galaxy

Galaxy Digital Research

In Search of Scaling: A Guide to Layer 2

SEPTEMBER 15, 2021

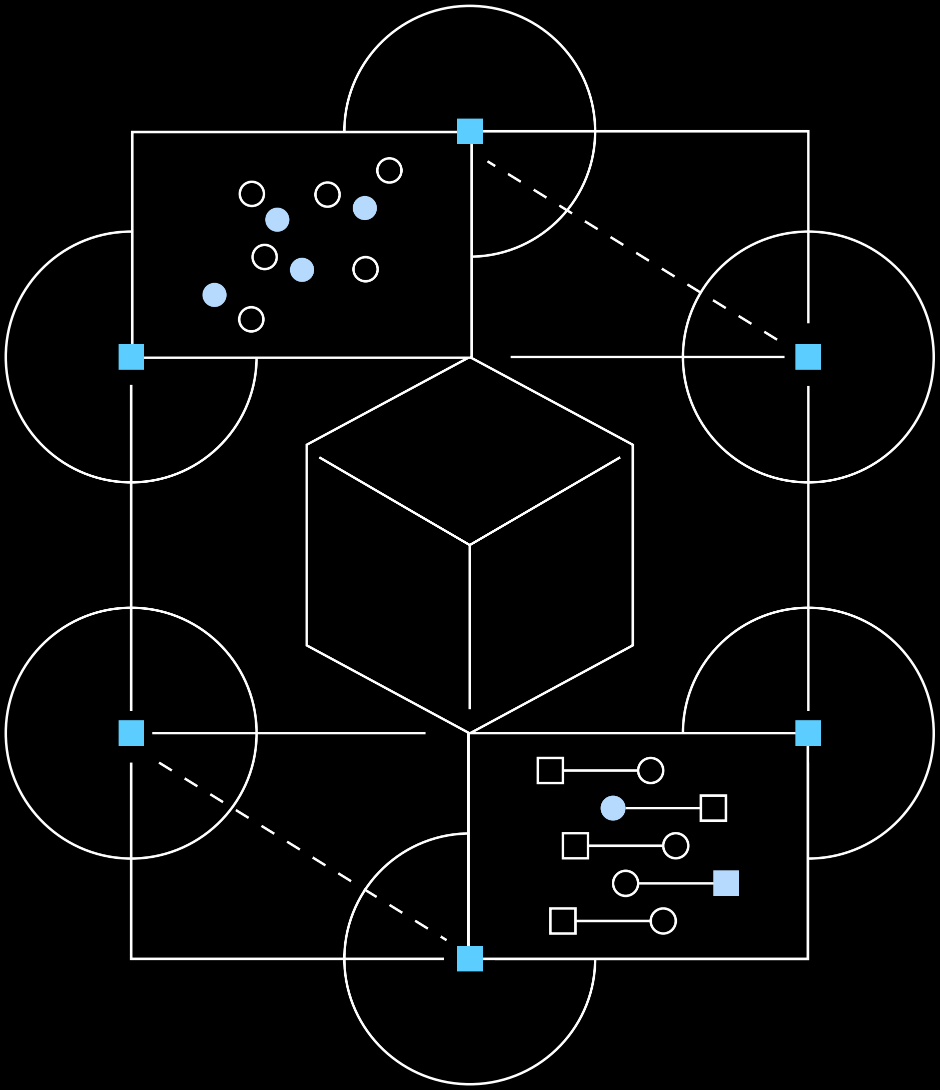

Author & Acknowledgements

Charles Yu, CFA
Research Associate
email: charles.yu@galaxydigital.io
twitter: @fullnodechuck

Alex Thorn
Head of Firmwide Research
email: alex.thorn@galaxydigital.io
twitter: @intangiblecoins

We would like to thank the following people for their support in researching and publishing this report:

- Gabriel Parker and Daniel Fishman, Galaxy Digital Research Summer 2021 Interns.
- Everyone at Galaxy Digital Research who helped make this report possible, including Walt Smith, Karim Helmy, and Sal Qadir.
- Our partners across Galaxy Digital including Michael Jordan, Will Nuelle, Jonathan Kol, Eva Casanova, Cat Orfanos, Beimnet Abebe, Trey Aslanian, Andrew Siegel, Rohan Muralidhar, Francesca Don Angelo, and Mike Novogratz.

This report is a product of Galaxy Digital Research, a research organization within Galaxy Digital, the leading provider of financial services in the digital assets, cryptocurrency, and blockchain technology sector. Galaxy Digital Research provides top-tier market commentary, thematic views, tactical insights, and deep protocol research.

View our publicly available research at www.galaxy.com/our-research. Contact us at research@galaxydigital.io.

Contents

Introduction

Public blockchains are distributed databases with the ability to reach consensus between an untrusted set of participants without the involvement of intermediaries or central parties. To achieving trustless consensus, blockchain design has necessarily made sacrifices on other fronts, including scalability, when compared to centralized solutions. Ensuring that a blockchain's nodes are widely distributed requires minimizing the burden on node operators, which necessitates limiting the amount of bandwidth, storage, and computation required to operate a node.

The two most important blockchains make different design decisions in this area. Bitcoin has optimized its development to create the most widely distributed node topology, while Ethereum has made some sacrifices in node operability to achieve additional scaling and functionality. Despite these different approaches, increased adoption of both networks has nonetheless necessitated demand for additional throughput and features that can't be achieved on the main blockchain (the "base layer," "main chain," or "layer 1") without accepting tradeoffs each community has deemed unpalatable.

Today, most agree that the most effective and sustainable path to scaling blockchains is to build other protocols atop the main chain in a layered approach, where higher-layer networks are introduced to increase functionality or throughput without compromising the fidelity of the base layer. This approach is favorable to expanding the footprint of the base layer feature-set because higher layers leverage the base layer's settlement assurances and security without negatively impacting its decentralization. In some ways, this is akin to how the internet scales, with HTTP built atop TCP/IP, HTML written on HTTP, and so on.

In this report, we examine several of the most prominent second layer ("layer 2" or "L2") scaling solutions, covering their designs, trade-offs, use cases, levels of adoption, and prospects for the future.

In search of scaling: why the base layer may not be enough

The primary motivations for scaling a blockchain are to bring more usability to the network. Growing adoption of crypto assets and an increasing universe of use cases has increased demand for block space on major blockchain networks. Blockchains must scale up to meet the growing levels of demand and to enable enhanced features or new applications that have not been possible on the base layer. Seeking to add more *throughput* to the network is perhaps the most common goal of those working on scaling blockchains, whether to reduce the settlement times of transactions, increase the count of transactions over a standard interval, or reduce the cost of making transactions.

But solutions are hard to implement at the base layer, and it is often not desirable or possible to implement the desired changes at the base layer. Enhancing the throughput of a blockchain often centralizes the network, which can weaken the value proposition of the system as a whole. Adding new features to a base layer can have unintended negative consequences, require a disruptive upgrade, or prove socially intractable. For example, the addition of more robust privacy to bitcoin's blockchain, resembling the features of ZCash or Monero, would be difficult to do while maintaining the credibility and transparency of the asset's monetary policy. On the other hand, zero-knowledge proofs could be employed at a higher layer without impacting the auditability of the base chain.

Let's quantify the need for scaling by looking at some on-chain data for the two most prominent blockchains: Bitcoin and Ethereum.

Fees

To prevent spam and incentivize miners, both Bitcoin and Ethereum require users to pay fees when submitting transactions to the network. But as user demand for block space increases, those with a high time preference opt to pay higher fees, increasing competition among users to transact on the network, resulting in higher transaction fees for all. While these fees ebb and flow with demand for block space, fees on both Bitcoin and Ethereum have been historically high several times over the last 18 months, to the point of causing significant disruption for users.

Median Tx Fee (USD)

Source: Galaxy Digital Research

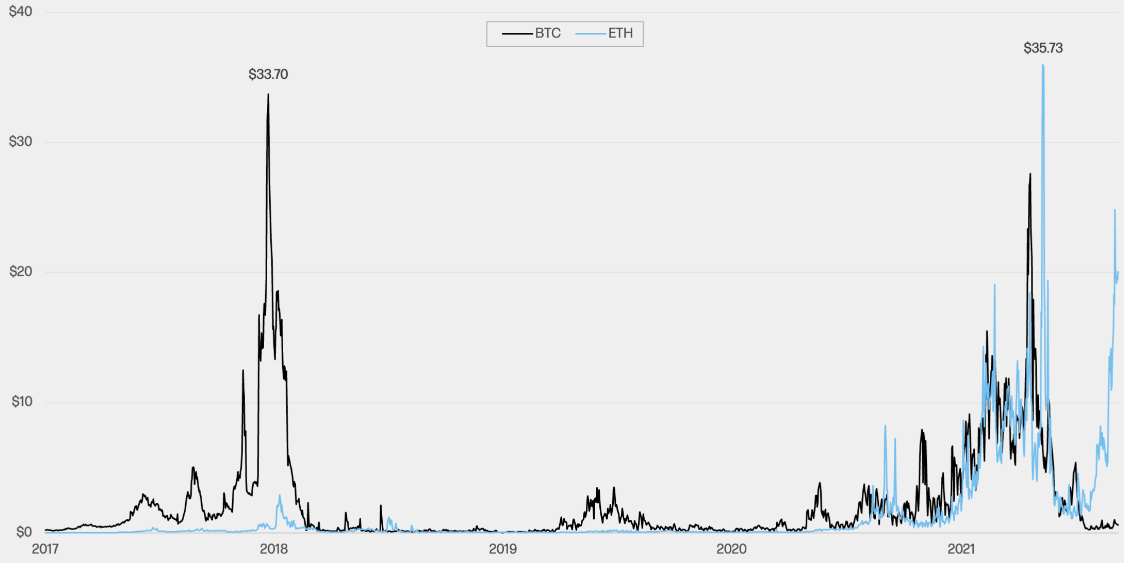

Data: Coin Metrics

Block Fullness

In 2017, Bitcoin's Segregated Witness ("SegWit") update changed how data in blocks is priced, performing calculations on *weight* rather than bytes and effectively increasing the maximum block size to 4MB. Bitcoin blocks have hit this limit fairly regularly since the upgrade's activation. ETH blocks have consistently been full since Summer 2020, in part due to the growing widespread use of increasingly complex transactions requiring more computation and more gas. ETH block sizes are periodically adjusted in network updates. Prior to the EIP-1559, ETH block sizes were hard-capped at 15 million gas. Implemented with the London upgrade during the first week of August, EIP-1559 enacted a variable block size that can temporarily be increased based on surges in demand, smoothing out gas prices. Read our report on EIP-1559 for more information. Generally, the world's two most valuable and prominent blockchains each have hit their throughput limits over the last 18 months.

Bitcoin - Daily Mean Block Weight

Source: Galaxy Digital Research

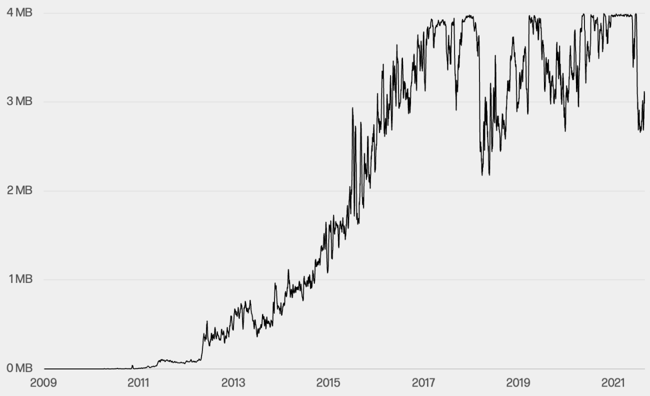

Data: Coin Metrics

Ethereum - Block Size vs. Network Utilization

Source: Galaxy Digital Research

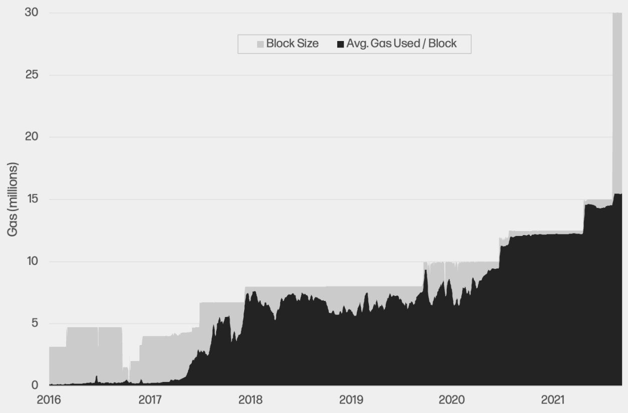

Data: Etherscan.io

Settlement Times

On average, Bitcoin blocks are published every 10 minutes, and Ethereum blocks are published every 13 seconds. However, in reality, transfer settlement times are longer (as most recipients will require several confirmations or blocks built on top of the block containing the transfer, before considering the transfer to have settled.) For example, Coinbase requires 3 confirmations (~30mins) for BTC deposits and 35 confirmations (~7.5mins) for ETH deposits. Thus, "settlement time" is different than "block time," and must take into account the robustness of the blockchain's security as well as it's "speed."

Transactions Per Second

As a function of both block times and block sizes, this metric can potentially be misleading because one transaction doesn't necessarily correspond to a single payment or deposit. One transaction can contain many outputs or represent an escrow that enables many off-chain payments with one on-chain transaction, which is true for many of the interactions between L1 blockchains and L2 networks discussed in this report. Transaction speed at the L1 level refers to the point that a transaction is irreversible; on L2, speed measures when transactions are committed and recorded - but they are not irreversible until they are finalized on the L1. Nonetheless, the number of transactions per time interval is a popular metric to compare the "speed" of different blockchains.

Certain use cases, like real-time payments or decentralized trading, are either prohibitively expensive for most users or outright impossible. Summer 2020 saw the first real growth of decentralized finance on Ethereum, known in the industry as *"DeFi Summer."* Concepts like algorithmic stablecoins and yield farming came to prominence and major applications like Yearn Finance, Aave, Curve, and Uniswap v2 launched. But as DeFi on Ethereum has proliferated, the platform's network constraints have often led to soaring gas fees. These costs have made DeFi on Ethereum prohibitively expensive for all but the largest transactions, pricing out average users and pushing some activity to alternative blockchains that may be less decentralized and come with higher security risks. Today, we see additional upward pressure on fees with congestion caused by waves of NFT drops.

BTC: Required Confirmations by Exchange

Source: Galaxy Digital Research

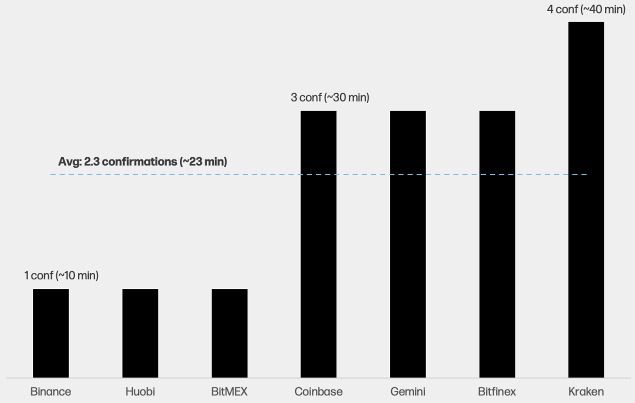

Data: Company websites

ETH: Req Confirmations by Exchange

Source: Galaxy Digital Research

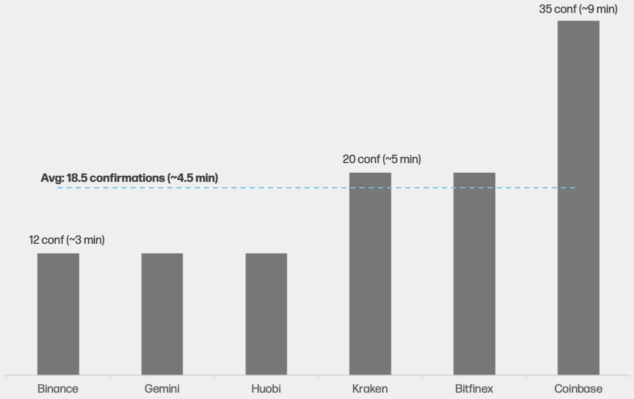

Data: Company websites

Tx Per Second

Source: Galaxy Digital Research

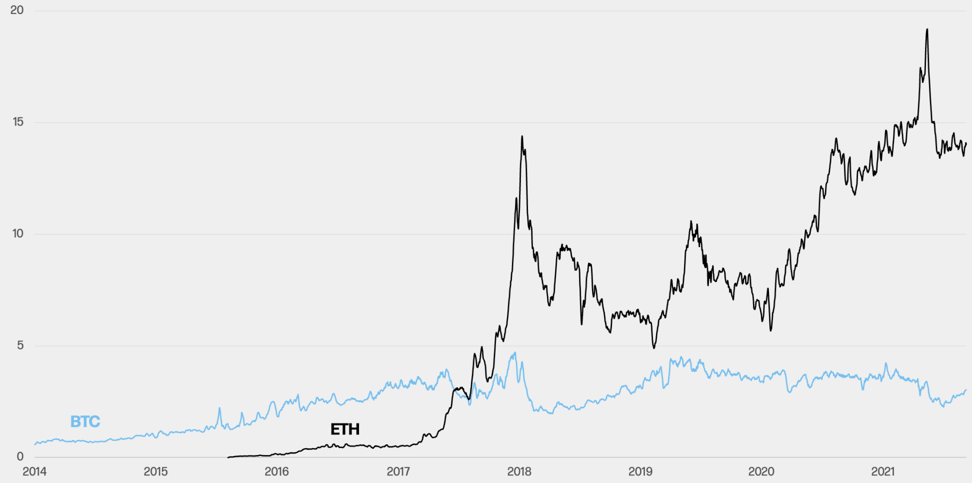

Data: Coin Metrics

Overview of Off-Chain Solutions

This year, we've seen accelerating growth on existing L2 protocols like Bitcoin's Lightning Network and the launch of several new L2 protocols on Ethereum that bring transaction costs down to enable quick and cheap payments, make DeFi more accessible to retail users, offer compliant solutions for various institutions, and enable new applications in derivatives and gaming verticals.

Off-chain scaling frameworks that are currently attracting developer activity largely fall into four different categories:

- **State channels**. State channels allow participants to transact off-chain, potentially an infinite number of times, and only commit the initial and final state to the mainchain when the channel is closed.
- **Sidechains**. Sidechains can operate in parallel with the base layer but are independent blockchains with their own consensus mechanism and security properties.
- **Optimistic rollups**. Rollups strip down and compress transaction data in rollup blocks to be submitted on-chain. Optimistic rollups assume in the best-case scenario that submitted transactions are valid by default; only if the submission is disputed will computations be executed on the mainchain to determine where the fraud occurred.
- **ZK-rollups**. In contrast with optimistic rollups, zero-knowledge rollups provide the proofs upfront for every state transition, which make it near impossible for operators to commit an invalid state.

Each of these scaling frameworks and each deployed protocol comes with different trade-offs along vectors like throughput, security, and decentralization. We've seen some protocols hybridize multiple frameworks in a search for the optimal level of trade-offs.

We describe some off-chain scaling models and their trade-offs in further detail below, along with their current implementation status and our outlook for each scaling design.

Source: Galaxy Digital Research

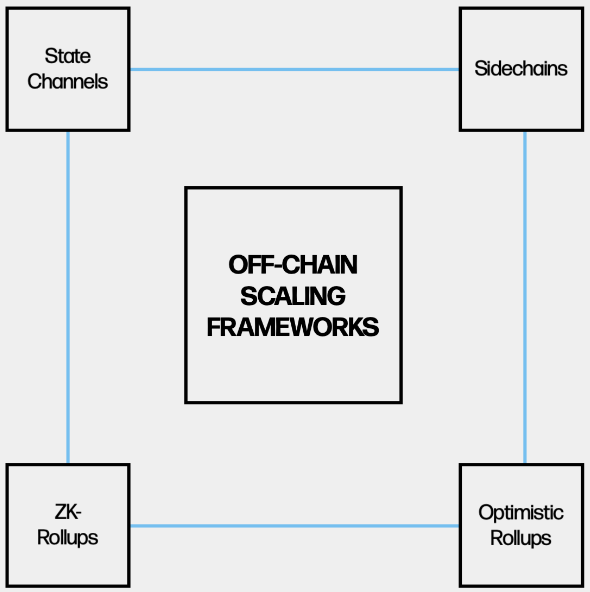

Evolution of Off-Chain Scaling Proposals and Ideas

As background, we take a view of how off-chain scaling designs have evolved over time. Each iteration can be viewed as a potential advancing evolution of past ideas, building off existing designs, addressing potential shortfalls, and picking which aspects to keep going forward.

Two scaling solutions for Bitcoin have seen meaningful usage: the **Lightning Network**, a layer-2 state channel protocol used for payments, including micropayments and cross-border remittances; and the **Liquid Network**, a sidechain used by exchanges and traders for faster settlement.

Lightning is the main implementation of a **state channel** scaling protocol and has recently seen a surge in interest. Recent growth can be attributed to the proliferation of accessible Lightning Network hardware and software solutions like Umbrel, Bluewallet, and Strike, as well as efforts to drive adoption of Lightning Network usage among individuals (#PlebNet) and nation states (El Salvador).

Timeline of L2 Proposals VS. Select Protocol Deployment

Source: Galaxy Digital Research

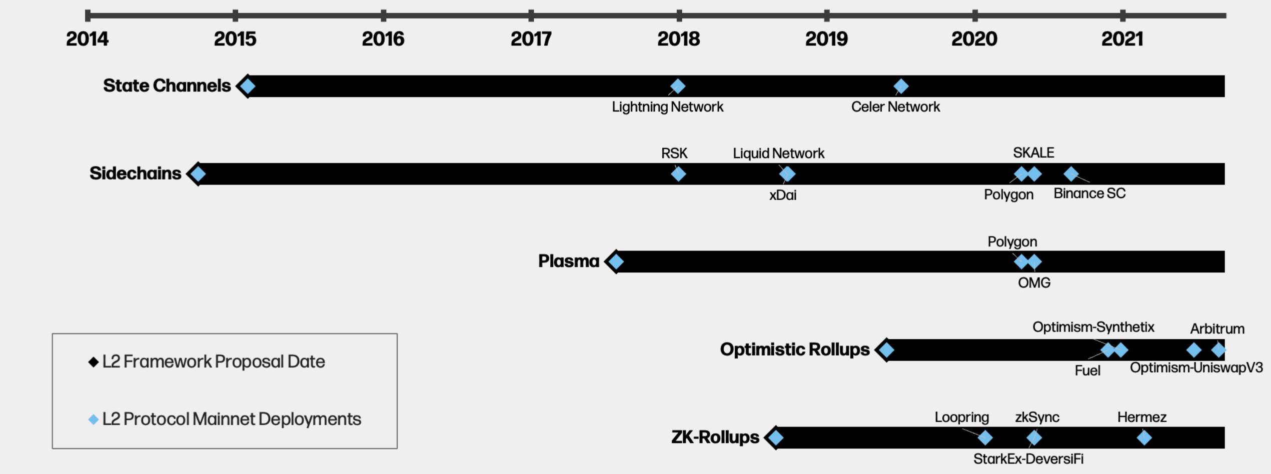

Initial L2 Framework Proposals

Source: Galaxy Digital Research

	Framework	Proposal Dates	Notes
STATE CHANNELS	State Channels	Feb 2015	Lightning Network white paper draft was published in February 2015.
SIDECHAINS	Sidechains	Oct 2014	Sidechains for bitcoin were discussed as early as 2012; framework more formulated when the Pegged Sidechains white paper was published in 2014.
PLASMA	Plasma	Aug 2017	Vitalik Buterin & Joseph Poon released the Plasma white paper in August 2017.
ORUs	Optimistic Rollups	Jun 2019	Arbitrum published working paper on scalable smart contracts in Jan 2018; Optimistic rollup framework solidified from John Adler's Min. Viable Merged Cons in June 2019.
ZKRUs	ZK-Rollups	Sep 2018	Vitalik published a post on using zk-SNARKs to verify blocks in September 2018. Barry Whitehat presented zero-knowledge proofs in late 2018.

Early Ethereum L2 protocols largely resulted in disappointment, as hopeful projects have seen their production timelines slip or failed to draw meaningful adoption after initial deployment. One such protocol is Raiden, an analog to Lightning on Ethereum. In 2016, Raiden CEO Heiko Hees contended that Ethereum was better equipped for state channels than Bitcoin and had targeted an alpha launch of the network in Q3 of that year. That timetable proved to be too ambitious, as the Raiden Network is still not production-ready, despite continual development.

Sidechains are blockchains that run in parallel to the parent blockchain. While assets can be moved between chains using bridges, sidechains are technically not considered a "layer-2" because they can employ their own consensus mechanisms and security properties. In short, sidechains are *other blockchains*, but whose purpose is to connect primarily to another L1 blockchain network. Sidechains have been operating in the wild for several years, offloading some network congestion on their main chains. The flexibility and ease of deployment allowed sidechains to be first widely adopted off-chain solution, but their security shortfalls have led some to search for a more optimal scaling framework with stronger guarantees.

L2 Protocol Deployment

Source: Galaxy Digital Research

	Mainnet Launch	Protocol	Notes
STATE CHANNELS	Sep 2018	Lightning Network	Lightning launched payment channels for microtransactions and x-border remittances
	Jul 2019	Celer Network	Celer Netowrk's alpha-mainnet launched; first generalized state channel network
SIDECHAINS	Jan 2018	RSK (BTC)	RSK (sidechain) enables execution of smart contracts with bitcoin as native asset
	Sep 2018	Liquid Network (BTC)	First production Bitcoin sidechain; used for exchanges, brokers, market makers, FIs
	Oct 2019	xDai	xDai stablechain created as a sidechain bridge with Dai stablecoin as native token
	May 2020	Polygon	Matic (now Polygon) launched with hybrid architecture Plasma + PoS commit chain
	June 2020	SKALE	Configurable network of sidechains; Phase 1 mainnet launch restriction for validators
	Sep 2020	Binance Smart Chain	Binance Smart Chain launched with Proof of Stake Authority consensus mechanism
PLASMA	May 2020	Polygon	Matic (now Polygon) launched with hybrid architecture Plasma + PoS commit chain
	June 2020	OMG	SYNQA's OMG Network (prev. OmiseGo) launched for high throughput transactions
ORUs	Dec 2020	Fuel	UTXO-based transaction system for payments; first optimistic rollup on mainnet
	Jan 2021	Optimism-Synthetix	Staking on Synthetix goes live on Optimistic Ethereum (OE) L2 mainnet
	May 2021	Arbitrum	Arbitrum mainnet opened to devs; deploys Arbitrum contracts on mainnet
	July 2021	Optimism-Uniswap	Uniswap V3 Alpha launch on OE Mainnet
ZKRUs	Feb 2020	Loopring	Exchange and payment protocol; first deployment of a ZK-rollup
	Jun 2020	StarkEx-DeversiFi	Starkware launched StarkEx on ETH mainnet to power DeversiFi's DEX
	Jun 2020	zkSync	V1.0 on mainnet for payments; smart contracts to be added at a later date
	Mar 2021	Hermez	Hermez Network launched for payments and token transfers

In August 2017, Vitalik Buterin teamed up with Joseph Poon, who co-authored the Lightning Network white paper, to propose the L2 successor to state channels and sidechains: **Plasma**. Plasma chains resemble sidechains but are non-custodial and rely on the security guarantees of the L1 mainchain, making them a true L2 solution. After years of development, some shortfalls with the framework became apparent: namely, that it suffers from data inaccessibility, as not all off-chain transaction data is reconcilable from the public information on the Ethereum blockchain. This eventually drove most of the prominent Plasma developers and the rest of the Ethereum community to largely abandon the framework and instead rally behind the idea of rollups.

Rollups aim to minimize the data footprint on L1 while concurrently preserving the ability for anyone to recreate the chain – solving the data availability limitations of Plasma. Rollups strip down the transaction data and compress the remaining essential components to be published on-chain so observers can keep up with the state of the network.

There are two main categories of rollups: **zk-rollups,** which leverage zero-knowledge proofs, and **optimistic rollups**, which take advantage of optimistic evaluation. The main difference between the two is when the cost of validation is paid: zk-rollups pay the cost of validation upfront while optimistic rollups delay the cost until after a dispute is raised. Zk-rollups were proposed before optimistic rollups as an L2 solution, but the latter flavor is further along in its production timeline for generalized smart contracts with several prominent optimistic rollup projects have either just launched or nearing deployment (e.g. Optimism and Arbitrum). Zk-rollups are more complex in design and have yet to launch in production for functions beyond simple payments and token swaps, but many observers and practitioners, including Vitalik Buterin, believe they will be the eventual winner of the L2 scaling wars.

State & Payment Channels

State channels were the first L2 scaling design to see meaningful adoption. State channels refer to more generalized, smart contract-based transactions while payment channels, such as the Lightning Network, are a specific subset of state channels.

State channels avoid mining every single transaction and only commit the initial and final state to the base layer when the channel is opened or closed. This setup allows participants to transact off-chain, potentially an infinite number of times at a minimal or zero cost with near-instant finality. The only involvement of the main chain is when channels are opened and closed, initiating or net-settling the channel for ultimate finality. In theory, state channels have a limitless throughput, capable of supporting millions of transactions per second.

Operating Framework

To transact using payment channels, participants must first open a channel between themselves – this can be a direct channel between the transacting parties or an indirect channel which is routed through middlemen connecting nodes. Participants must deposit liquidity into the channel, which entails locking up some amount of tokens at the base layer for the lifetime of the channel. To achieve bi-directional payments, meaning payments that can be sent both ways across a channel, requires both participants to commit liquidity to the channel.

From here, the transacting parties can conduct an unlimited number of payments between each other (as long as the net balance does not exceed the amount of liquidity locked in the channel). Users provide digital signatures with each transaction, serving as proof to prevent double spending, with each transaction resulting in an updated state between the participants. Once a participant is ready to settle the balance, they initiate a request to close the channel and submit a proof of the end state, or the final balances between the two parties. The counterparty would have to sign off on the submitted balance before funds are ultimately settled with one party paying their outstanding balance.

The 2-of-2 multi-sig setup requires both participants to remain online (*liveness* assumption). In addition to just the payor, the payee on the opposite end of the channel must remain online to sign transactions and to monitor the channel state for account balance accuracy. If there is any disagreement between the two parties, each may post the most recent timestamped signed ticket for the base layer to arbitrate and resolve (fitting the definition of an L2). Most forms of fraud on state channel networks come from a user who broadcasts an old state and requests an exit from the channel without notifying the other party who is offline (defined as a *non-cooperative closure*).

Security

To mitigate the impact of non-cooperative closures and to partially relieve users of the active monitoring requirements, the Lightning Network relies upon *watchtowers*, which are third-party surveillance nodes. Watchtowers handle fund recovery services when participants are disconnected for an extended period (either intentionally or unintentionally) by monitoring transactions broadcast to the mempool. Hash time lock contracts ("HTLC") provide a timestamp with each channel state update that is recognizable by the watchtowers and help prevent false state reports. When transactions are flagged as outdated contracts, watchtowers launch the fund recovery process by reverting the channel's history to the most recent state signed by both parties.

The main challenge facing state channels is building the network in a way that strikes the right balance between efficiency and decentralization. As payment channel networks are nondirected graphs, the larger the network grows, and the higher number of connections within the network, the more effective it is. However, it is unreasonable to open a channel between all network nodes and not ideal for a single entity to serve as the connecting hub between all nodes. Forming a hub & spoke model places significant reliance on a single entity, creating a centralized point of failure.

Payment Channel Token Flow: BTC Wallet

Source: Galaxy Digital Research

- STEP 1
  Funding transaction – lock up funds on mainchain
- STEP 2
  Open channel between nodes (pay on-chain fees)
- STEP 3
  Participants may send bidirectional payments, each resulting in an updated state
- STEP 4
  Either participant submits request to close channel
- STEP 5
  Both users must provide sign-off on final state
- STEP 6
  Final state is broadcast to mainchain

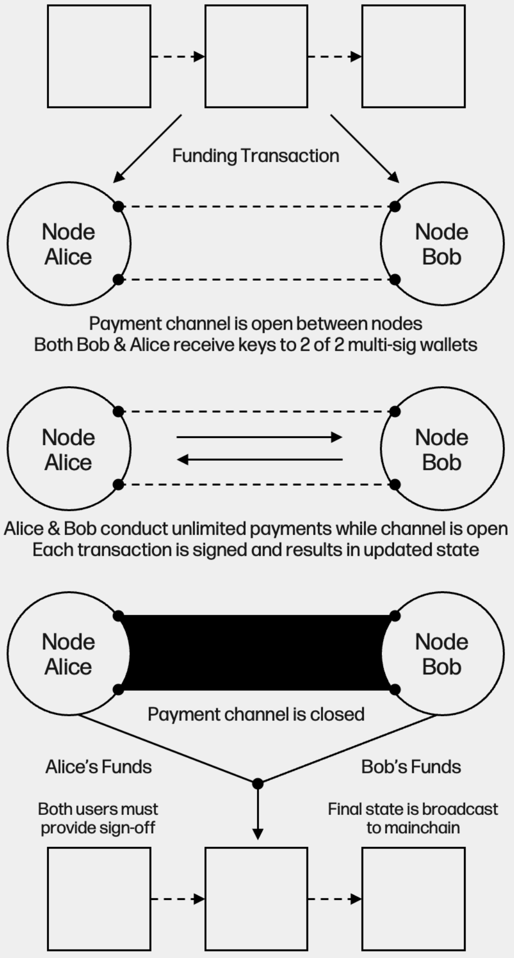
Funding Transaction
Node Alice
Node Bob
Payment channel is open between nodes
Both Bob & Alice receive keys to 2 of 2 multi-sig wallets
Node Alice
Node Bob
Alice & Bob conduct unlimited payments while channel is open
Each transaction is signed and results in updated state
Node Alice
Node Bob
Payment channel is closed
Alice's Funds
Bob's Funds
Both users must provide sign-off
Final state is broadcast to mainchain

State Channels Advantages

- **State channels have near-instant settlement finality (often settling in milliseconds)**, meaning as soon as both parties sign a state update, transactions can be considered final. The confirmation time once channels are closed, and the end balance is broadcast are dependent on the block time of the base layer (e.g. ~10 minutes for bitcoin; 13 seconds for Ethereum).
- **The throughput of state channel networks is theoretically infinite** and fees for off-chain transactions are near-zero (although it can vary based on routing fees specified by intermediary nodes).
- **State channels preserve privacy.** Transactions within a direct channel are visible only to the two participants transacting in the channel. The only information posted on-chain is the net channel balance once participants settle by closing their channel, and this is difficult to trace back off-chain. A transaction between two participants is routed through connected intermediary nodes still preserves privacy if onion routing protocols are used. Onion routing limits the information shared with routing nodes so that they only have visibility on where the transaction came from and where the transaction needs to be routed to but do not have information on the quantity of routing nodes used or the end recipient.

State Channel Disadvantages

- **Liveness requirement comes with security risks.** State channels require the payor, the payee, and potentially any routing nodes to remain online, which is demanding and undesirable for some users. Although the payee may delegate network monitoring responsibilities to watchtowers, the payee still must provide sign-off using his or her private key. The state channel framework is only suitable for applications with a defined set of participants given the state deposit contract must always know the participant addresses that are part of the channel. This constraint extends to the custody setup; LN users can't exactly store their coins on a hardware wallet due to the liveness requirement, so they must store funds on other third-party wallet software providers. State channels also do not automate routing fee adjustments based off network activity, requiring users to manually adjust fees themselves.
- **Capital inefficiency.** State channels do not completely bypass potentially high L1 transaction costs and must pay on-chain fees for the opening and closing of channels. In order to facilitate payments, all participants including payees and each routing node must lock up capital in each channel to provide sufficient liquidity (the amount being paid at a minimum). Given the minimal fees required to transact, state channels have limited incentives to distribute among network facilitators such as nodes, watchtowers, and routing nodes, which opens to potential exploits.

Adoption of Bitcoin's Lightning Network

Since its mainnet launch in early 2018, Bitcoin's **Lightning Network** has grown significantly with a sizable portion of the growth coming in recent months.

As of early September 2021, the network capacity measured in BTC has more than doubled during the year to over 2.3k BTC – when measured in dollars, the network capacity has surpassed $120m. The number of nodes with public channels on the network totals over 15k (+86% YTD). The total channel count stands at ~70k, outgrowing the node count at +89% YoY, and implying nodes have been more active now from earlier this year. This growth has been achieved without the introduction of a native, novel token or the "liquidity mining" incentives that have boosted the growth of DeFi protocols. Looking ahead, Lightning also has catalysts from El Salvador's bitcoin law going into effect and Twitter's planned integration of Lightning payments.

Lightning Network: Nodes w/ Channels

Source: Galaxy Digital Research

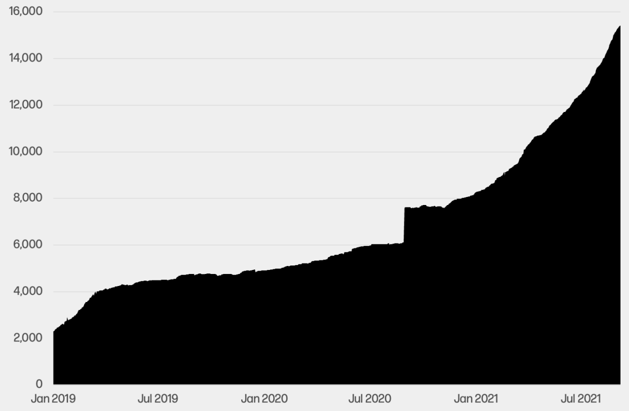

Data: Bitcoin Visuals

LN Channel Count vs. Network Capacity (BTC)

Source: Galaxy Digital Research

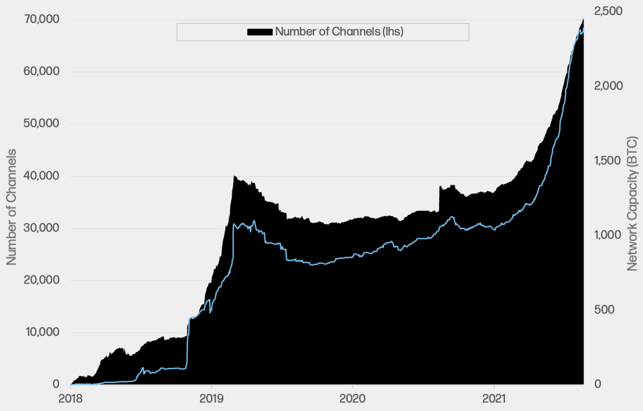

Data: Bitcoin Visuals

However, we note that these reported numbers may understate the true size of the Lightning Network since they only include public LN channels while some nodes may opt to open private channels. The privacy preserving design on the Lightning Network also prevents us from calculating other measures of engagement such as the transaction count per node or the actual network volume that is routed through channels. Standard measures of engagement such as number of transactions per node is unavailable given the nature of state channels only broadcasting the delta of payment transactions or the net result. Furthermore, major routing node operators report to us that they often turn the capital in their channels several times, meaning the total "usage" of the Lightning Network could be several multiples the visible "locked value."

State Channels on Ethereum

- **Raiden Network** is one of the earliest state channel projects and was intended to be the Ethereum-version of the Lightning Network but supporting ERC20 tokens instead of bitcoin transfers. According to the project's roadmap, the first phase of development, called µRaiden, has been live on mainnet since 2017 for unidirectional many-to-one payment channels. However, a functioning version of Raiden for many-to-many payment setups is not yet ready.
- **Celer Network** launched on the Ethereum mainnet in July 2019, becoming the first generalized state channel network to go live. Celer's State Guardian Network (SGN) guards off-chain states when users are offline – similarly to watchtowers in the Lightning Network. Celer delivered on its goal of becoming the world's first blockchain agnostic L2 by supporting both Ethereum and DFINITY. However, after general purpose computing was limited and failed to catch on, Celer repurposed its state channel technology to power cross-chain transfers and incorporate support for rollups (a playbook also followed by Connext, another interoperability protocol), and has launched products such as layer2.finance (rollup) and Celer cBridge to bring connectivity and accessibility across multiple chains.

Sidechains

Sidechains are independent blockchains, operating in parallel with the base layer through embedded connectivity but with their own separate operators, validators, and security mechanisms. The flexibility of sidechains allows for quick deployments to support high-throughput and low-latency transactions for users.

Since they employ their own consensus models, sidechains are not technically considered to be "Layer-2" by some critics who believe they function more as a separate, scaling L1. However, sidechains can be architected many ways and there should be a distinction between those that are properly aligned with and complementary to the base layer and those that aren't, although this distinction is not always clear.

Framework

The ability for sidechains to communicate and move assets between chains is a key aspect. Sidechains maintain connectivity and interchain messaging with the base layer through a two-way peg, which entails bridging assets by locking one's assets to a multi-sig address so that another useable version of that token can be unlocked on the sidechain. The two-way peg enables the interoperability When users want to move their assets back to the base layer ("peg-out"), the assets are typically burned from the sidechain and then the assets on the base layer are unlocked.

1 To move assets onto a sidechain, users send their tokens ("peg-in") to a multi-sig bridge address that is typically managed by the sidechain operator to create a 1:1 peg.

2. The proof of funds is then relayed to the bridge node, which then proceeds to unlock or send the equivalent amount of a useable version of the token to a corresponding sidechain-based wallet.
3. To exit funds back to the mainchain ("peg-out"), users will redeem the sidechain-based asset for the mainchain asset by sending funds to a bridge address/contract typically operated by a federation or trusted third-party who can verify the proof of funds.
4. After receiving the sign-off, the user's sidechain tokens may be burned or destroyed, and the equivalent amount is then unlocked on the mainchain and sent to the user's address.

Source: Galaxy Digital Research

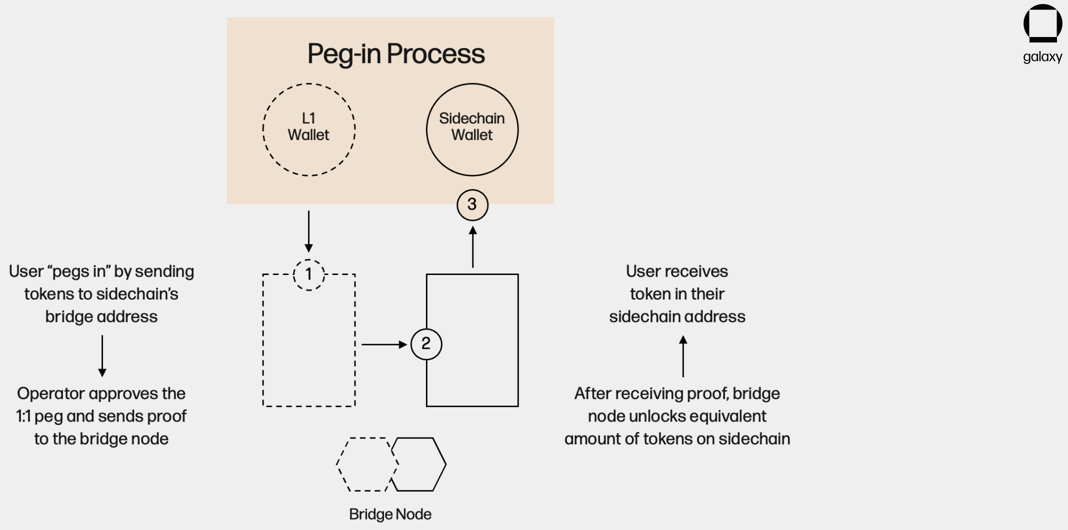
Peg-in Process
L1 Wallet
Sidechain Wallet
galaxy
3
1
2
User "pegs in" by sending tokens to sidechain's bridge address
Operator approves the 1:1 peg and sends proof to the bridge node
User receives token in their sidechain address
After receiving proof, bridge node unlocks equivalent amount of tokens on sidechain
Bridge Node

Bridge nodes are responsibility of receiving proof of locked tokens on the mainchain to release the equivalent sidechain version of the token to the user's sidechain wallet. The bridge nodes are the middleman facilitating the peg-in and peg-out process based on the checkpoints submitted by the network validators.

Sidechain Designs

Apart from the general two-way peg process to move assets between the base layer and the sidechain, sidechains can be architected in a variety of ways which can have meaningful considerations for the governance and security of the network.

Some of the biggest design decisions of sidechains are the consensus mechanism employed and the incentives in place to deter malicious behavior, which have important security considerations as it relates to censorship-resistance and fund ownership guarantees. The consensus mechanism and the custodial set up has implications for network participant groups including the transacting users, validating nodes, block creators, and the operators – although the flexibility of sidechains also means that the importance of these participant groups can relatively easily be changed.

The two most common examples of consensus mechanisms employed by sidechains include:

- **Proof of Stake** (PoS). Rather than being tied to computing power like PoW, PoS creates a different incentive model based on the financial value of one's stake in the network. The power that each validator holds is usually directly proportional to the number of tokens staked across the network. In the case of a fraud attempt or an attack on the network, malicious actors would see their stakes slashed. The idea is to create a mechanism to hold transaction validators accountable and ensure that they act in the best interests of the network. Issues with PoS include potential consolidation of stake that enables central permissioning. Another variation is Delegated Proof of Stake (DPoS) where network users vote on delegates by pooling tokens into a staking pool to elect the block producers. Instead of being solely tied to the monetary value staked, reputation becomes a factor when selecting validators in a DPoS system.
- **Proof of Authority (PoA).** PoA is similar to PoS but instead of staking tokens, participants stake their identity and reputation. PoA is a permissioned consensus mechanism where the set of validators are pre-determined and are intended to be identified and trusted. Benefits to PoA are efficiency through fast validation of transactions and low computing demand (eliminating the need for mining rewards). Of course, this means the network is very centralized and opens up to potential for corruption, manipulation, or common attack vectors like DDoS and 51% attacks. PoA sidechains are commonly used as controlled environments for testing different features and are used by three Ethereum testnets (Kovan, Goerli, and Rinkeby). Enterprises can also leverage PoA in private blockchain designs for internal transactions.

Pros and Cons

There are different implementations of each of these consensus mechanisms but generally, sidechains limit the number of entities that can validate transactions. This helps to maintain operators the flexibility to configure nodes and to power the network in an efficient manner. Sidechains are especially valued by developers for their usability with elastic support for smart contracts and their simplicity – some sidechains can be spun up in a day. They have low on-chain data requirements and offer high throughput.

On the other hand, the limited set of validators also make sidechains susceptible to collusion-based attacks. Sidechains also have relatively weak censorship-resistance properties and do not typically provide data availability guarantees or security guarantees around ownership of funds. Before jumping onto a sidechain, users should understand what the lock-up conditions are, who controls the exits, and how secure and trusted the operators are.

Bitcoin Sidechains

Some examples of Bitcoin sidechains include:

- **Liquid Network**. Liquid, created by Blockstream, is a settlement network used mostly by traders and exchanges for fast and confidential bitcoin transactions. Liquid is governed by a federation of 57 members handpicked by Blockstream including various exchanges, brokerages, wallets, and infrastructure providers. Federation members are responsible for membership, oversight, and technology of the network – including voting on future network upgrades, signing blocks, and maintaining the PAK list for peg-out transactions. Anyone of the 57 internal federation members poses a central point of failure but Liquid has an emergency recovery procedure, which consists of Blockstream having a set of three emergency keys to access all the funds on Liquid if the network is compromised. As of August-end, L-BTC in circulation totaled over 3,200.
- **RSK Network (Rootstock)**. Launching on mainnet in January 2018, RSK is Bitcoin's first general purpose smart contract platform and arguably the most secure smart contract platform in the world as it derives its security from Bitcoin miners – RSK is merge-mined with bitcoin typically with over half of the Bitcoin hash rate simultaneously mining both BTC and R-BTC. RSK maintains an open peg so that users do not have to go through an exchange or KYC process. RSK is operated by the PowPeg Federation, a uniquely designed group whose main purpose is to secure the two-way-peg and requires members to audit node software. As of August-end, RSK had roughly 2,000 RBTC locked in its 2-way-peg along with over 58k active accounts

Circulating Supply – Liquid & RSK Tokens

Source: Galaxy Digital Research

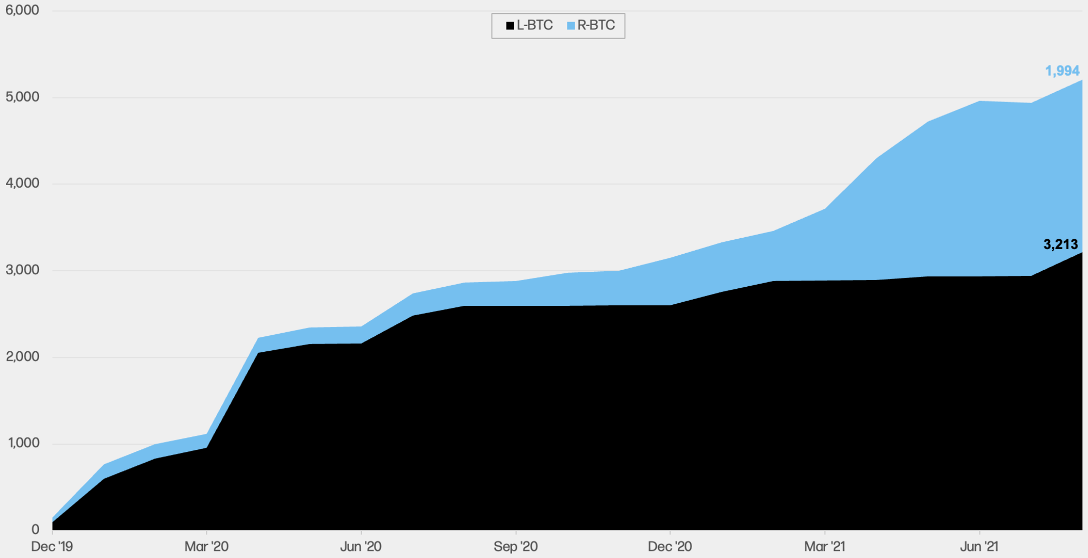

Data: Liquid.net, explorer.rsk.co

Ethereum Sidechains

Some examples of Ethereum sidechains include:

- **Polygon**. Matic (now **Polygon**) launched Matic Network Mainnet in May 2020 with a hybrid architecture consisting of its own implementation of Plasma and a PoS commit chain, which offers its users transaction fees well under one cent and can reportedly handle up to 7k TPS. Matic rebranded as Polygon in February 2021 as it shifted its strategic focus on creating a multi-chain ecosystem to support additional L2 solutions including rollups (discussed later). Note that some critics argue that the Polygon should not be classified as an L2 solution given the architecture primarily relies on the PoS commit chain – designed to offer more security measures compared to a more typical sidechain. L2 or not, Polygon has been beneficial to the Ethereum blockchain as one of the first smart contract platforms deployed with complete EVM-compatibility for portability of existing smart contract bytecode. At a time when Ethereum gas fees have priced certain users out, Polygon has attracted key Ethereum-native projects to its ecosystem and provided an outlet for those looking to those in search of lower fee environments.
- **xDai**. xDai uses DPoS with a pool of 19 validators and has block times of 5 seconds. xDai has been live since October 2018 and uses xDai as its native token, allowing for transactions to be paid without using ETH for gas. The xDai chain was built primarily for P2P payments but has since expanded to support other applications including in DeFi, DAO governance, and gaming.
- **SKALE**. SKALE is an elastic sidechain network protocol built to support thousands of independent blockchains and subchains tied to the Ethereum ecosystem. SKALE Mainnet launched in June 2020. Rather than relying on a small set of validators, SKALE uses a pooled validation model to attain a more collusion-resistant leaderless network compared to other sidechain protocols. Currently, there are 46 validator organizations running on SKALE.

Plasma

Plasma chains are an evolution of sidechains with an added level of security designed for Ethereum. Plasma was proposed in August 2017 by Vitalik and Joseph Poon, co-author of the Lightning Network white paper, as a non-custodial scaling framework, meaning users can recover funds back to the L1 mainchain in the event of an invalid chain on L2 or if the plasma chain operator goes offline. Rather than operating as a separate chain with its own (usually weaker) security properties, Plasma chains are a true L2 solution because they rely on the security guarantees of the underlying L1.

Transactions are aggregated by a plasma chain operator, responsible for batching the transactions and compressing them down to their Merkle root, which is then published to the mainchain for validation. Thousands of transactions can be executed off-chain while adding only a single hash to the Ethereum blockchain.

While Plasma improves upon the security of sidechains at a high level, the Plasma framework inherently suffers from its own set of problems. After several years of protocol development, three major limitations have become apparent:

- **Limited ability to execute smart contracts**. Plasma chains are limited in their ability to run the Ethereum Virtual Machine (EVM), Ethereum's runtime environment for smart contracts. This caps their usage to basic functions like token transfers and swaps, while other general computation use cases are not supported in a Plasma framework.
- **Difficulties in exiting**. Plasma chains are subject to liquidity constraints due to the fraud proof security mechanism – subjecting users to a lengthy challenge period (~1-2 weeks) to allow for a sufficient dispute window. If all users needed to exit at once in a worst-case bank run scenario, the latest valid state of the chain would have to be posted on the mainchain in a single challenge period. The exit process also requires participants to regularly be online to monitor the Plasma chain for exploits or to delegate the responsibility to another trusted actor.
- **Data availability**. Not every Plasma transaction is submitted to the mainnet by the operator. The operator may publish the block header, but the underlying off-chain transaction data is not publicly available to participants. This offline storage creates challenges in reconciling transactions and reconstructing the state of the chain, highlighting centralization concerns around immutability and censorship-resistance by the operator.

Plasma Usage / Activity

After Plasma initially attracted meaningful interest from developers, most protocols have either abandoned Plasma in favor of rollups or have adopted hybrid approaches incorporating other scaling frameworks. Data from L2Beat shows less than $5m of TVL locked in Plasma chains (Polygon not included).

- **Gluon** started out as a Plasma-based DEX operator, powering Leverj's DEX on the Ethereum mainnet since February 2019, before pivoting into rollups, citing scaling limitations and mass exodus security concerns with Plasma.
- As previously mentioned, **Polygon** was launched in May 2020 with a hybrid architecture consisting of its own implementation of Plasma and a PoS commit chain. Most of the growth of the Polygon ecosystem to date has been on the PoS commit chain, while growth on the Plasma chain has been more muted.
- SYNQA's **OMG Network** (previously OmiseGo) launched in June 2020 and is built using a Plasma implementation called More Viable Plasma (MVP). Genesis Block Ventures acquired parent company SYNQA in December 2020 with plans to grow the OMG Network adoption particularly in Asia. In May 2021, OMG Network added a new product group focused on smart contracts, OMGX Optimistic Rollup (now branded as the Boba Network), to complement the original OMG Plasma architecture used for high throughput transactions.
- The former researchers from the non-profit **Plasma Group** abandoned the project in January 2020 to form a new company focused on rollups called **Optimism** (discussed below).

Total Value Locked in Plasma - YTD ($USD)

Source: Galaxy Digital Research

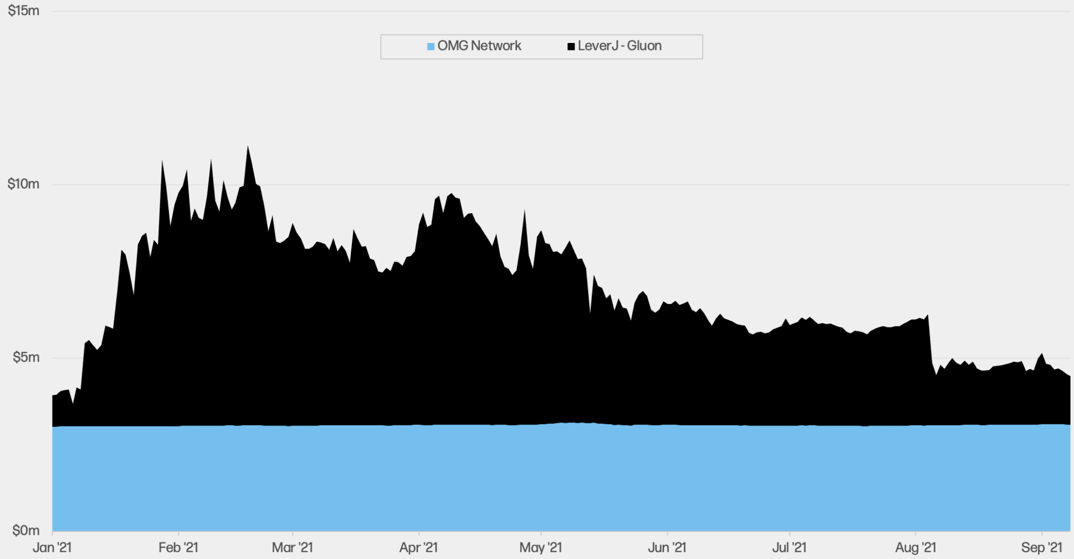

Data: L2Beat

Rollups

Rollups enable hundreds of transactions to be batched together and published together in a single block. User funds are stored in smart contracts on the main chain while state transition data is maintained in a separate off-chain state in a Merkle root.

The main proposition of rollups is to minimize the data footprint on L1 while still preserving the ability to check for fraud. Rollups build upon the data availability limitations of Plasma by publishing a compressed form of transaction data on-chain so that everyone can reconstruct the chain and keep up with the latest state of user account balances and contracts. In the event the sequencer disappears, a new sequencer may retrieve all the L2-related data from Ethereum, reconstruct the latest L2 state and continue from where their predecessor left off. Essentially, the main chain's smart contract containing enough data to reconstruct and prove the off-chain transactions are valid, but without storing all their data.

Rather than including the full transaction data on-chain, rollups create an index position for each address, where a subtree would then be added to the state to allow participants to map the indices to addresses [using the CALLDATA function in Ethereum]. Transactions are compressed to only the necessary components (the to/from addresses, transaction value, network fee, and nonce) with the other components are stripped out (account balances, code, internal memory of smart contracts).

As explained by Vitalik Buterin, this compression results in significant size reduction, and therefore gas fees (although to varying degrees depending on the transaction type):

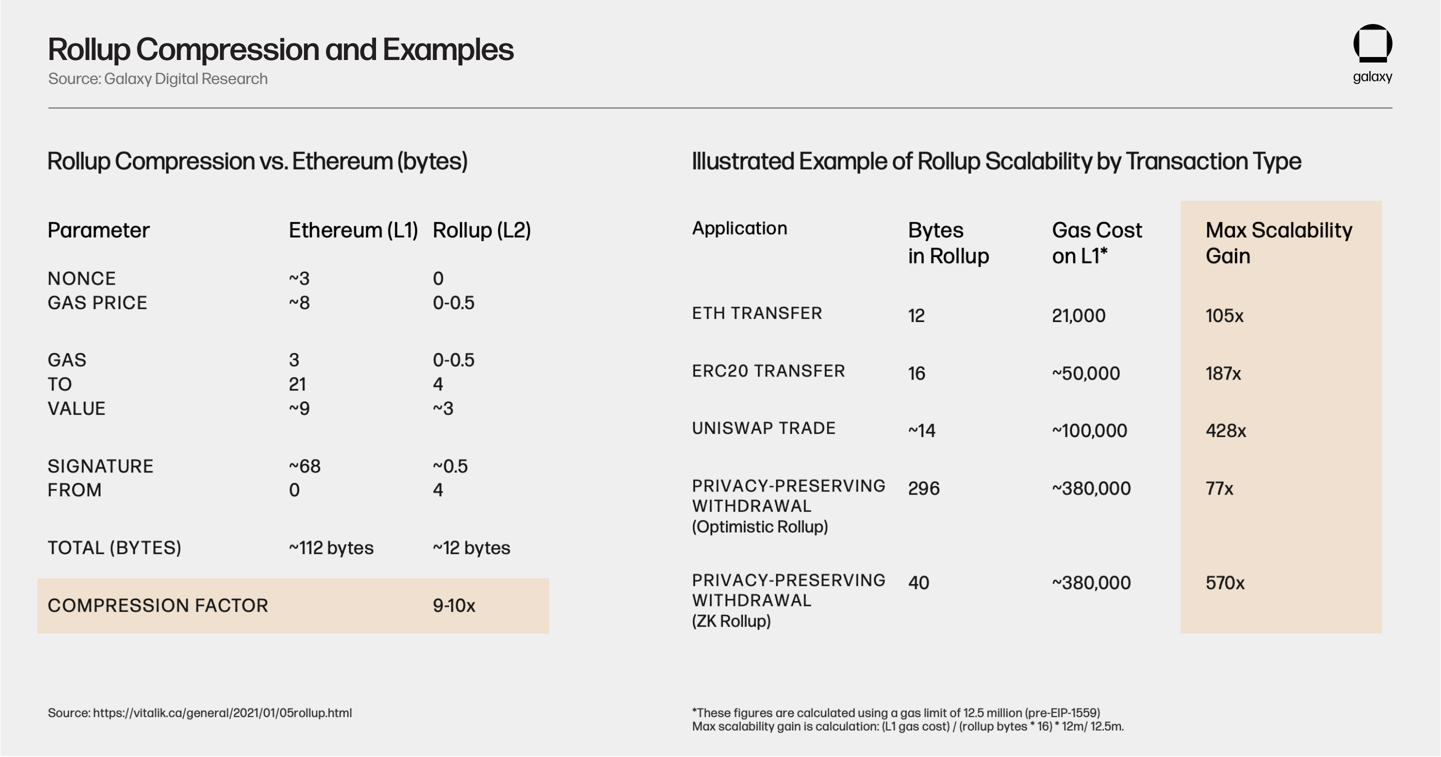

Rollup Compression and Examples

Source: Galaxy Digital Research

Rollup Compression vs. Ethereum (bytes)

Parameter	Ethereum (L1)	Rollup (L2)
NONCE	~3	0
GAS PRICE	~8	0-0.5
GAS	3	0-0.5
TO	21	4
VALUE	~9	~3
SIGNATURE	~68	~0.5
FROM	0	4
TOTAL (BYTES)	~112 bytes	~12 bytes
COMPRESSION FACTOR		9-10x

Source: https://vitalik.ca/general/2021/01/05rollup.html

Illustrated Example of Rollup Scalability by Transaction Type

Application	Bytes in Rollup	Gas Cost on L1*	Max Scalability Gain
ETH TRANSFER	12	21,000	105x
ERC20 TRANSFER	16	~50,000	187x
UNISWAP TRADE	~14	~100,000	428x
PRIVACY-PRESERVING WITHDRAWAL (Optimistic Rollup)	296	~380,000	77x
PRIVACY-PRESERVING WITHDRAWAL (ZK Rollup)	40	~380,000	570x

*These figures are calculated using a gas limit of 12.5 million (pre-EIP-1559)
Max scalability gain is calculation: (L1 gas cost) / (rollup bytes * 16) * 12m/ 12.5m.

Rollups typically come in two forms: (i) **optimistic rollups**, and (ii) **zero-knowledge** rollups. The primary difference between the two forms is their security models, which differ primarily around when the cost of proof generation/validation is paid. Optimistic rollups delay the cost of validation until after a dispute has been raised, while zk-rollups pay this cost upfront.

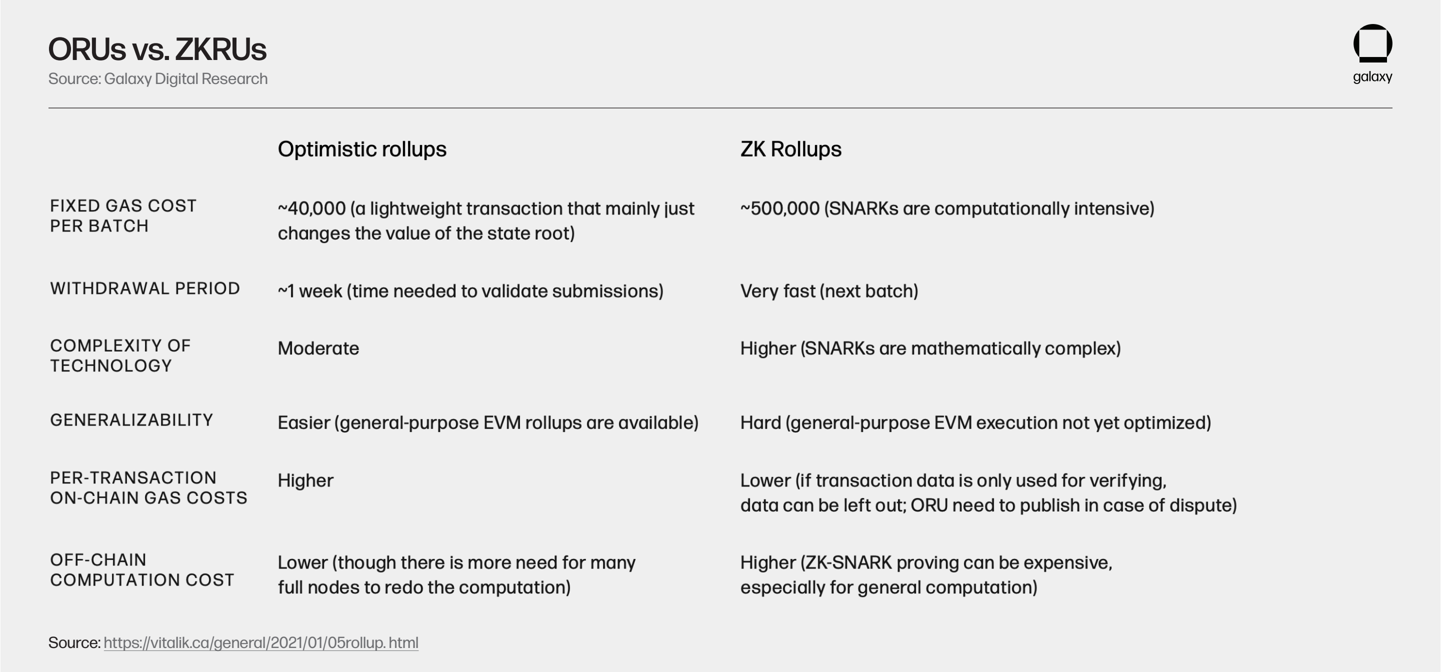

ORUs vs. ZKRUs

Source: Galaxy Digital Research

	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 (SNARKs are computationally intensive)
WITHDRAWAL PERIOD	~1 week (time needed to validate submissions)	Very fast (next batch)
COMPLEXITY OF TECHNOLOGY	Moderate	Higher (SNARKs are mathematically complex)
GENERALIZABILITY	Easier (general-purpose EVM rollups are available)	Hard (general-purpose EVM execution not yet optimized)
PER-TRANSACTION ON-CHAIN GAS COSTS	Higher	Lower (if transaction data is only used for verifying, data can be left out; ORU need to publish in case of dispute)
OFF-CHAIN COMPUTATION COST	Lower (though there is more need for many full nodes to redo the computation)	Higher (ZK-SNARK proving can be expensive, especially for general computation)

Source: https://vitalik.ca/general/2021/01/05rollup.html

Optimistic Rollups

Like the Plasma framework, optimistic rollups (ORUs) rely on **fraud proofs**, a security model where computation for transaction validation does not occur on the L1 mainnet unless the proof is disputed. Optimistic rollups derive their name because the batch of transactions submitted by the sequencer are *optimistically* assumed to be valid by default: only in the case of a dispute will the computation of each transaction included in the rollup block be executed on mainchain to determine whether fraud had occurred.

Depending on the protocol architecture, the aggregator can be a fixed central entity, a rotating selection, or a pool of bonded aggregators. As with most designs, centralizing this process into the hands of a single aggregator can provide a better user experience and faster confirmation times, but doing so can compromise the security and decentralization of the system.

Sequencers must run an ETH full node & a full L2 node to produce the L2 state. After a block is posted, verifiers have a dispute period to check the accuracy of the sequencer-published batch of state transitions (typically one week).

If no challenge is issued before the dispute window ends (the optimistic case), the published transactions are then finalized on-chain and can no longer be disputed.

If any of the state transitions within the batch are disputed or if any discrepancies are identified, then any participant may post a fraud proof against the non-finalized block. The correct proof is determined by executing the transaction of the state transition on-chain and comparing the correct state root against the root asserted by the sequencer. If they do not match, then the invalid state transition is cancelled and depending on the protocol implementation, the state transitions after the invalid one may also be cancelled or pruned.

Optimistic Incentives

Sequencers are required to put down a large deposit as a security bond attached with each submitted rollup block. Depending on who the honest actor is, the incentives will go towards either the sequencer or the verifier:

Optimistic case: If no challenge is issued, the security bond can then be returned to the sequencer, along with a reward from a portion of transaction fees, which mostly go towards covering the costs associated with the L1 “calldata.”

Non-optimal case: If verifiers If the sequencer is determined to be at fault, the sequencer’s security bond is slashed. A portion of the deposited collateral is also burned to prevent free miner griefing, and the remainder then goes to the honest actor for posting the fraud proof.

Fraud Proof Flow Chart

Source: Galaxy Digital Research

- STEP 1
  Sequencer batches compact transactions and submits rollup block to be posted on L1
- STEP 2
  7-day period for validators or network users to challenge the submitted block
- STEP 3A
  If no disputes occur, then the block is confirmed on L1
- STEP 3B
  If dispute is raised, then block enters "resolution" (re-executing transactions or interactive proving)
- STEP 4
  The fraudulent party's deposit is slashed, and the honest party receives a portion of the funds

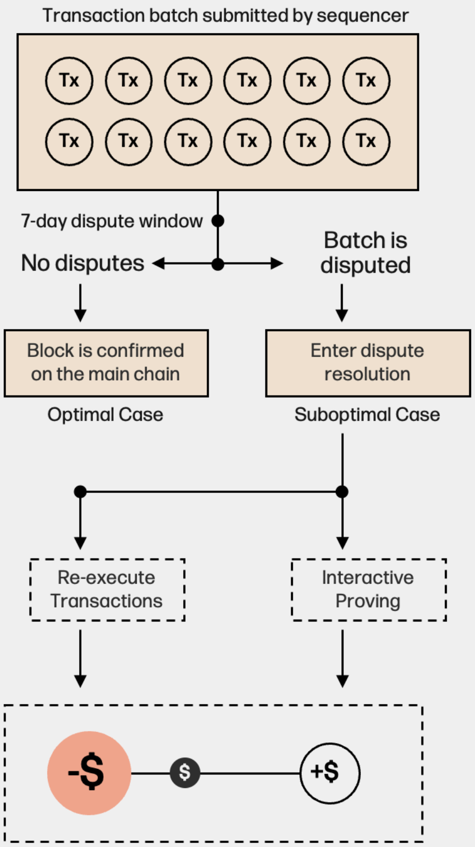

Malicious party loses their deposit, part of which is burned; the remainder goes to the honest actor.

Dispute Resolution

Current methods for handling disputes are **re-executing transactions** and **interactive proving**. Most ORU protocols use the re-execution method. Disputed rollup blocks will post a state claim for each transaction included in the block, which the L1 would then arbitrate by replaying the execution of an entire transaction to compare to the sequencer's state claim. Consequently, this setup in the unoptimistic case is more expensive than direct execution on the L1 and it requires imposing a lower gas limit than that on the L1.

The other fraud proof methodology—interactive proving—looks to implement a more efficient design with higher gas limits and is being explored by Arbitrum. The framework of interactive proving is to do as much off-chain work as possible to pinpoint the disputed execution step so that work is minimized for the L1. When a claim is disputed, the sequencer will go back-and-forth with the challenger. The sequencer will post two claims for each half of the initial claim; the challenger then picks one of the two to challenge, cutting the dispute size in half. The sequencer then posts two claims based on the latest claim chosen by the challenger, and this process continues until a single execution step is identified. Since the challenger can detect the validity of the sequencer's claim off-chain, the L1 doesn't have to replay the entire transaction and instead only re-executes one instruction. This process must be completed in the allotted time limit for maximal efficiency.

Benefits to interactive proving include lower costs with a smaller data footprint on L1, and a higher contract size limit. The downside is a much more burdensome resolution process for the involved participants. Since it is a multi-round process, interactive proofs generate the fraud proof slower than the re-executing method. It is a complex setup requiring the sequencer to reliably stay online to efficiently engage in the back-and-forth dispute process.

Ideally with ORUs, all submitted blocks would go undisputed to keep computation of fraud proofs off the L1. With an efficient dispute resolution design to detect malicious actors, attempts to commit fraud should be sparce and infrequent. The L1 network would then only have to witness the data, thereby preserving the L1 network capacity and increasing scalability.

Optimistic Rollup Advantages

- **Speed and security at low-cost in the optimal case.** ORUs provide fast confirmations (can be under one second) In the optimal operating environment (i.e. without any disputes), ORUs offer a low-cost solution that puts a low data and computing workload on-chain.
- **Maintains L1 security and protects data availability.** ORUs inherit the security of the L1 for arbitrating disputes while also preserving data availability so that any party can access and verify the data for off-chain results.
- **Equipped for general purpose smart contracts.** Having EVM-compatibility enables existing apps deployed on Ethereum to be migrated over to an ORU environment, enabling rapid growth of the ecosystem. ORUs also provide generalizability to handle general-purpose smart contracts, providing more functionality compared to other L2 frameworks that are limited in supporting general purpose computation.

Optimistic Rollup Disadvantages

- **Latency / lengthy fraud proofs and withdrawals.** The exit game for a user to withdraw funds may be relatively data-intensive and lengthy, limiting the usability of the network. For participants to withdraw their assets held inside of an optimistic rollup, they must wait a week to withdraw to provide sufficient time for one to validate or challenge the submitted transaction batches.
- **Complexity leads to centralization, raising other security issues.** ORUs assume there is always a live, honest validator. Given the complexities of the fraud proofs, most protocols will maintain some level of centralization to maximize efficiency. The game theory-based security model also introduces other potential attack vectors (e.g. role of the sequencer could theoretically be abused if they process enough blocks and become profitable enough to overcome having their stake slashed).
- **Throughput limitations.** Given the dispute resolution mechanism, the maximum throughput of ORUs is limited by the amount of data that can be published on L1. Using interactive proving for dispute resolutions would have a higher throughput limit in place compared to re-executing transactions.

Activity on Optimistic Rollups

ORUs saw moderate usage this year until the launch of Arbitrum in August. Before then, most of the TVL was locked up in derivatives liquidity protocol Synthetix on Optimism, which went live in January. Following the public launch of Arbitrum One, TVL skyrocketed to over $2bn.

- **Optimism** (fka the Plasma Group) was the first generalized rollup protocol to gain popularity. The Optimistic VM was designed to reuse much of the ETH tooling to closely resemble the EVM. Optimism has taken a gradual release process with whitelist restrictions rather than a complete public main launch and launched several notable projects (with limitations) in recent months including Uniswap V3, 1inch, and Lyra Finance – a protocol for trading options and the first Optimism-native project launched. Initially, the sequencer will be centralized but Optimism plans to adopt an auction methodology for choosing the sequencer and intends to use the captured MEV for public goods funding.
- Offchain Labs' **Arbitrum** is designed as a multi-round rollup for dispute resolutions for more compressed fraud proofs that put less data on the base layer, enabling higher transaction throughput. The Arbitrum team has been focused on supporting developers through compatibility with ETH tooling in several languages including YUL, Solidity, and Vyper. In contrast with Optimism, Arbitrum had a complete public mainnet launch, occurring at the end of August, rather than a gradual release process. The sequencer role will be centralized at start and Offchain Labs intends to progressively decentralize the role, although details are still pending. In less than two weeks since its launch, TVL on Arbitrum grew parabolic to over $2.2bn.
- Other ORU projects: Fuel, Metis, Celestia, Boba Network (fka OMGX), layer2.finance (Celer)

Total Value Locked in ORUs - YTD (USD)

Source: Galaxy Digital Research

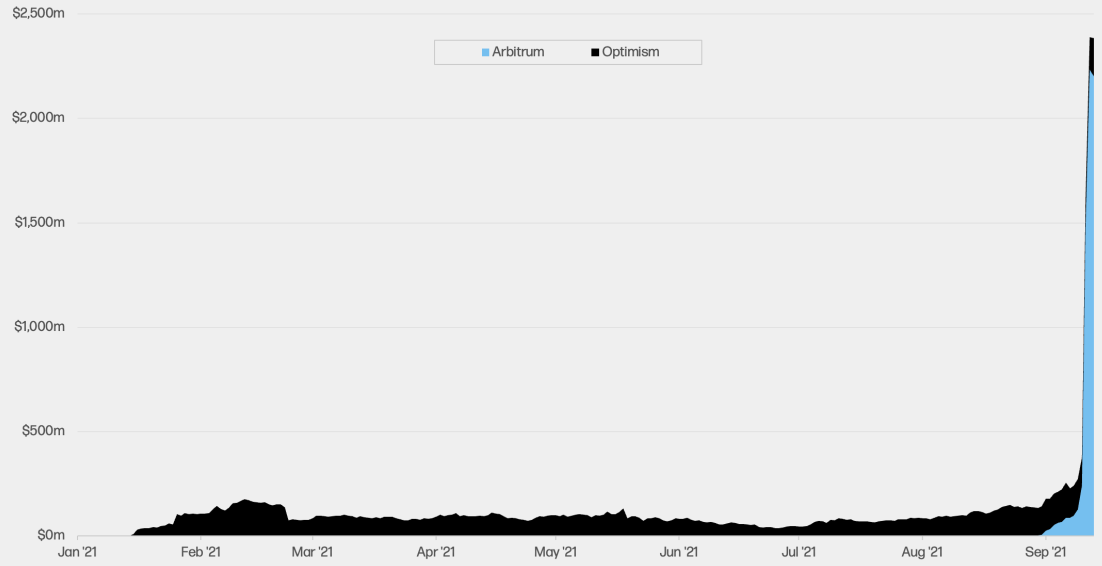

Data: L2Beat

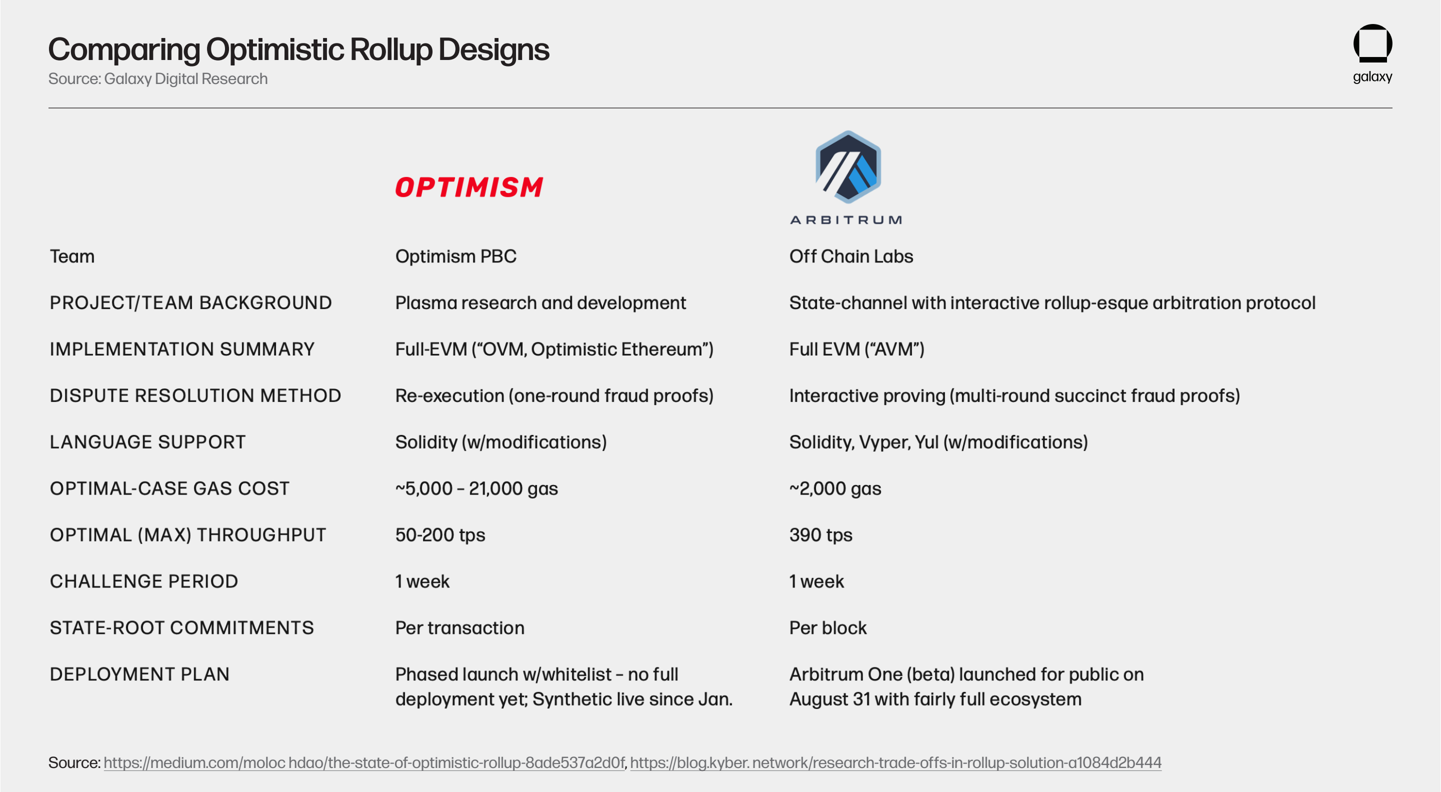

Comparing Optimistic Rollup Designs

Source: Galaxy Digital Research

galaxy

	OPTIMISM	ARBITRUM
Team	Optimism PBC	Off Chain Labs
PROJECT/TEAM BACKGROUND	Plasma research and development	State-channel with interactive rollup-esque arbitration protocol
IMPLEMENTATION SUMMARY	Full-EVM ("OVM, Optimistic Ethereum")	Full EVM ("AVM")
DISPUTE RESOLUTION METHOD	Re-execution (one-round fraud proofs)	Interactive proving (multi-round succinct fraud proofs)
LANGUAGE SUPPORT	Solidity (w/modifications)	Solidity, Vyper, Yul (w/modifications)
OPTIMAL-CASE GAS COST	~5,000 - 21,000 gas	~2,000 gas
OPTIMAL (MAX) THROUGHPUT	50-200 tps	390 tps
CHALLENGE PERIOD	1 week	1 week
STATE-ROOT COMMITMENTS	Per transaction	Per block
DEPLOYMENT PLAN	Phased launch w/whitelist - no full deployment yet; Synthetic live since Jan.	Arbitrum One (beta) launched for public on August 31 with fairly full ecosystem

Source: https://medium.com/moloc hdao/the-state-of-optimistic-rollup-8ade537a2d0f, https://blog.kyber. network/research-trade-offs-in-rollup-solution-a1084d2b444

zk-Rollups

Rather than going through this lengthy challenge game with fraud-proofs, ZKRUs provide a much quicker validation period through its **validity proof** security model, which generates the proof upfront as soon as blocks are submitted. The proof can then be quickly verified on the L1, allowing for fast user withdrawals.

Relayers (sometimes called *provers*) assume the role of the aggregator for ZKRUs. Relayers aggregate the rollup transactions to be submitted to the mainchain. In contrast to sequencers, relayers have the added responsibility of performing all the computations to generate the **zk-SNARK proof** (zero-knowledge Succinct Non-interactive ARgument of Knowledge), which only shows a portion of the resulting hash but not the actual data itself. The SNARK proof compares a snapshot of account values on blockchain before the transfers to a snapshot of the blockchain after the transfers.

The network verifiers can then validate only the submitted proof without the need to verify all of the embedded transactions (i.e. "zero knowledge" of the entire data is needed).

Generating a SNARK proof for every state transition makes it impossible for operators to commit an invalid or manipulated state, so funds cannot be stolen by operators. A user can be confident of mainnet verification and finalization of the proof after it has been submitted. The waiting period for users to withdraw their funds from L2 to the L1 is simply the time needed for the next batch submission.

However, depending on the setup, SNARKs usually require a trusted set up. Since SNARK proofs only represents the delta of the blockchain state, the initial setup is assumed to be a trusted state – but this cannot be proven by participants. Only a select group of the developers know with certainty, which undercuts the notion of decentralization. Other ZKRU implementations may rely on other proof designs to improve on the trusted setup in SNARKs including:

- **PLONK** (Permutations over Lagrange-bases for Oecumenical Noninteractive arguments of Knowledge) still requires a trusted setup but enables multiple parties to participate in the trusted setup, an improvement to the SNARK procedure. However, the added security demands a larger proof size relative to SNARKs.
- **zk-STARKs** (Scalable Transparent ARguments of Knowledge) remove the need for a trusted setup using hash functions to create trustless, verifiable computation systems. STARKs are a newer and more complex proof technology compared to SNARKs and require even fewer security assumptions than PLONK. The trustless setup comes at the cost of larger proof sizes, which require more gas and longer verifications.

Validity proofs are also employed by another framework called **Validium**, which uses a hybrid design combining aspects of ZKRUs and Plasma—basically Plasma with SNARKs or ZKRU with off-chain data. Recall, ORUs were an evolution of Plasma that primarily aimed to solve Plasma's data availability problems given potential risk of operators to freeze user funds. Validium revisits the idea of off-chain data storage to provide a more economical framework with lower costs and higher throughput compared to ZKRUs. While it still potentially subjects users to withholding of data, this model could be more fitting for certain use cases that are accepting of lower trust assumptions and requiring higher throughput capacity.

zk-Rollup Advantages

- **Short finality time and fast withdrawals.** With the proof submitted upfront, a user can be confident in verification after submitting a transaction. Given the quick assurances, the waiting period for users to withdraw their funds from L2 to the L1 is simply the time needed for the next batch.
- **Strong security guarantees and native privacy options.** If confirmed at initial setup, ZKRUs are always in a valid state. Operators cannot commit an invalid state and cannot steal user funds. ZKRUs inherently promote privacy through SNARK technology, also leveraged by privacy coin Zcash, which may be useful for trading strategy privacy.

zk-Rollup Disadvantages

- **Burdensome proof generation.** Operators generate SNARK proofs for every state transition. Proofs are computationally intense and come with a high fixed gas cost per batch, so the proof generation economics of ZKRUs still must be optimized. STARK and PLONK setups require even larger proof sizes.
- **Developer onboard difficulty.** Ethereum developers cannot immediately move their apps into ZKRUs without significant additional training since ZKRUs introduce new data structures and are not 100% EVM byte-code compatible. All variations of ZKRUs currently require rewriting contracts in a new language so developers need a higher degree of specification to write smart contracts.
- **General-purpose smart contract support still limited.** ZKRUs in their current form are not yet equipped for general computation, only supporting basic functions like payments and token exchange. Most ZKRU protocols are working on the compiling needed for EVM-bytecode but so far, no chain that has been deployed in production.
    - zkSync 2.0 launched with zkEVM testnet in May: "Our VM, zkEVM, is not an EVM 1:1 replica, but instead aims to be able to run 99% of contracts written in Solidity and maintain its same behavior, such as during reverts and exceptions. Simultaneously, the zkEVM is written to be efficient in a circuit to produce zero knowledge proofs." Once again, we are excited to announce, after months of hard work: the instruction set of the zkEVM has been finalized and implemented in circuit and in the execution environment."
    - Starkware uses its own native smart contract language Cairo—a Turing-complete STARK-friendly CPU architecture. The team at Nethermind just released its demo of an EVM-to-Cairo transpiler called Warp, bringing Solidity ERC20 contracts to StarkNet, Starkware's ZKRU product. The next milestone on the Warp roadmap is to compile an AMM, such as Uniswap, to StarkNet.

Activity on zk-rollups

As use cases of ZKRUs have mostly been limited mainly to payments and token transfers so far, TVL has been relatively low and concentrated in Loopring through the first half of the year. In recent months, TVL on ZKRUs grew to over $400m with most of it landing on dYdX for perpetual contracts.

Total Value locked in ZKRUs ($USD)

Source: Galaxy Digital Research

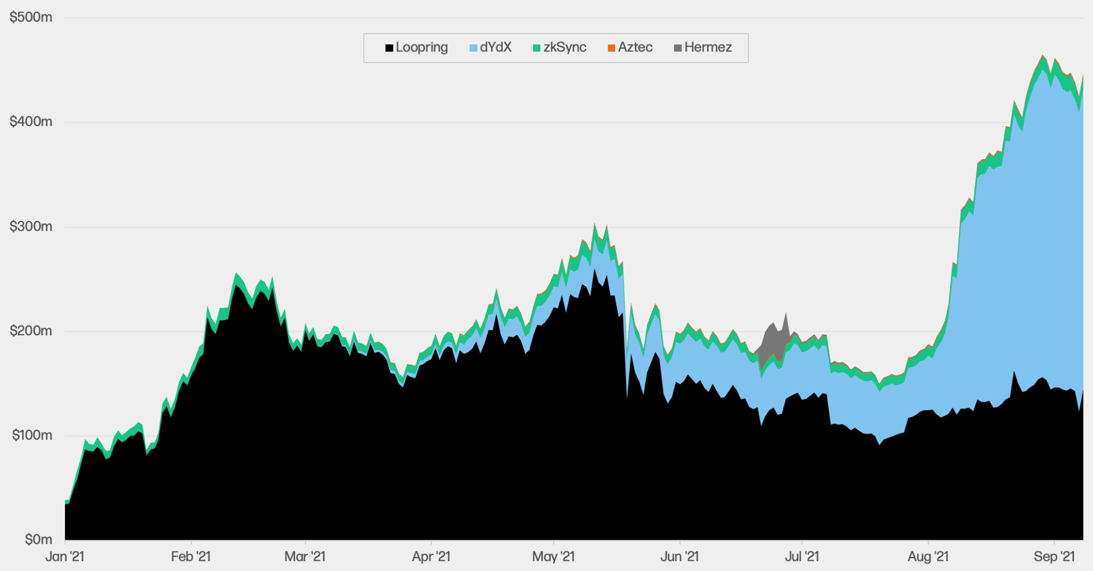

Data: L2Beat

- The first deployment of ZKRUs to the ETH mainnet was in February 2020 by Loopring, an exchange and payment protocol. **Loopring** uses a SNARK construction and claims the protocol can settle over 2k TPS. Loopring primarily offers two products, the Loopring Wallet and the Loopring Exchange, an L2 orderbook and AMM DEX, respectively. Loopring is collaborating with StarkWare for dAMM—a cross-L2 AMM solution for liquidity fragmentation—and recently added support for NFT minting and transfers.
- StarkWare's **StarkEx** deployed on mainnet in June 2020 with DeversiFi, a ZKRU-native DEX, and has two modes to support ZKRUs (on-chain data) or Validium (off-chain data). StarkEx, which is programmed using Cairo language, is largely used for DEXes and derivatives, and it has measured over 9k TPS for trades. dYdX, the perpetual swap protocol, launched on L2 with Starkware in February 2021 as another L2-native project. StarkWare has also teamed up with Immutable to build Immutable X, the first L2 scaling platform on Ethereum built strictly for NFTs.
- Matter Labs' **zkSync** is constructed with PLONK technology for a universal trusted set-up (instead of an application-specific trusted setup with SNARK). zkSync deployed on mainnet in June 2020 for simple payments and is working towards adding support for smart contracts. zkSync introduced zkPorter, a Validium-based system with off-chain data availability that complements the ZKRU side to achieve higher scalability with lower fees. Off-chain data availability in zkPorter is secured by "guardians" in a PoS setup.
- **Hermez** Network launched on mainnet in March 2021 for payments and token transfers. With Hermez, coordinators (like relayers/provers) collect and process the transactions that enter a rollup. The process for selecting a coordinator occurs via an auction and is decided for each 40-block period (about 10 minutes). In August, Polygon and Hermez announced they would be merging in the industry's first token merger of two blockchains. The Hermez team is committed to preserving decentralization and the integration with Polygon enables them to leverage the established Polygon platform including its brand and users so that the team can focus strictly on the technical development of their zkEVM.

L2 Comparison

Each of these off-chain scaling designs come with different trade-offs. There are nuances between each of the protocols within each classification, but broadly speaking:

- **State channels** maximize for transaction throughput, cost, and latency but have the drawbacks of capital inefficiency (requires users to fund each channel and pay on-chain fees to open/close channels), liveness assumption, and limited support beyond payments.
- **Sidechains** have maximized for flexibility to achieve faster innovation and to quickly deliver a usable scaling environment with full EVM-compatibility, which comes at the cost of giving up the security guarantees of the L1 and a higher trust requirement with a centralized operator.
- **Plasma** chains leverage the securities of the base layer and deliver fast and extremely low-fee transactions. However, Plasma overcompensates on cost-minimization and scalability at the expense of data availability and censorship-resistance while also suffering from lengthy withdrawals with fraud proofs.
- **Optimistic rollups** solve transparency/security/data availability inherent in sidechain and plasma chains but pay higher fees for it. They are the first true L2 to support generalized computing but have drawbacks including a lengthy withdrawal period from fraud proofs, liveness assumption, relatively low potential throughput, and depending on the protocol design, higher levels of centralization which may come with separate security considerations.
- **ZK-rollups** use validity proofs which are inherently privacy and security-preserving and avoid the lengthy withdrawals of fraud proofs. Relative to ORUs, ZKRUs have a larger L1 footprint, are tougher/costlier to implement, require a higher degree of developer specification, and are still largely under development.

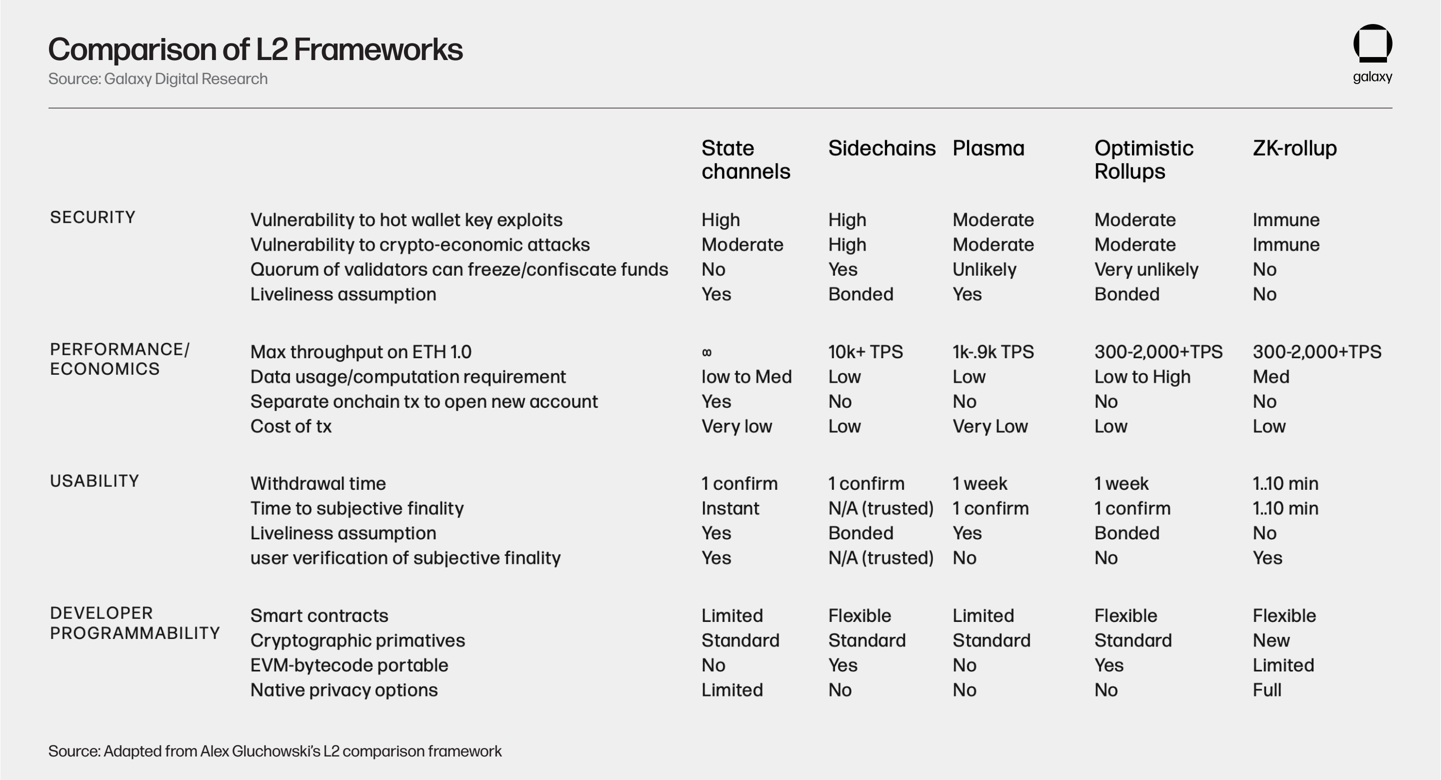

Comparison of L2 Frameworks

Source: Galaxy Digital Research

		State channels	Sidechains	Plasma	Optimistic Rollups	ZK-rollup
SECURITY	Vulnerability to hot wallet key exploits	High	High	Moderate	Moderate	Immune
	Vulnerability to crypto-economic attacks	Moderate	High	Moderate	Moderate	Immune
	Quorum of validators can freeze/confiscate funds	No	Yes	Unlikely	Very unlikely	No
	Liveliness assumption	Yes	Bonded	Yes	Bonded	No
PERFORMANCE/ ECONOMICS	Max throughput on ETH 1.0	∞	10k+ TPS	1k-.9k TPS	300-2,000+TPS	300-2,000+TPS
	Data usage/computation requirement	low to Med	Low	Low	Low to High	Med
	Separate onchain tx to open new account	Yes	No	No	No	No
	Cost of tx	Very low	Low	Very Low	Low	Low
USABILITY	Withdrawal time	1 confirm	1 confirm	1 week	1 week	1..10 min
	Time to subjective finality	Instant	N/A (trusted)	1 confirm	1 confirm	1..10 min
	Liveliness assumption	Yes	Bonded	Yes	Bonded	No
	user verification of subjective finality	Yes	N/A (trusted)	No	No	Yes
DEVELOPER PROGRAMMABILITY	Smart contracts	Limited	Flexible	Limited	Flexible	Flexible
	Cryptographic primatives	Standard	Standard	Standard	Standard	New
	EVM-bytecode portable	No	Yes	No	Yes	Limited
	Native privacy options	Limited	No	No	No	Full

Source: Adapted from Alex Gluchowski's L2 comparison framework

Finding the Optimal Use Cases

Different applications may want to optimize for different things such as transaction speed, transaction cost or security, but this usually means having to make a sacrifice in another factor.

For app developers, EVM-compatibility and programmability have proven to be the most desired traits so far. Developers behind existing Ethereum-native DeFi applications want to deploy across other platforms with simple portability as opposed to rewriting smart contracts in a new programming language or for new data structures in another protocol. This approach has found fertile ground among DeFi applications, for whom an immediate need has been scalability and offering their application in an environment with lower fees. That said, the greenfield opportunity for developers is massive and L2s have attracted new developer teams for applications that were not possible on L1s like options, derivatives, and gaming. We have already seen the launch of several rollup-native applications (e.g. Loopring, dYdX, DeversiFi, and Lyra Finance on Optimism) and expect to see many more to come.

With their programmability and ease of deployment, sidechains emerged as the first readily available off-chain scaling option equipped to handle general smart contracts from Ethereum, and they have met the immediate needs of the greater community. But now as rollups and other options emerge, each of these frameworks should see some degree of verticalization along various use cases. As potential examples:

- **State channels.** Lightning has proven to be effective for simple transactions (micropayments / commerce), cross-border remittances, and instances with multiple or recurring transactions (streaming, subscriptions, gaming).
- **Sidechains.** Security levels vary across protocols, but sidechains are generally good for small-value transactions which may not require the same guarantees as high-value transactions. Sidechains have also been employed by enterprises for internal transactions, such as small-cap centralized exchanges, or as testnets (detailed below).
- **Optimistic rollups.** ORUs are equipped to handle general computation, and transactions requiring strong safety guarantees. However, those looking for the cheapest transactions or fast liquidity for low-value transfers may find more fitting solutions on payment channels or sidechains.
- **ZK-rollups.** Only basic functions like payments and exchanges have seen any traction so far on ZKRUs while the technology is still under development. ZKRUs and Validium have been identified as a scaling platform for NFTs (per Immutable X and Loopring).

Centralized Exchanges are sidechains. Many large consumer-facing corporates have opted for off-chain scaling solutions to bypass the high gas fees of on-chain operations and to provide the most frictionless experience for users. This typically involves maintaining an internal ledger of transactions that can mirror the design of sidechains. For example, consider centralized exchanges (CEXes) offering customers offering users buy/sell/hold services for digital assets. With crypto-native exchanges like Coinbase and Gemini, user activity regarding account openings/closings or buying/selling levels is not reconcilable using on-chain data. Only when a user chooses to withdraw his or her funds from the exchange into an external wallet would the transaction be logged on-chain. Some fintechs (e.g. PayPal, Robinhood, SoFi, and Cash App) have followed similar playbooks to offer these trading services to their users through the enlistment of a digital asset custodian.

These off-chain operational models (e.g. maintaining internal ledgers) have been one of few economical ways to provide users these services while also preserving the same streamlined experiences that users are accustomed to. Users are presented with more straightforward fee schedules (vs. variable transaction costs based on network congestion for on-chain transactions) with the near-instant settlement times. This method achieves the required scalability and does not place undue reliance on a relatively untested blockchain network. However, these product offerings and user benefits do come at the expense of centralization and lack of transparency – which means that transactions can be censored or blocked at will, funds can be frozen and seized, and there is no data availability on-chain to revert to in the case of disputes.

Discussion

The L2 World Today

We are presently in a multi-chain reality

Until mid-March, over 95% of TVL across all chains belonged to Ethereum. By May, this number had dipped to under 75% after average gas prices on Ethereum had spiked dramatically to nearly $70, creating an uneconomical transaction environment and driving users to alternative and lower-cost smart contract platforms.

At this point, general purpose L2 platforms were not yet available so users in search of lower fees flocked primarily to Binance Smart Chain and Polygon, which were opportunistically ready to meet those needs. The rise of transaction count on Polygon coincided with a drop in transactions on BSC and ETH.

Although centralization is intentional for BSC / Polygon, users have been accepting of that. BSC and Polygon been able to deploy quickly to meet the needs of the Ethereum DeFi community as they have consciously architected a centralized design to optimize for scalability and useability. Of course, this means they are reliant on their own consensus code with lower security guarantees compared to the Ethereum base layer. The massive influx of users and activity onto these platforms suggests that the incentives were strong enough to overcome this trade-off.

TVL Across All chains (2021)

Source: Galaxy Digital Research

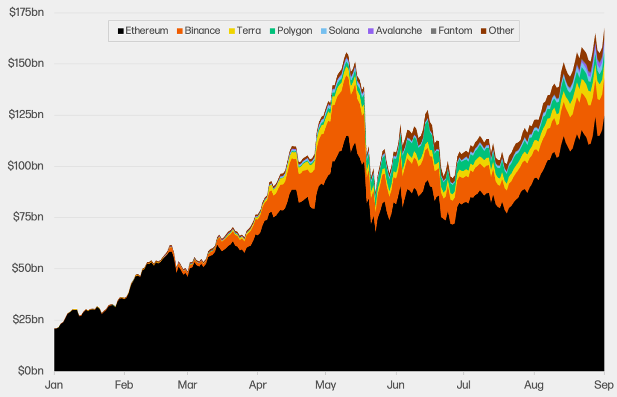

Data: DeFi Llama

TVL Share Across All chains (2021)

Source: Galaxy Digital Research

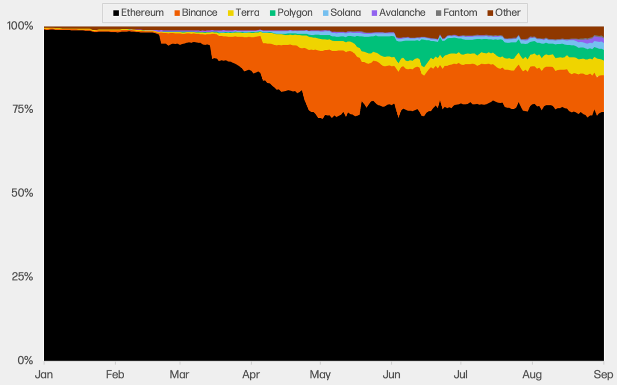

Data: DeFi Llama

Daily Tx - YTD

Source: Galaxy Digital Research

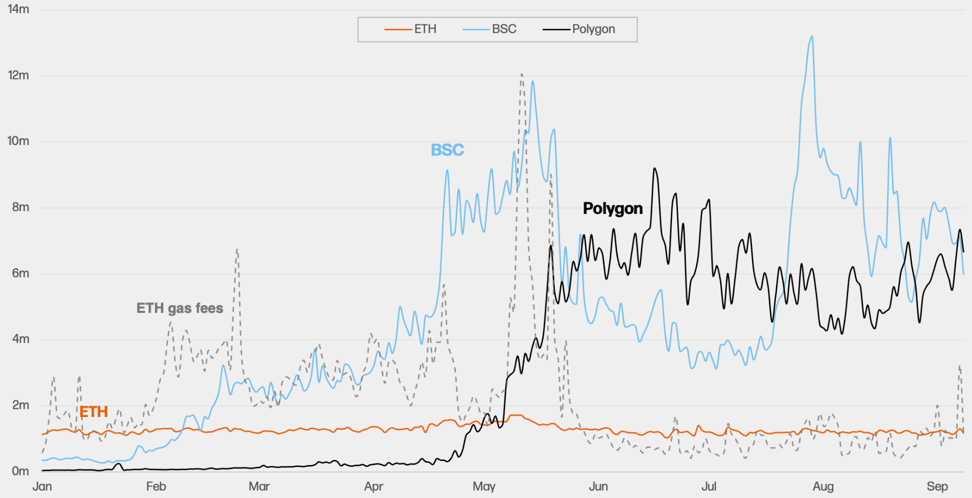

Data: Etherscan, PolygonScan, BscScan

Applications should strive to meet user demand across these chains

As competition across these scaling blockchains shakes out and ecosystems are being built across each execution environment, it is important for these applications to meet user demand wherever it may be. The open-source nature of projects has spurred a higher level of innovation, but this also entails more competitive risks, so application teams have additional defensive considerations. If committed to only one chain, applications risk losing users that migrate to other platforms or potentially getting forked and losing out on the revenue opportunity.

Some Ethereum-native applications have already committed to expanding across multiple chains including other L1s, L2s, and the sidechains/bridges in between. For example, Ethereum-native DeFi blue chips—like Aave, Uniswap, and Curve—have been deployed across Polygon, Optimism, or Harmony.

On average, dapps on our selected list average 3.8 deployments across different platforms (note: this does include committed deployments that are not yet fully live (e.g. Arbitrum has yet to fully open to the public while Aave and Curve have committed to Avalanche). TDEhis has also translated to more DeFi activity on non-Ethereum chains.

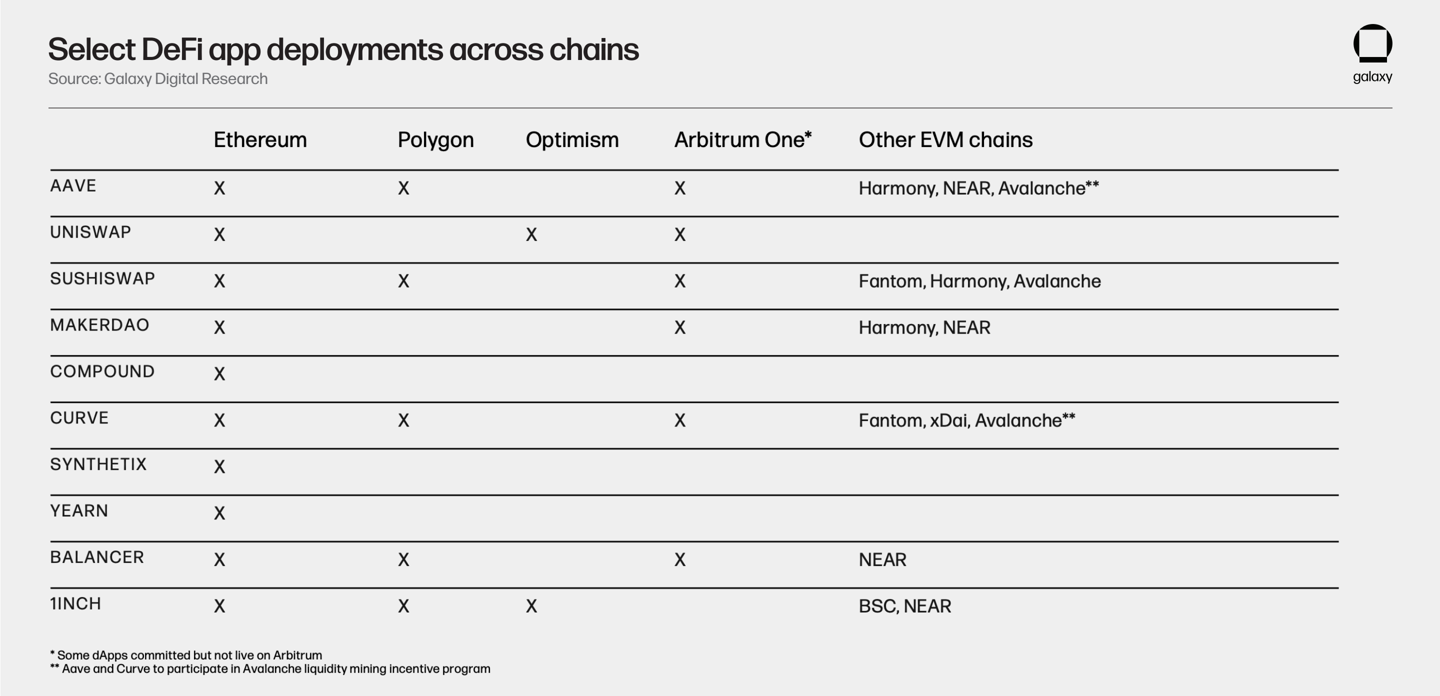

Select DeFi app deployments across chains

Source: Galaxy Digital Research

	Ethereum	Polygon	Optimism	Arbitrum One*	Other EVM chains
AAVE	X	X		X	Harmony, NEAR, Avalanche**
UNISWAP	X		X	X	
SUSHISWAP	X	X		X	Fantom, Harmony, Avalanche
MAKERDAO	X			X	Harmony, NEAR
COMPOUND	X				
CURVE	X	X		X	Fantom, xDai, Avalanche**
SYNTHETIX	X				
YEARN	X				
BALANCER	X	X		X	NEAR
1INCH	X	X	X		BSC, NEAR

* Some dApps committed but not live on Arbitrum
** Aave and Curve to participate in Avalanche liquidity mining incentive program

Rollups are a better technological path forward than sidechains. But today, they still require significant centralization to operate and, in general, are nascent.
Rollups are ideologically-motivated (decentralized, privacy preserving, censorship-resistant) – but at this point have centralized setups for sequencers and operators given the technical complexities around protocol development and certain safeguards in place for operators to throttle the network or to make planned network updates. But in contrast to BSC/Polygon, ordering transactions is a technically demanding process and requires an efficient operator that can quickly implement network updates or react to unexpected network interruptions.

Among rollups, ZKRUs have been identified as the technologically superior framework, but they aren't quite ready.
Vitalik agrees that solutions like Polygon and ORUs are filling an important and pressing void for DeFi, but believes eventually ZKRUs will win the L2 scaling wars: "In general, my own view is that in the short term, optimistic rollups are likely to win out for general-purpose EVM computation and ZK rollups are likely to win out for simple payments, exchange and other application-specific use cases, but in the medium to long term ZK rollups will win out in all use cases as ZK-SNARK technology improves."[2]

Consensus generally seems to agree with the Ethereum founder that ZKRUs represent a potential superior design compared to ORUs – the technology just still requires additional development:

- Polygon co-founder Mihailo Bjelic: "We consider ZK cryptography the single most important strategic resource for blockchain scaling and infrastructure development, and we have a clear goal of becoming the leading force and contributor in this field in years to come."[3]
- Matter Labs founder Alex Cluchowski: "Optimistic Rollup is great news for ZK Rollup. The transition to L2 scaling requires significant changes in wallets, oracles, dapps and user habits. Optimistic Rollup can help to prepare the ecosystem for this move, bringing scale to those dapps that cannot yet be built on ZK Rollup today. This will give ZK Rollup time to mature and make its adoption completely seamless, while maintaining Ethereum's growth momentum."[4]
- Digital asset derivatives exchange Interdax: "Optimistic Rollups can support both simple payments and complex smart contracts, and 80% of the Ethereum Virtual Machine (EVM) tooling can be transferred over. Given that most costs on Ethereum are complicated, Optimistic Rollups are seen as an immediate solution. On the other hand, it is more difficult to port over smart contracts seamlessly from Ethereum's main chain to ZK-Rollups. As a result, ZK-Rollups are viewed by Ethereum as a much more promising solution in the long term."[5]
- Ernst & Young: "Based on EY experience, ZK-Optimistic roll-ups are currently among the most effective in balancing security incentives and mathematical efficiency for running private transactions on the public Ethereum network. As we have in the past, we are again contributing this code into the public domain to speed up enterprise adoption of this technology."[6]

The L2 World Tomorrow

Protocol trade-offs should not be viewed in isolation; the future may not be ZK-dominated

While many view ZKRUs as the eventual L2 savior, the framework is far from proven and could see many difficulties in adoption. The technical requirements for writing existing smart contracts on ZKRUs is relatively difficult, while developers have already demonstrated their interest in straightforward migration of their code using more centralized solutions like Binance Smart Chain and Polygon. These platforms have made the intentional and pragmatic trade-off to be first to market with a useable scaling solution at the expense of decentralization and censorship-resistance. But having a centralized controlling entity provides the flexibility that is needed to quickly adapt to changes which is paramount in this rapidly evolving environment (e.g. Polygon-Hermez token merger). ZKRUs will have to be meaningfully better to justify the switching costs from the existing solutions. ZKRUs promise users a decentralized, privacy preserving, censorship-resistant platform – but it's not clear that ideologically-driven attributes will be sufficient to win over users that become accustomed to fast transactions for well-under one penny's worth. The longer ZKRUs remain non-fully compatible with EVM / Ethereum tooling and out of production, the harder it will be to get the DeFi ecosystem to migrate onto their technology, especially as L1s implement network upgrades, alternative L1s (e.g. Solana, Cosmos, NEAR, Avalanche, Terra, Fantom) continue to develop and attract users, and other participants bring more useability to other L2 solutions.

Therefore, the trade-offs between each of the protocols should not be viewed in insolation nor should they be viewed as a static metric. Protocols do not have to do all the heavy lifting on their own in addressing all the scaling limitations, on-ramp challenges, and UX difficulties with usage. With regards to ORUs and other fraud proof-based protocols, the week-long withdrawal period projects will be less of a headwind going forward as liquidity pools step in or as application-level bridges (e.g. Connext cBridge, Celer, Hop) are built to complement the protocol-level bridges. In addition, users may be less inclined to withdraw their funds over time, opting to park their funds in L2 as the ecosystem grows. ZKRUs must attract developer mindshare and user growth before the competitive advantages of the technology and switching incentives are eroded away, making the deployment timeline critical for longer-term success.

As multi-chain universe expands, cross-chain bridges are the next infrastructure frontier

As it stands today, most of these off-chain protocols have been operating somewhat in isolation from the L1 and from other L2s. The time required to bridge assets between L1<>L2 or to exit from L2<>L1 have been some of the highest points of user frictions. These time constraints may be doubled when moving from L2<>L2, which comes with the stepwise procedure of L1<>L21<>L1<>L22. But just as technological improvements in traditional payments have led to faster payments (e.g. checks, ACH, wire transfers) with more connectivity across network participants, crypto-based payments will grow increasingly faster across chains largely due to the development cross-chain and bridge infrastructure.

Over recent years, most of the x-chain efforts have been at the L1<>L1 led by Polkadot (relay chain / parachains), Cosmos (Hub & Spoke), and THORchain (cross-chain liquidity). EVM-compatibility will still be paramount as other L1s look to form bridges with Ethereum to transfer ERC20 tokens (e.g. new Avalanche Bridge launched last month; Neon Labs bringing an EVM solution to Solana testnet; along with Wormhole's mainnet launch as a x-chain messaging protocol supporting Solana, Ethereum, Terra, and BSC at start, and then Swim Protocol for x-chain transfers powered by Solana's Wormhole).

Now as L2s become more established, the x-chain attention can now shift to **L1<>L2** and **L2<>L2.** Just since we entered the second half of the year, we have seen several meaningful x-chain developments connecting both protocols and applications across different chains:

EVM-focused bridges

Source: Galaxy Digital Research

galaxy

Hop

Hop Protocol is a general token bridge for EVM-compatible blockchains. Hop went live in July, initially supporting Polygon and xDai. More recently, Hop launched Hoptism – a bridge connecting Optimism and intends to soon provide bridging with Arbitrum.

connext

Connext is another EVM-focused general token bridging protocol between L2s, providing non-custodial x-chain routing for ETH L2s. Connext launched Vector using generalized state channel technology and in august, introduced nxtp, a more contract-oriented system enabling contract-to-contract communications across chains that can then be applied towards virtual AMMs and route auctions.

CELER

Celer, the original generalized state channel protocol, went live with Celer cBridge in July to facilitate non-custodial, fast multi-hop value transfers between EVM-supported chains without having to go through the corresponding L1. cBridge supports transfers between Ethereum, Polygon, BSC, Arbitrum, Optimism among other chains.

The growing primacy of L2s that connect to multiple L1s and other L2s will be significant for several reasons:

- **The value that accrues from each platform layer to the L1 now also accrues to other platform layers.** With the infrastructure connecting each blockchain being built, the same way that value that accrues from each platform layer to the L1 now also accrues to other platform layers. This brings even more utility to the base layer assets and the protocol tokens operating at each level. Over time, the relationship between competing L1s will become less adversarial and more symbiotic as cross-chain connectivity is established – "ETH killers" may eventually turn into "ETH friends."
- **It frees up illiquid funds, creating a better UX and accelerating the velocity of money.** The underlying x-chain technology and bridges serve as the needed back-end to enable the front-end applications to abstract some of the cryptography that is unattractive to the average user for smoother onboarding.
- **Ultimately, as more users are onboarded across more blockchains, composability will be an increasingly important factor and is a necessary prerequisite to Web 3.0.** The value of NFTs might not make sense to most people now but what about when we become more immersed into Web 3.0 and the metaverse? Having the portability to move assets across platforms then closes the illiquidity discount assigned to these assets.

But we also do note that the technology is relatively immature and is not battle-tested. The designs of existing cross-chain protocols and cross-layer bridges differ dramatically as it relates to the custody arrangement, trust assumptions, swap design (AMMs using liquidity pools vs. lock-mint-burn), integrations with protocols and dapps, and other security or solvency assumptions. Bridges have typically implemented the lock-mint-burn design. These bridges have been targeted by hackers (e.g. Poly Network, THORChain #1, THORChain #2, AnySwap, ChainSwap). It speaks to how the technology still has to be more optimized at this point and with high value potentially at risk, users should be cautious. If security flaws are exposed or the protocol goes down, then there would be a negative impact with loss of capital, developer activity, and users. That said, users that want the functionality now may have to trade-off some levels of decentralization, security or cost.

Users will see higher rewards/incentives including in the form of covered L2 fees

UX has taken a backseat to DevEx, but with the back-end infra in place, that will change. So far at the off-chain protocol level, UX has generally taken a backseat to DevEx as protocols work through their testing phases – but now, as the back-end infrastructure is being established, more attention can be devoted to improving the experience for the end users. The same way that scaling smart contract platforms has catered to developers during the initial testing phases, these application developer teams must cater to the end users with the goal of providing the user-friendly UI/UX to onboard new cohorts of users.

Users benefit from faster and lower cost environments as applications are deployed across L2s. Some of the people that withheld from participating in DeFi because the base layer is too expensive will now experiment and become new users on L2s. We already mentioned how the lengthy withdraw period in fraud proof-based systems will be less of a concern over time as liquidity providers are stepping in as exit bridges and as ecosystems around each protocol are built out, reducing the need for users to withdraw funds out of the L2s.

Along with the lower fees and new use cases associated with L2s, users will also see a benefit from higher levels of incentives coming from both the L2 protocols and from the applications. Existing DeFi applications on the Ethereum base layer are already offering attractive incentives that are enough to overcome the onboarding UX challenges. With L2s, depending on the design of the protocol, users may receive incentives for participating on the network in methods that were too restrictive or uneconomical at the base layer (e.g. covered gas fees in the lower-fee environment as a marketing tactic to draw users onto their platform). Protocols with their own native tokens will have more flexibility around MEV design. These initial yield opportunities from both protocols and applications are likely to come down over time, but competition across protocols and applications should drive favorable rates to users for longer.

Similar to how Loopring/dYdX/DeversiFi/Lyra Finance deployed straight to L2, users will eventually bypass transacting directly on L1s altogether especially in the complexities are abstracted away.

L2 tech will also be adopted by / integrated into non-crypto-based (i.e. IRL) applications.

Non-crypto-based (i.e. IRL) applications are still adopting the Lightning Network for commerce and micropayments in reimagined business models that were not possible using existing payment rails. Content platforms that have relied on subscription-based models (e.g. Spotify or Time Magazine) could hypothetically charge users on a per-stream basis (potentially for under a penny's worth) by leveraging the Lightning Network. This can eliminate the need for inefficient subscriptions where users may be overpaying for content, and it can create new opportunities for content creators to earn a more equitable share of income earned on the platform. We have started to see green shoots of Lightning infiltration in the real-world as OpenNode partnered with content-platform Substack after integrating with BigCommerce earlier this summer to facilitate Lightning-based payments.

Certain games may leverage the Lightning Network to power digital economies using real currency or to reward players for completing in-game challenges. We note that blockchain-based gaming that can incorporate DeFi-concepts like yield farming and liquidity pools in a digestible format through regular gameplay – which can serve as a viable on-ramp to DeFi for the masses. Coupled with the prospects of NFT technology already showing up in art and music, these serve to greater incentivize L2 development and will accelerate the merge of the crypto and real worlds.

Conclusion

Layer 1 blockchains have not been able to scale significantly without sacrificing decentralization, a core feature that defines the value proposition for the entire cryptoeconomy. To scale without making unpalatable tradeoffs to the core Layer 1 blockchain necessitates building in layers. There have been many iterations of this concept, most of which we describe in this report. Ultimately, whether for scaling payments through state channels like Bitcoin's Lightning Network or computation through rollups like Ethereum's Arbitrum, we believe that a layered approach to scaling brings the most benefit with the least compromise on base-layer security, resiliency, and decentralization.

Legal Disclosure

1. See Zcash 2019 inflation bug. https://bitcoinist.com/zcash-inflation-bug-infinite-tokens/
2. https://vitalik.ca/general/2021/01/05/rollup.html
3. https://www.theblockcrypto.com/post/114479/polygon-hermez-merger-matic-hez-tokens-ethereum-projects
4. https://medium.com/matter-labs/optimistic-vs-zk-rollup-deep-dive-ea141e71e075
5. https://medium.com/interdax/ethereum-l2-optimistic-and-zk-rollups-dffa58870c93
6. https://www.ey.com/en_gl/news/2021/07/ey-contributes-a-zero-knowledge-proof-layer-2-protocol-into-the-public-domain-to-help-address-increasing-transaction-costs-on-ethereum-blockchain

This document, and the information contained herein, has been provided to you by Galaxy Digital Holdings LP and its affiliates ("Galaxy Digital") solely for informational purposes. This document may not be reproduced or redistributed in whole or in part, in any format, without the express written approval of Galaxy Digital. Neither the information, nor any opinion contained in this document, constitutes an offer to buy or sell, or a solicitation of an offer to buy or sell, any advisory services, securities, futures, options or other financial instruments or to participate in any advisory services or trading strategy. Nothing contained in this document constitutes investment, legal or tax advice. You should make your own investigations and evaluations of the information herein. Any decisions based on information contained in this document are the sole responsibility of the reader. Certain statements in this document reflect Galaxy Digital's views, estimates, opinions or predictions (which may be based on proprietary models and assumptions, including, in particular, Galaxy Digital's views on the current and future market for certain digital assets), and there is no guarantee that these views, estimates, opinions or predictions are currently accurate or that they will be ultimately realized. To the extent these assumptions or models are not correct or circumstances change, the actual performance may vary substantially from, and be less than, the estimates included herein. None of Galaxy Digital nor any of its affiliates, shareholders, partners, members, directors, officers, management, employees or representatives makes any representation or warranty, express or implied, as to the accuracy or completeness of any of the information or any other information (whether communicated in written or oral form) transmitted or made available to you. Each of the aforementioned parties expressly disclaims any and all liability relating to or resulting from the use of this information. Certain information contained herein (including financial information) has been obtained from published and non-published sources. Such information has not been independently verified by Galaxy Digital and, Galaxy Digital, does not assume responsibility for the accuracy of such information. Affiliates of Galaxy Digital own investments in some of the digital assets and protocols discussed in this document. This document provides links to other websites that we think might be of interest to you. Please note that when you click on one of these links, you may be moving to a provider's website that is not associated with Galaxy Digital. These linked sites and their providers are not controlled by us, and we are not responsible for the contents or the proper operation of any linked site. The inclusion of any link does not imply our endorsement or our adoption of the statements therein. We encourage you to read the terms of use and privacy statements of these linked sites as their policies may differ from ours. Except where otherwise indicated, the information in this document is based on matters as they exist as of the date of preparation and not as of any future date, and will not be updated or otherwise revised to reflect information that subsequently becomes available, or circumstances existing or changes occurring after the date hereof. The foregoing does not constitute a "research report" as defined by FINRA Rule 2241 or a "debt research report" as defined by FINRA Rule 2242 and was not prepared by Galaxy Digital Partners LLC.

For all inquiries, please email contact@galaxydigital.io.

galaxy

contact@galaxydigital.io
galaxy.com