Governance Paralysis Was The Victory

Josh · Secure Sovereign

Published February 2, 2026 · Updated September 20, 2026

Bitcoin is not captured. The implementation that writes the defaults is capturable.

Contents

Bitcoin is not captured. The implementation that writes the defaults is capturable.

Those are not the same thing. Bitcoin is a network, a set of consensus rules, and a fixed supply of 21 million coins secured by proof of work. Bitcoin Core is one software implementation, the only one anyone runs in meaningful volume, that enforces those rules. **Bitcoin Core's merge process is informal, but Bitcoin's properties are not broken.**

The problem is not Bitcoin. The problem is that Bitcoin runs on one dominant implementation, written by a small merge roster, and that single point of failure is capturable. The solution is not to leave Bitcoin. **The solution is to build a second implementation** that runs on Bitcoin's network, enforces Bitcoin's consensus rules, and can be checked against a written specification. Same chain, same 21 million. That is what implementation diversity on Bitcoin actually means. Everything else is abandoning the only asset that actually works. For the funding map, merge concentration, and adversarial cases, see *Who Controls Bitcoin*. For why the block size war consumed the bandwidth that diversity required, see *The Social Layer Is the Attack Surface, §V*. For the stalled-proposals evidence list and common monopoly defenses, see *What Bitcoin's Stalled Proposals Tell You* and *Bitcoin Core: The Biggest Fallacies*.

The Block Size War

In the early days of Bitcoin, Gavin Andresen, Satoshi's handpicked successor, called for multiple independent implementations of the Bitcoin protocol. The reasoning was straightforward: if one implementation has a consensus bug, others catch it, so you don't run a monetary network on a single codebase. Andresen set a threshold for when this work should be done, but that threshold was never met. The conversation moved on, and Bitcoin remained a single-implementation network, every node running variations of the same code and all sharing the same blind spots.

At the time Bitcoin was a $6 billion asset and the engineering problem was solvable. The conversation moved on, and Bitcoin remained a single-implementation network.

In 2015, the block size debate consumed Bitcoin's governance. On one side, increase block size to allow more transactions per block, enabling Bitcoin to scale as a payment network. On the other, keep blocks small to preserve decentralization, the property that actually makes Bitcoin valuable. This framing made it feel like a technical disagreement, but it wasn't. **It was a resource war.**

Exchanges needed throughput because every transaction on Bitcoin's network drove trading volume and trading volume drove revenue. Brian Armstrong, co-founder and CEO of Coinbase, was one of the most vocal advocates for block size increases. He publicly experimented with BitcoinXT and framed scaling as an election: "What's happening right now is an election in the bitcoin space." He pushed Coinbase to actively support scaling proposals behind the scenes. Armstrong's position didn't prevail and the early idealists held, but the damage wasn't in losing the block size argument. The damage was in what the argument consumed.

It is worth acknowledging that the small-block position was not purely defensive. Greg Maxwell and others argued, correctly, that keeping blocks small preserved the decentralization that made Bitcoin's security model work. Full nodes are the backbone of Bitcoin's trustless verification, and every kilobyte added to block size raises the hardware and bandwidth requirements to run one. If running a node becomes expensive, verification concentrates in the hands of exchanges and institutions, the exact entities whose interests don't align with Bitcoin's original design. The small-block position protected something real. **The problem was not that it won. The problem was that the war to win it consumed everything else.**

Bitcoin Cash forked off in 2017, but it wasn't a solution. Bitcoin Cash left Bitcoin's network, left Bitcoin's hash power, left Bitcoin's adoption, and carried the exact same single-implementation structure into a smaller, less secure chain. Its block size increase didn't fix governance. It just moved the same problem to a network that fewer people use and fewer miners secure.

CVE-2018-17144: The Bug That Proved Gavin Right

On September 17, 2018, a developer known as Awemany responsibly disclosed a vulnerability to Bitcoin Core developers. It was initially presented as a denial-of-service bug, but within hours, Bitcoin Core contributor Matt Corallo determined it was also an inflation vulnerability. A miner could create Bitcoin out of thin air by spending the same inputs multiple times within a single transaction.

The bug was introduced in Bitcoin Core 0.14.0 in March 2017 and had been present in the codebase for over a year and a half before anyone found it. It was created by an optimization that removed a costly duplicate-input check during block validation. The duplicate-input path was never mined. Replaying the live chain would have matched Core. The find was in the source. A second client would have seen it on a crafted transaction, not on those blocks.

Bitcoin Core developers chose to initially patch and release only the denial-of-service fix, deliberately withholding full disclosure of the inflation vulnerability to encourage upgrades before the attack surface was public knowledge. Eight months after the patch, more than half of Bitcoin's full nodes were still running vulnerable software. Luke Dashjr, Bitcoin Core developer, stated plainly: "The inflation bug is in practice a network-wide risk. It would allow a 51% miner attack to cause inflation, something such attacks can't normally do."

The bug was in Bitcoin Core, not in Bitcoin, and Bitcoin Core's monopoly is what made the bug potentially catastrophic for Bitcoin.

Who paid the seats

The block size war coincided with a period when Bitcoin's core development was funded through MIT's Digital Currency Initiative. On April 25, 2015, MIT Media Lab Director Joichi Ito emailed Jeffrey Epstein a primary source document later released by the House Oversight Committee. Ito forwarded his own internal message describing what had just happened:

"The way that Bitcoin is organized currently is that there are five core developers and around a hundred contributors to the core code. The five core developers are like Linus Torvalds of Linux. They decide what changes are made to the core code."

"Gavin, Wladimir and Cory were being paid out of a non-profit organization called the Bitcoin Foundation. A few weeks ago, it blew up when one of the board members declared the foundation bankrupt. Many organizations scrambled to step into the vacuum created by the foundation and take control of the developers. We moved quickly talking to all of the various stakeholders and the three developers decided to join the Media Lab. This is a big win for us."

Then Ito's note directly to Epstein: "Used gift funds to underwrite this which allowed us to move quickly and win this round. Thanks."

The language in Ito's email is not charity. "Take control." "Win this round." He is reporting that he has placed Bitcoin Core development inside an institution during the foundation's collapse.

Epstein's reply was "gavin is clever." Between 2002 and 2017, Epstein donated $850,000 to MIT. A portion of this funded the DCI, which employed Bitcoin Core's lead maintainer Wladimir van der Laan, Cory Fields, and Gavin Andresen. These were the three most important developers working on the protocol during the block size war, and they had no knowledge of the funding source. MIT deliberately concealed it, marking donations as anonymous to bypass vetting.

In July 2014, Peter Thiel emailed Epstein asking: "Do you think this is the first step in upping the anti-BTC pressure?" Epstein's reply mapped Bitcoin's internal contradictions, store of value versus currency versus property versus payment system, and the tension between anonymity and a transparent ledger. By October 2016, in the middle of the war, Epstein was emailing associates that he had "spoken to some of the founders of bitcoin who are very excited." He met with Brock Pierce and Larry Summers to discuss Bitcoin before it became mainstream.

Those are documented financial relationships, not a claim that the developers in the seats were conspirators. The war still consumed the bandwidth Gavin had said a second independently specified client required.

The PTSD: Good Ossification vs. Bad Ossification

The block size war ended and Bitcoin Cash forked off. The community "won" in the sense that block size didn't increase, but the war left a scar on how Bitcoin Core handles any change at all.

Bitcoin Core developers argue this conservatism is deliberate and necessary. The consequences of a consensus bug in a trillion-dollar asset are devastating, and extreme caution on consensus rule changes is correct. But none of the blocked improvements touch consensus rules. Separating wallet functionality from node functionality has been a recognized improvement for over twelve years. There is no technical objection to it, and it would allow Bitcoin to be used by a wider range of software without requiring full node operation. It hasn't happened, not because anyone disagrees, but because the informal merge process can no longer distinguish between "don't change the consensus rules" and "don't change anything."

The result is that good ossification, the properties that should never change like the 21 million cap and proof of work, looks the same as bad ossification, the improvements that everyone agrees on but that can't get through. And while Bitcoin Core sits frozen on the changes it should be making,

it is accumulating technical debt. Every improvement that can't get merged forces workarounds, workarounds breed complexity, and complexity breeds bugs. The codebase gets worse over time precisely because the process that is supposed to protect it can no longer process anything at all.

Bitcoin the network is fine. **Bitcoin Core the implementation is locked, and it is rotting from the inside.**

Lightning: A Partial Answer That Proves the Problem

The Lightning Network is the most successful scaling solution Bitcoin has produced. By early 2026 it processes millions of transactions with sub-second settlement and near-zero fees. Major exchanges have integrated it, and wallets like Strike and Phoenix have made it genuinely usable for everyday payments at small values. Lightning is real and works within its design constraints, which does not answer this essay's thesis.

Lightning solves throughput for micropayments, but it does not solve the bigger jobs stablecoins are capturing: large-value cross-border transfers, corporate treasury management, remittance corridors. These require on-chain finality that Lightning cannot provide. More importantly, Lightning does not address the governance problem at all. It is a layer built on top of a frozen base layer. If Bitcoin Core cannot implement UTXO set commitments, wallet-node separation, or formal verification, none of which touch consensus rules, then Lightning inherits those limitations. A second layer on a frozen first layer is still a frozen system.

The Market Bitcoin Should Own

Bitcoin educated hundreds of millions of people on what money could be. Then it didn't do much with that knowledge.

Stablecoins are filling the gap. As of February 2026, the total stablecoin market cap sits at $312 billion, up from $205 billion at the start of 2025, a 49% increase in twelve months. USDT dominates at roughly 60% market share, and total stablecoin transaction volume hit $33 trillion in 2025, surpassing Visa's annual payment volume.

Every dollar moving through a stablecoin is a dollar moving through a system with a central issuer, a kill switch, and blacklist capability. Tether can freeze funds and Circle can comply with government seizure orders. These are the exact properties Bitcoin was designed to eliminate, and yet the use cases Bitcoin was designed to own, payments, remittances, and cross-border settlement, are being captured by systems that have them. Not because of technical impossibility, since the throughput problem is solvable and Lightning proves it at small scale, but because governance paralysis prevents the base-layer changes needed to compete at scale.

Meanwhile, 137 countries representing 98% of global GDP are exploring CBDCs. China's e-CNY is the most advanced, having processed 3.48 billion cumulative transactions worth 16.7 trillion yuan by November 2025, with 230 million verified wallets. On January 1, 2026, China made the e-CNY interest-bearing, the first CBDC in the world to do so. Wallet balances now accrue interest, receive deposit insurance, and are treated as commercial bank liabilities. The e-CNY just moved from a pilot to a core financial instrument, and adoption is accelerating.

The GENIUS Act, signed in 2025, established the first federal regulatory framework for stablecoins in the US. Stablecoin issuers received provisional banking charters, and PayPal, Stripe, and Circle are integrating stablecoin rails into mainstream financial infrastructure. The regulatory environment is being built specifically for centralized digital money. Bitcoin, the asset that should have been the decentralized alternative to all of this, can't participate because Bitcoin Core can't adapt.

The answer is not a new chain. Ethereum runs multiple client implementations, including Geth, Nethermind, Lighthouse, and Teku, and this has genuinely improved its resilience. But Ethereum is not Bitcoin. It does not have Bitcoin's security model, Bitcoin's fixed supply, Bitcoin's credibility, or Bitcoin's role as digital money. Multi-client architecture works. It just needs to happen on Bitcoin.

The Path: Implementation Diversity in 2026

The call for implementation diversity is not new and not theoretical. The movement already exists, incompletely and with significant obstacles.

Bitcoin Knots, then maintained by Luke Dashjr, was the most visible example, a Bitcoin Core derivative with stricter transaction relay policies, filtering what it considers spam transactions including Ordinals and Runes data. In January 2024, Knots ran on fewer than 70 nodes, but by September 2025 it had surged to over 4,700, approximately 25% of all public Bitcoin nodes, driven by backlash to Bitcoin Core's planned removal of OP_RETURN data limits in version 30. That was the largest shift in Bitcoin node software outside of a hard fork event, and no chain split occurred. Core and Knots enforced identical consensus rules and diverged only on relay policy. That is how a policy exit on Bitcoin is supposed to work.

Luke's Knots later followed a BLAKE2b proof-of-work hard fork and left Bitcoin. The Core-derived forks that stayed on Bitcoin are still that Core lineage. They are still policy forks, not an architectural alternative. They share Core's codebase, Core's architectural assumptions, and by extension Core's structural limitations.

btcd, written in Go, has been in production since 2013 and deliberately separates wallet functionality from the node. That is the exact architectural improvement Bitcoin Core has failed to implement for twelve years. But btcd has a problem that undermines its entire value proposition:

it keeps getting the consensus rules wrong. In 2022, a developer intentionally crafted transactions that stalled btcd nodes twice in one month, first by exploiting a witness size parsing failure, then by triggering a script size limit that btcd still enforced after Taproot removed it. In 2024, two more consensus bugs were disclosed: a misimplementation of BIP 68 and BIP 112, and a FindAndDelete bug in legacy signature verification. Each one meant btcd would have accepted or rejected blocks differently from Core. The whole point of a second implementation is to catch exactly these divergences. btcd keeps being the thing that diverges.

libbitcoin, written in C++ and created by Amir Taaki, is now maintained by Eric Voskuil. It implements Bitcoin's protocol from scratch rather than forking Core, which is architecturally ambitious. But it abandons the UTXO model entirely. There is no UTXO set and no mempool. It stores transactions, headers, and a relationship table linking the two. To check whether an input has been spent, it looks up whether a spending transaction exists, the inverse of Core's approach, which maintains a live index of what is unspent. Every wallet, exchange, block explorer, and Lightning node in the Bitcoin ecosystem is built to query a UTXO set and Libbitcoin does not have one. It is a crystal castle, architecturally interesting, but the ecosystem cannot plug into it without being rebuilt almost entirely.

Bcoin, written in JavaScript and created by Christopher Jeffery, was built as backend infrastructure for Purse.io. Purse.io shut down in 2020. Bcoin has no known production deployment since then and commit activity has gone sparse. It is effectively orphaned.

A Core-derived fork is a derivative. btcd diverges. libbitcoin is incompatible. Bcoin is dead. Each fails for a different reason, but the failure is the same: none of them can prove they enforce the same consensus rules as every block Bitcoin has ever produced. Without that proof, a second implementation is just another codebase asking the network to trust it.

Bitcoin Commons is one public attempt to build that proof on the same chain, a from-scratch Rust client built from the Orange Paper, a human-readable mathematical specification of Bitcoin's consensus rules. Full-chain differential replay against Core is the scale test. `blvm-bench` targets on the order of 900,000 blocks. Those runs are operator-driven and resource-heavy. That is not a claim every CI job has zero-divergence proof to tip, so Commons has not yet passed the bar in the paragraph above. Methodology lives in the differential testing guide. Cryptographic merge authorization is designed and stays off until security review, key management, and community validation. No previous alternative shipped that mix of a readable spec, a Specification Lock, and a full-chain differential program.

The 2025 OP_RETURN controversy proved something important: that migration is latent demand, not apathy. The Bitcoin community will move when Core's direction conflicts with its values. It just needs somewhere worthy to move to.

The Only Move Left

Most people on both sides of the block size war believed what they were arguing. But the people who stood to benefit from Bitcoin's inability to adapt didn't need the argument to go their way. **They just needed it to consume enough time and energy to break coordination, and it did. Paralysis was the victory condition, and it was met.**

But paralysis is not permanent because it is a governance failure, not a protocol failure. Bitcoin’s consensus rules are sound, its fixed supply can’t be changed, and its proof-of-work security is unmatched. What Bitcoin lacks is a second implementation, written in different code, that can be checked against a specification rather than reverse-engineered from Core.

Bitcoin Core’s monopoly is not a law of nature but a historical accident, sustained by governance trauma and the absence of a viable alternative, and that alternative is now buildable. The tools exist, the community has demonstrated it will migrate, and the only question is whether someone builds the thing that Gavin called for fifteen years ago and proves it works before the window closes.

Stay on Bitcoin. Make the implementation layer cheaper to leave.

For why the adversarial culture that followed that trauma is still doing work, see *The Adversarial Default*. For what has to happen next, in what order, see *Making Core Irrelevant*.

Sources

- Epstein files, DOJ document releases: July 20–21, 2014 Epstein to Peter Thiel correspondence; October 13, 2016 Epstein to Alsabbagh/Alahmadi; Epstein-Pierce-Summers meeting (December 2025 release)
- Brian Armstrong, "Scaling Bitcoin: The Great Block Size Debate," Coinbase Blog, May 2016
- House Oversight Committee documents: Epstein to MIT/DCI/Core developer funding chain; Joi Ito email to Epstein, April 25, 2015 (November 2025 release)
- Bitcoin Core notice, CVE-2018-17144, September 20, 2018
- Luke Dashjr inflation bug quote, Cointelegraph, May 2019
- DefiLlama stablecoin market cap: live data, February 2026
- CoinGecko: stablecoin market cap cross-check (~$314B same date)
- DeFi Llama / Yahoo Finance year-end report: stablecoin growth $205B to $312B, December 30, 2025
- Bloomberg / Artemis Analytics; a16z State of Crypto 2025: stablecoin transaction volume ~$33T
- Yahoo Finance: GENIUS Act / stablecoin banking charters, December 30, 2025
- Atlantic Council CBDC Tracker: 137 countries / 98% GDP, July 2025
- People's Bank of China via gov.cn; BeInCrypto: e-CNY transactions, wallets, interest-bearing framework, December 29, 2025
- Bitbo; BitcoinNews.com; Bitfinex blog: Bitcoin Knots surge to ~25% of nodes, September 2025
- Bitfinex blog, September 5, 2025; Coin Dance: Knots trajectory 69 to 4,713 nodes
- Yellow.com research; CryptoSlate, August 31, 2025: Bitcoin Core v30 / OP_RETURN
- btcsuite/btcd on GitHub

SECURESOVEREIGN

Markdown is the canonical source: https://secsov.com/articles/governance-paralysis-was-the-victory/index.md

Formatted snapshot of https://secsov.com/articles/governance-paralysis-was-the-victory

Josh · Secure Sovereign