Twinstake

THOUGHT LEADERSHIP

Solana's Evolution to Institutional Grade Trading Infrastructure

By Matan Hamburger

20 July 2026

Executive Summary

Solana has undergone a profound infrastructure transformation over the past five years, evolving from a high-potential but operationally immature network into one of the most competitive venues for on-chain high-frequency trading. This article traces that journey chronologically: from the early protocol-focused years in which validator performance was largely indistinguishable, through the memecoin-driven demand shock of 2024 that exposed latency as the critical bottleneck, to the emergence of a sophisticated execution layer built around MEV extraction, transaction routing, and shred propagation.

It then examines the infrastructure upgrades spanning client diversity, XDP kernel-bypass networking, and the Alpenglow consensus overhaul that are positioning Solana to close the gap with centralised trading venues, and concludes with what this means for institutional participants.

Institutional implications

For financial institutions, the practical implication is that Solana infrastructure is no longer a purely technical consideration. Validator selection, transaction routing, staking strategy, and execution infrastructure are increasingly connected decisions that can affect performance, risk, and economics.

Institutions should consider the following implications:

- Execution quality depends on infrastructure choices. RPC providers, routing paths, validator relationships, and geographic location can influence transaction landing rates, latency, and slippage.
- Validator due diligence needs to go beyond uptime and commission. Institutions should assess client stack, MEV policy, geographic resilience, failover processes, and connectivity to key Solana infrastructure.
- MEV and priority fees require governance oversight. Institutions should understand how validators capture, distribute, and manage MEV-related revenue, and whether that approach aligns with their operational standards and ecosystem relationships.
- Staking and trading are becoming more interconnected. Mechanisms such as stake-weighted transaction prioritisation mean that validator relationships may increasingly influence execution outcomes.
- Solana is becoming an increasingly significant venue for HFT and latency-sensitive trading. As finality compresses and execution infrastructure improves, Solana becomes increasingly practical for strategies that depend on speed, predictable settlement, and high transaction throughput, including market making, arbitrage, liquidations, and on-chain order book trading.

In short, Solana's evolution means institutional participation increasingly depends on prudently and deliberately choosing infrastructure partners.

As Solana's infrastructure becomes more specialised, trusted institutional validators will play an increasingly important role. Twinstake provides institutional, non-custodial validator infrastructure and technical expertise to support institutions' own operational due diligence. References in this paper to transaction routing, MEV infrastructure and trading strategies describe developments across the wider Solana ecosystem and should not be understood as implying that Twinstake executes or arranges trades, routes client orders, operates a trading venue or provides investment advice.

Early Solana Infrastructure

In Solana's former years, the network's engineering effort was focused almost entirely at the core protocol layer. The primary objective was stability and throughput: building a network capable of processing thousands of transactions per second without collapsing under load. Slot times and throughput were the principal performance variable, and the team at Solana worked to compress them in the context of a globally distributed validator set.

During this period, validator performance was largely undifferentiated. Revenue came almost entirely from inflationary staking rewards, which were distributed proportionally to stake. There was little meaningful income from transaction fees or MEV (Maximal Extractable Value, the profit validators can earn by controlling the ordering of transactions within a block). The main edge a validator could hold was hardware quality and uptime. The network bottleneck sat firmly at the protocol layer, and until that was addressed, there was limited commercial incentive to optimise further up the stack.

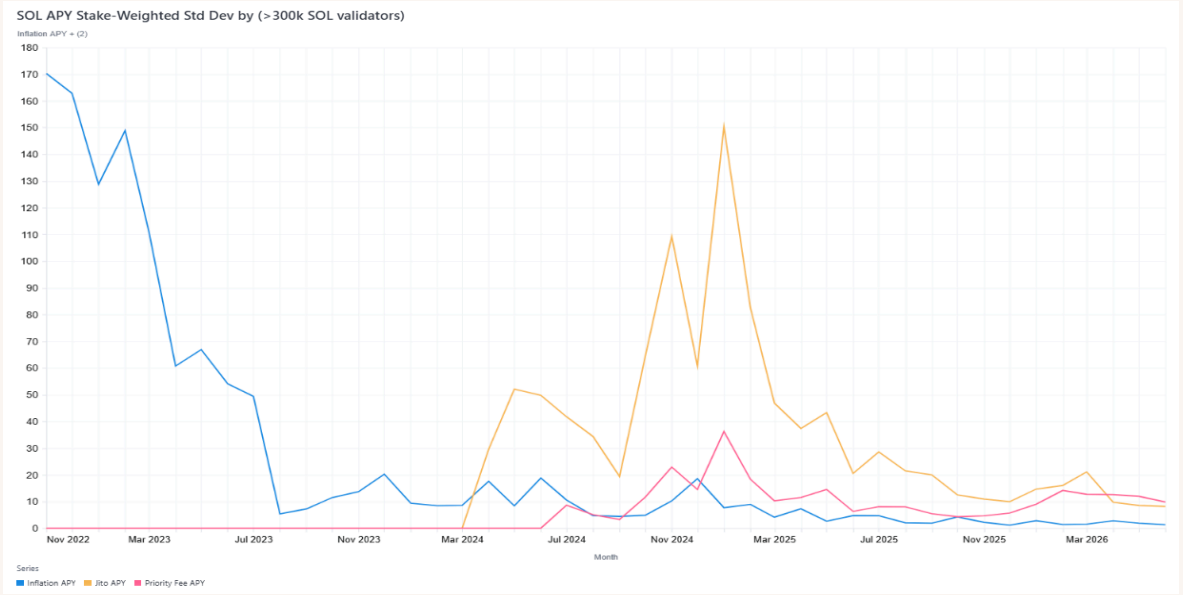

The graph above shows stake-weighted standard deviation of validator APR performance, a proxy for how much validators can differentiate on returns. Outperformance has shifted from inflationary rewards to execution-layer rewards (MEV and priority fees), which now account for nearly all variance in the set.

The Memecoin Demand Shock and the Rise of MEV

The landscape shifted dramatically when pump.fun launched in January 2024. Pump.fun is a token launchpad that allows anyone to create and immediately trade a new Solana token, triggering a memecoin frenzy that drove a step-change in on-chain activity. Within months it was responsible for over 70% of all new tokens minted on Solana and over half of all DEX transactions on the network. By March 2026, it had surpassed $1 billion in cumulative revenue, generating $124.7 million in Q1 2026 alone, representing 36% of all Solana application revenue for the quarter.

This explosion of retail activity had immediate consequences for infrastructure. Transaction volumes surged, fee revenue became material, and MEV opportunities multiplied. Latency, which had previously been a secondary concern, became the defining competitive variable. Users experienced failed transactions, poor landing rates, and adverse execution as the network struggled to handle demand efficiently. Existing inefficiencies in transaction routing and block propagation, which had been tolerable at lower volumes, were now actively harmful.

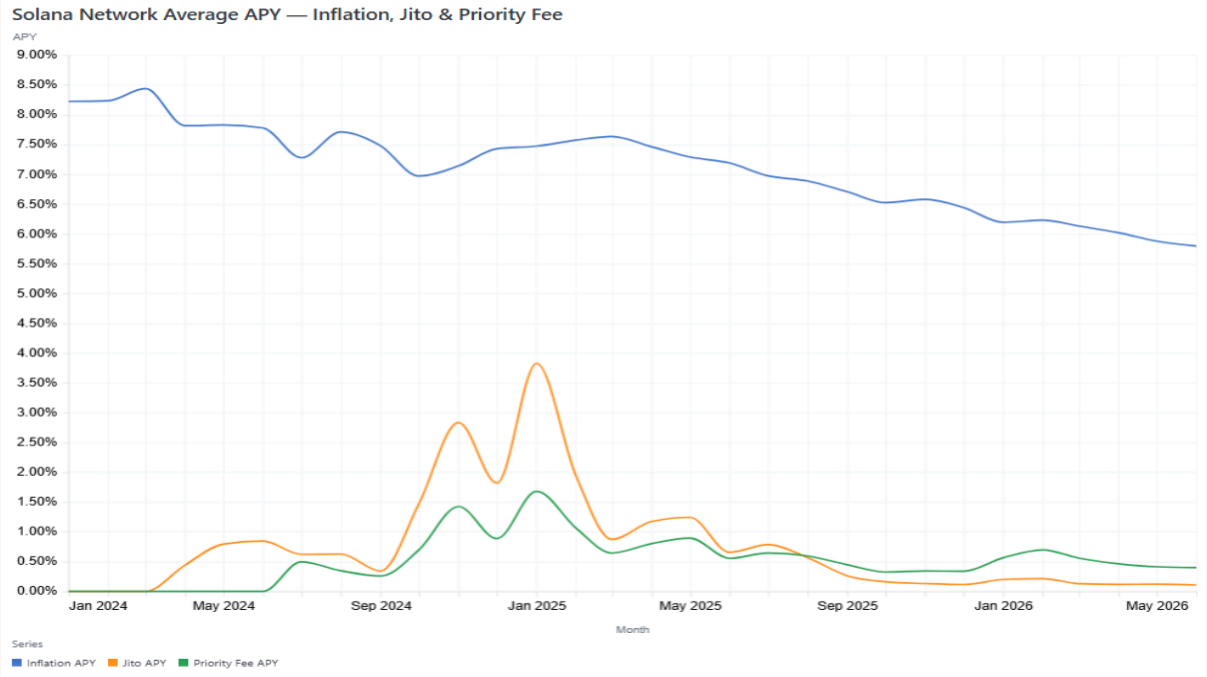

Pump.fun's emergence in late 2024 drove explosive growth in user fees on Solana, sharply increasing execution rewards (MEV and priority fees). Staker revenue from these sources has stayed elevated since.
(Data sourced from Twinstake's internal data)

Funded by this wave of on-chain activity, the broader node infrastructure expanded rapidly. RPC providers doubled down on Solana capacity. Validators invested in improved hardware and benefited from continued core protocol improvements, pushing block times below 400ms. But the more significant shift was in where validator outperformance now lived: not at the protocol layer, but at the execution layer.

This represented a structural shift in validator economics. In the inflation-dominated era, the path to outperformance was straightforward: maximise uptime, maintain hardware, and collect your proportional share of newly minted SOL. MEV and priority fees were negligible noise on top of that base. From 2024 onwards, that relationship inverted. Fee revenue and MEV tips began to rival and, for the most competitive validators, exceed inflationary rewards as the primary source of income. A validator's edge was no longer determined by whether it stayed online, but by how efficiently it could extract value from each leader slot it was assigned. This changed the incentive calculus entirely. Validators began optimising aggressively for transaction landing rates, bundle inclusion, and block packing efficiency, and that optimisation drive pushed them toward deliberate choices about client software, colocation strategy, and timing behaviour that had not mattered at all in the inflationary era.

During this period, Agave remained the core validator client underpinning the Solana network. Developed in Rust by Anza, the engineering firm spun out of Solana Labs, Agave was still the only full validator client, concentrating the network's software risk in a single codebase. This was a pragmatic choice for a protocol that was still maturing, but it created a structural vulnerability that the ecosystem would need to address. Jito emerged within this context as the dominant force in MEV and transaction routing. It achieved this by creating a modified version of Agave that integrates a block engine, allowing searchers, or sophisticated traders seeking profitable transaction-ordering opportunities, to submit bundles of transactions with attached SOL tips. Validators running Jito receive these tips in exchange for including the bundles, creating a structured marketplace for block space. As the revenue advantages became clear, Jito's share of the validator set grew rapidly.

With fee revenue now material, a less visible but economically significant behaviour emerged: timing games. Because the Solana leader schedule is known in advance, a validator assigned an upcoming slot has a narrow window in which it can delay block production to observe more incoming transactions, selectively including those with the highest tips or the most valuable MEV bundles. Intentional delay in block production may increase per-slot economics in certain circumstances, but can degrade network performance and may raise market-integrity, contractual, governance and regulatory concerns. The discussion in this paper is descriptive and should not be understood as endorsing manipulative, abusive or network-degrading conduct. The tension between individual validator revenue maximisation and network-level performance became one of the defining infrastructure debates of the period, and it has directly shaped the design priorities of Alpenglow, which embeds penalty mechanisms to make late block production economically costly.

Building the Execution Layer: Client Diversity, Shreds, and SWQoS

As transaction volume and MEV revenue scaled, attention shifted from the core protocol to the execution layer: the set of infrastructure choices that determine how quickly a validator can process transactions, produce blocks, and route order flow. The dominant question became: given the same protocol, how much performance can you extract from your stack?

This is the context in which the Solana client landscape diversified from a single-client network into a competitive ecosystem of specialised implementations, each reflecting a different philosophy about where the remaining performance gains lie.

Agave is the foundation. Developed by Anza and written in Rust, it is the direct descendant of the original Solana Labs validator and remains the most widely deployed client on the network. It handles the full validator lifecycle, covering

transaction ingestion, block production, consensus voting, and ledger replay, and serves as the reference implementation against which all other clients are measured. Most validators run some derivative of Agave, either in its base form or augmented with one of the specialised layers described below.

Jito-agave

Jito is an Agave fork that adds a block engine and bundle auction layer on top of the standard client. Searchers submit bundles of transactions with attached SOL tips through Jito's off-chain relay, and the Jito-Solana client gives the validator a mechanism to accept these bundles at the top of each block. This creates a structured MEV marketplace: rather than competing purely on tip size within the standard mempool, searchers bid explicitly for priority inclusion, subject to block-building, validator and network conditions in a specific block position. The economic advantages for validators are substantial. Jito tips routinely represent a significant component of per-slot revenue for high-stake operators, which is why the majority of stake now run Jito-Solana rather than vanilla Agave.

Harmonic

Harmonic is an open block-building marketplace for Solana. Rather than functioning as a single scheduler in the way Jito-agave does, it aggregates candidate blocks from multiple independent builders, including Jito, Temporal, Jito BAM, and Paladin, then lets each validator select the block that best matches its own preferences, whether that means maximising revenue or minimising toxic MEV. This shifts validators from passive recipients of a single builder's output into active participants who can choose across a competitive market each slot.

Frankendancer/ Firedancer

Frankendancer and Firedancer mark the next stage in this execution-layer race. Whilst Jito and Harmonic improve how validators capture and order transaction flow, Jump Crypto's clients focus on a more fundamental constraint: how quickly a validator can receive network traffic, process transactions, build a block, and propagate it to the rest of the network.

Frankendancer is the intermediate step. It keeps Agave's existing consensus system, but introduces Firedancer's faster networking and block production components. This gives validators a way to improve performance without fully replacing the infrastructure they already rely on. The practical benefit is lower overhead, faster packet handling, and more efficient block production during periods of heavy demand.

Firedancer goes further. It is a full rebuild of the Solana validator, designed from the ground up for speed and hardware efficiency. Rather than processing traffic through the standard software stack, Firedancer is built to handle packets closer to the network card and distribute validator tasks across dedicated CPU resources. The objective is to reduce latency, increase throughput, and make block propagation more reliable as Solana activity scales. It's worth noting that Firedancer is written entirely in C whereas the other Solana clients are written in Rust and is therefore a complete rewrite of the code base. It therefore further reduces any single point of failure risk within the Solana client ecosystem.

Rakurai

Rakurai is the newest entrant and the least widely deployed client at present. It is a fork of jito-solana that adds a proprietary heuristics-based transaction scheduler designed to maximise block rewards and TPS by efficiently selecting and processing high-value transactions. Because its primary differentiation sits at the scheduling layer, it occupies a different position in the client landscape from networking-focused implementations like Firedancer, representing an alternative route to performant validator operations.

Validator Configuration Design

In practice, for institutional operators like Twinstake, the choice of client is part of a broader multivariate optimization problem with conflicting ecosystem impacts. The core responsibility of validator operators like Twinstake is to protect the client's assets whilst trying to optimize rewards through technical configurations. However, the impact of client selection must also account for its downstream effect on Solana network's performance, Solana user experience, compliance requirements and ethical MEV practices.

Colocation strategy compounds these choices further: a validator co-located with Jito's block engine servers and major RPC providers in Frankfurt or Amsterdam will see better bundle flow and lower relay latency than one running identical software in a less connected region. The result is that validator economics, once determined almost entirely by inflationary performance, are now shaped by a combination of client selection, scheduler configuration, geographic placement, and network relationships that would have been unrecognisable to operators from the inflationary era.

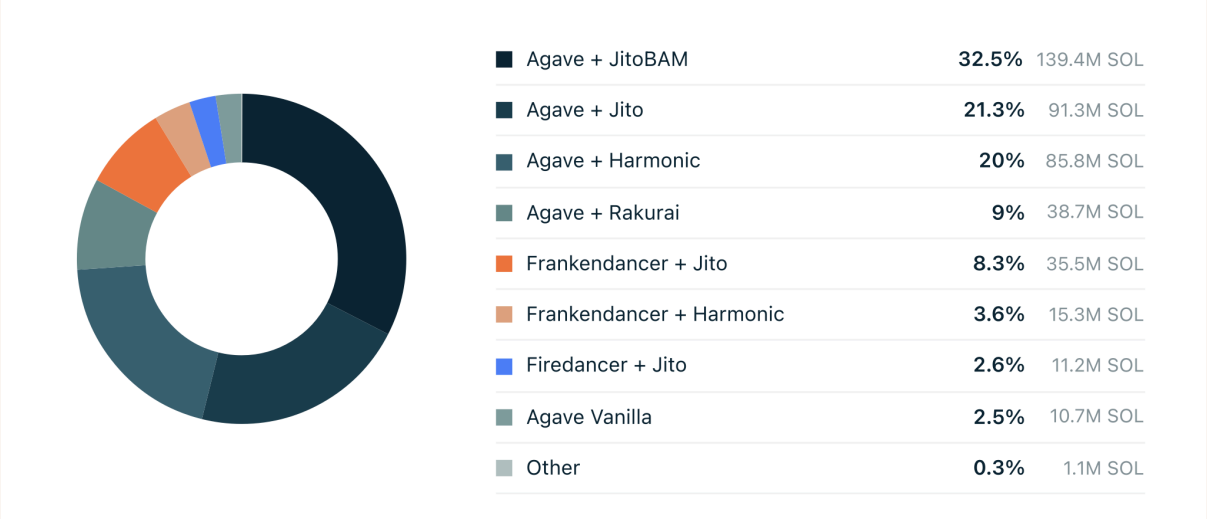

Market share of Solana Clients July 2026. (Data sourced from https://reports.firedancer.io/?tab=validators)

Solana's Geographic Diversification

Geography became another critical dimension for Solana's end users. Because Solana's leader schedule rotates through validators in a deterministic order, a trader's proximity to the current leader directly affects how quickly their transaction can land. Historically, Solana infrastructure concentrated around central European data centres, primarily Frankfurt and Amsterdam, where colocation providers such as Teraswitch have a particularly high density of validators. This centralisation made practical sense: low inter-validator latency reduces missed blocks and improves MEV extraction. But it also created geographic dependency and uneven user experiences for participants in other regions.

DoubleZero is addressing this through a dedicated fibre optic network built specifically for blockchain infrastructure (DoubleZero Network). Rather than routing traffic over the public internet, DoubleZero's network connects validators, RPC nodes, and MEV services via low-latency private links, reducing propagation times for participants regardless of where they are located (DoubleZero network dashboard). Double Zero has reinforced this with a delegation incentive programme, utilizing its foundation backed LST, that encourages validators to move to underrepresented geographies. Progress has been made, though the ecosystem still suffers from elevated latency for

isolated validators in Southeast Asia, South America, and Africa, regions where network infrastructure limits what is achievable regardless of software quality.

Independent capital is beginning to respond to this gap. One of the largest Solana digital asset treasury companies has announced plans to build dedicated high-speed infrastructure targeting market makers and high-frequency traders in Southeast Asia, a signal that the commercial case for low-latency infrastructure outside of Europe is increasingly hard to ignore.

High Frequency Data (Shreds) and Transaction Landing (SWQoS)

Two further infrastructure layers have become battlegrounds for latency-sensitive participants, and have simultaneously driven the geographic centralisation described above while intensifying the need to address it: shred propagation and Stake-Weighted Quality of Service (SWQoS).

Shreds are the small data fragments into which Solana breaks each block for propagation across the network via a protocol called Turbine. As the leader validator builds a block, it encodes transactions into shreds and broadcasts them to the rest of the validator set. The leader sees these shreds before anyone else, and any validator or RPC node with a direct connection to the leader can access them fractions of a second earlier than the rest of the network. For traders and market makers, earlier shred access means earlier visibility into what transactions are being included in the current block, which is directly valuable for strategies that depend on knowing the state of the order book. This has given rise to a dedicated ecosystem of ShredStream providers, validator-RPC colocation arrangements, and low-latency data feeds. Competition has grown so intense that a 1 to 2 millisecond advantage is considered meaningful.

SWQoS governs how the leader validator prioritises incoming transactions. Solana allocates a portion of the leader's incoming bandwidth proportionally to validators based on their stake. Transactions relayed by a high-stake validator node are treated as priority traffic and processed ahead of those arriving through standard RPC channels. In practice, traders often colocate their transaction-sending infrastructure with validator-controlled infrastructure to benefit from Stake-Weighted Quality of Service (SWQoS). Transactions relayed through validator-controlled infrastructure can benefit from the validator's stake-weighted priority, improving transaction landing rates. Achieving these benefits depends on low-latency connectivity and appropriate validator infrastructure configuration. The validator identity keys required to enable SWQoS remain under the validator operator's control and are not made available to third parties. The effect is significant: transactions sent through a well-positioned validator identity node land at materially higher rates than those sent directly. Every millisecond of additional latency in the relay path reduces landing probability, making infrastructure choices directly visible in execution quality.

XDP: The Kernel-Bypass Revolution

One of the most impactful recent infrastructure developments has been the integration of XDP (eXpress Data Path) into the Agave validator client. XDP is a Linux kernel technology that allows network packet processing to bypass much of the normal kernel networking stack. Instead of packets travelling through multiple layers of software before reaching the validator process, XDP runs small programmes directly on the network interface card driver, processing packets with dramatically lower overhead.

For Solana, the primary application is accelerating Turbine, the block propagation protocol. As block compute limits have increased towards 60 million compute units and are being pushed towards 100 million, the number of shreds a validator must send per slot grows correspondingly. Under conventional networking, this can become a bottleneck: highly-staked validators can approach 150,000 outbound packets per second under Turbine's fanout model, a load that strains standard Linux networking. XDP helps relieve this bottleneck by lowering unnecessary kernel networking overhead, including data copies and context switches, thereby reducing block propagation latency by up to 200x in early tests.

Alpenglow: The Consensus Overhaul

The most consequential upgrade on Solana's roadmap is Alpenglow: a complete replacement of the network's consensus layer, proposed by Anza in May 2025 and approved by validator governance in September 2025 with 98.27% of participating stake voting in favour. As of June 2026, it is live on a community test cluster with mainnet activation targeted for Q3 or Q4 2026.

To understand what Alpenglow changes, it is worth understanding what it replaces. Solana's current consensus runs on two mechanisms: Proof of History (PoH), a cryptographic clock that timestamps transactions before they reach consensus so validators do not need to argue over ordering; and TowerBFT, the voting mechanism validators use to agree on the state of the chain. Together they produce block times of roughly 400 milliseconds, but finality (the point at which a transaction is fully confirmed and irreversible) currently takes around 12.8 seconds due to the 32-round confirmation process TowerBFT requires. Additionally, the act of voting is recorded on-chain, meaning approximately 75% of all Solana transactions are validator votes rather than user activity, consuming significant block space.

Illustrative comparison of Solana protocol finality: current consensus and Alpenglow target:

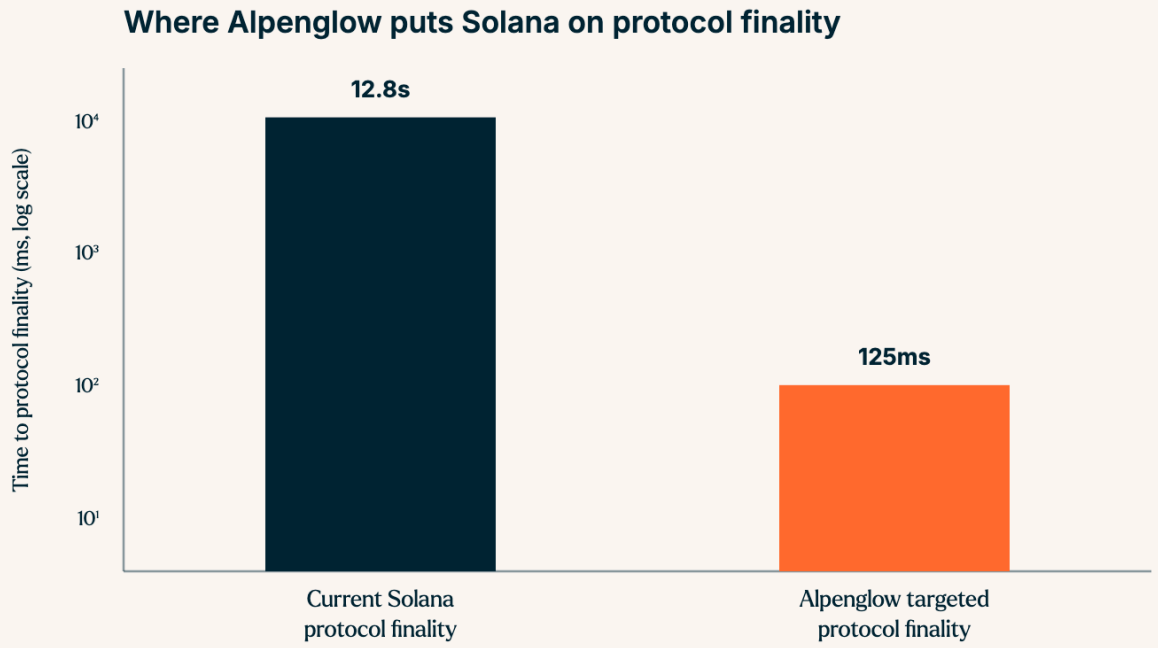

Sources: Solana current finality ~12.8s and Alpenglow target 100–150ms (Anza; SolanaFloor).

The figures shown relate to protocol-level finality and are not directly comparable with payment authorisation times, trading-venue execution latency or legal settlement finality. Alpenglow's target timings are based on current proposals and testing, remain subject to successful mainnet implementation and are not guaranteed.

Votor / Rotor Modify Consensus & Data Propagation across the Network

Alpenglow proposes an overhaul of both consensus and block propagation mechanics with two new components. Votor is a new voting protocol that collapses consensus into one or two rounds. Current projections for Votor, should it be adopted, when 80% or more of validator stake is online, a block can be finalised in a single round at approximately 100 milliseconds. When 60% of stake participates, two rounds are required, targeting 150 milliseconds. Critically, Votor moves voting entirely off-chain: validators exchange BLS-signed vote certificates directly with each other rather than publishing votes as on-chain transactions. Only the aggregated certificate, roughly 1,000 bytes, lands on-chain, replacing the roughly 500 kilobytes of vote data currently recorded per slot. The result is that block space freed from vote transactions becomes available for user activity, and validator operating costs may fall.

Rotor, Alpenglow’s second component, is positioned to replace Turbine as the block propagation mechanism. Where Turbine uses a multi-hop tree structure to fan block data out across the validator set, Rotor uses a one-hop broadcast model combined with erasure coding, a technique that divides block data into redundant fragments allowing recipients to reconstruct the full block even if some fragments are lost. Relay assignments under Rotor are deterministic and stake-weighted rather than random, making propagation more predictable. Simulations show Rotor completing block propagation across the network in as little as 18 milliseconds under typical bandwidth conditions.

The implications for trading infrastructure could be substantial if Alpenglow is implemented on mainnet as currently proposed. Targeted protocol-level finality of approximately 100 to 150 milliseconds could materially improve Solana's suitability for latency-sensitive trading applications. On-chain central limit order books, which have historically struggled to compete with centralised venues on latency, could become meaningfully more viable. Liquidation engines and arbitrage strategies that currently require probabilistic assumptions about finality could operate with significantly greater confidence in transaction confirmation before acting.

The above assessment assumes Alpenglow is activated on mainnet as proposed and performs as intended. These targets are based on published proposals, testing and simulations and are not guaranteed. It is also worthwhile noting that protocol finality is a technical network characteristic and is not equivalent to legal settlement finality; it does not eliminate counterparty, custody, operational, market, protocol or insolvency risk.

What This Means for the Future of On-Chain Trading

Viewed together, the progression from Solana’s early protocol-focused years to today tells a coherent story. A network that once competed primarily on throughput claims has systematically addressed each successive bottleneck: core protocol stability first, then MEV infrastructure and execution layer tooling, then geographic latency, client diversity, and networking efficiency. What is emerging is a stack that begins to approximate the performance characteristics of centralised trading venues while preserving the protocol-level finality characteristics of a public blockchain.

The transition has also increased the reliance on infrastructure providers like Twinstake to deliver enterprise-grade solutions that minimise latency while ensuring network reliability. Twinstake seeks to optimise staking rewards through specialist engineering and careful monitoring of protocol developments, while balancing network performance, resilience and security. As with all staking services, rewards are variable and depend on network and market conditions.

Alpenglow, when it activates on mainnet, will be the most significant single step in that journey. Sub-150 millisecond finality changes the category of application that can be built on Solana: real-time order books, on-chain perpetuals with institutional-grade latency, tokenised instruments that settle as fast as they trade. The expanded block space that follows from moving votes off-chain provides further headroom for user activity. And Firedancer's continued rollout, combined with mandatory XDP adoption across the validator set, means the hardware and networking foundations are being upgraded in parallel with the consensus layer.

The latency race will not end here. Alpenglow's faster finality will intensify competition at every layer below consensus: shred access, SWQoS positioning, and geographic proximity will all become more valuable, not less, as the performance ceiling rises. That competitive dynamic is itself a sign of maturity. Solana is no longer a fast blockchain in search of a use case. It is an evolving trading infrastructure being actively stress-tested by institutional participants and high-frequency strategies with real capital at stake.

For institutions, the practical implication is that infrastructure choices on Solana have become performance decisions, not operational details. Sub-second settlement with deterministic finality seem to help address some of the counterparty and reconciliation concerns that have historically kept institutional flow on centralised venues. The multivariate validator configuration optimization means that a participant's choice of validator relationships directly shapes both their own staking rewards as well as ecosystem-wide execution quality. Selecting a validator impacts the institution's direct staking operations, however it also contributes to network wide transaction routing performance and trading infrastructure. These considerations are converging into a single decision surface, and the institutions best positioned for this next phase will be those that treat validator selection with the same weight they apply to other critical vendors.

About Twinstake

Twinstake is an institutional-grade, non-custodial validator infrastructure provider, designed and built in collaboration with leaders in both traditional finance and the digital assets space. We serve a range of institutions including the largest crypto ETFs, VCs, asset managers and exchanges. Twinstake operates validator infrastructure across multiple proof-of-stake networks and supports institutional staking through secure, resilient and high-performance infrastructure.

Get in touch with our team to hear more about how Twinstake can support your institutional staking.

info@twinstake.io

Disclaimer: Twinstake does not provide staking services to retail customers. This briefing note is not intended as a promotion, offer, invitation or solicitation for the purchase or sale of any investment, nor is it intended to give rise to any other legal relations whatsoever and must not be relied upon for the purposes of any investment decision. It does not constitute financial, legal, or investment advice. If you do not have the relevant professional experience in matters relating to crypto asset investments, you should not consider this briefing note to be directed at you.

References to network performance, rewards, finality and planned protocol upgrades are based on information available as at 5 July 2026, may change and are not guaranteed. Historical reward data is not a reliable indicator of future results. Staking rewards are variable, and staking may involve slashing, downtime, liquidity, market, protocol, operational and other risks. References to blockchain or protocol finality are technical in nature and do not constitute representations regarding legal settlement finality or the elimination of counterparty, custody, insolvency or other risks.

This briefing note and the information in it are not directed at, or intended to be made available to, retail customers. It is directed only at persons who are professional investors (for the purposes of the Alternative Investment Fund Managers Directive (2011/61/EU) (known as 'AIFMD'); professional clients or eligible counterparties for the purposes of the Markets in Financial Instruments Directive (Directive 2004/39/EC) (known as 'MiFID'); if you are in the UK, to "Investment Professionals" or "High Net Worth Companies" as defined in s.19 and s.49 respectively of the Financial Promotions Order, or as otherwise defined under applicable local regulations and at whom this briefing note and the information in it may lawfully be directed in any relevant jurisdiction.

The appearance of any third-party hyperlinks or third-party reference in the briefing note does not constitute an endorsement, guarantee, warranty, or recommendation by Twinstake. Do conduct your own due diligence before deciding to use any third-party services.

Twinstake shall have no liability for any loss or damage that may arise directly or indirectly from the use of or reliance on the information provided herein or for any errors or omissions in the information.