



INTELLIGENT INNOVATION ACCELERATED

GREENLIGHT PLATFORM 2.0 · PARTNER SELECTION RESPONSE

A trusted platform for **trusted data**, built by a partner who already knows your world

riivo's response to The Greenlight Office Partner Selection Brief: our understanding of your mission, our Discovery approach, our initial technical point of view, and our commercial terms for Stage 1 and beyond.

Prepared for
The Greenlight Office (TGO)
Attn: Heather Jooste

Prepared by
riivo · riivo.io
Nic Bowman · nic@riivo.io · 083 995 1779

Submission
Final
30 July 2026

1 Greenlight context

Greenlight helps organisations move from intuition to evidence and from activity to impact. Platform 2.0 is the operational backbone that ambition now needs, and it has one job it can never fail at: protecting the **deep, contextual, trusted data** your whole value proposition rests on.

What we understand about your mission and methodology

South Africa has more than a quarter of a million registered non-profits. Funding flows to many that cannot show what actually changed for the people they serve, while some genuinely effective organisations struggle to raise money because they have no credible way to evidence their value or connect with the right funders. Greenlight sits precisely on that fault line. It gives funders confidence that money follows results, and gives strong organisations the evidence they need to grow. Platform 2.0 is how that capability scales: infrastructure, in a very real sense, for improving the impact of the NPO sector across South Africa.

The track record backs it up: TGO has grown into a partner to 175+ organisations across South Africa and Sub-Saharan Africa, holding over 110,000 household-level assessments. The methodology is multidimensional: participants self-diagnose against red, yellow and green indicators, and baseline, second and third survey cycles reveal not only what changed but why, and whether it will last.

We have also read your ambition carefully, because it shapes the architecture. TGO is explicit that the goal is not to become the country's largest measurement organisation, nor simply to produce more reports and complete more surveys. The goal is to help organisations build the capability to improve their impact over time, inside a learning ecosystem where implementing organisations, funders, practitioners and other stakeholders all contribute to and draw from a continuous cycle of evidence, learning and improvement. A platform serving that ambition cannot be a reporting tool with a portal bolted on; it has to carry the whole ecosystem.



The ecosystem Platform 2.0 serves. Every group above touches the platform; the architecture has to work for all of them at once.

We also understand that Greenlight has deliberately evolved beyond measurement. Your investment case is clear that you are now a **learning partner**: trusted evidence is the beginning of a contracted, collaborative learning journey over an eighteen-to-twenty-four-month lifecycle, not the end of a reporting process. That is why the member journey, not the survey, is the primary organising structure of Platform 2.0, and why journey management cannot remain outside the platform in Airtable.

Finally, we understand the sector reality: member organisations operate under real financial and operational pressure, facilitators work in low-connectivity environments on modest devices, and users span a very wide range of technological comfort. A platform that is technically excellent but high-friction would fail your Joyful Journey principle just as surely as bad data would fail Trusted Data.

Why riivo is well suited to this project



We already work inside your data

YEARS OF CONTEXT · NO RAMP-UP · LOWER DISCOVERY RISK

- ◆ We built the automated report preparation your previous provider could not deliver
- ◆ Extensive project work with your actual data model, survey structures, indicator logic and report templates has left us intimately acquainted with all of them
- ◆ That depth of insight into the entities Platform 2.0 must rebuild gives our Discovery a meaningful head start

Report automation built

Data model fluency

Known integrity issues mapped



AI-native delivery, grounded in real engineering discipline

CLAUDE PARTNER PROGRAM · MICROSOFT PARTNER · BUBBLE PARTNER

- ◆ Progressing through Anthropic's Claude partner program; certified Microsoft Partner and Bubble agency partner
- ◆ AI-built applications deliver more platform for the same budget, without compromising code, data or infrastructure integrity
- ◆ Paired with an outcome focus: fixed prices, we own the result, the target is the delivery, not the activity
- ◆ Part owned by Saratoga, one of Cape Town's largest enterprise software companies, for enterprise-grade depth on call

AI-accelerated build

Enterprise backing via Saratoga

Production systems at scale



We have already proven the partnership behaviours you are asking for

PARTNER, NOT SUPPLIER · PROVEN UNDER PRESSURE

- ◆ When your Microsoft licensing needed taking over during an awkward transition, riivo stepped in and stabilised it without drama
- ◆ Your brief asks for a thinking partner, not a vendor who executes instructions: that is how we already behave with TGO

Microsoft licensing steward

Swift, supportive delivery

Long-term relationships

Genuine commitment to the NPO sector

MISSION ALIGNMENT · LONG-TERM SUPPORT

- ◆ Giving back has always been core to our mission: we believe part of what good technology should deliver is uplifting the underprivileged
- ◆ Longstanding partnership with Fisantekraal Centre for Development: we migrated FCD from SharePoint onto the Power Platform, modernised their operation, and support them today
- ◆ Both FCD and TGO receive discounted project rates: our contribution to the sector, and a commitment that predates this proposal

FCD: modernised + supported

Discounted NPO project rates



The heart of our case is simple. Platform 2.0 will stand or fall on data integrity and journey architecture. riivo arrives already inside Greenlight's data, having built the automated preparation of Greenlight's reports, and already tested as a partner when things got awkward. Our Discovery starts at week three, not week one.

Aspects of the brief we would want to clarify in Discovery

Most open questions (anonymous surveys, iOS priority, multilingual, Power BI) are addressed in Sections 2 and 3; two we would pin down early:

- ◆ **Migration value versus effort per entity.** The brief invites a dialogue on the value of migration given current integrity. We would want to agree, per entity, whether we migrate, cleanse-and-migrate, or archive: for example, facilitator records are flagged as probably not migrated, and we would test whether all 18 million survey option rows carry forward or are consolidated.
- ◆ **Interim stewardship of the current platform.** The incumbent has resigned the account and management must move before 30 September 2026. riivo is willing and able to take this on as a separately quoted interim service, exactly as we did with your Microsoft licensing.


2 Discovery approach

Discovery has one purpose: to leave TGO with certainty on how this build will play out. A validated data model and migration plan. An architecture that has been demonstrated, not just described. Wireframes tested with your real users. A detailed delivery timeline. And a fixed-price build scope the board can commit to with confidence. Together these de-risk the entire build; we are rigorous where rigour de-risks, and iterative everywhere else.

How riivo runs Discovery

We run an **eight-week, three-phase Discovery**, sitting at the efficient end of your eight-to-twelve-week window. We can compress the calendar because we are not learning Greenlight from scratch: we already know your data model, your report structures and your team. Where most vendors would spend weeks one and two orienting, we spend them validating.

A distinctive feature of our Discovery: we do not only produce documents. By week five we put a **clickable prototype, including the journey configuration screens themselves**, in front of your staff, members and facilitators, because your brief rightly demands that configurability is demonstrated rather than inferred. We design with AI-assisted tooling (Claude and Lovable), which is why a tested prototype fits inside an eight-week Discovery at all, and wireframes are tested twice, with real users, before anything is priced.



The 8 weeks at a glance. Every artefact is tested with your people before it is priced.

Phase A · Immerse & validate (weeks 1-2)

Session	With whom	Produces
Kickoff, governance & success criteria	CEO, Heather, Kate, riivo leads	Charter, decision cadence, risk log
Member journey mapping workshops (x2)	Hub Admins, Member Journey Manager	As-is and to-be journey maps; stage-gate inventory
Member & facilitator voice sessions	2-3 member Champions; facilitators; Training Ambassador	Friction map across user roles and comfort levels
Airtable structural deep-dive	Heather, Kate, operations team	Full inventory of views, fields, automations to preserve
Data landscape validation	Heather + riivo data engineering	Profiling of the MySQL estate against the June 2026 audit; dedupe rule candidates

Phase B · Shape & demonstrate (weeks 3-5)

Activity	With whom	Produces
Target data model & integrity rules	riivo, validated with Heather	Single-source-of-truth model; ID uniqueness, validation and audit design
Journey architecture working sessions	Hub Admins, Kate	Template, stage, gate and cohort model; configurability spec
Report generation architecture	Heather + riivo	Automation design for all survey and report types; time-to-report targets
Wireframes & clickable prototype	AI-assisted design: Claude, Lovable	Staff console, member portal, mobile flows and the journey configuration screens

User testing round 1	Staff, Champions, facilitators	Scored usability findings feeding revision
----------------------	--------------------------------	--

Phase C · Validate & price (weeks 6–8)

Activity	With whom	Produces
Migration plan validated against profiled data	riivo data engineering	Entity-by-entity migration plan and risk assessment for the ~240GB estate, Python data and Airtable journeys, grounded in source profiling
Prototype refinement & user testing round 2	Same user panel	Validated UX baseline for build
Build backlog & fixed pricing	riivo, reviewed with TGO	Full backlog of user stories with acceptance criteria, point-sized; fixed build price; change control
Critical path & release plan	riivo + TGO leads	Sequenced plan to 1 March 2028 go-live
Board-ready readout	Board, CEO, project leads	Full Discovery pack handover; go / no-go basis for Stage 2

How Discovery outputs become a technical specification and build scope


Every workshop output lands in a single living specification, not a pile of slide decks. Journey maps become journey templates in the configuration model. Configurability scenarios become acceptance criteria. The prototype becomes the UX contract. The data profiling becomes migration rules with named owners. In week six we decompose the specification into a point-sized backlog of user stories, each with testable acceptance criteria, and price it. The result is the fixed scope, fixed price and critical path your brief requires: defensible because every line traces back to something your team saw, tested or signed off during Discovery.

UX design in Discovery versus build

In Discovery, UX goes deep enough to de-risk, not to decorate: journey and process maps for every role, low-fidelity wireframes for all primary flows, a clickable prototype of the highest-risk journeys (survey session lifecycle, report request, journey configuration), and two structured user-testing rounds across the full range of technological comfort. In the build phase, high-fidelity design is produced sprint by sprint from the validated wireframe baseline, with usability checks at each release. This keeps Discovery honest and the build fast.

Journey architecture and configurability in Discovery

Your proposed separation of Sales Journey, Programme Journey, cohorts and survey cycles is, in our view, exactly right, and it maps cleanly onto a configurable workflow engine. Discovery gives it a dedicated workstream: we model journeys as data (templates, stages, sub-stages, entry and exit rules, gates), then test the model by walking your staff through all six scenarios from your Discovery Input Pack on the prototype's configuration screens. Section 3 walks through each scenario in detail.



Your proposed hierarchy, as we would build it. Journeys are configurable data, not code: stage gates (the check above) block a report until its survey session is closed, and staff can add stages, cohorts and cycles themselves.

What we need from your team

Who	Involvement across the 8 weeks
Heather Jooste (Project Lead / Data SME)	Roughly one day per week, weighted to weeks 1-3 and 6-7
Kate Bain (Systems & Integration SME)	Roughly one day per week, weighted to journey and Airtable workstreams
CEO	Kickoff, two checkpoint sessions, board readout
Hub Admins / operations team	Four workshops plus two testing rounds
Member Champions & facilitators	One voice session and two prototype testing rounds each (we schedule around their availability)

TGO's brief commits dedicated resources, swift decisions and access to data and stakeholders; the plan above stays within that commitment.


Fixed fee and duration

riivo prices outcomes, not effort: we scope in build points, a measure of the size of a result rather than the time behind it, and every engagement in this proposal is fixed price. Our fixed fee for this Discovery is **R165 000 excluding VAT**. It follows our Discovery: Premium model, a senior, multi-stakeholder engagement producing a board-ready pack, extended here to eight weeks with a tested clickable prototype and a full user-story backlog because this project deserves them. In AI-accelerated delivery, the thinking up front matters more than the build that follows it; this Discovery is where the project is de-risked. Half is payable on commencement and half on handover of the artefacts, and if TGO proceeds to the build with riivo, a quarter of the fee (R41 250) comes back as a credit against the build. Duration is **eight weeks** from 1 September 2026, completing by 30 October 2026.

→ From Discovery to delivered: how riivo runs the whole journey



Indicative programme timeline



The pace is deliberate. The AI-accelerated build completes by mid-2027, months ahead of a conventional team; and since Airtable is already separate, the journey engine can go live and retire it in Q2 2027 without waiting for cutover. What sets the calendar after that is not our velocity: your team's absorption, training, migration rehearsals repeated until the reconciliation evidence is boring, and a cutover timed to survey-cycle boundaries. The window opens from October 2027; TGO picks the date, with 1 March 2028 the fully buffered ceiling.

3 Technical approach (high level)

What follows is our initial point of view, built from the brief, the Discovery Input Pack, the two audits and our working knowledge of your data. Discovery exists to test and refine it, and we hold it accordingly: firmly, but not preciously.

Technology stack: a full rebuild on a modern, licence-free core

The June 2026 audits are unambiguous: the current database is structurally unsound and the platform fails at 6-10 concurrent users. We agree with your conclusion to rebuild, not remediate, and we would recommend a full technology change rather than working within the Spring Boot / Angular / MySQL estate:

Layer	Our proposed direction	Why
Database	PostgreSQL, single well-architected store: Azure Database for PostgreSQL, or Supabase as a managed platform option	Strong constraint enforcement (unique IDs, foreign keys), row-level security, proven at hundreds of millions of rows. Supabase adds platform governance and managed auth for a modest subscription; Discovery weighs it against Azure-native
Backend & API	TypeScript service layer with a clean REST/JSON API	One API serves web, portal, mobile and the Python pipeline; API-first from day one
Web platform & portal	Next.js (React), mobile-responsive	Fast, accessible, WCAG-friendly, huge talent pool, no vendor lock-in

Mobile app	Expo (React Native) for a rebuild, or Flutter to evolve the current app	One codebase for iOS and Android either way; a rebuild keeps the whole stack TypeScript. Path confirmed in Discovery
Hosting	Microsoft Azure, South Africa region, auto-scaling, managed by riivo	Data residency for POPIA, auto-scaling kills the concurrency ceiling, and our Microsoft partnership keeps it well-governed
Delivery method	AI-accelerated engineering (Claude-based), human-reviewed, fully tested	More platform per rand; every line still owned by TGO, readable, documented and conventional

And this is not a paper stack: SurgiFlow (Section 4) runs on it in production today. Nothing above carries per-user licensing, so facilitator growth never triggers licence cost; the recurring platform costs are hosting and, if selected, the flat Supabase subscription. Either way TGO keeps full governance and South African data residency. And any competent successor vendor could take the codebase over, which is exactly how IP ownership should feel.

The core architectural challenge: a large legacy database with integrity failures

Our stance: **never build the new platform on the old data's assumptions**. We design the target model first, from the business domain (participants, households, cohorts, journeys, surveys, indicators), with integrity enforced in the schema itself: unique participant IDs, foreign keys everywhere, mandatory audit columns, validation at point of entry. The legacy database is then treated purely as a source to be mined, cleansed and mapped into that model, never as a template for it. Performance is addressed structurally: composite indexing around real query patterns, read models for dashboards, caching (absent entirely today, per the audit), and pagination and streaming so a 3,000-survey export never times out again.

Configurability without developer dependency

Configurability is your most important non-functional requirement, so here is our thinking against each of your six scenarios. The unifying principle: **journeys, surveys, views, rules, notifications and roles are data, not code**. Staff edit them through admin screens; nothing below requires a deployment. We have built precisely this class of engine before, most recently for a wealth-management client (delivered through a partner): triggers, stages and steps, with rules for which items carry over between stages and which must complete in strict order. Our recent survey design work for Nomios covered the other half: survey structures that hold up as they evolve.


Your scenario	How our architecture handles it
Add a 'Pre-contracting' stage with its own task list and entry criteria, plus a filtered pipeline view	Journey templates are editable records: staff insert the stage, attach a task checklist and entry criteria in the template editor, then save a filtered view of members by pipeline stage. Existing journey instances are versioned so history stays intact.
Change a stage-advance rule to 'session closed AND data quality score above 80%'	Stage gates are rule expressions built in a condition builder (field, operator, value, AND/OR). Data quality score is a first-class computed field, so this rule is a two-minute edit.

Add a new household indicator without a code deployment, syncing to the mobile app	Survey definitions live in the database with explicit versioning. Staff add the indicator in the survey designer; the mobile app pulls definitions on sync. Versioning rules (which we define with you in Discovery) protect longitudinal comparability.
Create a staff dashboard view: members in Baseline Survey, sorted by days overdue	A view builder over journey and cohort data: pick columns, filters, sort, save, share. Days-in-stage and overdue flags are computed fields available to every view. This is the Airtable muscle memory, preserved natively.
Reminder 14 days after a session opens if not closed	A notification rule builder: trigger event, delay, condition, audience, channel (email now, WhatsApp later). Rules are records; staff create and pause them freely.
Create a 'Strategic Partner' role with read-only access to a defined cluster of organisations	Role-based access control with organisation-cluster scoping is configuration: staff define the role, set permissions to read-only, and attach the cluster. Every permission change is audit-logged.

Power BI: retire it, and we say that as the people who built it

riivo built TGO's current Power BI geo-mapping and reporting dashboards, so this recommendation carries some self-awareness: move away from Power BI for Platform 2.0. It exists to compensate for what the platform cannot do, its data already disagrees with platform data, and it adds licensing, a sync surface and a second version of the truth. In the new platform, the movement map, Hub KPI views and reporting dashboards run natively against the single database, so the numbers on screen are the numbers in the reports, always. If a genuine self-service analysis need emerges beyond that, Discovery will size it on evidence; the requirement itself is answered natively, licence-free.

The Airtable question: our position



Five systems become one. Everything TGO runs on today lands in a single system of record.

We favour Option A, full consolidation at go-live, with a safety net. Consolidating your scattered tooling into one platform is, in our view, the single biggest structural win in this project, and Airtable is the centrepiece of it. Our reasoning:

- ◆ The member journey is Platform 2.0's central object and your strategic differentiator. Running it in two systems, even briefly, recreates the exact duplication risk your Trusted Data principle forbids. Option B's launch-time sync is the riskiest component you could build, and it is throwaway.
- ◆ The thing that actually must be preserved from Airtable is not Airtable: it is your team's ability to change stages, fields, views and dashboards on the fly. Our configurability architecture gives staff that same autonomy natively, and Discovery proves it on the prototype before you commit to the build.

- ◆ The safety net: Airtable is kept in read-only archive mode for a defined parallel-comfort window after go-live, at minimal subscription cost, so nothing is lost while confidence builds. It is a reference copy, never a second system of record.

And to be unambiguous about the perimeter: everything moves. CRM, sales pipeline, report request management and support requests come across along with journey tracking, because the Sales Journey triggers Programme Journeys and belongs in the same engine. Retiring Airtable does not even wait for the platform cutover: the journey engine can go live ahead of it, with the few stage gates that depend on survey-session state fed one-way from the legacy platform, or confirmed manually where a feed is not worth building, until cutover completes the automation. Journey data itself never lives in two places.

Data migration strategy

We treat migration as an engineering project with evidence, not a weekend cutover:

- ◆ **Profile and decide per entity.** Using the audit findings and our own knowledge of the schema, we classify every entity: migrate, cleanse-and-migrate, consolidate (the survey-option table, which your audit shows holds over 18 million denormalised rows, is the obvious candidate), or archive (facilitator records, clone tables).
- ◆ **Cleansing rules co-designed with TGO.** Duplicate participant IDs, null relationships and missing control fields get explicit, signed-off resolution rules; ambiguous records are quarantined for human review, never silently guessed.
- ◆ **Rehearse, reconcile, repeat.** The migration plan is grounded in source profiling during Discovery; full rehearsals then run repeatedly during the build. Go-live is simply the final, boring rehearsal.
- ◆ **Verify with evidence, not confidence.** Every pass produces a reconciliation report: source-to-target row counts per entity, checksums on key fields, field-level sampling, and a named list of every record that did not carry across and why. Where scale warrants it we run the pipeline on enterprise tooling such as Azure Data Factory. This is a proven pattern for us: we migrated Health Collective's clinical patient records, 100,000+ of them, from a legacy EHR system through a staging virtual machine and then into Dataverse via an ETL process, almost lossless, with the only loss being identified orphaned records, each one accounted for.
- ◆ **Airtable journey data** is migrated into the journey engine with stage history preserved, and the Python reporting data is reconciled so reports before and after cutover agree.

POPIA compliance

POPIA is designed in, not audited in afterwards: consent status captured as data on every participant record (already one of your quality flags); data minimisation in the model; encryption at rest and in transit; role-based access with full audit logging; South African data residency; a documented breach process; operator agreements; and independent penetration testing before go-live. Aggregated movement insights are anonymised by design so cross-member views can never leak participant identity.

iOS development

One codebase for iOS and Android either way; the real question is rebuild or evolve, and Discovery answers it. If the app is rebuilt, we would use Expo (React Native), which keeps the entire platform in TypeScript end to end; DAC Power, the offline-first field application we are building right now, is built exactly this way. If TGO prefers a substantial update of the current app instead, Flutter remains the right answer

and we work within it. The real engineering work is not the framework; it is **offline-first sync done properly**: a local queue, conflict-safe upload, resumable transfer for photo payloads, and automatic sync so facilitators never have to remember to press a button. Discovery validates the path, and the iOS priority itself, against your actual facilitator device base.

Maintainability, documentation and handover

One quiet failure of the current platform is that its knowledge lives in its provider: technical documentation was never maintained, so handover is now hard and slow. Our process makes that impossible for TGO. Documentation is generated as the build proceeds, not written up afterwards: technical docs, user guides and training artefacts are produced alongside the code, kept current release by release, and sit in a shared workspace TGO can open at any time. With the IP terms in Section 6 and a mainstream stack, TGO could hand the platform to any competent team at any point and lose nothing. The best proof of partnership is building as if you might leave, so you never have a bad reason to stay.

Ideas that push the mission further

One note covers all of them: the platform ships AI-ready, with a Claude integration layer built in, so each idea lands as configuration, not a rebuild. The simplest example is already in your pack: journey stage transitions creating Asana tasks and firing notifications automatically, so your team's time shifts from administration to the relational work only humans can do.

AI-assisted narrative interpretation

YOUR STATED FUTURE GOAL · RIIVO'S HOME GROUND

Report narrative is manually written today and flagged as a future automation goal. As a Claude partner program participant, riivo would design the reporting engine so an AI drafting layer can propose narrative interpretation from the quantitative results, always reviewed and finalised by your analysts.

Data quality copilot at the point of capture

TRUSTED DATA, ENFORCED UPSTREAM

Rather than flagging duplicates and GPS gaps at reporting time, the mobile app and platform validate at capture: possible duplicate participants surfaced to the facilitator, GPS sanity checks, consent prompts. Bad data is prevented, not detected.

Accessible multilingual, done with AI plus humans

FIXES THE 'TOO ACADEMIC' PROBLEM

An AI-assisted translation workflow with community reviewer sign-off, built on the survey versioning model, lets Zulu, Xhosa and Afrikaans land in phases with natural language, at a fraction of traditional localisation cost.



An AI ecosystem around the platform, not just inside it

DEEP ANTHROPIC EXPERIENCE · AI TO LEVERAGE HUMANS

The rebuild is the start, not the ceiling. Our Anthropic work lets us grow an internal AI ecosystem around TGO: knowledge bases holding your methodology and institutional memory, skills and plugins automating the routine work, assistants your team directs. The principle never changes: AI exists to leverage your people, so a team of ten can serve a movement of hundreds.

4 Relevant experience and references

Every capability Platform 2.0 needs is battle-tested somewhere in riivo's portfolio: scale, workflow engines, migration integrity, apps in the stores, and the rebuild motion itself. The four cases below each prove one of them; the matrix that follows maps the coverage.

Comparable builds



The Greenlight Office: automated report generation

YOUR DATA

YOUR REPORTS

BUILD COMPLETE

riivo designed and built the automated preparation of Greenlight's baseline and impact reports, the deliverable your previous provider could not automate: extraction from the current platform, transformation against your methodology, templated generation of the full packs. One honest caveat: the automation has never been widely rolled out, so reports are still generated manually today; Platform 2.0 is the chance to change that. The working knowledge of your data model that de-risks this whole project stands either way.

20-30pp

REPORT PACKS, GENERATED

Days → 0

PER-REPORT EFFORT AT ROLL-OUT



TTT Financial Group: custom build and platform at scale

NAMED

REFERENCED

LIVE METRICS

TTT runs its operation on systems riivo built and still evolves: a Power Platform core, end-to-end case management, a client portal, and Tina, an AI assistant living in WhatsApp, grounded in TTT's knowledge base and live client records. First go-live in three months.

10,000+

CLIENTS SERVED

879

TINA CASES, FIRST 10 DAYS

72.7%

RESOLVED, NO HUMAN

Referee: Roscoe Dekker, CEO



Health Collective: migration verified to the record

MIGRATION

RECONCILIATION EVIDENCE

Clinical patient data, captured across visits to doctors, migrated from a legacy EHR system through a staging virtual machine, then into Dataverse via an ETL process: reconciliation evidence on every pass, an almost lossless result, and every loss an identified orphaned record. The same discipline, scaled with enterprise tooling, is what Greenlight's 240GB estate gets.

100,000+

RECORDS MIGRATED

Zero

UNACCOUNTED LOSS



SurgiFlow: TGO's exact journey, live on the stack we propose

BUBBLE → NEXT.JS + SUPABASE + CLAUDE

LIVE + SCALING

Surgical Assistant's two-sided healthcare marketplace proved its model on Bubble; in 2026 riivo rebuilt it on the very stack proposed in Section 3, one of four such rebuilds this year, in service of a mission TGO will recognise: putting unemployed doctors back into theatre.

1,100+

VETTED PROFESSIONALS

60+

ACTIVE CUSTOMERS

R11.7m+

PAID TO SA DOCTORS

Referee: AD Schoeman, co-founder

Capability coverage, mapped to your evaluation

Capability TGO needs	Depth	Where riivo has proven it
Platforms of comparable complexity, at scale	●	TTT Financial Group: 10,000+ clients on a custom-plus-Power-Platform build; a listed insurer's core distribution system
Workflow and journey engines	●	A configurable workflow engine (triggers, stages, carry-over and strict-order rules) for a wealth-management client via a partner; TTT's end-to-end case management
Data migration and integration integrity	●	Health Collective: 100,000+ records, zero unaccounted loss; HR Department: 62 entities integrated with Xero via custom connections in one Dynamics app; legacy integration for the insurer above

Capability TGO needs	Depth	Where riivo has proven it
Client-facing apps, shipped to the stores	●	Teflon Connect (live events) and Everhot, both live on Apple and Google today
Offline-first mobile	🕒	DAC Power, in build now on Expo (React Native), offline-first by design
Bubble-to-AI-stack rebuilds	🔄	SurgiFlow, live at scale on the proposed stack after its 2026 rebuild, plus three further rebuilds this year; Platform 2.0 is this pattern at full scale
Reporting, dashboards and survey design	●	TGO's own automated report generation; Nomios survey design; management reporting across the portfolio
NPO sector, long-horizon stewardship	●	Fisantekraal Centre for Development: migrated from SharePoint onto the Power Platform, modernised and supported today, at discounted NPO rates
Our own medicine	●	riivo's delivery runs on internal applications and a client portal we built ourselves


● battle-tested in production 🔄 live and scaling on our AI stack 🕒 in build now

Why Power Platform and Bubble experience translates directly

A question worth pre-empting: much of riivo's delivered work runs on Power Platform or Bubble, so why does that qualify us for a custom build? Because the hard problems in Platform 2.0 are not language problems, they are **architecture problems riivo solves constantly**: configurable workflows and stage gates, multi-organisation security, staff-editable views, offline-sync capture, reporting automation. Power Platform and Bubble are professional workflow and configurability engines; years of building on them teaches a team how a configurable journey engine must behave, and the rebuild motion in the matrix above is how that experience now lands on our custom stack. Behind it all sits Saratoga, our part-owner, one of Cape Town's largest enterprise software development companies and an Andersen Consulting collaborating firm: senior review, QA capacity and continuity on call.

🔍 A project that ran into difficulty

Most projects hit difficulty somewhere; the test is what the partner does next. We built an early version of a core digital system for one of South Africa's largest insurers. As launch approached, it became clear the client's in-house data integrations feeding our system would not be ready: one internal developer carried that work, mid-Covid and hard to reach, and the go-live date could not move. The integrations were contractually not our problem; the failure of the project absolutely would have been.



The problem was upstream of our contract; we treated it as ours anyway. The system went live on the immovable date and kept a very large organisation functioning digitally through Covid. Trust is built in exactly these moments.

References

These clients have agreed to speak to TGO directly:

- ◆ **Roscoe Dekker**, CEO, TTT Financial Group · 082 300 7231
- ◆ **Sean Kelly**, CTO, Anisoptera · 082 300 7531
- ◆ **Alan Gardiner**, CEO, Teflon Connect (US-based) · 083 985 6095 / +1 561 309 0363
- ◆ **AD Schoeman**, co-founder, Surgical Assistant · 060 507 0849

5 Team and key personnel



Nic Bowman

SENIOR STAKEHOLDER · CEO

Owns: accountability, partnership, board interface

CEO of riivo since 2020, after heading operations and product development at Icon Digital Consulting and four years in investment banking at Absa/Barclays. Microsoft-certified Solution Architect (PL-600) and Power BI Data Analyst; Claude certified architect since June 2026.



Francois Lotter

DELIVERY LEAD

Owns: end-to-end delivery, architecture, technical quality

Head of Delivery. A certified Bubble developer with 20+ delivered projects; in 2026 he led the build-out of riivo's AI Build competency and its governance. Four production systems have been rebuilt onto the AI-built stack this year under his lead.



Natalia Luiz

PROJECT MANAGER

Owns: cadence, plan, risk, releases, communication

Certified ScrumMaster (CSM, A-CSM) who, as Agile Delivery Lead, took a multi-tenant SaaS platform from MVP to live production, onboarding its first customers and doubling team velocity. BCom Honours, UCT.



Dormehl Gertenbach

PRODUCT OWNER / BUSINESS ANALYST

Owns: workshops, spec, backlog, acceptance criteria

Joined riivo in 2025 from AgrigateOne, where he led the end-to-end integration lifecycle: complex data pipelines and data integrity across the fresh-produce supply chain. A former Dynamics 365 consultant with a BSc in Information Science; has led multiple successful riivo deliveries.



Lungile Willie

QUALITY ASSURANCE

Owns: test strategy, test packs, defect management, UAT

Senior Test Analyst with nine years of QA experience across insurance and wealth management, ISTQB certified, with a Master's in Information Technology. He owns riivo's test packs and defect discipline.

Key people and subcontracting

The people who run your Discovery lead your build; that continuity is part of what Discovery buys. If a key person must change, TGO hears first, meets the replacement, and pays nothing for the handover; project knowledge lives in shared artefacts, never in one head. Core delivery is riivo's own team, with one disclosed exception: Saratoga can add senior engineering or QA capacity under riivo's direction. Nothing else is subcontracted without TGO's prior written consent.

6 Commercial terms

One principle runs through everything below: we own the outcome. riivo commits to fixed prices; if work takes longer than planned, that is our cost, not yours.

Stage 1: Discovery fee

Discovery: Premium

Fixed fee · 8 weeks · senior multi-stakeholder engagement · board-ready pack

R165 000

FIXED · EXCL. VAT

THE FEE INCLUDES

- ◆ Validated data model & migration plan (risk-assessed)
- ◆ Full build backlog: user stories with acceptance criteria
- ◆ Fixed-price Stage 2 build scope & critical path
- ◆ Technical, journey & report architecture
- ◆ Wireframes, clickable prototype, two user-testing rounds
- ◆ Board-ready readout & full artefact handover

Payment structure for Stage 1: 50% on commencement, 50% on handover of the Discovery artefacts. If TGO proceeds to Stage 2 with riivo, **25% of the Discovery fee (R41 250) is credited against the build balance**, taking the net cost of Discovery to R123 750. All artefacts are TGO's property on final payment, whichever partner builds Stage 2: the pack is written to be buildable by any competent team.

Stage 2: indicative build range

Based on the brief, the Discovery Input Pack, both audits, and our existing working knowledge of your data model, our indicative order-of-magnitude for the full Platform 2.0 build is:



R2.3m to R2.8m excluding VAT for the complete Must Have scope: platform, member portal, iOS and Android mobile app, automated reporting engine, journey engine, full data migration, training and handover. Add the Discovery fee, net off its credit, and the full programme lands **comfortably inside your R2.5m-R3.5m envelope**, with headroom to spare.

How we got there. We decomposed the Must Have requirements into fifteen delivery areas and sized each in riivo build points, the same method that fixes the price after Discovery. The range covers all of it:

#	Delivery area	What it covers
01	Platform foundation	infrastructure, environments, authentication and role-based access, monitoring, security baseline
02	Data model & integrity	the new schema, validation engine, full audit trail
03	Data migration	profiling, cleansing rules, rehearsals, reconciliation evidence
04	Journey engine	templates, stages, gates, cohorts, cycles, notifications

#	Delivery area	What it covers
05	Survey management & designer	versioned definitions, no-code editing, mobile sync
06	Mobile app	offline-first field capture, iOS and Android
07	Reporting engine	automated generation across all survey and report types
08	Member portal	profiles, report access, survey history, resources
09	Staff dashboards & views	configurable KPI, pipeline and journey views
10	Integrations	maps, email, push notifications, Asana automation
11	Documentation & training	generated docs, user guides, onboarding flows
12	Structured testing	a written test pack per backlog item
13	UAT support	defect resolution through your acceptance rounds
14	Go-live & cutover	final migration, switch-over, rollback plan
15	Hypercare	post-launch stabilisation, handing into Momentum

Points are anchored to what each area is worth in traditional delivery terms; our AI-accelerated engineering is what fits this scope inside your envelope at all. Priced conventionally, we believe this scope exceeds your ceiling, which is consistent with what happened the last time it was built down to a price.

Confidence and headroom. Our confidence that this budget supports the requirements is high: our AI build competencies set the cost base, and we already know the data model, the report structures and the audits' findings first-hand. Per-entity migration decisions, the depth of native analytics at launch and the report automation target are confirmed in Discovery; if the fixed price exceeds this range we say so before you commit, drivers itemised. And rather than padding the build, we would direct the headroom under your envelope to Momentum: a funded runway of monthly improvements after launch, which is where selected Should Haves and early Phase 2 ambitions are best delivered.

A different way to structure the budget: Build + Momentum

You asked for our standard support model; we would go further. Our recommended shape: a fixed-price core build with formal change control, released first, then straight onto Momentum after roll-out, all within the stated envelope. From that point riivo owns support, tracks the metrics that matter (adoption, data quality, time-to-report), and responds to feature requests at speed, because the team that built the platform is the team evolving it. For TGO's size and roadmap we recommend **Momentum: Grow**. It replaces the static retainer you have today with a subscription that ships improvements monthly; Phase 2 never needs another RFP, its items simply become Momentum backlog. The hosting SLA is quoted with the fixed build price after Discovery.

Momentum: Launch

LEAN START

R30k /mo · 4 pts

- ✓ 4 build points / month
- ✓ Product Owner + Developer
- ✓ Fortnightly demo & review
- ✓ Fair-use support
- ✓ Roll over one month

RECOMMENDED FOR TGO

Momentum: Grow

THE CORE OFFERING

R50k /mo · 8 pts

Steady, consistent delivery for a product that is genuinely growing.

- ✓ 8 build points / month
- ✓ Product Owner + Developer
- ✓ Fortnightly demo & review
- ✓ Fair-use support
- ✓ Roll over one month

Momentum: Scale

SERIOUS GROWTH

R100k /mo · 16 pts

- ✓ 16 build points / month
- ✓ PO + two Developers
- ✓ Fortnightly demo & review
- ✓ Fair-use support
- ✓ Roll over one month

All tiers exclude VAT and hosting; tiers step up or down with a month's notice.



Ongoing hosting

Indicative hosting for the proposed architecture (Azure South Africa: auto-scaling application tier, managed PostgreSQL, storage, monitoring, backups, DR) is **R8 000 to R14 000 per month excluding VAT**, depending on the redundancy and DR posture chosen in Discovery. riivo manages hosting under an SLA quoted with the fixed build price; third-party API costs pass through at cost, no markup.

WORTH NOTING We have deliberately quoted a premium Azure posture. More cost-effective hosting options exist, and Discovery can right-size the architecture and the monthly cost to TGO's appetite.



Payment structure summary

Stage	Terms
Discovery	Fixed fee R165 000 ex VAT: 50% on commencement, 50% on handover of the artefacts. 25% of the fee (R41 250) is credited to the build balance if TGO proceeds with riivo.
Build	Fixed price set in Discovery, invoiced per completed milestone, never in advance. Each feature milestone completes in three signed steps: story sign-off, development, QC tested. Delivery milestones then follow for UAT, go-live and hypercare. No completed milestone, no invoice; formal change control governs any scope movement.
Ongoing	Momentum subscription monthly in advance; hosting monthly at the quoted SLA rate. If TGO prefers not to take Momentum, riivo offers standalone care and support packages so the platform is never unsupported.



IP ownership

Confirmed, without reservation. On delivery and payment, TGO owns all code, schemas, designs, documentation and assets, including everything produced in Discovery. The stack is open technology with no proprietary riivo runtime, so ownership is real, not theoretical: any competent team could take it over.

Defects and warranty, in plain language

If something we built does not work the way the signed-off specification says it should, that is a defect: **we fix it at our cost for 90 days after go-live**, no debate, no invoice. Critical defects (data integrity at risk, platform down, reports blocked) are worked immediately until resolved; the rest is prioritised with your team by impact. After 90 days, defect cover continues as part of Momentum for as long as the subscription runs. What warranty does not cover: new requirements, changes of mind, and third-party outages, which go through normal prioritisation. Data integrity issues are treated as critical at any time, because we understand exactly what trusted data means to Greenlight.

7 Assumptions

Everything in this proposal holds under the following assumptions; if any of them move, we say so immediately and re-baseline together.

- ◆ TGO provides the SME access, decision cadence and data access committed in the brief; Discovery starts 1 September 2026, and build mobilisation assumes a Stage 2 decision within two weeks of the readout (the go-live window absorbs a slower one).
- ◆ The indicative range covers the Must Have launch scope; Should and Could Haves are priced as options in the Discovery output.
- ◆ Hosting, third-party licences and API usage sit outside the project budget, per the brief, quoted separately.
- ◆ The current platform stays operational under existing or interim management during the build; riivo can quote interim stewardship separately (deadline 30 September 2026). All figures exclude VAT.



Next step. We would welcome the opportunity to present to the board why we believe riivo is the best long-term partner for TGO.

Nic Bowman · nic@riivo.io · 083 995 1779