

STRATEGIC ADVISORY — FINTECH & AGRICULTURAL ECONOMY

CAPITAL **DE** SINCLAIR

FEASIBILITY & FUNDABILITY ASSESSMENT

The *Amana* Score Dossier

An alternative credit rating and risk-scoring infrastructure for smallholder farmers and agribusinesses across Nigeria's agricultural value chain — technical feasibility, regulatory pathway, pre-mortem risk analysis, unit economics, and a three-phase launch roadmap.

Prepared by

Capital de Sinclair — Independent Strategic Advisory

Date

August 2026

CONFIDENTIAL DRAFT

Contents

SIX SECTIONS, ONE APPENDIX

01	Executive Summary & Verdict	Bottom-line assessment before the detail
02	Core Infrastructure & Technical Feasibility	Alternative data architecture and rural-ready tech stack
03	Regulatory & Licensing Framework	CBN Credit Bureau vs. risk-SaaS, NDPA, KYC/BVN
04	Limitations, Risk Factors & Pre-Mortem	Data scarcity, agro-macro shocks, bank risk aversion
05	Business Model & Unit Economics	Revenue streams, paying customer, final verdict
06	Execution Roadmap	Pilot → Integration → Scale
—	Appendix	Risk register & governing assumptions

Executive Summary & Verdict

Nigeria's agricultural sector employs roughly a third of the working population and contributes close to a quarter of GDP, yet agriculture receives a share of formal bank credit in the low single digits. The shortfall is not primarily an appetite problem — lenders want agric exposure to meet CBN sectoral-allocation expectations and DFI mandates — it is a legibility problem. The overwhelming majority of Nigeria's roughly 40 million smallholder farmers have no BVN-linked credit file, no titled land to pledge, and no digitised transaction history a loan officer can underwrite against. Amana Score, the illustrative platform modelled in this dossier, exists to manufacture that missing signal: a composite risk score built from satellite agronomy, weather history, mobile-money behaviour, cooperative repayment records, and off-taker delivery data, delivered as infrastructure to institutions that already have balance sheets, not as a new lender competing for the same borrowers.

BOTTOM LINE

The business is fundable, but only along one of the two paths examined. Pursuing a full CBN credit bureau license under the Credit Reporting Act 2017 is a capital-heavy, 18–24 month detour that competes head-on with three entrenched incumbents for a mandatory-reciprocity mandate this thesis does not need. Operating as a B2B alternative-scoring / agtech risk-SaaS provider — selling decision-support to licensed MFBs, commercial-bank agric desks, DFIs, and value-chain actors who already hold the credit risk — reaches market in roughly a third of the time and preserves the proprietary data moat the business is actually worth building.

The harder constraint is not the model. It is that a technically excellent score does not, by itself, change a Nigerian loan officer's incentives, a bank's funding-tenor mismatch, or a farmer's temptation to side-sell at harvest. Score quality is necessary and achievable within twelve to eighteen months; distribution into risk-averse lender balance sheets is the multi-year bet investors are actually underwriting.

This dossier works through the technical build, the regulatory choice, the specific ways this kind of venture fails in Nigeria's agricultural context, the revenue architecture and who actually pays for it, and a phased three-stage roadmap from cooperative pilot to national scale. Every section closes with a plain judgment rather than a menu of options, in keeping with the brief.

Core Infrastructure & Technical Feasibility

2.1 Alternative Data Architecture

No single data source is sufficient for an unbanked farmer; the model has to triangulate across independent evidence layers so that the absence of one (say, a smartphone) does not disqualify an applicant outright. Six layers make up the composite score:

LAYER	SOURCES	WHAT IT PROXIES
Geospatial & remote sensing	Sentinel-2 / Landsat NDVI & NDMI time series, NASRDA imagery, field-boundary delineation via GPS walk-the-plot	Actual farm size vs. declared size, planting consistency, in-season crop health, flood/drought exposure
Climate & agromet	NiMet station records, gridded rainfall reanalysis (CHIRPS) for ungauged areas, historical onset/cessation dates	Rain-fed yield probability, dry-spell frequency, exposure to ENSO-linked shock years
Soil & agronomic	National soil atlas (IITA/FMARD), crop-suitability and fertiliser-response zoning	Realistic yield ceiling for the stated crop and plot, input-response quality
Transactional / mobile money	Consented wallet inflow-outflow (OPay, PalmPay, Moniepoint, Paga rails), airtime top-up regularity, USSD savings/esusu history	Cash-flow rhythm, income stability, informal savings discipline
Off-taker & value-chain	Forward-purchase contracts and delivery logs (e.g. AFEX warehouse receipts, outgrower-scheme delivery records), agrodealer input-credit repayment history	Actual sales-side reliability and prior repayment on any existing input finance
Cooperative & identity	Cooperative membership tenure, group-guarantee repayment history, BVN/NIN deduplication, existing thin-file bureau hits	Social collateral, group discipline, identity certainty for fraud/duplication control

A supplementary psychometric layer — a short, USSD-deliverable questionnaire in the EFL Global style — is worth building only as a bridge signal for genuinely first-season entrants with no other trail. It is gameable and should never carry more than a minority weight in the composite score, flagged transparently to the lender as low-confidence input rather than blended invisibly into a single number.

2.2 Tech Stack & Rural-Ready Data Pipelines

The binding design constraint is not model architecture, it is connectivity. Most of the source population is reached through a field agent or a cooperative secretary, not a personal smartphone with a live connection, so the pipeline has to be offline-first by default rather than online-first with an offline fallback.

- **Field capture:** a lightweight Android agent app (offline SQLite store, background sync on connectivity) for cooperative secretaries and extension officers to log farm polygons, photos, and group attendance; a USSD/SMS channel (via an aggregator such as Africa's Talking) for farmers without smartphones to confirm consent and receive score-linked notifications.
- **Geospatial processing:** NDVI/NDMI time series pulled through the Sentinel Hub or Google Earth Engine APIs rather than hosting raw satellite tiles directly, which keeps storage and compute costs proportionate to a lean team.
- **Data lake & storage:** S3-compatible object storage, hosted in-region where possible given NDPA data-residency expectations for controllers of major importance; PostgreSQL for scored-entity records; Redis for high-frequency lookup caching.
- **Feature & model layer:** an Airflow-orchestrated ETL feeding a feature store (Feast) so that scoring features stay consistent across cohorts scored months apart; gradient-boosted trees (XGBoost/LightGBM) rather than deep learning, chosen deliberately for interpretability, robustness on small-N cohorts, and compatibility with the reason-code / adverse-action explanation obligation examined in Section 3.4.
- **Delivery:** a FastAPI scoring service exposing an authenticated REST endpoint and webhook to lender systems, returning a score band plus reason codes rather than a bare number.
- **Resilience:** most of the book is scored in weekly or seasonal batch refreshes rather than real time, with SMS-based score delivery to field officers in low-smartphone-penetration clusters — real-time scoring is reserved for the minority of transactions that genuinely need it, such as an off-taker's point-of-sale input-credit decision.

SECURITY NOTE

Every score pull is written to an immutable audit log, BVN/NIN are stored tokenised rather than in the clear, and access is role-scoped by lender account — this is not optional hardening, it is a direct NDPA and CBN-examination requirement covered in Section 3.4.

Regulatory & Licensing Framework

Nigeria offers two structurally different paths to market. The choice between them is the single highest-leverage decision in this feasibility study, because it determines both the capital required before revenue and the ceiling on what the business is ultimately allowed to sell.

PATHWAY A — CBN CREDIT BUREAU

- License required under the Credit Reporting Act 2017; CBN Governor approval, fit-and-proper review of directors and shareholders
- ₦500 million refundable minimum capital deposit to the CBN, plus a non-refundable application fee, before a license is even granted
- Three incumbents already hold the mandatory-reciprocity relationships with deposit money banks — CRC Credit Bureau, CreditRegistry, and FirstCentral
- A licensed bureau is compelled to serve all requesting lenders on comparable terms, which cuts against building a defensible, exclusively-monetisable agri-data model
- Realistic time to first revenue: 18–24 months

PATHWAY B — AGTECH RISK-SAAS

- Registers as an ordinary Nigerian company; sells scoring and portfolio-monitoring software to licensed lenders and value-chain actors who retain the actual credit decision
- No CBN bureau license needed *provided* the platform is consistently positioned and contracted as decision-support, not as the credit report itself
- Capital goes into data pipelines, satellite API contracts, and field operations rather than a refundable regulatory deposit
- Retains full commercial freedom over pricing and proprietary model IP
- Realistic time to first revenue: 6–9 months

SINGLE HIGHEST-PRIORITY PRE-LAUNCH ACTION

Before distributing the API to more than one or two pilot lenders, obtain an explicit written no-objection from CBN's Consumer Protection & Financial Policy Department confirming the platform's scope sits outside the Credit Reporting Act's bureau definition. Structuring alone will not settle this; a documented regulator position will. The two structural safeguards that keep the platform on the risk-SaaS side of that line are: the licensed lender always pulls its own formal bureau file and makes the final decision, and Amana Score's contracts and marketing never describe its output as "the credit report."

3.1 NDPA, NDPC & the Automated-Decision Obligation

The Nigeria Data Protection Regulation of 2019 has been superseded: the Nigeria Data Protection Act 2023 took effect in June 2023, established the Nigeria Data Protection Commission as the enforcing authority, and the NDPC's General Application and Implementation Directive, effective September 2025, formally retired the old NDPR framework. A platform processing farmer records at the scale this business requires will very likely be classified a data controller "of major importance," which carries specific, non-negotiable obligations:

- Registration with the NDPC and appointment of a named Data Protection Officer
- A maintained Register of Processing Activities and a Data Protection Impact Assessment completed before any automated-decision feature goes live
- A working right-to-object mechanism — the Act gives data subjects the right to contest automated decisions with significant effect on them, which means every farmer needs a human-reviewable appeal path attached to an adverse score, not just a number
- Layered, local-language consent (Hausa, Yoruba, Igbo, Pidgin) captured at cooperative enrolment, with a separate explicit consent step for any mobile-money or telco data pull

3.2 KYC, BVN & NIN Integration

BVN (via NIBSS) and NIN (via NIMC) are the practical identity backbone for deduplication and sanctions/PEP screening. The realistic constraint is that a meaningful share of the target population, disproportionately women and rural farmers in the north, still lack a BVN. Gatekeeping cooperative enrolment behind full BVN capture would exclude a large share of the exact population the platform exists to serve. The workable design is a tiered KYC model consistent with CBN's existing financial-inclusion tiered-KYC framework: Tier 1 enrolment through cooperative vouching and biometric capture, progressing to full Tier 2 BVN/NIN-linked KYC once a farmer transacts at a threshold that requires it.

Limitations, Risk Factors & Pre-Mortem

4.1 Data Scarcity & Fragmented Value Chains

Three structural facts about how Nigerian smallholder agriculture actually operates have to be designed around rather than assumed away.

INFORMAL LAND TENURE

Most smallholder land is held under customary or family arrangements, not formal title, so collateral-based underwriting is a dead end for the overwhelming majority of the target population. Amana Score has to function as a genuine collateral substitute — built on cooperative vouching, off-taker contract history, and repeat-season repayment behaviour — not as a collateral-verification layer bolted onto a conventional underwriting process.

CASH-BASED TRANSACTIONS

A large share of farm-gate trade still happens in cash at rural markets, invisible to both bank and mobile-money rails. Rather than assuming digitisation into existence, the model should reward farmers who route sales through a digitised off-taker or cooperative channel with faster review and a confidence bonus, treating digitisation as an incentivised outcome rather than a precondition.

THIN OR ABSENT FINANCIAL HISTORY

First-season entrants are a genuine cold-start problem. The mitigation is to score the cooperative first and individualise the score only after two to three seasons of observed behaviour, communicating explicitly lower confidence bands — and correspondingly lower initial credit limits — to lenders rather than presenting a first-season score with the same apparent certainty as a fifth-season one.

4.2 Agricultural & Macro Risk

RISK	EFFECT ON MODEL & PORTFOLIO	MITIGATION BUILT INTO THE DESIGN
Climate shocks (flood / drought)	Recurring, not tail-risk, for this population — a bad season is structural, not an outlier to be trained away	Climate-adjusted stress scenarios in the base model; bundle parametric weather-index insurance so a bad season does not become a portfolio-wide default event
Insecurity in farming belts	Banditry and farmer-herder conflict physically block extension access and ground-truthing in parts of the North-West and North-Central	A dynamic reduced-confidence zone layer feeding the score's confidence interval; expansion sequencing favours lower-insecurity belts first (see Section 6)
Currency devaluation	Naira volatility inflates imported fertiliser and agrochemical costs faster than farm-gate prices adjust, compressing margins the model priced in	Real-terms rather than nominal-naira margin modelling; quarterly, not annual, recalibration
Side-selling at harvest	Farmers contracted to one off-taker sell to a higher bidder, breaking sales-deduction repayment mechanisms	Side-selling propensity scored as its own feature using price-spread data and prior-season contract-vs-delivery ratios, not assumed away by the existence of a contract

4.3 Incentive Alignment — Why a Good Score May Still Be Refused

This is the pre-mortem's central finding, and the one most likely to be underestimated by a data-first team.

- **Cost-to-serve:** a well-scored ₦150,000 smallholder loan carries close to the same fixed underwriting and monitoring cost as a ₦15,000,000 SME facility, so unit economics push loan officers toward the larger ticket regardless of relative risk quality. A better score does not fix a structural cost-to-serve problem on its own.
- **Branch-level risk aversion:** loan officers are typically penalised for non-performing loans on their book far more than they are rewarded for agric-portfolio growth, so career incentives can override a favourable score. The mitigation is to bundle the score with a first-loss instrument — a NIRSAL Credit Risk Guarantee or comparable DFI cover — so the loan officer sees one decision object ("AA score, 70% guaranteed") rather than a score they still have to personally underwrite the tail risk on.
- **Funding-tenor mismatch:** many MFBs fund lending books with short-tenor deposits against agronomic cycles that run six to nine months for a single arable season and longer for tree crops. This argues for pointing the highest-conviction go-to-market motion at DFIs and blended-

finance vehicles with patient capital, not exclusively at deposit-taking banks whose balance sheets are structurally mismatched to the asset.

PRE-MORTEM CONCLUSION

The single biggest failure mode identified in this analysis is building a purely data-driven scoring product and assuming distribution follows automatically from model quality. It does not. The platform's defensible value proposition is *score plus de-risking instrument plus patient-capital access*, bundled as one underwriting decision — not a better number handed to an already risk-averse institution.

Business Model & Unit Economics

5.1 Revenue Streams

STREAM	VIABILITY	RATIONALE
API query / score-pull fees	High	Per-pull pricing (illustratively ₦500–₦2,000) mirrors how lenders already budget for existing bureau pulls – the easiest revenue line to sell because it needs no new budget category.
Origination / referral commission	Medium–High	0.5–1.5% of disbursed value aligns the platform's incentives with lender success, but must be firewalled from the model team – the scoring logic cannot be tunable by whoever owns the commercial referral relationship, or the score's independence (and its regulatory defensibility) is compromised.
Premium monitoring dashboards	High	Subscription pricing per cooperative cluster for off-takers and input distributors who carry direct commercial exposure through pre-financed inputs and forward contracts. Likely the more durable revenue base, since these customers want portfolio-level early warning continuously, not a one-off score.
Data monetisation / market intelligence	Medium	Aggregate yield, price, and risk-trend reporting sold to DFIs, insurers, and policy bodies is real but slower to build, and must respect NDPA anonymisation thresholds carefully to avoid re-identification risk in what are often small, geographically clustered cohorts.

5.2 Who Actually Pays

The farmer is not the paying customer and should not be modelled as one – willingness and ability to pay for a score are both near zero at this income level. The realistic sequencing runs in three waves. Input distributors and off-takers are the fastest first customer: they already extend informal credit through pre-financed inputs and forward contracts, they carry the commercial exposure directly, and the ROI story is the clearest to sell ("reduce input-credit default by a measurable margin"). MFBs and DFI-backed guarantee schemes form the second wave once a repayment track record exists to point to. Commercial-bank agric desks are realistically a third wave, arriving only once default-performance data from waves one and two has de-risked their adoption – banks in this segment move on evidence, not on model documentation.

5.3 Illustrative Unit Economics

Once the satellite and agent-network pipeline is built, the marginal cost of scoring one additional farmer — mostly amortised data-pipeline and agent-app sync cost — is modest relative to a ₦500–2,000 per-query price point. The economics only work, however, past a fixed-cost breakeven on the data pipeline, satellite API contracts, and field-agent network, which realistically needs a cohort in the range of 15,000–30,000 scored farmers to amortise meaningfully. This is the direct argument, developed further in Section 6, for a cluster-dense cooperative-first rollout rather than thin national coverage from day one. These figures are directional and illustrative for planning purposes; they are not an audited financial model.

5.4 Final Verdict on Fundability

FUNDABILITY ASSESSMENT

Fundable, but as an infrastructure and data-layer business under Pathway B — "the plumbing between agri value-chain data and formal capital" — not as a lender, and not as a bureau. Technically buildable within twelve to eighteen months. Commercially, the binding constraint is distribution into risk-averse lender balance sheets, not model quality, which means this should be underwritten by investors as a go-to-market and partnerships bet with a genuine data-science moat, not as a data-science bet with a go-to-market afterthought.

Given the multi-year path to lender-side monetisation, blended-finance and grant capital — AGRA, the IFC, FCDO, the Mastercard Foundation, FMO — is a realistically better fit for a first round than pure equity venture capital, which will want a faster line of sight to the commercial-bank third wave than this sector honestly offers in year one.

Execution Roadmap

Three phases, sequenced by proof rather than by calendar ambition — each phase has an explicit gate that must be cleared before the next begins.

PHASE 01

Pilot with Cooperatives

Months 1–9 · Target: 3,000–5,000 farmers

- Launch in two to three cooperative clusters within a single lower-insecurity agro-belt — the Oyo/Ogun cassava-cocoa corridor or the Benue/Kogi grain belt are the strongest starting candidates.
- Build the MVP pipeline: satellite NDVI overlay, NiMet weather integration, digitised cooperative repayment ledgers, and BVN/NIN deduplication.
- Ship the first interpretable scoring model with reason codes and validate it against actual repayment outcomes through one embedded pilot lender (an MFB or a NIRSAL-backed Bank of Agriculture branch).
- Secure the CBN no-objection confirmation on risk-SaaS positioning and complete NDPC registration and the first Data Protection Impact Assessment.

GATE TO PHASE 2

The model must show measurable lift over the pilot lender's existing (typically ad hoc) approval process, and at least one paying pilot contract — an input distributor or an MFB — must be signed.

PHASE 02

Financial Institution Integration

Months 10–20 · Target: 40,000–60,000 scored profiles

- Formal API integration with three to five licensed lenders, including at least one DFI-backed guarantee scheme (ideally NIRSAL), plus two to three off-taker or input-distributor accounts as parallel paying customers.
- Expand into four to six additional cooperative clusters across a second agro-ecological belt — Middle-Belt grains — while still avoiding high-insecurity LGAs.
- Formalise the parametric weather-insurance bundling partnership and lock in referral-commission and dashboard-subscription pricing.
- Build out the field-agent app and offline sync beyond MVP quality, and stand up the audit-log and adverse-action appeal workflow the NDPA requires.

GATE TO PHASE 3

Cost base approaches amortised breakeven and live default-performance data exists across at least two lender relationships.

PHASE 03

Scale

Months 21–36 · Target: 200,000+ scored profiles

- Multi-belt national coverage, excluding persistent-insecurity zones on a dynamically updated basis as ground conditions allow.
- Add commercial-bank agric desks as the third customer tier, now that Phase 2 default-performance data exists to de-risk their adoption.
- Layer in the data-monetisation and market-intelligence product for DFIs and policy bodies once the aggregate dataset is large and clean enough to anonymise safely.
- Evaluate a Series A raise on the strength of live default-rate performance — the only evidence that reliably moves risk-averse Nigerian lenders — rather than on model documentation alone.

OPTIONAL, NOT A COMMITMENT

Revisit CBN credit-bureau licensing only if proprietary scoring value and distribution scale eventually outweigh the capital lock-up examined in Section 3 — flagged explicitly as a possible Year 4+ option, not a near-term plan.

Risk Register & Governing Assumptions

RISK	LIKELIHOOD	IMPACT	MITIGATION
CBN deems risk-SaaS activity bureau-like	Medium	High	Written CBN no-objection before multi-lender API scaling; lender always pulls its own formal file
NDPC classifies platform "major importance" with enhanced obligations	High	Medium	Budget for DPO, DPIA, and ROPA from day one rather than retrofitting
Lenders adopt score but change nothing about approval behaviour	High	High	Bundle score with NIRSAL/DFI guarantee as a single decision object
Climate shock year breaks model calibration	High	Medium	Stress-scenario training data, quarterly recalibration, parametric insurance bundling
Insecurity blocks field ground-truthing in a cluster	Medium	Medium	Dynamic reduced-confidence zoning; sequenced expansion away from high-risk LGAs first
Fixed-cost breakeven not reached before capital runs out	Medium	High	Cluster-dense rollout (15–30k farmer cohorts) rather than thin national spread

This dossier is a strategic feasibility framework prepared for internal deliberation and investor circulation. It does not constitute licensed legal, regulatory, or financial advice; the regulatory citations in Section 3 reflect the Credit Reporting Act 2017, the Nigeria Data Protection Act 2023, and the NDPC's General Application and Implementation Directive as understood at time of writing, and should be confirmed with qualified Nigerian counsel before any licensing filing or lender contract is executed. Figures presented as illustrative are directional planning estimates, not an audited financial model. Prepared by Capital de Sinclair — Independent Strategic Advisory.