



Alacritous

AUTONOMOUS AGENT ORCHESTRATION

Full auditability, control, and data sovereignty for mid-market companies.
Not tools. Not dashboards. Outcomes.

alacritous.io · Seed 2026

BUILT BY OPERATORS

Jeremy Evans

15 years building revenue machines in B2B lead-gen and RevOps. We built an agency ecosystem around omnichannel leadgen, now fully automated via AI.

Thrice Agency

AI-first RevOps

AGENCY HUB

LinkedGuerilla

Automated LinkedIn appt setting

LEADER SERVICE

AIOutreach

Personalized messaging at scale

SAAS FOR OUTREACH AGENCIES

Naturally.io

Autonomous lead nurturing

SAAS FOR SALES DEPTS

VCNetworker.com

Connect with VCs for raises

SERVICE FOR SERIES A+ FOUNDERS

Neckbeard.agency

Reddit leadgen based on intent

PRIVATE SERVICE FOR GTM

Alacritous started as a means to automate omnichannel ABM + appt setting, now just one of many features.

First clients: LandVoice & The Good Brand Company — marketing began W1 Feb 2026.

\$100k

MONTH-1 REVENUE FOR A PPC AGENCY CLIENT

82

AVG QUALIFIED APPTS/MO PER CLIENT AT \$6K SPEND
— BEFORE AI MULTIPLIED THAT BY 2X

\$10M+

RAISED FOR CLIENTS VIA VCNETWORKER

This isn't a thesis. I ran the playbook on my own P&L first.

Performance Baseline

A 30-DAY CROSS-SECTION OF OUR OWN P&L

- We replace a material % of ops coordination for mid-market teams
- We deliver **new revenue + lower cost** in a **30-day pilot**
- Proof — our own ops, 30 days:

\$16k → \$6.5k

OPEX — 59% REDUCTION

18 → 37

APPOINTMENTS / MONTH

This is the only slide where we stack numbers. How it works comes next.

What It Is

AUTONOMOUS AGENT ORCHESTRATION

- **Fully autonomous ops** for any third party tooling: email, SaaS, CRM, ERP, etc
- **Immutable audit trail** — every action logged, every denial captured
- **Data sovereignty** ensuring your proprietary information never leaves your systems
- **Human-as-Approver** consent gates for high-stakes actions
- **GraphRAG** institutional memory — context that compounds

Teams of 5-25 · \$3,000/mo flat · No per-seat pricing

Why Now

**Everyone is building AI agents.
Almost nobody is making them work inside real businesses.**

11.5%

SBA 7(A) RATES — CHEAP MONEY IS OVER

1.5M+

B2B MID-MARKET BUSINESSES SEEKING EFFICIENCY

74%

OF ORGS LACK AI GOVERNANCE

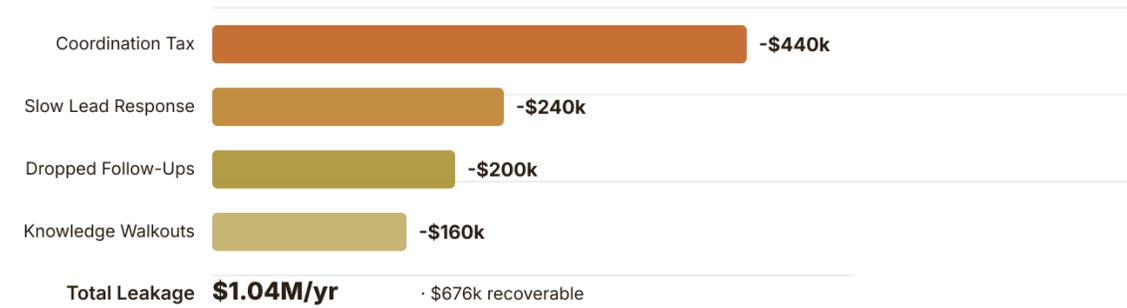
The SaaSocalypse proves per-seat pricing is dead.
Capital is expensive. Hiring is a trap. **Efficiency is mandatory.**

The window: True integration into a business will be bespoke until actual AGI.
We ship pipeline and ops autonomy, not just software.

[Cisco 2026](#) · [ESG Research](#) · [SBA current rates](#)

The Silent Bleed

A \$2M company loses 30–50% of potential revenue to failures nobody sees on the P&L:



Tools are stateless. Operations need memory.

[HBR \(lead response\)](#) · [McKinsey \(coordination\)](#) · Alacritous leakage model

The Market

AI AGENTS ARE SCALING RAPIDLY

\$10.9B

GLOBAL AI AGENTS MARKET (2026) TAM

40%

OF ENTERPRISE APPS EMBEDDING AGENTS

SAM: Mid-market businesses scaling workflows.

SOM: 10,000+ US marketing agencies (\$5M+ revenue) rapidly adopting AI.

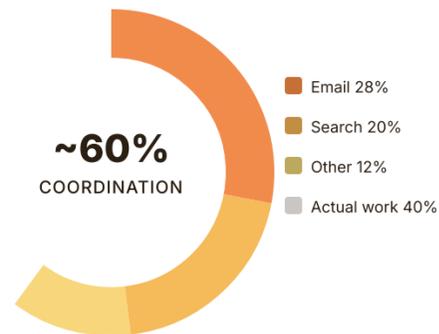
SOM selected based on [rapid AI adoption at marketing agencies](#) and Founder's experience selling into these markets as an agency owner himself.

Agencies using AI agents see an average 171% ROI .

[PlanetaryLabour \(ROI & Market Size\)](#) · [Gartner \(Adoption\)](#) · [BCG \(Marketing SOM\)](#)

The Hidden Cost

WHERE A 12-PERSON TEAM'S WEEK ACTUALLY GOES



\$4,875

LOST THIS WEEK ALONE (12-PERSON TEAM)

\$234k/yr

VANISHING INTO GLUE WORK

You can feel it in Slack, not in QuickBooks.

[McKinsey Global Institute](#) · Alacritous economics model

Big Idea

AUTONOMOUS AGENT ORCHESTRATION.

An MSP for autonomous AI ops — low-cost, high-impact deployment from day one.

Full auditability, control, and data sovereignty — **without compromising speed.**

- Bespoke to your workflows, not a template
- Audited and permissioned — every action governed
- Tuned by operators responsible for results

True integration into a business will be bespoke until actual AGI. That creates an industry, **Managed Agentic Providers.**

We're building the platform for it.

Year 1-2: Managed service — high-touch, learn the wedge.

Year 3+: Self-serve Hub platform where clients manage their own Alacritous instances.

The managed service is the GTM wedge that builds the platform. Customers stay for the knowledge graph, even when they move AI talent in-house. If SaaS survives, we build the best-informed self-serve product with a plurality of captured market share. If not, the managed service stands on its own.

OUR DIFFERENCE IS TWOFOLD.

1. THE MID-MARKET REALITY

Enterprise won't hand sensitive data to contractors — and won't tolerate inference-based data leakage.

SMBs will get cheap, good-enough AI from the model providers themselves (Anthropic, OpenAI). No room for us.

Only Mid-Market has the budget, the urgency, and the operational complexity to pay for bespoke AI ops — and our deployment SOPs get them live by W2.

2. THE BESPOKE IMPERATIVE

Mid-market business will require **bespoke modifications and management** to fit existing internal ops paradigms.

Off-the-shelf wrappers fail here. Any SaaS would require a responsible party in-house for outcomes, we bundle that.

The Pilot

\$1,250 · 30 DAYS · ONE REAL APPOINTMENT



- **CRM-Agnostic Automation:** Syncs perfectly with any CRM, notifies the sales team in their chat service, and automatically drafts or sends follow-ups.
- **Founder-Led Advantage:** 15 years selling leadgen and revenue systems gives us an edge selling a revenue-first pilot that is less invasive than pitching reduced opex from day one.
- **Revenue-First Wedge:** A booked meeting is immediate commercial proof, so the pilot is easy to approve, fast to validate, and naturally expands into deeper operational automation after trust is earned.
- **Negative CAC:** Using a profitable pilot at low cost delivers early value and ensures CAC is never a concern.

Early traction: LandVoice and The Good Brand Company onboarded W1 Feb 2026.

EXECUTION

The Math of GTM

A PREDICTABLE ACQUISITION ENGINE

60

LINKEDIN & REDDIT AVATARS

24k

OMNICHANNEL SENDS / MO

24

NEW PILOTS / MO

At a conservative 1% meeting conversion and 10% close rate.

BUDGET ALLOCATION	FOCUS
50% Marketing	Guerilla scaling (\$4k GTM budget yields 24 pilots)
20% TAMs	Technical Account Managers for onboarding & upsells
20% Founder & 10% Dev	Hub development for instance management

Unusual GTM: Guerilla outreach is our agency specialty and scales linearly by adding more outreach capacity; PPC is a support channel, not the dependency. Happy to unpack the full playbook.

Business Model

Managed service → accuracy compounds → automation scope creeps into overages.

Selling into a new industry, we believe the best provider will perform customer education without big budgets: profit-first pilots, platform licensing fees, and becoming "the AI guy" that customers go to when the latest model drops.

TIER	PRICE
Pilot	\$1,250 one-time
Founder	\$3,000/mo + \$150/h T&E overages
Enterprise	Custom scoped

3 FTEs
\$255k
traditional cost

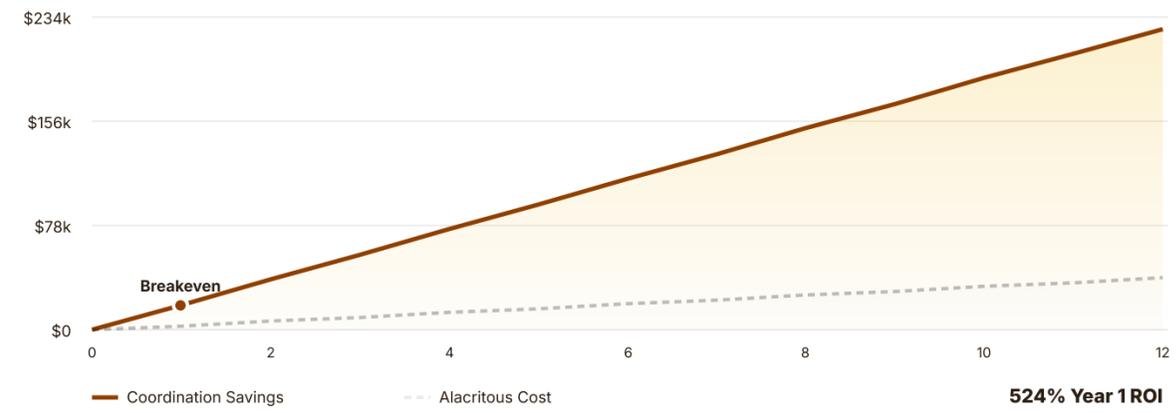
VS

Alacritous
\$36k
flat annual

\$100/day for a 24/7 AI operations team.

12-Month ROI

COORDINATION SAVINGS VS. ALACRITOUS COST



12-person team at \$65k avg salary · \$19,500/mo coordination tax · Month 1 payback

UNIT ECONOMICS

Month 3 Snapshot

CAPITAL EFFICIENCY AT SCALE

The Engine (Monthly):

- pilots at \$
- % convert to \$ /mo
- Setup cost: \$ / pilot
- Fixed Opex: \$ /mo
- Exp. Overage: \$ /mo

\$73.2k
M3 GROSS REV (\$30.0K PILOT + \$43.2K MRR)

-\$40.7k
M3 EXPENSES (\$18.0K SETUP + \$20.0K FIXED + \$2.7K MRR COGS)

\$32.5k
M3 NET PROFIT — 44% MARGIN

Projected Cash in Hand (M3):

STARTUP BENCHMARKS

Growth Metrics

HIGHLY CAPITAL-EFFICIENT GROWTH

-\$3,611

NET CAC / ACTIVATED ACCOUNT

\$8.0k

PILOT ACQUISITION SURPLUS / MO

80/72/108

LOGO / GRR / NRR (%)

94% / 88%

CONTRIBUTION / RAMP MARGIN

Acquisition Math:

- \$ **4000** GTM Spend → **24** Pilots
- Pilot Price: \$ **1250** per pilot
- Setup Cost: \$ **750** per pilot
- **30** % Pilot Conversion → 7.2 MRR Clients
- \$4,000 / 7.2 = \$556 gross acquisition / activated account
- *Excludes \$30.0k front-end revenue from 24 Pilots.
True Net CAC is effectively negative.

Service Economics:

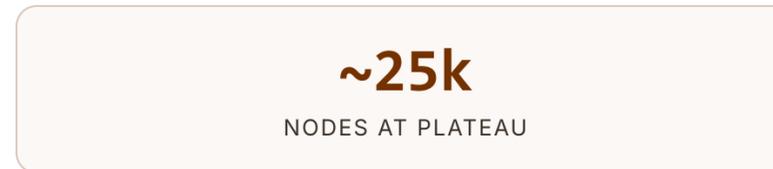
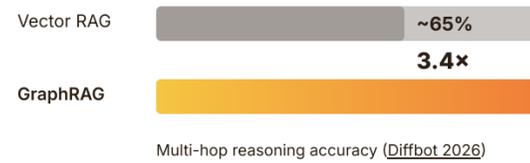
- TAM: \$ **6000** / **32** Accts = \$187.50
- \$ **3000** MRR - \$187.50 = \$2,812.50 Profit
- **94%** Contribution margin on recurring service revenue
- \$2,812.50 × **6** mo = \$16,875 LTV
- Retention assumptions: **80** % renewal / **10** % contraction / **50** % expansion
- Ramp months: **3**
- \$16,875 LTV / \$556 CAC = 30.4x Ratio

Managed AI service provider KPIs: **55% YoY** growth quality · **Immediate (pilot-funded)** payback · **98%** GTM repeatability

The Context Moat

INSTITUTIONAL MEMORY THAT COMPOUNDS

Standard AI has amnesia. We build a persistent knowledge graph from every touchpoint — Slack, email, CRM, project updates.



Entity resolution · Memory decay · Nightly pruning

Our moat is the 25,000-node knowledge graph we build for each client. It takes 90 days. No competitor ships with it. Ripping it out means starting over.

Competitive Map

WHAT WE'RE NOT

ALTERNATIVE	LIMITATION
Zapier / Make	Rigid, stateless triggers — break when conditions change
Salesforce Agentforce	Locked in the CRM walled garden
Microsoft Copilot	Favors M365 suite — not vendor-agnostic
OpenClaw	Insufficient memory solution for business, requires technical hire, CISO nightmare
Internal AI hires	Expensive, slow, hard to audit at scale

We win: **OUTCOMES + GOVERNANCE + MEMORY** ship together.

Vendor-agnostic · Protocol-native · Non-custodial · Outcome-responsible

STOP TYPING. START APPROVING.

The app is disappearing. You don't open a CRM —
you message the Sales Agent in Slack.

Your role shifts from Typist to Editor-in-Chief. Reclaim 30% of your day.

95%

ACCURACY WITHIN 60 DAYS

99%

ACCURACY WITHIN 90 DAYS

Accuracy compounds. The system gets better the longer it runs.

Two Loops

HOW EVERY ACTION FLOWS

AUTONOMOUS PROCESSES



HUMAN PROCESSES



**Agents are spawned with full Skill context and minimal viable toolsets (partial MCPs) to improve latency and adherence to requirements.*

The Orchestrator: Reviews memory before creating subagents, can spawn parallel agents or execute data-gathering rounds prior to delivering results, and immutably updates memory.

Approval Scope: Controlled at the admin level (via web portal or managed by us). Designed to prove which processes can go fully autonomous over time.

GOVERNANCE

The Control Plane

One place to monitor and approve everything your AI agents do.

- **Consent Gates** — high-stakes actions require human approval
- **Immutable Audit Trail** — what was accessed, proposed, approved, and when
- **Non-Custodial** — your data stays in your systems
- **Skills Versioning** — hash-pinned like a git commit

SOC 2 ready · EU AI Act compliant · HIPAA clean rooms · RBAC + sandboxing · Cloud or self-hosted

The Digital Workforce

DYNAMIC ORCHESTRATION, NOT RIGID ROLES



Skills, Not Data Dumps
STRUCTURED SOPS · AUDITABLE



Dynamically Spawned
EXACT TOOLING NEEDED



MCP-Native
60+ CONNECTORS



Inter-Agent Synergy
REASON TOGETHER

168 hrs/wk

AUTONOMOUS COVERAGE

<5 min

RESPONSE TIME, 24/7

1,500+

CODIFIED SKILLS · 15 DOMAINS

How This Fails



Agents make mistakes in high-stakes workflows



Security / compliance blocks deployment



Adoption fails — another tool nobody uses

These are the three real risks. Here's how we fence each one.

Fences

RISK	FENCE
Agent mistakes	Human-as-Approver gates + immutable audit trail; mistakes drop sharply after a 1-week memory-ingestion sprint and keep improving over time.
Security blocks	Non-custodial · self-host option · RBAC · sandboxing · per-request policy · SOC 2 ready
Adoption failure	Zero UI — lives in Slack. No dashboards to abandon. 72% dashboard abandonment is their problem, not ours.

Each fence is a design choice, not a roadmap item.

Seed Round

\$500K

Raising \$500k to dominate Autonomous Agent Orchestration by 2027.

Assumptions (from calculator defaults):

- 24 pilots/mo
- 30% pilot conversion
- \$1250 pilot, \$3000 MRR + \$750 avg client monthly overage
- \$4000 GTM/mo, \$750 setup/pilot (now impacts growth)
- TAM \$6000 / 32 accts
- Service retention: 80% renewal, 10% contraction, 50% expansion
- Ramp to steady-state: 3 months
- K-mult 42
- Log a 0.17 to 14, Gom b 3.6 c 0.145
- Piecewise r1/r2/r3: 11 / 6 / 3 %
- Bass p 2.8%, q 40%
- Reinvest 30%, spend lift 1.0x, CAC decay 4%, setup learning 12%
- Exit multiples: 5x to 10x ARR

Projected Exit Horizon (non-linear):

MODEL	M48 ARR	EXIT RANGE	EXIT (K/M)
Logistic S-curve	\$128,395,735	\$641,978,673-\$1,283,957,345	\$642.0M-\$1284.0M
Gompertz curve	\$128,353,012	\$641,765,062-\$1,283,530,125	\$641.8M-\$1283.5M
Piecewise scaling	\$29,004,314	\$145,021,568-\$290,043,137	\$145.0M-\$290.0M
Cohort unit economics	\$99,594,887	\$619,777,927-\$1,239,555,853	\$619.8M-\$1239.6M
Bass diffusion	\$128,792,312	\$643,961,562-\$1,287,923,124	\$644.0M-\$1287.9M

Average model output: \$102,828,052 ARR · \$514,140,260-\$1,028,280,520 EV

Exit window: 48 mo Preset: Base Probabilities D/B/U: 25 / 50 / 25 %

Probability-weighted EV: \$1,073,099,675-\$1,934,322,683 (\$1073.1M-\$1934.3M)

Logo retention / GRR / NRR 80% / 72% / 108%	CAC payback Immediate (pilot-funded)	Contribution margin 95% → 95%
Ramp margin 88% → 91%	ARR growth quality 43% YoY	GTM repeatability 99%

Assumption sources: [SaaS Capital retention benchmarks](#), [ChartMogul retention report](#), [Contraction benchmark context](#), [Expansion benchmark context](#), [Managed-services 30-90 day onboarding baseline](#), [BVP AI monetization playbook](#).

Month 6

116 active clients
\$491k MRR run-rate
TAM recruitment streamlined

Month 12

228 active clients
\$1.1M MRR run-rate
Hub v1 launched

Month 18

Series A readiness path
80% / 72% / 108% retention stack
\$1.9M MRR run-rate (68% of M24)

Capital fuels the Hub, 60-avatar scaling, and TAM hires. The math is scenario-based, not single-point optimism.



The coordination tax runs every week.

The payback period is weeks — not quarters.

**That's not a pitch.
That's arithmetic.**

alacritous.io