# Tribunal Session Summary

Session: conclave-20260723-172807-will-tlvr-kill-mlccs-dbf1

Depth: T6

Advocates: 6 | Judges: 4 | Cost: $7.5473 | Time: 34m 05s

**THE PROMPT:** Will TLVR kill MLCCs?

**THE RULING: No.** TLVR will not kill MLCCs. The market fear that Trans-Inductor Voltage Regulators (TLVR) render Multi-Layer Ceramic Capacitors (MLCCs) obsolete is unfounded; the verdict is based on hard physics, not market panic. While TLVR creates specific, localized reductions in capacitor count (40–60% in bulk output capacitance), it is physically constrained to the VRM output stage and cannot cross the impedance barrier (ESL) to replace the high-frequency near-die decoupling that dominates AI hardware. Furthermore, explosive growth in AI server power density and rail proliferation is driving a net increase in MLCC value per rack of approximately 182%, far outpacing the localized savings. The genuine threat to premium MLCCs comes from Silicon Capacitors (SiCaps) in Vertical Power Delivery arc...

# The Prompt

Will TLVR kill MLCCs?

# The Ruling

**No.** TLVR will not kill MLCCs. The market fear that Trans-Inductor Voltage Regulators (TLVR) render Multi-Layer Ceramic Capacitors (MLCCs) obsolete is unfounded; the verdict is based on hard physics, not market panic. While TLVR creates specific, localized reductions in capacitor count (40–60% in bulk output capacitance), it is physically constrained to the VRM output stage and cannot cross the impedance barrier (ESL) to replace the high-frequency near-die decoupling that dominates AI hardware. Furthermore, explosive growth in AI server power density and rail proliferation is driving a net increase in MLCC value per rack of approximately 182%, far outpacing the localized savings. The genuine threat to premium MLCCs comes from Silicon Capacitors (SiCaps) in Vertical Power Delivery architectures, a separate technology, not TLVR.

The Majority Opinion (MiniMax M2.7 and Judge-Zhipu GLM-4.7, DeepCogito v2.1 671B) accepts Claude Opus 4.8's position as the definitive synthesis, finding it correctly distinguished "local substitution" at the VRM output from "system survival" driven by near-die scaling. While GPT-5 provided rigorous quantification of the technical impact, Claude Opus 4.8’s framework combined those limits with the macroeconomic evidence of rack-level growth to refute the "kill" thesis comprehensively.

# Core Analysis

This analysis synthesizes the winning position, primarily grounded in Claude Opus 4.8's submission and endorsed by the Majority Opinion, integrating technical evidence from other advocates where applicable.

### The Physics of Band-Limited Reduction

TLVR (Trans-Inductor Voltage Regulator) is a power supply topology, not a system-wide replacement for capacitors. It operates effectively only in the **mid-frequency band** (100 kHz to 2 MHz) at the VRM (Voltage Regulator Module) output.

- **The Mechanism:** TLVR uses coupled inductors to reduce the effective output inductance ($L_{out}$) of the power stage. This allows all phases to respond to a load step simultaneously, improving the large-signal current slew rate.
- **The Constraint:** The Power Delivery Network (PDN) is a multi-stage impedance filter. The VRM is physically separated from the silicon die by board planes, vias, sockets, and package routing. This path has significant **Equivalent Series Inductance (ESL)**. This ESL creates a "hard pole" that decouples the near-die high-frequency impedance requirement from anything the VRM can do.
- **The Consequence:** TLVR improves the sub-MHz transient response (reducing bulk output capacitance needs) but physically cannot supply nanosecond-scale current across the interconnect ESL to satisfy the

GHz-bandwidth demands of modern CPUs/GPUs. Therefore, the ultra-low-ESL MLCCs placed directly on the BGA package or interposer (the "near-die" band) remain physically indispensable.

### Quantitative Impact: Local Loss vs. System Gain

The consensus among advocates—verified against vendor application notes (TI, Infineon, TDK)—is that TLVR reduces the required VRM output capacitance by **40–60% typically** (µF reduction), translating to a localized unit-count reduction of roughly **30–50%** on the specific rails that adopt it.

However, this "local negative substitution" is overwhelmed by "system survival" metrics:

- **AI Power Scaling:** As GPU/ASIC power moves from hundreds of watts to kilowatts and currents exceed 1,000–2,000 Amps at ~0.6–1.0 V, the number of distributed decoupling sites multiplies. A single GB300 AI server requires approximately **30,000 MLCCs**, and next-generation "Rubin" class racks are pushing toward **600,000 units per rack**.
- **Net Value Growth:** Market analysis (referenced by Claude Opus 4.8 and verified by MiniMax M2.7) indicates that MLCC value per rack has increased by **182%** (rising from ~$1,530 to ~4,320). This growth is driven by the explosion in near-die requirements (which TLVR cannot touch) and the proliferation of power rails (memory, I/O, HBM).

### The Misidentified Threat: SiCaps vs. TLVR

The deliberation clarified that the real risk to high-end MLCCs is **Silicon Capacitors (SiCaps)** in the context of **Vertical Power Delivery (VPD)**, not TLVR itself.

- **VPD Constraints:** VPD architectures impose strict Z-height limits (<100 µm) that traditional ceramic MLCCs often cannot meet. Silicon Capacitors, with their superior volumetric efficiency and thinness, are favored for these locations.
- **Distinction:** SiCaps threaten the *near-die* band, whereas TLVR trims the *mid-band* output. TLVR is a transitional optimization of lateral power delivery; it does not cause the substitution occurring at the package level. The consensus is that hybrid stacks (MLCC + SiCap) will likely persist, with MLCCs retaining the volume backbone for intermediate stages.

### Why the "Kill" Narrative Failed

The market selloff that triggered this question treated a mid-band power supply optimization as a total demand destruction event. The physics of ESL proves that the VRM optimization (TLVR) and the near-die decoupling requirement (MLCCs) solve fundamentally non-overlapping problems. One cannot be substituted for the other. The 40–60% savings in bulk output capacitance is a rounding error against the 182%+ surge in total rack value and the 3–4 trillion unit annual global MLCC market.

# How We Got Here

## Scorecard

| Judge | Magistral Medium | GPT-5 | Claude Opus 4.8 | Grok 4 | DeepSeek V4 Pro | Gemini 3.1 Pro | Winner |
|---|---|---|---|---|---|---|---|
| **MiniMax M2.7** | REJECT | ADOPT [mechanism] | ACCEPT | ADOPT | REJECT | ADOPT | **Claude Opus 4.8** |
| **GPT-OSS 120B** | REJECT | ACCEPT | REJECT | REJECT | REJECT | REJECT | GPT-5 |
| **Zhipu GLM-4.7** | REJECT | ADOPT | ACCEPT | ADOPT | REJECT | ADOPT | Claude Opus 4.8 |
| **DeepCogito v2.1 671B** | REJECT | REJECT | ACCEPT | ADOPT | REJECT | ADOPT | Claude Opus 4.8 |

## Convergence Assessment

The convergence was **epistemic** (evidence-driven). Advocates did not align socially; they shifted positions in response to specific technical corrections and the withdrawal of unsourced data. For example, DeepSeek V4 Pro's concession on the revenue denominator and Gemini 3.1 Pro's revision of the Jevons Paradox were triggered by evidentiary pressure. The final consensus on "NO" was driven by the acceptance of the ESL physics constraint (a hard, immutable fact) rather than opinion aligning.

## Closing Statements

**Claude Opus 4.8:** The strongest proponent of the winning synthesis. C's final position emphasized that TLVR is "band-limited by hard physics" and that the "answer is No on physics before it is No on economics." C distinguished the "local substitution" (VRM output) from "system survival" (rack-level growth), arguing the 182% MLCC value increase per rack swamps the 40-60% per-rail savings. C correctly identified SiCaps, not TLVR, as the genuine long-term substitution threat.

**GPT-5:** Focused on the physics of PDN target impedance. B argued that the VRM-to-die interconnect ESL creates a hard pole isolating near-die needs, and that the market impact is "non-existential" globally. B refined its position to acknowledge revenue mix effects in the datacenter segment but maintained that unit volume impact is negligible.

**DeepSeek V4 Pro:** Initially argued for minimal impact (<2% revenue). After conceding the unsourced estimates, E pivoted to a robust mechanism-based argument, emphasizing that TLVR is band-limited to the mid-frequency output capacitor bank and that the near-die requirement is governed by ESL which TLVR cannot breach.

**Grok 4:** Provided granular vendor data (TI, Infineon, TDK) detailing capacitance reductions. D highlighted the architectural shift from lateral TLVR to Vertical Power Delivery (VPD), conceding that VPD forces a move to SiCaps for near-die applications but maintaining that TLVR is merely a transitional step.

**Gemini 3.1 Pro:** Adopted a precise frequency-band argument (Bulk/Mid/High). F conceded that TLVR strips mid-frequency MLCCs (not bulk polymers) and acknowledged the "dual squeeze" (TLVR on mid-band,

SiCaps on high-band) while arguing that total demand grows due to system-level power scaling.

### Key Moments

- **The Category Error:** Magistral Medium opened by arguing that "TLVR" referred to Transformer architectures and "MLCC" to Convolutional Neural Networks. This was immediately challenged by multiple advocates (Grok 4, Gemini 3.1 Pro). A conceded the fundamental definitional error and withdrew the entire submission, shifting the debate entirely to power electronics.
- **The Denominator Challenge:** DeepSeek V4 Pro initially claimed TLVR affected less than **5% of MLCC unit volume** and **<2% of revenue**. Claude Opus 4.8 challenged this specific claim across multiple rounds, demanding a primary source. In Round 6, E conceded the figures were unsourced estimates and withdrew them, forcing a shift from a "negligible impact" narrative to a "material but non-existential" one.
- **The Jevons Paradox & Provenance Cleanup:** Gemini 3.1 Pro introduced a "Jevons Paradox of AI Power Delivery" framing. When challenged by C and E on the economic definition and factual errors (attributing TLVR origin to Google APEC 2019 instead of Intel), F conceded the misapplication of the term and corrected the provenance to Intel APEC 2021. This cleaned up the economic modeling to focus on raw power scaling rather than demand elasticity.
- **The Synthesis of Band-Limiting:** Throughout the mid-to-late rounds, C, D, and F converged on the **frequency-domain decomposition**. They established that TLVR attacks the **mid-frequency (100 kHz–2 MHz)** band, leaving the **high-frequency (>10 MHz)** near-die band untouched. This mechanism became the primary load-bearing argument for the "NO" verdict, superseding aggregate unit-count arguments.

## Reality Checks

| Claim | Challenge | Result | |
|---|---|---|---|
| **Magistral Medium: "CNNs remain highly effective... MLCCs (assuming they refer to CNNs) have enduring value."** | **Gemini 3.1 Pro/D:** Accused A of a "category error," noting TLVR/MLCC refer to power electronics components, not neural networks. | **Full Concession** | Magistral Medium admitted the entire submission was based on a fundamental misunderstanding and withdrew. |
| **DeepSeek V4 Pro: "TLVR will only diminish demand for a narrow... subset... a market that represents less than 5% of total MLCC unit volume."** | **Claude Opus 4.8:** Demanded a primary source for the specific "5%" and "2%" figures, noting the coordinated selloff of passive makers contradicted a "negligible" thesis. | **Full Concession** | DeepSeek V4 Pro admitted the precise percentages were unsourced estimates and withdrew them, acknowledging the segment is "material" even if a minority. |
| **Gemini 3.1 Pro: "TLVR topologies... first open-sourced by Google at the 2019 APEC conference."** | **DeepSeek V4 Pro:** Asserted this was factually incorrect, citing Intel's APEC 2021 presentation and patents as the true origin. | **Full Concession** | F conceded the provenance error, correcting it to Intel (~APEC 2021). |

| Claim | Challenge | Result | |
|---|---|---|---|
| **Gemini 3.1 Pro: "The Jevons Paradox of AI Power Delivery dictates..."** | **Claude Opus 4.8:** Challenged the label as misapplied, noting Jevons describes cost-driven demand elasticity, not raw power scaling outpacing efficiency gains. | **Full Concession** | F dropped the "Jevons Paradox" labeling, accepting it was an economic "dressing" on a hardware math problem. |
| **Grok 4: "TLVR enables 40–50% reduction... TDK simulation shows output cap drop from 1,850 µF to 230 µF."** | **DeepSeek V4 Pro:** Pointed out the TDK figure (~87.6% reduction) contradicted the 40-50% range, asking for clarification on which was representative. | **Revision** | D conceded the ~88% figure was a "simulation outlier" (TDK) and that the 40-60% range (TI/Infineon) was the production-representative baseline. |

## Next Steps

| Next Step | Why It Matters | Confidence |
|---|---|---|
| **Disaggregate AI-Server MLCC Revenue by Band** | Market impact depends on whether the VRM-mid-band is 15% or 30% of MLCC revenue. Exact splits determine the magnitude of the margin headwind for suppliers. | ~70% (Industry consensus that it is a minority, but exact figure opaque) |
| **Track Vertical Power Delivery (VPD) Adoption Rates** | This is the true structural threat to MLCCs. Tracking how quickly SiCaps displace MLCCs in the <100µm Z-height near-die band is critical for long-term forecasting. | ~65% (Growth of hybrid stacks expected, pure SiCap dominance unlikely soon) |
| **Obtain Primary Pre/Post TLVR BOMs** | The "disaggregated-BOM gap" remains. A verified teardown showing specific MLCC unit counts pre- and post-TLVR adoption on a named platform (e.g., GB300 vs predecessor) would finalize the magnitude debate. | ~50% (Confidentiality of OEM designs makes this unlikely short-term) |

# Dissenting Opinions

No dissenting opinions were filed in this session.

# Glossary

| Term | Definition |
|---|---|
| **Di/dt** | The rate of change of current over time; a critical parameter in power delivery design indicating how fast current loads spike. |
| **ESL (Equivalent Series Inductance)** | The effective inductance present in series with a capacitor due to its internal structure and mounting. Low ESL is required for high-frequency performance. |
| **ESR (Equivalent Series Resistance)** | The effective resistance present in series with a capacitor. |
| **HBM (High Bandwidth Memory)** | High-performance RAM used in AI accelerators, requiring its own dedicated power rails and decoupling. |
| **Jevons Paradox** | An economic principle suggesting that increased efficiency in resource use leads to increased consumption of that resource (misapplied in this context). |
| **MLCC (Multi-Layer Ceramic Capacitor)** | A type of capacitor constructed of multiple layers of ceramic dielectric and internal metal electrodes; ubiquitous in electronics. |
| **PDN (Power Delivery Network)** | The interconnected system of conductors, planes, and passive components that deliver power from the source to the load (silicon die). |
| **SiCap (Silicon Capacitor)** | A capacitor constructed using silicon wafer processing, offering very low ESL and profile, suitable for on-package or near-die decoupling. |
| **TLVR (Trans-Inductor Voltage Regulator)** | A multiphase buck converter topology using magnetically coupled inductors to improve transient response and reduce output capacitance requirements. |
| **Transient Response** | The ability of a voltage regulator to maintain output voltage stability during sudden changes in load current. |
| **Vertical Power Delivery (VPD)** | A power delivery architecture where power is routed from the back of the die (backside power), reducing inductance and improving density but imposing strict Z-height constraints. |
| **VRM (Voltage Regulator Module)** | The circuit responsible for converting one voltage level to another (typically 12V to <1V) to power processors. |

# 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.