FAB: Proof of Silicon

Paying verified GPU work in the tokenized equity of the chip's makers

FAB · fabrwa.xyz

Version 1.0 · September 2026

Abstract

FAB is a browser-native compute protocol on Robinhood Chain. A user runs a verified WebGPU workload for 15 to 60 minutes; FAB identifies the silicon lineage of the device (designer, foundry, memory supplier) and pays the user in tokenized shares of those companies from a public treasury contract. Rewards are denominated in verified compute units, converted to shares at settlement, and paid by voucher once per six-hour epoch. Payouts are bounded by treasury holdings through a rate governor. Mining is gated by holding the $FAB token. This paper describes the instruments, the verification protocol, the lineage map, the settlement and treasury contract, the token, measured performance on consumer hardware, the threat model, and the roadmap.

1 Introduction

Every graphics processor is the product of a short, concentrated supply chain: a designer (NVIDIA, Apple, AMD, Intel), a foundry (overwhelmingly TSMC), and a memory supplier (SK hynix, Samsung, Micron). The owner of the device paid all of them once, at purchase, and receives nothing from the device's later work.

Robinhood Chain, an Arbitrum Orbit layer-2 launched to public mainnet in July 2026, hosts over 2,000 tokenized US equities and ETFs as ordinary ERC-20 tokens (1 token = 1 share) that trade continuously on Uniswap. Early volume on the chain has been dominated by launchpad tokens rather than the stock tokens the chain was built for.

FAB connects the two facts. It turns a consumer GPU into a source of verified work and pays that work in the stock tokens of the companies that built the GPU. The design goals are: (i) no installation, the workload runs in a browser; (ii) verification strong enough that spoofed hardware earns no more than real hardware of the claimed class; (iii) rewards that can only be paid from assets the protocol already holds; and (iv) full public auditability of holdings, epochs, and claims.

2 System overview

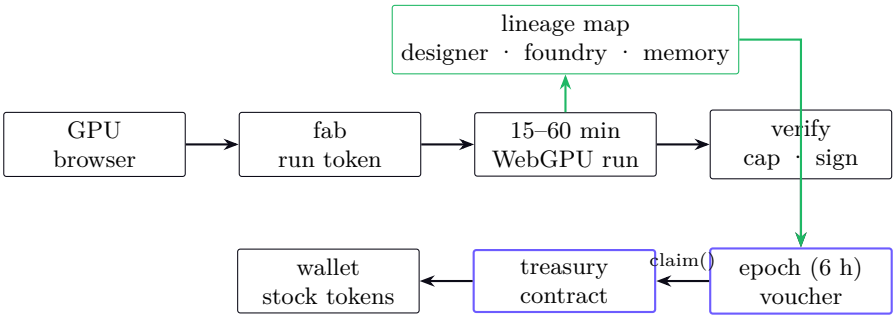

Figure 1: One run, end to end. The lineage map decides which tickers a wallet is owed; the epoch converts dollars to shares and signs a voucher; the treasury pays it.

A run has five stages. The browser requests a signed *run token* from the fab, which also records the wallet's gate status. The GPU executes one of five instruments for the chosen duration while the page displays live telemetry. On completion, the browser submits the result with the token; the server checks timing, caps the score at a per-vendor ceiling, computes the reward, and stores the run. Every six hours

the settler converts each wallet's accrued dollars into shares at the current price, signs one voucher per wallet, and anchors the epoch's report hash on chain. The wallet redeems the voucher against the treasury contract in a single transaction.

3 Instruments

Each instrument is a WebGPU compute program with an exact work accounting. All run in fp32, use the GPU's own timestamp queries for timing when the adapter exposes them, and adapt batch size so that each submission occupies the GPU for roughly 40–120 ms.

Program	Workload	Unit	Work counted
FAB-01 compute	dependent FMA chain, 2^{20} lanes × vec4	TFLOP/s	FLOP
FAB-02 memory	strided vec4 streaming read, 256 MB resident	GB/s	bytes moved
FAB-03 render	path tracer, 1024^2, 4 bounces, 7 spheres + ground	paths/s	camera paths
FAB-04 matrix	64×64 tiled SGEMM through workgroup memory, 1024^3	TFLOP/s	FLOP
FAB-05 inference	transformer forward pass, 12 layers, d=768, 12 heads, 128 tokens	TFLOP/s, tokens/s	FLOP

Table 1: Instruments. FAB-05 runs layernorm, QKV projection, causal attention with softmax, output projection, GELU MLP and residuals with synthetic weights; the work per token is 174 MFLOP.

Sustained, not peak. Runs are scored on throughput sustained over the whole run. This matters: the same Apple M2 Pro that peaks at 183 GB/s on FAB-02 averages 27 GB/s over 15 minutes as the package throttles, and a bandwidth-bound matrix kernel that starts at 0.5 TFLOP/s decays to 0.06 TFLOP/s within four seconds. A desktop card with a real cooler holds its number. The leaderboard is per GPU model and per instrument.

Instrument	Peak	Sustained	Note
FAB-01 compute	3.6 TF	3.4–3.5 TF	held for the full run
FAB-02 memory	183 GB/s	27 GB/s	package throttles within a minute
FAB-03 render	1.93 G paths/s	1.35–1.55 G paths/s	0.7 ms per 1024^2 frame
FAB-04 matrix	1.8 TF	0.3 TF	throttling, tiled kernel
FAB-05 inference	0.4 TF	0.4 TF	2,178 tokens/s, 58 ms per forward

Table 2: Measured on an Apple M2 Pro in Chrome (WebGPU), September 2026. Figures are the fab's own telemetry, timed by GPU timestamp queries.

4 Lineage map

On completion FAB resolves the adapter's vendor and architecture from the WebGPU adapter info and, where the browser exposes it, the WebGL renderer string. The exact model is shown only when the browser reveals it; otherwise the family ("RTX 40 series", "M-series") is used. The vendor selects a fixed reward map:

Vendor	Designer	Foundry	Memory	Other
NVIDIA	NVDA 48%	TSM 31%	SKHY 14%	SMH 7%
Apple	AAPL 46%	TSM 32%	SKHY 12%	SMH 10%
AMD	AMD 46%	TSM 32%	EWY 12%	SMH 10%
Intel	INTC 58%	TSM 22%	EWY 10%	SMH 10%
Qualcomm	QCOM 50%	TSM 30%	EWY 10%	SMH 10%
unknown	TSM 45%	ASML 25%	MU 20%	SMH 10%

Table 3: Reward map by vendor. SKHY is the SK hynix ADS token; EWY (iShares MSCI South Korea) stands in for Samsung, which has no stock token on the chain; SMH (VanEck Semiconductor ETF) covers packaging and tools. Weights are editorial and versioned with the protocol.

5 Verification and threat model

Run token. `POST /api/run/start` issues $\tau = \text{nonce} \,\|\, t_0 \,\|\, H(\text{ip}) \,\|\, \text{instrument} \,\|\, \text{duration} \,\|\, \text{mode} \,\|\, \text{HMAC}$. The nonce is stored server-side with a TTL and consumed once. A result is accepted only if it returns after $t_0 + \text{duration} - 5\,\text{s}$ and before $t_0 + \text{duration} + 30\,\text{min}$, from the same hashed IP.

Ceilings. Reported throughput s and total work W are clipped: $s' = \min(s, c_v)$, $W' = \min(W, c_v \cdot T)$ with c_v the vendor ceiling (Table 4) and T the run length. A spoofed result cannot exceed what the claimed silicon class can physically do.

Timing. GPU-side timestamp queries measure kernel time; wall-clock time is used only where the adapter lacks the feature. Main-thread jank therefore cannot inflate or deflate a score.

Rate limits and gate. Four runs per hour per address. Mining requires a *mining pass*: a message signed by the wallet, after which the server reads the wallet's $FAB balance on chain and requires at least $30 at the current price; passes last 24 hours. One free trial run per address.

Vendor	TFLOP/s	GB/s	M paths/s
NVIDIA	130	1,900	14,000
AMD	90	1,100	9,000
Apple	30	900	3,000
Intel	35	600	3,500
Qualcomm / Arm	8	150	800
unknown	40	600	4,000

Table 4: Per-vendor ceilings applied to every result.

What this does not prevent. A modified client can report a fabricated result up to the ceiling for its claimed vendor without doing the work. The economic exposure of that attack is bounded per address per hour by the ceiling, the rate limit, the 15-minute floor, and globally by the governor (Section 7). Closing the gap further requires either a proof-of-work-style challenge embedded in the kernels (a planned FAB-06 with server-issued seeds and spot-checked outputs) or trusted hardware attestation, which browsers do not expose. FAB states this openly rather than claiming cryptographic proof of computation.

6 Rewards and settlement

Compute units. One verified compute unit (VCU) is 40 TFLOP of verified FLOP, 800 GB of verified memory traffic, or 1.8×10^{10} verified camera paths. The constants are calibrated so that the same device earns similar units per minute across instruments.

Rate. A verified run earns

$$R = \big(0.10 + 0.0009 \cdot \text{VCU}\big) \cdot m \quad [\text{USD}],$$

where $m \in [0.1, 1]$ is the governor multiplier. On a 15-minute run this is about $1.05 for an RTX 4090 and $0.17 for an M2 Pro at $m = 1$.

Epochs. Epoch $e = \lfloor t/21600 \rfloor$ (UTC, six hours). During epoch e, rewards attached to wallet w accrue per ticker: $D_{e,w,k}$ dollars. At settlement, shares $a_k = D_{e,w,k}/p_k$ where p_k is the ticker's USD price at settlement (DexScreener for the pool with the most liquidity, else GeckoTerminal, else Yahoo Finance for stock tokens). The settler signs

$$\sigma = \text{Sign}_{\text{settler}}\big(\text{keccak256}(\text{chainId}, \text{treasury}, e, w, [\text{token}_k], [a_k])\big)$$

with the Ethereum signed-message prefix, and calls `closeEpoch(e, H(report))` on the treasury. The wallet calls `claim(e, w, tokens, amounts, σ)`; the contract verifies the signature against the settler,

marks (e, w) claimed, and transfers each token from its own balance. A claim reverts if the treasury lacks any of the tokens in the voucher.

7 Treasury

The treasury is `FabTreasury` at `0x223245E7C8A0Ea09C8a21d0EAfD45828b10Ca68f` on Robinhood Chain mainnet (chain id 4663), source verified. It holds the stock tokens listed in Table 5 and exposes `deposit(token, amount, memo)` for anyone, `closeEpoch` for the settler, `claim` for miners, `holdings` and `totalPaid` views, and `sweep` for the owner. The site reads holdings from the contract on every page load.

Governor. At every settlement the fab values the treasury V at current prices and sums what the trailing 24 hours of runs would have earned at full rate, F. The target multiplier is $m^\star = \text{clamp}(\tau V/F,\ 0.1,\ 1)$ with $\tau = 3\%$ per day, and the applied multiplier moves toward $m^\star$ by at most a factor of two per settlement. Deposits therefore raise the rate for every miner and new miners spread it thinner; at $\tau = 3\%$ a vault lasts about a month with no new deposits. The multiplier is applied to every run completed until the next settlement and is displayed on the treasury page. Rewards can slow down; they cannot be created.

Funding. Three sources, in expected order of importance: sponsor deposits (any party may deposit stock tokens with a public memo), protocol fees from jobs dispatched through the fab (roadmap), and founder capital. The launch-week treasury was funded by swapping ETH into stock tokens on Uniswap via aggregator routes and depositing them directly.

Ticker	Token	Address (Robinhood Chain)
NVDA	NVIDIA	`0xd0601CE157Db5bdC3162BbaC2a2C8aF5320D9EEC`
TSM	TSMC	`0x58FfE4a942d3885bAa22D7520691F611EF09e7AA`
SKHY	SK hynix ADS	`0x84CAb63bc87912E71ad199ff14A0bA45de68FeF8`
AAPL	Apple	`0xaF3D76f1834A1d425780943C99Ea8A608f8a93f9`
AMD	AMD	`0x86923f96303D656E4aa86D9d42D1e57ad2023fdC`
INTC	Intel	`0xc72b96e0E48ecd4DC75E1e45396e26300BC39681`
QCOM	Qualcomm	`0x0f17206447090e464C277571124dD2688E48AEA9`
ASML	ASML	`0x47F93d52cBeC7C6D2CfC080e154002370a60dAEA`
MU	Micron	`0xfF080c8ce2E5feadaCa0Da81314Ae59D232d4afD`
SMH	VanEck Semiconductor ETF	`0x072f979c2CAc8e1391B0162a87Fee094bF8744a0`
EWY	iShares MSCI South Korea	`0x7f0aBeF0C07280F82c6a08ead09dEd6BAE2C13Fc`

Table 5: Payout tokens. Symbols verified over RPC against the on-chain registry.

8 The $FAB token

$FAB (`0xc7db9a7034d061feca50f0a25636bcc83ca717d9`) has a fixed supply of 10^9, launched in full on the Pons v2 bonding curve with 0% allocated to the team. On graduation the curve's liquidity is placed in a Uniswap v4 pool that cannot be withdrawn by anyone. $FAB is a key, not a reward: holding at least $30 of it in the mining wallet is required to mine beyond the free trial. Rewards are never paid in $FAB and $FAB is never minted by the protocol. The price used by the gate is read from the pool with the most liquidity above a $5,000 floor, falling back to the indexer of the bonding curve.

9 Roadmap

- **Shipped, Sept 5–7 2026.** Five instruments; run verification; live feed and leaderboards; wallet claims; six-hour settlement with on-chain epoch anchoring; FabTreasury on mainnet; rate governor; holder gate with free trial; 15/30/60-minute runs; launch films and site.
- **Next.** A job queue: inference and render jobs submitted to the fab, dispatched to $FAB holders, verified with the same accounting, paid into the treasury by the buyer. Real model weights and prompts on

FAB-05. Treasury ownership moved to a multisig. Spot-checked kernels (FAB-06) with server-issued seeds.

- **Later.** Sponsor SDK for chip makers and issuers to deposit against their own ticker; mobile GPUs; per-model calibration of ceilings from observed distributions; a public dataset of sustained browser GPU performance.

10 Risks and disclaimers

FAB is early-access software operated by an independent team not affiliated with Robinhood, NVIDIA, TSMC, SK hynix, Apple, AMD, Intel, Qualcomm, Arm, ASML, Micron or Samsung. Nothing in this document is an offer, solicitation, or recommendation regarding any security or token. Tokenized equities are subject to their issuer's terms and to the laws of the holder's jurisdiction; availability varies. Rewards depend entirely on treasury holdings and can be reduced to zero by the governor. Reward maps, ceilings, and rates are protocol parameters and will change. Running a GPU at full load for extended periods stresses hardware; users are responsible for their devices.

Site: https://fabrwa.xyz · Treasury: https://robinhoodchain.blockscout.com/address/0x223245E7C8A0Ea09C8a21d0EAfD458
· Token: https://robinhoodchain.blockscout.com/token/0xc7db9a7034d061feca50f0a25636bcc83ca717d9