



AFRICA HEALTH TECH SUMMIT 2026 · KIGALI, RWANDA

AI Commons for Global Health Workshop

Read-out and key insights

This workshop convened contributors from across the AI for global health ecosystem to help shape the [AI Commons for Global Health](#).

The AI Commons for Global Health takes on the shared problems that no single organization is positioned to solve alone, building the common tools, standards and knowledge that AI needs to drive health impacts in low- and middle-income countries (LMICs). Its work is fueled by working groups, and four are active today.

The workshop opened with brief updates from each working group:

- **Data.** Health data that could make AI work in LMICs often exists but is hard to find, hard to get permission for, and risky to share. Rather than build another central repository, the group is making datasets discoverable and offering tiered contribution options, governance guidance and tooling that lower risk for data owners. Next, it is co-designing its [approach](#) with data owners.
- **Clinical decision support systems (CDSS).** LLM-enabled CDSS offers flexibility that was not possible with previous deterministic methods, but it needs new safety approaches, and innovators are learning the same lessons in isolation. The group has drafted, with innovators, a reference architecture of about 40 modular building blocks across eight layers, and is now testing whether there is demand for common blocks that can be built once and reused as a public good. There is an [innovator showcase](#) on October 8 sharing lessons learned from around the world.
- **AI evaluations.** Most evaluations are not designed for LMIC languages, clinical contexts or infrastructure. Hosted by the Open Health Stack Software Foundation (OHS-SF), the group is building the tooling that connects shared evaluation infrastructure to real products, from model selection through post-deployment monitoring. [Contributors welcome](#).
- **Procurement.** Many governments lack clear guidelines for buying AI, so selection often runs through networks and partnerships rather than evidence. The group is scoping work with WHO-SEARO, the Indian School of Business and Ashoka University to build a procurement playbook that will help health systems acquire best-fit AI solutions.

WHAT WE HEARD

Key insights

Across multiple breakout groups, workshop contributors gave input on the working group products and on how the AI Commons can best meet the needs of key stakeholders, from innovators and implementers to governments, clinicians and patients. Key insights are summarized below.

1. Evaluation and safety should be shared public infrastructure.

There was strong demand for clear guardrails and a standard set of evaluation metrics for each use case, adaptable to regional context. Innovators depend on academic and clinical partners for validation, which is slow; shared evaluation environments and better use of scarce domain experts' time could speed this up while keeping evaluation grounded in the setting where a tool will be used. Participants noted that metrics from engineers, clinicians and end users often disagree and need one framework; that evaluation must move beyond text to images, audio and multimodal inputs; and that evaluation of CDSS remains a specific gap. End users also need a short set of questions they can ask about any AI tool deployed to them, and safety guidance should be explicit about what can and cannot be given to a model, especially in sensitive areas such as mental health.

2. Procurement is where adoption gets decided, and the playbook has to meet governments where they are.

It is often unclear who the final decision-maker is. Ministries of health, ICT and finance, standards bureaus and even courts can shape or block adoption; a Kenyan High Court ruling against a teleradiology provider was one example. The playbook should start with AI basics for ministry staff, distinguish between solution types (a standalone chatbot is bought differently from AI embedded in lab equipment), and build on existing digital health and procurement playbooks rather than duplicate them. Open questions include how to price highly variable solutions, how to estimate total cost of ownership against efficiency savings, and how where buying happens (e.g., at national, state or pooled regional level) will shape the resource. Young innovators are an audience too, since knowing what governments will ask for shapes how they build.

3. Readiness and local context determine technology's success.

Several participants had seen deployments fail midway because infrastructure, connectivity or workforce was not in place; a CDSS that refers patients to a hospital seven hours away does not help. Governments want readiness assessments and help judging whether a use case needs AI at all, since deterministic tools may be safer and cheaper. Tools and playbooks built elsewhere need adaptation, and design must account for end users with low literacy, low incomes, disabilities or limited connectivity. Regulation adds friction: data-residency laws collide with cloud infrastructure, private companies are unsure how to share data legally, and software-as-a-medical-device rules are nascent and differ by country.

4. Data sharing depends on incentives and trust, not access alone.

Patient privacy and institutional responsibility were the top barriers to sharing, followed by difficulty getting permission, fear of losing control and a less-discussed one: researchers want to publish first. Participants asked for data-sharing agreement templates, governance boards and guidance on choosing licenses. A proposed data "shopping cart" drew a mixed response, since charging for access risks a paywall.

5. The Commons is a community, not only a technical resource.

Community was the most requested feature across groups: connecting data owners with collaborators, innovators with evaluators and clinical experts, and solutions with implementers and governments. Participants asked for local chapters, in-person learning and localized communities of practice, and suggested a dedicated end-user working group. They also want demand signals from innovators and governments to set Commons priorities, so that local needs shape the global agenda rather than the reverse.

6. Sustainability is unresolved and needs attention now.

Open-source tools, safety work and shared infrastructure need long-term financing. Participants pointed to approaches such as those used by OHS-SF and the Linux Foundation as ways to steward open tools over time. The Commons also needs a clear answer to why organizations would contribute valuable or proprietary data and tools. Funding and incentives need to be designed alongside the technical work, so that what the Commons produces stays maintained and useful.

GET INVOLVED

Next steps

- [Join the Commons](#). Sign up for the Community of Practice to engage with events and shape the future direction of the Commons. You can also use the form to indicate interest in joining specific working groups.
- [CDSS innovator showcase](#), **October 8**. Innovators share the design choices, successes and challenges behind their LLM-enabled CDSS.
- **Join the [AI Evals working group kickoff](#), October 20**. Hosted by OHS-SF, with fortnightly meetings to follow.
- **Procurement introductions**. The group is looking for national or state governments, or anyone planning to buy health AI at scale, who are about to make a decision. Willingness to try the playbook matters more than contract size. Indicate interest via the join form above.
- **Data co-design**. Data owners are invited to help co-design the Data Commons [site](#) and identify the barriers it should remove first. Sign up via the join form above.

The message from the room was consistent: the Commons will be most useful if it makes it easier to build on what already exists, in ways that are locally relevant, safe and sustainable.