author: steffenkd
date: 04.06.2026
wordcount: ~ 3300
reading time: ~ 25 minutes

The decentralization in Bitcoin

The word “decentralization” is inflationary used in the so called bitcoin and crypto space and in my opinion is one of the most misunderstood concepts in the whole space.
You may not be aware, but Satoshi himself hasn’t used the word “decentralization” even once in his whitepaper.
Since another mantra in the space is “don’t trust, verify”, I recommend you read The Bitcoin Whitepaper and verify my statement yourself.
Even though I admit, that the concept of “decentralization” can be read between the lines.
So maybe, just maybe, this concept of decentralization has a somehow different meaning as it is propagated by the majority of the so called bitcoin and crypto influencers for the last decade.
In this article I try to explain what I think the concept of “decentralization” in bitcoin means.
And where its boundaries and limits are.

Decentralization in the mathematical sense

Centralization means that something is centered in one point.
Therefore decentralization means that something is not centered in one point but exists in at least two points.
Mathematically “central” can be expressed through the number one (1).
This automatically means, that every number greater than one (n>1) is decentralized.
For the game theoretical nature of bitcoin to play out longterm, you would need at least three (3) nodes though.
Which would set the minimal amount of nodes for bitcoin to work properly longterm at three nodes (3).

To be clear: the original bitcoin is still being rolled out and during bootstrapping even one or two nodes would suffice for some time.
But this would be a discussion of its own.

So you could write a small program which compares if the amount of nodes is equal or greater than the number two (n != 2) and returns a boolean value (0 or 1) which means its either true or false.
It is either centralized or decentralized, simple as that.

Now you could write another program which checks if several values are greater than two (n>2) and order those values either by highest to lowest or lowest to highest.
This program could check if a value is more decentralized than another.
If we would put in the numbers 144, 2, 69, 88, 369, 3, 2016 and 1845 the program would order them in the sequence 2, 3, 69, 88, 144, 369, 1845, 2016, with 2016 being more centralized than 2.

This means something is decentralized, when it is larger than the number one and that you can have different levels of decentralization.

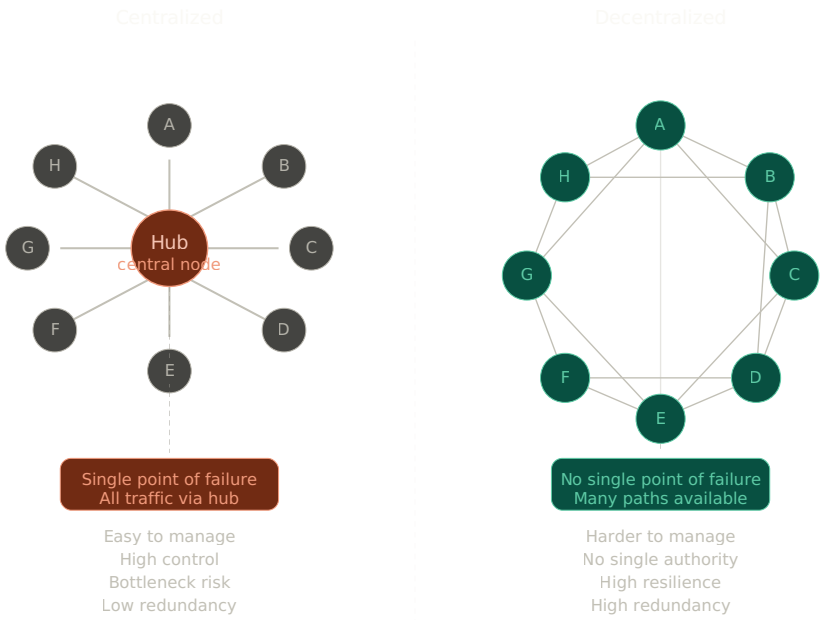

centralization_vs_decentralization

Picture: The two structures represent opposite philosophies.
In a centralized system, every node speaks only to one hub.
Simple to control and monitor.
But that one hub is both the bottleneck, the single point of failure and the single point of potential corruption.
Take this one point down or corrupt it and the whole network

collapses or can be abused.
In a decentralized network, every node can reach others via multiple paths.
Removing or corrupting any single node barely disrupts the system and the traffic just reroutes.
The trade-off is complexity: there's no one place to manage, influence, corrupt or audit everything.

The advantage and utility of decentralization

So decentralization means, that something exists in more than one place or in case of bitcoin, in more than two places.
This has certain effects, and introduces redundancy and backup capabilities into a system.
In a decentralized system you do not have one single point of failure, or, in case of power and control, not one single point of potential corruption.
Think about it with regards to politics, advertisement, big banks, big tech, big government, big media, lobbying and the extensive spending of tax payer money.
And about how Jeffrey Epstein has managed to corrupt many people in those positions of power.
Those are the advantages of decentralization.
You do not have one large pot of tax payers money where you only have to corrupt a few politicians or regulators to sell your useless product to the masses any longer.
Instead you would have to go from door to door to convince every citizen about the effectiveness and utility of your product, ideology or narrative.
It would be quite interesting for example, to see how many vaccinations would be sold in such a system.
Or if the citizens of two nations would actually vote for war.
I personally would bet against it!
The so called "representative democratic systems" world wide are nothing else than "man in the middle attacks" against humanity itself.
Compared to an entrepreneur in the economic system, where I have a claim of compensation or warranty if a false or broken product was delivered, a representative has close to no obligation to actually do what he has promised.
Which bares the question why why need those representatives at all.
Especially when having a peer-to-peer-electronic-cash-system like bitcoin which eliminates middlemen and intermediaries.

At some point the majority of people world wide may realize that they are being played, exploited and incited against each other - by big media.
People on both sides may vote and wish for peace and cooperation but they get war at the end.
Because lobbies in the background are corrupting those centralized points of power, namely politicians.
Instead of having to corrupt 80 million people in germany they only have to go for 630 politicians - actually only for half of them.
And voila, the armaments industry can squeeze every little drop out of the tax payers on both sides of the conflict.
Citizens on both sides suffer and a few powerful people in the background get massive profits.

Be aware, that you have to invest more energy and work to run and coordinate those different servers or miners though.
But the current big tech enterprises around the world already have decentralized their internal server structures.
Netflix for example doesn't have just one server where every human being on earth streams his movies from.
I do not know how many they have, but they probably already have a lot, since decentralization makes also sense from a "data delivery perspective".
At the moment we have google, netflix, amazon, microsoft, spotify, x and several other big tech enterprises who each runs several large server farms around the world.
Additionally a centralized entity like a bank can keep certain data secret, like the accounting of money or if it happens in their favour.

A short quote from Lord Acton:
> "Power tends to corrupt and absolute power corrupts absolutely."

And while we are on it, another one from Leopold Kohr's "Breakdown of Nations":
> "Even a confirmed thief will not steal if he has no chance of getting away with it.
On the other hand, even an honest man will misbehave if he has the opportunity, the power to do so."

And since I mentioned media I will bring in Malcolm X as well:
> "The media's the most powerful entity on earth.
They have the power to make the innocent guilty and to make the guilty innocent, and that's power.
Because they control the minds of the masses."

If you split up those centralized positions of power, which are just centralized positions of potential corruption you may end up with less temptation, less corruption, less misinformation, less money wasting and less resource misallocation.
You would reintroduce the law of supply and demand, the free market principle, the wisdom of crowds the gaussian normal distribution and the law of large numbers.
But more about those concepts in another article named “why bitcoin?”.

What Satoshi Nakamoto had to say about it

> The current system where every user is a network node is not the intended configuration for large scale. That would be like every Usenet user runs their own NNTP server. The design supports letting users just be users. The more burden it is to run a node, the fewer nodes there will be. Those few nodes will be big server farms. The rest will be client nodes that only do transactions and don’t generate..

> The existing Visa credit card network processes about 15 million Internet purchases per day worldwide. Bitcoin can already scale much larger than that with existing hardware for a fraction of the cost. It never really hits a scale ceiling. If you’re interested, I can go over the ways it would cope with extreme size..

Satoshi Nakamoto created the name bitcoin and defined its design and technical capabilities in the bitcoin whitepaper.
The name bitcoin therefore is linked to the design described in the whitepaper.
Bitcoin was released under a MIT-license which allows copying of the software and overall is very generous with using the source code.
You can essentially do whatever you want with it.
Crippling the original design and then passing it off to the masses as the original bitcoin brand is bad sportsmanship though.
And time will show since the truth can stand up straight on its own.

The essence of decentralization

The essence of decentralization in bitcoin is that auditors are auditing each other.
By that you have redundancy and backup capabilities in case of data loss.

And whoever makes a mistake or wants to cheat and enrich himself will get caught by the other auditors and users and will lose his initial investment.
Endless decentralization, which essentially is decentralization for the sake of decentralization is a questionable narrative though.
Backup capabilities and redundancy are useful.
But at a certain amount of backups the additional work and energy needed for coordination and keeping them up to date and in sync increases, while the benefits decline.

An example:
One data backup backup will add 100% utility to your case.
If your original data gets corrupted you have a copy to keep going - quite obvious and logical.
But a second backup may only add a 50% increase of utility already, since it is unlikely that the original and the backup fail at the same time.
A third backup will only add 25% utility, because the probability of the original data and the two backups failing at the same time decreases further.
Maybe the function is 100%, 99%, 98% and so on, maybe it is 99.99%, 66.66%, 33.33% and so on, but that is not the point.
The point is, that every additional backup is less useful than the previous one.
Meanwhile every additional update adds a little amount of time, energy or work needed for coordination and keeping it in sync.
If you backup your data onto a hard drive regularly it will be obvious that connecting and synchronizing a second hard drive takes more time, energy and overall work than syncing just one hard drive.
You may have automated the process, but this simply means that the computer is doing the work for you.
You may not see it, but the work is there.
It still costs computing power, bandwidth, hardware, energy and time to keep those backups in sync.
The lesson here is, that a potential infinite amount of decentralization, redundancy and backups becomes uneconomical at a certain point.
But where should be the threshold and who decides where the threshold should be?
You may be familiar with the value-cost-ratio or benefit-cost-ratio, where two functions are being laid over each other and the intersection of the graphs is where the sweet spot lies.
Above or before this intersection point it is profit, below or after this intersection point it is loss (or vice versa, depending on how you display the data and graphs).

With bitcoin miners it is similar.
A bitcoin miner simply has to calculate his “cost-benefit ratio”.
The metrics and data you need for calculation are your expenses and your expected returns.
That’s it.

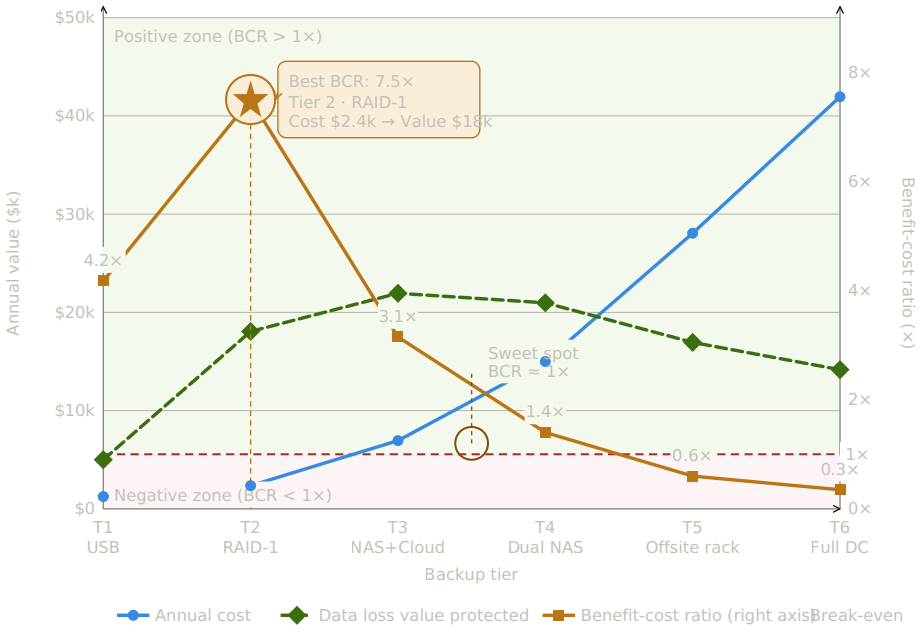

benefit_cost_ratio

Picture:
The best “benefit-cost-ratio” is marked with a star at Tier 2 (RAID-1), where $2.4k annual costs protect $18k of data value, yielding a 7.5× return.

The two kinds of decentralization in bitcoin

The misunderstood decentralization on the miner or server side - 2016 maximum

2016 is the maximum amount of blocks which can be mined during the two weeks “difficulty adjustment period”, after which the difficulty adjustment algorithm readjusts the difficulty.
After this period a bitcoin miner has to recalculate his benefit-cost-ratio.
This is how a profitable business is run - you supply a service, account for expenses and add a profit on top.
If a miner or node doesn’t find a block during this two week “difficulty

adjustment period” he essentially can’t get a return on his initial investment and therefore is a non profit enterprise.
His uptime and run conditions are not based on delivering a profitable service paid by the users.
If you are running an enterprise which supplies a service and has a cashflow with inputs and outputs, you are more embedded in the entangled economical network than someone who doesn’t.
This is not necessarily bad, but beware that profitable businesses will more likely attract entrepreneurs who want to run them, due to their profitable nature based on demand and supply.
In comparison, a non profit enterprise like a “non mining node” on small block BTC may attract less participants, due to its cash bleeding nature.
So the supply of non mining nodes on small block BTC is solely based on the emotions of their hosts and not on profit oriented market principles.
This can change any time due to a shift in sentiment or ideology.

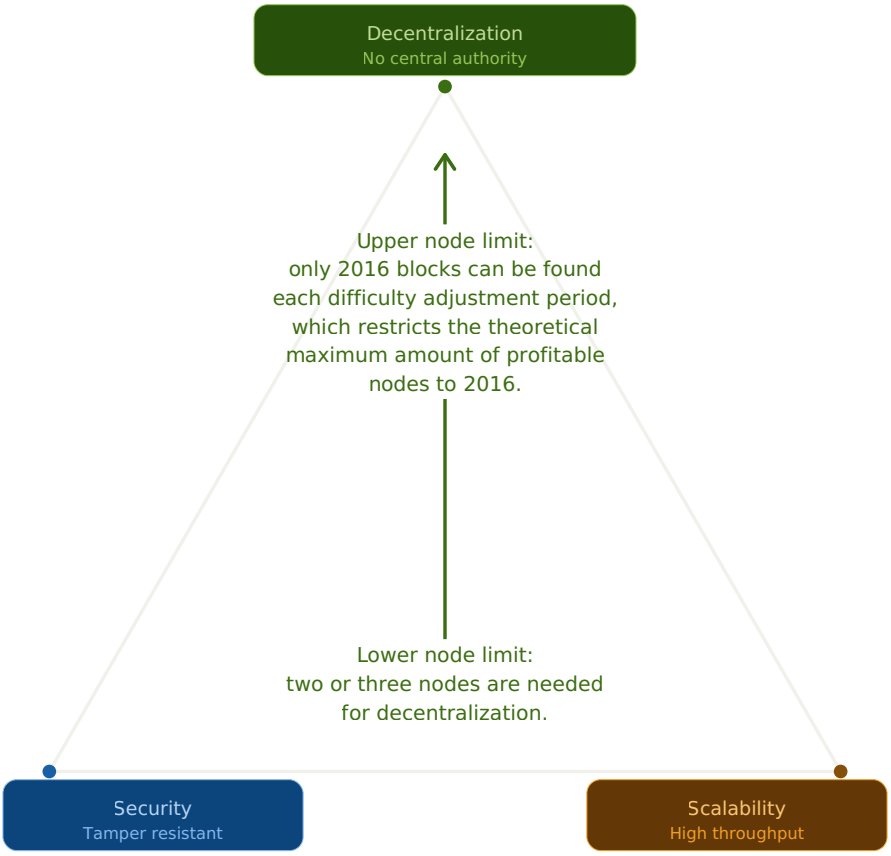

decentralization_3-2016

Picture:
Theoretically, only 2016 blocks can be found during the two week “difficulty adjustment period”, which therefore serves as the upper limit for economical nodes on the bitcoin network.

And if you do not find a block, you are not a miner.
The same way you are no car builder if you can’t produce a driving car, you are no marathon runner if you can’t finish a marathon and you are no chess player if you can’t finish a chess party.
If you enter a competition you agree to certain rules and in case of bitcoin the competition is about finding a block.
Therefore the amount of decentralization in bitcoin mining lies somewhere between 3 - 2016.
2016 nodes therefore is the theoretical upper limit for decentralization

on the bitcoin network.
In my opinion more than 100 nodes are highly unlikely due to the pareto-principle and the power-law though.

The ignored decentralization on the user side - infinite/unbounded

Like already stated in "what is bitcoin?" users which want to participate on the bitcoin network can create their user accounts themselves.
I highly recommend to read the article or at least the passage before reading on.
This account creation can be done due to the properties of public key cryptography.
By creating a pair of so called asynchronous keys a user is able to create his own identity without the need of centralized entities like big tech, big government or big banks.
Users then can verify each others identities and create a network of trust.
Since you are not dependent on centralized entities like google, meta, x, banks, governments or institutions alike, you essentially have a "decentralization of identity creation on the user side".
This "decentralization of identity creation on the user side" makes most sense, if users can write to the bitcoin blockchain database.
Otherwise it would be like everyone of the seven billion people one earth being able to build an electric water heater with two copper rods and a bucket, but there is only electricity for about 500000 people each day.
With around seven to eight billion people on earth this is laughable.
You would just have another PGP or GPG implementation.
Which is not necessarily bad, but PGP and GPG already exist and I do not think we need another copycat.
It already isn't used by a lot of people.
And Big Tech is already using asynchronous key pairs, but instead of you creating and managing your own keys on your device, big tech is doing it for you - the irony.

On small-block-BTC it would take around 40 years till everyone of the 7 billion people on earth has written to the bitcoin blockchain, because small-block-BTC is only capable of five transactions per second.
And it doesn't really matter if it is 30 years, 40 years or 45 years - everything above one day is impractical.
So on small-block-BTC you have uneconomical nodes, which do not

find any blocks and can't break even on their costs.
And at the same time the amount of transactions is so low, that the "decentralization of identity creation on the user side" is useless for 99.99% of users at scale.
99.99% of people who created their identity themselves won't be able to use their newly created identity on the small-block-BTC blockchain.
Or in small blocker language: You can be your own bank but you won't be able to transact.
Why should you run a node on a raspberry-pi, if you do not find any blocks and therefore can't break even on your costs?
And on top of that you won't be able to to transact on the network you are running a listening node for, because by design only 5 people per second can make a transaction.
Serious question: Why should anyone do this and where are the incentives?
I would consider this a double fail.
The good news is, that you can use your identity on all three major bitcoin versions, namely BTC, BCH and BSV.
Two of those bitcoin versions are restricted, one is unbounded.
You can use the sovereignity of your asynchronous key pair identity and decide independent and autonomous which bitcoin version you want to use.

The balance between user and server side decentralization

On an unbounded bitcoin version, the miners simply supply the demanded amount of transactions by the users, based on the universal law of demand and supply.
Only restricted by technical limitations like Moore's Law.
On the bitcoin versions with a restricted blocksize, the user-side-decentralization can't unfold its full potential.
Again: you can be your own bank but you won't be bale to transact.
Small block BTC is not peer-to-peer but peer-to-miner-to-peer.
And on Lightning you can't sign a utxo over to someone else, but you have to search for a path through intermediaries and middlemen again.
Those are no innovations, those are steps backward.
But more about that in another article.

Three different versions: BTC, BCH and BSV

Bitcoin version	People (n)	tx/s	time for 1 billion transactions
BTC	1 billion	5	> 6 years
BCH	1 billion	120	100 days
BSV	1 billion	1000000	17 minutes

On the BitcoinSV blockchain the blocksize is unbounded.
This means you have a potential network effect which can include every human being.
Since BitcoinSV is able to process around one million transactions per second, it means that it would take around 17 minutes for the network to process one billion transactions.

On BCH you have a blocksize limit of 32MB which means that you will end up with around 120 transactions per second.
This means the BCH blockchain has a capability of 10 million transactions per day.
With 1 billion people wanting to make a transaction, it would take the network around 100 days to process all the transactions.

On BTC you have a blocksize limit of 1MB (4MB blockweight), which means, that the network can only process around 5 transactions per second which are 432000 transactions per day.
With 1 billion people wanting to make a transaction, it would take the network over 6 years to process one billion transactions.
Self sovereignity and self custody is therefore not a realistic scenario.
It exists in theory, but in praxis a majority (99.99%) of people will never be able to transact onchain.

Thanks for reading!