

Stablecoins are now payment rails.

A field guide for teams shipping stablecoin products: what these rails actually consist of, why reading them isn't so straightforward, and what it takes to build with confidence.

Executive Summary	1
<hr/>	
Stablecoins are no longer a niche asset.	2
<hr/>	
Stablecoins are real payment rails now.	3
<hr/>	
Adoption is no longer speculative.	4
<hr/>	
Why adopt, and the real-world use cases at scale.	5
<hr/>	
Stablecoin rails 101, where chain data comes from.	6
<hr/>	
The hard part nobody sees: reading chain data reliably.	9
<hr/>	
What are the tradeoffs of building vs. buying?	13
<hr/>	
A closing note on regulatory tailwinds	14

Executive summary

Over the last two years, stablecoins quietly stopped being a crypto curiosity **and shifted into a real set of payment rails**. In 2025 they moved **\$33-35 trillion** on-chain, more than twice Visa's annual volume, though, as this report shows, most of that is automated churn rather than payments. Outstanding supply crossed **\$315 billion**, and **90% of surveyed financial institutions** now report they're live, piloting, or planning a stablecoin product. The question is no longer whether stablecoins matter. It's how fast can we ship.

While the above figures are often touted as signals of the meteoric rise stablecoins have seen, there has been far less attention paid to a harder truth: **transfers alone are the easy part**. On the back end of each transfer, every payments product has to answer three fundamental questions, continuously and, more importantly, correctly:

- *What is the sending account's balance, right now?*
- *Did this payment actually settle and will it stay settled?*
- *Does our ledger match the chain?*

The catch is that **a raw chain read answers none of them at the scale, freshness, and certainty a payments product needs**. A single balance is one cheap lookup, but you need every customer's balance kept current. Settlement is not a value the chain records; it's a probability you track until it hardens. Reconciliation is your ledger checked against the chain. All three are computed off-chain from a noisy, self-correcting stream of events. "Just read the chain" fails not because the data is missing, but because the answers your product books on are derived, and deriving them correctly at scale is the work.

This report covers the stablecoin opportunity in fintech terms, then explores what no other adoption report does: the **operational reality of the rails**, told through the use of live chain data.

Three numbers that frame the opportunity:

\$315B

stablecoin supply outstanding · *DefiLlama, Jun 2026*

\$33-35T

raw transfer volume in 2025 (+72% YoY), of which ~\$390B is genuine payment activity · *Artemis; McKinsey + Artemis, Jan 2026*

9 in 10

surveyed institutions live, piloting, or planning · *Fireblocks, Mar 2025*

Stablecoins are no longer a niche asset.

What was once a fringe asset class has quickly grown into a material store of value, and that growth keeps accelerating alongside adoption. As of June 2026, outstanding stablecoin supply sits at roughly **\$315 billion**, up from about \$160 billion in late 2024, and an amount **equivalent to more than 1% of all US dollars** now exists in tokenized form.

Taken together, **stablecoin issuers currently hold nearly \$200 billion in US Treasuries**, enough to rank them among the top 20 sovereign holders (US Treasury TIC data, March 2026; a16z State of Crypto 2025).

Across the issuer landscape, two issuers carry almost all of it: **USDT at ~59% (~\$187B)** and **USDC at ~24% (~\$76B)** representing **83% of total supply**.

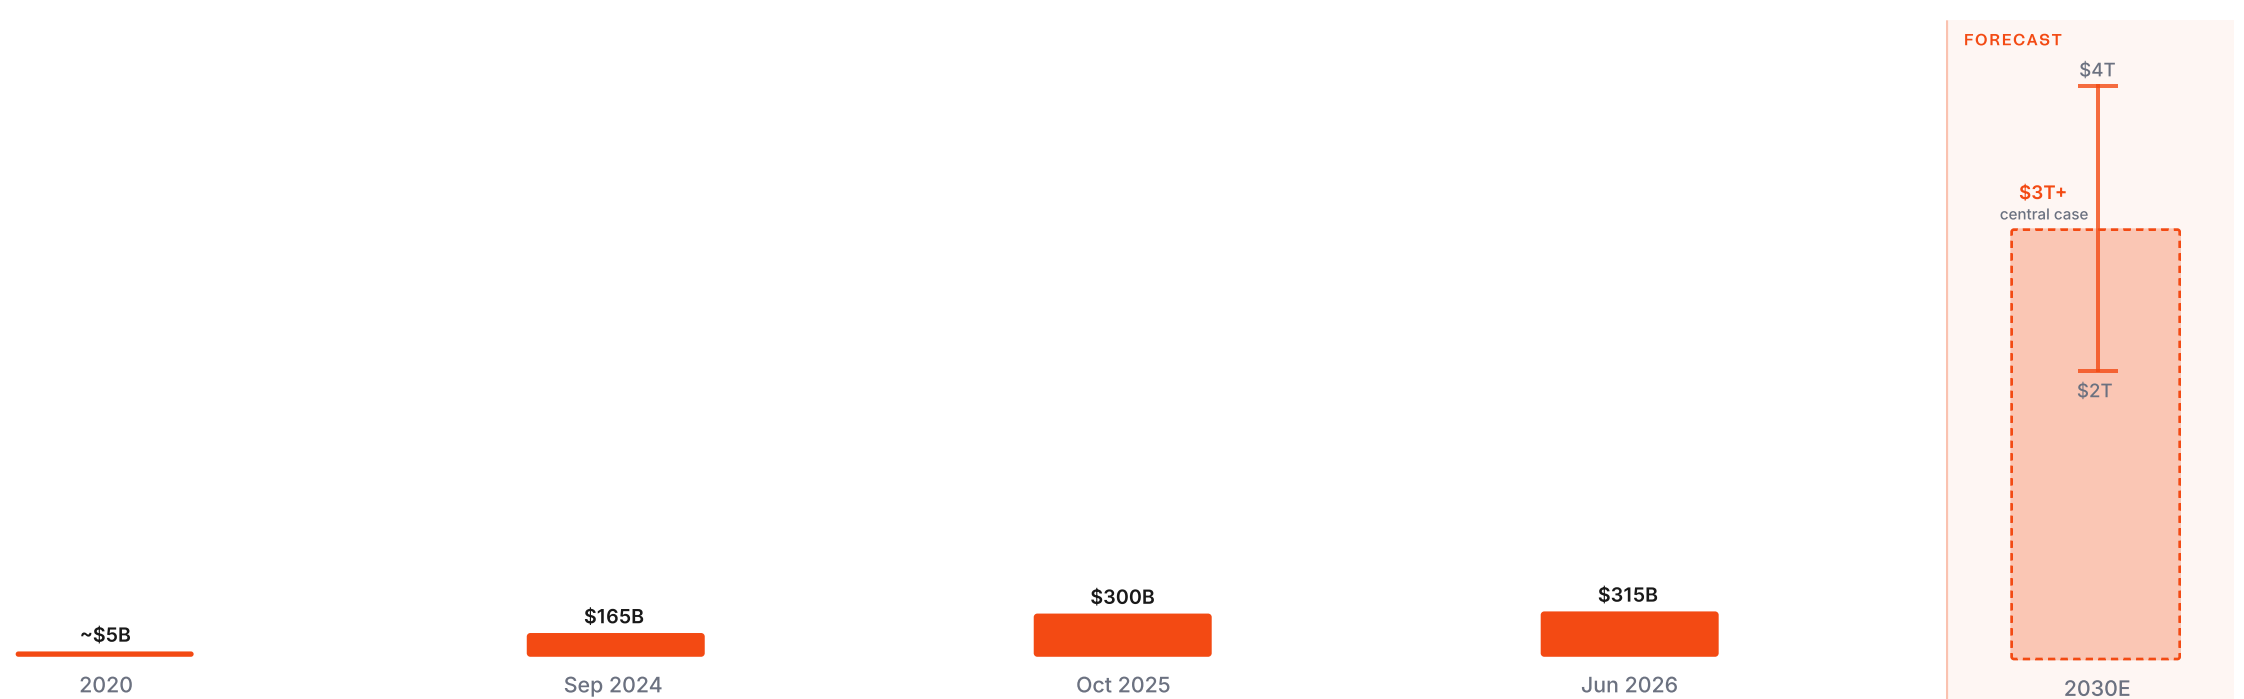
As the space grows, incumbents are starting to move in. Société Générale issues the MiCA-regulated EURCV and USDCV, Japan's three largest banks are piloting a yen stablecoin, and the banks behind Zelle have

unveiled ZelleUSD for cross-border use, slated for India and other markets. All of it runs on public rails. Running alongside them is a separate track of permissioned, bank-only settlement networks (Finality, Partior) built as an alternative to those public rails.

None of it has yet shifted the centre of gravity: USDT and USDC alone continue to hold ~83% of the ~\$315B market, a concentration that has only grown as more issuers pile in.

Recent growth and current size are worth noting, but **the real story lies in the growth trajectory we're on**. Analysts at a16z project supply growing **~10x to over \$3 trillion by 2030**, and they're not alone: Treasury Secretary Bessent (Nov 2025) and the research desks at Citi, Standard Chartered, and Bernstein all land in the **\$2-4 trillion** range. We're still in the early innings, as stablecoins become commonplace in global financial plumbing.

Figure 1 · Stablecoin supply outstanding, 2020 → 2030 (\$B) — actuals and forecast



Sources: DefiLlama (Jun 2026, \$315B); a16z State of Crypto 2025 (\$300B+); Castle Island / Brevan Howard (Sept 2024, ~\$165B), 2020 figure approximate. 2030 forecast: a16z projects ~10x growth to over \$3 trillion; the \$2–4 trillion range spans Treasury Sec. Bessent (Nov 2025), Citi, Standard Chartered, and Bernstein. Bar shows the ~\$3T central case; whisker marks the \$2–4T range.

Stablecoins are real payment rails now.

In 2025, stablecoins moved an estimated **\$33-35 trillion** on-chain, **up 72% year over year** (Artemis; McKinsey + Artemis, Jan 2026). Taken at face value that's more than twice Visa's annual volume. But face value is exactly the trap this report is about: **most of that figure isn't payments.**

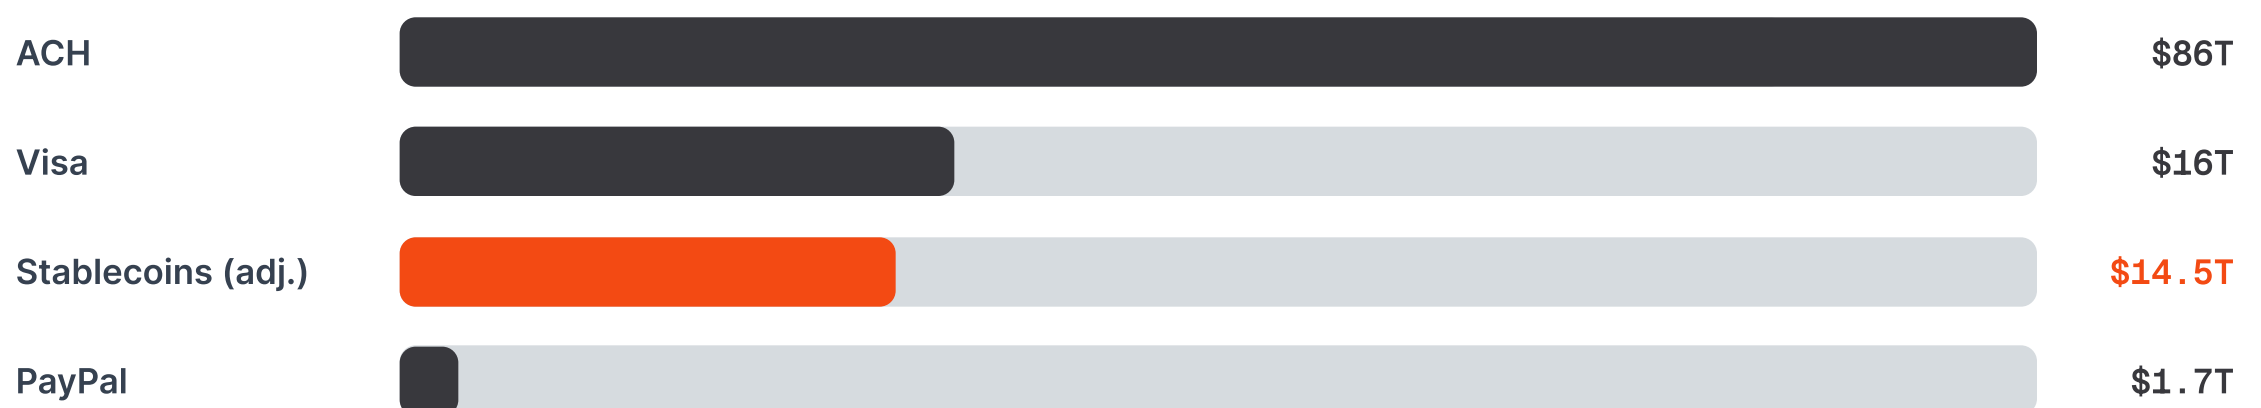
It's arbitrage, market-making, and intra-protocol churn, automated flows that never touch a customer. Strip them out and the genuine payment activity is roughly \$390 billion (McKinsey + Artemis): smaller, but growing far faster where it counts. **B2B is up 733% and card spend up 673% year over year** (the slice we size on Page 5). Each real transfer settles in seconds, for a fraction of a cent.

That gap, between **\$33-35 trillion** that moved and **~\$390 billion** that was actually a payment, is not an accounting footnote. It's the single most important fact about reading stablecoin data: **the headline number and the real number differ by nearly two orders of magnitude, and telling them apart takes infrastructure, not a dashboard.** Anyone can pull the big number.

Knowing which dollars are payments, which are a bot, and which are the same dollar counted twice is the entire topic of a the section starting on Page 9.

Despite this growth, not everyone is convinced this should be happening at all. The Bank for International Settlements, the central banks' central bank, devoted a chapter of its 2025 Annual Economic Report to arguing that stablecoins **fail the classic tests of sound money**: singleness (a dollar is always a dollar), elasticity (the system can expand credit on demand), and integrity (resistance to illicit use). Singleness and elasticity are fair critiques of stablecoins as a monetary system, a replacement for money itself. They land far softer against the use case we're discussing: **moving dollars that already exist, faster and cheaper than the wire that moves them today.** Integrity is the real exception. Illicit-use resistance is a genuine obligation, but an infrastructure one, not a reason the rail fails. A stablecoin doesn't have to be the future of money to be a better rail for the dollars we already have.

Figure 2 · **Annual payment volume by rail (\$T) — stablecoins vs incumbent rails**



Stablecoins settle more than 7X PayPal and on the same order as Visa — through ACH still dwarfs every card-era rail. The \$14.5T figure is the bot-adjusted series, not raw on-chain volume (~\$99T).

Methodology: raw and real-payment volume from Artemis and the McKinsey + Artemis joint analysis (Jan 2026). Raw = total on-chain transfer value; "real payments" strips trading, internal rebalancing, and contract loops. We cite both rather than a single "adjusted" figure because no adjusted number commands consensus.

Adoption is no longer speculative.

Powered by a combination of regulatory tailwinds, a superior settlement experience, and better unit economics, stablecoins are no longer a pilot-stage story. Their growth is being fuelled by accelerating institutional adoption and institutions your board benchmarks against have already shipped:

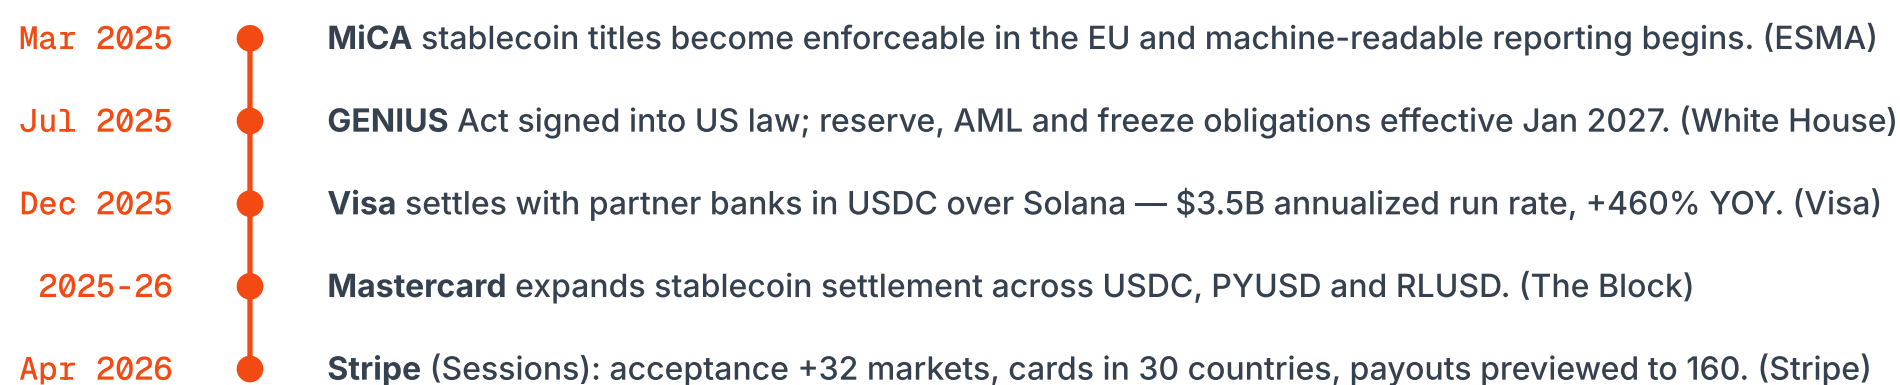
- **Visa** settles with partner banks in **USDC over Solana**. As of Nov 30, 2025, its monthly stablecoin settlement passed a **\$3.5 billion annualized run rate, up 460% year over year** (reportedly ~\$4.5B by Jan 2026). Stablecoins now make up roughly **1 in 5 dollars of Visa's crypto-card settlement**.
- **Stripe** expanded stablecoin **acceptance to 32 more markets**, issues stablecoin-backed cards in 30 countries (usable anywhere Visa is accepted), and **previewed payouts to 160 countries** in stablecoins, versus just over 100 in fiat.
- **Mastercard** expanded stablecoin settlement across USDC, PYUSD, and RLUSD, confirming this is a **both-networks shift, not a single-network bet**.

- **Citi** went live in October 2024 with **Token Services for Cash**, a tokenized-deposit rail that moves corporate dollars between its branches across four countries around the clock, with no cut-off windows. By November 2025 it had processed billions of dollars in transaction value: regulated, deposit-based settlement already in production, not a pilot.

Beyond these giants, the broader market is moving with them. In a survey of 295 senior financial-sector leaders (61% C-suite), **9 in 10 said they are live (49%), piloting (23%), or planning (18%)**. The **#1 cited benefit was faster settlement (48%)** ahead of liquidity (33%) and cost (30%).

Whatever the reason, there's a good chance your closest competitors and peers are already out there actively shipping real world use cases powered by stablecoins.

Figure 3 · Stablecoin payment milestones, 2025 → 2026



Sources: ESMA (MiCA); White House (GENIUS Act, Jul 2025); Visa press release (Dec 2025); The Block (Mastercard); Stripe Sessions (Apr 2026).

Why adopt, and the real-world use cases at scale.

The motivations for adopting here aren't speculative; they're the same problems payments teams have always struggled with, and in emerging markets they're **explicitly non-crypto**. In a five-country survey, users named **dollar access (47%), currency conversion (43%),** and yield (39%) as their top reasons for holding stablecoins. In Latin America, **71% of stablecoin users use them for cross-border payments** against 49% globally.

Strip out trading and internal flows and you can size the part that's actually retail and commercial payments: McKinsey and Artemis put **real payment activity at ~\$390 billion in 2025**. The growth is concentrated exactly where fintechs operate: **B2B payments hit \$226B, up 733% year over year,** and **stablecoin card spend reached \$4.5B, up 673%**. Among corporate users, **more than 2 in 5 report cost savings of 10% or more** (EY-Parthenon, 2025).

The benchmark they're beating is a familiar one: global remittances run **~\$944 billion a year at an average 6.4% cost** to send \$200. **Stablecoins move the same value for pennies, and settle in seconds.**

Whichever of these you build (remittance, neobank, wallet, custody, or trading), you hit the same problem the moment it goes to production: the chain doesn't hand you clean, settled data.

A primary cut from our own data:

82-99% of stablecoin transfers come in under \$3,000

Across every credible retail rail we measured, the size profile of retail and commercial payments, not institutional settlement.

Figure 4 · The benefits, by the numbers — stablecoins vs. traditional rails

Payment Rail	Settlement (to final)	Cost	Availability
Stablecoins	~seconds	<1¢ – a few ¢	24/7/365
ACH	1–2 business days	\$0.20 – \$1.50	Bank days only (~250/year)
Wire — intl (SWIFT)	~1 day, up to 5	\$25 – \$50+ in fees	Bank days only
Card (Visa)	Instant auth · T+1-3 payout	~1.5 – 3.5%	24/7 auth · bank-day settle
PayPal	Instant in-app · T+1-3 payout	2.9 – 3.5% + ~30 ¢	24/7 in-app · bank-day payout

Availability is only half the liquidity story: stablecoins settle T+0 with no pre-funded nostro or correspondent accounts. Freeing working capital that legacy rails leave idle (industry estimates put roughly \$27T locked in cross-border pre-funding worldwide).

Stablecoin rails 101, where chain data comes from.

This is a lay-of-the-land section, for readers who don't spend their days in a block explorer or are relatively new to the space. It's the one place in the report we slow down and explain how chain data is produced, where it's stored, and how it reaches your product. Everything after it assumes you can picture the stack.

The balance problem: solved once, reopened by scale

A blockchain keeps two things at once, and conflating them is where most teams go wrong. It keeps an **append-only** journal, every event that ever happened ("address A sent 100 USDC to address B in block 21,000,000"), and it keeps **state**: the current value of every account, updated as those events land. So the one number your product needs, how much USDC address B holds right now, **does exist on-chain**.

On Ethereum you call `balanceOf(B)` and get it back; you can even ask as of a past block. The old mental model, that you must replay every transfer from genesis to learn a balance, is how you'd rebuild it from the journal alone. But the token standard solved that years ago by keeping the balance in state.

So if one balance is a single cheap lookup, where's the hard part? **Scale, history, and currency: the three things a product actually runs on.**

- A stablecoin product never needs one balance once. It needs every customer's balance, **continuously correct to the chain tip**. You cannot poll millions of addresses on every new block; that doesn't scale. To keep balances current you have to know **which addresses changed in each block** and update only those, which means

decoding and ordering every transfer as it lands. The journal comes back, not because balances are missing but because keeping a whole population of them fresh is an indexing problem.

- State answers "*what is the balance now.*" It cannot answer "*show me every USDC payment this account received this month,*" or "*what was this balance the instant that disputed payment landed.*" History and analytics live in the journal, which has to be decoded, ordered, and stored before you can query it that way.
- And a balance true at the tip can be **revised** seconds later (see below). Correct-at-scale means correct *through* those revisions, not just at one instant.

This is the line between a node and an indexer, and it answers the obvious objection, *why not just call the chain yourself?* You can: for one address, once. (A node even keeps internal indexes for exactly those lookups.) What it will not hand you is your entire book of balances kept current, the month of history behind each one, reconciled to your ledger, across every chain you support. That gap is the whole job.

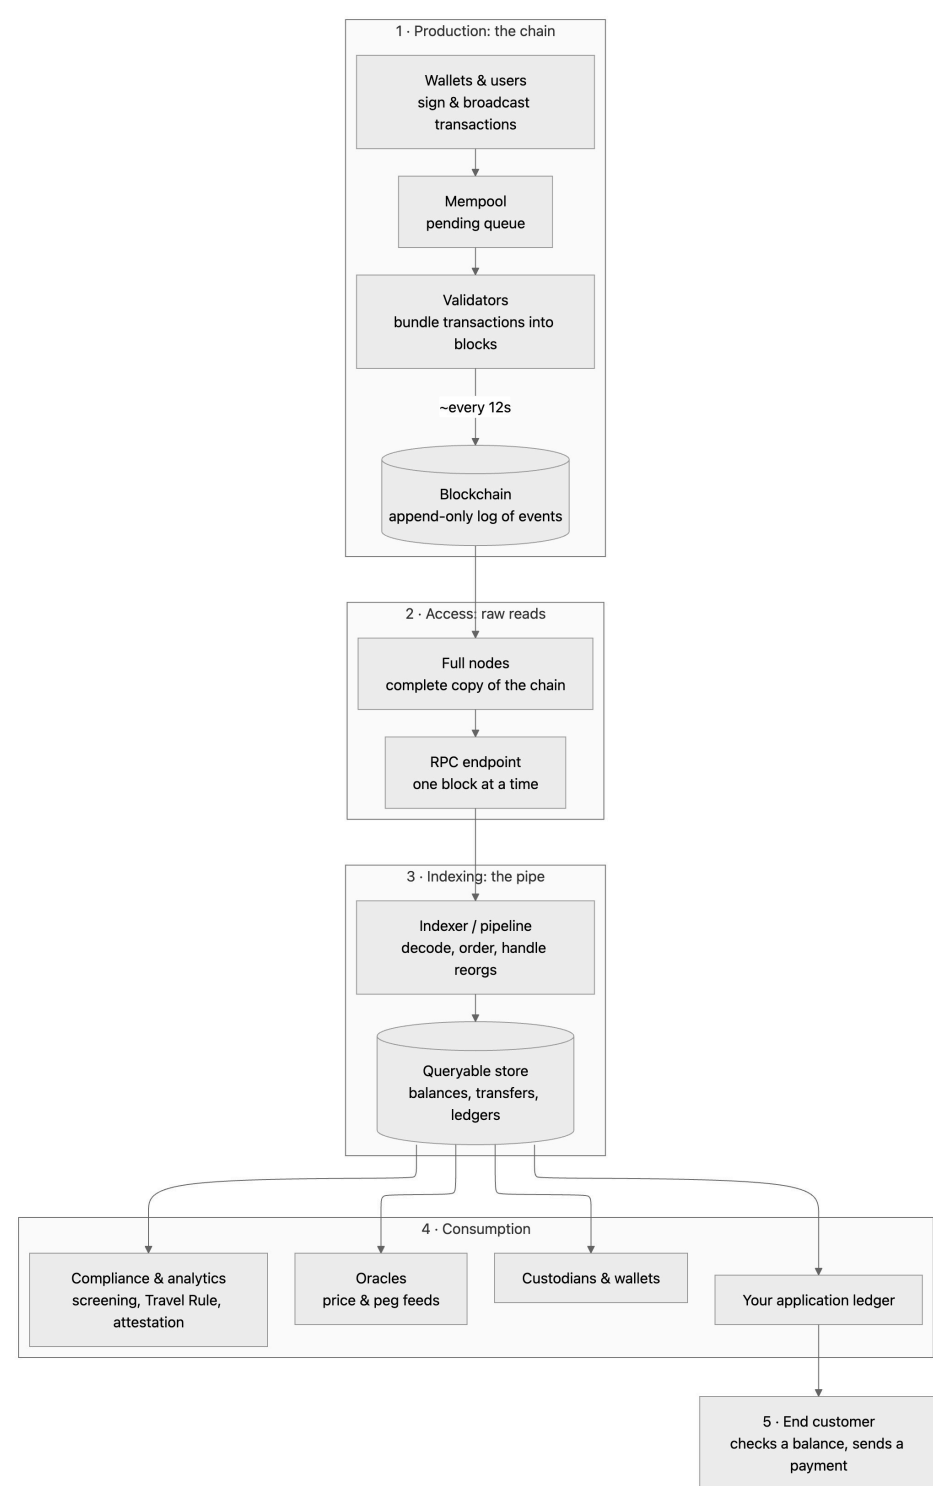
Those events are produced in batches. Roughly every twelve seconds on Ethereum, one participant, a *validator*, bundles the latest transactions into a **block** and appends it to the chain. The block is the batch; the individual transfers inside it are the *events* (also called "logs"). Stated otherwise, the validator is whoever posts the next page of a shared journal, and a stablecoin transfer is a single line item on that page.

The stack between the chain and your customer.

When it comes to blockchain data, you almost never read the chain directly and you wouldn't want to. Between the raw blocks and the balance a customer sees in your app sits a stack of infrastructure, each layer doing a job the one above it depends on:

- **The chain & validators** produce blocks. This is the source of truth, such as it is.
- **Nodes & RPC.** A node is a full copy of the chain and a remote procedure call (RPC) is the request interface you talk to it through. It answers one narrow question at a time ("give me block N"). That's fine for a point lookup, and an RPC node even keeps internal indexes for exactly those. What it has no answer for is the shape your product needs: "every USDC payment this account received this month," across millions of accounts, kept current. One narrow question at a time is not a query engine.
- **Indexers & pipelines** (the "pipes") read every block as it's produced, **decode** the raw events into something legible, **order** them, **correct** them when the chain revises itself, and write the result into a store you can actually query. **This layer is the entire subject of the next page.**
- **Data consumers** are everything that needs clean, current chain data: **custodians and wallets** (hold keys, must show accurate balances), oracles (feed prices and peg data between off-chain and on-chain, e.g. confirming a stablecoin is still worth a dollar), **compliance and analytics** (screening, Travel Rule, reserve attestation), and **your own application ledger** (the settlement state your product runs on).
- **End customers** sit on top, checking a balance or sending a payment. They never see any of the above infrastructure, just the end application and user experience (UX).

The stack between the chain and your customer:



Raw RPC is layer 2; the indexer at layer 3 is what turns one-block-at-a-time reads into the balances, transfers, and ledgers everything above it queries.

Why a payment isn't settled the second you see it.

Despite often claiming "instant" settlement, blockchains do not settle instantly. Speed is the headline selling point for stablecoins, so this often catches teams off guard when they arrive from card or bank rails, **until they notice the mechanics are familiar**. On card networks, every charge has two states: **authorization** (it shows as "pending") and **settlement** (it "posts," often a day or two later). A pending auth can still vanish. Blockchains have the same split just compressed into minutes and renamed:

- **Inclusion:** your transfer lands in a block. It looks done. This is the "pending" state.
- **Finality:** enough later blocks have stacked on top that the network can no longer change its mind. This is "posted."

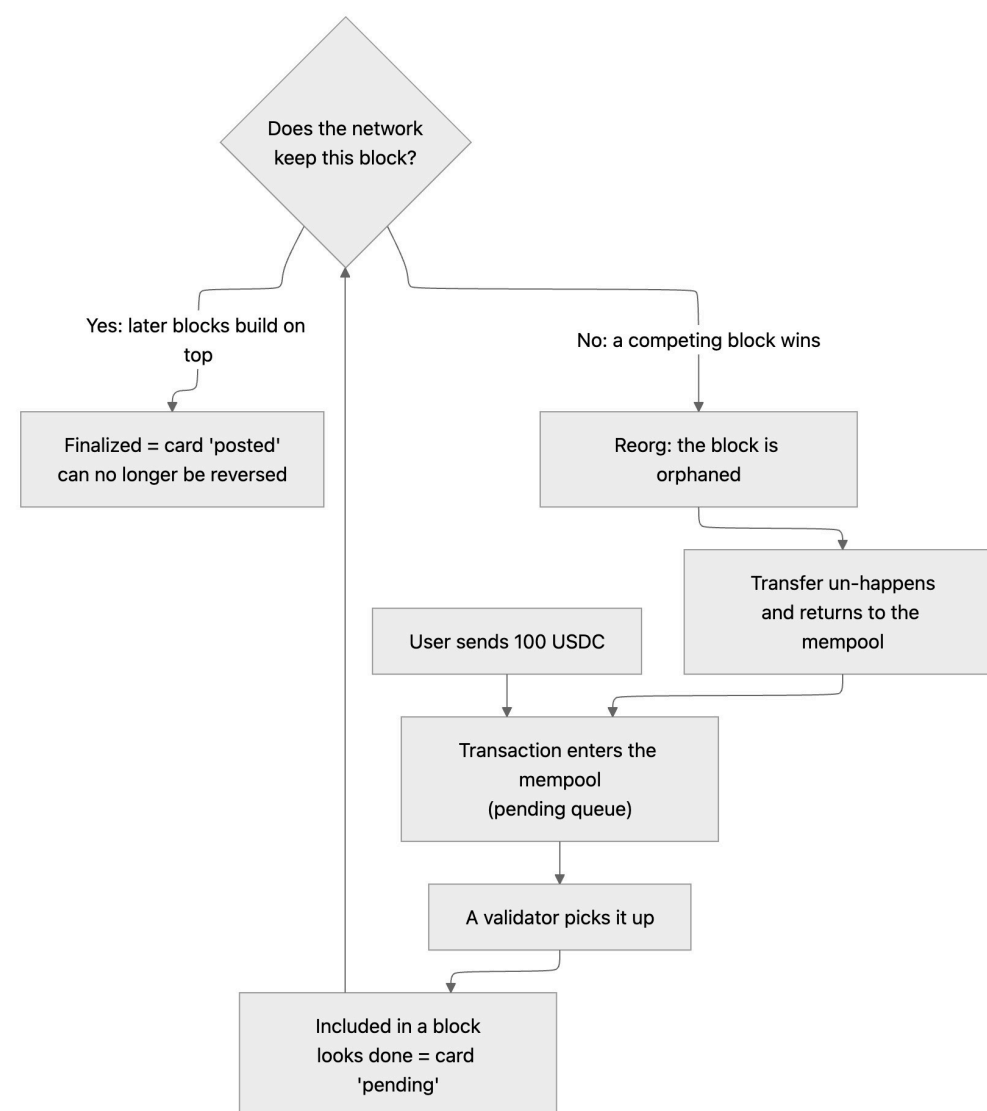
Inclusion and finality are different events anywhere from a couple of seconds to several minutes apart, depending on the chain. In between, the transfer can still be reversed. A ledger that treats "*I saw it in a block*" as "*it settled*" is booking entries the chain is still allowed to take back.

When the chain changes its mind: reorgs

Why would the chain take something back? Because occasionally **two validators produce a block at nearly the same moment**. The network briefly disagrees about which one is real, converges on a single winner, and discards the other. Every transfer in the discarded block is **un-happened**, reversed as if it never settled. In card terms, it's a pending charge that drops off before it ever posts. The crypto name is a reorganization, or reorg.

For a primer, the mechanic is the whole point: **inclusion is a strong signal, not a guarantee**. How often reorgs actually happen, and how much settled-looking money they quietly put at risk, is what the next section measures.

The life of a transfer, including the path nobody designs for:



Most payments fall straight through to finality. The dashed reality is the loop on the right: a transfer that looked settled gets reorged out and re-enters the queue. the failure mode the next section puts a dollar figure on.

The hard part nobody sees: reading chain data reliably.

What separates a demo from a reliable stablecoin product?

Reliable data.

As we said at the outset, shipping a stablecoin product means continuously answering the same three questions (balance, settlement, reconciliation), and, as Page 6 laid out, a raw chain read answers none of them at the scale and certainty a product needs. A balance is a lookup; your whole book of balances, kept current is not. Settlement isn't a recorded fact; it's a probability that hardens over time. Reconciliation is your ledger against the chain, by definition off-chain.

That's what makes everything below hard: the answers are derived, and the chain keeps changing the inputs underneath you.

When the chain takes it back

Every payments system has one promise at its core: **once money arrives, it stays arrived**. Blockchains quietly break that promise.

The primer described the mechanic: two blocks are produced at once, the network keeps one and discards the other, and every transfer in the discarded block is un-happened. **In payments terms, it's the ledger entry you already booked being silently retracted**. What the primer didn't do is put a number on it.

If your system booked any of those transfers as "settled" the moment it saw them on-chain, **your ledger was wrong** for a few seconds to a few minutes. You credited an account for money the chain then erased. For a payments company that isn't an abstraction: it's a reconciliation break, a possible double-credit, a support ticket, and, under the GENIUS Act's transaction-reporting regime, **a discrepancy you have to explain**.

What makes reorgs genuinely hard, not just inconvenient is that **a reorg is not an objective event**.

Whether you witness one depends on how closely you were watching. A system reading 50 seconds behind the tip sees a clean, settled history. A system reading

The stat nobody publishes:

In a 90-day window, 86,644 USDC and USDT transfers, carrying roughly \$6.7 billion in value, landed in Ethereum blocks that were later reorganized out of the chain. This happened across 544 distinct blocks on Ethereum alone.

This is gross transfer value briefly exposed: the dollars a naive, inclusion-based ledger would have booked as settled, concentrated in a handful of large transfers, not a realized loss. Captured by Goldsky's own indexing pipeline as it streams the chain; most data providers cannot report it, because by the time they read the chain the evidence is already gone.

at the tip, the only way to know a payment arrived now, sees the messy, self-correcting reality. That's why our \$6.7B figure is a floor, not a census: most databases overwrite reorged data the instant the correction lands. **Only infrastructure that records each retraction as it happens ever sees the number.**

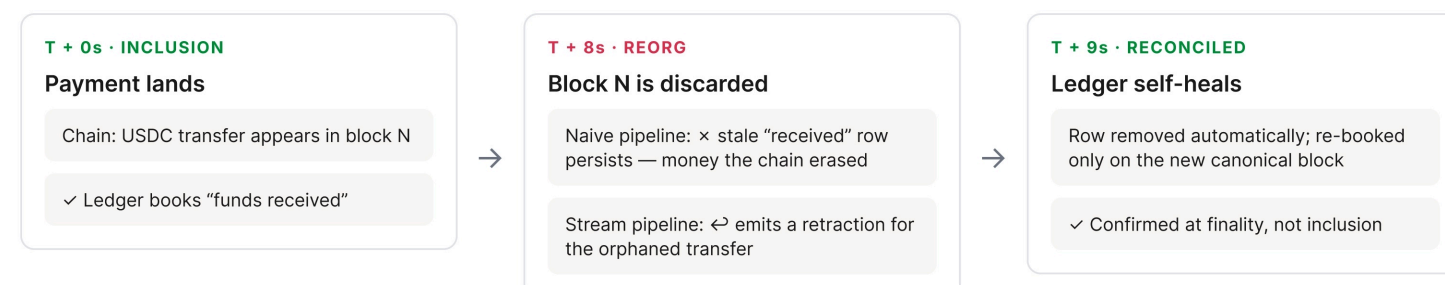
The fix is to **treat chain data as a stream that can correct itself**. Every block can be revised; when one is reorged out, the pipeline emits an explicit retraction for every transaction it contained, and the ledger heals automatically. And **settlement keys on finality, not block numbers**: a transfer isn't done because it's in a block, it's done when the network can no longer reorganize that block away. Inclusion and finality are different events that often resolve seconds to minutes apart.

How much any of this matters depends on the rail, which is its own argument for why multi-chain is hard:

Rail	Reorg exposure (90d)	Why
Ethereum	~\$6.7B / 544 blocks	PoS L1; short reorgs are routine consensus behavior, and we capture every retraction as it lands
BSC	~\$7.6M / 138 blocks	Reorgs occur on the same mechanism; smaller dollar value because transfers are smaller
Arbitrum, Optimism	None recorded	Single sequencer, no competing forks to retract; also not capturable in static tables. "None recorded," not a measured zero

A product on one chain can sometimes ignore reorgs. **A product on six chains cannot.** Each rail has a different answer, and "we'll just read the chain" has six different failure modes.

Figure 5 · **Inclusion ≠ finality — the retraction that heals the ledger**



Treating inclusion as settlement means occasionally crediting money the chain later deletes. The fix: process the chain as a correcting stream and settle on finality.


Raw data lies: the median that's off by nearly 90x

Ask a naive pipeline for the **median stablecoin transfer on Ethereum** and it returns **\$2.59**. That's not a rounding error: it's wrong by nearly 90x, and it would tell a product team that stablecoins are a micropayments rail. **The real median, once the data is clean, is \$227.**

The gap is address poisoning. Spam at industrial scale. Attackers blast out millions of near-zero-value transfers from throwaway wallets engineered to look like a victim's real counterparties (matching the first and last few characters of the address), betting that someone copies the wrong one out of their history. We confirmed it isn't ordinary dust: we found **100,137 distinct clusters of 5-or-more manufactured look-alike addresses**. In the most extreme case, a single real address was shadowed by **441 look-alikes**, a degree of clustering far beyond anything that arises by chance.

The result, over the 90 days to June 11, 2026: **~39% of all Ethereum USDC+USDT transfers were sub-cent, and at least 1 in 6 were provable poisoning spam**. Read the chain naively and one in six "payments" is an attack and your headline median is fiction. And it is **not uniform across chains**, which is the deeper point:

Figure 6 · Sub-cent ("poisoning") share of transfers, by chain — 90-day window



Ethereum reads \$2.59 median naively → \$227 once dust is stripped. Tron, the biggest USDT rail, is the cleanest. Sub-cent = value > \$0 and < \$0.01; bar length scaled to 40% for legibility; labels exact.

The chain everyone treats as the respectable institutional rail, **Ethereum, is the dirtiest**. The largest USDT rail in the world, **Tron, is the cleanest**, at 0.1%, despite being EVM-compatible under the hood. Poisoning spans **two orders of magnitude across rails**, so **there is no single "dust filter" that works everywhere**. A multi-chain product either segments per chain, or it publishes numbers that are quietly false.

Anatomy of a misleading number.

One more example to drive the point home here. A naive query reports **\$692 billion of USDC moved on Optimism in 90 days**, a number you could drop on a slide. Trace where it actually came from and it dissolves:

- **38.5%**: a single Velodrome liquidity pool (a DEX trade, not a payment)
- **17.3%**: the Across bridge relayer
- **9.7%**: Orbiter / Hop bridges
- **9.2%**: the Stargate bridge relayer

Three-quarters of it is one DEX and four bridges. Almost none of it is payments.

And when you try to trace the remainder by hand, the path doesn't even terminate. It loops through a Stargate factory ↔ pool funding cycle, **the exact point where a human needs a maintained entity graph** to make sense of the chain. Raw volume isn't payment volume, and naive tracing doesn't converge.

One model doesn't fit every chain.

The work above also **changes shape by chain family**. Dealing with multiple chains becomes a challenge for a few key reasons:

- **EVM chains** (Ethereum, Base, Arbitrum, Optimism, Polygon, BSC) share one event model, so a single decoder spans all of them.
- **Tron** rides that same EVM event model (the decoder is the same) but renders addresses in a different format (Base58, not hex). Cosmetic, until it isn't: it's the biggest USDT rail in the world, and that one difference quietly breaks address matching, dedup, and screening unless you handle it explicitly.
- **Stellar** is the opposite extreme: balances live as **native asset operations and as Soroban smart-contract tokens**, with several ways to represent the same balance, so on Stellar, "read the balance" isn't even a single question.
- **Solana** has no per-token transfer event at all; you reconstruct transfers from token-balance changes.

One event model still isn't one decoder, and two of these chains aren't EVM at all.

This is exactly the layer a managed indexer abstracts away, and it is exactly the layer a team on raw RPC rebuilds from scratch, per chain, forever.

One event model still isn't one decoder, and two of these chains aren't EVM at all.

This is exactly the layer a managed indexer abstracts away, and it is exactly the layer a team on raw RPC rebuilds from scratch, per chain, forever.

Methodology: all figures come from Goldsky-powered infrastructure. EVM chains and Tron run through our community dataset; Solana comes from CryptoHouse, the public blockchain analytics service Goldsky powers. Window: 90 days ending June 11, 2026. Counts are exact. Dollar figures are gross transfer-event value and directional, not economic volume, so we cite Artemis and McKinsey + Artemis for absolute figures. Reorged transfers are isolated by block-hash mismatch, with de-dup housekeeping filtered out (under 1%). All figures are lower bounds.

What are the tradeoffs of building vs. buying?

Building this in-house means **spending your highest-leverage engineering quarters on a layer that ships no product.**

A DIY indexer is not a weekend project. Amazon's own reference architecture for building one on AWS requires **five managed services stitched together; Cryo, Reth ExEx, Flink, MSK, and RDS all before a single line of business logic.** AWS itself notes that querying blockchain nodes directly is inefficient, and recommends evaluating existing indexers first.

And that's the happy path. The Hiro engineering team, building a production indexer, found their first design needed **six parallel global state snapshots just to handle reorg rollback**, the exact problem we were just discussing, and had to rearchitect from scratch. Their own description of building on an evolving chain: *"like building a plane while it's in flight."*

The cost that doesn't show up in the launch plan is the recurring one:

	Build in-house	Buy (managed)
Time to launch	6-12 months	< 1 week
Infra cost	\$5,000+/mo multi-chain nodes (ETH archive \$1k+/mo, Solana \$2.5-4k/mo)	Predictable subscription
Headcount	\$350k+/yr in dedicated senior engineering	None dedicated to infra
Maintenance	Permanent: reorgs, finality, backfills, schema drift, per-chain quirks	Vendor's problem, under SLA
Compliance posture	Build attestation / audit trails yourself	SOC 2, audit-ready data

A closing note on regulatory tailwinds

Historically, regulation has blocked crypto (including stablecoins) from operating as legitimate payment rails, so the instinct to read it as an obstacle makes sense. Over the last year that has inverted. Regulation is now a data requirement and every obligation on the list is a reason you need reliable infrastructure underneath you.

The GENIUS Act (signed Jul 2025, effective Jan 2027) made stablecoin issuers BSA financial institutions, mandated 100% liquid-asset backing with monthly public reserve disclosures, and requires issuers to have the technical capability to seize, freeze, or burn stablecoins when legally required. The EU's MiCA regime is already enforceable. And FATF has named stablecoins the primary on-chain illicit-finance vector (e.g., \$1.46 billion in DPRK-linked theft in 2025, primarily stablecoins). Underneath every one of these is the same demand: auditable, transaction-level chain data.

Five obligations that each require transaction-level chain data:

- 1. Reserve attestation:** daily-granularity auditable reserve data; monthly CPA-examined reports with CEO/CFO certification (OCC NPRM). Effective Jan 2027.
- 2. Travel Rule:** structured originator/beneficiary data linked to on-chain transfers, retrievable on demand (85+ jurisdictions; the EU has no de-minimis floor).
- 3. AML / sanctions:** real-time block / freeze / reject capability across primary and secondary markets; \$5,000 SAR threshold (FinCEN/OFAC NPRM).
- 4. MiCA Annex 4:** quarterly machine-readable (iXBRL) reserve + transaction reporting. Enforceable now.
- 5. FATF monitoring:** supervisors now expect stablecoin-specific transaction monitoring as table stakes.

One model doesn't fit every chain.

Here's what those obligations look like as data. The three Travel Rule regimes draw their lines in different places: the EU attaches data obligations to every transfer (zero floor), FATF guidance clusters around \$1,000, and the US BSA threshold is \$3,000, ignoring everything below it. So the share of transfers under \$3,000 is, precisely, the activity the US rule leaves uncovered that the EU rule captures.

Chain	Transfers (90d)	<\$1k	\$1k-\$3k	>\$3k
Ethereum	130.8M	85.1%	6.3%	8.5%
Tron (USDT)	218.2M	66.7%	15.2%	18.1%
BSC	1,072.9M	95.1%	3.0%	1.9%
Polygon	1,201.0M	98.8%	0.8%	0.4%
Arbitrum	104.3M	91.4%	4.7%	3.9%

82-99% of stablecoin transfers on every rail fall under \$3,000, and on four of the five rails, 85-99% fall under \$1,000.

Even Tron, the world's largest USDT rail and the one that carries the most large-value transfers, keeps two-thirds under \$1,000. That clustering is what makes the Travel Rule an infrastructure problem rather than a paperwork one. The lower your regime's floor, the more you must capture, and under the EU's zero floor that's all of it: originator and beneficiary data on hundreds of millions of transfers per chain per quarter.

The US rule looks lenient by comparison, but it doesn't spare you the work: you can't apply a \$3,000 threshold without valuing every transfer in real time to know which ones cross it. Those bands are raw, unfiltered counts (spam included), which is exactly right here: the obligation attaches to every transfer regardless of whether it turns out to be a real payment, dust, or an address-poisoning attempt. And because the pattern holds across architectures (EVM and Tron), it's a structural fact about stablecoin transfers, not an Ethereum quirk. But that surface is larger and noisier than a payment count implies. On Ethereum, a big share of these sub-\$1k transfers are address-poisoning spam (Figure 6), not payments, and a compliance system still has to screen every one. Whichever regime you operate under, the data surface is the same: every transfer, transaction-level, structured, and retrievable on demand. **Regulatory compliance is fundamentally an infrastructure problem.**

Figure 7 · Share of USDC+USDT transfers by size band, per chain — 90-day window







The "under \$1k" share is the activity the US \$3,000 threshold ignores but the EU zero-floor rule captures, and the 85-99% pattern holds across architectures (EVM, Tron, Solana). These are raw, count-based transfer shares and include non-economic dust; see Figure 6 for the spam share.

Methodology: Goldsky community dataset (public-blockchain data), USDC+USDT, 90-day window (Base 7d), selected rails shown. Optimism omitted; its USDC flow is DEX/bridge-dominated (see Page 12).

Ship stablecoin products faster with Goldsky.

Goldsky is the data layer for stablecoin products.

Everything in this report, settled balances, finality-aware ledgers, reorg-safe reconciliation, multi-chain decoding, audit-ready compliance data, is the layer your product **sits on top of**, not the product itself. We built it so your team doesn't have to.

-  150+ networks, one consistent data model
-  SOC 2 Type II, audit-ready data for attestation and Travel Rule reporting
-  99.99% uptime, stream-correct indexing that handles reorgs and finality for you
-  Live in under a week, not two quarters

Build the product. Let us handle the chain.

goldsky.com/fintech

Credits



Elliott Charbonneau CFA, *Business Operations Lead*



Håkon Åmdal, *Software Engineer Lead*