Coinsider

CRYPTOEQ

Ethereum Upgrade Guide 2022

CryptoEQ CORE+ Series.

by Mark Cole and CryptoEQ

The essential guide to the Layer 2 Ethereum Upgrade.

Ethereum Upgrade Guide 2022

I. Scalability

First, there was Bitcoin. Then, Ethereum. Now, the much-anticipated Ethereum upgrade. It might be helpful to think of Ethereum in two parts: the Ethereum consensus layer and the Ethereum execution layer—at least for the time being. But what does that mean?

Ethereum is a Layer-1 (L1) blockchain currently amid a 5+ year upgrade to satisfy future global demand while also improving security and decentralization. Previously, the Ethereum roadmap was planned in sequential stages that lent names like “Eth1”, “Eth1.x” and “Eth2.” However, that plan has been altered, making these terms—Eth1 and Eth2—irrelevant.

The old naming scheme suggested two issues—namely that “Eth1 comes first, and Eth2 only after” and that “Eth1 will cease to exist once Eth2 exists.” In reality, post-Merge, the chain(s) and their data will be seamlessly joined together. In regards to Ethereum’s next major upgrade, The Merge (~Q2 2022), the consensus layer (previously Eth2) will be merged with the execution layer (previously eth1), creating just one Ethereum again. Instead of referring to the chains as Eth1 or Eth2, the community has shifted to calling them the “consensus” and “execution” chains, respectively. The execution chain encompasses all the state (AKA data) associated with the user layer (dApps, account balances, tokens, etc.).

Consensus encompasses the proof-of-stake consensus mechanisms. This “base layer” is entirely focused on consensus and data availability. In a post-Merge environment, both of these layers coexist together.

Ethereum has been massively successful as an open network and computing platform where software and application developers can collaborate and innovate quickly, easily, and without requesting permission. The Ethereum network has become especially popular for games, digital collectibles, and person-to-person finance. But all of this use of the Ethereum network, however, has led to congestion, high user fees and surging electricity consumption.

1. What is "The Merge"?

The Merge is the term used for when Ethereum switches from proof-of-work (PoW) to a proof-of-stake (PoS) blockchain. Slated for Q2 2022, the merge will bring many benefits that were not previously possible with PoW.

A PoS structure purportedly removes the energy consumption often cited in the mainstream media. While PoW is not inherently a bad thing, it's inarguable that the world is highly critical of energy consumption, and now, with the transition to PoS, Ethereum will have eliminated this one enormous criticism. Estimates from Ethereum core developers hypothesize that Ethereum's energy use will drop by up to ~99%. With less need for physical mining hardware and infrastructure, Ethereum can become a more energy-efficient, geographically-distributed, and nimble blockchain.

Additionally, PoS is a predecessor for sharding, another critical Ethereum protocol change that will separate the chain into many concurrent threads. Finally, the PoS upgrade will reduce Ethereum's inflation rate from ~3.5% to nearly zero.

These changes will provide increased scalability for the Ethereum chain, which has regularly experienced periods of congestion and high network fees since 2020. The Merge, is just the first step in an enormous transformation for Ethereum. Below is the latest update to the roadmap (as of Q4 2021).

A traditional monolithic, do-it-all blockchain faces unavoidable limitations, by design, due to the inefficient nature of decentralized consensus. These limitations lead to inflating costs for its users as the chain becomes more widely used.

The costs occur because blocks and block space on the execution layer of a chain are scarce. There are only so many blocks that can be verified and added to the chain each second. Once demand outpaces this finite resource, the only recourse users have left to ensure their transaction gets into a block (and executed) is to pay more than the next person.

There are two ways a monolithic blockchain (a blockchain that provides its security, executes its transactions, and maintains its data availability) can scale: increase capacity at the base layer (on-chain) or move transactions to a second layer (off-chain).

On-chain scaling techniques are upgrades made to a blockchain's base layer to improve scalability. Ethereum's long-term, on-chain scaling solution, sharding, splits the base layer into 64 chains with shared security ensured by the Beacon Chain. Off-chain scaling refers to scaling solutions that use external execution layers (Layer 2s) rather than the base layer. Layer 2s or "L2s" are secondary layers that sit on top of the base layer, and, in the case of rollups, inherit the mainchain's security while providing more transactional capacity for the blockchain overall.

Ethereum is pursuing both off-chain and on-chain scaling strategies in the response to these challenges. It's a massive upgrade to the entire Ethereum ecosystem to accommodate continued growth and an increasing workload, consume less resources in its verification process, and secure itself from attacks. In essence, the Ethereum upgrade will make the network more scalable, sustainable, and secure.

Any highly successful human enterprise will eventually have to address how to keep up with demand. This is a good problem to have, but not an easy one to solve. Scaling can often challenge the core values that made the enterprise successful.

Example:
Imagine a craft coffee bean roaster: Quality Coffee Company. Quality painstakingly cares for every batch, from sourcing to roasting to packaging and shipping. Highly-trained and deeply-motivated professional coffee connoisseurs carefully control each step. Moving from limited runs of small batches to an automated operation shipping ever-larger volumes globally without destroying quality will be very challenging, perhaps even impossible. Quality Coffee Company will likely need to change its name to Big Coffee Company. Scaling up and maintaining original values are often at odds.

The same can be said with Ethereum and the value of its integrity in not allowing any small group of people the ability to change the Ethereum blockchain. The core commitment of Ethereum is decentralization and it is fundamental both technically and philosophically, but at the same time, Ethereum needs to do more work.

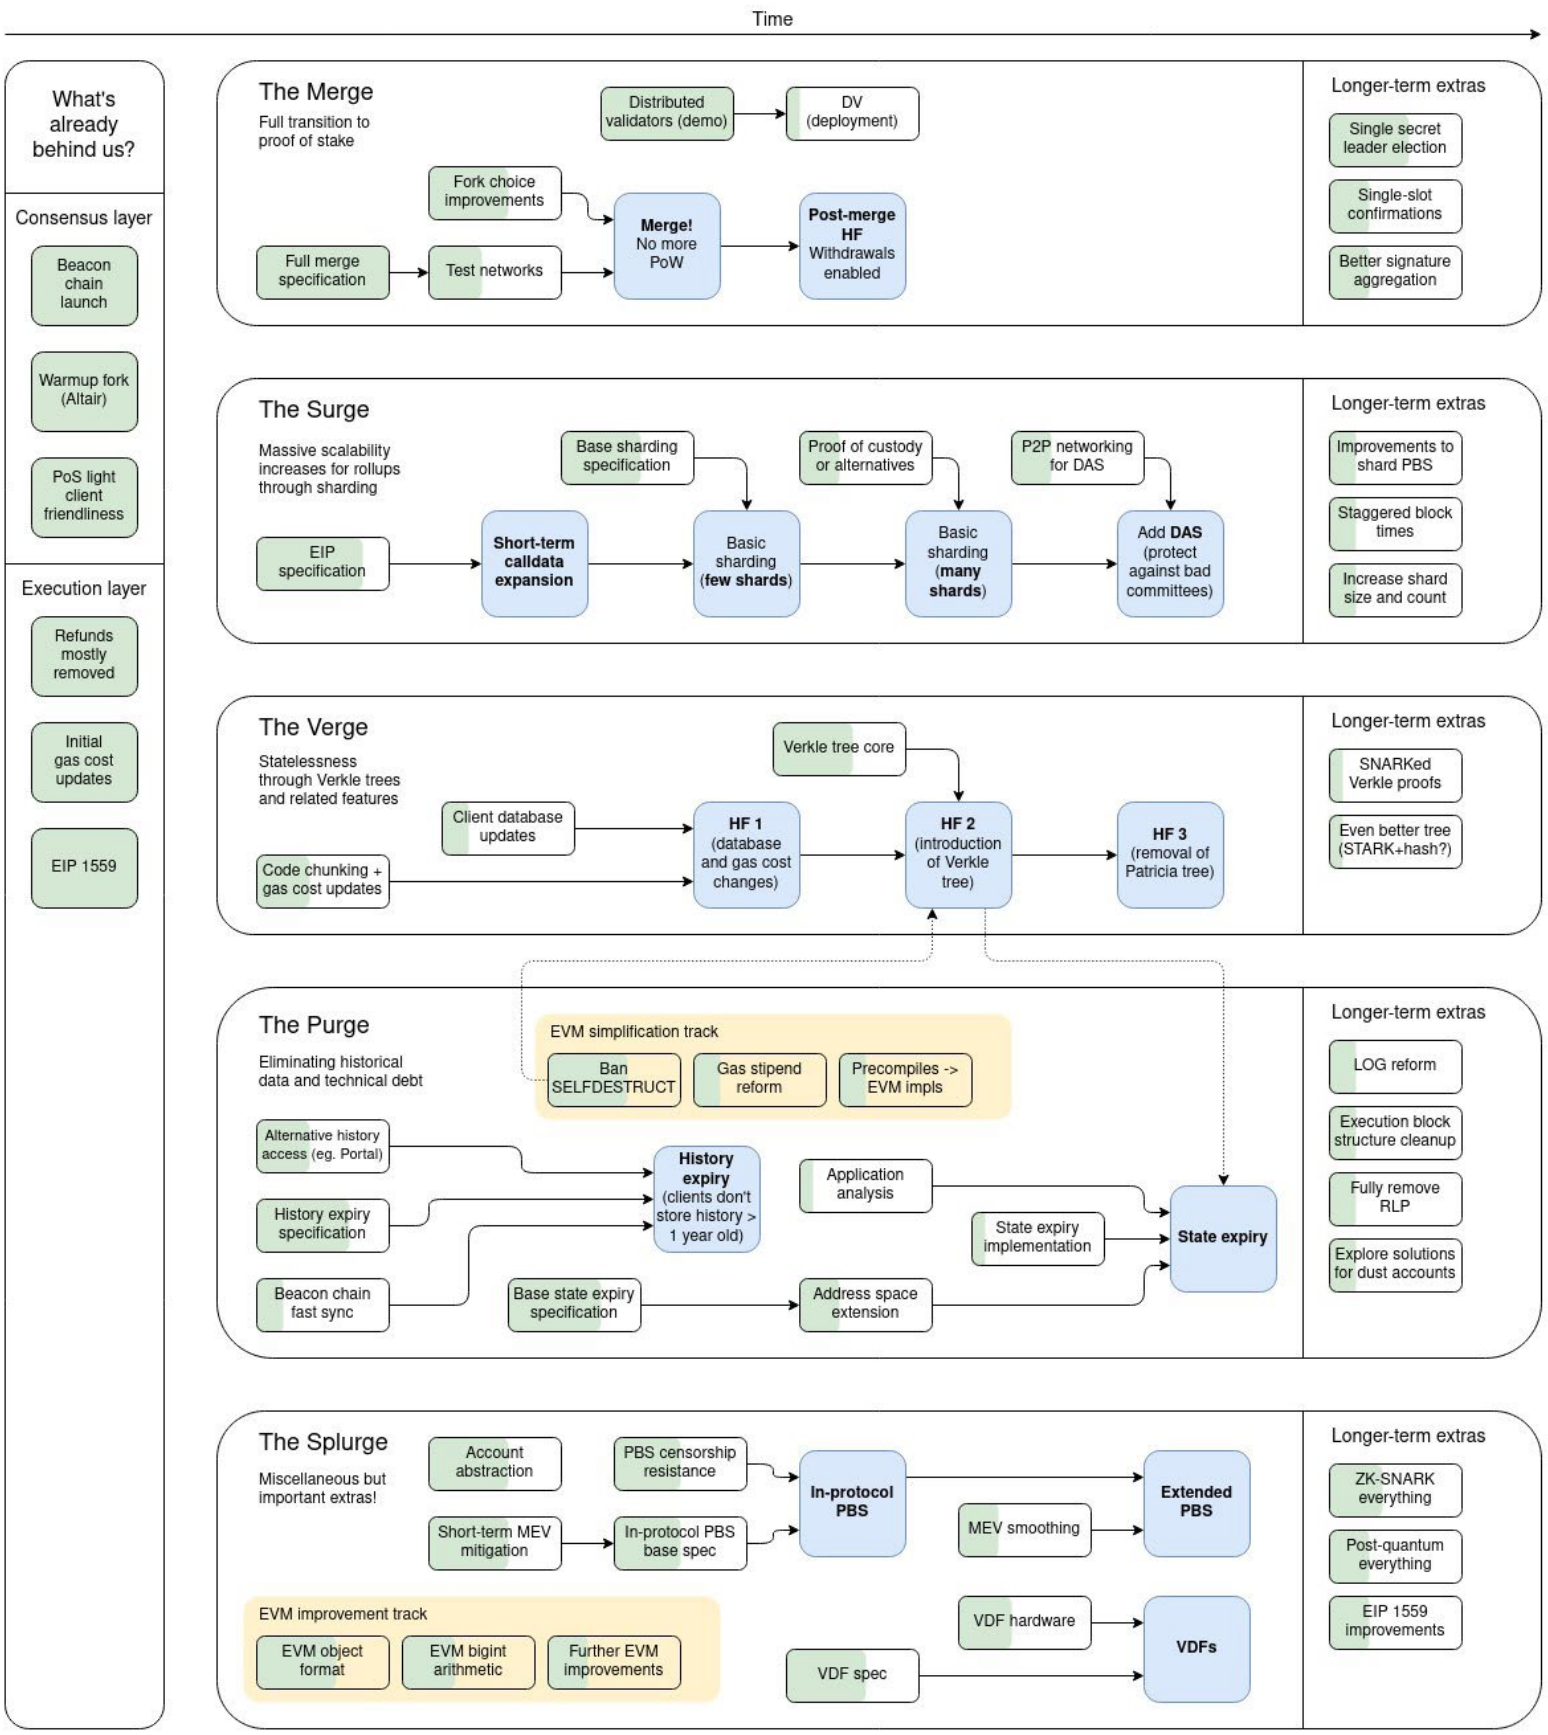

Ethereum Upgrade Roadmap **SOURCE**: Vitalik Buterin

The Issue

Transaction demand and smart contracts use has skyrocketed over the last two years. In 2021 alone, the number of DeFi users increased from ~150k to ~2M, while at the same time, gas fees grew 16 times faster. As of Q1 2022, the Ethereum mainnet routinely facilitates the transfer of tens of billions of dollars daily, with over $150B currently deposited in DeFi smart contracts.

As of Q1 2022, Ethereum can only process ~25 transactions per second (TPS) due to its design optimized for decentralization and security. Without processing more transactions, congestion on the network forces users to pay more to execute their transactions. This has led to extremely high (>$40) transaction fees for users due to the high demand for limited block space on the blockchain. Essentially, block space is the commodity that users, creators, and builders consume, making it the pulse of all cryptocurrency networks.

High network fees are a product of how blockchains process transactions. There is a cost associated with a global, decentralized, censorship-resistant financial settlement layer. All nodes across the decentralized network must agree for a transaction to be executed. All nodes on the network keep a full copy of the transactions to validate the transactions on the network.

Ethereum's ability to process transactions is (partially) constrained by computing power, bandwidth, and storage on the network. The scalability trilemma is a well-known issue among all blockchains.

Ethereum's ability to process transactions is (partially) constrained by computing power, bandwidth, and storage on the network. The scalability trilemma is a well-known issue among all blockchains.

A blockchain can achieve two of these traits, but at the expense of the third. Many alternative Layer-1 (L1) chains have sacrificed decentralization for scalability and security. However, it's important to remember why decentralization is necessary. It provides the chain anti-fragility, robustness, reliability, and censorship resistance.

The goal is to increase the number of transactions while retaining sufficient decentralization. What are the decentralization sacrifices (tradeoffs) other smart-contract L1s have made? Other chains typically make two sacrifices. They either increase the requirements to run a node to have more high-powered machines, which reduces the number of people who may participate in network consensus by pricing them out. Obviously, a network that can only be verified if you have X dollars in computing budget is not an ideal, permissionless system. Using a crude analogy, would be like making it harder for the average person to vote in an election.

The other tradeoff commonly conceded is for the network to use fewer nodes to achieve consensus faster and quicker. This makes the chain more vulnerable and centralized. It's easy to corrupt/destroy ten nodes all in one location rather than 10,000 all over the globe.

Although often discussed, blockchain scalability does not just pertain to TPS. Many L1s, like Binance Smart Chain (BSC), currently boast high TPS numbers but suffer from "chain bloat" and ever-increasing hardware requirements just to keep the chain running. L1s must process more transactions without creating more problems down the road. A node in a technically sustainable blockchain has to do three things:

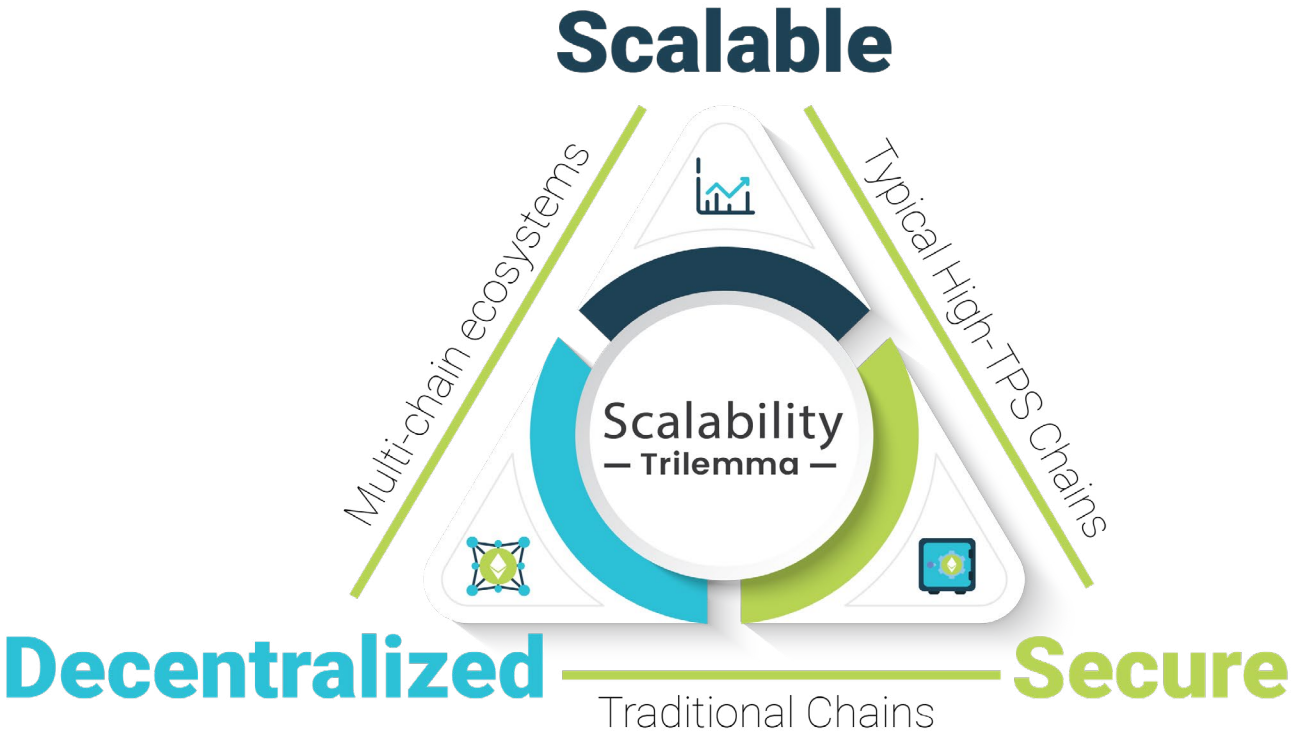

Node Requirements

- » Keep up with the tip of the chain (most recent block) while syncing with other nodes
- » Be able to sync from genesis in a reasonable time (days, instead of weeks)
- » Avoid state bloat

Requirement 1 is a physical limitation based on computing power (RAM, CPU, etc.) and bandwidth. These are bottlenecks for every node, which means there are upper, finite limits to how far you can push the network.

One way for Ethereum to increase its workload could be to increase the size of the computers participating in the Ethereum network (participating computers are called “nodes”). But larger, more expensive, and fewer computers in the network is a form of centralization. Having fewer bigger players involved in maintaining Ethereum is not Ethereum’s goal.

Fewer computers in the network also create security issues. A hack would be easier on fewer computers—or a single central computer—than on a vast number of computers. Just as with Bitcoin, more computers participating in the Ethereum network enhance the security and permanence of the data on the Ethereum blockchain.

2. Sharding

After switching to PoS, sharding is the next significant hard fork upgrade on Ethereum's roadmap. Like the merge, the sharding plan has evolved and may continue to change between now and implementation.

There are two main approaches to blockchain scaling:

1. Vertically, *i.e.*, making the network's nodes more powerful
2. Horizontally, *i.e.*, adding more nodes (with no performance improvement)

Because Ethereum prioritizes decentralization and security, it must be designed so that everyone has the option and ability to run their own node. This means the first approach, vertical scaling, which typically leads to more expensive and onerous hardware requirements, is not a viable option. Ethereum must keep the requirements to run a node low so that it is open to nearly everyone.

Sharding fits into a horizontal scaling approach. Sharding is the partitioning of a database (or blockchain) into smaller subsections. Rather than building layers atop one another (*e.g.*, L2s or Bitcoin's Lightning Network), sharding scales out horizontally without a hierarchy or layered structure. Doing so does not create more burden for the average user.

Shards will be divided among nodes so that every individual node is doing less work. But collectively, all of the necessary work is getting done—and done more quickly. More than one node will process each data unit, but no single node has to process all of the data anymore.

In Ethereum's vision of a sharded chain, a (pseudo) randomly-chosen committee of validators is randomly selected and assigned to specific shards. This means they are only responsible for processing and validating transactions in those specific shards, not the entirety of the network. The randomness of the validator selection process ensures it is (nearly) impossible for a nefarious actor to attack the network successfully.

Initially, before the breakthrough in rollups, Ethereum planned to do sharded computation. However, now with rollups providing the much-needed network scalability, sharding will focus on data available to provide throughput for the rollups. This is because the bottleneck for rollup scalability is data availability capacity rather than execution capacity. This will give L2s more space to store the chain's data and offer additional data capacity for rollups.

In a sense, shards will serve as data storage "buckets" for new network data storage demand from rollups. This enables tremendous scalability gains on the rollup execution layer. Just as significant, shards will also help avoid putting overly-onerous demand on full nodes, enabling the network to maintain decentralization.

Sharding will be released in a multi-step process to provide immediate data availability for rollups before releasing the more complex but ultimate vision. A small subset of data shards (4) will be released initially to keep complexity low *i.e.* a slow, controlled roll out.

Earlier, we outlined one reason why Ethereum transaction fees were so high was due to all nodes in the network having to process all transactions and reach consensus. Sharding is the answer to the question, “What if each node did not have to process every operation at the same time?” What if, instead, the network was divided into subsections (shards), that operated semi-independently until finally reaching consensus through a central hub (Beacon Chain)?

Shard 1 could process one batch of transactions, while Shard B processes another batch. This would effectively double the transaction throughput of a blockchain, since our limit is now what can be processed by two nodes at the same time. If we can split a blockchain into many different sections, then we can increase the throughput of a blockchain by many multiples.

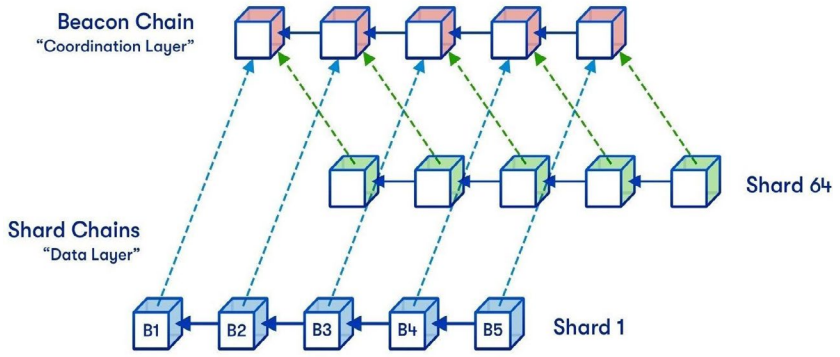

Sharded Ethereum **SOURCE**: Hsiao-Wei Wang

3. Layer 2 Solutions

Layer 2 is a broad, catch-all term used to describe scaling solutions built on top of an existing L1 blockchain. The primary advantage of using an L2 solution is the main chain remains untouched and unaffected by what is built atop it. Any issues that happen “up the stack” (*e.g.*, on another layer) will not compromise the base layer. L2s essentially function as smart contracts on the Ethereum mainnet that interact with off-chain software. Because of this, L1 serves as the security and consensus layer that cryptographically secures and stores data transactions on the immutable blockchain ledger.

L2s can further extend Ethereum’s utility, granting users increased scalability off of the blockchain that can still refer back to the main chain, as required
.

100,000

Transactions per second: Refers to how many transactions a payments system or cryptocurrency can successfully process per second. Currently, Ethereum Mainnet can process about 25 TPS. After zk-rollups and full sharding implementation this should grow to over 100,000 transactions per second.

Ethereum, as we know it today, won't scale. Meaning, the Ethereum L1 is designed to remain a highly decentralized, global settlement layer above all else. However, Ethereum's web of L2s will be responsible for scaling Ethereum and serving as its execution layer. These layers will absorb much of the existing value on Ethereum mainnet plus future inflows as Ethereum adoption grows. It's important to understand that Ethereum's web of L2s is a marketplace of independent projects competing to help scale Ethereum.

The original Ethereum upgrade roadmap was the response to these scaling challenges. It's a massive upgrade to the entire ecosystem accommodating continued growth and an increasing workload, to consume less resources in its verification process, and to be more secure from attacks. That is, the upgrade will make the network more scalable, sustainable, and secure. That's why Ethereum is upgrading.

The same security does not always apply to layers built "on top" of Ethereum as it does to on-chain operations, but these layers can still be sufficiently secure to be useful—especially if the user is comfortable with the tradeoff for low-value transactions.

Ethereum L2s allows builders to tailor their tooling to their needs, meaning they can decide for themselves where their product sits on the scalability trilemma.

Tradeoffs between speed, finality, and transaction cost can be developed, just like in competing alternative L1s. For the most valuable transactions, users can choose the main chain where security and censorship-resistance is highest and for low value transactions, a gaming sidechain may suffice. L2s allow the user to maintain control without compromising the underlying blockchain, preserving decentralization and finality.

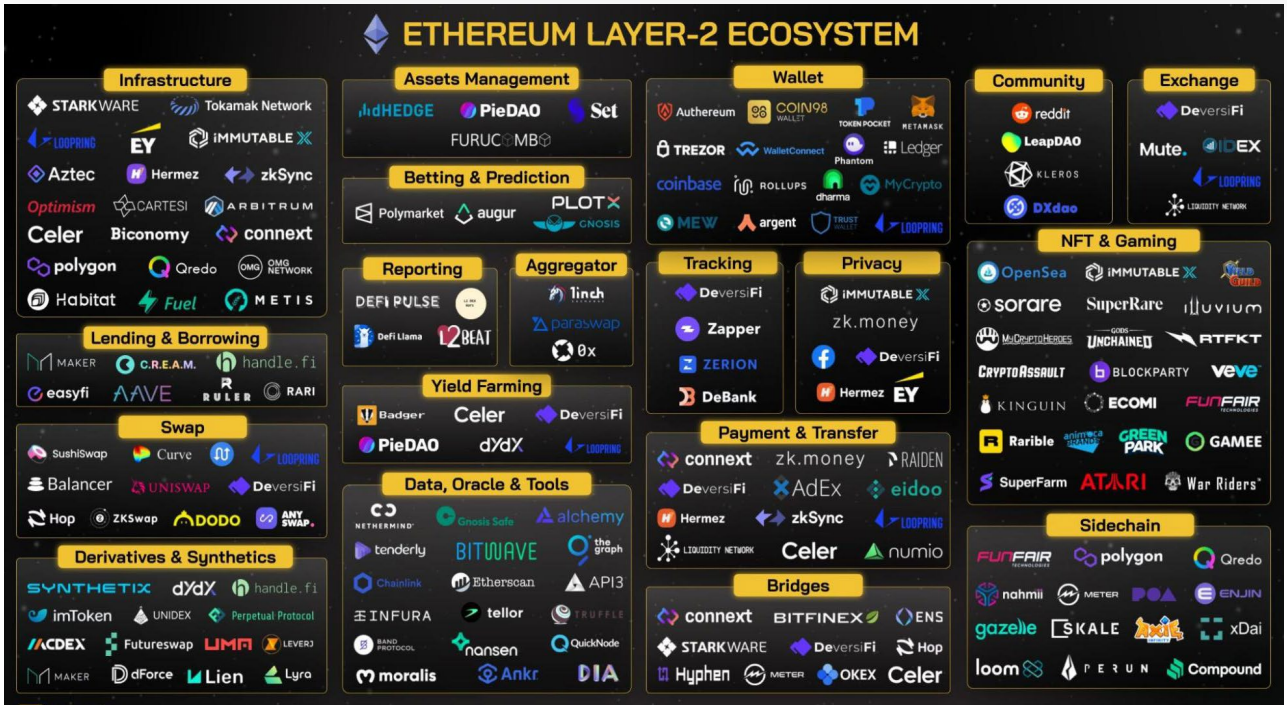

Ethereum Layer-2 Ecosystem **Source**: Coin98

4. Sidechains

In the context of Ethereum, sidechains are separate, Ethereum-compatible blockchains. Sidechains can be independent EVM-compatible blockchains, but more likely, they are application-specific blockchains catering to Ethereum users and use cases like Polygon or Ronin.

EVM stands for the Ethereum Virtual Machine and is the global network of computers that keeps Ethereum running. The EVM actually handles processing of every transaction on Ethereum. It is a Turing-complete virtual machine that is limited by the amount of gas provided by users.

Sidechains design themselves to be EVM-compatible to essentially copy and paste their code to easily interoperate with Ethereum and its infrastructure—wallets, block explorers, etc. Projects like Binance Smart Chain, Avalanche, Tron, Celo, and Fantom are all examples of competing L1 chains that have (more or less) launched an EVM-compatible chain, tweaked several parameters to increase TPS, and attached their own token to the project. Some proprietary architectural builds, like Avalanche's three-chain or Fantom's DAG-based consensus algorithm, have proved to be more innovative solutions as an alternative to Ethereum, but their use case and longevity are still to be determined.

Turing complete - In modern computing, a property that implies a computer, language, or protocol capable of computing anything that is at all computable (thus emulating a universal Turing machine, named for its inventor Alan Turing).

Most modern programming languages are Turing complete, as well as most other smart contract protocols, including Ethereum.

If a network goes down while a sidechain holds a user's funds (like Solana recently), the user's funds are stuck until the chain is brought back online. There is no way around this for the user. However, rollups contain immutable "escape hatches" that ensure a user can exit back to mainnet even if the rollup network is offline. Users can always manually submit transactions to the mainnet Ethereum rollup contract as you need, including exiting the rollup with your funds.

Some sidechains are purposely built to be complementary to Ethereum and offload some specific Ethereum use cases onto themselves. Because of this, sidechains increase Ethereum's scalability by serving as external execution layers for L1. However, it's important to remember that sidechains may not always provide the same amount of security as L1 Ethereum.

Ethereum L2 Scaling Solutions

	Scaling Solution	Sidechains	Plasma	ORUs	Validium	ZKRUs
Category	Examples	Skale, POA	OMG, Matic	OVM, Fuel	StarkEx	zkSync, Loopring
Security	Liveness Assumption	Bonded	Yes	Bonded	No	No
	Mass-exit Assumption	No	Yes	No	No	No
	Validator quorum can freeze funds	Yes	No	No	Yes	No
	Vulnerable to hot-wallet key exploits	High	Moderate	Moderate	High	Immune
	Cryptographic Primitives	Standard	Standard	Standard	New	New
Performance / Economics	Max throughput - ETH 1.0	10k+ TPS	1k...9k TPS	2k TPS	20k TPS	2k TPS
	Capital-efficient	Yes	Yes	Yes	Yes	Yes
	Cost of transaction	Low	Very low	Low	Low	Low
Usability	Withdrawal Time	1 confirmation	1 week	1 week	1...10 min	1...10 min
	Time to subjective finality	N/A (trusted)	1 confirmation	1 confirmation	1...10 min	1...10 min
Other Features	Smart Contracts	Flexible	Limited	Flexible	Flexible	Limited
	EVM-bytecode portable	Yes	No	Yes	No	No
	Native privacy options	No	No	No	Full	Full

a. Polygon

Technically, Polygon is its own blockchain (with its own token: MATIC), but was built to become Ethereum's internet of blockchains. Polygon provides the architecture that enables developers to create custom, application-specific chains that leverage Ethereum's security similar to the Cosmos hub-and-spoke model. It provides an interoperable layer that can bridge many different projects and scaling solutions such as zk-rollups, optimistic rollups, and sidechains.

Since Polygon is a separate chain, it must be secured by a separate proof-of-stake consensus mechanism where validators stake MATIC. However, MATIC is staked in smart contracts on the Ethereum main chain. Polygon connects to Ethereum through a bridge with the use of a lock-and-mint mechanism. Users deposit funds into the bridge which locks them in a smart contract on Ethereum and mints the equivalent amount on Polygon. Polygon also maintains a secure relationship with the Ethereum main chain through periodic checkpointing, posting state changes to Ethereum, leading the Polygon team to characterize it as a "commit chain." To withdraw funds, you will have to go back through the bridge. The bridge (and funds) are secured by a 5/8 multi-sig scheme making it incredibly more centralized than the Ethereum mainchain. Additionally, ~33% of MATIC staked is run by a node controlled by Binance. These centralization factors should be considered when weighing the cost of transacting on a Layer 2.

However, as of Q4 2021, Polygon's proof-of-stake (PoS) sidechain is an industry leader with ~$5 billion in total locked value (TVL) deployed over 100 DeFi and gaming applications.

45.18%

AAVE Dominance on Polygon: As of Q4 2021, Polygon's assets under management or total value locked (TVL) is led by Aave, enabling users to lend and borrow their cryptocurrencies with reduced transaction fees.

In Q2 2021, Polygon released the Polygon SDK, developer tooling for launching new blockchains as rollups or their own chain, and Avail, a data availability innovation for Polygon chains.

In November 2021, Polygon announced Polygon Miden, a zk-rollup implementation. Polygon previously acquired Hermez (another zk-rollup) and are positioning themselves as the premiere scaling solution for blockchains. They also have a $1 billion fund for zk-based solutions and research.

Finally, in December 2021, Polygon proved yet again that the project has big plans in t``he L2 and rollup space. Polygon made yet another crypto-acquisition, this time purchasing the zk-rollup project, Mir Protocol, for $400 million. Polygon claims Mir Protocol contains the "fastest" ZK-proof technology, generating proofs and verifying more transactions faster than other comparable technologies.

b. Plasma

A Plasma chain is a L2 scaling solution that utilizes fraud proofs like optimistic rollups, yet maintains data availability off-chain (unlike optimistic rollups). Plasma was one of the earliest areas of L2 research but failed to gain much traction, especially as the advantages of rollups became evident.

Plasma and dAppchains are childchains tethered to the Ethereum root chain. Plasma received significant attention following the release of the corresponding paper by Justin Poon and Vitalik Buterin in August 2017. Nonetheless, the increasing complexities around practical challenges when it comes to implementing Plasma become significant concern.

Plasma enables the creation of an unlimited number of transaction-processing child chains (Ethereum mainchain clones) using smart contract and Merkle Tree technology. It is an attempt to create a more flexible state channel that enables many-to-many asset transfers with complex logic, as opposed to just simple one-to-one transfers. Like State Channels, Plasma is completely separated from the Ethereum L1.

One downside of Plasma is the long withdrawal period for users who want to remove their funds from layer 2. Another is the 'data-availability problem'. Since Plasma and the child chains are entirely disconnected from the main chain, it creates game-theoretic issues when the Plasma chain and the base layer chain try to sync up about the state of truth. The main chain can never with 100% certainty know the state of any Plasma chain, and thus cannot export its security to any child-plasma chain.

$1.83B

Plasma TVL: Total Value Locked is the total value currently locked in smart contracts in USD. It is generally seen as an indicator to evaluate a scaling solutions' adoption.

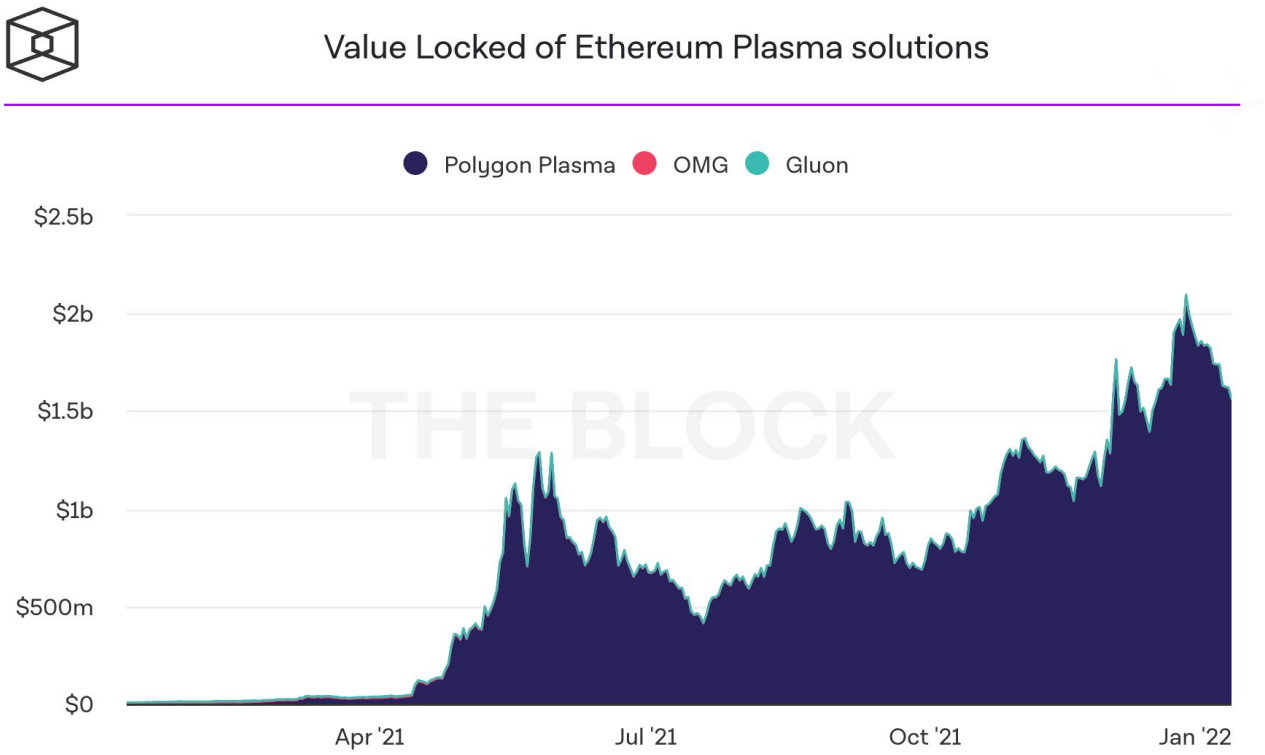

Ethereum Plasma Total Value Locked

SOURCE: The Block

5. Rollups

Rollups are a relatively new L2 solution being implemented on Ethereum that enable exponential scalability gains without sacrificing security. The primary innovation of rollups are that they move computation off-chain, while storing only the bare minimum of transaction data on-chain. The rollup chain handles all of the expensive and computationally-dense data processing, enabling exponential growth in Ethereum's ability to execute transactions. Again, in its simplest form, the rollup chain executes the Ethereum transactions on its own chain, "rolls" them up into one batch, and Ethereum receives and stores the results. However, in order to do so, the Ethereum mainnet needs some way to verify that the transactions that happen off-chain are valid. The answer is cryptographic proofs like validity proofs for zk-rollups (ZKRU) and fraud proofs for optimistic rollups (ORU).

To dig in a bit more, rollups generate a cryptographic proof (called a SNARK), and then only submit the proof to the base layer. The "batch" that is rolled up is periodically posted to mainnet Ethereum and contains the net outcomes of many different transactions as they occurred on the rollup layer. This data is verified and updated by the rollup operator every time the L2 advances its state. Therefore, L2 execution and L1 data update advance in lock step.

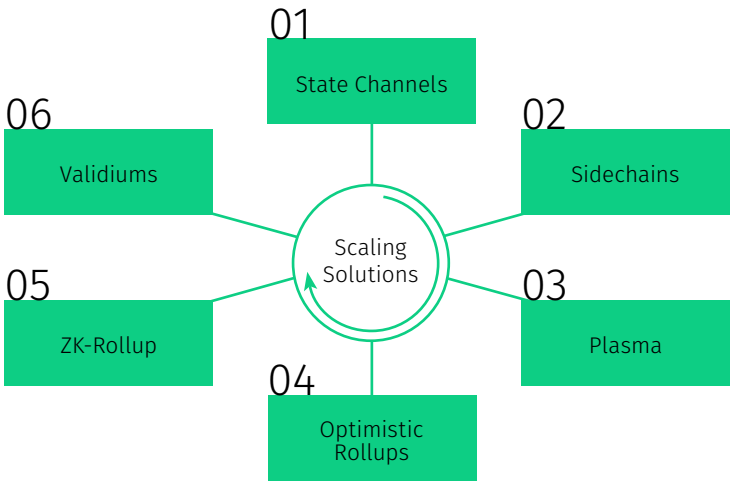

This removes the burden of data on layer 1 while also allowing Layer-2 transaction data to be available on layer 1 for validation. Rollups remove everything from being done on-chain (monolithic) to the Ethereum mainnet now serving as the settlement layer for off-chain L2 interactions (modular design).

A rollup needs orders of magnitude less validators than L1 to maintain its security. As long as a single honest validator does its job, the network will remain secure. Rollups can be thought of as branches off of the main trunk of Ethereum that increase the computation surface area of Ethereum.

With rollups, Ethereum can go from ~25 to 3,000+ TPS without sacrificing security. What makes rollups such an attractive scalability technology is the fact that users' funds cannot be stolen or censored (as is the case on some sidechains) and that no one can prevent users from exiting the rollup whenever they wish. Users can always access data on L1 to safely exit the rollup chain with their funds.

Technically speaking, a rollup is a single contract on the main L1 that holds funds and a cryptographic contractual commitment to a "sidechain" state. The sidechain/rollup is maintained by a small set of operators off-chain without significantly impacting L1 storage. The reason this "small set" does not introduce centralization is because a rollup's block producer could attempt to act nefariously, but if it does, the Ethereum L1 will simply reject the "bad batch" and financially penalize the bad actor.

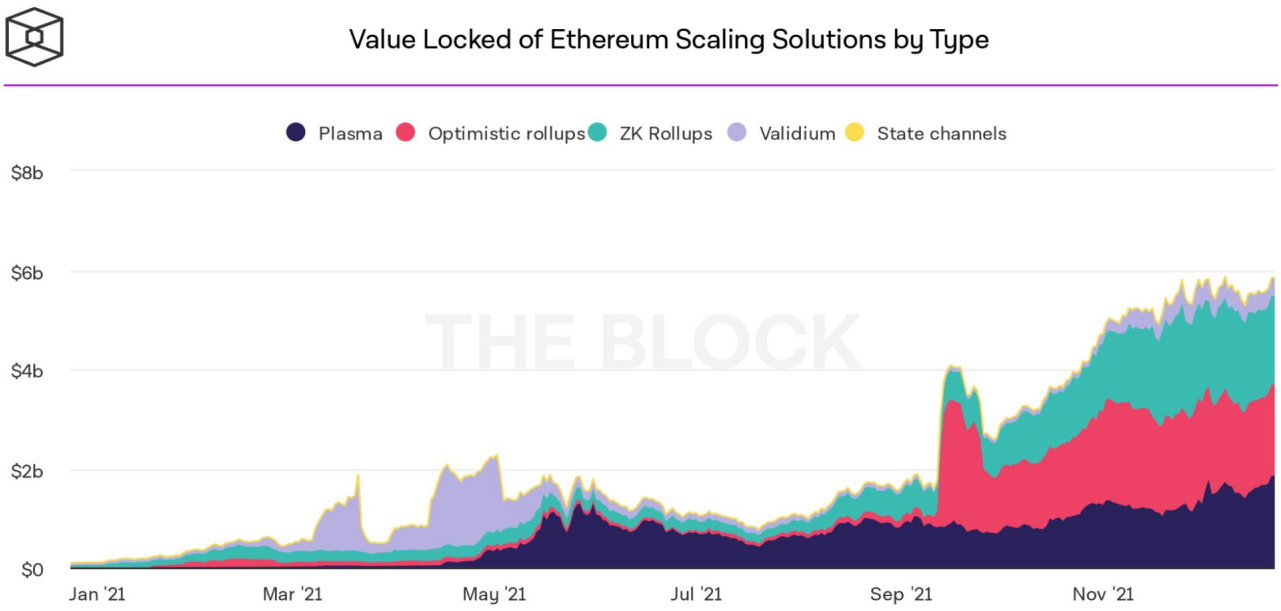

Ethereum Scaling Solutions Value Locked by Type

SOURCE: theblock.crypto

In previous sections, EVM-compatible sidechains and their potential benefits to Ethereum were discussed. Similarly, other alternative L1 blockchains like Polkadot, Solana, Cosmos, and NEAR could theoretically become rollups to Ethereum if they created a bridge that adheres to the rollup technical design pattern and post their data to Ethereum. This is a plausible future if alternative L1s fail to distinguish themselves and rollups on Ethereum become cheaper than competing chains.

Despite being extremely nascent, rollups are already significantly reducing fees for many Ethereum users. Ethereum founder, Vitalik Buterin writes in the Ethereum roadmap, "*l2fees.info frequently shows Optimism and Arbitrum providing fees that are ~3-8 times lower than the Ethereum base layer itself, and zk-rollups, which have better data compression and can avoid including signatures, have fees ~40-100 times lower than the base layer.*"

Scalability is improved on the base layer due to the lack of reliance on Layer-1 storage. The only factor on the scalability of a rollup is how much data the main chain can hold. This is why shards will complement rollups nicely as they increase the data availability of Ethereum (think 64 data centers vs just one). Once sharding is live (2023 or later), there will be an almost 20x increase in capacity, allowing rollups to operate cheaper and faster.

Rollups offer similar capabilities as Plasma, but without suffering from the "data availability problem." Layer-2 rollups batch users' transactions and post them on-chain via calldata. Posting the calldata on-chain is what allows Ethereum and its robust, decentralized network of nodes to "check the work" done off-chain. Instead of doing the computation, the calldata enables the Ethereum mainnet to quickly and easily verify that everything done off-chain was valid and accept the state changes *i.e.*, double-check the work. It also enables users to check the block explorer, like etherscan, and follow their transaction. Additionally, the availability of data on the Ethereum L1 means that any computation completed on a rollup can be redone by the Ethereum base layer, if needed. Without sufficient data availability, transaction execution becomes an opaque, black box unable to be audited by the L1.

Rollups already improve!

A new upgrade to Ethereum which targets rollups (EIP-4488) is currently being considered by the Ethereum community and would reduce the cost of posting this calldata onto mainnet. Rollups offer many-to-many transactions, smart-contract capabilities, and significantly reduced total L1 block space requirements, all while extending Ethereum's security assurance to the L2.

a. Embracing a “rollup-centric” future

Sharding is L1 scaling and is still years away from being fully implemented. It’s far more complex and riskier when compared to rollups because it’s altering the actual base layer. This means the ~$500 billion network is at risk of any bugs or miscalculations in the sharding rollout. Meanwhile, rollups are available now and possibly even more powerful. Optimistic rollups are a promising extant scaling technology that can be incorporated—and expanded upon—quickly. They offer developers an easy way to migrate their existing dApps to the rollup chain with a reasonable degree of security/scalability tradeoffs. This alleviates Ethereum congestion and high fees that already exist.

Additionally, the Ethereum community realized that rollups could provide immediate value and only improve once sharding is implemented. This means Ethereum scaling development is hyper-focused on rollups (plus some plasma and channels) as a scaling strategy for the near to mid-term future.

In the future, Ethereum users will primarily conduct their activity on an L2 rather than the L1 due to the cheap transaction fees and security guarantees. Meanwhile, the Ethereum mainnet will become a settlement layer for the L2 ecosystem and major ETH whales.

b. Rollup Flavors

There are two primary types of rollups: zk-rollups (ZKRU) and optimistic rollups (ORU).

ZK-rollups are theoretically faster and more efficient than optimistic rollups, but they suffer from friction and compatibility issues when migrating smart contracts to Layer 2. ZK-rollups submit transactions to the mainnet using a cryptographic verification method called a zero-knowledge proof, specifically a zk-SNARK. zk-SNARKs allow someone to prove they have a particular piece of information without revealing its content. Popularized by Zcash for enabling anonymous transactions, zero-knowledge-proof technology provides scaling efficiencies for state transitions on the rollup chain that are then submitted to the main chain. This approach is also called validity proofs, *i.e.*, using highly complex cryptography to ensure the validity of L2 transactions. The proof is submitted and checked by an on-chain L1 contract.

While validity proofs are complex and expensive relative to optimistic fraud proofs, verification by the L1 is simple, making them cheaper than a regular L1 transaction. However, due to the complex computations, special-purpose hardware may be needed to run a node, creating a centralizing effect and a less open network.

Ethereum L2 rollup scaling solutions include:

Arbitrum

Aztec

Boba

DeversiFi

Hermez Network

ImmutableX

Optimism

StarkWare

zkSync

Optimistic Rollups, contrastingly, are not secured by cryptographic zero-knowledge validity proofs. Instead, ORUs "optimistically" assume all transactions are valid but allow for/use dispute resolutions, a withdrawal period and crypto-economic incentives to maintain the integrity of the data. It is, essentially, an innocent-until-proven-guilty model with watchdogs in place.

Anyone may submit a rollup block. However, all other nodes can execute the duplicate transactions, essentially "checking the work" of the submitter. Only one honest actor is needed to submit the fraud proof and challenge any questionable block. This means fraud proofs are not sent with every batch of transactions. Instead, they are only used when an entity wants to dispute a transaction—*i.e.*, an attempt to prove any fraudulent transactions in a rollup batch.

They sacrifice some scalability for increased compatibility with the Ethereum Virtual Machine. Optimistic rollups run an EVM-compatible Virtual Machine called OVM (Optimistic Virtual Machine), which removes the compatibility issues in zk-rollups. This is exceptionally critical as composability is paramount in the Ethereum ecosystem, especially in DeFi. Using a virtual machine called the Optimistic Virtual Machine (OVM) allows developers to deploy code and projects on the secondary chains easily. On the other hand, there's no cryptographic proof that the state transition submitted to the main chain is correct.

6. Optimistic rollups

It is important to remember that while rollup technology can be quite technical, at its core, an optimistic rollup chain is simply a smart contract on mainnet Ethereum with some number of block producers that watch for transactions, batches them together into one string of data (rollup), and then posts it back to Ethereum mainnet with a signature attesting to their validity.

An optimistic rollup moves the heavy computation and data storage typically executed on L1 Ethereum off-chain to a new rollup network. Only a small portion of each batch of transactions is ultimately recorded on the mainnet, creating a much smaller computational impact on the L1. Since only one small data portion is registered on L1 and most of the computation is handled off-chain, fees can be significantly reduced when compared to a transaction executed entirely on L1.

By default, optimistic rollups "optimistically" assume submissions are valid. However, that's not always the case. Checks and balances are put into place to combat this seemingly reckless optimism. After withdrawals, there's a period where anyone can identify and dispute transactions they believe are incorrect or fraudulent. If the whistleblower can mathematically prove that fraud occurred by submitting the correct fraud proof, the rollup will revert the fraudulent transactions, penalize the fraud, and even reward the watcher.

Jun 2019

Minimal Viable Merged Consensus: The first publication that became what we know as optimistic rollups today. From this description, a solid technique was solid enough to build upon.

The ability to post L2 transaction data to the L1 is critical because it enables everyone to reconstruct the current and historical state of the rollup chain. Many other scaling technologies do not have this ability and therefore are less powerful to a user who has been wronged.

The drawback to this system is the delay when users move funds between the rollup and Ethereum and for transactions to be considered final. Because "watchers" need time to detect fraud, users' funds typically take a week to be withdrawn and available for further use. ORUs can only be considered safe with a ~ one-week challenge window. These dispute windows are expected to come down over time, and, in fact, some third-party solutions (HOP, Connext) already exist to remove this delay entirely. These are discussed further below.

Unlike the previously discussed sidechains, the breakthrough for rollups is simply increased scalability without sacrificing user security. ORU chains are secured by Ethereum L1. Users could be inconvenienced if a dispute or fraud situation arises, but their funds are always safe. Sidechains—such as Polygon—are secured by a separate validator set that is less secure than the Ethereum network. Additionally, the bridge connecting sidechains to Ethereum is typically highly centralized around just a few individuals. If less than ten people are compromised, all funds could be vulnerable.

One final advantage of ORU vs. ZKRU is the OVM: Optimistic Virtual Machine. OVM enables (almost) anything possible on the Ethereum mainnet to be possible in the ORU. Smart contracts, and therefore dApps, are easily transferable to the ORU because the OVM supports writing code in Solidity.

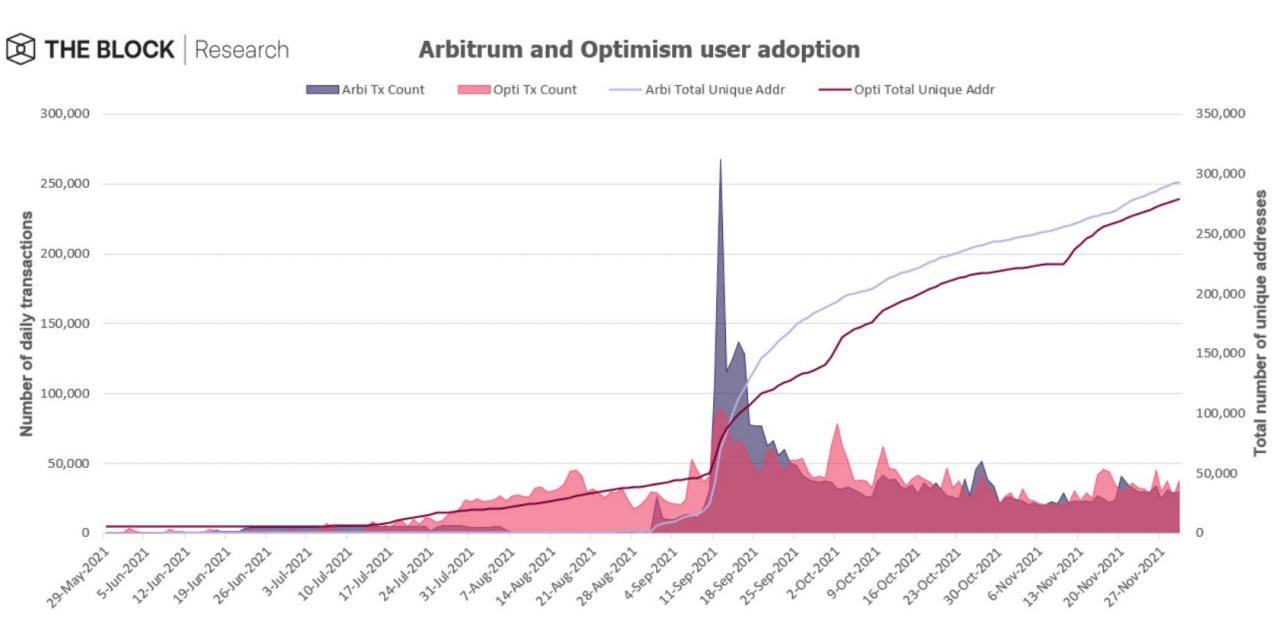

Arbitrum and Optimism User Adoption

SOURCE: theblock.crypto

In November 2021, Optimism PBC announced “EVM equivalence,” the complete compliance with Ethereum’s technical specification. This means that everything that currently exists and works on the Ethereum stack can now easily be integrated with Optimism’s ORU. This should drive tremendous network effects to Optimism as it’s now trivial for current projects to launch on the ORU. By reducing the friction, developers and users alike can now enjoy the benefits of ORU.

Arbitrum (by Off-chain Labs) and Optimistic Ethereum (by Optimism) are the two primary ORU projects on Ethereum. However, both implementations are still in their very early stages, with centralized companies primarily responsible for their success or failure. Both plan to decentralize over time, but any timeline estimate is simply a guess.

Both Arbitrum and Optimism launched in 2021, albeit with self-imposed limits and restrictions in case any bugs were encountered. Over time, more battle-tested and less constricted versions will be released, further reducing fees for users. Currently, neither Optimism nor Arbitrum One has implemented data compression, which, when fully released, could reduce fees by ~10x. A big step forward happened for Optimism, which launched its latest upgrade OVM 2.0, and Arbitrum’s next upgrade ‘Arbitrum Nitro’ promises to increase speed and reduce costs.

It is estimated that, once mature, optimistic rollups can offer anywhere from a 10–100x improvement in scalability and, at full scale, can possibly improve Ethereum transaction fees by ~50x.

However, as promising as rollup technology is, they are still very new technology that is not without risk. Arbitrum One, a specific kind of optimistic rollup discussed later, experienced downtime for around 45 minutes in September 2021 when a bug caused a large burst of transactions to overloaded the system. Optimism (OΞ), another optimistic rollup chain, also experienced a temporary outage (~one hour) in November 2021 in which its L2 transactions were halted. No funds were at risk during either issue (the beauty of L2s!) but processing new transactions was not possible making them useless until the matter was addressed.

One obvious note is that both Optimism and Arbitrum lack native tokens. It is not public knowledge whether either intend to eventually launch tokens but the general trend in the crypto industry woud suggest so. Regardless, both have had to try and bootstrap their rollups without lucrative airdrops or incentive programs (yield farming) . In an industry awash with 50%+ APY, five-figure airdops, and 8 figure incentive programs/funds, rollups, thus far, have chosen to try and grow without a token, making adoption an uphill battle.

~4900%

Transaction Fee Decrease: It is currently estimated that ORUs can reduce transactions fees by up to 50x once mature.

a. Optimistic Rollups Pros/Cons

Pros:

- Increase in scalability of ~2,000 TPS, reducing transaction costs by >5x
- Superior compatibility with Ethereum mainnet, less friction for developers to deploy projects (*e.g.*, EVM equivalence), can create and ship faster than ZKRU
- Flexibility in generalized computation (Turing-complete / EVM compatible)
- All data is available on-chain (no need to trust off-chain data providers)
- Computationally less expensive than ZKRU

Cons:

- Fewer TPS when compared to zk-rollups
- Relies on crypto-economic incentives and "watchers" rather than mathematically-certain security (fraud proof vs validity proof)
- Users (technically) need to wait 1+ week(s) for dispute period after a withdrawal from the rollup before being able to access funds

Additionally, ORUs and their challenge period are susceptible to 51% attacks. In this scenario, the attacker would try and introduce "bad" transaction data into the rollup and then attempt to censor any attempts to challenge it during the challenge period. The attacker is ultimately trying to corrupt the state of the rollup (with fraudulent data for their own self-interest) and stop anyone from challenging the submission.

This is why an adequately lengthy withdrawal/challenge period (1-2 weeks) is needed. An attacker may be able to censor or sneak a transaction through if the window was short enough but the longer the window, the harder it is to fool the rest of the chain.

Optimistic Rollup Tools

- Block explorer - Optimistic Etherscan
- Native bridge - Optimism Gateway
- User guide
- Live applications portal
- Network RPC config - Chainlist (search for Optimistic Ethereum)

b. Optimism OΞ

Optimism is a public benefit corporation (PBC) that created Optimistic Ethereum (OΞ), a leading optimistic rollup on Ethereum. Optimism aims to create a seamless L1-to-L2 developer experience by enabling (nearly) "copy and paste" code from one layer to the next, thanks to its OVM. OVM stands for Optimistic Virtual Machine and is the virtual machine that executes all transactions in rollup. OE is working towards "EVM equivalence" which enables the OVM to be an equivalent to the EVM in all technical aspects. Developer tools can be seamlessly built on/ported over to the new OVM 2.0 (the environment that enables EVM equivalence) from the tools already live on Ethereum mainnet.

Optimism launched with controlled rollout where a whitelisted group of dApps are approved to launch, most notably Uniswap, Synthetix, and 1inch. This limited release hampered Optimism adoption early on as it had onboarded only 6 dApps compared to ~60 for Arbitrum. However, on December 16th, 2021, the Optimistic team removed the developer whitelist for a full, open system which will allow all dApps to begin building on Optimism, if they so choose.

OPTIMISM

Optimism OΞ Pros/Cons

Pros:

- EVM-equivalence and better developer experience (existing tooling and programming languages)
- Easier dApp migration for existing dApps (L1 to L2), can create and ship faster than ZKRU
- All data available on-chain
- Computationally less expensive than ZKRU

Cons:

- Less theoretical/maximal throughput vs. ZKRU
- Centralized sequencer
- Longer withdrawal period
- Fraud proofs mechanism not published yet

c. Arbitrum

Arbitrum is an optimistic rollup L2 built by the Offchain Labs team. The currently-live implementation is called Arbitrum One and utilizes fraud proofs, on-chain calldata availability, a ~1 day withdrawal period, and a special type of node called a sequencer. Offchain Labs currently operates Arbitrum's sequencer, which has the ability to control the ordering of transactions. This early-stage centralization was mentioned previously and is not solely applicable to Arbitrum.

Arbitrum boasts a shorter ~1 day withdrawal period compared to Optimism 1-2 weeks but the trade-off is disputes on Arbitrum take longer to resolve. So, the majority of the time, Arbitrum withdrawals are quick and easy but in the rare occasion that a transaction is challenged, Arbitrum has some added complexities when compared to Optimism.

$1.78B

Total value locked in Abritrum contracts

4.96M

The total amount of transactions processed on Arbitrum platforms

$329B

Total market capitalization

While both optimistic rollups, Arbitrum has some key differences from its counterpart, Optimism. One critical difference is the Optimism OVM 2.0 is EVM-equivalent, running directly inside the EVM, while Arbitrum One is only EMV-compatible. This reduces code complexity and audit surface for Optimism. Arbitrum's AVM lacks EVM-equivalence because it consciously optimized for more compact fraud proofs but at the expense of implementation complexity.

Both are still incredibly easier to work with for developers than zk-rollups, but the Optimism EVM equivalence reduces all friction. However, while Optimism is still a permissioned network with only pre-approved projects, Arbitrum is completely public for any project.

Another critical difference is the reduced amount of data Arbitrum places on L1 as it executes transactions between postings. Optimism requires a posted state hash after every transaction, whereas Arbitrum executes several transactions before requiring a state hash to be posted. This can account for up to ~4x difference in on-chain storage.

Arbitrum One is currently the L2 network with the highest TVL. For an overview of the Arbitrum ecosystem of applications, see the Arbitrum Portal. Adding to its early adoption, Binance has open withdrawals to Arbitrum, becoming one of the first exchanges to open an on-ramp to Ethereum's layer 2. Crypto.com has also announced support for Polygon and Arbitrum, continuing the trend. Additionally, Arbitrum has partnered with Chainlink nodes and oracles to provide its validation services. This is a positive as Chainlink is already utilized in hundreds of Ethereum L1 projects and will bring the same security and composability to L2.

Arbitrum Pros/Cons

Pros:

- EVM compatibility
- ~1 day withdrawal period (under normal circumstances)
- No whitelisted rollout, enabling more dApps to be deployed early on
- Non-custodial and Ethereum wallet compatible

Cons:

- Currently uses centralized sequencer which carries front-running risk
- Less composability with EVM than Optimism
- Complexity switching between rollups and sidechains while guaranteeing high security

~1-2 weeks

Withdrawl Time Saved on Arbitrum: a single-currency transaction transferring a balance from your account to an external address, usually a wallet or an exchange.

Transactions are confirmed on-chain and can sometimes be overloaded. When this happens, higher transaction fees are generally required for the transaction to be confirmed faster, otherwise, it is stuck in the queue for long periods of time.

VI. Arbitrum Tools

- Block explorer - Arbiscan
- Bridge - native Arbitrum bridge
- AMM aggregator - 1inch
- Arbitrum bridge tutorial

ORU Honorable Mentions

Resources:

- Block explorer
- Boba Network Gateway (bridge)
- Developer portal

$471.95 M

Value locked in Boba Network as of Jan 1, 2022 (Zerion API)

d. Boba

Built by the OMG Foundation, Boba is an ORU scaling solution that originally began as a fork of Optimism and the OVM (optimistic virtual machine). Boba offers fast withdrawals backed by community-owned liquidity pools similar to other bridge solutions, reducing the challenge period from ~7 days to minutes, while incentivizing Liquidity Providers (LPs) with yield-farming opportunities. The team plans to completely rewrite the codebase for their upcoming v3 which is set to be rolled out on mainnet in the coming months. Boba is production-ready with a functioning bridge and a native dex called OolongSwap.

Resources:

- Chain explorer
- Developer Docs
- Charts

$283.07 M

Value locked in Metis as of Jan 1, 2022 (Zerion API)

e. Metis

Metis is an L2 scaling solution on Ethereum that is best described as a sharded optimistic rollup. The Metis Virtual Machine (MVM) contains various decentralized autonomous companies (DACs) with their own separate, application-specific computational and storage layers. Despite the separate execution layers, liquidity between the shards can flow frictionlessly due to the MVM cross-layer communication protocol. The goal is to scale horizontally with distinct, application-specific execution layers that are while also preserving the security of Ethereum via fraud proof submission to mainnet.

Pros:

Parallel sequencers
Withdrawal period could (theoretically) take minutes (rather than days)
Plans to inherit Optimism's EVM Equivalence

7. ZK-Rollups

ZK-rollups (ZKRUs) are separate blockchains networks with very few specific nodes (called provers). Sounds like other alternative L1 chains, right? However, ZKRUs have a cryptographic proof linking them to the Ethereum mainnet. This link prevents the rollup from censoring or stealing funds while maintaining the immutable properties of the Ethereum L1. This proof is called a validity proof, ensuring the validity of the off-chain transactions, making them instantly verifiable and removing the need for a withdrawal/challenge period.

ZKRUs improve scalability by moving computations and storage off-chain where computation is expensive. Zero-knowledge cryptographic proofs reduce the computing and storage resources for validating the block by reducing the amount of data held in a transaction; zero knowledge of the entire data is needed.

Remember, rollups batch together large amounts of off-chain transactions, compress them into a single transaction, and eventually find their way to the Ethereum L1. Because ZKRU do not assume all transactions are valid, validity proofs must be sent with every zk-rollup batch to cryptographically prove the validity of transactions. While a bit more technically cumbersome, this means that transactions are final once they are validated by the settlement layer.

To describe the process in detail,

- the highly-compressed batch of transactions are combined together with the current state root
- that combination is sent to an off-chain prover
- the prover computes the transactions, generating a validity proof of the results
- the prover then sends this to an on-chain verifier (Ethereum nodes)
- the verifier verifies the validity proof
- the smart contract on Ethereum's L1 that maintains the state of the Rollup is updated to the new state

Remember, in traditional L1 blockchains, more transactions lead to more expensive fees due to limited block space. However, for a ZKRU, the opposite it true! ZK-rollups work off of economies of scale, meaning more transactions makes the network cheaper to use. This is counterintuitive to a typical blockchain, but is possible because the costs are amortized across all participants. Verifying the validity proof on Ethereum has a certain cost and as the number of transactions insincluded in a rollup batch grows, the cost to

Jun 2019

Minimal Viable Merged Consensus: The first publication that became what we know as optimistic rollups today. From this description, a solid technique was solid enough to build upon.

verify grows slower than the number of transactions added (exponentially slower). Therefore, the more users, the more the cost is spread around.

On mainnet Ethereum, each transaction is executed by every node. With ZKRU, only one node needs to actually do the computation (the provers) and then produces a zero-knowledge proof of it. As mentioned prior, provers are a select set of nodes in charge of computing all the transactions and aggregating them into a zk-SNARK. Because of the complicated computations involved, the provers run on dedicated hardware, making them more centralized and opaque. The good news is that because of the validity proof, it is mathematically impossible for them to submit fraudulent data. The only trust involved is trust in the cryptography/mathematics.

Once the prover has submitted the proof, every other node (verifier) simply verifies this proof instead of having to do the full computation. The proof allows each node to verify the provided state is valid. Verifying the proof is much less intensive than actually computing it, which is where the scalability improvements are created. Therefore, verifiers don't need special high-end hardware to verify the proof. They simply use their existing hardware, creating no new stress or burden for current nodes. Only state transitions and a small amount of calldata need to be processed and stored by the nodes. With this system, nodes can easily agree on a common state and it puts the burden of execution on a single node instead of the whole network.

+$1.8B

Total Value Locked in ZK-Rollups: While still new to the market, zk-rollups are showing potential for strong adoption in the future.

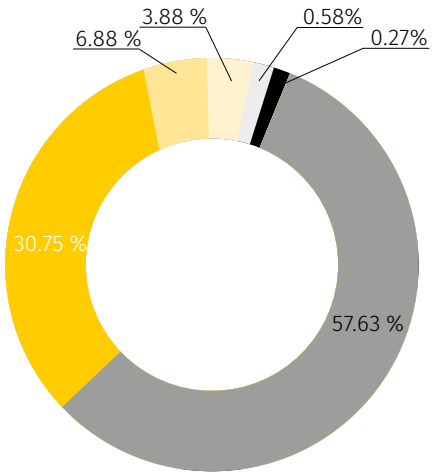

Market Share by Total Value Locked in ZK-Rollup

- dYdX
- Loopring
- zkSwap v2
- zkSync
- Aztec
- zkSwap

Source: Dune Analytics

Beyond simply the scaling benefits, ZKRUs are doubly-impressive due to their economic security guarantees. In a ZKRU, rollup operators must submit a Zero-Knowledge Proof (SNARK) for every state transition that then gets verified on the mainchain. This SNARK proves, by using world-class cryptography and math, that the batch of transactions (and their net state changes) are valid. Thus, it's impossible for the operators to commit an invalid or manipulated state.

It is not possible for operators to steal user funds or corrupt the rollup state. ZKRU relies on Layer 1's censorship-resistance, but not security, which means there is no need for anyone to monitor it. After a block has been verified, a user's funds are always guaranteed to be retrievable, even if the operators refuse cooperation.

a. ZK-Rollup Pros/Cons

Pros:

- Greater scalability and transaction cost reduction compared to optimistic rollups
- Transaction data reduction increases throughput and scalability
- No fraud dispute window required like in optimistic rollups, reducing withdrawal times from ~2 weeks to a few minutes
- Enables privacy by default

Cons:

- Zero-knowledge proof computing will require data optimization to get maximum throughput
- Initial set up relies on a centralized structure
- More difficult to initially build and integrate into the Ethereum network than optimistic rollups

Thus, ZKRUs embody the original ideals of cryptocurrencies and the cypherpunks that created them. They remove the need for trusted parties and replace them with cryptography and game-theoretical incentive alignment.

Another benefit of ZKRUs are that SNARKS can prove all of the computation is correct without having to actually reveal the details of the transactions! Zero-knowledge technology allows for someone to prove something while not revealing the contents of that information. For example, ZK-SNARKs enable Joe to verify Sally's banking information using a zero-knowledge cryptographic proof instead of revealing the confidential information to Joe.

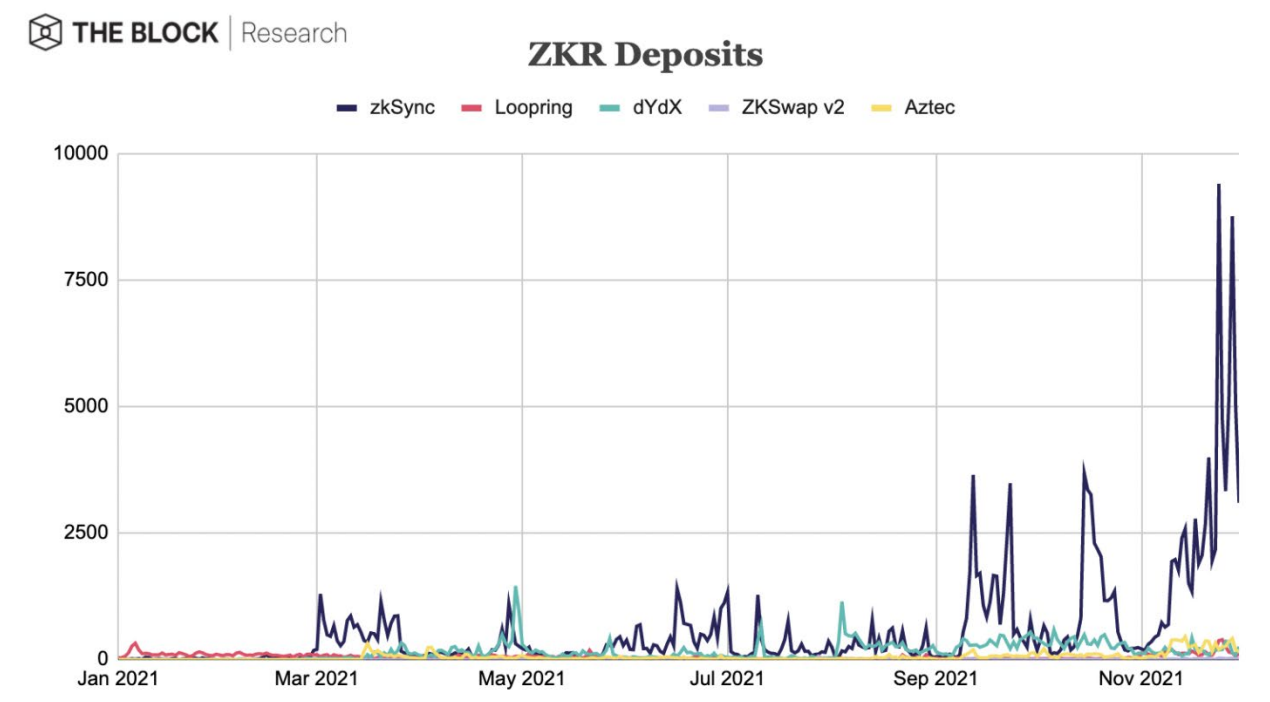

ZKRU Deposits **Source**: The Block

b. StarkWare & StarkNet

StarkWare is a leading firm that pioneered zk-rollups launching StarkEx in 2020, and again with StarkNet in 2021. The first iteration, StarkEx, supports the ability for smart contracts to run any arbitrary logic for specific use cases like trading and NFTs. As of Q1 2022, StarkEx has processed ~80 million transactions and a cumulative trading volume of ~$320 billion across the four protocols it hosts—dYdX, ImmutableX, DeversiFi, and Sorare.

StarkNet is StarkWare's next ZKRU iteration and the first to feature general smart contracts on a fully-composable network. Composability refers to the ability for applications to coordinate, build on top of one another, and interconnect—something for which StarkEx is not designed.

As another example of an Ethereum L2 ZKRU, StarkNet uses zero-knowledge proofs to minimze transaction times and hyper-scale without compromising security. An alpha version of StarkNet was launched in November 2021 with limited capabilities allowing developers to build on top of the protocol. StarkNet is designed to benefit from economies of scale *i.e.*, the greater the number of transactions in a batch, the less gas each participant in the batch must pay.

Under the hood, StarkNet compresses thousands of transactions into a single validity-proof called a 'STARK' that is submitted to the Ethereum L1. StarkWare's STARK technology has two primary advantages over zkSync's SNARKs; they do not require an initial trusted setup, and they are ~10x faster to compute than SNARKs.

Periodically, StarkNet transactions are batched and validity-checked in a STARK-proof by a sequencer on the Ethereum mainnet. The computational effort required to verify STARK-proofs is exponentially small when compared to proving the computation enabling StarkNet to scale Ethereum.

STARKWARE

StarkWare controls the sequencer and all of the transactions are verified by StarkWare cloud servers. Thus, StarkNet is not currently a permissionless system. However, StarkWare aims to create a decentralized sequencer set in the future.

It's important to note that StarkNet's transaction finality is tied to L1, meaning the L2 node must validate StarkNet and Ethereum simultaneously. StarkNet introduces a solution involving checkpoints to the Ethereum mainnet, enabling it to achieve effective finality on the rollup side very quickly. Therefore, all L2 nodes incorporate an L1 full node.

The data required to reconstruct the StarkNet state will be published on-chain. Anyone will be able to run a StarkNet node in the future, making StarkNet as secure and permissionless as Ethereum. Application deployment is permissionless, so anyone can write smart contracts and publish them on the testnet using Cairo, the native, Turing-complete programming language developed by them. A breakthrough in Cairo is enabling just one verifier to use a single proof to confirm the integrity of many different program executions. This has the effect of amortizing costs across separate dApps *e.g.*, a single proof that includes both dYdX trades and SoRare transactions.

StarkWare has commented, generally, that their plan with the StarkNet rollout will follow a similar path to that of Optimism (ORU): launch the network with a single sequencer and a limited whitelist of dApps early on to control the launch and limit any risks. A list of projects building on StarkNet can be found here. Ultimately, StarkWare hopes to grow the ecosystem into a Starknet "universe" while also decentralizing the network, nodes, and infrastructure.

The StarkWare team has also stated that while they do not currently have a token, it is their aim to decentralize StarkNet in the future. Launching a governance token similar to many other projects is one way in which they could do so.

As of Q1 2022, despite much fanfare, StarkNet remains in its early alpha phase with still much to prove at scale. Early disruptions and issues with its gradual rollout are likely. Despite that, StarkWare and OKEx announced in December 2021 a partnership designed to enable easy onboarding to StarkNet from OKEx sometime in 2022. Additionally, Argent, an Ethereum smart contrat wallet, also announced 'Agent X', the first wallet for StarkNet in Q4 2021.

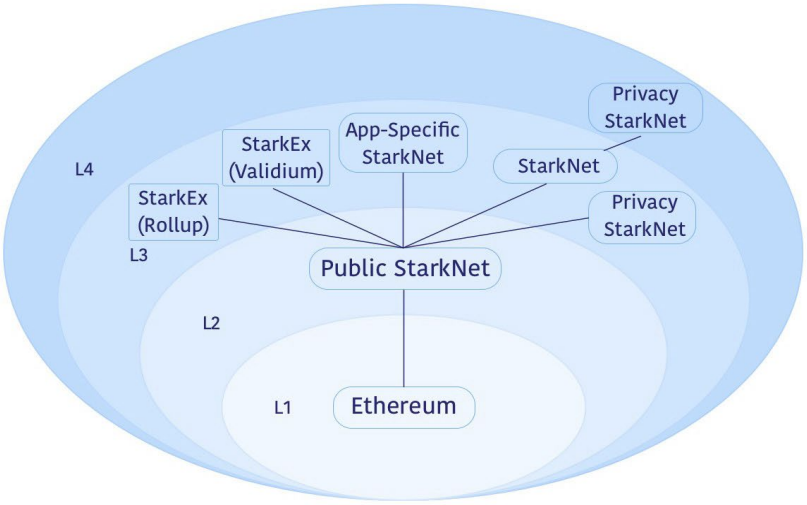

StarkWare Universe

Source: StarkWare.co

$1.19B

Total value locked in StarkEx projects

83M

The total amount of transactions processed on StarkEx platforms

$327B

Cumulative trading volume across all StarkEx platforms since 2020 launch

Source: StarkWare.co

Pros

- Increased TPS compared to ORUs (~9000+ TPS on Ropstein testnet)
- Faster withdrawals (no challenge period), enabling better capital efficiency and liquidity
- Volition (discussed below) unlocks even greater scalability gains for those that choose to make the trade-off on security

Cons

- Developer UX and dApps porting is more cumbersome and less friendly than ORU options
- Cairo language less popular among developers, meaning less talent pool to build on StarkWare products
- Technical challenges in solving data availability with Validium proudct. In particular, the trade-offs between transaction latency and cost and making data available on-chain.

c. Validium & Volitions

Nearly identical to a zk-rollup, Validium differentiates its mechanism with off-chain data availability. This means that zk-rollups post data on the L1 blockchain, while Validium's post on-chain validity-proofs but the data remains on a separate network. This enables Validium to achieve considerably higher throughput than ZKRUs or ORUs. By sending data off-chain rather than on-chain, it also reduces the cost of each transaction and increases the transactions per second (TPS).

Validiums also offer privacy benefits by keeping data off-chain as users' transaction and balance information is stored with the validium operator instead of publicly on the blockchain. However, because transaction data is not published on-chain, users are forced to trust an operator to make the data available when needed. This introduces new trust assumptions and centralization points where Validium operators could freeze users' funds.

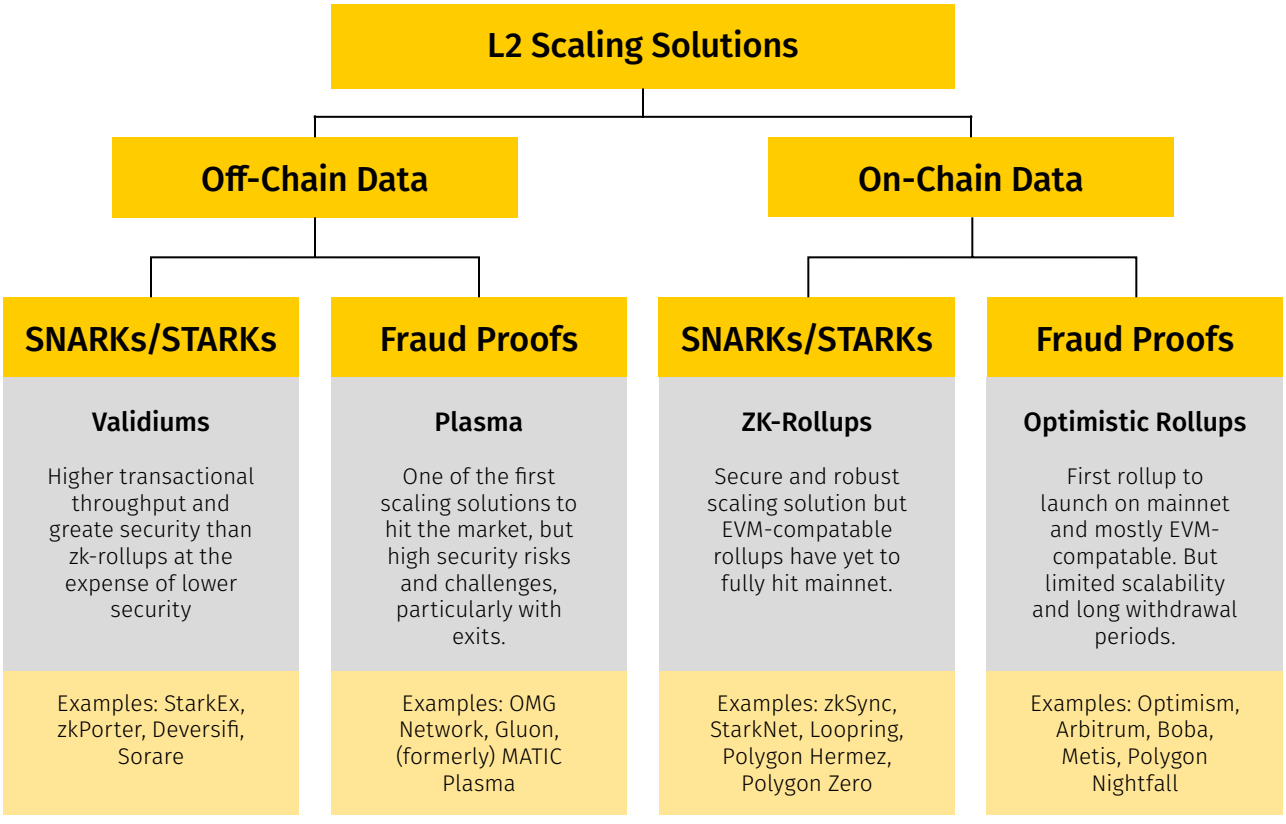

The trade-off for storing data off-chain requires trust in the third party who could prevent users from accessing their balances. StarkWare aims to solve this with a Data Availability Committee (DAC), a committee of 8 independent members that have their copy of the transactions made. They are also required to maintain this data by making it available at all times. If an operator prevents a user from accessing their funds, a committee member can override them to confirm their request if it is valid. Examples of where Validium is used: Loopring (LRC) and StarkWare.

Volitions are a zk-rollup and Validium hybrid solution that enables users to choose for data availability either on-chain or off-chain, *i.e.*, via Ethereum or through validiums.

	Validity Proofs		Fault Proofs
Data On-Chain	Volition	ZK-Rollup	Optimistic Rollup
Data Off-Chain		Validium	Plasma

Volition Chain Flexibility

Source: StarkWare

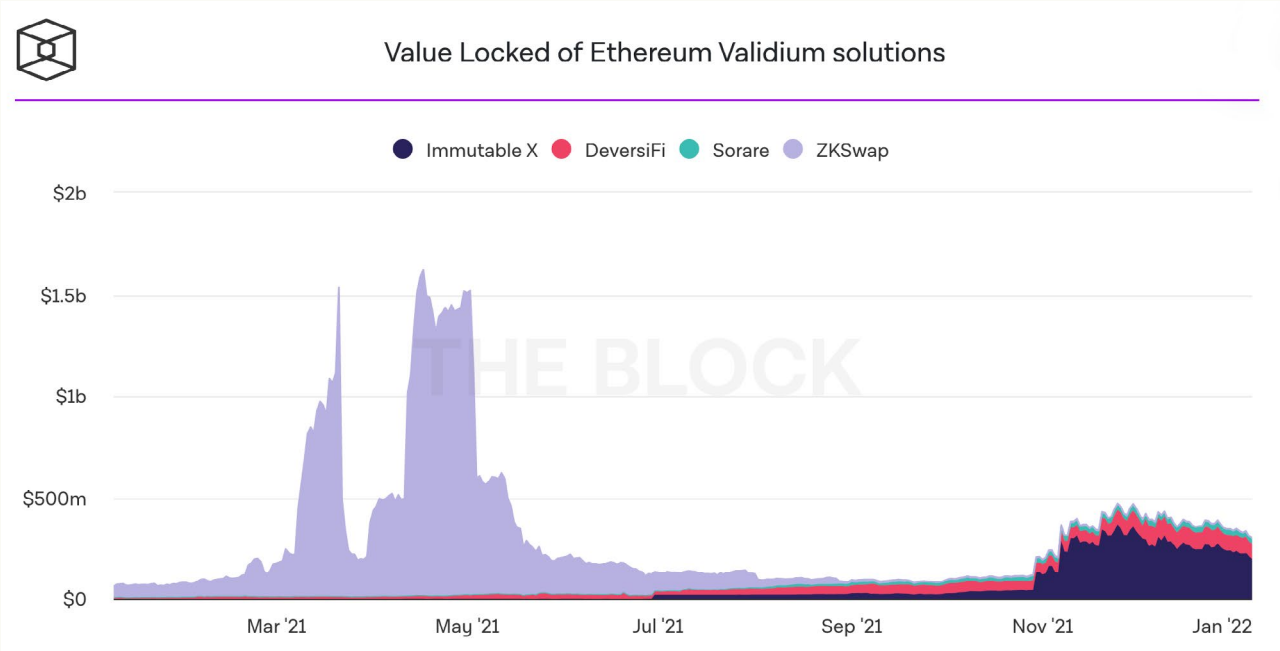

Value Locked of Ethereum Validium Solutions

Source: The Block

d. MatterLabs / zkSync

4 Million

Transaction Fees: As of Q1 2022, zkSync has processed over 4 million transactions with transfer feesless than $1.

Matter Lab's zk-rollup, zkSync, lets users deposit ETH onto the network and send payments between other zkSync accounts with lower transaction fees. It is a standard L2 zk-rollup scaling solution, in the sense that a smart contract holds all funds on Ethereum mainnet, computation and storage are performed off-chain, and every rollup block generates a zero-knowledge L1-verified proof. This methodology, paired with the power of SNARKs, means zkSync is unable to move or steal funds and is easier to integrate for EVM compatibility.

However, zkSync is slower than its StarkWare counterpart, in part, because it uses SNARKs—PLONK especially—and relies on a trusted setup at genesis. That means the entirety of the zkSync ecosystem depends on a trusted ceremony conducted in 2019. The good news is that the system is 100% provably secure if even just one participant was honest. The ceremony included many well-know and public crypto figures whose best interest were tied to the success of the launch. Therefore, this trusted setup is likely not an issue and uncompromised.

ZkSync launched its V1 in June 2020 with the only use case being for simple token transfers. However, in April 2021, the zkSync team announced zkSync 2.0 and its zkPorter technology, aiming to provide ~$0.01 transactions by moving transaction data off-chain (validium-style) and offering 20K transactions per second (TPS). It boasts a sharded-infrastructure design that interoperates seamlessly with zkSync.

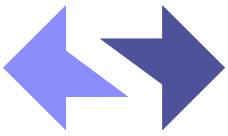

With zkSync 2.0, the L2 state will be divided into two distinct options: a zk-Rollup with on-chain data availability and the zkPorter option with off-chain data availability.

zkSync 2.0 is another L2 rollup that supports EVM programming languages like Solidity, Yul, and Vyper, and in the future, Rust and Zinc. This means developers can easily deploy EVM code onto zkSync 2.0, and for users, zkSync 2.0 offers instant withdrawals and objective finality limited only by batch frequency.

zkPorter will be part of the ultimate zkSync 2.0 vision. With zkSync 2.0, the L2 state will be divided into two distinct options: a zk-Rollup with on-chain data availability and the zkPorter option with off-chain data availability.

zkPorter is the internal consensus mechanism for data availability within zkSync 2.0, enabling the large TPS numbers. zkSync 2.0 can handle ~1,000 to 5,000 TPS as a standard ZKRU, but with zkPorter, it can accommodate ~20,000 to 100,000 TPS, depending on the complexity of each transaction. It should be noted that, when utilizing zkPorter, the user is relying on zkSync's internal consensus mechanism. This requires the user to trust Matter Labs and rely on a far less secure or decentralized rollup solution that leverages L1's consensus mechanism.

The good news is that users can choose either option based on their preferences and the trade-offs presented. zkPorter will offer negligible cost, but lower security for trivial transactions, and the zk-rollup mode provides maximum security. Both parts will be composable and interoperable: contracts and accounts on the zk-Rollup side will be able to seamlessly interact with accounts on the zkPorter side and vice-versa. The primary difference between zkPorter and StarkWare's Volition is that a user must choose by each zkPorter account whether to produce transactions with off-chain data availability. In Volition, a user can choose by each transaction within an account.

As of Q1 2022, zkSync has processed over 4 million transactions with transfers fees less than ~$1. Despite being relatively new, users began moving funds over to zk-rollup projects like Loopring and zkSync in 2021, especially in Q4, as the chart below illustrates. By November 2021, unique users increased by ~90,000, and deposits eclipsed ~$75 million. For zkSync, the wave of adoption can be attributed to its top projects, ZigZag Exchange and Gitcoin, a crowd-funding platform. According to L2fees, token swaps through ZigZag on zkSync have the lowest fees.

However, it should be noted that zkSync is currently highly centralized. Although the zkSync multi-signers have q-shared economic interests in the project's success, contracts can be upgraded anytime via the 9/15 multi-sig. Matter Labs claims "*the probability of bugs is significantly higher than a malicious collusion between the Matter Labs team and 9/15 members of the security council.*" They team has committed to develop the project, hitting future milestones including delivering V2 (however, no mainnet launch date has been set), supporting new exchanges, and decentralizing their security council council.

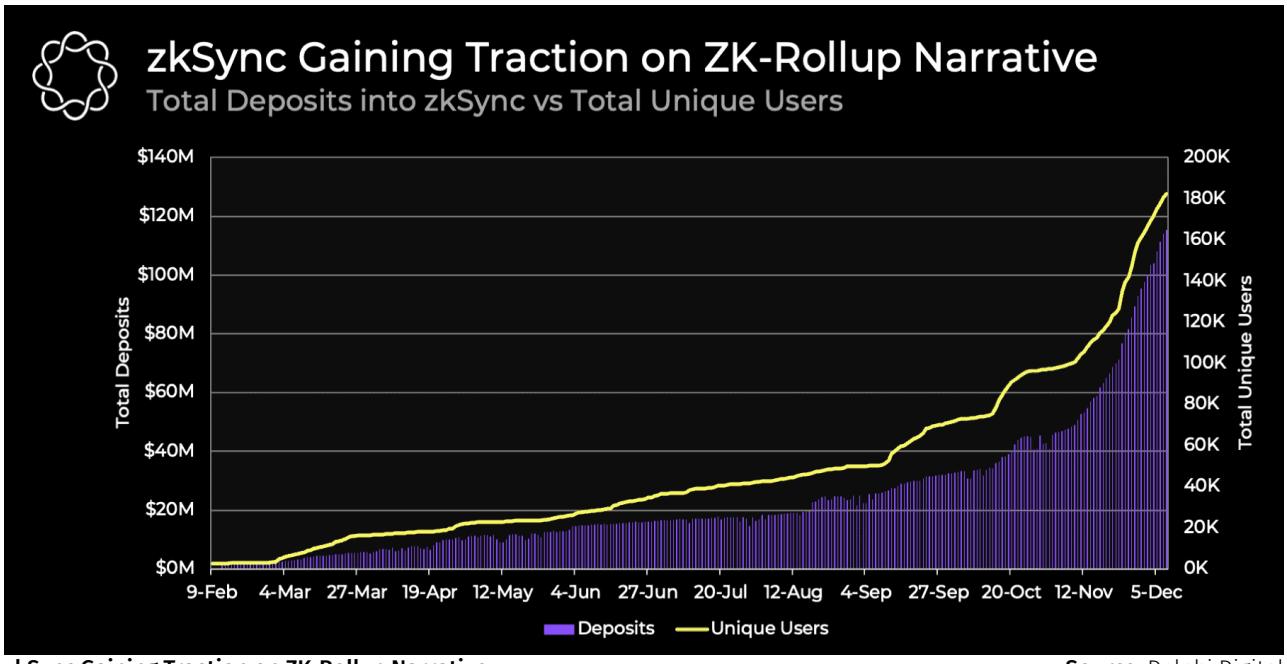

zkSync Gaining Traction on ZK-Rollup Narrative **Source**: Delphi Digital

Pros

- Less data contained in each transaction increases throughput and decreases fees
- No withdrawal periods and faster finality
- Inherent (and cheap) privacy

Cons

- Generalized smart contract support (similar to StarkNet) is not live or production-ready
- Initial trusted setup ceremony scares some, introduces trust
- New, less battle-tested cryptography

e. Polygon Hermez

Polygon Hermez is a zk-rollup that is the product of Polygon acquiring Hermez, merging it into the Polygon ecosystem. Polygon Hermez has announced its plans for full EVM-support (zkEVM) with a mainnet launch anticipated in Q2 2022. The Polygon Hermez protocol has an off-chain prover that validates transactions and generates a SNARK proof submitted to the on-chain verifier, just like other ZKRUs. However, it is not EVM-compatible, differing from most solutions previously discussed. Currently, Polygon Hermez is live and can be used by anyone as a payments platform, similar to zkSync v1.

Helpful Links

- Code
- Documentation

8. L2 Drawbacks

L2s are losing market share to L1s in recent months, despite the early stage migration of DeFi protocols from L1 to L2s. Most expected L2s to immediately become a hot-spot for developers and users priced out of Ethereum mainnet. But to the detriment of Ethereum, other L1s, especially EVM-compatible chains in which users can easily bridge over their ETH, stole the limelight. Ecosystems, like Polygon and Avalanche, that dedicated a portion of their token treasuries to user incentives were vital in making this happen.

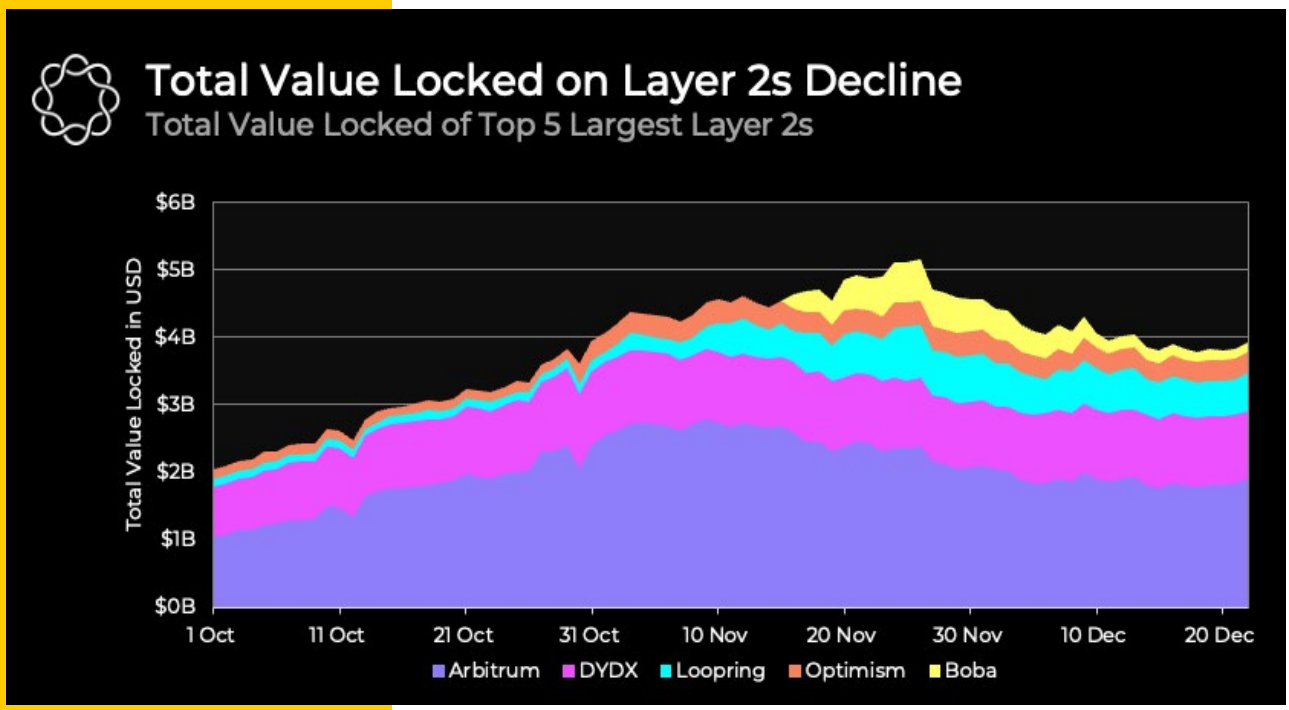

Total Value Locked on Layer 2s Decline **Source**: Delphi Digital

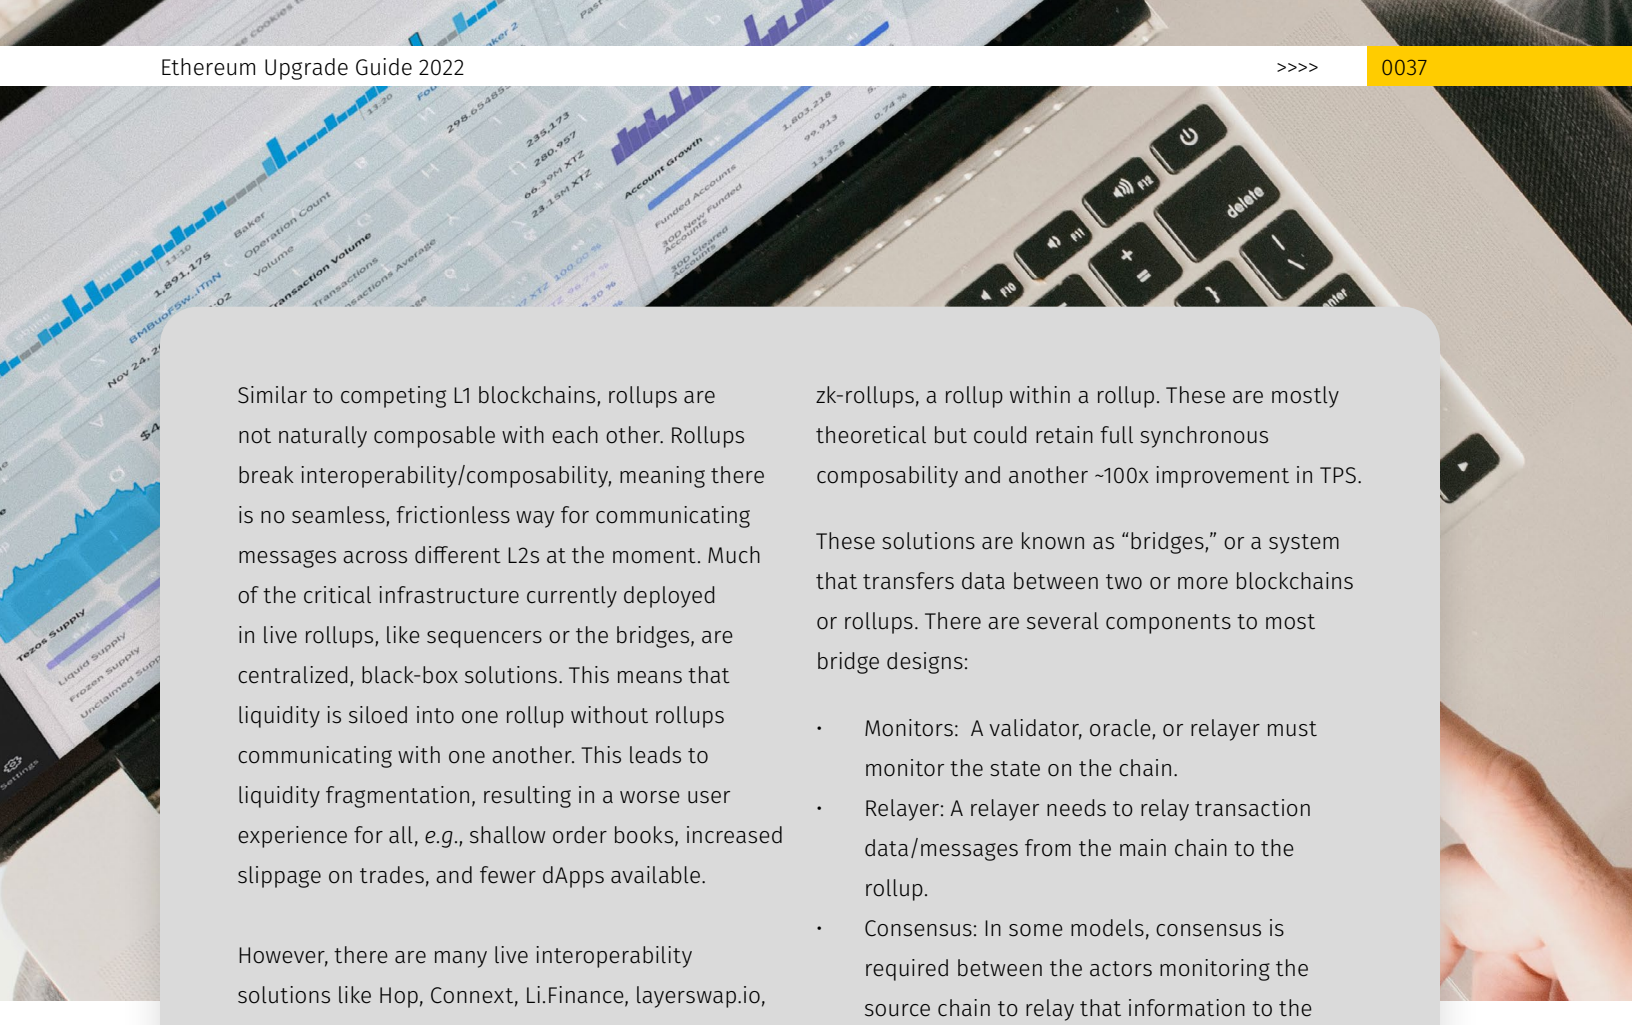

Similar to competing L1 blockchains, rollups are not naturally composable with each other. Rollups break interoperability/composability, meaning there is no seamless, frictionless way for communicating messages across different L2s at the moment. Much of the critical infrastructure currently deployed in live rollups, like sequencers or the bridges, are centralized, black-box solutions. This means that liquidity is siloed into one rollup without rollups communicating with one another. This leads to liquidity fragmentation, resulting in a worse user experience for all, *e.g.*, shallow order books, increased slippage on trades, and fewer dApps available.

However, there are many live interoperability solutions like Hop, Connext, Li.Finance, layerswap.io, cBridge, dAMM, and more that are already working to "bridge" liquidity and remedy this issue. In addition, projects are already working on internally-sharded zk-rollups, a rollup within a rollup. These are mostly theoretical but could retain full synchronous composability and another ~100x improvement in TPS.

These solutions are known as "bridges," or a system that transfers data between two or more blockchains or rollups. There are several components to most bridge designs:

- Monitors: A validator, oracle, or relayer must monitor the state on the chain.
- Relayer: A relayer needs to relay transaction data/messages from the main chain to the rollup.
- Consensus: In some models, consensus is required between the actors monitoring the source chain to relay that information to the destination chain.
- Signing: A participant needs to cryptographically sign the data sent to the destination chain.

Another key obstacle for L2 adoption is the user experience and cost onboarding to an L2. The obvious solution is fiat and exchange onramps directly to an L2 . As of Q1 2022, almost no centralized exchanges support native withdrawals to L2s. This means a user must first deposit to the L1 and then bridge over to the L2. This is costly and adds friction to the user experience. A current workaround is to use an exchange to withdraw to a sidechain like Polygon PoS which has sufficient liquidity in cross-chain (centralized) bridges like Hop or Connext.

9. Liquidity Bridge Solutions

Hop protocol

Helpful Links

- Hop Exchange
- Code

Hop Protocol is a rollup token bridge that relies on market makers—known as bonders—to provide liquidity for others in return for a fee. Hop allows users to send tokens from one rollup or sidechain to another almost immediately without having to wait for the network's challenge period. Hop Exchange is the front-end cross-L2 bridging protocol app built on top of the bridge system.

This automated market maker (AMM) system lets users send cross-chain transactions and provides an opportunity for liquidity providers to earn yield on their capital. However, there is no Hop token, yet. L2 canonical tokens are exchanged for Hop Bridge Tokens, which can be swapped for the underlying asset on L1 or a different L2 token.

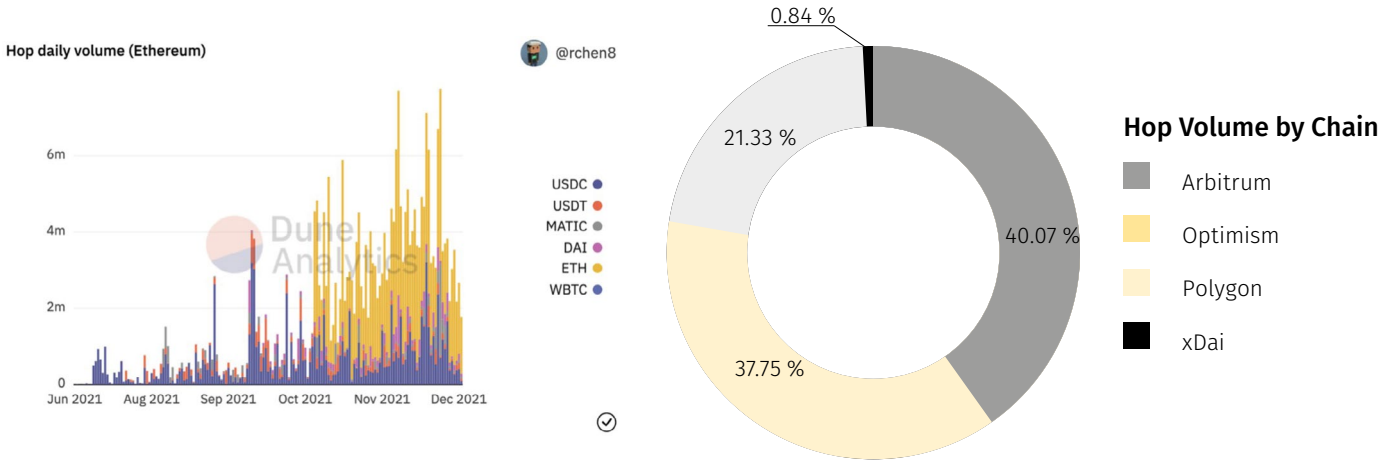

Source: Dune Analytics

Connext

Helpful Links

- Code
- Documentation

Connext is a cross-chain liquidity bridge—a network of pools on different L1 and L2 networks connecting liquidity to nine EVM-based chains. Users swap values between these pools, similar to AMM DEXes like Uniswap. It is a non-custodial and very capital-efficient way to bridge assets from one chain to another. Since their launch, Connext has facilitated over $350 million in volume, spread across 305k transactions.

In the month of December, Connext facilitated $153 million in volume across 135k transactions, averaging a daily volume of $4.9 million.
Routers act as the backbone of the Connext network, providing liquidity for user swaps and

earning fees in return. There are currently 22 active routers on Connext, providing $19 million in liquidity. The main priority for routers is rebalancing liquidity between chains, as chains with a lot of outflows concentrate liquidity while it becomes scarce at chains with a lot of incoming transactions.

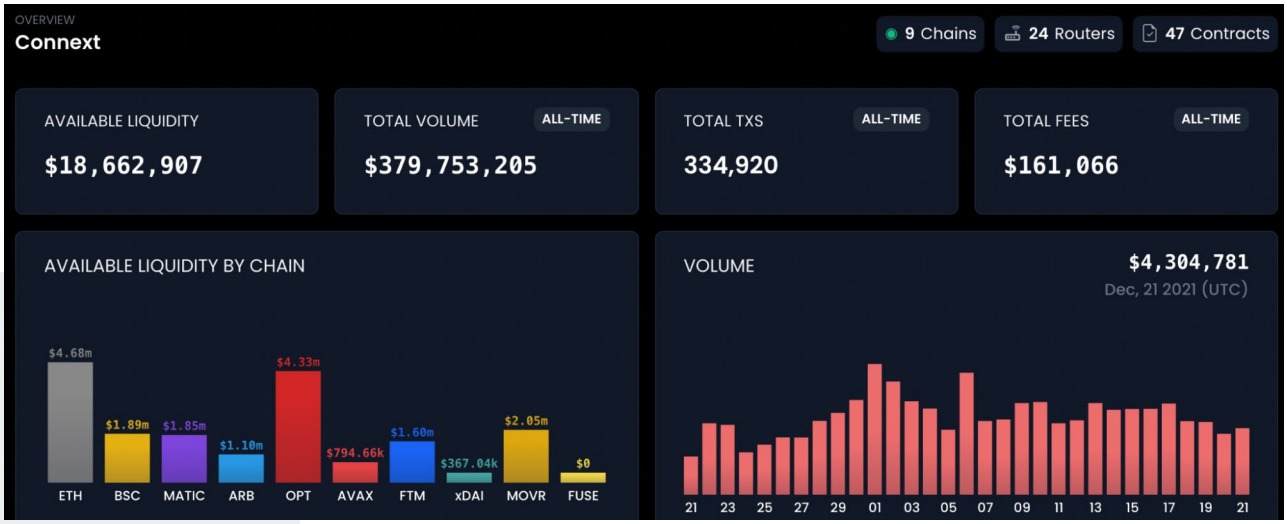

Connext Trading Data **Source**: ConnextScan.io

$1.12B

Total value locked in Synapse bridges

$472M

Total market Capitalization

$3.64B

Total bridged volume across all Synapse bridges

Source: SynapseProtocol.com

Synapse protocol

Synapse is a cross-chain layer-2 protocol powering interoperability between blockchains. Through decentralized, permissionless transactions between L1, L2 and sidechain ecosystems, Synapse aims to enable frictionless asset transfers, swaps, and generalized messaging with a fluent cross-chain functionality.

Secured through multi-party computation (MPC) validators operating with threshold signature schemes (TSS), the entire process is leaderless and maintains security by each parallel-task validator mirroring a process upon receiving on-chain events.

Synapse

Helpful Links

- Code
- Documentation

Celer cBridge

Celer's cBridge is a multi-chain network that lets users transer value between Layer-1 blockchains and different Layer-2 scaling solutions on top, *e.g.*, Ethereum to optimistic rollups, Polkadot to zk-rollups, or even Ethereum mainnet to Skale sidechain.

Helpful Links

- Code
- Documentation

$109.62M

Total value locked in Celer cBridges

235k

Total transaction count on cBridges

$1.97B

Total bridged volume across all cBridges

Source: cBridge-analytics

Source: cBridge-analytics

deBridge

deBridge is an interoperable liquidity transfer protocol that allows decentralized data transfer between various blockchains. Their smart contract cross-chain communication is powered by a network of independent oracles (validators) that are elected by the network's governance committee.

Helpful Links

- Code
- Documentation

The protocol enables transfers of assets between various blockchains via locking/unlocking of the asset on the native chain and issuing/burning the wrapped asset (deAsset) on secondary chains or L2s. Cross-chain communication between different blockchains is maintained by elected validators who run the deBridge node to perform validation of cross-chain transactions that pass between smart contracts of the deBridge protocol in different blockchain

dAMM (formerly Caspian)

dAMM is a decentralized, cross-L2 automated market maker (AMM) design developed jointly by Loopring and StarkWare. As a solution, dAMM allows liquidity to be bridged on L2 while remaining unfragmented on L1. dAMM functions with an off-chain operator that mimics the AMM contract logic by offered trade quotes offered and subsequent L2 trades at the beginning of a batch.

dAMM enables:

- ZK-based Layer 2s to share liquidity asyncronously (*e.g.*, DeversiFi, Loopring ...); this subjecs LPs to more trades
- Uncompromising scalability; LPs to simultaneously serve L1 AMM while partaking in L2 trading
- dAMMs utilize a permissionless L1 to mitigate against impermanent liquidity loss due to incompatable L2s

dAMM Architecture

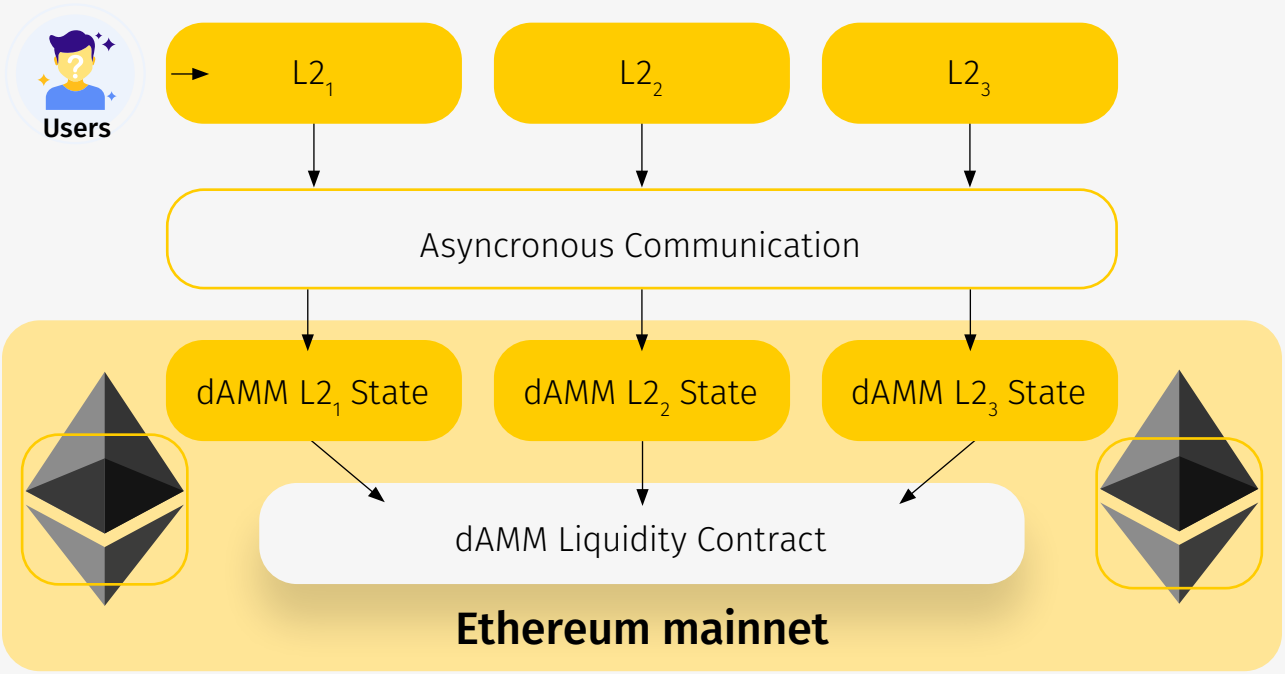

a. Optimistic vs. ZK-Rollups

Property	Optimistic Rollups	ZK-Rollups
Fixed gas cost per batch	~40,000 (a lightweight transaction that mainly just changes the value of the state root)	~500,000 (verification of a ZK-SNARK is quite computationally intensive)
Withdrawal period	~1 week (withdrawals need to be delayed to give time for someone to publish a fraud proof and cancel the withdrawal if it is fraudulent)	Very fast (just wait for the next batch)
Complexity of technology	Low	High (ZK-SNARKs are very new and mathematically-complex technology)
Generalizability	Easier	Harder (ZK-SNARK proving general-purpose EVM execution is much harder than proving simple computations, though there are efforts, working to improve on this, *e.g.*, Cairo.)
Per-transaction on-chain gas costs	Higher	Lower (if data in a transaction is only used to verify, and not to cause state changes, then this data can be left out, whereas in an optimistic rollup it would need to be published in case it needs to be checked in a fraud proof)
Off-chain computation costs	Lower	Higher (ZK-SNARK proving especially for general-purpose computation can be expensive, potentially many thousands of times more expensive than running the computation directly)

Source: vitalik.ca

b. Rollups You Can Try Now

The zk-rollup ecosystem is nascent but growing with multiple companies working on several implementations. Some prominent companies include StarkWare, Matter Labs, Hermez, and Aztec.

Arbitrum
Optimistic rollup

Optimism
Optimistic rollup

Loopring
ZK-rollup

Polygon Hermez
ZK-rollup

StarkEx
ZK-rollup

StarkNet
ZK-rollup

ImmutableX
ZK-rollup

dYdX
ZK-rollup

zkSync
ZK-rollup

Development

Value locked is generally regarded as an indicator to evaluate the level of public adoption. Heading into 2022, rollups remain very new but are primed for increased adoption.

+$1.8 B

Total value locked in zk-rollups (ZKRUs)

+$4.68 B

Total value locked in optimistic rollups (ORUs)

II. Sustainability

Climate change ranks near the top of most governmental and corporate agendas. Accordingly, large-scale human endeavors which consume large amounts of electricity or otherwise create emissions will be scrutinized by activists, investors, and regulators. This is the case regarding the energy-intensive verification process by which transactions are recorded and blocks produced onto the Bitcoin and Ethereum blockchains. This process (called "proof-of-work," discussed more below) is indisputably energy-intensive (but not necessarily bad!) and has therefore been criticized as power-hungry, unsustainable, and not green enough for today's world.

In response to this criticism, as part of the ongoing Ethereum upgrade, Ethereum developers are implementing a new verification process called "proof-of-stake." Researchers at the Ethereum Foundation estimate that the drop in electric power usage by the Ethereum network will be as much as 99%! If that is the case, the move to proof-of-stake should nullify this criticism from Ethereum and may position Ethereum as a greener and more sustainable Bitcoin as public awareness of both blockchains grow.

Proof-of-Work

A decentralized consensus mechanism using mathematical puzzles to validate block creation, confirm transactions, and mine tokens to increase the asset's supply. A new block is created when a miner discovers the randomly-generated target hash. A group of transactions are then confirmed to the new block. The miner is given a block reward in exchange for the private computing power used to solve the puzzle. The energy resources required to solve puzzles deter malicious transactions and spam.

1. Proof-of-Work (PoW)

A bit of history and background is helpful to understand the coming proof-of-stake sustainability upgrade. At launch (and currently), Ethereum, like Bitcoin, relies on the proof-of-work verification process where transactions are completed and made permanent on the blockchain only when they are confirmed by computers in the network by solving intricate math problems. In its simplest form, traditional proof-of-work requires miners to spend electricity in order to guess the hash of the previous block. The first miner to correctly guess/find this hash gets to submit the next block to the chain and receives bitcoin as rewards for its effort.

Miners are the international network of computers that:

- Bundle bitcoin transactions into the blocks (if the block gets filled, the remaining transactions will be added to the next block)
- Solve a cryptographic puzzle (the "proof-of-work")
- Send the blocks out over the network to be cross-checked and verified OR they will check other's incoming blocks for accuracy
- Propagate approved blocks across the network to let other nodes know and move on to the block of transactions

Miners are extremely important to the health of the network and the idea to include the process of PoW mining was one of Satoshi's key innovative ideas. In short, they are responsible for new Bitcoin block generation and adding blocks to the blockchain but beyond that mining aids in:

- Securing the network and preventing corruption from malicious actors
- Minting new bitcoin into circulation in a predictable, predetermined manner
- Maintaining a historical record so that the chain remains auditable and transparent allowing global consensus to be reached

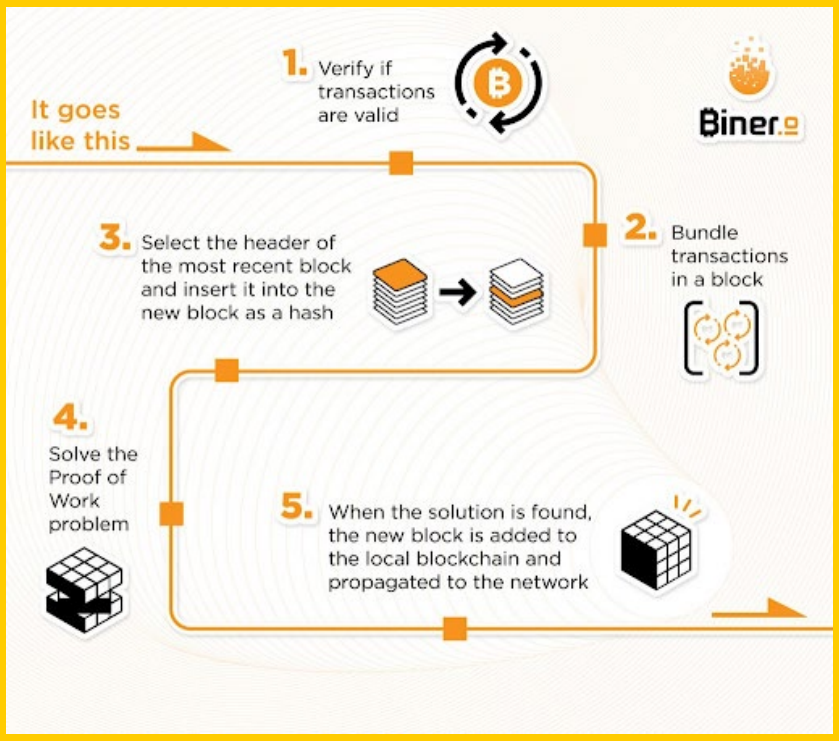

Mining Explained **Source**: Biner

The computers which solve these problems are rewarded with that blockchain's cryptocurrency and, for that reason, they are called "miners." "Mining" in this context is both an illuminating and misleading name. It's helpful because, just as in geologically mining the earth to extract valuable commodities, energy is expended, work is done, and a miner obtains the sought-after reward. But mining is also a misleading term in the blockchain context—a gold miner's energy expenditure and work does not relate to other previously completed mining operations or those still going on elsewhere in the world. If a miner mines a pound of gold, the work expended relates only to the gold sought. It does not in any way confirm or validate previous gold removed from the ground by other miners.

But in Bitcoin and Ethereum mining, this is exactly what is happening. Mining is the verification and confirmation process which keeps the Ethereum and Bitcoin blockchain continually updated. And for Bitcoin, this is not going to change. Mining and proof-of-work makes Bitcoin what it is: A global, decentralized, immutable, and secure digital asset. A miner is rewarded only when transactions are confirmed as valid, and thus, entered into the universally identical Bitcoin ledger. Mining is—and will always be—at the very heart of the consensus and verification process for Bitcoin. Without proof-of-work, Bitcoin is not Bitcoin.

But from the very beginning, the creators of Ethereum always understood that, while they would launch with proof-of-work, they would eventually implement a different transaction verification process that uses less computational power and electricity. This is a fundamental philosophical and operational difference between Ethereum and Bitcoin.

It's important to note that Bitcoin miners are growing increasingly more vigilant of the environmental impact associated with their electricity usage and the miners' continual quest for cheaper electricity is leading to the development of renewable

energy resources like wind, solar, hydro, and geothermal power. Bitcoin thought-leaders also point out that proof-of-work converts energy into financial value and the cost to secure the network is what gives it security. If there was no cost, then there would be no security. Even so, readers will often see headlines about Bitcoin using as much energy as a small nation. Well, that's because Bitcoin literally is a small, digital nation, facilitating commerce and settlement for millions of people. It's doing so without borders or a government, but is a top-10 currency in the world.

Perspective is also important because even with its current energy usage, Bitcoin uses a small fraction of the energy required to mine gold, run the traditional financial system, or secure and protect the US dollar. The energy consumption from Bitcoin mining should be considered alongside the product output resulting from the energy consumption. Does this energy consumption improve human civilization (like air conditioning, personal computers, refrigerators, and cell phones)?

How one answers that question will determine how one views proof-of-work. It seems that the criticism of proof-of-work's energy consumption often comes from people who don't believe that Bitcoin is a valuable product. And that is a different, though related, debate.

Additionally, based on a 2020 report from the Bitcoin Mining Council and past reports from the likes of CoinShares and Cambridge Alternative Finance, Bitcoin mining is already "greener" than most industries and vastly more so than the US electrical grid. This is contrary to the many hyperbolic headlines concerning Bitcoin's energy usage and provides a new perspective on the mining industry. Take a look at the numbers from each report and ask yourself, "Why should Bitcoin mining, specifically, be condemned when so many other industries are less scrutinized?"

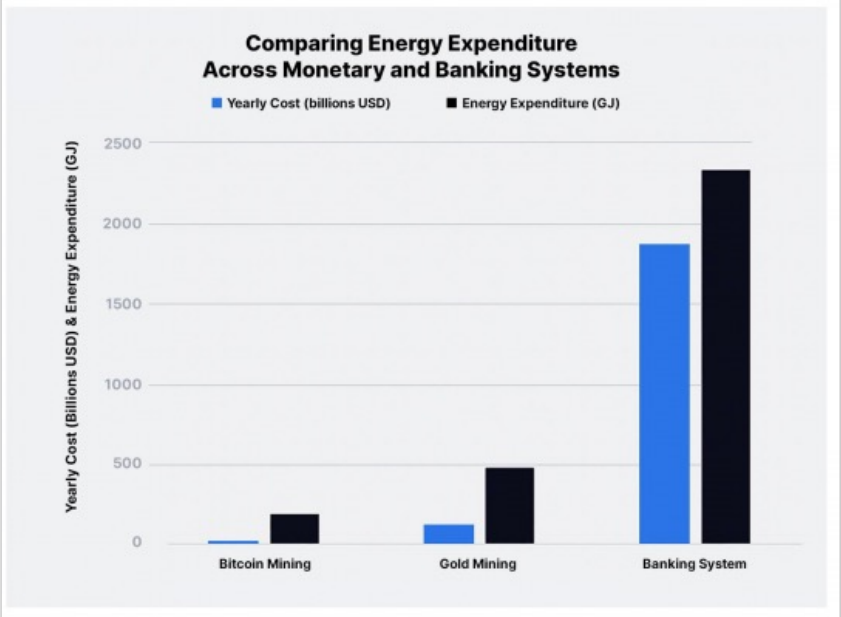

Comparing Energy Expenditure **Source**: Nasdaq

Proof-of-Stake

A decentralized consensus mechanism, proof-of-stake (PoS) uses network participation to validate block creation and confirm transactions. Participation in proof-of-stake requires nodes/network validators to commit a personal stake of the underlying asset to be locked for a stipulated time period, similarly to posting collateral. As opposed to proof-of-work, proof-of-stake does not require specialized mining hardware, only internet connectivity and a stake of the cryptocurrency.

2. Proof-of-Stake

In the future Ethereum blockchain (called "proof-of-stake") which is coming as part of the Ethereum upgrade, computers doing the confirmation work are called "validators" rather than "miners." Anyone is eligible to become

a validator after acquiring 32 ether (ETH). A validator puts those ETH at risk (or "stakes" them) as a guarantee of good behavior, as it were. Qualified validators (those that have "staked" their 32 ETH) are then chosen (pseudo) randomly to confirm transactions. Note, it is also possible to stake with less than 32 ETH through third-party pools and service providers which reduces the barrier to participating in the network and earning rewards.

In the staking model, there is no advantage to having more computational or electric power because validators are chosen randomly. Therefore, proof-of-stake eliminates the proof-of-work arms race for more electricity and computing power.

But what compels a proof-of-stake validator to do their job correctly? If a chosen validator erroneously confirms a transaction or colludes with other validators to confirm transactions falsely, their staked ETH will be taken ("slashed") and their validator reputation tarnished. If a validator confirms transactions correctly (along with other validators, until a consensus threshold is reached), they are then rewarded with more ETH. Good behavior rewarded, bad behavior decisively punished.

The Beacon Chain, which launched in December 2020, is the center of Ethereum's new PoS consensus mechanism. As the focal point of the PoS network, it's responsible for the liveness, veracity, and consensus of the Ethereum network. Future sharded layers (discussed previously) will all connect back to the Beacon Chain, beginning with just four shards and possibly growing to 1,000+ shards. The Beacon Chain will provide the foundation for hundreds of thousands of validators distributed across thousands of nodes globally. It'll organize validators into committees and apply the consensus rules that dictate the network.

How will all of this play out? If the dramatic drop in energy consumption emerges with proof-of-stake, then Ethereum should be immune from a criticism that will likely continue to be leveled at Bitcoin and its proof-of-work system.

3. Sustainable Scaling and Growth

Blockchains like Bitcoin and Ethereum strive for maximum decentralization and censorship-resistance while remaining totally open and inclusive networks. However, they also want to scale to accommodate billions of users. As they stand right now, their limited capacity to process transactions at the base layer (~7 and ~20 TPS, respectively) are in direct opposition to achieving that goal.

The question is "What is the best method of scaling a blockchain?" Nearly every new "next generation" blockchain since 2016 boasts sky-high transactions per second (TPS) as a selling point. However, the issue that persists is that TPS is not the sole metric in which to compare blockchain scaling. Generally, the truth is that the higher the TPS, the higher the cost (financially and computationally) to run the network. Given this, the question arises: Are these new "next generation" blockchains actually scaling, or just simply increasing TPS while shrinking the network in other regards?

The primary means to accomplish sustainable scaling are minimizing the hardware requirements needed to participate in the network and, also, ensuring the state of the network (data) does not balloon to unsustainable levels.

Network nodes are what enforce the rules of the chain and ensure noone is cheating the system. Therefore, having a robust, geographically-dispersed, and anti-fragile network of nodes is ideal for the decentralization and security of the network. In order to attain this system, the costs to run a node (hardware, bandwidth, energy, and storage) should be as little as possible. This allows the greatest number of people the option to join the network, if they so choose. Keeping costs low ensures no one is priced out and your network is not solely controlled by a wealthy, elite few.

The other variable to consider is state growth, *i.e.*, how quickly the blockchain's computational load grows. Full nodes store the network's entire history from genesis and must be able to validate the entirety of the network's state. Blockchains that scale by simply increasing the blockspace and throughput per unit of time (Binance Smart Chain and EOS), also greatly increase their state growth. Those chains are short-term solutions that lead to long-term unsustainable networks.

Blockchains like Solana, which are designed for greater TPS via specialized hardware, also run into state growth and centralization issues. To be fair, Solana did introduce some new technological innovations to improve sequencing like proof-of-history and a parallel execution environment. However, like the "Ethereum killers" of the 2017 era, this design is not long-term scalable/sustainable. Solana already boasts some of the most expensive and specialized hardware requirements of any top 20 cryptocurrency, and as Solana transactions and price increase, the hardware costs to run a node, be a validator, and process transactions also increases.

Hardware requirements:[1]

- Bitcoin: 350GB HDD disk space, 5 Mbit/s connection, 1GB RAM, CPU >1 Ghz. Number of nodes: ~10,000
- Ethereum: 500GB+ SSD disk space, 25 Mbit/s connection, 4–8GB RAM, CPU 2–4 cores. Number of nodes: ~6,000
- Solana: 1.5TB+ SSD disk space, 300 Mbit/s connection, 128GB RAM CPU 12+ cores. Number of nodes: ~1,200

Below is empirical data experienced by cryptocurrency and cybersecurity expert, Jameson Lopp, from a 2020 Bitcoin Node and 2021 Node Sync Tests. The table compares the time it takes to sync a full node of Bitcoin vs. Ethereum vs. Solana on an average consumer-grade PC.

In the Ethereum ecosystem, serving as a validator on the Beacon Chain requires staking 32 ETH (~$120,000 in Q4 2021). While this sounds quite expensive and exclusionary on the surface, relative to other chains top blockchains like Bitcoin, Avalanche, Solana, Binance Smart Chain, Ripple, and others, it removes the economy of scale that exists in PoW chains and with liquid

1 Requirements as stated in StarkWare's article "Redifining Scalability"

Blockchain	Throughput MB/hour	Lopp's node-sync time
Bitcoin	~6MB/hr	0 days : 5 hours – 3 days : 11 hours
Ethereum	~20MB/hr	2 days : 16 hours – 10 days : 2 hours
Solana	~2880MB/hr	not feasible to sync the full network state

staking services like Lido and RocketPool, users can participate with less than 32 ETH. Additionally, by removing hash power with randomness and a capped gas limit/block size, Ethereum enables any user with average hardware to profitably run an Ethereum validator.

Ethereum's state growth situation is also better than most chains (thanks to its lower gas limit) but could become problematic given enough time. As time passes and Ethereum adoption increases, the state grows in size and complexity. This ultimately increases the total time it takes for a full node to sync and the hardware requirements needed to run one. Fortunately, Ethereum has been designed to scale with rollups (discussed previously) which help reduce this state growth issue. As discussed at length, rollups handle enormous amounts of computation and transactions off-chain while only submitting a tiny "fingerprint" (proof) to the mainnet. This, coupled with sharding, enables exponential room for growth in a sustainable manner.

III. Security

If you talk to a cryptocurrency skeptic, one of the criticisms you hear may go something like this: "Your magic, internet money is fine until a twelve-year-old with a computer hacks you and takes it." In an era of increasingly sophisticated, destructive, and frequent cyberattacks, such a skeptic is on to something. Even so, this criticism fundamentally misunderstands the nature of blockchain-based digital assets like Bitcoin and Ethereum.

In the case of self-custodied assets, it's impossible to hack a single computer or server and steal these digital assets in the same way that your personal computer connected to the internet can be hacked and sensitive documents copied and removed. This exact scenario cannot happen because your digital assets are not stored in a single place. They are entries on the permanent digital ledger, continually updated on every computer participating in that blockchain's network.

Your Bitcoin is an entry on the Bitcoin blockchain; your Ethereum is an entry on the Ethereum blockchain. This is what cryptocurrency aficionados mean when they say digital assets have no single point of failure and that the blockchain is decentralized.

Blockchains are inherently more secure than traditional server and cloud-based computer architecture. But that doesn't mean that they are entirely invulnerable.

1. The 51% Attack

Public blockchains like Bitcoin and Ethereum are vulnerable to what is memorably called a 51% attack. If a party controls more than 50% of the computers doing the work of securing and updating the blockchain, then they could not only halt future transactions but could actually reverse old transactions. Basically, they could rework the blockchain and steal vast amounts of Bitcoin and Ethereum.

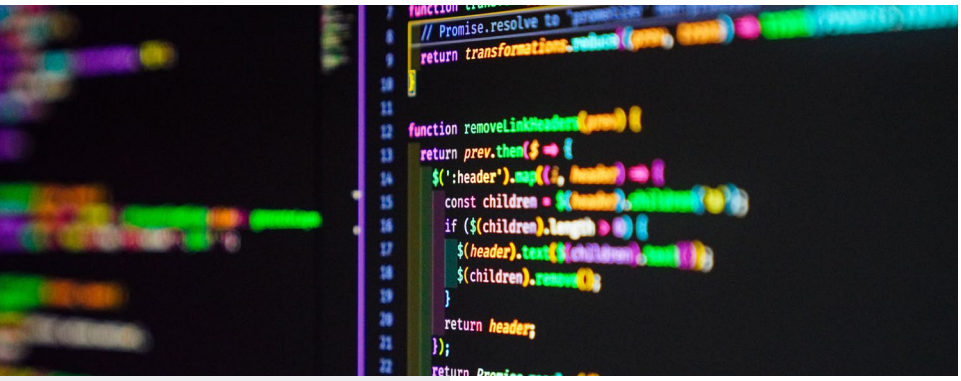

As mentioned in the article on Sustainability, the coming Ethereum upgrade will move from aproof-of-work to a proof-of-stake validation mechanism. Some Ethereum thought-leaders believe that proof-of-stake will be more secure than proof-of-work and have hypothesized that the Bitcoin blockchain is vulnerable to a state-sponsored 51% attack.

Interestingly, this scenario seems less likely now that China has shut down Bitcoin mining within its borders and that mining has been replaced by operations elsewhere. Even so, a 51% attack in a proof-of-work context is an interesting possibility and the proof-of-stake verification mechanism in the Ethereum upgrade does change the playing field in this regard by introducing a powerful disincentive.

In proof-of-stake, every validator puts their Ethereum tokens at risk in order to guarantee their good behavior as they perform the work of validating transactions on the Ethereum blockchain. So there is a way to punish bad actors. This disincentive does not exist in proof-of-work. Proof-of-Work only has carrots (mining rewards), not sticks (losing staked tokens).

Moreover, in proof-of-stake, validators will be randomly assigned transactions to confirm. This randomness eliminates much of the opportunity for validators to coordinate a planned attack. Imagine a game where four teams with different colored jerseys—red, yellow, blue, or green - play a game of elimination. Imagine also that a fundamental rule of the game is that each team gets its randomly-assigned jersey color just as it enters the field. The randomness of the jersey color assignment makes it very difficult for red, yellow, and blue to coordinate a unified attack on green.

And this cycle will continue. As the price of Ethereum rises, more validators will be incentivized to stake because the rewards are becoming increasingly valuable. More staked ether and more validators increase security still more and the upwards trend in price, number of validators, and security will continue.

51% Attack

A hostile offensive (attack) by one or more miners to gain control of blockchain transactions by controlling at least 51% of the hash rate, or computing power.

2. Lower Barrier of Entry

32 ETH

Full Validator Requirement:
To become a full validator, you'll need to operate a mainnet client and stake 32 ETH.

A validator is the equivalent to a miner for a proof-of-stake network. Validators collect transactions into blocks to add to the blockchain and are rewarded for adding valid blocks in proportion to the amount of currency they post ("stake") as collateral.

Other security threats might be reduced by proof-of-stake, as well. In theory, a large physical concentration of miners drawing from the same power source and doing proof-of-work could become a highly visible target for increasingly innovative cyberattacks and direct physical attacks.

At the minimum, as the price of Bitcoin climbs, mining operations in extensive facilities filled with noisy computers doing mining work will need more physical security (electronic, human, infrastructure). But proof-of-stake validation, by eliminating the arms race for more electricity and computing power, does not encourage physically concentrating computing resources but rather encourages a more physically-decentralized network.
This trend toward greater decentralization of validators should only accelerate because the only barrier to entry to proof-of-stake validation is the ownership of 32 Ethereum (by no means an insignificant investment if one is starting with nothing). The operating expenses (power, hardware, infrastructure, cooling) for proof-of-stake validation are significantly less than proof-of-work. In addition, since the computers doing the validating are less specialized, supply chain cost/delay risk is comparatively reduced for starting an Ethereum validation operation. Simply put, it will be easier to become a proof-of-stake validator than a proof-of-work miner.

It seems that with a comparatively lower barrier to entry, proof-of-stake incentivizes the continual growth of a geographically-dispersed and almost invisible army of validators, further strengthening and decentralizing the network.

None of this is to say that proof-of-stake or proof-of-work is more or less secure than the other. They are just different. Eventually, asking the question 'Which validation method is more secure?' is like asking the question 'Is my house more secure than Fort Knox? Or vice versa?' Or 'Is a cruise line or an airline more secure?' Answering those questions meaningfully needs much more context.

Every system and asset has security vulnerabilities. Many enterprises share some vulnerabilities, but some are also unique. Different techniques and assets have different functions and thus different vulnerabilities and security systems. So security comparisons between assets and systems are always only partial, and there is no one-size-fits-all solution. Moreover, nothing in life is completely secure, including Ethereum and Bitcoin.

Coinsider

Welcome to Coinsider! (What To Expect)

Aug 25, 2021 by Coinsider

Hey you, thanks for stopping by! We're Coinsider, your source for unique, insightful, and thought-provoking content in the crypto world. Check out our trailer for what we have in store for you and if you like what you see then hit that subscribe button & that bell icon to catch our future videos!

Check out our new website for important links: https://joincoinsider.com

Ethereum 2.0 Roadmap- What's Next For Investors?

Aug 30, 2021 by Coinsider

We're well on our way on the road to Ethereum 2.0, but it's still early. We just finished Phase 0 which was the Beacon Chain launch. We now have some other small upgrades post London Hard Fork (EIP 1559) and should have the Merge by late 2021 or early 2022 which finally makes Ethereum Proof of Stake. But then what comes after? And also how do all of these upgrades affect the price of ETH? In this video I'll explain and explore this important topic. So if you ...

Ethereum + ZK Rollups = DOMINATION!

Nov 23, 2021 by Coinsider

A lot of people have been railing on Ethereum lately because of its high gas fees and lack of scalability. But what if I told you that is all about to change? Because of a revolutionary tech innovation called ZK rollups? Well in this video I'm gonna share with you everything you need to know about ZK rollups, why they will help Ethereum WIN in the mid to long term, and how YOU can get a piece of this groundbreaking tech. It's a rollup-centric future and I'm all here for it!

Is Polygon's $MATIC Worth The HYPE?! Pros & Cons

Jun 9, 2021 by Coinsider

Polygon and their MATIC token has literally done a 100x in the span of a few months. But is it really Worth the Hype though? As one of the hottest Layer 2 solutions that's operational and live for Ethereum, Polygon has taken over the space by storm, onboarding a ton of top DeFi and NFT projects, and ...

Ethereum 2.0, A Deep Dive!

Oct 21, 2020 by Coinsider

What's the latest on Ethereum 2.0 and when will it finally arrive? With DeFi causing the ETH network to grind to a halt with massive congestion and high fees like back in the ICO days, the scalability issues of the main layer1 blockchain has yet again be thrusted into the forefront. Ethereum devs and the community is ...

X. About CRYPTOEQ

Mavericks & Thought-leaders

CryptoEQ™ is an independent cryptocurrency analysis and rating agency that provides unbiased, objective, and transparent research you can trust. We help people navigate their investment journey and trading decisions.

Spencer Randall
Principal & Co-Founder

- 8+ Years in System Architecture/ Implementation
- 5+ Years in Crypto Trading/Investing
- Bachelor of Science in Engineering

Brooks Vaughan
Head of Innovation & Co-Founder

- 17+ Years in Product Design/Management
- 8+ Years in Crypto Trading/Investing
- Bachelor of Industrial Design

Michael Thoma
Lead Analyst & Co-Founder

- 11+ Years in Technical Research/Analysis
- 5+ Years in Crypto Trading/Investing
- Master of Science in Geology

Company Statistics

+85%
Algorithm Win Rate

+175%
2020 Average CORE Report ROI

+300%
Q/Q Revenue Growth

+50,000
Total Active Users

We help you gain the market insights you need to refine your investing and trading strategies and efficiently manage your exposure across a variety of digital assets.

Our objective is to provide trusted information and analysis for quickly evolving blockchain technologies and make navigating cryptocurrency less intimidating for new investors. Our research and analysis encourage you to make smart decisions for both long term investments and short term trading strategies.

Our proprietary algorithms, exhaustive research and helpful community are key to our success as we follow strict principles and ethics to deliver honest information. We actively seek to identify scams and low quality nefarious projects relieving you of that burden.

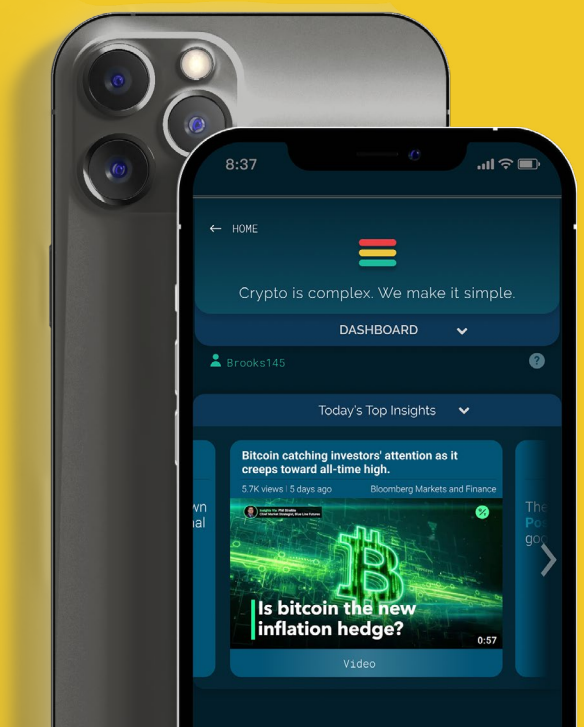

Platform Highlights

v1 launched in 2019

- » Signal over noise.
- » Direction over data.
- » Quality over quantity.
- » Usability over complexity.

XI. Final Words
Our Story

Like most disruptive tech startups, Crypto**EQ** started as a small group of like-minded individuals.

Each of the co-founders—Spencer Randall, Michael Thoma and Brooks Vaughan—was a cryptocurrency investor and trader before the crypto explosion of 2017. They met one another attending local crypto conferences and immediately began to admire their different perspectives and maverick approach to the assets available on the market. After some time getting to see each other in action, they each noticed a glaring hole in the crypto-asset market —truly unbiased, thorough insights and research.

We launched CryptoEQ v1 in July 2019 and acquired approximately 3,000 new users. Meeting our goal to be constantly launching, CryptoEQ v2 debuted in January 2020 with new features and an all-new sales funnel. Our third iteration, v3, launched in June 2020 with average quarterly revenue growth of over 300%. We also blew through our 5,000-user milestone. Recently, our v5 launch incorporated a new and intuitive user interface and exclusive one-on-one consulting sessions, pushing us past 50,000 users. And currently, we‘re tracking to exceed our next goal of 75,000 users within Q1 of 2022.

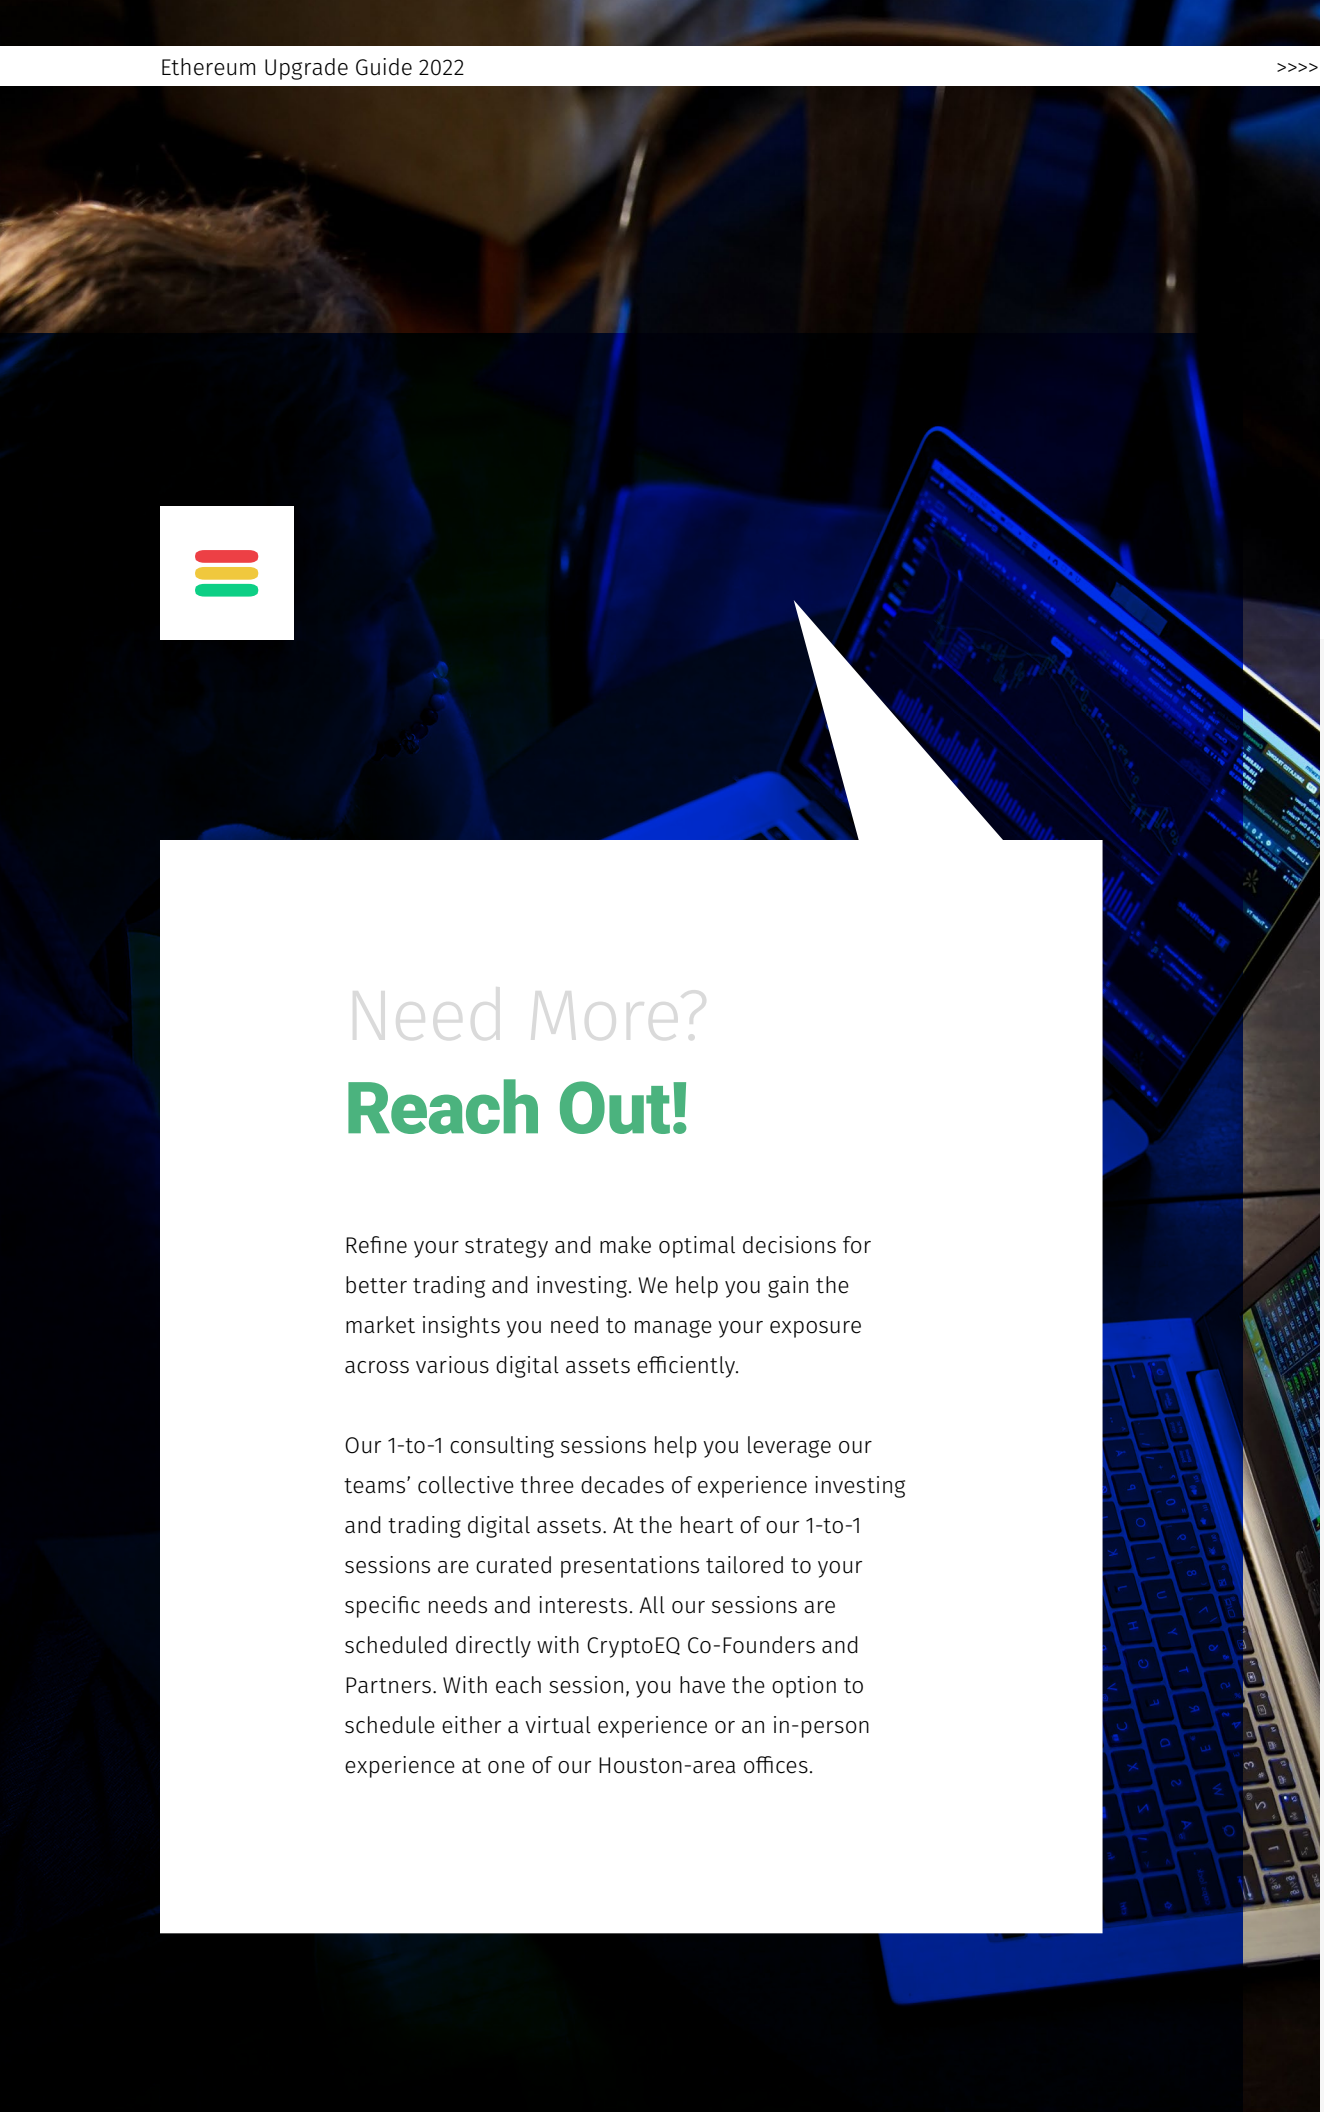

Need More?

Reach Out!

Refine your strategy and make optimal decisions for better trading and investing. We help you gain the market insights you need to manage your exposure across various digital assets efficiently.

Our 1-to-1 consulting sessions help you leverage our teams' collective three decades of experience investing and trading digital assets. At the heart of our 1-to-1 sessions are curated presentations tailored to your specific needs and interests. All our sessions are scheduled directly with CryptoEQ Co-Founders and Partners. With each session, you have the option to schedule either a virtual experience or an in-person experience at one of our Houston-area offices.

Coinsider

CRYPTOEQ

Ethereum Upgrade Guide 2022

CryptoEQ CORE+ Series.