PHAROS
RESEARCH

Speed Is Only the Baseline: Why Institutional Assets Really Choose the Pharos Public Blockchain

Table of Contents

Abstract

Taking Pharos as the main thread, and combining same-basis data from public dashboards with project disclosures, this article answers the three questions institutions and investors care about most: what Pharos's differentiated path is; what engineering and operational pain points it solves compared with mainstream high-performance chains; and how these technical choices translate into practical deliverability for payments and RWA.

As institutional capital moves into public-chain ecosystems, evaluation focus is shifting from peak TPS to verifiable finality, explainable security boundaries, deployable compliance and privacy capabilities, and the end-to-end cost structure for real-world assets and high-frequency workloads. Under a unified basis, this article visualizes and compares throughput, finality, decentralization indicators, and fee levels across representative networks; and, from a system-design perspective of parallel execution - storage pushdown - special processing networks (SPN), explains Pharos's full-stack parallel strategy and its institutional value.

Within this framework, we draw three conclusions. First, high-frequency payments and clearing depend more on the stability of confirmation-time distributions and failure semantics than on one-off benchmark peak TPS. Second, the long-run operability of RWA depends heavily on the data layer and the costs of indexing and reconciliation; storage and query paths often determine marginal on-chain costs. Third, compliance, privacy, and specialized computation should not remain as off-chain bolt-ons; they should be incorporated into a verifiable settlement closed loop, so that accountability boundaries can be contractualized, risks can be measured, and post-incident forensics can be evidenced.

Keywords: Pharos; Parallel Execution; Modularity; SPN; Finality; Fund Security; Institution-Grade Public Chain

01 / Introduction: What Are Institutions Really Buying on a Public Blockchain?

Over the past few years, stablecoins have become the base fuel for on-chain settlement, and the RWA narrative has evolved from “assets on chain” to “structured assets and cash-flow refactoring.” These demands impose engineering constraints that differ from pure DeFi: stronger reconciliation and auditability, lower final confirmation time, and more controllable data-layer costs. Pharos’s full-stack parallel route is a systematic answer proposed in this institutional context.

When institutions enter public-chain ecosystems, they are not buying a chain’s story; they are buying a portfolio of capabilities that business and risk-control systems can actually digest: deterministic settlement, auditable data access, controllable compliance boundaries, and predictable operating curves even under congestion and extreme market conditions. Concretely, they are buying an auditable, predictable, risk-manageable system: can it settle stably (Finality), can it carry peak load (Throughput), can it remain safe under extremes (Safety), and can it operate under real regulatory and commercial constraints (Compliance)?

The most typical misconception of the last few years has been treating high TPS as the decisive factor. For institutions, the real cost structure is often: on-chain fees are only the tip of the iceberg, while bigger spending comes from RPC, indexing, data availability, storage expansion, node operations, and security/compliance. In other words, institutions care about end-to-end throughput and end-to-end cost, not theoretical peak at a single layer.

1.1 Production constraints for institutions are not slogans

When institutions move assets or settlement flows onto a public chain, they care more about whether the system can run stably in production in an auditable, explainable, and replayable way. That introduces three hard constraints. First, audit and reconciliation constraints: “transaction completed” means internal books, custodian books, and on-chain books can be aligned and re-verified. Second, risk-exposure constraints: the more uncertain confirmation time is, the longer a buffer the business system must keep, and the more prefunding it must hold to cover uncertainty. Third, compliance and accountability-boundary constraints: which steps are guaranteed by the chain, which by external modules, and how to stop loss and produce evidence when something goes wrong - all must be expressible as terms and procedures.

The most typical trap is to treat high TPS as a silver bullet. Institutions’ real cost structure is often: chain fees are only a small part; the bigger portion comes from RPC, indexing, data availability, storage expansion, node operations, and security/compliance. Put differently, institutions want end-to-end throughput and end-to-end cost, not a theoretical peak of one layer.

1.2 Why end-to-end throughput/cost is the institutional KPI

If you treat a chain as a production system, TPS is only the surface output of the execution layer. Institutions actually consume the capabilities of a full pipeline: transaction generation, signing and broadcasting, ordering and consensus, execution and state persistence, RPC queries and indexing services, risk-policy write-back, and on-chain/off-chain book reconciliation. Any bottleneck in any step will reduce benchmark TPS into deliverable throughput. Likewise, institutional cost is inherently end-to-end: beyond on-chain fees, it includes high-availability RPC, indexing and historical backfill, audit retention, node operations and disaster recovery, custody and multisig, and compliance and security costs.

Pharos is worth discussing not because of a single-point performance number, but because it attempts to answer a more practical question: when payments, clearing, and real-world assets enter on-chain production, can execution performance, data access, and compliance capability be incorporated into one auditable, settleable responsibility chain - rather than being scattered across multiple off-chain components? In public materials, the project discloses that its testnet targets global payment scenarios, and highlights high throughput (e.g., 30,000 TPS) and about 1-second-level finality as goals/observations.[1]-[5] This article does not treat those disclosures as "facts already verified on chain." Instead, it treats them as Pharos's technical intent: what problems does it want to solve, and where is the key differentiation versus popular chains?

1.3 Understanding Pharos's system positioning

The key to understanding Pharos is to view it as a deliverable on-chain production pipeline: at the execution layer, a parallel framework improves stable throughput; at the storage layer, pushdown and structured queries reduce data and indexing costs; at the block stage, pipelining compresses confirmation-time jitter. When the business needs compliance, privacy, or specialized computation, SPN modularizes off-chain compute and returns to the mainnet to complete a verifiable settlement closed loop.

To avoid pseudo-comparisons caused by inconsistent definitions, this article uses Chainspect's unified basis for performance metrics (e.g., 1-hour-window TPS, Max TPS over 100 blocks, block time and finality), and DefiLlama's chain-level dashboard basis for assets and activity. Considering mindshare, ecosystem activity, and architectural diversity, we select Solana, Sui, Aptos, Sei, NEAR, Avalanche, and Ethereum as the main comparators. Ethereum serves as the security and decentralization baseline; the others represent mainstream high-performance and new-architecture directions. Metrics and sources include throughput (Real-time TPS / Max TPS), block time, finality, validator count, Nakamoto coefficient, and average transaction fees, primarily from Chainspect snapshots dated 2026-01-06.[6]-[12] Pharos is not yet fully tracked by the same dashboards; its performance points in charts are based on public disclosures.[4][5][13] Definitions are in the appendix.

02 / Panorama: The Distribution of Throughput - Finality - Decentralization

This chapter is not a "public-chain ranking." It answers six questions institutions most often ask when evaluating a new chain: can deterministic settlement support payments and clearing? Can the data layer and indexing costs remain controllable long-term? Can the security boundary be translated into risk-control clauses? Is it friendly to existing development and audit tooling? Does it reserve system-level interfaces for compliance, privacy, and specialized computation? And when scale arrives, can it still maintain a predictable operating curve? Accordingly, we place Pharos back into a coordinate system of performance × determinism × operability, use Solana, Sui, Aptos, and Ethereum as reference points, build an objective comparison under consistent bases, and then return to the engineering levers that explain "why."

Figure 1: Main Performance and Decentralization Metrics (snapshot: January 6, 2026)

	Real-time TPS(1H)	Max TPS(100 blocks)	Block Time(1H)	Finality/TTF	Validators / Nakamoto
Pharos (Test Network)	seventy-four point eight eight	30,000	~0.5s	⩽1s	32 / 11
Solana	1,042	5,289	0.4s	12.8s	796 / 19
Sui	eighty-eight point two seven	nine hundred and twenty-six point five	0.08s	<1s (Chaincheck shows 0)	126 / 19
Aptos	forty-four point five four	12,933	0.06s	<1s (Chaincheck shows 0)	138 / 17
Sei	seventy-seven point two five	two hundred and fifty-five point seven	0.51s	<1s (Chaincheck shows 0)	40 / 7
NEAR	twenty point four five	4,135	0.6s	0.6s	383 / 10
Avalanche	twenty-three point one five	one hundred and twenty-two point seven	1.39s	2s	727 / 29
Ethereum	twenty-three point six six	sixty-two point eight seven	12s	12m48s	988,900 / 2

Source: Pharos Research

(Note: Pharos indicators are from public testnet/material disclosures; other chains are from Chainspect page snapshots.)

As shown in Figure 1, performance, finality, and decentralization are placed within an institution-readable framework: throughput sets the upper bound, finality determines settlement certainty, and validator count and the Nakamoto coefficient more directly indicate whether the security boundary is acceptable to risk control. For Pharos, the point of a cross-chain comparison is not to win any single metric, but to demonstrate that its full-stack parallelization approach can simultaneously improve end-to-end experience and operational sustainability.

2.1 Throughput and finality

At a macro level, under the same-basis Max TPS (100 blocks) and finality (TTF) dimensions, how do mainstream chains distribute performance and trade-offs?

Figure 2. Throughput and finality comparison.

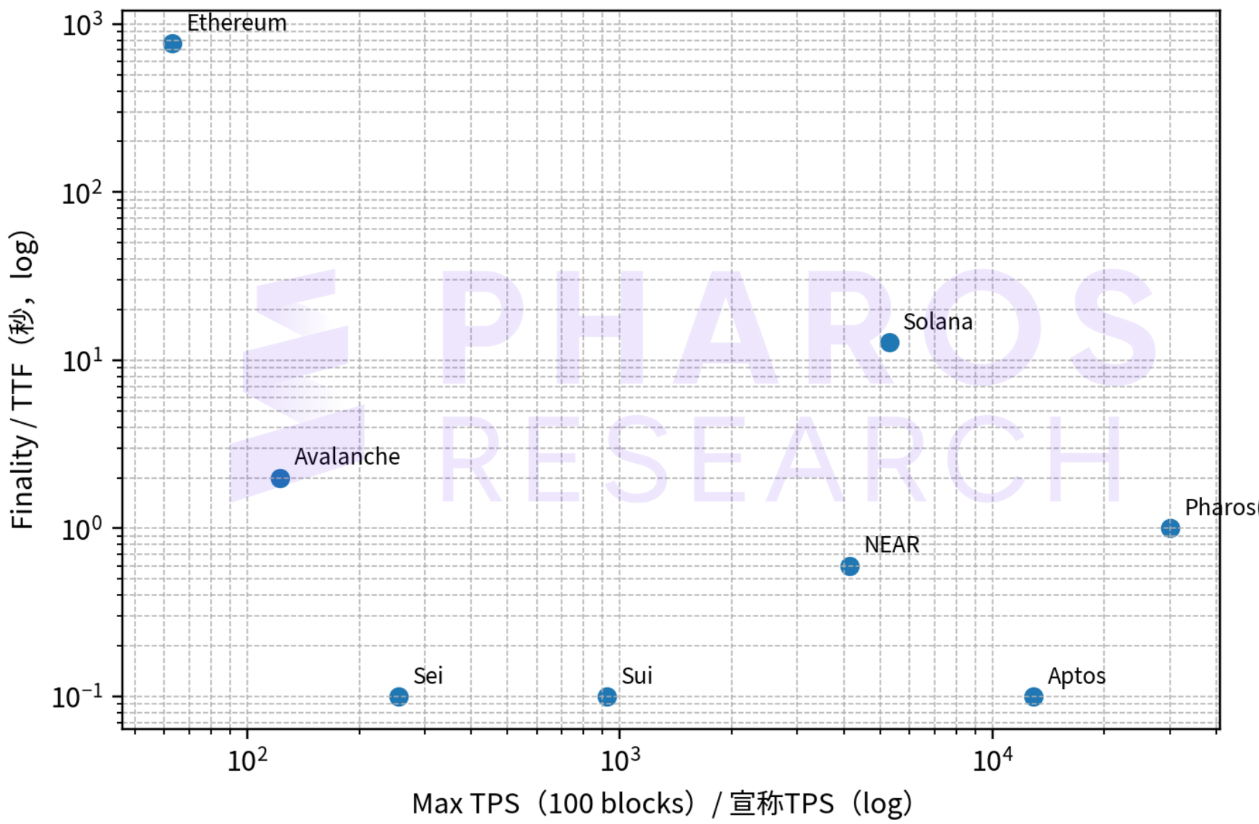

Source: Pharos Research

From an institutional perspective, throughput × finality is a map with more explanatory power than a TPS leaderboard: throughput determines the upper bound, while finality determines how fast business exposure can be closed. In practice, the predictability of confirmation time under congestion, jitter, and extreme markets often matters more than single-point performance numbers. Therefore, the value of Figure 2 is to bring different chains' performance narratives back to one question: where does the performance advantage come from, and can deterministic experience be maintained under high load?

As shown in Figure 2, mainstream high-performance chains each emphasize different points on the high-throughput/low-finality frontier, but few chains can remove bottlenecks simultaneously across consensus, execution, storage, and data access. Pharos's differentiation is built around this: using full-stack parallelism (execution parallelism + storage pushdown + pipelined stage parallelism) to trade for more stable end-to-end experience, not merely a single benchmark peak.

Putting throughput and finality on the same chart avoids the most common mistake in chain comparisons: staring only at TPS. Deterministic completion is often more critical than peak throughput, especially in payment clearing, collateral management, and RWA operations, where business systems need predictable confirmation time and controllable failure rates. If, after mainnet launch, Pharos can stably sustain the disclosed near-second-level finality together with high throughput, it would be better suited to high-frequency workloads.[4][5][12][6]

2.2 Decentralization indicators: validators and the Nakamoto coefficient

Institutional investors often use validator count and the Nakamoto coefficient as a first screening layer.

Figure 3. Validator scale comparison (log scale; Pharos is testnet data).

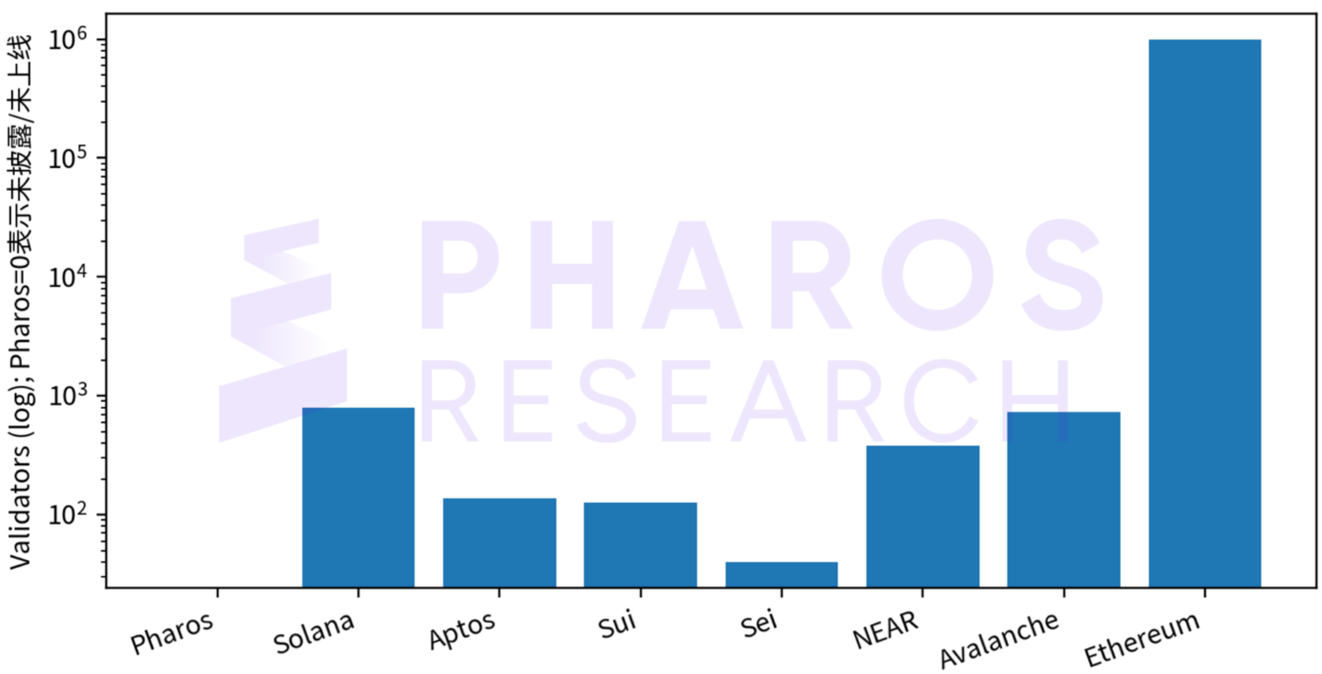

Source: Pharos Research

As Figure 3 suggests, mature networks often differ by orders of magnitude in validator scale. For a network like Pharos that is still expanding, institutions should focus more on route and transparency: whether node admission and incentives are clear, whether operation and failure data are continuously disclosed, and whether institution-facing infrastructure partners and custody/risk-control solutions exist - rather than one-voting the chain out early by absolute scale.

Validator count is unavoidable in institutional due diligence, but it should not be read simplistically as “more equals safer.” A more accurate understanding is that validator scale correlates strongly with censorship resistance, resilience to single points of failure, and ecosystem sustainability - while it must also be evaluated together with node independence, stake concentration, and client diversity.

What institutions truly need is to translate chain-level risk into executable risk-control actions. In practice, validator scale and the Nakamoto coefficient often feed into three mechanisms: admission clauses (higher exposures only after thresholds), limits and gradual ramp-up (scale increases after milestones), and disaster recovery and switching (degrade paths when confirmation distributions deteriorate or key infrastructure fails). This is why route and transparency matter more for early networks: clear node admission and incentives, continuous disclosure of operation and failure data, and accessible infrastructure partners and custody/risk-control solutions.

Figure 4. Nakamoto coefficient comparison (Pharos = 0 indicates not disclosed / not launched).

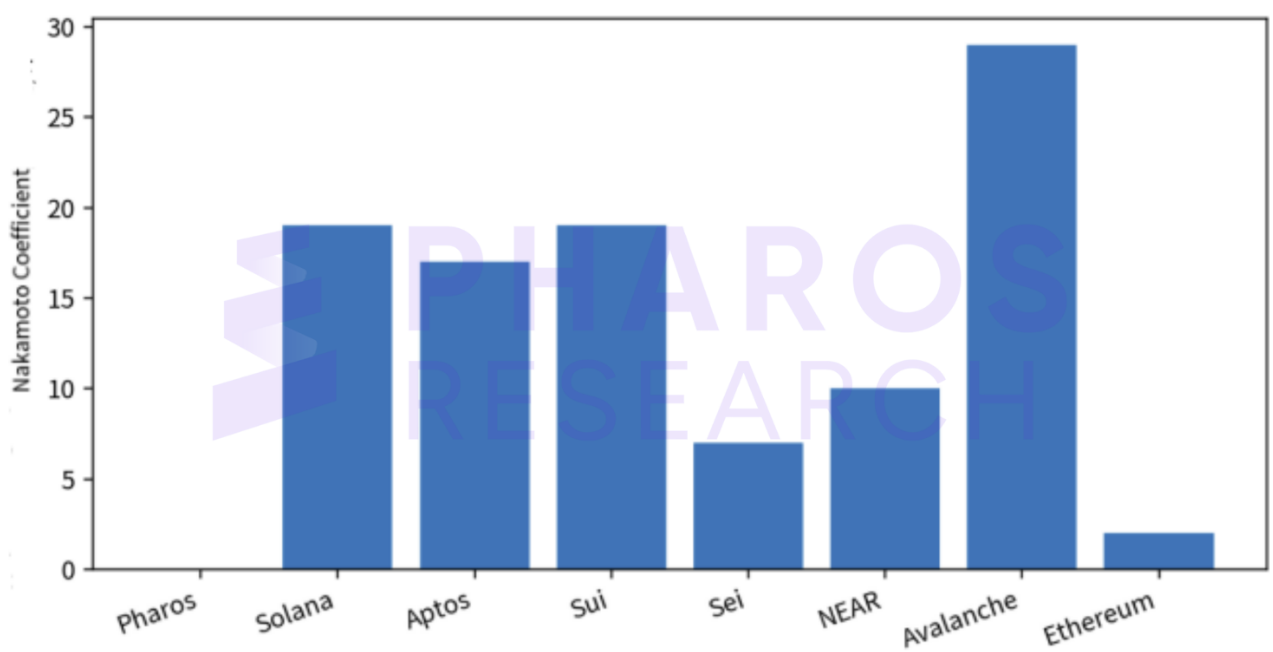

Source: Pharos Research

As Figure 4 shows, concentration risk differs significantly across networks. For Pharos, a near-term priority is to turn the Nakamoto coefficient into a trackable, disclosable, improvable engineering and governance indicator - including stake distribution, geographic distribution, operator diversity, and redundancy across key implementations. Institutions can also manage early concentration risk by "limits + gradual ramp-up," translating it into controllable exposure.

Compared with validator count, the Nakamoto coefficient is closer to the core question: how many independent entities must collude (or fail simultaneously) to affect the network's security boundary? It is a readable proxy for concentration risk. In the Chainspect snapshot, Solana is about 796/19; Avalanche about 727/29; Sui 126/19; Aptos 138/17; Sei 40/7; NEAR 383/10; Ethereum 988,900/2.[6]-[12] Pharos's mainnet validator scale and stake distribution will become key variables institutions need to track continuously.[13]

2.3 Fee levels: low gas is not low total cost

Average transaction fees are the most visible user cost, but institutions care more about end-to-end TCO. Figure 5 gives an order-of-magnitude comparison.

Figure 5. Average transaction fee comparison (Pharos is testnet data).

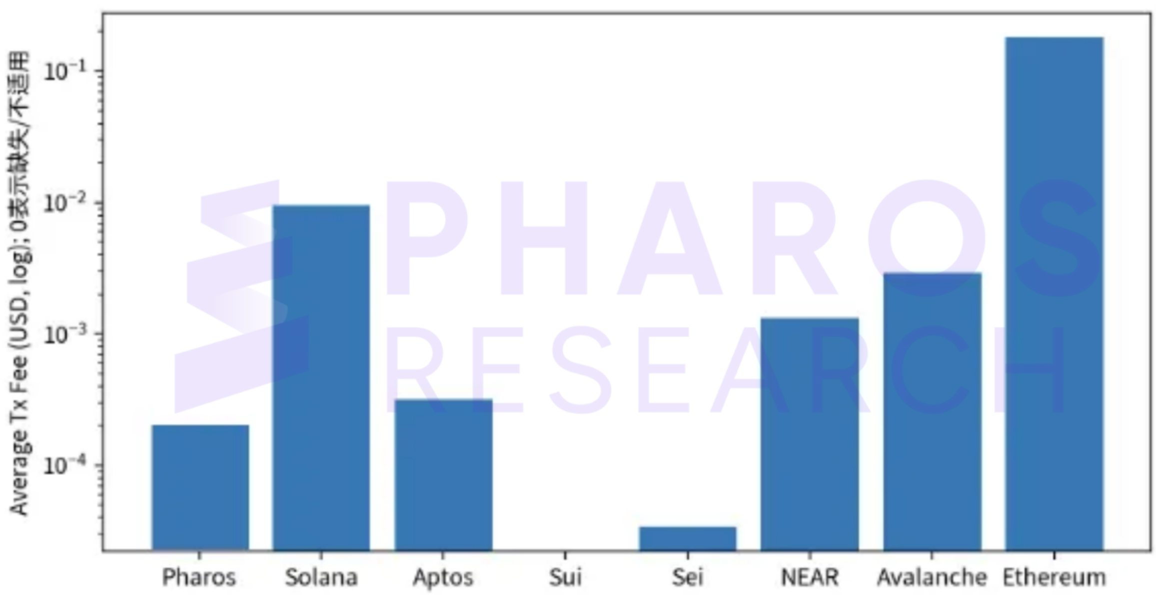

Source: Pharos Research

As Figure 5 indicates, average fees can differ widely across chains, but that does not directly imply which chain is better for institutions. Pharos emphasizes storage and pipelining because, once workloads become high-frequency and data-intensive, I/O, indexing, and proof generation become the real cost levers: they determine node hardware thresholds, audit and reconciliation difficulty, and long-term integration burden for institutional systems.

The most common trap in cost discussions is to treat gas price as total cost. In real operations, total cost also includes hidden losses from failures and retries, opportunity costs and slippage under congestion, and the engineering costs of on-chain data access and reconciliation. In payments and RWA, the most underestimated portion is often data access cost: balance proofs, transaction backfills, risk tracking, compliance reporting, and audit sampling all rely on stable RPC and indexing services. When state grows fast, query paths are heavy, and indexing updates are expensive, institutions pay higher infrastructure bills and carry more complex engineering maintenance beyond fees. Therefore, Figure 5 should be read only as a slice of write cost; the meaning of Pharos's storage-and-pipeline emphasis is to move long-term cost leverage from parameter subsidies to structural optimization of data paths.

2.4 Fund security: what institutions really fear is accountability boundaries and failure modes

If you interpret "fund security" only as hacker attacks, you underestimate the risks institutions actually face. They care more about failure modes: when a chain experiences congestion, halts, forks, client bugs, ordering attacks, or cross-domain module failures, do assets enter an uncontrollable state? Are accountability boundaries clear? Can loss be stopped quickly and audit evidence produced?

At the chain level, decentralization and concentration indicators determine whether the system is easy to influence by a small set of entities or prone to systemic stoppage; at the transaction level, the finality distribution determines when clearing exposure truly closes; at the operations level, data-layer and indexing costs determine whether institutions can sustain low-cost on-chain reconciliation and anomaly tracking. Pharos differs in that it attempts to make "fund security" a deliverable system capability with explicit interface boundaries: in many chain practices, compliance checks, privacy compute, risk rules, and even parts of indexing/monitoring are implemented via external services or application-side components - fast to ship, but the chain only guarantees ledger consistency while off-chain services guarantee business correctness. Once disputes arise (why accepted/rejected, why settlement deviated), institutions cannot form an on-chain verifiable evidence loop, and audit/compliance proof costs rise sharply with scale.

Pharos's proposed answer, via a mainnet + Store + SPN combination, is: on the settlement side, provide final/terminal states and clear failure semantics callable by risk controls; on the data side, provide low-cost, replayable query and backfill capabilities; on the extension side, modularize and isolate compliance/KYC/AML, privacy computation, and specialized risk modules via SPN, write back verifiable decisions/proofs to mainnet, and provide downgrade/stop-loss paths when modules malfunction - thereby converting failure modes into process-manageable risk.

In practice, institutional fund security typically becomes four families of controls: permissions, settlement, observability, and stop-loss. Multisig and tiered authorization control fund actions; confirmation strategies and limits control exposure; continuous monitoring of confirmation distributions, failure rates, and key infrastructure availability; and clear degrade paths (pause, raise thresholds, switch channels) that keep books reconcilable and audits replayable. By combining full-stack parallelism and SPN modularization, Pharos aims to turn the permissions-settlement-observability-stop-loss loop into configurable, observable, auditable infrastructure capability, and to convert it into verifiable production-grade determinism through continuous disclosure of runtime data - thereby differentiating itself as an institution-grade base layer.

03 / Technical Differentiation: Three Institution-Grade Infrastructure Levers

Pharos's route looks like designing for financial-grade infrastructure: combining consensus/execution, storage, and heterogeneous computation into a high-efficiency production pipeline usable by institutions.

3.1 Lever 1: chain-level parallel execution (Pharos VM: dual standards of EVM + WASM)

Pharos VM emphasizes supporting both EVM and WASM standards, aiming to converge the Solidity/EVM ecosystem with high-performance WASM ecosystems on one chain, while raising throughput via a parallel execution framework.[2] Parallel execution typically consists of conflict detection, parallel grouping, minimizing rollbacks/re-execution, and deterministic commit. Only by turning the parallel strategy into an evolvable engineering module and exposing clear constraints and tooling to developers can performance gains reliably translate into business experience. The value of dual standards is not only lower migration cost; it also means performance-sensitive workloads (matching, clearing, pricing, risk control) can progressively move down into a closer-to-the-metal execution environment without abandoning the assets and tooling of EVM.

The hard part of parallel execution is not "more threads," but ensuring different nodes produce consistent results over the same transaction set without sacrificing determinism and reproducibility. Institutions care about three production questions: when conflict density rises, does the system degrade gracefully or fall into rollback/replay storms; are failure semantics stable and capturable by business systems for automation; and can audit replay reproduce the same execution trajectory given deterministic inputs and block order. Therefore, in mainnet and pilots, Pharos should continuously disclose throughput curves under different conflict ratios, stability of failure/retry rates under load, and end-to-end confirmation time quantiles.

Figure 6. Core levers of parallel execution.

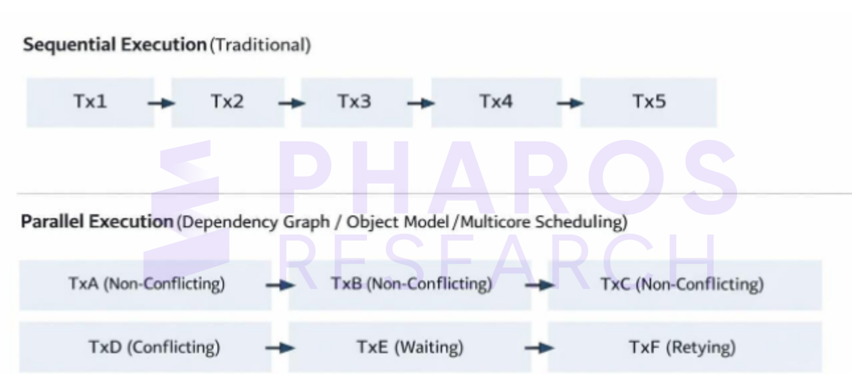

Source: Pharos Research

As Figure 6 illustrates, Pharos abstracts parallel execution into a scheduling/execution split, and offers compatibility and extensibility via Pharos VM (EVM + WASM): EVM lowers migration friction, WASM reserves interfaces for higher performance and broader language ecosystems. This means the chain can carry existing Solidity assets and audit experience, and introduce higher-performance execution paths when needed without replacing the entire stack.

3.2 Lever 2: storage pushdown (Pharos Store)

Storage and data costs are often the long-term bottleneck of high-performance chains. Pharos Store's idea is to push down some work traditionally handled by external databases into the storage layer itself (ADS/pushdown), achieving a better global optimum between throughput and cost. Public docs disclose that, under certain benchmarks, it can deliver a 15.8× throughput improvement and save 80.3% cost versus a traditional two-layer architecture.[3] These are more tangible to institutions because they directly target long-run operating expenses.

Figure 7. Publicly disclosed benefits of Pharos Store.

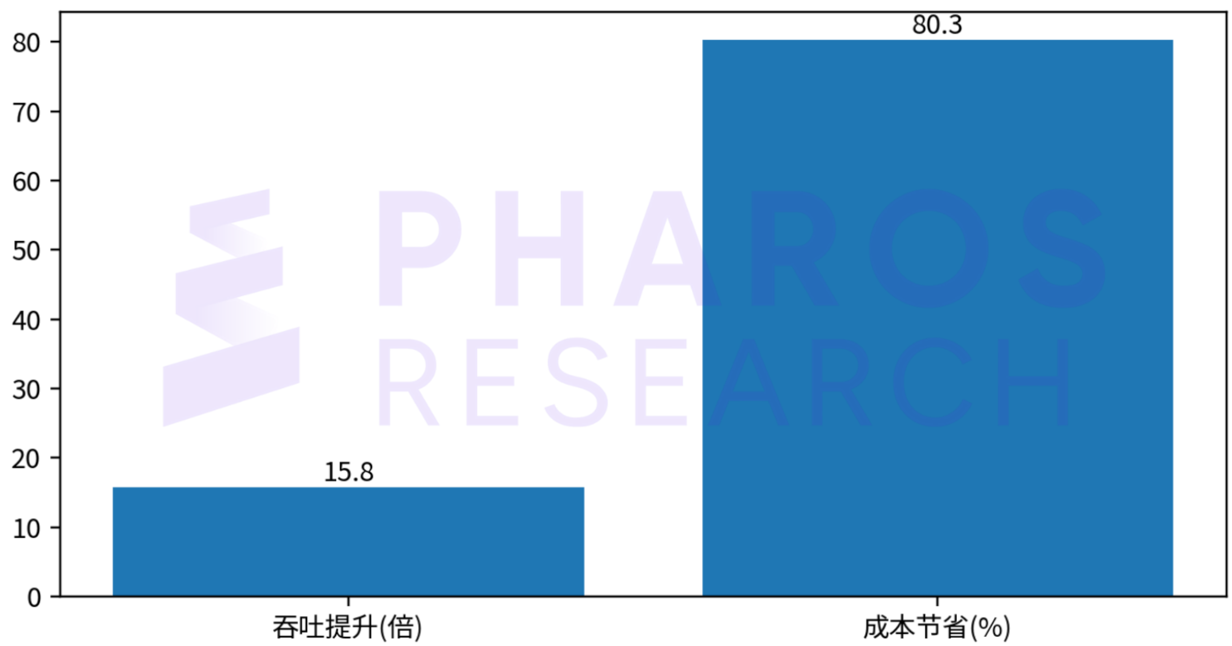

Source: Pharos Research

As Figure 7 shows, Pharos Store frames benefits as engineering metrics. The significance is not only benchmarks but long-run operations: more controllable node costs, more manageable state growth, and more efficient queries and reconciliation. Such data-layer changes often matter more than single-point TPS in determining whether a system is worth running in production.

Underestimating storage as “just persistence” ignores its decisive role in performance and operating cost. In high-frequency workloads, state updates can trigger read/write amplification, Merkle/proof-structure updates, and incremental index maintenance. The more finance-like a workload is, the more frequent reconciliation and audit sampling becomes, making such costs impossible to ignore. The value of Pharos Store should be translated into verifiable outputs: state growth curves, query latency distributions, index update time, and throughput degradation ratios under audit backfill scenarios.

Many chains’ bottlenecks are not CPU but storage read/write amplification: state writes, historical-version access, Merkleization, and index updates can quickly become the throughput ceiling under load. Especially for RWA, stronger traceability and more complex query/reconciliation are required, so storage and indexing engineering quality directly determines institutional operability.

3.3 Lever 3: SPN (Special Processing Networks)

In real business, zero-knowledge proofs, privacy computation, and AI inference often require heterogeneous hardware or specialized runtimes. Pharos proposes SPNs to connect such specialized computation as networked modules and then return to mainnet for settlement and verifiability.[1] The potential institutional value is to turn “off-chain compute” from a black box into a governable, composable, auditable module: compute is performed in a special network; results/proofs return to mainnet.

The key is not “more modules,” but bringing off-chain compute into a governable boundary:

(1) verifiable, with results written back as proofs or verifiable claims; (2) isolatable, with different business domains running in dedicated networks to reduce friction; (3) stop-loss capable, with explicit degrade strategies when modules malfunction (pause, raise thresholds, switch modules), without affecting core ledger security. This is the boundary Figure 8 should clarify: storage belongs to mainnet core; SPN is an external special-processing network that closes the loop via calls/settlement and proof/result write-back.

Figure 8. Pharos’s modular mainnet + SPN full-stack parallel approach.

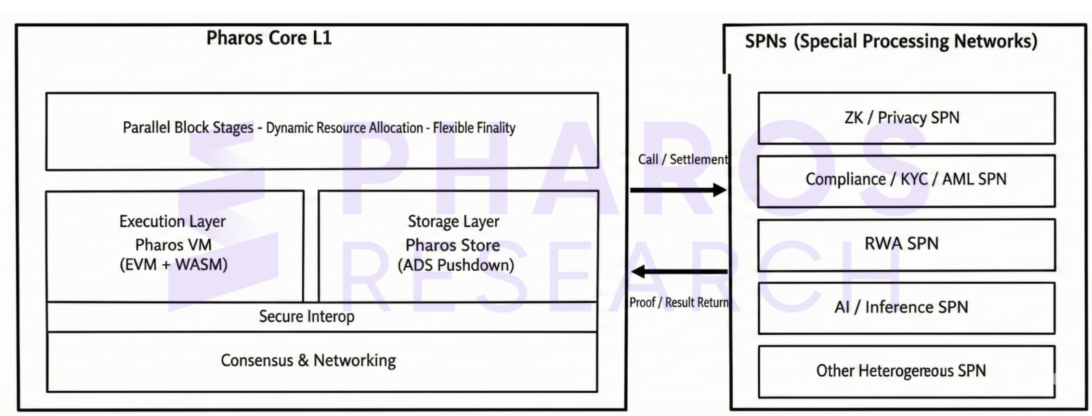

Source: Pharos Research

As Figure 8 shows, Pharos uses a “modular mainnet + SPN” approach: on top of shared security and a unified value network, it provides customizable, isolated special processing networks for different business domains and incorporates cross-domain interoperability by design.

3.4 SPN trust model and operating assumptions

Beyond whether SPN “can do it,” who runs SPN and whether its trust assumptions are explicit and controllable determines whether it can enter production risk frameworks. Unlike mainnet consensus, SPN is not assumed to be a fully trustless public execution environment; it is a constrained, governable special compute domain.

Based on currently disclosed information, SPN is closer to a modular-access special network: its node operators, consensus mechanism, and security assumptions can differ by business domain, rather than being forced to reuse the mainnet validator set. In practice, SPN can be operated by: (1) infrastructure operators within the Pharos ecosystem; (2) compliant third-party service providers or institutional consortia; (3) in specific scenarios, an application-side or custodian-designated node set. Each mode has explicit trade-offs among decentralization, performance, and compliance controllability.

Importantly, SPN’s security boundary is not equal to the mainnet’s. Its goal is not to eliminate trust assumptions, but to make them explicit and modular, and isolate them via verifiable outputs and mainnet settlement. Specifically, SPN is responsible for whether computation follows rules and completes; its output is submitted as a proof or verifiable claim; mainnet is responsible only for whether to accept the output into settlement state, and does not inherit SPN’s internal trust risks.

The key value of this layering is stop-loss: if SPN nodes misbehave, fail, or trigger compliance disputes, mainnet can pause acceptance, raise verification thresholds, or switch to backup modules to limit risk spillover, without affecting the core ledger’s consistency and asset safety. For institutions, this means SPN risk can be treated as a configurable, replaceable business-risk module rather than an uncontrollable chain-level systemic risk.

Therefore, in institutional deployments, SPN is better positioned as a governable special compute network within risk controls: operator identities, incentives/penalties, audit interfaces, and incident-handling processes should be explicitly written into business terms and technical appendices.

3.5 A capability matrix: Pharos’s native modular advantage

Pharos’s differentiation resembles native modularity: parallel execution, EVM compatibility, storage pushdown, and SPN are unified into one system from the network’s initial design stage.

Figure 9. Key capability coverage matrix (qualitative synthesis).

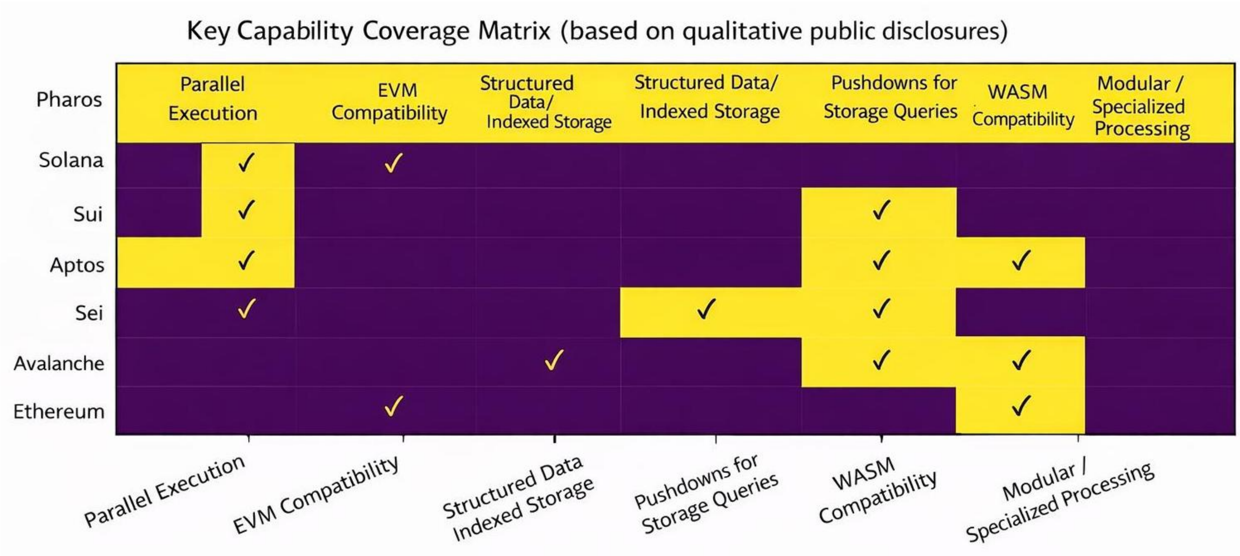

Source: Pharos Research

As Figure 9 indicates, when key capabilities are placed into one matrix, Pharos's claim can be summarized as a native modular advantage: within a single architecture, it advances execution parallelism, storage optimization, pipelining, and an SPN extension framework from the start. In institutional contexts, native integration has concrete value: less external "stitching," clearer accountability boundaries, and more controllable rollout cadence.

Native modularity is a governance and rollout constraint: each additional external component adds supplier risk, another audit object, another failure link, and more space for accountability disputes. When the base chain itself integrates execution, storage, and extension modules into one system, institutions can complete a more end-to-end on-chain production loop with fewer external parts, keeping key risks within an auditable boundary.

As evaluation moves from "performance" to "institutional deployment," dimensions expand quickly: compliance, privacy, identity, data availability, cross-domain interoperability, developer tooling, audit and risk interfaces, and more, can all become real blockers.

04 / Scenario Deployment: Why These Capabilities Matter to Institutions

Only by mapping the engineering levers back to deployable institutional workflows can Pharos's real target be seen: combining parallel execution, storage/indexing, and modular extension networks into an operable settlement-and-asset system. In short, institutions do not lack faster chains; they lack more deterministic, more reconcilable, more controllable on-chain workflows.

4.1 Payments and clearing: turning near-1-second finality into operable settlement experience

The core of payments and clearing is when risk exposure closes. The slower and more jittery finality is, the longer buffers and the higher prefunding institutions must hold, and the more complex reconciliation must be to cover uncertainty. Figure 2 shows that mainstream networks differ significantly across throughput and finality; for payments, the real KPI is often end-to-end confirmation time distribution (P50/P95/P99) and congestion-period failure/retry semantics.

For payments and clearing, the key is whether exposure closes quickly and whether confirmation time remains predictable under high load. When confirmation distributions are stable and failure/retry semantics are clear, business systems can shorten buffer cycles, reduce prefunding occupation, and incorporate on-chain settlement into automated operations.

Operationally, an institution-grade payment/clearing transaction usually includes: instruction generation (compliance checks and limits) → signing and broadcasting → on-chain confirmation → internal/external book alignment → exception handling (failure, delay, duplication, reversal). The chain's value is not only faster confirmation but more controllable exceptions: are failures explainable? are delays predictable? are retries idempotent? These determine whether institutions can integrate on-chain transactions into automated operations. In pilots, acceptance metrics should be defined at the operational level: confirmation time quantiles, failure/retry rates and stable error types, fee and availability curves under congestion, and recovery + reconciliation replayability.

4.2 RWA: compliance and privacy are product specifications, not add-ons

The hard part of RWA is transferring implicit institutional structures of traditional finance onto chain: who is eligible to hold, how cash flows are distributed, how disclosure and audits are done, and how sensitive information is minimized under compliance. Many projects treat compliance and privacy as "external services," causing trust assumptions to spill outside and system boundaries to blur.

In RWA, the key is not whether compliance/privacy is "supported," but whether the results can be incorporated into a settlement-and-audit closed loop. Only when compliance checks and risk judgments have verifiable results and are directly tied to asset settlement states can institutions explain every approval, rejection, and exception in audits. This makes compliance and privacy not procedural promises but auditable system behavior. Institutions should focus on SPN's trust model (who runs it, how penalties work, how faults are isolated), the verifiability boundary of proof systems, and stop-loss paths for mainnet assets/business when SPN fails.

Taking AquaFlux as an example, it decomposes underlying cash flows and risk tranches into multiple token forms, enabling composable strategies on the same underlying asset for different risk appetites.[26] Such structured models imply more complex state transitions, more frequent compliance checks, and heavier reconciliation needs. This is why Pharos combines Store, pipelining, and SPN into one storyline: it targets long-run on-chain operation of structured assets, not one-off issuance.

RWA complexity comes from lifecycle operations: compliance admission, cash-flow distribution, disclosures, exception handling, and audit evidence. The more complex the structured asset, the more frequent state updates are, the more query dimensions exist, and the denser compliance checks become. Here, Pharos Store's value is not only faster writes but lower total cost of sustained queries and reconciliation; SPN's value is not only "has KYC/privacy," but incorporating compliance/privacy computation results into a verifiable closed loop so institutions can explain "why approved/why rejected" in audits.

4.3 Trading and market making: parallel execution is the start; data layer and failure semantics are the key

Trading and market-making workloads usually demand simultaneously: high concurrency, low latency, predictable failure/retry semantics, and controllable data/indexing costs. Parallel execution can raise throughput, but as transaction density rises and state conflicts increase, what shapes market-making experience is often conflict detection/rollback strategy and the cost of state I/O and index updates.

Pharos's differentiation is to place the data layer into the core narrative: via storage pushdown and parallelized ADS/Merkle and indexing paths, it attempts to reduce read/write amplification and reconciliation burden under high-frequency workloads. For market-making and risk systems, such optimization directly impacts query latency, historical-state access cost, and recovery time in extreme markets. Institutional pilots should include these engineering indicators in acceptance, not only TPS.

Two stress-test types are suggested to validate real gains from parallel execution + data layer. First, hot-spot stress (concentrated access to the same state) to observe throughput degradation and stability of failure semantics. Second, query stress (high write plus heavy historical queries and index updates) to observe query-latency distributions and indexing cost curves. Only when these tests are repeatable and stable does the chain have the baseline to carry near-production trading systems.

In the end, the deployment test is simple: can Pharos convert verifiable parallelism into deliverable institution-grade product capability, including stable confirmation times, auditable data access, pluggable compliance/privacy modules, and clear operational and risk interfaces? If these are measurable in real workloads, network effects and partnership expansion can be ignited.

05 / Conclusion: The Key to Evaluating Pharos Is Verifying a Growth Curve

If competition among high-performance chains is framed as “processing transactions faster,” discussions easily fall into a TPS arms race. But from an institutional perspective, the real competition is settlement authority: who can, under compliance and audit constraints, turn stablecoin payments, clearing reconciliation, and RWA structured-asset workflows into a sustainably running system. Accordingly, Pharos’s differentiation should not be reduced to “faster,” but understood as institutional system engineering: its full-stack parallel route must ultimately be tested under real business loads for deterministic settlement, data operability, and accountable boundaries.

Narratives stand or fall not by how well they are described, but by whether they can be verified, reproduced, and operated. Therefore, the first set of key points for Pharos is end-to-end determinism: under different loads and congestions, are confirmation-time distributions (P50/P95/P99), failure rates, and retry semantics stable and predictable; under network jitter or extreme markets, can the system recover and reconcile in a clear, explainable way. Payments and clearing do not go live because a chain is fast on average; institutions need “controllable under bad cases.”

The second set of key points is the data layer and long-run TCO. High-frequency and data-intensive business eventually pushes cost pressure toward state storage, indexing, and historical access. By putting Store in its core narrative, Pharos is betting on moving bottlenecks forward and optimizing structurally, lowering node hardware thresholds and reducing data-access and audit-reconciliation burdens. Institutional evaluation should write data-layer metrics into acceptance: state growth curves, query latency, index update costs, and reconciliation complexity under real workloads - not just demo throughput.

The third set of key points is the accountability boundary of modular extensions. SPN’s value is isolating compliance checks, privacy computation, and specialized risk logic from the general execution path, while writing back “results” to mainnet settlement as proofs or verifiable claims.

At the same time, full-stack parallel routes can introduce new engineering and governance trade-offs while improving end-to-end performance and operability. More complex execution and data paths can raise hardware and ops thresholds for nodes, creating decentralization pressure in early phases; and multi-layer coordinated architectures raise the bar for stability, fault isolation, and replayability under extreme loads or upgrades. Evaluating Pharos should therefore consider node admission, runtime-data disclosure, and incident-handling paths to judge whether these risks can be quantified and managed.

For institutions and investors, the most robust path is metrics-driven limited pilots → staged acceptance → gradual ramp-up. As long as Pharos can continuously deliver three things - predictable deterministic settlement, controllable data-layer cost, and auditable/isolatable compliance-module boundaries - ecosystem expansion and partner growth can enter a sustainable positive loop, converting early growth dividends into a controllable return curve.

Appendix: Glossary

- Finality / TTF: the time scale at which a transaction is considered irreversible; used by institutions to estimate clearing exposure.
- Real-time TPS: real-time throughput per unit time (basis here is a 1-hour window).
- Max TPS (100 blocks): peak throughput over a recent ~100-block window.
- Nakamoto Coefficient: metric for the minimum number of entities required to influence consensus (basis depends on methodology).
- Parallel Execution: raising throughput by reducing global locks and conflict amplification.
- EVM: Ethereum Virtual Machine; the most mature smart-contract and tooling ecosystem.
- WASM: WebAssembly; a general high-performance execution environment supporting multi-language stacks.
- Modular Stack: a modular stack that composes execution, storage, and compute modules.
- Pharos VM: Pharos execution layer emphasizing EVM+WASM and a parallel execution framework.
- Pharos Store: Pharos storage layer emphasizing ADS/pushdown to reduce data and indexing costs.
- ADS: Augmented Data Structures; enhanced data structures for query and indexing efficiency.
- Pushdown: pushing filtering/aggregation logic closer to the data to reduce total cost.
- SPN: Special Processing Networks; special networks for heterogeneous compute such as ZK/privacy/AI.
- ZK: zero-knowledge proofs; proving a statement without revealing underlying information.
- MEV: maximal (miner/validator) extractable value; related to ordering, censorship, and transaction costs.
- Subnet: Avalanche subnets; isolated performance and governance domains for different businesses.
- Sharding: scaling by partitioning state and transactions (e.g., NEAR's route).
- Restaking: reusing existing security capital for new networks (materials mention EigenLayer/Babylon).
- TCO: total cost of ownership; comprehensive cost across fees + infrastructure + data + security + compliance.

References

[1] Pharos Docs – Pharos Modular Stack https://docs.pharosnetwork.xyz/architecture/pharos-modular-stack
[2] Pharos Docs – Pharos VM (EVM + WASM) https://docs.pharosnetwork.xyz/core-technologies/pharos-execution/pharos-vm
[3] Pharos Docs – Pharos Store (ADS pushdown) https://docs.pharosnetwork.xyz/core-technologies/pharos-store
[4] PRNewswire – Pharos launches Testnet (includes 30,000 TPS, ~1s finality claims) https://www.prnewswire.com/news-releases/pharos-launches-testnet-accelerating-global-payments-with-eigenlayer-and-babylon-restaking-302602484.html
[5] Pharos Blog – Pharos launches Testnet (performance and positioning) https://pharosnetwork.xyz/blog/pharos-launches-testnet-accelerating-global-payments-with-eigenlayer-and-babylon-restaking
[6] Chainspect – Solana metrics dashboard https://chainspect.app/chain/solana
[7] Chainspect – Sui metrics dashboard https://chainspect.app/chain/sui
[8] Chainspect – Aptos metrics dashboard https://chainspect.app/chain/aptos
[9] Chainspect – Sei metrics dashboard https://chainspect.app/chain/sei
[10] Chainspect – NEAR metrics dashboard https://chainspect.app/chain/near
[11] Chainspect – Avalanche metrics dashboard https://chainspect.app/chain/avalanche
[12] Chainspect – Ethereum metrics dashboard https://chainspect.app/chain/ethereum
[13] Chainspect – Pharos (listing/placeholder) https://chainspect.app/chain/pharos
[14] Foresight News – Pharos testnet news (30,000 TPS, 1s finality mention) https://foresightnews.pro/news/detail/39840
[15] The Block – Pharos seed round news (background) https://www.theblock.co/post/277212/pharos-raises-8-million-seed-round-led-by-lightspeed-faction-hack-vc
[16] Solana Docs – Architecture / runtime overview (background) https://docs.solana.com/
[17] Sui Docs – Sui architecture and object model (background) https://docs.sui.io/
[18] Aptos Docs – Aptos architecture and Move/parallel execution (background) https://aptos.dev/
[19] Sei Docs – Sei network overview (background) https://docs.sei.io/
[20] NEAR Docs – NEAR protocol overview (background) https://docs.near.org/
[21] Avalanche Docs – Avalanche consensus and subnets (background) https://docs.avax.network/
[22] Ethereum.org – Ethereum proof-of-stake and finality (background) https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/
[23] EigenLayer – Restaking overview (background for Pharos testnet positioning) https://docs.eigenlayer.xyz/
[24] Babylon – Bitcoin staking / restaking overview (background) https://babylonchain.io/
[25] MetaEra/Paragraph – AquaFlux on Pharos (reference reading) https://paragraph.com/@metera/aquaflux-zh
[26] AquaFlux Launches on Pharos Testnet, Advancing Structured RWA https://paragraph.com/@menews/aquaflux-launches-on-pharos-testnet-advancing-structured-rwa

Contributors

Authors: Huijie Tang (𝕏@web3sensen）

Reviewers: Colin Su, Grace Gui, NingNing, Owen Chen

Design: Alita Li

Disclaimer

Contact

Pharos Network is a next-generation public blockchain for Real-World Assets (RWA) and stablecoins, focused on asset tokenization and on-chain circulation. We connect traditional institutions with the Web3 ecosystem, enrich the types of on-chain assets, expand revenue sources, and meet the allocation needs of a broader range of investors. Meanwhile, we help traditional enterprises unlock sustainable value on-chain through customized solutions. Boasting profound professional expertise and top-tier technical capabilities, our team builds a secure, efficient, and scalable infrastructure, providing institutions with a comprehensive decentralized ecosystem for onboarding assets onto the blockchain. We welcome strategic partners with a long-term perspective to co-build an open, compliant, and sustainable RWA ecosystem. For industry exchanges with us, please contact: chris@pharoslabs.xyz

Pharos' Official Website: https://www.pharosnetwork.xyz/

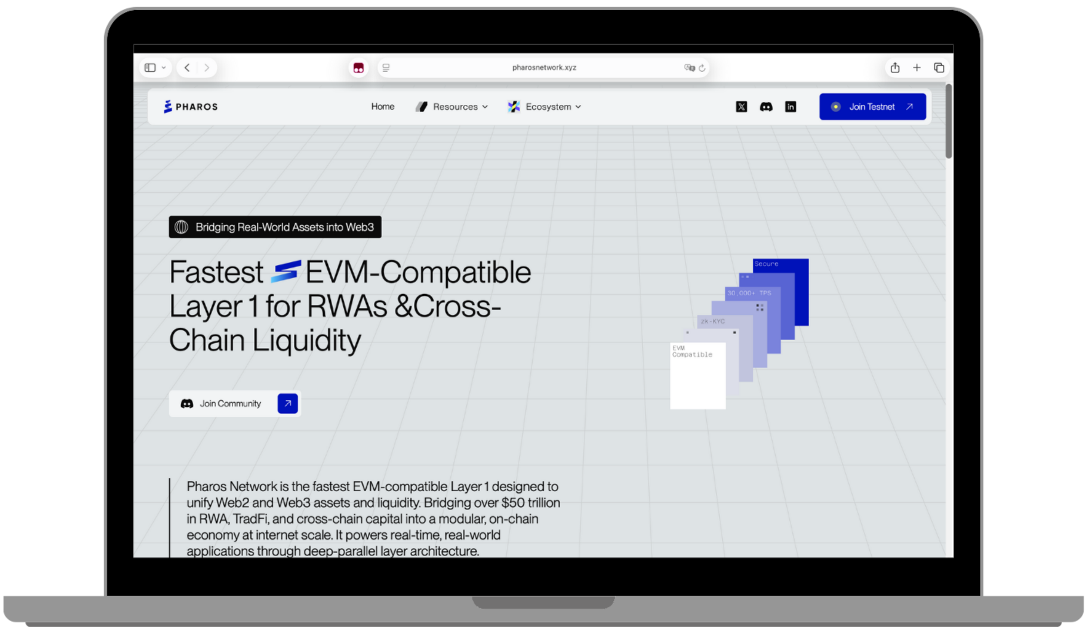

WeChat Official Account: Pharos Research

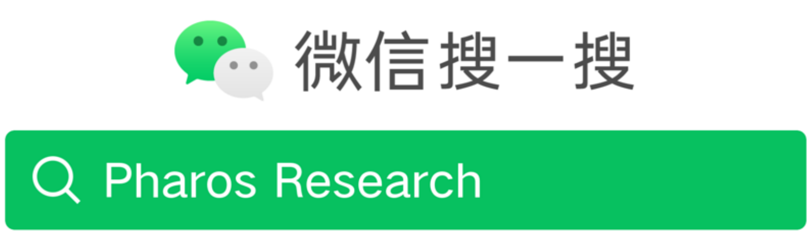

PHAROS
RESEARCH

From RWA to On-Chain Finance.

Mapping Real-World Value.