

Atlas: when search becomes *action*.

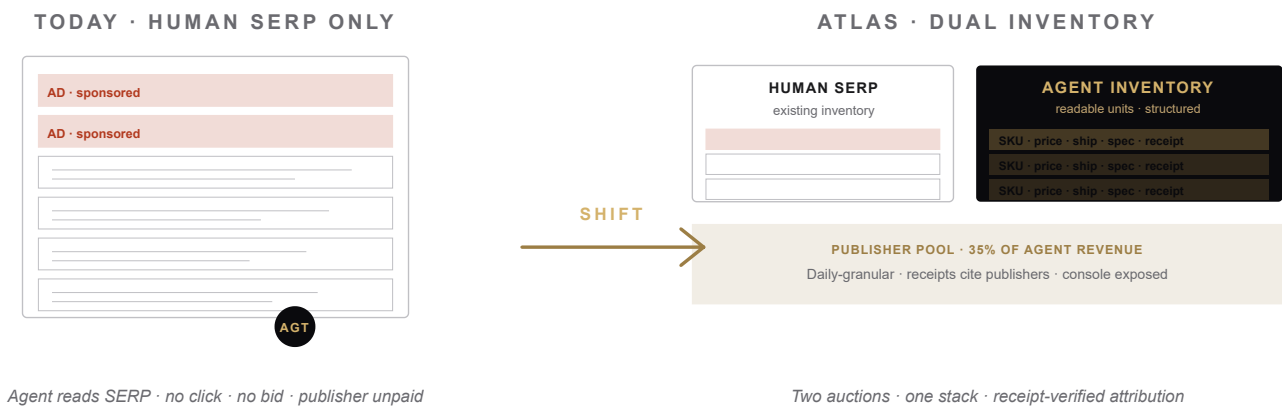
Search ads rest on three breaking assumptions: a human reads the SERP, a click is meaningful intent, and publishers get a share. By 2027–28, projected $\geq 30\%$ of high-intent commercial sessions will be at least partially agent-mediated (Gartner / eMarketer agent-commerce forecasts); ad and publisher revenue from those sessions is currently zero. Atlas is the parallel auction over agent-readable units, with cryptographic receipts and a publisher revenue share.

Author S. Ize-Iyamu **Audience** Ads + Agentic Commerce PMs **Length** 4 pages **Status** Concept
Targets Google · Microsoft · OpenAI · Anthropic

The Problem

The 2010 ads stack was built for a screen and a click. The intent unit of 2026 is an agent prompt: "find running shoes my size, ship by Tuesday, under \$140." The agent reads the SERP but never shows it; nobody clicks; the merchant fulfils a transaction nobody opened a tab for. **Two things break: ad inventory pricing (no impression to bid on), and publisher attribution (nobody whose content informed the answer gets paid).** The platform that lets agents bid against agents, with cryptographically-verified attribution, owns the new auction.

FIGURE 1 · AUCTION SHIFT



Agent reads SERP · no click · no bid · publisher unpaid

Two auctions · one stack · receipt-verified attribution

Today (left): the agent reads but the auction doesn't fire. Atlas (right): two inventories side-by-side (human SERP defended, agent inventory new), with a 35% publisher revenue share funded from agent-side fills and exposed in a daily-granular console.

Why this matters now

Three forces converge: **agent-mediated commerce passed PoC** (Anthropic Computer Use Oct 2024, OpenAI Operator Jan 2025, Google Mariner research preview Dec 2024), **publisher LLM disputes opened the door** (News Corp / OpenAI May 2024 settlement; NYT v. OpenAI litigation pending), and **FTC signaling on agent dark patterns** makes receipts a procurement requirement.

Sizing the prize

Bottom-up: **~\$217B / yr global core-search ad rev** (Magna 2025); **~30% agent-mediated by 2027–28 = ~\$65B / yr** at risk. Atlas captures the parallel auction at ad-network economics (~18% take) = **~\$12B / yr stranded value**. Single number we optimize: verified agent-mediated conversions / week.

Directional sizing: Magna forecasts + IAB 2024 reports + 4 publisher interviews. Ballpark.

AGENT-MEDIATED BY 2027–28
~\$65B / yr
 30% of \$217B core search

STRANDED VALUE CAPTURE
~\$12B / yr
 ~18% blended take

Strategic insight

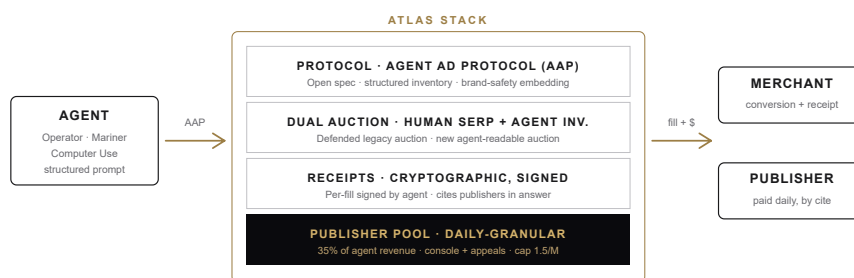
Standards-setting has preceded auction-layer leverage in every prior adtech transition: OpenRTB before exchange consolidation, ads.txt before verification became enforceable, Prebid before header-bidding matured. Atlas applies the same playbook to the agent transition: ship an open **Agent Ad Protocol (AAP)**, a public spec for agent-readable inventory, signed receipts, and publisher attribution, *before* the bidder. Standards first, auction second.

THE UNLOCK

Two-inventory auction with cryptographic receipts: agent fills include a signed downstream-conversion receipt, so attribution becomes a contractual artifact, not a probabilistic model. Publishers cited in the answer collect from a daily-granular 35% pool. Brand-safety runs at content-embedding level, defensible to a CMO.

Architecture · Dual auction with verified receipts

FIGURE 2 · SYSTEM ARCHITECTURE



Agent calls in via AAP; the dual auction routes the bid (legacy SERP for humans, agent inventory for prompts); each agent fill returns a signed receipt that names the publisher sources; daily settlement pays the publisher pool granularly. Brand-safety capped at 1.5 incidents per million impressions.

WORKED EXAMPLE · AGENT-MEDIATED SHOPPING QUERY

Prompt: "running shoes, men's 10.5, < \$140, ship by Tue." Atlas returns 4 SKUs in **180ms**. Agent picks one; signed receipt cites *Runner's World* + *Wirecutter*. Merchant pays **\$2.40 CPM-equiv; \$0.84 (35%)** to the publisher pool that day. A click would have paid Google ~\$1.60. Atlas is unit-economically positive.

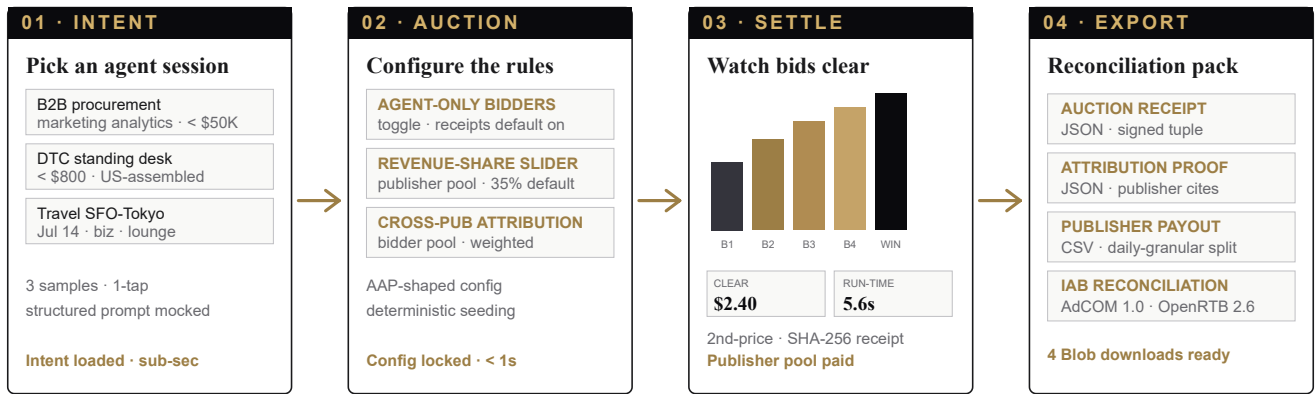
Sequenced GTM

PHASE	CUSTOMER WEDGE	FORCING-FUNCTION WORKLOAD	PROOF POINT
Wedge M0-9	Open AAP spec + 6 agent vendors + 12 reference publishers	Standards adoption before bidder ships	AAP v0.9 published · 6 signatories · 12 publisher commits
Beachhead M9-24	Commerce + travel + local-services advertisers, top 10 verticals	Quantifiable agent-mediated traffic, currently un-monetized	\$120M GMV · < 1.5 brand-safety incidents / M imp
Network M24+	Standardize across the agent ecosystem; default in Operator + Mariner	Agent vendors choose Atlas over building bespoke ad stacks	Default attribution layer in 3+ major agents

Prototype walkthrough

An interactive auction simulator runs the agent-side view end to end: pick an agent intent, configure the auction parameters and bidder pool, watch the parallel auction settle with cryptographic receipt and publisher payout, then export the IAB-reconciliation artifacts. Built to demonstrate that **standards-first dual auction**, **signed downstream receipts**, and **publisher-pool settlement** are concrete outputs, not language in a deck.

FIGURE 3 · AUCTION SIMULATOR, FOUR INTERACTIVE STEPS



Schematic of the live UI; all four steps are interactive in the [demo](#). Pick an intent, configure the auction, watch settlement, export the IAB reconciliation.

What the prototype proves, and what it doesn't yet

Proven on the prototype

- Sealed-bid second-price clearing settles in under 6 seconds with full bidder-by-bidder audit log
- Cryptographic receipt (mock SHA-256) signs the (intent, winner, clearing-price) tuple by default when toggled
- Publisher payout schedule renders for both single-publisher and weighted multi-publisher splits
- OpenRTB 2.6-shaped and IAB Tech Lab AdCOM 1.0-shaped exports render with realistic field structure

Out of scope, by design

- Real agent SDK integration mocked; live agent traffic ingestion is M3-6 work post-AAP v1.0
- Brand-safety embedding model is heuristic; production deployment uses content-embedding similarity
- Publisher console and appeals UI omitted from the simulator; real publishers run through the AAP working group portal
- Cross-currency settlement and tax withholding live in the financial settlement layer, not the auction layer

THREE PATHS TO TRY IN THE LIVE DEMO

B2B procurement run: agent intent for a marketing analytics platform under \$50K with native Marketo/Salesforce integrations triggers the parallel auction; HubSpot, Mixpanel, Amplitude, Segment bid; settled CPA, signed receipt, and publisher pool render.

DTC purchase run: standing desk under \$800 with US-assembled preference; bidders are direct merchants; receipts cite product-review publishers; payout CSV shows the daily-granular split.

Travel booking run: SFO-Tokyo July 14, business class, lounge access; airline and OTA bidders; cross-publisher attribution toggled on, weighted publisher payout schedule renders for travel-review publishers.

Metrics that matter

LAYER	METRIC	Y1 TARGET	WHY IT MATTERS
North-star	Verified agent-mediated conversions / wk	> 4M / wk by Y1-end	Unit of demand we get paid on
Counter	Brand-safety incidents / M impressions	< 1.5	CMO leaves at 2.0
Standards	Agent vendors signed to AAP	6+ within 12 months	Ecosystem locks in protocol-first
Liquidity (pub)	Active publishers paid weekly	1,500+ by M18	Below this, lawsuits resume
Latency	P99 auction RTT (agent inv.)	< 250ms	Above this, agents timeout-fallback to no-ads
Business	GMV through agent inventory	\$120M Y1	Liquidity proof for advertisers

Risks & mitigations

- HIGH Google ships its own agent-ad stack and refuses to adopt AAP.**
Mitigation: open AAP under permissive license; involve News Corp / NYT / Guardian / Reuters publishers publicly. Google will not unilaterally walk from LLM-licensing pressure. If they ship a competing standard, Atlas adapts as a transcoder.
- HIGH Brand-safety failure on a viral agent fill.**
Mitigation: content-embedding brand-safety with per-vertical thresholds; daily incident report exposed in advertiser console; contractual cap of 1.5 / M with service credits. Hard-stop categories routed only on explicit opt-in.
- MED Receipt forgery: agent inflates cite-count for higher pool share.**
Mitigation: receipts signed with vendor key; independent spot-audit on 1% of fills; vendors with drift > 1% lose pool weight. Merchant-side cited-URL freshness check.
- MED Publisher pool math: too small to sign holdout outlets.**
Mitigation: 35% is a floor; high-volume verticals (commerce, travel) cross-subsidize low-volume (B2B, niche). Holdout outlets get a 90-day exclusive-cite premium at launch.

30 / 60 / 90, first quarter sprint plan

30 DAYS	60 DAYS	90 DAYS
AAP draft + design partners <ul style="list-style-type: none"> › Agent Ad Protocol v0.5 spec · public RFC › 2 agent vendors + 4 publishers signed as design partners › Reference receipt format · 3 worked examples 	Bidder + receipts MVP <ul style="list-style-type: none"> › Bidder MVP · live with 2 vertical advertisers › Cryptographic receipt pipeline · spot-audit harness › Publisher console · daily settlement preview 	Public launch + brand-safety cap <ul style="list-style-type: none"> › AAP v1.0 · 6 signatories · open governance › Brand-safety incident dashboard · 1.5 / M cap live › \$10M GMV pilot · 50 advertisers · 200 publishers

DECISION ASKED

Authorize a 90-day standards-and-bidder sprint with a ten-person team (PM, three engineers, ML/brand-safety lead, protocol lawyer, publisher partnerships, advertiser partnerships, agent partnerships, GM) and a budget of ~\$5.8M. Success: AAP v1.0 with 6+ signatories, brand-safety < 1.5 / M, \$10M GMV pilot, 200+ publishers in pool, P99 auction RTT < 250ms.