Blockchain Scalability Trilemma

BSc Blockchain, Crypto Economy & NFTs

Course Instructor

Module A: Blockchain Foundations

[COMIC: A blockchain engineer at a carnival game with three prizes behind glass labeled "Decentralization," "Security," and "Scalability." Sign reads: "WIN ANY TWO! (Sorry, three is physically impossible)"—the engineer looks frustrated while holding only two tickets]

The Fundamental Constraint

- Bitcoin chose decentralization + security = 7 TPS
- EOS chose security + scalability = 21 block producers
- Traditional databases chose scalability + security = single point of failure
- Can we escape this constraint? Layer 2 and sharding try...

Vitalik Buterin coined "scalability trilemma" in 2017—still no one has solved it

Learning Objectives

By the end of this lesson, you will be able to:

- Explain the blockchain scalability trilemma
- Analyze trade-offs between security, decentralization, and scalability
- Understand Layer 1 scalability bottlenecks
- Compare throughput limitations across blockchains
- Evaluate vertical vs. horizontal scaling approaches
- Recognize emerging scalability solutions

Building on L10: Consensus Comparison

The Problem: How do we scale blockchains?

The Challenge

Blockchains face a fundamental limitation: achieving global-scale transaction throughput while preserving decentralization and security. Bitcoin processes 7 TPS while Ethereum handles 15 TPS—far below the requirements for mass adoption or competition with traditional payment systems.

Why It Matters

- Bitcoin: 7 TPS vs. Visa: 65,000 TPS—blockchains cannot serve billions of users
- Ethereum's 2017 CryptoKitties congestion halted the network, with gas prices spiking 1000%

What We Need

- Understanding design constraints
- Scale to handle global transaction volumes without sacrificing decentralization or security

The Cryptoeconomics Question

How do we optimize for competing goals: throughput, decentralization, and security?

Today's lesson: How the Scalability Trilemma framework addresses this challenge

What is the Scalability Trilemma?

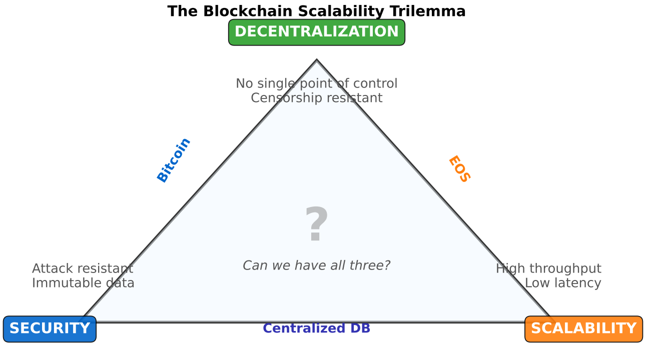

Coined by Vitalik Buterin: a blockchain can achieve at most two of three properties. | → Problem: How do we scale blockchains? — The trilemma is a fundamental constraint – every blockchain project must choose which properties to optimize

Why Can't Blockchains Have It All?

Three Properties:

- **Decentralization:** No single controller
- **Security:** Attack resistant
- **Scalability:** High TPS, low latency

Trade-Offs:

- BTC: Decentral + Secure = 7 TPS
- EOS: Scale + Secure = 21 nodes
- DB: Scale + Secure = centralized

Core Challenges:

- Large blocks = fewer validators
- Fewer validators = faster consensus
- Optimizing one degrades others

Key Insight:
No free lunch - trade-offs are fundamental

Compare the approaches shown above

Why Do These Three Properties Matter?

Decentralization

- Censorship resistance
- Fault tolerance
- Trustlessness
- No single point of failure

Security

- Immutability
- Double-spend prevention
- Attack resistance
- Data integrity

Scalability

- High throughput (TPS)
- Low latency (finality)
- Low transaction costs
- Mass adoption support

Why Blockchain Needs All Three:

- DeFi applications need high throughput AND security
- Mass adoption requires scalability without sacrificing decentralization
- Competition with traditional payment systems (Visa: 24,000 TPS)

Compare the approaches shown above

What Limits Layer 1 Throughput?

Block Size:

- Larger = more TXs, slower propagation
- BTC: 1-4 MB — ETH: gas limit

Block Time:

- Faster = more orphans/forks
- BTC: 10m — ETH: 12s — SOL: 0.4s

State Growth:

- BTC: 500 GB — ETH: 1 TB
- Full nodes need storage

Computation:

- Every node verifies all TXs
- Smart contracts expensive

→ *Problem: How do we scale blockchains? — Bitcoin's 7 TPS limit comes from 1MB blocks every 10 minutes – a deliberate security choice, not a technical limitation*

Layer 1 Throughput Comparison

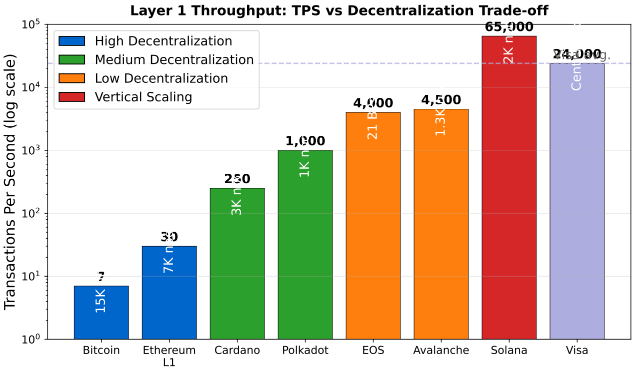

Inverse correlation: higher TPS generally means fewer nodes and lower decentralization.

Should Bitcoin Increase Block Size?

Big Block Camp:

- Increase to 8+ MB
- Higher throughput, low fees
- Risk: centralization

Small Block Camp:

- Preserve decentralization
- Scale via Layer 2
- Accept higher on-chain fees

Outcome (2017):

- BCH fork: 8 MB blocks
- BTC: SegWit (1-4 MB)
- Market: BTC ¿ BCH

Lesson:
Decentralization prioritized; Layer 2 for scale

Compare the approaches shown above

Vertical vs. Horizontal Scaling

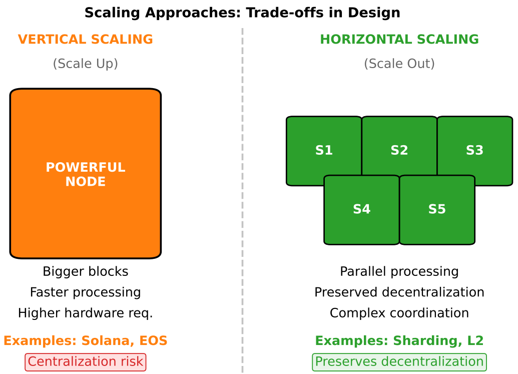

Most major blockchains pursue horizontal scaling to preserve decentralization.

How Should We Scale: Up or Out?

Vertical Scaling (Scale Up)

- Increase capacity of individual nodes
- Require more powerful hardware
- Examples: Solana, EOS

Advantages:

- Simpler implementation
- Immediate throughput gains

Disadvantages:

- Raises barrier to run nodes
- Centralization risk
- Hits physical limits

Horizontal Scaling (Scale Out)

- Distribute load across many nodes
- Parallel processing
- Examples: Sharding, Layer 2, Rollups

Advantages:

- Preserves decentralization
- Theoretically unbounded scaling

Disadvantages:

- Complex implementation
- Cross-shard communication overhead
- Longer development timelines

→ *Problem: How do we scale blockchains? — Vertical scaling (bigger blocks) centralizes; horizontal scaling (sharding, L2) preserves decentralization*

Problem Check-In: Scalability Trilemma

Recall Our Problem

How do we scale blockchains?

What We've Learned So Far

- The scalability trilemma states blockchains can optimize only 2 of 3 properties: decentralization, security, scalability
- Layer 1 bottlenecks include block size, block time, state growth, and computation limits
- Bitcoin chose security/decentralization (7 TPS); Solana chose speed/security (65k TPS but 33 validators)

Still to Address

- Layer 2 solutions (rollups) and sharding as approaches to escape the trilemma
- Can we achieve 10,000+ TPS without sacrificing decentralization?

Think About

- Based on what you've seen, how would *you* solve this problem?
- What trade-offs do you expect?

Pause and reflect: How does what we've learned so far address "How do we scale blockchains?"?

Layer 2 Scaling Solutions

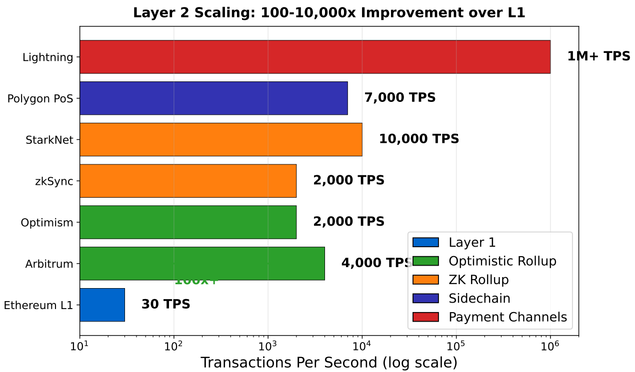

Layer 2 achieves 100-10,000x improvement while inheriting L1 security.

What Layer 2 Solutions Exist?

Payment Channels:

- Lightning Network
- 1 on-chain, unlimited off-chain
- Micropayments

Rollups:

- Bundle 100s of TXs off-chain
- Optimistic: fraud proofs (7-day challenge window)
- ZK: validity proofs (instant cryptographic verification)

Sidechains:

- Independent chain
- Own consensus
- Weaker security

Adoption: Ethereum L2 TVL $54B+ (Dec 2024, 205% YoY)

Compare the approaches shown above

Scaling Solutions Comparison

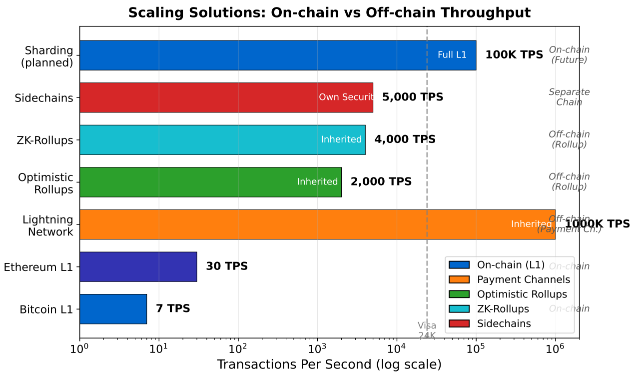

On-chain solutions inherit full L1 security; off-chain achieves higher TPS with varying trust assumptions.

State Growth Problem

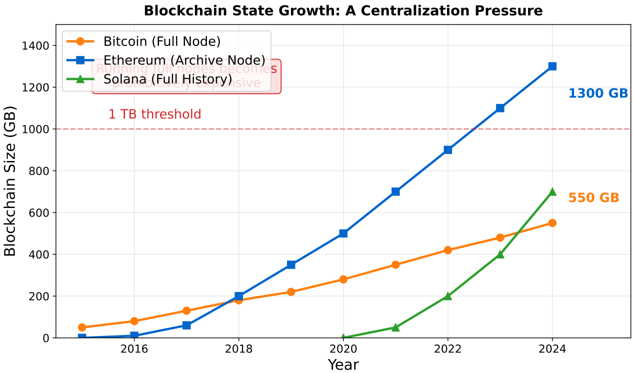

Running full nodes becomes prohibitively expensive as state grows.

How Do We Control State Growth?

Challenge:

- Unbounded growth
- ETH state: 100 GB
- Sync takes days/weeks
- Centralization pressure

Solutions:

- **State Expiry:** evict inactive, proof to recover
- **Pruning:** keep recent blocks only
- **Rent:** pay to store (Cosmos, EOS)

→ *Problem: How do we scale blockchains? — State growth is blockchain's hidden scaling problem – Ethereum's state grows 60GB/year, pricing out home validators*

How Does Sharding Scale at Layer 1?

Concept:

- Split into parallel shards
- Each processes subset of TXs
- Total TPS = shards x per-shard

Ethereum Roadmap:

- Danksharding (Ethereum's data availability approach using temporary "blobs" for cheap rollup data) for DA (Data Availability – the guarantee that transaction data is published and accessible)
- EIP-4844 (March 2024)
- Target: 100K+ TPS

Challenges:

- Cross-shard complexity
- Shard takeover attacks
- State management

Other Chains:
Zilliqa, NEAR, Elrond

Compare the approaches shown above

Trilemma Positioning

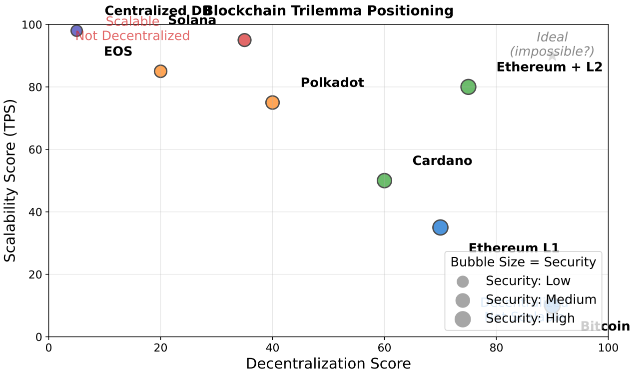

Ethereum + L2 approaches the "ideal" corner without sacrificing security.

Trilemma Trade-offs by Blockchain

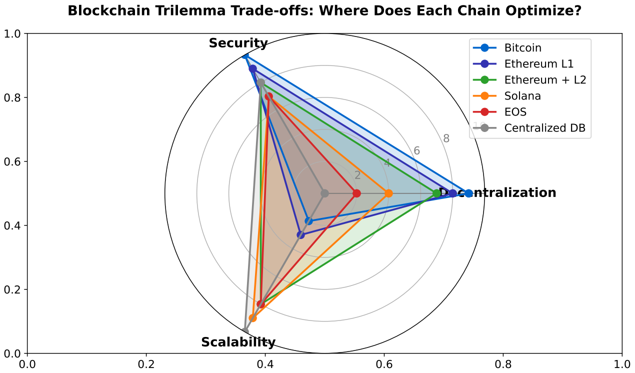

Radar chart shows each blockchain's optimization choices. Bitcoin maximizes security/decentralization; Solana prioritizes scalability.

Can We Break the Trilemma?

Optimistic:

- L2 separates layers
- Rollups keep security
- ZK compresses compute

Pessimistic:

- L2 adds trust
- More complexity
- DA (Data Availability) bandwidth limits

Consensus:

- Mitigate, not solve
- Modular design
- L1 security, L2 scale

Navigate the trilemma intelligently rather than "solve" it

Problem Solved? Scalability Trilemma

The Original Problem

How do we scale blockchains to handle global transaction volume?

How Scalability Solutions Address It

- **Layer 2 solutions (rollups, state channels):** Move computation off-chain, achieve 100-10,000x throughput improvement while inheriting L1 security
- Split state across parallel chains, enabling horizontal scaling to theoretically unbounded throughput
- Separate consensus, data availability, and execution layers to optimize each independently

Remaining Limitations

- **Layer 2 introduces new trust assumptions:** Sequencer centralization, bridge security risks, and withdrawal delays
- Increased complexity and fragmented liquidity in exchange for higher throughput

Open Questions

- Will rollups or sharding dominate as the long-term scaling solution?
- Risk: Over-optimization for single metric

Scalability Trilemma framework helps navigate trade-offs but doesn't eliminate fundamental constraints

The Economics of Scalability Trilemma

Incentive Structure

- Optimizing multiple competing goals
- Choose based on use case requirements
- Gain in one dimension, sacrifice in another

Economic Security

- Attack cost must exceed potential gain
- Honest behavior = Nash equilibrium

Key Economic Question

Who Pays, Who Earns?

Gain in one dimension, sacrifice in another

Design Principle

$$\text{Attack Cost} > \text{Potential Gain}$$

Cryptoeconomic security: Honest behavior must be the Nash equilibrium

Design Space: Scalability Trilemma

Alternatives Considered

1. Different points on trade-off frontier
2. Layer 2 solutions, hybrid approaches

Trade-offs Made

- Every design optimizes some properties
- ... at the expense of others

Design Questions

- What would YOU change?
- What's optimized? What's sacrificed?
- Are there other approaches?

Key Insight

No Perfect Solution

All blockchain designs involve trade-offs between decentralization, security, and scalability.

Every design is a trade-off. Understanding alternatives reveals the "why" behind choices.

Failure Modes

Failure Modes: Scalability Trilemma

Critical Failure Mode

- **What breaks:** Over-optimization for single metric
- **Why it happens:** Economic incentives misaligned

Root Cause

- Assumption violated
- Incentive structure broken
- External shock

Historical Context

- Multiple real-world failures documented
- Patterns repeating across protocols

Early Warning Signs

! Unusual economic behavior
! Incentive misalignment
! Centralization drift

Prediction: What could cause this to fail? How would you detect it early?

Blockchain throughput traffic jam

[COMIC: A massive traffic jam on a single-lane road (Bitcoin at 7 TPS) while parallel roads (Layer 2) and multi-lane highways (sharding) show alternative routes—some cars taking the toll road labeled "Roll up: $0.01 fee, 1000x faster"]

Scaling Without Compromise

- Layer 1: the secure but congested main highway (7-30 TPS)
- Layer 2 rollups: toll roads that inherit L1 security (1000+ TPS)
- Ethereum's strategy: L1 for settlement, L2 for execution
- Result: $54B TVL on L2s, 90%+ fee reduction post-EIP-4844

CryptoKitties 2017: gas fees hit $100—today's L2s: same tx for $0.01

Key Takeaways

Trilemma:

- Cannot maximize all three
- L1: block size, time, state
- Vertical = centralization
- Horizontal = decentralization

Next Lesson: L12 – Lab: Block Explorer

Solutions:

- L2: 100-1000x throughput
- State growth = risk
- Real TPS ¡ theoretical

Philosophy:
BTC: decentral — SOL: scale — ETH: L2 balance

Compare the approaches shown above

Discussion Questions

1. Why cannot a blockchain simply increase block size indefinitely?
2. How do Layer 2 solutions differ from sidechains in terms of security?
3. What are the implications of state growth for decentralization?
4. Can the trilemma be "solved" or only mitigated?
5. How does sharding introduce new security challenges?
6. Why do most blockchains have lower actual TPS than theoretical TPS?

Key point: Discussion Questions

Continued

Next Lesson Preview: L12 Lab - Block Explorer Analysis

Lab Activities:

- Navigate Etherscan, Blockstream
- Analyze TX details
- Trace mempool to confirmation
- Examine block structure
- Forensic analysis

Preparation:

- Review L6 TX structure
- Access Etherscan.io
- Access Blockstream.info

Compare the approaches shown above

Quiz Questions (1/4)

Q1. What is the blockchain scalability trilemma?
A) The choice between PoW, PoS, and DPoS B) Achieving only two of: decentralization, security, scalability
C) The trade-off between speed and cost D) The conflict between miners and validators

Answer: B – The trilemma states a blockchain can optimize at most two of three properties simultaneously.

Q2. Why does increasing block size lead to centralization?
A) It reduces miner rewards B) Fewer nodes can store/propagate large blocks
C) It makes attacks easier D) It decreases transaction fees

Answer: B – Larger blocks require more bandwidth and storage, raising barriers to running full nodes.

Q3. Bitcoin processes approximately how many transactions per second (TPS)?
A) 3 TPS B) 7 TPS C) 30 TPS D) 100 TPS

Answer: B – Bitcoin's design prioritizes decentralization and security over throughput.

Q4. What is vertical scaling in blockchain context?
A) Adding more nodes to the network B) Increasing capacity of individual nodes
C) Creating Layer 2 solutions D) Implementing sharding

Answer: B – Vertical scaling means requiring more powerful hardware per node.

Q5. Which blockchain outcome resulted from the 2017 block size debate?
A) Ethereum Classic fork B) Bitcoin Cash fork with larger blocks
C) Bitcoin unlimited blocks D) Complete consensus on 8 MB blocks

Answer: B – Bitcoin Cash forked with 8 MB blocks while Bitcoin implemented SegWit.

Quiz Questions (2/4)

Q6. What is the primary advantage of Layer 2 solutions?
A) Complete decentralization B) Higher throughput while inheriting L1 security
C) Lower hardware requirements D) Simpler code implementation

Answer: B – Layer 2 achieves 100-10,000x scalability while settling to secure Layer 1.

Q7. How does horizontal scaling differ from vertical scaling?
A) It uses better hardware B) It distributes load across many nodes/shards
C) It reduces security D) It eliminates the need for consensus

Answer: B – Horizontal scaling parallelizes processing to preserve decentralization.

Q8. What is sharding?
A) Compressing blockchain data B) Splitting blockchain into parallel processing shards
C) Encrypting transactions D) A type of consensus algorithm

Answer: B – Sharding divides the network to process transactions in parallel.

Q9. Which Layer 2 solution uses fraud proofs?
A) ZK-Rollups B) Payment channels C) Optimistic Rollups D) Sidechains

Answer: C – Optimistic Rollups assume validity and allow challenges via fraud proofs.

Q10. What is the state growth problem?
A) Blockchain size growing unbounded over time B) Decreasing transaction volume
C) Fewer validators joining D) Rising energy consumption

Answer: A – As blockchain accumulates data, storage requirements create centralization pressure.

Quiz Questions (3/4)

Q11. How do ZK-Rollups validate transactions?
A) Through fraud proofs B) Using cryptographic validity proofs
C) By running full nodes D) Through validator voting

Answer: B – ZK-Rollups use zero-knowledge proofs to guarantee transaction validity.

Q12. What was Ethereum's EIP-4844 proto-danksharding (March 2024)?
A) Full execution sharding B) Data availability sharding to reduce L2 costs
C) Transition to PoS D) Gas fee elimination

Answer: B – Proto-danksharding adds blob transactions for cheaper L2 data availability.

Q13. Which property does Bitcoin prioritize over scalability?
A) Speed and cost B) Decentralization and security
C) Smart contract flexibility D) Energy efficiency

Answer: B – Bitcoin's design maintains decentralization and security at the cost of throughput.

Q14. What is the main security concern with sidechains?
A) They are too slow B) They have separate validator sets (weaker security)
C) They cannot process smart contracts D) They require PoW consensus

Answer: B – Sidechains do not inherit L1 security guarantees like rollups do.

Q15. How does the Lightning Network achieve high throughput?
A) Larger blocks B) Off-chain transactions in payment channels
C) Faster block times D) More miners

Answer: B – Lightning opens channels with one on-chain TX, then unlimited off-chain payments.

Quiz

Quiz Questions (4/4)

Q16. What is state expiry in Ethereum's roadmap?
A) Deleting old transactions B) Evicting inactive state, requiring proofs to re-activate
C) Eliminating smart contracts D) Reducing gas fees to zero

Answer: B – State expiry removes inactive data to limit state growth while allowing recovery.

Q17. Why is cross-shard communication challenging?
A) It requires complex coordination between shards B) It eliminates security
C) It is impossible to implement D) It reduces decentralization

Answer: A – Transactions spanning shards need atomic execution and additional overhead.

Q18. Which blockchain has approximately 65,000 theoretical TPS?
A) Bitcoin B) Ethereum C) Solana D) Cardano

Answer: C – Solana uses vertical scaling with high hardware requirements to achieve high TPS.

Q19. What is the Total Value Locked (TVL) in Ethereum Layer 2 solutions as of late 2024?
A) $5 billion B) $50+ billion C) $200 billion D) $500 billion

Answer: B – Ethereum L2 TVL reached $54B by Dec 2024 (205% YoY growth).

Q20. Can the scalability trilemma be fully "solved"?
A) Yes, with sharding B) Yes, with Layer 2 C) No, only mitigated through trade-offs
D) Yes, Solana solved it

Answer: C – The trilemma represents fundamental trade-offs; solutions mitigate but don't eliminate them.