Research note · November 2026

Compute as consensus

AI × Crypto × Chips in 2026, starting from Pearl

Pearl · Useful work · DePIN and agents · ASIC economics · Process nodes · Staged budgets

Tan Shuai
tanshuai.com · i@tanshuai.com



THE VERDICT

The theorem holds; the incentive does not

01 Pearl works as engineering, not yet as useful work

About 90 EH/s of integer matrix multiplication secures the chain, but only about 1% of blocks carry real AI work.

02 Paid demand is in inference, data and privacy

Bittensor about $70 million a year of verifiable revenue, Grass $17 million in H1 2026, TEEs in production; agent payments about $5,000–11,000 a month.

03 An ASIC pays only with a large pot, stable rules and a second market

Its profit ceiling is R × (1 − 1/k). Today no network offers all three analysis.

≈90 EH/s

Pearl on November 13, 2026: about 122,000 H100s of integer throughput

≈1%

of Pearl blocks carry real AI work (Blockworks)

4–11×

Bittensor's emissions against its verifiable revenue

AI × crypto in six sectors: where paid demand is real

Useful-work chains 01

REAL BUT SMALL

≈$70M a year

Bittensor's verifiable external revenue

Signal: emissions 4–11× revenue; Pearl ≈1% AI blocks

The mechanisms work; usefulness needs payers

Decentralized compute 02

REAL BUT SMALL

$2.35M a year

Render's on-chain fees, the largest verifiable

Signal: market value ≈360× fees

Mostly resale of commodity inference

Training and inference 03

PROFIT IN EQUITY

$100M+ a year

Prime Intellect, an equity company on centralized GPUs

Signal: Gensyn's training mainnet not live

Progress is real; profit goes to equity

Verifiable AI 04

EARLY

398

paying users of Phala's TEE cloud, end of 2025

Signal: no disclosed paid zkML

TEEs in production; zkML early

Agent payments 05

BEFORE DEMAND

$5–11K a month

x402 payments actually initiated by agents

Signal: ≈200M transactions, mostly signalling

The rails came first; demand has not

Data networks 06

REAL BUYERS

$17M

Grass data revenue, first half of 2026 (self-reported)

Signal: market value ≈10× annualized revenue

One of the few sectors with real buyers

Subsidy is fuel for a cold start, not a business model. Crypto does distribution and settlement; equity takes the profits, at least in 2026.

One multiply-accumulate is one hash

- **From theorem to mainnet**

  Komargodski and Weinstein (April 2025): a proof of useful work for any matrix product at 1 + o(1) overhead. Pearl's mainnet started on April 27, 2026.

- **How it works**

  Miners multiply noisy matrices, hash each tile's accumulator into a lottery ticket, and prove the winning tile with a Plonky2 proof under 60 KB.

- **The unit**

  1 H is one credited INT7 × INT7 → INT32 multiply-accumulate: a blockchain's security budget priced in tensor throughput.

PEARL ON NOVEMBER 13, 2026

89.77 EH/s

network hashrate (independent full node)

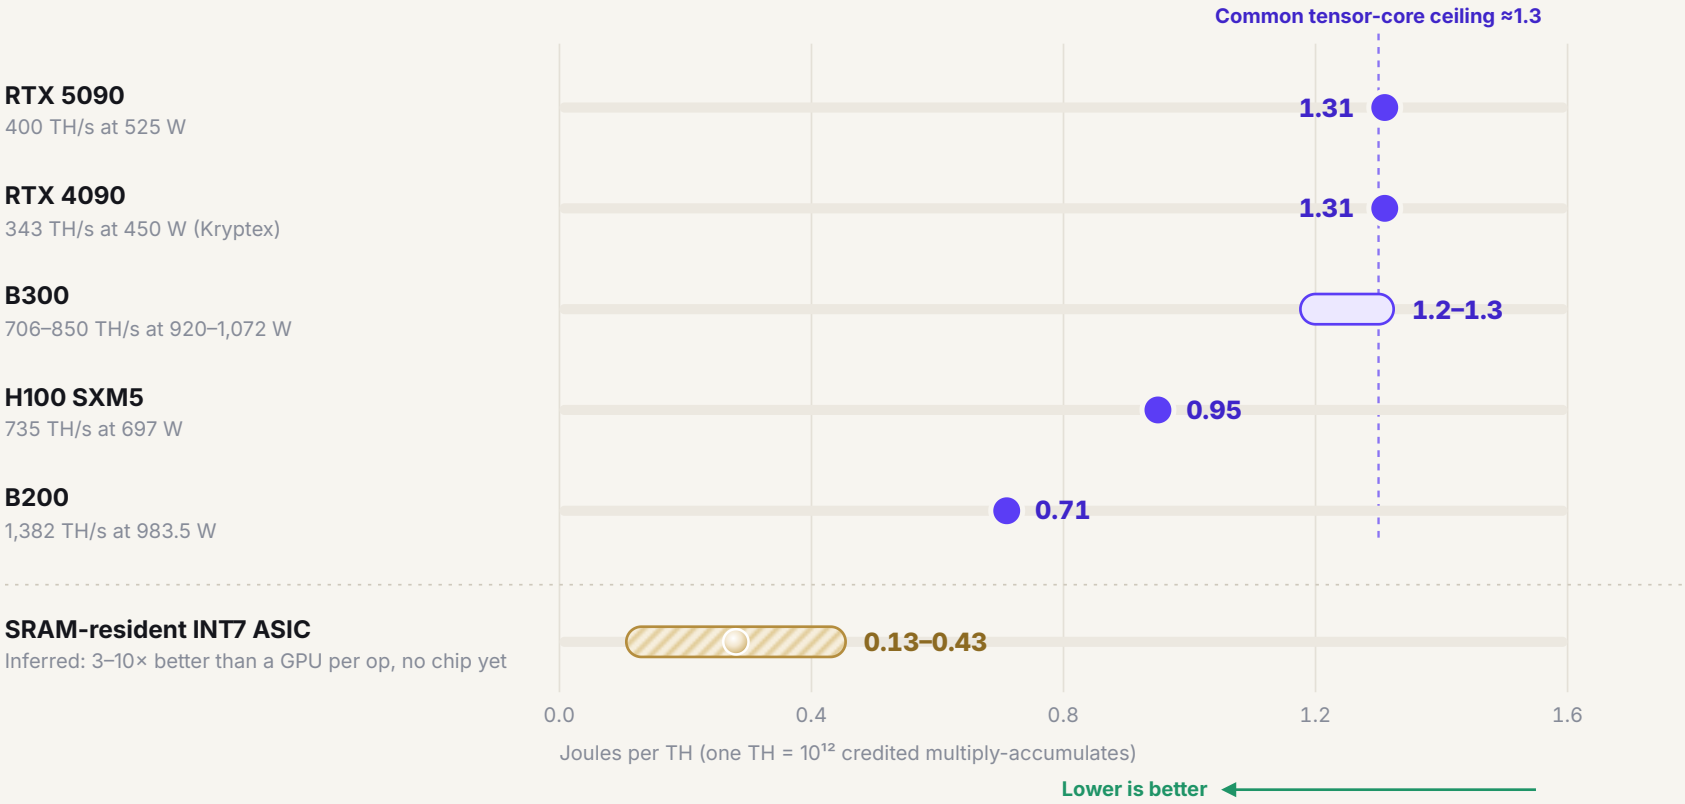
≈$500M

market value; fully diluted about $2.8 billion

+78%

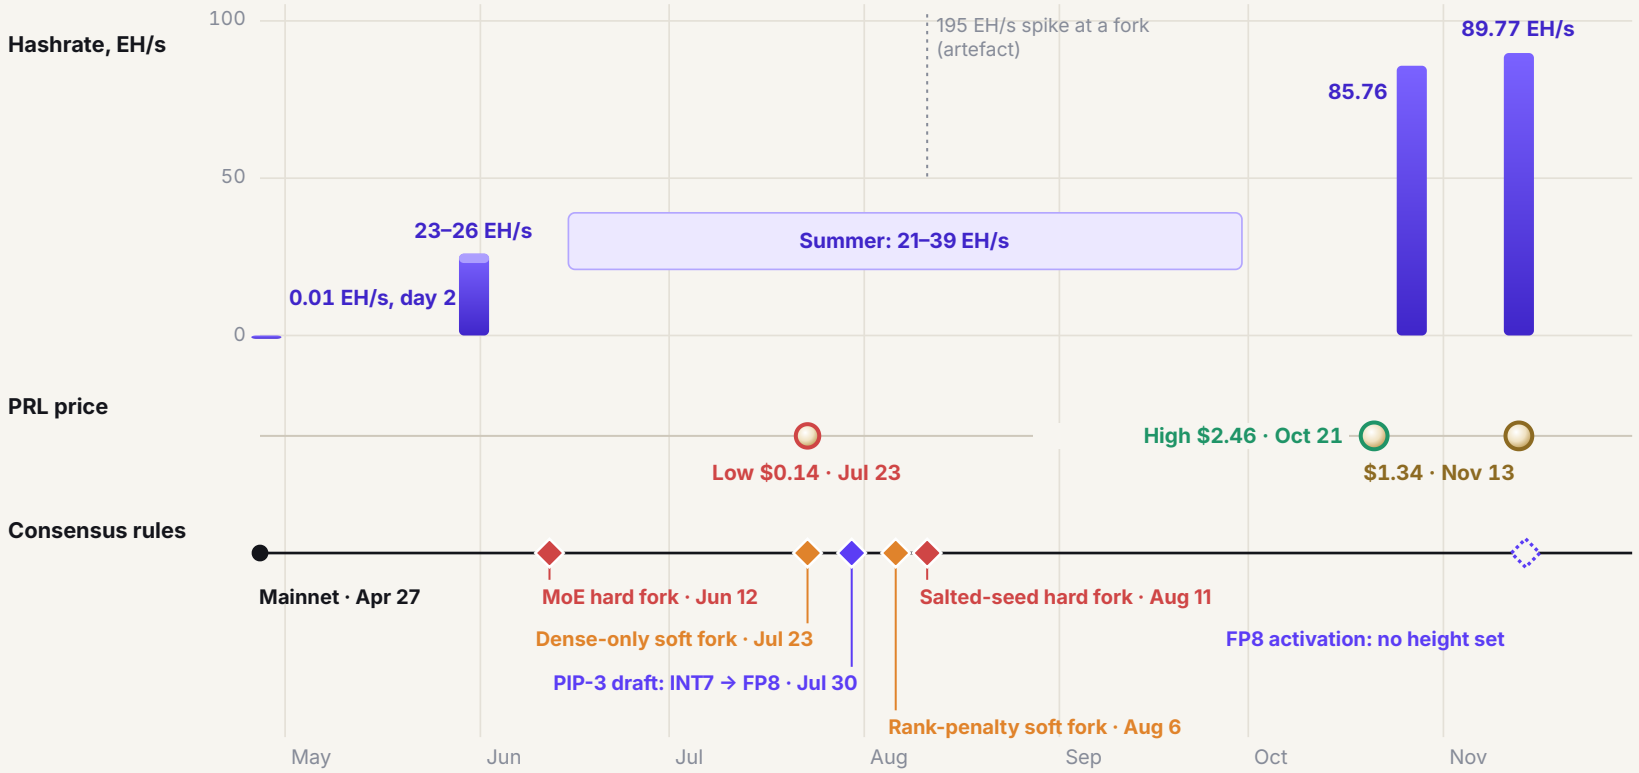
new supply over the next 12 months

On Pearl's integer work, GPUs share an energy ceiling



> Gaming cards and the B300 meet at about 1.3 J/TH; even the B200 reaches only 0.71. A chip that only multiplies 7-bit integers from SRAM could plausibly do 3–10 times better analysis.

Hashrate follows price; the rules keep changing



About 120 million coins, 32% of today's supply, were mined in the first six days. An RTX 5090 now earns about $5–6 a day; renting cards only to mine loses money.

The theorem holds; the incentive does not

WHAT USEFUL WORK COSTS

Inference side: small overhead

Pearl reports 5.08% more time for Llama 70B on four H200s; an independent test measured 9.2% less throughput on one Hopper GPU.

WHAT USEFUL WORK EARNS

Mining side: too little hashrate

A card serving inference mines at about 27% of a dedicated miner on prefill, 5.4% on chat, about 1% when decoding dominates.

So rational miners multiply random matrices. Usefulness is an off-chain commercial problem: someone has to pay for inference.

Four forks in two months, and FP8 next

June 12
Mixture-of-experts hard fork (block 71,935)

July 23
Dense-only soft fork (block 91,630)

August 6
Rank-penalty soft fork: some setups lose 50–87.5% of effective hashrate

August 11
Salted-seed hard fork: old miners produce only invalid shares

Not set
PIP-3: INT7 to FP8, only Hopper and Blackwell admitted

2 of 3

useful for AI, verifiable bit for bit, hardware-neutral: a number format gets at most two

13 bits

kept by Hopper's FP8 accumulation, with truncation: not reproducible across generations

≈82%

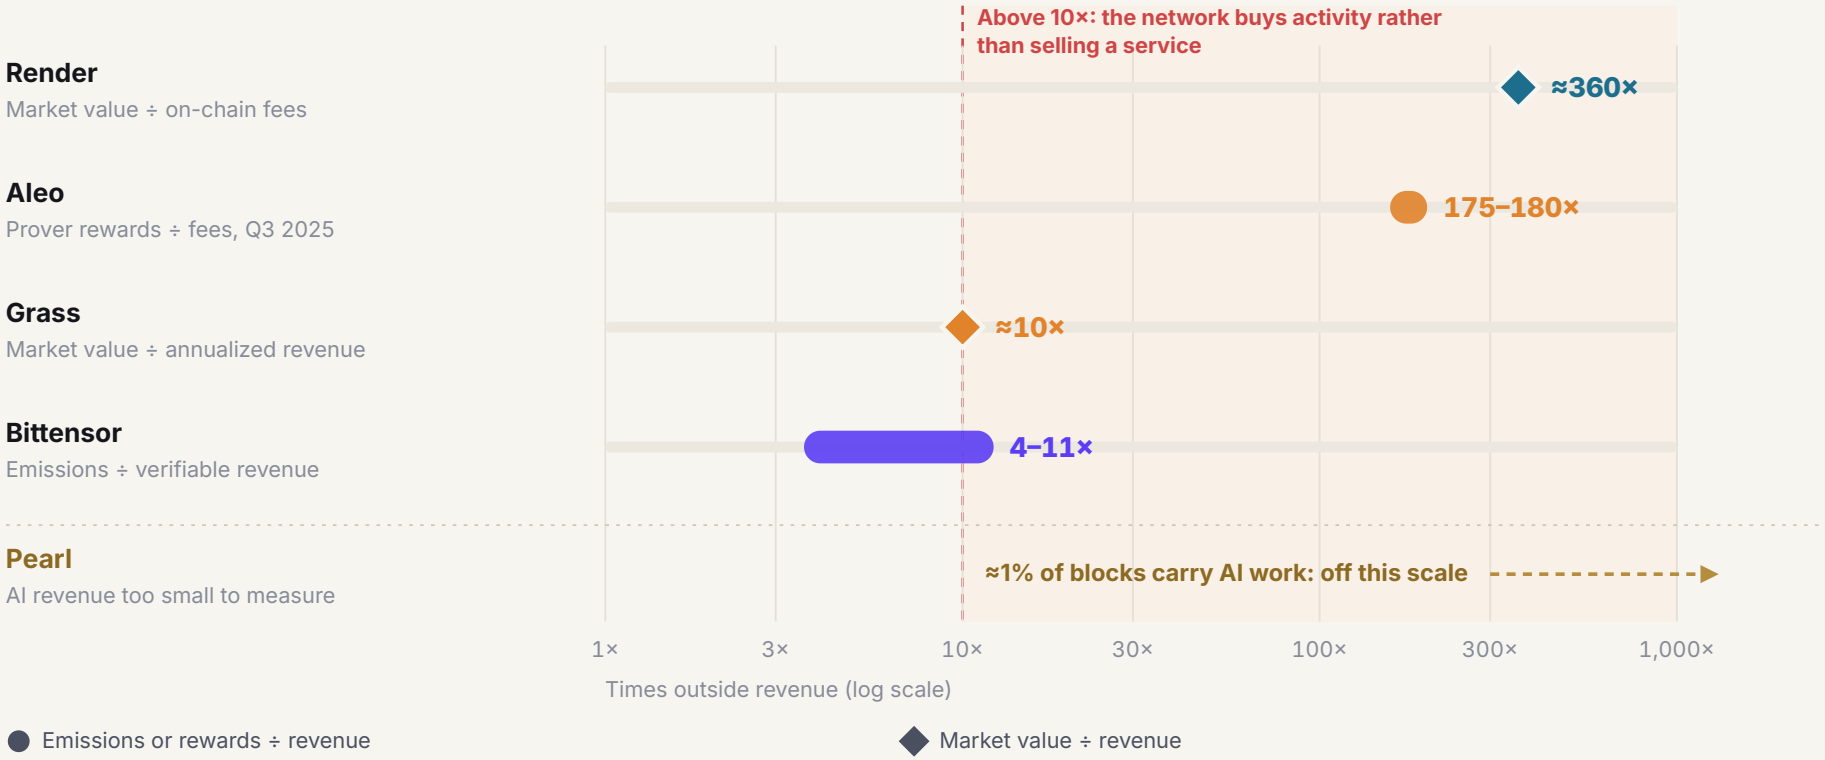
of hashrate on consumer cards that PIP-3 would not admit (low-confidence model)

Whoever controls the work function controls the fate of hardware investors analysis **.**

Neither matrices nor zero knowledge stop ASICs

Network	What happened	What it teaches
Kaspa	A 64 × 64 matrix product between hashes; ASICs after 16 months, about 1,000× GPU efficiency; the best now earns $0.86 a day	A matrix is neither useful nor ASIC-proof
Aleo	A synthetic SNARK puzzle; dedicated machines after five months; prover rewards 175–180× fees	Synthetic work runs on subsidy
Qubic	Claims AI training; its compute mined Monero (51% claimed) and now Dogecoin	General-purpose PoW can be captured from outside
Bittensor	Stake-weighted scoring; about $70 million a year of verifiable revenue against more than $300 million of emissions	The only one with outside revenue at scale
Monero, Sia, Ethereum	Forks bricked ASICs; after the Merge under 20% of miners found new work	Hardware with a second market survives

The subsidy multiple is the most honest health check



> Above about 10 times outside revenue, a network is buying activity rather than selling a service. Pearl's AI revenue is too small to place on the scale.

Real but small demand, and value moving to equity

- **Compute markets**

  Render's on-chain fees about $2.35 million a year; Akash's H100 at $2.73 an hour beats AWS and Google but not Vast.ai or RunPod.

- **Training**

  8–40 billion parameters trained over the internet; INTELLECT-3 went back to 512 H200s on InfiniBand.

- **Verifiable AI**

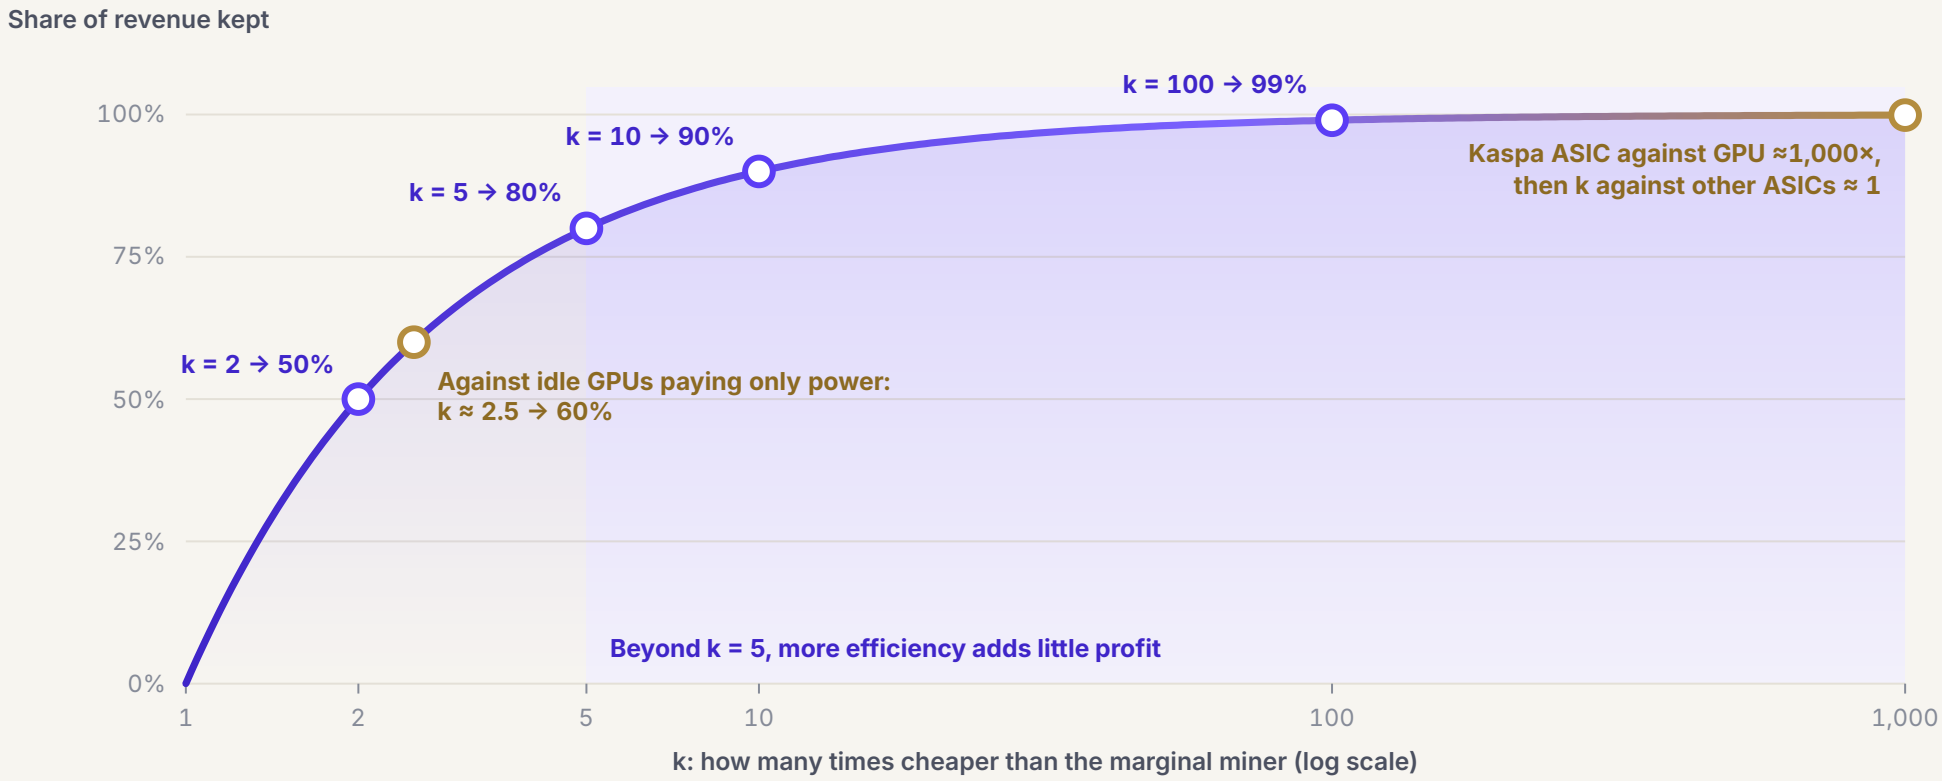
  TEEs add under 7% and run in production; TEE.fail broke attestation with a $1,000 interposer; zkML proves GPT-2 at 174 tokens a minute.

- **Agents and data**

  x402 agent-initiated flows about $5,000–11,000 a month; Grass $17 million of data revenue in H1 2026.

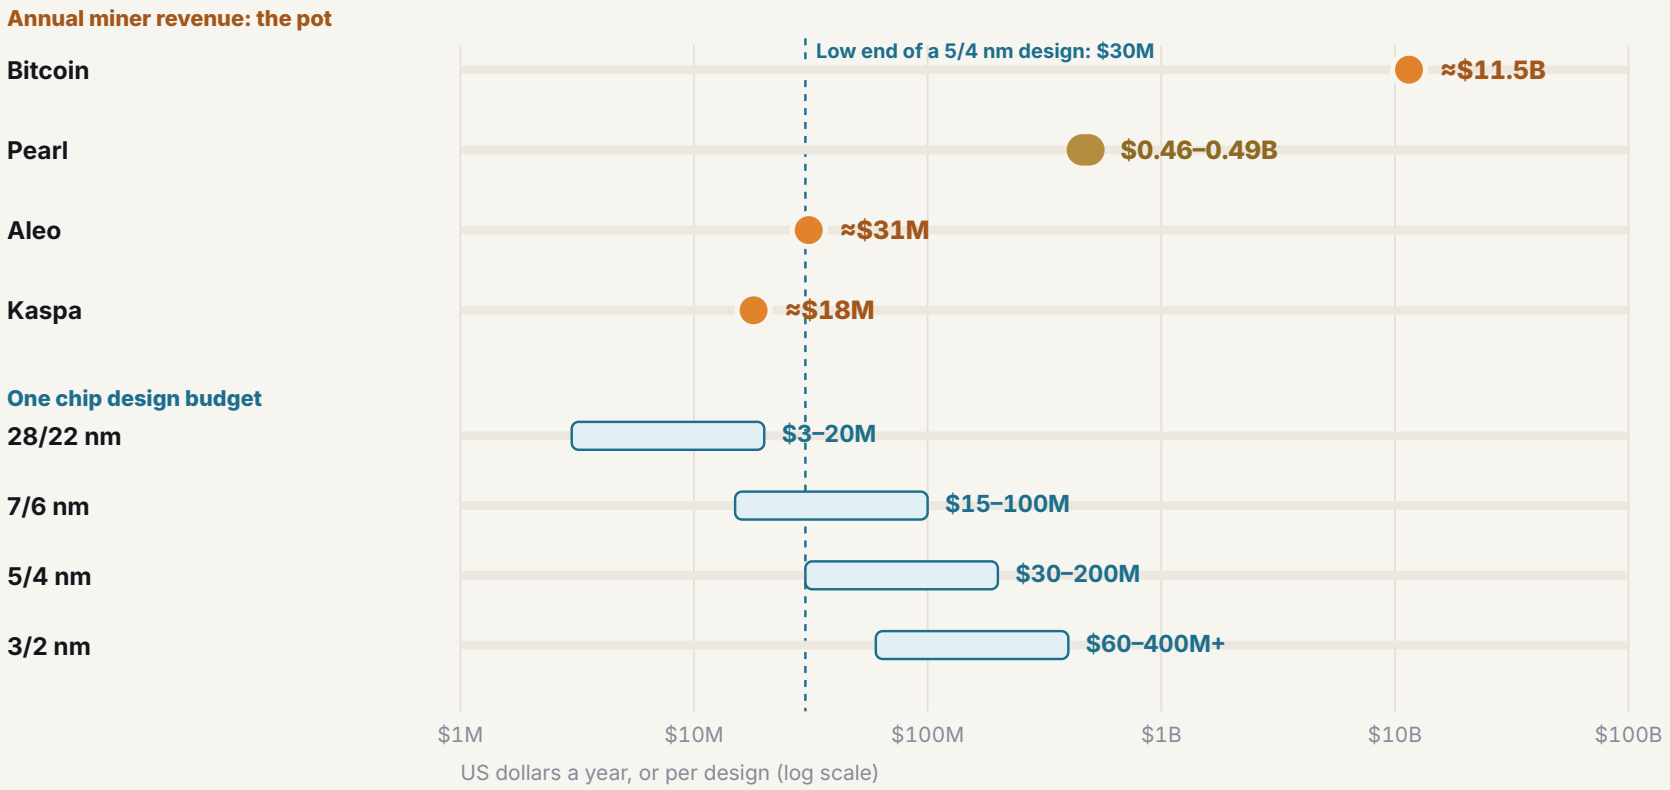
> **Four equity companies on centralized infrastructure are worth about $8.5 billion, about as much as all protocol- and application-layer AI tokens.**

Why efficiency stops paying after k = 5



Profit ceiling ≈ R × (1 − 1/k). Once ASICs become the marginal miners, k resets against other ASICs; against idle GPUs that pay only for power, k is about 2.5 estimate.

Most pots are smaller than one chip design budget



> Payback hashrate h* = NRE ÷ [R × (1 − 1/k) × T × p]. For Pearl at $20 million of NRE and k = 5: about 4.9 EH/s for a year if the rules hold, 9.8 EH/s if they hold with even odds **estimate**.

Five chip types, five answers

Pearl-style PoUW mining chip

NOT NOW

k: no HBM or general units: 3–10× per op (inferred)

Second market: AI inference, if general INT8/FP8 GEMM

Bet on an inference-first chip that can mine, not a dedicated INT7 part

Bitcoin SHA-256 mining chip

STAY OUT

k: ≤1 against the incumbents

Second market: almost none

Only with own power, a niche form, or as IP

AI training chip

IP, NOT CHIPS

k: 1.1–1.8× at the same precision

Second market: —

Packaging and HBM decide; build IP, chiplets, services

AI inference chip

BUILD

k: give up generality: 10–20× claimed

Second market: it is the main market

Sign the customer, then choose the lock-in

ZK proving accelerator

WAIT AND HEDGE

k: 10–100× GPUs claimed

Second market: limited outside ZK

Programmable MSM and NTT engines; FPGA first

FHE accelerator

POSITION EARLY

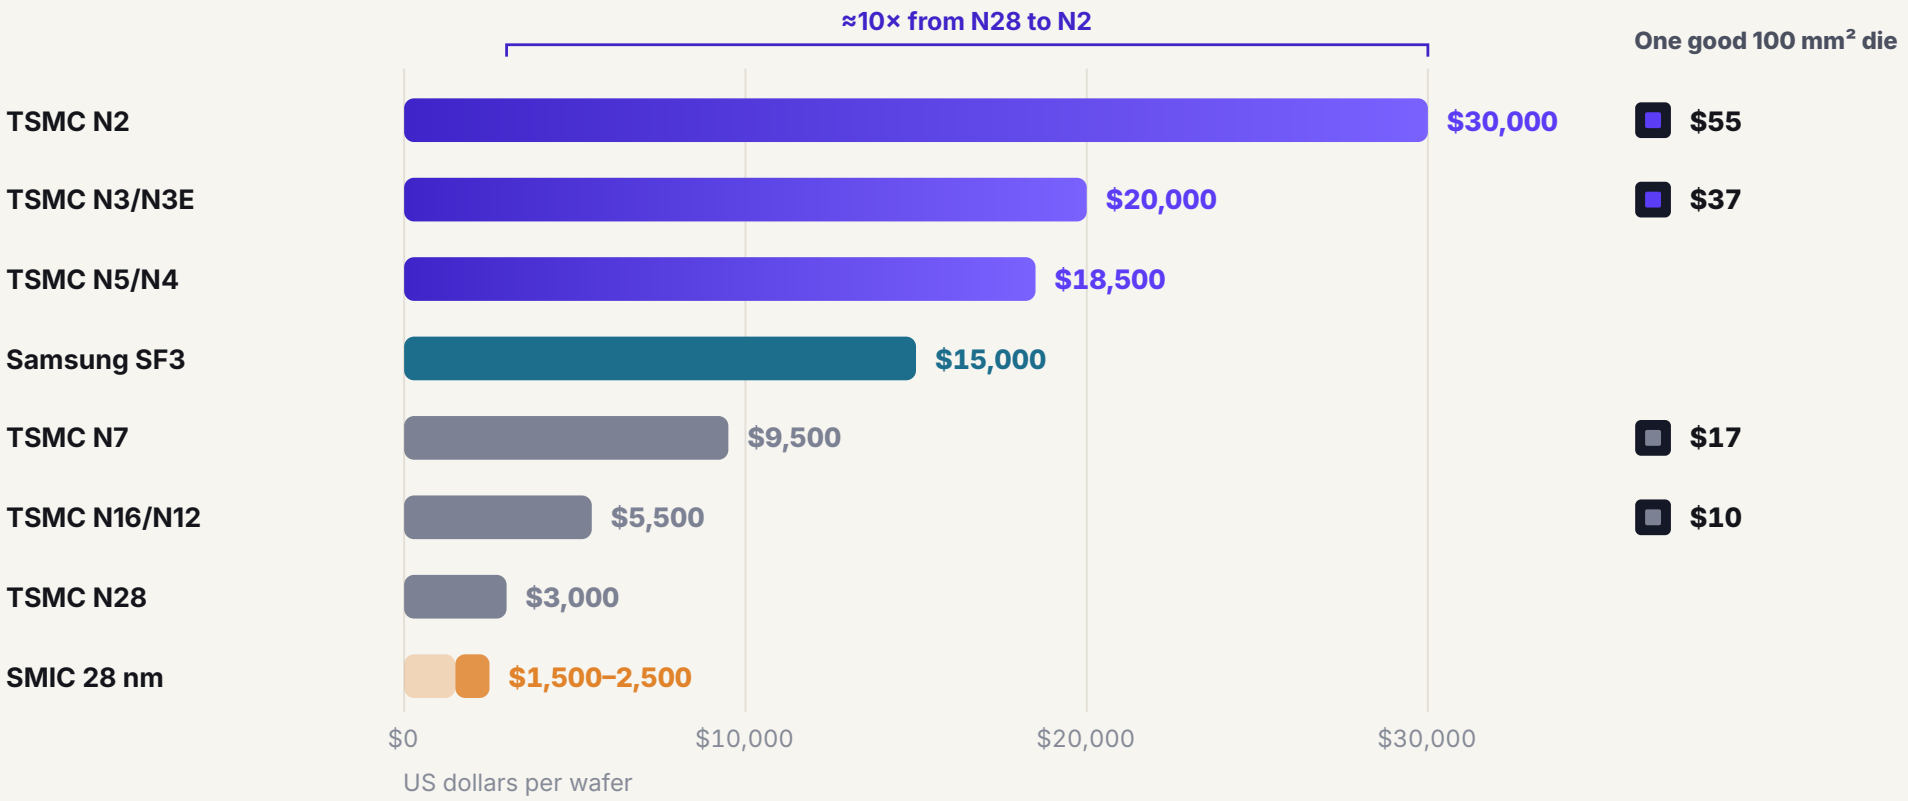
k: 1,074–5,547× a Xeon (Heracles)

Second market: encrypted computing, private AI

A 2027–2029 window; validate with light assets

> No network today offers a large pot, stable rules and k above 1 together: bitcoin's k belongs to the incumbents, Pearl's rules are the least stable, Kaspa's and Aleo's pots are too small.

A leading-edge wafer costs ten times a mature one



> Let the largest lifetime cost choose: energy for around-the-clock compute chips, packaging and HBM for training chips (63% of an H100's cost), the SRAM-or-HBM choice for inference chips.

PROCESS NODES

A node for each chip and each team

Chip	Small team	Scaled company	Mainland China path
PoUW mining chip	28 or 16/12 nm MPW; first volume N7/N6	N3/N3E or Samsung SF3	SMIC N+1/N+2 works, about 40–50% more per die
SHA-256 mining chip	Do not chase the frontier	3 nm class	Hard to match leading-node efficiency
AI training chip	IP, chiplets, design services	N3P then N2/A16 with CoWoS and HBM	SMIC N+2/N+3 with domestic HBM
AI inference chip	SRAM-centred at N6/N7 or N4/N5/SF4	N3/N2 with HBM, or SRAM chiplets	SMIC N+2 with LPDDR or GDDR
ZK and FHE accelerators	FPGA or 16/12 nm MPW first	N5/N3; FHE with HBM	ZK workable; FHE limited by HBM

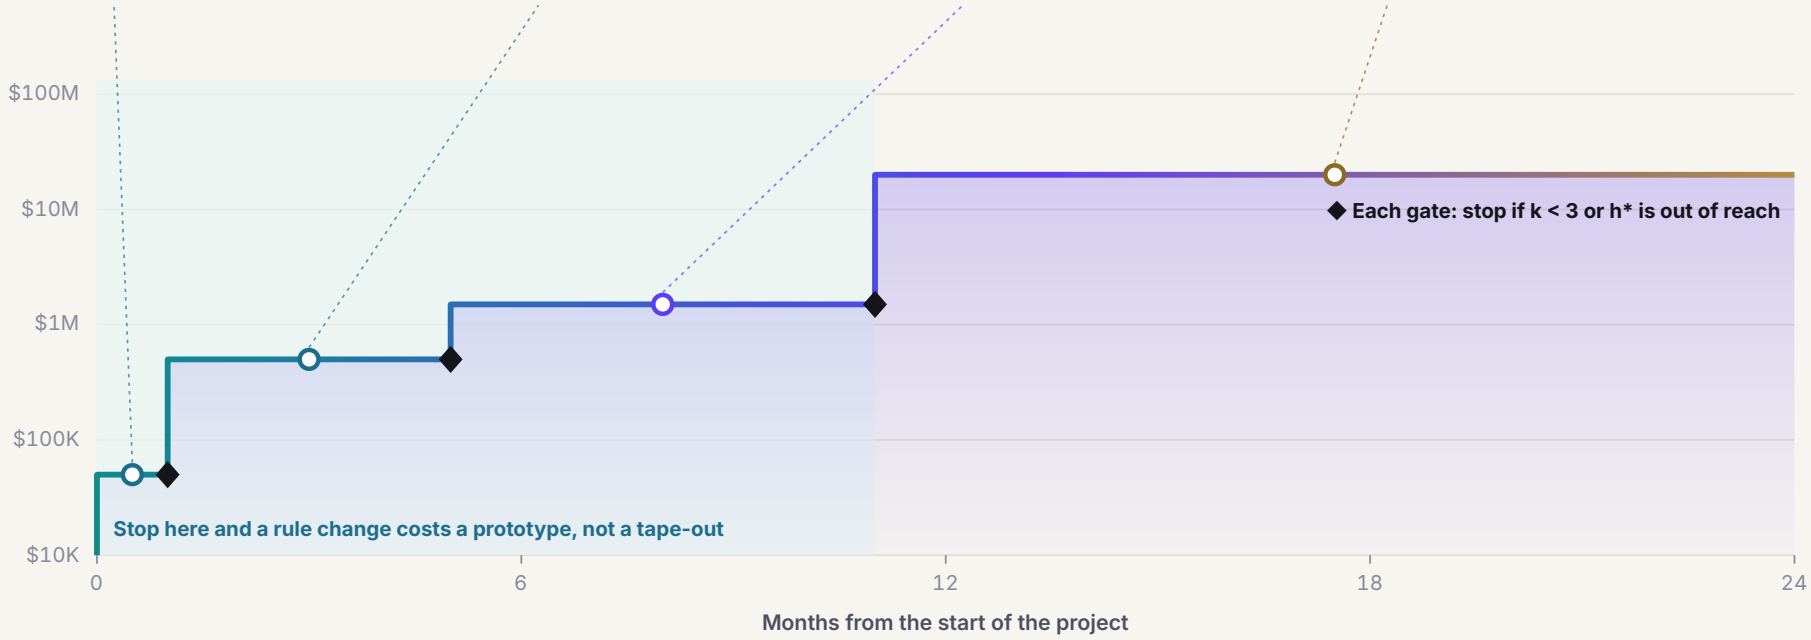
Buy evidence before masks

Gate 0 · Model
≈$50K · weeks
J/op and $/op from the arithmetic, SRAM and I/O

Gate 1 · FPGA
≈$0.5M · months
The real workload end to end: shares or tokens

Gate 2 · MPW
≈$1.5M · month 11
Measured silicon replaces the model

Gate 3 · Masks and first lot
≈$20M · month 24
N7 masks ≈$3.5M inside a $15–100M design



A model, an FPGA and an MPW reach measured silicon for about $1.5 million in eleven months; the mask set and first lot cost more than ten times as much. Stop before masks until the buyer or the protocol version is fixed.

OPPORTUNITIES

Crypto as the first customer, AI as the second market

Custom chips

- Demand is overwhelmingly AI: Broadcom about $58 billion of AI semiconductors in fiscal 2026
- Mining-chip skills transfer: low-voltage arrays, kilowatt power, fast tape-outs
- Win with an independent route to market, not mining channels

AI development

- Hours are a commodity: $25–49 an hour on Clutch
- The value is in the joints: kernels for custom chips, mining inside inference serving
- Verifiable inference: TEEs plus ZK spot checks

AI × crypto systems

- Layered verification: TEE or sampling, then disputes or ZK spot checks
- Distribution and liquidity are the barriers, not cost
- Sign the compute buyer before the token

WHAT I WOULD DO

Three paths, each with a small first step

Path	First step
A. An inference-first accelerator that can also mine	A four-week study of J/op and $/op for an SRAM-resident INT8/FP8 array against the RTX 5090 and H100, then FPGA or mature-node MPW
B. Custom-silicon services: crypto first, AI second	Turn this report's NRE and break-even model into paid feasibility studies with milestones and a fixed protocol version
C. An AI-native network with a buyer before a token	A design note on layered verification and a vendor-neutral canonical arithmetic, written with one named inference buyer

At most three paths, each tested before money is committed analysis.

WHAT TO WATCH

What to watch, and questions to ask

SIGNPOSTS

Any time	Pearl's FP8 activation height, and whether the chain splits
Monthly	Pearl's share of blocks carrying AI work
Quarterly	Bittensor's verifiable revenue against emissions; agent-initiated x402 flows
January 2027	TSMC's reported 3–6% price rise for 2 nm and 3 nm
Q1 2027	FP4 in Pearl, as reported from the founders' AMA
April 2028	Bitcoin's next halving

QUESTIONS TO ASK

01 Who controls the work function?

02 What is the subsidy multiple?

03 What does the next gate cost, and what stops you?

04 What is the payback hashrate, with p?

05 What is the chip's second market?

06 Who is the first paying customer, by name?

These are areas we can explore together, through technical exchange and technical collaboration: custom chips, AI development and AI × crypto systems.

i@tanshuai.com · tanshuai.com