

One-minute blocks: an analysis of the block time proposal

Leo Beltrán

October 28, 2026

Summary

CHIP-2025-03 proposes lowering Bitcoin Cash's target block interval from ten minutes to one. The problems it cites are real: confirmation variance, UTXO contention in DeFi, and dependence on services that require confirmations. But the change doesn't eliminate them — it rescales them. And to do so, it modifies one of the network's foundational constants, adds a permanent abstraction to the protocol, and shifts concrete costs — above all onto SPV clients — that have not yet been quantified.

My request is concrete: that the CHIP not lock in on November 15, 2026, and that it be discussed again for a future upgrade, once costs have been measured and specific solutions to each problem have been exhausted. It is not a “never.” It is a “not yet.”

Context

Bitcoin Cash currently produces one block every ten minutes on average. Because proof of work is a random process, the time between blocks follows an exponential distribution: some confirmations arrive within seconds and others take more than half an hour, even though the average holds.

CHIP-2025-03, *“Faster Blocks for Bitcoin Cash,”* written by *bitcoincashautist*, proposes reducing that average to one minute and introducing a system of “ticks” — units of idealized time — so that consensus rules that currently depend on block height keep working the same way.¹ The CHIP started in March 2025 proposing two minutes, and moved to one in April of that year after modeling the network's real latency.²

At the time of writing, the CHIP is listed as *“Proposed”* (last edited October 24, 2026) and is being discussed as a candidate for the May 2027 upgrade, whose lock-in date is November 15, 2026.^{3,4} Since September, a dedicated thread has been collecting statements of support and dissent from ecosystem participants.⁵ There are significant objections and conditional positions: some reject the change itself, as I do, and others reject activating it **now**. The author, for his part, argues that

¹CHIP-2025-03 (version of October 24, 2026), section *“Blockchain Height Abstraction”*: defines ticks as a unit of idealized time (600 per 10-minute block, 60 per 1-minute block).

²*bitcoincashautist*, [post #19](#), April 2, 2025: latency model (pings < 300 ms; > 80% of VPS providers with > 436 Mbit/s) and worst-case orphan estimate < 1.99%.

³CHIP-2025-03 (version of October 24, 2026), *header*: *“Status: Proposed,” “Last Edit Date: 2026-10-24.”* The activation schedule is in the *“Deployment”* section.

⁴kzKallisti, [post #57](#), “2027 protocol upgrade ideas” thread, August 8, 2026: candidate CHIPS for 2027 and lock-in date (November 15, 2026).

⁵Thread *“CHIP-2025-03 Faster Blocks: gathering statements of support (and dissent),”* opened September 15, 2026; last post reviewed: October 25, 2026. Positions are also tracked on Consensus.cash.

there will never be unanimity and that you “can’t hold the horse” for those who disagree.⁶ There are only a few days left before lock-in, which is why this essay ends with a concrete request.

1. What the CHIP is after

Before arguing against a proposal, it’s worth understanding it at its best. That isn’t easy here, because the CHIP itself defines its goal in a striking way: in response to the objection that its list of benefits kept changing over time, the author replied that **“the CHIP’s goal is to deliver faster blocks,”** and that the list of benefits describes how each kind of user would experience them, because “average user doesn’t exist.”⁷ Tom Zander read that reply as an admission: “There is no goal other than the change itself.”⁸

The arguments in favor are these:

1. **Faster first confirmation.** There are situations where 0-conf isn’t enough: risky transactions, irreversible cross-chain swaps (Thorswap, Sideshift), multi-currency payment gateways (BitPay, Bitgree), and multi-currency wallets that don’t show unconfirmed funds.⁹ There, block time is the visible speed limit. With one-minute blocks, 95% of first confirmations would arrive in under three minutes.¹⁰
2. **More predictable multiple confirmations.** For high-value operations, the wait would have much less dispersion.
3. **Faster recovery in DeFi.** When two users of a contract like Cauldron collide over the same UTXO, the wait to retry would drop from ten or thirty minutes to one or three.¹¹
4. **A “sliding scale.”** It’s not that a one-minute block is safer than a ten-minute one: a one-minute block accumulates, roughly, a tenth of the work. What you gain is the ability to choose more finely where to sit between speed and security. In rnbrady’s words: “a sliding scale between \$100 of security in 1-min, and \$1000 of security in 10-minutes.”^{12,13}
5. **Secondary benefits:** finer granularity for *locktime* contracts, more frequent payouts for small pools, a difficulty algorithm that reacts faster to hashrate drops, and faster recovery of the mempool’s minimum fee.¹⁴
6. **The cost of doing nothing.** The CHIP argues that keeping ten minutes is not neutral: new users form their first impression against faster chains, and the longer we wait, the more software will assume ten minutes as fixed.¹⁵

⁶bitcoincashautist, [post #59](#), “2027 protocol upgrade ideas” thread, August 8, 2026: “You will never get unanimous. [...] There will always be those who disagree, can’t hold the horse just for them.”

⁷bitcoincashautist, [post #156](#), May 29, 2026, replying to JohnGalt: “The CHIP’s goal is to deliver faster blocks”; “Average user doesn’t exist.”

⁸tom (Tom Zander), [post #158](#), June 1, 2026: “There is no goal other than the change itself.”

⁹bitcoincashautist, [post #157](#), May 29, 2026: real 1-conf users (Bitgree, BitPay, Coinsbee, Thorswap, Sideshift). “Nobody wants to be a 1-conf user, everyone wants to be a 0-conf user.”

¹⁰CHIP-2025-03 (version of October 24, 2026), section “[Summary](#)”: “95% of single-confirmation transactions clearing in under 3 minutes, compared to 25% exceeding 14 minutes today.” Developed in “[Instant Improvement in 1-Confirmation Experience](#).”

¹¹bitcoincashautist, [post #156](#), May 29, 2026: “Those users today wait 10-30 minutes when caught by a divergence event. With 1-minute blocks, the same users would wait 1-3 minutes.”

¹²rnbrady, [post #154](#), May 28, 2026: “you get a sliding scale between \$100 of security in 1-min, and \$1000 of security in 10-minutes.”

¹³BitcoinCashPodcast, [post #161](#), June 2, 2026: “faster blocks is sometimes a benefit to speed and sometimes a benefit to security.”

¹⁴CHIP-2025-03 (version of October 24, 2026), sections “[Blockchain Time Granularity](#),” “[More Responsive Mempool Fee Decay](#),” and “[Improved Difficulty Adjustment Stability](#).”

¹⁵CHIP-2025-03 (version of October 24, 2026), section “[Status Quo](#),” subsection “[The Cost of Inaction](#)”: “The status quo is not a neutral choice.”

The question is not whether these benefits exist, but whether they justify what they cost, and whether this is the right way to obtain them.

2. The problem is real, but the change doesn't eliminate it

2.1. Variance, with the right numbers

That block intervals follow an exponential distribution has two consequences. The first is that it is **memoryless**: whenever a transaction arrives, the expected wait until the next block is ten minutes. The second is that **the standard deviation equals the mean**. That's where the CHIP's own figures come from: roughly one in four transactions waits more than 14 minutes ($e^{-1.4} \approx 25\%$) and one in seven waits more than 20 ($e^{-2} \approx 13.5\%$).¹⁶

For someone who needs **one** confirmation, the change genuinely improves things. But for those who need **several** — the case of exchanges — the benefit isn't what people tend to imagine. If an exchange sets the number of confirmations it requires according to the security it wants, then with one-minute blocks it will multiply that number by ten.^{17,18} Kraken, for example, currently requires 15 BCH confirmations;¹⁹ with one-minute blocks it would need about 150 for comparable security. **The average wait would be the same: about 150 minutes.**

What does change is the dispersion. With six ten-minute confirmations, the probability of waiting more than 90 minutes is 11.6%; with sixty one-minute confirmations, it falls to 0.03%. In other words: for exchanges, the CHIP doesn't offer **speed**, it offers **predictability**. That is a much more modest benefit than the phrase "ten times faster blocks" suggests.

2.2. DeFi: mitigating a symptom is not solving the problem

UTXO contention in publicly accessible contracts is a legitimate problem: two users try to operate on the same UTXO, nodes see different versions of the mempool, and one of the two transactions fails. But one-minute blocks don't eliminate that problem: they only shorten the wait until it resolves.

JohnGalt quantified it: under adverse conditions, the window of divergence between nodes after each block would go from covering about 99% of the interval to about 91%. His conclusion: "Without mempool synchronisation, it doesn't matter if blocks are a minute or 10 minutes."²⁰ When a user reported waiting 115 minutes for a block and concluded that faster blocks were needed,²¹ rnbrady replied with something that sums up my position: with faster blocks, the error message would still appear, it would just go away sooner; **"We should solve the problem completely."**²²

If the problem is DeFi contention, it makes sense to look for a solution **specific to that problem**, in the layer where it arises, that doesn't depend on changing the block interval. There are paths worth exploring:

¹⁶CHIP-2025-03 (version of October 24, 2026), section "**Motivation**": "Even with steady hash power, about one in four transactions waits over 14 minutes, and one in seven waits over 20 minutes."

¹⁷ShadowOfHarbringer, [post #151](#), May 28, 2026: exchanges would scale to 10× confirmations.

¹⁸BitcoinCashPodcast, [post #153](#), May 28, 2026: "We are NOT trying to trick exchanges into 10x lower confirmations."

¹⁹rnbrady, [post #150](#), May 28, 2026: Kraken requires 15 confirmations; "I'm not aware of any exchanges that accept 1-conf."

²⁰JohnGalt, [post #149](#), May 28, 2026: "With 1-minute blocks, it is divergent for 55 seconds, about 91% of the interval"; "Without mempool synchronisation, it doesn't matter if blocks are a minute or 10 minutes."

²¹CashDragon, [post #162](#), June 29, 2026: a 115-minute wait (block 957,389).

²²rnbrady, [post #164](#), June 29, 2026: "We should solve the problem completely."

- **In services with their own infrastructure:** batch requests per application and order them so that conflicting transactions aren't produced.
- **In fully decentralized models:** use double-spend proofs (DSProofs) to warn the user of a possible collision, and even estimate a probability of success based on how many nodes saw each version.

Using a real problem to justify a change that doesn't solve it has an extra cost: the problem is still there after activation, and there is no longer a fresh justification for tackling it.

2.3. A permanent solution to a temporary problem

Many of the cases used to justify the CHIP depend on **centralized third parties**: exchanges that require confirmations, gateways that don't accept 0-conf, and multi-currency wallets that don't show unconfirmed funds. That is the temporary problem: the ecosystem's current dependence on entities it doesn't control.

At the same time, the network's recent upgrades — CashTokens, native introspection, more expressive contracts — point precisely the other way: toward Bitcoin Cash depending less and less on those intermediaries, with decentralized finance and self-custody as the engine of its use. There is something contradictory about permanently changing a protocol constant to better accommodate the very actors the network is trying to become independent from.

The CHIP replies that the change also favors self-custody and DeFi.²³ But in DeFi, as we saw, the change mitigates the problem without eliminating it.

3. The costs

3.1. SPV clients: the central cost

It's worth being precise about whom this cost affects, because not all light wallets work the same way:

- **True SPV wallets**, such as Flowee Pay, download the header chain and ask nodes for the filtered blocks that contain their transactions. They are the ones that verify each payment's inclusion on their own, without relying on a trusted server, and they are the most affected.
- **Wallets that use the Electrum protocol**, such as Electron Cash, also download headers, though only from a checkpoint rather than from the genesis block.
- **Many BCH mobile wallets** query Fulcrum servers and don't store the header chain. Their direct cost is lower, but they shift it onto those servers and depend on them more.

In other words: the cost falls hardest precisely on the wallets that depend least on third parties.

Each header takes 80 bytes. With ten times more blocks, annual growth goes from about 4.2 MB to about 42 MB, a figure on which the CHIP and its critics agree.²⁴ The problem has three dimensions, though they don't weigh the same.

1. Storage. The usual answer is that storage always gets cheaper. But that assumption is not guaranteed, as the 2026 surge in memory prices, driven by demand for AI servers, showed.^{25,26}

²³CHIP-2025-03 (version of October 24, 2026), section "[Reduced Reliance on Centralized Exchanges](#)": "*faster confirmations also strengthen self-custody as a practical alternative to leaving funds on CEXes.*"

²⁴CHIP-2025-03 (version of October 24, 2026), section "[Header Chain Growth](#)": "*Header chain growth rises 10×, from ~4.2 to ~42 MB/year*" (phrase from the "[Reader's Guide](#)").

²⁵The Register, "[DRAM prices expected to double in Q1 as AI ambitions push memory fabs to their limit](#)," February 2, 2026, with TrendForce data.

²⁶TrendForce, "[AI Server Demand Continues to Support Memory Prices in 3Q26](#)," July 3, 2026.

Even if temporary, it shows that continued cheapening can't simply be assumed to dismiss a cost that accumulates on every device, forever.

2. Synchronization. Tom Zander, who has developed Flowee's SPV wallet for about fifteen years, argues that SPV wallets scale because they separate the network's load from each wallet's load, and that asking a node for filtered blocks one at a time becomes slow when there are many more blocks. His conclusion is blunt: "Add one minute blocks and SPV will die"; "Making it 10 times as slow is a death sentence."²⁷ Zander also warns that laboratory tests are optimistic: the network has very slow nodes, mobile connections are very irregular, and latency can't be eliminated because a wallet can't act on a filtered block until it receives it.²⁸ In numbers: a wallet that has been offline for a month today has to catch up on about 4,300 blocks; with one-minute blocks, about 43,000.

3. Constrained devices. JohnGalt raises the case of a wallet that syncs headers over NFC at the moment of payment: a week of headers is about 80 KB today and transfers in a one-and-a-half-second tap; with one-minute blocks it would be about 800 KB and some fifteen seconds of contact — "not tap, but hold."²⁹ And this isn't a concern limited to opponents: kzKallisti, who doesn't oppose the CHIP in principle, said in August that this NFC card-payment problem was "alone enough to give me pause," and suggested pushing the change to 2028.³⁰

Supporters reply with mitigations: requesting blocks in larger batches,³¹ compact filters (BIP-157/158),^{32,33} Dogecoin's precedent,^{34,35} and CHIP-2026-02's header pruning.^{36,37} They all have something in common: they depend on work that isn't done yet or isn't deployed. Third-party wallet code has to change, compact filters have to be adopted, and another CHIP has to activate. If the SPV cost is solved by those mitigations, the reasonable thing is for them to arrive before one-minute blocks, not after.

3.2. Less work per block, and mining centralization

To Calin Culianu, lead developer of Bitcoin Cash Node, it "feels fundamentally wrong" to take a chain with about two hundred transactions per block and a tiny fraction of total SHA-256 hash, and divide it by ten: about twenty transactions per block and "far less secure blocks."³⁸ The

²⁷tom (Tom Zander), [post #41](#), "2027 protocol upgrade ideas" thread, August 1, 2026: "Add one minute blocks and SPV will die"; [post #44](#), August 2, 2026: "Making it 10 times as slow is a death sentence."

²⁸tom, [post #168](#), August 4, 2026: "Tests I did gave about 100 blocks per second"; "laboratory tests are very likely optimistic."

²⁹JohnGalt, [post #149](#), May 28, 2026: "one week becomes about 10,080 headers, roughly 800 KB"; "With 10x more headers it now becomes not tap, but hold."

³⁰kzKallisti, [post #62](#), "2027 protocol upgrade ideas" thread, August 8, 2026: "This issue alone is enough to give me pause on Faster Blocks"; "it might be worth pushing to 2028."

³¹bitcoincashautist, [post #48](#), "2027 protocol upgrade ideas" thread, August 2, 2026: "I changed this one line from 9 to 99 [...] the wallet jumped from syncing about 5 blocks / second to 50 blocks / second!"

³²nyusternie, [post #47](#), "2027 protocol upgrade ideas" thread, August 2, 2026: BIP-157/158 and compact filters.

³³nyusternie, [post #170](#), August 4, 2026: point-by-point reply to tom (post #168) on SPV sync, based on BIP-157/158.

³⁴CHIP-2025-03 (version of October 24, 2026), section "Reader's Guide": "Dogecoin has run 1-minute blocks since 2013."

³⁵bitcoincashautist, [post #61](#), "2027 protocol upgrade ideas" thread, August 8, 2026: "Dogecoin has been running 1-minute blocks the entire time (since 2014) without breaking SPV or exchanges or anything else."

³⁶CHIP-2025-03 (version of October 24, 2026), section "Reader's Guide": "storage is already mitigated by Electrum header merkle proofs, and header pruning via CHIP-2026-02 makes it fixed-size." Developed in "Header Chain Growth."

³⁷bitcoincashautist, [post #61](#), "2027 protocol upgrade ideas" thread, August 8, 2026: CHIP-2026-02, test implementation in BCHN, and checkpoint pruning in Electron Cash.

³⁸bitcoincashautist, [post #179](#), September 15, 2026, relaying Culianu: "it feels fundamentally wrong to take a chain with like ~200 txns per block and .025% hash divide it by 10x — ~20tx per block and far less secure blocks at 0.0025% hash." Reply: "grief is a rate problem, and the rate is fixed by hashrate, not by block time." The same phrase from Culianu appears in his statement (see the note on post #4 of the statements thread). The hash figures in the quote should be taken

author replies that the damage an attacker can do depends on their hashrate, not on block time.

In his September statement, Culianu disapproved of the CHIP “for now”: the code wasn’t ready and was likely to have new bugs.³⁹ His conditions — finishing the code, testing it more, and having it reviewed — are, in essence, the same ones this essay sets out.⁴⁰

A related risk is mining centralization. Large pools lose less to orphaned blocks than small ones, so the higher the orphan rate, the greater their relative advantage. The current version of the CHIP estimates that rate at between 0.41% and 1.94% depending on the modeled conditions, against a tolerable threshold of 2%,⁴¹ and the author cites Dogecoin’s experience, at around 0.15%.⁴² But the worst modeled case sits just six hundredths of a percentage point below the threshold the CHIP itself defines as tolerable. Cashrider asked for larger propagation margins and warned that other networks’ experience isn’t directly comparable to BCH’s.^{43,44} albaDsl, for his part, disapproved of the change because the implementation isn’t ready, and because he worries that too much safety margin is being given up: a system with shorter blocks is more rigid in the face of catastrophic events.⁴⁵ The model exists, but it is a model: the real rate will only be known once the network is running.

3.3. Permanent complexity

The CHIP technically solves the problem of contracts that depend on block height: it evaluates them against a normalized scale ($\text{tick} / 600$), so that they behave in real time exactly as they do today.⁴⁶ But the solution has a cost of its own: it introduces a **new, permanent abstraction** that anyone programming for Bitcoin Cash will have to understand and keep in mind. There are now block heights, ticks, confirmations measured in blocks, and confirmations measured in ticks. mainnet_pat warned that ticks seem to add complexity and will confuse newcomers to the ecosystem.^{47,48}

The strongest criticism on this point comes from the person who has to write the code. Culianu explained that, although ticks simplify the concept, for whoever implements them “it’s like a NUKÉ went off in the codebase. Everything got touched.”⁴⁹ He argued that the CHIP touches more than necessary and “marr[ies] lots of stuff now to PoW” that wasn’t tied to it before, including block size and $n\text{LockTime}/n\text{Sequence}$.⁵⁰ Of everything in the CHIP, what he likes least

as illustrative: according to CoinWarz, on October 28, 2026 BCH had about 4.4 EH/s ([chart](#)), around 0.5% of total SHA-256.

³⁹cculianu (Calin Culianu), [post #4](#), September 15, 2026: “Disapprove (for now). Rationale: code is not ready and likely to have new and exciting bugs.”

⁴⁰cculianu, [post #13](#), September 15, 2026: finish the code “ASAP w/ more testing” plus an additional review.

⁴¹CHIP-2025-03 (version of October 24, 2026), section “[Increased Orphan Rate](#)”: “Orphan rate rises modestly: 0.41–1.94% across our modeled conditions, against a 2% tolerable threshold.” On large pools’ advantage: “[Mining Centralization Risk](#).”

⁴²bitcoincashautist, [post #178](#), September 6, 2026: Dogecoin ~0.15% orphans; CHIP model 0.41–1.94%; chipnet ~6 months before mainnet.

⁴³Cashrider, [post #172](#), August 11, 2026, on propagation and centralization: “higher stale rates can create some advantage for miners with better connectivity and propagation infrastructure.”

⁴⁴Cashrider, [post #186](#), September 18, 2026: other networks’ experience is not directly comparable to BCH’s.

⁴⁵albaDsl, [post #99](#), September 30, 2026: “I disapprove mainly because the implementation is not ready, but I’m also wondering if we are giving up too much of our safety margin with 1-min blocks.”

⁴⁶CHIP-2025-03 (version of October 24, 2026), section “[Transaction Locktime and Input Sequence](#)”: “Height-mode $n\text{Locktime}$ and $n\text{Sequence}$ are normalized to the 10-minute scale”: a transaction is final at block b if $T(b) > 600 * L$.

⁴⁷mainnet_pat, [post #18](#), September 16, 2026: “Approve, but with caution”; “The ticks approach seems to bring more complexity.”

⁴⁸mainnet_pat, [post #184](#), September 16, 2026: ticks and API changes may confuse newcomers.

⁴⁹cculianu, [post #226](#), October 11, 2026: “To someone that has to implement this – it’s like a NUKÉ went off in the codebase. Everything got touched.”

⁵⁰cculianu, [post #228](#), October 11, 2026: “you touched more than needed to be touched”; “What you did was marry lots of stuff now to PoW in a way where it wasn’t explicitly married before. Including block size! Including $n\text{Locktime}/n\text{Sequence}$!”

is precisely the solution for contracts: redefining `nLockTime` so that height means an abstract height in an abstract space where blocks are still ten minutes, which, in his view, makes those fields very hard to reason about.⁵¹ And he summed up: “implementation is a nightmare for me. EVERY TEST BROKE.”⁵² His final assessment is that “compromises had to be made” and that, although ticks are presented as a simplification, they are a layer of abstraction that adds complexity to the code and to the node.⁵³

The author and `emergent_reasons` reply that, without ticks, the same places would have to be touched, but adjusting each rule by hand, with more risk of error.^{54,55} That’s a good argument for ticks if the interval is changed, but not for changing it. Perhaps it would have been simpler to have one-minute blocks from the start. But that wasn’t the parameter chosen, and changing it now requires a compatibility layer the network will carry forever.

3.4. Unquantified activation costs

The CHIP describes the activation costs — for nodes, mining, exchanges, light wallets, explorers, indexers, and DeFi applications — as “real but bounded.”⁵⁶ JohnGalt asked for them to be estimated in numbers, by consulting each team.⁵⁷ The author replied that most actual stakeholders — node teams, pool operators, exchanges, and wallet developers — are already reading the discussion or are at least aware of it.⁵⁸ But being aware of the proposal is not the same as having estimated what it costs to adapt to it. Cashrider raised a concrete risk that this signal doesn’t capture: third-party software that may keep running even though its assumptions about confirmations, block height, or security have changed **silently**, with nobody noticing until after activation. And he asked what readiness criterion would be sufficient.⁵⁹ The CHIP itself acknowledges that this is the main risk: that services crediting deposits by number of confirmations end up, without knowing it, requiring a tenth of the work.⁶⁰ And 1KBCH pointed out that the real improvement for users ultimately depends on the confirmation policies of Binance, OKX, and Kraken,⁶¹ which nobody controls from the protocol.

There is also a cost no CHIP accounts for: **opportunity cost**. The time developers, reviewers, and community members spend coordinating this change is not spent on other things, including the specific solutions in section 2.2.

⁵¹cculianu, [post #230](#), October 11, 2026: “Of all the things in this chip – the modification of `nLocktime` to have a changed meaning is the one I like the least.”

⁵²cculianu, [post #237](#), October 11, 2026: “implementation is a nightmare for me. EVERY TEST BROKE.”

⁵³cculianu, [post #243](#), October 11, 2026: “I’m not against the CHIP [...] I think it would be devastating at this point to the community if this were not to activate either now or 1 year from now. And I think we should do it.” Also: “compromises had to be made”; “there is more complexity added to the code and to the node.”

⁵⁴emergent_reasons, [post #222](#), October 11, 2026, quoting the author: without ticks, the change would be “the same change with the abstraction replaced by *N* ad-hoc re-derivations of it.”

⁵⁵bitcoincashautist, [post #227](#), October 11, 2026: “without ticks the same everything would’ve had to be touched – just touched in a different way.”

⁵⁶CHIP-2025-03 (version of October 24, 2026), section “[Activation Costs](#)”: “The activation costs of this proposal are real but bounded.”

⁵⁷JohnGalt, [post #149](#), May 28, 2026: “The CHIP should include or solicit third-party estimates: how much time node teams need, how much time mining pools need, how much time wallets need.”

⁵⁸bitcoincashautist, [post #156](#), May 29, 2026, on costs: “Most actual stakeholders (node teams, pool operators, exchanges, wallet developers) are reading or at least aware.”

⁵⁹Cashrider, [post #172](#), August 11, 2026: “A program that continues running while its assumptions about confirmations, block height, elapsed time or security have silently changed may be much harder to notice.”

⁶⁰CHIP-2025-03 (version of October 24, 2026), section “[Silent Assumptions](#)”: the risk is that crediting by confirmations ends up worth, without warning, “1/10th the chainwork” (phrase from the “[Reader's Guide](#)”).

⁶¹1KBCH, [post #51](#), September 23, 2026: “Real UX Gain also base on Binance/OKX/Kraken conf-count Policy.”

3.5. Improving one experience by worsening another

User experience has two sides that the change puts in tension: the **speed** at which users see their transaction confirmed, and the **security** they believe that confirmation represents. For many people, “confirmed on the blockchain” means “final.” One-minute blocks improve the first side at the expense of the second.

It’s also worth asking whose experience improves. The most visible benefit goes to people who use exchanges and services that require confirmations: a user who, in the view of many of us who build Bitcoin Cash, is not the user the network was designed for, nor the one the latest upgrades were made for. And the experience that may get worse is a merchant’s.

With one-minute blocks, a confirmation represents about a tenth of the proof of work it represents today, something the CHIP itself acknowledges.⁶² Joemar Taganna (Paytaca) put it this way: “A 1-minute first confirmation is faster, but it is not equivalent to today’s 10-minute first confirmation,” and accumulating the same work still takes the same time.⁶³ Add orphaned blocks to that: at the rate the CHIP estimates (0.41–1.94%) and 1,440 blocks per day, that’s six to twenty-eight orphaned blocks daily.

When a block is orphaned, its transactions return to the mempool and almost always make it into the next block. But:

- a merchant waiting for one confirmation may see their payment go from “confirmed” to “unconfirmed” and get confirmed again a minute later, which is, at the very least, disconcerting;
- reversing a one-confirmation payment costs an attacker about ten times less than today, so “one confirmation” stops meaning what it used to mean, without the merchant knowing.

The cost falls on whoever took the first confirmation as a synonym for security — which is precisely the expectation the network taught them to have. Even if each of these episodes is temporary, the experience of a user who isn’t the priority is improved at the expense of the trust of one who is.

4. Principles

4.1. Adding features is not the same as changing a constant

Bitcoin Cash has changed a lot since 2017: Schnorr signatures, native introspection, CashTokens, more expressive contracts. But those upgrades **added** capabilities without altering the constants that define the network’s base operation. Raising the maximum block size might look like an exception, but it isn’t: the 1 MB limit was not in Bitcoin’s original release; Satoshi added it in 2010 as a provisional anti-spam measure, not as a design parameter.

The CHIP replies that BCH has already replaced its “heartbeat” — the difficulty adjustment algorithm — three times, and that what hurt the network in the past was poor execution, not change itself.⁶⁴ But those replacements changed the **mechanism** that maintains the ten-minute interval, not the interval. They improved the way the constant is held; they didn’t redefine it.

⁶²CHIP-2025-03 (version of October 24, 2026), section “[Security Analysis](#)”: analysis of per-confirmation security; the CHIP acknowledges that each confirmation comes to represent “1/10th the chainwork.”

⁶³joemar_taganna (Joemar Taganna, Paytaca), [post #49](#), September 23, 2026: “A 1-minute first confirmation is faster, but it is not equivalent to today’s 10-minute first confirmation.”

⁶⁴CHIP-2025-03 (version of October 24, 2026), section “[Bitcoin Cash Image Risks](#)”: “BCH has already replaced its “heartbeat” (the DAA) three times; what hurt in the past was poor execution, not change itself.”

The ten-minute interval, by contrast, is at the heart of the original design. It determines issuance, difficulty adjustment, per-confirmation security, and all the documentation, literature, and software accumulated over more than fifteen years of Bitcoin. Changing it is not just another upgrade: it changes one of the coordination points on which everything else was built.

4.2. Who decides, and which parameters should be fixed

Today, with little activity on the network, technical decisions are in practice made by developers. If activity someday becomes enough for miners to live on fees alone, they will probably make their own decisions, on their own or by hiring developers to optimize their nodes. That's why we should be careful with the parameters a group of developers sets today and the whole network will inherit tomorrow.

Not all parameters weigh the same in that sense. Block size is, at bottom, something each miner's incentives can balance: how many transactions to process, how much is earned in fees, and how much is risked in orphaning. That's why a limit like ABLA, activated in 2024, can be relaxed later without invalidating anything built on top of it.

The block interval is a different kind of decision. It can't emerge from what each miner decides: it is necessarily a coordination parameter. And it isn't easy to reverse: once ticks, API changes, and adapted software are activated, going back would be as costly as the original change. Moreover, **the network's incentive model doesn't fundamentally change with the block interval**. If the change doesn't improve incentives and only changes the user experience in a subset of cases, the cost of touching a coordination parameter weighs more.

4.3. 0-conf: the answer we gave for years

For years, whenever someone objected that Bitcoin isn't instant because you have to wait ten minutes, those of us defending Bitcoin Cash always gave the same answer: 0-conf. Transactions are practically instant because they propagate in seconds and because miners' incentive is to include them in the next block. Bitcoin Cash built much of its identity on that answer.

The CHIP argues that 0-conf and confirmations are complementary: they cover different risk cases, and with faster blocks 0-conf's window of vulnerability closes sooner.⁶⁵ On paper, perhaps. People who don't follow the forum won't read the nuances: they'll see that the network that answered "you don't need to wait ten minutes" for a decade changed its protocol so as not to wait ten minutes. To a newcomer, a journalist, or a critic, that doesn't read as a distinction between risk cases, but as an admission that the answer wasn't enough. And that reading costs 0-conf credibility precisely in the cases where it does work, which are the ones we use most.

Bigplayermaster raised a related idea in the statements thread: that "we need to distinguish between two things: speed and the perception of speed," and that "money requires clear, predictable, and stable rules," because it flees assets whose rules seem to change.⁶⁶

0-conf

0-conf is a practical consequence of the protocol's rules and of certain assumptions about security and inclusion in the next block. It is not an explicit consensus guarantee.

⁶⁵CHIP-2025-03 (version of October 24, 2026), section "[Zero-confirmation Transactions](#)": "They are not alternatives; they are complementary."

⁶⁶Bigplayermaster, [post #55](#), September 24, 2026: "we need to distinguish between two things: speed and the perception of speed"; "money requires clear, predictable, and stable rules."

4.4. Identity

Part of the debate is not technical, and the CHIP itself acknowledges it: the concern that Bitcoin Cash might lose its “soul,” as kanat put it, or that the change might weaken its legitimacy claim as Bitcoin’s continuation, as BitcoinCashPodcast argued.^{67,68} Tom Zander, for his part, described the proposal from the start as “a solution searching for a problem.”⁶⁹ The CHIP’s author replied with irony — 20% of blocks already take two minutes or less: “are they soulless blocks?”⁷⁰ — but the objection was never about individual blocks; it was about the rule. In the statements thread, rnbrady, kanat, and Zander disapproved of the change with the same idea: “If it ain’t broke, don’t fix it.”^{71,72,73} A persistent social objection is, in itself, a coordination cost.

5. A secondary observation: merchant adoption

In the experience of those of us who work on direct Bitcoin Cash adoption with merchants, in Argentina and elsewhere, block time has never been a real obstacle when communication emphasizes the technology’s benefits. Between users, transactions are accepted immediately, and tools that support 0-conf, such as prompt.cash,⁷⁴ are used without trouble for online sales.

Where 0-conf works, the ten-minute wait is not the obstacle.

6. What can be done without changing the interval

Looking one by one at the problems the CHIP says it solves, the change is less favorable than it seems. In almost all of them, the CHIP mitigates without eliminating, and in most there are paths that don’t require touching a foundational constant.

Problem	What the CHIP does	Alternative
Slow and variable first confirmation (swaps, gateways, multi-currency wallets)	Speeds it up	DSProofs with risk scoring and amount-tiered policies, already used by Bitgree and Thorswap ⁷⁵
Multiple confirmations at exchanges	Gives predictability, not speed	Require confirmations by value and risk, not a fixed number
UTXO contention in DeFi	Shortens the wait; the collision remains	Order requests in the application, spread liquidity across several UTXOs, extend DSProofs to contracts

⁶⁷kanat, [post #14](#), March 7, 2025: “If BCH goes to 2-minute blocks, it might win some short-term attention but lose its “soul” in the process.”

⁶⁸BitcoinCashPodcast, [post #3](#), March 2, 2025: “a high risk endeavour that could seriously impact the legitimacy claim of BCH to “the real Bitcoin” legacy.”

⁶⁹tom, [post #6](#), March 3, 2025: “It is obviously written from the point of view of a solution searching for a problem.”

⁷⁰bitcoincashautist, [post #15](#), March 7, 2025: “20% of BCH blocks had durations shorter or equal than 2 minutes, are they soulless blocks?”

⁷¹rnbrady, [post #12](#), September 15, 2026: “Disapprove. If it aint broke, don’t fix it.”

⁷²tom, [post #20](#), September 16, 2026: “Disapprove. In balance doesn’t seem to be a good idea.”

⁷³kanat, [post #36](#), September 20, 2026: “Disapprove. If it aint broke, don’t fix it.”

⁷⁴[prompt.cash](#).

⁷⁵bitcoincashautist, [post #157](#), May 29, 2026: “Bitgree uses DSPs, but DSP-score-0 or DSP-fired transactions fall back to 1-conf”; Thorswap applies minimum confirmation tiers by amount.

Problem	What the CHIP does	Alternative
Hour-long-or-more blocks from hashrate drops	Shortens them ⁷⁶	Study the difficulty adjustment's response as a problem of its own
<i>Locktime</i> granularity	Improves it	Time-based <i>locktime</i> already exists ⁷⁷
New users' first impression	Indirectly	Wallets that show 0-conf payments as received; communication

The other predictable objection is that these alternatives depend on third parties who, in almost a decade, haven't adopted 0-conf. But the CHIP depends on the same thing: if an exchange doesn't update its configuration, its 15 confirmations end up worth a tenth of the work — what the CHIP itself calls “silent assumptions.”⁷⁸ The difference is in what is being asked of them. With the CHIP: adapt to a new protocol; if they don't, they become less secure without knowing it. With the alternatives: adopt tools that already exist; if they don't, nothing gets worse.

Hashrate drops are the case where one-minute blocks genuinely help. But they are the symptom of something else: the network has a minimal share of total hash and depends on miners who come and go. That problem deserves its own analysis, not a sideways fix that changes everything else.

7. The request: not in 2027

My request is concrete: that CHIP-2025-03 not lock in on November 15, 2026, and that it be discussed again for a future upgrade, once the alternatives have been exhausted. It is not a “never”: today the burden of proof has not been met. To activate a change of this kind, it would have to be shown that **the benefits clearly outweigh the sum of the costs**: those of implementing, coordinating, and maintaining the change, and the opportunity cost of not putting that effort elsewhere. At a minimum, before any activation we should:

1. **Quantify the activation costs** with estimates from the teams that will bear them (nodes, pools, exchanges, wallets, indexers, and DeFi applications), and define readiness criteria that catch third-party software that may break silently.
2. **Deploy the SPV mitigations first** — compact filters, batched sync, CHIP-2026-02 — and verify that they work with real wallets, not just in theory.
3. **Evaluate specific solutions** for each problem cited (DeFi contention, predictability for services that require confirmations), and show why they fall short before resorting to changing a foundational constant.
4. **Have a complete, tested implementation** on chipnet with enough time, and real orphan-rate data, not just models.
5. **Achieve broad support**, beyond the absence of technical objections among forum participants.

This is not an isolated position. Several participants in the statements thread condition their support the same way: SW9 supports the CHIP conceptually but doesn't think it necessary

⁷⁶CHIP-2025-03 (version of October 24, 2026), “[Reader's Guide](#)”: “Hashrate drops degrade gracefully: 10% of normal hashrate yields ~10-minute blocks instead of ~100-minute ones.” Developed in “[Improved Difficulty Adjustment Stability](#).”

⁷⁷cculianu (Calin Culianu), [post #230](#), October 11, 2026: “If some contract was using height as a proxy for real world time, when there already is a TIME mode, that's on the contract.”

⁷⁸CHIP-2025-03 (version of October 24, 2026), section “[Silent Assumptions](#)”: services that credit by number of confirmations end up requiring a tenth of the work: “a policy change, not a bug.”

to push it for the next cycle if the implementation isn't ready;⁷⁹ Fernando Pelliccioni (Knuth) approves of one-minute blocks but declares himself neutral on the November lock-in;⁸⁰ 2qx approves "if implementation is ready in time";⁸¹ and albaDsl conditions his vote on Culianu signing off on the implementation.⁸² ShadowOfHarbringer even proposed separating ticks from the interval change: introduce the abstraction first and decide the block time later.⁸³

Today these conditions are not met. Postponing one cycle has a bounded, recoverable opportunity cost; activating without meeting them, a permanent one.

8. Conclusion

The one-minute block proposal starts from real problems and is technically well crafted: it models the orphan rate, preserves existing contracts, and acknowledges its costs. But it doesn't eliminate the problems it cites: for exchanges, it offers predictability rather than speed; for DeFi, it shortens the wait for an error that keeps happening; for centralized services that require confirmations, it accommodates a dependence the network is trying to escape. In exchange, it modifies a foundational constant of Bitcoin, adds a permanent abstraction to the protocol, shifts onto SPV clients a cost whose mitigations aren't deployed yet, empties the first confirmation of meaning for the merchant who relied on it, and weakens — at least in the perception of those looking in from outside — the answer, 0-conf, that the community gave for a decade.

The alternative is slower and less flashy: solve each problem in the layer where it arises, and work on communication and wallet design where the problem is one of expectations. But it addresses causes, not symptoms.

A foundational constant should not be changed in a hurry. If you share these concerns, say so in the statements thread before November 15.⁸⁴ Postponing the proposal doesn't kill it: it gives it the time it needs to be done right.

That's why my answer is: not yet, and not in 2027.

⁷⁹SW9, [post #28](#), September 18, 2026: "I don't think it is necessary to push it for the next upgrade cycle if the implementation is not ready."

⁸⁰fpelliccioni, [post #89](#), September 28, 2026: "Approve (one-minute blocks); Neutral (November 2026 lock-in)."

⁸¹2qx, [post #22](#), September 16, 2026: "Approve of the concept, if implementation is ready in time."

⁸²albaDsl, [post #105](#), October 5, 2026: would change his position if Culianu signs off on the implementation, the tests, and the choice of parameter.

⁸³ShadowOfHarbringer, [post #6](#), September 15, 2026: "Would it be possible to separate the tick mechanism from the change to blocktime and make it two separate changes?"

⁸⁴To post a position: "[CHIP-2025-03 Faster Blocks: gathering statements of support \(and dissent\)](#)," on Bitcoin Cash Research.