Tech Explainer

SCALING WEB3: LAYER 2 SOLUTIONS

INTRODUCTION

When we scroll through posts on social media, few of us think of the transaction costs involved in uploading a post and allowing others to see and interact with it. We simply expect seamless and instantaneous interactivity. This is in large part because the transaction costs associated with interacting with web2 interfaces are currently negligible, while transaction speeds are kept high at all times. And yet, every action performed by every user in web2 entails a transaction that must be executed, the same way that every action on web3, built on blockchain technology, does. It just so happens that the transactions in web2 are executed unilaterally by companies and the costs are thus kept steadily low—and out of sight of users. However, blockchain infrastructure is run, and transactions on it are executed and validated, by a decentralized group of participants, and the costs associated therewith are passed directly on to users.

As a general purpose blockchain, Ethereum hosts a panoply of innovative projects. However, Ethereum's user experience can be slow and expensive. The blockchain currently processes around 15–20 transactions per second.[1] Monetary transaction costs borne by users, or *gas fees*, for each transaction on Ethereum rose to more than $40 on average during a bull market in November 2021 as users raced to outbid each other.[2] Certain applications running on blockchain, such as social media, gaming and micro-finance, can only operate smoothly in an environment featuring high transaction speeds and low gas fees because of the number of individual transactions involved in participant interactions on those applications.[3] Indeed, some nonfinancial applications on Ethereum have been forced to shut down, in part due to high gas fees that make it impracticable for users to engage with them.[4]

The wide range of projects on Ethereum also means that scalability is desperately needed to advance widespread adoption of web3. Over the years, to allow an exponentially larger scale of transactions to support these projects, the blockchain community has put forth creative and highly effective solutions to the problem of high transaction costs and low transaction speeds, while maintaining the security and decentralization features that are pivotal to blockchain. "Layer 2" (or "L2") solutions[5] represent one family of solutions to the scalability problem. L2 simply refers to separate blockchains built, or layered,

---

1 *Layer 2*, ETHEREUM, https://ethereum.org/en/layer-2/;
Plasma Chains, ETHEREUM, https://ethereum.org/en/developers/docs/scaling/plasma.

2 Jared Ronis, *Understanding Ethereum's Layer 1 and Layer 2: Differences, Adoption, and Drawbacks*, Wilson Center (Oct. 13, 2023), https://www.wilsoncenter.org/article/understanding-ethereums-layer-1-and-layer-2-differences-adoption-and-drawbacks.

3 Some argue that a main reason we have not yet seen nonfinancial applications take off on blockchain is due to high transaction fees. *See, e.g., Ethereum Has Blobs. Where Do We Go From Here?*, Mar. 28, 2024, https://vitalik.eth.limo/general/2024/03/28/blobs.html.

4 *See, e.g.*, Alex Van de Sande, *Out of Gas: We're Shutting Down UniLogin*, MEDIUM (Sept, 18, 2020), https://medium.com/universal-ethereum/out-of-gas-were-shutting-down-unilogin-3b544838df1a.

5 We define L2 broadly to refer to all scaling solutions that take place outside of L1 and where the bulk of the computation occurs outside of L1.

"on top of" the original, or "Layer 1" (or "L1"), blockchain, usually Ethereum.[6] As of the date hereof, there are more than a hundred L2 projects on Ethereum,[7] and the peak weekly number of distinct addresses interacting with one or more L2 projects had reached more than ten million.[8] Different types of L2 solutions provide different advantages, but all of them regularly "post" information about transactions that occur on L2 to the L1 blockchain and rely upon L1 to ensure some degree of security. This Tech Explainer will go over how L2 solutions actually work, and provide an overview of some of the different types of L2s.

LAYER 1 BLOCKCHAIN

Before we dive into L2, it is essential to understand why scalability is so hard to achieve. The two critical features of blockchain that make it a unique technological solution are decentralization (allowing individuals with basic technological setup to participate in the network) and security (making it difficult for attackers to alter information on the network). Layer 1 blockchains such as Ethereum achieve both of these. But what happens when a third feature, scalability (high transaction speeds and low transaction costs) is desired? Any solution that increases scalability will compromise one or both of decentralization and security. This problem is known as the "blockchain trilemma".[9] Layer 2 solutions are, in essence, creative attempts to increase scalability while minimizing decreases in decentralization and security.

Recall that a blockchain is a collection of computers, or *nodes*, that contains a copy of information about all transactions—or actions that take place on the blockchain—such that each individual computer can verify the validity of transactions without having to rely on a centralized authority. In a blockchain, a *transaction* is a cryptographically signed instruction sent by a user that contains information about a change of state they want to record. For financial applications, the simplest transaction is transferring cryptocurrency from one account to another (recorded as a decrease in the amount of cryptocurrency in one account and a corresponding increase in the amount of cryptocurrency in the other). *Blocks* are batches of transactions that are linked to the previous batch, or block, to ensure that the transaction information recorded in the blockchain cannot be altered without alerting the network, a feature known as "immutability". In addition to the new transactions being recorded, each block contains a reference to the block that came immediately before it, so the blocks are thus linked together to create a *blockchain*.

Blockchains each have a *consensus mechanism*, which is a protocol through which a decentralized network of validators determine whether a transaction is valid and should be recorded. In an L1 blockchain, security is ensured through a consensus mechanism that requires that each block—which records multiple new transactions—be processed by every full node in the network.[10] Ethereum has over 6,000

---

6 Bitcoin and Ethereum are both Layer 1 blockchains. However, Layer 2 projects are generally built on top of Ethereum, rather than Bitcoin, since Ethereum is a general purpose platform amenable to diverse ecosystem of projects.

7 L2BEAT, https://l2beat.com/scaling/summary.

8 GROW THE PIE, *Mastering Ethereum Layer 2s*, https://www.growthepie.xyz.

9 The term "blockchain trilemma" was first coined by Vitalik Buterin, the co-founder of Ethereum. Decypher Media, *The Data Availability Problem – Vitalik Buterin | Silicon Valley Ethereum Meetup* at 5:57, YouTube, https://www.youtube.com/watch?v=OJT _ fR7wexw.

10 Ethereum is working to provide *light nodes*, which are nodes that do not keep local copies of blockchain data, but instead request necessary data from a provider. A light node allows more individuals to participate in the blockchain and would also make rollups, discussed further down, more secure. *Light Clients*, ETHEREUM (Mar. 5, 2024), https://ethereum.org/en/developers/docs/nodes-and-clients/light-clients.

full nodes in its network,[11] and each node must download the entire transaction data from each block and verify the proposed new transactions included in that block. The new block is added only after two-thirds of the nodes agree that the transactions in the block are legitimate. Thus, every full node needs to download, verify and locally store transactions from each block, which means that any single node must have the capacity to execute all transactions.[12]

Individuals can become validators simply by downloading software onto their computer and staying online. In essence, the Ethereum consensus mechanism allows—and encourages—as many users as possible to independently participate in verification. In that way, no user needs to trust anyone else regarding the order and validity of transactions.[13] The Ethereum consensus mechanism boosts decentralization and security, as each full, active node is working to verify each transaction, and incentives are set to encourage as many users as possible to become validators. Ultimately, however, this protocol is not very scalable, as downloading and storing data to execute each transaction puts too much strain on each individual node, causing delays for the entire network. Lastly, it is important to understand that Ethereum operates three layers: *consensus* (nodes coming to an agreement regarding the inclusion and order of blocks), *execution* (the verification of blocks in real time) and *data availability* (the available record of past transactions).[14]

LAYER 2 BLOCKCHAIN

In an attempt to provide increased scalability while reducing compromises to security and decentralization, L2 blockchains take on some, but not all, functions of L1. Most L2s only perform the execution function—they validate individual transaction information and batch transactions. This means that they depend upon L1 for consensus and data availability, which is why they are blockchains layered over an L1 blockchain rather than independent blockchains. L2 is much faster and cheaper at processing transactions than L1 because L2 only focuses on executing transactions, instead of also coming to a consensus about the order and content of transactions and ensuring data from each execution is available on-chain.

After operators in L2 have executed multiple transactions, they batch the transactions together, compress the transaction data and then submit them to a *smart contract* in L1 for verification and posting. How does "batching" work? An L2 uses a cryptographic method called the *Merkle root* to compress information from many transactions into a single representation while still allowing users to verify that any particular small piece of data is part of the larger data set.[15] In other words, the Merkle root provides a way to prove that the list of transactions in the batch is valid, without needing to store all information about the entire batch. Ultimately, a large number of transactions on L2 will appear as a single transaction on L1, with only one transaction's worth of delay and L1 gas fees.

---

11 Etherscan, https://etherscan.io/nodetracker. While running node hardware is necessary for home staking, individuals without the requisite hardware who want to become a validator can also participate in "staking as a service" or "pooled staking." *Home Stake Your ETH*, Ethereum, https://ethereum.org/en/staking/solo/.

12 Ethereum's limit on block size—the number of transactions processed per block—is also a decentralizing force, as it caps the computing power needed to process a block.

13 "Don't trust, verify" is a central maxim of blockchain.

14 Importantly, data availability is not data storage. Data availability means making sure that data was available when it is published and available to be downloaded by everyone, while data storage stores data permanently.

15 The Merkle root is used extensively in L1 as well as L2 to promote efficiency. Other methods that build upon principles using the Merkle root have proliferated, such as the "Merkle-Patricia Trie" and future plans for a "Verkle Tree". *Merkle Patricia Trie*, https://ethereum.org/en/developers/docs/data-structures-and-encoding/patricia-merkle-trie.

Because L2 transactions are being recorded to L1, there must be a connection between the two layers of blockchain. These connections take the form of smart contracts on both L1 and L2. The smart contract (which automatically executes computer code once predetermined conditions are fulfilled) on L2 submits transaction information to a smart contract on L1. The smart contract on L1 verifies that the transactions submitted (in the form of the single representation of multiple transactions) are valid and, once it does, the transaction is finalized. This means that any modifications to an L2 transaction after finalization would require changing L1, which requires a large number of nodes in L1 to reach consensus.[16] This reliance on L1's consensus protocol provides security that takes effect once a batch of transactions is recorded onto L1. Depending on the type of L2, it is possible that the type of information posted onto L1, while compressed, still preserves data availability. Data availability, where every user can independently verify transactions if it so wishes, is crucial to decentralization and security.

It is important to note that *prior to the batched transactions being recorded onto L1*, the transactions executed on L2 do not benefit from the robust security that is the hallmark of L1. However, as discussed above, the decreased decentralization (and thus security) on L1 is what allows operations to be performed much faster or at lower costs, which is suitable for cases involving advanced interactivity such as social media and micro-decentralized finance.[17] And while L2s do not provide the exact same guarantees as L1, they can record transactions with sufficient finality and security to be useful.

It's important to note that there is a great deal of diversity in L2 blockchains and in the way each L2 chain handles consensus, execution and data availability. A number of different mechanisms have emerged for interactions between L1 and L2, where L2s are classified by the type of connection they have with L1. Some very popular examples of L2s include rollups—of which there are zero-knowledge rollups and optimistic rollups—validiums and plasma. Illustrated on the table below are three key differences between these types of L2s: computational burden, data availability and delay.

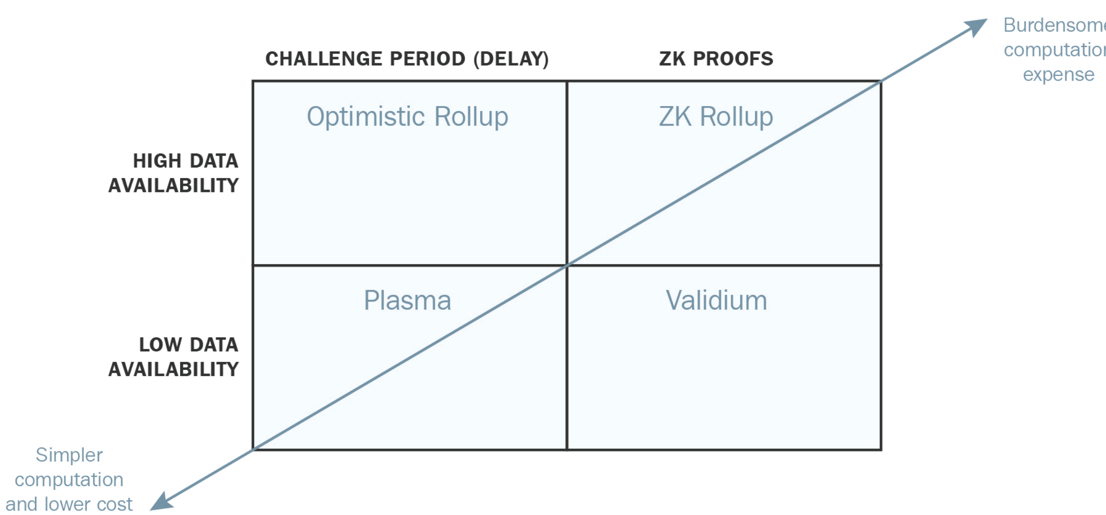

FIGURE 1

16 *Layer 2,* ETHEREUM, https://ethereum.org/en/layer-2.

17 Josh Stark, *Making Sense of Ethereum's Layer 2 Scaling Solutions: State Channels, Plasma, and Truebit*, MEDIUM (Feb. 12, 2018), https://medium.com/l4-media/making-sense-of-ethereums-layer-2-scaling-solutions-state-channels-plasma-and-truebit-22cb40dcc2f4.

As discussed further below, rollups post transaction data onto L1 and maintain high data availability, thus taking advantage of the security of L1's consensus mechanism. Plasma and validiums, on the other hand, post comparatively less data onto L1, providing only limited data availability, which means they are less secure than rollups. Optimistic rollups and plasma rely on a challenge period to ensure executed transactions are valid, which creates a significant delay between the initiation of a transaction and its final posting to L1, while zero-knowledge rollups and validiums provide instant finality through a computationally burdensome method called a "zero-knowledge proof". Posting more data onto L1 and generating zero-knowledge proofs both elevate the computational burden, and therefore the expense, for transactions on an L2.

These differences mean that each type of L2 has its own set of advantages and disadvantages that can be assessed depending on the type of L2 project being developed. Projects that prioritize community governance, such as decentralized autonomous organizations, for instance, would focus on decentralization more than security. On the other hand, applications with financial components, such as enabling trade of NFTs, would need to provide users with high degrees of security for their assets. Below is a brief explanation of how each of these mechanisms works to connect an L2 to L1.

Rollups

In a zero-knowledge rollup ("ZK rollup"), transactions are executed on L2. An L2 operator then submits a zero-knowledge proof ("ZK proof") along with the batch of transactions, to L1 for the L1 blockchain to verify whether the transactions are valid. A ZK proof is a cryptographic method that ensures a statement's correctness without revealing the statement itself, hence the name "zero-knowledge".[18] The inclusion of a ZK proof with every batch of transactions submitted to L1 allows for high security, since the ZK proof acts as a guarantee that the transactions executed on L2 are valid. However, generation of the proof is also computationally burdensome, a factor that has so far meant ZK rollups need to rely on centralized operators to generate and submit transactions to L1. The L2 operator is responsible for executing transactions, producing batches and submitting blocks to L1. In the next section, we discuss the challenge presented by having centralized operators, which includes the possibility of tampering with transaction records for their own benefit.

A close cousin of the ZK rollup is the *optimistic rollup*. Optimistic rollups execute transactions on L2 similarly to ZK rollups, but differ in how transactions are finalized on L1. Instead of using ZK proofs to show transactions are valid, optimistic rollups assume their validity—hence the name "optimistic". Optimisc rollups only post transaction data on L1, without an accompanying cryptographic proof. After a batch of transactions is submitted to L1, anyone can challenge the result of a rollup transaction by computing a *fraud proof* that reverifies the transaction during a *challenge period* (usually a week). If the batch is contested, it will not be finally recorded on L1. If the batch is uncontested after the challenge period, all transactions are deemed correctly executed and the transactions are finalized on L1. During the challenge period, assets subject to the transactions cannot be transferred out of the rollup because the transactions are not yet final. Thus, while in ZK rollups, funds subject to an L2 transaction are available for withdrawal as soon as a batch is submitted to L1, in optimistic rollups funds will be tied up until the

18 The ZK proof is much smaller than the data input, helping to save transaction costs, although it requires significant computational resources. Two types of proofs are ZK-SNARK (zero-knowledge succinct non-interactive argument of knowledge) and ZK-STARK (zero-knowledge scalable transparent argument of knowledge). For an explanation of zero-knowledge proofs, see, e.g., David J. Kappos, Sasha Rosenthal-Larrea, Carys J. Webb & Daniel M. Barabander, *Tech Explainers: Zero Knowledge Proofs*, CRAVATH, SWAINE & MOORE LLP (Oct. 2023).

challenge period elapses. The delay of the challenge period can be especially problematic in financial use cases where a user may want to have assets become immediately available on L1 for use on L1 or transfer to another L2.

On March 13, 2024, Ethereum launched the Dencun upgrade, a pivotal development for rollups by allowing them to become cheaper and more scalable. For rollups, data did not need to be available forever—it only needs to be stored for the time it requires to verify each transaction was valid. However, before the Dencun upgrade, rollup data was stored forever on L1 and subject to regular L1 gas fees. The Dencun upgrade introduced *proto-danksharding* (EIP-4844), which uses *data blobs* rather than blocks to contain rollup data that is deleted after approximately 18 days, which is sufficient for rollup purposes.[19] Importantly, the blob space market is decoupled with the gas fees market, so that rollups will pay (much lower) blob fees for making voluminous rollup data available, instead of gas fees.[20] Immediately after the Dencun upgrade, certain rollups saw a 99% reduction in overall transaction fees, from $6 to $0.04.[21] Proto-danksharding is an intermediate step to *full danksharding*,[22] which hopes to allow Ethereum to be able to record 100,000 L2 transactions per second and to make transactions on L2 as cheap as possible for users.[23]

Plasma, Validiums and Beyond

Unlike rollups, which provide excellent data availability, validiums and plasma are mechanisms for recording L2 transactions on L1 that do not provide data availability. Validiums, like ZK rollups, use ZK proofs and offer near-instant withdrawal of funds. However, validiums only publish the proof on L1, unlike ZK rollups that post both the proof and transaction data on L1. This means that validiums have decreased security compared to ZK rollups, since there is a chance, however small, that the ZK proofs are erroneous. Without transaction data available on L1 for anyone to verify, users will have to take the honesty of the L2 operator for granted. This can be especially problematic with centralized operators.[24] A chain controlled by a few parties might be able to influence the order of transactions, introducing perils such as fraud and censorship. In other words, a malicious centralized operator could post invalid states, thereby stealing funds from users, or censor certain users by refusing to include their transactions in a batch. At the same time, by not posting transaction data, validiums have a decreased computational burden, which means they are cheaper and faster than rollups. Further, decentralization in validiums is helped by the lower computational burden, since more users could potentially participate as operators.

Like optimistic rollups, plasma uses a challenge period that introduces delays of interacting with the L2. Despite large scalability gains, the lack of data availability resulting from plasma has raised security

---

19 Jon Charbonneau, *The Hitchhiker's Guide to Ethereum*, DELPHI DIGITAL (May 26, 2022), https://members.delphidigital.io/reports/the-hitchhikers-guide-to-ethereum#key-takeaways.

20 Instead of being stored inside blocks and competing with other projects for gas fees, as calldata did, blobs are "in sidecars" along the blocks. That is, blocks contain reference to the blobs, but do not actually contain blobs.

21 Prashant Jha, *Ethereum L1s Median Transaction Fees Decline As Much As 99% Post-Dencun Upgrade*, COIN TELEGRAPH (Mar. 14, 2024), https://cointelegraph.com/news/ethereum-l2s-fees-decline-99-post-dencun.

22 These concepts are not related to "sharding," which aims to split up the blockchain into multiple parts. Sharding is no longer part of the Ethereum development roadmap.

23 *Danksharding*, ETHEREUM (July 24, 2024), https://ethereum.org/en/roadmap/danksharding. *Merkle Patricia Trie*, https://ethereum.org/en/developers/docs/data-structures-and-encoding/patricia-merkle-trie.

24 L2BEAT, *Summary*, https://l2beat.com/scaling/summary. Most rollups are currently at "Stage 0", indicating a high degree of centralization.

concerns. Indeed, certain plasma projects have successfully transformed into optimistic rollups, which, as we recall, provide high data availability and the security that comes with it.[25] Further, as ZK-proof technology matures, plasma could see security gains by incorporating aspects of the technology. Some even posit that due to how plasma is structured, plasma infused with ZK proofs could offer more security than validiums.[26]

Rollups, validiums and plasma are only the more popular L2 solutions. There is also scaling technology beyond L2. Sidechains, for example, are independent blockchains that run in parallel to L1, with their own tokens and consensus mechanisms. From a user perspective, sidechains behave in the same way as L1 blockchains, instead of only performing partial functions of L1 in the case of rollups and relying on L1 for the rest. Sidechains do not rely on L1's consensus mechanisms, nor do they post transaction data back to L1, and so historically offered lower security compared to rollups.

While sidechains developed earlier than L2 solutions, some sidechains are now incorporating new developments to boost security. Polygon's most popular scaling solution, for instance, Polygon PoS, is a sidechain that is slated to upgrade into a zkEVM validium, which would rely on ZK proofs for security and be compatible with Ethereum.[27] This means it will no longer be a sidechain, but rather an L2 chain in its own right. This upgrade would allow users to enjoy the benefit of the existing blockchain, with its global ecosystem and complex smart contract functionalities, while also guaranteeing sufficient security. Polygon 2.0 also proposes linking separate L2 projects directly to each other without intermediation through L1 to further save on transaction delays and improve scalability.[28]

In addition, certain L1 blockchains can scale without the use of L2s or sidechains. Solana, as a Layer 1 blockchain, for instance, is much more scalable than Layer 1 Ethereum, but at the expense of lower decentralization[29] and decreased security.[30] Solana uses "Proof of History" alongside "Proof of Stake" to increase the number of transactions processed per second.[31] By assigning a timestamp to each block using a cryptographic function, "Proof of History" builds the precise time of each block into the blockchain.[32] This in turn allows validators to validate information much faster since they no longer have to allocate time and resources to agreeing on the order of each transaction, as in more traditional Proof of Stake protocols.[33] At the same time, Solana's architecture requires higher complexity and

---

25 X, https://x.com/Optimism/status/1516439487588364291.

26 *Exit Games for EVM Validiums: The Return of Plasma*, Nov. 14, 2023, https://vitalik.eth.limo/general/2023/11/14/neoplasma.html; PLASMA NEXT DOCS, *Overview*, https://intmax.gitbook.io/plasma-next-docs.

27 *Polygon 2.0: Polygon PoS -> ZK L2*, POLYGON LABS (June 20, 2023), https://polygon.technology/blog/polygon-2-0-polygon-pos-zk-layer-2.

28 *Aggregated Blockchains Fix the Layer 1 vs. Layer 2 Debate by Unifying Web3*, POLYGON LABS (Oct. 28, 2024), https://polygon.technology/blog/aggregated-blockchains-fix-the-layer-1-vs-layer-2-debate-by-unifying-web3; *see also* OPTIMISM, *Superchain Explainer*, https://docs.optimism.io/superchain/superchain-explainer.

29 Solana has around 1,400 validators, while Ethereum has more than one million. Solana Beach, https://solanabeach.io; Dune, https://dune.com/hildobby/eth2-staking.

30 *See, e.g., Solana vs. Ethereum: A Comprehensive Blockchain Comparison*, COINTELEGRAPH (Oct. 16, 2024) (noting more frequent outages in the Solana network compared to Ethereum).

31 Solana, *Solana Documentation*, https://solana.com/docs#understanding-the-architecture; Solana Foundation, *Proof of History: How Solana Brings Time to Crypto*, Nov. 29, 2021, https://solana.com/news/proof-of-history.

32 Solana Foundation, *Proof of History: How Solana Brings Time to Crypto*, Nov. 29, 2021, https://solana.com/news/proof-of-history.

33 Solana Foundation, *Proof of History: How Solana Brings Time to Crypto*, Nov. 29, 2021, https://solana.com/news/proof-of-history. Validators in Ethereum's Proof of Stake, for instance, rely on outside programs to assign a "median" timestamp, and validators subsequently rely on this timestamp to validate transactions in the order that they were received. *Id.*

compute power, which would contribute to decreased security (due to the heightened possibility of bugs) and decentralization (since less individuals can afford sophisticated hardware to participate as validators).[34] However, as Solana will likely be in a more prominent position going forward,[35] priorities for blockchain project developers could also shift accordingly.

LEGAL IMPLICATIONS

Scaling solutions promise to bring about more use cases of blockchain with higher transaction speed and lower costs. At the same time, these solutions necessarily sacrifice some parts of security or decentralization. A key issue is thus cybersecurity. For instance, a recent hack of Blast L2, an optimistic rollup, led to a $62 million loss.[36] A proposed solution was the rollup's core developers taking control of the network and reverting the unauthorized changes.[37] While this did not occur and the stolen funds were eventually returned without ransom, the Blast L2 hack once again brought the centralization debate to the forefront. As discussed above, most rollups currently rely on a single, centralized operator. To bypass centralization, rollups could allow users to submit transactions directly to the smart contract on L1 without having to rely on the operator's permission. Other rollups may rotate the operator role by choosing among prospective operators using a punishment and reward system.[38]

Further, as major permissionless blockchains start to include L2 solutions, regulators and standard-setting bodies would benefit from a knowledge of the types and functioning of L2 projects. Most L2s are in early stages of development. Building blockchains is an iterative process: most blockchains start with a core, centralized group of developers and, as they mature, become more decentralized as the founding group's work is replaced by and automated through smart contracts. In fact, there are evaluators that rate existing L2s according to features such as data availability so that users can understand more about the sophistication and security of each L2.[39] As the market matures, we expect baseline standards for network security and decentralization to eventually be developed, and that audits to ensure adherence to these standards will follow. It is important for developers and their counsel to fully understand the specific features and weaknesses of the various types of L2s, in order to select an L2—either to cooperate with or to build upon—that strikes an appropriate balance for the specific project at hand.

34 Unchained Crypto, *What Is Solana's Proof of History? A Beginner's Guide*, Apr. 1, 2024, https://unchainedcrypto.com/solana-proof-of-history.

35 *See, e.g.*, Krisztian Sandor, *TRUMP Token Frenzy Drives Solana Stablecoin Supply to $10B, Records DEX Volumes*, COINDESK (Jan. 24, 2025, 8:27 PM), https://www.coindesk.com/markets/2025/01/24/trump-token-frenzy-drives-solana-stablecoin-supply-to-usd10b-record-dex-volumes.

36 *Blast L2 Hack Prompts Debate Over Centralization of Ethereum Rollups*, CRYPTO NEWS (Mar. 27, 2024, 7:10 PM), https://cryptonews.net/news/security/28780789.

37 IMMUNEBYTES, *List of Largest Crypto Hacks / Exploits in 2024* (Mar. 6, 2024), https://www.immunebytes.com/blog/list-of-largest-crypto-hacks-in-2024.

38 *Zero-Knowledge Rollups*, https://ethereum.org/en/developers/docs/scaling/zk-rollups.

39 L2BEAT, https://l2beat.com/scaling/summary. In addition to data availability, L2Beat evaluates the development of L2 projects based on features such as the integrity of proofs and whether users can exit the L2 if an operator exploits their position and tampers with transactions.

CONCLUSION

As the web3 ecosystem continues to develop, different proposals have emerged that aim to enable blockchains to handle internet-scale activity. L2 solutions represent a major step forward in blockchain scalability and the ability to use blockchain as basic infrastructure. Within the confines of the blockchain trilemma, different L2 projects opt for different compromises. While this Tech Explainer reviews the current state of affairs, new technologies continue to emerge that mitigate the tradeoffs of decentralization, security and scalability. As L2 projects continue to develop, more general-purpose applications that enable seamless user experiences will be built on top of blockchains, bringing us closer to the future of a decentralized web.

Sasha Rosenthal-Larrea
+1-212-474-1967
srosenthal-larrea@cravath.com

Leslie Liu
+1-212-474-1297
lliu@cravath.com

CRAVATH, SWAINE & MOORE LLP

NEW YORK

Two Manhattan West
375 Ninth Avenue
New York, NY 10001
T+1-212-474-1000
F+1-212-474-3700

LONDON

CityPoint
One Ropemaker Street
London EC2Y 9HR
T+44-20-7453-1000
F+44-20-7860-1150

WASHINGTON, D.C.

1601 K Street NW
Washington, D.C. 20006-1682
T+1-202-869-7700
F+1-202-869-7600