BitMEX
Research

Complete Guide to Proof of Stake

11 April 2018

Abstract

In this piece we examine proof of stake (PoS) consensus systems. We look at their theoretical advantages and weaknesses. We then analyse the specific details of some of the most prominent and novel PoS systems attempted thus far, where we learnt that some pure PoS systems becomes increasingly complex, to the point which they became unrealistic. We review the latest Ethereum proposal, which we think is a significant improvement compared to previous attempts and it could provide net security benefits for the Ethereum network. However, the system may still be reliant on proof of work (PoW), which is still used to produce the blocks and at this point it is not entirely clear to us if the PoS element of the process contributes to ensuring nodes converge on one chain.

BitMEX Research

Filtering out the hype with unbiased, evidence-based reports on the crypto-coin ecosystem.

BitMEX Research is also active on Twitter and Reddit.

research.bitmex.com

Previous reports:

Bitcoin price correlation: Record high against the S&P 500
(28/03/18)

Update: SegWit transaction capacity increase compared to Bitcoin Cash
(22/03/18)

Tether: New financial data released by Puerto Rico
(19/03/18)

Diagram of a Bitcoin block: Covert versus overt AsicBoost
(01/03/18)

Tether
(18/02/18)

Introduction

Before diving into the specifics of Proof of Stake (PoS), it's important to clarify what one is trying to achieve when building these consensus systems. Essentially one is trying to construct a data structure with the following properties:

1. No one entity controls the content of the data (distributed storage and verification of the data is not sufficient);
2. The database can move forward, (Casper terminology: "Liveness"); and crucially
3. Participants agree on the content of the data i.e. nodes have a mechanism to decide between conflicting valid chains (Casper terminology: "Safety")

PoW uses the most accumulated work rule to decide between competing valid chains (fork choice rule). This is not only an apparent solution to criteria three above, but the PoW mechanism also inherently solves the block production and block timing issue. While total accumulated work is the fork choice rule, a block producer is also required to include an element of PoW in each block, a stochastic process, and therefore the issue of who produces each block and when each block is produced, is also be addressed by PoW.

PoS is the general concept of a fork choice rule based on the most accumulated stake (i.e. the chain with the most coins backing, voting or betting on it). However, unlike PoW, this does not necessarily directly address the issue of who produces each block or when blocks are produced. Therefore these issues may need to be addressed by alternative mechanisms. PoW is also a solution to the coin distribution problem, something which may also require an alternative solution in PoS based systems.

Theoretical overview of PoS

The byzantine generals problem

The Byzantine generals problem illustrates some of the main challenges involved when attempting to construct a data structure with the properties mentioned above. Essentially the issue is about timing and how to determine which updates to the ledger occurred first. Actually if one third or more of the actors are disruptive, the problem is provably unsolvable, from a mathematical standpoint, as Leslie Lamport proved in 1982.

> "It is shown that, using only oral messages, [reaching agreement] is solvable if and only if more than two-thirds of the generals are loyal; so a single traitor can confound two loyal generals"

(Source: The Byzantine Generals Problem (1982))

PoW can therefore be considered as an imperfect hack, which seems a reasonably strong Byzantine fault tolerant system, but certainly not a mathematically robust one. It is in this context, of imperfect systems, which one should analyse PoS alternatives, as like PoW, these systems will also have flaws.

In PoS there are two competing philosophies. One of which is derived from PoW. Coins based on this include Peercoin, Blackcoin and earlier iterations of Ethereum's PoS proposals. The second philosophy, is based more on Lamport's academic research from the 1980s and embraces the conclusion Lamport reached that a two-thirds majority is required to build a Byzantine fault tolerant system. Ethereum's current iteration of the Casper proposal adopts this second approach.

Advantages of PoS

PoS is typically looked at in the context of PoW, as an alternative which solves or mitigates against negative externalities or problems inherent in PoW based systems:

<u>More environmentally friendly</u>

Perhaps the most widely cited advantage of PoS systems is the absence of the energy intensive process which PoW requires. If PoS based systems can achieve the same useful characteristics as PoW systems, environmental damage can be avoided. This is a significant positive for PoS, although as we discussed in our piece on Bitcoin's energy consumption, the problem may be slightly overstated, due to the incentive to use lower cost or otherwise failed energy projects as a source of power, limiting environmental damage.

Stronger alignment of incentives

Another major problem with PoW based systems is that the interest of miners may not align with that of coin holders, for example miners could sell the coins they mine and then only care about the short term, not long term coin value. Another issue is that hashrate could be leased, with the lesee having little or no economic interest in the long term prospects of the system. PoS directly ties the consensus agents to an investment in the coin, theoretically aligning interests between investors and consensus agents.

Mining centralisation & ASICs

Another key advantage of PoS based systems is potentially improving decentralisation. PoW mining has a number of centralising forces which are not applicable to PoS:

- ASIC production is expensive and centralised (In Bitcoin Bitmain has a high market share);
- Chip foundries are expensive and centralised (TSMC, Intel, Samsung & SMIC are the only players with scale);
- ASIC related technologies can potentially be patented;
- There may be a limited number of cheap energy sources, with restricted access; and
- Many aspects of mining can have economies of scale, such as maintenance costs and energy costs, resulting in centralisation.

General and economic weaknesses of PoS

An incomplete solution

As we alluded to above, Satoshi's PoW systems appears to kill four birds with one stone:

- Chain selection (the fork choice rule),
- Coin distribution,
- Who produces blocks, and
- When blocks are produced.

PoS only appears to be a proposed solution to the chain selection problem, leaving the other problems open. Although these other issues could be less significant than the chain selection issue.

An “unfair” economic model

One of the most common criticisms of PoS systems is that they allocate new funds in proportion to the existing holdings. Therefore the “rich get richer” and it results in a few wealthy users holding a higher proportion of the wealth than the more egalitarian PoW alternative. If one invests in a PoS system at the start, you can maintain your share of the wealth, alternatively in a PoW system your wealth is diluted as new rewards are distributed to miners. Indeed, if rewards are allocated in proportion to the existing holdings, one could argue its not inflation at all and that the reward is economically equivalent to adding more zeros to the currency. Therefore one can even claim the reward system is pointless and does not provide an incentive at all. However this only applies if all users become PoS validators, while in reality some users will want to use the funds for other purposes.

Risk of a loss of funds

Another issue is that staking requires signing a message from a system connected to the internet. Therefore stakers are required to have a “hot wallet” which increases the risk that funds are exposed to theft from hackers. Although it may be possible to mitigate this downside by having a private key only entitled to stake for a short period of time, after which the balance reverts back to the owner. Although if there is a slashing rule (punishment for voting on two conflicting chains), a hacker could conduct action which destroys the funds even if this mitigation strategy is used. Another potential mitigation strategy could be the creation of specialist hardware for staking.

Technical & convergence weaknesses of PoS

Nothing at stake

Core to the consensus problem is timing and the order of transactions. If two blocks are produced at the same time, PoW solves the problem by a random process, whichever block is built on top of first can take the lead and then miners are incentivised to build on the most work chain. PoW requires energy, a finite real world resource and therefore miners have to decide which chain to allocate this resource to.

In contrast this process in PoS based systems is not entirely clear. If two blocks are produced at the same time, each conflicting block can build up stake. Eventually one block may have more stake than the other, which could make it the winner. The problem here is that if stakers are allowed to change their mind to back the winner, such that the system converges on one chain, why would they not use their stake on multiple chains?

After-all stake is a resource inherent to the chain and not linked to the real world, therefore the same stake can be used on two conflicting chains. Herein lies the so called "Nothing at stake" problem, which we view as the most significant issue facing PoS.

The "Nothing at Stake" problem

The Nothing at Stake problem

Stake does not add to the convergence of the system, since the same stake can be applied to multiple competing chains, which is a risk free way of stakers increasing their rewards. In contrast, in PoW based systems, energy is a real world finite resource and therefore the "same" work cannot be applied to multiple competing chains.

Defense 1

The issue can be avoided or mitigated against. The protocol can be adjusted such that if a staker uses the same stake on multiple chains, a third party can submit a proof of this to either chain, resulting in a punishment, such as the confiscation of the stake (slashing conditions). Alternatively instead of a punishment, the cheater could lose potential rewards or be excluded from the staker pool.

Response from PoS sceptic

The above defence is inappropriate and punishes what may be legitimate or necessary behavior. For example if a staker receives a block first, while the majority receives an alternative block first, it may be legitimate for that staker to change their mind and switch to follow the majority. Indeed the process of changing your mind and switching to the majority to ensure the network converges is the point of the consensus system. If this behavior is punished, how does the system converge?

Either the economic value of the punishment is higher than the rewards for switching to follow the majority, or it isn't. Therefore the nothing at stake problem means PoS systems can never contribute to system convergence and the idea is therefore fundamentally flawed.

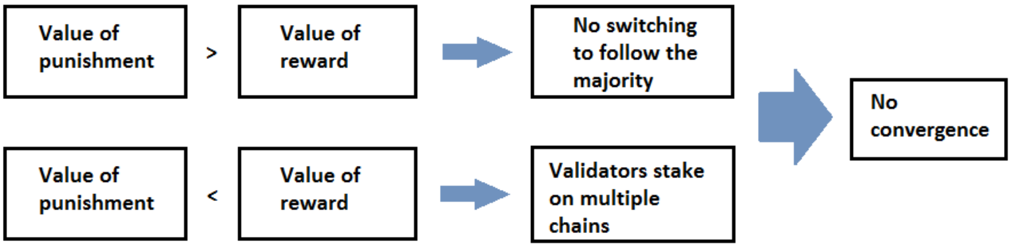

Defence 2

The apparent dilemma above can potentially be resolved in various ways. For instance:

- Earlier proposals from Casper used multiple rounds of staking. Changing one's mind in the early rounds can be legitimate and perhaps the punishment is small, while in later rounds the punishment for using the same stake in multiple competing chains increases, such that eventually users have a high degree of assurance over the finality of the system.

	• The most recent iteration of Casper aims to allow validators to change their minds, but only in “legitimate” scenarios and not when its “illegitimate”.
Response from PoS sceptic	By adding multiple rounds or criteria in which validators can change their minds one is increasing the complexity of the system. This is merely adding layers of obfuscation to conceal the inherent weaknesses illustrated by the nothing at stake problem, without solving the fundamental issue.
Defence 3	No system is perfect, indeed it’s mathematically impossible to construct a perfect system and therefore the nothing at stake problem is not solved, however the measures identified above mitigate the problem, such that these theoretical issues are unlikely to apply in the real world.

The long range attack consensus problem

Another potential issue with PoS is the so called “long range attack” problem. This is the idea that attackers could, for instance, buy a private key which had a large token balance in the past and then generate an alternative history from that point, awarding oneself more and more rewards based on PoS validation. Due to the large amount of rewards given to the attacker, one could then generate a higher stake chain than the existing chain and a large multi year chain re-organisation could be performed.

The solution to this problem is checkpointing, which is the process of locking in a certain chain state once a certain stake threshold has been met, such that it can never be re-organised. Critics argue that this solution requires one to keep their node online at all times, since an offline node cannot checkpoint. Some claim that if one goes offline, the security model therefore degenerates to “ask a friend”, since one is dependent on asking others for their checkpoints. Although in the past the Bitcoin reference implementation included checkpoints, the purpose of these was to speed up the initial sync, although the impact of this could be said to result in an “ask a friend” security model.

However, in our view this is a matter of different priorities. If one wants each individual user to fully verify all the rules and the state of the system, then relying in these checkpoints is insufficient. Indeed, the Satoshi’s original vision appears to imply that the ability of nodes to be switched off and then verify what happened when was gone is potentially important:

> “Nodes can leave and rejoin the network at will, accepting the proof-of-work chain as proof of what happened while they were gone”

(Source: Bitcoin Whitepaper)

Although the ecosystem is expanding, many businesses and exchanges operate 24×7 and are therefore required to keep a node running all the time, and can therefore do checkpointing. There are strong incentives preventing them from allowing a large chain re-organisation. To many, this is sufficient security and the risks posed by the long range attack problem are therefore irrelevant or too theoretical.

Stake grinding

In a pure PoS system, stakers also need to produce blocks. These systems have often worked by selecting a sequence of authorised block producers randomly from a pool, where the probability is proportional to the stake. The issue here is a source of randomness is required inside the consensus system. If the blocks themselves are used for generating the entropy, stakers could try to manipulate the content in blocks in order to allocate themselves future blocks. Stakers may then need more and more computing power to try more and more alternative blocks, until they are allocated a future block. This then essentially results in a PoW system.

In our view, the stake grinding problem is less of a fundamental problem with PoS, when compared to significant issues like the nothing at stake problem. All that is required to solve this problem is a source of entropy in the network and perhaps an Ethereum smart contract like the RanDAO, in which anyone can participate, can solve this problem.

Case studies – Peercoin & Ethereum's Casper

1 – Peercoin – 2012

Peercoin is a hybrid PoW and PoS system, built on the idea of coin age. The fork choice rules is the blockchain with highest total consumed coin age.

> "Coin age is simply defined as currency amount times holding period. In a simple to understand example, if Bob received 10 coins from Alice and held it for 90 days, we say that Bob has accumulated 900 coin-days of coin age"

(Source: Peercoin Whitepaper)

In Peercoin, some blocks were produced purely using PoW, whilst other blocks were produced using PoW where the difficulty adjusts based on the coin age destroyed by the miner in the transaction (the coinstake transaction as opposed to a coinbase transaction). "For example, if Bob has a wallet-output which accumulated 100 coin-years and expects it to generate a [PoS block] in 2 days, then Alice can roughly expect her 200 coin-year wallet-output to generate a [PoS block] in 1 day.

Analysis

Weakness	Summary
Nothing at Stake	The protocol aims to prevent miners using the same coins in a coinstake transaction on multiple chains by ignoring the second conflicting chain. However this is not sufficient and can result in nodes diverging, if they receive the conflicting blocks in a different order.
Block production	Solved by using PoW to produce the blocks
Long range attack	This was a critical vulnerability for Peercoin, an attacker can simply save up coin age by not spending their coins and then launch a re-organisation attack. This was solved by centrally broadcasting checkpoints several times a day. Peercoin was therefore a centralised system.
Stake grinding	This may not have been an issue, since there was no selection from a validator pool as PoW was always required and coin stake altered the PoW target.

Conclusion

At the time Peercoin was an interesting early novel approach, however the proposal resulted in a centralised system, not able to match the properties of PoW.

2 – Ethereum – Casper full PoS system – 2015

This is a full PoS proposal, based on “consensus by bet” methodology.

- Blocks are produced from a pool of block producers, a random number generator is used to select whose turn it is to produce a block and then the producer is given a time window in which they can produce a valid block.
- There is a set of bonded validators, one must be in the set to make or take bets on blocks.
- Validators can then make or take bets on block propositions, providing a probability each time, representing the return betters can make.
- After several rounds of betting, as the probability approaches 1 or 99%, the block is considered final.

(Source: Ethereum Blog)

Betting strategy

According to the Ethereum blog, betting should occur using the following strategies by default:

- “If the block is not yet present, but the current time is still very close to the time that the block should have been published, bet 0.5.
- If the block is not yet present, but a long time has already passed since the block should have been published, bet 0.3.
- If the block is present, and it arrived on time, bet 0.7.
- If the block is present, but it arrived either far too early or far too late, bet 0.3.
- Some randomness is added in order to help prevent “stuck” scenarios, but the basic principle remains the same.”

The default betting strategy had a formula (given below), to push the probability away from 0.5, such that the chain would move forward, with the probability expecting to either approach zero or one.

> “Let e(x) be a function that makes x more “extreme”, ie. pushes the value away from 0.5 and toward 1. A simple example is the piecewise function e(x) = 0.5 + x / 2 if x > 0.5 else x / 2”

If a validator bets when the probability is 99%, the return is very small (a 1% return used as a measure from which the reward is calculated), in contrast a winning bet placed placed with odds of 0.5, represents a return of 100%, which results in a higher return from the rewards pool.

The fork choice rule then is the sum of all the weighted probabilities, which have crossed a certain threshold, say 0.99. For instance a chain of five blocks, each with a probability of 1 will represent a score of 5. Any validator who changes their mind after the 0.99 threshold has been crossed, can be punished (slashed) for staking on multiple chains. While changing your mind before the threshold is considered legitimate and there is no punishment in that scenario.

Analysis

In our view, this proposal is highly complex, which we consider as the main downside.

Weakness	Summary
Nothing at Stake	The protocol aims to prevent miners using the same coins to bet on multiple chains by using a punishment mechanism, in which validators would lose their deposit. In our view, this could harm the convergence of the system, although betting formula may move the probability away from 0.5, which is designed to help mitigate the issue.
Block production	The RanDAO contract could be used to provide entropy to select the block producer. However, this only provides a time window in which blocks could be produced, it is possible there is a lack of consensus over whether the block was produced within the time window or not, after which the betting process is supposed to resolve the dispute.
Long range attack	The nodes checkpoint blocks once a certain probability threshold has been reached. The long range attack problem remains for periods in which nodes are switched off.
Stake grinding	The RanDAO contract may solve the stake grinding issue

Conclusion

The proposal was not adopted by Ethereum. In our view the proposal was never complete, as some parameters and aspects of the system lacked a specification. Although the consensus by bet approach was interesting, it seemed too complex and there were too many uncertainties. This approach illustrates the difficulties involved when constructing full PoS systems and how when one tries to address the weaknesses, it just results in more and more complexity, until the system becomes unfeasible.

3 – Ethereum – Latest version of Casper – The hybrid PoW/PoS System – 2018

The current Casper proposal represents a change in philosophy or a pivot, compared to some of the earlier PoS systems. It returns to the academic work of Lamport in the 1980s and Lamport's theorem that these systems work if and only if two-thirds of agents in the system are honest. Therefore the current version of Casper is less ambitious than before. PoS is no longer used to produce blocks or decide on the timing of blocks, which is still done by PoW miners. The PoS system is used as a checkpointing process. In our view, this proposal is superior to the more complex earlier iterations of Casper.

The system works as follows:

- The PoS system is only used every 100 blocks, to provide an extra layer of assurance over PoW, as a checkpointing system.
- Participants in the PoS process send their Ether into a "validator pool".
- Every 100 blocks validators put their stake behind a checkpoint block, whilst also referencing a previous checkpoint block. If two-thirds of the funds in the validator pool support a proposal, the block is considered "justified".
- Once a block is justified, it can be used as a reference for future votes. Once two-thirds of the stake use a justified block as a reference, this justified block is considered finalised and this finality takes precedence over PoW.
- Validators votes are only valid 12 confirmations after the last checkpoint block.
- If the two thirds threshold is not met, the chain continues to progress based entirely on PoW.
- If stakers do any of the following banned behaviors, in return for a small 4% fee, a third party can submit a proof of this, such that the cheater loses their entire stake/deposit (slashing):

    1. Votes for multiple conflicting blocks at the same height.
    2. Votes for multiple conflicting blocks at different heights, but using conflicting reference blocks, unless the new reference block has more height.

The Ethereum reward structure will be adjusted, such that PoS validators also receive a share of the rewards, in addition to the PoW miners. As far as we can tell, the details of this new allocation have not been decided yet.

Analysis

The latest iteration of Casper is a significant improvement from earlier versions, in our view, primarily because of lower levels of complexity and greater reliance on PoW mining.

In theory, there are only three problems with the new proposal:

1. Over one third of the stakers refusing to participate – in which case we are just back to a PoW based system
2. Stakers changing their mind after finality such that more than two thirds supports an alternative chain – the long range attack problem
3. Stakers reaching two-thirds majority support for a lower PoW chain than the current leading PoW chain, a new way of causing a re-organisation. We view this as the most significant downside of this proposal.

Core to the assumption behind this system is that its PoW which drives the chain forwards and that the PoS system only comes into play, once the PoW miners have decided on a chain, PoS votes are not even valid before 12 miner confirmations. Indeed, if the two thirds majority cannot be achieved then the chain continues on a PoW basis.

Therefore, we conclude, that the core characteristic of this latest Casper proposal is that the **PoW happens first**, and only after this does PoS potentially provide an extra assurance against a chain re-organisation, orchestrated deliberately by a hostile PoW miners. PoW therefore still provides computational convergence, with the PoS mechanism defending against the threat of a human/politically instigated miner re-organisation. Therefore although PoS provides this safety, as point three above indicates, it also provides extra risk, therefore its not clear if there is a net benefit.

Weakness	Summary
Nothing at Stake	Validators can vote on multiple chains, but not at the same height. This is designed to allow validators to change their mind, but only for “legitimate” reasons. For the hybrid version of the model, the convergence issue may be solved by relying on PoW mining.
Block production	PoW miners produce blocks and therefore there is no issue related to selecting the block producer.
Long range attack	Once two-thirds of the stake in the validator pool has used a block as a reference for voting, nodes finalize the block and there cannot be a re-organisation. The long range attack problem remains for periods in which nodes are switched off.
Stake grinding	PoW miners produce blocks and therefore there is no stake grinding issue.

Other potential unresolved issues

In the event of a contentious hardfork and chainsplit, if the new chain alters the format of the validator checkpoint votes, two-thirds of the validators could conduct destructive re-organisations on the original chain, while avoiding punishment (slashing) due to the new voting format. Validators could therefore destroy the original chain, while still moving forward on a new chain of their choice. The system could therefore be less resilient to being shut down.

Exclusive BitMEX Research Interview with Vitalik Buterin on the latest Casper proposal

Question 1 – Even though the PoS system may provide more assurance than before, prior to the 34% voting threshold being reached, re-organisation risk may be higher, since a re-organisation can occur in *more ways,* both via PoS and via PoW. Are you concerned about the negatives of this?

> "I would say no. There are plenty of reasons to believe that it should not negatively impact stability. The pre-finalization chain scoring rule is "highest finalized epoch + total difficulty * epsilon". There is a paper here that points out that any "monotonic" chain scoring rule is a Nash equilibrium; our scoring rule is clearly monotonic so it's a Nash equilibrium. Both miners and validators use the chain scoring rule, so miners and validators would both naturally help the chain grow, not try to revert it. Casper FFG was deliberately designed in this way, to "play nice" both with "chain-based" intuitions of consensus as well as BFT-theoretic concepts of finality.
>
> The only way in which "re-organisation risk may be higher" is either:
>
> - If validators are more likely than miners to be majority-dishonest
> - If the Casper-specific code has bugs
>
> We accept that if either of these are true then Casper FFG can add risks."

Question 2 – How do you expect users and exchanges to behave? Should exchanges modify their behavior before crediting deposits, for example 2 confirmations plus 34% of validator votes?

> "If I ran an exchange I would do something like "wait 12 confirmations for deposits up to $10k, and finality for anything higher"

Question 3 – Will there be an overall confirmation score metric, combining both the impact of PoW and PoS, which exchanges can use?

> "I suppose it's possible to create one. Here are a few distinct stages of confirmation that I can think of:
>
> - A transaction has been included into a block, which is the head
>   - Which is the Nth ancestor of the head
>   - Which is an ancestor of a checkpoint C which is an ancestor of the head. Validators have started voting on C.
> - Validators have justified C.
> - A child of C, C', exists, and validators have started voting on C' to finalize C
> - The child of C' has >1/3 votes. At this point, at least one validator needs to actually be slashed for the transaction to be revertedC is finalized."

Conclusion

This latest PoS proposal is the best proposal so far, in our view. We think it may be adopted by Ethereum and it could make a net positive contribution to the security of the system. However, the system remains reliant on PoW mining, at least at the interim stage. PoW is relied on to resolve any Byzantine faults first, before the PoS process occurs. Therefore the system relies on PoW for both block production and for the crucial property of ensuring the system converges on one chain. Although PoS mining may mitigate some risks (hostile PoW miners), it is unclear if it makes a net contribution to convergence or security. Critics of PoS could therefore argue that any rewards redistributed from PoW miners to stakers unnecessarily dilutes system convergence and security.

Although we think the current proposal could work, the nothing at stake problem could still be a significant challenge. The jury is still out on whether this new mechanism solves this problem. Therefore despite the plan to use this proposal as a stepping stone, as part of a gradual shift towards a full PoS system, this could be more difficult to achieve than some in the Ethereum community think.

Disclaimer

Transacting on BitMEX is not offered or available to any resident of (I) the United States of America, (ii) Cuba, Crimea and Sevastopol, Iran, Syria, North Korea, Sudan, or any other sanctioned jurisdiction, or (iii) any jurisdiction where the services offered by BitMEX are restricted.

This material should not be the basis for making investment decisions, nor be construed as a recommendation to engage in investment transactions and is not related to the provision of advisory services regarding investment, tax, legal, financial, accounting, consulting or any other related services, nor is a recommendation being provided to buy, sell or purchase any good or product.

Any views expressed are the personal views of the authors of the report. BitMEX (or any affiliated entity) has not been involved in producing this report and the views contained in this report may differ from the views or opinions of BitMEX.

The information and data herein have been obtained from sources we believe to be reliable. Such information has not been verified and we make no representation or warranty as to its accuracy, completeness or correctness. Any opinions or estimates herein reflect the judgment of the authors of the report at the date of this communication and are subject to change at any time without notice. BitMEX will not be liable whatsoever for any direct or consequential loss arising from the use of this publication/communication or its contents.

BitMEX