

Machine-native money for a useful compute economy.

Bitcoin K is a proof-of-work BlockDAG settlement network with a fixed four-year issuance window and a planned market for independently operated AI inference.

STATUS – L1 mainnet is operating in a disclosed guardian phase through checkpoint #271,828. L2 inference, liquidity routing and delegated agent wallets are experimental specifications, not live products. This paper distinguishes consensus rules from roadmap systems.

[DOWNLOAD PDF](#)[MONETARY POLICY](#)[COMBO](#)[GUARDIAN LAUNCH](#)[GPU ECONOMY](#)[AGENT WALLETS](#)[STATUS MATRIX](#)

1. Design principles

VERIFIABLE SCARCITY

All subsidy issuance is determined in integer SatK by consensus.

SEPARATE COMPUTE LAYERS

ASICs secure L1. GPUs sell useful work on L2 and cannot mint BTCK.

MACHINE-NATIVE ACCESS

Agents are intended to buy individual services under explicit wallet policies.

ZERO GENESIS REWARD

Genesis creates no spendable BTCK and imports no DEVNET balance.

DISCLOSED GUARDIAN PHASE

Controlled early hashrate is disclosed rather than described as an open launch.

VERIFY, DO NOT TRUST

Network data, artifacts, hashes and limitations must remain independently checkable.

2. Monetary policy

PARAMETER	CANONICAL VALUE
Name / ticker	Bitcoin K / BTC-K (UI shorthand: BTCK)
Hard cap	21,000,000.00000000 BTC-K
Base unit	1 BTC-K = 100,000,000 SatK
Target block interval	1 second
Subsidy window	1,460 days / 126,144,000 rewarded blocks
Genesis reward	0 SatK
PoW / ordering	kHeavyHash / GHOSTDAG BlockDAG

The theoretical reward curve is $R(t) = R_0 / (1 + 0.15t)$. Consensus never evaluates that floating-point expression. It uses a frozen table of 1,460 integer daily quotas and distributes each quota across 86,400 block positions using quotient and remainder. The final quota is adjusted so cumulative issuance equals exactly 2,100,000,000,000,000 SatK.

```
Canonical daily-quota SHA-256:
7896cf36303e123f799c71b4d72b74a684e0079c1e32e57eac34d4ed04913e59

Block 1 reward:      629,815,765 SatK
Final subsidy block: #126,144,000
Subsidy afterward:   0 SatK
```

3.L1 allocation

YEAR	MINER	COMBO	ECOSYSTEM	COMBO PERIOD
1	80%	15%	5%	12h / 43,200 blocks
2	70%	25%	5%	36h / 129,600 blocks
3	60%	35%	5%	108h / 388,800 blocks
4	50%	45%	5%	324h / 1,166,400 blocks
After subsidy	50% of fees	50% of fees	No new subsidy	324h / 1,166,400 blocks

The ecosystem output is enforced in coinbase validation. A block that omits or redirects it is invalid. Changing the hard cap, quota table, allocation or permanent 50/50 fee rule requires an explicit consensus fork; no pool operator or administrator key can modify it at runtime.

4. Combo selection and maturity

Each Combo window records valid L1 candidate blocks. One hundred blocks after the window closes, the beacon is derived from the future consensus block:

```
beacon = SHA256(block_hash>window_end + 100] ||
epoch_index_uint64_le)
winner = argmin(candidate_block_hash XOR beacon)
```

Every valid block has one candidate entry. Splitting rewards across many addresses does not create additional chances. Long-run probability remains proportional to valid block production.

YEAR	REGULAR MATURITY	COMBO MATURITY	INITIAL COMBO RELEASE
1	2 hours	48 hours	50%
2	6 hours	144 hours	25%
3	24 hours	432 hours	10%
4	72 hours	1,296 hours	10%

Pool anti-hopping uses persistent accepted-share records to decide eligibility for a Combo received by that pool. Private shares are never L1 consensus inputs. Solo blocks remain attributable by their on-chain coinbase output. The official pool requires at least 95% window presence and applies year-specific outage limits; those are pool distribution rules, not BlockDAG consensus.

5. Disclosed guardian launch

The first 271,828 blocks are mined under controlled project hashrate. This reduces cold-start exposure to immediate majority-hashrate reorganization on an algorithm with a large existing

hardware base. It is not described as an open-from-genesis fair launch.

- Genesis reward is zero and no DEVNET balance is imported.
- Every guardian coin is created by valid kHeavyHash work under the same year-one 80/15/5 rules.
- The hash of block #271,828 will be frozen into public node releases.
- The controlled hashrate, receiving outputs and reserve governance must be disclosed before public ports open.

6. Security after issuance

Blocks continue after the four-year subsidy ends. From block #126,144,001, actual transaction fees are divided permanently: 50% to the immediate block miner and 50% to the permanent Combo reserve. There is no tail emission and no renewed ecosystem subsidy.

6.1 The primary mission: an open Agent economy

Bitcoin K is not intended to stop at issuing a scarce proof-of-work asset. L1 is the neutral settlement and security foundation; the project's principal economic expansion is L2: a competitive market in which independent GPU operators load useful models and autonomous agents purchase inference one job at a time.

This changes the customer from a human clicking a subscription page into software that can discover a model, compare a signed quote, enforce its own budget, pay, verify delivery and continue working. A user who cannot pass a platform's account, region, card or KYC gate should still be able to fund a local agent and buy a narrowly scoped service without surrendering the agent's wallet keys.

L1: SECURITY

ASIC miners order transactions and make SatK settlement expensive to rewrite.

L2: PRODUCTIVE SUPPLY

GPU providers turn idle or dedicated hardware into model endpoints with measurable service quality.

AGENTS: DEMAND

Software purchases only the intelligence it needs under explicit owner policies.

6.2 One Agent call, end to end

01 Discover

The agent requests offers by model family, context size, jurisdiction, latency, privacy policy and maximum price.

02 Quote

Providers publish a KCU rate. The gateway converts task KCU through the disclosed BTC-K reference price, records mSatK quote precision and issues an integer-SatK invoice.

03 Authorize

The local policy wallet verifies service scope, per-call ceiling, daily budget and delegation expiry. Default is deny.

04 Challenge and pay

An HTTP 402 invoice binds amount, destination, nonce and expiry. The wallet signs locally; neither gateway nor provider receives a recovery phrase.

05

Route privately

The encrypted request is sent to the selected provider. A future relay profile can prevent the provider learning the client's home IP.

06

Prove delivery

The provider returns the result and a signed receipt containing job identity, model commitment, timing and metering evidence.

07

Settle and distribute

Paid SatK is divided by the active L2 schedule among the GPU provider, standby capacity, AI Combo, burn and ecosystem support.

This is the target production architecture. HTTP 402 settlement, provider discovery and DEX/liquidity routing are not live on mainnet today.

6.3 Private robot wallets


A robot wallet is not a hot wallet with unlimited authority. It is a constrained execution wallet controlled by a human or organization. The principal signs a delegation pact; the machine can spend only inside that pact. A delivery robot may buy charging but not exchange assets. A research agent may buy an approved model call but not transfer its entire balance.

THREAT	BTCK WALLET CONTROL	PRIVACY EFFECT
Robot is stolen	Non-exportable hardware key; revocable, expiring delegation	Master spending key is absent from application storage

THREAT	BTCK WALLET CONTROL	PRIVACY EFFECT
Model hallucinates a payment	Default deny, allowlist, per-call and rolling limits	Compromise is bounded by policy, not total balance
Merchant correlates purchases	Dual-key stealth recipient and one-time outputs	Repeated payments do not reuse a visible destination output
Chain observer follows the robot	Client scanning with the private view key	Public address alone cannot enumerate owned outputs
Network observer links location	Planned relay and Tor-compatible transport	Stealth addresses alone do not hide IP metadata
Gateway fakes settlement	Signed prepared receipt plus independent chain verification	A server cannot redefine pending as settled

The public stealth address contains only public spend and public view material. The private view key remains local and is required to recognize owned one-time outputs; the private spend key remains offline or inside an enclave. Human wallets, autonomous software and embodied robots can therefore share one settlement protocol without sharing custody.

6.4 The L2 economic flywheel



For every 100 SatK paid in year one, the proposed L2 schedule sends 60 SatK to the serving GPU worker, 10 SatK to standby capacity, 15 SatK to AI Combo, burns 5 SatK and assigns 10 SatK to ecosystem development. The proportions evolve toward larger Combo incentives and stronger

burn while ecosystem intervention declines to zero in year four. These are transfers of client-paid SatK; they never create new coins.

100 SatK SERVICE PAYMENT	YEAR 1	YEAR 2	YEAR 3	YEAR 4
Serving GPU provider	60	50	40	35
Ready-capacity pool	10	10	10	5
AI Combo	15	25	30	40
Provable burn	5	10	15	20
Ecosystem and open tooling	10	5	5	0

6.5 Cost advantage: a measurable target, not a slogan

Hosted APIs bundle research, global infrastructure, support, safety systems and margin. Bitcoin K does not claim that every open model will underprice every subsidized cloud model. Its structural opportunity is narrower and testable: let many operators bid unused hardware directly, eliminate mandatory subscriptions, settle per successful job and expose the full fee split.

REFERENCE ROUTE	PUBLIC LIST PRICE	ACCOUNT MODEL	BTCK RELEVANCE
OpenAI chat-latest	\$5 input / \$30 output per 1M tokens	Hosted account and billing	Premium hosted reference
Anthropic Claude Sonnet 5	\$2 input / \$10 output per 1M tokens	Hosted account and usage billing	Agent-quality hosted reference
Google Gemini 3.5 Flash-Lite	\$0.30 input / \$2.50 output per 1M tokens	Hosted free and paid tiers	Low-cost benchmark BTCK must compete against honestly

REFERENCE ROUTE	PUBLIC LIST PRICE	ACCOUNT MODEL	BTCK RELEVANCE
RunPod RTX 4090	\$0.74 per GPU-hour	User operates model and infrastructure	Public hardware-cost reference
Bitcoin K L2 target	Open KCU rate converted to SatK per job	No subscription; policy-wallet payment	Provider competition and transparent distribution

Reference prices were checked on 8 October 2026 from [OpenAI](#), [Anthropic](#), [Google](#) and [RunPod](#). Models are not equivalent; this is a market reference, not a quality-adjusted claim.

Provider raw output cost per 1M tokens = GPU hourly cost ÷ (output tokens/sec × 3,600 × utilization) × 1,000,000

Task mSatK = round(Task KCU × 10,000,000 ÷ BTC-K reference price)

Client quote = provider KCU rate + standby + AI Combo + burn + ecosystem allocation

KCU is the stable service-value bridge: 1 KCU = \$0.00010 of reference service value. A standard 8B route starts at 0.8 KCU / 1,000 weighted units . At the explicitly fixed bootstrap reference of \$0.05/BTC-K , 0.8 KCU converts to 160,000 SatK. If the reviewed reference price rises, the SatK amount falls; if it falls, the SatK amount rises. The bootstrap value is not represented as a live market price.

The launch objective for standard open-weight routes is an all-in median client quote between 20% and 33% of a genuinely comparable hosted route, subject to an electricity and hardware-cost floor. A public beta must publish model hash, prompt set, quality score, latency, throughput, hardware, utilization and every SatK allocation. If the measured route is not cheaper at comparable quality, the website must show that result rather than manufacture a discount.

IMPLEMENTATION STATUS – The disabled gateway prototype now derives its allocation year from canonical L1 height, verifies confirmed invoice-specific deposits, creates restart-safe withdrawals, sends the burn share to a frozen NUMS output, enforces

integer SatK conservation and authenticates GPU workers and results. The public gateway can operate with only the Treasury public address; spending keys remain outside that process. Automated fault injection covers RPC failure, wrong-chain data, shallow confirmation, reorg/eviction, duplicate broadcast and durable-ledger failure. Initial pricing uses a versioned KCU catalogue and an explicitly fixed \$0.05/BTC-K bootstrap reference, not a live market oracle. A dynamic reference requires audited multi-source median/TWAP, stale-data rejection, movement limits and an emergency stop. Production remains disabled pending a MAINNET small-value rehearsal, signed operating limits and an independent external security audit.

7. Two-layer compute economy

L1 and L2 are deliberately decoupled. Deterministic ASIC work establishes ordering and settlement. GPU inference is commercial work paid from already circulating SatK. A failed GPU job cannot change L1 consensus, and useful work never authorizes new issuance.

LAYER	PROOF / RECEIPT	VALUE SOURCE	STATUS
L1 ASIC	kHeavyHash block proof	Programmed subsidy and fees	Guardian mainnet live
L2 GPU	Job receipt, readiness and benchmark evidence	Client-paid SatK only	Experimental specification

Proposed L2 service-fee allocation

YEAR	MINER	STANDBY	AI COMBO	BURN	ECOSYSTEM
1	60%	10%	15%	5%	10%
2	50%	10%	25%	10%	5%
3	40%	10%	30%	15%	5%

YEAR	MINER	STANDBY	AI COMBO	BURN	ECOSYSTEM
4	35%	5%	40%	20%	0%

This table describes distribution of real L2 service payments. It is not current L1 issuance and is not active on mainnet.

8. GPU readiness and ecosystem reserve

The planned provider market uses unpredictable inference probes, job receipts and observable quality. Candidate signals include model identity, weight hash, quantization, time-to-first-token, throughput, completion rate, hardware class and quoted price. Slashing and reputation require a public testnet and adversarial review before production.

The 5% L1 ecosystem output is intended to support security review, open tooling and carefully disclosed GPU cold-start programs. No fixed 500-BTCK daily subsidy is active today. Any future program requires a published budget, objective recipient rules, reserve balance reporting and on-chain records. The protocol promises neither fiat returns nor guaranteed provider income.

9. Machine payments

The experimental HTTP 402 profile describes single-use invoices denominated in integer SatK. A challenge binds an invoice identifier, amount, destination, nonce and expiry. Payment proofs use canonical digests, signatures and idempotency keys. Direct balance mutation endpoints are disabled in production profiles.

```
WWW-Authenticate: L402 invoice="...", currency="SatK", amount="..."
X-Payment-Address: btck1s...
```

```
X-Payment-Network: BTCK-BlockDAG
X-Payment-Expires: unix_timestamp
```

The prototype records `1 SatK = 1,000 mSatK` quote precision but rounds the current invoice upward to integer SatK. On-chain consensus and the current settlement ledger remain denominated in integer SatK. Persistent sub-SatK accumulation is not enabled until its accounting path is independently audited.

10. Policy-bound agent wallets

BTC-K and SatK are two displays of the same L1 asset and use the same `btck1s...` wallet: `1 BTC-K = 100,000,000 SatK`. There is no separate L2 coin or second recovery phrase. The wallet can toggle the displayed unit without converting or moving funds. L2 balances below 10,000 SatK accumulate off-chain; withdrawals begin at 10,000 SatK and the network fee is paid from a separate system fee reserve.

The delegated wallet profile is designed for software and physical robots. A principal authorizes a narrowly scoped agent identity with allowed services, a maximum per payment, a rolling budget and an expiry. Default behavior is deny. Spending keys remain client-side or inside a hardware enclave.

- A server must never request a recovery phrase or private spend key.
- A prepared signature is not represented as settled until independently verified.
- Stealth outputs reduce address linkage; they do not hide IP metadata.
- Production use remains blocked pending wallet review and end-to-end settlement tests.

11. Implementation status

COMPONENT	STATUS	PUBLIC CLAIM
kHeavyHash, GHOSTDAG, adaptive DAA	Implemented	Guardian mainnet live
Integer emission and L1 allocation	Implemented	Consensus active
Pool accounting and dynamic VarDiff	Implemented and capacity-tested	Access pending checkpoint
Dual-key stealth wallet	Implemented, not independently audited	Guardian release only
Future-beacon Combo and anti-hopping	Implemented; evidence accumulating	Verify on chain and pool ledger
GPU market / readiness / slashing	Prototype and specification	Not live
HTTP 402 AI inference	Prototype	Not live
Delegated robot wallet	Prototype	Not production-ready
DEX, bridge, exchange adapters	Disabled	Not live

12. Risks and verification

Bitcoin K is experimental software. Proof-of-work networks face software defects, hashrate concentration, economic attacks, operational failure and regulatory uncertainty. Stealth addresses are not a substitute for a complete privacy audit. Token ownership does not promise profit, yield or future service access.

Public evidence should include source and release tags, reproducible build instructions, artifact hashes, an SBOM, genesis parameters, checkpoint #271,828, known limitations, reserve addresses and independent security findings.

NORMATIVE AUTHORITY – Consensus behavior is defined by released node code and canonical v2.4.1 integer tables. If descriptive prose conflicts with validated consensus code, the discrepancy must be disclosed and resolved through an explicit versioned release.