velociti

The AI Product Operating System™

Build the RIGHT product, every time.

AI made building faster. It didn't make knowing what to build easier.

PRE-SEED+ · Q4 2026

www.velocitipm.com

THE PROBLEM

Software product development is broken.

75%

of all software products fail (all size and stage companies)

$1T+

spent annually on software product development

THE REAL FAILURE MODE

Products don't fail because we can't build good software. They fail because we build the wrong thing.

Most of the $1T+ goes into building. Not enough goes into deciding what to build.

THE ROOT CAUSE

It was never a “how” problem. It’s a “what” problem.

Discovery is only the start. The answer to “what should we build, and why” is rarely captured, never kept current, and lost before it reaches the people and agents who build.

THE “WHAT” IS FOUR QUESTIONS EVERY BUILD DEPENDS ON

WHO

Who are we building for?

Personas and segments grounded in real customer evidence, not assumptions.

WHAT

Which problem matters most?

Problems ranked by evidence and impact, not the loudest voice in the room.

WHY

Why this, and why now?

A clear line from every initiative to strategy and business outcomes.

PROOF

How will we know it worked?

Success defined up front, with outcomes flowing back into the next decision.

WHAT HAPPENS WITHOUT IT

Discovery debt

The compounding cost of skipping discovery. Every unvalidated bet is harder to undo.

Perpetual build season

Teams never leave build mode long enough to ask if it’s the right thing.

The vibe coding trap

Building fast with AI toward a target nobody validated.

90%+ of PM time goes to reactive work: roadmap fires, stakeholder requests, status updates

80% of shipped features are rarely or never used. Most of what teams build misses.

Sources: Velociti customer research (90%+); Pendo, Feature Adoption Report, 2019 (80%). Concepts: Velociti, “discovery debt,” “perpetual build season,” “the vibe coding trap.”

© 2026 | VELOCITI

THE SHIFT: CODING IS A COMMODITY

AI solved coding. Coding was never the problem.

AI made writing code a commodity, but it fixed a problem the industry didn't have. The 75% failure rate was always about "what."

HOW TO BUILD

COMMODITIZED

90% of developers now use AI at work (DORA, 2025)

42% → 65% of committed code is AI-generated today, and expected by 2027 (Sonar, 2026)

Same tools Cursor, Claude Code, Codex, Copilot: every competitor has the same models and agents. Speed is table stakes.

WHAT TO BUILD

STILL UNSOLVED

75% of software products failed before AI coding existed. Teams could always build; they built the wrong thing.

80% of shipped features are rarely or never used (Pendo, 2019)

No context Agents know your codebase. They don't know your customers, your strategy, or why a feature matters.

The industry never had a coding problem. AI made the real one worse: teams now ship the wrong thing faster than ever.

Sources: Google Cloud DORA, State of AI-assisted Software Development, 2025 (n≈5,000); Sonar, State of Code Developer Survey, Jan 2026 (n=1,100+); Pendo, Feature Adoption Report, 2019.

© 2026 | VELOCITI

WHY IT'S URGENT

AI is pushing the failure rate from 75% to 85%+.

Faster building didn't fix the "what." It amplified it: more bets, shipped faster, and the misses found later.

SOFTWARE PRODUCTS THAT FAIL

BEFORE AI CODING	AI ERA (EARLY DATA)
75%	→ 85%+

When building took months, a bad bet cost a quarter. When building takes an afternoon, teams make a dozen bad bets before anyone sees the pattern.

That is discovery debt. AI didn't reduce it; it made it cheaper to take on.

THE EVIDENCE

81%	of enterprise tech leaders see more production issues from AI code *CloudBees, State of Code Abundance, 2026 (n=200+)*
79%	say agents write code much faster, yet release cycles are no faster *Undo / Coleman Parkes, 2026 (n=300)*
96%	of developers don't fully trust AI code is correct; only 48% always verify it *Sonar, State of Code Developer Survey, 2026 (n=1,100+)*
66%	of developers are frustrated by AI output that's "almost right, but not quite" *Stack Overflow Developer Survey, 2025 (n≈49,000)*
95%	of enterprise GenAI pilots show no measurable P&L impact *MIT NANDA, The GenAI Divide, 2025*

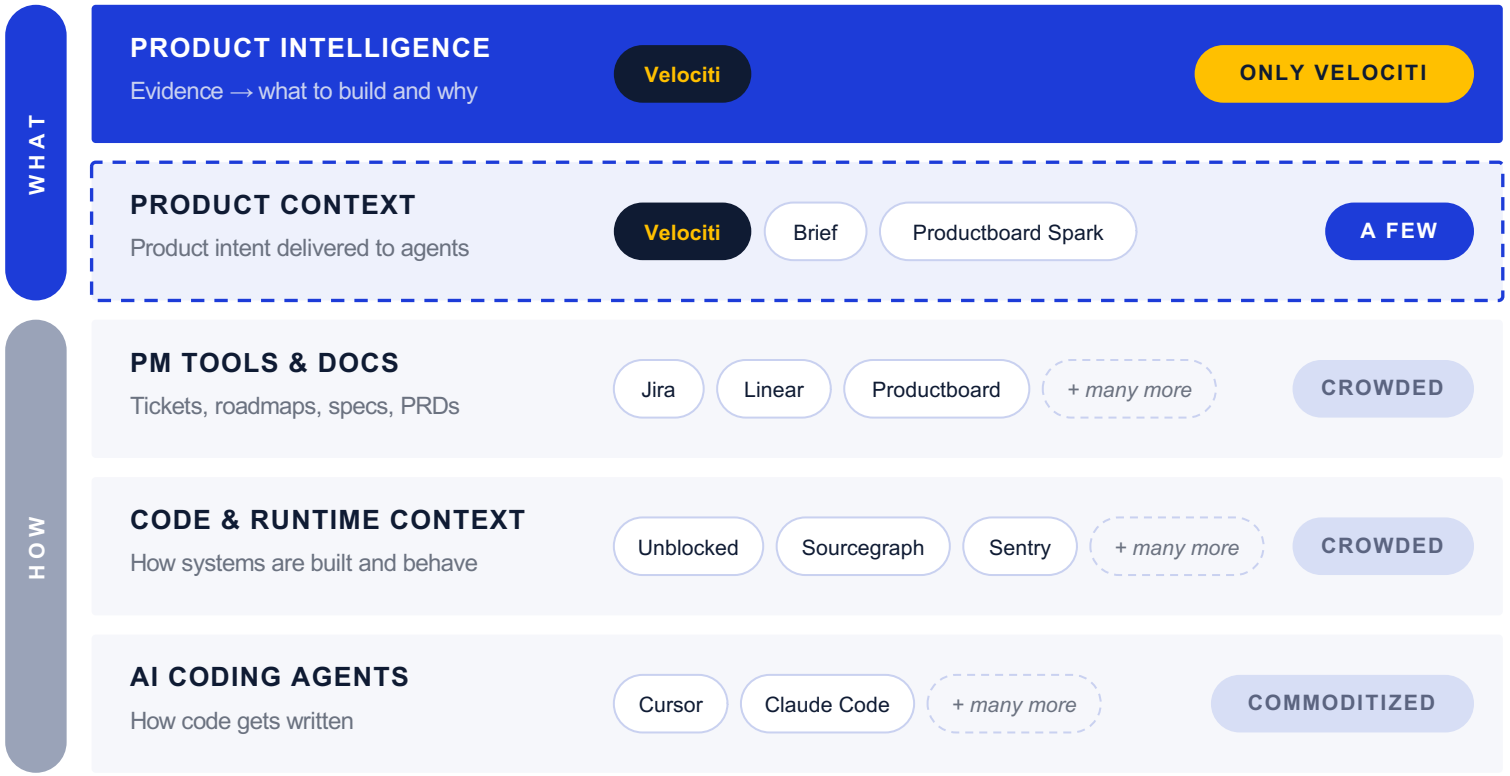
75% is the widely cited pre-AI industry failure rate. 85%+ is Velociti's directional estimate based on the AI-era evidence shown; see notes.

© 2026 | VELOCITI

THE MISSING LAYER

A crowded market, except where it matters.

Hundreds of companies compete on how software gets built. Only a few work on what should be built, and only Velociti generates it from customer evidence.



WHY "CROWDED" IS THE WRONG READ

The crowd is in "how."

Agents, code context, and ticketing are crowded and commoditizing. The "what" layer is nearly empty.

Context ≠ intelligence.

Context tools store decisions already made. Velociti generates the decision from customer evidence.

One layer, every agent.

Product intelligence sits above every coding agent and reaches all of them over MCP.

Assessment based on public product documentation, Oct 2026. MCP server count: Bloomberry analysis of 1,400 MCP servers (2026). Names shown are representative, not exhaustive.

© 2026 | VELOCITI

THE SOLUTION

Velociti is the AI Product Operating System™

One connected system, the Context Suite, turns customer evidence into product intelligence and agent-ready context.

EVIDENCE IN

- Interviews
- Sales calls
- Support & feedback
- Product analytics
- Experiments
- Market research

›

1

Context Layer

Persistent product intelligence that compounds and never resets

›

2

Context Loop

Discovery → Context → Strategy → Build, on repeat

›

3

Context Dispatch

Ready-to-build context sent to any coding agent via MCP

›

ANY AI AGENT

- Cursor
- Claude Code
- Codex
- Your agent

10 minutes

to run foundational discovery inside Velociti

Traceable

from strategy to shipped code, with no handoff lost

Agent-agnostic

works with whatever your team codes with

© 2026 | VELOCITI

HOW IT WORKS

Follow one customer signal to shipped code.

Every step is connected and traceable, so the “why” survives all the way to the agent, and the result flows back.

1

SIGNAL IN

Evidence captured

Signals extracted from interviews, calls, tickets, and analytics.

EXAMPLE

“We gave up during data import.”
14 interviews · 32 tickets

2

CONTEXT LAYER

Problem validated

Signals cluster into a problem, tied to a persona and its evidence.

EXAMPLE

Problem: activation stalls at import
Persona: Ops lead · 46 signals

3

CONTEXT LOOP

Bet prioritized

Ranked against strategy and outcomes, not opinion.

EXAMPLE

Initiative: guided import
Goal: activation +15% · Rank #1

4

CONTEXT DISPATCH

Sent to any agent

Requirements plus the “why” go to any coding agent via MCP.

EXAMPLE

→ Claude Code: requirements, acceptance criteria, evidence links

5

OUTCOMES

Result fed back

Shipped results flow back and sharpen the next decision.

EXAMPLE

Activation +11% measured
Learning added to Context Layer

WITHOUT VELOCITI Evidence scattered across tools, strategy in a deck, a vague prompt to the agent, and no way to tell if it worked.

The agent will run the loop either way. The question is what it’s running toward.

Illustrative example. Live demo available: youtu.be/KBj-IeFOOaA

© 2026 | VELOCITI

SIGNAL

Real signal. Virtually zero marketing spend.

$11K+
MRR, with virtually no paid marketing

~$1M
qualified pipeline, and growing

100+
paying customers

95%
customer retention since launch

2,500+
trial signups since MVP launch

Enterprise
pilots underway after moving upmarket to sales-led

The “aha” arrived with the full Context Suite: once discovery-to-build shipped end to end, initiative creation, session depth, and engagement jumped materially.

© 2026 | VELOCITI

VALIDATION

Validation from every direction.

MARKET PULL

YC asked for it

Y Combinator's Request for Startups called for a "Cursor for Product Managers." That's what we built.

COMPETITION

3× pitch finalist

Mayfield AI Garage (top 3%), Thomson Reuters Momentum AI (1 of 5), and Arizona Innovation Challenge (1 of 15).

CAPITAL

Pre-seed+ round

Anchor VC committed. Closing the round in Q4 2026.

RESEARCH

Academic thesis match

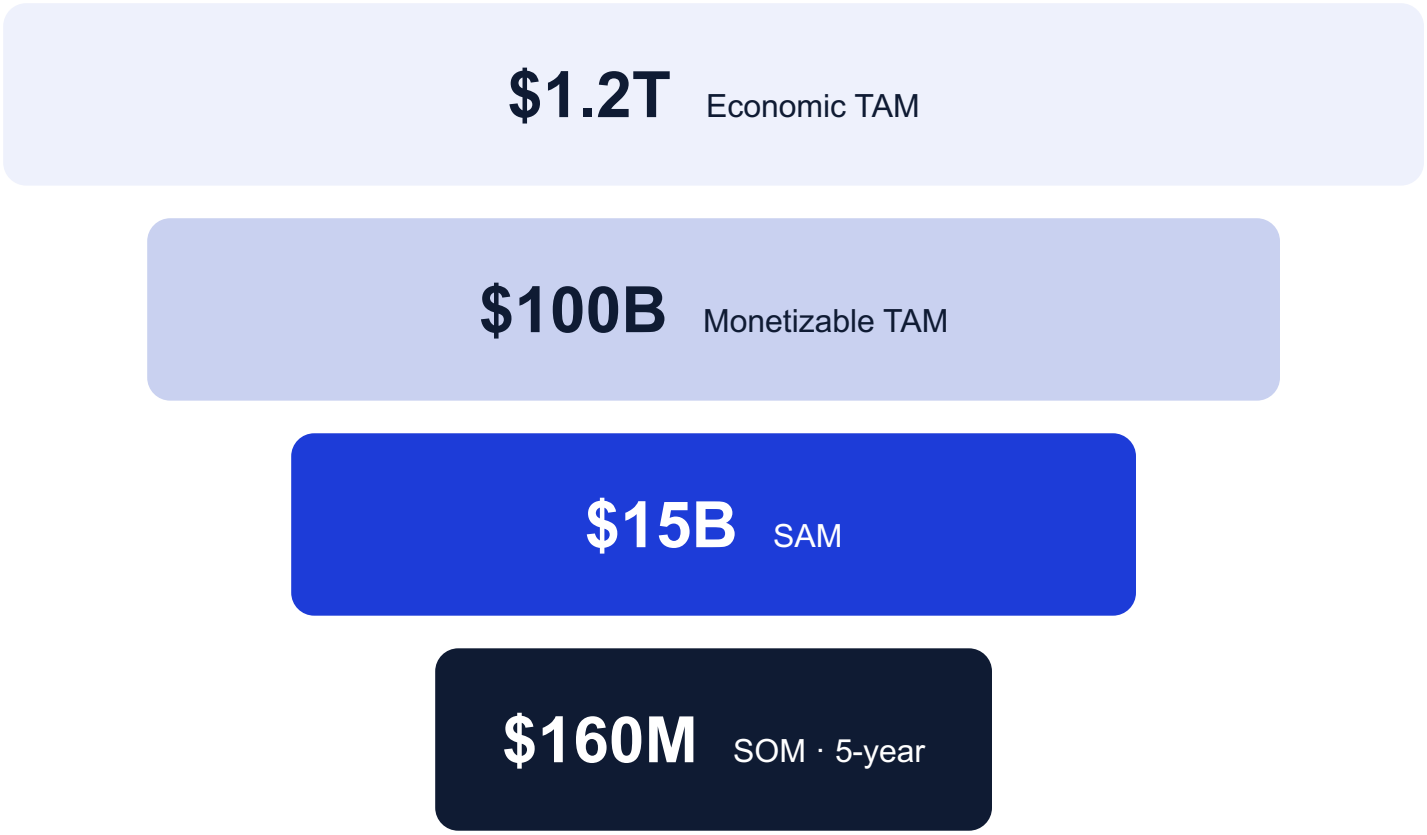
NYU's J.P. Eggers argues in HBR that AI makes building easy and choosing what to build harder. A joint study is in discussion.

WHY NOW

AI coding compresses build cycles from months to hours. The cost of building the wrong thing just went up 10x, and the team that owns the product-context layer owns the most valuable part of the AI product stack.

MARKET SIZE

A $1.2T spending problem. A $100B software opportunity.



HOW WE SIZED IT

Bottom-up: 36.5M developers × $85K average cost × 40% of effort on product work = $1.2T in annual product-development spend.

WHY IT'S BIGGER THAN "PM TOOLS"

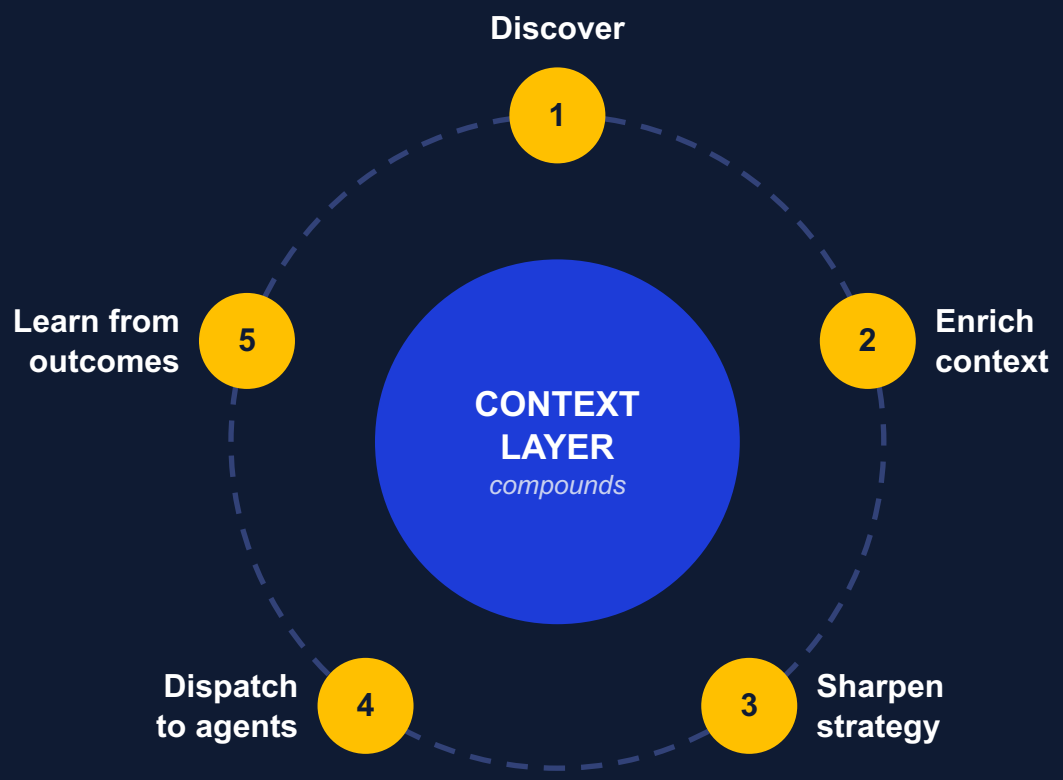
We don't compete for the PM-software line item. We sell against the cost of building the wrong thing, the largest line item in software.

ACV grows ~10x ($5K → $48K) as customers move from teams to enterprise.

Source: Velociti Addressable Market 2026. Full model, sensitivity ranges, segment ACVs, and sourcing available in the data room.

OUR MOAT

A moat that compounds with every cycle.



Proprietary, compounding data

Each customer's Context Layer becomes their system of record for product intent. It gets smarter every cycle and never resets.

Switching costs that grow over time

In month one, a team adopts a tool. By month twelve, leaving means abandoning their product's institutional memory.

Sharper output, every loop

Outcomes flow back into discovery, so each recommendation and dispatch is better than the last.

Discovery feeds the loop. The loop closes the gap. Velociti runs both.

© 2026 | VELOCITI

DEFENSIBILITY

Four layers of defensibility.

Why this is hard to copy, whether you're an incumbent, a frontier lab, or a team building it in-house.



Compounding Context Layer

Proprietary product intelligence that deepens with every cycle. It's data, not a feature, so it can't be cloned from a demo.



Encoded methodology

25 years of product discovery practice, built into agents. Customers run a proven system, not a blank prompt.



Neutral position in the stack

Agent-agnostic by design. A coding vendor's context serves only its own agent. Teams want one context layer for every agent they use.



Misaligned giants

Frontier labs monetize tokens and output. Velociti reduces rework and wasted generation, a problem they aren't incentivized to solve.

© 2026 | VELOCITI

COMPETITIVE LANDSCAPE

Others are moving toward “what.” Only we make it continuous.

	Velociti	Productboard (Spark + MCP)	Spec-driven dev (Spec Kit, Kiro)	Context tools (Brief)	DIY AI stack (LLM + docs)
Automated discovery from customer evidence	●	○	○	○	○
Persistent, continuously updated context layer	●	◐	◐	●	◐
Strategy & prioritization tied to outcomes	●	◐	○	◐	○
Requirements traceable to evidence	●	◐	○	◐	○
Dispatch to any AI coding agent	●	●	●	●	◐
Outcomes flow back into discovery	●	○	○	○	○

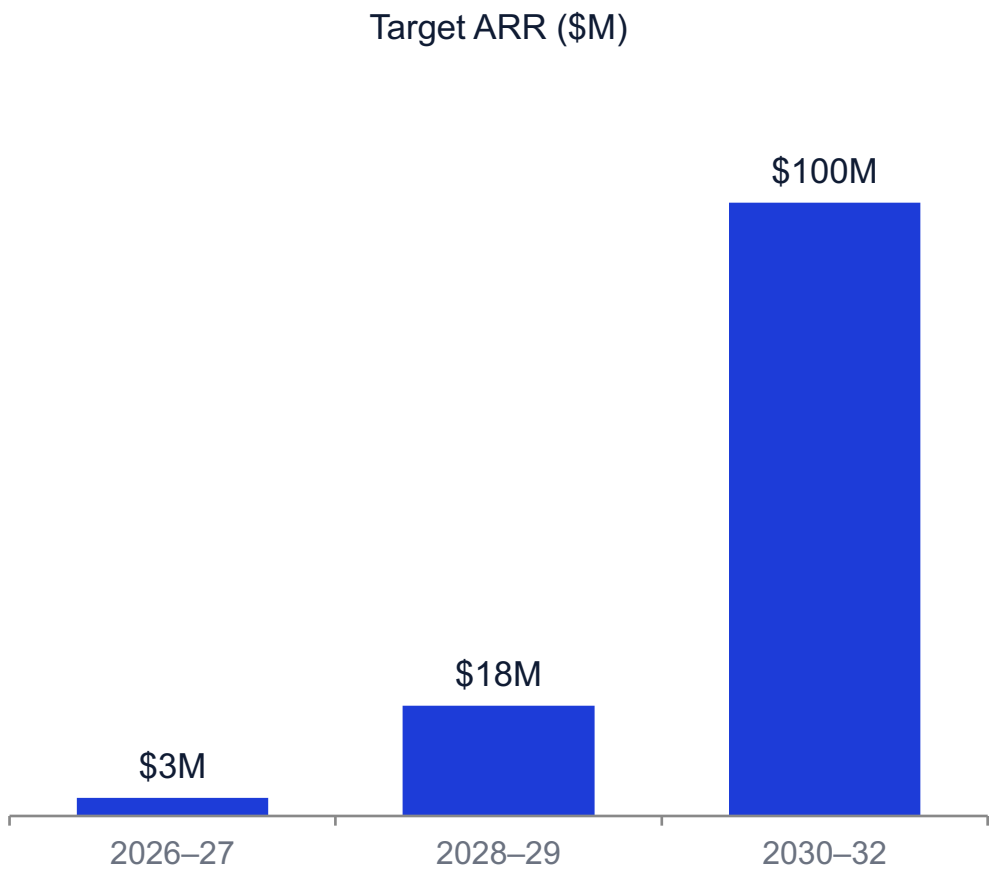
● Full ◐ Partial ○ None

Assessment based on public product documentation and release notes as of October 2026. Productboard shipped its MCP server for coding agents in June 2026; Brief captures decisions from Jira, Linear, Slack, and GitHub.

© 2026 | VELOCITI

BUSINESS MODEL

Path to $100M ARR.



2026–27

$3M ARR

650 customers
ACV $5K

Founder-led sales, first AE and Customer Success hires; convert pipeline and enterprise pilots.

2028–29

$18M ARR

1,500 customers
ACV $12K

Own the mid-market; expand accounts as the Context Layer deepens.

2030–32

$100M+ ARR

2,000 customers
ACV $48K

Enterprise sales plus PLG; own the AI Product OS category.

ACV grows ~10x as the Context Layer deepens and we move from teams to enterprise.

© 2026 | VELOCITI

LEADERSHIP TEAM

We've lived the pain our users feel.

We've built products that generated billions in revenue, managed hundreds of PMs and engineers, and talk to our ICP every day.



Anthony Argenziano

Founder & CEO

25 years building products, teams & companies
Amazon, eBay, Intel, 10+ Startups



Alexiz Hernandez

Co-Founder & CTO

12+ years architecting and building software
Amazon, Charles Schwab, 3 Startups



Kandice Wong

Head of Design & Marketing

12+ years design, UX, and product marketing
Akamai, WorkBoard, Kinectify

ADVISORS

Velociti Advisors.

We’re advised by seasoned leaders across Product, Marketing, Sales, and GTM



Dan Tyre

Founder & CEO

Tyre Angel

Microsoft HubSpot



Gary Arnold

Group Chair

Execute Mastermind

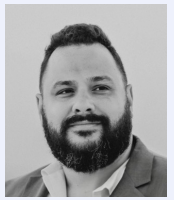
PayPal amazon



Graham Greytak

Founder & CEO

HireRising



Max Schwartz

Chief Executive Officer

Anvil Research Labs

glimmer



Jeff Booth

Director of Product

Helia Care

loadsmart helia

THE ASK

The product is proven.
Now we scale what's working.

$1.5M

PRE-SEED+

SAFE · $10M POST-MONEY CAP

USE OF FUNDS: SCALE WHAT'S ALREADY WORKING

GTM growth engine

First enterprise AE, a 2-call close, and the channels already working: LinkedIn, events, thought leadership.

Deepen the moat

Signal extraction, the Context Layer, and integrations that make us the product intelligence layer.

Pilots to contracts

Convert enterprise pilots to annual contracts and grow ACV from $5K toward $12K.

Core team

Key hires across sales, customer success, and engineering.

Join us to solve the most important problem in AI product development.

anthony@velocitipm.com · www.velocitipm.com

© 2026 | VELOCITI