BLOCKDAEMON

WHITEPAPER

From Fragments to Fabric

Macro Infrastructure Trends Across The ETH L2 Landscape

Authors:
Demetrios Skalkotos, Chief DeFi and Protocols Officer, Blockdaemon,
and the Blockdaemon Data Team.

Table of Contents

1 Introduction to Ethereum's Layer 2 Infrastructure

1.1 What is a Layer 2?

Ethereum launched in 2015 as a decentralized platform for programmable digital money. By 2020, Ethereum had achieved this goal, and followed an ambitious scaling strategy to drive further adoption. This scale would be achieved by Layer 2 (L2) solutions called rollups.

A Layer 2 (L2) is a collective term to describe a specific set of Ethereum scaling solutions. It is a separate blockchain that extends Ethereum and inherits the security guarantees of Ethereum.

As Ethereum founder Vitalik Buterin noted, "the Ethereum ecosystem is likely to be all-in on rollups (plus some plasma and channels) as a scaling strategy for the near and mid-term future."

Rather than relying mainly on Ethereum's base layer for transaction execution (limited to around 15 transactions per second), the ecosystem increasingly focused on Layer 2 networks with orders of magnitude higher transaction capacity. This allowed Ethereum's core network to evolve into a highly secure and decentralized settlement platform, optimized primarily for data availability and final settlement, while Layer 2 solutions handle greater transaction volumes.

The focus of this whitepaper will be on L2s that are public, permissionless networks.

		READ	
		Everyone	Restricted
WRITE	Everyone	Public & Permissionless (e.g. BTC, ETH, L2s)	Private & Permissionless (Not seen in practice)
	Restricted	Public & Permissionless (e.g. Institutional L2s)	Public & Permissioned (Consortia Blockchains)

Figure 1: Blockchain Types

1.2 What Are Rollups?

Rollups are L2 networks designed to process transactions efficiently off Ethereum's main chain. After processing, rollups publish transaction data and proofs back to Ethereum for verification. There are two main types of rollups: Optimistic rollups (e.g., Arbitrum, Optimism), which assume transactions are valid unless challenged, and Zero-Knowledge rollups (zkRollups) (e.g., zkSync, Starknet), which provide cryptographic proofs upfront.

Computation / Execution

Data Storage / Data Availability

The Layer 2 Two-by-Two	Zero Knowledge Validity Proofs	Interactive Deposit-Slashing Fraud Proofs
Everyone	zkRollup	Optimistic Rollup
Restricted	Validium	Plasma → Optimum

Figure 2: Source buildblockchain.tech

From their nascent beginnings in 2020 to today, these rollup L2s have processed billions of dollars worth of transactions, and accrued massive value. On Ethereum, gas limits are 15 million every 12 seconds. L2s aim for far higher throughput, with a notable example being the medium term goal for BASE to reach 1 gigagas per second (1 Ggas/s). This is 1 billion gas per second.

As noted in 2024, "Major applications have started to move over from L1 to L2, payments are starting to be L2-based by default, and wallets are starting to build their user experience around the new multi-L2 environment." The success here is clear, with Kraken, Sony, and Coinbase, as notable examples using L2s.

2 Top Layer 2 Networks by Numbers

Daily transactions processed by L2s have grown explosively, from approximately 1 million in early 2023 to more than 25 million by mid-2025. According to Dune, as of June 2025, the top three L2s by transaction share are:

1. **Base: 31%**
2. **Polygon: 17%**
3. **Arbitrum: 12%**

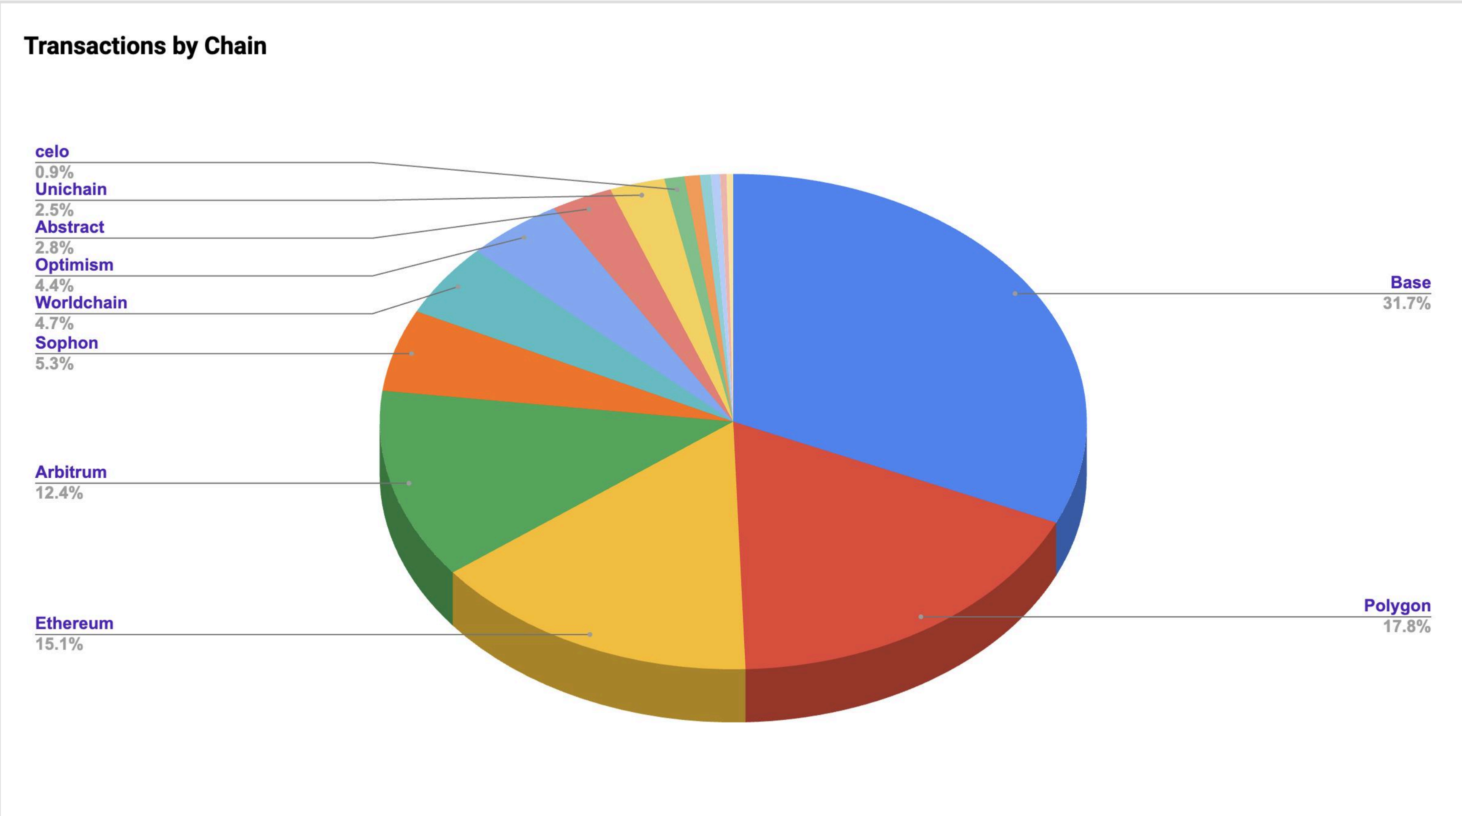

Figure 3: Transactions by Chain - June 2025. Source: Dune, 2025

Ethereum represented just 15% of total transactions, yet still handled 85% of all value transferred, proving that it remains the infrastructure settlement layer for a thriving L2 ecosystem.

Total Value Locked (TVL) across L2 networks reached a record high of $51.5 billion by November 2024, marking a 205% year-over-year increase. Prominent examples include Aave on Arbitrum, holding over 200,000 ETH, showing strong user confidence and substantial economic activity on L2 platforms.

3 ETH L2 Infrastructure Trends

3.1 Trend 1: Sequencers Go Decentralized

In the early days of Ethereum, the transaction lifecycle was relatively simple. An Ethereum user would send a transaction from their wallet of choice. It would then enter the mempool, a shared holding area for unconfirmed transactions. An Ethereum miner would then construct a block of these transactions from the mempool, prioritizing those with higher gas.

However, the growth of Ethereum's decentralized finance (DeFi) economy placed a premium on block space that offered a financial upside for successful transactions. In the context of these opportunities therefore, the concept of **MEV (maximum extractable value)** was introduced. MEV refers to the total amount of ETH stakers (previously miners) can extract from including, excluding, and changing the order of transactions in a block within a given timeframe.

MEV looks different on L2s compared to ETH L1 due to **sequencers**.

3.1.2 What are Sequencers?

Sequencers collect user transactions, organize them into batches, and compress them into ordered groups off-chain. These compressed, sequenced batches are then sent back to Ethereum's L1 as finalized blocks.

While rollups can technically function without a dedicated sequencer by relying directly on Ethereum's L1, doing so tends to be costly and inefficient. Consequently, most rollups choose to use centralized sequencers for lower fees and quicker transaction confirmations.

However, centralized sequencers pose risks. They have significant control over transaction ordering, which can lead to censorship or extracting MEV, potentially harming users economically. Additionally, centralized sequencers create single points of failure. If the sequencer fails, the entire rollup's operations could stall.

3.1.3 What are Decentralized Sequencers?

To address these issues, the industry is moving toward decentralized, shared sequencers. These decentralized solutions offer "decentralization-as-a-service," solving problems related to censorship, MEV extraction, and reliability. Shared sequencers also enable cross-rollup composability, unlocking innovative functionalities across multiple rollups.

Projects like Espresso, Astria, and Radius are at the forefront of decentralized sequencing innovation. Espresso has announced an ecosystem partnership with EigenLayer and plans to integrate restaking to help secure the Espresso Network; Astria uses Celestia for data availability, and Radius features an encrypted mempool to enhance security and privacy.

3.1.4 Based Rollups (early trend)

An emerging, longer-horizon direction is the rise of based rollups: L2s that delegate sequencing to Ethereum's L1 validator set, effectively using Ethereum itself as the decentralized sequencer. This approach targets the core pitfalls of centralized sequencers - censorship risk, MEV capture, and single-point-of-failure - while also aiming to reduce fragmented liquidity and improve cross-rollup composability by coordinating ordering at the base layer. Adoption is still early (Taiko is the notable exception with live progress), but research and coordination efforts such as Fabric - an initiative to accelerate connectivity for based-rollup infrastructure - suggest where the stack may be heading.

3.2 Trend 2: Data-Availability Layers Become Competitive

3.2.1 What is Data-Availability?

For rollups to function securely, Ethereum must guarantee Data Availability (DA), ensuring transaction data published by rollups is reliably stored and accessible. Ethereum's role is not executing these transactions but providing infrastructure for their authenticity and completeness.

3.2.2 Proto-Danksharding Reduced Costs

Ethereum's Dencun upgrade on March 13, 2024, significantly improved DA through "blobs" introduced by EIP-4844. Blobs are specialized data containers within Ethereum blocks dedicated exclusively to rollup transaction data. Unlike traditional Ethereum calldata, blobs are not processed by Ethereum's virtual machine (EVM), providing a cheaper dedicated lane for rollups to store data. This innovation initially reduced L2 transaction fees dramatically, by over 100 times in some instances.

Vitalik Buterin stated that blobs marked Ethereum scaling's shift from a "zero-to-one" foundational achievement to a "one-to-N" incremental approach. With blobs' core structure established in Dencun, Ethereum could gradually expand blob capacity without further fundamental changes.

In May 2025, the Pectra upgrade significantly increased the blob limit per block, although initially only two-thirds of the added capacity was utilized. Increased blob capacity drove fees down to near-zero (<$0.000001 each), substantially reducing blob-tip income and validator rewards in May 2025.

3.2.3 Further Innovations

Between late 2023 and 2025, modular DA layers like Celestia (launched Oct 31, 2023) and Avail significantly advanced rollup infrastructure. Avail's governance approved increasing block size to 4 MB in early 2025. Celestia maintains a 2 MiB per-transaction limit and is progressively raising maximum block capacity toward 8 MB via governance, with ambitions for further expansion.

EigenLayer introduced a novel restaking protocol in 2023, allowing Ethereum validators to lend staked ETH security to support Actively Validated Services (AVSs), including EigenDA - a specialized DA layer tailored for rollups. EigenLayer's Total Value Locked (TVL) peaked above $20 billion in 2024, fluctuating between approximately $11 billion and $16 billion throughout 2025, reflecting market conditions.

These innovations reshape rollup economics, reducing user transaction fees, fostering broader adoption, and enabling specialized blockchain architectures without compromising security or decentralization.

Date	Milestone
31 Oct, 2023	**Celestia Mainnet** First stand-alone data-availability (DA) layer goes live.
07 Nov, 2023	**Avail "Clash of Nodes" incentivized testnet** Kicks off Avail's public path to mainnet and positions it as an alternative DA layer.
13 Mar, 2024	**Ethereum "Dencun" hard fork (EIP-4844)** Introduces blobs. Layer 2 DA costs fall by roughly 100x.
29 Jan, 2025	**Avail block-size cap raised to 4 MB** Signals aggressive capacity scaling outside Ethereum.
07 May, 2025	**Ethereum "Pectra" hard fork** Doubles the blob target/limit; average blob fees drop to well under $0.000001.
Mid-2025 to 2026	**Celestia roadmap to 8 MB blocks** Governance proposals move maximum block size toward 8 MB, keeping non-Ethereum DA layers competitive.

3.3 Trend 3: Interoperability & Liquidity Networks

3.3.1 What is Interoperability?

In the context of web3, interoperability is the ability of one blockchain to communicate with another.

L2 blockchains use three main methods and technologies to achieve this goal.

3.3.1.1 Shared-stack federations

Shared-stack federations coordinate many chains that all run the same rollup stack under a common governance, canonical bridge, and security roadmap, so they can operate like one network. In the Optimism Superchain, OP-Stack chains agree to the Law of Chains, align on upgrade cadence, and rely on shared bridging to reduce liquidity fragmentation and standardize security guarantees. This model gives builders uniform tooling, users smoother cross-chain UX, and the ecosystem a fast path to rolling out security and efficiency upgrades at scale - but it also means that a single upgrade or incident can ripple across every member chain. For example, during the Bedrock upgrade (June 2023) all Superchain partners paused bridging for several hours; users experienced delayed deposits/withdrawals, yet benefited from a simultaneous fee reduction once the unified upgrade went live. The trade-off is tighter coupling (governance, upgrades, incident risk) across all member chains.

3.3.1.2 Liquidity / Proof-Aggregation Layers

Liquidity / proof-aggregation layers (e.g., Polygon's AggLayer) create a single liquidity & accounting surface across many heterogeneous chains, using ZK or "pessimistic" proofs to ensure asset correctness while minimizing wrapped-token sprawl.

Instead of forcing chains to share the same stack or governance, these layers verify (or pessimistically constrain) chain states and standardize cross-chain transfers so assets remain fungible and portable. The upside is dramatic UX simplification, deeper liquidity, and cryptographic assurances; the trade-off is reliance on a new coordination layer whose governance, proving economics, and failure modes must be carefully designed.

3.3.1.3 Cross Chain Messaging Standards

Generalized cross-chain messaging standards (IBC, Chainlink CCIP, LayerZero) give builders a chain-agnostic way to send messages or tokens between any two chains, with a menu of verification models ranging from native light clients (max trust-minimization) to oracle/committee-secured attestations (speed & simplicity). They enable rich cross-chain functionality - governance calls, intents, token movements, oracle updates - without imposing shared governance or a common stack. The upside is universal reach and flexibility. The downside is fragmented security assumptions and UX, requiring developers (and users) to explicitly understand who and what they trust for each route.

3.4 Trend 4: Rollup-as-a-Service (RaaS) & Hyperchains

3.4.1 What is Rollup-as-a-Service?

Rollup-as-a-Service (RaaS) refers to managed platforms that streamline the deployment of app-specific Layer-2 rollups. Providers like Ankr, AltLayer, and Zeeve offer packaged solutions including sequencing, data availability integrations, monitoring, wallets, and upgrade paths. This significantly cuts down the complexity and time required to launch L2 rollups, compressing deployment timelines from several months down to just weeks. The ease of deployment lowers entry barriers, enabling startups and enterprises to quickly access affordable, high-performance blockchain solutions.

The growth of RaaS shifts competitive advantage from merely launching infrastructure to differentiating via token strategies, liquidity solutions, user acquisition, and specialized product offerings built on standardized rollup infrastructure.

3.4.2 zkSync ZK Stack Hyperchains

zkSync's ZK Stack provides developers with an open-source toolkit to rapidly deploy customizable, zero-knowledge (ZK)-secured L2 rollups known as "hyperchains."

Hyperchains offer teams considerable flexibility, allowing customization of throughput, transaction fee structures, and even business logic, tailored to specific application needs. Crucially, these hyperchains maintain interoperability through shared provers, bridges, and aggregated proofs, enabling seamless integration within the broader Ethereum ecosystem. This approach fosters an environment rich in specialized, domain-specific chains, especially attractive for sectors such as trading, gaming, and real-world asset (RWA) applications.

3.4.3 Enterprise Adoption

Enterprise demand for dedicated L2 solutions is rapidly accelerating, driven by the unique benefits of rollups. Blockdaemon's data reveals that enterprise enquiries for custom rollups have doubled since 2023, underscoring significant growth in institutional interest. Enterprises prioritize dedicated L2s for their strong data isolation capabilities, customizable compliance frameworks, predictable transaction fees, and reliable performance assurances. Dedicated rollups allow enterprises to implement sophisticated blockchain solutions efficiently, without incurring the substantial overhead required by a fully independent blockchain.

Looking ahead, bespoke rollups are poised to become the standard enterprise blockchain deployment model. Providers will increasingly differentiate themselves through robust service-level agreements (SLAs), enhanced compliance modules, auditability features, disaster recovery capabilities, and advanced integration with custody, Know Your Customer (KYC), monitoring, and analytics solutions.

3.5 Trend 5: The Modular Thesis Matures

3.5.1 Shift to Modular Design

Modular blockchain architectures are redefining the industry by clearly separating execution (running applications), **consensus** (ordering transactions), and **data availability** (ensuring data is accessible).

Teams can independently choose, replace, or upgrade each layer, such as using Celestia or EigenDA for data availability, a shared sequencer for consensus, and a specialized execution environment tailored to their specific use cases. This design offers significant flexibility and scalability, enabling new innovations without requiring major disruptions like hard forks.

Applications can now precisely tailor security levels, adopting solutions such as light clients, zero-knowledge proofs, or restaked security based on their specific risk and trust requirements.

Modular Blockchain Stack

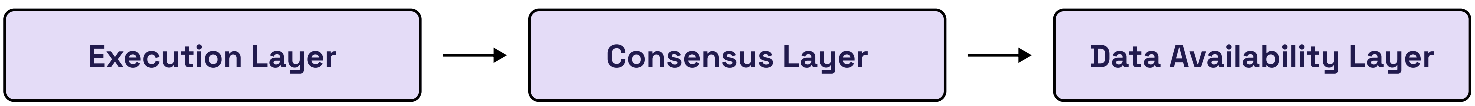

3.5.2 Cost and Velocity Benefits

Modular architectures substantially reduce infrastructure complexity and launch costs.

With Rollup-as-a-Service (RaaS) and modular data availability solutions, teams can quickly assemble proven components rather than building entire tech stacks from scratch. This accelerates deployment timelines and enables rapid iteration cycles.

Teams can easily adjust components like DA layers, sequencers, or provers to continuously optimize costs and performance without extensive codebase overhauls. Specialized rollups focused on specific niches - such as privacy, AI-driven applications, gaming, high-frequency trading, or real-world assets - can further refine economic models, state management, and performance settings to better match their unique needs.

3.5.3 Liquidity and Coordination Layers

Modular architectures, however, risk fragmenting liquidity, user identity, and application state across various specialized chains.

Emerging liquidity-coordination solutions, like Mitosis, address these challenges by re-bundling liquidity and composability. Coordination layers provide unified asset accounting, streamlined bridging, and aggregated proof systems. Users benefit from a seamless, unified network experience with simplified asset transfers and interactions, while developers maintain the freedom and flexibility of modular solutions beneath the surface. This balances modular freedom with cohesive, user-friendly experiences.

3.5.4 What the next Ethereum hard forks mean for L2s

Two upcoming L1 upgrades - Fusaka (next) and Glamsterdam (after) - aren't L2-specific, but they materially change L2 economics and UX.

Fusaka

The headline feature in Fusaka is **Peer Data Availability Sampling (PeerDAS)**, specified in **Ethereum Improvement Proposal (EIP) 7594**. PeerDAS allows nodes to sample rollup blob data rather than download it in full, which makes it safe to raise blob capacity - targets discussed are approximately 48 blobs per block. The result is lower L2 data costs and higher throughput. Fusaka's **meta-EIP** also tracks mechanisms such as blob-parameter-only hard forks and a blob base-fee bound, giving core developers levers to tune **Data Availability (DA)** capacity as demand grows.

Gas Limit Increases

Client teams have drafted an increase to the default L1 gas limit in **EIP-7935**, and there is active debate and testing around substantially larger increases. Meanwhile, validators are signaling toward 45 million gas on mainnet. Treat the "150 million" figures as exploratory for now; the reliable takeaway is that Fusaka includes gas-related EIPs and tuning work is underway.

Glamsterdam

A draft proposal, **EIP-7782**, would cut slot time from roughly 12 seconds to 6 seconds, improving confirmation latency for bridges, sequencer settlement, and other price-sensitive flows. In parallel, **EIP-7928** introduces **Block-Level Access Lists (BALs)** to enable parallel execution and validation with a modest ~40 KB average per-block overhead—creating headroom for future gas and throughput gains that L2s can compound.

Fusaka's meta-EIP explicitly omits the **EVM Object Format (EOF)** for now, keeping the roadmap focused on scalability levers - data availability, gas parameters, and blob economics - that most directly improve rollup fees and capacity.

4 Implications for Builders & Institutions

4.1 Cost curves & user experience

As mentioned, the March 2024 Dencun upgrade (Proto-Danksharding / EIP-4844) cut L2 data-posting costs by ~90%, driving rollup fees paid to Ethereum mainnet down to roughly 200–300 ETH per day. Median end-user fees on leading rollups - Base, Arbitrum, Optimism - now sit below one cent, matching high-throughput L1s like Solana while retaining Ethereum's settlement security. The result is retail-grade applications without performance trade-offs.

4.2 Security choices & risk matrices

Institutions must choose among native L1 security, restaked models, and emerging EigenDA-style frameworks. Each path carries distinct assumptions and risks across dimensions such as economic security, liveness and censorship resistance, smart-contract and implementation risk, governance centralization, and regulatory/compliance exposure. Selecting the right stack requires mapping these trade-offs to project goals, operational maturity, and jurisdictional requirements.

4.3 Liquidity strategy & governance

Superchains and agglayers compress liquidity fragmentation and boost capital efficiency, but they also introduce layered governance complexity. Coordinating liquidity across domains, setting shared upgrade and parameter policies, aligning MEV/fee-sharing rules, and managing dispute resolution become critical to preserve interoperability while maintaining decentralized control. Effective, credibly neutral governance and explicit liquidity-coordination mechanisms are now core design requirements, not afterthoughts.

5 Blockdaemon’s Research Lens

Blockdaemon works directly with protocol foundations and core developers to surface operational realities - performance bottlenecks, security edge cases, UX frictions - and then helps harden the stack without encroaching on protocol roadmaps or token economics. The mandate is simple: amplify what makes each network unique, and translate those strengths into measurable reliability, security, and user experience gains.

Blockdaemon runs **approximately 250,000 nodes across more than 60 blockchain networks**, giving it continuous insights on outages, latency patterns, client/version fragmentation, and usage spikes.

Blockdaemon is **actively operating shared-sequencer nodes for Espresso and Router, and data-availability validators for Celestia and Avail**, ensuring its models and recommendations are anchored to current network behavior.

Because Blockdaemon engages across every layer of the stack - execution, sequencing, DA, interoperability, and staking - it can act as an **ecosystem accelerator**: introducing protocols to the right partners, helping them productionize features, and smoothing the path from whitepaper to mainstream adoption. That combination of **technical depth and market reach** is the differentiator.

In short, Blockdaemon’s research lens fuses **collaborative protocol engagement, massive real-world observability, and hands-on piloting of next-gen components**.

6 Conclusion

Ethereum's post-Dencun L2 ecosystem is crystallizing around five reinforcing forces.

First, decentralization is no longer confined to base-layer validators. Sequencers, MEV supply chains, and shared prover/DA markets are fragmenting - and with that, revenue models are being redrawn.

Second, data-availability has become a competitive arena. Proto-Danksharding pushed on-chain DA costs toward zero, while restaked security markets (e.g., EigenLayer/EigenDA) and alternative DA layers (Celestia, Avail) introduce fluid cost and security trade-offs that will shape how future rollups are architected and priced.

Third, interoperability is finally turning into user experience, as cross-L2 messaging standards and liquidity networks make assets and data move seamlessly, reducing the "which L2 am I on?" friction that once defined the rollup era.

Fourth, modularity has matured from thesis to tooling. Teams can launch faster, specialize deeper, and compose more freely, compressing experimentation cycles and lowering capital intensity.

Finally, institutions are demanding audit-grade UX. Deterministic ops, risk-mapped security choices, and coordinated liquidity/gov processes. Meeting those requirements elevates observability, telemetry, and compliance-aware tooling to first-class primitives - areas where providers like Blockdaemon are already supplying essential infrastructure.

Put together, these currents imply a future where rollups resemble configurable platforms competing on fee curves, DA guarantees, finality speed, and governance assurances - interconnected by standard bridges and liquidity fabrics, secured by modular, opt-in cryptoeconomic layers. The winners will be the institutions who can continuously meet the highest security standards, minimize UX surface area for both retail and institutions, and operationalize credible, decentralized governance over increasingly shared components.

7 About Blockdaemon

Blockdaemon is the institutional gateway to Web3, securing over $110 billion in digital assets for 400+ institutions—including exchanges, custodians, crypto platforms, and financial enterprises. The company offers institutional-grade blockchain infrastructure spanning nodes, APIs, staking, MPC wallets and vaults. Since 2017, Blockdaemon's globally distributed infrastructure ensures unrivaled compliance and scalability.

For more information visit: Blockdaemon.com