LegalBlock Selected Responses to CFTC Request for Input 6351-01-P

*| * Presented by LegalBlock Ethereum Working Group * |*

On December 11th, 2018, the CFTC submitted a public "Request for Input" ("RFI"), which among other things, sought public comments and feedback concerning the development and usage of the Ethereum blockchain network ("Ethereum").

Seeking to provide additional clarity, these selected responses reflect topics of special interest to LegalBlock ("LB"), an online community collaborating to build actionable wisdom over blockchain-based applications and related policy decisions.

This submission is the result of discussions between members of the LegalBlock Ethereum Working Group ("LB-ETH") and the Ethereum development community, primarily aimed at adding clarity to perceived challenges over the network's popular use as a permissionless crowdfunding tool (as seen through "Initial Coin Offerings," or "ICOs") and its ability to scale to secure meaningful economic activity beyond mere digital token creation and transfers.

In substantial respects, LB-ETH wishes to incorporate the responses of EthHub to the RFI, which are comprehensive to the CFTC's request and can be found at the following link: https://docs.ethhub.io/other/ethhub-cftc-response/

For the convenience of CFTC staff, LB has excerpted topics and questions from the RFI (in bolded text) as they relate to LB-ETH input:

LegalBlock

Purpose and Functionality

3. How is the developer community currently utilizing the Ethereum Network? More specifically, what are prominent use cases or examples that demonstrate the functionalities and capabilities of the Ethereum Network?

More recently, the Ethereum developer community has recognized the growing usage of the network as a system for organizing “Open Finance” applications, which is a notable shift from prior themes touted for Ethereum enabling “unstoppable applications” and its projected goal of becoming the “world computer.”

Prominently and highlighting this “Open Finance” theme, the Ethereum development community has used the network to secure significant value related to (i) digital token fundraisers for kickstarting open-source software projects that, by virtue of not developing proprietary IP, might not otherwise be able to attract funding, (ii) managing bounty smart contracts that incentivize and coordinate open-source development through platforms such as Gitcoin and Bounties Network, (iii) exchanging tokenized value peer-to-peer through “decentralized exchanges” such as Kyber and Uniswap, and (iv) taking out loans against ETH as collateral, denominated in USD-pegged stablecoin “DAI” by leveraging the MakerDAO application and smart contracts.

Further, Ethereum projects aimed at governance applications, such as “decentralized organizations” provided by Aragon.org, are allowing developers to manage pooled resources and make group decisions with transparency ‘on chain’ as well as leverage smart contracts to secure obligations more trustlessly between largely ‘remote-first’ development teams. To hopefully help illustrate such governance applications, LB-ETH (a Working Group formed under LegalBlock) has deployed an instance of an Aragon decentralized organization on the *Rinkeby* testnet, located at the following url address: https://rinkeby.aragon.org/#/legalblocketh.aragonid.eth/, to among other things, help signal LB-ETH consensus on topics related to Ethereum and to trial internal reputation rewards and other collaborative incentives ‘on chain.’

5. What data sources, analyses, calculations, variables, or other factors could be used to determine Ether's market size, liquidity, trade volume, types of traders, ownership concentration, and/or principal ways in which the Ethereum Network is currently being used by market participants?

There is more to the size of digital asset markets than mere market capitalization. Better methods for judging their size may be found in studying underlying code development activity and certain market data sets more attuned to digital asset market development. These data sets can be extracted from various blockchain explorer or aggregator services like Blue Swan Grading, Coinmarketcap, Nomics, etc. The various open source project repositories supporting digital asset markets and free APIs from aggregator services can also be used to procure this kind of market 'snapshot.'

For liquidity analysis, normalised trading volume from exchanges, number of trades for volume per exchange, number of transactions on hot wallets, number of Ether locked into DAOs for stablecoins, derivatives, as well as 'cold storage' wallets of exchanges can be useful metrics.

To better determine market size beyond market capitalisation, the number of forks of code repositories (and their related market caps), number of active developers and commits being pushed to repositories, net lines of code, geographical volume of related digital asset trading (by tracking the transaction volume of exchanges, albeit less reliable due to suspected wash trading), can be helpful.

To gauge ownership concentration over digital assets markets, researchers can analyse wealth distributions by viewing the top 'wallets' by digital asset holdings via online sources such as Etherscan or Coin Dance, obtain KYC/AML data from exchanges, as well as track geographical volumes on dedicated timeframes from sources such as Blue Swan Grading.

Technology

8. Does the Ethereum Network face scalability challenges? If so, please describe such challenges and any potential solutions. What analyses or data sources could be used to assess concerns regarding the scalability of the underlying Ethereum Network, and in particular, concerns about the network's ability to support the growth and adoption of additional smart contracts?

Ethereum faces what many have called a 'scalability trilemma' for transaction processing in computer science, wherein blockchain systems are largely forced to choose two out of the following three options: (i) Decentralization (number of block producers), (ii) Scalability (transaction throughput), and (iii) Security (cost to attack network).

In other words, a highly decentralized and secure blockchain system will almost necessarily have to make sacrifices on its transaction processing capabilities - - for example, because of hard-coded limits on computation per block, Ethereum currently supports about 15 transactions per second versus the thousands processed by a provider like Visa.

In the case of Bitcoin, where the cryptocurrency might store large values, or Ethereum, where the network is valued for trustless computation, the permissionless Proof of Work consensus algorithm is chosen for its relatively high offerings in terms of network security and decentralization (proven difficulty to manipulate); however, systems that target different blockchain use cases, such as Steem for micropayments, might use Delegated Proof of Stake ("DPoS") as their consensus algorithm to offer more competitive processing speeds, wherein a limited delegation might mine blocks and have less overhead in reaching consensus, offering higher transaction throughput but losing aspects of decentralization and security in the same process.

For Ethereum, these scalability challenges are mostly present for transactions 'on chain' at the fundamental, protocol layer ('Layer 1') - - more immediately, 'off chain' or 'Layer 2' solutions have been presented for scaling Ethereum's processing capabilities by recording certain sets of transactions off of the blockchain. In a recent post (August 2018), Vitalik Buterin, Ethereum inventor and co-founder, argues:

"(A)s blockchains become more and more mature, layer 1 will necessarily stabilize, and layer 2 will take on more and more of the burden of ongoing innovation and change..."

While this brief overview will focus on these Layer 2 solutions as more imminent developments, there is active research on Layer 1 scaling solutions. "Sharding" is one such solution gaining significant traction within the Ethereum research community. Whereas the current network design requires all nodes to process every transaction (presenting bottlenecks for sake of decentralization), "sharding" proposes secure methods for separating various consensus functions to different participants, drastically reducing system load on Ethereum nodes and possibly enabling thousands of transactions per second.

Among emerging Layer 2 solutions for scaling the Ethereum Network, "Plasma" has attracted significant attention and collaboration from the research community (while this response focuses on Plasma, there are a number of promising Layer 2 scaling solutions, such as "State Channels" and "Truebit" that are seeing implementations and developer enthusiasm).

As described above, recording transactions on the root chain (in this case Ethereum), is relatively slow and costly, as each transaction, from micro to massive, is treated with the same level of security. Before Layer 1 scaling solutions such as ‘sharding’ are implemented, many see establishing relatively secure, parallel transaction methods ‘off chain’ as a promising path to more immediate adoption of blockchain applications.

On August 11, 2017, Vitalik Buterin and Joseph Poon published a whitepaper describing *Plasma*, a proposed framework for “incentivized and enforced execution of smart contracts” that is scalable to potentially millions of transactions per second (tps) and which would service a significant amount of decentralized financial applications.

At its core, Plasma seeks to resolve the tensions of the so-called scalability trilemma and deliver scalable blockchain applications without sacrificing security through a combination of smart contracts and cryptographic verification.

As a proposed ‘Layer 2’ scaling solution (i.e., not an upgrade to base blockchain layer), it targets applications where it is not necessary or even desirable to record every transaction to a blockchain, such as daily coffee purchases.

The general intuition within the Ethereum research community is that these more mundane blockchain transactions might be better recorded to application-specific ‘side chains,’ where assets are locked up on the main chain, replicated on a parallel blockchain, and then transacted under a more efficient consensus mechanism, such as DPoS; fallback to the root chain with less efficient consensus rules (but greater security) is then only utilized for high-value transactions or to establish finality after certain checkpoints.

The goal of the Plasma framework is to reap such massive scalability benefits of side chains while optimizing *fall back* safety to assets locked in the root chain in the event that side chain consensus fails or is overpowered (risking theft of user funds).

Employing an interactive exit mechanism to detect malicious behavior, 'plasma chains' are not quite side chains, which lose their state in the event of failure. When a plasma chain breaks, state is exited but remains intact; as a basic guarantee, the root chain utilizes mathematically verifiable methods to handle disputes and reward the correct party with their funds. Strong security guarantees are therefore a crucial and distinguishing feature of Plasma designs: digital assets that cannot be double-spent, withheld, and are always redeemable on the root chain.

In this way, solutions developed under the Plasma framework seek to offer blockchain security and finality for more mainstream use cases. For example, a cross-border payment networks or gaming platform might issue digital assets as *ERC721s* (non-fungible tokens) on the Ethereum root chain to take advantage of its network security but then get much greater day-to-day transaction throughput for these assets on a plasma 'child chain' running under, e.g., "Proof of Stake" or "Proof of Authority" consensus.

Elaborating on this framework for building more scalable blockchain applications, Ethereum researchers have already branched Plasma into a variety of specifications to serve different applications and project needs. As discussed, the essence of the framework is that a *Plasma Chain must be as secure as the root chain.* Beyond that, Plasma designs typically involve *exits* (user submits transaction history proving ownership of assets with collateral) and subsequent *challenge periods,* where others, incentivized to claim the exit collateral as a bounty, can challenge such exits by submitting a contrary proof.

On January 3, 2018, Vitalik Buterin, Joseph Poon and David Knott released specifications for a "minimal viable plasma implementation" (commonly, "Plasma MVP"). Essentially, the Plasma MVP specification aims to provide plasma's basic security properties "in a very simplified way," such as to enable scalable payments, "though it leans heavily on users being willing to immediately exit as soon as they detect any kind of malfeasance." While the Plasma MVP is designed for token transfers, it can be adapted for ERC721s and general state transitions and potentially scale to more than 1,000 tps (Ethereum researcher, Karl Floersch has noted that "later versions" of MVP design may scale to "millions" of tps).

On March 9, 2018, Vitalik Buterin introduced "Plasma Cash" at the Ethereum Community Conference, a specification which aims to increases Plasma security and usability with unique identifiers for funds deposited on a plasma chain, essentially turning each deposit into a non fungible "coin" with an independent serial number and transaction history. In other words, each deposit is treated like an indivisible bill, much like the familiar denominations of $10, $20, etc., for physical cash, and users only store data about coins they own.

Among other benefits, this construction allows for simpler fund withdrawals with “much less per-user data checking,” i.e., more compact proofs on a coin’s history by only requiring users to validate for the ones they own and are actively watching (and not entire chain of transactions for all, *contra* Plasma MVP). On March 14, 2018, Ethereum researcher Karl Floersch released a full specification for a Plasma Cash chain. Plasma researcher, Georgios Konstantopoulos, has also released a comprehensive document covering Plasma Cash topics and research on initial implementations (including gaming use-cases via LOOM Network development).

In June 2018, “More Viable Plasma” (commonly, “MoreVP”) was introduced on ethresear.ch by Kelvin Fichter and Ben Jones, a design which, among other things, seeks to make security and UX improvements to the MVP design by removing confirmation signatures and making withdrawals cheaper.

On November 1, 2018, Quantstamp, a blockchain security company, announced that it had completed a security audit of a Plasma MVP implementation designed by blockchain project, OmiseGO; a similar audit of a MoreVP implementation is also planned as of this response.

Current Plasma designs are not without drawbacks: for example, (i) in the worst case, Plasma MVP requires every plasma chain user to exit within a short period of time (limiting throughput, as number of UTXOs that can be safely withdrawn is also the number of UTXOs that can be safely supported on a Plasma chain); further, (ii) when users withdraw funds from a plasma chain, they’re required to wait for a period of time before those funds become available on the Ethereum root chain.

A number of proposals, however, are designed for these problems, and include protocols for “mass exits” that allow thousands of UTXOs to be exited at the same time, and “fast withdrawals” as a way for users to “sell” their withdrawals in order to avoid waiting.

EVM support on Plasma is an active area of research and discussion. Right now, it is difficult to create a Plasma chain that can run more general smart contracts like Ethereum for the following reasons (as outlined in an August 2018 post by Plasma researcher Kelvin Fichter):

(1) It’s not always clear who gets to move a contract from the Plasma chain to the root chain;
(2) If anyone can modify the state of the contract, then anyone can block an exit; and
(3) Validating EVM state changes inside the EVM is hard.

However, a potential (*but still highly preliminary*) solution known as “Plasma VM” involves breaking smart contracts down to a level where these issues might not matter as much.

Specifically, Plasma VM proposes reframing authority to move a contract from the Plasma chain to the root chain using “mini smart contracts”; in other words, instead of worrying about “who gets to move a contract from the Plasma chain to the root chain,” Plasma VM stipulates that it might be more clear who’s responsible for moving stuff to the root chain if, instead of moving the entire contract, everyone were moving their own “mini smart contracts” that can only be modified by their respective owners.

A related effort aimed at Plasma EVM-support, introduced as “Plasma Leap” by Johann Barbie on ethresear.ch in September 2018 and in active development by LeapDAO, seeks to resolve these issues for “Dapps to leap onto Plasma” by breaking smart contracts into smaller programs called spending conditions and into single state objects with clearly defined owners.

There are over 100 Plasma-related topics being actively discussed on ethresear.ch, and LearnPlasma.org provides comprehensive summaries of current research and implementations. Further, and for more practical reference, Plasma.Group provides provides resources to launch test plasma chains and transactions.

Governance

13. How is the governance of the Ethereum Network similar to and different from the governance of the Bitcoin network?

Governance in a blockchain context refers specifically to the process of maintaining and updating the software. This includes the protocol modification, deliberation and ultimate decision-making process. Governance processes here not only depend on the architecture of the network (i.e., public/private, permissioned/permissionless), but also on the general purpose of the project and the underlying ideology that motivates network participants.

Generally speaking, governance in permissionless networks is more complex than in permissioned ones. As a technical matter, closed networks are easier to coordinate, and politically, they have the benefit of fulfilling a highly specific purpose. On the other hand, public blockchain networks serve a more general purpose, allowing ideology to more readily evolve (as well as produce friction).

Blockchain governance processes can include ‘on chain’ and ‘off chain’ procedures. In ‘on chain’ governance, decisions are taken on the blockchain and token holders voters decide to endorse an update (which is automatically executed). ‘Off chain’ governance processes take place in the “real world” and can include more stakeholders from the broader network community. (Overall, miners decide whether to install a new software and thus exert agency.)

Both Bitcoin and Ethereum are open, permissionless networks and rely on ‘off chain’ governance processes known as Improvement Proposals; “Bitcoin Improvement Proposals” (BIP), in the case of Bitcoin, and “Ethereum Improvement Proposals” (EIP) in the case of Ethereum.

Although anyone can suggest changes changes to the protocol through BIPs or EIPs, only core developers exercise voice in proposing changes. These changes will then ultimately be implemented by the blockchain miners. This governance process is in many aspects informal with no clear, transparent rules in place regarding core developer appointment or removal. (This informality may rise to complications in the future, but efforts to formalize coordination and improve transparency are underway, as seen in community governance efforts such as the “Fellowship of Ethereum Magicians.”)

Ethereum governance largely resembles the Bitcoin network, wherein the general public can access and track core development upgrades via the EIP Github repositories. Ethereum is currently experimenting with a range of methods to incorporate more stakeholders in its core development in a bid to openness to match its relative liberal approach to upgrades, as seen in the variety of working groups collaborating on the ‘Constantinople Hard Fork Upgrade’ and Ethereum 2.0 research and design decisions.

Generally, governance towards the acceptance of proposals has been through “rough consensus” as gauged through coinvotes, as well as online surveys and polls between popular social media platforms such as Twitter and Reddit. Such amendments enter into force when they are adopted by miners with more than 50% of the mining power.

Ultimately, and similar to Bitcoin, users running nodes and supporting these networks can choose which code to run and updates to follow, serving as a check on core development that is not responsive to stakeholders. Therefore, as these networks grow and upgrades impact more economic activity, finding a healthy balance between introspective and considered development versus the desires of stakeholders will likely remain an ongoing project for Bitcoin and Ethereum governance and a source of competition (e.g., in Bitcoin development, there has been more observed conservatism among stakeholders over upgrades to base blockchain layer, due to network’s primary use as money, whereas Ethereum, which aims to serve a wider range of applications, has observed liberalism and upgrade forks are less contentious).

14. In light of Ether's origins as an outgrowth from the Ethereum Classic blockchain, are there potential issues that could make Ether's underlying blockchain vulnerable to future hard forks or splintering?

As stated by EthHub in their CFTC response, *"In July of 2016, following the exploit of a vulnerability found in a smart contract called 'The DAO' which resulted in substantial loss of Ether to a bad actor, the Ethereum blockchain was hard forked again to include an 'irregular state transition' that recovered the stolen Ether and returned it to the initial owners."* In this case, most of the stakeholders in the Ethereum Network decided to support the blockchain hard fork, and consequently, they upgraded their nodes.

Therefore, stating that "Ether origin is an outgrowth from the Ethereum Classic blockchain" might be a misrepresentation of the technical reality.

Now, with regard to Ethereum being vulnerable 'future hard forks,' the answer is simply "yes." Like other communities, the Ethereum blockchain can be vulnerable to future splintering in the event that upgrade decisions are controversial and/or not supported by the community and the nodes they run. However, the belief that blockchain forks are necessarily a bad thing is a misconception. The protocol is designed to allow hard forks.

Cyber Security and Custody

23. Are there security issues peculiar to the Ethereum Network or Ethereum- supported smart contracts that need to be addressed?

As noted by cryptocurrency and smart contract researcher Philip Daian in a recent presentation titled, "Smart Contract Security - Incentives Beyond The 🚀 or: How I Learned to Start Analyzing and Stop Building Inscrutability," Ethereum-supported smart contracts need more robust audits and 'holistic' security approaches beyond the domain of "launch once and walk away" ICO token smart contracts, as there are observed issues associated with more complex (yet arguably still *basic*) voting and exchange smart contracts, such as the vulnerability to miner frontrunning seen with decentralized exchanges (miner may profit by executing cancelled DEX orders with themselves as counterparty). However, there are promising trends towards more complete smart contract security in Ethereum, as seen in the growing (i) range of static analysis tools, (ii) open-source community of engaged developers, and (iii) sets of formal methods, tools, practices and standards.

******

LegalBlock

Thank you for providing this opportunity to share comments and for your attention to this matter of great significance to financial services and regulated markets around the world.

We look forward to engaging in future discussions related to Ethereum.

Sincerely,

Ross Campbell

Co-Chair | **LegalBlock Ethereum Working Group**
DAO | https://rinkeby.aragon.org/#/legalblocketh.aragonid.eth/

Member | **LegalBlock**
WEB | https://legalblock.co/