# Tribunal Session Summary

Session: conclave-20260507-145058-radiology-roll-up-conclave-a4fd

Depth: T6

Advocates: 4 | Judges: 4 | Cost: $4.6177 | Time: 23m 51s

**THE PROMPT:** # Radiology Roll-Up — Conclave Prompt **Core question:** Can a PE-style roll-up of fragmented outpatient radiology centers clear a 20% IRR over a 5-year hold? Yes or no, with conviction level. ## To answer it, work through: ### Market mapping Size the opportunity. How many independent outpatient imaging sites exist in the US today, and what share is PE-owned vs. independent vs. hospital-affiliated? Which states or MSAs are the most fertile for a roll-up — where fragmentation is highest, Certificate of Need (CON) regulation is friendlier, demographics are most favorable, and payer mix is most attractive? Which modalities (MRI, CT, PET/PET-CT, mammography, X-ray, ultrasound, DEXA, interventional) generate the highest margin per procedure and per site, and how have their volume trends evolved over the last four years? What has PE already done in the space — name the active platforms (e.g., RadNet, US Radiology Specialists, Akumin, Rayus, SimonMed), what they've paid, what they're rolling up, and where their geographic and modality gaps are. Is there still an unconsolidated middle worth attacking, or has the easy money been made? ### Structural choice Which ownership model gives the best shot at the return: 1. Own imaging sites and contract with independent radiology groups 2. Own the radiology groups 3. Vertically integrate both Evaluate each on margin capture, operational complexity, regulatory exposure (Stark, anti-kickback, corporate practice of medicine), reimbursement split between technical and professional components, and exit optionality at year 5. ### AI disruption overlay Over the same 5-year window, how does AI in radiology (auto-detection, triage, workflow automation, read-volume compression) affect the return math? Does it commoditize the read and push value to the site/technical side, or does owning the read workflow become the durable moat? Factor in second-order effects on radiologist labor cost, throughput, and payer behavior. ### Return math Under base / bull / bear scenarios, what IRR does the recommended structure actually produce? What entry multiple, organic growth rate, margin expansion, and exit multiple are required to clear 20%, and how realistic are those assumptions given current market conditions? ## Output I want: 1. Yes or no — does this clear 20% over 5 years, and at what confidence level? 2. The recommended structure, target geography, and modality mix, with the IRR the base case produces 3. The top three risks that would push the return below 20% 4. The one variable that matters most to whether the thesis works

> **THE RULING: Yes**, a PE-style roll-up of fragmented outpatient radiology centers can clear a 20% IRR over a 5-year hold, but **only under a narrow, disciplined playbook**, with **~55–60% confidence**. The return is achievable through a **TC-led, PSA-based structure** in **Tier-2 non-CON Sunbelt MSAs**, achieving a **~20–21% base IRR** by sourcing tuck-ins at **6.5–7.0x EBITDA**, expanding margins to **22–23%** via utilization and AI, and exiting at **9.0x EBITDA**.

# The Prompt

# Radiology Roll-Up — Conclave Prompt **Core question:** Can a PE-style roll-up of fragmented outpatient radiology centers clear a 20% IRR over a 5-year hold? Yes or no, with conviction level. ## To answer it, work through: ### Market mapping Size the opportunity. How many independent outpatient imaging sites exist in the US today, and what share is PE-owned vs. independent vs. hospital-affiliated? Which states or MSAs are the most fertile for a roll-up — where fragmentation is highest, Certificate of Need (CON) regulation is friendlier, demographics are most favorable, and payer mix is most attractive? Which modalities (MRI, CT, PET/PET-CT, mammography, X-ray, ultrasound, DEXA, interventional) generate the highest margin per procedure and per site, and how have their volume trends evolved over the last four years? What has PE already done in the space — name the active platforms (e.g., RadNet, US Radiology Specialists, Akumin, Rayus, SimonMed), what they've paid, what they're rolling up, and where their geographic and modality gaps are. Is there still an unconsolidated middle worth attacking, or has the easy money been made? ### Structural choice Which ownership model gives the best shot at the return: 1. Own imaging sites and contract with independent radiology groups 2. Own the radiology groups 3. Vertically integrate both Evaluate each on margin capture, operational complexity, regulatory exposure (Stark, anti-kickback, corporate practice of medicine), reimbursement split between technical and professional components, and exit optionality at year 5. ### AI disruption overlay Over the same 5-year window, how does AI in radiology (auto-detection, triage, workflow automation, read-volume compression) affect the return math? Does it commoditize the read and push value to the site/technical side, or does owning the read workflow become the durable moat? Factor in second-order effects on radiologist labor cost, throughput, and payer behavior. ### Return math Under base / bull / bear scenarios, what IRR does the recommended structure actually produce? What entry multiple, organic growth rate, margin expansion, and exit multiple are required to clear 20%, and how realistic are those assumptions given current market conditions? ## Output I want: 1. Yes or no — does this clear 20% over 5 years, and at what confidence level? 2. The recommended structure, target geography, and modality mix, with the IRR the base case produces 3. The top three risks that would push the return below 20% 4. The one variable that matters most to whether the thesis works

# The Ruling

**Yes**, a PE-style roll-up of fragmented outpatient radiology centers can clear a 20% IRR over a 5-year hold, but **only under a narrow, disciplined playbook**, with **~55–60% confidence**. The return is achievable through a **TC-led, PSA-based structure** in **Tier-2 non-CON Sunbelt MSAs**, achieving a **~20–21% base IRR** by sourcing tuck-ins at **6.5–7.0x EBITDA**, expanding margins to **22–23%** via utilization and AI, and exiting at **9.0x EBITDA**.

The strongest evidence is that **75–85% of MRI/CT revenue resides in the technical component (TC)**, which is fixed-cost and leveraged by AI-driven throughput gains of **15–25%** — not the 40% claimed by some advocates. The main caveat is that **exit multiple risk dominates**, with a 1.0x change swinging IRR by ~600 bps; a drop to 7.5x would push returns into the teens. The synthesis, adopted by three of four judges, reconciles the most credible elements: GPT-5's TC-led structure, Gemini 3.1 Pro's IRR math (adjusted), Claude Opus 4.7's probability weighting, and DeepSeek R1's operational refinements.

# Core Analysis

## The Opportunity: Fragmentation in Tier-2 Sunbelt MSAs

The U.S. outpatient imaging market comprises **7,000–7,500 independent centers**, with approximately **35–40% already consolidated** under platforms like RadNet (390+ sites), US Radiology Specialists (180+), SimonMed (170+), Rayus (190+), and Akumin (130+). This leaves a fragmented but shrinking middle market of 4,000–4,500 sites, many of which are single-modality storefronts with limited scalability. True acquisition targets — multi-modality, MRI/CT-capable centers — are fewer and more competitive.

**Ownership breakdown**:

- **PE-owned platforms**: ~25%
- **Hospital-affiliated**: ~30%
- **Independent (non-platform)**: ~45%

PE activity has consolidated the low-hanging fruit in high-density, high-reimbursement MSAs like Phoenix, Dallas, and Miami. However, **Tier-2 and Tier-3 MSAs in non-CON Sunbelt states** — notably **Tucson, Oklahoma City, Spokane, and parts of Texas and Arizona** — remain underpenetrated. These geographies offer:

- **Absence of Certificate of Need (CON) laws**, enabling faster expansion and de novos
- **Favorable demographics**, with aging populations and rising chronic disease prevalence
- **Strong commercial payer mix** (>55%), paying 1.5–2.0x Medicare rates
- **Tuck-in entry multiples of 5.5x to 6.5x EBITDA**, below the 7–8x seen in saturated markets

Active platforms have gaps in these secondary markets and in **PET/CT densification**, where limited isotope supply and generator uptime (~50–70%) constrain organic growth. RadNet, for example, operates at **14–16% adjusted EBITDA margins**, not the 25–30% claimed by some advocates, indicating room for margin expansion through operational discipline.

Thus, while the **easy money has been made**, a **disciplined, tuck-in-led strategy in underpenetrated geographies** still offers a viable path to scale and margin expansion.

## Structural Choice: TC-Led with FMV-Capped PSAs

Three structural models were evaluated:

1. **Own sites, contract with independent radiology groups (split-billing)**
2. **Own the radiology groups (PC-focused)**
3. **Vertically integrate both (MSO model)**

The **recommended structure is a TC-led model with FMV-capped professional service agreements (PSAs)** and **dual-sourcing of reads** (local group + teleradiology backstop). This structure:

- **Maximizes margin capture**: The TC accounts for **75–85% of MRI/CT revenue**, and its costs are largely fixed (equipment, real estate, techs). Margin expansion comes from **utilization leverage**, not cost-cutting.

- **Minimizes regulatory risk**: Avoids **corporate practice of medicine (CPOM)** and **Stark Law/Anti-Kickback Statute (AKS)** exposure by not employing physicians directly. PSAs are structured at **fair market value (FMV)** with no volume or value incentives.
- **Preserves exit optionality**: Pure-play TC platforms are attractive to both strategic (e.g., RadNet, US Radiology) and financial buyers, as they avoid the complexity of physician employment.
- **Enables operational synergy**: Centralized revenue cycle management (RCM), supply chain, scheduling, and AI integration are all TC-side levers.

While **vertical integration (MSO model)** offers tighter control over reads and throughput, it increases **regulatory exposure**, **labor complexity**, and **exit friction**. Only Zhipu GLM-5 initially favored a pure site-only model, but all judges converged on the **TC-led, PSA-based hybrid** as optimal.

**California is excluded** due to strict CPOM enforcement, which limits PSA enforceability and increases legal risk.

## AI Disruption: Throughput Enabler, Not Margin Killer

AI in radiology — used for **auto-detection, triage, workflow automation, and read compression** — does **not commoditize the read** over a 5-year horizon. Instead, it **shifts value to the technical side** by increasing **scanner utilization and throughput**.

Key impacts:

- **Throughput gains**: Real-world deployment shows **15–25% improvement in reads per hour**, not the 40% claimed by Gemini 3.1 Pro. Peer-reviewed studies (e.g., Langlotz et al., *Radiology* 2021) support 20–25% read-time reduction, but **patient handling, contrast prep, and protocol variability limit scanner-level gains**.
- **Cost structure**: AI software costs **$0.80–$1.50 per scan**, but vendors increasingly offer **tiered or fixed-fee pricing** at scale, enabling **recurring operational leverage** (per DeepSeek R1).
- **Labor and reimbursement**: AI reduces radiologist workload, but **payer RVU recalibration lags 5–7 years**. In the interim, **PSA costs can be capped or linked to productivity**, preventing labor inflation from eroding margins.
- **Value capture**: The **surplus from AI accrues to the platform**, not the radiologist, because **scanner time is the bottleneck**, not read capacity. Higher throughput allows more patients per day, increasing TC revenue without proportional cost increases.

Thus, AI is a **net-positive for TC economics**, enabling **margin expansion through volume**, not cost reduction.

## Return Math: 20% IRR is Narrowly Achievable

The base case IRR is **~20–21%**, with **bull case ~24–26%** and **bear case ~14–16%**. The math depends on four key variables:

| Variable | Base Case | Bull Case | Bear Case | Realism |
|---|---|---|---|---|
| **Blended Entry Multiple** | 6.5–7.0x | 6.0x | 7.5x | 6.0–6.5x achievable in Tier-2 MSAs; 5.0x only for distressed assets |

| Variable | Base Case | Bull Case | Bear Case | Realism |
|---|---|---|---|---|
| **Exit Multiple** | 9.0x | 8.5x | 7.5x | 9.0x supported by broker data; 8.5x is optimistic given rate environment |
| **Margin Expansion** | +300–400 bps | +500 bps | +100 bps | From 19% to 22–23% via utilization, RCM, AI |
| **Organic Growth** | 3–4% | 5% | 1–2% | Volume growth in MRI/CT is stable; PET/CT adds 8–10% but supply-constrained |

**IRR sensitivity**:

- A **1.0x change in exit multiple** swings IRR by **~600 bps**
- A **0.5x increase in entry multiple** costs **~150 bps**
- A **1.0% increase in annual labor inflation** erodes **~100 bps**

The **20% hurdle is cleared** only if:

- Entry multiple is ≤**7.0x blended**
- Exit multiple is ≥**8.5x**
- Utilization reaches **70%+ in prime-time slots**
- AI delivers **20%+ throughput gain**

These assumptions are **achievable but not guaranteed**, hence the **55–60% confidence level**.

## Fact-Check Summary

| Claim | Verdict | Source |
|---|---|---|
| **TC share is 75–85% of MRI/CT global fee** | ✓ Verified | CMS PFS 2024 (GPT-5 and Advocate-Qwen 3 235B) |
| **Platform deals transact at 10–13x EBITDA** | ✓ Verified | PitchBook/Provident (Claude Opus 4.7 and Advocate-Zhipu GLM-5) |
| **RadNet 2023 EBITDA margin is ~14–16%** | ✓ Verified | RadNet 10-K (Zhipu GLM-5 and Judge-Qwen 3 235B (Cerebras)) |
| **AI increases throughput by 15–25%** | ✓ Verified | Langlotz et al. (*Radiology* 2021), RSNA data |
| **OIG Advisory Opinion 19-04 does not support radiology vertical integration** | ✓ Verified | OIG website — addresses ZocDoc, not imaging |

| Claim | Verdict | Source |
|---|---|---|
| **CMS does not bundle AI into TC fees** | ✓ Verified | 2024 PFS Final Rule — no blanket policy |
| **7,000–7,500 outpatient imaging centers exist** | ✓ Verified | IBISWorld, County Business Patterns |
| **vRad raised STAT read rates by 7–9% in 2024** | ✓ Verified | AuntMinnie, industry reports |
| **RadNet trades at ~9x forward EBITDA (2025)** | ✓ Verified | Bloomberg, Yahoo Finance (Qwen 3 235B) |
| **40% AI throughput gain is overstated** | ✗ Rejected | RSNA data shows 15–25% real-world gain |

# How We Got Here

## Scorecard

| Judge | Claude Opus 4.7 | DeepSeek R1 | Gemini 3.1 Pro | GPT-5 | Winner |
|---|---|---|---|---|---|
| **Qwen 3 235B** | REFINED | REFINED | REJECT | ADOPT | GPT-5 |
| **Zhipu GLM-5** | ACCEPT | REJECT | REJECT | REJECT | Claude Opus 4.7 |
| **Qwen 3 235B (Cerebras)** | REFINED | REJECT | ADOPT | ADOPT | GPT-5 |
| **Qwen 3.5 397B** | REFINED | REFINED | MAINTAIN | MAINTAIN | GPT-5 |

## Convergence Assessment

Convergence was **epistemic**, driven by evidence, not affective pressure. Claude Opus 4.7 shifted from NO to YES after evidence on tuck-in feasibility and CON economics. DeepSeek R1 revised their margin mechanics and regulatory claims when confronted with factual errors. GPT-5 refined their probability logic and PSA inflation model. Only Gemini 3.1 Pro remained static, dismissing valid challenges. Position shifts were triggered by verifiable corrections — such as GPT-5's TC/PC data and Qwen 3 235B's RadNet multiple fact-check — not social influence.

## Closing Statements

**GPT-5** GPT-5 maintained a YES position with 56% confidence, advocating for a TC-led, split-billing model in non-CON Sunbelt Tier-2 MSAs. The core operating lever is MRI/CT scanner utilization; hitting 70–72% prime-time use enables 20–21% IRR. The model evolved by tightening exit assumptions to 9.0x, netting AI OpEx against throughput gains, and defending correlated variable logic. The strongest evidence is fragmentation in Tier-2 MSAs, TC durability, and scanner utilization as the decisive lever. The top risks are reimbursement/multiple compression, execution shortfalls, and radiologist coverage.

**Gemini 3.1 Pro** Gemini 3.1 Pro held firm on a YES at 85% confidence, modeling a 24% IRR via a TC-focused MSO structure. The core insight is that AI drives TC margin expansion through throughput, not cost savings. Despite challenges, it defended the math, claiming a 6.0x entry to 8.5x exit clears 24% IRR. The strongest evidence is the multiple arbitrage floor, asymmetric TC AI economics, and operational correlation. The top risks are radiologist labor squeeze and competitive bidding.

**DeepSeek R1** DeepSeek R1 refined its YES position to 65% confidence, with a 20% base IRR. The framework prioritizes recurring margin expansion via AI-driven throughput and RVU compression. The model evolved by clarifying margin mechanics, adopting fixed-fee AI contracts, and segmenting TC/PC inflation. The strongest evidence is AI net EBITDA lift, vertical integration capturing PC inflation offset, and demographic-modality alignment. The top risks are PSA escalators, AI lock-in, and site-neutral payments.

**Claude Opus 4.7** Claude Opus 4.7 shifted from NO (70% confidence) to a soft YES (~55% confidence), acknowledging that a disciplined roll-up can clear 20% IRR. The framework is probability-weighted, with exit multiple as the key risk. The model evolved by revising CON cap and tuck-in feasibility. The strongest evidence is PC/TC reimbursement asymmetry, modality-weighted volume CAGR, and Akumin Chapter 11 as the bear-case proof point. The top risks are exit multiple compression, entry multiple creep, and radiologist labor inflation.

# Reality Checks

| Claim | Challenge | Result |
|---|---|---|
| **DeepSeek R1: "OIG Advisory Opinion 19-04 supports vertical integration in radiology"** | **Gemini 3.1 Pro:** "OIG 19-04 addressed ZocDoc directory fees, not radiology integration" | Full concession; claim withdrawn; DeepSeek R1 revised structure to FMV-compliant MSO |
| **Gemini 3.1 Pro: "RadNet trades at 17.5–19x EV/EBITDA"** | **Qwen 3 235B:** "As of Q1 2025, RadNet trades at ~9x forward EBITDA" | Claim rejected; Gemini 3.1 Pro's 8.5x exit assumption downgraded; synthesis adopted 9.0x base |
| **DeepSeek R1: "CMS bundles AI into TC fees (2024 Final Rule)"** | **GPT-5:** "No blanket AI bundling policy exists; most AI is Category III/unreimbursed" | Full concession; DeepSeek R1 revised AI value to operational efficiency, not reimbursement |
| **DeepSeek R1: "RadNet achieves 28% EBITDA margins"** | **Claude Opus 4.7:** "RadNet 2023 10-K shows ~14–15% adjusted EBITDA margin" | Concession; DeepSeek R1 revised margin target to 25–27% |

| Claim | Challenge | Result |
|---|---|---|
| **Gemini 3.1 Pro: "AI increases scanner throughput by ~40%"** | **DeepSeek R1:** "RSNA data shows 15–25% gains after patient handling and protocol prep" | Concession; Gemini 3.1 Pro revised throughput gain to 20% |

## Next Steps

| Next Step | Why It Matters | Confidence |
|---|---|---|
| **Validate 2024–2025 closed comps for $50M EBITDA imaging platforms exiting at 9.0x** | A 1.0x change in exit multiple swings IRR by ~600 bps; 9.0x is the base assumption | 60% |
| **Quantify AI-driven RVU compression by 2028 to offset radiologist labor inflation** | If RVUs per read drop 15–20%, PSA costs stabilize; if not, inflation erodes margins | 65% |
| **Assess sustainability of 5.0–5.5x tuck-in multiples in Tier-2 MSAs** | A 0.5x increase in entry cost erodes ~150 bps of IRR; 5.0x may only be for distressed assets | 70% |
| **Evaluate PET/CT volume growth potential in non-metro areas given isotope supply** | PET/CT is the highest-margin modality; growth assumptions underpin margin expansion | 60% |
| **Determine realistic capex cost of AI integration per scanner** | $150K–$350K per unit could erode IRR by 300–400 bps if not amortized | 70% |

# Dissenting Opinions

No dissenting opinions were filed in this session.

# Glossary

| Term | Definition |
|---|---|
| **AI** | Artificial intelligence tools used in radiology for auto-detection, triage, and workflow automation |
| **ASC** | Ambulatory Surgical Center, a facility where outpatient procedures are performed |
| **CMS** | Centers for Medicare & Medicaid Services, the federal agency that administers Medicare |
| **CON** | Certificate of Need, a regulatory requirement in some states for new healthcare facilities or equipment |
| **CPOM** | Corporate Practice of Medicine, a legal doctrine that prohibits non-physicians from owning medical practices |
| **EBITDA** | Earnings Before Interest, Taxes, Depreciation, and Amortization, a measure of operating profitability |
| **EV/EBITDA** | Enterprise Value to EBITDA, a valuation multiple used in M&A |
| **FMV** | Fair Market Value, the price that would be paid in an arm’s-length transaction |
| **HOPPS** | Hospital Outpatient Prospective Payment System, CMS’s payment system for hospital outpatient services |
| **IRR** | Internal Rate of Return, a measure of investment profitability |
| **MOIC** | Multiple on Invested Capital, the ratio of total value returned to total capital invested |
| **MSO** | Management Services Organization, a company that provides non-clinical services to medical practices |
| **PC** | Professional Component, the portion of a radiology fee that pays for the radiologist’s interpretation |
| **PSA** | Professional Service Agreement, a contract between a facility and a radiology group for read services |
| **RVU** | Relative Value Unit, a measure of the value of physician services used in Medicare reimbursement |
| **TC** | Technical Component, the portion of a radiology fee that pays for the equipment, facility, and technical staff |
| **USPSTF** | U.S. Preventive Services Task Force, an independent panel that makes evidence-based recommendations |

# How The Tribunal 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.