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.

OpsMind

B2B SaaS (horizontal) · seed · US · raising $4,000,000

Score 63.4/100 — WATCH (conviction: medium)

Signal coverage: 44% company-specific (1 metrics parsed from the plan); remainder from sector priors.

Investment memo

OpsMind is a Watch, not a lead: we should engage but withhold a term sheet until two data gaps close. The strongest reason for is genuine, quantified pain relief—median triage cut from 22 to 6 minutes across 40 teams—in a large, growing AIOps market where the switching-cost moat, if earned, is durable. The strongest reason against is that OpsMind's differentiation is a feature incumbents already own: Datadog, PagerDuty (its own integration partner), and Splunk are bundling LLM root-cause copilots for free, so willingness-to-pay may erode just as seat-based pricing collapses under AI-reduced headcount. Compounding this, the MRR figure is internally contradictory ($34k vs $408k), and no RCA accuracy, LTV/CAC, inference-COGS, or NRR data was disclosed. Plan: offer to co-invest up to $2M for ~9.9% at a ~$16M pre, staged—release a first tranche only against verified top-1 RCA accuracy, cohort retention/NRR, and per-incident inference margins. Reserve $3M for pro-rata. Pass if pricing stays seat-based.

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

Entry strategy

Ticket: $2,000,000 target (range $1,000,000–$2,000,000)
Target ownership: 9.9%
Valuation band (pre-money): $8,321,600 / $16,206,000 / $32,412,000
Modelled return (not a forecast): 6.71x multiple (16.2x base) · horizon 7yr · loss prob 59%

Deployment schedule:

- 40% — Entry: On close, after founder + IP + cap-table diligence.
- 35% — Milestone: Product-market fit signal (retention cohort / first repeatable revenue).
- 25% — Pro-rata: Reserve for next priced round to defend ownership.

Portfolio: Size at ~2% of a diversified venture portfolio (fractional-Kelly, conviction-scaled). Reserve 3,000,000 USD for pro-rata follow-on.

Score breakdown

Market size & growth — 67/100 (weight 14%)
~$465B TAM, 13% CAGR (B2B SaaS (horizontal)).

Timing / tailwinds — 53/100 (weight 5%)
Sector growth 13% vs. 12% neutral baseline.

Moat / defensibility — 59/100 (weight 16%)
switching costs is the category's mature moat (ceiling 74), but ~61% realized at seed given disclosed traction — an unproven moat is discounted toward the 35 "no demonstrated defensibility" floor.

Unit economics potential — 73/100 (weight 15%)
~78% mature gross margin, capital intensity 40% (sector reference).

Team / execution signal — 64/100 (weight 28%)
Quantified traction: $408k MRR.

Scientific / tech feasibility — 55/100 (weight 7%)
usage-based pricing telemetry, embedded analytics, PLG instrumentation

Regulatory / legal headroom — 81/100 (weight 7%)
Regulatory intensity 30% (higher = more legal drag).

Competitive headroom — 44/100 (weight 8%)
Competitive intensity 80%. seat-based model pressure as AI collapses headcount-linked demand.

Analyst council

Research Scientist — Automated RCA is real but 55/100 is fair — LLM root-cause accuracy is the unproven core, not the plumbing

+ The 55/100 score conflates the wrong frontier. The listed frontier (usage-based telemetry, PLG instrumentation) is commodity SaaS plumbing and is low-risk; the actual scientific challenge is automated root-cause analysis (RCA) over heterogeneous logs/traces, which is materially harder and where I'd push the score DOWN toward 50 absent accuracy data.

+ Grounded in a credible research line: causal inference over distributed traces (e.g. Facebook/Google SRE literature, AIOps anomaly detection, log-parsing work like Drain, and retrieval over past incidents). LLM-over-observability is technically plausible today — but published AIOps RCA top-1 accuracy typically lands 40-70%, meaning ~1 in 3 proposed root causes is wrong.

+ Traction (22!'6 min median triage across 40 teams) is the strongest evidence and suggests the copilot adds value even when imperfect, because a ranked shortlist beats a blank page. This is a triage-accelerant, not an autonomous fixer — a defensible framing that de-risks the science by keeping a human in the loop.

+ The auto-generated rollback plan is the riskiest claim: an incorrect rollback executed 'within seconds' can cause a second incident. Feasibility hinges on read-only recommendation vs. write-action; the memo doesn't disclose which, and that distinction drives most of the technical risk.

! Root-cause hallucination: if top-1 accuracy sits at ~50-60%, on-call engineers may lose trust after a few confident-but-wrong diagnoses during a Sev1 — churn risk that no amount of latency polish fixes. No accuracy/precision metric was disclosed; this is the single biggest gap.

! Data moat is thin at seed. RCA quality depends on incident-history volume; 40 teams is a small corpus, and much of the model IP is a wrapper over foundation-model APIs that competitors (Datadog Bits AI, PagerDuty AIOps, Incident.io) can replicate — consistent with the 44/100 competitive headroom.

! Cold-start and generalization: root-cause patterns are highly environment-specific; accuracy demonstrated on design-partner stacks may not transfer to novel architectures, and the structural seat-based-pricing risk compounds if buyers expect the AI to reduce, not add, tooling spend.

Data Analyst — 73/100 unit-econ score is a sector-inferred placeholder — OpsMind disclosed zero CAC/LTV data to earn it

+ The 73 rests on a ~78% sector gross margin and 40% capital intensity — NOT company data. But an AI copilot running live inference on logs/traces per incident carries real COGS (GPU/LLM tokens per alert); if inference isn't priced through, true gross margin could sit 15-25pts below the 78% reference. Score should be treated as unverified, not earned.

+ Real traction is $34k MRR across 40 teams (~$850/team/mo, ~$408k ARR) — the model's '$408k MRR' is a parsing error conflating ARR with MRR; do not underwrite off it. No CAC, no payback, no LTV, no churn disclosed, so no LTV/CAC ratio can be computed and the stress test could not run.

+ TAM ($465B horizontal SaaS) is uselessly broad — the honest SAM is the AIOps/incident-response slice (~$3-5B, Datadog/PagerDuty/Splunk adjacency). No bottom-up TAM (# US eng teams x ACV) was supplied, so SOM is unmodeled.

+ Product signal is genuine: 22!'6 min triage (73% reduction) is a quantified, sellable ROI story — the strongest reason the unit economics COULD mature toward the sector ceiling if they convert design partners to expansion contracts.

! Structural seat-based-model erosion (competitive headroom 44/100): if OpsMind prices per-engineer while AI copilots reduce on-call headcount, ACV compresses exactly as the product succeeds — must migrate to usage/incident-based pricing to survive.

! Inference COGS could gut the assumed 78% margin; at 40 teams there's no data proving per-incident LLM cost scales sub-linearly with revenue. This is the single biggest kill-metric and is entirely undisclosed.

! Moat is unproven (59/100): incumbents Datadog, PagerDuty (its own integration partner) and Splunk can bolt on comparable root-cause AI, turning OpsMind's differentiation into a feature — displacement risk before switching costs entrench.

Economist — OpsMind rides real AIOps pain, but the $465B horizontal TAM overstates a defensible slice; moat and pricing model are the real questions.

+ The 67/100 market score is generous framing. The $465B horizontal-SaaS TAM is not OpsMind's addressable market — AIOps/observability tooling is a ~$3-6B serviceable segment growing ~20%+, faster than the 13% horizontal blend but far smaller. I'd hold the score near 67 on a bottom-up basis (real, growing niche) rather than the top-down $465B, which is a category-of-convenience number.

+ Demand is inelastic where it matters: incident MTTR reduction (22->6 min) maps directly to revenue-loss-per-minute-of-downtime, so buyers pay on ROI not seat count. That partly INSULATES OpsMind from the disclosed seat-based structural risk — if it prices on incidents/hosts/data volume (the sector frontier of usage-based telemetry), AI-driven headcount collapse is a tailwind, not a threat. Pricing

+ Traction inconsistency must be flagged: the model cites $408k MRR while the brief states $34k MRR (40 teams, ~$850/team/mo). At $34k MRR the team/execution and moat scores are inflated — this is a very early seed with unproven retention, not a $408k ARR-run-rate business. Underwrite to $34k.

+ Moat accrues via workflow embedding (PagerDuty + Slack + accumulated incident history creating a data flywheel). Realized switching cost is thin at 40 teams; the defensible rent is the proprietary incident/RCA corpus per

customer — durable IF net retention >110% and design partners convert to multi-year.
! Competitive headroom 44/100 is the binding constraint: Datadog, PagerDuty (native AIOps), New Relic, and Grafana are all shipping LLM RCA copilots as free/bundled features. A disclosed free-adjacent incumbent capability threatens OpsMind's willingness-to-pay — a standalone copilot competes against 'good enough and already in the bill.'
! Data-flywheel moat is unproven: at 40 teams there is no evidence of retention or that accumulated incident history compounds accuracy faster than incumbents with 100x more telemetry. Root-cause hallucination in a rollback plan carries asymmetric downside (bad automated rollback = outage), capping trust and automation depth.
! No LTV/CAC, no cohort retention, and the MRR figure is internally contradictory ($34k vs $408k). Stress test and TAM triangulation both un-runnable — we are underwriting a ~$400k ARR-equivalent narrative on a business that may be at $34k MRR. Insist on cohort retention and net-revenue-retention data before term sheet.

Corporate & Regulatory Lawyer — United States: Reg D 506(b) round; regulatory intensity 30% — legal headroom 81/100.

+ Round structure: Private placement under Reg D of the Securities Act — 506(b) (no general solicitation, accredited + "d35 sophisticated) or 506(c) (general solicitation permitted, all-accredited with verification); Form D filed within 15 days.
+ Sector licensing (B2B SaaS (horizontal)): No sector licence expected — focus on consumer-protection (FTC), IP freedom-to-operate, and CCPA/CPRA data handling.
+ Structure entry with pro-rata rights, information rights, and standard downside protection.
+ Data/AI exposure: No federal omnibus law — CCPA/CPRA (California) plus sectoral regimes (HIPAA, GLBA); state patchwork. · No federal AI statute — FTC Act §5 enforcement, NIST AI RMF (voluntary), emerging state laws (e.g. CO AI Act).
! Jurisdiction-specific compliance and IP freedom-to-operate not yet verified.
! Confirm each investor's accredited status (Rule 501) and that solicitation matched the chosen exemption.
! IP, data-privacy, and liability exposure require local counsel review — this is directional, not legal advice.

Market data sources

- Grand View Research (2025) — SaaS $464.7B in 2025 !' $1,109.2B by 2033 at 11.1% CAGR
  https://www.grandviewresearch.com/industry-analysis/saas-market-report
- Research and Markets (2025) — $281.8B (2024) !' $774.3B by 2030 at 18.3% CAGR
  https://www.statista.com/outlook/tmo/public-cloud/software-as-a-service/worldwide/

Assumptions & limitations

- Market size / growth for B2B SaaS (horizontal) is anchored to Grand View Research (2025): SaaS $464.7B in 2025 !' $1,109.2B by 2033 at 11.1% CAGR. Full citations are listed under "Market data sources".
- Signal coverage: ~44% of the score is backed by the plan's own disclosed metrics (1 quantified field); the remainder uses B2B SaaS (horizontal) sector priors — add financials to raise it.
- Stage norms reflect US-market seed deals; adjust for geography "US".
- 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.