity NREN integration (peering, advanced services, dedicated links)

F3. International connectivity for data-intensive collaboration
A. Frequently limiting (large transfers impractical)
B. Sometimes limiting
C. Generally adequate for large transfers
D. High-performance international connectivity (routine large-scale transfers)

F4. Firewall/proxy constraints impacting research software/package access
A. Severe restrictions; frequent blocks
B. Moderate constraints; workarounds needed
C. Minor constraints; exceptions process works
D. Minimal constraints; research-friendly network posture

Section G: Energy & Physical Infrastructure

G1. Power reliability for compute facilities
A. Outages daily or near-daily
B. Outages weekly
C. Outages rare (<monthly)
D. Data-centre grade uptime (targeting ≥99.9%)

G2. Backup power capability
A. None
B. UPS only (short bridging)
C. UPS + generator (covers typical outages)
D. Redundant backup (N+1) with fuel/maintenance plan and tests

G3. Cooling suitability for high-density compute
A. Office-grade cooling; not suitable
B. Server room retrofit; limited headroom
C. Purpose-designed facility adequate for current density
D. Scalable cooling designed for expansion (high-density readiness)

G4. Energy cost as a limiting factor
A. Severe constraint; expansion not feasible
B. Moderate constraint; limits scale/frequency
C. Manageable; included in operating model
D. Optimised/negotiated (tariffs/efficiency/PPAs); not a major constraint

Section H: Governance & Procurement Environment

H1. Typical procurement timeline for IT/compute equipment
A. >12 months
B. 6–12 months
C. 3–6 months
D. <3 months

H2. Hardware import barriers (accelerators/servers/networking)
A. Severe (frequent delays/high tariffs/uncertainty)
B. Moderate (delays common but manageable)
C. Minor (occasional delays)
D. Streamlined (predictable process)

H3. Institutional IT/HPC governance model
A. Unclear ownership; ad hoc decisions
B. Ownership exists but limited policy/process
C. Defined governance (roles, policies, allocation rules)
D. Mature governance (steering committee, SLAs, reporting, auditability)

H4. Alignment with national HPC/AI policy
A. No relevant policy / unknown
B. Policy discussions only
C. Draft/active policy exists but limited implementation
D. Active national strategy with funding mechanisms and institutional alignment

Section I: Economics & Sustainability

I1. Primary funding model for compute infrastructure
A. No stable funding (ad hoc/personal)
B. Short-term grants only
C. Mixed grants + institutional support
D. Dedicated institutional/national budget line(s)

I2. Multi-year sustainability plan for operations (staffing, power, maintenance, refresh)
A. None
B. Informal/implicit
C. Documented plan but not fully funded

D. Documented plan with committed funding and refresh cycle

I3. Cost recovery / access charging model (if applicable)

A. None and no mechanism possible
B. Considering options / informal recharge
C. Defined model (internal chargeback or project allocations)
D. Mature model with transparent pricing, reporting, and governance

I4. Institutional budget allocation for compute expansion (next 3–5 years)

A. None
B. Proposed but not approved
C. Partially approved / contingent
D. Approved and scheduled

Section J: Strategic Priorities & Engagement Interest

J1. Priority research domains requiring more compute (multi-select, select up to five primary domains)

Genomics & bioinformatics (e.g., sequencing, variant analysis)

- Epidemiological modelling & outbreak simulation
- Clinical data science & digital health analytics
- Drug discovery & molecular simulation
- Vaccine research & immunoinformatics
- Population health & health systems modelling
- Biomedical imaging & spatial omics

Artificial Intelligence & Data Science

- Machine learning model training (moderate scale)
- Large-scale AI / foundation model training
- Computer vision
- Natural language processing (including African languages)
- AI for agriculture
- AI for health
- AI for climate/environment

Climate, Environment & Earth Systems

- Climate modelling
- Weather forecasting
- Hydrological modelling
- Environmental monitoring (remote sensing)
- Oceanography & atmospheric science

Agriculture & Food Systems

- Crop modelling
- Precision agriculture analytics
- Soil and environmental modelling
- Food security simulation

Engineering & Physical Sciences

- Computational fluid dynamics
- Materials science simulation
- Structural/mechanical modelling
- Energy systems modelling
- Astrophysics / cosmology
- Nuclear / high-energy physics

Social & Economic Systems

- Economic modelling
- Agent-based modelling
- Urban systems modelling
- Policy simulation

Cross-Cutting / Data-Intensive

- Large-scale data integration platforms
- Knowledge graph construction
- Real-time analytics systems
- National data repositories

Other

- Other (please specify – short text)

J2. Planned upgrades next 3–5 years (multi-select + optional short text - select all that apply)

Compute Expansion

- Establish first institutional HPC cluster
- Expand existing CPU cluster
- Acquire institutional GPU cluster
- Upgrade GPU capacity (modern accelerators)
- Develop multi-node distributed AI capability

Data Infrastructure

- Expand storage capacity (>200 TB increase)
- Implement tiered storage architecture
- Establish data repository with FAIR compliance
- Establish a trusted research environment
- Implement cross-site backup or disaster recovery

Software & Operational Capability

- Establish Research Software Engineering (RSE) unit
- Adopt institutional Git and CI/CD standards
- Standardise containerised research workflows
- Deploy workflow management systems (e.g., Nextflow/Snakemake)
- Establish AI model management/MLOps capability

Connectivity & Network

- Upgrade institutional bandwidth
- Establish or upgrade NREN connection
- Improve international backbone connectivity
- Deploy research-friendly firewall/proxy policies

Energy & Physical Infrastructure

- Upgrade server room/data centre facility
- Improve cooling systems for high-density compute
- Install backup generation capacity
- Explore renewable energy integration for compute

Governance & Sustainability

- Develop institutional HPC/AI strategy
- Establish formal governance structure for compute resources
- Develop multi-year sustainability plan
- Introduce cost-recovery or allocation model

Regional / Federated Participation

- Join or host federated African HPC initiative
- Participate in shared GPU pools
- Contribute compute to regional consortium
- Host national data node

Other

- Other planned upgrade (short text)

J3. Interest in federated African HPC collaboration

A. Not now
B. Maybe
C. Yes
D. Yes + willing to contribute

J4. Willingness to participate in technical working groups

A. No
B. Maybe
C. Yes
D. Yes + nominate focal point

LIST OF COUNTRIES

1. Algeria
2. Angola
3. Benin
4. Botswana
5. Burkina Faso
6. Burundi
7. Cameroon
8. Cape Verde
9. Central Africa Republic
10. Chad
11. Comoros
12. Congo Brazzaville
13. Djibouti
14. DR Congo
15. Egypt
16. Equatorial Guinea
17. Eritrea
18. Eswatini
19. Ethiopia
20. Gabon
21. Ghana
22. Guinea
23. Guinea-Bissau
24. Ivory Coast
25. Kenya
26. Lesotho
27. Liberia
28. Libya
29. Madagascar
30. Malawi
31. Mali
32. Mauritania
33. Mauritius
34. Morocco
35. Mozambique
36. Namibia
37. Niger
38. Nigeria
39. Rwanda
40. Réunion
41. Senegal
42. Seychelles
43. Sierra Leone
44. Somalia
45. South Africa
46. South Sudan
47. Sudan
48. São Tomé and Pincipe
49. Tanzania
50. The Gambia
51. Togo
52. Tunisia
53. Uganda
54. Western Sahara
55. Zambia
56. Zimbabwe

SCIENCE FOR AFRICA FOUNDATION

Riverside Drive, Chiromo, Nairobi, Kenya
E-Mail: info@scienceforafrica.foundation
T: +254 705 199 199
www.scienceforafrica.foundation