Layer 2 Scaling Solutions

Rollups, ZK-Proofs, and the Cryptoeconomics of Scaling Ethereum

Prof. Dr. Jörg Osterrieder

BSc Blockchain, Crypto Economy & NFTs

Spring 2026

The Gas Fee Paradox

Summer 2021. DeFi is booming.

- You want to swap $50 of tokens on Uniswap
- The swap costs $3 in fees on a quiet day
- During the NFT craze of May 2021: same swap costs $200 in fees
- Result: only wealthy wallets can afford to interact with Ethereum

A network designed for everyone became accessible only to the few.

This is not an accident. It is a consequence of Ethereum's fundamental design trade-off: security and decentralisation came at the cost of throughput.

The Question This Lecture Answers

How do **Layer 2 scaling solutions** reduce fees by up to 80 times without sacrificing Ethereum's security?

And: what are the **economic and governance trade-offs** that come with moving activity off the main chain?

Central thesis: L2s are not free lunches. They shift trust assumptions from Ethereum's validators to new actors – sequencers, provers, and token holders.

Gas fees peaked above 500 gwei during the CryptoPunks and Bored Ape mint periods (2021). At ETH price of $3,000 and 21,000 gas per transfer, one transfer cost over $30.

Roadmap: Seven Sections, 90 Minutes

After this lecture you can:

1. Explain why Ethereum cannot scale on its own (Trilemma)
2. Describe how rollups batch transactions to reduce L1 cost
3. Distinguish optimistic rollups from ZK-rollups by trust model
4. Name the leading L2 protocols and their trade-offs
5. Identify the sequencer centralisation risk and its governance response
6. Apply the cryptoeconomics 6-question lens to any L2 design

Seven sections:

1. The Scaling Crisis: why fees got so high
2. Layer 2: the architectural solution
3. Optimistic Rollups: innocent until proven guilty
4. ZK-Rollups: prove it now, not later
5. The L2 Ecosystem: who is winning
6. Cryptoeconomics: the new trade-offs
7. Summary and key takeaways

No prior knowledge of Layer 2 assumed. All technical terms defined before use.

L2Beat (l2beat.com) tracks all major L2 protocols with live TVL, risk ratings, and upgrade keys. Recommended reference throughout this lecture.

Why Does Every Ethereum Transaction Have a Fee?

The block space auction:

- Every Ethereum block has a fixed gas limit: **30 million gas**
- Every operation uses gas: a simple transfer costs 21,000 gas
- That means at most $\approx$ 1,428 simple transfers per block
- Blocks arrive every 12 seconds: $\approx$ 7 transfers per second maximum
- When demand exceeds supply, users bid to get included

After EIP-1559 (Aug 2021):

- A base fee is algorithmically set and burned
- Users add a priority tip to incentivise validators

The Structural Limit

Ethereum processes **15 to 30 transactions per second** on the main chain.

Visa processes 24,000 per second.

This is not a bug to fix with better code. It is a deliberate trade-off for decentralisation.

Key term: Gas

Gas measures the computational work an Ethereum transaction requires. It is the unit of block space. More complex operations (DeFi swaps, NFT mints) use more gas than simple ETH transfers.

EIP-1559 introduced the base fee mechanism in August 2021, improving fee predictability. Users no longer need to guess the right bid; the protocol adjusts the base fee automatically.

The 2021 Gas Crisis: Five Years of Ethereum Fees

Four eras in the chart:

- **2020:** DeFi launches; fees rise from 12 to 55 gwei
- **2021 Q2:** DeFi Summer peak; fees above 130 gwei; swapping costs $200+
- **Sep 2022:** The Merge; Ethereum switches to Proof of Stake – fees fall, but throughput unchanged
- **Mar 2024:** Dencun upgrade (EIP-4844); fees drop sharply as L2s use cheap blob data

By 2025, mainnet fees average 4 gwei. L2 adoption changed the demand curve.

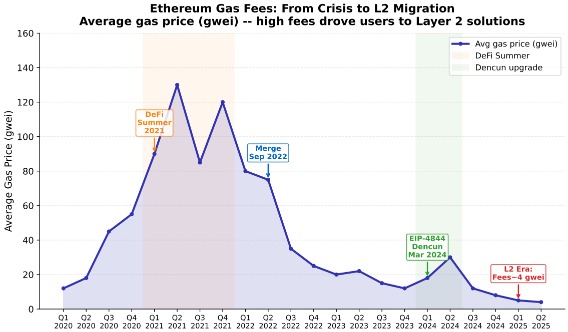

The Merge (September 2022) reduced energy use by 99% but did not increase throughput. EIP-4844 (Dencun, March 2024) added "blob" data, cutting L2 calldata costs by 10 to 100 times.

The Scalability Trilemma: Pick Any Two

Vitalik Buterin's Trilemma (2014):
A blockchain can achieve at most two of three properties simultaneously:

- **Decentralisation:** many independent nodes can verify all transactions
- **Security:** a majority of validators would need to collude to corrupt the chain
- **Scalability:** the chain processes thousands of transactions per second

Ethereum L1 chose decentralisation and security. Throughput was left as the hard constraint.

L2 rollups: the third path. Execute transactions off-chain; settle proofs on-chain. Inherit L1 security without consuming L1 block space.

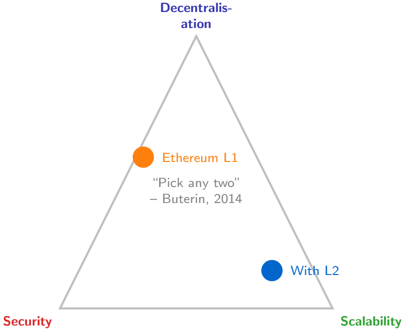

The L2 Insight

Rollups do not violate the Trilemma. They move *execution* off-chain while keeping *data and settlement* on-chain. Security is inherited, not sacrificed.

The Trilemma was popularised by Buterin but is not formally proven. Some researchers dispute whether it is a hard physical limit or an engineering constraint that better designs can relax.

What Is Layer 2? A Definition Before the Details

Layer 2 (L2): A separate blockchain or protocol that executes transactions off Ethereum mainnet (Layer 1), but **posts compressed transaction data or cryptographic proofs back to L1**. Security is inherited from Ethereum: even if the L2 operator disappears, users can always withdraw funds by interacting with L1 directly.

What L2 does:

- Executes thousands of transactions cheaply and quickly
- Groups (batches) many transactions together
- Submits a single compressed record to Ethereum L1
- Enforces correctness through fraud proofs or validity proofs

What L2 does NOT do:

- Replace Ethereum – it depends on L1 for finality
- Sacrifice security – an honest user can always exit to L1
- Remove trust entirely – new trust assumptions appear (sequencer, prover)

Analogy: A tab at a bar. Every drink is logged, but only one bill is settled at closing time – reducing the number of card swipes from 20 to 1.

L2s that post full transaction data to L1 are called "rollups" because they roll up many transactions into one. Variants that post less data (sidechains, validiums) offer lower security guarantees.

How Rollups Work: Batching Transactions to Save Block Space

The four-step flow:

1. **Users submit** transactions to the L2 network (fast, cheap)
2. **Sequencer orders** and executes the transactions off-chain
3. **Batch compressor** groups 100 to 1,000 transactions into one calldata or blob
4. **Batch posted** to Ethereum L1 as a single transaction

Cost arithmetic:

- 1,000 users share the cost of one L1 transaction
- Result: 100 to 1,000 times cheaper per user

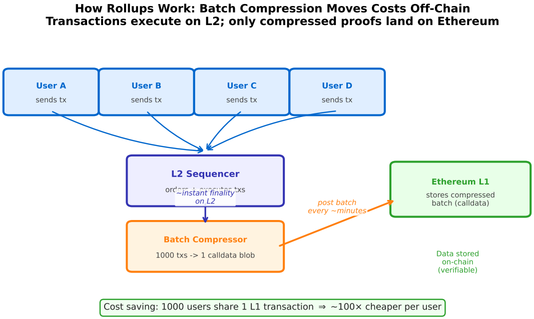

The bottleneck was L1 calldata cost until EIP-4844 (March 2024) introduced blobs – a new cheaper data space. L2 fees dropped by an additional 10 to 100 times after Dencun.

Data Availability: What Must Actually Land on Ethereum?

The core guarantee of rollups:
Anyone can reconstruct the full L2 state by reading Ethereum L1. This requires the L2 operator to post **enough data** so that independent observers can verify every transaction.

Optimistic rollups post:

- Full compressed transaction data (calldata / blob)
- Anyone can re-execute and detect fraud
- Higher data cost, stronger guarantee

ZK-rollups post:

- A validity proof (a few kilobytes) plus minimal state data
- No re-execution needed – the proof is the guarantee
- Lower data cost, mathematically guaranteed

The Validium Trade-off

Some systems (validiums, like Immutable X) post data **off-chain** (to IPFS or a committee).

This is cheaper still, but if the data committee disappears, users cannot reconstruct their balances from L1 alone.

Security hierarchy:

1. On-chain data rollup (strongest)
2. Validium with trusted committee (medium)
3. Pure sidechain (weakest – independent security)

Data availability is the most important security property distinguishing rollups from sidechains. L2Beat classifies every L2 protocol by its data availability model.

EIP-4844: How the Dencun Upgrade Made L2 Even Cheaper

The problem before March 2024:

- L2s posted compressed data as **calldata** to Ethereum
- Calldata is permanent storage: expensive for validators forever
- L2 transaction fees: typically $0.10 to $0.50

EIP-4844 solution: blobs

- New data type: **blobs** (Binary Large OBjects)
- Blobs are stored for only 18 days, then pruned
- Separate fee market: blob fees are much lower than calldata fees
- Result: L2 fees dropped to $0.001 to $0.05 per transaction

Before vs After Dencun (March 13, 2024)

Metric	Before	After
Base tx cost	~$0.25	~$0.01
Swap cost	~$0.50	~$0.05
Data type	calldata	blob
Data kept	forever	18 days

Why Blobs Are Safe

L2s only need blob data to be available *temporarily* – long enough to detect fraud or verify proofs. After that, the state root on L1 is the final record.

Proto-danksharding (EIP-4844) is a step toward full danksharding, which will allow hundreds of blobs per block and could reduce L2 costs by another 10 to 100 times in future Ethereum upgrades.

Optimistic Rollups: Innocent Until Proven Guilty

Optimistic rollup: A rollup that **assumes** all transactions are valid and posts them to L1 without immediate verification. A **challenge window** (typically 7 days) allows anyone to submit a fraud proof if they detect an invalid state transition.

How fraud proofs work:

1. Sequencer posts batch to L1 and claims a new state root
2. **Challenge window opens** (7 days)
3. Any verifier can re-run the batch and check the state root
4. If fraud is detected: submit a fraud proof on-chain
5. L1 contract adjudicates; fraudulent batch is rolled back
6. Sequencer loses a slashed bond as penalty

The optimism: why it works in practice

- Fraud proofs are rarely submitted – sequencers have strong incentives to behave
- The *threat* of challenge is enough to ensure honesty
- Only one honest verifier is needed in the entire network

The Cost of Optimism

The 7-day challenge window means **withdrawals from L2 to L1 take 7 days**. Fast bridges exist but introduce their own trust assumptions.

The fraud proof mechanism is analogous to a notary system: the notary stamps documents as valid, but any citizen can challenge a fraudulent stamp in court within a fixed period.

The 7-Day Withdrawal Wait: Feature or Bug?

Why 7 days?

- Enough time for verifiers worldwide to re-run all L2 transactions
- Enough time to submit a fraud proof even from a slow connection
- Long enough that no attacker can outmanoeuvre the challenge system

User experience impact:

- Moving funds from Arbitrum or Optimism to Ethereum mainnet takes 7 days natively
- Third-party bridges (Stargate, Hop, Across) offer instant exits
- Bridge liquidity providers take the 7-day risk and charge a small fee

The Bridge Trilemma

Fast bridges introduce new trust: you trust the bridge operator and its liquidity pool instead of Ethereum's validators.

Bridge hacks cost over $2.5 billion (2021–2023): Ronin ($625M), Wormhole ($320M), Nomad ($190M).

Key Insight

The 7-day delay is not a design failure. It is the security guarantee. Accepting a faster bridge means trading L1-level security for liquidity provider-level security.

ZK-rollups solve the 7-day problem: validity proofs are verified on L1 immediately, so withdrawals can settle in minutes. This is a key competitive advantage of ZK designs.

The Optimistic Camp: Arbitrum, Optimism, and Base

Arbitrum One

- Launched: Aug 2021
- Tech: Nitro (multi-round fraud proofs)
- TVL: largest optimistic L2 (~$9B, 2025)
- Token: ARB (governance)
- Users: DeFi-heavy (GMX, Aave, Uniswap)
- Key edge: fastest fraud proof system

Optimism (OP Mainnet)

- Launched: Dec 2021
- Tech: OP Stack (open-source)
- TVL: ~$4B (2025)
- Token: OP (retroactive public goods)
- Users: Synthetix, Velodrome
- Key edge: Superchain ecosystem (shared sequencer)

Base

- Launched: Aug 2023 (Coinbase)
- Tech: OP Stack fork
- TVL: ~$8B and growing (2025)
- Token: none (Coinbase-operated)
- Users: consumer apps, social (Farcaster)
- Key edge: Coinbase distribution; surpassed Ethereum mainnet in daily tx

The OP Stack (Optimism's open-source codebase) powers over 20 chains in 2025, creating a "Superchain" of interoperable rollups sharing the same fraud proof infrastructure.

Coinbase's Base reached over 2 million daily transactions in 2025, surpassing Ethereum mainnet. This is the clearest sign that L2s are no longer supplemental – they are the primary execution layer.

Zero-Knowledge Proofs: Proving Something Without Revealing It

The core idea (no mathematics required):

The Treasure Map Problem

Alice has a map showing the location of buried treasure. She wants to prove to Bob that she knows where the treasure is – without showing him the map.

A zero-knowledge proof lets Alice convince Bob she has the map (the "witness") by answering a series of challenges that only someone with the map could answer correctly – without Bob learning the map's contents.

Applied to L2 rollups:

- The "map" is all the transaction data processed by the sequencer
- The "proof" is a short cryptographic certificate (a few kilobytes)
- Ethereum L1 verifies the proof in milliseconds
- If the proof verifies, the new state root is accepted – instantly

Two main ZK proof systems:

zk-SNARK Succinct Non-interactive ARgument of Knowledge. Small proof, fast verify. Requires a trusted setup.

zk-STARK Scalable Transparent ARgument of Knowledge. Larger proof, no trusted setup. Used by StarkNet.

The Key Property

ZK proofs are **validity proofs**: they guarantee the new state is correct. No one needs to watch for fraud. The math is the enforcement.

zk-SNARKs require a "trusted setup ceremony" – a multi-party computation that generates public parameters. If all participants are compromised, the system is insecure. zk-STARKs avoid this requirement.

ZK-Rollups: Withdraw Today, Not in Seven Days

How a ZK-rollup finalises:

1. Sequencer collects and executes transactions off-chain
2. A **prover** (often a specialised GPU cluster) generates a validity proof for the entire batch
3. Proof and compressed state data are posted to Ethereum L1
4. L1 smart contract **verifies the proof** (takes milliseconds, costs ~500k gas)
5. New state root accepted immediately – **no challenge window**

Consequence: withdrawals from ZK-rollup to L1 take minutes, not 7 days.

Optimistic vs ZK compared:

	Optimistic	**ZK**
Finality	7 days	Minutes
Proof cost	Cheap	GPU-heavy
EVM support	Full	Growing
Trust model	Fraud proof	Validity proof
Key risk	Challenge window	Prover centralisation

The Future

Vitalik Buterin: “In the long run, everything will be ZK.” Proof generation costs are falling rapidly – zkEVM (ZK + full EVM) reached production in 2024.

Generating a ZK proof for a batch of 1,000 transactions can take several minutes on dedicated hardware. This means ZK-rollups have higher latency in proof generation, even though on-chain finality is faster.

The ZK Camp: zkSync, StarkNet, Polygon zkEVM

zkSync Era

- Launched: Mar 2023 (Matter Labs)
- Proof: custom zkEVM (SNARK)
- EVM: native zkEVM
- TVL: ~$3B (2025)
- Token: ZK (governance + gas)
- Key edge: largest ZK TVL; developer-friendly

StarkNet

- Launched: Nov 2021 (StarkWare)
- Proof: zk-STARK (no trusted setup)
- EVM: Cairo VM (not EVM-native)
- TVL: ~$1B (2025)
- Token: STRK (governance + fees)
- Key edge: most scalable proof system; games and DeFi

Polygon zkEVM

- Launched: Mar 2023 (Polygon Labs)
- Proof: recursive SNARKs
- EVM: bytecode-equivalent
- TVL: ~$0.5B (2025)
- Token: POL (governance)
- Key edge: full EVM compatibility; enterprise partnerships

EVM compatibility is the key battleground. Most DeFi code is written for the EVM. ZK-rollups that can run EVM bytecode without modification capture existing developer ecosystems without rewriting.

zkEVM projects differ in their compatibility level. Polygon zkEVM aims for bytecode equivalence; StarkNet uses a custom language (Cairo) and requires protocol-level transpilation.

Optimistic vs ZK: No Single Winner Across All Dimensions

Reading the chart:

- **EVM compatibility:** Optimistic rollups score higher today – they inherit the EVM directly
- **Withdrawal speed:** ZK wins decisively – validity proofs eliminate the challenge window
- **Proof cost:** Optimistic is cheaper – no GPU-intensive proving required
- **Security assumptions:** ZK wins – cryptographic guarantee vs game-theoretic guarantee
- **Ecosystem maturity:** Optimistic leads – Arbitrum and Base have larger user bases

Trend: ZK scores are improving every quarter as proving costs fall and zkEVM matures.

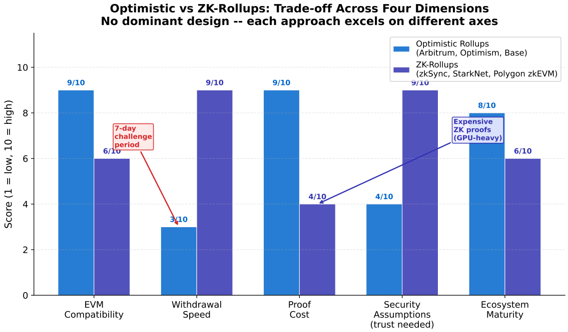

"Proof cost" here refers to the off-chain cost of generating proofs, not the on-chain verification cost. On-chain ZK proof verification (500k–3M gas) is already competitive with optimistic fraud proof adjudication.

Layer 2 Ecosystem: From $3 Billion to $30 Billion in Three Years

What the chart shows:

- **2022:** L2 ecosystem is niche, dominated by Arbitrum and Optimism
- **2023 H2:** Base launches; rapid user adoption from Coinbase
- **2024:** EIP-4844 (Dencun) cuts costs; TVL accelerates
- **2025:** Total L2 TVL exceeds $30B; Arbitrum and Base lead

Key milestone (2025):
Base processes more daily transactions than Ethereum mainnet. The centre of Ethereum activity has moved to Layer 2.

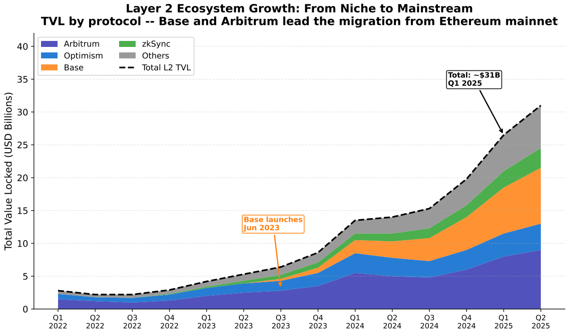

TVL (Total Value Locked) measures the assets deposited in L2 bridge contracts on Ethereum L1. It is a proxy for user adoption but does not capture L2-native assets or transaction volume.

20 to 80 Times Cheaper: The Transaction Cost Revolution

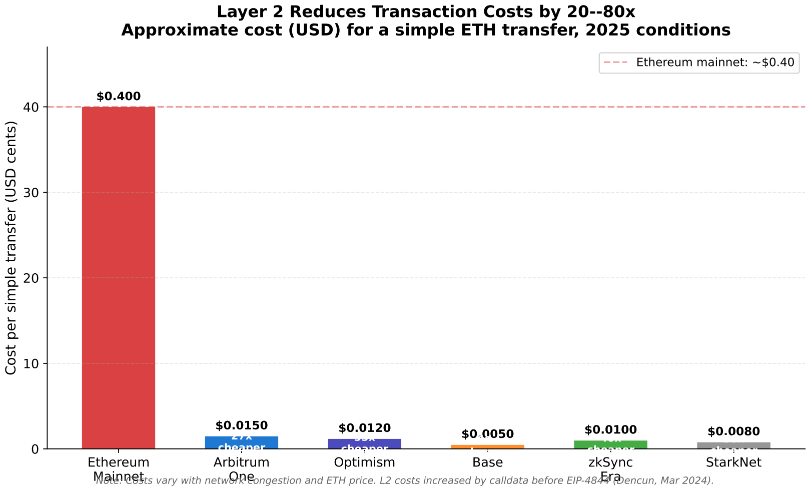

Costs shown are for a simple ETH transfer under normal network conditions, 2025. DeFi swaps cost 3 to 5 times more. L2 costs vary with L1 blob fees, which fluctuate with Ethereum mainnet demand.

The L2 Race: Metrics That Matter in 2025

Daily transactions (2025 estimates):

- **Base:** 5 to 8 million/day (surpassed Ethereum mainnet)
- **Arbitrum:** 2 to 4 million/day
- **Optimism:** 1 to 2 million/day
- **Ethereum mainnet:** ~1 to 1.5 million/day
- **zkSync Era:** 0.5 to 1 million/day

Developer activity (GitHub commits, 2024–2025):

- OP Stack and Arbitrum Nitro attract the most DeFi forks
- ZK stacks are growing fastest in new project starts

The Sequencer Revenue Question

Sequencers collect user fees but pay lower costs to L1 after EIP-4844.

The difference is **sequencer revenue** – essentially a private toll road on Ethereum settlement.

Arbitrum's sequencer earns tens of millions of dollars per month. Who controls this revenue is a live governance debate.

The L2 ecosystem is competitive: protocols that launched in 2021 have been overtaken by Base (2023) in transactions. Speed of user acquisition depends heavily on distribution channels (Coinbase has 100M+ users).

The Sequencer Problem: Speed Comes at a Cost to Decentralisation

What is a sequencer?
The sequencer is the entity that:

- Receives all user transactions on the L2
- Decides the order of execution
- Batches and submits to L1
- In most L2s today: a **single server** operated by the founding team

Sequencer risks:

- **Censorship:** sequencer can exclude specific addresses
- **Ordering manipulation:** sequencer can front-run users (MEV extraction)
- **Downtime:** if the sequencer crashes, the L2 halts
- **Force exit:** users can exit via L1 directly, but it is slow and expensive

The Escape Hatch

All major rollups include a **forced transaction** mechanism: if the sequencer ignores you for a set period (typically 24 hours), you can submit a transaction directly to the L1 bridge contract.

This preserves the "you cannot lose your funds" guarantee – but does not prevent censorship of non-exiting transactions.

Current Status (2025)

No major L2 has launched a decentralised sequencer in production. It is the most actively researched open problem in L2 design.

Espresso Systems, Astria, and others are building "shared sequencer" networks that multiple rollups can use, reducing the centralisation risk. Arbitrum has announced plans to move toward a decentralised sequencer.

L2 Token Economics: Governance, Incentives, and Legitimacy

Why L2s launch tokens:

- Governance: token holders vote on protocol upgrades and fee parameters
- Retroactive distribution: reward early users and developers
- Ecosystem funding: treasury for grants and public goods
- Revenue share: some designs direct sequencer fees to token stakers

ARB (Arbitrum):

- Airdropped March 2023 to 625,000 wallets
- ArbitrumDAO governs the protocol upgrade process
- Sequencer revenue flows to the DAO treasury

OP (Optimism):

- Two-house governance: Token House (OP holders) and Citizens' House (non-transferable NFT)
- Retroactive Public Goods Funding: rewards projects that created value for the ecosystem
- Part of sequencer revenue funds the RPGF pool

The Legitimacy Question

L2 tokens give their holders power over systems that process billions in user funds. Unlike Ethereum validators, L2 governance token holders are not staking collateral. The accountability mechanism is incomplete.

L2 upgrade keys are a major risk: if the founding team controls the upgrade key, they can change the bridge contract and drain user funds. L2Beat's risk assessment scores every protocol on upgrade key decentralisation.

MEV on Layer 2: The Same Problem in a New Venue

What is MEV? (recap)

Maximal Extractable Value: the profit a block producer can capture by **reordering, inserting, or censoring** transactions in a block.

MEV on L2 takes three forms:

1. **Sequencer MEV:** the centralised sequencer front-runs user swaps (sandwiching)
2. **Cross-domain MEV:** arbitrageurs exploit price differences between L1 and L2 (latency-based)
3. **Proposer MEV:** if decentralised sequencing is adopted, proposers compete like L1 validators

Scale of the Problem

Researchers estimate MEV on Arbitrum alone exceeds $50M per year. As L2s grow, MEV extraction will grow proportionally.

Mitigations being explored:

- **FCFS ordering:** First-Come-First-Served removes sequencer discretion
- **Encrypted mempools:** transactions sealed until execution
- **PBS on L2:** separating proposing and building roles

MEV is not unique to L2. Moving from Ethereum to L2 changes the extractor (from validators to sequencers) but does not eliminate extraction.

MEV on optimistic rollups is particularly under-studied because the sequencer is a single entity with full ordering control. ZK-rollups with decentralised provers face similar ordering games.

The Cryptoeconomics Lens: Six Questions for Layer 2

1. What problem does it solve?
Ethereum block space is scarce. Fees price out small users. L2s execute off-chain and settle on-chain to multiply throughput without sacrificing security.

2. What incentivises correct behaviour?
Optimistic: fraud proof slashing. ZK: mathematical validity proof. Sequencers: reputation and future revenue. Both: the escape hatch preserves user funds.

3. Who pays and who gains?
Users pay sequencer fees (10 to 80 times less than mainnet). Sequencer operators capture the spread between fees collected and L1 posting cost.

4. What can break it?
Sequencer centralisation, upgrade key abuse, bridge hacks, proving system bugs (ZK), and challenge period manipulation (optimistic).

5. What trade-offs were made?
Speed vs decentralisation (centralised sequencer). Compatibility vs security (EVM vs Cairo). Data cost vs availability (rollup vs validium).

6. How could it be improved?
Decentralised sequencing, shared sequencer networks, faster ZK proving, danksharding for data, and cross-L2 interoperability standards.

The cryptoeconomics of L2 is an active research area. Most of the design questions raised here have no settled answer – they are the frontier problems of blockchain infrastructure in 2025.

Key Takeaways: What Every Blockchain Professional Should Know About L2

1. **Ethereum cannot scale alone.** The Trilemma forces a choice: L2s move execution off-chain while inheriting L1 security.
2. **Rollups batch transactions.** 100 to 1,000 users share one L1 transaction, reducing per-user cost by the same factor.
3. **Two proof systems, one goal.** Optimistic rollups use fraud proofs (7-day window, cheap); ZK-rollups use validity proofs (instant, costly to generate).
4. **The ecosystem has tipped.** Base and Arbitrum now process more transactions than Ethereum mainnet. L2 is the default execution layer in 2025.
5. **New trust assumptions appear.** Centralised sequencers can censor and extract MEV. Upgrade keys are a systemic risk. Decentralisation is incomplete.
6. **Economics: follow the sequencer revenue.** Sequencer fees are the business model of every L2. Token holders govern the treasury funded by this revenue.

The L2 landscape changes rapidly. Always verify current TVL, fee levels, and governance status at L2Beat (l2beat.com) before making decisions based on figures from any static source.

References and Further Reading

Core references:

- Buterin, V. (2014). *Notes on Scalability.* ethereum.org
- Ethereum Foundation. (2024). *Layer 2 rollups.* ethereum.org/en/layer-2
- L2Beat. (2025). *L2 Risk Framework and TVL.* l2beat.com
- Optimism. (2022). *The Optimistic Vision.* optimism.io

Technical deep dives:

- EIP-4844 specification: eips.ethereum.org/EIPS/eip-4844
- Arbitrum Nitro whitepaper (2022): offchainlabs.com
- zkEVM Design (Polygon, 2022): polygon.technology/blog

Data and dashboards:

- **L2Beat** – TVL, risk ratings, upgrade keys for all L2s
- **DefiLlama** – cross-chain TVL by protocol
- **Dune Analytics** – on-chain metrics (gas, sequencer revenue)
- **GrowThePie** – L2 user and transaction statistics

Discussion questions:

- Should L2 sequencers be decentralised now, or wait until the technology matures?
- Is a 7-day withdrawal window an acceptable user experience trade-off?
- Who should capture sequencer revenue – users, token holders, or public goods?

This lecture focuses on rollup-based L2s. Payment channels (Lightning Network) and plasma chains are earlier approaches with different trade-offs; see the ethereum.org history of L2 for context.