Making Core Irrelevant

Josh · Secure Sovereign

Published August 22, 2026

What has to happen, and in what order

Contents

The diagnosis of the implementation monopoly and the no-spec moat is in *Who Controls Bitcoin* and the *Argument Map*. The verification work is in *Why Bitcoin Needs a Specification*. This article is what comes next. What has to happen, in what order, and what failure looks like.

I. The Problem Statement

Bitcoin does not have a governance problem in the sense most people mean when they use that phrase. It has an engineering problem. One implementation controls the network. That implementation's maintainers, funded by a small cluster of institutions with aligned interests, decide what gets merged. That implementation's undocumented behavior is the spec every alternative has to chase. That is a single point of failure dressed up as decentralization.

The numbers are public. Sixteen years of Bitcoin Core git history show a PR-weighted Gini of about 0.851. The top three merger accounts have executed roughly 80% or more of historical merges. In 2025 alone, 56% of all merges flowed through one person funded by Brink (Engineering Impact

Report). Five people still hold merge access. Core development drew $8.4 million in 2023 against a two trillion dollar market cap. The tables are in Bitcoin Governance Research and *Who Controls Bitcoin*.

None of this requires a conspiracy. People fund what they understand, hire from networks they trust, and select for demonstrated alignment. The capture is structural, not coordinated, which makes it harder to argue against and harder to fix. You cannot swap out maintainers and solve the problem. The problem is the monoculture. What keeps the monoculture in place is the absence of a formal consensus specification. That mechanism is the no-spec moat.

Without a spec, any alternative has to reverse-engineer undocumented behavior from Core itself. That ties every alternative to Core by definition. The no-spec moat is not an oversight. It is what makes the monopoly rebuild itself every time someone tries to compete. The moat dissolves when a credible formal specification exists, and not before.

II. The Actual Terrain

As of mid-2026, four pools controlled over 70% of Bitcoin's roughly 929 EH/s. Foundry held about 31%, AntPool 18%, ViaBTC 13%, and F2Pool 10% (source: mempool.space, D-Central H1 2026 Report). Three pools are enough to exceed half of all blocks. Those shares move week to week. The structure does not. A single pool in the high twenties is already large enough that a further swing matters.

The concentration that matters more is who builds the block, not who hashes it. Under Stratum V1, which still carries most of the network, the pool picks the transactions. Hashrate numbers describe who mines blocks. Template control describes who decides what goes in them. Those are different concentrations, and the second one is the censorship risk. Every miner pointing hashrate at Foundry, AntPool, F2Pool, or ViaBTC is handing transaction selection to a US or Chinese entity with its own regulators and its own business.

Stratum V2 with Job Declaration lets a miner build the template while staying in a pool so payouts stay even. The first live Job Declaration block was mined June 25, 2026 (block 955,318, DMND pool, GoMining template). Seven major pools, about 75% of global hashrate, joined the Stratum V2 Working Group in May 2026. Joining the working group is not the same as turning Job Declaration on. As of mid-2026, only DMND and Braiins run V2 natively in production. Ocean got the same result through DATUM before its leadership situation deteriorated. Most of the network is still on V1 with the pool picking the transactions.

The nodes tell the same story. Only about 16% of nodes are running Core v30 with the higher OP_RETURN limits, which is the version Core actually shipped (source: Clark Moody dashboard, May 2026). More nodes are running pre-v30 software or Knots than are running Core's own

controversial change. Knots peaked at roughly 25% of public nodes in September 2025 and has settled around 20% as of mid-2026 (source: Coin Dance). That is demand for a policy alternative. It is not implementation diversity. Knots is a Core-derived client on the same upstream lineage, which is why *Who Controls Bitcoin* treats reachable-node splits as demand, not as a break in the monoculture. The missing piece is not demand. It is a separately written client that an exchange or payment processor can actually run, and a clear answer to who is liable if that client splits the chain.

DCG owns Foundry, the largest pool, and also owns Grayscale, the largest Bitcoin asset manager. Fortitude Mining is trying to go public under that same roof. There is no documented funding tie between DCG and the organizations that pay Core developers. There does not have to be. Hashrate, asset management, and mining infrastructure already sit in one parent company.

III. Why Miner-First Strategies Fail

Miner campaigns without the entities that actually hold value show preference. They do not move the protocol. Protocol changes stick when the economic majority enforces them, because that is when refusing costs money. SegWit succeeded in 2017 not because miners wanted it but because economic nodes made non-compliance dangerous. A hashrate campaign that never gets that pressure fails even when the technical case is sound. People read those failures as hashrate problems. They are economic majority problems.

Hashrate matters. Pool decentralization matters. Neither moves the protocol unless economically significant nodes adopt the change. Push miners first and you get a preference poll, not a rule change. Implementation diversity is the precondition. Everything else waits on it. Consensus is the small part of a node. Close to 90% is not consensus. Wallet, policy, networking, index, RPC, and mempool sit in that remainder. A Core-derived client can change what it forwards. It cannot rewrite that surface. A separately written client can, but only if it can prove it still enforces the same money rules. That proof is what the spec makes possible. That is why the moat dissolves when the spec exists.

IV. The Two Tracks and Why Both Are Required

There are two problems. The order between them is the strategy.

The implementation problem is a specification problem first. Without a formal consensus specification, an alternative cannot prove it enforces the same rules as Core. Every serious alternative in Bitcoin's history has eventually been found to diverge. btcd accumulated documented consensus bugs across 2022-2024 after years as a clean-room Go client. Each bug was a case where the two programs would have accepted or rejected blocks differently. The point

of a second implementation is to catch exactly those gaps. btcd kept being the thing that diverged. That record is in *Governance Paralysis Was The Victory*. What a genuine alternative actually requires is in *Argument Map, Part XV*.

Satoshi warned against a second compatible implementation and called it a menace to the network. Take that warning seriously. The risk he named was a quiet consensus mismatch that splits the chain. Formal verification alone does not close that risk. You need formal verification, fuzzing, differential testing, and property tests. That stack is the engineering answer, not a dismissal of the warning.

Bitcoin Commons is the leading attempt at that path. It is a ground-up Rust node, written separately from Core, not copied from it.

The Orange Paper is the document that node is built from. It is Bitcoin's consensus rules written as mathematics a person can read: signature checks, UTXO accounting, script, and block structure, defined as objects with precise rules instead of as whatever Bitcoin Core happens to do. Any implementation can check itself against that document. That is why the Orange Paper is a public good even if Commons is slow. The no-spec moat dissolves when the spec exists, for the whole ecosystem, not only for Bitcoin Commons.

The BLVM consensus layer is the code that implements those rules. It is a bounded set of deterministic functions. Wallet, networking, and policy are not mixed into it. That isolation is what makes a proof about the money rules tractable.

The spec lock is the check that runs on every merge. It uses Z3, a theorem prover, to ask whether the Rust still satisfies the contracts taken from the Orange Paper. The spec does not write the code. People write the code. If a change would break a consensus property, the merge fails.

Differential testing then runs that code against Bitcoin Core across more than 900,000 blocks. Zero consensus divergence so far (source: docs.thebitcoincommons.org). No other project has published a comparable mix of a readable spec, a locked implementation, and full-chain differential testing. The usual objections about proofs and tests are answered in *Why Bitcoin Needs a Specification*.

Bitcoin Commons is not production-ready as of mid-2026. Governance is not activated. The software carries explicit unreleased warnings. What exists is that foundation. Getting an exchange or payment processor to run it still takes independent security audits, extended mainnet operation, operational hardening, and legal review. That work is ahead, not behind. First institutional adoption is not late 2027 on an optimistic calendar. It is more likely 2028-2029.

The first question a board will ask is liability. If a formally verified alternative and Core disagree, and the chain splits, who pays? Formal verification lowers the odds of that split. It does not tell counsel who is on the hook if it happens anyway. The first institution to run non-Core software

needs that answer in full, not a smaller probability. The likely package is an independent audit, insurance, a staged rollout with a fallback to Core, and a contract that names the split case. None of that exists yet, and all of it takes time.

The mining problem is simpler in one way. The product already worked. Ocean showed that miners will move for TIDES-style honest payouts, non-custodial architecture, and miner template control. What failed was the organization. A pool run by one founding figure, with that person also holding reserve management, transaction policy, and protocol advocacy, is a single point of failure. It failed the way that structure fails.

The successor needs the same product and a different organization. Reserve capital has to sit apart from day-to-day operators. Protocol advocacy has to sit apart from pool operations. The bar is governance that survives the removal or conduct of any one person. Ocean's collapse is the list of what not to repeat.

DMND offers V2-native job declaration in production. Braiins has run V2 with job declaration since before the working group existed. Neither fills Ocean's niche. DMND does not have the transparent pro-rata payout model. Braiins is not non-custodial. As of August 2026, nothing else combines non-custodial payouts, honest accounting, and miner template control.

The two tracks only work together. A verified alternative gives exchanges and payment processors an answer when the board asks why they are not running Core. Without that answer, a more decentralized pool is an ideology that institutions will not touch. With it, economically significant nodes can run non-Core software, and miner campaigns can actually move the protocol. Parallel work is what produces force that neither track produces alone.

V. The Scenario Map

Three scenarios cover the next three to five years.

Floor. Bitcoin Commons or a comparable formal-spec implementation is running on at least one economically significant node by 2029. A successor pool holds 3-5% of hashrate before the April 2028 halving. Stratum V2 Job Declaration reaches 10-15% of hashrate by the end of 2027, mostly because new ASIC firmware starts shipping it on. Core still dominates, but the claim that no credible alternative exists is gone. Making Core irrelevant is the destination. The floor is just the first proof that the alternative can be run.

Mid. One or two exchanges or payment processors are running a formally verified alternative by 2029. A successor pool holds 8-12% of hashrate. Implementation diversity then has real economic weight, pool share has real weight on what gets built into blocks, and Core's hold on protocol decisions is constrained for the first time.

Ceiling. A genuine multi-implementation Bitcoin by 2030-2031. Several exchanges running formally verified alternatives. No single pool near 25% of hashrate. Stratum V2 Job Declaration the default on new ASIC firmware. That is a decade of work, not a three-to-five year one.

Key Uncertainties

Five open questions decide which scenario shows up, and when.

How fast institutions adopt. A verified implementation that no exchange runs produces nothing. Selling node software into an exchange or payment processor takes quarters, not weeks. It needs an independent audit, legal review, insurance for a split, and board sign-off. The liability question lives inside that cycle. Formal verification lowers the odds of a split. It does not tell a board who pays if one happens. The software has to be ready before that cycle starts, and the cycle has to start well before the next protocol fight.

Cash for unlucky streaks. Most people understate what a successor pool needs in reserve. Ocean's $6.2 million seed covered launch, not running at scale after the halving when margins are thinner. Low-to-mid tens of millions is the realistic reserve for a pool that pays TIDES-style pro-rata and cannot default after a dry run of blocks.

When firmware actually ships Job Declaration. That depends on Foundry and AntPool turning the feature on, not on joining a working group. The large pools have a reason to take V2's encryption and efficiency and leave template control where it is. Watch for Foundry turning Job Declaration on. Working group membership is not that signal.

Whether Core writes a spec of its own. A credible formal-spec alternative with institutional backing could push Core-funded work to write more of the consensus rules down. If Core or a Core-funded organization publishes a competing spec, the ground changes. The no-spec moat dissolves no matter who dissolves it, which is good for Bitcoin. It also takes away the advantage of having the spec first.

Where Ocean's hashrate and Tether went. Tether committed hashrate to Ocean in April 2025 and has said nothing public about where that hashrate went after the August 2026 leadership situation. Do not count that volume for a successor until someone documents it.

VI. The Timeline and External Variables

The April 2028 halving is the largest external date. Block subsidy drops from 3.125 to 1.5625 BTC. Daily issuance falls from roughly 450 to roughly 225 BTC. Post-halving production cost lower bounds are modeled near $93,000 against mid-2026 break-even estimates near $65,000 (source: D-Central H1 2026). About 20% of miners are already unprofitable at current prices and difficulty. The efficient machines are hydro-cooled S23-class hardware at about 9.5 J/TH. Anything much worse than that is under shutdown pressure.

History is against a new small pool after a halving. Smaller pools have lost share every previous time, because the capital bar went up and only the well-reserved operators survived. That pattern holds unless this time is different, and it is, in one way. FPPS insurance gets more expensive when margins compress and luck hits harder. The gap between a guaranteed daily payout and what the blocks actually produce is not free. The pool eats it. When block rewards are smaller, that gap matters more to a miner on the edge. A TIDES-style model without that insurance overhead pays better expected returns to miners who will take variance. That relative advantage is widest in the 12-18 months after the halving, which means a successor has to be running and trusted before the halving, not after it.

F2Pool has a documented history of filtering. Foundry's compliance posture is institutional by design. Non-custodial template control gets more valuable as US-domiciled pools filter more, because it is the architecture that does not depend on the operator staying brave.

Implementation diversity is not only an engineering problem. Core's maintainers and funders have a tight network and a loud story. An alternative needs its own money, a review culture people trust, and a way not to be framed as an attack on Bitcoin. The capture critique is correct, and it is also the story Core's defenders will use against any alternative. The answer is not a louder governance argument. It is a running alternative with a clean audit trail and no chain splits.

VII. What Actually Has to Happen

Implementation work comes first. Nothing else has force without it.

The formal-spec path needs a release an exchange or payment processor can run. That means an independent security audit, extended mainnet operation without incident, operational hardening, and a prepared answer on liability before the first institutional meeting. The pitch is not a philosophy. It is 900,000 blocks of differential testing with zero divergence, the spec lock that checks the code against the Orange Paper, the fuzzing and property-test results, the audit, the insurance, and why this lowers their risk instead of raising it. That is a board conversation. Finding the people who can have that conversation is a different job from writing the software, and both have to run at once.

Do not start with the largest exchanges. Their legal overhead makes them the slowest. Start with mid-size exchanges and payment processors that have technical leadership, a reason to care about sovereignty, and a reason to stop waiting on Core's release cycle. The first one or two matter because they show that institutions can move.

On mining, the question is who can rebuild what Ocean built, with a better organization. The product is specific. TIDES-style transparent pro-rata payouts, non-custodial architecture, miner template control through DATUM or the equivalent, and a structure that keeps reserve capital, operations, and protocol advocacy in separate hands. The reserve needed to survive unlucky streaks at real scale is low-to-mid tens of millions. Ocean's $6.2 million seed covered launch, not that. A successor needs more capital, and an organization that survives the departure or controversy of any one founder.

VIII. What Failure Looks Like

It can fail in four places.

It fails on implementation if the verification stack exists and no economically significant node adopts the alternative before the next protocol fight. A clean implementation with no exchange running it has no weight. The governance critique can be right, the tests can be clean, and none of it matters if the economic majority never enforced anything.

It fails on mining if the Ocean successor is not standing before the April 2028 halving. After the halving, large institutional pools with deep reserves take share. The window for a non-custodial pool to attract enough miners to matter is widest before the halving raises switching costs.

It fails as an organization if the successor repeats Ocean. The first serious protocol fight the pool takes a side in will test whether reserve governance, operations, and advocacy are actually separate. If they are not, the pool splits at the moment it needs to hold.

It fails on the pitch if the campaign leads with ideology. The people who control economically significant nodes have boards, lawyers, and fiduciary duties. The argument that wins is risk management, with the verification stack and a direct answer on liability. Lead with Bitcoin philosophy and they will treat you as someone whose job they do not have.

IX. Conclusion

The implementation monopoly is a solvable engineering problem. The Orange Paper exists. The BLVM spec lock is running. Differential testing across 900,000 blocks has produced zero consensus divergence. Formal verification, fuzzing, differential testing, and property tests are what make a second client safe enough to run. That foundation has not existed before in Bitcoin. What remains is adoption, the liability answer, the audit, and the time those things take inside institutions.

The mining problem has a product that already worked, a market that already showed up, and a hole where the Ocean successor should be. That hole will not stay open. The halving closes it from one side. Regulatory pressure on US-domiciled pools opens it from the other.

The path is two tracks at once. Implementation diversity gives institutions an answer they can defend. Pools that do not hold keys and let miners build their own templates give miners a reason to move that is about money, not politics. When both exist, miner pressure on protocol questions has economic weight again for the first time since the blocksize war. The work is to get both in place before the next major protocol fight.

Related: Who Controls Bitcoin, Why Bitcoin Needs a Specification, Governance Paralysis Was The Victory.

Sources

- Bitcoin Governance Research. PR-weighted Gini, merge concentration, stalled-proposal dossiers
- Brink Engineering Impact Report 2025, March 26, 2026. 56% merge concentration
- mempool.space mining pools. Hashrate share
- D-Central, The State of Bitcoin Mining, H1 2026. Network hashrate, post-halving cost bounds
- Clark Moody Bitcoin dashboard. Core v30 node share
- Coin Dance. Knots node share
- Stratum V2 Working Group announcements, May 2026. Pool membership vs Job Declaration activation
- ocean.xyz. TIDES payouts, DATUM template control
- TFTC, first Stratum V2 Job Declaration block, June 25, 2026. Block 955,318
- Bitcoin Commons documentation. Orange Paper, BLVM spec lock, differential testing record
- Orange Paper / Bitcoin Commons consensus spec

SECURESOVEREIGN

Markdown is the canonical source: https://secsov.com/articles/making-core-irrelevant/index.md

Formatted snapshot of https://secsov.com/articles/making-core-irrelevant

Josh · Secure Sovereign