




The State of Solana: Network Health and Resilience

By Mia Galvin and Kevin Kelly

October 2026

From accelerating traditional finance (TradFi) adoption to reimagining consensus, 2026 marks a pivotal point for Solana. As the ecosystem enters a new stage of maturity, researchers from the Fidelity Center for Applied Technology (FCAT®) blockchain team evaluate network architecture, reliability, and performance history to measure Solana's health ahead of the Alpenglow upgrade and create a baseline for future analysis.



Introduction

As the Solana ecosystem matures, its value proposition is becoming increasingly defined by tangible utility, most notably through the growth of global payment and decentralized finance use cases. The ongoing migration of traditional finance on-chain signifies trust and reliance on blockchain infrastructure, even by non-crypto natives.

Among leading blockchain networks, Solana has experienced some of the fastest real-world asset (RWA) growth in the form of tokenized RWAs, with value rising roughly four times in one year, from ~\$1B to ~\$4.4B as of October 2026¹. The \$16B stablecoin supply on Solana is also substantial, with over \$2 trillion in volume processed during the first quarter of 2026 alone². Additionally, as of October 2026, around 8% of SOL is held by ETFs and Digital Asset Treasuries (DATs)³, signaling institutional confidence. These figures illustrate rapid market growth, but structural integrity remains a critical benchmark for the network's long-term viability.

For Solana to successfully scale and innovate, its software, infrastructure, and validator set must support ecosystem growth while preserving the network's permissionless design and resiliency. Foundational architectural components—such as rapid stake-weighted leader rotation, transaction forwarding, parallel execution, and stake-based consensus—enable speed and continuity. However, these components can create dependencies involving infrastructure centralization, network latency, leader predictability, and transaction-routing control.

This report evaluates Solana through three specific lenses:

- 01 **Network resilience:** Solana's ability to continue processing and finalizing transactions despite outages, software bugs, infrastructure failures, and malicious attacks
- 02 **Economic security and control:** Whether incentive, penalty, and governance mechanisms align participant behavior with the network's security, availability, and long-term health
- 03 **Performance and usability:** Whether Solana can consistently deliver fast, reliable, and cost-effective transaction execution under real-world operating conditions

Section I:

Network resilience

Can the Solana network continue to operate reliably in the face of technical failures, malicious actors, and adverse network conditions?

Validator count and stake distribution provide useful starting points to evaluate network resilience, but neither fully capture the shared dependencies that can cause otherwise independent validators to fail simultaneously. Therefore, a measure of Solana's resilience must also consider factors like validator participation, stake concentration, geography, infrastructure, and client diversity.

Validator participation and quality

Solana's monolithic architecture places substantial computational and networking demands on validators, creating a tradeoff between performance and accessible participation. Operators require high-performance hardware, reliable bandwidth, and technical support to process heavy workloads and maintain low-latency communication with the network.

Given these requirements, operators are generally advised to use dedicated bare-metal servers with substantial high-clock CPU capacity, high disk input/output performance, and large NVMe storage drives, supplemented by backup servers.

Annual operating costs can start at around \$60,000 and increase significantly depending on hardware costs, server redundancy, and SOL-denominated on-chain voting expenses⁴. These requirements present a meaningful barrier to entry that potentially limits the diversity of the validator set to well-capitalized entities.


As of Q3 2026, Solana had approximately 717 active validator identities. However, this figure measures vote accounts rather than fully independent operators. A single operator may control multiple validator identities, and separate operators may depend on common infrastructure, software, or sources of delegated stake. With this in mind, validator count should realistically be treated as a measure of participation, not a complete measure of independent control or operator diversity.

Stake concentration

Stake is distributed unevenly across the active validator set. As of Q3 2026, approximately 20 validator identities collectively controlled more than one-third of active stake, forming Solana's stake-based superminority. Under Solana's Byzantine fault-tolerant consensus model, a coordinated coalition above this threshold could prevent the remaining validators from reaching the two-thirds majority required for finality, threatening network liveness and potentially enabling censorship. This superminority group, however, would not gain unrestricted control over user assets or the ability to rewrite finalized transactions.

The one-third threshold is calculated at the validator-identity level, but several highly ranked vote accounts belong to the same publicly-identified operators. For example, Figment operates two of the top five validators within the superminority, while Kraken and Coinbase each operate multiple validators representing significant stake. As a result, aggregating related vote accounts by operator would present a different picture of consensus control than counting each account independently. Fortunately, these entities have significant economic exposure to the network, making collusion against their best interests.

In fact, institutional operators can strengthen the validator set through rigorous operational security controls, regulated operations, and external assurances (such as SOC 2 attestations). These characteristics may reduce certain operational and counterparty risks, but they also introduce a different form of correlated risk when large operators share legal jurisdictions, regulatory obligations, or institutional dependencies. Operators exposed to the same sanctions, court orders, or compliance requirements could respond similarly, reinforcing the importance of ownership and geographic diversity.



The figure above shows Solana validators logarithmically ordered by active stake, identifying the number of validator identities required to cross the one-third superminority and two-thirds supermajority thresholds. At around 80 validators, the supermajority could conduct an attack on the network and gain full control over consensus and finality.

Quantifying network decentralization

For Solana, the commonly reported value measures the number of validator identities required to control more than one-third of active stake. Under this method, recent estimates place the coefficient at approximately 20, equivalent to the number of validator identities in the superminority. The result may change when vote accounts are aggregated by controlling operators or stake ownership.

Cross-blockchain comparisons should be interpreted cautiously because studies may use different control thresholds and count a wide variety of entities, including validators, validator keys, mining pools, and operators. A Solana coefficient based on validator identities is not directly comparable with a Bitcoin mining-pool or Ethereum staking-pool coefficient without normalized methodology.


The stake-based Nakamoto Coefficient also measures only one form of concentration. As of Q3 2026, Solana's one-third stake threshold was concentrated across approximately 18 validators, two countries, two autonomous system numbers (ASNs), and two base-client implementations⁵. This suggests potential single points of failure at various levels, which would indicate that infrastructure and software diversity may be as important to network resilience as stake distribution.

The Nakamoto Coefficient indicates the minimum number of entities required to cross a network's critical control threshold that would disrupt finality.

A cost-to-attack analysis provides another measure of theoretical economic security. At current stake levels and SOL prices, acquiring one-third of active stake would represent approximately \$15 billion in market value. This is not a literal attack budget because large-scale acquisition would face liquidity constraints, price pressure, and operational requirements associated with deploying and managing the stake. An attacker might instead pursue less expensive methods, including credential theft, social engineering, or disruption of shared infrastructure. Cost-to-attack analysis is best understood as a measure of theoretical economic security rather than a definitive estimate of practical attack feasibility.

Geographic and infrastructure distribution

Latency considerations incentivize validators to co-locate near other nodes, encouraging the use of bare-metal servers in Tier 3 and Tier 4 European and U.S. data centers. As of Q3 2026, approximately 59% of network stake was located in Germany, the Netherlands, and the United States, including roughly 40% concentrated in Frankfurt and Amsterdam⁶. These estimates rely on IP geolocation and may not be entirely accurate, but they illustrate meaningful stake-weighted concentration in Frankfurt and Amsterdam alone.




Geographic and infrastructure concentration are distinct risks. Validators in different regions may rely on the same facilities, hosts, ASNs, or private communication infrastructure, while validators in the same region may use separate providers. Relevant correlated risks include:

- » Geography and legal jurisdiction
- » Physical data center and hosting provider
- » ASN or network provider
- » Private propagation and communication infrastructure

Specialized hosting providers like TerraSwitch and Latitude tailor hardware offerings to Solana’s workload, improving uptime and efficiency, while contributing to provider concentration. To address this risk, the Solana Foundation Delegation Program established an ASN concentration threshold below 25% and a data-center threshold of 15% for participating validators. Following previous work by the Foundation to reduce TerraSwitch dominance, an October 2026 snapshot placed the ASN at approximately 29% of network stake, still above the program’s ceiling⁷. This indicates a gap between the network’s desired risk profile and its actual infrastructure distribution.

In August 2026, a large operator experienced a routing failure that knocked approximately 90 validators offline for 33 minutes⁸. The affected stake was roughly 4.5 percentage points below the one-third threshold, narrowly avoiding a critical situation that would have halted transaction finality.



The chart above demonstrates Solana stake distribution by ASN. Multiple hosting companies may share the same provider, and some supply both hardware and network routing.

The risks discussed are partially mitigated through node-operations support, server failovers, backup power, and private network layers such as DoubleZero, which reduce dependence on the public internet. Continued infrastructure hardening and geographic diversification remain important, particularly against geopolitical, routing, and supply-chain disruptions.

Technical safeguards are also being supplemented by formal risk-management and incident-response processes. In early 2026, the Solana Foundation introduced the Solana Trust, Resilience and Infrastructure for DeFi Enterprises (or “STRIDE”) framework and the SIRN incident-response network. STRIDE evaluates threats across governance, applications, and infrastructure, while SIRN coordinates real-time responses to identified incidents.


Client diversity

Solana's core validator implementation is Agave, the Anza-maintained successor to the original Solana Labs client. Jito Labs subsequently forked the standard client to create Jito-Solana, which, at its peak, was adopted by more than 80% of active stake. Because Jito-Solana contains and extends Agave functionality, however, a defect in shared consensus, runtime, or execution components could affect both implementations. Adoption of two client configurations did not truly eliminate reliance on common Agave code.

Firedancer, developed by Jump Crypto on an independent C++ codebase, represents Solana's most significant move toward independent client diversity. Firedancer reduces the likelihood that a defect in Agave-derived code would affect the entire validator set. It first entered production through Frankendancer, a transitional architecture combining Firedancer components with portions of Agave, before the standalone implementation went live.

Mithril is another emerging implementation. Developed by Overclock Validator in Go, Mithril produced blocks on the Alpenglowl community test cluster in June 2026. Its design broadens Solana's software and developer base, targeting substantially lighter hardware than conventional validators: a move that could lower barriers to entry and increase node diversity. Mithril remains under development and does not yet represent a material share of mainnet stake.

Products such as Harmonic, Rakurai, and Jito BAM should be distinguished from independent base clients. These systems may introduce alternative schedulers, builder markets, maximal extractable value (MEV) integrations, or proprietary components, but they do not necessarily provide equivalent protection against shared consensus or execution-layer defects.



A Q3 2026 analysis classified approximately 63% of active stake on Jito, 26% on Harmonic, and less than 10% each on Firedancer, Rakurai, Agave, and integrated variations⁹. Although most configurations remained connected to Agave code at the base-implementation level, the emergence of Firedancer and other implementations represents impactful progress in software resilience.

The chart above shows Solana validator software distribution by active stake. As the oldest MEV-compatible client, Jito-related configurations retain the largest share, while emerging implementations are diversifying the software landscape.

Anza’s proposed Constellation architecture for Multiple Concurrent Proposers (MCPs) could provide further improvement by replacing the single leader per slot with multiple proposers, potentially reducing censorship and single-leader risks. Overall, Solana’s software stack is becoming more diverse, though contribution to the network’s resiliency varies materially by implementation.

Section II:

Economic security and control

Are network participants incentivized and governed in ways that promote secure, reliable, and decentralized network operation?

Solana's security model relies more heavily on rewarding productive behavior than directly penalizing harmful behavior: staking rewards compensate validators and delegators, while missed rewards, redelegation, and reputational consequences provide indirect penalties for poor performance. Economic security determines whether participants are motivated to act in the network's interest, while governance determines who can influence its future direction.

Validator revenue and profitability

Approximately 440 million SOL is staked as of Q3 2026, providing a substantial foundation for economic security¹⁰. Validators earn revenue from three primary sources:

- » **Inflation issuance:** Newly-issued SOL distributed to validators and delegators according to eligible stake
- » **Transaction fees:** Base and priority fees paid for transactions included in a leader's blocks
- » **MEV:** Tips paid by searchers or other transaction originators for favorable inclusion or ordering

Inflation represents the majority of validator and delegator rewards, although its relative contribution varies with network activity and each validator's fee and MEV revenue. Base transaction fees are fixed at 0.000005 SOL, while priority fees are dynamic and selected by the transaction submitter. MEV revenue has historically been associated with Jito bundles and the Jito Block Engine, although newer scheduling and block-building systems are creating additional approaches to transaction ordering.

Based on current inflation schedules and average operating costs, estimates put break-even at approximately 200,000 SOL in delegated stake, assuming a 5% validator commission on inflation and MEV rewards. Proposed reductions to the inflation schedule would likely increase this threshold by reducing staking rewards. Actual break-even levels vary with operating expenses, reward rates, fee and MEV generation, and SOL price. Since off-chain expenses are paid in fiat, but revenue is earned in SOL, sustained price decreases could leave smaller validators insolvent and threaten the diversity of the network.

Stake and revenue can create a reinforcing feedback loop. Greater stake provides more leader slots, transaction-fee opportunities, MEV revenue, and commission income. Stronger revenue performance may then attract additional delegation, while poor performance can result in missed revenue and lost stake. Questions surrounding the network's economic health persist, particularly whether aggregate rewards are adequate and if independent operators across the size spectrum can remain financially viable.

Incentives and penalty mechanisms

Solana does not currently rely on conventional protocol-level slashing that automatically removes principal stake for misconduct. Instead, its indirect deterrence mechanisms fall into three categories:

- » **Admission costs:** Hardware, networking, vote fees, and, eventually, the Validator Admission Ticket, which is intended to mitigate sybil identities
- » **Opportunity penalties:** Missed rewards caused by downtime, skipped slots, delayed votes, or poor performance
- » **Market penalties:** Delegation withdrawal, exclusion from stake pools or delegation programs, and reputational damage

Under pre-Alpenglow consensus, validators also pay approximately 1 SOL per day for on-chain vote transactions. These expenses discourage large numbers of low-stake voting identities, although they do not punish misconduct after entry.

Validators with strong uptime, low skip rates, and consistent performance earn more rewards and are better positioned to attract delegation. Stake pools and delegation programs can reduce or remove stake when validators fail to meet performance, commission, software, infrastructure, or decentralization criteria.

The Solana Foundation Delegation Program illustrates this model. As of Q3 2026, it delegated to approximately 51% of validators and accounted for approximately 4% of active stake¹¹. The program uses delegation eligibility and stake matching to encourage strong operational performance, timely upgrades, and infrastructure diversity.


These mechanisms address operational underperformance more effectively than deliberate misconduct. Missed rewards and redelegation may discourage poor uptime, but they may provide weaker deterrence against censorship or coordinated consensus attacks. Their effectiveness also depends on how quickly delegators and institutional allocators respond to performance failures or malicious behavior. [When a validator increases its commission immediately before the epoch ends in order to capture a larger share of rewards \(known as “commission rugging”\), mitigation requires delegators to notice and withdraw their stake.](#) Without a slashing mechanism, harmful validator behavior (like double signing) could theoretically be addressed through a coordinated software upgrade that removes the offending validator from the network. Ultimately, Solana’s enforcement model relies largely on stake mobility and delegator responses to validator behavior. This dynamic is also relevant in the context of liquid staking, which serves as an important mediator between delegations and validators.

Liquid staking

Approximately 68% of circulating SOL is staked, with liquid staking representing around 12% of all staked SOL¹². Liquid-staking tokens enable holders to retain a liquid, yield-bearing asset, while the underlying SOL remains delegated through stake-pool mechanisms.

Adoption of liquid staking protocols varies considerably across blockchains due to differences in staking design. On Ethereum, liquid staking accounts for roughly 50% of total staked ETH, compared to only 9% of staked AVAX on Avalanche¹³. From a network-health perspective, liquid-staking providers influence which validators receive stake, how long it remains delegated, and which performance, infrastructure, commission, or compliance standards operators must meet. Allocation methods may include performance scoring, auctions, decentralization criteria, and validator bonds.

The decentralization effect depends on both sides of the allocation process; how broadly the underlying stake is distributed across validators; and how concentrated authority over that distribution is among liquid staking token (LST) issuers, pool managers, scoring engines, or governance bodies.




LSTs introduce smart-contract and market risks beyond native delegation. At the same time, automated allocation may make stake more responsive to performance and infrastructure concentration than dormant native delegation. Their effect is neither inherently centralizing nor decentralizing, as LSTs can distribute underlying stake across a wider validator set while concentrating authority over that distribution among issuers, pool managers, scoring systems, or governance bodies. Pools also facilitate initiatives geared towards network health, such as the JitoBAM Subsidy Program for client diversity, the Marinade Stake Auction Marketplace for profit sharing, and DoubleZero Phase II for decentralization.

Block building and MEV economics

Solana's transaction-ordering landscape has evolved from default first-in, first out (FIFO) scheduling toward auction-based block construction and specialized schedulers. Under the Jito model, searchers submit bundles to a block engine, which simulates the transactions before forwarding selected bundles to the leader for execution and tip revenue.

Competition has since expanded to include Jito BAM, Harmonic, Rakurai, Firedancer-related implementations, and other scheduling systems. These approaches differ in builder selection, block auctions, proprietary scheduling, and transaction-routing models, but they broadly compete around two objectives: maximizing validator revenue and using increasing block capacity efficiently. Between revenue and slot time, small tradeoffs may be made to improve one objective over another. For example, Harmonic introduces a 50ms TPU delay required for its ordering that captures significantly more priority fees than other clients.

This competition may strengthen execution efficiency and validator economics, but it also introduces questions surrounding private order flow, transaction inclusion, transparency, and the concentration of block-building infrastructure.



The chart above illustrates median block revenue per 10 million compute units by client or scheduler configuration within a seven-day sample period. Specialized implementations like Harmonic generally earn greater block revenue than configurations that are not optimized for block rewards, such as Agave. This shows that software optimization remains a critical component of validator economics.

Governance

Historically, validators exercised governance authority associated with delegated stake. Delegators who disagreed with a validator's position generally had to redelegate to express that preference.

This model created potential conflicts when validators voted on proposals affecting their revenue or competitive position. One Solana Improvement Document (SIMD), SIMD-0228, proposed dynamic inflation and would have reduced validator rewards; it attracted unusually high validator participation before failing. SIMD-0096, which directed 100% of priority fees to validators, passed despite debate about validator and delegator alignment.

Solana Governance Proposals (SGPs) were introduced in July 2026 to formalize direct stake-weighted participation in major governance decisions. Delegators may now submit votes that override the default position taken by their validator, reducing reliance on redelegation as the primary means of expressing disagreement.

Validators with at least 100K staked SOL can propose their own SGP. This differs from the traditional SIMD structure in which only validator identities can vote directly. SGPs determine whether the network should pursue a major policy or protocol direction, while SIMDs specify how technical changes would be implemented. Direct voting may improve transparency and delegator participation, though validators and large institutional stakeholders may continue to exercise substantial influence if most delegators remain passive.

Economic reform

Several proposals are especially relevant to validator economics: SIMD-0550, which would accelerate disinflation; SIMD-0123, which addresses block-revenue distribution; and SIMD-0357, which specifies the Alpenglöw Validator Admission Ticket (VAT)¹⁴.

Inflation and security expenditure (SIMD-0550)

SIMD-0550, voted on and approved under SGP-0002, proposes doubling annual disinflation from 15% to 30%, accelerating Solana's movement toward its 1.5% long-term inflation floor. Modeling estimates that the proposal would reduce cumulative issuance by approximately 22.3 million SOL over six years and reach the inflation floor around 2029, roughly three years earlier than the current schedule.

Lower issuance would reduce dilution and nominal staking yield while decreasing revenue for operators dependent on inflation rewards. Larger validators with more leader slots, along with stronger fee and MEV revenue, may be better positioned to absorb the reduction than smaller operators. SIMD-0550 and the resource-based burn proposal SIMD-0553 remain under discussion.

Block-revenue distribution (SIMD-0123)

SIMD-0123 establishes a native mechanism for distributing block-derived revenue between validators and delegators. Under the existing system, certain transaction-fee revenue is not distributed through the same in-protocol mechanism as inflation rewards. Native distribution could make revenue sharing more transparent and improve alignment between validators and delegators.

The trade-off is increased competitive pressure on validator margins. Larger operators may be better positioned to offer lower commissions or higher revenue-sharing rates, while smaller operators may struggle to match those terms. The proposal could improve delegator alignment, but it may also contribute to a “race to zero” in validator commissions, potentially undermining the long-term financial viability of independent operators.

Validator admission ticket (SIMD-0357)

SIMD-0357 specifies the Validator Admission Ticket associated with Alpenglow’s transition away from on-chain vote transactions. As voting activity becomes more simplified, VAT preserves an economic participation cost by charging each voting validator 1.6 SOL per epoch.

Current voting expenses are estimated at approximately 2.14 SOL per epoch, so the VAT could reduce aggregate validator costs by around 25%. While the mechanism may improve validator sustainability overall via cost-cutting, the flat-fee structure creates a higher relative cost for smaller validators, potentially accelerating the exit of marginal operators and further concentrating the validator set.

Section III:

Performance and usability

Can Solana deliver fast, reliable, and affordable execution under normal and stressed conditions?

As discussed in the previous sections, reliance and economic security establish whether Solana can continue operating. Performance and usability determine whether continued operation translates into practical value through fast, reliable, and affordable execution from transaction submission to finality.


Network reliability

High theoretical throughput offers limited value if congestion or outages prevent a network from reaching consensus. During Solana's earlier years, bot-driven demand caused several network-wide disruptions. In one incident, exceptionally high request volume exhausted validator memory, disrupted voting, and required a network restart¹⁵. The network's performance since 2024 has been materially stronger, with no reported mainnet outage across two consecutive years. Fee-market changes, software improvements, and greater operational maturity have contributed to this progression.


Vote success and vote latency provide additional measures of validator synchronization. [Vote latency measures the difference between the slot being voted on and the slot in which the vote lands on-chain](#). During 2026, vote success averaged approximately 99.7%, while vote latency averaged 1.08 slots, indicating that validators generally remained closely synchronized¹⁶.

The figure below places this improvement in the context of growing network activity. Vote success was comparatively volatile during Solana’s earlier operating history, periodically falling below 80% and, in isolated instances, substantially lower. Beginning in late 2023, however, the rate stabilized near 100%.

Over the same period, non-vote transaction volume increased materially and reached several new highs during 2025 and 2026. These two series alone do not establish causation, but their combined direction indicates that greater user activity has not been accompanied by a sustained deterioration in consensus participation. Solana has instead maintained consistently high vote success while processing a substantially larger transaction load.



A complementary measure is skip-rate, which captures whether scheduled leaders successfully produce acceptable blocks. A skipped slot reduces effective block-production capacity and may delay transaction inclusion. Because validators are assigned leader slots in short, consecutive windows, many skipped slots are isolated to the scheduled validator rather than symptoms of a network-wide disruption.



As shown above, the ten-epoch, stake-weighted average skip rate remained below 0.15% throughout the observed period. The rate fluctuated between approximately 0.02% and 0.12%—including a temporary increase around epoch 1050—before declining to approximately 0.07% by epoch 1053. This variation indicates that leader performance is not perfectly uniform across epochs, but the consistently low absolute levels imply that more than 99.5% of stake-weighted scheduled production succeeded, even at the highest observed point.

At the latest point shown, the corresponding block production rate exceeded 99.9%. Together, vote success and skip rate measure complementary aspects of network reliability: validator participation in consensus and leader execution performance. Near-100% vote success and sub-0.5% stake-weighted skip rates indicate that Solana maintained reliable consensus and block production despite increased non-vote transaction activity, consistent with broader improvements in software, infrastructure, and validator operations.

The transaction lifecycle: from submission to finality

From a user perspective, performance is measured according to the time between sending a transaction and receiving confirmation of its completion. The lifecycle includes submission, forwarding to the upcoming leader, execution, confirmation by validators, and final settlement.

Short slots, stake-weighted leader scheduling, transaction forwarding, and parallel execution all reduce the time between transaction submission and execution. Solana historically targeted approximately 400-millisecond slots, but under SIMD-0525, the network reduces that target in staged increments: to 350, 300, 250, and, ultimately, 200 milliseconds. As of October 2026, the network is targeting 200-millisecond slots, with skip rates remaining low and stable. However, vote latency has increased in some regions¹⁷.

Reaching the 200-millisecond target reduces confirmation latency and narrows the period during which a single leader controls block production. However, shorter slots place greater execution and propagation demands on validators and do not guarantee twice the throughput. Results to date support further reductions but warrant a cautious approach.

Solana currently offers optimistic confirmation in sub-second timeframes, while finality generally occurs within approximately 13 seconds. Settlement speed is particularly important for financial applications because counterparties must know when a transaction can be treated as final.

Where traditional payment systems may provide rapid user feedback while relying on separate clearing and settlement processes, Solana provides execution, confirmation, and settlement through a continuously-operating ledger.

System	User confirmation	Settlement finality	Operating availability
Card payments	Seconds	1-2+ days	Business days only, jurisdiction dependent
Wholesale (RTGS)*	Seconds to minutes	Real-time	24/7, jurisdiction dependent
Solana today	Sub-second	~13 seconds	24/7, global
Solana with Alpenglow (target)	Sub-second	~150 milliseconds	24/7, global


* Real time gross settlement

If Alpenglow achieves its intended objectives, it would materially reduce the gap between initial confirmation and final settlement, moving Solana closer to providing near-instant feedback and finality.

Efficiency and congestion management

Solana’s approach to scaling differs from other blockchain architectures that rely heavily on separate execution layers or external scaling systems. The network instead expands the amount of computation processed directly on the base layer while improving how validators schedule and execute transactions. Network capacity is measured through both transaction counts and Compute Units (CUs), which represent the computational resources consumed during execution. CU limits constrain computational demand at both the transaction and block level, helping to prevent excessive resource consumption and limit block capacity.

The figure below provides an initial view of utilization following a network-level expansion to 100 million CUs per block, implemented in late July. Across a sample seven-day period post upgrade, requested CU utilization remained relatively stable at approximately 23% to 26%, while total block CU utilization ranged from roughly 30% to 45%. Block utilization peaked near 45% on August 4, 2026 before declining to approximately 37% on August 5, 2026. Although utilization varied daily, the network remained below half of available block capacity throughout this period, indicating substantial headroom immediately following the upgrade.



Additional capacity was introduced before the former limit became a persistent, network-wide constraint. During the year preceding the introduction of 100 million CU blocks, approximately 11.2% of blocks produced consumed at least 56 million CUs of the 60 million CU limit¹⁸. This increase provided room for additional growth while reducing the likelihood that temporary demand spikes would routinely exhaust available blockspace.

Unused capacity can help limit fee escalation, delayed inclusion, and transaction failures during periods of elevated demand. However, low average block utilization does not eliminate localized congestion. Transactions competing to modify the same accounts may still experience contention even when aggregate utilization remains moderate. Network-wide capacity and application-specific congestion should therefore be evaluated separately.


As CU limits increase, transaction scheduling and block construction become more important to execution quality. Competing schedulers and block-building systems may improve compute utilization and reduce contention, allowing validators to use available capacity more efficiently. Their economic effects are discussed in the previous section.

From a performance perspective, the relevant test is whether these systems maintain utilization, success rates, latency, and stable fees as demand increases.

Economic accessibility

Economic accessibility depends on both the upfront cost of creating on-chain state and the recurring cost of executing transactions. Solana combines a low base fee with dynamic priority pricing, allowing users to pay an additional amount when faster inclusion is particularly valuable. A typical transfer has a base fee of approximately 0.000005 SOL, while the total cost may vary with transaction complexity, network demand, and the priority fee selected by the user.

Applications and users may also fund rent-exempt deposits when creating accounts that store data on-chain. Under Agave 4.2, SIMD-0437 reduces Solana's rent parameter by 90%, lowering the deposit for a standard SPL token account from approximately \$0.159 to \$0.0159. While this deposit can be recovered when the account is closed, reducing the upfront requirement makes it more practical for applications to create accounts or sponsor user onboarding at scale.




Priority fees represent the largest component of aggregate transaction fee expenditure. As shown above, non-vote priority fees generally accounted for approximately 70% to 80% of total fees during the rolling one-year period. Non-vote base fees represented a much smaller share, while vote transactions accounted for most of the remainder. This composition indicates that users frequently pay for transaction prioritization, but fee composition alone does not establish whether the typical transaction is expensive.

Average fee data provides that additional context. Over the period shown above, the average priority fee generally remained around or below one cent per transaction relative to USD, though temporary spikes occasionally moved the average above that level.

Average SOL-denominated priority fees also declined from levels observed in late 2025 and early 2026 before increasing modestly later that year. The difference between the SOL and USD series reflects both changes in the native fee and market value of SOL.

Taken together, these figures show that priority pricing is widely used but remains inexpensive for the typical transaction. Users can pay more when timing is valuable without imposing the same fee on all network participants, while low base fees and reduced account deposits preserve accessibility for less time-sensitive activity and account-heavy applications.



Conclusion

Solana has made significant progress in terms of network health across resilience, economic security, and performance indicators. Solana’s resilience has improved through greater software diversity, infrastructure investment, and a prolonged period without a major outage. Economic security is supported by substantial staking participation and incentives that reward validator performance, while low latency, expanded capacity, and inexpensive transactions support practical use at scale.

Several areas surrounding decentralization and resilience do warrant continued observation. Stake concentrates among a relatively small number of validator identities and operators, while providers, jurisdictions, and shared software dependencies can create correlated failure risks. Economic reforms may also affect validator profitability and the long-term composition of the validator set, particularly where flat costs or lower issuance have a disproportionate effect on smaller operators.

Overall, the broader trajectory remains positive. Firedancer and emerging implementations (including Mithril) strengthen the outlook for software diversity, while Alpenglow, shorter slot times, expanded compute capacity, and infrastructure investments target further gains in performance and resilience. The next phase of Solana’s development will test whether these advances will come at the expense of decentralization or network security.

Works cited

- 1 RWA.xyz. "Solana." RWA.xyz, <https://app.rwa.xyz/networks/solana>.
- 2 "Solana Q1 2026: Stablecoin Transfers Hit \$2.1 Trillion, and What's Changing Underneath." Raiku, 20 May 2026, <https://raiku.com/blog/solana-q1-2026-stablecoin-transfers-hit-2-1-trillion-and-whats-changing-underneath>.
- 3 "Solana Analytics." Blockworks Research, <https://blockworks.com/analytics/solana>.
- 4 "How Much Does a Solana Validator Cost in 2026? (\$60K Breakdown)." The Good Shell, 16 Apr. 2026, <https://thegoodshell.com/solana-validator-cost/>.
- 5 "Solana Network Decentralization Analysis." Sonda, <https://sonda.network/?cluster=mainnet-beta>.
- 6 "Solana Network Decentralization Analysis." Sonda, <https://sonda.network/?cluster=mainnet-beta>.
- 7 "Data Centers." Validators.app, https://www.validators.app/data-centers?locale=en&network=mainnet&sort_by=asn.
- 8 "TeraSwitch Routing Failure Brought Solana Within 4.5 Points of a Finality Halt." Solana Compass, <https://solanacompass.com/news/teraswitch-routing-failure-brought-solana-within-45-points-of-a-finality-halt>.
- 9 "Solana Network Decentralization Analysis." Sonda, <https://sonda.network/?cluster=mainnet-beta>.
- 10 "Solana Supply, Staking and Validators." Blockworks Research, <https://blockworks.com/analytics/solana/solana-supply-staking-and-validators>.
- 11 Solana Foundation. "Delegation Dashboard." Solana Foundation, <https://solana.org/delegation-dashboard>.
- 12 "Solana Analytics." Blockworks Research, <https://blockworks.com/analytics/solana>.
- 13 "Chain Comparison." Blockworks Research, <https://blockworks.com/analytics/chain-comparison>.
- 14 Solana Foundation. "Solana Improvement Documents." GitHub, <https://github.com/solana-foundation/solana-improvement-documents>.

- 15 Solana Foundation. "9-14 Network Outage Initial Overview." Solana, 20 Sept. 2021, <https://solana.com/news/9-14-network-outage-initial-overview>.
- 16 Syndica. "Deep Dive: Solana Onchain Activity." Syndica, 27 Apr. 2026, <https://blog.syndica.io/deep-dive-solana-onchain-activity/>.
- 17 Natale, Umberto. "Slot Time Reduction Effects." Solana, 28 Sept. 2026, <https://solana.com/news/slot-time-reduction-effects>.
- 18 Solana Foundation. "100M CU Blocks." Solana, <https://solana.com/upgrades/100m-cu-blocks>.

Digital assets are speculative and highly volatile, can become illiquid at any time, and are only for those investors willing to risk losing some or all of their investment and who have the experience and ability to evaluate the risks and merits of an investment.

This information is for informational purposes only and is not intended to provide investment or any other advice and should not be construed as an offer to sell, a solicitation of an offer to buy, or a recommendation for any security or other assets. The opinions provided are those of the authors and not necessarily those of the Fidelity Center for Applied Technology (FCAT) or its affiliates. FCAT does not assume any duty to update any of the information. FCAT and any other third parties are independent entities and not affiliated. Mentioning them does not suggest a recommendation or endorsement by FCAT. Third-party trademarks and service marks are the property of their respective owners. All other trademarks and service marks are the property of FMR LLC or its affiliated companies.

© 2026 FMR LLC. All Rights Reserved. 1284544.1.0



An FCAT Publication