

Public Blockchains and Regulated Financial Institutions

Key Challenges, Layered Solutions, and a Path Forward

kinexys by J.P.Morgan

digital currency initiative

mit media lab

Contents

Executive Summary 4

Background 5

Key issues in the use of public blockchains 7

Front-running 7

Transaction omission or censorship 8

Receiving unsolicited tokens 9

Gas fees paid to sanctioned entities 10

Others: Quantum Resilience, Privacy, Dealing with Forks, Liveness, & Finality 10

A framework for solutions at different layers 17

Addressing at Application and Smart Contract Layers 20

Participating in Consensus 20

Private Pools 20

Addressing Front-Running and Censorship 23

Addressing the Receipt of Unsolicited Tokens 24

Addressing Gas Fees Being Paid to Sanctioned Entities 26

Addressing at Blockchain Network and Network Governance Layers 28

Addressing Front-Running and Censorship 28

Addressing Receipt of Unsolicited Tokens 31

Addressing Gas Fees Being Paid to Sanctioned Entities 31

Addressing at Regulation Layer 33

Recognize Front-Running of Financial Transactions on Blockchain as Market Manipulation 33

Regulatory Clarity on Receiving Unsolicited Tokens 34

Regulatory Clarity on Sanctions Compliance and Safe Harbor Provisions 34

Next Steps 36

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Acknowledgements

Authors:

Wee Kee, Toh
Global Head of Insights
Kinexys by J.P. Morgan

Sai Valiveti
Head of Engineering for Digital Payments and Liink, Kinexys by J.P. Morgan

Neha Narula
Director
MIT Digital Currency Initiative

Michael Maurer
Software Engineer
MIT Digital Currency Initiative

Raunak Rajpuria
Head of Architecture for Digital Payments and Liink, Kinexys by J.P. Morgan

Contributors:

Umar Farooq
J.P. Morgan

Emma Landriault
Kinexys by J.P. Morgan

Oliver Harris
Kinexys by J.P. Morgan

The MIT Digital Currency Initiative collaborates with external stakeholders to explore practical ideas and design options for digital assets built for the public good. Our work is intended to inform public dialogue about the benefits and risks associated with various technologies and architectures. This report refers to certain products and services offered by J.P. Morgan, as well as products and services offered by other companies, which we draw on as concrete examples. Discussions of these products and services should not be interpreted as endorsements of their design or use or as validation of their performance.

Kinexys by J.P. Morgan is the firm's blockchain business unit focused on building the next generation of financial infrastructure using blockchain technology, and has deployed financial products on both private permissioned and public blockchains. Blockchain Deposit Accounts (BDAs) are deposit accounts recorded on a private, permissioned blockchain ledger, that enable domestic and cross-border payments between BDAs on a 24/7, near real-time basis. Deposit tokens (JPMD) are recorded on blockchains external to J.P. Morgan, including public blockchains. They serve as a digital representation of commercial bank money on public blockchains, and enable institutional clients to make native on-chain payments.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Executive Summary

Public blockchains present a compelling opportunity for financial institutions (FIs) to access open, interoperable infrastructure that can support innovation across payments, tokenization, and capital markets. However, their adoption in regulated financial services remains constrained by a set of structural risks and concerns that are inherent in their design. These risks are not easily mitigated by individual institutions and instead require coordinated action across the technology stack and regulatory environment.

This report identifies key challenges FIs face in the use of public blockchains, and focuses on four that appear straightforward, but remain unresolved: transaction front-running, transaction omission or censorship, the receipt of unsolicited tokens, and the risk of gas fees being paid to sanctioned entities. These challenges persist due to differences in perspective, risk tolerance, and design philosophies between FIs and the broader public blockchain ecosystem.

The report also identified other challenges, many of which are also experienced by non-institutional public blockchain participants. These tend to be better understood, with solutions under active development. However, even in these cases, FIs often have more stringent and nuanced requirements. For such challenges, we focus on highlighting the considerations of FIs to help ensure solutions can better address institutional needs.

To evaluate potential responses, this report introduces a layered framework spanning application, smart contract, token standard, blockchain network, network governance, and regulatory layers. A key insight is that solutions that are most accessible to FIs, such as those at the application and smart contract layers, primarily mitigate symptoms rather than solve the root causes. By contrast, more effective and durable solutions reside at the protocol and governance layers, where FIs have limited direct control. Therefore, meaningful progress depends on collaboration between FIs, protocol developers, and policymakers.

In the near term, FIs can adopt a range of practical mitigations, including engaging with private transaction pools, designing smart contracts with enhanced controls, and implementing operational processes to manage token flows and compliance risks. While these measures can reduce exposure, they cannot fully eliminate the underlying challenges.

Over the medium to long term, more fundamental improvements are required at the blockchain network and governance layers. These include innovations such as encrypted mempools to reduce information leakage, mechanisms to improve transaction inclusion and prioritization, and adjustments to incentive structures to better align network behavior with desired outcomes. However, implementing such changes in decentralized systems is complex and requires broad ecosystem support.

Regulations are also critical. Certain risks, particularly those related to front-running, sanctions exposure, and the treatment of unsolicited tokens, cannot be fully addressed through technical means alone. Clear regulatory guidance is needed to provide legal certainty and define acceptable practices.

Ultimately, the adoption of public blockchains by FIs will depend on the ability of the ecosystem to bridge the gap between decentralized infrastructure and regulatory requirements. No single stakeholder group can resolve these challenges independently. Progress requires coordination. This report aims to support ongoing dialogue among FIs, blockchain developers, and regulators to enable secure and compliant adoption of public blockchain infrastructure.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Background

In a traditional centralized design, a financial institution (FI) typically has control over the full technology stack, either through direct ownership of the components, or contractual governance arrangements that provide them with a high degree of control even if the components are outsourced to third parties. This is different in a public blockchain environment which is designed to be decentralized, and where no single or small group of parties can have complete control over the network.

As such, FIs, as users of the network, have limited control over its underlying infrastructure and governance. Without the high degree of control they are accustomed to, FIs cannot design public blockchains according to their requirements, and must instead adapt to their evolution, which is driven by the broader community. Consequently, they are unable to fully address infrastructure issues when they arise, and are primarily limited to managing associated risks and impacts instead. These challenges, and approaches to mitigating them, have also been highlighted in an industry-led research to create a risk mitigation framework for the use of public blockchains by FIs, facilitated by GBBC and Oliver Wyman.[1]

Solving these challenges requires collaboration between FIs, protocol developers and regulators. However, such collaboration has historically been limited due to differing interests and priorities across stakeholders. Recent discussions, including a roundtable on the use of public blockchains in banking at the Singapore FinTech Festival in November 2025, suggest increasing interest in collaboration across stakeholders.[2] The discussion referenced earlier research by Kinexys and MIT DCI on payment token design,[3] and the above-mentioned risk mitigation framework.

In the discussions, FIs highlighted similar concerns and a shared interest in addressing them. Moreover, there was strong interest across the ecosystem to collaborate on solving such fundamental issues, particularly to enable greater institutional usage of public blockchains. Participants from the public blockchain ecosystem, particularly blockchain protocol developers, expressed eagerness to address barriers to institutional adoption, recognizing that increased institutional usage would broaden network participation and growth. Similarly, policymakers were keen to understand how regulations could evolve to support safe and compliant expansion of institutional adoption.

Motivated by the ecosystem's willingness to collaborate, we are taking a first step in explaining the issues that FIs face when using public blockchains and describing potential solutions that can be applied at different layers. We hope this paper will serve as a starting point for policymakers, blockchain protocol developers, and FIs to collaboratively develop solutions that remove barriers to the use of public blockchains in the financial industry.

This report is intended to be broadly applicable across public blockchain networks in general. However, to provide concrete and practical illustrations, we use the Ethereum Layer 1 (L1) as the primary example of a public blockchain reference point throughout. Ethereum is selected due to its wide adoption, and its support for stateful, general-purpose smart contracts. Unless otherwise specified, references to Ethereum pertain to the network following the Fusaka upgrade, and relevant active improvement proposals.

Where considerations are specific to other blockchain platforms (such as Bitcoin), it will be indicated explicitly. In addition, this report examines Layer 2 (L2) networks where relevant, with a focus on Ethereum L2s currently in deployment. L2s exist in many different varieties, and may differ significantly from one another, as well as from L1s.

1. https://assets.ctfassets.net/so75yocayyva/66juWKzjrjQoWgV6a7gfo3/5d960a9ce90c393c0ef20b95aa00c4e2/Phase_2_Proposed_Risk_Mitigation_Framework_for_Non-Financial_Risks_of_Blockchain_Infrastructure_2-2.pdf
2. https://www.fintechfestival.sg/agenda/agn61-public-blockchains-in-banking-technology-risks-and-token-design-for-safety-and-resilience
3. https://www.jpmorgan.com/kinexys/documents/designing-payment-tokens-for-safety-integrity-interoperability-usability.pdf

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

In many cases, they can more readily address some of the challenges highlighted in this report. However, they often rely on different trust assumptions as compared to L1s. As such, they may introduce new or alternative risks, which are also considered in this report.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Key issues in the use of public blockchains

In this section, we examine the challenges that FIs encounter when using public blockchains. While public blockchains have existed for over a decade, their adoption by FIs, particularly for non-crypto use cases, has only begun to scale in recent years. As a result, many of the issues faced by FIs are still emerging and not yet fully addressed.

Some of these challenges are not unique to FIs and are also experienced by non-FI users. These issues are generally better understood, with solutions already proposed or included on protocol developers' roadmaps. Examples of these shared challenges include the inherent tension between transparency and privacy, where the public nature of blockchain data can expose sensitive transaction information; limitations in network liveness, where congestion, censorship, or adverse conditions may delay or prevent transaction inclusion; and uncertainty around transaction finality, particularly in systems with probabilistic settlement guarantees.

However, even in these cases, FIs often have more stringent or specific requirements and solutions deemed sufficient for non-FIs may still not fulfil FIs' requirements. Articulating these nuances is important to ensure that broadly designed solutions can adequately support institutional use cases. The report describes the issues and highlights requirements specific to FI to ensure that their requirements are well captured and understood.

In this report, we examine four challenges for FIs that appear straight-forward, but remain unresolved. We believe that this gap is driven, in part, by differences in perspective, risk tolerance, and priorities between FIs and the broader public blockchain ecosystem. Four such challenges that we focused on are front-running of transactions, transaction censorship, receipt of unsolicited tokens, and payment of gas fees to sanctioned entities.

These issues may appear relatively minor or manageable to non-FI participants but are treated with significantly greater caution by FIs. This is often due to the need to align with existing regulatory frameworks, the tendency to apply analogies from traditional banking which may not always be appropriate or fully applicable in a public blockchain context, or the adoption of conservative interpretations in areas where regulatory guidance remains unclear. As a result, mitigating these issues can introduce substantial compliance and operational overheads for FIs.

In discussions with public blockchain ecosystem participants, some expressed surprise that these were considered issues for FIs. It is precisely because of this divergence in understanding that has led to limited focus on addressing institution-specific challenges. By outlining these issues and the workarounds currently employed by FIs, this report aims to encourage more targeted efforts across the ecosystem to develop effective and scalable solutions.

Finally, the issues discussed in this section are intended to highlight risks and challenges that may emerge as the use of public blockchains by financial institutions scales over time. While some of these issues are already observable in today's environments, others are forward-looking in nature and reflect how current design characteristics of public blockchains may interact with the requirements of regulated financial institutions at greater levels of adoption. The intention is to contribute constructively to ongoing dialogue among financial institutions, protocol developers, and policymakers, by articulating potential considerations in the context of future, scaled usage scenarios.

Front-running

One of the key challenges in using public blockchains is the potential for front-running of transactions, in which a party gains knowledge of transactions being processed and inserts their own transactions to be processed ahead

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

of others.[4,5]

Front-running is most commonly associated with trading activity, particularly on decentralized exchanges (DEXs). Block builders with knowledge of pending trades may insert their own trades earlier in a block than a third-party's trade, to take advantage of subsequent market movements. This can result in financial losses for the affected parties and undermine market fairness.

For FIs, particularly those enforcing on-chain controls, there is concern that knowledge of incoming administrative activity may be used to evade sanctions or other enforcement actions. For example, a party could monitor for transactions that place their blockchain addresses on a denylist or attempt to seize their tokens, and upon detection, initiate transactions to spend or otherwise move the tokens, such as swapping them for other tokens on a decentralized exchange.

The primary cause of front-running stems from two related technical designs at the protocol level. First, the typical processing flow involves transactions being disseminated in plaintext across nodes to be considered for future blocks. This lack of transaction privacy allows nodes to view unconfirmed transactions. In contrast, most FIs are used to transactions only being visible to a select set of trusted entities during processing.

Second, on public networks, transactions included in a block are often not required to be ordered by time, i.e. a First-In-First-Out (FIFO) model. Instead, they can be reordered into any valid arrangement, often based on extractable value, i.e. transactions more valuable to the block proposer tend to be processed first. This enables a party to observe transactions being broadcasted and then submit their own transactions with higher gas fees, resulting in their transactions being placed and processed ahead of the targeted transactions.

There are also secondary factors that enable such behavior. The application of existing market manipulation laws to front-running on public blockchains is an unsettled area of law, and further regulatory and judicial guidance may be needed to clarify how traditional doctrines map onto on-chain conduct. From a network governance perspective, validators are incentivized to prioritize and reorder transactions to maximize the amount they earn. This practice, known as Maximal Extractable Value (MEV) extraction may result in unfair outcomes for users, but benefits miners by increasing miner rewards, which incentivises miner participation.[6] If FIs deploy assets that are available to trade in DEXes, and FI clients become more heavily involved in on-chain trading activity, solving for MEV-style front-running will likely be necessary for both consumer protection and regulatory compliance.

Transaction omission or censorship

In decentralized ledgers, transactions may be omitted for various reasons. Block proposers may omit transactions because they simply are not aware of them, or if they do not offer enough economic benefit for inclusion. They also may elect to omit transactions for any other reason, including strategic or malicious considerations, which is commonly referred to as censorship.

We view this primarily as a network governance issue, where there are no defined prioritization rules and block proposers are not obligated to order transactions according to prescribed rules. Instead, transactions are ordered based on the gas fees earned, rather than the priority or criticality of the transactions. Gas fees are an integral part of managing the supply and demand for computational resources on a blockchain. However, prioritizing processing

4. https://cow.fi/learn/what-is-frontrunning
5. https://en.wikipedia.org/wiki/Front_running
6. MEV also is sometimes referenced as "Miner Extractable Value"; See: https://cow.fi/learn/is-mev-good-or-bad

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

solely based on gas fees can effectively lead to inadvertent censorship under certain conditions.

For example, if there is a "run" on a token issuer, a surge of transactions to sell or redeem tokens may temporarily exceed the processing capacity of the blockchain network. Token holders are incentivized to pay increasingly higher gas fees to ensure their transactions are processed. In such an event, a token issuer may wish to initiate specific administrative actions, such as a global pause to suspend further transactions involving the token.[7] The economics of gas fees and the run scenario could prevent the token issuer from successfully initiating these transactions.[8] In many systems, different priorities are assigned to different types of transactions, and administrative transactions are often given the highest priority. However, such granular prioritization mechanisms are not always present in public blockchains.

Receiving unsolicited tokens

FIs are also concerned about receiving unsolicited tokens or having tokens associated with blockchain addresses for which they hold the keys. This issue is exacerbated when the tokens originate from addresses on sanctions lists or are known to belong to illegitimate parties. Such parties sometimes distribute tokens to a large number of blockchain addresses to confuse blockchain analytics tools, while others "airdrop" tokens to many addresses as a marketing or promotional strategy. The root issue lies with the protocol – specifically, how token transfers may be unilaterally performed on a blockchain.

A transfer occurs when a smart contract updates the balances of addresses. The address a party uses to transact in one particular token is valid not only for that token but also for all other tokens on the same blockchain, and even for tokens on other blockchains that use the same address scheme. For example, a party with a blockchain address intended for transacting in JPMD on Base could, with the same address, receive other tokens on Base as well as other tokens on public Ethereum.

Currently, parties have limited ability to prevent themselves from receiving tokens or interacting with other addresses and/or smart contracts, except through functions provided by the smart contracts. However such functions are not standard and are not widely adopted. As the transfer functions are implemented at the smart contract level rather than a standard primitive at the protocol level, there is no universal mechanism to block incoming transfers.

A secondary issue is the lack of regulatory clarity regarding the need to address, and the methods for dealing with, tokens received from unknown or potentially sanctioned entities. As a result, FIs often take a conservative approach, drawing an analogy to receiving funds in their accounts, where the response is to return the tokens to the sender. Given that tokens may exist across multiple blockchains and smart contracts, FIs must monitor a broad set of networks and all emitted events to detect if they have received unsolicited tokens.

Because transfer functions are implemented by smart contracts and may differ across contracts, FIs usually conduct due diligence audits and reviews of the smart contracts before using the transfer function to return the tokens. This creates unnecessary effort simply to respond to the receipt of unsolicited tokens. Furthermore, many of the tokens received are promotional, such as newly issued tokens, and often have negligible or no value. A more fitting analogy is that of junk mail, which then an FI must spend time and effort to analyze the contents and their origin before returning it.

7. This ability may help preserve confidence in the system, similar to a stock market circuit breaker https://www.investor.gov/introduction-investing/investing-basics/glossary/stock-market-circuit-breakers
8. MakerDAO's "Black Thursday" was made significantly worse due to on-chain price oracles not getting updated in a timely manner due to network congestion https://medium.com/linum-labs/black-thursday-makerdaos-multi-collateral-dai-exploitation-and-the-plan-to-recover-c083c0b81875

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Gas fees paid to sanctioned entities

Gas fees are essential for maintaining public blockchain networks, serving as a financial incentive for block proposers to process transactions and secure the network. The concern for FIs is not the payment of gas fees itself, but rather about the recipients of these fees, and the possibility that they may be sanctioned entities, potentially resulting in inadvertent sanctions violation.

The root cause lies in the network governance of public blockchains, where there are typically few or no rules or onboarding requirements for block proposers. This issue is further complicated by the existence of mining and staking pools, where the identity of the block proposer may be known, but the identities of the underlying entities that comprise the pools are often unknown. While some layer 2 (L2) solutions have stricter network governance, significantly reducing the likelihood of gas fees being paid to sanctioned entities, L2s still rely on layer 1 (L1) blockchains. Periodically, data and transactions are "rolled up" and posted to the L1. In these cases, there remains a possibility of gas fees being paid by the L2 to sanctioned entities on the L1. As a result, there remains a residual possibility of an FI indirectly paying gas fees to sanctioned entities, although the amounts involved would be exceedingly low due to the batching of transactions.

However, sanctions are a matter of strict liability, and parties can be held liable regardless of intent. Clearer guidance regarding gas fee payments would help, but in its absence, FIs tend to adopt an extremely cautious approach, treating even a remote possibility of exposure to sanctioned entities as a material risk.

Quantum Resilience

On public blockchains, assets are secured on transparent, public ledgers and are associated with persistent identifiers. By design, everyone can see all of the information required to validate a transaction's legality. This is both possible and safe due to modern cryptography, such as elliptic curve cryptography. However, the assumptions we can today safely make about many modern, efficient cryptography schemes do not hold in the presence of a quantum computer.[9]

In Ethereum, externally-owned addresses (EOAs) are derived from hashes of ECDSA public keys, and valid transactions require signatures generated using the corresponding private keys.[10] Control of the funds is therefore dependent on knowledge of the corresponding private keys.

However, for certain widely used schemes like ECDSA, a powerful quantum computer could derive a private key from a known public key.[11] If such capabilities were realized without corresponding upgrades to the underlying cryptography, a large portion of existing accounts on networks such as Ethereum and Bitcoin could become vulnerable to compromise.

The most immediate concern is a "long-range attack," where public keys exposed on-chain (e.g., through transaction activity) could be targeted by even a moderately powerful quantum computer. As quantum capabilities advance, there is also concern about "short-range attacks," where the time required to compute the corresponding private key for a newly revealed public key becomes sufficiently short to impact recently broadcast transactions (i.e. those sitting

9. https://quantumai.google/static/site-assets/downloads/cryptocurrency-whitepaper.pdf
10. https://dev.singularitynet.io/docs/products/DecentralizedAIPlatform/UsedTechnologies/ethereum-address/
11. Elliptic curve private and public keys are mathematically related, but we can assume one cannot determine the corresponding private key for a public key due to the hardness of the elliptic curve discrete logarithm problem (ECDLP). With a quantum computer, the hardness assumption does not hold. This is true for nearly all widely deployed public-key cryptosystems. Many post-quantum public-key systems however are under development or are becoming adopted: https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

in mempools).

Significant research and development efforts are underway to address these risks, but there remains considerable uncertainty regarding the appropriate mitigation strategies. There is ongoing debate about how to make these platforms quantum-resilient, and relatedly, about the extent to which users will maintain confidence in these systems during and after such a transition. Addressing these issues is inherently more complex in decentralized systems, where coordinating upgrades is often difficult,[12] than in centralized environments, where upgrades can be deployed at the behest of a small set of operators, and often with minimal delay. In public blockchains, changes require broad consensus, and may lead to unintended forks.[13]

Within the Ethereum ecosystem, several approaches have been proposed to improve quantum resilience. For example, smart contract-based accounts (as opposed to EOAs) can define flexible authorization logic, including the use of alternative, potentially quantum-resistant signature schemes.[14]

Proposals such as EIP-7702 further expand this flexibility by enabling EOAs to adopt alternative validation logic defined by account code. More broadly, there is an ongoing shift toward account abstraction, with some ecosystem participants advocating for the eventual deprecation of EOAs in favor of smart contract-based accounts.[15]

Beyond account cryptography, Ethereum needs to upgrade consensus-layer signatures, commitments used for data availability,[16] and some rollups and applications in the Ethereum ecosystem will need to upgrade to post-quantum proof systems.[17]

Nonetheless, post-quantum signature schemes have not yet been subjected to the same level of long-term scrutiny and real-world testing as established classical cryptosystems. These schemes are also currently larger and slower than what is used today,[18] which will likely increase the hardware requirements to run a validating node and/or degrade overall network throughput.[19] As a result, it remains possible that even after initial adoption, further iterations or transitions in cryptographic standards will be required. This highlights the importance of maintaining robust coordination mechanisms and upgrade pathways to support future cryptographic migrations.

Privacy

Most public chains have a very different privacy profile compared to traditional banking systems. Transactions on Ethereum are publicly visible, and the effects are apparent as soon as they are observed. This transparency also contributes to challenges such as front running. At scale, complete public visibility may be untenable for certain use-cases. Many market participants value privacy highly, and businesses are often dependent on confidentiality of order flow.[20]

12. To some, this is a feature rather than a bug. Resistance to change gives users confidence that the system will not abandon its commitments, and that it prevents any single developer or corporation from changing the protocol without sufficient community acceptance. https://a16zcrypto.com/posts/article/why-decentralization-matters-incentivizing-decentralization-incentives/
13. Different networks also have different challenges in rolling out post-quantum fixes: https://medium.com/@gwrx2005/post-quantum-roadmaps-for-blockchain-ecosystems-af9e77a6fe8b
14. https://ethereum.org/roadmap/future-proofing/quantum-resistance/#eoa-signatures
15. https://eips.ethereum.org/EIPS/eip-7702
16. For blob data for example https://ethereum.org/roadmap/danksharding/#how-are-blobs-verified
17. https://x.com/VitalikButerin/status/2027075026378543132; https://ethereum.org/roadmap/future-proofing/quantum-resistance/#four-vulnerable-areas
18. Signatures, for example, become much larger. One Dilithium2 (a PQ scheme) signature is about 38x the size of an Ed25519 signature. This heavily impacts storage and network speed requirements for validators in the network: https://pqcvisualizer.com/?algos=ed25519,dilithium2
19. https://www.akamai.com/blog/security/digital-signatures-traditional-post-quantum-cryptographic
20. See Chart 4: https://www.ecb.europa.eu/pub/pdf/other/Eurosystem_report_on_the_public_consultation_on_a_digital_euro~539fa8cd8d.en.pdf

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

There are different degrees of privacy that can be achieved. A very basic degree of privacy can be achieved by not allowing personal identifiable information (PII) to be stored on-chain, and to instead rely on pseudonymous identifiers, such as Ethereum addresses which are not publicly linked to real world identities.[21]

However, given that the transaction graph on a network like Ethereum is publicly visible, an astute observer can infer the identities of participants.[22] To achieve a stronger degree of privacy on-chain, it is important to reduce the linkability between sending and receiving addresses, thereby limiting the information revealed when one or more addresses are identified. Additionally, hiding the value of transactions is important for privacy, as transaction amounts with some external context may allow much to be inferred.[23]

For institutional users, privacy is particularly important as businesses typically seek to avoid revealing sensitive information such as cash flows and pricing that could disadvantage them competitively.

Different approaches exist for dealing with privacy on different ledgers. For example, Zcash supports shielded transactions where transaction details are cryptographically hidden, with the help of zero-knowledge (zk) proofs. These details can only be viewed with a particular cryptographic key,[24] while the validity of the transaction can be publicly verified via succinct zero-knowledge proofs.[25] However, this imposes strong limitations on programmability; encoding general-purpose computation as a zero-knowledge circuit is complex and resource-intensive, and supporting arbitrary smart contracts requires either a circuit per computation or a zkVM.[26] Mixers are another commonly discussed technique (particularly on Ethereum) but they do not conceal transaction occurrence, and instead primarily reduce the ability to link the source and destination of funds.

Dealing with Forks

A distinguishing characteristic of public blockchains is their ability to fork. A fork occurs when subsets of validators disagree on the state of the chain, creating two divergent states. A blockchain fork can be compared to a fork in the road. Up to the point of divergence, the ledger shares a common history. Beyond that point, two distinct chains exist, each with its own state.

Forks can and often do occur naturally, such as when two block proposers broadcast valid blocks at the same height connected to the same parent. In Bitcoin, mining a valid block is in effect a random process, and as a result occasionally two miners will discover and broadcast a valid block around the same time, potentially creating a fork. In many cases, one branch quickly becomes dominant while the other loses support and is eventually abandoned.

However, there are notable instances where both branches continue to persist with meaningful adoption, often after an intentional fork. This occurs as validators do not agree on what the state or rules should be for the network, and is often called a “hard fork”. A hard fork splits the network, while also “duplicating” the assets, as they now exist on both networks.

For example, the split between Bitcoin and Bitcoin Cash, caused by disagreement over what Bitcoin’s blocksize should be, resulted in two actively maintained networks, which have some shared history but now have incompatible states

21. https://www.dci.mit.edu/projects/enhancing-the-privacy-of-a-digital-pound
22. https://dl.acm.org/doi/pdf/10.1145/2504730.2504747
23. https://dl.acm.org/doi/abs/10.1145/3428332
24. See “Viewing Keys”: https://zcash.readthedocs.io/en/latest/rtd_pages/basics.html
25. https://www.di.ens.fr/~nitulesc/files/Survey-SNARKs.pdf; https://z.cash/learn/what-are-zk-snarks/
26. zkVMs are an active area of research: see https://eco.com/support/en/articles/11803106-what-is-risc-zero-zkvm-verifiable-computation; https://www.lita.foundation/blog/zero-knowledge-paradigm-zkvm

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

and consensus rules.[27] Similarly, Ethereum has experienced several notable forks: the Ethereum Classic fork in 2016, which reflected a fundamental philosophical divide within the community,[28] and the emergence of EthereumPoW after "The Merge", where some validators chose to retain a proof-of-work variant of the network.[29] Software bugs can also lead to forks, as was observed due to bugs in some Ethereum client software in 2021.[30]

Today, Ethereum regularly has hard forks as part of rolling out upgrades. However, since changing to proof of stake, Ethereum stakers may incur significant monetary losses if they do not fulfil their duties on the more valuable side of the fork. As a result, there is stronger economic alignment towards maintaining a single canonical chain.[31]

When both sides of a fork remain viable, a unique situation emerges in which assets effectively exist on multiple chains. For example, following the Bitcoin Cash fork, holders of Bitcoin prior to the split also ended up with equivalent assets on the Bitcoin Cash network, with both assets independently transferable within their respective ecosystems.

This dynamic presents challenges for tokenized assets that are intended to represent claims on real-world value. In such cases, a fork can lead to duplicate on-chain representations of a single underlying asset. One common mitigation strategy is for the issuing institution to designate a single supported chain, treating assets on the alternative chain as invalid. For instance, Circle elected to support only the proof-of-stake Ethereum chain following the Merge, ensuring that USDC retained a single canonical representation.[32]

However, this approach introduces operational complexity and requires continuous monitoring of network conditions to detect and respond to forks. The challenge is particularly pronounced for assets designed to minimize issuer intervention, as parallel markets across competing chains can rapidly create ambiguity.[33] In addition, because forks can occur unpredictably, FIs may need to monitor multiple competing chain states until a canonical outcome emerges, either through market adoption or issuer designation. More broadly, this highlights a fundamental tension: while decentralized systems enable open participation and resilience, certain use cases, particularly those involving real-world asset backing, introduce points of centralization in determining which representation is authoritative.

Liveness

In this report, public blockchains effectively take on the role of traditional centralized infrastructure. A key feature of centralized systems is the presence of an accountable operator, typically bound by a service-level agreement (SLA), which defines acceptable performance standards and remedies in the event of service degradation. This includes scenarios where the network becomes unavailable or loses liveness.

In contrast, users of public blockchains do not enter into formal SLAs with the network itself. Instead, liveness is maintained through a combination of protocol-level incentives and system design choices.

At a high level, decentralized public ledgers rely on economic incentives to encourage participation. Validators or

27. This was in response to the SegWit upgrade https://en.wikipedia.org/wiki/Bitcoin_Cash#History
28. https://help.coinbase.com/en/coinbase/getting-started/crypto-education/eth-hard-fork
29. https://decrypt.co/110004/ethereum-merge-hard-forks-reveal-competing-proof-of-work-projects
30. https://www.theblock.co/post/115822/bug-impacting-over-50-of-ethereum-clients-leads-to-fork
31. This comes with the tradeoff that there is more friction to propose blocks in PoS Ethereum than a PoW system. It also increases friction when trying to exit the system and stop validating blocks.
32. https://www.circle.com/blog/usdc-and-ethereums-upcoming-merge
33. This is a challenge as well for ETFs tracking cryptocurrencies; see "Forks" https://www.sec.gov/Archives/edgar/data/1980994/000143774923033524/bit20231201_s1a.htm

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

miners are rewarded for proposing and validating blocks that are accepted by the network, and these rewards can be financially significant. This creates a strong baseline incentive for participants to maintain network operation.

On Ethereum, additional mechanisms exist to reinforce liveness. Validators are expected to actively participate in block proposal and attestation duties; failure to do so results in missed rewards. In more severe cases, such as when the network is unable to finalize blocks, Ethereum activates an "inactivity leak" mechanism. If finality is not achieved for a defined period (currently four epochs), inactive validators begin to incur penalties, gradually reducing their stake and restoring the network's ability to reach consensus.[34] This mechanism provides a strong economic incentive for validators to remain online and responsive during periods of instability.

Layer 2 systems introduce a different set of liveness considerations. Many L2 networks rely on a sequencer responsible for ordering and batching transactions before submission to Layer 1. In several widely used implementations, this sequencer is operated by a single entity, creating a potential central point of failure. If the sequencer becomes unavailable, transaction processing may halt, and users may be unable to interact with the network through standard interfaces.

In such cases, users can typically fall back to interacting directly with L1 smart contracts to recover funds or force transaction inclusion. However, this process is often slower and more complex. For example, in optimistic rollup designs, withdrawals may be subject to a challenge period that can last several days (commonly around seven days), delaying access to funds.

Bitcoin provides a contrasting model. It does not impose explicit penalties for inactivity, as participation in mining is entirely permissionless. Instead, liveness is maintained through market-driven incentives. Miners are rewarded for successfully mining blocks, but face no direct penalty for going offline. The network adjusts to changes in participation through its difficulty adjustment mechanism: if overall hash rate decreases and blocks are produced more slowly, mining difficulty is reduced, making block production easier and incentivizing new or returning miners to participate.[35]

Taken together, these systems demonstrate that while public blockchains do incorporate mechanisms to preserve liveness, they differ fundamentally from traditional legal or contractual guarantees. In some cases, particularly with centralized L2 sequencers or permissioned blockchain environments where there are known operators, it may be possible to introduce more conventional legal arrangements. FIs may find such models more familiar and operationally comfortable, although publicly standardized SLA frameworks for these systems remain limited at present.

Finality

For FIs operating within traditional financial systems, the notion of payment finality is relatively well defined and predictable. The point at which a transaction becomes final is typically governed by legal frameworks, regulatory requirements, or platform-specific rules. As such, finality can be relied upon, provided there is confidence in the legal system and its ability to enforce these rules in the event of disputes.

In decentralized networks, however, the concept of transaction finality is less straightforward. FIs and legal scholars often question whether robust finality can be achieved in the absence of a legal framework.[36] This uncertainty stems

34. https://eth2book.info/capella/part2/incentives/inactivity/
35. https://www.lightspark.com/glossary/difficulty-adjustment
36. Probabilistic Settlement Finality in Proof-of-Work Blockchains: Legal Considerations" https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4022676

in part from the differing ways by which blockchain systems define and achieve finality.[37]

In proof-of-work (PoW) systems such as Bitcoin, finality is probabilistic rather than absolute. A commonly cited heuristic is that a transaction may be considered "final" after six confirmations, meaning that the block containing the transaction is followed by six additional blocks on the canonical chain. This convention originates from the Bitcoin whitepaper and reflects an estimate of the number of confirmations required to reduce the probability of reversal below an acceptable threshold under certain assumptions about adversarial hash power.[38]

Importantly, probabilistic finality is sensitive to the distribution of hash power. A sufficiently large adversary could, in principle, reorganize the chain and reverse transactions. While a 51% attack represents a critical threshold at which an attacker could consistently rewrite history, even lower levels of control can increase the time required to achieve a desired level of confidence. Accordingly, users must calibrate their required confirmation depth based on their risk tolerance and assumptions about the security of the network.[39]

The Ethereum ecosystem typically depends on economic finality guarantees, not probabilistic ones. Under its proof-of-stake (PoS) consensus mechanism, blocks are considered finalized once validators controlling a supermajority (two-thirds) of staked ETH attest to a sequence of checkpoints.[40] Reverting a finalized block would require an attacker to control or compromise a substantial portion of the total staked ETH and incur significant economic penalties, including the potential loss of a large share of the stake. In practice, checkpoints occur at regular intervals, and finality is typically achieved within minutes under normal network conditions.

Before finalization, however, Ethereum exhibits similar properties to PoW systems, where recently proposed blocks may still be subject to short-range reorganizations.[41] While such reorgs are typically shallow, they highlight the distinction between block inclusion and finality.[42] Ongoing research efforts, such as proposals for single-slot finality, aim to reduce this window further.[43]

Layer 2 systems introduce additional nuance. Finality on L2s is multi-layered: users may trust the sequencer and choose to treat a transaction as final upon sequencer confirmation,[44] or wait until the corresponding state transition is committed and finalized on the underlying L1. In many architectures, transaction data is periodically batched and posted to L1, after which finality depends on the settlement guarantees of the base layer.[45] This results in a spectrum of finality definitions with differing trade-offs between speed and security.

Overall, while these notions of finality differ from those in traditional finance, the gap may be narrower than it initially appears. Economic finality provides strong guarantees under reasonable assumptions about network security and incentives. However, unlike legal finality, it ultimately depends on the continued value and integrity of the underlying staked asset. In adversarial scenarios, particularly in systems without strong economic deterrence or coordination mechanisms, confidence in finality may degrade. For FIs, these distinctions remain a key consideration

37. This is also closely related to a blockchain's ability to fork.
38. https://bitcoin.org/bitcoin.pdf
39. https://blog.lopp.net/how-many-bitcoin-confirmations-is-enough/
40. https://ethereum.org/developers/docs/consensus-mechanisms/pos/gasper/
41. https://eth2book.info/latest/part3/forkchoice/
42. https://ethresear.ch/t/post-pectra-effects-on-ethereum-reorg-rate-propagation-and-block-size/23018
43. https://ethereum.org/roadmap/single-slot-finality/
44. This is typically a delay of ms to a few seconds unless congestion is particularly high. The three largest L2s by total value secured (https://l2beat.com/scaling/summary) have block times of 250ms (Arbitrum One: https://research.arbitrum.io/t/the-power-of-faster-blocks/9609), and 2s (Base, without Flashblocks https://base.blockscout.com/stats, OP Mainnet: https://explorer.optimism.io/)
45. See L1 Batch Inclusion / Finality: https://docs.optimism.io/op-mainnet/network-information/transaction-finality#finality-for-op-mainnet-l2-transactions

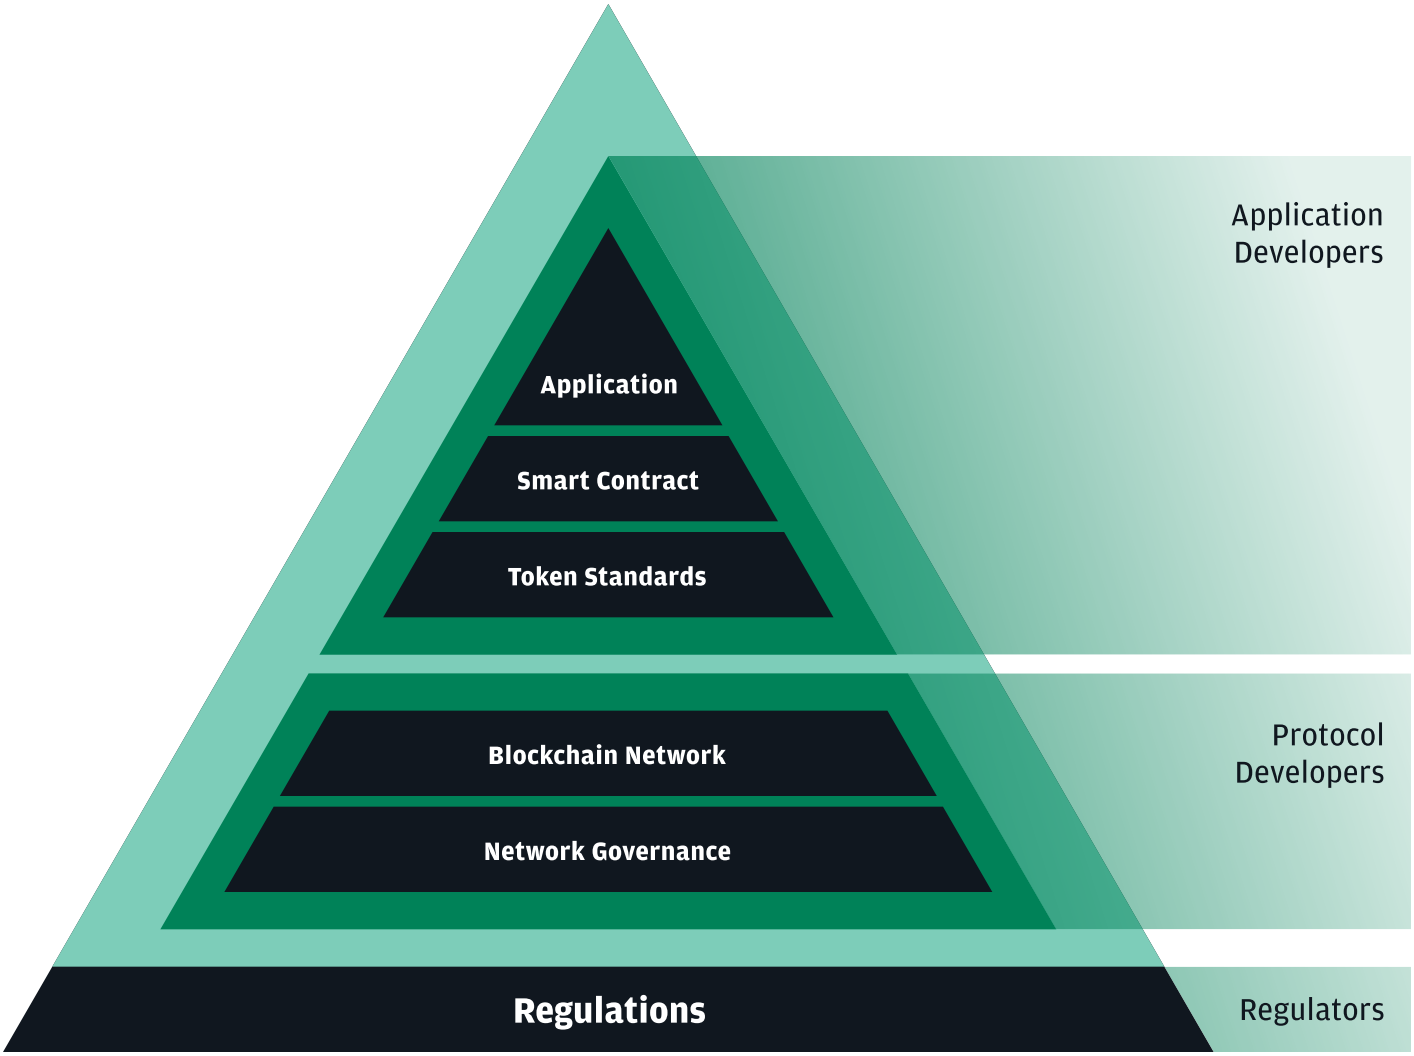
 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

when evaluating the use of public blockchains. Where necessary, complementary legal or regulatory frameworks for on-chain assets may help bridge remaining gaps.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

A framework for solutions at different layers

We defined a framework to evaluate how solutions can be applied at different layers. The framework comprises both technical and governance layers. The different layers, and how they relate to each other, can be viewed as: **Application**, using **Smart Contracts**, built to **Token Standards**, deployed on a **Blockchain Network**, subjected to **Network Governance**, with the entire arrangement subjected to **Regulations**. The individual layers are described as follows.



 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

- **Application:** What the user uses or interacts with, and which subsequently interacts with a decentralized ledger and smart contracts as part of the application's logic. This includes special-purpose applications such as those provided by the FIs, or general-purpose applications like blockchain wallets that could interact with a large number and variety of tokens.
- **Smart Contract:** The logic as written and deployed on-chain.
- **Token Standards:** A commonly agreed set of specifications that is shared by many contracts deployed on-chain, allowing for cooperation between many smart contract developers and interested parties.
- **Blockchain Network:** The platform on which smart contracts are deployed and executed, and where transactions, state updates, and other data are recorded. This network is formed by a set of nodes that communicate and agree with each other on the current global state in a process referred to as consensus. The nodes adhere to the protocol specifications that have been developed and agreed to by the community.
- **Network Governance:** Rules established to ensure orderly functioning of a network. In a decentralized network without a central operator, rules are typically in the form of social norms rather than formal agreements enforceable in a court of law. This includes incentive or disincentive structures within the protocol or at the community level.
- **Regulations:** Rules established by a competent authority, typically to ensure stability, integrity, and fairness, and to protect consumers or users. This includes incentive structures at the legal or governmental level, external to the decentralized ledger itself.

Different groups of parties will have different degrees of control over the different layers, hence the solutions that they can implement to address the issues will also differ. Users of a blockchain network, including application developers, will have the greatest control over the Application and Smart Contract layers, with a degree of influence on Token Standards. Protocol developers will have the greatest control over Blockchain Network and Network Governance, while regulators will have the greatest control over regulations.

While financial regulations are primarily applied directly to FIs, they often include provisions on third-party outsourcing and technology risk management. As such, while protocol developers and operators are not directly subjected to financial regulations, FIs' use of the blockchain networks are often subject to such regulations, which they would in-turn apply as requirements for blockchain networks that they use.

Table 1 summarizes potential solutions to the four issues, and how they could be applied at the different layers. In cases where the root issues of a problem are at a lower level, application-level solutions are typically imperfect workarounds or mitigations, and solutions on the lower foundational layers tend to be more effective. The solutions which we find to be most effective are marked in green.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Table 1

	1. Front-running	2. Transaction Censorship	3. Receiving unsolicited tokens	4. Gas fees paid to sanctioned entities
Application	• Participating in consensus, by self-hosting nodes and publishing blocks • Using private pools			
		• Offering attractive fee rates		
Smart Contract	• Contract-imposed liveness enforcement • Logic obfuscation		• Allow/deny-listing on transfers	• Block reward address introspection
Token Standards			• Incorporate allow/ deny-listing into token standards	
Blockchain Network	• Enable encrypted mempools • Support block position introspection	• Congestion Reduction • Prioritization of certain state updates	• Incorporate allow/ deny-listing into network-level capabilities	• Manage block reward
Network Governance	• Update incentive design to discourage frontrunning / MEV	• Prioritization of certain state updates from the governance layer		• Set and enforce rules on who can be validators
Regulations	• Recognize front-running of financial transactions as a form of market manipulation		• Provide regulatory clarity on receiving tokens from unknown or sanctioned entities and how to deal with such tokens	• Provide regulatory clarity on requirements to prevent transacting with sanctioned entities, and safe harbor if such transactions still happen

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Addressing at Application and Smart Contract Layers

As users of blockchain networks, FIs develop and deploy smart contracts on the blockchains, as well as applications that interact with the smart contracts. Solutions at these layers are hence within their direct control and would be easiest to implement, imperfect as they are. As one of the root issues for FIs is the lack of control over the foundational layers, a natural solution for users is to regain control.

This could be through hosting their own nodes and participating in consensus, where they regain some degree of voting power and blockspace on the blockchain network. As a decentralized system, no one party would have complete control of the network, but they would have some degree of control that corresponds to their portion of mining or voting power on the network.

Larger mining pools or block builders are likely to have a greater degree of control than individuals running nodes. FIs can utilize such pools where they can exercise governance control through a commercial arrangement with a third party with a higher degree of voting power.

Participating in Consensus

As a user or institution looking to avoid transactions being front-run, censored, or executed in unfavourable circumstances, the most straightforward mechanism available is to participate in consensus, by running a node and publishing blocks. In fact, this is one of the main proposed benefits of public decentralized ledgers.[46] This inherently gives any participating user the ability to include transactions they want on the ledger when they get the chance to create blocks. This reduces the likelihood of transaction censorship and also allows the user to order transactions. Processing transactions through self-hosted nodes also ensures no tangential relationship with an unwanted entity. That is to say, if the financial institution is publishing the blocks containing its transactions, it retains control on who the recipient of the block fees is.

However, the proportion of transactions a user can self-publish is proportional to the amount of mining or voting power the user has.[47] For many users, it may be infeasible to acquire enough voting power to frequently publish transactions. For example, on the Bitcoin ledger, solo-miners only published 22 blocks in 2025.[48]

Running an Ethereum staking node requires at minimum 32 ETH to be staked,[49] and a minimum staking node on average gets to propose a block about once every five months.[50] FIs could also form a pool with other like-minded institutions to act collectively and increase their mining or staking power on a public blockchain.

Private Pools

Preventing censorship and front-running by self-hosting nodes may not always be feasible, either due to budgetary constraints or laws and regulations. In such cases, users can consider establishing relationships with mining pools or

46. https://bitcoin-takeover.com/10-reasons-you-should-run-your-own-bitcoin-full-node/amp/
47. In Bitcoin, this is the proportion of total hashrate a user has. In Ethereum, this is related to the proportion of stake a user controls.
48. https://cryptoslate.com/how-solo-bitcoin-miners-hit-22-blocks-hit-in-12-months-as-the-global-energy-lottery-gets-hyped-up/
49. Close to $75K at the time of writing.
50. Roughly every 5 months, per https://www.figment.io/insights/beyond-the-basics-understanding-rewards-on-ethereum/

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

entities that publish a significant proportion of blocks. Oftentimes, these nodes will establish relationships with users and offer to process transactions without distributing them to other nodes. In these arrangements, the user may rely on legal agreements, commercial terms, or reputation risk, to dissuade censorship or front-running.

Private order flow of this nature is fairly common on Ethereum, comprising 30–40% of all transactions in blocks.[51] This pattern may be more comfortable for FIs, in particular because they can establish relationships with known and addressable entities, as is done with other parts of the technology stack in modern architectures.[52] In Ethereum, this has also become easier since proposer-builder separation was formalized. However, a downside is that large private pools have a centralizing effect on the networks, putting control in a smaller subset of entities.

Participating in Consensus and Publishing Blocks

There are different consensus mechanisms, such as Proof of Work and Proof of Stake, and this affects the process of how blocks are published.

In Bitcoin, which uses Proof of Work, the primary rule to block acceptance and inclusion is that the block must have sufficient "work" associated with it. Work is computational effort, and Bitcoin consensus rules require that the next block added to the chain meets the work threshold (difficulty) specified by the network. Nodes do "work" by searching for a unique value (a nonce) which, when combined with the proposed block, hashes to a value meeting the difficulty target. The only known way to do this search today is by generating hashes with new random values repeatedly, and so meeting a difficulty target is expected to take a certain number of attempts. In order to publish blocks more frequently, a miner needs to find more random number and block combinations exceeding the difficulty target, and as a result, miners or mining pools with more "hash power" end up publishing more blocks to the network.[53]

Ethereum, on the other hand, is a Proof-of-Stake system.In the Ethereum protocol, the validator set is known[54] and the protocol pseudo-randomly selects a given validator from the set to propose a block for a given "slot." The likelihood of being chosen as the block proposer is correlated to the validator's amount of stake. Solana, Sui, and others use a similar stake-based scheme to decide block proposers.

Further different designs exist in this space. Stellar's variant of Federated Byzantine Agreement leans on agreement between trusted nodes to determine which transactions to include, so the barrier to getting transactions included in the ledger depends on whether other nodes trust the proposing node.[55] Hyperledger Sawtooth uses a Proof of Elapsed Time, relying on TEEs to force a node to wait a random amount of time to

51. https://ethresear.ch/t/expanding-mempool-perspectives/22022 comments that 30-40% of transactions are not visible in the mempool databases of Flashbots and Xatu
52. For example, if you cannot run servers on your own, you can always establish a relationship with a cloud service provider instead.
53. https://en.wikipedia.org/wiki/Proof_of_work
54. Though not necessarily identified.
55. https://developers.stellar.org/docs/learn/fundamentals/stellar-consensus-protocol

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

propose a block, with the shortest wait time "winning".[56] The likelihood of getting transactions included in the ledger is dictated by these rules and should be well understood for any system on which an FI intends to be a major participant.

Proposer-Builder Separation (PBS) in Ethereum

In Ethereum, there are different parties who FIs can work with to get transactions included on the ledger. This is due to Proposer-Builder Separation where the block proposer and block builder could be different entities.[57]

In such a workflow, the proposer may solicit a block to propose from a builder, who essentially buys the blockspace of the proposer. Often this is economically appealing for validators, particularly those with limited available hardware. Generally speaking, these blocks specify the builder as the recipient of the transaction fees, and include the bid (payment to the validator from the builder) within the block itself. There is a third participant in PBS currently, known as the relay, who essentially brokers trust to facilitate the relationship between the builder and proposer. ePBS, proposed in EIP-7732, looks to remove the relay and allow this relationship to exist with only two parties.[58]

Over 90% of recent blocks have participated in the “MEV-Boost” protocol, which is the most common variant of blocks proposed through the PBS schema. Of these blocks, they were proposed by fewer than 100 unique builders, and in January 2026, 45% of these blocks were built by one builder (Titan Builder).[59] This highlights the potential for institutional relationships with block builders in getting timely transaction inclusion to the network.

Importantly, validators are not required to use PBS. They may still build and propose “vanilla” blocks, which helps to decentralize decisions around what transactions exist on the network.

56. https://medium.com/kokster/understanding-hyperledger-sawtooth-proof-of-elapsed-time-e0c303577ec1
57. https://ethereum.org/roadmap/pbs/
58. https://eips.ethereum.org/EIPS/eip-7732
59. https://mevboost.pics/

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Addressing Front-Running and Censorship

Offering Attractive Fee Rates

High fee rates reduce the likelihood of censorship. If a transaction is highly profitable to mine, it is unlikely that it will be ignored, as most block proposers act in their best economic interest.[60] With respect to preventing front-running or MEV extraction, high fee rates are less effective on their own, since they cannot typically prescribe where in the block a transaction will be executed. In some cases, if a transaction is distributed between many nodes and is therefore publicly visible, there is another front-running vector in which a third party attempts to offer an even higher fee for a transaction that would be invalidated after the already high-fee transaction in question is executed. For example, if an FI wanted to add Eve to the deny-list and offered a high fee hoping to get the transaction included quickly, Eve may spot this transaction during the gossip phase and create an even higher fee transaction moving all of her funds before the deny-list update is processed.[61]

Contract-Imposed Liveness Enforcement

One specific concern with front-running is where a token issuer wants to pause a smart contract, and holders try to front-run the pause transaction with other transactions to use or withdraw their tokens.

In such controls -oriented scenarios, it may be better to fail-secure: if the FI cannot interact with the contract or make state updates, the contract pauses or limits what actions can be taken until the FI un-pauses the system. A dead man's switch is a typical example of fail-secure, where a machine stops when it loses inputs from a user.

Heart-beat Locking as a Fail-secure Mechanism

A smart contract can be set to pause at a pre-set future time or block height, with a periodic "heart-beat" to set it to a further time or block height. This acts like a countdown timer that is periodically reset, and in effect, creates a dead man's switch, where a smart contract will pause and stop working shortly after the "heart-beat" stops.

The drawback is that the financial institution must have a transaction included in every N-th block for the contract to work, which introduces a stronger dependency on the FI to stay online and incurs higher fees. It is also likely that this will contribute to a large on-chain state footprint. While the solution is likely excessive for resolving the front-running issue as N has to be fairly small to be effective, it may be useful as a general fail-secure mechanism when connectivity is lost for an extended time period.

60. https://eprint.iacr.org/2015/261.pdf
61. In this case, Eve would not even necessarily need to offer higher fees since a block builder looking for maximum profits would see that they may legally execute Eve's transaction before the FI's, but not the other way around. So the fee would merely need to be high enough for inclusion.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Logic Obfuscation

In a scenario where we wish to perform an administrative update and want to conceal the nature of the update from passive observers until it is included in a block on-chain, we may consider obfuscating the logic or function calls that result in this state update.

For example, consider once again the scenario where an address is being added to the deny-list. We may not want this information to be public until it is committed to, and so we can try to discourage passive observers from suspecting that an administrative change is coming by embedding the state change in a function that is also used for other purposes, so that it blends into similar transactions.

The major downside to this approach is that a node executing transactions could determine, upon execution, that such a state update will occur. Also, astute observers reviewing the deployed smart contract code may be able to determine under what circumstances side effects will be triggered and act accordingly. As such, this solution does not solve the general front-running problem, but it may slow down information propagation for naïve adversaries. In general, it is not recommended to rely on obfuscation.

Addressing the Receipt of Unsolicited Tokens

Allow/Deny Listing on Transfers

At the smart contract layer, denylisting or allowlisting can be implemented to impose restrictions on who is allowed to interact with the smart contract and who can perform transfers to whom. This could be implemented in a way that allows a user to set up controls on its blockchain address such that it will only accept token transfers from allow-listed addresses, and that unsolicited token transfers will be rejected by the smart contract. The limitation is that this would only be specific to a given token and smart contract.

On Ethereum, smart contract accounts can help work around this by defining the account's behavior when native asset is received.[62] At this level, a user or FI could reject incoming payments from addresses they want no interaction with.

Incorporate Allow/Deny Listing into Token Standards

A user will only be able to use allow/deny-listing capabilities if provided by the token issuer. Broader provisioning of such capabilities is likely if they are incorporated into token standards. Even then, it would be a challenge, from both an effort and cost perspective, for users to apply controls across all the different tokens on different blockchain networks. Furthermore, "dusting" or "airdrops" are often performed to market newly issued tokens, which are new to users who would not have applied such controls for those tokens, and popular networks do not support tools like this for generic network assets.

Token standards of course do not help in cases where the asset does not comply with the standard, but for FIs, compliance with such a standard (even a standard as simple as having an allow/deny list) may make on-chain assets more available to FIs.

62. In Solidity, this could be done by defining the **receive()** function

On networks where there are no stable accounts (e.g. Bitcoin), or token-based networks in particular, this is likely not a problem in the first place. Even for large entities who need to publicly list addresses there is a clear distinction between assets and certain tokens received could clearly be marked as unwanted or "not to be spent".

Transferring Unsolicited Tokens from EOA and Smart Contract Addresses

The final remediation step after having received unsolicited tokens is likely transferring them, either returning to the original sender, sending to a designated address or burning them by sending to an invalid address. While we typically refer to a blockchain address or wallet as "holding" tokens, a more technically precise explanation is that the token's ledger, as a set of smart contracts, records balances attributed to a blockchain address. Correspondingly, transferring a token works by initiating the token's transfer function in its smart contract, which updates the balances recorded within it.

With an Externally Owned Account (EOA) address, the owner holds the private keys to the address, and can digitally sign a transaction that instructs the token's smart contract to transfer the tokens. With a Smart Contract address, no one holds private keys that correspond to the address, hence a different mechanism would be required.[63]

The Smart Contract would need to include an additional function that could be used to initiate the transfer function of the unsolicited token's smart contract. Since the function needs to be included at initial deployment and support yet-to-be-known tokens, it has to be sufficiently flexible to initiate a smart contract with an undetermined smart contract address, function name, and parameters.

Flexibility has to be balanced with security, and one possible path to this trade-off would be to restrict to a small set of functions and parameters as on well known contracts, rather than any arbitrary asset. Naïvely calling methods on unknown contracts is dangerous given that the code could do anything under the hood of a method call.

Also, controls will need to be imposed on such a function to limit its scope. Smart contracts may already hold tokens as part of their normal functioning. Examples include Automated Market Makers and Escrow contracts. A generic function that allows transfer of tokens held by the smart contract could have unintended consequences, such as allowing tokens to be transferred out of the escrow before conditions are met. The function must hence be designed with appropriate logic to prevent such actions.

Finally, there remains a question of whether such measures are even required. If a smart contract address has

63. It is theoretically possible for a collision, where the private keys corresponding to the smart contract address are found, but this is computationally infeasible. See: https://csrc.nist.gov/glossary/term/preimage_resistance.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

no ability to interact with the unsolicited tokens attributed to it, should the "owner" of the smart contract even be considered the "owner" of the tokens?[64] If not, then this is unnecessary.

Addressing Gas Fees Being Paid to Sanctioned Entities

Block Reward Address Introspection

In some networks, the recipient of the fee can be known at time of execution. As such, a smart contract could be written to revert if the transaction were to be processed by an undesirable block proposer (i.e. a sanctioned entity whose address is on a deny-list).

However, in all popular modern protocols, the block proposer may specify the recipient of the block reward to be any arbitrary blockchain wallet address. Hence, a sanctioned block proposer could specify the beneficiary address as a different address that they control, which is not on the deny-list. This, however, is a broader issue that is not specific to block rewards. Illicit actors can generate new wallet addresses that are deemed as "clean" to receive tokens in those addresses, before using them directly or consolidating them into their other existing addresses.

COINBASE opcode in EVM

In EVM-based systems, the COINBASE opcode allows the caller to identify which address receives the rewards from mining the block.[65]

It should be noted that gas fees will still be charged for processing when a transaction reverts, though this will often be significantly lower than if the entirety of the smart contract code for the transaction was processed. Hence, the signer of the transaction would still have paid a small fee, which may constitute a sanctions violation, even as they acted to minimize their exposure to the sanctioned entity.

64. Smart contract code cannot be modified once deployed, but can be executed. There is typically logic within the smart contract code that allows certain administrative actions to be performed, and the party that holds the credentials to perform such actions are typically viewed as the owner of the smart contract.
65. https://www.evm.codes/?fork=osaka#41

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Base Fees and Priority Fees in Ethereum

Since the London Hardfork and EIP-1559, the fee paid for a transaction on Ethereum has been split up into two parts: "Base fees" and "Priority fees".[66] If G gas is consumed by a transaction, the amount paid is **G * (BaseFee + PriorityFee)**.

The base fee is determined by how full recent Ethereum blocks have been. When the blocks are more full, or congestion is higher, then the base fee increases. The inverse is true when blocks are less full. Transactions which pay less than the base fee are not consensus valid. The base fee is also burned as a part of the transaction.

The priority fee is paid to the block proposer (or whoever is specified as the `COINBASE` of the block) as a tip. A user who wants to more strongly incentivize inclusion in a block will typically specify a higher priority fee.

Transactions can be legally included in a block while also paying no priority fees, but this is less compelling to a block proposer. Many Ethereum client implementations by default prioritize transactions by fees. In cases where an FI is concerned about paying a sanctioned entity, broadcasting a zero-priority-fee transaction really has no downside except for that the average block proposer may pass it by for rapid inclusion.

Two parameters that are included in EIP-1559 transactions include **maxFeePerGas** and **maxPriorityFeePerGas**. If **maxFeePerGas** is lower than the block base fee, the transaction will be invalid. The amount of priority fee paid for the transaction will also never be higher than **maxPriorityFeePerGas**.[67]

Prior to the London hardfork, the entirety of the fee was paid to the block proposer and none was burned. This change makes Ethereum more deflationary in times of high network traffic.

66. https://ethereum.org/developers/docs/gas/
67. https://eips.ethereum.org/EIPS/eip-1559

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Addressing at Blockchain Network and Network Governance Layers

Much of the issues have their root causes in these layers, and hence true solutions, rather than just workarounds, typically have to be implemented here. These layers are within the control of blockchain protocol developers, and their support would be critical to implementing such solutions. As decision making for public blockchains is often decentralized, achieving support to implement such solutions in existing blockchains may be challenging and long-drawn. As new blockchains emerge, some of which are specifically targeting financial use-cases, these solutions could be implemented to make them more suitable for use by regulated institutions and enterprises.

Addressing Front-Running and Censorship

Encrypted Mempools

A proposed mechanism for preventing front-running is requiring block creators to validate and commit to the transactions in the block, as well as their order, before the actual transactions are revealed. This technique is often referred to as encrypting the mempool. In theory, this idea may help increase decentralization in the network, since if block proposers cannot see what these transactions are, then they cannot as easily exploit them.

However, there is a great deal of additional complexity required to achieve this, and it would be a real challenge to deploy in the largest public ledgers today. In practice, this may require block creators and validators to use Trusted Execution Environments (TEEs), or techniques like threshold decryption. TEEs are challenging because they impose a specific hardware requirement and would likely require some mechanism to validate that TEEs are being used properly by the block creators. Also, they have been shown to not always be secure and trustworthy,[68] so a revocation scheme is likely required. The use of specialized hardware could also lead to consolidation and re-centralization risks, affecting the node diversity and hence resiliency of the network.

Another option is to enable encrypted mempools via threshold decryption. This is well understood cryptographically but does add another layer of complexity and requires the committee of validators evaluating a given block to be known ahead of time. Additional complexities arise, such as determining a schema to validate transactions without executing them, which becomes more complicated with greater expressivity. There are techniques that block builders could use to still try to speculate about MEV-strategies, though these are almost certain to be less effective than strategies that can see the full picture.[69]

Block Position Introspection

Another option may be to enable users to specify where in the block they would like to be executed and add the ability to invalidate the transaction if this positioning is not respected. If a transaction truly needs to be executed first, one approach is to enable smart contracts to specify that a given action is only valid if it occupies the first position in the block.

68. https://tee.fail/
69. https://a16zcrypto.com/posts/article/limits-encrypted-mempools/

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

This comes with the problem where all users target the most desirable spot, although it could be mitigated with well-designed incentive schemes.

Updating Incentive Criteria

The reason that MEV extraction plays such a large role in decentralized ledgers is because it is profitable to try and identify the most advantageous arrangement of transactions in a block, which may be beyond the user's influence. For example, a high transaction fee may not be nearly as high value as what can be derived from a smart contract exploit. This is not necessarily a unique problem to decentralized ledgers; however, the incentives that can be applied within the system are different from what is done in traditional systems, such as imposing a legal penalty for "bad behavior".

In the decentralized case, systems choose different incentives and disincentives for good and bad behavior. In the case of Bitcoin, the incentive to mine blocks correctly and keep the chain moving is the block reward and associated fees. The disincentive for incorrect behavior is the wasting of resources.[70] In Ethereum, the reward for behaving correctly is a proportion of stake, and the penalty for misbehaving is forfeiting a proportion of stake.

However, "correct" behavior is determined by the protocol, and the notion of good or bad behavior could be modified. A protocol designed to discourage MEV may modify its incentive system accordingly. These modifications in some systems must be defined by code updates or forks, and in others may be achievable at the social layer.

Congestion Reduction

On the concern of "high priority" or "urgent" transactions not being processed in a timely manner, one reasonable goal is to reduce network congestion. This could be through increasing maximum allowable blocksize and having each block contain more transactions, or reducing block time and having more blocks created in a given time period, both of which would increase total throughput of the network.

Since the London hardfork and EIP-1559, Ethereum has supported what is often referred to as "elastic blocksize". Before this time, blocks had a fixed gas limit. Post EIP-1559, blocks using up to 2x the network's target gas-per-block are valid, allowing room to handle additional load.

Moreover, the target gas of the network does change over time.[71] Blocks proposing gas limits within **1/1024** of the parent block's gas limit can be consensus valid, and so if a block proposer wants to nudge the gas limit of the network higher, they may propose a block with a higher gas limit.[72] In early 2025, the block gas limit increased from 30M to 36M.[73] At the social layer, the network encourages changes like this through signaling desire for a higher gas limit.[74] Today, the gas limit sits around 60M. If congestion increases further, encouragement can be made at the social layer to further increase the gas limit, however this can be slow (on the order of months) due to delays *at* the social layer.[75]

70. Ignoring network level disincentives like peer denylisting
71. See https://etherscan.io/chart/gaslimit for gas limit changes over time
72. See Eq. 54 of https://ethereum.github.io/yellowpaper/paper.pdf
73. The target gas in this case increased from 15M to 18M https://www.theblock.co/post/334951/30-of-ethereum-validators-signal-higher-block-gas-limit-data-shows
74. This can be seen on pages like https://gaslimit.pics/ highlighted by initiatives like: https://pumpthegas.org/
75. The fastest timeline that the 30M to 36M increase could have occurred is 38 minutes. In reality the network took several months to determine if it wants and then to adopt this increase. https://www.bitget.com/news/detail/12560604525694

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Increasing the speed at which blocks can be produced and agreed upon is a good mechanism for greater throughput, and many blockchains propose new mechanisms allowing for small blocktimes.[76] On some other networks, throughput is increased by modifying how blocks (or more simply, transactions) need to be linked, with the goal of allowing multiple concurrent transactions to be validated at one time, which also increases throughput.[77]

There is a class of congestion reduction techniques which shifts a proportion of transactions to a parallel network, which is often referred to as a Layer 2 (L2). In popular constructions, these networks periodically commit to the original network (L1),[78] or at least leverage the L1 for added security or dispute resolution.[79] This model is popular in the Ethereum ecosystem, and helps offload work to where there is cheap bandwidth while maintaining a connection to the Ethereum L1.

A pattern we tend to see in higher throughput systems is that the designs often have centralizing effects. The operators of the system may be a smaller set due to lofty hardware requirements, or if the system design elects to constrain the acceptable time to consensus then a node may need a lower latency connection to its peers. The consensus scheme may also choose a small set of trusted participants in the network, which can improve performance but creates a larger reliance on a smaller number of entities. Many L2 designs have this architecture, in which there is a single (or rotating) sequencer, though certain types of misbehaviors can be remedied on the L1.[80]

Transaction Prioritization

Even with higher throughput, there may still be periods of congestion when transaction traffic spikes. Some of this volume may be profitable, but interferes with more urgent or "natural" economic activity.[81] On platforms like Ethereum, this can often be trading activity that blocks out basic transfers. Transaction prioritization is the concept of enabling certain types of transactions to be processed first or at priority even in times of high congestion. For FIs, it is a compelling idea for systems to define administrative actions as high priority.

One option is to reserve or add space for specific types of transactions or data. For example, the first transaction of each block in Bitcoin is reserved for the "coinbase" transaction which is the payment to the miner. Similarly, Ethereum in the Dencun upgrade added the notion of "blobs" and allocated dedicated space in the network within the consensus layer.[82] This allowed for rollup data, which tends to be larger than standard transaction data, to be stored separately and differently from calldata handled by the execution layer, which has a much higher unit cost.

Given the use case for blobs and their temporary nature, the approach is appropriate and the trade-offs are acceptable. In cases where the transaction data is similarly long-lived, it must reside within the execution layer, which would require partitioning of the space in the execution layer for different types of transactions, i.e. "priority" and "non-priority" space. Therefore, one challenge is for the protocol to determine what constitutes a "priority" transaction for the network. If too much space is reserved for priority activity, then it is wasted when that activity is not occurring at a high rate.

76. This includes parallel execution (https://www.dci.mit.edu/projects/parallelized-architecture-for-scalably-executing-smart-contracts-parsec, https://docs.monad.xyz/monad-arch/execution/parallel-execution), as well as new data models (e.g. https://docs.sui.io/develop/sui-architecture/object-model) and cryptographic protocols (e.g. https://www.guibibeau.com/blog/proof-of-history-explained).
77. DAG-based systems are an example.
78. See "Rollups": https://vitalik.eth.limo/general/2021/01/05/rollup.html
79. This is the technique used in Lightning https://lightning.network/lightning-network-paper.pdf
80. https://docs.optimism.io/op-stack/protocol/components#sequencing
81. https://blockworks.com/news/ordinals-impact-on-bitcoin-network
82. https://a16zcrypto.com/posts/article/understanding-dencun-upgrade-protodanksharding-surge-merge/

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Some in the Ethereum community are considering inclusion lists to indicate what transactions ought to be included in upcoming blocks. EIP-7805, referring to Fork-choice enforced Inclusion Lists, or FOCIL, proposes a technique in which some validators form an inclusion list committee, and each member of the committee gossips an inclusion list to the block proposer and attesters of the next slot (and presumably any other interested listener). If adopted, attesters would only vote for the block if it contains transactions from all inclusion lists for the slot.[83] This protocol is designed primarily to guarantee timely transaction inclusion.

For an FI, getting a transaction included in very short order (e.g. the next slot) only requires contact with a member of the inclusion list committee for a slot. This is likely to reduce the expected wait time for transaction inclusion on-chain. Different validators may create unique inclusion list criteria, and some may elect to prioritize administrative operations. Once enough validators adopt similar criteria, FIs may enjoy a greater sense of safety on-chain. A similar protocol can be achieved without requiring consensus amongst a committee, if one is willing to introduce TEEs to the system. However, this negatively affects node diversity.

Addressing Receipt of Unsolicited Tokens

Global Allow/Deny List

A feature that could be included in a protocol is the ability to define a global allowlist or denylist for an account.[84] This would allow a user to specify if they wish to not receive funds from a certain set of addresses or for certain tokens, or similarly if they would only like to receive funds from a defined set. This would prevent dusting attacks in the native asset of the network, and if this list were respected in the execution layer, it would avoid this situation for second-order assets as well.

However, this would contribute to a larger global state and would likely open up censorship avenues that are currently avoided, which would not be the case if this is only applied at the smart contract layer. Different approaches to this exist on different networks. For example, on Algorand, assets (Algorand Standard Assets) must be opted into before they may be received, which works as the asset level but not necessarily on a per-sender level.[85]

Addressing Gas Fees Being Paid to Sanctioned Entities

Managing Block Rewards

If a transacting party can, during execution, determine who the beneficiary of the block reward is and check against a sanctions list, they could avoid cases in which the transaction directly pays fees to sanctioned entities. However, this requires a greater degree of introspection than is currently available in several networks, such as Bitcoin. Ethereum enables this through the COINBASE opcode. This would also be possible on Solana if the leader schedule is known; however, this information would have to be part of the transaction itself, and the caller may not be privy to the sanctions list.

83. https://eips.ethereum.org/EIPS/eip-7805, this description also taken in large part from the abstract of the following. See this page for a nice abstract and discussion of FOCIL https://ethereum-magicians.org/t/eip-7805-committee-based-fork-choice-enforced-inclusion-lists-focil/21578
84. This would be more challenging to implement in a non-account based system.
85. https://dev.algorand.co/concepts/assets/asset-operations/#opting-in-and-out-of-assets

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

In the case of reverting transactions on Ethereum, some fee may still go to the recipient of the block reward. This makes the usage of COINBASE to detect the recipient only partially effective, as priority fees are still paid out after execution reverts. It may be possible for the rules and logic to be modified such that fees will not be paid to the beneficiary in the case of a revert. However this will likely require further analysis and validation of the crypto-economics. For example, even if fees are not paid to beneficiaries, they might still be incurred (i.e. burned) to manage usage and network congestion.

Validator Participation Rules

Another option to prevent interacting with sanctioned entities is to design a protocol that enforces that the validator set is known or has some sort of onboarding process. Some protocols such as Sui have a pre-defined number of validators, and therefore they must be known. Ethereum L2s like Base are run by known entities, which has a similar impact. Ethereum does have a known number of validators, but entry into the validator set is public and permissionless, so “sanctioned entities” cannot be barred from participating.

The primary concern with this technique is the centralizing effect it has on the network. In the case of Ethereum, this may not be an acceptable tradeoff. However, alternative networks targeting FIs in particular may adopt this strategy.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Addressing at Regulation Layer

Regulatory clarity plays a critical role in enabling the use of public blockchains by FIs. While technical solutions at the application, smart contract, and protocol layers can mitigate many of the issues identified in this paper, some challenges cannot be fully resolved through technical means alone. This is particularly the case where the root issues stem from the inherent design of public blockchains, such as the permissionless nature of validator participation or the inability to prevent unsolicited token transfers, and where the residual risks carry legal or compliance implications.

In these cases, regulatory guidance is essential to establish clear expectations, define proportionate obligations, and provide legal certainty for FIs operating in good faith. Regulations are within the control of central banks, financial regulators, and policymakers, and their support would be critical to creating an environment in which FIs can use public blockchains with confidence. As the financial industry's use of public blockchains grows, regulations that are informed by the technical realities of these networks will be important in striking the right balance between managing risks and enabling innovation.

It is also important that regulations evolve collectively and cohesively across jurisdictions. Public blockchains operate across a global, borderless environment, where participants, infrastructure, and activities are distributed across multiple jurisdictions. Regulatory frameworks governing such activities are currently evolving at different paces across jurisdictions, with some authorities providing greater clarity on the treatment of blockchain-based activities, while others are still in earlier stages of development.

These differences in regulatory approaches and jurisdictional boundaries may give rise to variations in how activities are governed and supervised, which, in certain cases, may create incentives for participants to operate from jurisdictions with less defined or less stringent regulatory frameworks. As these participants transact on a shared network, their actions or behaviours may create negative spillover effects for other participants. This is not unique to public blockchains, but may be more pronounced given the open and globally accessible nature of these systems.

Finally, as highlighted earlier, financial regulations are primarily applied directly to FIs. However, they often include provisions relating to third-party out-sourcing and technology risk management. As such, while protocol developers and operators are not directly subject to financial regulations, FIs' use of the blockchain networks are often subject to such regulations, which they would in-turn apply as requirements for blockchain networks that they use. Regulations are therefore relevant for all ecosystem stakeholders, even if they are not directly regulated entities.

Accordingly, the observations in this report aim to highlight areas where continued coordination, clarity, and alignment across jurisdictions may support more consistent outcomes, particularly as institutional participation scales. The report is intended to contribute constructively to ongoing dialogue among regulators, financial institutions, and ecosystem participants.

Recognize Front-Running of Financial Transactions on Blockchain as Market Manipulation

In traditional financial markets, front-running is widely recognized as a form of market manipulation and is subject to legal penalties. However, it is currently unclear whether front-running financial transactions on public blockchains constitutes market manipulation under existing laws and further regulatory guidance may be needed to clarify how such conduct will be treated. This ambiguity creates an environment in which front-running and MEV extraction can occur without legal consequence, even when the transactions involved are financial in nature.

Regulators could address this gap by providing clear guidelines on treatment of front-running of financial

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

transactions on public blockchains, and where required, extend existing legal frameworks or establish new ones to cover such activity. Clear legal repercussions would serve as a deterrent and complement the technical solutions proposed at other layers, creating a more comprehensive approach to addressing the issue. Nonetheless, given that these networks are global and decentralized, this may be challenging to attempt.[86]

Regulatory Clarity on Receiving Unsolicited Tokens

FIs currently face significant uncertainty regarding their obligations when they receive unsolicited tokens, particularly when those tokens originate from unknown or potentially sanctioned entities. For example, after Tornado Cash was sanctioned, many accounts were "dusted" with tainted funds, and though addresses that were dusted largely did not have criminal charges enforced against them, there is concern that a less friendly executor may take a different route in the future.[87] In the absence of clear regulatory guidance, institutions tend to adopt an extremely conservative approach, treating the receipt of unsolicited tokens similarly to receiving funds in a bank account and undertaking costly due diligence and return processes.

As discussed earlier in this paper, this creates a disproportionate operational burden, especially given that many unsolicited tokens are promotional in nature and have negligible or no value. Regulators could provide much-needed clarity on several fronts: whether the mere receipt of unsolicited tokens constitutes a regulatory concern, and what obligations, if any, FIs have to return or dispose of such tokens, and what constitutes an acceptable and proportionate response. This is particularly important, as even merely interacting with an unwanted asset or counterparty could introduce significant operational overhead. Such guidance would enable FIs to adopt more practical and risk-proportionate approaches, reducing unnecessary operational costs while still maintaining compliance with sanctions and anti-money laundering requirements.

Regulatory Clarity on Sanctions Compliance and Safe Harbor Provisions

On most public blockchains, FIs cannot fully control who processes their transactions or receives the associated gas fees. As discussed in this paper, even with the best available technical mitigations, there remains a residual risk that gas fees may be paid to sanctioned entities, particularly on permissionless layer 1 networks where validators can participate without onboarding or identity verification. Given that sanctions are a matter of strict liability, this residual risk raises concerns for FIs, often leading to cautious approaches that may limit their ability to use public blockchains altogether.

Regulators could address this by providing clearer requirements on what constitutes reasonable efforts to prevent transacting with sanctioned entities in the context of public blockchains, recognizing that the technical architecture of these networks differs fundamentally from traditional financial infrastructure. Additionally, establishing safe harbor provisions for FIs that have taken demonstrable and reasonable steps to avoid such transactions, but where inadvertent interaction nonetheless occurs due to the inherent design of public blockchains, would provide the legal certainty needed for institutions to participate more confidently in public blockchain ecosystems.

86. After the OFAC-sanctioning of Tornado Cash for example, several builders continued to include sanctioned transactions https://blog.toniwahrstaetter.com/builder-censorship.html. However remediation of market manipulation likely manifests very differently than enforcing sanctions.
87. https://decrypt.co/109635/treasury-wont-punish-dusted-celebs-will-allow-users-to-recover-funds-from-tornado-cash

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Layer 2 Networks

Many of the issues highlighted earlier can be more readily addressed when there is a centralized operator. Layer 2 networks inherently solve many of these issues as they are often operated and maintained by a single, defined entity. This centralization allows for clearer accountability and operational control. It allows FIs to establish and maintain strong relationships with the L2 operator, resulting in legal and commercial arrangements that are broadly comparable to those governing traditional infrastructure providers.

For FIs, this model may be a pragmatic entry point into decentralized finance, and represents a meaningful shift from the status quo. Nevertheless, it remains one step removed from operating within a fully decentralized environment. While there are numerous open questions in the design and development of L2 solutions, they are not explored in depth in this report. In many cases, the considerations involved are conceptually similar to those encountered in traditional financial systems. Additionally, the diversity of L2 architectures currently in development and deployment makes it challenging to cover them comprehensively within the scope of this report.

At the same time, many L2 networks are actively exploring approaches to increase decentralization.[88] For example, rollup-based designs typically incorporate mechanisms that allow users to mitigate L2 censorship or operator unavailability. In such cases, users can unilaterally withdraw funds or force transaction inclusion through interactions with the underlying Layer 1 network.[89] However, this reliance on L1 reintroduces many of the challenges discussed in this report, even if they are partially abstracted at the L2 level.

Further analysis of the evolution of L2 architectures, particularly in relation to decentralization, is a valuable avenue for future work. L2 networks may also emerge as leading candidates for adopting features required by FIs as they progress along the decentralization spectrum. In addition, efficient interoperability between L2 networks presents a significant and active challenge.[90] Given that one of the key value propositions of decentralized systems for FIs is the ability to move liquidity more efficiently and achieve greater interoperability, this area warrants focused attention.

88. For example, the L2 Aztec has a decentralized sequencer https://aztec.network/
89. https://hackernoon.com/blockchains-need-escape-hatches-in-layer-2 https://thebitcoinmanual.com/articles/unilateral-exit/
90. Shared infrastructure for example is a good starting point https://www.alchemy.com/overviews/optimism-superchain-op-stack-guide

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Next Steps

Public blockchains present a meaningful opportunity for FIs, but their adoption at scale will depend on the ecosystem's ability to address structural risks and align with regulatory expectations. As this report highlights, many of the most critical challenges remain unresolved.

These challenges require coordinated action across stakeholders, reflecting differences in incentives, design philosophies and priorities, and accountability across the public blockchain ecosystem. Near-term priorities for FIs include continuing to adopt practical mitigations at the application and smart contract layers. Technical experiments were conducted as part of this research and source code to software developed will be published to facilitate further developments of the mitigation techniques.

In the longer term, there is a need for the ecosystem to work together to ensure that FIs' needs are supported in their use of public blockchains. Greater engagement is needed between FIs and protocol and infrastructure providers to articulate institution-specific requirements, and to influence design choices and roadmap development.

Regulatory clarity will be a key enabler of progress. Policymakers have an important role to play in defining acceptable practices, setting proportionate expectations, and providing legal certainty in areas where technical solutions alone are insufficient. Ongoing dialogue between regulators and industry will be essential.

Looking ahead, meaningful progress will depend on sustained collaboration across FIs, blockchain protocol developers and regulators. No single stakeholder group can resolve these challenges independently. Bridging the gap between decentralized infrastructure and institutional requirements will require coordinated advances in technology design, market practices and policy frameworks.

It is our hope that this report has contributed to that effort by clarifying key challenges and identifying areas for further collaboration, supporting a more secure, compliant and resilient integration of public blockchain infrastructure into the financial system.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative

Disclaimer

The information in this report, or on which this report is based, has been obtained from sources that the authors believe to be reliable and accurate. However, such information has not been independently verified and no representation or warranty, express or implied, is made as to the accuracy or completeness of any information obtained from third parties.

The information and conclusions are provided as at the date of this report and are subject to change without notice, and MIT and J.P. Morgan undertake no obligation to update or revise any information or conclusions contained herein, whether as a result of new information, future events, or otherwise. The information and conclusions provided in this report take no account of any relevant persons' individual circumstances, should not be taken as specific advice on the merits of any investment decision, product, or service and should not be deemed to be a reasonably sufficient basis upon which to make an investment decision or undertake any product or service.

This report is not intended to provide and should not be relied on for, accounting, legal, or tax advice, or investment recommendations. Please consult your own tax, legal, accounting, or investment advisor concerning such matters. MIT, J.P. Morgan and their respective affiliates accept no liability for any loss arising from any action taken or refrained from as a result of information and conclusions contained in this report or any reports or sources of information referred to herein, or for any consequential, special or similar damages even if advised of the possibility of such damages. This report has been provided solely for information purposes and does not constitute a recommendation, advice, or an offer or solicitation to buy or sell any securities or financial instruments or of any product or service. It should not be so construed, nor should it or any part of it form the basis of, or be relied on in connection with, any contract or commitment whatsoever.

Further, this report shall not be considered advice on the merits of acquiring or disposing of any particular investment or as an invitation or inducement to engage in any investment activity or other product or service. By accepting this report, you agree to be bound by the foregoing limitations. J.P. Morgan is a marketing name for the payments business of JPMorgan Chase Bank, N.A. and its affiliates worldwide. JPMorgan Chase Bank, N.A., is organized under the laws of USA with limited liability.

 © 2026 Kinexys by J.P. Morgan and MIT Digital Currency Initiative