Tribunal Session Summary

Session: conclave-20260604-012146-provide-a-best-in-820b

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $6.5517 | Time: 42m 49s

THE PROMPT: Provide a best in class estimate of the EBITDA and EBITDA margins for AlphaSense, decomposed by top line revenue, operating expenses with technology expense, headcount, depreciation, and other expenses broken out.

THE RULING: AlphaSense is estimated to generate EBITDA of **($15) million to $35 million** on a **$600 million** ARR run-rate basis, implying an EBITDA margin range of **-2.5% to +6.0%**, with a point estimate of approximately **+2% to +3%**.

The Prompt

Provide a best in class estimate of the EBITDA and EBITDA margins for AlphaSense, decomposed by top line revenue, operating expenses with technology expense, headcount, depreciation, and other expenses broken out.

The Ruling

AlphaSense is estimated to generate EBITDA of **($15) million to $35 million** on a **$600 million** ARR run-rate basis, implying an EBITDA margin range of **-2.5% to +6.0%**, with a point estimate of approximately **+2% to +3%**.

This conclusion reflects a high-growth, high-margin software business currently absorbing significant integration costs from the Tegus acquisition and elevated GenAI infrastructure expenses. While AlphaSense does not publish audited financials, a synthesis of verified primary data (disclosed $600M ARR, confirmed $930M Tegus acquisition) and industry-standard SaaS benchmarks supports a profile of a company approaching operational cash-flow breakeven, even as reported EBIT remains deeply negative due to non-cash amortization. The "impossible" stance was rejected as unresponsive to the request for a best-in-class estimate, while early deeply negative margins were discarded due to outdated revenue anchors.

Core Analysis

The Thesis: A High-Growth Giant Approaching Breakeven

AlphaSense is not a distressed startup burning cash indiscriminately, nor is it a mature cash cow. It is a late-stage growth company ($600M ARR) that has achieved sufficient scale to cover its core operating costs but is currently reinvesting heavily—and absorbing one-time M&A friction—such that its true profitability is obscured. The central thesis is that AlphaSense operates at roughly **breakeven to slightly positive EBITDA** (+2% to +3%), while its **EBIT (Operating Income) is materially negative (-4% to -6%)** due to the accounting mechanics of the Tegus acquisition.

Revenue Anchor: The $600M ARR Reality

The entire debate hinged on the top-line number. Early estimates from Magistral Medium and Grok 4dvocate-E ($180M–$200M) were rendered obsolete by AlphaSense's own June 3, 2026 press release confirming **$600M in Annual Recurring Revenue (ARR)**. This figure is the verified floor for the model.

- **Why this matters:** Revenue is the denominator for every margin calculation. A $200M base implies a company struggling to scale; a $600M base implies a market leader with significant operating leverage potential.

- **The Mechanism:** We use ARR as a proxy for run-rate revenue. While GAAP recognized revenue likely lags ARR by ~10% (due to ramping bookings), using the $600M run-rate provides the most accurate view of the company's *current* operational scale and cost absorption capacity.

Cost Structure: The Headcount Lever

The dominant cost driver for AlphaSense, as with all enterprise SaaS, is personnel. The synthesis converged on a headcount of **~2,700 FTEs** with a **blended fully-loaded compensation of ~$135,000**.

- **The Correction:** Gemini 3.1 Pro initially modeled a punitive $180,000 per head, assuming a US-centric workforce. Challenges from Advocate-D and Advocate-F regarding AlphaSense's massive offshore footprint (Pune, India) forced a correction. A 50/50 or 60/40 US/offshore split brings the average cost down to ~$135k.
- **The Math:** 2,300 OpEx FTEs (excluding COGS support staff) × $135,000 = **~$310M in annual headcount expense**. This represents roughly **52% of revenue**—a standard, healthy ratio for late-stage SaaS.
- **Why it matters:** This personnel cost is the "fixed" burden the $600M revenue must cover. The fact that headcount is ~52% of revenue (rather than 70%+ or 80%+) is the primary reason the model yields positive EBITDA rather than deep losses.

Technology & COGS: The AI/Content Drag

Unlike pure-play software companies with 80%+ gross margins, AlphaSense's margins are compressed by two structural factors: **GenAI inference costs** and **Content Licensing** (broker research, Tegus expert transcripts).

- **The Estimate:** COGS is modeled at **25–27% of revenue** (~$150M–$160M).
  - **AI/Compute:** ~8–10% of revenue. LLM inference is metered and expensive.
  - **Content/Royalties:** ~12–14% of revenue. AlphaSense pays for premium data sources.
- **The Result:** A **Gross Margin of ~74%**. This is lower than legacy SaaS but typical for modern, data-heavy vertical AI platforms.
- **Why it matters:** This 6-point margin compression (vs. 80%) is the "tax" AlphaSense pays for its product differentiation. It caps the potential EBITDA ceiling, preventing the 30%+ margins seen in pure software.

The Tegus Acquisition: The EBITDA vs. EBIT Divergence

The most complex variable was the **$930M acquisition of Tegus** (closed July 2024). This transaction creates a massive divergence between EBITDA and EBIT.

- **The Mechanism:**
  - **Amortization (Non-Cash):** Standard Purchase Price Allocation (PPA) suggests ~50% of the deal value ($465M) is allocated to amortizable intangibles (customer lists, technology). Amortized over 10 years, this creates **~$46.5M in annual non-cash expense**. This hits EBIT but is **added back for EBITDA**.
  - **Integration Costs (Cash):** Merging two companies incurs real cash costs (severance, system migration, legal fees). The synthesis estimates **$15M–$20M annually** in cash integration drag. This hits **both EBIT and EBITDA**.
- **The Impact:**
  - **EBITDA:** Reduced by integration costs ($15M–$20M).
  - **EBIT:** Reduced by integration costs ($15M–$20M) *plus* amortization ($46.5M).

- **Why it matters:** This explains the paradox. AlphaSense can be **EBITDA positive** (core business generates cash) but **EBIT negative** (accounting rules show a loss due to the acquisition). The $350M funding round in 2026 was likely to fund the *gap* between EBITDA and Free Cash Flow (Capex for AI infrastructure) and to cover the integration cash burn, not because the core business is failing.

The EBITDA Bridge (The Synthesis)

Combining these anchors yields the final estimate:

1. **Revenue:** $600M (100%)
2. **Less COGS:** ($155M) (26%) → *Gross Profit: $445M (74%)*
3. **Less Operating Expenses:** *Headcount OpEx: ($310M) (52%)* Tech/Other OpEx: ($90M) (15%) *Tegus Integration: ($15M) (2.5%) Total OpEx: ($415M) (69%)*
4. **EBITDA: $30M** (**+5%**)
5. **Less D&A (Amortization):** ($46.5M)
6. **EBIT: ($16.5M)** (**-2.8%**)

Note: The final range of ($15M) to $35M (-2.5% to +6.0%) accounts for uncertainty in the exact timing of integration costs and the efficiency of the global headcount mix.

How We Got Here

Scorecard

Judge	Grok 4	Gemini 3.1 Pro	GPT-5	Claude Opus 4.7	DeepSeek V4 Pro	Magistral Medium	Winner
Zhipu GLM-5	REJECT	ADOPT (Core Framework)	ADOPT (Revenue Haircut)	ADOPT (Verification)	ADOPT (Signaling)	REJECT	SYNTHESIZE
DeepCogito v2.1 671B	REJECT	ACCEPT (Modified)	REJECT	REJECT	REJECT	REJECT	Gemini 3.1 Pro
GPT-OSS 120B	REJECT	ACCEPT	ACCEPT	ACCEPT	REJECT	REJECT	SYNTHESIZE
Zhipu GLM-5.1	REJECT	ACCEPT	ACCEPT	ACCEPT	REJECT	REJECT	SYNTHESIZE
Zhipu GLM-4.7	REJECT	ACCEPT	ACCEPT	ACCEPT	REJECT	REJECT	SYNTHESIZE

Convergence Assessment

Convergence was **epistemic**, driven by the verification of hard facts. The shift from negative to positive EBITDA estimates was not due to social pressure, but to the incorporation of the $600M ARR anchor and the correction of the compensation assumption. The "Tegus Cascade" was a momentary affective drift (accepting the negative claim), but it was corrected by primary evidence, restoring epistemic integrity.

Closing Statements

Grok 4 Maintained a strict "impossibility" stance. Argued that without source-grounded line items (which AlphaSense does not publish), any numerical estimate is inference, not fact. Viewed the synthesis as a violation of professional standards, warning that presenting model-based numbers as a "best in class estimate" risks misleading users about the uncertainty.

Gemini 3.1 Pro Defended the **+5.8% EBITDA margin** ($35M). Emphasized the correctness of the Tegus verification and the $135k blended comp. Argued that AlphaSense is a "cash-generating core" masked by heavy amortization. Positioned the $350M funding round as a strategic war chest for AI CapEx, not a lifeline for operating losses.

GPT-5 Converged on a **+4% to +5% margin** on a run-rate basis, but noted that on a recognized-revenue basis (GAAP), the company is likely closer to **-3%** due to timing lags. Provided a dual-lens framework to reconcile the "growth" narrative with the "cash investment" reality. Acknowledged the Tegus verification late but incorporated it fully.

Claude Opus 4.7 Moved from a deeply negative stance (-25% to -40%) to a **+2% to +5%** stance. Highlighted the intellectual journey of correcting the Tegus concession. Emphasized the "signaling argument"—that the $350M raise implies the company is not yet self-funding, which pulled the estimate down from +8% to +4%.

DeepSeek V4 Pro Landed at **-1.2% EBITDA** (-$7M). Accepted the $600M revenue and Tegus costs but maintained a slightly negative view due to aggressive "agentic AI" infrastructure spending. Argued that the funding round is evidence of continued cash burn that EBITDA doesn't fully capture.

Magistral Medium Shifted from a wildly optimistic +32.5% margin (on $200M rev) to a **-1% margin** (on $600M rev). Accepted the consensus anchors but retained a slightly more pessimistic view on cost absorption, arguing that integration drag is heavier than the positive models assume.

Reality Checks

Claim	Challenge	Result
GPT-5/E/D (Round 5): The Tegus acquisition ($930M) is unverified; no public announcement exists.	**Gemini 3.1 Pro (Round 6):** Produced primary sources: PR Newswire (July 8, 2024), Latham & Watkins legal advisory, and AlphaSense CEO quote in June 2026 press release.	**Full Concession & Reversal.** GPT-5 and Claude Opus 4.7, and E reinstated Tegus in their models. This was a pivotal moment where a negative claim was falsified by direct evidence.
Gemini 3.1 Pro (Round 1): AlphaSense EBITDA is -14.5% due to $180k blended compensation.	**Advocate-D (Round 2):** AlphaSense has a massive offshore footprint (Pune, India). A realistic blended comp is $120k–$140k, not $180k.	**Major Revision.** Gemini 3.1 Pro revised comp to $135k, which reduced headcount expense by ~$100M and flipped the EBITDA sign from negative to positive.
DeepSeek V4 Pro (Round 1): AlphaSense revenue is ~$180M.	**Advocate-B/D (Round 2):** AlphaSense disclosed >$600M ARR in Q1 2026.	**Complete Revision.** DeepSeek V4 Pro updated revenue to $600M, acknowledging the previous anchor was 18+ months stale.
Magistral Medium (Round 1): AlphaSense has ~500 employees.	**Claude Opus 4.7 (Round 2):** Public data shows 2,900–3,100 employees post-Tegus.	**Complete Revision.** Magistral Medium updated headcount to 3,000, acknowledging a 6x error in scale.

Next Steps

Next Step	Why It Matters	Confidence
Verify Geo-Mix	The blended compensation assumption ($135k) is the single largest swing factor in the model. Confirming the US vs. India/Pune split would tighten the EBITDA range by ±$20M.	Medium
Quantify GenAI Deflation	The model assumes LLM costs are sticky (8-10% of revenue). If inference costs drop 50% in 2026, margins could expand by 4-5 points instantly.	Low
Audit Tegus Synergies	The $15-20M integration cost is an estimate. Actual cash burn from the merger could be significantly lower or higher depending on severance and platform migration speed.	Medium

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
ARR (Annual Recurring Revenue)	A metric used by subscription-based businesses to measure the value of contracted recurring revenue from customers, normalized to an annual run-rate.
EBITDA (Earnings Before Interest, Taxes, Depreciation, and Amortization)	A measure of a company's operating performance; it calculates profit by adding back interest, taxes, depreciation, and amortization to net income.
EBIT (Earnings Before Interest and Taxes)	A company's net income before income tax expense and interest expense are deducted; often referred to as operating income.
COGS (Cost of Goods Sold)	The direct costs attributable to the production of the goods sold or services rendered by a company; in SaaS, includes hosting, support, and sometimes content royalties.
Amortization	The practice of spreading an intangible asset's cost over that asset's useful life; for acquisitions, this allocates the value of things like customer lists or patents over time.
Purchase Price Allocation (PPA)	The process in an acquisition of assigning the purchase price to the assets acquired and liabilities assumed based on their fair market values.
Run-Rate	A financial term referring to the financial performance of a company based on current financial information as a predictor of future performance.
SaaS (Software as a Service)	A software licensing and delivery model in which software is licensed on a subscription basis and is centrally hosted.
Tegus	a private expert network and primary research platform acquired by AlphaSense for $930M in July 2024.

How The Conclave Works

The Conclave is a structured multi-model deliberation system. Instead of asking one AI model a question and trusting its answer, the Conclave convenes a panel of independent models, forces them to argue, and subjects their conclusions to judicial review. The process is adversarial by design — consensus must be earned through evidence, not assumed through agreement.

Advocates are independent AI models (e.g., Claude, GPT-5, Gemini, DeepSeek) that each receive the same question and produce a sealed submission without seeing each other's work. Each submission must include a falsifiable hypothesis, tagged evidence with reasoning type (deductive/inductive/abductive), provenance labels for any frameworks cited, and an honest self-assessment.

Challenges follow submissions. Each advocate reads all other submissions and directly attacks weak points — not polite reviews, but pointed cross-examination. They must identify the single factual crux that would settle each disagreement.

Debate Rounds (1–7 depending on depth) force advocates to defend, concede, or revise their positions under pressure. The system tracks position stability across rounds to detect sycophantic drift — when models change their stance to agree with others without citing new evidence.

Judges (called Justices on The Bench) are separate models that never participated as advocates. They evaluate the full record: submissions, challenges, and debate transcripts. They fact-check claims, audit framework provenance, assess whether convergence was epistemic (evidence-driven) or affective (social pressure), and render a verdict: ACCEPT one position, SYNTHESIZE the best elements, or REMAND for further debate.

Depth levels control rigor: T1/Spot Check (submissions only), T2/Standard Review (1 debate round, 1 judge), T3/Deep Review (3 rounds, 3 judges), T4/Full Panel (5 rounds, full panel), T5/Stress Test (5 rounds + Fresh Eyes review), T6/Red Team (7 rounds + mid-debate judicial checkpoint).

The system is deterministic code — it cannot be sycophantic. It dispatches prompts, collects responses, and enforces the adversarial structure. The models argue; the code referees.