



# Quarterly Network Insights

**PREPARED BY** \_\_\_\_\_

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

Q1  
22

# Table of Contents

01

Growing the  
Chorus One  
Research Team

03

Announcing  
Chorus Ventures

04

Quicksilver Ramping  
Up For Launch

06

Our Framework For  
The Interoperability  
Space

09

Our Picks For The  
Metaverse In 2022

13

Ethereum Post-merge:  
An Economic Overview

16

Solana MEV Outlook:  
Are Sandwich Attacks  
Happening On Solana  
DEXs?

19

On-chain bot-spotting  
and delegator  
whale-watching on  
the Cosmos Hub

24

The Problem With  
Client Diversity

# Growing the Chorus One Research Team

Q1 2022 saw our research team grow to a team of 4 for the first time. We started off the quarter with Vish Modali departing Chorus One to become a Co-Founder at Quicksilver, a project that was incubated within the company. Fortunately for us, we were able to recruit three highly talented and capable researchers throughout the quarter to rebuild and reinvigorate our research team.

Umberto Natale was the first to join us during the quarter. Umberto completed a PhD in Physics at Ferrara University and has previously been working as a Research Associate in Cosmology at Cardiff University for the past 2 years. Umberto's academic research background adds a new element of masterfulness to Chorus One's research team. Umberto's current focus at Chorus One is researching MEV and cryptoeconomics; we recently published his first research piece, a scenario analysis and summary of eth2 rewards and penalties from the perspective of a validator on [ethresearch.ch](https://ethresearch.ch).

Next to join us was Gabriella Sofia. Gabriella studied Computer Science and Engineering, finishing her degree as recently as this year. Gabriella decided to do her thesis on distributed databases, which is what led her down the crypto rabbit-hole. As a result, Gabriella educated herself further in crypto by completing a blockchain developer program organised by Bloom (now Future Proof). Gabriella also has experience in design/UI amongst many of her other multitude of talents. Gabriella's current focus at Chorus One is on researching Ethereum and other EVM networks and understanding the business viability of a node product the company is working on.

In March, Michael Moser joined us. Michael is a crypto OG, first entering the space back in 2013. Michael graduated from HEC Paris with an MSc in Management & Public Policy, doing a thesis on equilibria in commodity markets. After completing his Master's degree, Michael led a culture institute associated with the Austrian embassy in Moscow for a year. In 2020, Michael started dabbling in crypto again, specifically in algorithmic trading. Michael's current focus at Chorus is to inform our internal strategy and market understanding by working in large part on-chain to analyze validator and delegator data across a variety of networks.

Finally, at the end of the quarter, we had crypto-native Luis Nunez join us. Luis is about to complete a Civil and Environmental engineering degree at Imperial College London. In his time studying, Luis joined the crypto research team at Imperial, where he spent time interviewing some of the biggest founders in Ethereum on the topic of governance. Recently, Luis joined [Occam.fi](#) (Cardano's launchpad) as a DeFi analyst before joining us. Luis' interesting background in sustainability engineering and experience mixing with renowned crypto founders on top of his knack for writing eloquently about complex crypto mechanisms makes him a great addition to our team.

# Announcing Chorus Ventures

Over the past 3 and a half years, we have developed unparalleled expertise in decentralised networks through launching and operating some of the most used protocols in the space. We have experienced what has worked for networks and what has not. We have provided input on token economics, mechanism design, node operation and key management solutions, validator relations, governance structures and GTM strategies. We value projects that we work with and always make sure to contribute in an impactful way. For this reason, we [announced Chorus Ventures](#), which will be a \$30m fund that is investing in outstanding decentralised networks, protocols and related products over 3 years. Chorus Ventures will be managed by our Research and Ventures Lead ([Xavier Meegan](#)), with oversight from [Felix Lutsch](#) and [Brian Crain](#). To date, we have made 26 investments in this space.

We believe that at least 10% of humanity's economic activity will run through PoS in 10 years time bringing freedom and prosperity through open and transparent decentralised networks. We leverage our multidimensional relationships with network ecosystem participants and expertise in node infrastructure operations to connect and secure our portfolio of networks.

Chorus Ventures is investing in Proof-of-Stake networks, interoperability, and middleware protocols and products. Both The Block and Coindesk covered the Chorus Ventures announcement, which can be read [here](#) and [here](#).

# Quicksilver Ramping Up For Launch

A silver lining from a few core team members from Chorus One departing last quarter was the fact that 3 of them teamed up to work on a new liquid staking protocol. Quicksilver is a liquid staking zone in the Cosmos ecosystem that is planning on launching mainnet in Q3 2022. We are excited to be leading [Quicksilver's seed round](#) and have been working closely with them as they prepare to go-live.

Apart from the well-known advantages of using liquid staking such as improved liquidity, advanced financial products and enhanced user experience that is solved by existing liquid staking solutions such as Lido, Quicksilver goes one step further to also enable governance by proxy, airdrop qualification, [inter-chain accounts \(ICA\)](#), and permissionless validator sets. In essence, Quicksilver is moving into uncharted territory with its network. There is \$5.5bn of staked ATOM that is currently locked-up, which will likely be the first asset that Quicksilver brings liquidity to. Should Quicksilver decide to onboard the next 7 biggest Cosmos networks by staked assets (excluding Terra), there is another ~\$4.5bn locked up that the network could bring liquidity to. Putting two and two together, there is ~\$10bn locked-up in staking across the top Cosmos networks right now, if Quicksilver can convince delegators that their solution is superior to regular staking, the Quicksilver TVL could grow, and fast.

Quicksilver will go-to-market as one of the first zone's to utilise [Interchain Security](#) and is also one of the first networks in the Cosmos ecosystem to leverage Interchain Accounts. In other words, Quicksilver is testing a lot of new features of the Cosmos SDK when it goes live, which is exciting in itself. For the first time, delegators can change their validator preferences across multiple networks inside one network (Quicksilver), whilst simultaneously holding onto their governance rights (i.e. being able to vote in other

networks and being able to reap the benefits of liquid staking). As for product market fit, there is no doubt that delegators in Cosmos appreciate governance flexibility. [Juno proposal 16](#) had a 98.45%(!) voter turnout as token holders voted on whether or not to punish the whale that gamed the original JUNO airdrop.

# Our Framework For The Interoperability Space

As economic activity increasingly moves across networks, interoperability protocols are growing at breakneck speed. This was recently best illustrated by Layer Zero's Stargate realizing TVL above \$3bn in less than a week after launch, but also insane exploits for a combined value of almost \$1bn that recently hit Wormhole and Ronin bridge. Cross-chain DeFi applications are hungry for collateral from other networks, and protocols such as Wormhole have helped them in their success by moving liquidity across the sometimes architecturally very differently designed networks.

Bridge building is one of the core challenges in crypto, and for the longest time centralized exchanges have served as the de facto bridge between the various isolated networks. But these are limited to token transfers and don't really fit into the interconnected, decentralized vision of the blockchain ecosystem.

From contributing to protocols like Cosmos IBC, integrating Lido stSOL into Anchor, and operating nodes for various interoperability networks, Chorus One has gathered a lot of experience in this realm. We know bridges are some of the most complex and fragile infrastructures that underpin the crypto ecosystem, which is why we are choosing our supported projects in this space carefully.

In this quarterly insight, we want to present our current thoughts and a short framework of how we mainly evaluate interoperability projects. We plan to release more extensive research on this topic in the coming weeks given the importance of it.



## Trust Assumptions

Not all bridges are the same, economic guarantees and risks can vary widely based on the soundness of cryptoeconomic design and its implementation. The high growth and fast-moving nature of the space exacerbates issues in the bridge space as user due diligence tends to boil down to checking the APY they can earn. Teams building bridges compete on getting as many integrations to different ecosystems done as quickly as possible, leading to protocols being entrusted with billions of dollars of collateral in a short amount of time.

Hailing from the Proof-of-Stake BFT ecosystem, we have confidence in the security of the established delegation model and capability of node operators in this industry provided the maturity of this space and underpinning protocols. Thus, we believe designs leveraging years of research into BFT and cryptography, as well as the experience of node operators (such as [Axelar](#)) are well-suited as a solution to the existing demand for interoperability between protocols with different consensus mechanisms allowing the ecosystem to move away from centralized bridges.

We believe trust-minimized protocols not relying on a separate consensus set will ultimately play a major role in the interoperability ecosystem, but also realize that designs in this space are currently largely limited to their native ecosystems. From our team's experience building a [light client for Tendermint in Solidity](#), we can posit that connecting networks with different consensus mechanisms come with complex engineering requirements and often necessitate severe compromises on UX (e.g. with respect to gas costs or finality, a point we will return to in the next paragraph). Many also still need to figure out how to decentralize core components. Even Cosmos IBC, one of the arguably most mature and tested "trustless" interoperability standard and protocol, has yet to figure out an incentive model for its relayer ecosystem. We are actively monitoring this space and will be placing our strategic bets with our capital and operational capabilities according to how we believe this ecosystem will evolve.

## User Experience

Another major factor that needs to be considered when evaluating interoperability protocols is looking at the user experience for bridgoors. If a protocol is hard to use, its adoption will be severely hampered. This applies to the experience on the front-end and associated latency, e.g. in the case of optimistic bridges such [Nomad](#), where workarounds to achieve instant finality are needed.

UX also matters a lot when considering what happens to your tokens when moving across chains. Are you receiving the native token on the other end? A wrapped version? A different wrapped version depending on which route you went through? A lot of this creates a highly confusing user experience, which combined with the notorious rehypothecation of DeFi protocols can lead to something that looks more like [that famed Welsh town](#) than anything you might consider a good investment opportunity. These issues and the plethora of bridging solutions also leads to fragmented liquidity, an issue for adoption in and of itself:

## Integrations and Liquidity

A core part of evaluating bridges is how well the associated team and supporters are set up to establish integrations and liquidity within networks and DeFi protocols. Interoperability is an exercise connecting communities and there are only a few teams with the capability, long-term vision, and hustle to pull off what is required for lasting success here.

We hope our framework will help you reason about interoperability protocols and wish you good bridging! Make sure to stay tuned and reach out if you are actively building, investing, or interested in our services in this space.

# Our Picks For The Metaverse In 2022

2021 was the year that NFTs took off. There was reportedly [~\\$25bn worth of NFT sales](#) that occurred throughout the year. Taking a step back from the enormous amount of sales and looking at the bigger picture of what it means, we can decipher the impact NFTs might have on the crypto industry.





For now, a key application of NFTs is to simply trade with other buyers and sellers. To some extent, the parabolic run we experienced in 2021 could be largely attributed to the rise of NFT buying and selling as it introduced a number of people to crypto that had never transacted in the space before. There were 2.3 million crypto wallets that traded NFTs in 2021, up from 75,000 the year before (3066.66% increase). Although NFT buyers increased significantly last year, the fact of the matter is that there are still only 2m people trading NFTs for what could be a global addressable market (a [Morgan Stanley report](#) estimated the NFT market to grow to \$300bn by 2030).

One of the key reasons NFTs increased so much in popularity last year was because the concept of an NFT is easy to understand. An NFT in its purest form, is a token that is non-divisible, which can be minted onto a blockchain any discretionary amount of times. In its simplest form, an NFT is just a non-fungible token (e.g. ERC-721, i.e. a token that is non-divisible) on a blockchain, which has a field to include an external link to a JSON file stored on an external server (e.g. IPFS). One of the most common ways to leverage the unique values and scarcity of NFTs is to include a JSON link to a JPEG hosted on an external server. It has never been easier for artists to upload and sell their work and never been easier for collectors to own art that has provenance on a blockchain. It is no wonder then, that art/collectible NFTs were the main driver of NFT sales in 2021 ([link](#)).

So where are we going with all of this? Well, if 2021 was the year of NFT trading, then we

expect 2022 to be the year of NFT displaying. What better way then, to display NFTs, than in a metaverse. The concept of a metaverse is another non-crypto native friendly concept (such as NFTs), which is relatively easy to understand. This is in large part due to monolithic corporations such as Meta that are now attempting to pivot into metaverse products. A key difference of course between web2 companies like Meta and web3 companies building their own metaverses, is that metaverse users own their own land and data in web3 metaverses, whereas Meta owns all the land and data in their web2 metaverse. For this reason, we believe that web3 projects can outcompete their peers in the web2 metaverse arena. In particular, we believe that Stargaze, Sintra and Passage are in pole position to make their mark in the metaverse vertical.

Stargaze is an interchain NFT marketplace that solves many problems that exist in NFT marketplaces today. Stargaze is a marketplace that exemplifies security, decentralisation, transparency and flexibility, which differentiates it from any existing competition from NFT marketplaces such as Opensea and MagicEden. In our view, Stargaze will be the go-to marketplace for NFTs in the Cosmos ecosystem, in large part due to the advantages above-mentioned.

#	Zone	IBC transfers 	Total Txns 	IBC share % 
1	 Osmosis 	<b>203,617</b>	<b>682,253</b>	<b>22,98%</b>
2	 Stargaze 	<b>68,862</b>	<b>66,581</b>	<b>50,84%</b>
3	 Cosmos 	<b>39,747</b>	<b>62,474</b>	<b>38,88%</b>

Source: [Map of Zones](#)

Stargaze had a successful 'genesis mint' in Q1 2022. There were 20 projects to mint on Stargaze during the 'genesis mint'. 19 out of 20 sold the total amount of NFTs available in their collection. The average amount of STARS collected per collection was 461,205 STARS (\$124,055.38). Every Friday night, the community votes for a number of collections to launch on Cosmos that STARS holders can mint. There are currently 88 projects in the pipeline. Using the current ~\$124,000 average raise from collections, we can assume >\$10m to be raised in coming weeks. Stargaze currently has a 10% mint fee that creators earn on their collections. We can use the 10% mint fee as an input from expected \$10m sales to find \$1m

in mint fees to be earned by creators. Stargaze introduced a mechanism to burn 50% of the 10% mint fees earned by collectors, meaning \$500k in STARS could be burnt if \$10m in sales is done in the next few weeks. The 50% fee burn from collector minting fees is an interesting economic experiment done by Stargaze, which could have a powerful impact when the network starts doing more sales volume.

We have previously released related content on [Stargaze launch](#), [Stargaze Liquidity Bootstrapping Pool \(LBP\)](#) and [What to Consider Before Minting an NFT on Stargaze](#). We would not be surprised to see development of a metaverse in Stargaze during the course of this year for people to display their NFTs acquired.

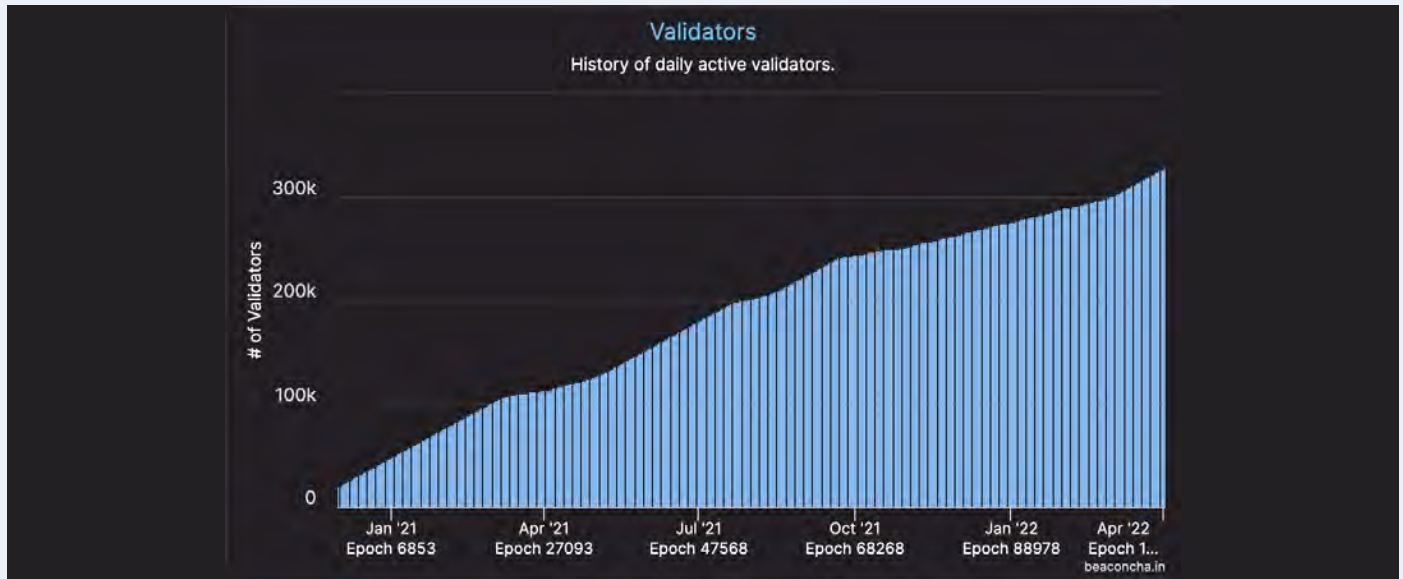
Another metaverse project we have been keeping a keen eye on is [Sintra](#). We have a strong relationship with the founder of the project, who has garnered an immense amount of experience in developing web2 products across Facebook and Amazon. Sintra is an NFT marketplace with added data analytics, social and security. There are numerous problems with leading NFT marketplaces such as OpenSea and MagicEden and we believe the race to be the number one NFT marketplace is far from over. The edge Sintra has on data analytics and security is immense compared to its competitors. We are excited to see the alpha version be built in Q2 2022.

Finally, Passage is another metaverse project that has captured our attention so far in 2022. Passage is a network launching in the Cosmos ecosystem that empowers virtual world builders with tools to create their own virtual worlds. Think of it like the Sims, only you own the scarce 3D assets that you use in your worlds. The first virtual world to go-live in Passage is the game Strange Clan. NFTs have already been minted for Strange Clan and currently have a floor price of 85 ATOM on the marketplace (\$2568). A reason for the escalated floor price of [Strange Clan](#) NFTs might be from the confirmed future airdrop of 5,700 \$PASG tokens to all Strange Clan holders. The NFTs, available on the Strange Clan marketplace, will be used as in-game characters whereby players can farm, quest and loot inside the game by themselves or with friends. Strange Clan architecture and design and gameplay is somewhat similar to an open-world RPG such as Runescape. One of the great things about

Passage is the way in which it leverages multiple networks across the Cosmos ecosystem so it can specialise on building virtual worlds, without having to concern itself with other components such as cloud (using Akash for GPU to ensure games can be used in browser) and NFT marketplace (using Juno to ensure NFTs can be purchased in ATOM). Passage has released multiple live demos of their product such as this video [here](#). The token economics of \$PASG are also unique given the amount of utility the token has, which will be announced publicly in upcoming months. We plan to onboard Passage network in the near future.


# Ethereum Post-merge: An Economic Overview

2022 seems to be the year of the long-awaited transition from Ethereum's Proof-of-Work (PoW) to Proof-of-Stake (PoS), and this transition is only one of [Serenity's upgrades](#) scheduled into Ethereum's roadmap. All pieces of Serenity's upgrades have the purpose of bringing Ethereum into the mainstream and serving all of humanity. It inevitably requires a scalable, secure, and resilient network. Improving scalability means making decentralized applications faster and cheaper to use. Following from scalability follows adoption. And as adoption grows, the protocol needs to become more secure against all forms of attack, as now important processes and economic activity rely on it. According to Ethereum's design principles, the network should be scalable and secure, but all of this without renouncing decentralization. Achieving these 3 qualities is a problem known as the scalability trilemma, and Ethereum's upgrade [aims to solve this trilemma](#). This is not an easy task, and solving it comes with a bunch of challenges. Precisely, to guarantee the safety of the network, a complex consensus is needed, and accounting for all possible scenarios, by achieving the best outcomes in all of them, requires to state stringent rules. As the network is designed to survive several types of attacks, the complexity of the implemented reward and penalty mechanisms increases. Complexity, sometimes, means high exposure for validators. Despite this, the number of active [validators on the Beacon Chain](#) is showing a clear up-trend, promoting high usage and liveness of the network.



Understanding how the Ethereum consensus protocol works and how its implementation affects validators is a key point needed to assess the efficiency of decision-making strategies. Further, understanding how the network works is crucial to guarantee its functionality: this is always our first objective at Chorus One.

Deep diving into Ethereum's cryptoeconomics, ended up with a [report](#) published on [ethresearch.ch](#). We established what a validator that respects the consensus design should expect in terms of annualized revenue.





In this figure we show the probability density function of reward distribution in a year for a single validator assuming a total of 291,362 active validators considering both an ideal case and a more realistic scenario, thereafter ideal and conservative cases. The ideal case does not take into account the participation rate and the probability of a validator being caught to be offline while being called to pursue his duties. The conservative case assumes that the participation rate is, on average, 99.25% and that the validator under scrutiny is online 99.99 % of the time. We find that the expected annualized reward for an ideal validator is 5.44%, a value that decreases to 5.4% if we take into account a more realistic case. This provides also a guideline to check if the status of our validators is respecting what the theory predicts.

We also provide a forecast of what happens if a validator is caught committing a slashable offense. This is done to quantify the risk of trying to override the rules, aiming to push other validators to make the right decisions - again privileging the liveliness of the network.


It is part of our Q2 2022 research plan to use what we found here to produce a guide to different scenarios that may happen in the context of client diversity. At Chorus One, we believe that knowledge is the key to guarantee longevity of a network, and we will push towards a deep analysis of what we should expect from the future.

# Solana MEV Outlook: Are Sandwich Attacks Happening On Solana DEXs?


Solana is a young blockchain, and having a complete picture of what is happening on-chain is a difficult task - especially due to the high number of transactions daily processed. The current number of TPS is around 2,000, meaning that we need to deal with  $\sim 10^8$  transactions per day. Figure below from [solana.fm](https://solana.fm) shows the daily number of transactions vs time.



The situation slightly improves if we consider only program activity. By only considering what happens on DEX, we are talking about 4107 transactions per day (source by [solana.fm](https://solana.fm)). This makes it clear that a big effort is required to assess which type of MEV attack is taking place and who is taking advantage of it, even because tools like [Flashbots](#) do not exist on Solana.



One of our first attempts is to extrapolate the frequency of [sandwich attacks](#). By considering only 5h of transactions on Orca, we can immediately see the preference of the user to select lower values of slippage. Further, we didn't find any trace of sandwiching happening on purpose.



Only 1% of the set of swap transactions we employed are surrounded, with 96.75% of them surrounded by swap performed by accounts that are doing most of the swaps. If it is a hint of a sandwich attack or a coincidence is hard to say at the moment, even if the likelihood points toward the second option. This is an accessible scenario due to the lower price fee on Solana, which raises the likelihood that a sandwich attack can happen by chance spamming transactions in the pool with the highest volume.

Part of Chorus One's research team goals in Q2 2022 is to better frame the statistical significance of this type of attack, extending the analysis to other types of MEV. A complete understanding of what's happening on-chain matters for the healthiness of the network, which we are committed to improve.

# On-chain bot-spotting and delegator whale-watching on the Cosmos Hub

The goal of this section is to give a brief introduction of our general approach to on-chain analysis, and to share some illustrative insights, for Cosmos Hub.

Chorus One's infrastructure holds in excess of \$3bn, delegated to us by a heterogeneous group, across almost 30 networks – we are on an ongoing quest to better understand what matters most to our delegators.

To meaningfully approach this question, we set out to understand who typically delegates to Chorus One.

To this end, we have been exploring novel ways of looking at delegation transactions, with a view towards identifying actionable patterns that hold for a significant share of delegators.


Initially, we were looking for a way to reliably identify 'malicious', automated delegations (e.g. delegations intended to game airdrops).

In a typical scenario, this would reflect as an unexpected but highly regular inflow of very small (smaller than 0.1 ATOM) delegations, in anticipation of a particular event.

A general assumption is that the presence of transaction costs imposes some kind of purpose on every transaction, i.e. if such a pattern were to occur, there is a rationale for undertaking further analysis.

In the case of automated transactions, there are several telltale elements we have identified. An initial way of honing in on unusual patterns is to cluster delegations by size and note abnormal increases (e.g. for very small delegations). These inflows can then be plotted – in the case of automated delegations, we would expect to see either a sharp increase or an unusually regular increase. The former case is a clear-cut indicator of automation.

In the latter case, it is worth looking out for further regularities. For example, there could be a certain time of day that sees a disproportionate amount of inflows. In our case, a corresponding chart, for Cosmos Hub, appears as follows (rendered via the excellent Plotly python library):




The blue line shows inflows to Chorus One, the dotted yellow line maps all delegations on Cosmos Hub. While there is a spike, it is in line with a general pattern and not outside the expected bound for our sample size.

With this as a starting point, we may extend the analysis to yield insights on ‘benevolent’ delegators. This is an area which implies more inference, as it covers a diverse group with unique behavioral patterns (i.e. in contrast to a single bot operator).

An assumption is that delegations drop off between 12AM and 6 AM, for a given time zone. By looking at delegation patterns during this time interval for different time zones, we may infer the relative ‘importance’ of the corresponding geographies to our Cosmos Hub validator.

Plotting this shows a tangible pattern:





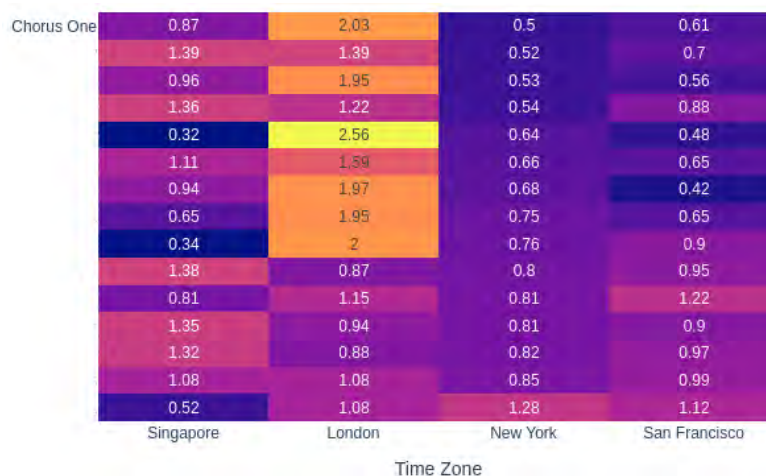
The heat map shows the relative delegations for Chorus One and some peer validators, between 12 AM and 6 AM in different time zones – sorted by “London”.

- Each row akin to the row labeled “Chorus One” represents a validator; we have anonymized other validators.
- The columns represent time zones (- represented by cities).
- The numbers indicate the change in delegations for a validator between 12AM and 6AM in a timezone, relative to other regions. A number below 1 implies a drop in delegations, vice versa above 1. This is color-coded.

Across the entire set, delegations typically drop off most during European nights and hold up best during Asian nights. Exceptions to the rule (e.g. a drop-off of 0.82 during nighttime in Singapore) can generally be qualitatively traced to a strong marketing presence, in a given region.

In our case, we see a (best-in-class) 2% drop-off in delegations during European nights. Overall, we are geographically well diversified, with a total absolute divergence from 1 of 0.2 – this is in the top 10% of large validators on Cosmos Hub.

Filtering by the top 10% of validators and sorting by “New York” demonstrates that most large delegations tend to happen during US working hours:





It is worth mentioning that further limiting the delegator set to the top 0.5% of accounts indicates an approximate parity for Asia and the US, with Europe lagging behind.

Integrating these and similar insights allows us to better understand network dynamics, communicate with delegators in a more targeted manner, and to further increase the resilience of our technical services.

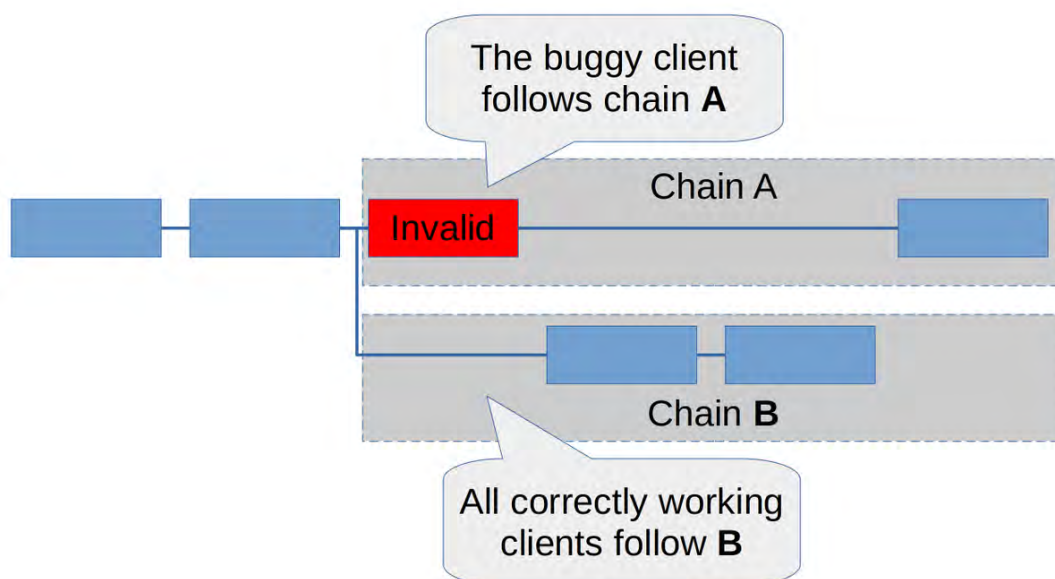
# The Problem With Client Diversity

One of the biggest stories for Ethereum this quarter that might have passed under the table in the last few days. The supermajority (>2/3) that Prysm used to hold as the most popular Ethereum consensus client **was dissolved**. This sets the network up for a safe Merge, as the worst case scenarios are fully averted. To understand how this became a cause for concern, we must go back in time and remember years ago, when the Ethereum Foundation offered grants to different client implementation teams for the Ethereum PoS blockchain. This early cohort would evolve, and by early 2020 it had narrowed down to four: Nimbus, Lighthouse, Teku, and Prysm. Prysm quickly became very popular amongst users, among other things by promoting their frontend-heavy testnet and growing a large Discord community.

The design principles behind the chosen multi-client architecture fall into the buckets of safety and liveness. Deciding on a set of standards to be implemented by independent teams in a variety of programming languages was a feature, not a bug. It was meant to reduce the fault in any individual client, cross-check to ensure correctness in the implementation and to limit the effects of an attack or failure to a portion of the network. But these strengths, unique to Ethereum, can only be realized if the chain is distributed as close to evenly as possible amongst the available clients. We could make several cases to explain why this is a priority, but let's explore the two that are most relevant to node providers: the "corrupted network" case and the "financial" case.

First, the "corrupted network" case. When a client holds a supermajority they alone can decide what block is considered canonical. If there are not enough validators on a different client to counter a bug or incorrect block (according to the Ethereum specs), then the chain will fork to an invalid chain that continues to operate, with no realistic way of reversing to a

different state. A corrupted client with  $>2/3$  of the total stake is able to finalize its own chain, and effectively create Ethereum's history. At a lower scale, a faulty client that still has  $1/2$  of the total stake could lead to an unrecoverable fork of the Beacon Chain. This would take a significant governance effort to solve, and in the worst case, the community could stay divided (similar to the fork that brought Ethereum Classic).





<https://dankradfeist.de/ethereum/2022/03/24/run-the-majority-client-at-your-own-peril.html>

If the client was used by  $>1/3$  of the validators, then there is no supermajority, but there would still be economic consequences for the network as a whole, with an activation of the inactivity leak. This penalty drains the stake of the offline validator for each time a slot is missed, thus reducing their voting weight so that the remaining validators can continue to finalize. Considering this, we should be aiming for no client to have more than 33% of client share. In such a case, the consequences would be negligible for the network and the individual validators, while the chain continues to make correct attestations and finalizing as normal. This is just an overview of some of the risk scenarios with single-client dominance.


The other case to consider for validators is a financial one. Client diversity is directly pushed by the [rewards and penalties protocol](#) in Ethereum's Beacon Chain. In order to encourage diversification of clients, we have correlated penalties, which means that penalties are light for isolated incidents but severe when many validators are penalized in a short time period, up to a total loss of stake. A 2/3 supermajority of corrupt validators, would result in its own forked chain, where the validators would not be able to return to the real chain without being slashed. And if 66% of the chain gets slashed simultaneously, the [penalty](#) would be the whole 32 ETH. Minor penalties are incurred on the ideal, <1/3, scenario, and they increase in severity and in burned stake when the inactivity leak comes into effect. The potential monetary loss should raise alarms for responsible stakers to take client diversity seriously.

We have mentioned consensus clients at length, but execution clients are not exempt from the risks mentioned above. These risks apply to both consensus and execution clients equally. In fact, the problem is even more pointed at the execution level, where there is a single point of risk with the Geth client. This risk has actually become a point of failure in the past, as exemplified by the [2016 Shanghai DoS attack](#). Similar examples where client diversity could have rescued the Ethereum consensus testnets have been observed. In 2020, the [Medalla testnet](#) suffered a big drop in active validators due to an issue in Prysm client's clocks, causing chaos and leading to a large number of slashings. As recent as the past weeks, the [Merge](#) of the Kiln testnet came close to the worst case scenario, when Prysm validators were unable to produce valid blocks.



At this point, it would be interesting to inquire about the differences between the major clients. Any analysis in this space will be impacted by the type of hardware, and type of user, for each of these clients. For example, Nimbus has been promoted as having minimal hardware requirements (able to run on phones), while on the other end of the spectrum, Teku is regarded as an enterprise solution that is more typically run on sophisticated hardware. One such analysis found out that there are in fact clear differences to be observed between the clients, but also that across metrics different clients came at the top for performance, proving that not all clients are made the same. Prysm, Lighthouse and Teku all showed strengths in different areas, which could impact the choice of one client above others.

Finally, let's take a look at how much client diversity of the Beacon Chain has changed in the past few months. The supermajority scenario has been avoided, and now the next step is ensuring no 50% scenario.



For many years, uptime has been regarded as the golden standard for measuring staking competence in PoS networks, especially Ethereum. But now we can see that another very pressing measure for competence would be a commitment to client diversity. With the rise of liquid staking platforms we have also seen a new standard for transparency, with Lido publishing [quarterly updates](#) on the status of their node operators, showing a very healthy distribution of clients; a group Chorus One is excited to be a partner of. For users that do not wish to run their own validators, these services are a route to contribute to better client diversity.

As we continue to be in this client diversity challenge, many development teams have voiced their stance, with statements from the people behind [Prysm](#), [Lighthouse](#), [Nimbus](#) and [Lodestar](#). Prysm in particular has made clear their opinion that “client diversity should be accomplished by focusing on software interoperability and standardization, rather than forcing it on stakers, which goes against the ethos of decentralization and free-and-open-source software”, concluding that competition between the clients is also a part of Ethereum. However, we are still treading on dangerous waters. If there's one metric we must pay close attention to in the run up to the Merge, it's the distribution of clients across the network.