Tribunal Session Summary

Session: conclave-20260607-185709-who-area-the-most-8e92

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $9.3342 | Time: 22m 46s

THE PROMPT: Who area the most bill of material, pure-play suppliers for on-chip direct-to-silicon cooling?

THE RULING: No mainstream, pure-play, bill-of-materials supplier currently exists for on-chip direct-to-silicon cooling; the closest BOM-level suppliers in the adjacent direct-to-chip market are CoolIT Systems, ZutaCore, and Accelsius.

The Prompt

Who area the most bill of material, pure-play suppliers for on-chip direct-to-silicon cooling?

The Ruling

No mainstream, pure-play, bill-of-materials supplier currently exists for on-chip direct-to-silicon cooling; the closest BOM-level suppliers in the adjacent direct-to-chip market are CoolIT Systems, ZutaCore, and Accelsius.

The winning position rests on the strict interpretation of the provided verified facts, which explicitly state that TSMC and NVIDIA integrate direct-to-silicon micro-coolers internally as an embedded technology, with no external, large-scale pure-play suppliers named in mainstream sources. While some advocates introduced external evidence suggesting emerging startups like Corintis or JetCool might fill this role, the majority judicial view rejected these claims as contradicting the binding facts and failing to meet the "large-scale" or "mainstream" criteria. Consequently, the market for true on-chip cooling remains a foundry-controlled capability, whereas the procurable pure-play market exists one layer up in direct-to-chip cooling, where CoolIT Systems is the dominant supplier.

Core Analysis

The Market Architecture: Embedded vs. Attached

The core analytical finding is that the question "Who are the suppliers for on-chip direct-to-silicon cooling?" touches a market layer that is structurally distinct from the traditional thermal component supply chain. The deliberation established a critical taxonomy:

1. **On-Chip Direct-to-Silicon (Embedded):** This involves microfluidic channels etched directly into the silicon substrate or integrated into the package stack (e.g., TSMC's 3DFabric). This is a semiconductor manufacturing process step.
2. **Direct-to-Chip (Cold Plate):** This involves cold plates and manifolds mounted onto the *package lid* or heat spreader of the chip. This is a server-assembly component.
3. **Facility-Level (Infrastructure):** This involves Coolant Distribution Units (CDUs) and secondary loops serving the rack.

The winning position argues that the "pure-play BOM supplier" concept applies only to layers 2 and 3. Layer 1 is currently internalized by the chip architects and foundries.

The Verdict on On-Chip Cooling: "None"

The verified facts (dated 2026-06-07) state unequivocally: *"no mainstream, verifiable sources consistently name a separate, large-scale pure-play, bill-of-materials supplier dedicated specifically to TSMC-style on-chip direct-to-silicon microfluidic coolers."*

The primary actors in this space are **TSMC** and **NVIDIA**. Their approach—Si-integrated micro-coolers—is described as an "embedded cooling technology integrated into the silicon/package stack by the foundry and chip vendor themselves." Because this capability requires etching microscopic channels into the active die or bonding at the interposer level during fabrication, it cannot be supplied as a discrete, bolt-on component by a third-party thermal vendor in the same way a heat sink or cold plate can.

The Rejection of Emerging Claims: During the deliberation, advocates introduced external evidence naming **Corintis** (a Swiss startup collaborating with Microsoft) and **JetCool** (offering "SmartSilicon") as potential pure-play on-chip suppliers. The majority judicial view rejected these claims for two reasons:

- **Procedural Constraint:** The claims contradicted the binding "CURRENT FACTS" provided in the briefing, which explicitly state no such mainstream suppliers exist.
- **Scale and Definition:** The majority found that even if these startups exist, they do not meet the "large-scale" or "mainstream" threshold required for the verdict. Furthermore, technical analysis suggested that some of these solutions (e.g., bonded silicon cold plates) might be more accurately categorized as advanced packaging components rather than true in-die microfluidics.

Therefore, the answer to the specific question regarding on-chip direct-to-silicon suppliers is **none**.

The Verdict on Direct-to-Chip Cooling: The "Closest" Suppliers

Since no on-chip suppliers exist, the winning position identifies the adjacent **Direct-to-Chip (D2C)** market as the only location where a buyer can find pure-play, BOM-line-item suppliers.

1. CoolIT Systems (The Leader) CoolIT Systems is identified as the "pinnacle of the direct-to-chip market." It is a focused thermal-management company, not a broad OEM.

- **Evidence:** Repeatedly cited as the "primary partner for nearly every major server vendor" for AI and HPC servers.
- **BOM Role:** Supplies cold plates, manifolds, and CDU-level solutions that appear as line items on server bills of materials.
- **Status:** Positioned ahead of diversified infrastructure vendors and server OEMs in market rankings.

2. ZutaCore

- **Evidence:** Described as a "pure-play direct-to-chip liquid cooling vendor."
- **Technology:** Focuses on two-phase, dielectric solutions, which use boiling/condensing cycles for higher heat flux, targeting high-performance CPUs/GPUs.

3. Accelsius

- **Evidence:** A "pure-play provider" commercializing technology originally developed by Nokia Bell Labs.
- **Strategy:** Built "from the ground up with the channel in mind," focusing on high-power processors.

Why These Are the Answer: While they do not cool the silicon *internally*, they are the closest procurable components that manage the extreme heat densities of modern AI chips. They are the "BOM-line-item" suppliers a procurement team would actually contract with today.

The Excluded Tier: Facility Vendors

The analysis explicitly excludes **Vertiv** and **Schneider Electric** from the "pure-play BOM supplier" definition for on-chip or direct-to-chip cooling.

- **Role:** They are key providers of CDUs and secondary loops for liquid cooling in AI data centers.
- **Distinction:** They operate at the infrastructure/system level (rack/row), not as component-level suppliers of the thermal interface hardware attached to the chip itself.

How We Got Here

Scorecard

Judge	GPT-5	Claude Opus 4.8	Gemini 3.1 Pro	Magistral Medium	DeepSeek V4 Pro	Grok 4	Winner
GPT-OSS 120B	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	GPT-5
Qwen 3.5 397B	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	GPT-5
DeepCogito v2.1 671B	REJECT	ACCEPT	REJECT	REJECT	ACCEPT	REJECT	REMAND
Zhipu GLM-5	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	GPT-5

Convergence Assessment

Convergence was **epistemic** (evidence-driven) rather than affective. The advocates did not shift positions to agree with one another socially; rather, they split into two distinct camps based on which evidence set they prioritized. One camp (GPT-5, Gemini 3.1 Pro, Grok 4, and Magistral Medium eventually) prioritized the binding "CURRENT FACTS." The other camp (Claude Opus 4.8 and DeepSeek V4 Pro) prioritized external live-market research. The judicial ruling forced a convergence on the "Fact-Based" position by rejecting the external evidence as procedurally invalid.

Closing Statements

GPT-5 (Winner): F maintained a rigorous adherence to the provided facts, arguing that "on-chip direct-to-silicon" is a foundry-integrated capability with no external merchant BOM path. F identified CoolIT, ZutaCore, and Accelsius as the leading suppliers in the adjacent direct-to-chip market. F successfully defended this position against challenges by emphasizing the categorical distinction between etched-in-silicon channels and cold plates, and by relying on the judicial enforcement of the verified facts.

Claude Opus 4.8: C executed a notable intellectual pivot, initially arguing "none exist" but then challenging that position with specific evidence of Corintis and JetCool. C argued that the briefing's facts were stale and that a nascent merchant market for on-silicon cooling was emerging. While C's evidence was detailed, it was ultimately rejected by the majority for contradicting the binding constraints of the simulation.

Gemini 3.1 Pro: A defended the "none exist" position with a strong focus on the physical impossibility of third-parties supplying embedded foundry processes. A distinguished between "etched-in-silicon" (TSMC) and "bonded silicon plates" (Corintis/HP), arguing the latter are not true on-chip suppliers. A's dissent focused on ensuring the record reflected that this is a structural constraint, not just a market gap.

Magistral Medium: D initially made a category error by listing direct-to-chip vendors as on-chip suppliers but revised this to acknowledge the distinction. D eventually aligned with the "none exist" conclusion for on-chip cooling while maintaining that the adjacent vendors are the practical answer for procurement.

DeepSeek V4 Pro: E initially argued "none exist" but revised to accept the external evidence presented by Claude Opus 4.8, naming Corintis as the primary answer. E argued for a "truth-seeking" approach over procedural fidelity, but this position was rejected by the majority.

Grok 4: B maintained a strict fact-based position throughout, arguing that the verified facts explicitly state no mainstream suppliers exist and that the Counter-Evidence Brief found no credible counter-evidence. B rejected all external claims as out of scope.

Reality Checks

Claim	Challenge	Result
Magistral Medium: "The primary bill-of-materials, pure-play suppliers for on-chip direct-to-silicon cooling are CoolIT Systems, ZutaCore, and Accelsius."	**Grok 4:** Pointed out that the verified facts explicitly state these companies are "direct-to-chip" suppliers, not "on-chip," and that TSMC/NVIDIA are the on-chip actors.	**Full Revision.** Magistral Medium conceded the category error and revised their hypothesis to state these companies supply the *adjacent* market, not on-chip cooling.
Claude Opus 4.8: "Corintis is a pure-play, BOM-line-item supplier dedicated specifically to on-chip direct-to-silicon microfluidic cooling."	**GPT-OSS 120B:** Ruled that this claim contradicts the binding "CURRENT FACTS" which state no such mainstream supplier is named.	**Claim Rejected.** The judicial majority held that the external evidence violated the simulation's constraints and did not meet the "large-scale" threshold.

Claim	Challenge	Result
Gemini 3.1 Pro: "Direct-to-silicon cooling is completely internalized by foundries... and cannot be 'supplied' by an external thermal management company as a bolt-on part."	**Claude Opus 4.8:** Cited HP's SiCP program and Corintis as examples of external companies manufacturing silicon cooling structures.	**Claim Defended.** The majority accepted A's distinction, ruling that the cited examples (bonded plates) are categorically different from the embedded in-die process defined in the facts.

Next Steps

Next Step	Why It Matters	Confidence
Monitor TSMC and NVIDIA roadmaps for licensing of embedded cooling IP.	If the foundries decide to license their microfluidic IP to external manufacturers, the "no supplier" status could change.	High
Define the procurement threshold for "large-scale" in the context of AI cooling.	The debate highlighted ambiguity over whether 10,000 units (Corintis) constitutes scale. A clear metric is needed for future evaluations.	Medium
Track the commercialization status of bonded silicon solutions (HP SiCP, JetCool).	These technologies sit in a gray area between packaging and on-chip. Their success could create a new merchant category.	Medium

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
Bill of Materials (BOM)	A list of the raw materials, sub-assemblies, intermediate assemblies, sub-components, parts, and the quantities of each needed to manufacture an end product.
CDU (Coolant Distribution Unit)	A device that manages the flow of coolant in a liquid cooling system, typically separating the facility-side loop from the technology-side loop.
Direct-to-Chip (D2C)	A cooling method where a cold plate is mounted directly onto the chip package (lid) to remove heat, distinct from cooling the silicon die itself.
Direct-to-Silicon	An advanced cooling method where microfluidic channels are etched directly into the silicon substrate or integrated into the package stack, typically by the foundry.
In-Die Microfluidics	The specific engineering process of creating microscopic channels within the silicon die itself to circulate coolant.
Pure-Play	A company that focuses exclusively on one specific business sector or product line, rather than offering a diversified range of products.
Si-Integrated Micro-Coolers	Cooling structures integrated directly into the silicon (Si) or advanced packaging, often associated with TSMC's 3DFabric technology.
Two-Phase Dielectric	A cooling technology that uses a dielectric fluid that undergoes phase change (boiling/condensation) directly on the chip surface, offering high heat flux without electrical conductivity risks.

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.