

How mature is your crypto compliance operation?

Nine capability pillars, five maturity levels, and a score from 0 to 100




Industry Recognition 2025–2026


Recognised by leading analyst firms, government bodies and independent research organisations




Ranked Exceptional in Liminal's 2026 Age Estimation Index



**Ranked Exceptional in Liminal's
2026 Age Verification Index**



**Broadest global reach in the 2025
KuppingerCole Extended IDV report**



Recognised as a Leader across four
G2 Summer 2026 reports



Differentiated by Gartner on document diversity and country coverage



240+
Countries
& Territories

10k+
Document
Types



**First EU company to achieve iBeta
Level 3 Conformance to ISO/IEC
30107-3, with 0% APCER**

**iBeta Level 3
Conformance**

0% APCER



**Ranked Top 5 in the DHS RIVR 2025
for identity validation**

0%
Extraction
Failure

0.67%
Worst-Case
FNMR



What is inside this guide

01	The framework in brief	03
02	Why compliance is the floor	04
03	What the model is	06
04	The five maturity levels	07
05	How a level is claimed	09
06	The Maturity Index	10
07	Reading your results	11
08	The nine capability pillars	14
09	Rate your own operation	15
10	Pillar reference	16
11	Indicative metric anchors	02
12	What higher maturity buys	26
13	Architecture and consolidation	27
14	What the model does not cover	30
15	Glossary	31
16	Methodology and limitations	32
17	About Shufti	33

The pillar reference pages carry one page per capability pillar, in the order the domains appear: Govern (pillars 01 to 03), Know (04 to 06), Prove (07 and 08), and the Foundation (09).

The framework in brief

Regulatory compliance is the floor of a crypto compliance operation. What sits above it decides cost, resilience, and how far the operation can grow. This framework measures that.

THE PROBLEM

The Financial Action Task Force reports Travel Rule legislation in force or in progress in 102 of 109 surveyed jurisdictions. A control almost every material market has legislated no longer marks an operation out as advanced, it marks it as present. Two providers can both satisfy the rule and run completely different operations: one clears a compliant transfer in minutes at a few cents of cost, the other takes a week and adds headcount for every rise in volume.

WHAT THE FRAMEWORK DOES

It rates an operation across nine capability pillars, grouped in three domains on a shared data foundation, against five maturity levels from Reactive to Adaptive. Every level from 2 upwards carries a binary gate and an evidence test, so a rating can be failed rather than argued, and an internal audit function can check it.

WHAT IT PRODUCES

A rating on each of the nine pillars, which is the primary output. An Index from 0 to 100 as the shorthand for a board, an auditor, or a partner. And a flag on any pillar sitting far enough below the rest to cap the whole operation.

THE CENTRAL FINDING

The profile matters more than the score. Two operations can share an Index of 44 and look nothing alike, one with a single critical gap hiding behind a respectable average, the other strong in two pillars and weak in seven. Any pillar two or more levels below the overall level caps the operation, whatever the number says.

WHAT TO DO WITH IT

Close the critical gaps, then raise every pillar to the level above, then take on the next level's organizing principle. That cycle repeats from Level 1 to Level 5, and it is the roadmap.

Who this is for

Heads of Compliance, Money Laundering Reporting Officers (MLROs), Chief Compliance Officers, and the operations, risk, and engineering leaders who carry the same responsibility at virtual asset service providers (VASPs), crypto-asset service providers (CASPs), exchanges, and financial institutions.

Why is compliance the floor rather than the finish line?

In its [seventh targeted update of 15 July 2026](#), the Financial Action Task Force (FATF) reported that 91 of 109 surveyed jurisdictions had Travel Rule legislation in force, with a further 11 in progress, taking the combined figure to 102 of 109. Adoption at that scale changes what the rule signals, because a control almost every material market has legislated no longer marks an operation out as advanced. It marks it as present.

Legislating the rule and supervising it are not the same thing. The same update found that of the 91 jurisdictions with Travel Rule legislation in force, 55 have not yet issued findings or directives or taken enforcement action against VASPs on Travel Rule compliance.

Two virtual asset service providers can both satisfy the Travel Rule and run completely different operations. One clears a compliant transfer in minutes at a few cents of cost, while the other takes a week, leans on spreadsheets, and adds headcount for every rise in volume. Both meet their obligations, but only one is built to scale.

That gap in compliance operations is what this framework measures. The market has spent years asking crypto businesses which protocol they support and whether they are compliant, and those questions are close to binary. The questions that predict cost, resilience, and growth are harder. How much manual effort does a compliant transfer take? How quickly can the operation absorb a new regulation or a new market? How well do the compliance systems actually work together? This framework turns those into a score.

FIGURE 01 · THE TRAVEL RULE ACROSS FATF, THE EU, AND THE US

REGIME	INSTRUMENT	BINDS	THRESHOLD	STATUS
FATF (global standard)	Recommendation 16, via the Interpretive Note to Recommendation 15	VASPs	USD/EUR 1,000 for occasional transactions	Travel Rule extended June 2019. Recommendation 16 revised June 2025 , in effect by end 2030
European Union	Regulation (EU) 2023/1113 (Transfer of Funds Regulation)	CASPs	No de minimis for crypto transfers	Applies from 30 December 2024
United States	31 CFR 1010.410(e) and (f) (Bank Secrecy Act)	Financial institutions	USD 3,000	In force. 2020 proposal to lower the cross-border threshold not finalized

Regulatory position verified against primary sources on 12 August 2026. Thresholds, effective dates, and scope vary by jurisdiction and change over time. Confirm the current position for your own markets before relying on any figure here.

ENFORCEMENT EXAMPLE

The gap is not theoretical.

On 11 October 2022, Bittrex settled with both [OFAC](#) and [FinCEN](#) over conduct between 2014 and 2018. FinCEN found the exchange relied on as few as two employees to manually review every transaction for suspicious activity, at volumes that reached [an average of 23,800 transactions a day worth about \\$97.9 million](#). OFAC found that Bittrex collected IP address and physical address data from every customer at onboarding and did not screen it against sanctioned jurisdictions, letting through 116,421 apparent violations worth \$263.5 million. Bittrex was a licensed, operating exchange with a compliance program. The program did not scale, and its systems did not talk to each other. The penalties came to roughly \$29.3 million.

A program built to clear one obligation, in one jurisdiction, at one threshold, on one messaging rail can be fully compliant today and structurally fragile tomorrow, because it breaks the first time a transfer crosses a border, meets a different threshold, or lands on a counterparty using a different protocol. An operation that meets the baseline is in the game. That says little about whether that operation can carry the load as volume, jurisdictions, and attack patterns multiply.

What is the Crypto Compliance Maturity Model?

The Crypto Compliance Maturity Model is a self-assessment framework that rates a crypto compliance operation across nine capability pillars and five maturity levels, then summarizes the result as a single score from 0 to 100, the Crypto Compliance Maturity Index.

The framework exists to answer the harder questions, the ones about manual effort, speed of adaptation, and how well the systems work together. It gives a compliance leader a structured way to benchmark where an operation sits today, see where the biggest gaps are, decide what capability to invest in next, and understand how to grow without cost and headcount rising at the same rate. It is built for Heads of Compliance, Money Laundering Reporting Officers (MLROs), Chief Compliance Officers, and the operations, risk, and engineering leaders who carry the same responsibility at virtual asset service providers (VASPs), crypto-asset service providers (CASPs), exchanges, and financial institutions.

What it produces

Three things, in order of importance. A rating on each of the nine pillars, which is the primary output. An Index from 0 to 100, which is the shorthand and the language to use with a board, an auditor, or a partner. And a flag on any pillar sitting far enough below the rest to cap the whole operation, because a single weak link limits real-world resilience no matter how strong everything around it looks.


What it is not

A qualitative self-assessment instrument rather than an empirical study or a validated benchmark. Every level from 2 upwards carries a binary gate and an evidence test, which makes a rating harder to inflate and easier for an internal audit function to check. The limits are set out in the methodology note at the end.

What are the five maturity levels?

The framework places a compliance operation on a five-level curve, and every level answers one question. What is actually doing the work?

FIGURE 02 · THE FIVE MATURITY LEVELS



Bar length shows maturity, not how many operations sit at each level. This framework makes no claim about how maturity is distributed across the market.

THE FIVE LEVELS IN FULL

LEVEL 1

Reactive

The person is the system. Capability lives in individuals and static documents, and work starts when something outside prompts it, whether a regulator, a bank, an alert, or a complaint. Costs are high, visibility is low, and the only way to handle more volume is to hire.

LEVEL 2

Enabled

The capability exists, but nothing joins its parts. Tooling is in place and the basics work. Travel Rule messaging runs on at least one rail, screening happens, obligations are documented, and someone is accountable. The pieces do not talk to each other, so a person still does the joining. Compliant, and not yet efficient.

LEVEL 3

Intelligent

Policy decides, and people handle exceptions. Risk-based decisioning clears the routine and routes only genuine exceptions to a human. Customer, counterparty, wallet, and jurisdiction signals feed decisions, and policy changes are made by configuration, not engineering. The operation is consistent, explainable, and noticeably leaner.

LEVEL 4

Optimized

One data layer, tuned against measured outcomes. The stack behaves as one platform on a single authoritative record, and the operation improves against numbers rather than opinion. Automation is high, exceptions are low, and audit-readiness is a standing state rather than a scramble.

LEVEL 5

Adaptive

Change is absorbed as an input. New rules, rails, jurisdictions, and risk typologies are taken on without re-architecting, often ahead of enforcement dates, and compliance scales with volume without a matching rise in headcount. Level 5 should be rare.

How is a level claimed?

Every cell carries three things.

- **Descriptor.** How the operation works at that level.
- **Gate.** The one binary sentence that decides the rating.
- **Evidence test.** The artifact, measurement, or behavior that proves the gate.

All three sit inside the pillar reference pages later in this guide. The interactive assessment online shows the gate on each option, and its detail view opens the descriptor and evidence test for all five levels at once.

Four rules for claiming a level

Levels are cumulative. An operation sits at the highest level whose gate it passes, and whose lower gates it also passes, so passing the Level 4 gate while failing Level 3 scores Level 2. Start at Level 2 and climb, then stop at the first gate you fail and score the level below it. Level 1 sits at the bottom of the scale, carries no gate, and uses a recognition test instead.

Compound gates require every clause. Several gates carry more than one condition, and a failure on any single clause fails the whole gate. Partial satisfaction scores the level below, and the partial state should be recorded, because it is usually the most useful finding on that pillar.

A gate whose triggering event has not occurred is not met. Some gates ask what happened the last time the operation launched a market, onboarded a rail, or met a new typology. If that event has never arisen, score the level below and record the pillar as untested, not assumed. An untested capability is a real finding, not an administrative gap.

Instance-based gates test routine practice. Where a gate asks for a documented instance in the last twelve months, that instance is evidence the practice is routine. One deliberate exercise staged to pass the assessment does not satisfy it.

How to score honestly

Score the level the operation has fully achieved today, not the highest level for which some evidence exists somewhere, and not the level the roadmap will deliver. Partial achievement of a level scores the level below. Where the honest answer sits between two levels, take the lower one, because the framework is more useful when it is uncomfortable. Score the operation as a whole, and where a capability is mature in one business line and absent in another, score the weaker one and note the split. That split is itself the finding.


How does the Crypto Compliance Maturity Index work?

The Index summarizes nine pillar ratings into one number from 0 to 100. It is a directional summary rather than a validated benchmark, and its value depends on the honesty of the inputs. The nine-pillar profile is the primary output, and the Index is the shorthand for it.

Rate the operation on each pillar from 1 to 5, average the nine ratings, then convert to a 0 to 100 scale.

$$\text{Index} = ((\text{average pillar level} - 1) / 4) \times 100$$

FIGURE 03 · THE INDEX BAND SCALE, 0 TO 100



An operation sitting at Level 1 across the board lands at 0, and one at Level 5 across the board reaches 100. The Index is rounded to the nearest whole number before it is banded, and a half value rounds to the higher whole number. All nine pillars carry equal weight in this version.

Overall level throughout this framework means the band the rounded Index falls into, rather than the rounded or floored average of the pillar levels. The critical-gap rule below is measured against it.


How do you read your results?

The framework gives three things to act on, and the order matters.

The Index and level are the headline. They place the operation overall and give a shared language for the board, auditors, and partners. A number and a named level travel further in a governance conversation than a folder of process documents.

The pillar profile is the diagnosis. Viewed as a nine-spoke radar, it shows where the operation leads and where it lags. Two operations can share an Index and look nothing alike, which is why the profile matters more than the number. Three short examples make the point, and the first and last share an Index of 44.

FIGURE 04 · THREE PROFILES



THE THREE PROFILES READ AGAINST THE SAME INDEX

— A · The hidden gap

INDEX 44

CRITICAL GAP

Eight pillars at Level 3, wallet and transaction assurance at Level 1. Average 2.78.

The Index reads as a solid Level 3. The wallet pillar two levels below everything else is a critical gap, and it exposes the business where a crypto operation can least afford it. The number flatters the operation, and the profile tells the truth.

— B · The even climber

INDEX 50

NO FLAG

All nine pillars at Level 3.

A clean Level 3 with no critical gap and no uneven-profile flag. This operation is genuinely level, and its next move is to raise every pillar rather than patch a hole.

— C · The uneven build

INDEX 44

UNEVEN PROFILE

Regulatory readiness and counterparty operations at Level 5, reporting at Level 3, the remaining six pillars at Level 2.

The same Index as Profile A. No pillar sits two or more levels below the overall level, so no critical gap fires. The spread of three levels between the strongest and weakest pillars triggers the uneven-profile flag instead. This operation has invested hard in two capabilities and left seven behind.

The flags are the sequence. Any pillar two or more levels below the overall level is a critical gap, and a weak link caps real-world resilience no matter how strong the rest of the operation looks. In compliance, a single hole, such as no real wallet screening, is a risk that strength elsewhere cannot offset. Where no critical gap fires but the highest pillar sits three or more levels above the lowest, the profile is flagged as uneven instead, which says the operation has invested unevenly.

For example, an operation with mature onboarding, screening, and reporting but a Level

1 wallet pillar will verify a customer to a high standard, clear them through sanctions screening, then send funds to an address it has never screened. The strong pillars will record that transfer accurately. None of them can stop it.

The most efficient way up is usually to lift the lowest pillars to the median before pushing the whole operation to the next level, because filling a gap tends to reduce more risk and free more capacity per unit of effort than improving a pillar that is already strong. Close the gaps, then raise every pillar. That sequence repeats at every level, and it is the roadmap.

FIGURE 05 · THE ROADMAP FROM LEVEL 1 TO LEVEL 5




The cycle repeats at every level. Close any critical gap, raise every pillar to the level above, then take on the next level's organizing principle.

What are the nine capability pillars?

Maturity is not one thing, so the framework measures nine capabilities, grouped in three domains. The reference pages that follow rate each one, and every level there carries the full descriptor, gate, and evidence test.

FIGURE 06 · NINE PILLARS IN THREE DOMAINS, ON ONE FOUNDATION



Rate your own operation

Nine pillars, one at a time. Mark the highest level whose gate the operation passes today, and whose lower gates it also passes. The pillar reference pages that follow carry the descriptor, gate, and evidence test for every cell.

	CAPABILITY PILLAR	L1	L2	L3	L4	L5	LEVEL
01	Governance, accountability and people	☐	☐	☐	☐	☐	_____
02	Risk assessment, policy and decisioning	☐	☐	☐	☐	☐	_____
03	Regulatory and jurisdictional readiness	☐	☐	☐	☐	☐	_____
04	Customer due diligence and onboarding	☐	☐	☐	☐	☐	_____
05	Counterparty and network operations	☐	☐	☐	☐	☐	_____
06	Wallet and transaction assurance	☐	☐	☐	☐	☐	_____
07	Screening, monitoring and suspicious activity	☐	☐	☐	☐	☐	_____
08	Reporting and auditability	☐	☐	☐	☐	☐	_____
09	Systems integration and data unification	☐	☐	☐	☐	☐	_____

CALCULATE THE INDEX

Add the nine levels, divide by nine for the average, then apply the formula. Round to the nearest whole number before banding.

Sum of nine _____

Average _____

Index _____

CHECK THE FLAGS

Any pillar two or more levels below the overall level is a critical gap. Where no critical gap fires but the highest pillar sits three or more levels above the lowest, the profile is uneven.

Overall level _____

Critical gaps _____

Uneven _____

Index = ((average pillar level - 1) / 4) × 100. An interactive version of this worksheet, which scores itself, is on the guide page online.

PILLAR 01

DOMAIN A · GOVERN

Governance, accountability and people

COVERS named accountability for financial crime risk, the resourcing and competence of the compliance function, board and executive oversight, training, independent testing, oversight of outsourced providers, and the controls over internal financial crime risk including bribery, employee collusion, privilege abuse, and the whistleblowing channel.

1 Reactive

DESCRIPTOR No single named owner holds financial crime risk, or the role exists on paper and carries no authority. The function is under-resourced and reacts to events, training is ad hoc, and the board hears about compliance when something has already gone wrong. Third parties are onboarded without compliance oversight.

TEST No individual holds documented, board-acknowledged accountability for financial crime risk, or the role exists and has not reported to the board or a board committee in the last twelve months.

EVIDENCE Ask three people who is accountable. You get different answers, or a name with no terms of reference behind it. Or the name holds, and no board or committee minute from the last twelve months records a report from them.

2 Enabled

DESCRIPTOR An MLRO or compliance officer is appointed with a documented remit, and policies and a training program exist. The function is separated from the business it oversees. Oversight is periodic and narrative: the board sees activity, not measures, and independent testing is infrequent. Vendor oversight stops at the contract.

GATE A named MLRO or equivalent holds documented accountability and reports to the board or a board committee at least annually.

EVIDENCE Terms of reference for the role, plus minuted board or committee papers from the last twelve months.

3 Intelligent

DESCRIPTOR Responsibilities are defined across the lines of defense, so each capability has a named owner and defined triggers drive escalation. Board and executive reporting is scheduled and carries metrics, not only narrative. Training is role-specific and independent testing follows a plan.

GATE Every capability in this framework, or in the operation's own equivalent taxonomy, has a named owner, and escalation to the MLRO happens on defined triggers rather than at an individual's discretion.

EVIDENCE A current responsibility map covering all nine pillars, plus an escalation policy stating the triggers, plus the most recent board pack showing compliance metrics.

4 Optimized

DESCRIPTOR Governance runs on the same data the operation runs on, so board and executive reporting is drawn from live measures. Control effectiveness is tested against outcomes, not existence. The operation plans resourcing against measured workload, and monitors third parties continuously against defined risk indicators.

GATE The compliance metrics in the most recent board or committee pack were produced without manual re-keying or spreadsheet assembly, and changing a control because measurement showed it was not working is routine practice, evidenced by at least one documented instance in the last twelve months.

EVIDENCE The reporting pipeline itself, plus a documented instance of a control changed on the strength of measured outcomes, recording the measurement that prompted the change and the measurement that followed it.

5 Adaptive

DESCRIPTOR The governance model adapts. New risks, products, markets, and regulatory expectations trigger a defined reassessment of accountabilities, resourcing, and competence. Nothing is absorbed into existing roles by default. Compliance is consulted before product and market decisions, and the function's authority is real enough to stop things.

GATE A new product, asset, or market triggers a documented reassessment of accountabilities, resourcing, and competence before launch rather than after it, and compliance holds a documented right of hold over that launch.

EVIDENCE The pre-launch reassessment record for the most recent product or market, plus the documented terms of the hold right showing who may exercise it and on what grounds. Where a hold has been exercised, the recorded decision, which is the strongest available evidence but is not required. If nothing has launched, the gate is not met and the pillar is recorded as untested.

Risk assessment, policy and decisioning

COVERS the business-wide risk assessment, the customer, asset, and jurisdictional risk models built from it, the policy that expresses them, and whether that policy reaches the point of decision without a human in the way.

1 Reactive

DESCRIPTOR There is no current business-wide risk assessment, or one exists as a document that drives nothing. Decisions are made case by case, so outcomes vary by who is on shift. Policy lives in documents, not in the system. Risk is not tiered, so high-risk and low-risk activity gets similar handling.

TEST There is no documented business-wide risk assessment, the current one predates a material change in the business, its products, or its markets, or customers do not carry a risk rating derived from it.

EVIDENCE The date on the current assessment set against the date of the last material change to the business, whether any live control can be traced back to a risk the assessment identified, and whether a sample of customers carries a risk rating traceable to that assessment.

2 Enabled

DESCRIPTOR A business-wide risk assessment exists and is refreshed periodically, and policy is documented. Some rules are automated, such as basic thresholds and simple blocks, but most decisions still route to a human, and policy changes need reconfiguration or engineering work. Risk is loosely tiered, yet a large share of low-risk activity still consumes analyst time.

GATE A current business-wide risk assessment exists, and every customer carries a risk rating derived from it.

EVIDENCE The assessment, the customer risk-rating methodology, and a sample showing ratings actually applied to live customers.

3 Intelligent

DESCRIPTOR Policy is expressed as a configurable rule set that drives risk-based decisioning. Low-risk activity clears automatically and analysts see genuine exceptions. The score blends signals across customer, counterparty, wallet, jurisdiction, asset, and behavior. Decisions are consistent and explainable.

GATE A compliance user can change a policy rule and put it into production without an engineering release.

EVIDENCE A change log entry showing a policy change made by a named compliance user, with the date it took effect in production.

4 Optimized

DESCRIPTOR Decisioning is tuned against outcomes, so the exception rate falls without raising risk. Policies are versioned, testable, and deployed centrally, and rules and models are monitored for drift and false-positive burden. Risk models are validated against what actually happened.

GATE Rule and model performance is measured on an ongoing basis, and retiring or retuning rules on the strength of that measurement is routine practice, evidenced by at least one documented instance in the last twelve months.

EVIDENCE Rule performance reporting, plus the change record for the retired or retuned rule showing the measurement that drove it.

5 Adaptive

DESCRIPTOR The decisioning layer improves continuously under documented governance covering testing, approval, and audit, as new risks, typologies, products, and regulations appear. Automation covers the overwhelming majority of decisions at high accuracy, and the system surfaces where policy itself should evolve without waiting to be told.

GATE Newly identified risk typologies reached live decisioning within a service level that was documented and approved before the event, under a documented model governance process, on each of the last three occasions.

EVIDENCE The model governance standard, plus a dated example tracing a new typology from identification to live decisioning, with the elapsed time against the stated service level. If no new typology has arisen, the gate is not met and the pillar is recorded as untested.

PILLAR 03

DOMAIN A · GOVERN

Regulatory and jurisdictional readiness

COVERS keeping pace with FATF Recommendation 16 and each market's rules and thresholds, and how fast the operation adapts when they change.

1 Reactive

DESCRIPTOR Obligations are tracked informally, often per market and in someone's head or a static document. New rules or threshold changes surface late, usually when a regulator, bank, or auditor raises them, and implementing them is a scramble. No clear owner watches the regulatory horizon.

TEST Obligations are not documented and mapped to controls for every market the operation serves, or no one holds documented responsibility for monitoring regulatory change in those markets.

EVIDENCE Ask who watches the horizon. Either no name comes back, or the named person holds no record of what they reviewed and when. Or the name holds and there is no obligations register covering every live market.

2 Enabled

DESCRIPTOR Core obligations for the primary markets are documented and mapped to controls. The operation covers where it operates today, but standing up a new market is a heavy manual project, and cross-border differences are handled case by case. Change tracking is periodic, not continuous.

GATE Obligations for every market the operation currently serves are documented and mapped to the specific controls that satisfy them.

EVIDENCE An obligations register naming the instrument, the requirement, and the control, with no live market missing from it.

3 Intelligent

DESCRIPTOR Requirements are maintained as structured, jurisdiction-specific rules that drive the operation. New jurisdictions onboard from a repeatable playbook, and threshold and scope differences are configured, not hard-coded. A named owner watches for change, so change is anticipated.

GATE Every jurisdictional difference the operation handles, thresholds and scope included, is held as configuration rather than written into code or process documents, and a documented playbook exists for onboarding a new market.

EVIDENCE The configuration itself, the playbook, and the date of the most recent market onboarded using it.

4 Optimized

DESCRIPTOR Jurisdictional rules are centrally managed and pushed into the transaction flow automatically, so a threshold change updates behavior everywhere without re-engineering. A new market goes live quickly and at low risk. Regulatory change is monitored continuously and assessed for impact before it takes effect.

GATE A threshold change is made once and takes effect across every affected flow without a code release, and the elapsed time from regulatory decision to live effect is measured as a standing metric.

EVIDENCE A change record showing a single jurisdictional change propagating to production, plus the elapsed-time metric tracked across changes rather than calculated for this one. If no threshold change has arisen, the gate is not met and the pillar is recorded as untested.

5 Adaptive

DESCRIPTOR Regulatory change is a normal, continuous input. New rules, jurisdictions, and thresholds are absorbed rapidly, often ahead of enforcement dates. Coverage extends into new markets as a configuration choice. A new jurisdiction becomes a commercial decision rather than a compliance constraint.

GATE Regulatory changes are implemented ahead of their enforcement dates without a dedicated project, evidenced on each of the last three changes, and the most recent market launch required no new engineering.

EVIDENCE The dated implementation record set against the published enforcement date, plus the elapsed time and engineering effort for the most recent market launch. If fewer than three such changes have arisen, or no market has launched, the gate is not met and the pillar is recorded as untested.

Customer due diligence and onboarding

COVERS identifying and verifying customers at onboarding, whether individuals, businesses, or institutional and nested relationships such as another provider onboarded as a customer, establishing beneficial ownership, applying enhanced due diligence to higher-risk relationships, screening customers at onboarding, keeping that picture current across the life of the relationship, and exiting it, covering de-risking, offboarding, and account termination.

1 Reactive

DESCRIPTOR Identity details are collected but verification is light or manual, and documents are eyeballed rather than checked. Business customers are onboarded without establishing beneficial ownership. There is no enhanced due diligence path for higher-risk relationships, and once a customer is onboarded the record is never revisited unless something goes wrong.

TEST Customer identity is not verified against an authoritative or tamper-evident source before the account can transact, or corporate customers are onboarded without their beneficial owners recorded.

EVIDENCE Trace one recent onboarding end to end. If verification was a person looking at an uploaded image, with no document authentication or database check behind it, Level 1 applies. Trace one corporate onboarding as well. If no beneficial owners were recorded, Level 1 applies.

2 Enabled

DESCRIPTOR Identity verification runs on a tool, with document authentication and a liveness or biometric step for individuals and registry checks for businesses. Beneficial ownership is collected for corporate customers. Enhanced due diligence exists as a documented process but runs manually and inconsistently. Refresh is calendar-driven at best.

GATE Every customer is identity-verified through a system before transacting, and every corporate customer has beneficial owners recorded.

EVIDENCE The onboarding flow showing verification as a hard block rather than a step that can be skipped, plus beneficial ownership records for a sample of corporate customers.

3 Intelligent

DESCRIPTOR The onboarding path is driven by the customer risk rating, so standard customers clear automatically and higher-risk relationships route to enhanced due diligence on defined triggers. Sanctions and politically exposed person screening happens inside the same flow. The customer record updates on events, not only on a calendar.

GATE The depth of due diligence applied is determined automatically by the customer's risk rating, and enhanced due diligence triggers are defined in the system rather than left to the analyst.

EVIDENCE The rule set mapping risk rating to due diligence path, plus a sample of enhanced due diligence cases each showing the trigger that fired.

4 Optimized

DESCRIPTOR Onboarding, screening, and the customer risk rating draw on the same record, so a change anywhere updates the picture everywhere. The operation measures and tunes onboarding performance, tracking drop-off, pass rate, and manual-review rate, then acting on what they show. Scheduled periodic review is supplemented or replaced by review driven by change in the underlying data.

GATE Customer risk ratings change outside scheduled reviews when underlying data changes, and onboarding conversion and manual-review rates are measured, with acting on them routine practice evidenced by at least one documented change in the last twelve months.

EVIDENCE An instance of a risk rating changing on an event rather than a date, plus the onboarding metrics reporting and the change made in response to it.

5 Adaptive

DESCRIPTOR New customer types, jurisdictions, document types, and verification methods are added without rebuilding the flow. A new market or product cohort onboard as a configuration change. Verification methods adapt as fraud patterns emerge, including synthetic identity, presentation attacks, and injection attacks.

GATE The most recent market, customer type, verification method, or document type was added to the onboarding flow as configuration without an engineering release, and a countermeasure for a newly emerged attack type was added the same way.

EVIDENCE The configuration record for the most recent addition, showing no engineering release, plus the dated attack type and the configuration record for the countermeasure. If neither has arisen, the gate is not met and the pillar is recorded as untested.

Counterparty and network operations

COVERS identifying the provider on the other end of a transfer, connecting across Travel Rule protocols and networks, and exchanging originator and beneficiary data successfully.

1 Reactive

DESCRIPTOR There is no systematic way to identify the provider on the other end. Travel Rule data, if it moves at all, is handled ad hoc, by email or spreadsheet or not at all, and counterparty checks happen only when a bank or regulator asks. Most outbound transfers carry little or no originator and beneficiary exchange.

TEST The operation is not connected to any Travel Rule protocol, cannot exchange originator and beneficiary data with counterparties on one, or has no systematic method for identifying the provider controlling a destination address.

EVIDENCE Ask how the last ten outbound transfers to unfamiliar addresses were handled. If the answer is case by case, by email, or not at all, or no protocol membership can be produced, Level 1 applies.

2 Enabled

DESCRIPTOR The operation is connected to at least one Travel Rule protocol or network and can exchange data with counterparties on that same rail. Identification works for in-network cases, but anything on another protocol, or to a self-hosted wallet, drops to manual handling. Interoperability is the limit.

GATE The operation is connected to at least one Travel Rule protocol and can exchange originator and beneficiary data with counterparties on it.

EVIDENCE Protocol membership, plus a sample of completed exchanges with timestamps.

3 Intelligent

DESCRIPTOR Counterparty discovery is systematic. The operation resolves which provider controls a destination address and routes across multiple protocols. Self-hosted-wallet transfers follow a defined path, and each counterparty carries a risk read across sanctions, licensing, and jurisdiction that informs whether to send or hold.

GATE Counterparty discovery runs automatically across more than one protocol, and self-hosted-wallet transfers follow a defined path rather than being handled as exceptions.

EVIDENCE The discovery configuration listing more than one protocol, the documented self-hosted-wallet path, and a sample showing it actually followed.

4 Optimized

DESCRIPTOR Interoperability spans the major rails without manual routing, and delivery success is measured and improved. Counterparty intelligence is enriched and wired into the stack, so a send or hold decision draws on the same data as sanctions and business verification. Unmatched-counterparty and sunrise cases have real fallback workflows.

GATE Travel Rule delivery success rate is measured continuously, and acting on that measurement is routine practice, evidenced by at least one documented action in the last twelve months.

EVIDENCE Delivery success reporting with a time series, plus the change made and its measured effect.

5 Adaptive

DESCRIPTOR Counterparty and network operations run without manual routing and improve against their own measurements. New protocols, networks, and jurisdictions onboard quickly without re-architecting. Delivery stays high and predictable even when the counterparty is not specified, using analytics-based fallback. Reach and risk intelligence are an advantage rather than a cost.

GATE The most recent Travel Rule protocol or network was onboarded without changes to the existing integrations, and unmatched-counterparty and sunrise cases have automated fallback rather than dropping to manual handling.

EVIDENCE The date and engineering effort for the most recent protocol onboarded, plus the fallback logic and the share of unmatched cases it resolves without a human. If no protocol or network has been onboarded since the first, the gate is not met and the pillar is recorded as untested.

Wallet and transaction assurance

COVERS verifying self-hosted wallet ownership, screening wallets and transactions for risk and sanctions exposure, and whether those two operate as one assured step or as two disconnected checks.

1 Reactive

DESCRIPTOR Self-hosted wallet transfers are largely taken on trust, with little or no proof of ownership, and wallet or transaction risk is not screened before funds move. Any check is manual and after the fact, prompted by an alert or a complaint. Exposure to sanctioned or illicit addresses goes largely unseen.

TEST Wallet or transaction risk is not screened before funds move, or self-hosted transfers proceed without any documented proof of ownership.

EVIDENCE Trace a recent outbound transfer to a self-hosted wallet. If no screening result predates the transfer timestamp, or no documented ownership proof method was applied to it, Level 1 applies.

2 Enabled

DESCRIPTOR Basic ownership proof exists for some flows, and wallet screening runs, but the two are applied inconsistently and never joined. Coverage is patchy across asset types and channels, and edge cases fall to manual review. Screening is point-in-time, not continuous.

GATE Wallet screening runs before funds move on every in-scope transfer, and a documented ownership proof method is applied to self-hosted transfers.

EVIDENCE Screening timestamps that precede transfer timestamps across a sample with no gaps, plus the documented ownership proof method.

3 Intelligent

DESCRIPTOR Ownership verification is applied systematically, with more than one method to fit the situation, and risk scoring across illicit exposure, sanctions, and source-of-funds signals feeds the decision before a transfer completes. Risky wallets are held or escalated on defined rules.

GATE More than one ownership verification method is available and selected by case, and a risk score drives the send, hold, or block decision through defined rules rather than analyst discretion.

EVIDENCE The method list and its selection logic, the rule set connecting risk score to outcome, and a sample of holds each showing the rule that fired.

4 Optimized

DESCRIPTOR Ownership verification and risk screening operate as one assured step, with the method chosen automatically by context and results integrated with sanctions and counterparty data. Screening extends beyond onboarding into ongoing exposure monitoring, and the operation actively reduces false positives. Manual touch is the exception.

GATE Wallet exposure is monitored on an ongoing basis after the transfer rather than only at the point of transfer, and false-positive burden is measured, with acting on it routine practice evidenced by at least one documented tuning change in the last twelve months.

EVIDENCE Alerts generated after a transfer completed, plus false-positive measurement showing the tuning change and its effect.

5 Adaptive

DESCRIPTOR Assurance runs continuously, and new evasion patterns and address typologies are absorbed as they emerge. Ownership and risk are assessed together with high automation and low friction, and assurance extends to new asset types or chains without rebuilding. Illicit-exposure risk is managed predictively.

GATE The most recent chain or asset type was brought under assurance without changes to the existing screening pipeline, and newly identified evasion typologies reach live screening within a service level documented and approved before the event, evidenced on the last three occasions.

EVIDENCE The onboarding record for the most recent chain or asset with its engineering effort, plus the dated typology and the elapsed time to live coverage. If no new chain or asset type has been added, or no new typology has arisen, the gate is not met and the pillar is recorded as untested.

Screening, monitoring and suspicious activity

COVERS sanctions and watchlist screening of customers and related parties, sanctions actioning including blocking, rejecting, freezing, and reporting blocked property to the competent authority, jurisdictional access controls such as restricted-market blocking and proxy or VPN detection, ongoing transaction monitoring for suspicious behavior, alert investigation, the decision to file a suspicious activity or suspicious transaction report, and what follows it, covering law enforcement and regulator requests, production and freezing orders, and tipping-off controls.

1 Reactive

DESCRIPTOR Sanctions screening happens at onboarding if at all, and is not repeated as lists change. There is no systematic transaction monitoring, so suspicious behavior surfaces through complaints, bank enquiries, or law enforcement requests. A report, on the rare occasion one is filed, is assembled by hand and improvised.

TEST Customers are not rescreened against sanctions and watchlist updates after onboarding, or alerts are not managed in a queue with recorded dispositions.

EVIDENCE The date of the last full rescreen following a list update, plus the alert queue with its recorded dispositions. If nobody can produce either, Level 1 applies.

2 Enabled

DESCRIPTOR Screening runs on a tool and repeats on list updates. Transaction monitoring exists with vendor-default or basic rules, generating alerts that a team works through a queue. Alert handling and filing decisions are documented but inconsistent, and the rules have never been tuned to this business.

GATE Sanctions screening repeats automatically on list updates, and monitoring rules generate alerts into a managed queue with recorded dispositions.

EVIDENCE The rescreening schedule and its last run, plus the alert queue showing dispositions and the named person who made each one.

3 Intelligent

DESCRIPTOR Monitoring scenarios derive from the business's own risk assessment and typologies rather than vendor defaults, and are tiered by customer risk. Alerts route by risk and type, escalation to the MLRO follows defined triggers, and the path from alert to investigation to filing decision is a defined workflow with rationale recorded at each step.

GATE Every monitoring scenario is traceable to a typology identified in the business-wide risk assessment, the alert-to-filing decision follows a defined workflow that records the rationale for closing or escalating, and documented workflows with defined response-time targets exist for law enforcement and regulator requests, freezing and restriction orders, and post-filing tipping-off controls.

EVIDENCE The mapping from risk assessment typology to monitoring scenario, one closed alert and one filed report each showing the recorded rationale, and the documented workflows with their response-time targets. Where such a request has arisen, the response-time record for the most recent one.

4 Optimized

DESCRIPTOR Monitoring, screening, and case management draw on the same customer and transaction record, so an investigator sees the whole picture without assembling it. Alert quality is measured, with false-positive rate and alert-to-report conversion tracked and used to tune scenarios. Investigation time and backlog are measured and managed.

GATE Alert quality is measured, including false-positive rate, and retuning or retiring scenarios on the strength of that measurement is routine practice, evidenced by at least one documented instance in the last twelve months.

EVIDENCE Alert quality reporting over time, plus the scenario change record and its effect on alert volume and quality.

5 Adaptive

DESCRIPTOR New typologies, including crypto-specific ones such as mixing, chain-hopping, and emerging laundering patterns, become monitoring scenarios quickly and under governance. Detection adapts as behavior changes, not on an annual tuning cycle, and new reporting formats or regulatory asks are absorbed without a project.

GATE Newly identified typologies reach a live monitoring scenario within a service level documented and approved before the event, under documented tuning governance, evidenced on each of the last three occasions.

EVIDENCE The tuning governance standard, plus a dated example tracing a typology from identification to a live scenario, with the elapsed time against the stated service level. If no new typology has arisen, the gate is not met and the pillar is recorded as untested.

Reporting and auditability

COVERS producing regulator-ready evidence, audit trails, and records on demand, including responses to regulator and law enforcement information requests.

1 Reactive

DESCRIPTOR A report or an audit trail takes manual, painful assembly, with data pulled from several systems and joined by hand, often after the request lands. Records are incomplete or inconsistent, and reconstructing why a transfer cleared is slow and uncertain. Audit preparation is a fire drill.

TEST Not every standard regulatory report the operation owes can be produced from its systems, or reconstructing why a specific transfer cleared means pulling data from more than one system by hand.

EVIDENCE Ask for the report inventory. If any report the operation owes has no system that produces it, Level 1 applies. Then pick a transfer from six months ago and ask for its full decision trail. Time how long it takes and count how many people are involved.

2 Enabled

DESCRIPTOR Standard reports and records exist and can be produced, but often with manual effort and per-system stitching. An audit trail covers core actions yet has gaps, and pulling a full history takes time. Reporting is periodic and backward-looking.

GATE Every standard regulatory report the operation owes can be produced from its systems.

EVIDENCE The report inventory, with the manual steps documented for each report.

3 Intelligent

DESCRIPTOR Key decisions and actions are logged consistently with the data behind them, so reports and audit trails are generated on demand. Filing and record-keeping obligations are mapped and systematized, and a transfer's decision path can be traced end to end.

GATE The full decision path for any transfer or customer can be produced on demand without engineering help.

EVIDENCE Ask a compliance user, not an engineer, to produce a complete decision trail for a randomly chosen case. They either can or they cannot.

4 Optimized

DESCRIPTOR Reporting and evidence are automated from the same unified data the operation runs on, so what a regulator sees matches what happened, with full lineage. Dashboards give real-time visibility into posture, and audit-readiness is a standing state. Overrides are captured with their rationale.

GATE Regulatory reporting is generated from the same data the operation runs on, with lineage from any figure in a report back to the underlying records, every manual override is captured with its rationale, and override volume and reporting exceptions are measured as standing metrics.

EVIDENCE A published report traced back to source records, a sample of overrides each showing who made it, when, and why, and the override and exception measures across at least two periods.

5 Adaptive

DESCRIPTOR Auditability is continuous and complete. Every decision is evidenced, traceable, and query-ready in real time, and reporting adapts quickly to new formats or regulatory asks. The operation can demonstrate compliance proactively, and audit effort scales with volume without scaling headcount.

GATE The most recent new reporting format or regulatory data request was satisfied from existing data without a new data collection project, and audit effort per unit of transaction volume has been measured across at least two periods and has not risen.

EVIDENCE The most recent new reporting requirement with the elapsed time to satisfy it and confirmation that no new data collection was needed, plus the audit-effort-per-volume measure across at least two periods. If no new reporting format or data request has arisen, the gate is not met and the pillar is recorded as untested.

Systems integration and data unification

COVERS how well the compliance systems work as one on shared data, versus a stack of disconnected point systems.

1 Reactive

DESCRIPTOR Compliance runs on disconnected point systems and manual hand-offs, with data re-keyed or exported between systems. There is no single view of a customer or a transfer, so analysts stitch context together by hand. Every new system adds another silo.

TEST Data is re-keyed or manually exported between compliance systems.

EVIDENCE Name the systems, then ask how data moves between any two of them. If the answer involves a person, a spreadsheet, or a file export, Level 1 applies.

2 Enabled

DESCRIPTOR Core systems are integrated at a basic level, with some data flowing between them, but gaps remain and reconciliation is still partly manual. A customer or transaction view can be assembled, but not instantly or completely. Integrations are point-to-point and brittle.

GATE Core compliance systems exchange data through system-to-system integrations rather than re-keying or file export.

EVIDENCE The integration inventory, alongside an explicit list of the remaining manual reconciliation points.

3 Intelligent

DESCRIPTOR Key functions share data through defined integrations, ideally a common interface, so identity, screening, counterparty, and wallet signals meet on one record. Analysts work from a largely unified case view, and adding a capability means connecting to the layer rather than building a silo.

GATE An analyst can see identity, screening, counterparty, and wallet signals for a customer or transfer on one screen without opening a second system.

EVIDENCE Sit with an analyst and watch them work a live case. Count the systems they open.

4 Optimized

DESCRIPTOR The stack behaves as one platform on unified data, with a single operational view across the transfer lifecycle feeding decisions, reporting, and audit. A new data source or capability onboards routinely and with little disruption. Integration overhead stops being a bottleneck to change.

GATE Every function reads the customer and transfer record from one authoritative source rather than holding its own copy, decisions, reporting, and audit are produced from that same source, and completeness of that record is measured as a standing metric.

EVIDENCE A data map showing a single authoritative record rather than per-system copies, one decision and one report traced back to it, and the record-completeness measure over at least two periods.

5 Adaptive

DESCRIPTOR Systems and data are unified and extensible by design. New systems, chains, rails, or data sources connect without re-architecting, and every function draws on the same real-time record. The architecture enables change rather than constraining it.

GATE The most recent new chain, rail, or data source was connected without changes to the systems already connected.



EVIDENCE The change records and release history for the most recent connection, showing which systems were modified, plus the integration interface or contract that made connecting without those modifications possible. If nothing new has been connected, the gate is not met and the pillar is recorded as untested.

Indicative metric anchors

Three pillars carry a metric you can use as a second check on your rating. Take the highest level whose band the operation meets. The bands start at Level 2 because a Level 1 operation has no consistent process to measure. Its numbers land in the Level 2 band.

These are reasoned anchors derived from the logic of the levels, not measured industry benchmarks, and where an anchor and a gate disagree the gate wins. A gate is a structural fact about how the operation works, whereas a metric can be distorted by product mix, customer base, or volume.

For example, an operation serving mostly low-risk retail customers can show a manual-review rate under 15% while still routing every exception to a person by hand. The metric reads Level 4. The gate, which asks whether policy makes the decision, does not.




Band widths are drawn to scale on each metric's own axis. Time uses a logarithmic scale.

What does higher maturity actually buy?

The nine pillars describe what an operation builds. They do not, on their own, prove the building was worth it. That proof lives in a second layer of measures the framework tracks but does not score. An operation does not become efficient by aiming at efficiency. It becomes efficient because the capabilities beneath it matured.

FIGURE 07 · THE OUTCOME LAYER



If an operation climbs the pillars and these numbers do not move, it has added sophistication without adding value, and the framework has surfaced a finding that matters. Maturity that does not show up in cost, scale, and customer experience is not yet doing its job.

Two of these metrics also appear as indicative bands inside the pillars they most directly reflect. Manual-review rate anchors Pillar 2, and time to stand up a new jurisdiction anchors Pillar 3. Those bands are reasoned anchors to calibrate against rather than measured industry benchmarks, and where a band and a gate disagree, the gate wins. The gate is a structural fact about how the operation works, whereas the metric is an outcome that product mix, customer base, or volume can distort.

Does higher maturity require a single platform?

Trace any pillar from Level 1 up to Level 5 and the same pattern emerges. Maturity is in large part a story about consolidation. Reactive operations are fragmented, with many systems, many manual seams, and no shared view of a customer or a transfer. Adaptive operations run on one unified data layer, where identity, screening, monitoring, counterparty, wallet, and messaging draw on the same record and inform the same decisions.

That pattern helps explain why so many operations stall between Enabled and Optimized. The problem is usually not a missing point system. It is that the point systems do not operate as one. Every disconnected system can add another silo to reconcile, another integration to maintain, and another seam where risk hides and audit trails break.


For instance, an operation that adds a dedicated wallet-screening tool to a stack where onboarding, sanctions screening, and case management already sit apart now has a fourth risk signal an analyst must open a fourth window to see. The capability is real. The decision it should inform still gets made without it. Past a certain point, more capability on a fragmented stack can lower maturity rather than raise it, because it adds coordination cost faster than it adds coverage.

The pillars, though, score outcomes rather than architecture. What the framework rewards is unified data, low coordination cost, and decisions that draw on shared context. A single platform is one way to reach that. A well-governed modular estate, with disciplined integration and a genuine shared data layer, can reach the same place. Consolidation is the common route to high maturity, not the only one, so read the pillars as a description of the destination and choose the architecture that gets there.

FIGURE 08 · A MATURE OPERATING MODEL, NINE PILLARS ON ONE SHARED LAYER



FIGURE 09 · HOW A SINGLE TRANSFER MOVES THROUGH A MATURE DECISIONING LAYER



What does this model not cover?

This framework covers the financial crime compliance operation of a virtual asset service provider, crypto-asset service provider, exchange, or financial institution handling crypto. Several adjacent disciplines sit outside it deliberately.

Out of scope, deliberately

AREA	WHY IT IS EXCLUDED
Custody, safeguarding of client assets, proof of reserves	Prudential and operational obligations rather than financial crime compliance, with a different maturity curve.
Capital, liquidity, and financial resilience	Prudential, and outside the financial crime remit entirely.
Cyber security and operational resilience	Adjacent and important, with its own established maturity frameworks. This framework would duplicate that mature work poorly.
Consumer protection, marketing, financial promotions	Conduct regulation rather than financial crime.
Tax reporting regimes	A separate reporting obligation with its own architecture.
Payment fraud and scam prevention	Related, and often run by the same team, but a distinct discipline.

Glossary

VASP, virtual asset service provider

a business that exchanges, transfers, or safekeeps virtual assets on behalf of others.

CASP, crypto-asset service provider

the equivalent term used in EU regulation.

Travel Rule

the obligation to obtain, hold, and transmit originator and beneficiary information alongside a transfer.

IVMS101

the shared data format for the originator and beneficiary information exchanged under the Travel Rule.

Counterparty discovery

resolving which provider, if any, controls the wallet on the other end of a transfer. A self-hosted wallet has no provider on the other side, so ownership is proven instead.

Self-hosted wallet

a wallet controlled directly by an individual rather than by a service provider.

Sunrise problem

the situation where the counterparty or its jurisdiction is not yet set up to receive Travel Rule data at all, which differs from two systems being unable to interoperate.

Business-wide risk assessment

the documented assessment of the financial crime risks a business faces across its customers, products, channels, and markets, from which its policy and controls derive.

Critical gap

any pillar rated two or more levels below the overall level.

Unified data layer

a shared record that lets separate compliance functions draw on the same customer and transaction context.

Methodology, scope, and known limitations

Methodology and scope

This is a qualitative self-assessment framework rather than an empirical study. The five levels and nine pillars are an analytical model, and the Crypto Compliance Maturity Index is a computed self-assessment instrument rather than a validated metric. An operation's Index is the average of its nine self-rated pillar levels, each on a 1 to 5 ordinal scale, normalized to a 0 to 100 range with the formula stated above and rounded to the nearest whole number before banding. All nine pillars are weighted equally. Equal weighting is a deliberate starting choice that keeps the instrument neutral and avoids implying a precision the inputs do not support, and a later version may re-weight if evidence supports it. The scoring space was checked by enumerating every one of the 1,953,125 possible nine-pillar combinations: the Index takes 37 distinct values and none of them lands exactly on 20, 40, 60, or 80, so no score falls ambiguously on a band boundary.

Known limitations

The inputs are ordinal self-ratings, so the Index carries the usual limits of self-assessment. The binary gates and evidence tests improve consistency between raters and make a rating harder to inflate, and they do not make the instrument validated. It has not been tested for inter-rater reliability and is not calibrated against a benchmark distribution, which is why the framework leads with the pillar profile and treats the number as a directional summary. Equal weighting is a starting choice, not an evidenced claim that the pillars matter equally. The framework carries no proprietary dataset and makes no empirical claim about how maturity is distributed across the market. The indicative metric bands on Pillars 2, 3, and 5 are reasoned anchors derived from the logic of the levels, not measured industry benchmarks. Statements about the effect of maturity on risk, cost, and customer experience are propositions of the model, offered as reasoning rather than measured findings.

HOW TO CITE THIS GUIDE

Shufti (2026). The Crypto Compliance Maturity Model: A Self-Assessment Framework for Benchmarking and Scaling Crypto Compliance Operations. Version 1.0, 27 July 2026.

This framework is offered as a decision aid, not as legal or regulatory advice. The regulatory statements in this guide were verified against primary sources on 12 August 2026. Thresholds, effective dates, and scope change over time, so confirm the current position for your own jurisdictions.



Score Your Crypto Compliance Operation, Then Close the Gaps

The framework rates nine capability pillars across five maturity levels and returns an Index from 0 to 100, your pillar profile, and any gap that is capping the operation. Shufti's identity verification, AML screening, KYB, and wallet assurance capabilities support the pillars this framework measures. Walk through the platform configuration before a vendor decision.

[Request a Demo](#)

[Take the Assessment Online](#)

Shufti is based in the United Kingdom and operates offices worldwide.



United States

16200 SW Pacific
Hwy, Suite H PMB
1053,
Tigard, OR



Cyprus

Arch. Makarios III
Avenue 229,
Limassol



Dubai

Unit 507, Level 5,
Gate District
Precinct
Building 03, DIFC,
Dubai



Hong Kong

8 Queen's Road
East, Wan Chai,
Hong Kong



Pakistan

322-G3 Main Blvd,
Johar Town, Lahore,
Pakistan



Singapore

68 Circular Road
#02-01, Singapore



United Kingdom

Office 408,
Coppergate
House, 10 Whites
Row, London E1
7NF

shuftipro.com

sales@shuftipro.com