AEVION QVenture — Investment Memo

Generated 2026-11-26 · AEVION AI Investment Analyst · not investment advice

FICTIONAL EXAMPLE: the company in this example is made up; any resemblance of names is coincidental. All figures are model output on invented inputs.

FarmYield

AI Applications (vertical SaaS) · idea · IN

Score 54.1/100 — PASS (conviction: low)

Signal coverage: 0% company-specific (0 metrics parsed from the plan); remainder from sector priors.

Investment memo

Pass, with the door open. FarmYield attacks a genuine problem with credible technology, and its single strongest asset is a rare structural tailwind—a 36% CAGR sector plus a domain-validated prototype showing real engagement across 200 farms. But the decisive reason against is a fatal incentive conflict that no analyst could reconcile: monetizing through input-supplier referral fees pays the company to recommend *more* fertilizer and water, the scientific and economic opposite of the optimization it promises—which would destroy the very trust and retention that constitute its only traction signal. Compounded by zero disclosed unit economics, an unlaunched product (team signal 30/100), and well-capitalized or free public-good incumbents, the 54.1 composite is a pass. No ticket now. This is a watch-list name: we re-engage only if a *paid* pilot demonstrates real revenue-per-farm and CAC under a subscription or outcome-based model that resolves the referral conflict. On a re-score "e55, enter at ~$3.1M pre for ~1.3% of a diversified book—never single-name concentration.

Narrative engine: model anthropic (per-part source not recorded)

Entry strategy

Ticket: $184,428 target (range $92,214–$295,085)
Target ownership: 4.8%
Valuation band (pre-money): $1,012,300 / $3,073,800 / $6,147,600
Modelled return (not a forecast): 7.34x multiple (30.7x base) · horizon 9yr · loss prob 76%

Deployment schedule:

- 0% — Initial: Do not deploy. Add to watchlist.
- 100% — Re-entry: Only after a materially improved re-score ("e55) with fresh evidence.

Portfolio: Pass for now. If re-scored "e55 after new traction, size at ~1.3% of a diversified venture book — never single-name concentration at this stage.

Score breakdown

Market size & growth — 57/100 (weight 14%)
~$70B TAM, 36% CAGR (AI Applications (vertical SaaS)).

Timing / tailwinds — 100/100 (weight 5%)
Sector growth 36% vs. 12% neutral baseline.

Moat / defensibility — 44/100 (weight 16%)
switching costs is the category's mature moat (ceiling 74), but ~22% realized at idea given disclosed traction — an unproven moat is discounted toward the 35 "no demonstrated defensibility" floor.

Unit economics potential — 69/100 (weight 15%)
~70% mature gross margin, capital intensity 35% (sector reference).

Team / execution signal — 30/100 (weight 28%)
plan states the product has not launched / has no users

Scientific / tech feasibility — 98/100 (weight 7%)
agentic workflows, domain eval harnesses, retrieval + tool orchestration

Regulatory / legal headroom — 74/100 (weight 7%)
Regulatory intensity 40% (higher = more legal drag).

Competitive headroom — 41/100 (weight 8%)
Competitive intensity 85%. thin wrapper risk — value must accrue above the model layer.

Analyst council

Research Scientist — Ag-advisory science is real but 98/100 overstates feasibility; agronomic signal-to-action gap is the true technical risk

+ Push back on 98/100. The frontier tags (agentic workflows, retrieval orchestration) barely apply here — this is not an LLM product, it's a geospatial/agronomic inference product. The hard science is remote-sensing crop-water estimation (NDVI/NDWI, ET0 via Penman-Monteith, thermal soil-moisture proxies), where 10-30m Sentinel-2 pixels versus <2ha smallholder plots is a real resolution mismatch. Cre

+ Evidence base is thin: 200-farm prototype with an agronomist partner is validation of workflow, not of yield/water outcomes. The literature (e.g. ICRISAT/CGIAR digital-advisory RCTs, Precision Agriculture for Development SMS trials) shows advisory lifts input efficiency but effect sizes are modest and highly heterogeneous; there is no disclosed causal measurement of whether the advice moved yields

+ Fertilizer-timing advice needs soil-test/nitrogen data that satellite+weather alone cannot supply at plot level — the model can proxy but not measure N status. This bounds accuracy and creates liability if advice is wrong. Local-language SMS delivery is solved engineering, not a moat.

+ Structural thin-wrapper risk is misframed: the defensible asset here would be a proprietary labeled dataset (plot boundaries + realized outcomes across seasons + local crop calendars), not the model layer. An eval harness measuring recommendation accuracy against measured field outcomes would be the de-risking breakthrough.

! Agronomic accuracy at smallholder plot scale (<2ha) is unproven at 10m satellite resolution; wrong irrigation/fertilizer timing directly harms farmer income and referral trust, and no outcome-measurement RCT is disclosed.

! Incentive conflict: monetizing via input-supplier referral fees biases advice toward more inputs, which is scientifically opposite to water/N optimization — this undermines the product's credibility and could invalidate the retention signal once farmers detect it.

! Zero quantified plan metrics and no revenue/LTV-CAC means feasibility score rests on sector priors, not FarmYield-specific evidence; the 'strong WhatsApp retention' claim is unmeasured and likely selection-biased in a hand-picked pilot.

Data Analyst — 69/100 unit-economics score is a sector-prior mirage — referral-fee model on smallholder margins is unproven and likely far thinner

+ The 69 UE score inherits a ~70% SaaS gross margin prior, but FarmYield is NOT a subscription — it monetizes input-supplier referral fees. That's a take-rate business, not software margin. On smallholder input baskets of maybe $100-300/season with 3-8% referral rates, expected revenue is ~$3-15/farm/year — gross-margin % may be high but absolute contribution per farm is tiny. I'd push the score dow

+ TAM is misapplied: the $70B vertical-SaaS TAM is irrelevant. Bottom-up SOM = (reachable smallholders) x (referral revenue/farm). India has ~120M farm households; even 5M reached at $8/farm = ~$40M revenue — a real but modest ceiling, and the referral model creates an incentive conflict (advising more inputs = more fees) that can erode the trust the whole retention signal depends on.

+ CAC is the make-or-break unknown: SMS delivery is cheap, but acquiring/verifying/onboarding smallholders in local languages via agronomist channels is field-heavy. With ~$5-10 revenue/farm/year, payback fails unless CAC stays under ~$3-5 — implausible without a distribution partner (co-ops, lenders, telcos) subsidizing acquisition. No CAC/LTV disclosed; stress test not run.

+ Comps: agri-advisory referral models (e.g. DeHaat, CropIn, Plantix) monetize thinly and lean on B2B/enterprise or transaction rails, not pure advisory. Referral-only monetization has historically forced pivots into input distribution/fintech. Structural thin-wrapper risk applies — the satellite/weather model is commoditizable; defensibility must come from proprietary agronomic outcome data and dis

! Monetization-trust conflict: referral fees reward recommending inputs, which can bias advice and destroy the retention that is the only real traction signal — a fatal misalignment if not addressed with outcome-based or subscription pricing.

! Zero disclosed unit economics (no ARPU, CAC, LTV, referral rate) at idea stage on a low-ARPU population — the composite 54/pass and my UE sub-score both rest on sector priors, not company data. Thesis is unconfirmable until a paid pilot shows revenue/farm and CAC.

! Reachability/willingness-to-pay-for-inputs: smallholders are credit-constrained; if farmers can't act on fertilizer timing advice due to cash or supply access, the referral funnel converts near zero regardless of engagement.

Economist — Real agronomic value, but the $70B vertical-SaaS TAM is the wrong denominator for a referral-fee model serving sub-$2/day farmers.

+ The 57/100 TAM score is generous and mis-anchored: it applies a global AI-vertical-SaaS $70B/36% CAGR prior to a business whose actual monetizable pool is Indian agri-input referral commissions on smallholder purchases. India has ~120M smallholders, but ARPU is the binding constraint — fertilizer/seed spend per marginal farm is $100-300/yr, and a plausible 3-8% referral take implies $3-24 revenue/

+ Demand economics are favorable on the cost side (SMS delivery ~near-zero marginal cost, ~70% gross margin prior holds) but the willingness-to-pay is structurally zero at the farmer — value must be extracted from the input supplier, which caps pricing power and aligns FarmYield's incentives toward pushing purchases rather than optimizing yield (agency/trust conflict).

+ Moat at 44 is fair-to-slightly-generous: switching costs on free SMS advice are near nil, and the satellite+weather layer is a thin wrapper over commoditized inputs (Sentinel/ISRO data, public weather models). Durable rents, if any, accrue from the input-supplier distribution relationships and proprietary farm-level yield-response data — not the AI.

The 30/100 team score (no launch, no revenue) is
+ Macro/agri sensitivity is high: referral revenue is pro-cyclical with input spend, exposed to monsoon variance, MSP/subsidy policy shifts, and fertilizer-subsidy regime changes that dampen the elasticity FarmYield is trying to monetize.
! Monetization-model conflict: referral fees reward input sales, not water/fertilizer reduction — the pitch's core value prop (advising *less* input at optimal timing) directly reduces the revenue base. This unresolved incentive misalignment is the central economic flaw and is not addressed in the plan.
! TAM is unvalidated (triangulation not run, 0 quantified fields): the addressable pool is likely 1-2 orders of magnitude below the $70B prior once you condition on referral-only monetization and smallholder ARPU. Undisclosed LTV/CAC means unit economics remain a black box.
! Competitive intensity 85% with well-capitalized/subsidized incumbents (DeHaat, Cropin, government e-NAM / Kisan advisories, and free telco/NGO SMS services) — a free public-good alternative compresses any pricing rent and raises CAC in a trust-dependent, low-literacy channel.

Corporate & Regulatory Lawyer — 74/100 legal headroom is roughly right but the label 'Other/Unspecified' understates India-specific data, telecom-marketing and agent-referral drag

+ Push back on the generic score: this is India, not 'unspecified.' The Digital Personal Data Protection Act 2023 (rules being phased in 2024-25) governs the geolocation + farm data collected via satellite/SMS; smallholder farmers are 'Data Principals' whose consent must be verifiable in local languages — a real compliance build, not a footnote. Net, 74 is defensible but for the wrong reason (low se
+ Monetization is the sharpest legal exposure, not privacy: input-supplier referral fees for fertilizer/pesticide 'timing' advice risk being read as promotion of agri-inputs. Fertilizer (Fertiliser Control Order 1985) and pesticides (Insecticides Act 1968, and Pesticides Management Bill regime) are regulated; steering farmers toward specific SKUs for a fee invites conflict-of-interest, mis-selling a
+ SMS/WhatsApp channel triggers TRAI commercial-communication rules (DLT registration, sender-ID, consent scrubbing against DND) — cheap to comply with but a hard gating item before scale; unregistered bulk SMS is routinely blocked/fined.
+ Deal structure: the standard Delaware/Singapore flip is the right instinct, but India's FEMA/RBI pricing, DPIIT startup-recognition and (post-2020) Press Note 3 scrutiny on cross-border ownership add friction. Insist on Singapore/Delaware topco holding the Indian opco, with founder IP assignment executed at incorporation, plus data-processing terms flowing down to the agronomist partner and any sa
! Referral-fee model may be legally recharacterized as unlicensed input promotion or generate mis-selling/consumer-protection liability if AI advice causes crop loss on subsistence farms — reputational and class-exposure risk disproportionate to a pre-revenue, undisclosed-raise idea.
! Data provenance gap: satellite/weather models and WhatsApp pilot data have no documented DPDPA consent trail; retrofitting verifiable local-language consent across illiterate/low-literacy smallholders is operationally hard and a diligence red line.
! Honest counter-argument to my own concern: at idea stage with no revenue and 0% quantified plan metrics, all legal exposure is theoretical — none of these regimes bind a 200-farm prototype today, so the 7%-weight 74 score arguably overstates its relevance versus the 30/100 team signal that actually drives the 'pass.'

Market data sources

- Global Market Insights (2026) — Generative AI $83.3B in 2026 !’ $988.4B by 2035 at 31.6% CAGR
  https://www.gminsights.com/industry-analysis/generative-ai-market
- Fortune Business Insights (2026) — Generative AI $161B in 2026 !’ $1,260.2B by 2034 at 39.6% CAGR
  https://www.fortunebusinessinsights.com/generative-ai-market-107837

Assumptions & limitations

- Market size / growth for AI Applications (vertical SaaS) is anchored to Global Market Insights (2026): Generative AI $83.3B in 2026 !’ $988.4B by 2035 at 31.6% CAGR. Full citations are listed under "Market data sources".
- Signal coverage: ~0% of the score is backed by the plan's own disclosed metrics (0 quantified fields); the remainder uses AI Applications (vertical SaaS) sector priors — add financials to raise it.
- Stage norms reflect US-market idea deals; adjust for geography "IN".
- Score is a screening signal, not a substitute for legal, financial, and technical due diligence.

This memo is generated by an AI screening tool for research purposes and is not investment advice, an offer, or a solicitation. Figures are model estimates, not guarantees.