Making Decentralized Economic Policy

BIP 100 - Theory and Discussion, v0.8.1 - draft

Jeff Garzik <jgarzik@gmail.com>

Bitcoin block size, a change in economic policy

The block size limit debate. This is the first major change in bitcoin economic policy since inception, Jan 2009. There is no perfect answer; as with most modern policy changes, any action - including inaction - generates winners and losers.

A solution is presented which updates bitcoin with a simple solution that addresses hard fork risk in a simple, predictable manner, transitioning the system to a more free-market-based approach to security and scalability.

What is bitcoin? An existential question to answer.

Is bitcoin a commodity ownership transfer system? Is bitcoin a payment system? Is bitcoin a secure service for timestamping and notarizing documents? Is bitcoin a currency for 7 billion world citizens to personally use every day?

To make an informed decision on this change, one must first decide where you want to go, and how to best map the blockchain technology to that goal.

Does bitcoin natively provide an instant, secure, decentralized, trustless, egalitarian currency system? Is that a goal? Is that possible? If we can “turn a knob” and make a leap towards that goal, should we? What other uses are we aiming for? NASDAQ stock settlement, millions of trades per day on the blockchain, major non-currency applications competing for space with, and possibly crowding out, currency-based blockchain applications?

The future of blockchain tech is certainly bright - yet also unclear. The block size limit invites consideration of the best path forward.

High hurdle: Risky “hard fork” required to change policy

Economic policy in the bitcoin system is intentionally hard to change. According to the rules of the bitcoin consensus system, any deviation from the existing rules will cause immediate separation from the main network, as if an immune system attacking cancerous cells. Bitcoin users that fail to follow bitcoin's protocol with 100.0% precision will find themselves banned from Bitcoin Island (until they return to 100.0% compliance).

Changing bitcoin's rules requires a risky operation known as a hard fork, an abrupt global upgrade. There are many policy and engineering risks related to a hard fork. In many ways, it is akin to a Constitutional Convention in the US: Anything about the foundation of the bitcoin system may be changed in a hard fork. Nothing is sacred.

One of the largest risks is social: If two or more populations of users choose to follow different rules, one of bitcoin's key axioms is broken: the same bitcoin may be spent twice (one in each user population). Each population has an enormous incentive to seek consensus and end this situation. If such a split occurs, bitcoin payment processors, exchanges and other businesses must either Pick A Side or simply shut down until the confusion is resolved, severely damaging confidence.

Censorship resistance: an engineering requirement

Preventing the double-spending of bitcoin tokens requires an underlying engineering system that cannot be altered by any party, ever, once transactions are published to the blockchain. This censorship resistance is required for reliable operation of the system, and enables a core principle also found in the Internet itself: permissionless innovation. Apps are more easily built upon top of open platforms such as bitcoin, or the Internet.

This is the core security service "backing" bitcoin. This irrevocable security service may be used to secure transfer currency tokens (BTC), physical assets (smart property), stocks, bonds, and more.

Eroding the censorship resistance of the system erodes bitcoin's long term value. There can only be trust achieved in the system as a whole when bitcoin is viewed as a neutral actor.

What is the block size limit?

A hidden yet crucial metric in the bitcoin system is the $-cost to flood the system with garbage. Absent a speed limit, of sorts, attackers bog down the entire system. In the early days of

bitcoin, when the dollar cost to flood the system was low, originator Satoshi added a limit on the total number of transactions validated every 10 minutes.

When the number of transactions in the past 10 minutes exceeds this speed limit, a bidding war occurs. Bitcoin transactions are prioritized based on the transaction fee attached. Higher fee paid by the user, the more likely their transaction will be prioritized before a lesser-paying user.

In short, with this limit, a block is an economically constrained resource for which bitcoin users bid space according to the laws of supply and demand. Your bitcoin transaction has very little security until it is validated and placed in a block. This system discourages garbage floods (spam) - it costs increasing amounts of money to do so. This system incentivizes conservation and efficiency.

Why change?

The primary reason is removing a roadblock to bitcoin growth. The theoretical limit with the current block size is 7 transactions/second. Any responsible business projecting capacity usage into the future sees the system reaching an absolute maximum capacity, with this speed limit in place. Increasing or removing this limit will encourage businesses to view bitcoin as scalable and capable of supporting millions of new users.

On a more practical level, as blocks get full and the bidding war ensues, the bitcoin user experience rapidly degrades to poor. In part due to bitcoin wallet software's relative immaturity, and in part due to bitcoin's settlement-based design, the end user experience of their transaction competing for block size results in erratic, and unpredictably extended validation times. The user is left waiting around for their payment, without knowing much why, or how to avoid this situation.

What are the consequences of inaction?

In a minority view, bitcoin is a commodity to which an inconsequential payment network is attached. These users view bitcoin as purely as a commodity whose storage cost is far less than the equivalent value of storage. To these users, increasing transaction volume on the network, scaling to support millions of users is simply folly. As the system approaches the speed limit, transaction fees rise. This userbase generates little transaction volume, and higher transaction fees are still offset by the value gained from the network's protection of their commodity. These users likely feel inaction, keeping the speed limit intact, is the most rational step for their situation.

Conversely, a prevailing view is that this speed limit presents a severe constraint on growth and adoption. The speed limit, in the face of userbase growth, will lead to higher fees, ultimately reaching an equilibrium when the system capacity of 7 transactions/second is reached. This will cause a ripple effect where blockchain-centered business plans are shelved, or never started in the first place, for lack of perceived scalability.

Higher fees then drive a preference for centralized services which can aggregate those fees across multiple bitcoin users, or strike bulk rate deals. The use of bitcoin becomes a rather exclusive club.

What are the consequences of action?

Lower fees incentivize use of bitcoin at a key time in its young life. The cost to experiment with new and exciting bitcoin applications remains low. Noname and Big Name developers make plans to build atop bitcoin. Economic activity remains near frictionless. A virtuous cycle is created.

All very positive - yet there are costs.

This change creates some technical, policy, political and social fallout. This action is willfully - artificially? - keeping bitcoin transaction fees low. Thus, much like the US Federal Reserve's goal of keeping inflation low, a more actively managed policy is being put into place. While true the consequences of inaction are painful, the parallels between Fed policy and the proposed solution are inescapable.

Keeping fees low (versus 1MB limit) has additional fallout. The natural economic state of a block is to be 100% full at all times. The only reason why this does not occur today is an anti-spam economic policy in the bitcoin software, which ignores extremely-low-value "dust" transactions which are accompanied by extremely low (or zero) fees. Lowering fees in the bitcoin system thus relies more heavily on the anti-spam system to filter out "junk" The definition of "junk" is subjective and not entirely technical.

Larger blocks push marginal actors off the network, having centralizing effects on the overall network. Techniques exist to mitigate this to some extent, but it remains true. Large blocks are more sensitive to network latency. Network bottlenecks, delays and cyberattacks, unintentional and malicious, are amplified at larger block sizes. The opportunities for bad actors to "push the system around" increases, as marginal actors leave, and decentralization leaves right along with it. 1MB is too small. How much is too much? Good answers remain elusive.

Overall system security may be impacted. Increased traffic on the bitcoin network will most likely exacerbate negative trends in metrics such as bitcoin network size (5,000 and falling), one key metric of network security health. As these metrics fall, increasing network load is shouldered by a decreasing number of actors.

Further, executing a hard fork to enact this change introduces the risks noted above.

If the number of validating actors in the system reaches an extreme minimum, bitcoin's core value of censorship resistance, the core system security is easily eroded by government or private competition.

All depends on the degree of action. Proposal: Rather than simply removing the speed limit, a radical and overly risky move, raise the limit conservatively, simulate, field test and observe the results. Build a process that smooths future upgrades, to avoid the pain points of the current 1MB hard fork.

A concrete proposal: BIP 100

On balance, increasing block size limit above 1MB is needed.

It is clear that the current speed limit, 1,000,000 bytes, is too low and will retard growth and lead to user pain and expense if not increased. What is unclear and under-researched is the opposite side of the equation: How much is too much? Some Chinese pools and exchanges, for example, warn they will be at a disadvantage and ultimately pushed off the network at the originally proposed 20MB size.

Further, it is important to separate the hard fork risk away from the block size change events as much as possible. It is also important to plan ahead for change… yet not plan too far ahead. Be flexible and build a framework for rapid iteration as informed by user input, field experience and market signals.

Protocol changes proposed:

1. Hard fork, to
2. Remove static 1MB block size limit.
3. *Simultaneously*, add a new floating block size limit, set to 1MB.
4. The historical 32MB limit remains.
5. Schedule the hard fork on testnet for September 1, 2015.
6. Schedule the hard fork on bitcoin main chain for January 11, 2016.
7. Changing the 1MB limit is accomplished in a manner similar to BIP 34, a one-way lock-in upgrade with a 12,000 block (3 month) threshold by 90% of the blocks.
8. Limit increase or decrease may not exceed 2x in any one step.
9. Miners vote by encoding 'BV'+BlockSizeRequestValue into coinbase scriptSig, e.g. "/BV8000000/" to vote for 8M. Votes are evaluated by dropping bottom 20% and top 20%, and then the most common floor (minimum) is chosen.

This creates a framework whereby the network may increase the block size by consensus, a lower and less politically risky hurdle than hard fork. Sizes beyond 1MB may be chosen without a flag-day network upgrade. A small size increment limits the potential for unexpected harm to bitcoin network security, and gives the network time to test, prepare and adjust overall behavior.

Other, more complex solutions such as extension blocks - a speed limit workaround - are rejected in favor of this one-time, highly simple change that greatly reduces the need for future hard forks in this area.

This BIP accomplishes several goals:

G1: Demonstrates change is possible; the bitcoin protocol can be upgraded.

G2: Eliminate 1MB limit as impediment to adoption.
G3: Get hard fork risk out of the way early.
G4: KISS solution, in terms of code changes.
G5: Upgrade path, yet constrained until problem & solution better understood.

This introduces friction into the block size increase process - making it changeable, yet giving participants in the system sufficient time to observe system behavior and change course. Ultimately moving towards a system where the market decides the best block size.

Users - network node operators - exercise their voice twice: One hard fork at the introduction of BIP 100, and a 2nd hard fork at 32MB, assuming users choose to scale that high. The 32MB hard fork is largely coincidental - a whole-network upgrade at 32MB was likely needed anyway, for historical reasons unrelated to this proposal.

Discussion: Low signal, high noise

Major policy changes in an open source project are messy. What would, in a central bank, be a closed door policy discussion is instead held out in the open, with all parties airing their opinions in an open forum. Transparency wins, even if sometime a painful process.

Yet the amount of actual financial engineering involved in this debate has been disappointing. Simulations are practically non-existent. Data on the change has often lacked (or held constant) one or more key variables associated with today's network. The proponents of a block size change have sometimes been poor debaters, exaggerating the arguments of opponents to a ludicrous degree.

In toto, this is a major change to bitcoin. Changing a $3B economic system - even if to keep it alive and growing and healthy - deserves more in-depth attention and research than has been given to date. The 1MB limit is a relic from an earlier time, always intended to be removed - yet paradoxically it forms the core of our economics today and cannot be changed lightly or without consequence.

Discussion: Accepting engineering reality

In the drive to "growth at any cost" in some bitcoin business models, some axioms tend to get glossed over.

Bitcoin is fundamentally a settlement based design. The nature of distributed consensus, known in database circles as eventual consistency, is that it takes time for the network to come to a consensus on the "right" historical timeline (blockchain). The blockchain takes time to bake.

Answering some the existential questions opened above, bitcoin as engineered today will not be the Star Trek ideal we all want: instant, secure, borderless, egalitarian, non-volatile currency directly accessible by 7 billion.

That should not be misinterpreted as the Star Trek ideal being impossible. Rather, systems will have economic and engineering incentives to use the bitcoin blockchain for its strength as an *anchor*, the root chain in a forest of chains. Individual chains, thus secured by the "main" chain, will offer unique features such as Ethereum's smart contract language, or serve specific user communities.

Putting all the world's coffee transactions, and all the world's stock trades, and all the world's Internet-of-Things device samplings, on the bitcoin blockchain seems misguided. The block size limit does serve a useful function, discouraging wholesale dumping of multi-gigabyte files into the blockchain. Already "blockchain as my video dropbox" apps have been proposed, much to collective dismay.

A bitcoin transaction today is analogous to an irreversible, one-way message. In the world of computer networking, the lowest layers of networking are irreversible, one-way messages. Sometimes these messages, are lost, corrupted, need to be retransmitted or updated. IP, the Internet Protocol, requires additional, upper-layer protocols TCP and HTTP to be useful. Yet IP forms the foundational layer of the Internet.

Bitcoin protocol and network today is that foundational layer. It is value transfer network. Beyond that, it is a core, backbone security service securing contracts, physical and digital property, equities, bonds, robot A.I. and an enormous wave of applications which have not yet been conceived. Inventing bitcoin-the-currency, securing bitcoin-the-token, was the Minimum Viable Product, the necessary first step towards building a universe of secure, decentralized services.

Do not try to stuff every feature into the bitcoin protocol. Let it do what it does best. Build systems on top of bitcoin which use its strengths, the ultra-secure blockchain service & bitcoin ownership registration service (aka secure value transfer).

Discussion: Doesn't this just hand control to the miners?

This change moves from centralized management of the fee market to free-market control, improving bitcoin's governance by removing a hardcoded policy control from the software.

Let's review the fee market from a block size perspective. Consider a highly constrained block size of 200k - a speed limit below today's bitcoin traffic levels. Under such tight constraints, competition for block space is fierce, and fees will be high. Tight supply (assuming inelastic demand for the sake of example).

Consider the opposite case, a block size of 200 megabytes - a speed limit well above today's traffic levels. There is no competition for space at all. Fees will be zero (assuming for ex. no anti-spam filtering), as block space is always available.

This block size speed limit embeds **centralized economic policy planning** choices for transactions fees into the bitcoin software. That is bad news. BIP 100 proposes a solution, by slowly migrating this centralized policy out to the free market at large.

Consider three conflicting or opposing viewpoints, *all of which are equally valid* from their individual points-of-view as Rational Economic Actors:

1. Early Adopter: Do not increase 1MB speed limit. I am happy to pay high fees to secure my bitcoin. I make 1-2 transactions per year.
2. Cautious Miner: Only increase the 1MB speed limit a little. Enough for adoption, not enough to reduce my fee income.
3. Funded Startup: Scale bitcoin to VISA rates in 12 months. Keep fees near zero to subsidize adoption. On-board 1 billion users in 2 years. No speed limit.

Should a Chief Scientist or Technical Advisory Board be making such a fundamentally *economic, market-shaping* choice?

In 2015, in the block size debate, some Chinese firms complained that a too-large block size would put them at a disadvantage vis their western counterparts in US and EU. **The block size speed limit choice can disadvantage one nation before another**. Is an informal group of software developers best suited to choose?

When the phrase "picking winners and losers" is used in the block size debate, that refers to picking an arbitrary block size - 1MB, 2MB, 20MB - which, in turn, dictates the shape, economics and actor behavior of the bitcoin transaction fee market. This advantages some in the bitcoin market and disadvantages others in the bitcoin market (or discourages actors from entering in the first place). The bitcoin protocol's 1MB limit - added an ancient times to prevent garbage from overwhelming the system - has now unintentionally become a policy tool.

Changing this policy tool - block size limit - is seductively simple, a one-line change to the software source code. Trivial.

Who changes the software? Economic actors that wish to see the speed limit at X or Y - thus dictating the fee market - will lobby the Chief Scientist and other "core" developers, individually,

in private, in public, with carrots, and with sticks. When bitcoin market cap achieves 10x or more, the lobbying is even more intense. Yet there is no single human or commitment on the planet capable of picking a good speed limit.

Who chooses the level of decentralization? Some of the smartest folks in the bitcoin block size debate highlight the link between block size and decentralization, which directly feeds back into bitcoin's core security value. If mining actors send a dynamic block limit to 1GB, forcing all-but-3 actors out of bitcoin, the resulting system cannot be described as decentralized. The speed limit acts to contravene that centralizing factor. At lower block sizes, more are able to participate in the bitcoin network, sharing the load and reducing centralization - and thus potential manipulation. At extremely high block sizes, very few are able to directly participate on the bitcoin network, the opposite of a censorship resistant, egalitarian, neutral "backed by math" actor lauded by so many.

There is no right answer for a speed limit which acts as a Magic Lever of Decentralization. One level of decentralization for one market actor - versus a sought gain in bitcoin users and transaction volume - will be different from another. Software - and software developers - should not be taking that choice out of the hands of users. Today, we make that choice for you, by virtue of the block size speed limit.

A key goal, therefore, is to transition the speed limit from software control to market control. Remove the policy wedge. Let the free market decide the long term shape of bitcoin's transaction fee market, level of security and level of decentralization.

Miner voting was chosen for BIP 100 as a "lesser of the evils" Stakeholder voting is appealing, but the devil is in the details (what if they're not online? lacks bias for active users? biased towards early adopters? miner self-dealing?). Sampling validating nodes - a key rung in the bitcoin security ladder, validating the work of miners - remains difficult. Measuring, and perhaps compensating, these key actors to remain in the network is even more difficult. All of these factors are subject to various methods of gaming, such that the metric sought - "what speed limit do bitcoin users want?" - is never truly discovered.

Dynamic algorithms for block size are appealing. The 5+ year solution for block size speed limit may indeed be a dynamic algorithm. Having a dynamic rather than static speed limit is nevertheless still an active economic policy choice, with all the developer lobbying and subjective judgements and coercions that entails. There remains no one right answer in the market tension between decentralization, security and user/volume growth.

Ultimately, the block size is dictated by market actors, as bound by the laws of physics and the inefficiencies of software. Miner voting is viewed as a temporary, slow, friction-ful method of transitioning away from the block size speed limit being chosen by software (and software developers).

The threshold of change is intentionally months long. Voting occurs in public. Collusions, cartels and attacks via/on an overwhelming super-majority of hashpower must be sustained for similar lengths of time. Bitcoin users have months to plan and act in the markets themselves, possibly re-instituting an absolute limit via hard fork if it comes to that.

This does not give the miners “complete power.” Miner income is in the long run aligned with users in a cooperative market relationship, even if short term counter-incentives exist. Users signal economically via the market - or directly via email and social media! - their views on miner behavior - which must be signalled up front, months in advance of any change. While imperfect, miner voting works; it has been field tested (BIP 34).

BIP 100 is a slow, paced transition to free market control.

Open challenge: Decentralized governance

Many of these questions would seem to have answers if only one can “ask the users” This is easier said than done. Large bitcoin stakeholders may be offline for months, years. Quantifying users and enabling voting are a challenge. Secret cartels easily undermine public votes.

Open source projects are by their nature opt-in, join and contribute as your energy, funds and time permit. There is no natural constituency to handle marketing, governance, testing, as there is with a company building a software product. All those tasks must be collectively organized by the bitcoin community. Free riding is easy at multiple levels. The entire community must be proactive in contributing time and resources to upkeep of this “automaton”

How can we create a technical advisory board, potentially enshrining its members as a bitcoin priesthood, without the board becoming as corrupt and easy to manipulate as FIFA? Technical decisions have economic consequences in bitcoin. Imagine the economic impact of code changes at 10x, 100x+ scale.

Ultimately I have confidence that the free market, transparency and user choice will sort it all out. Bitcoin is as much about the humans that use the network, as it is about the technology. If the technology suddenly breaks, users have over $3B reasons to work towards fixing the problem.

Full Disclosure

The preceding text is my own personal, vendor neutral opinion. I have many formal and informal associations with bitcoin businesses, notably BitPay, most of whom support a block size increase. Those associations are held separate from this opinion.