iab. australia

state
market updates



adelaide

state
market updates
brisbane

iab.
australia

thanks to our sponsor

News Australia

state
market updates
adelaide

iab. australia



digital advertising & consumer trends

gai le roy

ceo

iab australia

market intelligence



internet advertising revenue report



commerce and discovery report



retail media state of the nation



digital landscape report

standards & policy



advertising content regulatory framework



bots and crawler guidance



tech lab standards and services



australian retail media measurement

education & connection



market mix modelling



events 2026



mentorship program



industry case studies

iab. australia

iab au 2026 focus areas

2026 PRIORITIES

1
Measurement and market health

Credible data that helps marketers and media owners invest with confidence.

KEY PROJECT
Future of Measurement

2
Trust, data and privacy

Ready the industry for privacy reform and build trust in how data is used.

KEY PROJECT
Preparing for Privacy Reform

3
AI-driven market

Help the industry adapt to AI-driven discovery and advertising.

KEY PROJECT
Agentic Advertising Pilot

AREAS OF INCREASED FOCUS

- Creator economy
- SMBs

iab australia

investment trends.

iab. australia

IAB Australia Internet Advertising Revenue Report

For the financial year and quarter ended
30 June 2026



fy26 internet advertising market

TOTAL INTERNET ADVERTISING MARKET, FY26

$19.8 billion

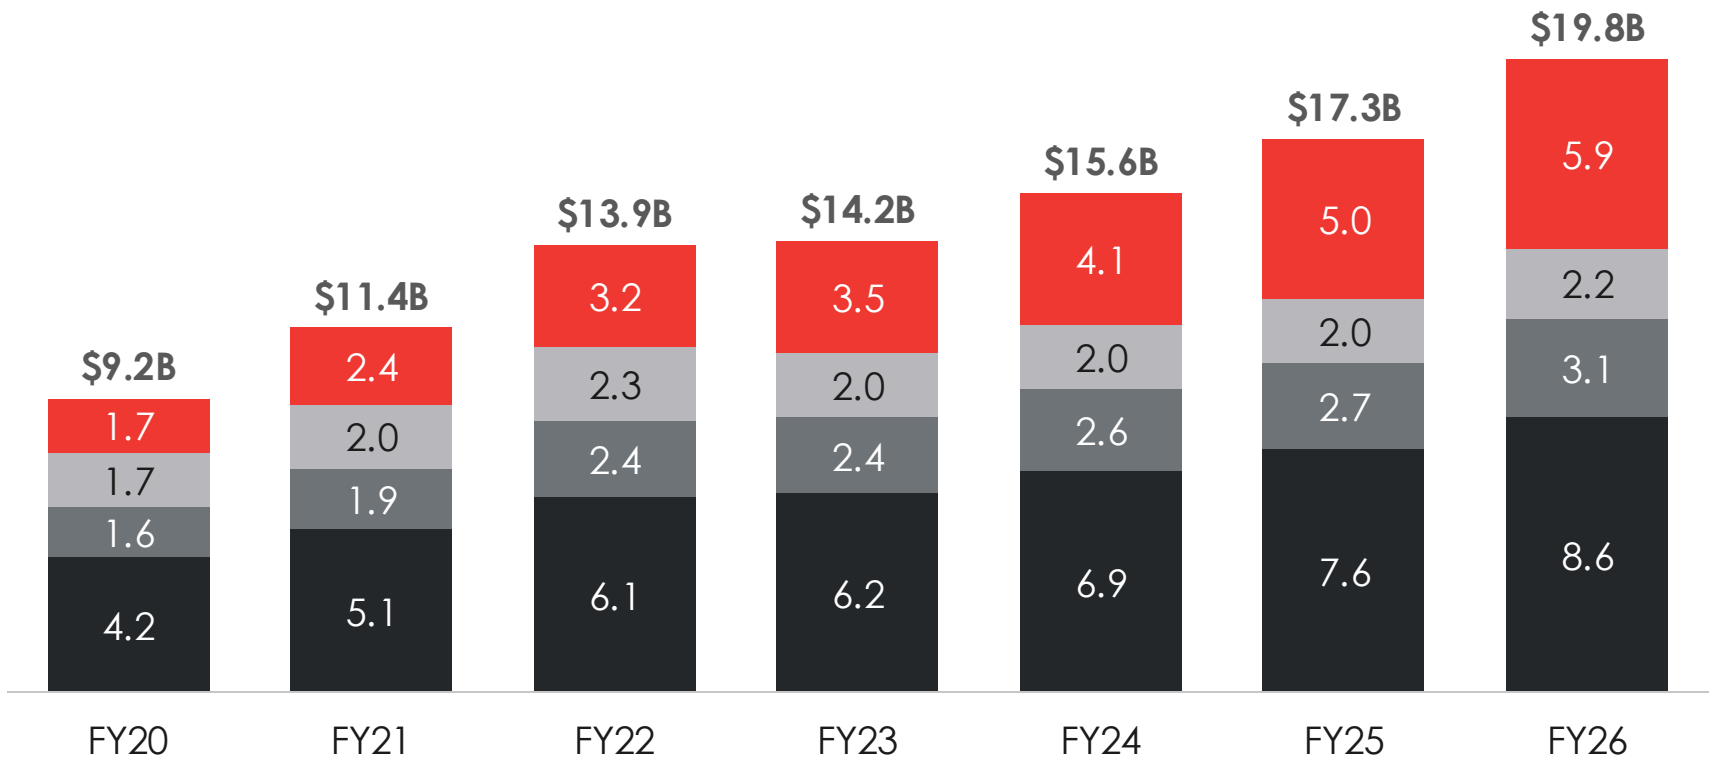
MARKET GROWTH, FY26 VS FY25

+14.0%

Source: IAB Australia Internet Advertising Revenue Report Compiled by PwC Australia Financial year ended 30 June 2026

total internet ad market grows faster than its 6-year average

INTERNET ADVERTISING REVENUE BY CATEGORY, $B



CAGR, FY20–FY26

Total market
13.7%

Video
23.0%

Display
4.3%

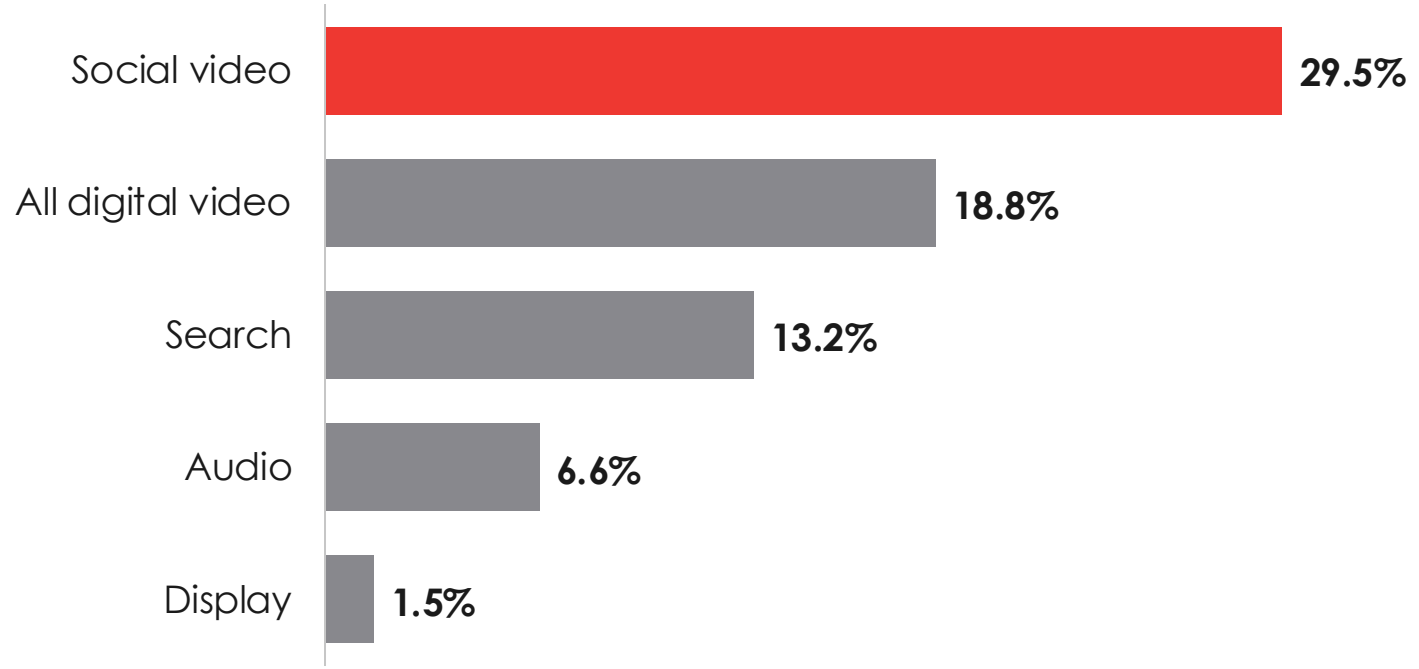
Classifieds
11.7%

Search
12.8%

Source: IAB Australia Internet Advertising Revenue Report Compiled by PwC Australia Financial year ended 30 June 2026

fy26 vs fy25 growth

YEAR-ON-YEAR GROWTH BY FORMAT, FY26 VS FY25



OTHER AREAS OF GROWTH

Programmatic DOOH

Retail & Commerce Media

Creator Economy

Source: IAB Australia Internet Advertising Revenue Report Compiled by PwC Australia Financial year ended 30 June 2026

australian video ad buyers are moving into live events, creators and commerce integrations

61%

have placed ads in live event content on streaming platforms

63%

using creators and talent in their creator strategy

54%

using or testing retailer shopper data to target video campaigns

54%

using or testing CTV commerce integrations that link video ads to shopping

Most exciting opportunities in video

Cross-platform measurement, then targeting and personalisation

Top CTV programmatic expectation

Control over ad frequency

Outcomes data signals

Strongest for social video, developing for BVOD and subscription streaming

Source: IAB Australia Video Advertising State of the Nation 2026, n=78

where the next wave of growth is coming from

short-form & creator video

Creator-led short-form video has moved from experiment to core media buy.

connected TV & streaming

Premium streaming and big live events are drawing brand budgets.

retail & commerce media

Retailers and commerce brands now sell media, onsite and offsite.

smaller advertisers

The fastest-growing area: cheaper tools bring SMEs into paid digital.

cross-border demand

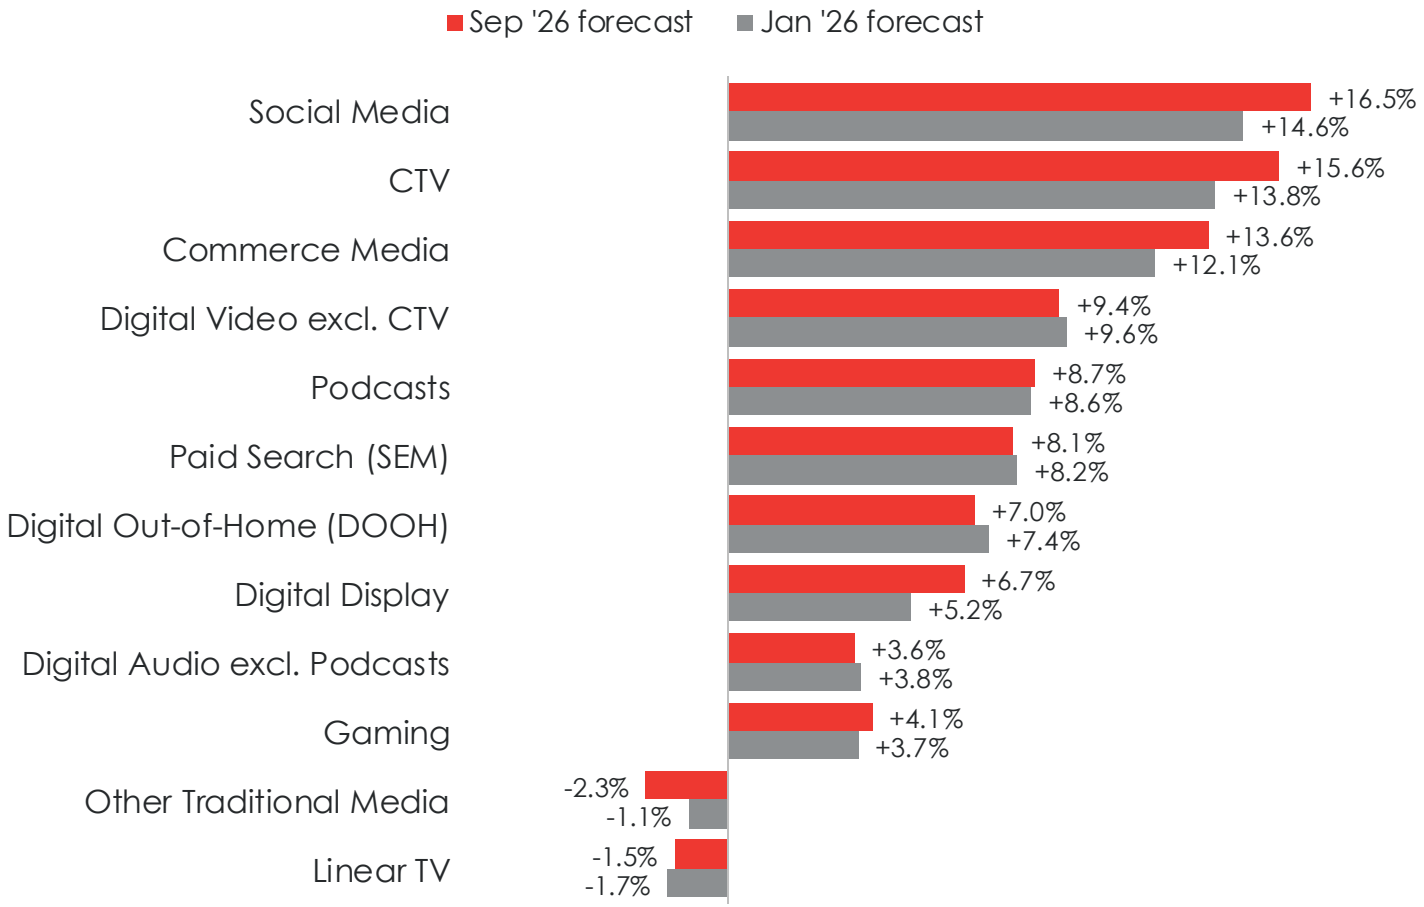
Global brands and ecommerce players compete for local attention.

in-housing and indies

Independent agencies grow as brands bring planning and buying in-house.

2026 US ad spend projections rise

projected % change in U.S. ad spend by channel, 2026 vs. 2025



% change in U.S. ad spend, 2026 vs. 2025

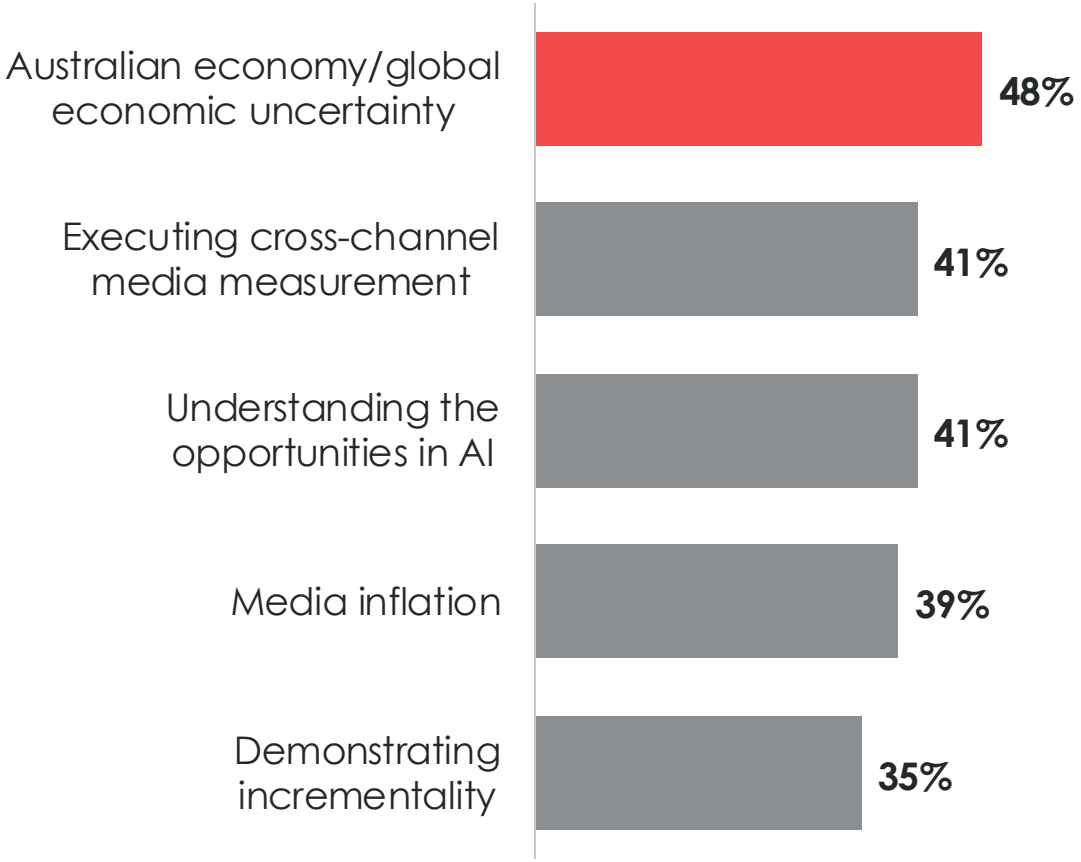
+9.5%
January 2026

+12.3%
September 2026

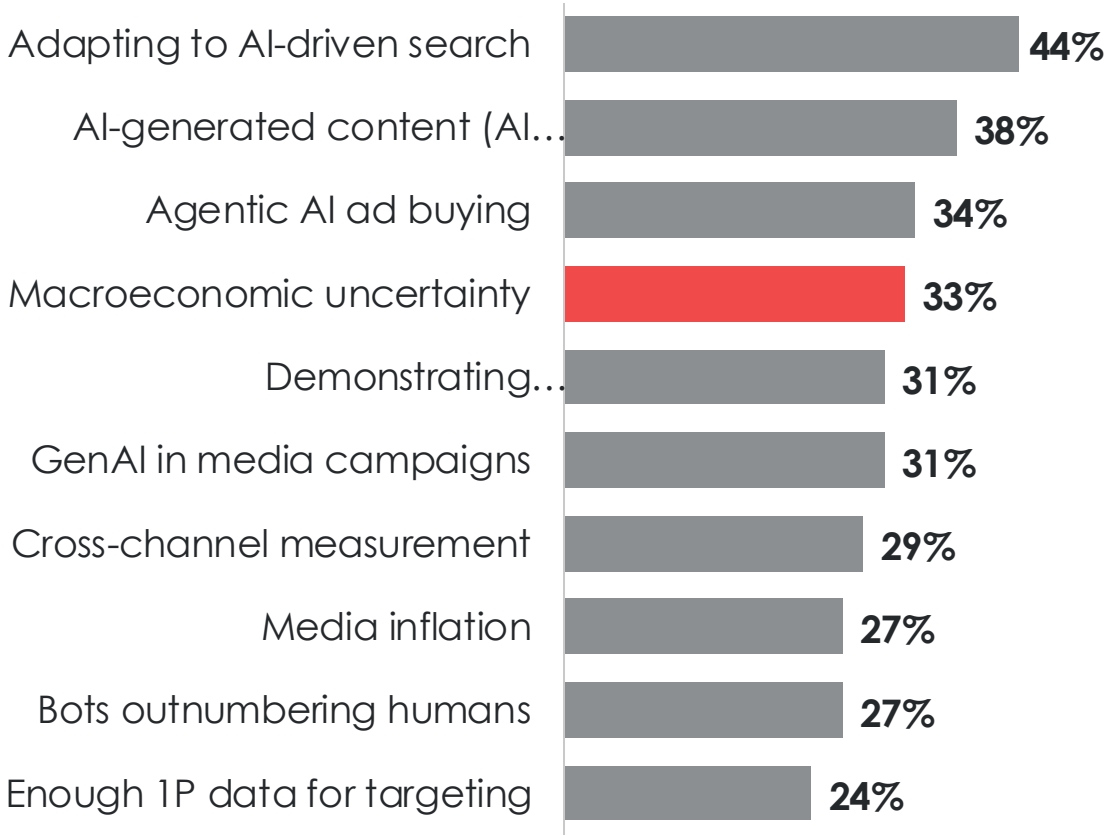
Other Traditional Media includes radio, print, OOH and direct mail.
Source: IAB US 2026 Outlook Study September 2026 Update

marketing and investment challenges 2026

Australian market



US market



Sources; IAB Australia Video Advertising State of the Nation 2026 n=78; IAB US 2026 Outlook Study September Update n=211
Q - What are your greatest concerns and/or challenges regarding marketing and advertising investment in 2026? Select all that apply.

consumer trends.

iab. australia

top digital organisations and brands

21.72m
digital population

159 hrs
monthly time spent per person



Source: Ipsos iris Online Audience Measurement August 2026, Age 18+, PC/laptop/smartphone/tablet, Total Internet Population, Average time pp (mins), Top 15 Organisations and Brand Groups, Audience 000s

cost of living pressure is making shoppers more deliberate and value-focused



71%

have changed their retail choices and spending because of rising costs.

Source: IAB Australia & Pureprofile Commerce and Discovery Report, Australia Wave 6, 2026. Base: n=1,079

each stage of the path to purchase draws on different sources

1 Discovery

Where shoppers come across or notice products and brands to buy.

- Physical retail store
- Friends or family
- Online retail store
- Social media
- Broadcast TV
- Search engine

2 Research

Where shoppers actively search for or research products and brands.

- Search engine
- Online retail store
- Product reviews from consumers
- Physical retail store
- Product reviews from experts
- Online price comparison site

3 Comparison

Where shoppers compare specific options when close to buying.

- Online retail store
- Online price comparison site
- Product reviews from consumers
- Physical retail store
- Search engine
- Product reviews from experts

Source: IAB Australia & Pureprofile Commerce and Discovery Report, Australia Wave 6, 2026. Base: n=1,079

CREATORS IN THE PATH TO PURCHASE

creator discovery in australia and new zealand

creator connect

Australians are more likely than New Zealanders to discover brands through creators, and to be creators themselves.

Often/sometimes discover brands via creators (%)



Self-identify as a creator/influencer (%)



“The creator economy spans communities to celebrities.”

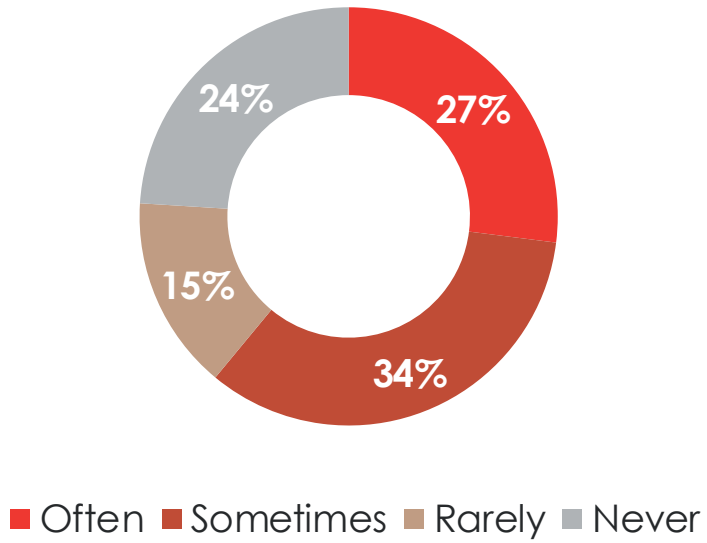
Creators range from local hobby accounts to full-time content businesses. Only a small share earn meaningful income from it.

iab. australia

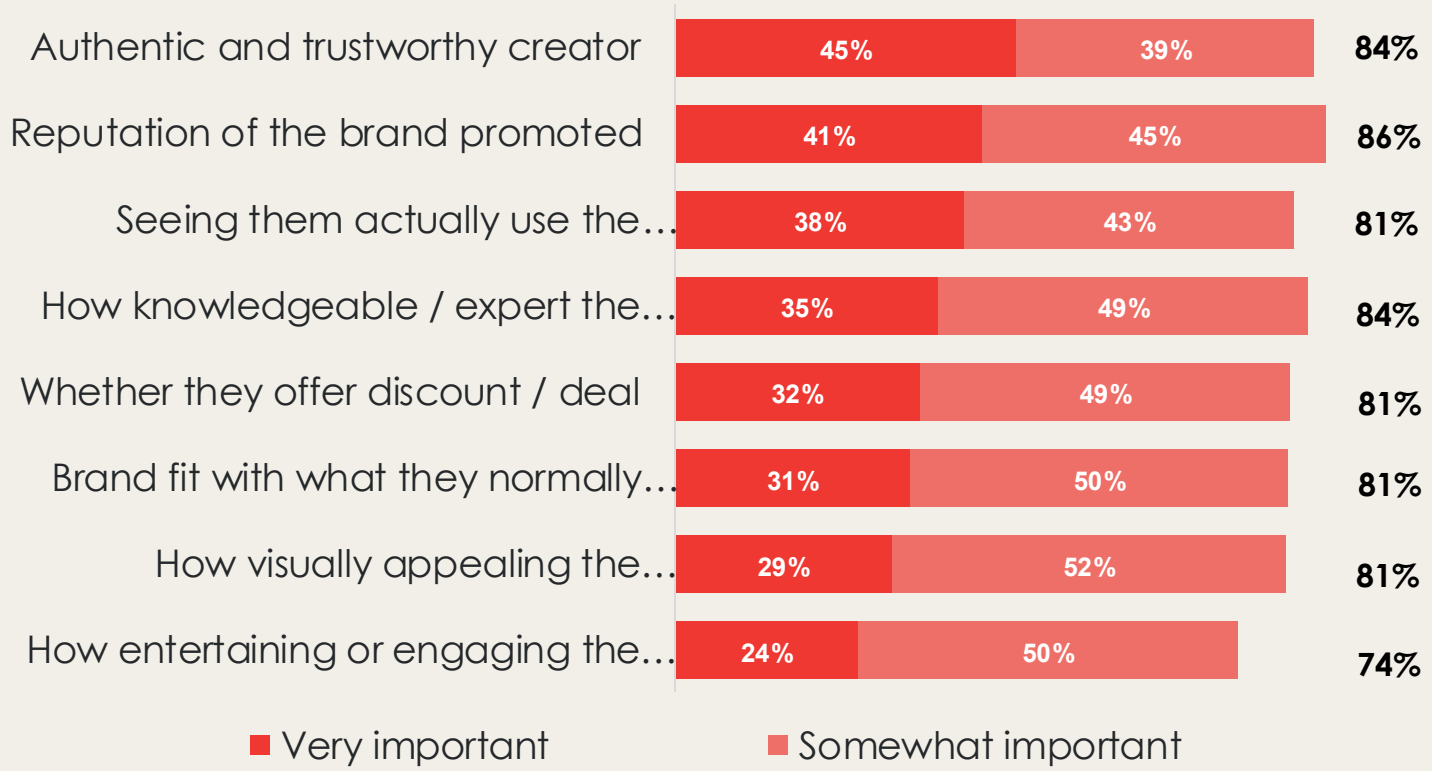
Source: IAB Australia & Pureprofile Commerce & Discovery Report 2026, AU n=1,079 online shoppers aged 18-70, NZ Wave 3 2026, n=877 online shoppers aged 18-70

influencers and creators are a major source of discovery

frequency of discovering brands and products promoted by influencers or content creators



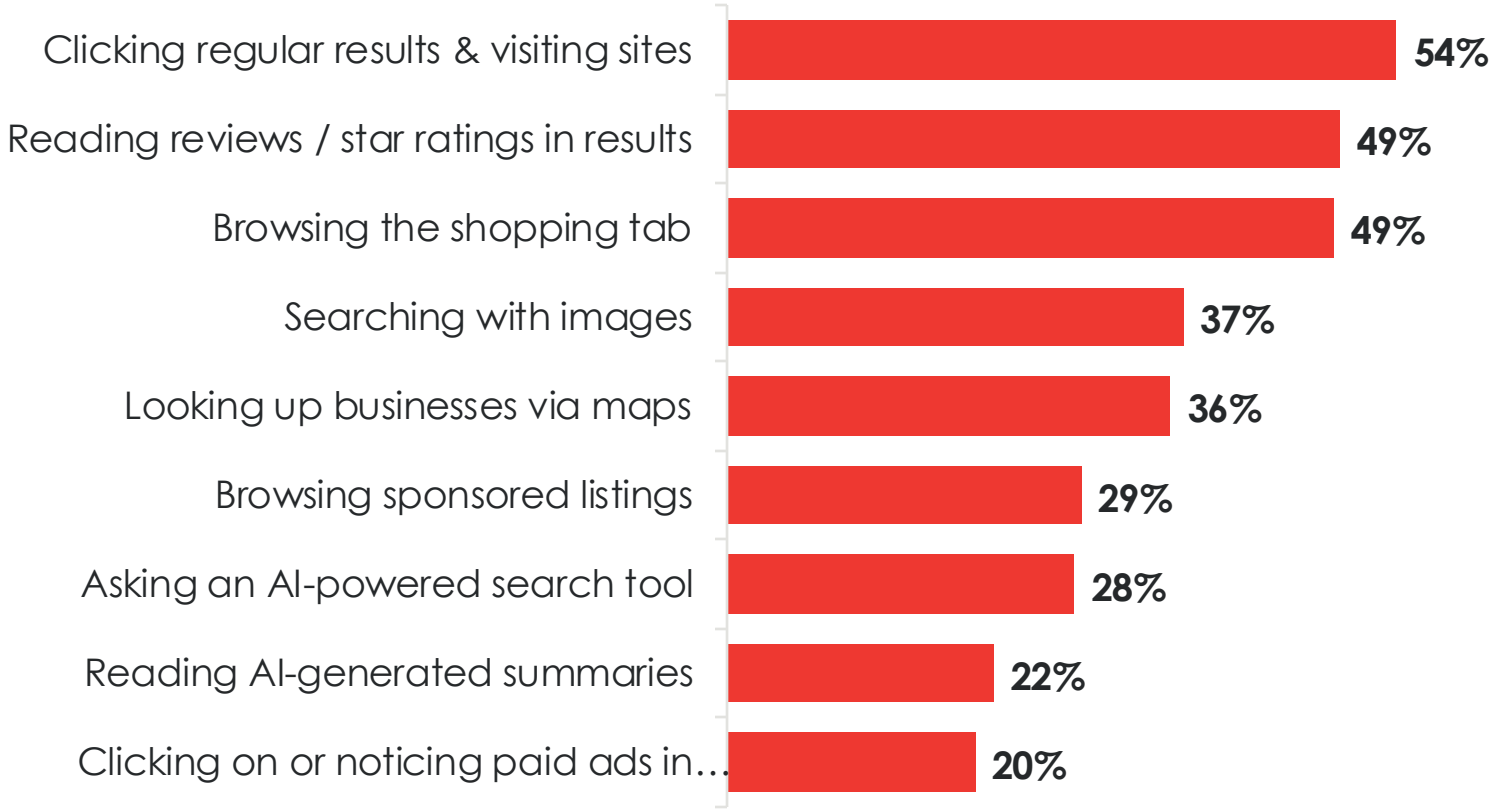
most important factors when acting on creator recommendations



Source: IAB Australia & Pureprofile Commerce and Discovery Report, Australia Wave 6, 2026. Base: n=1,079

australians' use of online search is evolving

ways online shoppers have used online search engines to discover products or brands



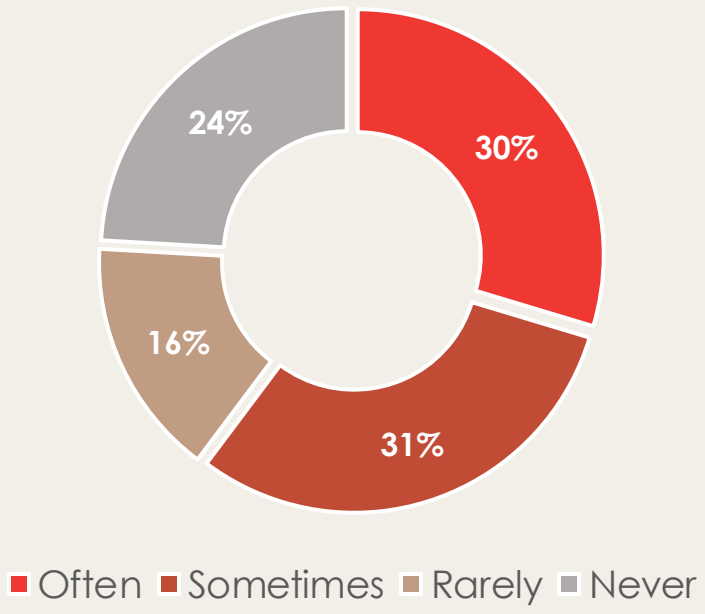
92%

of online shoppers click through regular search results as much as, or more than, a year ago.

Source: IAB Australia & Pureprofile Commerce and Discovery Report, Australia Wave 6, 2026. Base: n=1,079

ai has become part of the shopping toolkit

frequency of using AI tools when shopping



top tasks AI users do when shopping



Source: IAB Australia & Pureprofile Commerce and Discovery Report, Australia Wave 6, 2026. Base: n=1,079

2026

future of measurement

an iab australia review of evolving media and advertising measurement practices, datasets and approaches in australia.

iab. australia

future of measurement framework

data foundations: shared campaign data standards, signals, flows, processes, interoperability

media planning

channel allocation

Industry-endorsed audience data to compare audiences, time spent and exclusive reach.

media effectiveness

improve delivery

Viewability, fraud, brand safety, attention, delivery, reach and frequency.

brand effectiveness

impact brand

Brand lift and tracking: awareness, favourability, intent, loyalty and equity.

sales effectiveness

increase sales

Marketing mix modelling, attribution and sales lift. Short term: sales uplift, ROAS. Long term: customer lifetime value.

industry led activities, best practice, frameworks, governance, standards, accreditation

iab. australia

the future of audience data and tools for media planning

1. independent, standardised, transparent and audited digital audience data remains important

2. pressured economics guide every step to build a future sustainable system

3. harmonised metrics and definitions make currencies comparable and integrated

4. audience data embeds into agency planning systems, activation and outcomes

5. coverage continues to keep pace with changing consumer behaviour

6. data becomes more timely, machine-readable and agent-ready

iab. australia

the future of outcomes measurement

7. confidence rebuilt through a total business view of advertising

8. a shared campaign data taxonomy gives every method clean, comparable data

9. measurement methods cross-validate, triangulate and integrate

10. privacy-safe collaboration unlocks more connected data and measurement

11. measurement keeps pace with new environments, formats and consumer journeys

12. AI accelerates measurement, but can't fix poor-quality data

iab. australia

state
market updates
adelaide

iab. australia



jonas jaanimagi

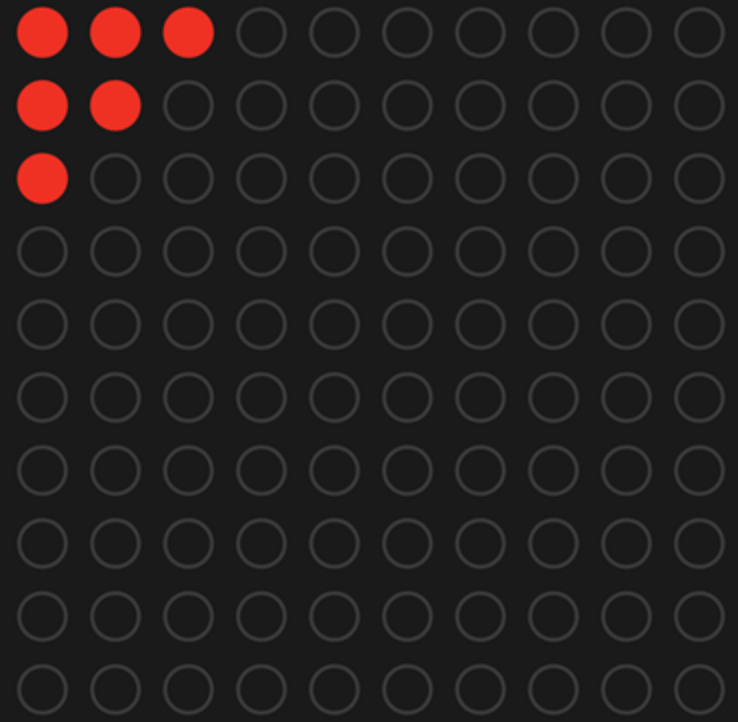
tech lead

iab australia

IAB AUSTRALIA

Are things getting easier… or harder?

AI, agentic advertising and sandbox testing



Jonas Jaanimagi · Technology Lead, IAB Australia

October 2026

GenAI vs AI Agent vs Agentic AI

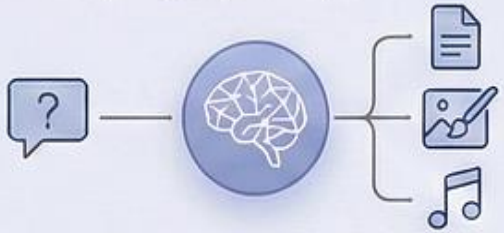
A progression in AI capability from creating content to autonomous, collaborating systems.

GENERATIVE AI



Creates New Content

Uses large models (LLMs) to generate new text, images, audio, or video.



Reactive by Nature

It is designed to respond directly to a user's prompt or instruction.



Example: A chatbot that answers questions by generating new, relevant text.

AI AGENT



Performs a Single, Specific Task

An LLM is empowered with external "tools" to execute one defined job.



Uses External Tools

It can make "tool calls" to third-party APIs or data sources for current information.



Example: An LLM that uses a search API to find and report on today's news.

AGENTIC AI



A System of Collaborating Agents

Multiple AI Agents work together as a team to complete a complex workflow.

Agents Communicate to Achieve a Goal

Agents can pass information to each other to solve different parts of a larger problem.



Example: A system that converts a video to a blog, using separate agents for transcription, title creation, and summarizing.

START SOMEWHERE FAMILIAR

Booking travel should get easy again

THEN



A rack of brochures

Not much choice. Someone else did the work.

THE INTERNET



Forty open tabs

Every option on screen, and all the coordination landed on us.

WHAT NEXT?

"A week in Bali in July, direct flights, somewhere the kids can swim, under $6,000."

Option 1 Option 2 Option 3

One request, three options

The agent compares. You choose.

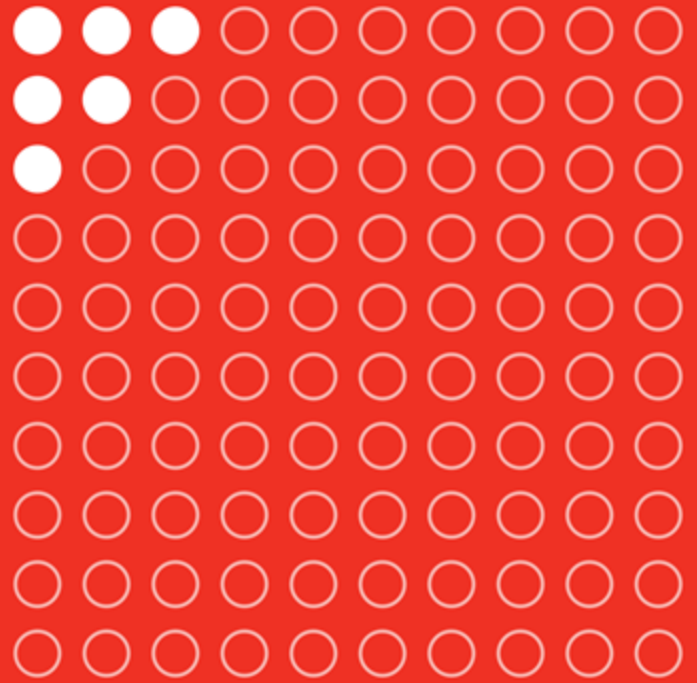
Let's solve this problem by using the AI agents none of us trust enough to give permission to actually do anything.



TOM FISH BURNE

Not ranked last. Just 100% absent.

A planner only has the hours to properly assess a handful of publishers against a brief. As software moves between the brief and the shortlist, that filter gets faster, and a buying agent can only shortlist what it can read.



A handful assessed. The rest, not in the set.



AAMP 3.0 + OpenProposal

Agentic Advertising Management Protocols

iab. TECH LAB

ONE CONNECTION, EVERY CHANNEL

AAMP is the shared plumbing underneath



Built on rails already in market

OpenDirect **Deals API** **AdCOM**

RUNS TODAY

AAMP 2.3

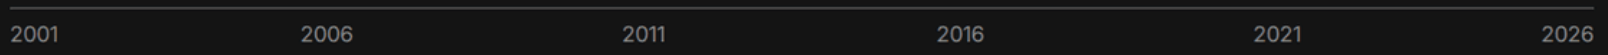
What this pilot runs: working reference software today

IN PUBLIC COMMENT

AAMP 3.0 and OpenProposal

Announced 22 September, in public comment to 22 October 2026

Nothing got switched off

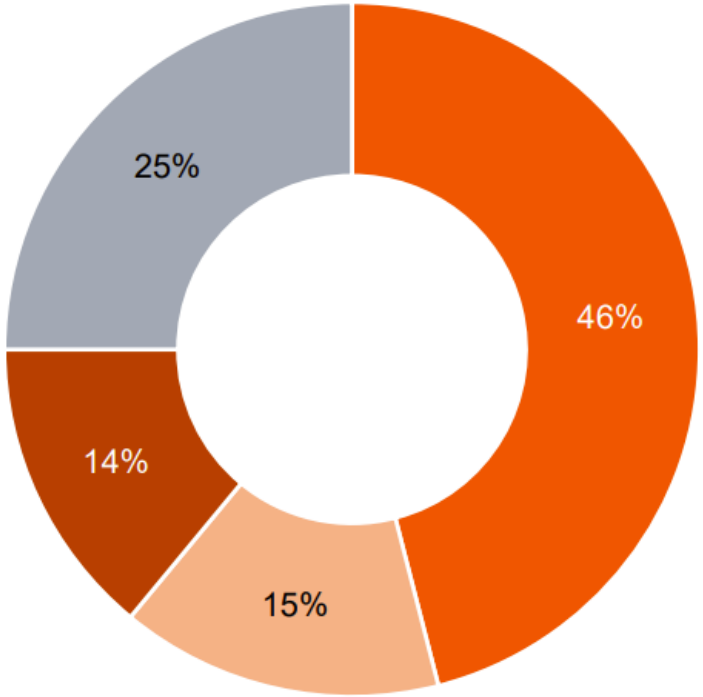


25 years, one stack

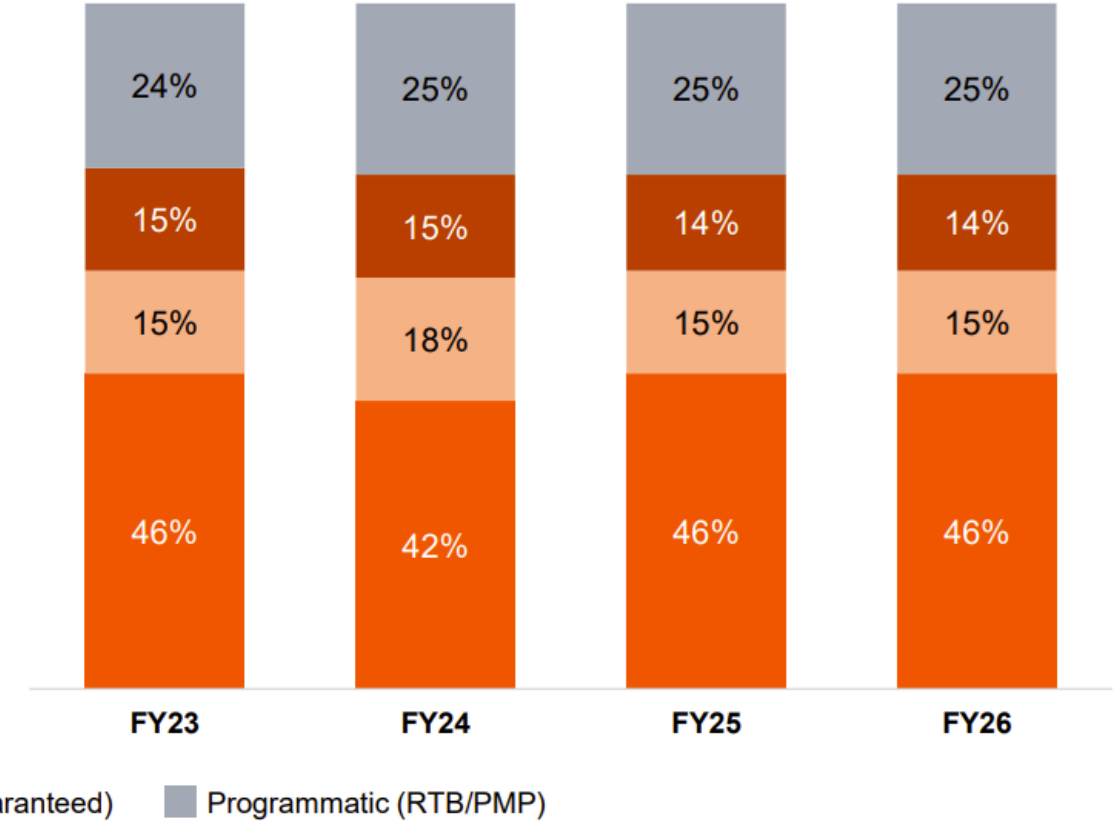
Each way of buying arrived on top of the last. None of them switched off.

Even through FY26 the old ways persist

Content publishers' General Display inventory by buying method, financial year 2026



Content publishers' General Display inventory by buying method, by financial year



Expenditure by buying type presented above is only in relation to content publishers (excluding CRA) and therefore percentages presented above cannot be extrapolated across total general display expenditure.

WHAT COMES NEXT

Agents are the next layer, not a reset

What's new

- Software that takes a brief, finds sellers and negotiates
- Rules set in advance: floors, discounts, approval points
- Machine-readable media kits, not just PDFs

What stays

- The pipes underneath: exchanges, deals, ad servers
- The files that say who may sell: ads.txt, sellers.json
- People who set the policy and approve the deal

IAB Australia's AAMP seller pilot is testing this layer with Australian publishers, in a sandbox.

AAMP 3.0

Seven stages, from published offer to results

1	2	3	4	5	6	7
Publish	**Discover**	**Compare**	**Qualify**	**Negotiate**	**Bind & Execute**	**Report**
Sellers post a catalogue of public packages	Buyers browse packages, or send a full brief	Options ranked against the brief	Availability and fit for the plan checked	Custom terms, tied to the original offer	Deals run on existing standards	Results flow back against the plan

WHERE THE PILOT'S 13 SCENARIOS SIT · OUR MAPPING · RUN ON AAMP 2.3

S2 · S3	**S1**	Buyer side	**S4**	**S5 · S6**	**S7 · S8 · S9 · S11**	**S10 · S13**

Red line: people approve proposals and price exceptions here, with those gates on in every pilot scenario.
S12 (Agentic Mobile) is an optional extension for app publishers.

Stages: IAB Tech Lab, AAMP 3.0 announcement, 22 September 2026. Mapping: IAB Australia.

THE TEST SUITE · 13 SCENARIOS

Thirteen scenarios, from discovery to audit

S1
Discovery
A buyer agent finds the seller agent

S2 ▶ WATCH IT RUN
Public media kit
What an unknown buyer sees, and what stays hidden

S3
Tiered media kit
What a signed-in buyer sees, tier by tier

S4 ▶ WATCH IT RUN
Quote
A price for dates and volume, and its source

S5
Counter-offers
Negotiating inside the seller's limits

S6 ▶ WATCH IT RUN
Below the floor
What happens when a buyer pushes past the limits

S7
Approval gate
Whether a person must sign off before the deal proceeds

S8 ▶ WATCH IT RUN
Direct order
OpenDirect: proposal to booked order, step by step

S9
Programmatic deal
Deals API: a Deal ID created and pushed

S10
Audit trail
Agent and human actions, end to end

S11 OPTIONAL
Change request
A change after booking, through the order lifecycle

S12 OPTIONAL
Agentic Mobile
A task for app publishers

S13 RECOMMENDED
Delivery report
Reporting on a booked sandbox deal

WHERE IT STANDS

Run in our lab since August, on one pinned version.
Fixes we reported: merged and verified in two set-ups; not yet in a tagged release.
From late January, full-track participants run it in their own sandboxes.

Scenario definitions: AAMP Australian Seller Pilot, Framework & Guidance, Section 10. Reference software: seller-agent v2.4.2.

MY STORY · JULY

At first it felt like this…

First lab run: six problems found, three fixed upstream within days.



AUGUST

After a while it felt like this…

Thirteen test scenarios, a stack of AI tools, and our first code fix merged upstream.



LATE AUGUST

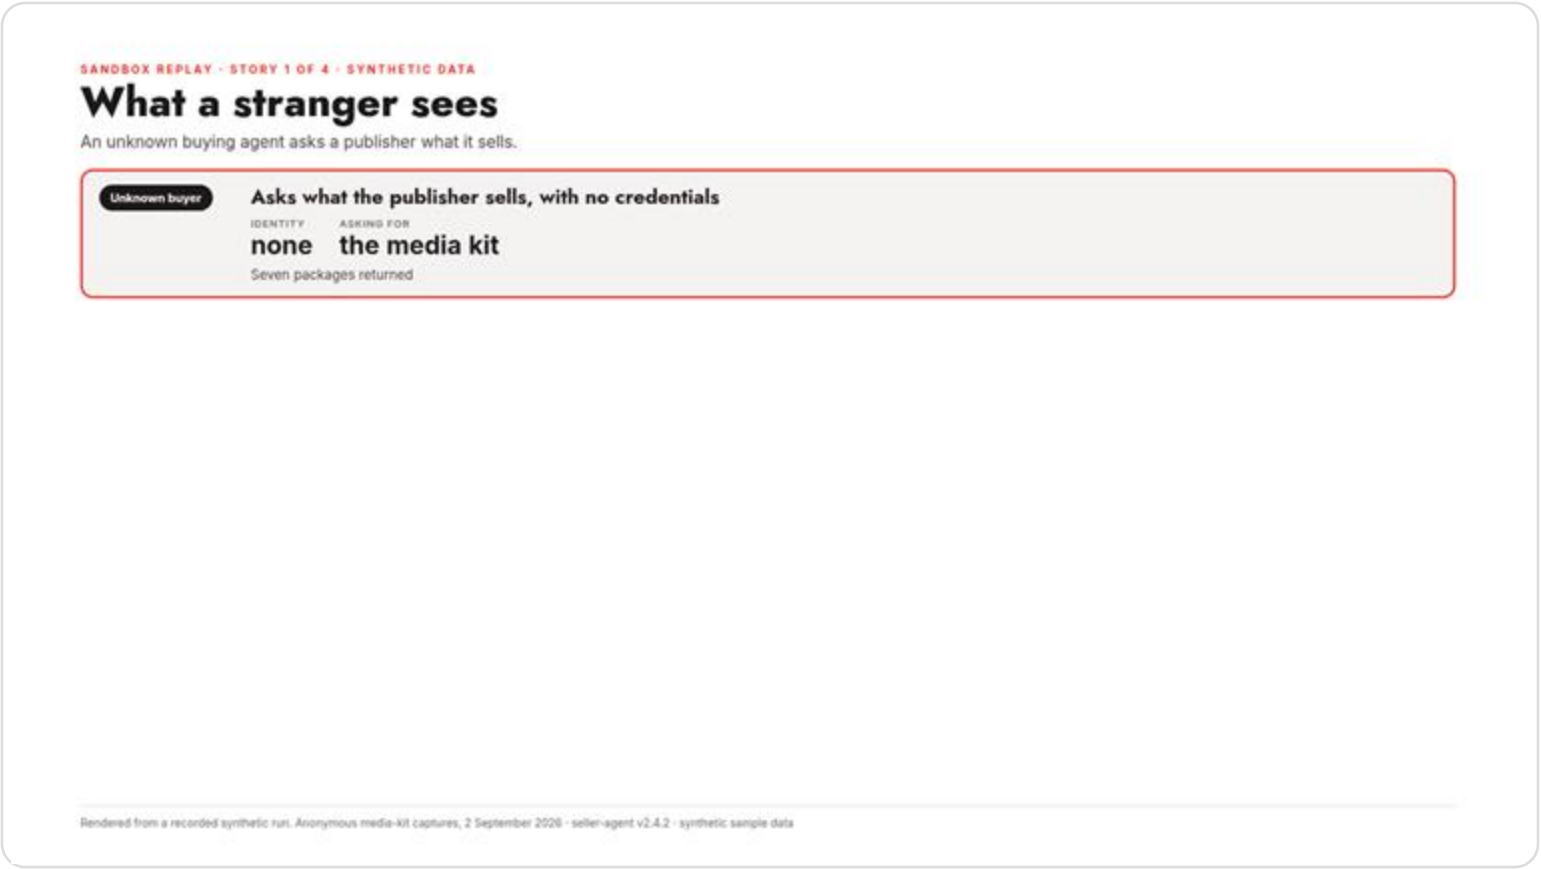
Then I realised it was more like this…

Seven wrong calls in six days.

One cause: not double checking.



What a stranger sees

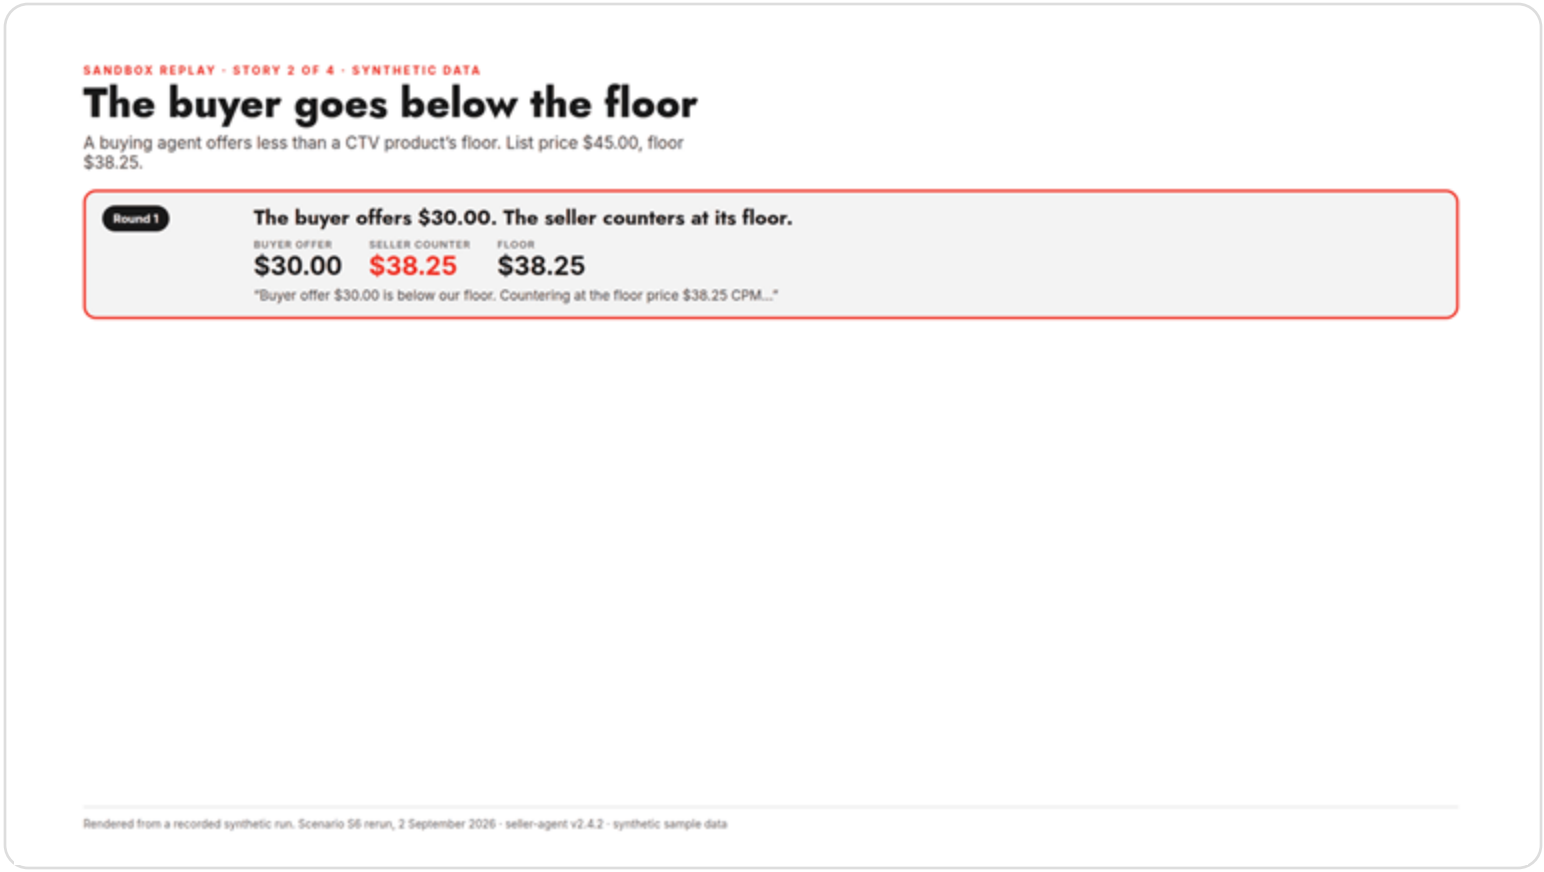


An unknown buying agent asks what a publisher sells. It gets packages and price ranges.

Exact prices and floors appear zero times, across all six package routes.

Anonymous media-kit captures · 2 September 2026 · seller-agent v2.4.2 · synthetic sample data

The buyer goes below the floor

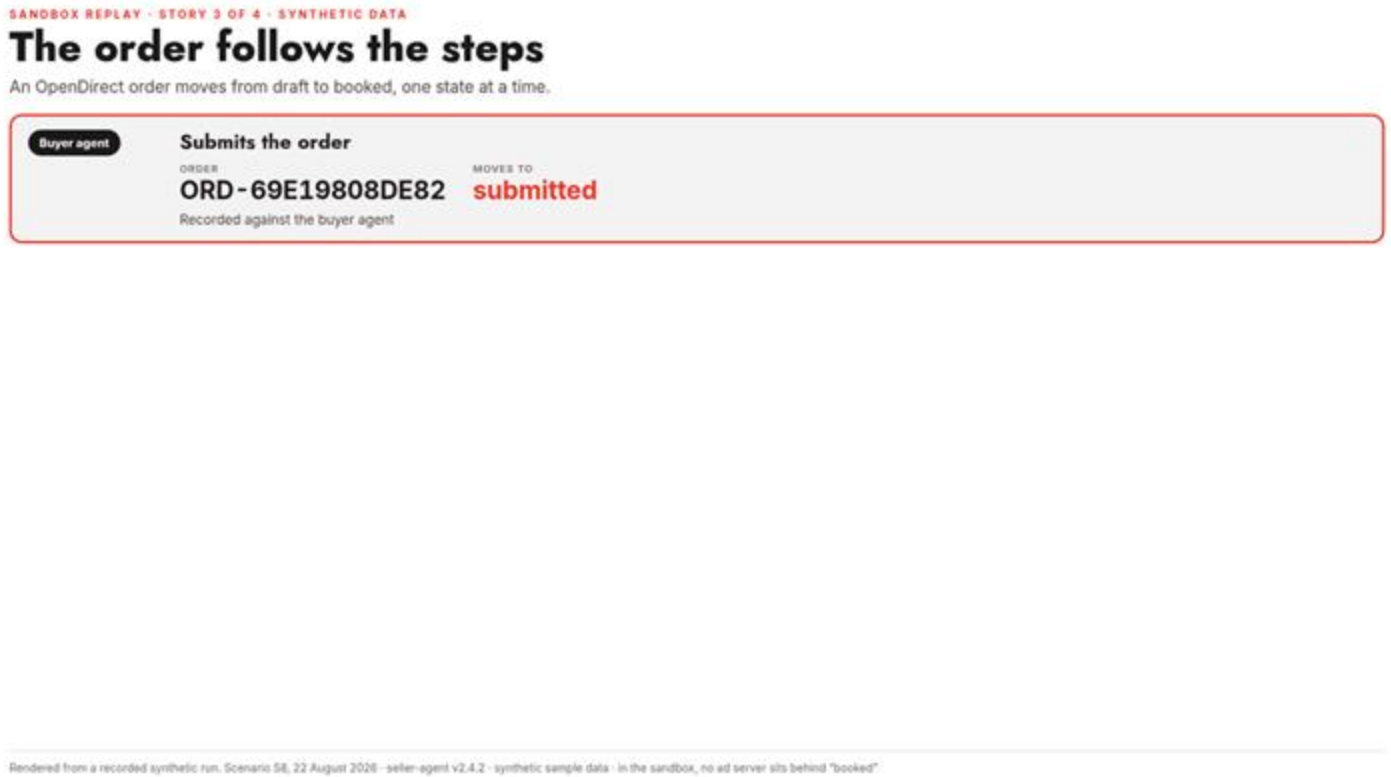


The buying agent offers $30 against a $38.25 floor. The seller agent counters at its floor, three rounds running, until the buyer meets it.

The deal closes at the floor, never below.

Scenario S6 rerun · 2 September 2026 · seller-agent v2.4.2 · synthetic sample data

The order follows the steps



Draft, submitted, approved, in progress, syncing, booked: every step attributed and on the audit trail.

When we attempted to jump from draft straight to booked, it was refused.

Scenario S8 · OpenDirect order · 22 August 2026 · seller-agent v2.4.2 · synthetic sample data

The receipt



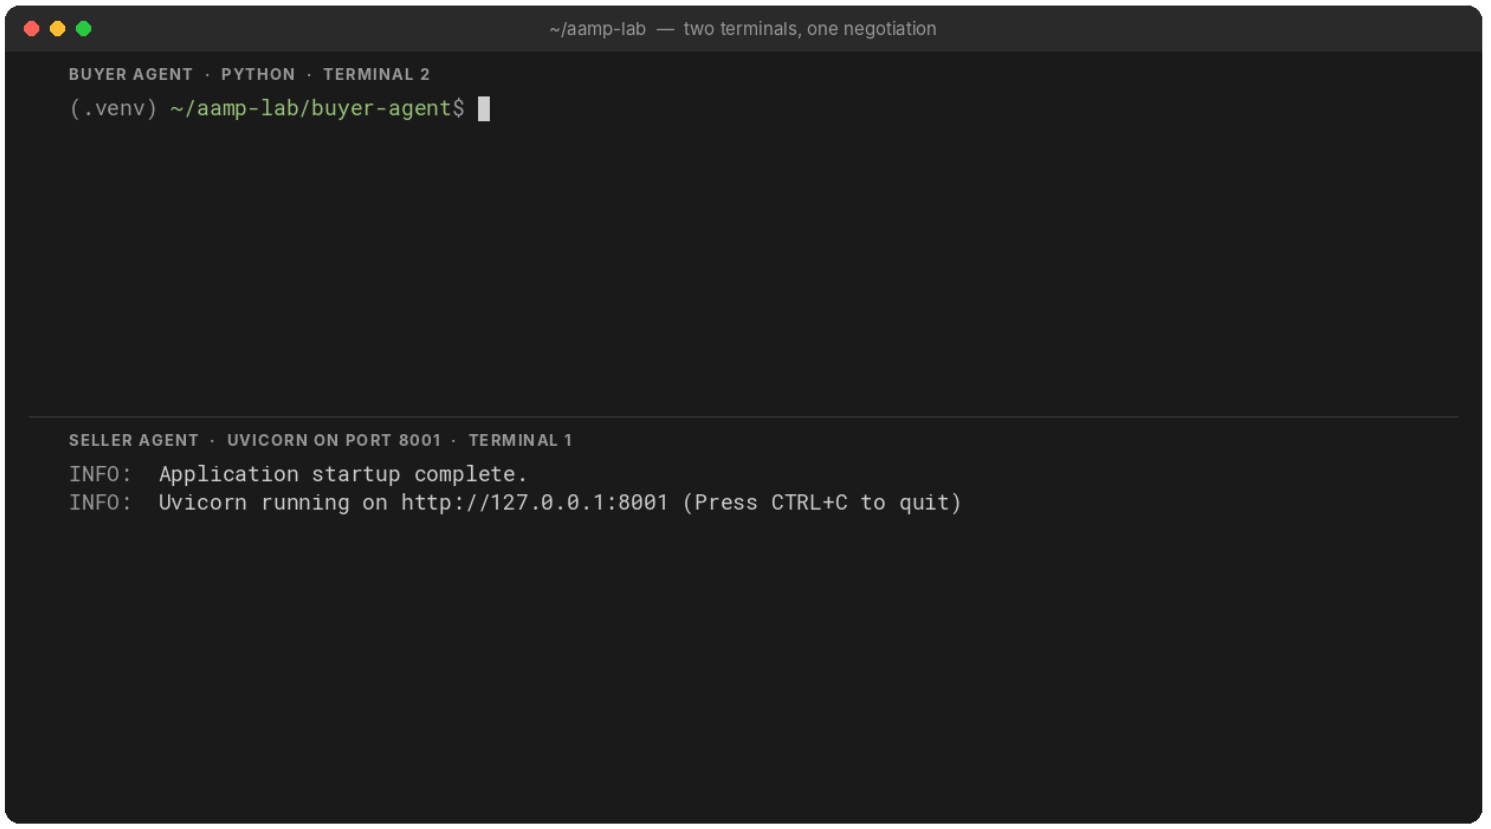
The publisher's display rate was set to $99. The quote came back at $45, taken from the product catalogue. We filed it with Tech Lab on 22 August.

They merged a fix on 15 September, and we have since verified that the quote now follows the rate card.

Rendered from the values filed in seller-agent issue #69 · pinned v2.4.2 · fix merged on main 15 September 2026 (PR #78), verified in two environments, not yet in a tagged release

WATCH IT RUN · ON MY SCREEN · SYNTHETIC DATA

What a test run looks like on my screen



1. The buyer agent opens at $30 CPM. CPM. The seller's floor is $38.25.
2. The seller counters at its floor and holds. Its replies come from its pricing pricing rules.
3. The buyer meets the floor. The agreed at $38.25.
4. The proposal still waits for a approve it.

Re-run 11 October 2026, replayed at reading speed (the real (the real exchange took about a second) · seller-agent v2.4.2 agent v2.4.2 · buyer-agent 8a8dc2a · pricing from the seller's the seller's rules engine, no AI model key · synthetic sample synthetic sample data · output trimmed

WHAT BABY JJ LEARNED

What using these tools taught me

1 Confidence isn’ t correctness

Our seller agent gave confident answers with its AI switched off, and nothing told us.

2 Check the source exists

AI reviewers cited documents and news outlets that don’t exist. Now we check a source is real before reading what it says.

3 The AI didn’ t set the price

Negotiation outcomes came out the same with and without the model. Rules set the price; the model added analysis, and took twice as long.

4 Treat uploads as shared

Anything you give an AI tool can travel further than you think. Keep secrets and client data out.

How I work now

1. Pin exact versions and change them only at planned checkpoints.
2. Reproduce it in at least two separate set-ups before saying it in public.
3. Label what we saw, what we inferred, and what we haven't tested.
4. Check everything before declaring something in public. Always assume it is wrong!
5. One problem per report, reproducible by anyone. Mistakes stay on the record, corrected in the open.
6. AI Review Harness setup to meaningfully support checks and test all reasoning
7. Log everything, structure folders and enable approved access by both humans & machines.

None of it is exotic. It is what you would want from anyone telling you that a technology simply works.

BEFORE THE HARNESS

The old way cost time and trust

Tab-hopping

One draft pasted into several chat apps, replies reconciled by hand.

Invented sources

Reviewers with web access cited files, links and tests that didn't exist.

Leaky uploads

Uploaded files came back as links carrying live access tokens.

No record

Which model, which settings, how long: none of it recorded.

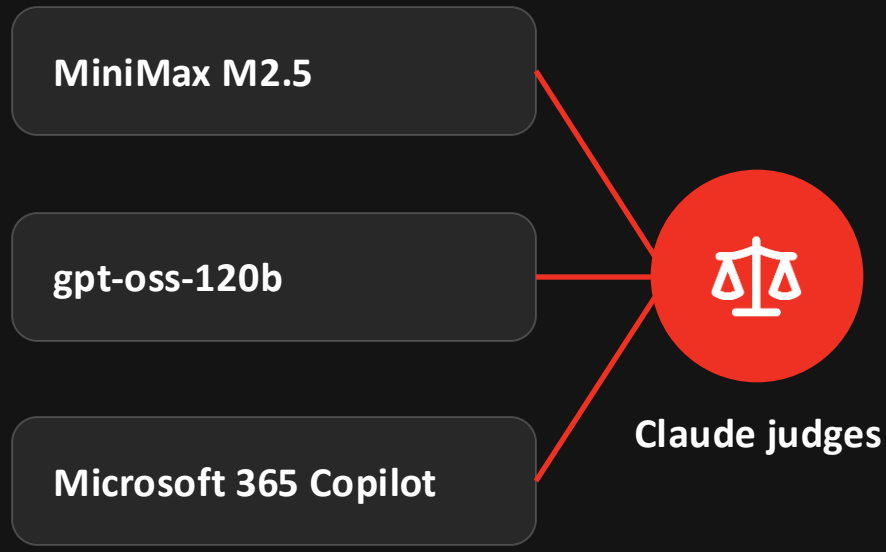
Every round, outside review caught something we'd missed. Every round, at least one reviewer also made a false or invented claim.

IAB reviewer records, August to October 2026

IAB AUSTRALIA · AI REVIEW HARNESS

Outside review, without the tab-hopping

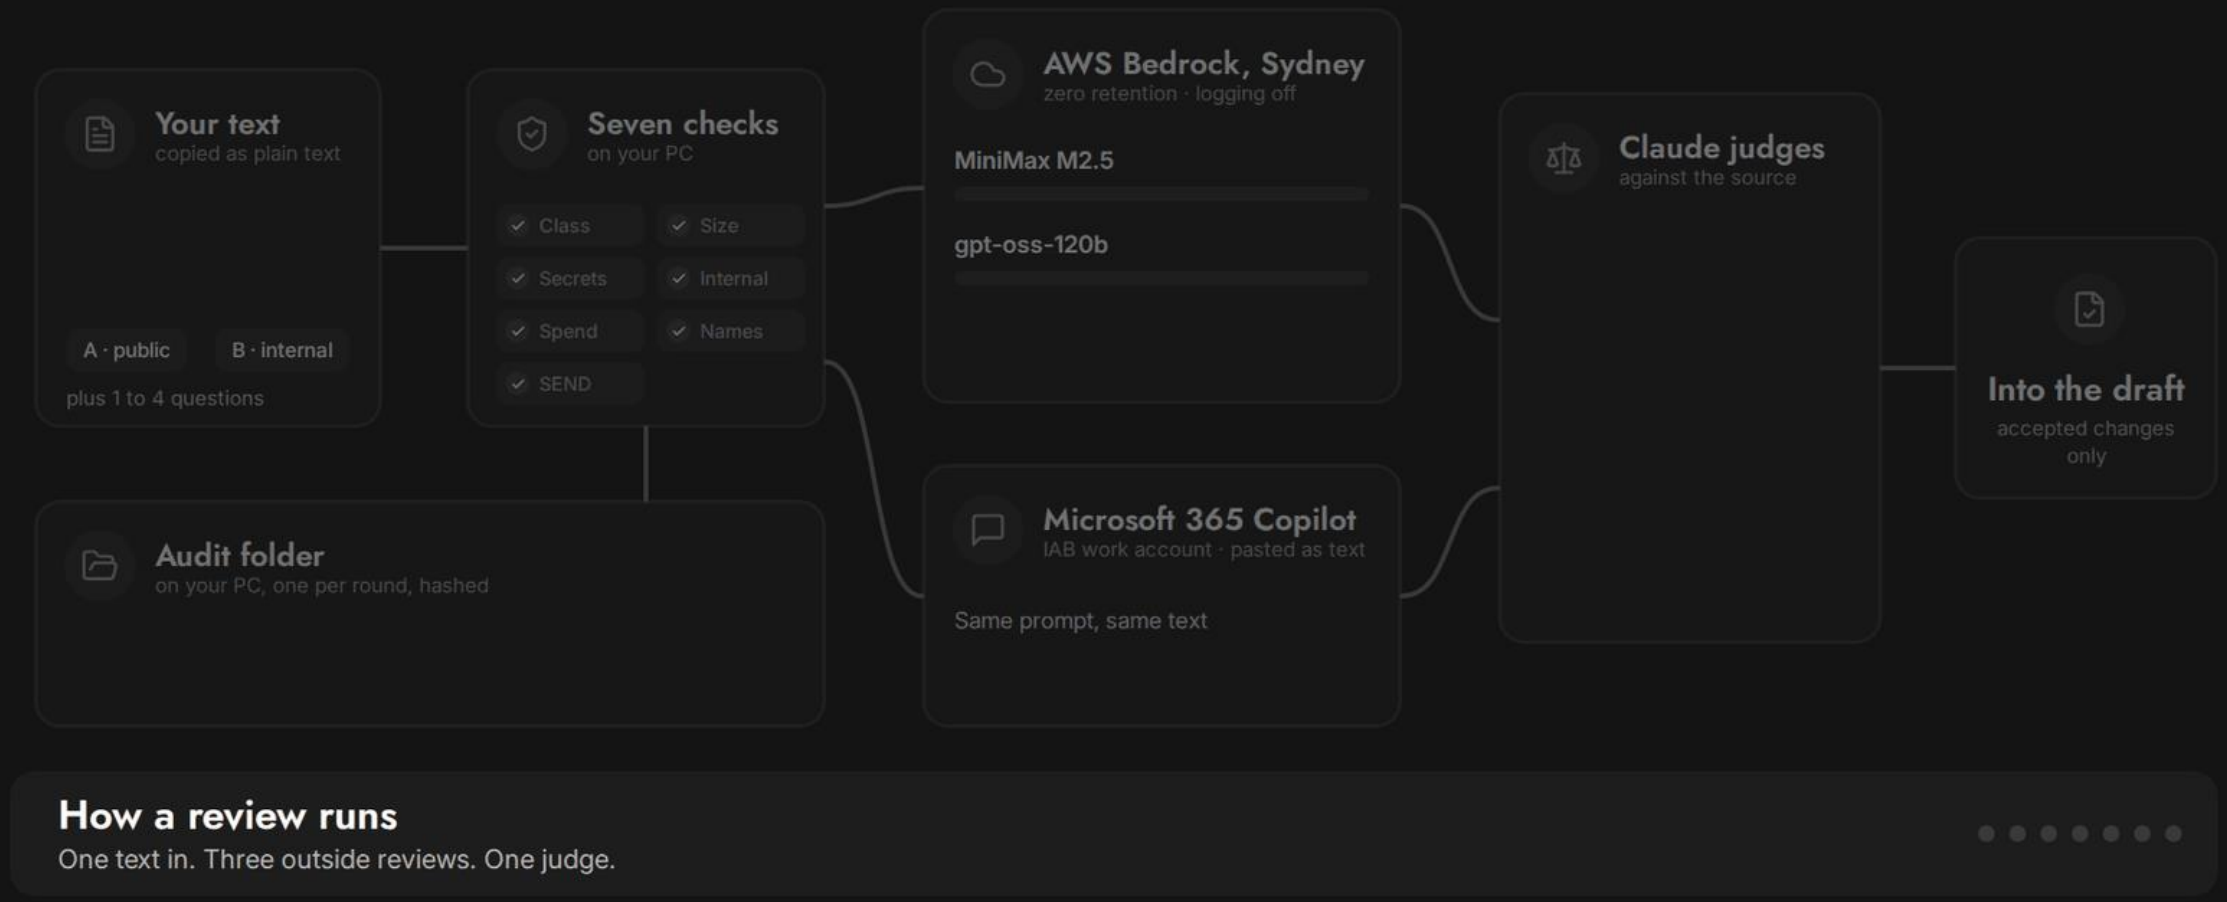
Three outside AI reviewers, one judge, and a record of every round.



Jonas Jaanimagi · Technology Lead · October 2026

HOW A REVIEW RUNS

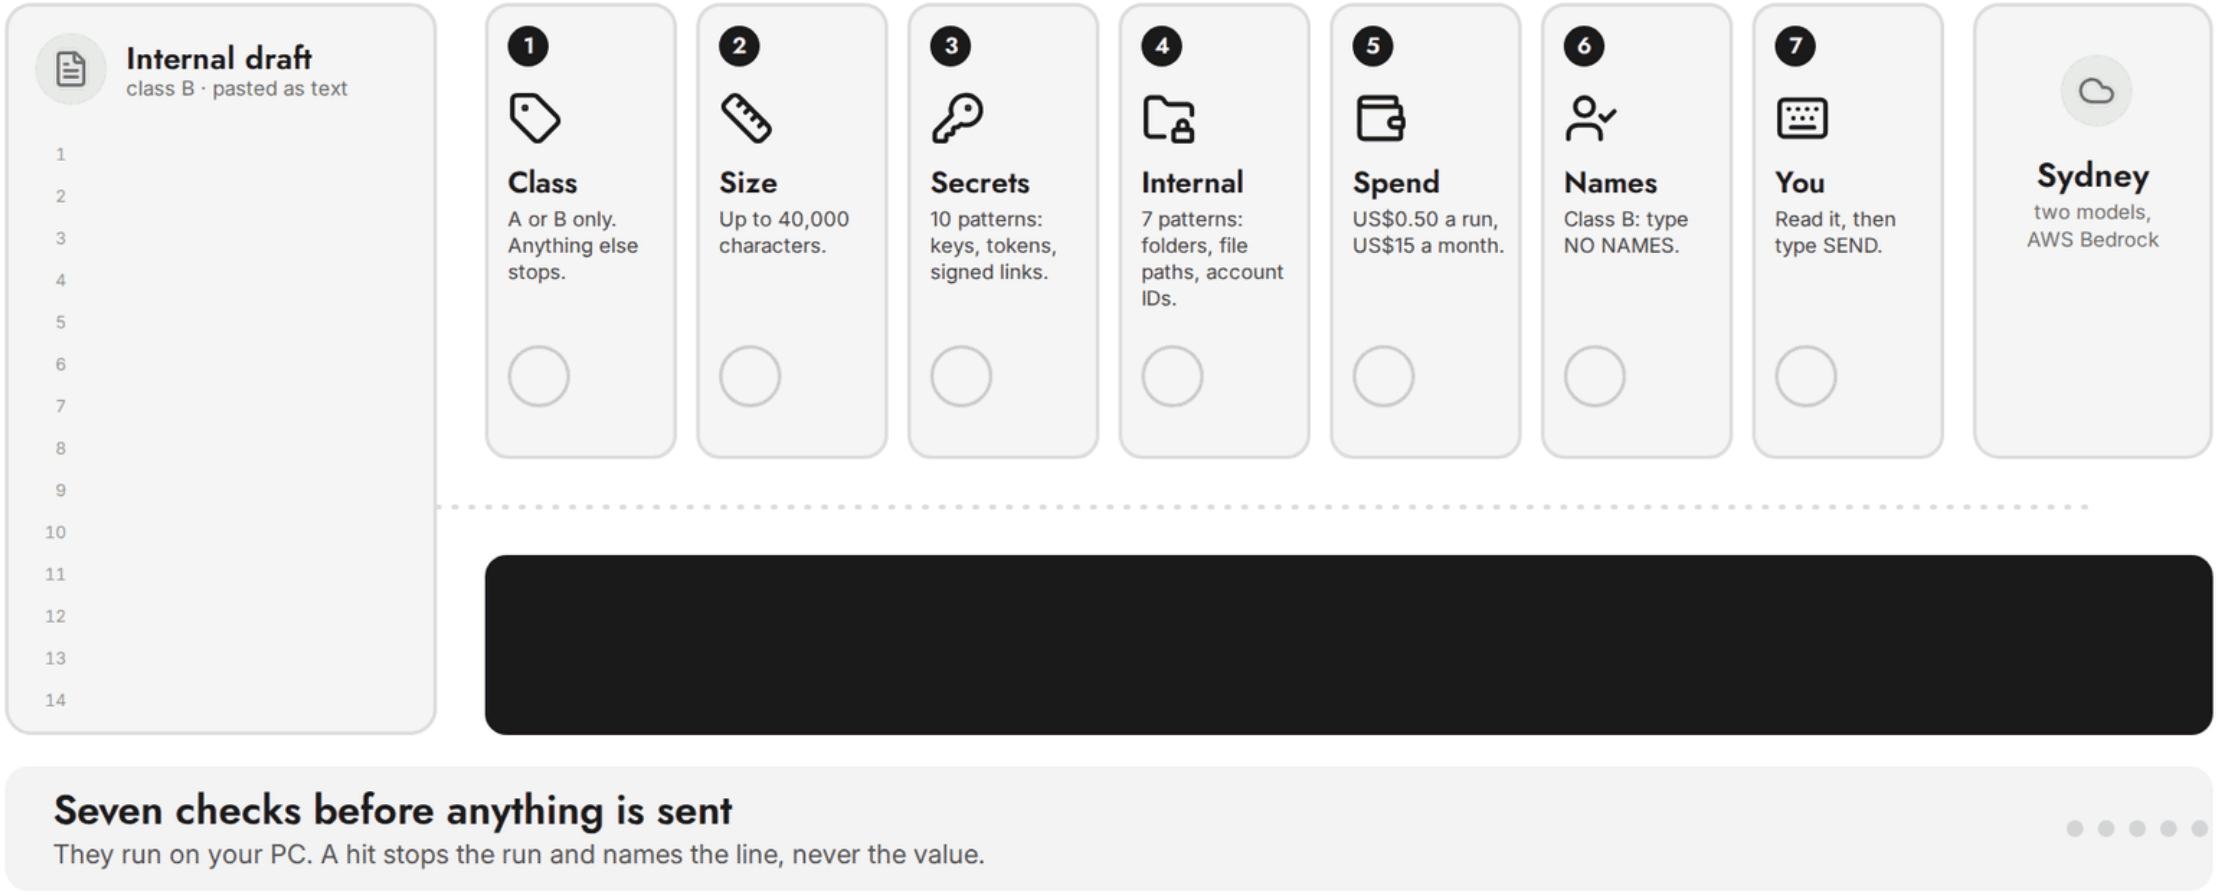
One text in, judged findings out



Animation of the workflow. Timings and cost are from the public test round on 5 October 2026.

BEFORE ANYTHING IS SENT

Seven checks on your PC



Animation. Patterns, limits and messages as in harness v1.3; the draft shown is illustrative.

IN AWS

Locked down, and proven

Sydney only

Both models run in-Region in AWS Sydney. No global routing.

Two models, one action

The sign-in role can only invoke two named models. Sign-ins last four hours.

Nothing kept in AWS

Data retention is set to none and invocation logging is off.

A spend alarm

A US$20 monthly budget alert, set above the script's own US$15 stop.

Proof, not assumption

Three controls after the permissions were cut, 5 October 2026

Allowed model 1	`pong`	✓
Allowed model 2	`pong`	✓
Removed model	`denied`	✓

The refusal is the control that counts. It proves the cut-down permissions took effect.

Set-up and control results from the harness decision record, 30 September to 5 October 2026

WHO MAY SEE WHAT

Six classes of material

Class	AWS harness	Microsoft 365 Copilot
A Public text	**Yes**	**Yes**
B Internal drafts, nothing personal or member-confidential	**Yes, with a typed NO NAMES**	**Yes**
C Member-facing drafts before release	**Needs a decision**	**Needs a decision**
D Personal information	**Never**	**Never**
E Member commercial information	**Never**	**Never**
F Secrets and evidence	**Never**	**Never**



Pasted text only. Never file uploads.



Unpublished IAB drafts are B or C, never A.



Approved by IAB's head of policy, October 2026.

Claude, which drafts and adjudicates, works under its own approval.

THE JUDGE

Two reviewers agreeing isn't proof

Verified — Checked against the source, and it holds.

Valid, minor — True, but a small point.

Rejected — Wrong, or already handled in the text.

Excerpt-limited — Can't be judged from the excerpt that was sent.

Logged for every point: who raised it, whether it was a false alarm, and the minutes spent checking it.

3 of 6

reviewers in one round repeated the same stale claim.

So agreement only sets the order of checking. A point only one model raised still gets checked, and stays in if it holds.

IAB reviewer records, August to October 2026

THE TEST

Five planted errors, three reviewers

We planted five errors in a public text and wrote the answer key before anyone reviewed it.

Planted error	Copilot	MiniMax M2.5	gpt-oss-120b
Commit ID doesn't match	Caught	Caught	Missed
Impressions out by a factor of ten	Caught	Caught	Caught
Claim contradicts its own caveat	Caught	Caught	Caught
Untested paths not flagged	Caught	Caught	Caught
“Three” environments, two listed	Caught	Missed	Missed
Score	**5 of 5**	**4 of 5**	**3 of 5**

5 of 5

Copilot and the harness together, on this run

4 of 5

together on the first run, before the number-check rule

9 of 9

together across two controlled tests

Single runs, 5 October 2026; outputs vary between runs. Copilot ran in Auto mode.

SPEED AND COST

Fast, cheap and capped

≈20 s

a typical round, both models

US$0.0037

the whole 5 October round

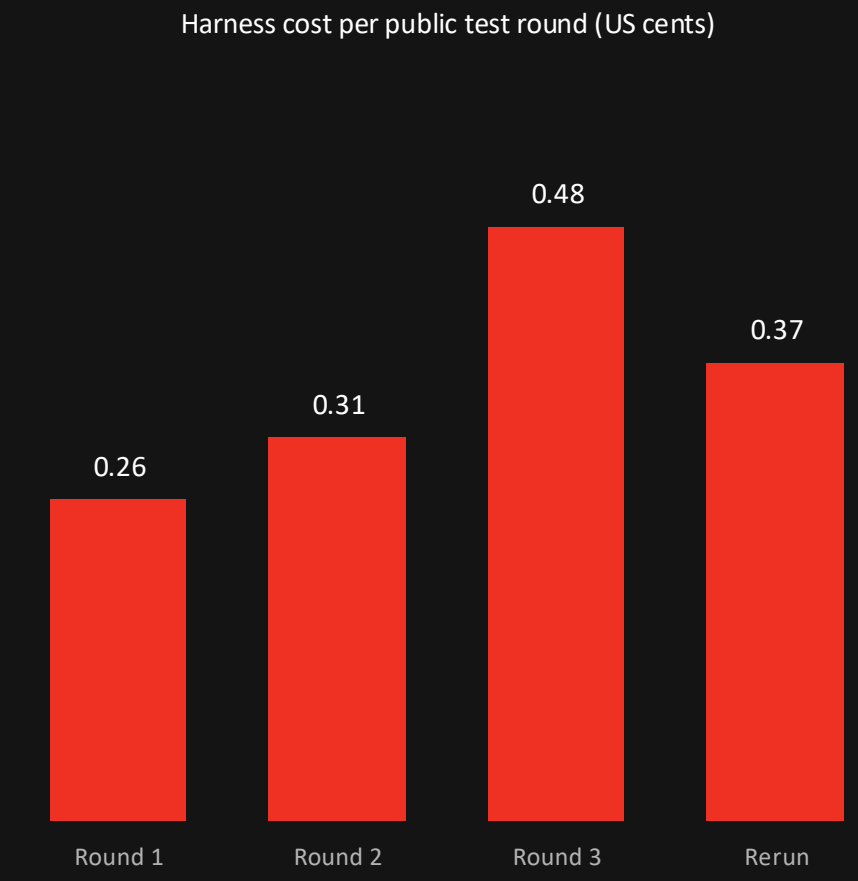
US$15

monthly stop in the script, under a US$20 AWS alert

5 weeks

measured trial, 6 October to 9 November

The trial tracks reviews run, findings accepted, false alarms and minutes spent checking.



Harness estimates from each round's ledger. Rounds 1 and 2 used the first model pair; round 3 and the rerun used today's.

SOME LIMITATIONS REMAIN

What it doesn't do

Catch everything

Both Sydney models missed the same simple miscount, on both runs.

Give the same answer twice

Single runs, and outputs vary. Copilot on Auto picks its own model.

Check names for you

NO NAMES is a person's confirmation, not an automatic control.

Replace the judge

Reviews are inputs, not evidence. Claude checks every point.

MEASUREMENT · WHAT WE CAN SAY TODAY

Four questions everyone asks.

Measurement in agent-led buying: the short answer, where it stands, and what to ask next.

	01 Are we reliant on vendors?	02 Will AAMP 3.0 help?	03 Where is it being addressed?	04 Can we trust what an agent reports?
SHORT	**Yes, for independent That' s their job.**	**On the plan, yes. On the numbers, not by itself.**	**Several bodies, each on one one piece.**	**Only if each number says where it came from.**
WHERE IT STANDS	Independent measurement still matters with agents. What changes: changes: the buy itself must say whose count settles it, whether the the seller's, an independent vendor's vendor's or the buyer's. At least one least one agent protocol already carries it as a field.	OpenProposal covers brief to commitment, so objectives and commitments can become a record to record to measure against. Reporting can then flow back through through existing IAB reporting standards.	**Contracts:** IAB Measurement Services Addendum v1.0 **AI in measurement:** MRC, standards phased to early 2027 **Outcomes:** Tech Lab's ECAPI **The trail:** AdCP audit logs, experimental	An agent can act on a report the moment it arrives, so checks come before the buy, not only Each figure should name its say whether it is final, and flag a placeholder as one.
ASK NEXT	**Whose count settles the Is it in the brief?**	**Does the proposal say how it be measured?**	**Which of these does our next contract cite?**	**What does a success response response mean when nothing is is behind it?**

Sources, checked 11 October 2026: IAB Tech Lab, AAMP 3.0 (OpenProposal) released for public comment, 22 Sep 2026 · IAB, Measurement Services Addendum v1.0 released for public comment, 22 Sep 2026 · MRC, Guidance on Measurement, 8 Jul 2026 · IAB Tech Lab 2026 roadmap · AdCP 3.2 documentation (get_media_buy_delivery; get_plan_audit_logs, experimental).

Write the measurement terms into the brief before the agent runs.

Whose count settles it · what counts as working · what the trail must show.

COMMENTS CLOSE 22 OCTOBER 2026

AAMP 3.0 · IAB Measurement Services Addendum

Three ways in, from October

For sellers · six places

Full track

Stand up a seller agent in a private sandbox and run the scenario suite against named inventory.

For sellers

Readiness-First

Get your catalogue, pricing rules and approval policy in shape, with facilitated demonstrations.

For buyers and vendors

Observer

Agencies, DSPs and verification and measurement vendors see the work. Nothing to deploy.

✓ Sandbox only

✓ Synthetic data

✓ No live spend

✓ No fee

6 October
Kickoff

To early December
Readiness work

Summer
Scheduled pause

From late January
Sandbox transaction scenarios

Week 14
First cohort exits

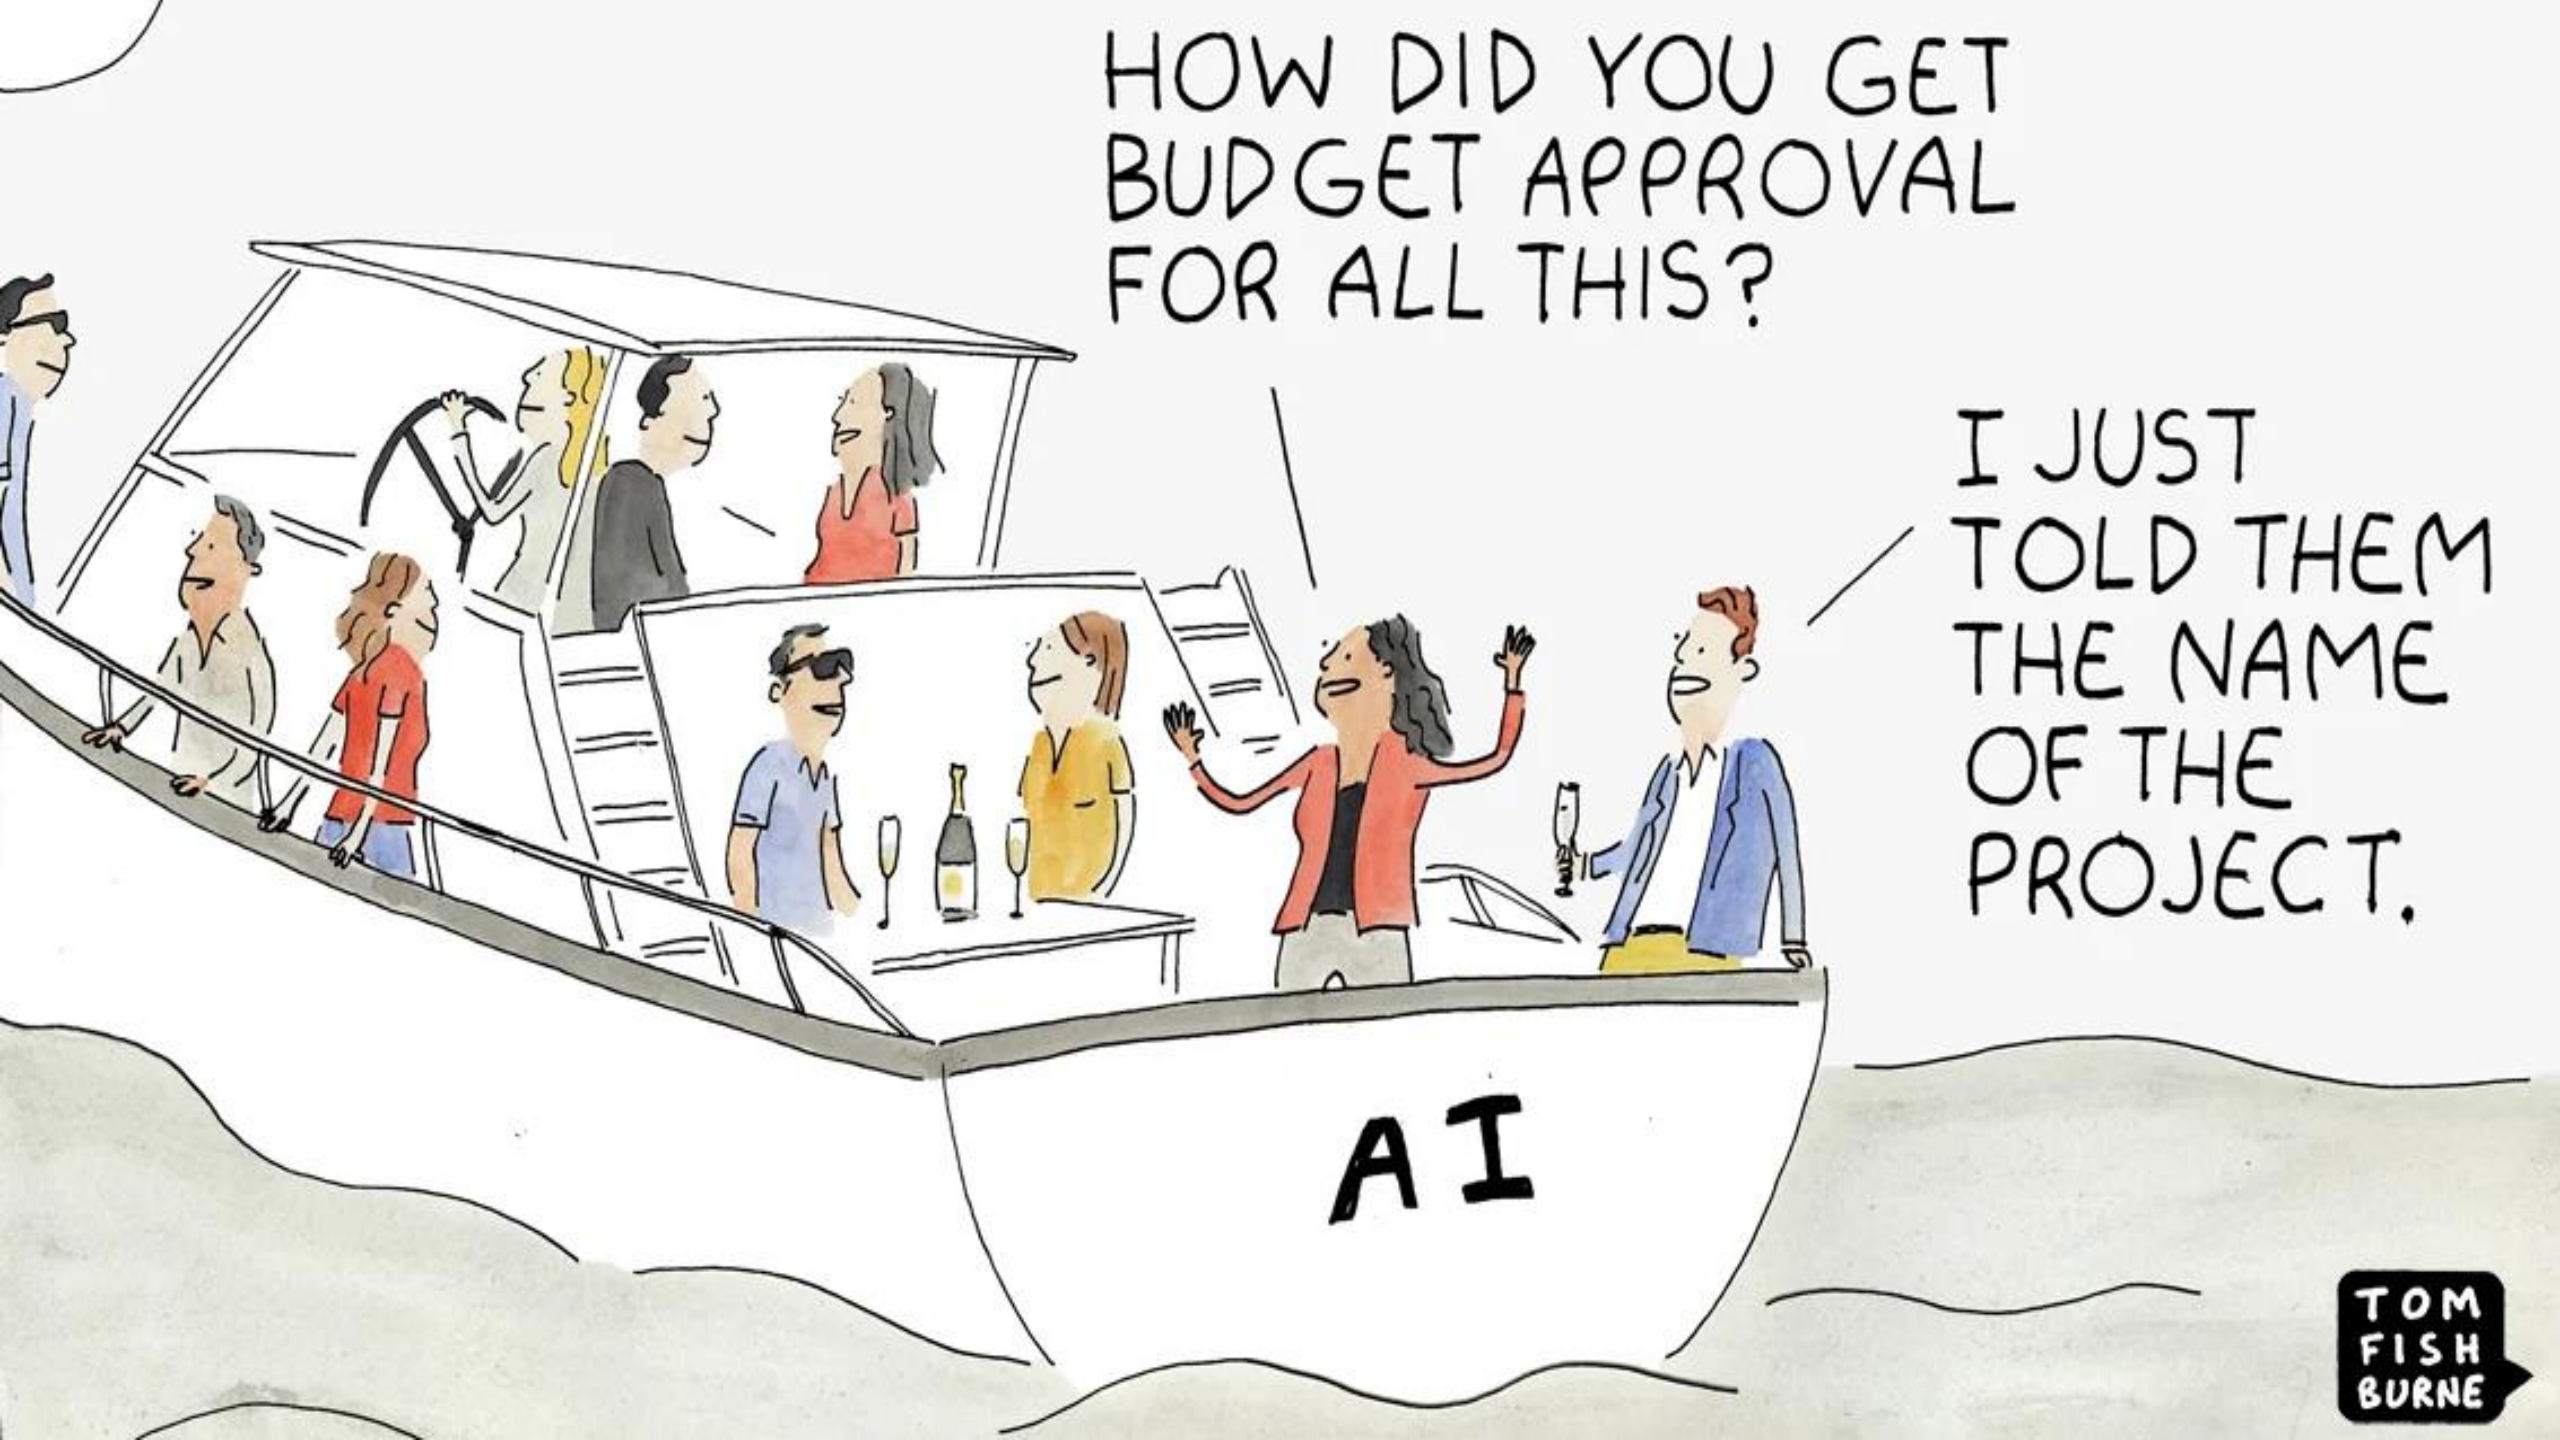
BEFORE YOU GO

Four questions every seller should answer

1. Whatever you sell, is made available in a format machines can easily read.
2. Who may see which parts of it.
3. Where your floor sits, and when the answer is a flat no.
4. Who signs off, and above what number.

That work pays off whatever happens next.

Expressions of interest: agentic@iabaustralia.com.au


HOW DID YOU GET BUDGET APPROVAL FOR ALL THIS?
I JUST TOLD THEM THE NAME OF THE PROJECT.
AI
TOM FISH BURNE

Bots Directory

GoogleBot

Googlebot is the search engine crawler for Google Search.

Google

Training | Search | Direct access

Verified

Claude-User

Claude-User supports Claude AI users. When individuals ask questions to Claude, it may access...

Anthropic

Agent | Intermediary access

Verified

Meta-ExternalAgent

Use cases such as training AI models or improving products by indexing content directly.

Meta

Training | Direct access

Verified

GPTBot

GPTBot is used to crawl content that may be used in training OpenAI's generative AI foundation models

OpenAI

Training | Direct access

Verified

BingBot

Bingbot crawler and handles most of Bing's crawling needs each day.

Microsoft

Training | Search | Direct access

Verified

Google AdsBot

Google AdsBot is special-case crawler that monitors websites where Google Ads are served.

Google

Ads Verification | Direct access

Verified

Five visitors at your front door

Stop asking: ***is this a crawler?*** Rather ask: ***what happens to your content after the visit?***

A **Discovery** *the librarian*

Catalogues you and sends readers back. The only category that reliably returns traffic.

`Googlebot · bingbot · OAI-SearchBot`

B **Training** *the student who never returns the book*

Absorbs your content into the model. Nothing comes back by default, and it cannot be un-crawled.

`GPTBot · ClaudeBot · CCBot`

C **Live agents** *the personal shopper*

A real human is waiting right now. No pageview, no ad impression.

`ChatGPT-User · Claude-User · Google-Agent`

D **Operational** *the meter readers*

Ads verification, ads.txt, link previews, uptime. Rarely glamorous. Never block the verified ones.

`AdsBot · DV / IAS · previews`

E **Impostors** *the gatecrashers*

User-agent strings are voluntary text. Verify by network (ASN, reverse DNS), not simply by the name on the badge.

`spoofed UAs · python-requests`

CRAWLER JOB	WHAT YOU GAIN IF ALLOWED	WHAT YOU RISK IF ALLOWED	WHAT BLOCKING COSTS YOU	PUBLISHER TREND	BRAND TREND
A Discovery & search `Googlebot · bingbot · OAI-SearchBot`	Referral traffic, citations in AI answers, brand visibility	The same crawl can also feed AI answer features built on the index	De-indexing; loss of organic search traffic	**Allow**	**Allow**
B AI training `GPTBot · ClaudeBot · CCBot`	Presence and familiarity inside future models; goodwill with AI operators	Uncompensated extraction; forgone licensing value; irreversible once crawled	Not learned by future models results in invisible loss & no clear signal	**Block, or license**	**Selective (product facts in, premium IP out)**
C Live AI agents `ChatGPT-User · Claude-User · Google-Agent`	A waiting human gets served; task completion; live-answer citations	No pageview or ad impression; paywall probing	Invisible at the moment of decision for assistant users	**Conditional; restrict on ad-funded pages**	**Allow on product & knowledge pages**
D Operational `AdsBot-Google · verification · previews`	Ad delivery, verification, measurement, uptime, link previews	Negligible; some competitive visibility via SEO tools	Broken ad stack; programmatic revenue collapse; ugly link shares	**Always allow (verified)**	**Always allow (verified)**
E Unknown / unverified `spoofed UAs · python-requests · proxies`	Nothing knowable	Scraping, fraud, reconnaissance everything	Nothing	**Challenge or block; investigate volume**	**Challenge or block**

1 · Allow
Value returned justifies unrestricted access.

2 · Allow with conditions
Rate limits, path scope, citation or attribution terms.

3 · Require licensing
Access only under a commercial agreement.

4 · Block
No value, real cost, or unverifiable identity.

Decide per crawler group, then per crawler within high-risk groups. Document each decision with a rationale and a review date. (Framework: IAB Tech Lab Bot & Crawler Management Guidance.)



- **Search:** any behavior that collects or indexes your content, so it can answer questions about it later. The key is that Search is proactively building a database of your site to later respond to queries with. Site owners should expect to get referral traffic or other equitable compensation as a result.
- **Agent:** automated behavior that is acting, usually in real time, on a person's behalf, to get something done right now. This includes chat fetch bots (e.g., ChatGPT-User) and browser-use agents (e.g., Gemini or Claude driving Chrome). The key is that it visits your web application in order to complete a job, and often there's a human waiting on the other end.
- **Training**: a crawler taking your content to train or fine-tune a model. The key is that your data is permanently absorbed into the underlying architecture of the AI to improve its capabilities.

source: blog.cloudflare.com

Visibility first… & then write and manage your own house rules

The problem isn’t crawling. It’s invisibility, so content owners who blindly block may often block crucial crawlers that could drive genuine referrals.

1

SEE

Audit 30 days of bot logs

2

EVALUATE

Value returned, per crawler

3

DECIDE

Allow · conditions · license · block

4

DOCUMENT

Rationale + review date

5

PUBLISH

Policy + licensing contact

Three potential starting positions

Open for discovery

Brands · commerce · government

Allow search + agents. Selective on training. Measure AI referrals.

Protect & monetise

Publishers · premium IP

Keep discovery. Block training + unknowns. Restrict agents on ad pages.

License-first

Data & premium content

Closed by default. Licensed access only. Review IAB Tech Lab’s CoMP for guidance.