

ChipCoin (CHIP)

Whitepaper — Comprehensive Edition v1.0

Own the Silicon Age



ChipCoin cover — a coin rendered as a microchip resting on an iridescent silicon wafer

Document control

Title	ChipCoin (CHIP) Whitepaper
Version	v1.0 (supersedes internal Draft v0.1)
Status	Pre-launch design document — not an offer of securities
Date	August 2026
Note	Items marked [Proposed] are recommended defaults adopted for this edition; they remain subject to final governance decision and legal review, and must be ratified and published before launch. Items marked [TBD] remain open decisions.

Executive Summary

Compute has become the defining scarce resource of this decade. Artificial intelligence, autonomous systems, and the modern digital economy all compete for the same limited

supply of advanced semiconductors, and every credible forecast indicates that demand for AI compute continues to outpace what fabs and data centers can deliver. The global semiconductor market is projected by Gartner to exceed **\$1.3 trillion in 2026** — the fastest growth in two decades — with the World Semiconductor Trade Statistics (WSTS) Spring 2026 forecast going further, projecting **\$1.51 trillion** and a **\$1.9 trillion** market in 2027, driven overwhelmingly by AI infrastructure demand ^{1 2}. Yet participation in this boom remains concentrated in a handful of mega-cap equities, private data-center deals, and allocated GPU supply chains, leaving ordinary participants with no native, open, community-owned asset for the theme itself.

ChipCoin (CHIP) is a community token built around the semiconductor and AI-compute narrative. It begins as a cultural and thematic asset — a way for anyone to hold “a piece of the silicon age” — and follows a staged, gated roadmap toward practical utility within decentralized compute (DePIN) ecosystems, where tokens already coordinate real GPU power at prices **50–70% below hyperscaler rate cards** ³. The token’s design philosophy is deliberately honest: community and narrative first, utility grown deliberately on top, with each utility phase announced only when it is real.

This edition of the whitepaper does three things beyond the original draft. First, it grounds the narrative in verified market research — the semiconductor super cycle, the hyperscaler capex wave, the memory shortage, and the maturing DePIN sector — with full citations and a research appendix. Second, it resolves the draft’s open design questions into a complete, clearly-marked set of **[Proposed]** defaults covering chain selection, supply, allocation, vesting, and liquidity policy. Third, it presents the full system design: ecosystem architecture, the four-phase utility roadmap, the Chip Watch narrative engine, launch sequencing, regulatory posture, and a candid risk framework.

CHIP makes no claim to be a chip company, a fund, or a security. It does not represent equity, debt, revenue share, or any claim on semiconductor companies or hardware. It is a narrative-first community asset with a transparent, phased plan to become useful — and with the discipline to say so plainly.

1. Background: The Age of Scarce Compute

1.1 The compute crunch


The rise of large-scale AI has transformed semiconductors from a cyclical industry into a strategic bottleneck. After reaching approximately **\$805 billion in 2025**, global semiconductor revenue is forecast by Gartner to grow **64% in 2026 to over \$1.32 trillion**, with AI semiconductors alone accounting for roughly **30% of total industry**

¹ Gartner. <https://www.gartner.com/en/newsroom/press-releases/2026-04-08-gartner-forecasts-worldwide-semiconductor-revenue-to-exceed-us-dollars-one-point-3-trillion-in-2026>

² Semiconductor Industry Association. <https://www.semiconductors.org/global-semiconductor-sales-increase-11-month-to-month-in-april/>

³ BuildMVPFast. <https://www.buildmvpfast.com/blog/depin-gpu-inference-io-net-akash-below-aws-2026>

revenue Error! Bookmark not defined.. WSTS's Spring 2026 forecast is even more aggressive, projecting **90% growth to \$1.51 trillion** in 2026 after monthly sales ran 93.9% above prior-year levels in April 2026, and approximately **\$1.9 trillion** in 2027 ⁴ **Error! Bookmark not defined..** Deloitte estimates the AI chip segment specifically at roughly **\$500 billion in 2026**, following an upward revision driven entirely by AI demand ⁵.



Global semiconductor revenue trajectory, 2023–2027F

Figure 1 — Global semiconductor revenue, 2023–2027F. Gartner's April 2026 forecast (bars) with the WSTS Spring 2026 projection (gold markers) shown for comparison; the memory segment (purple line) nearly triples in a single year. Sources: Gartner (Apr 2026), WSTS/SIA (Jun 2026).

The crunch is not abstract — it is priced into every layer of the stack. Conventional DRAM contract prices surged **90–95% quarter-over-quarter in Q1 2026**, with NAND flash up 55–60%, and Gartner projects full-year 2026 price increases of **125% for DRAM and 234% for NAND** — a condition analysts have named “memflation” ⁶ **Error! Bookmark not defined..** DDR4 spot prices rose 683% in eight months, automotive memory lead times have stretched to 52 weeks, and industry watchers expect the shortage to persist until at least **2028**, when new fab capacity reaches volume production ⁷ ⁸. High-bandwidth


⁴ WSTS. <https://www.wsts.org/76/Recent-News-Release>

⁵ Deloitte. <https://www.deloitte.com/us/en/insights/industry/technology/technology-media-telecom-outlooks/semiconductor-industry-outlook.html>

⁶ Utmel. <https://www.utmel.com/blog/news/semiconductor/the-2026-memory-super-cycle-navigating-the-500-surge-in-dram-and-nand-flash-prices>

⁷ siliconanalysts.com. <https://siliconanalysts.com/analysis>

memory (HBM) — the binding constraint on every new GPU generation — is sold out through 2026, with Micron forecasting the HBM market growing from roughly **\$35 billion in 2025 to \$100 billion by 2028** at a ~40% compound annual rate **Error! Bookmark not defined.**



The memory supercycle: memory vs non-memory revenue, 2025–2027F

Figure 2 — The memory supercycle. Memory nearly triples in 2026 to \$633 billion and becomes 48% of the total chip market. Source: Gartner semiconductor revenue forecast, April 2026.

1.2 The capex supercycle


Behind the chip numbers sits the largest corporate capital expenditure cycle in recorded history. The four largest hyperscalers — Amazon, Alphabet, Microsoft, and Meta — collectively plan approximately **\$725 billion of capital expenditure in 2026**, a ~77% increase over 2025's record \$410 billion, with Amazon alone projecting \$200 billion ⁹. Including Oracle, the five largest US cloud and AI infrastructure providers have committed **\$660–690 billion for 2026**, and every one of them reports being **supply-constrained rather than demand-constrained** — Microsoft alone disclosed an \$80 billion backlog of Azure orders it cannot fulfill due to power constraints ¹⁰. Goldman Sachs projects a combined **\$5.3 trillion** of hyperscaler capex from fiscal 2025 through fiscal 2030, with a

⁸ bandwidth memory v4 supply takes shape. <https://www.blocksandfiles.com/ai-ml/2026/01/21/memory-semiconductor-supercycle-set-to-run-through-2028/4090501>

⁹ yahoo.com. <https://finance.yahoo.com/sectors/technology/article/meta-microsoft-amazon-and-alphabet-are-about-to-spend-a-shocking-amount-of-money-to-dominate-the-ai-era-115359575.html>

¹⁰ futurumgroup.com. <https://futurumgroup.com/insights/ai-capex-2026-the-690b-infrastructure-sprint/>

baseline aggregate estimate of \$7.6 trillion between 2026 and 2031 across compute, data centers, and power **Error! Bookmark not defined..**



Hyperscaler and AI-infrastructure capex, 2025 vs 2026E

Figure 3 — Hyperscaler capex, 2025 actual vs 2026 guidance. Sources: company guidance as compiled by Futurum Research (Feb 2026) and press reports (Jun 2026); Alphabet, Meta, and Microsoft shown at guidance midpoints.

Two features of this wave matter for the CHIP thesis. First, the constraint is migrating rather than disappearing: as GPU availability improved through 2025–2026, the bottleneck shifted to power and grid connections, with Gartner projecting **40% of AI data centers will be power-constrained by 2027** and grid approval timelines running 24–36 months in major markets ¹¹. Second, the spending is concentrated to a historic degree — a structural feature of the market that directly motivates the access problem described next. The build-out is real, it is enormous, and its benefits accrue to a very short list of balance sheets.

1.3 The concentration problem

The value created by the silicon boom accrues to a narrow set of actors. NVIDIA holds roughly **80% of the AI accelerator market by revenue**, generating \$193.7 billion of fiscal-2026 data-center sales — approximately 90% of its total revenue — and consumes about 60% of TSMC’s advanced CoWoS packaging allocation **Error! Bookmark not defined.** ¹². TSMC in turn commands approximately **70–72% of global foundry revenue**,

¹¹ spheron.network. <https://www.spheron.network/blog/ai-data-center-power-constraints-2026/>

¹² trendxinsights.com. <https://trendxinsights.com/blogs/ai-chip-wars-2026-deep-dive/>

meaning a single company manufactures the leading-edge silicon for nearly every competitor in the industry ¹³. Memory is a three-player oligopoly of Samsung, SK hynix, and Micron, whose pricing discipline through the current super cycle has been absolute ¹⁴. Across five leading firms, combined AI and data-center revenue reached \$318.2 billion in the most recent fiscal year — with NVIDIA alone capturing 61% of that pool **Error! Bookmark not defined.**

The barriers protecting this concentration are capital walls that no community can climb. A single leading-edge fab production line costs **\$15–25 billion**, EUV lithography scanners run \$200–350 million each, and a 7nm chip tape-out can exceed \$30 million for a first-time design team ¹⁵. Even the United States government, deploying \$33.7 billion of CHIPS Act grants against projects totaling hundreds of billions, is only expected to lift US fab capacity 203% by 2032 — from a base that had fallen from 37% of global capacity in 1990 to 10% in 2022 ^{16 17}. For individuals, “exposure to compute” effectively means buying shares in a handful of mega-cap companies. There is no open, permissionless, community-owned asset for the theme itself. That gap is the one CHIP is designed to fill — not as a claim on these companies, but as a symbol and, eventually, a coordination tool for the people the boom has left watching from the outside.

1.4 The geopolitical layer

Semiconductors have become instruments of statecraft, which keeps the theme permanently on the front page — free, perpetual narrative fuel. Successive US administrations have imposed escalating export controls on advanced AI chips, from the original A100/H100 restrictions through the April 2025 addition of NVIDIA’s H20 and AMD’s MI308 to the control list ¹⁸. The bipartisan GAIN AI Act would go further, requiring US chipmakers to satisfy domestic demand before exporting advanced integrated circuits — framing GPU access explicitly as a matter of national security and economic stability ¹⁹. China has responded with an accelerating self-reliance drive organized around Huawei’s domestic supply chain, while tariff proposals of 25–100% on imported semiconductors have added a trade-policy dimension to the shortage **Error! Bookmark not defined.**

For a narrative asset, this matters enormously: “who has the chips” is now a recurring headline question, not an industry niche. Every export ruling, fab subsidy announcement, and earnings surprise is a news event that refreshes public attention on the exact theme CHIP represents. A token built on this narrative does not need to manufacture relevance; it

¹³ binance.com. <https://www.binance.com/en/square/post/323039274995793>

¹⁴ uncoveralpha.com. <https://www.uncoveralpha.com/p/every-memory-cycle-ends-the-same>

¹⁵ Prepared Data Center Campus in Indonesia. <https://www.irecruit.co/insights/cost-per-square-foot-semiconductor-fab-construction-2026>

¹⁶ ayna.ai. <https://www.ayna.ai/insights/can-the-us-regain-its-semiconductor-edge>

¹⁷ substack.com. <https://marklapedus.substack.com/p/will-the-us-chips-act-succeed>

¹⁸ balsilliepapers.ca. <https://balsilliepapers.ca/bsia-paper/the-geopolitical-rivalry-behind-us-chip-export-controls-and-its-implications-for-canadas-semiconductor-autonomy/>

¹⁹ The American Action Forum. <https://www.americanactionforum.org/insight/the-gain-ai-act-in-context-export-controls-and-u-s-ai-competitiveness/>

needs the discipline and machinery to convert a continuous real-world news cycle into community energy. That machinery — the Chip Watch engine — is specified in Section 6.

1.5 The decentralized answer

In parallel with the concentration of centralized compute, decentralized physical infrastructure networks (DePIN) have matured from speculative experiments into real, revenue-generating infrastructure businesses. Messari's State of DePIN 2025 report values the sector at roughly **\$10 billion in circulating market capitalization generating an estimated \$72 million of FY2025 on-chain revenue**, with leading networks trading at 10–25× revenue — rational by technology standards — and projects the sector's addressable market at up to **\$3.5 trillion by 2028**²⁰. The ecosystem now spans over **650 projects with 8.8 million active devices across 199 countries**, having peaked near \$19 billion in market cap in September 2025, and attracted roughly \$1 billion of private capital in 2025 alone²¹ **Error! Bookmark not defined..**

The decentralized GPU compute segment demonstrates why this matters economically: by aggregating idle and underutilized hardware with token incentives and verification primitives, networks like io.net (~50,000 aggregated GPUs), Akash, Render, and Aethir deliver H100-class compute at **\$1.40–1.60 per GPU-hour against \$4.20–5.00 at the hyperscalers** — a structural 50–70% discount that saved-workload math makes impossible for cost-sensitive AI builders to ignore^{22 23}. The sector collectively managed an estimated \$200 million in annualized protocol revenue by early 2026, Render alone processed over \$40 million of annualized compute jobs, and the investment case is no longer purely narrative: Bittensor generated \$43 million of revenue from AI customers in Q1 2026, and Grayscale filed for a Bittensor Trust in December 2025 — the first institutional vehicle for a DePIN-adjacent asset^{24 25 26}.

²⁰ Messari. <https://messari.io/report/state-of-depin-2025>

²¹ BlockEden.xyz. <https://blockeden.xyz/blog/2026/03/21/depin-march-2026-reality-check-650-projects-19b-market-cap-revenue/>


²² Prompt20 Blog. <https://blog.prompt20.com/posts/decentralized-gpu-compute/>

²³ io.net. <https://io.net/blog/io-net-vs-akash-network-comparing-gpu-cloud-pricing-and-features>

²⁴ Yellow Exchange. <https://yellow.com/research/ai-compute-demand-crypto-gpu-networks-gap-2026>


²⁵ weex.com. <https://www.weex.com/news/detail/best-ai-crypto-coins-2026-top-7-tokens-ranked-by-data-u6p58bez0nza5hxvqn1iga6n>

²⁶ Phemex. <https://phemex.com/blogs/ai-tokens-profitable-crypto-sector-q1-2026>



H100 GPU hourly pricing by provider, mid-2026

Figure 4 — H100-class GPU pricing by provider, mid-2026. Decentralized networks (teal) run 50–70% below hyperscaler on-demand rates (gray). Sources: provider rate cards and independent comparisons, Q2 2026.



Q1 2026 returns: AI-compute tokens vs majors

Figure 5 — Q1 2026: AI-compute tokens were the only profitable crypto sector while Bitcoin and Ethereum fell sharply. Source: Phemex sector scoreboard, March 2026.

These networks prove the central premise that CHIP's roadmap builds on: **tokens can coordinate real hardware**. They also define the ecosystem CHIP intends to plug into — not as a competitor spinning up yet another GPU network, but as a community layer that routes attention, users, and eventually treasury incentives toward the decentralized compute economy. The distinction is important for honesty and for strategy: CHIP does not need to out-engineer io.net or Akash. It needs to become the cultural front door through which their future users arrive.

2. Vision: Flag First, Tool Second

2.1 Layer 1 — The Flag

CHIP's first layer is symbolic, and the project does not apologize for that. Every era's defining resource eventually gets a symbol people can hold: gold coins for scarce metal, oil majors for the automotive age. The silicon age — a \$1.3-trillion-and-accelerating industry **Error! Bookmark not defined.** — has lacked a symbol that belongs to everyone rather than to those with fab-scale capital. Holding CHIP is a statement: compute is the new gold, and its story belongs to the crowd, not only to the balance sheets of five hyperscalers. Symbols with genuine cultural resonance are not a lesser form of value in crypto; they are the primary form, as the sector's own history demonstrates daily.

The Flag layer is engineered with the same seriousness as the utility layers that follow. A fixed, immutable supply; a transparent, verifiable launch; renounced mint authority; and a published, never-silently-changed tokenomics policy. The symbolism works only if the asset itself is clean — a flag nobody can trust is just a banner ad. This is why Section 4 specifies supply integrity and liquidity policy before it specifies anything else, and why the draft's insistence that "final tokenomics must be published before launch and never changed silently" is preserved here as a hard design constraint rather than a footnote.

2.2 Layer 2 — The Tool

Over time, CHIP aims to become a bridge between its community and real compute. The direction is set by what the DePIN sector has already proven: decentralized GPU marketplaces route real workloads to real hardware at 50–70% below hyperscaler pricing, using tokens as both currency and coordination mechanism **Error! Bookmark not defined.** CHIP's Tool layer connects holders to that economy through three mechanisms, each corresponding to a roadmap phase: token-gated access to compute-market analytics and partner offers (Phase 2), swap-and-stake mechanics converting CHIP into compute credits on partner networks (Phase 3), and a holder-governed Foundry Treasury that funds builders and compute providers — a community-directed micro "sovereign compute fund" (Phase 4).

The order matters, and the gating matters more. Community and narrative come first because they are the only layers a new token can deliver with certainty on day one. Utility follows deliberately, with each phase announced **only when it is real** — partnerships signed, integrations functional, dashboards live. This sequencing discipline is the project's core credibility mechanism: in a market where most theme tokens promise utility they never ship, CHIP's refusal to market unbuilt utility is itself the differentiator.

2.3 Design principles

Five principles govern every design decision in this document. **Narrative honesty**: never claim live compute utility, hardware backing, or named partnerships before they exist; Phases 2–4 are framed as exploration goals everywhere they appear. **Transparency by default**: supply, allocation, vesting contracts, treasury wallets, and LP policy are published and verifiable on-chain, not described in prose alone. **Fair launch orientation**: no insider allocation mechanics hidden in launch parameters; the largest single allocation goes to liquidity and launch, not to the team. **News-cycle leverage**: the project's content engine is architected to ride the semiconductor news cycle rather than fight for attention against it. **Progressive decentralization**: control migrates from a launch multisig toward holder governance as the phases advance, with the treasury designed from day one for eventual DAO operation.

These principles are not marketing language; they resolve real trade-offs encountered in the design. For example, the tension between “move fast on narrative momentum” and “never market unbuilt utility” is resolved by the phase-gating rule in Section 5.5. The tension between “treasury funds growth” and “treasury must not become an insider honeypot” is resolved by multisig control, published quarterly reporting, and the Phase 4 migration to holder governance. Where this document marks a parameter **[Proposed]**, the principle framework above is the rationale; where it marks **[TBD]**, the decision is genuinely open and flagged for the community and legal counsel to close.

3. Market Context: Lessons from Narrative Assets

3.1 The survival statistics of theme tokens

Any honest whitepaper for a narrative token must begin with the survival math of its own category, because that math is brutal. Over **36,000 new meme and theme tokens launch daily** on platforms like Pump.fun; according to aggregated industry data, **97% fail**, 60% are dead within 24 hours, and fewer than 2% ever graduate from launchpad bonding curves to major DEX listings²⁷. The average meme token survives roughly one year — a third of the lifespan of other crypto projects **Error! Bookmark not defined.** Yet the survivors demonstrate the category's ceiling: the meme sector commanded a **\$37.6 billion market cap** in 2026 with billions in daily volume, and tokens like PEPE, WIF, and BONK

²⁷ Medium. <https://medium.com/@chainpeak/why-some-meme-coins-go-viral-while-others-die-in-72-hours-the-marketing-playbook-that-actually-434a66334e5d>

have sustained billion-dollar valuations and multi-year communities ²⁸ **Error! Bookmark not defined..**

The autopsy of the failures is remarkably consistent, and each failure mode maps to a specific CHIP design countermeasure. Tokens die because launch is treated as the finish line (CHIP: the launch plan in Section 7 is explicitly the start of operations, with the Chip Watch engine designed for daily cadence); because a logo plus paid KOL tweets substitutes for culture (CHIP: Proof of Silicon and The Foundry build identity around hardware culture, not influencer rotation); and because there is no narrative refresh mechanism after the first hype wave (CHIP: the semiconductor news cycle provides perpetual, externally generated narrative fuel, specified in Section 6.3). CHIP cannot guarantee survival — nothing can — but every named failure mode has a designed answer rather than a hope.

3.2 What survivors do differently

The tokens that endure share three traits, all visible in the 2026 market data. First, they attach to a narrative with independent real-world momentum: in Q1 2026, AI-compute tokens were the **only profitable crypto sector**, with Bittensor up ~40%, Render up 32%, and Fetch.ai up 67% while Bitcoin fell 23% and Ethereum fell 32% — capital rotated toward tokens whose underlying story (compute demand) was verifiably true in the physical economy **Error! Bookmark not defined..** Second, they develop revenue or utility floors over time: Bittensor's \$43 million of Q1 protocol revenue and Render's \$40 million of annualized compute jobs gave those tokens support that pure speculation could not **Error! Bookmark not defined. Error! Bookmark not defined..** Third, they build institutions rather than audiences — Grayscale's Bittensor Trust filing and over \$620 million of subnet staking show what it looks like when a narrative asset crosses into infrastructure **Error! Bookmark not defined..**

CHIP's positioning synthesizes these lessons into a hybrid: it launches with the cultural energy and accessibility of a theme token, on the single strongest macro narrative of the decade, with a credible staged path to the revenue-adjacent utility that separated 2026's winners from the graveyard. The DePIN sector's own metrics — +24.95% sector performance against a \$9.4 billion market cap in 2026 narrative tracking, and AI tokens at \$22.6 billion — confirm that the compute theme is where narrative attention and real-economy verification already intersect **Error! Bookmark not defined..** CHIP does not need to convince the market that compute matters; it needs to be the cleanest, most honest vehicle for a belief the market already holds.

3.3 Competitive positioning

CHIP occupies a deliberately distinct position relative to the assets its holders will compare it against. It is not a DePIN protocol token (those are potential partners, not competitors), not a meme coin with no roadmap (the phased utility path is the differentiator), and not a

²⁸ KuCoin. <https://www.kucoin.com/blog/en-top-10-crypto-sectors-in-2026-a-data-driven-comparison-of-market-trends-and-growth-opportunities>

security-like claim on chip company cash flows (explicitly disclaimed). The table below maps the positioning.

Asset class	Example	Value driver	Utility floor	CHIP's relationship
Meme/theme tokens	PEPE, WIF, BONK	Culture, attention	None	Same accessibility; CHIP adds gated utility phases
DePIN protocol tokens	RENDER, AKT, IO, TAO	Protocol revenue, hardware coordination	Real and growing Error! Bookmark not defined.	Integration partners for Phase 3, not rivals
AI-narrative tokens	FET/ASI, agent tokens	AI story + protocol usage	Mixed Error! Bookmark not defined.	Adjacent narrative; CHIP is the compute-resource pure play
Mega-cap chip equities	NVDA, TSM	Cash flows, fabs	Full	The concentration CHIP exists to complement for retail
CHIP	—	Silicon-age narrative + staged compute utility	Built through Phases 2–4	Community layer for the compute economy

The strategic conclusion is that CHIP's real competition is attention, not other protocols. Its moat, if one forms, will be the quality and honesty of its narrative engine plus the credibility earned by disciplined phase-gating. Both are execution-dependent, which the risk section states plainly. What the design can and does control is that CHIP enters the market with a complete, published, verifiable structure — something the 97% failure cohort overwhelmingly lacks **Error! Bookmark not defined..**

4. The CHIP Token

4.1 Specification

The table below sets out the complete token specification for this edition. Values adopted as recommended defaults are marked **[Proposed]** and remain subject to final ratification and legal review before launch; genuinely open items are marked **[TBD]**.

Item	Detail	Status
Name	ChipCoin	Final
Ticker	CHIP	Final
Chain	Solana	[Proposed]
Standard	SPL	[Proposed]

Item	Detail	Status
Total supply	1,000,000,000 CHIP, fixed forever	[Proposed]
Mint authority	Renounced at launch	[Proposed]
Freeze authority	Renounced at launch	[Proposed]
Decimals	9 (Solana default)	[Proposed]
Launch venue	Fair launch via reputable launchpad, liquidity migrated to Raydium/Jupiter	[Proposed]
LP policy	LP tokens locked ≥ 12 months or burned at migration; policy published pre-launch	[Proposed]
Treasury control	Multisig (≥ 3 -of-5) from day one; migrates to DAO governance in Phase 4	[Proposed]
Jurisdictions	Restricted jurisdictions per legal review	[TBD]

The specification’s intent is that every trust-relevant property of the token be verifiable on-chain rather than asserted in prose. A fixed supply with renounced mint and freeze authority removes the entire class of “team prints more tokens” risks at the protocol level. Publishing the LP policy before launch — and executing it visibly at launch — removes the most common rug mechanism in the theme-token category. None of these choices is novel; all of them are, collectively, what separates a clean launch from the daily graveyard documented in Section 3.1.

4.2 Chain selection rationale

Solana is the **[Proposed]** base chain for four evidence-based reasons. First, liquidity and attention: as of mid-2026, Solana is the established leader for new token launches, offering the best mix of sub-cent fees, sub-second finality, launchpad infrastructure (Pump.fun and successors), and DEX routing depth (Raydium, Jupiter) — the chain where new narrative assets consistently find the deepest speculative liquidity²⁹. Second, economics: deployment and liquidity costs on Solana run to roughly 12–14 SOL all-in for a basic launch, versus \$20–100 per transaction in gas on Ethereum mainnet — a decisive factor for a token whose community will include small first-time holders^{30 31}. Third, ecosystem fit: the DePIN networks CHIP’s roadmap targets are heavily Solana-native — Render migrated to Solana, io.net is Solana-based — minimizing future integration friction for Phase 3 compute-credit mechanics **Error! Bookmark not defined.** Fourth, category conventions: the chain’s launchpad ecosystem, despite its noise, provides the fairest available launch mechanics for a community-first distribution³².

²⁹ Meet Parse. <https://parse.gl/prompts/p/as-a-meme-coin-trader-which-chain-solana-vs-base-vs-bsc-vs-eth-currently-offers-the-best-mix-of-low-fees-high-volume-and-reliable-liquidity-for-new-launches--ca2b9658-1c5c-4efb-9bb0-472de782d6b7>

³⁰ troniextechnologies.com. <https://www.troniextechnologies.com/blog/meme-coin-development-cost>


³¹ Crypto News. <https://cryptonews.com/cryptocurrency/best-meme-coins/>

³² Join the world’s oldest Bitcoin advocacy organization.. <https://bitcoinfoundation.org/news/altcoins/best-memecoin-launchpads/>

Base was evaluated as the strongest challenger — Coinbase distribution, low fees, and fast-growing meme liquidity make it genuinely attractive — and BNB Chain offers the largest raw retail base **Error! Bookmark not defined.** Solana wins the recommendation on the combination of DePIN-native ecosystem alignment and launch liquidity depth. This remains a **[Proposed]** default: if launch timing coincides with a structural shift in launch-venue liquidity, the decision should be revisited with fresh data before contract deployment, and the rationale published either way.

4.3 Allocation

The **[Proposed]** allocation preserves the draft’s structure, which was well-designed: community-oriented, with the team allocation deliberately small. The governing rule — published before launch, never changed silently — is retained as a binding commitment.



CHIP token allocation, proposed model

Figure 6 — CHIP allocation, proposed model. Fixed supply of 1,000,000,000 CHIP.


Allocation	%	Tokens	Notes	Status
Liquidity & launch	40%	400,000,000	Paired at launch; LP locked ≥12 months or burned	[Proposed]
Community & rewards	30%	300,000,000	Airdrops, campaigns, Chip Watch incentives; 24-month linear vest	[Proposed]
Foundry Treasury (DAO)	15%	150,000,000	Ecosystem grants, DePIN partnership incentives; multisig-controlled, DAO in Phase 4	[Proposed]

Allocation	%	Tokens	Notes	Status
Team & contributors	10%	100,000,000	12-month cliff, then 24-month linear vest; on-chain vesting contracts	[Proposed]
Marketing & listings	5%	50,000,000	Exchange listings, market-making, promotions; 12-month linear	[Proposed]

Three properties of this allocation deserve emphasis. At 40%, the liquidity-and-launch tranche is the largest single allocation, and at 10% with a full-year cliff, the team tranche is at the conservative end of market norms — the inverse of the insider-heavy structures that precede most theme-token collapses. The 30% community tranche is the fuel for the Chip Watch and Proof of Silicon programs, vested over 24 months so that campaign rewards cannot flood the market in the opening weeks. The 15% Foundry Treasury is the seed of the Phase 4 “micro sovereign compute fund,” and its DAO migration path is specified at launch rather than promised for later.

4.4 Vesting and unlock schedule

Predictability of supply expansion is a core trust property for any new token. The **[Proposed]** schedule below makes the entire 36-month unlock trajectory public before launch and enforceable on-chain through audited vesting contracts.



Proposed CHIP unlock schedule, 36 months

Figure 7 — Proposed unlock schedule. 40% of supply (the liquidity tranche) is circulating from day one by construction; all other tranches vest on published linear schedules, with the team tranche fully locked for the first 12 months.

Tranche	Cliff	Vesting	Fully unlocked
Liquidity & launch	—	Unlocked at TGE (paired into LP)	M0
Community & rewards	—	Linear, months 0–24	M24
Foundry Treasury	—	Linear, months 0–36; spend additionally DAO-gated	M36
Team & contributors	12 months	Linear, months 12–36	M36
Marketing & listings	—	Linear, months 0–12	M12

The shape of the curve is the message: no insider tranche moves for a year, community emissions taper smoothly rather than cliff-dumping, and total circulating supply reaches roughly 79% by month 24 and 100% only at month 36. Because the liquidity tranche’s 40% is circulating from launch by design, the effective *new* supply entering the market in year one is limited to the marketing tranche and half the community tranche — a materially gentler emissions profile than most launchpad tokens, which typically unlock team and private allocations within months.

4.5 Liquidity policy


Liquidity policy is specified separately because it is where theme tokens most often fail their holders. The **[Proposed]** policy has three components. At launch, the full 40% liquidity tranche is paired against the launch proceeds on a reputable Solana launchpad and migrated to a Raydium/Jupiter pool at graduation, following the ecosystem’s standard fair-launch path **Error! Bookmark not defined. Error! Bookmark not defined..** At migration, LP tokens are either burned or locked for a minimum of 12 months via a public locker, with the choice and the lock transaction published; this removes the single most common rug mechanism in the category. Thereafter, the Foundry Treasury may add liquidity under multisig approval, and all treasury wallet addresses are published in the launch documentation and tracked in quarterly transparency reports.

A candid note on market-making: the 5% marketing-and-listings tranche contemplates professional market-making on any centralized exchange listings, because thin books punish community holders more than they punish insiders. Any market-making engagement will be disclosed in aggregate terms (venue, not strategy) in the quarterly reports. The policy’s spirit is that everything material to holder trust is either on-chain, published, or both — and that “trust us” appears nowhere in the token’s operational design.

5. Utility Roadmap

5.1 Overview and gating rule

The roadmap converts CHIP from flag to tool in four phases. Every phase after the first is an **exploration goal, not a promise**, and the governing rule — repeated deliberately from the draft — is that each utility phase is announced only when it is real: partnerships signed, integrations functional, deliverables live. This is the project’s core honesty mechanism, and it is stated on every surface where the roadmap appears.



CHIP utility roadmap, four phases

Figure 8 — The four-phase roadmap from Flag to Foundry DAO, with the gating rule between phases.

5.2 Phase 1 — Flag (launch)

Phase 1 delivers everything a new token can deliver with certainty on day one: a clean launch, a token-gated community called **The Foundry**, and the **Chip Watch** content engine. Membership roles (“Fabricators”) are tied to participation, not only holdings, so that contribution — creating content, running campaigns, welcoming new members — earns standing that money cannot simply buy. Proof of Silicon culture programs (rig showcases, GPU lore, chip-history content) give the community an identity rooted in the hardware the token celebrates, rather than in price talk alone.

Phase 1 is also where the token’s narrative reflexes are built. The semiconductor news cycle in 2026 delivers a trigger event almost weekly — memory price surges, export rulings, fab groundbreakings, earnings surprises **Error! Bookmark not defined. Error! Bookmark not defined.** Chip Watch converts each into briefings, campaigns, and memes within hours, training the community to experience real-world chip news as *their* news.

The phase's success metrics are community health metrics — active Fabricators, briefing reach, campaign participation — not price.

5.3 Phase 2 — Access (post-launch)

Phase 2 adds the first token-gated utility: compute-market analytics and partner perks. The deliverables are a holder-gated dashboard tracking GPU spot prices across centralized and decentralized providers (the data underlying Figure 4, made live), memory and component price indices, and DePIN network metrics — genuinely useful tooling for an audience of AI builders, GPU owners, and narrative traders, and a concrete reason to hold CHIP beyond the trade. Partner offers — discounts, priority access, or credits negotiated with compute and DePIN projects — extend the Access layer outward.

The phase is deliberately scoped to what the team can build without external dependencies: dashboards and analytics are in-house work, gated by token holdings through standard Solana tooling. Partner perks require counterparties, which is why they are listed second within the phase and announced only when signed. The strategic purpose of Phase 2 is to establish the pattern of shipped utility before any ask for patience on the bigger Phase 3 integration — credibility is built in ascending order of difficulty.

5.4 Phase 3 — Compute Credits (roadmap goal)

Phase 3 is the roadmap's centerpiece and the point where “own the silicon age” becomes literal: integrations allowing CHIP to be swapped or staked for compute credits on partner decentralized GPU networks. The target partners are the networks validated in Section 1.5 — io.net, Akash, Render, Aethir, or their successors — which already settle real workloads for real customers at \$1.40–2.50 per H100-hour **Error! Bookmark not defined. Error! Bookmark not defined.** Mechanically, the integration pattern most likely to succeed is a credit-rail model: CHIP is swapped or staked through a partner-facing interface, the holder receives compute credits denominated in dollars or GPU-hours, and the partner network receives a new demand channel plus treasury-funded incentives from the Foundry Treasury.

The honesty requirements for this phase are the strictest in the document. It is explicitly contingent on third-party partnerships that may not materialize; it is announced only when agreements are signed; and no material — website, pitch, or community briefing — may describe compute-credit functionality as live, imminent, or guaranteed before that point. The phase's fallback, if integrations stall, is not silence: the Access layer's dashboards already orient the community toward decentralized compute, and treasury incentives can flow to providers and builders through grants without a formal swap integration.

5.5 Phase 4 — Foundry DAO (long-term)

Phase 4 completes the arc: the Foundry Treasury migrates from multisig stewardship to holder governance, and the community becomes what the concept document calls a micro “sovereign compute fund.” Treasury priorities target the real compute economy — grants to indie AI builders, incentives for GPU providers joining partner networks, funding for compute-market public goods (open benchmarks, price oracles, developer tooling).


Governance uses the token the community already holds, with proposal and voting mechanics ratified before activation.

The phase's design deliberately borrows from the DePIN sector's hard-won lesson that token incentives can bootstrap supply but cannot subsidize demand forever **Error! Bookmark not defined.** Foundry DAO grants are therefore oriented toward demand-side builders and provider onboarding rather than open-ended emissions, and quarterly transparency reports continue under DAO operation so that treasury flows remain auditable by the same community that directs them. The end state is a community that doesn't merely hold a symbol of the compute economy but allocates capital within it.

6. Ecosystem Architecture & Community Design

6.1 Architecture overview

The full system stacks four layers: a public base chain, the protocol layer (token core, governance, transparency), the phased utility layer, and the external compute economy the project plugs into. The architecture diagram below presents the proposed target state; layers above the base activate progressively through the roadmap phases.



CHIP ecosystem architecture

Figure 9 — CHIP ecosystem architecture. The base layer and token core exist at launch; utility modules activate by phase; external DePIN networks and the semiconductor news cycle are the two external systems the architecture is designed to channel.

Two architectural choices warrant explanation. First, CHIP does not build its own chain, its own GPU network, or its own marketplace — it is a community and coordination layer that integrates with proven external infrastructure, which keeps the attack surface, regulatory surface, and engineering burden minimal. Second, the news cycle is drawn as an architectural input on par with DePIN networks, because for a narrative asset it is exactly that: an external system that continuously injects energy, which the Chip Watch module is designed to capture.

6.2 The Foundry

The Foundry is the holder community, framed as a collective “fab” minting the future. Its design borrows from the survivor playbook of Section 3.2: identity before speculation, participation before holdings. Fabricator ranks are earned through contribution — producing Chip Watch briefings, running Proof of Silicon showcases, onboarding new members, building tooling — with roles recorded and recognized publicly. The community allocation (30% of supply, vested over 24 months) funds these programs on published campaign schedules, so rewards are predictable rather than discretionary.


The Foundry also carries the project’s culture code: confident, techno-optimist, slightly playful, “in on the trade” energy — but never financial advice, never price promises, never utility claims ahead of reality. Community guidelines codify the honesty constraints of this whitepaper at the moderator level, because a community that hypes unbuilt utility would destroy the project’s core differentiator faster than any competitor could. The mascot (“Chippy”) exists for meme energy in community spaces, kept visually separate from the serious wafer-and-gold identity of official documents such as this one.

6.3 The Chip Watch flywheel

Chip Watch is the project’s growth engine, and it is designed as a flywheel rather than a marketing calendar. Real-world semiconductor events — a memory price surge, an export ruling, a fab opening, a blowout earnings report **Error! Bookmark not defined. Error! Bookmark not defined.** — trigger rapid-response briefings; briefings feed campaigns and memes (Proof of Silicon, the Global Chip Count tracker); campaigns grow the community; a larger, more engaged community deepens liquidity and signal strength; and that engagement makes the next news-triggered campaign reach further. The semiconductor cycle, in effect, does the marketing — the flywheel just needs to keep turning.

The Chip Watch Narrative Flywheel

The semiconductor news cycle does the marketing



The Chip Watch narrative flywheel

Figure 10 — The Chip Watch flywheel: external news triggers convert to content, campaigns, community growth, and compounding attention.

The flywheel's operating requirements are modest but real: a small editorial rota with authority to publish within hours of a trigger, templated formats (briefing card, meme pack, thread skeleton) prepared in advance for the five most common event types, and campaign budgets drawn from the community tranche on published schedules. The failure mode this design avoids is the one that kills most theme tokens after week one: the silence that follows when the launch hype fades and there is no externally-driven reason to talk about the project. For CHIP, the reasons arrive on the wires every week, courtesy of the largest capex cycle in history **Error! Bookmark not defined..**

6.4 Builder alignment

The treasury's alignment strategy gives the token a credible reason to exist beyond the trade. Rather than generic ecosystem grants, Foundry Treasury spending is targeted at the compute economy the narrative celebrates: DePIN node operators bringing GPU supply online, indie AI builders who are the decentralized networks' natural customers, and open compute-market infrastructure (benchmarks, dashboards, price oracles). Grant criteria,

amounts, and recipients are published quarterly — the same transparency standard applied to every treasury flow.

This alignment is also strategically defensive. A theme token whose only relationship to its theme is a name has no defense when the narrative rotates; a token whose treasury has funded dozens of compute-economy builders has relationships, reputation, and integration surface that survive sentiment cycles. The DePIN sector's own evolution — where capital now flows selectively to projects with real revenue and enterprise customers **Error! Bookmark not defined.** — shows the direction of travel: credibility compounds for projects that fund real work, and the Foundry Treasury is designed to earn exactly that compound.

7. Launch Plan

7.1 Pre-launch checklist

The launch sequence begins only when every gate below is cleared. The order is deliberate: legal and tokenomics finalization precede any public artifact, because a narrative asset that has to retract or restate its structure after publication has destroyed its one asset.

#	Gate	Owner	Status
1	Final tokenomics ratified (all [Proposed] items resolved)	Core team + community review	Open
2	Legal review: jurisdictions, disclaimers, participation restrictions	External counsel	Open [TBD jurisdictions]
3	Contract deployment + third-party review of token and vesting contracts	Engineering + auditor	Open
4	Publication pack: contract address, supply, allocation, vesting contracts, LP policy, treasury wallets	Core team	Open
5	Site + socials live with wafer/foundry identity	Design	Open
6	Chip Watch templates and editorial rota operational	Community lead	Open

7.2 Launch sequence

Launch itself follows five steps, synchronized with the industry news cycle for maximum narrative leverage. Step one: deploy the token and vesting contracts, publish the full publication pack, and verify every on-chain claim (renounced authorities, locked vesting, treasury multisig) against the published documentation. Step two: fair-launch the liquidity tranche through a reputable Solana launchpad, migrating to Raydium/Jupiter at graduation and burning or locking LP as published **Error! Bookmark not defined.** Step three: activate The Foundry with opening campaigns — Proof of Silicon showcases and the Global

Chip Count tracker — timed to ride a live chip-industry news moment under the “Own the silicon age” banner.

Step four: sustain the first ninety days on the Chip Watch cadence — briefings on every material semiconductor event, weekly community programming, monthly transparency notes — because the post-launch quarter is where 97% of the category dies **Error! Bookmark not defined..** Step five: begin Phase 2 utility announcements only after launch stability is demonstrated (orderly markets, published reporting cadence, community health metrics), and only for deliverables that are actually live. The sequence’s meta-rule overrides any calendar pressure: a delayed announcement beats an unbuilt one, every time.

7.3 Post-launch operations

Steady-state operations run on three published cadences. **Quarterly transparency reports** cover treasury balances and flows, vesting contract status, grant recipients, and aggregate market-making disclosures — the ongoing evidence that the structure published at launch remains the structure in operation. **Monthly community programming** rotates Proof of Silicon events, builder spotlights, and Foundry governance discussions, keeping participation pathways open to new Fabricators. **Continuous Chip Watch operations** convert the news cycle per Section 6.3, with a quarterly retrospective on which trigger types produced the strongest engagement so the engine tunes itself.

Operational security is specified at the same standard: treasury multisig signers are doxxed or institutionally anchored, signing policies require out-of-band confirmation for transactions above threshold, and community channels run standard anti-phishing hygiene (verified links, no-DM policies, slow-mode during volatility events). These are unglamorous details, but they are the ones whose absence fills the post-mortems of failed launches, and they are published here so the community can hold the operators to them.

8. Regulatory & Compliance Posture

The regulatory environment for digital assets in 2026 is clearer than at any prior point, and the project is designed to operate within it conservatively. In the United States, the GENIUS Act — signed July 18, 2025 — established the federal stablecoin framework, with agency rulemaking in progress and broader market-structure legislation (the CLARITY Act) still pending in the Senate ³³. In the European Union, MiCA is fully in force, with the transition period having ended July 1, 2026; it comprehensively covers utility tokens and CASP licensing **Error! Bookmark not defined..** CHIP is designed as a community token with no yield, no revenue share, no claim on assets, and no promise of profit from the efforts of others — and its public materials are written to keep it that way: the roadmap’s utility phases are framed as exploration goals, and this document carries explicit non-security, non-advice disclaimers.

³³ T Recommendation H.263. <https://tech-insider.org/genius-act-vs-mica/>

Concrete compliance measures follow from that posture. The project will obtain jurisdiction-specific legal review **before launch** and restrict participation where required — the jurisdictions item remains [TBD] pending that review. Launch communications avoid price speculation, return language, and any characterization of CHIP as an investment; the quarterly transparency reports are factual operational disclosures, not performance marketing. Because Phase 3 compute-credit mechanics involve exchanging tokens for services, that phase will receive its own legal review at announcement time, in each jurisdiction where it is offered, before activation. The project's rule is simple: counsel reviews structure before launch, phases before activation, and jurisdictions before access.

9. Risk Factors & Disclaimers

Participation in CHIP involves substantial risk, including the possible loss of the entire contribution. The principal risks are set out below, candidly and without mitigation language that would dilute them.

Risk	Description
Market risk	CHIP is a highly volatile, narrative-driven asset. Theme tokens as a category fail at a ~97% rate, and 60% die within 24 hours of launch Error! Bookmark not defined.. Price can go to zero, and past narrative strength guarantees nothing.
Narrative-timing risk	Theme tokens live and die by narrative timing. The semiconductor cycle is strong now — record capex, a memory super cycle forecast through 2028 Error! Bookmark not defined. — but cycles turn; an AI capex pullback or memory oversupply (a realistic 2028–2029 scenario per industry analysts Error! Bookmark not defined.) would weaken the narrative engine itself.
Roadmap risk	Phases 2–4 are intentions, not commitments. Compute-credit integrations depend entirely on third-party partnerships that may not materialize, on timelines the project does not control.
Partner/ecosystem risk	DePIN is an early sector: \$72 million of FY25 on-chain revenue against a ~\$10 billion circulating market cap means most projects are valued on growth assumptions, and individual partner networks may fail, be exploited, or lose competitiveness to centralized clouds Error! Bookmark not defined. Error! Bookmark not defined..
Smart-contract & operational risk	Token, vesting, and future integration contracts may contain defects; multisig signers may be compromised; launch and liquidity operations may be disrupted. Audits and published procedures reduce but do not eliminate these risks.
Liquidity risk	Even with locked LP, secondary liquidity can evaporate in risk-off markets, and holders may be unable to exit positions at any

Risk	Description
	reasonable price.
Regulatory risk	Digital-asset regulation varies by jurisdiction and continues to evolve (GENIUS Act rulemaking in progress; CLARITY Act pending; MiCA fully applicable) Error! Bookmark not defined.. Future rules could restrict the token's distribution, trading, or planned utility in some jurisdictions.
No claims	CHIP does not represent equity, debt, revenue share, or any claim on semiconductor companies, GPU hardware, compute capacity, or any other asset. Holding CHIP confers no rights except those encoded in its contracts.

This document is not investment advice, is not an offer to sell or a solicitation of an offer to buy securities in any jurisdiction, and does not constitute financial, legal, or tax advice. It is a design document for a community token, prepared for general informational purposes. Forward-looking statements — including every roadmap phase, market projection, and integration goal — are subject to change without notice and may never be realized. Anyone considering participation should consult independent professional advisers and participate only with funds they can afford to lose entirely.

10. Conclusion

Every era's defining resource eventually gets a symbol people can hold. The silicon age — a \$1.3 trillion industry in 2026 growing at its fastest pace in two decades, fueled by a \$725-billion-a-year hyperscaler capex wave and a memory super cycle with no relief in sight before 2028 **Error! Bookmark not defined. Error! Bookmark not defined.** — has lacked one. At the same moment, the crypto ecosystem has finally produced the missing counterpart: decentralized compute networks that prove tokens can coordinate real hardware for real customers, at prices the hyperscalers cannot match **Error! Bookmark not defined..**

ChipCoin is a community's claim on that story — starting as a flag, engineered with discipline to grow into a tool. Its design bets that honesty is a competitive advantage in a category defined by overpromising: publish everything, gate every utility phase behind reality, vest insiders longest, and let the most powerful macro narrative of the decade supply the energy. The world will never stop needing more compute. **CHIP. Own the silicon age.**

Appendix A: Market Research Digest

The tables below consolidate the external research underlying this whitepaper's market claims. Full source links are provided in the footnotes attached to each figure and data point in the body text.

A.1 Semiconductor market trajectory

Year	Total market	Memory segment	Source
2023	\$526.8B	—	SIA ³⁴
2024	\$626.0B	—	Gartner Error! Bookmark not defined.
2025	\$805.3B	\$216.3B	Gartner (Apr 2026) Error! Bookmark not defined.
2026F	\$1,320.2B (+64%)	\$633.3B (~3×)	Gartner (Apr 2026) Error! Bookmark not defined.
2026F (alt)	\$1,510B (+90%)	>\$800B (+250%)	WSTS Spring 2026 Error! Bookmark not defined.
2027F	\$1,554.5B	\$748.1B	Gartner (Apr 2026) Error! Bookmark not defined.
2027F (alt)	~\$1,900B (+27%)	+32%	WSTS Spring 2026 Error! Bookmark not defined.

A.2 AI infrastructure capex

Provider	2025 capex	2026E capex	Source
Amazon	\$125B	\$200B	Company guidance Error! Bookmark not defined.
Alphabet	\$91B	\$175–185B	Company guidance Error! Bookmark not defined.
Meta	\$72B	\$115–135B	Company guidance Error! Bookmark not defined.
Microsoft	~\$90B	\$110–120B+	Company guidance ³⁵
Oracle	~\$21B	~\$50B	Company guidance Error! Bookmark not defined.
Big Four total	\$410B	\$725B (+77%)	Goldman Sachs / press Error! Bookmark not defined.
Cumulative 2025–2030	—	\$5.3T	Goldman Sachs Error! Bookmark not defined.

A.3 Memory super cycle indicators

Indicator	Value	Source
DRAM contract price, Q1 2026 QoQ	+90–95%	TrendForce via industry press Error! Bookmark not defined.

³⁴ New Market Pitch. <https://newmarketpitch.com/blogs/news/semiconductor-market-size>

³⁵ Data Center Richness. <https://datacenter richness.substack.com/p/hyperscalers-plan-630-billion-in>

Indicator	Value	Source
NAND contract price, Q1 2026 QoQ	+55–60%	TrendForce via industry press Error! Bookmark not defined.
Full-year 2026 DRAM / NAND price forecast	+125% / +234%	Gartner Error! Bookmark not defined.
DRAM spot price, Q4 2025 YoY	~3×	Market analysis Error! Bookmark not defined.
HBM TAM 2025 → 2028E	~\$35B → ~\$100B (~40% CAGR)	Micron Error! Bookmark not defined.
HBM market 2026E	\$54.6B (+58%)	Bank of America ³⁶
Supplier DRAM inventories, Q3 2025	~3.3 weeks (2018-cycle lows)	Market analysis Error! Bookmark not defined.
Shortage expected duration	Into 2028	Industry consensus Error! Bookmark not defined.

A.4 Decentralized compute economics

Metric	Value	Source
DePIN sector circulating market cap	~\$10B	Messari Error! Bookmark not defined.
DePIN FY2025 on-chain revenue	~\$72M	Messari Error! Bookmark not defined.
DePIN projects / devices	650+ / 8.8M across 199 countries	Messari / DePINScan Error! Bookmark not defined.
DePIN addressable market, 2028E	Up to \$3.5T	Messari projection Error! Bookmark not defined. Error! Bookmark not defined.
Decentralized GPU annualized protocol revenue	~\$200M (early 2026)	DeFiLlama/Dune via press Error! Bookmark not defined.
H100 GPU-hour: hyperscalers	\$4.20–5.00	Rate cards, Q2 2026 Error! Bookmark not defined.
H100 GPU-hour: decentralized	\$1.40–1.60	io.net, Akash, Vast.ai, Q2 2026 Error! Bookmark not defined. Error! Bookmark not defined.
Bittensor Q1 2026 protocol revenue	\$43M	Sector reporting Error! Bookmark not defined.
Render annualized compute jobs	\$40M+ (Q4 2025)	Sector reporting Error! Bookmark not defined.

³⁶ SK hynix Newsroom. <https://news.skhynix.com/en/2026-market-outlook-focus-on-the-hbm-led-memory-supercycle/>

A.5 Theme-token category statistics

Metric	Value	Source
New meme/theme token launches	>36,000/day	Launchpad data Error! Bookmark not defined.
Category failure rate	~97%; 60% dead within 24h	CoinLaw via industry press Error! Bookmark not defined.
Meme sector market cap, 2026	~\$37.6B	DefiLlama narrative tracker Error! Bookmark not defined.
AI-token sector market cap / Q1 2026 performance	\$22.6B / only profitable sector	DefiLlama / Phemex Error! Bookmark not defined. Error! Bookmark not defined.

Appendix B: Proposed Default Parameters — Summary

All [Proposed] values in one place, for ratification before launch:

Parameter	[Proposed] value	Section
Chain / standard	Solana / SPL	4.1–4.2
Total supply	1,000,000,000 CHIP, fixed	4.1
Mint / freeze authority	Renounced at launch	4.1
Allocation	40 / 30 / 15 / 10 / 5 (liquidity / community / treasury / team / marketing)	4.3
Community vesting	24-month linear	4.4
Treasury vesting	36-month linear, DAO-gated spend	4.4
Team vesting	12-month cliff + 24-month linear	4.4
Marketing vesting	12-month linear	4.4
LP policy	Burn or ≥12-month lock, published pre-launch	4.5
Treasury control	≥3-of-5 multisig → DAO in Phase 4	4.1, 5.5
Launch venue	Reputable Solana launchpad → Raydium/Jupiter	4.2, 7.2
Restricted jurisdictions	[TBD] — external legal review	4.1, 8

Appendix C: Figure Index

Figure	Title	Type
Cover	ChipCoin — coin as microchip on silicon wafer	Brand artwork
1	Global semiconductor revenue, 2023–2027F	Data chart
2	The memory supercycle, 2025–2027F	Data chart
3	Hyperscaler capex, 2025 vs 2026E	Data chart
4	H100 GPU hourly pricing by provider, mid-2026	Data chart

Figure	Title	Type
5	Q1 2026 returns: AI-compute tokens vs majors	Data chart
6	CHIP token allocation, proposed model	Tokenomics chart
7	Proposed CHIP unlock schedule, 36 months	Tokenomics chart
8	Four-phase utility roadmap	Architecture diagram
9	CHIP ecosystem architecture	Architecture diagram
10	The Chip Watch narrative flywheel	Infographic

Chart data reflects the sources cited in each figure's caption; figures 6–8 depict [Proposed] parameters as described in Section 4. Brand artwork and diagrams are original to this document.

ChipCoin (CHIP) Whitepaper v1.0 — August 2026. Supersedes Draft v0.1. Contact: [TBD]. Website: [TBD]. This document is for general informational purposes only and does not constitute investment, legal, or tax advice.