STATE OF THE MARKET

2025 REPORT

Routescan

THE 2026 INFRASTUCTURE PLAYBOOK

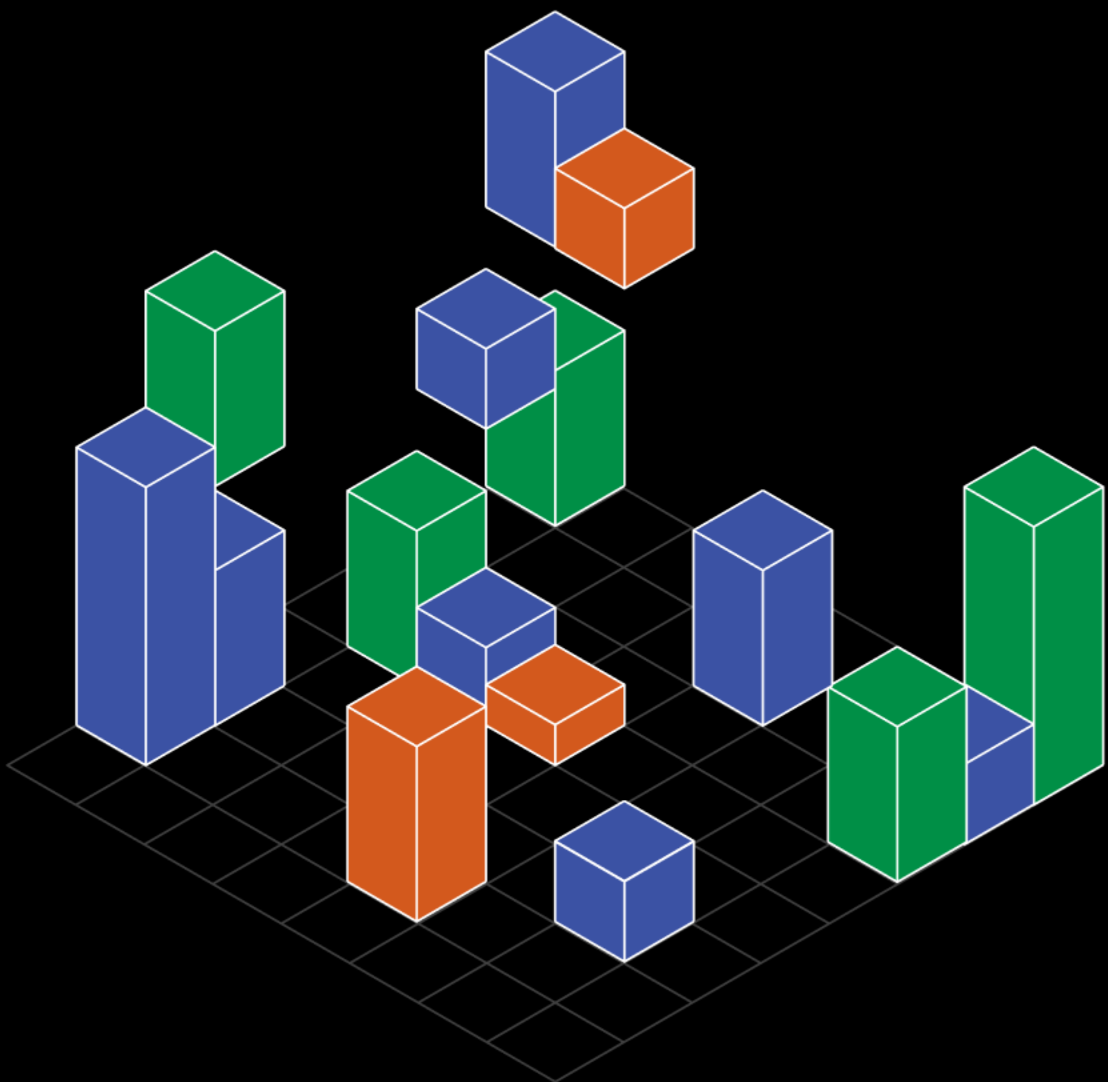

Index

Routescan | Research

Introduction

2025 confirmed the industry's shift toward production-ready blockchain infrastructure, with clear consolidation around a small set of execution environments and rollup architectures.

This report analyzes 3,473 validated blockchain networks as a year-end 2025 snapshot, including a layer-classified subset of 2,133 networks (L1/L2/L3).

The analysis focuses on operational signals, mainnet vs. testnet deployment, execution environment standardization, and infrastructure dependency formation, validated through RPC-level checks, manual verification, and structured metadata extraction.

Where coverage depends on classification availability (e.g., ecosystems, VM identification), metrics are computed on explicitly stated subsets rather than extrapolated over the full validated universe.

Imagine you need to navigate tens of emerging blockchain technologies

—how do you quickly identify the most robust, promising, and truly reliable solutions for your project or investment strategy?

Do you fully understand how the consolidation of EVM and the rise of RaaS providers could transform the industry in the coming months?

What is Routescan?

Routescan is a multichain explorer and data platform used by 160+ blockchains to make onchain data auditable, searchable, and operationally useful across environments.

It provides unified explorer infrastructure and developer-grade tooling, including rapid explorer deployments (e.g., 48-hour setup) and standardized integrations across chains.

Routescan's research methodology mirrors this operational focus: networks are validated directly at the protocol level (RPC testing), cross-checked against multiple sources, and enriched with metadata to support reliable infrastructure classification.

Who should read this report—and why?

- **Protocol & infrastructure founders (L1/ L2):**
Use the dataset to reduce stack-selection risk (framework/rollup architecture), validate execution-environment tradeoffs (EVM vs WASM vs COSVM), and define an explorer/ observability strategy aligned with time-to-production and provider dependencies.
- **Developers & blockchain engineers:**
Use VM, rollup, and explorer distributions to choose compatible tooling, understand infrastructure dependencies (RaaS/ explorers), and benchmark production readiness using mainnet adoption patterns.
- **Market Analysts and Researchers:**
Use denominator-explicit subset analytics (layer-classified, ecosystem-mapped, VM-identified) to analyze concentration vs fragmentation without extrapolating beyond mapping coverage.
- **Investors and Venture Capital Firms:**
Use validated infrastructure signals (production readiness, dependency concentration, rollup/VM standardization) as inputs for risk assessment and market-structure theses.
- **Institutional and Business Leaders:**
Use a validated network universe and clearly stated subsets to inform adoption strategy, vendor selection (RaaS/explorers), and partnership evaluation.

What key questions does this report answer?

- Which execution environment choices show the strongest production conviction (mainnet adoption), and where is EVM still the default vs where non-EVM is meaningfully deployed?
- Which rollup architectures are converging into standards (Optimistic vs ZK), and what does that imply for 2026 stack selection?
- How concentrated is managed infrastructure (RaaS) vs self-hosted deployment, and where do dependency risks cluster?
- Which ecosystems/frameworks dominate where mappings exist, and how should subset coverage affect interpretation (no "total market" inference from partial mappings)?
- What explorer/observability patterns correlate with maturity by layer, and what is the practical explorer strategy at launch?

Why is this research worth your attention?

- **Precision and Accuracy:**
Routescan's data collection methods are rigorously validated and manually verified, reducing reliance on unverified claims and improving technical trustworthiness.

- **Comprehensive and Unique Coverage:**
The dataset covers **3,473** cataloged networks, with **2,133** networks categorized by layer (L1/L2/L3) to support cross-layer infrastructure analysis.

- **Strategic Relevance**: The report connects execution environment choices (**18 VM types**) and scaling architectures (**6 rollup types**) to operational infrastructure decisions, including RaaS provider reliance (**10 providers tracked**).

- **Market Understanding:** The report frames ecosystem distributions explicitly as a **subset** view (**1,038 chains across 24 ecosystems**), avoiding misleading "total market" conclusions from partial mappings.

RESEARCH OBJECTIVES

This report evaluates how VM distribution, rollup-architecture adoption, and infrastructure dependencies reflect the industry's shift toward production-ready deployment patterns in 2025, and what this implies for 2026 infrastructure decisions.

Central research question:
How does the 2025 distribution of virtual machine types, measured on the VM-identified subset of networks, together with rollup architecture concentration and layer-specific production readiness (mainnet adoption), reflect the evolution toward standardized yet modular blockchain ecosystems, and what actionable implications follow for 2026 stack selection and infrastructure strategy?

How to use this report:
This report can be read end-to-end for a full understanding or consulted by section depending on the decision problem (execution environment, rollup architecture, RaaS reliance, explorer strategy).
Unless specified, percentages use the validated network universe (3,473). Layer metrics use the layer-classified subset (2,133). Ecosystem distributions use the ecosystem-mapped subset (1,038). VM and rollup distributions use the VM-identified and rollup-identified subsets respectively, with denominators stated in their dedicated sections. When interpreting ecosystem distributions, treat them strictly as subset analytics (1,038 mapped chains) rather than projections over the full validated universe.

Important Notice:
Methodology is evolving
This analysis reflects year-end 2025 infrastructure based on a validation-first methodology that continued to evolve throughout 2025 as discovery, validation, and categorization standards improved. Because stricter criteria can both add newly discovered networks and remove networks that fail operational checks, stakeholders should prioritize directional signals (architecture

adoption, dependency concentration, production readiness) over naive period-to-period comparisons of absolute counts. All reported percentages are rounded; small discrepancies from 100% are expected due to rounding and do not affect interpretation.

Chapter I: Data Overview

How can stakeholders trust infrastructure analysis when the blockchain landscape evolved significantly throughout 2025, with new networks launching and others becoming inactive?

The answer lies in making the methodology auditable: clearly documenting validation criteria, classification rules, denominators, and subset boundaries so that every metric can be interpreted and challenged without requiring distribution of the underlying dataset. This chapter establishes the methodological foundation underlying every insight, trend analysis, and strategic recommendation presented throughout this report.
Our approach addresses a fundamental challenge in blockchain research: **the tension between comprehensive coverage and analytical precision in a rapidly evolving technological landscape.**

1.1
Data Acquisition Methodology

Important Methodological Disclaimer

Critical Context for Data Interpretation: This report reflects year-end 2025 blockchain infrastructure. Our methodology continued to evolve throughout 2025, with enhanced validation criteria and refined ecosystem categorization standards.
The total count of qualified networks analyzed stands at 3,473 catalogued blockchain networks, with 2,133 blockchains categorized by layer (L1/L2/L3). This represents the most comprehensive and rigorously validated blockchain infrastructure dataset available.

Factors Contributing to Dataset Evolution:

- **Enhanced Discovery Processes:** Throughout 2025, improved network discovery techniques surfaced previously unmapped blockchain infrastructure, expanding the validated universe.
- **Refined Classification Standards:** Ecosystem and component categorization continued to evolve, improving technical alignment between what networks claim and what they actually run in production.
- **Infrastructure Maturation:** Some networks operating as separate entities in earlier periods consolidated, merged, or clarified technical relationships, impacting counting and classification outcomes.
- **Data Quality Improvements:** Enhanced validation processes exclude networks that fail operational or technical criteria, trading raw counts for stronger reliability.

Strategic Implication: While these numbers provide valuable insight into the blockchain infrastructure landscape, stakeholders should focus on qualitative trends, technical patterns, and ecosystem dynamics rather than treating numerical variations as direct market growth indicators. The value lies in understanding architectural evolution, adoption patterns, and strategic positioning within the mapped infrastructure.

The Manual-First Validation Framework

Routescan's data collection methodology prioritizes manual verification and direct blockchain validation over scalable but less reliable automated approaches. This reflects a deliberate trade-off: sacrificing collection speed to achieve higher data accuracy and reliability.
Our comprehensive dataset encompasses:

3,473
unique blockchain networks discovered and cataloged

1,621
testnet

1,852
mainnet

2,133
blockchain integrated within established layer ecosystems (L1/L2/L3)

2,155
block explorer integrations manually validated

18
distinct virtual machine types identified and classified

6
rollup architecture categories mapped and validated

10
distinct RaaS providers tracked

1,621
primary ecosystems categorized

8
frameworks analyzed

591
unique explorer provider organizations categorized and analyzed

RPC-Level Validation Process

Direct blockchain validation forms the cornerstone of our methodology. Rather than relying on third-party APIs or documentation that may be outdated, validation is grounded in direct checks of network status, configuration, and operational parameters.

Validation Pipeline:

- **Network Discovery:** Systematic identification through multiple source aggregation
- **Chain ID Verification:** EVM compatibility and identifier confirmation

Quality Assurance Standards:

- **Multi-Source Cross-Validation:** Each network validated against multiple authoritative sources
- **Technical Specification Verification:** Direct validation of VM types, consensus mechanisms, and architectural choices

Direct Explorer Probing

Block explorer validationprovides critical infrastructure intelligence beyond basic network discovery. Our methodology includes systematic evaluation of 2,155 explorer integrations, provider relationships, and implementation quality across production and staging environments.

Explorer Analysis Framework:

- **Provider Identification:** Manual verification of explorer development and hosting organizations
- **Technical Implementation Assessment:** Feature completeness, update frequency, and reliability analysis
- **Network Coverage Validation:** Confirmation of accurate blockchain data representation
- **Infrastructure Relationship Mapping**: Provider-network affiliations and strategic partnerships

Metadata Extraction Standards

Comprehensive metadata collection enables nuanced ecosystem analysis beyond surface-level network counting. Our approach systematically captures technical specifications, governance structures, and operational parameters that inform strategic insights.

Data Categories Systematically Collected:

- **Technical Architecture:** VM type, consensus mechanism, rollup configuration
- **Network Classification:** Layer classification (L1/L2/L3), environment type (mainnet/testnet)
- **Ecosystem Affiliation:** Parent chain relationships, framework utilization, provider dependencies
- **Operational Metrics:** Launch status, activity levels, development momentum
- **Infrastructure Dependencies:** RPC providers, RaaS providers, explorer services

Data Quality Disclaimer:
All percentages reported throughout this study are calculated from validated network data and rounded to enhance readability. Ecosystem and explorer coverage percentages are calculated based on subsets of networks meeting specific classification criteria, which may differ from total layer counts due to exclusion of networks with ambiguous technical specifications or insufficient operational data. Minor discrepancies in percentage totals (which may not sum exactly to 100%) result from decimal rounding and do not affect the validity of the analysis. Raw numerical data remains available for detailed verification.

DeFi Warhol
@Defi_Warhol

> “We are currently in a market where launching a chain is often cheaper than actually operating one, so validation at the RPC and explorer layer is the only reliable way to separate real infrastructure progress from a simple marketing strategy”

1.2
Reevaluating Ecosystem Classification

Refined Classification Framework

Ecosystem classification is one of the most methodologically challenging aspects of blockchain infrastructure analysis. The rapid evolution of modular architectures, rollup frameworks, and multi-chain solutions throughout 2025 required continuous refinement of categorization approaches to maintain analytical precision.

Classification Challenges Addressed:

- **Overlapping Ecosystem Boundaries**: Networks often reuse multiple frameworks or present hybrid stacks that obscure primary technical dependencies.
- **Evolution of Technical Standards**: Frameworks and rollup stacks can change materially across a network lifecycle, creating midstream classification risk.
- **Parent-Child Relationship Complexity**:L2/L3 networks introduce multiple dependency layers that complicate “ecosystem ownership” and settlement attribution.
- **Marketing vs. Technical Reality:** Classifications must reflect implementation reality rather than positioning language.

Consolidation of Overlapping Ecosystems

Ecosystem consolidation eliminates analytical confusion created by networks claiming multiple affiliations or using hybrid architectures. The methodology prioritizes primary technical dependencies over secondary integrations or marketing associations.

Consolidation Principles:

- **Primary Framework Precedence:** Classification based on core technical architecture rather than secondary integrations
- **Parent Chain Authority:** L2/L3 networks categorized by immediate settlement layer rather than ultimate L1 dependency
- **Technical Implementation Priority:** Actual code utilization over ecosystem partnership announcements
- **Operational Dependency Analysis:** and economic relationships influence classification decisions

Differentiation of Technical Components

Modular blockchain architecture requires classification that distinguishes execution environments, settlement dependencies, data availability choices, and rollup implementations. This report treats VM type, rollup architecture, and infrastructure provider dependencies (RaaS/explorers) as first-class classification dimensions to reflect technical reality.

Technical Component Classification:

- **Execution Environments:** Virtual machine implementation and smart contract capabilities
- **Settlement Layer Dependencies:** Where transaction finality and dispute resolution occur
- **Data Availability Architecture:** Onchain vs. offchain data storage and validation approaches
- **Rollup Type Implementation:** Optimistic vs. ZK proof systems with architectural subcategories

Refined Categorization Results:

24	**Ecosystems** tracked where ecosystem mappings are available, explicitly treated as a subset view (1,038 mapped chains)	6	**Rollup Architecture Types** distinguished by proof mechanism and deployment model
18	**Virtual Machine Implementations** categorized by compatibility and feature set, including EVM, WASM, and COSVM	3	**Layer Classifications** (L1, L2, L3) applied to the layer-categorized subset (2,133 chains)

Exclusion of Inactive and Ambiguous Networks

Network exclusion criteria ensure analytical focus on operationally significant blockchain infrastructure while maintaining transparency about methodology limitations. Exclusions include technically inactive networks, ambiguous technical classifications, pre-mainnet development-stage entries, and duplicate implementations without meaningful differentiation.

Exclusion Categories:

- **Technically Inactive Networks:** No RPC response, outdated explorer data, or confirmed shutdown
- **Ambiguous Technical Classification:** Insufficient technical documentation or conflicting architectural claims
- **Development Stage Networks:** Pre-mainnet projects without operational infrastructure
- **Duplicate or Fork Networks:** Identical technical implementations without differentiated utility

Quality Threshold Standards:

- **Operational Verification:** Networks must demonstrate active RPC endpoints and live blockchain data
- **Technical Documentation:** Sufficient architectural information for accurate classification
- **Infrastructure Maturity:** Block explorer availability or alternative verification mechanisms mechanisms
- **Ecosystem Contribution:** Meaningful differentiation from existing network implementations

Strategic Implications of Methodology

Methodological rigor directly impacts the strategic value and reliability of all subsequent analysis presented in this report. Our validation-focused approach creates several competitive advantages for stakeholders utilizing this intelligence:

For Technical Decision-Makers:

- **Confidence in Infrastructure Choices:** Validated network status and technical specifications reduce selection risk
- **Architectural Understanding**: Clear component classification enables informed technology stack decisions
- **Ecosystem Positioning:** Accurate relationship mapping supports strategic partnership and integration planning

For Investment and Market Analysis:

- **Market Reality vs. Hype:** Manual verification distinguishes operational networks from speculative projects
- **Growth Trend Accuracy:** Validated data enables reliable quarter-over-quarter comparison and forecasting
- **Competitive Landscape Clarity:** Precise ecosystem classification reveals genuine market concentration and opportunities

This methodological foundation ensures that every data point, trend analysis, and strategic insight presented throughout this report meets the highest standards of accuracy and reliability demanded by professional blockchain infrastructure analysis. professional blockchain infrastructure analysis.

The comprehensive dataset and validation methodology described in this chapter underpins all subsequent analysis. Readers can reference specific technical validation criteria and data collection standards in our methodology appendix for additional transparency and verification.

Chapter II: Executive Summary

The Strategic Challenge Facing Blockchain Infrastructure

How do you make informed infrastructure decisions when the blockchain ecosystem encompasses 3,473 validated networks across 18 different VM types?
This fundamental question shaped the strategic challenge facing protocol builders, investors, and institutional decision-makers throughout 2025. The year highlighted how quickly "default choices" emerge (EVM compatibility, Optimistic rollups, managed RaaS) while non-default execution environments (WASM and COSVM) proved they can reach production at meaningful scale.

This report addresses that challenge by providing validated infrastructure data across all layers and distilling actionable insights.

What This Report Reveals

Infrastructure Maturation Evidence

Blockchain infrastructure maturation throughout 2025 is evidenced by four key shifts:
1. **Production-First Achievement:** Mainnet deployments now represent 53.3% of total infrastructure (1,852 networks) versus 46.7% testnet (1,621 networks), marking the first definitive production-first tilt across the entire ecosystem.2.

2. **EVM Dominance with Continued Erosion:** EVM maintains **81.52% market dominance (2,721 chains)** but demonstrates below-average production readiness at 46.9% mainnet adoption, 6.4pp below the 53.3% industry baseline. This gap exposes a structural weakness: EVM's scale is driven by compatibility defaults and low deployment friction, not production conviction. In contrast, WASM (85.6% mainnet, +38.7pp vs EVM) and COSVM (69.1%, +22.2pp) show materially stronger production signals, indicating that alternative VM selection correlates with deliberate technical differentiation strategies rather than experimental deployment.3.

3. **WASM Breakthrough as Credible Alternative:** WASM emerges as the first viable EVM alternative, capturing 11.23% market share (375 chains) with 85.6% of WASM chains deployed on mainnet (vs 46.9% for EVM), signaling that teams selecting WASM do so with materially higher production conviction.

4. **Rollup Architecture Standardization:** Optimistic Rollups continue to dominate with 57.68% rollup market share (154 chains), while ZK Rollups achieve meaningful scale at 26.59% share (71 chains). Optimistic and ZK rollups collectively account for 84.3% of the rollup ecosystem, establishing a mature two-standard market structure.

Strategic Themes 2025

- **Production Maturation Achieved**

The industry crossed the production-first threshold with 1,852 mainnet deployments (53.3%) outpacing 1,621 testnet environments (46.7%). This marks the first time production infrastructure definitively surpasses experimental deployment across the entire validated dataset, indicating genuine enterprise and production adoption rather than speculative development.

- **L1 Dominance Persists, L2/L3 Scale Achieved**

L1 infrastructure comprises 72.2% of categorized layer deployments (1,540 networks) with strong 58.1% mainnet adoption, while L2 reached meaningful scale at 24.0% share (512 networks) with 51.8% mainnet rate. L3 remains experimental at 3.8% share (81 networks) but demonstrates exceptional 63.0% mainnet adoption, the highest across all layers, indicating application-specific production focus.

- **WASM Emerges, EVM Production Weakness Exposed**

EVM maintained 81.52% market share but exhibits below-average production conversion with 46.9% mainnet adoption, 6.4 percentage points below the 53.3% industry baseline, while WASM reached 11.23% share with exceptional 85.6% mainnet deployment rate (38.7pp above EVM). COSVM achieved 5.72% share with strong 69.1% mainnet adoption (22.2pp above EVM).

This divergence reveals a counterintuitive pattern: teams selecting alternative VMs demonstrate materially stronger production conviction than the average EVM deployment. EVM dominates by volume, but alternative VMs dominate by production commitment, signaling that non-EVM adoption correlates with technical differentiation requirements and higher launch readiness.

- **L1 Layer Specialization**

On L1 specifically, EVM's share compressed to 61.1%, while WASM captured 25.3% and COSVM 12.6%, with WASM+COSVM combined representing 38.0% of L1 VM distribution, indicating that foundational sovereignty layers increasingly select specialized execution environments.

- **Ecosystem Consolidation Around Four Frameworks**

Ethereum-based (373 chains, 35.93%), Polkadot (149 chains, 14.35%), Cosmos (117 chains, 11.27%), and Avalanche (108 chains, 10.40%) collectively control 71.97% of categorized ecosystem infrastructure, demonstrating that framework provision, not individual chain deployment, determines strategic positioning. The top three ecosystems alone control 61.56% of categorized chains.

- **RaaS Market Consolidation with Stable Distribution**

RaaS provider distribution shows market maturation with two dominant players: Caldera RaaS (27.34% share, 35 chains with 77.1% mainnet focus) and Conduit RaaS (26.56% share, 34 chains with 82.4% mainnet focus), followed by Gelato (15.63% share, 20 chains). The top three providers control 69.53% of the RaaS market.

- **Rollup Architecture Dual-Standard Emergence**

Optimistic Rollups maintain clear leadership at 57.68% share (154 chains), while ZK Rollups achieved meaningful scale at 26.59% share (71 chains), establishing a dual-architecture standard rather than single-winner dominance. Alternative rollup types (Optimium 10.11%, Validium 4.87%) remain niche but production-viable.

Strategic Framework

Our analysis evaluates infrastructure choices through three strategic lenses:

- **Technical Capability Alignment**
  - Does the VM/rollup architecture match your performance and compatibility requirements?
  - What is the production readiness (mainnet adoption %) of your chosen technology?
  - Are you building for compatibility (EVM 81.52%) or production maturity (WASM 85.6% mainnet)?
- **Business Model Optimization**
  - Does ecosystem consolidation (Ethereum-based/ Polkadot/Cosmos/Avalanche) create network effects you can leverage?
  - Should you prioritize broad developer access (EVM) or demonstrate technical conviction through specialized VMs?
  - Can your architecture benefit from RaaS managed infrastructure (69.53% top-3 control) or require self-hosted autonomy (13.3% L2)?
- **Competitive Positioning**
  - Are you building in a production-mature segment (WASM 85.6% mainnet) or developer-friendly but lower-conviction market (EVM 46.9% mainnet)?
  - Does your infrastructure strategy align with layer-specific production patterns (L1 58.1%, L2 51.8%, L3 63.0% mainnet)?

Critical Guidance:
Focus on qualitative trend identification and architectural patterns rather than absolute numerical comparisons. The value lies in understanding where the industry is moving, not where it stood at a single moment.

Chapter Guide

Chapter 3: Layer 1 Analysis	Examines L1's sustained dominance (72.2% of categorized layers), WASM's breakthrough adoption in L1 infrastructure (25.3% layer share with 86.1% mainnet), and ecosystem consolidation patterns around Ethereum-based (13.51%), Polkadot (9.35%), Cosmos (7.53%), and Avalanche (7.01%).
Chapter 4: Layer 2 Analysis	Analyzes L2 scaling maturation (24.0% layer share, 51.8% mainnet), near-total EVM dominance (95.94% L2 VM share), stable RaaS provider distribution, and the persistent 13.3% self-hosted deployment rate that challenges "managed infrastructure dominance" narratives.
Chapter 5: Layer 3 Analysis	Explores L3 application-specific deployment patterns (3.8% layer share), exceptional production readiness (63.0% mainnet, highest across all layers), EVM concentration (94.20%), and the strategic question of whether L3 represents a sustainable architecture layer or experimental frontier.
Chapter 6: Cross-Layer Comparison	Synthesizes production maturation patterns, VM distribution evolution, rollup architecture adoption, and ecosystem positioning strategies across all three layers for comprehensive strategic planning.
Chapter 7: Conclusion & 2026 Outlook	Distills 2025 findings into actionable 2026 predictions and strategic recommendations for protocol builders, investors, and infrastructure providers.

Chapter III: Layer 1 (L1) Analysis

Layer 1 blockchains are the base infrastructure layer where sovereignty, security, and consensus define the boundary conditions for everything built above.

The 2025 dataset includes 1,540 L1 networks, confirming L1 as the dominant deployment layer and the primary arena where teams still pursue execution-environment differentiation beyond EVM. Across L1, the data shows the strongest concentration of non-EVM execution environments at meaningful scale, with WASM and COSVM acting as the only two alternatives that consistently map to production deployments.

3.1 Virtual Machine (VM) Distribution

The Foundation Layer Execution Environment Question

How do Layer 1 networks balance compatibility requirements with performance optimization when selecting execution environments?

> On L1, the “compatibility vs sovereignty” tradeoff is measurable: EVM remains the default for reach, but the strongest production conviction clusters on non-EVM execution environments. To keep this signal clean, VM shares in this section are computed on the VM-identified subset for L1 (1,472; 95.6% coverage of total L1 networks), meaning percentages sum to 100% within the set of networks where VM identification is validated.

VM Distribution Analysis

On L1, EVM remains the majority execution environment at 61.14% (900 chains), but its dominance is materially lower than the global EVM share, signaling that L1 is still where differentiation happens.
At the same time, EVM’s mainnet adoption on L1 is only 43.3%, which sits well below the L1 average mainnet rate (57.6%) and implies that many EVM L1 deployments remain pre-production or exploratory.
WASM represents the first credible non-EVM alternative at scale on L1: 25.34% share (373 chains) with 86.1% mainnet adoption, indicating strong production conviction among adopters.
COSVM follows with 12.64% share (186 chains) and 69.9% mainnet adoption, reinforcing the pattern that alternative VM selection correlates with higher production readiness on L1.
Combined, WASM + COSVM account for 38.0% of L1 VM distribution, making L1 the layer where non-default execution environments are most concentrated and most production-mature.
Beyond the “big three” (EVM/WASM/COSVM), L1 includes a long tail of niche VMs (11 total VM types on L1), some of which show perfect mainnet rates but at negligible network counts.

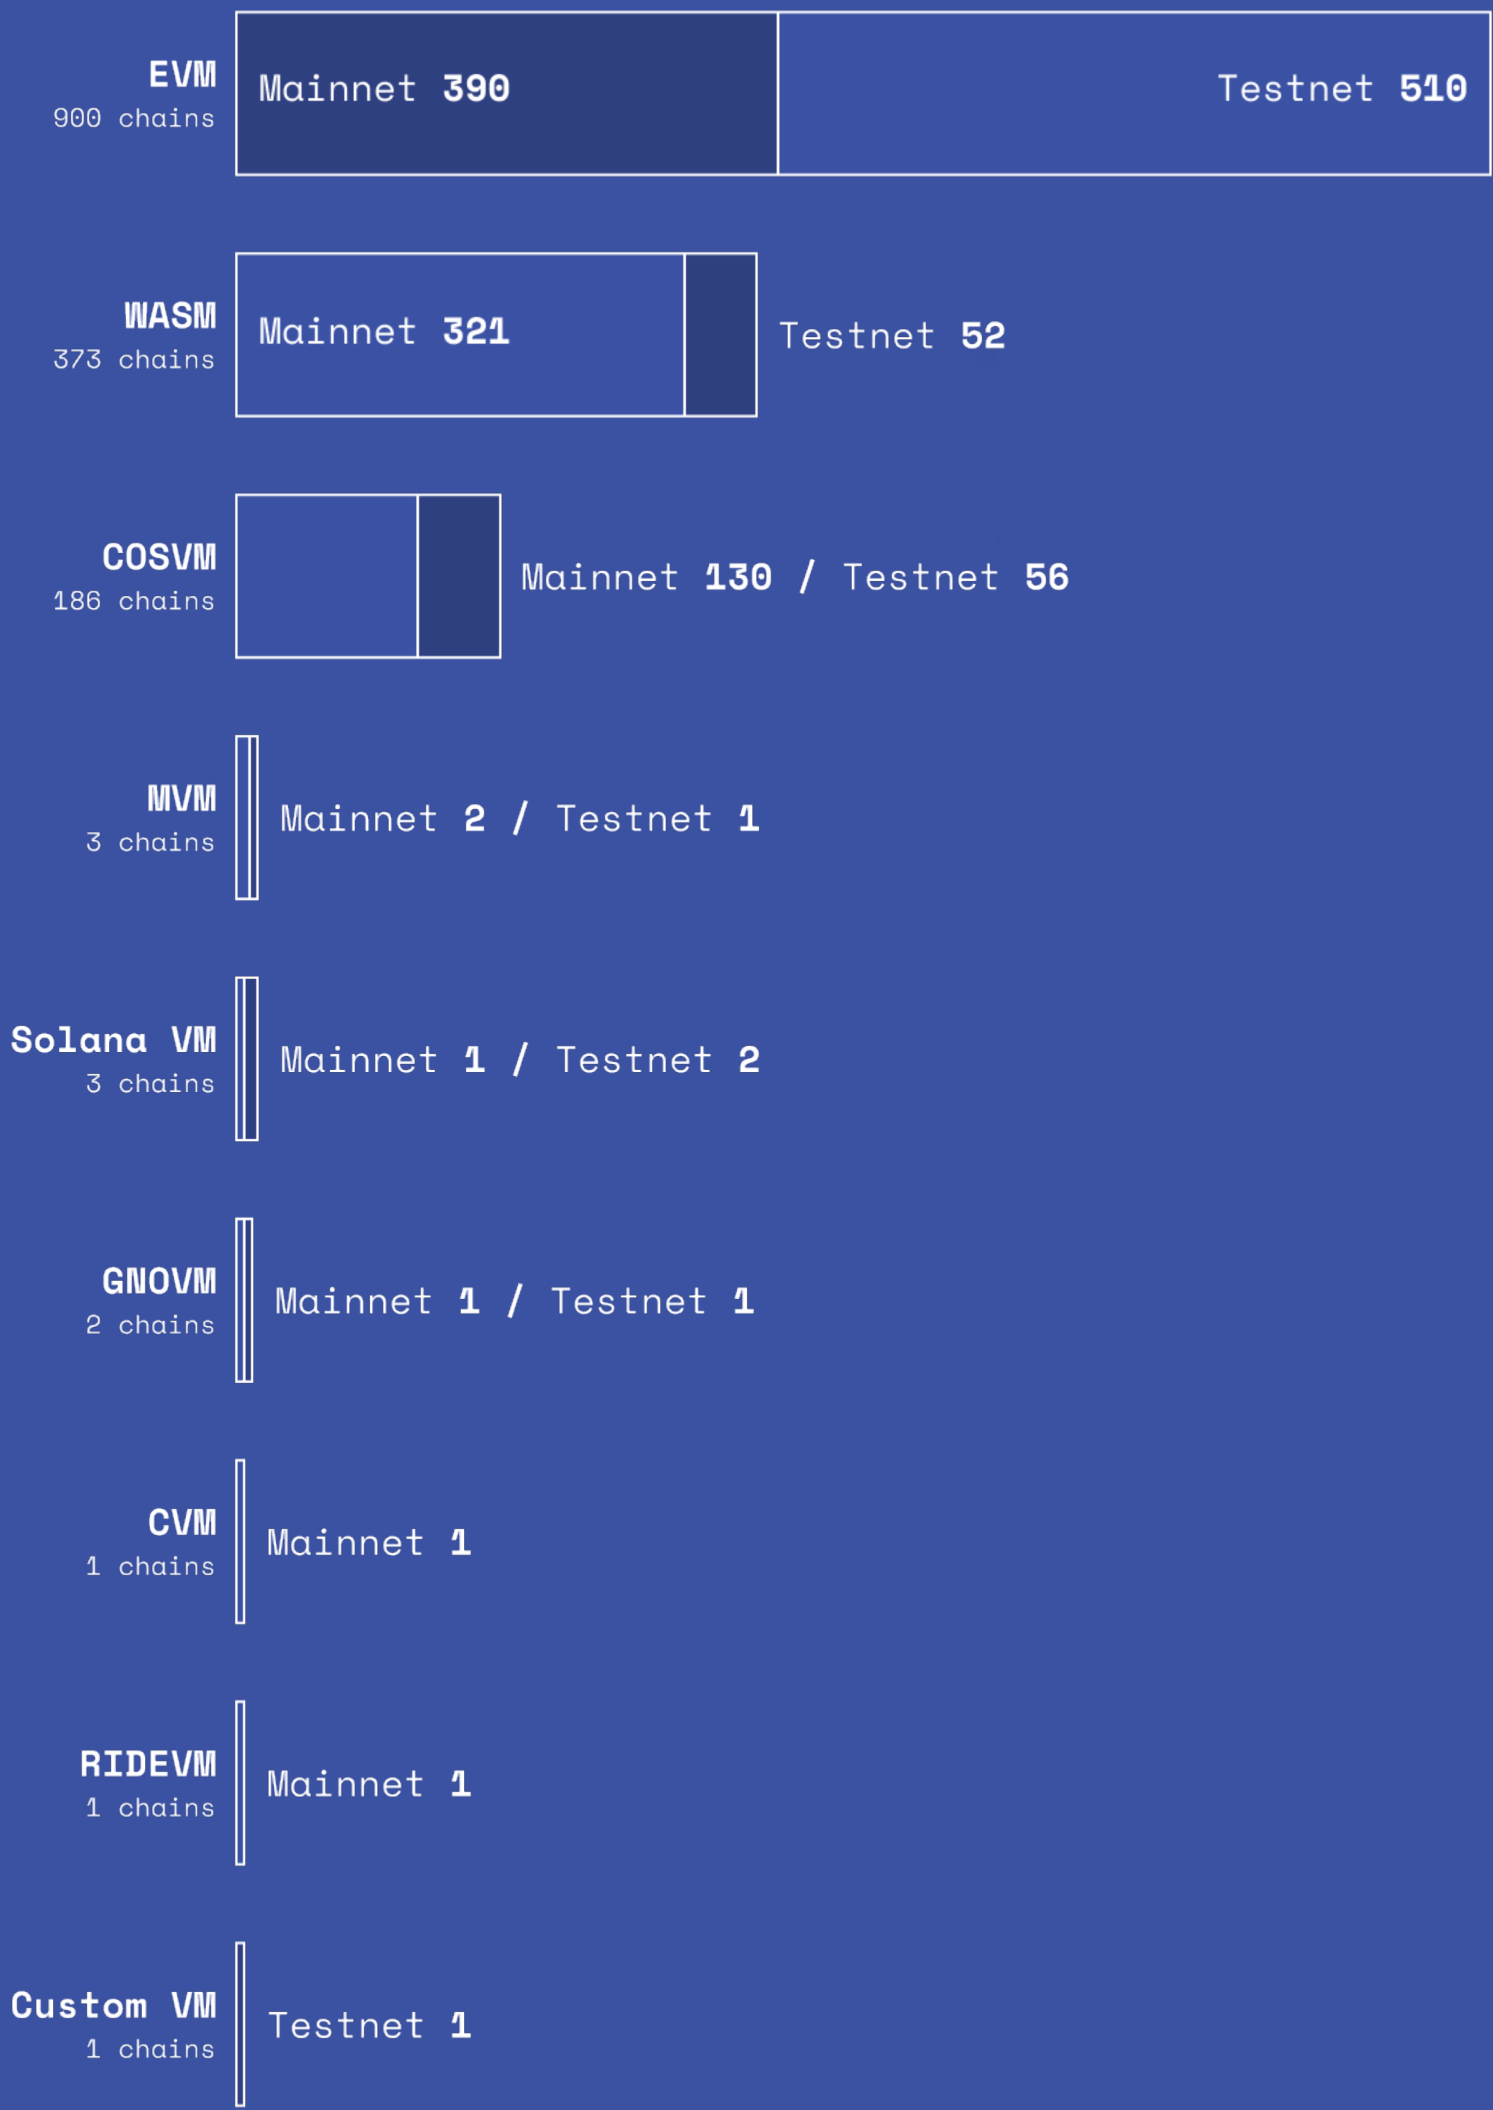
EVM
900 chains
Mainnet 390
Testnet 510
WASM
373 chains
Mainnet 321
Testnet 52
COSVM
186 chains
Mainnet 130 / Testnet 56
MVM
3 chains
Mainnet 2 / Testnet 1
Solana VM
3 chains
Mainnet 1 / Testnet 2
GNOVM
2 chains
Mainnet 1 / Testnet 1
CVM
1 chains
Mainnet 1
RIDEVM
1 chains
Mainnet 1
Custom VM
1 chains
Testnet 1

3.2 Network Type Distribution

Production-First Infrastructure Development

L1 networks achieved strong production majority in 2025, with mainnet deployments significantly outpacing experimental environments.

> This section uses the full L1 dataset (1,540) to classify network environment (mainnet vs testnet), rather than the VM-identified subset used in Section 3.1. Environment status can be validated for all L1 networks, so using the full L1 total maximizes coverage and avoids filtering results based on whether a VM was identified (which isn't required for this metric).

Production Networks

895

mainnet deployments (58.1%)

Staging/Testing Networks

645

testnet environments (41.9%)

Environment Distribution:
This 58.1% mainnet rate exceeds the overall industry average (53.3%) by 4.8 percentage points, indicating L1 infrastructure tilts more toward production deployment than the network universe overall. The remaining 41.9% testnet share indicates continued experimentation on L1, but the balance is decisively production-led rather than testnet-heavy proliferation.

3.3 Ecosystem Distribution

Architectural Framework Consolidation

> Ecosystem shares are computed on the ecosystem-mapped L1 subset (657), rather than the full L1 total (1,540), because ecosystem attribution requires a verified framework mapping that is not available for every network. Restricting the analysis to the mapped subset keeps percentages meaningful (summing to 100% within the classified sample) and avoids extrapolating ecosystem dominance to unmapped L1 networks.

L1 ecosystem distribution reveals strategic positioning around four dominant technical frameworks: Ethereum-based, Polkadot, Cosmos, and Avalanche, which collectively control 87.67% of categorized L1 ecosystem infrastructure. Ethereum-based leads L1 ecosystem categorization with 31.66% share (208 chains), reflecting Ethereum's evolution from a singular Layer 1 toward a settlement anchor and technical standard for numerous sovereign L1s built using its core stack.

Polkadot captures 21.92% share (144 chains), representing Substrate-based chains that leverage its shared security and interoperability framework.

Cosmos achieves 17.66% share (116 chains), with this figure including CosmWasm/Cosmos SDK- based chains that share technical dependencies.

Avalanche maintains 16.44% share (108 chains), driven by its L1s (formerly “subnets”) expansion strategy enabling application-specific sovereign chains.

Kusama (Polkadot’s canary network) represents 10.81% share (71 chains), indicating strong experimental deployment on production-grade Substrate infrastructure.

Other ecosystems like **Ripple, Oasys, Bitcoin, Iota, Nexus,** and **Filecoin** represent distinct, specialized approaches but maintain minimal L1 footprint in categorized data.

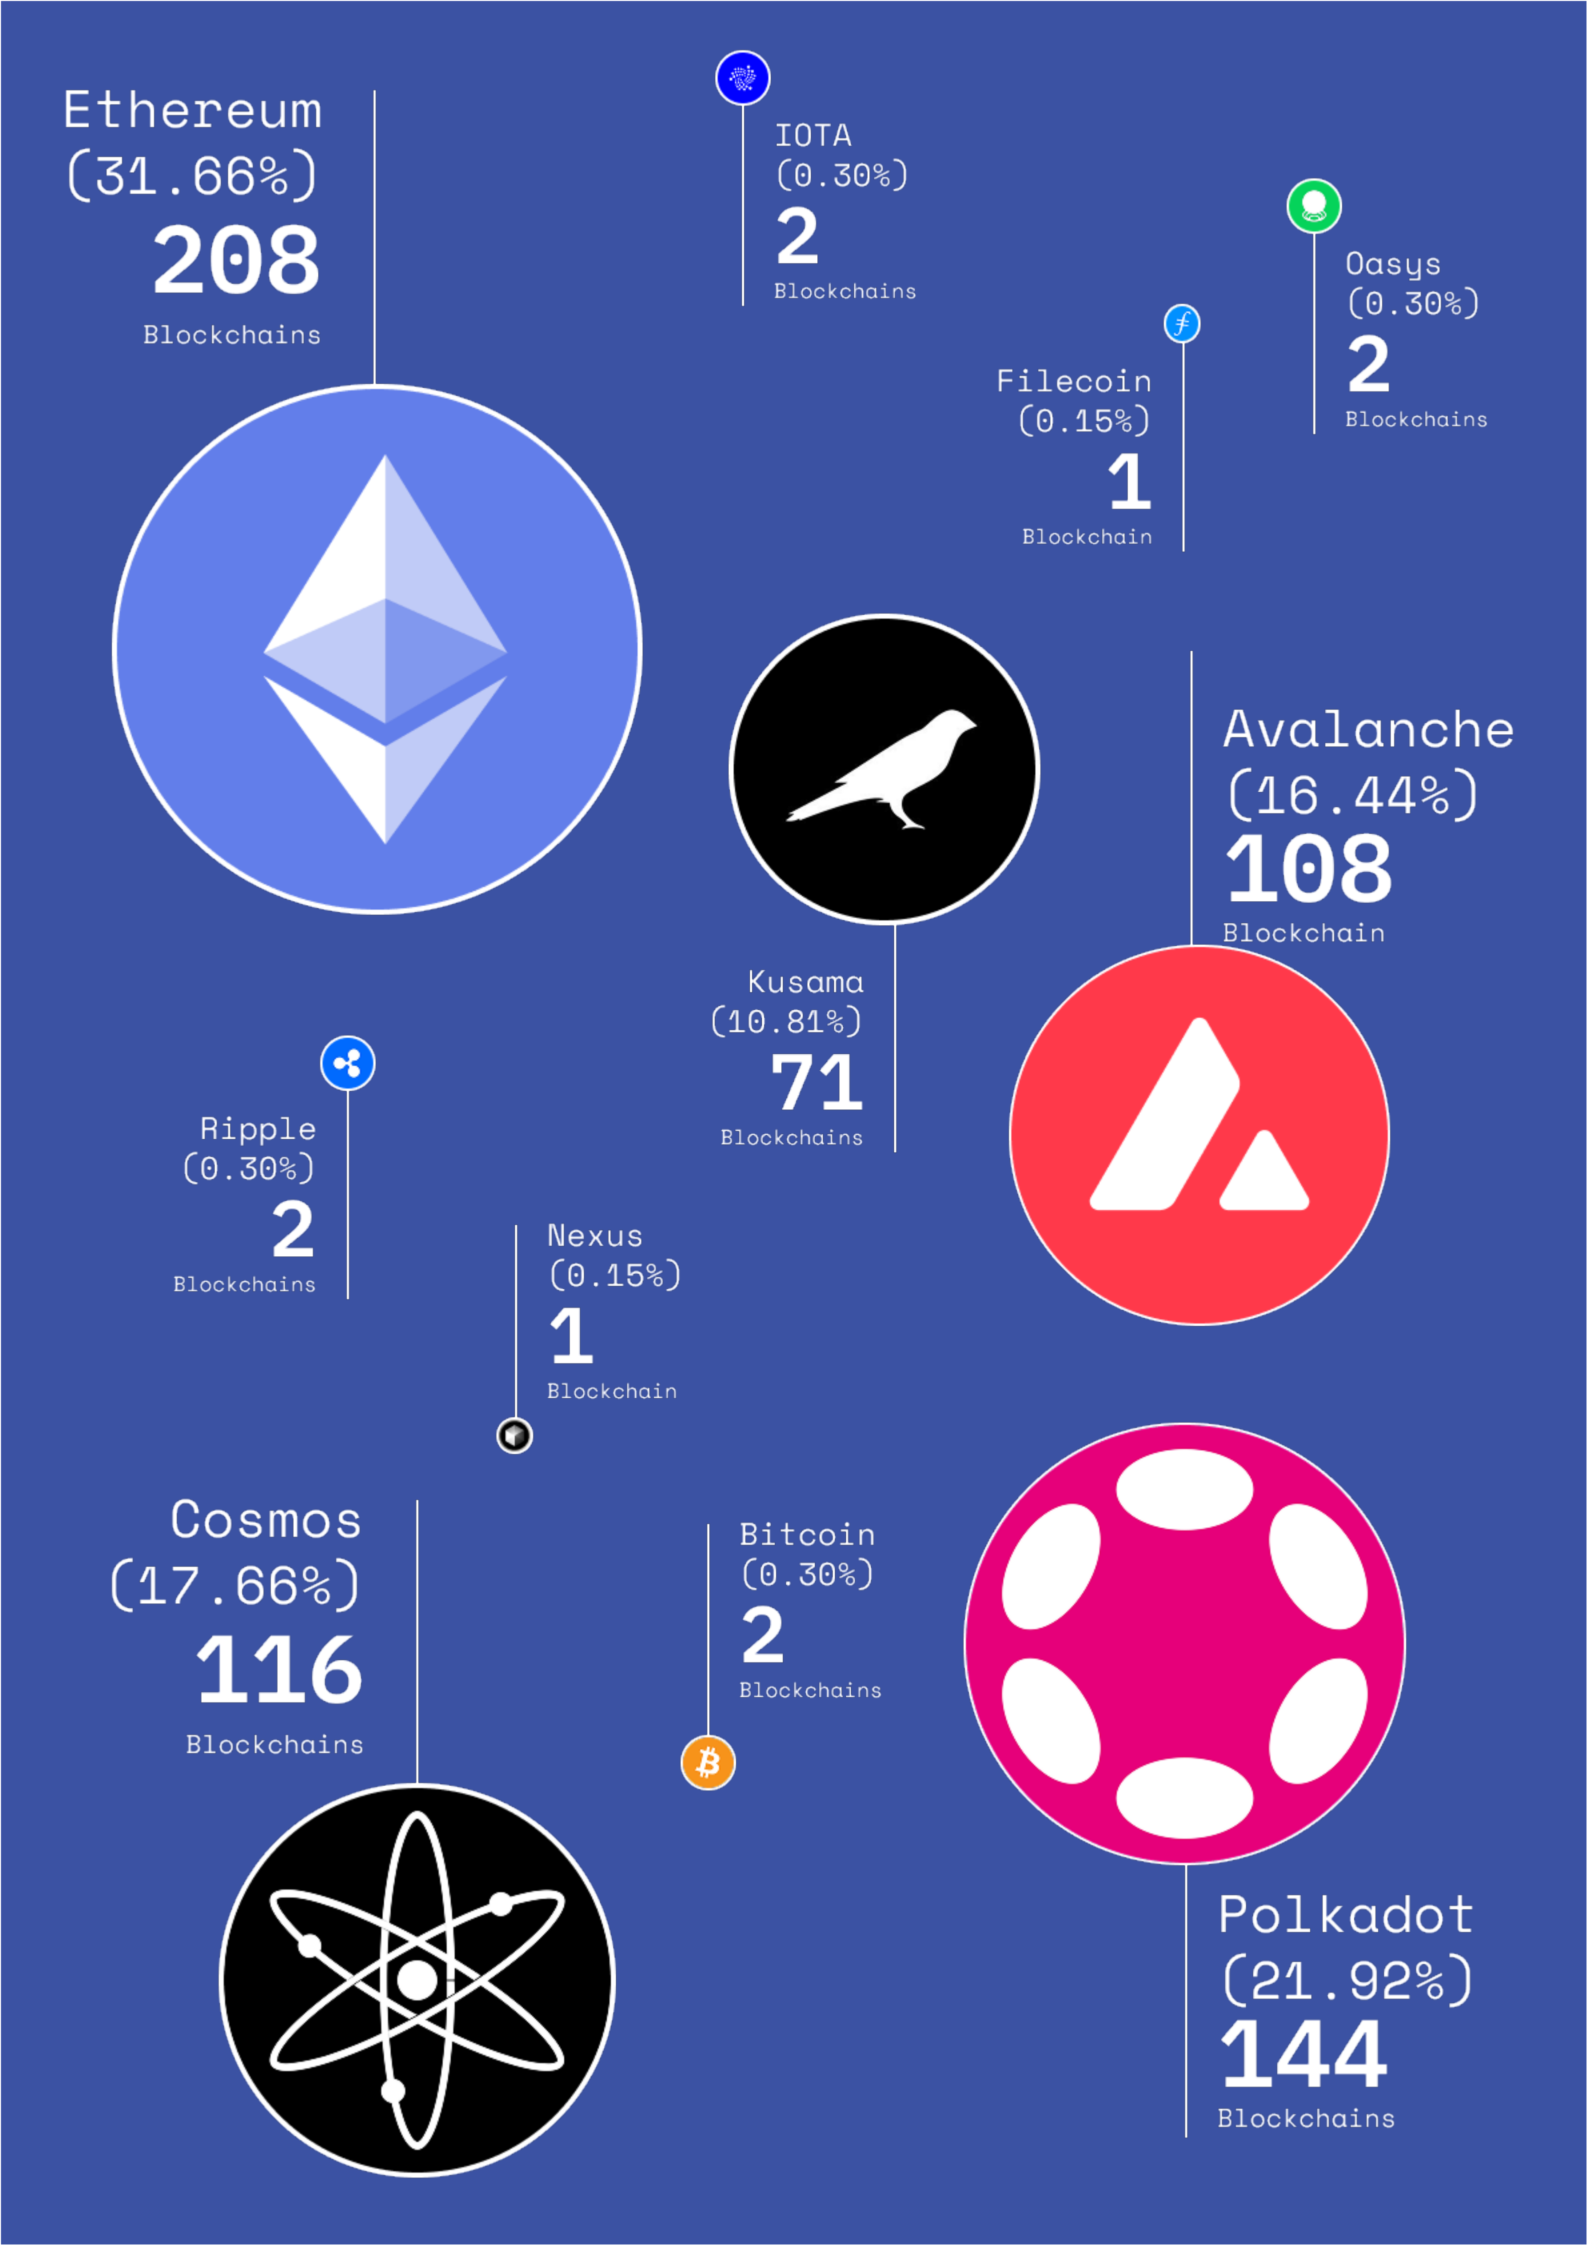

Ethereum (31.66%) 208 Blockchains
IOTA (0.30%) 2 Blockchains
Oasys (0.30%) 2 Blockchains
Filecoin (0.15%) 1 Blockchain
Avalanche (16.44%) 108 Blockchain
Kusama (10.81%) 71 Blockchains
Ripple (0.30%) 2 Blockchains
Nexus (0.15%) 1 Blockchain
Cosmos (17.66%) 116 Blockchains
Bitcoin (0.30%) 2 Blockchains
Polkadot (21.92%) 144 Blockchains

3.4 Role of Block Explorers

Block explorer distribution reveals how observability infrastructure correlates with L1 network maturity and ecosystem development.

> Explorer distribution is measured on validated L1 explorer integrations (1,162), not on the total number of L1 networks (1,540), because this section analyzes observability coverage and provider usage at the integration level. Not every L1 has a validated explorer integration, and some L1s have multiple integrations, so using the integration set provides a consistent denominator for share and adoption patterns.

Explorer Market Leadership Analysis

Blockscout leads L1 explorer support with 325 integrations (195 mainnet, 130 testnet), demonstrating its versatility across diverse L1 architectures and maintaining production-first deployment patterns.

Avalabs captures 117 integrations with heavy testnet focus (92 testnet vs. 25 mainnet), aligning with Avalanche's L1s expansion strategy that prioritizes rapid deployment over immediate production maturity.

Specialized explorers like **Ping Dashboard (88 integrations), Routescan (59 integrations), and Subscan (42 integrations)** indicate that L1 networks with unique consensus mechanisms or ecosystem-specific requirements demand tailored, production-focused observability solutions.

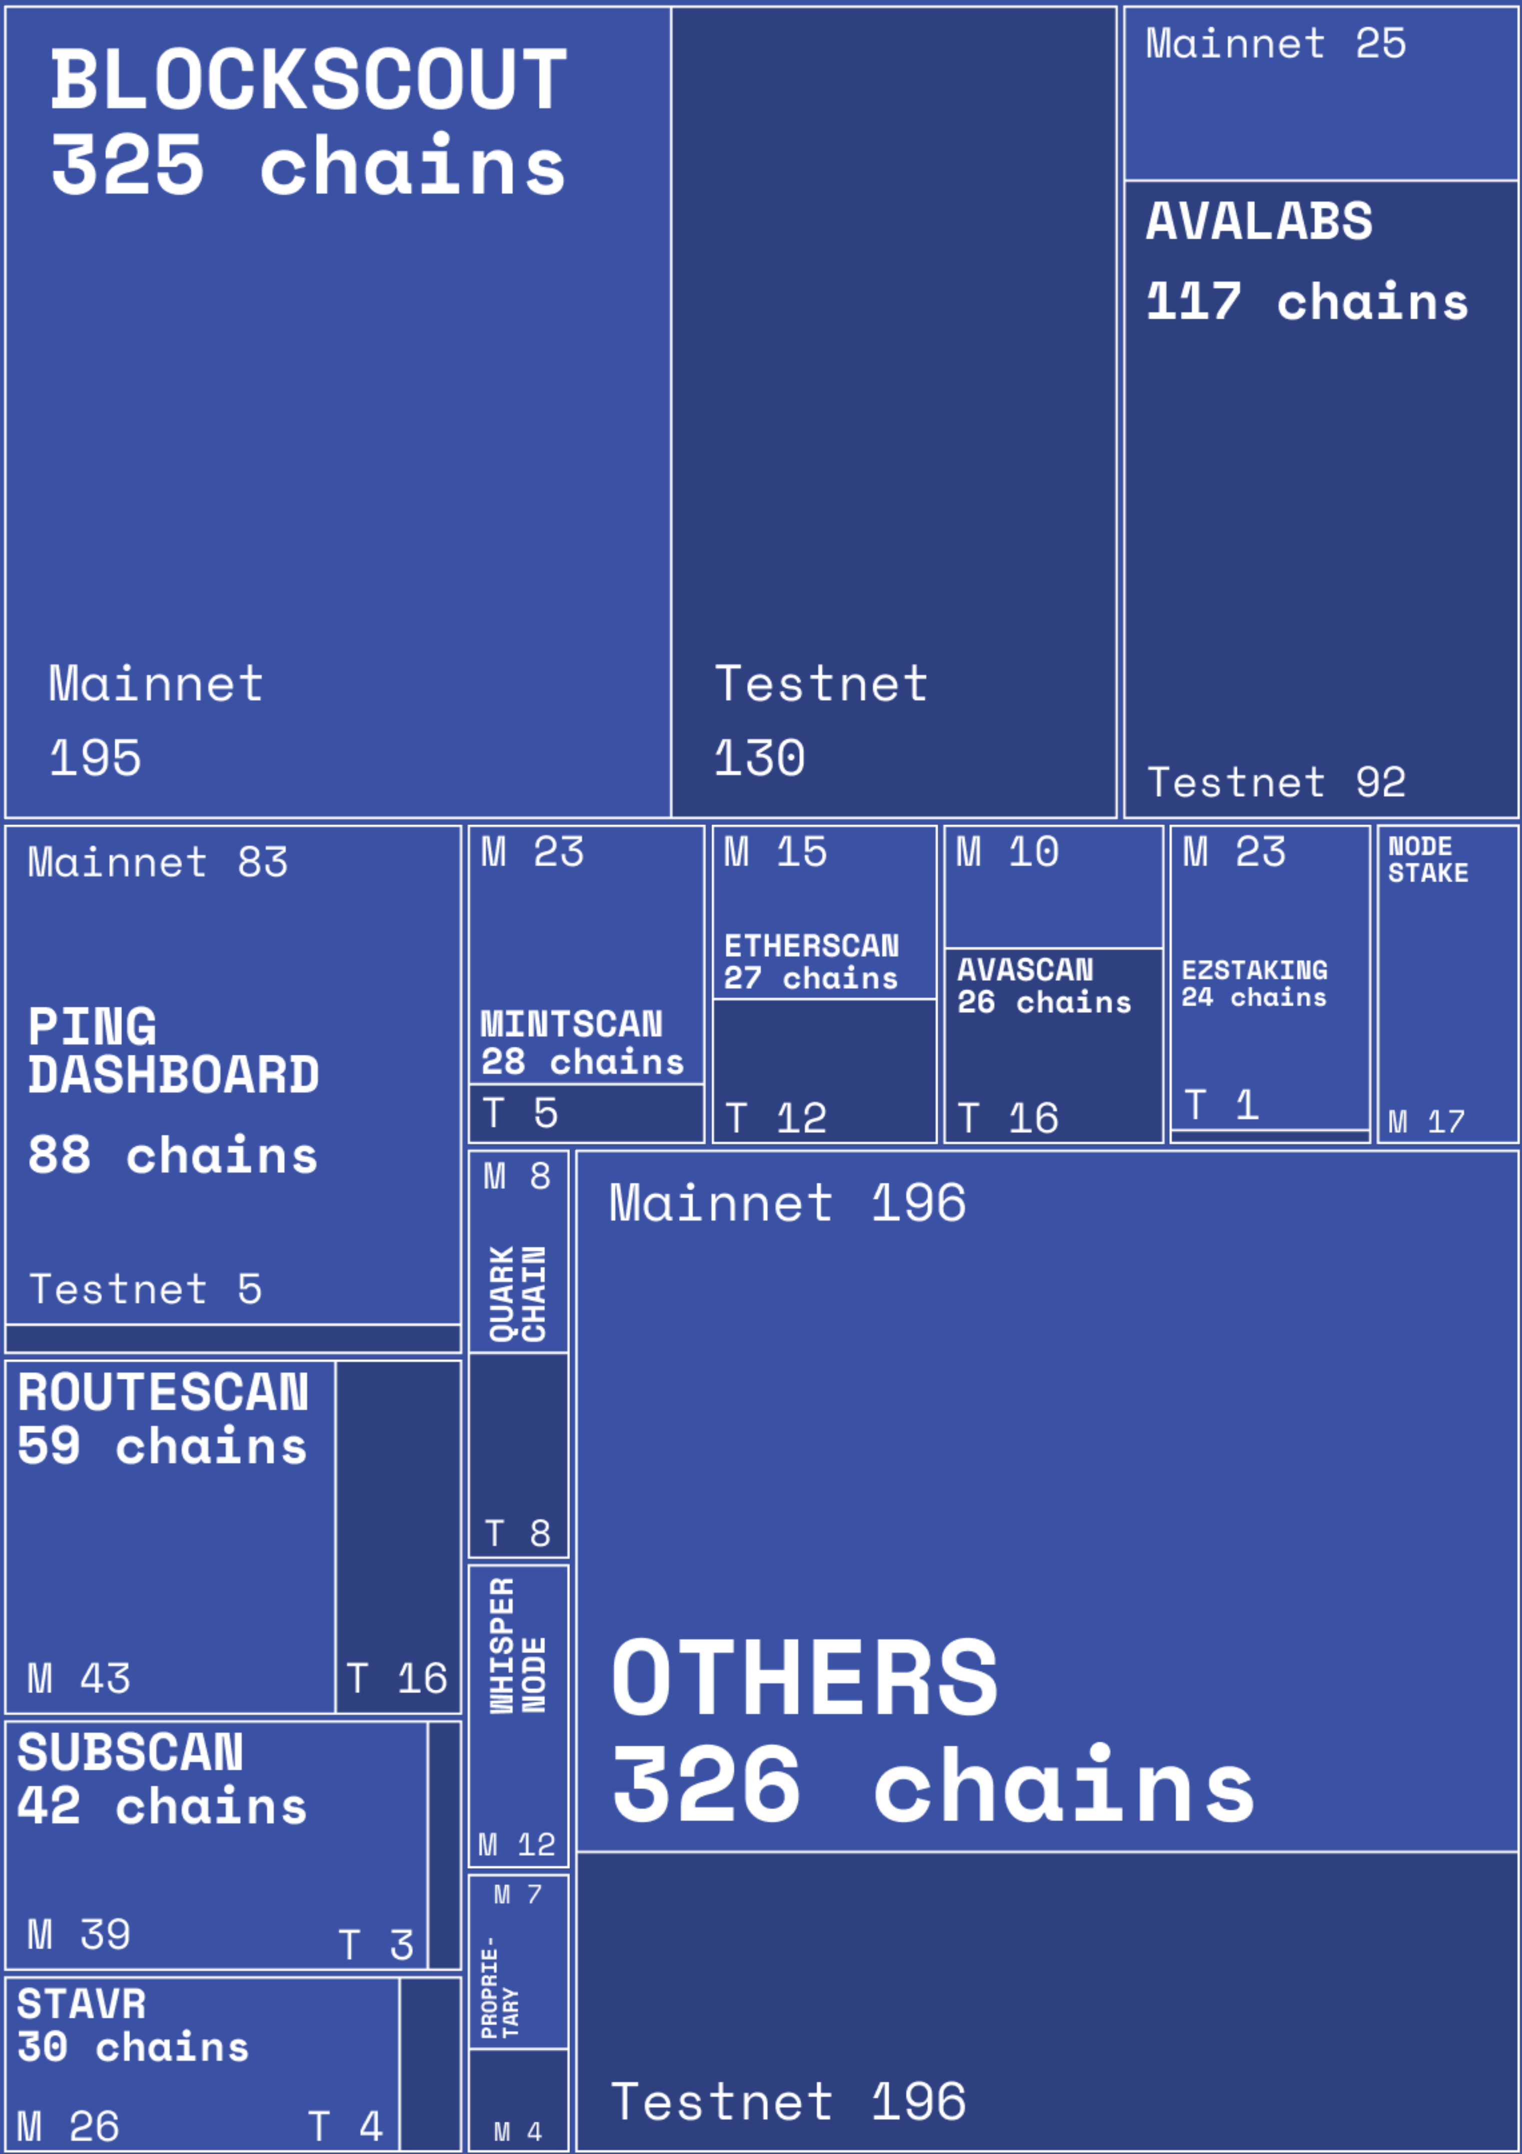
BLOCKSCOUT
325 chains
Mainnet
195
Testnet
130
Mainnet 25
AVALABS
117 chains
Testnet 92
Mainnet 83
PING DASHBOARD
88 chains
Testnet 5
M 23
MINTSCAN
28 chains
T 5
M 15
ETHERSCAN
27 chains
T 12
M 10
AVASCAN
26 chains
T 16
M 23
EZSTAKING
24 chains
T 1
NODE STAKE
M 17
ROUTESCAN
59 chains
M 43
T 16
SUBSCAN
42 chains
M 39
T 3
STAVR
30 chains
M 26
T 4
M 8
QUARK CHAIN
T 8
WHISPER NODE
M 12
M 7
PROPRIE-TARY
M 4
Mainnet 196
OTHERS
326 chains
Testnet 196

3.5 Key Insights: L1 as Sovereignty Infrastructure

L1 shows the clearest "sovereignty vs compatibility" split in 2025: within the VM-identified L1 subset (1,472), EVM is still the majority execution environment (61.14%; 900 chains), but it also has comparatively weak production conversion (43.3% mainnet), sitting 14.3pp below the overall L1 mainnet rate within that same VM-identified sample (57.6%).

In contrast, alternative VMs signal higher production conviction : within the same VM-identified subset, WASM and COSVM together account for 38.0% of L1 VM distribution, and both show much stronger mainnet adoption (86.1% for WASM; 69.9% for COSVM).

On the ecosystem side, concentration is measured only where ecosystem attribution is verified (the ecosystem- mapped L1 subset, 657): within that mapped sample, Ethereum-based (31.66%), Polkadot (21.92%), Cosmos (17.66%), and Avalanche (16.44%) collectively represent 87.67% of categorized L1 ecosystems, which is a framework-level consolidation signal, not a claim about all 1,540 L1 networks.

Explorer distribution is tracked at the integrations level (validated L1 explorer integrations, 1,162 across 269 explorers), because observability is an infrastructure dependency that can vary per network and is not a 1:1 “one chain, one explorer” relationship.

Taken together with the layer totals (L1 = 1,540 out of 2,133 layer-classified networks; 58.1% mainnet on L1), the picture is consistent: L1 remains the primary layer for architectural differentiation, while production-ready adoption clusters disproportionately around non-default execution environments and the frameworks that package them.

Why Ethereum-based Now Leads the L1 Ecosystem Distribution

In the refined 2025 methodology, networks are categorized by primary technical dependencies and framework affiliation, so "Ethereum-based" here is a technical-stack bucket, not a claim that Ethereum "owns" most L1s. Within the ecosystem-mapped L1 subset (657), i.e., only the L1s where ecosystem attribution is validated, 208 chains (31.66%) fall under "Ethereum," reflecting its shift from a standalone L1 to a settlement/security anchor and technical standard that many sovereign L1s still build on or depend on. Unlike Cosmos or Polkadot, where interoperability is an explicit native design (zones/parachains), the "Ethereum" L1 category mainly captures standardization and deep stack reliance; that's why the "lead" should be read as dominance
within the mapped subset, not across the full L1 universe (1,540).

Chapter IV: Layer 2 (L2) Analysis

Layer 2 blockchains are scaling layers built to increase throughput and reduce costs while inheriting security guarantees from an underlying L1.
Our 2025 snapshot includes 512 L2 networks, reaching meaningful scale and showing a balanced production split (51.8% mainnet). L2 is also the layer where execution environments converge the hardest on EVM compatibility, reinforcing “distribution > differentiation” as the dominant strategy.

4.1 Virtual Machine (VM) Distribution

The Scaling Layer Execution Environment Question

How do L2 networks trade off compatibility versus performance specialization when choosing an execution environment for scaling?

> VM shares in this section are computed on the VM-identified L2 subset (467), meaning the analysis only includes the L2 networks where VM attribution is validated; this corresponds to 91.2% coverage of total L2 networks (467/512). The remaining 44 L2 networks are still part of the overall L2 count (512), but they are excluded from VM share calculations because their VM could not be reliably identified at this snapshot, so percentages here should be read as “within the VM-identified subset” (i.e., they sum to 100 inside that subset).

VM Distribution Analysis

The 2025 data shows L2 has the highest EVM concentration across layers at 95.94% (450 chains), prioritizing developer compatibility and migration paths. Specialized L2 VMs appear at very small counts but often ship straight to mainnet (e.g., CAIROVM, ARBITRUMVM, SOLANAVM show 100% mainnet within their tiny sample sizes). This pattern implies that “non-EVM on L2” is typically a deliberate, high-conviction engineering choice rather than broad ecosystem experimentation.
Overall, L2 still shows some VM diversity (11 VM types identified), but it is statistically dominated by EVM to a degree that makes “EVM-first tooling” the operational default.

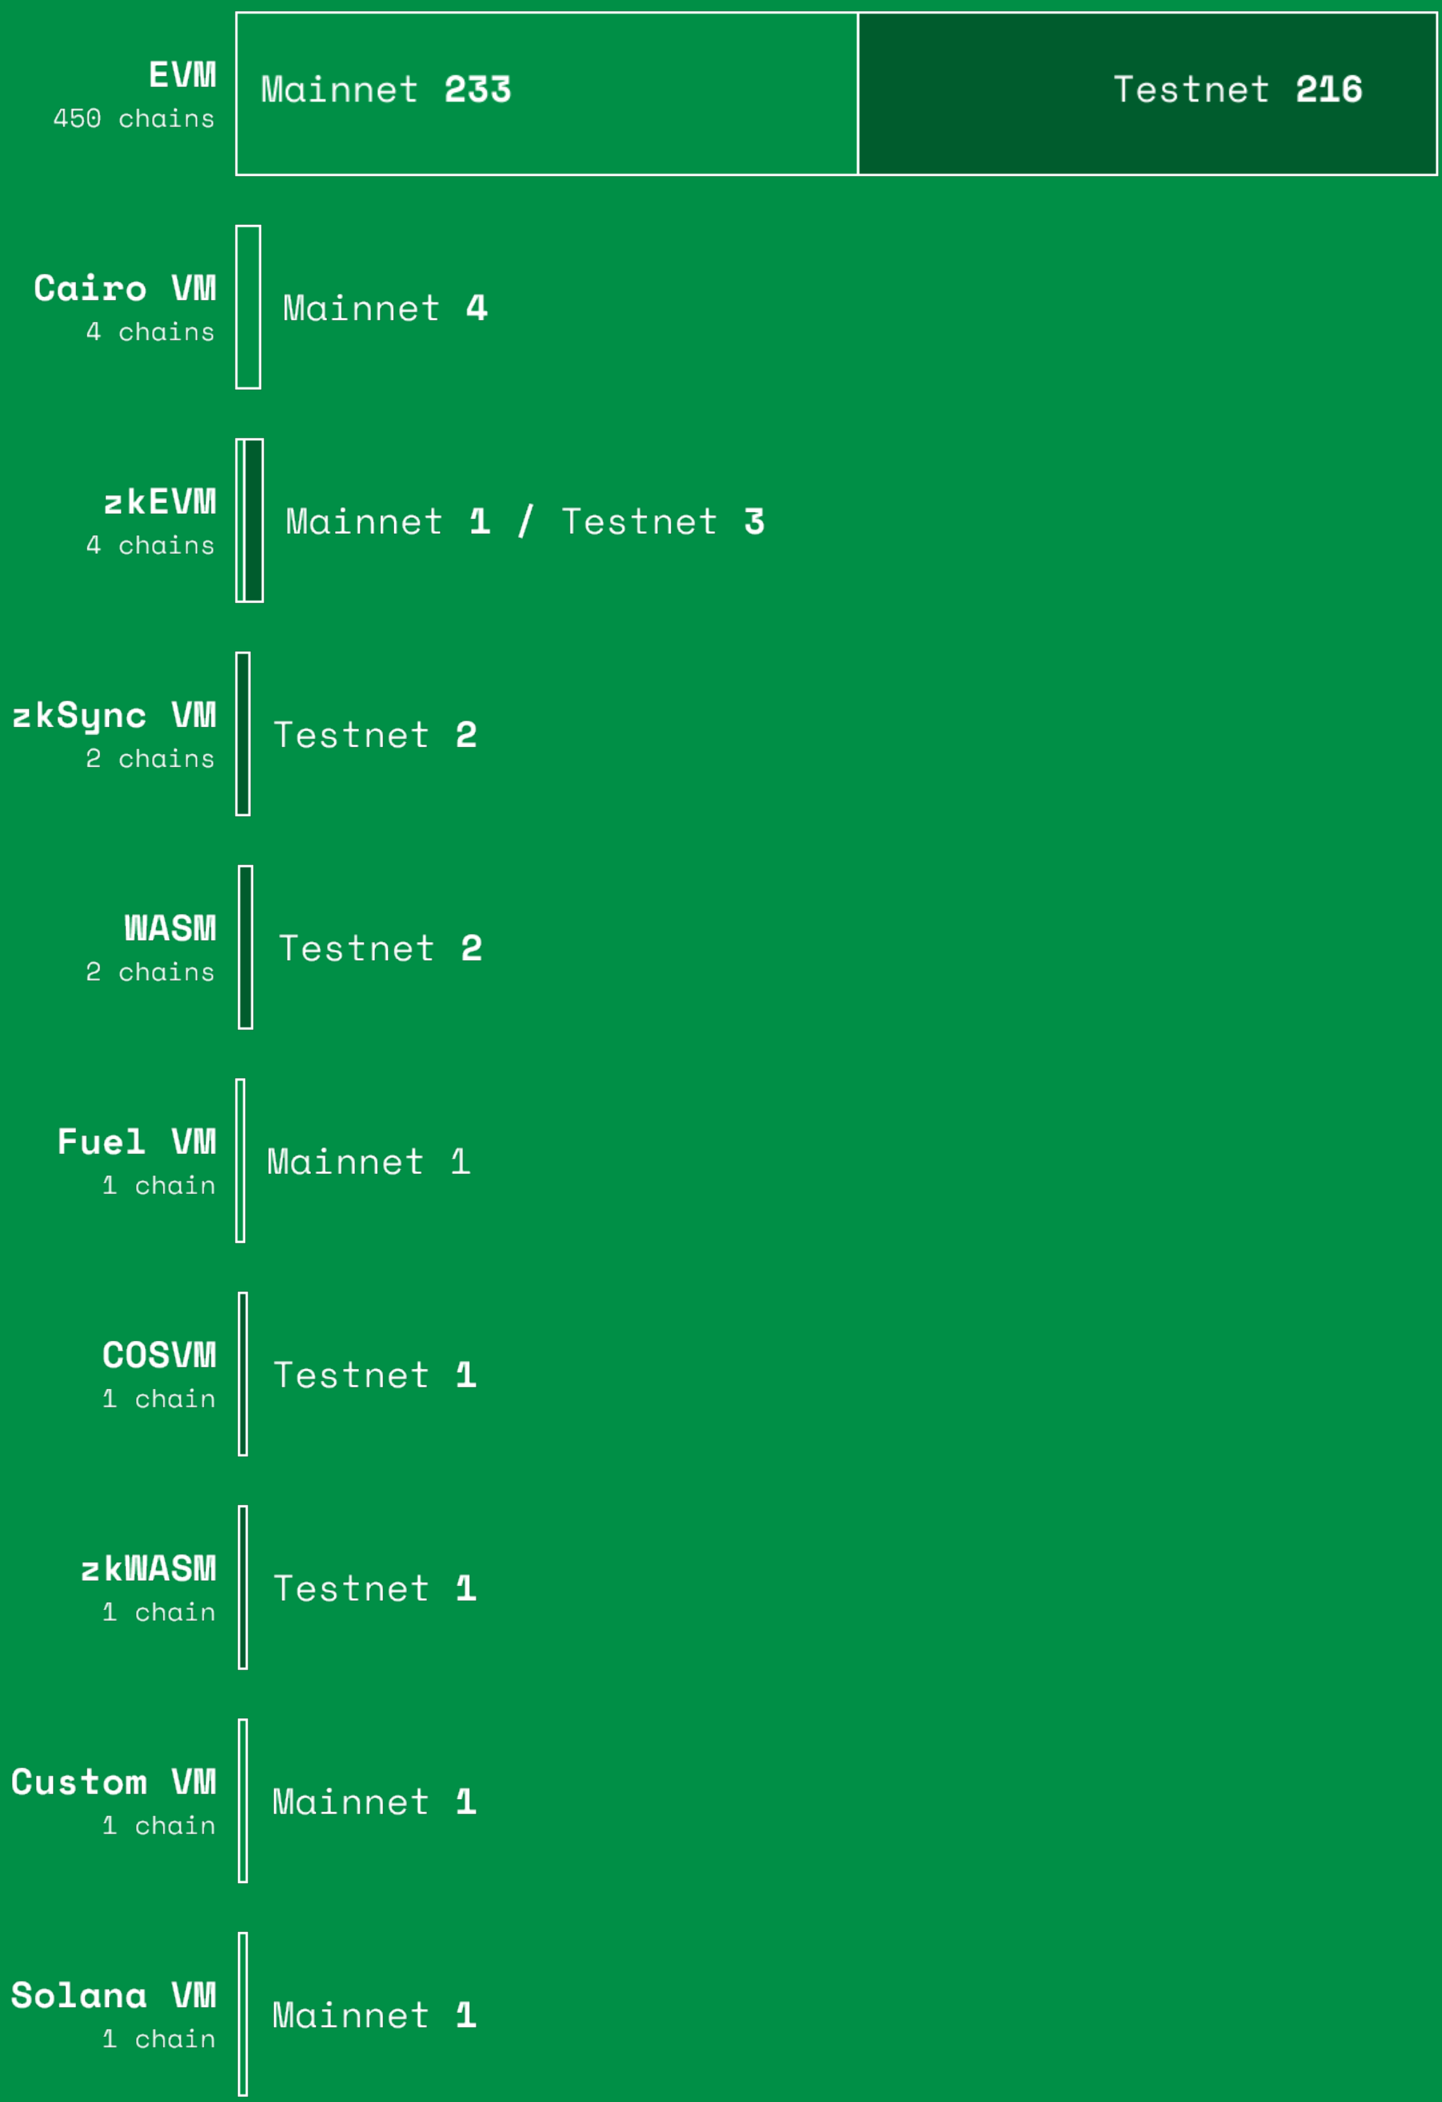

EVM
450 chains
Mainnet 233
Testnet 216
Cairo VM
4 chains
Mainnet 4
zkEVM
4 chains
Mainnet 1 / Testnet 3
zkSync VM
2 chains
Testnet 2
WASM
2 chains
Testnet 2
Fuel VM
1 chain
Mainnet 1
COSVM
1 chain
Testnet 1
zkWASM
1 chain
Testnet 1
Custom VM
1 chain
Mainnet 1
Solana VM
1 chain
Mainnet 1

4.2 Rollup as a Service (RaaS) Provider Distribution

The Rise of Managed RaaS Infrastructure

How much of the L2 deployment landscape is consolidating around managed infrastructure providers, versus teams choosing to self-host and operate their own stack?

> RaaS market shares here are computed on the RaaS-deployed subset (128), only L2 networks where a managed RaaS relationship is identified and attributable to one of 10 tracked providers. These percentages answer “within RaaS-run L2s, which provider dominates?”, not “what share of all L2s uses RaaS?”, because the full L2 universe is n512 and using it would mix managed deployments with self-hosted and any cases not attributed to a tracked provider.
>
> For that reason, managed vs self-hosted is anchored on the full L2 set: providers serve 128 networks; self-hosted is 13.3% (68 chains), while provider competition is evaluated only within the 128-network subset.

RaaS Market Distribution Analysis

Caldera and **Conduit** emerge as co-leaders in production-grade RaaS deployment:

- **Caldera** (27.34% share, 35 chains) demonstrates strong production focus with 77.1% mainnet (27 mainnet / 8 testnet)
- **Conduit** (26.56% share, 34 chains) shows even stronger production commitment at 82.4% mainnet (28 mainnet / 6 testnet)

Gelato (15.63%, 20 chains) forms a clear third-tier operator with balanced deployment mix (13 mainnet / 7 testnet, 65% mainnet).

Top-3 RaaS control is 69.53%, this represents meaningful concentration without monopolistic lock-in, and critically, all three leaders are production-oriented, unlike previous market structures where volume leaders were testnet-heavy.

Beyond the top tier:

- **AltLayer** (10.94%, 14 chains) and Alchemy (7.03%, 9 chains) form a second cluster of meaningful operators
- Remaining providers (**Zeeve, Lumoz, Ankr, Gateway, Node Infra**) serve narrow, specialized demand with 1-9 chains each.

Strategic implications

Self-hosted L2 deployments (68 chains, 13.3% of L2) remain the real counterweight to RaaS consolidation. These represent teams explicitly paying ops cost to retain sequencing/control and avoid provider lock-in. The 69.53% top-3 control creates a balanced dependency landscape: concentration is real but not absolute, and unlike earlier periods where volume leaders were testnet-heavy, the current top tier is production-focused. Caldera (77.1% mainnet) and Conduit (82.4% mainnet) both demonstrate materially stronger production readiness than the 51.8% L2 average.

Provider risk remains structural but distributed: with near-parity between Caldera and Conduit, the RaaS market exhibits duopoly dynamics at the production tier rather than single-vendor dominance. This creates competitive pressure on pricing, SLAs, and feature velocity while still concentrating dependency risk within a small operator set.

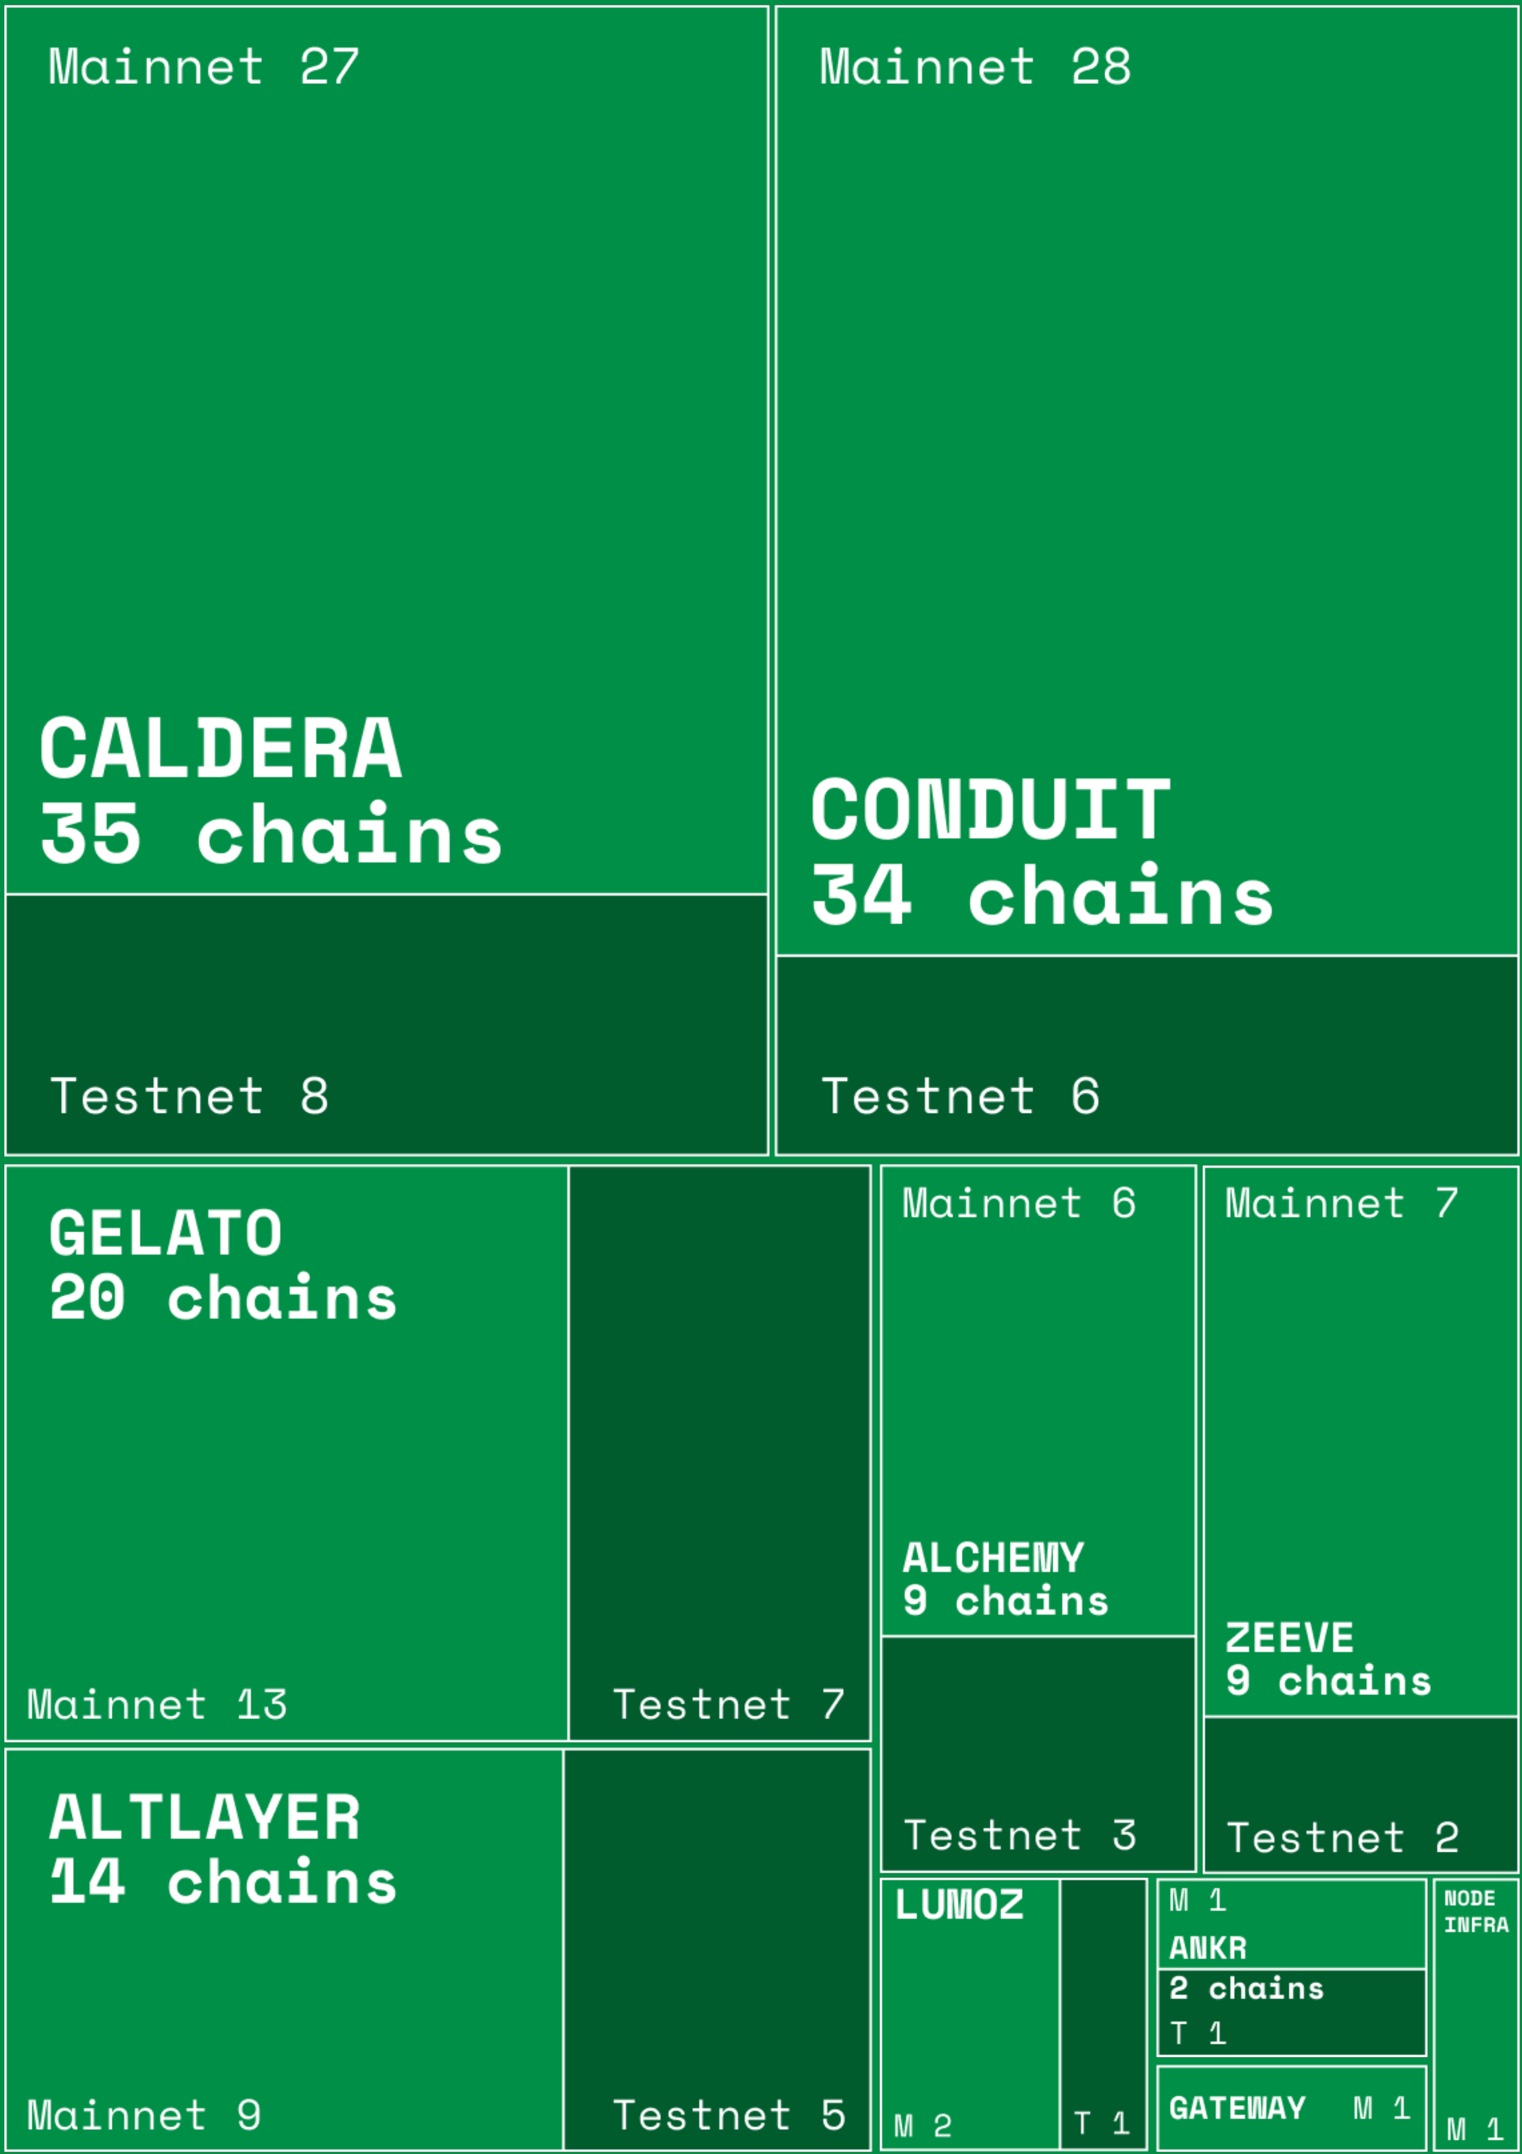

Mainnet 27
CALDERA
35 chains
Testnet 8
Mainnet 28
CONDUIT
34 chains
Testnet 6
GELATO
20 chains
Mainnet 13
Testnet 7
Mainnet 6
ALCHEMY
9 chains
Testnet 3
Mainnet 7
ZEEVE
9 chains
Testnet 2
ALTLAYER
14 chains
Mainnet 9
Testnet 5
LUMOZ
M 2
T 1
M 1
ANKR
2 chains
T 1
GATEWAY
M 1
NODE
INFRA
M 1

4.3 Network Type Distribution

Production readiness vs experimentation

> This section deliberately switches back to the full L2 total (512), because mainnet vs testnet is a baseline environment attribute that can be classified even when VM identification is not available (which is why Section 4.1 uses the VM-identified subset, 468). Put differently: 4.1 measures what the L2 runs (only where VM attribution is validated), while 4.3 measures what stage the network is in (mainnet/testnet), so it can include all L2s without introducing VM-coverage bias.

L2 shows a near-balanced split between production and staging deployments, with 51.8% mainnet and48.2% testnet across 512 networks. Operationally, this means L2 is scaling in production while still maintaining a large experimentation surface, so "L2 maturity" is real, but it's not a finished market.

Production Networks

265

mainnet deployments (51.8%)

Staging/Testing Networks

247

testnet environments (48.2%)

4.4 Ecosystem Distribution

Partial consolidation, long tail remains

> Ecosystem shares in this section are computed on the ecosystem-mapped L2 subset (344 across 18 ecosystems), because ecosystem attribution is only reported when a network's primary settlement/framework affiliation can be validated with enough confidence to avoid mislabeling hybrid or ambiguously described stacks.

The L2 ecosystem distribution is heavily concentrated around Ethereum-aligned stacks, with the top three (Ethereum, Optimism, Arbitrum) accounting for 74.13% of categorized L2 chains. This indicates real settlement/framework consolidation rather than superficial tagging noise.

L2 ecosystem breakdown

Ethereum leads with 46.80% (161 chains), confirming L2 as primarily "Ethereum-settled" scaling infrastructure. Optimism (17.73%, 61 chains) and Arbitrum (9.59%, 33 chains) follow as the main rollup families absorbing most of the mapped L2 supply. Superchain (6.98%, 24 chains) appears as a distinct bucket, separating OP-related collective infrastructure under its own label.

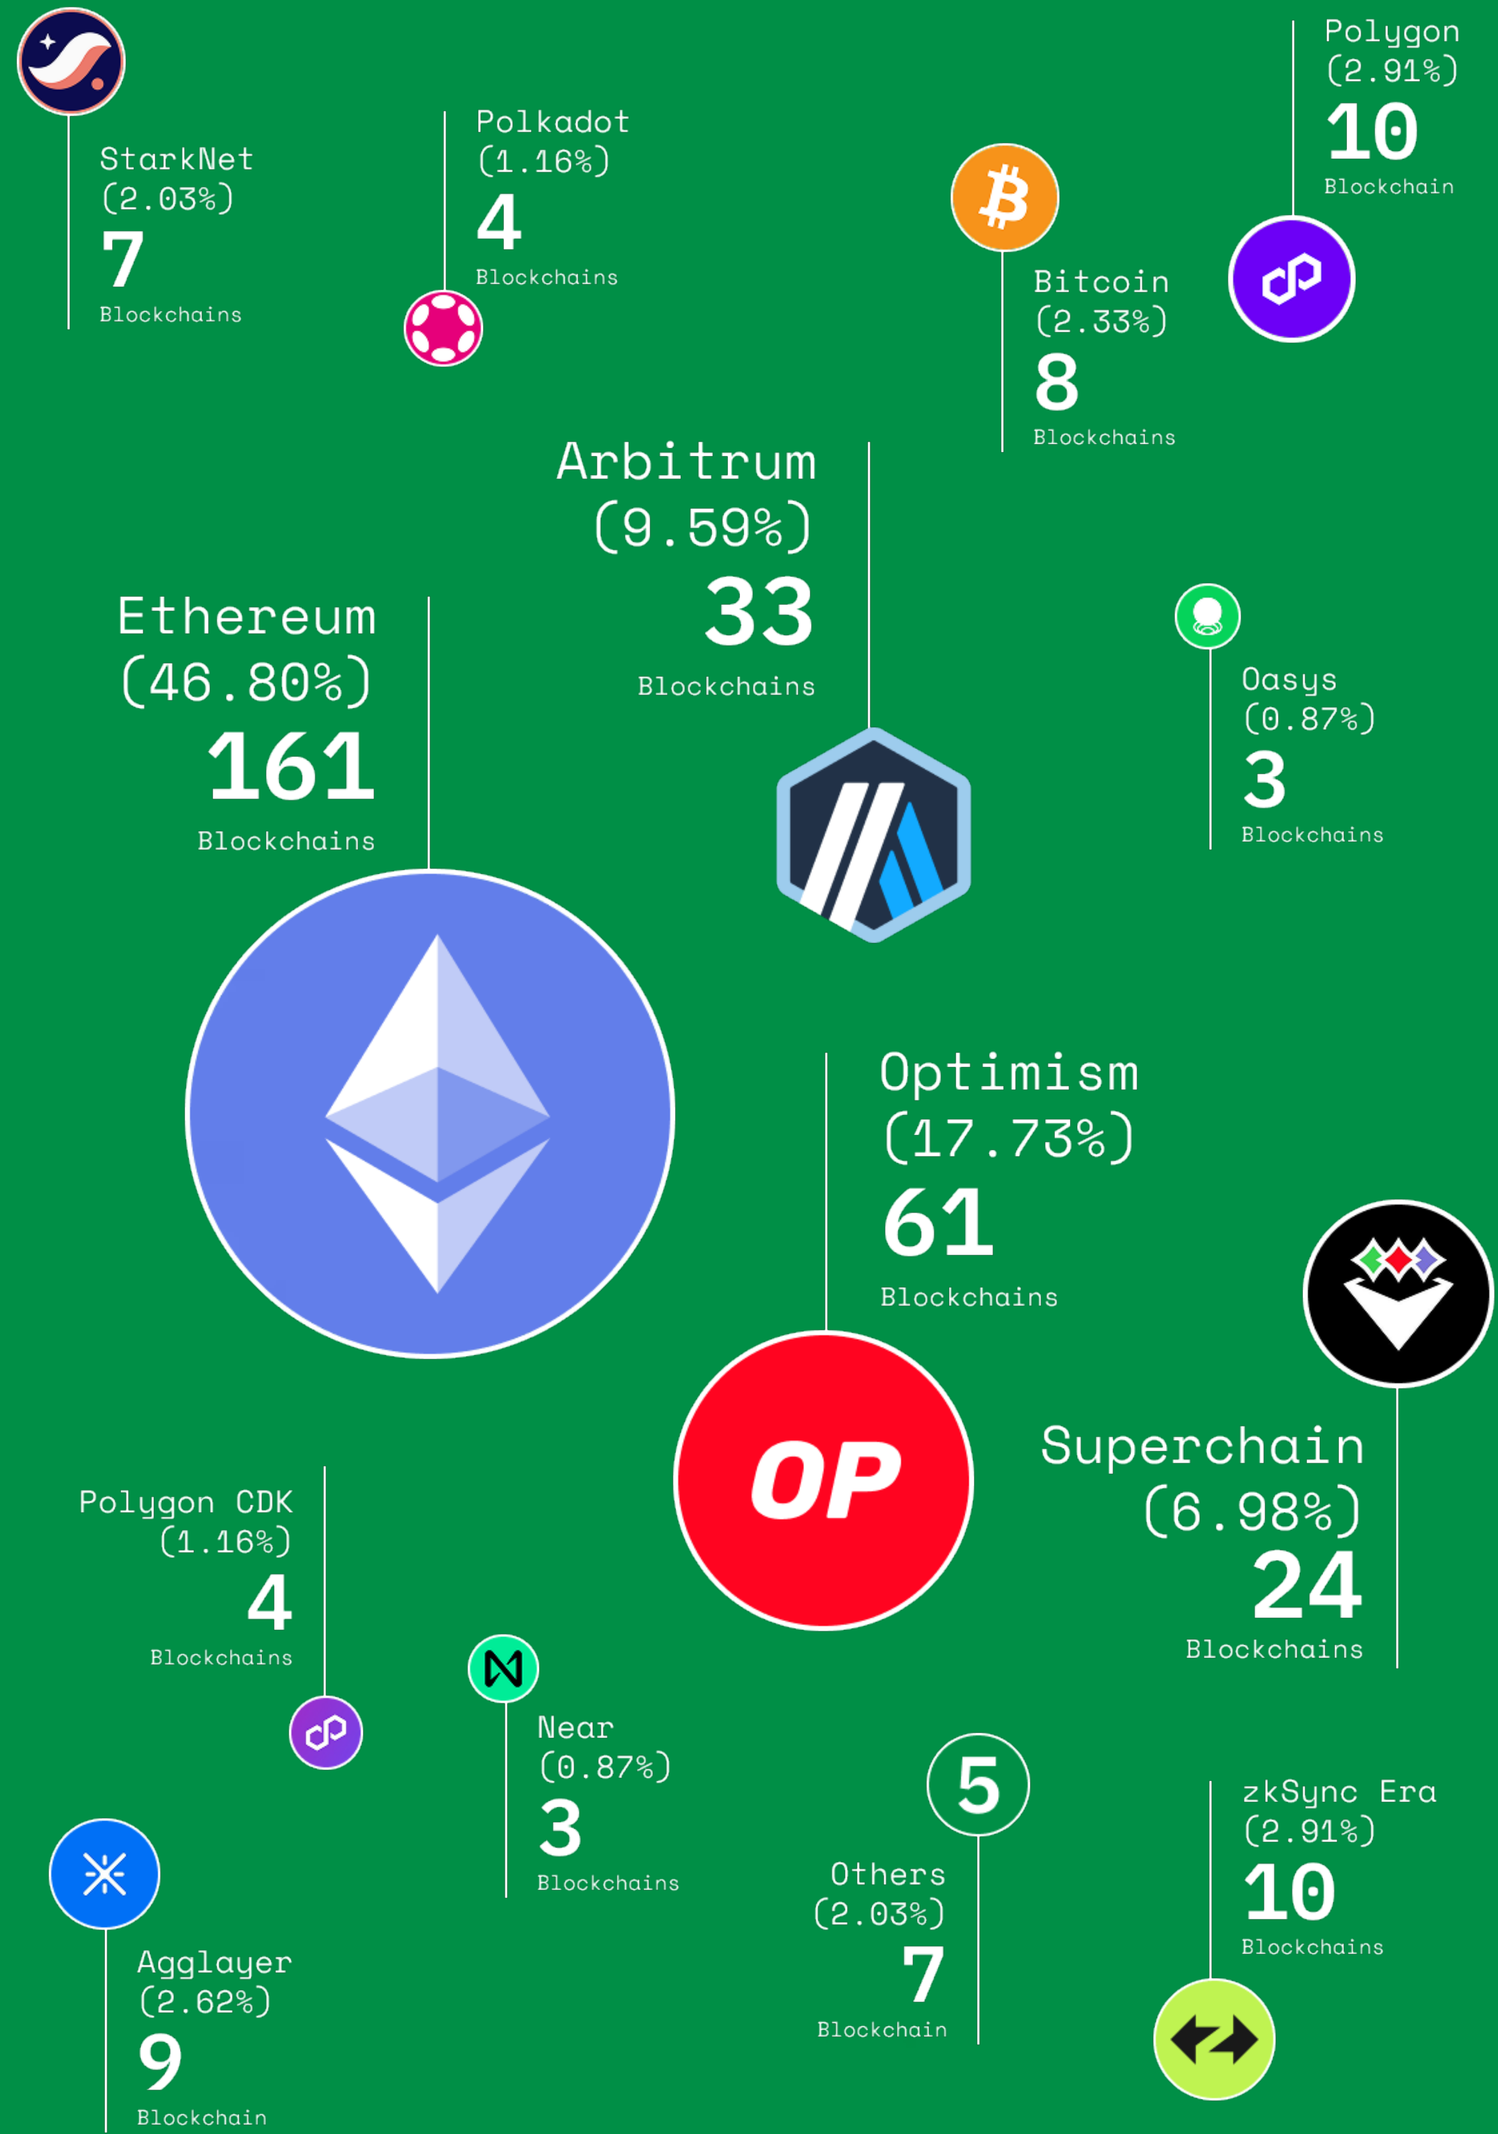

StarkNet
(2.03%)
7
Blockchains
Polkadot
(1.16%)
4
Blockchains
Bitcoin
(2.33%)
8
Blockchains
Polygon
(2.91%)
10
Blockchain
Arbitrum
(9.59%)
33
Blockchains
Ethereum
(46.80%)
161
Blockchains
Oasys
(0.87%)
3
Blockchains
Optimism
(17.73%)
61
Blockchains
OP
Superchain
(6.98%)
24
Blockchains
Polygon CDK
(1.16%)
4
Blockchains
Near
(0.87%)
3
Blockchains
5
Others
(2.03%)
7
Blockchain
zkSync Era
(2.91%)
10
Blockchains
Agglayer
(2.62%)
9
Blockchain

4.5 Framework Distribution

Standardized stacks, not bespoke engineering

> Framework shares in this section are computed on framework attributions identified across L2 networks (661 across 8 frameworks), not on the total number of L2 networks (512). This is intentional: a single L2 can map to more than one framework/component (e.g., a base framework plus a specific rollup stack), so counting
>
> attributions captures stack standardization without forcing a misleading "one network = one framework" assumption. As a result, percentages here should be read as "share of identified framework attributions within L2," rather than "share of L2 networks."

The framework layer is concentrated: Substrate (220 attributions, 33.28%), CosmWasm (175, 26.48%), and Cosmos SDK (130, 19.67%) account for 79.43% of framework attributions, while OP Stack (63, 9.53%) is the largest rollup-oriented framework bucket, followed by Orbit (29, 4.39%) and Nitro (27, 4.08%).

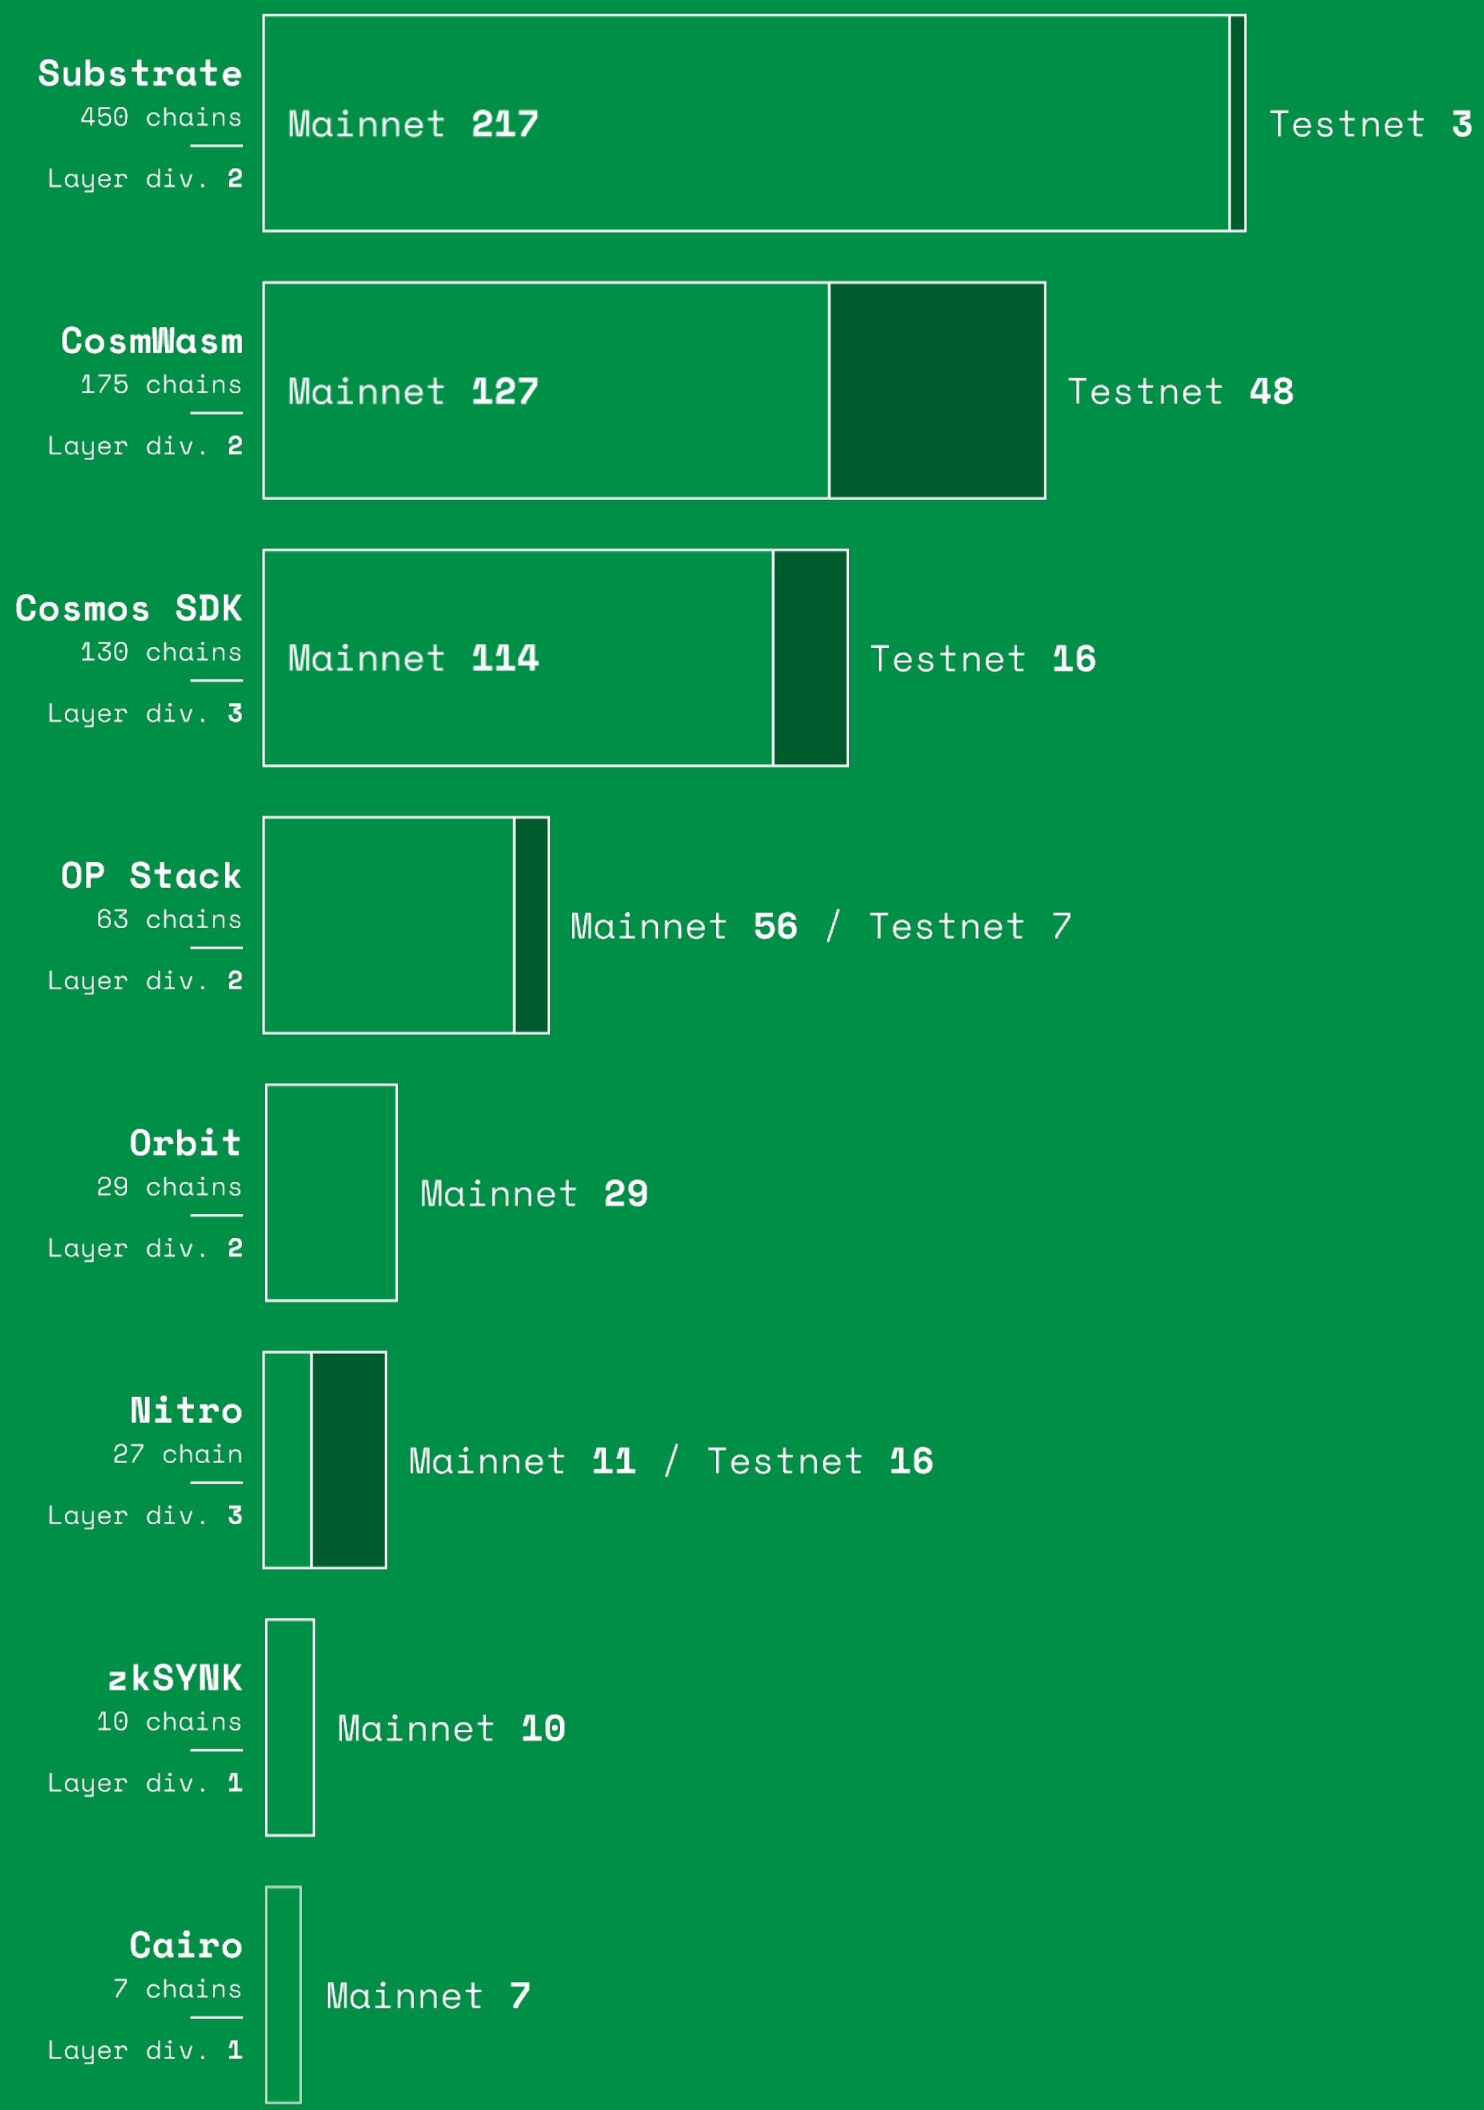
Substrate
450 chains
Layer div. 2
Mainnet 217
Testnet 3
CosmWasm
175 chains
Layer div. 2
Mainnet 127
Testnet 48
Cosmos SDK
130 chains
Layer div. 3
Mainnet 114
Testnet 16
OP Stack
63 chains
Layer div. 2
Mainnet 56 / Testnet 7
Orbit
29 chains
Layer div. 2
Mainnet 29
Nitro
27 chain
Layer div. 3
Mainnet 11 / Testnet 16
zkSYNK
10 chains
Layer div. 1
Mainnet 10
Cairo
7 chains
Layer div. 1
Mainnet 7

4.6 Role of Block Explorers

Observability is still centralized

Block explorer distribution reveals how L2 scaling infrastructure demands specialized observability solutions optimized for rollup architectures and multi-layer data availability.

> The L2 market is served by 114 different explorers across 453 total integrations, and this section intentionally uses
>
> integrations (453) as the unit of analysis because "explorer adoption" is not a 1:1 mapping with "number of L2 networks." One explorer can support many L2s, and a single L2 can also have multiple explorer integrations (e.g., different environments or parallel implementations), so using total L2 networks (512) would blur provider concentration and misstate how observability is actually distributed.

Explorer Market Leadership Analysis

Blockscout leads L2 explorer support with 216 integrations (108 mainnet, 108 testnet), demonstrating perfect production-development balance and maintaining its position as the most versatile multi-layer explorer platform.
Routescan captures 68 integrations (41 mainnet, 27 testnet) with production-first focus, indicating strong positioning in production L2 observability.
Etherscan achieves 45 integrations (22 mainnet, 23 testnet) with balanced deployment, leveraging its Ethereum ecosystem brand recognition. Together, the top three explorers account for 329 of 453 L2 integrations (72.6%), so "fragmentation" exists at the provider count level but not at the integration share level.
The remaining 111 explorers represent specialized solutions, proprietary implementations, and ecosystem-specific tools, indicating significant fragmentation in L2 observability infrastructure.

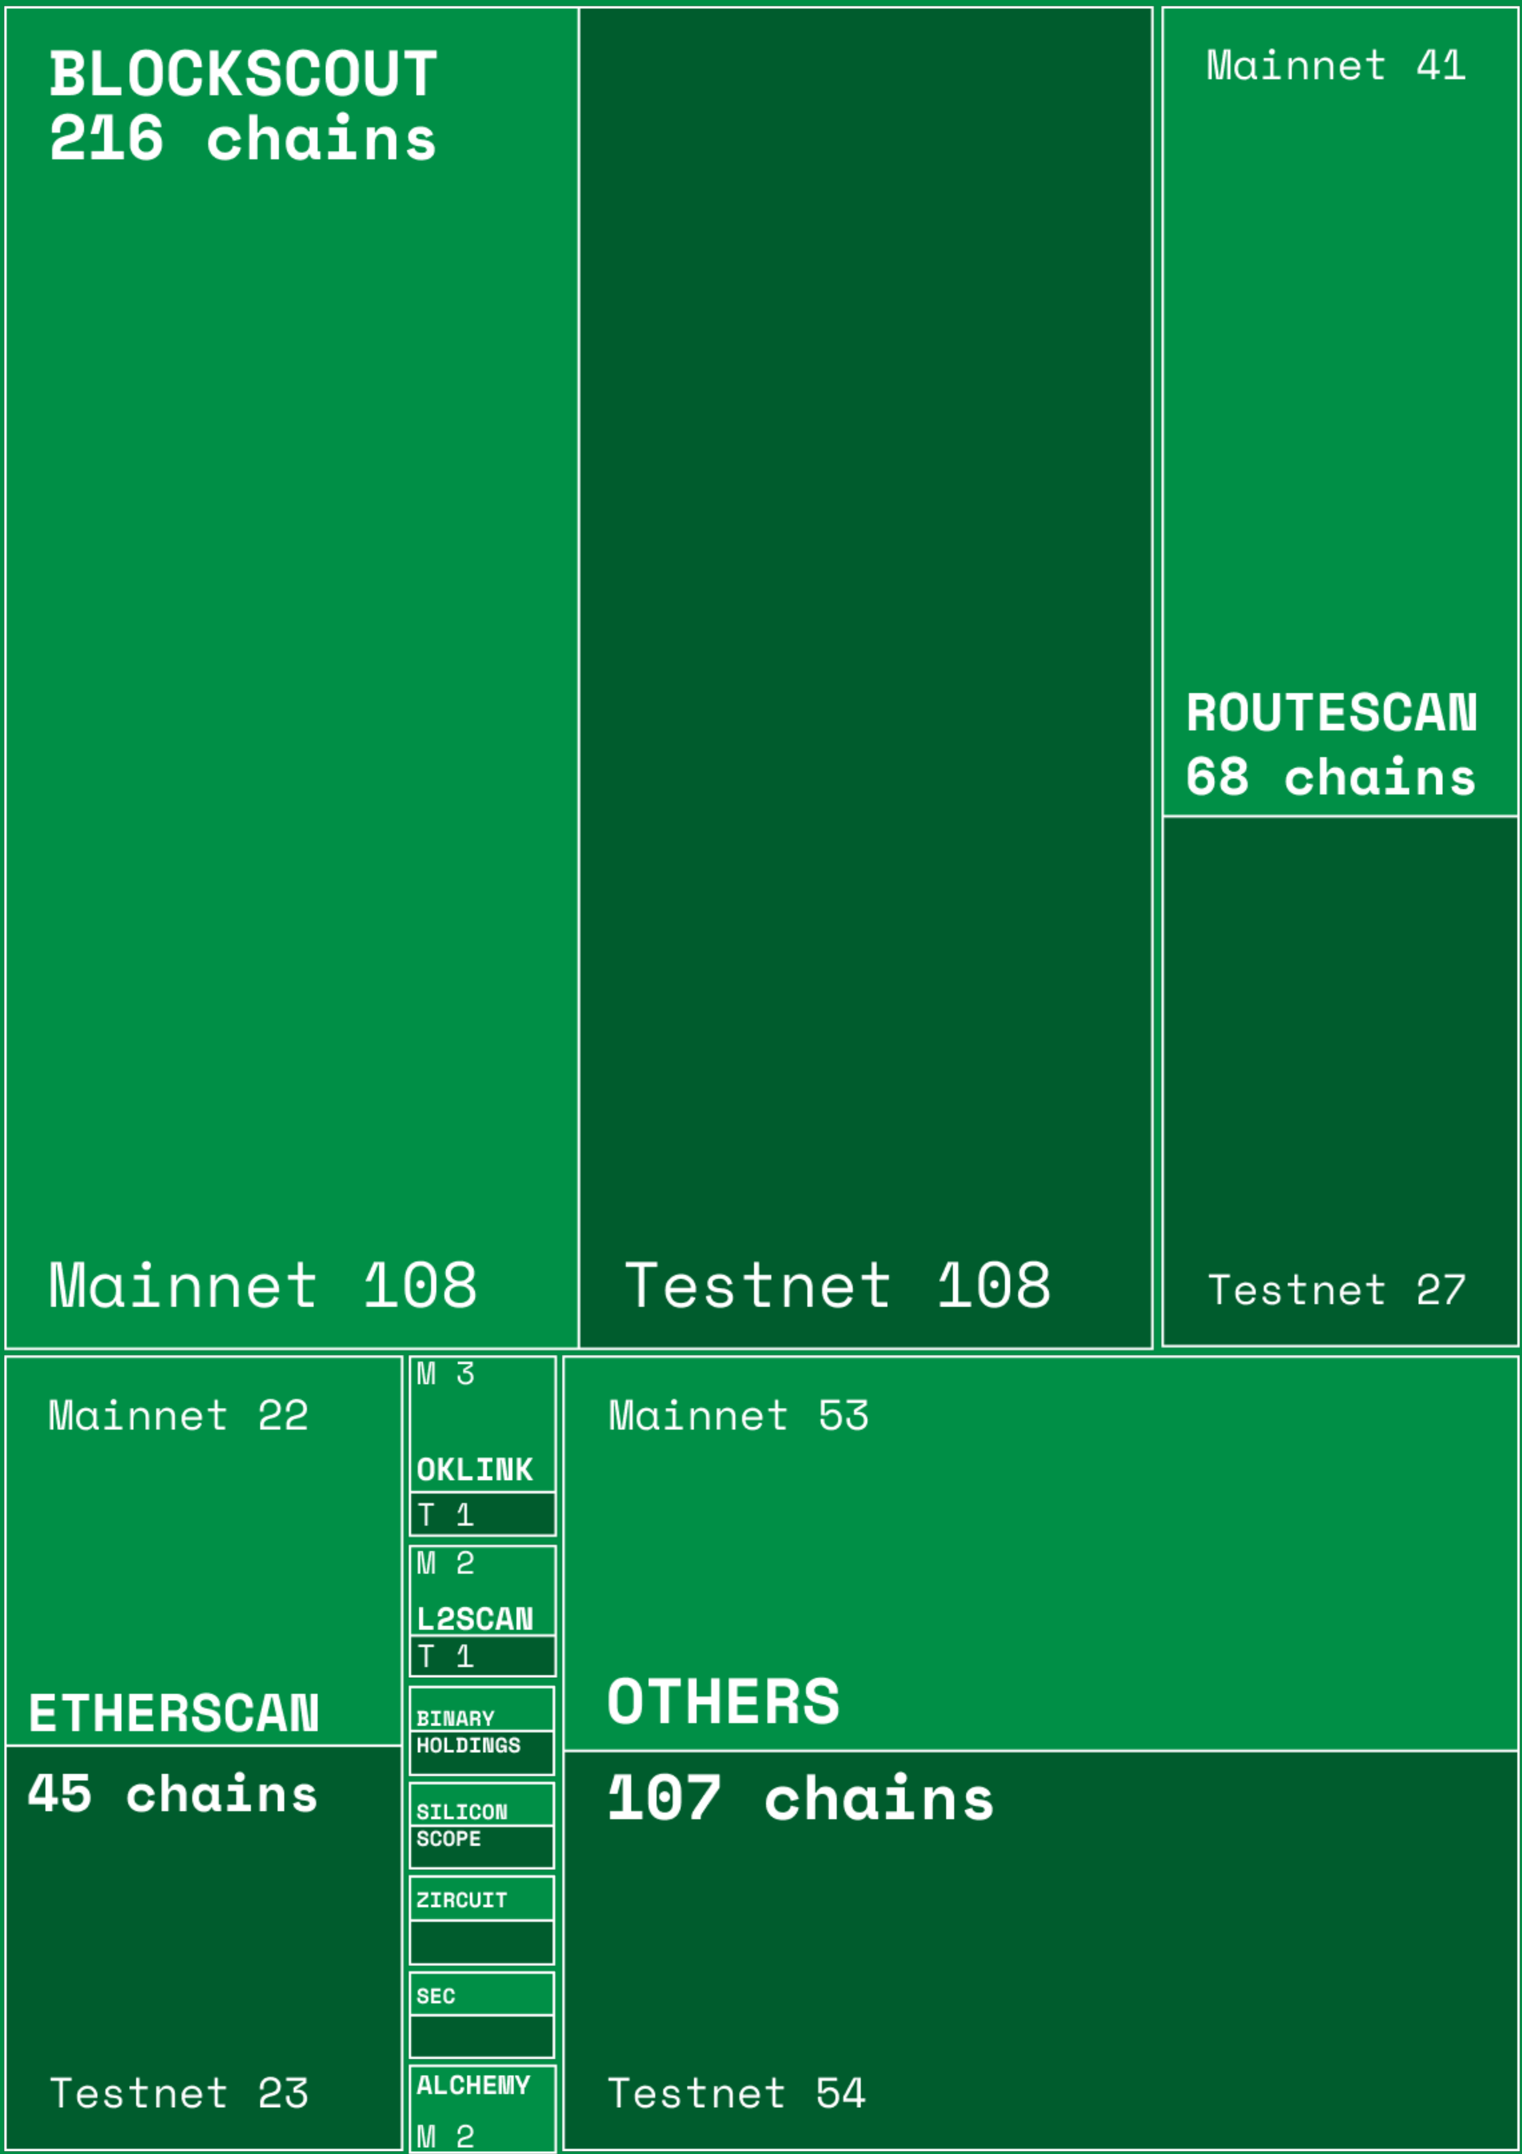
BLOCKSCOUT
216 chains
Mainnet 108
Testnet 108
Mainnet 41
ROUTESCAN
68 chains
Testnet 27
Mainnet 22
ETHERSCAN
45 chains
Testnet 23
M 3
OKLINK
T 1
M 2
L2SCAN
T 1
BINARY
HOLDINGS
SILICON
SCOPE
ZIRCUIT
SEC
ALCHEMY
M 2
Mainnet 53
OTHERS
107 chains
Testnet 54

4.7 Key Insights: L2 as EVM-Native Scaling Standard

L2 infrastructure converged hard on EVM in 2025, prioritizing compatibility and developer distribution over execution-environment innovation. The layer still maintains a near-even production/development split (51.8% mainnet across 512 L2 networks), so “mature” does not mean “done.”

Execution & production signals
Specialized VMs (e.g., CAIROVM, ARBITRUMVM, FUELVM) show 100% mainnet rates in their tiny counts, implying non-EVM L2 choices are high-conviction bets rather than broad experimentation.

RaaS concentration vs autonomy
RaaS shows structural concentration: the top three providers (Caldera, Conduit, Gelato) control 69.53% of deployments, with both Caldera and Conduit exceeding 75% mainnet rates, well above L2’s 51.8% average. This shift from volume-led to production-led consolidation signals market maturation, as teams now prioritize operational reliability over deployment velocity. Self-hosted L2 (13.3%, 68 chains) remains significant, indicating persistent demand for sequencing control despite managed convenience. The near-parity between Caldera (27.34%) and Conduit (26.56%) creates duopoly dynamics, maintaining competitive pressure while concentrating systemic risk within a narrow operator base.

Rollups and ecosystem alignment
Rollup architecture shows a dual-standard: Optimistic Rollups lead at 57.68% (154 chains) and ZK Rollups follow at 26.59% (71 chains), with smaller categories (Optimium 10.11%, Validium 4.87%) remaining niche but present.

Chapter V: Layer 3 (L3) Analysis

Layer 3 blockchains represent application-specific execution environments designed to optimize for specialized use cases, custom economics, and tailored governance while inheriting security from underlying L2 infrastructure. Our analysis of 81 Layer 3 networks shows an experimental yet production-focused segment (3.8% of categorized layer infrastructure) with the highest mainnet adoption across layers at 63.0%. This production tilt suggests L3 deployment is typically driven by deliberate app requirements rather than speculative infra proliferation.

5.1 Virtual Machine (VM) Distribution

The Application-Specific Execution Environment Question

How do Layer 3 networks balance application-specific optimization with ecosystem interoperability when designing execution environments for vertical integration?
The 2025 data reveals L3 infrastructure demonstrates near-total EVM concentration at 94.20% (65 chains), slightly lower than L2's 95.94% but more production-mature with 58.5% mainnet adoption. Specialized VMs achieve perfect mainnet adoption (ARBITRUMVM, COSVM, ZKEVM all at 100%), indicating that the small minority of L3 teams selecting non-EVM execution environments possess absolute production conviction.

> Because some classifications require explicit technical confirmation, not every metric in this chapter uses the same denominator: totals like "mainnet vs testnet" use the full L3 count (81), while VM shares are computed only on the VM-identified L3 subset (69), i.e., the L3 networks where VM attribution is validated; percentages in the VM section therefore sum to 100 within that subset, not across all L3s.

VM Distribution Analysis

EVM's 94.20% L3 dominance (between L2's 95.94% and L1's 61.14%) demonstrates that application-specific layers overwhelmingly prioritize compatibility and developer ecosystem access. The 58.5% mainnet adoption rate, significantly above EVM's global 46.9% average and L2's 51.9%, indicates that L3 EVM deployments reflect genuine production use cases rather than experimental infrastructure.
Specialized VM perfection: ARBITRUMVM (2 chains, 100% mainnet), COSVM (1 chain, 100% mainnet), and ZKEVM (1 chain, 100% mainnet) all achieve perfect production rates, demonstrating that L3 teams selecting non-EVM architectures possess exceptional clarity on technical requirements and production pathways. The presence of only 4 distinct VM types across L3 infrastructure, the lowest diversity across all layers, suggests that application-specific deployments converge on proven execution environments rather than experimenting with novel VM architectures.

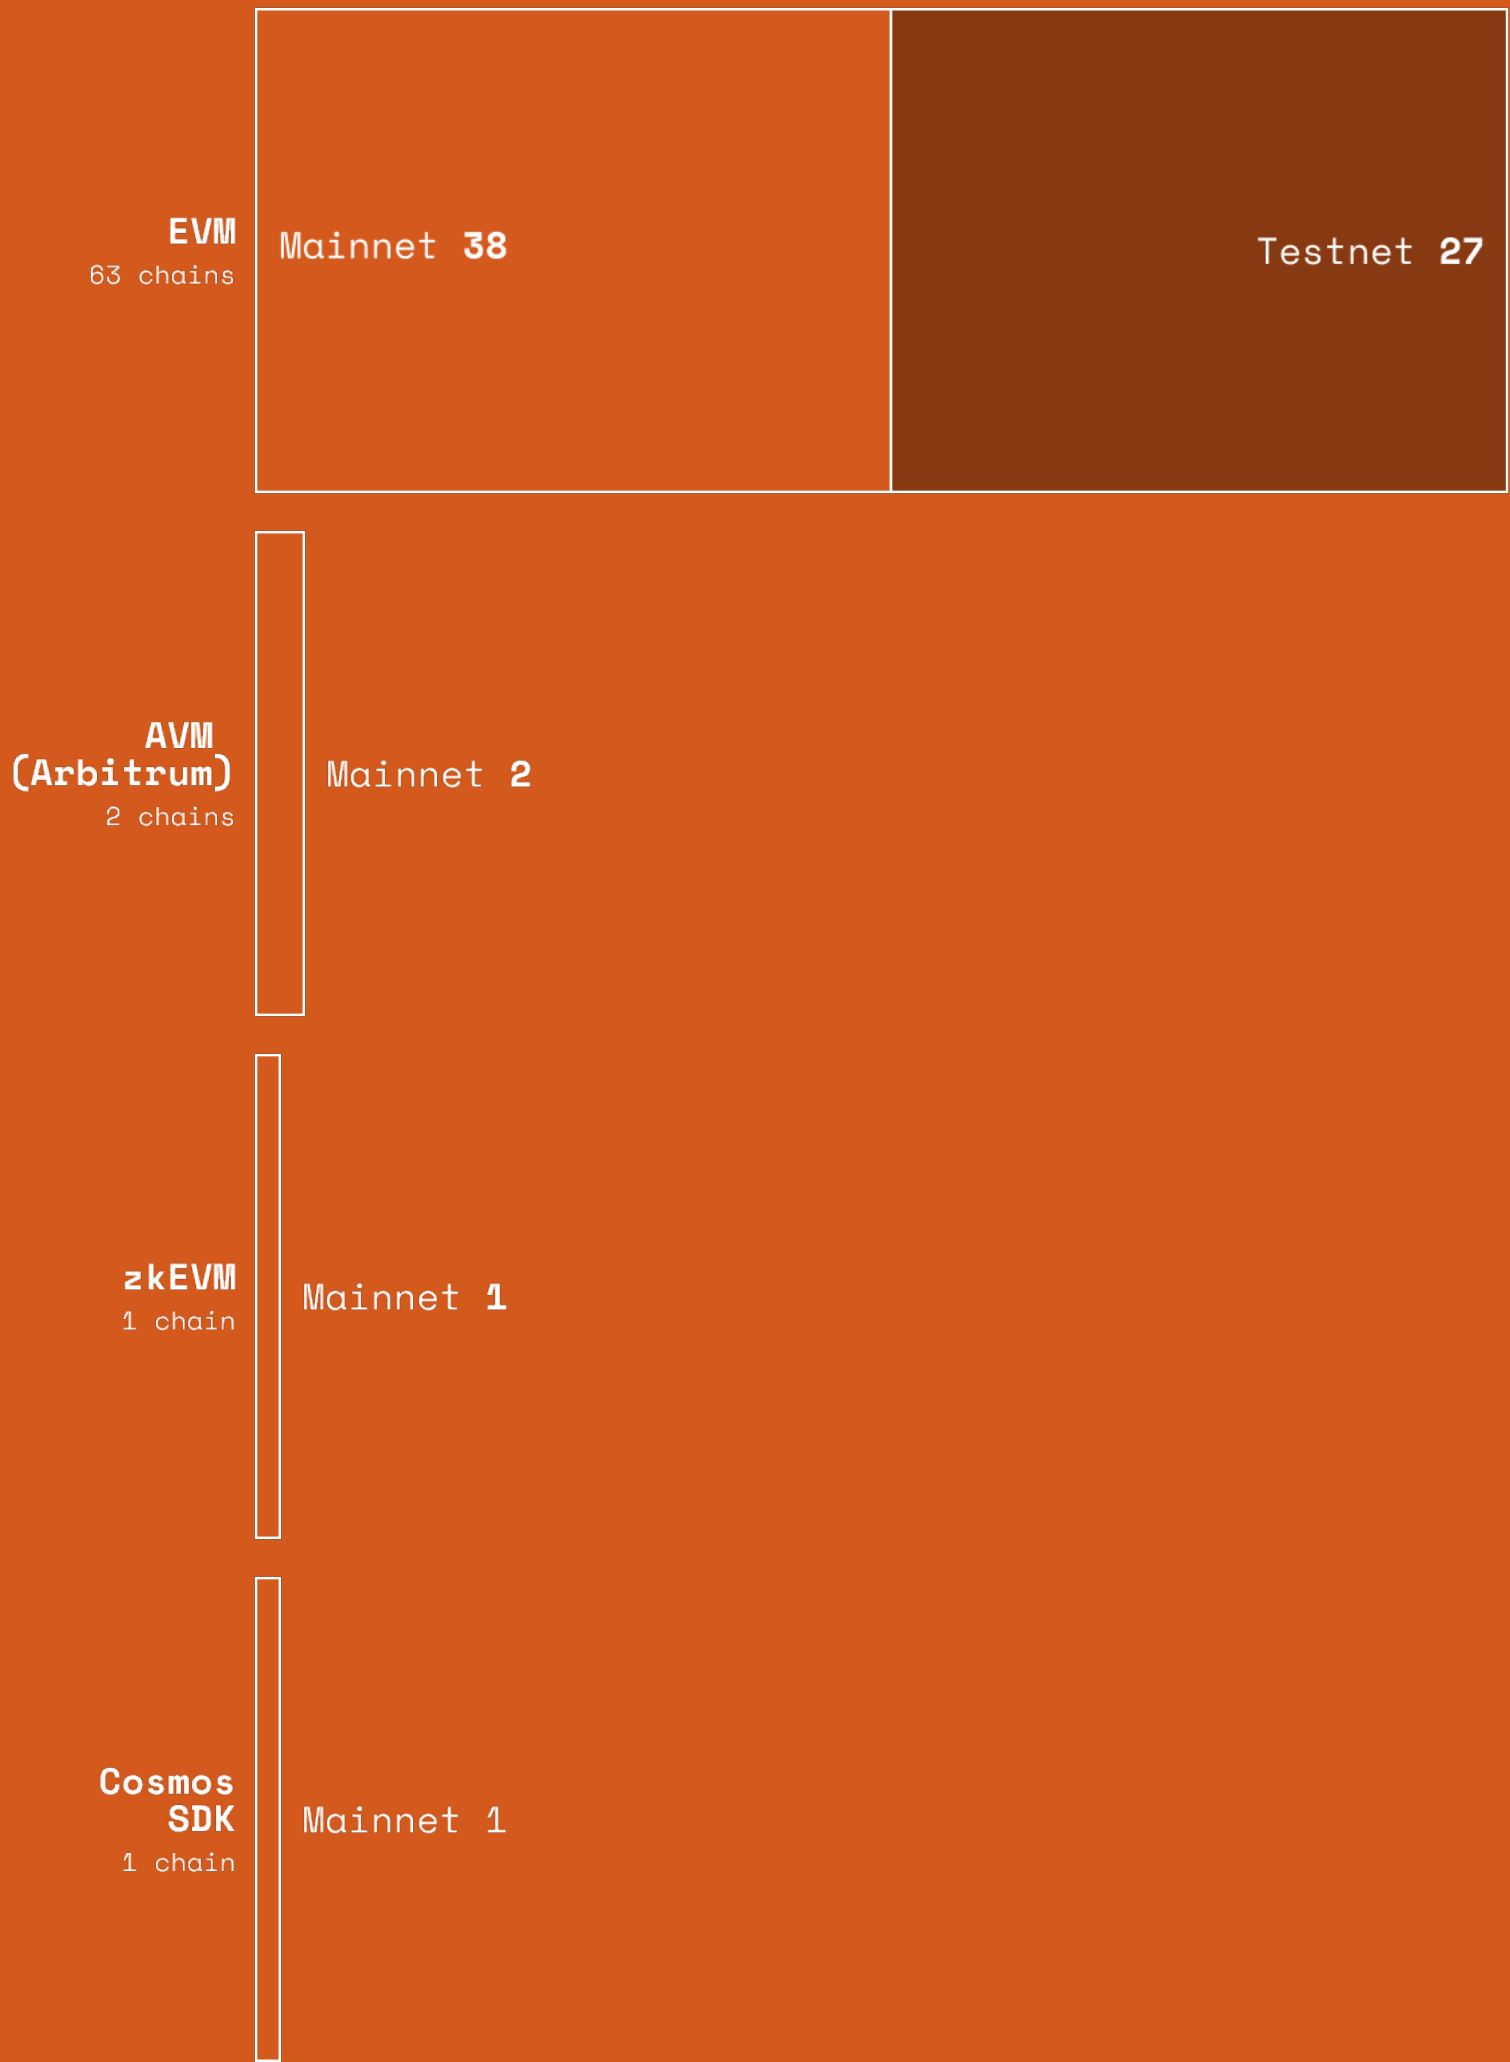
EVM
63 chains
Mainnet 38
Testnet 27
AVM (Arbitrum)
2 chains
Mainnet 2
zkEVM
1 chain
Mainnet 1
Cosmos SDK
1 chain
Mainnet 1

5.2 Network Type Distribution

L3 blockchains demonstrate the strongest production focus across all layers, with mainnet deployments significantly outpacing experimental environments.

Environment Distribution:
This 63.0% mainnet rate represents:

- + 9.7pp above industry average (53.3%)
- + 4.9pp above L1 (58.1%)
- + 11.2pp above L2 (51.8%)

Highest production readiness across all layers.
The exceptional mainnet tilt indicates that L3 deployment reflects deliberate application-specific production requirements rather than experimental infrastructure proliferation. Teams deploying L3s possess clear business models, defined user bases, and production conviction, validating L3 as an application-focused architecture rather than speculative technology demonstration.

Production Networks

51

mainnet deployments (53.3%)

Staging/Testing Networks

30

testnet environments (37.0%)

5.3 Ecosystem Distribution

Arbitrum Dominance in Application-Specific Deployment

> Ecosystem shares in this section are computed on the ecosystem-mapped L3 subset (28), not on total L3 networks (81), because ecosystem attribution is only reported where a L3's primary settlement/framework affiliation can be validated with enough confidence to avoid misclassifying hybrid or ambiguously positioned stacks. In practice, this means the percentages below answer "within the L3s we can reliably map to an ecosystem, how is L3 supply distributed?", and they should not be read as a projection over all L3 networks.

L3 ecosystem distribution reveals extreme concentration around Arbitrum's Orbit framework, which controls 78.57% of categorized L3 infrastructure (22 out of 28 categorized chains).
Arbitrum's 78.57% L3 dominance reflects Orbit's strategic positioning as the first production-grade L3 deployment framework, providing turnkey application-specific chain infrastructure with inherited Arbitrum L2 security.
Alternative L3 ecosystems include Ethereum (2 chains, 7.14%), Base (2 chains, 7.14%), and Optimism (2 chains, 7.14%), representing experimental L3 deployments on OP Stack and other frameworks.

The extreme ecosystem concentration (78.57% Arbitrum) versus L2's more distributed pattern (46.13% Ethereum, 17.48% Optimism, 9.46% Arbitrum) indicates that L3 represents an Arbitrum-led architectural innovation rather than a universal multi-ecosystem scaling paradigm.

Arbitrum
(78.57%)
22
Blockchains

Ethereum
(7.14%)
2
Blockchains

Base
(7.14%)
2
Blockchains

Optimism
(7.14%)
2
Blockchains

5.4 Role of Block Explorers

Layer 3 Observability Infrastructure Fragmentation

Block explorer distribution reveals how L3 application-specific infrastructure demands customized observability solutions, often proprietary or specialized implementations. The L3 market is served by 36 different explorers across 61 total integrations, the highest explorer-to-integration ratio across all layers, indicating significant fragmentation.

> In L3, explorer analysis is best done at the integration level (61), because "how many explorers exist" maps to deployments/instances rather than to the total number of L3 networks (81). A single L3 can have multiple explorer setups (e.g., separate mainnet/testnet instances or parallel implementations), while many L3s may share the same underlying explorer provider, so using "total L3 networks" as the denominator would blur provider concentration and misrepresent how observability is actually distributed.

Explorer Market Leadership Analysis

Blockscout leads L3 explorer support with 25 integrations (14 mainnet, 10 testnet), maintaining its position as the most versatile multi-layer explorer platform even in the highly fragmented L3 market.
Etherscan and Routescan each capture 2 integrations, representing minimal L3 presence for major explorer brands.The remaining 33 explorers serve individual or small groups of L3 networks, indicating that application-specific chains often require proprietary observability infrastructure tailored to custom business logic, governance models, or data structures. The high explorer provider count (36) relative to total integrations (61) suggests that many L3 deployments utilize custom-branded or Caldera/Conduit-hosted explorer instances rather than relying on generic multi-chain platforms.

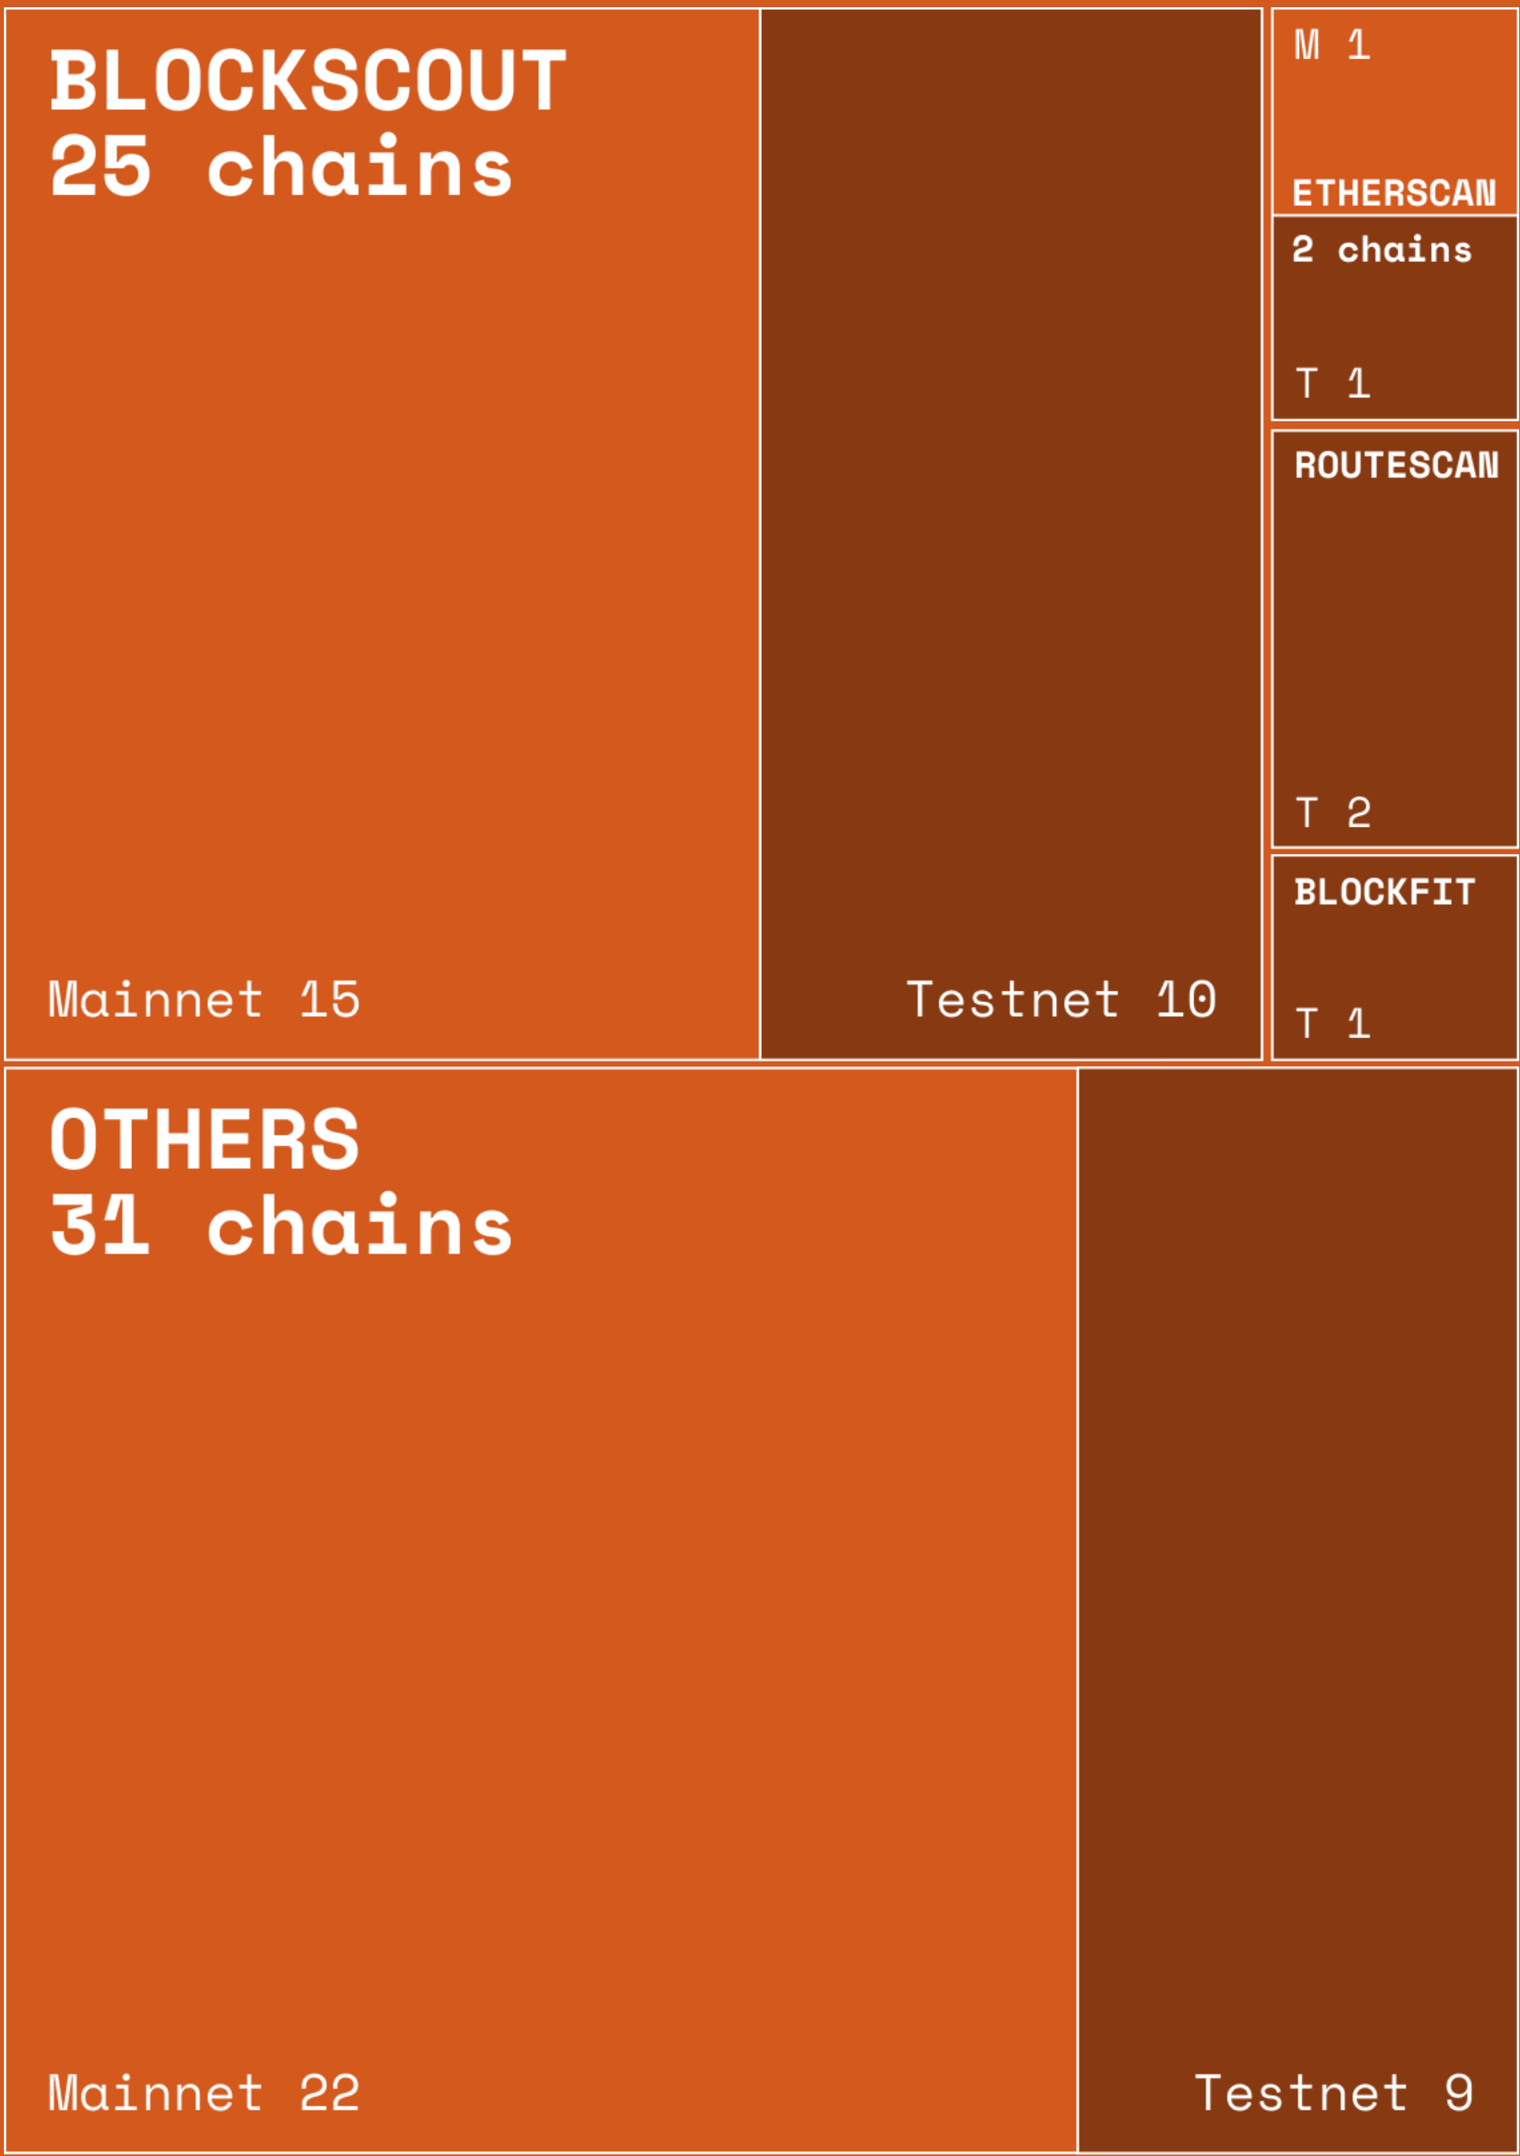
BLOCKSCOUT
25 chains
Mainnet 15
Testnet 10
M 1
ETHERSCAN
2 chains
T 1
ROUTESCAN
T 2
BLOCKFIT
T 1
OTHERS
31 chains
Mainnet 22
Testnet 9

5.5 Key Insights: L3 as High-Conviction Application Layer

Application-Specific Infrastructure Maturation Indicators

L3 infrastructure demonstrates exceptional production maturity (63.0% mainnet, highest across all layers, +9.7pp above industry average) despite minimal scale (3.8% layer share, 81 chains), validating L3 as a narrow but high-conviction application-specific deployment layer rather than a universal scaling paradigm.

Extreme Arbitrum Orbit concentration (78.57% of categorized L3) indicates framework provision monopoly, with EVM dominance (94.20%) and progressive mainnet adoption (L3 EVM 58.5% vs L2 EVM 51.9% vs L1 EVM 43.3%) confirming that L3 teams standardize on proven execution environments to minimize fragmentation while optimizing for application-specific economics.

L3 EVM share is 94.20% (69), marginally below L2's 95.94% (468), confirming that both scaling layers are effectively EVM-standardized while L1 remains the only layer with meaningful VM diversification.

The 1.69 integrations-per-provider ratio (36 explorers serving 61 integrations) reinforces that L3 chains often require custom-branded or RaaS-hosted explorer instances tailored to application-specific business logic rather than generic observability platforms.

Arbitrum's structural positioning as the only scaled production-grade L3 framework suggests L3 growth correlates directly with Orbit availability rather than organic multi-ecosystem demand, with alternative ecosystems (Ethereum, Base, Optimism) representing small 2-chain deployments.

Sustainable L3 expansion likely requires Orbit-equivalent frameworks targeting specific verticals rather than generic "L3 for everything" positioning.

DeFi Warhol
@Defi_Warhol

> "Layer 3 chains remind me more of product launches rather than infrastructure experiments. L3s ship when there is already a user, a revenue model, and a reason (which happens rarely)"

Chapter VI: Cross-Layer Comparison

The multi-layer blockchain architecture captured in the dataset shows clear differentiation across L1, L2, and L3 infrastructure in production maturity, VM standardization, ecosystem mapping coverage, and observability patterns. This chapter synthesizes those signals into a cross-layer view of how each layer behaves operationally, not narratively.

6.1 Production Maturity Across Layers

Layer-Specific Production Readiness

Cross-layer production maturity shows a clear pattern: L3 has the highest mainnet adoption rate (63.0%) despite being the smallest layer, while L1 sits in the middle (58.1%) and L2 is the lowest (51.8%).

Strategic Interpretation:

The inverse relationship between layer scale and production readiness implies that deployment motivations differ by layer: L3 tends to be shipped when there is already an application-driven need, while L2 keeps a heavier testnet pipeline, and L1 balances production with sovereignty-led experimentation.

	Layer 1	Layer 2	Industry
Total Chains	1,540	512	3,473
Mainnet	895	265	1,852
Testnet	645	247	1,621
Mainnet (%)	58.1%	51.8%	53.3%
Mainnet (%)	+4.8pp	-1.5pp	baseline

6.2

Growth Dynamics and Layer Distribution

L1 Dominance Persists
On the layer-categorized subset (2,133 networks), L1 remains dominant at 72.2% (1,540 chains), while L2 reaches 24.0% (512) and L3 remains at 3.8% (81).
This directly challenges any simplistic "everything migrates to L2/L3" assumption when measured by infrastructure footprint.

Strategic Implications:
The 72.2% L1 concentration validates that sovereignty, composability, and technical differentiation remain compelling value propositions despite scaling layer availability. The data challenges simplistic modular scaling narratives suggesting inevitable migration to L2/L3 architectures.

	Layer 1	Layer 2	Layer 3	Total
Total Chains	1,540	512	81	2,133
Layer Share (%)	72.2%	24.0%	3.8%	100%
Categorization Notes	Dominant architecture	Meaningful scale achieved	Specialized/ experimental	Out of 3,473 total

6.3

Virtual Machine Distribution Patterns

EVM Concentration Peaks on L2 (Not L3)

Cross-layer VM distribution is non-monotonic: within the VM-identified subsets, EVM concentration is lowest on L1 (61.14%), peaks on L2 (95.94%), and remains comparably high on L3 (94.20%). This pattern is consistent with a clear structural dynamic in the dataset: scaling layers converge most aggressively on EVM compatibility, while execution-environment diversification is primarily an L1 phenomenon tied to sovereignty and base-layer differentiation. The correct interpretation, therefore, is not "higher layer --> more EVM," but rather "EVM dominance peaks at L2 and persists at L3 in this snapshot".

L1 Alternative VM Breakdown:
On L1, WASM represents 25.34% (373 chains) and COSVM represents 12.64% (186 chains) of the VM-identified L1 subset (1,472), for a combined 37.98% WASM+COSVM share (~38.0%).

	Layer 1	Layer 2	Layer 3
VM-identified	1,472	468	69
EVM Share (%)	61.14	95.94	94.20
EVM Mainnet (%)	43.3	51.9	58.5
Non-EVM Share (%)	38.86	4.06	5.80
VM Types	11	12	4

Layer 1
In the VM-identified subset, EVM is less dominant on L1 (61.14% on 1,472; 11 VM types), confirming that execution environment diversity concentrates primarily on the sovereign layer.

Layer 2
On L2, standardization is nearly total: EVM at 95.94% on 468 (12 VM types), so the architectural choice is clearly driven by compatibility and tooling-first portability.

Layer 3
On L3, diversity collapses further (4 VM types on 69) but convergence remains very high (EVM 94.20%), indicating that even when the layer is "application-specific" the prevailing strategy is to minimize fragmentation and operational risks by standardizing on EVM.

DeFi Warhol
@Defi_Warhol

> "EVM dominance essentially reflects its distribution supremacy, not production confidence. When teams accept higher engineering cost by choosing WASM or COSVM, they overwhelmingly do so with a clear path to mainnet."

6.4 Mainnet Adoption by VM and Layer

Layer-Specific EVM Mainnet Adoption

EVM mainnet adoption increases with layer height: 43.3% on L1, 51.9% on L2, and 58.5% on L3. This suggests EVM choice on L1 includes a larger experimental tail, while on L2/L3 it aligns more with production-oriented compatibility requirements.

1. EVM shows inverse progression: EVM mainnet adoption increases from 43.3% (L1) to 58.5% (L3), indicating that EVM selection on L1 often reflects experimental deployments, while on higher layers it correlates with compatibility requirements for production use cases.

2. Alternative VMs demonstrate superior production conviction: WASM reaches 86.1% mainnet on L1 (+42.8pp vs EVM L1), COSVM maintains 69.9% (+26.6pp), and specialized VMs (CAIROVM, ARBITRUMVM, FUELVM) reach 100.0%, non-EVM selection is a reliable signal of technical differentiation strategy.

3. VM diversity concentrates on L1: within the VM-identified subsets, alternative VMs represent ~38.0% of L1 (WASM+COSVM), versus 4.06% on L2 and 5.80% on L3, validating that VM experimentation concentrates at sovereignty layers while scaling layers optimize for compatibility.

	Layer 1	Layer 2	Layer 3	Weighted Avg Mainnet %
EVM	43.3%	51.9%	58.5%	46.9%
WASM	86.1%	0%	-	85.6%
COSVM	69.9%	0%	100%	69.1%
CAIROVM	-	100%	-	100%
ARBITRUMVM	-	100%	100%	100%
FUELVM	-	100%	-	100%
ZKEVM	-	25%	100%	50.0%
ZKWASM	-	-	-	-

6.5 Ecosystem Consolidation: Framework Provision vs. Chain Deployment

Framework Provision Determines Strategic Positioning

Ecosystem mapping coverage differs significantly by layer: 42.7% for L1 (657/1,540), 67.2% for L2 (344/512), and 34.6% for L3 (28/81).
As a result, ecosystem "leaders" should be read as "leaders within the mapped/categorized subset," not as universal truth across all chains in each layer.

Critical Pattern (categorization coverage by layer)

Ecosystem categorization rates increase from L1 to L2, then decline at L3: L1 is 42.7 categorized (657/1,540), L2 is 67.2 (344/512), and L3 is 34.6 (28/81).
This indicates that L2 scaling layers exhibit the strongest ecosystem alignment signal (highest mapping coverage), while L1 and L3 still contain substantial "uncategorized" infrastructure in this dataset, either genuinely independent deployments or entries where ecosystem mapping is not available/assigned.

L1 Ecosystem Leaders (share of total L1, with 42.7% mapped): Within the categorized L1 subset, these four ecosystems account for 576/657 = 87.67%, indicating strong framework consolidation where classification exists.

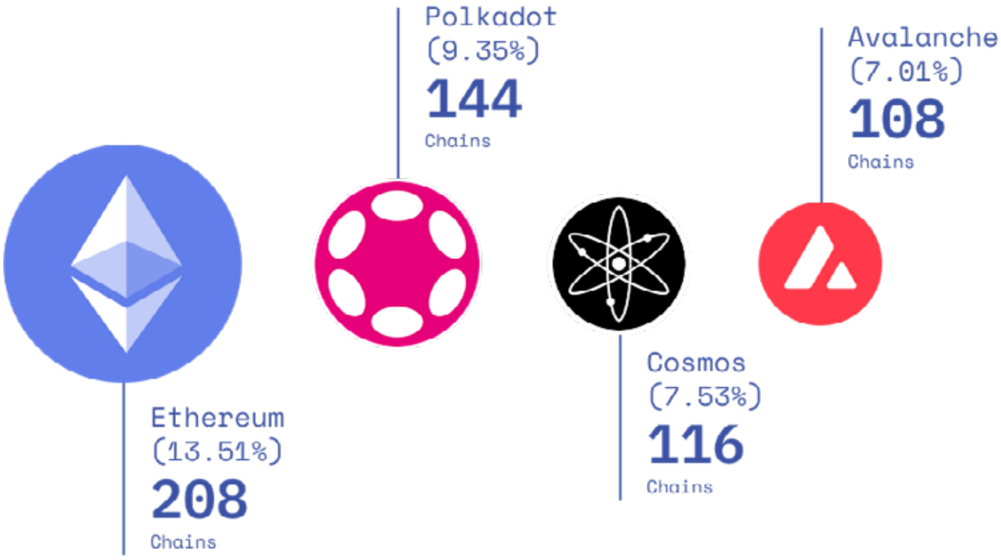

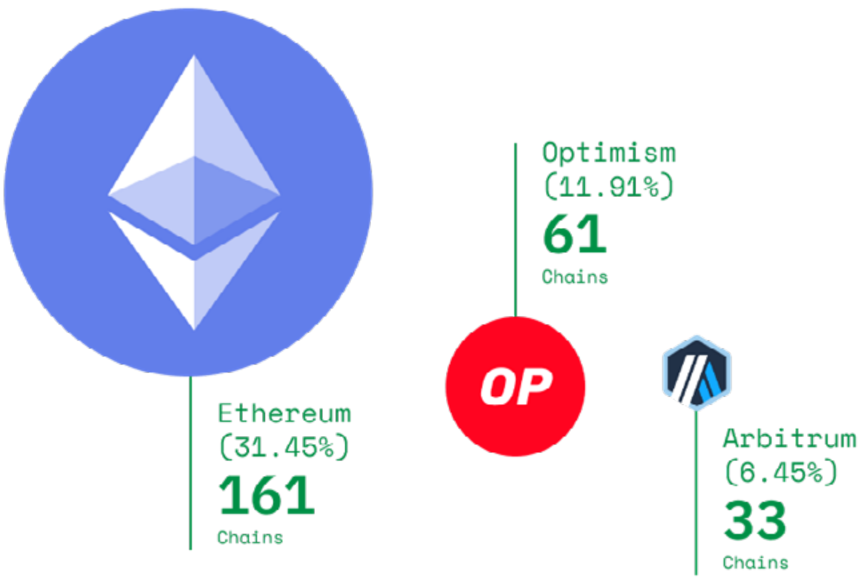

L2 Ecosystem Leaders (share of total L2, with 67.2 mapped): Within the categorized L2 subset, the top three represent 255/344 = 74.13

L3 Ecosystem Leaders (share of total L3, with 34.6% mapped): Within the categorized L3 subset, Arbitrum represents 22/28 = 78.57%, showing extreme concentration where ecosystem mapping exists.

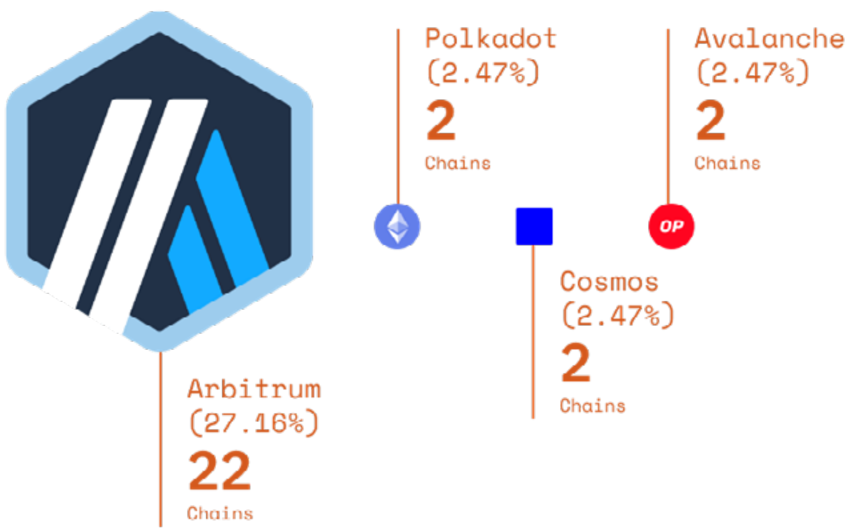

Strategic Implication (success patterns differ by layer)

Success patterns vary fundamentally by layer, because "what matters" operationally shifts: sovereignty and differentiation on L1, settlement/framework alignment on L2, and strong single-ecosystem gravity on L3 (within the mapped subset).

- L1: When ecosystem classification exists, it is highly consolidated (Ethereum/Polkadot/Cosmos/Avalanche account for 576/657 = 87.67% of the categorized L1 subset) so a credible L1 strategy typically requires either alignment with these dominant frameworks (within the mapped universe) or a differentiation thesis strong enough to justify being outside them.
- L2: With the highest categorization coverage (67.2%), L2 shows clearer ecosystem "anchoring" signals; within the categorized L2 subset, the top three ecosystems (Ethereum 161, Optimism 61, Arbitrum 33) represent 255/344 = 74.13%, which implies that ecosystem/stack alignment is a primary go-to-market and distribution constraint for L2s.
- L3: With lower mapping coverage (34.6%) but extreme concentration where mapped, Arbitrum represents 22/28 = 78.57% of the categorized L3 subset, implying that L3 adoption (in the categorized slice) is structurally dependent on a dominant ecosystem anchor rather than being evenly multi-ecosystem.

6.6 Explorer Infrastructure and Observability Patterns

Layer-Specific Explorer Density

Explorer integration density is highest on L2 (0.88 integrations/chain), while L3 shows the highest provider fragmentation (only 1.69 integrations/provider). This matches incentives: scaling layers compete on UX/tooling maturity, while application-specific L3 deployments frequently rely on custom-branded or provider-hosted explorer instances.

	Layer 1	Layer 2	Layer 3	Total (Global*)
Total Chains	1,540	512	81	3,473
Integrations	1,162	453	61	2,155
Explorers	269	114	36	591
Integrations/ Chain	0.75	0.88	0.75	0.62 Global*
Integrations/ Provider	4.32	3.97	1.69	3.65 Global*

*Layer rows use the layer-classified subset (2,133 chains = L1+L2+L3); the 'Global' row uses the full validated universe (3,473) and is not additive across layers.

Strategic Interpretation

L2 Highest Integration Density (0.88):Scaling layer infrastructure prioritizes user-facing observability with 88 integrations per 100 chains, reflecting competitive multi-explorer support driven by production-focused deployment and ecosystem emphasis on developer tooling accessibility.

L1/L3 Lower Density (0.75): Foundational and application-specific chains demonstrate 75 integrations per 100 chains, indicating reliance on single-explorer or proprietary solutions where generic multi-chain platforms provide insufficient architectural specificity.

L3 Observability Fragmentation (Integration-Level Analysis):L3 explorer analysis uses the integration level (61) rather than total network count (81), since individual L3s may deploy multiple explorer instances while many share the same provider. 36 explorers serve 61 integrations, yielding 1.69 integrations per provider, the lowest ratio across all layers, signaling extreme fragmentation and customization. Blockscout leads with 24 integrations (14 mainnet, 10 testnet), Etherscan and Routescan each have 2, and the remaining 33 explorers mostly serve individual or small L3 clusters, consistent with application-specific chains requiring proprietary or heavily customized observability rather than standard multi-chain platforms.

Blockscout Multi-Layer Dominance

Blockscout maintains leadership across all layers
with 565 total integrations (26.2% global share): 325 L1 integrations (28.0% L1 share), 216 L2 integrations (47.7% L2 share), and 25 L3 integrations (39.3% L3 share). This cross-layer consistency demonstrates Blockscout's position as the most versatile multi-architecture explorer platform, supporting diverse VM types, consensus mechanisms, and rollup architectures.

Layer-Specific Explorer Leadership:

- L1: Blockscout (28.0%), Avalabs (10.1%), Ping Dashboard (88 integrations), Routescan (59), Subscan (42)
- L2: Blockscout (47.7%), Routescan (15.0%), Etherscan (9.9%)
- L3: Blockscout (39.3%), followed by fragmented proprietary/RaaS-hosted solutions

The increasing Blockscout market share from L1 (28.0%) to L2 (47.7%) to L3 (39.3%) validates its competitive advantage in scaling layer observability, while alternative explorers like Avalabs (L1-focused), Routescan (L2 strength), and Etherscan (Ethereum-aligned) demonstrate ecosystem-specific specialization patterns.

6.7 Strategic Layer Specialization Patterns

Cross-layer data points to a clear specialization pattern: L1 concentrates technical differentiation and alternative VMs, L2 converges on EVM-centric scaling and ecosystem alignment, L3 focuses on application-specific production atop a single dominant framework.

- L1, which accounts for 72.2% of categorized infrastructure (1,540 chains), shows the highest VM diversity and the strongest presence of alternative execution environments, with WASM and COSVM together representing 38.0% of L1 VM distribution and achieving significantly higher mainnet adoption (86.1% and 69.9%) than L1 EVM at 43.3%. This confirms L1 as the sovereignty layer where teams pay the cost of non-standard choices to secure architectural differentiation, ecosystem positioning across Ethereum/Polkadot/Cosmos/Avalanche, and long-term autonomy.

- L2, which represents 24.0% of categorized infrastructure (512 chains), exhibits near-total EVM convergence (95.94%) and the highest explorer integration density (0.88 integrations per chain), while ecosystem distribution is heavily skewed toward Ethereum and Ethereum-aligned stacks (46.80% Ethereum, 17.73% Optimism, 9.59% Arbitrum). This configuration validates L2 as the EVM-native scaling fabric where compatibility, tooling, and settlement alignment dominate over experimentation, with RaaS consolidation further reinforcing standardized deployment patterns.

- L3, despite representing only 3.8% of categorized infrastructure (81 chains), delivers the highest production readiness (63.0% mainnet) and extreme ecosystem concentration (78.57% under Arbitrum). Combined with 94.20% EVM share and 39.3% Blockscout penetration, this shows L3 as a narrow but high-conviction application layer: teams deploy L3s when they have specific business logic, economics, or UX requirements that justify an additional execution layer, and they overwhelmingly standardize on Orbit and EVM to minimize tooling and liquidity fragmentation.

Overall, the specialization is unambiguous: L1 is where differentiation and alternative VMs live, L2 is where Ethereum-aligned EVM scaling standardizes, L3 is where high-conviction apps carve out customized execution while inheriting security and liquidity from lower layers. Any credible infrastructure strategy must position itself explicitly within this division of roles rather than treating L1/L2/L3 as interchangeable abstractions.

6.8 The Modular Scaling Narrative: Data vs. Rhetoric

Prevailing narratives suggest an inevitable migration from monolithic L1s toward modular L2/L3 architectures, but the 2025 dataset indicates a more stratified reality where each layer serves distinct operational roles rather than representing an evolutionary ladder. As shown throughout Chapter 6 (especially Sections 6.2–6.7), the key question is not whether modular scaling exists, but where it applies, and where it does not.

Key Data Contradictions to Modular Scaling Orthodoxy:

1. L1 dominance persists (in footprint). The layer-categorized subset shows L1 retains the majority of deployments, contradicting any simplified "everything migrates upward" narrative when measured by infrastructure footprint rather than discourse.

2. Execution diversification concentrates at sovereignty layers.VM diversity remains primarily an L1 phenomenon, while L2/L3 converge on compatibility-first execution choices, limiting the practical scope of "modular flexibility" claims.

3. L3 is not "the future of everything." L3 reads as an application-driven niche layer rather than a universal destination, with production behavior driven by deliberate app requirements rather than generalized infrastructure experimentation.

4. L2 standardizes aggressively. L2 behaves as a compatibility and tooling-driven scaling fabric, which is consistent with near-total convergence around standardized execution and operational patterns rather than architectural plurality.

5. Ecosystem gravity intensifies upward (within mapped subsets). Ecosystem concentration patterns, especially where mapping exists, indicate that higher layers tend to cluster around fewer dominant rails, which cuts against the expectation that modularity yields broad, evenly distributed proliferation.

Synthesis:
The data supports a narrower, more accurate framing: modular scaling is a deployment pattern optimized for specific constraints (throughput/cost, compatibility, and operational leverage), not a universal architectural destination. In parallel, sovereignty requirements and differentiation incentives keep L1 structurally relevant, while L3 adoption appears when teams have strong application-specific drivers that justify additional execution complexity

Conclusion: Strategic Layer Specialization, Not Universal Modularity

What emerges from 2025 is not a clean migration story, but a specialization story: each layer persists because it is optimized for a different job. Modular scaling is one of the strategies teams deploy, not the endpoint that replaces everything else, L1, L2, and L3 remain concurrently relevant because their trade-offs are structurally different.

Strategic Implication for 2026

This section establishes the analytical premise for 2026 planning, layer behavior is tier-specific and constraint-driven. The actionable 2026 decision framework is consolidated in Chapter 7 to avoid duplicating conclusions across chapters, and to keep all recommendations in one place.

Chapter VII: Conclusion & 2026 Outlook

2025 made it clear that L1, L2, and L3 are not evolutionary "steps," but distinct operational tiers with different incentives and failure modes. Chapter 6 already did the quantitative cross-layer synthesis and the "data vs. rhetoric" debunking, so this chapter converts those signals into 2026 execution choices: stack selection, go-to-market constraints, and, most importantly, dependency risk management across frameworks, RaaS, and observability. The core thesis for 2026 is straightforward: winners won't be the teams that "pick the right layer" in the abstract, but the teams that choose the tier that matches their product constraints and then minimize lock-in and fragility across the supply chain (framework → deployment → explorer/observability).

7.1

Key Findings Synthesis

Finding 1 — Production-first is real, but uneven
The “production-first” signal exists at a macro level, but it is not uniform across layers, so in 2026 “mainnet count / mainnet share” cannot be treated as a generic quality proxy. Operationally, where a tier shows higher churn and heavier staging pipelines (as established in Chapter 6), upgrade cadence, incident response, and support costs must be planned as part of the product, not as overhead.

Finding 2 — L1 dominance is structural, not inertia
L1 remains a structural component of the infrastructure landscape, not “inertia,” so in 2026 it should be treated as an active market: sovereignty, differentiation, and autonomy still justify dedicated deployment and tooling decisions. Concretely, infrastructure/tooling providers should avoid positioning as “scaling-only,” and builders should not dismiss L1 as “legacy” when core risk is governance/control or meaningful execution divergence.

Finding 3 — Execution diversification is an L1 phenomenon; scaling converges on EVM
Execution-environment diversification concentrates where sovereignty exists (primarily L1), while scaling layers converge on compatibility and standardization (EVM-first), consistent with the cross-layer evidence. A 2026 decision rule follows: “multi-VM strategy” is largely an L1 problem; on L2/L3 the real competition is distribution, tooling, settlement alignment, and integrations—not VM novelty.

Finding 4 — Framework provision determines market structure
Where mapping is available, concentration around a small set of framework rails is a market constraint, not a descriptive footnote. In 2026, framework alignment should be treated as a go-to-market constraint because it shapes

tooling compatibility, liquidity/integration pathways, and the cost of acquiring developers and partners.

Finding 5 — Managed infra is the default; dependency risk is the tax

Managed infrastructure is concentrated: the top three RaaS providers (Caldera, Conduit, Gelato) control 69.53% of deployments, with Caldera (27.34%) and Conduit (26.56%) at near-parity. Concentration has shifted from volume-led to production-led: Caldera (77.1% mainnet) and Conduit (82.4% mainnet) both exceed the L2 average (51.8%), proving managed infrastructure now powers production systems, not just testnets. The 2026 trade-off is operational-economic: time-to-market versus migration risk, managed convenience versus sequencing autonomy, duopoly pressure versus systemic dependency.

Hard rule: define the exit path before launch, sequencer/infra migration plan, runbooks, explorer/indexing compatibility, because provider failures become systemic events when 69.53% of supply depends on three operators.

Finding 6 — L3 is a niche production layer, not a universal endpoint

L3 shows up as an application-specific, high-conviction layer, but not as a universal scaling endpoint. In 2026, L3 makes sense when constraints are strong enough (economics, governance, UX, execution requirements) to justify added complexity and when the project can ride a dominant rail rather than fighting fragmentation alone.

Finding 7 — Observability is a strategic dependency, not a checkbox

Explorer/observability is not a final deliverable; it is part of the product and part of operational trust, and provider concentration/fragmentation patterns differ by layer as shown earlier in the report. In 2026, explorer strategy must be defined alongside deployment strategy because it impacts DX, incident response, ecosystem credibility, and downstream integrations (wallets, bridges, analytics).

7.2

Study Limitations & Future Research Directions

This analysis remains deliberately denominator-bound: layer-classified, VM-identified, ecosystem-mapped, rollup-identified, and explorer-integration views have different coverage, so percentages must be read strictly within their stated subsets. As a result, the best 2026 decisions come from structural signals (convergence, concentration, production tilt) rather than aggressive interpretation of a single metric out of context.

Priority research directions that would materially improve 2026 decision quality:

- Longitudinal tracking of alternative VM adoption/churn to separate durable trends from cohort effects.
- Measuring real “graduation/migration” dynamics across layers (what actually moves, when, and at what cost).
- Comparing operational outcomes of RaaS vs self-host (incident rate, upgrade cadence, migration cost), since concentration is already structural.
- Expanding/refining ecosystem/framework mapping coverage to reduce the “uncategorized” surface and improve comparability.

Conclusion: Strategic positioning in a stratified landscape

In 2026, strategy should start from the assumption that infrastructure is stratified: each tier optimizes different trade-offs (sovereignty/differentiation, compatible scaling, application-specific execution). The useful question is not "which layer wins," but "which specialization, and which rail, minimizes risk and maximizes distribution for this product," given that concentration and dependencies (frameworks, RaaS, observability) are already structural features of the market.

Routescan

Glossary

Glossary

Aa

Arbitrum Orbit	Framework for deploying custom Layer 3 chains using Arbitrum technology.
Ascraeus / Huygens	Specialized VM variants used in experimental Layer 2 networks.
AVM (Arbitrum Virtual Machine)	VM for Arbitrum optimistic rollups, executing Ethereum-compatible smart contracts.

Bb

Block Explorer	Web interface to view blockchain data like transactions, blocks, and addresses. Essential for network transparency and observability.

Cc

CairoVM	VM used in Starknet, optimized for zero-knowledge proof (ZKP) computation using the Cairo programming language.
Clipboard Hijacking	A type of malware that replaces copied wallet addresses or seed phrases with attacker-controlled values.
Consensus Mechanism	Protocol that allows distributed blockchain nodes to agree on the network state (e.g. PoS, PoW, etc.).
CosmWasm	WASM-based smart contract platform within the Cosmos ecosystem.
Cosmos SDK	Modular blockchain development framework for building application-specific blockchains in Cosmos.
COSVM (CosVM)	An execution environment positioned as a Cosmos-based chain that is compatible with the Ethereum Virtual Machine (EVM), enabling deployment of Ethereum-compatible smart contracts and dApps (including Solidity-based development) while interfacing with the broader Cosmos ecosystem (e.g., via IBC-style cross-chain connectivity).
CVM	Custom Virtual Machine often used in enterprise or privacy-focused chains; high production adoption rate.
Custom VM	Any virtual machine designed outside of standard frameworks (like EVM or WASM) for specialized use cases.

Dd

DEX (Decentralized Exchange)	A blockchain-based platform for peer-to-peer cryptocurrency trading without intermediaries.
DA (Data Availability)	Refers to how and where transaction data is stored and retrievable—onchain (validium, rollups) or off-chain.

Ee

Ecosystem	A grouping of chains, tools, and infrastructure providers sharing governance or technical compatibility.
EVM (Ethereum Virtual Machine)	Default smart contract execution environment for Ethereum and most L2 chains.
Ethereum-based (L1 context)	Classification category for Layer 1 networks built using the Ethereum stack or technically dependent on Ethereum settlement/security architecture. In the L1 ecosystem distribution analysis, "Ethereum-based" represents a technical-stack bucket (208 chains, 31.66% of categorized L1s) rather than referring exclusively to the Ethereum mainnet L1 itself. This includes sovereign L1s that adopt Ethereum's core stack, EVM compatibility standards, or structural dependencies on Ethereum infrastructure.
Execution Environment	The VM and runtime context in which smart contracts are processed.
Explorer Probing	Method of validating explorer integrity, features, and ecosystem coverage through direct testing.

Ff

FVM (Filecoin Virtual Machine)	Execution environment for smart contracts in the Filecoin network, combining storage and compute.
Fraud Proof	A mechanism used in optimistic rollups to verify the correctness of state transitions.
FUELVM (FuelVM)	VM type label appearing in the report's VM tables. It refers to Fuel's execution environment (FuelVM) as a non-EVM smart contract runtime.

Gg

GnoVM	Functional virtual machine using Go-based smart contracts, mostly experimental.

Hh

Hybrid VM	VMs combining components of multiple architectures (e.g., EVM + WASM) to optimize for interoperability or performance.

Ii

IBC (Inter-Blockchain Communication)	Protocol that enables token and message transfers between Cosmos-based chains.
Infrastructure Maturity	Degree of development of a blockchain, classified as Early Stage, Developing, or Mature.

Interoperability Ability of different blockchain networks to interact and share data or assets securely.

Kk

Keylogger Malware that records user keystrokes, often used to steal seed phrases or passwords.

Ll

L1 (Layer 1) Base blockchain protocol layer that ensures security, consensus, and settlement (e.g., Ethereum, Solana).

L2 (Layer 2) Protocols built on top of L1s to scale transactions while preserving L1 security guarantees (e.g., Optimism, Arbitrum).

L3 (Layer 3) Application-specific chains built on L2s, optimized for particular use cases or industries (e.g., gaming, DeFi).

Lumoz Rollup-as-a-Service provider with limited market share and single-layer focus.

Mm

Mainnet The live production version of a blockchain network, where real assets and economic activity occur.

Metadata Extraction The process of collecting structured, protocol-level information about blockchain networks.

MoveVM Move Virtual Machine used by blockchains implementing the Move programming language, originally developed by Meta for Diem and now adopted by networks like Aptos and Sui for secure smart contract execution.

Oo

OP Stack Open-source toolkit developed by Optimism to build Layer 2 rollup chains.

Optimistic Rollup L2 scaling method that assumes transactions are valid unless proven otherwise via fraud proofs.

Optimium A rollup variant that stores transaction data off-chain to reduce costs, used in some OP Stack configurations.

Pp

Polygon CDK (Chain Development Kit) Framework for building Ethereum-compatible, zero-knowledge-powered chains.

Production Environment A live blockchain environment processing real transactions and value.

Rr

RaaS (Rollup-as-a-Service) Managed service providers offering plug-and-play infrastructure for deploying rollups.

RPC (Remote Procedure Call) A communication protocol allowing applications to interact with blockchain nodes.

Routescan Unified multichain block explorer and analytics platform offering RPC-level validation and metadata-based infrastructure analysis.

RIDEVM (RideVM) VM type label appearing in the report's VM tables. It refers to the execution environment associated with Ride-based smart contracts (commonly linked to the Waves ecosystem).

Ss

SVM (Solana Virtual Machine) High-performance VM designed for parallel execution on the Solana blockchain.

Seed Phrase A list of words used to recover a crypto wallet. If lost or stolen, access to funds is irreversibly compromised.

Self-hosted A deployment model where the blockchain infrastructure is run by the project team or enterprise rather than a third-party service.

Side chain A blockchain that runs independently but connects to a main chain via a bridge for asset or message transfers.

Superchain A network of interoperable L2s sharing security and governance, proposed by Optimism.

Tt

Testnet A simulated blockchain environment for testing and development, using valueless tokens.

TVM Tron Virtual Machine, used primarily in Tron and compatible environments.

Vv

Validium A zero-knowledge-based scaling method storing transaction data off-chain while ensuring validity via proofs.

VM (Virtual Machine) Execution layer within a blockchain responsible for processing smart contracts in a deterministic and isolated way.

Ww

WASM (WebAssembly) A low-level binary format optimized for fast execution and portability; increasingly adopted as a VM in newer chains.

Zz

ZK Rollup (Zero-Knowledge Rollup) A scaling technique that uses zero-knowledge proofs for validity and finality of transaction batches.

zkEVM A virtual machine compatible with Ethereum that supports generation of zero-knowledge proofs.

zkSync Era A Layer 2 protocol using zkEVM to deliver scalable, Ethereum-compatible transactions with enhanced privacy and speed.

zkMVM / zkWASM Experimental zero-knowledge VMs combining ZKP validation with MVM or WASM architectures.

Note:
This glossary reflects technical terminology as used in the 2025 State of the Market blockchain infrastructure analysis. Definitions are provided for clarity and may encompass broader or more specific meanings in different contexts.

Routescan | Research