

The agent proposes. Ava decides.

The authority layer between AI agents and capital. A mandate set before the money moves, a signer the agent never touches, and a receipt anyone can check.

GETAVA.XYZ KAMAL NAYAN SINGH · AARAV RAJ

02 — WHAT CHANGED

Your agent can write the transaction. It still can’t sign it.

Claude Code, Cursor and Codex already research venues, encode calldata and simulate the result. MCP made tools universal. The last step is the one nobody made safe to hand over: the signature.

```
agent › supply 100 USDC to the safest lending market on Base

✓ compared lending markets on Base
✓ chose Morpho Blue
✓ encoded approve + supply calldata
✓ simulated: succeeds

✗ no signing key. stopping here.

Or worse: PRIVATE_KEY=0x… in a .env it can read.
```

A key held by a model is a key held by whatever it last read.

$175K

drained from an AI-linked Bankr wallet in May 2026 by an instruction hidden inside a post. The model read it. The wallet obeyed.

81%

hijack success after NIST's red team tailored attacks to a newer model, up from 11%

23.9%

of high-stakes tests where even the safest agent in ToolEmu failed severely

24%

of InjecAgent attacks that hijacked a ReAct GPT-4 agent

Prompt-level limits do not survive prompt injection. OpenAI's own guidance is to constrain the impact, because filters miss developed attacks. The limit has to live outside the model.

LEDGER ACADEMY AND METAMASK SECURITY REPORT, MAY 2026 ($175K–$204K REPORTED) · NIST AGENT-HIJACKING EVALUATION · TOOLEMU, ARXIV 2309.15817 · INJECAGENT, ARXIV 2403.02691

04 — HOW PEOPLE DO IT TODAY

Three ways to let an agent spend. All three fail.

UNBOUNDED

Hand it the key

Delegates everything, forever. One injected instruction reaches the whole treasury, and nothing records what it was allowed to do.

NO AUTONOMY

Approve every transaction

Safe and useless. A human clicking yes on each call is not an agent, and approval fatigue turns the click into a reflex.

NO MANDATE, NO PROOF

An agent wallet with per-tx caps

Checks whether this transaction looks safe, not whether this agent was allowed to make it. Caps split across parallel sessions, and the log is the wallet's word.

Builders converge on the same wish: set a scoped, cumulative limit once, let the agent run inside it, and get a record they can trust afterwards. Visa, Mastercard, Stripe and AWS describe the same primitives for card payments. On-chain, nothing we surveyed ships them as one layer.

05 — AVA

One contract between an untrusted agent and your money.

01 **Propose**
YOUR AGENT, OVER MCP
Untrusted. May ask for anything.

02 **Authorize**
YOUR MANDATE
Capital, venues, limits, expiry. EIP-712 signable.

03 **Reserve**
DETERMINISTIC POLICY
Budget and nonce claimed atomically.

04 **Sign**
TURNKEY ENCLAVE
The key never reaches the agent or a disk.

05 **Finalize**
CHAIN VERIFIER
The fill is read from the chain, not reported.

06 **Prove**
SIGNED RECEIPT
Hash-chained. Anyone can re-check it.

The model is outside the trusted computing base.

MCP is transport, not authority.

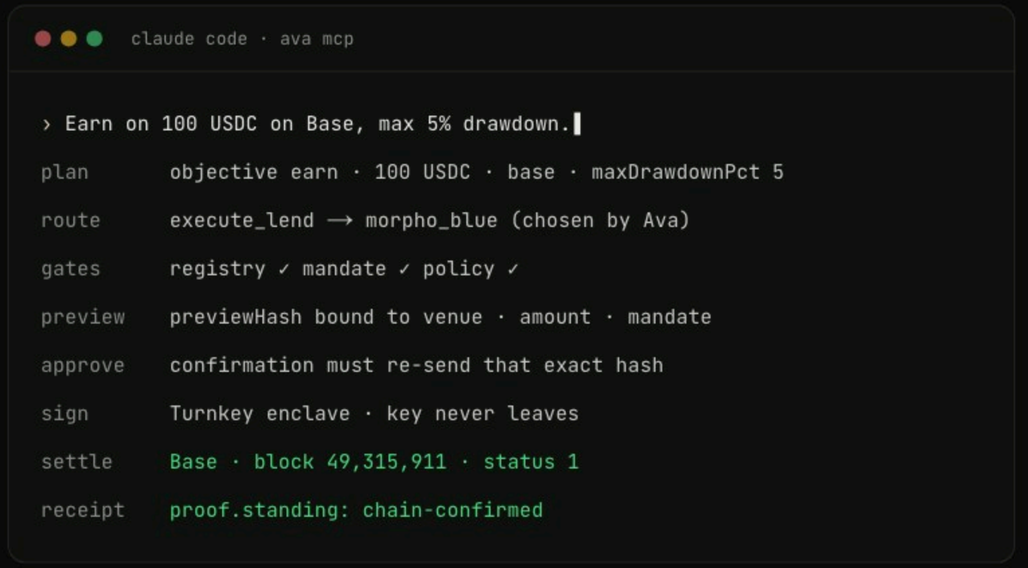
A tx hash is evidence of broadcast, not proof of effect.

06 — IT RUNS

One sentence in. A receipt out.

No venue named, no calldata written by hand, nothing Ava-shaped in the request. Ava picks the venue, checks it against a mandate written first, signs inside an enclave and returns a receipt a stranger can re-read from the chain.

COMPOSITE OF TWO RECORDED RUNS: THE SENTENCE → PLAN RUN, AND THE 0.1 USDC BASE SUPPLY THAT SETTLED.



07 — THE REFUSAL IS THE PRODUCT

Ava refuses more than it executes. On purpose.

A refusal costs nothing and names exactly what was wrong, with the offending values attached. A system that will execute anything is not one you would hand a key to.

157

distinct, named refusal codes
between an agent and a signature

```
POST /v1/live/approve
409 Conflict

{
  "code": "LIVE_PREVIEW_MISMATCH",
  "message": "The approved previewHash does not match
    the preview on file. Nothing was signed.
    Re-preview and approve the current numbers."
}
```

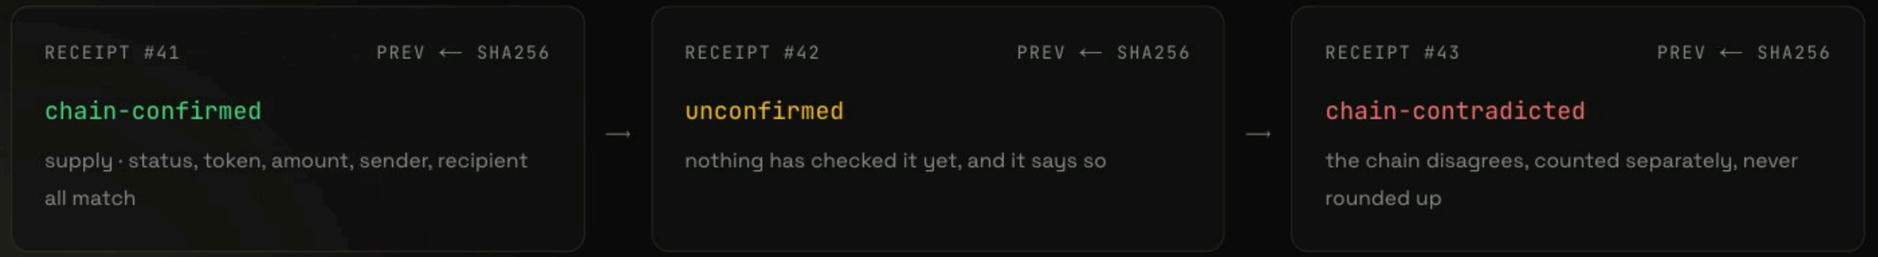
The quote moved after you approved it. Ava noticed before the signer did.

MANDATE_AMOUNT_EXCEEDS_CUMULATIVE
LIVE_PREVIEW_MISMATCH
MANDATE_WRONG_CHAIN
POLICY_DAILY_NOTIONAL_EXCEEDED
MANDATE_SIGNER_NOT_USER_CONTROLLE
MANDATE_EXPIRED
MANDATE_WRONG_AGENT
POLICY_APPROVAL_REQUIRED
MANDATE_TERMS_MISMATCH
MANDATE_PAUSED
MANDATE_SIGNATURE_INVALID
MANDATE

08 — THE RECEIPT

Every action ends in a receipt a stranger can check.

Not a log line written by the thing being checked. A canonical, signed record, chained to the one before it and compared against the chain itself, with no credential needed to verify it.



RFC 8785 canonical JSON · SHA-256 hash chain · Ed25519 signature · EIP-712 mandates · chain heads anchored on Base · public /verify

09 — TRACTION

Real money. Real chains. Checkable by anyone.

7

mainnet settlements signed by Ava's own enclave

4

mainnets: Base, Avalanche, BNB Chain, Monad

25

adapter packages behind one capability registry

236

test files guarding the gates and the signer path

BASE	Morpho gtUSDCp · supply USDC	0x97b78540…36bbba
BASE	Morpho gtUSDCp · withdraw USDC	0xce4109e4…6b0a7f
AVALANCHE	Aave v3 · supply USDC	0x12990c11…f4209f
BNB CHAIN	Aave v3 · supply USDC	0x7737cf5c…08b9de
MONAD	Aave v3 · supply USDC	0x50940c95…8715cf
MONAD	Aave v3 · supply USDC	0x1b23c35d…628af7
BNB CHAIN	SushiSwap via LI.FI · swap, no venue adapter	0x14850931…de0542

`npx @getava-xyz/connect` · session, MCP config and a Base address in one command

Listed on the official MCP registry as `xyz.getava/ava`

Works with Claude Code, Cursor, Codex and OpenClaw

EVERY ROW IS RE-READ FROM THAT CHAIN'S OWN NODE BEFORE THE SITE MAY BUILD · **VERIFY A RECEIPT YOURSELF** · SMALL AMOUNTS ON PURPOSE.
NO OUTSIDE USERS YET.

The rails shipped. The authority layer didn't.

97M

monthly MCP SDK downloads. Every serious agent now calls tools the same way.

138M

agent payments made over x402. Agents already pay for things on-chain.

$290B

in stablecoins: the dollars an agent would actually move.

112

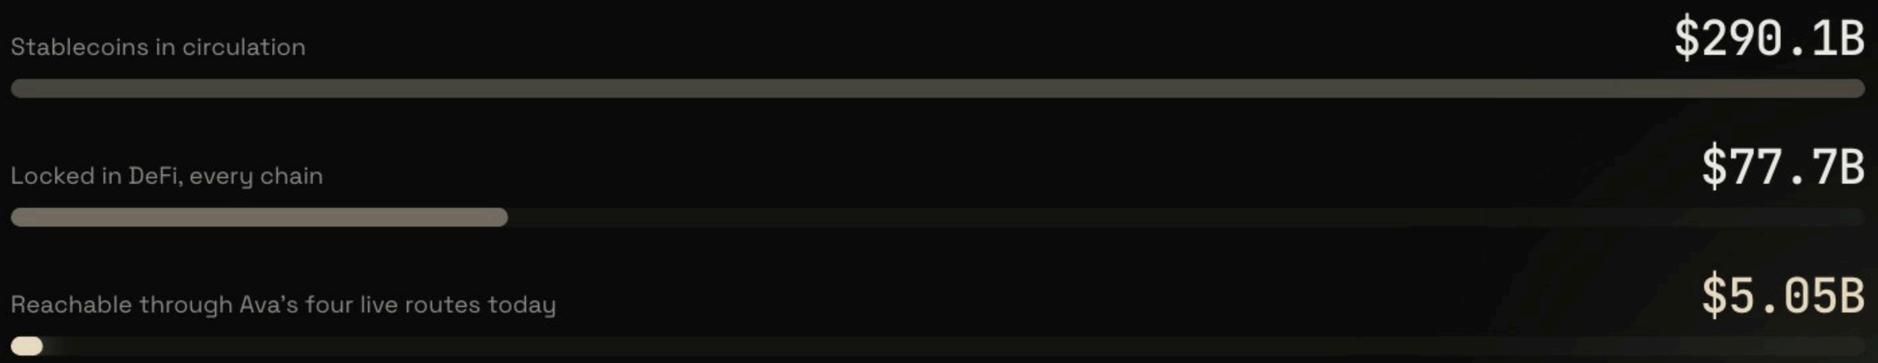
startups building agent payments, $115M raised in 2026 alone.

Agents can pay, and wallets can sign. What is missing is the layer that says this agent may do this, up to here, until then, and proves afterwards that it did nothing else.

ANTHROPIC, MAR 2026 · CHAINALYSIS, OCT 2026 · DEFILLAMA, 11 NOV 2026 · DEALROOM, 2026 YTD

Start where the money already is.

Bottom-up, from routes that already settle. Not a total addressable market: the share of this that comes to be run by agents is the market, and we think almost none of it will be without a receipt.



Next $0.99B: Aave v3 on Arbitrum and Base, wired and held behind the operator switch until each settles once on mainnet.

Beyond DeFi: Morgan Stanley puts US agentic e-commerce at $190B–$385B by 2030. Same problem, card rails instead of chains.

DEFILLAMA, 11 NOV 2026: AAVE V3 ON AVALANCHE, BNB AND MONAD PLUS MORPHO BLUE ON BASE · MORGAN STANLEY, DEC 2025

Wallets check the transaction. Ava checks the agent.

Agent wallets answer “is this transaction safe?” Marketplaces answer “who arbitrates a dispute?” Ava answers “was this agent allowed to do this, and can you prove it?” Wallets are partners: external wallets can already be linked, and Turnkey is the default signer.

OUR READ, FROM EACH PRODUCT’S PUBLIC DOCS, JUL–AUG 2026



We swapped through SushiSwap. We have no SushiSwap code.

A BNB Chain swap routed through SushiSwap via LI.FI settled on mainnet with no venue adapter, because Ava bounds the transaction instead of recognising the protocol. That generic route is wired into the API and held behind the operator switch until its pre-sign contract is complete.

25

adapter packages written the old way, one venue at a time

4

of them proven on mainnet after all that work

1

generic path, reaching any venue an aggregator can route to

Receipts compound into reputation. Every verified action adds to an agent's public track record: a history built on chain-checked proof, not ratings.

Mandates are a switching cost. Once an owner has set limits, history and receipts in one place, they move every new agent there, not off it.

We get paid when we can prove it.

A settlement fee on every receipt the chain confirms. A refusal costs nothing, and so does a receipt the chain has not confirmed or has contradicted. An executor paid per call has a reason to let one through. Ava does not.

FREE

0.10%

per confirmed settlement

$0 · 1 agent wallet

Mandates, policy gates, preview hash, public receipts, MCP and REST.

MOST OWNERS

PRO

0.05%

per confirmed settlement

$29 / mo · 10 agents

Isolated budgets per agent, standing authorizations, webhooks, CSV export.

FUND

0.02%

per confirmed settlement

$299 / mo · unlimited

Retained receipt archive, signed audit export, team seats, signed mandates on every agent.

ENTERPRISE

Custom

per confirmed settlement

Custom · dedicated

Own signer organisation, on-chain envelope module, RPC quorum, SLA.

On the preview. Shown before approval and bound into the previewHash.

Only on proof. Charged on chain-confirmed receipts. Everything else is $0.

Never skimmed. Invoiced from the receipt log, so Ava stays out of the funds path.

Verify is free. Anyone can check any receipt, with no credential and no charge.

PROPOSED PRICING · NO FEE CODE, BILLING OR PAYING USER YET · MIN $0.05, MAX $25 PER SETTLEMENT

One owner, five agents, $49 a month.

WORKED EXAMPLE · PRO

5 agents × $2,000 × 4 moves a month	$40,000 settled
× 0.05% settlement fee	$20
+ Pro plan	$29
Ava earns	$49 / month

The default 0.10% is about one week of yield on a lending market paying 5%.

$49K

1,000 Pro owners at $49 a month

$49.9K

100 Funds at $299 plus 0.02% on $1M a month each

$1.19M

annual recurring revenue from those 1,100 accounts. Illustrative, not a forecast.

Aggregator integrator fees. On swaps, LI.FI and Jupiter carry the fee in the quote instead of the invoice, never both.

Proof API. Metered bulk verification for auditors and allocators checking an agent before funding it.

Ecosystem funding. Chains pay to make their route live. Counted as funding, never as revenue.

We publish our own gaps. Here is what closes each one.

01 **Human confirmation binds content, not a person**	A user-key EIP-712 signature over the preview hash, so an agent cannot confirm on a human's behalf.
02 **Lend receipts report filled at inclusion**	A reconciler and per-chain finality depth before anything may read filled.
03 **Limits are enforced off-chain, by us**	An ERC-7579 validator module, so caps, allowlists and expiry hold even if Ava is compromised.
04 **Dependent cross-chain legs refuse today**	Two-leg evidence: both sides finalized, bridge message bound, timeout enforced.
05 **One RPC endpoint per read**	Agreement across independent endpoints, or a light client.

Built today: mandate and policy gates, preview-hash binding, an atomic single-flight claim that reserves budget and nonce, per-chain finality tables and the receipt chain. The model of all of it is written down in a 48-page technical paper.

The team that already put it on mainnet.

KN

Kamal Nayan Singh

CO-FOUNDER & CEO

Built Ava from its 2025 hackathon agent to the mainnet settlements in this deck.

AR

Aarav Raj

CO-FOUNDER & CTO

Co-author of Ava's technical paper and its architecture documentation.

1st

Cookie DeFAI Hackathon, Avalanche track, with Ava v1

212

votes and 6 attestations in ETHDenver 2025 quadratic voting

Demo

presented at the Sui AI Agent Hackathon demo day

Raising a pre-seed to put the first outside agents on real, mandated capital.

01

Close the gaps

Signed confirmation, finality reconciler, on-chain envelope bounds.

02

Get audited

An external review of the gate, the claim and the signer path.

03

Open the routes

Flip proven routes on, wire the generic path, then cross-chain.

04

First outside agents

Real owners, real mandates, receipts anyone can check.

kamal@getava.xyz

x.com/avatheaiagent t.me/avatheaiagent linkedin.com/company/avatheaiagent



The agent proposes. Ava decides.

Every number in this deck, and where it came from.

Settlements and chains	Re-read from each chain's node at build time, @/lib/settlements ↗
Bankr wallet drain, May 2026	Ledger Academy; MetaMask Security Report ↗
ToolEmu 23.9%	arXiv 2309.15817 ↗
138M x402 agent payments	Chainalysis, Oct 2026 ↗
112 startups, $115M	Dealroom, 2026 YTD ↗
Architecture	Components, workflows, data and trust ↗

Refusal codes, adapters, tests, limits	Counted from the repository at commit eff936c ↗
Agent hijacking, 11% → 81%	NIST technical blog ↗
InjecAgent 24%	arXiv 2403.02691 ↗
Stablecoins, DeFi TVL, reachable TVL	DefiLlama API, 11 Nov 2026 ↗
US agentic e-commerce by 2030	Morgan Stanley ↗
Live receipt verification	No credential needed ↗