BLOG

DEVELOPMENT

Zenith: Zano's Move to Pure Proof of Stake

Zano's next consensus is pure proof of stake, developed with Common Prefix and built on Zarcanum's privacy foundations. What Zenith is, why we chose it, and what it changes for users, node operators, and holders.

Ravaga
Jul 16, 2026 • 11 min read

Contents:

Subscribe

Consensus is the part of a blockchain most users never think about, and it's the part everything else depends on. It decides who is allowed to add the next block, how the network resolves disagreements over which version of history is correct, whether it keeps working if it's attacked, and when a transaction can be treated as reliably confirmed.

For a privacy chain, the problem is harder. The protocol has to reach agreement without learning the balances, identities, or activity of the people securing it. A network shouldn't have to pick between strong consensus and strong privacy, and Zano never accepted that trade.

After more than 12 years of building and running blockchain networks, repeated consensus upgrades, and a long research collaboration with Common Prefix, we're ready to present the direction of Zano's next major evolution: Zenith, a pure proof-of-stake consensus protocol built on the privacy foundations of Zarcanum.

This isn't a pivot. Zenith is where a very long line of work finally converged.

Why Consensus Deserves More Attention Than It Gets

A project can have sophisticated ring signatures, hidden amounts, stealth addresses, and other advanced privacy mechanisms. None of that secures the network by itself if an attacker can reorganize the chain, manipulate block production, keep honest participants from extending the ledger, or exploit the rules that decide which version of the chain wins when there's a disagreement. Privacy and consensus solve different problems, and a serious privacy network has to solve both.

There's still a common belief in this space that proof of work is inherently sufficient and proof of stake is inherently suspect. Some of that traces back to legitimate concerns about early PoS designs. A lot of it is tribalism, terminology, or treating every PoS protocol as though it were the same protocol.

PoW and PoS are broad families of mechanisms, not two fixed designs with fixed properties. A PoW network isn't secure just because it burns computation, and a PoS network isn't secure just because it locks up stake. What matters is the specific construction: what it assumes an attacker can do, how it decides which chain wins, how it picks who produces the next block, its privacy model, incentives, network assumptions, and behavior under realistic attack.

Our views on this weren't formed in the abstract. Across 12+ years of running blockchain projects we've been attacked directly, handled adversarial behavior in production, and watched the wider industry learn its lessons through repeated successes and failures. That kind of experience teaches things papers can't: how attackers behave when real economic value is on the table, how components that look independent interact under pressure, and how assumptions that seem reasonable on paper fail on a live network.

It also teaches that consensus is never a finished component just because the chain is running. Threat models evolve. Markets for rented hash power mature, hardware becomes more specialized, network conditions change, attacker tooling improves, and the value available to exploit keeps growing.

A privacy-preserving PoS protocol must solve one additional problem on top of all this: the network must determine whether someone is eligible to produce a block without publicly revealing that person's stake, balance, or transaction history. Zano pioneered this with Zarcanum. Zenith takes the next step by extending private staking into a complete, pure-PoS consensus design.

We hope this work contributes to a broader reassessment across the privacy ecosystem. Moving away from PoW shouldn't be ideological, and adopting PoS shouldn't be fashion. The question is which protocol gives a particular network the strongest practical security. For Zano, our research says that's now a carefully constructed privacy-preserving PoS.

Where This Started

The starting point goes back to 2013, when Zano's lead developer built his first pure proof-of-work blockchain. Around 2015, that experience turned into a broader question: could PoW and PoS be combined into a network more resilient than either mechanism alone? Several years of research and engineering on that question fed directly into the creation of Zano. When the mainnet launched in 2019 with a hybrid PoW/PoS design that mixes proof of work (mining) and proof of stake (staking), few privacy projects were taking proof of stake seriously at all.

The protocol has been through several substantial revisions since. The biggest was Zarcanum, which made it possible to stake while keeping amounts and staking activity private.

Each revision improved the network and raised new questions. A few years ago we opened another research phase focused on the next generation of Zano consensus and brought in Common Prefix, a research and engineering team specializing in provably secure blockchain protocols. They analyzed our hybrid consensus, studied realistic attack scenarios against it, and worked with us on designs for a pure, privacy-preserving PoS protocol suited to CryptoNote-style payments. That collaboration ended up producing two research directions and two papers.

The first is *StakeNote: A Proof-of-Stake Protocol for CryptoNote Payments*. It presents security and privacy arguments, along with a proof-of-concept implementation, demonstrating that private CryptoNote-style payments and modern PoS consensus can be combined efficiently.

The second is *Zenith*, built directly on Zarcanum's cryptographic foundations. Zenith is the design we're implementing in Zano.

Both designs hold up. We chose Zenith for practical engineering and integration reasons: it extends the private-staking machinery already native to Zano instead of replacing it with a separate architecture. StakeNote is not a discarded idea; it stands as an independent research result, and other CryptoNote-based systems may find it directly useful. For Zano, Zenith was simply the clearest route from the current architecture to a fully private, pure-PoS network.

How Trying to Fix Hybrid Consensus Led Us to Drop It

We didn't set out to remove proof of work.

This research phase began as another attempt to improve Zano's hybrid consensus. The plan was evolutionary: keep the PoW/PoS structure, work through the weaknesses and limits we'd identified, ship a stronger version of the existing model.

It kept hitting dead ends.

Not because the current hybrid consensus is suddenly broken; it isn't. But every attempt at a substantial improvement inherited structural constraints from the PoW component. Solving one problem introduced a new dependency. Strengthening one part of the fork-choice logic exposed an edge case somewhere else.

After enough impasses, we changed the question. Instead of endlessly repairing the interaction between PoW and PoS, what does the network look like if that interaction simply isn't there?

The early pure-PoS work was exploratory rather than a fixed plan. We proposed constructions, tore them apart, reworked them, and abandoned some entirely. Out of that came the two research paths described above. We ran both for a while; Zenith and StakeNote take different approaches to privately selecting the next block producer, consensus security, and integration with CryptoNote-style confidential transactions. Zenith eventually emerged as the practical candidate for Zano, and we put it through deeper analysis covering

security properties, operational behavior, privacy guarantees, and compatibility with the existing network. The results are presented in the Zenith whitepaper.

Zenith allows Zano to remove the PoW component while preserving the privacy properties Zarcanum already provides. Stakers participate without exposing the amounts they control, and producing blocks doesn't create an obvious, permanent link between their outputs and their identity.

That part is non-negotiable. A conventional transparent PoS, with visible validator balances and traceable staking histories, would gut the guarantees that define Zano. The goal was never just pure PoS. It was pure PoS without giving up private staking.

One thing we refuse to do is describe any consensus design as invulnerable. No responsible protocol team should. Security depends on clearly stated assumptions, correct implementation, extensive testing, and continued review. What we can say is this: Zenith resolves the architectural problems that repeatedly blocked our hybrid-consensus research, and, within the model and conditions analyzed so far, we are not aware of any exploitable vulnerability in its construction. For a protocol that also keeps its stakers private, that's a real result for both Zano and the wider effort to reconcile strong privacy with modern PoS security.

Faster Blocks, Faster Confirmations

Security is the main reason Zenith exists, but it's not the only thing users will notice. The new protocol supports a much faster block cadence. In short:

- Block time: about 15 seconds, down from Zano's current one minute.
- Confirmations needed: 4 to 6 blocks, down from 10 today.
- Typical confirmation time: about 60 to 90 seconds, down from several minutes.

And it's not just faster: six Zenith confirmations will carry far stronger security guarantees than the ten confirmations we wait for today. These parameters still have to survive implementation testing and network-level validation, but they're the operating range we're designing toward.

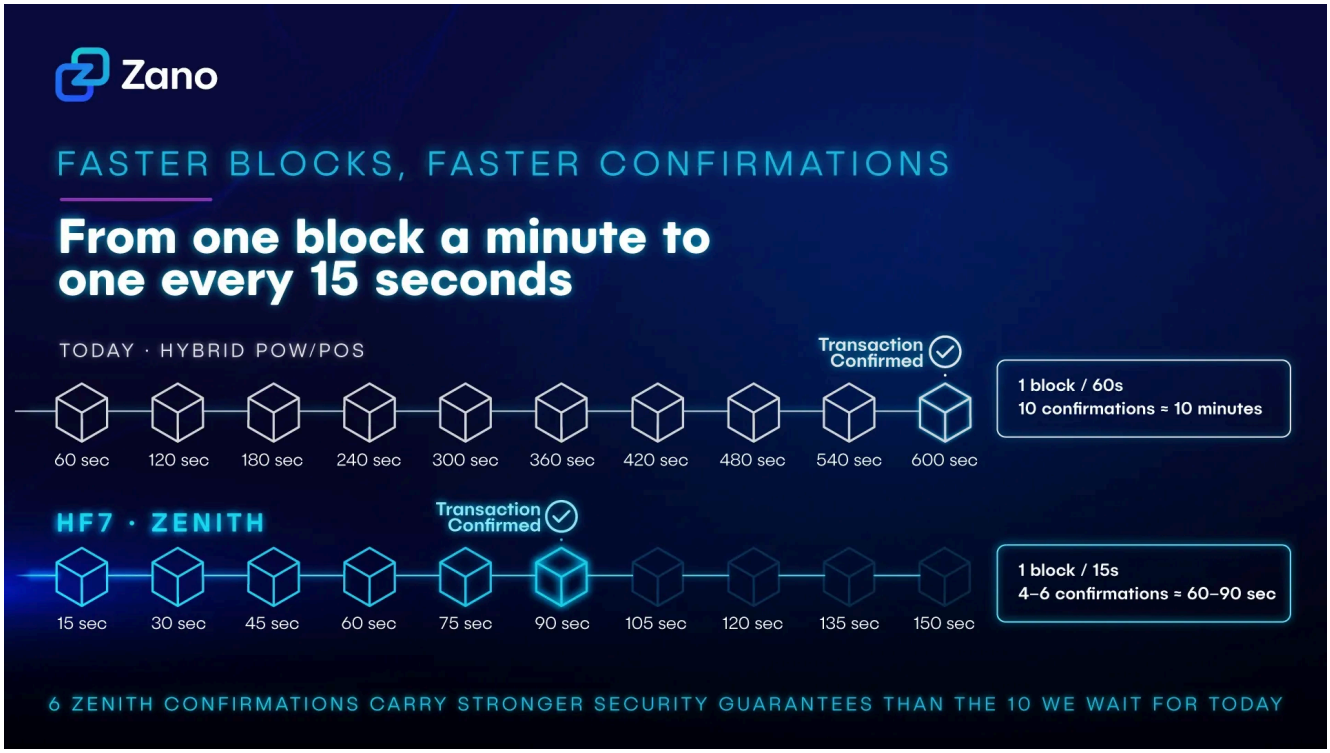

Fast confirmation stopped being a luxury some time ago. For an everyday payment, several minutes of uncertainty is the difference between a working checkout and an impractical one. For decentralized trading, cross-chain transfers, merchant systems, automated applications, and ordinary wallet-to-wallet transfers, latency accumulates at each stage, and a delay at the consensus

layer propagates through the entire workflow. Paying privately shouldn't mean accepting a user experience from 2014.

A short block interval is only worth having if the consensus underneath can carry it safely, though. Cut block time without dealing with propagation delays, competing chains, and fork resolution, and you get the appearance of speed with weaker confirmations behind it. Zenith is designed around the higher cadence rather than having it bolted on. The point is to give users rapid, repeated evidence that the network keeps building on their transaction, without quietly relying on weaker security assumptions to do it.

In practice: a transaction lands in a block quickly, confidence accumulates over the next several blocks, and the recommended confirmation threshold is reached in about a minute to a minute and a half. For private payments, Confidential Assets, decentralized trading, bridges, and application integrations alike, that responsiveness matters.

Ephemeral Blocks: More Blocks Without Heavier History

Four times as many blocks raises an obvious question: does the chain become four times heavier to store and synchronize?

Under a conventional design, it could. Zenith approaches historical data differently, through what we call ephemeral blocks.

The observation underneath is that a block isn't equally valuable forever. Near the tip of the chain, frequent blocks do real work: they

confirm transactions quickly, demonstrate continued agreement among participants, and let the network react to competing histories in real time. Once an epoch is finalized and sits far enough behind the tip, that work is done. The transactions and resulting state are settled, and keeping every intermediate block in its original full form forever adds very little marginal security.

Under the currently proposed parameters, an epoch spans about 10 minutes and contains 40 blocks produced at 15-second intervals. During the live epoch, all 40 are propagated, validated, and used normally. That's the dense block flow that makes confirmations fast.

Here's what happens to those blocks as they age:

- While the epoch is live: every block is used in full, exactly as blocks are today.
- Once the epoch finalizes and the retention period passes: most of the 40 blocks are pruned down to a compact cryptographic skeleton, preserving the hash-chain structure and the essential commitments.
- One exception per epoch: a single block reward, chosen randomly, is permanently recorded in the network's master list of outputs (what we call the global output index), so it stays directly referenceable. The other 39 reward records never need that permanent entry.

Worth being precise here: ephemeral blocks are not ignored blocks, and pruning is not erasing the evidence of consensus. Every ephemeral block matters while the network is agreeing on the current tip. What changes later is the form in which that history has to be kept. The security of an old, finalized epoch doesn't depend on

every node storing all 40 original block bodies forever; it's carried by the finalized commitments, the retained chain skeleton, and the subset of historical data needed to verify the resulting state.

That asymmetry is the whole point. In real time, the network gets 40 block opportunities per 10-minute epoch. During historical sync, a new node doesn't process those as 40 complete, permanently indexed, reward-bearing blocks. Today, 10 minutes of history means roughly 10 full one-minute blocks. Under the proposed model, it means about one durable reward-index entry plus compact skeletal data for the rest. Raising the live block rate doesn't have to slow synchronization down at all; it may actually make sync faster than it is today, because old epochs become far cheaper to represent and process.

The chain ends up with a data lifecycle that matches how the data is used: detailed near the active tip, where detail matters, and compact for history the network has already settled. Users get fast confirmations, and node operators don't pay for them with permanently multiplying storage and sync costs. For anyone who wants the complete record anyway, an optional archival mode will store the block data that would otherwise be pruned, separately and safely, for future generations.

Lower Emission, Less Forced Selling

The last major consequence of removing proof-of-work is economic rather than technical.

Under the current hybrid model, new coins are created by both PoW mining and PoS staking, with a fixed reward of 1 ZANO per block split

between the two. The PoW side pays miners about 720 ZANO per day:

720 × 30 = 21,600 ZANO per month
720 × 365 = 262,800 ZANO per year

At a reference price of $9.25 per ZANO, the monthly figure is roughly $200,000.

Not every mined coin gets sold right away, and we're not claiming it does. But PoW mining carries operating costs that never stop: electricity, hardware, cooling, hosting, and maintenance. For most miners, selling a meaningful share of rewards is an operational necessity rather than a choice. Stakers sell too, sometimes, but they aren't servicing a monthly power bill. The incentive structures are simply different.

Once Zenith activates, PoW mining ends, and the PoW share of block rewards ends with it in its current form. What happens to that share of emission afterward is one of the parameters we're still finalizing: it could be retired, partly redirected into staking rewards, or restructured some other way. We'll publish the exact emission schedule once those numbers are settled, closer to activation.

To be equally clear about what this doesn't mean: it doesn't guarantee anything about price. Price depends on demand, liquidity, adoption, and broader conditions the protocol has no say in, so none of this should be read as a prediction. The narrow, concrete claim is about who earns new coins and why they sell. Today, the PoW share of issuance (roughly 21,600 ZANO a month) goes to miners, most of whom have to sell part of it to cover costs that never stop. After the

transition, whatever final shape the emission schedule takes, the network's security budget shifts toward participants who actually hold a stake in Zano, rather than leaking out through external electricity bills.

The End of a Long Road, and the Start of Another

Zenith closes a loop that opened in 2013 with a pure proof-of-work chain, ran through several generations of hybrid consensus, produced private staking with Zarcanum, and now lands on a privacy-preserving pure-PoS design.

We got here by a process, not a narrative. We tried to improve hybrid consensus but failed repeatedly, for structural reasons. So we changed the question, investigated multiple pure-PoS constructions, developed two of them into full research directions, and chose the one with the strongest practical path for Zano.

There's a long way between a research result and a live network. Consensus sits under everything else, and a change of this size will get the care it demands: rigorous implementation, testing, analysis, and review. But after years of investigation, the direction is set.

Zano's next era is pure proof of stake: private by design, fast enough for everyday payments, and built on twelve years of keeping real networks alive.

We'll be sharing much more as implementation progresses. Questions about Zenith? Come find us — every Zano channel is at links.zano.org.

Cheers, zAnons!

Sign up for more like this.

Enter your email Subscribe

Zano Monthly Project Update #21 - June 2026

HF6 preparations are finished and the activation height is locked in. Plus security hardening across Companion, Trade, the...

Gonbatfire
Jul 9, 2026 • 10 min read

The Countdown to Hard Fork 6 Has Begun!

After more than a year of work, one of Zano's most significant upgrades yet has arrived. Hard Fork 6 opens Zano up to the...

Quinten van Welzen
Jun 29, 2026 • 3 min read