HOC

House of Chimera Research: Decentralizing RPC Services - Challenges, Risks, and Solutions

This research delves into the current state and challenges of decentralized RPC services, emphasizing the critical role they play in the Web3 ecosystem. It explores the inherent risks of centralization within supposedly decentralized systems and evaluates the industry's competitive landscape.

May, 2024

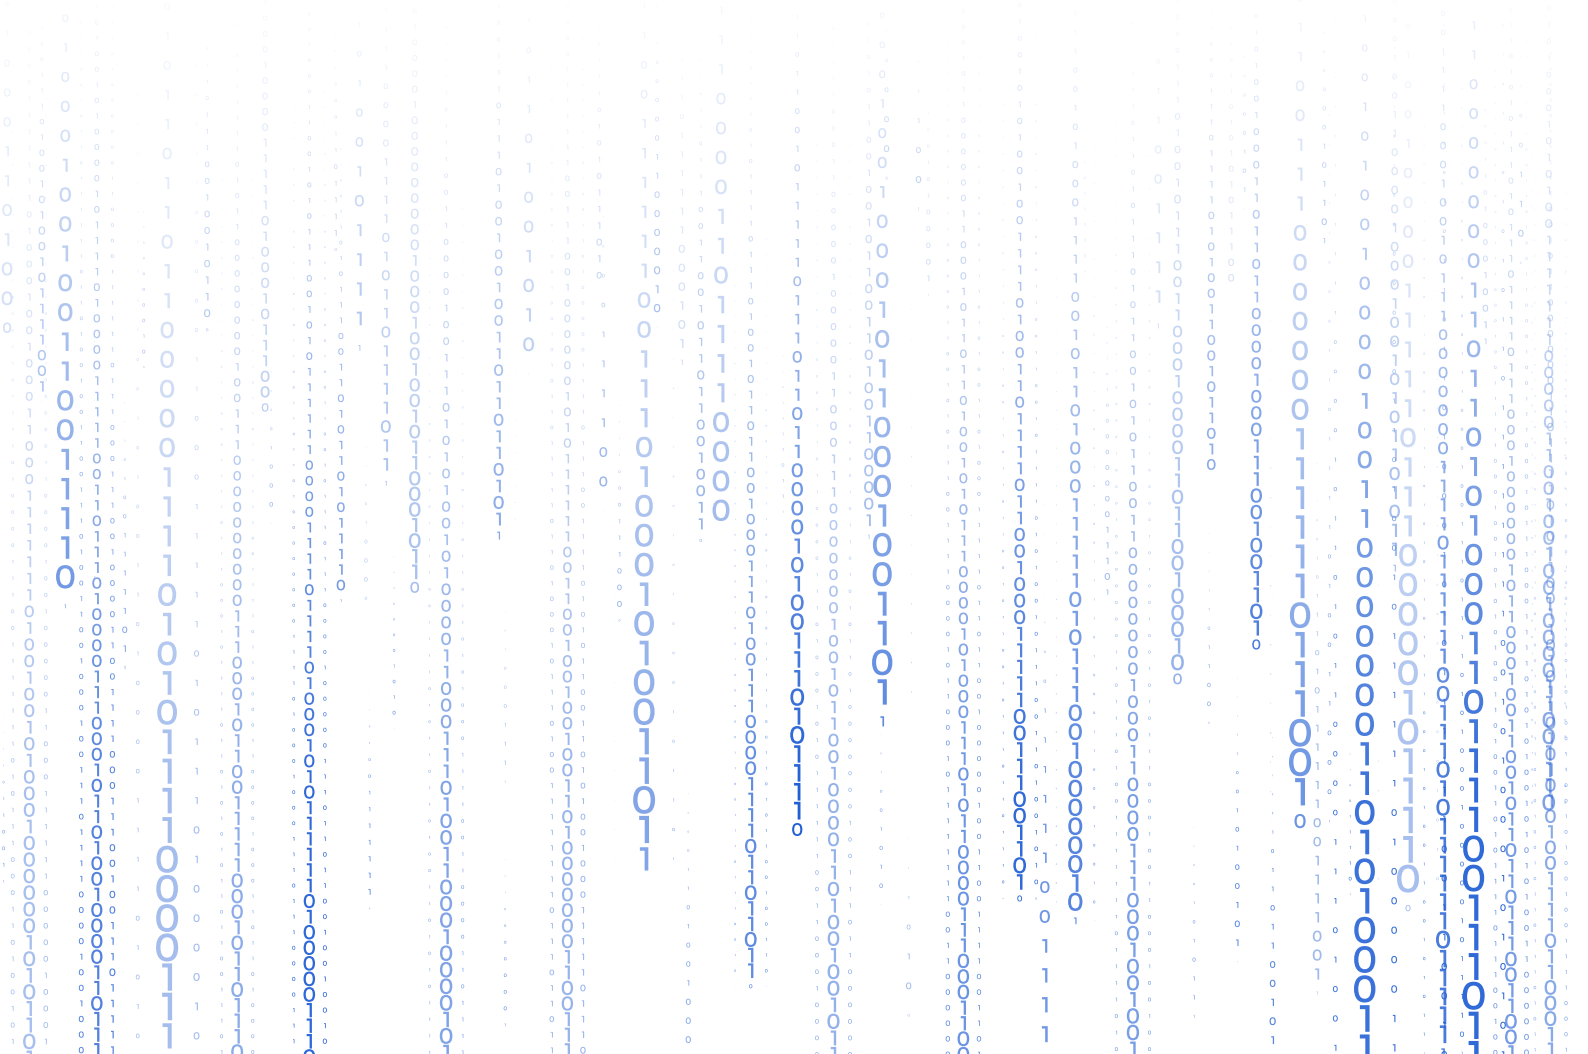

The content is for informational purposes only, and you should not construe any such information or other material as legal, tax, investment, financial, or other advice. Nothing contained in the research paper constitutes a solicitation, recommendation, endorsement, or offer by House of Chimera or any third party service provider to buy or sell any securities or other financial instruments in this or any other jurisdiction in which such solicitation or offer would be unlawful under the securities laws of such jurisdiction. All content of the research paper is information of a general nature and does not address the circumstances of any particular individual or entity. Nothing in the research paper constitutes professional and/or financial advice, nor does any information on the research paper constitute a comprehensive or complete statement of the matters discussed or the law relating thereto. House of Chimera is not a fiduciary by any person's use of or access to the research paper. You alone assume the sole responsibility of evaluating the merits and risks associated with using any information or other content of the research paper before making any decisions based on such information. In exchange for using the research paper, you agree not to hold House of Chimera, its affiliates, or any third-party service provider liable for any possible claim for damages arising from any decision you make based on information or other content made available to you through the research paper.

House of Chimera is an independent blockchain research and advisory firm, committed to integrity and transparency. We are fully transparent about our holdings and personal interests within SubQuery. House of Chimera has a financial stake in SubQuery through our investment in the SubQuery native token, SQT. Our integrity remains uncompromised in researching SubQuery, as the SubQuery team had no influence on our research at any stage.

HOC

Contents

HOC

Foreword

In an era where decentralization promises to reshape the digital landscape, this research paper examines the pivotal role of RPC services in the Web3 ecosystem. By exploring the complexities and risks associated with centralization within these services, we shed light on the challenges faced by developers and users alike. The study not only evaluates the competitive dynamics of the industry but also identifies emerging market opportunities. With a focus on SubQuery's innovative solutions, this paper aims to provide a comprehensive understanding of the future of decentralized RPC services.

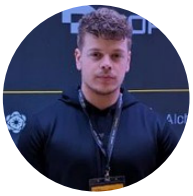

Diederick Jacobs

Founder

House of Chimera

HOC

Photo by NASA on Unsplash

Beyond Centralization: The Future of RPC Services in Web3

Exploring the growing demand for high-performance computing resources and the rise of io.net as a decentralized solution in the AI-driven landscape.

One of the main selling points of Web3 is decentralisation, meaning that there is no central authority that holds most of the power. At the heart of this transformation is the need to break away from the chains of centralization, by dismantling centralized structures of Web2, where mainly Web2 companies thrive by controlling data sources and thus monetization. Web3 lures users by permissionless and censorship resistant blockchain networks, consequently giving users back control, thus, changing how users interact digitally. The power flows back to the 'people' as the idea stands, allowing them to govern the decentralized applications (dApps) and thus, creating a truly decentralized sphere. At least, that is the vision perpetuated by developers by creating open-source protocols which are indeed governed by it's users. However, the first cracks in the truly decentralized and permisionless dream are showing through the Web3 dream.

RPC services allow dApps to perform functions or execute procedures on remote systems as if they were local, bridging the gap between the user's device and the blockchain. This is crucial because most dApps do not run their blockchain nodes but instead rely on third-party RPC services to interact with the blockchain network. The main issue is that running many blockchain nodes in either internet browsers or on mobile devices is currently impossible. This is not expected to change anytime soon due to several technical limitations, including processing power, memory, and storage constraints. Thus, RPCs facilitate seamless communication across different machines and systems, enhancing user experience by simplifying interactions with the blockchain. Users can, for instance, send transaction requests via their digital wallets, which are then processed by the RPC services, executed on the blockchain, and the results are communicated back to the users.

The foundational concept of RPCs isn't new. Developed initially by Xerox Parc in 1980[1], RPC protocols have significantly evolved. The rise of the internet and digital services saw the development of protocols like JSON-RPC in the 2000s, which uses JSON for data encoding. Alongside Representational State Transfer (REST), these protocols have become mainstays in web development due to their flexibility and ease of integration, prevalent in both traditional (Web2) and blockchain-based (Web3) applications.

Users who hold Ether in their non-custodial wallets interact with it in various ways, such as swapping it or transferring it to another network, buying NFTs, or using it within a DeFi protocol. All these activities depend on the RPC protocol.If the RPC protocol fails to operate, then the access and required data to support these interactions would be unavailable, essentially making it impossible to interact with these applications. Given the ever-growing blockchain states (e.g. 1.7TB for an Ethereum Full Node sync) the RPC services play an increasingly important role in the dApp ecosystems, scaling dApp clients to low-end mobile devices and web browsers.[2] Despite its importance, RPC services are relatively centralized compared with a blockchain network (i.e. Ethereum) and therefore more vulnerable against a single point of failure, which could lead to a full collapse of the dApp ecosystem. Recently, the RPC services of Manta Network were victim of a significant Distributed Denial-of-Service (DDoS) attack, as malicious actors were flooding the RPC nodes. According to Kenny Li, co-founder of P0xeidon Labs, the development team behind Manta Network, the RPC nodes were flooded with over 135 million RPC calls on a

[1] Birrell, A. D., & Nelson, B. J. (1984). Implementing remote procedure calls. ACM Transactions on Computer Systems (TOCS), 2(1), 39-59.

[2] Li, K., Chen, J., Liu, X., Tang, Y. R., Wang, X., & Luo, X. (2021). As Strong As Its Weakest Link: How to Break Blockchain DApps at RPC Service. In NDSS, 1-18.

single day.[3] Thus, essentially disrupting the communication channel between dApps and the blockchain infrastructure.

However, the protocols themselves are not the issue. The problem lies with the major companies offering centralized RPC services, which highlights one of the "cracks" in the ideal of a fully decentralized Web3. Despite the vision of decentralization, many dApps depend on these somewhat centralized services to function effectively. This reliance is a poignant reminder of the ongoing challenges in achieving true decentralization. It shows that while the technology has advanced, the infrastructure still leans on components that somewhat contradict the decentralization ethos.

In the realm of Web3 development, most third-party RPC providers are centralized. Due to the high costs and significant complexity involved in running their own RPC nodes, developers often opt for these third-party providers instead. The main issue with these providers is that they are centralized, leading to a single point of failure by nature, posing security and reliability challenges. To provide some perspective, imagine using a centralized internet service provider (ISP) versus creating and maintaining your own private ISP infrastructure. While ISPs provide the necessary connectivity for accessing the internet, Web3 RPC providers maintain node infrastructure for decentralized applications. Reliance on a single ISP can lead to outages and vulnerabilities, and dependence on centralized RPC providers introduces similar risks for decentralized applications. When utilizing a centralized ISP, internet access is contingent upon the stability and security of that single provider. Any disruption, whether due to technical failures, cyberattacks, or policy changes, can result in a loss of connectivity. Similarly, in the context of Web3, centralized RPC providers are responsible for facilitating communication between decentralized applications and the blockchain. Should these providers experience downtime, cyberattacks, or other issues, it can significantly impact the functionality and reliability of the decentralized applications that depend on their services.

SubQuery offers a decentralized alternative to the conventional centralized RPCs, enhancing the blockchain ecosystem by decentralizing the underlying infrastructure and addressing the problem of single points of failure. In addition to providing decentralized RPC services, SubQuery acts as middleware that facilitates easier interactions between complex blockchain databases and decentralized applications (dApps). It achieves this by indexing blockchain data, transforming it into a structured and easily queryable format. This service is particularly valuable for developers who want to build responsive and efficient dApps without investing heavily in data management. This paper aims to provide context and insight on SubQuery, while also evaluating the industry and competition.

[3] Peshkar, P. (2024). Breaking News: Manta Network Hit by RPC Flood Attack during Token Launch. What's Next? Crypto Ticker. https://cryptoticker.io/en/manta-network-rpc-flood-attack.

Centralization Risks in Blockchain RPC Services

The RPC landscape has an oligopoly market structure, whereby a few centralized providers dominate the industry, as developers and users often rely on major RPC providers for stable and reliable access to blockchain networks. Additionally, some major wallet providers (e.g. Metamask) do not allow you to change RPCs, essentially locking users in their centralized ecosystem. Metamask, the largest Ethereum non-custodial wallet, with over 30 Monthly Active Users (MAUs) is entirely reliant on Infura as its sole RPC provider.[4] Infura and Metamask are both created by Consensys, a blockchain software company focused on the Web3 industry.[5] According to recent research, the majority of cryptocurrency wallets do not disclose their RPCs nor allow you to modify RPCs, leading to significant centralization issues.[6] Less than a third of the used cryptocurrency wallets allow you to modify RPCs, by adding new RPCs, only 6 out of 28 show which RPCs they are using, and 5 of these 6 are centralized (e.g. Infura, Nodereal, and Alchemy).[7]

This is an issue as the centralization of RPC services has caused many security and censorship risks. In March 2022, Infura cut off Ukranian users for policy reasons, essentially censoring a part of their userbase. However, due to a too broad geographic IP block, other users out of different jurisdictions were also affected.[8] In another instance, Alchemy and Infura blocked access to Tornado Cash, a cryptocurrency mixer privacy service, thus preventing users from accessing the application.[9] It is also well known that Infura tracks users' wallets and IP addresses to enable geographical blocking of IPs, thereby defeating the core principles of decentralization and censorship resistance. RPC nodes, for the most part, serve as your gateway to the blockchain ecosystem.

Therefore, governments can potentially disrupt supposed dApps by coercing RPC providers to disable services or restrict access to blockchain networks within their borders. This is already being done on a certain extent, as leading RPC providers are currently censoring transactions from addresses sanctioned by the Office of Foreign Asset Control (OFAC), further degrading the censorship resistance nature of blockchain networks.

In addition to censorship concerns, centralized RPCs pose a risk by opening up a wide range of potential attack vectors that could be exploited for financial gain through user censorship. One scenario involves an auction (such as an NFT sale), where a bidder might delay the transactions of competitors by disrupting the RPC service they rely on, aiming to secure the auction item at an unfairly low price. This could be achieved through a DDOS attack, effectively silencing certain users. Another example concerns a user attempting to deposit into a hash-time-lock contract (HTLC), a commonly used mechanism in dApp development for facilitating cross-chain swaps. In this case, an attacker could delay the user's ability to withdraw their deposit after the lock period expires by disrupting the RPC service tied to

---

[4] Consensys (2024). MetaMask Reveals 55% Surge in Users, introduces Default Security Alerts to Drive Wider Adoption and prevent Billions Lost to Fraud. Consensys. https://consensys.io/blog/metamask-reveals-55-surge-in-users-introduces-default-security-alerts-to.

[5] See https://consensys.io/products.

[6] Yan, K., Zhang, J., Liu, X., Diao, W., & Guo, S. (2023). Bad Apples: Understanding the Centralized Security Risks in Decentralized Ecosystems. In Proceedings of the ACM Web Conference 2023, 2274-2283.

[7] Ibid.

[8] Benson, J. (2022). Ethereum's Infura cuts off Users to Separatist Areas in Ukraine, accidentally blocks Venezuela: Infura says it accidentally changed its settings too broadly after Venezuelans complained of being blocked. Decrypt. https://decrypt.co/94315/ethereum-infura-cuts-off-users-separatist-areas-ukraine-accidentally-blocks-venezuela.

[9] Sun, Z. (2022). Alchemy and Infura Block Access to Tornado Cash as Vitalik Buterin weighs in on Debate. Cointelegraph. https://cointelegraph.com/news/alchemy-and-infura-block-access-to-tornado-cash-as-vitalik-buterin-weighs-in-on-debate.

the withdrawal process, temporarily freezing the fund.[10]

Additionally, there are clear privacy concerns. In a centralized RPC system, all requests between the user's application (like a wallet or a decentralized application) and the blockchain pass through servers controlled by a central entity. This concentration of data flow creates a single point of control, allowing the service provider to monitor, log, and analyze all incoming and outgoing traffic. As a result, the entity can potentially track user activities, including tracking which addresses a user queries, the transactions they initiate, and their interaction patterns with different smart contracts. The ability to monitor user interactions with the blockchain opens the door to data harvesting. Centralized RPC providers can collect vast amounts of metadata, such as IP addresses, transaction timestamps, and frequency of interactions with specific blockchain services. Over time, this data can be aggregated to profile users, inferring their habits, preferences, and possibly even their real-world identities. Such profiles are not only a privacy concern but could also be leveraged for targeted advertising, sold to third parties, or even handed over to authorities upon request. Using a centralized RPC service requires users to place a significant amount of trust in the provider. They must trust that the provider will not only maintain service availability but also safeguard their privacy and act in the users' best interests. This reliance on trust is at odds with the trustless nature of blockchain technology, where systems are designed to operate without needing to trust any single party.

Figure 1 Overview of Basic RPC (Remote Procedure Call) Operations

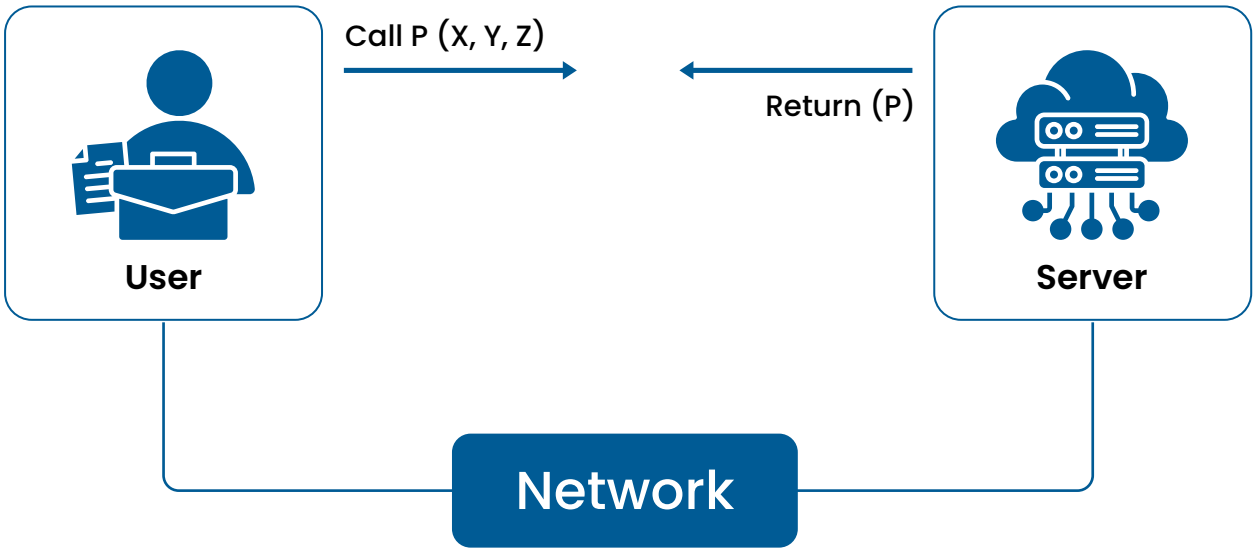

[10] Li, K., Chen, J., Liu, X., Tang, Y. R., Wang, X., & Luo, X. (2021). As Strong As Its Weakest Link: How to Break Blockchain DApps at RPC Service. In NDSS, 1-18.

How does a blockchain work?

Truly decentralized ecosystem with verifiable, self governing, permissionless payments and mechanisms, allow anyone to access services equally, as there is no personal data required to be part of the ecosystem. There is no hierarchy, and therefore there is no centralization within any of the used components, meaning that ownership is distributed among its builders and users. As every service is decentralized, there is no single point of failure, as there are multiple nodes.

Blockchain technology serves as the foundational infrastructure for decentralized ecosystems, enabling the construction of additional services atop its framework. Ethereum stands out as the most prominent relatively decentralized platform at present, supporting the cryptocurrency Ether (ETH) along with a myriad of dApps. In contrast, centralized exchanges (CEXs) are platforms that facilitate the exchange between fiat currencies and cryptocurrencies. They implement rigorous know-your-customer (KYC) policies and are subject to regulatory oversight according to their operational protocols, similar to traditional institutions (i.e. Web2). CEXs are the most Web2-like element in the Web3 space, increasingly taking on characteristics similar to traditional financial institutions.

Web3 or cryptocurrency wallets provide a bridge between conventional financial systems and the digital currency landscape, offering management of digital assets and facilitating interaction with dApps. DApps are independent applications that combine a backend powered by smart contracts with a user-friendly frontend interface. Ethereum utilizes high-level programming languages like Solidity for creating smart contracts, which are then converted into bytecode executables on the Ethereum Virtual Machine (EVM). Ethereum clients offer a variety of RPC commands allowing decentralized services to engage with the blockchain through activities such as data retrieval and transaction submissions. Due to the lack of blockchain nodes operated by most decentralized services, they often depend on external RPC providers to process transactions, which are then relayed to the broader blockchain network for recording by validators.

RPCs in Web3 Gaming

In a blockchain-based game, when a player wins a tournament, the game utilizes a blockchain RPC endpoint to update the player's balance. Initially, the application queries the player's current balance via the RPC endpoint. Following the tournament win, it requests the blockchain to increment the player's balance, triggered by a transaction from the game's account.

Figure 2 Flow of Interactions in Blockchain Ecosystems.

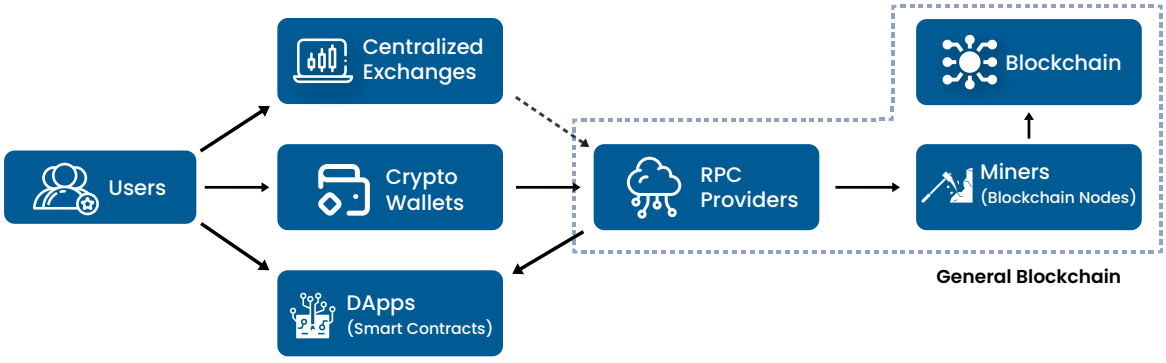

The graphical overview illustrates a Proof of Work (PoW) blockchain ecosystem.

Decentralizing RPCs in Blockchain Networks

The evolution of RPCs from centralized to decentralized models marks a pivotal shift in the blockchain ecosystem, aiming to address the significant security, censorship, and user experience issues prevalent in centralized RPC protocols. These issues not only pose a threat to the integrity and economic stability of blockchain networks but also limit user autonomy in terms of RPC modification and choice, particularly evident in the interaction with cryptocurrency wallets. Centralized RPC providers can censor users, and additionally, there are significant security risks involved in using centralized RPCs.

Decentralized RPC systems aim to change the centralized landscape by distributing user requests, whether for querying blockchain data or submitting transactions, across a vast network of nodes instead of a singular endpoint. This model empowers users to partake in network operations (e.g. providing node services), often requiring them to hold native tokens and operate specialized hardware. Such tokens serve dual purposes: acting as a deposit to safeguard against malicious activities through a slashing mechanism, and incentivizing global user participation, thereby enhancing the network's robustness through geographical dispersion.

Decentralized RPCs strive to hand control back to users, allowing them to dictate the terms of their data sharing and interactions, thus fostering a more ethical, sustainable, and user-centric Web3 industry. Privacy and security receive a significant boost through advanced data handling techniques such as data packet splitting and routing through multiple nodes. These methods ensure user data remains fragmented and obscure across the network, considerably enhancing privacy and security compared to centralized alternatives.

Furthermore, decentralized RPC networks employ advanced routing algorithms that take into account factors like geographic location, node performance, and current load to efficiently process requests with minimized latency. Despite these advancements, the challenge of latency remains, especially in applications requiring real-time data, as requests must navigate through multiple network hops, potentially leading to delays. Nonetheless, the distributed nature of these networks offers unparalleled reliability and uptime, mitigating the impact of individual node failures and ensuring a resilient global infrastructure.

Figure 3 Overview of the Decentralized RPC Operations of the SubQuery Network

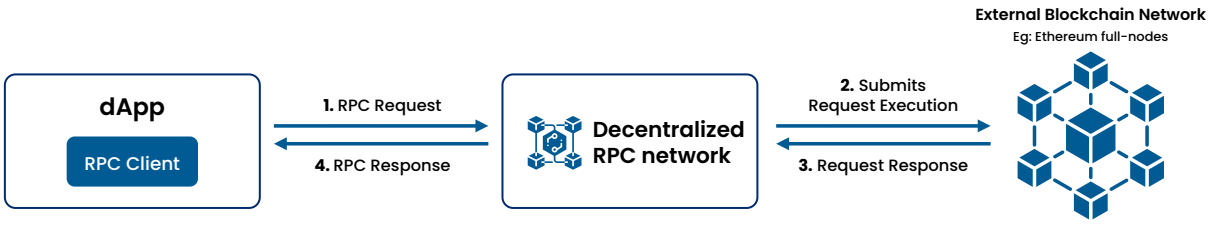

Drawbacks of decentralized RPCs

Despite the transformative potential of decentralized RPCs in creating a more open, transparent, and resilient digital world, the journey is fraught with complex challenges inherent to the technology. These challenges include scalability and performance issues that could hinder network throughput, as well as the complexities involved in developing sustainable economic and incentive models to ensure the ecosystem's viability. Latency emerges as a critical concern, particularly when accessing data from nodes situated far from the request origin, potentially degrading the user experience in real-time applications. Solutions such as proximity-based and dynamic routing offer a potential solution by improving response times through intelligent request distribution, enhancing user experience for time-sensitive operations. However, these solutions must navigate the delicate balance of optimizing performance without compromising user privacy or exposing sensitive location data.

Ensuring consistent quality of service across a decentralized network, characterized by diverse node capabilities, presents another challenge. Implementing performance standards, reputation systems, and adaptive routing mechanisms are potential strategies to address service variability, prioritise higher-performing nodes, and adjust to real-time network conditions. Decentralized RPC networks lower the barriers to entry for new participants because anyone with the necessary hardware and network capabilities can join as a node provider. This contrasts with an oligopoly, where high entry barriers can limit the number of competitors.

If we look further into the decentralized RPCs industry there are only a few ways they can compete, as the products they provide are similar but not substitutes. As decentralized RPCs, while the core service (blockchain data access) is mostly standardized, differentiation can come from factors like query speed and data reliability. As the main way to compete is through price, the demand is very elastic, meaning that consumers are very responsive to price changes. In this scenario, the primary consumers are dApps, which utilize RPCs to establish connections with the blockchain. Consumers will change from one RPC to another RPC, solely on price increases, assuming that the quality remains similar. Additionally, the decentralized model shifts the onus of node maintenance and upgrades onto individual operators, necessitating aligned incentives to foster network health.

The role of economic incentives in decentralized RPC networks cannot be overstated. These incentives are multifaceted, designed to motivate node operators to not only participate in the network but to also ensure the highest standards of performance and reliability. This is crucial because, in a decentralized setting, the network's health and scalability depend largely on the collective efforts of individual node operators. Incentives typically come in the form of transaction fees and token rewards. Transaction fees ensure that operators are compensated for the computational resources they dedicate to processing queries and transactions. Meanwhile, token rewards can serve multiple purposes: they can incentivize participation, reward longevity and reliability, or even penalize malicious service provider behavior. For instance, operators might earn more tokens for maintaining uptime or providing faster responses, aligning their interests with those of the network's users.

However, these incentive mechanisms also carry potential pitfalls. If not carefully calibrated, they can lead to adverse outcomes such as token inflation. Excessive

issuance of tokens as rewards can dilute the value of the currency, undermining the financial stability of the project and eroding stakeholder trust. This was observed in the case of Pocket Network, where a significant allocation of revenue to supply-side incentives led to substantial inflation, negatively impacting its ecosystem.[11] Additionally, poorly designed incentive models, can inadvertently favor larger operators as they can benefit from economies of scale, leading to a more centralized network structure. Smaller operators, feeling marginalized or unable to compete, might push for changes or fork the network to create a more equitable environment. Moreover, while penalties for poor performance or malicious acts are necessary, they must be balanced to avoid disincentivizing participation. The challenge lies in designing an incentive model that motivates node operators to improve their infrastructure and contribute positively to the network while avoiding excessive punishment that might lead to a decrease in participation or even network forks.

Is decentralized, truly decentralized?

The cryptocurrency industry champions the ideal of decentralization, aiming to distribute control away from central authorities. Yet, a significant paradox lies at its core, particularly within infrastructure projects and dApps (e.g. DeFi). These initiatives, while decentralized in intent, often rely on centralized cloud providers such as Amazon Web Services, Azure, or Google Cloud for their operation. Recent studies highlight this contradiction, revealing that around 70% of DeFi applications opt for centralized hosting solutions like Cloudflare, which, along with AWS, supports the infrastructure of protocols managing roughly $90 billion in funds.[12] This reliance poses risks, as demonstrated by historical outages on platforms like Cloudflare, underscoring the potential for reputational and financial damage despite no direct loss of funds.[13] Users can still access their assets via alternative interfaces or direct smart contract interactions, but the ease of use significantly diminishes, questioning the extent of decentralization in these projects.

The centralized nature of the cloud provider market worsens this issue, with a few dominant players exerting significant control and influence, leading to higher prices, customer lock-in, and heightened censorship risks, among other concerns. Decentralized cloud providers offer an alternative, promising a more distributed approach to hosting. However, scaling these decentralized networks presents its own set of challenges. Coordinating a diverse network of individually owned nodes to deliver consistent, reliable service is inherently more complex than managing centralized data centers. This results in variability in service quality, as the decentralized nature of these systems introduces inconsistencies in hardware, network connections, and node availability.

Despite these challenges, the movement toward decentralized cloud services represents a shift on the spectrum of decentralization, away from the centralized status quo. This transition is not without its issues, but it suggests a path forward where, over time, as more participants join the decentralized ecosystem, the benefits of reduced intermediary costs could outweigh the economies of scale enjoyed by centralized entities. This potential for cost savings, alongside the ideological drive for decentralization, fuels the hope that decentralized cloud providers will gradually overcome their scaling hurdles, leading to a more distributed, resilient, and user-empowered internet.

---

[11] Casella, M., & Grigore, M. (2023). State of Pocket Q2 2023. Retrieved April 25, 2024, from: https://messari.io/report/state-of-pocket-q2-2023.

[12] Winter, P., Lorimer, A. H., Snyder, P., & Livshits, B. (2021). Security, Privacy, and Decentralization in Web3, 1-11.

[13] Graham-Cumming, J. (2019). Details of the Cloudflare Outage on July 2, 2019. Retrieved April, 25, 2024, from: https://blog.cloudflare.com/details-of-the-cloudflare-outage-on-july-2-2019/.

Figure 4 Estimated Funds at Risk During Hosting Provider Outages

HOSTING PROVIDER	AFFECTED (USD)
CLOUDFLARE	61.6B
AWS	28.1B
FASTLY	2.4B
DIGITAL OCEAN	0.9B
OTHER	0.8B

The table is based on 2021 data.

"The paradox of decentralization: cryptocurrency relies on centralized cloud providers. Overcoming this could lead to a more resilient, distributed, and user-empowered internet."

Chapter Summary

One of the main selling points of Web3 is decentralization, eliminating central authorities that hold most power. This transformation seeks to dismantle centralized structures of Web2, where companies control data and monetization. Web3 entices users with permissionless and censorship-resistant blockchain networks, empowering users to govern decentralized applications (dApps) and create a truly decentralized sphere. However, the ideal of decentralization faces challenges, as cracks in the dream appear.

RPC services enable dApps to function remotely, bridging the user's device and the blockchain. This is crucial because most dApps rely on third-party RPC services to interact with blockchain networks, as running many blockchain nodes on browsers or mobile devices is currently impossible due to technical limitations. RPCs enhance user experience by facilitating seamless communication across different machines and systems. Users can send transaction requests via their digital wallets, processed by RPC services, executed on the blockchain, and communicated back to the users.

Users with Ether in non-custodial wallets interact with it in various ways, such as swapping or transferring it, buying NFTs, or using it within DeFi protocols. These activities depend on the RPC protocol. If the RPC protocol fails, access to required data and interactions would be unavailable. Given the growing blockchain states (e.g., 1.7TB for an Ethereum Full Node sync), RPC services play an increasingly important role in scaling dApp clients to low-end mobile devices and web browsers. However, RPC services are relatively centralized compared to blockchain networks, making them vulnerable to a single point of failure, potentially leading to the collapse of the dApp ecosystem.

In Web3 development, most third-party RPC providers are centralized due to the high costs and complexity of running their own RPC nodes. Centralized RPC providers pose security and reliability challenges, leading to outages and vulnerabilities. For example, Infura and Metamask's reliance on centralized services exposes users to risks. SubQuery offers a decentralized alternative, enhancing the blockchain ecosystem by decentralizing underlying infrastructure and addressing single points of failure. SubQuery acts as middleware, indexing blockchain data and transforming it into a structured format, valuable for developers building responsive and efficient dApps without investing heavily in data management.

Overall, while Web3's vision of decentralization promises a significant shift from Web2's centralized structures, the reliance on centralized RPC services poses challenges. Decentralized solutions like SubQuery offer promising alternatives, but achieving true decentralization requires overcoming significant technical and infrastructural hurdles.

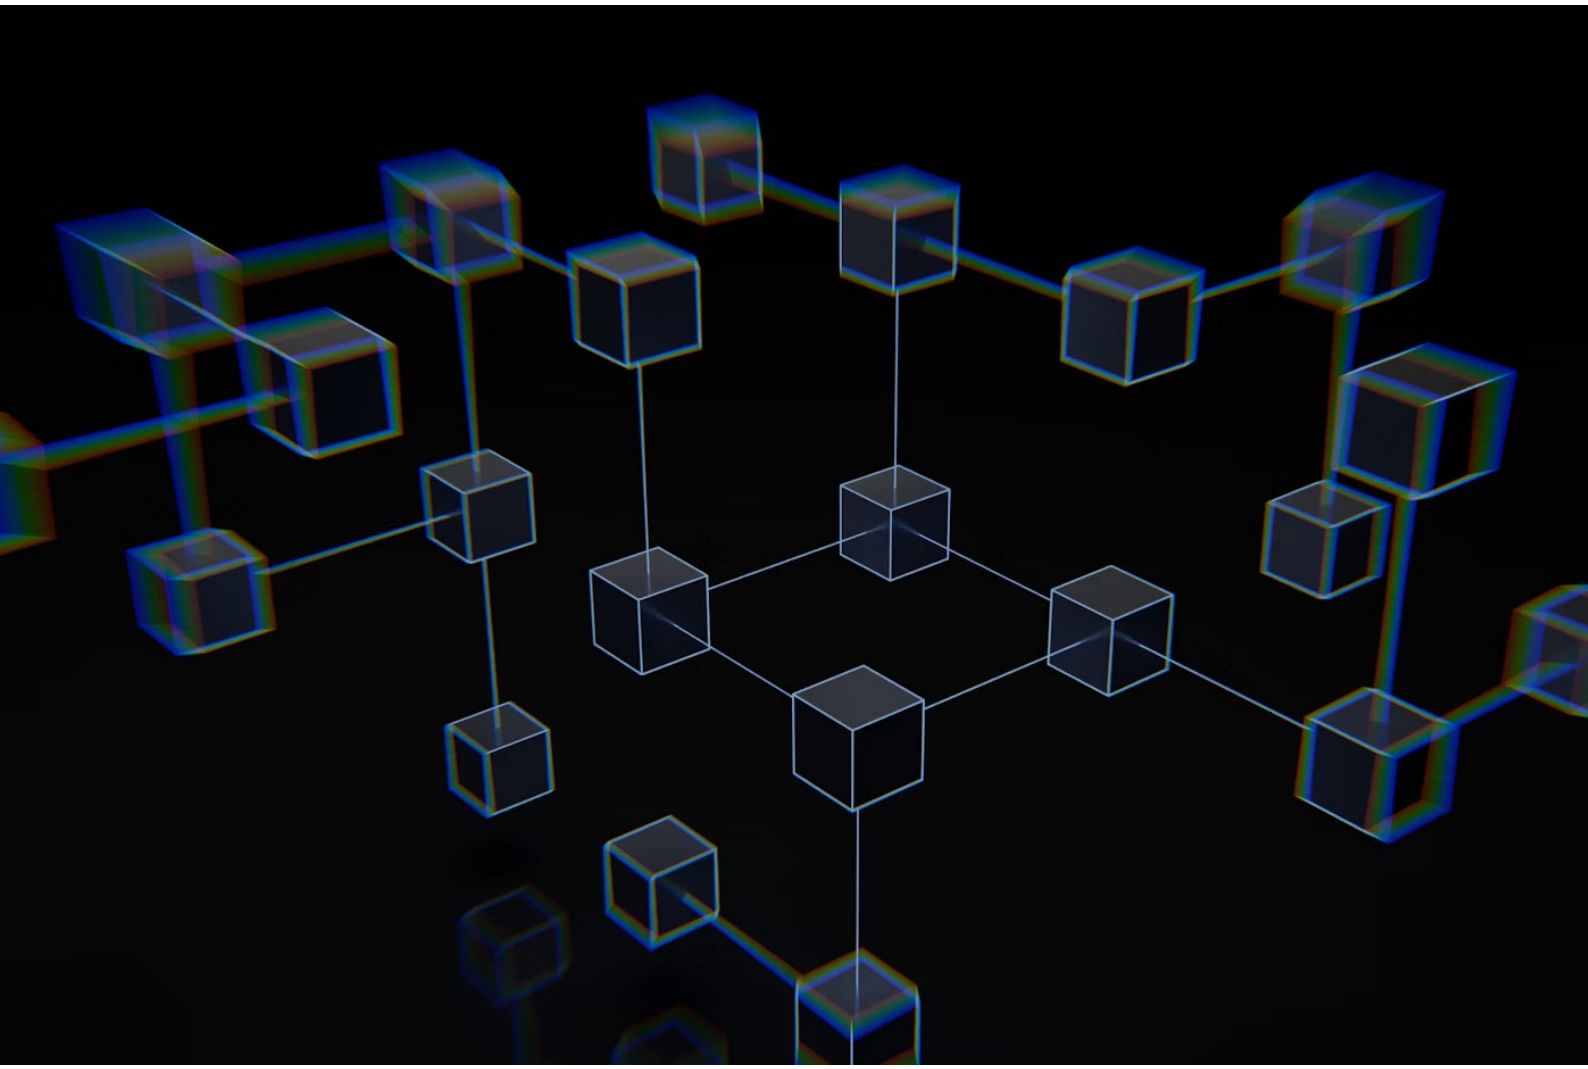

Photo by Shubham Dhage on Unsplash

SubQuery: Enhancing Blockchain Data Access

A detailed analysis of SubQuery's middleware architecture, decentralization strategies, and reward mechanisms for improving blockchain data querying and interaction.

SubQuery is developed to mitigate the challenges associated with processing and querying blockchain data, positioning itself as middleware that simplifies the interaction between complex blockchain databases and dApps. By indexing blockchain data, SubQuery transforms it into a structured and easily queryable format, addressing a crucial need for developers who aim to create responsive and efficient dApps without dedicating extensive resources to data management.

The architecture of SubQuery is designed with a focus on decentralization and scalability, aligning with the foundational principles of blockchain technology. The system incorporates two types of nodes: Data Indexer nodes and RPC nodes, each serving a specific function within the ecosystem. To participate as a node operator, users are required to stake 200,000 SubQuery Tokens (SQT) as a security deposit, which is approximately 55,000 USD as of writing.[14] This deposit acts as a deterrent against malicious behavior, with the possibility of being partially forfeited in the event of such actions. Should the deposit fall below the required threshold due to a slashing event, the operator must replenish it to the initial 200,000 SQT to continue their participation. Consequently, the moral hazard decreases as there is a significant punishment for malicious actions.[15]

Data Indexer nodes are responsible for collecting, processing, and storing raw blockchain data. This process involves organizing the data according to predefined schemas, making it readily accessible for quick retrieval. The indexed data is made available to developers through an Application Programming Interface (API), allowing for efficient queries using standard query languages. This setup facilitates the real-time retrieval of data, which is crucial for the performance of dApps. Moreover, SubQuery offers flexibility for projects with specific data requirements, enabling them to define custom indexing criteria to ensure the data is tailored to their applications' needs. On the other side, RPC nodes are designed to handle real-time interactions with the blockchain, processing requests from dApps to execute blockchain functions. These nodes serve as a bridge between dApps and the blockchain, managing request execution and returning the results to the dApps. This role is critical for applications that rely on up-to-the-minute data and interactions with the blockchain. As stated before, anyone can setup a node as long they hold the threshold of tokens and have access to the required hardware.

One of the major benefits of SubQuery is its ability to consolidate data into a single, easily accessible marketplace, all under one token. This integration significantly enhances accessibility within the ecosystem by centralizing data access, making it readily available to users and developers alike. By reducing the fragmentation of information, SubQuery effectively minimizes data silos, allowing for seamless and efficient retrieval of relevant information. This unified approach not only streamlines the process of accessing data but also fosters a more cohesive and user-friendly environment. Developers no longer need to navigate multiple sources or deal with disparate datasets, thus saving time and resources. Furthermore, the data marketplace supports better interoperability among different applications and services within the ecosystem, promoting a more interconnected and efficient Web3 infrastructure.

---

[14] See https://www.coingecko.com/en/coins/SubQuery-network

[15] Baron, D. P., & Besanko, D. (1987). Monitoring, Moral Hazard, Asymmetric Information, and Risk Sharing in Procurement Contracting. The RAND Journal of Economics, 509-532.

The SubQuery network implements a scoring system for node operators, which is derived from metrics such as uptime statistics, the amount staked, and results from security assessments. This scoring mechanism provides users with insights into the reliability of node operators for their required services. As previously discussed, operators who fail to meet reliability standards, or are found to be malicious, risk being penalized through slashing. SubQuery plans to enhance this framework by incorporating additional metrics, such as historical performance data, aiming to reduce the information gap between consumers and providers. This reduction in information asymmetry means that adverse selection is reduced, empowering consumers with the knowledge they need to make well-informed choices. This aligns with the principles of a properly functioning market, where informed economic participants (i.e. consumers) make decisions leading to the efficient distribution of resources (i.e. capital). However, adverse selection complicates this ideal scenario by creating situations where one party possesses more or superior information compared to the other, rendering price signals less effective. This discrepancy leads to prices that don't truly represent the value or risks of the goods or services exchanged, thus elevating transaction costs and resulting in less efficient market outcomes.

SubQuery employs an open-source public gateway as an intermediary between requesters (i.e. consumers) and service providers (i.e. node operators). This setup ensures consumer requests are routed through the gateway before reaching node operators, effectively anonymizing the interaction by obscuring the direct link between consumers and service providers. Initially, the gateway was closed-source, and our research highlighted the risks associated with this, emphasizing how it could perpetuate the status quo. As a result, the SubQuery team has expedited the process of making the gateway open-source, enabling users to run their own private gateways for the network.[16]

Alternatively, consumers have the option to directly route their requests to the GraphQL API using SubQuery's Software Development Kit (SDK). This SDK enables the direct interaction with the GraphQL API via Apollo Link, a comprehensive toolkit crafted specifically for GraphQL API operations. As a result, this approach reduces latency since requests bypass a centralized gateway. However, it also means that the consumer's location is potentially exposed to the node operator that the request is routed to. On the other hand, consumers can utilize their own logic for node operator selection with the SDK, and can potentially also utilize other mechanisms to hide their location.

Currently, enhancing the gateway's algorithm is a key focus for the SubQuery team. Planned upgrades include the introduction of geographical selection, which prioritizes nodes closest to the consumer based on latency, price, stability, and other criteria, without relying on precise locations. This approach ensures that users' locations are not disclosed, thereby minimizing their counterparty risk. This approach aims to boost the ecosystem's efficiency and competitiveness by reducing latency, a critical challenge in decentralized RPC networks.

[16] Github (2024, April 24). Commits. https://github.com/subquery/network-gateway-service/commits/main/.

Figure 5 SubQuery Network Schematic Overview: Indexers and RPC Providers

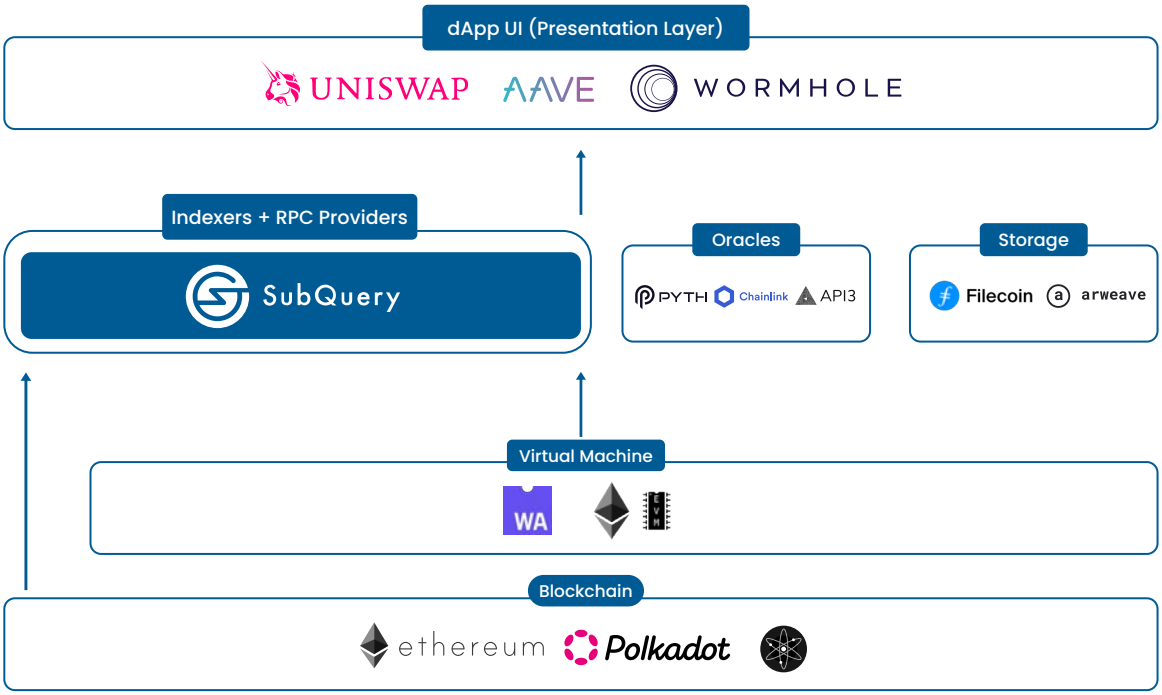

The Reward Mechanisms of SubQuery

In the SubQuery network, Node Operators are essential contributors who ensure the availability and quality of data services. To participate in the network and qualify for rewards, Node Operators must initiate their service status with an on-chain transaction and commit a minimum stake in SQT. This staking mechanism serves as both a security measure and a financial commitment to the network's integrity and performance. Rewards for Node Operators are structured around several key activities: providing data services, hosting deployments, and meeting specific service agreements. At the core of the rewards system is the Cobb-Douglas production function, which allocates rewards based on the volume of serviced requests and the amount of SQT staked. The function describes how various inputs, usually labor and capital, affect output production. For SubQuery, these inputs are staked tokens and the volume of service requests. Essentially, a Node Operator's income from the pool is influenced by their share of responded requests relative to others and the size of their SQT stake compared to peers. This mechanism incentivizes operators to maximize productivity while increasing their total stake in SQT tokens. Operators can increase their stake through additional delegated stakes or acquire more SQT tokens directly. This approach ensures an equitable distribution of rewards, incentivizing both the provision of high-quality services and significant token staking. Delegators augment the system by staking their tokens with Node Operators, earning a share of the rewards generated. This relationship enhances the Node Operators' potential reward pool, as a higher collective stake attracts more rewards. Node Operators can adjust the reward percentage for delegators every era (approximately every seven days), further incentivizing delegation. Node Operators' overall stakes, including both their tokens and those delegated by others, play a crucial role in determining their share of rewards. A higher stake indicates a greater

commitment to the network, qualifying operators for a more significant portion of the rewards. This approach motivates Operators to bolster their stake, thereby strengthening the network's security and reliability.

For each deployment managed by a Node Operator, a distinct reward pool exists. The size of this pool is influenced by the demand for the data or services provided, targeted network incentives, and overall network inflation. The rewards system is divided into three main categories: Closed-Agreement Rewards for services rendered under specific contracts, Pay-As-You-Go (PAYG) rewards based on actual usage of the services, and Inflation Rewards that incentivize the staking of tokens towards a deployment.

Closed Agreements Rewards involve a direct financial agreement between a single node operator and a consumer, bypassing the use of the Cobb-Douglas Production Function. This setup mirrors traditional Business-to-Business transactions in the Web2 space, where services are delivered within a specific timeframe for a predetermined fee. This model ensures predictable income for node operators, fostering financial stability and operational reliability. To mitigate potential monopolistic pricing by node operators, agreement offers can be placed on SubQuery projects by consumers, inviting more node operators and thus enhancing competition. This increase in supply, against a steady demand, is aimed at driving prices down to achieve more favorable market rates. A SubQuery project in the SubQuery network is essentially an indexing project created by a participant known as an Architect. These projects enable Node Operators to efficiently index and query the data. Architects can develop these projects by either creating entirely new datasets or by transforming existing datasets, such as Subgraphs, into the SubQuery format. This facilitates easier and more effective data handling and retrieval within the network.

PAYG, offer a pay-per-use model where consumers pay according to their actual consumption. Node operators list their services and pricing for specific SubQuery Projects or RPC endpoints. Consumers must pre-lock the required tokens in a state channel, which are then allocated to the node operators at the end of each Era based on their contributions, as determined by the Cobb-Douglas production function. By locking SQT to boost a project, consumers not only motivate node operators but also earn network incentives(i.e. free queries), which can be utilized to compensate node operators through state channels. This mechanism positions PAYG as an optimal choice for most consumers, enabling them to effectively manage their expenditures on data queries from node operators, especially when sufficient SQT is allocated to boost a project.

Figure 6 Comparison of SubQuery Payment Models

	PAY AS YOU GO	CLOSED AGREEMENT
COST PER REQUEST	Relatively High	Lower at High Volume
REWARD DISTRIBUTION	Cobb-Douglas Pool	Direct
SOURCE OF FUNDS	From Wallet or Booster Reward	From Wallet

In addition to these rewards, the network also distributes "Inflation Rewards" to support projects considered "Public Goods", such as essential network infrastructure and SubQuery projects without direct financial backing. These inflation rewards ensure Node Operators are compensated for their contributions to these valuable but not directly monetized projects. Inflation rewards consist of newly minted tokens amounting to 1.2% of the total supply annually. From this amount, 1% is allocated to support Projects, while the remaining 0.2% goes to the SubQuery treasury. Consumers can direct inflation rewards towards specific Node Operators by boosting their projects with SQT. This mechanism increases the financial incentives for Node Operators to prioritize certain projects, effectively directing network resources toward areas with heightened consumer interest. Boosting also benefits the consumers by granting them free queries on the projects they support, encouraging active participation in the ecosystem.

The network's reward mechanism is multifaceted, incorporating productive work, inflation rewards, and consumer-driven incentives. The SubQuery Foundation plays a pivotal role in sustaining this ecosystem, providing grants, from the SubQuery Treasury, to architects who create or migrate valuable Indexing projects to the network. These grants, funded by a portion of the network's inflation, reward architects for their contributions and encourage the development of projects that meet consumer demand. Architects can also directly sell their services, leveraging their expertise to fulfill the specific needs of the market.

Figure 7 SubQuery Ecosystem: Reward Distribution and Staking Flow

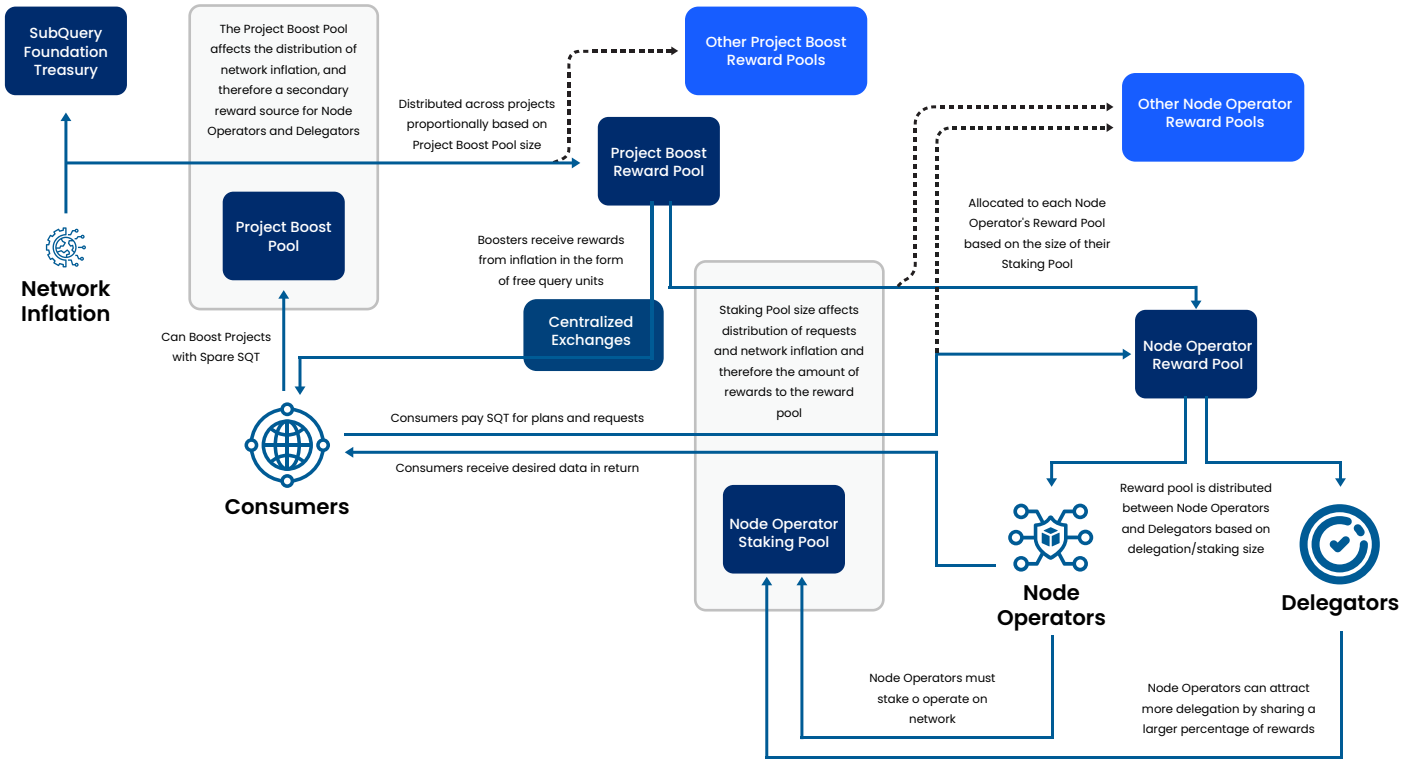

What is the SubQuery Token?

The SubQuery Native ERC-20 Token serves as a utility token on the Ethereum and Base Network, primarily incentivizing node participation and functioning as a medium of exchange within the SubQuery Network (see appendix for token distribution). Its utility drives demand, with consumers requiring SQT for network transactions to access node operator services, who in return receive compensation. Additionally, operators can leverage their tokens to enhance their network stake, potentially boosting their rewards from a pool comprising Inflation Rewards and PAYG. The token also attracts delegators seeking to augment their delegation and, consequently, their portion of the rewards from this pool. Looking ahead, the SubQuery Network's move towards greater decentralization includes introducing governance mechanisms, and potentially utilizing SQT for voting or council membership, based on token-holding thresholds. Despite the intention for equitable governance, challenges persist, notably the disproportionate influence of venture capitalists and early investors, capable of skewing voting outcomes—a disparity evident in unbalanced vote distributions, a scenario atypical in democratic processes.

In response, innovative voting designs like the Quadratic Voting System have been suggested, where votes cost quadratically, making dominant influence harder to achieve, thus diluting the potential for vote manipulation (Buterin et al., 2019; Lalley & Weyl, 2018; Weyl, 2017). However, this system isn't foolproof, as a malicious actor could distribute tokens across multiple wallets to amplify their voting power, a tactic known as a Sybil Attack. To counteract this, digital identities have been proposed, providing a means to verify the uniqueness and authenticity of each participant without compromising privacy. The Quadratic Voting System has been used by Subquery in April 2024.[17]

Another governance approach is Commitment or Bond voting, where votes are weighted by the duration a voter agrees to lock their tokens if the vote passes, emphasizing long-term commitment to governance and mitigating against both Sybil attacks and the rise of plutocracy (Berg, Davidson, & Potts, 2020; Mohan et al., 2022).

Figure 8 Understanding Quadratic Voting: Cost Implications

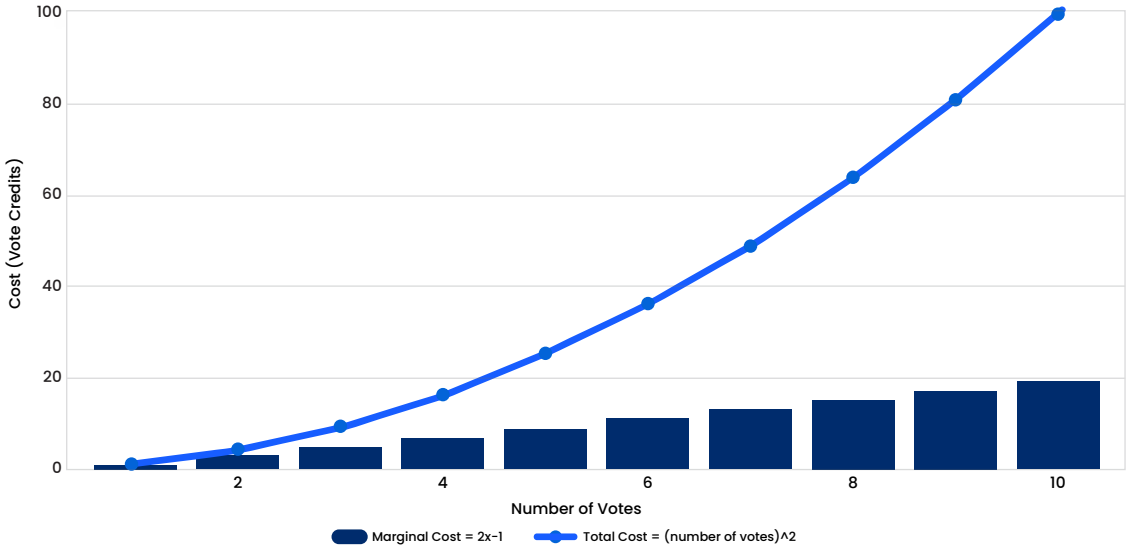

This graph illustrates the Quadratic voting mechanism, where the cost of votes increases quadratically, enhancing fairness and reducing potential for vote manipulation.

[17] SubQuery Network (2024, April 11). SubQuery Foundation Executes First Governance Vote. Medium. https://subquery.medium.com/subquery-foundation-executes-first-governance-vote-1610d0c4554f.

What is the SubQuery Foundation?

The SubQuery Foundation plays a crucial role in the governance and strategic direction of the SubQuery Network within the Web3 landscape. Its primary objective is to establish a robust governance framework that fosters an inclusive, transparent, and decentralized ecosystem. The Foundation oversees key governance decisions and guides the strategic development of the network, aiming to ensure alignment with the principles of decentralization.

The SubQuery Foundation's responsibilities include managing a Treasury of SQT tokens to fund various ecosystem initiatives. This management is vital for maintaining financial stability and supporting innovation within the network. By administering grants programs, the Foundation incentivizes developers and users to engage with the SubQuery Network, thus promoting its usage and fostering a diverse ecosystem of partners and users.

Additionally, the Foundation develops educational programs and materials to support developers at all levels, ensuring they have the resources needed to effectively utilize the SubQuery ecosystem. It also supports community-driven initiatives and organizes hackathons to build a robust network of partners and enhance community participation.

Operationally, the SubQuery Foundation is expected to handle day-to-day tasks related to business development, marketing, and product management. This includes organizing community events, managing marketing strategies, and overseeing the development roadmap. The Foundation encourages contributions to the SubQuery SDK and Network contracts from various organizations and individuals, promoting continuous development and innovation.

To streamline governance, the Foundation proposes a structured decision-making process involving subcommittees responsible for grants, technical oversight, and marketing. This structure aims to improve the efficiency and effectiveness of governance within the SubQuery Network.

As of the time of writing this research paper, the SubQuery Foundation has not yet been fully implemented, though the initial steps towards its integration have been taken. This research has accelerated this process by highlighting the risks associated with maintaining a centralized governance structure, particularly in critical areas such as identifying and addressing malicious behavior. Consequently, the SubQuery team has introduced a roadmap to Community Governance on the 8th of May to expedite the integration process.

Team Overview

SubQuery boasts a robust team (See Appendix) of highly skilled engineers, supported by a well-structured Business Development (BD) and Operations team. Headquartered in Singapore, the team has key representation in New Zealand, Lisbon, and Sydney.

The engineering team specializes in blockchain technology, decentralized infrastructure management, and automated code deployments. They have extensive experience working on large-scale infrastructure projects and contributing to major blockchain platforms and technologies, including Bitcoin, Ethereum, Tezos, and CENNZnet.

The operations team ensures the efficient functioning of the organization by managing day-to-day activities and strategic planning and execution, drawing on extensive experience from previous roles in Web2 tech companies.

The commercial team is committed to providing an excellent onboarding experience for SubQuery customers. Their focus includes delivering technical consultancy, developer-oriented content, and technical education to support customers in developing and scaling their projects.

SubQuery's leadership team brings a wealth of experience from various high-profile tech companies and successful blockchain projects. Their expertise spans product engineering, technical consultancy, business development, digital marketing, and community engagement. Overall, SubQuery's team combines deep technical expertise with strong operational and commercial service capabilities, positioning the company as a strong contender in the blockchain and decentralized infrastructure space. No negative press or OFAC alerts were found for any C-level team members.

"SubQuery's skilled engineers, efficient operations, and dedicated business development teams position it as a strong contender in blockchain and decentralized infrastructure, with global representation and extensive experience."

Chapter Summary

SubQuery is designed to address the complexities of processing and querying blockchain data, serving as middleware that simplifies interactions between blockchain databases and decentralized applications (dApps). By indexing blockchain data, SubQuery converts it into a structured and easily queryable format, enabling developers to create efficient dApps without extensive resource dedication to data management.

SubQuery's architecture focuses on decentralization and scalability, featuring two types of nodes: Data Indexer nodes and RPC nodes. Data Indexer nodes collect, process, and store raw blockchain data, making it accessible through an API for efficient querying. RPC nodes handle real-time blockchain interactions, executing requests from dApps and returning results. Node operators must stake 200,000 SubQuery Tokens (SQT), acting as a security deposit to deter malicious behavior. The network employs a scoring system for node operators based on uptime, staked amount, and security assessments, providing users insights into node reliability. SubQuery aims to enhance this framework by incorporating additional metrics to reduce information asymmetry, empowering consumers with better decision-making tools. SubQuery uses an open-source public gateway to anonymize interactions between consumers and service providers. Consumers can also directly route requests to the GraphQL API using SubQuery's SDK, which reduces latency but may expose their location to node operators.

A major advantage of SubQuery is its ability to consolidate data into a single, accessible marketplace, reducing information fragmentation and enhancing interoperability within the ecosystem. This unified approach streamlines data access, saving developers time and resources while fostering a more cohesive environment.

The reward mechanisms for Node Operators include providing data services, hosting deployments, and meeting service agreements. The Cobb-Douglas production function allocates rewards based on the volume of serviced requests and the amount of SQT staked, incentivizing productivity and significant token staking. Additional rewards include Closed-Agreement Rewards, Pay-As-You-Go (PAYG) rewards, and Inflation Rewards for supporting non-monetized projects.

The SubQuery Native ERC-20 Token (SQT) is used within the network for transactions and incentivizing node participation. Future plans include introducing governance mechanisms where SQT holders can participate in decision-making. The Quadratic Voting System, implemented in April 2024, is expected to be used for future governance decisions. The SubQuery Foundation plays a crucial role in governance and strategic direction, managing a Treasury of SQT tokens to fund ecosystem initiatives and support innovation. It oversees key governance decisions, administers grants programs, develops educational resources, and organizes community events to promote network engagement.

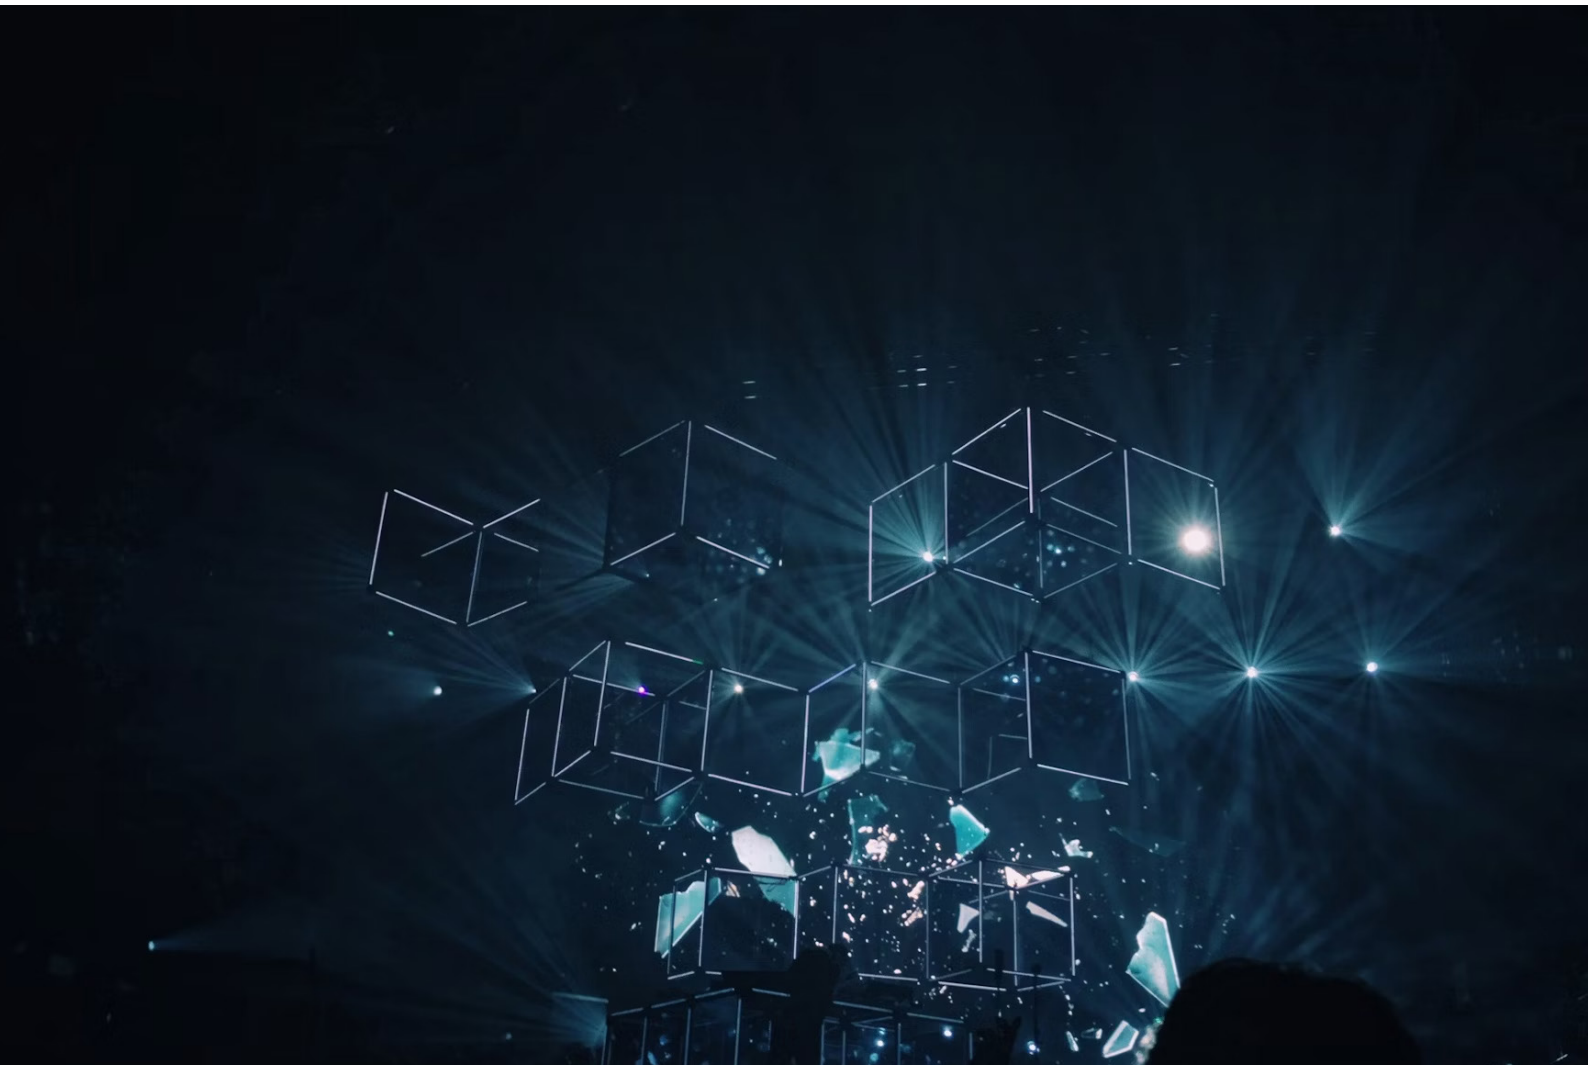

Photo by fabio on Unsplash

Navigating the RPC Industry: Centralized vs. Decentralized

Analyzing the competitive dynamics, challenges, and market opportunities in the RPC industry, with a focus on the transition from centralized to decentralized services.

The RPC industry is currently marked by intense competition involving both centralized and decentralized entities. Centralized players, however, hold a notable competitive edge, primarily due to lock-in strategies. These strategies revolve around creating a closed ecosystem, preventing users from switching between RPC providers. This situation presents significant challenges, as discussed earlier in this research paper, and fosters an inefficient market characterized by diminished competition due to these lock-in mechanisms. Consumers often find themselves with no alternative but to commit to a specific RPC provider, trapped by their initial choice of wallet. An open market could potentially foster genuine competition, provided there is no emergence of a monopolistic environment. This presupposes that wallet providers will eventually allow users the flexibility to switch RPC services. Nevertheless, a monopolistic market could still develop if there are limited participants in the market (i.e., RPC providers) or if information is not adequately disseminated. Such circumstances could lead to economies of scale for a handful of providers, eventually establishing a monopolistic market where the likelihood of disrupting the current balance is low. This scenario would notably disadvantage consumers by diminishing their surplus.

The scale of centralized RPC providers plays a crucial role, making it difficult for new entrants to capture a significant portion of the market due to high barriers to entry. Specifically, establishing a scalable decentralized RPC service requires a substantial network of decentralized nodes. As previously discussed, many nodes that are purportedly decentralized actually rely on centralized cloud services for cost efficiency at scale. Moreover, drawing in node operators is challenging without offering sustainable financial incentives. These incentives must be structured within a closed system, where the total issuance of tokens is capped to avoid substantial inflation. In such a system, rewards for node operators are derived from existing tokens paid by consumers, rather than minting new tokens for compensation. This approach prevents long-term unsustainability, which could lead to considerable token inflation, thereby exerting downward pressure on token value and damaging credibility.

The primary competition among RPC providers often revolves around pricing, a metric that may not support sustainable long-term growth. Eventually, prices could stabilize near the base cost of processing an RPC request, theoretically fostering a fiercely competitive environment beneficial for consumers and innovation but challenging for less efficient providers. Another critical competitive aspect is latency, where some companies, like SubQuery, are making notable strides, especially in indexing performance. Competing through technological innovation offers a more sustainable path by enhancing service quality for consumers without necessarily compressing profit margins. SubQuery's recent launch of SDK 4.0 exemplifies this approach, boasting significantly faster block indexing capabilities than many competitors.The key enhancements include the introduction of multi-threading, which allows the system to handle multiple requests in parallel, significantly improving response times and overall efficiency. Additionally, the updates reduce the frequency of database write operations, minimizing the load on the system's storage resources. Furthermore, the SubQuery Data Node has improved its indexing efficiency through the integration of a filtering RPC API. This enhancement enables more precise targeting of data queries, which reduces the retrieval of irrelevant data. As a result, this lowers the overall volume of requests and decreases potential data overhead, making the system more streamlined and effective in

handling user queries. Currently, SubQuery achieves indexing speeds approximately 3.9 times faster than The Graph, one of the leading Web3 Indexing providers.

Figure 9 Comprehensive Overview of Existing RPC Solutions

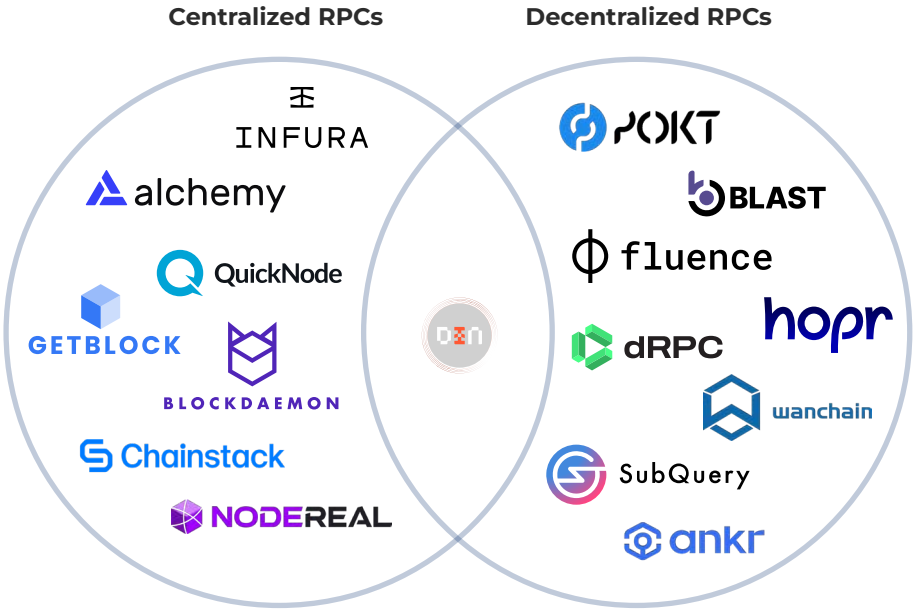

Decentralization exists on a spectrum, meaning that not all decentralized RPCs (dRPCs) are equally decentralized.

The Value Drivers of the RPC Industry

The rise of decentralized RPC services has disrupted the market. Consequently, SubQuery can capitalize on numerous inherent value drivers within this segment. Firstly, Decentralization plays a pivotal role in ensuring fault tolerance and network resilience by distributing the network across multiple nodes, reducing reliance on a single point of failure. This enhances accessibility to data and services even in the face of compromised or offline network segments. Additionally, the inherent security features of decentralized networks, coupled with improved privacy through anonymized requests, offer heightened protection against malicious attacks and data breaches.

Furthermore, scalability and performance are crucial attributes of decentralized RPCs. These services can dynamically scale with the market demands of the network, effectively managing high query volumes and potentially ensuring faster access to blockchain data. This capability meets the demands of a growing user base. As highlighted earlier in this research, latency is exceptionally important and can be minimized by having nodes closer to the end user. Excessive latency can result in unacceptable wait times for users, leading to a potential loss of user engagement.

Cost Efficiency is yet another compelling value driver. It is achieved through competitive pricing facilitated by a distributed network of providers through market mechanisms and token incentives. One of the key features is the Censorship Resistance inherent in decentralized RPC services ensures uninterrupted access to blockchain data, vital for maintaining information freedom in regions where censorship may be prevalent as constant access to blockchain data is guaranteed. Last but not least, Interoperability and adherence

to open standards foster a more inclusive blockchain ecosystem, enabling seamless interaction and data exchange across different blockchain platforms, ultimately unlocking new possibilities for developers to build cross-chain applications unlocking collaboration and innovation.

Overall, decentralized RPC services offer numerous benefits driven by key factors such as decentralization, security, scalability, cost efficiency, censorship resistance, and interoperability. These services ensure fault tolerance, enhance security and privacy, scale dynamically, offer competitive pricing, resist censorship, and promote collaboration across blockchain platforms. As they continue to evolve, decentralized RPC services can play a crucial role in shaping the future of decentralized technologies, empowering consumers and developers with greater control, security, and accessibility not relying on centralized entities (e.g., cloud service providers)

Threat of New Entrants

RPC protocols serve as a vital bridge between consumers and blockchains, pivotal for executing transactions or interacting with dApps. Thus, such an infrastructure project must secure a sufficiently robust base for its operations to become a proficient RPC provider. Particularly, a reliable cloud service provider is necessary, which can be either expensive and exposed to single points of failure, in case of selecting a centralized solution, or potentially not scalable enough affecting the consumer's project functionality speaking of decentralized cloud service providers. Furthermore, established centralized providers like Infura, utilized by industry giants like Metamask, maintain hegemony in the market as they have built trust in the space aiding them to lock in consumers who are unlikely to often change providers. This dominance creates formidable barriers to entry that discourage aspiring competitors. Nevertheless, the decentralized RPC solutions have partially disrupted the high entry barriers even though some of these providers grapple with performance issues, curtailing their ability to attract significant market share. Overall, the decentralized RPC service segment opened the market to newcomers, although one still needs sturdy infrastructure to enter and build a reputation among the strong market leaders. Therefore, an assumption can be made that the threat of the new entrants is marginal.

Rivalry Among Existing Competitors

Competition in the RPC provider market is fierce with both centralized and decentralized entities striving for market share. Many rely on centralized providers despite facing criticism for their susceptibility to single-point attacks and failures, inherent in centralization, potentially compromising the reliability and security of blockchain networks. Conversely, some decentralized providers, for the time being, struggle to match the performance of their centralized counterparts, limiting their competitiveness; However, that is not necessarily the case for all decentralized providers. Currently, the market has an oligopolistic structure, where a few centralized providers dominate, narrowing the room for competition (e.g., price, marketing), thus limiting innovation. Be that as it may, decentralized counterparts present a serious threat to centralized providers as the latter have only a few benefits. Still, the market leaders can leverage their track record and other lock-in mechanisms; hence, an assumption can be made that the competition is severe with the decentralized providers to stay.

Bargaining Power of Suppliers

The RPC providers require the necessary infrustrustructure to operate. This can be provided by either well-known cloud service providers or decentralized providers may decide to run the network via independent node operators. Firstly, many RPCs mandate

the centralized cloud operators since they are more resource-efficent at present. However, the cloud service market is dominated by a few companies (i.e., Google, Microsoft, and Amazon) considering the necessary large investments (e.g., data centers) to provide the service. Thus, providers have leeway to adjust the price or quality of the service to a certain extent without risking customer loss. Further, the switching costs for projects are significant, even though the costs can vary based on the extent of reliance on the cloud provider, the specific features they use and the overall architecture of their application. Last but not least, some decentralized RPC providers utilize independent node providers who are usually incentivized to participate in securing network's operations profoundly reducing the switching costs for RPCs compared to using centralized solution. Overall, the suppliers hold rather significant bargain power given, their presence in the RPC market segment, the cloud market's olygopolistic structure and the rather high switching costs. However, these power might be significantly offset by the choice of running the request through a vast network of node operators.

Bargaining Power of Users and Consumers

In the RPC market, users of dApps and other blockchain networks, wield very limited influence given only a handful of wallets permits RPC change, fostering a lock-in mechanism built on unconditional trust and limited options. Secondly, speaking of the consumers (i.e., dApps) the selection of providers does not solely depend on price but rather on their performance, track record, partners, reliability and other factors. However, price may play a pivotal point in choosing a provider for developers. In conclusion, the choice of a wallet provider severely constrains users' bargaining power. The dApps and developers are not strictly restrained by the choice of a wallet and enjoy the product differentiation. Hence, it can be inferred that dApps and developers possess some bargaining power in the RPC market.

The Threat of Substitute Products

RPC providers face the threat of substitution by developers opting to run their own RPC nodes. However, high costs and technical intricacies involved in this approach restrict its viability as a widely adopted alternative. Furthermore, it is worth mentioning, that the decentralized RPC solutions offering an alternative to centralized providers are rather a direct competitor than a substitute. Therefore, an assumption can be made that the threat level is marginal.

In summary, SubQuery operates within a challenging market environment characterized by high barriers to entry, fierce competition from especially established centralized providers, and high bargaining power of suppliers in case of not selecting the less common decentralized option. Navigating these complexities strategically will be crucial for SubQuery to establish its position and succeed in this dynamic market.

by blockchain and decentralized RPCs create a trustless environment, where transactions and data exchanges are verifiable and secure, fostering greater enterprise collaboration and innovation.

The Market Opportunity of Decentralized RPCs

The journey of decentralized RPCs into enterprise adoption is marked by the promise of heightened security and privacy features. Traditional enterprise networks, often present centralized points of failure that can be exploited by malicious entities. By leveraging the decentralized nature of blockchain, enterprises can avoid the pitfalls of centralized systems, such as single points of failure and vulnerability to cyber-attacks. Industries like finance, healthcare, and supply chain are especially poised to benefit from adopting decentralized RPCs, utilizing them to safeguard financial records, patient information, and logistics data, respectively. The immutable record-keeping, transparent and censorship resistant operations enabled by blockchain and decentralized RPCs create a trustless environment, where transactions and data exchanges are verifiable and secure, fostering greater enterprise collaboration and innovation.

The integration of decentralized RPCs with the Internet of Things (IoT) and Decentralized Physical Networks (DePin) heralds a new frontier for smart technologies. DePin, with its emphasis on creating decentralized networks for physical assets while having the ability to track and verify, synergizes with IoT's needs for secure, scalable, and resilient communication protocols. This combination is set to shift industries requiring real-time monitoring and verification of physical assets, from logistics and manufacturing, agriculture and energy management to computing. Through decentralized RPCs, IoT devices can interact within a decentralized, tamper-resistant framework, ensuring data integrity and continuity even in adverse network conditions, thereby unlocking new levels of efficiency and reliability for smart applications.

Decentralized RPCs are the backbone of the rapidly growing DeFi sector, which promises to extend financial services to the unbanked and underbanked populations around the globe. By facilitating direct interactions between users and financial services on the blockchain, without the need for traditional banking intermediaries, DeFi platforms are making a wide range of financial services accessible to anyone with internet access. Decentralized RPCs ensure these transactions are not only secure but also transparent and immutable, embodying the principles of financial democracy and inclusion. This paradigm shift in finance, powered by blockchain and decentralized RPCs, is dismantling barriers to economic participation, offering everyone a stake in the global financial system.

The vision for a decentralized web—Web3—is predicated on the robustness of its infrastructure, where decentralized RPCs play a pivotal role. By facilitating distributed data storage and processing across the blockchain, these protocols enhance the internet's resilience against attacks, censorship, and central points of failure. The adoption of decentralized RPCs is instrumental in building a more secure, private, and user-centric internet, aligning with the foundational goals of Web3.

Chapter Summary

The RPC industry faces intense competition from both centralized and decentralized entities. Centralized players hold a competitive edge due to lock-in strategies that create closed ecosystems, preventing users from switching providers and leading to market inefficiencies. An open market could enhance competition, but there is a risk of monopolistic environments if few RPC providers dominate.

High barriers to entry make it difficult for new players to capture market share. Establishing a scalable decentralized RPC service requires a substantial network of nodes, often relying on centralized cloud services for efficiency. Sustainable financial incentives are crucial to attract node operators without causing token inflation, which can damage credibility.

Price competition is fierce, potentially driving prices down to the base cost of processing requests. Technological innovation, such as SubQuery's SDK 4.0, which improves indexing performance through multi-threading and optimized database operations, offers a more sustainable competitive advantage. SubQuery achieves faster indexing speeds, significantly outperforming competitors like The Graph.

Decentralized RPCs offer several value drivers, including fault tolerance, enhanced security, scalability, cost efficiency, censorship resistance, and interoperability. These attributes make decentralized RPCs attractive by distributing the network across multiple nodes, reducing single points of failure, and offering dynamic scaling with market demands. Additionally, decentralized networks ensure uninterrupted access to blockchain data, which is crucial for maintaining information freedom and privacy.

The market is characterized by significant entry barriers, fierce competition, and high bargaining power of suppliers due to the dominance of established cloud service providers. Consumers, primarily dApps and developers, have some bargaining power based on performance, reliability, and price. However, users of dApps face limited influence due to the lock-in mechanisms of wallet providers.

Decentralized RPCs present a viable alternative to centralized services, offering greater security and resilience. Industries like finance, healthcare, and supply chain can benefit from decentralized RPCs for safeguarding sensitive data. Integration with IoT and Decentralized Physical Networks (DePin) opens new opportunities for real-time monitoring and verification of physical assets.

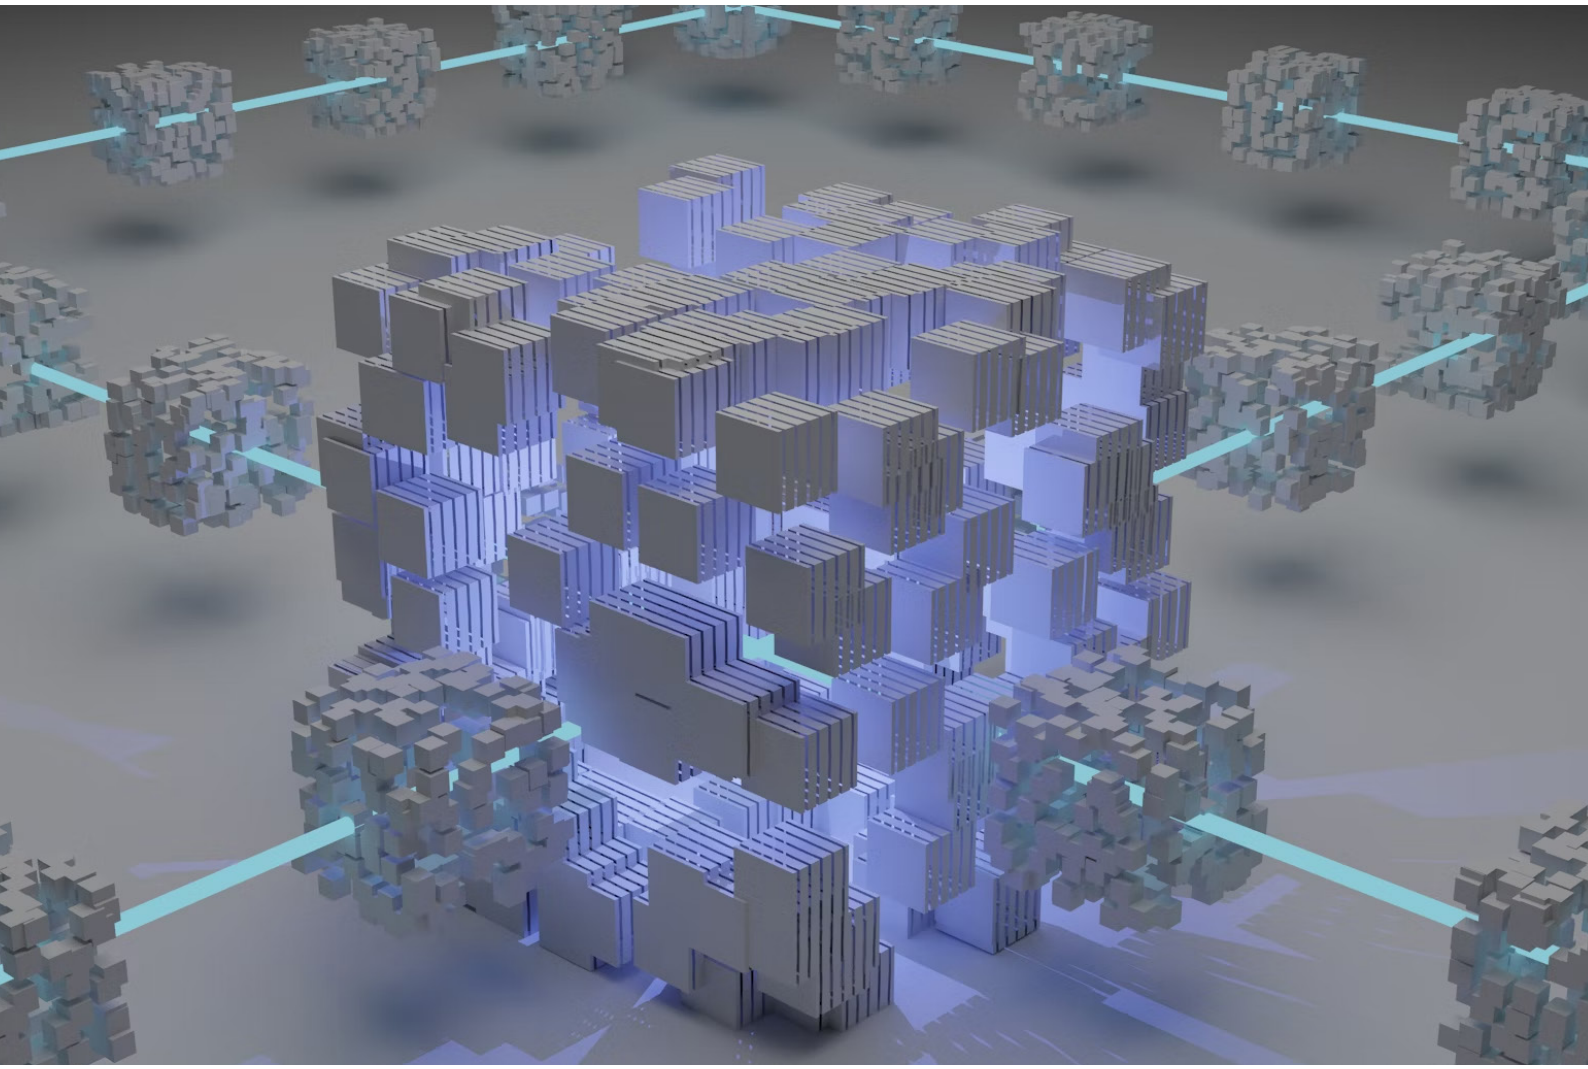

Photo by Shubham Dhage on Unsplash

Risks in Decentralizing RPC Services

Exploring the challenges and market penetration risks of decentralized RPCs, including reliance on centralized cloud services and the complexities of gaining user trust and technological scalability.

Governance Risk of SubQuery

Decentralized RPCs present an innovative leap towards a more open and decentralized web; however, they encounter notable challenges that impede their full decentralization. A primary concern lies in the reliance of many decentralized RPC nodes on underlying centralized cloud computing services. This reliance subtly reintroduces centralization into an ecosystem that is fundamentally aimed at distributing control and reducing single points of failure. Furthermore, the operational frameworks within many decentralized RPC projects, such as SubQuery, reveal additional layers of centralization. As it currently stands, SubQuery operates under the oversight of a node committee, which is entirely managed by the SubQuery team. However, after considering the findings of this research paper, the SubQuery team has introduced a roadmap to implement Community Governance within the SubQuery ecosystem. Therefore, aiming to provide an effective solution to the risk. In it's current form, the committee plays a pivotal role in the ecosystem, tasked with identifying and addressing nodes that exhibit malicious behavior. The repercussions for such behavior can range from partial to complete slashing of the node's stake, a punitive measure that, while necessary for maintaining network integrity, also centralizes judgment and control within a limited group.

Penetration Risk of Decentralized RPCs

The market penetration risk for decentralized RPCs lies in the challenge of capturing market share from centralized providers entrenched with lock-in mechanisms. In core infrastructure projects, reputation plays a critical role, establishing trust and reliability in the eyes of potential users. Centralized RPC providers, having established a foothold in the market, benefit from high user adoption barriers that deter switching. These barriers are not only technological but also psychological, as users may hesitate to transition to decentralized solutions due to unfamiliarity or perceived risk.

Technologically, decentralized RPCs face challenges related to scalability, latency, and compatibility with existing systems. For instance, ensuring low latency and high availability in a decentralized setup requires sophisticated algorithms and extensive infrastructure, which are often resource-intensive to develop and maintain. Additionally, integrating decentralized solutions with existing platforms that are designed around centralized architectures can be complex and costly.

Psychologically, users may be reluctant to switch due to a lack of familiarity with decentralized systems and concerns over potential risks, such as security vulnerabilities or inconsistent performance. This hesitation is compounded by the fact that centralized providers often offer well-established, reliable services with robust customer support, making the perceived risk of switching even higher.

Overcoming these barriers requires decentralized RPCs to not only match but exceed the performance, reliability, and ease of use offered by centralized services. This involves not only technological advancements but also strategic efforts to build a strong reputation and trust within the Web3 community.

Chapter Summary

Decentralized RPCs aim to create an open web but face challenges. Many rely on centralized cloud services, reintroducing centralization. For example, SubQuery's node committee, managed by the SubQuery team, adds a layer of centralization. SubQuery plans to implement Community Governance to address this, but currently, control is centralized within a small group responsible for monitoring and penalizing malicious nodes.

Decentralized RPCs struggle to capture market share from established centralized providers with strong lock-in mechanisms. Centralized providers benefit from high user adoption barriers due to their established reputation and reliability. These barriers are technological and psychological, as users may hesitate to switch to decentralized solutions due to unfamiliarity and perceived risks.

Technologically, decentralized RPCs face scalability, latency, and compatibility challenges. Ensuring low latency and high availability in a decentralized setup requires advanced algorithms and extensive infrastructure, which are resource-intensive. Integrating decentralized solutions with existing platforms built around centralized architectures can be complex and costly.

Psychologically, users may be reluctant to switch due to unfamiliarity with decentralized systems and concerns about security vulnerabilities or inconsistent performance. Centralized providers offer reliable services and robust customer support, making the perceived risk of switching higher.

To overcome these barriers, decentralized RPCs must match or exceed the performance, reliability, and ease of use of centralized services. This involves technological advancements and strategic efforts to build a strong reputation and trust within the Web3 community. Addressing these challenges is crucial for the broader adoption and success of decentralized RPCs.

Appendix

Figure 1 SQT Token Distribution

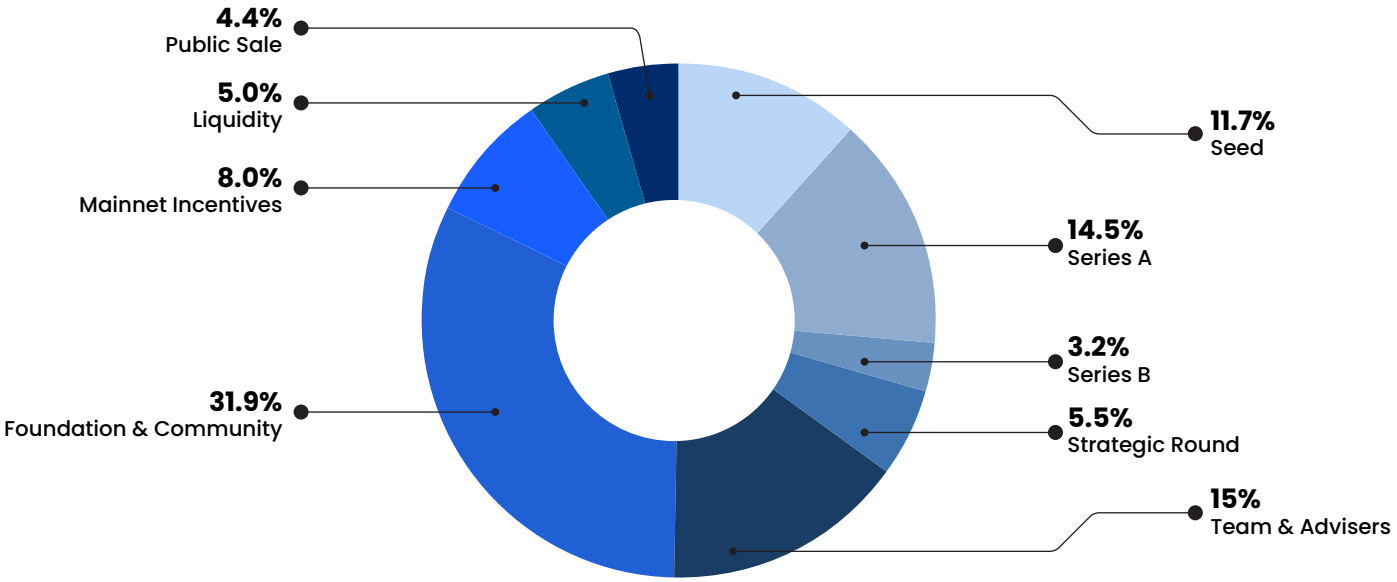

Figure 2 SQT Token Vesting

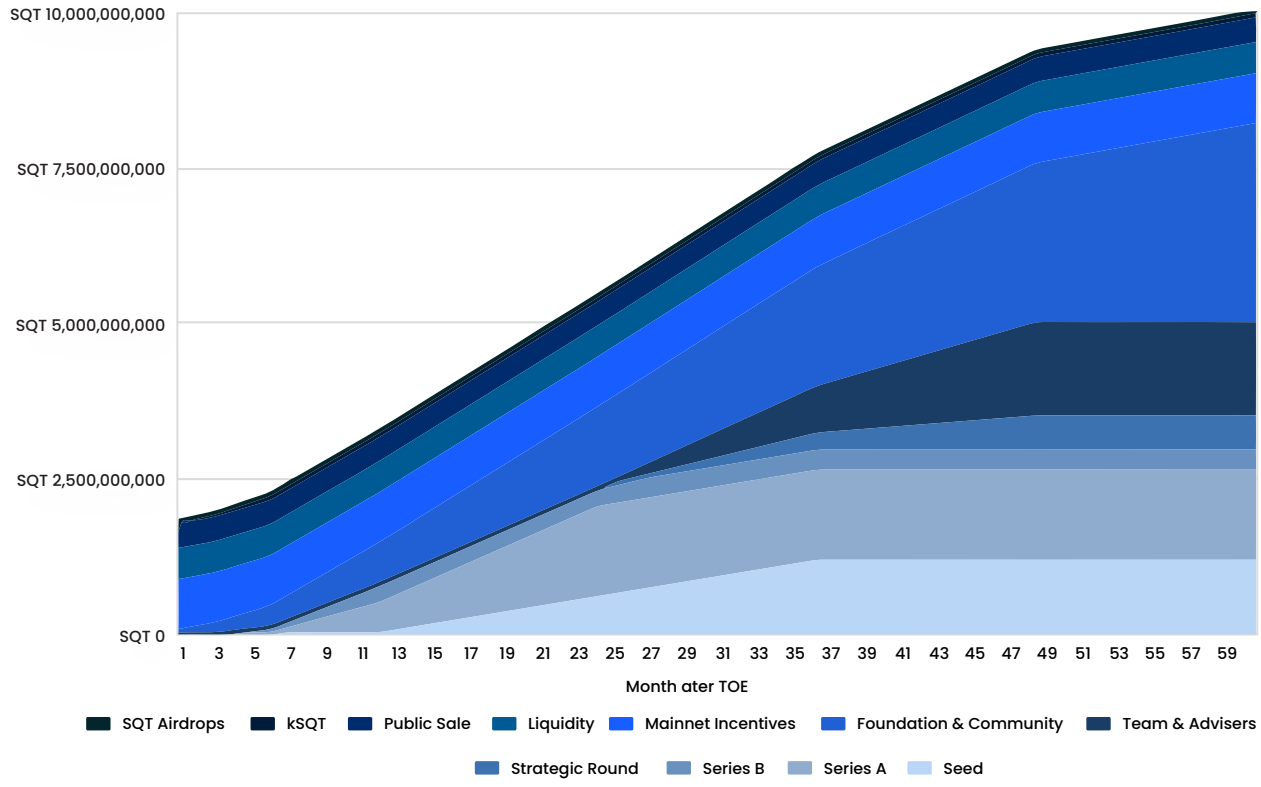

Figure 3 C-level Overview and Team Experiences of SubQuery

Sam Zou

CEO and Founder

Ian He

CTO and Co-Founder

James Bayly

COO and Co-Founder

Scott Twiname

Principal Dev of SDK

Marta Adamczyk

Head of Business Development

Brittany Seales

Head of Marketing

Previous Experiences

Alipay onfinality sylo

termius HAYS

May 2024

Designed by House of Chimera

HouseOfChimera.com