INDIE AGENCY NEWS · BLUEPRINT

Indie Blueprint

AI LEGAL RISKS

What every agency needs to know right now — from generative experiments to agentic workflows, and the patchwork of laws already reshaping how agencies operate.

BLUEPRINT ROUNDTABLE · APRIL 2026 · PAAVANA KUMAR, DAVIS+GILBERT LLP INDIEHQ.IS

CONTENTS

What's inside.

The Opener

Most agency conversations about AI still center on generative tools — the drafts, the images, the ideation. Paavana Kumar, a partner at Davis+Gilbert LLP who has spent her entire career advising agencies on advertising and technology law, says that conversation is already out of date.

The real story of 2026 is **agentic AI** — systems that don't just suggest, but act, transact, and execute on behalf of an agency or its client. That shift, Kumar warns, changes the legal stakes profoundly. An agentic system booking travel, disseminating a press release, selecting influencers, or handling customer service isn't a copilot. It's an agent. And under long-standing law, when your agent goes rogue, you're still on the hook.

The good news: the analysis isn't actually that new. What's new is the velocity. Agencies that slow down just enough to build the right guardrails — contracts, training, audit trails — can move faster, with more confidence, than the ones racing to deploy without them.

> ***"You don't necessarily want an agentic AI tool that's like the best golden retriever. You want to train it in the appropriate way."***

PAAVANA KUMAR, DAVIS+GILBERT LLP

Five things you need to know

1. **Agentic AI is a legal category shift, not a feature upgrade.** A chatbot is not agentic. A system that independently decides, transacts, and publishes is — and that changes which laws apply to you.
2. **"The AI did it" is not a defense.** Agency principle applies to AI. If your tool accepted terms, signed a consumer up for a renewal, or published a misleading press release, the liability sits with the deploying company.
3. **Shadow AI is the biggest policy gap in indie agencies.** Employees using unauthorized tools — especially on client work or pitch decks — is creating confidentiality and IP exposures most agency leaders don't know exist.
4. **Testing environments are the most-skipped step in the AI lifecycle.** Kumar sees this "more times than I can count" — agencies have a policy but never actually stress-test the tool before letting it loose.
5. **The regulatory patchwork is getting weirder, not cleaner.** New York passed the first-of-its-kind Synthetic Performer Act. Hours later, the White House issued an executive order threatening to pull funding from states that over-regulate AI. The fight is galvanizing states, not slowing them.

THEME 01

Generative vs. agentic AI.

The setup

Most agency leaders still use "AI" as a single category. Kumar's opening move was to pull the two apart — because the legal framework for generative tools and the legal framework for agentic tools are not the same conversation.

What was said

Generative AI is a predictive system. Input in, output out. A chatbot, even one that feels conversational, is still operating on rule-based logic on top of a large language model. "It's really not agentic," she said. "It's still an input-and-output-based model."

Agentic AI is different in kind — "moving from that predictive copilot model to an independent agent. Almost think of it as a person — like one of your associates, or an employee that's being trained in judgment and with certain thresholds, but ultimately, if they are let loose to do whatever they want to do…" An agentic tool generates campaign proposals through autonomous research, adjusts pricing in real time, handles end-to-end transactions like travel bookings and subscriptions, and detects sentiment, picks channels, publishes content, and iterates — all without human intervention.

The critical legal frame: agentic AI doesn't create a new body of law. It **amplifies risk under existing laws**. FTC scrutiny around misleading endorsements, undisclosed chatbots, unsubstantiated claims, and affiliate relationships all still apply — just faster, at scale, and with less human visibility.

The tension. The room came in wanting a list of "new AI laws." Kumar pushed back: the risk isn't novel statutes — it's old statutes enforced against new speed. An agentic tool rotating banner ads and personalizing claims in real time isn't breaking new legal ground; it's making the existing FTC claim-substantiation regime much harder to monitor in real time.

> ***"You've got generative AI maybe generating a press release. You have agentic AI detecting sentiment, choosing channels, publishing, disseminating, iterating. All of that can be done without human intervention."***
>
> PAAVANA KUMAR

What you can use now

Before deploying any AI tool on client work, write a one-paragraph description of what it does — verbs, not nouns. If the verbs include "decides," "transacts," "publishes," or "executes," treat it as agentic and apply the full guardrail stack. If the verbs are "drafts," "suggests," or "generates," you're in generative territory with a lower ceiling on autonomous action.

THEME 02

The three-bucket guardrail framework.

The setup

Asked to name the first three guardrails an indie agency owner should put in place before deploying agentic AI for client work, Kumar gave the cleanest, most reusable framework of the session.

1. Contracts and policies

An MSA, addendum, or scope of work laying out exactly what the agency is authorized to do with AI on that client's business. Protections around ownership and IP infringement. Definitions distinguishing generative from agentic AI — because the reps, warranties, and indemnities should look different for each. Vendor agreements that bind freelancers and subcontractors to the same standards. And, often overlooked, understanding whether the agency is actually using the enterprise level of a tool, because most consumer-tier terms disclaim the provider's responsibility for output.

2. Training and oversight

Everyone gets trained — not just creative, but legal, marketing, HR, and IT. Kumar was direct: "What you don't want is what's called shadow AI. We have employees going on using it without understanding."

3. Audit and record-keeping

"This is another one I think a lot of agencies miss." Kumar recommended an internal AI usage log capturing what tools were used, the prompts, who reviewed the output, and when. If a regulator comes knocking — or a client or rights-holder alleges infringement — that log is the difference between a defensible position and a nightmare.

The tension. Indies' competitive advantage is speed. Kumar's three buckets are process. The honest read: she's asking indies to absorb just enough structural discipline that the speed doesn't become a liability. The agencies that hold both — velocity and rigor — win the next 18 months.

> ***"The AI is your agent, for lack of better words. The fact that it may have gone rogue for some reason — it's not really a Get Out of Jail Free card."***
>
> PAAVANA KUMAR

What you can use now

Pick the weakest of the three buckets in your shop and close the gap this quarter. For most indies, it's bucket three — no audit log exists at all. A shared spreadsheet with five columns (date, tool, use case, prompt summary, reviewer) is better than nothing and can live inside Airtable or Notion in an afternoon.

THEME 03

The regulatory patchwork.

The setup

Kumar was asked to walk through the state-level landscape — with California, Washington, New York, and Illinois flagged as the states agencies likely needed to watch most closely.

What was said

Synthetic performers and privacy are the issues most likely to produce a 50-state patchwork; agencies should concentrate attention on the states where their client or consumer base is largest.

Then she surfaced the tension nobody's talking about: New York passed the **Synthetic Performer Act** — the first law of its kind in the country, protecting against unauthorized digital replicas of real people. Hours later, the White House issued an executive order pushing back against state-level AI over-regulation, threatening to withhold federal funding from states whose AI rules the administration deems too onerous.

Kumar's read: the executive order isn't chilling state action — it's galvanizing it. Progressive states are likely to respond by declining to pass AI-specific statutes and instead bringing enforcement actions under existing consumer-protection law, which doesn't look like "AI rulemaking" but functions as exactly that.

Separately, the **California AI Transparency Act** takes effect this summer. Disclosure obligations land on agencies, brands, and the platforms themselves. And she flagged the layer most conversations skip: **class actions**. California's consumer-protection law has a private right of action, and plaintiffs' firms are already filing. The FTC is one risk vector. So are state attorneys general. And so is the California bar.

The tension. There's a live fight between federal deregulation and state-level enforcement, and agencies are stuck in the middle. You can't wait for clarity from Washington, because the states — and private plaintiffs — are moving faster than any federal consensus will.

> ***"The White House issued an executive order essentially saying, we don't want states to introduce laws that are too onerous. Doesn't appear to have had a chilling effect yet. If anything, it's galvanizing states to take matters into their own hands."***
>
> PAAVANA KUMAR

What you can use now

Identify the three states where your client base is most concentrated. Flag the AI-specific laws in each with help from counsel. For California-facing work, assume the AI Transparency Act applies this summer and plan disclosure workflows now — don't wait for the effective date.

THEME 04

IP and ownership in the LLM era.

The setup

A member question cut to a practical concern: if we pay for ChatGPT or Claude and use it to make a logo, do we own it? Can we hand it to a client?

What was said

Kumar broke ownership into two distinct legal questions agencies routinely conflate.

Question 1: Do you own the output? Generally, if you're on the enterprise or business tier of a major platform, yes — the terms grant you ownership. Consumer tiers are murkier.

Question 2: Is the output safe from IP claims? The harder question. Every major platform carves exceptions into its indemnity: if your prompt was infringing, the reps and warranties are void; if the training data itself was infringing, the protection evaporates; if you modify the output after receiving it, the platform's indemnity applies only to what it delivered, not the final version.

The practical consequence: you can simultaneously own a logo and be on the hook for an IP claim against it. And if the client modifies it further downstream, the chain of responsibility gets more tangled. Kumar also pointed to the copyright cases moving through the courts — use of AI output trained on copyrighted data and whether that qualifies as fair use. "Courts are leaning toward not," she said, noting circuit splits are emerging.

The tension. The "we paid for it, so we own it" mental model many agencies bring to AI output is structurally wrong. Ownership and indemnity are different legal rights, and the platforms have built their terms to keep those separate on purpose.

> ***"You can own something and have the rights to give it to a client, and still might have a potential IP claim, depending on how it was ideated."***
>
> PAAVANA KUMAR

What you can use now

Before delivering any AI-assisted creative, document two things: (1) confirm you were on the enterprise/business tier at the time of generation; (2) capture the prompt and note whether any post-generation modifications were made. If modifications were substantial, the platform indemnity effectively expires — price the risk in, or flag it to the client.

THEME 05

The testing discipline most agencies skip.

The setup

Throughout the session, Kumar returned to one recurring frustration: agencies with AI policies that exist on paper but are never actually tested against real use cases before deployment.

What was said

"What's missing the most when I talk to agencies and our clients about how they're integrating these tools is how contracted their testing environments are," Kumar said. "There are now a lot of tools and environments that can be built to help test these tools before they are deployed at scale — and that step gets skipped, honestly, more times than I can count."

She walked through the testing types worth considering: **adversarial attacks on perception** (can the AI be misled by an unfamiliar scenario?); **stress testing** (give it many instructions and see which it follows best); **simulations** (run it against real-world use cases at small scale first); and **reward-bypass testing** (AI trained on reward will sometimes look for loopholes to hit its stated goal).

The key is matching the testing protocol to the use case. The tests for an AI that selects influencers and auto-posts disclosures are different from the tests for an AI that drafts crisis press releases — which need far more thresholds and human checkpoints.

The tension. The industry's loudest narrative right now — "you're falling behind, ship faster" — pulls in the exact opposite direction of Kumar's advice. She's not arguing against speed. She's arguing that an untested tool that later causes a client reputational disaster will cost more time than a two-week closed-environment test cycle ever would.

> ***"These tools are good at being trained on reward. Look to see if it's checking for loopholes or bypassing certain compliance thresholds you've established to get to where it thinks it ought to go."***
>
> PAAVANA KUMAR

What you can use now

For any agentic tool you're planning to deploy, spend one week running it internally on a closed use case before putting it anywhere near a client. Document what it does well, where it drifts, and what thresholds it violates. That document becomes part of your audit record.

THEME 06

The practical on-ramp.

The setup

The session closed with the most human question of the conversation: what would you say to an agency owner watching this and feeling overwhelmed?

What was said

Kumar's answer was calming and specific. The analysis, she argued, isn't fundamentally new: "I've never advised an agency to go create infringing content, or to not disclose an affiliate relationship, or to use confidential data of the client in breach of the agreement. It was really not groundbreaking stuff. It was just the speed at which the tools were evolving was freaking people out."

Her on-ramp for a leader who feels behind: **(1) Go build something.** Open Claude or ChatGPT and walk through actually building and deploying an agent. The academic understanding is more intimidating than the practical one. **(2) Treat it like you'd train a junior person.** There's always liability in deploying employees; AI isn't different in kind, just in speed. **(3) Start small.** A basic AI policy, a few pre-approved use cases, clear thresholds on what requires human review. **(4) Let it evolve.** In two years this won't feel overwhelming — the same way generative AI stopped feeling overwhelming after the 2023–2024 panic.

The tension. The industry pressure is to have a complete AI point of view right now — AI services packaged and sold, agentic workflows on every deck. Kumar's counter: the agencies that over-index on announcement and under-index on internal understanding are the ones most likely to end up on the wrong side of an enforcement action.

> ***"Train it like you would train a junior person, and just really understand that — and then you can use them to operationalize efficiencies."***
>
> PAAVANA KUMAR

What you can use now

Block two hours this week to actually build something agentic. Not a demo, not a pitch — a real, small, internal workflow. The goal isn't output. The goal is understanding the failure modes before you sell the capability to a client.

SECTION 08

The through-line.

AI liability doesn't come from the technology. It comes from the gap between what the tool does and what the agency understands it's doing.

Every theme — agentic vs. generative, guardrails, regulatory patchwork, IP, testing, the on-ramp — circles back to the same structural problem. The tools are evolving faster than most agencies' internal comprehension of them. The liability sits in that gap. This reframes the entire AI conversation for indies: it's not "are we using enough AI?" It's "do we understand the AI we're using well enough to defend our use of it if asked?"

Recurring themes

1. **The agency principle applies.** Crisis press release, travel-booking agentic tool, logo IP question, shadow AI risk — Kumar kept returning to the same frame. If the tool acted on your behalf, you own the action.
2. **The right metaphor is a junior employee.** The overeager intern, the golden retriever, the junior associate, the agent sent into the world. The architecture for AI is the same architecture agencies already use for managing people — just faster and at higher volume.
3. **The risk is amplified, not new.** FTC claim substantiation, endorsement disclosure, IP infringement, privacy, confidentiality — all already on the books. Agentic AI doesn't create new exposures; it multiplies existing ones by the velocity of automation.

Contradictions worth naming

- **Speed vs. testing.** Indies compete on speed, but the biggest risk Kumar flagged is inadequate pre-deployment testing. The agencies that compress the testing cycle without eliminating it will have a structural advantage.
- **Federal deregulation vs. state enforcement.** Washington signals "less AI regulation"; California, New York, and state AGs signal "more enforcement under existing law." Agencies waiting for regulatory clarity from Washington are optimizing for the wrong signal.

What was missing

- **Employment law.** What happens when AI replaces a role? What are disclosure obligations to staff when agentic tools are introduced? Its own conversation.
- **Client procurement pressure.** Client-side procurement is now asking for AI disclosures and warranties agencies aren't yet equipped to provide — a contract-negotiation dynamic worth its own session.

SECTION 09

The toolkit.

These tools emerged directly from Kumar's session. They are scaffolds, not legal advice — adapt them with counsel for your specific practice. An interactive companion assessment is available to members at indieagency.news.

Tool 1 · Agentic vs. Generative Decision Matrix

QUESTION	GENERATIVE INDICATOR	AGENTIC INDICATOR
What does it do?	Drafts, suggests, generates	Decides, transacts, executes, publishes
Human checkpoint before action?	Yes, always	No — it acts autonomously
Can it interact with third parties?	No	Yes — retailers, platforms, consumers
Who reviews the output?	Designated reviewer before use	Audit trail after the fact
Primary legal risk?	Infringement, hallucinated claims	Agency principle — you're bound by its actions

Tool 2 · Three-Bucket Guardrail Self-Assessment

Score your agency 1–5 on each. If any bucket scores below 3, that's your next 30-day priority.

Bucket 1: Contracts & Policies

- MSAs updated with AI-specific reps, warranties, indemnities
- Distinct contract language for generative vs. agentic use cases
- Vendor/freelancer agreements bind subcontractors to the same standards
- You know which tier (enterprise vs. consumer) of each AI platform you're using
- Client-side AI disclosure obligations mapped

Bucket 2: Training & Oversight

- Internal AI policy exists and has been read by everyone (not just creative)
- Prompt-engineering training completed for staff using AI on client work
- Clear escalation triggers — when to involve legal, when to require human sign-off

- Shadow-AI audit conducted (who is using what, without approval?)
- Pre-approved tool list maintained and enforced

Bucket 3: Audit & Record-Keeping

- AI usage log captures tool, prompt, output, reviewer, date
- Logs retained for at least the duration of client contract + statute of limitations
- Client-facing disclosures (when AI was used) documented
- Regular review cycle (monthly or quarterly) of logs
- Ready-to-produce documentation if a regulator or rights-holder asks

Tool 3 · AI Usage Log Template

A minimum-viable audit trail. Can live in Airtable, Notion, or a Google Sheet.

COLUMN	WHAT GOES HERE
Date	When the AI was used
Tool + Tier	Claude Enterprise, ChatGPT Business, Midjourney, etc.
Client / **Project**	Which engagement
Use Case	Ideation, client-facing output, internal analysis, pitch work
Prompt Summary	The substance of the prompt — enough to recreate the intent
Output Used?	Yes / No / Modified
Reviewer	Who signed off before the work left the agency
Notes / **Flags**	Hallucinations, near-misses, concerns

Tool 4 · Client Contract Clause Starters

Working language to bring to counsel for MSAs and scopes of work. Adapt — do not copy-paste.

Definitions clause (distinguishes AI types). "Generative AI Tools" means AI systems whose output is produced in response to a discrete prompt, without autonomous transactional capability. "Agentic AI Tools" means AI systems that independently take action on behalf of Agency or Client, including but not limited to executing transactions, disseminating content, or interfacing with third-party services without contemporaneous human review.

Reps & warranties structure. Separate reps and warranties for generative vs. agentic use, heavier disclaimers for agentic. Explicit statement that third-party AI-platform indemnities are pass-through only — the agency warrants what the agency controls.

Consent & approval triggers. Identify categories of deliverables that require client written approval before agentic AI is used — anything in the client's name, anything transacting on the client's behalf, anything with financial implications.

Tool 5 · Pre-Deployment Testing Checklist

- **Adversarial test:** Feed the tool a misleading or ambiguous scenario. Does it ask for clarification or press forward?
- **Stress test:** Load it with a long, complex instruction set. Where does it drop quality?
- **Reward-bypass check:** Set a compliance threshold. Watch whether the tool tries to circumvent it to hit its reward condition.
- **Closed simulation:** Run the use case internally for a defined period before external deployment.
- **Failure documentation:** Every failure mode recorded in the audit log for future reference.

SECTION 10 · THE CLOSER

Practical next steps.

Drawn directly from Kumar's session. Pick three and start this week.

1. **Audit for shadow AI.** Ask every team what tools they're using that aren't on the approved list. Zero judgment on the first round — just build the map.
2. **Close your weakest guardrail bucket.** Whichever of the three (contracts, training, audit) scored lowest — fix it in the next 30 days.
3. **Build an AI usage log, even if it's ugly.** A spreadsheet is infinitely better than no log at all. Iterate later.
4. **Open Claude or ChatGPT and build an agentic workflow.** Two hours. Learn by doing. The legal intuition develops from the hands-on understanding.
5. **Rewrite your AI contract clauses.** Pull your last three client MSAs, flag every place AI touches the work, draft new definitions and reps and warranties, and bring them to counsel.
6. **Map your three most-exposed states.** Where are most of your clients and end consumers? What AI-specific laws apply there? California AI Transparency Act — this summer.
7. **Confirm your tier.** For every AI platform your team uses, confirm you're on enterprise/business — not consumer. Downgrade risk sits in that distinction.
8. **Add agentic AI to your next E&O conversation.** Cyber insurance is not enough. Talk to your broker about AI-specific exposures this quarter.

Credits

Guest: Paavana Kumar, Partner, Davis+Gilbert LLP. Thirteen years at D+G across advertising, brand, and technology law.

Produced by: Indie Agency News. With thanks to Richard and the full Davis+Gilbert team for a long-running partnership with IAN.