

*** Public Markets ***

Crypto ended June with a \$3.4T market cap (+24% QoQ), outperforming both equities (S&P: +11%) and physical commodities (Gold: +6%), taking the title of best performing global asset class of the quarter. Bitcoin alone was responsible for more-than-all of the outperformance, returning +30% as an explosion of [Bitcoin treasury vehicles](#) drove new speculative capital flows amounting to over \$10B. Excluding Bitcoin and Ethereum, the rest of crypto (+9% QoQ) and DePIN specifically (+6% QoQ) lagged even equities.

In our prior letter, we discussed where the next set of inflows into crypto would come from: institutional allocators, retail speculators, or protocol buybacks. While we were most bullish on the latter, all three channels saw significant inflows in Q2'25 and gave rise to at least one of the quarter's top-performing assets: Microstrategy (+40%), Robinhood (+125%), and Hyperliquid (+200%) respectively.

One view is inflows will come from **TradFi allocators**, in which case the [15+ cryptoassets](#) with active **ETF applications** are likely to outperform. A second view is that **retail speculators** will re-enter crypto, driven by falling rates and mainstream fintechs like Robinhood [deepening](#) their crypto offerings. In this version of the world, a [2021-like market structure](#) emerges whereby **front-running listings** is the dominant strategy: the meta-game becomes predicting which tokens the major fintech platforms are most likely to list first. Anecdotally, we'd estimate 80% of investors cite one of the above two reasons for being bullish on alts.

The third, and least popular, view is that protocols will generate **sizeable onchain revenues** that drive organic demand for tokens via **buybacks or dividends**. [Hyperliquid](#), whose native token \$HYPE is up [+100%](#) since launching in late November, generates ~\$1m/day in onchain trading fees and directs half that revenue (~\$500k/day) towards token buybacks. \$HYPE is valued at a 25x multiple to those buybacks

EV3's Q1'25 investor letter

Institutional allocators drove \$25B of new money into Bitcoin, and to a much lesser extent, Ethereum, last quarter. The majority (60%) went to unlevered low-fee [ETFs](#). The remainder (40%) was corporate treasury vehicles, led by [Microstrategy](#), adding BTC to their balance sheets. Microstrategy's success - it trades at a [>110% premium to NAV](#) - has inspired another 20+ publicly-listed entities around the world to replicate its playbook, not only for BTC but also increasingly for cryptoassets like [Solana](#) and [Bittensor](#). If you're not already familiar with these vehicles, you might be surprised to hear investors today are willing to pay a full dollar for a share of stock that represents ownership of less than fifty cents worth of BTC.

The reason for this is two-fold: 1) investors expect Microstrategy to earn BTC-denominated yield by [selling convertible debt to hedge funds](#) seeking to harvest Bitcoin's volatility, and 2) a class of investors who want crypto exposure, but are not willing or able to transact on crypto exchanges, pay a premium to access equity-exchange listed products. We [recently interviewed](#) a fund betting on these vehicles and pressed them on their investment case—ultimately, they admitted underwriting 20-30% premiums to NAV. In other words, everyone knows today's valuations are nonsense, but no one knows how long the music will last, and there are reasons to believe we are closer to the beginning than the end of the mania.¹ As volatility subsides and more of these vehicles launch, premiums will collapse and the pace of new capital raised by these companies will fall precipitously. Investors who bought shares on the secondary markets will underperform holding crypto, but pre-launch investors - i.e. those with existing crypto exposure who contribute tokens in-kind to start up these vehicle - could still make money, assuming the NAV premium settles at a level that more-than-offsets fees. Fund I invested in one Bittensor-focused treasury company, [xTAO](#), in Q1'25, which received preliminary approval for a listing on the TSXV in Toronto [last week](#). Despite our skepticism on the longevity of the "treasury trade" overall, the structure, terms, team and timing behind xTAO make it an attractive opportunity to re-deploy a portion of Fund I's TAO position,

¹ Back in 2022, seven full months passed between SBF's infamous "[magic box](#)" interview in April and FTX's collapse in November.

Perhaps the quarter's biggest shift - and one that we didn't foresee - was the **insatiable demand for crypto-related equities** exemplified by Circle **blowing the IPO window wide open** and returning +500% in its first month of trading. Circle's business is simple: they hold ~\$60B of dollar reserves on their balance sheet and earn roughly the risk-free rate in gross yield (~4% in Q1'25). Circle has wholesale distribution agreements with crypto exchanges which are entitled to the majority (~60% in Q1'25) of their gross yield. For every USDC outstanding, Circle earns 4 cents per year of interest, and keeps 1.6 cents after revenue sharing to distribution partners, but before salaries, admin and technology costs. Today, the market values Circle at almost 75 cents: a multiple north of 50x revenues, or 65x profits. When the business of holding *someone else's* dollar is itself worth 75 cents, even the **CNBC talking heads** start sounding the alarms.


	Today	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10
USDC Outstanding	\$62,000	\$80,600	\$104,780	\$136,214	\$177,078	\$230,202	\$299,262	\$389,041	\$505,753	\$657,479	\$854,723
(x) Gross Yield	4%	4%	4%	4%	4%	4%	4%	4%	4%	4%	4%
Gross Revenue	\$2,480	\$3,224	\$4,191	\$5,449	\$7,083	\$9,208	\$11,970	\$15,562	\$20,230	\$26,299	\$34,189
(x) Revenue Share	60%	55%	50%	45%	40%	35%	30%	25%	20%	20%	20%
Net Revenue	\$992	\$1,451	\$2,096	\$2,997	\$4,250	\$5,985	\$8,379	\$11,671	\$16,184	\$21,039	\$27,351
(x) Margin	70%	70%	70%	70%	70%	70%	70%	70%	70%	70%	70%
EBITDA	\$694	\$1,016	\$1,467	\$2,098	\$2,975	\$4,190	\$5,866	\$8,170	\$11,329	\$14,728	\$19,146
(x) Multiple	65x										15x
Enterprise Value	\$45,000										\$287,187
IRR	25%										

What would it take for Circle to generate a 25% IRR at today's prices? Based on a **napkin DCF**:

- USDC supply grows at a 30% CAGR over a decade to \$850B.
- This would make Circle the **second-largest holder** of US treasuries behind Japan.
- US treasury yields remain at 4% long-term... relative to a **trailing 20-year average** of below 1%.
- Circle negotiates revenue-share paid to distribution partners down to 20%... from 60% today.
- Circle scales into operating leverage... achieving best-in-class 70% operating margins at scale.
- Circle is valued at 15x terminal EBITDA multiple... in line with high-growth asset managers.


The takeaway: equity markets are hungry for, and willing to overlook a lot to get, exposure to stablecoins.

The second source of inflows is **new retail speculators**, buying crypto through fintechs like Robinhood and Nubank **aggressively deepening** their crypto trading products. These two companies alone account for 150M users with \$250B of assets, roughly half the size of Binance plus Coinbase, but list only a very limited number of tokens. We believe their untapped retail distribution will crown many of the winners and losers in alts over the medium-term. Today, a Robinhood or Nubank listing is worth much more than a Coinbase or Binance listing. At the same time, liquidity and **price discovery** for top assets is shifting to decentralized exchanges like Hyperliquid. CEXs realize their oligopoly is under threat and are finally beginning to accept lower-margin business. We see evidence of this in our own portfolio with companies negotiating listing fees 20-40% lower than has been possible in the past. As the exchange landscape gets more competitive, TGE playbooks will mature and more closely resemble IPOs in costs and structure.²



² See page 8 of our **Q1'25 investor letter** for a more detailed discussion on TGEs vs IPOs.


The billion-dollar question in the crypto alts market today is identifying which types of assets these fintech platforms will prioritize listing. Nubank primarily lists the leading L1/L2 tokens ranked top-25 by market cap, plus a small number of DeFi “blue-chips” (e.g., Aave, Chainlink, Ripple). Robinhood has a similar approach, but with a penchant for low-market cap memecoins (e.g., Moodeng, Popcat, Mew). Hamstrung by uncertainty surrounding the treatment of utility tokens under US securities laws, they have focused on listing crypto-assets with the lowest risk of being classified as securities, i.e. memecoins. This is, in our view, a **temporary aberration driven by over-regulation**: long-term, all else equal, users will prefer to speculate on, and exchanges will prefer to list, assets with fundamental value over those with none.



Number of crypto-assets listed by Nubank and Robinhood


This dynamic is set to change over the next eighteen months. The Trump administration’s top crypto priority is stablecoin legislation, with the [GENIUS](#) act [passing](#) the Senate in [June](#). Their second priority is market structure, i.e. the [CLARITY](#) act, which defines the parameters by which tokens are categorized as SEC-regulated securities vs CFTC-regulated commodities. Both bills are [up for vote in the House](#) this week. Assuming they pass, we expect platforms like Robinhood to begin spot listings for crypto-assets that fall under CFTC’s jurisdiction in 2026, followed by derivative products in 2027.

There are currently eight DePINs ranked in the top-200 tokens by market cap and trading volumes - Bittensor, Render, Filecoin, Helium, Walrus, Theta, Bittorrent and Arweave - which will likely be first in line for listings and disproportionately benefit from the first wave of new capital flows. [the two category-leading tokens on the list with strong momentum and asymmetric upside: \\$TAO in decentralized AI and \\$HNT in decentralized wireless.](#) As these networks re-rate, crypto market participants will extrapolate the new “ceiling” valuations across the entire sector and go hunting for exposure among the 40+ live DePIN tokens with sub-\$100M market caps.



Source: Coingecko.

The last channel for capital flows into crypto, and the one we are most confident in sustaining long-term, are **onchain buybacks funded by protocol revenues**. We are already seeing evidence of thesis playing out: two of the five top-performing crypto-assets in Q2'25 - [HYPE](#) and [VIRTUAL](#) - generated \$175m and \$9m, respectively, in quarterly fees (the other three were memecoins). With top-line revenue multiples of 20-30x, these two protocols trade in line or at a discount to leading AI startups in the private markets. Within crypto, revenue-generating tokens are outperforming by ~15% YTD per the team at [Pantera](#).



Per [DeFiLlama](#), crypto protocols generated \$18B of onchain fees in the past thirty days (\$215B run-rate) through trading fees (53%), gas and MEV-related fees (29%), and lending or staking fees (18%). DePIN, with only ~\$75M of annualized onchain fees today, isn't even in the race yet.

But not every dollar of revenue is equal: **revenue quality** matters. Most VCs, trained on a decade of SaaS, fintech, social and marketplaces, have a certain framework for thinking about revenue quality: net retention, capital intensity, lifetime value, and paybacks. Crypto is a special beast: the inherent reflexivity of token-based models reduces the predictive power of traditional metrics down to near zero. The seminal example is Axie Infinity, which - after reporting world-class 90-day retention rates of [~40% in Q3'21](#) - churned [more than eighty percent](#) of its users in 2022 as its token price declined by almost 95%. At EV3, our mental model for revenue quality is inspired by George Soros' [theory of reflexivity](#):

"Stock market valuations have a direct way of influencing underlying values: through the issue and repurchase of shares and options and through corporate transactions of all kinds-mergers, acquisitions, going public, going private, and so on. There are also more subtle ways in which stock prices may influence the standing of a company: credit rating, consumer acceptance, management credibility, etc."

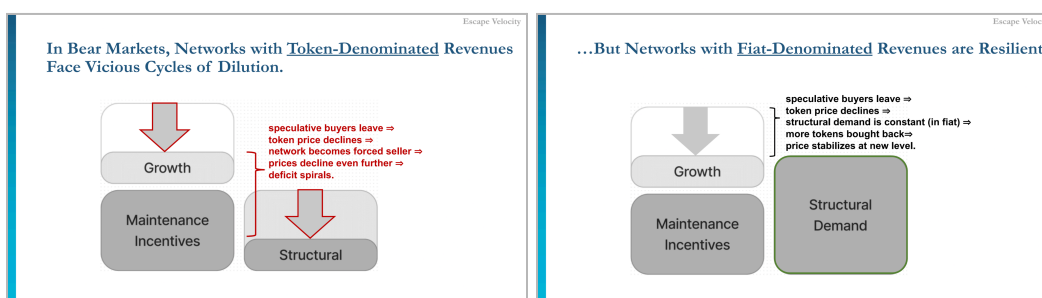
There are two types of onchain revenues. Financial use cases that represent >99% of onchain fees today are *token-denominated*: algebraically, a product of token prices.³ DePIN flips the script, generating *fiat-denominated* fees that are independent of token prices. When markets turn for the worse, only the latter bucket are resilient to a downward spiral of reflexivity.

Token Denominated	Trading fees	=	\$ token price	X	# of tokens traded	X	% swap fee
	Gas fees		\$ token price		# of transactions		avg gas fee
	Lending fees		\$ token price		# of tokens borrowed		% net interest
Fiat Denominated	Wireless fees	=	\$ / GB	X	# of GBs used	X	
	Energy fees		\$ / kWh		# of kWhs used		
	Compute fees		\$ / compute-hr		# of compute-hrs		

³ This actually understates the reflexive nature of token-denominated fees. The second factor - volume of tokens traded, borrowed or staked, or volume of transactions processed - is usually both highly cyclical and highly correlated to token prices.

DePINs exhibit a level of reflexivity akin to Soros' model for stocks. Inflationary token incentives are worth more in bull markets, accelerating growth, and less in bear markets, decelerating it. Token price volatility therefore impacts growth, but has **no direct impact on the existing revenue base**—assuming demand isn't being subsidized. This dynamic is similar to a corporation that relies on primary equity issuances to fund its growth capex, but not maintenance capex. We call this type of behavior *first-order reflexivity*.

Protocols with token-denominated revenues exhibit *second-order reflexivity*, akin to Soros' model for free-floating currencies.⁴ Token price volatility impacts both growth and **also the existing revenue base**. In bull markets - in Soros' parlance, when the prevailing trend is positive - revenues grow super-linearly; but in bear markets revenues decline rapidly in a reflexive downward spiral with price. We've been writing about this framework [since 2023](#), but crypto markets continue to systematically undervalue the resiliency of fiat-denominated revenues such as DePIN. To be fair, DePINs were not generating meaningful onchain revenues during the last crypto bear market in 2022-2023, so there's not yet hard evidence to back us up. Only in the next bear will we know empirically if DePIN weathers the storm better than the rest of crypto.




EV3's [Scaling Software with Web3](#)

To illustrate the point, we ran a napkin math [thought experiment](#) comparing HYPE and HNT. We assume:

- Four bull and three bear markets over the next fifteen years, each lasting two years.
- Hyperliquid trading volume grows +100% YoY in bull markets, but falls -75% in the first year of a bear market and recovers +50% in the second. By year 15, Hyperliquid does over \$150B in ADTV, bigger than NYSE or Binance today.
- Helium data traffic growth decelerates from 100% to 40% YoY. Traffic is flat in the first year of a bear market, but never declines. By year 15, Helium serves 4-5% of US adults with mobile phone plans, comparable to TracFone or Boost.⁵
- Both tokens are valued at a 15x terminal multiple, i.e. at a 25% discount to Hyperliquid's current valuation. Both tokens increase their circulating supply by 50% over the next fifteen years, i.e. 33% dilution for today's tokenholders.

Hyperliquid			Helium		
	Year 0	Year 15		Year 0	Year 15
ADTV	\$8,000	\$162,000	GBs per Day	33,000	12,113,330
(x) Fees	0.026%	0.026%	(x) \$ per GB	\$0.50	\$0.15
Annual Revenue	\$750	\$15,196	Annual Revenue	\$6	\$663
(x) Multiple	21x	15x	(x) Multiple	89x	15x
Market Cap	\$15,468	\$234,924	Market Cap	\$538	\$10,253
Dilution-Affected	\$15,468	\$156,616	Dilution-Affected	\$538	\$6,835
IRR	24%		IRR	23%	




For roughly the same payout, would you rather hold HYPE through three -75% downturns to bet on it becoming the “next Binance”? Or hold HNT for consistent growth through cycles and bet on it becoming the “next TracFone”? In our view, both of them are attractive bets, and we own both networks in different funds. However, with only one of them can we sleep through the night without checking Coingecko.

⁴ While it's out of scope here, interested readers can refer to [Alchemy of Finance](#) Ch 3: Reflexivity in the Currency Market.

⁵ We also assume Helium's per-GB pricing gradually falls from \$0.50 to \$0.20 by year six and thereafter. See our [DeWi research](#).

With the words “DePIN”, “revenue” and “fees” appearing nearly fifty times in the past eight pages, we couldn’t end this section without addressing the elephant in the room—**where are the DePIN revenues?**

At the end of Q1’25, DePINs were annualizing \$37M in onchain revenues and \$10M in onchain buybacks. Onchain revenues doubled to a \$75M run-rate in Q2’25, primarily driven by:

- **IO.Net (\$IO)** grew GPU marketplace revenues from \$20M (offchain) to \$27M (onchain) ARR, up +35% QoQ. More importantly, their revenue is now [verifiable onchain](#) albeit after more than a year since the TGE.
- **Helium (\$HNT)** grew wireless revenues from \$2.4M to \$4.4M ARR, up +80% QoQ. This excludes \$3M in non-recurring fees and an additional \$6M in offchain revenue currently (but [not permanently](#)...) earned by Nova Labs, implying a total annualized run-rate for Helium of \$13M.
- **Dimo (\$DIMO)** began reporting \$1.0M ARR of [verifiable onchain revenues](#) from developer subscriptions for its automotive data platform. They also launched partnerships to expand into [Japan](#) and [Chile](#).
- **Natix (\$NATIX)** grew mapping and subscription revenues from \$0.7M to \$1.0M ARR, up +40% QoQ. The company recently started shipping its new mapping product, the [VX360](#), specifically for Tesla owners.
- **Silencio (\$SLC)** began reporting \$0.3M ARR of [verifiable onchain revenues](#) from noise data sales, in-app subscriptions, and burning of unclaimed airdrops.
- **Impossible Cloud (\$ICNT)** [launched](#) its token and began self-reporting \$6M offchain ARR, up +15% QoQ, from enterprise cloud contracts in Europe. The team has indicated revenues will be moved onchain in Q3.
- While decentralized cold storage seems to be a failed experiment, decentralized **hot storage** feels like the complete opposite. Decentralized CDNs including [Titan](#), [Aro](#), [Pipe](#), [Gradient](#), [Mawari](#), and [Rilla](#) each raised between \$1-10M from VCs and are privately reporting generating a similar amount of annualized revenues. If their momentum continues, CDNs could soon begin to rival GPUs as the highest-grossing DePIN product.
- 

While DePIN’s top-line growth was strong, not every development was positive:

- **Filecoin (\$FIL)** grew storage revenues from \$2M to \$3.5M, up +70% QoQ. However, more-than-all of this growth was driven by slashing penalties (+140% QoQ), whereas paid storage demand actually shrank (-15% QoQ). Arweave generates <\$20K in annualized fees. There is no demand for decentralized cold storage.
- Across nearly **every subsector of DePIN**, there was a network like Akash (\$AKT) in compute, Glow (\$GLW) in energy, Hivemapper (\$HONEY) in mapping, Braintrust (\$BTRST) in labor, or WiFi Map (\$WIFI) in wireless whose **onchain revenues declined** quarter over quarter. The reality is that it’s extremely rare for DePINs to achieve product-market fit on the first try.⁶ In the interim period, revenue momentum tends to dissipate.
- **Grass (\$GRASS)** has not yet disclosed further details about their “mid eight figures” of offchain ARR.
- The early cohort of DePINs, like Helium and Filecoin, made it possible to **analyze revenue data at the line item level**; for example, to delineate between recurring vs one-time or demand- vs supply-side fees. The common practice for DePINs now is to **lump all revenues into a single address**, making it impossible to disaggregate revenue streams. Silencio and Natix, for example, have a single burn address that combines proceeds from data sales to enterprises, in-app consumer subscriptions, e-commerce affiliate commissions, early withdrawal and other slashing penalties, and even burning of unclaimed airdrops.

Based on what we know today, we believe DePIN will finish the year at \$100-150M onchain ARR.

⁶ Helium tried IoT and pivoted to mobile. Akash tried cold storage and pivoted to GPUs. Glow tried carbon credits and pivoted to solar subsidies. Natix tried smartphones and pivoted to autonomous vehicles. Braintrust tried a freelancer marketplace and pivoted to AI recruiting. Dimo tried OBD readers and pivoted to API connections. Hivemapper tried WiFi-enabled single-camera dashcams and pivoted to LTE-enabled devices with stereo cameras and advanced GNSS. WiFi Map and Silencio tried aggregating WiFi and noise data, respectively, and pivoted to general-purpose data collection. Grass tried reselling browser extension IPs and pivoted to scraping data and LCR via desktop hardware and software nodes. The notable exceptions are GPUs (\$IO, \$RENDER, \$LPT) and positioning (\$GEOD) networks: each generates \$1M+ ARR with substantially the same business model as when they launched.

*** DePIN's Path to \$1B+ ARR ***

What will it take for DePIN to generate \$1B+ of onchain revenues?

1. **Cover the full spectrum of digital infrastructure.** There are obvious shots on goal for venture-scale outcomes like fixed wireless, hot storage, and residential solar, that are not represented in onchain revenue today. These projects exist, and we own sizable stakes relative to fund size in many of the category-pioneering ones already, but none of them are live on mainnet just yet.
2. **Unify, or at least transparently align, the interests of tokenholders and shareholders.** As crypto institutionalizes, investors are increasingly wary of buying tokens that resemble [junk equity](#) at the very top of the cap stack with essentially no rights or protections. We are watching the potential merger of Nova Labs and Helium Foundation as a precedent for other projects. The CLARITY act will hopefully provide a legal framework towards which deals can be structured.
3. **Debt financing to unlock the next 10x in scale.** DePIN has been almost entirely equity-funded to date, which is a nonsensical way to finance a global-scale infrastructure buildout. Lower-cost financing, either in the form of loans from external lenders to miners, or from miners lending to the protocol itself, will be a powerful accelerant for DePINs with proven, scalable product-market fit.
4. **Build stronger ties with AI companies.** The majority of DePIN revenues come directly from the pockets of venture-backed AI companies, and yet we find few early-stage DePIN founders today designing products and protocols with the specific needs of these customers in mind.
5. **Expand the scope to User-Generated Services.** True DePIN, i.e. hardware networks, is a subset of the broader category of crypto-enabled business models that leverage crypto-incentives to aggregate and coordinate resources to power a useful service. We believe this umbrella category, User-Generated Services, can generate \$1B ARR by YE'30 (60% CAGR).

1. First, DePIN needs to **cover the full spectrum of digital infrastructure**. Of eight categories, only five have achieved some level of product-market fit and are now in the phase of scaling onchain revenues. While skeptics will point to the fact that every category has product flops, it's also true that many end-markets within each category are still greenfield. Fixed internet, hot storage, residential solar, pluralistic identity, wearables, and food delivery are all \$100B+ markets with zero DePIN products live on mainnet today. Emerging technologies like visual positioning systems (VPS) and LIDAR are small markets today (<\$10B ARR) but could be direct beneficiaries of tailwinds indexed to robotics growth. With many of the biggest opportunities yet to be explored in earnest, notwithstanding those we haven't even considered yet, DePIN's onchain revenues today (<1% of crypto) are not representative of its long-term potential.

Category	Post-PMF	No PMF	Pre-PMF
Wireless	Mobile (HNT)	IoT (HNT)	Fixed (DAWN)
Compute	GPUs (IO, AKT)	Cold storage (FIL, AR)	Hot storage (ARO, PIPE)
Positioning	RTK (GEOD)	BLE (NODLE)	VPS (AUKI, OVR)
Mapping	Stereo (NATIX, HONEY)	Mono (NATIX, HONEY)	LIDAR (ROVR)
Energy	--	Carbon Credits (GLW, AKRE)	VPPs & DR (Daylight)
Identity	Monolithic (WLD)	Federated (BRIGHT)	Pluralistic (zkTLS)
Health	--	Step counters (SWEAT, GMT)	Wearables (Cudis, Pulse)
Logistics	--	Rideshare (Teleport)	Food delivery (LOCAL), Parcel (Heale)

DePIN products by category and PMF stage.


2. DePINs need to **unify - or at the very least, transparently align - the interests of tokenholders and shareholders**. This problem is [not new](#) nor is it [unique to DePIN](#), but it's particularly problematic because the DePIN revenues are typically mediated through a centralized "Labs" company that repackages the core commodity produced by the network into enterprise- and/or consumer-grade products. The crux of the issue is that both the revenue collected from customers and the offchain costs incurred by Labs are invisible to the protocol, i.e. there is no way to verify them onchain, yet.⁷ Some DePINs, like [Aethir](#), exploit the ambiguity to inflate onchain "revenues" (reportedly \$14M per *month*) by using pre-mined tokens to "purchase" services on the network, often from miners that are owned by or affiliated with the core team.

More broadly, by the time a DePIN launches its token, VCs own 25-40% of the shares of Labs and the rights to 15-25% of the token supply. VCs on the same cap table often own different proportions of shares vs tokens depending on which round they invested in, and founders' ownership is almost always skewed towards shares. Some investors are incentivized to use the cash on Labs' balance sheet to **subsidize growth of protocol revenues**, while others are incentivized to prioritize initiatives that **drive profits from the value-added services** that Labs earns as a margin on top of the protocol, depending on their relative ownership and their relative bullishness on crypto vs equity capital markets.⁸ At EV3, our philosophy is to align our ownership with founders - the people actually creating the value - as much as we reasonably can in the private markets, but in liquid markets we are at the mercy of the options the market gives us.⁹

The Helium team, as usual, is at the forefront of pushing the space forward with [recent discussions](#) about the possibility of [buying out](#) shareholders with newly-issued HNT, effectively merging Nova Labs and the Helium Foundation. With HST (Helium Security Tokens) [phased-out](#) last year, this move would fully-align the interests and governance of the entire ecosystem behind HNT. Our latest [DeWi research](#) covers what a merger could mean for HNT tokenholders in more depth, but at a high-level we have a constructive view of the potential deal (pending more details) and hope to see it come to fruition over the next few quarters.

In cases where unification of interests is not possible or practical, the industry is at the very least finally demanding a minimum level of standard disclosures from protocols, led by Blockworks' [token transparency framework](#). We hope to see more DePINs adopt this or similar frameworks, in addition to providing more granular, verifiable data about company revenues, expenses and agreements. On [DePIN Pulse](#), the successor of DePIN.Ninja, we've added warnings flagging or suspicious revenue data. Without the crypto industry taking initiative to self-police itself with respect to transparent disclosures, an entire class of investors accustomed to reading S-1s and 10Ks will remain locked out of the alts market.


3. As protocols achieve product-market fit and are able to reward miners independent of token subsidies, **DePIN needs debt financing to hyperscale the next 10x in growth**. Simply put, financing a global infrastructure buildout with equity (or an equity-like token) is not a viable strategy—the math doesn't work. US telcos and utilities combined own \$1.5T of assets and generate \$100B in annual operating profits, i.e. a core asset yield of 7%. Even with a relatively modest 12% cost of capital, an entirely equity-funded infrastructure business would "lose" \$75B per year after implied financing costs. It just doesn't work.



DePINs can support at least three forms of debt: two exogenous and one endogenous. Exogenous debt is issued from external lenders directly to miners, collateralized by some combination of future earnings streams and residual hardware value. Debt capital can be sourced both onchain ([Metastreet](#) is leading here, working with several Fund I portcos) and offchain ([Tether](#) is beginning to lend to DePIN miners). [Top-performing](#) Helium hotspots are generating 1000%+ IRRs in dollars based on data transfer alone, i.e. before inflationary token incentives; they are more than capable of supporting debt on the basis of future earnings. For networks with commodity hardware (most of wireless and compute), devices have residual value and can be resold to ISPs or datacenters; that said, collections remains an unsolved problem, so most lenders take recourse to a corporate entity as well. Historically, the limiting factor on debt has been the scale of DePIN's revenues: with the biggest protocols generating \$3-5M ARR, and the biggest miners representing at most 10-20% of each network, individual miners can only support \$0.5-1M of annual debt service which implies loan ticket sizes between \$2-3M—too small to get institutional lenders out of bed.

The other form of debt for DePINs is **endogenous, or in-protocol debt**: miners provide implicit financing to the protocol by accepting mining rewards in a new token that is redeemable for a fixed, \$-denominated amount of the original token in the future—effectively a convertible debt instrument. It's a win-win: miners gain certainty in their ability to cover \$-denominated costs with mining rewards (assuming the protocol doesn't default), and the protocol lowers its cost of capital by borrowing from counterparties that are already aligned with its long-term success. Of course, the devil is in the details: protocols need thoughtful market-based interest rate mechanisms and redemption gating to avoid a “run on the bank” where miners redeem claims en masse and sell the proceeds, driving a reflexive vicious cycle of leverage.¹⁰ More pragmatically, following the collapse of Terra's UST in 2022, exchanges are extremely reluctant to engage with projects with any sort of similar dual-token two-tier capital structure. But once one successful project pioneers this approach, we believe it will quickly become part of the standard DePIN playbook.

4. DePIN needs to build stronger ties with AI companies, which have proven to be the most eager customers for most of the core commodities - namely compute and data - that DePINs produce at a scale far greater and a cost far lower than centralized incumbents. By our estimates, a majority of DePIN's current revenues come from AI companies. StabilityAI (\$180M raised) and LeonardoAI (\$30M raised) rent compute [from Render and IO.Net](#), respectively. Volkswagen (\$50B market cap) and Grab (\$20B market cap) buy mapping data [from Hivemapper and Natix](#), respectively. Grass, for scraping online data, alongside a cohort of projects that raised VC rounds in H1'25 to build multi-modal data collection for robotics use cases - [Bitrobot](#), [Reborn](#), [PrismaX](#) and [Eidon](#) - are reportedly selling data to the likes of OpenAI and Anthropic. Even Geodnet, the high-precision positioning network, has [rebranded to focus on robotics](#). We don't think it's an exaggeration to say **there are only three proven, scalable business models in DePIN today**: 1) renting compute to AI companies, 2) selling data to AI companies, and 3) selling WiFi capacity to telcos. Yet the vast majority of early-stage DePIN teams lack the talent and insights to design their networks in a way that caters to the specific needs of cash-rich AI companies.





¹⁰ For interested readers, we [wrote but never published](#) a letter on DePIN financing (see page 10).

5. Finally - and perhaps most controversially - we need to **retire the term DePIN**. In the literal sense, Decentralized *Physical* Infrastructure networks are only a subset of the projects we've discussed above. Braintrust doesn't need hardware: its software aggregates and coordinates disparate pools of employers and freelance talent around the world. Neither do CDNs or VPNs: their software simply routes traffic over existing internet rails in a way that optimizes for latency or privacy. Both types earn revenues denominated in fiat ("first-order reflexivity") and we consider both to be within the universe where EV3 has a right to win vs top investors both within the crypto universe and outside it. Therefore, going forward, you'll hear us use the umbrella term **User-Generated Services** to describe the broader set of **crypto-enabled businesses** that leverage tokenization in some manner to **aggregate and coordinate large, disparate sets of user-owned resources** to provide a useful, **fiat-denominated** service or commodity.

DePINs, then, are the subset of user-generated services that require aggregating and coordinating user-owned (or user-operated) **hardware devices** in order to power their core service. **Crypto-enabled advertising networks**, like our Fund I portcos **Daisy** and **EarnOS**, are the subset that aggregate user attention and coordinate **user behavior** in a way that unlocks new forms of ad inventory for brands. **Value-added bandwidth** networks aggregate and coordinate **existing internet-connected computing resources** to more efficiently or performantly route network traffic and compute loads through the grid. Pluralistic identity (zkTLS) networks, like our Fund I portco **Opacity**, aggregate and coordinate **users' account data** in order to provide developers with a tamper-proof, yet privacy-preserving alternative to traditional tracking and KYC. Freeing "DePIN" from the shackles of being the frontman term for every fiat-denominated, crypto-enabled business on earth makes the term more intuitive and more useful.

Since our love language at EV3 is sharing falsifiable predictions, we wanted to share some with you today as to what the User-Generated Services landscape looks like in 2030 with \$1B in annualized revenues:

- We think compute will grow and remain the biggest sector, but not by much—wireless and energy will grow even faster, each generating ~\$200M of annualized revenues in 2030.
- Identity and ads is our dark horse sector—we think it'll surprise everyone with hockey-stick revenue growth to >\$100M ARR while building defensible moats, e.g. an identity graph.
- Mapping, labor and logistics we think will end up in similar places by 2030, whereas wearables will take longer to scale, and positioning will grow in volumes but offset by margin compression.



*** Private Markets: Adtech ***

[REDACTED] we'll take this time to briefly discuss our emerging adtech thesis, on which we published a [long-form research piece](#) in June. If you had asked us in 2022, or really at any point in our investing careers, if we planned to invest in adtech, you would have heard two emphatic no's. Thankfully, over the years we've learned to tune out the left side of our brain and trust our instincts when choosing the founders we partner with—and then we learn from them.

In adtech, we specifically have [Ray Lee](#), co-founder of [Daisy](#), and [Phil George](#), co-founder of [EarnOS](#), to thank. From them, we discovered three insights that we would not have gotten with any number of hours spent at crypto conferences or using AI research tools and eventually led to us double-clicking on adtech.

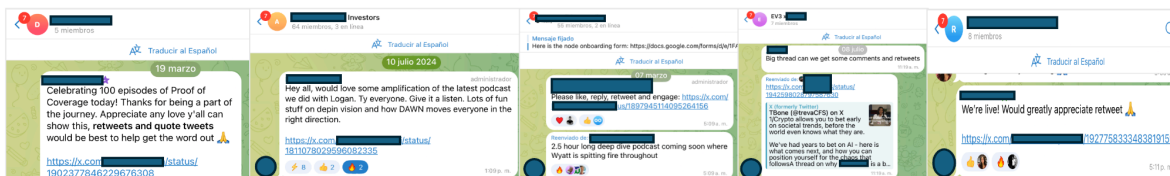
1. Performance marketing agencies are fantastic customers for VC-backed startups. They control massive brand marketing budgets and strategically allocate a portion of this spend for experimenting with new channels. This high-risk bucket provides agency execs the mandate to unilaterally pilot new adtech products without needing “buy-in” from colleagues. Agencies are not price sensitive because their core business is spending OPM (other people's money). They don't bat an eyelash at 20-30% take rates so long as campaign results net of fees are better than what they can get elsewhere. If a product performs better than their existing channels on well-understood metrics, they don't hesitate to roll it out across the rest of their budgets—not just for the one client, but across their entire book of business. Said otherwise, marketing agencies hit the trifecta for the ideal startup customer: they're relatively easy to acquire (low CAC), relatively price-insensitive (high margin), and there's little friction for them to shift massive amounts of spend to new channels (i.e., high net retention). Plus, there's no shortage of them: the US alone counts 50K+ digital marketing agencies and another 50K+ “analog” ones.

It didn't take long for us to notice this phenomenon in the portfolio. [REDACTED] By the end of the summer, agencies and brands were spending \$[REDACTED] mo boosting their sponsored content through [Daisy](#). By the end of 2024 that figure topped \$[REDACTED] mo, and in less than twelve months since putting the initial version of the product into the hands of customers Daisy surpassed \$[REDACTED] in monthly gross revenues. EarnOS is a similar story: we led their [preseed round](#) in December 2024, the team shipped a public beta in [May](#), and in the past two months, through partnerships with top agencies, EarnOS facilitated over 25M verified user missions for [brands including](#) Doordash, Uber, North Face, Amazon, Temu, L'Oreal, Adidas, Disney, Coinbase, Binance, Tinder, Xbox, Vodafone, Marriott, Barbie, Call of Duty, and the Chicago Bulls. After waiting three years for DeWi's fabled [AT&T contract](#), the urgency with which marketing agencies are willing to adopt new adtech tools and test new channels is admittedly a breath of fresh air.

2. Even the simplest-sounding behaviors - when coordinated at scale - can quickly snowball into high-growth, high-margin adtech businesses in today's attention economy. Daisy is the most salient example in the portfolio. We were actually first [pitched the idea for Daisy](#) by Hersh Patel, co-founder of Opacity, a few months after writing the first check into his company in Q4'23. Hersh, with a few customer discovery calls, keenly picked up on an interesting emergent behavior in influencer marketing circles:

“Currently what happens is one influencer makes a tiktok for some product, and other influencers in the group will comment/like/share within 10 minutes to boost it in the algorithm. They use screenshots and the honor system to make sure everyone is contributing. The payment is then split b/w the Tiktok creator and the influencers who boosted it.”

Like many of you, we were already very familiar with an adjacent behavior pattern from portcos and partners asking us to engage with (i.e., like, comment, repost) their announcements on social media.



Examples of “requests for boosts” within the EV3 ecosystem.

However, we never considered the potential of simplifying and systematizing this emergent behavior with software and crypto rails. That is, until we met [Ray](#), [Ian](#) and [Vincent](#), co-founders of Daisy.

Daisy’s [mobile app](#) is extremely simple, with just one core function: **getting influencers paid for liking, commenting and reposting sponsored content**, as quickly as possible. There are hundreds of existing influencer marketing software tools out there, but there’s one trait they all share in common: they’re slow. From the time an influencer onboards to a new platform, finds a compatible deal, executes a campaign, verifies the results, and finally sees cash land in their bank account, weeks if not months have passed by. Daisy leverages crypto rails - [zkTLS](#) for real-time attribution and stablecoins for real-time payments - to reduce influencer’s time from onboarding to cash in hand to less than 24 hrs, and faster in the future.¹¹ Fast payouts means influencers create sticky habits around the app.¹² In the words of one influencer:

“With Daisy, I make \$50 every time I brush my teeth.”


Today, Daisy’s core user base participates in several campaigns per week and earn thousands of dollars per month through the app.







All

from an app that simply makes it easier for influencers to collaborate in ways they were already doing.

We see evidence of this phenomenon outside the portfolio too. In March, [Whop.io](#) launched a [clipping marketplace](#) that made it easier to do what a set of users were already doing informally: “clipping” long-form streams from larger accounts and reposting the content as short-form viral clips on social media. At launch, Whop ran a campaign with [Druski](#) (10M IG followers) to reward clippers of his new video with a total of \$5K in rewards. The campaign generated 40M views, or a [CPM of \\$0.12](#) (per 1K views), roughly a [40-60x improvement](#) from the average Meta or Google ad. Whop’s clipping product alone is rumored to be doing \$800K/mo (\$10M run-rate) GMV after a quarter and the company overall reported \$108M/mo (\$1.3B run-rate) GMV and was recently [valued at \\$800M](#) in a round led by Bain Capital Ventures. On the back of Whop’s momentum, a fresh cohort of (pre)seed companies like [Tunni](#), [Sploot](#), [Blockbook](#) and [Virality](#) are building next-gen clipping platforms. All this from such a simple idea! Video clips.



We've made the [painful mistake](#) of dismissing early-stage ideas as *too simple* in the past. When we met the [Grass](#) team  we thought a “browser extension that resells your bandwidth” was not defensible enough - *too simple* - to work. Even if the network scaled to millions of browser extensions, there are tens of thousands of existing apps out there with similar-sized install bases that could integrate the same functionality overnight if they wanted, let alone existing universities and ISPs sitting on hundreds of millions of “clean” (dormant) IPv4 addresses... or so we thought. Last year, Grass launched arguably the most successful [DePIN token](#) of the cycle to date - with a valuation 50x+ higher than at the time of our initial decision - and at least eight of our portcos have followed in their footsteps and launched browser extensions. We severely underestimated the size of the opportunity to build value-added products like data scraping and live context retrieval on top of a commodity network.

Now, when we see networks where **users contribute in an extremely simple** and low-friction way, yet the network is solving a complex coordination problem that's **valuable for a broad set of use cases**, we lean in rather than out. With Daisy and Whop, while the individual user behavior of clipping a video or reposting a story is simple, the problem of coordinating hundreds or thousands of influencers with the right followers to do it on a single post is highly complex, and solving it is valuable for every brand that advertises on social media. 



3. Measurement is the moat. It's relatively easy to get marketing departments and agencies as pilot customers, but the only way to graduate from the 5% experimental bucket to the 95% core bucket is to prove your channel drives concrete results—in advertising circles, they call this *attribution*. We believe we are in the early phases of the third wave of adtech¹³, where [privacy-preserving cryptography](#) replaces the current intrusive web2 tracking technologies, giving users the choice to opt-in to sharing their data, while at the same time providing ever-more granular (and therefore valuable) data to advertisers.

Both EarnOS and Daisy use [Opacity's SDK](#) to measure hyper-granular data on the results delivered by every incremental dollar of marketing spend, and we think many more companies will join them in the future. EarnOS requires users to submit a verifiable cryptographic proof of engagement before claiming their rewards. Without Opacity, brands would only be able to verify users' engagement on their own first-party digital real estate (website, app, etc). With Opacity, brands are able to measure and verify - with users' consent - engagement on third-party digital real estate, opening up the design space for attribution to data that was previously impossible to access, gated behind logins and walled gardens. We sometimes call this a “vampire attack”, because many early users of Opacity's SDK wanted to identify and recruit high-value users from competitors with, e.g. with outsized cash or token incentives. With Opacity, any data available to users (even behind a login) is made available to 3rd-party developers: the challenge left for adtech startups is to analyze and present the data in a way advertisers understand, and ideally in a way - i.e., faster and/or with more granularity - than they get from their existing channels.

¹³ The first wave of attribution was analog, i.e. firms doing manual surveys or panels to estimate each campaign's impact. The second wave was digital, i.e. advertisers storing a pixel (cookie) in your browser that tracks your session across multiple sites. The third wave is cryptographic, i.e. advertisers prove, cryptographically, that they catalyzed certain user actions without compromising users' privacy. [Liveramp](#) (\$2B market cap) is an early winner of the third wave. Their [RampID](#) product allows websites to hash user login details (e.g., email, phone number) and resolve that hash against other websites in Liveramp's network, effectively tracking users across sites and devices where they use the same login. Liveramp owns the resolver and generated \$745M revenues in 2024.

Daisy is a great example. You may have noticed that social media sites provide more [granular analytics](#) on your own posts vs posts from other accounts, both explicitly in their [analytics dashboards](#) and also implicitly, e.g. only the poster of an Instagram story can see which specific accounts viewed and liked it. Daisy takes login-gated data shared by its influencers, combines it with public data about the post, and uses the resulting engagement graph both to validate campaign results to brands and agencies and also to reward influencers commensurate to the value of the engagement they bring. Because Daisy “sees” the combined overlapping engagement graph “through the eyes” of every participating influencer, they’re able to deliver an unprecedented level of detail about the types of users that brands are reaching, and do so in a fraction of the time marketing executives are accustomed to. To lead the charge in deepening this moat, Daisy recently recruited two [data scientists](#) from TikTok’s video analytics team.

We could write twelve more pages about attribution, but we’ll leave it to interested readers to check out our 20-page ‘initiating coverage’ report, [Advertising from First Principles](#). Venture funding for adtech has fallen 80-90% over the past decade, which makes us all the more hungry to hunt for overlooked founders building in the sector. Until the world catches up to the ramifications of verifiable, interoperable data via zkTLS (it seems to be [catching up](#) on stablecoins already...), there will be opportunities for outsized returns by backing the most ambitious founders exploring new business models around this technology. Except us to keep writing about and investing in crypto-enabled adtech as we further develop our thesis.

*** Portfolio Update: DAWN ***

This week, our Fund I portfolio company [Dawn](#) (built by [Andrena](#)) unveils a new product: the [Black Box](#).

The Black Box is best understood with an analogy. Some of you may have solar panels on your homes. The economics are pretty simple: the typical residential solar install costs \$25K and saves homeowners from spending \$150/mo in electricity bills, paying for itself on a cash basis in 15 years. For the rest of the system’s useful life, typically another decade, electricity is effectively “free”.¹⁴

The Black Box creates similar **solar-like economics for wireless**. Retailing for \$2K, it provides reliable, high-speed internet and allows homeowners to replace their existing \$50-75/mo internet bill with a \$30/mo payment to DAWN. Resulting savings of \$20-45/mo means the Black Box pays for itself on a cash basis in 4-8 years - twice as fast as solar - and provides “free” internet thereafter for the rest of its useful life.

That’s not all the Black Box does. Equipped with industry-standard WiFi, CPU, GPU, and SSD modules, the Black Box is the mother of all DePIN devices, capable of mining across dozens of networks in parallel. At launch the device will mine a diversified basket of rewards from 10 leading DePINs, with more to come. All users need to do is plug it in; set and forget. Boosted by incremental yield from mining on third-party networks, we believe the cash payback on the Black Box will be measured in **months rather than years**.

At a scale of 10K+ devices, the Black Box becomes a **crypto platform**. New DePINs will use it as a launchpad, allocating rewards to BBs to bootstrap their network capacity and tokenholder base. DeFi protocols will collateralize a basket of mining rewards into innovative new lending products for miners. The Black Box “app store” will allow miners to search and discover hundreds of mining and DeFi apps. Many projects have attempted parts of this vision, but no one yet has had the team and scale to execute.

¹⁴ US figures. Most homeowners use free government financing in the form of tax credits and low-cost financing from consumer solar lenders to fund the majority of installation. Interest rates on solar loans, depending on homeowners’ credit profile and cost of energy, are typically roughly equal to the system’s unlevered yield. In any case, it’s not relevant for the purposes of the analogy.

At a scale of 1M+ devices, the Black Box becomes a **sovereign computing** platform. Its computing resources will no longer be used solely for mining crypto; users will host their private personal data on its local storage and run *sovereign apps* on its local compute. Humanoids will navigate your home with 3D maps hosted on the BB, a device within your control, rather than on Tesla's servers. Your intimate medical questions will be answered by LLMs running on the BB, rather than on OpenAI's servers. As a sovereign computing platform, the BB offers users something arguably better than cryptographic security: physical security. If you don't like what it's doing, you simply unplug the box... No fancy cryptography needed.

The vision for a sovereign computing network is not [new](#) or [unique](#), but all prior attempts focused on deploying on existing machines and controlling the operating system or application layer. No one yet has had the combination of talent, capital, and distribution to take a real swing at a **hardware-first approach**. DAWN's team, led by [Neil Chatterjee](#), have built and deployed proprietary wireless devices in production at [Andrena](#) for 5+ years, driving efficiencies that have supported industry-leading gross margins despite undercutting incumbents by 30-40%. On the capital front, Dawn has raised >\$30M since we first invested, making it one of the best-capitalized projects in DePIN (top-5 by capital raised post-'21). On distribution, Dawn has grown its [browser extension](#) to an install base of >2.5M and its [X](#) following to >500K in a year.

For most Americans, and many other people elsewhere around the world, the **Black Box makes financial sense on the basis of wireless savings alone**. For crypto investors, it earns yield from a diversified basket of leading DePINs, providing **index exposure to the entire sector**. If animal spirits ever return to DePIN like they did in 2021, a huge audience of speculators - the same people who rushed to deploy Helium IoT hotspots back then - will see it as a **box that prints money**. There's never been a device with this value proposition, and we expect sales growth will be lumpy as the company experiments with new channels, messaging and markets. Over the next 18-36 months, the focus is on scaling the BB first to a hundred thousand devices, then a million. If you'd like to order one, click [here](#).

*** *Outro* ***

The landscape today is different from when we started EV3. The underlying crypto rails perform much better, yet the public sentiment towards crypto is materially worse. The market is more skeptical of tokens, demanding revenues and disclosures, yet indiscriminately bids crypto-exposed equities. Regularity clarity is (almost) finally here for projects with fundamental value, yet memecoins dominate in trading volumes, market cap and attention. The deepest and most liquid exchanges, where price discovery happens, are now the onchain ones. DePIN went from being too narrow of a term, to too broad of one.

Your partners,
Sal & Mahesh