

Build a business that *thinks*.

Marketing agency that builds and grows your business in an
AI-first world.



Build a business that thinks.

A business was only ever as good as its best person. We build yours to operate like a room full of experts, one for every part of it.

For decades, software waited for instructions. You opened it, you did the work, you closed it. That era is ending. The businesses pulling ahead now run on systems that think for themselves, do the work, and keep getting better while everyone else is asleep. Before long this is simply how a serious company runs. The only real question is whether you are early or catching up.

Fallback Labs builds you the early version. We give your business a brain that knows it and nine experts that act on it, one each for your website, marketing, generative AI, local SEO, social, email, paid ads, sales, and security. The outcome is a company that markets itself, sells for you, defends itself, and improves every week, with real expertise in every seat and no need to hire for every seat.

Here is why moving first pays. The brain gets sharper the longer it runs, so the advantage compounds. A business that started a year ago is a year of learning ahead, and that lead widens every month it keeps working. Start now and you are not just adopting the new normal, you are banking a head start no one can buy back later.

And the result is yours. It runs on infrastructure in your name, it keeps working if we ever part ways, and everything it learns about your business belongs to you. We are founders who lived the pain of establishing, building, and growing a company in the UAE, so we built the tool we wish we had. Establish, build, and grow, all in one place. The businesses that think are about to be the ones that win. We make yours one of them, first.

TRUSTED BY



and we are just getting started



Nine brains. Five *teams*.

A brain that knows your business, one for every part of it, and full teams of specialists that do the work. Establish, build, and grow, all in one place.

■ SECTION 01 · THE BRAINS

01	Website Brain	MEMORY	P. 05
02	Marketing Brain	STRATEGY	P. 06
03	Gen AI Brain	INTELLIGENCE	P. 07
04	Local SEO Brain	LOCAL PRESENCE	P. 08
05	Social Media Brain	SOCIAL	P. 09
06	Email Marketing Brain	MESSAGING	P. 10
07	Paid Ads Brain	REACH	P. 11
08	Sales Brain	PERSUASION	P. 12
09	Cybersecurity Brain	DEFENSE	P. 13

■ SECTION 02 · THE TEAMS

01	SEO Team	SEARCH	P. 15
02	Content Team	CONTENT	P. 16
03	Ads Team	ADVERTISING	P. 17
04	Email Team	EMAIL	P. 18
05	Video Team	VIDEO	P. 19
■ AND			
Straight answers (FAQ)			P. 20
Contact			P. 22

01

The nine *brains*.

What your business knows. A brain that reads and remembers who you are, and nine of them, one for every part of the job, so nothing that follows comes out generic.

01 Website

02 Marketing

03 Gen AI

04 Local SEO

05 Social Media


06 Email Marketing

07 Paid Ads

08 Sales

09 Cybersecurity

Website Brain



● MEMORY

■ WHAT IT IS

The foundation. It points at any website, yours, a client's, a competitor's, and comes back with a complete, connected memory of that business: every page, product, and promise, plus the brand itself. Everything the other eight Brains do is grounded in it, so nothing comes out generic.

■ HOW IT'S BUILT

It studies the live site the way a brand designer would, then goes further: it measures your true colors and fonts by how often you actually use them, collects and organizes every image with duplicates removed, and locks the result into one reusable brand profile that other tools consume directly. Its only source is the site itself, captured once, then usable fully offline. Before the capture counts, six independent reviewers grade it, and it ships with the grade attached.

■ WHERE YOU'D USE IT

Onboarding a new client in minutes instead of weeks. Capturing a site before a redesign so nothing is lost. Studying a competitor. Generating on-brand assets without re-explaining your brand every time.

■ WHAT IT PRODUCES

A complete brand kit (palette, fonts, logo, image style) plus a readable profile of what the business says and does, an organized image library, and the shared memory every other Brain reads from. One capture, many on-brand outputs.

■ WHY IT BEATS A CLAUDE/CODEX PROMPT

Connect a crawler and a prompt can read your site too. It skims, guesses your palette, gives a different answer every time, and starts blank next session. This captures everything once, measures your colors and fonts by actual usage, clears every duplicate image, grades its own capture with six independent reviewers, and keeps it all. Connects to: every other Brain and every Team reads this memory first.

■ IN CONVERSATION

Capture our website into a brain.

Done. I read all 63 pages, organized 196 images, and measured your true brand colors and fonts from how your site actually uses them. Six independent reviewers graded the capture: A, 92 out of 100.

✓ **SAVED TO MEMORY** · full site capture, brand kit, image library

What are our exact brand colors?

Sky blue on near-black. Not a guess: confirmed in your site's own styling and sampled from your hero image. It's locked in your brand kit, so everything we make for you matches.

GRADED BY 6 REVIEWERS

92/100 REFERENCE BUILD

COLORS MEASURED, NOT GUESSED

WORKS OFFLINE AFTER CAPTURE

NO TRACKING

FEEDS ALL 8 BRAINS

■ BY THE NUMBERS

Manual audit: days. Agency: weeks. This: minutes. Reference build: 63 pages, 196 organized images, about 1,700 internal links, 21 entity notes, graded A (92/100) with zero blockers. Crawl once, then work fully offline. No lock-in, no tracking.

~95K

EST. TOKENS SAVED
PER TASK

Marketing Brain



● STRATEGY

■ WHAT IT IS

Your strategist. It studies your market and competitors with live search data and turns a vague goal into a cited 30/60/90-day execution plan: which channels, which messages, which pages, in which order.

■ HOW IT'S BUILT

It follows Google's own published standards for helpful, trustworthy content and a proven method called FLOW (Find, Leverage, Optimize, Win). Research runs on live search data pulled at the moment of analysis, with every raw result saved as evidence and hard cost caps on every run. Each keyword is scored with the same opportunity formula, a second specialist proofs the sheet for errors, and the finished plan is fixed at ten parts where every number traces to a source or is marked unconfirmed. It is white-hat and read-only by design.

■ WHERE YOU'D USE IT

A real 30/60/90 plan. Choosing what to build next and in what order. Seeing where a competitor is beatable. Replacing gut-feel channel decisions with sourced ones.

■ WHAT IT PRODUCES

The flagship strategy plan (3,000 to 5,000 words, ten parts), a scored keyword workbook, a keyword-to-page map, a cannibalization check of pages competing with each other, and a competitor gap analysis, all tuned to one of six business types, from local services to online stores to software.

■ WHY IT BEATS A CLAUDE/CODEX PROMPT

A prompt with live search tools pulls keyword data, then improvises the plan: different every run, confident even where data is missing. This runs one fixed method: every pull saved as evidence, one scoring formula, a proofing pass, and a ten-part plan where every number cites a source or says unconfirmed. It refuses guaranteed-rankings talk entirely. Connects to: hands its plan to the SEO Team and Content Team to execute.

■ IN CONVERSATION

Build our 90-day plan.

Plan ready: ten parts, built from live research pulled today. Every number traces to a saved source; where data was missing I wrote "unconfirmed" instead of guessing.

✓ **SAVED TO MEMORY** · plan filed with its evidence, marked "ready for your review"

Can you guarantee us #1 on Google?

No, and no honest strategist can. What I can show you is where you can realistically win, with the evidence attached.

LIVE SEARCH DATA

COST-CAPPED RESEARCH

EVERY NUMBER SOURCED

6 BUSINESS TYPES

WHITE-HAT, READ-ONLY

NO RANK GUARANTEES

■ BY THE NUMBERS

A ten-part plan of 3,000 to 5,000 words. Six business types. Live research with strict cost caps per run. Every figure traced to a source or marked unconfirmed. Read first, write second, verify third.

~120K

EST. TOKENS SAVED
PER TASK

Gen AI Brain



● INTELLIGENCE

■ WHAT IT IS

The expert reference for building with today's most capable AI. It is a complete, verified library of how to build on Google's generative-AI stack (Gemini, Vertex AI, image and video generation), so anything built for you follows the right pattern, done the right way.

■ HOW IT'S BUILT

One strict rule: cite the source. All 318 of Google's official builder playbooks were distilled in advance into one connected library, with every claim tied to the exact official example it came from. Then it was independently audited by a 66-reviewer verification pass: 95 out of 100, effectively zero invented facts, and a release check that blocks on a single dead link.

■ WHERE YOU'D USE IT

Designing or reviewing any AI feature on Google's stack. Finding the right model or pattern in minutes instead of opening 318 notebooks. Onboarding an engineer. Grounding the other Brains whenever a build touches Gemini or Vertex.

■ WHAT IT PRODUCES

686 connected reference pages: 317 source pages, 20 concept maps, 20 best-practice guides, and 67 model pages, with 494 real architecture diagrams preserved. In practice, it is the foundation the team builds from, so what ships for you is current and correct rather than recalled from memory.

■ WHY IT BEATS A CLAUDE/CODEX PROMPT

Ask a prompt to be a Vertex AI expert and it names tools from memory, some stale, some invented. This distilled all 318 of Google's official builder playbooks in advance, tied every claim to its exact source, and was independently audited at 95/100 with effectively zero invented facts. Answers arrive in minutes without opening a single notebook. Connects to: grounds every AI feature we build, including the Gemini-powered image generation the Teams use.

■ IN CONVERSATION

What's the right Google model for generating product photos?

Imagen, and here's the official pattern for using it. That answer isn't from memory of the internet: I hold all 318 of Google's own builder playbooks, distilled and cross-checked, with every claim linked back to its source.

✓ **FROM MEMORY** · 73 official playbooks on this exact topic · independently verified 95/100

318 OFFICIAL PLAYBOOKS

686 CONNECTED PAGES

95/100 VERIFIED

ZERO INVENTED FACTS

0 DEAD LINKS

WORKS OFFLINE

■ BY THE NUMBERS

318 official playbooks distilled into 686 pages with 494 diagrams. Independently verified 95/100 by a 66-reviewer audit, effectively zero invented facts, zero dead links. Works fully offline with zero credentials.

~450K

EST. TOKENS SAVED
PER TASK

Local SEO Brain



LOCAL PRESENCE

WHAT IT IS

It plants a local business on the map so nearby customers find it first: rankings, Google Business Profile, reviews, and listings, run as one system with memory. Its own line says it best: the brain produces the recipe; you cook.

HOW IT'S BUILT

It follows a three-step local plan (site first, then Business Profile, then reviews, listings, and links) on top of a proven "Core 30" content strategy and 59 chapters of current local-search knowledge. It works from your real local data, Business Profile insights, review exports, listing reports, and block-by-block map-rank scans, each stamped with its source and date. And it enforces hard refusals: no fake reviews, no "guaranteed number one," no profile changes without your sign-off.

WHERE YOU'D USE IT

Agencies running 5 to 50 client locations. Multi-location operators. Service businesses (plumbers, HVAC, law, dental) competing on "near me" searches. Anyone losing weeks to spreadsheet audits and screenshot dumps.

WHAT IT PRODUCES

A map of where you rank across your whole service area and where you are invisible. A locked, consistent version of your business details with every wrong listing flagged. A review picture with response priorities. A health scorecard, action roadmap, and weekly report, each recommendation carrying its source.

WHY IT BEATS A CLAUDE/CODEX PROMPT

Hand a prompt your rankings and reviews and it reads them once, then forgets. Rank trackers show dashboards with no judgment; templates give you blank pages. This keeps a permanent record per business, computes the map math a prompt would eyeball (your top-3 share, whether distance is the real blocker), cites every recommendation, and refuses fake reviews and "guaranteed number one." Connects to: the SEO Team's local mode runs the hands-on audits.

IN CONVERSATION

Scan our map rankings.

Scanned 49 points across your area. You average #7, but on the north side you're invisible, and only a third of your area sees you in the top three. Three fixes are queued for your approval.

✓ **SAVED TO MEMORY** · scan results per keyword, this week's fix list

Can we just buy some reviews?

No. Fake reviews are a hard refusal here; they get businesses removed from Maps. There are three clean ways to lift your rating, and they're already in your plan.

BLOCK-BY-BLOCK RANK SCANS

59 KNOWLEDGE CHAPTERS

EVERY CLAIM DATED

NO FAKE REVIEWS, EVER

YOUR SIGN-OFF REQUIRED

DEMO IN 30 SECONDS

BY THE NUMBERS

59 knowledge chapters, 37 ready prompts, 1,271 cross-links. A three-phase plan plus the Core 30. Map Pack: 60 to 70 percent of local clicks. Demo runs in under thirty seconds. A year of client history stays under 200 MB.

~85K

EST. TOKENS SAVED
FALLBACKLABS.COM

Social Media Brain



SOCIAL

WHAT IT IS

It keeps you present and ranking across seven platforms without you living inside the apps. The goal is not to post more. Its own motto: rank, not just post.

HOW IT'S BUILT

Most social playbooks freeze the day they ship and rot within a quarter. This one refreshes itself: it researches each platform's current ranking rules against a strict source ladder with the platforms' own statements at the top, dates every claim, sets a recheck date, and re-verifies monthly. A second pass opens the sources and drops anything that does not hold up. It optimizes every piece three ways at once (feed reach, in-app search, and AI citations) and folds your own results back into a scorecard, so it learns what works for you specifically.

WHERE YOU'D USE IT

A solo creator staying consistent. A founder building an audience. A content team or agency running multiple accounts. Anyone tired of undated "expert tips" that were true two algorithms ago.

IN CONVERSATION

What's actually working on Instagram right now?

Watch time, likes per reach, and shares per reach; Instagram's own head named them as the top signals. That claim sits in my memory with its source and date, and it gets rechecked monthly. The moment it stops being true, it leaves your playbook.

✓ **FROM MEMORY** · Instagram playbook · every claim dated and sourced · next recheck July 10

7 PLATFORMS

3 RANKING LAYERS

MONTHLY RECHECKS

429 CLAIMS RE-VERIFIED

ZERO FABRICATED NUMBERS

CITE OR OMIT

BY THE NUMBERS

7 platforms, 3 ranking layers, a 10-part ranking playbook per platform. Monthly rechecks on every claim. Latest audit: 429 claims re-verified, zero fabricated numbers, 95/100 trust.

~110K

EST. TOKENS SAVED
PER TASK

Email Marketing Brain



MESSAGING

WHAT IT IS

The email department's memory and rulebook: strategy, lifecycle flows, retention, deliverability, and compliance, tuned to your exact business, with a drafting layer that never sends on its own.

HOW IT'S BUILT

It plans with proven copy formulas and full lifecycle logic across 20 industry playbooks, from B2B software to online stores, creators, nonprofits, and local businesses. Every claim cites a library of 364 dated sources, zero stale: the inbox providers' actual sender rules and the real laws with real penalties, including CAN-SPAM at \$53,088 per email. Execution is gated: eight hard checks stand between a plan and any send, and a live send needs three separate switches plus your confirmation.

WHERE YOU'D USE IT

Standing up welcome, retention, and win-back programs. Planning a compliant launch across regions. Keeping a list warm without hiring a lifecycle marketer. Automating safely.

WHAT IT PRODUCES

Sequence and nurture programs per industry, 44 ready templates, 15 audited automation workflows chosen trap-free from a field of 80, a compliance matrix for your regions, a deliverability risk check, and ready-to-review drafts in your inbox. Where an industry lacks benchmarks it says "not benchmarked." Never a fabricated KPI to fill a gap.

WHY IT BEATS A CLAUDE/CODEX PROMPT

Connect a prompt to your email and it will draft and send with nothing checking consent or the law. Here, eight enforced gates stand before any send, consent is verified for every single contact, a kill switch stops everything, and it is draft-first always. Its advice cites 364 dated sources: the providers' actual rules and the real laws, not the model's memory of them. Connects to: hands the deep technical checks to the Email Team.

IN CONVERSATION

The campaign looks good. Send it to everyone.

Held it. Part of your list has no recorded permission, and I won't email anyone without one; that's the law, and it's my rule. Everyone else is drafted in your inbox. You review, you send.

✓ **SAVED TO MEMORY** · campaign plan, permission check results, drafts logged

364 DATED SOURCES

20 INDUSTRY PLAYBOOKS

8 SEND GATES

CONSENT CHECKED PER CONTACT

KILL SWITCH

NEVER AUTO-SENDS

BY THE NUMBERS

364 dated sources, zero stale. 20 industry playbooks, 44 templates, 15 audited automations. 8 enforced gates before any send. Current penalty exposure cited precisely, including CAN-SPAM at \$53,088 per email.

~130K

EST. TOKENS SAVED
PER TASK

Paid Ads Brain



REACH

WHAT IT IS

Persistent memory for your paid media across 12 platforms. It ingests your authorized account data, grounds every finding in evidence, and produces client-ready reporting and next actions. It is not an autonomous ads agent, and that is deliberate: it is the operating system a serious ads practice runs on.

HOW IT'S BUILT

Evidence first, always. Your authorized exports become tamper-evident evidence files; findings cite them or do not ship. It knows the difference between what got credit and what caused the sale, and it treats "unknown" as different from "zero." Account changes are switched off outright, and it even refuses to output a health score until scoring is properly grounded, because its own rule is that maturity is evidence-based, not a claim in a document.

WHERE YOU'D USE IT

Agencies and freelancers running client accounts. Founders and in-house teams who want week-over-week truth. Anyone who has been burned by confident numbers that turned out to be invented.

IN CONVERSATION

What's our account health score?

I'm not giving one yet. Your tracking isn't verified, and a score without evidence would be a guess. What I can prove from your own data: money is going to searches that never turn into sales. The cut list is waiting for your approval, and I never touch your account myself.

✓ **SAVED TO MEMORY** · this week's findings, wasted-spend list, approval queue

12 PLATFORMS

WEEK-OVER-WEEK MEMORY

EVERY FINDING SOURCED

SAYS "NO DATA," NEVER GUESSES

CANNOT TOUCH ACCOUNTS

APPROVAL-FIRST


BY THE NUMBERS

12 platforms. 10 standing deliverables. Every artifact fingerprinted for integrity. Account changes: disabled. Guarantees of ROAS: refused on principle.

~150K

EST. TOKENS SAVED
PER TASK

Sales Brain



PERSUASION

WHAT IT IS

Evidence-based selling for your exact market. It arms every pitch with proof, routes each deal to the right method, and builds industry playbooks that accumulate instead of evaporating after each conversation.

HOW IT'S BUILT

It distills 100 classic and modern sales books plus 22 dated primary sources (regulators, peer-reviewed studies, named industry data), and labels every claim as proven, contested, or myth. It routes deals by explicit rules: large, multi-stakeholder, long-cycle deals get MEDDPICC and Challenger; mid-size get MEDDIC and Gap Selling; simpler ones get BANT and SPIN. And it actively debunks the folklore, including the famous "78 percent buy from the first responder" line, which has no primary source.

WHERE YOU'D USE IT

Founders selling their own product. Sales teams that want one evidence-backed playbook instead of tribal knowledge. Agencies selling retainers. The stakes are real: 69 percent of reps miss quota, and 17 percent of reps generate 81 percent of revenue. The gap is method.

IN CONVERSATION

They said we're too expensive. What do I say?

Not the "feel, felt, found" script; my research found no real evidence it works, so it's marked folklore in your playbook. A deal this size routes to a proper enterprise method instead. Here's the response, with two proof points, both cited.

✓ **SAVED TO MEMORY** · objection and answer added to your library for next time

100 BOOKS DISTILLED

PROVEN / CONTESTED / MYTH LABELS

DEAL-SIZE ROUTING

14 INDUSTRY PLAYBOOKS

7-COUNTRY COMPLIANCE CHECK

NO INVENTED DATA


BY THE NUMBERS

100 books distilled. 22 dated primary sources. 14 industry playbooks. Compliance across 7 countries. Named methods: MEDDIC, MEDDPICC, SPIN, Challenger, Gap Selling, BANT. Released market-ready only after a 90-plus audit.

~9M

EST. TOKENS SAVED
PER TASK

Cybersecurity Brain



DEFENSE

WHAT IT IS

It guards your site, product, and hosting, and shuts the doors before anyone walks through them. Timely, too: AI-written code is shipping vulnerabilities at record rates, and the common attacks were just reshuffled by the industry's own 2025 risk rankings.

HOW IT'S BUILT

It checks your actual setup against the recognized standards (OWASP's Top 10 risks, the NIST framework, CIS baselines) and scores severity with the official vulnerability formula, automatically escalating anything on the government's actively-exploited list. Behind it sit 267 dated, verified sources, refreshed on schedule, hardened by three adversarial fact-check passes that corrected 63 claims before release. Every recommendation carries at least two official sources and an undo step. It is read-only and defense-only: no exploit instructions, ever.

WHERE YOU'D USE IT

A pre-launch review. Hardening hosting and configuration. Ongoing watch on real-world exploits. A defensive posture record an agency can hand to a client. Built for indie hackers, web developers, agencies, founders, and security engineers.

WHAT IT PRODUCES

A four-area security review (website, hosting, product, compliance) with pass or fail verdicts and evidence per control. A hardening checklist tailored to your stack. A priority-ranked fix list with actively-exploited threats first. A compliance evidence map for PCI DSS, SOC 2, GDPR, and HIPAA. A weekly threat digest.

WHY IT BEATS A CLAUDE/CODEX PROMPT

A prompt lists best practices, and even connected it lacks discipline: which standard covers this, how severe is that really, what is being exploited this week? This maps every finding to OWASP and NIST, scores severity with the official formula, auto-escalates actively-exploited threats, requires two official sources per recommendation, and tracks findings over time so your posture improves instead of resetting. Defense-only by hard rule.

IN CONVERSATION

Anything urgent this week?

One thing. A critical flaw affecting your setup is being actively exploited right now, so it jumps to "patch now." The fix comes with an undo step, like every change I recommend. Two smaller items can wait for the sprint.

✓ **SAVED TO MEMORY** · finding logged, priority set, tracked until closed

267 VERIFIED SOURCES

100/100 RELEASE AUDIT

OFFICIAL SEVERITY SCORING

ACTIVE-EXPLOIT ESCALATION

2 SOURCES PER RECOMMENDATION

DEFENSE-ONLY

BY THE NUMBERS

153 interlinked security notes across 4 areas. 267 verified dated sources. Released only after a 100/100 structural audit. Severity scored on the official CVSS system with active-exploit escalation. Defense-only, never touches production.

~200K

EST. TOKENS SAVED
PER TASK

02

The five *teams*.

If the Brains are the memory and judgment, the Teams do the work. Each Team is a group of AI specialists doing one job end to end: same fixed method every run, live data, hard quality gates, and a scored result. Where a chat gives you its first draft, a Team gives you its checked one.

01 SEO

02 Content

03 Ads

04 Email

05 Video

SEO Team



Team 01 visual
or [browse files](#)

SEARCH

WHAT IT IS

A complete SEO department that gets your site found on Google and in AI answers. One command runs a full audit; what comes back is a scored, prioritized, evidence-backed plan. Its motto: falsifiable, not promotional.

HOW IT'S BUILT

It reads your real site, up to 500 pages, and puts up to 15 specialists on it at once across 25 capabilities: technical health, content quality, markup, speed, local, store pages, links, and AI search. It works from live data (real rankings, real-world speed from actual Chrome users) and Google's current standards, including the May 2026 AI guidance and the September 2025 quality guidelines. Every recommendation ships with a built-in test: the observation it rests on, and a "how would we know this failed?" check with a leading indicator to watch. It is deterministic (ask twice, get the same rigor twice) and fully local: your data never leaves your machine.

WHERE YOU'D USE IT

Agencies: the client health-score email drops from 4 hours to 12 minutes, and quarterly deep audits become weekly at the same team capacity. In-house: run it 24 hours before the quarterly review to catch what dashboards hide. Freelancers: a real score on the discovery call, so the proof of value happens during the call, not after the proposal.

IN CONVERSATION

WHAT IT PRODUCES

A 0-to-100 health score across seven weighted areas. A prioritized fix list, ready-made markup, sitemaps, content briefs, and topic clusters. A client-ready PDF report of roughly 32 pages. And a drift watch that catches regressions between audits.

WHY IT BEATS A CLAUDE/CODEX PROMPT

Ask a connected prompt for an audit and you get a different one every time, covering whatever occurs to it that day. This runs the same 25-capability catalog every run with up to 15 specialists, knows exactly what changed and when (FAQ rich results retired May 2026), debunks the myths prompts repeat (Google ignores llms.txt), and attaches a "how would we know this failed?" check to every recommendation. Connects to: takes strategy from the Marketing Brain, runs local mode with the Local SEO Brain, and keeps its own Pro memory vault.

Write us a new page for "business setup in Dubai."

Thinking... Checking your keyword ledger: that search already has an owner page. Reading the search-results evidence saved July 21. Two pages on one keyword would compete with each other.

Don't build a new page. You already own that search with your business-setup page, and a second one would split your strength. Better move: refresh the owner page. And here's how we'll know it worked: impressions should move within 28 days, or we revisit.

FULL AUDIT IN 10-15 MIN

25 CAPABILITIES

15 PARALLEL SPECIALISTS

326 AUTOMATED TESTS

FALSIFIABLE RECOMMENDATIONS

FULLY LOCAL, NO LOCK-IN

BY THE NUMBERS

A full audit in 10 to 15 minutes, versus 4 to 8 hours of senior time manually, 1 to 3 weeks and \$2k to \$15k+ from an agency, or \$99 to \$999 a month for tools that typically run 6 to 12 months behind. 25 capabilities, 18 specialists, 326 automated tests. Real proof: a site ~~FA~~ started in March 2026 and run on this workflow showed steady click growth in its first three months.

FALLBACKLABS.COM

02

Content Team



Team 02 visual
or [browse files](#)

● CONTENT

■ WHAT IT IS

A full content department that writes and optimizes articles to rank on Google and get cited by AI (ChatGPT, Perplexity, AI Overviews, Gemini). Its core promise: you are never the first reviewer.

■ HOW IT'S BUILT

It routes any request across 31 capabilities and 12 article templates, with five specialists handling research, writing, SEO, review, and translation. Statistics come only from vetted source tiers, and a fact-check opens every cited link to confirm the claim really appears on the page; the workflow is designed to block invented numbers. Every piece is graded on a 100-point system and nothing ships below 90 or with a single critical issue; it revises itself up to three times before you ever see it. It is also built on how AI actually cites: three-quarters of ChatGPT's most-cited pages were updated within the last 30 days, and fresh content is three times more likely to be cited, so freshness is engineered in, not hoped for.

■ WHERE YOU'D USE IT

Solo founders shipping serious posts without losing days to SEO plumbing. Teams and agencies keeping quality consistent across brands and languages. Refreshing decaying content before rankings slip.

■ WHAT IT PRODUCES

Finished articles delivered as document, web page, PDF, and hero image, with markup, charts, licensed images, and a quality scorecard. Plus briefs, calendars, topic clusters, social versions, multilingual editions, and even audio narration.

■ WHY IT BEATS A CLAUDE/CODEX PROMPT

A connected prompt researches and writes, then ships its first draft: half-remembered statistics, AI-sounding phrasing, nothing checking it. This never ships below 90 out of 100: vetted sources only, every cited link opened and confirmed, AI-tell patterns screened out, and up to three self-revisions before you see anything. You are never the first reviewer. Connects to: works from the Marketing Brain's plan and the Website Brain's voice, with its own bundled content memory.

■ IN CONVERSATION

Write the article on AI marketing agents.

Thinking... Reading your brand file and voice profile. Confirming no existing page owns this topic. First draft scored 87, below the 90 bar. Revising and rescoreing.

Delivered after two rounds at 93 out of 100. Every statistic was checked at its source, and it's formatted to rank on Google and get quoted by AI. You're never the first reviewer; the quality bar is.

BLOCKS BELOW 90/100

SOURCES VERIFIED AT THE LINK

AI-DETECTION SCREENED

31 CAPABILITIES

252 AUTOMATED TESTS

YOUR FILES STAY LOCAL

■ BY THE NUMBERS

31 capabilities, 12 templates, 5 specialists, 252 automated tests. A 100-point grade with a hard floor at 90. Everything runs locally; your brand files never leave your machine. Signal it is built on: pages cited by AI Overviews earn about 120 percent more clicks per impression.

Ads Team



Team 03 visual
or [browse files](#)

● ADVERTISING

■ WHAT IT IS

A paid-ads department across eight platforms (Google, Meta, YouTube, LinkedIn, TikTok, Microsoft, Apple, Amazon): it audits accounts, cuts waste, keeps you inside the law, and generates on-brand creative.

■ HOW IT'S BUILT

Seven specialists audit in parallel (Google, Meta, creative, tracking, budget, ad policy, privacy law) against a catalog of 300 named, test-verified checks, weighted by severity with one consistent formula. A dedicated specialist tracks the five hard regulatory clocks of 2026, including the EU AI Act's August baseline and December watermarking deadlines with penalties up to 15 million euros or 3 percent of turnover. A four-person creative crew then generates images and copy from your saved brand profile, with a quality gate that auto-regenerates anything scoring under 6 out of 10. Platform claims are verified and honestly labeled; vendor-supplied case studies get flagged as vendor-supplied.

■ WHERE YOU'D USE IT

Agencies with 5+ accounts: audit every account weekly instead of once a quarter, same time budget. In-house teams on 4+ platforms: a second pair of eyes before exec reviews, no human bias. Freelancers: win the engagement before you spend an hour on diagnostics.

■ IN CONVERSATION

Thinking... Reading your business profile: e-commerce, sales goal. Comparing with last month's audit: one issue still open. One ad set breaks the 3x rule: triple your target cost, zero sales.

Pause that ad set today; it's spending three times your target with nothing to show. And the issue from last month's audit is still open: brand and non-brand traffic are mixed together. Separating them is this week's quick win, worth about \$420 a month.

How are the campaigns doing?

300 VERIFIED CHECKS

8 PLATFORMS

AUDIT IN 10-15 MIN

READS BEFORE IT EVER WRITES

2026 LAW CLOCKS BUILT IN

SECURITY-AUDITED 94/100

■ BY THE NUMBERS

A single-account audit in 10 to 15 minutes, versus 4 to 6 hours of senior PPC time, 1 to 2 weeks and \$2k to \$10k+ from an agency, or \$99 to \$799 a month for tools. 300 verified checks (Google 95, Meta 72, LinkedIn 46, TikTok 46, Microsoft 41). 11 specialists.

FALLBACKLABS.COM Independently security-audited at 94/100.

FALLBACKLABS.COM

04

Email Team



Team 04 visual
or [browse files](#)

● EMAIL

■ WHAT IT IS

A complete email department: it writes campaigns and sequences, scores them before they send, and verifies your emails can actually reach the inbox. It never sends anything without your say-so.

■ HOW IT'S BUILT

Six jobs (triage, writing, review, deliverability, sequences, strategy) with four specialists behind them. It runs real checks on your domain's authentication records and the actual blocklists, then keeps two 0-to-100 scores: one for whether your setup can land in the inbox, one for whether the email is any good. Copy follows proven formulas (PAS, AIDA, and peers), and the thresholds are the inbox providers' own: Gmail's 5,000-a-day bulk-sender rules, the 0.30 percent spam ceiling, and one-click unsubscribe. A hard gate physically blocks a send that is missing legally required pieces, and volume is capped without your approval.

■ WHAT IT DRAWS ON

Real, live checks of the behind-the-scenes records that decide inbox versus spam. The inbox providers' own 2024 to 2026 sender rules (the ones Google began hard-enforcing in late 2025).

■ WHERE YOU'D USE IT

Campaigns that actually land. Sequences that run themselves. A deliverability audit before the big send. Compliance across regions without a lawyer on retainer.

■ WHAT IT PRODUCES

Mobile-ready emails with three scored subject-line options each. Full automated sequences with timing built in. A 90-day plan. A deliverability score with the exact records to fix. A compliance report with the real penalty exposure, up to \$53,088 per email under CAN-SPAM.

The laws that govern email (CAN-SPAM, GDPR, CCPA), with the actual penalty exposure. And proven copy formulas rather than freestyle.

■ WHY IT BEATS A CLAUDE/CODEX PROMPT

A connected prompt writes and sends; it will not run the checks that decide inbox versus spam unless you know to ask, and nothing blocks a bad send. This runs the real record and blocklist checks, scores the copy against proven formulas, physically blocks a send that is missing legally required pieces, and caps volume without your approval. Connects to: pairs with the Email Marketing Brain. The brain plans the program; this team makes it land.

■ IN CONVERSATION

Send the launch email.

Thinking... Reading your profile: software business, EU customers, friendly voice. Checking your domain's inbox setup: one record missing. Checking the email: no unsubscribe link. Blocking.

Held at the gate. It's missing the legally required unsubscribe link, and your domain has one setting hurting inbox delivery; the exact fix is included. The draft is ready with three scored subject lines. Nothing goes out until you say so.

REAL DNS + BLOCKLIST CHECKS

TWO 0-100 SCORES

HARD PRE-SEND GATE

MAX 20 EMAILS WITHOUT APPROVAL

ONE-CLICK UNSUBSCRIBE ENFORCED

DRAFT-FIRST, ALWAYS

■ BY THE NUMBERS

6 jobs, 4 specialists, two 0-to-100 scores. Real checks against the actual records and blocklists. Hard caps: never more than 20 emails in a session without explicit approval. House rule worth framing: design for Apple Mail, test in Gmail, assume Outlook

05

Video Team



Team 05 visual
or [browse files](#)

● VIDEO

■ WHAT IT IS

A video-production department (built on our Walt system) for any creator or brand: it plans the video most likely to perform, packages it title-and-thumbnail first, scripts it to hold viewers, and gets sharper from your own numbers. A producing brain, not an editor.

■ HOW IT'S BUILT

An eight-step process with a gate at every stage: idea, packaging, script, storyboard, produce, audio and captions, quality check and publish, then review and refine. The order is the method: packaging before script, style before production, measure before you refine. It mines your real comments and analytics, works from current 2026 algorithm research sourced from the platform itself (watch time and satisfaction, not just clicks), and draws hooks from a library of 18 proven patterns.

■ WHAT IT DRAWS ON

Your own channel: your real comments and your analytics. Current algorithm research taken straight from the platform's own sources, not stale advice. Keyword demand. And a library of

■ WHERE YOU'D USE IT

Planning a video with a real shot at performing. Fixing hooks that lose viewers in ten seconds. Testing titles and thumbnails properly instead of by feel. Learning from your own channel history instead of generic advice.

■ IN CONVERSATION

Thinking... Reading your channel profile: friendly, concise, no hype. Checking your performance log: this video is inside your normal range. Rule: never change the formula on one video. One data point is noise.

Keep the formula. Against your own channel history this isn't a flop, it's normal variance. Your best-performing hook is already banked in memory, and the next script opens with it.

PACKAGING BEFORE SCRIPT

18 HOOK PATTERNS

TESTS SETTLED AT 95% CONFIDENCE

3-VIDEO REFINEMENT GATE

LEARNS YOUR CHANNEL

RUNS ENTIRELY LOCAL

Yesterday's video flopped. Should we change the format?

■ BY THE NUMBERS

8 steps, each gated. 18 hook patterns. 3 scored titles per video. Tests settled only at 95 percent confidence. A 37-column performance record per video. Zero outside dependencies, runs entirely on your machine.

Straight *answers.*

Is this a template with AI branding on it?

No. Every build gets its own architecture, design, and development, and it ships on infrastructure in your name. Add the Website Brain and it reads and remembers your entire business, every page, product, and promise, so what it does is grounded in who you actually are. The tools are off the shelf: TypeScript, Next.js, Vercel, Supabase. What we build with them is not.

Does it actually run itself, or is that just a demo?

It does, but not on day one and not on everything. Every agent starts in draft mode: it does the real job next to your team, writing what it would send instead of sending it. You see every action and approve it, or you don't. It earns full handover when the drafts stop needing your edits, not before. After that it runs unattended, nights and weekends included. On the work that matters, it stays in draft for good, by design.

What do you actually mean by a "brain?"

A brain is a knowledge base of your business that the software reads before it acts and writes back to after, so it gets sharper the longer it works for you. There are nine, each on one part of the job: website, marketing, gen AI, local SEO, social, email, paid ads, sales, and cybersecurity. The Website Brain holds the memory the other eight build on. The brain structure is ours, licensed for use in your business; what goes into it and everything it learns about your business are yours.

How is this different from just paying for ChatGPT?

It works from your business, not a blank slate. The Website Brain reads and remembers your entire business, every page, product, and promise, and the other eight brains build on that memory. Your data trains your brain only, and we never mix one client's data into another client's systems. We will not tell you what a general subscription can or cannot do, only what ours does.

AI writes the code. How do we know it is any good?

AI writes a lot of it, and that is the point: it is why we build faster. But you are never handed raw model output. Before we hand a build over, a person reviews the code and it goes through testing and QA, so what you get is checked, working software. That review is part of every build we deliver, whatever tier you are on.

How long does it take?

It depends on scope, and we would rather say that than invent a number. Here is our record instead: six products shipped, the fastest in two weeks, the longest in three months. Those

Do we own what you build, or do you?

You own it. The code and the agent are yours, running on infrastructure in your name, with no lock-in, and they keep working if we ever part ways. The brain is the one nuance: its structure is licensed to you and stays Fallback Labs' property, but everything inside it, your data and everything it learns about your business, is yours. That split is the same on every engagement.

were full custom builds, not a typical Launch site, so treat them as the range, not your quote. Bring your scope to a build call and we will give you the real window for it.

Bring your scope to a build call and we will give you the real window for it.

fallbacklabs.com

Establish, build, and grow, *all in one place.*

Bring your scope to a build call and we will give you the real window for it.

WEBSITE

www.fallbacklabs.com

PHONE

+971 55 414 1234

VISIT

in5 Tech, Dubai Internet City,
Dubai, UAE

FOLLOW

@fallbacklabs

Instagram · YouTube · TikTok · LinkedIn