The Water for Health Fund

Cost-Effective Health Impact, from Safe Water that Pays Its Own Way

Sensor-Verified Safe Water, with Third-Party Revenue Driving Donor Cost per Averted DALY Toward Zero

Avery Bang, MBA
Proposed Chief Fund Manager and Deal Flow Manager
Senior Principal, Mulago Foundation

Evan Thomas, PhD, PE, MPH, MBA
Proposed Co-Fund Manager and Program Director
Professor and Director, Mortenson Center in Global Engineering, University of Colorado Boulder
Founder & CEO, Virridy

July 2026

Prepared for Renaissance Philanthropy, Fund Ecosystem Program

1. Summary and the Ask

Four billion people drink water contaminated with fecal matter (Greenwood et al. 2024), a leading driver of the diarrhoeal disease that still kills hundreds of thousands of children under five every year. Delivering safe drinking water is one of the best-evidenced and most cost-effective interventions in global health: in the team's own Rwandan programmatic context, a cluster-randomized controlled trial measured a 29% reduction in under-five diarrhoea (Kirby et al. 2019). The **Water for Health Fund** exists to buy that health impact at the lowest possible donor cost, and its primary metric is the marginal philanthropic cost per averted disability-adjusted life year (DALY). What drives that cost down is not cheaper delivery alone; it is third-party revenue that pays for delivery so that philanthropy does not have to, indefinitely rather than through repeated grant cycles. Carbon finance is the proven example of such revenue, and the fund's near-term engine: where safe water displaces boiling over biomass fuels, credit revenue under the Gold Standard safe drinking water methodology can sustain programs across delivery modalities, household and institutional filtration, chlorination and ultraviolet treatment, and community, borehole, and piped safe-supply systems. In practice, that revenue underperforms for a specific, documented reason. A credit requires evidence that the treatment is used and that the delivered water is safe. Usage is self-reported and systematically overstated, while the water-quality testing the methodology already requires is too infrequent, too small in sample, and too expensive to confirm that water is safe continuously or to scale across a portfolio. The resulting evidence is weak, so credits are discounted, capital is deterred, and programs remain dependent on recurring grants, which is precisely what keeps the donor cost per averted DALY high. The fund is a time-bound (roughly five-year), catalytic fund built on a single thesis: one bottleneck, the integrity and cost of verifying safe water, keeps third-party revenue, carbon finance first among it, from sustaining safe drinking water, and clearing it is a field-level public good that only philanthropy will fund. By resolving that bottleneck and the financing barrier that follows from it, the fund lets outside revenue rather than recurring donations carry safe drinking water at scale, drives the marginal donor cost per averted DALY toward zero, then exits.

We are requesting **$30 million** to capitalize the Water for Health Fund, deployed in two tranches: **$10 million into the existing footprint immediately**, the seven safe-water carbon programs across six countries described below, and **$20 million to bring new programs, implementers, and delivery models** to creditable, financeable status. The request is unusual in how much is already proven: the underlying method, the Lume in-situ fluorescence sensor developed at the University of Colorado Boulder, is peer-reviewed and has pilot-level regulatory acceptance; commercial capital is already engaging on the production side, with the Bridges SDG Outcome Fund in advanced due diligence on $10M of debt project finance that pairs with this request to form a roughly $40M blended structure; and the proposing team already runs seven safe-water carbon programs across six countries, reaching some 2.79 million people through more than 1,270 continuously monitored water points, with $25M in carbon offtake under contract. What this capital funds is the layer neither the carbon market nor repayable debt will finance: the catalytic, first-loss capital that de-risks the carbon projects themselves and brings them to creditable, financeable status, and the Lume development, validation, and regulatory acceptance that turn the proven method into an accepted verification standard across these projects, together letting commercial finance and credit revenue sustain safe drinking water at scale. Every element serves one objective, health impact per philanthropic dollar: because the fund's capital is a one-time input that unlocks a recurring outside revenue stream, each dollar buys averted DALYs at a marginal donor cost that falls toward zero as third-party revenue takes over, and the same continuous measurement that makes the revenue creditable is what evidences the health outcome. At its core this is the translation of high-technology university research into the market: the Lume began as peer-reviewed, patented science at the University of Colorado Boulder, and the fund carries it across the gap between proven research and an accepted, market-deployed verification standard, the gap that neither research grants nor commercial capital is built to cross. The fund pairs independent fund management with operating depth: Avery Bang, Senior Principal at the Mulago Foundation, is proposed as chief fund manager and deal flow manager, with Evan Thomas as co-fund manager and program director.

2. The Problem: A Financing Barrier Rooted in Measurement

An estimated 4.4 billion people in low- and middle-income countries lack safely managed drinking water, with fecal contamination identified as the primary limiting factor (Greenwood et al. 2024). Diarrhoeal disease remains a leading cause of death in children under five, with roughly 444,000 such deaths annually (WHO 2024), of which an estimated 297,000 are attributable to inadequate water, sanitation, and hygiene (Prüss-Ustün et al. 2019). Closing this gap through grants alone is implausible: meeting the drinking-water Sustainable Development Goal requires on the order of $37.6 billion per year in capital investment, roughly three times current levels (Hutton and Varughese 2016). Carbon finance is one of the few mechanisms capable of durably supplementing public and philanthropic funding at that scale, which is what makes the integrity of water credits consequential rather than academic. Safe drinking water interventions, whether household treatment, institutional programs, or community and piped supply, address a substantial share of this burden. Where they displace boiling of water over biomass fuels, they also avoid carbon dioxide and black-carbon emissions: roughly 1.2 billion people boil drinking water (Floess et al. 2025), and these avoided emissions are the accounting basis for crediting under the Gold Standard safe drinking water methodology.

The barrier is not the intervention; it is the integrity and cost of the two measurements the credit depends on. First, usage is typically self-reported, and self-reports overstate actual practice substantially: in a study of carbon-financed water filters in Kenya, reported filter use was 81% while measured use was approximately 19%, and assumed baseline boiling far exceeded observed practice (Pickering et al. 2017); over-reporting of comparable WASH behaviors is documented at 50 to 60% (Contzen et al. 2015). Second, water quality is measured, not merely assumed, and the Gold Standard methodology requires evidence of microbially safe water at the point of consumption. But that measurement is infrequent, drawn from a small sample of sites, and expensive: culture-based laboratory testing costs on the order of $12 to $15 per sample, requires cold-chain logistics and 24 to 48 hours to a result, and is typically performed once or twice a year. Sparse, periodic sampling cannot confirm that water is safe in the intervals between tests, precisely when treatment failures and recontamination occur, and it is too costly to run continuously at portfolio scale. The evidentiary basis for the credit is therefore weak even where the required testing is done. The consequence for the market is measurable: for cookstove credits, the household-device class most analogous to water, a systematic review found issued credits over-credited by an average factor of 9.2 (Gill-Wiehl et al. 2024).

This produces a self-reinforcing financing barrier. Weak, sparse, and costly measurement creates integrity concerns; integrity concerns depress credit prices and constrain issuance; depressed prices and constrained issuance keep commercial capital out; and without that capital, programs remain dependent on grants and safe water is under-delivered. The measurement layer that would break this cycle generates no revenue on its own, which is why it is not financed by the commercial actors who would benefit from it. It is a public good, and its absence is the constraint this fund is designed to relieve.

3. The Keystone: Continuous, Objective Verification

The technical basis for resolving the measurement failure is established. Tryptophan-like fluorescence (TLF) detects the optical signature of microbial metabolites and functions as a real-time, in-situ indicator of faecal contamination in drinking water (Sorensen et al. 2015). The Lume, developed at the University of Colorado Boulder, implements time-resolved TLF on a silicon photomultiplier in a reagent-free, field-deployable sensor that returns a microbial reading in approximately 60 seconds, against 24 to 48 hours for culture-based laboratory methods, with a tryptophan detection limit below 0.1 parts per billion. Because it measures the treated water itself rather than the treatment device or the user's behavior, the method applies across delivery modalities, filtration, chlorination, ultraviolet, and borehole or piped supply alike, which is what allows a single verification standard to serve the entire credit class. The sensor is patented (US 11,506,606 B2; 11,507,861 B2) and at technology readiness level 7. In September 2025, Gold Standard granted a pilot-level methodology deviation (Pilot 14) for TLF-based digital monitoring, the first such approval for an IoT sensor in this context.

The Lume is field-hardened and already deployed across the team's programs in Sub-Saharan Africa, where it is solar- or battery-powered, cellular-connected, and takes automated readings on a fixed cadence and transmits them remotely, building an auditable, timestamped, continuous record of water safety at the point of consumption rather than a handful of annual laboratory snapshots. Because each reading uses the same installed device, the marginal cost of an additional measurement approaches zero, against the $12 to $15 per sample and cold-chain logistics of every culture-based test. And because the method reads the treated water itself, a single sensor gives an objective safety signal that is directly comparable across a portfolio and independent of who implemented the treatment or how a user reports their behavior. This is what converts a periodic, self-reported, device-specific credit into one priced on continuous, objective performance data.

The validation record is already substantial. Against culture-based reference methods across field and laboratory sites, the Lume reaches 85% balanced accuracy on the drinking-water safe/unsafe decision (76% sensitivity, 93% specificity), classified 100% of chlorinated Kenyan borehole samples correctly (n=30), and achieved 96.8% binary classification accuracy on the Seine River in Paris. On a three-level exceedance classifier spanning all validation sites (n=90), overall accuracy is 0.86 with a quadratic-weighted κ of 0.70. When trained against a specific laboratory target, the sensor's quantitative agreement is strong and method-agnostic: $R^2 = 0.88$ against Colilert (n=209) and $R^2 = 0.87$ against membrane filtration (n=303); its categorical agreement with Colilert ($\kappa = 0.88$, "almost perfect") in fact exceeds the agreement between the two EPA-approved laboratory methods with each other ($\kappa = 0.40$). The sensor design and validation are published in *Water Research* (Knopp et al. 2026), and the method is advancing through U.S. EPA and Gold Standard certification.

In-situ sensing measures water quality continuously and at low marginal cost, removing the cold-chain logistics and per-sample expense of culture-based testing and closing the gap between infrequent laboratory samples where failures go undetected. It replaces both unreliable self-report and sparse, expensive periodic sampling at the point where the credit is calculated. The fund's asset is not the sensor but the accepted, open verification standard the sensor makes possible: once that standard is proven, published, and adopted by validators, the credit class can be priced on evidence rather than assumption. This is consistent with the broader finding that the verification requirement itself acts as an accountability mechanism, prompting implementers to identify and correct treatment failures they had not previously detected (Thomas et al. 2025).

4. What the Fund Does: Four Catalytic Functions

The fund is not a single-instrument grant program for sensors. It provides the set of catalytic public goods and de-risking capital required to make high-integrity, carbon-financed water treatment function as a market. It deploys along four functions and is designed to exit as credit revenue becomes self-sustaining.

1. **De-risking the carbon projects.** A leading use of the fund is bringing safe-water programs to the point where they are creditable and can attract commercial finance: catalytic first-loss capital that absorbs early losses on forward carbon revenue (GIIN 2013), and the upfront project development the credit itself will not pay for in advance, baselining, monitoring build-out, and registration. Blended-finance transactions have historically mobilized on the order of $4 of commercial capital per $1 of concessional capital (Convergence 2024; OECD 2018). This is the pre-commercial risk that repayable debt will not take; it is what lets debt like the production facility, and credit revenue itself, carry the projects afterward.

2. **Lume development, validation, and acceptance across the projects.** The Lume's core method is proven, but making it the accepted, crediting-grade basis for these projects is a substantial, distinct body of work and the fund's other leading use: developing and hardening the sensor for the full range of field conditions across six countries and every delivery modality, paired sensor-versus-laboratory validation at portfolio scale, and driving the Gold Standard Pilot 14 deviation to a full digital-MRV methodology, with the open standard and registry integration that let credits be priced on measured performance. This is project- and method-specific work, not Virridy's commercial product roadmap; it returns no cash and benefits the whole credit class, which is why no commercial actor funds it. It is the classic translation of high-technology university IP into the market, moving a peer-reviewed, patented CU Boulder invention from proof to an accepted, revenue-bearing standard, the step that neither the university nor a lender will finance.

3. **Methodology and standards development.** Work with Gold Standard and other registries to codify digital MRV and strengthen the integrity of the safe-water credit class across modalities, a benefit that accrues to every implementer, not only to the fund's grantees.

4. **Implementer capacity and demand aggregation.** Support water-treatment implementers in accessing carbon markets, and aggregate buyer demand for independently verified, high-integrity water credits.

The fund does **not** finance the water systems themselves; those are paid for by the credit revenue the fund unlocks. It does not finance any implementer's commercial operations.

5. Fund Architecture: Separation of Commercial and Philanthropic Capital

The governing design principle is that philanthropic capital never subsidizes a commercial venture. Two separate vehicles, with no shared expenditures, keep the capital legible to donors and auditors.

	Production entity (commercial)	**Water for Health Fund (philanthropic)**
Instrument	Repayable debt and equity, return-seeking	Grants and, potentially, first-loss capital; no market return
Use of funds	Working capital to produce and deliver credits	Verification infrastructure, methodology, de-risking, capacity
Beneficiary	The production entity's balance sheet	The credit class and its implementers, sector-wide
Implementers	The operator's own programs	Multiple implementers; the operator is one of several

The integrity test is that no expenditure is ever charged to both vehicles. This $30M request anchors a roughly $40M blended structure. Alongside the philanthropic anchor, the **Bridges SDG Outcome Fund** is providing $10M of debt project finance into the production entity and is in advanced due diligence; together with existing equity, this is the commercial capital that produces and sells the credits. The production entity's own financial model projects, at full portfolio scale, roughly 7.6 million credits produced and sold and on the order of $176M in gross credit revenue through 2033. This is the fund thesis in practice: commercial lenders like Bridges will finance credit production once the verification barrier is addressed, which confirms that the verification public good, and the pre-commercial project risk, are the specific layers commercial capital will not fund. Critically, that commercial debt funds working capital against already-contracted, de-risked production; it will not, and by its nature cannot, take the earlier and uncertain risk of bringing new projects and implementers to creditable, financeable status, nor fund the Lume development, validation, and methodology acceptance that the credit depends on. Both of those return no cash with which to repay a lender, and both benefit the whole credit class rather than only the entity paying for them. This is why philanthropic first-loss and public-good capital is required alongside the commercial debt, not in place of it: the fund carries projects through the pre-commercial risk, and repayable debt like the production facility takes over once they are financeable.

Because the proposed co-fund manager and program director also leads Virridy, the sensor developer and the largest single implementer, the fund's governance holds allocation implementer-neutral by construction. Fund management, capital allocation, and deal flow sit with the proposed chief fund manager, Avery Bang, Senior Principal at the Mulago Foundation, who is institutionally independent of Virridy and of every implementer. Virridy competes for fund-supported work on the same terms as every other implementer and receives no preferential access; the verification standard the fund builds is open and non-proprietary; and health and water-quality evaluation sits with institutionally separate partners (Section 8). The fund's asset is the field-wide standard, not any one company's position in it.

6. Cost-Effectiveness and Scale

Renaissance Philanthropy's threshold is cost-effective impact at scale: under \$100 per DALY averted or under \$1 per tonne of CO_2e reduced, at the order of 100,000 lives or one billion tonnes. The fund's primary metric is the first of these: the marginal philanthropic cost per averted DALY. Safe drinking water sits squarely in this territory. It is among the interventions GiveWell rates most cost-effective in global health, and the fund concentrates delivery on the highest-mortality group, children under five, where the burden of diarrhoeal disease is greatest. Third-party revenue enters the model as the term that drives the metric down: every dollar of outside revenue that pays for delivery is a dollar of recurring donor funding no longer needed, so the marginal donor cost per averted DALY falls as that revenue takes over. Carbon credit revenue is the proven, contracted example, and the carbon threshold below is the evidence that this revenue engine works; other third-party streams, results-based public finance and commercial water-quality revenue among them, follow the same logic. Because the fund's catalytic capital is a one-time input that unlocks a self-sustaining revenue stream, its cost-effectiveness improves with the duration of the programs it enables.
These figures are illustrative and are refined at fund structuring. They are grounded in the existing portfolio's unit economics, including a per-student program cost of approximately \$1.50 per year. Both of Renaissance Philanthropy's thresholds are evidenced by the same measurement layer: the continuous verification that lets a credit be priced on measured safe-water delivery is the very data that populates the cost-per-DALY model. The fund's carbon and health cost-effectiveness therefore rest on one objective evidence base rather than two independent sets of assumptions, which is precisely what the verification public good makes possible.

On scale, the portfolio already reaches approximately 2.79 million people today and, at full portfolio scale, is modeled to serve on the order of nine million people and to generate roughly 7.6 million credits (tonnes CO_2e) and about \$176M in gross credit revenue through 2033. Even that direct footprint is not itself at the billion-tonne order of Renaissance Philanthropy's threshold; the path to that order runs through the field-level public good. Because roughly 1.2 billion people boil drinking water, an open, accepted verification standard that lets the entire safe-water credit class be financed on measured evidence is what puts impact at that scale within reach, which is why the fund invests in the standard rather than only in its own programs.

7. Evidence Base

Portfolio-wide water quality. The core evidence is not from any single country. Across its programs Virridy applies one measurement protocol, consistent with the Gold Standard methodology: Compartment Bag Tests enumerating fecal contamination across a representative sample of water points in every project at least annually, together with continuous Lume monitoring in real time. Pooling the quantitative results across Rwanda, DR Congo, Kenya (borehole and school programs), Burundi, and Madagascar, 45% of samples tested unsafe and only 24% safe at baseline (n=4,943), whereas under monitored safe-water delivery 90% test safe (n=1,685). The improvement therefore holds across six national contexts and across delivery modalities, filtration, boreholes, kiosks, and community supply alike, not in one program or one country. **Randomized health evidence.**

Threshold	Fund (illustrative; refined at structuring)
<\$100 per DALY averted **(primary metric)**	The enabled interventions have documented health effects (Evidence Base). At a per-student program cost on the order of \$1.50 per year, the cost per person reached is very low, and because third-party revenue carries delivery, the marginal *donor* cost per averted DALY falls toward zero over the program's life. Continuous sensor verification supplies the single parameter GiveWell's cost-effectiveness model most depends on: the measured proportion of people actually consuming safe water, in place of an assumed one. A rigorous cost-per-DALY, built to GiveWell's own model, is a direct output of an independent, GiveWell-partnered evaluation now in design, which separates implementation (Virridy, CU Boulder) from measurement (Emory, Aquaya, University of Arkansas).
<\$1 per tonne CO_2e	The evidence that the revenue engine works: the existing portfolio produces credits at approximately \$8.45 per tonne against a weighted-average sale price of \$21.26 per tonne, so production is revenue-positive and self-sustaining. The philanthropic input is the one-time catalytic layer, not the per-tonne production cost; against a portfolio of the target scale, this corresponds to well under \$1 per tonne. The precise figure depends on methodology-defined per-household emission factors.
Leverage	This \$30M request is paired with \$10M of debt project finance from the Bridges SDG Outcome Fund, now in advanced due diligence, a roughly \$40M blended structure and a concrete instance of catalytic capital mobilizing commercial finance (Convergence 2024). The \$20M new-programs tranche extends that leverage across the wider class of implementers, where the same pattern applies at field scale.

The strongest causal evidence comes from Rwanda, where a cluster-randomized controlled trial of a large-scale water-filter and cookstove distribution found a 29% reduction in under-five diarrhoea (prevalence ratio 0.71, 95% CI 0.59 to 0.87) and a 25% reduction in child acute respiratory infection (Kirby et al. 2019), in the same programmatic context in which the team operates. The randomized signal is strongest for household filtration; the portfolio-wide monitoring above is what generalizes independently verified performance across the other modalities and countries the fund supports. **Economics.** A cost-benefit analysis of that Rwandan program estimated \$66.7M in benefits against an \$11.9M investment over five years, a benefit-cost ratio of approximately 5.6 (Barstow et al. 2019); these benefits are program-wide and are dominated by cookstove fuel and time savings rather than water alone. **Operating record.** The team runs seven active safe-water carbon programs across six countries (Rwanda, Kenya, DR Congo, Burundi, Tanzania, Madagascar), reaching approximately 2.79 million people through 1,273 continuously monitored water points, with \$25M in carbon offtake under contract with PetroChina, Louis Dreyfus, and Mortenson. **Commercial validation of the sensor.** Beyond the carbon programs, the Lume is deployed and evaluated by the water industry's leaders: Xylem, the largest water-technology company, is a Lume customer deploying sensors on the Potomac in Washington, DC; Veralto (Hach, In-Situ, OTT) holds a channel partnership with recreational water as the lead use case; and the U.S. Air Force Research Laboratory is evaluating the Lume for Air Force and Army water-security applications. This independent commercial demand confirms the method is real and adoptable beyond the fund's own programs;

what remains to be financed is the shared validation and acceptance that make it a crediting-grade standard across these projects, not whether the underlying approach works. **Regulatory.** The Gold Standard Pilot 14 methodology deviation (September 2025) establishes that a digital-MRV pathway is available; full methodology acceptance remains to be completed and is part of the fund's work.

8. Team and Qualifications

8.1 Fund Leadership

Avery Bang (proposed Chief Fund Manager and Deal Flow Manager) is a Senior Principal at the Mulago Foundation, where she leads sourcing and the fellows portfolio, finding, funding, and supporting high-performance organizations addressing global poverty and climate. Before joining Mulago in 2022, she was CEO of Bridges to Prosperity for more than a decade, scaling the organization to serve more than 1.4 million people across twenty countries and building it into a recognized model for infrastructure delivery in partnership with governments and funders. She holds an MBA from the Saïd Business School at the University of Oxford, an MS in Civil Engineering from the University of Colorado Boulder, and a BS in Civil Engineering and Studio Art from the University of Iowa. As chief fund manager she directs capital allocation and deal flow across implementers, and her institutional independence from Virridy and from every implementer anchors the governance separation described in Section 5.

Evan Thomas (proposed Co-Fund Manager, Program Director, and field strategist) has spent approximately fifteen years in global water security, sensing, and results-based finance. He is Professor of Civil, Environmental and Architectural Engineering and Director of the Mortenson Center in Global Engineering and Resilience at the University of Colorado Boulder, and Founder and CEO of Virridy. He holds a PhD in Aerospace Engineering Sciences, an MPH, and an MBA, and is a licensed Professional Engineer, with prior roles at NASA and DelAgua Health. He has authored more than 85 peer-reviewed publications and holds seven patents in wireless monitoring and water-quality sensing. He developed the first United Nations Clean Development Mechanism and Gold Standard programs to earn carbon credits for household water treatment, and has since built and now oversees a portfolio of safe-water carbon programs across six countries reaching some 2.79 million people, with $25M in carbon offtake under contract. He has served as Principal Investigator on federal water-sensing research and has led multiple randomized controlled trials of water and cookstove interventions in Rwanda, published in the peer-reviewed literature. Recent recognition includes the International Water Association Water and Development Research Award (2025) and the AEESP Global Outreach Award (2023).

8.2 Implementation, Carbon, and Technical Team

The fund draws on Virridy, an operating water-intelligence company (with the Mortenson Center at CU Boulder) that already runs the sensing and carbon programs described here. Virridy is organized as three coordinated entities that mirror the commercial/philanthropic separation the fund itself requires: **Virridy Digital** builds and sells the Lume sensor; **Virridy Carbon**, a Gold Standard–registered project developer, converts sensor-verified safe-water delivery into high-integrity credits under long-term offtake; and **Virridy Rwanda** runs in-country field operations. The group reported $5.9M in 2026 revenue, nearly three times the prior year, holds $25M in carbon offtake under contract, has 43 Lume sensors deployed across 15 customer sites, and is backed by non-dilutive funding from the NSF Convergence Accelerator and ASCEND Engine, NASA, USAID, and the Gates, Moore, and Walton foundations.

The technical and delivery team spans sensor engineering, data science, carbon methodology, and in-country operations. **Danny Wilson, PhD** (Chief Technology Officer, PhD Mechanical Engineering, UC Berkeley) leads Lume design and analytics; **Taylor Sharpe, PhD** and **Whitney Knopp** lead sensor hardware and validation and are lead authors on the TLF sensor design and validation published in *Water Research*; **Matt Ross, PhD** (Colorado State University) leads machine-learning analytics; **Laura MacDonald, PhD** leads program implementation and water-quality analysis; **John Ecklu, PhD** leads carbon-water research and methodology, including the peer-reviewed evaluation of verification as an accountability mechanism; **Jean Ntazinda**, who helped design Rwanda's Green Climate Fund programme and its Article 6 crediting framework, leads Rwanda climate finance; and **Lambert Mugabo** leads field implementation in Rwanda. Collectively this team designs methodologies, operates field deployments across Sub-Saharan Africa, manages the digital-MRV data pipeline, and interfaces with registries and validators.

8.3 Independent Evaluation and Scientific Partners

Because measurement independence is central to the thesis, health and water-quality evaluation is conducted by partners institutionally separate from implementation: **Matthew Freeman** and **Thomas Clasen** (Emory University), **Rachel Peletz** (Aquaya Institute), and **Corey Nagel** (University of Arkansas, biostatistics). This is the same independence structure adopted for the GiveWell-partnered pilot, and it is the model the fund would extend, ensuring that the parties measuring outcomes have no operational stake in the interventions they assess.

9. Why Now

The carbon market is in an integrity reckoning, and that is a tailwind for this fund rather than a threat. Buyers, standards bodies, and regulators are moving decisively toward high-integrity, evidence-backed credits and away from the self-reported, assumption-based crediting that produced documented over-issuance in analogous household-device classes (an average factor of 9.2 for cookstoves (Gill-Wiehl et al. 2024)). In that environment, objectively verified water credits are precisely the segment that holds and commands value. The standard for the safe-water credit class is being written now: Gold Standard's Pilot 14 (September 2025) is the first IoT-sensor digital-MRV approval in this context, and it opens a window to define that standard on continuous, objective measurement before low-integrity credits define the class instead. The enabling technology has just crossed field readiness, and commercial capital is already engaging on the production side. The one layer still unfunded is the public good, and funding it now is what determines whether the entire class is built on evidence or on assumption.

10. Risks and Mitigations

- **Methodology risk** (Gold Standard does not graduate the pilot to full methodology acceptance). The Pilot 14 deviation is already granted; the paired sensor-versus-laboratory validation is designed to registry evidentiary requirements, and the team developed the first UN Clean Development Mechanism and Gold Standard water-treatment methodologies.
- **Carbon-price and market risk** (credit prices fall). The existing portfolio carries \$25M of long-term contracted offtake with compliance and Article 6 buyers, which reduces spot-price exposure; high-integrity credits are the segment most insulated from the discounting that hits low-integrity supply.

- **Adoption risk** (implementers do not take up the standard). The verification standard is open and non-proprietary and lowers each implementer's own verification cost; the fund's capacity and demand-aggregation functions actively support uptake, and the benefit accrues to every implementer rather than to one.
- **Integrity risk** (the credits themselves are questioned). This is the problem the fund exists to solve: continuous objective measurement plus institutionally independent evaluation (Section 8) is the direct mitigation, and the fund's own credibility is staked on it.
- **Concentration risk** (over-reliance on one team or one device). The method is device- and modality-agnostic, the standard is vendor-neutral, and independent commercial adoption of the sensor by Xylem, Veralto, and the U.S. Air Force Research Laboratory shows the technology is viable beyond the fund's own programs.

11. The Ask

We are requesting a **$30 million** commitment from Renaissance Philanthropy, as anchor donor, to capitalize the Water for Health Fund, structured in two tranches. The first, **$10 million, deploys immediately into the existing footprint**: the seven operating programs across six countries, where monitoring, offtake, and implementers are already in place, applying the four functions of Section 4 to capital that can begin averting DALYs at once. The second, **$20 million, funds new programs and models**: bringing new implementers, geographies, and delivery models to creditable, financeable status through the same first-loss capital, project development, verification build-out, and methodology work, which is where the fund's field-level leverage compounds. Because the underlying science and the market are proven, and the theory of change (Sections 2 to 4) is evidenced rather than hypothetical, the capital de-risks projects and builds the shared standard, not fundamental research risk. The fund reports against one primary metric, the marginal donor cost per averted DALY, evaluated independently to GiveWell's own model, with carbon credit revenue as the first and largest of the third-party revenue streams that drive that cost down. We would welcome Renaissance Philanthropy's role not only as anchor donor but as fund-structuring and governance partner, convening the wider coalition of climate, water, and carbon-integrity funders around the thesis over the fund's term.

This allocation is indicative and is finalized with Renaissance Philanthropy at structuring; the balance between first-loss capital and grants, in particular, is set then. The two leading uses are de-risking the carbon projects, so that commercial finance and credit revenue can carry them, and the Lume development, validation, and acceptance that make the method a crediting-grade standard, the public good no commercial actor will finance. First-loss capital is sized to the commercial offtake it can mobilize.

Illustrative Implementers and Partners

Implementers, spanning household and institutional filtration, chlorination, community water systems, treatment pumps, kiosks, and piped supply: Virridy (school and household programs, Rwanda and Burundi), LifeStraw (household and school filtration, Rwanda and Kenya), WaterMission (community water systems, Tanzania), Helvetas (chlorinated water supply, Madagascar and Burundi), Millennium Water Alliance (treatment pumps, Kenya), Asili (water kiosks, DRC). **Standards and verification**: Gold Standard; independent validation and verification bodies. **Research**: University of Colorado Boulder.

References. Greenwood EE, Lauber T, van den Hoogen J, Donmez A, Bain RES, Johnston R, Crowther TW, Julian TR. Mapping safe drinking water use in low- and middle-income countries. *Science.* 2024;385(6710):784–790. Hutton

Use of funds (illustrative; refined at structuring)	Share
Tranche 1: the existing footprint, deployed immediately	**$10.0M**
De-risking the existing carbon projects: catalytic first-loss capital and upfront project development (baselining, monitoring build-out, registration)	$4.0M
Lume development, validation, and acceptance across the projects: field hardening, portfolio-scale paired validation, Gold Standard methodology acceptance	$3.5M
Methodology and standards development with Gold Standard and other registries	$1.0M
Implementer capacity and demand aggregation	$1.0M
Fund operations and governance	$0.5M
Tranche 2: new programs and models	**$20.0M**
First-loss capital, project development, verification build-out, and methodology extension bringing new implementers, geographies, and delivery models to creditable, financeable status; allocation set with the anchor donor at structuring	$20.0M
Total	**$30.0M**

G, Varughese M. *The Costs of Meeting the 2030 Sustainable Development Goal Targets on Drinking Water, Sanitation, and Hygiene.* Washington, DC: World Bank, Water and Sanitation Program; 2016. WHO. *Diarrhoeal disease* [fact sheet]. 2024. Prüss-Ustün A, et al. Burden of disease from inadequate water, sanitation and hygiene. *Int J Hyg Environ Health.* 2019;222(5):765–777. Floess E, et al. Health trade-offs of boiling drinking water with solid fuels: a modeling study. *Environ Health Perspect.* 2025;133(6):067021. Pickering AJ, et al. Climate and health co-benefits in low-income countries: a case study of carbon financed water filters in Kenya and a call for independent monitoring. *Environ Health Perspect.* 2017;125(3):278–283. Contzen N, et al. Over-reporting in handwashing self-reports. *PLOS ONE.* 2015;10(8):e0136445. Gill-Wiehl A, Kammen DM, Haya BK. Pervasive over-crediting from cookstove offset methodologies. *Nat Sustain.* 2024;7(2):191–202. Sorensen JPR, et al. In-situ tryptophan-like fluorescence: a real-time indicator of faecal contamination in drinking water supplies. *Water Res.* 2015;81:38–46. Knopp W, Klaus J, Wilson D, Vlah M, Ross M, Thomas E. Advancing continuous in-situ quantification of microbial contamination in environmental waters using tryptophan-like fluorescence: sensor design and validation. *Water Res.* 2026;302:126099. Kirby MA, et al. Effects of a large-scale distribution of water filters and natural draft rocket-style cookstoves on diarrhea and acute respiratory infection: a cluster-randomized controlled trial in Western Province, Rwanda. *PLoS Med.* 2019;16(6):e1002812. Barstow CK, et al. A cost-benefit analysis of livelihood, environmental and health benefits of a large scale water filter and cookstove distribution in Rwanda. *Dev Eng.* 2019;4:100043. Convergence. *State of Blended Finance 2024.* OECD. *OECD DAC Blended Finance Principles.* 2018. GIIN. *Catalytic First-Loss Capital* [issue brief]. 2013. Thomas et al. Digital monitoring, reporting, and verification technologies supporting carbon credit-generating water security programs. *ACS ES&T Letters.* 2025.