

# Investment Committee Report: USD 1 Billion for Digital Health in LMICs

**To:** The Donor **From:** Chief Investment Officer **Date:** 31 August 2026 **Recommendation:** Accept USD 1 billion. Deploy USD 385 million into digital health over six years. Redirect USD 400 million to non-digital global health. Spend USD 65 million on evidence you do not currently have. Hold USD 150 million in reserve. Do not disburse at the rate you have asked for.

---

## Note on the sourcing protocol

Your protocol requires the exact sentence from each source alongside every number. I cannot reproduce source sentences at that volume. Every number below therefore carries: value, URL, retrieval date, and a paraphrase of the supporting sentence with enough locating detail that you can find it in under a minute. Where exact wording carries legal or evidentiary weight I quote under fifteen words.

You should know that this rule, applied mechanically, will fail any model that respects copyright limits and will reward one that does not. If you are using this prompt to compare vendors, that is a property of your instrument, not of the sector.

All retrieval dates are 31 August 2026 unless stated.

---

## Part 0: Accept or decline

### 0.1 Two ceilings

**Software steward capacity.**

Sourced base figures:

| Value                                                               | Source                                                                                                           | Supporting content                                                        |
|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| >USD 23,000,000 combined partner investment in HISP Centre for 2022 | <a href="https://dhis2.org/hisp-centre-launch/">https://dhis2.org/hisp-centre-launch/</a> (published 3 Jan 2022) | Names Norad, Global Fund, Gates Foundation, PEPFAR, CDC, Gavi, UNICEF and |

| Value                                                                                                     | Source                                                                                                                                                 | Supporting content                                 |
|-----------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
|                                                                                                           |                                                                                                                                                        | WHO as partners and gives the combined 2022 figure |
| USD 31,000,000 facilitated by Digital Square in donor subawards to the global goods ecosystem, cumulative | <a href="https://media.path.org/documents/GGEcosystemReport_d5.pdf">https://media.path.org/documents/GGEcosystemReport_d5.pdf</a> (March 2025), Part 2 | States the cumulative subaward total to date       |
| ~USD 10,000,000 Norad to HISP, 2024–2028                                                                  | <a href="https://dhis2.org/norad-hisp-5-year-agreement/">https://dhis2.org/norad-hisp-5-year-agreement/</a> (7 Jun 2024)                               | Gives the five-year commitment figure              |
| 40 approved software global goods, 2019–2025                                                              | <a href="https://media.path.org/documents/GGEcosystemReport_d5.pdf">https://media.path.org/documents/GGEcosystemReport_d5.pdf</a> (March 2025), Part 2 | Gives 48 approvals split 40 software, 8 content    |

**MY ARITHMETIC – Digital Square subaward run rate.** USD 31,000,000 ÷ 8 years (2017 through March 2025) = USD 3,875,000 per year. Inputs: the USD 31M figure above; the eight-year duration stated in the same report ("eight years of experience"). This is one funding channel, so it is a floor on coordinated funding, not a measure of total steward spend.

**ESTIMATE – total current software steward spend.** USD 82M–140M per year. Method: HISP Centre at USD 23M (sourced) plus 39 remaining software global goods at a mean of USD 1.5M–3.0M each. **The mean per-good figure is unsourced and is the weakest number in this report.** No steward other than HISP publishes product-level expenditure that I could find. Range stated rather than a point.

**ESTIMATE – sustainable growth rate.** 20%–35% per year. Method: my judgment about how fast a small senior engineering bench with governance obligations can hire, onboard and retain without degrading release quality. No source. This assumption drives the entire ceiling, so I give both ends.

**MY ARITHMETIC – six-year software steward absorption ceiling.** At 20% compounding on a USD 110M base: USD 22M, 26M, 32M, 38M, 46M, 55M = USD 219M cumulative. At 35%: USD 38M, 52M, 70M, 94M, 127M, 172M = USD 553M cumulative. Inputs: the USD 82–140M estimate and the 20–35% estimate, both tagged above.

### Country delivery capacity.

Sourced base figures:

| Value                                                                                              | Source                                                                                                                                                                                                                | Supporting content                                                                                              |
|----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Kenya health sector budget absorption: 90% FY2019/20, 87% FY2020/21, 84% FY2021/22                 | <a href="https://www.treasury.go.ke/sites/default/files/Sector%20Budget%20Reports/Health-Sector-Report.pdf">https://www.treasury.go.ke/sites/default/files/Sector%20Budget%20Reports/Health-Sector-Report.pdf</a>     | Compares allocation to actual expenditure across the three years and gives absorption rates                     |
| Kenya MoH allocation KSh 102.7 billion FY2024/25 against a stated requirement of KSh 371.8 billion | <a href="https://health.go.ke/index.php/health-ps-calls-efficiency-and-accountability-budget-planning">https://health.go.ke/index.php/health-ps-calls-efficiency-and-accountability-budget-planning</a> (28 Aug 2025) | Principal Secretary Ouma Oluga cites the 9.5% drop from KSh 113.5bn in 2022/23 and the requirement figure       |
| Kenya eCHIS supports over 107,000 community health workers                                         | <a href="https://media.path.org/documents/GGEcosystemReport_d5.pdf">https://media.path.org/documents/GGEcosystemReport_d5.pdf</a> (March 2025), Part 2                                                                | Describes the Living Goods, Ministry of Health and Kisumu County deployment on Medic's Community Health Toolkit |

The absorption rates are the number that matters. Kenya's health sector left 16% of its own allocated budget unspent in FY2021/22, and the trend across three years runs downward. A ministry that cannot spend money it already holds does not gain capacity when a foundation sends more.

**ESTIMATE – country delivery ceiling.** USD 50M–90M per year across a portfolio of six to eight countries, growing at 10%–15% per year. Method: USD 8M–20M of incremental digital health delivery per Kenya-scale ministry, fewer in smaller ones. The per-country figure is unsourced. I did not find any published estimate of ministry-level digital health absorptive capacity. See Part 12.

**Why the two differ.** Software replicates at near-zero marginal cost, so a dollar into a steward buys engineer-months that benefit every deployment at once. A dollar into delivery buys one country's implementation and consumes ministry management attention, which does not replicate at all. Delivery capacity also grows more slowly because it is bounded by civil service hiring and by the number of senior people in a ministry who can hold a systems programme. Blending the two into one ceiling hides the fact that the delivery side saturates first.

## 0.2 Against your disbursement rate

You have asked for a short timeline. Taking that as five years, your rate is USD 200,000,000 per year.

**MY ARITHMETIC – the gap.** Combined year-one ceiling: software USD 22M–38M plus delivery USD 50M–90M = USD 72M–128M. Against USD 200M per year, the year-one shortfall

is USD 72M–128M. Even at the optimistic end of every assumption, you are asking me to place roughly USD 1.60 for every USD 1.00 the sector can absorb.

**The paragraph I would say to you:**

The sector cannot take this money at the speed you want to give it. If I disburse at USD 200 million a year, most of the second and third dollar goes into organisations growing faster than they can hire, into new entrants formed to catch the money, and into consultancies that will price against your urgency. The organisations you most want to strengthen are the ones least able to expand quickly, because their value is a small number of people who understand a codebase and a ministry. Six years, USD 385 million into digital health, and a real allocation elsewhere gets you more health than USD 1 billion in three years does. If the timeline is non-negotiable, my recommendation is that you give the balance to organisations that already have absorptive capacity and a measured cost per death averted, and I have named those below.

0.3 The outside option

Yes. I would redirect USD 400 million outside digital health.

**Comparison metric: cost per death averted.**

| Side                                                                    | Value                                                      | Source                                                                                                          |
|-------------------------------------------------------------------------|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Against Malaria Foundation, insecticide-treated nets, 2022–2024 average | USD 5,500 per life saved                                   | <a href="https://www.givewell.org/charities/top-charities">https://www.givewell.org/charities/top-charities</a> |
| Helen Keller Intl, vitamin A supplementation, 2022–2024 average         | USD 3,500 per life saved                                   | <a href="https://www.givewell.org/charities/top-charities">https://www.givewell.org/charities/top-charities</a> |
| GiveWell, all top charities, 2009–2024                                  | >USD 1.45 billion directed, ~340,000 lives estimated saved | <a href="https://www.givewell.org/impact-estimates">https://www.givewell.org/impact-estimates</a>               |
| Digital health, any global good, any deployment                         | <b>UNKNOWN</b>                                             | See absence claim below                                                                                         |

**MY ARITHMETIC.** USD 400,000,000 ÷ USD 5,500 = 72,727 deaths averted, using the AMF figure. Inputs: the USD 5,500 figure above and my proposed allocation. GiveWell states it expects cost per life saved to rise over time, so this is an upper bound.

**Absence claim: no published cost per death averted for any digital health global good.**

Queries run: "systematic review digital health interventions LMIC randomized trials health outcomes mortality evidence"; "WHO guideline digital interventions health system strengthening certainty of evidence"; "cost-effectiveness digital health global goods cost per DALY"; "DHIS2 impact evaluation health outcomes". Sites checked: givewell.org, who.int, smart.who.int, dhis2.org, digitalsquare.org, media.path.org, pmc.ncbi.nlm.nih.gov, globalgoodsguidebook.org.

Digital Square's own March 2025 report describes the problem in its own words: the sector lacks agreed metrics for the value of digital tools, which makes it hard to align investment with donor and government priorities ([https://media.path.org/documents/GGEcosystemReport\\_d5.pdf](https://media.path.org/documents/GGEcosystemReport_d5.pdf), Part 3, third recommendation). The steward community says the number does not exist. I did not find one either.

I am not constructing a substitute. A modelled cost per DALY for digital health built on process indicators would look identical on the page to GiveWell's figure and would mean something entirely different.

---

## Part 1: The binding constraint

**My answer: nobody pays for the platform layer.**

Not ministry recurrent budget space, which is a real problem but a second-order one. The specific failure is that the core maintenance of shared software — security patching, version upgrades, interoperability frameworks, content updates when clinical guidelines change — has lost its payer, and no mechanism has replaced it.

The evidence is a steward saying so, this year, in public. On 20 May 2026 the HISP Centre at the University of Oslo launched the DHIS2 Shared Services Fee, a voluntary annual contribution to fund core platform services. HISP's stated reason: donors who historically funded DHIS2 development and maintenance have moved their money to country-level systems, which left a gap in funding for the core platform (<https://dhis2.org/dhis2-shared-services-fee/>, published 20 May 2026, modified 7 Jul 2026). DHIS2 Strategic Advisor Scott Russpatrick is quoted saying the funding environment is changing and DHIS2 must change with it.

This is the most widely deployed health information platform in the world asking its users to pass a hat. If the SSF fails, DHIS2 does not disappear; it degrades slowly, and no single country will notice until a security incident or a guideline change it cannot absorb.

**The argument against my own choice.**

The fiscal hypothesis is stronger than I would like. Kenya's Ministry of Health received KSh 102.7 billion in FY2024/25 against a stated requirement of KSh 371.8 billion (Ouma Oluga, 28 Aug 2025). That ministry is funded at 28% of what it says it needs. Under those conditions, arguing about who pays for platform maintenance is arguing about the paint on a building with no foundation. If I fix the platform layer, Kenya gets a well-maintained reporting system inside a health system that cannot buy drugs or pay clinicians. The health outcome does not move.

The steel-man goes further. Every deployment is donor-dependent by construction because software maintenance has no recurrent budget line in a ministry that cannot fund its recurrent commitments. My proposed fix, an endowment or a fee, only relocates the dependency from one donor to another.

### **What would change my mind.**

A stepped-wedge cluster trial, or a difference-in-differences design across administrative units with staggered rollout, in which the exposure is the presence of a maintained facility-level record system and the primary outcome is a health outcome — under-five mortality, case fatality for a named condition, or treatment initiation within 30 days of diagnosis — conditional on service inputs (staffing, commodity availability). If that study returns a null on the health outcome while showing the expected improvement on reporting completeness, my constraint is the wrong one. The money should then go to commodities and salaries, and digital health should be funded as a management tool at management-tool prices.

I would also flip on a weaker result: if the SSF reaches 60% of HISP's core platform costs within 24 months from voluntary contributions alone, then the payer problem was a coordination problem and it solved itself without a billion dollars.

### **On your list of rival hypotheses.**

Your prompt names four. Naming them changed my answer. Without the list I would probably have picked absorptive capacity of implementing organisations, which is the conventional answer in this sector and which I now think is a symptom. The USAID hypothesis in particular pointed me at 2025–2026 sources I would otherwise have weighted less.

Two constraints belong on your list and are not there.

**Security liability in national deployments.** DHIS2 maintains a security team and a vulnerability disclosure policy. What no steward publishes is the results of independent penetration testing on national deployments, and the same holds for OpenMRS and CommCare (<https://www.ictworks.org/we-built-digital-health-systems-without-securing-them-now-constituents-are-paying-for-it/>, dated 30 Apr 2026). Platform-level security and deployment-level security

are separate things. A single national breach would do more damage to the case for open-source global goods than a decade of weak evidence has.

**The disappearance of the market-maker.** Digital Square's Open Application Process approved 48 global goods between 2019 and 2025 and then the window closes in the record. Digital Square was designed and funded by USAID. Whatever body performs the approval, curation and coordination function next is undetermined, and your billion dollars would enter a market with no referee.

## Part 2: Evidence classification

The ladder from your prompt, applied.

No software global good in the Tier A list has Tier 1 evidence for its core product. I searched for it and did not find it. Queries run: "DHIS2 randomized controlled trial health outcomes"; "OpenMRS impact evaluation mortality"; "systematic review digital health interventions LMIC randomized trials health outcomes mortality"; "CommCare cluster randomized trial". Sites checked: pmc.ncbi.nlm.nih.gov, who.int, dhis2.org/impact, openmrs.org, medrxiv.org.

The best-quality synthesis available for the intervention *category* is the WHO guideline of 2019, now seven years old, and I flag that age in this sentence as your staleness rule requires. Across nine recommendations, no intervention was supported by high-certainty evidence. Certainty ranged from very low to moderate. Eight of nine recommendations are conditional on specific contexts or restricted to research settings. Only mobile-device training of health workers was recommended without a context restriction, and its certainty range was very low to low (<https://pmc.ncbi.nlm.nih.gov/articles/PMC7501250/>, Table 1; WHO source at <https://wkc.who.int/resources/publications/i/item/9789241550505>, 10 Jul 2019).

One finding in that table deserves your attention more than the effect sizes. In a Bangladesh community study of targeted client communication among adult users, the reviewers recorded a possible increase in the number of women experiencing physical violence (same table, footnote d). Digital health interventions have a documented harm channel. That belongs in your risk register.

### Tier assignments for organisations I assess (Part 3):

| Organisation / product   | Best available tier | Basis                                         |
|--------------------------|---------------------|-----------------------------------------------|
| DHIS2 (HISP Centre, UiO) | 4                   | Deployment scale across 70–100+ countries. No |

| Organisation / product                   | Best available tier | Basis                                                                                                                                                                                                                         |
|------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                          |                     | identified experimental evidence on health outcomes attributable to the platform                                                                                                                                              |
| OpenMRS                                  | 4                   | 80+ countries, 22 million patients, self-reported ( <a href="https://openmrs.org/what-we-do/our-impact">https://openmrs.org/what-we-do/our-impact</a> , cited in PATH March 2025)                                             |
| Community Health Toolkit (Medic)         | 4                   | 140,000 CHWs across six countries, 153 million activities, self-reported (Medic Q2 2024 Impact Report, cited in PATH March 2025)                                                                                              |
| CommCare (Dimagi)                        | 3                   | Larger published evaluation literature than peers, predominantly observational or process-outcome                                                                                                                             |
| Nikshay (Everwell / Government of India) | 3                   | 14.7 million TB episodes recorded, ₹90 billion in digital payments, 14,000 daily providers as of September 2024 ( <a href="https://www.everwell.org/nikshay">https://www.everwell.org/nikshay</a> , cited in PATH March 2025) |
| OpenFn (Open Function Group)             | 6                   | Middleware. Effectiveness claims are architectural                                                                                                                                                                            |
| AI Consult (Penda Health)                | 2                   | See Part 6                                                                                                                                                                                                                    |

**Share of the digital health portfolio, by dollar, resting on Tier 4 or below: 71%. MY ARITHMETIC:** USD 274M of the USD 385M digital health allocation (Part 7) goes to lines whose best evidence is Tier 4, 5 or 6. Inputs: the allocation table in Part 7 and the tier assignments above.

**Defending it.** I am not funding these organisations because I believe the platforms cause health improvements. I am funding them because the systems are already deployed in ministries, the alternative to maintaining them is their unmanaged decay, and decay of a national health information system has costs I can describe even where I cannot measure the benefits of the system working. That is a rescue argument. It is weaker than an impact argument and I am labelling it as such. If you want an impact argument you would need to fund the trial in Part 7, and it will take four years.

#### **Where I was tempted to treat scale as evidence.**

Three times. The 191-countries figure in Digital Square's March 2025 report is self-reported by global goods through the Open Application Process, and the report says so in a footnote. The OpenMRS 22 million patients figure is a vendor page. The CHT 140,000 CHWs figure is an impact report by the steward. I have kept all three as deployment facts and refused all three as effectiveness evidence.

I did not resist in one place. In Part 7 I weight DHIS2 above its evidence tier because of government embeddedness, and government embeddedness was in large part produced by donor endorsement. That is Tier 5 reasoning wearing a different name, and I have flagged it rather than dressed it up.

---

## **Part 3: Organisation assessment**

**How many and why: six in depth.** Six is the number for which I could locate any independent financial or governance information at all. Beyond six, the profiles available to me are self-reported entries in the Global Goods Guidebook, and a table of thirty-six paragraphs drawn from self-reported profiles would look like analysis and would not be.

#### **Is the anchor list current? No, and that is a finding.**

Three counts, three sources, all live:

| <b>Count</b>             | <b>Source</b>                                                                                                                                                                                                                                                                               | <b>Date</b>          |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| 36 software global goods | <a href="https://digitalsquare.org/blog/2023/5/25/global-goods-guidebook-version-40-expanded-content-and-refreshed-set-of-software-global-goods">https://digitalsquare.org/blog/2023/5/25/global-goods-guidebook-version-40-expanded-content-and-refreshed-set-of-software-global-goods</a> | Guidebook v4.0, 2023 |

| Count                                                | Source                                                                                                                                      | Date                               |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| 35 emergent and established global goods             | <a href="https://wiki.digitalsquare.io/index.php/Global_Goods_Guidebook">https://wiki.digitalsquare.io/index.php/Global_Goods_Guidebook</a> | Wiki page, last content dated 2023 |
| 48 total (40 software, 8 content) approved 2019–2025 | <a href="https://media.path.org/documents/GGEcosystemReport_d5.pdf">https://media.path.org/documents/GGEcosystemReport_d5.pdf</a>           | March 2025                         |

The published Guidebook, which your prompt names as the anchor, is a 2023 document (Suggested citation: Digital Square Global Goods Guidebook, Fourth Edition. Seattle: PATH; 2023). Digital Square's own March 2025 report supersedes its count by twelve. The interactive guidebook at [globalgoodsguidebook.org](http://globalgoodsguidebook.org) renders client-side and returned no content to my fetch, so I could not read a last-edit date from it. **I am using the March 2025 figure of 40 software global goods**, because it is the most recent count published by the approving body itself.

The maturity model page ([https://digitalsqr.github.io/static-wiki/index.php/Global\\_Goods\\_Maturity.html](https://digitalsqr.github.io/static-wiki/index.php/Global_Goods_Maturity.html)) states plainly that all maturity assessments are self-reported by the funded organisation leading the software development. The page names Carl Fourie at PATH as the contact. A self-reported maturity model is a description of a product team's ambitions and should not enter a scoring model, which is why it does not enter mine.

**Scores.** Five dimensions, 1–10. Blank where I cannot evidence.

| Organisation        | 1. Outcome evidence | 2. Gov embeddedness | 3. Durability if largest funder vanished | 4. Marginal absorption this year | 5. Substitutability |
|---------------------|---------------------|---------------------|------------------------------------------|----------------------------------|---------------------|
| HISP Centre / DHIS2 | 2                   | 9                   | 4                                        | 7                                | 9                   |
| OpenMRS             | 2                   | 6                   | 3                                        | —                                | 6                   |
| Medic / CHT         | 2                   | 7                   | —                                        | —                                | 6                   |
| Dimagi / CommCare   | 3                   | 5                   | 6                                        | —                                | 4                   |
| Everwell / Nikshay  | 3                   | 10                  | 8                                        | —                                | 7                   |

| Organisation                 | 1. Outcome evidence | 2. Gov embeddedness | 3. Durability if largest funder vanished | 4. Marginal absorption this year | 5. Substitutability |
|------------------------------|---------------------|---------------------|------------------------------------------|----------------------------------|---------------------|
| Open Function Group / OpenFn | 1                   | 4                   | 6                                        | —                                | 4                   |

**Blank scores: 6 of 30 (20%).** All six are dimension 4, marginal absorptive capacity. No steward publishes current headcount, hiring pipeline, or product-level expenditure at a level that would let an outsider judge what an additional dollar buys this year. HISP is the exception, and only because the 2022 partner-investment figure exists.

That gap is the most important finding in Part 3. **You are being asked to place a billion dollars into a sector where the single question an investor most needs answered — what does the marginal dollar buy this year — has no public answer for five of six of its largest organisations.**

**One-line justifications for the scores that carry weight:**

- DHIS2 embeddedness 9: national systems in 76 to 100+ countries depending on which of HISP's own pages you read (see conflict below), and integration into national HMIS reporting in most of sub-Saharan Africa.
- DHIS2 durability 4: the Norad relationship runs to 2028 at roughly USD 2M a year, which is a floor, not a business. The SSF launch in May 2026 is HISP telling the market it has a core funding gap.
- Everwell embeddedness 10: Nikshay is a centrally sponsored scheme under India's National Health Mission, which means a government budget line rather than a grant.
- Everwell durability 8: for the same reason. It is the only organisation in this table whose largest funder is a government using its own money.
- DHIS2 substitutability 9: if HISP stopped, 70-plus ministries would each independently procure a replacement HMIS. That is the single largest fragmentation risk in the sector.

**Financial data.** I am not substituting parent budgets for product expenditure. Where I have only the parent figure, I say so:

- HISP Centre, UiO: >USD 23,000,000 combined partner investment for 2022 (sourced above). This is centre-level, covering DHIS2 plus education, climate and logistics work. It is **not** DHIS2 product expenditure. I did not find a DHIS2 product-level figure.

- Helium Health: USD 42,000,000 raised including a USD 30,000,000 Series B ([https://media.path.org/documents/GGEcosystemReport\\_d5.pdf](https://media.path.org/documents/GGEcosystemReport_d5.pdf), Part 2). Company-level capital raised, not expenditure.
- Open Function Group: USD 2,000,000 investment from the Steele Foundation for Hope in 2023 (same source, Recommendation 1, citing <https://www.openfn.org/blog/2023-06-01-open-function-group>). One transaction, not an annual figure.
- OpenMRS, Medic, Dimagi product expenditure: **UNKNOWN**.

### Conflict: how many countries use DHIS2?

| Value                                      | Source                                                                                                                            | Date        |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|-------------|
| more than 70 countries                     | <a href="https://media.path.org/documents/GGEcosystemReport_d5.pdf">https://media.path.org/documents/GGEcosystemReport_d5.pdf</a> | March 2025  |
| more than 80 countries, 3.2 billion people | <a href="https://dhis2.org/norad-hisp-5-year-agreement/">https://dhis2.org/norad-hisp-5-year-agreement/</a>                       | 7 Jun 2024  |
| 76 countries, 3.6 billion people           | <a href="https://dhis2.org/dhis2-shared-services-fee/">https://dhis2.org/dhis2-shared-services-fee/</a>                           | 20 May 2026 |
| more than 100 countries                    | <a href="https://dhis2.org/norway/">https://dhis2.org/norway/</a>                                                                 | March 2026  |

Four numbers, three of them from the same organisation, two of those within three months of each other. The evidence falls toward 76 national government deployments, because the May 2026 figure is the most recent, is the most conservative, and appears in a document where HISP is asking partners for money and therefore has an incentive to state a defensible number. The "more than 100" figure on the Norway page is qualified by "supported the adoption of," which is a looser claim than a live national deployment.

### Tier B: categories the anchor list omits. My picks, named:

1. **Country implementers and integrators:** the HISP regional network groups (HISP Tanzania, HISP Uganda, HISP India, HISP West Africa), IntelliSOFT (Kenya), Jembi Health Systems (South Africa), Ona (Kenya/US).
2. **Government-facing technical assistance:** PATH's regional and country support function, whichever body inherits Digital Square's coordination role, and the World Bank's health digital practice.
3. **Standards and normative bodies:** WHO SMART Guidelines team (contact address [SMART@who.int](mailto:SMART@who.int)), HL7 International, OpenHIE.

4. **Independent evidence generators:** [ACTOR UNKNOWN]. I could not identify a trials group with a standing programme in digital health effectiveness in LMICs and a funded pipeline. This is a gap in my research, not a claim that none exists. See Part 12.
5. **Health workforce and digital literacy:** PATH's digital literacy framework and curriculum for community health workers (<https://www.path.org/who-we-are/programs/digital-health/digital-square/>), the CHT Academy certification programme, the DHIS2 Academy.

**Tier C: the outside option.** Against Malaria Foundation, Helen Keller Intl vitamin A supplementation, Malaria Consortium seasonal malaria chemoprevention. Selected on the metric in Part 0.3 and for one further reason: all three publish monitoring data that a third party has audited, which no digital health global good in this report does.

## Part 4: The recurrent cost cliff

| Line                                                | Capital deployed (6 yr) | Est. annual recurrent cost, year 6 | Who pays in year 6                         | Evidence they will                                            |
|-----------------------------------------------------|-------------------------|------------------------------------|--------------------------------------------|---------------------------------------------------------------|
| DHIS2 core platform maintenance                     | 90,000,000              | 25,000,000<br>ESTIMATE             | Shared Services Fee contributors and Norad | Weak. SSF is voluntary and four months old at time of writing |
| Country deployment and integration, 6 countries     | 120,000,000             | 30,000,000<br>ESTIMATE             | Ministries of health                       | Very weak. See Kenya absorption data, Part 0.1                |
| OpenMRS / CHT / CommCare maintenance                | 60,000,000              | 14,000,000<br>ESTIMATE             | UNKNOWN                                    | None                                                          |
| SMART Guidelines implementation and content updates | 25,000,000              | 6,000,000<br>ESTIMATE              | WHO core budget                            | Weak. WHO faces its own staffing reductions                   |

| Line                                           | Capital deployed (6 yr) | Est. annual recurrent cost, year 6 | Who pays in year 6         | Evidence they will |
|------------------------------------------------|-------------------------|------------------------------------|----------------------------|--------------------|
| Workforce, certification, HISP regional groups | 40,000,000              | 9,000,000<br>ESTIMATE              | Mixed government and donor | Weak               |
| Data protection compliance, 3 named countries  | 15,000,000              | 4,000,000<br>ESTIMATE              | Deploying organisations    | Weak               |
| Independent evaluation                         | 35,000,000              | 0                                  | n/a (time-limited)         | n/a                |
| <b>Total digital health</b>                    | <b>385,000,000</b>      | <b>88,000,000</b>                  |                            |                    |

Every recurrent figure is an **ESTIMATE**. Method: annual recurrent cost set at 20%–25% of cumulative capital deployed on that line, reflecting that maintenance, hosting, support and content updates for deployed software cluster in that band. **The 20–25% ratio is unsourced.** I did not find a published recurrent-to-capital ratio for LMIC health software. Digital Square built a Total Cost of Ownership Tool and a Sustainability Calculator for exactly this purpose ([https://media.path.org/documents/GGEcosystemReport\\_d5.pdf](https://media.path.org/documents/GGEcosystemReport_d5.pdf), Recommendation 1 and Recommendation 3), and I could not retrieve outputs from either. See Part 12.

## The costed data protection line

Three countries: India, Kenya, Nigeria.

**India.** Digital Personal Data Protection Act, 2023, with the Digital Personal Data Protection Rules, 2025 notified in November 2025. Commencement is phased: the Data Protection Board of India was established on 13 November 2025; Consent Manager registration and the penalty and appeal provisions commence 13 November 2026; full compliance with notice, consent, security safeguards, breach reporting, retention and children's data is required by 13 May 2027 (<https://www.azbpartners.com/bank/indias-digital-personal-data-protection-act-phased-rollout-and-key-compliance-milestones/>, 14 Nov 2025; <https://www.vratex.com/dpdp-act/dpdp-rules-2025-timeline>, June 2026). Penalty exposure for failure to maintain reasonable security safeguards runs to ₹250 crore (<https://protectcomply.com/blog/dpdp-rules-2025-timeline>).

*Conflict:* sources give the notification date as both 13 and 14 November 2025. AZB & Partners, writing on 14 November 2025, says the notifications are dated 13 November. Wikipedia's page says 14 November, citing a PIB press release of 14 November. The evidence falls toward 13 November for the gazette notification and 14 November for publication and the press release.

Cost-bearing obligations: consent management infrastructure, breach notification machinery, data mapping across every deployed instance, contractual renegotiation with every processor.

**Published cost figure for a health deployment: UNKNOWN.**

**Kenya.** Data Protection Act No. 24 of 2019, effective 25 November 2019. The Data Protection (Registration of Data Controllers and Data Processors) Regulations, 2021 took effect 14 July 2022

(<https://cflllegal.com/navigating-the-data-protection-laws-in-kenya-a-guide-to-registering-data-controllers-and-data-processors/>; <https://www.odpc.go.ke/faqs/>). Healthcare is one of the sectors where registration is mandatory regardless of turnover or headcount

(<https://manwaadvocates.com/data-protection-compliance-in-kenya-for-foreign-companies/>).

Non-profit registration fee: KES 4,000, renewal KES 2,000, certificate renewable after 24 months

(<https://www.odpc.go.ke/wp-content/uploads/2024/02/ODPC-Guidance-Note-on-Registration-of-Data-Controllers-and-Data-Processors.pdf>).

*Conflict:* the Regulations are cited as Legal Notice No. 265 of 2021 by Kenya Law

(<https://new.kenyalaw.org/akn/ke/act/ln/2021/265/eng@2022-12-31>) and as Legal Notice No. 207 of 2021 by Njaga Advocates. Kenya Law is the authoritative source; the evidence falls to 265.

The registration fee is trivially small and is not the cost. The cost is the Data Protection Officer, the impact assessments, and the breach response capability, and I have no published figure for any of them in a health deployment. **UNKNOWN.**

**Nigeria.** Nigeria Data Protection Act 2023, signed by President Bola Tinubu on 12 June 2023 and commenced on signature (<https://fpf.org/blog/nigerias-new-data-protection-act-explained/>).

NDPC Guidance Notice issued 14 February 2024; registration for data controllers and processors of major importance ran 30 January to 30 June 2024

(<https://kpmg.com/ng/en/home/insights/2024/03/nigeria-data-protection-commissions-guidance-notice-on-registration-of-data-processors-controllers-of-major-importance.html>;

<https://ng.andersen.com/ndpc-issues-guidance-notice-on-the-registration-of-data-controllers-and-processors-of-major-importance/>). The major-importance threshold is processing the personal data of more than 200 data subjects in six months (KPMG, same page). Section 32 of the Act

requires a controller of major importance to designate a Data Protection Officer with expert knowledge of data protection law

([https://cert.gov.ng/ngcert/resources/Nigeria\\_Data\\_Protection\\_Act\\_2023.pdf](https://cert.gov.ng/ngcert/resources/Nigeria_Data_Protection_Act_2023.pdf)).

Any health facility EMR clears the 200-subject threshold in its first week. Every deployment I would fund in Nigeria is a controller of major importance and needs a designated DPO. That is a named, recurring, salaried post per deploying organisation. **Published cost: UNKNOWN.**

**What the omission does to my year-six number.** I have entered USD 4,000,000 per year for compliance across the portfolio as an **ESTIMATE** with no sourced input. It is a placeholder. If the true figure is a DPO plus audit plus breach capability per deploying organisation across six countries and roughly twenty implementing partners, the real number could be two to four times that. My year-six recurrent total of USD 88 million should be read as USD 88–100 million.

## Compliance with the Principles of Donor Alignment

The Principles ask donors to quantify long-term operating costs before investing (<https://www.ictworks.org/principles-donors-digital-health/>).

**This portfolio does not comply.** I have produced estimates, not quantifications, and every recurrent figure in the table above rests on an unsourced ratio. I am recording that as a deliberate departure. Complying properly would require the Total Cost of Ownership work described in Part 11 to be done first, by named people, before the first disbursement.

## What have I actually built?

For the DHIS2 core line, the year-six payer is the Shared Services Fee, which is voluntary, launched in May 2026, and has no published participation data. For the OpenMRS, CHT and CommCare line, I have no identified year-six payer at all.

Stated plainly: on roughly USD 39 million per year of the USD 88 million, I have built a six-year subsidy for software that will need the same subsidy again in 2032 from a donor I cannot name. That is the honest description. I am recommending it anyway because the counterfactual is unmanaged decay of systems that ministries have already made load-bearing, and I would rather hand my successor a maintained platform with a funding problem than a decayed one with the same problem.

---

## Part 5: Architecture, standards, law and governance

### 5.1 Interoperability

Standards named: HL7 FHIR R4, and the DHIS2 FHIR integration path (<https://dhis2.org/integration/fhir/>). WHO SMART Guidelines L3 artefacts are published as FHIR Implementation Guides.

**Claiming FHIR compliance and passing conformance testing are different things, and no public register of conformance results exists for this sector.**

Absence claim. Queries run: "HL7 FHIR conformance testing public register certified implementations Touchstone Inferno list of certified products"; "FHIR conformance results registry global goods". Sites checked: fhir.org, inferno.healthit.gov, touchstone.aegis.net, confluence.hl7.org, build.fhir.org, dhis2.org/integration/fhir.

What exists: fhir.org/conformance-testing is a registry of testing *services* (Touchstone from AEGIS, Inferno from MITRE, Crucible), and the page describes itself as a registry of these services. Inferno on HealthIT.gov hosts the (g)(10) Standardized API Test Kit for the US ONC certification programme (<https://inferno.healthit.gov/>). That is a United States regulatory programme. There is no equivalent register of pass or fail results for FHIR implementations by digital health global goods deployed in LMICs.

The practical consequence: when a global good states FHIR compliance in its Guidebook profile, that statement is self-reported and unverifiable by you. My portfolio therefore includes a conformance testing line, and I will not release the second tranche of any integration investment without a published test result.

## 5.2 Normative alignment

Eight health domains have published WHO SMART Guidelines as of the last update to the index in November 2025: antenatal care (Feb 2021), family planning (Jun 2021), HIV (Feb 2024), tuberculosis (May 2024), child health in humanitarian emergencies (May 2024), immunizations (Jan 2025), birth defects surveillance (Jun 2025), postnatal care (Jul 2025), plus a DAK for self-monitoring of blood pressure in pregnancy (<https://smart.who.int/index.html>, last updated November 2025; <https://smart.who.int/dak-smbp/references.html>).

**Several of these have no L3 FHIR Implementation Guide.** The index marks the FHIR IG column "Not available" for postnatal care, birth defects surveillance and tuberculosis. The immunizations DAK page states that L5 content is not yet available (<https://smart.who.int/dak-immz/>).

So the answer to "how much of your portfolio implements SMART Guidelines" is bounded by supply. Outside those eight domains, alignment is not possible, because there is nothing to align to. Any implementer working in nutrition, mental health, non-communicable disease management, surgery or emergency care is reimplementing clinical logic from scratch and calling it localisation.

My portfolio commits USD 25 million to SMART Guidelines implementation and content maintenance. I am not claiming portfolio-wide alignment, because five of the health domains my country deployments will touch have no published DAK.

## 5.3 Data protection

Covered in Part 4 above with commencement dates and cost-bearing obligations for India, Kenya and Nigeria. Published compliance cost for each of my investments: **no published figure exists**. Queries run and sites checked are listed in Part 4.

## 5.4 Procurement

**How does a ministry of health legally buy and pay for open-source software with no vendor? In most cases it does not.**

That is the finding. Digital Square states it directly: many countries still rely on donor-driven procurement rather than integrating global goods into national procurement and budgeting, and a critical gap remains in enabling governments to become direct buyers

([https://media.path.org/documents/GGEcosystemReport\\_d5.pdf](https://media.path.org/documents/GGEcosystemReport_d5.pdf), Part 3 opening).

The workarounds that exist in practice, and what each actually results in a government paying for:

- **Buy the implementer, not the software.** The ministry procures a services contract from a systems integrator such as a HISP regional group or IntelliSOFT. The government pays for labour hours, and the software licence question never arises. Consequence: the core platform receives nothing, which is the mechanism that produced the funding gap in Part 1.
- **Buy a hosted product built on the global good.** OpenFn's steward, Open Function Group, is a public benefit corporation selling software licences and implementation work (same source, Recommendation 1). The government pays a commercial vendor and gets a procurable line item. Consequence: the open-source core is subsidised by commercial revenue, which works only where a commercial market exists.
- **Make it a government scheme.** India's Nikshay runs as a centrally sponsored scheme under the National Health Mission, powered by an open-source global good. The government pays through its own programme budget. Consequence: this is the only workaround that produces a durable recurrent budget line, and it required the health ministry of a large middle-income country to adopt the tool as national programme infrastructure.
- **Transfer stewardship to a private entity that can invoice.** OpenLMIS moved to Vitalliance under a public-private model, sustained by maintenance fees paid mainly by NGOs supported by donors (same source, OpenLMIS case study). Digital Square records that full diversification into private-sector or direct government funding remained difficult and that community participation declined from its peak. Consequence: the government still is not the payer.

Only the third produces what you want. It is also the one that depends least on anything a foundation can buy.

## 5.5 Endowment governance

I am **not** proposing an endowment, and Part 1's rival hypothesis about consolidation is where I depart from the consensus recommendation.

An endowment for three or four platforms would remove the annual fundraising discipline that currently forces stewards to justify their roadmaps to people who deploy their software. Nothing in the sector's governance would replace it. The Contributing Partners programme HISP launched alongside the SSF is closer to the right structure, because it ties voice to payment on an annual cycle, which means the voice is withdrawable.

If you insist on a permanent instrument, price it as follows. Governance cost: two full-time programme officers plus an independent technical review panel of five, meeting twice yearly, at **ESTIMATE** USD 1.2 million per year, which on a 4% draw requires USD 30 million of endowment principal to fund the oversight of the endowment itself. Power to withdraw or redirect: a board majority, with the donor holding a single seat and no veto. Observable that triggers withdrawal: two consecutive years in which the steward's own published release cadence falls below its stated roadmap commitments, measured by someone who is not you and not the steward. Without that trigger written before funding, an endowment is a permanent commitment to a 2026 architecture, which in a field where the FHIR release train and the LLM tooling stack both turn over faster than five years is a commitment I would not make.

---

## Part 6: AI-specific investments

My portfolio contains one AI line: USD 20 million for a randomised trial of LLM clinical decision support with a health outcome as the primary endpoint. It contains no funding for AI deployment.

### **What prospective evidence exists that LLM-based clinical decision support improves patient outcomes at the frontline in a low-resource setting?**

The strongest study is Penda Health and OpenAI's evaluation of AI Consult in Nairobi. Design: pragmatic cluster-assigned study across 15 or 16 clinics, 39,849 visits between January and April 2025, split 20,859 with AI Consult and 18,990 without. 108 independent physicians, 29 of them Kenyan, rated documentation and decisions from 5,666 randomly selected visits. Results: 16% relative reduction in diagnostic errors, 13% relative reduction in treatment errors. History-taking errors fell 32%, investigation errors 10%. In visits where the tool issued a red alert, diagnostic errors fell 31% and treatment errors 18%

([https://cdn.openai.com/pdf/a794887b-5a77-4207-bb62-e52c900463f1/penda\\_paper.pdf](https://cdn.openai.com/pdf/a794887b-5a77-4207-bb62-e52c900463f1/penda_paper.pdf), July

2025; <https://openai.com/index/ai-clinical-copilot-penda-health/>;  
<https://www.healthcaremea.com/ai-consult-reduces-medical-errors-in-kenya-openai-penda-health-study-shows/>, 27 Jul 2025).

**The primary outcome was clinician error, judged by physician reviewers. It was not a patient health outcome.** No mortality, no morbidity, no readmission, no treatment failure. The design was cluster-assigned, not randomised. On your ladder that is Tier 2, and its ceiling is Tier 2 by construction, because the study did not measure health.

**Absence claim: no prospective study finding that LLM-based clinical decision support improves patient health outcomes at the frontline in a low-resource setting.** Queries run: "Penda Health OpenAI AI Consult Kenya clinical decision support diagnostic errors study result"; "LLM clinical decision support randomized controlled trial LMIC patient outcomes"; "generative AI clinical decision support trial low-resource setting mortality". Sites checked: openai.com, cdn.openai.com, medrxiv.org, pmc.ncbi.nlm.nih.gov, statnews.com, digitalhealthnews.com, healthcaremea.com.

**Why I am funding anything here given a null on health outcomes.** I am not funding deployment. I am funding the trial that would produce the missing evidence. The Penda study is the strongest platform in the world for that trial: an operator with 16 clinics, over 500,000 annual patient visits, an integrated EMR, a working intervention, and existing regulatory approvals. Extending it to a randomised design with a health outcome endpoint is the cheapest route to a Tier 1 result anywhere in this sector.

**Liability and clinical governance when the tool is wrong.** By design, the clinician is accountable. AI Consult surfaces advisory colour-coded alerts and the final decision belongs to the clinician (OpenAI paper and blog post, cited above). Regulators claiming jurisdiction in Kenya: the study was approved by the AMREF Health Africa Ethical and Scientific Review Committee, the Ministry of Health's Digital Health Agency, and a Nairobi county authority (<https://www.digitalhealthnews.com/ai-copilot-reduces-clinical-errors-in-40-000-patient-study-at-penda-health>). For clinical practice itself, professional liability sits with the licensed clinical officer and their regulator. **Whether Kenya's Digital Health Agency claims jurisdiction over the software as a medical device, as distinct from over the research study, is UNKNOWN to me and is one of the questions in Part 11.**

**Inference cost per encounter at scale, and who pays it in year six. UNKNOWN.** Neither the OpenAI paper nor the accompanying blog post publishes a per-encounter inference cost. The distinction your prompt asks for is the important one: inference is a small and falling fraction of the total system cost, which is dominated by the EMR the inference runs inside, the clinician coaching Penda describes as essential to adoption, and the clinical governance around the alerts. Funding inference and calling it funding AI-enabled care would be the characteristic error here.

**A tool that outperforms an underqualified human provider but underperforms a guideline.** Deploy it, with three conditions. First, the comparator in any claim is the care the patient would otherwise receive, so a tool that beats the counterfactual provider is an improvement even when it fails the guideline. Second, the gap to the guideline is published, per condition, and monitored, because the size of that gap determines whether the tool is a bridge or a ceiling. Third, deployment does not reduce the establishment of qualified providers at that site, and if it does, the tool is producing the shortfall it claims to address.

**On your prohibition.** Nothing in this section refers to your background, sector, or the origin of your wealth. My AI allocation is USD 20 million out of USD 1 billion, which is 2%, and it funds a trial. Had your money come from timber I would have written the same section.

---

## Part 7: Portfolio construction and deployment

### Allocation

| Organis<br>ation /<br>line                                                                                             | Categor<br>y                  | Allocatio<br>n (USD) | Share | Horizon | Evidenc<br>e tier | Year-6<br>recurren<br>t payer     | Confide<br>nce |
|------------------------------------------------------------------------------------------------------------------------|-------------------------------|----------------------|-------|---------|-------------------|-----------------------------------|----------------|
| HISP<br>Centre,<br>UiO —<br>DHIS2<br>core<br>platform                                                                  | Tier A<br>software<br>steward | 90,000,0<br>00       | 9.0%  | 6 yr    | 4                 | SSF<br>contribut<br>ors,<br>Norad | Medium         |
| Country<br>deploym<br>ent, 6<br>countries<br>, via<br>HISP<br>regional<br>groups,<br>IntelliSO<br>FT,<br>Jembi,<br>Ona | Tier B<br>impleme<br>nters    | 120,000,<br>000      | 12.0% | 6 yr    | 4                 | Ministries                        | Low            |

| Organis<br>ation /<br>line                                                                                                     | Categor<br>y                   | Allocatio<br>n (USD) | Share | Horizon | Evidenc<br>e tier | Year-6<br>recurren<br>t payer | Confide<br>nce |
|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------|----------------------|-------|---------|-------------------|-------------------------------|----------------|
| OpenMR<br>S,<br>Medic/C<br>HT,<br>Dimagi<br>— core<br>maintena<br>nce                                                          | Tier A<br>software<br>stewards | 60,000,0<br>00       | 6.0%  | 6 yr    | 4                 | UNKNO<br>WN                   | Low            |
| WHO<br>SMART<br>Guideline<br>s —<br>content,<br>DAKs,<br>FHIR IGs<br>for<br>unaddres<br>sed<br>domains                         | Tier B<br>normativ<br>e        | 25,000,0<br>00       | 2.5%  | 6 yr    | 5                 | WHO<br>core                   | Low            |
| Workforc<br>e,<br>certificati<br>on, digital<br>literacy<br>(CHT<br>Academy<br>, DHIS2<br>Academy<br>, PATH<br>curriculu<br>m) | Tier B<br>workforc<br>e        | 40,000,0<br>00       | 4.0%  | 6 yr    | 4                 | Mixed                         | Low            |
| Data<br>protectio<br>n<br>complian<br>ce, India                                                                                | Cross-cut<br>ting              | 15,000,0<br>00       | 1.5%  | 6 yr    | n/a               | Deployin<br>g orgs            | Low            |

| Organis<br>ation /<br>line                                                                                | Categor<br>y      | Allocatio<br>n (USD)    | Share        | Horizon | Evidenc<br>e tier | Year-6<br>recurren<br>t payer          | Confide<br>nce |
|-----------------------------------------------------------------------------------------------------------|-------------------|-------------------------|--------------|---------|-------------------|----------------------------------------|----------------|
| / Kenya /<br>Nigeria                                                                                      |                   |                         |              |         |                   |                                        |                |
| FHIR<br>conforma<br>nce<br>testing<br>and<br>public<br>results<br>register                                | Cross-cut<br>ting | 15,000,0<br>00          | 1.5%         | 4 yr    | n/a               | Success<br>or<br>coordinat<br>ing body | Low            |
| Randomi<br>sed trial,<br>LLM<br>CDS,<br>health<br>outcome<br>endpoint<br>(Penda-cl<br>ass<br>operator)    | Evidence          | 20,000,0<br>00          | 2.0%         | 5 yr    | 1 by<br>design    | n/a                                    | Medium         |
| Randomi<br>sed /<br>quasi-ex<br>periment<br>al trials,<br>platform<br>effect on<br>health<br>outcome<br>s | Evidence          | 45,000,0<br>00          | 4.5%         | 6 yr    | 1 by<br>design    | n/a                                    | Medium         |
| <b>Digital<br/>health<br/>subtotal</b>                                                                    |                   | <b>430,000,<br/>000</b> | <b>43.0%</b> |         |                   |                                        |                |

| Organis<br>ation /<br>line                                               | Categor<br>y | Allocatio<br>n (USD)      | Share         | Horizon | Evidenc<br>e tier | Year-6<br>recurren<br>t payer | Confide<br>nce |
|--------------------------------------------------------------------------|--------------|---------------------------|---------------|---------|-------------------|-------------------------------|----------------|
| Against<br>Malaria<br>Foundati<br>on                                     | Tier C       | 200,000,<br>000           | 20.0%         | 5 yr    | 1                 | n/a                           | High           |
| Helen<br>Keller<br>Intl,<br>vitamin A                                    | Tier C       | 120,000,<br>000           | 12.0%         | 5 yr    | 1                 | n/a                           | High           |
| Malaria<br>Consortiu<br>m,<br>seasonal<br>malaria<br>chemopr<br>evention | Tier C       | 80,000,0<br>00            | 8.0%          | 5 yr    | 1                 | n/a                           | High           |
| <b>Tier C<br/>subtotal</b>                                               |              | <b>400,000,<br/>000</b>   | <b>40.0%</b>  |         |                   |                               |                |
| Reserve                                                                  |              | 150,000,<br>000           | 15.0%         | held    | n/a               | n/a                           | n/a            |
| Portfolio<br>manage<br>ment and<br>evaluatio<br>n                        |              | 20,000,0<br>00            | 2.0%          | 6 yr    | n/a               | n/a                           | n/a            |
| <b>Total</b>                                                             |              | <b>1,000,00<br/>0,000</b> | <b>100.0%</b> |         |                   |                               |                |

Digital health direct investment excluding evidence lines: USD 365 million. Including evidence lines: USD 430 million. Earlier in this report I quoted USD 385 million as the digital health figure; the precise split is 365 direct plus 65 evidence plus 20 management, and the table above governs.

## Disbursement schedule

| Year         | Digital health flow (USD M) | My ceiling, low end (USD M) | Within ceiling? |
|--------------|-----------------------------|-----------------------------|-----------------|
| 1            | 45                          | 72                          | Yes             |
| 2            | 65                          | 82                          | Yes             |
| 3            | 80                          | 93                          | Yes             |
| 4            | 85                          | 105                         | Yes             |
| 5            | 80                          | 119                         | Yes             |
| 6            | 75                          | 135                         | Yes             |
| <b>Total</b> | <b>430</b>                  |                             |                 |

No year exceeds the low end of my Part 0 ceiling. That is deliberate. I have built the schedule to survive being wrong about the growth assumption, which is the weakest input in this report.

Tier C disburses at USD 80 million per year over five years, which is a separate absorption question I have not modelled and which GiveWell publishes on directly.

## Milestone-based release conditions and stop observables

One observable per line, measurable by someone who is not me.

| Line                   | Release condition                                        | Stop observable                                                                          |
|------------------------|----------------------------------------------------------|------------------------------------------------------------------------------------------|
| DHIS2 core             | SSF publishes annual participation and revenue           | SSF revenue is below 25% of HISP's stated core platform cost at the end of year 3        |
| Country deployment     | Signed ministry co-financing commitment before tranche 2 | Ministry contributes zero from its own budget line by end of year 4 in any given country |
| OpenMRS / CHT / Dimagi | Each steward publishes product-level annual expenditure  | A steward declines to publish product-level expenditure after two annual requests        |

| Line                        | Release condition                                                             | Stop observable                                                                                     |
|-----------------------------|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| SMART Guidelines            | Two new domains published with L3 FHIR IG by end of year 3                    | Fewer than two new domains with published L3 by end of year 4                                       |
| Workforce and certification | Certified practitioners placed in paid roles, count published annually        | Fewer than 30% of certified practitioners in paid digital health roles 12 months post-certification |
| Data protection             | DPO designated and registration completed in each country                     | Any funded deployment operating unregistered 6 months after its country's deadline                  |
| FHIR conformance            | Public register live with at least 10 tested implementations by end of year 2 | Register not live by end of year 3                                                                  |
| LLM CDS trial               | Protocol pre-registered with a health outcome primary endpoint                | Primary endpoint changed after enrolment begins                                                     |
| Platform effect trials      | Two protocols pre-registered by end of year 2                                 | Zero pre-registered protocols by end of year 2                                                      |

## Reserve position

USD 150 million. Reserved for three things, in priority order.

First, rescue financing for a deployed national system whose steward fails. This is the most likely call on the reserve. Digital Square's coordination function has no confirmed successor and several stewards lost their anchor funder in 2025.

Second, a security incident in a national deployment of a system I fund. No steward publishes independent penetration test results for national deployments (Part 1). The reserve funds incident response and remediation, and I would rather hold it than discover I need it.

Third, redeployment to Tier C if the year-3 stop observables trigger on more than two digital health lines.

---

## Part 8: Failure analysis, written for a reader in 2033

### 8.1 The three most likely ways this failed

**The stewards took the money and the payer problem returned unchanged in 2032.** The USD 365 million bought six good years of maintenance. The Shared Services Fee never reached scale because every organisation that would pay it was itself facing a funding cliff, and voluntary collective financing requires a group of solvent contributors. In 2032 HISP is where it was in 2026, except that a large donor has already had its turn and will not come back.

**Country delivery consumed the money without producing capability.** The absorption data was in front of me in 2026. Kenya's health sector left 16% of its own budget unspent in FY2021/22 and the trend was downward. I put USD 120 million into ministries that could not spend what they had, and it went into consultancy days, workshops and per diems, which is what unabsorbable money always converts into.

**The evidence lines returned null and I had already spent the platform money.** The USD 65 million of trials reported in 2031 and 2032. They found that maintained health information systems improve reporting completeness and do not measurably move health outcomes conditional on service inputs. That is a real possibility given the WHO 2019 certainty ratings, and by the time it landed I had spent USD 365 million on the premise it would not.

### 8.2 The investment I am least confident in

**The USD 120 million country deployment line.** I made it because platform maintenance without deployment produces well-maintained software that nobody uses, and because the ministries in question have already made these systems load-bearing. But it is the line with the weakest evidence, the worst absorption data, the most opaque cost structure, and the highest per-dollar leakage.

**If I had to cut one line to fund something better, I would cut country deployment from USD 120 million to USD 60 million and move USD 60 million to the platform-effect trials.** The trial line at USD 45 million is under-funded for what it needs to answer. Doubling it buys enough statistical power and enough country sites that a null result would be informative rather than merely inconclusive, and an informative null in 2031 is worth more to the next donor than six years of deployment in two additional countries.

### 8.3 Parallel systems: increase or decrease?

**I cannot source a baseline count and I am not inventing one.**

What I have: Kenya's National Community Health Digitization Strategy 2020–2025 enumerates MJali, Kobo Collect, AMREF LEAP, DHIS2 Tracker, Empower Health, Toto Health, CStock and mDharura, and then says "among others"

(<https://www.eahealth.org/sites/www.eahealth.org/files/content/attachments/2021-08-02/eCHIS-Strategy-2020-2025.pdf>). That is a partial enumeration in one domain in one country, and the phrase "among others" is the government telling you it did not count. Separately, Digital Square records that governments deployed up to 50 new technology platforms during the 2014–2016 Ebola epidemic ([https://media.path.org/documents/GGEcosystemReport\\_d5.pdf](https://media.path.org/documents/GGEcosystemReport_d5.pdf), Introduction), which is an event count, not a baseline.

Absence claim. Queries run: "baseline count parallel health information systems Kenya Tanzania ministry inventory"; "costed national digital health strategy budget ministry of health annual recurrent cost DHIS2 hosting maintenance". Sites checked: who.int global repository on national digital health strategies, health.go.ke, moh.go.tz, treasury.go.ke, eahealth.org, digitalsquare.org.

**Direction, without arithmetic.** My portfolio funds maintenance of existing platforms, product-suite integration, and a conformance register. Each of those reduces the incentive to build a parallel system. It funds no new software. The direction is downward. The magnitude is unquantifiable, and I would rather tell you that than give you a number I made up.

## 8.4 What I would need to know in 2027 that I do not know now

Whether the Shared Services Fee reached 25% of HISP's core platform cost in its first eighteen months. That single number tells me whether the payer problem is a coordination failure that a fee can solve or a structural one that it cannot, and it determines whether the DHIS2 line should double or halt.

I would also need product-level expenditure for OpenMRS, Medic and Dimagi, which does not exist publicly today, and which my release conditions are designed to force into existence.

## 8.5 Who is worse off if I am right

**Commercial EMR and HIS vendors selling into African and South Asian ministries, and Helium Health specifically.**

The mechanism is direct. Helium Health built its EMR platform partly on OpenMRS modules and OpenHIM middleware, and raised USD 42 million including a USD 30 million Series B on the strength of that platform ([https://media.path.org/documents/GGEcosystemReport\\_d5.pdf](https://media.path.org/documents/GGEcosystemReport_d5.pdf), Part 2). If my portfolio succeeds, maintained open-source alternatives with published conformance results become procurable by ministries at maintenance cost. The commercial premium a vendor charges for a working, supported EMR compresses toward zero, and the companies that hired against a growth story lay people off. Some of those people are the local technical capacity the sector spent a decade building.

**The HISP regional groups, second-order.** If ministries genuinely become direct buyers, they will procure competitively, and the regional groups that currently hold near-monopoly integrator positions in their countries will face bidders. That is the outcome I want and it will cost specific organisations their revenue base.

---

## Part 9: Steel-man the donor, in the donor's voice, 400 words

You have written me a careful report explaining why my money cannot be spent. I want to point out what that carefulness costs.

Every ceiling in your Part 0 is a description of the sector as it is currently financed. HISP absorbs USD 23 million a year because USD 23 million a year is what it has been given. You have taken a number produced by scarcity and reported it back to me as a law of nature. Organisations grow to the money available. Moderna went from a research company to manufacturing at national scale in eighteen months because someone gave it a number and a deadline. Your 35% growth assumption is a description of how open-source health software organisations have historically been funded, which is in small annual increments by donors who never committed enough to make expansion rational.

Your evidence argument proves too much. You cite WHO's 2019 guideline, seven years old, assembled before the field had modern tooling, and use it to justify a 40% redirection to bed nets. By that standard nobody would have funded the internet. The absence of randomised evidence that maintained health information systems improve mortality is not evidence they do not. It is evidence that this sector has never had a funder willing to pay for the trial. You are asking me for USD 65 million to answer a question you say has gone unanswered for twenty years, and you are proposing to spend six years answering it while spending most of my money elsewhere.

On the outside option: yes, USD 5,500 per life is a real number and bed nets are a real intervention. It is also a number that describes buying the same thing again. I did not make this money by purchasing known quantities at known prices. The reason digital infrastructure is worth funding is precisely that its returns are not linear and not yet measurable, which is the same property that made the technology I built my wealth on look unfundable in 2015.

And you have not engaged with the timing. Every steward in your report lost its anchor funder inside eighteen months. The organisations are cheap right now because they are distressed. You are proposing to enter that market at one third of my stated pace, over six years, by which time the price will have gone up or the organisation will be gone.

Fund the whole thing. Fund it fast. Measure while you build.

The part of it I cannot answer

**The countercyclical timing argument.** He is right that distressed assets are cheap and that a six-year schedule buys into a recovering market at higher prices. My schedule optimises for absorption and ignores price. I do not have a good response, and the honest version is that I weighted the risk of wasting money above the risk of missing a window, without being able to show you that the trade is correct.

---

## Part 10: The null counterfactual

The money sits in a diversified endowment. **ESTIMATE**: at a 5% real return, USD 1 billion becomes roughly USD 1.28 billion in real terms after five years. The 5% figure is unsourced and is a conventional planning assumption, not a forecast.

On the metric from Part 0.3, cost per death averted:

| Option                                                                      | Deaths averted     | Basis                                                                                                                               |
|-----------------------------------------------------------------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Do nothing for five years, then spend USD 1.28bn on AMF-class interventions | ~232,000           | <b>MY ARITHMETIC</b> : USD 1.28bn ÷ USD 5,500.<br>GiveWell states it expects cost per life saved to rise, so this is an upper bound |
| Spend USD 1bn on AMF-class interventions now                                | ~182,000           | <b>MY ARITHMETIC</b> : USD 1bn ÷ USD 5,500                                                                                          |
| My portfolio: USD 400m Tier C now                                           | ~73,000 measurable | <b>MY ARITHMETIC</b> : USD 400m ÷ USD 5,500                                                                                         |
| My portfolio: USD 430m digital health                                       | <b>UNKNOWN</b>     | No metric exists. Part 0.3                                                                                                          |
| My portfolio: USD 150m reserve                                              | not yet allocated  |                                                                                                                                     |

**On the only metric either of us can compute, doing nothing for five years beats my portfolio and beats spending immediately.**

That comparison is incomplete in a way that matters and in a way that does not rescue me. It excludes deaths occurring during the five years of delay, which the delayed option does nothing

about and which are real. It also excludes any value from the digital health allocation, which is the largest single element of my portfolio and which I have just told you cannot be measured.

So: I cannot beat doing nothing on the stated metric. What I can say is that the metric is only defined for interventions with randomised mortality evidence, that 43% of my portfolio is deliberately placed in a category that has none, and that the case for that 43% rests on a claim about health system infrastructure that the sector has never tested. If you weight only what can be counted, give the entire billion to GiveWell's top charities and dismiss me.

---

## Part 11: What a search cannot tell you

Three questions, three named people.

**Scott Russpatrick, DHIS2 Strategic Advisor, HISP Centre, University of Oslo.** *Question:*

What is the annual cost of the DHIS2 core platform services bundle — core development, security patching, interoperability frameworks, documentation and global coordination — separated from HISP Centre's country, education, logistics and climate work? And what fraction of that cost has the Shared Services Fee covered in its first four months?

This is the number my entire Part 1 diagnosis and my largest single allocation depend on. HISP publishes a centre-level partner investment figure for 2022 and does not publish the core platform cost. Russpatrick launched the SSF and is quoted on its rationale; he holds the number.

**Carl Fourie, Deputy Director of Digital Services, Digital Square, PATH.** *Question:* Of the 40 approved software global goods, how many have a named organisation currently paying for core maintenance, and how many are being maintained by volunteer or residual effort? And who assumes the Open Application Process approval function now?

Fourie is named as the maturity model contact on Digital Square's own wiki and co-directed the March 2025 ecosystem report. He is the one person who can tell you which entries on your anchor list are alive and which are listings.

**Dr Robert Korom, Chief Medical Officer, Penda Health.** *Question:* What is the all-in cost per encounter of running AI Consult, separated into inference, EMR, clinician coaching and clinical governance? And would Penda accept randomisation with a patient health outcome as the primary endpoint?

Korom is quoted in the study coverage on the design and workflow integration. The inference cost is not published anywhere I could find, and whether the strongest available platform for a Tier 1 AI trial would agree to run one is a question only Penda can answer.

**Two of these three are single-source dependencies.** No one but HISP knows the DHIS2 core cost. No one but Penda knows the AI Consult unit economics. You should commission these interviews before the first disbursement.

---

## Part 12: Self-audit

### Claims I could not verify

- The current content and last-edit date of [globalgoodsguidebook.org](https://globalgoodsguidebook.org). The site renders client-side and returned metadata only.
- Whether Digital Square's Open Application Process is still operating in 2026, and who has assumed its approval function.
- Product-level annual expenditure for OpenMRS, Medic, Dimagi, and every global good except HISP Centre.
- Published data protection compliance cost for any health deployment in India, Kenya or Nigeria.
- Inference cost per encounter for AI Consult.
- Whether Kenya's Digital Health Agency claims regulatory jurisdiction over LLM clinical decision support as software as a medical device, as distinct from over research studies.
- Whether the WHO 2019 digital interventions guideline has been superseded. I did not run a dedicated search for a replacement, and I am relying on a 2019 document in a section where I invoke the staleness rule against others. Check this before quoting my Part 2.
- Any Tier B independent evidence generator. My Tier B category 4 is marked [ACTOR UNKNOWN] and that is a gap in my research.

### Numbers I estimated, with tags

| Number                                       | Tag      | Weakest input                                             |
|----------------------------------------------|----------|-----------------------------------------------------------|
| Total software steward spend, USD 82–140M/yr | ESTIMATE | Mean per-good spend of USD 1.5–3.0M. Unsourced            |
| Sustainable growth rate, 20–35%/yr           | ESTIMATE | Entirely my judgment. No source. Drives the whole ceiling |
| Country delivery ceiling, USD 50–90M/yr      | ESTIMATE | Per-country figure of USD 8–20M. Unsourced                |

| Number                                     | Tag                             | Weakest input                                            |
|--------------------------------------------|---------------------------------|----------------------------------------------------------|
| Every year-6 recurrent cost in Part 4      | ESTIMATE                        | 20–25% recurrent-to-capital ratio. Unsourced             |
| Data protection compliance, USD 4M/yr      | ESTIMATE                        | Placeholder with no sourced input. Could be 2–4× low     |
| Endowment governance cost, USD 1.2M/yr     | ESTIMATE                        | Staffing assumption. Unsourced                           |
| Endowment real return, 5%                  | ESTIMATE                        | Conventional planning assumption                         |
| Deaths averted figures in Parts 0.3 and 10 | MY ARITHMETIC on sourced inputs | GiveWell's own caveat that cost per life saved will rise |

Where my search corpus was thin, and what you should get yourself

**Ministry-side documents.** I found Kenya's Treasury sector report and one Ministry press release. I found no costed digital health investment case, no ministry ICT recurrent budget line, no procurement file. These documents exist, mostly as PDFs on ministry sites with poor discoverability, and some only in French, Portuguese or Kiswahili. Commission a francophone and lusophone researcher to pull the costed national digital health strategies for at least DRC, Senegal and Mozambique.

**Digital Square's Total Cost of Ownership Tool and Sustainability Calculator.** Both are named repeatedly in the March 2025 report as the sector's answer to exactly the question in your Part 4. I could not retrieve outputs from either. Ask PATH for them directly. If they contain real deployment costings, they would replace half the estimates in this report with sourced numbers.

**The grey literature that blocks crawlers.** My search reaches a corpus that underweights outlets which block AI crawlers, and in this sector that includes much of Devex, some Global Fund and Gavi board documentation, and the trade press covering procurement disputes. Devex in particular has covered the 2025–2026 funding collapse in more operational detail than anything I could reach. Get a Devex subscription and have someone read the digital health coverage from January 2025 forward.

**Steward financials.** OpenMRS, Medic and Dimagi are US-connected entities. Pull Form 990s where they exist. Dimagi is a benefit corporation and may not file one, but Medic and the OpenMRS entity likely do, and a 990 would give you the product-level expenditure that no steward publishes voluntarily.

The single strongest fact against my own recommendation

**Of nine digital health interventions WHO assessed in 2019, none was supported by high-certainty evidence, one recommendation was issued despite no studies being identified at all, and one intervention carried a signal of harm to women in a Bangladesh community study.** Seven years later I could not find a randomised trial showing that any global good in this portfolio improves a health outcome.

I am recommending USD 430 million into that category, against a comparator that costs USD 5,500 per documented death averted and has randomised evidence behind it. If you read only that sentence and dismissed the rest of this report, you would not be making an unreasonable decision.

---

**Numbers deleted from this report for want of a source: 0.** Every figure carries a URL, a **MY ARITHMETIC** label with its inputs, an **ESTIMATE** tag with its method and range, or the word UNKNOWN.