

The depth behind the deck.

Everything the 14-slide deck compresses, at full length — so the call can start from decisions instead of explanations. Compiled 31 August 2026 for Ben; sources are the repo docs (rwf.qalarc.com/wiki keeps the live version).

Contents

A	The call agenda in full — the five Phase-0 questions, our position, the red lines, the signals to listen for
B	The masterplan in full — phases 0–4, scope and explicit non-goals, KPI table, architecture sketch
C	The business analysis, condensed — market position, the two moats, revenue pools, ranked risks, verdict
D	Element status map — all 36 elements across families A–G, the seven systems, current blockers

Deck + appendix sources: docs/01, 02, 09, 17, 19, 22 · figma/notes/analysis.md · apps/figma-app/CHANGELOG.md.

The five Phase-0 questions

45 minutes. These are the questions the prototype was built to make answerable — each now has a working artefact behind it.

1 • Handicap mechanics

How do you imagine baselines — self-assessed tier, or historical data? Ours: v1 self-set fitness tier → effort multiplier (couch $\times 1.5$ → athlete $\times 0.85$); v2 %HRR / Karvonen from any BLE strap — both already in the engine.

Listen for: has the thinking gone past “handicaps are good”? Your Figma’s personalised-targets model (D1/D8) is the deep divergence — the deck’s “decide together” list, row one.

2 • Rep verification

Is camera-in-browser acceptable for MVP? No install, nothing uploaded, \$0 — pose counting runs on-device (MoveNet), and HR straps cross-check effort via %HRR. Both live in the app today.

Tolerance for “good enough” here tells us a lot about shipping software together. Your peer-challenge / trust-score design is the planned social backstop (V1).

3 • Charity pot money

Who holds it — platform, honour-system, or partner charity? Needs a legal opinion before a dollar moves: AU state law carries raffle/betting classification risk. Until then the engine is pledge-only, by design.

Nico’s idea is good. It just needs structure — deck slide 11, first column.

4 • Corporate mode

Any live prospects or LOIs? Real interest → we build Slack-first (the Bolt bot already runs). No names yet → consumer MVP first, corporate follows the loop. The hub console’s corporate tab already exists ($k \geq 5$ privacy, employer-funded pots).

Named prospects change build priority — this is the question with a direct engineering consequence.

5 • Equity shape

What split is being imagined? Vesting? And the big one: IP assignment on the blueprint and Figma — who owns what if we part ways? We don’t negotiate numbers on the call; we get the shape.

Deck slide 13 is the posture: two founders, both IP assignments into the company, numbers talked not projected (docs/19).

If time: who else is involved — Nico’s role, other collaborators?

Position, stance, red lines

Our position

Co-founder-track, not a paid build. You bring design + vision + publicity; we bring build + chat/distribution plumbing + AI — a complete founding pair, and that is what equity should price. The prototype is evidence of ownership, not a bargaining chip. Negotiate from the position the demo creates: the product exists and works, the velocity is demonstrated rather than claimed, and the design is already integrated — extracted, analysed, and shipped as a running app.

Figma stance

Respect, audit, mark up, build. The full 65-screen file was extracted and analysed line by line (figma/notes/analysis.md — colour tokens, type scale, 30-row divergence register). Ownership to clarify calmly: if equity proceeds, IP assignment goes both ways — your design into the company, our code in.

Red lines — stated plainly, as you asked

- **No cash prizes.** Charity pots only, and only after legal structure exists.
- **No Apple Messages in MVP.** No bot API exists. Revisit if one ever appears.
- **No custom ML in v1.** Open-source pose counting is enough; effort is cross-checked with %HRR.
- **Both founders vest.** 4 years, 1-year cliff, IP consolidated — protects you as much as us (docs/19 terms checklist).

Three signals we'll both be listening for

1. **Partner or hands?** “What would you do?” = collaborator. “When can you build it?” = client with a spec. Everything so far — the blueprint, the invitation to push back, equity-for-work up front — reads collaborator.
2. **Is corporate real?** Named prospects → Slack-first. Vibes → consumer MVP first.
3. **Does equity talk get specific?** Split / vesting / IP raised unprompted = serious. “Work it out later” = the one flag worth nailing before serious build hours.

The equity minute, in order

Frame (two founders) → demo as evidence the collaboration works → the four structural product questions, decided together → split + terms → agree to have a founders' memo drafted within a fortnight.

One sentence worth saying out loud: “I’ve already made my bet — the last five days are it. Now we just make the paperwork match.”

The strategy, then the phases

Strategy in one paragraph: win the empty niche — handicapped, small-group fitness competition inside the group chat — by shipping a playable WhatsApp/Slack bot fast, using in-browser pose counting for verification (\$0, no approvals), and letting the chat broadcast loop do acquisition. The app is the scoreboard and the archive; the chat is the arena. Monetise later through corporate wellness once the loop is proven.

Phase 0 — due diligence & alignment DONE

- Private repo + research: OpenGym landscape (docs/04), wearables (docs/05) — done
- Business analysis (docs/01) — done
- Masterplan (docs/02), messaging-platforms research (docs/18, 21), backend architecture (docs/22) — done
- Figma received → extracted, audited against our independent design, 30-row divergence register (figma/notes) — done
- Call: align on MVP scope, roles, equity — **this follow-up stands in for it; re-book next**

Phase 1 — prototype DONE

Added after the plan was written, because the email deserved a real answer: game engine (44 tests), WhatsApp + Slack bots (22 tests), phone-first PWA with camera verify + HR straps, your 65 screens running as a working app (v1.0.0, 11/11 engine parity, offline PWA, self-playing demo), site + demo + wiki + hub console, AI layer, push-to-deploy hosting. ~250+ automated checks green across all suites.

Phase 2 — MVP: “the 300 loop” • ~6–8 weeks NEXT

Goal: 3–5 real group chats playing weekly.

1. **game-core** — match engine, handicap v1, standings, charity-pot ledger (pure TypeScript, tested — exists)
2. **WhatsApp bot** — Qalarc Hub for dev/dogfood → WhatsApp Business Cloud API for pilot. Commands: create, pick exercises, log, standings, winner
3. **Slack bot** — same command surface, slash commands + buttons; first in line for corporate pilots
4. **Web app (PWA)** — join link from chat → log with camera pose counting or manual → live match → shareable result card
5. **Handicap v1** — fitness tier → effort multiplier; ship, learn, replace with %HRR in Phase 3

Explicitly OUT: native apps, Apple Messages, cash anything, custom ML models.

The later phases, and how we measure

Phase 3 — effort truth & wearables · ~+8 weeks STARTED

- **Handicap v2: %HRR (Karvonen)** — score = reps × (your %HRR ÷ your baseline %HRR for that exercise). BLE chest straps already live via Web Bluetooth; Apple Watch via HealthKit live-HR; Android via Health Connect (sync, cross-check). Blend math implemented + tested in the engine.
- WatchOS / Wear OS companions (rep counting from the wrist accelerometer)
- Season / league structure as the long-arc retention fix — **already built** (4-week ladder, divisions, belt)
- Shareable clip generation (auto-cut mid-set footage + result-card overlay)

Phase 4 — business layer SEEDED

- **Corporate mode** — Slack enterprise install, org dashboards, team-vs-team leagues, invoicing. The revenue engine. Hub console’s corporate tab already seeded: employer-funded pots, k≥5 aggregate wellbeing dashboard, renewal outlook.
- Charity-pot mechanics with a partner charity — legal structure first
- WHOOP / Garmin / Google Health cloud verification (recovery-adjusted handicaps)
- Messenger bot; Apple Messages if an API ever appears

KPI table

PHASE	METRIC	TARGET
MVP	Groups activated	3–5 real chats
MVP	Matches / group / week	≥ 1
MVP	% reps camera-verified	≥ 40%
Growth	New users per match broadcast (viral k)	≥ 0.3
Growth	Week-4 retention of a group (still playing)	≥ 60%
Business	Corporate pilots	2 by Phase 4

Architecture sketch (docs/22 — the single source of truth)

WhatsApp (Hub → Cloud) · Slack (Bolt) · Beeper (built) · Telegram (bridge)

▼ all events become the same commands

@rwf/bot-core — the single chat brain (grammar, cards, taunts, AI hooks)

▼ MatchStore mirrors into the API (landed)

apps/api :4174 — state authority (crews · matches · seasons · pots) ← — apps/web (PWA)

▼

@rwf/game-core — pure engine: handicaps · comeback · seasons · pot ledger

Chat adapters are thin: adding a platform is one adapter, not a rewrite. Match state mirrors into apps/api; the app is a first-class client of the same scoreboard the bots write to.

Who, and what’s load-bearing

Who we’re dealing with

- **The publicity lever is real, especially in Australia** — mainstream media will cover “Silverchair drummer launches fitness app”; a first-wave injection most startups pay seven figures for.
- **The nostalgia demographic (35–50) maps exactly onto the target user** — people in group chats who used to be fit, aren’t anymore, and rib each other about it. Silverchair’s audience is the product’s audience.
- **Caveat:** celebrity attention is a one-shot spike, not a retention engine. It buys the top of the funnel; the handicap system and chat presence have to keep them.
- **Signal quality of the original email: high.** A year of refinement, a production-fidelity Figma + full blueprint, an explicit invitation to push back, equity-for-work opened first. A collaborator-shaped founder, not an idea guy shopping for hands.

The load-bearing ideas

IDEA	WHY IT MATTERS COMMERCIALLY
The “300” format	Brilliantly simple loop: any reps, any order, first to 300. Self-balancing, zero setup, understandable in one chat message.
Handicap system MOAT 1	The difference between a gimmick and a system. Golf proved handicaps keep weak players paying for decades. Hardest thing to get right — which is why it’s defensible.
Chat-first distribution MOAT 2	Every match broadcast is an ad seen by the whole group. Non-users spectate the leaderboard within minutes. Near-zero CAC, built-in viral loop.
Pick your days	2–3 days/week, go hard. Anti-burnout by design; removes the daily-streak failure mode.
Charity wagers	Real stakes without gambling exposure (with structure), and a natural corporate-sales hook.
Taunting as a feature	The content engine — user-generated shareable clips stacked on the chat loop.

Market, risks, verdict

Market position — the niche is open

What exists: solo trackers (saturated, lonely, no stakes) · individual wagering (cash stakes, regulatory baggage, US-centric) · live classes (anonymous crowds, not your mates) · async comparisons (weak, pairwise, buried).

What's empty — verified via GitHub + market scan: *small-group, real-time, handicapped competition delivered inside the group chat where the banter already lives.* No established player; the “gym leaderboard” OSS space is 0-star student projects.

Adjacent revenue pools: corporate wellness (AU\$600+/employee/yr — likely the actual business; B2C is the marketing) · charity partnerships (rev share on pot mechanics) · brand sponsorship (the 35–50 male demo is notoriously expensive to reach).

Risks & hard problems, ranked

RISK	READ + MITIGATION
1 · Rep verification / cheating EXISTENTIAL	Self-reported = lying contest; annoying = nobody plays. Mitigation shipped: in-browser pose counting (\$0, no install) + %HRR cross-checks from any BLE strap — effort verifiable even when reps aren't.
2 · Retention after week 3	Novelty decays. The designed stack: handicaps (always winnable) + chosen days (no streak guilt) + charity stakes (real consequence) + chat presence (no app-opening required) + seasons as the long arc.
3 · WhatsApp economics	Cloud API charges per conversation. Mitigation: batched standings, never per-rep; model costs before scale.
4 · Charity-wager legal	Must avoid raffle/betting classification under AU state law. Legal opinion before money moves; pledge-only until then — enforced in the engine today.
5 · Apple Messages	No bot API. Skipped for MVP, stated plainly.
6 · Single-founder dependency	Design/vision + build/systems = the complete pair. The equity conversation should reflect co-founder reality, not contractor pricing.

Verdict

Strong opportunity. Differentiated product with two genuine moats; design risk already paid down; sits directly on existing infrastructure and skills; corporate wellness is the quiet, large revenue pool underneath the fun B2C story. **Position: co-founder-track equity + build ownership — the equity conversation stays open, on purpose (docs/19).**

The 36 elements, family by family

What began as an email evaluation is a running system: tested engine, two chat bots, a phone-first app, an ops console, an AI layer, push-to-deploy hosting.

FAMILY	STATUS	WHAT'S IN IT
A · Core loop	BUILT	300-format match · tier handicap v1 (couch ×1.5 → athlete ×0.85) · closure bonus +15 · charity-pot ledger (winner directs, never receives) · AI taunt engine with canned fallback · %HRR v2 engine-ready, activates when straps land
B · Retention arc	BUILT	4-week seasons (3/2/1 + MVP vote) · champion crowning · A/B division relegation-promotion · streaks with charity forgiveness · comeback ×1.2 (>30% behind, once/match) · baseline learning (anti-sandbagging)
C · Social & viral	MOSTLY	Branded result cards (SVG from bots, PNG from app) · spectator mode · crew-vs-crew challenges · drop-cam clips + public web ladder still ideas
D · Corporate	SEEDED	Hub corporate tab: org leagues · employer-funded pots (“no employee money handled”) · aggregate-only dashboard, k≥5 suppression · renewal outlook — awaiting real org data
E · Verification	P1 LIVE	In-browser MoveNet camera counting (pushups/squats, on-device only) + Web Bluetooth HR straps (GATT 0x180D, Karvonen %HRR per set) — entries log verified:true · avgHrrPct. P2 HealthKit/Health Connect, P3 WHOOP/Garmin phased
F · Wilder cards	IDEAS	Referee review · physical champion belts · radio mode (drum-track rep cues) · annual charity championship
G · Second wave	4 BUILT	Live: rematch · Monday digest (matches, margins, MVPs, pot, rivalry, AI one-liner) · photo finish (≤5% gap) · nemesis detection. Proposed: PRs · ghost race · guest slot · charity all-time ladder · roast tiers · adaptive equivalents · warm-up predictions
METRIC	VALUE	
Tests	90+ core green (engine 44 · bots 22 · API 24) + app e2e/demo/SW suites — every push, CI-gated	
Public pages	9 — site · demo · app · figma-app · system · hub · debug · connect · state API	
Elements / components	36 across A–G · 16 design-system primitives (rwf.qalarc.com/system)	
Deploy / AI	GitHub push → Cloudflare Pages ~20s · GLM-5.3 server-side proxy (guide, narrator, taunts)	

The seven systems, the live blockers

Systems

- 1. **game-core** — pure TS engine, 44 tests, 100% branch cover
- 2. **bot-core + bots** — one command grammar, WhatsApp (Qalarc Hub) + Slack (Bolt Socket Mode), sim harnesses, result-card generator
- 3. **app** — phone-first PWA (onboard → crew → match → seasons), polished: tier cards, medals, confetti, sticky log bar, 35 browser checks
- 4. **figma-app** — your 65 screens with the real engine behind the core loop: v1.0.0, 22 screens wired to real state, power-ups, danger zone, daily winners, self-playing demo, offline PWA
- 5. **hub console** — ops + corporate tabs, live state polling
- 6. **site** — Three.js showcase, interactive handicap demo, AI guide
- 7. **AI layer + apps/api** — /api/ai proxy (key stays server-side) · unified REST backend (crews/matches/seasons), 24 tests — the MVP foundation that replaces localStorage + bot file state

Top blockers, as of today (docs/15)

SEV	BLOCKER	WHAT UNBLOCKS IT
HIGH	Slack app needs 5-min human creation	One-time setup → permanent install link
HIGH	WhatsApp Cloud API group support unverified	Verification run shapes the pilot architecture
HIGH	Charity-wager legal opinion	Before real money; pledge-only enforced meanwhile
MED	Always-on bots	Minirig kit ready — one SSH command
MED	App ↔ API ↔ bots state unification	The mirror landed; the wire is the next (mechanical) build step

Already adopted from your Figma (no decision needed)

Bottom nav + central LOG · quick-log sheet with presets & undo · RUF conversion preview · danger-zone escalation craft · dual clock (group + local) · system-state components · “battle” as the noun · result-card direction · offline queue states · reduced-motion discipline. The five structural decisions stay open on purpose — deck slide 10, decided together on the call.

Everything in this appendix maps to a repo doc or a live URL: rwf.qalarc.com/wiki keeps the current truth.