Tribunal Session Summary

Session: conclave-20260513-204209-which-three-companies-is-86a4

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $5.3955 | Time: 46m 38s

THE PROMPT: Which three companies is NVIDIA most likely to acquire next? Include a probability estimate for each potential acquisition and discuss the strategic reasons behind these choices.

THE RULING: NVIDIA is most likely to acquire AI21 Labs, Ayar Labs, and WEKA next, with probability estimates of 65%, 30%, and 25% respectively. This reflects a strategic pivot from horizontal hardware consolidation to vertical integration of critical AI infrastructure layers—software talent, photonic I/O, and storage optimization—while avoiding antitrust triggers from acquiring direct competitors. The strongest evidence is the verified "advanced talks" with AI21 Labs, the physical necessity of co-packaged optics to overcome Rent's Rule, and the distinct storage bottleneck addressed by WEKA's parallel filesystem, all of which survived rigorous adversarial testing.

The Prompt

Which three companies is NVIDIA most likely to acquire next? Include a probability estimate for each potential acquisition and discuss the strategic reasons behind these choices.

The Ruling

NVIDIA is most likely to acquire AI21 Labs, Ayar Labs, and WEKA next, with probability estimates of 65%, 30%, and 25% respectively. This reflects a strategic pivot from horizontal hardware consolidation to vertical integration of critical AI infrastructure layers—software talent, photonic I/O, and storage optimization—while avoiding antitrust triggers from acquiring direct competitors. The strongest evidence is the verified "advanced talks" with AI21 Labs, the physical necessity of co-packaged optics to overcome Rent's Rule, and the distinct storage bottleneck addressed by WEKA’s parallel filesystem, all of which survived rigorous adversarial testing.

The ruling synthesizes the most robust elements from multiple advocates: the in-flight deal logic for AI21 Labs (Gemini 3.1 Pro, -D), the photonic I/O imperative for Ayar Labs (Kimi K2.6, -E), and the storage bottleneck case for WEKA (GPT-5, -F). It rejects hardware-centric picks like SambaNova due to fatal antitrust math (HHI ≥ 6,400), community-sensitive targets like Hugging Face due to governance barriers, and channel-conflicting plays like Together AI. The synthesis was necessary because no single advocate presented a flawless case—Gemini 3.1 Pro’s Lumentum fabrication and Grok 4’s HHI error were critical judicial red lines—but the convergence on infrastructure tuck-ins was epistemic, driven by technical and regulatory realities rather than social pressure.

Core Analysis

The Post-Groq Acquisition Strategy: Vertical Integration Over Horizontal Consolidation

NVIDIA’s acquisition of Groq for $20 billion in May 2026 marked a decisive shift from pure GPU dominance to full-stack AI infrastructure control. With ~$60 billion in cash reserves and ~80% market share in AI inference, NVIDIA now faces two existential constraints: the laws of physics (the I/O wall) and the laws of governments (antitrust scrutiny). Its next acquisitions will not be horizontal competitors but vertical enablers—companies that solve critical bottlenecks in the AI stack without triggering regulatory or ecosystem backlash. The synthesis identifies three such targets: AI21 Labs for software talent, Ayar Labs for photonic I/O, and WEKA for storage optimization.

1. AI21 Labs: Securing Israeli AI Talent for Inference Software (65% Probability)

Strategic Rationale: AI21 Labs is the highest-probability next acquisition because it is already in advanced talks with NVIDIA, as confirmed in the Current Facts. The $2–3 billion deal targets AI21's ~200-person AI research and model development team, reinforcing NVIDIA's push into proprietary agentic reasoning models that will run exclusively on its hardware. This aligns with NVIDIA's broader strategy of acquiring top-tier AI talent, particularly in Israel, where it has already completed three major acquisitions (Mellanox, Deci, Run:ai). By owning the software layer, NVIDIA can lock in inference economics, much like it did with CUDA for training.

Probability Derivation: The 65% probability is a calibrated estimate based on the typical closure rate for advanced-stage tech M&A (~70%) adjusted downward for post-Groq antitrust scrutiny. Historical precedents like the failed NVIDIA-Arm deal (abandoned 2022) and Adobe-Figma deal (abandoned 2023) show that even in-advanced talks can collapse due to regulatory, valuation, or retention issues. The 65% figure acknowledges this risk while recognizing that AI21 is a software/talent acquisition, not a horizontal hardware play, making it less likely to trigger a regulatory block.

Why It Survived Debate: The AI21 Labs thesis was challenged on procedural grounds—whether "advanced talks" constitute a prediction or a transcription. The tribunal ruled that it is a valid prediction because the deal is not yet completed. The probability was further tested against antitrust concerns, but the software/talent nature of the deal was deemed low-risk compared to hardware consolidation. No advocate presented a non-AI21 target with a higher mathematically sound probability.

2. Ayar Labs: Solving the Photonic I/O Wall (30% Probability)

Strategic Rationale: Ayar Labs addresses the literal physical limit of copper interconnects in next-generation GPU architectures. As compute density scales, Rent's Rule—a well-established semiconductor scaling law—dictates that I/O bandwidth must scale proportionally to sustain performance. Electrical I/O (e.g., NVLink over copper) cannot meet the bandwidth demands of post-Blackwell systems without melting or consuming unsustainable power. Co-packaged optics (CPO), which integrate photonics chiplets directly into the GPU package, are the only credible path forward. Ayar Labs possesses the proprietary silicon photonics IP required to break this "I/O wall."

Probability Derivation: The 30% probability reflects the technical necessity of CPO versus the uncertainty of whether partnership suffices. NVIDIA has a partnership with TSMC on the COUPE platform for optical integration, but ownership of Ayar Labs would provide tighter control over performance, margins, and roadmap execution—similar to the strategic rationale behind the Mellanox acquisition. The 30% figure accounts for a ~60% confidence that acquisition is necessary, given the deep integration required at the die-package level.

Why It Survived Debate: The Ayar Labs thesis was challenged on timing and partnership viability. Critics argued that NVIDIA could rely on TSMC's consortium work rather than acquire. The tribunal found that CPO sits permanently inside NVIDIA's proprietary silicon die package, making partnership insufficient for optimal performance and margin control. The Rent's Rule argument, tagged as [TEXTBOOK], was verified by multiple judges and survived all counterattacks.

3. WEKA: Optimizing the Storage Bottleneck for Inference Clusters (25% Probability)

Strategic Rationale: WEKA solves the critical data-feeding bottleneck in AI clusters. As model training scales into exascale regimes, the ability to feed data rapidly to GPUs becomes a limiting factor. WEKA's parallel filesystem is optimized for GPU workloads, offering high-throughput metadata and file I/O that outperforms general-purpose storage solutions. Unlike block-storage providers like Excelero or object-storage platforms like VAST, WEKA operates at the filesystem layer, enabling POSIX-compliant, low-latency access to training data. This makes it a natural complement to NVIDIA's DGX and HGX platforms.

Probability Derivation: The 25% probability accounts for the mechanistic fit and precedent of owning full-stack infrastructure (e.g., Spectrum for networking) but is discounted for partner-alienation risk. NVIDIA's strong partnership with VAST Data and WEKA's need for neutrality reduce acquisition urgency. The 25% figure reflects a ~65% confidence in the technical imperative, offset by ~30% weight for hyperscaler retaliation risk.

Why It Survived Debate: The WEKA thesis was challenged on redundancy with existing solutions (Excelero, VAST). The tribunal found that WEKA's parallel filesystem addresses a distinct bottleneck—metadata scaling and namespace contention—that block and object storage do not solve. The technical distinction was verified by multiple judges and survived all counterattacks.

Rejected Alternatives

- **SambaNova, Tenstorrent, Lightmatter**: Rejected due to antitrust risk. With NVIDIA's dominant market share (~80%) in AI inference and the recent Groq acquisition, acquiring another inference-focused hardware competitor would face near-certain regulatory scrutiny. Grok 4's HHI miscalculation (claiming 2,800 vs. the correct 6,400+) further undermined this thesis.
- **Hugging Face**: Rejected due to governance and structural barriers. Despite strategic fit in open-source AI, Hugging Face's decentralized governance, AWS stake, and community-driven model make acquisition unlikely.
- **Modular**: Rejected due to platform conflict and fork risk. While Modular's Mojo language could lock developers into the NVIDIA stack, acquisition risks community backlash and a fork, as seen with Oracle-MySQL.
- **Ampere Computing**: Rejected as a "falling-knife" scenario. Hyperscalers are exiting the merchant CPU market, and Ampere conflicts with NVIDIA's Grace CPU strategy.
- **Recursion Pharmaceuticals**: Rejected due to pattern mismatch. NVIDIA has historically taken minority stakes in biotech firms (e.g., Recursion, Schrödinger), not full acquisitions.

Fact-Check Summary

Claim	Verdict	Source
NVIDIA in advanced talks to acquire AI21 Labs for $2–3B	✓ Verified	Current Facts (2026-05-??), stocktwits.com, tradingview.com
NVIDIA acquired Groq for $20B in cash	✓ Verified	Current Facts (2026-05-??), ad-hoc-news.de

Claim	Verdict	Source
NVIDIA holds ~$60B in cash reserves (Q1 2026)	✓ Verified	Current Facts (2026-05-??), ad-hoc-news.de
NVIDIA Riva provides native voice AI capabilities	✓ Verified	Qwen 3 235B (Cerebras) and Judge-Qwen 3.5 397B (PRIMARY_SOURCE)
Hugging Face is France-based	✗ Incorrect	DeepCogito v2.1 671B and Judge-Qwen 3.5 397B (US-headquartered)
$2B NVIDIA investment in Lumentum (March 2026)	✗ Incorrect	GPT-OSS 120B and Judge-Qwen 3.5 397B (UNATTRIBUTED, admitted fabrication)
HHI for AI inference market is ~2,800	✗ Incorrect	Qwen 3 235B (Cerebras) and Judge-Qwen 3.5 397B (80% share implies ≥6,400 HHI)
Rent's Rule predicts I/O growth limits in GPUs	✓ Verified	DeepCogito v2.1 671B and Judge-GPT-OSS 120B, Qwen 3 235B (Cerebras) (TEXTBOOK)
WEKA provides a parallel filesystem distinct from block storage	✓ Verified	Qwen 3.5 397B and Judge-Qwen 3 235B (Cerebras) (SECONDARY, technical whitepapers)
NVIDIA invested $3.5M in SoundHound (2021)	✓ Verified	GPT-OSS 120B (PRIMARY_SOURCE, SEC filing)
Run:ai was acquired by NVIDIA for $700M	✗ Incorrect	Judge-F (not in Current Facts)

How We Got Here

Scorecard

Judge	Grok 4	DeepSeek R1	Kimi K2.6	GPT-5	Gemini 3.1 Pro	Claude Opus 4.7	Winner
DeepCogito v2.1 671B	REJECT	REJECT	ADOPT (partial)	ADOPT (partial)	REJECT	ADOPT (partial)	Kimi K2.6
GPT-OSS 120B	REJECT	REJECT	REJECT	REJECT	ACCEPT	REJECT	Gemini 3.1 Pro

Judge	Grok 4	DeepSeek R1	Kimi K2.6	GPT-5	Gemini 3.1 Pro	Claude Opus 4.7	Winner
Qwen 3.5 397B	REJECT	REJECT	REJECT	ADOPT (partial)	REJECT	ADOPT (partial)	GPT-5
Qwen 3 235B (Cerebras)	REJECT	REJECT	ADOPT (partial)	ADOPT (partial)	ADOPT (partial)	ADOPT	Claude Opus 4.7
MiniMax M2.7	PARTIAL	PARTIAL	PARTIAL	PARTIAL	PARTIAL	PARTIAL	Claude Opus 4.7

Convergence Assessment

The convergence was epistemic, driven by verifiable corrections rather than social pressure. Key evidence includes Grok 4's concession on the HHI math error, Gemini 3.1 Pro's retraction of the Lumentum investment, and Claude Opus 4.7's abandonment of Together AI due to channel conflict. These shifts were triggered by factual and logical challenges, not by groupthink. Position stability data shows that advocates with the highest sycophancy risk (e.g., DeepSeek R1 and Grok 4dvocate-C) made no changes, while those who changed (e.g., Grok 4 and Grok 4dvocate-E, Claude Opus 4.7) did so in response to evidence. The tribunal successfully filtered out flawed positions through adversarial testing.

Closing Statements

Grok 4: After rigorous debate, my evolved position is that NVIDIA is most likely to acquire Lightmatter, Tenstorrent, and SambaNova Systems next, with revised probability estimates of 50%, 35%, and 15% respectively. These choices reflect NVIDIA's strategic imperative to consolidate its dominance in AI hardware by targeting complementary innovators in emerging compute paradigms, while navigating heightened antitrust scrutiny in concentrated markets. This represents a refinement from my opening hypothesis, which initially prioritized SambaNova at 65% due to its inference overlap with the recent Groq acquisition, followed by Tenstorrent at 50% for talent and efficiency, and Lightmatter at 40% for photonic innovation. The evolution stems from a critical concession on market concentration metrics, which exposed overoptimism in regulatory clearance for high-overlap deals, prompting me to demote SambaNova and elevate Lightmatter as a lower-risk vertical play. My framework for assessing likelihood weights strategic alignment (60%)—how well a target fills gaps in NVIDIA's AI stack—against financial feasibility (20%) and regulatory viability (20%), calibrated via a Bayesian approach updating priors from NVIDIA's historical acquisition patterns (e.g., 80% closure rate on advanced talks, per ad-hoc-news.de) with debate-tested evidence.

DeepSeek R1: After rigorous debate testing, I maintain that NVIDIA's next three most likely acquisition targets are Hugging Face (20% probability), WEKA (20%), and Recursion Pharmaceuticals (20%). This equal weighting reflects NVIDIA's balanced strategic trifecta: consolidating AI developer ecosystems (Hugging Face), solving infrastructure bottlenecks (WEKA), and capturing high-value vertical markets (Recursion). My framework prioritizes three non-negotiable dimensions: strategic synergy (50% weight), acquisition feasibility (30% weight), and market insulation (20% weight). This evolved stance—from an initial Hugging Face-centric view—recognizes NVIDIA's post-AI21 Labs pivot toward vertical domination and

infrastructure hardening. Hugging Face remains critical for AI mindshare but now shares priority with WEKA's data-layer solutions and Recursion's biopharma monetization. NVIDIA isn't merely acquiring companies; it's assembling an end-to-end stack where hardware (Grace/Groq), software (Hugging Face/BioNeMo), and infrastructure (WEKA) converge to control AI's full value chain.

Gemini 3.1 Pro: Nvidia's overarching strategic imperative is no longer about winning the merchant GPU market; it is about transitioning from a silicon vendor into a vertically integrated, full-stack AI datacenter platform. At a $4.6 trillion valuation with an estimated 80% market share in AI compute, the only two forces capable of halting Nvidia's momentum are the laws of physics (the memory bandwidth and power walls) and the laws of governments (horizontal antitrust scrutiny from the FTC and the European Commission). My framework for predicting Nvidia's next acquisitions rests entirely on navigating these two immovable constraints. Any viable acquisition target must solve a critical physical or software bottleneck in the AI factory without triggering regulatory vetoes or alienating key customers. Early in this deliberation, I proposed a "Vertical Tri-Modal Defense Strategy," positing that Nvidia would strictly bypass horizontal compute consolidation to evade antitrust enforcement. I argued they would instead target elite agentic reasoning talent, internalize the optical networking supply chain, and capture adjacent software orchestration layers. While my core analytical framework remained intact throughout the debate, my specific targets underwent a rigorous, humbling, and ultimately necessary evolution. I initially proposed AI21 Labs, Lumentum, and Covariant. I briefly flirted with Together AI in the middle rounds. Today, my final and definitive slate stands as: AI21 Labs (65%), Ayar Labs (30%), and WEKA (25%).

GPT-5: After testing my initial slate against hard challenges, I stand by the view that NVIDIA's most likely next acquisitions over the coming 12 months are: (1) AI21 Labs (65% probability), (2) WEKA/WekaIO (25%), and (3) Ayar Labs (20%). These probabilities are independent, not a sum-to-100 distribution. My framework is simple: NVIDIA's deal logic clusters around three imperatives—lock down AI inference economics (software, models, services), eliminate I/O bottlenecks end-to-end (from fabric to storage to runtime), and secure top technical talent with fast time-to-integration (notably in Israel). All three targets score high on these axes while staying within regulatory and ecosystem guardrails, and within the balance-sheet capacity NVIDIA still has after the Groq agreement. Why these weights? First, base rates matter. Current facts confirm NVIDIA is in advanced talks with AI21 Labs, and that the company has been explicitly pursuing Israeli talent through multiple prior acquisitions. On strategy, NVIDIA has been moving quickly up the software stack—NIM, Triton, TensorRT-LLM—and just agreed to buy Groq to dominate inference throughput. AI21 is a clean extension: enterprise-grade models and tools that can be embedded into NIM and tuned to run best on NVIDIA hardware. The risk is sequential scrutiny after Groq and the general non-triviality of turning "advanced talks" into a closed deal. That is precisely why this is a real prediction and not a transcription: late-stage deals routinely fail at valuation, diligence, or remedies. I trimmed AI21's probability from 70% to 65% to reflect that friction while keeping it the top pick because its strategic fit, geography, and size all line up.

Claude Opus 4.7: After six rounds of adversarial testing, my final slate of NVIDIA's three most likely next acquisitions is: WEKA (~22%), Modular (~15%), and Ayar Labs (~12%). I want to be candid about what these numbers mean. None of my three picks clears 25%. That is not hedging — it is the honest output of the framework I now hold. The framework. NVIDIA's M&A behavior after Groq ($20B) and AI21 (~$2-3B in-flight) is constrained by three binding forces operating simultaneously: antitrust window compression, channel preservation, and seller willingness and digestibility. The three picks that survive these filters are all infrastructure-layer plays with no antitrust footprint, no channel cannibalization, and structurally weak independent paths. WEKA, Modular, and Ayar Labs each occupy a chokepoint in the AI compute stack where NVIDIA today rents capability that it could own. The intellectual journey. My opening submission

predicted Together AI (35%), Recursion (22%), and SambaNova (18%) — a "vertical integration + inference cloud consolidation" thesis derived directly from Williamsonian margin-capture logic. That thesis was theoretically sound but empirically naïve about the antitrust and channel constraints that bind NVIDIA in 2026. Under pressure from D, E, and B, I systematically dropped each of my original three picks: SambaNova fell to my own counter-attack on Grok 4's HHI math. If NVIDIA has 80% inference share, HHI from that single firm alone is 6,400 — far above the 2,500 presumptive-challenge threshold. Acquiring a second inference-silicon firm is regulatorily implausible. I could not defend SambaNova while holding that arithmetic against A. Together AI fell to GPT-5's channel-conflict challenge in the final round. I had already conceded CoreWeave overlap as a risk; holding Together at 15% was internally incoherent. Recursion fell earlier (Round 6) when the pharma-partnership integration complexity (Bayer, Roche contractual entanglements) was shown to be a multi-year integration friction NVIDIA would not absorb in the current antitrust window. What replaced them — WEKA, Modular, Ayar Labs — are objectively less ambitious picks. They are also better-calibrated picks. That is the trade I made.

Reality Checks

Claim	Challenge	Result
Grok 4: SambaNova's SN40L chip achieves 1.5x faster inference than NVIDIA's H100 on certain LLMs (SambaNova benchmarks, 2025).	**GPT-5:** This appears vendor-self-reported and not independently validated under standardized conditions. Explain why you’re elevating SambaNova’s marketing claims to decision-grade evidence, and cite third-party corroboration or apples-to-apples MLPerf-style data.	**Grok 4:** I concede that the 1.5x benchmark is primarily self-reported and lacks sufficient independent validation, as highlighted in the counter-evidence brief, which notes no broad third-party corroboration and potential bias in hand-picked workloads. This weakens the evidence for SambaNova's technical edge as a primary driver for acquisition.
Gemini 3.1 Pro: In March 2026, NVIDIA injected $2 billion into Lumentum.	**Claude Opus 4.7:** This appears unverifiable because the current facts (verified 2026-05-13) make no mention of a Lumentum investment, and your deductive premises on copper limits cite no sources tying it directly to NVIDIA's roadmap. Explain.	**Gemini 3.1 Pro:** You are absolutely right. I failed the basic standard of evidence by inventing the Lumentum investment. My deductive chain crumbles without it. Lumentum is retracted.

Claim	Challenge	Result
DeepSeek R1: NVIDIA’s DRIVE platform lacks native voice capabilities, making SoundHound a strategic fit.	**GPT-5:** How do you defend this given NVIDIA Riva (ASR/TTS/NLP) has long shipped as NVIDIA’s native speech stack and integrates with DRIVE/Jetson? Your premise is factually wrong; reconcile this with the case for buying SoundHound.	**DeepSeek R1:** This is a critical factual error. NVIDIA Riva does provide native ASR/TTS capabilities integrated with DRIVE, as evidenced by Riva’s documentation and automotive deployments (e.g., Mercedes-Benz at GTC 2025). SoundHound’s "100B+ queries" metric is overshadowed by Riva’s maturity, making acquisition redundant. I retract SoundHound as a top-3 target.
Grok 4: The FTC and European Commission will allow NVIDIA to buy another top inference-silicon competitor within 12 months of Groq.	**Claude Opus 4.7:** How does FTC/EC clearance work for the dominant AI accelerator firm buying its #2 and #3 inference-silicon rivals in the same year? Quote a single precedent.	**Grok 4:** I revise SambaNova's probability downward to 45%, acknowledging that the antitrust risks are higher than I initially weighted, especially in light of the counter-evidence brief's strong emphasis on regulatory scrutiny for horizontal consolidation in inference silicon.
Gemini 3.1 Pro: AI21 Labs is the most likely next acquisition with a 90% probability.	**Claude Opus 4.7:** How do you defend 90% closure probability from “advanced talks”? Deals at this stage routinely fail on diligence, price, or regulatory issues. Provide a base-rate or precedent that justifies a near■certainty figure rather than 60–70%.	**Gemini 3.1 Pro:** I concede that a 90% probability for an unclosed M&A deal ignores the ~30-40% baseline failure rate of advanced tech M&A due to diligence or regulatory roadblocks (e.g., Nvidia-Arm). I revise it downward to 65%.

Next Steps

Next Step	Why It Matters	Confidence
Confirm the closure of the AI21 Labs acquisition.	This is the highest-probability next deal; its closure would validate the vertical integration strategy.	65%
Monitor NVIDIA's progress on co-packaged optics through TSMC's COUPE platform.	This will determine whether an Ayar Labs acquisition is necessary or if partnership suffices.	60%

Next Step	Why It Matters	Confidence
Assess the performance of WEKA's parallel filesystem in NVIDIA's DGX Cloud reference architectures.	This will indicate whether WEKA's technology is critical enough to warrant acquisition.	50%
Evaluate the risk of community fork if NVIDIA acquires Modular.	This will determine whether the value of a portable runtime can be preserved post-acquisition.	50%
Determine the timeline for the electrical I/O wall in post-Blackwell GPU architectures.	This will affect the urgency of an Ayar Labs acquisition.	40%

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
AI21 Labs	Israel-based AI startup focused on enterprise LLMs and agentic reasoning, currently in advanced acquisition talks with NVIDIA.
Ampere Computing	U.S.-based developer of ARM server CPUs, deployed in Oracle Cloud Infrastructure and Microsoft Azure, considered as a potential acquisition target for NVIDIA.
Ayar Labs	U.S.-based startup specializing in co-packaged silicon photonics to solve the electrical I/O wall in AI systems.
CoreWeave	U.S.-based GPU cloud provider in which NVIDIA holds a strategic stake, serving as a preferred inference cloud partner.
Groq	U.S.-based AI hardware company acquired by NVIDIA for $20 billion in May 2026, specializing in low-latency inference via its LPU architecture.
HHI	Herfindahl-Hirschman Index, a measure of market concentration calculated as the sum of squared market shares, used to assess antitrust risk.
Lumentum	U.S.-based optical components company falsely claimed to have received a $2 billion investment from NVIDIA in March 2026.
Rent's Rule	A semiconductor scaling law predicting that required off-chip I/O grows sub-linearly with compute transistor count, creating an I/O gap.
Riva	NVIDIA's native speech AI stack providing ASR/TTS/NLP capabilities, integrated with DRIVE and Jetson platforms.
SambaNova Systems	U.S.-based AI hardware company specializing in full-stack AI systems with custom chips for enterprise inference workloads.
SoundHound AI	U.S.-based voice AI company processing 100B+ annual queries, previously considered as a potential acquisition target for NVIDIA.
Tenstorrent	Canadian-based AI chip startup led by Jim Keller, focusing on RISC-V based processors for efficient AI training and inference.
VAST Data	U.S.-based storage company providing object storage solutions, partnered with NVIDIA for AI workloads.
WEKA	U.S.-based data platform company with significant R&D in Israel, providing a GPU-optimized parallel filesystem for AI clusters.

How The Conclave Works

The Tribunal is a structured multi-model deliberation system. Instead of asking one AI model a question and trusting its answer, the Tribunal 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.