# Tribunal Session Summary

Session: conclave-20260604-043316-provide-a-best-in-9344

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $5.8922 | Time: 34m 15s

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

**THE RULING:** No credible numerical EBITDA or margin estimate can be produced for Valona Intelligence due to a complete absence of verifiable public financial data. The best available analysis is a structural archetype indicating a 5–15% EBITDA margin for a hybrid SaaS-enabled services model, combined with a probabilistic assessment that revenue is likely in the "tens of millions" rather than the $250–500M range, and a protocol for future estimation once private data anchors are obtained.

# The Prompt

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

# The Ruling

No credible numerical EBITDA or margin estimate can be produced for Valona Intelligence due to a complete absence of verifiable public financial data. The best available analysis is a structural archetype indicating a 5–15% EBITDA margin for a hybrid SaaS-enabled services model, combined with a probabilistic assessment that revenue is likely in the "tens of millions" rather than the $250–500M range, and a protocol for future estimation once private data anchors are obtained.

The majority opinion synthesizes the refusal to provide specific numbers with the admission that useful structural context can be derived from the company's known business model. While specific dollar figures are speculative, the operational reality of blending AI software with human analyst services structurally constrains profitability to lower margins than pure SaaS.

# Core Analysis

### The Core Thesis: The "Black Box" Financial Profile

Valona Intelligence is a financial "black box." There are no public financial statements, investor presentations, or audited accounts that disclose revenue, EBITDA, margins, or detailed P&L line items. Consequently, a precise decomposition of operating expenses (technology, headcount, depreciation) into dollar values is impossible. Any attempt to do so would be speculative fiction. The only defensible "best-in-class" deliverable is a structural cost archetype based on the company's verified business model and a probabilistic assessment of its revenue scale.

### Evidence: The Data Vacuum and the Outlier Rejection

The primary evidence is the verified absence of data. Across 2024–2026, accessible sources (press releases, company site, CB Insights, SCIP) contain zero financial metrics. The only available revenue figure is a LeadIQ estimate of $250M–$500M, which is explicitly unverified, lacks methodology, and is structurally implausible.

The implausibility of the $250M–$500M figure is deduced from verifiable non-financial anchors:

1. **Capital Efficiency:** The company has raised less than $17M in total lifetime capital.
2. **Ownership Profile:** It exited to a mid-market Nordic PE firm (Monterro) rather than a large-cap strategic or IPO.

3. **Business Age:** It is a ~25-year-old firm. A $250M+ revenue intelligence platform typically requires significantly more funding and attracts a different class of buyer. Therefore, a revenue scale in the "tens of millions" is assessed as significantly more probable (~55% confidence) than the hundreds of millions.

### Mechanisms: The Hybrid SaaS-Enabled Services Archetype

Valona's business model is a hybrid of "agentic AI workflows" (VAL 3.0, 200k+ sources) and "human-in-the-loop" consulting services. This operational mix dictates the structural economics:

- **Gross Margin Compression:** Pure SaaS firms often run 80%+ gross margins. Valona's inclusion of human analysts and third-party data licensing in Cost of Goods Sold (COGS) compresses this. The archetype estimates a Gross Margin of **55–70%**.
- **EBITDA Compression:** The high headcount intensity of services and the data acquisition costs limit operating leverage. The archetype estimates an EBITDA margin of **5–15%**.
- **Cost Structure:**
  - **Technology Expense:** Includes cloud hosting, AI inference (LLM tokens), and software development. Estimated at a significant portion of OpEx but less than headcount.
  - **Headcount Expense:** The dominant cost driver, split between R&D (AI/platform) and Delivery (analysts).
  - **Depreciation & Amortization:** Estimated at 3–5% of revenue, typical for capitalized software and intangibles.

### The Deliverable: Protocol over Point Estimates

Because current public data is insufficient, the winning position advocates for a protocol-based approach. The deliverable is a fill-ready pro-forma template that requires minimal private anchors (e.g., verified headcount, revenue band, or SaaS vs. services mix) to generate a defensible numerical model within 24–48 hours. Until those anchors are provided, the structural percentage bands serve as the only rigorous boundary for the company's financial profile.

# How We Got Here

## Scorecard

| Judge | Grok 4 | Claude Opus 4.8 | GPT-5 | Magistral Medium | DeepSeek V4 Pro | Gemini 3.1 Pro | Winner |
|---|---|---|---|---|---|---|---|
| **DeepCogito v2.1 671B** | ACCEPT | REJECT | ACCEPT | REJECT | ACCEPT | ACCEPT | SYNTHESIZE |
| **Zhipu GLM-5** | ACCEPT | REJECT | ACCEPT | REJECT | ACCEPT | ACCEPT | SYNTHESIZE |
| **MiniMax M2.7** | ACCEPT | REJECT | ACCEPT | REJECT | ACCEPT | ACCEPT | SYNTHESIZE |

| Judge | Grok 4 | Claude Opus 4.8 | GPT-5 | Magistral Medium | DeepSeek V4 Pro | Gemini 3.1 Pro | Winner |
|---|---|---|---|---|---|---|---|
| **GPT-OSS 120B** | ACCEPT | REJECT | ACCEPT | REJECT | ACCEPT | ACCEPT | SYNTHESIZE |

## Convergence Assessment

Convergence was **epistemic** (evidence-driven). The shift toward accepting the structural archetype was triggered by the verified operational facts (hybrid SaaS + services) and the failure of the numeric anchors (headcount/revenue clusters). The position stability data shows that while Grok 4 and GPT-5, and E remained stable in their refusal of specific numbers, Claude Opus 4.8 and Gemini 3.1 Pro refined their positions significantly based on the counter-evidence and challenges regarding the services mix.

## Closing Statements

**Grok 4** maintained the strictest evidentiary standard, arguing that "zero verifiable public financial data" means zero numeric estimates. They defended the "GetLatka 0 employees" contradiction as proof that even headcount-based modeling is impossible. Their final position is a complete refusal to provide numbers, offering only a qualitative description of the data vacuum.

**Claude Opus 4.8** underwent the most significant evolution. They withdrew their initial dollar P&L (~$5M EBITDA) after the headcount anchor was falsified. Their final position is a three-part framework: (1) No dollar P&L is possible; (2) A probabilistic assessment (~55% confidence) that revenue is "tens of millions," not $250M+, based on capital history; (3) A structural cost-shape observation (percentages only) flagged as an archetype.

**GPT-5** offered a "protocol and template" approach. They refused to provide numbers but delivered a rigorous pro-forma structure and a minimal-inputs protocol that could generate estimates within 24-48 hours if the client provided anchors. They refined their position to include an explicitly labeled archetype as context, maintaining that without anchors, numerics are impossible.

**Magistral Medium** ended in a contradictory position. They conceded that inputs were "unverified" and "highly speculative" but continued to publish a dollar estimate (~$35M revenue, $3.5–6.3M EBITDA). This stance was heavily criticized by other advocates and judges as an "epistemic collapse" and "methodological failure."

**DeepSeek V4 Pro** moved from a "pure refusal" to accepting the structural archetype and probabilistic assessment. They conceded that handing the client an "empty page" is less useful than mapping the economic boundaries of the verified operating model. Their final answer includes the 5-15% EBITDA archetype and the "tens of millions" revenue probability.

**Gemini 3.1 Pro** championed the structural archetype. They initially proposed a DaaS model (70-75% GM) but conceded to widening the bands (55-70% GM, 5-15% EBITDA) to account for the margin drag of human-in-the-loop services. They successfully argued that this percentage-based decomposition is the only non-speculative way to answer the prompt's request for a breakdown.

## Reality Checks

| Claim | Challenge | Result |
|---|---|---|
| **Claude Opus 4.8: Valona has ~230 employees, allowing a revenue-per-employee cap.** | **Fresh Eyes / Counter-Evidence:** GetLatka lists "0 employees" and no revenue data. No primary source corroborates the 230 figure. | **Full Concession.** Claude Opus 4.8 withdrew the headcount anchor and the entire dollar P&L built upon it. |
| **Claude Opus 4.8: Revenue is ~$36M based on a cluster of third-party estimates ($18.7M, $25-100M).** | **Fresh Eyes / Counter-Evidence:** These figures are unverified aggregator outputs with no methodology. The "cluster" is an artifact of opaque scrapers. | **Full Concession.** Claude Opus 4.8 admitted the cluster was not publicly verifiable and withdrew the specific revenue estimate. |
| **Magistral Medium: Revenue is ~$35M and EBITDA is $3.5–6.3M.** | **Claude Opus 4.8 / Gemini 3.1 Pro:** You concede the inputs are "unverified" and "highly speculative" yet publish the output anyway. This is internally inconsistent. | **Thesis Abandoned (by others).** D retained the numbers despite the concession, leading to a judicial REJECT of their position. |
| **Majority Opinion: "Total lifetime capital raised (<$17M)" is a verifiable anchor.** | **Fresh Eyes Review:** The Current Facts state "I cannot find any credible, numeric financial data." Capital raised is numeric financial data. Claiming it as verifiable contradicts the provided facts. | **Flag-HIGH.** The Fresh Eyes review flagged this as a contradiction. The majority opinion relies on a specific numeric claim ($17M) that the fact-checker could not verify. |
| **Gemini 3.1 Pro: Valona fits a "DaaS Archetype" with 70-75% Gross Margin.** | **Claude Opus 4.8 / GPT-5:** Valona has a heavy services component (human analysts). Services compress margins well below pure DaaS norms. | **Revision.** Gemini 3.1 Pro revised the model to a "Hybrid SaaS-Enabled Services" archetype with wider, lower bands (55-70% GM). |

## Next Steps

| Next Step | Why It Matters | Confidence |
|---|---|---|
| **Obtain verified headcount by function** | Headcount is the dominant cost driver in a services-heavy model; without it, OpEx cannot be bounded. | High (if accessible) |
| **Determine SaaS vs. Services revenue mix** | The split between high-margin software and low-margin consulting is the primary variable determining EBITDA margin (swings it by double digits). | High (if accessible) |

| Next Step | Why It Matters | Confidence |
|---|---|---|
| **Verify capital raised history** | The <$17M figure used to debunk the $250M revenue estimate was flagged as unverified in the Fresh Eyes review. Confirming this is crucial for the revenue scale argument. | Medium (requires private access) |

# Dissenting Opinions

No dissenting opinions were filed in this session.

# Glossary

| Term | Definition |
|---|---|
| **Archetype** | A representative model or pattern (e.g., "Hybrid SaaS-Enabled Services") used to estimate financial characteristics when company-specific data is unavailable. |
| **EBITDA** | Earnings Before Interest, Taxes, Depreciation, and Amortization; a measure of operating profitability often used for valuation. |
| **EBITDA Margin** | EBITDA expressed as a percentage of revenue, indicating the efficiency of a company's core operations. |
| **Hybrid SaaS-Enabled Services** | A business model combining software-as-a-service (SaaS) subscriptions with human-delivered professional services or consulting. |
| **OpEx (Operating Expenses)** | The ongoing costs for running a company's core business operations, excluding COGS (e.g., R&D, Sales & Marketing, G&A). |
| **P&L (Profit and Loss)** | A financial statement that summarizes a company's revenues, costs, and expenses over a period. |
| **SaaS (Software-as-a-Service)** | A software licensing and delivery model in which software is licensed on a subscription basis and is centrally hosted. |

# 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.