ANTLER

ANTLER INTELLIGENCE | APRIL 2026

Before the Big Idea, There Was the Founder

What 1,800+ companies tell us about where innovation is heading

AUTHOR: ROSALIND BAZANY, PARTNER

Six months ago I recorded a podcast with a founder called Jeppe. His company, InReality, had a deceptively simple thesis: detecting AI-generated content is a losing game. Generation is improving faster than detection. Within a few years, possibly already, you will not be able to reliably tell. So rather than trying to identify what is fake, InReality digitally signs real content at the point of capture. The signature travels with the content. Viewers can verify it. The line between real and fake does not disappear. It just moves.

I remember thinking this was not a product idea but a diagnosis of how information infrastructure will have to be rebuilt.

Jeppe was at pre-seed. No revenue yet, but a handful of customers already in early conversations. The kind of stage where most institutional investors are not yet paying attention.

That is the thing about pre-seed. The ideas arrive before anyone is ready to validate them.

When a16z published their Big Ideas 2026 series in December, I read through all three parts with a sense of recognition. For Antler, they largely described territory we had already been mapping for eighteen months. Content authenticity infrastructure. Agent-native identity systems. New financial rails for autonomous AI actors. We had seen it, backed it, and watched it iterate.

Let me be clear, that observation is not a criticism of a16z. It is simply a function of vantage point.

The gap between when an investor writes about a category and when Antler first encounters the founder building it is not a competitive advantage. It is a different type of data entirely.

Two different questions

A quick clarification, because the framing matters. a16z invests across all stages, including at seed. The Big Ideas series does not represent a late-stage view of the world.

But here is the question worth sitting with: what kind of signal do you get from a company that has already raised a round, found early customers, and attracted institutional conviction? You get a high-quality signal about what is already working. That is useful, but it is also, by definition, a signal about the recent past.

Antler's portfolio spans pre-seed through to Series C, with regional funds investing at the earliest stages and global vehicles providing follow-on capital as companies scale. Most companies enter the platform before product-market fit, often when the problem itself is still being defined. The signal is noisier and the data is messier, but it arrives earlier.

That distinction is what makes a portfolio of more than 1,800 companies across 27 markets interesting. Not because every one of them will succeed. Most will not. But because at that scale, in that many markets, you start to see things that no other data source shows you: which problems are attracting serious founders independently and simultaneously, which infrastructure gaps are being discovered by multiple teams who have never heard of each other, which categories are forming before they have a name.

That is precisely what early-stage data is for. It does not predict which companies will win, but it does reveal which problems the next generation of founders cannot ignore.

The shift nobody published a newsletter about

Here is the data point that surprised us most when we looked back at it.

In 2022, infrastructure and developer tools accounted for 1% of all AI investments in our portfolio. By 2024, that figure was 21%.

One percent to twenty-one percent in two years.

Founders had collectively decided to stop building on top of AI and start building the layer underneath it. And it happened before the institutional market had named it as a category worth backing.

The way this shows up in practice: across five countries, six separate founding teams independently arrived at the same architectural problem between 2023 and 2024. The problem was that identity systems were built for humans. When AI agents start accessing SaaS platforms as autonomous actors, the entire identity and access management stack breaks. None of these teams were reading each other's pitch decks. They were solving the same problem because the problem was real, not because the narrative was compelling. That kind of independent convergence across geographies, without coordination, is one of the most reliable signals that a real category is forming.

Similarly, by late 2024 we had five separate teams across four countries building payment infrastructure specifically for AI agents: programmable guardrails, verifiable agent identity, full auditability for autonomous transactions. The infrastructure was already taking shape before most investors had decided whether the application layer was real.

When six teams across five countries independently converge on the same unsolved problem without coordination, you are no longer looking at a coincidence. You are watching a category take shape before it has a name.

What we were already building

Content authenticity infrastructure. A16z's Part 1 identifies synthetic media as one of the defining challenges of the next decade. Their framing centres on detection: better tools to identify what is fake. InReality, backed in our Nordic cohort, had reached a different conclusion entirely. Detection is a losing race. The solution is provenance at point of capture, not forensics after the fact. That design decision was embedded in product before it appeared in any tier-1 VC publication. We backed it because Jeppe was right, not because a category existed for it yet.

Agent identity. One of our Australian portfolio companies, first backed in late 2024, is building what they call Agent Customer Identity: typed identities, scoped tokens, revocable sessions, and delegated access flows built specifically for agents operating across company boundaries. They identified a $21.5B identity management market that had not yet recognised the problem existed, because the problem had not been properly named.

Financial rails for autonomous agents. A UK company we backed in 2024 is constructing payment infrastructure for B2B AI agents to send and receive money across payment rails, with programmable guardrails, verifiable agent identity, and full auditability. Their founding pitch described a gap in European infrastructure that no venture-backed company had yet moved to fill.

AI observability. A16z's Part 1 identifies agent-native infrastructure as the defining build challenge of 2026, focused on the throughput and concurrency problem: systems not architected for agent-speed workloads. A Netherlands company we backed in late 2023 had

already identified the adjacent problem: not whether the infrastructure could handle agent workloads at scale, but whether anyone could tell if those workloads were producing reliable outputs. LLM evaluation, observability, and agent testing for enterprises shipping AI products. The category now has a name. In 2023 it did not.

The numbers that change how you read everything else

In 2022, 40% of Antler's portfolio had AI at its core, broadly defined*. By the 2025 cohort, that figure is 79%, based on classified companies. The shift represents a deeper change in what founders believe is worth building.

**Includes company classifications 'Infrastructure and Developer Tools AI' and 'Application Layer AI' where AI is the core product. 'AI Enabled' covers companies where AI is a meaningful component but not the defining proposition. We include this category in broad figures for completeness but treat it separately from AI-native companies.*

The write-off differential tells a sharper story. Non-AI companies in our portfolio carry a historical write-off rate of approximately 29%. Application layer AI: approximately 9%. AI infrastructure and developer tools companies: zero write-offs across 104 companies.

Data in this article uses the November 2025 dataset, predating the a16z Big Ideas publication in December 2025, to reflect what was visible in our portfolio before those predictions were made. As of the November dataset, the zero write-off figure was accurate. Subsequent data through March 2026 records five write-offs across an expanded pool of 134 companies (<4%), consistent with the directional gradient described here. The vintage caveat applies: the majority of these companies are 2024-2025 investments with limited time to reach write-off outcomes.

That is not primarily a story about founder quality. It reflects something structural: investors are funding AI-native infrastructure companies with more conviction and longer runway. But it also tells you something about founder clarity. The people building the layer underneath tend to have an unusually precise understanding of the problem they are solving. AI classification, in our data, turns out to be a surprisingly good proxy for whether a founder knows exactly what they are building and why.

There is also a pricing shift worth noting. Usage-based monetisation in our portfolio moved from 7% of companies in 2024 to 11% in 2025. Traditional SaaS dropped from 34% to 27% over the same period. A16z's Part 2 identifies the shift from SaaS to consumption-based pricing as a 2026 big idea. Our founders were making that pricing decision before they had a paying customer.

What this data cannot tell you

Early-stage company data has a specific failure mode: survivorship bias in reverse. Instead of only seeing the winners that survive to later stages, we see the full population of attempts. The write-off rates in our broader portfolio reflect that honestly. InReality and the companies described above are compelling. But they operate in a market where most pre-seed bets do not return capital.

What early-stage data is good at is pattern recognition at scale: which problems are attracting serious founders, which infrastructure gaps are being independently discovered by multiple teams in different markets, which categories are receiving concentrated attention before institutional capital has priced them in.

The a16z Big Ideas series answers a different question: what has a smart, well-resourced investor decided is going to matter, based on what has already found traction. Both perspectives are valuable. They just operate at different points on the same timeline. One tells you what is already working. The other tells you what is being assembled.

Jeppe did not need a16z to validate the problem he was solving. He had already diagnosed it and just needed someone to back him before the category existed.

When the next Big Ideas series lands, some of the founders in our portfolio will already be the companies it describes.

ABOUT ANTLER INTELLIGENCE

Antler is one of the most active early-stage venture firms globally, with US$1.3B+ AUM, operations across 27+ cities, and a portfolio of 1,800+ companies spanning every major technology market. We invest from pre-seed and seed, where the majority of our portfolio sits, through to Series C via regional early-stage funds and global follow-on vehicles. That structure gives us something few investors have: a longitudinal, cross-geography dataset of founder behaviour, category formation, and technology adoption that begins before most institutional investors are in the room.

Antler Intelligence is our research and market insight function designed for an LP audience. The views expressed here are the author's own, informed by Antler's proprietary portfolio data. They represent an analytical perspective on market trends, not the firm's investment recommendations or the position of any Antler fund. We publish this analysis because the early stage vantage point is underrepresented in market discourse, and because our LPs and ecosystem partners benefit from understanding what we are seeing before the market names it.

DATA AND DISCLOSURE

This article may contain forward-looking statements based on current assumptions, which are subject to change. No obligation is undertaken to update this material. It is not intended for distribution in any jurisdiction where such distribution would be contrary to applicable law.

The data in this piece is from November 2025, before this article was written and before a16z published their Big Ideas 2026 series in December. We have not updated it for publication. The analysis stands on what we saw then, not on what would have been convenient to see. All portfolio data referenced in this piece reflects Antler's internal dataset as of November 2025 across 1,429 active companies. AI classification analysis excludes 66 companies (4.6% of the active portfolio) not yet categorised at time of reporting; the 2025 cohort figure of 79% reflects 289 of 346 companies classified at that date. Write-off rates are drawn from Antler's full historical dataset including inactive companies and reflect cohort-level outcomes within a specific time window; they are not indicative of future results. Usage-based and SaaS pricing figures reflect business model classifications at point of investment and are subject to change as companies iterate.

Early-stage companies operate in conditions of high uncertainty: they pivot, raise capital on compressed timelines, and occasionally cease operations without public notice. Figures presented here are directional indicators of portfolio-level trends and should not be treated as audited performance data, point-in-time valuations, or predictive of future returns. Portfolio companies named in this piece (InReality) have consented to being referenced. Companies described in anonymised form are active Antler investments; identifying details are withheld. This document does not constitute an offer or invitation to invest in any Antler fund or vehicle. Recipients should not rely on this analysis for investment decision-making without independent verification.

Podcast: Verifying Reality in the Age of AI by Ros Bazany in conversation with Jeppe Nørregaard, co-founder of InReality. Available on Spotify.

For further discussion on portfolio composition or Antler's research, contact investor.relations@antler.co or visit antler.co.