

EdgePointe Weekly Research Update — Monday, 2026-09-07

Prepared for: Ed Smith · EdgePointe IT Solutions **Cadence:** weekly deep pass (the daily Game-Changer Watch runs beneath this) **Window:** 2026-08-31 → 2026-09-07 **Register:** all findings landed in `Opportunity-Register.md` (updated same date)

The one-paragraph read

Two of this board's oldest hand-rolled projects became vendor features in the same 48 hours, and a paper published on the last day of the window says our design-critic rubric may be scoring itself. On **Sept 3** Anthropic shipped `ant apply` — agents, environments, **skills**, memory stores and deployments declared in files and applied from a lockfile ([Anthropic, Sept 3](#)). On **Sept 4** Claude Code shipped `/skill-doctor`, which reports **unused loaded skills and their context cost** ([Claude Code v2.1.261](#)). Between them they answer the buildable half of **#57** (the repo-level deny rule, act-now for **nine consecutive weeklies**) and the whole of **#5's** skill-trim, across a ~50-skill portfolio nobody has audited. **Neither is a project any more; both are commands.** Meanwhile [arXiv 2609.02942](#), submitted **Aug 31**, finds that **classifiers trained only on rubric text — with no access to any response — predict judge outputs at nontrivial accuracy**, and that judges often fail to change their verdict when the criterion is *reversed*. That is the sixth consecutive amendment to **#1/#3** and the first that questions whether the instrument measures anything. On our own stack: **Browser Run's `/crawl` now enforces the Content Signals use directive** (Aug 31) — a hard gate landing directly on `edgepointe-hero-cloner`; **Workers bundles went to a flat 64 MiB uncompressed on every plan** (Sept 4); and **Cloudflare's Pay Per Crawl has been rebranded Pay Per Use with named launch partners**, paying when content is *used in an answer* rather than when a bot fetches it. **#32 is now 8 days out and the five-minute check has still never been run — third consecutive week at the top of the research stack.** Correcting the 09-04 daily on one point: **Webflow Source is a limited research preview enrolling select partners, not a GA product** — the urgency is real, the calendar is softer than the daily implied.

0. Method finding: the 09-04 header promised a row that does not exist

The 09-04 daily header states **"THREE NET-NEW ROWS (#85-#87)."** The register contains **#85** and **#86** as numbered rows. **#87 — the Claude Fable 5.1 breaking changes — was written into the risks table as an unnumbered entry and never given a row.** A grep for `^| 87 |` returns nothing; the string `#87` appears exactly once in the file, in the header that promises it.

Why this matters and why it is not pedantry: row numbers are **stable for life** and are how every later report refers back. A number announced in prose but absent from the table is a dangling reference — the next run that writes "see #87" points at nothing, and the run after that may reuse 87 for something else, which silently rewrites history. This is the same class of failure as the 08-24 scan-list drift and the 08-31 surface error: **the process was followed in narrative and skipped in the artifact.**

The rule: a row number is created by writing the row, not by announcing it. Before finishing, a run greps its own claimed new numbers and confirms each returns a table row. **This run added that check and fixed #87. Impact □ / Effort S. → #87 created as a proper row this run.**

1. Claude → elite websites (models, Code, Agent SDK, skills, MCP)

□ `ant apply` turns agents, environments, skills, memory stores and deployments into declared files with a lockfile. This is the shape #57 has been asking for, and it arrived as a CLI command.

What it is: **ant CLI v1.30.0, Sept 3**, adds `ant apply`, which "creates and updates agents, environments, skills, memory stores, and deployments from files in your repository." You describe each resource in a file, run `ant apply`, and **approve the plan it prints**. It writes a `claude-lock.json` lockfile that you commit, "so that later runs, on your machine or in CI, update the same resources instead of creating new ones" ([Anthropic release notes, Sept 3](#)).

Why it matters: three separate rows have been circling this for months. **#57** wants a repo-level constraint that survives a session and cannot be forgotten — a plan-then-approve step with a committed lockfile is that constraint, expressed by the vendor. **#49** and **#30** carry the ~50-skill portfolio, which is currently a directory of files with no declared state and no CI reconciliation; `ant apply` makes the portfolio a reviewable diff. And the **"approve the plan it prints"** step is precisely the *AI proposes, a person posts* discipline already

adopted for Google Ads — the same guard rail, now native.

How EdgePointe applies it: one sitting, and it starts small. Declare **two** skills — one Vantage skill and one client-facing skill — as files, run `ant apply`, commit `claude-lock.json`, and confirm a second run updates rather than duplicates. If that holds, the portfolio migrates incrementally rather than in a rewrite. **⚠ Do not migrate all fifty in the first pass** — a lockfile that mis-maps a resource is worse than no lockfile, and the failure mode is silent duplication. **Impact □ / Effort M. → Register #88; substantially reframes #57, #49, #30.**

□ `/skill-doctor` reports unused loaded skills and what they cost in context. #5 has carried "trim the 50+ skill context" as an aspiration since July; this is the instrument.

What it is: **Claude Code v2.1.261**, published **Sept 4**, adds `/skill-doctor` to "show unused loaded skills and their context cost." The same release adds an **Organization policy** line to `/status` and `claude doctor` explaining **why a policy failed to load** (for example, a proxy not passing the endpoint through), and reports helper timeouts as timeouts rather than generic failures ([Claude Code changelog](#), [how to use `/skill-doctor`](#)).

Why it matters: every skill loaded is context spent before a single useful token is generated, and with ~50 skills the tax is invisible and unmeasured. **#5** ("cost-optimize the build loop") has always had this as its second half and has never had a number attached. **⚠ One methodological trap, and it is the whole difference between a measurement and a vibe:** a single quiet run tells you which skills *that run* did not use, which is not the same as which skills are dead. The correct use is **a repeatable canary task** — the same brief, run the same way — so successive `/skill-doctor` results are comparable. One run is an anecdote; the delta across a fixed task is evidence.

How EdgePointe applies it: pick one canary — a Vantage tenant edit, or a small clone-polish pass — run `/skill-doctor` after it, record the list and the cost, and only then decide what to unload. **Impact □ / Effort S. → Register #89; closes the measurement half of #5.**

□ A skill-deletion semantic changed on Aug 27 and it is destructive by default. In **Python SDK 1.2.0 / TypeScript 0.122.0 / Go 1.68.0 / Java 2.59.0 / Ruby 1.67.0 / C# 12.44.0**, `client.beta.files` and `client.beta.skills` stop sending the `files-api-2025-04-14` and `skills-2025-10-02` beta headers and return the non-beta shapes. **With that change, `client.beta.skills.delete()` deletes a Skill together with all of its versions**, and the Messages type `BetaSkill` is renamed **BetaContainersSkill** ([Anthropic, Aug 27](#)).

Why it matters: a call that used to remove one version now removes the lineage. Against a ~50-skill portfolio with any scripted cleanup, that is an irreversible loss wearing the same method name. **Run the grep, do not assert the zero** — the standing rule since 08-22. If `skills.delete` appears anywhere in tooling, pin the SDK or send the beta header until the call site is reviewed. → **added to the risks table.**

□ **Note, and it holds the line on last week's read: #52 stays closed.** Sonnet 5 at \$2/\$10 is confirmed permanent; **September 1 passed and nothing happened.** Fourth confirmation. Do not resurrect the countdown.

MCP: no new dated migration deadline this window. Standing action unchanged — pin SDK versions before touching Ask Vantage's MCP surface.

2. Elite wedding-venue web design

Venue software: quiet for the third consecutive week, and this run says which surfaces it checked. **Perfect Venue**, **Tripleseat** and **HoneyBook** were each checked by name; no dated August-or-September feature release surfaced for any of them, and the results returned were comparison and review content rather than vendor announcements ([Software Advice comparison](#), [EverBridal Tripleseat review](#), [EverBridal HoneyBook review](#)). △ Per the 08-31 rule, that is a **weaker read than a changelog** — these three vendors do not publish a machine-readable product-updates surface comparable to [webflow.com/updates](#) or [duda.co/product-updates](#), so "quiet" here rests on absence-of-news, not absence-from-changelog. Recorded as such.

The consequence is unchanged and it is now the load-bearing claim in the venue pitch: nobody bundles the website with the venue software. The **#18** gap survives another week. It is worth saying what that gap is *not* — it is not a technology gap, since any of the three could add a site builder. It is a **buyer gap**: the venue software is sold to an events manager and the website is bought by an owner, and no vendor has decided to sell to both. That is the sentence, and it is more durable than any feature.

#33's 08-31 re-sequencing holds. The 73% aggregator-concentration figure and the 36% couple-adoption figure are unchanged; nothing this week moves them. **#4 (aggregator NAP consistency) stays ahead of #33's on-site half.**

3. Cloudflare (our stack) — deprecations flagged

Changelog index read directly this run. Window Aug 31 → Sept 7; newest entry is **Sept 4**.

□ **NEW — Browser Run's `/crawl` endpoint now enforces the Content Signals `use` directive. This lands squarely on the cloner and it is a hard 400, not a warning.**

What it is: on **Aug 31**, the `/crawl` endpoint began respecting the **`use` directive of the Content Signals standard**, "*letting site owners express the maximum level at which their content may be used.*" A new `contentUse` parameter takes **`reference`** or **`full`** (default `full`), and **if a target site's `robots.txt` sets a use level more restrictive than your declared `contentUse`, the crawl request is rejected with a 400 error** (Cloudflare, Aug 31).

Why it matters, and it matters in two directions at once. **Operationally:** `edgepointe-hero-cloner` and `edgepointe-clone-polish` exist to capture a prospect's live site. If any capture path routes through Browser Run `/crawl` and the target publishes a restrictive `use` signal, **the capture now fails with a 400 rather than degrading** — a silent workflow break waiting for the first prospect whose agency set the directive. **Commercially and ethically:** this is a machine-readable statement of what a site owner permits, and "*we honour the owner's stated content signal*" is a defensible line in a pitch that opens by cloning somebody's website. The standing rule that prospect sites are untrusted content already applies; this adds that they are also **governed** content.

How EdgePointe applies it: (a) grep the capture pipeline for Browser Run `/crawl` and set **`contentUse`: "reference"** explicitly wherever capture is for analysis and re-skin rather than republication — declaring the narrower level is what keeps requests inside the owner's permission; (b) handle the 400 as a **named, reportable outcome** ("this site's owner has restricted AI use of its content") rather than a generic failure, because that message is itself a credible thing to say to a prospect; (c) record whether the pipeline uses `/crawl` at all — if capture is entirely Chrome-driven, exposure is nil and the row closes in five minutes. **Impact □ / Effort S. → Register #90; updates #44, #68.**

□ **Worker bundles went to a flat 64 MiB uncompressed on every plan, and the compressed ceiling was removed outright.**

What it is: **Sept 4**. Previously Cloudflare checked the **compressed** bundle and rejected deploys over **3 MB (Free) / 10 MB (Paid)**. "*That limit has been removed. Cloudflare now only checks the uncompressed size of your bundle, which is 64 MiB across all plans.*" Check with `wrangler deploy --outdir bundled/ --dry-run` and read `Total Upload` (Cloudflare,

Sept 4).

Why it matters: the render engine, the owner dashboard, the reach tools and the reports module all ship inside one Vantage worker, and **"will this still fit"** has been an unstated tax on every feature decision. It is now roughly a **6× headroom increase on Paid** and the constraint moves from *compressed* (which you cannot easily predict) to *uncompressed* (which `--dry-run` prints). **△ This is headroom, not permission.** Bundle size still costs cold-start and parse time, and a worker that grows because it can is a worker nobody can reason about. *How EdgePointe applies it:* run the `--dry-run` once, **record the current number as a baseline in the deploy cheatsheet**, and stop treating size as a live constraint on feature scoping. **Impact □ / Effort S. → Register #91.**

□ The Images binding learned to rasterize text and to accept uploads without an API token. Both are real Vantage features, not platform trivia.

What it is: **Sept 2.** The Images binding adds `.text()` to "rasterize text with the Images binding" — output a string as its own image or draw it over another, styled with `font`, `size` and `color`, with `cf.image`'s `draw` array now accepting a `text` key. It also adds `.signedUrl()` for server-side signing of a private image, `.createDirectUpload()` to hand a client a Direct Creator Upload link "so that a client can upload an image to your storage", metadata filtering on `.list()`, and a `headers` option on `.response()` (Cloudflare, Sept 2).

Why it matters: two live problems. **First, social/OG cards.** Every Vantage tenant page wants a share image with the venue's name on it, and today that is either hand-made per tenant or absent. `.text()` over a tenant hero, generated at the edge from `theme.json`, is a **per-tenant OG image with no design step** — the kind of thing that looks like polish and costs one function. **Second, owner photo upload.** The standing lesson "*portal upload ≠ real photo — look at the pixels*" came from an upload path that was harder than it should be; `.createDirectUpload()` gives the owner a signed upload link **without exposing an API token to the browser**, which is the correct shape for a portal. *How EdgePointe applies it:* prototype the OG card first — it is self-contained, visible, and demoable in a pitch. **Impact □ / Effort S-M. → Register #92; touches #6 (R2/media) and the portal upload path.**

□ Four smaller items, each of which answers a question rather than opening one.

- **The Next.js WAF rules merge into the originals tomorrow — Sept 8.** The two **beta** rules announced Sept 1 (Image Optimizer RCE via crafted AVIF, and CVE-2026-75604) are scheduled to merge into the existing rules on **2026-09-08 in Log mode** (Cloudflare scheduled WAF release). **This is the follow-through on #76** and it is a reason to make the client Next.js inventory call *this week* rather than *next*: a rule change landing on a live RCE is a natural, non-salesy pretext for

the phone call.

- **A new SQLi detection moved from Log to Block on Sept 1** — "SQLi - WHERE Comparison With WITH Clause" in the Cloudflare Managed Ruleset ([Cloudflare, Sept 1](#)). New detections that go straight to Block are the classic source of a false positive on a legitimate query. No Vantage exposure expected (D1 access is parameterised), but if any client app builds SQL from user input behind Cloudflare, this is the week it might start 403-ing.
- **AI Gateway invoices no longer break out input and output tokens** (Sept 1) — a single total per model, and model names standardised to `provider/model` ([Cloudflare, Sept 1](#)). **⚠ This is a small regression for us specifically. #63** is a three-line voice/AI cost model that depends on separating input from output; the invoice will no longer show that split, so **the breakdown must come from AI Gateway logs, not the bill**. Record it before the model is rebuilt on a number that no longer exists.
- **Python Workers now support WSGI and ASGI frameworks** — Django, Flask, FastAPI, Starlette, via `wsgi.entrypoint` / `asgi.entrypoint` ([Cloudflare, Sept 2](#)). Note only: Vantage is TypeScript and stays TypeScript. Worth knowing the day a client arrives with a small Flask app and a hosting bill.

Deprecation clock:

Item	Date	Exposure
Cloudflare AI-bot defaults	Sept 15 — 8 DAYS	#32 — THIRD consecutive week at the top of the research stack. Scope still unverified. See below
Next.js WAF beta rules merge into originals	Sept 8 — TOMORROW	Follow-through on #76 . Client-side only; Vantage runs no Next.js
Azure Functions Sentinel connector → CCF connector	Sept 14 — 7 days	Enterprise only. No known EdgePointe exposure — record the answer either way
NEW — <code>client.beta.skills.delete()</code> now deletes ALL versions of a Skill	Aug 27 (shipped)	UNVERIFIED — run the grep. Destructive change to a method that kept its name
Legacy KV namespace routes retire	Oct 15	VERIFIED CLEAR — zero

Item	Date	Exposure
		occurrences in repo
Fable 5.1 / Mythos 5.1: <code>tool_choice</code> <code>any/tool</code> return 400	Sept 1 (shipped)	#87 — now a proper row. Also: text watermark, 30-day retention, thinking-block prefix binding
D1 free-tier daily limits enforced	Sept 1 (LIVE)	#81 — the 60-second Workers- plan check is STILL unrecorded
Anthropic Python SDK v1.0 breaking	Aug 20 (shipped)	Low — <code>httpx</code> = 0 in repo
~~Sonnet 5 \$2/\$10 → \$3/\$15~~	~~Sept 1~~	☐ CANCELLED — the date passed, nothing happened. Fourth confirmation. Closed for good

☐ **#32 — 8 days, third consecutive week at the top, and this week the policy itself changed shape.**

Scope is unchanged and re-confirmed: the new defaults apply to **new Cloudflare customers, new sites created by existing customers, and all existing free-plan customers; existing paid customers keep their current settings**; and mixed-use crawlers get blocked **on pages that host ads**, with Search left allowed ([TechCrunch](#), [Technology.org](#)). **Two facts still decide our entire exposure and neither is recorded: plan tier per tenant zone, and ad presence** — and Vantage tenants are venue and small-business marketing sites that do not run ads, which is very likely to make the answer "no exposure." **Likely is not recorded.** It is a dashboard check and it is Ed's.

NEW **What is new this week is the commercial half: Pay Per Crawl has been rebranded Pay Per Use, and the billing trigger moved.** Launch partners are named — **Ceramic.ai** paying publishers when content appears in search results, **You.com** paying when premium content is accessed by an agent — and the model "*pays publishers when AI actually uses their content in an answer, not when a bot fetches the page*" ([Cloudflare Sept 15 deadline coverage](#), [MLQ](#)).

Why that is worth its own row and not a footnote: a venue site earns nothing from a crawl and nothing from an impression — but **"you get paid when an AI uses your content to**

answer a question" is a sentence a venue owner understands immediately, and it sits directly on top of the AEO work in **#33** and **#69**. It is almost certainly not economic at one venue's traffic. **Its value is as a framing device that makes AI visibility feel like an asset rather than an expense**, and it costs nothing to say. Δ Do not promise revenue. **Impact $\square$ / Effort S (a sentence, and a check on whether tenant zones can even opt in).** → **Register #93.**

4. AI landscape (only what passes "so what for EdgePointe")

$\square$ **The sixth consecutive amendment to #1/#3, and the first one that asks whether the instrument measures anything at all.**

What it is: **"Judging LLM-as-a-Judge: Concerning Rubric Artifacts in LLM-based Automated Text Generation Evaluation"** ([arXiv 2609.02942](https://arxiv.org/abs/2609.02942), submitted **Aug 31, 2026**). The finding, stated plainly: **classifiers trained only on rubric text — with no access to any evaluated response — achieve nontrivial predictive performance on judge outputs.** The authors conclude that *"rubric formulations encode recoverable evaluative signals, allowing scores to be partially anticipated independently of model outputs."* Alongside it: **judges often fail to reliably update their decisions when either the candidate response or the rubric criterion is reversed** ([HTML version](#)). A companion line of work, **RuVerBench** ([arXiv 2606.29920](https://arxiv.org/abs/2606.29920)), makes the same point for **agentic** scenarios specifically, where long complex outputs are exactly what our clone-polish loop produces.

Why it matters to us specifically, and it is worse than last week's finding: the 08-31 amendment worried that judges might **agree for shallow reasons**. This says something sharper — **the rubric itself leaks its own answer**. If a score can be predicted from the wording of the criterion before anything is looked at, then the criterion is not measuring the page, it is announcing an expectation. And the reversal result is the killer: a criterion that scores the same when you invert it **is not a criterion**, it is decoration. The full sequence now reads: 08-03 *score the trajectory* → 08-10 *close the checklist* → 08-17 *test the criteria* → 08-24 *cite the evidence* → 08-31 *vary the evidence kind* → **09-07: prove the criterion responds to reality at all.**

How EdgePointe applies it — and this is the smallest possible version of the test: add a **null control** to the design-critic rubric. Run every criterion against **(a) a blank page** and **(b) the criterion inverted**. Any criterion that returns a materially unchanged score in either case is an artifact and gets deleted. That is a delete-list, not a build, and it can be

run against the rubric that exists today. **This is the lab hour — see below. Impact □ / Effort S. → Register #94; amends #1, #3.**

AI moat — the count is now nine. #10 — the Vantage owner-edit feedback loop — is **unbuilt for the ninth consecutive week**. This week's channel reading (§6) says the market is done treating AI as an add-on and is folding it into the platform. When AI is native to every platform, **the only thing that distinguishes one platform's AI from another's is the data it learns from**. #10 is that data. Said again because it keeps being true: **#10 is the moat; everything else is packaging**.

5. Competitive + integrations

□ **Correction to the 09-04 daily, and it is a favourable one: Webflow Source is a limited research preview, not a shipping product.**

What it is: Source "is available as a limited research preview," currently **enrolling select agency partners and customers in a phased research preview**, with sign-up at webflow.com/source ([Webflow Conf 2026 announcements](#), [GlobeNewswire](#), Sept 2).

Campaigns is in beta and "will be available for purchase next month." Agent Presence ships "for all customers later this month."

Why it matters: the 09-04 daily read Campaigns as a dated, imminent competitive event and reordered the stack accordingly. **The substance of that read stands** — Campaigns does describe our capture-and-build workflow, it does wire to Google Ads, and it is dated. **The calendar is softer than the daily implied**. Source is gated behind partner enrolment, which means the *platform* is not arriving in October; one *product* on it is. *How EdgePointe applies it:* **#85** (the "why doesn't Vantage run my ads" paragraph, owed before the **09-11** proposal) is **unchanged and still owed** — the question a prospect asks does not wait for GA. **#86** (the Campaigns-vs-cloner bake-off) can be **honestly timed against a beta**, which also means it can be joined to **#84**'s Duda test as one sitting rather than rushed as two. → **updates #85, #86**.

□ **Duda's Vibe is no longer a website builder. On Aug 31 it shipped built-in user authentication, and on Sept 4 it shipped SEO/AI-search optimization and functional forms — inside the same product.**

What it is: two dated entries from duda.co/product-updates, read directly. **Aug 31** — "Add built-in user authentication to your Duda Vibe projects by simply asking for it" ([Duda](#)). **Sept 4** — "Generate fully functional forms, and optimize your project for SEO & AI search

- *right within Duda Vibe*" ([Duda](#)).

Why it matters: **authentication is the line between a website and an application.** A white-label builder that generates a site *with a login* is generating something with the shape of an owner portal — which is the half of Vantage that has never been contested. Combined with the Aug 26 Copilot-in-editor item already logged as **#84**, Duda has now shipped, within twelve days: conversational control of the whole site, AEO plumbing, forms, and auth. **Do not overstate it.** Vibe generating a login screen is not the same as a multi-tenant owner dashboard wired to a lead inbox, an AI receptionist, reviews and analytics — **the bundle is still ours.** But the phrase "they can't build a portal" has quietly stopped being true, and it should come out of the pitch before a prospect discovers it.

Impact □ / Effort S (a positioning edit). → Register #95; updates #84, #21.

□ Squarespace and Wix: checked by name, nothing dated in the window.

Squarespace's product-updates surface and Wix's release coverage returned no September item; the most recent material for both is the Wix Harmony / vibe-coding line from January and Squarespace's AEO scanner already closed as **#83** ([Squarespace product updates](#)). Consistent with #83's conclusion: **AEO at the builders is settled as table stakes; stop re-scanning it.**

□ AI receptionist pricing — a second comparable landed, and it is much higher than the channel's \$99. White-label partners can price **a single AI voice agent at \$75-\$150/mo at 70%+ margins**, which is the figure already in **#71/#60**. But the same research puts **home-services resale at \$297-797/mo with 60-80% margins**, against wholesale voice minutes of roughly **\$0.08-\$0.24** ([Ringlyn white-label pricing](#), [Trillet home services](#)).

Why it matters: **#71** framed our top tiers as "2.5-9× the channel price" and treated that as a positioning problem. **A 4-8× spread between two verticals for the same underlying technology says the number is set by the vertical's cost-of-a-missed-call, not by the technology.** A missed HVAC emergency and a missed wedding-venue tour enquiry are both expensive; a missed generic SMB call is not. That reframes **#71** from "*justify our price*" to "*pick the vertical where the price is obvious*" — and venues are already that vertical. **Impact □ / Effort S. → updates #71, #60, #63.**

6. MSP channel + service-delivery automation

Surfaces checked by name this run: pia.ai (blog + news + automations), ChannelPro Network, CRN — plus the M&A multiples sources the acquisition

thesis rests on.

□ **Pia's Q3 2026 release is the clearest buy-vs-build benchmark #8 has ever had, and one feature in it is a direct instruction.**

What it is: Pia's **Q3 2026 release** ships an **AI Automation Builder** that "*turns plain-language prompts into complete, ready-to-deploy automation packages*"; **self-serve SmartForms** so end users help themselves; an **Automation Hub with 48+ ready-to-deploy PiaPacks**, including automations for travelling users and O365 mail tracing; **18 new API endpoints** across Automate and Configure; and a **smarter Pia Chat** that turns **Microsoft Teams into a real self-service channel** — users can *trigger real fixes like restarting services or pulling reports, mark issues resolved, and view open tickets and known issues, inside Teams* ([Pia product update](#), [Pia news](#), [Channel Insider](#)). △ **Vendor's own material — treat capability and ROI claims as marketing until corroborated.**

Why it matters: the standing read of Pia is that **its catalog is a published list of what a funded vendor believes is worth automating** — i.e. the highest-volume SMB service-desk work. **48+ packs** is that list getting longer, and the two named examples (travelling users, O365 mail tracing) are both **EdgePointe-shaped tickets**. But the item that should change behaviour is **Teams as the service desk**. EdgePointe's clients are M365 clients; **Teams is already open on every desktop**; and a ticket raised and resolved without a phone call or a portal login is the cheapest ticket there is. That is not an AI feature — it is a **channel** decision, and it is buildable without any of Pia's automation engine.

How EdgePointe applies it: **#8** stops being "point the engine at Managed IT" in the abstract and gets a first concrete target: **a Teams-native intake for the top three ticket types**, with the Pia catalog used as the evidence for which three. Do not buy Pia to find out what to automate — **read the catalog, which is free. Impact □ / Effort M. → Register #96; gives #8 its first specific scope.**

□ **ChannelPro's read of the market is an argument against the way #60 and #71 are currently framed.**

What it is: ChannelPro's channel reporting this window: "*MSPs are moving past the AI add-on conversation.*" The framing is that **AI is increasingly treated as a native capability inside platforms and tools rather than a separate layer**, and the question has shifted from *whether to offer AI services* to **how AI changes the way services are delivered, monitored, secured, and packaged** ([ChannelE2E channel brief](#), [ChannelPro Network](#)).

Why it matters: **#60** and **#71** both assume AI is a thing you price and sell as a line. If the

channel is folding AI into the base platform, then **an "AI tier" starts to read as an upsell for something competitors include** — the same trap Duda and Webflow just sprang on AEO. The defensible position is the opposite: **AI is included, and what you pay for is the outcome** (calls answered, leads captured, tickets closed). △ This cuts against a natural instinct to price the expensive part visibly, and it should be argued rather than adopted — but it is the third independent signal this quarter pointing the same way.

How EdgePointe applies it: when **#71's** positioning answer gets written for Bob, write it as **outcome pricing with AI included**, not as an AI tier. → **updates #60, #71, and the service-catalog work.**

□ **GTIA's Managed Intelligence Alliance now has a date and a membership list. #75 gets a calendar.** Executive Members are **The 20, New Charter** and **Pax8**; the alliance is **forming in Q3 2026** and holds its **first official member meeting in October 2026**, alongside the Texas A&M GCRI Annual Summit ([Channel Insider](#), [ChannelPro](#), [GTIA press release](#)). The alliance is deliberately housed at **GTIA rather than Pax8** to keep it vendor-neutral. **#75's practical question — whether an accreditation exists that EdgePointe can hold before competitors do — now has a checkable date rather than an open watch.** → **updates #75.**

□ **The M&A numbers moved, and the movement is in the security line, not the recurring line. This is #50's most useful update since it opened.**

What it is: 2026 MSP M&A data. **Sub-\$1M revenue** MSP/IT services trade at **3.0-5.0× SDE**; **\$3-10M revenue** at **6.0-9.0× adjusted EBITDA**; **\$25M+ platform-scale MSSPs** at **12.0-16.0×**. **66% of investment bankers** say **recurring revenue** will be the most important characteristic to acquirers in 2026. And the headline: **the MSSP premium over a comparable MSP has widened from roughly 1.5× adjusted EBITDA in 2020 to roughly 2.5× at the Q2 2026 midpoint** ([CT Acquisitions multiples guide](#), [Eight-M MSP multiples](#)).

Why it matters: **#50** currently says *recurring-revenue share drives the exit multiple* and treats that as the largest financial lever on the board. **That is still true and the 66% figure confirms it.** What is new is that **security has become a second, independent lever, and it is widening** — a full turn of extra premium accrued in six years, and the movement is recent. For EdgePointe that is unusually actionable because the **cybersecurity service line already exists**; the question is whether it is packaged and measured as a **security practice** with its own recurring revenue, or dissolved into Managed IT where it earns the MSP multiple rather than the MSSP one. △ These are broker-published ranges — directional, not a valuation. *How EdgePointe applies it:* when

the acquisition thesis is next reviewed with Bob, **carry two lines: recurring share, and security-attributable recurring revenue as a percentage of it. Impact □ / Effort S (a reporting change).** → Register #97; updates #50.

Top 5 act-now (impact vs effort)

#	Item	Why now	Impact / Effort
1	#32 — run the five-minute zone check	8 DAYS. Third consecutive week at the top. Plan tier + ad presence per tenant zone. The answer is almost certainly "no exposure" — record it and close the row	□ / S
2	#85 — write the "why doesn't Vantage run my ads" paragraph	Owed before the 09-11 proposal — 4 days. Webflow Campaigns is beta not GA (\$5), which softens the calendar but not the question a prospect asks	□ / S
3	#89 + #88 — run /skill-doctor on one canary task, then declare two skills with ant apply	Two commands replace two long-standing projects (#5, #57). One sitting, and the first is fifteen minutes	□ / S-M
4	#90 — grep the capture pipeline for Browser Run /crawl	New 400 on restrictive Content Signals. Either it is a silent break waiting for a prospect, or exposure is nil and the row closes today	□ / S
5	#76 — make the Next.js client inventory call	The WAF rules merge TOMORROW (Sept 8). A live rule change on a live unauthenticated RCE is the cleanest outbound pretext this board has produced	□ / S

Standing, unranked, and the reason all five above are small: #72 — ship what is already built. The AI receptionist answers a phone number and Phases 1-3 are code complete. Nothing on this list outranks shipping finished work.

Delta vs the previous weekly (2026-08-31)

Area	08-31 said	09-07 says
Method	<i>The vendor was right, the surface was wrong — every vendor carries a named URL</i>	Followed, and it worked (Duda and Cloudflare read from changelogs; Webflow corrected). New failure found one layer down: #87 was announced in prose and never written as a row. A number is created by writing the row
Evals (#1/#3)	Add <code>evidence_kind</code> ; agreement may be shared shallowness	Sharper and worse: the rubric leaks its own answer (2609.02942). Criteria that survive inversion or a blank page are artifacts. This week's test is a delete-list, not a build
#57 / #5 / #49	<i>Check whether Claude Code restricted mode expresses #57 natively before hand-rolling</i>	Answered, and better than expected: <code>ant apply</code> (Sept 3) + <code>/skill-doctor</code> (Sept 4). Both are commands, not projects
Webflow (#85/#86)	Webflow inside ChatGPT/Codex as skills (#74)	Source is a limited research preview, Campaigns beta purchasable next month. 09-04's substance stands; its calendar was tighter than the facts
Cloudflare	Emergency WAF release for the Next.js RCE pair (#76)	Those rules merge into the originals TOMORROW. Plus a hard new gate on crawling (#90), 64 MiB bundles (#91), text rasterization (#92)
#32	15 days, scope unverified, second consecutive week	8 days, scope still unverified, THIRD consecutive week. New: Pay Per Crawl → Pay Per Use , paid on <i>use in an answer</i> (#93)
MSP channel (#60/#71)	GTIA alliance gives the category a standards body	Alliance dated — first meeting October. And ChannelPro says AI is becoming native, not an add-on , which argues

Area	08-31 said	09-07 says
		against pricing an AI tier at all
#50 (acquisition)	Recurring-revenue share drives the multiple	Still true (66% of bankers) — plus a second lever: the MSSP premium widened 1.5× → 2.5× EBITDA. Security is now its own multiple
#10 (moat)	Unbuilt, eighth week	Unbuilt, NINTH week. The channel folding AI into every platform is the strongest argument yet that the only differentiator left is proprietary data
#52 (Sonnet pricing)	<i>Tomorrow is Sept 1; do not resurrect the countdown</i>	Sept 1 passed. Nothing happened. Fourth confirmation. Permanently closed

Lab hour — 30 minutes, this week

The null-and-inversion control on the design-critic rubric.

The last six weeks have added rules to the rubric. This week **removes** criteria — and it is the first experiment on this board whose success condition is a shorter file.

Run it in three steps against the rubric that already exists:

1. **Null control.** Feed the rubric a **blank page** (an empty `<html><body></body></html>`, or a page stripped to a single `<div>`) and score every criterion. Record each score.
2. **Inversion control.** Take each criterion and **negate its wording** — "*the hero states the venue's location above the fold*" becomes "*the hero omits the venue's location above the fold*" — and re-score the **same real page** (use Hidden Creek, which has a known-good result).
3. **The delete-list.** Any criterion that (a) returns a **confident, non-neutral score on the blank page**, or (b) returns a **materially unchanged score when inverted**, is an artifact. Delete it.

The prompt to run it with:

*You are scoring this page against the rubric below. For every criterion return {score, evidence, evidence_kind} where evidence is a selector, computed value, literal string, or layout measurement actually found in the page, and evidence_kind names which of those four it is. **If the page does not contain evidence for a criterion, return score: null and evidence: null. Do not infer, do not average, do not score from the criterion's wording.** Return nothing but the JSON.*

What each outcome means:

- **Criteria score confidently on the blank page** → the rubric is scoring its own wording. This is exactly the 2609.02942 result reproduced on our own instrument, and those criteria are dead weight that has been inflating every mirror score we have taken.
- **Scores barely move on inversion** → the criterion is decoration. Delete without ceremony.
- **Most criteria return null on the blank page and flip cleanly on inversion** → **the rubric is sound**, the six weeks of amendments worked, and #1/#3 can finally move from "amend the method" to "run the method." **That would be the most valuable possible result and it should be recorded as loudly as a failure.**

Total: about 30 minutes, no new code, no new model. The output is a delete-list and a one-line verdict on whether the design critic measures the page or itself.

Method note: this run read the Cloudflare changelog index, duda.co/product-updates and Webflow's Conf announcements directly; Anthropic items are dated from the release-notes feed against docs.claude.com/en/release-notes/overview.md. Perfect Venue, Tripleseat, HoneyBook, Squarespace and Wix were checked by name and returned no dated item in the window — recorded as a weaker read, since none publishes a changelog surface equivalent to the above. pia.ai, ChannelPro and CRN were checked by name; Pia's own material is treated as vendor marketing. No compensation or equity material appears in this report.