Twinstake

TWINSTAKE RESEARCH

Solana Alpenglow Upgrade: The Impact on Institutional Staking

What changes, why it matters, and what it means for institutions staking SOL

14th September 2026

by Matan Hamburger

Contents

Report Context and Purpose

Over the coming months, Solana is set to undergo a significant consensus upgrade: Alpenglow. Approved by validator governance in September 2025 with overwhelming support, Alpenglow replaces two of Solana's founding technical components, Proof of History and TowerBFT, with a new architecture built for deterministic, near-instant finality.

This report explains what changes, why it matters, and what it means specifically for institutions staking SOL, whether directly or through an institutional provider such as Twinstake. Unlike a routine client update, Alpenglow touches the assumptions institutional stakers rely on most: how quickly a transaction becomes irreversible, how much stake would need to misbehave or fail before the network's safety guarantees are tested, the liquidity considerations around staking, what it costs a validator to participate in consensus and how validator rewards are generated. This report focuses on the staking-relevant dimensions of the upgrade rather than its implications for application developers or traders, though we touch on the latter where it shapes the institutional case for Solana as a venue. We also discuss some related SIMDs that are not included in Alpenglow directly but are related.

The core Solana Improvement Documents (SIMDs) relevant to Alpenglow are:

- **SIMD-0326 (Alpenglow):** The core proposal replacing Proof of History and TowerBFT with two new components, Votor and Rotor, targeting a reduction in transaction finality from approximately 12.8 seconds to roughly 100–150 milliseconds.
- **SIMD-0357:** The Validator Admission Ticket (VAT), which replaces continuous on-chain vote transaction fees for validator operators with a single, predictable per-epoch admission cost. Partially live (admission gating only) since 22 July 2026; the fee itself activates alongside full Alpenglow consensus.

- **SIMD-0387 (BLS Pubkey Management in Vote Account):** The validator-side prerequisite for Alpenglow, requiring every validator to register a new cryptographic key type. Live on mainnet since 8 July 2026.
- **SIMD-0525 (Reduce Slot Times):** A related, independently progressing proposal that shortens Solana's block time in stages from 400ms toward 200ms, with direct implications for liquidity modelling and validator timing-game incentives (discussed in Section 2). The first stage of this proposal has already been implemented on 21st August reducing target slot times to 350ms.

As with any upgrade of this scope, the impact is more nuanced than the headline "100x faster" framing suggests. It changes the cost structure of running a validator, introduces a new, and in one specific respect lower, security threshold in exchange for materially better real-world uptime. We also present a Twinstake view framing the upgrade as a foundational piece of Solana's broader ambition to become settlement infrastructure and an "Internet Capital Markets."

For the convenience of the reader, this report is split into three sections:

- **Introduction:** A holistic view of the upgrade and a brief explainer of each new component.
- **Deep Dive & Institutional Impact Analysis:** A description of how Votor, Rotor and the Validator Admission Ticket work, how block-time reductions interact with validator timing-game incentives, and what this means for reward predictability and settlement risk.
- **Twinstake Opinion Piece**: Twinstake's view on the security trade-offs involved, framed within the context of Solana's ambition to host on-chain financial infrastructure.

Section 1: Introduction

1.1 Why Solana is replacing its consensus engine

Solana's speed advantage has always rested on two components working together: Proof of History (PoH), a verifiable cryptographic clock, and TowerBFT, a voting protocol adapted from classic Byzantine Fault Tolerant (BFT) designs. This combination has worked well, but two problems compound as the network has grown.

First, true finality (the point at which a transaction is cryptographically irreversible) has taken approximately 12.8 seconds (32 slots at Solana's historical 400ms slot time), even though users typically perceive transactions as “confirmed” within one to three seconds through optimistic client settings[1] . Second, validator votes are themselves ordinary on-chain transactions; at Solana's transaction volumes, these votes have historically consumed a significant amount of available block space, resulting in an overhead in bandwidth, transaction fees and processing load.

Alpenglow is designed to solve both problems at once.

1.2 SIMD-0326 (Alpenglow): Votor and Rotor at a glance

At its core, Alpenglow replaces Proof of History and TowerBFT with two new components: Votor, a new voting and finality mechanism, and Rotor, a new block propagation mechanism. Both are described in technical detail in Section 2. In plain terms, Votor lets validators agree that a block is final in one or two rounds of off-chain voting rather than the slower, on-chain process used today, while Rotor changes how a block's data physically reaches every validator on the network, cutting the number of network “hops” required. When combined, Solana's speed and throughput can be significantly improved.

The governance outcome for SIMD-0326 was one of the most decisive in Solana's history: the proposal passed on 2 September 2025 with 98.27% in favour, 1.05% against and 0.69% abstention, from roughly 52% of total staked SOL participating in the vote. [1,17]

1.3 SIMD-0387 & SIMD-0357: registering for the new system

Two further proposals handle the validator-side mechanics of the transition. SIMD-0387 introduces a new type of cryptographic key (a BLS public key) that every validator must register on-chain before they can participate in Alpenglow consensus. This is a technical pre-requisite to Alpenglow and adds a new public key to validator vote accounts for participation in post-Alpenglow voting. The SIMD-0357, the Validator Admission Ticket (VAT), replaces today's continuous vote-transaction fees with a single, predictable fee charged once per epoch.

A note on rollout sequencing

It is important to distinguish what is live today from what activates later. As of 9 September 2026, validators can register BLS keys, and an admission-gating rule already excludes validators without one from the active consensus set. The VAT fee itself, and the associated removal of on-chain vote transactions, does not take effect until Alpenglow's full consensus switch activates, expected in Agave 4.3 around October 2026. Several press reports have conflated these stages; this report treats them separately throughout.

1.4 Where things stand today

(9 September 2026) [8]

MILESTONE	DATE	STATUS
SIMD-0326 governance vote passes	2 Sept 2025	Complete
Community validator testing begins	11 May 2026	Complete
SIMD-0387 (BLS pubkey registration) live on mainnet	8 July 2026	Complete
SIMD-0357 (VAT) admission gating live on mainnet	22 July 2026	Partial: gating only, fee not yet collected
Agave 4.2 ships full Alpenglow codebase, switch left off	Week of 17 Aug 2026	Complete
SIMD-0525 first slot-time reduction (400ms → 350ms)	21 Aug 2026	Complete
Agave 4.3: full Alpenglow consensus + VAT fee activation	Targeted October 2026	Pending

Section 2: Deep Dive & Institutional Impact Analysis

2.1 Votor: how blocks get finalised

Votor replaces TowerBFT as Solana's voting and finality mechanism. Two paths to finality run concurrently, and whichever completes first determines the outcome:

- Fast path: if 80% of stake is actively participating, a block finalises in a single round of voting. The technical name for this route is ‘fast-finalization.’
- Slow path: if only 60% of stake is responsive, finality still completes, but requires two rounds. The technical name for this route is ‘slow-finalization.’

Voting itself moves off-chain: rather than submitting individual vote transactions, validators exchange compact cryptographic signatures (BLS signatures) that are aggregated into a single certificate anchored on-chain. Instead of Proof of History driving the schedule, validators use a fixed block time (independently being reduced under SIMD-0525 from 400ms to 200ms, see Section 2.5) measured against their own local clocks, issuing a “skip vote” if a block does not arrive before a timeout.

The combined effect is a reduction in cryptographic finality from approximately 12.8 seconds today to roughly 100–150 milliseconds, an improvement of roughly 80–100 times. [1,2,3,13]

Why this matters for institutions

Any institution that treats "finality" as the moment risk formally transfers (an exchange crediting a deposit, a fund settling a redemption, a custodian confirming a transfer) currently has to account for a multi-second window of residual uncertainty that most traditional risk frameworks were never designed to tolerate on a public blockchain. Alpenglow compresses that window to below the latency of many conventional payment rails.

2.2 Rotor: how blocks get distributed

Rotor replaces Turbine, Solana's existing block-propagation system. Turbine distributes block data through a multi-layer relay tree whereby data is distributed to the highest stake first before disseminating to lower staked validators. Rotor replaces this with a single-hop relay broadcast, using the same erasure-coding principle Turbine already relies on, but reducing the number of network hops required for the data to reach the full validator set. [3,4]

It is important to note that Rotor is planned for implementation under its own, separate SIMD on a later timeline than Votor. "Alpenglow" as a term therefore arrives in two distinct technical waves: the consensus and voting change first, with the propagation change following at a later date.

A related consequence of Rotor concerns the market for third-party shred-streaming services, such as bloXroute's Optimized Feed Relay and several similar providers, which sell clients early access to raw shreds ahead of standard network delivery. These services exist because Turbine's current multi-hop propagation tree creates a real latency gradient: well-staked validators typically receive shreds directly from the leader or in the earliest fanout layers, while smaller validators and ordinary participants sit further down the tree and receive the same data measurably later, often by hundreds of milliseconds. Commercial shred-streaming providers sell paying customers, principally high-frequency trading firms and MEV searchers, a way to access shreds ahead of this queue. [10]

Rotor is designed to compress precisely this gradient. Replacing Turbine's multi-hop tree with a single layer of stake-weighted relays that broadcast to the whole network in one round means blocks are expected to reach effectively all validators within a couple of milliseconds of each other, at a latency close to the physical network's own floor, rather than in a long stake-ranked sequence. If baseline propagation converges this tightly for everyone, the multi-hop queue that today's shred-streaming services are built to skip past narrows substantially, and with it much of the commercial rationale for paying a premium for early shred access. This does not mean stake stops mattering to propagation: Rotor's relay set is itself selected in proportion to stake, so well-staked validators continue to play a preferential role in distributing data. It is the specific business of monetising a slow, sequential queue that becomes harder to sustain once the queue itself is this short. Jito's decision to wind down its own ShredStream service during 2026, directing customers instead toward dedicated network infrastructure such as DoubleZero, is an early sign of this market already shifting ahead of Rotor's activation. [16]

2.3 The "20+20" resilience model: the trade-off

Alpenglow changes the assumptions under which Solana remains both safe (the chain never finalises two conflicting versions of history) and live (the chain keeps producing blocks). This is a significantly controversial trade-off inherent in the upgrade.

SECURITY PROPERTY	TOWERBFT (TODAY)	ALPENGLOW ('20+20')
Adversarial (Byzantine) stake tolerated while remaining safe	Up to 33%	Up to 20%
Additional offline / non-responsive stake tolerated while remaining live	None beyond the 33% figure: the chain can stall entirely if more than a third of stake is unresponsive, for any reason	A further 20% on top of the above (40% combined mixed-failure tolerance)

In other words, Alpenglow accepts a lower ceiling against a purely coordinated, malicious attack in exchange for a meaningfully higher tolerance for the kind of failure that has actually caused Solana outages historically: client bugs, network congestion, and correlated but non-malicious validator crashes. The Solana Foundation's own governance documentation for SIMD-0326 is explicit that this is a deliberate trade-off, made because one-round voting and 40% combined crash-failure resilience were judged worth the reduction in pure adversarial tolerance. [1]

We return to this trade-off, and our own view of it, in Section 3.

2.4 The Validator Admission Ticket: what it costs, and what's live today

Before any reduction in target slot times under SIMD-0525, validators pay ongoing transaction fees every time they vote, a cost that was roughly 2.1 SOL per epoch (each epoch is approximately two days at current slot times), or somewhere in the region of 360–390 SOL per year per validator.

SIMD-0357 replaces this system with a flat Validator Admission Ticket: a single, predictable fee of 1.6 SOL per epoch at a 400ms slot time, burned rather than paid to any party. Because SIMD-0525 (Section 2.5) is independently shortening epochs, the VAT fee is designed to scale down at each slot-time stage so that the effective daily cost stays roughly constant at approximately 0.8 SOL per day throughout the transition. [6,7] It's worth noting that the slot reductions under SIMD-0525 have already begun, however the Validator Admission Ticket has not, resulting in temporarily increased voting costs for validators.

SLOT TIME	EPOCH LENGTH	VAT PER EPOCH	APPROX. VAT PER DAY
400ms (baseline)	48 hours	1.6 SOL	~0.8 SOL
350ms (already live)	42 hours	1.4 SOL	~0.8 SOL
300ms (already live)	36 hours	1.2 SOL	~0.8 SOL
250ms	30 hours	1.0 SOL	~0.8 SOL
200ms	24 hours	0.8 SOL	~0.8 SOL

Source: SIMD-0525, official VAT scaling table. [7]

What's actually live today (9 September 2026)

Only the admission-gating half of the VAT is currently active: validators without a registered BLS key are excluded from an admitted validator set (currently capped at 2,000 validators). The fee itself (the 1.6 SOL deduction) and the corresponding removal of ordinary on-chain vote transactions do not take effect until Alpenglow's consensus switch activates in Agave 4.3. Institutions should not assume validator cost savings, or reduced blockspace consumption from votes, are already reflected in current network conditions.

2.5 Block times, leader windows and validator timing games

This section addresses a question we are asked often by institutional stakers: does Alpenglow change validators' ability to strategically time block production for extra reward capture, and if so, how does that affect the rewards institutions should expect to receive?

“Timing games” refers to a validator's ability, under the current system, to exploit the gap between when a validator is assigned a slot and when the network actually requires the resulting block, delaying transmission or adjusting which transactions are included in order to capture additional fee or MEV revenue at the margin. This behaviour has been documented [11] and quantified by our own research into validator and MEV incentives on Solana today, discussed further below.

A related but independent proposal, SIMD-0525, is shortening Solana's slot time in stages: from 400ms to 350ms (already live, since 21 August 2026), and onward through 300ms, 250ms and eventually 200ms. Because a validator's leader window is fixed at four slots, this schedule directly compresses the wall-clock time any single leader controls: 1.6 seconds today, falling to 800 milliseconds once 200ms slots are reached. [7]

Two distinct mechanisms drive this, and they are worth separating clearly. The more fundamental one is Votor's fixed timeout. Today, a late block has room to manoeuvre: the network's tolerance for lateness is loose and probabilistic, so a leader who holds a block back to capture a better fee or MEV opportunity risks little. Under Votor, each slot has a fixed deadline set in advance, and a block that misses it is simply skipped, meaning the leader forfeits the whole slot's reward, not just the marginal opportunity they were chasing. That flips the arithmetic of delay: the potential upside from waiting a little longer is small, while the downside of waiting too long is total.

The second, compounding mechanism is SIMD-0525's leader-window compression. A leader's window is fixed at four slots, and as slot time falls from 400ms toward 200ms, the wall-clock time available within that window shrinks from 1.6 seconds to 800 milliseconds. Even a leader willing to risk the timeout simply has less absolute time in which to delay, reorder, or selectively include transactions before the next leader takes over. The proposal's own stated motivation is explicit on this point: shortening the window is intended to reduce “the worst-case time a leader can delay, reorder, or selectively include transactions before the next leader has an opportunity to produce a block.” [7]

Together, these narrow the surface area for timing-game strategies considerably, though not to zero. The same proposal flags a tradeoff worth noting: a shorter, harder-to-target window can also make it more difficult for legitimate infrastructure, market makers and searchers included, to reliably reach the correct leader in time, which may push some participants toward broadcasting to a wider set of nodes rather than one precise connection. That same pressure should also push serious participants toward lower-latency infrastructure generally, RPC providers included, arguably a net improvement to the quality of Solana's infrastructure stack even where it was not the primary goal.

Block Proposal Timing Strategies

Given the first two phases of this SIMD have already been implemented on mainnet, the below section provides some analysis as to how this has impacted Solana slot times as well as the timing game behaviour across different operators and clients including:

- *Historic Epoch lengths*
- *Distribution of slot times across clients over the first two phases*
- *Distribution of slot times across validators over the first two stages*
- *Qualitative description of different proposal timing strategies*

Distribution of slot times between validators and clients

Network-wide slot time held flat for most of the past year, drifting from the ~390ms to the ~420ms through mid-2026. SIMD-0525's first two stages went live resulting in two step changes with slot times reaching ~366ms average at epoch 1020 (350ms target, live 21 August) and ~315ms at epoch 1024 (300ms target, live 30 August). Both steps landed close to their nominal target with dispersion in mean slot times across both validator operators and the clients they are run on. Measured mean slot time has remained a few per cent over whatever the target is throughout the transition, which is expected given inherent latency combined with timing game behaviour throughout the validator set. However the excess in slot times did widen slightly at the second stage: about 5.7% over 300ms against 4.5–4.8% over 350ms in the prior stage.

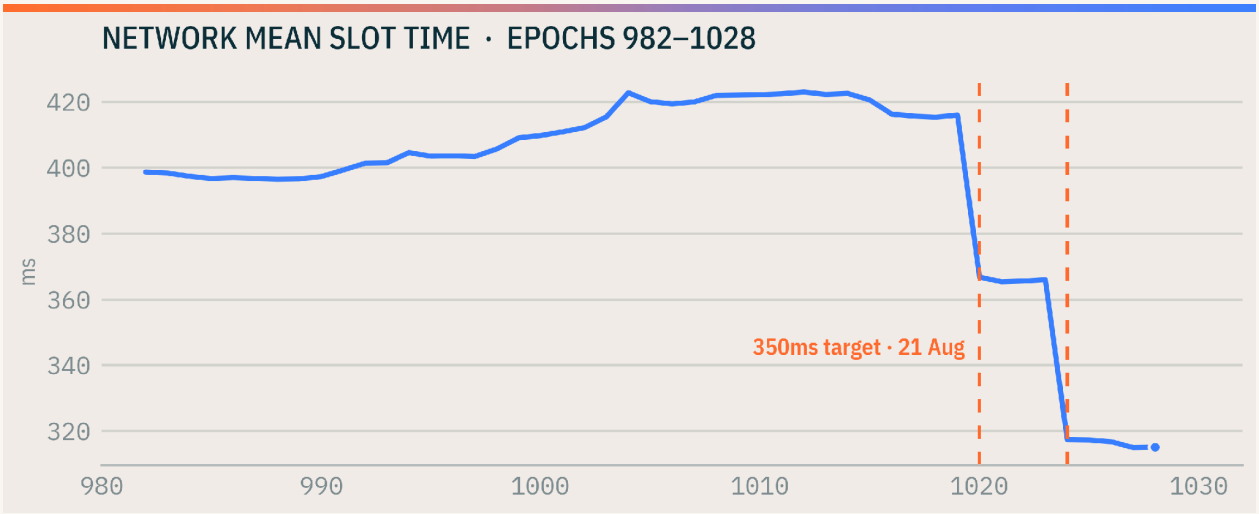

Figure 1: Implied network slot time per epoch, derived from epoch duration ÷ 432,000 slots.

Analysing slot times over the transition across different clients also provides useful insight. All validator client averages, across both the Agave and Firedancer/Frankendancer codebases, stepped down at both epoch 1020 and epoch 1024 as expected given the upgrade, and the relative lag differentials between them barely moved. Firedancer-Jito's raw speed advantage over Agave (non-vanilla) narrowed in absolute terms as slot times fell (about 34ms pre-1020, 24ms post-1024) but has held close to 7–8% of slot time throughout. Conversely, the Firedancer/ Frankendancer and Harmonic combination consistently had the slowest slot times. It is worth noting that this data is not necessarily indicative of client performance itself. As an example, firedancer is a newer architecture when compared to agave. Therefore it is expected that operator sophistication is positively correlated to Firedancer usage, as well as to the likelihood that they are implementing timing game, a manual configuration independent of client performance. Therefore there is a logical causation between client selection and slot times beyond purely attributing it to client performance.

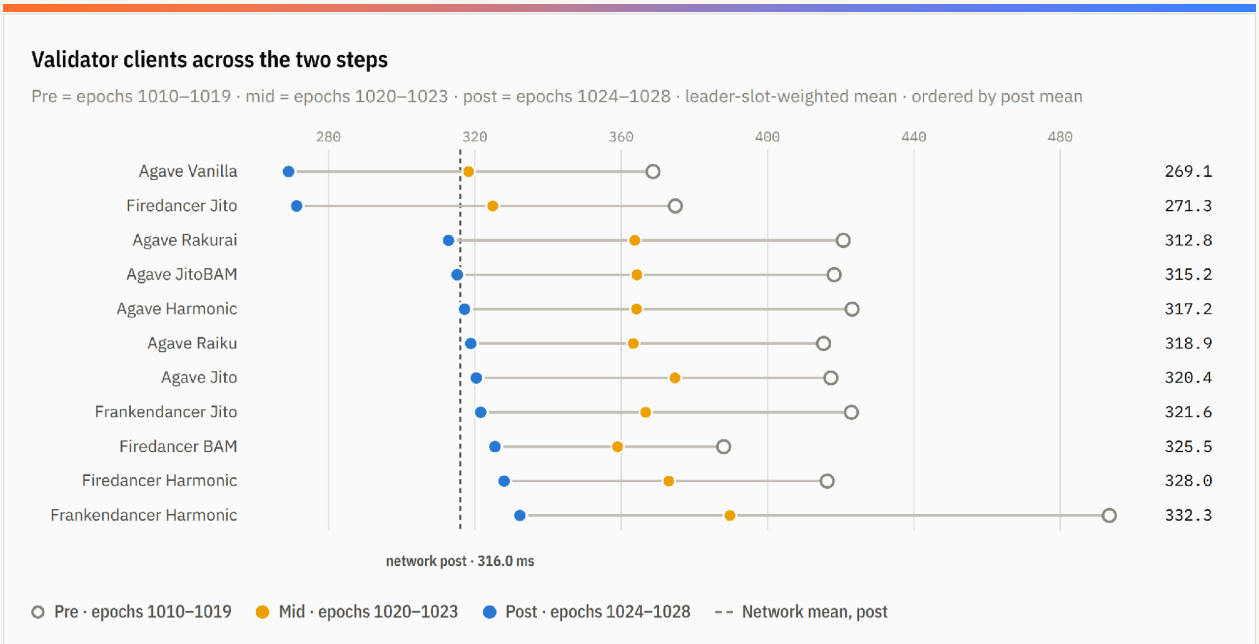

Figure 2: Mean slot time for four representative client configurations, pre/mid/post the two SIMD-0525 steps.

At the individual validator level, the most recent complete epoch (1028) had 660 of the network's 692 reporting vote accounts recording at least one leader slot. Three-quarters of these sit within 10ms of the network mean, with a long right tail of mostly low-stake validators averaging above 350ms, with some outliers lagging beyond 600ms. It is natural that the skew will be positive given that the minimum slot time is bound by network latency and client enforced minimum slot times whereas the maximum is theoretically unbound. It is slightly surprising however, that some validators are hitting 250ms averages, well below the target.

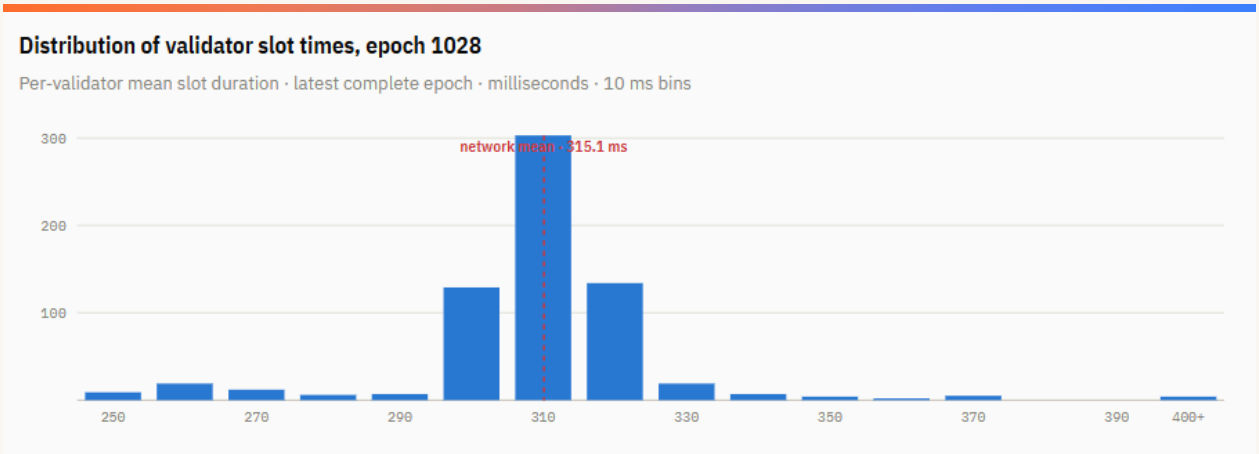

Figure 3: Per-validator mean slot duration, epoch 1028, 10ms bins.

Different slot timing strategies (delaying all slots vs. first or last slot)

Mean slot time doesn't capture the full range of within-window timing strategies. Each validator is assigned a four-slot leader window and can either spread any delay evenly across all four slots or concentrate it in just the first or last while keeping the rest near the 300ms target; the second approach lifts the mean only slightly while the median stays close to target, and is arguably worse for users, since transactions in the delayed slot are held back far more than the average implies. Our dataset measures per-validator mean slot time across a full epoch, not where within the window that time falls, so distinguishing these strategies would need slot-level data keyed to leader-window position, which is beyond the scope of this report.

How timing games have changed over the recent slot time reductions

Before epoch 1020 the peer set of providers we track was loosely spread around the network's 419.4ms mean. In the period after the first slot reduction, most of the dispersion tightened onto a narrow band around the new 365.9ms floor, nine of eleven displayed validators were within about 11ms of each other. After epoch 1024 the same group stepped down again, to within about 14ms of each other. Two providers did not compress at either step: they picked up a structural premium of roughly 50ms and 43ms respectively above the network mean back in August, before either reduction took effect, and have carried close to that same absolute lag through both steps since, even as the network floor itself moved twice.

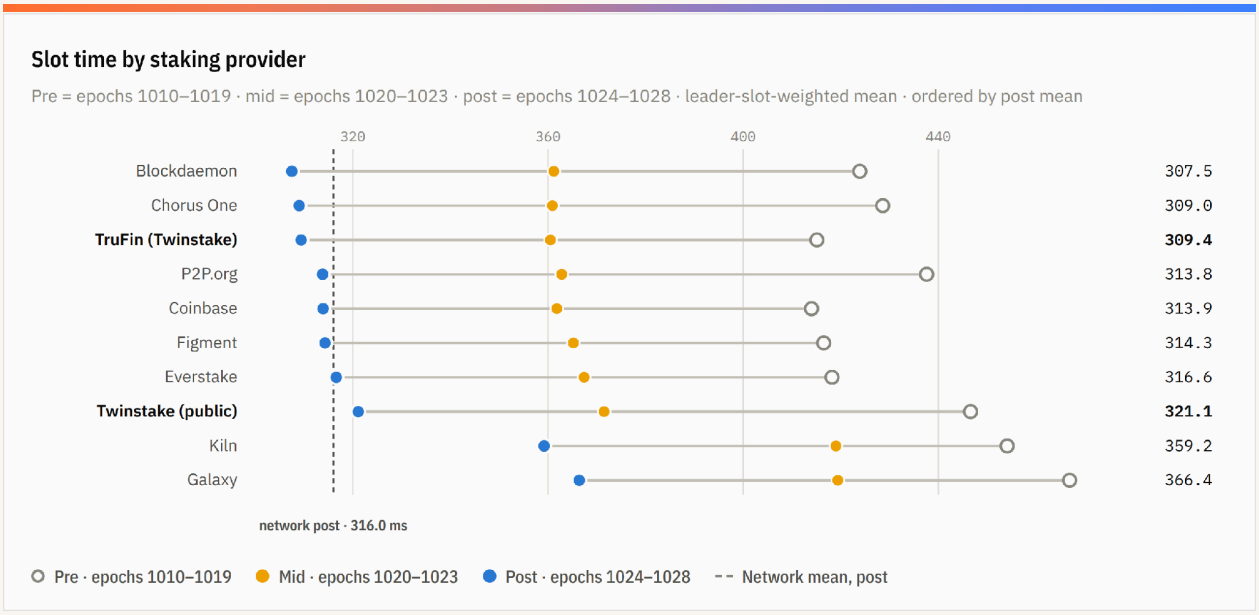

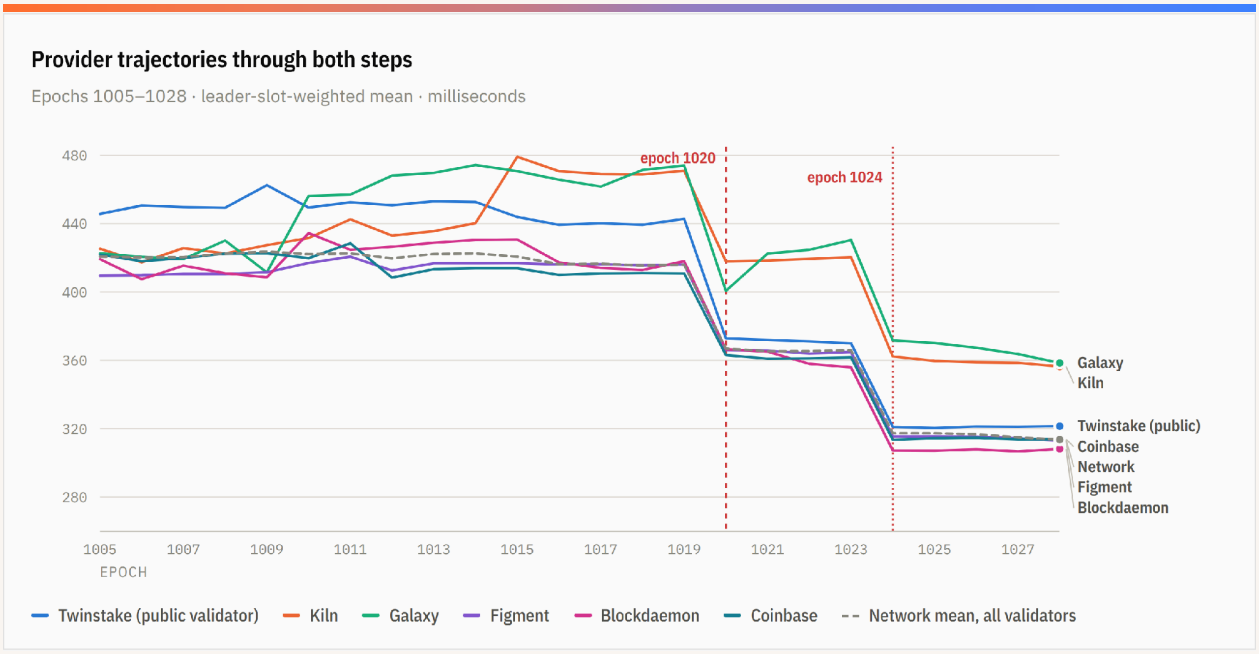

Figure 4 and 5: Provider proposal block times during epochs 1005–1028. Taken from public data.

This is consistent with the rationale in Section 2.5: the mechanisms narrowing timing-game incentives appear to be taking hold at the protocol level, without yet reaching the handful of validators carrying the largest fixed latency lag strategies. It also shows the variety in timing strategies across operators.

A final point regarding timing games. A validator's staking reward performance is a stochastic function of various parameters beyond just timing games. Hardware, geolocation, client configurations and block building strategies form some of these parameters. Twinstake has consistently delivered rewards consistently above the network average throughout the past 4 years, including throughout the SIMD-0525's transition. Over the past 60 days, Twinstake's flagship validator has delivered an APY of 7.81% compared to the network average of 7.51%.

2.6 Epoch length, liquidity management and redemption constraints

SIMD-0525 does not only shorten leader windows. Because epochs remain fixed at 432,000 slots throughout the staged rollout, reducing slot time shrinks epoch length by the same proportion: from approximately 48 hours today, down to roughly 24 hours once the final 200ms stage is reached. This matters well beyond timing games, because epoch length is also the mechanism that governs how quickly staked SOL becomes liquid again after an unstake request.

Solana's stake activation and deactivation only take effect at epoch boundaries. A validator or delegator who submits a deactivation request must wait for the following epoch boundary before the position finishes cooling down, and a separate withdrawal transaction then moves the SOL back to a liquid balance. In practice, at today's roughly 48 hour epoch, this typically means a total wait of one to just over two days depending on when in the epoch the request is submitted and the length of that epoch.

There is also a network-wide throttle worth noting for large redemptions: no more than 25% of total active stake can change state (activate or deactivate) across the entire network in a single epoch. Above that threshold, the excess simply rolls into the following epoch. This is mostly a tail-risk consideration, relevant to correlated, market-wide redemption events rather than routine flows, but it is the scenario in which shorter epochs matter most, since even a throttled, multi-epoch wait compresses in step with the shorter epoch length.

Halving epoch length therefore roughly halves the worst-case time needed to turn staked SOL back into liquid SOL, for the entire staked base, not just for new entrants. This has two institutional applications worth spelling out.

Applicability to staked ETF structures and T+1 redemption

US-listed ETFs generally settle investor redemptions on a T+1 cycle. Staked crypto ETFs need to be able to source that liquidity reliably, and the design of the underlying network's unstaking mechanism has turned out to be a first-order constraint on how those products are built. Ethereum is the clearest illustration: its unbonding period is queue-based rather than fixed, and has ranged from under two weeks to beyond fifty days depending on how much stake is simultaneously trying to exit. Because that duration is both long and unpredictable, staked Ethereum ETFs cannot stake their full holdings; issuer filings show typical staked ratios in the region of 50 to 95% (depending on their available alternate lines of liquidity), with the unstaked remainder held purely as a liquidity buffer to fund redemptions, directly diluting the yield investors receive.

Solana's epoch-bound cooldown is already a materially better starting point: short, and bounded by a known, protocol-fixed schedule rather than an open-ended queue. Whilst epochs are not fixed temporally, they have statistically confident time intervals based on typical slot times, at which they are extremely likely to end. This is likely why Solana ETFs typically have higher staking utilization rates of the AUM of the fund, with many staking between 70-100% of the AUM, a structural difference from its Ethereum counterparts rather than a marginal one.

As SIMD-0525 progresses, the worst-case liquidity turnaround for a fully staked structure tightens further, from roughly two days toward roughly one day at the final stage. A one day epoch makes it materially easier to justify staking higher proportions of the ETF, and improves the reward rate the fund can pass on to its investors. Beyond existing SOL ETFs, future asset managers considering launching similar products will be more incentivised to set up a Solana ETF knowing the operational challenges of managing liquidity is lower, and the reward rate they can offer investors is higher.

Applicability to exchange omnibus account management

Exchanges and custodians that pool customer funds into omnibus accounts (with the option to stake funds), face a related but distinct problem: individual customers expect to be able to withdraw on a normal exchange timescale, which is typically much faster than the underlying stake can actually be deactivated. To bridge that gap, the operator holds an internal liquidity buffer of unstaked SOL, sized against the worst-case time needed to unwind enough underlying stake to replenish it after a wave of customer withdrawals. These buffers are evaluated on a frequent basis, accounting for client redemptions and new client holdings.

A shorter, more predictable epoch-bound cooldown changes the considerations around the buffer an operator needs to hold against a given expected redemption volume, potentially freeing a larger share of the pooled book to be productively staked rather than held idle, which improves blended reward rate across the omnibus pool. It also reduces the tail-risk scenario described above: a market-wide redemption spike large enough to trigger the network's 25% per-epoch throttle is precisely the moment an operator's own liquidity buffer is under the most strain, and a shorter epoch shortens that specific worst case as well.

2.7 Reward predictability for institutional stakers

Even once Alpenglow is fully live, individual validator rewards are unlikely to become perfectly deterministic, though the specific picture here is still developing rather than settled. Rotor's design intends for relays, the stake-weighted subset of validators resampled each slot to redistribute shreds, to be compensated for that bandwidth cost, with heavier relayers expected to receive a larger share. That said, neither the Alpenglow whitepaper nor SIMD-0326 itself defines a concrete mechanism for calculating or distributing these rewards: SIMD-0326 explicitly preserves today's existing vote-reward mechanics unchanged and defers new economic mechanisms, including how relays would eventually be paid, to a future, separate economics-focused SIMD that has not yet been published. What is already confirmed is that the relay role itself is not evenly distributed, since a fresh stake-weighted subset is sampled each slot rather than every validator relaying every time. If a concrete relay-reward mechanism is specified, this sampling may contribute to additional reward variance.

For institutions running or delegating to a small number of validators, this variance is more pronounced. Large, multi-validator operators, including institutions delegating across a diversified validator set as is typical practice at Twinstake, see realised rewards converge much closer to the network-wide average. This mirrors the reward-variance dynamics we highlighted in our earlier Pectra research on Ethereum proposer and MEV variance: scale reduces variance, but does not eliminate the underlying randomness in the reward mechanism itself.

2.8 Settlement risk and regulatory relevance

CONFIRMATION TYPE	TOWERBFT (TODAY)	ALPENGLOW, FAST PATH (80% STAKE RESPONSIVE)	ALPENGLOW, SLOW PATH (60% STAKE RESPONSIVE)
Optimistic confirmation	~1–3 seconds (client-dependent)	Broadly unchanged	Broadly unchanged
Full cryptographic finality	~12.8 seconds (32 slots)	~100–150 milliseconds	~150–180 milliseconds

Sources: [1] [13]

This has direct regulatory relevance. The SEC's March 2026 interpretive guidance on the application of federal securities laws to crypto staking arrangements turns on whether staking rewards are generated by transparent, protocol-determined processes rather than the discretionary efforts of an intermediary, the central test under Howey. Faster, deterministic finality strengthens precisely this argument: it is a protocol-level property, not a service provided by any intermediary, and it reduces the residual uncertainty that has historically complicated how staking arrangements are characterised for these purposes. [12]

It is also directly relevant to custody and fund-structure design. Deterministic, sub-200ms finality speaks to the operational criteria custodians and ETF issuers apply when evaluating validator infrastructure for redemption processing, NAV calculation timing, and audit reporting, areas where a multi-second window of residual uncertainty has historically been a genuine, if often unstated, source of operational risk.

Section 3: Twinstake Opinion Piece – Solana as Internet Capital Markets Infrastructure

Solana's own core contributors (Anza, Jito, DoubleZero, Drift and Multicoin Capital, writing jointly in July 2025) describe Alpenglow as one of three pillars required to deliver on Solana's stated ambition to become the backbone of what they call "Internet Capital Markets": a single, always-on, global venue where every asset class trades on the same infrastructure as everyday payments. Solana Foundation President Lily Liu reiterated this framing publicly at Consensus Hong Kong in February 2026, arguing that blockchain's core advantage lies in building open, tokenised capital markets rather than in general-purpose Web3 experimentation. [13]

We think this is the right lens through which institutional stakers should evaluate Alpenglow: not as a speed upgrade in isolation, but as the specific piece of Solana's roadmap addressing the properties that matters most for financial infrastructure: speed, predictability and reliability.

Speed and reliability are prerequisites for financial infrastructure

Financial infrastructure has a far lower tolerance for unpredictability than most other blockchain use cases. A payments network or trading venue that is fast most of the time, but occasionally takes twelve or more seconds to become final, or periodically stalls under load, cannot become the settlement layer that serious financial institutions build businesses on top of, regardless of how impressive its throughput headline numbers are the rest of the time.

Anza's own framing of this problem is useful: they distinguish time-to-inclusion (how quickly a transaction lands on-chain) from time-to-finality (how quickly it becomes irreversible). They argue that shrinking both is what allows market makers to quote tighter, safer prices on-chain without taking on "gap risk" from stale quotes that the market has already moved past [13]. Alpenglow is precisely the half of that equation aimed at finality; the bandwidth and latency improvements referenced elsewhere in Solana's roadmap (and in DoubleZero's network infrastructure, which we have covered in separate research) address the other half.

Twinstake recently released a report tracking 'Solana's Evolution to Institutional Grade Trading Infrastructure [11],' which follows a chronological account of how Solana has undergone a profound infrastructure transformation over the past five years. The article focuses on high frequency applications in particular, although the need for low latency systems with high throughput and reliable inclusion forms the foundation for any capital market infrastructure. Alpenglow continues this trajectory and, if successful, will bring Solana another step closer to achieving its mission.

For institutional allocators, this is the practical translation of "Internet Capital Markets": not a slogan, but a specific, measurable reduction in the operational uncertainty that has kept the most risk-conscious participants (banks, asset managers, and regulated custodians) evaluating Solana from the sidelines rather than building on it directly.

Addressing the BFT trade-off directly

The reduction in adversarial fault tolerance from 33% to 20%, discussed in Section 2.3, is the part of this upgrade that deserves the most direct scrutiny, and we do not think it should be waved through on the strength of the network's other improvements alone. The Solana Foundation's own proposal documentation is explicit that this is a deliberate choice, accepted in exchange for faster, single-round finality and stronger resilience to crash failures. The question worth asking directly is whether that exchange is a good one for the institutions this report is written for.

The case for it starts with what actually causes Solana to stop producing blocks. The network's historical liveness failures have overwhelmingly come from congestion, client bugs, and correlated but non-malicious validator crashes, rather than from coordinated attacks. It is precisely this failure mode that the 20+20 model is designed to tolerate far better than the system it replaces, which could stall entirely once a third of stake became unresponsive for any reason at all, malicious or otherwise.

It also holds up on a straightforward cost basis. Solana currently has roughly 404 million SOL staked, worth approximately $41.7 billion at recent prices. Twenty percent of that figure, the stake an attacker would need to control to threaten the network's safety guarantee, is approximately $8.3 billion: a fully attributable, slashable, and economically self-destructive undertaking. The practical deterrence lost by moving the threshold from 33% to 20% is smaller than the percentage-point difference alone suggests, because both figures sit well above what an economically rational attacker would risk at current stake values.

Put together, an institution assessing Solana as candidate financial infrastructure is better served by weighing demonstrated, real-world uptime against a purely hypothetical and extremely costly attack scenario, rather than treating a bare percentage-point comparison between BFT thresholds as the whole story.

None of this makes the trade-off costless. A lower Byzantine threshold is a real change, and institutions with a mandate to model worst-case adversarial scenarios specifically, rather than realised operational risk, should still factor the 20% figure into their own frameworks. Weighing both sides, our own conclusion is that Alpenglow is a net positive for Solana's suitability as financial infrastructure, precisely because it targets the failure mode that has actually caused problems historically, rather than the one that makes for a cleaner whitepaper comparison.

Our view, in summary

Alpenglow, taken together with the slot-time reductions in SIMD-0525 and Solana's broader Internet Capital Markets roadmap, meaningfully strengthens the case for Solana as venue infrastructure for institutional capital. Not because it makes Solana faster alone, but because it removes specific, identifiable sources of unpredictability (slow finality, uncertain validator costs, and exploitable timing-game incentives) that have been genuine, rather than perceived, obstacles for the most risk-conscious allocators.

Twinstake continues to track Agave 4.3's rollout closely through the remainder of 2026 and will update clients directly as validator-side readiness milestones are confirmed.

References & Further Reading

[1] SIMD-0326: Alpenglow (Solana Improvement Documents, GitHub): https://github.com/solana-foundation/solana-improvement-documents/blob/main/proposals/0326-alpenglow.md

[2] Solana Status, official X account confirming SIMD-0326 governance vote results: https://x.com/SolanaStatus/status/1962875262938628136

[3] TechTimes, "Solana Confirms Alpenglow Activates in October: Agave 4.3 Carries Switch": https://www.techtimes.com/articles/324490/20260814/solana-confirms-alpenglow-activates-october-agave-43-carries-switch.htm

[4] Anza, "Alpenglow: A New Consensus for Solana": https://www.anza.xyz/blog/alpenglow-a-new-consensus-for-solana

[5] Helius, "Alpenglow": https://www.helius.dev/blog/alpenglow

[6] Anza, "The Internet Capital Markets Roadmap": https://www.anza.xyz/blog/the-internet-capital-markets-roadmap

[7] Solana.com, Alpenglow upgrade page: https://solana.com/upgrades/alpenglow

[8] Helius, "Solana Shreds: Winning the Millisecond Game": https://www.helius.dev/blog/solana-shreds

[9] Jito, "Low Latency Block Updates (ShredStream)", service deprecation notice: https://docs.jito.wtf/lowlatencytxnfeed/

[10] SIMD-0357: Validator Admission Ticket (Solana Improvement Documents, GitHub): https://github.com/solana-foundation/solana-improvement-documents/blob/main/proposals/0357-alpenglow_validator_admission_ticket.md

[11] SIMD-0525: Reduce Slot Times (Solana Improvement Documents, GitHub): https://github.com/solana-foundation/solana-improvement-documents/blob/main/proposals/0525-reduce-slot-times.md

[12] Twinstake, "Solana's Evolution to Institutional Grade Trading Infrastructure": https://www.twinstake.com/reports/solanas-evolution-to-institutional-grade-trading-infrastructure

[13] King & Spalding, "Staking out the Perimeter: SEC Outlines Limits of Jurisdiction Over Digital Assets" (covering SEC Release No. 33-11412, 17 March 2026): https://www.kslaw.com/news-and-insights/staking-out-the-perimeter-sec-outlines-limits-of-jurisdiction-over-digital-assets

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 report 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 report to be directed at you.

References to network performance, rewards, finality and planned protocol upgrades are based on information available as at 24 August 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. Alpenglow has been approved through validator governance but is not yet active on mainnet, and the performance figures in this report are targets drawn from published proposals, testing and simulations rather than guarantees. 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 report 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 II (Directive 2014/65/EU) (known as 'MiFID II'); if you are in the UK, to "Investment Professionals" or "High Net Worth Companies" as defined in Articles 19 and 49 respectively of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005, or as otherwise defined under applicable local regulations and at whom this report 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 report 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.