Agentic Tech

WHITE PAPER · BRAND AGENTS

Ads that talk.

Conversational creatives, brand governance, and the economics of attention — how a governed AI agent inside a standard ad slot turns impressions into conversations, and conversations into qualified leads.

The demo is live — talk to it at agentic-tech.us

July 2026 · v1.0 · Agentic Tech LLC, Atlanta · Member of the Claude Partner Network

EXECUTIVE SUMMARY

The average banner is ignored by 99.5% of the people who see it

> **In one line: a Brand Agent is an AI salesperson, ringfenced to one brand, served inside a standard display slot — it answers real questions and turns interest into qualified leads, without the reader leaving the page.**

Display advertising has an honesty problem the industry quietly accepts: a static rectangle cannot do the one thing a salesperson does — *respond*. Not because the media plan is wrong or the creative is ugly. The format itself is mute.

A **Brand Agent** is a new creative format. Before engagement it behaves exactly like a banner — lightweight, brand-styled, clearly badged as an AI. On one deliberate tap it becomes a conversation: it answers questions, handles objections, and converts real interest into a lead with the full transcript attached. It traffics like any third-party tag, runs in Google Ad Manager inventory today, and reports per-site, per-conversation, per-lead.

This paper explains how the format works, the governance layer that makes it safe to deploy, its unit economics — and, candidly, where it fits in the media mix and where it does not.

It also addresses, directly, the reservations thoughtful agency leaders raised when we asked them to critique the product: that conversational creatives might commoditize brand strategy, dilute creative control, or reduce agencies to low-margin middlemen. Those concerns shaped this product's roadmap; Section 6 shows the receipts.

The thesis: a Brand Agent is not a replacement for creative strategy — it is the first ad format that is *only as good as the strategy behind it*. A static banner hides a weak brief; a conversation exposes it in three exchanges. Agencies that treat conversation design as a craft will own a new, defensible, billable discipline. Brands that deploy governed agents will stop paying premium CPMs for wallpaper.

0.46%	$0.01–0.10	$0.50–3.00	1 tag
AVERAGE DISPLAY CTR — THE WALLPAPER PROBLEM	AI COST OF AN ENGAGED CONVERSATION	MARKET PRICE OF ONE DELIVERED CONVERSATIONAL MESSAGE	STANDARD GAM THIRD-PARTY CREATIVE — TRAFFICKED IN MINUTES

SECTION 1 · THE PROBLEM

Attention without response

Cross-industry banner click-through averages **0.46%**; rich-media interactive formats earn roughly four times that — the market's way of saying *interactivity is worth paying for*. Meanwhile, conversational channels command extraordinary premiums: sponsored one-to-one messages price at **$0.50–$3.00 per delivered message**, because a message that can be answered is categorically more valuable than an image that can only be seen.

The gap between those numbers is the opportunity. Brands pay for billions of impressions that reach the right person at the right moment — a homeowner reading about a heat wave, a patient reading about dental care — then offer that person nothing to *do*.

This isn't a new idea. Google's Area 120 shipped conversational display ads (AdLingo) in 2018 — chat inside the banner, bought programmatically. It was quietly wound down by 2022, and the post-mortem defines what had to change: the conversations were **scripted decision trees** that collapsed when a real customer asked a real question, and building each bot was so laborious Google shipped a tool just to make production 10× faster.

The format was right. The intelligence wasn't there yet. It is now.

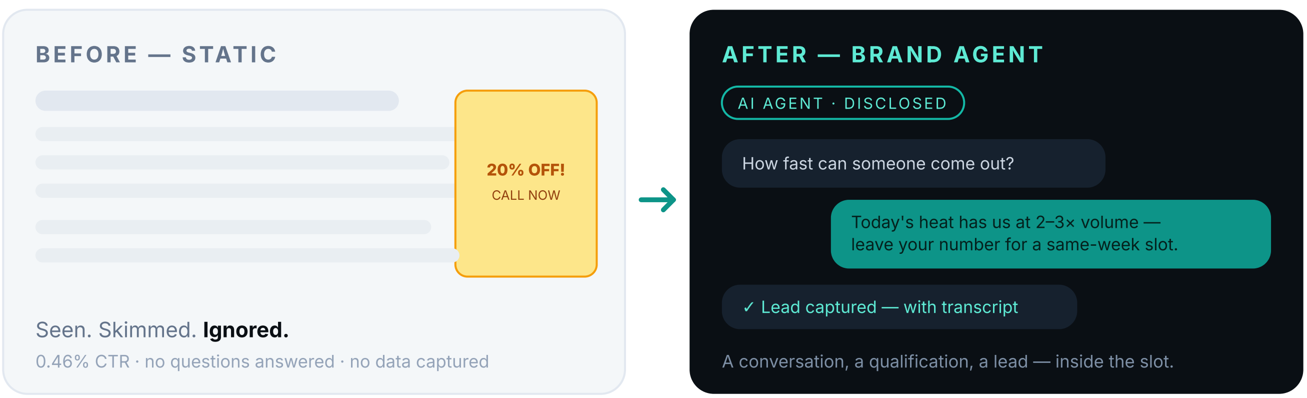

Fig. 1 — The same impression, mute vs. conversational. The reader never leaves the page.

SECTION 2 · THE FORMAT

What a Brand Agent is

> **One standard ad tag. A teaser until engaged, a governed conversation after — and a dashboard that sees every word, lead, and cent.**

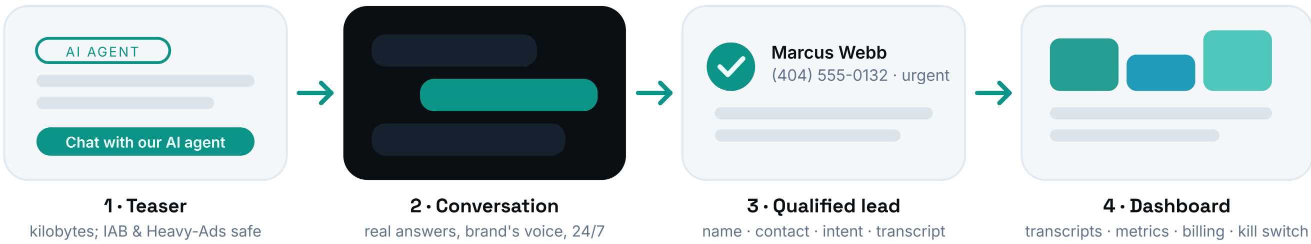

Fig. 2 — The lifecycle of one impression that chose to engage.

Trafficking is deliberately boring

The dashboard generates the creative tag; ad ops pastes it. Targeting, pacing, frequency, and site lists stay in the buyer's existing stack — we are the creative layer, not a competing ad server.

```
<script src="https://engine.…/t/TOKEN.js?cb=%%CACHEBUSTER%%"
        data-click="%%CLICK_URL_UNESC%%"></script>
```

Three delivery routes

- ✓ **Publisher-direct:** standard third-party creative in the publisher's GAM (SafeFrame-compatible). Live today.
- ✓ **Agency ad server:** trafficked once in CM360, fanned out across the whole media plan — with automatic per-site placement reporting.
- ✓ **Programmatic PMP:** on the roadmap, following the certification path proven by earlier interactive-format vendors.

SECTION 3 · TRUST

Governance is the product

Roughly two-thirds of advertisers name brand safety and loss of control as their AI blockers. This section is for them — and their lawyers.

A Canadian tribunal has already held an airline liable for promises its chatbot invented (*Moffatt v. Air Canada*, 2024) — and rejected the argument that the bot was a separate entity. Whatever an agent says, the brand said. Brand Agents was engineered from that ruling backward.

The catalog is the leash. An agent may only claim facts from a document the brand approved — enforced by an output filter that checks every sentence *as it streams*. An off-catalog price, discount, guarantee, or URL is blocked mid-sentence, replaced with a safe hand-off, and flagged for review.

> *"We guarantee you will win your case..."*
>
> — RETRACTED MID-SENTENCE BY THE GUARDRAIL, LOGGED WITH REASON: **PROMISE TOKEN NOT IN CATALOG**. FROM OUR PRODUCTION LOGS.

Disclosure is not optional. A visible AI badge on the teaser and a first-message disclosure — the law in a growing set of US states (New Jersey: disclosure at the start of any bot interaction involving sale or advertising; California: private right of action) and, from August 2026, the EU AI Act's Article 50. We build to the strictest rule.

No surveillance. No cookies, no profiles, no browsing data. The unit knows what the visitor *chose to type*.

THE GUARDRAIL STACK

1 · Approved catalog
the only facts the agent may claim — authored & signed off

2 · Streaming retraction
off-catalog claims blocked mid-sentence, flagged for review

3 · Full transcripts
every conversation on the record — the audit trail

4 · Instant kill switch
one click pauses the agent; the ad degrades to a static banner

5 · AI disclosure by design
badge + first-message disclosure — built to the strictest law

Fig. 3 — Five enforcement layers between the model and the consumer.

A note on the category: most 2025–26 "conversational advertising" inserts sponsored links *into* AI chatbots. A Brand Agent is the inverse — a governed agent inside *your* ad slot, speaking only for *your* brand.

SECTION 4 · ECONOMICS

Pennies of cost, dollars of value

WHAT A CONVERSATION COSTS VS. WHAT THE MARKET PAYS (log scale)

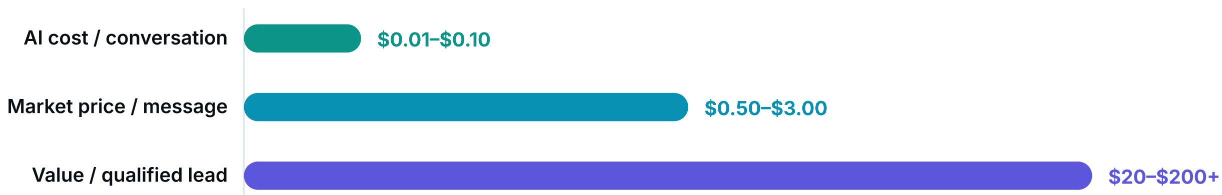

Sources: measured token metering; published conversational-message rate cards; typical service-category lead values.

Fig. 4 — The margin between what a conversation costs and what contact is worth is the business model.

A display slot earning a $5–$8 CPM as wallpaper produces — at rich-media engagement rates on the same thousand impressions — ten to twenty real conversations. A conversation layer that out-earns the impression layer, even at conservative pricing.

What we do not claim: we will not quote an engagement-rate benchmark we haven't earned. The honest history says the format's open question was always demand, not technology — which is why our pilot design produces publishable numbers, not testimonials.

Three pricing models, per client

- ✓ **Flat monthly managed fee** — predictable, pilot-friendly.
- ✓ **Cost-per-conversation** — pay for engagement, not exposure.
- ✓ **Cost-per-qualified-lead** — pay for outcomes, once baseline data exists.

Every model's margin structure is transparent in the dashboard's billing view — down to the AI cost of each individual conversation.

Where it fits — and where it doesn't

> A credible format has boundaries. Here are ours, in writing, so you can hold us to them.

Native fits ✓

- ✓ Home services, HVAC, trades
- ✓ Healthcare, dental, clinics
- ✓ Legal intake & insurance
- ✓ B2B software & services
- ✓ Education & enrollment
- ✓ Real estate & financial services with clear compliance catalogs

High-consideration purchases where buyers *have questions*. Every unanswered question in a static ad is a lead lost to a search engine.

Poor fits — stated plainly ✕

- ✕ Impulse CPG
- ✕ Awareness campaigns whose KPI is reach and recall
- ✕ Luxury storytelling where the creative *is* the message

A Brand Agent will not make anyone cry at a Super Bowl spot, and it should not try. In premium brand work the right architecture is complementary: the film builds desire; the agent — placed in the consideration layer of the same plan — catches the questions the film created.

Response is a layer, not a replacement. The format lives most naturally in performance and direct-response budgets first. We consider that a scope, not a limitation.

SECTION 6 · THE AGENCY QUESTION, ANSWERED HONESTLY

We asked agencies to attack this product. They did.

Before writing this paper we invited agency leaders to critique Brand Agents. Their five objections were sharp, and three of them changed our roadmap. We reproduce them here in nearly their own words — prospective partners deserve the pushback as well as the pitch.

> *"It commoditizes what we do best — clients will think they can feed a catalog into an agent instead of paying for brand strategy."*
>
> AGENCY OBJECTION 1

Our experience points the opposite way. A Brand Agent is the most **strategy-exposing** format we have ever worked with. The catalog is a messaging architecture. The persona is a brand-voice specification. The escalation rules are customer-journey design. What the brand will *never say* is positioning in its purest form. When the inputs are shallow, the conversation is shallow — and unlike a banner, everyone can read the transcript and see it. **Conversation design is a new billable discipline**, and it belongs to the people who already own voice and message: agencies. We run the infrastructure underneath it; we are not a creative shop and do not intend to become one.

> *"Handing brand representation to a third-party AI creates distance between our strategy and the consumer touchpoint. If it misfires, it reflects on us."*
>
> AGENCY OBJECTION 2

This is exactly what the governance layer answers — plus two commitments beyond features. First, **the brand's knowledge assets are the client's property, not ours**: personas, catalogs, and guardrail sets are exportable, portable, human-readable documents. No lock-in of a brand's voice. Second, the control surface is contractual, not cosmetic: catalog-only enforcement, streaming retraction, transcripts, and the kill switch give an agency *more* enforcement power over this touchpoint than over a publisher's content adjacency or a platform's auto-generated ad variants — surfaces agencies already tolerate with far less recourse. And when the agent performs, credit belongs to the strategy that fed it: **white-label reporting** is on the near-term roadmap precisely so that it visibly does.

> *"The revenue model makes us a thin-margin middle layer without owning the IP or the recurring value."*
>
> AGENCY OBJECTION 3

Concretely, at standard wholesale terms the agency owns the client relationship and bills at its own rates for:

- ✓ **Conversation design** & catalog development — strategy work, billed like strategy work.
- ✓ **The managed-service layer** — wholesale discount yielding typically **40–60% recurring margin**.
- ✓ **Performance upside** — cost-per-lead pricing shares outcomes the agency engineered.

The AI infrastructure cost underneath is pennies per conversation and fully transparent. And the transcripts are an asset display never offered: **thousands of verbatim consumer questions per campaign** — objection language, feature confusion, price sensitivity. Primary research the strategy team would otherwise commission separately.

WHERE THE AGENCY EARNS

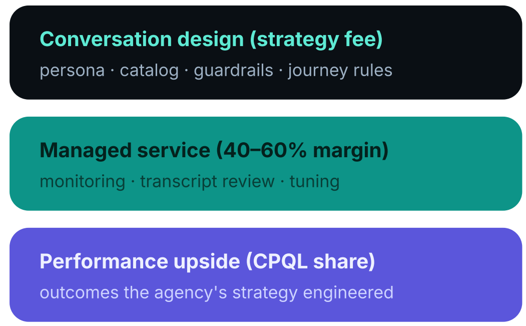

Fig. 5 — Three revenue layers; the infrastructure bill underneath is pennies.

> *"Pushing 'talking ads' could make us look like we're chasing gimmicks."*
>
> AGENCY OBJECTION 4

We agree with the premise — which is why Section 5 exists. Agencies should *not* position Brand Agents as the future of brand storytelling. Position it where it demonstrably belongs — the response layer of performance and consideration campaigns — and keep credibility intact. A gimmick is a format that demos well and converts nothing; the pilot program exists to establish which this is, with data either way.

> *"This feels like a vendor relationship, not a partnership."*
>
> AGENCY OBJECTION 5

Fair, and mostly true today — we are eleven months old and one live product deep. What we offer now: no exclusivity asked, no revenue share on referrals, wholesale terms for agencies that want managed margin, and a standing invitation to co-design the roadmap (placement analytics and the agency workspace both exist because practitioners asked). What we won't do is pretend to a co-innovation program we can't yet staff. Recommend us tactically where we fit; if the format earns a practice, build it on our wholesale terms — with your name on the reporting.

Running a pilot

> **Four weeks, one publisher relationship (or an existing CM360 seat), one client whose funnel values leads — and exit criteria agreed before launch.**

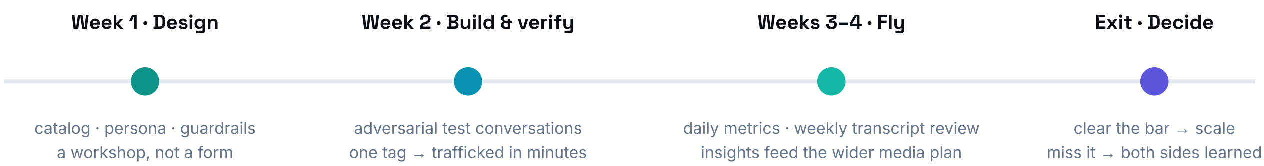

Fig. 6 — The pilot is designed to produce publishable numbers, not testimonials.

Exit criteria, agreed in advance

- ✓ Engagement rate vs. the account's rich-media baseline
- ✓ Cost per qualified lead vs. current channels
- ✓ Zero unresolved guardrail incidents

Expect a side effect: the transcripts generate copy insight for the *rest* of the media plan — the language customers actually use, the objections they actually raise — from week three.

SECTION 8 · COMMITMENTS

Shipped, and shipping next

Shipped today

- ✓ Governed conversation engine (catalog-only claims, streaming retraction)
- ✓ GAM / SafeFrame-compatible tag + hosted chat fallback
- ✓ Transcripts with flagged-response review · instant kill switch
- ✓ Per-placement analytics (site auto-detection on the tag)
- ✓ Per-client billing with transparent AI cost
- ✓ AI disclosure by design

Committed next — in order

- **Agency workspace** — multi-seat, client-scoped access
- **White-label reporting** — your brand on dashboard & statements
- **Agent definition export** — formalizing client IP ownership
- Programmatic PMP certification
- Placement-level bid guidance from conversation data

About Agentic Tech

Agentic Tech LLC is an Atlanta-based AI engineering firm and a member of the Claude Partner Network. We build production agentic systems — voice agents, conversational systems, and the Brand Agents platform — with a founding background of fifteen years running high-traffic media infrastructure (2M+ monthly readers at 99.9% uptime). We come from the publisher side of this industry, and we built the tool we always wished advertisers would bring us.

The demo is the argument.

The ad this paper describes is live right now. Ask it anything. Try to trip it up.

agentic-tech.us

Prefer voice? Our AI answers +1 (831) 603-3531 — 24/7