



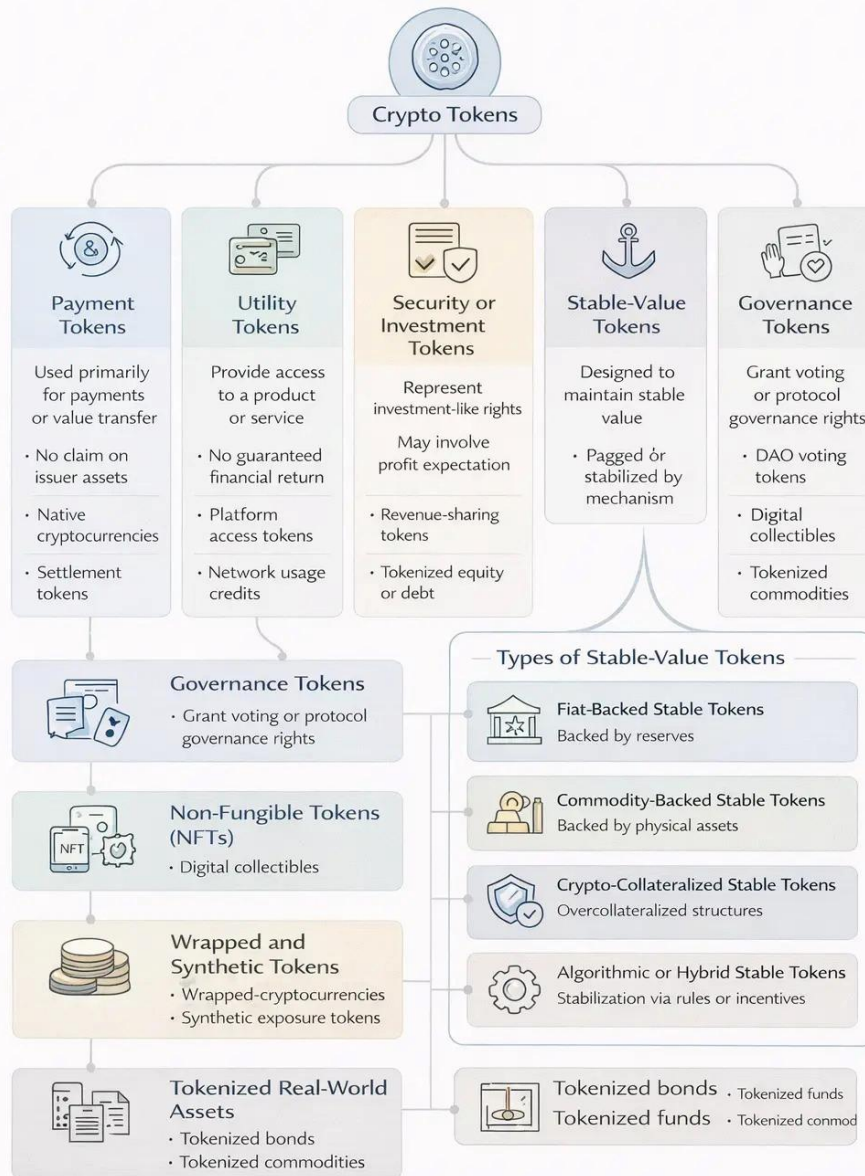
Crypto Regulatory Landscape

Abstract

This paper maps the current regulatory environment for crypto-assets and the services built around them. It treats crypto regulation as a boundary-setting exercise before it is a rule-writing exercise: what counts as a crypto-asset, which activities count as regulated services, and how token types and market functions are sorted into legal buckets. The analysis begins with scope and definitions, then builds a practical asset taxonomy that can be used across regimes and across products that blur older categories.



Crypto Token Taxonomy Used by Regulators



All rights reserved by the Blockchain Council.

The paper then frames regulation through objectives and a risk model. It argues that, despite differences in statutes and agency mandates,

most jurisdictions converge on a shared logic: identify the economic function (issuance, trading, custody, transfer, lending, staking, stable-value issuance), map that function to a harm model (who is harmed, how, and how quickly), and apply outcomes comparable to traditional finance where risks are comparable—while adjusting for crypto-specific drivers such as key custody concentration, remote cross-border access, code change risk, and irreversible transfers.

The paper places domestic regimes within a global baseline layer shaped by international standard setters. It explains how AML/CFT standards, market conduct expectations, and stability-focused guidance create a common floor that pulls jurisdictions toward similar outcomes through evaluation, peer pressure, and cooperation needs, even where national rules remain different.

Finally, the paper draws practical implications for policy design and supervision. It sets out a repeatable method for classification, a supervision approach that scales with risk and reach, and a coordination agenda that treats cross-border enforcement, data sharing, and group-wide oversight as necessary conditions for effectiveness. It concludes that the core challenge is no longer drafting definitions alone, but making those definitions work in markets that move fast, cross borders easily, and evolve through code, governance, and business structure.

Keywords

Crypto-assets; virtual assets; regulation; classification; asset taxonomy; market conduct; investor protection; financial stability; AML/CFT; sanctions; travel rule; custody; trading venues; stable-value tokens; decentralized finance; cross-border supervision; supervisory colleges; enforcement.



Introduction

Crypto-assets entered financial policy as a mix of technology, markets, and ideology. They arrived as software, but quickly grew into trading venues, custody businesses, payment instruments, and credit products. They also arrived with a vocabulary that was not built for legal drafting: “tokens,” “protocols,” “wallets,” “stablecoins,” “governance,” “staking,” “bridges,” and “decentralization.” Much of the early debate assumed that crypto regulation would turn on one question—whether a token is a security. That question matters, but it is only one part of the regulatory problem.

Regulators have learned, often the hard way, that the biggest fights begin earlier: at the boundary stage. What is being regulated? A token? A firm? A market function? A payment promise? A custody arrangement? A contract deployed on a public chain? Or a group of people who control upgrades and fees? The answers decide everything that follows: which agency leads, which license is required, what disclosure is demanded, what controls are tested, and what enforcement tools can be used.

This paper treats scope, definitions, and classification as the organizing spine of crypto regulation. It does so for a simple reason. In crypto markets, legal consequences flow from classification, and classification depends on economic substance. A token’s marketing label is easy to change. The economic function is harder to disguise for long. A trading venue that matches orders for the public is a trading venue, even if it calls itself a “platform.” A service that holds customer keys is custody, even if it is wrapped in app design. A stable-value arrangement that invites money-like use creates money-like expectations, even if the issuer claims it is “just a token.”



The regulatory environment is complicated by overlap. A single crypto business can touch securities law, payments law, consumer protection law, prudential oversight, and AML/CFT rules at the same time. A single transaction can look like investment activity, payment activity, and cross-border value transfer all at once. A single product can shift category when its use changes: a token sold as access can become a traded instrument; a token built for settlement can become collateral for leverage; a governance token can start to look like a claim on fees.

These overlaps are not noise. They are features of how crypto markets are built. Crypto markets have grown through a dense mix of vertical combination (multiple functions housed in one group), global remote service provision (customers served across borders with little physical presence), and technology-driven fragility (security failures, contract bugs, admin key misuse, and operational outages). These features do not invent new forms of harm. They change the speed of harm, the difficulty of recovery, and the places where supervision must focus.

The central thesis of this paper is that most modern crypto regimes—whether they present as “crypto-specific” statutes or as extensions of existing law—are converging on a functional approach with three moving parts.

First, they sort tokens and activities into categories that carry legal consequences. Those categories differ in name and in statutory detail, but the recurring buckets are familiar: payment-style tokens, access-style tokens, security-like tokens, stable-value arrangements, and service-provider roles such as exchange operation, brokerage, custody, and transfer.



Second, they anchor regulation in objectives and a harm model. When agencies speak in different languages, they still tend to pursue the same outcomes: fair and orderly markets, protection of users from fraud and abuse, safeguarding of client assets, containment of spillovers, and reduction of laundering and sanctions risk.

Third, they treat supervision and enforcement as the real test. A rulebook that cannot be enforced against offshore service provision does not create a meaningful perimeter. A licensing regime without data collection and inspections does not change behavior. A transfer transparency rule that cannot be implemented between counterparties does not produce usable traceability. The paper therefore gives weight to how regimes are implemented, not only to how they are drafted.

The paper is written for regulators, compliance leaders, legal analysts, and researchers who need a coherent framework for comparing regimes and for reasoning about new products. The goal is not to argue for one country's model over another's. The goal is to provide a reusable method: a way to classify assets and services, map them to risk and objectives, and anticipate the likely regulatory response.

The analysis proceeds in four steps.

Chapter 1 sets the foundation: scope, definitions, and asset taxonomy. It explains why perimeter decisions decide regulator choice, licensing tracks, and disclosure burdens. It provides a practical vocabulary and a cross-regime taxonomy that treats labels as secondary to economic substance.

Chapter 2 builds the objectives and risk model that explains why regimes, despite surface differences, often end up imposing similar controls. It separates market integrity goals, user protection goals, stability goals, financial crime goals, and policy goals around



competition and responsible market development. It then describes how risk-based supervision scales obligations and supervisory intensity to the parts of the market where harm can be largest.

Chapter 3 explains the global baseline layer: the way international standards shape domestic rules and push convergence through evaluation and cooperation needs. It also explains why gaps persist at the frontier and why cross-border coordination remains a weak point even as domestic rulebooks mature.

The later chapters (not reproduced here in full) apply this foundation to the major regulatory stress points that keep returning across jurisdictions: licensing of crypto service providers; market structure and conflicts in vertically combined firms; custody, segregation, and insolvency treatment of client assets; token issuance disclosure and marketing controls; stable-value arrangements and reserve governance; credit, yield, and staking products; decentralized finance and the search for control points; and enforcement strategy in a borderless market.

Throughout, the paper uses a simple discipline. It distinguishes between the asset (the token or claim), the activity (what service is provided), and the arrangement (how rights, controls, and risks are structured). That discipline prevents two common mistakes: treating all tokens as one product category, and treating the absence of a traditional intermediary label as proof that no intermediary exists.

The stakes are practical. Crypto markets can support useful payment and settlement functions, faster access to markets, and new ways to coordinate software-based services. They can also produce rapid consumer harm through fraud, custody failure, and high-risk leveraged trading. They can create spillovers through stable-value



runs, forced selling, and links to core finance. They can enable laundering and sanctions evasion. Regulators are therefore not deciding whether crypto is “good” or “bad.” They are deciding how to prevent the worst outcomes while allowing activity that can be conducted safely.

A serious research paper must also acknowledge what crypto regulation cannot do. Regulation cannot remove volatility from speculative assets. It cannot prevent every self-custody mistake. It cannot eliminate all protocol vulnerabilities. It cannot make cross-border enforcement easy. What regulation can do is make intermediaries accountable where they exist, make product claims testable, make custody safer, make market abuse harder to run, and make financial crime controls real at the main gateways.

Those are the measures by which crypto regimes should be judged. The rest of this paper provides the concepts and structure needed to make that judgment in a consistent way.

1) Scope, definitions, and asset taxonomy

1. Why scope and definitions sit at the center of crypto regulation

In crypto regulation, the hardest disputes usually come before anyone gets to “rules.” They happen at the boundary-setting stage: what counts as a crypto-asset, who counts as a service provider, and which tokens fall into which legal bucket. Get those boundaries wrong and the rest of the framework wobbles: the wrong agency claims authority, the wrong license is applied (or none at all), and disclosure duties either land on the wrong party or never attach.

Scope and definitions decide three practical outcomes.



First, they decide which public body has primary responsibility. In most jurisdictions that means a mix of authorities with overlapping mandates: a securities regulator, a payments or central bank function, a banking supervisor, a commodities or derivatives regulator, a consumer protection agency, and the financial intelligence unit for anti-money laundering and counter-terrorist financing. Crypto routinely crosses those lines. A token can look like a security when marketed, like a payment instrument when used, and like a commodity when traded as a spot asset in a market. The “answer” is rarely singular; the framework must still pick a lead regime for each activity and product.

Second, scope and definitions decide which authorization track applies. That matters because licensing is the system’s gate. A firm might need a securities intermediary license, a payments institution license, an e-money authorization, a crypto-asset service provider authorization, a registration tied to anti-money laundering duties, or some combination. The compliance burden, supervisory intensity, and permissible business model shift with that choice.

Third, scope and definitions decide which disclosure duties turn on. Depending on classification, the law may require prospectus-style disclosure, a token “white paper,” reserve and redemption disclosure for stablecoins, standardized risk warnings, product governance rules, and market integrity controls.

That is why many regimes are built around classification. The legal consequences follow the classification. The classification should follow economic substance: what rights the token gives, what promises sit behind it, how it is issued, how it is marketed, and what activities are performed around it.



Two regulatory styles appear again and again.

One is a perimeter-based approach. The regulator asks whether the token or activity already fits inside an existing category: a security, a derivative, e-money, deposit-taking, collective investment, or a payments instrument. If yes, existing law applies. If not, the activity may sit outside the perimeter until lawmakers extend it, or a crypto-specific regime is created for the remainder.

The other is an activity-based approach. The regulator focuses on what the intermediary does: custody, exchange, brokerage, operation of a trading venue, order execution, transfer, staking, lending, or issuance support. If the activity resembles an activity regulated in traditional finance, the regulator tries to reach similar protective results, even if the technology differs. The common slogan is simple: same activity, same risk, same regulatory result.

This section builds the definitional and taxonomy foundation for later analysis of authorization, conduct rules, disclosure, stablecoins, decentralized systems, and enforcement. It does not try to “solve” classification debates in the abstract. Instead, it sets a method for doing the work consistently.

2. What this section covers and what it does not

This section sets the scope for “crypto regulation” as used throughout the larger paper. It focuses on two things:

1. The objects of regulation: crypto-assets and related instruments.
2. The subjects of regulation: service providers and other actors who perform regulated functions around those assets.



It uses “crypto-asset” as a broad descriptive label for tokens and token-like instruments recorded on distributed ledgers or similar systems, including stable-value tokens, access tokens, governance tokens, and tokenized claims on off-chain assets. In later chapters the paper will separate (a) the token as an object, from (b) the activity performed around it, because modern crypto markets often create the same risk through different combinations of tokens and services.

This section does not attempt to rank jurisdictions or compare political choices. It also does not treat “crypto” as a single product category. The point is the opposite: crypto is a collection of design choices that can map to very different legal categories.

3. Core terms and why the labels matter

Regulatory texts rarely use one universal term. They use overlapping labels that are close in everyday language but not identical in legal effect. The differences are not cosmetic. Each label is tied to a trigger in a statute or rulebook.

3.1 Crypto-asset, virtual asset, digital asset: similar words, different triggers

“Crypto-asset” is widely used in European policy and is central to the European Union’s crypto-specific regime. It is built around a broad notion: a digital representation of value or a right that can be transferred and stored electronically using distributed ledger technology or similar technology. The definition is intentionally wide. The narrowing happens in exclusions and subcategories.

“Virtual asset” is used by the global anti-money laundering standard setter for the purpose of attaching anti-money laundering and counter-terrorist financing duties. The term is meant to be broad enough to catch new instruments that can move and be exchanged, even if they



do not fit neatly into classic categories. This is not a securities-law concept. It is a financial crime concept.

“Digital asset” is often used in public debate, especially in the United States, as a catchall phrase. Sometimes it appears in guidance or agency statements rather than a single statute. In practice, “digital asset” often operates as a signpost rather than a classification. Regulators then apply older definitions (investment contract, commodity, money transmission, banking activity) to the facts.

These terms lead to different consequences.

Under an EU crypto regime, “crypto-asset” may pull an issuer and service provider into a market access and disclosure system, unless the instrument is already a regulated financial instrument, in which case the financial instrument regime applies.

Under a global anti-money laundering standard, “virtual asset” is used to bring service providers into customer due diligence, recordkeeping, suspicious activity reporting, and information-sharing duties for transfers.

Under a “digital asset” umbrella, older categories are applied case by case, often producing overlapping authority and disputes over who regulates what.

A consistent research approach treats these labels as entry points, not conclusions. The conclusion depends on how the token works and how it is used.

3.2 Token, coin, and the false precision of market vocabulary

Market actors often use “coin” to suggest decentralization and “token” to suggest issuance on top of an existing platform. Regulators rarely rely on that distinction because it can be gamed and it does not



track legal rights. A so-called “coin” can still have identifiable promoters, controlled supply schedules, and ongoing managerial decisions. A “token” can be widely dispersed with no meaningful issuer control. A legal taxonomy should therefore focus on rights, promises, and functions, not marketing terms.

3.3 On-chain record vs legal right

A token is an on-chain record. A legal claim is a right recognized by law and enforceable against someone. Sometimes the token and the claim align neatly: tokenized shares where the token is legally treated as the share itself, or at least as a record of entitlement. Often they do not. Many tokens grant access to software, signal voting preferences, or serve as receipts in a system where the real enforceable right sits elsewhere.

This distinction is central to classification. A token can trade like a financial instrument without clearly granting enforceable rights. Regulators then face a choice: regulate the token as a financial product because of how it is sold and traded, regulate the intermediary activities around it, or do both.

4. The service-provider concepts that carry most compliance weight

Across jurisdictions, the greatest compliance burden typically lands on intermediaries rather than token issuers. That is not an accident. Most consumer harm, market abuse, and financial crime risk concentrate at the points where retail users enter, exit, or store value: exchanges, brokers, custodians, wallet providers, and issuers of stable-value tokens used as a settlement asset.

Two acronyms dominate cross-border discussions.



“VASP” (virtual asset service provider) is used in the anti-money laundering context. It is aimed at any person or firm that, as a business, conducts certain covered activities involving virtual assets, including exchange, transfer, and custody or administration.

“CASP” (crypto-asset service provider) is used in the EU crypto market regime and in many policy discussions as a label for entities providing crypto services professionally. The list of services is typically specified: custody, exchange, execution, trading venue operation, order transmission, advice, portfolio management, and transfers.

The labels differ, but the functional target is similar: gatekeepers to customer funds and market access.

The research implication is straightforward. Any serious regulatory analysis must keep separate:

- the classification of the token as an object, and
- the classification of the activities performed around that token.

A token might sit outside securities law, while a platform listing it still faces strict duties on custody, conflicts, and anti-money laundering controls.

5. What counts as a crypto-asset or virtual asset in reference frameworks

A useful taxonomy needs a perimeter: a working answer to what is “in scope” and what is not.



5.1 The EU crypto perimeter and the financial-instrument carve-out

The EU's crypto regime provides a broad definition of crypto-assets and then builds tailored rules for key subtypes. It also preserves a major carve-out: where a crypto-asset qualifies as a regulated financial instrument under existing financial services law, that existing regime applies instead of the crypto regime. This is not a minor technicality. It is the hinge that determines whether the primary rulebook is “crypto” or “securities.”

Within the EU crypto regime, the major subtypes are defined in ways that track common market designs:

- stable-value tokens referencing one official currency (often treated like e-money),
- stable-value tokens referencing a basket or other assets, and
- access tokens meant to provide entry to goods or services supplied by the issuer and accepted only by that issuer.

This structure signals two regulatory judgments.

First, stable-value tokens require their own rules because they can function as payment and settlement assets. The main risks are not only investor loss from price swings but also safeguarding, redemption, reserve quality, operational resilience, and the risk that a stable-value token becomes a widely used rail.



Second, access tokens are treated differently from tokens sold as investments. The legal system tries to avoid forcing full securities-style disclosure on software access rights, while also preventing issuers from abusing the “access” label to avoid investor protection.

5.2 The global anti-money laundering perimeter

The anti-money laundering framework uses “virtual asset” to cast a wide net over transferable digital value used for payment or investment. It is defined broadly but excludes assets already captured as fiat or traditional financial assets under other parts of the anti-money laundering framework. The goal is coverage, not reclassification.

The associated “VASP” definition captures covered business activities:

- exchange between virtual assets and fiat currencies,
- exchange between virtual assets,
- transfer of virtual assets,
- safekeeping or administration of virtual assets or instruments enabling control over virtual assets, and
- participation in and provision of financial services related to an issuer’s offer or sale of a virtual asset.

Three interpretive problems show up repeatedly in this perimeter.



First is control. Who has the ability to move customer assets? In crypto, control is often private-key control, but it can also be practical control through account systems, smart contracts, or administrative privileges.

Second is the “as a business” test. Many actors participate in crypto networks: developers, validators, liquidity providers, governance voters, and people running front ends. The framework generally aims to capture professional services to others, not ordinary use.

Third is the boundary with decentralized systems. If a system has no identifiable intermediary, regulators struggle to attach duties, and sometimes respond by targeting the on-ramps, off-ramps, or any entity that provides user-facing access.

5.3 A market integrity perimeter for investor protection

Investor protection and market integrity frameworks often avoid prescribing one rigid global taxonomy. Instead they use economic substance and functional analysis. The central idea is that many crypto markets recreate familiar problems: conflicts of interest, weak custody practices, thin disclosure, price manipulation, and abusive sales tactics. The fact that a token is not a share or bond does not remove those risks.

A functional perimeter typically treats crypto market intermediaries as the focal point: trading venues, brokers, dealers, custodians, and entities offering lending or staking services to the public. The same basic questions appear: who holds customer assets, who sets trading rules, who has access to non-public order flow, who makes markets on their own venue, and what protections exist against abuse.

6. Why classification drives regulator choice, authorization, and disclosure



Taxonomy is not an academic sorting exercise. It is a routing system for legal obligations and supervision.

6.1 Classification allocates lead authority

A single token can touch multiple mandates. A classification approach must therefore be explicit about whether the goal is to find:

- one dominant classification for all purposes, or
- multiple classifications depending on the activity and risk.

Many jurisdictions effectively do the second, even if they describe the result as one answer. A token might be “not a security” as an object, yet trading venues, brokers, and custodians dealing with that token might still face a full authorization regime. A stable-value token might be treated as a payments instrument for issuance and redemption, while derivatives based on it fall under derivatives supervision.

Because authority overlaps, the classification work must also explain what counts as:

- an issuance activity,
- a trading activity,
- a custody activity, and



- a transfer activity.

Each may attach to a different regulator or a different part of the same regulator's remit.

6.2 Classification decides the authorization pathway

Authorization is where the framework becomes real. It controls market entry and sets baseline conduct rules.

A security-like token generally triggers securities offering rules and trading rules: disclosure tied to public offerings, requirements for intermediaries, market abuse controls, custody and client asset protections, and often trading venue obligations.

A stable-value token used for payments generally triggers safeguarding, reserve management, redemption, and operational resilience duties. The issuer's business model starts to resemble regulated money issuance. Supervisors tend to care about asset segregation, the quality and custody of reserves, and how quickly users can redeem at par (if par redemption is promised).

A payment-style token or exchange token may sit outside a securities perimeter, yet the intermediaries around it commonly face authorization or registration due to anti-money laundering duties and consumer protection rules. The token itself may be "unregulated," but the market infrastructure is not.

This is why the same token can be sold legally in one venue but not another: authorization and conduct rules attach to intermediaries, not only to the asset.

6.3 Classification decides disclosure burdens and who must speak



Disclosure is a core policy tool because crypto markets often run on information gaps. The legal question is who must speak and what must be said.

Securities-style disclosure is issuer-focused. It assumes a business entity raising funds from the public and requires information about that entity: governance, financial condition, use of proceeds, conflicts, and ongoing reporting.

Crypto-specific disclosure is often product- and protocol-focused. It requires information about how the token works: supply schedule, token distribution, governance controls, key technical dependencies, custody arrangements, and for stable-value tokens, reserve assets and redemption.

The taxonomy must therefore separate:

- disclosure about the issuer or promoter,
- disclosure about the token's rights and economics,
- disclosure about the service provider's business practices (custody, conflicts, order execution), and
- disclosure about risk that arises from the market structure itself (thin liquidity, price manipulation, leverage).



A taxonomy that fails to assign these disclosure duties leaves regulators with only one tool: enforcement after harm occurs. Most regimes try to avoid that by making classification decisions that are clear enough to attach up-front obligations.

7. A comparative taxonomy of token types

Most regulators and standard setters converge on a small set of token categories. The labels differ, but the functional buckets repeat. This section sets a working taxonomy that can be used throughout the paper.

7.1 Payment-style tokens

Core function: a token primarily used as a medium of exchange, a unit of account within a network, or a store of value, typically without an enforceable issuer promise to deliver a service or pay returns.

Common features include:

- no claim on an issuer's assets,
- no contractual right to profits or interest,
- open transferability, and
- value driven by supply and demand rather than redemption.



Regulatory pattern:

- often not treated as securities by default,
- frequently captured by anti-money laundering rules through the activities of exchanges, brokers, and custodians,
- often relevant to consumer protection rules when marketed to retail users,
- frequently connected to derivatives oversight when used as an underlying for futures, options, or perpetual contracts.

Key edge cases:

- a payment-style token bundled with yield promises, buybacks, revenue sharing, or managerial profit schemes can shift into securities-like treatment,
- tokens that start as a payment-style asset can become an investment product through marketing, distribution, and platform design, even if the on-chain rights are minimal,



- payment-style tokens used as collateral in lending and leverage systems can create systemic-like spillovers without ever becoming “money” in a legal sense.

A research taxonomy should not treat “payment token” as a safe harbor. It should treat it as a starting point, then check whether the surrounding activity turns it into something else for legal purposes.

7.2 Utility or access tokens

Core function: a token intended to provide access to a platform, application, or service.

Common features include:

- a right to use a digital service,
- limited or no claim on issuer profits,
- a link between token holding and service access, and
- sometimes restrictions on acceptance (for example, accepted only by the issuer).

Regulatory pattern:



- often treated outside securities rules when the token genuinely functions as an access right and is usable as such at issuance,
- still subject to anti-money laundering duties when it is widely used for payment-like purposes or traded through regulated intermediaries,
- subject to consumer protection and unfair marketing rules when promoted to retail.

Key edge cases:

- “utility” is one of the most abused labels in crypto markets,
- if the token is sold mainly on the expectation that its price will rise through promoter efforts, regulators may treat it as an investment product even if it also has an access feature,
- if the token grants rights tied to fees, revenue, buybacks, or treasury assets, it can become security-like in substance,



- if the “service” is vague, undeveloped, or dependent on future managerial work, the access story may be weak.

The legal analysis should therefore focus on the token’s real function at the time of sale, not the issuer’s aspirational language.

7.3 Security-like tokens

Core function: a token that confers rights and obligations that resemble shares, bonds, derivatives, or units in a pooled investment.

Common features include:

- rights to cash flows, interest, or dividends,
- rights to repayment or redemption tied to issuer performance,
- rights to profits from a common enterprise,
- contractual claims tied to underlying assets,
- trading designed to mirror securities markets.

Security-like tokens appear in two broad forms.



The first is tokenized traditional instruments: shares, bonds, funds, or derivatives issued in a token form. Classification is usually straightforward because the underlying is already a regulated financial instrument, and tokenization is treated as a change in record-keeping and settlement, not a change in substance.

The second is “native” security-like tokens: instruments created on-chain that replicate investment rights through code, contractual terms, or a mix. These are harder because the rights may be poorly documented or hard to enforce.

Regulatory pattern:

- securities offering rules and secondary market rules often apply,
- intermediaries need appropriate authorization,
- market abuse controls and surveillance become central,
- custody rules tend to be strict because client asset loss is a high-risk failure mode.

Key edge cases:

- governance tokens that effectively control a revenue stream may drift into this category,



- tokens with “soft” promises (marketing claims, informal buyback practices) can create disputes about whether there is a real enforceable right or merely an expectation shaped by promotion.

A research taxonomy should treat this bucket as grounded in economic reality: if the token is sold as an investment and buyers rely on others to create value, regulators may treat it as a security-like product even if the formal rights are thin.

7.4 Stable-value tokens (stablecoins and related designs)

“Stablecoin” is not one legal category. It is a family of arrangements designed to keep the token’s price near a reference value. Different designs create different risks and tend to draw different regulatory responses.

A practical taxonomy distinguishes stable-value tokens by how stability is achieved and what rights holders have.

Type A: fiat-backed stable-value tokens with clear reserves

- The issuer holds reserves intended to match outstanding tokens.
- The issuer may promise redemption at par.



- The stabilizing mechanism is redemption against the reserve.

Key regulatory concerns:

- quality, liquidity, and custody of reserves,
- segregation of reserve assets from the issuer's own assets,
- redemption timing and conditions,
- operational resilience of issuance and redemption,
- governance and controls over reserve management.

Type B: asset-backed stable-value tokens with broader reserve assets

- The reserve may include assets beyond cash and short-term government instruments.



- The system may reference a basket, multiple currencies, or other assets.

Key regulatory concerns:

- market risk in the reserve,
- valuation methods and disclosure,
- liquidity under stress,
- concentration and credit risk.

Type C: crypto-collateralized stable-value tokens

- Stability is attempted through over-collateralization with volatile crypto-assets.
- Smart contracts, liquidations, and oracles become core.

Key regulatory concerns:



- liquidation risk and cascade effects,
- oracle and smart contract risk,
- governance control over risk parameters,
- user understanding of the difference between “stable target” and “stable guarantee.”

Type D: algorithmic or endogenous designs

- Stability is attempted through incentives, supply adjustments, or linked tokens rather than strong external reserves.

Key regulatory concerns:

- fragility under stress,
- reliance on market confidence rather than redemption,



- reflexive feedback loops.

Stable-value tokens receive special attention because they can become payment rails. Once used for everyday transfers, they move beyond investor loss and raise broader concerns: payment integrity, consumer protection for users who treat the token as money, and spillovers when the token becomes embedded as collateral across markets.

A good taxonomy also separates stable-value tokens used mainly inside trading platforms (as a trading pair and settlement asset) from those used in general commerce. The same design can pose different public-policy risks depending on reach and use.

7.5 NFTs and tokenized collectibles

Non-fungible tokens are often treated as a separate category because they represent unique items or unique rights. Yet uniqueness alone does not answer the legal question. The key issue is function. NFTs can operate as:

- collectibles with limited financial features,
- access rights (tickets, memberships),
- receipts for off-chain assets,



- or financial interests when fractionalized or pooled.

A practical classification method for NFTs asks:

1. Is the token truly unique in a meaningful way, or is it one of a large series of similar tokens sold under a common scheme?
2. Does holding the NFT give a claim to cash flows, revenue shares, or pooled profits?
3. Is the NFT marketed primarily as an investment, with a focus on resale gains?
4. Does the structure include fractional interests that look like securities or fund units?
5. Is there an intermediary that offers custody, exchange, brokerage, or pricing services that create market integrity risks?

Many NFT arrangements sit outside classic securities classification because they look like digital collectibles. Still, consumer harm can be large: fraud, misleading marketing, wash trading, and custody



losses. That is why the activity-based lens matters. The market infrastructure around NFTs can be regulated even when the NFT itself is treated as a collectible.

7.6 Governance tokens, staking tokens, and DeFi-linked instruments

This category sits at the center of modern disputes because it can combine features of access, voting, and investment.

Governance tokens

These allocate voting rights over protocol parameters: fees, treasury spending, upgrades, and risk controls. The legal question is whether governance is meaningful in an economic sense. A token that votes on cosmetic changes is different from a token that can turn on fee distributions to token holders or direct a treasury that generates revenue.

Key issues include:

- concentration of voting power,
- delegation and control by founders or funds,
- the link between voting and economic benefits,
- the existence of a group that proposes and executes changes in practice.



Staking and liquid staking

Staking can be a technical network function. It can also be packaged as a yield product offered to retail users through intermediaries. Liquid staking tokens represent a staked position and can be traded or used as collateral elsewhere.

Key issues include:

- whether the customer is relying on an intermediary to run validators and manage risks,
- how rewards are calculated and distributed,
- slashing and operational risk,
- disclosure of fees and conflicts,
- how liquid staking tokens behave under stress.

Liquidity provider tokens and pool interests

In automated market maker systems, users provide assets to pools and receive pool share tokens. Those pool interests can resemble fund



interests: users place assets into a shared pool with rules set by code and governance, and their returns depend on trading volume, fees, and price movements.

Key issues include:

- whether users understand the risks (impermanent loss, oracle manipulation, pool drains),
- whether the pool is marketed as an investment product,
- whether there are identifiable promoters or managers who set parameters and collect fees.

The taxonomy should treat these as instruments that can drift between buckets depending on design and marketing. A governance token can be an access tool in one system and a revenue-linked investment in another. A staking product can be a technical service when self-run, and a regulated investment or custody service when offered as a packaged product to the public.

7.7 Wrapped tokens, bridged assets, and synthetic exposures

Crypto markets are full of representations: tokens that point to something else. These instruments complicate classification because the legal claim is often unclear.

Wrapped tokens

A wrapped token usually represents a claim to an underlying asset



held or locked elsewhere. Sometimes a custodian holds the underlying and issues the wrapped token. Sometimes a protocol locks the underlying in a smart contract.

Key classification questions:

- who holds the underlying and on what terms,
- whether the holder has redemption rights,
- whether the arrangement looks like a deposit, custody service, or security-like receipt,
- what happens in insolvency of the custodian or failure of the protocol.

Bridged assets

Bridges move representations across chains. The bridged token's value depends on bridge security and governance.

Key classification questions:

- whether the bridge operator is a service provider with custodial control,



- whether users have enforceable rights against anyone if the bridge fails,
- whether the bridge's governance creates a centralized point of failure.

Synthetic exposures

Some protocols create tokens that track the price of external assets via collateral and price feeds. These can mimic derivatives.

Key classification questions:

- whether the token is a derivative-like exposure,
- whether the arrangement is marketed as a substitute for regulated products,
- whether leverage is embedded,
- how pricing and settlement are handled, and by whom.

A taxonomy that includes these categories helps later chapters explain why certain failures produce legal disputes: users may think they hold



an asset, but legally they may hold only an on-chain record with no enforceable claim.

7.8 Tokenized real-world assets

Tokenization of off-chain assets is often promoted as a way to bring traditional finance onto distributed ledgers: real estate interests, invoices, commodities, fund shares, and government securities. The classification work here is less about “crypto” and more about whether tokenization changes legal rights.

Key questions include:

- what asset is being tokenized, and what legal form does the underlying right take,
- whether the token is the legal right itself or merely a record of ownership,
- who ensures the link between token and underlying asset (custodian, trustee, issuer),
- how transfers are recognized legally,
- what disclosures and safeguards exist for custody and insolvency.



In many cases, tokenized real-world asset products look like traditional securities or fund products and will be treated under existing regimes. The novelty lies in settlement, custody, and operational risk, not in the underlying legal type.

8. Hybrid tokens and the “dominant feature” problem

Crypto-assets often combine features that point in different regulatory directions. A token can grant access to a service, be tradable on markets, and be promoted as an investment. A stable-value token can be used for payments and also used as collateral for leveraged trading. A governance token can vote on technical parameters and also direct a treasury that funds buybacks or rewards.

Hybrid design creates a recurring analytical task:

1. Identify the features that could trigger different regimes (investment rights, redemption rights, payment function, pooled management, derivative-like exposure).
2. Determine which features are central in practice: how the token is marketed, how most users use it, and which rights are real rather than aspirational.
3. Decide which regime takes priority where multiple regimes plausibly apply.



Regulators sometimes use a “dominant feature” concept, even if the phrase is not formal. The idea is simple: do not let a token escape investor protection merely because it also has an access function, and do not treat an access token as a security simply because people speculate on it, unless the token’s sale and structure create the kind of reliance and profit expectation that securities rules are meant to address.

Hybrid analysis also forces regulators to confront a deeper issue: crypto markets can create financial risk without formal legal rights. A token may have weak claims, yet be used in leverage systems that produce cascade failures. That does not necessarily change the token’s classification, but it does strengthen the case for regulating the activities and intermediaries around it.

9. A repeatable classification method for the rest of the paper

A research paper needs a method that can be reused across chapters and across token types. The method below is designed to be applied to any crypto-asset and any service arrangement.

Step 1: Identify the legally relevant promises and rights

Start with the simplest question: what does a holder get, in legal terms?

Key prompts:

- Does the holder receive rights to cash flows: interest, fees, dividends, profit shares, buybacks, distributions?



- Does the holder have a claim on an issuer's assets or a right to repayment?
- Does the holder have redemption rights, especially at par in a fiat currency?
- Does the token grant governance rights that control economically meaningful decisions (treasury, fee switches, distributions, risk settings)?
- Are these rights legally enforceable, or are they merely described in marketing?

Documenting rights matters. In crypto markets, rights are often implied rather than stated. A careful analysis distinguishes:

- enforceable contractual rights,
- rights implemented by code (which can still be altered through upgrades),



- informal practices that buyers may rely on but that are not legally binding.

If the token includes cash-flow rights or redemption promises, the analysis should treat securities-like or payments-like regimes as a serious possibility.

Step 2: Identify the stabilization mechanism and reference value (if any)

If the token claims price stability, classify the stabilization mechanism.

Key prompts:

- What is the reference value (one currency, multiple currencies, a basket, another asset)?
- What is the mechanism (reserve-backed redemption, collateral and liquidation, incentives and supply changes)?
- Who controls the mechanism (issuer, governance group, smart contract with admin keys)?
- What rights do holders have if the mechanism fails?



This step prevents a common error: treating all stable-value tokens as similar. A reserve-backed token with clear redemption differs from a token that targets stability through incentives without firm redemption.

Step 3: Identify transferability, fungibility, and market structure

Next, look at how the token can move and trade.

Key prompts:

- Is the token freely transferable, or limited to a closed network?
- Is it fungible, semi-fungible, or non-fungible?
- Is it commonly traded on secondary markets, and if so, where?
- Is the token used as collateral, and in what types of arrangements?
- Is leverage offered around it, directly or indirectly?

Market structure can change the regulatory interest. A token that is mainly used as an access right inside a closed system raises different



risks than a token that is widely traded on retail platforms with leverage.

Step 4: Identify the actors who sit in the flow

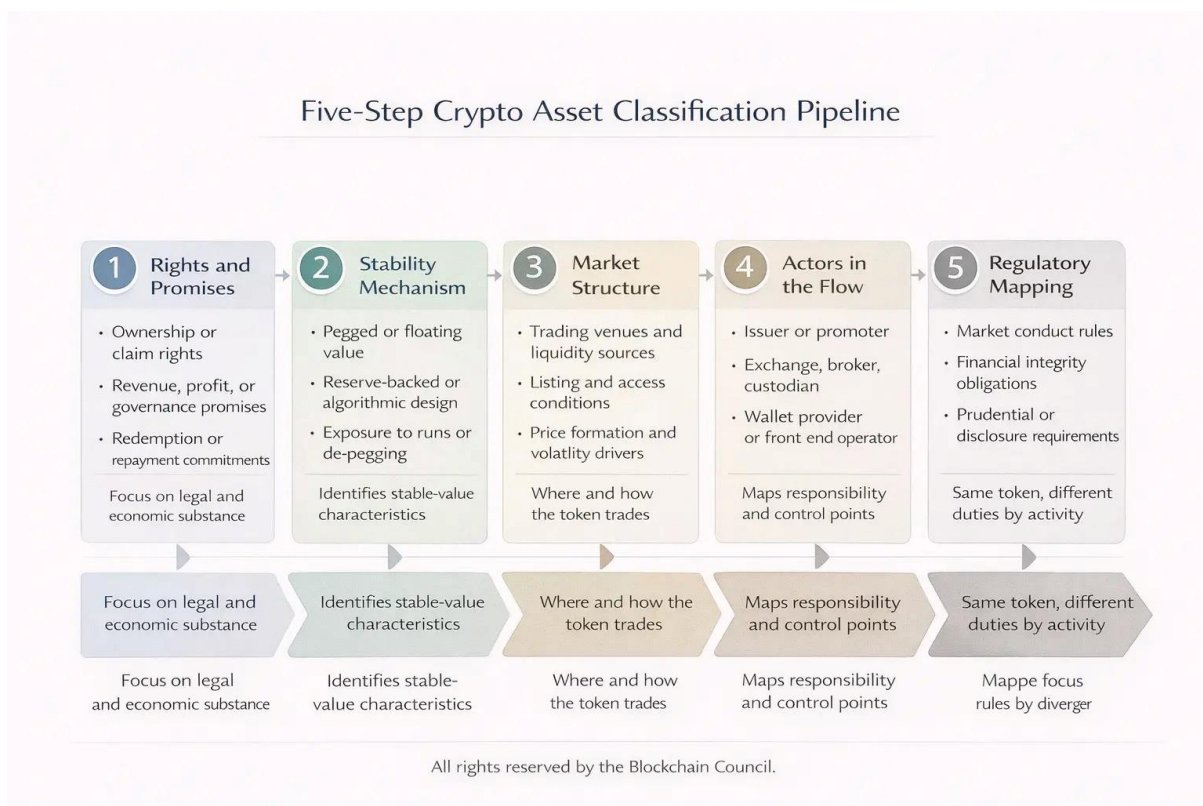
The next step focuses on who does what.

Key prompts:

- Who issues the token and sets its terms?
- Who markets it, and what claims are made to buyers?
- Who provides custody, and who can move customer assets?
- Who runs the trading venue, sets listing rules, and controls market operations?
- Who provides brokerage, order routing, or market making?
- Who provides lending, staking, or yield products tied to the token?

This step shifts analysis from “is the token a security?” to “who is doing regulated work?” Even if the token is not classified as a security, custody and exchange activities can still be regulated heavily.

Step 5: Map to legal regimes and choose the applicable perimeter for each activity



Only after the first four steps should the analysis assign regimes.

A practical mapping can be described as follows:

- If the token’s sale and structure create a clear investment product, apply securities-like analysis and treat securities rules as the primary perimeter.



- If the token is a stable-value instrument with redemption promises and reserve backing, apply payments and money issuance analysis, alongside any securities analysis that might apply to reserve management or profit schemes.
- If the token is a payment-style asset without issuer claims, treat the main regulatory focus as the service providers: anti-money laundering duties, authorization for trading and custody services where required, consumer protection, and market conduct controls.
- If the token is a collectible or access instrument, treat consumer protection and intermediary conduct as key, while remaining alert to financialization features (fractionalization, pooled profits, revenue rights).

This mapping does not force one universal answer. It forces an explicit answer for each regulated activity.

10. Service provider categories and why they form the enforcement hinge

A practical regulatory system cannot supervise every smart contract and every wallet. It supervises people and firms that provide services to others. That is why service provider definitions matter so much: they are the hinge for authorization, supervision, and enforcement.

10.1 Core service types that recur across regimes



Across jurisdictions, the same set of service types appears, even when the labels differ.

Custody

Holding or controlling assets on behalf of clients, including safekeeping of private keys or other means of control. The core risk is loss or misuse of client assets. Custody also creates conflicts: a custodian might lend client assets, stake them, or use them for its own purposes unless rules prevent it.

Exchange and trading venue operation

Providing a market where orders meet, prices form, and trades are executed. The core risks include manipulation, conflicts of interest (especially when the venue also trades on its own platform), weak listing standards, and unfair access to order flow.

Brokerage and order execution

Receiving and transmitting orders, executing orders for clients, and sometimes acting as principal. The core risks include conflicts, poor execution, hidden fees, and sales practices that mislead retail users.

Transfer services

Moving crypto-assets on behalf of clients. The core risks include financial crime, sanctions evasion, and operational errors that lead to irreversible loss.

Issuance and distribution support

Running token sales, placing tokens with buyers, advising issuers, marketing and promotion. The core risks include misleading disclosure, mis-selling, conflicts, and fraud.

Lending, staking, and yield products

Providing products that promise returns tied to crypto markets. The



core risks include maturity mismatch, leverage, rehypothecation, opaque risk-taking, and runs.

A taxonomy of service providers should treat these functions as building blocks. Real firms combine them. That combination creates conflicts. A common crypto business model is vertical integration: the same entity operates a trading venue, brokers retail flow, makes markets, holds custody, and offers leverage and yield products. Traditional finance often separates these roles or subjects them to strict conflict rules. Crypto markets frequently began without such separation, making service provider classification central to investor protection.

10.2 The control test and why it dominates custody debates

In crypto, “custody” is often described through private-key control. That is useful but incomplete. Control can exist through:

- private keys held by the provider,
- multi-signature arrangements where the provider is one signer,
- smart contract admin privileges that allow freezing or moving assets,
- account systems where the provider can approve or reject withdrawals,



- technical dominance over infrastructure that determines whether a user can access funds.

For regulation, the question is often practical: can the provider move customer assets or block movement? If yes, the provider is in a custodial role for risk purposes, even if the legal paperwork avoids the word “custody.”

10.3 The “as a business” threshold and decentralized systems

Many frameworks use “as a business” language to avoid capturing ordinary users. Yet decentralized systems blur the line between users and service providers.

A useful research approach distinguishes:

- people who use protocols for themselves,
- people who provide services to others (front ends, custodial products, advisory services, brokerage),
- groups that exercise effective control over key functions (admin keys, protocol upgrades, governance execution),



- and entities that earn fees from facilitating transactions for the public.

This approach avoids false binaries. A system can be technically decentralized yet have a small group that controls key parameters and earns fees. Conversely, a system can have clear promoters but still allow self-custody and peer-to-peer use. The classification question is not “is it decentralized?” but “who does what for whom, and who holds power over user assets and market rules?”

11. Comparative taxonomy snapshot across common regulatory approaches

Because terminology differs across jurisdictions, a research paper benefits from a simple crosswalk. The goal is not to claim uniformity but to show functional convergence.

Payment-style tokens

Often described as cryptocurrencies, exchange tokens, or payment tokens. Typically treated as non-securities as an object, while exchanges, brokers, and custodians face anti-money laundering and conduct duties.

Access tokens

Often described as utility tokens. Often treated outside securities rules when they function as access rights at issuance, with warnings that the label does not decide the result.

Security-like tokens

Often described as security tokens, investment tokens, asset tokens, or tokenized securities. Typically treated as within securities rules, with full intermediary and market integrity duties.



Stable-value tokens

Often split into tokens referencing a single currency and those referencing multiple assets or baskets, or separated by reserve-backed versus algorithmic designs. Common focus is on reserves, redemption, and systemic spillovers.

Service providers

Common focus is on exchanges, custodians, brokers, trading venues, and entities offering yield products.

This crosswalk supports later chapters that compare authorization models without getting trapped in local vocabulary.

12. Where definitional choices shape market design

Definitions do not merely label markets. They shape them. Firms design around the edge of the perimeter.

12.1 Regulatory arbitrage is usually definitional arbitrage

Firms and projects structure tokens and services to fall into the lightest category available.

If “utility token” language avoids securities oversight, issuers will lean into access narratives and downplay investment features. If stable-value tokens can operate under lighter rules when framed as “crypto” rather than “money,” issuers will avoid redemption promises that trigger money issuance duties, even if users treat the token as money in practice. If anti-money laundering registration is lighter than full market authorization, firms may locate in jurisdictions where only the registration layer applies, while serving customers elsewhere through remote means.



This is why definitions must be tied to substance and activity, not labels. A perimeter that can be avoided by changing marketing language invites weak disclosure and mis-selling.

12.2 Disclosure is enforceable only when the perimeter is clear

Disclosure obligations attach to defined products and activities. If the perimeter is vague, enforcement becomes uncertain. That can produce two failure modes:

- **Overreach:** regulators attempt to stretch definitions in a way that courts reject, creating uncertainty and uneven treatment.
- **Underreach:** harmful products remain effectively outside enforceable disclosure duties, leaving retail users dependent on platform marketing and informal claims.

A taxonomy that assigns disclosure duties clearly, and ties them to activities that can be supervised, supports predictability. It also supports better enforcement because regulators can prove breaches of defined duties rather than argue about the nature of the technology.

12.3 Enforcement choices depend on which definition can be proved

In enforcement, agencies often select the theory that is most provable.

- A market authority may focus on whether the token fits an investment category, or whether the intermediary performed regulated dealing or trading venue functions.



- A financial crime authority may focus on whether the entity provided exchange, transfer, or custody services as a business without meeting anti-money laundering duties.
- A payments authority may focus on whether a stable-value token issuer made redemption promises or issued a money-like instrument without authorization.

Definitions are therefore tools of litigation and supervision, not neutral words.

13. A working asset taxonomy for later chapters

To support the rest of the paper, this section ends with a working taxonomy that links assets to typical regulatory concerns. It is not a legal code. It is a research framework.

13.1 Category 1: Non-claim, non-stable tokens (payment-style assets)

What they are:

- freely transferable tokens without issuer claims and without stability targets.

Typical risks:



- price volatility and retail losses,
- fraud and manipulation on trading venues,
- custody failures at intermediaries,
- financial crime through transfers,
- spillovers through leverage and collateral use.

Typical regulatory focus:

- authorization and conduct rules for exchanges, brokers, and custodians,
- anti-money laundering duties,
- restrictions or warnings for retail marketing and leverage.



13.2 Category 2: Access tokens (service entry rights)

What they are:

- tokens tied to access or consumption of services.

Typical risks:

- misleading marketing when access is speculative or undeveloped,
- resale markets that turn access into investment-like trading,
- consumer protection issues around delivery of the promised service.

Typical regulatory focus:

- disclosure about the service and token function,
- marketing and sales conduct rules,



- anti-money laundering coverage where relevant,
- platform conduct rules if widely traded.

13.3 Category 3: Investment and claim tokens (security-like assets)

What they are:

- tokens that grant profit rights, repayment rights, pooled investment exposure, or derivative-like payoffs.

Typical risks:

- the full set of investor protection and market integrity risks familiar from securities markets,
- mis-selling and unsuitable distribution,
- conflicts at trading venues and brokers,



- custody and segregation failures.

Typical regulatory focus:

- securities offering rules,
- authorization for intermediaries,
- market abuse rules,
- custody and client asset protections,
- ongoing reporting duties where appropriate.

13.4 Category 4: Stable-value instruments (money-like tokens and settlement assets)

What they are:

- tokens designed to stay near a reference value.

Typical risks:



- reserve shortfall, liquidity stress, and runs,
- redemption failures and consumer harm,
- operational failures that disrupt payments,
- contagion if widely used as collateral or settlement.

Typical regulatory focus:

- reserve requirements, safeguarding, and redemption rules,
- governance and risk management for issuers,
- oversight of critical service providers (custodians, administrators),
- limits on use or distribution where risks are high.



13.5 Category 5: Unique tokens and collectibles (NFTs and similar instruments)

What they are:

- tokens representing unique items or unique rights.

Typical risks:

- fraud, wash trading, misleading promotions,
- custody and marketplace risks,
- financialization through fractionalization and pooled schemes.

Typical regulatory focus:

- consumer protection and market conduct for marketplaces and custodians,
- securities-like treatment where fractionalization or profit schemes are present,



- anti-money laundering coverage where marketplaces function like exchange services.

13.6 Category 6: Representations and linked exposures (wrapped, bridged, synthetic)

What they are:

- tokens that represent, mirror, or track another asset or value.

Typical risks:

- unclear legal claims and weak user rights,
- bridge and custodian failure,
- derivative-like risk hidden in simple wrappers,
- governance concentration over critical functions.

Typical regulatory focus:



- custody and operational risk controls,
- disclosure about claim structure and redemption,
- treatment as derivatives or securities-like products where payoff structure warrants it.

14. Practical guidance for using this taxonomy in later analysis

The rest of the paper should treat taxonomy as a tool for consistent analysis, not a static checklist. The same asset can move categories depending on changes in design, marketing, and use. Likewise, the same activity can trigger different duties depending on whether it is offered to the public as a business, whether it involves custody or control, and whether it creates reliance.

To keep the analysis disciplined, later chapters should follow these rules:

1. Start with rights and promises, not labels.
2. Separate asset classification from activity classification.

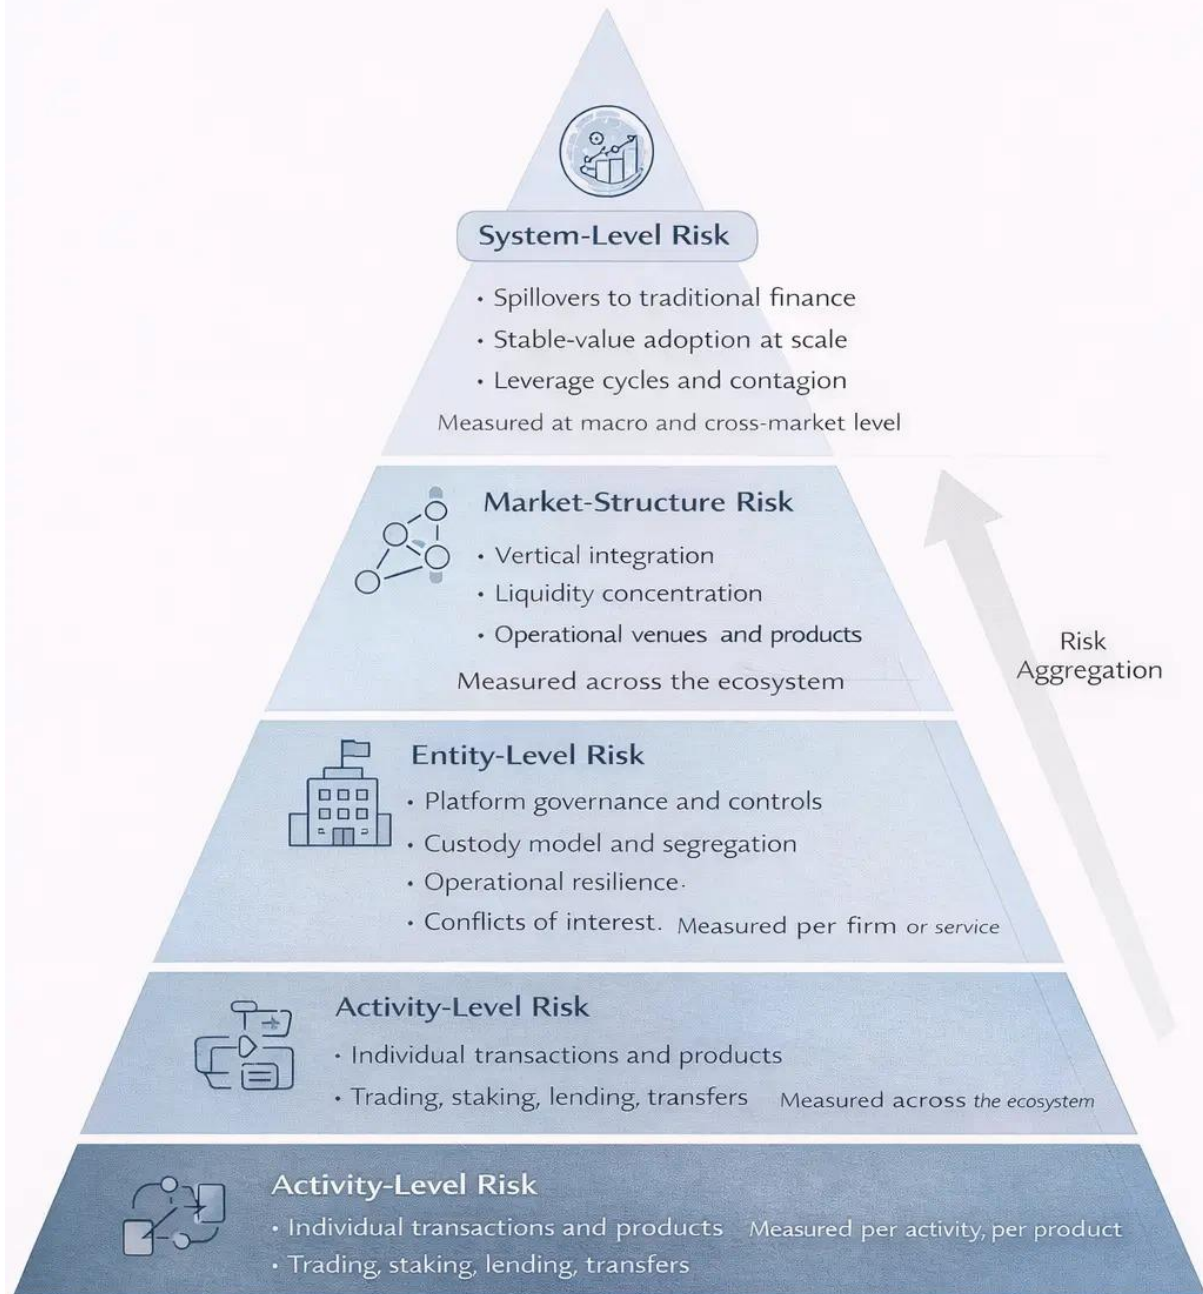


3. Identify control points where duties can realistically attach.
4. Treat stable-value claims as a special case because they shape user expectations and payment risk.
5. Treat vertical integration as a core risk driver and assess conflicts as a first-order issue.
6. Be explicit about what is unknown: enforceability of rights, governance control, reserve composition, and operational dependencies often require facts that marketing materials omit.

2) Regulatory objectives and risk model



Four-Layer Risk Model for Crypto Regulation



Measured at macro and cross market level



Crypto regulation looks messy because different agencies arrive with different mandates and different habits. A securities regulator starts with disclosure, market abuse, conflicts, sales practices, and custody of client assets. A banking supervisor starts with capital, liquidity, concentration, and group risk. An AML authority starts with customer due diligence, sanctions controls, and information that follows value across borders. A payments regulator starts with redemption, settlement, outages, and consumer complaints.

That mix can feel like a pile of unrelated rulebooks. It is not. Step back far enough and most frameworks follow the same basic logic.

They begin by naming the economic function of what is happening: trading venue, broker, custodian, issuer, payment instrument, lending or borrowing, staking service, stable-value arrangement, or a blend. They then ask a short set of questions: what can go wrong, who carries the loss, how fast can the loss spread, and what can be done to prevent it or limit it. Finally they pick regulatory outcomes that match the harm. Where a crypto activity recreates familiar risks from traditional finance, the outcome looks familiar too. Where crypto changes the risk drivers (code risk, governance risk, remote access, irreversible transfers), the controls shift.

This chapter sets out that objectives-and-risks spine in a way that can be used later in the paper. The aim is not to argue that every jurisdiction has reached the same answer. The aim is to show why the same questions keep appearing, why the same pressure points keep drawing supervision, and why regulators keep returning to functional regulation and proportional oversight.

2.1 Regulatory objectives as the anchor for scope and design



Regulatory objectives do more than justify intervention. They shape the architecture of a regime.

They shape the perimeter. If the main objective is investor protection, the perimeter tends to follow investment-like selling and trading. If the main objective is payments safety, the perimeter tends to follow redemption promises and settlement use. If the main objective is stopping illicit finance, the perimeter tends to follow control points where customer checks can be done.

They shape licensing thresholds. Some objectives push toward licensing (market integrity and consumer protection), while others can be met through registration plus program requirements (AML). Objectives also shape who is expected to be supervised as a firm rather than as a product. Many regimes choose the firm because products move and mutate quickly.

They shape conduct and prudential obligations. Disclosure, custody rules, segregation, governance, conflicts management, capital or liquidity buffers, and audit and reporting duties all flow from objectives. The same activity can trigger different requirements if lawmakers have chosen different objectives as the lead.

They shape supervisory intensity. A regime with high concern about retail harm will stress promotions, sales practices, complaints, and custody controls. A regime with high concern about systemic spillovers will stress group risk, liquidity stress, stable-value arrangements, and links to banks and payment rails. A regime with high concern about sanctions evasion will stress screening, travel rule compliance, and information quality.

They shape cross-border cooperation. A regulator that sees crypto primarily as a global market integrity problem pushes for trade data,



listing information, and enforcement help. A regulator that sees crypto primarily as an illicit finance channel pushes for information that follows value between service providers and for joint investigations. A regulator that sees crypto primarily as a stability problem pushes for shared monitoring of large stable-value arrangements, group exposures, and crisis response playbooks.

Although the objectives overlap, they do not always point to the same policy choice. That tension is not a bug. It is the design problem.

A workable chapter on objectives and risk therefore needs two things. First, a clean statement of the main objectives that show up across regimes. Second, a risk model that explains how those objectives translate into priorities, tools, and supervisory choices.

Across jurisdictions the objectives tend to cluster into five buckets:

1. Market integrity
2. Investor and consumer protection
3. Financial stability and containment of system-wide spillovers
4. Illicit finance and national security risks (AML/CFT, sanctions, fraud as predicate crime)
5. Innovation, competition, and broader economic policy goals

These buckets are not equal in every jurisdiction, and they are not always equal in the same jurisdiction over time. Many regimes start with financial crime controls because the legal path is clear, then add consumer protection and market integrity rules once activity grows. Others start by licensing trading venues because retail exposure is high. Others start with stable-value arrangements because those touch payments.



The rest of this chapter treats each objective in turn, then pulls them together into a practical risk model.

Regulatory Objectives That Shape Crypto Rules

Objective	What Regulators Want	How Rules Show Up	What "Compliance Proof" Looks Like
Market Integrity	Fair and orderly markets	Market abuse surveillance	Trade monitoring logs
Investor and Consumer Protection	Clear risks and honest disclosures	Disclosure and conduct standards	Risk statements and customer records
Financial Stability	No spillover into the real economy	Reserves, liquidity, and stress testing	Redemption data and exposure reports
Illicit Finance Prevention	Reduced abuse without blanket bans	AML controls and reporting	KYC files and alert trails
Innovation and Fair Competition	Room to experiment safely	Regulatory sandboxes and thresholds	Pilot approvals and usage limits

All rights reserved by the Blockchain Council.

2.2 Market integrity

2.2.1 What regulators mean by market integrity in crypto

In finance, market integrity is a plain idea: markets should be fair, orderly, and built on information that is not rigged. Prices should not be the product of fraud, fake volume, or inside knowledge misused. Customers should not be steered into trades through hidden conflicts. Rules should apply to the venue, the intermediaries, and the people who make markets.

In crypto markets, those aims run into features that make classic misconduct easier to run and harder to spot.



First, crypto trading is cross-border by default. A retail user in one country can trade on a platform based somewhere else, using an affiliate that is registered in a third place, with liquidity that is split across many venues. That means price formation is global, and the weakest venue can set the tone for the whole market.

Second, crypto trading often sits inside groups that combine roles that traditional finance separates. A single firm or group may run the venue, broker retail orders, act as a dealer, make markets, offer margin, hold client assets, and run the stable-value token that serves as the main settlement unit on the venue. When those roles sit under one roof, conflicts are not occasional. They are built into the structure.

Third, the microstructure is partly visible and partly hidden. On-chain transfers can be seen by anyone, but much of the trading that sets price happens off-chain in internal order books. Many venues internalize orders, route flow to affiliates, or match customers against a house desk. Even when trade prints are published, the identity of the real actor can be masked through layered accounts.

Fourth, listing decisions are fast and often commercial. Tokens can be admitted to trading with thin disclosure and weak due diligence. A listing can also be a promotional event. That creates incentives to list early, list often, and deal with problems later.

Fifth, retail participation has often been a large share of flow. Retail can be a healthy part of markets when protections are in place. It becomes a problem when most users are inexperienced, trading is gamified, and risk is hidden behind app design.

These features do not create new forms of dishonesty. They change the odds and the payoff.

The misconduct patterns regulators focus on are familiar:



- Fraud and misrepresentation in token offerings, listings, and promotions
- Manipulation such as wash trading, spoofing, coordinated pump-and-dump groups, and price marking around liquidations
- Misuse of non-public information, including inside knowledge about listings, delistings, contract upgrades, or large forced liquidations
- Conflicts of interest that push customers into worse prices, higher fees, or products that mainly benefit the venue
- Abusive liquidation rules and unfair margin mechanics
- Misleading fee schedules and hidden spreads
- Weak surveillance and weak response that allow abuse to persist

What changes in crypto is that these problems can run at high speed and at global scale.

2.2.2 Integrity controls regulators tend to require

Market integrity is not achieved by a single rule. It is built through a set of controls that cover the life of a trade: admission of the product, access to the market, order handling, execution, settlement, custody, and post-trade reporting.

A workable integrity program for a trading venue usually includes at least five building blocks.

Trading rules and surveillance. The venue needs clear rules on order types, trading hours, outages, and error handling. It needs surveillance that can spot wash trades, self-trading, spoofing, layering, and coordinated manipulation. That surveillance has to cover



not only one venue but also patterns that jump between venues or run through affiliates.

In crypto, surveillance also has to account for the link between trading and on-chain transfers. A pump-and-dump group may move tokens between wallets to create the story of demand. A manipulator may use on-chain moves to create fear, then trade on a venue where those moves are not linked to identity.

Fair access and fair execution. A venue should be able to explain how orders are routed, how matching works, and whether some customers get faster access or better information. In traditional markets, regulators spend a lot of time on best execution and fair access because small differences add up. In crypto, those differences can be larger because many venues are also dealers.

Fair execution also means rules around internalization and matching against the house. If a venue or its affiliate trades against customers, it has to manage conflicts, disclose the practice, and put limits in place so customers are not used as a source of easy profit.

Controls on conflicts of interest. Conflicts are the core integrity problem in many crypto markets. The risk is not only that a conflict exists, but that it is not managed.

Common conflict channels include:

- The venue listing tokens while holding positions or receiving token allocations
- The venue or an affiliate acting as a market maker in a token it has just listed
- House trading in the same products sold to retail users



- Using customer order flow to trade ahead
- Steering customers into products with high fees and fast liquidation

Regulators respond in different ways. Some require separation of certain functions, such as custody from trading, or broker functions from venue operation. Others allow combination but impose governance requirements, conflict policies, and independent oversight. The design choice varies, but the logic is the same: integrity fails when a venue can quietly pick winners.

Listing and delisting discipline. Token admission standards matter because listings shape retail exposure. A venue that will list anything becomes a factory for harm. A venue that applies clear standards can reduce fraud and manipulation risk.

A listing framework typically asks:

- What does the token do, and what rights does it claim to give holders?
- Who controls upgrades and key parameters, and what is the change process?
- How concentrated is the supply, and when do large allocations unlock?
- Is there a clear public description of supply changes, fees, and governance?
- Are there known security weaknesses or admin privileges that could be abused?
- Is the token likely to be used mainly for speculation, and if so, is it prone to manipulation due to thin liquidity?



Ongoing monitoring matters as much as the initial listing. Token risk can change quickly after a contract upgrade, a governance vote, a hack, or a large unlock.

Handling of margin and liquidation. A large amount of retail harm in crypto markets has come from margin products with harsh liquidation rules and unclear fees. Even where spot tokens are the main focus, venues often tie spot and derivatives in a single account, letting customers move between them with one click.

From an integrity standpoint, liquidation is a danger zone. The venue sets the rules, the venue often runs the liquidation engine, and the venue may have incentives to liquidate quickly if it earns fees or captures spreads. A credible integrity framework therefore treats liquidation mechanics as a regulated feature, not a private design choice.

Controls in this area can include:

- Clear disclosure of liquidation triggers, fees, and the order of operations
- Guardrails against liquidating customers at off-market prices
- Policies that limit the venue's ability to profit from customer liquidations
- Review of margin settings and automatic changes during stress
- A clear process for handling outages that prevent customers from reducing exposure

2.2.3 Why market structure keeps coming up



Crypto integrity debates often circle back to structure rather than single acts of fraud. That is because structure shapes the incentives and the odds of being caught.

Four structural features matter.

Fragmentation. When trading is split across many venues, and customers can move quickly between them, manipulation can be done in one place and cashed out in another. Fragmentation also makes it hard to build a full view of the market. A single venue may see only a slice of order flow.

Remote provision. When an offshore platform serves domestic customers through the internet, the home regulator has fewer tools. It may be able to regulate local marketing and local on-ramps, but it cannot easily inspect the offshore venue.

Group complexity. Many crypto firms use multiple entities: one for the venue, one for custody, one for marketing, one for derivatives, one for stable-value issuance, and one for global operations. That structure can be normal business planning. It can also be a way to hide risk and move liability.

Vertical combination of roles. When the same group controls trading, custody, and house dealing, the market becomes dependent on a small set of private rule-makers. In that setup, integrity depends less on the behavior of retail users and more on the conduct of the venue and its affiliates.

The regulatory response to these features is not always the same, but it tends to move in the same direction: tighter authorization for trading venues and custodians, stronger conflict rules, clearer disclosure, and more cross-border cooperation.



2.3 Investor and consumer protection

2.3.1 Why protection looks broader in crypto

In securities markets, investor protection is mainly about investment decisions: getting enough information to choose, ensuring fair dealing, avoiding conflicts, and having rights when things go wrong.

Crypto pulls in that set of concerns and adds another layer. Many crypto products are not only investment-like. They are also consumer tech products: wallets, apps, custody accounts, payment features, staking dashboards, and customer support systems. The harm is not only a bad trade. It can be the loss of the asset itself due to a hack, an address error, a freeze, or a platform failure.

A practical protection model in crypto has to cover a wide harm spectrum:

- Losses from volatility, manipulation, or predatory trading mechanics
- Losses from custody failure, commingling, or misuse of client assets
- Losses from hacks, phishing, and account takeover
- Losses from user mistakes that cannot be reversed
- Confusion about product features such as fees, spreads, redemption rights, and lockups
- Lack of redress when a platform fails or refuses to help
- Sales and promotion that create false expectations



This is why many frameworks speak of both "investor" and "consumer" protection. Some crypto customers are investing. Some are using. Many are doing both at once.

2.3.2 Disclosure as a protection tool

Disclosure is often the first tool regulators reach for because it is a way to shift a market from persuasion to information.

But crypto disclosure has to deal with a basic problem: what, exactly, is being disclosed?

In a typical share offering, there is a company with accounts, managers, and legal duties. In crypto, many tokens are linked to a project but do not give a legal claim on that project. Some tokens are issued by a company. Some are issued by a foundation. Some are distributed by code. Some have no single issuer that can be sued.

A workable disclosure model therefore separates four things.

First, disclosure about the token itself: supply, distribution, vesting, governance controls, and how changes are made.

Second, disclosure about the system the token is used in: what the protocol does, what risks users take on, what happens in a hack or an upgrade, and what parts depend on third parties.

Third, disclosure about the service provider selling or holding the token: fees, spreads, conflicts, order handling, custody practices, and what the provider may do with customer assets.

Fourth, disclosure about rights and recourse: what legal claim a holder has, what the complaint path is, what happens in insolvency, and what losses the customer bears.



Regulators often require risk warnings because crypto markets have repeated a pattern: products are sold as simple while the risks are hidden in the details. Risk warnings are not a cure, but they can reduce the harm caused by surprise.

A useful research point is that disclosure only works when it is paired with enforceable duties. If disclosure is only a marketing exercise, it becomes a tool for misdirection. If it is paired with liability for false statements and duties of fair dealing, it can improve behavior.

2.3.3 Sales practices, suitability, and product controls

Crypto markets have often reached retail users through a mix of social media promotion, referral schemes, influencer marketing, and app-based gamification. That distribution channel raises questions that securities regulators know well: is the customer being sold an unsuitable product, is risk being downplayed, and are incentives distorting advice?

Not every regime uses the same tools. Some rely mainly on disclosure and promotion rules. Some impose appropriateness tests for higher-risk products. Some restrict certain products for retail users. Some restrict incentives. The choices differ, but they all sit under the same objective: reduce harm caused by mis-selling.

A practical protection model tends to treat some products as higher risk by design:

- Products with margin and forced liquidation
- Products with complex payoffs, including perpetual contracts and structured returns
- Yield products that depend on rehypothecation, lending, or opaque trading



- Tokens marketed as stable when stability is not backed by firm rights

Retail protection also includes cooling-off features, limits on incentives, and clearer fee presentation. Those are not only consumer comfort measures. They shape behavior by changing how easy it is to trade on impulse.

2.3.4 Custody, safeguarding, and the difference between loss and loss of access

Custody is the center of retail harm in crypto markets.

In traditional finance, custody failures matter, but account systems and legal frameworks often provide ways to unwind errors or compensate losses. In crypto, a mistaken transfer can be final. A stolen private key can mean the asset is gone. A platform hack can wipe out customer assets, and customers may discover too late that the platform used client assets for its own purposes.

A custody and safeguarding regime typically tries to prevent four failure modes.

Commingling and misuse. If a platform mixes client assets with its own, customers become unsecured creditors in insolvency. If a platform uses client assets for lending or trading without clear consent and controls, customers bear hidden risk.

Weak key management. If private keys or signing devices are poorly protected, a single breach can drain assets. A custody regime therefore expects access controls, separation of duties, limits on withdrawals, cold storage practices, and incident response.

Operational failure and outages. Customers can be harmed even without theft if a platform freezes withdrawals, loses access to keys,



or suffers outages during stress. Protection in this area is about change management, business continuity, and clear communication.

Forks, airdrops, and protocol changes. Crypto assets can change by code. A fork can create two assets. An airdrop can create a new claim. A contract upgrade can change token behavior. A custody regime needs rules for how a custodian handles these events and whether customers have rights to resulting assets.

A research paper should also separate "loss" from "loss of access." In crypto, users can lose assets because they are stolen, or they can lose access because a custodian freezes accounts, loses keys, or is shut down. Both produce the same result for the user: the asset is unusable.

2.3.5 Redress, complaints, and the reality of cross-border firms

Consumer protection is not only about preventing harm. It is also about what happens after harm.

In crypto markets, customers often have limited access to redress. Firms may be offshore. Contracts may contain broad disclaimers. Customer support may be slow. Legal claims may be unclear.

Regimes that take consumer protection seriously therefore tend to require basic building blocks:

- A clear complaints process and response times
- Records that allow disputes to be investigated
- Clear disclosure of what losses are covered, if any
- In some cases, access to ombuds or alternative dispute resolution
- In some cases, compensation arrangements or insurance, with clear limits



Even where formal compensation is not offered, requirements around complaint handling can change firm behavior. It forces firms to keep records, measure customer harm, and face supervisory questions when patterns appear.

2.4 Financial stability

2.4.1 When crypto becomes system-relevant

Financial stability authorities intervene when a set of losses can spread beyond the original customers and become a wider stress event. The trigger is not price volatility by itself. Volatility is common in many assets. The trigger is the channel that turns losses into broader dysfunction.

Those channels have become clearer as crypto markets have grown and as links to traditional finance have increased. They include:

- Borrowed trading and forced liquidation cascades
- Stable-value token runs that force asset sales in reserve portfolios
- Interconnectedness through collateral chains and rehypothecation
- Links to banks, payment firms, funds, and market infrastructure
- Operational failures that disrupt settlement or payment rails
- Rapid scale combined with weak governance and weak controls

A stability lens therefore asks a different set of questions than an investor protection lens. The investor lens asks who loses money and whether they were treated fairly. The stability lens asks whether the loss creates feedback loops and threatens the functioning of markets or payment systems.



2.4.2 Functional regulation in stability terms

Stability frameworks have pushed a simple idea: match the oversight to the function.

If an entity performs an exchange-like function at scale, it should face exchange-like expectations on governance, risk controls, and resilience.

If an entity offers deposit-like promises, it should face bank-like prudential expectations, or it should be prohibited from making such promises.

If an entity issues a money-like instrument used for payments, it should face strong requirements around redemption, reserves, and operational resilience.

This is not about forcing crypto into old boxes for the sake of it. It is about preventing the creation of bank-like risk outside bank rules.

2.4.3 Stable-value tokens: run risk and payment adjacency

Stable-value tokens sit closest to payment systems. That closeness changes the risk model.

A token that claims stability invites users to treat it as cash-like. Users may park funds in it, use it for transfers, and rely on it to settle trades. If the token fails to hold value or fails to redeem, the harm is not only a trading loss. It can be a payments disruption.

The stability risks in a stable-value arrangement can be grouped into three parts.

Reserve risk. If the reserve assets are risky, illiquid, or concentrated, the issuer may not be able to meet redemptions under stress. If the



reserve is held through weak custody arrangements or commingled with issuer assets, the reserve may not be available when needed.

Run dynamics. If users doubt the reserve or the issuer, redemptions can surge. Early redeemers get paid, late redeemers get less. That creates a rush to exit. The rush can force the issuer to sell reserve assets quickly, putting pressure on those markets.

Operational risk. Issuance and redemption depend on rails: banking partners, payment processors, custody providers, and the issuer's own systems. A failure in those rails can freeze redemptions even if reserves are sound.

Stable-value designs also differ. A fully reserved, redeemable token creates one set of risks. A crypto-collateralized token with liquidation mechanics creates another. An algorithmic design creates a different set again, often with sharper run dynamics.

From a stability perspective, it matters less what the token is called and more what it promises and how redemptions work.

2.4.4 Borrowed trading, collateral chains, and forced sales

Many stress events in crypto markets have been driven by borrowed trading and by collateral chains.

Borrowed trading turns price moves into forced trades. When prices fall, margin calls rise. If customers cannot post more collateral, positions are closed out. Those closures can push prices down further, triggering more closures.

Collateral chains deepen the problem. A user deposits one asset, borrows another, buys a third, then pledges that third as collateral elsewhere. The same underlying value supports multiple exposures. When values fall, the unwind can run through many linked positions.



Rehypothecation adds another layer. If firms reuse customer assets as collateral for their own borrowing, a liquidity problem at the firm can turn into a run. Customers rush to withdraw, but the assets are locked elsewhere.

A stability-focused framework therefore pays attention to:

- How much borrowing and margin is offered to retail and to large traders
- How collateral is valued, including the role of price feeds and thin markets
- Whether collateral is reused, and on what terms
- Whether firms have liquidity plans for stress
- Whether the same group runs both trading and lending, creating incentive to keep risk hidden

2.4.5 Links to banks and payment firms

Even if most crypto losses fall on crypto users, stability authorities care about links to the wider financial system.

Banks may provide custody, settlement accounts, credit lines, and payment rails. Funds may hold crypto exposures. Payment firms may use stable-value tokens in products.

The core stability question is not whether banks "like" crypto. It is whether exposures are sized, managed, and disclosed in a way that prevents surprise.

Prudential tools in this area include:

- Capital treatment of crypto exposures
- Limits on concentration and large exposures



- Expectations around risk management and governance for custody and third-party reliance
- Disclosure to support market discipline

These tools aim to keep crypto from becoming a hidden amplifier inside core finance.

2.4.6 Fragmentation as a stability risk

Fragmentation is not only a market integrity problem. It is also a stability problem.

When activity is split across borders and across many venues:

- Firms can choose the weakest oversight point.
- Supervisors may see only local pieces of group risk.
- Crisis response can be slow because information is scattered and legal powers stop at borders.
- Authorities can trip over each other, or leave gaps, because mandates overlap.

This is why stability-focused work often stresses cross-border cooperation, shared monitoring of large arrangements, and crisis planning.

2.5 Illicit finance: AML/CFT, sanctions, and predicate crime

2.5.1 Why illicit finance sits at the base of many regimes

Illicit finance is not treated as a side issue in crypto. It is often the first reason governments act.

Crypto can move value across borders quickly. It can be used in fraud, ransomware, and scams. It can be used to move proceeds of crime. It can be used to evade sanctions.



Those risks are not unique to crypto, but crypto changes the mechanics. Transfers can be fast and final. Wallet addresses can be created easily. Value can be routed through many hops.

Because of that, many jurisdictions have chosen to apply AML/CFT obligations to crypto service providers even when other parts of crypto remain lightly regulated.

2.5.2 The risk-based approach and why it matters

A risk-based approach in AML/CFT starts with a simple idea: not every activity carries the same risk, so controls should vary with risk.

In crypto, that matters because the ecosystem includes many different roles.

A fiat on-ramp that takes bank transfers and lets customers buy crypto carries a different risk than a non-custodial wallet software provider. A custodian that holds keys for retail users carries a different risk than a protocol developer writing code. A broker that routes trades carries a different risk than a miner or validator.

A workable risk-based approach therefore:

- identifies which roles can realistically implement customer checks and monitoring
- applies stronger measures where risk is higher
- applies lighter measures where risk is lower, within clear boundaries

In practice, the highest AML/CFT attention usually falls on fiat gateways, exchanges, brokers, custodians, and stable-value issuers, because those roles sit at points where money enters and exits and where large volumes are concentrated.



2.5.3 Core AML/CFT controls in crypto

AML/CFT programs in crypto service providers tend to rely on a familiar set of controls, adapted to crypto specifics.

Customer due diligence. Identifying the customer, understanding ownership and control for entities, and assessing risk. In crypto, that also means understanding how the customer intends to use the service: trading, transfer, payments, or business activity.

Transaction monitoring. Looking for patterns that suggest laundering, fraud proceeds, or sanctions evasion. Crypto adds a layer: on-chain tracing and wallet risk assessment. The quality of that work depends on tools and on human judgment.

Sanctions compliance. Screening customers and, where feasible, screening wallets and counterparties. Because wallets can be created easily, sanctions controls rely on a mix of list screening, behavior monitoring, and response to law enforcement alerts.

Information that follows transfers. Many regimes require information about the originator and beneficiary to accompany transfers between service providers. The aim is to prevent value from moving through the regulated sector with no identifying information.

Suspicious activity reporting. Reporting activity that looks like laundering, fraud, or sanctions breaches.

Crypto changes the operating environment for these controls. A criminal can move funds quickly. They can use mixers, chains of wallets, or bridges. They can use decentralized exchanges and self-custody to bypass intermediaries. That does not make controls pointless. It makes the identification of control points more important.

2.5.4 The collision with decentralization narratives



A recurring design challenge is when actors claim they are "not intermediaries" while still enabling high-risk flows.

Regulators often respond by asking a practical question: who can implement controls?

If there is an operator that runs a front end, holds admin privileges, sets fees, or controls upgrades, the operator may be treated as a service provider for AML purposes.

If there is no operator, regulators may focus on gateways: fiat on-ramps, custodial services, and venues that list the assets.

If controls are not feasible, authorities may restrict certain activities, require stronger measures at the gateways, or push for new reporting methods.

This is not always clean. The boundary between software, service, and business is contested. But the objective is clear: stop the regulated sector from becoming a blind conduit.

2.5.5 Fraud as predicate crime and the retail harm loop

Illicit finance in crypto is not limited to classic laundering. Fraud is a major driver.

Many crypto frauds are retail-facing: fake investment schemes, fake token launches, phishing, account takeover, and social engineering. The proceeds then move through the same venues used for legitimate trading.

That creates a loop between consumer protection and AML. A platform that fails to stop scams creates victims. Those victims generate proceeds for criminals. The criminals then use exchanges and bridges to move value.



A serious AML regime therefore intersects with consumer protection in practical ways:

- controls on account takeover and unauthorized access
- warnings and friction for high-risk transfers
- stronger checks for new payees and new withdrawal addresses
- rapid response to fraud reports

These are not only customer service features. They are crime controls.

2.6 Innovation, competition, and responsible growth

2.6.1 Innovation as a conditional objective

Many policymakers want financial innovation. Few treat it as absolute.

The usual stance is conditional: allow new products and new market structure, but only with safeguards that match risk. That stance is both political and practical. If rules are too strict, activity moves offshore or into informal channels. If rules are too light, fraud and failure erode trust and attract harsher backlash.

A responsible growth approach therefore often includes:

- sandboxes or limited pilots with clear limits
- staged authorization, where a firm can start with a narrow permission set and expand after it proves controls
- time-limited transitional arrangements for existing firms
- guidance channels and published expectations so firms can build toward compliance



These tools are not acts of leniency. They are ways to bring activity into view while it is still manageable.

2.6.2 Competition and market power concerns

Competition is a quieter objective, but it matters.

Crypto markets often show strong network effects. Liquidity tends to cluster where there are the most traders. Stable-value tokens can become default settlement units. Custody services can lock in users through friction and switching costs. If one firm or one small group controls key rails, the market can become dependent on private rule-making.

Competition concerns show up in several places:

- dominant exchanges that set fee terms and listing standards for the whole market
- stable-value issuers whose token becomes the default unit of account on many venues
- wallet providers and custodians that make it hard to move assets out
- affiliated market makers that get preferred access

A competition objective can justify rules that look like consumer protection rules: clear fee disclosure, limits on incentives, controls on conflicts, and portability of accounts and assets.



Investor and Consumer Protection in Crypto Markets

Harm Spectrum Map for Retail Users

 Market-Driven Harms	<ul style="list-style-type: none">• Volatility losses• Sudden price swings• Liquidity gaps
 Conduct and Information Harms	<ul style="list-style-type: none">• Market manipulation• Misleading claims• Unclear fees
 Operational and Custody Harms	<ul style="list-style-type: none">• Custody failure• Platform hacks• System outages
 User Error and Irreversibility	<ul style="list-style-type: none">• Irreversible transfer errors• Smart contract mistakes• Address input errors
 Redress and Recovery Gaps	<ul style="list-style-type: none">• Limited dispute resolution• No guaranteed recovery• Weak complaint handling

All rights reserved by the Blockchain Council.

It can also justify scrutiny of mergers and of group structures that concentrate too much control over market access.

2.6.3 Why innovation and protection are not opposites

It is common to talk as if protection and innovation sit on opposite ends of a scale. In practice, weak protection can crush innovation.

Markets that become known as scams do not attract serious builders or long-term users. Payment use does not grow if users fear they cannot redeem or that a token will freeze. Institutional involvement does not grow if custody and governance are weak.

A well-designed framework can support innovation by lowering uncertainty and by giving firms a clear path to operate legally.

2.7 A practical risk model for crypto supervision



Objectives tell regulators what they want. A risk model tells them where to focus.

A useful risk model for crypto has to do two jobs at once.

First, it has to be functional. It should start with activity types rather than token labels. Token labels are easy to game. Activities are harder to hide.

Second, it has to be layered. Crypto risk is not only in a single firm. It can sit in a market structure, in a protocol dependency, or in a cross-border group.

A practical model therefore treats risk at four levels:

1. Activity level: what function is being performed
2. Entity level: governance, resources, controls, and resilience of the firm
3. Market structure level: concentration, role combination, and fragmentation
4. System level: links to traditional finance, payment rails, and macro channels

Each level adds context. A custody activity carries more risk when performed by a thinly capitalized firm with weak controls. A trading venue carries more integrity risk when it also runs a dealer desk and sets rules for liquidations. A stable-value arrangement carries more stability risk when it becomes widely used for payments.

2.7.1 Activity-level risk: what can go wrong in each function

An activity-based risk map can be built around the main functions that show up across crypto markets.



Issuance and distribution. Main risks: false disclosure, mis-selling, insider allocations, market manipulation around listings, and conflicts where promoters sell into the hype.

Trading venue operation. Main risks: manipulation, weak surveillance, unfair access, conflicts from house trading, and fragile systems that fail during stress.

Brokerage and order handling. Main risks: steering, hidden fees, poor execution, mis-selling, and conflicts with referral and affiliate programs.

Custody. Main risks: theft, loss of keys, commingling, misuse of client assets, and freezes or outages that block customer exits.

Transfer services. Main risks: fraud and irreversible error, laundering, sanctions breaches, and operational errors.

Lending and borrowing products. Main risks: maturity mismatch, rehypothecation, liquidity runs, and hidden exposure to market shocks.

Staking services. Main risks: unclear disclosure of lockups and slashing, custody of staked assets, conflicts in validator selection, and packaging of technical network functions as yield promises.

Stable-value issuance. Main risks: reserve shortfall, redemption failure, run dynamics, and operational failure of issuance and redemption rails.

Each function can be present in a centralized firm, a distributed protocol, or a blend. The risk model therefore focuses on where the function is offered "to others" and where someone has control.

2.7.2 Entity-level risk: governance, controls, and resilience



Even when two firms offer the same activity, their risk differs.

Entity-level risk factors include:

- governance quality: who makes decisions, how conflicts are handled, and whether compliance has independence
- financial resources: ability to absorb losses, cover claims, and fund controls
- internal controls: segregation of duties, access controls, audit trails
- operational resilience: uptime, change management, incident response
- group structure: reliance on affiliates, related-party transactions, and intra-group guarantees
- culture: willingness to comply, respond to supervisors, and remediate issues

Crypto markets have shown that weak governance is often the root cause of failures. Losses are not always caused by code bugs. They are often caused by human decisions: taking customer assets, hiding risk, delaying redemptions, or running a market with conflicts unmanaged.

A risk-based supervisor therefore asks for evidence. Policies alone are not enough. The supervisor wants proof that controls work: access logs, incident reports, audit results, and records of how decisions were made.

2.7.3 Market structure risk: concentration, role combination, and fragmentation



Market structure can raise risk even when individual firms look compliant.

Concentration risk appears when a small number of venues or custodians handle most volume. A failure at one of them becomes a market-wide event. Concentration can also exist in stable-value tokens used for settlement.

Role combination risk appears when a single group combines trading venue, broker, dealer, custodian, and lender roles. In such a setup, conflicts can be managed, but they require heavy governance and close supervision.

Fragmentation risk appears when activity is scattered across borders and across lightly supervised venues. Fragmentation increases the payoff of jurisdiction shopping and weakens enforcement.

A good risk model treats these features as multipliers. They do not create risk alone. They magnify it.

2.7.4 System-level risk: spillovers to core finance and payments

System-level risk is about channels.

The most common channels are:

- stable-value token arrangements that interact with banking and payment rails
- large-scale borrowed trading that forces asset sales and price spirals
- links to banks through custody, lending, and settlement accounts
- links to funds through holdings and derivative exposures



- operational dependencies on a small number of service providers such as custody tech, price feeds, and infrastructure providers

A system-level lens asks: if this fails, does it stay contained, or does it hit payment flows, bank balance sheets, or core market infrastructure?

2.7.5 Core risk categories that recur across regimes

Across agencies, the language differs, but the risk categories recur.

Conduct and integrity risk. Fraud, manipulation, conflicts, poor execution, and misleading promotions.

Consumer and investor harm risk. Poor disclosure, mis-selling, custody loss, and lack of redress.

Prudential and stability risk. Liquidity stress, runs, collateral spirals, and interconnectedness.

Operational and technology risk. Cyber incidents, code vulnerabilities, key management failure, outages, and third-party dependence.

Governance and accountability risk. Unclear responsibility lines, weak control environment, and opaque group structures.

Illicit finance and national security risk. Weak AML/CFT programs, sanctions exposure, fraud proceeds laundering.

These categories overlap. A custody failure can be operational risk, consumer harm, and illicit finance risk at once. A stable-value token can be a consumer product and a stability concern. A trading venue can be a market integrity concern and an illicit finance conduit.

A risk model is valuable because it forces prioritization. A supervisor cannot inspect everything at once. It must decide what to inspect first.



2.7.6 Crypto-specific multipliers that change speed and severity

Crypto does not invent every risk. It changes how risks behave.

Several features make harm faster or harder to contain.

Irreversible transfers. Many transfers cannot be reversed. That makes mistakes and fraud harder to fix.

Custody concentration. A single key compromise can drain a large amount of customer assets. Custody is therefore a single point of failure.

Upgradeability and governance. Protocols can change. Admin privileges can change parameters. A governance vote can alter risk in ways retail users do not track.

Composability and dependency chains. Protocols depend on other protocols. A failure in one can cascade into another through shared collateral and shared price feeds.

Token distribution and unlocks. Concentrated holdings and cliff unlocks can create sudden sell pressure and manipulation risk.

Cross-border reach. Firms can serve customers in places where they have no local presence, making supervision and redress harder.

A risk model that ignores these multipliers will understate harm.

2.8 Risk-based supervision and proportionality across jurisdictions

Risk-based supervision is how objectives become real oversight. It is also how proportionality becomes more than a slogan.



The core idea is simple: focus attention and resources where risks are higher, and apply lighter obligations where risks are lower and can be shown to be lower.

This approach shows up in both financial regulation and AML supervision. It shows up in how regimes choose to license trading venues first, then expand into other services. It shows up in tiered permissions and thresholds.

2.8.1 Proportionality mechanisms in crypto regimes

In practice, proportionality shows up through several common mechanisms.

Tiered authorization and scoped permissions. A firm may be allowed to provide custody but not trading. Or it may be allowed to broker orders but not hold client assets. Permissions allow supervisors to match obligations to the risk of the service.

Activity-specific requirements. Stable-value issuance tends to carry heavier reserve and redemption requirements than other token issuance. Custody tends to carry heavier safeguarding requirements than pure software services.

Threshold-based intensification. Requirements often tighten as a firm reaches certain scale, as retail exposure grows, or as a token becomes widely used for payments.

Product-based restrictions. Higher-risk products such as high-margin derivatives and complex yield products are often subject to stricter rules or retail limits.

Group and conflict controls. Where a firm combines multiple roles, supervisors often require stronger governance, more reporting, and clearer separation of functions.



Supervisory tools that scale with risk. Higher-risk firms face more frequent reporting, more inspections, and deeper reviews.

None of these tools is perfect. Together they create a way to supervise a fast-changing market without either crushing it or ignoring it.

2.8.2 The political economy of proportional oversight

Proportionality is not only technical. It is also political economy.

Rules that are too strict can push activity into less visible channels, which can increase fraud and illicit finance risk. Rules that are too loose can attract scams and failures, which can also push activity into less visible channels after trust collapses.

A proportional framework gives regulators a way to say: activity is allowed, but the burden rises with the risk and the scale.

That is also why proportionality often includes clear triggers. Firms can plan. Supervisors can explain why they tighten rules. The public can understand why a large stable-value issuer is treated differently from a small software project.

2.8.3 Patterns in how regimes roll out oversight

Although the details differ, rollouts often follow patterns.

Some jurisdictions begin with AML registration and promotion controls, then move toward full licensing of trading venues and custodians.

Some begin with licensing of trading venues because retail access is the main harm channel.

Some begin with stable-value arrangements because payment use is the main concern.



These choices are often shaped by legal constraints as much as policy. It is easier to extend AML duties than to rewrite securities law. It is easier to regulate promotions than to supervise offshore venues. A risk model helps explain why these step-by-step choices still fit a coherent objective set.

2.9 Supervisory toolkit: turning risk models into enforcement reality

A chapter on risk-based supervision needs to describe tools, not only principles.

Supervisors rely on a mix of off-site and on-site work, plus enforcement and cooperation.

2.9.1 Off-site supervision

Off-site supervision is the ongoing stream of information that allows a supervisor to see risk building before it becomes a failure.

In crypto, off-site information often includes:

- periodic financial statements and capital or liquidity metrics where required
- client asset reports: holdings, segregation methods, and reconciliation results
- incident reporting: hacks, outages, key compromise events, failed transfers, and material changes
- market information for venues: volumes, concentration, suspicious trading alerts, and listing changes
- margin and liquidation metrics where relevant: number of forced closures, stress events, and changes to margin parameters



- AML program metrics: customer risk profiles, monitoring alerts, sanctions hits, and remediation timelines
- third-party reliance: key vendors, custody partners, banking partners, and critical infrastructure dependencies

The goal is not to drown the supervisor in reports. The goal is to obtain enough information to triage. A risk-based supervisor uses off-site information to decide where to go on-site.

2.9.2 On-site supervision

On-site supervision is where a regulator tests whether controls are real.

In crypto firms, on-site work often focuses on:

- governance: board oversight, decision-making process, and independence of compliance and risk functions
- custody controls: key management, access controls, segregation, withdrawal processes, and audit trails
- conflicts management: related-party trading, house activity, token allocations, and fee design
- customer outcomes: complaints handling, customer support capacity, and dispute resolution process
- AML program: quality of customer checks, monitoring scenarios, escalation, training, and audit findings
- resilience: incident response drills, change management practices, and business continuity plans

On-site supervision also looks at group structure. Many crypto firms operate through multiple entities. A supervisor needs to know which



entity holds assets, which entity contracts with customers, and which entity makes key decisions.

2.9.3 Perimeter work and enforcement posture

A risk model is only as strong as enforcement.

If a jurisdiction has rules but cannot enforce them against offshore actors serving local customers, the rules may mainly affect the firms willing to comply. That can create unfair competition and can push customers toward weaker venues.

Supervisors therefore use a mix of tools:

- licensing and authorization gatekeeping
- enforcement against unauthorized activity, including local marketing and local affiliates
- sanctions for breaches of custody and conduct duties
- public warnings and restrictions
- cooperation with other regulators and law enforcement

Perimeter work is not only a legal exercise. It is a risk control. It determines whether high-risk activity stays outside supervision.

2.9.4 Cross-border cooperation as a core tool

Crypto markets are cross-border. Supervision that stops at national borders will miss group risk, market abuse that jumps venues, and flows that cross service providers.

Cooperation tools include:

- information sharing on firms and group structures
- joint investigations of fraud and manipulation



- shared approaches to licensing and fit-and-proper checks
- coordination on crisis response for large stable-value arrangements

Cooperation is hard. Legal barriers, data protection rules, and different definitions get in the way. Still, without cooperation, regimes invite jurisdiction shopping.

2.10 Closing synthesis

Crypto regulation is often described as chaotic because many agencies are involved. The better description is that crypto regulation is multi-objective.

The objective set is fairly stable across jurisdictions: fair markets, protection of customers, containment of spillovers, control of illicit finance, and room for responsible market development.

The risk logic that turns those objectives into rules is also fairly stable. Identify the function, map it to what can go wrong, and apply outcomes comparable to traditional finance where the risks match. Adjust controls for crypto-specific risk drivers such as irreversible transfers, custody concentration, code change risk, and cross-border reach.

Risk-based supervision and proportionality are the methods that make this workable. They allow supervisors to focus where harm is most likely and most severe, while leaving lower-risk activity under lighter obligations.

Stable-value arrangements sit close to payments and can create run dynamics, so they tend to draw heavier oversight. Trading venues and custodians sit at the main retail harm points, so they tend to be licensed and closely supervised. AML duties often reach first because



they can be applied through service provider definitions and control points.

The final test of any framework is not only the text of the rules. It is whether supervision is real, enforcement is credible, and cross-border cooperation is strong enough to prevent the market from escaping the perimeter whenever it is convenient.

3) Global standards and coordination (the “baseline” layer)

Crypto regulation is often described as fragmented. That description is partly true in the narrow sense that rulebooks differ across borders, agencies argue about perimeter, and enforcement strength varies. But the word “fragmented” hides a more important structural fact: most jurisdictions are not starting from zero. They are building domestic rules on top of an existing layer of global expectations that already shapes how financial markets are policed.

This baseline layer is not a single global statute. It is a stack of shared expectations set by standard-setting bodies and reinforced through peer review, mutual evaluation, and routine cooperation between supervisors and law enforcement. The stack is “soft law” in form, yet hard in effect. Countries are judged against it. Firms are judged through it. Financial institutions use it to decide whether to do business with each other. Domestic regulators use it as a reference point when they draft or revise national rules.

In practice, the baseline layer does four things.

First, it defines minimum outcomes that jurisdictions are expected to achieve. A country does not need to write the same words as its peers, but it is expected to achieve a recognizable result: gatekeepers must apply customer checks; market operators must deter manipulation;



large stable-value arrangements must not operate without meaningful reserve and redemption controls.

Second, it creates convergence pressure even without treaty-level enforcement. Peer review, public comparison, and the reputational cost of being tagged as weak all push jurisdictions toward similar outcomes.

Third, it shapes domestic perimeter decisions. When lawmakers argue about whether a token is a security or whether a protocol is “decentralized,” they often end up answering a different question: what control points must sit inside supervision if the country wants to meet baseline expectations on financial crime, market conduct, and stability.

Fourth, it provides cooperation rails. Regulators do not need to invent cross-border tools from scratch. They can reuse, extend, and adapt tools built for banking, securities markets, and financial intelligence work: memoranda of understanding, supervisory colleges, information exchange channels, and joint work plans.

This chapter focuses on three baseline pillars that dominate most crypto policy debates.

- The anti-money laundering and counter-terrorist financing baseline, set through global financial integrity standards.
- The investor protection and market integrity baseline, shaped through securities regulator coordination.
- The financial stability baseline, shaped through a framework for crypto activities and stable-value arrangements.



The chapter then turns to the coordination problem itself: how baseline standards are made operational across borders, where the hardest gaps remain, and what tools are being used to close them.

3.1 The baseline architecture: who sets standards, and why it matters

International financial regulation is not run by a global legislature. It is shaped by a network of standard-setting bodies, forums, and institutions. Each body has a different mandate and a different set of members. Each produces a different kind of output: guidance, principles, recommendations, and peer review reports. Together they create a baseline that many domestic regimes follow.

For crypto-assets, three bodies dominate the baseline conversation.

The financial integrity standard setter. This body sets the minimum expectations for anti-money laundering and counter-terrorist financing controls. It defines who must sit inside the AML/CFT perimeter, what customer checks are expected, and how information should follow value transfers. It also runs mutual evaluations and follow-up processes that rate implementation.

The securities regulator coordination body. This body frames crypto markets as markets in which investors can be harmed through fraud, conflicts, custody failure, and market abuse. It focuses on outcomes that securities regulators care about: fair dealing, orderly markets, custody safeguards, and controls that deter manipulation.

The financial stability forum. This forum focuses on how crypto activities can create spillovers: runs, forced selling, links to banks, and payment system adjacency through stable-value arrangements. It issues high-level recommendations and uses peer reviews to identify gaps.



Other bodies are also part of the baseline stack, even when they are less visible in public debate.

- Banking supervisors set standards for how banks should treat crypto exposures, including how those exposures appear in capital planning and risk limits.
- Payment and market infrastructure bodies shape expectations for settlement, custody, and operational resilience.
- Macroeconomic institutions track cross-border spillovers, capital flow issues, and the interaction between crypto activity and monetary and financial conditions.
- Tax cooperation work sets expectations for how crypto activity should be reported for tax compliance.
- Financial intelligence cooperation networks create secure rails for cross-border exchange of information about suspicious activity.

A research paper can treat the baseline architecture as a set of overlapping circles. One circle is financial crime controls. One is market conduct. One is stability. The overlap matters because the same firm may face all three at once.

3.1.1 Why “soft law” is hard in practice

It is tempting to dismiss global standards as non-binding guidance. That underestimates their force.

The baseline layer shapes domestic law because governments do not want to be publicly rated as weak. A poor rating on AML/CFT implementation can affect the willingness of foreign banks to offer correspondent services. Weakness in market conduct can affect



investor confidence and the willingness of firms to list or operate locally. Weakness in stability oversight can affect macro credibility when a large stress event occurs.

The baseline layer also shapes firm behavior because large cross-border firms do not want to operate under a patchwork of incompatible expectations. They tend to build to the highest common requirement among the markets that matter to them. A firm that wants access to large pools of customers and banking rails will often adopt baseline controls even before a local law requires them.

This is not altruism. It is risk management. A cross-border firm wants to avoid being cut off from banks, payment processors, and institutional counterparties. Those counterparties will often require evidence of compliance with baseline expectations.

3.1.2 How the baseline layer is enforced without a global police force

Global standards become real through several channels.

Peer review and mutual evaluation. Countries are assessed against shared criteria. The assessments become public and create political pressure. They also become practical inputs for other decisions, such as whether financial institutions treat a jurisdiction as high risk.

Market access and de-risking pressure. Banks and payment firms respond to perceived regulatory weakness by reducing exposure. That can be blunt and can create its own policy issues, but it is part of why governments care about AML/CFT evaluations.

Domestic statutory hooks. Many domestic laws explicitly incorporate global standards by reference or use them as interpretive



guides. Even when not explicit, domestic agencies cite them as a rationale for guidance and enforcement.

Cooperation reliance. Regulators need each other. A market regulator that wants help with a cross-border case depends on another regulator's willingness and legal ability to share information. That willingness is shaped by trust. Trust is shaped by whether both sides meet baseline expectations on confidentiality, due process, and supervisory competence.

Reputational competition. Jurisdictions compete for legitimate activity. Being seen as a place with weak controls can attract scams and short-term volume, but it can also repel long-term institutional participation.

The net effect is that global baselines create convergence even without treaties.

3.1.3 Why the baseline layer is a special issue in crypto

Crypto markets are cross-border by design. A token can be issued in one place, traded on venues in several other places, held in custody by a firm somewhere else, and used as collateral in a protocol that has no clear physical base. If domestic rules differ widely, risk will tend to flow toward the weakest perimeter.

That is why baseline standards matter so much. They are the nearest thing to a shared floor.

At the same time, the baseline layer cannot remove all differences. Domestic law still matters, especially in areas where global bodies set high-level expectations but leave details to local choice: retail access rules, civil liability, product bans, and the boundary between securities law and crypto-specific regimes.

The baseline layer is therefore best understood as a common language that allows domestic regimes to coordinate, compare, and cooperate. It reduces the cost of building rules and the cost of working together, even when laws are not identical.

3.2 Global AML/CFT standards: the minimum floor for financial integrity

3.2.1 The AML baseline in one sentence

The AML/CFT baseline for crypto is direct: jurisdictions must identify and mitigate laundering and terrorist financing risks linked to crypto activity by bringing relevant actors inside a regulated perimeter, requiring preventive controls, supervising compliance, and enabling domestic and cross-border cooperation.

This baseline is built around two definitional moves.

First, it defines a broad category of transferable digital value used for payment or investment.

Second, it defines a category of service providers that perform exchange, transfer, custody, and issuance-related services as a business.

Once those definitions apply, a familiar control stack follows: customer due diligence, ongoing monitoring, recordkeeping, suspicious reporting, sanctions controls, and information that accompanies transfers.

3.2.2 Technology-neutral definitions and the “as a business” line

A central design choice in the AML baseline is that definitions are meant to be function-based, not chain-based.



If a service allows customers to exchange between fiat and crypto, exchange between crypto assets, transfer crypto for others, or hold crypto on behalf of others, it is treated as part of the gatekeeping system. The definition is meant to stop technical form from becoming a loophole.

The hardest interpretive line sits in the phrase “as a business.” That phrase tries to separate ordinary users from professional service providers. In crypto, it also forces regulators to define what counts as a service, not merely code.

A user who sends tokens from a self-held wallet to another self-held wallet is not a service provider. A firm that holds customer keys and transfers tokens on request is.

The hard cases sit between.

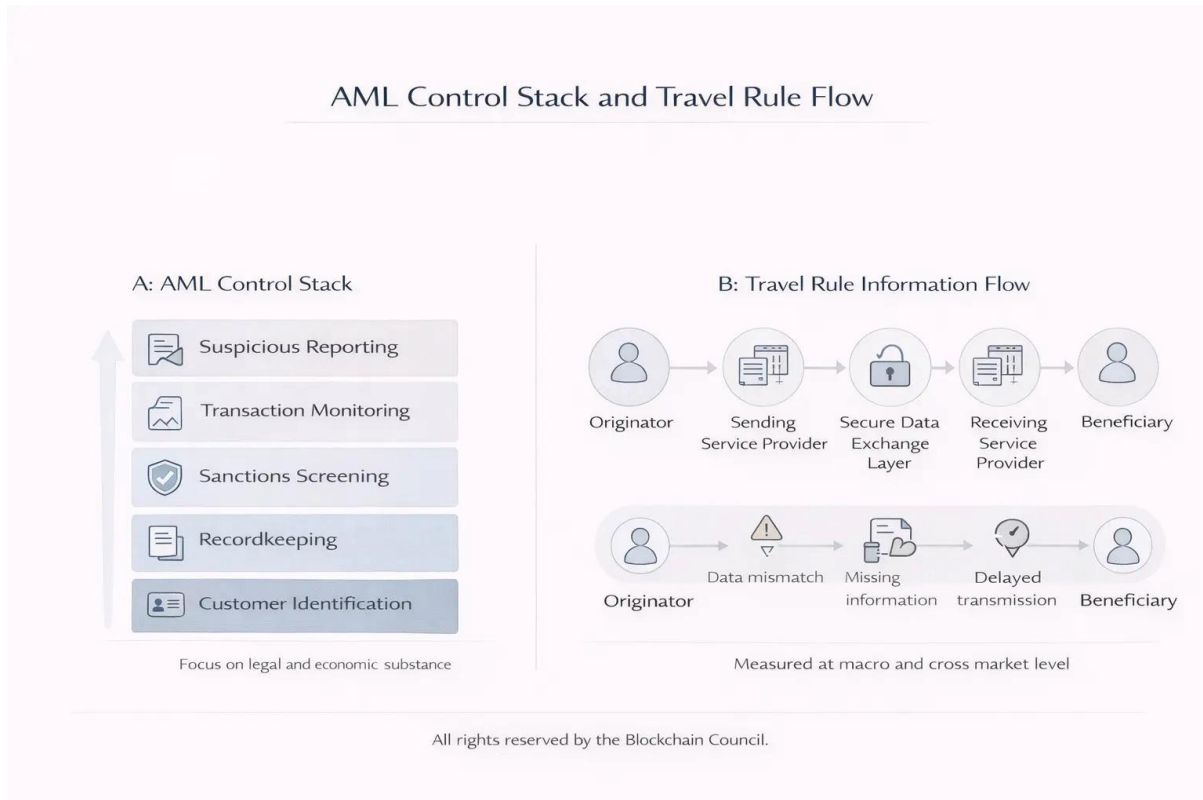
- A group that runs a user-facing front end for a protocol, collects fees, and controls upgrades looks more like a service provider.
- A developer who publishes open-source code and does not control its use looks less like a service provider.
- A governance group that can change fee settings, control admin keys, and direct treasury funds may fall somewhere in between depending on how much control it has in practice.

The baseline approach pushes countries to look for control points. Where there is an actor that can apply checks and monitoring, authorities are expected to bring that actor inside the perimeter.

This is one reason the AML baseline has become a key driver of perimeter expansion. Even jurisdictions that are hesitant to build full market conduct regimes often implement AML registration and supervision for crypto service providers.



3.2.3 The AML control stack expected for crypto service providers



An AML law is not only a statute that declares obligations. It is a stack of policies, procedures, technology, supervision, and enforcement that produces usable outcomes.

A mature control stack includes at least the following components.

Risk assessment and a risk-based program. Firms are expected to assess their risks and apply stronger controls where risk is higher. A firm serving retail users through bank transfers and offering fast withdrawals faces different risk than a firm serving institutional clients with strict onboarding and lower-speed transfer limits.

Licensing or registration. The baseline expects jurisdictions to require service providers to be licensed or registered, so supervisors can identify who is in the market and can enforce standards.



Fit-and-proper checks and ownership clarity. A registration regime that does not identify beneficial ownership or the people controlling a firm provides little protection.

Customer due diligence. Firms should identify customers, understand control for entities, and apply higher scrutiny to higher-risk cases. In crypto, this often includes understanding whether the customer is acting for themselves or as a business.

Ongoing monitoring. Firms should monitor activity for patterns consistent with laundering, fraud proceeds movement, or sanctions evasion. That includes monitoring at account level and, where feasible, monitoring linked on-chain activity.

Recordkeeping and audit trails. Supervision depends on records. Without them, firms can claim compliance without proof.

Suspicious reporting. Firms should report activity that appears linked to crime.

Sanctions controls. Firms should screen customers and, where feasible, apply controls to prevent service to sanctioned parties and to stop transfers linked to sanctioned wallets.

International cooperation readiness. Firms should be able to respond to lawful requests and to preserve records.

Each component has a crypto-specific twist.

- Transfers can be final and can occur at any time.
- Wallet addresses are not names.
- Customers can move assets across chains through bridges.
- Some activity bypasses custodians through self-held wallets.



The baseline does not require supervisors to control every peer-to-peer transfer. It requires supervisors to control the regulated gatekeepers and to apply extra measures when gatekeepers face higher risk.

3.2.4 The Travel Rule: a baseline requirement that turns into a systems problem

The Travel Rule is the clearest example of a global expectation that forces operational coordination.

At its core, the rule requires that certain information about the sender and recipient travel along with a value transfer between service providers. The policy goal is old: payment transparency. The implementation is new because crypto transfers do not carry identity information by default.

In practice, Travel Rule compliance requires firms to do several things at once.

- Collect required information from the originator.
- Validate or at least make reasonable checks on information quality.
- Send that information to the recipient institution or service provider.
- Receive information when acting as the recipient.
- Apply controls when information is missing or does not make sense.
- Retain records and make them available to supervisors.

Two implementation realities follow.



First, it requires interoperability. A Travel Rule program is only as strong as the ability of one firm to exchange the required data with another. If a firm cannot identify the counterparty service provider, or cannot send information in a mutually usable format, compliance becomes uneven.

Second, it requires supervision. Legislation alone does not produce compliance. Supervisors must test whether firms are actually exchanging data, whether data quality is acceptable, and whether firms apply controls when counterparties fail to comply.

Travel Rule implementation also raises design questions that are not present in older wire transfer systems.

- How should service providers identify whether a counterparty is a service provider or a self-held wallet?
- How should firms handle transfers that involve multiple hops or that route through aggregators?
- How should firms handle cross-border transfers where the counterparty is in a jurisdiction that has not implemented the rule?
- How should data be protected when it includes personal information and moves across borders?

Because of these issues, Travel Rule implementation has become partly a coordination exercise and partly an engineering exercise.

One recent global update on implementation reported that among jurisdictions that did not prohibit crypto service providers, a large majority reported passing Travel Rule legislation. The same update also observed that enforcement experience remained limited. In other words, many countries have adopted the rule on paper, while fewer



have tested it in the field or taken action against weak implementation.

For research purposes, this matters because it shows a gap between adoption and effect. It also shows why cross-border alignment matters: a rule that depends on counterparties cannot work well when only half the counterparties are in scope.

3.2.5 Common Travel Rule implementation models

Although laws are written at a high level, firms tend to implement Travel Rule compliance through a small number of operational models.

Bilateral exchange. Firms agree to exchange required information directly. This is workable when a firm has a small number of counterparties and a stable set of relationships. It becomes less workable when the firm's customer base sends transfers to a long tail of destinations.

Shared messaging networks. Firms join a network that provides a common format and routing layer for messages. This model improves reach but creates questions about governance of the network, confidentiality, and reliance on a third party.

Directory and counterparty identification tools. Some systems focus on solving the “who is the counterparty” problem. The firm can identify whether the recipient address belongs to a known service provider and can route required information accordingly.

Risk-based gating. Some firms apply stricter controls to outbound transfers when they cannot exchange information with the counterparty. Controls include transfer delays, enhanced customer checks, and restrictions on high-risk destinations.

Each model interacts with domestic law on privacy and data transfer. A law that restricts cross-border transfer of personal information can slow Travel Rule implementation unless a legal basis exists for compliance-related transfers.

The baseline layer therefore pushes countries to coordinate not only on rules but also on data governance.

3.2.6 Unhosted wallets, peer-to-peer transfers, and the limits of entity-based supervision

The AML baseline is centered on service providers. That is a strength because supervisors can enforce obligations on firms. It is also a limit because some crypto activity occurs without intermediaries.

Self-held wallets allow users to transfer value directly, without a firm controlling the transfer. Some protocols allow users to exchange assets directly through contract code, without a custodial intermediary.

These realities create two policy pressures.

First, supervisors push stronger controls onto gateways: fiat on-ramps, custodians, and venues where users exchange between fiat and crypto. If peer-to-peer activity cannot be directly supervised, the next best control point is the point where value enters and exits the regulated sector.

Second, supervisors expect regulated firms to manage exposure to self-held wallets through risk-based controls. That can include collecting information about the owner of a self-held wallet in certain cases, applying transfer limits, and monitoring patterns that suggest laundering.



There is no single global answer for unhosted wallet exposure. But the baseline expectation is that jurisdictions should not ignore it. They should assess the risk and require reasonable controls from the regulated sector.

3.2.7 DeFi and the control-point approach

Decentralized finance is often presented as an obstacle to AML rules because it can operate without a traditional intermediary. The baseline layer does not accept that framing as a reason for inaction. It treats DeFi as a perimeter stress test.

The key question is the same: is there an actor providing a covered service as a business?

If there is an actor that runs a front end, controls admin keys, collects fees, markets the service, or can change the system, regulators will ask whether that actor should be treated as a service provider.

If there is no such actor, regulators often shift focus to the gateways: the exchanges and custodians that allow users to acquire and cash out assets used in the protocol.

A recent global survey-based update reported that only a small number of jurisdictions had in practice registered or licensed DeFi entities as service providers, and only a small number reported enforcement actions against DeFi entities that qualify as service providers. That is a data point that matters because it suggests the baseline is clear, but implementation at the frontier remains limited.

The reasons are practical.

- Some DeFi systems are hard to map onto entity-based licensing.



- Some jurisdictions lack legal tools to treat a front end or governance group as a regulated firm.
- Some supervisors lack resources to investigate and prove control.
- Some projects can relocate quickly or distribute operations across borders.

These limits do not remove the baseline expectation. They help explain why implementation gaps persist.

3.2.8 Implementation gaps as a baseline feature, not an exception

Global AML/CFT work does not assume perfect implementation. It is built around the idea that implementation will be uneven and must be monitored.

Updates and evaluations repeatedly point to gaps.

- Many jurisdictions remain only partially compliant with crypto-related AML expectations.
- Many struggle to identify unregistered service providers operating domestically.
- Many have passed Travel Rule legislation but have limited enforcement activity.
- DeFi perimeter work remains thin.
- A meaningful number of jurisdictions do not respond to surveys or provide evidence of progress, suggesting capacity or priority gaps.

For a research paper, the right conclusion is not that the baseline failed. The right conclusion is that the baseline is a floor that



jurisdictions reach at different speeds. That speed difference creates arbitrage risk.

It also creates a policy problem for jurisdictions that do implement: how to stop their regulated firms from losing business to weakly supervised offshore competitors.

That question leads directly to the coordination tools discussed later in this chapter.

3.3 Securities regulator standards: investor protection and market integrity outcomes

The market conduct baseline for crypto is shaped through the coordination of securities regulators. The core argument is that many crypto markets recreate the same basic risks seen in securities markets: fraud, manipulation, conflicts, misuse of customer assets, and weak disclosure.

Even when a token is not classified as a security, the way it is sold and traded can create securities-like harm. The baseline response is to focus on activities and outcomes.

3.3.1 Outcomes first: why form is a weak guide in crypto markets

Crypto projects can change labels quickly. A token can be called a “utility” token even when it is sold mainly as a speculative bet. A venue can claim it is a technology platform rather than a market operator. A lending product can be packaged as a simple “earn” program rather than a credit product.

A market conduct baseline that relies on labels will be easy to evade. That is why securities regulator coordination work tends to stress outcomes: fair dealing, conflict control, custody safeguards, and deterrence of market abuse.



This does not mean every crypto activity should be regulated as a securities market. It means regulators should not accept a label as a substitute for analysis.

A research paper can treat this as a perimeter principle: treat the substance of the activity as the starting point, then apply controls that match the risk.

3.3.2 The control domains that keep appearing in securities regulator work

Securities regulator coordination work in crypto has produced a set of policy recommendations that cover a broad set of control domains. Rather than list recommendations as a checklist, it is more useful to group them into domains that map to harm channels.

Governance and accountability. Regulators expect firms that run markets, hold assets, and sell high-risk products to have clear governance: responsible people, clear decision-making, internal controls, and the ability for supervisors to hold someone accountable.

Conflicts of interest. Crypto firms often combine roles that create conflicts: operating a venue while also trading on it, holding custody while also lending, listing tokens while receiving allocations, selling products while also acting as principal. Regulators expect conflicts to be managed through a mix of structure, rules, and disclosure.

Custody and client asset safeguards. Custody failures have been a major driver of losses. A market conduct baseline treats custody and segregation as central investor protection outcomes.

Disclosure and communications. Regulators expect customers to receive information that allows them to understand what they are buying, what rights they have, and what risks they take. They also



treat marketing as part of market conduct. A market can be “orderly” in a narrow sense and still be distorted by misleading promotions.

Market abuse controls. The baseline expects controls that deter manipulation, insider dealing, and other abusive conduct. That includes surveillance, investigation capability, and coordination with law enforcement.

Market operation and trading controls. Trading rules, fee clarity, outage handling, listing standards, and order execution practices sit inside the market integrity frame. These are not mere operational details. They shape who wins and who loses.

Operational resilience and cyber controls. Because crypto markets are technology-heavy, resilience is a conduct issue. Outages can trap customers. Breaches can drain assets. Weak change management can alter market rules without warning.

Retail protection for high-risk products. Where markets offer leverage, complex derivatives, or yield products with embedded risk, regulators expect extra safeguards: risk warnings, limits, appropriateness checks, and product governance.

These domains map to recurring failure modes in crypto markets. That is why they keep appearing across jurisdictions even when statutes differ.

3.3.3 Vertical combination of roles as the central conduct stress point

A key difference between many crypto market structures and traditional securities markets is vertical combination. In many crypto groups, one firm or a set of affiliates can be all of the following at once:



- a market operator
- a broker
- a dealer
- a custodian
- a lender
- a stable-value issuer
- a staking service

Traditional finance can contain some combinations, but it usually surrounds them with strict conflict rules, capital rules, and oversight that assumes conflicts are present.

In crypto markets, vertical combination has often been treated as normal business structure rather than a risk factor. The market conduct baseline challenges that assumption. It treats vertical combination as a risk multiplier.

This matters for coordination because a global firm can move roles across borders. It can locate the venue in one jurisdiction, custody in another, and the trading desk in a third. Without cross-border coordination, no single supervisor sees the full conflict map.

3.3.4 Listing standards and token admission as a baseline issue

Token admission is often treated as a private business decision. In crypto markets, it becomes a public policy issue because listings shape retail exposure and because listing events can be tied to manipulation and insider abuse.

A market conduct baseline therefore pushes for a disciplined listing process.



A research-friendly way to describe this is to treat token admission as a risk gate with four main questions.

1. **Information quality.** Is there enough public information to allow users to understand supply, governance, upgrade control, and major risks?
2. **Concentration and unlock risk.** Are holdings concentrated? Are there large unlocks that could crash price and harm retail users?
3. **Code and admin risk.** Does the token rely on admin keys or upgrade mechanisms that allow a small group to change core behavior?
4. **Manipulation susceptibility.** Is liquidity thin? Is the token prone to wash trading? Are there known patterns of insider allocation and dumping?

Listing discipline connects to investor protection and market integrity at the same time. It also connects to coordination because a token listed on one major venue can quickly become a global retail product.

3.3.5 How the baseline deals with different domestic classifications

Domestic law still differs on classification. Some jurisdictions treat many tokens as securities. Some treat most as outside securities law unless they look like shares or funds. Some create a crypto-specific regime for non-security tokens.

A market conduct baseline does not require every jurisdiction to adopt the same classification. It pushes for a common set of outcomes where the activity creates comparable risk.

That is why an activity-based framing is useful for research.



- If a firm operates a trading venue for retail users, it should face baseline expectations for fair dealing, surveillance, conflict control, and resilience.
- If a firm holds client assets, it should face baseline expectations for segregation, key security, and recordkeeping.
- If a firm sells high-risk leveraged products, it should face baseline expectations for clear disclosure, product controls, and oversight.

This framing allows jurisdictions with different statutory definitions to still move toward similar market outcomes.

3.3.6 Implementation progress and the recurring gaps

The market conduct baseline has also moved from recommendations to implementation reviews. Those reviews tend to tell the same story.

Progress exists. Many jurisdictions have begun licensing trading venues or applying registration systems. Many have issued guidance on custody standards, conflict policies, and risk disclosures.

At the same time, recurring gaps remain.

- Cross-border information sharing remains uneven.
- Regulators face legal limits on sharing non-public information.
- Definitions differ enough to slow cooperation.
- Offshore provision allows firms to reach customers without local authorization.
- Data collection is inconsistent, making it hard to compare markets and detect group-wide risk.



In other words, adoption is moving. Coordination and enforcement remain the hard part.

3.4 Financial stability standards: the system-wide lens

The financial stability baseline for crypto is shaped through a framework that focuses on system-wide spillovers. It treats crypto markets as part of the wider financial system when they become large, interconnected, or tied to payment rails.

This baseline matters even in jurisdictions where crypto is not yet large. The reason is simple: crypto markets can grow quickly, and cross-border firms can import stress.

3.4.1 The stability problem is about channels, not labels

A stability lens starts with a question: can crypto activity create problems that spill beyond the original users?

Volatility alone does not answer that question. Many assets are volatile. What matters is the channel.

Common channels include:

- leveraged positions that force sales during price drops
- liquidity mismatch in lending and borrowing products
- stable-value token runs that force reserve asset sales
- links to banks and funds through custody, lending, and investment exposures
- payment and settlement reliance on stable-value tokens
- operational failures at large venues that trap customers and freeze market function



A stability baseline therefore pushes for functional oversight that matches these channels.

3.4.2 High-level recommendations and perimeter breadth

The stability baseline is set through high-level recommendations for crypto activities and markets.

Those recommendations share several themes.

Comprehensive perimeter. Authorities should not assume crypto sits outside regulation. Where activities create stability risk, the perimeter should expand to include them.

Functional regulation. Similar activities should face similar outcomes. If a product looks like credit, it should face credit-like oversight. If a stable-value arrangement looks like money issuance, it should face money-issuance style expectations.

Proportionality. Oversight should scale with risk, size, complexity, and system relevance.

Group-wide view. Supervisors should be able to understand group structure and exposures across affiliates.

Data and monitoring. Authorities should collect enough information to monitor market development, identify leverage, and see concentrated dependencies.

Cross-border cooperation. Authorities should cooperate because the same firm and the same stable-value token can affect many markets.

A useful research point is that stability work often leaves detailed consumer and market conduct rules to other bodies. That is not a gap; it is division of labor. The stability baseline is about system-wide monitoring and perimeter coverage.



3.4.3 Stable-value arrangements as the main stability bridge

Stable-value arrangements are the clearest bridge between crypto markets and the wider financial system.

A stable-value token that is widely used can become a unit of account inside crypto markets, a settlement asset, and a payment instrument for commerce. That creates expectations of reliability and quick redemption.

A stability baseline for stable-value arrangements focuses on three core functions.

Issuance and redemption. How tokens are created and destroyed, what rights holders have, and what assets back redemption.

Transfer and settlement. How tokens move between users and what operational systems support movement.

User access and storage. How users store tokens, how custodial and non-custodial solutions interact, and what happens during outages.

The stability baseline for stable-value arrangements tends to focus on the following control areas.

- Reserve assets: quality, liquidity, custody, segregation, and valuation.
- Redemption: clarity, timeliness, conditions, and what happens in stress.
- Governance: who controls reserve management, who can change rules, and how conflicts are managed.
- Operational resilience: uptime, incident response, and reliance on third parties.



- Disclosures: clear information on reserves, redemption terms, and risks.
- System relevance: extra expectations when an arrangement reaches scale.

The label “stable” is treated with caution. The baseline view is that the word can create false expectations. A stable-value arrangement should be judged by its design and controls.

3.4.4 Peer reviews and what they reveal about real coordination limits

Stability work also uses peer review to identify gaps.

Peer reviews tend to point to uneven implementation and to the limits of cross-border cooperation.

A repeated theme is that fragmented implementation creates arbitrage opportunities and complicates cross-border oversight.

Another repeated theme is that cooperation tools are often ad hoc and not designed for stability-related information sharing. Traditional cooperation tools are often designed for enforcement cases or for routine supervision in banking. Stability monitoring requires different data: leverage, liquidity dependencies, reserve exposures, and group-wide flows.

Peer reviews also point to data gaps.

- Authorities do not always collect comparable information about crypto market activity.
- Stable-value arrangements can be large, but reserve data may not be collected or shared.



- Group structures can be complex, and supervisors may lack the full map.

A research paper can use these observations to explain why stability oversight is harder to coordinate than AML cooperation. AML cooperation has long-established rails through financial intelligence units. Stability cooperation is less mature because it requires sharing supervisory data that is often confidential and because no single supervisor has a full picture of global exposures.

3.4.5 Interaction with banking and payment infrastructure standards

The stability baseline interacts with other standards.

Banking supervisors set expectations for how bank exposures to crypto should be treated. That includes how exposures are counted, how capital is planned, and how concentration is managed.

Payment and market infrastructure bodies shape expectations for settlement, resilience, and the handling of operational outages.

These bodies are not “crypto bodies,” but their standards shape crypto regulation when banks custody crypto, when payment firms support stable-value redemption, or when stable-value tokens are used in settlement.

3.5 The wider baseline stack beyond the three headline pillars

A full baseline picture includes more than AML, market conduct, and stability.

3.5.1 Banking supervision and capital treatment



When banks gain exposure to crypto, whether through custody, trading, lending, or investment, bank supervisors focus on two broad aims.

First, make sure exposures are measured and disclosed. Hidden exposures create surprise during stress.

Second, make sure capital and risk limits reflect the risk profile of crypto exposures.

This is not an attempt to ban crypto from banks by default. It is an attempt to avoid a repeat of past episodes where new asset exposures built up without a clear prudential view.

For crypto markets, bank supervision matters because banks provide critical rails: custody services, settlement accounts, and payment processing for fiat on-ramps and redemptions.

If banks pull back because of uncertain prudential treatment, crypto firms may lose access to core rails. If banks move in without proper controls, bank exposure can become an amplifier.

3.5.2 Payment systems and market infrastructure expectations

Crypto markets rely on infrastructure that can look like market infrastructure even when it is not labeled that way.

Trading venues match orders and produce prices. Custodians hold assets. Stable-value arrangements can function as payment instruments.

Payment and infrastructure standards bring a focus on operational resilience, settlement finality, and continuity under stress.



Even where a stable-value token is not formally designated as payment infrastructure, the same operational questions matter when it is widely used: what happens if transfer systems fail, if redemption rails break, or if a critical service provider goes offline.

3.5.3 Macroeconomic surveillance and policy coordination

Macroeconomic institutions look at crypto through a different lens.

They focus on cross-border flows, capital controls, currency substitution risk, and the interaction between crypto activity and monetary policy.

This layer matters most when stable-value tokens are used in ways that resemble dollarization, or when crypto activity becomes a meaningful part of financial access.

Even in jurisdictions where crypto is mainly speculative, macro institutions track spillover risk, especially when leverage and stable-value arrangements grow.

3.5.4 Tax reporting expectations

Tax compliance is another baseline layer that increasingly touches crypto.

Tax authorities care about reporting of crypto transactions and holdings, especially when crypto provides a way to hide taxable gains or to move value across borders.

Tax reporting standards do not replace market conduct rules or AML rules, but they add a separate pressure for service providers to keep records and share information.



From a coordination standpoint, tax reporting can reinforce AML recordkeeping and can create another reason for service providers to register and maintain a compliant customer base.

3.6 How domestic regimes build on the baseline layer

Global baselines shape domestic law, but domestic law still makes the core design choices.

A research paper can describe three common pathways by which countries build rules on top of the baseline.

3.6.1 Pathway one: fit crypto into existing categories

Some jurisdictions extend existing securities, payments, and banking rules to cover crypto where possible. Under this pathway, the main task is classification.

If a token is treated as a security, securities offering and trading rules apply.

If a stable-value token is treated as money-like, payments or e-money rules apply.

If a firm holds customer assets, custody and client asset rules apply.

This pathway can work well where existing statutes are flexible. It can also produce uncertainty when older definitions do not map cleanly onto newer designs.

3.6.2 Pathway two: create a crypto-specific market regime

Some jurisdictions create a tailored regime for crypto assets and crypto services. This is often done for the part of the market that does not fit neatly into securities or payments categories.



Under this pathway, lawmakers define crypto service types, set authorization requirements for those service types, and impose conduct and disclosure duties tailored to crypto markets.

The advantage is clarity. The downside is that it must be kept aligned with existing financial law so firms do not face inconsistent obligations.

3.6.3 Pathway three: hybrid perimeter with layered controls

Many jurisdictions end up with a hybrid model.

- AML registration and supervision applies to service providers.
- A crypto service authorization regime applies to market-facing services.
- Securities law applies to security-like tokens.
- Payments rules apply to money-like stable-value arrangements.

This layered model is common because it follows how baselines are layered. AML is often implemented first. Market conduct rules follow. Stability rules tighten as activity grows.

3.6.4 Baselines as drafting templates and as enforcement arguments

Baselines influence domestic regimes in two ways.

They influence drafting because lawmakers can borrow language and structure from global expectations. That speeds up legislation and reduces the risk of missing a key control area.

They influence enforcement because domestic agencies can cite global expectations as evidence of what “reasonable” controls look



like. Even when a statute is high-level, global baselines help agencies interpret and justify detailed guidance.

3.6.5 Domestic perimeter disputes that the baseline layer tends to settle

In many domestic debates, two issues dominate: whether DeFi is outside regulation, and whether non-custodial activity should be unregulated.

The baseline layer does not settle every detail, but it pushes debates toward control points.

- If a front end is run as a business and controls access, regulators ask whether it should be within perimeter.
- If a venue offers services to the public, regulators ask whether it should be authorized.
- If a stable-value arrangement invites money-like use, regulators ask whether reserve and redemption controls should apply.

This is why baselines often influence perimeter more than they influence specific product rules.

3.7 Cross-border coordination: why it is required and how it is done

Global baselines assume cross-border coordination because crypto markets are not domestic markets with a few foreign links. They are cross-border markets first.

Without coordination, firms can avoid oversight by shifting operations. Without coordination, a supervisor sees only part of a group's risk. Without coordination, enforcement actions become slow and easy to evade.



3.7.1 The three facts that drive coordination needs

Three structural facts keep bringing coordination back to the center.

Borderless service provision. A crypto service provider can serve users in one country while being authorized elsewhere or nowhere.

Contagious risk events. A failure at a major venue, custodian, or stable-value arrangement can spread quickly through liquidity and settlement dependencies.

Fragmented data. Without information sharing, regulators see only local fragments: a local marketing campaign, a local affiliate, a slice of customer accounts. The real risk map sits across the group.

3.7.2 Cooperation is not one tool but a toolkit

Cross-border cooperation is often described as a single thing, but it is a toolkit. Different tools support different objectives.

Securities regulator tools. These include memoranda of understanding, information exchange agreements, and alert networks that share information about firms that may be operating without authorization.

Financial intelligence tools. These include secure channels for financial intelligence units to share information about suspicious activity and to support investigations.

Supervisory colleges. These are structured groups of supervisors who share an interest in the same cross-border firm or arrangement. A college creates routine coordination rather than ad hoc requests.

Crisis coordination tools. These include shared contact lists, incident reporting protocols, and plans for handling a failure at a large firm or stable-value arrangement.



A mature coordination system uses several tools at once.

3.7.3 Memoranda of understanding and information exchange agreements

Memoranda of understanding (MoUs) are a standard tool in financial regulation. They set out how supervisors share information, how they protect confidentiality, and how they handle requests.

In crypto, MoUs matter because:

- firms operate across borders
- market abuse schemes cross borders
- supervisors need background on group structure and control

MoUs also matter because they reduce friction. A supervisor with an MoU in place can request information without negotiating terms from scratch each time.

MoUs are not a cure. They depend on the legal powers of each party. A supervisor cannot share information it does not have. A supervisor cannot share information if domestic law blocks sharing. MoUs work best when domestic law is written with cross-border cooperation in mind.

3.7.4 Supervisory colleges: from requests to routine coordination

A supervisory college is a structured approach to supervising a cross-border firm or arrangement.

A college typically includes:

- a lead supervisor, often where the main entity is authorized
- host supervisors in jurisdictions where key services are offered



- sometimes central bank or stability authorities when stable-value arrangements are involved

Colleges can:

- share non-public supervisory information under confidentiality rules
- coordinate inspection plans
- discuss group risk, conflicts, and dependencies
- align on key supervisory priorities
- coordinate response to incidents

Colleges help solve a common crypto problem: group structures are complex, and no single supervisor sees the full picture.

A stable-value arrangement college can also include supervisors who focus on reserve management, redemption rails, and payment linkages.

3.7.5 Alert networks and warning systems

Another coordination tool is the sharing of alerts about firms that appear to be operating without authorization or that pose high risk.

Alert networks can help regulators warn each other about:

- clone firms using similar names
- platforms that target retail users across borders
- suspected manipulation campaigns tied to listings
- fast-moving scam patterns

In crypto markets, alerts are useful because scams and unauthorized platforms can appear and grow quickly through online promotion.



3.7.6 Cooperation on the financial crime side

Financial crime cooperation often runs through financial intelligence units and law enforcement. These channels are older and often more mature than stability cooperation.

They include:

- secure information exchange about suspicious activity
- joint investigations
- sharing of typologies and emerging methods

These channels matter for crypto because many major frauds and laundering schemes are cross-border by default.

3.7.7 Why cooperation is harder for stability information

Stability cooperation requires sharing different kinds of information.

A stability authority may need to know:

- the size and composition of a stable-value arrangement's reserves
- redemption flows during stress
- liquidity needs at a large venue or lender
- group-wide exposures and intra-group transfers

This information is often confidential and can be market-moving.

Domestic law may not permit sharing it.

Even when law permits sharing, agencies may be cautious.

That is why peer reviews often conclude that stability cooperation tools are not yet as mature as financial crime cooperation tools.



3.8 Persistent baseline stress points

The baseline layer creates convergence, but it also has stress points. These are areas where implementation is hard, where definitions are contested, or where existing tools do not fit.

3.8.1 Perimeter ambiguity at the frontier

The frontier problems are familiar by now.

- DeFi arrangements with unclear control points.
- Self-held wallets and peer-to-peer transfers.
- Cross-chain bridges that move value across ecosystems.
- Protocol governance that is spread across many holders but still guided by a small group.
- Yield products that blur lending, staking, and trading.

Baselines push regulators to find control points, but control points are not always easy to prove.

A regulator may suspect a group controls a system, but proving that control in a legal process can be hard.

A regulator may identify a front end, but the front end can move or change.

A regulator may target on-ramps, but users can route around them.

This frontier ambiguity is part of why adoption lags at the edges.

3.8.2 Adoption without field testing

A second stress point is the gap between passing rules and making them work.



This appears most clearly in Travel Rule work, where many jurisdictions report adopting legislation while enforcement and field testing remain limited.

It also appears in market conduct work, where licensing regimes may exist but supervision resources are thin.

Crypto supervision is resource-intensive. It requires technical understanding, data analysis, and the ability to investigate cross-border structures.

A jurisdiction can adopt a baseline rule quickly. Building the supervisory machinery takes longer.

3.8.3 Coordination gaps and data gaps

A third stress point is coordination and data.

Even when domestic rules are strong, cross-border firms can route activity through weaker points.

Without shared data standards, supervisors collect different information and cannot easily compare.

Without routine information exchange, supervisors may not see group-wide exposures until stress hits.

This is why peer reviews keep returning to coordination and data as the main remaining gaps.

3.8.4 The incentives problem: why firms seek weak points

A baseline creates a floor, but firms still face incentives.

If a firm can reach customers in a strict jurisdiction while booking operations in a weak jurisdiction, it can cut compliance costs and avoid oversight.



That creates a race problem: strict jurisdictions risk pushing activity offshore, while weak jurisdictions attract activity that can harm users.

Baselines aim to stop this race by creating convergence pressure. But convergence is slow, and crypto markets move quickly.

3.9 How baseline standards shape domestic coordination choices

The baseline layer does not only push countries toward similar rules. It also pushes them toward similar coordination choices.

3.9.1 Group-wide supervision as a design goal

A cross-border crypto firm can split its business into affiliates: one holds custody, one runs trading, one runs lending, one runs marketing. A domestic supervisor might license only one affiliate.

Baselines push supervisors to look at the group.

Group-wide supervision requires:

- clarity on which entity contracts with customers
- clarity on where customer assets are held
- clarity on related-party flows
- the ability to request information across affiliates

When supervisors cannot get a group-wide view, failures become more likely.

That is why coordination tools like colleges are increasingly discussed for crypto firms that are system-relevant.

3.9.2 Equivalence, recognition, and the limits of “passporting” logic



Some jurisdictions consider equivalence or recognition models: allowing firms authorized elsewhere to operate locally if the foreign regime is judged comparable.

This can reduce duplication and can support cross-border business.

It can also create risk if equivalence judgments are too broad or not kept up to date.

A workable equivalence model needs:

- clear criteria for what counts as comparable
- ongoing monitoring of the foreign regime
- a plan for what happens when the foreign regime changes or enforcement weakens

Baselines help here because they provide shared reference points. Two jurisdictions can compare their rules against baseline expectations rather than against each other's statutes word for word.

3.9.3 Extraterritorial enforcement pressure

Another coordination choice is extraterritorial enforcement: taking action against offshore firms that serve domestic users without authorization.

This can include:

- restricting marketing
- blocking access through domestic intermediaries
- issuing public warnings
- pursuing local affiliates and promoters



Extraterritorial enforcement is difficult. It depends on legal powers and on cooperation from other jurisdictions.

Baselines support extraterritorial work indirectly by creating shared definitions and shared expectations about what a regulated service is.

3.10 Practical coordination design: what a mature baseline layer would look like

A research paper benefits from moving beyond “cooperate more” and describing what a mature coordination setup would require in practice.

A mature baseline layer would include at least six elements.

3.10.1 Shared minimum data elements

Cooperation improves when supervisors collect comparable information.

For crypto markets, shared minimum data elements could include:

- basic firm identifiers and group structure maps
- custody arrangements and where keys are controlled
- client asset segregation methods and reconciliation results
- incident reports, including breaches and outages
- reserve and redemption data for stable-value arrangements
- leverage and liquidation metrics for venues offering margin
- major related-party exposures and flows

This does not require a single global database. It requires that each supervisor collects a core set in a comparable way, so information can be shared and understood.



3.10.2 Routine incident reporting and cross-border escalation paths

Crypto incidents move fast. A breach can drain assets in minutes. A stable-value token can face a redemption wave within hours. A venue can freeze withdrawals and trigger panic.

A mature coordination setup would include:

- clear incident reporting thresholds
- contact points for rapid cross-border notifications
- agreed templates for what information is shared
- plans for public communications where needed

This is common in banking supervision. It is less mature in crypto supervision.

3.10.3 Colleges for system-relevant firms and arrangements

Not every firm needs a college. Colleges are resource-heavy. They should be reserved for:

- firms with large retail reach across borders
- firms that combine multiple roles (trading, custody, lending)
- stable-value arrangements with broad use

For these cases, colleges can reduce blind spots and allow supervisors to align on priorities.

3.10.4 A clearer boundary between supervisory sharing and enforcement sharing

Supervisors share information for two reasons.

- Routine supervision: assessing governance, controls, and risk.



- **Enforcement:** investigating suspected breaches and taking action.

The legal basis for sharing can differ between these reasons.

A mature setup would make this clear and would provide legal bases for each, with confidentiality protections.

3.10.5 Interoperable AML information exchange

Travel Rule work shows that AML information exchange is partly a technical coordination problem.

A mature setup would support interoperability across firms and jurisdictions. It would also address privacy and data transfer issues clearly so compliance does not depend on uncertain legal interpretation.

3.10.6 Capacity building and shared tooling

Finally, baseline implementation depends on capacity.

Not every supervisor has the same resources. Not every jurisdiction has the same ability to collect market data or investigate complex protocols.

A mature baseline layer therefore includes capacity building: training, shared typologies, and shared tools.

This is not charity. It reduces arbitrage. If weak links remain, risk will route through them.

3.11 Chapter synthesis: why the baseline layer drives convergence and why gaps persist

Global standards and coordination form the hidden backbone of crypto regulation.



The baseline layer is a stack. Financial integrity expectations bring service providers into AML supervision and require that customer information and transfer information be available. Market conduct expectations push toward strong custody safeguards, conflict controls, market abuse deterrence, and clear communications to customers. Stability expectations push toward oversight that matches function, scales with system relevance, and covers stable-value arrangements as a priority bridge to payment and liquidity risk.

The baseline layer drives convergence through peer review, reputational pressure, and the practical needs of cross-border supervision.

At the same time, gaps persist for three reasons.

First, the frontier is hard. DeFi, self-held wallets, and cross-chain systems do not map neatly onto entity-based rulebooks.

Second, adoption is faster than supervision. A law can be passed quickly. Building supervisors, data systems, and enforcement practice takes time.

Third, coordination is uneven. Cooperation tools for financial crime are mature. Cooperation tools for stability monitoring are less mature. Data collection is not yet consistent.

These gaps matter because crypto markets move fast and cross borders easily. Weak links invite arbitrage and make enforcement difficult.

For the rest of this research paper, the baseline layer should be treated as the floor that shapes domestic choices. It explains why many domestic regimes look similar in structure even when their statutes differ. It also explains why the hardest policy questions are



increasingly about making rules work in practice: supervision capacity, enforcement reach, and cross-border coordination that goes beyond one-off requests.

3) Global standards and coordination (the “baseline” layer)

Crypto regulation is often described as fragmented. That description is partly true in the narrow sense that rulebooks differ across borders, agencies argue about perimeter, and enforcement strength varies. But the word “fragmented” hides a more important structural fact: most jurisdictions are not starting from zero. They are building domestic rules on top of an existing layer of global expectations that already shapes how financial markets are policed.

This baseline layer is not a single global statute. It is a stack of shared expectations set by standard-setting bodies and reinforced through peer review, mutual evaluation, and routine cooperation between supervisors and law enforcement. The stack is “soft law” in form, yet hard in effect. Countries are judged against it. Firms are judged through it. Financial institutions use it to decide whether to do business with each other. Domestic regulators use it as a reference point when they draft or revise national rules.

In practice, the baseline layer does four things.

First, it defines minimum outcomes that jurisdictions are expected to achieve. A country does not need to write the same words as its peers, but it is expected to achieve a recognizable result: gatekeepers must apply customer checks; market operators must deter manipulation; large stable-value arrangements must not operate without meaningful reserve and redemption controls.

Second, it creates convergence pressure even without treaty-level enforcement. Peer review, public comparison, and the reputational



cost of being tagged as weak all push jurisdictions toward similar outcomes.

Third, it shapes domestic perimeter decisions. When lawmakers argue about whether a token is a security or whether a protocol is “decentralized,” they often end up answering a different question: what control points must sit inside supervision if the country wants to meet baseline expectations on financial crime, market conduct, and stability.

Fourth, it provides cooperation rails. Regulators do not need to invent cross-border tools from scratch. They can reuse, extend, and adapt tools built for banking, securities markets, and financial intelligence work: memoranda of understanding, supervisory colleges, information exchange channels, and joint work plans.

This chapter focuses on three baseline pillars that dominate most crypto policy debates.

- The anti–money laundering and counter–terrorist financing baseline, set through global financial integrity standards.
- The investor protection and market integrity baseline, shaped through securities regulator coordination.
- The financial stability baseline, shaped through a framework for crypto activities and stable-value arrangements.

The chapter then turns to the coordination problem itself: how baseline standards are made operational across borders, where the hardest gaps remain, and what tools are being used to close them.

3.1 The baseline architecture: who sets standards, and why it matters



International financial regulation is not run by a global legislature. It is shaped by a network of standard-setting bodies, forums, and institutions. Each body has a different mandate and a different set of members. Each produces a different kind of output: guidance, principles, recommendations, and peer review reports. Together they create a baseline that many domestic regimes follow.

For crypto-assets, three bodies dominate the baseline conversation.

The financial integrity standard setter. This body sets the minimum expectations for anti-money laundering and counter-terrorist financing controls. It defines who must sit inside the AML/CFT perimeter, what customer checks are expected, and how information should follow value transfers. It also runs mutual evaluations and follow-up processes that rate implementation.

The securities regulator coordination body. This body frames crypto markets as markets in which investors can be harmed through fraud, conflicts, custody failure, and market abuse. It focuses on outcomes that securities regulators care about: fair dealing, orderly markets, custody safeguards, and controls that deter manipulation.

The financial stability forum. This forum focuses on how crypto activities can create spillovers: runs, forced selling, links to banks, and payment system adjacency through stable-value arrangements. It issues high-level recommendations and uses peer reviews to identify gaps.

Other bodies are also part of the baseline stack, even when they are less visible in public debate.

- Banking supervisors set standards for how banks should treat crypto exposures, including how those exposures appear in capital planning and risk limits.



- Payment and market infrastructure bodies shape expectations for settlement, custody, and operational resilience.
- Macroeconomic institutions track cross-border spillovers, capital flow issues, and the interaction between crypto activity and monetary and financial conditions.
- Tax cooperation work sets expectations for how crypto activity should be reported for tax compliance.
- Financial intelligence cooperation networks create secure rails for cross-border exchange of information about suspicious activity.

A research paper can treat the baseline architecture as a set of overlapping circles. One circle is financial crime controls. One is market conduct. One is stability. The overlap matters because the same firm may face all three at once.

3.1.1 Why “soft law” is hard in practice

It is tempting to dismiss global standards as non-binding guidance. That underestimates their force.

The baseline layer shapes domestic law because governments do not want to be publicly rated as weak. A poor rating on AML/CFT implementation can affect the willingness of foreign banks to offer correspondent services. Weakness in market conduct can affect investor confidence and the willingness of firms to list or operate locally. Weakness in stability oversight can affect macro credibility when a large stress event occurs.

The baseline layer also shapes firm behavior because large cross-border firms do not want to operate under a patchwork of incompatible expectations. They tend to build to the highest common



requirement among the markets that matter to them. A firm that wants access to large pools of customers and banking rails will often adopt baseline controls even before a local law requires them.

This is not altruism. It is risk management. A cross-border firm wants to avoid being cut off from banks, payment processors, and institutional counterparties. Those counterparties will often require evidence of compliance with baseline expectations.

3.1.2 How the baseline layer is enforced without a global police force

Global standards become real through several channels.

Peer review and mutual evaluation. Countries are assessed against shared criteria. The assessments become public and create political pressure. They also become practical inputs for other decisions, such as whether financial institutions treat a jurisdiction as high risk.

Market access and de-risking pressure. Banks and payment firms respond to perceived regulatory weakness by reducing exposure. That can be blunt and can create its own policy issues, but it is part of why governments care about AML/CFT evaluations.

Domestic statutory hooks. Many domestic laws explicitly incorporate global standards by reference or use them as interpretive guides. Even when not explicit, domestic agencies cite them as a rationale for guidance and enforcement.

Cooperation reliance. Regulators need each other. A market regulator that wants help with a cross-border case depends on another regulator's willingness and legal ability to share information. That willingness is shaped by trust. Trust is shaped by whether both sides



meet baseline expectations on confidentiality, due process, and supervisory competence.

Reputational competition. Jurisdictions compete for legitimate activity. Being seen as a place with weak controls can attract scams and short-term volume, but it can also repel long-term institutional participation.

The net effect is that global baselines create convergence even without treaties.

3.1.3 Why the baseline layer is a special issue in crypto

Crypto markets are cross-border by design. A token can be issued in one place, traded on venues in several other places, held in custody by a firm somewhere else, and used as collateral in a protocol that has no clear physical base. If domestic rules differ widely, risk will tend to flow toward the weakest perimeter.

That is why baseline standards matter so much. They are the nearest thing to a shared floor.

At the same time, the baseline layer cannot remove all differences. Domestic law still matters, especially in areas where global bodies set high-level expectations but leave details to local choice: retail access rules, civil liability, product bans, and the boundary between securities law and crypto-specific regimes.

The baseline layer is therefore best understood as a common language that allows domestic regimes to coordinate, compare, and cooperate. It reduces the cost of building rules and the cost of working together, even when laws are not identical.

3.2 Global AML/CFT standards: the minimum floor for financial integrity



3.2.1 The AML baseline in one sentence

The AML/CFT baseline for crypto is direct: jurisdictions must identify and mitigate laundering and terrorist financing risks linked to crypto activity by bringing relevant actors inside a regulated perimeter, requiring preventive controls, supervising compliance, and enabling domestic and cross-border cooperation.

This baseline is built around two definitional moves.

First, it defines a broad category of transferable digital value used for payment or investment.

Second, it defines a category of service providers that perform exchange, transfer, custody, and issuance-related services as a business.

Once those definitions apply, a familiar control stack follows: customer due diligence, ongoing monitoring, recordkeeping, suspicious reporting, sanctions controls, and information that accompanies transfers.

3.2.2 Technology-neutral definitions and the “as a business” line

A central design choice in the AML baseline is that definitions are meant to be function-based, not chain-based.

If a service allows customers to exchange between fiat and crypto, exchange between crypto assets, transfer crypto for others, or hold crypto on behalf of others, it is treated as part of the gatekeeping system. The definition is meant to stop technical form from becoming a loophole.

The hardest interpretive line sits in the phrase “as a business.” That phrase tries to separate ordinary users from professional service



providers. In crypto, it also forces regulators to define what counts as a service, not merely code.

A user who sends tokens from a self-held wallet to another self-held wallet is not a service provider. A firm that holds customer keys and transfers tokens on request is.

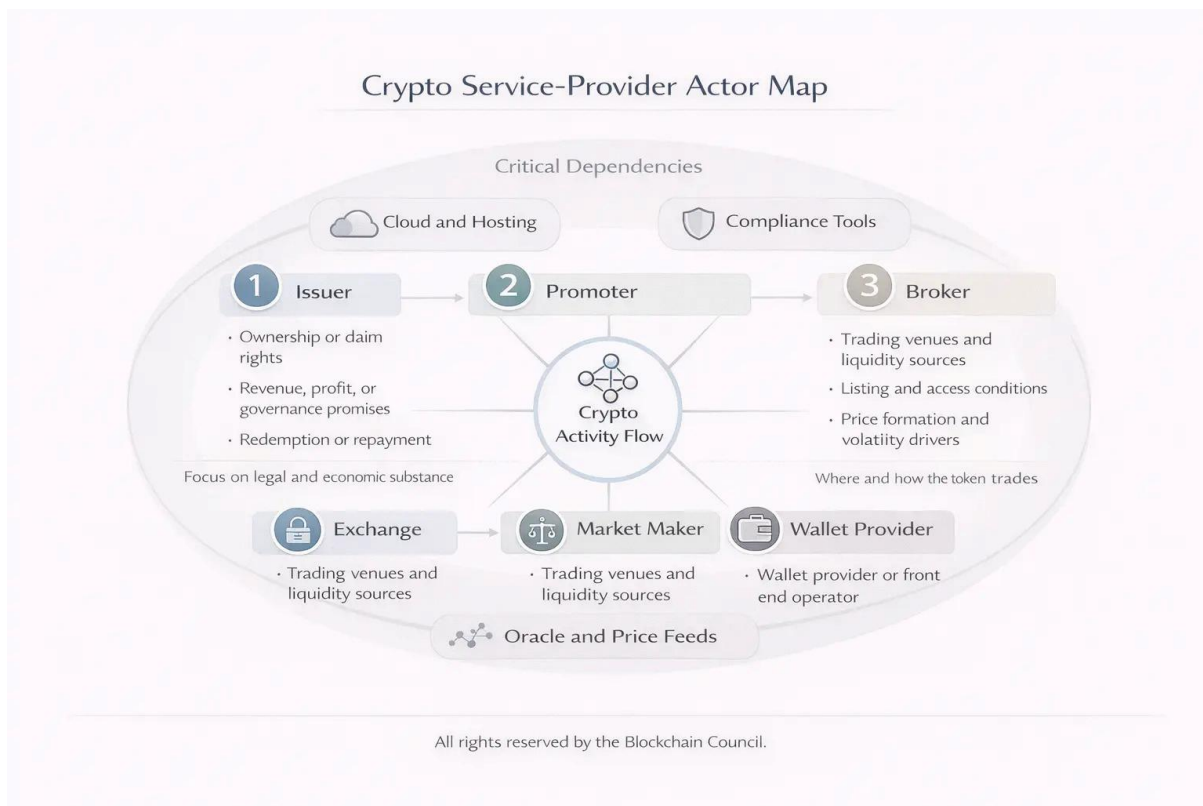
The hard cases sit between.

- A group that runs a user-facing front end for a protocol, collects fees, and controls upgrades looks more like a service provider.
- A developer who publishes open-source code and does not control its use looks less like a service provider.
- A governance group that can change fee settings, control admin keys, and direct treasury funds may fall somewhere in between depending on how much control it has in practice.

The baseline approach pushes countries to look for control points. Where there is an actor that can apply checks and monitoring, authorities are expected to bring that actor inside the perimeter.

This is one reason the AML baseline has become a key driver of perimeter expansion. Even jurisdictions that are hesitant to build full market conduct regimes often implement AML registration and supervision for crypto service providers.

3.2.3 The AML control stack expected for crypto service providers



An AML law is not only a statute that declares obligations. It is a stack of policies, procedures, technology, supervision, and enforcement that produces usable outcomes.

A mature control stack includes at least the following components.

Risk assessment and a risk-based program. Firms are expected to assess their risks and apply stronger controls where risk is higher. A firm serving retail users through bank transfers and offering fast withdrawals faces different risk than a firm serving institutional clients with strict onboarding and lower-speed transfer limits.

Licensing or registration. The baseline expects jurisdictions to require service providers to be licensed or registered, so supervisors can identify who is in the market and can enforce standards.



Fit-and-proper checks and ownership clarity. A registration regime that does not identify beneficial ownership or the people controlling a firm provides little protection.

Customer due diligence. Firms should identify customers, understand control for entities, and apply higher scrutiny to higher-risk cases. In crypto, this often includes understanding whether the customer is acting for themselves or as a business.

Ongoing monitoring. Firms should monitor activity for patterns consistent with laundering, fraud proceeds movement, or sanctions evasion. That includes monitoring at account level and, where feasible, monitoring linked on-chain activity.

Recordkeeping and audit trails. Supervision depends on records. Without them, firms can claim compliance without proof.

Suspicious reporting. Firms should report activity that appears linked to crime.

Sanctions controls. Firms should screen customers and, where feasible, apply controls to prevent service to sanctioned parties and to stop transfers linked to sanctioned wallets.

International cooperation readiness. Firms should be able to respond to lawful requests and to preserve records.

Each component has a crypto-specific twist.

- Transfers can be final and can occur at any time.
- Wallet addresses are not names.
- Customers can move assets across chains through bridges.
- Some activity bypasses custodians through self-held wallets.



The baseline does not require supervisors to control every peer-to-peer transfer. It requires supervisors to control the regulated gatekeepers and to apply extra measures when gatekeepers face higher risk.

3.2.4 The Travel Rule: a baseline requirement that turns into a systems problem

The Travel Rule is the clearest example of a global expectation that forces operational coordination.

At its core, the rule requires that certain information about the sender and recipient travel along with a value transfer between service providers. The policy goal is old: payment transparency. The implementation is new because crypto transfers do not carry identity information by default.

In practice, Travel Rule compliance requires firms to do several things at once.

- Collect required information from the originator.
- Validate or at least make reasonable checks on information quality.
- Send that information to the recipient institution or service provider.
- Receive information when acting as the recipient.
- Apply controls when information is missing or does not make sense.
- Retain records and make them available to supervisors.

Two implementation realities follow.



First, it requires interoperability. A Travel Rule program is only as strong as the ability of one firm to exchange the required data with another. If a firm cannot identify the counterparty service provider, or cannot send information in a mutually usable format, compliance becomes uneven.

Second, it requires supervision. Legislation alone does not produce compliance. Supervisors must test whether firms are actually exchanging data, whether data quality is acceptable, and whether firms apply controls when counterparties fail to comply.

Travel Rule implementation also raises design questions that are not present in older wire transfer systems.

- How should service providers identify whether a counterparty is a service provider or a self-held wallet?
- How should firms handle transfers that involve multiple hops or that route through aggregators?
- How should firms handle cross-border transfers where the counterparty is in a jurisdiction that has not implemented the rule?
- How should data be protected when it includes personal information and moves across borders?

Because of these issues, Travel Rule implementation has become partly a coordination exercise and partly an engineering exercise.

One recent global update on implementation reported that among jurisdictions that did not prohibit crypto service providers, a large majority reported passing Travel Rule legislation. The same update also observed that enforcement experience remained limited. In other words, many countries have adopted the rule on paper, while fewer



have tested it in the field or taken action against weak implementation.

For research purposes, this matters because it shows a gap between adoption and effect. It also shows why cross-border alignment matters: a rule that depends on counterparties cannot work well when only half the counterparties are in scope.

3.2.5 Common Travel Rule implementation models

Although laws are written at a high level, firms tend to implement Travel Rule compliance through a small number of operational models.

Bilateral exchange. Firms agree to exchange required information directly. This is workable when a firm has a small number of counterparties and a stable set of relationships. It becomes less workable when the firm's customer base sends transfers to a long tail of destinations.

Shared messaging networks. Firms join a network that provides a common format and routing layer for messages. This model improves reach but creates questions about governance of the network, confidentiality, and reliance on a third party.

Directory and counterparty identification tools. Some systems focus on solving the “who is the counterparty” problem. The firm can identify whether the recipient address belongs to a known service provider and can route required information accordingly.

Risk-based gating. Some firms apply stricter controls to outbound transfers when they cannot exchange information with the counterparty. Controls include transfer delays, enhanced customer checks, and restrictions on high-risk destinations.



Each model interacts with domestic law on privacy and data transfer. A law that restricts cross-border transfer of personal information can slow Travel Rule implementation unless a legal basis exists for compliance-related transfers.

The baseline layer therefore pushes countries to coordinate not only on rules but also on data governance.

3.2.6 Unhosted wallets, peer-to-peer transfers, and the limits of entity-based supervision

The AML baseline is centered on service providers. That is a strength because supervisors can enforce obligations on firms. It is also a limit because some crypto activity occurs without intermediaries.

Self-held wallets allow users to transfer value directly, without a firm controlling the transfer. Some protocols allow users to exchange assets directly through contract code, without a custodial intermediary.

These realities create two policy pressures.

First, supervisors push stronger controls onto gateways: fiat on-ramps, custodians, and venues where users exchange between fiat and crypto. If peer-to-peer activity cannot be directly supervised, the next best control point is the point where value enters and exits the regulated sector.

Second, supervisors expect regulated firms to manage exposure to self-held wallets through risk-based controls. That can include collecting information about the owner of a self-held wallet in certain cases, applying transfer limits, and monitoring patterns that suggest laundering.



There is no single global answer for unhosted wallet exposure. But the baseline expectation is that jurisdictions should not ignore it. They should assess the risk and require reasonable controls from the regulated sector.

3.2.7 DeFi and the control-point approach

Decentralized finance is often presented as an obstacle to AML rules because it can operate without a traditional intermediary. The baseline layer does not accept that framing as a reason for inaction. It treats DeFi as a perimeter stress test.

The key question is the same: is there an actor providing a covered service as a business?

If there is an actor that runs a front end, controls admin keys, collects fees, markets the service, or can change the system, regulators will ask whether that actor should be treated as a service provider.

If there is no such actor, regulators often shift focus to the gateways: the exchanges and custodians that allow users to acquire and cash out assets used in the protocol.

A recent global survey-based update reported that only a small number of jurisdictions had in practice registered or licensed DeFi entities as service providers, and only a small number reported enforcement actions against DeFi entities that qualify as service providers. That is a data point that matters because it suggests the baseline is clear, but implementation at the frontier remains limited.

The reasons are practical.

- Some DeFi systems are hard to map onto entity-based licensing.



- Some jurisdictions lack legal tools to treat a front end or governance group as a regulated firm.
- Some supervisors lack resources to investigate and prove control.
- Some projects can relocate quickly or distribute operations across borders.

These limits do not remove the baseline expectation. They help explain why implementation gaps persist.

3.2.8 Implementation gaps as a baseline feature, not an exception

Global AML/CFT work does not assume perfect implementation. It is built around the idea that implementation will be uneven and must be monitored.

Updates and evaluations repeatedly point to gaps.

- Many jurisdictions remain only partially compliant with crypto-related AML expectations.
- Many struggle to identify unregistered service providers operating domestically.
- Many have passed Travel Rule legislation but have limited enforcement activity.
- DeFi perimeter work remains thin.
- A meaningful number of jurisdictions do not respond to surveys or provide evidence of progress, suggesting capacity or priority gaps.

For a research paper, the right conclusion is not that the baseline failed. The right conclusion is that the baseline is a floor that



jurisdictions reach at different speeds. That speed difference creates arbitrage risk.

It also creates a policy problem for jurisdictions that do implement: how to stop their regulated firms from losing business to weakly supervised offshore competitors.

That question leads directly to the coordination tools discussed later in this chapter.

3.3 Securities regulator standards: investor protection and market integrity outcomes

The market conduct baseline for crypto is shaped through the coordination of securities regulators. The core argument is that many crypto markets recreate the same basic risks seen in securities markets: fraud, manipulation, conflicts, misuse of customer assets, and weak disclosure.

Even when a token is not classified as a security, the way it is sold and traded can create securities-like harm. The baseline response is to focus on activities and outcomes.

3.3.1 Outcomes first: why form is a weak guide in crypto markets

Crypto projects can change labels quickly. A token can be called a “utility” token even when it is sold mainly as a speculative bet. A venue can claim it is a technology platform rather than a market operator. A lending product can be packaged as a simple “earn” program rather than a credit product.

A market conduct baseline that relies on labels will be easy to evade. That is why securities regulator coordination work tends to stress outcomes: fair dealing, conflict control, custody safeguards, and deterrence of market abuse.



This does not mean every crypto activity should be regulated as a securities market. It means regulators should not accept a label as a substitute for analysis.

A research paper can treat this as a perimeter principle: treat the substance of the activity as the starting point, then apply controls that match the risk.

3.3.2 The control domains that keep appearing in securities regulator work

Securities regulator coordination work in crypto has produced a set of policy recommendations that cover a broad set of control domains. Rather than list recommendations as a checklist, it is more useful to group them into domains that map to harm channels.

Governance and accountability. Regulators expect firms that run markets, hold assets, and sell high-risk products to have clear governance: responsible people, clear decision-making, internal controls, and the ability for supervisors to hold someone accountable.

Conflicts of interest. Crypto firms often combine roles that create conflicts: operating a venue while also trading on it, holding custody while also lending, listing tokens while receiving allocations, selling products while also acting as principal. Regulators expect conflicts to be managed through a mix of structure, rules, and disclosure.

Custody and client asset safeguards. Custody failures have been a major driver of losses. A market conduct baseline treats custody and segregation as central investor protection outcomes.

Disclosure and communications. Regulators expect customers to receive information that allows them to understand what they are buying, what rights they have, and what risks they take. They also



treat marketing as part of market conduct. A market can be “orderly” in a narrow sense and still be distorted by misleading promotions.

Market abuse controls. The baseline expects controls that deter manipulation, insider dealing, and other abusive conduct. That includes surveillance, investigation capability, and coordination with law enforcement.

Market operation and trading controls. Trading rules, fee clarity, outage handling, listing standards, and order execution practices sit inside the market integrity frame. These are not mere operational details. They shape who wins and who loses.

Operational resilience and cyber controls. Because crypto markets are technology-heavy, resilience is a conduct issue. Outages can trap customers. Breaches can drain assets. Weak change management can alter market rules without warning.

Retail protection for high-risk products. Where markets offer leverage, complex derivatives, or yield products with embedded risk, regulators expect extra safeguards: risk warnings, limits, appropriateness checks, and product governance.

These domains map to recurring failure modes in crypto markets. That is why they keep appearing across jurisdictions even when statutes differ.

3.3.3 Vertical combination of roles as the central conduct stress point

A key difference between many crypto market structures and traditional securities markets is vertical combination. In many crypto groups, one firm or a set of affiliates can be all of the following at once:



- a market operator
- a broker
- a dealer
- a custodian
- a lender
- a stable-value issuer
- a staking service

Traditional finance can contain some combinations, but it usually surrounds them with strict conflict rules, capital rules, and oversight that assumes conflicts are present.

In crypto markets, vertical combination has often been treated as normal business structure rather than a risk factor. The market conduct baseline challenges that assumption. It treats vertical combination as a risk multiplier.

This matters for coordination because a global firm can move roles across borders. It can locate the venue in one jurisdiction, custody in another, and the trading desk in a third. Without cross-border coordination, no single supervisor sees the full conflict map.

3.3.4 Listing standards and token admission as a baseline issue

Token admission is often treated as a private business decision. In crypto markets, it becomes a public policy issue because listings shape retail exposure and because listing events can be tied to manipulation and insider abuse.

A market conduct baseline therefore pushes for a disciplined listing process.



A research-friendly way to describe this is to treat token admission as a risk gate with four main questions.

1. **Information quality.** Is there enough public information to allow users to understand supply, governance, upgrade control, and major risks?
2. **Concentration and unlock risk.** Are holdings concentrated? Are there large unlocks that could crash price and harm retail users?
3. **Code and admin risk.** Does the token rely on admin keys or upgrade mechanisms that allow a small group to change core behavior?
4. **Manipulation susceptibility.** Is liquidity thin? Is the token prone to wash trading? Are there known patterns of insider allocation and dumping?

Listing discipline connects to investor protection and market integrity at the same time. It also connects to coordination because a token listed on one major venue can quickly become a global retail product.

3.3.5 How the baseline deals with different domestic classifications

Domestic law still differs on classification. Some jurisdictions treat many tokens as securities. Some treat most as outside securities law unless they look like shares or funds. Some create a crypto-specific regime for non-security tokens.

A market conduct baseline does not require every jurisdiction to adopt the same classification. It pushes for a common set of outcomes where the activity creates comparable risk.

That is why an activity-based framing is useful for research.



- If a firm operates a trading venue for retail users, it should face baseline expectations for fair dealing, surveillance, conflict control, and resilience.
- If a firm holds client assets, it should face baseline expectations for segregation, key security, and recordkeeping.
- If a firm sells high-risk leveraged products, it should face baseline expectations for clear disclosure, product controls, and oversight.

This framing allows jurisdictions with different statutory definitions to still move toward similar market outcomes.

3.3.6 Implementation progress and the recurring gaps

The market conduct baseline has also moved from recommendations to implementation reviews. Those reviews tend to tell the same story.

Progress exists. Many jurisdictions have begun licensing trading venues or applying registration systems. Many have issued guidance on custody standards, conflict policies, and risk disclosures.

At the same time, recurring gaps remain.

- Cross-border information sharing remains uneven.
- Regulators face legal limits on sharing non-public information.
- Definitions differ enough to slow cooperation.
- Offshore provision allows firms to reach customers without local authorization.
- Data collection is inconsistent, making it hard to compare markets and detect group-wide risk.



In other words, adoption is moving. Coordination and enforcement remain the hard part.

3.4 Financial stability standards: the system-wide lens

The financial stability baseline for crypto is shaped through a framework that focuses on system-wide spillovers. It treats crypto markets as part of the wider financial system when they become large, interconnected, or tied to payment rails.

This baseline matters even in jurisdictions where crypto is not yet large. The reason is simple: crypto markets can grow quickly, and cross-border firms can import stress.

3.4.1 The stability problem is about channels, not labels

A stability lens starts with a question: can crypto activity create problems that spill beyond the original users?

Volatility alone does not answer that question. Many assets are volatile. What matters is the channel.

Common channels include:

- leveraged positions that force sales during price drops
- liquidity mismatch in lending and borrowing products
- stable-value token runs that force reserve asset sales
- links to banks and funds through custody, lending, and investment exposures
- payment and settlement reliance on stable-value tokens
- operational failures at large venues that trap customers and freeze market function



A stability baseline therefore pushes for functional oversight that matches these channels.

3.4.2 High-level recommendations and perimeter breadth

The stability baseline is set through high-level recommendations for crypto activities and markets.

Those recommendations share several themes.

Comprehensive perimeter. Authorities should not assume crypto sits outside regulation. Where activities create stability risk, the perimeter should expand to include them.

Functional regulation. Similar activities should face similar outcomes. If a product looks like credit, it should face credit-like oversight. If a stable-value arrangement looks like money issuance, it should face money-issuance style expectations.

Proportionality. Oversight should scale with risk, size, complexity, and system relevance.

Group-wide view. Supervisors should be able to understand group structure and exposures across affiliates.

Data and monitoring. Authorities should collect enough information to monitor market development, identify leverage, and see concentrated dependencies.

Cross-border cooperation. Authorities should cooperate because the same firm and the same stable-value token can affect many markets.

A useful research point is that stability work often leaves detailed consumer and market conduct rules to other bodies. That is not a gap; it is division of labor. The stability baseline is about system-wide monitoring and perimeter coverage.



3.4.3 Stable-value arrangements as the main stability bridge

Stable-value arrangements are the clearest bridge between crypto markets and the wider financial system.

A stable-value token that is widely used can become a unit of account inside crypto markets, a settlement asset, and a payment instrument for commerce. That creates expectations of reliability and quick redemption.

A stability baseline for stable-value arrangements focuses on three core functions.

Issuance and redemption. How tokens are created and destroyed, what rights holders have, and what assets back redemption.

Transfer and settlement. How tokens move between users and what operational systems support movement.

User access and storage. How users store tokens, how custodial and non-custodial solutions interact, and what happens during outages.

The stability baseline for stable-value arrangements tends to focus on the following control areas.

- Reserve assets: quality, liquidity, custody, segregation, and valuation.
- Redemption: clarity, timeliness, conditions, and what happens in stress.
- Governance: who controls reserve management, who can change rules, and how conflicts are managed.
- Operational resilience: uptime, incident response, and reliance on third parties.



- Disclosures: clear information on reserves, redemption terms, and risks.
- System relevance: extra expectations when an arrangement reaches scale.

The label “stable” is treated with caution. The baseline view is that the word can create false expectations. A stable-value arrangement should be judged by its design and controls.

3.4.4 Peer reviews and what they reveal about real coordination limits

Stability work also uses peer review to identify gaps.

Peer reviews tend to point to uneven implementation and to the limits of cross-border cooperation.

A repeated theme is that fragmented implementation creates arbitrage opportunities and complicates cross-border oversight.

Another repeated theme is that cooperation tools are often ad hoc and not designed for stability-related information sharing. Traditional cooperation tools are often designed for enforcement cases or for routine supervision in banking. Stability monitoring requires different data: leverage, liquidity dependencies, reserve exposures, and group-wide flows.

Peer reviews also point to data gaps.

- Authorities do not always collect comparable information about crypto market activity.
- Stable-value arrangements can be large, but reserve data may not be collected or shared.



- Group structures can be complex, and supervisors may lack the full map.

A research paper can use these observations to explain why stability oversight is harder to coordinate than AML cooperation. AML cooperation has long-established rails through financial intelligence units. Stability cooperation is less mature because it requires sharing supervisory data that is often confidential and because no single supervisor has a full picture of global exposures.

3.4.5 Interaction with banking and payment infrastructure standards

The stability baseline interacts with other standards.

Banking supervisors set expectations for how bank exposures to crypto should be treated. That includes how exposures are counted, how capital is planned, and how concentration is managed.

Payment and market infrastructure bodies shape expectations for settlement, resilience, and the handling of operational outages.

These bodies are not “crypto bodies,” but their standards shape crypto regulation when banks custody crypto, when payment firms support stable-value redemption, or when stable-value tokens are used in settlement.



Global Crypto Regulatory Baseline Architecture

“Baseline Stack” Layered Model



All rights reserved by the Blockchain Council.

3.5 The wider baseline stack beyond the three headline pillars

A full baseline picture includes more than AML, market conduct, and stability.

3.5.1 Banking supervision and capital treatment

When banks gain exposure to crypto, whether through custody, trading, lending, or investment, bank supervisors focus on two broad aims.

First, make sure exposures are measured and disclosed. Hidden exposures create surprise during stress.

Second, make sure capital and risk limits reflect the risk profile of crypto exposures.



This is not an attempt to ban crypto from banks by default. It is an attempt to avoid a repeat of past episodes where new asset exposures built up without a clear prudential view.

For crypto markets, bank supervision matters because banks provide critical rails: custody services, settlement accounts, and payment processing for fiat on-ramps and redemptions.

If banks pull back because of uncertain prudential treatment, crypto firms may lose access to core rails. If banks move in without proper controls, bank exposure can become an amplifier.

3.5.2 Payment systems and market infrastructure expectations

Crypto markets rely on infrastructure that can look like market infrastructure even when it is not labeled that way.

Trading venues match orders and produce prices. Custodians hold assets. Stable-value arrangements can function as payment instruments.

Payment and infrastructure standards bring a focus on operational resilience, settlement finality, and continuity under stress.

Even where a stable-value token is not formally designated as payment infrastructure, the same operational questions matter when it is widely used: what happens if transfer systems fail, if redemption rails break, or if a critical service provider goes offline.

3.5.3 Macroeconomic surveillance and policy coordination

Macroeconomic institutions look at crypto through a different lens.

They focus on cross-border flows, capital controls, currency substitution risk, and the interaction between crypto activity and monetary policy.



This layer matters most when stable-value tokens are used in ways that resemble dollarization, or when crypto activity becomes a meaningful part of financial access.

Even in jurisdictions where crypto is mainly speculative, macro institutions track spillover risk, especially when leverage and stable-value arrangements grow.

3.5.4 Tax reporting expectations

Tax compliance is another baseline layer that increasingly touches crypto.

Tax authorities care about reporting of crypto transactions and holdings, especially when crypto provides a way to hide taxable gains or to move value across borders.

Tax reporting standards do not replace market conduct rules or AML rules, but they add a separate pressure for service providers to keep records and share information.

From a coordination standpoint, tax reporting can reinforce AML recordkeeping and can create another reason for service providers to register and maintain a compliant customer base.

3.6 How domestic regimes build on the baseline layer

Global baselines shape domestic law, but domestic law still makes the core design choices.

A research paper can describe three common pathways by which countries build rules on top of the baseline.

3.6.1 Pathway one: fit crypto into existing categories



Some jurisdictions extend existing securities, payments, and banking rules to cover crypto where possible. Under this pathway, the main task is classification.

If a token is treated as a security, securities offering and trading rules apply.

If a stable-value token is treated as money-like, payments or e-money rules apply.

If a firm holds customer assets, custody and client asset rules apply.

This pathway can work well where existing statutes are flexible. It can also produce uncertainty when older definitions do not map cleanly onto newer designs.

3.6.2 Pathway two: create a crypto-specific market regime

Some jurisdictions create a tailored regime for crypto assets and crypto services. This is often done for the part of the market that does not fit neatly into securities or payments categories.

Under this pathway, lawmakers define crypto service types, set authorization requirements for those service types, and impose conduct and disclosure duties tailored to crypto markets.

The advantage is clarity. The downside is that it must be kept aligned with existing financial law so firms do not face inconsistent obligations.

3.6.3 Pathway three: hybrid perimeter with layered controls

Many jurisdictions end up with a hybrid model.

- AML registration and supervision applies to service providers.



- A crypto service authorization regime applies to market-facing services.
- Securities law applies to security-like tokens.
- Payments rules apply to money-like stable-value arrangements.

This layered model is common because it follows how baselines are layered. AML is often implemented first. Market conduct rules follow. Stability rules tighten as activity grows.

3.6.4 Baselines as drafting templates and as enforcement arguments

Baselines influence domestic regimes in two ways.

They influence drafting because lawmakers can borrow language and structure from global expectations. That speeds up legislation and reduces the risk of missing a key control area.

They influence enforcement because domestic agencies can cite global expectations as evidence of what “reasonable” controls look like. Even when a statute is high-level, global baselines help agencies interpret and justify detailed guidance.

3.6.5 Domestic perimeter disputes that the baseline layer tends to settle

In many domestic debates, two issues dominate: whether DeFi is outside regulation, and whether non-custodial activity should be unregulated.

The baseline layer does not settle every detail, but it pushes debates toward control points.

- If a front end is run as a business and controls access, regulators ask whether it should be within perimeter.



- If a venue offers services to the public, regulators ask whether it should be authorized.
- If a stable-value arrangement invites money-like use, regulators ask whether reserve and redemption controls should apply.

This is why baselines often influence perimeter more than they influence specific product rules.

3.7 Cross-border coordination: why it is required and how it is done

Global baselines assume cross-border coordination because crypto markets are not domestic markets with a few foreign links. They are cross-border markets first.

Without coordination, firms can avoid oversight by shifting operations. Without coordination, a supervisor sees only part of a group's risk. Without coordination, enforcement actions become slow and easy to evade.

3.7.1 The three facts that drive coordination needs

Three structural facts keep bringing coordination back to the center.

Borderless service provision. A crypto service provider can serve users in one country while being authorized elsewhere or nowhere.

Contagious risk events. A failure at a major venue, custodian, or stable-value arrangement can spread quickly through liquidity and settlement dependencies.

Fragmented data. Without information sharing, regulators see only local fragments: a local marketing campaign, a local affiliate, a slice of customer accounts. The real risk map sits across the group.

3.7.2 Cooperation is not one tool but a toolkit



Cross-border cooperation is often described as a single thing, but it is a toolkit. Different tools support different objectives.

Securities regulator tools. These include memoranda of understanding, information exchange agreements, and alert networks that share information about firms that may be operating without authorization.

Financial intelligence tools. These include secure channels for financial intelligence units to share information about suspicious activity and to support investigations.

Supervisory colleges. These are structured groups of supervisors who share an interest in the same cross-border firm or arrangement. A college creates routine coordination rather than ad hoc requests.

Crisis coordination tools. These include shared contact lists, incident reporting protocols, and plans for handling a failure at a large firm or stable-value arrangement.

A mature coordination system uses several tools at once.

3.7.3 Memoranda of understanding and information exchange agreements

Memoranda of understanding (MoUs) are a standard tool in financial regulation. They set out how supervisors share information, how they protect confidentiality, and how they handle requests.

In crypto, MoUs matter because:

- firms operate across borders
- market abuse schemes cross borders
- supervisors need background on group structure and control



MoUs also matter because they reduce friction. A supervisor with an MoU in place can request information without negotiating terms from scratch each time.

MoUs are not a cure. They depend on the legal powers of each party. A supervisor cannot share information it does not have. A supervisor cannot share information if domestic law blocks sharing. MoUs work best when domestic law is written with cross-border cooperation in mind.

3.7.4 Supervisory colleges: from requests to routine coordination

A supervisory college is a structured approach to supervising a cross-border firm or arrangement.

A college typically includes:

- a lead supervisor, often where the main entity is authorized
- host supervisors in jurisdictions where key services are offered
- sometimes central bank or stability authorities when stable-value arrangements are involved

Colleges can:

- share non-public supervisory information under confidentiality rules
- coordinate inspection plans
- discuss group risk, conflicts, and dependencies
- align on key supervisory priorities
- coordinate response to incidents

Colleges help solve a common crypto problem: group structures are complex, and no single supervisor sees the full picture.



A stable-value arrangement college can also include supervisors who focus on reserve management, redemption rails, and payment linkages.

3.7.5 Alert networks and warning systems

Another coordination tool is the sharing of alerts about firms that appear to be operating without authorization or that pose high risk.

Alert networks can help regulators warn each other about:

- clone firms using similar names
- platforms that target retail users across borders
- suspected manipulation campaigns tied to listings
- fast-moving scam patterns

In crypto markets, alerts are useful because scams and unauthorized platforms can appear and grow quickly through online promotion.

3.7.6 Cooperation on the financial crime side

Financial crime cooperation often runs through financial intelligence units and law enforcement. These channels are older and often more mature than stability cooperation.

They include:

- secure information exchange about suspicious activity
- joint investigations
- sharing of typologies and emerging methods

These channels matter for crypto because many major frauds and laundering schemes are cross-border by default.

3.7.7 Why cooperation is harder for stability information



Stability cooperation requires sharing different kinds of information.

A stability authority may need to know:

- the size and composition of a stable-value arrangement's reserves
- redemption flows during stress
- liquidity needs at a large venue or lender
- group-wide exposures and intra-group transfers

This information is often confidential and can be market-moving.

Domestic law may not permit sharing it.

Even when law permits sharing, agencies may be cautious.

That is why peer reviews often conclude that stability cooperation tools are not yet as mature as financial crime cooperation tools.

3.8 Persistent baseline stress points

The baseline layer creates convergence, but it also has stress points. These are areas where implementation is hard, where definitions are contested, or where existing tools do not fit.

3.8.1 Perimeter ambiguity at the frontier

The frontier problems are familiar by now.

- DeFi arrangements with unclear control points.
- Self-held wallets and peer-to-peer transfers.
- Cross-chain bridges that move value across ecosystems.
- Protocol governance that is spread across many holders but still guided by a small group.



- Yield products that blur lending, staking, and trading.

Baselines push regulators to find control points, but control points are not always easy to prove.

A regulator may suspect a group controls a system, but proving that control in a legal process can be hard.

A regulator may identify a front end, but the front end can move or change.

A regulator may target on-ramps, but users can route around them.

This frontier ambiguity is part of why adoption lags at the edges.

3.8.2 Adoption without field testing

A second stress point is the gap between passing rules and making them work.

This appears most clearly in Travel Rule work, where many jurisdictions report adopting legislation while enforcement and field testing remain limited.

It also appears in market conduct work, where licensing regimes may exist but supervision resources are thin.

Crypto supervision is resource-intensive. It requires technical understanding, data analysis, and the ability to investigate cross-border structures.

A jurisdiction can adopt a baseline rule quickly. Building the supervisory machinery takes longer.

3.8.3 Coordination gaps and data gaps

A third stress point is coordination and data.



Even when domestic rules are strong, cross-border firms can route activity through weaker points.

Without shared data standards, supervisors collect different information and cannot easily compare.

Without routine information exchange, supervisors may not see group-wide exposures until stress hits.

This is why peer reviews keep returning to coordination and data as the main remaining gaps.

3.8.4 The incentives problem: why firms seek weak points

A baseline creates a floor, but firms still face incentives.

If a firm can reach customers in a strict jurisdiction while booking operations in a weak jurisdiction, it can cut compliance costs and avoid oversight.

That creates a race problem: strict jurisdictions risk pushing activity offshore, while weak jurisdictions attract activity that can harm users.

Baselines aim to stop this race by creating convergence pressure. But convergence is slow, and crypto markets move quickly.

3.9 How baseline standards shape domestic coordination choices

The baseline layer does not only push countries toward similar rules. It also pushes them toward similar coordination choices.

3.9.1 Group-wide supervision as a design goal

A cross-border crypto firm can split its business into affiliates: one holds custody, one runs trading, one runs lending, one runs marketing. A domestic supervisor might license only one affiliate.

Baselines push supervisors to look at the group.



Group-wide supervision requires:

- clarity on which entity contracts with customers
- clarity on where customer assets are held
- clarity on related-party flows
- the ability to request information across affiliates

When supervisors cannot get a group-wide view, failures become more likely.

That is why coordination tools like colleges are increasingly discussed for crypto firms that are system-relevant.

3.9.2 Equivalence, recognition, and the limits of “passporting” logic

Some jurisdictions consider equivalence or recognition models: allowing firms authorized elsewhere to operate locally if the foreign regime is judged comparable.

This can reduce duplication and can support cross-border business.

It can also create risk if equivalence judgments are too broad or not kept up to date.

A workable equivalence model needs:

- clear criteria for what counts as comparable
- ongoing monitoring of the foreign regime
- a plan for what happens when the foreign regime changes or enforcement weakens



Baselines help here because they provide shared reference points. Two jurisdictions can compare their rules against baseline expectations rather than against each other's statutes word for word.

3.9.3 Extraterritorial enforcement pressure

Another coordination choice is extraterritorial enforcement: taking action against offshore firms that serve domestic users without authorization.

This can include:

- restricting marketing
- blocking access through domestic intermediaries
- issuing public warnings
- pursuing local affiliates and promoters

Extraterritorial enforcement is difficult. It depends on legal powers and on cooperation from other jurisdictions.

Baselines support extraterritorial work indirectly by creating shared definitions and shared expectations about what a regulated service is.

3.10 Practical coordination design: what a mature baseline layer would look like

A research paper benefits from moving beyond “cooperate more” and describing what a mature coordination setup would require in practice.

A mature baseline layer would include at least six elements.

3.10.1 Shared minimum data elements

Cooperation improves when supervisors collect comparable information.



For crypto markets, shared minimum data elements could include:

- basic firm identifiers and group structure maps
- custody arrangements and where keys are controlled
- client asset segregation methods and reconciliation results
- incident reports, including breaches and outages
- reserve and redemption data for stable-value arrangements
- leverage and liquidation metrics for venues offering margin
- major related-party exposures and flows

This does not require a single global database. It requires that each supervisor collects a core set in a comparable way, so information can be shared and understood.

3.10.2 Routine incident reporting and cross-border escalation paths

Crypto incidents move fast. A breach can drain assets in minutes. A stable-value token can face a redemption wave within hours. A venue can freeze withdrawals and trigger panic.

A mature coordination setup would include:

- clear incident reporting thresholds
- contact points for rapid cross-border notifications
- agreed templates for what information is shared
- plans for public communications where needed

This is common in banking supervision. It is less mature in crypto supervision.



3.10.3 Colleges for system-relevant firms and arrangements

Not every firm needs a college. Colleges are resource-heavy. They should be reserved for:

- firms with large retail reach across borders
- firms that combine multiple roles (trading, custody, lending)
- stable-value arrangements with broad use

For these cases, colleges can reduce blind spots and allow supervisors to align on priorities.

3.10.4 A clearer boundary between supervisory sharing and enforcement sharing

Supervisors share information for two reasons.

- Routine supervision: assessing governance, controls, and risk.
- Enforcement: investigating suspected breaches and taking action.

The legal basis for sharing can differ between these reasons.

A mature setup would make this clear and would provide legal bases for each, with confidentiality protections.

3.10.5 Interoperable AML information exchange

Travel Rule work shows that AML information exchange is partly a technical coordination problem.

A mature setup would support interoperability across firms and jurisdictions. It would also address privacy and data transfer issues clearly so compliance does not depend on uncertain legal interpretation.



3.10.6 Capacity building and shared tooling

Finally, baseline implementation depends on capacity.

Not every supervisor has the same resources. Not every jurisdiction has the same ability to collect market data or investigate complex protocols.

A mature baseline layer therefore includes capacity building: training, shared typologies, and shared tools.

This is not charity. It reduces arbitrage. If weak links remain, risk will route through them.

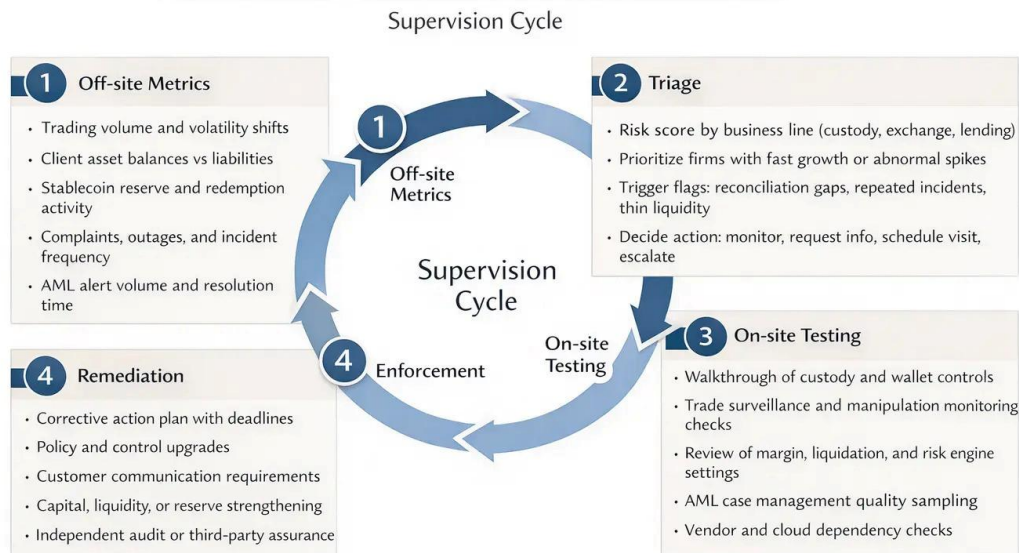
3.11 Chapter synthesis: why the baseline layer drives convergence and why gaps persist

Global standards and coordination form the hidden backbone of crypto regulation.

The baseline layer is a stack. Financial integrity expectations bring service providers into AML supervision and require that customer information and transfer information be available. Market conduct expectations push toward strong custody safeguards, conflict controls, market abuse deterrence, and clear communications to customers. Stability expectations push toward oversight that matches function, scales with system relevance, and covers stable-value arrangements as a priority bridge to payment and liquidity risk.



Supervisory Toolkit for Crypto Market Oversight



All rights reserved by the Blockchain Council.

The baseline layer drives convergence through peer review, reputational pressure, and the practical needs of cross-border supervision.

At the same time, gaps persist for three reasons.

First, the frontier is hard. DeFi, self-held wallets, and cross-chain systems do not map neatly onto entity-based rulebooks.

Second, adoption is faster than supervision. A law can be passed quickly. Building supervisors, data systems, and enforcement practice takes time.

Third, coordination is uneven. Cooperation tools for financial crime are mature. Cooperation tools for stability monitoring are less mature. Data collection is not yet consistent.



These gaps matter because crypto markets move fast and cross borders easily. Weak links invite arbitrage and make enforcement difficult.

For the rest of this research paper, the baseline layer should be treated as the floor that shapes domestic choices. It explains why many domestic regimes look similar in structure even when their statutes differ. It also explains why the hardest policy questions are increasingly about making rules work in practice: supervision capacity, enforcement reach, and cross-border coordination that goes beyond one-off requests.

Conclusion

Crypto regulation has moved past the phase where the main question is whether governments will regulate at all. The current question is how to build a perimeter that holds in a market that is global, fast-moving, and shaped by software.

The first lesson is that definitions are policy. The hardest disputes are often not about the fine print of a rule, but about the category a token or service falls into. Once a token is treated as a security-like instrument, an e-money-like instrument, or a non-security crypto-asset, an entire set of obligations follows. Once an actor is treated as a service provider, a full set of governance, conduct, and AML duties follows. That is why scope, definitions, and taxonomy are not introductory filler. They are the control panel of the whole regime.

The second lesson is that most jurisdictions are converging on function. Different statutes still use different labels, and agencies still argue about classification, but the underlying move is similar: identify the economic function and impose outcomes that match the harm model. Markets must be fair. Customers must not be misled. Client



assets must not be mixed with firm assets. Custody must be safe enough to keep losses from becoming routine. Stable-value arrangements must not invite money-like use without money-like governance of reserves and redemption. Credit-like products must not hide leverage and maturity mismatch behind simple “yield” branding. AML/CFT controls must apply to gatekeepers who can implement them.

The third lesson is that market structure matters as much as product type. Crypto markets have often grown through vertical combination: venue operation, brokerage, house trading, custody, lending, stable-value issuance, and staking services housed in one group. This structure concentrates conflicts and concentrates operational risk. It also concentrates power: a small number of private actors can set key rules for market access, listing, liquidations, and custody terms. A regime that focuses only on token labels and ignores market structure will miss the main drivers of harm.

The fourth lesson is that custody is the point where theory meets loss. The largest waves of harm in crypto markets have not come only from bad trades; they have come from loss of access, misuse of client assets, hacks, and insolvencies where customers learn too late that their claims were weak. Any regime that seeks durable legitimacy must make custody duties clear, enforce segregation, require auditability, and force firms to be honest about what happens in a failure.

The fifth lesson is that stable-value arrangements sit at the boundary between crypto markets and payments. That boundary changes the risk model. When a product invites money-like use, users expect quick redemption and low value variation. If those expectations are not backed by reserve quality, clear redemption terms, sound



governance, and operational resilience, the result can be a run. A run can force reserve asset sales and transmit stress beyond the crypto market itself. Stable-value oversight is therefore not a niche issue; it is the core bridge where crypto can become system-relevant.

The sixth lesson is that financial crime controls provide the earliest and widest perimeter, but they are not self-executing. AML/CFT rules can be written quickly, yet effectiveness depends on supervision capacity, sanctions for weak programs, and operational coordination between firms and jurisdictions. Travel Rule compliance, in particular, shows the gap between adopting a rule and building a working system that allows counterparties to exchange required information safely and consistently.

The seventh lesson is that cross-border coordination is the real constraint. Crypto firms can serve customers remotely. They can split operations across affiliates in different jurisdictions. Liquidity and price formation are global. As a result, a domestic regime can be well drafted and still be undermined by offshore service provision and group structure games. Baseline standards help by creating a common floor and a shared vocabulary, but cooperation still needs working tools: information exchange agreements, alert networks, structured colleges for system-relevant firms and stable-value arrangements, and shared minimum data that allows supervisors to see group-wide exposures.

These lessons point to a practical agenda for regulators and policymakers.

- Build definitions and taxonomy that follow economic substance and can be applied repeatedly to new products.



- Regulate by activity as well as by asset type, so that core services—trading, brokerage, custody, transfer, lending, staking, stable-value issuance—are supervised where they are offered to the public.
- Treat market structure and conflicts as first-order risks, especially where firms combine venue operation, dealing, custody, and lending.
- Make custody standards concrete, testable, and enforceable, including segregation, key control, reconciliation, incident response, and insolvency clarity.
- Calibrate stable-value oversight to real redemption and reserve design, not to marketing labels.
- Treat AML/CFT controls as a full stack—licensing or registration, supervision, enforcement, and operational capability—rather than as a paper rule.
- Invest in cross-border coordination tools and shared data elements so that supervision can match the market’s cross-border reality.

Crypto regulation will continue to evolve because the market will continue to evolve. New token types will appear. Market functions will be packaged in new ways. Some activity will migrate toward more direct on-chain execution. At the same time, the core policy aims are not mysterious. They are the same aims financial regulation has pursued for decades: reduce fraud, control conflicts, protect customers, keep markets orderly, stop financial crime, and contain system-wide spillovers.



The difference is the operating environment. Crypto markets move faster, cross borders more easily, and rely more heavily on software and private governance. That difference does not require regulators to invent a new philosophy. It requires regulators to be clear about scope, disciplined about function, and serious about supervision and coordination. If those elements are in place, the market can develop under rules that are legible, enforceable, and capable of protecting users without pretending that every risk can be removed.

