

# Layer 2 Scaling Solutions

30-minute condensed lecture

Prof. Dr. Jörg Osterrieder

BSc Blockchain, Crypto Economy & NFTs

Spring 2026

**Core question:** Why do Layer 2 rollups reduce Ethereum fees by 20 to 80 times – and what do users give up in return? Four key concepts: the Trilemma, rollup mechanics, optimistic vs ZK, and sequencer economics.

## The core problem:

- Ethereum processes 15 to 30 transactions per second (Visa: 24,000)
- Every block has a fixed gas limit: 30 million gas
- When demand exceeds supply, fees spike through an auction
- DeFi Summer 2021: a simple token swap cost over \$200 in fees

## The Scalability Trilemma:

A blockchain can achieve at most two of: **decentralisation**, **security**, **scalability**.

Ethereum L1 chose the first two.

## Layer 2: The Solution Architecture

**Definition:** A Layer 2 (L2) is a separate system that:

1. Executes transactions cheaply off-chain
2. Posts compressed data or proofs back to Ethereum L1
3. Inherits Ethereum's security for settlement

## Key Insight

L2s do not violate the Trilemma. They move *execution* off-chain while *settlement* remains on Ethereum.

---

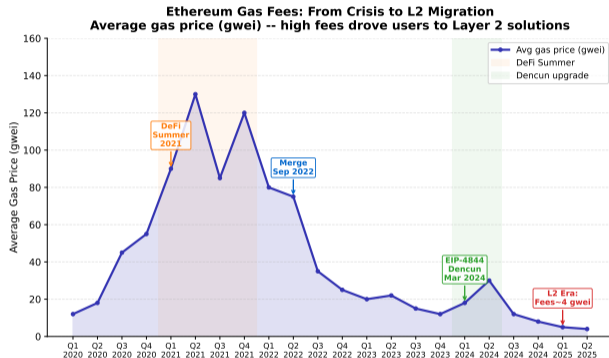
EIP-1559 (Aug 2021) improved fee predictability but did not increase throughput. The Merge (Sep 2022) reduced energy use by 99% but also left throughput unchanged.

# Five Years of Ethereum Gas Fees: The Rise and Fall

## Four eras:

- **2020:** DeFi launches; fees rise to 55 gwei
- **2021 Q2:** DeFi Summer; fees above 130 gwei
- **Sep 2022:** The Merge; fees moderate
- **Mar 2024:** EIP-4844 (Dencun); L2 blob fees drop 10 to 100 times
- **2025:** Mainnet averages 4 gwei; L2 average below 1 gwei

**L2 adoption changed the demand curve for Ethereum block space.**



EIP-4844 added “blob” data – a separate, cheaper data type stored for only 18 days. L2s use blobs instead of calldata, dramatically reducing their L1 posting cost.

# How Rollups Work: Batch Compression Saves Block Space

## The four-step flow:

1. Users submit transactions to L2 (fast, cheap)
2. Sequencer orders and executes off-chain
3. Batch compressor: 100 to 1,000 txs become one blob
4. Batch posted to Ethereum L1 as one transaction

## Cost arithmetic:

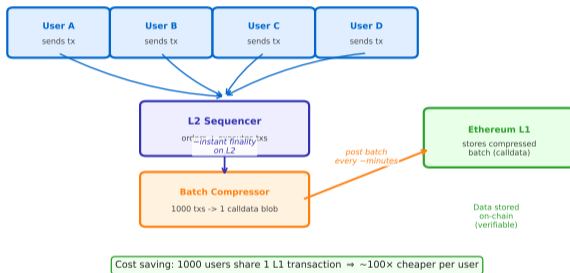
1,000 users share one L1 posting cost.  
Result: 100 to 1,000 times cheaper per user.

## Rollup variants:

- **Optimistic:** assumes validity; fraud proofs if challenged
- **ZK:** posts validity proof; mathematically verified on L1

The sequencer (the entity that orders transactions) is currently centralised in all major L2s. This is the key governance risk: a single entity controls transaction ordering on chains processing billions in daily volume.

## How Rollups Work: Batch Compression Moves Costs Off-Chain Transactions execute on L2; only compressed proofs land on Ethereum



# Two Proof Systems: Optimistic Rollups vs ZK-Rollups

## Optimistic rollups:


- Assume transactions are valid
- Anyone can submit a fraud proof within 7 days
- Withdrawal to L1: 7 days (native)
- EVM-compatible out of the box
- Examples: Arbitrum, Optimism, Base

## ZK-rollups:

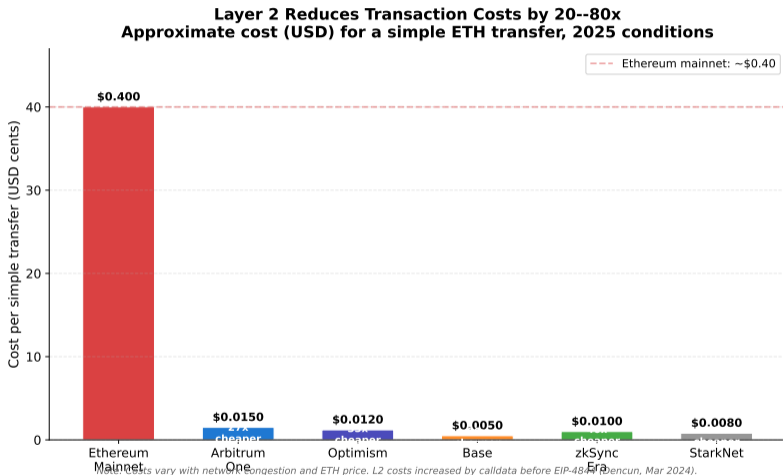
- Prove validity mathematically before posting
- Withdrawal to L1: minutes (no challenge window)
- EVM support: improving (zkEVM reached production 2024)
- Higher off-chain proving cost (GPU-intensive)
- Examples: zkSync Era, StarkNet, Polygon zkEVM

**Trend: ZK is the long-run direction, optimistic leads today.**

Vitalik Buterin: "In the long run, everything will be ZK." Proving costs for ZK-rollups fell by roughly 10 times between 2022 and 2024. Full zkEVM equivalence reached production in 2024 with Polygon zkEVM and zkSync Era.



# The Cost Revolution: 20 to 80 Times Cheaper on Layer 2



Post-EIP-4844 (Dencun, March 2024), Base transactions fell below \$0.001 for a simple transfer. L2 fees fluctuate with Ethereum mainnet blob demand, but remain a fraction of mainnet cost even at peak.

# The L2 Ecosystem in 2025: Who Leads?

## Arbitrum One

- Launched Aug 2021
- TVL: ~\$9B
- Tech: Nitro (optimistic)
- Token: ARB
- Strengths: DeFi, largest TVL

## Base

- Launched Aug 2023
- TVL: ~\$8B
- Tech: OP Stack (optimistic)
- Token: none (Coinbase)
- Strengths: most daily txs, consumer apps

## zkSync Era

- Launched Mar 2023
- TVL: ~\$3B
- Tech: zkEVM (ZK)
- Token: ZK
- Strengths: fastest finality, ZK leader

**Key milestone 2025:** Base now processes more daily transactions than Ethereum mainnet. The centre of Ethereum activity has shifted to Layer 2. L2 TVL total exceeds \$30 billion – up from under \$3 billion in early 2022.

L2Beat ([l2beat.com](https://l2beat.com)) tracks every major L2 with live TVL, risk scores, and upgrade key status. It is the authoritative reference for comparing L2 security models.

## 1. Sequencer Centralisation

- One server controls tx ordering on all major L2s
- Can censor specific addresses
- Can front-run users (MEV)
- If it crashes, L2 halts
- Escape hatch: force-exit via L1 (slow but safe)

## 2. L2 Token Economics

- ARB and OP: governance tokens
- Token holders vote on protocol upgrades
- Sequencer revenue funds DAO treasury
- Retroactive funding rewards past contributors
- Risk: upgrade keys allow instant bridge changes

## 3. MEV on L2

- Sequencer can reorder txs for profit
- Sandwich attacks on DEX trades
- Cross-domain MEV (L1/L2 arbitrage)
- Mitigations: FCFS, encrypted mempool, PBS
- MEV on Arbitrum: estimated \$50M+ per year

Moving from Ethereum to L2 **changes the trust model**: from distributed validators to centralised sequencers. Security is preserved for fund custody; governance and ordering are not yet decentralised.

L2Beat classifies every L2 protocol on a risk matrix covering: data availability, sequencer failure handling, proposer failure handling, state validation, and upgrade key permissions.

## Key Takeaways: Layer 2 Scaling Solutions

1. **Ethereum's bottleneck is block space.** 15 to 30 tx/s is a design choice, not a flaw. L2s work around it by batching execution.
2. **Rollups batch transactions.** 100 to 1,000 users share one L1 transaction, cutting per-user cost proportionally. EIP-4844 made this even cheaper.
3. **Two proof approaches, different trade-offs.**  
Optimistic: cheap, 7-day wait. ZK: costly proving, instant finality. ZK is the long-term direction.
4. **The L2 ecosystem has tipped.** Base surpassed Ethereum mainnet in daily transactions. L2 is no longer experimental – it is the primary execution layer.
5. **New trust actors appear.** Centralised sequencers control ordering. Upgrade keys allow protocol changes. Token holders govern fee distribution. Decentralisation is incomplete.
6. **Follow the sequencer revenue.** The fee spread between user payments and L1 posting cost is the L2 business model. Who captures this revenue is the central governance question.

---

For live data: [l2beat.com](https://l2beat.com) (risk and TVL), [growthepie.xyz](https://growthepie.xyz) (transactions), [defillama.com](https://defillama.com) (cross-chain TVL), [dune.com](https://dune.com) (sequencer revenue metrics).

### Consumer applications:

- **Base:** Farcaster social feeds, friend.tech, onchain gaming
- **Arbitrum:** GMX (perpetuals), Aave, Uniswap V3
- **Optimism:** Synthetix, Velodrome, Worldcoin
- **StarkNet:** Immutable X (NFT gaming), dYdX V4

### Payments and remittances:

- USDC transfers on Base for under \$0.001
- Stripe and PayPal integrating USDC on L2
- Cross-border micro-payments now economically viable

### Enterprise and institutional:

- Polygon zkEVM: enterprise NFT and loyalty programs
- Arbitrum One: institutional DeFi and structured products
- Base: Coinbase's retail bridge for 100M+ users

### Discussion question

If Ethereum is the settlement layer and L2s are the execution layer, who has the power? The L1 validators (security) or the L2 sequencers (user experience and revenue)?

---

The analogy: Ethereum mainnet is like the central bank, and L2s are commercial banks. Settlements flow through the centre; everyday transactions happen at the edges.

## Quiz: Layer 2 Scaling (Part 1 of 2)

**Q1.** What is the approximate maximum throughput of Ethereum mainnet?

- A) 100 tx/s   B) **15–30 tx/s**   C) 1,000 tx/s   D) 7,000 tx/s

## Quiz: Layer 2 Scaling (Part 1 of 2)

**Q1.** What is the approximate maximum throughput of Ethereum mainnet?

A) 100 tx/s   B) **15–30 tx/s**   C) 1,000 tx/s   D) 7,000 tx/s   **Answer: B** – Ethereum processes 15 to 30 transactions per second. The block gas limit caps throughput regardless of validator speed.

**Q2.** What is the main mechanism by which rollups reduce per-user costs?

A) A faster consensus algorithm   B) Removing gas fees entirely   C) **Batching many txs into one L1 transaction**   D) Using a different VM

## Quiz: Layer 2 Scaling (Part 1 of 2)

**Q1.** What is the approximate maximum throughput of Ethereum mainnet?

A) 100 tx/s B) **15–30 tx/s** C) 1,000 tx/s D) 7,000 tx/s **Answer: B** – Ethereum processes 15 to 30 transactions per second. The block gas limit caps throughput regardless of validator speed.

**Q2.** What is the main mechanism by which rollups reduce per-user costs?

A) A faster consensus algorithm B) Removing gas fees entirely C) **Batching many txs into one L1 transaction** D) Using a different VM **Answer: C** – 100 to 1,000 users share the cost of one L1 posting. Cost savings scale with batch size.

**Q3.** The challenge window for optimistic rollups is typically: A) 1 hour B) 24 hours C) **7 days** D) 30 days

## Quiz: Layer 2 Scaling (Part 1 of 2)

**Q1.** What is the approximate maximum throughput of Ethereum mainnet?

A) 100 tx/s B) **15–30 tx/s** C) 1,000 tx/s D) 7,000 tx/s **Answer: B** – Ethereum processes 15 to 30 transactions per second. The block gas limit caps throughput regardless of validator speed.

**Q2.** What is the main mechanism by which rollups reduce per-user costs?

A) A faster consensus algorithm B) Removing gas fees entirely C) **Batching many txs into one L1 transaction** D) Using a different VM **Answer: C** – 100 to 1,000 users share the cost of one L1 posting. Cost savings scale with batch size.

**Q3.** The challenge window for optimistic rollups is typically: A) 1 hour B) 24 hours C) **7 days** D) 30 days **Answer: C** – 7 days gives verifiers worldwide time to detect and prove fraud. This also means native withdrawals take 7 days.

**Q4.** What does a ZK-rollup post to Ethereum L1 to prove correctness?

A) Full transaction data plus a fraud proof B) **A validity proof and compressed state data** C) Only a state root hash D) The full transaction database

## Quiz: Layer 2 Scaling (Part 1 of 2)

**Q1.** What is the approximate maximum throughput of Ethereum mainnet?

A) 100 tx/s B) **15–30 tx/s** C) 1,000 tx/s D) 7,000 tx/s **Answer: B** – Ethereum processes 15 to 30 transactions per second. The block gas limit caps throughput regardless of validator speed.

**Q2.** What is the main mechanism by which rollups reduce per-user costs?

A) A faster consensus algorithm B) Removing gas fees entirely C) **Batching many txs into one L1 transaction** D) Using a different VM **Answer: C** – 100 to 1,000 users share the cost of one L1 posting. Cost savings scale with batch size.

**Q3.** The challenge window for optimistic rollups is typically: A) 1 hour B) 24 hours C) **7 days** D) 30 days **Answer: C** – 7 days gives verifiers worldwide time to detect and prove fraud. This also means native withdrawals take 7 days.

**Q4.** What does a ZK-rollup post to Ethereum L1 to prove correctness?

A) Full transaction data plus a fraud proof B) **A validity proof and compressed state data** C) Only a state root hash D) The full transaction database **Answer: B** – A validity proof (a few kilobytes) plus minimal state data. The proof is verified on-chain instantly, eliminating the challenge window.

**Q5.** Which L2 was launched by Coinbase in August 2023?

A) Arbitrum B) Optimism C) **Base** D) zkSync

## Quiz: Layer 2 Scaling (Part 1 of 2)

**Q1.** What is the approximate maximum throughput of Ethereum mainnet?

A) 100 tx/s B) **15–30 tx/s** C) 1,000 tx/s D) 7,000 tx/s **Answer: B** – Ethereum processes 15 to 30 transactions per second. The block gas limit caps throughput regardless of validator speed.

**Q2.** What is the main mechanism by which rollups reduce per-user costs?

A) A faster consensus algorithm B) Removing gas fees entirely C) **Batching many txs into one L1 transaction** D) Using a different VM **Answer: C** – 100 to 1,000 users share the cost of one L1 posting. Cost savings scale with batch size.

**Q3.** The challenge window for optimistic rollups is typically: A) 1 hour B) 24 hours C) **7 days** D) 30 days **Answer: C** – 7 days gives verifiers worldwide time to detect and prove fraud. This also means native withdrawals take 7 days.

**Q4.** What does a ZK-rollup post to Ethereum L1 to prove correctness?

A) Full transaction data plus a fraud proof B) **A validity proof and compressed state data** C) Only a state root hash D) The full transaction database **Answer: B** – A validity proof (a few kilobytes) plus minimal state data. The proof is verified on-chain instantly, eliminating the challenge window.

**Q5.** Which L2 was launched by Coinbase in August 2023?

A) Arbitrum B) Optimism C) **Base** D) zkSync **Answer: C** – Base is built on the OP Stack (Optimism's open-source codebase) and leverages Coinbase's 100M+ user distribution.

## Quiz: Layer 2 Scaling (Part 2 of 2)

**Q6.** What is the primary risk of a centralised sequencer?

- A) Increases gas fees   B) Cannot process ZK proofs   C) **Can censor transactions and extract MEV**   D) Requires a trusted setup

## Quiz: Layer 2 Scaling (Part 2 of 2)

**Q6.** What is the primary risk of a centralised sequencer?

A) Increases gas fees   B) Cannot process ZK proofs   C) **Can censor transactions and extract MEV**   D) Requires a trusted setup

**Answer: C** – A centralised sequencer controls transaction ordering, enabling censorship and front-running. Users can force-exit via L1 but cannot prevent sequencer MEV.

**Q7.** EIP-4844 (Dencun, March 2024) reduced L2 costs by introducing: A) A new proof system   B) **Blob data with a separate cheaper fee market**   C) Faster block times   D) Decentralised sequencing

## Quiz: Layer 2 Scaling (Part 2 of 2)

**Q6.** What is the primary risk of a centralised sequencer?

A) Increases gas fees   B) Cannot process ZK proofs   C) **Can censor transactions and extract MEV**   D) Requires a trusted setup

**Answer: C** – A centralised sequencer controls transaction ordering, enabling censorship and front-running. Users can force-exit via L1 but cannot prevent sequencer MEV.

**Q7.** EIP-4844 (Dencun, March 2024) reduced L2 costs by introducing: A) A new proof system   B) **Blob data with a separate cheaper fee market**   C) Faster block times   D) Decentralised sequencing **Answer: B** – Blobs are stored for only 18 days (not permanently like calldata), enabling a separate, much cheaper fee market for L2 data posting.

**Q8.** The Scalability Trilemma states that a blockchain can achieve at most two of: A) **Decentralisation, security, scalability**   B) Privacy, scalability, governance   C) Finality, throughput, consensus   D) Anonymity, censorship resistance, performance

## Quiz: Layer 2 Scaling (Part 2 of 2)

**Q6.** What is the primary risk of a centralised sequencer?

A) Increases gas fees   B) Cannot process ZK proofs   C) **Can censor transactions and extract MEV**   D) Requires a trusted setup

**Answer: C** – A centralised sequencer controls transaction ordering, enabling censorship and front-running. Users can force-exit via L1 but cannot prevent sequencer MEV.

**Q7.** EIP-4844 (Dencun, March 2024) reduced L2 costs by introducing: A) A new proof system   B) **Blob data with a separate cheaper fee market**   C) Faster block times   D) Decentralised sequencing **Answer: B** – Blobs are stored for only 18 days (not permanently like calldata), enabling a separate, much cheaper fee market for L2 data posting.

**Q8.** The Scalability Trilemma states that a blockchain can achieve at most two of: A) **Decentralisation, security, scalability**   B) Privacy, scalability, governance   C) Finality, throughput, consensus   D) Anonymity, censorship resistance, performance **Answer: A** – Buterin (2014): optimising for all three simultaneously requires unacceptable trade-offs. L2s “cheat” by moving execution off-chain.

**Q9.** Which rollup type offers faster withdrawal from L2 to L1?

A) Optimistic rollups   B) **ZK-rollups**   C) Both are the same   D) Depends only on congestion

## Quiz: Layer 2 Scaling (Part 2 of 2)

**Q6.** What is the primary risk of a centralised sequencer?

A) Increases gas fees   B) Cannot process ZK proofs   C) **Can censor transactions and extract MEV**   D) Requires a trusted setup

**Answer: C** – A centralised sequencer controls transaction ordering, enabling censorship and front-running. Users can force-exit via L1 but cannot prevent sequencer MEV.

**Q7.** EIP-4844 (Dencun, March 2024) reduced L2 costs by introducing: A) A new proof system   B) **Blob data with a separate cheaper fee market**   C) Faster block times   D) Decentralised sequencing **Answer: B** – Blobs are stored for only 18 days (not permanently like calldata), enabling a separate, much cheaper fee market for L2 data posting.

**Q8.** The Scalability Trilemma states that a blockchain can achieve at most two of: A) **Decentralisation, security, scalability**   B) Privacy, scalability, governance   C) Finality, throughput, consensus   D) Anonymity, censorship resistance, performance **Answer: A** – Buterin (2014): optimising for all three simultaneously requires unacceptable trade-offs. L2s “cheat” by moving execution off-chain.

**Q9.** Which rollup type offers faster withdrawal from L2 to L1?

A) Optimistic rollups   B) **ZK-rollups**   C) Both are the same   D) Depends only on congestion **Answer: B** – ZK validity proofs are verified on L1 immediately; optimistic rollups require a 7-day challenge window before funds can be released.

**Q10.** L2Beat classifies L2 protocols primarily on: A) Transaction speed and token price   B) **Security model, data availability, and upgrade key control**   C) Number of users   D) Gas savings vs mainnet

## Quiz: Layer 2 Scaling (Part 2 of 2)

**Q6.** What is the primary risk of a centralised sequencer?

A) Increases gas fees   B) Cannot process ZK proofs   C) **Can censor transactions and extract MEV**   D) Requires a trusted setup

**Answer: C** – A centralised sequencer controls transaction ordering, enabling censorship and front-running. Users can force-exit via L1 but cannot prevent sequencer MEV.

**Q7.** EIP-4844 (Dencun, March 2024) reduced L2 costs by introducing: A) A new proof system   B) **Blob data with a separate cheaper fee market**   C) Faster block times   D) Decentralised sequencing **Answer: B** – Blobs are stored for only 18 days (not permanently like calldata), enabling a separate, much cheaper fee market for L2 data posting.

**Q8.** The Scalability Trilemma states that a blockchain can achieve at most two of: A) **Decentralisation, security, scalability**   B) Privacy, scalability, governance   C) Finality, throughput, consensus   D) Anonymity, censorship resistance, performance **Answer: A** – Buterin (2014): optimising for all three simultaneously requires unacceptable trade-offs. L2s “cheat” by moving execution off-chain.

**Q9.** Which rollup type offers faster withdrawal from L2 to L1?

A) Optimistic rollups   B) **ZK-rollups**   C) Both are the same   D) Depends only on congestion **Answer: B** – ZK validity proofs are verified on L1 immediately; optimistic rollups require a 7-day challenge window before funds can be released.

**Q10.** L2Beat classifies L2 protocols primarily on: A) Transaction speed and token price   B) **Security model, data availability, and upgrade key control**   C) Number of users   D) Gas savings vs mainnet **Answer: B** – L2Beat’s risk framework focuses on the properties that determine whether users can always recover funds, regardless of operator behaviour.