



Quarterly Insights

PREPARED BY _____

Gabriella Sofia, Umberto Natale
Luis Nunez, Felix Lutsch, Xavier Meegan
Thalita Franklin, Michael Moser

Q2
22

Table of Contents

01

Ethereum -
Merge update

09

Vertical Scaling is
Coming to the Cosmos

14

Solana MEV -
A first outlook

18

Network and Validator
Supply Dynamics

20

Terra (Incognita):
Retail & Institutions
on Terra pre- and
post-collapse

26

The State of
Avalanche Subnets
in Q2 2022

Ethereum - Merge update

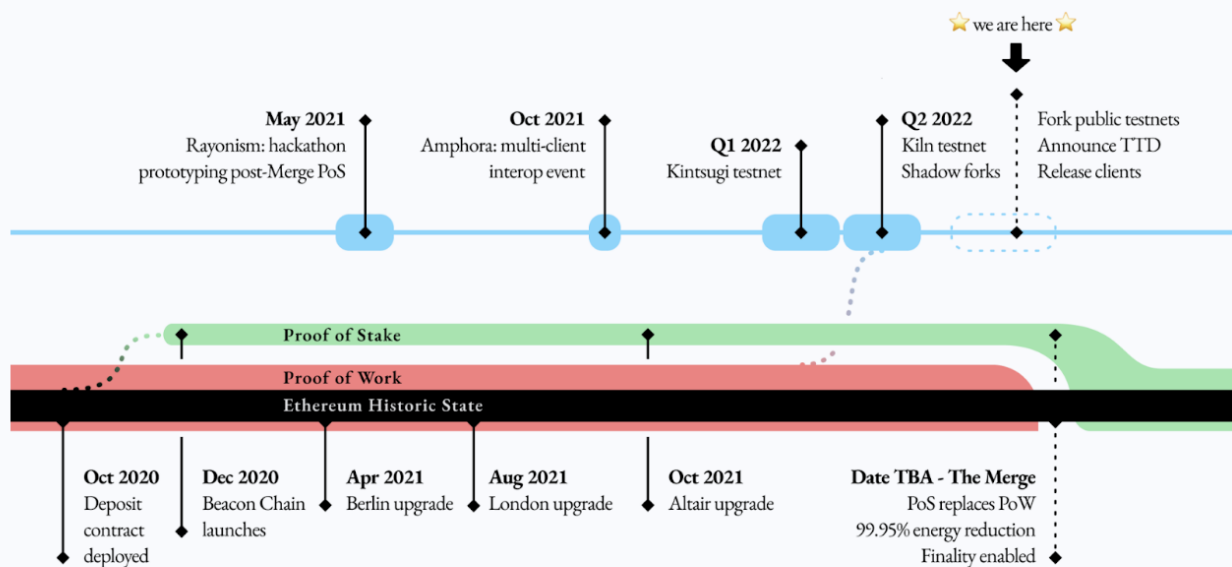
Very soon, the current Ethereum mainnet chain will Merge with the Beacon Chain Proof-of-Stake system, marking the end of the Proof-of-Work era for Ethereum and setting the stage for future scaling upgrades.

To be somewhat sceptical on the final timing of the Merge would be understandable. Ethereum has undergone several network upgrades in the years since genesis, tweaking the parameters of what is known as the Difficulty Bomb (a code adjustment that makes it much harder for miners to verify transactions and earn a reward under PoW) on more than a few occasions: **Byzantium, Constantinople, Muir Glacier, London** and **Arrow Glacier** are the name of those upgrades, with the latest being **Gray Glacier** just this past June, the third hard fork containing exclusively a delay of the bomb that threatened to cause network issues by the end of the month. At this point, it is worth revisiting some of the concepts and facts behind this monumental task and the reasons why we believe we are very close to seeing it materialized.

The Path to the Merge

Offchain Onchain

May 4 2022 - @trent_vanepps



Why TTD

A network upgrade is a change to the underlying Ethereum protocol, creating new rules to improve the system. The decentralized nature of blockchain systems makes a network upgrade more difficult. Network upgrades in a blockchain require cooperation with the community, as well as with the developers of the various clients in order for the transition to go smoothly.

Unlike any other protocol update, the Merge does not get activated by a specific block number, but rather a mechanism known as Difficulty Bomb, and it comes into effect after reaching an accumulated difficulty (Total Terminal Difficulty = TTD), which can be influenced by the behaviour of the bomb and miner activity. At the moment of the transition, validators on the consensus chain are watching the PoW chain for this terminal condition to be satisfied, so they can begin in the next block to reference this terminal PoW block as the parent of its execution layer.

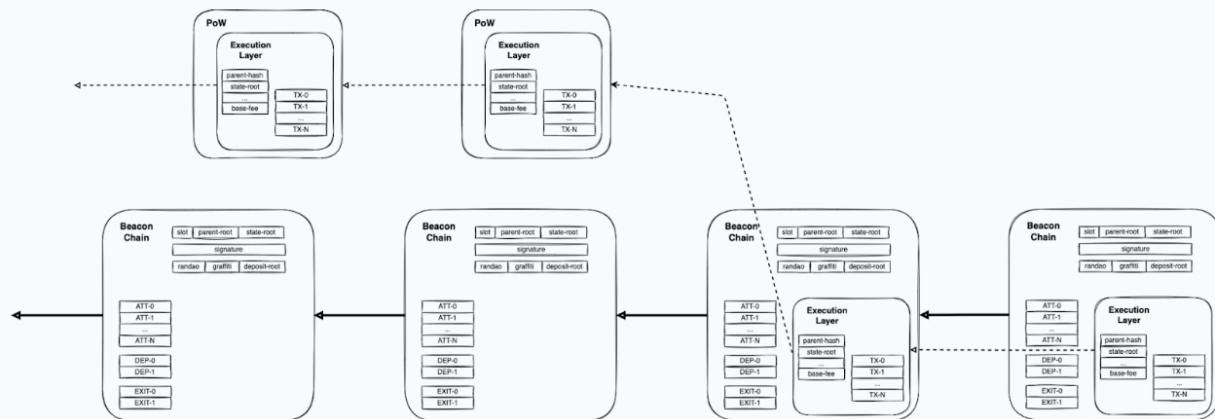
The reason to move forward with a TTD and not block height comes from a few considerations: first, because as validators we want to take over from a specific state (the canonical chain) and not any other state that could be easily manipulated. For example, an attacker could (potentially) use a minority of hash power to build a malicious chain fork that would satisfy the block height requirement. In that sense, a more secure point of activation is needed to ensure that that singular event is correct. Also, the use of a difficulty bomb is meant to quickly and effectively dissuade miners to continue on the PoW chain by making it economically inviable.

How it happens

The existing Ethereum mainnet is the blockchain we're all used to: it contains every transaction, smart contract, and balance since it began to operate in July 2015. In a parallel manner, the Beacon Chain has been working since December 2020, not to process transactions, but instead to reach consensus on its own state by agreeing on active validators and their account balances. This PoS network is not a thing of the future, it has been operating with 100% reliability for the past year and a half, through the work of staking operators such as **Chorus One**.

This “liveliness” of the chain is the main reason why it has been so difficult to reach the final specification for the transition. At the Merge, we will see the last block under the PoW mechanism followed immediately by the first PoS block, in a continuous way and with no noticeable downtime to an observer. The execution layer continues to back reference itself, now surrounded by the validation properties of the consensus layer, and without any change to the application layer.

Post-merge, Beacon blocks contain the transactional payload that is inside current PoW blocks. This is called the ExecutionPayload in the specs. The following diagram illustrates what it would look like:



Once again, when the last PoW block is produced, subsequent Beacon Chain blocks include the transactional data. Roughly, the validation process would look something like this: first, a validator is elected to propose a block. This validator asks its execution layer, via the Engine API, to send him an ExecutionPayload. The EL returns the payload which contains the most profitable set of valid transactions to the consensus layer. The CL proposes a block which includes this payload and propagates it on the Beacon Chain network. Finally, other validators attest to the block and, if valid, propagate it on the Beacon Chain network.

Shadow forks

There are several testing efforts going on currently to get us to the point of delivering a Merge on mainnet. A new innovative testing method has been the inclusion of “shadow forks” to the main network. Essentially, this represents a new devnet where a small number of nodes are configured to fork off from an Ethereum network at a certain point and Merge with a Beacon Chain that is launched only for testing purposes. By modifying these nodes, which are set to run through the Merge at an earlier point than the entire network, it allows developers to test how the upgrade would have happened in similar conditions to the shadow-forked network, but without the vast majority of the nodes being aware this has happened and with a similar (simulated) throughput that can contain any transaction seen on the canonical chain. These are the harshest conditions in which we can test without creating a disruption.

As of the writing of this report, Ethereum developers have gone through 8 shadow forks of the main network, with a 9th mainnet shadow fork to test the Merge **coming up**.

Testnet merges

The second testing method, and arguably the best indicator of progress, is the merging of the testnets. These are networks used by protocol developers or smart contract developers to test their software in a production-like environment before deployment to mainnet (similar to production versus staging servers).

So far, we have successfully completed the Merge to the **Ropsten** testnet on June 8th. With a 99.2% participation rate, healthy proposals and sync committee performance right before the activation, we saw a drop of just

~13% participation, which turned out to be a configuration issue on the Nimbus client and not a merge-related problem. After the fixes were in place, participation and proposal rates jumped back towards ~99%. Most recently, we completed the Merge on **Sepolia** (where most of its PoW validators are public validators, unlike other PoA testnets) which hit TTD on July 7th. This chain is currently being monitored after roughly 25-30% of the validator set experienced issues with configuration, related to the **Lighthouse** client. This is, however, not expected to affect the timing of the Merge.

Ropsten will not be maintained long term, and should be **deprecated** by EOY. On the other hand, Sepolia will be maintained by client developers with a closed validator set. Because it is fairly new, its state and history are both quite small which means the network is quick to sync to and that running a node on it requires less storage. This is useful in contrast to a robust testnet like Goerli.

After years of work to bring proof-of-stake to Ethereum, this final stage of testing gives us more confidence in estimating the timing of the Merge.

Estimating the Merge

No specific date has been set for the Merge. Once the client teams are confident with the state of testing, this will be announced alongside the releases to the client upgrades. "According to latest estimates, the Goerli/Prater Merge would be hitting around August 11th," the last step before going into mainnet. Alongside the mentioned delay in the Difficulty Bomb (set to start showing back up by the middle of September),

one could estimate the most optimistic outcome to be an announcement to Merge during the month of September 2022. And most recently, the development teams finally mentioned September 19th as a tentative Merge date, confirming suspicions." However, this is entirely dependent on the outcome from the continuing tests.

There are some things to keep in mind as we approach the final months of PoW Ethereum:

To address the elephant in the room first, the Merge is a change of consensus mechanism, not an expansion of network capacity and will not result in lower gas fees. Smaller block and slot times do increase available blockspace, but not significantly. And available blockspace is only one input which influences fees, the other being demand. Also, though some slight changes exist, transaction speed will mostly remain the same on Layer 1. In fact, blocks will arrive exactly every 12 seconds post-merge, only a small difference from today, when they arrive in a poisson distribution around ~13 seconds.

Another point has to do with withdrawals. It is widely known this upgrade will not implement withdrawing or transferring of staked ETH. This feature will be included in the Shanghai upgrade, the next hard fork that is planned to follow the Merge. However, it is incorrect to say that validators will not receive any liquid ETH rewards until Shanghai, as fee tips/MEV will be credited to a mainnet account controlled by the validator, immediately following the transition. So the rewards will continue, with a predicted increase of about 50% in APR.

Finally, how will the Merge impact ETH supply?

We can break the supply of ETH into two primary forces: issuance and burn. The issuance of ETH is the process of creating ETH that did not previously exist. The burning of ETH is when existing ETH gets destroyed, removing it from circulation. To make a long story short, currently we get an issuance of ~13,000 ETH/day in mining rewards and ~1,600 ETH/day in staking rewards. After the Merge, only the second one of these numbers will remain, dropping total new ETH issuance by about 90%. As a counterpart, at least 1,600 ETH gets burned every day in base fees thanks to EIP-1559 (at an average gas price of at least 16 gwei), which would effectively bring net ETH inflation to zero in this scenario.

Ethereum is the only major protocol building a scalable unified settlement and data availability layer, and although this event feels in some ways like a conclusion, there's still so much to be done to achieve their ambitious roadmap. In a way, it is just the beginning.

Vertical Scaling is Coming to the Cosmos

Cosmos has historically been an ecosystem that has promoted horizontal scalability, as opposed to vertical scalability. The Cosmos ecosystem has been able to scale horizontally more efficiently than any other ecosystem as a result of having the most mature interoperability protocol and software development kit in cryptocurrency, known as the Inter-Blockchain Communication Protocol (IBC) and Cosmos Software Development Kit (Cosmos SDK). The power of having the flexibility and optionality to create your own blockchain in the Cosmos allows the ecosystem to scale 'horizontally'. Now, for the first time in Cosmos history, there are vertical scaling solutions being built in the Cosmos ecosystem that complement existing horizontal scaling solutions that already exist within the ecosystem. In particular, there are four promising vertical scaling solutions being worked on in the Cosmos right now, which includes Cosmos Hub Interchain Security, Dymension, Celestia and Saga. The Cosmos ecosystem is unique in that each vertical scaling solution being worked on intrinsically scales horizontally as well, thanks to the flexibility facilitated by the modularisation of Cosmos. There is now a chance for the Cosmos ecosystem to become the world's most scalable and secure blockchain as the ecosystem leverages both horizontal and vertical scaling.

As of time of writing, there is **an ongoing governance proposal** to bring **Interchain Security** 'to-market' in the Cosmos Hub. **Neutron** and **Quicksilver** will be the first two 'consumer chains' to borrow security from

Cosmos Hub. Neutron is going-live as the first DeFi zone consumer chain on Cosmos Hub and Quicksilver is going-live as the first liquid staking consumer chain zone on Cosmos Hub. Although both Neutron and Quicksilver are targeting different types of assets and audiences, both zones are alike in that they borrow security from the underlying 'provider chain', Cosmos Hub. Interchain security provides teams with a new option to deploy applications directly onto consumer chains that are intrinsically linked to the Cosmos Hub, which yields a variety of benefits such as smoother upgradability and more-developed tooling for their applications versus deploying as a sovereign chain elsewhere. Cosmos Hub will be an attractive option for teams to deploy applications to in the future, given it is the most secure network within the Cosmos. A feature that is little-known about the Cosmos Hub Interchain Security upgrade is that synchronous calls will be possible (i.e. composability). Not to forget, composability was the catalyst for DeFi summer back in 2020. There is nothing stopping Cosmos Hub becoming a centre for financial activity within the Cosmos ecosystem once Interchain Security is enabled.

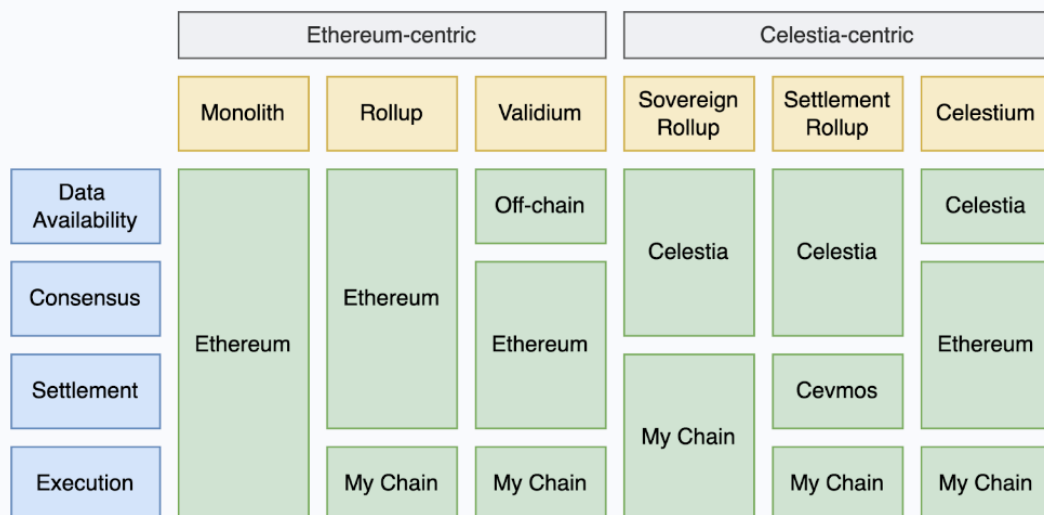
Another vertical scaling solution being worked on in the Cosmos is **Dymension**. Dymension is taking a very similar approach to Ethereum's current vertical scaling roadmap. The main difference that Dymension is taking compared to Ethereum is the level of customisation and flexibility on offer versus what is available in Ethereum. Dymension is working on creating a Rollup Development Kit (RDK). The RDK takes inspiration from the Cosmos SDK and can be tweaked effortlessly by any team, depending on their needs. The best way to think of Dymension is like Ethereum's current settlement and execution layer design (e.g. ORUs borrowing executing tx off-chain and then writing state to the 'settlement' layer), only Dymension

inherits many properties that makes Cosmos networks so dynamic, such as native interoperability, PoS and a developer framework to easily spin-up rollup chains. Considering the amount of zones that have gone-live as an outcome of the Cosmos SDK being an understandable framework to build blockchains with, it could be expected that a similar amount of roll-ups will go-live on Dymension if the network manages to keep the RDK simple and extensible.

Related to Dymension but also with its own unique design that is a vertical scaling solution going live in the Cosmos is **Celestia**. In a nutshell, Celestia is a 'data availability network'. Breaking this down, Celestia validators guarantee that state (data) is available for verifiers to verify themselves that execution has been done properly off-chain in order to mitigate any need for a challenge period on the 'settlement layer'. Celestia network itself does not execute any transactions. It is merely a network that has the latest state of an L2 that can be leveraged by verifiers to determine whether or not data is available (and therefore can reconstruct the previous state to check if execution has been done appropriately in different intermediate states). Essentially, Celestia is blockchain-agnostic and provides consensus over data availability within an execution layer. This is a powerful concept and Celestia's importance could transpire across both Cosmos and Ethereum in the near future.

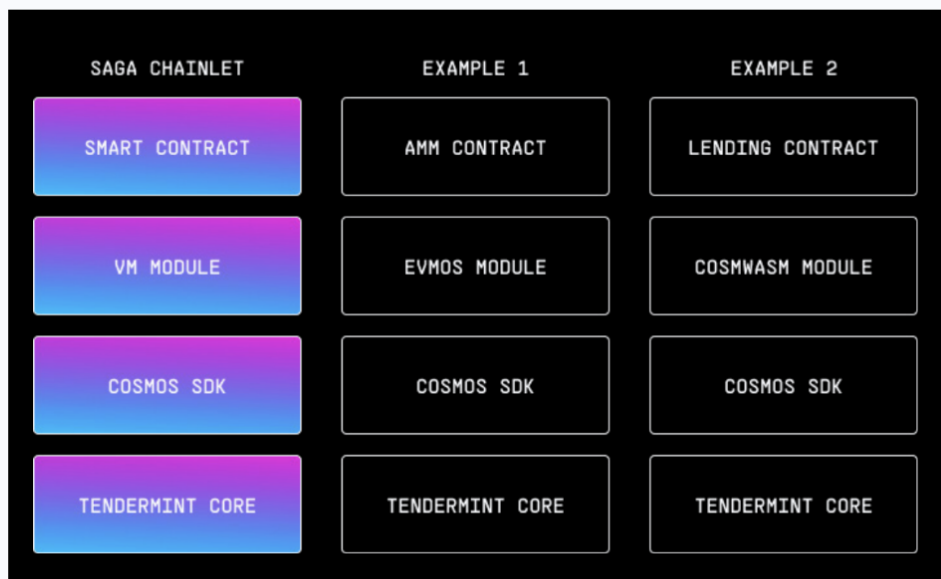
There is a chance that Celestia will be the nail in Ethereum's coffin when it comes to vertical scaling. As it stands, Ethereum-centric rollups (e.g. Arbitrum) uses Ethereum for consensus, settlement and data availability. However, it is becoming less practical for Ethereum-centric rollups to use Ethereum for data availability and settlement because data availability is

becoming increasingly expensive on Ethereum and settlement is unnecessary if an execution layer removes the requirement to only accept bridged assets from another chain (e.g. Ethereum). In the past, execution layers had no choice but to accept the inefficiencies of using Ethereum at different levels of the stack. With the advent of Celestia, execution layers now have the option to use Celestia instead of Ethereum at the consensus and data availability level of the stack. The benefits of using Celestia over Ethereum for data availability and consensus are profound. For example, execution layers that choose to leverage Celestia will have a larger degree of flexibility, customisation and autonomy and reduce ‘rent-seeking’ expenses significantly. It might not be long before execution layers that exist on Ethereum become enticed by the advantageous properties that Celestia offers and choose to migrate to Celestia for consensus and data availability. If this happens, it would significantly reduce the value proposition of Ethereum itself, which is fully-focused on becoming a data availability layer in the future.



Source: Twitter @ptrwttts

Finally, another vertical scaling solution being built in the Cosmos is Saga. Saga is a network that is purpose-built to give each application that launches on its network its own execution environment. This means there could potentially be hundreds / thousands of ‘chainlets’ running on Saga. A core value proposition of Saga is that execution environments are customisable, an application has the flexibility to choose its own execution environment depending on its needs. The power of each individual application having its own execution environment is that resources can be managed in a more efficient way. Not only that but Saga plans to break down applications into little components. If an application becomes too burdensome in one execution environment, it can be broken up and split modularly into different execution environments. This type of sovereign execution environment per application design is unique in the Cosmos and we are excited to support the network in the future.




Source: Saga Litepaper

Solana MEV - A first outlook

The first conceptual exploitation of the Maximum Extractable Value (MEV) appeared 8 years ago in a pmgoohan [reddit post](#). However, the whole problem was framed only in 2019 by [P. Daian et al](#) in the “Flash Boys 2.0” paper, where the word MEV was used for the first time. Now it represents a fundamental concept in cryptoeconomics, highly affecting permissionless blockchains. Indeed, not all MEV can be considered benign since some MEV represents an invisible tax on the user, see e.g. [Solana MEV Outlook](#). In general, MEV can also be an incentive for consensus instability, see e.g. [time bandit attack](#).

However, not all forms of MEV are malignant, and an example of this is arbitrage. It consists in the simultaneous purchase and sale of the same asset in different markets in order to profit from differences in the asset's listed price. From the definition it is clear why arbitrage represents an AMM health medicine. Imagine that some user swaps a huge amount of token A on a specific AMM (huge with respect to the total amount in the pool), and that this transaction creates a \$5,000 arbitrage opportunity. All users that swap tokens in the same pool and same direction will see their output lowered with respect to the actual value. Thus, who exploits this MEV opportunity will also bring the market back to parity with the true price. This will make the AMM more efficient without harming its users in the process.


On Solana, MEV still represents a dark forest since no one has pointed a solid lighter on it. This is because Solana is a much younger blockchain with respect to Ethereum, implying that products like **Flashbots** are still not consolidated. One of the most interesting projects that is moving in this direction is **Jito Labs**, which recently delivered the **first MEV Dashboard for Solana**, representing an explorer aimed at illuminating MEV - see [here](#) for an introduction. However, it is not the only one. For example, at Chorus One, we are pointing lights on some Solana Decentralized Exchanges (DEXs) in order to illuminate the dark forest. MEV is a consequence that will be a crucial factor for the future of PoS networks and we are continually looking for the best way to ride it. You can explore our Solana MEV dashboard [here](#)



By lighting the Dark Forest, we discovered a searcher which uses an address with clear reference to **Jito, jito...3s8**, extracted 118,764 USDC in 40 days, mostly profiting during the LUNA crash time-windows. An example of two profitable transactions [here](#) and [here](#), with a total of 71,977.63 USDC.

Extracted MEV from same pool Orca x Raydium


Compute the extracted MEV from same USDC-connected pools on Orca & Raydium



If we consider what happened after the LUNA crash, by tracking only a few of the possibilities on Solana, we estimate that the total extracted MEV in the month of June was higher than 109,464 USDC. What is interesting is that this number takes into account only 3 DEXs - Orca, Raydium, and Creme Finance - and was extracted during a period of low volume exchanged. By only using the imbalance between the same Orca's and Raydium's pools, searchers were able to extract 43,969 USDC. This shows how dense the forest is in terms of MEV opportunities when we consider only the same pools on different DEXs.

Extracted MEV from same pool Orca x Raydium by accounts

Compute the extracted MEV from same USDC-connected pools on Orca & Raydium



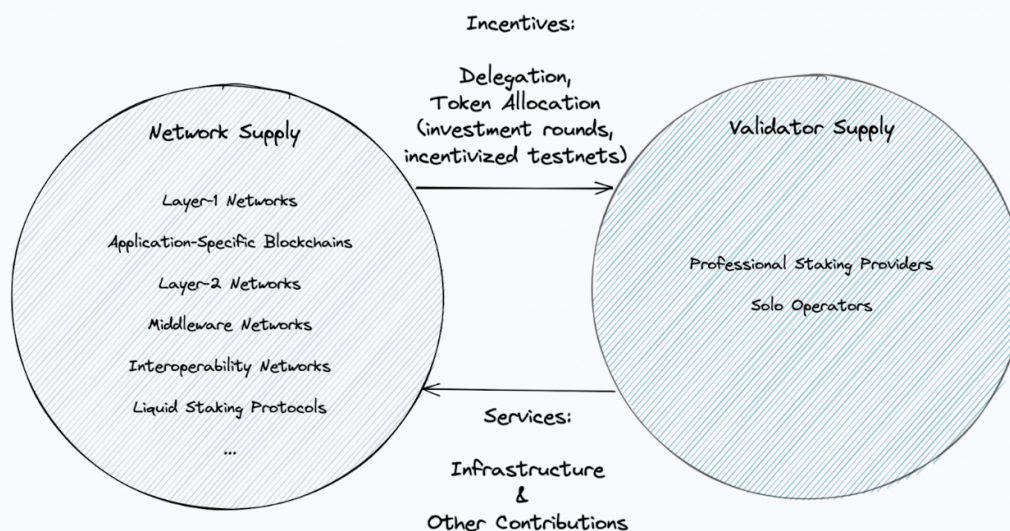
Deep diving, we discovered that only one searcher, [F6j...22q](#), is taking 77% of the total extracted MEV. Here [jito...3s8](#) extracted 13% of it. This result shows how much the extracted MEV can be centralised among few searchers, even on Solana. MEV in general can point toward a centralised system, increasing the challenge for infrastructures like Jito to bring democratisation again on users.

Now at **Chorus One** we are starting to explore MEV, trying to identify all its forms and understand how it works. This is because MEV represents an unavoidable consequence of the protocol engine, but is in our best interest to understand how to bring back power to users, truly democratising it. In addition, we are closely monitoring the upcoming upgrades that are occurring on Solana, such as QUIC, QoS and Fee Markets. All three of these proposed changes will impact MEV on Solana in the future. We recently published an article that explored the three upcoming upgrades planned for Solana and unravelled what they mean for the network, which is available [here](#). We are now actively exploring the source code on Solana to determine the exact way in which MEV will be impacted with these proposed changes.

Network and Validator Supply Dynamics

At **Chorus One**, one of the core parts of our business is to adequately distribute our resources across the networks we support. This choice comes down to due diligence of the underlying protocol, economics, team, vision, as well as maturity and processes surrounding the software that operators need to run. A prime concern for any node operating team is that the incentives provided by the network need to justify the expenditure in terms of CapEx and OpEx involved in operating the network. And beyond that enable upside from participating in the early days, and/or contributing to the protocol via governance participation, tooling, or in other ways.

Over the past 4 years, we have spoken to roughly 200 teams building a network that involves some kind of node operating role and ended up onboarding 32 of them in the process, with roughly a dozen in our pipeline that we're in the process of, or likely going to, onboard in the course of the year. This means we reject about 80% of the networks we look into or receive requests to onboard. This serves to illustrate how even we, among the larger teams in the node operator ecosystem, need to be relatively restrictive in terms of which projects we engage with.



One of the core tools in the toolbox of protocol designers that have to incentivize node operators to join and improve their network is to delegate tokens to those continuously participating and/or adding value to the ecosystem in other ways if their goal is to create a decentralized network.

In [one of our latest articles](#) on the [Chorus One blog](#), we're looking into stake delegation programs comparing the criteria that 7 of the most prominent projects utilizing this strategy use to incentivize validators to join their protocol. For this, we look at the various criteria, breaking them up into 3 buckets:

- Performance & Participation
- Ecosystem Contributions
- Security & Decentralization

Read the full article [here](#) to see how delegation strategies from projects including Solana, Polkadot, and Lido compare.

Terra (Incognita): Retail & Institutions on Terra pre- and post-collapse

In this spot in the last quarterly insights, we highlighted idiosyncratic risk(s) of Anchor and Terra, for validators specifically.

Since then, we have decided against validating any Terra-Chain going forward. For this issue, we'll venture into Terra Incognita by comparing the pre-collapse delegation landscape to the status quo.

We will address two questions:

- Who is still there (and continues to foot the bill)?
- How has the validator set changed, and is this consistent with our initial risk assessment?

In summary, we argue that:

- There is proportionally larger retail participation on the current Terra chain ('Terra 2.0') vs. Terra Classic.

- Next to price impact, the demise of Anchor has significantly altered the delegation structure, materialising an idiosyncratic risk for top validators.
- In practice – 10 of the top 20 validators on the original Terra chain have continued their operations on Terra 2.0. Of these, 6 dropped out of the top 20.
- The absence of Anchor has led to a re-emergence of ‘validator clustering’ by firm strategy (e.g. exchanges; retail focus). This is consistent with other chains.

As a recap, we noted that delegations on the original, pre-collapse, Terra chain were disproportionately top-heavy.

For the original Terra Blockchain: 23 addresses delegating in excess of 1m Luna (roughly \$100m at the time of writing) accounted for 62% of all delegations, and for 78% of delegations to the top 20 validators.

Consider the following graph:




Figure 1. Stake by delegation size brackets, for the top 20 validators. Terra pre-collapse (April 2022).


Typically, validators form clusters on a plot showing the number of delegations and the total stake on the axes respectively. For example, validators run by investment firms and exchanges may control a lot of stake, delegated by a few wallets controlled by the same entity. An inverse effect can be observed for validators with a primary focus on retail, e.g. a strong social media presence, with little institutional sales acumen.

On Terra, these clusters were not particularly pronounced. This was due to the disproportionate role of delegations driven by Anchor, which would strongly influence the composition of the top 20 validators. Anchor (via bLuna / Lido) alone accounted for 27% of all staked Luna; by contrast, the single largest delegator Cosmos Hub (an exchange wallet) accounts for 'only' 6% of all staked ATOM.

Next to a depeg risk, the idiosyncratic risk for validators was thus an unavoidable over-reliance on delegations from Anchor. This was another potential entry point for a vicious cycle, which at a minimum would have had guaranteed impact on validators – a superior solution to Anchor emerges, mercenary capital leaves, a two-pronged value loss hits validators: a loss of delegations, and a decrease in the asset price.

With the absence of Anchor, the delegation landscape on Terra 2.0 is structured differently. The new LUNA token has heavily declined in price since launch (~\$ 2.1 at time of writing vs. a ~\$ 5.15 launch price). An interesting framing is – who is footing the bill? Are institutions still in?


Repeating the exercise above yields:



Total stake by delegation size brackets, for the top 20 validators. Terra post-collapse (July 2022).

The increased stake share of lower stake brackets is an indicator of proportionally higher retail participation on Terra post-collapse vs. pre-collapse.


This, and the impact of Anchor's demise for validators, comes across best by re-calibrating the chart to show the delegation structure by validator. Pre-collapse:



Stake by size brackets, by validator, for the top 20 validators. Terra pre-collapse (April 2022)

The upshot is that most validators in the top 20 depended heavily on very large delegations (yellow > 1m LUNA) and, to a lesser extent, 'large' delegations (pink > 100k LUNA), i.e. tickets outside the retail range.

The contrast with the status quo is striking:



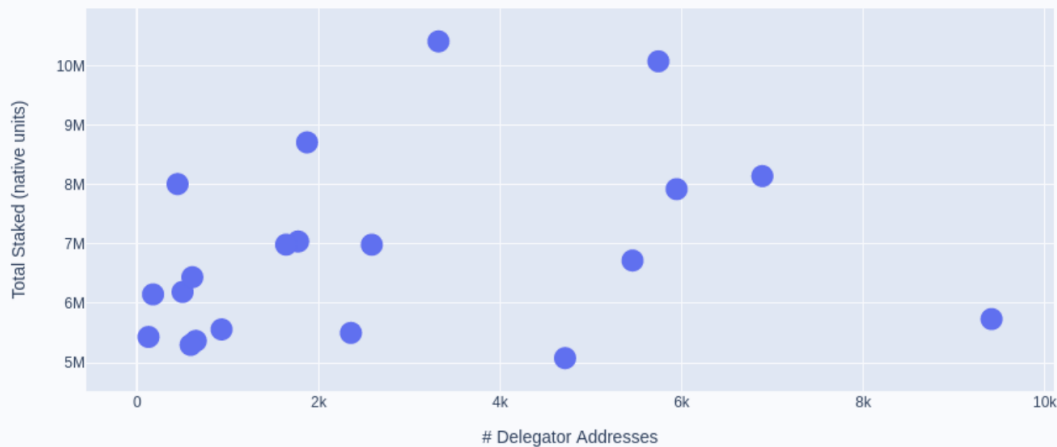
Stake by size brackets, by validator, for the top 20 validators. Terra 2.0 (July 2022)

For a validator, the decreased impact of very large delegations (driven by 1. Anchor 2. VC firms) corresponds to a change in strategy – pursuing individual parties is no longer the best path to success.

Significantly, 10 of the top 20 validators on the original Terra chain have continued their operations. Of these, only 4 are in the current top 20.

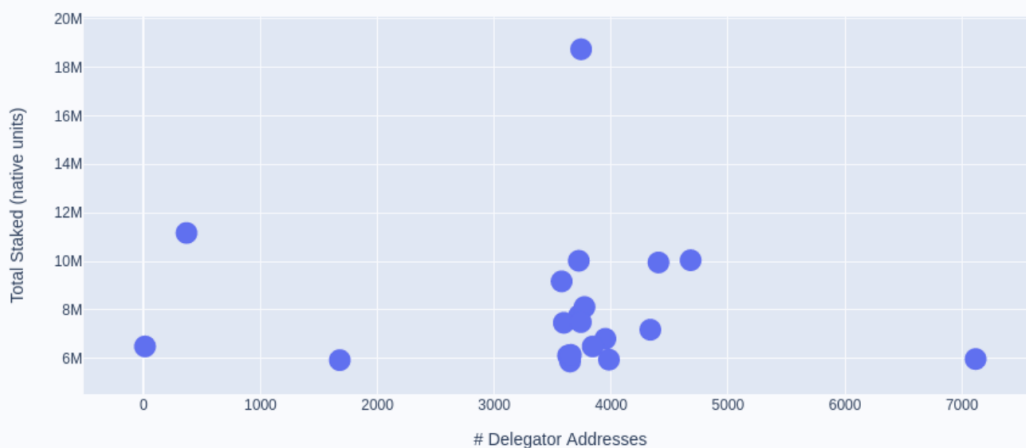
Plotting the entire top 20 set by the number of delegators and total stake shows the re-emergence of peer clusters in the absence of Anchor. These can be observed on most chains:

Pre-collapse (removing one outlier):



Top 20 Validators by total delegators and stake. Terra pre-collapse (April 2022).

By contract, post-collapse (again, removing one outlier):



Top 20 Validators by total delegators and stake. Terra post-collapse (July 2022).

For example, the second dot from the left (few delegators; significant stake) corresponds to an exchange wallet (the sole exchange validating Terra 2.0). The rightmost validator indicates disproportionate retail participation (many delegators; little stake) - in actuality, it is run by a YouTuber. Overall, the absence of Anchor is reflected by the re-emergence of firm-specific strategies.

The State of Avalanche Subnets in Q2 2022

Although the term "subnet" is used interchangeably and synonymously with blockchains, subnets are a bit more complex than that. "The definition of a subnet is technically: A Subnet is a dynamic set of validators working together to achieve consensus on the state of a set of blockchains, according to Avalanche's [FAQ](#) page." The unleashing and unlocking of subnets is an event of great importance in the wider web 3 ecosystem. Crabada and DeFi Kingdoms were the first Subnet successful tryouts. Crabada migrated from Avalanche C-Chain to its own subnet with surprising results. Crabada **gas fees** dropped from 80-100 gwei to 26-27 after the launch of its own subnet. The change also lowered the demand on C-Chain, lowering **gas fees** for all network applications. Currently, over 50 new subnets focused on the gaming sector are **in the making**. The high customizability also facilitates the development of institutional DeFi with native KYC functionalities.

A subnet may validate a large number of very customizable blockchains using various VMs. Avalanche is a blockchain platform that allows anybody to build highly customized blockchains. Subnets are thus an ideal platform for developers to construct customized games. The developer may design a subnet that is permissioned, semi-permissionless, or permissionless to

match their particular needs, as well as define sophisticated rule sets. Alternatively, one may simply join an existing subnet and incentivize validators to validate their blockchain, therefore sharing security with blockchains within that subnet. Entering the validator set of these games will be a key revenue stream for validators. There is also the option of selecting any virtual machine for each blockchain, such as the Ethereum VM, WASM, Solana's VM, and so on. Allowing current DAPPs to be simply transferred over and used in an environment for which the developer already has experience.

These are some of our favourite picks:

Shrapnel, the world's first blockchain-enabled AAA first-person shooter game, has announced that it would use the Avalanche network as its foundation for its impending release. They want to establish a subnet devoted to the game using the Avalanche Subnet capabilities. Shrapnel is creating a novel AAA experience for gamers that puts competitive multiplayer, creative tools, and genuine digital ownership front and center. Activision, creator of the famous Call of Duty games have **100 million** active monthly users according to their recently published financials. Even so, Call of Duty has lost **1 in 3** players compared to last year due to emerging competition. With an increasing adoption of the metaverse, we bet that Shrapnel will take a large share in the war gaming space. SHRAP is the unique ERC-20 token with multifaceted use on the Shrapnel platform, acting as both a governance token and a utility token. The SHRAP **token** will be utilized for governance, staking, as a gas token on the Avalanche subnet, and as a medium of exchange token on the Shrapnel marketplace.

Wildlife Studios has already established itself as one of the top ten largest mobile game producers in the industry, and it is now extending that expertise into the Web3 gaming landscape by integrating Castle Crush into an Avalanche network. **Castle Crush** is a free-to-play iOS/Android game that has had over **75 million** downloads and **200k** daily active users. Players must destroy their opponent's fortress while defending their own. To accomplish so, players utilize real-time cards (these will be NFTs) to react to their opponent's tactics. Players must have at least one NFT card to collect tokens. The Ascension Crystals (\$ACS) token may be obtained through daily in-game awards (expected returns of **8.67%** per month) as well as other activities such as tournaments and the selling of NFTs. Players will be able to spend \$ACS on card upgrades and participation in in-game events.

Increasingly, we are seeing a movement towards the multichain (app-specific-chain) environment. Networks like Avalanche and Cosmos support this. As more web3 decentralized apps migrate to subnets in order to scale, enhance user experience, and provide greater usefulness for their underlying currencies. It won't be long until web2 native teams and enterprises on the edge take action and begin supplementing their goods with web3 benefits and enabling their sovereign economies by bridging the simple bridge subnets enable. Just like the world's largest taxi firm does not own any vehicles, **Avalanche aims** to be the world's largest web3 infrastructure provider that does not own any servers.

Connect with Chorus One

Reach out to our Ventures team
ventures@chorus.one

Contact our Research team
research@chorus.one

Speak to our Sales Team
sales@chorus.one

Visit our website
chorus.one
