

News Release

India, Italy and Africa unite to bring voice AI to the world's most underserved communities

Trilateral AI Hub partnership showcases offline voice AI for agriculture and fintech at GITEX Kenya 2026, powered by European supercompute, India's digital public goods and African language innovation



Nairobi, Kenya | 20 May 202: A landmark [trilateral partnership between India, Kenya, and Italy](#) — anchored by the G7-endorsed AI Hub for Sustainable Development that is co-led by the Ministry of Enterprises and Made in Italy and the United Nations Development Programme — reached a major milestone this week at GITEX Kenya. Showcasing offline voice AI capabilities designed for agriculture and fintech use-cases in low-connectivity communities across Africa, it marked the most visible proof-point yet

of what the partners call an “AI Diffusion Pathway”: a replicable model for scaling AI adoption where it can make the most transformative impact.

AI diffusion for the next adopter: Orchestration, compute and language

Each partner contributes to creating AI diffusion pathways to solve the most complex challenges, demonstrating 30 percent tech and 70 percent non-tech focus of diffusion pathways through use-cases. The diffusion pathway offers the architecture — not by pushing more products into communities that lack the conditions to use them, but by building the conditions themselves. The language. The compute. The deployment infrastructure. The trust.

India: Orchestration and open-source Infrastructure

India provides the intelligence layer that binds the system together. EkStep, the open-source technology non-profit co-founded by Nandan Nilekani, contributes the VoicERA platform — a production-ready, open-source voice AI orchestration stack hosted on GitHub. VoicERA integrates the Pipecat real-time voice pipeline framework and a Use Case Adoption Framework that has already been deployed at scale across India’s public digital infrastructure. Critically, VoicERA’s codebase has now travelled to Africa — and Africa’s on-device quantised voice AI capabilities have in turn been shared back with EkStep’s Commercial Open Source Software (COSS) team, who had not previously come across voice AI at that level of compression running on a mobile phone.

Italy: Supercompute power through CINECA Leonardo

Italy provides what no cloud credit can replicate: sustained, uninterrupted access to Europe’s leading supercomputing infrastructure. CINECA Leonardo — the EuroHPC AI Factory facility located in Bologna — has been the training ground for the partnership’s African language voice models, including Crane AI Labs’ Luganda text-to-speech model with zero-shot voice cloning, trained entirely on CINECA’s A100 clusters. The team describes the quality difference starkly: “Cloud credits let you experiment. CINECA lets you build.” Training documentation produced on CINECA will be made available to future cohorts of Hub innovators requiring supercomputer access — lowering the barrier for the next generation of African AI builders.

Africa: Language Models, Edge Deployment, and Enterprise Roots

Africa provides the cultural intelligence, language models, and on-the-ground deployment capability that give the pathway its real-world meaning. Two African-built AI companies anchor the technical contribution. Crane AI Labs (Uganda) contributes an

entirely open-source, MIT-licensed, model-agnostic offline voice stack — hosted on HuggingFace — with Luganda automatic speech recognition (ASR) and text to speech (TTS) models trained on CINECA, alongside the Swahili-Gemma language model that powers reasoning in the demo pipeline. MsingiAI Kenya, founded by Gilbert Korir, has developed what is recognized as the best-published Swahili ASR and TTS globally through the Sauti models — and these Sauti models are the speech recognition and synthesis layer that powered the trilateral demonstration onstage today in GITEX Kenya. Enterprise pilots getting ready: Hello Tractor and Shamba Records are among the agriculture and agritech partners bringing voice AI to smallholder farmers across Africa, with use cases partnership under discussion with Equity Bank in Kenya and Rwanda.

The Africa Compute Fund: Building the continent's own AI infrastructure

While European supercompute provides critical training capacity today, the partnership's longer-term vision is captured in a simple phrase: “train on European supercompute today, deploy on African compute tomorrow.” The Africa Compute Fund is emerging as the African infrastructure bridge for that transition, connecting trained AI models to the compute, distribution, metering, and monetization rails needed for real-world deployment across the continent.

Through Monarch, its model hosting and distribution platform, Africa Compute Fund is building the layer through which African-language AI models can be catalogued, hosted, accessed, metered, and deployed by developers, enterprises, public institutions, and sector partners. Monarch has been in direct dialogue with African language AI builders — including MsingiAI Kenya, Crane AI Labs, technical institutions, research groups, and startups across the continent — on how models such as Sauti Swahili ASR and TTS, Luganda ASR/TTS systems, offline voice stacks, edge-optimized models, and other African-language AI systems can be hosted, distributed, monetized, and made accessible beyond their original training environments.

Monarch is designed for the economics of AI deployment in Africa: pay-per-inference access, prepaid usage flows, offline-capable payment infrastructure, usage tracking, and monetization pathways that allow model builders and application partners to earn from adoption at scale. The pathway is not complete when a model is trained. It becomes real when that model can be deployed, paid for, integrated, governed, and used at scale in African contexts.

Technical deployment alone, however, is not what makes African AI bankable. What the continent still needs — and what the next phase of this work will require — is a **continental infrastructure intelligence layer**: the connective architecture that links real-economy infrastructure, AI infrastructure, capital, policy, and project pipelines into structured, investable mandates. Without that layer, even the strongest models and best compute risk staying as pilots. With it, inference becomes procurement-ready, standards-aligned, and institutionally bankable. The partnership operates against a clear continental thesis in the [AI10](#) Billion, that no single capital instrument is adequate to the scale or diversity of AI infrastructure investment in Africa that will emerge in AI diffusion pathways responding to Africa's opportunities. The infrastructure, regulatory frameworks, data systems, and institutional capacities being assembled today will define that distinction for a generation.

The pathway playbook: Documentation for replication

A diffusion pathway is only useful if multiple other adopters can re-use, build on top, and scale. Alongside the technical, deployment, and use case work, the trilateral is producing a body of structured documentation designed to allow other African AI builders, government partners, and Global South ecosystems to replicate, adapt, and extend the model — not just consume it. The documentation suite under development includes:

- Use case documentation — agriculture advisory and financial services templates that any country, ministry, or implementing partner can adapt to local context, languages, and regulatory environment.
- African UX requirements — a living document capturing the channel patterns (WhatsApp, telephony, Unstructured Supplementary Service Data - USSD, on-device), offline-first design principles, code-switching behaviours, and cultural trust dynamics that distinguish African voice AI deployment from Indian or Western contexts, including a section on what Africa can teach India.
- CINECA training and fine-tuning playbook — step-by-step documentation of how African language models are trained on EuroHPC infrastructure, structured as a public resource for future cohorts of African AI builders requiring supercompute access.
- Sovereign deployment pathway — emerging documentation of how models trained on European supercompute move to African

sovereign compute, including hosting, metering, and monetization through partners such as the Africa Compute Fund.

- VoicERA integration patterns — documentation of how African speech models plug into the India-built VoicERA orchestration stack, enabling other African startups to adopt the framework with their own language models.

These documents are produced not for the partnership alone, but for the next country, the next pathway, and the broader goal of 100 AI Diffusion Pathways by 2030.

A coalition built for scale: 100 AI Diffusion Pathways by 2030

Voice AI adoption is central to the global goal of creating 100 AI Diffusion Pathways by 2030. The partnership's coalition of supporting partners includes EkStep, the Observer Research Foundation (ORF), the Gates Foundation, Anthropic, Qhala, the Harmonic Innovation Group and others — each contributing expertise, funding, or policy engagement to the shared goal of closing the AI impact gap through enabling adopters at the last mile. The reusable assets generated by the India-Africa-Italy pathway are already circulating: VoicERA's codebase is now deployed in Africa; Crane AI Labs' edge quantisation techniques are informing India's next deployment cycle; and CINECA training documentation is being structured as a public resource for future cohorts of the AI Hub for Sustainable Development. The partnership's longer-term ambition extends beyond technology delivery to ensuring that the continental architecture for converting AI infrastructure into investable, institutionally bankable opportunities is built alongside it.

What was demonstrated at GITEX Kenya

The GITEX Kenya demo showcased:

- Offline voice AI running entirely on-device in Swahili and Luganda — no internet connection required.
- Swahili speech recognition and synthesis powered by MsingiAI Kenya's Sauti ASR and TTS models — recognized among the best-published Swahili speech AI globally — integrated through the VoicERA orchestration stack to deliver a complete Swahili voice pipeline live on stage.
- Agriculture use-cases that enable smallholder farmers to access market prices, agronomic advice, and financial services through voice in their native language.

- Fintech use-cases integrating with prepaid inference and offline wallet Software Development Kits, enabling pay-per-use AI without cloud dependency.
- Zero-shot voice cloning in Luganda — the first such model trained on CINECA Leonardo for an African language.
- A live deployment of VoicERA, demonstrating the India-originated orchestration stack running on African edge hardware.