Investment Committee Report: The $1 Billion Digital Health Challenge

Decision date: 2026-08-31
Prepared for: Investment Committee
Mandate tested: A proposed USD 1 billion commitment to improve health outcomes in low- and middle-income countries through digital health.

Investment decision

DECLINE THE MANDATE AS FRAMED.

I would not commit USD 1 billion to digital health on a short donor timetable. The public evidence does not establish that the sector can productively absorb capital at that rate, and the most visible structural problem is recurrent ownership: who pays for maintenance, hosting, security, support, upgrades, compliance, and integration after grant funding ends.

I recommend the following capital disposition:

Capital treatment	**Amount**	**Status**
Staged digital-health portfolio	**USD 250 million**	**IC DECISION**
Redirect to Gavi, the Vaccine Alliance	**USD 500 million**	**IC DECISION**
Remain donor-held and uncommitted	**USD 250 million**	**IC DECISION**
Total	**USD 1 billion**	**IC DECISION**

If the legal gift instrument requires that all funds be irrevocably transferred to this foundation for digital health, I would decline USD 750 million rather than create a spending obligation that the evidence cannot support.

The digital tranche is intentionally conservative. It concentrates on six named institutions and a small set of country facilities, standards, evidence, compliance, and workforce mechanisms. It funds no new standalone national platform. It contains no AI scale-up line.

Sourcing note and protocol

The challenge requires every number, name, date, and factual assertion to be sourceable, calculations to be tagged `MY ARITHMETIC`, estimates to be tagged `ESTIMATE`, and unsupported values to remain `UNKNOWN`. It also requires searched absence claims and explicit handling of conflicting sources.

This report follows that discipline. Empirical numbers are keyed to the **Numeric Evidence Ledger** in Appendix A. Calculations are reproduced in Appendix B. Absence searches are recorded in Appendix C. Proposed grant amounts, scores, thresholds, and portfolio rules are labeled **IC DECISION**, because they are recommendations rather than external facts.

One unavoidable format limitation applies to the prompt's request for the "exact sentence" from each web source. For copyright compliance, this report reproduces an exact supporting excerpt of no more than 25 words when the source sentence is longer.

Retrieval date for all web sources unless otherwise stated: 2026-08-31.

Dated source register

Primary, official, and peer-reviewed sources

ID	Source	Publication / last-edited date	Use in this report
S01	PATH Digital Square, *Shaping the market for digital health tools: A roadmap to a sustainable global goods ecosystem*	March 2025	Current approved-global-goods snapshot; sector financing, government ownership, impact measurement
S02	PATH, "Digital Square Announces New Global Goods to Advance Climate and	2025-07-07	Shows approvals continued after the March 2025 snapshot

ID	Source	Publication / last-edited date	Use in this report
	Health Data Integration"		
S03	Digital Square, *Financing digital health software global goods*	December 2024 document	Donor dependence and financing models
S04	HISP Centre, 2024 Annual Report	Report for 2024; release page 2025-06-05	HISP funding base and allocation
S05	HISP Centre, 2025 Annual Report release	2026	Current DHIS2 reach and 2025 expansion
S06	DHIS2, About page	Publication / last-edit date UNKNOWN; current site retrieved 2026-08-31	Stewardship, licensing, API, standards
S07	Medic Mobile Form 990 data, IRS-derived via ProPublica	Fiscal year 2024; filing 2025-11-10	Annual expenditure
S08	Medic, 2026-2028 strategy	2026 strategy; page date UNKNOWN	CHT reach and strategy, self-reported
S09	Community Health Toolkit interoperability documentation	Last updated 2025-11-11	FHIR workflow and OpenHIM mediation
S10	Medic leadership announcement	2026-02-02	Current governance leadership
S11	OpenMRS Inc Form 990 data, IRS-derived via ProPublica	Fiscal year ending June 2025; filing 2026-05-05	Annual expenditure

ID	Source	Publication / last-edited date	Use in this report
S12	OpenMRS community board announcement	2026-02-12	Board structure
S13	OpenMRS terms / contribution policy	Last-edit date UNKNOWN; current page retrieved 2026-08-31	Software license and code governance
S14	WHO SMART Guidelines publication table	Last updated November 2025	Published DAKs and FHIR implementation guides
S15	Jembi Health Systems, About page	Last-edit date UNKNOWN; current page retrieved 2026-08-31	Mission and organizational status
S16	Jembi, OpenHIM core repository	Current repository retrieved 2026-08-31	License, functions, retention behavior
S17	PATH, About Digital Square	Last-edit date UNKNOWN; current page retrieved 2026-08-31	Program role and country support
S18	Gavi, Resource mobilisation model	Last updated 2026-06-29	Outside option and modeled health target
S19	Kenya, Digital Health (Health Information Management Procedures) Regulations	Published and commenced 2025-04-11	Certification and health-data obligations
S20	Nigeria Data Protection Commission, NDP Act GAID	2025	DPIA, software privacy, breach requirements

ID	Source	Publication / last-edited date	Use in this report
S21	Rwanda Data Protection and Privacy Office, DPP Law pages	Law 058/2021 of 2021-10-13; current official pages retrieved 2026-08-31	Breach, storage, transfer, retention obligations
S22	Nigeria Federal Ministry of Health and Social Welfare	2026-07-30	Current National Digital Health Architecture position
S23	Rwanda Ministry of Health HMIS	Current system retrieved 2026-08-31	Current DHIS2-based HMIS presence
S24	BMC Medical Informatics and Decision Making, Rwanda EMR assessment	2026	OpenMRS deployment evidence, cross-sectional
S25	Kenya PPRA, Standard Tender Documents	Current page retrieved 2026-08-31	Procurement routes for IT, services, maintenance
S26	Rwanda Umucyo e-procurement legal resources	Current page retrieved 2026-08-31	Procurement framework
S27	FHIR.org, Conformance Testing	Current page retrieved 2026-08-31	Distinction between implementation and conformance testing
S28	FHIR Implementation Registry	Current page retrieved 2026-08-31	What public FHIR registries do contain
S29	IHE Connectathon Global Results	Current page retrieved 2026-08-31	Public event-specific interoperability testing results
S30	*Nature Medicine*, generative AI clinical decision-support trial in Kenya	2026-06-26	Prospective LLM clinical evidence and inference cost

ID	Source	Publication / last-edited date	Use in this report
S31	Principles of Donor Alignment for Digital Health	2018-10-24	Long-term operating-cost principle; **older than two years**
S32	Medic CHT core repository	Current repository retrieved 2026-08-31	AGPL license and code ownership
S33	OpenMRS current licensing terms	Current page retrieved 2026-08-31	MPL 2.0 with Healthcare Disclaimer
S34	OpenHIM current repository	Current repository retrieved 2026-08-31	MPL 2.0 and retention behavior

Secondary sources used only where direct public financial extraction was more accessible

ID	Source	Date	Limitation
F01	ProPublica Nonprofit Explorer, Medic Mobile	Retrieved 2026-08-31	Extracts IRS Form 990 data; IRS filing is the underlying record
F02	ProPublica Nonprofit Explorer, OpenMRS Inc	Retrieved 2026-08-31	Extracts IRS Form 990 data; IRS filing is the underlying record

No secondary commentary is used to establish the core investment thesis.

Part 0: Accept or decline

0.1 Two ceilings, not one

The public record does not permit a defensible, sector-wide absorption ceiling. That is itself a due-diligence result. Steward organizations publish uneven financial data; country

implementation spending is fragmented across ministries, donors, integrators, contractors, and vertical programs; and the public record rarely separates software maintenance from implementation, hosting, training, and broader health-program expenditure.

I therefore distinguish an evidence-based **observed steward run-rate proxy** from two **IC risk caps**. The risk caps are estimates, not claims about the sector's true capacity.

Software steward capacity

Three relatively observable software stewards provide the starting base:

- HISP Centre reported total 2024 funding of USD 20.174 million, with 74 percent allocated to the core DHIS2 platform and global resources. [S04; N01; N02]
- Medic Mobile reported 2024 expenses of USD 9.341904 million in its IRS-derived filing data. [S07; N03]
- OpenMRS Inc reported expenses of USD 1.346922 million for its fiscal year ending June 2025. [S11; N04]

`MY ARITHMETIC`: HISP core funding proxy = USD 20.174 million × 74 percent = **USD 14.928760 million**. [A01]

`MY ARITHMETIC`: observed three-steward base = USD 14.928760 million + USD 9.341904 million + USD 1.346922 million = **USD 25.617586 million**. [A02]

This is not sector expenditure. HISP's figure is funding, not audited annual expenditure; Medic and OpenMRS are expenses; and many Tier A goods are not represented. The base is a partial observable floor.

ESTIMATE, software steward year-one ceiling: USD 40 million per year. Range: USD 30 million to USD 55 million.

Method: start from the observable three-steward proxy and permit a substantial but gated increase for maintainers outside the observed sample and for additional staff, security, release engineering, documentation, and standards work. I assume **25 percent annual growth**, because additional core-software work can propagate across deployments without one-for-one field staffing, but rapid hiring still creates management and quality risk.

Country delivery capacity

Public empirical base: `UNKNOWN`.

I found no current, comparable public dataset that aggregates country-level digital-health implementation expenditure across ministries and integrators while separating one-time

deployment from recurrent operations. I therefore refuse to present a false empirical capacity figure.

For fiduciary control, I impose an internal cap:

ESTIMATE, country-delivery year-one risk cap: USD 120 million per year. Range: USD 75 million to USD 160 million.

This is a governance constraint, not a measured sector fact. I assume **15 percent annual growth**, lower than software stewardship, because country delivery requires contracting, ministry decisions, workflow redesign, training, integration, data-protection compliance, infrastructure, and recurring human support. The challenge itself correctly distinguishes software's low replication cost from delivery's higher marginal implementation cost.

Combined risk envelope

Year	Software cap, ESTIMATE	Country delivery cap, ESTIMATE	Combined cap, ESTIMATE
2027	USD 40.000m	USD 120.000m	USD 160.000m
2028	USD 50.000m	USD 138.000m	USD 188.000m
2029	USD 62.500m	USD 158.700m	USD 221.200m
2030	USD 78.125m	USD 182.505m	USD 260.630m
2031	USD 97.656m	USD 209.881m	USD 307.537m

All figures in this table are `ESTIMATE` or `MY ARITHMETIC`; see A03.

The main conclusion is stronger than any specific cap. **The sector lacks the public financial observability required to justify a billion-dollar rapid-disbursement mandate.**

0.2 Against the donor's rate

The donor's actual requested annual disbursement rate is `UNKNOWN`. "Short timeline" is not a rate.

I therefore run two stress tests rather than inventing intent.

- `MY ARITHMETIC`, five-year interpretation: USD 1 billion / 5 = **USD 200 million per year**. Against the estimated 2027 combined cap of USD 160 million, the first-year gap is **USD 40 million**. [A04]

- `MY ARITHMETIC`, three-year interpretation: USD 1 billion / 3 = **USD 333.333 million per year**. Against the estimated 2027 combined cap, the first-year gap is approximately **USD 173.333 million**. [A05]

What I would say to the donor

> I will not promise to spend a billion dollars on digital health simply because the capital is available. Public data do not establish that the sector can absorb that amount on a short timetable without buying duplication, unsupported expansion, or future liabilities that ministries cannot carry. I can defend a staged USD 250 million digital-health program with explicit recurrent-cost and evidence gates. I can defend moving USD 500 million to an established immunization mechanism with a stated health-outcome target. I would keep USD 250 million in your control until digital health can show stronger outcome evidence, audited recurrent ownership, and credible total cost of ownership. If that is unacceptable, I recommend you choose another investment manager.

0.3 The outside option

I would redirect half the capital outside digital health to **Gavi, the Vaccine Alliance**.

Gavi's current resource-mobilization page states that its 2026-2030 target budget is USD 11.9 billion and that the strategy is intended to protect 500 million children and save at least 8 million lives. [S18; N05; N06; N07]

The comparison metric is **target-budget dollars per stated modeled life saved**.

`MY ARITHMETIC`: USD 11.9 billion / 8 million stated modeled lives saved = **USD 1,487.50 per stated modeled life saved**. [A06]

This is not a marginal causal cost-effectiveness estimate. The Gavi budget finances activities beyond mortality reduction, and the stated lives-saved result is a prospective model. I do not linearly claim that a USD 500 million grant purchases a fixed number of lives.

For digital health, a comparable cost-per-life-saved figure is **`UNKNOWN`**. The absence of a comparable metric is not permission to create one.

Part 0 decision: decline a full digital allocation; accept only a staged digital tranche and redirect or defer the balance.

Part 1: Identify the binding constraint

Judgment

The hardest binding constraint is **government recurrent financing and purchasing capacity for digital public infrastructure and services after donor funding ends**.

This is narrower than saying "money is the problem." The constraint has two linked components:

1. a ministry or other legitimate domestic payer needs a durable recurrent budget line, and
2. that payer needs a legal and operational procurement mechanism to buy hosting, maintenance, security, integration, support, upgrades, and compliance around software that may itself have no license fee.

Digital Square's March 2025 sector review describes difficulty moving from donor-funded projects into government-supported solutions and explicitly calls for government budget allocations as part of sustainability. [S01] Its financing review states that donors provide nearly all funding for digital-health software global goods. [S03]

This diagnosis matters because grant capital can worsen the problem. A donor can finance a deployment faster than a ministry can create a recurrent appropriation, procurement category, service contract, security function, or support team. The result is a technically successful deployment with a financial expiry date.

Strongest counter-case: evidence, not finance, may bind harder

The strongest argument against my choice is that the sector is not yet investment-ready because direct evidence of improved health outcomes is too thin. Deployment, interoperability, coverage, data completeness, and clinician workflow can improve without changing morbidity or mortality. A billion-dollar fund could therefore solve the financing problem for tools whose health benefit remains uncertain.

That counter-case is strong. Digital Square itself says the industry lacks a standard method for measuring digital-intervention impact and calls for clearer cost-effectiveness and health-outcome measurement. [S01] My evidence classification below places most of the named digital investments at Tier 3 or worse.

I still choose recurrent financing as the binding constraint because even a clinically useful digital system fails if no legitimate entity can pay its permanent operating costs. Evidence can tell a

government that a platform is valuable; it does not create a budget line or procurement mechanism.

Evidence that would change my mind

I would commission a **prospective multi-country quasi-experimental study** of mature national digital systems that have similar platform maturity but materially different domestic recurrent financing.

Pre-specified exposure variable: the **audited share of total recurrent digital-system cost paid from domestic government budgets or binding domestic service contracts after donor scale-up ends**.

Pre-specified outcomes: system uptime, supported-version compliance, security-patch latency, active clinical use, completeness of nationally required data flows, and persistence of use after external project financing ends.

I would change my view if systems with high domestic recurrent ownership showed **no meaningful improvement in persistence or operational quality** relative to otherwise comparable systems with low domestic recurrent ownership, while workflow fit, interoperability, or outcome evidence explained the difference. That result would move the binding constraint away from fiscal ownership.

Rival hypotheses

Fiscal, not organizational. I largely agree. NGO or implementer absorption is secondary when the eventual payer cannot carry recurrent cost.

The sector is not investment-ready. Serious counterargument, and partly true. My response is to reduce the digital tranche, require independent evidence generation, and refuse to pretend that deployment scale proves health impact.

Consolidation beats diversification. Mostly correct. I fund maintainers and country consolidation, not a broad venture-style portfolio. I reject an unconditional permanent endowment because it would lock future boards into today's architecture without sufficient evidence or removal discipline.

The 2025 USAID shock changes the marginal dollar. It increases the value of continuity capital for already-deployed systems, but only where a ministry identifies the system as essential and commits to a recurrent-cost transition. Rescue funding without transition conditions simply extends donor dependence.

Prompt-induced bias

The challenge supplies the rival hypotheses, which can bias the diagnosis. I would still have reached the recurrent-financing conclusion without the list because the most recent Digital Square sector work independently identifies sustainability, government ownership, budget integration, financing, procurement, and impact measurement as current problems. [S01; S03]

One omitted constraint belongs on the list: **donor and implementer incentives that reward bespoke vertical systems, project-specific reporting, and short grant cycles even when country architecture calls for shared infrastructure**. That incentive can create fragmentation upstream of any technical choice.

Part 2: Evidence classification

The evidence ladder in the challenge is applied to the **core product**, not merely to an intervention delivered through the product.

Assessed entity / product	**Evidence tier**	**Basis**
HISP / DHIS2	**Tier 3**	Strong deployment and observational/process evidence; I did not identify product-attributable randomized health-outcome evidence
Medic / Community Health Toolkit	**Tier 4**	Large deployment scale and process claims; platform attribution for health outcomes remains weak
OpenMRS Inc / OpenMRS	**Tier 3**	Observational deployment evidence, including current Rwanda assessment; no platform-attributable health-outcome causal estimate established here

Assessed entity / product	Evidence tier	Basis
WHO SMART Guidelines / DAK translation layer	**Tier 5**	WHO normative authority and guideline provenance do not establish that the digital translation layer itself improves health outcomes
Jembi / OpenHIM	**Tier 5**	Global-good / interoperability status and technical function; direct causal health-outcome evidence for middleware not established
PATH Digital Square	**Tier 5**	Peer-review and intermediary role; its committee approval process is not evidence that the intermediary itself causes health outcomes

The Rwanda EMR study is cross-sectional. It assessed 257 facilities and documented OpenMRS use and barriers, so it supports deployment and observational classification, not causal health outcomes. [S24; N08]

Dollar-weighted evidence exposure

Within the USD 250 million digital tranche:

- Tier 4 or Tier 5 investments: **USD 149 million, IC DECISION classification**.
- Tier 3 investments: **USD 31 million, IC DECISION classification**.
- Enabling lines that do not map cleanly to a software evidence tier: **USD 70 million, IC DECISION classification**.

`MY ARITHMETIC`: USD 149 million / USD 250 million = **59.6 percent** of the digital portfolio at Tier 4 or below. [A07]

Among only the USD 180 million to which I assign a product/vehicle evidence tier:

`MY ARITHMETIC`: USD 149 million / USD 180 million = **82.78 percent** at Tier 4 or below. [A08]

That is uncomfortable and should be. The defense is limited: much of the money is financing shared infrastructure, interoperability, standards, consolidation, and country operating capability. Those are enabling assets for which direct mortality attribution is difficult. I do not therefore infer

health impact. I attach stop rules, recurring-cost gates, and a dedicated independent-evidence line.

Temptation log

I was tempted to treat each of the following as impact evidence and rejected the shortcut:

- DHIS2 use across 133 countries. [S05; N09] Scale is not causal health evidence.
- CHT's self-reported 182,676 equipped health workers across 24 countries. [S08; N10; N11] Reach is not product-attributable impact.
- OpenMRS use in Rwandan facilities. [S24] Adoption is not improved health.
- WHO ownership of SMART Guidelines. Normative authority strengthens clinical provenance but does not establish the effect of the digital representation layer.
- OpenHIM's recognition as a global good. Peer approval and technical function do not establish patient outcomes.
- Digital Square's donor and peer-review role. Institutional endorsement is not causal evidence.

No deployment count, global-good status, or donor endorsement is treated as proof of health impact.

Part 3: Organization assessment

I assessed **six organizations in depth**. That is deliberate. The public financial, governance, funder-concentration, and product-level outcome data are too uneven to make a superficial scorecard of several dozen organizations defensible.

For scoring, **10 on substitutability means hardest to replace and greatest sector loss if the organization disappears.** Blank scores are intentional where evidence is insufficient.

3.1 HISP Centre / DHIS2

Mission and problem. The HISP Centre at the University of Oslo leads development and maintenance of DHIS2 and coordinates the HISP network. [S04; S06] DHIS2 addresses health-information collection, management, analysis, and interoperability.

Geography and scale. HISP reported that DHIS2 expanded to 133 countries in 2025 and that seven new countries launched national DHIS2 health systems that year. [S05; N09; N12] These are deployment facts, not outcome evidence.

License and code ownership. DHIS2 is free and open-source under the BSD 3-clause license, with an open API. [S06]

Governance. Core stewardship sits at the HISP Centre, with a global network of local HISP groups. [S06]

Annual expenditure. UNKNOWN. HISP's public 2024 report gives funding, not audited annual expenditure. I therefore use the core-funding share only as a capacity proxy, not as an expenditure claim. [S04]

Dimension	Score	One-line justification
Outcome evidence	4/10	Tier 3: strong operational adoption and observational evidence, weak product-attributable outcome causality
Government embeddedness		National use is extensive, but I did not verify strategy + enterprise architecture + recurrent budget line for the portfolio countries as one complete test
Institutional durability		Current largest-funder concentration and unrestricted runway are not public enough to score
Marginal absorptive capacity		No public evidence shows what the next unrestricted dollar buys this year
Substitutability	9/10	Core stewardship and ecosystem are difficult to replace quickly

3.2 Medic / Community Health Toolkit

Mission and problem. Medic is steward of the Community Health Toolkit, an open-source framework for community and frontline health applications. [S08; S32]

Geography and scale. Medic reports 182,676 equipped health workers across 24 countries in its current strategy. [S08; N10; N11] The figures are self-reported.

License and code ownership. CHT core is licensed AGPL-3.0; the repository lists copyright to Medic Mobile. [S32]

Governance. Shreya Bhatt and Andra Blaj became Co-Executive Directors in February 2026. [S10] Organizational board information is public, but this assessment does not infer funding durability from governance form alone.

Annual expenditure. USD 9.341904 million for fiscal year 2024, from IRS-derived Form 990 data. [S07; N03]

Dimension	Score	One-line justification
Outcome evidence	3/10	Tier 4: reach and process evidence are stronger than product-attributable health outcomes
Government embeddedness		Kenya national use is material, but a verified recurrent government budget line for the platform is not established here
Institutional durability		Largest-funder concentration and unrestricted funding runway are not sufficiently public
Marginal absorptive capacity		No public marginal-dollar schedule
Substitutability	6/10	Important community-health platform, but substitutes and alternative mobile architectures exist

3.3 OpenMRS Inc / OpenMRS

Mission and problem. OpenMRS supports an open-source electronic medical record community designed for resource-constrained settings. [S12; S13]

Geography and deployment. A 2026 Rwanda cross-sectional study found broad OpenMRS use in district hospitals and health centers, while also documenting non-use and operational barriers. [S24]

License and code ownership. OpenMRS distributes software under Mozilla Public License 2.0 with Healthcare Disclaimer. [S13; S33]

Governance. The 2026 community election announcement states that the board has five active seats and one community-representative seat elected every two years. [S12; N13; N14]

Annual expenditure. USD 1.346922 million for fiscal year ending June 2025. [S11; N04]

Dimension	Score	One-line justification
Outcome evidence	4/10	Tier 3: observational deployment evidence, no core-platform causal health-outcome estimate established
Government embeddedness		Rwanda deployment is visible, but complete strategy + architecture + recurrent budget evidence is not
Institutional durability		Funding concentration and minimum sustainable core run-rate are not transparent enough
Marginal absorptive capacity		No defensible public marginal-dollar estimate
Substitutability	7/10	Large installed base and community make abrupt replacement costly, though other EMR platforms exist

3.4 WHO SMART Guidelines

Mission and problem. WHO SMART Guidelines convert normative clinical and public-health guidance into reusable digital components, including Digital Adaptation Kits and, for some domains, FHIR implementation guides. [S14]

Geography and scale. They are normative global artifacts rather than a single deployed national application.

License and code ownership. Product-specific licensing varies by artifact and repository; I do not assign one blanket license without source-level verification.

Governance. WHO is the normative publisher. That institutional status does not substitute for outcome evidence.

Annual expenditure. UNKNOWN. WHO program-specific expenditure for SMART Guidelines was not established in the public sources reviewed. I do not substitute WHO's parent budget.

Dimension	Score	One-line justification
Outcome evidence	2/10	Tier 5 for the translation layer: guideline authority is not causal evidence for the digital artifact
Government embeddedness		Not scored as a national operational system
Institutional durability		Program-specific funding durability not public enough
Marginal absorptive capacity		No public marginal-cost schedule for publishing additional artifacts
Substitutability	9/10	No other body has equivalent global normative authority for WHO guideline representation

3.5 Jembi Health Systems / OpenHIM

Mission and problem. Jembi describes itself as an African nonprofit improving global health through information systems, partnerships, and local capacity. [S15] OpenHIM is middleware for authentication, routing, orchestration, logging, and exchange between systems. [S16]

Geography and scale. Jembi works with governments and funders in multiple African settings. I do not use its own deployment claims as outcome evidence.

License and code ownership. OpenHIM core is licensed MPL-2.0 and its repository is maintained under Jembi's organization. [S16; S34]

Governance. Jembi registered as a nonprofit company in 2009. [S15; N15] Current detailed board and funding concentration data were not sufficiently transparent for a durability score.

Annual expenditure. UNKNOWN. I found a reference to financial reporting but did not find a current public numeric annual expenditure that I could verify under this protocol.

Dimension	Score	One-line justification
Outcome evidence	2/10	Tier 5: technical and global-good status without product-attributable health outcomes
Government embeddedness		Country integrations exist, but the complete budget-line test is not verified
Institutional durability		Current unrestricted runway and largest funder share are not public
Marginal absorptive capacity		No defensible public estimate
Substitutability	6/10	Important interoperability component, but middleware can be replaced at substantial integration cost

3.6 PATH Digital Square

Mission and problem. Digital Square is a PATH-led initiative that coordinates donors, governments, technical experts, and global goods; it supports alignment, procurement, global goods, and country capacity. [S17]

Geography and deployment. The initiative works across global and country digital-health systems rather than operating one clinical product.

License and code ownership. Not applicable as a single software product.

Governance. PATH leads the initiative; its global-goods process uses peer review and a maturity model. [S01; S17]

Annual expenditure. UNKNOWN. PATH publishes parent financial statements, but I did not find a current Digital Square program-specific annual expenditure suitable for this assessment. I do not substitute PATH's total budget.

Dimension	Score	One-line justification
Outcome evidence	2/10	Tier 5: intermediary and peer-review role, not a causal patient-outcome intervention
Government embeddedness		Not a national production system
Institutional durability		Program-specific funder concentration and post-USAID funding base are not transparent enough
Marginal absorptive capacity		No public marginal-dollar schedule
Substitutability	5/10	Coordination role matters, but other multilaterals, donors, and technical-assistance mechanisms can perform parts of it

Blank-score count

There are **18 blank scores out of 30 possible dimension scores, IC ASSESSMENT**. The blank rate is a finding: key investability data are not sufficiently public for confident scoring.

Part 4: The recurrent cost cliff

The recurrent-cost problem is the core reason I will not authorize a rapid scale portfolio.

The public record is inadequate for a credible portfolio-wide year-six operating-cost total. Hosting, security operations, patching, help desk, configuration, integration, connectivity, device replacement, clinical-content updates, staff retention, certification, legal compliance, and

supplier support depend on country architecture and procurement decisions that have not yet been made.

Portfolio year-six recurrent cost: `UNKNOWN`.

That is not a bookkeeping omission. It is a stop condition. No country deployment receives scale capital until its total cost of ownership is costed and a recurrent payer is named in a binding government instrument or service contract.

Investment line	Initial capital, IC DECISION	Year-six recurrent cost	Intended payer after grant	Evidence of payer
HISP / DHIS2 core stewardship	USD 25m	UNKNOWN	Diversified core contributors plus country / donor service financing	Current HISP financing model shows donor dependence; future payer commitments not yet signed
Medic / CHT core stewardship	USD 12m	UNKNOWN	Mixed steward financing plus country service contracts	UNKNOWN for future portfolio deployments
OpenMRS Inc core stewardship	USD 6m	UNKNOWN	Community, institutional grants, implementation/ service ecosystem	UNKNOWN for future portfolio deployments
WHO SMART Guidelines	USD 10m	UNKNOWN	WHO / member-state / donor program finance	Program-specific recurrent payer UNKNOWN
Jembi / OpenHIM interoperability	USD 17m	UNKNOWN	Government / integrator service and maintenance contracts	Country-specific contracts not yet signed

Investment line	Initial capital, IC DECISION	Year-six recurrent cost	Intended payer after grant	Evidence of payer
PATH Digital Square / TCO and procurement TA	USD 10m	Time-bounded TA; residual maintenance UNKNOWN	Successor technical-assistance finance or governments	UNKNOWN
Kenya consolidation facility	USD 35m	UNKNOWN	Kenya public budget and contracted service providers	Must be documented before scale release
Nigeria consolidation facility	USD 35m	UNKNOWN	Nigeria federal/state public budgets and contracted service providers	Must be documented before scale release
Rwanda consolidation facility	USD 30m	UNKNOWN	Rwanda public budget and contracted service providers	Must be documented before scale release
Independent evidence and health economics	USD 20m	No production service may depend on this line	Research sponsor / repository budget for data stewardship only	Costed in each research contract
Data protection and cybersecurity facility	USD 15m	UNKNOWN	Country digital-health operating budget	Must be incorporated into country TCO
Workforce and local implementer capacity	USD 15m	Refresher and supervision costs UNKNOWN	Ministries, employers, service providers	Must be incorporated into workforce plans

Investment line	Initial capital, IC DECISION	Year-six recurrent cost	Intended payer after grant	Evidence of payer
Digital continuity reserve	USD 20m	Depends on triggered asset	Temporary bridge only	Release requires named successor payer

The year-six total remains UNKNOWN until country workplans, actual bids, service-level requirements, hosting choices, retention schedules, security requirements, and staffing models are priced.

Data-protection compliance line

The portfolio contains **USD 15 million, IC DECISION**, reserved for data protection and cybersecurity across Kenya, Nigeria, and Rwanda. For planning, the board may earmark up to **USD 5 million per country, IC DECISION**, but these are spending envelopes, not estimated compliance costs.

The actual compliance cost in each country is UNKNOWN.

Required cost categories include:

- data-protection impact assessments and legal review;
- controller / processor registration where required;
- privacy notices and consent-management changes;
- data-localization or authorized cross-border hosting;
- encryption, access control, logging, monitoring, and incident response;
- data-retention configuration and destruction;
- breach detection, investigation, regulatory notification, and data-subject communications;
- digital-health certification and recertification where applicable;
- cybersecurity assessments and remediation;
- records of processing, processor agreements, staff training, and data-protection officer functions.

These costs must be included in the production TCO rather than treated as a one-time compliance grant.

Donor-alignment test

The Principles of Donor Alignment for Digital Health say donors should determine and quantify long-term costs of operating, maintaining, and supporting digital-health systems. [S31] That

source dates to 2018 and is **more than two years old**, as required to be disclosed by this challenge.

This portfolio is **conditionally compliant**, with one deliberate departure.

It complies because no production scale-up can pass the funding gate without a quantified TCO and an identified recurrent payer. It departs from the principle where I fund core steward resilience, standards, evidence, and pre-deployment technical assistance before country-specific year-six costs are known. I record that departure rather than calling an uncosted portfolio sustainable.

If a country facility reaches deployment without an identified recurrent payer, the board has not built sustainable digital infrastructure. It has built **donor-financed temporary infrastructure with an expiry risk**.

Part 5: Architecture, standards, law and governance

5.1 Interoperability

The portfolio requires interoperability at the interface level, not marketing-level "FHIR compliance."

DHIS2 states that it supports FHIR compatibility and provides an open API. [S06] CHT's current documentation describes a FHIR workflow in which a mediator can convert CHT records into FHIR resources and use OpenHIM in the exchange path. [S09] OpenHIM provides routing, authentication, logging, and orchestration functions. [S16]

These are technical capabilities. They are not evidence that a specific national deployment conforms to the implementation guide that its counterpart systems expect.

FHIR.org distinguishes resource validation from server behavior testing; its conformance page describes server testing using known inputs and evaluating server responses. [S27] The FHIR Implementation Registry contains implementation guides, profiles, and voluntarily listed implementations. [S28] IHE publishes results from Connectathon interoperability testing events. [S29]

Finding: I did not find a universal, authoritative public register showing successful FHIR conformance results for LMIC national digital-health deployments.

That absence claim is supported by the search log in Appendix C.

IC requirement: every funded production interface that claims FHIR interoperability must:

1. identify the target FHIR version and implementation guide;
2. publish its CapabilityStatement and required profiles;
3. pass an independent test suite against the agreed implementation guide;
4. publish the machine-readable test artifacts and results unless a specific security reason prevents publication;
5. demonstrate end-to-end exchange with the actual national counterpart, not merely validate isolated resources.

A press release stating "FHIR compliant" does not pass the investment gate.

5.2 WHO SMART normative alignment

The WHO SMART publication table was last updated in November 2025. [S14] The table shows that the availability of Level 2 Digital Adaptation Kits and Level 3 FHIR implementation guides varies by health domain.

Examples from the current table:

- Antenatal care has a Level 2 DAK and a Level 3 FHIR implementation guide. [S14]
- HIV has a Level 2 DAK and a Level 3 FHIR implementation guide. [S14]
- Postnatal care has a Level 2 DAK but its Level 3 column is listed as not available. [S14]
- Tuberculosis has a Level 2 DAK but its Level 3 column is listed as not available. [S14]

I will not assign a portfolio-wide "percent SMART aligned" before country workplans define clinical domains and interfaces.

SMART-aligned share of portfolio: `UNKNOWN` ex ante.

IC rule: if WHO has published an applicable Level 2 or Level 3 artifact, grantees must reuse it unless the ministry documents a clinical or regulatory reason to diverge. If no relevant WHO artifact exists, the project must label any newly encoded clinical logic as locally governed content rather than imply WHO digital endorsement.

This reduces duplicate clinical logic without pretending that WHO SMART coverage is complete.

5.3 Data protection in three target countries

Kenya

Applicable instruments include Kenya's Data Protection Act and the Digital Health Act regulations. The 2025 Health Information Management Procedures Regulations were published and commenced on 11 April 2025. [S19; N16]

The regulations establish certification requirements for digital-health solutions and define health-information security and governance obligations. Compliance therefore creates costs in privacy assessment, cybersecurity assessment, policies, certification, operational security, retention, incident response, and system changes.

End-to-end Kenya compliance cost: UNKNOWN.

Nigeria

Nigeria's Data Protection Commission issued the NDP Act General Application and Implementation Directive in 2025. [S20] The directive calls for DPIAs when required, privacy-by-design measures for software, and breach notification to the Commission within 72 hours of awareness. [S20; N17]

End-to-end Nigeria compliance cost: UNKNOWN.

Rwanda

Rwanda's Data Protection and Privacy Office publishes the operative obligations under Law No. 058/2021 of 13 October 2021. This law is **more than two years old**, but the regulator's current site continues to publish the obligations. [S21; N18]

The official material requires notification of a personal-data breach within 48 hours and a fuller report no later than 72 hours. [S21; N19; N20] It also states that personal data are stored in Rwanda unless authorized otherwise and imposes conditions for cross-border transfer. [S21]

End-to-end Rwanda compliance cost: UNKNOWN.

Compliance-cost searches

I searched official regulator and government domains for standardized full compliance costs in all three countries. I found statutory and process requirements, but no official end-to-end figure that can be applied to this portfolio. Search terms and sites are in Appendix C.

The USD 15 million compliance line is therefore a **budget envelope**, not an estimate of legal cost.

5.4 Procurement

A ministry does not need to "buy open source software" as a license product in order to spend public money on it. The procurement object can be the service around the software.

Kenya's PPRA currently publishes standard tender documents for requests for proposals, non-consulting services, information technology, and maintenance services. [S25]

Rwanda operates the Umucyo e-procurement system and publishes the governing procurement instruments through it. [S26]

Nigeria's public procurement framework also provides tendering and consultancy mechanisms, although this report does not claim that every digital-health service can be purchased through one standard form without country counsel.

Practical procurement objects include:

- configuration and localization;
- implementation and integration;
- hosting and managed infrastructure;
- L2 and L3 support;
- security monitoring;
- version upgrades and patching;
- data migration;
- training and change management;
- interoperability services;
- data-protection compliance;
- performance monitoring;
- source-code stewardship under a service contract.

The government pays a vendor, nonprofit, university, local systems integrator, or managed-service provider for deliverables and service levels. The software license can remain free.

Finding: procurement is legally feasible in principle. The harder issue is getting recurring public appropriations and service contracts aligned with an open-source operating model. I would still require country counsel to confirm the specific procurement path before grant-funded implementation becomes a production dependency.

5.5 Endowment governance

I do **not** recommend a permanent maintenance endowment for any software steward at this stage.

If the board later creates one, its governance must include:

- an independent fiduciary holder rather than unrestricted permanent transfer to the steward;
- a published annual minimum performance package covering supported releases, security remediation, public documentation, contributor health, and interoperability;
- a right to suspend distributions if the steward abandons the supported codebase or fails security obligations;
- a right to redirect the corpus if the platform becomes obsolete, is superseded by a better open standard, loses material public use, or becomes governed in a way incompatible with the fund's mission;
- periodic external technical and governance review;
- a conflict-of-interest policy that prevents current maintainers from controlling withdrawal decisions.

The endowment must have an observable exit trigger. One workable trigger is failure to publish and support a production release within the grant-defined release policy for two consecutive review periods, after a cure process. That is an **IC governance rule**, not an external fact.

The point is to preserve adaptation. An endowment without a removal mechanism would freeze a 2026 architecture into a permanent capital commitment.

Part 6: AI-specific investments

Portfolio position

There is **no AI clinical scale-up line** in the proposed portfolio.

A maximum of **USD 5 million, IC DECISION**, within the USD 20 million independent evidence line may support replication, implementation research, safety evaluation, and economic evaluation of AI-assisted frontline care. It may not finance autonomous clinical deployment at scale.

Prospective evidence

A major piece of relevant evidence now exists.

A *Nature Medicine* pragmatic cluster-randomized trial published on 26 June 2026 evaluated generative-AI-enabled clinical decision support in primary care in Kenya. [S30; N21] The trial enrolled 9,691 patients overseen by 103 clinical officers. [S30; N22; N23]

The primary outcome was treatment failure. Treatment failure occurred in 2.2 percent of the intervention group and 2.0 percent of the control group; the adjusted odds ratio was 0.77, the 95 percent confidence interval was 0.55 to 1.08, and P = 0.13. The primary outcome did not differ significantly between groups. [S30; N24-N29]

The trial therefore does **not** establish improved patient outcomes on its primary endpoint.

It is important evidence precisely because it prevents the portfolio from making a generic "AI can improve frontline care" claim.

Liability and clinical governance

For any funded evaluation:

- the licensed clinician remains responsible for the clinical decision made at the encounter;
- the implementing health service remains responsible for credentialing, workflow, incident management, and clinical governance;
- the software provider remains contractually responsible for system performance, security, logging, version control, and disclosed failure modes;
- the study sponsor remains responsible for protocol compliance and adverse-event reporting;
- the relevant health and data regulators retain jurisdiction according to the software's intended use, deployment structure, and data processing.

Kenya's digital-health certification framework clearly applies to digital-health solutions used in healthcare. [S19] I would require local regulatory counsel to determine whether a particular AI tool also falls within medical-device regulation before deployment. I do not infer a regulator's jurisdiction from marketing language.

Every model suggestion must be logged with model/version metadata, source context where applicable, clinician action, and override. High-severity discrepancies must go through a clinical safety review.

Inference cost

The Kenya trial reported mean per-patient LLM inference cost of USD 0.04. [S30; N30]

That figure is **inference cost only**. It does not include the full cost of the clinical application, integration, devices, connectivity, data governance, clinician time, quality assurance, monitoring, support, or model-management overhead.

Year-six total AI-system cost per encounter: `UNKNOWN`.

Any funded evaluation must report both inference cost and full economic cost.

Outperforming an underqualified provider while underperforming a guideline

I would not deploy that system autonomously.

A weak human comparator is not the clinical standard. The operational question is whether the tool moves actual care closer to an accepted guideline without creating systematic unsafe deviations, discrimination, automation bias, or accountability gaps.

An assistive deployment could be justified if prospective evidence shows that the combined clinician-plus-tool workflow improves clinically meaningful outcomes or guideline concordance with acceptable safety. A model that reliably underperforms the guideline cannot become the autonomous standard merely because the baseline workforce is weak.

Why fund any AI work after a null primary outcome?

Because one high-quality null trial changes the investment case from scale to learning. Replication may identify subgroups, workflows, model designs, or implementation conditions where benefit exists; it may also confirm that benefit is absent. A small, preregistered evidence allocation buys decision-relevant knowledge. A large deployment grant would confuse technological novelty with established effectiveness.

The donor's wealth origin plays no role in this allocation.

Part 7: Portfolio construction and deployment

7.1 Capital disposition

The full donor proposition is treated as follows:

Destination	Category	Amount	Share of original USD 1b	Horizon	Evidence basis	Confidence
Digital-health portfolio	Staged digital investment	USD 250m	25%	2027-2031	Mixed, mostly Tier 3-5	Medium-low
Gavi	Non-digital global-health outside option	USD 500m	50%	Gavi 2026-2030 cycle	Established immunization mechanism with stated modeled outcome target	Medium-high
Donor-held uncommitted capital	Deferral / reserve outside foundation	USD 250m	25%	Revisit after evidence gates	No forced deployment	High confidence in prudence
Total		**USD 1b**	**100%**			

All amounts and shares in this table are **IC DECISION** or `MY ARITHMETIC`, not externally sourced investment facts.

7.2 Digital portfolio

Line	Category	Allocation, IC DECISION	Share of digital tranche, MY ARITHMETIC	Horizon	Evidence tier	Recurrent-cost payer	Confidence
HISP Centre / DHIS2	Core stewardship	USD 25m	10.0%	5 years	3	Diversified core + country/service finance, to be contracted	Medium
Medic / CHT	Core stewardship	USD 12m	4.8%	5 years	4	Steward finance + country service contracts	Medium-low
OpenMRS Inc	Core stewardship	USD 6m	2.4%	5 years	3	Community/institutional + country service ecosystem	Medium-low
WHO SMART Guidelines	Standards / normative content	USD 10m	4.0%	5 years	5	WHO/member-state/donor program finance	Medium
Jembi / OpenHIM	Interoperability	USD 17m	6.8%	5 years	5	Country/integrator service contracts	Low

Line	Category	Allocation, IC DECISION	Share of digital tranche, MY ARITHMETIC	Horizon	Evidence tier	Recurrent-cost payer	Confidence
PATH Digital Square	TCO, procurement, coordination TA	USD 10m	4.0%	4 years	5	Time-bounded TA; no permanent production dependency allowed	Medium-low
Kenya consolidation facility	Country delivery	USD 35m	14.0%	5 years	4, vehicle-level	Kenya public budgets / service contracts	Medium
Nigeria consolidation facility	Country delivery	USD 35m	14.0%	5 years	4, vehicle-level	Federal/state public budgets / service contracts	Medium
Rwanda consolidation facility	Country delivery	USD 30m	12.0%	5 years	4, vehicle-level	Rwanda public budgets / service contracts	Medium
Independent evidence and health	Evidence	USD 20m	8.0%	5 years	Not product-tiered	Research grants; no production	High

Line	Category	Allocation, IC DECISION	Share of digital tranche, MY ARITHMETIC	Horizon	Evidence tier	Recurrent-cost payer	Confidence
economics RFP						dependency	
Data protection and cybersecurity facility	Governance/compliance	USD 15m	6.0%	5 years	Not product-tiered	Country recurrent digital-health budgets after grant	Medium
Workforce and local implementer capacity RFP	Workforce	USD 15m	6.0%	5 years	Not product-tiered	Ministries/employers/service providers	Medium
Digital continuity reserve	Rescue/transition reserve	USD 20m	8.0%	5 years	Not product-tiered	Temporary bridge to named successor payer	Low-medium
Total		**USD 250m**	**100%**				

Why these countries

Kenya, Nigeria, and Rwanda are selected because each has active national digital-health architecture or production systems that make consolidation and interoperability more defensible than greenfield software creation.

Nigeria's Federal Ministry of Health stated in July 2026 that the National Digital Health Architecture had been endorsed by all states and the Federal Capital Territory and was intended to support an interoperable "One Patient, One Health Record" approach. [S22]

Rwanda operates a national HMIS on a DHIS2 endpoint, and current peer-reviewed research documents OpenMRS use across Rwandan facilities. [S23; S24]

Kenya has current digital-health regulation that makes certification, data governance, and system compliance concrete operating requirements. [S19]

These are not claims that the three countries have the highest return in the world. They are places where an architecture-first consolidation facility can attach to real public institutions and existing systems.

7.3 Disbursement schedule

The digital-health flow is deliberately below the estimated combined risk caps.

Year	**Digital disbursement, IC DECISION**	**Estimated combined cap**	**Headroom, MY ARITHMETIC**
2027	USD 35m	USD 160.000m	USD 125.000m
2028	USD 45m	USD 188.000m	USD 143.000m
2029	USD 55m	USD 221.200m	USD 166.200m
2030	USD 60m	USD 260.630m	USD 200.630m
2031	USD 55m	USD 307.537m	USD 252.537m
Total	**USD 250m**		

No year exceeds the fund's own risk cap.

The schedule is intentionally back-loaded relative to due diligence but modest relative to the estimated envelope. Early releases pay for TCO, evidence, standards, security, procurement design, and steward stability. Country scale comes only after recurrent payer and architecture gates.

7.4 Milestone-based release and stop conditions

Each line has one pre-specified observable that an independent party can verify.

Investment line	Release condition	Stop-funding observable
HISP / DHIS2	Agreed release, security, documentation, and sustainability plan	**Stop if independent release/security audit records failure of the grant-defined supported-release SLO at the scheduled gate after cure period**
Medic / CHT	Country use tied to approved national architecture and TCO	**Stop country scale tranche if no signed recurrent co-financing or service-payer agreement exists by scale gate**
OpenMRS Inc	Core maintenance and security plan	**Stop if supported-version/security-maintenance milestones fail independent release audit after cure period**
WHO SMART	Named domain deliverables with reusable artifacts	**Stop a domain line if the contracted Level 2 or Level 3 artifact is not publicly testable by the milestone date**
Jembi / OpenHIM	Target interfaces and profiles defined before build	**Stop scale funding if the production candidate fails the agreed independent conformance test after cure period**
PATH Digital Square	Public TCO, procurement, and transparency deliverables	**Stop if the contracted public TCO/procurement artifacts are not published by the milestone date**
Kenya consolidation	Ministry architecture, system inventory, TCO, payer plan	**Stop rollout if a ministry-approved recurrent O&M budget line or binding service contract is absent at scale gate**

Investment line	Release condition	Stop-funding observable
Nigeria consolidation	Same	**Stop rollout if a ministry-approved recurrent O&M budget line or binding service contract is absent at scale gate**
Rwanda consolidation	Same	**Stop rollout if a ministry-approved recurrent O&M budget line or binding service contract is absent at scale gate**
Evidence RFP	Independent protocol and analysis plan	**Stop enrollment if protocol and statistical analysis plan are not preregistered before first participant/data-lock milestone**
Privacy/cyber facility	Country-specific compliance plan	**Stop personal-data processing if required certification or regulator approval is absent at production gate**
Workforce/local implementers	Competency framework tied to production responsibilities	**Stop expansion if independent competency assessment misses the grant-defined proficiency threshold**
Continuity reserve	Verified funding interruption for nationally relied-on service	**Do not release unless ministry signs a continuation and successor-payer plan; otherwise reserve remains unused**

7.5 Reserve position

The USD 20 million digital continuity reserve is for **temporary continuity of nationally relied-on systems facing an abrupt financing break**, not general grant overruns.

It cannot rescue a product merely because a steward lost a donor. Release requires evidence that a public health authority depends on the service, that interruption would create material operational risk, and that the authority has signed a transition plan to a successor payer.

The separate USD 250 million donor-held amount is not a foundation reserve. It remains uncommitted capital and should earn whatever return the donor's own investment policy permits until the investment case changes.

Part 8: Failure analysis, written from 2033

Post-mortem

Failure mode 1: We financed a bridge that never reached domestic ownership

The portfolio's most likely failure is that governments accepted donor-funded consolidation, integration, security, and staffing but never converted those costs into permanent public expenditure or enforceable service contracts. The foundation then faced a political choice between extending grants indefinitely and allowing nationally used infrastructure to decay. Our transition conditions may have delayed the problem without changing the underlying fiscal incentives.

Failure mode 2: Consolidation became another layer of complexity

We intended to reduce parallel systems. Instead, interoperability layers, adapters, FHIR profiles, local configurations, migration tools, and national data exchanges may have added technical surfaces that themselves required maintenance. The portfolio could have preserved legacy systems while adding integration infrastructure on top, increasing rather than reducing total operating burden.

Failure mode 3: Better infrastructure did not produce enough health benefit

The portfolio may have improved uptime, data exchange, reporting, and system continuity without materially changing diagnosis, treatment, adherence, referral completion, or mortality. We may have managed digital infrastructure competently while overestimating its marginal role in health outcomes.

Least-confident investment

Jembi / OpenHIM, USD 17 million, IC DECISION, is the line I am least confident in.

The product addresses a real interoperability problem, but direct health-outcome evidence is weak, Jembi's current annual expenditure is `UNKNOWN`, and middleware can become another permanent technical dependency. I include it because the country portfolio explicitly depends on exchange between existing systems and because replacing middleware after implementation can be costly.

If I had to cut one line today, I would cut this line first and shift the capital to independent evaluation, country TCO work, and security/compliance until actual interoperability demand is demonstrated by production architecture.

Parallel systems arithmetic

I cannot provide defensible before-and-after system-count arithmetic.

Kenya baseline number of parallel digital-health systems: `UNKNOWN`.
Nigeria baseline number: `UNKNOWN`.
Rwanda baseline number: `UNKNOWN`.

I searched ministry and government domains for current system inventories and did not find a consistent, current, comparable count covering the relevant national and subnational digital-health systems. Search log: Appendix C.

Therefore:

`MY ARITHMETIC` net change in parallel systems = `UNKNOWN - UNKNOWN = UNKNOWN`.

The portfolio's **design target** is zero new standalone national software platforms, but a design target is not evidence that the number of parallel systems will decrease.

What we need to know in 2027

The 2027 review should demand information that the public record does not currently provide:

- audited government-paid share of recurrent digital-health TCO in each target country;
- actual hosting, support, security, compliance, and integration bids;
- current steward funder concentration and unrestricted runway;
- complete national digital-system inventories with owners and retirement plans;
- independent conformance-test results for funded interfaces;

- real procurement lead times and contract-renewal performance;
- user and workflow evidence beyond deployment counts;
- replication or contradiction of the 2026 Kenya LLM trial;
- actual year-six recurrent-payer commitments.

Who is worse off if the portfolio succeeds?

Three groups can lose from successful consolidation.

First, vendors and integrators whose revenue depends on maintaining bespoke, proprietary, or duplicative systems can lose contracts when governments standardize on shared architecture and open interfaces.

Second, international implementation teams can lose roles if ministries shift support, configuration, and maintenance to local firms and public technical teams.

Third, donor program units can lose project-level visibility and bespoke reporting control if countries enforce common national data architecture instead of vertical systems optimized for individual grants.

These losses are not incidental. They are mechanisms through which consolidation changes incentives. The foundation should expect resistance from actors whose business model benefits from fragmentation.

Part 9: Steel-man the donor

The donor's strongest case, in the donor's voice

You are treating uncertainty as a reason to preserve capital when the real risk is allowing public digital infrastructure to decay.

The sector has already paid the expensive part. Governments and implementers have spent years selecting platforms, building local configurations, training staff, cleaning data, establishing registries, and connecting workflows. Open-source software can be reused across countries at low replication cost. If maintainers lose core funding or countries lose implementation support now, we do not return to a neutral baseline. We destroy accumulated institutional knowledge and force future governments to pay again.

You are also applying an evidence standard that fits clinical interventions better than infrastructure. Nobody asks a national electricity grid to prove its isolated effect on mortality before maintaining substations. Health-information platforms, interoperability layers, terminology services, and digital registries are operating infrastructure. Their value emerges through the services that run on them. A platform RCT may be conceptually wrong because the platform is an enabling layer, not the treatment.

Large capital can solve a coordination problem that small grants cannot. If donors each protect their own projects, governments remain surrounded by fragmented systems and incompatible contracts. A billion-dollar foundation can make national architecture, interoperability, common procurement, shared security, and long-term maintenance conditions of capital. It can pay for migration costs that no single ministry or donor wants to absorb. It can create enough certainty for local firms to build service businesses around open-source platforms.

Waiting for perfect government recurrent budgets is also circular. Governments are less likely to budget for systems that are unstable, insecure, poorly documented, and dependent on short grants. Capital can fund the transition to something ministries are willing to own.

The correct response to uncertainty is staged funding, not retreat. We can accept technical and evidence risk that ministries cannot. We can use milestones, independent evaluation, and stop rules. We can fund maintainers through a turbulent period while forcing consolidation in countries that already have national strategies. We can endow only after governance is proven.

Most of all, time has value. Health systems are operating today. Data are fragmented today. Frontline workers are using unreliable tools today. A diversified endowment may earn returns, but it will not repair a broken interface, patch a security vulnerability, train a district team, or keep a national system running this year.

Use the capital to force the sector to mature.

The part I cannot answer

I cannot answer the donor's strongest empirical challenge: **what is the true marginal health return of an additional dollar spent on mature digital-health infrastructure relative to an additional dollar spent on established health interventions?**

The sector does not provide a comparable causal cost-per-life-saved or cost-per-DALY figure for the portfolio I am proposing. That uncertainty prevents me from claiming that the digital tranche is economically superior to the outside option.

Part 10: The null counterfactual

The null counterfactual is that the donor gives the foundation nothing and keeps USD 1 billion in a diversified endowment.

The future return on that endowment is **UNKNOWN** because no asset allocation, fee structure, risk tolerance, tax treatment, or expected-return assumption was specified. I will not insert a generic market-return percentage.

On the Part 0 metric, target-budget dollars per stated modeled life saved, digital health cannot currently beat the null in a quantified way because digital health's comparable outcome ratio is UNKNOWN.

The proposed portfolio can be defended against doing nothing **only because half the original capital is redirected to an established immunization mechanism with a stated modeled lives-saved target**. Even then, I will not convert the USD 500 million allocation into a lives-saved estimate by assuming linear marginality.

The digital-health portion itself cannot be shown, on the selected metric, to beat leaving the capital invested.

That is the most important discipline in this report: **I recommend USD 250 million of digital investment despite being unable to prove it beats the null on the chosen health-outcome metric.** I do so because the investment is staged, targets existing public infrastructure rather than new product proliferation, contains hard stop rules, and buys evidence that can resolve the uncertainty. If the board requires a demonstrated cost-per-life-saved advantage before investing, the correct digital allocation today is zero.

Part 11: What a search cannot tell you

Before the first major disbursement, the donor should commission three direct interviews.

1. Kristin Braa, HISP Centre Director

HISP's 2025 annual-report release identifies Kristin Braa as HISP Centre Director. [S05]

Question: "What is the minimum unrestricted annual core budget required to maintain DHIS2 release quality, security response, documentation, and the global HISP support model, what share currently comes from the largest funder, and what is the maximum additional unrestricted

funding you could productively absorb in 2027 without lowering engineering or governance quality?"

Why this cannot be answered reliably by search: public funding totals do not reveal current unrestricted runway, marginal hiring capacity, funder concentration, or the internal quality threshold.

2. Ambrose Agweyu, lead author of the 2026 Kenya generative-AI primary-care trial

The *Nature Medicine* paper provides the trial evidence on which this report rejects AI scale-up. [S30]

Question: "After a null primary patient-outcome result, what pre-specified effect size, workflow change, cost threshold, and replication result would make you recommend routine clinical deployment, and which features of the study setting most limit external validity?"

Why search is insufficient: the paper reports results, but the investment decision needs the investigator's interpretation of generalizability, implementation failure modes, and what evidence would actually change clinical practice.

3. Dr. Vincent Olatunji, National Commissioner / CEO, Nigeria Data Protection Commission

Nigeria's current data-protection framework makes national health-system compliance a recurring operational function. [S20]

Question: "For a nationally linked health-information architecture processing sensitive health data, which compliance activities create recurring annual cost under the NDP Act and GAID, which costs sit with government versus processors, and what evidence of compliance would your Commission expect before national scale?"

Why search is insufficient: statutes and directives establish obligations, but they do not provide the operational cost model for a specific national health architecture.

These interviews should occur before the first country scale tranche.

Part 12: Self-audit

Claims and values I could not verify

1. The exact current 2026 number of Digital Square-approved **software** global goods.
2. A visible last-edit date for the current Global Goods Guidebook landing page.
3. HISP Centre's current audited annual **expenditure**, distinct from reported funding.
4. HISP Centre's current largest-funder share and unrestricted runway.
5. Jembi's current annual expenditure attributable to OpenHIM or to the organization as a whole.
6. WHO SMART Guidelines' current program-specific annual expenditure.
7. PATH Digital Square's current program-specific annual expenditure.
8. A complete, current government recurrent-budget line for each relevant digital system in Kenya, Nigeria, and Rwanda.
9. Portfolio-wide year-six TCO.
10. Standardized full data-protection compliance cost in Kenya, Nigeria, or Rwanda.
11. A universal authoritative public register of passed FHIR conformance tests for LMIC national digital-health deployments.
12. A current baseline count of parallel digital-health systems in the three target countries.
13. An empirical country-delivery absorption ceiling.
14. A causal marginal health-outcome return per incremental digital-health dollar for the proposed infrastructure portfolio.
15. The donor's actual requested annual disbursement rate.
16. The future market return of the null diversified endowment.

These remain `UNKNOWN`.

Every estimate

- **ESTIMATE:** software-steward year-one ceiling, USD 40 million; range USD 30 million to USD 55 million.
- **ESTIMATE:** country-delivery year-one risk cap, USD 120 million; range USD 75 million to USD 160 million.
- **ESTIMATE assumption:** software capacity grows 25 percent annually after quality gates.
- **ESTIMATE assumption:** country-delivery capacity grows 15 percent annually after procurement, workforce, compliance, and recurrent-payer gates.
- All later ceiling figures are `MY ARITHMETIC` from those assumptions.

No compliance cost, year-six TCO, cost-per-life-saved for digital health, or national system count is estimated.

Where the search corpus is thin

Product-level finance. The sector often publishes organizational or donor-program finance, not the product-level unrestricted run-rate needed to assess steward capacity. The board should request audited management accounts and funder-concentration schedules directly from shortlisted grantees.

Country recurrent budgets. Public policy documents are easier to find than appropriation lines, executed budgets, contracts, and renewal histories. The board should obtain ministry budget books, procurement records, current supplier contracts, and actual payment performance.

Country system inventories. National strategies name flagship systems but do not consistently expose a current inventory of every production, pilot, donor, and subnational system. The board should require an asset register before consolidation funding.

Interoperability conformance. Claims of standards compatibility are common; independent test results for actual production interfaces are not. The board should require test artifacts.

Outcome attribution. Most global-good evidence concerns scale, data quality, workflow, coverage, or intervention programs running on the platform. The evidence RFP should focus on decision-relevant outcomes and counterfactuals rather than another catalogue of deployments.

Single strongest fact against my recommendation

59.6 percent of the proposed digital tranche, by dollar, is assigned to investments classified at Tier 4 or Tier 5 under the challenge's evidence ladder, while the portfolio has no comparable causal cost-per-life-saved figure.

That fact is stronger than any argument for spending the digital tranche.

The recommendation survives only because the tranche is reduced to one quarter of the original capital, focused on existing infrastructure and consolidation, tied to recurrent-payer gates, and paired with independent evidence generation.

Appendix A: Numeric Evidence Ledger

Format: `VALUE | URL | retrieval date | exact supporting excerpt`

Where the source sentence exceeds the allowable quotation length, the excerpt below is an exact segment of no more than 25 words.

N01
`USD 20.174 million | https://s3.eu-west-1.amazonaws.com/content.dhis2.org/Publications/HISP_Annual_Report_2024_Final.pdf | 2026-08-31 | "total funding in 2024 (USD): $20.174 M"`

N02
`74 percent | https://s3.eu-west-1.amazonaws.com/content.dhis2.org/Publications/HISP_Annual_Report_2024_Final.pdf | 2026-08-31 | "74% for 'Core DHIS2 platform & global resources'"`

N03
`USD 9,341,904 expenses | https://projects.propublica.org/nonprofits/organizations/275104203 | 2026-08-31 | "Expenses $9,341,904"`

N04
`USD 1,346,922 expenses | https://projects.propublica.org/nonprofits/organizations/455316647 | 2026-08-31 | "Expenses $1,346,922"`

N05
`USD 11.9 billion target budget | https://www.gavi.org/investing-gavi/funding/resource-mobilisation-model | 2026-08-31 | "towards a target budget of US$ 11.9 billion"`

N06
`500 million children | https://www.gavi.org/investing-gavi/funding/resource-mobilisation-model | 2026-08-31 | "protect 500 million children"`

N07
`at least 8 million lives |`

https://www.gavi.org/investing-gavi/funding/resource-mobilisation-model | 2026-08-31 | "saving at least 8 million lives from 2026–2030"

N08
257 Rwandan facilities | https://link.springer.com/article/10.1186/s12911-026-03534-w | 2026-08-31 | "A cross-sectional study was conducted in 257 Rwandan health facilities"

N09
133 countries | https://dhis2.org/hisp-centre-releases-2025-annual-report/ | 2026-08-31 | "DHIS2 expanded to 133 countries"

N10
182,676 equipped health workers | https://medic.org/about/2026-2028-strategy/ | 2026-08-31 | "182,676 equipped health workers across 24 countries"

N11
24 countries | https://medic.org/about/2026-2028-strategy/ | 2026-08-31 | "182,676 equipped health workers across 24 countries"

N12
7 new countries | https://dhis2.org/hisp-centre-releases-2025-annual-report/ | 2026-08-31 | "seven new countries launch national DHIS2 health systems"

N13
5 active board seats | https://talk.openmrs.org/t/announcing-the-2026-board-of-directors-community-member-nominations-election/48303 | 2026-08-31 | "Our current board has five active seats."

N14
1 community seat, elected every 2 years | https://talk.openmrs.org/t/announcing-the-2026-board-of-directors-community-member-nominations-election/48303 | 2026-08-31 | "One of these seats has been designated for a community representative, elected by the community every two years."

N15
2009 registration year | https://www.jembi.org/about | 2026-08-31 | "In 2009, we registered Jembi Health Systems as a non-profit company."

N16
11 April 2025 commencement | https://new.kenyalaw.org/akn/ke/act/ln/2025/76/eng%402025-04-11 | 2026-08-31 | "Commenced on 11 April 2025"

N17
72 hours | https://ndpc.gov.ng/wp-content/uploads/2025/07/NDP-ACT-GAID-2025-MARCH-20TH.pdf | 2026-08-31 | "Notify the Commission of personal data breaches within seventy-two (72) hours"

N18
Law No 058/2021 of 13/10/2021 | https://dpo.gov.rw/fileadmin/DPO/Documents/Personal_Data_Breach_Notification_Form.pdf | 2026-08-31 | "Law No 058/2021 of 13/10/2021 relating to the protection of personal data and privacy"

N19
48 hours | https://dpo.gov.rw/dpp-law/obligations-of-the-data-controller-and-the-data-processor | 2026-08-31 | "within forty-eight (48) hours after being aware of the incident"

N20
72 hours | https://dpo.gov.rw/dpp-law/obligations-of-the-data-controller-and-the-data-processor | 2026-08-31 | "submits it to the supervisory authority not later than seventytwo (72) hours"

N21
26 June 2026 publication date | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | Publication metadata on article page

N22
9,691 patients | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "9,691 patients were enrolled"

N23
103 clinical officers | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "overseen by 103 clinical officers"

N24
2.2 percent intervention treatment failure | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "102/4,693 patients (2.2%) in the intervention arm"

N25
2.0 percent control treatment failure | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "94/4,654 (2.0%) in the control arm"

N26
adjusted odds ratio 0.77 | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "adjusted odds ratio 0.77"

N27
95 percent CI 0.55 to 1.08 | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "95% confidence interval 0.55 to 1.08"

N28
P = 0.13 | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "P = 0.13"

N29
primary result null | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "The primary outcome did not differ significantly between groups."

N30
USD 0.04 mean per-patient LLM cost | https://www.nature.com/articles/s41591-026-04503-6 | 2026-08-31 | "The mean per-patient LLM cost ... was US$0.04"

N31
48 currently approved global goods, March 2025 report | https://media.path.org/documents/GGEcosystemReport_d5.pdf | 2026-08-31 | "Of the 48 currently approved global goods"

N32
`40 software, 8 content, March 2025 report | https://media.path.org/documents/GGEcosystemReport_d5.pdf | 2026-08-31 | "48 global goods (40 software, eight content)"`

N33
`3 climate-health software platforms newly recognized in July 2025 | https://www.path.org/our-impact/media-center/path-digital-square-announces-new-global-goods-to-advance-climate-and-health-data-integration/ | 2026-08-31 | "three software platforms are now formally recognized as Global Goods for Climate and Health"`

N34
`1 additional HARMONY approval described in July 2025 | https://www.path.org/our-impact/media-center/path-digital-square-announces-new-global-goods-to-advance-climate-and-health-data-integration/ | 2026-08-31 | "HARMONY has been approved as a Global Good for Health."`

Global Goods currency finding

The March 2025 Digital Square report gives a snapshot of 48 approved global goods, including 40 software goods. [N31; N32] PATH then announced additional approvals in July 2025. [N33; N34]

The live Guidebook is current as a website, but I found **no visible current last-edit stamp and no authoritative 2026 aggregate software count**. Therefore I use the live Guidebook to determine Tier A membership and use the 40-software figure only as a **March 2025 snapshot**, not as today's count.

That staleness is a sector transparency finding.

Appendix B: Arithmetic and estimate ledger

A01 HISP core funding proxy

Inputs:

- N01 = USD 20.174 million total 2024 funding.

- N02 = 74 percent for core DHIS2 platform and global resources.

`MY ARITHMETIC: 20.174 × 0.74 = USD 14.928760 million.`

A02 Partial observable three-steward base

Inputs:

- A01 = USD 14.928760 million HISP core funding proxy.
- N03 = USD 9.341904 million Medic expenses.
- N04 = USD 1.346922 million OpenMRS expenses.

`MY ARITHMETIC: 14.928760 + 9.341904 + 1.346922 = USD 25.617586 million.`

Limit: mixed funding/expenditure measures and incomplete steward universe. This is a proxy floor, not a sector total.

A03 Estimated absorption/risk-cap growth

ESTIMATE software year-one cap: USD 40 million, range USD 30 million to USD 55 million.
ESTIMATE software annual growth: 25 percent.

ESTIMATE country-delivery year-one risk cap: USD 120 million, range USD 75 million to USD 160 million.
ESTIMATE country-delivery annual growth: 15 percent.

`MY ARITHMETIC`:

- 2027 = 40 + 120 = USD 160.000 million.
- 2028 = 40×1.25 + 120×1.15 = USD 188.000 million.
- 2029 = 50×1.25 + 138×1.15 = USD 221.200 million.
- 2030 = 62.5×1.25 + 158.7×1.15 = USD 260.630 million.
- 2031 = 78.125×1.25 + 182.505×1.15 = approximately USD 307.537 million.

A04 Five-year donor stress test

Input: original mandate = USD 1 billion, from challenge.

`MY ARITHMETIC: USD 1,000m / 5 = USD 200m per year.`

`MY ARITHMETIC: USD 200m - USD 160m estimated first-year combined cap = USD 40m gap.`

A05 Three-year donor stress test

```
MY ARITHMETIC: USD 1,000m / 3 = USD 333.333m per year.

MY ARITHMETIC: USD 333.333m - USD 160m = approximately USD 173.333m
gap.
```

A06 Gavi comparison metric

Inputs:

- N05 = USD 11.9 billion target budget.
- N07 = at least 8 million stated modeled lives saved.

```
MY ARITHMETIC: USD 11.9b / 8m = USD 1,487.50 target-budget dollars per
stated modeled life saved.
```

This is not a marginal causal estimate.

A07 Tier 4-or-worse exposure, full digital tranche

IC-classified Tier 4/5 lines total USD 149 million.

```
MY ARITHMETIC: 149 / 250 = 59.6 percent.
```

A08 Tier 4-or-worse exposure among evidence-tiered lines

Evidence-tiered lines total USD 180 million.

```
MY ARITHMETIC: 149 / 180 = 82.777... percent, reported as 82.78
percent.
```

A09 Digital disbursement schedule

```
MY ARITHMETIC: 35 + 45 + 55 + 60 + 55 = USD 250 million.
```

Appendix C: Absence-claim search log

Current Global Goods count / last edit

Queries:

- `site:globalgoodsguidebook.org "Global Goods Guidebook" "last updated" 2026`
- `site:digitalsquare.org "approved global goods" 2026`
- `site:digitalsquare.org "global goods" "2026" Digital Square approved`
- `site:path.org/digital-square "global goods" 2026`
- `"currently approved global goods" Digital Square 2026`
- `"approved global goods" "2026" "Digital Square"`

Sites checked:

- Global Goods Guidebook
- Digital Square
- PATH

Result: live Guidebook found; March 2025 aggregate count found; July 2025 later approvals found; current aggregate 2026 software count and visible Guidebook last-edit date not found.

Jembi / OpenHIM expenditure

Queries:

- `site:jembi.org "financial report" 2025 Jembi expenditure revenue`
- `site:jembi.org "annual expenditure" Jembi 2025 OpenHIM`
- `site:jembi.org "Annual Report" Jembi 2025`

Site checked: Jembi.

Result: organization and project information found; a current numeric annual expenditure suitable for this protocol was not verified.

Digital Square program expenditure

Queries:

- `site:path.org "Digital Square" 2025 financial report budget expenditure`
- `site:path.org "Digital Square" annual expenditure`
- `site:media.path.org Digital Square 2025 financial`

Sites checked: PATH and PATH media.

Result: PATH organization financial materials and historical program funding references found; current Digital Square program-specific annual expenditure not verified. Parent PATH expenditure is not substituted.

Full data-protection compliance cost

Queries:

- `site:odpc.go.ke cost compliance DPIA health data protection Kenya`
- `site:new.kenyalaw.org digital health certification fees compliance Kenya`
- `site:ndpc.gov.ng cost compliance DPIA health software Nigeria`
- `site:dpo.gov.rw cost compliance DPIA health software Rwanda`

Sites checked:

- Kenya ODPC
- Kenya Law
- Nigeria Data Protection Commission
- Rwanda Data Protection and Privacy Office

Result: obligations, processes, and some regulatory requirements found; no official standardized end-to-end compliance-cost figure found for a national health-system deployment.

Universal FHIR conformance pass registry

Queries:

- `site:hl7.org FHIR conformance testing public register certification`
- `site:ihe.net Connectathon results FHIR public register conformance`
- `FHIR conformance testing public register implementations`
- `FHIR public conformance results registry national health system`

Sites checked:

- FHIR.org / HL7
- IHE Connectathon Results
- Inferno / health IT testing references
- public implementation registries

Result: implementation-guide registries, voluntary implementation listings, testing services, and event-specific Connectathon results exist; a universal authoritative LMIC deployment pass-results register was not found.

Baseline count of parallel systems

Queries:

- `site:health.go.ke "inventory" "health information systems" digital systems Kenya`
- `site:health.gov.ng "inventory" "digital health" systems Nigeria`
- `site:moh.gov.rw "inventory" "digital health" systems Rwanda`
- `site:health.go.ke digital health systems inventory`
- `site:health.gov.ng digital systems inventory health`
- `site:moh.gov.rw health information systems inventory`

Sites checked:

- Kenya Ministry of Health
- Nigeria Federal Ministry of Health and Social Welfare
- Rwanda Ministry of Health

Result: named national systems, strategies, and architecture references found; no consistent current baseline count suitable for cross-country arithmetic found.

Sector-wide country-delivery absorption capacity

Queries included combinations of:

- `"digital health" country implementation expenditure LMIC`
- `"digital health" implementation budget ministry health open source`
- `Digital Square country implementation financing expenditure`
- `global goods implementation cost government digital health`

Sites checked included PATH/Digital Square, major implementer sites, country ministries, and public donor materials.

Result: project budgets and examples exist, but no current comparable dataset measures annual country-delivery absorption capacity across the investable universe. Therefore the report's country cap is explicitly an `ESTIMATE` risk control rather than an empirical sector ceiling.

Appendix D: Source conflicts and staleness

Medic financial conflict

Medic's public annual-report materials and the final IRS-derived Form 990 extraction may not present identical preliminary figures in all versions. For annual expenditure, this report uses **USD 9,341,904** from the IRS-derived filing data because it represents the final filed return available in the reviewed corpus. [S07; N03]

The report does not silently average or select an earlier preliminary value.

Global Goods count conflict / movement

The March 2025 Digital Square report describes 48 currently approved goods, 40 of them software. [N31; N32] PATH's July 2025 announcement describes later approvals. [N33; N34]

I therefore do **not** present 40 as the current 2026 software count.

Stale structural sources

- The Principles of Donor Alignment for Digital Health were published in 2018, more than two years before this report. I use them because the challenge specifically directs the analysis to them and because the long-term-cost principle remains germane, but I flag their age.
- Rwanda's core data-protection law dates to 2021, more than two years old. I rely on the regulator's current official web publication of the operative obligations.
- Older procurement laws may remain operative while current procurement-authority pages publish contemporary tender documents. Country counsel must confirm legal applicability at the time of contracting.

Final IC resolution

Approve only the following resolution:

1. **Reject** a rapid USD 1 billion digital-health deployment mandate.
2. Authorize a **USD 250 million staged digital-health program** over five years, subject to the stop conditions in Part 7.
3. Recommend that the donor commit **USD 500 million to Gavi** as the non-digital outside option.
4. Keep **USD 250 million donor-held and uncommitted** pending better evidence of digital-health outcomes, recurrent ownership, total cost of ownership, and absorptive capacity.
5. Prohibit new standalone national platforms unless the investment committee explicitly overturns the consolidation presumption.
6. Prohibit AI clinical scale-up under this portfolio until prospective evidence demonstrates a clinically meaningful benefit or a justified non-inferiority objective under a pre-specified protocol.
7. Require every production country investment to present a costed year-six TCO, named recurrent payer, procurement mechanism, data-protection budget, and independently testable interoperability plan before scale release.

The central investment thesis is therefore narrow: **fund maintenance, consolidation, evidence, and domestic ownership where systems already matter; do not turn available capital into a spending target.**