Load Assumptions

Ethereum as Fundamental Digital Infrastructure

An Assessment of the Current and the Target State

Henrik Asmus

Garbsen, July 2026

Contents

Preface

2025 ranks among the most eventful years the crypto market has witnessed to date. A summer bull run — one that even US President Donald Trump publicly praised — was followed by the events of 10 October, the largest liquidation day in recorded history. On that day, more than nineteen billion dollars in leveraged positions were liquidated, and more than one and a half million accounts were wiped out. The months that followed were grinding ones, and the aftershocks reach into 2026. I have been following the market for roughly four years, and after leaving school I had, for the first time, the chance to give sustained attention to a subject on which I could conduct my own research, develop the way I think about systems, and sharpen my judgment against a genuinely demanding object.

This year in particular has forced a question on me that I, as an investor, would otherwise easily pass over: what does all of this actually rest on? Enormous sums are moving through these systems, grand narratives circulate about payment, identity, and the management of one's own assets without an intermediary — and yet it remains strangely unclear what the thing is at its core. If it were ultimately only about payment and autonomy over one's own keys, it would be hard to account for the valuations many of these networks command. And if their value drew solely from scarcity, as with gold, there would be little new about them. That is the question driving the work: what does such a cryptographic system make possible beyond decentralized payment, and does it even need to go beyond it?

To test this question against something concrete, I needed one system that could stand in for the rest and on which the question could actually be answered. I chose Ethereum. It is the first Layer 1 with native, general programmability for Smart Contracts and has the second-largest market capitalization after Bitcoin.

According to Electric Capital 2025, it has the largest active developer base[1] and, by total value locked, the largest application ecosystem. It is the largest settlement layer for institutional and decentralized financial activity, holds the largest share of Stablecoin value held on-chain, and is being developed along a wide-ranging roadmap. Ethereum is frequently described as infrastructure, and that word is what sharpens the broad opening question into a testable one: does Ethereum's architecture meet the requirements that must be placed on fundamental digital infrastructure?

I intend to give myself and the reader a clear account of the current state of such an infrastructure and of the state it could reach under full implementation of its roadmap. My own assessment comes at the end. Beside it, and weighing just as much, is the aim of giving the reader what they need to form their own, independently of mine. I work in this field, writing and publishing on Web3, and I make this the basis for an explicit commitment to neutrality — for the closer one stands to the subject, the more one owes the reader that distance. The work also accompanies my professional applications, which I name so that the reader knows the interest against which they are reading it. It was produced with the help of digital tools, including AI-assisted systems, which I used for research, structuring, and linguistic review. The research question, the evaluation framework, and every judgment that follows from them are my own responsibility. The question of utility inevitably brings a question of returns with it. I treat the latter as a consequence of the analysis and hold my judgment to account by what proves load-bearing in the long run. Whether Ethereum meets the requirements of fundamental digital infrastructure is decided on the pages that follow.

Garbsen, July 2026

Henrik Asmus

[1] Electric Capital: Developer Report 2025. https://www.developerreport.com

Chapter 1 Introduction

1.1 Problem Statement and Relevance

Ethereum has been referred to as infrastructure for several years now — within the ecosystem itself, in the analyses of consultancies, and increasingly in the explanatory materials of regulators. Objections to this attribution do arise, pointing to the volatility of valuations, the complexity of the system, and the speculative character of a market in which billions in positions can unwind within hours. The objection from speculation bears directly on the question of infrastructure suitability, because in Ethereum the valuation of the native asset is tied to the security of the system, and no such linkage exists among established infrastructures. The objection cannot be resolved, however, as long as no one specifies what degree of independence from its own valuation a system must exhibit in order to qualify as infrastructure. The attribution therefore stands largely unchallenged, even though its examination has no point at which to begin.

The term encompasses a bundle of properties that goes beyond the designation itself. What is meant is durability over decades, neutrality toward participants, the impossibility of unilateral shutdown, and suitability for uses that no one anticipated when the system was built. Whether the word is used makes no difference. What does make a difference is that actors already behave as though these properties were in place. J.P. Morgan launched a tokenized money market fund in May 2026 that operates exclusively on Ethereum and grew from two hundred million to nearly seven hundred million dollars in roughly one month. Crédit Agricole issued, through its custody subsidiary, a MiCA-compliant euro stablecoin on the same chain, in productive operation and not as a pilot. Institutional investors orient their allocations to the expectation that the network will still exist in ten years, citing a decade of uninterrupted availability across fif-

teen protocol upgrades. Legislators, in turn, draft regulatory frameworks such as the European regulation on markets in crypto-assets, which already contain a decision as to whether a given thing constitutes an application, a platform, or an infrastructure. All of these actors examine carefully, but their examination extends only to the current state of the system. The assumption about its permanent continued existence eludes examination and is made by each actor at their own discretion.

The costs of a mistaken assumption fall on both sides. If the trust proves unfounded, the failure strikes all accumulated positions simultaneously at every level, while the escape routes have already been abandoned in the course of integration and the interconnection achieved can be reversed only at considerable cost. If the rejection proves unfounded, access is lost to a coordination layer that cannot simply be built anew under network effects, and the design of its rules, its conditions of access, and its neutrality guarantees falls to others. Both misjudgments are costly, and neither can be avoided without a standard.

Such a standard does not yet exist. The term "fundamental digital infrastructure" circulates in all these debates as a predicate that is granted or withheld, without anyone ever having determined what properties it requires. The present work first constructs the standard and then applies it to a system that continues to develop actively.

1.2 Gap

The state described in the preceding section might be dismissed as a peculiarity of public discourse, were the research to supply the missing standard. It does not. There exists a developed theory of infrastructure, there exists an extensive body of blockchain research, and there exists an established practice for assessing infrastructure systems — but each of these three bodies of knowledge does something other than what the research question requires.

Infrastructure theory has been applied to blockchain without its conceptual apparatus having been systematically applied to any individual system. The relevant works describe their subject using the vocabulary of infrastructure research and provide the groundwork on which an assessment could build.[2] They do not

[2] Cf. Ølnes, Svein / Jansen, Arild (2018): Blockchain Technology as Infrastructure in Public Sector: An Analytical Framework, and id. (2021): Blockchain Technology as Information Infrastructure in the Public Sector; also Grimmelmann, James / Windawi, A. Jason (2023): Blockchains as Infrastructure and Semicommons, Dimitropoulos, Georgios (2020): The Law of Blockchain, and Nabben, Kelsie (2023): Web3 as 'Self-Infrastructuring'. Full references are in the footnotes to Section 2.1.

arrive at an assessment because they lack the means to do so. What is missing are operationalized criteria, what is missing are thresholds, and what is missing is a procedure that forms a judgment from individual findings.

Blockchain research possesses a mature toolkit for measurement, but one that targets different properties. Its methods capture the degree of decentralization of a system or the question of whether an organization should consider adopting this technology.[3] That decentralization does not coincide with infrastructure suitability is shown by the global financial messaging system SWIFT, which is operated by a single organization and yet counts as fundamental infrastructure of international payments. A method that measures decentralization thus captures a property from whose degree suitability cannot be inferred.

The established practice for assessing infrastructure systems, finally, presupposes what the subject does not offer. Its criteria are directed at an operating organization that takes on investment commitments and is subject to oversight. Chapter 3 examines this presupposition against seven reference infrastructures and shows that it is the rule there, but not a condition of the infrastructure function. A system whose rules emerge from the consent of independent participants and for whose continued existence no authority stands guarantor escapes these criteria, without this implying anything about its suitability.

The gap is thus defined. Theory supplies concepts without an evaluative standard, blockchain research supplies standards for different properties, and established practice supplies criteria for a type of system to which the subject does not belong.

1.3 Question

The question to which the present work responds reads as follows:

Does Ethereum's architecture meet the requirements that must be placed on a fundamental digital infrastructure?

Anyone who asks this way presupposes that the requirements are fixed against which a system is to be measured. They are not fixed, and the determination of the requirements therefore precedes the assessment of Ethereum. Only once they are in place can the question be decided at all.

The question targets the structure of the system and what that structure accomplishes. It asks how Ethereum is built, what follows from the way it is

[3] Cf. Ovezik, Christina et al. (2025): SoK: Measuring Blockchain Decentralization; Srinivasan, Balaji S. / Lee, Leland (2017): Quantifying Decentralization. Wachter, Victor / Ross, Omri (2021): How Decentralized is the Governance of Blockchain-based Finance. Full references are in the footnotes to Section 2.1.

built, and whether that meets the requirements. Whether Ethereum prevails over other systems is not part of this question, because a system's prevailing depends on circumstances that remain external to its structure. An architecture can meet every requirement and still fail to prevail. This work makes no statement about Ethereum's establishment.

The system to be assessed is changing while the assessment is under way. Ethereum is being developed along a roadmap whose proposals emerge from a process in which independent developers negotiate, revise, and occasionally reject them in an open procedure, without any authority issuing them. The protection of cryptographic methods against quantum computers, whose computing power could break the signatures in use today, has been under research for years. New resource estimates, among them a white paper by Google Quantum AI from March 2026, have given the threat additional weight.[4] The Ethereum Foundation has maintained its own post-quantum research team since January 2026 and is testing initial building blocks in early implementations.[5] The post-quantum transition is running as a parallel long-term strand with a target horizon around 2029, while the roadmap for the coming years remains driven by scaling and usability. What the proposals will change concerns the core of the system: how participants reach agreement on a shared state, how the system executes code, and how an individual user can verify whether the result is correct.

A judgment about the system as it exists today loses its validity to the extent that these proposals are implemented, because they bear on the very thing that was judged. A judgment about the future system, conversely, says nothing about what a user can rely on today, because it presupposes conditions that do not yet exist. The work therefore answers the research question twice — once for the current state and once for the target state that the system would reach under full implementation of its roadmap. It names neither a point in time for full implementation nor a probability.

The order in which the two answers are given follows from the subject. The transition builds on what exists, and anyone who does not understand how Ethereum works today cannot assess what the proposals will change. The assessment of the current state therefore accomplishes two things: it answers the

[4] Google Quantum AI / Ethereum Foundation / Stanford University: Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities. Resource Estimates and Mitigations, 30 March 2026. Publication occurred three days after the cutoff date of the current state assessment.

[5] Ethereum Foundation: post-quantum security team, established in January 2026 under the direction of Thomas Coratger. https://pq.ethereum.org

research question for the present, and it describes the system in such a way that the assessment of the target state can build on that description.

1.4 Methodological Approach

Since no existing procedure makes the requirements of fundamental digital infrastructure amenable to examination, the work constructs the standard itself. It derives twelve criteria from the concept of fundamental infrastructure — from the question of what a system must exhibit in order for others to be able to rely on it lastingly. The derivation follows an established procedure of concept formation, which develops each criterion individually from the concept and derives its validity from the derivation.[1]

The twelve criteria are fixed before they are tested against seven established infrastructures: the power grid, the Internet, the road network, the financial messaging system SWIFT, the satellite navigation system GPS, cloud infrastructure, and the legal system. These systems reveal whether a criterion holds up against actual infrastructure and whether it captures properties that real infrastructures share. The probative force of the examination rests on the diversity of the seven systems, for a criterion that holds up against the physical power grid as well as the institutional order of law has been tested across the breadth of these cases.

Each of the twelve criteria is assessed on a four-level scale ranging from full compliance to an open finding. The twelve criteria stand in a hierarchy in which some function as thresholds. If such a criterion fails, this already limits the achievable overall judgment, however well the remaining eleven perform. Where a load-bearing requirement fails, strength elsewhere does not offset it. From the twelve individual findings, the overall judgment emerges through fixed rules whose application a second examiner can reproduce with the same material and arrive at the same result.

Three features of the answer follow from this design. The answer to the research question is graduated and does not reduce to a simple yes or no, because the overall judgment recognizes a spectrum of categories. The answer for the current state rests on measured states of the system as it runs today. The answer for the target state rests on specifications, prototypes, and research designs whose implementation is still pending. The target state, finally, is described in terms of the capabilities that the fully realized system would provide, because such a description outlasts changes to the roadmap, whereas a list of planned features would grow stale.

□ *Jabareen, Yosef (2009): Building a Conceptual Framework. Philosophy, Definitions, and Procedure.* International Journal of Qualitative Methods*, 8(4), pp. 49–62.*

1.5 Position, Neutrality, and Scope

The work maintains analytical neutrality toward its subject, to which it explicitly commits. Negative findings are as welcome as positive ones, and its procedure is open to inspection. This serves a dual purpose. The end result is a judgment, and the work simultaneously gives the reader the means to judge for themselves. The work therefore remains useful even when the reader does not follow its judgment.

It addresses corporate decision-makers, institutional investors, technical architects, political decision-makers, and critical academics who, as a rule, do not know Ethereum from the inside. This choice determines the register of the work. It explains what such an audience needs explained and presupposes no prior knowledge that circulates only within the ecosystem itself. It is self-published and without external peer review, which is acknowledged openly here. Precisely because no reviewing authority stands between the work and its reader, the work makes its procedure transparent at every point and leaves the reader to judge its cogency.

The standard the work applies extends as far as what can be read off the system itself. What is measurable within the protocol, it can assess. What lies outside remains out of reach. This includes the quality of the data that enter the system from outside, whose accuracy the protocol itself does not ensure, as well as the regulatory environment and social acceptance, which co-determine actual use. This limitation is intentional, because the work assesses what it can measure and states what it cannot measure. Equally, it examines the infrastructure itself and not the applications built on top of it. The line between the two follows a single question: does an application affect the suitability of the infrastructure? A service whose share in securing the network touches the thresholds of the consensus mechanism belongs in the assessment. A service that builds on the infrastructure without acting back upon it does not. Section 2.3.4 draws the line through examples.

The work also separates the current state from the future state and answers the research question for each separately, for the reasons the section on the research question has given. The current state has a data cutoff date: 27 March 2026. The cutoff binds only the assessment data, while the framing information in the introduction extends to the editorial deadline. The figures describing it are snapshots of a system that is continually changing, while its structure and the structural

findings that follow from it remain more stable. For the argument, therefore, what matters is the order of magnitude in which a figure falls and the direction it indicates, not its precise value on a given day.

From all of this follows what the work is not. It is not a technical manual explaining how to build on Ethereum, and not investment advice. Nor is it a forecast of the price of the native asset or of the future adoption of the system. Why prevailing lies outside the question is explained by the section on the research question.

1.6 Structure of the Work

The answer to the research question is built up across the following six chapters, each taking up where the previous one left off. First the standard is constructed, then it is applied to the present and to the future, and finally both findings are brought together. The reader thus has before them the sequence of the following chapters, which the next paragraphs describe in turn.

The first two chapters construct the standard. Chapter 2, Theoretical Framework and Methodology, first clarifies what makes an infrastructure fundamental and develops from that clarification the evaluation framework, the twelve criteria. Chapter 3, Infrastructure Contextualization, holds the already derived criteria against seven established infrastructures, tests whether the criteria withstand confrontation with actual infrastructure, and makes explicit what the procedure cannot accomplish. At the end of these two chapters stands a tested instrument that can be applied to Ethereum.

The two following chapters apply it, twice to the same subject at different points in time. Chapter 4, Current State Assessment, describes today's Ethereum in a coherent account and then assesses it along the three dimensions of the evaluation framework — Structural Foundation, Qualitative Load-Bearing Capacity, and Resilience and Sovereignty — arriving at an overall judgment for the present. Chapter 5, The Target State, applies the same instrument to the Ethereum that would result from full implementation of its roadmap, and likewise arrives at an overall judgment. Both chapters work with the same standard, so that their findings can be directly compared.

The final two chapters synthesize and look ahead. Chapter 6, Delta, Limits, and Overall Judgment, places the two findings side by side, examines the difference between the current and the future state, assesses whether and which limits remain with the system across both points in time, and subjects the evaluation framework to scrutiny for what no criterion captures. In this final step, the work

turns its own standard against itself and makes visible where it reaches its limit. At the end of this chapter, the research question is answered. Chapter 7, Implications, steps out of the evaluative role of the preceding chapters and thinks the answer further — toward its consequences for those who build on the system, trust it, or regulate it, and toward an outlook whose tone opens up. The outlook takes in once more the time of writing and research. Thus the work returns at its close to where the foreword began: to the voice of the one who posed the question.

At the end stands the answer to the question of whether Ethereum's architecture meets the requirements of a fundamental digital infrastructure. How it turns out is decided by the chapters that lead to it.

Chapter 2 Theoretical Framework and Methodology

This chapter sets out the theoretical and methodological foundation of the work. Section 2.1 locates the research question within infrastructure theory. Section 2.2 explains why a separate evaluation framework must be developed, documents the research gap, and demarcates against existing approaches. Section 2.3 presents the evaluation framework in full, from the evaluation logic through the twelve criteria to the scope boundaries. The work pursues a dual goal: it delivers a reasoned judgment on the research question and simultaneously documents evidence and conclusions with sufficient transparency for the reader to arrive at an independent, potentially divergent judgment.

2.1 What Is Infrastructure?

The research question of this work is whether Ethereum's architecture meets the requirements that must be placed on a fundamental digital infrastructure. Before this question can be answered, it is necessary to clarify what the term infrastructure means in an academic context. In political and media discourse, the term has become inflated: broadband networks are described as digital infrastructure, education systems as social infrastructure, charging stations as charging infrastructure, and in the crypto industry every protocol with more than a handful of users goes by the name of infrastructure. Everyday usage follows an implicit heuristic, according to which any system that appears important and is used by many deserves the title of infrastructure. For an academic assessment this heuristic does not suffice, because it draws no line between infrastructure and non-infrastructure and thereby renders the research question unanswerable. The following theoretical perspectives supply the conceptual precision that makes such a line possible.

The Economic Definition

In 2012, Brett Frischmann presented, in *Infrastructure: The Social Value of Shared Resources*, a foundational economic definition of infrastructure.[6] His central insight shifts the analytical focus: infrastructure is defined by what it enables, not by what it is. Whether a system is physical or digital, large or small, expensive or inexpensive tells us nothing, according to Frischmann, about whether it is infrastructure. What matters are three functional criteria. The first criterion is non-rivalry in consumption up to a capacity threshold: use by one actor does not preclude simultaneous use by others, as long as the system does not reach its capacity limit. The second criterion concerns the demand structure: demand for the system is driven primarily by productive activities that require it as an input, with the aggregate social benefit exceeding individual benefit. The third criterion is input character: the system serves as a factor of production for further activities and is not an end product.

The road network illustrates the three criteria. It is non-rival as long as there is no congestion, because one user's journey does not preclude another's until the road's capacity limit is reached. It generates productive demand, because trade, labor mobility, and social participation build on a functioning road network and the aggregate social benefit exceeds individual travel benefit. It has input charac-

[6] Frischmann, Brett M. (2012): *Infrastructure: The Social Value of Shared Resources*. Oxford University Press.

ter, because without roads neither a logistics sector nor widespread retail nor a labor market with commuter flows could exist. Frischmann's three criteria thus describe a specific economic role: infrastructure is an enabling structure whose value is realized in the activities it makes possible.

From this insight Frischmann draws a regulatory consequence that is directly relevant to the subject of this work. Because infrastructure generates its value through enabling and that value increases with the breadth of use, commons management — that is, non-discriminatory open access — is in many cases economically more efficient than restricting access through private property rights.[7] Ethereum's permissionless design, which allows every actor to use the network without the permission of a gatekeeper, is a technical implementation of this argument. The connection between Frischmann's infrastructure theory and blockchain systems is not an analogical inference drawn by this work: James Grimmelmann and A. Jason Windawi showed in 2023, in the *William & Mary Law Review*, a respected American law review, that blockchains meet Frischmann's infrastructure criteria and simultaneously function as semicommons — that is, as systems with intertwined private and shared resource use.[8][9] The present work builds on this academically established connection.

---

[7] Frischmann's argument builds on the empirically substantiated finding of Elinor Ostrom (1990) that communities can successfully manage shared resources themselves, without privatization or state control. Ostrom's eight design principles for sustainable commons governance are applicable to Ethereum's governance model. A systematic application to blockchain systems has been undertaken by Rozas et al. (2021). — Ostrom, Elinor (1990): *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press. — Rozas, David / Tenorio-Fornés, Antonio / Díaz-Molina, Silvia / Hassan, Samer (2021): When Ostrom Meets Blockchain: Exploring the Potentials of Blockchain for Commons Governance. *SAGE Open*, 11(1), pp. 1–14. DOI: 10.1177/21582440211002526.

[8] Grimmelmann, James / Windawi, A. Jason (2023): Blockchains as Infrastructure and Semicommons. *William & Mary Law Review*, 64(4), pp. 1097–1129.

[9] Cf. also Dimitropoulos, Georgios (2020): The Law of Blockchain. *Washington Law Review*, 95(3), pp. 1117–1200, who conceives of blockchain as "infrastructural commons".

The Sociological Extension

Frischmann's definition is economic and works for systems whose value creation is measurable. Susan Leigh Star introduced in 1999, in *The Ethnography of Infrastructure*, a perspective that goes beyond the economic determination of function.[10] Her central insight is: infrastructure is relational. It arises not from objective properties of a system, but in relation to organized practices. What for the engineer maintaining the fiber-optic network is an object of work is, for the user conducting a video conference over that network, invisible infrastructure. What for the inhabitant of an industrialized country is self-evident infrastructure can, for the inhabitant of another country, be an unattainable good. Star identified a further property that belongs to the standard repertoire of infrastructure research: infrastructure becomes visible upon breakdown. As long as it functions, it is invisible. No one thinks about the power grid until the electricity fails, and no one thinks about the Domain Name System until a website becomes unreachable.

Star's perspective alters the research question of this work in a way that shapes the entire evaluation framework. If infrastructure is relational, then the question of whether Ethereum is infrastructure cannot be answered in binary terms, because the answer depends on for whom and in what application context it is posed. For a DeFi developer whose entire business logic is built on Ethereum smart contracts, Ethereum is already infrastructure today in Star's sense, because the system is an invisible precondition for their organized practice. For a user making a single transaction, it is more of a tool. This context-dependency is precisely what the meta-pattern operationalizes that shapes the entire evaluation logic of this work: the degree of fulfillment must match the claim. The evaluation framework does not assess whether Ethereum is infrastructure in an absolute sense, but whether it meets the requirements that must be placed on a fundamental digital infrastructure — that is, context-sensitive and calibrated to the claim.

Digital Infrastructure as a Category in Its Own Right

The theory discussed so far describes predominantly physical infrastructure, and transferring it to digital systems requires a conceptual foundation of its own. Ole Hanseth and Kalle Lyytinen defined information infrastructures in 2010 as shared, open, heterogeneous, and evolving sociotechnical systems, and identified two fundamental design problems that distinguish digital from physical

[10] Star, Susan Leigh (1999): The Ethnography of Infrastructure. *American Behavioral Scientist*, 43(3), pp. 377–391.

infrastructure.[11] The bootstrap problem describes the challenge that an infrastructure must be immediately useful to early users in order to gain the momentum that enables later network effects. The adaptability problem describes the challenge that local design decisions must accommodate unlimited growth and functional uncertainty, because the future forms of use of the infrastructure are not foreseeable at the time of design. Tilson, Lyytinen, and Sørensen argued in the same year that digital infrastructure constitutes its own analytical category, with its own paradoxes — in particular the tension between openness and control — and its own research questions arising from the intertwining of technical and institutional dynamics.[12]

Both problems manifest in the subject of this work. Ethereum had to start in 2015 — financed through a crowdsale of 18.3 million US dollars — with a concrete use case that was immediately useful to early users: the possibility of deploying and executing smart contracts. The EVM, Ethereum's virtual machine, provided from day one a programmable execution environment that went beyond pure payment transactions and thus created the foundation for the network effects that sustain the system today. The adaptability problem manifests in the roadmap complexity that Chapter 5 of this work describes in detail: the system must evolve to meet growing demands for scaling, data availability, and user access, without breaking the existing user base and the applications built on the system.

The Conceptual Ground of the Evaluation Criteria

The three perspectives together define the analytical framework within which the research question stands. Frischmann supplies the economic foundation: infrastructure is an enabling structure whose open access can be justified economically. Star supplies the sociological deepening: infrastructure is relational and context-dependent, which rules out a binary answer to the research question and requires an assessment calibrated to the claim. Hanseth and Lyytinen supply the information-technological specification: digital infrastructure is subject to its own design problems that must be taken into account when assessing a concrete system. Against this backdrop the question of whether Ethereum is fundamental digital infrastructure can be posed as a question that is academically tractable: it asks whether a concrete system exhibits the properties that real infrastructures

[11] Hanseth, Ole / Lyytinen, Kalle (2010): Design Theory for Dynamic Complexity in Information Infrastructures: The Case of Building Internet. *Journal of Information Technology*, 25(1), pp. 1–19.

[12] Tilson, David / Lyytinen, Kalle / Sørensen, Carsten (2010): Research Commentary — Digital Infrastructures: The Missing IS Research Agenda. *Information Systems Research*, 21(4), pp. 748–759.

share. The three perspectives thereby supply the conceptual ground from which the evaluation criteria are analytically derived. The interplay of the economic, the relational, and the information-technological determination of infrastructure is what constitutively characterizes fundamental digital infrastructure. Chapter 3 then tests the criteria derived in this way against seven real infrastructures. The theory supplies the framework and the derivation. The seven reference systems supply the empirical test.

For a work written in German, the tradition of large technical systems is additionally relevant. Renate Mayntz conceived large technical systems in 1993 as components of function-specific infrastructure systems and identified the tension between control and self-organization as a central research problem.[13] The tension describes a fundamental conflict inherent in every infrastructure: the larger and more complex the system, the more difficult central control becomes, and the stronger the system's own dynamics. Ethereum's governance model, which operates in a decentralized manner without a formal governing authority and rests on the principle of Rough Consensus, is an extreme case of this tension. There is no central authority that can mandate protocol changes. Every change must be voluntarily adopted by a majority of node operators. Whether this model meets the requirements of adaptive governance that must be placed on fundamental infrastructure is one of the twelve assessment dimensions of the evaluation framework of this work.

An institutional-economics perspective complements the framework. Sinclair Davidson, Primavera De Filippi, and Jason Potts argued in 2018 that blockchain represents an institutional innovation that extends Oliver Williamson's governance framework beyond the dichotomy of markets and hierarchies.[14] Ethereum can be read in this perspective as a system that realizes a third form of governance: protocol-based coordination through algorithmic rules, which are neither negotiated in markets nor ordered hierarchically, but arise from the consensus of network participants.

[13] Mayntz, Renate (1993): Große technische Systeme und ihre gesellschaftstheoretische Bedeutung. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 45(1), pp. 97–108.

[14] Davidson, Sinclair / De Filippi, Primavera / Potts, Jason (2018): Blockchains and the Economic Institutions of Capitalism. *Journal of Institutional Economics*, 14(4), pp. 639–658. DOI: 10.1017/S1744137417000200.

2.2 Why a Separate Evaluation Framework?

The Research Gap

The infrastructure theory set out in the preceding section offers a mature conceptual apparatus for posing the question of infrastructure suitability. For blockchain systems in general and Ethereum in particular, this apparatus has not yet been systematically applied. Svein Ølnes and Arild Jansen transferred in 2018 and 2021 Star and Ruhleder's relational infrastructure definition and Hanseth and Lyytinen's information infrastructure theory to blockchain in general, and showed that blockchain systems can be conceived as information infrastructures.[15][16] Their work identifies properties that blockchain shares with the information infrastructures described in IS research — such as openness and heterogeneity of users — and concludes from this that the vocabulary of infrastructure research is applicable to blockchain. The works describe blockchain as infrastructure. Whether a concrete system meets the requirements of fundamental infrastructure, they do not examine: what is missing are operationalized criteria, thresholds, and an assessment mechanism. Kelsie Nabben analyzed Web3 in 2023 as self-infrastructuring — that is, as a process in which the community itself produces and continuously transforms the infrastructure it uses.[17] Her analysis works with the vocabulary of infrastructure studies and provides insights into the emergence dynamics of decentralized infrastructure, but aims at process description rather than suitability assessment. Paul Dylan-Ennis, Donncha Kavanagh, and Luis Araujo examined Ethereum in 2023 as a sociotechnical project with three imaginaries — three imaginative worlds that drive the Ethereum community: Ethereum as a world computer, as a decentralized financial infrastructure, and as a coordination protocol.[18] Their sociological analysis sheds light on the self-description and internal tensions of the community. It does not deliver an assessment of the architecture or of infrastructure suitability on the basis of verifiable criteria.

---

[15] Ølnes, Svein / Jansen, Arild (2018): Blockchain Technology as Infrastructure in Public Sector: An Analytical Framework. In: *Proceedings of the 19th Annual International Conference on Digital Government Research (dg.o '18)*. ACM, pp. 1–10. DOI: 10.1145/3209281.3209293.

[16] Ølnes, Svein / Jansen, Arild (2021): Blockchain Technology as Information Infrastructure in the Public Sector. In: Reddick, C.G. / Rodríguez-Bolívar, M.P. / Scholl, H.J. (eds.): *Blockchain and the Public Sector*. Public Administration and Information Technology, vol. 36. Springer, Cham, pp. 19–46. DOI: 10.1007/978-3-030-55746-1_2.

[17] Nabben, Kelsie (2023): Web3 as 'Self-Infrastructuring': The Challenge Is How. *Big Data & Society*, 10(1). DOI: 10.1177/20539517231159002.

[18] Dylan-Ennis, Paul / Kavanagh, Donncha / Araujo, Luis (2023): The Dynamic Imaginaries of the Ethereum Project. *Economy and Society*, 52(1), pp. 87–109. DOI: 10.1080/03085147.2022.2131280.

The gap that this work fills is thus specific: there exists no framework-based assessment that defines operationalized criteria, sets transparent thresholds, and applies a reproducible evaluation procedure to examine whether Ethereum meets the requirements of a fundamental infrastructure. The works mentioned contribute preliminary work on which this work builds, but none of them answers the research question.

Demarcation against Existing Blockchain Assessment Approaches

The need for a separate evaluation framework becomes clearer when one considers the existing assessment approaches from the blockchain literature. What the approaches share is a common thrust: they measure decentralization or organizational suitability, but they do not assess infrastructure properties in the infrastructure-theoretical sense.

In 2017, Balaji Srinivasan and Leland Lee proposed the Nakamoto Coefficient, which measures the degree of decentralization of a system as the minimum number of actors that would need to collude to compromise the system.[19] The coefficient captures an important property, but says nothing about infrastructure suitability: SWIFT, the global financial messaging system, has a Nakamoto Coefficient close to one, because it is operated by a single organization, and is nonetheless fundamental infrastructure of international payments. In 2018, Adem Efe Gencer and colleagues presented an empirical decentralization measurement via mining distribution and network topology that quantifies the operational degree of decentralization, but captures neither governance dimensions nor institutional embeddedness.[20] The most recent systematization of the research strand, a SoK study (Systematization of Knowledge) on measuring blockchain decentralization published in 2025, inventories all existing decentralization metrics and makes visible that the entire research strand is oriented toward decentralization measurement.[21] The question of whether a system meets the requirements of infrastructure is addressed by none of the metrics.

Alongside decentralization measurement there are approaches that assess blockchain suitability from an organizational perspective. In 2023, Spencer-

[19] Srinivasan, Balaji S. / Lee, Leland (2017): Quantifying Decentralization. Preprint/Blogpost. URL: https://news.earn.com/quantifying-decentralization-e39db233c28e.

[20] Gencer, Adem Efe et al. (2018): Decentralization in Bitcoin and Ethereum Networks. In: *Financial Cryptography and Data Security*. Springer, pp. 439–457.

[21] Ovezik, Christina / Karakostas, Dimitris / Milad, Mary / Kiayias, Aggelos / Woods, Daniel W. (2025): SoK: Measuring Blockchain Decentralization. arXiv:2501.18279. Also published in: Fischlin, M. / Moonsamy, V. (eds.): *Applied Cryptography and Network Security (ACNS 2025)*. Lecture Notes in Computer Science, vol. 15825. Springer, Cham. DOI: 10.1007/978-3-031-95761-1_7.

Hicken and colleagues developed a Blockchain Feasibility Assessment for organizations that evaluates whether a company should deploy blockchain technology. It thus makes an adoption decision and does not assess infrastructure suitability.[22] In 2021, Jensen and colleagues examined the governance decentralization of DeFi protocols, which operate at the application layer, not at the base layer that this work assesses.[23] In 2023, Pincheira and colleagues analyzed the infrastructure costs of blockchain systems, using the term infrastructure in the IT-technical sense — as a generic term for servers, storage, and network resources — rather than in the social sense of an enabling structure in Frischmann's sense.[24]

The conclusion from the stocktaking is clear: the existing approaches ask how decentralized a system is or whether blockchain is suitable for a particular organization. The research question of this work is whether a system meets the requirements that must be placed on fundamental infrastructure. This work answers the question by means of an evaluation framework that operationalizes infrastructure properties, grounded in infrastructure theory and validated against real infrastructures. An evaluation framework that originates from the blockchain literature and aggregates decentralization metrics cannot answer the research question, because it measures the wrong property.

The Methodology of Constructing the Evaluation Framework

The work validates its evaluation framework against seven reference infrastructures: power grid, Internet (TCP/IP), road network, SWIFT, GPS, cloud infrastructure, and legal system. The seven cases were chosen because they cover five infrastructure types: physical (power grid, road network), digital-open (Internet/TCP-IP), digital-closed (SWIFT, GPS), hybrid (cloud infrastructure), and institutional (legal system). This heterogeneity represents the breadth of the infrastructure concept and strengthens the validation force of the selection, because a criterion that is found across such diverse systems can be regarded with greater plausibility as constitutive of infrastructure.

The analytical approach was deliberately chosen because there is no established evaluation framework for assessing cryptographic infrastructure. Existing assessment approaches from telecommunications or the energy sector are

[22] Spencer-Hicken, Scott et al. (2023): Blockchain Feasibility Assessment: A Quantitative Approach. *Frontiers in Blockchain*, 6. DOI: 10.3389/fbloc.2023.1155125.

[23] Jensen, Johannes R. / von Wachter, Victor / Ross, Omri (2021): How Decentralized is the Governance of Blockchain-based Finance: Empirical Evidence from four Governance Token Distributions. In: *Proceedings of the 29th European Conference on Information Systems (ECIS 2021)*.

[24] Pincheira, Miguel et al. (2023): An Infrastructure Cost and Benefits Evaluation Framework for Blockchain-Based Applications. *Systems*, 11(4), 184. DOI: 10.3390/systems11040184.

tailored to institutional infrastructure operated by central organizations and embedded in regulatory frameworks, such that their transfer to a permissionless system without a central operating organization would, in the work's assessment, produce categorical misassumptions. The work therefore develops its evaluation framework independently, by means of analytical deduction. It decomposes the concept of fundamental digital infrastructure — as determined by the theoretical framework of the preceding section — into the constitutive properties that a system must exhibit in order to satisfy the claim. It operationalizes each of these properties as a verifiable criterion with its own definition, its own indicators, and its own thresholds.

The procedure follows the Conceptual Framework Analysis described by Yosef Jabareen. This qualitative method for framework construction proceeds in systematic phases from the identification of relevant conceptual and theoretical sources through the extraction of the load-bearing concepts to their integration into a coherent instrument.[25] The test of an evaluation framework constructed in this way is not its origin from a set of cases, but its internal coherence, the completeness of its dimensions, and its analytical productivity — that is, the question of whether it makes the distinctions that matter for the assessment. Each of the twelve criteria is self-sustaining through its own conceptual derivation, and none claims validity on the basis of the frequency with which a property appears in a sample of real systems.

Because an analytically derived evaluation framework runs the risk of being constructed past the reality of established infrastructure, Chapter 3 validates the twelve criteria. It tests them against seven structurally heterogeneous reference infrastructures and asks whether the deductively obtained requirements withstand confrontation with systems that have historically established themselves as fundamental infrastructure. This test is not a retrospective illustration, but the empirical stress test of the evaluation framework, whose results are fully documented in Chapter 3.

2.3 The Evaluation Framework

The evaluation framework is the instrument by which the research question is answered.

[25] Jabareen, Yosef (2009): Building a Conceptual Framework: Philosophy, Definitions, and Procedure. *International Journal of Qualitative Methods*, 8(4), pp. 49–62.

2.3.1 Evaluation Logic

Three fundamental principles determine the architecture of the evaluation framework. The first principle is the conditional structure of the assessment. The research question is formulated conditionally: it asks whether the architecture fulfills the claim it raises. The result is a spectrum that ranges from "Suitable" through "Suitable with Conditions" and "Suitable under Substantial Conditions" to "Conditionally Suitable" and "Not Suitable". Each category is defined, and assignment to a category follows from explicit rules laid down in the evaluation mechanism (M3 cascade).

The second principle is the meta-pattern introduced in Section 2.1: the degree of fulfillment must match the claim. It is confirmed in the validation against the seven reference infrastructures in Chapter 3. GPS offers maximum inclusivity, because any receiver worldwide can use the signal, but no neutrality, because the system is under the unilateral control of the US government, which can selectively switch off or throttle the signal. SWIFT offers global coordination in financial messaging, but no open access and no censorship resistance, because states can exclude individual participants from the system. Both systems are fundamental infrastructure, even though neither of them fulfills all conceivable infrastructure properties at maximum degree. The meta-pattern prevents the fallacy that Ethereum must fulfill every criterion at maximum degree in order to qualify as infrastructure. At the same time it prevents the reverse fallacy that any arbitrary degree of fulfillment is acceptable: the degree of fulfillment must match the specific claim, and the claim that Ethereum makes is that of a neutral, permissionless, censorship-resistant infrastructure.

The third principle is a deliberate design decision: the assessment is hierarchical-rule-based, not additive. The work does not use a weighted scoring model in which each criterion receives a percentage value and the overall assessment follows from the sum of the weighted individual values. Instead, a rule-based hierarchy operates in which Critical Conditions function as a threshold: if a Critical Condition is not met at the required level, this limits the best possible overall assessment, regardless of how well the other criteria perform. In an additive model, a system that stands at "Open" on Neutrality — a Critical Condition — but reaches "Met" on all other eleven criteria could still achieve a good overall result. This contradicts the infrastructure logic: a system that is not neutrally accessible cannot fulfill the claim of a fundamental infrastructure, regardless of its performance in other dimensions. The hierarchy reflects this logic. It trades accumulation sensitivity — the ability to aggregate many small

limitations into a worse overall result — for reproducibility and transparency: a second researcher assessing the same material arrives at the same result, because the rules are explicit and require no implicit weighting decisions.

2.3.2 The Three-Level Logic

Each of the twelve criteria is examined at three levels. The protocol level asks what the architecture makes possible. The operational level asks what of that is actually realized. The infrastructure level asks whether the interplay of architectural possibility and operational reality fulfills the infrastructure claim. This three-level logic is the central analytical tool of the entire work. It prevents the assessment from stopping at the architecture and ignoring operational reality, which would lead to overassessment, because every well-designed protocol exhibits the right properties at the architectural level. Conversely, it rules out assessing operational reality without taking architectural potential into account, which would lead to underassessment, because operational deficits may be remediable through protocol changes. The infrastructure level synthesizes both perspectives and forces the question that is decisive for the infrastructure claim: is what is architecturally possible and operationally realized sufficient for the claim?

An example makes the three levels tangible. Censorship resistance — one of the criteria in the Critical Conditions category — operates at all three levels with different findings. At the protocol level, Ethereum's architecture enables any valid transaction to be included in a block: the protocol itself contains no mechanism that discriminates transactions by content or sender. At the operational level, a divergent reality appears: the large majority of blocks is produced by a small number of builders, who could theoretically exclude transactions, and a considerable share of blocks is built in compliance with OFAC sanctions lists. At the infrastructure level, the decisive question arises: does the protocol-level possibility of inclusion suffice, or must operational censorship resistance be present for the system to fulfill the claim of a neutral infrastructure? The answer, which Chapter 4 justifies in detail, depends on whether the operational discrepancy undermines the infrastructure claim or whether mitigation mechanisms — such as protocol-level inclusion lists (FOCIL) — can close the gap between protocol design and operational reality. The final assessment follows the principle of the weakest link: the level with the lowest degree of fulfillment determines the overall assessment of the criterion.

2.3.3 The Three Building Blocks of the Evaluation Framework

The evaluation framework consists of three interlocking building blocks: the Implementation Status Taxonomy (M1), the thresholds for the twelve criteria (M2), and the Criteria Hierarchy and Assessment Mechanism (M3 cascade).

Implementation Status Taxonomy (M1)

The taxonomy classifies all Ethereum features according to their implementation maturity in five stages. Features with the status IMPL (implemented) are active on mainnet and in productive use, stable for at least three months and free of critical bugs. Features with the status DEPL (in deployment) are implemented in client releases and active on testnets, with an announced mainnet date. Features with the status PLAN (planned) have an accepted EIP, are in Review or Last Call status, are on the official roadmap, and have an estimated time horizon of less than 24 months. Features with the status RES (research) are mentioned in roadmap documents and are the subject of active research, but without a final specification. Features with the status DEP (deprioritized) are explicitly deferred or abandoned.

The boundary between DEPL and PLAN is the central epistemic dividing line of the work. IMPL and DEPL are treated as present or secured, because their implementation is operationally confirmed or immediately imminent. PLAN and RES are treated as conditional, with explicit uncertainty marking, because their realization depends on future governance decisions, technical feasibility, and community consensus. Everything below DEPL is projection, and the taxonomy indicates the degree of projection. Five stages were chosen because they map the relevant epistemic gradations without generating unnecessary granularity: three stages (implemented, planned, not planned) would flatten the difference between a feature in testnet deployment and a feature in early research, which is decisive for the robustness of the assessment. The taxonomy produces a certainty gradient that determines the entire work and is operationalized in the integrated robustness assessment in Section 6.1 according to the maturity level of the load-bearing features.

Thresholds: Twelve Criteria in Three Categories (M2)

The twelve criteria are divided into three categories that play different roles in the hierarchy of the evaluation framework. Each criterion is assessed on a four-stage scale. "Met" means the requirement is fully implemented and there are no structural limitations. "Met with Qualification" means the requirement is fundamentally met, with identified limitations that do not fundamentally undermine the claim. "Conditionally Met" means fulfillment depends on conditions that are not yet secured, but for which a recognizable path exists. "Open" means the requirement is not met and no operational path to fulfillment exists. The criterion assessment results from the argumentative synthesis of the indicators along the three-level logic. The indicators provide the evidence base, but the assessment does not follow a mechanical counting rule, because the indicators of a criterion carry different weight for the infrastructure claim. A counting rule that would treat, say, three out of four indicators at "Met" as sufficient for an overall "Met" would treat all indicators as equivalent, which is substantively untenable. The argumentative synthesis instead enforces a structured justification that is transparently documented in each individual assessment. The reader can assess the evidence and the conclusion independently. If an indicator does not change between two assessment time points, its level is retained. An unchanged indicator cannot fall below the level it achieved in the initial state.

The following paragraphs introduce the twelve criteria in the order of their hierarchical function. First come the Critical Conditions, which form the foundation of the infrastructure claim and whose non-fulfillment limits the best possible overall assessment. Then follow the Structural Conditions, which determine the load-bearing capacity of the system. Finally come the Qualitative Criteria, which differentiate the degree of suitability within the category reached.

Security and Trust Load (I.2) operates on two dimensions. The first dimension asks whether attacks on Ethereum are economically infeasible. Economic security is measured by the cost of a 34-percent attack, for which the threshold for "Met" stands at more than 30 billion US dollars, derived from the CoinMetrics study "Breaking BFT" of 2024. The 34-percent threshold follows directly from the blocking minority of the Casper-FFG consensus algorithm: an attacker controlling 34 percent of staked ETH can block finalization. The second dimension asks how much residual trust in third parties the user must place in order to use the system securely. This dimension is complementary to criterion II.3 (Independent Verifiability): whereas II.3 asks whether the possibility of verification exists, I.2 asks how much residual trust remains despite that possibility. The answers

can diverge, because a system that enables verification can nonetheless generate a high trust load if the majority of users access the system through centralized intermediaries such as RPC providers.

Minimal Load-Bearing Guarantees (I.4) asks whether Ethereum offers a minimum set of guarantees that are maintained even under stress. These include Finality — the irreversibility of finalized transactions — Liveness — the ability to continue producing blocks under stress — a degradation mode in the event of a validator outage, and censorship resistance under attack, which is co-assessed via a cross-reference to criterion II.1. The time-to-finality threshold of under 15 minutes for "Met" derives directly from the Beacon Chain specification, which provides for two epochs of 32 slots each at 12 seconds per slot — approximately 12.8 minutes — for finalization.

Neutrality and Censorship Resistance (II.1) is the criterion with the largest indicator set in the evaluation framework. Seven indicators examine neutrality at two levels: at the transaction level, whether individual transactions are censored or delayed by builders, relays, or regulation, and at the system level, whether a state or group of actors can instrumentalize Ethereum as a whole. The OFAC compliance rate of blocks, builder concentration, the existence of a protocol-level censorship resistance mechanism, the maximum censorship delay, the geopolitical jurisdictional diversity of validators, the market share of the largest liquid staking provider, and the cumulative concentration of the three largest staking providers — regardless of whether they appear as a liquid staking protocol or as a custodial service — are each measured individually and synthesized in the three-level logic. The 50-percent threshold for OFAC-compliant blocks as the boundary to "Met" goes back to Uri Klarman's watershed analysis. The 33-percent and 66-percent thresholds for staking concentration are derived from the blocking minority and the supermajority of the Casper-FFG consensus algorithm.

Functional Irreplaceability (I.1) targets depth of entrenchment, measured by network effects, liquidity concentration, and ecosystem lock-in. It explicitly does not ask about substitutability by alternative systems, because the work assesses Ethereum in isolation against its own claim, not against competitors. Three indicators examine it: dominance in DeFi Total Value Locked across all Layer-1 systems, the share of global stablecoin issuance, and the size of the developer ecosystem as measured by public repository activity. The thresholds are set such that dominance above 50 percent reaches "Met" and a drop below 30 percent triggers "Open". The third indicator — the developer ecosystem — is transparently normatively grounded, because the data source, the Electric Capital Developer Report, measures public repository activity as a proxy for developer network

effects: an off-chain metric with greater measurement uncertainty than the on-chain verifiable TVL and stablecoin data.

Coordination Function (I.3) asks whether Ethereum provides a coordination service that could not be rendered without the system, or only at significantly higher cost. The criterion name and the overarching pattern from the infrastructure contextualization are oriented toward the function of real infrastructures: power grids coordinate supply and demand, SWIFT coordinates financial messages between banks, the Internet coordinates packet routing between networks. Three indicators operationalize this question. The first is the coordination scope, which measures the breadth of the coordination service, from pure settlement function to settlement, execution, verification, and data availability. The second is Layer-2 settlement dependency, which measures how many Layer-2 systems use Ethereum as a coordination point. The third is cross-Layer-2 composability, which captures the coherence of the coordination, because Ethereum's strongest coordination property — atomic composability within a single transaction — is lost at the Layer-2 level and thereby fragments the coordination service.

Long-Term Stability (III.1) asks whether Ethereum can operate over decades without inherent design decisions destabilizing the system. This criterion holds a special position within the Structural Conditions as the most critical long-term condition, because without a solution for State growth and client diversity, long-term load-bearing capacity cannot be guaranteed. The five indicators measure client diversity on the execution layer, client diversity on the consensus layer, the existence of a State growth mitigation solution, the stability of the last five protocol upgrades, and the concentration of the ten largest validators. The boundary to "Met" is formed by the 33-percent threshold of the blocking minority from Casper FFG.

The Qualitative Criteria differentiate the degree of suitability within the category reached. They decide the how, not the whether.

Open Generativity (II.2) asks whether anyone can build and innovate on Ethereum permissionlessly — whether, that is, new applications, protocols, and business models can build on the system without a gatekeeper. The three indicators measure permissionless deployment (can anyone deploy without permission?), atomic composability (can smart contracts interact with one another within a single transaction?), and open-source tooling (are the most important development tools open-source?). The threshold between "Met" and "Met with Qualification" lies at documented Gas limitations and cross-Layer-2 restrictions that fragment composability.

Independent Verifiability (II.3) asks whether the correctness of smart contracts and protocol State can be verified without trust in third parties. This criterion represents the supply side to Security and Trust Load (I.2), whose complementarity was explained there. The indicators measure the availability of formal verification tools, the existence of a formal EVM specification, and the prevalence of security audits as standard practice before the publication of DeFi protocols.

Low-Threshold Inclusivity (II.4) asks whether participation in Ethereum as a user, developer, validator, or verifier is possible without prohibitive barriers. With five indicators, this criterion carries the second-largest indicator set in the evaluation framework alongside Long-Term Stability (III.1), because inclusivity operates at the intersection of several barrier types: the hardware requirements for operating a full node, the hardware requirements for operating a validator, staking access below the 32-ETH threshold, Layer-2 transaction costs compared to conventional transaction costs, and usability without cryptographic prior knowledge. The last indicator is tied to the implementation status of Account Abstraction (EIP-7702) and measures the complexity barrier: can a user who is able to operate online banking and app stores, but has never interacted with cryptocurrencies, use Ethereum? The grounding of this indicator is transparently normative, which means the reader may weight its significance differently.

Adaptive Governance (III.2) asks whether Ethereum can adapt to changed requirements without the adaptation process itself destabilizing the system. The three indicators measure the functionality of the EIP process as a documented governance mechanism, the capacity for hotfixes in the event of critical bugs within 48 to 72 hours, and the ability to resolve contentious decisions without a chain split, which has been empirically demonstrated since the DAO fork of 2016.

Sovereign Portability (III.3) asks whether Ethereum is free from proprietary dependencies that could produce vendor lock-in. The indicators examine whether the complete blockchain history can be exported, whether the system is entirely open-source with multiple independent clients, and whether assets can be moved between layers without trust in individual bridges. The threshold between "Met" and "Met with Qualification" lies at practical barriers to data export and at trust assumptions in bridge design.

Hardware Agnosticism (III.4) asks whether Ethereum can be operated on heterogeneous hardware without specific hardware requirements limiting participation. The three indicators measure cloud independence (no single provider above 50 percent), geographical distribution (nodes in at least ten countries each with more than one percent), and architecture support (all major clients on x86

and ARM). This criterion holds a special position among the twelve. None of the seven reference systems exhibits the specific risk of cloud concentration in the form that exists with Ethereum: the power grid and the road network are physically distributed. The Internet is architecturally decentralized. SWIFT and GPS operate on proprietary infrastructure. Ethereum, by contrast, runs on consumer hardware and cloud servers, and a concentration of node infrastructure on a small number of cloud providers creates a physical dependency that undermines decentralization at the logical layer. The criterion is motivated by the specific claim: it exists because the subject requires it, and it functions as a precondition for several other criteria, because client diversity (III.1) presupposes hardware independence and validator operation on consumer hardware must remain possible in order to secure inclusivity (II.4) and jurisdictional diversity (II.1). The work marks this criterion transparently as a special case whose justification derives from the architectural requirements of the subject, which no analogue in any of the seven reference infrastructures exhibits.

Criteria Hierarchy and Assessment Mechanism (M3 Cascade)

The M3 cascade converts the twelve individual assessments into an overall judgment. The twelve criteria are organized into three hierarchical levels that determine the evaluation logic. At the first level stand the three Critical Conditions (I.2, I.4, II.1), which form the foundation of the infrastructure claim. At the second level stand the three Structural Conditions (I.1, I.3, III.1), which determine the load-bearing capacity of the system. At the third level stand the six Qualitative Criteria (II.2, II.3, II.4, III.2, III.3, III.4), which differentiate the degree of suitability.

The cascade examines in an explicit sequence of five steps. In the first step it is checked whether at least one Critical Condition stands at ”Open”. If so, the overall judgment is capped at ”Conditionally Suitable”, regardless of all other assessments. In the second step it is checked whether at least one Critical Condition stands at ”Conditionally Met”. If so, the overall judgment is limited to ”Suitable under Substantial Conditions”. In the third step the Structural Conditions are examined according to the number of their weak ratings. A condition is considered weak if it stands at ”Met with Qualification”, ”Conditionally Met”, or ”Open”. If none of the three Structural Conditions is weak, the top category is attainable, provided the Critical Conditions also stand at least at ”Met”. If exactly one stands weak, the judgment is limited to at most ”Suitable with Conditions”. If two or more stand weak, it is limited to at most ”Suitable under Substantial Conditions”. The capping checks of the first three steps take precedence

over the positive definitions of the categories: if a cap applies, it determines the highest possible category, and the positive definition assigns the judgment within that upper limit. In the fourth step the Qualitative Criteria determine the degree within the suitability category reached, according to a degree scale of Good (all six qualitative criteria at least at "Met with Qualification") through Satisfactory (four or five) and Adequate (two or three) to Poor (fewer than two). In the fifth step the overall judgment is formulated. Identified tensions between criteria enter the cascade by capturing for each tension the mitigation path with its implementation status according to the M1 taxonomy, and the degree of fulfillment of the affected criterion thus determined flows into the responsible step.

The cascade rule in the second step was added during the application of the evaluation framework. In the original version, the cascade provided for only two states for critical conditions: "Open" as the capping threshold and all levels above it as passable. The application of the evaluation framework to the current state identified a situation that the binary distinction did not capture. A Critical Condition at "Conditionally Met" is qualitatively a different state from one at "Met with Qualification", because the core property of the infrastructure claim is not yet secured at the required level on the protocol side. The hierarchy defines Critical Conditions as the foundation of the infrastructure claim. A foundation at "Conditionally Met" means that the path to fulfillment is recognizable but not secured, while a foundation at "Met with Qualification" means that the property is fundamentally present, with documented limitations. The cascade must reflect this distinction, and the addition does so through an independent judgment category. The addition is formulated generically: it applies to every application of the evaluation framework in which a Critical Condition stands at "Conditionally Met", regardless of which criterion is affected and which system is being assessed.

The cascade produces five possible judgment categories, whose definitions clarify what a category means and what it excludes. The degree is determined for each category reached by the Qualitative Criteria. It differentiates suitability within the category reached without shifting the category itself. "Suitable" means that all Critical and Structural Conditions stand at least at "Met". It expressly does not mean operational readiness for deployment, global scaling, or completeness, because these factors describe market and adoption dynamics, not infrastructure properties. "Suitable with Conditions" means that at least one Critical Condition or exactly one Structural Condition exhibits a limitation, as long as no capping rule of a higher level applies. The limitation qualifies the claim without fundamentally undermining it. "Suitable under Substantial Conditions" means

that at least one Critical Condition stands at "Conditionally Met": the core property of the infrastructure claim is not yet secured at the required level, and suitability is tied to named, verifiable conditions. "Conditionally Suitable" means that at least one Critical Condition stands at "Open" and a fundamental requirement is missed, for which, however, a recognizable path to fulfillment exists. "Not Suitable" means that fundamental requirements are missed without a recognizable path to fulfillment.

2.3.4 Scope and Boundaries

The evaluation framework operates within explicit boundaries that define its informative value and demarcate its applicability.

The first boundary is the oracle boundary. The evaluation framework assesses the on-chain state in full — that is, all properties that are measurable within the protocol. What lies outside the protocol cannot be assessed by the evaluation framework in principle. The quality of data that enter the system is a central example: a smart contract can correctly execute the logic of an insurance payout, but if the data input that reports the triggering event is manipulated, the correct execution is of no use. The regulatory environment in which Ethereum operates, and the social acceptance that determines its actual use as infrastructure, are further variables that are not measurable within the protocol and thus lie outside the scope of assessment. The oracle boundary is a restriction in favor of methodological rigor: the evaluation framework assesses what it can assess and documents what it cannot.

The second boundary is the application-layer dividing line. Applications built on Ethereum are treated only insofar as they influence the infrastructure itself. The dividing line follows a question: does the application influence infrastructure suitability? A liquid staking provider whose market share touches the consensus thresholds of the protocol is relevant, because its concentration affects the neutrality of the infrastructure. An NFT marketplace is not. The dividing line is applied where a finding decides its location — for example, at the oracle boundary in Chapter 4.

The third boundary is the current state/target state dichotomy. The work assesses two temporal planes with the same evaluation framework. The current state describes what Ethereum is in the first quarter of 2026, and rests on operational evidence — that is, features with the taxonomy status IMPL. Features with the status DEPL are included insofar as they already shape the running system, but carry no assessment on their own. The target state describes what Ethereum can become on the assumption of complete roadmap implementation,

and rests on architectural coherence with graduated implementation maturity — that is, features with the status PLAN or RES. The dichotomy is a methodological decision: Ethereum is an evolving system whose current properties differ from the properties the roadmap holds out. Without the dichotomy, the work would either have to assess the current state and ignore the roadmap — which would obscure the system's dynamics — or take the target state as given, which would conceal the epistemic uncertainty. The dichotomy allows the reader to distinguish between what is operationally confirmed today and what the roadmap holds out, and to assess the robustness of this undertaking on the basis of the taxonomy stages. It shows that the leap between the current state and the target state is tied to concrete implementations whose robustness Chapter 6 analyzes separately.

Infrastructure suitability is a multidimensional property that resists reduction to a single metric. The work has therefore chosen a qualitative approach. A quantitative assessment would require weightings between the criteria that would either have to be set arbitrarily or statistically derived from a larger sample. Both are not possible for the subject: arbitrary weightings would produce a false precision that conceals the actual uncertainty, and statistical derivation fails because the population of cryptographic infrastructure systems is too small for statistical methods. The completeness of the twelve criteria is reflected on in Section 6.3, where four properties are identified that the evaluation framework cannot measure: permissionless composability, cryptographic system State verification, decentralized self-development capacity, and privacy as an architectural infrastructure property. These four documented boundaries of the evaluation framework mark the point at which validation against the reference material reaches its limit and the subject exhibits properties that none of the seven infrastructures shows in comparable form.

That the criterion assessment rests on argumentative synthesis and does not follow a counting rule was justified in the section on evaluation logic. The methodological justification for this procedure lies in the nature of qualitative application of the evaluation framework: the strength of a rule-based but interpretive assessment lies in the theory-guided weighing of the indicators, not in their mechanical aggregation into a point value. The transparency of the procedure enables the reader to independently trace each individual assessment on the basis of the documented evidence and the reasoned conclusion, and, if desired, to arrive at a divergent judgment. This dual goal — introduced in the chapter introduction — is a methodological strength that defuses a frequent criticism of

qualitative research: the work is useful even if the reader does not agree with the judgment, because it provides the foundation for an independent judgment.

A final methodological boundary is the survivorship bias in the validation base: the seven reference infrastructures are recognized systems. Failed ones are not included, and the validation therefore supports only the reading of the twelve criteria as necessary conditions. Section 3.3 elaborates on this boundary.

Chapter 3 documents the validation of the twelve criteria and the methodological boundaries of this procedure. The complete individual analyses of the seven infrastructures are in Section 3.1, the consolidation and validation results in Section 3.2.

Chapter 3 Infrastructure Contextualization: Validation of the Evaluation Criteria against Seven Reference Infrastructures

Chapter 2 set out the theoretical framework of this work and analytically derived the evaluation framework with its twelve criteria from the concept of fundamental digital infrastructure. A deductively constructed evaluation framework must, however, prove itself against the reality of established infrastructure, demonstrating that its requirements do not miss the mark. The present chapter provides this proof. It confronts the twelve criteria with seven structurally heterogeneous reference infrastructures and examines whether the analytically obtained requirements find purchase there, whether they can distinguish between the systems, and where a criterion finds no counterpart in the reference systems. The examination is therefore validation, not origin. It tests the criteria rather than generating them.

The selection of the seven reference infrastructures and their assignment to five infrastructure types was typologically justified in Chapter 2. The individual justifications follow here. The selection followed two guiding principles: each infrastructure type should be represented by at least one system, and the chosen systems should be sufficiently documented to permit a source-based analysis. Within physical infrastructure, the power grid was chosen because, as a regulated natural monopoly, it is the paradigm case of public infrastructure, and the road network because, through common use right, it exhibits a legally unique form of open access. The Internet stands without alternative as digital-open infrastructure on account of its protocol-based architecture and its generative capacity. SWIFT and GPS were chosen as digital-closed systems because, despite their global reach, they exhibit fundamental restrictions in control and access, and thus function as counterexamples to open architecture. Cloud infrastructure represents the hybrid type because it possesses de facto infrastructural sig-

nificance without public-law status. The legal system, finally, was included as the only non-technical infrastructure because the question of coordination and enforcement would remain incomplete without an institutional reference.

Each of the seven infrastructures is examined through three guiding questions: what is the core function of the system and why is it socially relevant, which properties enable that function, and what happens when the system fails or is manipulated? The answers make it possible to test each of the twelve criteria against the concrete manifestation of the system. This chapter conducts the examination in three steps. It characterizes the seven infrastructures compactly along the guiding questions. It holds the twelve criteria against the systems and shows in what form and with what variance each requirement appears there. From this confrontation it draws the balance sheet for the load-bearing capacity of the evaluation framework.

3.1 The Seven Reference Infrastructures

3.1.1 Power Grid

The power grid is the physical infrastructure for the transmission and distribution of electrical energy from generation facilities to end consumers. In Germany, the transmission grid is operated by four operators under the supervision of the Bundesnetzagentur, and the transport layer exhibits the properties of a natural monopoly: high sunk costs and economies of scale make the parallel construction of a second grid physically and economically pointless.[26] The core function of the power grid — the permanent provision of electrical energy for all downstream sectors — is not substitutable by any alternative system. No other infrastructure sector functions without electricity: telecommunications, water supply, transport systems, healthcare, and financial markets are all electrically dependent, making the power grid a meta-infrastructure whose failure produces cascading effects beyond its own sector.[27]

The infrastructural properties of the power grid are shaped by two principles. The layer separation through unbundling, which Europe has progressively implemented since 1996, structurally separates generation, transport, and retail from one another and ensures that every generator can feed in without having to

[26] Europäische Richtlinie (EU) 2019/944 über gemeinsame Vorschriften für den Elektrizitätsbinnenmarkt (successor to Directives 96/92/EC, 2003/54/EC, and 2009/72/EC). The unbundling principle was first introduced in Directive 96/92/EC of 1996.

[27] Fraunhofer IEE: Systemsicherheit und Netzstabilität, 2023.

own or operate the grid.[28] The grid is thereby source-neutral: whether electricity comes from a nuclear power plant, a wind turbine, or a solar installation is irrelevant to the transport layer. Regulation by the Bundesnetzagentur, which monitors grid tariffs, investments, and system stability, is constitutive for the functioning of the grid, because the monopoly character of the transport layer would, without external control, tilt into discrimination or exploitation.[29] In normal operation the system is extraordinarily stable: the average outage time in Germany is 10 to 15 minutes per year, one of the best values worldwide.[30]

The infrastructure-specific observation of the power grid concerns the physically enforced real-time coordination without margin for error. Electrical energy cannot be stored in the grid. Generation and consumption must be in exact equilibrium at every moment, measured by the grid frequency of 50 Hz in Europe. Deviations above a certain threshold produce cascade failures, because the system admits no gradual loss of quality. In the blackout in the northeastern United States and Canada in 2003, a single unpruned tree led to a cascade failure that affected 55 million people and caused estimated economic damage of 4 to 10 billion US dollars.[31] None of the other seven systems examined exhibits this zero tolerance for desynchronization. The Internet transports packets asynchronously and tolerates latency fluctuations. The road network organizes itself in a decentralized manner and absorbs local disruptions through alternative routes. The power grid, by contrast, operates permanently at the limit of its physical tolerance, and society becomes aware of this property only when the limit is exceeded. This observation illustrates Star's insight about visibility upon breakdown.[32]

3.1.2 Internet (TCP/IP)

The Internet is a global network of approximately 70,000 autonomous systems communicating via the shared protocol TCP/IP.[33] It comprises the physical transport layer of fiber optics, copper, and radio technology, the logical layer of protocols (TCP/IP, BGP, DNS), and the application layer (HTTP, SMTP, and their successors). The core function consists in vendor-independent, technology-neutral networking: TCP/IP abstracts the physical transport layer, so that for the application level it is irrelevant whether data are transmitted via fiber optics,

[28] Europäische Richtlinie (EU) 2019/944, op. cit.
[29] Bundesnetzagentur: Bericht zur Eigenkapitalverzinsung, 2024. https://www.bundesnetzagentur.de
[30] Bundesnetzagentur: SAIDI-Wert Deutschland, Monitoringbericht 2024.
[31] U.S.-Canada Power System Outage Task Force: Final Report on the August 14, 2003 Blackout, 2004.
[32] Star, Susan Leigh (1999): The Ethnography of Infrastructure. *American Behavioral Scientist*, 43(3), pp. 377–391.
[33] APNIC: BGP Routing Table Analysis, 2024.

copper, or radio. Before TCP/IP, network communication was dominated by proprietary, mutually incompatible systems. The open protocol replaced these monopolies with a freely available standard that anyone can implement.

The outstanding infrastructural properties of the Internet lie in the combination of open protocol architecture and decentralized governance. The Internet has no owner and no central regulator. Instead, the IETF develops technical standards on the basis of "rough consensus and running code," while ICANN administers the Domain Name System and IP address allocation.[34] 98 percent of international data traffic flows through more than 500 submarine fiber-optic cable systems with a total length of 1.5 million kilometers, with resilience geographically unevenly distributed: Europe and North America have massive redundancy, while island states can be cut off by a single cable break.[35][36] In failure behavior the Internet shows a fundamentally different pattern from the power grid: it typically loses performance gradually, gains latency, and loses bandwidth, while BGP automatically calculates alternative routes.[37]

What distinguishes the Internet from the other infrastructures is the end-to-end principle as an architectural decision that has enabled permissionless innovation. Saltzer, Reed, and Clark formulated it in 1984: intelligence sits at the endpoints, the network itself is "dumb" in the sense that it transports packets without knowing or deciding what they contain.[38] Anyone who wanted to build a new application needed no permission from a network operator. Van Schewick showed in 2010 that this architectural property had direct economic consequences. The end-to-end principle reduced innovation barriers to nearly zero and enabled the emergence of the World Wide Web, e-commerce, cloud computing, and social media, without any central authority having to authorize or even anticipate such developments.[39] None of the other infrastructures examined exhibits this architecturally anchored permissionlessness in comparable form: the power grid qualifies feeders technically, SWIFT ties access to compliance, and even the road network, beyond common use, has access restrictions in the form of special use permits.

---

[34] IANA Functions Stewardship Transition, 2016. https://www.icann.org

[35] TeleGeography: Submarine Cable Map, 2025. https://www.submarinecablemap.com

[36] Internet Society: Submarine Cable Resilience Report, 2015.

[37] Kentik: Red Sea Cable Disruptions Impact Analysis, 2024–2025.

[38] Saltzer, Jerome H. / Reed, David P. / Clark, David D. (1984): End-to-End Arguments in System Design. *ACM Transactions on Computer Systems*, 2(4), pp. 277–288.

[39] Van Schewick, Barbara (2010): *Internet Architecture and Innovation*. MIT Press.

3.1.3 Road Network

The road network in Germany encompasses more than 830,000 kilometers of traffic routes, from motorways through federal and state roads to municipal roads, and is the most extensive man-made infrastructure in physical terms.[40] Its core function lies in comprehensive territorial coverage: it reaches every location, every building, every settlement. Neither rail nor waterway nor air routes provide this universal connectivity. The last mile of every transport chain is almost always a road. The economic significance of this coverage is empirically documented: Aschauer estimated a positive output elasticity for public capital, and particularly for transport routes, thereby making visible the contribution of infrastructure to overall economic productivity.[41]

The infrastructural properties of the road network are shaped by the combination of public provision and decentralized operation. The provision and maintenance of transport infrastructure is, in German law, a public service obligation and thus a direct task of the public authorities, distributed vertically among the federal government, the states, and municipalities.[42] Unlike the power grid with its control centers or the Internet with its BGP routing, the road network has no central operational control: every road user navigates independently, and coordination is taken over by decentralized rules and physical infrastructure such as traffic lights and signs. In failure behavior the road network resembles the Internet: local outages are typically absorbed by alternative routes without the overall system collapsing, as the collapse of the I-35W bridge in Minneapolis in 2007 exemplarily demonstrated, leading to local traffic diversions rather than a systemic collapse of the road network.[43] The real threat to the road network lies in the creeping decay caused by underinvestment. It unfolds over decades, and the OECD puts the global infrastructure financing gap at approximately one trillion US dollars annually.[44]

The unique property of the road network lies in common use right as the legal form of open access. German road law recognizes common use as the designated, typically free-of-charge, and equal use by all within the framework of traffic regulations.[45] Access to the road is a subjective public right arising from

[40] Bundesministerium für Digitales und Verkehr: Verkehr in Zahlen, 2024.
[41] Aschauer, David Alan (1989): Is Public Expenditure Productive? *Journal of Monetary Economics*, 23(2), pp. 177–200.
[42] Bundesfernstraßengesetz (FStrG) §§ 3 and 5; Straßengesetze der Länder.
[43] National Transportation Safety Board: Collapse of I-35W Highway Bridge, Minneapolis, Minnesota, August 1, 2007. Highway Accident Report NTSB/HAR-08/03, 2008.
[44] OECD: Global Infrastructure Outlook, Infrastruktur-Finanzierungslücke, 2023.
[45] Bundesfernstraßengesetz (FStrG) §7, op. cit.

the statutory dedication and revocable only by changing that dedication. This construction differs fundamentally from all other forms of access in the sample: the power grid offers regulated access tied to technical prerequisites and contractual relationships. The Internet requires a connection contract with a provider. SWIFT restricts access to authorized financial institutions. Common use right is one of the oldest forms of institutionally secured inclusivity, and it defines access to infrastructure neither technically nor economically but juridically as a fundamental right. Frischmann's argument that open access to infrastructure can be economically efficient finds its long-established legal counterpart in common use right.[46]

3.1.4 SWIFT

SWIFT (Society for Worldwide Interbank Financial Telecommunication) is a messaging network for standardized communication between financial institutions. It transmits payment instructions, securities transactions, and treasury operations, but does not itself move money. The actual value transfer is handled by correspondent banks and payment systems such as T2 in Europe or Fedwire in the United States. Around 11,000 financial institutions in 212 countries and territories are connected, at an estimated daily transaction volume of approximately 5 trillion US dollars.[47] The core function of SWIFT lies in the standardization and transmission of financial messages, and its quasi-monopoly position rests on the network effect of universal adoption: China's CIPS processed only 6.6 million transactions in 2023 compared to SWIFT's approximately 11 billion annually, even though it was conceived as a political alternative.[48]

SWIFT is organized as a Belgian cooperative owned by its member banks, supervised by G10 central banks with the Belgian National Bank as lead overseer. Access is *permissioned by design*: only authorized financial institutions can participate, private individuals have no direct access, and AML, KYC, and CTF compliance are mandatory prerequisites. The actual settlement runs through bilateral correspondent relationships, a federated architecture of independent payment systems in various jurisdictions. The chain of trust is long and opaque for the end user: whether a transfer has been correctly processed is not something the end user can verify independently, and cross-border payments take 1 to 3 days with stacked fees and opaque intermediate states.

[46] Cf. Frischmann 2012.
[47] SWIFT: Annual Review 2024. https://www.swift.com
[48] People's Bank of China: CIPS Annual Report, 2023.

SWIFT exhibits a property that appears in none of the other systems in the sample in this form: the combination of network effect as the monopoly base and vulnerability to geopolitical instrumentalization. The system's lack of alternatives rests on network-effect lock-in, even though the SWIFT protocol would be technically replicable. Virtually every financial software system worldwide processes SWIFT formats, and the message standards established by SWIFT (MT messages, ISO 20022) have developed a life of their own that extends beyond the network. The geopolitical dimension shows in the sanction disconnections: in 2012, Iranian banks were cut off from the network by EU decision. In 2022, Russian banks were cut following the invasion of Ukraine.[49] This instrumentalization reveals a tension that is central to the infrastructure question: SWIFT claims neutrality as messaging infrastructure, but selectively abandons that neutrality under geopolitical pressure. Control by G10 and EU central banks means de facto western control over global financial messaging. The direct consequence is the fragmentation risk: the existence of CIPS (China) and SPFS (Russia) as alternative systems shows that an infrastructure whose neutrality is perceived as selective triggers sovereignty responses and loses its universal validity.

3.1.5 GPS

The Global Positioning System consists of 31 satellites in six orbital planes orbiting Earth at an altitude of 20,200 kilometers and permanently transmitting position and timing signals.[50] GPS provides three functions (positioning, navigation, and timing) and was launched in 1978 by the US Department of Defense as a military system. Civil use began in 1983 as a response to the shootdown of Korean Air Lines Flight 007 and was extended to full civil accuracy in 2000 by deactivating intentional signal degradation (Selective Availability) under President Clinton.[51][52] The entire civil GPS infrastructure — on which financial transactions, mobile networks, power grids, and precision agriculture now depend as a time reference — is an unintended byproduct of a military investment.

A fundamental design difference shapes the infrastructural properties of GPS: satellites transmit signals but receive none from users. There is no interaction, no registration, no authentication. The system scales perfectly, because capacity is independent of the number of users — whether one or eight billion receivers.

[49] EU-Verordnung (EU) 2022/328 zum Ausschluss russischer Banken von SWIFT, 2022.
[50] GPS.gov: Space Segment, U.S. Government. https://www.gps.gov/systems/gps/space/
[51] Presidential Statement on the Use of GPS for Civilian Aviation Safety, September 1983, in response to the shootdown of KAL 007.
[52] Presidential Decision Directive (Clinton), May 2000: Deactivation of Selective Availability.

At the same time, this architecture means that civil GPS signals are unauthenticated: the receiver cannot cryptographically prove that the received signal is genuine, and spoofing is possible and documented. Only the military M-code offers cryptographic authentication, which creates a two-class system on the same physical infrastructure: civil receivers with 5 to 10 meters accuracy and no spoofing protection, military receivers with approximately 1 meter accuracy and encryption.[53]

GPS exhibits a property that produces the maximum asymmetry in the sample: the read-only architecture under unilateral control. Anyone can receive; no one can transmit; one party controls. GPS is operated entirely by the US Department of Defense and run by the US Space Force. Free access is a political decision that could be revised at any time, and the US government explicitly reserves the right to locally jam GPS in crisis zones. The geopolitical consequence of this asymmetry is duplication: the European Union developed Galileo as an explicit sovereignty measure, Russia operates GLONASS, China BeiDou.[54] The RAND Corporation identifies GPS, owing to its cross-sector dependency, as an "attractive target for adversaries."[55] GPS thereby exhibits a property directly relevant to the research question: an infrastructure used by the entire world but subject to the sovereign control of a single state provokes duplication that fragments the universal character of the infrastructure.

3.1.6 Cloud Infrastructure

Cloud infrastructure provides compute, storage, and networking as a service. Three providers control 63 percent of the global market, which exceeds 400 billion US dollars annually: AWS with 30 percent, Azure with 20 percent, and Google Cloud with 13 percent.[56] Financial platforms, healthcare systems, government services, AI applications, and communications services run on this infrastructure. Cloud requires massive capital investments in global data centers, submarine cables, and proprietary hardware, creating barriers to market entry that are barely surmountable for new competitors. Customers who deeply integrate their applications into provider-specific services are bound by vendor lock-in that arises from ecosystem depth, because the technical integration extends far beyond formal contractual commitment.

[53] RAND Corporation: GPS Vulnerabilities and Alternative PNT Capabilities, 2020.
[54] European Space Agency (ESA): Why Europe Needs Galileo. https://www.esa.int/Applications/Satellite_navigation/Galileo/Why_Galileo
[55] RAND Corporation, op. cit.
[56] Synergy Research Group: Cloud Infrastructure Market Share, Q3 2025.

The infrastructural properties of the cloud are shaped by a tension that is unique in the sample. The economies of scale of the large providers enable lower prices, broader services, and faster innovation, but simultaneously create single points of failure. In the US-East-1 incident in December 2021, a network error caused cascading outages at DynamoDB, S3, and other AWS services, demonstrating that a single failure at one provider in one region can have global consequences.[57] Cloud providers are private-sector companies without infrastructure regulation. There is no Bundesnetzagentur for cloud, no statutory dedication, no common use claim. AWS, Azure, and GCP can reject customers, terminate services, and change prices. Independent verification of infrastructure integrity is not possible for the user: whether the underlying infrastructure is functioning correctly is communicated through SLAs and provider-owned status dashboards, not through independent verification mechanisms.

The infrastructure-specific observation of the cloud is the discrepancy between de facto infrastructural significance and the absence of public-law status. AWS, Azure, and GCP are de facto infrastructure on which a substantial portion of the digital economy runs, but they have no infrastructure status, no public interest obligation, and their terms of service can change at any time. This discrepancy has triggered regulatory responses: with the Data Act and the Digital Markets Act, the EU increasingly acknowledges the systemic significance of cloud providers and regulates data sovereignty through portability requirements.[58] The market concentration mentioned lies in the hands of three US companies. Europe is attempting to create a sovereign alternative with initiatives such as Gaia-X.[59] The parallel to the GPS situation (US control provokes a European sovereignty response) is direct. The historical comparison is close at hand: the situation of cloud resembles the state of the power grid before unbundling, when the transport layer was still in the hands of vertically integrated companies and the regulatory separation of grid and use had not yet been accomplished.

[57] AWS: Summary of the AWS Service Event in the Northern Virginia (US-EAST-1) Region, December 2021. https://aws.amazon.com/message/12721/

[58] Europäische Kommission: Data Act (Verordnung (EU) 2023/2854), 2023; Digital Markets Act (Verordnung (EU) 2022/1925), 2022.

[59] Gaia-X European Association for Data and Cloud. https://gaia-x.eu

3.1.7 Legal System

The legal system encompasses the totality of laws, courts, registers, and enforcement mechanisms that enable binding agreements between parties. It is the only infrastructure in the sample that is institutional rather than technical in nature. Its core function solves the fundamental problem of human cooperation: how can strangers enter into binding agreements? Avner Greif has shown that before the development of institutional law, transactions functioned only within clan, family, or ethnic networks, because reputation-based enforcement does not scale beyond small groups.[60] Only state-enforced contract law makes cooperation between millions of strangers possible. The effectiveness of this infrastructure shows in how rarely courts are actually called upon: contracts function "in the shadow of the law," as Mnookin and Kornhauser have formulated it, because the mere existence of the enforcement possibility suffices to stabilize cooperative behavior.[61]

The legal system manages the state of the real world through registers: who owns which parcel of land, which company exists with which liability, which agreement applies. The register function is the institutional equivalent of what technical systems manage as State, as Hodgson has worked out in his institutional-economics analysis of capitalism.[62] Verification is institutional: court proceedings are in principle public, judgments are published and contestable through appeal and revision. Access to the legal system is formally universal, but in practice unequally distributed. The so-called justice gap — the discrepancy between the formal right of access and actual accessibility — is documented in all countries.[63] Globally, approximately 200 legal systems exist with different traditions that are not interoperable. There is no universal protocol that automatically mediates between jurisdictions, and cross-border contracts must clarify which law applies.

The legal system differs from all other infrastructures through institutional enforcement by means of the state's monopoly on the use of force. No other infrastructure in the sample possesses physical enforcement power: the power grid can cut off electricity, the Internet can block packets, SWIFT can exclude

[60] Greif, Avner (1994): Cultural Beliefs and the Organization of Society. *Journal of Political Economy*, 102(5), pp. 912–950.

[61] Mnookin, Robert H. / Kornhauser, Lewis (1979): Bargaining in the Shadow of the Law: The Case of Divorce. *Yale Law Journal*, 88(5), pp. 950–997.

[62] Hodgson, Geoffrey M. (2016): *Conceptualizing Capitalism: Institutions, Evolution, Future*. University of Chicago Press.

[63] OECD / Open Society Foundations (2019): *Legal Needs Surveys and Access to Justice*. OECD Publishing, Paris. DOI: 10.1787/g2g9a36c-en.

participants, but none of these systems can actively intervene in reality. The legal system can, and this property makes it the meta-infrastructure within which all others operate: the power grid needs concession contracts, the Internet needs property rights to cables, SWIFT needs banking regulation, cloud providers need the enforceability of their terms of service. The question of what enforcement means in a system without a monopoly on the use of force is one of the open questions that Ethereum's infrastructure claim raises, and one that the evaluation framework takes up through the Coordination Function (I.3). The legal system does not "fail" like a power grid. It erodes through corruption, political interference, underfunding, and loss of trust, as numerous historical and contemporary examples attest.[64] Economic development correlates directly with the quality of rule-of-law institutions, which empirically underscores the infrastructural significance of the system.

3.2 The Twelve Criteria in the Mirror of the Reference Infrastructures

The twelve criteria are organized into three dimensions that simultaneously structure the examination. Structural Foundation (I) asks whether a system brings the load-bearing prerequisites of an infrastructure. Qualitative Load-Bearing Capacity (II) asks whether it delivers them in a neutral, open, verifiable, and accessible manner. Resilience and Sovereignty (III) asks whether it holds up under pressure and over time. Within each dimension a digit locates the individual criterion, such that Neutrality and Censorship Resistance, as the first criterion of dimension II, carries the identifier II.1. The following examination holds each criterion against the seven systems and asks a dual question: whether the required property finds purchase in established infrastructure, and with what variance it appears. Precisely this range demonstrates that a criterion draws a real distinction and does not elevate any single system form to the norm.

Functional Irreplaceability (I.1) finds purchase in all seven systems, with remarkably different justification. Each infrastructure provides a function that no alternative system can comparably perform, yet the basis of irreplaceability varies. It is physical in the natural monopoly of the electricity transport layer, economic in SWIFT's network effect, institutional in the comprehensive presence of the road network, and sovereign in the legal system's monopoly on the use of force. The Internet owes its irreplaceability to universal protocol adoption:

[64] Cf. Greif 1994, who shows that the reversion to clan-based coordination is a direct consequence of weak legal institutions.

TCP/IP is so established as a standard that an alternative protocol would be technically conceivable but practically impossible to introduce. GPS remains without a functional equivalent as a timing reference, because Galileo, BeiDou, and GLONASS do replicate the function, but only partially replace GPS owing to its installed base (cf. 3.1.5). Cloud infrastructure, finally, is irreplaceable at the level of the economies of scale that simultaneously produce quality and concentration of power. This range confirms a structural anchoring that takes many forms.

Security and Trust Load (I.2) and Independent Verifiability (II.3) are tested by the same question: on what the users' trust in the correct functioning of the system rests. The reference systems show why the evaluation framework carries two criteria here, where a superficial consideration would expect one. One requirement asks about the demand side — the trust load of the user — the other about the supply side — the verification possibilities of the infrastructure — and the two can diverge. The Internet shows the split most clearly: with TLS it offers cryptographic verification — the supply side would be satisfied — yet the chain of trust depends on centralized certificate authorities that the end user must trust. For SWIFT and cloud, the verification possibility is absent entirely, so that both requirements together yield a negative result. The power grid delivers the third constellation, in which verification is possible locally at the calibrated meter but remains reserved to the operator at the system level. That the supply and demand sides diverge for the Internet and coincide for SWIFT is the actual finding of this test: it demonstrates that the separation of the two criteria is not a conceptual duplication but a distinction that real systems enforce. The supply side of Independent Verifiability is examined separately in the qualitative dimension, because it depends on its own prerequisite: the transparency of operational parameters.

Coordination Function (I.3) finds purchase in all seven systems, because every infrastructure provides a coordination service that without it would be impossible, or only possible at considerably greater effort. Here too the spectrum is wide. It spans from the physically enforced real-time coordination of the power grid, where generation and consumption must remain in equilibrium second by second, through the decentralized self-organization of the road network, to the institutional enforcement of the legal system, which makes binding agreements between strangers possible. The Internet coordinates packet routing between 70,000 autonomous systems via BGP. SWIFT mediates financial messages between 11,000 institutions in 212 countries and territories via standardized formats. GPS delivers timing signals as a time reference for financial transactions, mobile communications, and power grids. Cloud infrastructure controls the allocation of compute, storage, and network resources across global data cen-

ters. The type of coordination depends on the infrastructure type. The requirement itself is universal: without coordination performance, there is no infrastructure claim.

Minimal Load-Bearing Guarantees (I.4), which an infrastructure must maintain even under stress, are tested by the failure behavior of the seven systems. The requirement finds purchase in six of them; the legal system constitutes the exception that sharpens the focus of the criterion. The examination reveals three distinct failure patterns. The first is cascade failure, in which a local error propagates through the system, as in the 2003 blackout in the power grid (cf. 3.1.1). The second is gradual degradation, in which the system loses performance but remains functional. The Internet loses bandwidth and gains latency with cable breaks while BGP automatically reroutes. The road network absorbs local outages through alternative routes, but with traffic jams and overloading knows a threshold beyond which the degradation endangers the infrastructure claim, because the coordination service collapses. The third is redundant stability with local vulnerability, in the case of GPS: 31 satellites compensate for individual failures, but the extreme signal weakness from 20,200 kilometers' altitude makes the system vulnerable to ground-based jammers. The legal system explains why the requirement applies in only six systems: it does not fail technically but erodes over decades through corruption, underfunding, and loss of trust — a pattern without analogue in the technical infrastructures. This exception confirms that an infrastructure that, under stress, transitions to total failure rather than controlled degradation does not fulfill the claim.

Neutrality and Censorship Resistance (II.1) finds purchase in six of the seven systems, with GPS as an instructive counterexample. The examination yields two findings, the second of which is the essential one. The first concerns the source of neutrality, which is not tied to any particular mode of origin: for the Internet it follows architecturally from the end-to-end principle — which must, however, be enforced regulatorily against the interests of network operators — while for the road network it is legally anchored in common use right. The second and more important finding is that neutrality can diverge at two levels, and SWIFT demonstrates this most clearly (cf. 3.1.4): at the transaction level it is neutral, because every message is technically treated equally, but under geopolitical pressure it abandons this neutrality at the system level as soon as states exclude individual participants. GPS provides the counterexample: in reception it remains neutral, because everyone receives the same signal, yet it stands under unilateral state control that can selectively throttle it. That even the architectural neutrality of the Internet can be physically overridden was shown by the network shut-

down in Iran in September 2022. The examination must therefore not stop at the architecture but must evaluate operational controllability and both levels of neutrality simultaneously.[65]

Open Generativity (II.2) denotes the capacity of an infrastructure to sustain unforeseen innovation without the permission of a gatekeeper. Its conceptual origin lies in the end-to-end principle and the permissionless-innovation line that the theoretical framework already carries. The examination finds it most pronounced in the Internet, whose architecture reduced the innovation barrier to nearly zero and enabled the World Wide Web, e-commerce, and cloud computing without any central authority having to authorize them. The road network generates economic generativity, because open access enables economic activities that would not exist without the infrastructure. GPS, as an unintended byproduct of military investment, has produced a global ecosystem of receivers and applications, because the signals are receivable free of charge and the specifications are open. And even SWIFT's message formats have become an industry standard that extends beyond the network. That the property appears in pronounced form in only four systems is not a deficiency of the criterion but confirms that generativity is a distinguishing requirement that established infrastructure fulfills to very different degrees.

Independent Verifiability (II.3), whose complementary relationship to Security and Trust Load was already shown in the structural examination, depends on its supply side closely on the transparency of operational parameters. This appears in all seven systems, but is not a binary state. It is a spectrum whose manifestation correlates with the governance model. Regulated infrastructures are obligated to transparency: for the power grid, grid tariffs, operational metrics, and investment plans are disclosed by regulation. For the road network, condition and investment volumes are publicly accessible. Open infrastructures are architecturally transparent, because the protocol specifications of the Internet are fully documented and anyone can view the routing tables of BGP. Closed infrastructures are operationally opaque, because SWIFT communicates only aggregated transaction statistics and cloud providers publish SLAs and status dashboards but keep the underlying infrastructure proprietary. GPS is a special case: its signal specifications are public, so that any receiver manufacturer can implement them, while the control decisions — when and where the signal is throttled — remain under the authority of the US Department of Defense. This

[65] NetBlocks: Internet Shutdown in Iran during Mahsa Amini Protests, September 2022. https://netblocks.org/reports/internet-disrupted-in-iran-amid-mahsa-amini-protests-X8qVE8Ap

transparency is a prerequisite for independent verifiability. Its variance confirms that the supply side of trust applies only where the operational parameters are accessible in the first place.

Low-Threshold Inclusivity (II.4) finds purchase in all seven systems, but the criterion measures precisely the location on a spectrum from maximum inclusivity to maximum exclusivity. GPS stands at the most inclusive end, with access that is read-only, without an account, free of charge, and anonymous. The road network offers legally secured open access through common use right. The Internet is open but requires a connection contract with a provider. The power grid offers regulated access tied to technical prerequisites and contractual relationships; the legal system is formally universal but in practice income-dependent; and SWIFT stands at the most exclusive end — permissioned, compliance-bound, and without individual access. The requirement is thus not met or missed in binary fashion, but demands a location. For an assessed system, the question is whether the remaining access thresholds are compatible with the claim it makes.

Long-Term Stability (III.1) over decades finds purchase in all seven systems, and the examination shows that each infrastructure is exposed to its own long-term risk that does not impair its current function but determines its future. The risks fall into clearly different categories, which is the essential finding. The power grid carries a system-immanent risk — the transformation from central-unidirectional to decentralized-bidirectional — whose management without loss of stability is the central problem of the energy transition. GPS carries an externally determined risk, because its continued existence depends on the budget priorities of a single state for satellite renewal. SWIFT faces the fragmentation risk through CIPS and SPFS, which could erode universal connectivity. The legal system shows that load-bearing capacity over time requires institutional stability beyond technical robustness, because democratic erosion can hollow out functionality from within. The variance across these categories confirms that the criterion must ask about the specific long-term threat to the assessed system and not about a uniform standard.

Adaptive Governance (III.2) finds purchase in all seven systems, and the examination yields a finding constitutive for the criterion: no system shows the same governance model. The spectrum ranges from the state monopoly of the legal system through the regulated private monopoly of the power grid, the public service provision of the road network, the multistakeholder consensus of the Internet, the cooperative governance under state supervision at SWIFT, and the unilateral state control at GPS, to the unregulated oligopoly of the cloud. This heterogeneity is the essential finding: it shows that no particular model is required for

infrastructure, and therefore the criterion asks not about the model but about the capacity for adaptation without the adaptation process destabilizing the system. That the power grid is transforming itself through the energy transition (cf. III.1) and the Internet completed the transition of IANA functions from the US Department of Commerce to an independent multistakeholder organization demonstrates this adaptive capacity in two very different governance forms.

Sovereign Portability (III.3) finds a counterpart in five of the seven systems, with the manifestation varying considerably. The examination shows that portability must be understood as multi-layered, because a system can be free on one layer and locked in on another. The Internet marks the portable end of the spectrum, because TCP/IP implementations are vendor-independent and a change of provider requires no adaptation of the applications built on top. The legal system marks the other end, because a change of jurisdiction requires the complete renegotiation of the legal basis, since the approximately 200 legal systems are not interoperable. The layer dependency appears most clearly with GPS: on the receiver side it is maximally portable, because any receiver processes standardized signals; on the control side not at all, because a change of provider is out of the question owing to lack of alternatives. SWIFT and cloud finally show how network effects and ecosystem depth produce a lock-in that does not prohibit switching but makes it prohibitively expensive. This range confirms that the criterion must examine the absence of proprietary dependency at multiple layers simultaneously.

Hardware Agnosticism (III.4) is the one of the twelve criteria that finds no counterpart in the reference infrastructures, and this absence is itself the finding. None of the seven systems is a permissionless decentralized network whose decentralization at the logical layer could be undermined by a hardware monoculture at the physical layer. The risk that this criterion captures therefore arises only with Ethereum's class of system. Its derivation thus rests not on the reference systems but on the concept of decentralized digital infrastructure itself. The examination confirms, precisely through the absence of a counterpart, that established infrastructure does not know this specific risk, while it is constitutive for a system with Ethereum's architecture.

The examination finally reveals a property that deliberately does not become a criterion. Five of the seven systems function as meta-infrastructure whose failure cascades beyond their own sector: electricity forms the physical base layer, the Internet the digital, the road the intermodal connecting layer, GPS the timing base layer, and law the institutional base layer. This status is, however, an emergent property that a system acquires through adoption and the formation

of dependencies, and one for which it cannot be designed. It therefore does not belong in an evaluation framework that examines which constitutive properties distinguish a system as infrastructure. The observation is recorded here without any claim to bearing on the assessment.

The overarching result confirms the meta-pattern that Chapter 2 introduced as the second fundamental principle of the evaluation logic: the degree of fulfillment must match the claim. None of the seven infrastructures fulfills all requirements at maximum degree, and this finding prevents the fallacy that a system must fulfill every criterion maximally in order to qualify as infrastructure. For a system that makes the claim of a neutral, permissionless, censorship-resistant infrastructure, the requirements for neutrality, access, and censorship resistance are higher than for a system without this claim. Ten of the twelve criteria find purchase in established infrastructure, in different form and with a variance that demonstrates their discriminating power. The twelfth proves to be a claim-specific requirement that applies precisely where the reference systems offer no counterpart.

3.3 Methodological Boundaries and Positioning

The documented derivation and validation of the evaluation criteria is subject to four boundaries that must be transparently named.

The first boundary is the survivorship bias in the validation base. The seven reference infrastructures are systems that have historically established themselves as fundamental infrastructure, while failed or displaced infrastructures — such as proprietary network protocols displaced by TCP/IP, or Telex as a formerly global communication system — are not included. The consequence is specific: the validation can confirm that the criteria capture properties shared by all successful reference infrastructures, and thereby supports the reading of the criteria as necessary conditions of infrastructure suitability. Whether their fulfillment is sufficient for actual establishment as infrastructure the examination cannot answer, because the validation base contains no failed systems from which sufficient conditions could be read off.

The second boundary concerns the theoretical framing of the derivation of criteria. The criteria were deduced from the concept of fundamental digital infrastructure, and the derivation was guided by the research question, because the concept itself was developed with a view to a cryptographic, decentralized system, such that properties such as neutrality, censorship resistance, independent verifiability, and hardware independence carried greater weight than properties that do not bear on the claim. The framing is a normal component of concept-guided construction of the evaluation framework and not a deficiency, but it must be disclosed, because it raises the question of which properties were not formulated as criteria and why. Traditional infrastructure properties such as regulatory embeddedness in state oversight structures, the existence of institutional maintenance organizations, or universal service obligations characterize the established reference infrastructures without appearing uniformly across all of them. They were not formulated as criteria, and the decision follows a specifiable distinguishing criterion rather than the contours of the subject, with the exception of Hardware Agnosticism, which is separately identified under the fourth boundary. A property is constitutive if its absence would prevent the system from fulfilling its infrastructure function at all, while a property is operating-model-specific if it describes how a particular historical class of infrastructure is organized and regulated, without the function itself depending on it. By this criterion, regulatory oversight is operating-model-specific, because the Internet fulfills its coordination function through its protocol architecture without a central supervisory authority, while adaptive governance is constitutive, because a sys-

tem that cannot adapt to changed requirements loses its function over time. The application of this criterion remains a researcher's decision, which is here transparently documented and provides the reader with the basis for an independent judgment.

The third boundary concerns the scope of the validation. Seven reference systems test the criteria not statistically but exemplarily. Their informative value rests not on a case count but on their heterogeneity: a criterion that proves itself against physical, digital, hybrid, and institutional infrastructures simultaneously is more rigorously tested than one examined against only one infrastructure type, because the structural diversity of the systems increases the probability of seeing an unsustainable criterion fail against at least one of them. The validation therefore makes no claim to completeness, and further infrastructures such as the Domain Name System, the rail network, or telecommunications infrastructure could sharpen or supplement the examination without their absence calling into question the load-bearing capacity of the tested criteria.

The fourth boundary concerns the positioning of the evaluation framework as an instrument. The twelve criteria do not claim to formulate a complete theory of fundamental infrastructure. They operationalize properties analytically derived from the concept of fundamental digital infrastructure and tested against the reference systems, supplemented by subject-specific requirements such as Hardware Agnosticism, whose relevance follows from the claim of the assessed system. The evaluation framework is constructed for the specific purpose of answering the research question of this work, and its applicability to other systems or other research questions would be an independent methodological question that this work does not address. The criteria are calibrated to the claim: an infrastructure with a different claim profile from Ethereum's — for instance, a purely national, regulated infrastructure — would require different or fewer criteria.

The twelve evaluation criteria have been tested against seven reference infrastructures, their validation is documented, and their boundaries are named. Chapter 4 applies the criteria to the current state of Ethereum — to the architecture, the operational reality, and the question of whether the interplay of the two fulfills the infrastructure claim that the system makes.

Chapter 4 Current State Assessment: Ethereum in Operational State Q1 2026

The preceding chapters have laid out the theoretical framework and established the empirical basis on which the assessment of the concrete subject can be built. Chapter 2 introduced the evaluation framework with its twelve criteria in three hierarchical levels, documented the evaluation logic with its four-tier scale and the M3 cascade, and defined the scope boundaries of the investigation. Chapter 3 validated the twelve criteria, which Chapter 2 derived analytically from the concept of infrastructure, against seven structurally distinct reference infrastructures. They are now applied. The criteria, ordered into Critical Conditions, Structural Conditions, and Qualitative Criteria, form the grid through which the subject is examined. The reader thus has the instrument at hand, but not yet the subject: Ethereum as a system has so far only been named in relation to the research question, not described.

This chapter closes this gap in two steps. Sections 4.1 through 4.6 describe Ethereum's architecture, the transaction lifecycle, the security economy, the scaling model, the governance structures, and the institutional embedding of the system in the operational state of the first quarter of 2026. This description is not an enumeration of technical features, but an attempt to make the system comprehensible as a coherent whole, in which each component derives its function from its interplay with other components. Sections 4.7 through 4.10 apply the twelve criteria to this described subject, assess each criterion individually on the four-tier scale, synthesize the individual assessments into a dimension-based profile, and formulate the M3 overall assessment of the current state.

The scope note from Chapter 2 applies without change: the current state encompasses features that are implemented and operational on the Ethereum Mainnet (status IMPL) or that carry status DEPL under the taxonomy of Section 2.2. Features with status PLAN or RES are assessed in Chapter 5, which

analyzes the target state of the Ethereum Roadmap. The reference date for all network data is 27 March 2026, unless otherwise specified. At the end of this chapter, the reader will have a complete picture of the system, a reasoned assessment of each individual criterion, and a current-state profile that makes the strengths and limitations of the system equally visible.

Alongside this substantive scope, a temporal qualification applies that is equally constitutive for the reading of the following sections. All quantitative figures in this chapter reflect the state of the first quarter of 2026, measured on the stated reference date, and they are to be understood as snapshots of a living infrastructure whose market shares, validator counts, and value aggregates continuously shift. The underlying architecture and the structural findings derived from it, by contrast, possess considerably greater permanence, so that in the further course of the argument it is not the exact figure of any individual number that carries the argument, but the order of magnitude and the direction it marks. A subsequent update of these figures would thus refresh the empirical basis without affecting the logic of the assessment.

4.1 Architectural Foundations

Historical Classification and Development Path

Ethereum was launched as a public network on 30 July 2015, financed through a crowdsale in 2014, and differed from the outset from Bitcoin, the then-dominant blockchain technology, through a central design decision: it was not conceived as a payment system, but as a programmable platform on which arbitrary code could be executed.[66] The white paper that Vitalik Buterin published in 2014 described a system in which smart contracts operate as self-executing programs on a shared state database maintained by a decentralized network of equal nodes. The technical specification of this system was presented by Gavin Wood in the same year in the Yellow Paper, which formally defined the Ethereum Virtual Machine and still serves today, in its current Berlin version of 2024, as the reference document for protocol mechanics.[67]

In its first seven years, Ethereum operated under a Proof-of-Work consensus mechanism, in which miners competed through the deployment of computing

[66] Buterin, Vitalik (2014): A Next-Generation Smart Contract and Decentralized Application Platform. Ethereum Whitepaper. URL: https://ethereum.org/en/whitepaper/

[67] Wood, Gavin (2014/2024): Ethereum: A Secure Decentralised Generalised Transaction Ledger (Yellow Paper). Current version: Berlin Version, 2024.

power for the right to produce new blocks. On 15 September 2022, the network executed the transition to Proof of Stake with the so-called Merge, a consensus mechanism in which validators acquire the right to produce blocks by depositing capital. The fundamental idea behind Proof of Stake is to establish network security through economic stake rather than computing power: validators deposit capital as collateral and risk its loss upon misconduct, so that economic risk replaces energy expenditure as the security mechanism. The consequence is a system in which security grows proportionally to the capital locked, not proportionally to electricity consumption, and in which the cost of an attack becomes directly measurable in the amount of capital to be risked. The Merge was a technical event of considerable complexity, because a live system with a total value of more than one hundred billion US dollars was converted in operation without a single transaction being lost or the network pausing, and reduced the energy consumption of the network by 99.95 percent.[68] The Beacon Chain, which implements the Proof-of-Stake system, had already launched on 1 December 2020 and had operated in parallel with the existing Proof-of-Work system for over 21 months before both layers were merged. The Merge was not the first critical moment in the network's history: in June 2016, barely a year after launch, the first decentralized autonomous organization on Ethereum, known simply as The DAO, was exploited through a vulnerability in the smart contract code, with approximately 60 million US dollars in ETH withdrawn. The Ethereum community opted for a hard fork that reversed the transaction, a decision that split the network into Ethereum and Ethereum Classic and triggered a fundamental debate about the immutability of blockchain transactions. For the subject of this work, the DAO fork is relevant for two reasons: it shows that Ethereum's community was capable of action in an existential crisis, and it marks the historical moment after which the community in practice decided that social intervention in the state of the system is possible, even if it has never been repeated since.

Since the Merge, Ethereum has undergone four protocol upgrades whose cadence illustrates the system's capacity for continued development. Shapella in April 2023 enabled for the first time the withdrawal of staked ETH, which had been frozen since the launch of the Beacon Chain in December 2020, and thereby eliminated a risk that stakers had borne for over two years. Dencun in March 2024 introduced Blob transactions, a new transaction type that enables Layer-2 Rollups to anchor data on the base layer at a fraction of previous costs, and

[68] Ethereum Foundation Blog (2022): The Merge. URL: https://blog.ethereum.org/2022/09/15/the-merge. For the energy consumption comparison cf. Digiconomist: Ethereum Energy Consumption Index. Cambridge Centre for Alternative Finance (2022): Cambridge Blockchain Network Sustainability Index.

thereby opened up the architectural foundation for the Rollup-centric scaling strategy described in Section 4.4. Pectra on 7 May 2025 raised the maximum effective balance per validator from 32 to 2,048 ETH, reduced deposit finalization from twelve hours to approximately 13 minutes, and introduced with EIP-7702 the first protocol-native form of Account Abstraction. Fusaka on 3 December 2025 implemented PeerDAS, a peer-to-peer-based Data Availability Sampling that reduces bandwidth requirements on validators by approximately 85 percent and extends capacity for Blob data. Each of these upgrades was executed as a hard fork, in which all node operators had to update their software to remain in the network.

The Two-Client Architecture

The Ethereum network consists of 14,339 nodes distributed across the public internet that collectively maintain the state of the system.[69] Not every node is a validator: a validator is a node whose operator has deposited at least 32 ETH as collateral and thereby acquires the right to participate in the consensus process, propose blocks, and submit attestations. The 964,768 active validators described in detail in Section 4.3 constitute a subset of the network. The remaining nodes validate transactions and maintain the system state without directly participating in consensus. The term node designates a computer that runs the Ethereum protocol software and communicates with other nodes to propagate transactions, validate blocks, and synchronize the current system state. Each node simultaneously operates two software components: a Consensus Layer Client, which manages the Proof-of-Stake consensus, coordinates validator assignments, and communicates with other Consensus Clients via the libP2P network protocol, and an Execution Layer Client, which executes smart contracts, maintains the state of accounts and contracts, and propagates transactions via the DevP2P protocol. Both clients are connected via the Engine API, an internal interface on port 8551 authenticated by JWT tokens, through which the Consensus Client instructs the Execution Client which blocks to execute and which state changes to adopt. The Consensus Client decides which block is canonical. The Execution Client executes the transactions in this block and returns the results for verification.

The architecture provides that multiple independent client implementations exist for each of the two layers. Five Consensus Layer Clients are available: Light-

[69] Etherscan: Ethereum Node Tracker. URL: https://etherscan.io/nodetracker (accessed 27 March 2026). For the Ethereum Foundation documentation of the client architecture cf. https://ethereum.org/en/developers/docs/nodes-and-clients/

house, written in Rust and developed by Sigma Prime; Prysm in Go by Prysmatic Labs; Teku in Java by ConsenSys; Nimbus in Nim by the Status organization; and Lodestar in TypeScript by ChainSafe. Likewise five Execution Layer Clients: Geth in Go, the longest-established and most extensively tested client; Nethermind in C#; Besu in Java with a focus on enterprise requirements; Erigon in Go with optimized storage efficiency; and Reth in Rust as the most recent implementation with a modern codebase. All ten implementations are open source, written in different programming languages, and developed by independent teams. This multi-client philosophy is not an organizational preference, but a core security property of the system: a bug in one software implementation can only threaten the network if the affected client serves more than one third of the network, because above this threshold the finalization of blocks can be blocked.

The December 2025 incident at Fusaka provides empirical validation of this design principle. A bug in the Consensus Client Prysm caused network participation to fall to 75 percent, but finalization remained uninterrupted because Prysm's market share at that time stood at 22 percent.[70] Had the same bug occurred in early 2022, when Prysm still served 68 percent of the network, participation could have fallen below the critical two-thirds threshold and brought finalization to a halt. Client diversity acts as structural fault isolation, whose protective effect depends on the actual distribution of market shares. The network's history shows that this distribution is the result of active community work, driven by awareness campaigns, economic incentives, and repeated public discussion of the dangers of supermajority concentration, and is not an automatic equilibrium state.

The geographic and infrastructural distribution of nodes provides contextual data relevant for the subsequent assessment. 39 percent of nodes are located in the USA, 14.5 percent in Germany, and 14 percent in China, reflecting a concentration on the North American, European, and East Asian regions that creates jurisdictional risks for the system. Of the hosted Execution Layer nodes, 59 percent run on three cloud providers: AWS with 35.5 percent, Hetzner with 13.8 percent, and OVHcloud with 9.7 percent.[71] This figure has been declining since 2022, when cloud concentration was even higher, but remains a relevant factor for the question of how independently the network can operate from individual

[70] Cointelegraph (2025): Ethereum sees 25% validation drop post-Fusaka as Prysm bug affects network participation. December 2025. For historical client market shares cf. clientdiversity.org.

[71] Ethernodes: Ethereum Node Distribution. URL: https://ethernodes.org (accessed early 2026). The cloud percentages refer to hosted EL nodes, not the totality of all nodes, as a portion of nodes is operated on private hardware.

infrastructure providers. A coordinated failure or a regulatory measure against these three providers would affect the majority of hosted nodes, even though the network would continue to operate as long as the total number of active nodes remains above the threshold required for consensus.

The EVM as Programmable Execution Environment

What distinguishes Ethereum from a pure payment chain like Bitcoin is the Ethereum Virtual Machine. The EVM is a deterministic 256-bit stack machine: every node in the network executes the same code with the same inputs and necessarily arrives at the same result.[72] The 256-bit word size was chosen to allow Keccak-256 hashing and the secp256k1 cryptography that Ethereum uses for signatures to be processed in single operations. The stack holds a maximum of 1,024 elements, and a call depth limit prevents infinite recursion. This determinism property is constitutive for the functioning of the system, because it forms the basis for all nodes being able to agree on the correct state without having to coordinate with each other about which result of a calculation is correct. The agreement follows from the fact that identical inputs on an identical machine must produce identical outputs.

The EVM operates with a three-tier storage model that separates responsibilities by persistence and cost. The stack is a volatile data structure following the last-in-first-out principle, serving as working memory for the immediate operands of the opcodes and cleared after the end of a function call. Memory is a volatile, byte-addressable storage area serving as a temporary workspace within a transaction, whose costs grow quadratically with size—a design decision that economically penalizes excessive memory consumption. Storage, finally, is the persistent storage whose contents are anchored in the Merkle Patricia Trie and are preserved across transactions. Storage operations are the most expensive in the entire opcode set, because every written value must be stored by all nodes indefinitely. This cost hierarchy is an economic translation of technical reality: what permanently burdens the network costs more than what exists only temporarily.

Smart contracts are programs executed on the EVM. They are deployed on the network as bytecode, a machine-level representation of program code, and are thereafter immutable: the code can no longer be modified, and the rules the code implements apply for as long as the network exists. Immutability is a consequence of the data structure in which the code is stored: the code hash of a

[72] Wood, Gavin (2014/2024): Ethereum: A Secure Decentralised Generalised Transaction Ledger (Yellow Paper). Current version: Berlin Version, 2024.

contract account is part of the state trie and is protected by the consensus mechanism like any other state entry. In practice, upgradeability patterns exist in which a proxy contract forwards calls to a replaceable implementation, but these patterns operate at the application layer and do not change the deployed bytecode itself. Deployment is permissionless. The only prerequisite is Gas fees paid to the network. Any actor can deploy code without requiring permission, and any other actor can call that code as long as they pay the fees. The execution architecture is sequential: all transactions within a block are processed strictly in order, one transaction after another, without parallelization. This property is a deliberate design decision in favor of determinism, because parallel execution can create conflicts with simultaneous accesses to the same storage area that would threaten determinism. The price of this decision is limited throughput on the base layer: at a slot time of 12 seconds and the current Gas Limit of 60,000,000, Ethereum processes on Layer 1 an average of 15 to 30 transactions per second at mixed transaction complexity, with a theoretical maximum of 238 transactions per second for simple ETH transfers and a practical throughput at the lower end of this range for complex DeFi interactions. These capacity limits are the architectural reason for the Rollup-centric scaling strategy that Ethereum has pursued since 2020: the base layer prioritizes security and decentralization, while Layer-2 systems multiply throughput by executing transactions outside the base layer and anchoring the results on Layer 1.

Ethereum operates as a fully transparent system: every transaction, every account balance, and every smart contract state is visible to every network participant. There are no native privacy mechanisms at the protocol level. Neither sender nor recipient nor transaction amount is obscured. This transparency is a consequence of the verification architecture, which requires that all nodes can trace the same state, and it creates a property that is a strength for some use cases—above all auditability and traceability—and a fundamental limitation for others, particularly for institutional actors with data protection requirements and users with privacy needs.

Within a block, the EVM enables a property that has no equivalent in traditional financial systems: atomic composability. Arbitrarily many smart contract interactions can be chained within a single transaction, and the entire chain is either fully executed or fully reversed, leaving no partial states. Flash Loans illustrate this property: a user can in a single transaction borrow millions of US dollars, deploy the borrowed funds in an arbitrage operation between decentralized exchanges, and repay the loan in the same transaction, without having to post collateral and without an intermediary approving the transaction. If the repayment

at the end of the transaction fails, the entire transaction is reversed as if it had never occurred. This property is an emergent consequence of the sequential execution architecture and the atomic state changes of the EVM, and it forms the basis for the Decentralized Finance ecosystem. Decentralized Finance, abbreviated DeFi, designates the ecosystem of decentralized financial protocols that operate as smart contracts on Ethereum and enable financial services such as lending, trading, derivatives, and asset management without institutional intermediaries. In this ecosystem, protocols build on each other without needing to enter into bilateral agreements, because compatibility is established at the level of the shared execution environment. A lending protocol can use the liquidity pools of a decentralized exchange protocol as a price source, a derivatives protocol can accept positions in a lending protocol as collateral, and a yield aggregator can orchestrate all three in a single transaction. Composability creates network effects at the protocol level: each new smart contract that builds on existing protocols increases the utility of all participating protocols, because it opens new combination possibilities. The downside of this interconnection is systemic risk, because a flaw or an economic exploit in a central protocol can, through composability chains, produce cascading effects on dependent protocols, as several DeFi exploits in the years 2020 to 2023 demonstrated.

The determinism property of the EVM enables a further infrastructurally relevant practice: the formal verification and systematic auditing of smart contracts. Tools such as Certora Prover, Slither, Echidna, and Mythril allow automated inspection of smart contract code for security vulnerabilities and logical errors, and for Tier-1 DeFi protocols, an independent audit by specialized security firms before deployment is industry standard. The EVM specification itself is formalized: the Yellow Paper defines the semantics, and KEVM provides a machine-verified formalization, which creates the basis for machine verification and constitutes a relevant prerequisite for the later assessment of independent verifiability (II.3).

Account Abstraction

Ethereum has historically recognized two types of accounts. Externally Owned Accounts are controlled by a single cryptographic key that the user holds themselves: a string of 64 hexadecimal characters whose loss means the irrecoverable loss of all assets on the associated account. Contract Accounts are controlled by code running on the EVM whose behavior is determined by the program logic. This separation created considerable usability barriers in practice: the management of the seed phrase, a human-readable representation of the private key, became the central challenge for end users. Every transaction required payment of Gas in ETH, even if the user only held other tokens and no ETH. More complex access patterns such as multi-signatures, time-limited permissions, or automated transaction authorizations were only implementable via workarounds at the application layer.

EIP-7702, activated on the Mainnet with the Pectra upgrade on 7 May 2025, addresses these limitations at the protocol level by allowing standard accounts to temporarily adopt the logic of a smart contract. In operational practice, this means that a user can configure their account so that session keys are issued granting an application limited transaction rights for a defined period without revealing the master key. Social Recovery becomes implementable, in which a user can restore their access via a network of trusted contacts if the master key is lost. Gas Sponsoring becomes possible, in which a wallet provider or an application covers the transaction costs for the user, so that users can execute transactions without holding ETH themselves. The significance of these functions for the infrastructure claim of the system is direct: a system that requires its users to manage cryptographic keys without error, to hold Gas fees in a specific currency, and to risk the irreversible loss of assets with every mistake, sets access requirements that are too high for a fundamental infrastructure with a broad user base. Account Abstraction lowers these requirements to a level closer to the user experience of traditional financial services, without relinquishing the self-custody property that distinguishes Ethereum from centralized systems. The EIP is deployed and active on the Mainnet; integration into the major wallet applications such as MetaMask and Coinbase Wallet is in progress. The adoption dynamics of smart account functionality can already be read from the older ERC-4337 standard, which implements the same functionality without protocol changes at the appli-

cation layer: over 25.5 million smart accounts and 132 million UserOperations document the demand for programmable accounts.[73]

State as State Database

Everything that Ethereum as a system stores resides in a data structure called the Merkle Patricia Trie. This structure consists of four nested tries: the World State Trie, which contains all accounts with their balances, nonces, and references to the code and storage; the Account Storage Trie, which maps the persistent storage of each individual smart contract and is referenced as a sub-trie in the World State Trie; and the Transactions Trie and the Receipts Trie, which store transaction data and execution results per block. The trie structure enables cryptographic verifiability: the World State Root, a 32-byte hash that summarizes the entire state of the system, is anchored in the header of every block, and a Merkle Proof can demonstrate with logarithmic effort that a specific state entry is present or absent in the trie, without having to search the entire trie.

The State amounts to 150 to 200 GB in the compressed storage form of the clients. The total size of a full node is 1,579 GB, with the difference consisting of historical block data, transaction receipts, and chain history.[74] The hardware requirements for operating a full node currently include at least 2 TB SSD storage, 16 GB RAM, and a stable internet connection with at least 25 megabits per second—requirements that are trivial for professional operators, but raise the entry barrier for solo operators and particularly for users in regions with limited infrastructure. Every new smart contract, every new wallet, every new token transfer accumulates state data that must be stored permanently, because Ethereum has no native mechanism for the expiry or deletion of unused state entries. A smart contract deployed in 2016 that has not been called for years occupies the same storage space as a contract processing thousands of transactions per day, because the protocol does not distinguish between active and inactive state entries. The permanent growth of State is the central long-term problem of the Ethereum architecture: rising hardware requirements raise the entry barrier for new node operators, extend synchronization times for nodes wishing to join the network, and increasingly concentrate the operation of nodes among professional providers with appropriate infrastructure. The planned solutions,

[73] Dune Analytics: ERC-4337 Account Abstraction Dashboard (accessed 27 March 2026). For EIP-7702 cf. Ethereum Foundation: Pectra Upgrade Specification.

[74] Etherscan / YCharts: Ethereum Chain Data Size (accessed 27 March 2026). The figure of 1,579 GB refers to the full-node size; the state trie comprises 150 to 200 GB in compressed client storage. Chapter 5 measures the uncompressed total state (approximately 430 GiB), so both figures quantify the same resource under different measurement definitions.

in particular the migration from the Merkle-Patricia-Trie structure to Verkle Trees, which enable more efficient state proofs, and the introduction of History Expiry, which releases historical state data from nodes' mandatory retention after a defined period, carry status PLAN and are assessed in Chapter 5.

4.2 The Transaction Lifecycle

A single transaction traverses the entire system on its journey from creation to finalization. At each station of this journey, a component of the architecture reveals itself that, in combination with the others, determines the properties of the system. The following section traces this journey and explains at each station the mechanisms involved, their design logic, and the tensions they create.

Creation and Gas

The lifecycle of a transaction begins with the user, who creates a signed message describing the desired state transition: a transfer of ETH to another address, a call to a function of a smart contract, or the deployment of a new contract on the network. The signature is performed via the ECDSA algorithm on the secp256k1 curve, which Bitcoin also uses, and cryptographically proves that the sender controls the private key of the sending account without revealing that key. The transaction contains, in addition to the signature, the target address, the value to be transferred, the input data for the smart contract call, a sequential nonce that prevents double execution, and the parameters for fee calculation.

Every operation on the EVM consumes a defined quantity of Gas, the unit of measurement for computational costs in the Ethereum network. The cost structure reflects the actual burden that an operation imposes on the network: a simple addition costs 3 Gas, a Keccak-256 hash 36 Gas, loading a storage value from persistent Storage on first access 2,100 Gas, and a write operation into persistent storage filling an empty slot 22,100 Gas. Storage operations are the most expensive, because they expand the permanently growing State and thereby create a burden that all future node operators must bear. The Gas Limit per block, the upper bound of total computational work a block may consume, stands at 60,000,000 and was doubled in the course of 2025 from the previous 30,000,000 through validator signaling without a hard fork.[75] The signaling procedure, in which validators indicate a preferred Gas Limit in their blocks and the actual

[75] Etherscan: Ethereum Average Gas Limit Chart. URL: https://etherscan.io/chart/gaslimit (accessed 27 March 2026). The increase occurred through gradual validator signaling over the course of 2025 without requiring a protocol upgrade.

value gradually approaches the weighted average, illustrates a governance mechanism that enables capacity changes without formal protocol changes.

The fee model EIP-1559, active since August 2021, divides the transaction fee into two components and replaces the earlier auction model, in which users had to submit bids without being able to reliably assess the current market situation.[76] The Base Fee is determined algorithmically and adjusts dynamically to block utilization: if the preceding block consumed more than half its Gas Limit, the Base Fee rises by up to 12.5 percent. If utilization falls below that, it falls by the same maximum factor. This mechanism produces predictable fee dynamics, because a user can derive the Base Fee of the next block from the current block with high accuracy. The Base Fee is burned—that is, the corresponding quantity of ETH is permanently removed from circulation. The Priority Fee is freely chosen by the user and flows to the block's proposer as an incentive for preferential transaction inclusion. In periods of low network utilization, the Base Fee moves in the range of 1 to 2 Gwei—one billionth of an ETH. In periods of high utilization, it can rise to 20 to 50 Gwei and above.

Propagation in the Peer-to-Peer Network

The signed transaction is propagated via the peer-to-peer network to the nodes of the Ethereum network. The Execution Layer uses the DevP2P protocol for transaction propagation, a gossip-based procedure in which each node forwards received transactions to a selection of its connected peers, who in turn forward them to their peers, until the transaction has permeated the entire network. The Consensus Layer uses libP2P, a more modular network protocol that also supports topic-based publish-subscribe and is used for the propagation of blocks, attestations, and sync committee messages. The separation of network protocols between the two layers mirrors the architectural separation of the clients and allows the communication patterns of each layer to be optimized independently.

There is no single central mempool in which all pending transactions are collected and ordered; instead, each node maintains a local view of the transactions it has received from its peers. This decentralized mempool architecture means that different nodes can see different sets of transactions at any given time, depending on network topology, propagation times, and individual filtering rules of the nodes. The heterogeneity of mempool views is a direct consequence of the decen-

[76] Buterin, Vitalik / Conner, Eric / Dudley, Rick / Slipper, Matthew / Norden, Ian / Bakhta, Abdelhamid (2019): EIP-1559: Fee market change for ETH 1.0 chain. Ethereum Improvement Proposal. Activated in the London upgrade, August 2021. Cf. Ethereum Foundation Blog for the technical documentation.

tralized architecture and creates an informational asymmetry that is constitutive for the MEV phenomenon: builders who have privileged connections to more nodes or receive private order flows have a more complete view of pending transactions than average nodes and can derive an economic advantage from this. The transaction waits in the local mempool of one or more nodes for inclusion in a block, with the waiting time depending on the user's willingness to pay fees and the current utilization of the network.

Block Building and MEV

Who decides which transactions are included in a block in what order, and why is this order economically relevant? The answer to this question reveals a dimension of the system that was neither foreseen in the white paper nor in the original protocol specification, and that only emerged as an independent phenomenon with the growth of decentralized financial markets. Specialized block builders collect transactions from public mempools and private order flows sent directly to them by search agents and trading protocols, and construct blocks optimized for the extraction of Maximal Extractable Value.[77]

MEV arises from the fact that the ordering of transactions within a block influences the economic outcome. A builder can place arbitrage transactions between decentralized exchanges that exploit price differences between markets—buying cheaply on one exchange and selling at a higher price on another—and insert this transaction at the moment when the price difference is greatest. Liquidations in lending protocols, which become due when the value of deposited collateral falls below a defined threshold, also offer extraction opportunities, because the liquidator receives a premium. Sandwich attacks construct a scenario in which a large user transaction is embedded between two builder transactions: the first buys ahead of the user and drives the price upward, the user buys at the higher price, and the second transaction sells after the user at the increased price. The economic analysis of this phenomenon by Daian and colleagues in 2020, which coined the term Miner Extractable Value, showed that the order-dependence of transaction processing creates a structural incentive for the manipulation of block composition that can destabilize the consensus protocol.

[77] Daian, Philip / Goldfeder, Steven / Kell, Tyler / Li, Yunqi / Zhao, Xueyuan / Bentov, Iddo / Breidenbach, Lorenz / Juels, Ari (2020): Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In: *IEEE Symposium on Security and Privacy (S&P)*, 2020. DOI: 10.1109/SP40000.2020.00040.

The cumulative MEV extraction since the Merge is estimated at 1.5 to 2 billion US dollars, with approximately 93 percent of these values flowing as bids to the validators who proposed the block.[78] For the individual user, MEV manifests as a hidden cost component beyond Gas fees: a swap on a decentralized exchange can experience a worse execution price through sandwich attacks without the user immediately noticing. Economic research has shown that MEV is a structural property of the system that arises from the order-dependence of transaction processing and cannot be eliminated without architectural changes to the block production pipeline.

The market structure of block production as of 27 March 2026 shows considerable concentration: Titan Builder controls 51.2 percent of blocks, BuilderNet—a decentralized multi-operator network that Flashbots, Beaverbuild, and Nethermind launched in December 2024—builds 25.7 percent, and Quasar builds 16.4 percent.[79] BuilderNet represents the attempt to break through the concentration tendencies of the builder market through a collaborative architecture: multiple independent operators contribute transactions and share the profit instead of competing against each other, which is intended to reduce the bundling advantages of individual large builders. Despite this innovation, three builder addresses control 93.3 percent of Ethereum block production. The Herfindahl-Hirschman Index of this market structure stands at 3,554, a value that clearly exceeds the threshold of 2,500 above which the US Department of Justice classifies a market as highly concentrated. Block production, the most critical operational path of the system, thus exhibits a concentration that stands in marked contrast to the decentralized architecture of the consensus protocol, where 964,768 independent validators validate the blocks. The cause of this concentration lies in the economies of scale in block building: a builder with more private order flows can construct more profitable blocks, pay higher bids to proposers, and thereby expand its market position—a mechanism that tends toward natural concentration and resembles the market pattern of cloud infrastructure that Chapter 3 analyzed as a reference.

---

[78] Flashbots: MEV-Explore and MEV-Boost Dashboard (accessed 27 March 2026). The estimate of cumulative MEV extraction is based on proposer payments documented via MEV-Boost since the Merge. Cf. Daian et al. (2020), ibid., for the theoretical foundation.

[79] relayscan.io: Builder and relay market shares (accessed 27 March 2026). The OFAC compliance data are drawn from MEV Watch and relayscan.io.

Proposer-Builder Separation

The architectural response to the question of how block production can be separated from block validation is Proposer-Builder Separation. Builders send completed blocks with bids to Relays—off-chain intermediaries such as Ultra Sound Relay, Titan Relay, or BloXroute—that act as trusted intermediaries between builders and validators. The Relays verify the validity of the blocks and the correctness of the bids before forwarding them to validators, thereby preventing a builder from submitting an invalid block and causing the proposer to lose their slot. The proposer selected for the slot—a validator from the network—blindly selects the highest bid and proposes the associated block without knowing its contents. The blindness of the proposer is a design decision intended to prevent the proposer from subsequently manipulating the transaction order or extracting MEV themselves. 90 percent of blocks use MEV-Boost, the sidecar software that connects validators to this builder marketplace.

This architecture has a direct consequence for the infrastructure claim of the system: the block production pipeline, the most critical path from the moment of block assembly to presentation to the consensus protocol, depends on off-chain intermediaries that are not disciplined by protocol rules, but operate on the basis of reputation and economic self-interest. A Relay that fails or is manipulated can withhold or selectively forward blocks without the protocol offering any corrective mechanism. The cryptographic and economic security guarantees provided by the consensus protocol do not extend to the Relay layer. ePBS (EIP-7732), the planned protocol-native solution that would embed Proposer-Builder Separation into the consensus protocol and eliminate the Relay dependency, carries status PLAN and is scheduled for the Glamsterdam upgrade in the second half of 2026.

The question of censorship resistance at the block production level can be read from concrete data. 15 percent of blocks run through Relays that operate explicitly OFAC-compliant and exclude transactions from sanctioned addresses: Flashbots Boost with 3.8 percent and BloXroute Regulated with 11.4 percent.[80] The largest builder, Titan, explicitly operates without censorship. The historical development shows a marked decline: in November 2022, a few months after the Merge, 79 percent of blocks ran through OFAC-compliant Relays, a figure that caused considerable concern in the community, because it meant that a government sanctions list was de facto controlling transaction processing in a per-

[80] relayscan.io: Builder and relay market shares (accessed 27 March 2026). The OFAC compliance data are drawn from MEV Watch and relayscan.io.

missionless system. The decline to 15 percent occurred without protocol intervention, driven by market dynamics: through the growing dominance of non-censoring builders and Relays, through the economic attractiveness of processing all transactions—since censoring builders must forgo the fee revenue from excluded transactions—and through community pressure on censoring Relays. The trajectory from 79 to 15 percent within three years is an empirical datum of considerable significance for the infrastructure assessment: it shows that the system established a state of de facto censorship resistance without formal governance intervention, but it also shows that this censorship resistance rests on market dynamics rather than a protocol guarantee, and is therefore reversible if market conditions or the regulatory environment change. The maximum delay that a sanctioned transaction can currently experience is limited to a few blocks, because within this period a non-censoring builder is very likely to be next in line. FOCIL (EIP-7805), a validator-committee-based inclusion list system in which a committee of validators designates transactions that must be included in the next block, is intended to replace this emergent censorship resistance with a protocol guarantee and carries status PLAN.

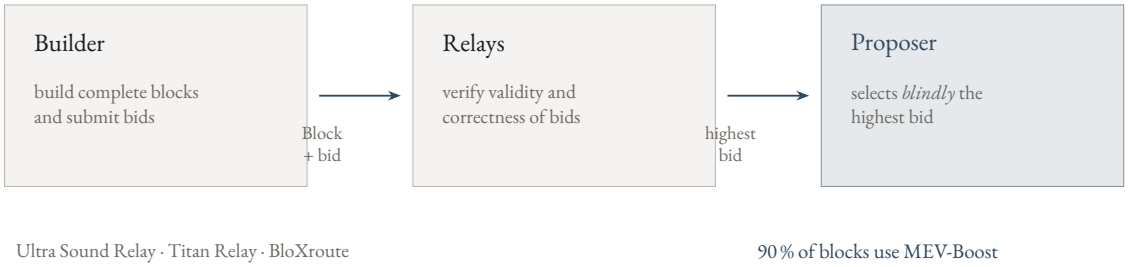

Figure 4.5 Proposer-Builder Separation. The path of a block from the Builder via the Relay to the Proposer, who blindly selects the highest bid.

Consensus and Finality

The proposer selected for a slot proposes the block to the network, and the 964,768 active validators attest via the Gasper protocol whether they recognize the block as valid.[81] Gasper is a combination of two complementary mechanisms, formally described in two separate academic papers and functioning as a unit in operational practice.[82]

Casper FFG, the Friendly Finality Gadget, is a Byzantine Fault Tolerant finalization protocol that establishes economic Finality. The Ethereum network is divided into epochs of 32 slots each, with each slot lasting 12 seconds and each epoch thus spanning 6.4 minutes. At the end of each epoch stands a checkpoint, and the validators attest over pairs of checkpoints: the source checkpoint, which is the most recently justified checkpoint, and the target checkpoint, which designates the current epoch-boundary block. When validators representing together more than two thirds of the total staked ETH attest a supermajority link between source and target, the target checkpoint is marked as justified. When the subsequent checkpoint is likewise justified and references the previous one as its source, the previous checkpoint is finalized. This two-epoch cascade explains the finalization time of 12.8 minutes: 64 slots at 12 seconds each. The two-thirds threshold derives from the theory of Byzantine fault tolerance: if two thirds of the validators vote for two conflicting checkpoints, at least one third must have voted twice and is therefore slashable, which defines the economic cost of an attack. A finalized block is cryptographically and economically irreversible, because a reversion would require more than one third of the total staked ETH to be slashed—a sum that, at current price levels, amounts to approximately 26 billion US dollars.

LMD-GHOST, Latest Message Driven Greedy Heaviest Observed SubTree, handles the complementary task of real-time block selection between finalization points. Starting from the last finalized checkpoint, the algorithm follows at each fork the subtree that has accumulated the most stake weight, counting only the most recent attestation of each validator. The restriction to the most recent message prevents an attacker from accumulating old attestations to build an alternative chain history, and the greedy traversal algorithm ensures that the network follows the path supported by the majority of stake weight.

---

[81] beaconcha.in: Ethereum Beacon Chain Explorer (accessed 27 March 2026). The figure of 964,768 reflects the post-Pectra consolidation effect of EIP-7251; prior to consolidation the validator count stood at approximately 1.07 million.

[82] Buterin, Vitalik / Griffith, Virgil (2019): Casper the Friendly Finality Gadget. arXiv:1710.09437. For the combined Gasper specification cf. Buterin, Vitalik / Hernandez, Diego / Kamphefner, Thor / Pham, Khiem / Qiao, Zhi / Ryan, Danny / Sin, Juhyeok / Wang, Ying / Zhang, Yan X (2020): Combining GHOST and Casper. arXiv:2003.03052.

The operational management of nearly one million simultaneously attesting validators requires an infrastructure that keeps the communication overhead manageable. In each epoch, the active validators are pseudo-randomly distributed across committees assigned one epoch in advance through a RANDAO-based shuffle algorithm. Each slot can contain up to 64 committees, and each committee comprises between 128 and approximately 2,000 validators depending on the total number of validators. Within a committee, designated aggregators aggregate the individual attestations of their members and publish them via the global subnet. BLS signature aggregation on the BLS12-381 curve makes it possible to combine hundreds of individual signatures into a single 96-byte aggregate signature that can be verified with constant effort. Without this cryptographic compression, processing nearly one million attestations per epoch would be computationally infeasible. Optimal inclusion of an attestation—that is, inclusion in the immediately following block with a delay of one slot—maximizes the validator's rewards. Delayed inclusion progressively reduces the reward.

Each attestation a validator submits contains three simultaneous votes: the head vote, indicating which block the validator considers the current chain head and feeding into the LMD-GHOST algorithm; the source vote, naming the last justified checkpoint and feeding into Casper FFG; and the target vote, naming the checkpoint of the current epoch and likewise relevant for Casper FFG. The interweaving of both mechanisms in a single message is the central architectural decision of the Gasper protocol: Casper FFG delivers Finality, which guarantees irreversibility, and LMD-GHOST delivers the real-time continuation, ensuring that the network can keep working between finalization points without waiting for the next checkpoint.

The empirical record of the consensus mechanism since the Merge is remarkably stable. Of 287,000 produced epochs, 13 were not finalized, a Finality rate of over 99.99 percent.[83] All 13 non-finalized epochs occurred in May 2023 in two incidents within 24 hours, triggered by bugs in the Consensus Clients Prysm and Teku in the handling of old attestations. The May 2023 incident was simultaneously the first and to date only Mainnet test of the Inactivity Leak, an automatic degradation mode that is activated when finalization is absent for more than four epochs, i.e., 25.6 minutes. In the Inactivity Leak, regular attestation rewards are suspended, and inactive validators progressively lose stake, with penalties scal-

[83] beaconcha.in: Epoch Finality Statistics. The 13 non-finalized epochs occurred on 11 and 12 May 2023. Cf. Etherscan Blog (2023): Battle-Testing Ethereum's Finality.

ing quadratically with the duration of inactivity: a validator that is completely offline loses approximately 40 percent of its stake during the leak, while a validator that remains 90 percent online loses only 0.4 percent. The mechanism aims to restore finalization by reducing the stake of inactive validators sufficiently that the remaining active validators again reach the two-thirds threshold. The self-healing in May 2023 took 96 minutes, without any external intervention being required, and the cumulative penalties amounted to approximately 28 ETH. The system thus demonstrated under stress exactly the behavior that the protocol specification prescribed.

Settlement and State Update

After finalization, the transaction is irreversibly written into the state of the network, and no actor, regardless of their economic or political power, can reverse this inscription without bearing the economic costs of a 34-percent attack. The fee distribution follows the EIP-1559 mechanism: the Base Fee is burned, and cumulatively since the introduction in August 2021, over 4.6 million ETH have been removed from circulation, which influences the monetary dynamics of the overall system analyzed in Section 4.3. The burning mechanism creates a direct connection between network activity and supply dynamics: in periods of high utilization, the burning exceeds new issuance, and ETH temporarily becomes deflationary. In periods of low utilization, issuance dominates and supply grows. The Priority Fee and any MEV payments flow to the proposer of the block, who shares a portion of the MEV payments with the builder who constructed the block, in accordance with the bid submitted in the PBS process. The State Transition Function applies the changes triggered by the transaction—such as altered account balances, updated storage values, newly created accounts, or emitted event logs—to the World State Trie and computes a new State Root, which is anchored in the header of the block and from then on serves as the cryptographic anchor for the entire system state.

System Properties of the Lifecycle

The journey of a transaction through the system reveals four architectural properties that are central to the assessment in the later sections of this chapter. The modularity of the architecture, manifested in the separation of Consensus and Execution Layer, in the multi-client philosophy, and in the layered architecture, enables fault isolation and independent further development of individual components, but creates complexity in the coordination between layers. The economically anchored security, implemented through the staking system and the Slashing mechanisms, creates attack costs that protect the system, but are coupled to the market value of ETH and thus fluctuate with the ecosystem. The off-chain dependency of block production, visible in the Relay-based PBS architecture and the builder concentration, creates a structural vulnerability in the most critical path of the system, which is not addressed protocol-side by the current architecture. The cryptographically enforced Finality, established through Casper FFG, offers an irreversibility guarantee that is stronger than the probabilistic guarantees of Proof-of-Work systems, but with a finalization time of 12.8 minutes is slower than the instant finality of centralized systems. Each of these four characteristics carries strengths and limitations in equal measure.

4.3 The Security Economy

Why do 964,768 validators stake their capital in a system that can confiscate that capital upon misconduct, and what happens if they stop doing so? The economic logic behind this commitment carries the security architecture of the system and simultaneously reveals the tensions that put this architecture under pressure in the long run.

Economic Incentives and Consolidation

Validators receive rewards for correct behavior in the consensus protocol, distributed across two sources: protocol issuance—newly created ETH that the protocol distributes to correctly attesting validators—and MEV rewards that flow from the block production pipeline to the proposer of a block. The reward structure of the protocol weights different forms of correct behavior: the correct target vote, which supports Casper FFG finalization, receives 40.6 percent of the total weight; the source vote and the head vote receive 21.9 percent each; participation in the sync committee receives 3.1 percent; and block production receives 12.5 percent. Issuance scales inversely with the square root of the total staked ETH, which means that the yield per validator falls the more capital is staked in the system. The effective staking rate stands at approximately 2.6 percent, down from approximately 13 percent when the staking share was still low.[84] Between 30 and 31 percent of the ETH total supply of 120,693,582 ETH are staked, a volume that reflects the economic anchoring of the consensus mechanism but simultaneously explains the declining yield: the more ETH is staked, the less each individual validator receives.

The Pectra upgrade in May 2025 triggered a structural shift in the validator landscape. EIP-7251 raised the maximum effective balance per validator from 32 to 2,048 ETH, and large staking providers used this opportunity to consolidate many individual 32-ETH validators into a few high-balance validators.[85] The total number of active validators subsequently fell from 1.07 million to 964,768, without any change in the quantity of staked ETH. The same economic security is now carried by fewer but larger validators, which reduces the network load because fewer attestations per epoch need to be processed, but produces a shift in the distribution of power that becomes relevant in the analysis of staking concentration. The consolidation is not a loss of confidence, but an architectural rationalization whose long-term effects on the power distribution in the network will not be fully visible until the coming years.

[84] CoinDesk (18 February 2026): Ethereum Staking Rate Reaches 30.8% of Total Supply. The effective APR figure of approximately 2.6 percent reflects the decline arising from the growing validator base under inverse square-root scaling of issuance.

[85] beaconcha.in: Ethereum Beacon Chain Explorer (accessed 27 March 2026). The figure of 964,768 reflects the post-Pectra consolidation effect of EIP-7251; prior to consolidation the validator count stood at approximately 1.07 million.

Attack Economics and Slashing

The central security threshold of the system is the 34-percent blocking minority, which derives directly from the Casper FFG consensus mechanism: an attacker controlling more than one third of the staked ETH can block finalization by withholding or contradictorily submitting their attestations, and an attacker with more than two thirds could theoretically finalize a conflicting chain.[86] With a staked volume of 30 to 31 percent of the total supply and an ETH price that stood at approximately 2,100 US dollars at the reference date of late March 2026, the static capital requirement for a blocking attack amounts to approximately 26 billion US dollars, and the attacker risks the complete loss of their stake through Slashing. To put this figure in context: the attack would require the acquisition of more than 12 million ETH on a market that already shows significant price movements at purchases in the range of hundreds of thousands of ETH, which would drive the actual costs of an attack far above the static calculation. The economic security is thus a function of stake volume, ETH price, and market liquidity, which prevents a step-free accumulation of the required position in practice.

Slashing is the economic punishment mechanism that protects the consensus protocol against two specific forms of attack. Equivocation occurs when a validator proposes two different blocks for the same slot, or submits two different attestations for the same target epoch. Surround voting occurs when a validator submits an attestation whose source-target range encompasses one of their earlier attestations, behavior that can indicate an attempt to support an alternative chain. The penalty structure scales with correlation: the initial penalty for a single slashed validator is small—approximately one four-thousandth of the effective balance—but the more validators are slashed within the same 36-day window, the higher the correlation penalty on the 18th day of the exit period, up to complete loss of stake in a coordinated attack in which one third or more of the validators are involved. This correlation scaling is a game-theoretic design decision: a single validator that double-signs due to a configuration error is mildly penalized, because the misconduct is presumably not malicious. A coordinated group that is slashed simultaneously experiences a penalty that makes the attack economically ruinous.

The empirical record of the Slashing mechanism confirms a system that operates through deterrence. Over more than five years of the Beacon Chain, since

[86] Buterin, Vitalik / Griffith, Virgil (2019): Casper the Friendly Finality Gadget. arXiv:1710.09437. For the combined Gasper specification cf. Buterin, Vitalik / Hernandez, Diego / Kamphefner, Thor / Pham, Khiem / Qiao, Zhi / Ryan, Danny / Sin, Juhyeok / Wang, Ying / Zhang, Yan X (2020): Combining GHOST and Casper. arXiv:2003.03052.

1 December 2020, 525 validators have been slashed out of over 2.2 million ever created—a rate of 0.024 percent.[87] The largest single event affected 75 validators of the provider Staked in early February 2021 and stemmed from a disabled Slashing protection database that permitted double-signing. The most recent larger event affected 39 to 40 validators in September 2025 due to infrastructure problems at SSV Network and Ankr, an incident attributable to faulty configuration of the distributed validator infrastructure. The overwhelming majority of Slashing events stem from operational errors, in particular the parallel operation of identical validator keys on multiple nodes—an error that can occur during failover or migration between server environments. Malicious attacks are not documented in the entire Slashing history. The parallel to the legal system, which Chapter 3 analyzed as an institutional reference infrastructure, is instructive: both systems deliver their coordination service in the shadow of the sanction, whose effectiveness lies precisely in that it rarely needs to be enforced. The existence of the penalty, not its application, produces the behavioral conformity that stabilizes the system.

Staking Distribution

The distribution of staked ETH across different staking categories reveals a landscape in which the access pathways to the system are structured differently and carry different implications for the decentralization of the network. Liquid Staking dominates with 31.1 percent of staked ETH: users delegate their ETH to a protocol such as Lido, Rocket Pool, or ether.fi and receive in return a tradeable token representing the staked share, which can be deployed in decentralized financial markets as collateral, liquidity provision, or a trading instrument. The Liquid Staking tokens, particularly Lido's stETH, are deeply integrated into the DeFi ecosystem: they serve as collateral in lending protocols, as liquidity in decentralized exchanges, and as the underlying asset for further derivatives. This integration creates an economic advantage for Liquid Staking over solo staking, because the user simultaneously earns staking yield and deploys their capital in DeFi applications—a phenomenon known as rehypothecation, which explains the natural tendency toward concentration of staked ETH with Liquid Staking providers. At the same time, the integration creates systemic risks: a depeg event, in which the Liquid Staking token loses its exchange rate to the underlying ETH, can trigger cascading liquidations in DeFi protocols that accept the

[87] beaconcha.in: Slashing Statistics (accessed 27 March 2026). Cf. Migalabs / CryptoSlate (11 September 2025) for the analysis of the SSV Network/Ankr slashing event.

token as collateral. Staking through centralized exchanges accounts for 24 percent, with Coinbase and Binance as the largest providers. This form of staking delegates validator control to companies directly subject to national regulation, which in the event of regulatory pressure may be forced to censor certain transactions or restrict staking services. Staking Pools, which bundle deposits below the 32-ETH threshold without issuing a tradeable token, hold 16 percent. Solo stakers—individuals who independently operate a validator and thereby represent the most decentralized form of network security—hold less than 1 percent of staked ETH, a figure that reflects the effect of the economic entry barrier: the minimum deposit of 32 ETH equates at the reference date price to approximately 67,000 US dollars, an amount that is prohibitive for the majority of potential participants. The four categories cover approximately 72 percent of staked ETH.

Lido, the largest Liquid Staking protocol, holds 22.8 to 23 percent of staked ETH, a share that has declined from 32 percent in 2023.[88] The decline results from the interplay of three factors: the falling staking yield, which has dropped from approximately 5 percent in 2023 to 2.6 percent in the first quarter of 2026, makes Lido's offering less attractive relative to alternative capital uses, particularly in a market environment where traditional interest-bearing products offer comparable returns. The growing competition from newer protocols such as ether.fi, which are gaining market share with differentiated product offerings like Native Restaking, and from staking offerings of centralized exchanges, which are low-threshold for their existing user bases, erodes Lido's dominance. Community pressure, which characterized Lido's market share as a systemic concentration risk and publicly called for stake diversification, also contributed to the decline. The Lido DAO rejected a formal self-limitation to 22 percent by governance vote in 2022—proposed by members of the Ethereum community—but as a structural countermeasure integrated Distributed Validator Technology, which distributes the operation of individual validators across multiple independent operators so that no single operator controls a significant share of Lido validators. 547,968 ETH run on DVT, a growth of 57 percent relative to the previous quarter.[89] Lido currently stands below the critical 33-percent threshold above which a single staking provider could block finalization.

The cumulative staking concentration of the three largest providers—Lido, Coinbase, and Binance—stands at 40 to 45 percent of staked ETH, while no

[88] Dune Analytics / CCN (5 March 2026): Lido Staking Market Share. The decline from 32 percent (2023) to 22.8 to 23 percent reflects relative market dynamics, not necessarily a decline in Lido's absolute staking volume.

[89] Lido DAO: Tokenholder Update (26 February 2026). DVT adoption and QoQ growth. For the governance decision on self-limitation cf. Lido DAO Snapshot Vote, June 2022.

individual provider reaches the 33-percent blocking minority threshold. The concentration is high enough that a hypothetical coordinated action by these three providers could potentially block finalization, but low enough that no single provider alone reaches the critical thresholds. The question of whether such coordination is realistic depends on factors beyond the pure staking distribution: regulatory pressure simultaneously affecting multiple providers could compel de facto coordination without the providers intending it. A concrete scenario would be a regulatory order to US-based staking providers to cease validating certain transaction types. Coinbase and other US-regulated providers would be compelled to comply, and the question of whether their cumulative stake reaches the blocking minority would shift from theoretical to practical. The DVT integration at Lido and the diversification of the builder market through BuilderNet are countermeasures that mitigate this risk, but the structural tension between regulatory reachability and decentralized claims persists.

Three Tensions of the Security Economy

The first tension concerns the inflationary dynamics of the overall system. ETH is slightly inflationary in the operational state Q1 2026, with an annual supply change of approximately 0.5 percent. The mechanics behind this dynamic can be precisely stated: the L2 migration has absorbed more transaction activity from the base layer than Blob fees return in burn volume. Layer-2 Rollups, described in Section 4.4, pay for Blob space on Layer 1, but the Blob fees are near zero because Blob capacity after the Dencun and Fusaka expansions exceeds current demand. Before the introduction of Blob transactions, Rollups paid for calldata on Layer 1, which led to substantial Base Fee contributions. After Dencun, L2 costs fell by 80 to 95 percent, and with them the contribution of L2 activity to ETH burning. In parallel, direct L1 transaction activity—and thus Base Fee burning—has declined, as users and applications increasingly migrate to Layer-2 systems that offer lower fees and higher throughput. The tension lies in the circularity: the L2 scaling that strengthens Ethereum as infrastructure by expanding the user base, lowering fees, and multiplying the capacity of the overall system simultaneously weakens the economic mechanism that drives the burning of ETH and is thereby intended to help finance L1 security in the long term. Whether this tension is structural or represents a transitional phase that resolves with growing L2 demand once Blob capacity again hits its limits is an open question.

The second tension concerns the procyclical risk of the security architecture. The attack costs in the order of 26 billion US dollars are a function of the ETH price, and the ETH price is a market variable that fluctuates with the broader

crypto market and macroeconomic conditions. In the 2022 crypto winter, the ETH price fell by over 80 percent, and the economic attack costs fell proportionally. The system survived this stress test empirically: no attack was attempted, validator numbers continued to rise despite falling yields, and the consensus protocol operated throughout in normal mode. The security guarantee is, however, not a physical constant like the load-bearing capacity of a bridge, but a market-dependent quantity that is weakest at the very moment the ecosystem is under the greatest stress and an attack would have the greatest effect. The systematic analysis of this procyclical risk, including scenarios in which attack costs could fall below critical thresholds and the mechanisms that could mitigate or exacerbate such a scenario.

The third tension concerns the long-term financing of network security and is discussed within the Ethereum ecosystem under the term security budget sustainability. If Layer-2 systems absorb the bulk of transaction activity and L1 fee revenue remains permanently at low levels, the question arises whether network security must be progressively financed through new ETH issuance, which permanently raises inflation and undermines the deflationary mechanism of EIP-1559. The alternative would be for L2 Rollups to return substantial fee revenue to the base layer over the long term through rising Blob demand—a scenario that the current Blob capacity expansion through PeerDAS makes unlikely in the medium term, because capacity is growing faster than demand. A third possibility, increasingly discussed in the community, is the reduction of issuance itself: if 30 to 31 percent of supply is already staked and security thereby stands at a level far more than sufficient to make even state-funded attacks economically pointless, issuance could be reduced without substantially compromising security. The issuance reform debate within the Ethereum ecosystem is active and unresolved: several proposals for adjusting the issuance curve are under discussion, including models that couple issuance to a target staking rate and progressively lower issuance when the staking share exceeds a defined threshold, as well as models introducing an absolute cap on issuance. None of these proposals has reached the status of a formal EIP, and the debate is politically charged, because any issuance change affects the interests of staking providers, ETH holders, and node operators differently. The question of whether Ethereum's security economy operates in an equilibrium sustainable over decades, or whether structural adjustments are required that recalibrate the relationship between issuance, fee revenue, and security budget, is one of the open questions that the current-state assessment in the evaluation section of this chapter will address.

4.4 The Scaling Architecture

The security economy has shown how the system is financed and protected. The next architectural question concerns capacity: whether the base layer can serve the transaction demand of a system that claims to function as fundamental infrastructure for global coordination.

The Gas Limit of 60,000,000 per block described in Section 4.2 limits the transaction capacity of the base layer to 15 to 30 transactions per second for complex operations. The sequential execution architecture of the EVM, which Section 4.1 explained as a deliberate design decision in favor of determinism, rules out parallelization at the base layer in the current state. To contextualize these figures: Visa processes over 65,000 transactions per second at peak times, and even the average load of global payment traffic exceeds Ethereum's L1 capacity by several orders of magnitude. The limitation is not a temporary bottleneck that could be resolved by hardware improvements, but a consequence of the fundamental design that requires every node to trace every transaction in order to independently verify the system state. The next planned capacity jump to 200,000,000 Gas with parallel transaction processing (EIP-7928) is scheduled for the Glamsterdam upgrade in the second half of 2026, but carries status PLAN and is assessed in Chapter 5. The architectural question arising from this limitation is: if the base layer can process only a few dozen transactions per second, how is the system to serve millions of users?

The Rollup-Centric Strategy

The answer that the Ethereum ecosystem has pursued since October 2020 is a division of labor between layers. Vitalik Buterin's blog post "A rollup-centric ethereum roadmap" formulated the strategic course: the base layer concentrates on two functions, settlement and data availability, while independent Layer-2 systems take over transaction execution.[90] This decision was the result of a multi-year research process in which alternative scaling approaches such as sharding of the execution layer and Plasma chains were examined and rejected, because they would either have destroyed the composability of the base layer or created unsolvable problems with data availability. The Rollup-centric strategy unites the insights of both research streams: sharding solved in 2019 the problem of Data Availability Sampling, and Rollups needed precisely the high data bandwidth at the base layer that a sharding design could provide.

[90] Buterin, Vitalik (2020): A rollup-centric ethereum roadmap. Blogpost, October 2020. URL: https://vitalik.eth.limo/general/2020/10/08/rollup.html

The division of labor can be illustrated by an analogy to the power grid, which Chapter 3 analyzed as a reference infrastructure: the high-voltage grid transports energy reliably over long distances and provides the voltage on which the entire system builds, but it does not need to reach every socket directly. Distribution grids handle the last mile to the end user. Ethereum as the base layer provides, in this analogy, the voltage—that is, the security guarantee and data availability on which Layer-2 systems build—while Layer-2 systems process the transactions with which end users interact. The analogy has limits: distribution grids in the power system are regulated, interoperable, and mutually coordinated, properties that Ethereum's L2 ecosystem possesses only to a limited extent in its current state.

The shift fundamentally changes Ethereum's operational role. From a platform on which users directly execute transactions, the system becomes an infrastructure layer providing security and data availability for systems built on top of it. This transformation is not without precedent in the history of digital infrastructure: the Internet underwent a similar layer separation when the transport protocols (TCP/IP) and the application layer (HTTP, SMTP, DNS) functionally differentiated, with the transport layer becoming invisible to the end user. In practice, this means that a user executing a transaction on Arbitrum or Base interacts with the Layer-2 system, and the base layer appears only as the anchor that guarantees the correctness of the Layer-2 state. Whether this shift in role strengthens the infrastructure claim—by positioning Ethereum as a settlement layer carrying numerous building systems—or weakens it—by delegating the direct user experience to intermediaries—is an assessment question that Section 4.8 addresses.

Data Availability as a Prerequisite

Before the mechanics of Rollups can be described, a concept must be clarified that carries the entire scaling architecture: data availability. The problem is precisely formulable: if a Layer-2 system executes transactions outside the base layer, it must be ensured that the transaction data are accessible to independent auditors, so that they can verify the state of the Rollup or demonstrate fraud. Without guaranteed data availability, a Rollup operator could suppress transactions or falsify states without this being demonstrable, because the data that would prove the fraud would not be retrievable. The base layer assumes this guarantee: it temporarily stores the transaction data in the form of Blobs, so that any auditor can access them within a defined time window. The data availability guarantee is thus the architectural precondition for Layer-2 systems to be able to inherit the security properties of the base layer, and it marks the dividing line between genuine Rollups, which use Ethereum Blobs, and alternative constructions such as Validiums, which store their data outside the base layer and offer weaker guarantees.

Blobs, introduced through EIP-4844 in the Dencun upgrade of March 2024, are 128 KB temporary data packets held on Beacon nodes for approximately 18 days and then automatically deleted.[91] The temporary storage is a deliberate design decision: the data must be available long enough for a Fraud Proof to be submitted within the challenge period, but they do not need to be stored permanently, because that would accelerate State growth at the base layer. The Blobs operate in their own fee market, separate from execution Gas, so that Layer-2 data do not compete with regular EVM transactions for block space. The Blob capacity stands after the Blob-parameter-only forks of the Fusaka upgrade at a target of 14 and a maximum of 21 Blobs per block, and PeerDAS enables further expansion to significantly higher Blob counts through peer-based Data Availability Sampling. The cost reduction since the introduction of Blobs was substantial: transaction costs on leading Layer-2 systems fell after Dencun by 80 to 95 percent, causing transaction volume on systems such as Base to rise by over 200 percent. In operational terms, this means that a simple token transaction on Arbitrum, Base, or Optimism costs a median of less than 0.01 US dollars—a level below the fees of a credit card transaction or bank transfer, economically opening access to the system for users in low-income regions. After Pectra, Blob-related costs fell by a further 51 percent. In the current state, Blob capacity exceeds

[91] EIP-4844: Shard Blob Transactions. Ethereum Improvement Proposal, activated in the Dencun upgrade, March 2024. For cost reduction data cf. L2BEAT: Transaction Costs Dashboard (accessed 27 March 2026).

demand, and Blob fees move near the minimum—a fact identified in Section 4.3 as a driver of inflationary dynamics.

Alternative data availability solutions such as Celestia, EigenDA, and Avail offer higher throughput at lower costs, but operate with their own validator sets and security models that provide weaker guarantees than the Ethereum base layer. Celestia uses its own Proof-of-Stake system with a Nakamoto coefficient of 6, meaning six actors would suffice to compromise the system. Ethereum's 964,768 validators are by contrast a raw count that overstates actual decentralization, because a significant proportion is operated by a few large providers. A Nakamoto coefficient and a validator count measure different quantities and cannot be directly compared. EigenDA is based on a restaking model in which Ethereum validators provide additional security services, but operates with a Data Availability Committee that does not make the data publicly verifiable. Layer-2 systems using these alternative DA solutions lose the security guarantees of Ethereum and are classified by L2BEAT as Validiums or Optimiums—a distinction signaling that full security inheritance is only present when Ethereum Blobs are used as the data availability layer. All five leading Layer-2 systems by total value secured—Arbitrum One, Base, OP Mainnet, Starknet, and zkSync Era—use Ethereum Blobs as their primary data availability layer, reflecting the preference for the stronger security guarantees of the base layer even when cheaper alternatives are available.

How a Rollup Works

A Rollup delegates transaction execution to an independent system but anchors the results and transaction data on the base layer in order to inherit its security guarantees. The operational flow follows a recurring pattern: a sequencer—a specialized operator steering the Layer-2 system—collects transactions from L2 users out of the local mempool of the L2, executes them on the L2 execution environment, bundles the results into batches, and posts compressed data as Blobs on Layer 1. The sequencer determines the order of transactions on the L2, publishes the resulting state change as a State Root on the base layer, and is thus the actor controlling the interface between the two layers. The compression is substantial: hundreds to thousands of L2 transactions are packed into a single Blob, and the L2 users pay only a fraction of L1 costs, because the costs of Blob anchoring are distributed across all transactions in the batch.

Transaction confirmation on the L2 occurs in two stages with different security properties. The first stage is the sequencer confirmation, occurring within seconds or sub-seconds, signaling to the user that their transaction has been included in the L2 chain. This confirmation rests on trust in the sequencer and carries no Finality guarantee: the sequencer could theoretically reverse the transaction as long as it has not yet been anchored on Layer 1. The second stage is L1 finalization, which occurs when the batch has been posted on Layer 1 and finalized through the consensus mechanism, plus the verification period—which for Optimistic Rollups is seven days and for ZK-Rollups is the proof generation time. Only after this second stage does the transaction inherit the full security guarantees of the base layer.

The verification of State Roots follows one of two approaches, whose difference lies in the burden of proof. Optimistic Rollups assume that the submitted states are correct and grant a seven-day challenge period in which any participant can submit a Fraud Proof refuting an erroneous state claim. The security of this model rests on the 1-of-N honest verifier assumption: it suffices that a single honest auditor demonstrates the fraud within the deadline, and the system remains secure as long as this condition is met. The price is the seven-day delay for native withdrawals from L2 to L1—too long for many use cases and compelling the use of liquidity bridges that introduce their own trust assumptions. ZK-Rollups pursue the opposite approach: they mathematically prove the correctness of the submitted state before its acceptance through a validity proof, which a verifier contract on Layer 1 checks in constant time regardless of how many transactions the batch contains. Withdrawals can be finalized immediately after verification

of the proof, but generating the proofs is computationally intensive and creates costs passed on to L2 users. Arbitrum One, Base, and OP Mainnet are Optimistic Rollups. Starknet and zkSync Era are ZK-based, with the boundaries increasingly blurring as projects such as BOB combine optimistic execution with ZK proofs on demand.

The L2 Ecosystem in Operational State

All leading Layer-2 systems operate with centralized sequencers—a fact that contradicts the decentralization claim of the overall system and represents the most important limitation of the current scaling model. Sequencer control lies with the developer teams: Offchain Labs operates the Arbitrum sequencer, Coinbase the Base sequencer, the Optimism Foundation the OP Mainnet sequencer, StarkWare the three Starknet sequencers. This centralization creates a triple risk profile: the sequencer constitutes a single point of failure for the liveness of the L2, because its failure interrupts the entire operation of the system. It can arbitrarily exclude or delay transactions, because it has sole decision-making power over transaction inclusion. Its monopoly on transaction ordering enables MEV extraction at the L2 level. Documented sequencer failures—including 78 minutes at Arbitrum in December 2023, 33 minutes at Base in August 2025, and over five hours at Starknet in September 2025—illustrate the liveness risk in practice. The Linea incident of June 2024, in which the development team deliberately stopped the sequencer and censored attacker addresses following an exploit, illustrates the censorship risk and shows that the capacity for censorship is not merely theoretical, but is actually exercised when circumstances require it—even if justified as preventing harm to the ecosystem. The MEV dynamics on Layer-2 systems differ fundamentally from those on Layer 1: centralized sequencers operate with private mempools, which prevents classic mempool-based sandwich attacks, but sequencer operators hold a monopoly on transaction ordering that they can exploit economically. Base generated approximately 70 percent of all Rollup profits through sequencer fees in August 2025 alone—a volume that illustrates the economic incentives working against decentralization of the sequencer layer.

The L2BEAT Stage framework measures the degree of decentralization of Rollups on a three-tier scale.[92] Stage 0 requires self-identification as a Rollup,

[92] L2BEAT: Stages Framework and Risk Analysis. URL: https://l2beat.com/scaling/summary (accessed 27 March 2026). The stage classification follows the framework introduced in June 2023, which measures decentralization and trust minimization. For Vitalik Buterin's analysis of stage transitions cf. Buterin, Vitalik (2025): Stages as a framework for evaluating rollup maturity. Blogpost.

data availability on L1, and open-source software. Stage 1 requires a functioning proof system, a Security Council with at least eight participants of whom at least half are external, and a 75-percent threshold for Council interventions. Stage 2 requires a permissionless proof system, at least 30-day exit windows for all upgrades, and a restriction of the Security Council to on-chain-verifiable bugs. Arbitrum One (Stage 1, total value secured 17.5 billion US dollars) has the BoLD proof system and a publicly staffed Security Council. Base (Stage 1, 11 billion US dollars), operated by Coinbase, and OP Mainnet (Stage 1) share the Optimism proof infrastructure. Starknet has reached Stage 1 with the Stwo proof system. zkSync Era remains at Stage 0 because its Boojum proof system was paused. Stage 2, full decentralization, exists only for immutable projects with minimal transaction volume such as Aztec v1 with a TVL of 2.95 million US dollars. The distance between the current Stage-1 status of the leading Rollups and the Stage-2 target is considerable: it requires 30-day exit windows instead of the current zero to 17 days, restriction of the Security Council to on-chain-verifiable bugs, and a permissionless proof system that operates without privileged actors. The combined transaction capacity of all active L2 systems exceeds 10,000 transactions per second, thereby extending the L1 capacity by more than two orders of magnitude, with approximately 65 percent of all new smart contracts now deployed directly on L2 systems rather than on the base layer.

Forced Inclusion—a mechanism allowing users to send transactions directly to Layer 1 bypassing the sequencer, thereby ensuring their transaction is processed even in the event of sequencer censorship—exists at Arbitrum with a maximum delay of 24 hours and at Optimism with 12 hours, but is absent at Starknet. More recent research has identified the "hand-off problem," in which the sequencer can see Forced Inclusion transactions in the queue and cause them to fail through state manipulation, a limitation that reduces the censorship protection of the mechanism in practice. Exit windows—the time frames available to users to withdraw their funds after an upgrade proposal—vary drastically: 17 days at Arbitrum, but zero days at Base and OP Mainnet, meaning that an upgrade without prior notice can change the rules under which users operate. An exit window of zero days means that the Security Council can activate an upgrade that changes the rules of the system without granting users a deadline to withdraw their funds.

The decentralization of the sequencer layer is being pursued through two approaches that differ in their architectural ambition. Shared Sequencing—where multiple Rollups share a decentralized sequencer network—promises improved censorship resistance and the possibility of establishing cross-Rollup

atomicity, but suffered a setback with the discontinuation of the Astria project in 2025, which revealed the technical and economic fragility of this approach. Based Rollups, a concept introduced in 2023, pursue a more radical approach: the Ethereum validators themselves take over transaction ordering, which maximizes decentralization because the entire validator base of the base layer acts as the sequencer. Taiko is the only productive Based Rollup, and the approach creates a trade-off between decentralization and user experience, because the block times of the Based Rollup are tied to the 12-second slot time of the base layer, while centralized sequencers can deliver sub-second confirmations.

Broken Cross-L2 Composability

The atomic composability that Section 4.1 described as an emergent property of the EVM does not exist at the Layer-2 level. Over 50 active Layer-2 systems operate with isolated State, and a transaction on Arbitrum cannot interact atomically with a transaction on Base, even though both systems settle on the same base layer. The isolation arises necessarily from the architecture: each L2 executes its transactions in its own execution environment, maintains its own State, and posts its own batches on Layer 1, without the base layer enforcing coordination between the L2 States. Cross-L2 transfers require bridges—specialized protocols that lock assets on one chain and release them on another—which introduce their own trust assumptions beyond the security guarantees of the base layer. The security record of these bridges is problematic: cumulative damage from bridge hacks exceeds 2.8 billion US dollars. The largest crypto theft in history—the Bybit hack of February 2025 with 1.5 billion US dollars—affected the wallet infrastructure of an exchange, however, not a bridge.[93] The attack vectors are diverse: compromised multisig wallets, faulty smart contract logic, manipulated oracles, and social engineering attacks on bridge infrastructure operators. The alternative to bridges is the native withdrawal mechanisms of the Rollups, which at Optimistic Rollups require a seven-day waiting time and are thus too slow for most use cases.

The practical consequences of fragmentation affect users and protocol developers equally. A user wishing to move assets from Arbitrum to Base must either use a bridge—accepting fees, slippage, and a security risk—or choose the native withdrawal path via Layer 1 and wait a week. The liquidity that was concentrated

[93] For cumulative bridge hack damages cf. DeFiLlama: Hacks Dashboard and Chainalysis: Crypto Crime Report 2025. The Bybit hack of February 2025, which affected hot wallet infrastructure, is classified by the industry as the largest crypto theft in history.

on the base layer and enabled the atomic composability of the DeFi ecosystem is distributed across dozens of isolated systems, reducing market efficiency and increasing spreads. Protocol developers must decide on which L2 to deploy, and if they wish to be present on multiple L2s, they must maintain separate deployments that do not interoperate with each other.

The Superchain initiative, which brings together 34 OP-Stack-based chains under a common interoperability standard and represents approximately 66 percent of total L2 TVL, has made progress in intra-stack fungibility with the SuperchainERC20 standard—that is, the seamless movement of tokens between chains using the same software stack. Cross-stack atomicity—the possibility of atomic transactions between an OP-Stack chain and an Arbitrum chain—does not exist and is not foreseeable. L2 fragmentation is the central architectural tension of the scaling model: the capacity of the overall system grows with each new Layer-2 system, but the coherence of the ecosystem decreases with each isolated State, liquidity is distributed across an ever-growing number of systems, and the user experience is fragmented by the necessity of moving assets between systems. The question of whether Ethereum is a coherent system or a collection of loosely coupled systems sharing a common settlement layer is addressed in the assessment of the coordination function.

Stablecoin Anchoring and Substitution Risk

Ethereum dominates global stablecoin issuance with a market share of 52 to 54 percent, corresponding to a volume of approximately 166 billion US dollars.[94] USDC, USDT, and DAI are primarily issued on Ethereum, and the DeFi integration of these stablecoins—particularly as collateral in lending protocols, as base liquidity in decentralized exchanges, and as a reserve medium in treasury strategies of institutional actors—creates an anchoring depth that goes beyond mere issuance. This anchoring must, however, be weighed against a concrete substitution risk: since December 2025, USDC transaction volume on Solana has been higher than on Ethereum. Solana is an independent Layer-1 blockchain with lower fees and sub-second transaction times, not an Ethereum L2, and its growing share of stablecoin volume shows that transaction activity migrates wherever the user experience is most favorable and fastest, regardless of the anchoring depth of the issuing chain.

[94] CoinGecko / DefiLlama: Stablecoin Market Cap by Chain (accessed 27 March 2026). The figure of 52 to 54 percent refers to Ethereum's share of the total stablecoin market including L2 issuance.

The differentiation this work observes lies in the distinction between anchoring depth and transaction volume: issuance, institutional collateral use, and deep DeFi integration remain Ethereum-concentrated, while transaction volume—particularly for retail transfers and payment applications—follows the path of lowest fees. Whether this differentiation is stable in the long term or whether sustained volume migration erodes anchoring depth—because issuers relocate their primary issuance to where volume lies—is an open question relevant for the assessment of functional irreplaceability. The parallel to SWIFT's situation from Chapter 3 is instructive: SWIFT dominates international payment messaging despite technically superior alternatives such as CIPS, because the installed base and institutional integration create switching barriers that go beyond technical functionality. Ethereum possesses with DeFi integration and stablecoin issuance an analogous installed base, but the switching barrier is lower than at SWIFT, because stablecoins can exist on multiple chains in parallel and an issuer can relocate primary issuance without coordinated system change. BNB Chain holds relevant consumer market shares, but operates with a fundamentally different decentralization compromise that opens a different risk category for the infrastructure assessment, because the chain is controlled by a single commercial entity.

The GENIUS Act, signed on 18 July 2025 as the first US federal law for stablecoins, regulates stablecoin issuers with regard to reserve requirements, transparency obligations, and licensing, but does not address the protocol layer.[95] The law creates regulatory clarity at the application layer by defining what requirements a stablecoin issuer must meet to operate legally in the USA, thereby stabilizing the framework for the largest use case built on Ethereum's infrastructure. Reserve requirements ensure that every issued stablecoin dollar is backed by qualified reserves, and transparency obligations require regular reports on the composition and scope of the reserves. For the infrastructure assessment, the GENIUS Act is a relevant datum because it shows that the regulatory framework concentrates on the application layer and leaves the protocol layer untouched: the law regulates Circle and Tether as issuers, but regulates neither Ethereum as a protocol nor the smart contracts that manage USDC and USDT on the base layer. For DeFi overall—that is, for the decentralized protocols operating on the infrastructure that have no identifiable issuers—no comparable regulatory framework exists, and the question of how the decentralized financial layer will be regulated,

[95] GENIUS Act: Guiding and Establishing National Innovation for U.S. Stablecoins Act. U.S. Congress (2025), signed 18 July 2025.

whether through adaptation of existing regulatory instruments or through new frameworks adapted to the decentralized architecture, remains open.

The division of labor described—in which the base layer provides settlement and data availability while Layer-2 systems take over execution—has been supplemented since early 2026 by a strategic shift. Vitalik Buterin formulated in a blog post the objective of scaling the base layer itself to a multiple of its current capacity, through parallel transaction processing (EIP-7928), a Gas Limit of 200,000,000, and in the long run a thousandfold capacity increase. This shift, discussed in the community as the L1-first pivot, changes Ethereum's positioning: from a pure settlement layer on which end users execute transactions only indirectly via L2s, to a system that should also enable direct usage on the base layer at larger scale. The features intended to implement this shift carry status PLAN and are assessed in Chapter 5.

4.5 Governance and System Evolution

How does a system without central authority continue to develop? The preceding sections have described Ethereum's architecture as a technical system. This section describes the processes through which this system is changed, the actors that carry these changes, and the tensions arising from the interplay of decentralized claim and operational reality.

The EIP Process

Ethereum Improvement Proposals are the formal standardization mechanism for protocol changes, inspired by Python's PEP system and Bitcoin's BIPs.[96] The process follows a defined sequence: a proposal is submitted as a Draft, with any actor able to submit a proposal regardless of institutional affiliation. The Draft undergoes a review in which technical correctness, specification completeness, and compatibility with existing protocol elements are checked; reaches Last Call status, in which the community receives a final objection period; and is marked Final upon acceptance. In 2025, 230 EIPs were submitted, of which 37 were accepted—an acceptance rate of 16 percent that reflects the selectivity of the process and shows that the majority of proposals are either technically immature, redundant, or too controversial for inclusion. EIP Editors check the formal correctness and completeness of proposals. Substantive decisions about

[96] EIP-1: EIP Purpose and Guidelines. URL: https://eips.ethereum.org/EIPS/eip-1. For the 2025 EIP statistics cf. Ethereum Magicians Forum and EIPs.ethereum.org (accessed 27 March 2026).

the inclusion of an EIP in a concrete upgrade are made in the AllCoreDevs Calls—biweekly video conferences of the client developer teams that are publicly streamed, recorded, and documented. The recordings are a key transparency feature of the process: any interested party can trace which arguments were made for and against the inclusion of an EIP and who held which position.

The procedure is transparent, documented, and selective, but operates without formal authority: no body can force a decision, no voting mechanism can bind a majority, and no institutional mandate legitimizes the decision-makers beyond their technical competence and their commitment to the ecosystem. The decision-making power lies with the client teams that release the software and with the node operators who decide which version to run. An upgrade implemented by the client teams but rejected by a significant minority of node operators would lead to a chain split—a risk that acts as a disciplining element and increases the pressure for consensus ahead of any decision. Ethereum has deliberately not implemented on-chain governance—that is, token voting mechanisms that would legitimize protocol changes through votes of token holders. The decision-making principle follows the model that the Internet from Chapter 3 has shaped: Rough Consensus and Running Code, a formulation of the Internet Engineering Task Force that replaces formal voting with iterative technical consensus-finding.[97] The only governance instrument that validators exercise directly is Gas Limit signaling, through which the doubling of the limit from 30,000,000 to 60,000,000 occurred without a hard fork, as Section 4.2 described.

The Ethereum Foundation and Decentralized Financing

The Ethereum Foundation, founded in 2014 as a Swiss foundation, is the oldest and financially most significant organization in the Ethereum ecosystem. Its treasury is estimated at 850 to 950 million US dollars, of which 70,000 ETH are staked and approximately 5,800 ETH are invested in DeFi vaults.[98] The Foundation finances research, development, education, and community initiatives, but has deliberately no protocol control: it cannot mandate protocol changes, cannot force upgrades, and cannot censor transactions. The ”subtraction” philosophy advocated by the Foundation’s current leadership circle provides that the EF

[97] The formulation ”Rough Consensus and Running Code” originates with David Clark, who coined it at an IETF meeting in 1992. Cf. RFC 7282: On Consensus and Humming in the IETF. Internet Engineering Task Force, 2014.

[98] Ethereum Foundation: Report 2024. URL: https://ethereum.foundation/report-2024.pdf. For current treasury estimates cf. Arkham Intelligence: Ethereum Foundation Wallet Tracking (accessed 18 March 2026). The figure of $850 to $950 million is an estimate based on published holdings and market developments since the last report.

progressively withdraws from operational functions and restricts its role to areas the ecosystem cannot finance on its own—particularly long-term fundamental research, formal verification, and coordination of client development.

The limitations of this structure are recognizable and openly discussed within the ecosystem. The Foundation is controlled by a Board subject to no protocol-level accountability. Its decisions on fund allocation, personnel policy, and strategic priorities are voluntary transparency acts documented in the annual report, but not protocol-required disclosures that governance mechanisms could compel. The question of who sets the research agenda and which projects are funded has direct effects on the direction of protocol development, even if the Foundation formally has no decision-making authority over protocol changes. The concentration of a substantial treasury in a single organization creates a centralization risk in tension with the protocol's decentralized claim, particularly because the Foundation in practice co-determines the direction of protocol development through its financing decisions. In 2024 and 2025, the EF undertook multiple internal restructurings that triggered discussions about the governance of the Foundation itself—a sign that the tension between centralized organizational structure and decentralized ecosystem is also perceived within the institution. Protocol Guild, a collective financing instrument for core developers with over 190 members, has received more than 50 million US dollars in voluntary donations from the ecosystem since its founding and represents a decentralized counterpoint to EF financing.[99] The instrument distributes received funds time-weighted to members actively working on protocol development, thereby creating an incentive for long-term engagement that operates independently of the Ethereum Foundation. The financing through Protocol Guild is, however, neither protocol-anchored nor long-term plannable, but dependent on the continuing willingness to donate of an ecosystem acting from economic self-interest, and has not entered into a contractual obligation to continue. The overall picture of protocol financing shows a multi-tier structure: the Ethereum Foundation as the largest single funder, Protocol Guild as a collective mechanism, individual companies such as ConsenSys and Paradigm financing client teams, and grant programs of various L2 ecosystems supporting specific development work. The dependence of protocol development on voluntary financing—secured neither through protocol fees nor through institutional obligations—distinguishes Ethereum from the reference infrastructures of Chapter 3, which uniformly dis-

[99] Protocol Guild: Documentation and Membership. URL: https://protocol-guild.readthedocs.io (accessed 27 March 2026).

pose of regulated fee models, tax-financed budgets, or institutional membership contributions.

Upgrades as Risk and Evidence

The DAO fork of 2016 documented in Section 4.1 remains the only case in which the community undertook a state intervention, and it has never been repeated since. The four post-Merge upgrades—Shapella, Dencun, Pectra, and Fusaka—all proceeded without chain splits, empirically demonstrating the maturation of the governance process: the community has coordinated, implemented, and activated four complex protocol changes without the network splitting or being interrupted. The upgrade cadence of approximately one major fork per year since the Merge shows a system actively developing, and each upgrade has introduced substantial improvements, from staking withdrawals through Blob transactions to Account Abstraction.

Complexity accumulation is a growing risk: Pectra bundled 11 EIPs in a single release—the most complex single-release scope in Ethereum's history—and the Prysm bug at Fusaka described in Section 4.1 was a direct consequence of the interaction between new protocol elements and existing client implementations. Coordination costs rise with the number of teams involved and the complexity of the changes: each EIP must be correctly implemented in ten independent implementations, the interactions between EIPs within an upgrade must be consistent across all ten implementations, and the testing phase on the testnets must cover enough edge cases to avoid Mainnet errors. The Ethereum community actively debates whether the current bundling strategy—combining many EIPs in one release—should be replaced by smaller, more frequent upgrades in order to reduce the error risk per release. The question of whether the governance model, which rests on Rough Consensus and voluntary cooperation, is capable of managing the growing complexity over the long term without either slowing the upgrade cadence or increasing the error rate, is open and is addressed in the assessment of Long-Term Stability.

The governance system's capacity to respond to acute security incidents has historical precedents beyond the DAO fork. In September 2016, a targeted denial-of-service attack against the Geth client forced two emergency hard forks within a matter of weeks — Tangerine Whistle and Spurious Dragon — an exceptionally rapid response that demonstrated the governance system's ability to act in acute crises. The limit of this capacity lies in the absence of a pause mechanism: if a critical bug is deployed at the protocol level, no party can halt the network. The only response is an emergency hard fork, which requires the coordination of

all client teams and can take weeks. For a system securing assets in the three-digit billions, this is a relevant constraint that is taken into account in the assessment of Adaptive Governance.

4.6 Emergent Properties: What the Architecture Enables

Sections 4.1 through 4.5 have described Ethereum's architecture as a system of components: the EVM, the consensus protocol, the security economics, the scaling architecture, and governance. Each of these components has specific strengths and constraints identified in the sections. The properties that underpin the infrastructure claim, however, cannot be located in any single component. They emerge from the interplay of architectural elements and generate capabilities that exceed the sum of the individual parts. The analogy to infrastructure theory from Chapter 2 is direct: Star's relational perspective showed that infrastructure arises in relation to organized practices, and the following five properties describe precisely these relations — the capabilities the system generates for its users that empirically ground the infrastructure claim. Five such properties can be identified, each arising from the interaction of specific components, and whose infrastructural relevance is disclosed by the theoretical framework of Chapter 2.

The first emergent property is Permissionless Deployment. From the interplay of the EVM as an open execution environment, permissionless access to the network, and the Gas market as the sole access requirement, a system emerges in which any actor can deploy code without having to seek authorization. Frischmann's argument that open access to infrastructure is economically more efficient than restriction through private property rights — because the aggregate social benefit rises with the breadth of use — finds a technical implementation here.[100] The parallel to the end-to-end principle of the Internet, which Chapter 3 described as the architectural foundation of the network's innovative capacity, is direct: van Schewick showed that the end-to-end principle enabled the emergence of the World Wide Web, e-commerce, and cloud computing by lowering innovation barriers at the edges of the network.[101] The EVM generates an analogous dynamic for programmable financial logic and decentralized applications, and the breadth of activity built on Ethereum illustrates the generative capacity of open access. The 31,869 active developers documented by the

[100] Cf. Frischmann 2012 and the presentation in Chapter 2, Section 2.1.

[101] Van Schewick, Barbara (2010): Internet Architecture and Innovation. MIT Press. Cf. the reference in Chapter 3, Section 3.1.2.

Electric Capital Developer Report for September 2025 are working on protocols that encompass financial markets, identity systems, supply chain tracking, decentralized governance, and digital art, and each of these application domains arose from the initiative of individual actors who used the infrastructure without seeking authorization.[102] The DeFi Total Value Locked of approximately 100 billion US dollars on Ethereum L1 and L2 combined quantifies the economic scope of productive activities built on this infrastructure. The constraint lies in Gas costs, which on Layer 1 are prohibitive for many applications and effectively shift deployment to Layer 2 systems, where the infrastructure properties — in particular decentralization and composability — are more limited.

The second emergent property is atomic composability. From the interplay of the EVM with the shared state and sequential block execution, the possibility arises of chaining any number of protocol interactions in a single transaction, as Section 4.1 illustrated using Flash Loans. The infrastructural implication reaches beyond the technical mechanism: protocols can build on one another without needing to conclude bilateral agreements, because compatibility is established at the level of the shared execution environment. A developer creating a new DeFi protocol can integrate the liquidity pools of Uniswap, the lending markets of Aave, and the stablecoin logic of MakerDAO into their own application without a contractual relationship to any of these protocols, without needing to obtain authorization, and without the operators of the used protocols needing to know about the integration. DeFi as an ecosystem — whose economic scope the preceding paragraph quantified — is the result of this property: an emergent layer of productive activity that builds on the infrastructure without having been planned or anticipated by it. This is Frischmann's concept of the input character of infrastructure in its most precise form: the infrastructure creates the enabling space, and the productive activities of users create the value that justifies the infrastructure. The strength of composability lies in its permissionlessness. Its weakness lies in its limits. At the Layer 2 level, this composability is broken, as Section 4.4 documented, and fragmentation across more than 50 isolated systems undermines the very property that produced the ecosystem.

The third emergent property is independent verifiability of the system state. From the interplay of the multi-client architecture, the Merkle Patricia Trie, and the peer-to-peer network, the possibility arises of verifying the entire system state without trusting an intermediary. Technically, this verification is possible on con-

[102] Electric Capital: Developer Report, September 2025. The figure of 31,869 active developers encompasses public repository activity in the Ethereum ecosystem. DefiLlama: Ethereum DeFi TVL (accessed 27 March 2026).

sumer hardware with a 2 TB NVMe SSD and 16 to 32 GB RAM, at a total cost of 500 to 1,500 US dollars. A user running a full node can independently verify every transaction, every account balance, and every smart contract state without having to trust an external source, because their node has replicated all blocks since the genesis block and independently computed every state transition. Operationally, approximately 70 percent of users rely on RPC providers such as Infura and Alchemy, which act as centralized intermediaries between users and the network and whose failure disrupts user access without the protocol itself being affected. The Infura outage of November 2020, during which MetaMask and numerous decentralized applications were unreachable for hours even though the Ethereum protocol continued to operate, illustrated the discrepancy between protocol-level decentralization and operational centralization at the access layer. Decentralized RPC alternatives such as Pocket Network and dRPC exist, but hold only a single-digit percentage of market share, because institutional users prefer the service-level agreements and latency stability of centralized providers. The possibility of independent verification is given. Its use is opt-in, with economic and technical barriers. This distinguishes Ethereum from traditional systems in which the correctness of state is asserted by institutional authority — banks, clearinghouses, or central registries — and in which the user has no technical means to verify the assertion independently.

The fourth emergent property is cryptographic ownership. From the interplay of private keys, smart contracts, and the settlement finality described in Section 4.2, a form of ownership control emerges that is exercised through cryptographic proofs, without any institution needing to grant it. A user who controls their private key has unconditional control over their assets at the protocol level, without a bank, an exchange, or an authority being able to grant or revoke that control. The unconditionality of this control is a direct consequence of the cryptographic architecture: the private key is the only information required to sign a valid transaction, and no instance in the system can block a correctly signed transaction as long as the user pays Gas fees and the transaction complies with the protocol rules. Self-custody is the default. Delegation to an intermediary such as a centralized exchange is a deliberate choice by the user, reversible at any time. The flip side of this property is its irreconcilability: key loss is irreversible, because no instance can issue a new key or restore access, and the Account Abstraction described in Section 4.1 addresses this barrier only partially, because wallet integration is still being implemented and the majority of users continue to operate with classical Externally Owned Accounts. The empirical reality reveals a tension: a growing share of ETH holdings is custodied on centralized exchanges, whose

users effectively forgo the self-custody property — whether out of convenience or ignorance of the alternative. The protocol-level possibility of self-custody and its actual use diverge, a discrepancy relevant to the assessment of security burden.

The fifth emergent property is programmable enforceability. From the interplay of smart contracts, the cryptographically enforced finality, and the Slashing mechanism, an enforcement layer for agreements emerges that operates without any external authority. Grimmelmann and Windawi have shown that blockchains function as semicommons — in the sense of the definition set out in Chapter 2, as systems in which private and shared resource use are intertwined and in which the rules are enforced through code.[103] Smart contracts enforce agreements automatically: a lending protocol liquidates an undercollateralized position without a court order as soon as the value of the pledged collateral falls below the threshold defined in the code. A decentralized exchange protocol executes a swap at exactly the terms the algorithm specifies, without counterparty risk. The Slashing system described in Section 4.3 penalizes validator misconduct without any authority and without a physical monopoly on force. Enforcement is algorithmic and deterministic: when the defined conditions are met, the prescribed consequence is executed, regardless of whether the affected party consents or whether an authority orders the enforcement. This property does not replace the legal system that Chapter 3 described as institutional meta-infrastructure, but it generates a complementary enforcement layer for cases in which institutional enforcement is too slow, too costly, unavailable across jurisdictions, or not envisaged for the specific use case. The practical relevance of this property is evident in DeFi use: lending protocols with a cumulative volume of billions of US dollars operate without a single legal agreement between the parties, and the security of the agreement rests on the correctness of the code and the finality of the base layer. The limit of this property lies in the expressive capacity of the code: what the smart contract cannot represent — in particular disputes over intentions, interpretations, negligent conduct, and non-contractual claims — remains dependent on institutional enforcement, and the interface between code-based and institutional enforcement is legally largely unresolved.

DeFi protocols rely on external data sources — so-called Oracles — for price information; these operate outside the Ethereum protocol. Oracle networks such as Chainlink deliver price feeds from external markets to smart contracts, thereby creating an additional trust layer not secured by the consensus protocol. The evaluation framework deliberately excludes the Oracle layer from evaluation through

[103] Cf. Grimmelmann/Windawi 2023, pp. 1097–1129, and the presentation in Chapter 2, Section 2.1.

the Oracle boundary defined in Chapter 2, because Oracles are application-layer components, not infrastructure-layer components.

These five properties ground the infrastructure claim that the following sections examine. The assessment asks whether the current manifestation of these properties meets the requirements of a fundamental digital infrastructure — for the existence of a property and the adequacy of its manifestation are two distinct questions, and the answer to the second requires systematic evaluation against the twelve criteria. A system can enable Permissionless Deployment and still fail at the access threshold if Gas costs effectively restrict access. It can offer Trustless Verification and yet generate a high trust burden if the majority of users access it through centralized intermediaries. It can establish programmable enforceability and still fall short of coordination requirements if L2 fragmentation undermines the coherence of the system. The assessment must systematically capture these tensions between possibility and reality — and that is precisely what the twelve criteria were designed for.

Four properties — two of them described in this section — cannot be fully operationalized with the indicators of the evaluation framework. Permissionless composability — the ability to create and atomically execute arbitrary protocol combinations without authorization — would require formal modeling of the interaction possibilities between smart contracts. Cryptographic system-state verification — the ability of every user to independently verify the entire network state — would require an empirical survey of actual verification practice. Decentralized self-developmental capacity — the system's capacity to evolve without a central authority — would require a long-term institutional analysis beyond the process description documented here. Finally, privacy is not anchored at the protocol level, which limits use cases with confidentiality requirements — in particular institutional actors with regulatory data protection obligations. These four dimensions appear in the synthesis (Section 6.3) as a research agenda, because they are relevant to a complete infrastructure assessment but cannot be adequately operationalized with the methods and indicators of this study.

4.7 Dimension I: Structural Foundation

The preceding sections have described Ethereum as a coherent system: its architecture, the transaction lifecycle, the security economics, the scaling model, the governance structures, and the institutional embedding. The assessment that now follows applies the twelve criteria of the evaluation framework to this described subject. It references the architecture. It does not recount it. Each criterion is evaluated on the four-level scale, and each evaluation makes the three-level logic visible in the flow of argument: what the protocol enables, what of that is operationally realized, and whether the interplay suffices for the infrastructure claim.

The first dimension examines whether Ethereum brings the structural prerequisites that can sustain a fundamental infrastructure. Four criteria operationalize this question. Two of them — Security and Trust Load (I.2) and Minimal Viable Guarantees (I.4) — are anchored as Critical Conditions, meaning that a finding of "Open" on either would cap the overall verdict through the M3 cascade at "Conditionally Suitable," regardless of performance across all other criteria. Two of the three Critical Conditions in the entire evaluation framework thus lie within this dimension.

I.1 Functional Irreplaceability

The first question put to a system asserting the infrastructure claim concerns its anchoring: is Ethereum structurally positioned to function as a Schelling point for decentralized coordination — the point, that is, on which actors converge without explicit coordination?

The protocol does not produce this anchoring through technical lock-in mechanisms. There is no enforced vendor commitment, no exclusive API, and no proprietary network stack. The open architecture described in Section 4.1 creates the preconditions for network effects but does not enforce them. In Frischmann's infrastructure terminology, Ethereum is accordingly a system that establishes its indispensability through the breadth and depth of productive activities built upon it, not through the access barriers it erects.[104] The implication is direct: enforced lock-in would be incompatible with the neutrality claim, and emergent anchoring is the only form compatible with the claim of a global and neutral infrastructure. Cloud infrastructure — one of the reference infrastructures from Chapter 3 — illustrates the pattern: AWS is de facto indispensable

[104] Cf. Frischmann 2012, pp. 61–96, on the characterization of infrastructure as an enabling structure whose value is realized in the activities it makes possible.

despite the absence of vendor lock-in, because the accumulation of tooling, integrations, and expertise generates switching costs that no technical lock-in mechanism could achieve.

The usage reality shows an anchoring that exceeds the threshold in three dimensions. The DeFi TVL dominance documented in Section 4.6 amounts to roughly 100 billion US dollars under Ethereum security and consistently represents over 60 percent of all on-chain-secured DeFi capital, with the concentration within the L2 ecosystem following a power-law distribution: Base holds approximately 46.6 percent of L2 DeFi TVL, Arbitrum approximately 30.9 percent, and the remaining more than 48 L2s share less than 25 percent. This concentration on Ethereum-native systems is an expression of liquidity network effects, because DeFi protocols access a depth of liquidity not comparably available on any other platform. Stablecoin issuance on Ethereum amounts to roughly 166 billion US dollars, representing 52 to 54 percent of the total market, with the three largest stablecoins — USDC, USDT, and DAI — having their deepest DeFi integration on Ethereum.[105] The strategic significance of this anchoring is amplified by the GENIUS Act (18 July 2025) as the first US federal law on stablecoins, which creates the regulatory framework within which stablecoin issuance operates and increases the institutional legitimation of the dollar-denominated instruments issued on Ethereum. The developer ecosystem comprises 31,869 active developers (as of September 2025), the historical high and over 70 percent of all blockchain developers.[106] Developer anchoring is the most stable of the three dimensions, because the Ethereum-specific toolchain generates high switching costs: a developer who has invested two years in EVM, Solidity, and Foundry will not migrate to another platform without a substantial productivity loss, and this inertia is a structural barrier independent of market cycles.

The shock resilience of this anchoring is empirically documented. In the 2022 market cycle, DeFi TVL collapsed in absolute terms by more than 75 percent, yet Ethereum's relative dominance position remained intact — indicating that the anchoring extends beyond the price mechanism and is held by network effects from liquidity, composability, and toolchain. Constraints are nonetheless documented. The stablecoin volume migration to Solana described in Section 4.4 shows that Ethereum remains the primary issuance platform but that transaction volume is increasingly migrating to other networks. The L2 autonomy — manifested in independent treasuries (Arbitrum DAO with over 3.5 billion US dol-

[105] Data on stablecoin market capitalization on Ethereum: DefiLlama, Stablecoins Dashboard, accessed early 2026.
[106] Electric Capital (2025): Developer Report, September 2025. The methodology counts active developers as persons who have contributed code to an Ethereum-related open-source repository within the past 30 days.

lars), proprietary business models (Base as a Coinbase product), and their own governance structures — means that technical attachment to Ethereum coexists with growing organizational independence. The multi-chain drift of institutional actors, exemplified by BlackRock BUIDL's deployment across seven chains since March 2025, relativizes the exclusivity of Ethereum anchoring.

In Frischmann's terminology, the Schelling point character is an enabling function: the anchoring generates its value by enabling productive activities, and the breadth of those activities constitutes the indispensability. The anchoring is emergent and thus in principle reversible, even though the historical evidence documents considerable persistence. The reversibility is not merely theoretical: L2 autonomy shows that ecosystem actors are actively seeking to reduce their Ethereum dependency. Functional Irreplaceability therefore stands at: Fulfilled with qualification. The anchoring exceeds all quantitative thresholds and has passed a stress test, but the emergent character of the anchoring, the L2 autonomy, and the multi-chain drift qualify the finding.

The anchoring shows that Ethereum functions as a center of gravity. The next question follows necessarily: is this center of gravity also secure?

I.2 Security and Trust Load

This criterion is anchored as a Critical Condition in the evaluation framework: a finding of "Open" would cap the overall verdict at "Conditionally Suitable." The examination operates on two dimensions — economic security and residual trust burden — which stand in a complementary relationship: one asks whether attacks are infeasible, the other how much trust the user must still place in the system despite this security.

The Casper FFG architecture, whose mechanics Section 4.2 described, anchors economic security in a mathematically binding way: validators who double-sign or attest to contradictory checkpoints lose their stake through Slashing, with the three critical thresholds of BFT mathematics — 33 percent as the blocking minority, 51 percent for reorgs before finality, and 66 percent as the supermajority — implemented in the protocol and not deactivatable.[107] In parallel, the protocol anchors the technical possibility of trustless access: full nodes verify the entire chain history independently, and Merkle Patricia Trie structures enable cryptographic state verification with logarithmic efficiency.

[107] Buterin, Vitalik / Griffith, Virgil (2019): Casper the Friendly Finality Gadget. arXiv:1710.09437. The BFT thresholds are implemented in the Beacon Chain specification and have been operative since the Merge (September 2022).

The data on the operational security profile are strong. With 964,768 active validators and approximately 30 percent of the ETH total supply of 120,693,582 ETH staked, the static cost of a 34-percent attack amounts to roughly 26 billion US dollars — calculated from approximately 34 percent of the 35.8 million staked ETH multiplied by the ETH price on the reference date of approximately 2,100 US dollars. This value is in the range of the threshold of over 30 billion US dollars, falling below it at the reference date's price level, so that the economic security unfolds its effect through the enforced Slashing mechanics, the scale of the attack, and market illiquidity — not through the static threshold alone. The calculation follows BFT logic: an attacker would need to control roughly one-third of staked ETH and would be compulsorily exposed to Slashing in a finality attack, causing the controlled stake to be entirely lost regardless of whether the attack succeeds. The Slashing record of 525 validators since genesis, documented in Section 4.3, demonstrates that the mechanism operates and that the deterrence signal is effective — the low count across nearly one million validators shows that rational actors avoid the Slashing risk. In the assessment window of the past 24 months, no finality loss spanning more than four epochs has occurred, and the finality rate exceeds 99.99 percent: of roughly 287,000 epochs since the Merge, only 13 were not finalized — all in May 2023 — and since that point, not a single failure has occurred.[108]

The procyclical constraint deserves a differentiated classification, because it concerns the nature of the security guarantee itself. The attack cost is a function of the ETH price, which stands in a circular dependency with the DeFi ecosystem: a declining price weakens the security guarantee precisely when it is most needed, and in a stress scenario this effect is amplified — roadmap complexity can produce bugs, bugs reduce confidence, declining confidence depresses the ETH price, and a declining price further reduces the attack cost. In the 2022 market cycle, the ETH price fell by more than 80 percent, which proportionally reduced the attack cost, yet the system survived empirically without a security incident — documenting robustness even under adverse conditions. This market-price dependency distinguishes Ethereum from traditional infrastructures: the physical security of an power grid does not decline with the electricity price, and the institutional security of SWIFT does not depend on the valuation of its member banks. The current attack cost of approximately 26 billion US dollars nevertheless surpasses conventional security guarantees in one qualitative respect: Slash-

[108] Finality data: beaconcha.in, 27 March 2026. The 13 non-finalized epochs in May 2023 were the result of two consecutive attestation-handling bugs in Prysm and Teku.

ing is mechanically enforced and depends neither on law enforcement nor on institutional control.

The residual trust burden reveals a discrepancy between protocol claim and usage reality that must be captured in its cumulative effect. Trustless access is technically possible: a full node on consumer hardware, as documented in Section 4.6, enables independent verification without an intermediary, and the protocol does not enforce any intermediary. The operational reality, however, has produced a layering of trust dependencies, where each layer adds a dependency that the infrastructure claim ostensibly precludes, and where the cumulation of layers generates a trust burden qualitatively different from examining each layer individually.

Wallets such as MetaMask operate as closed-source components with unilateral update mechanisms, and the user trusts the wallet provider to construct the correct transaction and use the right RPC endpoint. On the next layer, Infura and Alchemy dominate approximately 70 percent of RPC traffic, with the Infura outage of November 2020 empirically demonstrating that the systemic dependency is operationally manifest: Ethereum was temporarily inaccessible to the majority of users even though the protocol continued to operate throughout. Above that lies the L2 sequencer layer: 91.5 percent of L2 TVL operates on Stage 1 rollups with security council multisigs and centralized sequencers, with the Linea incident of June 2024 — in which the team deliberately stopped the sequencer and censored attacker addresses — demonstrating the censorship capability of this architecture. At the outermost layer, cross-L2 transfers require trust-based bridges with cumulative hack damages exceeding 2.8 billion US dollars, or seven-day challenge periods, with exit window conditions varying drastically: Base has an exit window of zero days, Arbitrum One of 17 days. The cumulation of these four layers means that a typical user accessing an L2 DeFi service via MetaMask traverses four trust layers, none of which is secured by the Ethereum protocol itself.

The consequence of this layering can be stated precisely: the cryptographic trustlessness that Ethereum claims as infrastructure applies in the current state primarily to a small group of users consisting of full node operators, on-chain protocols, and L1 users with direct protocol access. For the mass user and for the majority of L2 capital, Ethereum is an infrastructure with institutional trust layers. The qualitative difference from traditional systems remains, because the opt-out is always available: any user can run a full node and verify the entire chain themselves, and this right is anchored at the protocol level. Quantitatively, few exercise this opt-out, and the barriers — hardware, operation, maintenance, tech-

nical knowledge — are real. Security and Trust Load therefore stands at: Fulfilled with qualification. The economic security dimension satisfies all technical indicators. The residual trust dimension documents an operational discrepancy between protocol possibility and usage reality that qualifies the claim without fundamentally undermining it.

Economic security is given. The next question is: what is this system secure for? A secure system without coordination capacity would be a vault without contents.

I.3 Coordination Function

This criterion examines, following the definition established in Chapter 2, whether Ethereum provides a coordination capacity that could not be provided without the system, or only at significantly higher cost. The reference infrastructures from Chapter 3 show the pattern: power grids coordinate supply and demand in real time, SWIFT coordinates financial messages among over 11,000 banks, the Internet coordinates packet routing between autonomous networks. The question is whether Ethereum provides a comparable coordination capacity, and if so, of what kind.

The protocol operationally implements three of the four coordination primitives described in Sections 4.2 and 4.4: Settlement as finalized, irreversible transaction processing after approximately 12.8 minutes; Execution as deterministic smart contract processing in the EVM with atomic composability; and Data Availability as KZG-commitment-based blob availability for L2 rollups since Dencun (March 2024). The fourth primitive — protocol-side verification of arbitrary off-chain computations, carried by zkEVM verification and exposed through the EXECUTE precompile of Native Rollups (EIP-8079) — has the status RES and is therefore not part of the current-state assessment. Three of four primitives with the status IMPL document a coordination scope that exceeds the threshold for "Fulfilled with qualification."

The absolute figures merit a contextualization that goes beyond the internal blockchain comparison. Roughly 100 billion US dollars are coordinated in a permissionless system without an institutional trustee, and this volume is processed around the clock, 365 days a year, without maintenance windows and without holiday calendars. Fedwire, the largest US settlement system, settles approximately 4 trillion US dollars daily, but operates with an institutional sponsor (the Federal Reserve), geographic restriction (the US dollar space), and access restriction (Federal Reserve membership required). Ethereum's volume cannot be equated with Fedwire's, because the systems serve different functions and dif-

ferent user bases, but the coordination property — permissionless, global, without intermediary — is structurally distinct and cannot be replicated by scaling up a traditional system: no institutional settlement system can operate permissionlessly, because institutional sponsorship implies access control. Atomic composability, as described in Section 4.1, enables Flash Loan-based operations within a single block that would require multi-day settlement latency in traditional systems, and the stablecoin issuance of approximately 166 billion US dollars on Ethereum shows that the coordination capacity already includes institutional actors issuing dollar-denominated values on the Ethereum infrastructure.[109]

The L2 settlement dependency documents the operational reach. All top-5 rollups by TVL — Arbitrum One, OP Mainnet, Base, zkSync Era, and Starknet — settle on Ethereum and use blob transactions for Data Availability, with hybrid exceptions such as Arbitrum Nova or Mantle holding no significant TVL share. This binding is technical, because no alternative offers comparable security guarantees at comparable cost. The DA decisions illustrate the dependency: all leading rollups use Ethereum blobs as their primary DA source, and the Fusaka upgrade (December 2025) described in Section 4.4 further increased blob capacity through PeerDAS, strengthening the economic attractiveness of Ethereum DA relative to alternative DA solutions. The entire block production pipeline operates via MEV-Boost, with the Proposer-Builder Separation described in Section 4.2 organizing coordination between validators and builders through off-chain relays that themselves bear no protocol-level accountability and operate on reputation and economic self-interest. The relay dependency is part of the coordination architecture and is therefore relevant to I.3: the system coordinates block production through a layer that lies outside the protocol.

Cross-L2 composability, however, is the qualitative constraint that determines the overall finding. Ethereum's strongest coordination property — atomic composability — is structurally absent at the L2 level: each L2 processes transactions in its own, asynchronous state spaces, and cross-L2 transfers require trust-based bridges or multi-day challenge periods. The Superchain with its 34 OP-Chains has made progress in intra-stack interoperability, with SuperchainERC20 enabling cross-chain token transfers within the OP Stack, but atomic composability across stack boundaries — for instance between OP Mainnet and Arbitrum — remains unavailable. Monthly bridge volumes of approximately 11.2 billion US dollars show the scale of coordination demand handled through

[109] For comparison of settlement speeds: SWIFT settlement takes 1–3 days, ACH 2–3 days, SEPA transfers typically 1 business day. Ethereum's 12.8-minute finality and atomic composability enable coordination patterns that are structurally impossible in these systems.

intermediaries, and bridge hack damages exceeding 2.8 billion US dollars document the costs of this intermediation. The liveness risks of centralized sequencers reinforce the picture: Arbitrum experienced a 78-minute outage in December 2023, Base a 33-minute outage in August 2025, each caused by single points of failure in the sequencer infrastructure.

Measured against the infrastructure claim, a two-tier nature of coordination emerges that is decisive for the assessment. For L1 users, coordination is trustless and protocol-guaranteed: Settlement is irreversible, Execution is deterministic, and the interaction requires no intermediary. For L2 users — who account for the majority of actual activity — coordination is constrained by centralized sequencers, relay dependencies, and bridge trust. An infrastructure that claims universal coordination without institutional trust, but handles the majority of its coordination through institutional intermediaries, exhibits a claim-reality discrepancy. The discrepancy is qualitatively different from traditional infrastructures: SWIFT coordinates throughout via institutional intermediaries and asserts no trustlessness, whereas Ethereum asserts the claim and delivers it on L1, but relativizes it on L2. The finding for Coordination Function is: Fulfilled with qualification. The coordination capacity is real, quantitatively substantial, and qualitatively unique, but L2 fragmentation undermines the coherence of coordination that the system must provide as infrastructure.

Coordination is real and unique. The decisive follow-on question is whether it holds under stress — for a system that coordinates under normal conditions but fails under disruption is not an infrastructure.

I.4 Minimal Viable Guarantees

This criterion is the second Critical Condition in this dimension and asks whether Ethereum provides a minimum set of guarantees that are maintained even under stress. The criterion has a special function in the evaluation framework because it examines the system's robustness under adverse conditions — conditions constitutive of the infrastructure claim: each of the seven reference infrastructures from Chapter 3 had to absorb stress situations, from power grid blackouts through DNS attacks to SWIFT sanctions, and the question of whether the system degrades or collapses is decisive for infrastructure suitability. Four guarantees are under examination: Finality, Liveness, Degradation Mode, and censorship resistance under attack.

The time-to-finality is 12.8 minutes, calculated from two epochs of 32 slots at 12 seconds each, and thus falls below the threshold of 15 minutes.[110] This value is anchored in protocol mathematics and has been consistent since the Merge. For a system asserting the infrastructure claim, the quality of this finality is decisive: SWIFT settlement takes one to three days, ACH two to three days, SEPA transfers typically one business day. The 12.8-minute finality is qualitatively superior, because it is carried by cryptographic irreversibility requiring no institutional settlement. Liveness is documented by a finality rate above 99.99 percent, with not a single failure occurring since May 2023. The December 2025 incident, in which a Prysm bug temporarily reduced network participation to 75 percent, empirically demonstrated that improved client diversity functions as a safety buffer: had the same bug occurred in 2022, when Prysm's share stood at approximately 68 percent, participation could have fallen below the 67-percent threshold. The fact that the system absorbed this bug without any user noticing an effect illustrates the robustness of the finality guarantee.

The Degradation Mode is the property that distinguishes Ethereum from all seven reference infrastructures analyzed in Chapter 3, and it warrants a contextualization that goes beyond the technical description in Section 4.2. The Inactivity Leak described there was tested on Mainnet in May 2023, when two consecutive finality failures pushed network participation below the two-thirds threshold. The protocol responded automatically: inactive validators progressively lost stake until the remaining active validators again represented more than two-thirds of the set, and self-healing occurred after approximately 96 minutes — without manual intervention, without coordination outside the protocol, and

[110] The Beacon Chain specification defines an epoch as 32 slots of 12 seconds each. Finality requires the justification and finalization of two consecutive checkpoints by more than two-thirds of the validator set.

without interruption to block production. The infrastructural significance of this mechanism is sharpened by comparison with the reference infrastructures: when part of the DNS root server network fails, coordination is required — typically through manual escalation between root server operators. When SWIFT nodes fail, institutional emergency procedures engage. When a partial power outage occurs, restoration requires physical intervention by grid operators. In each of these cases, recovery depends on external coordination that in turn presupposes communication infrastructure, organizational structures, and human decision-making processes. Ethereum's Inactivity Leak requires none of this. The protocol returns the system to functionality autonomously, and the Mainnet proof makes this finding robust.[111]

The fourth guarantee — censorship resistance under coordinated attack — does not reach the level of the first three. No protocol mechanism compels a block proposer or builder to include a specific valid transaction. FOCIL (EIP-7805), which would close this gap, has the status PLAN. The OFAC compliance rate, fully analyzed under criterion II.1, stands at approximately 15 percent — well below the 50-percent threshold — but this performance rests on market dynamics, not on a protocol guarantee, which means it is reversible under changed market conditions.[112] The builder concentration documented in Section 4.2 — which receives full analysis as the main topic under criterion II.1 — means for I.4 that censorship resistance is emergent and not protocol-guaranteed. This single indicator stands at "Conditionally fulfilled," because the path to protocol-level assurance is discernible, but not secured in the current state.

The aggregation of the four guarantees produces a picture that precisely locates the system's strengths: three guarantees stand at "Fulfilled," one at "Conditionally fulfilled." The three fulfilled guarantees cover the core of the minimum set, with Finality and Liveness securing the base function and the Degradation Mode providing a property that goes beyond what the reference infrastructures from Chapter 3 offer. The conditional guarantee concerns censorship resistance under coordinated attack, which is functionally given as long as market conditions remain favorable but is not secured by the protocol against unfavorable conditions. The stronger of the two possible labels is justified, because the three fulfilled guarantees form the core and the fourth has a clear path to closure. Minimal Viable Guarantees therefore stands at: Fulfilled with qualification. Three of

[111] The technical details of the May 2023 Inactivity Leak event are documented in Section 4.2. The assessment here concerns the infrastructural significance: automated self-healing without external coordination.

[112] OFAC compliance data: MEV Watch / relayscan.io, 27 March 2026. The historical trajectory from the 79-percent peak documents the market-driven normalization.

four guarantees are protocol-anchored and operationally proven. The fourth is functionally given, but not secured by the protocol.

Synthesis Dimension I

The assessment profile of the first dimension shows a consistent pattern: all four criteria stand at "Fulfilled with qualification." The homogeneity is remarkable, because no criterion falls below this and none achieves "Fulfilled" without qualification. Functional Irreplaceability (I.1) is characterized by strong but emergent and thus reversible anchoring. Security and Trust Load (I.2) documents robust economic security alongside an operational residual trust burden from the four-layer trust stack. Coordination Function (I.3) evidences a quantitatively and qualitatively unique coordination capacity, constrained by L2 fragmentation and relay dependency. Minimal Viable Guarantees (I.4) shows three strong and one conditional guarantee, with the Degradation Mode representing the outstanding single property of the entire current-state assessment.

The two Critical Conditions — I.2 and I.4 — stand at "Fulfilled with qualification." In the cascade logic of Chapter 2, this means that the core property is fundamentally given, with documented limitations, and this state is passable for the next cascade stage. These two Critical Conditions do not constrain the overall verdict downward. The pattern that characterizes the entire dimension can be condensed into a single sentence: the architectural foundation is load-bearing, and the constraints lie consistently in the discrepancy between what the protocol enables and what operational reality delivers from it. Four constraints are manifestations of the same underlying pattern. The four-layer trust stack reduces the cryptographic trustlessness to an institutional trust dependency for the mass user. The fragmented L2 composability restricts the system's strongest coordination property to the base layer. The emergent censorship resistance depends on market dynamics rather than being guaranteed by the protocol. The reversible anchoring, finally, rests on network effects rather than technical binding. The question of whether the path from possibility to reality is secured by the roadmap belongs to the target-state assessment of the following chapter.

The structural foundation is given. The next dimension examines the qualitative suitability of the system for the infrastructure claim — and that is where the sharpest single finding of the entire current-state assessment lies.

4.8 Dimension II: Qualitative Viability

The second dimension examines whether Ethereum possesses the qualitative properties that substantively carry the infrastructure claim. Four criteria are under assessment: Neutrality and Censorship Resistance (II.1) as the sole Critical Condition of this dimension — and the third and final one of the entire evaluation framework — followed by Open Generativity (II.2), Independent Verifiability (II.3), and Low-Threshold Inclusivity (II.4). The asymmetry between the Critical Condition — which carries the sharpest single finding of the entire current-state assessment — and the three Qualitative Criteria — which show a consistent profile without fundamental deficits — defines the overall picture.

II.1 Neutrality and Censorship Resistance

Of the twelve criteria in the evaluation framework, this is the one with the sharpest finding in the current state. It is anchored as a Critical Condition and contains the only indicator standing at the level "Open." The examination asks: does Ethereum operate neutrally toward all users, and can every valid transaction be included?

The protocol anchors neutrality through Permissionless Participation without authorization by any central authority, content-neutral block validation through Casper FFG, and fork-choice neutrality without whitelist or blacklist. What the protocol does not anchor is an inclusion obligation: no mechanism compels a block proposer or builder to include any specific valid transaction. This gap — described in Section 4.2 — is the protocol-level basis of censorship vulnerability.

Seven indicators operationalize the examination question, and their assessment produces a heterogeneous picture that shows the tension between operational performance and structural assurance more sharply than any other criterion. The OFAC compliance rate of approximately 15 percent of blocks falls well below the 50-percent threshold and documents a positive trajectory from the 79-percent peak in November 2022, with normalization occurring without protocol intervention through market dynamics: non-OFAC-compliant builders gained market share because they had no restrictions on transaction screening and could thus build more profitable blocks. The Van Loon v. Treasury ruling of November 2024 and the OFAC delisting of Tornado Cash in March 2025

reduced regulatory pressure.[113] Even at the 79-percent peak, sanctioned transactions experienced only incremental delays of a few blocks, not permanent exclusion. The maximum censorship delay is operationally a matter of a few blocks, because approximately 85 percent of builder slots are built by non-OFAC-compliant actors. The performance is, however, emergent and not protocol-guaranteed, which means it is reversible under changed market conditions.

The structural defect finding lies in builder concentration. Titan with 51.2 percent, BuilderNet with 25.7 percent, and Quasar with 16.4 percent together control 93.3 percent of all blocks, with a Herfindahl-Hirschman Index of approximately 3,554 — substantially exceeding the DOJ threshold for highly concentrated markets of 2,500.[114] This concentration is not coincidental, as Daian et al. showed in 2020: builders implementing the most efficient MEV extraction strategies systematically win the competition for the most profitable blocks, generating a self-reinforcing concentration dynamic.[115] Titan alone produces more than half of all Ethereum blocks. If Titan and BuilderNet coordinate, they control 76.9 percent of block production — sufficient for de facto censorship of the majority of all MEV-Boost blocks. The protective factors against this coordination are economic in nature: censorship reduces builder revenue, and proposers can build locally — but with 2.66- to 5.57-fold lower yield. These protective factors are real, but they are market incentives, not protocol guarantees, and the experience of the 79-percent OFAC peak shows that regulatory pressure can override economic incentives. This indicator stands at "Open" and is the only Open finding of the entire current-state assessment.

FOCIL (EIP-7805) is the specified solution that would oblige proposers at the protocol level to adopt transactions from a in a decentralized way assembled inclusion list: a randomly selected validator committee of 16 validators assembles the list, and the proposer of the next block is protocol-bound to include these transactions provided they are valid and profitable. The mechanism would transform the emergent market performance into a protocol guarantee. The status PLAN with a clear timeline documents a discernible path, but no result secured in the current state, because governance decisions and technical feasibility are still pending. Geopolitical jurisdictional diversity shows moderate concentration:

[113] Van Loon v. Department of the Treasury, 5th Circuit Court of Appeals, November 2024: immutable smart contracts are not "property" within the meaning of the IEEPA. OFAC delisting of Tornado Cash: 21 March 2025.

[114] Builder market shares: relayscan.io, 27 March 2026. HHI calculated as the sum of squared market shares. The DOJ Horizontal Merger Guidelines of 2010 define an HHI above 2,500 as a highly concentrated market.

[115] Daian, Philip et al. (2020): Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In: *IEEE Symposium on Security and Privacy*, pp. 910–927.

approximately 39 percent of nodes operate in the USA and approximately 53 percent in two countries, with the infrastructure contextualization from Chapter 3 putting the pattern in perspective: SWIFT has been geopolitically instrumentalized, GPS is under unilateral US control, and cloud infrastructure is subject to the jurisdiction of the operator's country. Ethereum's situation is not comparable to the unilateral control of GPS, but the risk implication is real, because institutional staking through US-regulated entities amplifies jurisdictional exposure.[116] The cloud concentration of approximately 59 percent of hosted Execution Layer nodes on three providers — with AWS alone hosting 35.5 percent — expands US jurisdictional exposure beyond mere node geography.

The staking concentration indicators are within acceptable ranges. Lido holds 22.8 to 23 percent of staked ETH — well below the 33-percent blocking minority threshold and declining from approximately 32 percent in 2023 — with the DVT integration of 547,968 ETH and a quarter-on-quarter growth of 57 percent reducing the single-point-of-failure risk in the operator set.[117] The cumulative top-3 staking provider concentration stands at 40 to 45 percent and exceeds the 33-percent blocking minority threshold as a coordinated sum, while no individual provider reaches it.

The overall picture is asymmetric. The transaction layer is operationally sound, the LST indicators are within the acceptable range, but builder concentration as the single Open indicator constitutes a structural defect. The system's censorship resistance results from market dynamics, not from protocol design, and for a system claiming that every valid transaction can be included, dependence on three commercial entities is a structural gap that cannot be closed by pointing to the positive OFAC trajectory: the OFAC normalization shows that the market produces censorship resistance under favorable conditions, but the builder concentration shows that the structure of the market undermines the prerequisites for this performance. The overall assessment of the criterion is nonetheless not "Open," because the aggregation of all seven indicators shows that the transaction layer functions, regulatory pressure is easing, LST concentration is declining, and a clear path to protocol-level assurance exists. Neutrality and Censorship Resistance therefore stands at: Conditionally fulfilled. Fulfillment depends on conditions not secured in the current state — in particular the implementation of FOCIL or a functionally equivalent mechanism. This is the weakest sin-

[116] Node distribution data: Ethernodes, early 2026. Cloud concentration estimate: approximately 59 percent of hosted EL nodes on AWS (35.5%), Hetzner (13.8%), and OVHcloud (9.7%).

[117] Lido Tokenholder Update, 26 February 2026: 547,968 ETH on DVT (Distributed Validator Technology), +57 percent QoQ.

gle finding of the entire current-state assessment and the only one in which a Critical Condition stands below "Fulfilled with qualification."

The neutrality gap is documented. The very next property — open generativity — is the system's strongest. The tension between the most critical weakness and the greatest strength of the system reflects a system that produces maximum openness but does not secure the neutrality of that openness at the protocol level.

II.2 Open Generativity

The architecture described in Section 4.1 implements Permissionless Deployment in full: any Ethereum address can deploy smart contracts without a whitelist, approval requirement, or registration obligation, and the EVM as a quasi-Turing-complete machine excludes no category of application logic at the protocol level. Atomic composability enables interaction patterns such as Flash Loans that have no equivalent in any traditional system: a user can, in a single transaction, borrow millions of US dollars, deploy them for arbitrage, and repay them — without collateral and without an intermediary — and if the repayment fails, the entire transaction is reversed. The ERC standard family has established itself as an emergent interoperability mechanism without any central authority, with standards such as ERC-20, ERC-721, and ERC-4626 defining the interfaces between protocols without any central body prescribing them. The interplay of these three properties — permissionless access, atomic interaction, and emergent standards — generates the generativity that produced the DeFi ecosystem: protocols build on one another without needing to conclude bilateral agreements, because compatibility is established at the level of the shared execution environment.

Practice confirms the protocol property. Permissionless Deployment is fully realized, approximately 65 percent of all new smart contracts are deployed directly on L2s, and in 2025, 230 EIPs were submitted, of which 37 were accepted — documenting that the innovation process itself is permissionless: anyone can propose standards, and acceptance follows a community-driven process without a central gatekeeper. The only de facto deployment constraint on L1 is the Gas Limit of 60,000,000, which in practice presents no obstacle for the majority of application logics, as even complex smart contract deployments do not exhaust the block gas limit.

The toolchain merits independent consideration because it opens a quality dimension for generativity that goes beyond Permissionless Deployment. The entire core development infrastructure stands under liberal open-source licenses (MIT, Apache 2.0): Foundry as a testing and deployment framework, Hardhat as

a build environment, Ethers.js as a JavaScript library for blockchain interaction, Slither and Echidna as security tools. No single provider controls the toolchain or can restrict access, and the tools are maintained by independent teams — ensuring the resilience of the development infrastructure against the failure of any single provider. This is relevant to generativity because a developer needs not only the ability to run code on the network but also freely available tools to write, test, and verify that code. Van Schewick showed for the Internet that open tools and open protocols together generate the innovation dynamic, because they shift the innovation barrier from the infrastructure layer to the edges of the network. The same pattern holds for Ethereum at the smart contract layer, where any developer with access to open-source tools can innovate without seeking permission from the protocol operator.[118]

The constraint lies in L2 fragmentation, and it strikes precisely the property that produced Ethereum's generativity. The composability fragmentation documented in Section 4.4 means that a developer needing liquidity across multiple L2s cannot use L1's atomic composability and is dependent on bridge mechanisms with trust assumptions. Flash Loans — the paradigmatic example of the generative power of atomic composability — function only within a single block on a single chain and are structurally impossible across L2 boundaries. The richest composability property is available on the layer where only a small fraction of transaction activity takes place, while the L2 layer — where most activity occurs — develops its own, internal composability patterns that do not interact across stack boundaries. For the generativity claim, this is a real constraint, because the network effects arising from composability are weakened by fragmentation: each L2 ecosystem develops its own liquidity pools, its own DeFi protocols, and its own user groups, and the connection between them requires intermediaries that introduce trust assumptions. Open Generativity therefore stands at: Fulfilled with qualification. The permissionless deployment model and the open toolchain constitute the system's structurally strongest property. L2 composability fragmentation is the documented constraint.

If the system is generative — if it can be built upon — the complementary question arises: can the results also be verified?

[118] Cf. van Schewick, Barbara (2010): *Internet Architecture and Innovation*. MIT Press. Van Schewick's argument that the end-to-end architecture of the Internet enables innovation at the edges is transferable to Ethereum's permissionless smart contract layer.

II.3 Independent Verifiability

The supply side of verification complements the residual trust burden (I.2): while I.2 asks how much trust remains despite the possibility of verification, this criterion asks whether the possibility of verification exists at all.

The deterministic EVM, described in Section 4.1, guarantees that every full node computes the identical state from identical input. Merkle proof verifiability enables cryptographic verification of individual state entries without retracing the entire state — relevant for light clients and resource-constrained verifiers. The EVM specification is formally defined via the Yellow Paper and machine-verified through KEVM, creating a foundation for machine verification that has no equivalent in traditional financial systems: the correctness of a SWIFT transaction is institutionally asserted and confirmed through bilateral reconciliation between sender and recipient bank, whereas the correctness of an Ethereum transaction can be mathematically demonstrated — any full node independently computing the state change.[119] Full-node verification without third parties is technically possible, and the protocol requires no intermediary for state verification, creating a strong foundation on the supply side.

The operational evidence is substantial. At least six productively used verification tools exceed the threshold: Certora Prover for formal verification of invariants, Foundry/Forge for fuzz testing and property-based testing, Echidna as a fuzzer, Mythril for symbolic analysis, Slither for static analysis, and KEVM as a machine-verified reference implementation. The EVM specification is substantial but has documented gaps at precompiles and a lag in formalizing recent EIPs, meaning that newer protocol features are at times less completely specified than older ones. The audit culture is standard for Tier 1 DeFi protocols, with protocols such as Aave, Compound, and Uniswap regularly undergoing multiple independent audits by specialized firms such as Trail of Bits, OpenZeppelin, and Spearbit before new versions are deployed. The universality of the practice remains limited, however, because Permissionless Deployment means that anyone can deploy unaudited contracts, and the long tail of protocols lacks both the resources and the incentive for comprehensive audits. The Lazarus Group attacks of 2025, with damages of 2.02 billion US dollars, demonstrate that even audited systems remain vulnerable through attack vectors such as private key compro-

[119] The Yellow Paper (Wood 2014/2024) formally defines the EVM semantics. KEVM (Hildenbrandt et al. 2018) provides a machine-verified formalization in the K Framework.

mise and Oracle manipulation — underscoring the status of audits as a necessary but not sufficient security condition.[120]

Verification of Layer 2 states adds a layer of complexity beyond L1 verification, because a system claiming cryptographic verifiability as a core property must deliver on this promise at all layers. For Optimistic Rollups, the state is finalized after the seven-day challenge period, with security resting on the assumption that at least one honest prover submits a fraud proof within the deadline. This verification model is permissionless — anyone can submit a fraud proof — but requires the prover to be able to fully replicate the L2 state, which demands specialized infrastructure and technical knowledge. For ZK-Rollups, verification is immediate and mathematical: an on-chain STARK or SNARK proof can be verified by any L1 node. Creating and understanding these proofs, however, requires cryptography expertise substantially beyond standard EVM verification — from STARK mathematics through SNARK trusted setup assumptions to ZK circuit logic — and the number of actors capable of performing this verification independently is small. Native Rollups (EIP-8079, status RES) would return verification to the base layer and thereby unify it, with the EVM itself serving as the verification environment for off-chain computations, but they are not part of the current-state assessment. The growing L2 verification complexity is a constraint that co-determines the finding of II.3 at the operational level.

The RPC concentration of 70 percent across two providers — assessed under I.2 as a residual trust indicator — shows that the majority of users do not avail themselves of the verification possibility. This is not strictly an II.3 problem, because II.3 examines the supply side, but the infrastructural implication is real: a system that appeals to cryptographic verifiability must ensure that verification remains practically accessible, and the growing full-node hardware requirements (Section 4.1 documents a full-node size of 1,579 gigabytes with an upward trend) could constrain the supply side in the medium term if verification increasingly requires specialized hardware. The assessment for Independent Verifiability yields: Fulfilled with qualification. The verification offering is protocol-anchored and substantive on the tooling side, but L2 verification complexity is growing, specification gaps exist, audit culture is not universal, and the practical accessibility of verification is challenged in the long term by rising hardware requirements and L2 layering.

[120] Lazarus Group data: Chainalysis Report, February 2026. The attacks occurred primarily through private key compromise and cross-chain message spoofing at audited protocols.

If building and verifying are possible, the closing question of the qualitative dimension remains: can participation also occur?

II.4 Low-Threshold Inclusivity

This criterion has the second-largest indicator set in the evaluation framework — five indicators — because inclusivity operates at the intersection of several barrier types: hardware, capital, cost, and knowledge.

Full-node hardware requirements are in the consumer range, with costs of 500 to 1,500 US dollars for a setup of NVMe SSD, 16 to 32 gigabytes of RAM, and a broadband internet connection. This is above standard consumer hardware — a typical laptop with a 512-gigabyte SSD does not suffice — but below datacenter-grade hardware, with resource-efficient clients such as Nimbus enabling operation on a Raspberry Pi 5. The state growth problem documented in Section 4.1 is the central long-term threat: with a full-node size of 1,579 gigabytes and a weekly growth of approximately 14 gigabytes, requirements rise progressively, and projections for a Gas target of 200 million yield 5 terabytes by 2028 and 9 terabytes by 2030. The bandwidth requirement of at least 25 megabits per second excludes household networks in regions with poor infrastructure, constraining the global inclusivity claim.

The 32-ETH minimum deposit for solo staking is the deficit that shapes the entire inclusivity assessment, and it merits a contextualization that goes beyond the bare data point. At the price level of the first quarter of 2026, 32 ETH corresponds to approximately 67,000 US dollars at the reference date's price, and solo staking has shrunk to less than one percent of total staking volume.[121] The hardware barrier for validator operation is low, but the economic barrier excludes the broad majority of potential participants. The consequence is a usage distribution in which Liquid Staking holds 31.1 percent of staked ETH (Lido dominant), CEX Staking 24 percent, Staking Pools 16 percent, and solo stakers under one percent — the four categories covering approximately 72 percent of staked ETH. This distribution means that decentralized validator control — which the infrastructure claim of a permissionless infrastructure requires — is de facto concentrated among institutional entities, while the broad user base participates through delegated models.

This deficit is the only one in the entire current-state assessment for which the roadmap contains no operational path. EIP-7251 (MaxEB, Pectra) increased

[121] L2BEAT documents the distribution of staking methods. Solo stakers under one percent: estimate based on Dune Analytics, Q1 2026.

the maximum effective balance to 2,048 ETH but targets the efficiency of large stakers, not the access threshold. Rainbow Staking — which would enable various staking types with different capital requirements — has the status RES without an implementation commitment. For other constraints in the current-state assessment, a feature with the status PLAN or DEPL marks the path to closure: FOCIL for censorship resistance, EIP-4444 for state growth mitigation. The 32-ETH barrier and the L2 verification gap share the property that their only named resolution path has the status RES — Rainbow Staking in one case, and the return of verification to the base layer through Native Rollups in the other. Unlike the L2 verification gap, behind which an active research program with a draft EIP stands in the case of Native Rollups, the 32-ETH barrier has no such program, and this difference qualifies the deficit in a particular way. Liquid Staking through providers such as Lido, Rocket Pool, or Coinbase resolves access to staking yield from amounts as low as 0.01 ETH, but the user thereby delegates validator responsibility to third parties and gains yield, not control. Economic access is inclusive. Access to decentralized validator control remains behind the barrier.

L2 transaction costs document the strongest improvement relative to earlier states of the system. Median costs on the leading L2s have fallen below 0.01 US dollars per transaction after EIP-4844 (Dencun, March 2024) and the Pectra upgrade — orders of magnitude below the reference values of credit card fees (two to three percent) and bank transfers (0.50 to 5.00 US dollars). Before Dencun, Optimism fees were approximately 1.00 US dollar per transaction. The percentage figures in Section 4.4 refer to different survey periods and system baskets and cannot be chained arithmetically with this value pair. The reduction is a direct consequence of the blob architecture described in Section 4.4. A user in sub-Saharan Africa or Southeast Asia can participate in the global financial system at under one cent per transaction — representing substantial inclusivity for transactional use. L1 transactions remain expensive, typically 2 to 15 US dollars for complex DeFi operations during high-load phases, but since most activity has migrated to L2s, the L1 fee problem is not a critical inclusivity issue for typical usage scenarios.

Usability without cryptographic prior knowledge has improved substantially through EIP-7702 (Pectra, May 2025). Account Abstraction enables session keys, social recovery, and Gas sponsoring at the protocol level: a user can configure their account so that an application is granted limited transaction rights, a lost key can be recovered through trusted contacts, and transaction costs can be covered by wallet providers — allowing users without their own ETH holdings to

interact. The EIP is deployed, but integration into the major wallet applications is still being implemented, and the typical user still encounters seed phrases and Gas management as the primary interaction patterns. The adoption dynamic through the older ERC-4337 standard — documented by over 25.5 million smart accounts and 132 million UserOperations — shows the demand for programmable accounts. The UX transformation is, however, not complete, and the transition phase in which the old model (EOA with seed phrase) and the new model (Smart Account with social recovery) coexist tends to generate more rather than less complexity for the typical user, because wallet providers must serve both paradigms in parallel and the onboarding experience remains fragmented.

The profile of the five indicators shows no finding below "Fulfilled with qualification," with L2 transaction costs as the only indicator standing at "Fulfilled." The 32-ETH barrier as the only current-state deficit without a roadmap answer is the most significant constraint. Low-Threshold Inclusivity therefore stands at: Fulfilled with qualification. Participation is substantially inclusive for transactional use on L2, possible but endangered by rising state size for full-node operation in the consumer range, and prohibitively constrained for solo staking by the 32-ETH barrier.

	Criterion	Hierarchy Level	Current-State Assessment
I.1	Functional Irreplaceability	Structural	Met with Qualification
I.2	Security and Trust Load	Critical	Met with Qualification
I.3	Coordination Function	Structural	Met with Qualification
I.4	Minimal Load-Bearing Guarantees	Critical	Met with Qualification
II.1	Neutrality and Censorship Resistance	Critical	Conditionally Met
II.2	Open Generativity	Qualitative	Met with Qualification
II.3	Independent Verifiability	Qualitative	Met with Qualification
II.4	Low-Threshold Inclusivity	Qualitative	Met with Qualification

Seven of the eight criteria carry the same label. II.1 is the only outlier and is simultaneously critical in the cascade.

Figure 4.10 Current-State Assessment Profile Dimensions I and II — the eight individual assessments with hierarchy level and assessment label

Synthesis Dimension II

The assessment profile of the second dimension shows a pronounced asymmetry. The Critical Condition (II.1) stands at "Conditionally fulfilled," and the three Qualitative Criteria (II.2, II.3, II.4) stand consistently at "Fulfilled with qualification." The weakest finding of the entire current-state assessment lies in this dimension, driven by the Open indicator of builder concentration. The dimension contains both the system's most critical vulnerability — the absence of a protocol-level neutrality guarantee — and one of its structurally strongest properties — open generativity. The tension between the two findings is not a contradiction: it reveals a system that produces maximum openness but leaves the neutrality of that openness operationally to market dynamics.

The three Qualitative Criteria show no fundamental deficits, and their constraints follow the basic pattern of the entire current-state assessment: the protocol level delivers the right properties; the operational level realizes them incompletely. Open Generativity (II.2) combines Permissionless Deployment with an open toolchain, constrained by L2 composability fragmentation — which undermines precisely the property that produced the DeFi ecosystem. Independent Verifiability (II.3) is protocol-anchored and substantive on the tooling side, with growing L2 verification complexity and specification gaps as documented constraints. Low-Threshold Inclusivity (II.4) has improved substantially through blob transactions and Account Abstraction, with the 32-ETH barrier as the only current-state deficit without a roadmap path — blocking access to decentralized validator control while transactional access offers substantial inclusivity.

II.1 as a Critical Condition carries a weight that exceeds the weight of the Qualitative Criteria. "Conditionally fulfilled" at a Critical Condition is qualitatively a different state from "Fulfilled with qualification," because the core property of the infrastructure claim — the protocol-level neutrality guarantee — is not yet secured at the required level. The path to assurance is discernible because FOCIL exists as a specified mechanism, but it is not secured in the current state because governance decisions, technical feasibility, and community consensus are still pending. The finding is, however, not "Open," which means the cascade treats this state differently from a fundamental failure. Which cascade stage this finding triggers and how it interacts with the findings of the third dimension is set out in the overall synthesis (4.10).

4.9 Dimension III: Resilience and Sovereignty

The first two dimensions examined whether the system brings the structural and qualitative prerequisites of a fundamental infrastructure. Their assessment forms a profile that places seven of eight criteria at "Fulfilled with qualification" and one — Neutrality and Censorship Resistance — at "Conditionally fulfilled." The third dimension shifts the time horizon: away from the question of whether Ethereum is suitable today, toward the question of whether it can remain suitable — stable across decades, adaptable without destabilization, free of proprietary dependencies, and operable on generic hardware. Four criteria operationalize this question — one Structural Condition (III.1) and three Qualitative Criteria (III.2, III.3, III.4) — with III.4 being a claim-specific special case that exists without a corresponding pattern from the infrastructure analysis in Chapter 3. No Critical Condition lies in this dimension, which means its results cannot cap the overall verdict downward and instead differentiate within the suitability level determined by the Critical Conditions.

III.1 Long-Term Stability

Long-Term Stability is the sole Structural Condition of this dimension and asks whether Ethereum can operate stably over decades — economically sustainable, technically maintainable, without accumulation of systemic risks that call the continuation of the system into question. The criterion has a special status because it must make visible the tension between the system's remarkable short-term stability and the identifiable long-term vectors that could undermine it.

The protocol upgrade track record is the strongest indicator. Since the Merge, Ethereum has completed four successful hard forks — Shapella, Dencun, Pectra, and Fusaka — at a cadence of approximately six to twelve months. None of these forks produced a chain split. For a global, decentralized protocol without a central release authority, this consistency is remarkable, because each fork requires coordination across ten independent client teams. The flip side of this cadence is complexity accumulation: Pectra bundled eleven EIPs in a single release — the most complex single-release scope in Ethereum's history — and the December 2025 Prysm bug at Fusaka was a direct consequence of this complexity. A bug in the Prysm client temporarily reduced network participation to approximately 75 percent, without interrupting finality, because the remaining clients continued to operate. Backward compatibility has been maintained throughout the entire runtime: contracts from 2017 run unchanged, with the only systematic excep-

tion being Gas repricing EIPs, which have historically broken contract functionality because some contracts hardcoded Gas cost assumptions.

State growth is the critical long-term factor threatening the system's stability, and its significance extends beyond III.1, because it directly affects the assessments of II.4 (inclusivity) and III.4 (hardware agnosticism). The full-node size of 1,579 gigabytes documented in Section 4.1 grows at approximately 14 gigabytes per week — a rate that measures the full node including history and therefore exceeds the growth rate of the state alone — and the protocol has no native mechanism to limit this growth. Every deployed contract, every new wallet, every stored storage slot accumulates permanently on all full nodes, without inactive entries expiring or being compressed. The projections from Section 4.8 show the trajectory. At a Gas target of 200 million — defined as a goal for the Glamsterdam phase of the roadmap — full-node size reaches 5 terabytes by 2028 and 9 terabytes by 2030.[122] This projection is not a pessimistic assumption but a linear extrapolation of documented growth rates accounting for the planned Gas limit increase. If full-node operation becomes restricted to datacenter hardware, Ethereum loses the physical decentralization foundation that the inclusivity assessed in II.4 and the hardware agnosticism assessed in III.4 presuppose. The Ethereum core developer consensus classifies state growth as the most pressing unresolved scaling-related centralization risk. Verkle Trees (EIP-6800, status RES, Stagnant) and History Expiry (EIP-4444, Phase 1 DEPL since July 2025, Phase 2 PLAN) address parts of the problem — Verkle Trees enabling stateless clients that need not store the state locally, and History Expiry reducing storage requirements for historical data. State Expiry itself — the concept of allowing inactive state entries to expire after a defined period — has the status RES without an implementation commitment and is conceptually highly complex, because the delineation between active and inactive state, proof generation for expired state, and backward-compatibility implications are open research questions.[123]

Economic sustainability poses the second long-term question. The issuance model described in Section 4.3 generates a slight inflation of approximately 0.5 percent annually, with the EIP-1559 base fee burn providing a deflationary counterpart under high utilization. The L2 migration has shifted this equilibrium: as L2s absorb the majority of transaction activity, L1 fee revenue declines, and with it the deflationary counterpart to issuance. In the long term, the question is

[122] Projections based on documented growth rates and the planned Gas target of 200 million. Cf. Section 4.8 for the data basis and scenario analysis.

[123] Verkle Trees: EIP-6800, status RES, Stagnant. History Expiry: EIP-4444, Phase 1 DEPL since July 2025, Phase 2 PLAN. State Expiry: RES, without implementation commitment.

whether L1 fee revenue as a settlement layer for L2s suffices to finance the security budget that carries the economic security of I.2. The issuance reform debate — documented in Section 4.3 as active and unresolved — addresses this question, but has not produced a solution path in the current state: no feature with the status PLAN or DEPL addresses the long-term security budget question. The debate itself is a governance test case, because the interests of stakers (higher yield) and the interests of the network (sustainable security) stand in a tension that the informal consensus process has not yet resolved.

The post-quantum migration constitutes the third long-term question. The cryptographic foundations of the system — ECDSA for account signatures, BLS12-381 for validator attestations, and KZG commitments for Data Availability — must be migrated to quantum-resistant algorithms in the long term, because a cryptographically relevant quantum computer could derive private keys from public keys. The migration has the status RES, with over ten client teams coordinating devnets for post-quantum cryptography, and it is constitutive for the durability of the cryptographic guarantees on which the security assessed in I.2 and I.4 rests. In the current-state time horizon, the risk is not acute, because cryptographically relevant quantum computers are not expected to become operational within the five-to-ten-year horizon according to current research, but as a long-term question the migration belongs in the stability assessment because it concerns the foundations of the security model.

The differentiation by time horizon is decisive for the assessment and mirrors a pattern that the reference infrastructures of Chapter 3 also show: every infrastructure has long-term risks that do not impair its current functionality but determine its future. In the short to medium term — on a horizon of five to ten years — Ethereum is stable: the upgrade track record is consistent, economic security is in the range of 26 billion US dollars, and the validator infrastructure is intact. In the long term — on a horizon of twenty years and beyond — state growth, the open issuance question, and the pending post-quantum migration threaten the system's viability. State growth is the structurally heavier problem, because it concerns physical hardware requirements that cannot be resolved through economic incentives, while the issuance question is a governance problem that could in principle be addressed through a protocol change, and the post-quantum migration is a technical challenge affecting the entire blockchain sector for which Ethereum has prepared a migration path through Account Abstraction. Long-Term Stability therefore stands at: Fulfilled with qualification. Short- and medium-term stability is operationally documented. Long-term stability depends on developments not secured in the current state.

The system is long-term stable, but can it also evolve? Stability without adaptability would be stagnation, and the roadmap complexity described in Section 4.5 requires a governance system capable of delivering.

III.2 Adaptive Governance

This criterion poses the complementary question to III.1: can Ethereum respond to technical requirements, security vulnerabilities, and changed user needs without the adaptation process itself destabilizing the system? The tension between stability and adaptability is present in every infrastructure, as Mayntz showed in 1993 for large technical systems, but in a decentralized system without formal governance authority it takes a particular form: no body can mandate changes, and no voting mechanism can be binding.[124]

The EIP process is transparent, selective, and public: of 230 EIPs submitted in 2025, 37 were accepted — an acceptance rate of 16 percent documenting a selective decision culture in which the majority of proposals do not pass the quality filter of the process. The All Core Devs calls operate as a biweekly coordination body, publicly streamed and minuted, in which representatives of all client teams coordinate the content and timing of hard forks. The transparency of this process is relevant for a system asserting the infrastructure claim because it allows any observer to trace governance decisions and makes the exercise of power by individual actors visible. The process functions, but without formal authority: the decision ultimately lies with the client teams that release the software and with the node operator community that decides which version to run. This is governance by Rough Consensus — a principle that picks up the IETF parallel from Chapter 3: decisions are made when no substantial objection remains, and the IETF has shown that this model can function for the governance of a global protocol standard — the TCP/IP stack having been developed over decades without formal voting mechanisms.

The hard fork track record demonstrates operational capability. All post-Merge forks were completed without chain splits, and the historical emergency governance capability — documented in Section 4.5 with the September 2016 DoS emergency forks Tangerine Whistle and Spurious Dragon — shows that the system is capable of acting under acute pressure. The December 2025 Prysm bug required no emergency fork because improved client diversity absorbed the bug,

[124] Cf. Mayntz 1993, pp. 97–108. The tension between steering and self-organization that Mayntz identifies as the central research problem manifests in Ethereum's informal governance model as the tension between Rough Consensus and formal authority.

documenting structural governance resilience: the best governance response is the one that is not needed, because the system processes the shock autonomously. The limit lies in the absence of a pause mechanism: in the event of a zero-day consensus bug at the protocol level, the response takes weeks, because it requires coordination of all client teams and an emergency fork — and for a system securing over 100 billion US dollars, this response time is a constraint.

The Ethereum Foundation as a governance actor is a documented centralization tension relevant to the infrastructure claim of a decentralized system. Board-controlled, without protocol-level accountability, and with a treasury estimated at 850 to 950 million US dollars, the EF is the single most influential actor in the ecosystem — because it influences the financing of most core developer teams and, through its employment of Geth developers, directly controls the dominant Execution Layer client.[125] The subtraction philosophy that the EF has articulated as its guiding principle aims to systematically reduce its own role, and Protocol Guild — with over 190 members and more than 50 million US dollars in received contributions — is a decentralized alternative to EF financing that reduces the ecosystem's dependence on a single organization. Both mechanisms are, however, not protocol-anchored: neither the EF nor Protocol Guild has a protocol-side guarantee securing their financing or governance role, and the sustainability of Protocol Guild depends on the continued willingness to contribute of an ecosystem acting from economic self-interest — an obligation that has not been contractually committed and that can decline during downturns.

The issuance reform debate is the active test case for the governance system's capacity to handle controversial decisions, and its outcome will show whether the informal consensus process also functions when the economic interests of governance participants are directly affected. The question of whether and how ETH issuance should be adjusted touches the direct economic interests of stakers — who control approximately 30 to 31 percent of the ETH supply and benefit from higher yields — and stands in tension with long-term security sustainability, which depends on a sustainable issuance policy. The debate has not produced consensus in the current state, documenting both the difficulty of the decision and the limits of the informal governance model in distributional conflicts. The finding for Adaptive Governance is: Fulfilled with qualification. The governance system is operationally strong and has a track record that no comparable decentralized system can match, but the informality provides no mechanical guaran-

[125] Protocol Guild: 190+ members, >$50 million received (as of January 2026). EF treasury estimate: $850–950 million based on the Ethereum Foundation Report 2024 and ETH market developments.

tees for extreme conflicts, and EF centralization stands in tension with the decentralized claim.

The system can evolve. The next question examines whether in doing so it creates new dependencies: is the openness that II.2 identified as the system's strongest property also realized at the infrastructure layer?

III.3 Sovereign Portability

This criterion asks whether Ethereum is free of proprietary dependencies that place individual actors in uncontrollable positions of control. The answer differs depending on which layer of the system is examined.

The protocol layer is entirely open source and proprietary-free — and this distinguishes Ethereum fundamentally from all seven reference infrastructures in Chapter 3. All ten clients — five on the Execution Layer and five on the Consensus Layer — stand under liberal licenses, from LGPL-3.0 for Geth through Apache 2.0 for Besu to MIT for numerous tools. The specifications — from the Yellow Paper through the Execution Layer Specs to the Consensus Specs — are publicly documented and freely available. There is no Ethereum Inc. holding IP rights, and the Ethereum Foundation as a Swiss foundation does not own the protocol and cannot legally control it. SWIFT is a Belgian cooperative under member governance, GPS is under the US Department of Defense, cloud infrastructure is proprietary: none of these reference infrastructures reaches the degree of proprietary freedom that Ethereum's protocol layer exhibits. The EVM standard itself has diffused into the broader blockchain ecosystem as a de facto standard without proprietary control: the OP Stack, Arbitrum, zkSync, Polygon, and BNB Chain have all adopted the EVM as an execution environment, generating a standardization dynamic that facilitates migration between EVM-compatible systems and strengthens portability at the application layer.

The multi-client architecture reinforces this finding. The historical Geth dominance — which stood above 70 percent of the Execution Layer from 2020 to 2023 — is the documented concentration risk of this dimension, because Geth bugs at a market share exceeding the 33-percent blocking minority could cause consensus splits. The current state shows a clear improvement: Geth has declined to approximately 42 percent, Nethermind holds approximately 24 percent, Besu approximately 16 percent, and no Consensus Layer client exceeds 34 percent.[126]

[126] Client diversity data: clientdiversity.org / supermajority.info, as of early 2026. Execution Layer distribution: Geth approximately 42 percent, Nethermind approximately 24 percent, Besu approximately 16 percent, Erigon approximately 11 percent, Reth approximately 7 percent.

The fact that users can freely choose and combine five EL and five CL clients is a portability property that has no equivalent in centralized systems: a SWIFT participant cannot replace the SWIFT software with an alternative. An Ethereum node operator can choose from 25 client combinations. The cryptographic standards — ECDSA, BLS12-381, Keccak-256, and KZG commitments — are open standards without proprietary binding, with quantum resistance a long-term risk that is not acute in the current-state time horizon, because cryptographically relevant quantum computers are not expected to become operational within the five-to-ten-year horizon according to current research.

The operational infrastructure layer presents a different picture. The RPC concentration of 70 percent across Infura and Alchemy — documented in I.2 — both proprietary SaaS services, creates a de facto dependency at the access layer that represents the operational standard for the majority of users and developers. Decentralized alternatives such as dRPC and Pocket Network exist but have limited uptake, because institutional users prefer the SLAs of established providers. The L2 proliferation has created new soft lock-ins operating on a layer that the Ethereum protocol itself does not control: the OP Stack is Apache 2.0 licensed, but the Optimism Foundation acts as de facto standard-setter. zkSync and Starknet use partially proprietary prover technologies, and a protocol deployed on Arbitrum that uses Arbitrum-specific features such as Stylus or specific precompiles is not straightforwardly migratable to another stack. This new layer of dependencies qualifies the proprietary-freedom claim at the usage level without impairing the proprietary freedom of the base layer.[127]

Sovereign Portability therefore stands at: Fulfilled with qualification. The protocol layer is proprietary-free to a degree that none of the reference infrastructures achieves. The operational layer shows concentrations and L2 lock-ins that qualify the claim at the usage level.

Software sovereignty is given. The final question concerns hardware: can the system be operated on generic, commodity hardware?

[127] On the L2 lock-in problem: OP Stack (Apache 2.0, Optimism Foundation), Arbitrum Nitro (MIT, Offchain Labs governance), zkSync Era and Starknet (partially proprietary provers). The fragmentation creates a new layer of soft lock-ins that does not exist at the base layer.

III.4 Hardware Agnosticism

This criterion has a special position in the evaluation framework because it exists without a corresponding pattern from the infrastructure analysis in Chapter 3 — as the definition established in Chapter 2 notes: none of the seven reference systems exhibits the specific risk of cloud concentration in the form it takes with Ethereum. The criterion examines whether the hardware independence that Proof of Stake enables architecturally is operationally realized.

The architecture is strong. The transition from Proof of Work to Proof of Stake described in Section 4.1 eliminated ASIC dependency and reduced energy consumption by 99.95 percent — a fundamental change in hardware requirements: Proof of Work required specialized ASICs or high-performance GPUs that generated an entire industrial sector with supply chain dependencies, electricity cost arbitrage, and hardware obsolescence cycles, while Proof of Stake relies on generic CPU computing capacity and SSD storage. A validator node requires approximately 100 watts — compatible with consumer energy and enabling operation on solar power or uninterruptible power supply in regions with unstable infrastructure. All ten clients run natively on ARM-64 and x86-64 architectures, with Nimbus explicitly optimized for low-power ARM devices. The protocol has no hardware preference: a validator on consumer hardware has the same voting weight as a validator on a datacenter server, because the voting is key-pair-based and independent of computing power. This structural property is a post-Merge achievement that distinguishes Ethereum fundamentally from ASIC-dependent systems such as Bitcoin and creates the prerequisite for the physical decentralization that the infrastructure claim requires.

Cloud concentration counteracts this protocol property. Approximately 59 percent of hosted Execution Layer nodes run on three cloud providers — AWS with 35.5 percent, Hetzner with 13.8 percent, and OVHcloud with 9.7 percent. The trend has been declining since 2022, documenting a positive development, but the concentration remains substantial: if AWS, Hetzner, and OVHcloud were to shut down in a coordinated fashion or be regulatorily decoupled, a majority of the hosted node infrastructure would fail. The AWS outage of October 2025 demonstrated that the network could absorb the failure of an AWS region without losing finality — documenting robustness — but the corridor between normal operation and finality loss narrowed temporarily. Home staking — operation on one's own hardware — stands at under 15 percent of staked ETH and includes solo stakers from Section 4.3 as a smaller subset, which means the majority of validator infrastructure is physically located on commercial cloud servers,

and the decentralization that the protocol enables at the logical layer is undermined at the physical layer by commercial dependencies.[128]

The geographic distribution reinforces the picture: 39 percent of nodes operate in the USA, 14.5 percent in Germany, 14 percent in China, and approximately 53 percent in two countries. Beyond these three countries, the remaining node share is distributed across numerous other jurisdictions — more than ten of them, including Singapore, Canada, Japan, Australia, the Netherlands, and Switzerland, each holding over one percent of nodes. The geographic indicator of the criterion — at least ten countries each with over one percent — is thereby fulfilled. Cloud concentration multiplies the jurisdictional exposure, because US-based providers such as AWS expand it beyond mere node geography. For a system claiming to resist geopolitical instrumentalization, the combination of geographic and infrastructural concentration is a constraint visible both in II.1 under the jurisdictional dimension and here in the hardware dimension.

The single largest provider, AWS, carries 35.5 percent — below the threshold of "Open" of over 50 percent on a single provider — while the aggregate of 59 percent is distributed across three providers, and the declining trend since 2022 documents a positive development. Hardware Agnosticism therefore stands at: Fulfilled with qualification. The protocol-level hardware agnosticism is fully realized — without ASIC dependency, ARM-compatible, and consumer-energy-compatible — but the operational cloud concentration on three providers creates a physical dependency that undermines decentralization at the logical layer, and the low home-staking share of under 15 percent shows that the architectural possibility of consumer-hardware operation is not utilized by the majority of validators.

Synthesis Dimension III

The assessment profile of the third dimension shows a consistent pattern: all four criteria stand at "Fulfilled with qualification" — identical to Dimension I. The pattern is, however, differently configured. The three dimensions form a triptych of discrepancy: Dimension I shows the gap between protocol promise and operational delivery in the here and now; Dimension II shows the tension between maximum openness and absent neutrality guarantee; and Dimension III shows the gap between current functionality and long-term viability. Each facet has its own paths to resolution, and the overall synthesis will show which of these facets determines the overall verdict. State growth, the open issuance question, and the

[128] Cloud concentration estimate: Ethernodes, early 2026. The 59 percent refers to hosted EL nodes. The home-staking share of under 15 percent is an estimate based on the node and cloud distribution from Section 4.1.

pending post-quantum migration are not acute problems that threaten today's operation but long-term vectors unfolding over years, challenging the structural viability of the system in the medium term. Governance is operationally strong and has a track record that no comparable decentralized system can match, but the informality provides no mechanical guarantees for the event that social consensus breaks. Proprietary freedom is stronger at the protocol layer than in any reference infrastructure but is qualified at the operational layer by RPC concentration and L2 lock-ins. Hardware agnosticism is architecturally fully realized but counteracted by cloud deployment patterns that create a physical dependency that the protocol design rules out.

The Structural Condition III.1 stands at "Fulfilled with qualification" and thus counts as weak in the third cascade step: long-term stability is fundamentally given, with documented risks that qualify the claim without undermining the viability within the assessed current-state time frame. No Critical Condition lies in this dimension, but III.1 enters the third cascade step as a weak Structural Condition and contributes to the cap there, while the suitability level is differentiated by the six Qualitative Criteria. All twelve criteria have been assessed and the complete profile is available. The overall synthesis brings the individual findings together into the current-state profile and formulates the overall verdict through the M3 cascade.

4.10 Overall Synthesis and Current-State Verdict

The twelve individual assessments of Sections 4.7 through 4.9 yield a profile that is as revealing in its consistency as in its single deviation. No criterion achieves "Fulfilled" without qualification — meaning the system meets the requirements in no dimension fully and without structural constraints — and this simultaneously shows that the infrastructure requirements derived in Chapter 2 and validated against seven reference infrastructures in Chapter 3 are ambitiously calibrated for the subject. Eleven of twelve criteria stand at "Fulfilled with qualification" — meaning the requirement is fundamentally fulfilled, with identified limitations that qualify the claim without fundamentally undermining it. One — Neutrality and Censorship Resistance (II.1) — stands at "Conditionally fulfilled" — meaning fulfillment depends on conditions not secured in the current state, in particular the implementation of a protocol mechanism to enforce transaction inclusion. None stands at "Open" — meaning no criterion fundamentally fails a requirement for which no operational path would be discernible. The numerical profile 0-11-1-0 describes a system that addresses the infrastructure requirements in breadth, but in which operational reality has not yet fully delivered on the protocol-level possibilities.

Figure 4.11 Current-State Profile of All Twelve Criteria — assessment label by hierarchy level and dimension

The M3 cascade defined in Chapter 2 determines the overall verdict on the basis of the hierarchical grading of the criteria. The three Critical Conditions form the first examination stage. Security and Trust Load (I.2) stands at "Fulfilled with qualification": the economic security carries through the mechanically enforced Slashing sanction and the scale of an attack whose static capital requirement on

the reference date stands at approximately 26 billion US dollars — the actual costs driven far beyond this by market illiquidity. The constraint lies in the market-price dependency of this security and in the operational residual trust burden of the four-layer trust stack, which qualify the claim without fundamentally undermining it. Minimal Viable Guarantees (I.4) likewise stands at "Fulfilled with qualification": three of four guarantees are protocol-anchored and operationally proven, with the Degradation Mode constituting a property without equivalent in the reference infrastructures, and the fourth guarantee — censorship resistance under attack — functionally given but not protocol-secured. Both Critical Conditions pass the cascade threshold: the core properties of security and guarantees are fundamentally given. Neutrality and Censorship Resistance (II.1) does not pass the threshold unconditionally. The builder concentration of 93.3 percent among three actors marks a structural defect, and no protocol mechanism enforces transaction inclusion. The system's censorship resistance is an emergent market performance, not a protocol guarantee, and for an infrastructure claiming neutrality as its core claim, this dependency is a condition not yet delivered on in the current state. The finding is "Conditionally fulfilled," not "Open" — meaning the path to fulfillment is discernible, but depends on the implementation of FOCIL or a functionally equivalent mechanism. No Critical Condition stands at "Open," which lifts the overall verdict above "Conditionally Suitable." One Critical Condition stands at "Conditionally fulfilled," which sets the cascade to the level "Suitable under substantial conditions."[129]

The three Structural Conditions form the second examination stage. Functional Irreplaceability (I.1), Coordination Function (I.3), and Long-Term Stability (III.1) all stand at "Fulfilled with qualification." The structural foundation is load-bearing: Ethereum functions as a center of gravity for decentralized coordination with a TVL dominance exceeding 60 percent and a developer anchoring exceeding 70 percent of all blockchain developers. It provides a quantitatively and qualitatively unique coordination capacity that coordinates approximately 100 billion US dollars permissionlessly, and it is stable in the short to medium term — with a consistent upgrade track record and a finality rate above 99.99 percent. The constraints — the emergent anchoring, the L2 fragmentation, and the long-term risks from state growth and issuance — qualify the viability without fundamentally undermining it. No Structural Condition stands at "Condi-

[129] The cascade logic in the second step was classified in Chapter 2 as a methodological supplement: "Conditionally fulfilled" at a Critical Condition is qualitatively a different state from "Fulfilled with qualification," because the core property of the infrastructure claim is not yet secured at the required level. Cf. Section 2.3 for the definition of verdict categories and the M3 cascade logic.

tionally fulfilled" or "Open," but all three stand at "Fulfilled with qualification" and are thereby counted as weak, so the third cascade step sets the same cap as the second.

The six Qualitative Criteria differentiate the degree within the achieved suitability level. All six — Open Generativity (II.2), Independent Verifiability (II.3), Low-Threshold Inclusivity (II.4), Adaptive Governance (III.2), Sovereign Portability (III.3), and Hardware Agnosticism (III.4) — stand at "Fulfilled with qualification." The degree scale from Chapter 2 defines the degree as "Good" when all six Qualitative Criteria stand at least at "Fulfilled with qualification" — which is the case here. The constraints are distributed across different dimensions and show no clustering in a single problem area: L2 composability fragmentation and L2 verification complexity concern the scaling architecture; the 32-ETH solo staking barrier as the only deficit without a roadmap path concerns the access threshold; informal governance without mechanical guarantees concerns adaptability; operational RPC concentration and L2 lock-ins concern proprietary freedom; and cloud deployment patterns concern physical decentralization. No single Qualitative Criterion shows a fundamental failure, and the breadth of the constraints confirms the basic pattern: the constraints are operational in nature, not protocol-level.

The overall verdict is: **Suitable under substantial conditions, degree Good.** Chapter 2 defined this category as a state in which at least one Critical Condition stands at "Conditionally fulfilled" and suitability is tied to named, verifiable conditions — with the core property of the infrastructure claim not yet secured at the required level at the protocol side, but the path to assurance remaining discernible.[130] The central condition that determines the label is the implementation of FOCIL (EIP-7805) or a functionally equivalent mechanism that transforms censorship resistance from an emergent market performance into a protocol guarantee. As long as transaction inclusion depends on the market structure of builders rather than being enforced by the protocol, the neutrality guarantee that the infrastructure claim asserts remains operationally delivered but structurally unsecured. The verdict "Suitable under substantial conditions" explicitly does not mean "Conditionally Suitable": there is no "Open" finding failing a fundamental requirement, and the path to delivering on the central condition is discernible and specified.

[130] Cf. Section 2.3.3 for the definition of the five verdict categories and the M3 cascade logic. "Suitable under substantial conditions" is the middle category of the spectrum and was defined generically in Chapter 2, independent of which criterion is affected and which system is being assessed.

A pattern runs through all twelve assessments, and it can be stated in one sentence: the discrepancy between what the protocol enables and what operational reality delivers is the basic pattern of the current state. Ethereum offers, at the protocol level, the prerequisites of a fundamental infrastructure — from economic security through open generativity to cryptographic verifiability, from atomic coordination to automated self-healing. The operational usage reality has not fully realized these possibilities: the trust stack generates institutional dependencies, L2 fragmentation undermines composability, builder concentration threatens neutrality, state growth threatens long-term decentralization, and governance operates without mechanical guarantees for extreme conflicts. Each of the three dimensions reveals its own facet of this discrepancy: Dimension I documents the tension between protocol possibility and operational use; Dimension II the asymmetry between the strongest property (open generativity) and the most critical vulnerability (absent neutrality guarantee); and Dimension III the tension between short-term stability and long-term risks. The profile 0-11-1-0 is the numerical condensation of this threefold discrepancy. The strength of the system lies in that the protocol anchors the right properties. The constraint lies in that the operational layer — the trust stack, the L2 architecture, the builder economics, and the cloud infrastructure — has not yet delivered on this anchoring in breadth.

The current-state verdict defines the starting point for the target-state assessment. Chapter 5 will examine whether the Ethereum roadmap addresses the identified constraints and whether the target-state profile can shift the overall verdict. The question is not whether Ethereum is suitable as infrastructure — the current-state verdict confirms fundamental suitability under conditions — but whether the roadmap can deliver on the conditions to which suitability is tied. The gap between "Suitable under substantial conditions" and the next higher level "Suitable with conditions" is defined primarily by the cascade-determining condition: the transformation of emergent censorship resistance into a protocol guarantee through FOCIL. The eleven constraints at "Fulfilled with qualification" — from the trust stack through L2 fragmentation and state growth to cloud concentration — do not determine the overall verdict but the degree of maturity within the achieved suitability level, and the target-state assessment will examine which of these constraints are addressable through roadmap features. The roadmap claims to address the cascade-determining condition through FOCIL, to close L2 composability through Native Rollups and Based Sequencing, to mitigate state growth through Verkle Trees and History Expiry, and to reduce the trust stack through Native Rollups and decentralized alternatives. The target-

state assessment will examine whether the features intended to close these gaps meet the epistemic requirements of the taxonomy: PLAN features with a clear timeline and EIP acceptance, or RES features in early research stages. Whether the roadmap delivers on the conditions in a robust way is the subject of the target-state assessment.

Chapter 5 The Target State: Presentation and Assessment of the Complete Roadmap

Table of Contents

5.1-A The Strategic Pivot

Chapter 4 presented, explained, and assessed Ethereum in its current state as of the first quarter of 2026 against the criteria introduced in Chapter 2. The study arrived at a qualified verdict, according to which the system is suitable as fundamental digital infrastructure, albeit under substantial conditions. With Chapter 5, the study turns its gaze and describes what Ethereum would look like if the Roadmap were fully implemented over a horizon of ten or more years. The protocol is henceforth presented and treated as a system conceived as complete, whose properties emerge from the interplay of the planned overhauls, with the target state remaining the consistent subject of discussion rather than the development status of individual components.

The transition shifts not only the point in time of the analysis but also the logic by which the system is read. Chapter 4 tested each criterion against measured states — validator distributions, client shares, fee trajectories, and so on — while Chapter 5 reads a system design whose components exist as specifications, prototypes, and lines of research, which are here assembled into a coherent picture. The presentation takes the components in the final state described by their specifications and research designs, and keeps the question of the probability and sequence of their implementation outside the system description. The following sections show an Ethereum whose base layer would carry its own execution at the scales today reserved for Rollups, whose consensus would finalize

in seconds rather than minutes, and whose guarantees extend across connected layers. Before sections 5.2 through 5.6 unfold this system layer by layer, however, the target state requires its frame — because the multiplicity of overhauls only becomes readable as a coherent program once the strategic premise from which they derive is named.

The premise set out here takes its starting point in 2025. Over that year, the Layer 1 substantially expanded its execution capacity, while the Rollups, to which the scaling of the platform had been assigned in the previous architecture, continued to be ordered by a single centralized Sequencer each, which alone determines the acceptance and ordering of transactions. From the tension between these two observations arose the strategic reorientation whose result the present chapter shows. The target state described here rests on a fundamentally altered system architecture. Layer 1 expands its own execution capacity again going forward. Rollups serve as integrated, specialized execution environments (e.g., for privacy) that process complex applications and whose data are ultimately secured on Layer 1.

What the reorientation has displaced is the rollup-centric Roadmap that Ethereum had followed since 2020. In it, the Layer 1 remained a deliberately lean layer for consensus, settlement, and data availability, with deliberately limited execution capacity, while the processing of transactions was outsourced to the Rollups. The division of labor rested on a dual assumption: that the Rollups inherit the security guarantees of Layer 1 because they anchor their data and proofs there, and that they would simultaneously free themselves over time from the control of their operators. How far the real system matched that assumption in the first quarter of 2026 was examined in detail in Chapter 4.

Two Findings Force the Revision

Two empirical findings, each of its own weight, forced the revision. The first concerns the outsourcing promise itself: the decentralization of the leading Rollups proceeded considerably more slowly than the Roadmap had assumed. This is evident in that the large networks continue to be ordered by centralized Sequencers belonging to their respective operators, that substantial value lies behind the multi-signature keys of Security Councils, and that the organizational ties of the operators to Ethereum are looser than the technical connection of their systems.

The second finding bears the weight in the opposite direction, because the Layer 1 demonstrated that it can expand its own execution. Over the course of 2025, the Gas Limit doubled from 30 to 60 million without the protocol needing to change, since the adjustment ran through the block-by-block voting of the validators and was coordinated by social consensus alone.[131] The effect reached into individual application projects, as with the Ethereum Name Service, which abandoned its planned rollup of its own in favor of the now cheaper Layer 1.[132] What is demonstrated is a path, not yet a state. At 60 million Gas, the Layer 1 processes 15 to 30 transactions per second. The claim that the base layer scales on its own is not yet realized as a full property in the current system. The finding establishes the direction of expansion, not its scale.

Vitalik Buterin brought both findings together on 3 February 2026 in a post on X in which he announced the departure from the previous strategy.[133] He identified two facts. The decentralization of the second layer had proceeded far more slowly and laboriously than originally anticipated, and Layer 1 was by then scaling on its own. From this he drew the conclusion that the original vision of Rollups as an outsourced scaling layer no longer made sense and demanded a new path. The post articulates the reorientation and grounds it in the two findings, but leaves open what form the relationship between the layers will take in the future. This openness is what the new model closes.

[131] Gas limit doubling from 30 to 60 million over the course of 2025, coordinated through the block-by-block validator voting (adjustment of approximately 0.1 percent per block) without a protocol change. The default target is formalized in EIP-7935. Data basis: Ethereum Foundation, cf. the data basis in section 5.2.

[132] Discontinuation of the planned ENS-native rollup Namechain in favor of Layer 1, justified by an approximately 99 percent reduction in registration costs on Layer 1. ENS Blog (nick.eth), 6 February 2026, secondary reporting: CoinDesk, The Block.

[133] Vitalik Buterin, X post of 3 February 2026 (pivot articulation with the two facts of overly slow L2 decentralization and the self-scaling Layer 1, and the conclusion that the original vision no longer made sense and a new path was needed). Secondary anchors: CoinDesk, Decrypt.

The New Model: Two Directions

The new model has two directions, the first of which concerns the Layer 1 itself. It resumes the scaling of its own execution. It rests on a graduated path of further Gas Limit increases and on an enshrined zkEVM, under which validators verify blocks solely through validity proofs. The load of verification would be decoupled from the system's capacity, so that capacity could grow without making the requirements for participation in consensus grow with it. The validator voting that carried the doubling of 2025 would be merely the beginning in this scheme, because the later expansions would remain bound to the proof architecture. Which stages the path passes through and on which architecture the proofs rest is developed in section 5.2.

The second direction redefines the relationship to the Rollups. In place of a homogeneous scaling layer, a spectrum emerges that is ordered along two dimensions. The degree of binding ranges from full convergence with Layer 1 to the sovereign chain with minimal binding. The differentiation describes what a Rollup grounds its existence on — privacy-focused execution environments, application-specific efficiency, low latency rather than mere throughput, or extreme scaling beyond the capacity of the Layer 1. The idea of a uniform role for Rollups disappears in the spectrum. Each network positions itself along both dimensions and bears the consequences of its position. At the outer end stands the sovereign chain, which gains full design freedom with its own data availability and its own consensus while simultaneously operating outside Ethereum's guarantees. Within the spectrum, the choice of trust level lies with users and applications themselves. The redline for connection to Ethereum's guarantees is Stage 1, that level of maturity from which a functioning proof system limits the operator's rights of intervention to emergencies. The mechanism of full convergence is to be provided by the native-rollup precompile, through which a converged Rollup would directly use the execution logic of the Layer 1 and automatically adopt every protocol upgrade — a construction whose substance section 5.6 develops. The privacy direction gains its own shape in the closing section of 5.1.

The new strategy takes dimensional shape in five target pictures that the research community leads as North Stars and that Justin Drake bundled in the Strawmap in February 2026. The Strawmap is a coordination document without formal governance status that documents the direction of work without setting a binding target horizon. Buterin publicly categorized it as important orienta-

tion.[134] The five target pictures span the dimensions of the target state, from execution through data to consensus, cryptography, and privacy. Fast L1 designates the concentration of finality, which would fall from today's roughly 13 minutes into the seconds range and in the realized target state would lie at three-slot finality of roughly 36 seconds, whose consensus-side substance the later course of the chapter develops. Gigagas L1 names the target of roughly 10,000 transactions per second on Layer 1, which would become achievable over the long term via the graduated expansion path of execution, and whose stages likewise appear in section 5.2. Teragas L2 designates the decoupling of second-layer throughput from Layer 1 via Data Availability Sampling, whose mechanics the data section of 5.2 develops. Post-Quantum L1 describes the target state of fully hash-based cryptography, whose substance the following section and section 5.3 develop, and Private L1 names native shielded ETH transfers at the protocol level as a long-term direction, whose substance that very closing section of 5.1 addresses.

In sum, the target state of the chapter describes a division of labor in which Layer 1 again carries execution on its own and Rollups choose their place through degree of binding and differentiation. The guarantees of the base layer extend through convergence to where the connected networks wish to draw on them. That the ambition of the reorientation reaches beyond the protocol is documented by the Ethereum Foundation's reconnection to Cypherpunk principles under the heading DeFipunk, the establishment of the dAI Department for the economy of autonomous AI agents together with the standard ERC-8004 for Trustless Agents, and Buterin's New Year framing of Ethereum as civilizational infrastructure.[135] Whether a program of such reach can be realized depends on a design philosophy that makes the complexity of the protocol itself its subject, which the following section develops.

5.1-B Lean Ethereum as Design Philosophy

With each year of its development, the protocol has become harder to survey. What began as a manageable set of rules has grown into an entanglement of specifications, client implementations, and cryptographic procedures whose maintenance demands specialized knowledge and whose auditing becomes more expensive with the volume of the code. The circle of those who can still survey the system in its entirety has grown smaller over the years, and every new capability

[134] Justin Drake, Strawmap (strawmap.org), published in February 2026, explicitly named as a strawman and roadmap without formal governance status, a coordination instrument for researchers and client teams.

[135] Vitalik Buterin, New Year's post of 1 January 2026 (framing of Ethereum as Civilisational Infrastructure), evidence for the ambition level of ecosystem signals.

the protocol acquires adds to it another layer of code and attack surface. The new strategy outlined in section 5.1-A enlarges the entanglement further, and the Roadmap ties its manageability to a design philosophy that makes the complexity of the protocol itself its subject.

Lean Ethereum is the program of simplification, and it rests on the assumption that the security and durability of a protocol decrease with every avoidable component that an attacker can study and a developer can misunderstand. The Ethereum Foundation has identified accumulated complexity as its own burden since around 2024, visible in the code complexity of the clients as well as in the exhaustion of the EIP consensus process.[136] The consensus process can resolve only a small selection of competing proposals per hard fork and plays each additional function against the next. The elements of the program range from the reduction of client code through the transition to SSZ as a unified serialization format to a formally verified consensus client and hash-based primitives. Added to these are a new execution environment in the form of leanVM and the RISC-V line as its foundation.[137]

Post-quantum resistance would emerge from the principles of this simplification as an integral property. The migration from the currently deployed components would then constitute only the transitional problem. The actual construction effort lies elsewhere. Hash-based primitives stand by default in place of number-theoretic procedures, from Poseidon through leanXMSS to the STARKs, whose security rests solely on the collision resistance of hash functions. The ECDLP-dependent components — from BLS aggregation through Pedersen Commitments and IPA to KZG — yield to hash-based alternatives, whereby the attack surface for a sufficiently large quantum computer disappears. Formal verifiability at all levels would be added, making correctness provable rather than merely empirically familiar. Lean Ethereum thus appears as cryptographic simplification with emergent post-quantum resistance. It follows structurally from the architecture beyond 2029.

At the consensus layer, the Lean Ethereum target state would take the form of the Beam Chain, now called Lean Consensus — a complete redesign of the existing Beacon Chain that would absorb five years of operational experience and take over its task without inheriting its cryptographic foundation. Hash-based validator signatures from the leanXMSS line would replace the BLS signatures, constructively post-quantum-resistant. In place of algebraic BLS aggregation would

[136] EF complexity diagnosis: Vitalik Buterin, „Possible futures of the Ethereum protocol" (blog series from October 2024, in particular „The Purge"), continued in the post-quantum roadmap of 26 February 2026.

[137] leanVM and RISC-V: reference to section 5.2-E for the substance.

come hash-based aggregation that condenses many valid signatures through hash-native proof procedures such as STARKs and hash-based SNARKs into a single verifiable proof.[138] Post-quantum security would rest solely on the hardness of the hash function. The Beam Chain would be formally verified and its correctness mathematically proven. Where finality is concerned, a three-slot finality of roughly 36 seconds stands as the target state, which would replace today's finality after several minutes, with reference to section 5.5 for the mechanics. Methodologically, the Beam Chain remains a research subject without a fork date. It depends on already deployed prerequisites, among which the higher maximum Effective Balance (MaxEB) introduced with Pectra facilitates hash-based aggregation.[139]

leanVM on the execution environment and Beam Chain on the consensus layer would share the same three properties — hash-based primitives, formal verifiability, and constructive post-quantum resistance. The Lean Ethereum design philosophy would take the same shape on both layers as an integral consequence, whose execution substance section 5.2-E unfolds and whose operational consensus migration section 5.5 describes. The target state shown here is a design state and applies under the same conditions to both layers.

The design philosophy names the direction of the Roadmap and sets no deadlines for it. leanVM and Beam Chain lie beyond the horizon of 2029 without being specified in final form. Concrete research work and an EF-internal consensus carry them, naming the components without yet determining their final interlocking. The interlocking with the RISC-V line remains open in the primary sources and will occupy the coming rounds of specification. Whether the design philosophy holds will be decided beyond the protocol architecture at the operational conditions of its implementation — at governance, client diversity, and privacy, which the following section develops.

[138] Beam Chain: Justin Drake, Beam Chain (Devcon SEA 2024), ethresear.ch follow-ups; Vitalik Buterin, post-quantum roadmap (February 2026).
[139] leanXMSS line, Beam Chain as RES, MaxEB (EIP-7251, DEPL, Pectra).

5.1-C The Operative Expectation

The operational level is the place where a technically coherent protocol design can still fail, because an ecosystem must carry, maintain, and defend it against erosion over years. A correctly specified protocol would disintegrate if the client base became monocultural, if further development depended on the engagement of specific individuals, or if an announced property never found its way into the protocol. All three conditions, which the present section develops, share the feature that the design presupposes them without being able to guarantee them itself. The first is plannable governance, the second an actively maintained client diversity, the third a privacy that moves from the exceptional case to the default. The last of these conditions simultaneously shows where the evaluation framework encounters its own limit.

Plannable Governance

The governance target picture is a plannable, institutionalized protocol development that would run as a sustained process independently of any given leadership. Its most visible anchor is the envisaged six-month fork cadence — two hard forks per year at a rhythm that the ecosystem can plan for years in advance and that reaches to 2029.[140] For an infrastructure system, what counts is less the speed of change than its predictability, because dependent systems fix their own development to the known fork windows. The rhythm establishes the dates without fixing the content, which is why the cadence can be maintained even when individual features move to a later fork. Only the perpetuation over several years transforms a release plan into an institutional expectation that no longer depends on the disposition of individual actors.

What makes development robust against the turnover of individual persons is its anchoring in a documented process that guides every proposal through the same stages regardless of its author. A named technical contact accompanies a proposal from submission to possible inclusion in a fork. Per upgrade and layer, the process selects a headliner feature, while the remaining proposals are gathered until a deadline and evaluated as a batch.[141]

[140] On fork cadence: Strawmap (Justin Drake, February 2026), two hard forks per year in a six-month cycle through 2029.

[141] On the institutional carriers: EF, „Protocol Priorities Update for 2026" (18 February 2026), which names predictable delivery as a stated ambition. The fork process selects one headliner per upgrade and layer and assesses the remaining proposals as a batch in the All-Core-Developers process with status tracking via Forkcast. The Protocol Support Team provides a documented path for each proposal via the EIP Champions handbook (EF Blog, Checkpoint April 2026).

The status of every proposal is publicly traceable in graduated stages from initial consideration to firm scheduling, which detaches coordination from the internal knowledge of specific participants. For a system on which others build, a readable pipeline is itself a service, because dependent teams can assess the maturity of a planned change without participating in internal conversations. The process thus makes the progress of a proposal depend on its documented status rather than on the attention of a particular leadership. A change at the top therefore does not interrupt the cadence. The reorganization of protocol work into three tracks in February 2026 aligned the process with a few clear priorities rather than orienting it around a growing feature list. As a second indicator of plannable long-term coordination, the Strawmap is added, reaching beyond the individual fork-by-fork coordination.[142]

Maintained Client Diversity

Client diversity remains an actively maintained property of the ecosystem rather than a self-sustaining one, because market dynamics press toward a standard implementation whose prevalence ties the network to a single codebase. Maintenance requires ongoing incentives for minority clients, because the convenient choice of the most widely deployed implementation would otherwise reinforce concentration on its own. The target picture would be a distribution of execution and consensus layer clients in which no single implementation exceeds the 33 percent threshold. Above the threshold, a bug in a single client could threaten network finality. Below the threshold, the diversity of the remaining clients catches the error of any single one, so that an implementation bug remains local and does not escalate to a consensus failure across the entire network.[143] As a long-term simplification, the leanVM harmonization would be added, deriving client consistency from a formal specification, whose three dimensions section 5.2-E unfolds.[144]

---

[142] Strawmap (February 2026) here in its governance function as coordinated multi-year planning beyond the fork-by-fork decision-making. The North Star source appears in 5.1-A.

[143] On the 33 percent threshold and staking concentration: the largest LST provider would remain at approximately 23 percent in the target state, and no single client should exceed the blocking-minority threshold. Stateless Clients carry the status RES; validator consolidation proceeds via MaxEB (EIP-7251, DEPL).

[144] leanVM as the target state of the execution environment (RES, horizon beyond 2029). The three leanVM dimensions and the formal specification are unfolded in 5.2-E.

Privacy as Limit and Target Picture

The twelve-criteria framework of the study does not capture privacy, thereby marking a limit of the evaluation instrument itself, whose reflection section 6.3 carries. The limit arises from the fact that the twelve criteria were deductively derived from an infrastructure concept that did not carry confidentiality as constitutive, and that verification on systems without their own confidentiality claim provided no occasion for correction. The instrument can describe the transparency of the current state, but holds no criterion that would register a movement toward confidentiality as progress or regression. Even a protocol-native shielded pool would not resolve end-to-end privacy on its own, because the encryption of the mempool, anonymity at the network layer, and the design of wallet UX lie outside the protocol.[145] Protocol privacy would be a necessary foundation without on its own securing the confidentiality of the entire interaction.

As a target picture on the horizon beyond a decade, privacy-by-default would reverse today's ordering. Currently all transfers are public, and confidentiality is the concern of individual applications and Rollups. The protocol does not provide it. Anyone seeking confidentiality today must actively establish it and becomes visible simply by moving into a protected path, while a default would shift this burden onto disclosure. As long as confidentiality remains distributed across individual applications, it fragments into small and weak anonymity sets in which participants can barely conceal themselves. The Strawmap North Star Private L1 would embed native shielded ETH transfers in the protocol and would understand confidentiality as the protocol's default.[146] The horizon extends beyond the 2028 North Star mark, until privacy would count as an architectural property of the infrastructure and would outgrow today's binding to individual applications. Confidentiality is then a reliable foundational property on which every application rests, rather than remaining an add-on service of individual providers.

EIP-8182 concretizes the direction with a protocol-managed shielded pool and ZK verification as a protocol feature, so that private transfers to any ordinary address would be possible without requiring a dedicated privacy address

[145] On pre-inclusion privacy as a distinct layer alongside the shielded pool: Encrypted Mempool (EIP-8105/LUCID, RES) as evidence that protocol privacy does not solve end-to-end privacy alone. The reflection of the evaluation framework is carried by section 6.3.

[146] Strawmap North Star Private L1 (February 2026), native shielded ETH transfers at the protocol level, privacy as default rather than opt-in (RES).

format.[147] The pool would sit as a system contract at a fixed address, without a privileged key and without a pause mechanism, so that no party could subsequently lift the confidentiality or interrupt the traffic. The core lies in the unified anonymity set, because every integrating wallet would feed the same shared pool and anonymity would thereby arise from a single pool. With every wallet that uses the shared pool, the anonymity set grows, so that confidentiality increases with the spread of integration. The proposal has been submitted for Hegotá and is at the status of a draft in PFI-review — the Proposed-for-Inclusion stage of the ACD process — whose inclusion is not secured. Buterin's 2025 call for a native shielded balance from which sending happens by default carries the logic of such a default.[148] Both rest on curve-based ZK cryptography and therefore belong to a different horizon than the hash-based post-quantum target state outlined in the previous section.

The following sections examine the individual directions of the target state in sequence. Section 5.2 addresses the scaling of execution, 5.3 verification and access, 5.4 state management, 5.5 neutrality, and 5.6 the architecture of Layer 2. Each section takes up one of the directions that the framework has named and examines how far the design carries it in the target state.

5.2 Execution Scaling

5.2-A The Dependency Chain: Parallelization, Decoupling, Convergence

Every capacity increase beyond 100 million Gas forces a decision that the current protocol does not offer: higher throughput or a broader verification base. At 60 million Gas, Ethereum processes between 15 and 30 transactions per second, depending on transaction complexity.[149] For infrastructure intended to carry financial systems, supply chains, and institutional settlements, this is a structural limitation: Visa processes on average more than 8,000 transactions per second, at a tested peak capacity of 65,000 TPS.[150]

[147] EIP-8182 „Private ETH and ERC-20 Transfers" (Tom Lehman, May 2026), protocol-managed shielded pool with ZK verification, submitted for Hegotá and in PFI review, inclusion not secured. The verification is curve-based (Groth16 over BN254) and not post-quantum.

[148] Vitalik Buterin, call for native wallet privacy with send-from-shielded-balance active by default (April 2025).

[149] At a Gas limit of 60M, Ethereum processes 15–30 TPS depending on transaction complexity. Simple ETH transfers consume 21,000 Gas (approximately 2,800 TPS theoretically), complex DeFi transactions 200,000–500,000 Gas (approximately 12–30 TPS effectively).

[150] Visa: 257.5 billion processed transactions in fiscal year 2025 (Visa Q4 2025 Earnings Release, ending 30 September 2025), corresponding to an average of approximately 8,165 TPS. The frequently cited 65,000 TPS refers to

The comparison in terms of transaction counts understates, however, the performance class in which Ethereum already operates. A single Ethereum transaction can settle value of several hundred million US dollars with programmable finality while consuming the same 21,000 Gas as a transfer of 0.01 ETH. In the dimension relevant for institutional settlement — value throughput per unit of time — Ethereum L1 already moves in the range of global payment infrastructure.[151] The structural limitation thus concerns the number of simultaneous operations that L1 can carry and the breadth of use cases beyond high-volume individual transactions.

The exponential scaling path that EIP-9698 formulates as a target would dissolve this limitation by increasing L1 capacity by a factor of 100 over four years (tenfold increase every two years over two cycles).[152] Such an increase would bring Ethereum from the niche of specialized crypto applications into the reach of global infrastructure loads, provided that the decentralization of verification is maintained. Exactly here lies the tension: three features must architecturally break the trade-off between capacity and decentralization before the expansion can proceed safely. They address different bottlenecks but form a directed dependency chain whose shared infrastructural prerequisite is Enshrined Proposer-Builder Separation.

---

the tested VisaNet network capacity, not the actual peak value. The comparison serves as an order-of-magnitude reference. Visa and Ethereum are not directly architecturally comparable.

[151] Stablecoin transfer volume on Ethereum: over 8 trillion US dollars in the fourth quarter of 2025 (Token Terminal, January 2026). The bot share of approximately 70 percent comprises arbitrage, paymaster transactions, and other automated transfers across all chains. The Ethereum-specific share may differ (CEX.io, Stablecoin Report Q3 2025). The gross figure excludes native ETH transfers and DeFi settlements. Visa in fiscal year 2025: 17 trillion US dollars at 329 billion transactions (Visa Annual Report, ending September 2025). The figure differs from the 257.5 billion in the preceding note because both reports use different transaction concepts. The metrics are not directly commensurable in methodology, because Ethereum measures value movements between addresses while Visa records purchase transactions.

[152] EIP-9698 (Draft, author: Dankrad Feist, April 2025) formulates an exponential Gas limit path with the goal of a 100-fold capacity increase over four years (tenfold increase every 164,250 epochs, two cycles, G0 = 50M from Epoch 369,017 / approximately 1 June 2025, final value approximately 5 billion Gas). The EIP changes no consensus rules, only the voting default behavior of the clients. It is backward-compatible.

Parallelization as Gatekeeper

The sequential verification of all transactions by every node becomes a temporal bottleneck with increasing Gas Limit. In the current model, a validator checks every transaction in sequence, the ordering being mandatory as long as the node cannot determine which transactions compete for the same storage slots. At a Gas Limit below 100 million, the verification time remains manageable within the 12-second slot.[153] At higher values it grows linearly, until the slot no longer suffices. Block Access Lists (EIP-7928, SFI Glamsterdam, i.e. Scheduled for Inclusion) solve this time problem through a new field in the block header: the Builder declares at block construction all State accesses of every transaction, including the post-execution values.[154] Receiving nodes can then parallelize disk accesses and transaction validation, because the dependencies between transactions are known in advance.[155] The additional data overhead remains at an average of 70 KiB per block at a Gas Limit of 60 million — below the maximum calldata size — and does not burden the protocol with a new data problem.[156]

BALs are therefore not an optimization but the gatekeeper of the entire L1 execution scaling. Without them, the Gas Limit cannot safely be raised above 100 to 150 million, because the sequential verification time would burst the slot. The governance status underpins feasibility: SFI for Glamsterdam, three client teams (Besu, Geth, Nethermind) with ongoing prototype implementations.[154] Parallelization solves, however, exclusively the time problem on the verification side. The cost problem remains: despite parallel processing, every node continues to fully re-execute every transaction. Hardware requirements for validators grow linearly with the Gas Limit, and this linear coupling has no upper bound. Beyond a certain threshold, the requirements become so high that only professionally operated nodes can keep pace, which erodes the verification base. Exactly this erosion must a second mechanism prevent. BALs also unfold a second effect that reaches beyond the validator side: the dependency graph they provide is simultaneously

[153] At a Gas limit of 36M, the average validation time is approximately 400 ms. At 60M it rises to approximately 1–2 seconds. From approximately 100M, the validation time approaches the slot boundary, at which new prerequisites (BALs, ePBS) become necessary. Source: EF Research, Execution Layer Performance Benchmarks, 2025.

[154] EIP-7928. Authors: Toni Wahrstätter, Dankrad Feist, Francesco D'Amato, Jochem Brouwer, Ignacio Hagopian. Status: Draft since March 2025, SFI for Glamsterdam. Prototype implementations: Besu, Geth, Nethermind.

[155] BALs enable, in addition to parallel disk reads and transaction validation, also parallel state root computation and executionless state updates in ePBS contexts.

[156] The average BAL size is approximately 35 KiB per block at 36M Gas and approximately 70 KiB at 60M Gas (EIP-7928, Overhead Analysis). In the worst case it remains below the maximum calldata size.

the prerequisite for parallel witness generation in the zkEVM proving pipeline.[157] Since 60 to 80 percent of transactions in a block use disjoint storage slots, provers can parallelize execution across independent transaction clusters and distribute proof generation across multiple machines.[155] Without the dependency graph known in advance, serial witness generation would remain the bottleneck limiting zkEVM scaling at high Gas Limits.

Decoupling through Cryptographic Verification

The linear coupling between capacity and hardware requirements can be quantified. At a Gas Limit of 60 million, full verification of a block requires one to two seconds.[158] At 600 million, it would be ten to twenty seconds with sequential processing. Even with the parallelization enabled by BALs, the verification time in this range reaches the boundary of the 12-second slot. At 5,000 million — the projected endpoint of the EIP-9698 path — verification within a 12-second slot would be impossible.[159] The consequence is clear: every increase in network capacity simultaneously raises the minimum hardware requirements for validators. This stands in direct conflict with the decentralization goal, because participation in verification becomes more expensive with every capacity level. The zkEVM (EIP-8025, active EF Roadmap since January 2026) breaks this coupling by fundamentally changing the mode of verification.[160]

The paradigm shift consists in replacing universal re-execution with cryptographic proof verification. Instead of every validator recalculating every transaction, a specialized prover generates a zero-knowledge proof that mathematically proves the correct execution of all transactions.[160] The prover is thereby a standalone role that the EIP-8025 design separates from Proposer and Builder: the Proposer proposes the block, the Builder constructs it, and the Prover delivers the cryptographic proof of correct execution. Validators verify the proof in milliseconds, regardless of how many transactions the block contains. A block with

[157] In a ZK context, the latency bottleneck is frequently the serial execution for witness generation. BALs supply the deterministic dependency graph that enables parallel execution and sharding of proof generation. Cf. HackMD, „A Deep Dive into EIP-7928 and Block-Level Access Lists", January 2026.

[158] Validation times: 60M Gas ≈ 1–2 s, 600M ≈ 10–20 s. At 5,000M, validation within a 12-second slot would no longer be possible. EF Research, L1-zkEVM Roadmap, 26.01.2026.

[159] EIP-9698 (Draft, author: Dankrad Feist, April 2025) projects a tenfold increase every 164,250 epochs (approximately two years) over two cycles, i.e., a total increase by a factor of 100 within four years, starting at G0 = 50M from Epoch 369,017 (approximately 1 June 2025). Projected path: 50M (June 2025) → approximately 500M (mid-2027) → approximately 5,000M (mid-2029). The EIP changes no consensus rules, only the voting default behavior of the clients. It is backward-compatible. The projection describes the maximum possible path under continuous validator voting, not the operationally expected trajectory.

[160] EIP-8025. Authors: Alex Stokes, Justin Drake, Ethereum Foundation Research. Paradigm: Optional Execution Proofs. The EF published an implementation roadmap with six sub-themes on 26 January 2026.

60 million Gas and a block with 5,000 million Gas create the same effort for the verifier. The decisive design principle is optionality: nodes can fall back to full re-execution at any time if they do not trust the proof.[161] Proving latency fell between July and December 2025 from 16 minutes to 16 seconds, with a 45-fold cost reduction.[162] This threshold has architectural significance: for the first time it becomes realistic to prove a block within a single slot, so that a validator on commodity hardware verifies a block without executing it. This changes the verification architecture of the network.

The Security Question as Open Point of Friction

The proving speed is resolved. The mathematical underpinning is outstanding. Fast proving times rest on STARK-based proof procedures that rely on unproven cryptographic assumptions, primarily regarding the collision resistance of certain hash functions and the security of the Fiat-Shamir transformation that converts interactive proof systems into non-interactive ones.[163] Work from the EF research group and independent cryptographers began during 2025 to formally question individual such assumptions.[163] Should they fall, the affected proof systems would be mathematically unsecured, and the 128-bit security that the protocol sets as its target would not be achievable with current procedures. The Ethereum Foundation responds with graduated milestones: provable 100-bit security by Glamsterdam, 128 bits by end of 2026.[164]

For the scaling path, two things follow from this. Should proving times increase by a factor of three to five under mathematically secured security, the decoupling shifts temporally backward. Single-slot proving would then first be achievable with the next generation of proving hardware or optimized proof systems, which could delay the complete scaling path by one to two years. Network security is not affected by this, because optionality acts as a fallback: as long as the cryptographic underpinning is outstanding, nodes verify through full re-exe-

[161] The optionality requires no hard fork and is backward-compatible. EIP-8025 defines a novel governance status: protocol-side integrated, but without enforcing usage.

[162] EF Blog, 18.12.2025: Performance sprint completed. Proving latency from 16 minutes to 16 seconds, 45-fold cost reduction. Target: 99 percent of all mainnet blocks provable in under 10 seconds, on hardware in the range of $100,000, fully open-source, at 128-bit security.

[163] STARK-based proof systems rely in part on conjectures regarding the collision resistance of specific hash functions and the security of the Fiat-Shamir heuristic. In the course of 2025, both EF-internal work and independent cryptographic research groups began systematically investigating the formal robustness of these assumptions. Source: EF zkEVM Team, L1-zkEVM Roadmap, 26.01.2026.

[164] EF milestones: 100-bit provable security by Glamsterdam (target Q2 2026, but expected only in the second half of 2026 per EF Checkpoint #9 of April 2026), 128-bit by end of 2026. Source: EF Blog, „Shipping an L1 zkEVM #2: The Security Foundations", 18 December 2025.

cution. The question of unproven assumptions is thus a timeline risk for decoupling.

A second security layer reduces the remaining risk further. The proposed 3-of-5 threshold model foresees that five independent client implementations each generate their own proofs. A block is considered valid if three of them are verified.[165] A cryptographic flaw in a single proof system does not compromise the network as long as the other implementations remain intact. Client diversity, which in the current state represents best practice without protocol enforcement, becomes through this model a cryptographic security guarantee.

Why Building Remains Specialized

The zkEVM decouples verification from the Gas Limit. State remains tied to capacity. Block builders must continue to hold the full State, which grows with increasing capacity.[166] The architecture thus deliberately separates three roles. Verification is democratized, because proof verification is independent of block size and functions on resource-constrained hardware. Building remains the role of specialized operators, because a block builder must hold the full State (as of Q1 2026, roughly 430 GiB) in fast access, and this load grows with every Gas Limit increase. Proving establishes itself as a third role on its own hardware basis: specialized GPU clusters with high energy load, without State retention, with its own market conditions and concentration risks.[167] Section 5.3 elaborates both dimensions: Stateless Clients show the democratization of verification. The cryptographic infrastructure layer addresses there the hardware and market conditions of the Prover role.

---

[165] 3-of-5 threshold mechanism: attesters accept a block as valid when three of five independent proofs from different client implementations are verified. EF zkEVM Team, 2026.

[166] The zkEVM shifts the decentralization bottleneck from "running an Execution Layer client" to "maintaining full state." The soundness of the zkEVM is cryptographically guaranteed by the ZK proof verification; the 1-of-N model in EIP-8025 refers to liveness: a single honest prover suffices to keep the proving pipeline running. The economic backing of this minimum operation — i.e., the incentives for a sustainable number of active provers — remains an open coordination task.

[167] EF Blog, 18.12.2025: target for proving hardware in the range of approximately $100,000 (cf. note 32).

ePBS as Convergence Point

The zkEVM solves the cost problem of verification but creates a new requirement in doing so: proof generation requires time that the current slot design does not provide. Without structural change, the prover is left after block propagation and attestation with a window of only one to two seconds, which is insufficient for real-time proving at practically relevant Gas Limits.[168] Enshrined Proposer-Builder Separation (EIP-7732) resolves this bottleneck through a reorganization of the slot sequence.[169] ePBS separates block verification into two phases: consensus verification in slot N and execution verification in slot N+1. This creates a proving window of six to nine seconds in which the prover can complete the cryptographic proof before the attestation of the following slot begins.

The same separation of Proposer and Builder is simultaneously the prerequisite for BALs. The Builder creates the Block Access List in parallel with Proposer selection. Without the protocol-level formalization of this role separation, the procedure cannot be safely implemented.[170] Two features with fundamentally different functions — parallelization in BALs and decoupling in the zkEVM — converge in the same infrastructural mechanism. This convergence is a consequence of the architectural change: as soon as the protocol formally separates block construction and block verification, both scaling paths become possible simultaneously.

The bottleneck is resolved on the governance side. ePBS has SFI status for Glamsterdam, all five consensus layer clients are in active development, and the first generalized devnet was started in April 2026.[171] ePBS unfolds beyond the capacity function described here a second effect: the elimination of the Relay chokepoint and its consequences for censorship resistance. Section 5.5 addresses these.

---

[168] Without ePBS, the proving window is approximately 1–2 seconds. ePBS extends it through block pipelining to 6–9 seconds. Cf. EIP-8025 / EF L1-zkEVM Roadmap, 26.01.2026.

[169] EIP-7732 (Enshrined PBS). Status: SFI Glamsterdam (activation 2026, Q3/Q4 more realistic than H1 per EF Checkpoint 9, April 2026). CL headliner since August 2025.

[170] BALs require ePBS for block pipelining. The zkEVM requires ePBS for the proving window (Slot N+1). Both dependencies are documented in the EF roadmap.

[171] ePBS development status: Prysm (active development), Lighthouse (rebasing devnet4), Teku (tests on dev branch), Lodestar (FOCIL prototype on ePBS branch), Nimbus (in progress). epbs-devnet-0 in preparation.

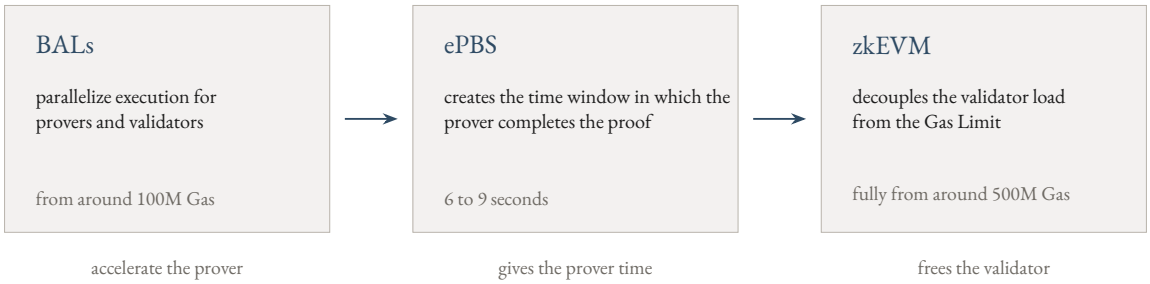

The chain is directed. An undirected relationship would be incorrect.

Figure 5.2-A.1 The Dependency Chain of L1 Execution Scaling. BALs parallelize execution for provers and validators (from roughly 100M Gas). ePBS creates the time window in which the prover can complete the proof (6–9 seconds). The zkEVM decouples validator load from the Gas Limit (from roughly 500M Gas, fully). The chain is directed: BALs accelerate the prover → ePBS gives the prover time → zkEVM frees the validator. Graduated Gas thresholds: <100M (validator voting), 100–150M (BALs + ePBS), 150–500M (Gas repricings), 500–1,000M (zkEVM), >1,000M (open prerequisites: proving economics and State).

The Interplay of the Three Features

The three features form a directed dependency chain in which each link enables the next. BALs provide the dependency graph that first enables parallel witness generation and distributed proving. ePBS reorganizes the slot sequence so that the prover can use the resulting time window of six to nine seconds. The zkEVM translates the output of this chain into a constant verification time for validators, independent of the Gas Limit. If any of the three links is absent, scaling breaks at a different point: without BALs, proof generation becomes the bottleneck; without ePBS, the prover lacks the time window; without the zkEVM, validator load grows linearly with capacity.[172]

The governance status staggers implementation along this chain. BALs and ePBS are confirmed as SFI headliners for Glamsterdam and create the infrastructural foundation on which the zkEVM can build. The zkEVM itself follows an active EF Roadmap with graduated security milestones, whose complete implementation depends on cryptographic underpinning. The complete hundredfold scaling depends on prerequisites beyond execution scaling, which are addressed in sections 5.3 and 5.4. What the chain achieves as a whole, none of the features could achieve alone: a capacity expansion that architecturally safeguards the decentralization of verification by formally separating the roles of Builder, Prover, and Validator and assigning to each role the scaling prerequisites suited to its function.

[172] Temporal sequencing: BALs + ePBS (SFI Glamsterdam, H2 2026) → zkEVM (active EF roadmap since January 2026, Draft EIP-8025, no fork date; graduated security milestones through end of 2026) → Gas limit >1,000M (dependent on state solution and further prerequisites).

5.2-B The Graduated Scaling Path

For Ethereum to fulfill the criteria of Functional Irreplaceability and Coordination Function operationalized in Chapter 3 at the level of global infrastructure, its L1 capacity must substantially exceed today's level of 15 to 30 transactions per second at 60 million Gas. The role separation described in 5.2-A is the architectural answer to this requirement. The graduated scaling path along this role separation begins, however, before its formalization. The Gas Limit doubled from 30 to 60 million between February and November 2025 — in roughly ten months — without the protocol needing to change.[173] This doubling shows that Ethereum can expand its L1 capacity without sacrificing the decentralization of verification. The question for infrastructure assessment is how reliably the prerequisites of each further stage are secured[174] and what new bottlenecks each stage creates. The scaling path divides into five thresholds, with each threshold revealing a specific technical limit that must first be addressed before the next stage becomes achievable.

---

[173] Gas limit: 30M (through February 2025) → 36M (February 2025) → 45M (mid-2025) → 60M (November 2025). Triggered by Vitalik Buterin's public endorsement in November 2024. No consensus change, no fork, no protocol change; the Fusaka hard fork (3 December 2025) standardized the already-achieved 60M value via EIP-7935 as the client default. Cf. EIP-9698 (Informational), author: Dankrad Feist. Validator voting: a decentralized voting process in which each block proposer can adjust the Gas limit by approximately 0.1 percent; consensus emerges through convergence.

[174] The thresholds mentioned in this section (approximately 100M, 150–500M, 500–1,000M, >1,000M) are orders of magnitude derived from client benchmarks of execution-layer teams, EF communication, and public research documents. They mark ranges in which specific bottlenecks become binding and are not fixed protocol boundaries. The exact transitions may vary depending on client implementation and transaction mix.

The Operative Proof

The doubling of the Gas Limit from 30 to 60 million (cf. section 5.1-A) expanded L1 transaction capacity by a factor of two.[173] In validator voting, validators converged on the new target value within a few months, with each block proposer able to adjust the limit by approximately 0.1 percent per block. The Fusaka hard fork in December 2025 anchored two accompanying protective measures. The Gas Limit Default Target (EIP-7935, DEPL) formalizes the preset target value for client software. The Transaction Gas Cap (EIP-7825, DEPL) limits the maximum Gas consumption of an individual transaction to 16.78 million.[175] Blob parameter optimizations progressively raised the Blob Target to 14 and the maximum to 21.[176]

The same mechanism that carried the doubling to 60 million can push the limit to approximately 100 million, because the sequential verification time in this range remains manageable within the 12-second slot.[153] Only above this threshold does the growing verification time compel new protocol prerequisites. In practice, three mechanisms act as sequential thresholds that couple the Gas Limit to actual protocol maturity. When blocks exceed the processing capacity of validators, missed attestations and associated reward losses increase — a self-correcting economic incentive that ties the limit to operational reality. Client teams and the Ethereum Foundation take public positions on the prerequisites of each increase level, as evidenced by Buterin's endorsement of the doubling in November 2024.[173] Large staking operators only raise their targets once their own infrastructure can carry the new requirements. The doubling from 30 to 60 million was accordingly not an emergent process: the EF had communicated the prerequisites, client teams had optimized execution, the hardware was ready. The actual dynamic is sequential: features are deployed, their operational stability is confirmed, and only then does validator voting enable the next capacity level.

[175] EIP-7935 (Gas Limit Default Target, DEPL) formalizes the target value for validator voting. EIP-7825 (Transaction Gas Cap, DEPL, Fusaka) sets a maximum of 16.78M Gas per individual transaction as DoS protection.

[176] Fusaka (December 2025, DEPL) includes DoS hardening as a prerequisite for aggressive Gas limit increases. Blob parameter optimizations: BPO1 (9 December 2025, Blob Target 10 / Max 15), BPO2 (7 January 2026, Target 14 / Max 21).

From Validator Voting to Hard Fork

From 100 million Gas, sequential verification runs into a temporal limit that the current protocol cannot resolve: verification time grows linearly with the Gas Limit and exceeds the 12-second slot once the node can no longer determine which transactions are independently processable. BALs and ePBS, whose mechanisms section 5.2-A presented, address exactly this bottleneck. BALs reduce verification time through parallel processing, ePBS creates the extended propagation window for larger blocks. Both require protocol changes that only a coordinated hard fork can deliver.[177] The transition from autonomous voting to coordinated fork development follows from the complexity of the next stage: BALs change the block header, ePBS reorganizes the slot into two phases, and both require that all clients adopt the new structures simultaneously. Ethereum's fork cadence demonstrates the viability of this coordination: Dencun (March 2024), Pectra (May 2025), Fusaka (December 2025), and Glamsterdam (expected 2026) follow in an increasingly shorter rhythm, which has accelerated from annual to twice yearly and continuously structures the scaling path. Both features have SFI status. Implementations are running in three execution layer clients for BALs and five consensus layer clients for ePBS. The next step is thus specification-ready and in advanced implementation work, with ePBS development proving more complex than originally anticipated according to EF Checkpoint 9 (April 2026).[178] Tomasz K. Stańczak, then EF Co-Executive Director, specified the graduated thresholds in December 2025: first 100 million Gas after Glamsterdam, 200 million Gas after fully operational ePBS, and long-term an L1 target of approximately 10,000 transactions per second.[179] The Soldøgn developer consensus of 2 May 2026 affirmed these stage values as operational orientation.[180]

The threshold at 150 to 500 million Gas creates a second bottleneck: the Gas costs of individual EVM operations no longer reflect actual hardware costs, so

[177] BALs (EIP-7928) and ePBS (EIP-7732): both SFI Glamsterdam (H2 2026). BALs address validation time (parallel transaction processing), ePBS the propagation window (extended distribution time for larger blocks). Hard fork: a coordinated protocol upgrade in which all nodes simultaneously switch to new rules.

[178] BALs: prototype implementations in Besu, Geth, and Nethermind (cf. note 24); BAL devnets making steady progress per EF Checkpoint #9 (April 2026). ePBS: all five CL clients in active development (cf. note 41); the same checkpoint rates ePBS implementation as more complex than originally expected, because the two-party coordination between proposer and builder touches every part of the stack. Status of both features: SFI per EF Checkpoint #9 (April 2026).

[179] Stańczak clarification, public communication of the EF executive team, December 2025. The values are authorized, not speculative. They mark the first stage of a graduated path whose target state (factor 100 over four years per EIP-9698, see note 22) lies far beyond these thresholds.

[180] Soldøgn consensus of core developers of 2 May 2026: reaffirmation of the 100 million threshold after Glamsterdam and the 200 million threshold after fully operational ePBS as binding operational guidance. The consensus marks the transitional phase between Fusaka and Glamsterdam.

that an attacker can use low Gas expenditure to force disproportionately expensive operations. The Gas Repricings (EIP-7904 and EIP-8038, Meta-EIP 8007, CFI Glamsterdam, i.e. Considered for Inclusion) recalibrate costs to current hardware realities.[181] This type of recalibration is established governance practice: EIP-1884 (Istanbul, 2019) and EIP-2929 (Berlin, 2021) had already adjusted Gas costs to changed hardware and altered attack profiles in earlier forks.[182] The details and DoS risks that the current recalibration must address are treated in section 5.2-C.

Up to that threshold, the scaling path is governance-anchored and operationally prepared: each prerequisite has an SFI or CFI status and a named fork. The first near-tenfold increase in L1 capacity, from 60 to approximately 500 million Gas, moves within the range of proven governance mechanisms. In parallel with the execution bottleneck, however, the scaling path generates a correlated State growth from that threshold, whose sustainability section 5.4 quantifies and which places the assessment of the first tenfold increase under a qualification not resolved here.

The Second Tenfold: Active Research, Open Specification

From approximately 500 million Gas, the scaling path enters a stage in which prerequisites are actively being researched but not yet specification-ready. The specific bottleneck lies in full re-execution: from approximately 600 million Gas, verification time exceeds the slot even with parallel processing, so that the zkEVM described in 5.2-A becomes a mandatory prerequisite for decoupling verification and capacity. The zkEVM has a draft EIP (EIP-8025), a published EF Roadmap, and a dedicated research team, but no fork date.[183] A feature without a fork date is the expected state at this development stage: BALs and ePBS likewise had none before becoming specification-ready and receiving SFI status. The remaining uncertainty lies in the precise timing of protocol integration. The Conjectures problem described in 5.2-A acts as a timeline risk that can delay the path. Should the cryptographic assumptions of STARK-based procedures prove untenable, the delay would be more substantial than the one to two years estimated in 5.2-A. The sequential logic of the path bounds this risk: the Gas Limit will not reach 500

[181] Gas repricings: EIP-7904, EIP-8038; Meta-EIP 8007. All CFI Glamsterdam. Recalibration of Gas costs to current hardware realities, independent of Verkle-specific Gas changes.

[182] EIP-1884 (Istanbul, December 2019) increased Gas costs for SLOAD and BALANCE to reflect changed state access times. EIP-2929 (Berlin, April 2021) introduced differentiated Gas costs for cold and warm storage accesses. Both addressed DoS vectors arising from discrepancies between Gas costs and actual execution costs.

[183] EIP-8025 (zkEVM). Status: Draft EIP, active EF roadmap since January 2026 with six sub-themes and graduated milestones (cf. notes 10, 14). No fork date.

million before BALs, ePBS, and Gas Repricings are deployed, and the zkEVM will have had several years of development time by then.

At 600 to 800 million Gas, distributed proving becomes a prerequisite because a single prover can no longer maintain the ePBS window.[184] The economic coordination of such a proving network has neither an EIP nor a Roadmap. The comparison with Builder economics provides a countermodel. MEV auctions, Builder APIs, and the entire PBS infrastructure emerged organically when economic incentives were present, without an EIP specifying them in advance.[185] The analogy has a structural limit, however: MEV extraction is a value extraction game with endogenous profit sources, while proving is primarily an infrastructure service whose compensation must be fed from protocol subsidies or priority fees. The analogy assumes that proving economics generates comparable incentives, which depends on the protocol-level anchoring of the zkEVM and the transaction volume at this capacity level. Execution Tickets (RES), discussed for Building, could prospectively move the coordination of Building and Proving into the protocol. Proving economics is a future coordination problem whose difficulty can be measured against a real analogy but not predicted.

>1,000 Million Gas as Projection

Above 1,000 million Gas lies the range for which the Roadmap holds no concrete answers. The soundness of the zkEVM is cryptographically guaranteed through ZK proof verification. The 1-of-N model ensures liveness as long as at least one honest prover is active. The economic safeguarding of this minimum operation — that is, the incentives for a sustainable number of active provers — remains open as a coordination task.[184] The State problem, which already acts as a growing bottleneck from approximately 500 million Gas and at this threshold becomes the binary limiting factor, is the subject of section 5.4.[186]

Decisive for the infrastructure assessment is the classification of thresholds in operational orders of magnitude. At 60 million Gas per block and 12-second

[184] At approximately 600 to 800M Gas, the proof generation time exceeds the ePBS window of a single prover. Distributed proving becomes a prerequisite. The soundness of the zkEVM is cryptographically guaranteed by the ZK proof verification itself; faulty blocks cannot produce a valid proof. The 1-of-N model in EIP-8025 refers to liveness: a single honest prover suffices to keep the proving pipeline running. The economic coordination of this minimum operation is not addressed by the model and is subject to research.

[185] MEV-Boost (Flashbots, September 2022) established the builder market without a protocol specification. The entire PBS infrastructure, including builder APIs and the relay network, emerged organically in response to economic incentives.

[186] The state problem with increasing Gas limits is quantified and analyzed in section 5.4. Tiered State as a proposed solution has the status RES.

slots, Ethereum processes approximately 5 million Gas per second. At 500 million Gas per block, this would be approximately 42 million Gas per second — sufficient for roughly 2,000 simple ETH transfers or roughly 200 Uniswap swaps per second. The actual capacity depends on the transaction mix. The order of magnitude illustrates the leap from a niche of specialized crypto applications into a performance class in which institutional settlement loads could operate on L1. At 1,000 million Gas, the values double again. The first near-tenfold increase, which this study assesses as governance-anchored and operationally prepared, addresses the infrastructure requirements formulated in Chapter 2. The second moves on a research timeline, whose maturity corresponds to its position in the pipeline. Beyond that begins open terrain that does not concern the infrastructural capacity needs of the relevant planning horizon.

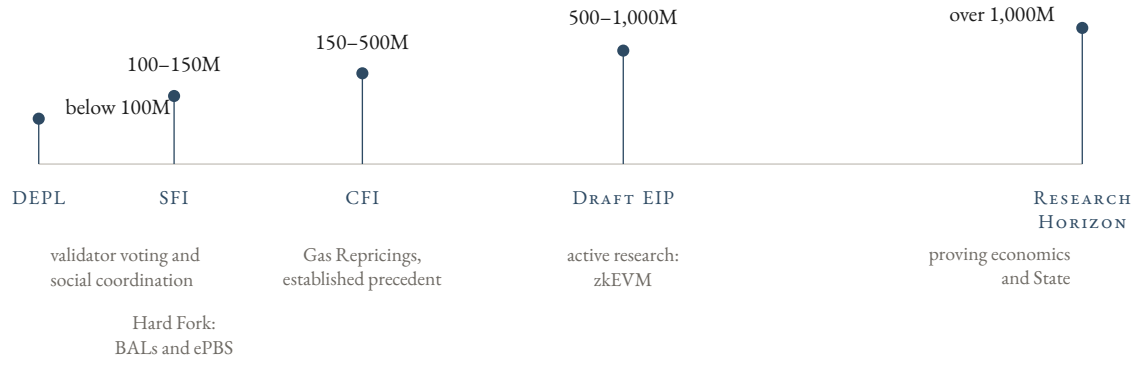

The increasing distance is the point. Scaling does not become uniformly harder, but advances in leaps.

Figure 5.2-B.1 The Scaling Path as Sequential Pipeline. Five thresholds with increasing maturity gaps: <100M (DEPL, validator voting + social coordination), 100–150M (SFI, hard fork: BALs + ePBS), 150–500M (CFI, Gas repricings, established precedent), 500–1,000M (draft EIP, active research: zkEVM), >1,000M (research horizon: proving economics + State). Per threshold: governance status, bottleneck type, de facto thresholds, and sequential dependencies.

The Path as Governance Achievement

Each threshold of the scaling path creates a specific bottleneck that the next protocol change addresses: the sequential verification time to 100 million, the slot congestion to 150 million, the DoS vulnerability of miscalibrated Gas costs to 500 million, the linear coupling between capacity and verification effort to 1,000 million. The sequential dynamic of the path — in which features are deployed before the Gas Limit reaches the next stage — is the actual governance achievement. The practical mechanisms that secure this rhythm are effective even without protocol formalization: economic incentives via missed attestations, public EF positioning on each increase level, and the infrastructure readiness of staking operators. Should a stage be delayed, the path blocks at that point, because validator voting only releases the next threshold once operational reality sustains it. The threshold at 150 to 500 million Gas depends on a recalibration of Gas costs that has so far only been mentioned as a prerequisite. What exactly this recalibration encompasses and which DoS risks it must address the following section clarifies.

5.2-C The Recalibration of the Pricing System

The recalibration that carries the scaling path to 500 million Gas corrects a divergence between Gas prices and real hardware costs. Gas is the pricing system by which the protocol accounts for every operation against the load on validator hardware. The calibration must be chosen so that even the most expensive constructible block can be fully verified within the 12-second slot. If this fails, the honest validator misses its vote on which block is added to the chain and loses part of the reward associated with the slot. Today's Gas prices date from Ethereum's early period. Computation has since become considerably cheaper to execute because processor performance has grown strongly. Storage access, by contrast, has become more expensive. The blockchain's State storage has grown to roughly 430 GiB, so that locating individual entries takes correspondingly longer.[187] The Gas prices for both categories of operations remained unchanged, so that a portion of computation is today overpriced relative to grown processor performance,

[187] State size Ethereum: approximately 430 GiB per Q1 2026, with a range of 410–450 GiB depending on client implementation and compression mode. Sources: Maria Silva (ethresear.ch/t/23476, November 2025) measures 340 GiB in May 2025 (uncompressed, Geth node, dedicated to state). Vitalik Buterin (ethresear.ch/t/24052, February 2026) gives current growth at approximately 100 GB per year. The Q1 2026 estimate follows from the May 2025 baseline plus two-stage growth: approximately 73 GiB per year at a Gas limit of 36 million (through June 2025), approximately 124 GiB per year at a Gas limit of 60 million (from December 2025). Compute performance per processor has grown substantially since the original Gas calibration due to hardware development, while access time to state has structurally increased with state growth. The asymmetry between compute, state, and data is treated as a distinct problem in section 5.4.

while other computation and storage accesses remain underpriced relative to their actual hardware load.

At 60 million Gas per block, this divergence is too small to noticeably strain validator capacity. From 150 to 500 million Gas it becomes a practical problem. At this stage an attacker can with a low Gas budget construct blocks that generate considerably higher hardware load for validators than they pay for. At greater block size, more underpriced operations fit simultaneously into a block, so that hardware load can cumulatively exceed the slot boundary. A graduated capacity path thus places a requirement on the pricing system. Relative prices must approach the real hardware load closely enough that an attacker with little Gas can no longer generate oversized validator load.

The Glamsterdam repricings raise the Gas costs of three systematically underpriced categories of operations: computation, State access, and State creation. The Compute Gas Cost Increase (EIP-7904, CFI Glamsterdam) raises the costs of 13 underpriced computation operations and precompiles whose actual execution is slower than their current pricing reflects. Cryptographic precompiles such as POINT_EVALUATION rise from 50,000 to 89,363 Gas, basic arithmetic operations such as DIV from 5 to 15 Gas.[188] The State Access Gas Cost Increase (EIP-8038, CFI Glamsterdam) raises the costs of underpriced State accesses. A read operation such as EXTCODESIZE, which triggers two database reads on the current State, is recalibrated in this direction.[189] As a third correction, the State Creation Gas Cost Increase (EIP-8037, SFI Glamsterdam) raises the costs for creating new storage slots and accounts. The long-term effects of new State entries on the overall size of State were not priced into the original calibration.[190] The question of why State growth remains an unresolved structural problem despite the adjustment is addressed in section 5.4. The three corrections affect

[188] EIP-7904 (Compute Gas Cost Increase). Status: PLAN, CFI Glamsterdam. Raises the Gas costs of 13 underpriced compute operations and precompiles whose actual execution on validator hardware is slower than their current pricing reflects. Example increases: DIV 5 to 15 Gas, MOD 5 to 12 Gas, SDIV 5 to 20 Gas, KECCAK256 30 to 45 Gas, BLAKE2F 0 to 170 Gas, BLS12_G1ADD 375 to 643 Gas, ECADD 150 to 314 Gas, POINT_EVALUATION 50,000 to 89,363 Gas. Raising the costs of underpriced operations is a prerequisite for the planned block limit increase, because bottleneck-forming operations must first be repriced before they no longer constrain capacity at larger blocks. POINT_EVALUATION (KZG commitment verification), BLS12_G1ADD, and ECADD are structurally part of the on-chain precompile layer, whose systematic classification section 5.3 elaborates in the treatment of the cryptographic infrastructure level. Source: EIP-7904 original text.

[189] EIP-8038 (State-Access Gas Cost Increase). Status: PLAN/Draft, CFI Glamsterdam. Increase of costs for underpriced state accesses. Example evidence: operations with two database reads per call whose I/O costs the current Gas prices no longer reflect at today's state size.

[190] EIP-8037 (State Creation Gas Cost Increase). Status: PLAN, SFI Glamsterdam. Raises the Gas costs for creating new storage slots and accounts. In the original Gas calibration, the long-term impact of new state entries on total state size was not adequately priced in. Source: Meta-EIP 7773 (Glamsterdam Meta-EIP) and EIP-8037 original text.

application classes with differing intensity: protocols with a high density of cryptographic operations or State-creating deployments bear stronger cost increases than protocols with lightweight transaction profiles. Section 5.5 addresses the neutrality implications of the differential effect.

Reduce Intrinsic Transaction Gas (EIP-2780, CFI Glamsterdam) lowers as a flanking component the fixed 21,000 Gas intrinsic costs that every standard transaction carries regardless of its content.[191] Meta-EIP 8007 centrally consolidates all EIPs belonging to the repricing package and thereby gives the ecosystem an overview of which corrections belong together. The substantive coordination between the individual EIPs is the work of the Glamsterdam repricings calls of the Ethereum core developers.[192]

Two mechanisms already active in Fusaka limit the room that an attacker can exploit per transaction and per fork. The Transaction Gas Cap (cf. section 5.2-B) acts in the recalibration as a boundary condition. At a block size of 500 million Gas, even the most expensive permissible transaction can thus occupy only 3.4 percent of the block. An attacker would consequently need to combine many transactions to push a block beyond the slot, which substantially increases the cost of attack. The Gas Limit Default Target formalizes the target value by which validators orient their Gas Limit voting. Validators can increase or decrease the Gas Limit slightly per block. The Default Target gives them the common reference value at which individual decisions coordinate. The corrections and the flanking mechanisms address the price–hardware divergence on three levels: the repricings calibrate the price per operation, the Transaction Gas Cap limits the effect per transaction, and the Gas Limit Default Target coordinates the adjustment per fork.

---

[191] EIP-2780 (Reduce Intrinsic Transaction Gas). Status: PLAN, CFI Glamsterdam. Lowers the fixed 21,000 Gas intrinsic costs of a standard transaction. A flanking component alongside EIP-7904, EIP-8038, and EIP-8037. Source: Meta-EIP 7773 (Glamsterdam Meta-EIP).

[192] Meta-EIP 8007. The EIP text declares itself as "purely informational" and explicitly states that it takes no active role in the governance process. The meta-EIP centrally aggregates the individual EIPs of the repricing package. Content coordination takes place in the Glamsterdam Repricings Calls of the Ethereum core developers.

The Friction of Every Recalibration

The recalibration of the pricing system follows in Ethereum an established practice whose friction point structurally belongs to the mode. The Istanbul recalibration of December 2019 adjusted the Gas costs of a triad of State-access opcodes to changed State access times.[193] The effect can be read from SLOAD, whose costs were raised from 200 to 800 Gas to carry the increased I/O load. Roughly 680 deployed Aragon contracts became non-functional after this adjustment because their internal transfer logic implicitly rested on the old Gas costs. ETH deposits from other smart contracts to Aragon organizations before version 0.8 could no longer be retrieved after the fork. The affected range extended beyond Aragon to Kyber, Gods Unchained, and various OpenZeppelin proxies. This is a structural cost category of recalibration: every price increase for underpriced operations hits contracts that carry the same assumption. For the current Glamsterdam package it follows that the price increases there will likewise hit such contracts. This friction is the price of every recalibration of the pricing system. The Berlin hard fork confirmed the logic 16 months later with another calibration of State accesses, which followed as the second major adjustment within that period.[194] The fork cadence has since become shorter: 14 months between Dencun and Pectra, 7 months between Pectra and Fusaka, roughly six months aimed at between Fusaka and Glamsterdam.[195] The semi-annual cadence is the stated goal. Governance responds iteratively; the ecosystem absorbs the costs of adjustment in those contracts built on the old calibration.

On several levels simultaneously, the components of the correction address the pricing side of the threshold between 150 and 500 million Gas. In the interplay of elements, the compute increase corrects the relationship between opcode

[193] EIP-1884 (Istanbul, 8 December 2019). Recalibration of the Gas costs of a triad of state-access opcodes (including SLOAD, BALANCE, EXTCODEHASH) to changed state access times. SLOAD from 200 to 800 Gas. Breaking effect: approximately 680 Aragon contracts (per Aragon CTO Jorge Izquierdo) became non-functional because their internal transfer logic implicitly relied on the old Gas costs; ETH deposits from other smart contracts to pre-0.8 Aragon organizations could no longer be retrieved after the fork (Aragon Help Desk). Also affected: Kyber, Gods Unchained, and OpenZeppelin AdminUpgradeability Proxies. Sources: EIP-1884 original text, Aragon Help Desk, CoinDesk.

[194] EIP-2929 (Berlin, 15 April 2021). Introduction of a differentiated cost model for state accesses that prices first accesses differently from repeated ones. Second major recalibration 16 months after EIP-1884. Source: EIP-2929 original text.

[195] Fork cadence: Dencun (March 2024) → Pectra (May 2025), 14 months; Pectra → Fusaka (December 2025), 7 months; Fusaka → Glamsterdam, approximately six months targeted. The six-month cadence is the stated goal of Ethereum core developers. Per EF Checkpoint 9 (April 2026), however, ePBS is proving more complex than originally expected, so activation in Q3 or Q4 2026 is more realistic than H1. The cadence discipline is a structural statement about governance maturity (criterion III.2), not about the time horizon of the target state, which extends far beyond this fork.

costs and hardware load. The State access increase removes an attacker's ability to generate disproportionate I/O load with little Gas, while the State creation increase binds the creation of new State entries to their long-term costs. The Transaction Gas Cap holds the effect of an individual transaction at the 3.4 percent of the block cited above. The Default Target anchors the path on the governance side by formalizing the target value around which the coordination of all subsequent adjustments orients. A further Glamsterdam component shifts the structural boundary of contract size: EIP-7954 (Max Contract Size Increase, CFI Glamsterdam) raises the current 24 KB limit per contract to 64 KB, thereby relieving larger zkVM verifiers and more complex DeFi architectures.[196] The price safeguard thus stands alongside the parallelization and verification mechanisms described in 5.2-A. These elements jointly address the first near-tenfold capacity increase envisioned in the Roadmap.

What the Recalibration Does Not Solve

The safeguard remains within a structural boundary that Glamsterdam shifts but does not remove. The PLAN corrections operate within a pricing system that combines computation time, storage access, and bandwidth in a common Gas budget. A price shift on one side thus continues to feed back onto the other sides. Two features at RES status aim at a structural redesign of this system. EIP-8011 (Multidimensional Gas Metering, RES/DFI Glamsterdam, i.e. Declined for Inclusion) proposes replacing the one-dimensional Gas budget with separate counters for multiple resource categories, which the EIP's author bundles as block execution time, block network time, short-term memory, and long-term storage. Each resource category thereby receives its own Gas budget, so that a price shift on one side no longer distorts the other sides.[197] EIP-8057 (Inter-Block Temporal Locality Gas Discounts, RES/DFI Glamsterdam) pursues a different direction and discounts Gas costs for accesses to State entries that have been touched within a window of the last 32 blocks. The idea draws on empirical

[196] EIP-7954 (Max Contract Size Increase). Status: CFI Glamsterdam. Raises the EIP-170 limit from 24,576 bytes (24 KB) to 65,536 bytes (64 KB) per contract, thereby relieving applications whose bytecode presses against the old threshold: in particular zkVM verifiers, more complex DeFi architectures with embedded libraries, and Solidity compiler output without aggressive optimization. The increase is situated in the context of Gigagas scaling, because greater block capacity should also support more complex individual transactions.

[197] EIP-8011 (Multidimensional Gas Metering). Status: RES/DFI Glamsterdam (declassified from the Glamsterdam inclusion list per EIP-7773, still in research status). A conceptual development that replaces the one-dimensional Gas budget with separate counters for multiple resource categories. The EIP specifies six technical dimensions (compute, access, size, memory, state, history) and clusters them in the rationale section into four conceptual resource categories: block execution time, block network time, short-term memory, and long-term storage. The presentation in the main text follows this four-cluster grouping.

cache observations: recently touched entries are likely still in the client cache and incur lower actual costs on re-access than cold entries.[198] Both options carry EIP numbers but are assigned to no fork and have no fork date. They mark research directions intended to structurally advance the pricing side beyond the point corrections of the PLAN components.

The reach of the overall price safeguard remains limited in two respects. First, the Glamsterdam package calibrates prices for the threshold between 150 and 500 million Gas. Higher thresholds will generate further recalibrations that the established governance mode can carry, but which are not delivered in this fork. Second, the correction addresses the pricing side, not State growth itself. The structural asymmetry between the processing, data, and storage resource thus persists. The data resource is addressed in section 5.2-D, State growth in section 5.4.

The Roadmap thus aims to develop the pricing system into an adaptive layer of the protocol that aligns with the hardware reality of network resources. The corrections at PLAN status are intended to address the divergence between price and hardware load within the existing one-dimensional Gas budget. The research directions at RES status — namely multidimensional Gas metering and the temporally local discounting of recently touched State entries — aim at a structural redesign of the pricing system itself, whose concrete implementation remains open. The goal is for the network to price its resources as the hardware actually loads them, and to readjust this pricing along the graduated capacity levels.

The pricing side of the execution resource is addressed in the Roadmap through PLAN corrections and RES research directions. The protocol scales not only processing, however. L2 Rollups publish their transaction data on L1, so that the capacity of the second direction becomes a standalone scaling task.

[198] EIP-8057 (Inter-Block Temporal Locality Gas Discounts). Status: RES/DFI Glamsterdam (declassified from the Glamsterdam inclusion list per EIP-7773, still in research status). A conceptual option that discounts Gas costs for accesses to state entries touched within a rolling window of the last 32 blocks (approximately six minutes). The discount follows a smoothstep decay function that decreases with temporal distance from the last access. Background: recently touched entries are likely still in client caches and incur lower actual I/O costs on re-access than cold entries.

5.2-D The Scaling of Data Availability

The recalibration of the pricing system from section 5.2-C calibrates the execution resource. The second resource that the scaling path addresses follows its own logic. The data that L2 Rollups publish on L1 is not subject to the same bottlenecks as execution itself. Before Proto-Danksharding (EIP-4844, DEPL since March 2024)[199] these data competed for the same block space as regular transactions and thereby directly determined L2 cost overhead. Proto-Danksharding resolved this competition by introducing Blobs as a separate data type with its own fee market.[200] Blob Gas follows its own EIP-1559-analogous price dynamic, decoupled from execution Gas. The data availability load of L2 Rollups is thereby separated in price from the L1 execution load.

Before Dencun, the large L2 Sequencers together paid more than 15,000 ETH per month for calldata on L1 — an order of magnitude of roughly 34 million US dollars.[201] After Dencun, Blob costs at normal utilization move near the minimum. Transaction costs on the major Optimistic Rollups fell from an average of roughly 0.35 US dollars before Dencun to below 0.01 US dollars after Dencun.[202] This corresponds to a reduction of 95 to 99 percent, with deviations depending on the Rollup and utilization. For an L2 operator this means a change in the cost model: data availability was one of the largest line items before Dencun. After Dencun it is no longer a structural operating load.

This decoupling harbors a vulnerability at the lower end of the price corridor. At below-target utilization — when fewer Blobs are submitted per block than the target value specifies — the Blob price converges rapidly toward zero, because the EIP-4844 fee formula reduces the price exponentially in each block. The data availability resource is thus factually cost-free at low utilization, even though it consumes L1 block space and validator bandwidth. The protocol carries the

[199] EIP-4844 „Shard Blob Transactions" was activated on the Ethereum mainnet with the Dencun hard fork on 13 March 2024. Source: eips.ethereum.org/EIPS/eip-4844.

[200] Technical parameters per EIP-4844 „Shard Blob Transactions": blob size 131,072 bytes (128 KB) = 4,096 field elements × 32 bytes on the BLS12-381 scalar field. Retention period 4,096 epochs (≈ 18 days) in the Consensus Layer, per MIN_EPOCHS_FOR_BLOB_SIDECARS_REQUESTS. After the retention period expires, only the KZG commitments to the blob contents remain permanently on-chain. Source: eips.ethereum.org/EIPS/eip-4844 (Final, activated with Dencun on 13 March 2024); Ethereum Consensus Specs, Deneb fork specification.

[201] Aggregated pre-Dencun calldata expenditures of major L2 sequencers (in particular Arbitrum, Optimism, Base) on L1, order of magnitude converted at average ETH rate Q1/2024. Source: on-chain data via Dune Analytics, L2BEAT (as of March 2024).

[202] Arbitrum: average fee of approximately $0.37 before Dencun to $0.012 after Dencun (−97%). Optimism: $0.32 → $0.009 (−97%). Base: $0.31 → $0.0005–0.001 (−96% to −99%). Starknet: $6.80 → $0.02–0.04 (−99%). The average value cited in the main text of approximately $0.35 before Dencun and under $0.01 after Dencun is an average across the major Optimistic Rollups; values vary with blob utilization and rollup-specific compression. Source: L2BEAT, growthepie.xyz.

resource economically without drawing any counter-revenue from it. The Blob fee share of validator revenues approaches zero, and the L1 security budget weakens accordingly. Blob Fee Stabilization (EIP-7918, DEPL with Fusaka)[203] couples the minimum Blob base fee to execution Gas costs, so that the price assumes a floor that reflects the block space costs of Blob submission. The Blob fee share of the L1 security budget is thereby maintained even at below-target utilization. The complete security budget argument is developed in section 5.5. At this point it suffices to note that EIP-7918 economically stabilizes the data resource.

PeerDAS: Verification via Sampling

The price layer is economically calibrated. The second architectural layer of the data resource addresses the verification side. Until Fusaka, every full node downloaded every Blob in full, so that validator bandwidth scaled linearly with Blob capacity. PeerDAS (EIP-7594, DEPL with Fusaka in December 2025)[204] replaces this effort with Data Availability Sampling — a probabilistic procedure. Each node checks only a small portion of the Blob data. An attacker withholding data becomes statistically detectable. Technically, the Blobs are extended horizontally via Reed-Solomon coding, and the extended matrix is divided into 128 column subnets. Each node deterministically downloads only the roughly 12.5 percent of data assigned to it.[205] The architecture follows the same principle as the zkEVM on the execution side: where the zkEVM replaces re-execution with proof verification, PeerDAS replaces full data retention with sampling verification. Both mechanisms decouple the resource requirement per node from the total load of the network.

The capacity effect of this decoupling can be read from the Blob parameter path of the recent hard forks. Since Dencun, Blob capacity has been increased in several steps. The current status after BPO2 is Target 14 and Maximum 21 Blobs per block.[206] The theoretical 1D maximum of the PeerDAS architecture is

[203] EIP-7918 „Blob base fee bounded by execution cost". Introduced with the Fusaka hard fork. Couples the minimum blob base fee to execution Gas costs to avoid convergence of the blob price toward the absolute minimum of 1 wei (10^{-18} ETH) at low blob utilization. The EIP-4844 fee formula reduces the blob price at below-target utilization exponentially in each block per BLOB_BASE_FEE_UPDATE_FRACTION = 3,338,477. Source: eips.ethereum.org/EIPS/eip-7918, eips.ethereum.org/EIPS/eip-4844.

[204] EIP-7594 „PeerDAS — Peer Data Availability Sampling". Activated on mainnet with the Fusaka hard fork in December 2025, with implementation in all six Consensus Layer clients (Prysm, Lighthouse, Teku, Nimbus, Lodestar, Grandine). Source: eips.ethereum.org/EIPS/eip-7594, ethereum.org Forkcast.

[205] PeerDAS parameters: 128 column subnets in the extended blob matrix; deterministic assignment of a subset per node based on node ID; average download load corresponds to approximately 1/8 = 12.5 percent of total blob data. Source: EIP-7594, CL specs.

[206] Blob parameters over time: Dencun (March 2024): Target 3 / Max 6. Pectra (May 2025): Target 6 / Max 9. Fusaka + BPO1 (December 2025): Target 10 / Max 15. Fusaka + BPO2 (January 2026): Target 14 / Max 21. Source: ethereum.org Forkcast, EF Blog „Fusaka BPO Schedule".

64 to 128 Blobs, and thus a multiple of the current level.[207] At a Rollup-specific compression of a few hundred to a few thousand L2 transactions per Blob, the current level reaches the order of magnitude of a few thousand L2 TPS, while the 1D maximum reaches several tens of thousands.[208] The mechanism that enables these increases is Blob-Parameter-Only Forks: isolated parameter changes that exclusively affect Blob capacity parameters without requiring a coordinated hard fork. The previous BPO rounds demonstrate the implementation maturity of the PeerDAS mechanics under production load. Blob capacity can thus be expanded independently of the general fork calendar as long as network stability sustains the next step. The execution resource with its hard fork dependencies knows no such governance mode. Operational validation under adversarial conditions will be demonstrated in the coming Blob increases — specifically in network stability at higher Blob counts and in sync behavior at column subnet unbalance.

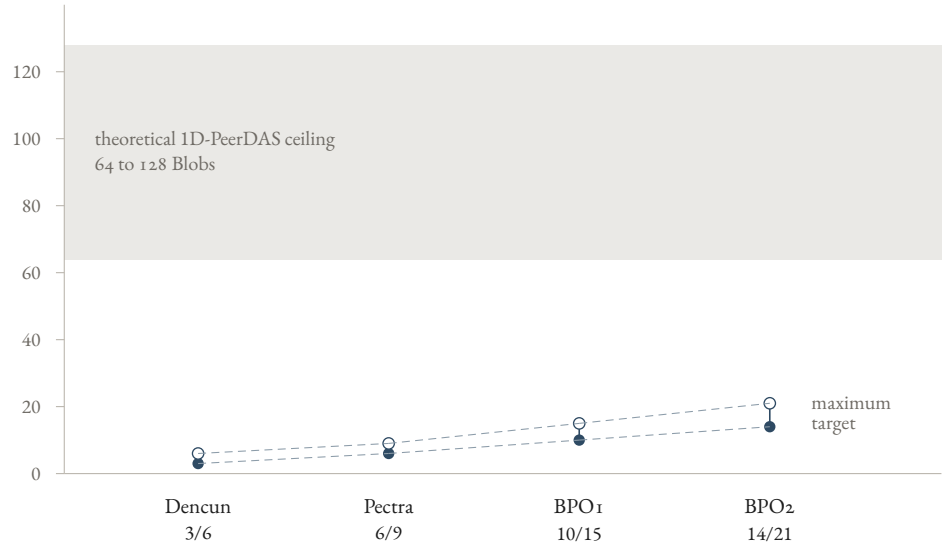

The gap between the achieved value and the ceiling conveys the point.

Figure 5.2-D.1 Blob Parameter Path — Dencun (3/6) → Pectra (6/9) → BPO1 (10/15) → BPO2 (14/21), with theoretical 1D PeerDAS ceiling at 64–128 Blobs

[207] The theoretical 1D PeerDAS maximum of 64–128 blobs per block is derived from the bandwidth and networking parameters of the 1D sampling architecture. Source: ethresear.ch „PeerDAS Capacity Analysis", EF Research Blog.

[208] Blob compression values vary between rollups and with the transaction composition: Arbitrum and Optimism report compression values on the order of several hundred transactions per blob, highly optimized rollups (Starknet, zkSync) approach four-digit values. At a slot interval of 12 seconds and a blob target of 14, compression values between 500 and 2,000 L2 transactions per blob yield an aggregated L2 TPS capacity on the order of a few thousand; at the 1D ceiling (64–128 blobs), correspondingly several tens of thousands. The values are orders of magnitude, not precise projections. The load-bearing statement is the decoupling of L2 capacity from the L1 Gas limit. Source: L2BEAT blob compression data, growthepie.xyz, Ethereum Foundation Research.

Beyond the 1D maximum, the long-term path of the data resource lies in a two-dimensional extension of the same architecture. Full Danksharding (RES, no fork assignment) extends PeerDAS from one-dimensional to two-dimensional sampling by adding a vertical Reed-Solomon extension to the horizontal one. The reconstructibility of data scales quadratically rather than linearly, and a node grasps Blob completeness from a considerably smaller sample fraction. In this matrix, nodes download individual cells (matrix cells) rather than entire columns, so that the data volume per node remains constant as Blob count grows. Validator requirements are structurally decoupled from the scaling objective.[209] The capacity projection is 64 Blobs Target and 256 Maximum, corresponding to roughly 32 MB of data throughput per slot.[210] The technical prerequisite is efficient bivariate polynomial interpolation that reconstructs the complete matrix from a partial subset of samples. This reconstruction has been mathematically defined for some time. What remains open is the production-ready implementation, because the computational load per reconstruction currently exceeds the hardware budgets of prover and validator infrastructure.[211] Full Danksharding thus depends on an engineering bottleneck whose solution is expected but not scheduled. Under the L1-first pivot, the path is not critical as long as PeerDAS carries the foreseeable L2 development. Accelerated adoption through institutional Rollups, ENS Namechain, or Based Sequencing can overtake the baseline assumption.

Alongside the engineering dimension, the target state harbors a cryptographic migration question. The KZG commitments that today guarantee Blob correctness without full download rest on elliptic curve pairings, which a sufficiently large quantum computer would break via Shor's algorithm. The Roadmap therefore foresees the transition to STARK- or lattice-based commitments. The problem is that STARKs do not possess the mathematical linearity property (homomorphism) of KZG. Since linearity constitutes the mathematical foundation for two-dimensional Data Availability Sampling — the efficient verification of data

[209] Full Danksharding extends the Reed-Solomon encoding from a purely horizontal structure (row extension) to a two-dimensional one (row + column extension). Nodes sample cells rather than columns. The sampling bandwidth per node is thereby decoupled from the blob count and approaches a constant value. Source: Vitalik Buterin, „Proto-Danksharding FAQ" and „Dankrad Feist on 2D DAS" (notes.ethereum.org).

[210] Full Danksharding target specification: 64 blobs target / 256 blobs maximum per slot, corresponding to approximately 32 MB data throughput per slot. Source: ethereum.org „Danksharding Roadmap", Dankrad Feist Research Notes.

[211] The bivariate polynomial interpolation for 2D matrix reconstruction is implemented in algorithmic variants, in particular as FFT-based methods and GPU-accelerated Reed-Solomon encoders. The open question is production-ready efficiency under the latency and hardware budgets of a slot-synchronous prover and validator infrastructure. Source: EF Research publications on 2D-DAS algorithms, Dankrad Feist Research Notes (notes.ethereum.org).

structures via partial samples — the switch leaves an as yet unresolved downstream problem in the scaling architecture.[212]

State: The Direction without Breakthrough

While the data resource permits a graduated expansion, the two breakthrough mechanisms on the execution and data sides raise a question that reaches beyond the data resource. The three scaling resources of the protocol — Execution, Data, and State — stand at three different maturity levels that follow from the architecture. The execution resource has with the zkEVM an identified breakthrough mechanism whose governance path and Gas threshold assignments sections 5.2-A and 5.2-B unfolded. For the data resource, PeerDAS provides a deployed breakthrough mechanism and Full Danksharding a long-term option. The storage resource, by contrast, has no comparable breakthrough mechanism standing at the same maturity level as these two.

The breakthroughs in the first two directions drive the growth of the third. Every additional execution capacity can generate State, and every additional data capacity enables L2 activity whose settlement on L1 again produces State. With every scaling stage, the asymmetry sharpens: what in earlier Roadmap formulations appeared as a long-term risk becomes, under the L1 Gas Limit path, a medium-term bottleneck.

The role separation from 5.2-A solves verification and proving, but not Building itself. The zkEVM decouples the validator from Gas capacity, PeerDAS additionally from Blob bandwidth. Distributed proving decouples the prover from block volume. ePBS separates the Proposer from Building and simultaneously eliminates the Relay chokepoint (cf. section 5.5). The role of Block Builder remains tied to State. A Builder must hold the full State to construct a valid block. With every capacity level, its State synchronization load grows. The entry barrier for Block Building thus rises with each capacity level, even as the verification and proving barriers fall. The decentralization of verification is thereby architecturally sustained. The decentralization of Building remains tied to the unresolved State resource. Section 5.4 examines this State dependency as the deepest unresolved bottleneck of the Roadmap.

The execution resource thus stands with an architectural chain, a graduated path, and the prospective decoupling through the zkEVM. The data resource

[212] On the KZG quantum vulnerability and migration candidates: Vitalik Buterin, „Possible Futures of the Ethereum Protocol — Post-Quantum Roadmap" (26 February 2026); ethereum.org/roadmap/future-proofing/quantum-resistance/. The linearity trade-off between KZG and STARK commitments for two-dimensional data sampling is made explicit in Vitalik's roadmap.

stands with a deployed sampling mechanism and a two-dimensional long-term path. State remains the open dimension and is the subject of section 5.4. The Blob capacity of the data resource simultaneously unfolds an effect beyond the L2 cost model. It is the L1 prerequisite for Rollups to shift their security guarantees from an institutional to a cryptographic basis. Section 5.6 addresses this shift as a separate architectural question. Up to this point, the design state and the interplay of features that the Roadmap provides for capacity expansion in the two scaling directions have been established. Beyond the design state, a development emerges that redesigns not the capacity of the execution layer but its execution environment itself.

5.2-E Post-Design State: RISC-V, leanVM, and the EIP-Invariant Role Separation

The execution environment whose further development the preceding section marked is the Ethereum Virtual Machine: the runtime environment in which all the scaling mechanisms described above run. The zkEVM, which 5.2-A introduced, proves the correct execution of this existing machine with its historically grown opcodes and State transitions. A deeper further development replaces not individual mechanisms but the instruction set architecture itself: the foundational catalog of instructions that the runtime environment interprets and executes. Every smart contract code is ultimately translated into these instructions. Buterin's proposal of May 2025 sketches this direction: to supplement or replace the EVM in the long term with a RISC-V-based or similarly ZK-native execution environment.[213]

RISC-V and the Proving Advantage

RISC-V (RES, exploratory, no governance anchoring) stands for Reduced Instruction Set Computing, fifth generation. The philosophy behind the Reduced Instruction Set is a minimalist foundational decision in computer architecture: few, simple instructions that each execute only one computation step, while complex macro-instructions decompose internally into many sub-steps. RISC-V is the open standardization of the architecture, developed at UC Berkeley in 2010 and since available under a license-free specification. For Ethereum, neither the academic origin nor the license-freedom is the decisive point. What is decisive is what the simplicity of the instructions means for ZK proving: every

[213] Vitalik Buterin, „A simplified version of the EVM: Long-term L1 execution layer proposal with RISC-V", Ethereum Magicians Forum, May 2025. The proposal carries no EIP status and no fork date. It is published as a research direction, not as an implementation plan.

RISC-V instruction executes only a single computation step, without hidden internal complexity, and can therefore be mapped with few mathematical equations. A ZK circuit constraint is such an equation that enforces the correct execution of an instruction: it specifies which values must hold before and after the instruction for the computation to be provably valid. Each instruction is represented by a set of such equations. A RISC-V instruction operates on a small number of registers — the fast internal storage cells of a CPU — and performs on them a single arithmetic operation: addition, multiplication, comparison, shift. Side effects — a simultaneously modified stack state, a growing memory area, an accounting of resource consumption — are absent. An EVM opcode, by contrast, typically combines several such basic operations internally and additionally carries a Gas accounting, a growing memory area, and stack operations with it. All of this must be mapped in ZK verification. This is the source of the order-of-magnitude difference: a RISC-V instruction is typically associated with single-digit to low double-digit equations, an EVM opcode execution with several hundred to thousands. The performance improvement for smart contract execution in ZK provers derived from this lies, according to Buterin's proposal, at a factor of 100 or more. The practical consequence: the proving effort, which today sets significant latency and cost barriers, would fall to levels sufficient for realtime proving within a slot window.[214] The figure is a projection from the proposal. It has been measured on no running system.

The Roadmap expects three consequences from this architectural change that increase accessibility and provability at different levels. The first level is proving economics. As a direct consequence of the constraint counts in the preceding paragraph, specialized proving infrastructure is needed today: an EVM block contains millions of opcode executions, which together generate several billion constraints. A proof system that must process this quantity in acceptable time requires massively parallel hardware — typically GPU clusters or specialized ASICs. With a RISC-V-based execution environment, the infrastructure loses its structural necessity: the constraints per instruction fall by one to two orders of magnitude, the total effort becomes manageable on commodity hardware, the prover market opens to a broader actor base. The second level is developer generativity. Worldwide there are between 70,000 and 80,000 developers with Solidity

[214] ibid. The 100× projection refers to proving performance in ZK provers, not to native execution performance. The figure is a guideline from the primary source, not a measurement of a running system. The derivation is based on the reduction in constraints per instruction. The consequence level mentioned in the main text (latency and cost constraints, realtime proving within a slot window) is a qualitative classification based on proving development in 2024/2025; concrete time anchors for comparing EVM proving latency vs. RISC-V proving latency are not quantified in Vitalik's proposal.

experience. Rust, C, and Go together are the system languages from which Linux kernels, browser engines, cloud infrastructure, databases, and operating systems emerge. Smart contracts would be writable for the same developer pool from which the infrastructure itself comes. The special community around Solidity would persist but be expanded by an order of magnitude. Smart contract development would no longer be a specialized discipline with its own language and own community — it would become a domain of general software development. The third level is formal verification. The RISC-V Sail specification is a complete, academically reviewed formal description of the execution environment, while the EVM has only partially formalized semantics for individual components: Runtime Verification models for individual opcodes, K-Framework formalizations of selected subsets, but no unified formal foundation. Simpler execution environments with more clearly specified instruction sets allow faster formal verification of their implementations: the mathematical proof that concrete software does exactly what the specification prescribes. This proof is operationally relevant for Ethereum because the EVM specification is carried by several independent client implementations — software in different programming languages, such as Geth in Go, Nethermind in C#, Besu in Java, Reth in Rust. Each of the implementations must behave exactly the same way for the network to achieve consensus. Today, extensive testing and mainnet operation ensure the correspondence. A formally verifiable execution environment would allow the correspondence to be mathematically proven. Consensus bugs through divergent client interpretations would thereby be structurally excludable, whereas tests and operation today only make them unlikely.[215]

The horizon of the potentials lies in the post-Strawmap period after 2029. RISC-V stands without governance anchoring, without an EIP, and without a fork date. Temporally the proposal lies behind the zkEVM: only when the latter is deployed and stabilized does the question of a RISC-V-based further development become operative.

[215] On the Sail-based formal specification, see RISC-V International, „The RISC-V Instruction Set Manual, Volume I: Unprivileged ISA", with accompanying Sail modeling. On partial EVM semantics formalizations, cf. Runtime Verification (EVM-as-Automaton) and the K-Framework modelings of individual opcodes. On the classification within the Lean Ethereum philosophy, cf. Vitalik Buterin, „Possible futures of the Ethereum protocol", February 2026.

leanVM as Target State of Execution

The EVM stands within a larger protocol framework whose accumulated complexity the Ethereum Foundation has identified as its own burden since around 2024. This burden directly affects the EVM.[216] Buterin and others have repeatedly identified the greatest structural task of the coming years as not further expansion but systematic simplification: less code, clearer semantics, formally specified components whose implementations provably comply.[217] Lean Ethereum is the program that carries this simplification, and section 5.1-B develops it as the overall framework. Within this program, leanVM (RES, horizon beyond 2029) marks the postulated target state of the execution environment with three defined properties. The precise interlocking with the RISC-V line remains open in the primary sources.

leanVM defines itself through three dimensions that in their interplay change the trust model of the protocol. The first dimension is harmonized client implementations. Client diversity in today's Ethereum is an active ecosystem effort, won through separate teams, different programming languages, and years of specification coordination. Bugs in individual clients must therefore be caught by the diversity of other clients. The leanVM architecture would derive client diversity from a formal specification. A formal specification is a document that precisely defines every instruction in a mathematical language without room for interpretation by implementers. Different teams continue to build different implementations, in different languages, with different optimizations, but their semantics — their input-output behavior — would be mathematically provably identical to the specification. The risk of a consensus split through divergent interpretation disappears. The second dimension is formal verifiability. Today's EVM is empirically stable. Its correctness is evidenced by years of operation, extensive testing, and bug bounties, not by mathematical proof. A formally verified target state would derive implementation correctness from specification proofs. Whole classes of bugs would consequently be architecturally excluded. The third dimension is ZK nativity. The EVM uses primitives efficient on classical CPU hard-

[216] Justin Drake (Ethereum Foundation) characterizes the ACD process (All Core Devs) as burdensome for the developer community: per hard fork, typically three to five EIPs are resolvable, while ten to thirty proposals compete. Lean Ethereum, in Drake's account, understands itself as a governance batching strategy that combines longer R&D phases with subsequently bundled EIP packages. Source: Bankless interview „Ethereum Beast Mode — Scaling L1 to 10k and Beyond", 12 November 2025.

[217] The complexity diagnosis is prominent in Vitalik Buterin, „Possible futures of the Ethereum protocol", six-part blog series from October 2024, in particular the parts „The Purge" and the posts on the Lean Ethereum vision. It is continued in Vitalik's post-quantum roadmap of 26 February 2026 and in Lean Ethereum research posts on ethresear.ch in 2025 and 2026.

ware: KECCAK256 hashing for State tree accesses, 256-bit integer arithmetic for most computations, Gas accounting for resource control. All of this is optimally chosen for CPU execution but expensive in ZK circuits: KECCAK256 alone generates thousands of constraints per hash operation. Gas accounting is carried along with every instruction. The leanVM would from the outset choose primitives that represent favorably in ZK circuits. Hash-based functions such as Poseidon generate orders of magnitude fewer constraints in ZK circuits than KECCAK256. Field arithmetic works in the algebraic structure of the proof system rather than in 256-bit integer operations. Data structures dispense with running accounting. The costs of subsequent ZK adaptation, which today make up a large portion of total proving costs, thereby disappear. Together, these three dimensions shift the trust model: away from empirically established stability, toward mathematically secured correctness of the runtime environment. The Roadmap postulates the target state without presenting a concrete specification.

What leanVM is for the EVM, Beam Chain is for the Beacon Chain. Justin Drake sketched the consensus layer target state at Devcon SEA 2024, which has since been further developed on ethresear.ch and in Buterin's Roadmap of February 2026: hash-based cryptography instead of grown complexity, post-quantum security as design principle, and formal verifiability at all levels.[218] The substantive elaboration is carried by section 5.1-B. Both target states lie beyond 2029 and stand for the same design philosophy of simplification to a formally manageable minimum.

While leanVM and Beam Chain are still outstanding, the direction of development already shows in the current governance process. Two EVM extensions that on first glance appear as technical improvements stand under the caveat that a long-term RISC-V replacement would make them redundant. EOF (EVM Object Format, RES, status open) would have reformed EVM bytecode structure through code validation, static jumps, and function boundaries. In April 2025 the proposal was removed from Fusaka for several reasons, including the increasingly questionable long-term benefit under the RISC-V perspective.[219] EVM-MAX (RES) would make certain cryptographic basic operations such as elliptic curve multiplication and modular exponentiation considerably more efficient in the EVM, making smart contracts with signature verification or ZK proof veri-

[218] Justin Drake, „Beam Chain", lecture Devcon SEA 2024; follow-up discussions on ethresear.ch from late 2024; Vitalik Buterin, roadmap update February 2026.

[219] Fusaka Interop Testing Call, 28 April 2025. In addition to the RISC-V redundancy perspective, the ACDE discussion also documented the complexity costs of double-maintenance of legacy EVM and EOF as well as the lack of application-layer consensus as reasons. Cf. EIP-7692 (Meta-EIP, status stagnant); the individual EIPs (including EIP-3540, EIP-3670, EIP-4200, EIP-4750, EIP-5450) remained in Draft status.

fication cheaper to execute. This approach too, however, stands under the same RISC-V redundancy caveat. The post-design state already shapes the prioritization of current EIPs.

Ethereum's Answer: Role Separation

Other blockchain architectures have chosen different answers to the scaling question. One direction raises the hardware requirements for the validator role and thereby accepts a higher access threshold to network participation. Another reduces the number of validators and concentrates consensus formation on few actors. Solana and Monad are prominent representatives of the first direction, BNB Chain of the second.[220] Ethereum's answer is a third: role separation. The scaling by a factor of 100 that section 5.2 has unfolded only becomes possible because no single role must accomplish it alone. The core of this role separation, as 5.2-A anchored it in the protocol, can be captured in three parallel statements. The Validator verifies without recalculating. The Builder constructs without validating. The Prover proves without proposing blocks — thereby establishing itself as a standalone infrastructure layer whose elaboration section 5.3-A′ undertakes. None of these roles could carry the scaling alone. Only their interplay generates capacity growth without the decentralization of any single role falling below the hardware threshold. Decisive here is the load distribution: Builder and Prover may become capital-intensive because their outputs are verified by every Validator with minimal hardware. Validator and Proposer remain on consumer hardware because they alone verify proofs. This role separation is the EIP-invariant principle that holds section 5.2 as a whole together.[221]

The limit of the principle was already marked by 5.2-D, and the role separation reaches as far as it carries, to proposing (cf. section 5.2-D). This has a structural consequence that points beyond the capacity question. While verification and proving become architecturally democratizable through the architectures of 5.2-A through 5.2-D, Block Building remains tied to State retention — and thereby to hardware requirements that grow with every scaling level. The partic-

[220] Solana operates with approximately 1,000 validators, of which more than half, per Justin Drake (Ethereum Foundation), are concentrated in two data centers located a few dozen kilometers apart. Monad operates with approximately 100 validators. BNB Chain has operated with 41 active validators in a two-tier structure since 2024, formerly 21. Source: Bankless interview „Ethereum Beast Mode — Scaling L1 to 10k and Beyond", 12 November 2025.

[221] The consolidation of the four roles (validator, builder, proposer, prover) as an EIP-invariant principle is an analytical synthesis of the individual mechanisms anchored in the protocol in 5.2-A: EIP-7732 (ePBS) for the proposer/builder separation, EIP-7928 (BALs) for builder parallelization, EIP-8025 (zkEVM) for the prover role and the decoupling of the validator from Gas capacity. The triad as a rhetorical condensation does not appear in this form in the primary sources. It is an analytical condensation by the author.

ipation architecture of the system becomes asymmetric: those who verify need less. Those who build need more.

Scaling raises a question that goes beyond capacity: two questions remain open — the one concerning the verifiability of the scaled system, and the one concerning participation in it. The answer lies in the architecture of verification, in the cryptographic basis that carries it, and in the access threshold it sets.

5.3 Verification and Access

Validating means: actually recalculating computations. A validator validating a block executes the block's transactions itself, with the same EVM logic, the same Gas costs, the same State transitions that the block builder followed in construction. Verifying means something more general: checking whether a block is correct without necessarily calculating oneself. In Ethereum's historical architecture, both concepts coincided, because the protocol knew no other form of verification than re-execution by each individual verifier. Anyone wishing to check a block therefore needed the hardware that re-execution required, and that hardware grew with the Gas Limit.

The zkEVM breaks this identity. It replaces re-execution with the verification of a cryptographic proof whose effort remains independent of block size. Verification becomes possible even where the verifier neither wishes nor can follow the computations. A second coupling remains in place, however. Even zk-based verification needs State as input. The proof confirms: if the initial State was X and transactions Y were applied, the final State is Z. But one who does not know that the initial State was X cannot embed the proof's statement in the world. The verifier needs access to the State values that the block reads and writes. As long as they must hold the full State locally to access these values, their participation depends on the storage capacity that State occupies. This second coupling is the subject of section 5.3.

The section develops the answer in three steps. First, a new State data structure reduces the volume of data a verifier needs at all. Then, post-quantum-resilient procedures secure the cryptographic basis on which this reduction rests. Finally, a more flexible account layer lowers the access threshold from which an ordinary actor can actually use the system. The sequence is mandatory. The cryptographic basis first protects a reduced data volume. The reduction first becomes reliable through a long-term stable cryptographic basis. The access threshold decides whether anyone outside professional infrastructure actually uses the ver-

ification. The first step begins with the data structure in which Ethereum holds its State.

5.3-A The New State Data Structure

Ethereum's State encompasses all account balances, contract storage values, and bytecodes of the network. The protocol stores these values as a tree with several hundred million leaves. Each leaf contains a concrete value — for example, the token balance of an account or an individual storage slot of a smart contract. The root of this tree is the State Root, a cryptographic summary of all leaves. The block header carries this State Root. Anyone wishing to check an individual value — for example, the token balance of a specific account — does not need the whole tree. They need a witness: a proof path from the corresponding leaf to the State Root. A verifier checking the witness against the State Root can establish the correctness of the value without knowing the other leaves.

The size of a witness depends on two properties of the tree structure. First, the branching factor — the number of children each node in the tree has. With a tree of branching factor 16, each node has 16 children, and the path from root to leaf must carry all 15 sibling nodes at each branch, otherwise the node computation cannot be reproduced during verification. Second, the commitment scheme: the procedure that forms from the values of a node the cryptographic anchor of that node. The commitment scheme determines the form in which the sibling nodes must be carried at all.

Ethereum's current data structure, the Merkle Patricia Trie (MPT), has a branching factor of 16 and uses Keccak-256 hashes as its commitment scheme. The average witness size per account access is roughly three kilobytes, in the worst case roughly nine.[222] Verkle Trees (EIP-6800, RES, Stagnant) change both properties. They increase the branching factor to 256 and replace the Keccak hashes with Pedersen Commitments with Inner Product Arguments. Pedersen Commitments combine many State values into a single short anchor. Anyone wishing to check an individual value can do so against the anchor without needing to know or provide the other values. Concretely the commitments operate on the Bandersnatch curve, an elliptic curve over the BLS12-381 scalar field. Inner Product Arguments generate a single compact proof for the path to the State Root instead of listing all sibling nodes individually. The witness size no longer

[222] Witness sizes for the Merkle Patricia Trie: approximately 3 KB per account access on average, in the worst case up to approximately 9 KB; block witness in the worst case up to approximately 18 MB. Cf. Ethereum Foundation Research, Verge Documentation, and EIP-6800, Rationale section.

depends on the depth or breadth of the tree. The average witness size per account access falls to roughly 200 bytes, the block witness in the worst case from roughly 18 megabytes to roughly 1.2 megabytes. This corresponds to a reduction by a factor of fifteen to forty-five.[223]

The witness reduction decouples verification effort from State volume. What a verifier must check no longer depends on the size of the State tree but remains bounded for each individual value to the compact path to the State Root. What hardware consequences a Stateless Client draws from this, the next section addresses. Verkle Trees play a dual role beyond the verification function. They are the State data structure on which later architectural refinements build, and they are the witness form that Native Rollups need for their EXECUTE precompile, because MPT witnesses would be prohibitively large for the purpose.

Verkle is deployment-ready. Kaustinen-7 has been running since late 2024 as a dedicated Verkle testnet in multi-client configuration, with Geth, Nethermind, Besu, and EthJS on the execution side and Lighthouse, Lodestar, and Teku on the consensus side.[224] Reth and Prysm are absent — these are gaps in the multi-client picture, but not blockers of the path. EIP-7612 and EIP-7748 provide named migration strategies with measured benchmarks. When governance assigns the headliner slot, Verkle can be activated in one of the next forks. This path is stalled, however: Verkle carries the designation Stagnant because the Ethereum Foundation has prioritized the long-term direction toward Binary Trees and deployment readiness has not yet translated into a fork assignment.

Verkle or Binary: The Quantum Question

This implementation readiness conceals, however, a cryptographic weakness that carries no weight under classical hardware. Pedersen Commitments and Inner Product Arguments rest on the ECDLP over the Bandersnatch curve. The ECDLP states: given a point P and a point Q on the curve, find the number k such that Q equals k times P. Classical computers cannot find this number for the curve sizes used in any practicable time. A sufficiently powerful quantum computer can compute it directly using Shor's algorithm, in a runtime that grows polynomially with key length rather than exponentially. The ECDLP falls, and

[223] Witness sizes for Verkle Trees: approximately 200 bytes per account access on average, block witness in the worst case approximately 1.2 MB. The reduction by a factor of 15–45× refers to the comparison avg-to-avg and worst-to-worst between MPT and Verkle. Sources: EIP-6800; EF Verkle Tree Documentation, as of February 2026.

[224] Implementation status Kaustinen-7 per February 2026: Geth (work in progress), Nethermind (Stateless Client complete), Besu (Java libraries and Bonsai interface), EthJS (implementation active); on the consensus side Lighthouse, Lodestar, Teku. Reth and Prysm without Verkle implementation; Erigon implemented but not active on testnet. Source: EF Verkle Implementation Tracker.

with it fall Pedersen Commitments and IPA. A research paper from the Dagstuhl context records that Verkle Trees are not secure against quantum computers, and the design of dynamic vector commitments that would be simultaneously asymptotically optimal, practically efficient, and post-quantum-resilient remains an open research question.[225] Buterin characterizes Verkle against this background as an interim solution that buys time until STARK-based and related procedures are mature.[226]

Binary Trees (EIP-9257, RES) replace Bandersnatch and IPA with STARK-compatible hash functions. Hash functions cannot be accelerated by Shor's algorithm. Effective security halves at most under Grover's algorithm, which can be compensated for by sufficiently long hashes. The implementation status lies considerably further behind, however. There is no multi-client testnet at the level of Kaustinen-7 and no migration EIPs at comparable maturity. Estimates from the research community calculate that a switch to Binary Trees would shift the verification and L2 guarantee strand by twelve to eighteen months.[227]

The EF Protocol Priorities Update of 18 February 2026 names for the first time "a move to binary trees and statelessness" as long-term direction. The Ethereum Foundation thereby marks a strategic prioritization toward Verkle. Guillaume Ballet, one of the EIP-6800 authors, subsequently characterizes Binary Trees as "a serious alternative."[228] Two scenarios are open. In one scenario, Verkle is deployed as a bridging solution, with long-term architecture on Binary Trees basis. In the other, the Roadmap revises the Verkle deployment decision in favor of a direct Binary path. The tension of this decision can be concretely named. Verkle is today deployment-ready but does not carry long-term stability because its cryptographic basis falls before quantum computers. Binary Trees carry long-term stability but are today not deployment-ready, because their implementation lags the Verkle path by twelve to eighteen months. The Roadmap thus decides on the relationship between short-term availability and

[225] Boneh, Bünz, Fisch et al., Dagstuhl reports on vector commitments and post-quantum security; cf. also Nomos Research, „Verkle Trees and Quantum Resistance", 2024. Paraphrased: Verkle Trees are not post-quantum secure, and the construction of dynamic vector commitments that are simultaneously asymptotically optimal, practically efficient, and post-quantum resilient remains open.

[226] Vitalik Buterin, „Possible futures of the Ethereum protocol — The Verge", 2024; roadmap update of February 2026. Verkle is characterized therein as a transitional technology whose function is to enable statelessness early while STARK-based and post-quantum-secure schemes continue to mature.

[227] Estimate from the EF research environment; see ethresear.ch discussion on Binary Trees vs. Verkle Trees, February 2026. The range of twelve to eighteen months is a consolidated assessment from multiple posts, not an officially quantified roadmap surcharge, and should be read as an order of magnitude.

[228] Ethereum Foundation, Protocol Priorities Update, 18 February 2026. Guillaume Ballet, ethresear.ch post on Binary Trees as a serious alternative to Verkle, February 2026.

long-term robustness. This decision fits into the Lean Ethereum program, which understands post-quantum security as a design principle.

The Migration without Downtime

The migration path to the new State data structure is conceived as a bundle of three EIPs that remains largely invariant with respect to the Verkle-vs.-Binary choice. EIP-7612 (Overlay Tree, RES, without fork assignment) solves the problem of the State switch without mainnet downtime. Upon fork activation, the protocol activates the new Tree as an overlay parallel to the existing MPT, and the MPT is frozen as read-only. All State writes run from this point into the new Tree. Reads search there first, fall back to the MPT if needed, and migrate the found value on the next write. The big-bang switch, which would have required an offline period of roughly one month, was discarded.[229]

EIP-7748 (State Conversion, RES, without fork assignment) governs the background migration. Per block, the protocol converts a fixed number of key-values from the MPT into the new Tree. Benchmarks on an AMD Ryzen 7 5800U measure roughly 300 milliseconds for 10,000 keys per block, from which a total conversion of approximately fifteen days follows. The migration is temporary but plannable.[230] EIP-4762 (Statelessness Gas Cost Changes, RES, without fork assignment) maps the new access costs in the Gas schedule. Per code chunk and per storage slot, an access costs 200 Gas, with adjacent storage optimization for consecutive slots in the same 256-group. The average Gas increase is six to twelve percent.[231]

The three EIPs solve migration as a separate problem. Which commitment scheme is used is decided by the Tree choice, while the migration bundle determines how the switch remains manageable. Both questions are separately solvable. The complexity of this switching procedure is considerable. Guillaume Ballet has called the Verkle migration "on the scale of the Merge if not worse in terms of complexity."[232] This complexity lies in the implementation of the migration.

[229] EIP-7612 (Verkle State Transition via an Overlay Tree). Authors: Guillaume Ballet, Ansgar Dietrichs, Tanishq Jasoria, Gajinder Singh, Ignacio Hagopian. The rejected big-bang migration would have required approximately one month of offline conversion at mainnet state size.

[230] EIP-7748 (State Conversion to Verkle Tree). Benchmark: 10,000 keys per block in approximately 300 ms on AMD Ryzen 7 5800U. At a 12-second slot cadence this yields a total conversion in approximately 15 days.

[231] EIP-4762 (Statelessness Gas Cost Changes). Code chunk access (31 bytes) and storage slot access cost 200 Gas each. The adjacent-storage optimization reduces subsequent costs within the same 256-slot group. Average Gas increase: six to twelve percent (EIP-4762 impact analysis).

[232] Guillaume Ballet (Ethereum Foundation), cited in the ACDE discussion on the Verkle migration; cf. All Core Devs Call Notes, 2024–2025.

The State data structure is the first step of the verification architecture that section 5.3 develops. It reduces the data volume a verifier needs from the full State to individual witnesses. Which data structure Ethereum concretely deploys — Verkle or Binary — decides the long-term cryptographic robustness. The architectural logic of the reduction remains unaffected by this. Pedersen Commitments, Inner Product Arguments, and the Bandersnatch curve are not special equipment of the State migration. They belong to a broader layer of specialized cryptographic components that the protocol provides above the EVM. This layer the next section unfolds.

5.3-A′ The Extension Layer above the EVM

In the Ethereum protocol there are operations that are cryptographically expensive: generating proofs, verifying signatures, checking Blob data against commitments, and soon also aggregating quantum-resistant signatures. They cannot run efficiently in the EVM, because the EVM knows no instruction for a 256-bit modular multiplication or an elliptic curve point addition. It would have to build such operations from thousands of basic components whose Gas costs add up. A native implementation in the client software of the same operation, by contrast, completes in microseconds, since the CPU knows a direct instruction for multiplication of wide numbers. Exactly this bottleneck between EVM speed and native speed the extension layer addresses. Precompiles and provers execute expensive operations outside the EVM and bind their correctness back to the EVM via cryptographic proofs or consensus guarantees. Four functions of the target picture rest on it: zkEVM verification, the post-quantum migration of signatures, Native L1 Privacy, and the verification of the reduced State from 5.3-A. Where 5.3-A discusses two alternative Tree options, 5.3-A′ unfolds two cooperating components of the extension layer for the computational work on the reduced data volume. The term extension layer is an analytical designation of this study and not an established term in the sources. The primary sources discuss the on-chain precompiles and the off-chain prover pipeline predominantly separately. Here they are consolidated because they fulfill a common architectural function.

The extension layer consists of two parts: the on-chain precompiles and the off-chain prover pipeline. Off-chain here means: outside the consensus protocol, so that the prover is not incorporated into the slot cycle and its proofs flow into the protocol only after their generation. The on-chain precompiles extend the EVM with operations it cannot itself efficiently perform — for example, the verification of a BLS signature or a KZG commitment. The off-chain prover pipeline

generates for each block a cryptographic proof of its correct execution. A verifier can check this proof on-chain without having to repeat the entire block execution. Both pursue the same goal: to extract expensive operations from EVM bytecode execution. Precompiles do this for individual cryptographic operations, the prover pipeline for the entire block execution. The verification of a prover proof on the on-chain side itself uses precompiles as cryptographic building blocks. The two components have matured to different degrees. Today the precompiles are active on mainnet, while the prover pipeline in the design state of the Roadmap is not yet anchored by hard fork. The Prover as actor follows from the role separation of 5.2: who should prove the blocks is someone other than who builds, proposes, or validates them.

The precompiles are established today in a first stage and grow through future hard forks. Three operations are currently active and cover the cryptographic classes needed by the EVM in the current state. BLS signatures carry the validator aggregation of the Beacon Chain (EIP-2537, IMPL Pectra). KZG commitment verification checks the Blob data of the Rollups (EIP-4844, IMPL Cancun). secp256r1 verification opens passkey authentication for accounts (EIP-7951, IMPL Pectra). Two further components are planned in the Roadmap. EVMMAX (EIP-6690, RES) accelerates modular arithmetic, with the RISC-V caveat from 5.2-E. A planned family of post-quantum precompiles (RES, ETH2030 line) encompasses the lattice-based operations and STARK verification components.

The accumulation of the layer proceeds asymmetrically between protocol and application. Hard forks activate new precompiles discretely, in sharp jumps every six to nine months. Application in smart contracts diffuses by contrast continuously over years. A precompile can exist in the protocol for years before reaching the most widely used contracts. This is the normal speed at which smart contract adoption develops. The extension layer thus emerges as a slow concentration over several hard fork cycles.

The Prover as New Actor

The prover pipeline introduces a new economic actor into the protocol whose hardware requirements lie above the validator level and below the data center level. The EF profile for realtime proving fixes the threshold through two values: capital outlay of at most 100,000 US dollars and power consumption of at most 10 kilowatts, each for an on-premises setup.[233] The range lies above the validator setup, whose investment costs are in the low four-digit range. It lies below a hyperscaler data center, whose investments reach into the seven-digit range. The Prover thus belongs neither to the validator class nor to the data center class, but forms its own economic layer between the two. The threshold is not static, since the zkVM teams have reduced their GPU requirements by roughly three quarters in recent months. Investment costs shift into the range of actual home prover setups.[234] Two points still separate the prover pipeline today from the design state: the realtime proving threshold and the economic architecture of prover compensation. 99 percent of mainnet blocks should be proven in ten seconds, but mainnet reality today still lies above the target range. For compensation, the Roadmap specifies neither a reward in the protocol nor a complete solution outside the protocol. Section 5.5-B picks up the deficit under the aspect of MEV and Builder economics.

To avoid the chokepoint of a single zkVM implementation, EIP-8025 (zkEVM Optional Execution Proofs, active EF Roadmap) anchors an acceptance threshold. Three proofs from five different zkVM implementations must be present for a block before it counts in the full security level.[235] The threshold targets software diversity in zkVM implementation, while the actor diversity of the provers remains unaffected. A single prover can generate its proofs with three different implementations and thereby satisfy the threshold. What matters is not the market concentration of provers but the diversity of the software — such as SP1, RISC Zero, Pico Prism, or Zisk. The logic of today's client diver-

[233] EF profile for realtime proving: Ethereum Foundation Blog, „Shipping an L1 zkEVM #1: Realtime Proving", 10 July 2025 (Sophia Gold). The specification encompasses, in addition to the two constitutive values mentioned in the main text, three further parameters: a latency requirement of at most ten seconds for 99 percent of mainnet blocks, a proof security of at least 128 bits, and a proof size of at most 300 KiB without trusted setups. All values apply as minimum requirements for realtime proving in production.

[234] Hardware trajectory of the RTP cohort: Ethproofs 2025/2026 (HackMD, 6 December 2025, Will Corcoran). Pico Prism reduces its requirement from 64 to 16 RTX-5090 GPUs for 99 percent proving in under twelve seconds, lowering GPU costs from approximately $128,000 to approximately $32,000. The three-quarters reduction is not representative for all implementations of the RTP cohort but shows the upper bound of implementation progress achieved so far.

[235] EIP-8025 and 3-of-5 multi-prover mechanics: anchored in the EIP-8025 design (zkEVM Optional Execution Proofs) and developed in the Ethereum Magicians discussion on the L1-zkEVM roadmap (26 January 2026, Kevaundray Wedderburn, EF zkEVM team) as a diversity architecture.

sity thereby transfers to the prover pipeline without intervening in the market structure of provers. A bug in a single zkVM does not distort consensus, provided the majority of independent proofs do not reproduce it. Liveness remains unaffected, since in the 1-of-N model a single honest prover suffices for proof generation. Here too two points remain open: the maturity level of the zkVM implementations and the hard fork activation of EIP-8025 itself. The EF security milestones for 2026 provide three stages in which security and performance requirements are progressively tightened. Full 3-of-5 viability depends on the last stage.[236] EIP-8025 remains an architecture statement until hard fork activation and does not take effect in today's consensus rules.

Many Proofs, One Master Proof

The extension layer contains an aggregation mechanic that consolidates many individual proofs into a master proof. STARK proofs have a recursive property for this purpose: a superior proof packages the verifier logic of several subordinate proofs within itself. The verification effort of the master proof thereby remains approximately independent of their number. At the mempool level, a second aggregation stage is added, where the protocol collects the signatures of multiple transactions in a 500-millisecond cycle and consolidates them into a single aggregation proof.[237] The post-quantum migration of signatures first rests on this mechanic, since its unaggregated verification costs would otherwise exceed the transaction budget. The Roadmap quantification names the range from ECDSA to STARK: 3,000 Gas for a classical ECDSA signature, 200,000 Gas for an unaggregated hash-based individual signature, 10 million Gas for a STARK without aggregation. In the aggregation batch, the amortized per-signature costs return into the transaction budget.[238] Validation Frames (EIP-8141, PLAN, CFI Hegotá Non-Headliner) bring the aggregation vehicle into the transaction format and make it available to all application classes needing such ZK wrappers. Beyond the PQ migration, three further application classes rest on the same mechanic: Native Privacy wrappers, Data Availability Sampling proofs,

[236] Three security milestones 2026 (February, May, December): Ethereum Foundation Blog, „Shipping an L1 zkEVM #2: Security Foundations" (18 December 2025). The improvement in proof latency achieved over the three milestones is recorded in EF data as a reduction from 16 minutes to 16 seconds, corresponding to a 45-fold cost reduction.

[237] 500-millisecond tick as the mempool aggregation cadence: Vitalik Buterin, „Possible Futures of the Ethereum Protocol — Post-Quantum Roadmap", 26 February 2026. The source names the tick as a design anchor without specifying the exact interplay with the hard forks.

[238] Gas cost range ECDSA / unaggregated hash-based signature / STARK without aggregation: Vitalik Buterin, PQ roadmap, ibid. The amortized per-signature costs in the aggregation batch are described in the source qualitatively as "near-zero long-term" without giving a concrete per-signature quantification.

and application-layer ZK proofs of external protocols. The architecture is fully described in the design state but depends for the on-chain verification of the aggregated quantum-resistant proofs on the not-yet-activated PQ precompile family.

The extension layer follows a glue-and-coprocessor architecture in which the EVM coordinates data and calls while precompiles and provers handle the expensive operations. The picture from the hardware level is the relationship between CPU and GPU. The EVM behaves like a universal CPU: it can execute any operation but is not optimal for expensive specialized tasks. Precompiles and provers behave like specialized chips that execute quickly those operations that would be slow in the EVM. Buterin empirically substantiated the scheme in 2024 using an EVM trace. Of the 46,924 Gas of a simple ENS hash update, roughly 73 percent fell on a small number of structured expensive operations such as storage reads, writes, logging, and cryptography.[239] From the pattern, the direction of Roadmap development can be read. Efficiency work shifts outward: to precompile extensions for expensive cryptography and to the prover pipeline for expensive verification. A deeper opcode restructuring of the EVM itself is omitted. The withdrawal of EOF from Fusaka follows the shift, since an EVM-internal bytecode-format refactoring can be expected to yield less effect than the parallel development of the coprocessors.

The extension layer is thereby described as the architectural foundation on which the post-quantum migration from 5.3-D and the Native L1 Privacy statement from 5.1-C rest. The PQ migration itself is a cross-cutting phenomenon with four vulnerability classes that section 5.3-D individually unfolds. For the immediately following section 5.3-B it provides a central prerequisite. 5.3-A showed how the protocol reduces the data volume of a block: witnesses on Verkle or Binary basis replace the State tree path with a compact proof path. 5.3-A′ showed how the protocol reduces the computation volume of a block: STARK proofs replace re-execution of the block with constant verification. Both reductions work together, and their common consequence is an architectural shift. Anyone wishing to validate a block need neither download the complete State nor repeat the complete block execution. They verify a compact witness and check a short proof. This is exactly the prerequisite for Stateless Verification, whose hardware consequence section 5.3-B unfolds. The extension layer carries

[239] Glue-and-coprocessor architecture and EVM trace: Vitalik Buterin, „Glue and Coprocessor Architectures", September 2024. The trace refers to a transaction updating the IPFS hash of a blog ENS entry. The 73 percent figure covers EVM execution alone. Extended with base costs, the share is approximately 85 percent.

not only the cryptographic substance of the late Roadmap functions but also the computational prerequisite for the Stateless Clients of 5.3-B.

5.3-B The Hardware Consequence of Stateless Verification

A Stateless Client verifies blocks without holding State locally or repeating block execution — and thereby verification moves from a server task to a function running on consumer hardware. The prerequisite is provided by the two reductions of the preceding sections: witness reduction decouples verification from local State storage, proof verification decouples it from re-execution. The factor that hardware-side dominates full verification today thereby disappears — the obligation of local State storage created the hardware load of execution in the first place, over and above light consensus verification. The Stateless Client knows this hardware load no longer, because it receives State per block as a witness and checks correctness against the State Root without locally tracing the computation path.

The resulting hardware profile shifts the requirements of a verification node by three to four orders of magnitude. Each of the four hardware factors derives from one of the two reductions whose technical mechanics were presented in the preceding sections. Storage requirements fall from 2 to 4 terabytes to below 100 megabytes — a reduction by a factor of 20,000 to 40,000.[240] RAM falls from 16 to 64 gigabytes to below 1 gigabyte, by a factor of 16 to 64. The compute requirement falls from 4 to 8 CPU cores to a single core, because proof verification with constant cost magnitude replaces linear re-execution. Hours to days consume the initial synchronization today — the build-up of the local State database from the network — constituting the largest entry barrier for new node operators. It collapses to seconds, provided the client no longer builds historical State. The hardware class shifts thereby from professional server configuration to the spectrum of consumer devices: smartphone, browser tab, Raspberry Pi, and home server.

[240] Hardware reduction through Stateless Clients: storage from 2–4 TB to under 100 MB (factor 20,000–40,000), RAM from 16–64 GB to under 1 GB (factor 16–64), CPU from 4–8 cores to 1 core, sync time from hours/days to seconds. The reduction factors refer to the transition range between the current full-node profile and the projected stateless profile under Verkle or Binary Tree migration. They presuppose that the client no longer builds historical state locally and receives the state per block as a witness.

Only One of Three Roles

The reduction concerns, however, only one of the three roles that the protocol carries on the hardware side. The verifier in the strict sense follows the profile of the preceding paragraph and operates their node on consumer hardware, since they neither produce blocks nor order transactions. The Validator, who in the design state as zkAttester verifies proofs, falls in their compute requirement to the same profile. They retain, however, a hardware dimension that escapes the reduction: bandwidth — the network connection for block propagation. It scales linearly with the Gas Limit, since the block size itself scales. At a hypothetical limit of 5,000M Gas, blocks grow to an order of magnitude of 100 to 500 megabytes per slot.[241] Such a scale requires a network connection that exceeds current home connections beyond the top tier. The current validator recommendation is roughly 25 megabits per second — a connection that the scaling path shifts into the range of several gigabits. The compute requirement is thus decoupled from the Gas Limit, the bandwidth scaling with the Gas Limit persists, and it defines the hardware boundary of the design state for the attestation role. The Block Builder finally remains in the data center class. They hold the full State to extract MEV and compose block candidates in parallel, competing with other builders in a millisecond latency race. The competitive infrastructure of the role is today estimated at roughly 200,000 US dollars per month.[242] The naive reading that Stateless Clients lower hardware requirements across the board thereby overlooks the structural differentiation. The reduction applies to the verification side fully, to the attestation side with an open bandwidth question, to the builder side not at all.

Today the hardware layering of the three roles rests on an out-of-protocol construction that ePBS moves into the design state. MEV-Boost assumes the separation between Proposer and Builder hardware as an out-of-protocol sidecar and

[241] Bandwidth scaling with Gas limit: at a hypothetical Gas limit of 5,000M, blocks grow to an order of magnitude of 100–500 MB per slot. Current validator recommendation: approximately 25 megabits per second in the mainnet profile. The scaling path shifts this requirement into the range of several gigabits. The bandwidth requirement is the remaining hardware scaling question for validators, since the compute load is decoupled via zkEVM attestation, but bandwidth grows linearly with block size.

[242] Builder latency race and infrastructure costs: competition between builders operates in the late-slot phase in the millisecond range. Empirical data show a peak of winning bids between 2,000 and 2,500 milliseconds into the slot, with near-complete reorg safety for bids before 1,500 milliseconds and frequent orphan events for bids after 2,500 milliseconds. Monthly infrastructure costs for competitive builders are estimated at approximately $200,000, driven by geographic proximity to relays and proposers, high-performance compute clusters, and private order-flow acquisition. The hardware class of the builder role is therefore not defined by a single compute requirement. It arises from the combination of local state storage, compute capacity, low network latency, and private order-flow infrastructure.

as of Q1 2026 routes roughly 96 percent of all mainnet blocks.[243] ePBS relocates the same separation into the protocol and makes it a binding element of the slot sequence.[244] The hardware consequence of the relocation is the institutional safeguarding of what the architecture requires regardless. The Proposer function remains supportable on the consumer hardware profile, the Builder function remains in the data center class. The protocol step sequence thereby anchors the separation structurally, where today it depends on social pressure and market coordination. The institutional relocation shifts the decentralization requirement to the Proposer layer, while the Builder layer operates as a specialized market whose neutrality is secured elsewhere.

The Cloud Concentration Remains

With the reduced verification profile, the possibility of participation shifts from the data center to home hardware, without the current hyperscaler concentration thereby being resolved. In the current state, roughly 59 percent of hosted execution layer nodes run their service on the infrastructure of AWS, Hetzner, and OVHcloud.[245] The M2 threshold value of at most 50 percent applies to a single provider. The 59 percent are distributed across three providers and thereby do not violate it. The Roadmap contains no mechanism that directly addresses the concentration. Neither do protocol penalties for cloud deployment exist, nor does the system have an incentive layer that economically rewards geographically distributed home nodes. The concentration rests on operational advantages that the protocol cannot eliminate because they lie outside its reach. These include the guaranteed uptime of hyperscaler SLAs, the automated monitoring of validator infrastructure, the low latency to large builder nodes, and the convenience of a

[243] MEV-Boost share in the current state: approximately 90 percent of all mainnet blocks are routed through the MEV-Boost sidecar per Q1 2026, compared to 11 of 20 blocks at launch in September 2022. MEV-Boost assumes, as an out-of-protocol implementation of proposer-builder separation, today's separation between proposer hardware and builder hardware without protocol anchoring. Builder concentration of the three largest actors stands at approximately 93 percent in the same survey period. This finding belongs methodologically to the neutrality discussion in 5.5 and is not further elaborated here.

[244] ePBS as institutional separation of hardware classes: the ePBS mechanics (separation of block validation into two slot phases, consensus validation in slot N and execution validation in slot N+1, resulting proving window of 6–9 seconds) were fully presented in 5.2-A. In 5.3-B they appear as functional resonance after R7d, with focus on the hardware consequence of the separation: proposer role (consumer hardware-capable) and builder role (data-center class) are structurally separated at the protocol level. Sources: 5.2-A (ePBS as convergence point). The neutrality consequences of the separation (relay elimination, MEV internalization) are taken up in 5.5.

[245] Cloud concentration in the current state: approximately 59 percent of Ethereum nodes are operated on the infrastructure of AWS (approximately 35.5 percent), Hetzner (approximately 13.8 percent), and OVHcloud (approximately 9.7 percent). M1 threshold for III.4 cloud independence: no single cloud provider exceeding 50 percent. The data basis comes from the Ethernodes/Migalabs snapshots of the current-state survey period. The exact provider distribution fluctuates between snapshots. The order of magnitude of the concentration is stable. The roadmap contains, as of February 2026, no measure that directly addresses this concentration.

click setup compared to the manual operation of a home node. Stateless Clients lower the hardware threshold for home staking to a level that is technically achievable. They do not automatically translate, however, into a changed distribution of operations, since the operational frictions of home operation persist independently of the hardware profile. The architecture provides the prerequisite. The operational consequence is decided outside the protocol, in the choices of operators. The hardware consequence belongs to the diagnosis with the same clarity as the reduction factors of the verification profile, because otherwise the architecture statement would be confused with an operations statement that it does not carry.

For the Hardware Agnosticism of the design state, an asymmetric improvement results. On the verification side, the spectrum of hardware classes that can cryptographically verify an Ethereum block is radically expanded. Smartphone, browser tab, and embedded systems take their place alongside the professional server class, and the spectrum of acceptable verification hardware is in the design state no longer bounded by the Gas Limit. On the operations side, the concentration of node infrastructure on a few hyperscalers persists, since no protocol measure outweighs the operational convenience of the cloud against the friction of home operation. The reach of the protocol ends at the boundary of its mechanism, and the choice of operating environment remains a decision of actors that the protocol at most facilitates but does not prescribe. The hardware prerequisite for broader participation is thereby clarified. The question of where a Stateless Client obtains its witness and proof without the gained trustlessness being lost again through a new central data channel is carried by the next section.

5.3-C The Trustless RPC Architecture

The Stateless Client from 5.3-B can cryptographically verify blocks on a smartphone because it neither holds State nor repeats execution. What it does not itself produce, however, are witness and proof per block — so it must receive the data from the network, and the receive side is where the viability of the entire hardware democratization depends. The access channel today is the RPC endpoint, a server that operates an Ethereum full node and accepts requests via the JSON-RPC protocol. As of February 2026, roughly ninety percent of Ethereum RPC traffic runs through three providers: Infura from the ConsenSys environment, the independent Alchemy, and QuickNode.[246] When a wallet receives an RPC response, it relies on the truthfulness of the provider without its own cryptographic verification. This trust relationship cannot be dissolved by the cryptographic verification of the Stateless Client. The client verifies the data against proofs, but both — data and proofs — come from the same provider, so that a manipulated provider could internally forge both consistently. With the viability of verification on consumer hardware, the question of trustlessness thus shifts to data provisioning.

The Trustless RPC Architecture replaces the central provider relationship with three cooperating out-of-protocol components whose combination fulfills the function of an RPC endpoint with cryptographic verification. The Ethereum Foundation has listed the combination in its Protocol Priorities 2026 under exactly the name Trustless RPC Architecture.[247] Helios carries verification on the end device through cryptographic checking of Beacon Chain headers. The Portal Network handles the decentralized distribution of historical and State data that the client needs for verification. The witness availability from the State tree architecture, whose mechanics 5.3-A fully covers, finally provides the cryptographic attestations against which checking takes place. The three components each carry their function only on the presupposition of the others. Helios verifies only the block header and needs the attestations that the Portal Network supplies. The Portal Network transports only the data that become transmissible

[246] RPC provider concentration in the current state. As of February 2026, approximately ninety percent of all Ethereum interactions are routed through three to four commercial providers, dominantly Infura (ConsenSys environment), Alchemy, and QuickNode. The concrete trust vector arises doubly: the provider can, first, manipulate the delivered witnesses and state roots so that the client executes a cryptographic verification against a compromised input state, and, second, filter the client's view of the mempool and of incoming transactions. Source: EF Protocol Priorities 2026.

[247] Naming of the "Trustless RPC Architecture." The convergence of Stateless Clients, Portal Network, and light clients culminates in an overarching goal that the Ethereum Foundation names in its Protocol Priorities 2026 as the 'Trustless RPC Architecture'. The English proper name is retained in this work because it is used in the primary sources as a fixed designation and is not an analytical coinage of this work.

through the witness structure of 5.3-A at all. The witnesses carry only insofar as Helios has cryptographically verified the header status. The architecture thereby resolves the question of where the client obtains its data without the trust load being newly centralized.

Helios: Verification on the Device

Helios resolves the trust problem by making the authenticity of data locally verifiable so that the data server itself need not be trustworthy. The client downloads light client updates from any Beacon Chain endpoint, and each update contains the signed block header together with the aggregated signature of the sync committee. The signature is checked locally against the known committee membership, so that authenticity is established independently of the supplying server and the trust relationship shifts from the data server to verification in the end device. The implementation written in Rust and compiled to WebAssembly requires approximately two seconds for initial synchronization and has a binary size of thirteen megabytes, making Helios operable on a smartphone, in a browser tab, and on Raspberry Pi class hardware without requiring any blockchain storage of its own.[248] The trust model rests on the assumption that at least two thirds of the sync committee act honestly, with the committee consisting of 512 validators drawn randomly from the total pool every 27 hours who sign Beacon block headers via BLS aggregation.[249] The reach of verification ends at the consensus layer. Helios checks that the block header is signed by the sync committee. What Helios does not check is whether the transactions within the block were actually correctly executed — that is, whether the State transition follows from the inputs according to protocol rules. This execution correctness rests in current Helios on the assumption that validators fulfill their duty and have not signed a manipulated result — that is, on a social assumption about validator honesty without its own cryptographic guarantee. In the design state it is closed by a combination of two layers: the proof pipeline from 5.3-A′ supplies the correctness proof of the State transition, the Stateless Client architecture from 5.3-B verifies it on the end device. The implementation has been productively in use since 2024.

[248] Helios specifications. Light client from the a16z crypto research unit, implemented in Rust with compilation to WebAssembly, runnable in the browser and on mobile devices. Initial synchronization in approximately two seconds, binary size at thirteen megabytes, no local storage requirements. Multichain support for Ethereum Mainnet and OP-Stack-based Layer 2. Status DEPL, productive since 2024 and actively maintained.

[249] Sync committee mechanics. 512 validators are randomly selected from the total pool for a term of 256 epochs (approximately 27 hours). They sign every beacon block header during their term. BLS aggregation enables efficient verification of the aggregated signature. Data volume for Helios on the order of 25 kilobytes per period, security budget of sync committee verification 512 times 32 ETH equals 16,384 ETH.

The Portal Network complements the verification layer through decentralized distribution of historical and State data, whose necessity arises from the history reduction introduced with EIP-4444. EIP-4444 (DEPL since July 2025) allows Execution Clients to discard historical data older than one year, enabling full node operation on a two-terabyte hard drive.[250] The structural consequence of the reduction is an availability gap, because no individual node need hold the complete history any longer and the network shifts data responsibility from the individual to the collective. The Portal Network distributes historical and State data over a decentralized peer network in which each node holds part of the data. A discovery system routes requests to the respectively responsible nodes. Four sub-protocols cover different data classes. The History Network carries block bodies and receipts, the State Network account data and bytecode. The Beacon Network supplies the sync committee data on which Helios verifies, while a Transaction Gossip subnet ensures mempool access without local State. Four productive implementations (Trin, Fluffy, Ultralight, Shisui) are operationally active as of February 2026, with cross-client interoperability, without the Portal Network being anchored as a binding part of the consensus protocol.[251] Data responsibility thus lies with a distributed network outside the protocol core.

So that fully decentralized data distribution also holds up as the Gas Limit rises, EIP-7975 (PLAN/CFI Glamsterdam) upgrades the eth wire protocol to eth/70 and introduces pagination of block receipt lists, thereby lifting the ten-mebibyte transmission limit per devp2p message for architecture scaling. The devp2p protocol on which peer-to-peer communication of Execution Clients rests today permits messages up to ten mebibytes. The limit is not sustainable as the Gas Limit rises, since block receipts grow with transaction count. The computational threshold for exceeding the ten mebibytes according to the EIP text is 83 million Gas; the practical threshold documented in Erigon tests of January 2026 already at roughly 75 million Gas.[252] Anyone re-synchronizing historical blocks

[250] EIP-4444 (History Expiry Phase 1). Status DEPL since 8 July 2025, all five execution clients (Geth, Nethermind, Besu, Erigon, Reth) support pre-merge history pruning. Reduces storage requirements by 300 to 500 gigabytes per node and enables full-node operation on a two-terabyte disk. Phase 2 (rolling window for post-merge data) is designated PLAN, without timeline.

[251] Portal Network. Status IMPL. Discovery via discv5 (UDP-based protocol with ENR records per EIP-778). Content distribution follows a Kademlia-inspired DHT with XOR distance metric, in which nodes hold those contents whose IDs are closest to their own node ID. Four productive implementations: Trin in Rust (focus on History Network and State Network in alpha stage), Fluffy in Nim (History Network, Beacon Network, and Portal Bridge), Ultralight in TypeScript (browser and mobile focus, State Network in R&D), Shisui in Go (ecosystem diversification). Cross-client interop tests via Hive framework, network monitoring via GlaDOS. As of February 2026, not anchored in the consensus protocol as a binding mainnet component, but operationally out-of-protocol.

[252] EIP-7975 and the transition to eth/70. Status PLAN/CFI Glamsterdam (per April 2026, Checkpoint 9). EIP-7975 upgrades the eth wire protocol from eth/69 to eth/70 and introduces the pagination of block receipt lists as an integral component of the protocol update. There is no separate networking EIP for the pagination. The computational

from the network can no longer transfer the associated receipts in a single message, causing synchronization to fail at a pure transmission boundary. EIP-7975 resolves the risk by allowing a node to retrieve the receipt list of a block in multiple sub-requests — for example, the first hundred receipts in one request and the next hundred in a separate one — at which point the wire protocol reassembles the sub-transmissions into a complete receipt list. With this hardening, the scaling capacity from 5.2 becomes genuinely enforceable in the RPC path — meaning Gas thresholds above today's 60 million. Otherwise it would collapse at a transmission bottleneck long before the hardware reduction from 5.3-B takes effect.

Available, but not Integrated

The Trustless RPC Architecture is technically available today. Its effect on the real trust load depends on adoption by wallets and frontends, actors the protocol cannot compel. As of February 2026, none of the three market-dominant wallet solutions — neither MetaMask, Trust Wallet, nor Ledger — integrates Helios as the default for trustless verification.[253] The infrastructure is available. Integration into end-user software is outstanding. The diagnosis follows a pattern already seen in the Hyperscaler concentration in 5.3-B. Both phenomena affect operational layers above the protocol. The protocol can facilitate them but cannot dictate how they are populated. They differ in their subject matter. The Hyperscaler concentration concerns the choice of operator infrastructure by validator operators; the RPC provider diagnosis concerns the choice of frontend integration by wallet and dApp developers.

With the distribution of the trust question across three out-of-protocol components, and the resolution of the scaling question through wire protocol hardening, the verification architecture of the scaled network rests entirely on cryptography at the endpoint device. The cryptographic load-bearing layer of Helios in turn relies on the BLS signatures of the Sync Committee, whose security rests on the same ECDLP that was identified in 5.3-A as non-quantum-resistant for

threshold for exceeding the ten-mebibyte limit per devp2p message lies, per EIP text, at 83 million Gas; the practical threshold documented in Erigon implementation tests from January 2026 is already approximately 75 million Gas. eth/69 was deployed as a networking update for Fusaka (December 2025). Eth/70 is linked to Glamsterdam (H2 2026 probable). Sources: ethereum/EIPs repository, EIP-7975 specification; erigontech/erigon issue January 2026 on receipt sizes above 75 million Gas.

[253] Wallet adoption per February 2026. MetaMask, Trust Wallet, and Ledger support the account extensions deployed in Pectra, but do not integrate a light client for cryptographic verification of RPC responses as the standard connection path. Frontend adoption for light-client verification is a market phenomenon outside protocol design and lies beyond the reach of ACD governance.

the IPA commitments of Verkle Trees.[254] The verification architecture of Stateless Clients and the cryptographic foundation from 5.3-A thus share the same assumption and face the same migration requirement. 5.3-D unfolds the post-quantum roadmap as a migration that concerns both pillars of the architecture built so far.

5.3-D The Post-Quantum Migration of Verification

The preceding sections have shown how a block can be verified on ordinary consumer hardware — on a phone or laptop — without local state and without one's own re-execution. This verification is carried by the signatures and commitments the block contains, whose authenticity the verifier recomputes. Today's signatures rely on elliptic curve cryptography, whose security rests on the discrete logarithm problem. The private key allows the corresponding public value to be computed in fractions of a second, while the reverse computation becomes exponentially more expensive for classical computers with each additional digit of the key. The security of every signature protecting a block today rests on the asymmetry between easy computation and practically impossible inversion.

A sufficiently large quantum computer shifts the boundary at which inversion becomes impossible. Shor's algorithm breaks the ECDLP in polynomial time (cf. Section 5.3-A). Against growing classical computing power, a longer key has so far sufficed, but against a polynomial runtime this measure fails, because doubling the key length only moderately increases the attacker's effort. The quantum computer thereby invalidates the very assumption on which the security of these procedures rests, and the break reaches beyond any single one of them.

Not every form of cryptography rests on such an algebraic structure. A hash function conceals no solvable problem. Its security lies solely in that an attacker would have to try all possible inputs. Grover's algorithm makes this search faster for a quantum computer, but it remains a search — the computer continues searching rather than computing the answer. A faster search can be countered by enlarging the search space and extending the output of the hash function. This is

[254] BLS signatures and IPA commitments under the discrete-log assumption. BLS12-381, the pairing-friendly elliptic curve of validator aggregation and sync committee verification, rests on the discrete-logarithm problem in pairing-friendly elliptic curve groups. The IPA commitments of the Verkle Trees rest on Bandersnatch, a non-pairing-friendly curve whose security likewise rests on a discrete-logarithm problem on elliptic curves. Both assumptions belong to the same cryptographic class and are attackable in polynomial time by Shor's algorithm on a sufficiently large quantum computer. They are not identical, however, because BLS uses the pairing property and IPA does not.

precisely the difference from Shor's attack on elliptic curves, which computes the answer directly and therefore cannot be countered by any larger search space.[255]

The target state of verification therefore replaces the vulnerable procedures with hash-based cryptography. The late roadmap treats resistance as a design principle and introduces it under the name Lean Ethereum (Section 5.1). For the verification layer that Section 5.3 has built upon the restructured data structure and the Trustless RPC Architecture, the transition poses a twofold question. What remains to be examined is whether the architecture remains cryptographically load-bearing under it, and what a delayed migration would entail.

Four Vulnerability Classes

The threat extends beyond verification, and the roadmap divides it into four vulnerability classes, each affecting a different primitive at a different location in the protocol. The BLS signatures of consensus are replaced by hash-based signatures with STARK aggregation, the substance of which Section 5.5 develops. Equally affected are the KZG commitments of data availability, which follow a separate migration path with an unresolved downstream problem addressed in Section 5.2-D. The ECDSA signatures with which every account signs its transactions today transition to hash-based procedures via an opt-in vehicle. Section 5.3-E unfolds the mechanics. The fourth class, the proofs of the application layer, rests on the aggregation mechanics that recurs below as a hinge. The map shows a cross-cutting phenomenon, whose sections each bear one class.[256]

Of the four classes, verification is most heavily exposed, because it has purchased its trustlessness at the cost of two efficiency gains, both of which rest on the vulnerable class. The first is the compact Witness from 5.3-A, which compresses the proof path to the State Root into a short proof and makes the Stateless Client viable on consumer hardware. The second is the aggregated Sync Committee signature from 5.3-C, which allows Helios a single check against the entire committee rather than verifying each signature individually. Both rest on the

[255] The characterization of the two algorithms follows Vitalik Buterin, „Possible Futures of the Ethereum Protocol — Post-Quantum Roadmap", 26 February 2026, and established post-quantum cryptography. Shor's algorithm reduces the factoring and discrete-log problems to polynomial runtime, which is why longer keys provide no effective protection; Grover's algorithm delivers only a quadratic advantage for unstructured search, which halves effective hash security and is compensated by doubling the output length. The EF treats post-quantum security within the Lean Ethereum program as a design principle. Its substance is carried by section 5.1.

[256] Four vulnerability classes: Vitalik Buterin, PQ roadmap, ibid. The four classes comprise the consensus BLS signatures (home 5.5), the KZG commitments of data availability (5.2-D), the ECDSA signatures of accounts (5.3-E), and the application-layer proofs. The KZG migration is distinct insofar as the switch to STARK-based erasure coding loses the homomorphic property of KZG commitments and therefore maximizes one-dimensional sampling rather than supporting two-dimensional sampling. It is partly unsolved and is treated in 5.2-D as a distinct migration with its own open problem.

same discrete logarithm — the Witness via the Inner Product Arguments of the Bandersnatch curve, the signature via the pairing property of BLS12-381. Precisely what makes the smartphone a verifier is thus the most quantum-vulnerable part of the architecture.[257]

The hash-based alternative is robust against the quantum attack, but expensive for signatures. In place of BLS come the hash-based signatures of the leanXMSS line (RES), whose substance Section 5.5 unfolds. A single such signature unaggregated consumes more than sixty times the Gas of an ECDSA signature, and its storage requirement is orders of magnitude higher. If the protocol had to verify each signature individually once hash-based procedures move from the exception to the rule, these costs would burst the budget of a block.[258]

This is where the aggregation mechanics from 5.3-A′ intervene as a hinge and render the expensive path affordable again. Aggregation combines the proofs (cf. Section 5.3-A′), so that the per-signature prohibitive costs in aggregate fall to a manageable share per block. The same mechanics that carries the proofs of the application layer thereby also makes the hash-based consensus signatures viable. The hash-based State Tree achieves its quantum resistance by its own route — the switch to Binary Trees — whose cost lies in the larger Witness and which no signature aggregation reduces. The mechanics are fully described in the design state but, for the verification of aggregated proofs, depend on a family of precompiles not yet activated by a hard fork.[259]

---

[257] Verkle commitments and Binary Trees: the ECDLP vulnerability of the IPA commitments in Verkle is treated separately from the four vulnerability classes, which is why their grouping with the BLS signatures of the Sync Committee is a coinage of this section; both belong to the same cryptographic class, BLS12-381 via the pairing property (cf. 5.3-C) and the Bandersnatch curve via the ECDLP (cf. 5.3-A), without being identical. The Ethereum Foundation frames the switch to Binary Trees primarily as a scaling step and cites quantum resistance as a complementary effect.

[258] Cost range: the Gas values follow the roadmap quantification carried in 5.3-A′ (Vitalik Buterin, PQ roadmap, ibid.): approximately 3,000 Gas for a classic ECDSA signature compared to approximately 200,000 Gas for an unaggregated hash-based individual signature. The increased storage requirement of quantum-secure signatures (ML-DSA, SLH-DSA) is one to three orders of magnitude greater than that of an ECDSA signature (approximately 2.4 to 4.6 KB for ML-DSA and 7.9 to 49.9 KB for SLH-DSA against 64 bytes).

[259] Maturity of the aggregation mechanics: the mechanics are, per 5.3-A′, fully described in the design state, but for the on-chain verification of aggregated quantum-resistant proofs are bound to the family of post-quantum precompiles treated there, not yet activated by hard fork. Status RES.

Exposure Is Not Risk

Maximum exposure is not maximum risk, because what carries Ethereum through a delayed migration is the willingness of consensus to keep running without finality — not a cryptographic procedure. The Inactivity Leak, which enforces this willingness, is a mechanism proven on Mainnet for years: if the network can no longer reach finality, the non-participating validators gradually lose their stake until the remaining ones again constitute two thirds. The mechanism was designed for the case of validators going offline, and it accommodates the quantum case as well, without ever having been designed for it. If a procedure breaks before migration is complete, the chain continues producing blocks, and finality returns as soon as two thirds of the stake have migrated to the new scheme — an economic process without new cryptographic preconditions. The network thus loses the finality guarantee while the chain keeps running, in Buterin's formulation "keeps chugging along."[260]

The design state drastically compresses finality without abandoning the mechanism that secures the transition. Single-Slot Finality finalizes a block in the same slot in which it is proposed, but explicitly retains the Inactivity Leak.[261] The consensus layer finds its end state in the Beam Chain, whose substance Section 5.5 develops and which lies beyond the current decade.[262] The reform of the Leak for the migrated procedure belongs to this horizon, while the mechanism that secures the transition already stands on Mainnet in its current form.

The verification architecture of Section 5.3 thus rests on a cryptographic base that is exposed but not undefended. Witness and Sync Committee signature share the quantum fragility of their discrete-log foundations, but their hash-based replacement becomes affordable through aggregation, and the load-bear-

[260] Inactivity Leak and robustness statement: the Inactivity Leak is tested and operational on mainnet in the current state (status DEPL of the mechanism). It progressively reduces the stake of non-participating validators until the remaining ones regain a two-thirds majority. The robustness statement that the chain keeps running without a finality guarantee ("keeps chugging along") comes from Vitalik Buterin's PQ roadmap and is multiply attested. A Metaculus estimate referenced in the same post puts the probability of cryptographically relevant quantum computers before 2030 at approximately 20 percent.

[261] Single-Slot Finality and the Inactivity Leak: in the current state, the Leak engages as soon as the Beacon Chain has not finalized for more than four epochs (MIN_EPOCHS_TO_INACTIVITY_PENALTY = 4), and reduces the weight of inactive validators until the active ones again hold two-thirds of the stake (ethereum.org, „Proof-of-stake rewards and penalties"; eth2book, section 2.8.6). Single-Slot Finality finalizes a block in the same slot in which it is proposed, but explicitly retains the Inactivity Leak as a fallback mechanism that keeps the chain running if more than one-third of validators fail (Vitalik Buterin, „Epochs and slots all the way down", 30 June 2024; ethereum.org, „Single slot finality"). SSF carries the status RES. The exact adaptation of the Leak to slot finality is not specified.

[262] Beam Chain: complete redesign of the consensus layer with SNARK-based aggregation and finality shortened to a few slots, which absorbs Single-Slot Finality and MaxEB (Justin Drake, Devcon 2024; RES). The reform of the Inactivity Leak for the post-quantum-migrated scheme is listed as a sub-item of this redesign. Its substance and the consensus-side BLS migration are carried by section 5.5, the consensus end-state by section 5.1.

ing capacity of the transition lies in the end less in cryptography than in a consensus that keeps running without timely migration. What the account layer contributes to this picture — the programmable validation logic through which a user switches their authentication from ECDSA to a quantum-resistant procedure — is unfolded in the concluding section of 5.3.

5.3-E Account Abstraction and the Migration of the User Signature

The preceding sections of this chapter have shown how the network verifies an entire block on cheap hardware, without central services, and even when a quantum computer breaks today's cryptography. Each of these checks begins, however, one layer below, with the individual transaction, whose validity the protocol measures against the signature of the sending account. The final section of 5.3 turns to the lowest-level check — and thus to the user themselves, who maintains their account and signs their transactions. An account is today bound to a single secret, a private key that its owner must safeguard. Anyone sending a transaction proves ownership by signing it with the key, and the signing procedure, ECDSA, is the same for every account and firmly anchored in the protocol.

This rigidity carries a price borne by every user. If they lose the key, they lose the account and everything in it, with no possibility of recovery, since no one else can prove ownership. They cannot change their signing method, even if a more secure procedure were available, and they can only pay a transaction's fee in the network's own currency. The account is bound to a single key and a single procedure, and the fixed binding is the real barrier to broad use.

Account Abstraction dissolves the binding. It removes from the protocol the prescription of how an account validates a transaction, and gives it to the account itself, making the validation rule freely programmable. Everything else rests on this single shift in the protocol: the wallets and applications built on top can unlock an account with a fingerprint, recover it after loss, or migrate its signature to a quantum-secure procedure. How far programmability changes participation in Ethereum for the ordinary user, and where the gained ease carries its own price, is the question this section addresses.

The Device Becomes the Wallet

For the ordinary user, participation in the target state begins with a gesture they have long known from their phone. They unlock their account and confirm a transaction via fingerprint or passkey — the same procedure their device already uses to secure logins without a password, and whose key never leaves the device's secure element. In place of the character string the account holder previously had to write down and safeguard for years — and whose loss was irreversible — comes an authentication the user already commands from everyday interaction with their device. The account demands no new knowledge about keys and custody, and fits into a handling that requires no further explanation. The device they carry anyway thus becomes the wallet, and the separate installation of specialized software that currently stands between the user and the network is eliminated. The first contact with Ethereum no longer presupposes that the user has understood the logic of private keys and practiced their safekeeping.

If the user loses the device on which their key rests, their account remains accessible. They have previously bound it to multiple trusted parties of their choice — other individuals, their own secondary devices, or a specialized service — of which a predetermined majority suffices to jointly restore access, without any single one of them having sole control over the balance. They pay the fee for a transaction in any currency in their account, or not at all themselves, because the application they are using covers it on their behalf, so that acquiring the network's currency as a prerequisite for first use is eliminated. Multiple actions that the protocol today treats as individually confirmed transactions — such as approving an amount and then using it — are batched by their account into a single, atomic confirmation whose parts either all succeed or all fail. Once they have granted an application a narrowly defined authorization — for instance for recurring payments within a limit they set — their account acts within those bounds thenceforth, without each individual action requiring renewed confirmation. The sum of these conveniences yields a participation that, in its handling, approaches familiar services of the open internet, without adopting their centralized custody.

Each of the conveniences rests on the single shift the protocol makes, yet emerges only on the layer above it. The protocol relinquishes the fixed prescription that an account validates its transaction solely via an ECDSA signature, and treats the account instead as a Smart Contract whose logic for checking validity lies in its own code. Before executing a transaction from this account, the protocol hands it to the account's validation function, which decides according

to freely chosen logic whether the accompanying data authorize the transaction. This data no longer needs to be an ECDSA signature. The account's code can verify a signature generated via fingerprint, require the collaboration of multiple keys for recovery, or charge the fee to a third party — and which rules apply is determined solely by the account, not the protocol. What the user perceives as the experience is therefore shaped by the wallets and applications that write and maintain the code for them, while the protocol alone guarantees the freedom within which they do so. The boundary between what Ethereum itself delivers and what the layer above makes of it runs exactly at that point: the protocol provides programmable validation, and the wallet turns that into the account operable by the user.[263]

The Price of Ease

The gained ease carries its price at the same seam at which it arises. Binding recovery to trusted parties provides security against loss and simultaneously creates a dependency on their availability and integrity, whose collaboration is precisely what carries access in a critical situation. The assumption of fees by a third party requires their continued willingness. Managing the account through a particular application relies on its validation code having been written correctly and kept maintained. The convenience thus shifts part of what the user is solely responsible for in the current state onto providers whose selection the user makes themselves but whose behavior they subsequently have only limited ability to control. How far the low entry barrier ultimately holds depends on whether this layer reliably conceals the complexity it introduces — and that part of the answer lies outside what the protocol itself can guarantee.[264]

[263] The currently active foundation of Account Abstraction: ERC-4337 (Account Abstraction via a separate mempool, DEPL since March 2023) establishes smart contract accounts with programmable signature verification, Gas sponsoring, and batching at the application level. Over 40 million such accounts have been deployed. EIP-7702 (DEPL with Pectra, May 2025) elevates programmable validation to the protocol level by allowing an existing EOA to temporarily delegate to smart contract code; in the first week after activation, mainnet recorded over 11,000 authorizations, with wallet support from MetaMask, Trust Wallet, Ambire, and Ledger. Passkey authentication rests on the secp256r1 precompile (EIP-7951, DEPL in the Fusaka cycle, December 2025), which makes natively available the curve used by FIDO2/WebAuthn and the Apple Secure Enclave.

[264] The shift to the application level is built into the evaluation framework of this work: criterion II.4 explicitly frames net inclusivity as dependent on whether wallet providers and dApp developers abstract the grown conceptual complexity for the end user, which lies outside the protocol. The extent of current fee sponsoring is shown by ERC-4337 adoption: in 2024, approximately 87 percent of all UserOperations were sponsored by a paymaster, with a sponsored Gas volume of approximately 3.4 million US dollars.

The Voluntary Migration of the Signature

The same programmable validation that makes the account easier to manage permits it a further and more significant step: the switch of the procedure by which it forms its signature. The ECDSA signature to which every account is today bound rests on that elliptic curve cryptography whose security a sufficiently large quantum computer would break, making it the single largest quantum vulnerability directly affecting the user's own assets. Via a dedicated transaction format, the Validation Frames, an existing account can migrate its signature to a quantum-secure procedure — for example a hash- or lattice-based signature, whose only materially relevant property here is its significantly larger key.[265] The protocol does not prohibit the old signature. It places the new one alongside it as a voluntary option, so that each account chooses the timing of its migration itself. The per-signature initially prohibitive costs of the quantum-secure procedure remain affordable solely through the aggregation mechanics from 5.3-A′, which combines many such signatures into a single proof and reduces their verification overhead in aggregate to a small share per block.[266] The voluntary nature of this migration is the same one that keeps access low-threshold from the outset, because the account here too determines for itself which validation logic it carries.

Section 5.3 closes the arc from the verification of an entire block on ordinary hardware to the lowest-level check at which the protocol measures the signature of an individual account, and it shows the same movement at every level. The restructured data structure reduced the proof required for verification, and the unburdened hardware brought this verification onto the device of the ordinary user. The Trustless RPC Architecture freed it from centralized data services; the post-quantum migration secured its cryptographic foundation against the quantum transition. The lowest of the checks, that of the user signature, the programmable validation finally placed in the hands of the user themselves. At every level the protocol guarantees a load-bearing and quantum-safely migratable foun-

[265] Frame Transactions (EIP-8141, PLAN, listed per Checkpoint #9 of 10 April 2026 as CFI Hegotá Non-Headliner and placeholder commitment for native Account Abstraction) bundle in a new transaction type, among other things, programmable validation logic in place of the fixed ECDSA check and the native switch from ECDSA to quantum-secure schemes such as ML-DSA, SLH-DSA, or STARK-based signatures. They build on complete native Account Abstraction (EIP-7701, PLAN). A generic alternative path that admits secondary signature schemes at the protocol level is available in EIP-7932 (Secondary Signature Algorithms).

[266] The manageable costs of the quantum-secure signature rest on the aggregation mechanics unfolded in 5.3-A′. The roadmap quantification puts a classic ECDSA signature at approximately 3,000 Gas and an unaggregated hash-based individual signature at approximately 200,000 Gas, whose amortized per-signature share in the aggregation batch falls back to a manageable value. Source: Vitalik Buterin, „Possible Futures of the Ethereum Protocol — Post-Quantum Roadmap", 26 February 2026; cf. 5.3-A′.

dation, while usefulness for the user only arises on the layer above, which makes the foundation usable for people.

The voluntary migration of the user signature with which the section has closed carries, however, a consequence that reaches beyond verification. A quantum-secure key such as that of an ML-DSA signature is orders of magnitude larger than today's ECDSA key, and every account that adopts it deposits the larger key permanently in the network's state.[267] If many accounts migrate, the larger keys aggregate into a perceptible expansion of the State that no signature aggregation reduces — because aggregation reduces the verification overhead of signatures, not the permanently stored volume of data. This very expansion of state, which the migration of verification generates as its side effect, is the subject of the following section.

[267] On key and signature sizes and their state consequences: quantum-secure schemes (ML-DSA, SLH-DSA) produce significantly larger signatures, on the order of 2.4 to 4.6 kilobytes for ML-DSA and 7.9 to 49.9 kilobytes for SLH-DSA against 64 bytes of an ECDSA signature, as well as a higher State Bloat of approximately 59 times for ML-DSA, arising from the public key of the migrated account permanently stored in state. Tiered State and the state-tree migration mitigate the storage consequence, while the bandwidth question remains an open research question.

5.4 State and Storage Persistence

5.4-A Why State Is Different

The expansion of state to which the migration of the user signature pointed at the end of the preceding section is only the most recent pressure on a resource whose difficulty runs deeper than any individual load upon it. The structural finding formulated in 5.2-D — Execution has a breakthrough mechanism, Data has one, State has none — reveals a second level of asymmetry: the nature of the resource itself.[268] Execution and Data can be delegated: the zkEVM shifts the verification load onto a proof, PeerDAS shifts data availability onto a sampling procedure (cf. Sections 5.2 and 5.2-D). Both mechanisms allow individual nodes to perform only a fraction of the total work while still verifying completely. State does not permit this delegation. A Builder seeking to produce a valid block must hold the complete state — every account balance, every storage slot, every contract bytecode. A lower Gas Limit reduces individual blocks. The accumulated State is unaffected by this. Buterin puts this asymmetry in a formula: for Execution there is the zkEVM, for Data there is PeerDAS, for State there is no comparable instrument.[269]

Two bottlenecks limit State scaling, and only one is in principle solvable. The short-term bottleneck lies in database efficiency: today's client databases are not designed for multi-terabyte states, and every State write operation requires logarithmically growing tree updates.[270] New database designs and Block Access Lists can alleviate this bottleneck. The long-term bottleneck eludes this kind of technical solution. Synchronizing a State of several terabytes requires, even at perfect bandwidth efficiency, significant time, because the entire State Tree must be traversed and verified. With every doubling of State size, the number of actors willing and technically able to hold the complete State falls approximately to half. No database design can dissolve this relationship as long as a single Builder must hold the entire State locally.

This synchronization limit has direct consequences for the decentralization of block building. If only operators with server hardware can synchronize the

[268] Execution: zkEVM enables factor-1,000 scaling (cf. 5.2). Data: PeerDAS (DEPL) and prospectively Full Danksharding (RES) enable factor-500 scaling. State: short-term factor 5–30 through database optimizations and Block Access Lists, long-term no comparable mechanism.

[269] Vitalik Buterin, „Hyper-scaling Ethereum state by creating new forms of state", ethresear.ch/t/24052, 5 February 2026. Paraphrased rendering of the formulation.

[270] State writes require log(n) tree updates each with log(n) database operations. Once state size exceeds available RAM, I/O costs rise sharply due to accesses on secondary storage.

State, the circle of Block Builders narrows. This narrowing strikes exactly the role of the Block Builder that the target state anchors in the protocol through the separation of Proposer and Builder. The scaling hierarchy thereby reveals an asymmetry that places the decentralization assumptions of the scaling path above 500 million Gas under qualification.

What the State Consists Of

Ethereum's State comprises as of Q1 2026 approximately 430 GiB, but what is decisive for the building question is the composition of this volume.[271] 81.7 percent falls on Contract Storage, with ERC-20 token balances at 27.2 percent and ERC-721 data at 21.6 percent forming the largest individual categories.[272] The overwhelming part of this State is not actively used: approximately 80 percent has been untouched for more than a year. 63 percent of all storage slots were written only once.[273] More than half of all contracts have no storage slots at all yet still occupy space in the State Tree. The protocol makes no distinction between active and inactive State: both occupy the same storage, require the same synchronization, and generate the same holding costs for every Builder. A Builder seeking to produce a valid block synchronizes the entire State, including the 80 percent that no one has accessed in years. At a constant Gas Limit of 60 million, the State grows by approximately 100 to 125 GiB per year, yielding a volume of approximately 830 to 930 GiB by 2030.[274] This remains within the range of commercially available 2 TB SSDs. The critical threshold, however, comes earlier: empirical stress tests by the Bloatnet initiative identify a performance cliff at approximately 650 GiB, at which State access times increase by 40 percent and sync times extend exponentially.[275] According to the scaling model of Maria Silva, this

[271] State size approximately 430 GiB per Q1 2026, with a range of 410–450 GiB depending on client implementation and compression mode. Sources: Maria Silva (ethresear.ch/t/23476, November 2025) measures 340 GiB in May 2025 (uncompressed, Geth node, dedicated to state). Vitalik Buterin (ethresear.ch/t/24052, February 2026) gives current growth at approximately 100 GB per year. The Q1 2026 estimate follows from the May 2025 baseline plus two-stage growth per Gas limit pipeline (36M through June 2025, 45M July–November 2025, 60M from December 2025).

[272] Paradigm (2024): Contract Storage 81.7 percent, Accounts 14.1 percent, Bytecodes 4.3 percent. Within Contract Storage: ERC-20 token balances 27.2 percent, ERC-721 data 21.6 percent.

[273] Han / EF Research: approximately 80 percent inactive state (>1 year untouched), 63 percent single-use storage slots, 54 percent stateless contracts (without storage slots), ≥7.4 percent dormant state (abandoned protocols).

[274] Growth rates 100–125 GiB/year at 60M Gas limit. Sources: Vitalik Buterin (ethresear.ch/t/24052, February 2026) gives approximately 100 GB/year. Maria Silva's scaling model (ethresear.ch/t/23476, November 2025) yields 124 GiB/year (342 MiB/day times 365 days). The range reflects measurement differences between clients and compression modes. The historical Paradigm estimate of 31–72 GiB/year (March 2024, at 30M Gas limit) is no longer directly comparable with the current 60M Gas limit level.

[275] Bloatnet Initiative (cperezz.github.io/bloatnet-website), Carlos Pérez et al., 2025–2026. Empirical stress tests on a dedicated testnet identify a performance cliff at approximately 650 GiB: state access times increase by approximately 40 percent, memory consumption increases exponentially, sync times lengthen substantially. Maria

threshold is exceeded at a constant Gas Limit already in mid-2027; according to a linear extrapolation of the actual growth rate, around the turn of 2027/28. The roadmap does not foresee constancy.

The Gas Limit path presented in 5.2 transforms the calculation fundamentally. The exponential Gas Limit path under EIP-9698 would, with continuous validator voting, drive the limit to 500 million by mid-2027. By mid-2029 it would rise to five billion. Operationally, the actual thresholds depend on feature maturity and client benchmarks, so the timeline can shift by one to two years.[276] The correlation between Gas Limit and State growth is not linear: a historically observed correction factor of approximately 1.5 dampens the effect, because not every additional unit of Gas flows into State write operations.[277] Even with this dampening, the cumulative State reaches approximately 2.7 to 7.0 TB by 2030 without countermeasures. Already in 2028, in the conservative estimate, it exceeds one terabyte — in the aggressive estimate, nearly two.[278] 2 TB SSDs are no longer sufficient for Block Builders from that point on. Block Builders would need server infrastructure with 8 to 16 TB, of the kind common in data centers today and absent from the home offices of independent Block Builders.

A differentiated State model introducing new, more restrictive State types while leaving existing State untouched could limit permanent State to 1.2 to 1.5 TB.[279] Before the path there lies a dead end that must first be understood.

Silva's scaling model (ethresear.ch/t/23476, November 2025) projects that this threshold will be exceeded by mid-2027 in all three Gas limit scenarios (conservative 686 GiB, base 859 GiB, aggressive 1,080 GiB).

[276] EIP-9698 defines the exponential Gas limit increase path. Cf. 5.2 for the full presentation.

[277] Paradigm (2024): historically observed correction factor approximately 1.5 between Gas limit doubling and state growth increase, because part of the additional Gas flows into execution and pure calldata.

[278] Author's own projection based on the updated growth rates (note 7) and the EIP-9698 path (cf. 5.2). The projection sums annual state growth across the EIP-9698 increase steps, each corrected by the correction factor of approximately 1.5 described in note 10 (a Gas limit doubling produces approximately 1.5 times state growth, not 2 times). Starting point is the Q1 2026 state size of approximately 430 GiB (cf. note 4). The conservative estimate (2.7 TB) uses the lower growth rate (100 GiB/year at 60M Gas, scaled over EIP-9698), the upper estimate (7.0 TB) a growth rate of 125 GiB/year in the same scaling. Hardware requirement from approximately 2 TB: 8–16 TB SSD (server class).

[279] Vitalik Buterin, „Hyper-scaling Ethereum state by creating new forms of state", ethresear.ch/t/24052, 5 February 2026.

History as a Resolved Contrast

History shows, in contrast to State, what a resolved storage problem looks like in practice. History Expiry under EIP-4444 (cf. Section 5.3-C) enables all five Execution Clients to delete historical blocks, transactions, and receipts after a defined period.[280] The savings amount to 300 to 500 GB per node, giving Full Nodes on 2 TB disks room to spare again.[281] Before pruning they required over a terabyte. Archive Nodes in the Geth implementation reached 15 to 20 TB.[282] The Portal Network archives the deleted data in a decentralized manner across three independent clients, an architecture described in 5.3-C as a building block of the trustless verification infrastructure.[283] Those who need historical data retrieve it from the network. Rolling Window Expiry as the next phase extends this model to ongoing pruning and is anchored in the governance process.[284]

The contrast between History and State sharpens the finding for the overall architecture. History could be delegated: historical data can be pruned and retrieved from the network when needed, because no running process needs to hold it permanently. State is the exact opposite — it is needed at every block, by every Builder, in its entirety. The only storage problem Ethereum has so far structurally resolved was not a State problem. A direct attempt at solving State did exist: State Expiry was intended to let unused State lapse after a defined period. The attempt failed at a fundamental technical obstacle.

---

[280] EIP-4444: history pruning for pre-merge data. All five execution clients (Geth v1.16.0+, Nethermind v1.32.2+, Besu v25.7.0+, Erigon v3.0.12+, Reth v1.5.0+) have supported history pruning since July 2025.

[281] Pre-merge data savings: 300–500 GB. Full nodes after pruning operable on a 2 TB disk.

[282] Archive nodes (Erigon-optimized): 3–3.5 TB vs. 15–20 TB (Geth).

[283] Portal Network: DHT-based, three independent clients (Trin/Rust, Fluffy/Nim, Ultralight). Cf. 5.3-C for the role in the verification stack.

[284] History Expiry Rolling Window (EIP-4444 Phase 2): PLAN status in the governance process.

5.4-B State Expiry and Tiered State

The fundamental obstacle lies in the requirement for non-existence proofs.[285] When a user creates new State — at an address or a storage slot — the protocol must verify that nothing ever existed at that location in Ethereum's entire history, not only in the current period. Without State Expiry, this check is trivial: what is not in the current State Tree does not exist. As soon as the protocol allows State to expire, this check loses its foundation. A slot that appears empty could contain expired State that would be reactivatable at any time. Distinguishing between "never occupied" and "once occupied, then expired" requires a historical proof that the protocol cannot produce efficiently.

For accounts, a partial solution approach exists: CREATE3 would bind addresses to creation periods, so that an address can demonstrably only exist from year N onwards.[286] For storage slots, there is no comparable safeguard: an ERC-20 balance for Account X is stored under keccak256(..., X), a hash that is opaque to the protocol.[287] The protocol sees a slot. It does not see that the slot represents a token balance belonging to a concrete user. State Expiry for storage slots is therefore not implementable in a backward-compatible manner without completely rewriting ERC-20 logic. All existing token contracts — whose balances together account for 27.2 percent of the State — must be migrated.[288]

In September 2025, Buterin drew the consequence: State Expiry should be abandoned in favor of Partial State Nodes, which would be functionally similar, require no Consensus Layer changes, and be considerably more flexible.[289] The Ethereum Foundation kept the conceptual option open in December 2025, but a workstream, team, and discussion point in the All Core Developers Calls are missing.[290] State Expiry is de facto abandoned.

---

[285] The problem of non-existence proofs is identified in Vitalik Buterin, „Hyper-scaling Ethereum state by creating new forms of state", ethresear.ch/t/24052, 5 February 2026, as the central obstacle. Additionally: EF Research, State Expiry Design Space, 2024.

[286] CREATE3 (Address Period Mechanism): addresses that can demonstrably only be created from a certain period onward. Conceptual approach without EIP finalization.

[287] ERC-20 storage slots are computed via keccak256(abi.encode(address, slot)). The result is semantically opaque to the protocol: it sees a 256-bit key, not its meaning.

[288] Paradigm (2024): ERC-20 token balances 27.2 percent of Contract Storage.

[289] Vitalik Buterin, 19 September 2025, paraphrased: "Don't do state expiry, do partial state nodes." Partial state nodes require no Consensus Layer changes.

[290] EF Blog, December 2025: state expiry conceptually not excluded, but no active workstream.

Tiered State: The Reversed Lever

The failure of State Expiry explains why the only remaining approach takes a fundamentally different path. State Expiry was meant to automatically remove unused State from the active dataset. This proved impossible precisely as long as the protocol cannot distinguish between active and expirable State. Tiered State (RES, ethresear.ch post February 2026) starts at the opposite point: existing State remains entirely untouched. Instead, new, cheaper but more restrictive State types are introduced to channel future growth.[279]

The architecture opts for two extremes rather than a middle path and places three State types side by side. Permanent State encompasses accounts, core contracts, and all data that must be available completely and without additional retrieval costs at every block. Temporary Storage is reset to zero at the beginning of each period, intended for short-lived data with a monthly reset cycle. UTXOs drive this principle to the extreme: their expiry period is zero — they expire immediately upon consumption.[291] Existing contract code, existing balances, existing DeFi positions remain in the permanent tier. The growth limitation arises through the incentive structure: new applications choose the more restrictive types once permanent persistence is not required. No deletion of existing data takes place.

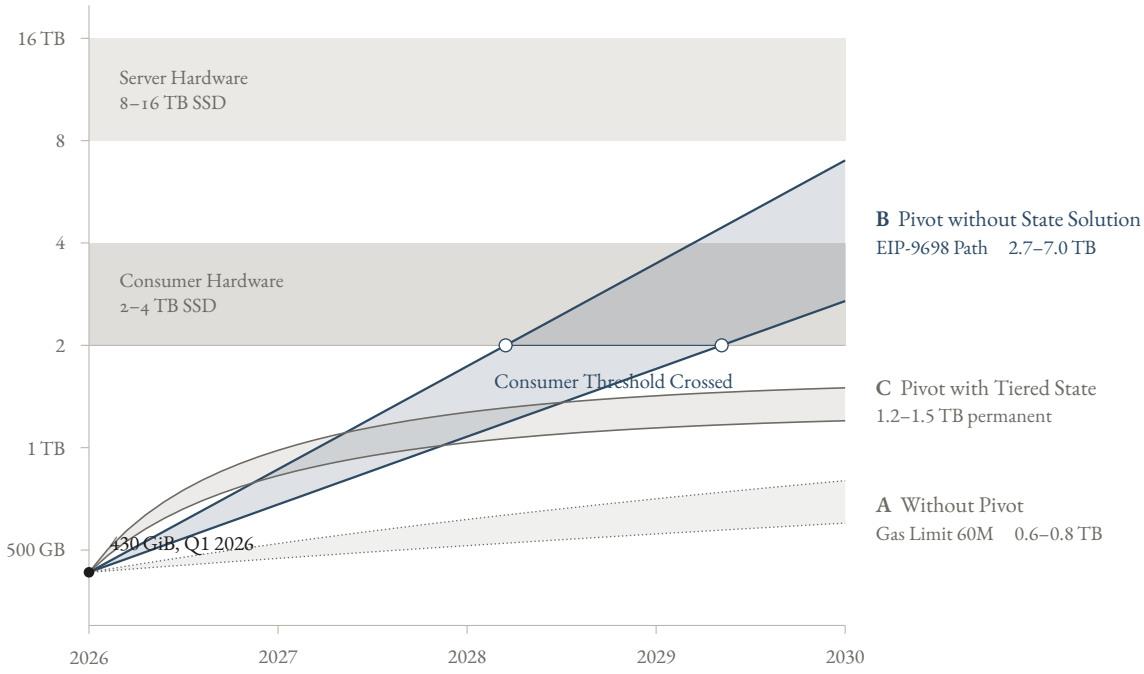

[291] ibid. Three state types: Permanent (accounts, core contracts), Temporary Storage (monthly reset), UTXOs (immediate expiry). Vitalik's fourth type (Strong-Stateless Storage Tree) was rejected in the same post.

Figure 5.4-B.1 State Growth Scenarios. Three projections: (A) Without pivot, constant Gas Limit 60M → 600–800 GB by 2030. (B) Pivot without State solution, EIP-9698 path → 2.7–7.0 TB by 2030. (C) Pivot with Tiered State → 1.2–1.5 TB permanent State by 2030. Reference lines: consumer hardware (2–4 TB SSD), server hardware (8–16 TB SSD).

The largest State consumers would be the first candidates for the more restrictive tiers. ERC-20 token balances claim 27.2 percent of the current State, ERC-721 data a further 21.6 percent.[292] Both could be assigned to the UTXO or Temporary tier, as could individual DeFi positions.

The question "in which tier do I store this data?" does not exist in today's Ethereum. It would become the central design decision for every new contract, comparable to the choice between calldata and Blob storage that Proto-Danksharding introduced for L2 operators. The calldata-Blob decision concerned a handful of Rollup teams. The tier decision would concern every Smart Contract developer in the ecosystem. No protocol upgrade can compel this migration: it requires the voluntary adoption of new token standards, the adaptation of existing toolchains, and the coordination of an ecosystem that has settled on permanent State as its default assumption.

A Pricing Model That Does Not Exist

Even if the ecosystem could manage this coordination, Tiered State would face unresolved technical dependencies. The multi-tier system presupposes an efficient State Tree structure, such as provided by Verkle Trees (EIP-6800, RES, Stagnant) or Binary Trees (EIP-9257, RES).[293] This State separation alone is insufficient: the system requires separate Gas dimensions for different tiers so that the economic incentives produce the desired allocation. Multidimensional Gas (EIP-8011, RES/DFI Glamsterdam) and Inter-Block Temporal Locality Gas Discounts (EIP-8057, RES/DFI Glamsterdam) are the technical prerequisites for this price differentiation, analogous to the Blob pricing template established by EIP-4844 for data blobs.[294] Blob pricing builds on a functioning EIP and a deployed mechanism. For State tiers, the pricing model exists so far only con-

[292] Paradigm (2024): ERC-20 27.2 percent, ERC-721 21.6 percent of Contract Storage. Cf. note 5 in 5.4-A.

[293] Verkle Trees (EIP-6800, RES, stagnant) or Binary Trees (EIP-9257, RES) as prerequisites for state separation in the Tiered State model. The EF Protocol Priorities Update of 18 February 2026 signals Binary Trees as the long-term direction. Verkle is deployment-capable but marked as a transitional solution due to the quantum vulnerability of the Pedersen commitments. Cf. 5.3-A for the full presentation of the tree transition.

[294] EIP-4844 (Proto-Danksharding) introduced separate blob Gas pricing with the Dencun activation in March 2024 and qualifies, per Chapter 2 definition (stable for over two years), as IMPL. Tiered State requires an analogous extension to state tiers. EIP-8011 (Multidimensional Gas Metering) and EIP-8057 (Inter-Block Temporal Locality Gas Discounts) were both listed for Glamsterdam as CFI and have been declassified per EIP-7773 as RES/DFI Glamsterdam. They remain research directions without fork assignment.

ceptually. EIP draft, reference implementation, and simulation of the market dynamics between tiers are missing. Tiered State presupposes Multidimensional Gas. Multidimensional Gas presupposes a specified pricing model. The pricing model does not exist.

A smaller precursor to this pyramid underscores the distance between concept and implementation. EIP-8032 (Size-Based Storage Pricing, RES/DFI Glamsterdam) would have coupled storage costs to actual storage size, delivering a first building block of the storage repricing architecture that Tiered State presupposes.[295] The EIP was scheduled for the Glamsterdam inclusion list and was downgraded via EIP-7773. The argumentative position is thus clear: the Tiered State architecture requires multiple repricings, and even the smallest intended building block did not reach the first fork cycle.

The ethresear.ch discussion from February 2026 shows that no consensus exists even within the research community.[296] The most conservative objection concerns sequencing: user MicahZoltu in the Ethereum Magicians forum demands that State expansion be allowed only once at least one user-operated RPC node can serve the complete State.[297] This condition is not met today, and State expansion itself would additionally hinder its fulfillment. Even if the sequencing were clarified, an architectural trade-off would remain: different tiers generate different access patterns, from which inferences about the use of a contract can be drawn. User igor53627 identifies therein a privacy risk that does not exist in current Ethereum.[298] The most radical counterargument questions the premise itself: user kladkogex argues that parallelization over node clusters would trivialize the State problem and that a multi-tier system therefore introduces unnecessary complexity.[299]

These objections meet a proposal at the earliest conceivable stage of development. Tiered State is an ethresear.ch post and is neither listed as an EIP draft nor on the agenda of the All Core Developers Calls. Research team, advocates in the client ecosystem, and a fork timeline are missing. Of the features addressing a core resource, none is at a comparably early stage. Tiered State is simultaneously

[295] EIP-8032 (Size-Based Storage Gas Pricing, RES) scales storage costs with a contract's state footprint and makes larger state entries more expensive. The EIP was intended for the Glamsterdam inclusion list and was moved to the Declined-for-Inclusion list per EIP-7773. Re-inclusion in Hegotá or a later fork is conceivable.

[296] ethresear.ch/t/24052, discussion thread, February 2026. Reviewers include: Guillaume Ballet, Marius van der Wijden, Justin Drake.

[297] MicahZoltu, comment in ethresear.ch/t/24052, February 2026, paraphrased.

[298] igor53627, comment in ethresear.ch/t/24052, February 2026, paraphrased.

[299] kladkogex, comment in ethresear.ch/t/24052, February 2026. Argument: horizontal scaling via node clusters as an alternative.

the only feature that offers any prospect of a long-term answer to the scaling-hierarchy asymmetry described in 5.4-A.

5.4-C The Timing Contradiction

Solution and problem thus move on fundamentally different timescales. The Gas Limit path under EIP-9698 drives the State from 2027 into territory where a management solution becomes structurally necessary. The only proposed solution has neither a specification nor a timeline. The Gas Limit, however, has both.

The Gas Limit is governed by validator voting, a decentralized voting process that can converge faster than the coordinated development of complex protocol changes. The most recent increase illustrates this speed differential: the Gas Limit doubled within a year (cf. Section 5.1-A), while Fusaka with PeerDAS as its core feature required the same period for specification, implementation, and deployment.[300] The EIP-9698 path provides for an exponential increase that would push the limit past 500 million by 2028. Validator voting for that requires no new specification, no client implementation, and no coordinated hard-fork deployment. Tiered State requires all of this and has not begun any of these steps.

The consequence is a twelve-to-twenty-four-month window in which L1 capacity overtakes the infrastructure for sustainable use.[301] Within the window, fragility increases: higher capacity is used, the State grows accordingly, but the mechanisms for differentiating and limiting this growth are absent. ML-DSA, the leading post-quantum signature candidate (RES), increases the pressure further: its signatures and public keys are several times larger than those of ECDSA.[302] Where the data is stored in State — for instance in Smart Contract wallets under Native Account Abstraction — this increases the State pressure, which the post-quantum migration described in 5.3-D and 5.3-E additionally intensifies. The window closes only once the prerequisites are delivered. The prerequisite chain from 5.4-B remains unchanged, alongside the voluntary adoption of new token standards by the entire ecosystem.

The temporary centralization within the window affects precisely the role that the target state anchors in the protocol through the separation of Proposer

[300] Gas limit 30M→60M: period February to November 2025, in four stages via validator voting (30M→36M→45M→60M). Fusaka: DEPL December 2025. Cf. 5.2-B for the full governance dynamics and EIP-7935 for the formalized default.

[301] Derived from the EIP-9698 path (Gas >500M by 2027–2028) and the earliest possible Tiered State deployment (no EIP, no timeline). The range reflects the uncertainty of both variables.

[302] ML-DSA (NIST FIPS 204) produces signatures on the order of 2.4 to 4.6 kilobytes against 64 bytes for ECDSA, i.e., approximately 38 to 72 times. The State Bloat is approximately 59 times, because the quantum-secure public key is permanently stored in state, for example under Native Account Abstraction. Cf. 5.3-D and 5.3-E for the full presentation of the post-quantum migration paths.

and Builder. When the State grows to two to three terabytes before a management solution takes effect, only operators with server infrastructure can synchronize the complete State.[303] The circle of Block Builders narrows at the very moment the protocol intends to open it through ePBS. The verification layer, decoupled from State growth through Stateless Clients and zkEVM, is unaffected. ePBS accepts professional building as a deliberate trade-off and protects decentralization via the Proposer set and FOCIL (EIP-7805, SFI Hegotá Headliner). The timing contradiction thus concerns censorship resistance: a Builder market with two to three dominant actors raises the question of whether the inclusion lists that FOCIL secures on the Proposer side can structurally withstand the building bottleneck.

SSZ as Incremental Contrast

The SSZ migration (EIP-6404, EIP-6466, EIP-6493, EIP-7807, all DRAFT) shows, by contrast, that protocol simplification at the State representation layer is achievable when the change proceeds incrementally.[304] SSZ replaces the informally specified RLP serialization with a typed, Merkle-provable format, thereby improving interoperability between the Consensus Layer and the Execution Layer. The Consensus Layer has used SSZ since December 2020. The EL migration is a multi-stage project with a time horizon of 2025 to 2026 and beyond.[305] SSZ does not resolve State growth. It reduces the complexity of State representation, eliminates RLP parsing as a source of error, and enables direct state proofs. The simplification is making progress because it is backward-compatible and does not alter any existing developer experience.

The contrast with Tiered State sharpens the finding. Tiered State is neither incremental nor backward-compatible in the same sense. It introduces a new design decision that every Smart Contract developer would have to make for every new contract. Added to this is the open prerequisite chain from 5.4-B. Ethereum can simplify State. Scaling State is a qualitatively different challenge that goes beyond serialization formats and intervenes in the core logic of the ecosystem.

[303] Conservative projection based on the scenarios in 5.4-A. Cf. note 11 for the full range (2.6–6.9 TB).

[304] SSZ EL migration (EIP-6404, EIP-6466, EIP-6493, EIP-7807, all in DRAFT status): transfers the serialization of the execution layer from RLP to SSZ (Simple Serialize), a typed, Merkle-provable format for all EL data structures.

[305] The Consensus Layer has used SSZ since December 2020. The EL migration is a multi-stage project with a time horizon of 2025 to 2026 and beyond and is considered a prerequisite for protocol simplification and the leanVM vision (cf. 5.1).

The interplay of Tiered State, Multidimensional Gas, and differentiated persistence levels would for the first time make it possible to couple State costs to actual usage patterns: permanent data at higher cost, temporary at lower, consumable at the cheapest rate. The economic logic would correspond to the Blob pricing template (cf. Section 5.4-B), extended to the resource the protocol has until now treated without differentiation. This interplay is the most fragile of all system combinations in Chapter 5, because its central component stands at the beginning of its specification. The fragility lies in the distance between design and implementation.

The prospect of Tiered State thus lies in giving State the architectural answer that makes the growth of a core resource sustainably load-bearing in the long term — an answer already delivered for History through pruning. As long as this answer is absent, State remains the dimension in which the roadmap cannot demonstrate a reliable solution.

Capacity scales, verification is democratized. The neutrality of the system under this expanded capability — in particular censorship resistance in concentrated block building — is not secured: whether the protocol guarantees against censorship and discrimination hold when economic incentives shift is answered by Section 5.5. The open flank of State enters that analysis. A system that expands its capacity faster than its protection mechanisms must actively secure neutrality.

5.5 Neutrality and Fairness

5.5-A The Three-Pillar Censorship Resistance System

Up to this point the chapter has examined the base layer from three sides, each of which has exposed a part of the target state. Scaling has shown that the layer can deliver far more than it does today, because it transforms the load of execution into a proof and replaces the holding of all data with a sampling procedure, so that a single block carries a multiple of its current volume (Section 5.2). Verification has shown that this expanded layer remains verifiable by everyone — on ordinary hardware, without central services, down to the signature of the individual account (Section 5.3). And State management has exposed the flip side: State is the one resource without a breakthrough mechanism, the bottleneck at which block production narrows to a few professional Builders, while verification distributes across thousands of ordinary devices (Section 5.4). So what remains at the end is a system that is more capable and at the same time generally verifiable, but whose production bundles in few hands. Up to this point the question was what the layer can do. Now the question is what it permits its most powerful actors. This section examines the neutrality and the fairness of the concentrated layer: its neutrality by the standard of whether a single powerful producer can still exclude a valid transaction or displace an already confirmed block retroactively from the chain. Its fairness by the standard of whether they can extract the value of block building exclusively for themselves. The answer of the target state shifts neutrality from the integrity of the operators to the rules of the protocol, which demands from them what it previously entrusted to them. The most fundamental of the three dangers against which the protocol enforces neutrality in this way is the exclusion of a valid transaction.

The exclusion meets a boundary that begins where verification ends, because it shifts the question of neutrality from the correctness of a block to its completeness. The cheap verification that the preceding part of the chapter brought to the device of the ordinary user reliably catches the invalid block, because thousands of validators recompute it and immediately reject any error in the state transition. It does not catch, however, the valid block that simply omits a particular address — because a fully correct block can bypass a disfavored transaction without ever violating a rule the protocol could verify. Verification ensures that a block contains nothing false, and it says nothing about whether it includes everything valid. What forces the inclusion of a valid transaction against the producer's discretion is the layer this section describes. It stands alongside verified correctness as its

necessary complement, because a concentrated producer could exploit the gap between the two that the protocol currently leaves open.

Three Points of Censorship

A valid transaction can be excluded from a block at three points, and the target state closes each of them through its own mechanism. The first point lies before the chain, at the off-chain relay through which the overwhelming majority of blocks run today, and at which transactions can be filtered according to external directives such as OFAC lists, even though the chain itself knows no such filtering. Eliminating the relay is accomplished by the Enshrined Proposer-Builder Separation, which lets the Builder direct its bids immediately and without a trusted intermediary to the chain, thereby removing the single point at which filtering outside the protocol can apply at all. Were the relay to remain, every protocol-anchored inclusion rule would be rendered ineffective, because its effect could be circumvented at the intermediary — so the separation of Proposer and Builder, scheduled for the earlier of the two upcoming forks, is less a building block of neutrality than its precondition.[306] The two remaining points lie within the protocol, at the Proposer's choice of block content and at the visibility of transaction content, and each demands its own natively anchored mechanism, which the following paragraphs develop.

The second point concerns the choice the Proposer makes over the content of their block, and it is closed by a forced inclusion list that prescribes which transactions they must include. This coercion is provided by the Fork-choice Enforced Inclusion Lists (FOCIL, EIP-7805), in which a committee of 16 randomly chosen validator-Includers, together with the Block Proposer, compiles in every slot — based on their own mempool view — a list of transactions to be included, with each list capped at 8 KiB. The guarantee rests on a 1-of-N assumption: a single honest Includer among the 16 suffices for an unjustly omitted transaction to make it onto the list and for the Proposer to be required to include it. Enforcement of the list is carried out by the Attesters, who withhold their vote from a block that excludes valid list transactions without the reason of insufficient space, so that disregarding the list costs the block its confirmation. The maximum delay with which an initially bypassed transaction nevertheless enters the chain falls

[306] Enshrined Proposer-Builder Separation (EIP-7732, PLAN, SFI Glamsterdam headliner; activation H2 2026 per EF Checkpoint #9, April 2026). It anchors role separation in the protocol and makes the relay obsolete by having the builder submit signed bids on-chain. In the current state, approximately 90 percent of blocks run through trusted relays; the OFAC compliance rate reached a peak of 79 percent in November 2022 and stands most recently at approximately 15 percent. The detailed treatment of the capacity and convergence function of ePBS stands in 5.2-A; here only the neutrality function counts.

from an empirical average of approximately 29 to 12 to 24 seconds. The strength of the 1-of-N construction is evident under regulatory pressure: even if that pressure compelled the majority of Builders to conform to a blocklist, a single non-conforming Includer forces the inclusion of the affected transaction. Among the three mechanisms, forced inclusion is the most mature, scheduled as the sole set Consensus Headliner for the upcoming Hegotá fork.[307] A formal safeguard of this forced inclusion against retroactive reorgs is still outstanding, so its protection for now rests on the general finality guarantees of the protocol and not on a rule specifically set for it.[308] How far a confirmed block can be displaced by such a reorg at all depends on Single-Slot Finality, to which this section turns later.

The third point is the visibility of the transaction content itself, because what the Builder can read before inclusion it can also deliberately bypass — and this information remains open once the first two mechanisms have taken effect. Visibility is addressed by a procedure that keeps the content of every transaction encrypted until its inclusion in the block. The Builder thus decides on inclusion without knowing the content, and cannot censor what it cannot see. It is carried forward under the name Universal Enshrined Encrypted Mempool (EIP-8105). Only after the order of transactions is fixed is their content decrypted, so that neither the Builder can make a content-based selection nor an observer can anticipate an order. The encryption closes a second gap, because with the concealed content, front-running and sandwich attacks also disappear — both of which are made possible today by reading the public mempool in advance. The draft from the Shutter environment and the approach developed under the name LUCID by EF research converge technically on the same Threshold Encryption. Among the three mechanisms, the encrypted mempool is the least mature, proposed for the Hegotá fork but without a fixed place in it, having failed to prevail against forced inclusion as the Headliner proposal.[309]

[307] Fork-choice Enforced Inclusion Lists (EIP-7805, PLAN, SFI Hegotá headliner; final ACD decision on 23 February 2026, the only confirmed consensus headliner of the fork, expected second half of 2026, with possible slip into early 2027 if Glamsterdam is delayed). A committee of 16 randomly selected validator includers and the block proposer (17 actors) produces inclusion lists of at most 8 KiB each per slot; the 1-of-N assumption requires only one honest includer; attesters deny votes to blocks that omit valid list transactions without a space reason; the maximum inclusion delay decreases from empirically approximately 29 to 12 to 24 seconds. Author: Thomas Thiery (EF Robust Incentives Group).

[308] Available Attestation (EIP-7942, RES, DFI Glamsterdam, moved to the Declined-for-Inclusion list per EIP-7773; re-inclusion in Hegotá or a later fork conceivable). The procedure adds an availability voting element to validator attestations, so that the chain only follows forks on which at least one-third of stake has attested data availability, making deep reorgs and long-range attacks practically impossible without substantial validator collusion. Critical for I.4 (protection against finality rollbacks) and II.1.

[309] Universal Enshrined Encrypted Mempool (EIP-8105 or LUCID, RES). Transaction content remains encrypted until inclusion and is only decrypted after the order is fixed, thereby also eliminating front-running and sandwich attacks. The mechanism proposed for Hegotá received no fixed fork slot and did not prevail as a headliner

The guarantee of forced inclusion rests on precisely the decentrality that Section 5.5-C marks as unsecured, because the 1-of-N assumption and the enforcing Attesters hold only as long as the Validator set remains sufficiently diverse. With the largest staker holding approximately 23 percent, the assumption holds comfortably — and how durably the precondition holds in the long run is examined in that section.[310] How far forced inclusion goes beyond earlier attempts is shown by the contrast with the Original Inclusion Lists (EIP-7547), the first and now discarded attempt at forced inclusion, whose rigid model lacked precisely the flexibility that the later 1-of-N committee provides.[311]

What the target state places alongside verified correctness is, with the three mechanisms, a forced completeness that the protocol demands from the producer where today it depends on their good conduct. It answers the first of the three dangers arising from concentrated production, while the extraction of value and the retroactive displacement of already confirmed blocks are reserved for two later sections.

5.5-B Value Extraction and the Decoupling of Validator Economics

The same concentrated producer whose censorship power the three mechanisms of the preceding section have curtailed continues to earn from the blocks it must now include a profit that none of these mechanisms captures. Forced inclusion prescribes which transactions enter the block, and leaves open the order in which the producer arranges them, even though the largest part of its profit comes precisely from the ordering. This section answers the question that begins at the open point: whether the producer can use the profit so gained to couple the economics of the validators to its concentration.

The profit in question arises solely from the selection and ordering of transactions, and the ecosystem refers to it as MEV (cf. Chapter 4).[312] As long as the profit accrues to the producer alone and grows with the volume of production, it

proposal against FOCIL; the Shutter design (EIP-8105) and the LUCID approach of EF research have technically converged (threshold encryption, including BLS threshold).

[310] The 1-of-N assumption and the enforcing attesters rest on the diversity of the validator set. The largest LST provider, Lido, holds approximately 23 percent of staked ETH in the current state. The M1 threshold for the highest neutrality level stands at under 33 percent for the largest individual provider, and this tolerance is the tightest of all II.1 indicators. The full treatment of staker concentration takes place in 5.5-C.

[311] Original Inclusion Lists (EIP-7547, DEP) were the first approach to enforced transaction inclusion and were considered too restrictive and inflexible. They were superseded by FOCIL with its 1-of-N model.

[312] Maximal Extractable Value (MEV). The value that the choice and ordering of transactions brings to a block producer beyond the regular fees. Front-running and sandwich attacks, made possible by reading the public mempool in advance, were treated in 5.5-A at the encrypted mempool; the present section treats the distribution of this value.

connects the concentration of building with an economic incentive that acts on the validators. The preceding section touched on the profit at one point, because the Encrypted Mempool deprived the Builder of the view of transaction content, and with the content-based selection, front-running was also eliminated. This section asks about its distribution: where it ends up after the build and whether its distribution binds the economics of the validators to the concentration of production.

The first answer of the target state is that building blocks does not become fair, because the incentives that bring a few professional Builders to the top of production persist and the protocol resolves none of them. A single Builder provides in the baseline state approximately half of all blocks, and the production market, by conventional concentration measures, is among the most highly concentrated of all.[313] Three structural causes keep this concentration stable, and none lies within reach of any single protocol rule, because they follow from the professional, vertically integrated, and multi-domain character of the market.

The Source of Builder Concentration

The first cause is the latency with which a Builder submits its bid in the auction. Builder competition runs as an auction within the fixed time window of a slot, in which bids are sharpened until the last moment, and the largest part of MEV comes from CEX-DEX arbitrage, whose price differential moves until the final instant. A Builder closer to the auction and trading infrastructure bids later and with more current prices and still arrives in time, while a distant Builder must bid earlier with stale valuations. The advantage is secured by low-latency infrastructure co-located at the trading venue, whose costs lie in the six-figure US dollar range per month.

The second cause is cross-domain arbitrage between the base layer and Layer 2 networks. The same assets trade at different prices across these domains, and a Builder or sequencer active on multiple simultaneously can atomically extract both sides of the differential, while a Builder confined to one domain sees only one side. This coordination requires operator rights or privileged access on multi-

[313] Builder concentration in the current starting state, reference point for the target-state assessment. A single operator — the commercial entity Titan with its own proprietary infrastructure — built approximately half of all blocks in February 2026 and approximately 47.6 percent in March 2026. The share measures the build identity of one market participant, not client software used by many validators, so the concentration falls on a single operator and is to be distinguished from the separate question of client diversity. The Herfindahl-Hirschman Index was determined at a Beaverbuild-Titan duopoly constellation in March 2025 at approximately 3,892, far above the DOJ threshold of 1,800. Individual values fluctuate, the classification as highly concentrated remains stable. Current-state finding: T. Wahrstätter (EF), ACDT #69 (Feb. 2026); CoinMetrics, State of the Network #356 (March 2026). The target-state assessment records that the concentration persists in the target state.

ple networks and advantages the already vertically integrated actors who already dominate the base layer. Their returns grow superlinearly with the number of domains controlled, because an entire class of atomic opportunities only opens once an actor controls multiple domains simultaneously.

The third cause is exclusive order flow, the privately routed stream of transactions. Approximately 54 percent of block value stems from transactions that bypass the public mempool and go directly to a single Builder via private RPCs, available only to them. Their source — a wallet, a trading bot, or an aggregator — entrusts this flow exclusively to one Builder against a paid arrangement and simultaneously protects its users from front-running, because the transaction never becomes public. Access is self-reinforcing, because a larger private inflow generates higher block valuations, more frequent auction wins, and greater profit margins, which in turn attract further flow.[314]

That production remains concentrated is not a goal the target state fails to achieve. The protocol treats concentrated block building as an accepted fact, because none of its measures attacks the three causes: ePBS eliminates trust in the relay and leaves the concentration incentives intact. Even the most prominent attempt to decentralize building nevertheless documents only its limit. The Builder network BuilderNet from the Flashbots environment lets multiple operators work together on a block without any of them seeing the others' contributions, and has thereby achieved a market share of approximately 27 percent — but operates outside the protocol and eliminates none of the three causes.[315] The concentration becomes a problem only where the value extracted from building enters the economics of the validators and binds the broad mass of stakers to the advantage of the concentrated Builder. The target state interrupts this transition at precisely two points, at which it separates the Builder's profit from the reward of the validators and returns its base amount to the commons.

[314] The three structural barriers of the current starting state that persist in the target state. Latency advantages (last-minute bids, geographic proximity to proposers, infrastructure costs in the six-figure US dollar range per month), cross-domain arbitrage (superlinear returns from stake across multiple domains), and exclusive order flow. An empirical study of Ethereum blocks from January 2023 to May 2024 puts the share of privately routed transactions in block value at 54.59 percent and demonstrates the self-reinforcing effect: builders with higher private flow value blocks more highly, win more frequently, and retain larger profit shares. The protocol addresses relay trust, not this value capture. Current-state finding: Private Order Flows and Builder Bidding Dynamics, WWW '25 (arXiv:2410.12352); on artificial latency arXiv:2312.09654.

[315] BuilderNet (DEPL, outside the protocol). A block-building network from the Flashbots environment in which multiple builders collaboratively work on a block with TEE isolation without seeing each other's contributions; market share approximately 27 percent. It demonstrates the operational possibility of decentralized building, does not remove the structural barriers, and lies outside the protocol.

Decoupling and Socialization

The first point is the decoupling, which separates the extracted value from the reward of the validator. In the baseline state, the randomly selected validator proposing a slot decides on the content of their block and simultaneously receives the value extracted from it, so that the value reaches the staker economy through them. The decoupling now separates the right to determine and propose the content of a block from the validator's remaining duties and auctions it off to a dedicated market of buyers. The ecosystem refers to this separation as Attester-Proposer Separation (APS). It takes two forms, discussed as Execution Ticket and Execution Auction, which are alike in awarding the right to the highest bidder rather than allocating it to the randomly selected validator. Their goal is to protect the broadly distributed mass of validators against the centralizing effect of the extracted value, without intervening in the build itself. Because the extracted value thereafter remains with the buyer of the right and no longer flows through the validator, the validator receives after the build solely the regular consensus reward — and the advantage of the concentrated Builder no longer reaches the staker economy in either form. APS builds on the protocol-level separation of Proposer and Builder that ePBS (EIP-7732) anchors in consensus.[316]

The second point is socialization, which returns the base amount of extracted value to the commons. The protocol burns the proceeds of the auction, thereby passing them on to no individual actor. It thus extends the fee burning in effect since 2021 under EIP-1559 to the extracted value.[317] Because the burned amount reduces the supply of ETH, it mathematically benefits every holder, while being withheld from the Proposer and the Builder. APS and MEV Burn interlock and do not compete: APS auctions the right to determine the block, and MEV Burn burns the proceeds of the auction.

The decoupling simultaneously eliminates an incentive that previously additionally favored concentration. In the baseline state, the Proposer's reward grows with the extracted value, and this value grows with every additional second during which trading takes place. A Proposer therefore has reason to delay publish-

[316] Attester-Proposer Separation (APS), in the forms Execution Ticket and Execution Auction (RES, part of „The Scourge", carried by the EF Robust Incentives Group; ePBS-dependent, expected from 2027). APS separates the right to determine and propose the content of a block from the other validator duties and auctions it, so that the validator reward becomes independent of the extracted value and predictable. It presupposes the protocol-level proposer-builder separation of ePBS (EIP-7732, PLAN, SFI Glamsterdam), whose capacity and convergence function 5.2-A treats; here only the decoupling of the extraction counts.

[317] MEV Burn (RES, ePBS add-on; from the environment of Justin Drake / EF). The auction proceeds are burned and not distributed to any single actor, analogous to the fee burning from EIP-1559 (DEPL since 2021). The measure smooths MEV spikes, makes the proposer reward predictable, and reduces the ETH supply.

ing their block in order to extract greater value. Such delay strategies — known in the ecosystem as Timing Games — shift the temporal dynamics of consensus in favor of actors with the greatest latency advantage.[318] Once the extracted value no longer flows through the Proposer, the reason for delay disappears and the economic cause is eliminated. The structural cause — the time window of the slot itself — is not eliminated until Single-Slot Finality, which a later section of 5.5 addresses.

What the target state achieves lies in the distribution of value and in the economics of the stakers, while production itself remains concentrated. The Builder continues to extract value, but can use it neither to harm the commons — whose share is burned — nor to make the validators dependent on its concentration — whose reward is separated from its advantage. The concentration of production remains a fact of the baseline state, but its effect on neutrality is interrupted at both transitions.

Both mechanisms stand differently than forced inclusion, which the preceding section carried. APS and MEV Burn are research projects without a fixed place in a named fork, supported by the EF Robust Incentives Group and the research around Justin Drake. They are thus the most immature of the responses the section assembles. They belong to the roadmap phase addressing systemic centralization risks, which the ecosystem refers to as "The Scourge." Their effect is therefore a direction with established mechanics and an open timeline.

The decoupling prevents production from passing its profit on to the staker economy, and does not reverse the concentration that already exists in the baseline state. The largest single staking provider holds approximately 23 percent of staked ETH, and whether this threshold holds durably is examined in the following section.[319] What MEV Burn touches beyond fairness is the security budget of the chain, because the burned value benefits the protocol rather than increasing individual actors' rewards, which intervenes in the economic sustainability

[318] Timing games. Strategies in which a proposer delays block publication to extract more value, since the extractable value rises monotonically with time in the slot. They are considered an economic neutrality violation in favor of actors with latency advantages. The decoupling of the extraction removes the economic cause; the structural cause via Single-Slot Finality and the ePBS delivery obligations (builder staking, PTC attestation) is treated in 5.5-D. In the short term, containment remains referred to coordination among participants outside the protocol, and research acknowledges that the current design does not fully exclude them.

[319] Staker concentration and security budget as a forward reference to 5.5-C. The largest individual staking provider (Lido) holds approximately one quarter of staked ETH in the current starting state. The February 2026 figure from Lido is 23 percent, Dune data cite approximately 24 percent, down from a peak of over 32 percent (2023). The M1 threshold for the highest neutrality level stands at under 33 percent for the largest individual provider, and this tolerance is the tightest of all II.1 indicators. The security budget is fed by staking rewards, MEV, and transaction fees; MEV Burn redirects MEV proceeds to the protocol and thereby intervenes in staking economics, the synthesis of which 5.5-C carries (gaps 7.2 and 7.10). Current-state finding: Lido tokenholder update (Feb. 2026); Dune (hildobby).

whose synthesis the same section provides. The target state thus answers the second of the three dangers arising from concentrated production, by separating its profit from the economics of the stakers and directing its base amount to the commons. The retroactive displacement of an already confirmed block, the last danger, remains reserved for the concluding section.

5.5-C Staking Economics and the Limit of Distribution

The guarantees of the two preceding sections depend on an assumption neither has examined: the Validator set remains sufficiently diverse. Forced inclusion needs only a single honest Includer in the committee for this (Section 5.5-A), and the decoupling of rewards protects an existing distribution of stakers without creating it (Section 5.5-B). The preceding sections examined the concentration of production — the question of who builds the block. Now the question is the concentration of stake — who validates at all and how many providers hold the staked capital.

Two questions follow, which diverge in their assessment. The first is economic and concerns the lasting sustainability of staking, on which the affordable security of the network depends. The second concerns distribution and asks whether the protocol prevents control over consensus from concentrating in few hands. To the first question the protocol answers with a stable yes. To the second, it does not. The first part shows how the economics holds. The second, why distribution remains open.

Economically, the protocol above all must sustain the security budget — the reward to stakers that keeps the cost of acquiring an attack share higher than the gain from it. Its level is composed of three quantities: issuance, transaction fees, and burn. Issuance pays for security, fees supplement it, and burn maintains the value of the paid-out ETH through scarcity.

Of the three quantities, the protocol sets only issuance itself, and only for it is the future level open. Every newly created ETH distributes the existing supply across more units, devaluing those already held — which is why the protocol should create only as much as security requires. Under the current rule, however, the distributed quantity grows the more ETH is staked in total, so that emission can grow without limit. The reform discussed under the heading Minimum Viable Issuance changes the rule: once the staked share exceeds a target, the annual return per validator falls, and further stake loses its incentive.[320] The

[320] Issuance Reform / Minimum Viable Issuance (RES, without governance anchoring; carried by Anders Elowsson and others). A reformed emission curve reduces validator returns at high stake ratios and targets an equilib-

reform seeks the share that is high enough for security and low enough to keep emission small. Its status remains open, since research is still debating the right target. Added to this is the fact that the narrative of scarce, deflationary ETH — which the ecosystem calls Ultrasound Money — no longer held in 2024 (and subsequently 2025) when the declining burn rate allowed supply to grow again.[321]

The other two quantities — fees and burn — stand on firmer ground than issuance. On the fee side, Blob Fee Stabilization (EIP-7918) secures a minimum price for the data that Layer 2s deposit on Layer 1.[322] It prevents Blob fees from falling toward zero and preserves the fee component for validators. On the burn side, the protocol already burns the base price of every transaction today, and in the target state additionally the MEV. MEV fluctuates heavily, since rare blocks with large liquidations or arbitrages yield several times that of an ordinary block. MEV Burn destroys it via the auction for block determination, so that no single block provides a Proposer with exceptional returns and validator rewards remain steady.[323] It counts here on the burn side of the budget. In Section 5.5-B it appears under fairness.

Returns exceed operating costs only with efficient validation, and here MaxEB (EIP-7251) takes effect, delivered since the Pectra upgrade in spring 2025. The upgrade raised the maximum Effective Balance of a validator from 32 to 2,048 ETH.[324] Large operators can since then consolidate thousands of smaller nodes into fewer larger ones, reducing their costs and steadying their returns. For the economics of staking, node consolidation counts as a stabilizing factor. What consequences the consolidation has for distribution is clarified in the second part.

---

rium issuance with a target stake ratio as a stationary state in which the security budget balances among stake, fees, and burn. The dispute over the correct target size is open; the "Ultrasound Money" narrative around the deflationary ETH burned through burn no longer held in 2025, after the burn decreased as activity shifted to Layer 2 networks and supply grew again.

[321] Attributable to the Dencun upgrade (March 2024), which drastically reduced transaction fees on Layer 1 and caused the ETH burn rate to collapse. For detailed statistics on supply growth, see: https://ultrasound.money.

[322] Blob Fee Stabilization (EIP-7918, DEPL with Fusaka, 3 December 2025). A minimum blob base fee, coupled to execution costs and set at approximately one-sixteenth of the execution base fee, causes the blob price to converge toward an economically rational floor at target utilization and keeps it above the absolute minimum of one wei, which supports the fee component of the budget against the fee-alignment problem.

[323] MEV Burn (RES, ePBS add-on, from the environment of Justin Drake / EF). The present section treats only the economic function as a component of the security budget, after 5.5-B has derived the socialization of MEV. The mechanism smooths MEV spikes and makes the proposer reward predictable.

[324] MaxEB (EIP-7251, DEPL, Pectra, May 2025). The increase of the maximum effective validator balance from 32 to 2,048 ETH allows the consolidation of many small nodes into fewer large clusters, reduces operating costs, and makes rewards more predictable. It acts on operational efficiency and on the reduction of the validator set (prerequisite for SSF), not on stake distribution.

From the interplay of the three quantities follows a self-stabilizing state. The staking rate settles at a value at which issuance just covers security. Burn offsets issuance to the extent that the network need not buy its security at ever-higher expense. This is the basis for the high economic load-bearing capacity of the network. An attack on a third of the stake requires acquiring approximately 11.5 million ETH. The buying pressure of such a quantity would drive the price high enough to make the attack uneconomical.[325] The threshold holds even if the reform modestly reduces the staked share, because the price driven by acquisition counts more than the quantity to be purchased. Security is thus load-bearing, but its dollar value depends on the ETH price and usage volume and is not guaranteed by the protocol alone.

The Unresolved Concentration

Everything described so far stabilizes the economics of staking and does not touch its distribution. No mechanism in the roadmap reduces the share that a single provider holds of the staked capital. Every existing mechanism operates on the economics or the operation of validation, and none shifts the market shares of providers. The limit of the section becomes visible as soon as the mechanisms are individually examined for what they address.

A technological approach to risk mitigation is provided by *Distributed Validator Technology* (DVT). This decentralizes the operation of an individual validator by distributing its tasks across multiple physical machines and different operators. The *Single Point of Failure* is thereby eliminated: neither the failure of individual servers nor the loss of cryptographic keys endangers validation or leads to Slashing. Implementations of the technology are already deployed via protocols such as Obol and SSV Network, including within the market leader Lido. DVT does not, however, affect market-structural centralization. Since the technology merely distributes the technical infrastructure behind the validation keys, the percentage market share of the staking providers in total capital remains unaffected.

For network distribution, MaxEB unfolds a dynamic contrary to its economic benefit. While raising the maximum validation balance optimizes operational efficiency, it simultaneously promotes consolidation of nodes by large operators,

[325] 34 percent attack and security budget in the target state. The security budget from staking rewards, MEV, and fees must incentivize staked capital on the order of 100 billion dollars. An attack on one-third of stake requires the acquisition of approximately 11.5 million ETH and remains uneconomical under realistic liquidity and price conditions. The static indicator value falls below the M1 threshold of 30 billion dollars on the reference date, but actual attack costs lie above this due to market illiquidity and slippage in acquisition. The indicator is met with qualification, price-sensitive, and not secured by protocol changes alone.

which has a structurally centralizing effect. A shift in market shares is not to be expected from the modified emission rate either, as this only marginally reduces the incentive for dominant actors. Furthermore, the minimum threshold of 32 ETH remains untouched, as MaxEB solely reforms the upper limit of the effective balance. This economic entry barrier continues to block the growth of solo staking, which still shows a market share of under one percent of aggregated stake.

A direct limitation of market shares could be realized through a protocol-level upper limit (*Staking Cap*) that anchors the maximum share of a single provider at the protocol level. Although such a cap is discussed within the community, it constitutes neither a planned nor a decided measure in the current target state and is consequently not in the official roadmap.[326] Such a cap would directly intervene in free-market competition, which corresponds to a regulation with which the Ethereum protocol has deliberately left the market to itself for reasons of neutrality. Herein lies the actual deficit of network distribution, since the problem of centralization is one of regulatory policy. Since the only structurally effective countermeasure — in the form of a protocol-level maximum share — is not anchored in the governance rules of the protocol, Ethereum currently lacks the protocol-level tool to prevent a monopoly-like concentration of staking power.

Furthermore, the approach of *Rainbow Staking* is still in the conceptual phase. The model aims to reactivate solo staking by dividing the validation system into two classes: capital-strong actors take on complex network security, while private small participants are involved as a control instance without expensive hardware requirements, which significantly lowers the financial and technical entry barrier.

This advancing market concentration pushes the protocol-side guiding principle of protocol neutrality to its conceptual limit. While the protocol can technically enforce certain behaviors — such as guaranteed inclusion of a transaction (Section 5.5-A) or the economic decoupling of rewards (Section 5.5-B) — it cannot stop the accumulation of market power by a shrinking number of staking providers. The technical measures thus protect the network from the acute misuse of existing power structures, but do not effect a structural redistribution of power. Beyond the internal capital allocation, there is also an external, geograph-

[326] Protocol-level staking cap and Rainbow Staking. A cap that would fix a maximum share per provider is discussed in the community, carries neither PLAN nor DEPL status in the target state, and thus remains a discussion without an implemented roadmap status. Here lies the governance-side gap in distribution. Rainbow Staking (RES) would fan out validator roles into light (attestation) and heavy (block building) tasks and lower the entry barrier, but is a research concept rather than an implemented mechanism.

ical limit of neutrality, since approximately 39 percent of all nodes are located in the United States.[327] The causes of geographical clustering lie outside the protocol's influence.

Precisely that decentralization which the protocol itself cannot actively bring about forms, however, the foundation for the security guarantees of the preceding sections. The mathematical assumption of forced transaction inclusion thus presupposes a sufficiently diverse participant network, while the decoupling mechanism protects an existing distribution but cannot artificially reproduce it. The conceptual neutrality of the entire protocol is thus based on a structural diversity of actors that is indispensably presupposed in the target state, but not guaranteed by the system itself.

5.5-D The Price of Finality

Between the vote of the Attesters and the definitive confirmation by the protocol, approximately 13 minutes currently elapse — a period during which a block is considered confirmed but not yet irreversible. During this period there are two problems, which the preceding section resolved economically and referred to this section for technical resolution. The first problem is the retroactive displacement of an already confirmed block through a chain reorganization (Reorg). This is the third and final main danger arising from concentrated production. The second problem is the artificial delay of block publication (Timing Games). The economic incentive for this has been removed from the Proposer by the separation of Proposer and Builder, but the temporal possibility still exists. Both problems exploit the same tolerant temporal structure of the network, which makes delays and retroactive changes technically possible at all. The following section brings the topics together through Single-Slot Finality, since immediate, second-grained block finalization removes the operational basis from both chain reorganizations and artificial publication delays.

The full mechanics of the delay — which Section 5.5-B merely named — belongs here. The value a Proposer extracts from the ordering of transactions grows with every second in the slot, as temporal shifts in market prices create additional arbitrage opportunities and enable further liquidations. A Proposer

[327] Geographic node distribution as the second, exogenous boundary of neutrality. In the current state, the US share of nodes stands at approximately 39 percent per current survey (Etherscan, early 2026, from 14,339 recorded nodes; older surveys cited up to 46 percent), so the US share lies, depending on the survey, in a range of 39 to 46 percent and constitutes the most pronounced geographic concentration of node distribution. Stateless Clients structurally lower the hardware barrier but do not address the non-technical causes (regulation, infrastructure costs, technical knowledge). No direct roadmap measure exists. The full treatment as a resilience topic takes place in 6.2. Source: Etherscan (early 2026).

who publishes their block later extracts more MEV (these delay strategies, cf. Section 5.5-B). They reward actors with the greatest latency advantage, who bid late enough yet still arrive in time, and shift the temporal dynamics of consensus in favor of fewer operators.[328] The decoupling introduced with APS deprived the delay of its direct reward by separating the extracted value from the Proposer. It left, however, the slot's discretionary window intact — within which the delay takes place — so that Timing Games merely shifted to the level of Builders. As long as the protocol permits a temporal tolerance within a slot, the structural cause for the emergence of Timing Games formally remains.

The risk of a chain reorganization (Reorg) exploits the same time window but operates at a different point. A block is considered confirmed as soon as the Attesters have voted for it. It becomes irreversible, however, only once the protocol finalizes it after two consecutive epochs (checkpoints). In the interim, a reorganization of the chain could theoretically displace the block and replace it with another. A market-dominant actor with a sufficiently large share of validation votes could thereby exchange an already confirmed block. In this way, a transaction could be retroactively removed from the chain even after it had just been included through the inclusion list (from Section 5.5-A). With the replacement of the block, the transaction it contained would also disappear. The complete inclusion of transactions that the first protection mechanism of this chapter was meant to enforce would remain unsecured against such retroactive displacement. The underlying separation of Proposer and Builder (ePBS) already binds the Builder to a firm commitment. Mechanisms such as Builder Staking and dedicated validator confirmations from the Payload Timeliness Committee (PTC) further deprive it of the ability to simply withhold a completed block. As a result, retroactive modification of the blockchain is limited to at most a single slot. This risk nevertheless persists as long as the first confirmation of a block and its irreversible finalization remain temporally separate in the protocol.

[328] Timing games and their structural cause. Strategies in which a proposer delays block publication because the extractable MEV rises monotonically with time in the slot. They are considered an economic neutrality violation that destabilizes the temporal structure of consensus in favor of latency-advantaged actors. The decoupling of the extraction removes the economic cause (5.5-B); the ePBS delivery obligations via builder staking and PTC attestation make strategic withholding unprofitable and simultaneously address the Free Option Problem (up to six percent empty slots on volatile days). The structural cause — the loose time window of the slot — is only removed by Single-Slot Finality.

Single-Slot Finality as the Answer

Single-Slot Finality is the roadmap's goal for fully closing the temporal gap between the first confirmation and the irreversible finality of a block. The concept promises, however, more than the currently favored design actually delivers. The research around Justin Drake, Buterin, and others currently places the practically achievable finalization at three slots, which is why the concept must be understood in substance as three-slot finality. Finalization within a single slot would require spreading and combining the attestations of the entire Validator set into a proof within a single interval — which exceeds the available time for propagation and aggregation. Three slots give the network this time and remain short enough to deny reorganizations and delays their room to maneuver. At the current tempo of twelve seconds per slot, the period between confirmation and finality falls from approximately 13–15 minutes to approximately 36 seconds. The accelerated finalization merely shortens the time until the security threshold is reached. It does not lower that threshold. Once finality is reached after 36 seconds, reversing the chain is economically just as expensive — and thus practically just as impossible — as in the current system. Independently, the time to first confirmation falls to a few seconds via so-called Fork-Choice Confirmations. Section 5.6-C develops the mechanics. Finalization on the short timescale of slots thus eliminates at a single point both the risk of retroactive chain reorganizations and the remaining technical cause of Timing Games.

At the level of the entire Validator set, finality requires one of the most profound architectural adaptations of the blockchain altogether. Currently, only a subgroup (*Committee*) of validators attests per slot. Single-Slot Finality (SSF) requires instead that the attestations of the entire Validator set be condensed into a verifiable proof within a few seconds. No current consensus layer is designed for this enormous load of signature aggregation. The roadmap provides for the redesign under the concept of *Lean Consensus* and locates the end state in the Beam Chain — corresponding to a complete redesign of the consensus layer whose implementation is a research state beyond 2029 without a fork deadline.[329] The raising of the maximum Effective Balance (MaxEB), which has been driving the consolidation of validator nodes since the Pectra upgrade, reduces the

[329] Lean Consensus and Beam Chain (RES). The Beam Chain is the complete redesign of the consensus layer with three-slot finality and SNARK aggregation, replacement of today's Beacon Chain as a research state beyond 2029 without a fork date. It absorbs Single-Slot Finality and MaxEB and depends for SNARK aggregation on zkEVM maturity, carried as the endpoint of Lean Ethereum. The EF post-quantum team has existed since January 2026; the reform of the Inactivity Leak for the migrated scheme is a sub-item of the redesign (cf. 5.3-D). MaxEB (EIP-7251, DEPL, Pectra) reduces the validator set via node consolidation and is a prerequisite for fast aggregation. Stake distribution consequences are treated in 5.5-C.

absolute number of validators and is thereby a direct functional prerequisite for aggregation. The efficient compression of the signatures of the entire set further depends on the market maturity of advanced ZK proofs. The endpoint of rapid finality thus coincides technologically with the most demanding building blocks of scaling. The requirements for Single-Slot Finality thus reach from the fine temporal structure of the slot deep into the fundamental architecture of the consensus layer.

The cryptographic basis of the challenge is the signature procedure on which aggregation relies today. The BLS signature procedure natively combines the signatures of many validators into a single one, and this property makes fast aggregation computationally and capacity-wise highly efficient. The hash-based procedure of the leanXMSS line, intended to replace BLS in consensus, produces signatures that are orders of magnitude larger. In addition, it carries no native aggregation capability, so the switch massively complicates aggregation technically.[330] The new procedure is chosen as a post-quantum cryptography measure, whereby the consensus layer anticipates the same transition to hash-based procedures that Section 5.3-D derived for verification. The technological requirements of Single-Slot Finality overlap here with the necessities of quantum security, without either challenge inherently resolving the other.

The larger signatures of the new procedure are handled by an integrated zkVM that creates a compact validity proof for them. Via a recursive SNARK, the instance referred to in the roadmap as leanVM folds them into a single proof. This reduces the data load by approximately 250-fold, so that finality remains within the time window despite the larger individual signatures.[331] The same aggregation mechanics that 5.3-A′ invoked for the quantum-resistant proofs of the application layer thereby also makes the consensus signatures scalable again. Closing the last time window in consensus thus means rebuilding it down to its signature procedure — and the new procedure is chosen quantum-secure from the outset.

With this, the third and final danger is answered: the retroactive displacement of an already confirmed block — after the exclusion of a transaction and the

[330] BLS and leanXMSS. BLS (DEPL) aggregates consensus signatures natively into a single one, keeping fast aggregation cheap; the hash-based scheme of the leanXMSS line (RES), which replaces BLS in consensus, produces signatures larger by orders of magnitude and does not aggregate natively. The quantum question itself — its framework and migration logic — is handled in 5.3-D and is touched here only in its function as the chosen security level. Implementation statuses as leanSig (Rust) and leanSpec (Python).

[331] leanVM and SNARK aggregation (RES). A minimal zkVM folds the larger leanXMSS signatures via a recursive SNARK into a single proof and compresses them by approximately 250 times, so that finality remains fast despite the larger individual signatures. The mechanics share the aggregation principle that 5.3-A′ carries for quantum-resistant proofs.

extraction of its value. These three answers protect the neutrality of the consensus layer against its most powerful actors, without dissolving the concentration of block production itself. They also all rest on a diversity of the Validator set that Section 5.5-C has marked as unsecured. What weight the unsecured diversity carries against the achieved guarantees is analyzed in the target-state assessment in 5.7, from which the chapter derives its overall verdict in 5.8.

5.6 L2 Architecture and Settlement

5.6-A The Second Security Class and Its Convergence Point

The preceding section addressed the neutrality of Layer 1 against its most powerful producers at the protocol level, while one layer higher the two-class problem of Layer 2 persists. Whoever sends a transaction on Layer 1 trusts solely the protocol, whose consensus and cryptography rejects every rule violation. Whoever sends the same transaction on a Rollup trusts additionally the company that operates the Rollup. The operator controls two services themselves today: the correct execution of transactions (Execution) and the order in which they are processed (Sequencing). From Layer 1, a Rollup so far takes only data availability, whose guarantee Section 5.2 developed. The target state reverses the verification relationship of Execution: today the operator proves to Layer 1 the correctness of their system; in the future Layer 1 checks it with its own State Transition Function. This section addresses Execution, the following one Sequencing. From the convergence of both follows the basis of the Settlement role: Ethereum as the layer on which other networks and applications definitively record their state changes.

Behind the second security class stand identifiable actors, since the operator of a major Rollup is in each case a single company — Offchain Labs for Arbitrum, Coinbase for Base, Matter Labs for zkSync Era. The state follows structurally from the rollup-centric scaling that shifted execution onto Rollups and left its security to the operators. The concrete problem lies in a verification gap in the protocol: a Rollup publishes its transaction data and the resulting State Root on Layer 1, yet Layer 1's consensus possesses no rule of its own connecting the two. Whether the claimed State Root actually follows from the execution of the published transactions is verified solely by a Verifier Contract — a Smart Contract that the operator develops, deploys, and can modify via its own governance. Layer 1 executes the contract according to its rules but does not vouch for the correctness of its verification logic, and a bug in the proving circuits can allow an

invalid state transition to pass as cryptographically correct. The operator thereby replicates a function that Layer 1 has long performed natively for its own blocks: the verification of EVM state transitions.

From the gap follow the three self-provided services on which the correctness of a Rollup today rests. A standalone proof system attests that the operator correctly advances the state of its users, and a Security Council — a body with privileged intervention rights — can halt or modify the system in an emergency. Added to this is a rollup-specific governance that decides on upgrades to which users are subject. Optimism relies on the fraud-proof system Cannon, Arbitrum on BOLD, zkSync Era on the ZK system Boojum.[332] Each of the three systems requires its own codebase and its own audit, and every line of proprietary proof logic increases the probability of failure against which the Security Council must stand ready. For the user, the security of their funds thus depends on the competence and integrity of identifiable companies and bodies — a dependency that Layer 1 has eliminated at the protocol level.

The counter-concept to operator-based security is Trustlessness — the state in which the validity of every operation follows solely from protocol rules and cryptographic verification and requires no trust in identifiable actors. How much power an operator still holds over its users can be categorized in three development stages, which the analytics platform L2BEAT — a continuous documentation of the degree of decentralization of Rollups — runs as a Stage schema. At Stage 0, the operator can modify the system or freeze user funds at any time. At Stage 1, a functioning proof system is present and the Security Council is limited to emergencies. At Stage 2, every privileged intervention is excluded and mathematical verification alone applies.

No Rollup at Stage 2

As of early 2026, no major Rollup has reached Stage 2: Arbitrum, Base, Optimism, Starknet, and Scroll stand at Stage 1, zkSync Era and Linea at Stage 0.[333]

[332] Proprietary proof systems of rollups in the current starting state (delivered per rollup). Optimism operates the interactive fraud-proof system Cannon, Arbitrum the dispute protocol BOLD, zkSync Era the validity-proof system Boojum. Additionally security councils, multisig upgrade keys, and rollup-native governance serve as institutional safeguards. This proprietary infrastructure carries the second security class and simultaneously forms the vendor lock-in source that the target state dissolves through Native Rollups.

[333] L2BEAT stage system and L2 decentralization status in the current starting state (survey early 2026). Stage 0 allows the operator to intervene at any time up to freezing user funds, Stage 1 requires a working proof system and restricts the security council to emergencies, Stage 2 prohibits any privileged intervention and allows only mathematical verification. No major rollup has reached Stage 2. Arbitrum, Base, Optimism, Starknet, and Scroll stand at Stage 1, zkSync Era and Linea at Stage 0, and several operators have publicly left Stage 2 open, among other reasons regulatory. All major rollups run a single-operator sequencer (Offchain Labs, Coinbase, Optimism Labs, Matter Labs, StarkWare), while the leading shared-sequencer attempt Astria was discontinued in 2025; Vitalik Buterin describes

Several operators have left the relinquishment of the last intervention rights publicly unanswered for regulatory reasons. In the second direction a comparable concentration exists, because every major operator deploys the Sequencer — the instance determining the order of transactions — as a single, self-controlled component. The leading attempt at a shared Sequencer, Astria, was discontinued in 2025, and the expectation that operators would incrementally relinquish their position of their own accord has not been fulfilled in either direction. Buterin drew the consequence publicly in February 2026: progress toward Stage 2 had been far slower and more difficult than originally expected.[334] The strategic reorientation that follows from this is carried by Section 5.1.

The Stage classification describes the existing state; the trajectory of the target picture leads both Rollup classes to the same end state. The distinction between Optimistic Rollups with a downstream challenge period and zkRollups with an upfront Validity Proof dissolves to the extent that Validity Proofs become cheaper and more general in their reach. At the end, zkRollups achieve proof of complete EVM execution, and Optimistic Rollups abandon the seven-day challenge period once a cheap Validity Proof replaces it. At the end of both movements stands the same cryptographic Settlement security, at which no user needs to trust the operator for correctness. It is fully achievable for neither class as long as the operator maintains the proof system itself — and here the target state intervenes: the verification of state transitions is relocated into the protocol of Layer 1.

Verification Relocates into the Protocol

Native Rollups (EIP-8079, RES) relocate verification via a single protocol-level call onto Layer 1.[335] The core is the EXECUTE precompile, which extends

progress toward Stage 2 as far slower and more difficult than originally expected — a finding among the empirical triggers of the strategic pivot.

[334] Empirical trigger of the strategic pivot (February 2026). Vitalik Buterin names in X posts of 3 and 5 February 2026 the progress of rollups toward Stage 2, which proceeded far more slowly and with greater difficulty than expected, as one of the two findings behind the reorientation of the L2 role. The strategic interpretation — the spectrum of trust models with Stage 1 as the minimum threshold for rollups with Ethereum-native assets and the role of protocol-side verification — is treated in section 5.1. Sources: Vitalik Buterin, X (3 and 5 February 2026); cf. 5.1.

[335] Native Rollup Precompile (EIP-8079, RES; Draft in Standards Track Core since 13 November 2025, authors Luca Donno of L2BEAT and Justin Drake of the Ethereum Foundation, without CFI confirmation and without fork headliner status). The EXECUTE precompile exposes the state transition function of Layer 1 and accepts three inputs: the pre_state_root, the post_state_root, and the witness trace with the read and written state values. It depends on the compact witness form of the state restructuring and on protocol-side integration via enshrined Proposer-Builder Separation, while SNARKification presupposes the maturity of the Layer 1 zkEVM. Vitalik Buterin signaled explicit support on 19 January 2026; a proof of concept by the Ethrex client team with the Ethereum Foundation and L2BEAT of 11 March 2026 documents verification via re-execution, explicitly as a research milestone and not as a deployment decision. The conceptual origin is Justin Drake's ethresear.ch post on Native Rollups of January 2025; the EIP

the precompile layer introduced in Section 5.3-A′ with an entry point into Ethereum's own State Transition Function. The operator submits a pre_state_root, a post_state_root, and a Witness Trace, and the precompile checks whether Layer 1 computes the same end state from the initial state as the operator claims. The Witness Trace — a compact collection of precisely those State values that the Rollup block reads and writes — enables verification without Layer 1 needing to hold the complete Rollup State. The operator's claim is thus checked by the same thousands of validators that verify every Layer 1 transaction, and the correctness of a Rollup block moves from a contract rule of the operator to a rule of the protocol. In the proof-of-concept, verification proceeds by re-execution; prospectively by a SNARK proof — continuing the verification infrastructure that on Layer 1 separated recomputation from checking.

For the user of an EVM-equivalent Rollup — whose execution corresponds exactly to the specification of Layer 1 — the trust foundation would be different. The validity of Rollup blocks would be decided by the same consensus as a Layer 1 transaction, so that an attack on their funds would require an attack on Layer 1 itself. From the inherited client diversity, a multi-prover system would arise without additional infrastructure. Because different Execution Clients use different proving procedures, a bug in any single procedure cannot finalize a false state — the deviation would be visible in consensus and would find no majority. The Security Council of the operator becomes superfluous, and automatic compatibility with the hard forks of Layer 1 replaces operator-specific rollup governance that would otherwise need to be kept in sync. The Bridge — the connection layer for deposits and withdrawals between the levels — disappears as a trust layer, because Withdrawals would be secured by the same consensus as an ordinary Layer 1 transaction, without Security Council and without challenge period. On the side of future operators, the entry threshold shifts from the million-dollar effort for proof system, audit, and Security Council to the configuration of a single call. The permissionless access that Layer 1 grants for individual transactions extends to the creation of entire Rollup networks.

The gain remains tied to two conditions, of which the first concerns the costs of verification. As the second condition, the Witness Trace must be available on-chain and reaches the magnitude of five to ten times the Rollup batch — an expenditure that remains manageable only because the data availability capacity

motivation explicitly names the redundancy, since EVM-equivalent rollups maintain complex proof systems solely to reproduce what Layer 1 already provides. Source: Justin Drake, ethresear.ch (January 2025).

of Layer 1 has grown structurally.[336] The Witness remains compact only under the restructured State organization from Section 5.4. With today's tree structure it would be two orders of magnitude larger and prohibitive. Added to this is a throughput limit, because in the first phase every node on Layer 1 fully recomputes every invocation of the precompile. The number of invocations per block therefore remains limited until SNARKification replaces the repeated execution. The second condition concerns the scope of resolution, because the exposed State Transition Function supports neither custom opcodes nor proprietary precompiles or divergent transaction types. Any deviation from EVM equivalence therefore leads back into self-managed settlement, in which the user again trusts the operator's institutions. Most Rollups today are EVM-compatible but not EVM-equivalent, and the resolution is initially open only to that part of the ecosystem whose execution matches the specification of Layer 1.

In the target state, execution would thus be brought into the first security class, insofar as a Rollup remains EVM-equivalent and the compact Witness form and protocol-level verification are available. A strand of research investigates the generalization of the EXECUTE precompile to a VM-agnostic basis — an approach belonging to the Lean Ethereum direction and standing here as an outlook. For existing Rollups with their own proof system, the Bridge with its trust assumptions would remain a transitional phenomenon diminishing in significance with the spread of Native Rollups. One power the operator retains regardless: the order of transactions, over which the centrally controlled Sequencer continues to decide alone. The following section addresses the relocation of Sequencing to Layer 1 as well. From the convergence of both directions emerge the complete Trustless state with its Finality Stack and the Settlement role, in which a heterogeneous ecosystem of chains definitively records its state changes on Ethereum.

[336] DA overhead of the witness trace in the target design state. A typical rollup batch lies between 125 and 750 kilobytes, while the witness trace to be provided on-chain in the compact witness form is on the order of three megabytes per EXECUTE invocation, yielding an overhead of five to ten times the batch. With today's Merkle Patricia structure the witness would reach approximately 300 megabytes and would be prohibitive. The overhead remains manageable solely because the data availability capacity of Layer 1 has structurally grown.

5.6-B The Shared Order: Based Sequencing and Based Preconfirmations

Resolving the Execution resource leaves the operator one remaining service: the order in which the transactions of their Rollup are processed. The Sequencer, which every major operator deploys as a single, self-controlled component, bundles three authorities. It can omit a transaction, enabling a local censorship at the Rollup level that Layer 1 addresses at the protocol level for its own blocks (Section 5.5). It determines the ordering of transactions and can extract the value that follows from it. And its failure halts the entire Rollup, because liveness depends on the infrastructure of a single company. For the user this means that the inclusion of their transaction and the timing of their withdrawals depend on the availability of a component that no protocol binds. The three authorities structurally repeat what the preceding section negotiated for Layer 1, but they lie in the hands of a third actor — the operator — who joins the Builders and stakers of the base layer. The target state hands the ordering to Layer 1: Based Sequencing lets the Proposer of the slot order the Rollup transactions, and Based Preconfirmations give the user a confirmation in approximately 2 seconds for this.

The mechanism rests on a property that every Rollup already possesses: its execution follows deterministically from the order of its transactions. Once the sequence is fixed, every Rollup node computes the same state from it, and whoever fixes the order has done everything decisive, without executing themselves. Based Sequencing uses this property and dispenses with a dedicated ordering instance: the Proposer of the Layer 1 slot includes the next Rollup block as a data packet in their own block, and with this inclusion the order of the Rollup transactions is binding.[337] The user continues to send their transaction to the Rollup; it travels as part of the packet through the Layer 1 block without becoming a Layer 1 transaction itself, and the Rollup nodes then execute the fixed sequence. In place of the single company server steps the entire Validator set of Layer 1; the liveness of the Rollup is identical with the liveness of Layer 1; and participation in the ordering is permissionless. With the Sequencer, the surrounding infrastructure disappears: separate consensus, escape hatch, and proprietary ordering rule

[337] Based Sequencing (DEPL via Taiko mainnet, conceptually without its own EIP). Proposed by Justin Drake in March 2023 (ethresear.ch, Based rollups as L1 sequencing); Based Preconfirmations followed in November 2023. The proposer of the L1 slot, in collaboration with builders, includes the next rollup block in the L1 block; the rollup inherits the validator set of Layer 1 (approximately 1 million validators) as sequencing infrastructure, without its own sequencer component, without additional consensus, and without an escape hatch, and rollup rules follow Layer 1 hard forks automatically. Of the three modules Layer 1 provides a rollup — Inclusion (data availability), Ordering (sequencing), and Execution (settlement) — a Based Rollup also uses the second. The third is fetched by native verification (cf. 5.6-A).

beyond Layer 1 are eliminated. A Based Rollup is no longer a network with its own ordering — its blocks arise as part of the Layer 1 blocks, and what remains is an execution environment over the shared ordering.

Because all Based Rollups use the same ordering, the same Proposer orders their blocks within the same slot — a basis whose significance for the interplay of Rollups Section 5.6-D develops. The neutrality layer of the preceding section accordingly extends to the Rollup, to the same extent and with the same qualification with which it was addressed there. One building block is still missing from the procedure: whoever is to give a commitment regarding the forthcoming order must be established in advance. This is provided by the Deterministic Proposer Lookahead (EIP-7917, DEPL since Fusaka): the protocol computes in advance and bindingly which validator proposes the block in which of the upcoming slots — a publicly visible sequence of the next Proposers.[338] Before this anchoring, the sequence was manipulable, because a validator could shift the allocation by changing their effective balance — a path the deterministic computation closes.

The price of the handover is speed, because a central Sequencer confirms in approximately 0.25 seconds, while a Based Rollup without further mechanics does so only with the next Layer 1 block — after up to 12 seconds. On this difference the operators' decision for their own Sequencer had so far depended. The target state's answer is Based Preconfirmations, hereafter Preconfs. A subset of validators registers voluntarily as Preconfirmers and deposits additional stake as security. Whoever sends a transaction receives from the next responsible Preconfirmer — identified by the Lookahead sequence — an immediate commitment regarding the inclusion and position of their transaction.[339] If the Preconfirmer fails to honor the commitment, they lose a portion of the deposited stake through Slashing. The commitment reaches the user after approximately 2 seconds and is economically secured. Cryptographic finality remains absent — a remainder that the Finality Stack of the following section puts in context. Based Sequenc-

[338] Deterministic Proposer Lookahead (EIP-7917, DEPL since Fusaka, 3 December 2025; authors Lin Oshitani, Nethermind, and Justin Drake, Ethereum Foundation). The proposer assignment for upcoming slots is computed deterministically and protocol-side, making the identity of the upcoming L1 proposer predictable and enabling preconf guarantees to be given; at the same time the grinding of effective balance as an attack on the prediction is eliminated. Primary assignment 5.6-B as a prerequisite for Based Sequencing and Based Preconfirmations.

[339] Based Preconfirmations (DEPL partial, Taiko Phase 1). A portion of validators registers via a registry contract with deposited stake as preconfirmers; the user identifies the next responsible preconfirmer via the lookahead, sends the transaction with a preconfirmation request, and receives a commitment in the range of approximately 100 milliseconds to 2 seconds, whose breach is sanctioned via slashing. The guarantee is economically secured; its classification as the first layer of the finality stack is accomplished by 5.6-C.

ing thus forms the second building block of the strategic reorientation whose framework Section 5.1 carries.

The proof of concept for the architecture has been running on Mainnet since May 2024: Taiko has operated since then as the first Based Rollup and has processed over 900 million transactions without downtime as of April 2026.[340] Since August 2025, Preconfs have been running on Mainnet in a first phase via a permissioned list of Preconfirmers, with confirmation times of approximately 2 seconds. The proof carries a qualification: the Preconf layer runs permissioned, and its opening is on the roadmap for 2026 and thus still outstanding. As an Ethereum-equivalent Rollup, Taiko also satisfies the condition on which the inherited protection of the Execution resource depends — and the candidate for the convergence of both directions is named.

Three Open Limits

Three limits remain open, of which the first is economic: with the ordering, its revenues migrate too — the value of ordering flows to the Proposers of Layer 1, and the operator loses a revenue source. The value lands where it finances security, as the revenues feed the security budget whose composition Section 5.5-C developed. On this loss hangs the reluctance of existing operators, and the operational evidence follows the pattern: Taiko started as a Based Rollup, and no existing major operator has made the switch. The second limit concerns the prerequisite itself, because the deterministic Lookahead that enables Preconfs makes the future Proposer publicly predictable and thus a reachable target for a targeted denial-of-service attack. A line of research under the name Single Secret Leader Election pursues the opposite and keeps the identity of the Proposer secret until their slot is executed, to shield them from the attack. The advance commitment of Preconfs and the protection of the Proposer thus currently stand as architectural choices against each other.[341] The third limit is the separation of directions: the shared ordering does not make a Rollup trustless as long as its Execution

[340] Taiko Alethia as operational proof of existence (DEPL). Mainnet since May 2024 as the first Based Rollup, over 900 million processed transactions without downtime through April 2026; preconfirmations on mainnet since 11 to 13 August 2025 in Phase 1 via a permissioned whitelist of preconfirmers with approximately 2 seconds confirmation time; roadmap: fully decentralized preconfirmations and formal protocol specification Q1 2026, sub-second latency and Stage 2 path Q2 2026. Taiko identifies itself as an Ethereum-equivalent (Type-1) ZK-EVM rollup. Source: Taiko announcement and documentation (August 2025).

[341] Lookahead-SSLE tension. The Single Secret Leader Election (SSLE, RES) keeps the identity of the proposer secret until the execution of its slot, protecting it from targeted denial-of-service attacks, relevant for I.4 (Liveness under attack) and II.1; the Deterministic Proposer Lookahead (EIP-7917) pursues the opposite with deterministic prediction. The relationship represents an architectural decision between L2 preconf enabling and DoS protection and is unresolved in the target state.

remains with its own Verifier Contract, and the resolution of the preceding section and that of this section only take effect together.

In the target state, the two services a Rollup today performs itself would individually be relocated to Layer 1: verification of execution via the protocol-level check, ordering via shared Sequencing. The operator who at the start of the section held both directions would hold neither. The following section assembles from both directions the complete picture: the Rollup as an execution environment whose ordering and correctness Layer 1 itself carries and whose confirmations a three-tier Finality Stack arranges.

5.6-C The Convergence: The Trustless Rollup and Its Finality Stack

The Rollup at the end of the preceding section has handed over its Sequencer and is nonetheless not trustless. Its operator (the Rollup company) can no longer corrupt any transaction, and no failure of its infrastructure halts the Rollup any longer. The order is fixed by the Proposers of Layer 1. Yet what the ordered sequence brings about is still checked by the operator's Verifier Contract. If it accepts a faulty State Root — for instance through a bug in the proving circuits — the users' funds are incorrectly recorded, even though the ordering was correct. A Rollup that conversely only uses EXECUTE verification is checked by the consensus of Layer 1, but its operator still orders and can omit or fail. The two-class problem — the finding from the beginning of the section — is thus in both cases only halved, and only the combination of both mechanisms closes both gaps simultaneously. The section first describes what such a Rollup is, and then the three confirmation levels that a user receives there, from the first commitment to final finality.

A Rollup combining both mechanisms — hereafter the "converged Rollup" — uses all three modules of Layer 1: data availability (Section 5.2), ordering of transactions (Section 5.6-B), and verification of Execution (Section 5.6-A). After the Rollup nodes have computed the new state from the fixed sequence, the EXECUTE verification confirms the transition through the same consensus that also checks the blocks of Layer 1. The operator's Verifier Contract — on which the funds were still hanging just moments ago — becomes superfluous, and with it the risk of a faulty State Root. A converged Rollup therefore carries no trust assumption beyond Layer 1. What remains of it is an execution environment: its own fee market, its own state, and the applications on it — without its own Sequencer, without its own proof system, without a Security Council, and without its own governance. In the discourse of the ecosystem the converged Rollup

bears the name Ultrasound Rollup, and with Taiko Gwyneth an implementation is in development that explicitly aims at the combination of both mechanisms.[342] Creating a Native Rollup has shrunk to one invocation of the EXECUTE precompile, and because the ordering is open to all validators, such a Rollup would be the first to be permissionlessly replicable. The consequences of the interplay for the relationship of Rollups to one another are developed in Section 5.6-D.

At the protocol level, both mechanisms rely on the enshrined Proposer-Builder Separation that the chapter has already carried as a capacity building block (Section 5.2) and as a neutrality condition (Section 5.5). The separation of building from proposal gives the Proposer the position from which it orders the Rollup blocks. And it opens the time window in which EXECUTE verification is integrated into the Layer 1 block.[343] The verification itself flows into the infrastructure that Section 5.2 developed for Layer 1: the L1 zkEVM, which validates a block via proof and replaces recomputation. The EXECUTE verification carried via SNARK proof reverses the current relationship of proof systems. In the baseline state, every Rollup maintains its own; in the target state, a single one — the zkEVM of Layer 1 with its multiply implemented provers — checks the state transitions of all EVM-equivalent Rollups.[344] The timelines of both projects are explicitly coupled, since the maturity of the zkEVM is simultaneously the prerequisite for protocol-level Rollup verification.

[342] Convergence Native plus Based (target image, dependent on the research status of the EXECUTE precompile, EIP-8079, RES). The combination forms the maximum degree of coupling: Inclusion, Ordering, and Execution all come from Layer 1, the system is fully trustless. In ecosystem discourse the construction is referred to as Ultrasound Rollup. Taiko Gwyneth is in development as a Based Rollup with synchronous composability to Ethereum and explicitly targets the convergence.

[343] Enshrined Proposer-Builder Separation as convergence point (EIP-7732, PLAN, SFI Glamsterdam; cf. 5.2-A on capacity and 5.5-A on neutrality function, here the third and final encounter). ePBS forms the convergence point of the roadmap that conditions, among other things, the EXECUTE integration of Native Rollups and the proving window of the zkEVM (Slot N+1).

[344] L1-zkEVM / Optional Execution Proofs (EIP-8025, PLAN, EF implementation roadmap with six sub-themes of 26 January 2026, without hard fork requirement and thereby of novel governance status). Validation through proof rather than re-execution, with 3-of-5 multi-prover via client diversity; the realtime proving target (99 percent of blocks under ten seconds) was reached in December 2025, with an improvement from 16 minutes to 16 seconds proving latency and approximately 45-fold cost reduction, while formal security hardening is outstanding in graduated milestones (100-bit by Glamsterdam, 128-bit by end of 2026). Vitalik Buterin explicitly couples the timelines: the adoption of ZK on Layer 1 and the introduction of Native Rollup precompiles proceed essentially in sync. The primary treatment of the zkEVM is carried by 5.2; here the L2 effect via the SNARKified EXECUTE check counts.

The Tiered Finality Stack

What the user of a converged Rollup receives as confirmation is tiered across three layers with three different types of guarantee. The first layer delivers after approximately 2 seconds the Preconf — the economically secured commitment of a registered Layer 1 validator. The second layer delivers the Fast Confirmation Rule (PLAN), an attestation-based rule without a hard fork, specified as a pull request in the consensus-specs and in client implementation.[345] It declares a Layer 1 block deterministically confirmed once sufficiently many attestations support it — after approximately 13 seconds. It presupposes that at least three quarters of the stake attest honestly and network latency remains below three seconds. With it, the Layer 1 confirmation falls from approximately 13 minutes to approximately 13 seconds, accelerating exchange deposits and bridging by approximately 98 percent.

The third layer is finality itself. The roadmap goal is Single-Slot Finality, but the practically converging design — whose architecture and price Section 5.5-D developed as three-slot finality — reduces irreversible finalization in the target state to approximately 36 seconds.[346] The stack tiers three types of guarantee in ascending strength: the secured commitment after approximately 2 seconds, the deterministic confirmation after approximately 13 seconds, and the cryptographic finality after approximately 36 seconds. For the user of a converged Rollup, each waiting time corresponds to a nameable guarantee, and the strongest guarantee no longer costs minutes.

The reach of convergence remains limited, and its limits are those of both directions. It applies to Rollups that remain EVM-equivalent and adopt Based Sequencing. Both are decisions of the operators, and the ecosystem therefore retains a spectrum of attachment levels, from full convergence to the sovereign chain with minimal binding to Layer 1. The Stack itself is unequally mature: the middle layer is active, the first runs partially and permissioned, the third is a research subject beyond 2029. Added to this are two qualifications of the carriers. The zkEVM currently achieves its speed on security assumptions whose for-

[345] Fast Confirmation Rule (consensus-specs PR #4747, PLAN, pure client feature without hard fork, client implementation ongoing with rollout in the coming months). Attestation-based confirmation: at at least 75 percent honest stake and a network latency under three seconds, a block is considered deterministically confirmed after approximately 13 seconds; deposits at exchanges and bridging from Layer 1 to Layer 2 accelerate by approximately 98 percent. The rule is independent of ePBS and of Single-Slot Finality and forms the second layer of the finality stack between Based Preconfirmations and SSF.

[346] Single-Slot Finality as the third layer (RES, horizon after 2028; the full treatment including the substantive classification as three-slot finality is carried by 5.5-D). In the target state, irreversible finalization drops to approximately 36 seconds (three slots); in the present section only the stack function as cryptographic finality above the deterministic confirmation counts.

mal hardening is pending in graduated milestones. And the deterministic confirmation rests on the honesty of three quarters of the stake — a precondition that refers back to the unsecured distribution from Section 5.5-C.

In the target state, the two-class problem would thereby be resolved for the converged part of the ecosystem. A Rollup there is an execution environment of Layer 1, ordered by its Proposers, checked by its consensus, confirmed in its Stack. What loads such a layer can carry once other systems make binding entries on it is developed in the two following sections: first the Settlement of tokenized values, then the role for a heterogeneous ecosystem of chains.

5.6-D The Relationship of Rollups and the Settlement of Tokenized Values

The Rollup that at the end of the preceding section stood as an execution environment of Layer 1 does not exist in the target state alone. Whoever moves funds from one Rollup to another today crosses the border between two foreign security systems, with their own Bridge, their own deadlines, and their own trust assumptions on both sides. The convergence of both mechanisms changes the border itself, because with the security system of the individual Rollup, the foreignness between Rollups also dissolves. Both questions of the section depend on the same property: what the Rollups can guarantee one another is simultaneously what the layer offers external users. This section therefore first develops what follows from the shared ordering and the shared finality horizon for the relationship of Rollups to one another. It then asks what loads the layer carries once not only its own ecosystem makes entries on it. It finds the answer with the tokenized capital of the traditional financial world, whose booking presents the most demanding requirement profile.

The new relationship begins with the ordering, because the handover of Sequencing has effects beyond the individual Rollup: the same Proposer orders the blocks of all Based Rollups within the same slot (Section 5.6-B). The relative order of transactions of two Rollups thus arises in a single act, holds across the Rollup boundary, and no longer depends on the coordination of independent systems. Whoever trades on two Rollups simultaneously plans against a single binding sequence rather than against two mutually independent orderings whose relationship no one guarantees. Because the same verification decides on the state transitions of all converged Rollups, their entries also share a common finality horizon. A transfer between two such Rollups would be finalized on both sides in the same Stack, according to the same deadlines and under the same assumptions (Section 5.6-C). A Bridge with its own trust assumptions —

whose disappearance Section 5.6-A developed for the relationship with Layer 1 — would no longer be needed. For an application spanning multiple Rollups, the decisive uncertainty would shift from trust in foreign systems to the waiting time until finality.

Composability — the ability of applications to interlock across system boundaries — reaches its strongest form in atomic composition, where two operations stand in the same block and either both succeed or both fail.[347] The atomic form remains unachievable in the target state, because it would require synchronous communication between Rollups within a single slot — which neither the shared ordering nor the protocol-level verification delivers. What would be achievable is the second level of the work's four-level logic: asynchronous Composability with deterministic guarantees, tendentially below a slot. Compared to the baseline state in which most connections run via the third level, the gain would lie less in speed than in the type of guarantee. In place of trust in third-party systems would come a nameable deadline with cryptographic backing. The preceding section announced a second consequence, and it concerns the growth of the ecosystem. Because a converged Rollup could be permissionlessly replicated and every copy would arise under the same ordering, each new Rollup would enter the described relationship from the outset. The ecosystem would grow as a set of execution environments over a common foundation, and its diversity would remain a question of specialization rather than security.

This relationship becomes practically usable through a layer that relieves the user of managing the borders. An Intent is the declaration of a desired outcome rather than a prescribed execution path. The user specifies what value should arrive on which Rollup, and a Solver — a market actor specialized in this — takes on the path and execution for a fee. The standard ERC-7683 (Cross-Chain Intents, DEPL) normalizes the form of such Intents across Rollups, and the Open Intents Framework (DEPL) assembles the associated ecosystem infrastructure.[348] The deterministic confirmation of Layer 1 shortens the deadline after

[347] Cross-L2 composability in the four-level M1 tier logic: atomic composability (same block, no trust assumptions), asynchronous composability under one hour, trust-bridge-based composability, no structural composability. The current state operates for most L2 interactions at level 3; the target state reaches level 2 via the combination of Based Sequencing (DEPL) and native verification (RES), asynchronous with deterministic finality guarantees and tending below one slot; level 1 remains unreachable in the target period because it would require a shared execution environment or synchronous cross-L2 communication within a single slot. Taiko Gwyneth describes itself as a "based rollup synchronously composable with Ethereum"; the claim refers to the synchronous composability of a single Based Rollup with Layer 1, not atomic composability between rollups, and remains a project claim of an implementation in development.

[348] Cross-Chain Intents (ERC-7683, DEPL; audited code since January 2025) and Open Intents Framework (OIF, DEPL; launched by the Ethereum Foundation in February 2025, over 30 teams). ERC-7683 standardizes order structures and settler interfaces for cross-L2 intents, with over 70 integrating projects and cross-chain execution in approx-

which the settlement of a Solver is fixed — a function of the Finality Stack from the preceding section. Solvers operate where liquidity and Settlement security are highest, and therefore route demand to the layer that combines both. The limit lies below the standard: ERC-7683 makes Intents portable, but Solvers maintain protocol-specific liquidity positions and proprietary routing logic, so that the execution environment unifies more slowly than the form of Intents.

Binding Settlement of Tokenized Values

Binding Settlement presents a requirement profile that goes beyond throughput. An entry must become final within a nameable deadline, its execution must favor no one, its correctness must be verifiable without permission, and its existence must be claimable for decades. The converged layer would deliver each of the four requirements from the inventory of the chapter. Finality would be carried by the tiered Finality Stack (Section 5.6-C), equal treatment by the protocol-level neutrality (Section 5.5), verifiability by the universally accessible verification (Section 5.3), and lasting claim by the differentiated State management (Section 5.4). Here equal treatment is more than a fairness property, because a booking location that can delay or omit individual entries is unsuitable as neutral infrastructure between competitors. Added to this is a property no single module delivers: the entry depends on no custodian whose continued existence the holder must verify — it would depend on the protocol itself. For institutional users, the Stack would translate into a guarantee whose strength grows with each of its layers and whose strongest form no longer costs business days.

The demand for such a layer is not a postulate of the target state. Tokenized values — referred to in the financial discourse as Real World Assets — already exist in the double-digit billions on the Ethereum Mainnet, the largest single booking location of the category.[349][350] The American market infrastructure opened the category regulatorily at the end of 2025, but maintains finality

imately 2 seconds. As of April 2026, productive endpoints were running at, among others, Across, UniswapX, CoW Protocol, and Eco. Unification of the solver layer is proceeding more slowly than that of the standard layer, because solvers maintain protocol-specific liquidity pools and routing logic. The Fast Confirmation Rule accelerates the settlement of intents via the deterministic confirmation of Layer 1 (cf. 5.6-C).

[349] Tokenized values (Real World Assets, RWA) on Ethereum, primary data status February 2026: over $17 billion on mainnet, approximately 315 percent annual growth compared to approximately $4.1 billion in the prior year; market share depending on measurement approximately 34 percent (The Block) to well over 50 percent (rwa.xyz); stablecoin holdings on mainnet over $175 billion. Verification per June 2026: $16.6 billion and 52.85 percent market share per rwa.xyz; total market over $65 billion with Ethereum at approximately 33 percent per The Block. The divergence of the share values is methodologically conditioned (different delineation of the category and the captured networks) and is stated here, not resolved. Sources: The Block (February and May 2026); rwa.xyz/Token Terminal (June 2026).

[350] BlackRock USD Institutional Digital Liquidity Fund (BUIDL) as the benchmark product of the category: tokenized money market fund, launched March 2024 with Securitize on Ethereum; approximately $2.5 billion AUM as of mid-May 2026, largest tokenized Treasury fund; issuance on multiple chains with Ethereum as the

in its pilot architecture at the central custodian, while the converged layer would carry it as a property of consensus.[351] What today lies there uses the guarantees of the baseline state, and the requirement profile of these users is simultaneously the standard against which the target state must measure itself. One requirement the layer does not fulfill itself: the confidentiality of business transactions — because the public booking exposes positions and counterparties that institutional actors cannot disclose — and the protocol-level target picture for that is developed in Section 5.1.

The section began with a structural problem and ends with a property that follows from its resolution. A layer that would carry its own Rollups without institutional trust — thereby resolving the two-class problem for the converged part — would offer external capital precisely the guarantees on which binding Settlement depends. Incoming capital deepens liquidity, and deepened liquidity attracts Solvers and further capital. Each level would raise the economic and security-related switching costs. A booking location whose departure becomes ever more expensive approaches the state in which its function is no longer replaceable. The qualification rests not on any single innovation. It rests on the interaction of the architectures the chapter has developed from capacity to neutrality. The finding thus stands: two mechanisms, each bringing one direction to Layer 1, would yield a layer — internally deterministically ordered, externally a booking location without institutional precondition. Whether this role reaches beyond its own ecosystem to chains that are not Rollups is addressed in the concluding section of 5.6.

majority carrier of volume. In May 2026, BlackRock filed two additional tokenized money market funds with the SEC. Sources: BlackRock/Securitize (fund documentation); Token Terminal (AUM status May 2026).

[351] No-Action Letter from the SEC Division of Trading and Markets to the Depository Trust Company of 11 December 2025: three-year pilot ("Preliminary Base Version") of the DTCC Tokenization Services, mapping security entitlements on US Treasuries, Russell-1000 equities, and selected index ETFs as tokens on approved, including permissionless, blockchains; start scheduled for the second half of 2026. Caveats: tokenized entitlements receive no settlement or collateral value in the pilot, the DTC remains the source of settlement finality and retains the final record; the first announced implementation runs with Digital Asset on the Canton Network. The letter opens the category regulatorily without assigning it to a network. Source: SEC/DTCC (December 2025).

5.6-E Research Horizon: Heterogeneous Chains and Autonomous Agents

The question with which the preceding section closed reaches beyond the target state, and that changes the status of the answer. No specification in the roadmap answers whether the converged layer also carries systems that are not Rollups, or serves actors who are not humans. Both extensions are regarded in the roadmap's environment as a goal of the layer, and precisely for that reason their presentation requires the explicit marking of their status. This section therefore distinguishes two quantities: the architecture developed in the preceding sections and here only applied, and the adoption by new systems and actors — which carries no maturity level and is carried as a research horizon. The verdict of 5.6, which closed with the preceding section, depends on none of the following statements.

The first question concerns chains with their own consensus that would wish to bind themselves to the layer without becoming Rollups. Structurally, the weaker part of the offering would be open to such a system: it could periodically record its block history in the layer and would gain a finality that its own consensus does not produce. Disputes over its own past can thus be resolved against an external standard whose neutrality the system need not provide itself. What it does not gain is the property of the converged Rollup, because verification of its state transitions remains its own, and with it every assumption of its consensus. An economic binding is also conceivable — for instance through collateral deposited on the layer — but this would supplement rather than replace the assumptions of the foreign consensus. The binding of heterogeneous systems therefore remains economic or institutional, while Rollups receive a cryptographic one, and the two-class problem that the section closed for Rollups would return at their external boundary in another form. The ecosystem's discourse nonetheless carries the layer as the foundation of entire chain ecosystems — claiming the scope of a civilizational infrastructure whose significance reaches beyond any single application.[352] What mechanisms could carry such a role and whether foreign consensus systems would adopt it is not laid out in the target state and remains a subject of open research.

[352] The framing of the layer as civilizational infrastructure traces to Vitalik Buterin's New Year's post on X (@VitalikButerin, 1 January 2026), which frames Ethereum beyond short-term market and tokenization narratives as a durable, neutral, and censorship-resistant foundation for finance, identity, and governance across decades, and measures it against the standard of the walkaway test. The extension to the layer as the foundation of entire chain ecosystems is a position of the broader ecosystem discourse, which the body reports while also limiting, not a statement of this work. Sources: V. Buterin, New Year's post on X (1 January 2026). CoinDesk, Vitalik Buterin on the two goals Ethereum must meet to become the world computer (1 January 2026).

Autonomous Agents as Users

The second question concerns an actor class whose needs present the requirement profile of binding Settlement in sharpened form. Autonomous agents — software systems that commission, provide, and pay for services without human approval — are excluded from bank accounts, signing of contracts, and legal entities. A booking location without institutional precondition (Section 5.6-D) would not be the better option among several for such actors — it would be the only one. Machine counterparties could extend no credit from relationship to each other. Every service would require settlement delivery against delivery, whose guarantee is verified rather than negotiated. The amounts are small, the frequency high, and the counterparty changing — a profile for which institutional payment infrastructure provides neither accounts nor prices. An agent would hold its own balance, provide collateral from its own holdings, and require no custodian acting on its behalf for either. The path for this is not outstanding: the Intent layer (Section 5.6-D) carries the coordination and payment between agents without any human intermediation between commission and settlement. The layer records what satisfies its rules, regardless of who declares a result, and precisely therein lies the accessibility for actors without legal form.

The guarantees of the layer would be machine-readable, their deadlines nameable, and their access condition exhausted in the possession of a key — a combination no institutional settlement system offers. That confirmation is available in seconds and deterministically — a function of the Finality Stack (Section 5.6-C) — matches the deadline to the frequency of machine interaction. Added to this would be requirements that human users do not pose: an identity tied to no person, a reputation that is machine-verifiable, and a verification of rendered services without courts and without contractual partners. For the first of these, the answer already stands in the chapter, because the account architecture from Section 5.3 binds authorization to programmable rules rather than to a person. Draft ERC-8004 normalizes identity, reputation, and validation registers on the layer for this purpose — an explicitly early reference point documenting a direction.[353] What the draft does not carry is the question of liability and over-

[353] ERC-8004 („Trustless Agents"), Draft standard since August 2025, emerged in the Ethereum Foundation environment with participation from MetaMask, Google, and Coinbase; builds on the Agent-to-Agent protocol (A2A) and supplements it with three on-chain registries: identity (as ERC-721), reputation, and validation. The registries are deployed. At DevConnect in November 2025 first implementations were shown, the number of participating developers was reported at 1,000 to 2,000 by Davide Crapis (Ethereum Foundation). Classification by this work: Draft with deployed registries, above pure research status, below productive maturity; the standard documents a development direction, not an established demand from the actor class. Sources: ERC-8004 specification (Ethereum Magicians, August 2025); DevConnect (November 2025).

sight, since an actor without legal form is also an actor against whom no claim is enforceable. The structural finding does not depend on this maturity level: the same layer carries two demand classes under the property bundle that the preceding section derived for tokenized capital, and the machine actors demand the same bundle for the same reasons — only without the alternative. Whether the actor class becomes a demand class depends on developments outside the protocol, and no statement in the chapter gains or loses with it.

Both questions share the same status: they lie beyond the target state specified by the roadmap, and beyond the maturity levels with which this work substantiates its findings. The section records them because the architecture implies them, and it does not answer them because the assumption rests on no specification against which an answer could be verified. The open variable in this context is their adoption by systems and actors over which the protocol has no authority. For the overall picture of the section, the architecture already connects both dimensions of the settlement role. The target image is in this respect architecturally complete. What remains open is only whether the second demand class materializes. The assessment in the following section therefore applies to the target state developed in sections 5.2 through 5.6, not to the horizon this section marks.

5.7 Target-State Assessment

5.7-A Target-State Assessment of the Critical Conditions

The preceding sections have unfolded the target state along its technical directions, from execution through verification and state to neutrality and the architecture of Layer 2. The following assessment shifts register and maps what has been unfolded onto the twelve criteria of the evaluation framework, beginning with the three Critical Conditions whose rating caps the overall verdict. It applies to the target state as an architecturally coherent design state, not to the operational current state. Throughout, it carries the maturity levels of the taxonomy: from delivered deployment maturity (DEPL, IMPL) through planning anchoring (PLAN) to unproven research (RES), since no classification may exceed its maturity level. The assessment is simultaneously reduced to the architectural level, as the current-state comparison, the full three-level synthesis, and robustness against partial realization belong in Chapter 6.

I.2 Security and Trust Load

Security and Trust Load (I.2) reaches the rating Met with Qualification in the target state.[354] The relay elimination through ePBS unfolded in 5.5 carries the trust-load reduction on the consensus layer, so that the relay trust assumption of the current state — through which the majority of blocks are routed today — falls away in the target state. The zkEVM verification developed in 5.3 reduces the validator trust assumption by proving the correctness of execution cryptographically rather than assuming it socially, while the Stateless Clients from the same verification architecture dissolve the dependency on commercial RPC providers. Mapping the aforementioned mechanisms to the criterion yields a systematic reduction of trust assumptions spanning nearly every protocol layer, whose epistemic robustness is graduated from the delivered inter-layer simplification through to the planning-anchored ePBS activation.

The rating remains below Met because two residual risks are not reached by protocol changes. The first lies in the zkEVM itself, whose proving security rests on mathematical assumptions with proofs still open. Verification shifts trust to the integrity of the underlying assumptions rather than dissolving it. The second lies on the access layer, since the trust load of wallets, frontends, and DNS resolution lies outside the protocol and is not reached by any roadmap component, because the majority of users interact through intermediary applications rather than their own nodes. Both risks mark the boundary of protocol-side trust reduction and carry the qualification of the rating.

I.4 Minimal Load-Bearing Guarantees

The broadest of the three criteria is the Minimal Load-Bearing Guarantees (I.4), which encompasses four aspects: Finality, Liveness, a Degradation Mode, and censorship resistance under attack. The four aspects are concentrated in the preceding sections — Finality and Liveness in the finality architecture from 5.5, the Degradation Mode in the Inactivity Leak from 5.3-D, and censorship resistance in the neutrality block from 5.5 — so that a brief synthesis forward suffices for them. Two further substances cut across the four aspects and are carried outside the group of four as supplementary material: state persistence, which 5.4 unfolds as a coherent section, and post-quantum robustness, which is distributed across

[354] Target-state assessment I.2. The rating Met with Qualification follows from the graduated trust-load reduction spanning nearly every protocol layer, carried by relay-trust elimination via ePBS (PLAN/SFI Glamsterdam), validator-trust reduction via zkEVM (active EF roadmap), and RPC-trust reduction via Stateless Clients and Helios (DEPL/IMPL). The two residual risks not reachable by the protocol are the zkEVM proving conjecture and the access layer (wallets, frontends, DNS).

several sections and requires its own collection to make the material visible. The greater depth of the post-quantum collection follows from the distribution of its material.

The finality architecture unfolded in 5.5 reinforces the guarantee in the target state in multiple layers, graduated by maturity level. At the Layer 2 level, Based Preconfirmations deliver confirmations in the sub-two-second range, the Confirmation Rule (FCR) unfolded in 5.6 shortens consensus confirmation to approximately 13 seconds, and three-slot finality binds the definitive Layer 1 finality to approximately 36 seconds.[355] The M2 threshold of under 15 minutes is already far undercut by the more mature layers: the deployed Based Preconfirmations and the Confirmation Rule, specified as a consensus-specs pull request and currently in client implementation, bring Layer 1 confirmation into the seconds range, independently of the maturity level of three-slot finality. The long-term direction toward Single-Slot Finality remains nameable as a consensus end-state without determining the near-term target value. Against the criterion, the gradation means that the finality guarantee is carried redundantly across layers of differing maturity — from the deployed Based Preconfirmations through the Confirmation Rule in client implementation to three-slot finality in research status — rather than depending on a single mechanism.

Liveness — the ability to continue producing blocks under stress — is reinforced in the target state through the same finality architecture, because the graduated confirmation layers no longer bind block production to a single mechanism, and Single-Slot Finality explicitly retains the Inactivity Leak unfolded in 5.3-D as a fallback mechanism that keeps the chain running if more than one-third of validators fail. The Degradation Mode rests on that same Inactivity Leak, which is tested and operational on mainnet and carries the orderly fallback in the event of non-participating validators; it operates more granularly under Single-Slot Finality, so the guarantee is maintained and sharpened in the target state. Censorship resistance is protocol-side enforced through FOCIL, whose assessment as a mechanism stands in II.1 and is credited here only as a minimum guarantee. The cross-reference is methodologically necessary because I.4 defines the existence of the guarantee as a requirement and II.1 evaluates its fulfillment.

Across the four assessed guarantees — Finality, Liveness, Degradation Mode, and censorship resistance — I.4 reaches the rating Met, because all four are

[355] Based Preconfirmations approximately 2 seconds (DEPL, Layer 2 level) and three-slot finality approximately 36 seconds (RES, horizon beyond 2029, without fork date) per the finality architecture carried in 5.5. The Confirmation Rule (FCR) approximately 13 seconds is specified as consensus-specs PR #4747 (March 2026), is in client implementation with rollout in the coming months, and therefore carries the status PLAN, not IMPL.

strengthened or maintained in the target state and no structural limitation caps the rating.[356] State persistence and post-quantum robustness do not carry the rating.

II.1 Neutrality and Censorship Resistance

Neutrality and Censorship Resistance (II.1) carried the rating Conditionally Met in the current state as the sole Critical Condition, constituting the most serious individual finding of this work. In the target state its foundation changes, as censorship resistance rises from an emergent market property to a protocol-side enforced, three-layered system. The inclusion enforcement through FOCIL unfolded in 5.5, the relay elimination through ePBS, and the pre-inclusion privacy through the Encrypted Mempool interact complementarily. The 1-of-N honesty model of FOCIL carries the formally stronger guarantee, because a single honest committee member enforces inclusion. Mapping the three-layered architecture to the criterion elevates the rating from Conditionally Met to Met with Qualification, and the step is necessary because it removes the cap imposed by the second cascade step.[357]

The cap at Met with Qualification has three reasons, all of which lie outside the roadmap's reach. Builder concentration remains an operational reality whose structural causes — from latency advantages to exclusive order flow — ePBS does not address, because relay elimination removes the intermediaries without restructuring the economics of block production. Without an implemented protocol cap, the concentration of liquid staking tokens, led by Lido holding approximately 23 percent of staked ETH today, is carried forward into the target state with a declining tendency, constituting a structural limitation.[358] The geographic node distribution — with approximately 39 percent of nodes in the United States — improves technically through the hardware accessibility of Stateless Clients, without the roadmap reaching the non-technical causes of the concentration. The three limitations form the bridge to Chapter 6, which car-

[356] Target-state assessment I.4. The four I.4 aspects (Finality, Liveness, Degradation Mode, and censorship resistance under attack) are strengthened or maintained in the target state and assessed in place; state persistence and post-quantum robustness cut across the four aspects, the former in the assessment of Long-Term Stability, the latter inventoried in the closing collection as material. FOCIL carries PLAN/SFI Hegotá headliner.

[357] Target-state assessment II.1. Current-state rating Conditionally Met as the sole Critical Condition, target-state rating Met with Qualification. Maturity levels: ePBS PLAN/SFI Glamsterdam, FOCIL PLAN/SFI Hegotá headliner, Encrypted Mempool RES.

[358] At the infrastructure level of II.1: Lido approximately 23 percent of staked ETH, declining from the 2023 peak of over 32 percent, without implemented protocol cap, geographic node distribution approximately 39 percent United States.

ries the persistent boundaries, and they justify the rating without being further assessed here.

With the assessment of the three Critical Conditions, the section shifts into an inventory mode that makes the post-quantum material distributed across five sections visible as a body of evidence. The post-quantum migration is, in the design state, a distributed architectural movement across the three layers of Consensus, Data, and Execution. At the Consensus layer, the hash-based replacement of BLS validator signatures lies in 5.5-D, with the Beam Chain from 5.1-B as the long-term end-state of SNARK-based consensus proofs. At the Data layer, 5.2-D presents the transition from curve-based KZG commitments to hash-based schemes in the Full Danksharding end-state. The migration path at the Execution layer moves away from ECDSA through Native Account Abstraction, with the Validation Frames from 5.3-E as the primary opt-in vehicle for the ECDSA exit. As a common load-bearing layer, 5.3-A′ locates the hash-based aggregation — a recursive proof aggregation that is to be treated as a compute and cost problem, not to be confused with signature aggregation in live consensus, whose character remains a problem of network coordination.

The maturity level of the distributed movement can be inventoried by status without promising a date. Deployment-capable is the secp256r1 precompile (EIP-7951) — curve-based and itself not quantum-resistant — which diversifies the signature schemes through FIDO2 and WebAuthn Passkeys and provides the infrastructure for the migration path. The Helios light client stands at the same maturity level, as a verification precursor. The Validation Frames (EIP-8141), by contrast, stand at an earlier stage — the specification phase of the Hegotá fork as a placeholder commitment per Checkpoint #9 — and mark the ECDSA migration path.[359] The preparation is anchored in the Ethereum Foundation's post-quantum research team, founded in January 2026 and working with biweekly ACD calls, published research prizes, and live testnets. Added to this are the generic precompile position EIP-7932 in research status, Binary Trees as a STARK-compatible state-tree architecture, and the Emergency Response Plan signaled as a backstop.[360] The gaps remain visible nonetheless: the full-stack migration across the Consensus Layer, the replacement of KZG commitments with STARKs, and the transition to Binary Trees together constitute a generational project whose Beam

[359] PQ preliminary material 5.7-A. secp256r1 precompile EIP-7951 (DEPL), Helios light client (DEPL), Frame Transactions EIP-8141 (CFI Hegotá Non-Headliner, placeholder commitment per Checkpoint #9).

[360] PQ preliminary material 5.7-A. Post-quantum research team of the Ethereum Foundation founded January 2026, EIP-7932 (DFI Glamsterdam, still in research status), Binary Trees (RES, signaled in the EF Protocol Priorities Update of 18 February 2026), Emergency Response Plan from the Vitalik PQ roadmap post.

Chain timeline lies beyond 2029 without a fork date. The choice of hash function is moreover an upstream decision affecting multiple layers simultaneously, and a protocol-wide migration sequence remains uncoordinated. The signature sizes of post-quantum schemes — ranging from 2.4 to 4.6 kilobytes for ML-DSA and from 7.9 to 49.9 kilobytes for SLH-DSA against 64 bytes for ECDSA — create an open state question with a State Bloat on the order of 59 times for ML-DSA.[361] The Metaculus estimate referenced in Buterin's post of 26 February 2026 puts the probability of cryptographically relevant quantum computers before 2030 at approximately 20 percent, cited as a maturity-level context rather than as the work's own probability assessment.[362]

The Degradation Mode unfolded in 5.3-D is added as robustness material. The Inactivity Leak progressively withdraws stake from non-participating validators until the participating ones regain a two-thirds supermajority. It is tested and operational on mainnet. Under a quantum attack arriving faster than the migration, the stress mode would engage, allowing the network to continue producing blocks while compromised validators exit through the Leak, with Single-Slot Finality enabling the trigger to engage more granularly and earlier. The availability of the stress mode is inventoried here only as present and enters as material into the I.4 robustness assessment in Chapter 6.

The inventory names the distributed architecture, its graduated maturity level, and the available Degradation Mode, without ruling on its adequacy. Section 6.1 decides whether the migration suffices, whether the timeline holds, and whether the gaps are critical, not by the present section. The collection provides the material for that judgment.

5.7-B Target-State Assessment of the Structural and Qualitative Criteria

Following the three Critical Conditions, whose rating caps the overall verdict, the Structural Conditions are added. They determine the system's load-bearing capacity beyond the Critical Conditions and constrain the overall verdict in the third cascade step once two or more of them stand weak. They are assessed in the order I.1, I.3, and III.1.

[361] Strawmap (Post-Quantum L1 North Star), referenced in PQ preliminary material 5.7-A. Signature sizes ML-DSA 2.4 to 4.6 kilobytes and SLH-DSA 7.9 to 49.9 kilobytes against 64 bytes for ECDSA, State Bloat on the order of 59 times for ML-DSA.

[362] Vitalik PQ roadmap post of 26 February 2026, referenced Metaculus estimate. Cited maturity-level context, not the work's own probability forecast.

The Structural Conditions

Ethereum's irreplaceability as a settlement layer reaches the highest rating on all three anchor indicators in the target state, which is why I.1 is to be assessed as Met.[363] The Native Rollup binding unfolded in 5.6 supports the assessment asymmetrically. It couples L2 systems to Ethereum at the protocol level and creates an anchoring depth that leaves the system as the Schelling point of the settlement layer with no technical switching option remaining.[364] At the infrastructure level, the settlement attraction through ERC-7683 likewise described in 5.6 reinforces the binding, because it draws solver networks economically toward the place where liquidity and settlement security coincide. The capacity scaling unfolded in 5.2 further supports the anchoring by widening the space in which this activity takes place.

One limitation remains nameable without pushing the rating down. The cloud concentration of node infrastructure is not addressed directly by the roadmap: approximately 59 percent of hosted Execution Layer nodes are distributed across three providers, with approximately 35.5 percent on AWS, approximately 13.8 percent on Hetzner, and approximately 9.7 percent on OVHcloud, with a declining tendency since 2022.[365] The finding qualifies the depth of the anchoring but does not change the rating. I.1 tests the anchoring as a Schelling point, and the concentration on external providers does not negate this anchoring. It shifts part of the dependency onto an infrastructure layer that the criterion does not measure, which is why it appears as a qualifying factor.

The Coordination Function expands qualitatively in the target state and fulfills I.3 with a broader reach than the starting state.[366] What previously coordinated primarily settlement and data availability for L2 networks now extends to execution, verification, and the governance of MEV economics, as sections 5.5 and 5.6 unfold. In doing so, the protocol absorbs functions that currently require external trust assumptions. The relay coordination of block production shifts protocol-side into ePBS; the verification of L2 state transitions into the Native Rollup binding to Ethereum's own execution.

---

[363] The three anchor indicators of Functional Irreplaceability (DeFi TVL dominance, stablecoin issuance, developer ecosystem) each reach the highest anchor level above the 50 percent threshold in the target state.

[364] On the mechanics of Native Rollups (maturity level RES) and the settlement attraction via ERC-7683, cf. 5.6; on capacity scaling, cf. 5.2.

[365] Status of Execution Layer node hosting: approximately 59 percent on three providers (AWS approximately 35.5%, Hetzner approximately 13.8%, OVHcloud approximately 9.7%), declining since 2022. Cf. the survey in Chapter 4.

[366] On the addition of protocol-side verification as the fourth coordination primitive alongside settlement, execution, and data availability, cf. 5.5 and 5.6. Maturity levels: ePBS PLAN/SFI (Glamsterdam), Native Rollups RES.

The coordination capacity remains bound to the specific Ethereum ecosystem through network effects, so that switching to a competing system would entail the loss of accumulated network effects. In the target state, all top-10 L2s settle on Ethereum, which keeps the indicator at the highest level, and the Native Rollup implementation raises switching costs structurally by removing from L2s the technical option of changing their settlement layer.[367] I.3 is Met, with expanded reach, without a qualification pushing the rating down.

The target state responds to Long-Term Stability with three decouplings that fulfill III.1.[368] The zkEVM decouples verification from execution capacity, Stateless Clients decouple verification from state size, and Tiered State decouples permanent state growth from total capacity. Each of the decouplings arises at the location of its mechanism, which sections 5.2, 5.3, and 5.4 respectively unfold.

The finding is supported by the upgrade track record. Shanghai, Dencun, Pectra, and Fusaka were delivered without mainnet-critical incidents, and with Fusaka the semi-annual hard fork schedule was met for the first time.[369] State growth, client diversity, and upgrade robustness are all carried in the target state, so the rating Met is reached.

III.1 also carries a named limitation that does not push the rating down. Validator concentration does not change structurally, and no roadmap feature addresses it directly, so Lido is expected to remain at approximately 23 percent, clearly below the 33 percent threshold, with a declining tendency.[370] The factor is named but does not carry the rating. III.1 tests Long-Term Stability against state growth, client diversity, and upgrade robustness, and these three properties are carried by the target state, while stake distribution touches a different question that the criterion does not decide here.

---

[367] All top-10 L2s settle on Ethereum in the target state (highest level). The Native Rollup implementation removes from implementing L2s the technical option of changing their settlement layer.

[368] The three decouplings arise at the location of their mechanism: zkEVM (active EF roadmap, 5.2), Stateless Clients via the Verkle or Binary Trees path (both RES. The migration bundle EIP-7612/7748/4762 remains invariant to the tree choice, 5.3), Tiered State (RES, 5.4).

[369] Shanghai, Dencun, Pectra, and Fusaka were delivered without mainnet-critical incidents. With Fusaka the semi-annual hard fork schedule was met for the first time. State relief is additionally provided by History Expiry (EIP-4444, Phase 1 DEPL) and the SELFDESTRUCT limitation (EIP-6780, DEPL).

[370] Validator concentration remains structurally unchanged: without a protocol-level staking cap, Lido is expected to remain at approximately 23 percent and clearly below the 33 percent threshold, with a declining tendency. No roadmap feature addresses the concentration directly.

The Qualitative Criteria

Following the Critical and Structural Conditions, the six Qualitative Criteria differentiate the degree within the reached suitability category without capping the overall verdict downward. Open Generativity (II.2) provides the first response. The target state expands the generative capacity of the system on three levels, whose mechanics the preceding sections have unfolded. They range from computational capacity via zkEVM and Full Danksharding, through user experience via Native Account Abstraction, to the generation of new Layer 2 networks via Native Rollups. The composability guarantee and EVM stability are preserved, protecting existing developer investments. The qualification rests on two open points that hold the rating: RISC-V as an alternative execution runtime remains an exploratory research path without governance anchoring, and the foregoing of EOF omits an improvement to the developer experience that would have been structurally possible. In the target state, Open Generativity stands at Met with Qualification.[371]

Independent Verifiability (II.3) fulfills all three indicators in the target state, and the formal specification basis improves structurally, because the Execution Layer Specification as an executable reference implementation forms the operational specification standard for new protocol features. The one critical qualification that holds the rating arises from the zkEVM itself: it shifts the verification task from re-execution to the checking of cryptographic proofs, and the formal correctness of the prover — which becomes the new trust foundation of the stack — is not yet formally provable in the target state. A complete prover circuit exceeds, in its complexity, the capacity of current formal methods, so that the verification offering remains anchored at the protocol level but carries an unsolved problem at its most critical point. Independent Verifiability stands at Met with Qualification.[372]

In Low-Threshold Inclusivity (II.4), the target state lowers participation barriers for three of the four roles. Users benefit from Native Account Abstraction and Passkey authentication via the secp256r1 precompile, verifiers from Stateless Clients that no longer need to maintain the full state locally, and developers from a stable toolkit protected by the EVM guarantee. The qualification lies

[371] On maturity levels: the zkEVM stands on the active EF roadmap, Full Danksharding on RES, Native Account Abstraction via EIP-8141 on PLAN as CFI Non-Headliner for Hegotá, with the already activated EIP-7702 (DEPL) as operational evidence, Native Rollups on RES. RISC-V as a VM alternative is an exploratory research path (RES) without governance anchoring, EOF was withdrawn from the Fusaka roadmap.

[372] The Execution Layer Specification is an executable reference implementation of EVM semantics and is the operational specification standard for new features in the target state. Formal verification of the zkEVM prover circuit stands at RES and exceeds the capacity of current formal methods.

in the fourth role: the 32-ETH threshold for solo staking remains in place, and the target state mitigates it only indirectly through hardware relief for validators, without lowering the economic barrier itself. Low-Threshold Inclusivity stands at Met with Qualification.[373]

Adaptive Governance (III.2) is the first of the Qualitative Criteria to carry the full rating. The governance model demonstrates adaptive capacity on all three levels, and the strategic pivot unfolded in 5.1 is the strongest evidence — a paradigmatic change of direction within the existing model without crisis, fork, or formal vote. Because the property is structurally and historically documented and does not depend on a future roadmap delivery, Adaptive Governance stands at Met.[374]

Sovereign Portability (III.3) carries the same type of finding, because the target state systematically dissolves the proprietary lock-ins of the starting state: the relay dependency through ePBS, the L2 proof systems through Native Rollups, and the RPC concentration through Trustless RPC. The EVM remains a bilateral structural lock-in that binds developers to Ethereum and Ethereum to the EVM, without pushing the rating down, so Sovereign Portability stands at Met.[375]

Hardware Agnosticism (III.4) requires an explanation. In the target state, the architectural hardware agnosticism of the protocol is complete — without ASIC dependency and with ARM compatibility across all major clients. Verification is decoupled from state size by Stateless Clients, so it remains feasible on consumer hardware up to and including smartphones and browsers. The one qualification that remains is cloud concentration, with approximately 59 percent of hosted Execution Layer nodes on three providers — a value declining since 2022 that the roadmap does not address through any direct feature. Decisive for the rating is that the target state does not worsen the value: Stateless Clients lower the hardware barrier for home staking without removing the economic drivers of cloud usage, so the concentration remains unchanged. An unchanged indicator cannot fall below its starting rating, and since the current-state assessment in Chapter 4

[373] User-side maturity levels: EIP-7702 (DEPL) and the secp256r1 precompile via EIP-7951 (DEPL) for Passkeys. The Stateless Client migration via Verkle Trees stands at PLAN. The 32-ETH threshold is not directly lowered by any roadmap feature; hardware relief occurs indirectly via zkEVM-based attestation.

[374] The strategic pivot from a rollup-centric to an L1-oriented scaling strategy (February 2026) is demonstrated and counts as an established fact. The other governance properties are structurally and historically documented. Cf. section 5.1.

[375] Maturity levels of lock-in dissolution: ePBS at PLAN (SFI, Glamsterdam), Native Rollups at RES, Trustless RPC at RES/PLAN. The EVM remains a structural lock-in of the protocol.

carries Hardware Agnosticism at Met with Qualification, the target state maintains this rating rather than falling to Conditionally Met.[376]

With the assessment of all twelve criteria in the target state, the complete target-state profile is visible.

5.7-C The Target-State Profile

Target-State Profile Map

Stufe	Kriterien	Anzahl
Erfüllt	I.1, I.3, I.4, III.1, III.2, III.3	6
Erfüllt mit Einschränkung	I.2, II.1, II.2, II.3, II.4, III.4	6
Bedingt erfüllt	keine	0
Offen	keine	0

Table 5.1 The target-state profile map: the four assessment levels with their associated criteria and their number.

Target-state profile: 6-6-0-0.

All six Qualitative Criteria stand in the target state at at least Met with Qualification, which, according to the degree scale from Chapter 2, carries the degree Good.[377] The complete profile and its evaluation through the M3 cascade to the overall verdict is carried by the following section.

[376] The current-state assessment of Hardware Agnosticism stands in the finalized Chapter 4 at Met with Qualification; cloud concentration there at approximately 59 percent of hosted Execution Layer nodes across three providers (AWS 35.5%, Hetzner 13.8%, OVHcloud 9.7%), declining since 2022 (Ethernodes, early 2026). Since the target state neither worsens nor directly improves cloud concentration, the criterion maintains the current-state rating. Architectural support (x86 and ARM) and Stateless Client verification are treated in sections 5.3 and 5.7-B.

[377] The degree scale from Chapter 2 defines the degree Good when all six Qualitative Criteria stand at at least Met with Qualification. The current-state assessment in Chapter 4 awards the same degree at an identical criterion position. Cf. section 2.3 and Chapter 4, overall synthesis.

5.8 Overall Verdict

The assessment of the twelve criteria in the preceding section yields the complete target-state profile, which the M3 cascade now leads to the overall verdict. Six criteria stand at Met in the target state, the other six at Met with Qualification, and no criterion falls to Conditionally Met or Open. Among the three Critical Conditions, the Minimal Load-Bearing Guarantees (I.4) achieve full Met, while Security and Trust Load (I.2) and Neutrality and Censorship Resistance (II.1) stand at Met with Qualification. On the structural side, Functional Irreplaceability (I.1), the Coordination Function (I.3), and Long-Term Stability (III.1) carry Met throughout, so the numerical profile reads 6-6-0-0.[378]

The cascade examines this profile in five successive steps, the first three of which can cap the overall verdict downward. The first check addresses the question of whether any of the three Critical Conditions stands at Open, which would limit the verdict to Not Suitable. None of the three stands there, since I.2 and II.1 carry Met with Qualification and I.4 carries Met. The cap at Not Suitable does not apply.

The second check asks whether a Critical Condition stands at Conditionally Met, which would qualify the verdict as Suitable under Substantial Conditions. In the current state, Neutrality and Censorship Resistance (II.1) carried the rating Conditionally Met as the sole Critical Condition and triggered the cap at the middle suitability category. In the target state, II.1 rises to Met with Qualification, because the three-layered system unfolded in 5.5 elevates censorship resistance from an emergent market property to a protocol-side enforced guarantee. The second check therefore no longer applies, and the cap that held the current-state verdict falls away.

The third check asks whether two or more Structural Conditions stand weak. Since I.1, I.3, and III.1 all stand at Met, it does not apply either. Since none of the three capping checks applies, the verdict lies above the middle category: neither at Not Suitable nor at Suitable under Substantial Conditions. Since I.2 and II.1 remain below full Met, simple Suitable is not reached. What remains is Suitable with Conditions.[379]

Within the reached category, the six Qualitative Criteria determine the degree. According to the degree scale from section 2.3.3, this carries the degree Good, and

[378] Cf. the profile map in 5.7-C. Met is carried by I.1, I.3, I.4, III.1, III.2, and III.3; Met with Qualification by I.2, II.1, II.2, II.3, II.4, and III.4. Conditionally Met and Open remain unoccupied.

[379] On the M3 cascade and its five examination steps, cf. section 2.3. In the current state, the second check set the verdict to the middle category because II.1 was the sole Critical Condition standing at Conditionally Met (cf. Chapter 4, overall synthesis).

the finalized current-state assessment in Chapter 4 awards the degree Good at an identical threshold position, so the target state carries the same degree.[380] The degree remains the same even though two Qualitative Criteria improve: Adaptive Governance (III.2) and Sovereign Portability (III.3) rise from Met with Qualification to full Met. Good is the highest level of the scale. Both improvements therefore leave no expression in the overall verdict, because Qualitative Criteria alone determine the degree and it was already secured in the starting state. The categorical comparison with the current state reads precisely: one verdict step higher at the same degree. The progress of the roadmap lies in the foundation of suitability, not in the fine quality of the assessed properties.

A finding emerges between the rating-bearing individual assessments, accumulates from their collective view, and belongs to none of them, which is why the cascade does not produce it on its own. The structural concentrations persist because no roadmap feature directly removes them. Builder concentration in block production is not reached by ePBS at its economic root, and the concentration of liquid staking around Lido remains without a protocol cap. The informal governance sharpens the finding, because its emergency response rests on documented action plans rather than on protocol-enforced mechanisms. Together, the three factors form a crisis-response gap that appears in no individual rating-bearing criterion and cuts across Neutrality (II.1), Long-Term Stability (III.1), and Adaptive Governance (III.2).[381] The system's capacity to respond to urgent concentration and crisis risks rests on social coordination that the protocol does not enforce, and its depth and its weight for suitability is carried by Chapter 6.

The overall verdict Suitable with Conditions, Degree: Good, rests on architectural coherence at graduated implementation maturity. The cryptographic load-bearing capacity shows this gradation most clearly. The post-quantum migration is architecturally distributed by design, with the verification layer as its anchor and the Beam Chain as its long-term end-state, and this distribution supports the load-bearing capacity of the architecture without securing it.[382] Whether the roadmap delivers on the conditions robustly, whether partial realization holds, and whether the gaps are critical is decided by the integrated robustness assess-

[380] The finalized current-state assessment in Chapter 4 awards the degree Good when all six Qualitative Criteria stand at at least Met with Qualification. The target state fulfills them at an identical threshold position. Cf. Chapter 4, overall synthesis.

[381] The concentration values are carried by Chapter 4 (builder concentration in block production) and 5.7 (Lido approximately 23 percent of staked ETH, declining, without protocol cap). On the Emergency Response Plan as a documented rather than protocol-enforced emergency mechanism, cf. 5.7-A.

[382] On the distributed post-quantum architecture with the verification layer as its anchor and the Beam Chain as its long-term end-state, and on its graduated maturity level, cf. the inventory in 5.7-A. Its adequacy is assessed by Chapter 6.

ment in section 6.1, which at the same time carries the persistent boundaries. The present section carries the verdict whose robustness section 6.1 examines.

Chapter 6 Delta, Limits, and Overall Assessment

6.1 Delta: Current State → Target State

The overall assessment of the target state stands one suitability category above the finding that Chapter 4 reached for the operational baseline. The baseline carries Suitable under Substantial Conditions with a current-state profile of 0-1 1-1-0. The target state carries Suitable with Conditions with a target-state profile of 6-6-0-0. The distance between the two assessments measures what the roadmap moves in terms of suitability.[383] The comparison that makes the answer to the work's central question out of the assessment reached in 5.8 was withheld from the target-state evaluation and assigned to Chapter 6. The following juxtaposition of the twelve criteria carries it out and simultaneously lays bare what each improvement depends on.

[383] Current-state profile 0-1 1-1-0 and overall assessment Suitable under Substantial Conditions: Chapter 4, overall synthesis. Target-state profile 6-6-0-0 and overall assessment Suitable with Conditions: 5.7-C and 5.8. The grade Good follows in both states from the threshold defined in Chapter 2 and applied in Chapter 4, under which all six Qualitative Criteria stand at a minimum of Met with Qualification.

Kriterium	IST (Kapitel 4)	SOLL (5.7)	Kaskaden-Rolle
I.1 Funktionale Unersetzbarkeit	Erfüllt mit Einschränkung	Erfüllt	Strukturell
I.2 Sicherheits- und Vertrauenslast	Erfüllt mit Einschränkung	Erfüllt mit Einschränkung	Kritisch
I.3 Koordinationsfunktion	Erfüllt mit Einschränkung	Erfüllt	Strukturell
I.4 Minimale tragfähige Garantien	Erfüllt mit Einschränkung	Erfüllt	Kritisch
II.1 Neutralität und Zensurresistenz	Bedingt erfüllt	Erfüllt mit Einschränkung	Kritisch
II.2 Offene Generativität	Erfüllt mit Einschränkung	Erfüllt mit Einschränkung	Qualitativ
II.3 Unabhängige Verifizierbarkeit	Erfüllt mit Einschränkung	Erfüllt mit Einschränkung	Qualitativ
II.4 Niedrigschwellige Inklusivität	Erfüllt mit Einschränkung	Erfüllt mit Einschränkung	Qualitativ
III.1 Langfristige Stabilität	Erfüllt mit Einschränkung	Erfüllt	Strukturell
III.2 Adaptive Governance	Erfüllt mit Einschränkung	Erfüllt	Qualitativ
III.3 Souveräne Portabilität	Erfüllt mit Einschränkung	Erfüllt	Qualitativ
III.4 Hardware-Agnostik	Erfüllt mit Einschränkung	Erfüllt mit Einschränkung	Qualitativ

Table 6.1 The twelve criteria in the transition from the current state of Chapter 4 to the target state of Section 5.7, with their role in the cascade.

Current-state profile 0-1 1-1-0, target-state profile 6-6-0-0.

Seven of the twelve criteria improve, five remain constant, and four movements carry the elevation of the overall assessment. The seven improvements encompass Functional Irreplaceability (I.1), the Coordination Function (I.3), Minimal Load-Bearing Guarantees (I.4), Neutrality and Censorship Resistance (II.1), Long-Term Stability (III.1), Adaptive Governance (III.2), and Sovereign Portability (III.3). Of these, only Neutrality and Censorship Resistance (II.1) leaves the zone of Conditionally Met and rises to Met with Qualification. The remaining six move from Met with Qualification to Met. Three of them — Functional Irreplaceability (I.1), the Coordination Function (I.3), and Long-Term Stability (III.1) — are Structural Conditions, and with this ascent they lift the ceiling of the third cascade step. The rise of II.1 lifts the ceiling of the second cascade step; the ascent of I.1, I.3, and III.1 lifts that of the third. Only both together advance the overall assessment by one category. The three remaining movements lift neither the category nor the threshold-based grade, which was already secured at Good in the baseline. As noted in 5.8: one assessment step higher at the same grade, carried by four movements.

The Change of Foundation

The carrying movement draws its substance from the change of the foundation on which censorship resistance rests. In the baseline it was an emergent market property, with OFAC-compliant block production falling from 79 percent in November 2022 to roughly 15 percent — market-driven, without protocol-level enforcement. The assessment therefore classified it as reversible and the criterion as Conditionally Met. In the target state, the three-layer system unfolded in 5.5 — FOCIL, ePBS, and Encrypted Mempool — enforces inclusion at the protocol level. The guarantee shifts from a market condition to a property sustained by the 1-of-N Honesty model. The tier rises to Met with Qualification and remains below Met because ePBS's Relay elimination removes the intermediaries of block production without remedying their economic concentration. Protocol-level enforcement changes the nature of the guarantee while leaving its structural cause untouched.

Alongside the carrying movement, the remaining delta entries quantify how far the target state shifts the foundation, and the sharpest of them concerns Finality under Minimal Load-Bearing Guarantees (I.4). The baseline finalizes via Casper-FFG in 12.8 minutes; the target state binds the definitive finality of Layer 1 via Three-Slot Finality to roughly 36 seconds, shortening it by a factor of approximately 21. A maturity-tiered stack of Based Preconfirmations, the Confirmation Rule, and Three-Slot Finality itself carries the target-state value.[384] Single Slot Finality remains nameable as the consensus end-state beyond 2029 without a fixed fork date, without shifting the near-term target-state value. The three structural improvements follow the same pattern, because each resolves a dependency that the baseline still carried. Functional Irreplaceability (I.1) rises to Met because the Native Rollup binding couples the L2 layer to Ethereum's execution at the protocol level. The Coordination Function (I.3) extends its reach from Settlement, Execution, and data availability to protocol-level verification. Long-Term Stability (III.1) decouples State growth, validation costs, and verification load from total capacity.[385]

The improvements are offset by the five constant criteria, which remain at Met with Qualification because the roadmap neither removes nor worsens their

[384] Current-state Finality 12.8 minutes via Casper-FFG: Chapter 4. Target-state finality stack comprising Based Preconfirmations (roughly 2 seconds, DEPL), Confirmation Rule (roughly 13 seconds, PLAN), and Three-Slot Finality (roughly 36 seconds, RES): 5.7-A. Single Slot Finality remains the consensus end-state beyond 2029 without a fixed fork date and does not determine the near-term target-state value.

[385] On the three structural elevations cf. 5.7-A and 5.7-B: Native Rollup binding and Settlement attraction via ERC-7683 (I.1), extension of coordination to execution and verification (I.3), three decouplings via zkEVM, Stateless Clients, and Tiered State (III.1).

underlying constraints. Security and Trust Load (I.2) still carries two residual risks that are not reachable by the protocol — the zkEVM proving assumption and the access layer of wallets, frontends, and DNS through which the majority of users interact. Open Generativity (II.2), Independent Verifiability (II.3), and Low-Threshold Inclusivity (II.4) hold their tier because the target state meaningfully strengthens them without resolving their respective core constraints, which range from the unresolved prover formalization to the 32-ETH threshold for solo staking. For Hardware Agnosticism (III.4), the constancy follows its own logic: cloud concentration of roughly 59 percent of hosted Execution Layer nodes across three providers is neither improved nor worsened by the target state. The principle established in 5.7-B — an unchanged indicator does not fall below its baseline tier — keeps Hardware Agnosticism at Met with Qualification.[386] The five constants secure the grade, as none of them falls below the level reached in the baseline.

What the Target-State Assessment Depends On

With the complete delta in view, the question of what the target-state assessment depends on comes to the fore. The target-state assessment rests on the assumption that the roadmap delivers, and this assumption that forms the largest attack surface of the evaluation. The integrated robustness assessment responds to this objection with a statement about the current maturity of each carrying component. The assessment of Neutrality and Censorship Resistance (II.1) rests on FOCIL and ePBS, both of which carry Plan status, making it the least secured of the three Critical Conditions. The assessment of Minimal Load-Bearing Guarantees (I.4) presupposes, in its finality layer, already-deployed Based Preconfirmations and the Confirmation Rule currently in client implementation; since L2 confirmation via Based Preconfirmations already operates in the seconds range while L1 confirmation hinges on the Confirmation Rule, it holds even in a conservative reading. Maturity sorting ranks the improvements without dating their arrival.

The largest concentrated dependency in the roadmap lies in the zkEVM, which on the active EF roadmap stands between unproven research and planning-level anchoring and simultaneously carries five criteria. It reduces the validator trust load (I.2), adds protocol-level verification as a fourth coordina-

[386] Cloud concentration in the baseline: roughly 59 percent of hosted Execution Layer nodes across three providers (AWS roughly 35.5 percent, Hetzner roughly 13.8 percent, OVHcloud roughly 9.7 percent), declining since 2022: Chapter 4. On the non-deterioration logic of Hardware Agnosticism in the target state cf. 5.7-B.

tion primitive alongside Settlement, Execution, and data availability (I.3), reinforces the execution determinism of the minimum guarantees (I.4), decouples validation costs from execution capacity (III.1), and reduces validator hardware demands (III.4). Three of the five criteria are improvements that would not reach Met without the zkEVM: the Coordination Function (I.3), whose extended reach additionally depends on ePBS and the Native Rollup binding; Minimal Guarantees (I.4); and Long-Term Stability (III.1). The other two — Security and Trust Load (I.2) and Hardware Agnosticism (III.4) — remain constants whose tier the zkEVM helps sustain at Met with Qualification. A delayed or absent zkEVM would therefore weaken the robustness of the three elevations without jeopardising any improvement for the two constants, as no elevation is at stake there. No other single feature bundles comparable cross-cutting impact, which is why the concentrated dependency is the most precise technical expression of the greatest implementation risk.[387]

The maturity logic extends beyond the zkEVM and sorts the seven improvements by the maturity of their weakest carrying component, with the range spread wide. Most robust is Adaptive Governance (III.2), whose foundation is a purely demonstrated process, carried by the strategic L1 pivot of February 2026 and a hard-fork cadence without chain splits. Adaptive Governance does not fall back even in the most pessimistic partial-delivery scenario. Weakest is Long-Term Stability (III.1), whose central element — the Tiered State System — rests on unproven research without its own EIP and makes it the taxonomically most questionable improvement in the delta. Between the poles lie the carrying components of intermediate maturity: ePBS with its planning-level anchoring but no guaranteed fork slot, and Native Account Abstraction, which builds incrementally on the already-deployed EIP-7702.[388] Closer to the weak pole lie Functional Irreplaceability (I.1) and the Coordination Function (I.3), whose elevation depends on Native Rollups (EIP-8079). These stand in draft-EIP research status, more weakly anchored than ePBS's planning tier but more firmly than the Tiered State System without its own EIP. The distribution of maturity and the distribution of assessment weight run against each other. Adaptive Governance, the most robust of the seven improvements, does not carry the overall assessment because, as a Qualitative Criterion, it affects only the grade, which

[387] On the cross-cutting impact of the zkEVM across I.2, I.3, I.4, III.1, and III.4 cf. the target-state assessments in 5.7-A and 5.7-B. Maturity: active EF roadmap between unproven research and planning-level anchoring.

[388] On the strategic L1 pivot (February 2026) as a demonstrated fact cf. 5.1 and 5.7-B; on the Tiered State System resting on unproven research without its own EIP cf. 5.4 and 5.7-A; on ePBS (planning-level anchoring, Glamsterdam) and Native Account Abstraction (building on the deployed EIP-7702) cf. 5.5 and 5.3.

was already secured in the baseline. What is assessment-carrying — alongside Neutrality and Censorship Resistance (II.1) — are precisely those movements at the weak pole: Functional Irreplaceability (I.1), the Coordination Function (I.3), and Long-Term Stability (III.1). Two of the three suffice, because the third cascade step caps below the reached category only when two weak outcomes are present. The target-state assessment thus rests on one indispensable movement and on two of three further movements tied to research lines without finalized specifications, and it is more dependent on unfinished work than the maturity sorting alone reveals. What the roadmap does not structurally resolve passes to the sections that follow.

6.2 Persistent Limits: The Crisis-Response Gap

What the roadmap does not structurally resolve accumulates at three points that survive its full delivery because no feature addresses their structural cause. At the top of block production a single Builder remains; in liquid staking a single protocol; in emergency response a documented plan without an enforcing mechanism. The target state responds to urgent concentration and crisis situations through social coordination rather than through protocol-enforced mechanisms. This gap is not an oversight of the evaluation but its anchored finding. No feature of the examined roadmap introduces such a mechanism, and the risk cuts across three criteria.

Block production carries the sharpest of the three concentrations, and it sits on the most critical path of the system. Titan builds 51.2 percent of blocks, BuilderNet 25.7, and Quasar 16.4, so three builder addresses control roughly 93 percent of production. The Herfindahl-Hirschman Index of this structure stands at 3,554 and clearly exceeds the US Department of Justice (DOJ) threshold for highly concentrated markets of 2,500. If Titan and BuilderNet coordinated their 76.9 percent, that would suffice for de-facto censorship of the majority of MEV-Boost blocks. The cause of the concentration is economic: more private order flow allows more profitable blocks and higher proposer bids, and the market position of the most efficient Builder reinforces itself. This is precisely where ePBS does not intervene (cf. 5.7-A on Relay elimination). Latency advantages and exclusive order flow persist as structural drivers. The countermeasure lies in BuilderNet, which is intended to break the concentration through a collaborative multi-operator architecture. Local building remains open to every proposer, though at 2.66 to 5.57 times lower yield. The protection is real, but it consists of market incentives rather than protocol guarantees, and the experience of

the OFAC peak of 79 percent shows that regulatory pressure can override economic incentives. For responsiveness, this means a defense that holds as long as the market sustains it, and no mechanism that would dissolve the concentration at the protocol level in a crisis.[389]

In liquid staking the pattern repeats with the opposite sign, because the concentration is declining here and it is precisely the decline that carries the finding. Lido holds roughly 23 percent of staked ETH, down from a peak of roughly 32 percent in 2023 and well below the blocking-minority threshold of 33 percent. The cumulative top-3 concentration stands at 40 to 45 percent and exceeds the 33-percent blocking-minority threshold only as a coordinated sum, while no single provider reaches it alone. No roadmap feature introduces a protocol cap, so the concentration remains structural even where it is falling. The most obvious countermeasure — a formal self-limit to 22 percent — was rejected by the Lido DAO in 2022 by vote, making the cap socially rejected rather than protocol-set. In its place came Distributed Validator Technology with 547,968 ETH and a quarterly growth rate of 57 percent. Distributed validation spreads operator activity and reduces the single point of failure in the validator set. The decline itself stems from the falling yield of roughly 5 to 2.6 percent between 2023 and the first quarter of 2026, competition from newer protocols, and public community pressure — a market- and socially-driven movement without a protocol mechanism. For responsiveness, the implication is that even the favorable trend of concentration rests on forces that the protocol neither generates nor could reliably invoke in a crisis.[390]

Informal governance sharpens the finding (cf. 5.8): its emergency response rests on documented plans rather than enforced mechanisms. The track record is strong. It extends from the social intervention of the DAO Fork and the two emergency hard forks of 2016 against the Geth DoS series that began in September, to four post-Merge upgrades without chain splits. The limit lies in the absence of a pause button: in the event of a zero-day consensus bug, no one can halt the network. The only answer is an emergency hard fork, which requires the coordination of all client teams and can take weeks. For a system that secures

[389] Builder market shares (Titan 51.2 percent, BuilderNet 25.7 percent, Quasar 16.4 percent), Herfindahl-Hirschman Index 3,554, DOJ threshold 2,500, and OFAC peak 79 percent: Chapter 4 (relayscan.io, 27 March 2026). On the yield disadvantage of local building (2.66 to 5.57 times), Relay elimination via ePBS, and BuilderNet as a collaborative countermeasure cf. 5.5.

[390] Lido share (roughly 23 percent, declining from a peak of roughly 32 percent in 2023), cumulative top-3 concentration (40 to 45 percent), and the 33-percent blocking-minority threshold: Chapter 4. On the Lido self-limit rejected in 2022, Distributed Validator Technology (547,968 ETH, plus 57 percent in the quarter), and yield development (from 5 to 2.6 percent between 2023 and the first quarter of 2026) cf. 5.5.

more than 100 billion US dollars, such a response time is a constraint. The Emergency Response Plan, signalled as a backstop, remains a documented action plan and not a protocol-enforced mechanism. The counterexample was provided by the Prysm bug of December 2025, which required no emergency fork because client diversity absorbed it — the best response was the one that was not needed. But this diversity is itself the result of active community work and not an automatic equilibrium, so even the successful absorption rests on social coordination. The system's strongest safeguard lies where a response becomes unnecessary.[391]

A Finding Across Three Criteria

The three factors each remain below the threshold at which a tier-carrying criterion changes tier, and emerge as a single finding only in their combined view. The shared finding cuts across Neutrality (II.1), Long-Term Stability (III.1), and Adaptive Governance (III.2) without tipping any one of the three.[392] Each factor carries a real countermeasure — BuilderNet and local building, Lido's decline alongside distributed validation, the fork track record alongside client diversity — yet each of these countermeasures is itself market- or socially-driven. The classification confirms the finding: the system, even where it responds successfully, responds socially and does not invoke an enforcing mechanism. Responsiveness therefore rests on social coordination, not on protocol-level enforcement.

6.3 Limits of the Evaluation Framework: What No Criterion Measures

The crisis-response gap of the preceding section is a limit that the evaluation framework still captures, because it cuts across three of its criteria and enters as their shared finding. A second limit concerns the reach of the instrument rather than the system, because four properties of Ethereum lie outside what any of the twelve criteria targets. The assessment of suitability as infrastructure has not rated them weakly — it has not addressed them at all, and it can say nothing about what no criterion measures.

---

[391] DAO Fork 2016, four post-Merge upgrades without chain splits, the emergency hard forks of October and November 2016 against the Geth DoS series, and the absorption of the Prysm bug of December 2025 through client diversity: Chapter 4.

[392] On the target-state assessments of the three affected criteria — Neutrality (II.1), Long-Term Stability (III.1), and Adaptive Governance (III.2) — across which the shared finding cuts without tipping any individual tier, cf. 5.7.

Four Properties Outside the Grid

Permissionless Composability appears in the evaluation framework, but only in a narrowed form. The second criterion of the second dimension (II.2) tests atomic composability as a binary indicator — whether Smart Contracts can interact with one another within a single transaction. The criterion supplements the indicator with permissionless deployment and open tooling. Outside the measurement lies composability as a combinatorial system property: the exponential and permissionless interplay of arbitrarily many contracts producing constructions that no one foresaw. In the so-called Money Legos of DeFi protocols this interplay takes concrete form. The criterion captures whether contracts can interact, not the generative power of their open combinability as a quality in its own right. The assessment makes no statement about this power, because the criterion reduces it to a yes-or-no question that presupposes its existence without measuring its extent.

Cryptographic system-state verification stands closer to the evaluation framework and yet remains beyond its reach. The third criterion of the second dimension (II.3) tests whether verification is possible at all. The criterion reads this from formal verification tools, a formal EVM specification, and the prevalence of audits — as the supply side of the trust load (I.2). Beyond the bare possibility lies the proof-based form of verification, in which the system mathematically proves its own aggregate state — via zkEVM and Stateless Clients — rather than securing it through institutional trust. The criterion measures whether a contract can be audited, not that the system carries a cryptographic proof of its own correctness. The assessment says nothing about the shift from institutional to mathematical trust, because the criterion captures the possibility of verification and not the ground on which it rests.

Decentralized self-development finds in the evaluation framework a criterion that reaches past it. The second criterion of the third dimension (III.2) tests whether governance is adaptive — whether the system can evolve and respond to new requirements. Outside adaptivity lies self-development without a center: evolution through rough consensus without any instance that could direct changes. Every change must be voluntarily adopted by a majority of node operators. Chapter 2 itself calls this model an extreme case of the Mayntz tension between steering and self-organization, because the system evolves without a central authority directing it. The criterion captures whether adaptive governance is present, not centerless self-development — so the assessment does not bring the absence of a directing authority into view.

Privacy falls outside the evaluation framework, and unlike the preceding three limits, no criterion narrows it, because none targets it at all. None of the three dimensions — neither Structural Foundation nor Qualitative Load-Bearing Capacity nor Resilience and Sovereignty — contains a question about confidentiality. With the first three limits, a criterion narrows the property to an indicator. Here the measurement point is entirely absent. The gap means there is simply no criterion. Outside the evaluation framework lies confidentiality as a protocol property, built into the architecture through Encrypted Mempool and pre-inclusion privacy. Privacy is the only one of the four limits that has a correlate in the examined roadmap, anchored in the Lean Ethereum program (5.1-B and 5.1-C). The assessment nonetheless says nothing about a property that is genuinely being built into the target state, because the instrument provides no measurement point for it.

The four properties are not a scattered remainder that the criteria happened to miss; they share a common position, each lying outside what the evaluation targets at all. Each of them lies where the system carries a property and the instrument asks no question. That there are four and not a single isolated case reveals a pattern whose interpretation points beyond the reach of this assessment and is taken up by Chapter 7.

6.4 Synthesis and Answer to the Research Question

Ethereum today, as a Settlement and data availability layer, secures capital on the order of 100 billion US dollars, representing more than 60 percent of all on-chain-secured DeFi capital. On the same infrastructure, Stablecoins worth approximately 166 billion US dollars are issued — more than half of the global Stablecoin market — and roughly 65 percent of all new Smart Contracts are deployed on the Layer-2 systems whose Settlement and data availability the base layer carries. The system binds more than 31,000 active developers, representing more than 70 percent of all blockchain developers worldwide. This is the concrete substance at stake in the evaluation, before the question of its suitability can be answered. Does this architecture meet the requirements that must be placed on fundamental digital infrastructure? The answer is yes, in both examined states — but in both only conditionally.

Almost all requirements the system meets already today, each with qualifications, and the one that pulls it most sharply downward is censorship resistance at the level of block production. Currently, roughly 15 percent of blocks pass through Relays that exclude transactions from sanctioned addresses, while Titan as the largest Builder openly refrains from doing so. In November 2022 the share stood at 79 percent, meaning that a government sanctions list effectively controlled which transactions a permissionless system processed. This market-driven decline occurred without any intervention in the protocol — solely because non-censoring Builders operate more profitably and the market has shifted toward them. That is the substantial condition in a single image: censorship resistance functions, but it rests on market structure, and a market structure can reverse.

It is this condition that the roadmap addresses, and two proposals would resolve it at the protocol level. FOCIL, a validator-committee-based inclusion-list system, would enforce the inclusion of every named transaction at the protocol level. ePBS, the protocol-native anchoring of Proposer-Builder Separation, would remove the Relays as off-chain intermediaries through which OFAC filtering currently takes place at all. Both carry Plan status — designed but not delivered. Therein lies the actual change, because the nature of the condition itself shifts: today, suitability depends on how the Builder market behaves. After full implementation, it would depend on whether the planned mechanisms deliver what their design prescribes. The roadmap thereby exchanges a market dependency for an implementation dependency, and the condition does not become smaller — it becomes of a different kind.

The Conditionality Cannot Be Discharged

Even full implementation of the roadmap does not dissolve the conditionality, and a single missing mechanism makes this tangible: Ethereum has no pause button. As shown in 6.2, in the event of a critical bug no one can halt the network. The only answer remains a costly emergency hard fork. The missing pause button is not a feature that a later roadmap adds — it is the flipside of centerlessness itself. Anyone wishing to dissolve the conditionality would have to introduce an authority capable of intervening, halting, or bearing liability in a crisis — and would thereby surrender precisely the property that the work has been evaluating Ethereum against. As long as the system remains without a center, its response to acute crises remains bound to social coordination. The conditionality is therefore not an incompleteness that a forthcoming roadmap corrects. It is a consequence of the architectural decision against a center.

A Qualification Regarding the Instrument

A qualification belongs to the assessment, and it concerns the instrument rather than the subject. The evaluation grid is constructed on a rule-based rather than an additive basis, so minor constraints below its thresholds do not propagate into the overall assessment. Two systems of very different substance can therefore carry the same assessment, and the baseline and target state demonstrate this against each other: far apart in substance, only one step apart in formal assessment. Anyone reading only the assessment misses how much substance lies between the two states. This is not a weakness of the assessment — it is a property of the instrument that the reader must carry alongside it.

The assessment therefore stands: Ethereum is suitable as fundamental digital infrastructure — today under substantial conditions, after full implementation of the roadmap with conditions, in neither state unconditionally. What conditional suitability means beyond the evaluation result — for the actors who build on the infrastructure, place their trust in it, or regulate it — is taken up by Chapter 7. The final word rests with the assessment, not with the qualification.

Chapter 7 Implications

After the Assessment

The research question has been answered: Ethereum is suitable as fundamental digital infrastructure — both as of the reference date of March 2026, which I have designated the current state, and in the state under full implementation of the roadmap (the target state). Under a closer criterion-by-criterion comparison of the two states, the one under the roadmap fulfills the stated conditions more completely, as expected, but remains bound to conditions in equal measure. That is the result of the work, and it is no longer the subject of this chapter.

No assessment is made here, and accordingly no judgment is rendered. From here, the register returns to that of the preface, because the answer provided in Chapter 6 now allows two questions to be answered that have remained open since the beginning of this work. The first concerns the actors who interact with the network: those who are actively building on the infrastructure (institutions, and with them the developers not separately mentioned in 1.1), those who put their money behind it (investors and users), and those who regulate it (politics and law). Section 1.1 noted that they conduct their due diligence carefully, but that in doing so they act under an assumption that eluded examination because there was no standard to test it against. The question that follows from this is, in substance: on what does the assumption of permanent existence rest — the assumption under which those who build on the system, trust it, or regulate it already act — when no authority stands behind it? The second is the naively posed question that originally inspired me to conceive this work: "What does such a cryptographic system make possible beyond decentralized payment, and does it even need to go beyond it?" The resolution of both questions is the subject of Chapter 7.

The Picture When Everything Is Built

To answer the questions and maintain a clear thread, I first sketch the features of the target state as a whole. From them I can derive a simplified definition that makes the questions easier to answer.

A block that carries a multiple of today's volume costs the verifier exactly as much as an empty one. This sentence is the pivot of the picture that emerges when the roadmap is fully built. Validators no longer recompute transactions individually. They verify a proof of correct block execution generated by a prover — a specialized machine that performs the execution once on behalf of all. This verification takes only a few milliseconds, regardless of how many transactions the block contains. Network capacity scales without increasing verification costs. Today Ethereum processes between 15 and 30 transactions per second at 60 million Gas, depending on transaction complexity. The target for the base layer is roughly 10,000, reachable over the long term via a staged expansion path for execution. At the same time, participants keep their sovereignty: anyone who does not trust the cryptographic proof can continue to execute and verify transactions in full themselves, as before.

This verification runs directly on ordinary end devices such as laptops or smartphones. The storage requirement for a verifying node falls from 2 to 4 terabytes to under 100 megabytes, while the required working memory is reduced to under 1 gigabyte and the compute load to a single processor core. The initial synchronisation, which currently still takes hours or days, shortens to a few seconds. One example is the client Helios, which runs in a browser tab and on Raspberry Pi-class hardware without holding its own copy of the blockchain.

This stands in sharp contrast to current practice, in which roughly 70 percent of users access the network through centralized infrastructure providers such as Infura or Alchemy and must rely on their data without independent verification. In the target state, what is now a necessity becomes a choice: anyone who wants to verify can do so on their own device, and the data required for this is drawn from a decentralized peer-to-peer network. No single provider stands between them any longer.

The account itself also changes in nature: where the protocol previously prescribed how an account authorises a transaction, it will in the future delegate that task directly to the account itself. Its validation rules become freely programmable.

For users this means that they can unlock their account with a fingerprint, making the end device itself the wallet. If the device is lost, several trusted par-

ties they have chosen themselves can restore access together, without any one of them being able to access the balance alone. Network fees can moreover be paid in any currency such as USDT, or the application covers them directly in the background. The signature scheme can also be switched to a quantum-resistant logic.

Under the same precondition — the fully implemented roadmap — the confirmation time of a payment is no longer measured in minutes. Already after roughly 2 seconds a secured commitment from a validator is available; after approximately 13 seconds the deterministic confirmation on Layer 1 follows; and after roughly 36 seconds irreversible finality is reached. By comparison, this currently takes 12.8 minutes. Each of the waiting times represents a clearly defined security guarantee for the user, with even the strongest guarantee secured in seconds.

The decision over which transactions are included in a block no longer lies with a single actor. Per slot, 16 randomly selected validators jointly compile with the proposer a binding list of transactions to be included (Inclusion List), with a single honest validator among the 16 being sufficient to ensure that a bypassed transaction still makes it onto the list. A proposer who ignores the list loses the votes of the attesters and thus the entire block.

Since Builders will in the future submit their bids directly to the blockchain, Relays also become redundant — through which today roughly 15 percent of blocks are routed and which apply censorship filters such as OFAC lists in the process. In addition, transaction content remains encrypted until final sequencing, so that no one can selectively exclude transactions whose content they do not know. Censorship resistance is thereby anchored in the protocol. Previously it depended on the good conduct of operators. It holds as long as the validator network remains broadly distributed.

The security guarantees of the base layer do not end at its system boundaries. If a Rollup delegates the sequencing of its blocks to the proposers of Layer 1 and has its state transitions validated by their consensus, it requires no dedicated sequencer, no separate proof system, no security council, and no governance of its own. What remains is a pure execution environment with its own fee market and State.

Transactions between two such Rollups no longer require bridges with additional trust assumptions. This applies to Rollups that remain EVM-equivalent and fully cede their block ordering to the base layer. Since implementation ultimately rests with the operators, a wide range of integration levels remains — from

full convergence to the sovereign chain that shares almost no interfaces with the base layer.

What It Is at Its Core

That is the picture of the target state, assembled. From here it can be simplified into a general definition of the subject.

Ethereum is at its core a state — a ledger of who owns what and which contractual rules for changing it are binding. Four properties define it.

It is shared: all participants maintain the same state and not each their own version with their own truth. It is verifiable: every entry can be checked against the rules of the protocol without a third party needing to attest to its correctness. It is binding: an entry that once holds can no longer be annulled by any subsequent authority. And it is centerless: no entity exists that could maintain, alter, or shut it down unilaterally. This description holds for the current state as much as for the state after full implementation of the technical roadmap. The four properties are present today, but they are not fully realized everywhere. Full implementation of the roadmap makes them scalable, accessible, and protocol-enforced. Today their realization rests on market conditions.

The evaluation framework itself did not capture four properties, because no criterion targeted them. The twelve criteria derive from infrastructures whose operation is owned by someone and whose financing comes from outside. Where Ethereum departs from this, the standard has no measurement point.

A ledger in itself is nothing unusual, as institutions like banks have maintained registers for centuries. That this state reaches far beyond a mere account book rests on three fundamental principles.

The ownership of an entry rests on a cryptographic proof. Anyone who can furnish this proof has absolute control over the entry at the protocol level. No bank, no exchange, and no authority can grant or revoke that disposal. This is a different kind of ownership from that which is conferred. It rests on the proof.

Furthermore, the rules written into the state enforce themselves. As soon as the defined precondition occurs, the associated consequence executes within the system automatically, without any ordering, supervising, or enforcing authority being required. What is enforced is the pure execution of the rule. The protocol executes a flawed rule with the same reliability as a flawless one.

Finally, every entry can be freely combined with every other, without the parties needing to know one another or ask permission. In their strongest form, two operations on the same layer happen in a single, indivisible step, so that they

either succeed together or fail together. Across the boundary between different layers this no longer holds, but the principle of free combinability persists.

A pure payment system uses only the simplest of the three — the mere transfer of balances. Before the picture, two questions remain open: what the assumption of permanent existence rests on — the assumption under which those who build on this state, trust it, or regulate it already act — and what it makes possible beyond decentralized payment.

Who Builds, Who Trusts, Who Regulates

Institutional actors have long been acting under this assumption. J.P. Morgan's tokenized money market fund operates exclusively on Ethereum, and Crédit Agricole issued a MiCA-compliant euro Stablecoin on the same chain through their custody subsidiary — both as documented in 1.1. Both institutions subjected the system to a detailed practical due diligence beforehand. The more far-reaching assumption, however, eluded examination: that the blockchain will still run reliably in ten years, treat all participants equally, cannot be shut down unilaterally, and will enable future, as yet unknown, forms of use. What the actors presuppose in practice are permanence, neutrality, the impossibility of unilateral shutdown, and openness. These are the four properties of the state, named from the perspective of practice. The first group acting under these premises builds directly on the system.

Anyone developing on Ethereum benefits from the architecture from the very first moment. A developer creating a new DeFi protocol can integrate the liquidity pools of Uniswap, the lending markets of Aave, and the Stablecoin logic of MakerDAO into their own application — without requiring a contractual relationship with these protocols, without their approval, and without their operators even needing to learn of the integration. In this structure the classical counterparty risk disappears. Permanence shows in practice in that Smart Contracts from 2017 continue to execute unchanged to this day. The only systematic exceptions are protocol upgrades adjusting transaction fees, which in rare cases affected code that assumed fixed fee values. An institution anchoring an application here receives neutrality and interface stability without contracts, negotiations, or renewals.

These guarantees apply primarily to the base layer, however, on which the fewest actors build directly. In practice there is a hurdle. users accessing a DeFi service on a Rollup via MetaMask move across four layers of trust that the Ethereum protocol does not directly secure. Security conditions vary considerably between

them. While the Base network has an exit window of zero days, Arbitrum One's stands at 17 days. The use of platform-specific functions such as Stylus or specific precompiles also makes switching to alternative systems more difficult. The original expectations of the blockchain shift to an upstream layer with its own operators, deadlines, and switching costs. Even after full implementation of the technical roadmap, the choice of layer remains decisive, as its implementation rests with the operators and ranges from fully adapted to entirely independent networks. Chapter 4 describes this constraint as a discrepancy between the theoretical possibilities of the protocol and actual operational practice. For developers this describes the real decision: the chosen layer determines the actual security.

Investors, by contrast, make fewer technical decisions but place great trust in the system. Institutional actors trust in the network's continued existence over the next ten years and point to uninterrupted availability despite fifteen major upgrades in the past. This study puts a time frame on that trust. Over a five-to-ten-year horizon Ethereum looks stable: the upgrade process runs reliably, economic security covers roughly 26 billion US dollars, and the validator infrastructure is intact. Over a twenty-year timeframe, unresolved issues — uncontrolled chain data growth, the unsettled long-term reward structure for security, and the outstanding transition to quantum-resistant cryptography — jeopardize sustainability. The three risks are not equally far along. For data growth, theoretical solutions exist that, according to Chapter 6, rest on unproven research without a concrete implementation plan. For the long-term financing of network security there is as yet no agreed concept, while the transition to quantum resistance is already being tested in testnets. The reference to past stability thus supports only the medium-term outlook. The expectation of decade-long permanence, as Section 1.1 describes it, goes beyond the currently demonstrable foundations. One feature without parallel in the seven reference systems examined is that the value of the native currency directly determines the security of the entire system. Investors hold no claim against a company. With their capital they directly underpin the integrity of the system.

End users rely on the same assumptions without knowing it, even though Section 1.1 does not name them explicitly. They benefit from direct, inviolable control over their digital assets through their own keys, without dependence on intermediaries such as banks or authorities. Resilience showed vividly in May 2023, when after two successive disruptions network participation fell below the critical two-thirds threshold. The protocol responded autonomously, incrementally reducing the stakes of inactive validators and restoring full functionality after 96 minutes without human intervention or external coordination. Classical sys-

tems such as the Internet address directory, the SWIFT network, or power grids require manual intervention, contingency plans, or physical interventions in the event of failures. Ethereum's protocol requires no such external assistance.

Despite the technical strengths, there is a gap between what is possible and what is used. While any user can in principle run their own node and independently verify the chain, in practice hardware, maintenance, and technical knowledge usually stop them. Many users voluntarily outsource custody to centralized exchanges. This is particularly pronounced when it comes to securing the network. The minimum requirement of 32 ETH corresponded, at the reference-date rate, to roughly 67,000 US dollars, and the share of privately operated, independent validators stands below one percent of total volume. Service providers lower the barrier to amounts as small as 0.01 ETH, but take control of validation in return. Economic access is thus open, while effective control over the infrastructure remains behind technical and financial barriers. According to Chapter 6, the constraints — including the 32-ETH threshold and the upstream user interfaces — survive full implementation of the roadmap, because no feature addresses their structural cause.

Among government regulators, a different picture emerges, as the point of access is at the application layer. The US GENIUS Act, signed on 18 July 2025, regulates Stablecoin issuers with regard to reserve requirements, transparency obligations, and licensing, but excludes the underlying protocol layer. It regulates Circle and Tether as issuers while affecting neither the Ethereum protocol itself nor the underlying Smart Contracts governing USDC and USDT. For decentralized applications without identifiable issuers, no comparable framework yet exists. This decision limits legal access to the application layer.

While private actors delegate responsibility to the protocol, regulators seek concrete points of contact. Since the protocol itself offers no addressable party, the focus shifts to central interfaces. Three commercial actors control roughly 93 percent of block production, led by Titan with 51.2 percent. The Herfindahl-Hirschman Index stands at 3,554 and clearly exceeds the US threshold for highly concentrated markets of 2,500. Node operation is also concentrated among a few cloud providers: roughly 59 percent of hosted Execution Layer nodes run on three of them. Geographically, 39 percent are in the US, 14.5 percent in Germany, and 14 percent in China, so roughly 53 percent fall in just two countries. State interventions therefore target the physical concentration points.

The weight of the concentration showed in November 2022, when 79 percent of blocks filtered transactions from sanctioned addresses. The restriction, however, caused only minimal delays of a few blocks. At present the share stands

at roughly 15 percent, as the largest actor, Titan, deliberately refrains from filtering. This decline occurred without protocol changes, driven purely by market incentives. The Van Loon v. Treasury ruling of November 2024 and the OFAC delisting of Tornado Cash in March 2025 fell in the same period. Such regulatory access remains unstable, however, as it rests on changeable market structures. According to Chapter 6, high concentration at the top of block production will persist even after full implementation of the technical roadmap.

The system operates under the law but diverges from its classical mechanisms of access. The interface between programmatic and state enforcement remains largely unsettled. Chapter 3 illustrates this in the comparison of infrastructures: none of the examined platforms possesses physical enforcement power. The power grid can cut supply, the Internet can block data, and SWIFT can exclude participants — but none of these systems can exercise coercion against persons. The legal system, by contrast, retains its full power over persons, companies, and locations as the overarching authority. On the decentralized blockchain, however, legal access has no direct address to reach for.

The absence of a central pause button illustrates the independence. In the event of critical errors the system cannot be stopped. The only answer remains an emergency hard fork requiring the coordination of all client teams. The Prysm bug of December 2025 showed that it can be avoided. The diversity of client software absorbed the effects. The stability rests, however, on ongoing social coordination and does not represent an automatic equilibrium. There is no pause button, because it contradicts centerlessness. Anyone introducing such a control authority would destroy the very foundation of the system. The permanent existence of the infrastructure thus depends on no single actor. It rests on the interplay of the centerless structure and the ongoing social coordination that sustains it. With this the first core question is answered, while the second examines the possibilities for action beyond pure value transfer.

It Was Never Only About Payment

I noted in the preface that enormous sums are in motion, that grand narratives about payment, identity, and the management of one's own assets without an intermediary are circulating, and that it nonetheless remains strangely unclear what the thing is at its core. What I was looking for was the core. Payment was the word I used to outline it, and what lies beyond it remained open. The state now stands named, together with its three principles. The answer begins with the word the question carries with it without ever clarifying it.

For me, payment meant payment transactions and autonomy over one's own keys. Both rest on the first of the three principles. Anyone who can furnish the cryptographic proof has absolute control over the entry at the protocol level, without a bank, an exchange, or an authority granting it. The mere transfer of balances is the simplest action derivable from such disposal. It changes a number in two places and is done. It demands little of the state. The state must recognize the proof and bindingly record the change — a payment asks no more of it than that.

What I called payment is one of three principles, and it uses that principle in its simplest form. With this it is answered what the thing is at its core, and the answer is: not payment. The core is the state. Payment is one way of using it. Anyone who describes Ethereum as a payment system describes a use it permits, and leaves two of the three principles unmentioned. The question of what lies beyond payment is therefore not a question about add-on functions. It is the question of what problem the two remaining principles solve that would otherwise remain unsolved.

What remains are self-enforcement and free combinability. Both address a problem older than any technology — one that Chapter 3 attributed to the legal system. How do strangers arrive at binding agreements? Economic history shows that before the development of institutional law, transactions functioned only within clan, family, or ethnic networks, because reputation-based enforcement does not scale beyond small groups. Chapter 3 elaborated this through Avner Greif's study. Only state-enforced contract law opened cooperation to millions of strangers. Anyone wishing to bind themselves previously had to know one another, and the reach of an agreement ended at the reach of the group that could punish a breach. The legal system, which appeared as enforcement power in the preceding section, stands here on the other side: as the thing that makes agreements between strangers possible in the first place. It differs from all other infrastructures in that behind its enforcement stands the state's monopoly on the use

of force. Binding commitment among strangers previously existed only under an authority capable of intervening when necessary.

Self-enforcement generates binding commitment without such an authority. As soon as the defined precondition occurs, the associated consequence executes within the system automatically, without anyone ordering, supervising, or enforcing it. The example lies in the protocol itself. The Slashing system penalises validator misconduct without any authority being involved. The enforcement is algorithmic and deterministic. When the defined conditions arise, the prescribed consequence follows — whether or not the affected party consents.

Free combinability continues where self-enforcement begins. Every entry can be freely combined with every other, without the parties needing to know one another or ask permission. Chapter 6 captures the combinatorial property in its sharpest form as the exponential and permissionless interplay of arbitrarily many contracts producing constructions that no one foresaw. Both principles together yield what lies beyond the transfer of balances: binding cooperation between strangers without an authority to enforce it. What contract law gained in reach, it gained at the price of an authority with access. The state holds the reach and omits the price.

Two limits accompany this achievement. The first is that only the execution is taken over. The correctness of the rule is not decided by the protocol — it executes a flawed one with the same reliability as a flawless one. The second limit lies in the expressive capacity of code. Disputes over intentions, interpretations, negligent conduct, and non-contractual claims remain reliant on institutional enforcement. The legal system is not replaced. What arises is a complementary enforcement layer for cases where institutional enforcement is too slow or too costly, does not reach across jurisdictions, or was never designed for the use case at all.

The second half of the question remains. Does the system need to go beyond payment? It is the question of an investor wanting to know whether they carry complexity that no one needs.

For payment the answer is no — the extra is ballast. Anyone who only transfers needs no self-enforcement, because a transfer enforces nothing beyond itself. They need no combinability, because a transfer connects with nothing. Both principles would lie in the state and remain unused. A payment system that carries them along carries them for nothing, and the investor's suspicion is on point here.

For infrastructure the answer flips. Chapter 2 established, following Brett Frischmann, that infrastructure defines itself by what it enables, not by what

it is. It serves as a factor of production for downstream activities and is not an end product. A system that only pays is an end product. Self-enforcement and combinability are the condition for the concept of infrastructure to apply at all. What a payment system would leave unused is precisely what a builder reaches for. Without self-enforcement there would be nothing to anchor; without combinability nothing to connect to.

My naively posed question was thereby exactly the distinction that Chapter 2 later derived as conceptual work: end product or input. The preface states that it is at this concept that the broad opening question sharpens into a testable one. The system does not decide the distinction — it permits both. Anyone asking the question decides it in the asking, because in asking they indicate what use they intend to make of the system.

Both questions are answered. One hung on an assumption that eluded examination for lack of a standard; the other on a word I used without knowing what it referred to. What shows in them applies not to Ethereum alone. I chose the network as a representative because the question had to be tested against something concrete, and the answer applies to the class of such systems. It does not apply to the question of which system ultimately prevails. What the writing has done to my own perspective belongs in the final lines of this work.

More Than Software: A Few Last Words

I am writing this final section of the final chapter after almost seven months of work, research, reading, learning, writing, iteration, conversations, and a great deal of reflection. To be honest, I do not know how one truly concludes a work like this. These lines will probably turn out to be one further round of the countless iterations, after I have asked an AI to situate what is written structurally and in content.

This self-initiated work has taught me much. The very fact of approaching such a topic and finding a starting point was a task in itself. Finding a logical thesis and an entry point from a banal question and a certain uncertainty — one that during a brutal market phase nearly tipped into frustration — felt like that meme: a small knight with a drawn sword facing an impossibly large one.

However impenetrable such a start without a framework handed to you from outside may be, it has enabled me to become independent. The sheer act of reading and consuming the learning section of ethereum.org within a few days — which ranged from absolute beginner articles to first more complex concepts and architectures — already indicated that the work might be longer than expected. The partly newly acquired information that I could not situate and that did not answer my motivating question led me, precisely for that reason, to set a genuine framework and redefine the question. In this way I was able to discover how to approach a work that is with no academic supervision and was not to be limited to the usual eight to fifteen pages.

With a clearly defined question testable against a fundamental framework, a horizon and a structure suddenly opened up in which one could work with certainty. From that point the research extended across theories of infrastructure as well as physical, philosophical, political, and geographical topics and much else. Alongside the research and the various perspectives, a chapter structure that evaluates a state provided a clear thread to follow, one that made the research appear even more ordered with a certain equanimity. For the subject develops continuously, with pivots, iterations, and refinements occurring almost weekly. Without the division into the reference-date state (in the work: ”current state”) and the emerging state (”target state”), one could hardly write about it in peace without everything being overturned the next day. The approach taught me bitter lessons and at the same time introduced me to instruments I came to value useful not only for researching and writing, but above all for one’s own clarity of thought and for connecting relationships.

And so to the continuously evolving subject itself. When examining such a dynamic system, it is easy to assume that its evolution always follows a comprehensively structured, wide-ranging plan and rests on the clear consensus of all involved. During the writing of this work, over a period of just under seven months, I saw at first hand, however, that even for the most informed actors this is by no means the reality.

A striking example that illustrates the discrepancy between rational planning and chaotic reality is the recent debate around the Ethereum Foundation. It centered on how the network should position itself in the future. While the Foundation wanted to hold to the ideological, almost romantic, cypherpunk ideals, some users, part of the community, and members of the team called for pragmatic, market-oriented adjustments in the face of stiff competition from other chains. The debate revealed that even in one of the networks regarded as the most technologically advanced, no automatic consensus exists on the path of further development. Progress is instead the result of painful friction, internal disputes over direction, and the constant struggle between idealism and pragmatism.

Alongside the debate itself, announced features and architectures changed, shifted, or were even dropped in the course of the months. The most far-reaching — also for the focus of this work — was the reorientation toward direct L1 scaling initiated in February 2026. With the pivot, the research environment, set in motion by Vitalik Buterin, corrected the strategy that had dominated for years of offloading scaling primarily to external Layer-2s. Instead, increasing the capacity of Layer 1 itself moved back into focus. Such an abrupt paradigm shift confirms the chosen division of this work. It shows that the "current state" can be precisely captured while the "target state" remains a dynamic construct that eludes rigidly predetermined roadmaps and realigns itself repeatedly in the consensus process.

The realignments are attributable in part to the uncertainties of the developers, researchers, and other involved actors, who in part operate independently of the Ethereum Foundation's focus and remit. They led to restructurings within the Foundation and beyond. At the same time they make clear that no fully defined picture of the future can exist when even the researchers themselves iterate monthly to annually. This does no harm to the work and above all none to the network. I consider it entirely normal and appropriate, because Ethereum in particular must change with the times in the face of continuously growing technological possibilities and threats such as artificial intelligence, quantum computers, or advances in pure mathematics — and planning and research have to be adjusted accordingly.

To get a picture of the research and the otherwise predominantly digital exchange, the work has taken me to various gatherings and events — specifically to Berlin Blockchain Week. There I was able for the first time to meet people who work on the project and in its environment. Not that the trip changed my assessment; rather the opposite: the assessment — or rather the work and the understanding gained through it — made me curious to experience the otherwise purely digital communicating, learning, and getting to know one another also at a physical level. I am grateful that this piece of writing opened up such experiences and insights and paid for the ticket to the conferences.

So much for the personal documentation and what I got to learn and experience over the course of these seven months. What remains is the reference back to the preface and the account promised there of what I think.

When one analyses the technical infrastructure of Ethereum and then, as in Berlin, suddenly sees it alive before one, the impulse arises to radically widen one's view. For this is about more than software.

At the moment Ethereum is still understood by the wider public mostly in pragmatic terms, as a technological alternative to our existing, centralized financial systems. I interpret the development, however, more fundamentally: in my view the blueprint of our reality brings forth, at every pole, a counter-pole — just as there is a below for every above, or a reaction for every action.

From this perspective Ethereum is no accidental invention. The structural template for a decentralized network lay ready in logical space, and Ethereum has occupied it. It is this polarity that explains why the common comparisons of the present underestimate the phenomenon.

The gap in pure payment is enormous, because the overwhelming share is settled via centralized fiat money and only a vanishingly small fraction via the entire decentralized sector. Since it is, as already established in this work, not about pure payment processing, however, a purely percentage-based comparison falls short. It merely sets the in a decentralized way settled payment volume against the centralized one and leaves the actual dimension beside it untouched.

Even on the narrow field of pure balance transfer, Ethereum has occupied the counter-pole, even though that is the discipline in which it departs least from a classical payment system. The actual, unoccupied field lies, however, where it was never about payment alone: in everything that people agree to bindingly with one another without knowing each other and without a higher authority to step in in a dispute. This space cannot yet be expressed in percentages today. The reason is that it has for the most part not yet been entered. What can be said is the smaller and at the same time decisive thing: the infrastructure that opens it exists.

Whether and how quickly it fills is a question for regulation, habit, inventiveness, ease of use, and the circumstances under which people begin to trust an order without a center more than one with one.

Figures

Tables

Quellen

Die Arbeit belegt überwiegend aus Primärquellen: Spezifikationen, Vorschlagstexten und öffentlich abrufbaren Daten. Das Verzeichnis führt sie in drei Abteilungen. Alle Angaben stehen wortgleich in den Fußnoten des Haupttextes.

Literatur

Aschauer, David Alan (1989): Is Public Expenditure Productive? Journal of Monetary Economics, 23(2), pp. 177–200.

Buterin, Vitalik (2014): A Next-Generation Smart Contract and Decentralized Application Platform. Ethereum Whitepaper.

Buterin, Vitalik (2020): A rollup-centric ethereum roadmap. Blogpost, October 2020.

Buterin, Vitalik (2025): Stages as a framework for evaluating rollup maturity. Blogpost.

Buterin, Vitalik / Conner, Eric / Dudley, Rick / Slipper, Matthew / Norden, Ian / Bakhta, Abdelhamid (2019): EIP-1559: Fee market change for ETH 1.0 chain. Ethereum Improvement Proposal. Activated in the London upgrade, August 2021. Cf. Ethereum Foundation Blog for the technical documentation.

Buterin, Vitalik / Griffith, Virgil (2019): Casper the Friendly Finality Gadget. arXiv:1710.09437. For the combined Gasper specification cf. Buterin, Vitalik / Hernandez, Diego / Kamphefner, Thor / Pham, Khiem / Qiao, Zhi / Ryan, Danny / Sin, Juhyeok / Wang, Ying / Zhang, Yan X (2020): Combining GHOST and Casper. arXiv:2003.03052.

Buterin, Vitalik / Hernandez, Diego / Kamphefner, Thor / Pham, Khiem / Qiao, Zhi / Ryan, Danny / Sin, Juhyeok / Wang, Ying / Zhang, Yan X (2020): Combining GHOST and Casper. arXiv:2003.03052.

Daian, Philip / Goldfeder, Steven / Kell, Tyler / Li, Yunqi / Zhao, Xueyuan / Bentov, Iddo / Breidenbach, Lorenz / Juels, Ari (2020): Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In: IEEE Symposium on Security and Privacy (S&P), 2020. DOI: 10.1109/SP40000.2020.00040.

Davidson, Sinclair / De Filippi, Primavera / Potts, Jason (2018): Blockchains and the Economic Institutions of Capitalism. Journal of Institutional Economics, 14(4), pp. 639–658. DOI: 10.1017/S1744137417000200.

Dylan-Ennis, Paul / Kavanagh, Donncha / Araujo, Luis (2023): The Dynamic Imaginaries of the Ethereum Project. Economy and Society, 52(1), pp. 87–109. DOI: 10.1080/03085147.2022.2131280.

Frischmann, Brett M. (2012): Infrastructure: The Social Value of Shared Resources. Oxford University Press.

Gencer, Adem Efe et al. (2018): Decentralization in Bitcoin and Ethereum Networks. In: Financial Cryptography and Data Security. Springer, pp. 439–457.

Greif, Avner (1994): Cultural Beliefs and the Organization of Society. Journal of Political Economy, 102(5), pp. 912–950.

Grimmelmann, James / Windawi, A. Jason (2023): Blockchains as Infrastructure and Semicommons. William & Mary Law Review, 64(4), pp. 1097–1129.

Hanseth, Ole / Lyytinen, Kalle (2010): Design Theory for Dynamic Complexity in Information Infrastructures: The Case of Building Internet. Journal of Information Technology, 25(1), pp. 1–19.

Hodgson, Geoffrey M. (2016): Conceptualizing Capitalism: Institutions, Evolution, Future. University of Chicago Press.

Jabareen, Yosef (2009): Building a Conceptual Framework: Philosophy, Definitions, and Procedure. International Journal of Qualitative Methods, 8(4), pp. 49–62.

Jensen, Johannes R. / von Wachter, Victor / Ross, Omri (2021): How Decentralized is the Governance of Blockchain-based Finance: Empirical Evidence from four Governance Token Distributions. In: Proceedings of the 29th European Conference on Information Systems (ECIS 2021).

Mayntz, Renate (1993): Große technische Systeme und ihre gesellschaftstheoretische Bedeutung. Kölner Zeitschrift für Soziologie und Sozialpsychologie, 45(1), pp. 97–108.

Mnookin, Robert H. / Kornhauser, Lewis (1979): Bargaining in the Shadow of the Law: The Case of Divorce. Yale Law Journal, 88(5), pp. 950–997.

Nabben, Kelsie (2023): Web3 as 'Self-Infrastructuring': The Challenge Is How. Big Data & Society, 10(1). DOI: 10.1177/20539517231159002.

Ostrom, Elinor (1990): Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge University Press.

Ovezik, Christina / Karakostas, Dimitris / Milad, Mary / Kiayias, Aggelos / Woods, Daniel W. (2025): SoK: Measuring Blockchain Decentralization; Srinivasan, Balaji S. / Lee, Leland (2017): Quantifying Decentralization; Wachter, Victor / Ross, Omri (2021): How Decentralized is the Governance of Blockchain-based Finance. Full references are in the footnotes to Section 2.1.

Pincheira, Miguel et al. (2023): An Infrastructure Cost and Benefits Evaluation Framework for Blockchain-Based Applications. Systems, 11(4), 184. DOI: 10.3390/systems11040184.

Rozas, David / Tenorio-Fornés, Antonio / Díaz-Molina, Silvia / Hassan, Samer (2021): When Ostrom Meets Blockchain: Exploring the Potentials of Blockchain for Commons Governance. SAGE Open, 11(1), pp. 1–14. DOI: 10.1177/21582440211002526.

Saltzer, Jerome H. / Reed, David P. / Clark, David D. (1984): End-to-End Arguments in System Design. ACM Transactions on Computer Systems, 2(4), pp. 277–288.

Van Schewick, Barbara (2010): Internet Architecture and Innovation. MIT Press. Cf. the reference in Chapter 3, Section 3.1.2.

Spencer-Hicken, Scott et al. (2023): Blockchain Feasibility Assessment: A Quantitative Approach. Frontiers in Blockchain, 6. DOI: 10.3389/fbloc.2023.1155125.

Srinivasan, Balaji S. / Lee, Leland (2017): Quantifying Decentralization. Preprint/Blogpost.

Star, Susan Leigh (1999): The Ethnography of Infrastructure. American Behavioral Scientist, 43(3), pp. 377–391.

Tilson, David / Lyytinen, Kalle / Sørensen, Carsten (2010): Research Commentary — Digital Infrastructures: The Missing IS Research Agenda. Information Systems Research, 21(4), pp. 748–759.

Wood, Gavin (2014/2024): Ethereum: A Secure Decentralised Generalised Transaction Ledger (Yellow Paper). Current version: Berlin Version, 2024.

Ølnes, Svein / Jansen, Arild (2018): Blockchain Technology as Infrastructure in Public Sector: An Analytical Framework, and id. (2021): Blockchain Technology as Information Infrastructure in the Public Sector; also Grimmelmann, James / Windawi, A. Jason (2023): Blockchains as Infrastructure and Semicommons, Dimitropoulos, Georgios (2020): The Law of Blockchain, and Nabben, Kelsie (2023): Web3 as 'Self-Infrastructuring'. Full references are in the footnotes to Section 2.1.

Ølnes, Svein / Jansen, Arild (2021): Blockchain Technology as Information Infrastructure in the Public Sector. In: Reddick, C.G. / Rodríguez-Bolívar, M.P. / Scholl, H.J. (eds.): Blockchain and the Public Sector. Public Administration and Information Technology, vol. 36. Springer, Cham, pp. 19–46. DOI: 10.1007/978-3-030-55746-1_2.

Ethereum Improvement Proposals

Zitierte Vorschläge mit den Seiten ihres Vorkommens. Die Texte selbst stehen unter https://eips.ethereum.org.

Web-Quellen und Daten

Summary of the AWS Service Event in the Northern Virginia (US-EAST-1) Region, December 2021. https://aws.amazon.com/message/12721/

https://blog.ethereum.org/2022/09/15/the-merge

https://eips.ethereum.org/EIPS/eip-1

https://ethereum.foundation/report-2024.pdf

For the Ethereum Foundation documentation of the client architecture cf. https://ethereum.org/en/developers/docs/nodes-and-clients/

https://ethereum.org/en/whitepaper/

https://ethernodes.org

https://etherscan.io/chart/gaslimit

https://etherscan.io/nodetracker

Gaia-X European Association for Data and Cloud. https://gaia-x.eu

https://l2beat.com/scaling/summary

Internet Shutdown in Iran during Mahsa Amini Protests, September 2022. https://netblocks.org/reports/internet-disrupted-in-iran-amid-mahsa-amini-protests-X8qVE8Ap

https://news.earn.com/quantifying-decentralization-e39db233c28e

post-quantum security team, established in January 2026 under the direction of Thomas Coratger. https://pq.ethereum.org

https://protocol-guild.readthedocs.io

For detailed statistics on supply growth, see. https://ultrasound.money

https://vitalik.eth.limo/general/2020/10/08/rollup.html

Bericht zur Eigenkapitalverzinsung, 2024. https://www.bundesnetzagentur.de

Developer Report 2025. https://www.developerreport.com

Why Europe Needs Galileo. https://www.esa.int/Applications/Satellite_navigation/Galileo/Why_Galileo

Government. https://www.gps.gov/systems/gps/space/

IANA Functions Stewardship Transition, 2016. https://www.icann.org

Submarine Cable Map, 2025. https://www.submarinecablemap.com

Annual Review 2024. https://www.swift.com