Haym @SalomonCrypto

Sep 11 • 25 tweets • SalomonCrypto/status/1569023606029697024

(1/24) Blockchain Scaling: Optimistic Rollups

Wondering why folks are comfortable bridging assets to @arbitrum? Curious what's going on under @optimismFND's hood? Need to understand how @MetisDAO secures your $ETH?

Your guide to the today's premier blockchain scaling solutions.

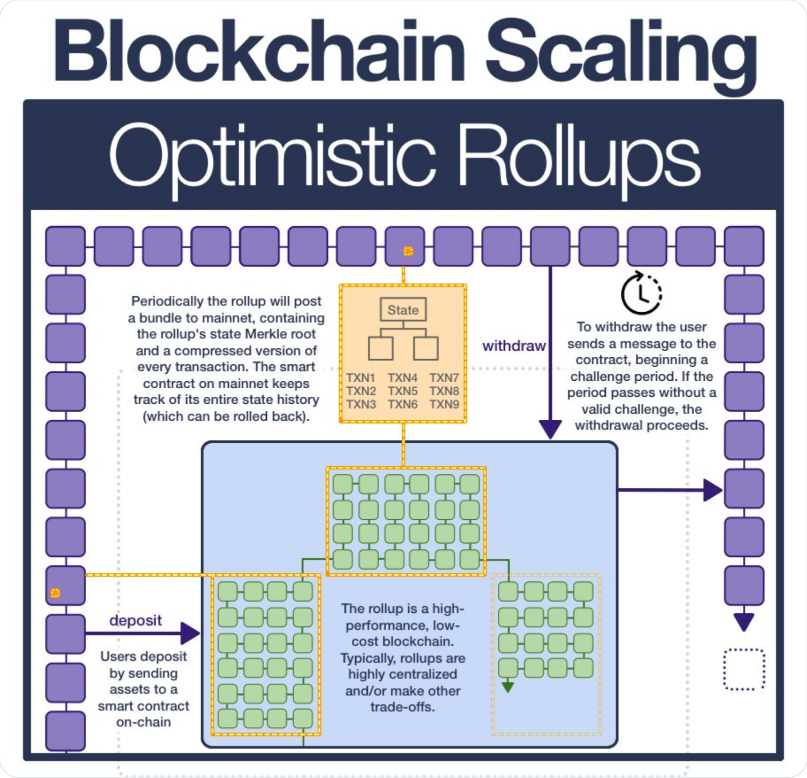

Blockchain Scaling

Optimistic Rollups

Periodically the rollup will post a bundle to mainnet, containing the rollup's state Merkle root and a compressed version of every transaction. The smart contract on mainnet keeps track of its entire state history (which can be rolled back).

To withdraw the user sends a message to the contract, beginning a challenge period. If the period passes without a valid challenge, the withdrawal proceeds.

Users deposit by sending assets to a smart contract on-chain

The rollup is a high-performance, low-cost blockchain. Typically, rollups are highly centralized and/or make other trade-offs.

(2/24) In 2008, modern finance failed.

As Satoshi Nakamoto watched the world burn around him, he had a vision: @Bitcoin.

7 years later, @VitalikButerin delivered on the promise that Satoshi first gave us. @ethereum was born.

(3/24) @ethereum is the World Computer, a decentralized, globally shared utility. Unfortunately, the World Computer is slow…

…or at least it was born slow.

Fortunately, we have scaling solutions!

(4/24) The first group of scaling solutions are state channels.

To open a channel, the users fund a smart contract where the funds are held in on-chain-escrow. The participants can transact off-chain as much as they want. When finished, the smart contract settles the channel.

Haym
@SalomonCrypto · Follow

(1/14) Blockchain Scaling: State Channels

@Bitcoin, @ethereum and all (good) blockchain computers share one important quality: they are SLOW. State channels are the first attempt at changing this and bringing blockchain to scale.

Your guide to the original scaling tech.

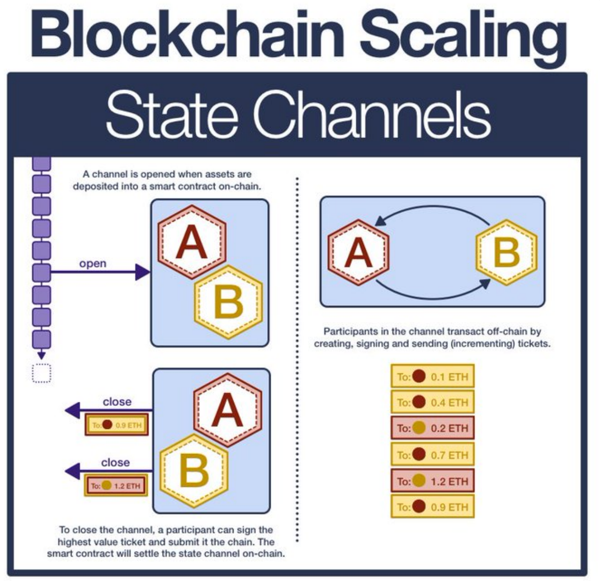

Blockchain Scaling

State Channels

A channel is opened when assets are deposited into a smart contract on-chain.

Participants in the channel transact off-chain by creating, signing and sending (incrementing) tickets.

To close the channel, a participant can sign the highest value ticket and submit it the chain. The smart contract will settle the state channel on-chain.

3:25 PM · Sep 10, 2022

Read the full conversation on Twitter

158 Reply Copy link

Read 9 replies

(5/24) State channels are powerful, but they have some serious limitations:

- participants must opt-in, cannot send funds to people who are not in channel
- requires large amount of capital to be locked
- cannot represent objects without a clear owner (eg @Uniswap)

(6/24) The next scaling solution was created to deal with (some of) these weaknesses.

Plasma (aka plasma chains) are independent blockchains that are anchored to @ethereum mainnet.

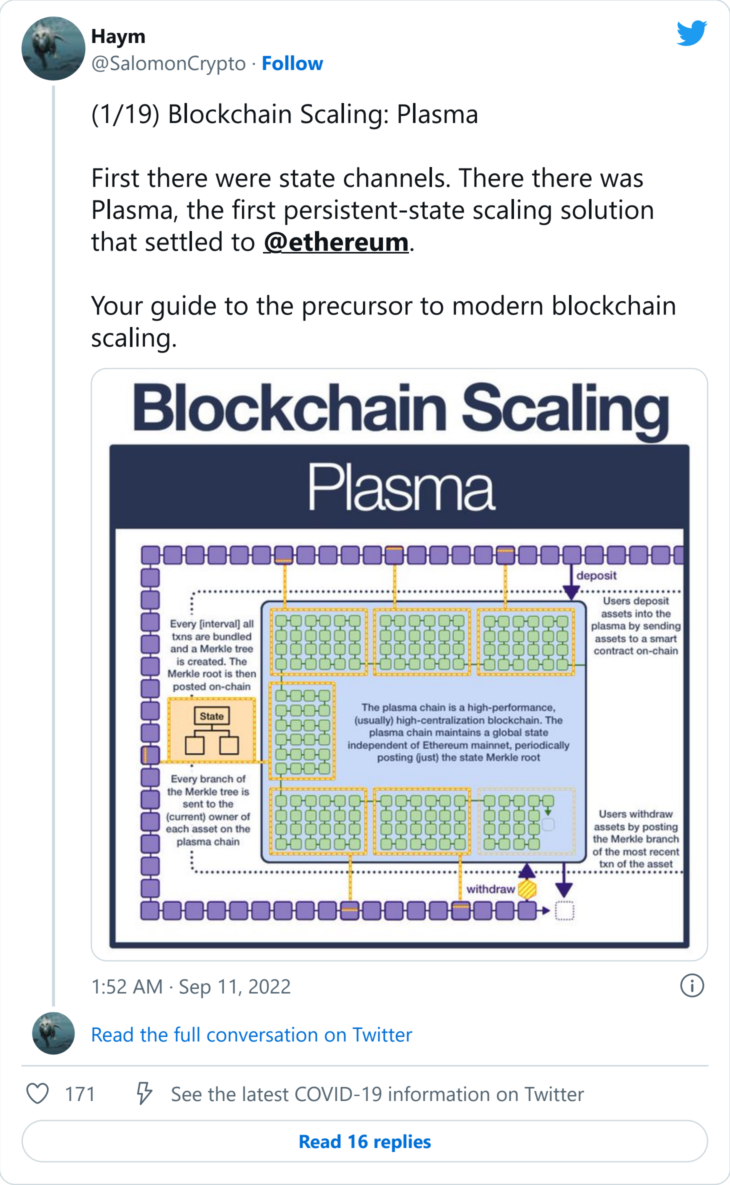

Haym
@SalomonCrypto · Follow

(1/19) Blockchain Scaling: Plasma

First there were state channels. There there was Plasma, the first persistent-state scaling solution that settled to **@ethereum**.

Your guide to the precursor to modern blockchain scaling.

Blockchain Scaling

Plasma

Every [interval] all txns are bundled and a Merkle tree is created. The Merkle root is then posted on-chain

Every branch of the Merkle tree is sent to the (current) owner of each asset on the plasma chain

The plasma chain is a high-performance, (usually) high-centralization blockchain. The plasma chain maintains a global state independent of Ethereum mainnet, periodically posting (just) the state Merkle root

Users deposit assets into the plasma by sending assets to a smart contract on-chain

Users withdraw assets by posting the Merkle branch of the most recent txn of the asset

1:52 AM · Sep 11, 2022

Read the full conversation on Twitter
171
See the latest COVID-19 information on Twitter
Read 16 replies

(7/24) Just like @ethereum, a plasma chain has its own virtual machine state - the status/balance of every user & smart contract

This state can be represented as a Merkle tree - a data structure that allows a huge amount of data to be compressed into a single line (Merkle root)

Haym
@SalomonCrypto · **Follow**

Replying to @SalomonCrypto

(9/13) A Merkle proof is established by providing the specific data and the intermediate hashes which allow a verifier to recreate the Merkle tree.

If the newly computed root node matches root node of the dataset, the verifier can be certain the data is in the dataset.

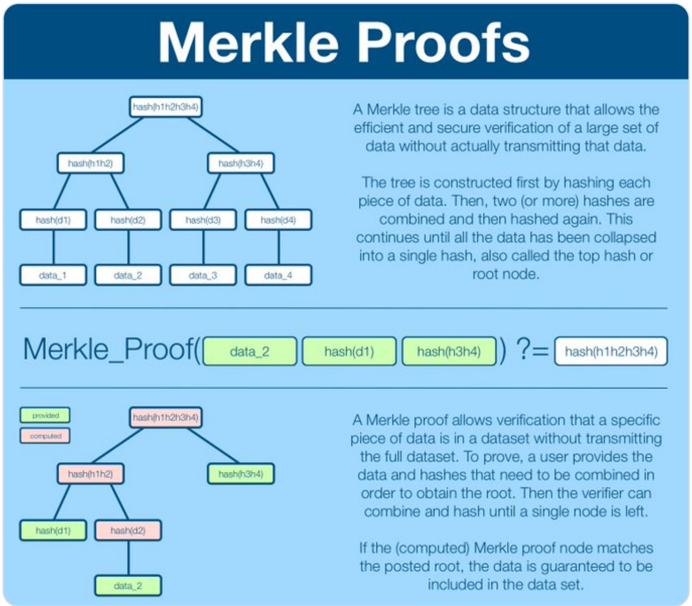

Merkle Proofs

A Merkle tree is a data structure that allows the efficient and secure verification of a large set of data without actually transmitting that data.

The tree is constructed first by hashing each piece of data. Then, two (or more) hashes are combined and then hashed again. This continues until all the data has been collapsed into a single hash, also called the top hash or root node.

Merkle_Proof(data_2 hash(d1) hash(h3h4)) ?= hash(h1h2h3h4)

A Merkle proof allows verification that a specific piece of data is in a dataset without transmitting the full dataset. To prove, a user provides the data and hashes that need to be combined in order to obtain the root. Then the verifier can combine and hash until a single node is left.

If the (computed) Merkle proof node matches the posted root, the data is guaranteed to be included in the data set.

10:17 PM · Sep 7, 2022

10 Reply Copy link

Read 1 reply

(8/24) Once every [interval] the plasma chain will build the state Merkle tree and post the Merkle root to mainnet. Thus, a (compressed) record of the entire plasma chain exists on mainnet.

In order to withdraw assets, a user submits the Merkle proof for that asset.

(9/24) While plasma delivers improvement on state channels, they share a (huge) weakness: both require engaged ownership. If an owner does not care about an asset, then an "invalid" outcome involving that asset may result.

Good luck implementing a DEX... or really any dApp.

(10/24) But there is an even bigger issue: data availability.

A plasma chain only posts the state Merkle root to mainnet; it does not post any of the transaction or block data needed to generate fraud proofs.

Rollups are the answer: post everything to mainnet.

(11/24) Users interact with an optimistic rollup by depositing assets into a smart contract on @ethereum. The rollup operator then mints an equivalent amount of assets on the rollup chain and gives it to the depositor.

On mainnet, the assets remain in escrow.

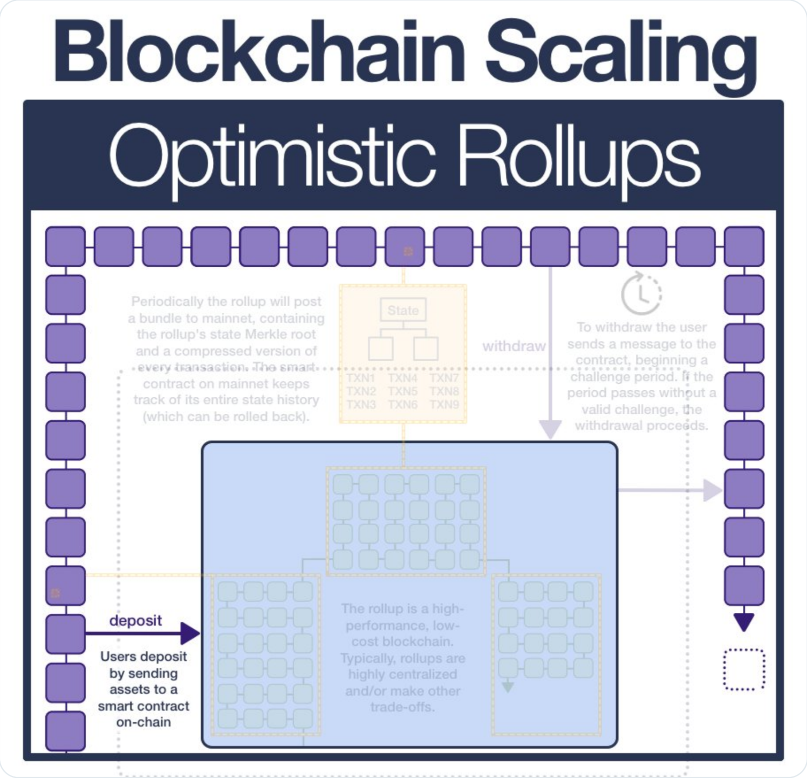

Blockchain Scaling

Optimistic Rollups

Periodically the rollup will post a bundle to mainnet, containing the rollup's state Merkle root and a compressed version of every transaction. The smart contract on mainnet keeps track of its entire state history (which can be rolled back).

To withdraw the user sends a message to the contract, beginning a challenge period. If the period passes without a valid challenge, the withdrawal proceeds.

Users deposit by sending assets to a smart contract on-chain

The rollup is a high-performance, low-cost blockchain. Typically, rollups are highly centralized and/or make other trade-offs.

(12/24) Because rollups rely on @ethereum for decentralized property rights and settlement, the rollup operator(s) can be much more centralized.

From a user perspective, execution times and gas costs are SIGNIFICANTLY cheaper than using mainnet.

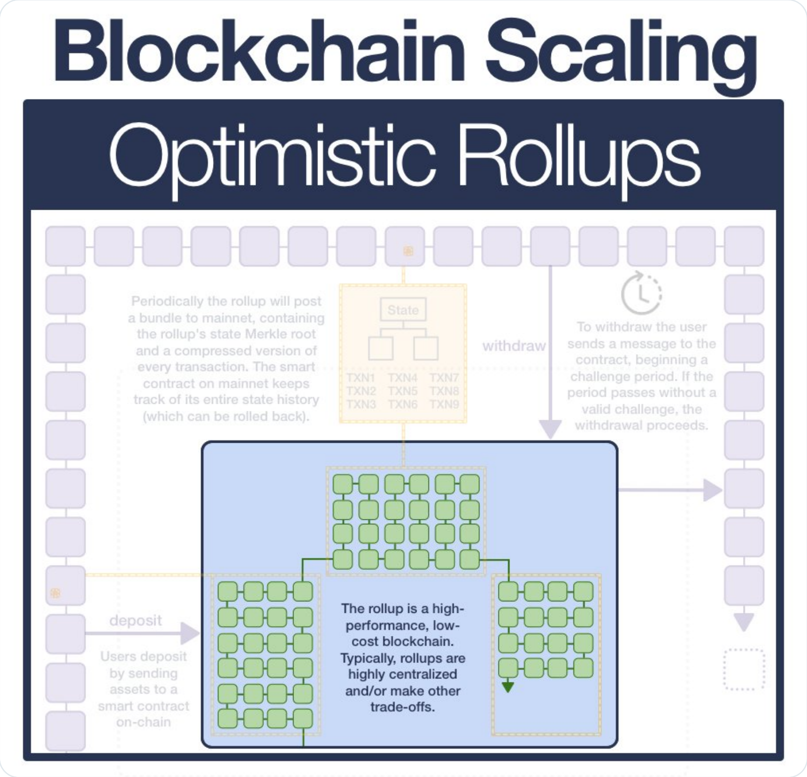

Periodically the rollup will post a bundle to mainnet, containing the rollup's state Merkle root and a compressed version of every transaction. The smart contract on mainnet keeps track of its entire state history (which can be rolled back).

To withdraw the user sends a message to the contract, beginning a challenge period. If the period passes without a valid challenge, the withdrawal proceeds.

The rollup is a high-performance, low-cost blockchain. Typically, rollups are highly centralized and/or make other trade-offs.

Users deposit by sending assets to a smart contract on-chain

(13/24) Through here, rollups and plasma are basically the same thing. Both are high performance, centralized blockchains anchored to @ethereum via escrow smart contracts.

Plasma post the Merkle root to mainnet once every [interval], writing an unimpeachable record to mainnet.

(14/24) Plasma chains stop here; the only post the Merkle root to mainnet.

The Merkle root is a single line that can be used to prove a transaction. You can only prove a txn was in a Merkle tree, you cannot search the txns in a Merkle tree just from its root.

(15/24) During normal operation, this isn't an issue. Every time the plasma operator posts a new Merkle root, they also send every asset owner the necessary Merkle branches.

Users must rely on the operator to provide block data if they need to create fraud proofs

(16/24) But what if one day they just... stopped.

A malicious operator could easily make an invalid transaction and hide the data necessary for creating the fraud-proof.

This wouldn't be possible if the operator was required to make txn data available on mainnet.

(17/24) Rollups are the solution to plasma's data availability problems

Putting ALL data on-chain allows anyone to locally process all the operations in the rollup if they wish to, allowing them to detect fraud, initiate withdrawals, or personally start producing txn batches

(18/24) Operators publish a batch of transactions (including the previous and new state root) to the contract

The contract checks the submitted previous state root matches its copy of the current state root and, if so, switches the contract state root to the supplied state root

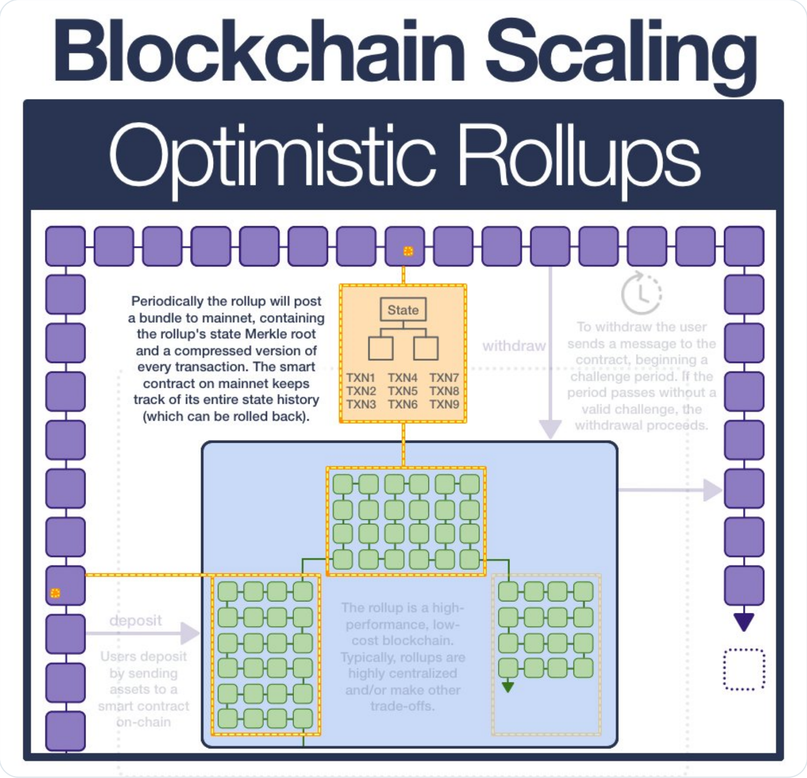
Blockchain Scaling

Optimistic Rollups

Periodically the rollup will post a bundle to mainnet, containing the rollup's state Merkle root and a compressed version of every transaction. The smart contract on mainnet keeps track of its entire state history (which can be rolled back).

To withdraw the user sends a message to the contract, beginning a challenge period. If the period passes without a valid challenge, the withdrawal proceeds.

Users deposit by sending assets to a smart contract on-chain

The rollup is a high-performance, low-cost blockchain. Typically, rollups are highly centralized and/or make other trade-offs.

(19/24) Rollups scale @ethereum in 2 ways:

First, by moving computation to a high-performance chain, gas costs related to execution are drastically reduced

Second, they are able to post txns in a highly compressed form, greatly reducing the gas costs of posting data to mainnet

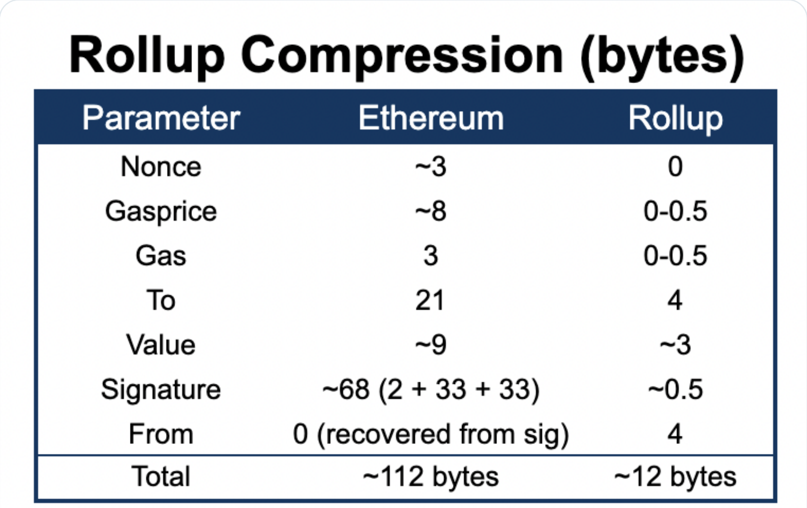

Rollup Compression (bytes)

Parameter	Ethereum	Rollup
Nonce	~3	0
Gasprice	~8	0-0.5
Gas	3	0-0.5
To	21	4
Value	~9	~3
Signature	~68 (2 + 33 + 33)	~0.5
From	0 (recovered from sig)	4
Total	~112 bytes	~12 bytes

Rollups leverage superior encoding (and some tricks) to significantly decrease the amount of data needed to store a transaction

(20/24) But there is one glaring issue: how does the smart contract know that the new state roots are correct?

An OPTIMISTIC rollup assumes all batches are valid... BUT it leaves open a challenge window.

(21/24) Anyone who was keeping up with the chain and detects fraud can publish a fraud proof, proving the batch is invalid and should be reverted.

Thus, the rollup assumes good behavior, but relies on the economic incentives of untrusted actors to maintain its integrity.

(22/24) To withdraw, a user sends the contract a message and initiates a withdrawal txn. This changes the state of the rollup; a new batch/state root is posted on-chain.

After the challenge window passes, the txn is finalized and the user can withdraw their assets

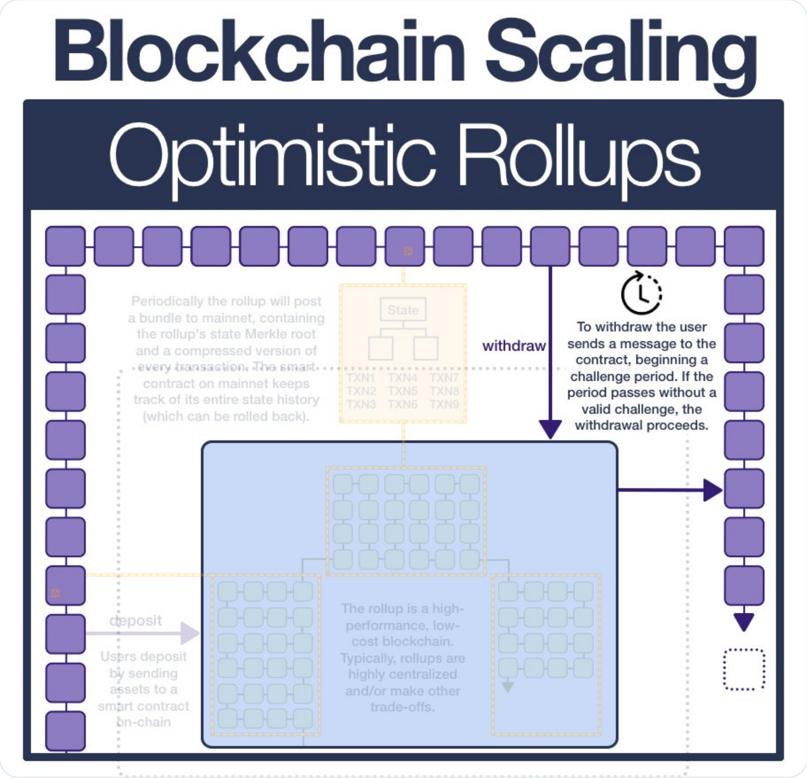

Blockchain Scaling

Optimistic Rollups

Periodically the rollup will post a bundle to mainnet, containing the rollup's state Merkle root and a compressed version of every transaction. The smart contract on mainnet keeps track of its entire state history (which can be rolled back).

To withdraw the user sends a message to the contract, beginning a challenge period. If the period passes without a valid challenge, the withdrawal proceeds.

The rollup is a high-performance, low-cost blockchain. Typically, rollups are highly centralized and/or make other trade-offs.

Users deposit by sending assets to a smart contract on-chain

(23/24) Both plasma and rollups are built on the same principle: offload execution while anchoring settlement to @ethereum. But rollups make a huge leap forward by solving data availability.

The most important result: assets no longer needs owners.

Rollups can run an EVM.

(24/24) A EVM-compatible/equivalent rollup is the dream: all the properties of @ethereum, just cheaper and faster.

...but we can do better, can't we? What if instead of being "optimistic" we wanted to settle things instantly?

Is verification possible with Zero-Knowledge?

Like what you read? Help me spread the word by retweeting the thread (linked below).

Follow me for more explainers and as much alpha as I can possibly serve.

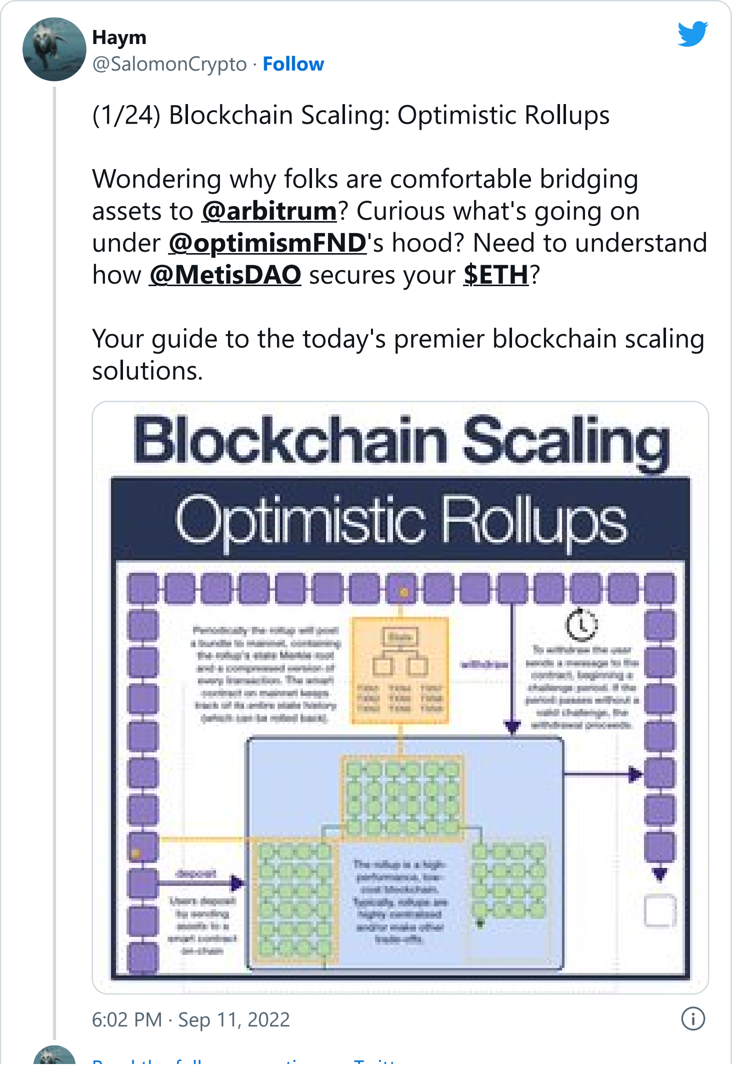

Haym
@SalomonCrypto · Follow

(1/24) Blockchain Scaling: Optimistic Rollups

Wondering why folks are comfortable bridging assets to @arbitrum? Curious what's going on under @optimismFND's hood? Need to understand how @MetisDAO secures your $ETH?

Your guide to the today's premier blockchain scaling solutions.

Blockchain Scaling

Optimistic Rollups

Periodically the rollup will post a bundle to mainnet, containing the rollup's state Merkle root and a compressed version of every transaction. The smart contract on mainnet keeps track of its entire state history (which can be rolled back).

To withdraw the user sends a message to the contract, beginning a challenge period. If the period passes without a valid challenge, the withdrawal proceeds.

The rollup is a high-performance, low-cost blockchain. Typically, rollups are highly centralized and/or make other trade-offs.

Users deposit by sending assets to a smart contract on-chain

6:02 PM · Sep 11, 2022

Read the full conversation on Twitter

• • •