

Q2 2026 Security & Compliance Report

Architecture of Trust — how security, compliance and risk intelligence are reshaping digital asset markets.

Contributing Partners

 Chainlink

BYBIT

 Sui

MOODY'S

ACM

DA
Insured

 1MONEY

HOWDEN

SVRN

 | EISNERAMPER

 Allium

 Stellar



"Digital assets are increasingly exposed to Public Markets: the counterparties that will attract institutional capital are those that treat security, compliance, and risk intelligence as one ongoing discipline and can produce verifiable evidence that they remain safe today, not that they were safe when last reviewed.



Dyma Budorin
Group CEO &
Co-Founder

In Q3, trust is proven continuously, or not at all."

"Q2's turbulence creates chaos, but also space: to concentrate on the challenges that matter and to back the founders and institutions building for the long term. For leadership teams entering regulated finance, the opportunity is to make security and resilience a standing capability and to lead with evidence where others still lead with reputation.



Yev Broshevan
Board Member &
Co-Founder

That is what earns the confidence of institutional counterparties, and it is the advantage that will separate the market in Q3."

Table of Contents

01	Executive Summary
02	Q2 Exploits & Vulnerability Trends
03	Compliance Landscape Updates
04	Architecture of Trust
05	Risk Intelligence for Asset Managers
06	Conclusions & Recommendations

Executive Summary

Q2 2026 was defined by a small number of very large operational failures. Roughly \$763,971,791 was stolen across 67 incidents, and most of it was traced not to flawed code but to compromised people, keys, and infrastructure. Two DPRK-associated attacks alone accounted for three-quarters of the total loss. Meanwhile, the trust signals the market has long relied on, a prior audit, years of history, value locked, failed to predict who would be exploited.

1

~\$763,971,791 across 67 incidents. That is down 25.9% from Q2 2025's \$1.03 billion and up 58.3% from Q1 2026's \$482.7 million.

2

Operational and infrastructure failures accounted for 88.3% of stolen value, or \$674,866,100 across 21 incidents.

3

Both DPRK-associated attacks exploited operational (access control) surface, accounting for \$577m in losses and 75.5% share of Q2 total lost.

4

Smart contract bugs were the most frequent type of incident, at 44 of 67, but they represented only about 11% of total losses (\$87.7 million)

5

MiCA transition ended on July 1. Just ~210-215 CASPs were authorized out of 1,200+ pre-MiCA registrations.

6

USDC is the only MiCA-compliant stablecoin out of the top-10 by market cap.


7

First case of AI prompt injection resulting in ~\$174,000 malicious transaction.

Q2 2026 demonstrated that institutional adoption is no longer constrained by blockchain technology itself, but by the industry's ability to establish and continuously verify trust. Security, compliance, governance, transparency, monitoring, and risk intelligence are increasingly converging into a unified Trust framework that regulators, exchanges, asset managers, stablecoin issuers, and financial institutions use to evaluate digital asset ecosystems.

In this report, Hacken researchers, together with contributors from Moody's Ratings, Bybit, Chainlink, Abraxas, and other industry leaders, introduce the Architecture of Trust framework, showing how modern digital asset issuers must move beyond old trust metrics of audits and compliance exercises toward integrated assurance models that support institutional participation in Web3 markets at scale.

Q2 2026: Scale and Composition of Losses




Q2 2026 vs Q1 2026 Crypto Losses

Q2 2026 was the worst quarter for Web3 since Q2 2025. Across 67 incidents, smart contract exploits, access control breaches, cross-chain infrastructure compromises, and one DNS hijack happened. The total is 26% lower than Q2 2025's \$1.03 billion, but the composition changed.

Q1 2026 was driven by a \$282M individual-wallet phishing theft. Q2 was driven by two attacks of nearly \$290M each against protocol infrastructure, both attributed to North Korean actors.

The same pattern from Q1 held in Q2: most of the money was lost to operational failures. Of the \$763,971,791 stolen, 88.3% traced to compromised keys, signers, and infrastructure.


At the same time, smart contract bugs were the most common exploit reason, accounting for 44 of 67 incidents, but they represent just about 11% of the losses. The two most expensive failures of the quarter were a multisig-signer compromise at Drift and a single-validator cross-chain bridge at KelpDAO. Neither would have been caught by a conventional code audit.



Crypto Losses by Quarter

Q2 2026 becomes the worst quarter for Web3 since Q2 2025, driven mainly by operational controls compromise, where 75.5% of losses are associated with just two incidents, both are linked to state-actors of DPRK.

Total losses by quarter (USD)




*In the Hacken Q1 2026 Security & Compliance report, Drift protocol was attributed to the Q1 losses because of its significance. In the Q2 Security & Compliance, Drift is excluded from Q1 and counted towards Q2 losses.

DeFi Losses

DeFi protocol losses of \$750,908,891 are the highest in this two-year window, up 152.2% year over year. The figure is dominated by the two April infrastructure attacks. The May and June base losses are ~\$131.6M, consistent with prior quarters.

DeFi protocol losses by quarter (USD)



Losses were heavily concentrated in April, which accounted for 82.5% of the quarter (\$619,324,750). Almost all of that came from the KelpDAO and Drift breaches, happening seventeen days apart.

Top 5 Incidents of Q2 2026 by Dollar Value

KelpDAO (rsETH) – \$292 Million (DPRK-Attributed)

Date	April 18, 2026
Attack vector	RPC poisoning against single-verifier bridge
Threat actor	DPRK-linked operators
Assets stolen	~116,500 unbacked rsETH released; ~\$292M extracted via Aave borrowing into ETH
Chain	Ethereum and connected chains
Recovery	~\$71M frozen downstream; a separate backstop later covered protocol bad debt

Drift Protocol – \$285 Million (DPRK-Attributed)

Date	April 1, 2026
Attack vector	~6-month social engineering; pre-signed durable-nonce approvals to a zero-timelock multisig
Threat actor	DPRK-linked operators
Assets stolen	~\$285M from protocol treasury and vaults
Chain	Solana
Detection delay	Drained quickly once enough signer devices were compromised

Humanity Protocol – \$31 Million

Date	June 9, 2026
Attack vector	Malware-infected device, private key compromise
Assets stolen	\$H drained from 17+ wallets; 100M additional \$H minted
Chain	Ethereum, BNB Chain
Detection delay	Token fell 80–90% as unbacked supply hit the market

MEV Bot (JaredFromSubway) – \$15 Million*

*Operator-stated
Independently verified figure ~\$7.5M

Date	June 20, 2026
Attack vector	Smart contract vulnerability, automated approval-generation abuse
Assets stolen	WETH, USDC, USDT swept in a single transaction (~\$15M, operator-stated)
Chain	Ethereum
Detection delay	Single-transaction drain; 50% bounty offered afterward

Grinex – \$13.7 Million

Date	April 15, 2026 (exact date undisclosed)
Attack vector	Compromised exchange infrastructure
Assets stolen	~\$13.7M from exchange systems
Chain	Multi-chain
Detection delay	Disclosed after drainage; attributed by the exchange to external actors

DPRK Attribution: Drift and KelpDAO

In recent years, nearly all crypto thefts attributed to North Korean actors have resulted from social engineering and operational compromise rather than smart contract exploitation, with the DangerousPassword, BlueNoroff, and TraderTraitor clusters taking roughly \$2.04 billion in 2025, about 52% of all losses that year. Q2 2026 followed the same pattern at a larger scale. In Q1, the main DPRK incident was the \$40M compromise of an executive's device at Step Finance. In Q2, the same approach was applied to protocol multisigs and cross-chain infrastructure.



Drift Protocol – \$285 million (April 1)

The Drift breach was a result of a six-month social-engineering campaign against the protocol's team, a longer-horizon version of the DangerousPassword approach seen in Q1.

Rather than targeting one executive, the operators spent roughly six months working the protocol's security council, ultimately obtaining pre-signed durable-nonce transaction approvals, standing authorizations that Solana lets a holder submit for execution later, without further signer involvement. With those approvals in hand, the attackers deployed a fake governance token, used it to inflate an oracle price, and executed the durable-nonce approvals to drain the protocol's vaults in roughly twelve minutes.

The failure was architectural as much as operational: the multisig carried no timelock, so once a durable-nonce approval existed, there was no enforced delay between execution and effect that the team could have used to catch and stop the drain.



KelpDAO – \$292 million (April 18)

In KelpDAO case, instead of compromising signers, operators exploited the LayerZero layer responsible for authorizing cross-chain token releases that Kelp's rsETH relied. By forging that release path, the attackers got the protocol to treat unbacked rsETH as legitimately bridged, minting and releasing roughly 116,500 rsETH with no collateral behind it. They then borrowed large amounts of WETH against that rsETH on Aave and sold into ETH, turning unbacked accounting entries into liquid value before the gap was caught, and leaving roughly \$177M of bad debt.

Kelp's cross-chain setup made token releases depend on a single validator, LayerZero's own infrastructure, so compromising that single verifier was enough to forge releases, with no independent second attestation required to flag a fraudulent message.

Together, these two operations come to \$577,000,000, 75.5% of the quarter losses. This confirms the Q1 finding, that North Korean operators continue to target people and infrastructure rather than focusing on smart contracts.

Smart Contract Vulnerabilities

Smart contract exploits were the most frequent category in Q2, accounting for 44 of 67 incidents, but they also caused a loss of \$87.7 million, which is about 11% of the quarter’s total. It’s almost the same as in Q1, and 67% less than Q2 2025’s \$263M. In Q2, the pattern was a high count of small exploits and very few large ones.

Vulnerability Type	Incidents	Loss (\$)		% of SC Loss
Bridge / Cross-chain Verification	7		\$32,780,000	37.4%
Logic / Validation Flaw	11		\$28,464,379	32.4%
MEV / Approval Abuse	2		\$15,200,000	17.3%
Arbitrary Call / In-Contract Access	3		\$3,344,000	3.8%
ZK Proof Verification	1		\$2,100,000	2.4%
Oracle / Price Manipulation	3		\$1,409,750	1.6%
Signature / Verification Bypass	2		\$1,385,519	1.6%
LP-Mint / Token-Standard Flaw	3		\$1,380,000	1.6%
Burn-Mechanism / Reserve Manipulation	6		\$972,293	1.1%
Inflation / Accounting Inflation	1		\$359,000	0.4%
Accounting / Decimal	3		\$265,750	0.3%
Reentrancy	2		\$70,000	0.1%

Bridge and cross-chain hacks in Q2 shared a pattern: In most cases a verifier failed to check if a transfer’s claimed value matched its real value. Therefore attackers initiated fake transfers stealing real assets. Taiko was the most recent case: ~\$1.7M was drained from the ERC vault on June 22, where fake cross-chain proofs were accepted on Ethereum even though nothing matching happened on the Taiko chain.

Business-logic and validation bugs caused \$28,464,379 in losses across eleven incidents in Q2. This was the quarter’s long tail. Hacken research team highlights three biggest smart contract vulnerability exploit patterns worth paying attention to:

- integer sign errors that let negative values be accepted where only positive ones should exist (Dango, Aftermath);
- circular swap paths that bypassed bad swap output checks (Rhea Finance, \$18.4M, the biggest pure smart contract loss),
- bad redemption calculation math in older vaults (Thetanuts).

Audit Findings vs Exploit Outcomes, Q1–Q2 2026

The Q1–Q2 2026 audit findings and the Q2 2026 exploit data describe two different stages: what auditors find during contract review, and what attackers turned into losses.



The dominant finding is inverted logic and calculation errors (22%). It results in code with the right structure but with an inverted conditional, a mismatched token index, or a reversed operand order that sends the value in the wrong direction. These findings pass the surface-level review because the code looks correct from the structure perspective.

Front-running/MEV is #2 at 18%. Order hijacking, recipient manipulation, and missing slippage checks are now standard risks which makes privacy-preserving protocols more exposed than transparent ones, since deterministic commitment construction makes transactions linkable.

Zero-knowledge circuits are a new category, accounting for ~12% of findings and invisible to standard Solidity review. Issues include underconstrained circuits that enable arbitrary minting, missing nullifier checks that enable replay, and manipulated merkle-depth inputs that skip verification.

Access control (10%) declined to #5, where the risks also changed. Access control vulnerabilities show up as role-renunciation bypass, pause-mechanism gaps, and sender-side (not recipient-side) validation, not the old missing-onlyOwner pattern.

Access Control and Operational Compromises

Access control and operational compromises were the defining category of Q2 2026: 21 incidents and \$674,866,100, or 88.3% of all losses. This is the reverse of Q1, where phishing led, though the actor behind both quarters is the same. The two largest incidents, Drift and KelpDAO, both sit in this layer, and both are attributed to North Korean operators who compromised signing and infrastructure.

Project	Loss	Root Cause	Chain
 KelpDAO	\$292,000,000	RPC poisoning against single-verifier bridge	Ethereum
 Drift Protocol	\$285,000,000	Multisig governance takeover via compromising teammates	Solana
 Humanity Protocol	\$31,000,000	Private key compromise	Multi-chain
 Grinex	\$13,000,000	Compromised exchange infrastructure	Multi-chain
 THORChain	\$10,700,000	Progressive threshold-signature key-material leak	Multi-chain
 DxSale	\$7,300,000	Private Key compromise	BNB Chain
 TrustedVolumes	\$6,700,000	Missing allowlist access control (self-added order signer)	Ethereum
 SUPERFORTUNE (GUA)	\$5,660,000	Multisig destination-address tampering	Ethereum
 Gravity Bridge	\$5,400,000	Compromised signing	Multi-chain
 Wasabi Protocol	\$5,000,000	Compromised deployer/admin key -> malicious proxy upgrade	Multi-chain
 Volo Protocol	\$3,500,000	Private key / operator key compromise	Sui
 StabIR (USDR/EURR)	\$2,800,000	Multisig key compromise (1-of-3 mint key)	Ethereum

The incidents fall into three patterns: compromised multisig signers and private keys (Drift, Humanity Protocol, TESSERA, SUPERFORTUNE, StabIR, Volo); compromised bridge and cross-chain infrastructure (KelpDAO, THORChain, Gravity, Alephium); and operational backend compromises (Grinex exchange systems, DxSale liquidity-locker ownership, OpenMonero server).

Deprecated Infrastructure Exploits

Aztec – \$4M

Thetanuts – \$2.1M

Raydium – \$1.34M

When a protocol migrates to a new version, the previous contracts are typically removed from the interface and documentation but remain deployed, callable via direct transactions, and, in some cases, still holding user funds. Deprecation therefore, changes a contract's security coverage rather than its exposure: subsequent audits are scoped to the current version, monitoring is configured for the active deployment, and bug bounty programs cover the documented surface, while the deployed legacy code and its balance remain reachable.

Q2 provided examples of the two main failure patterns. In Raydium's case, deposits into legacy contracts were disabled at the interface level but remained callable on-chain, so the restriction applied to product users rather than direct callers. In Thetanuts' case, a flaw in the redemption calculation was exploited in an older vault whose logic had not been reviewed since its original audit. In both patterns, the exploited code was outside the scope of every current security control while remaining functionally active.

AI on Attack and Defense

\$174,000

first case of agent maliciously moving funds on-chain

4 y.o.

Zcash Orchard bug identified with LLM

0

forensically confirmed AI-enabled smart contract/protocol exploits in Q2

As LLM capabilities advance, the industry widely assumes that adversaries are weaponizing them. A stronger version of that claim circulated in Q2: that the release of flagship models directly enabled more exploits and explained the quarter’s incident frequency. Our incident data does not support it. No Q2 exploit has been forensically attributed to AI tooling, and the correlation between model releases and incident counts remains speculation.

What the evidence does support is narrower. Stephen Ajayi, security researcher at Hacken, identifies three uses of LLMs already established in offensive practice:

- phishing, impersonation, and social fraud at scale;
- autonomous search for documented vulnerabilities across deployed code;
- and acceleration of exploit development through model-written scripts.

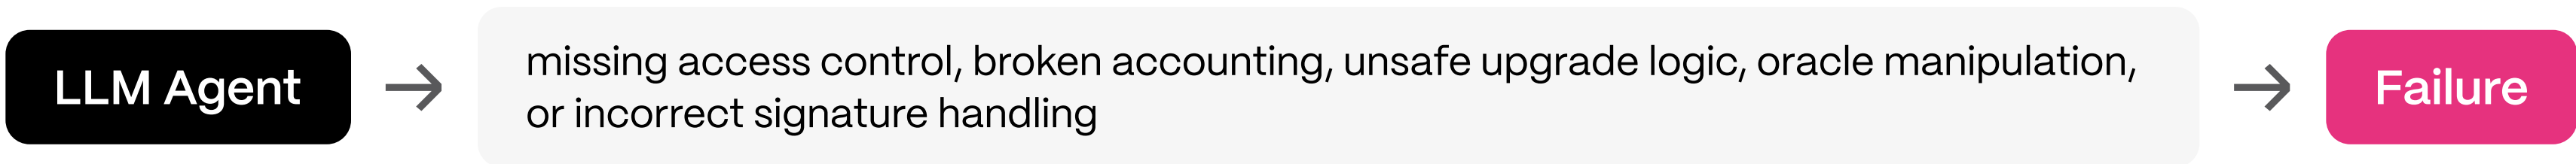
All three compress attacker timelines; none has yet been shown to produce an attack class that did not exist before.

How prevalent is unreviewed AI-generated code in deployed protocols?

AI-generated code is likely already common in deployed protocols, but no reliable industry-wide figure exists, and none can currently be produced. Development tools rarely preserve provenance, code is edited and squashed before release, and audit reports classify vulnerabilities by type, not by whether a model originally wrote the affected function. In Web3, attribution is further degraded by copied contracts, forks, and reused libraries. Any exact percentage offered today is an estimate.

In practice, AI authorship becomes visible only after an incident, when post-mortem analysis finds an entry point trivial enough that no experienced reviewer would have let it be deployed.

Is AI the root cause of exploits, or a contributing factor?



Primary exploit reasons: inadequate review, missing invariants, or weak testing

A contributing factor, in almost every case that can be assessed. The technical root cause remains something concrete: missing access control, broken accounting, unsafe upgrade logic, oracle manipulation, or incorrect signature handling. AI may have introduced the defect, accelerated development, or created false confidence in unreviewed output, but the underlying failure is inadequate review, missing invariants, or weak testing. The point most analyses miss is that even when a model wrote the vulnerable line, the organizational root cause remains unchanged: security-critical code was deployed without prior independent review.

Agents as the Attack Surface

SVRN

SVRN perspective



Q2 2026 produced two documented attack patterns against agentic systems that fall outside conventional crypto security. The first is the malicious intermediary. An April 2026 study of LLM routers, the third-party services that sit between an agent and the model it calls, showed these routers injecting malicious instructions into agent tool calls, and in at least one case draining a wallet, because the intermediary has full plaintext access to everything passing through it and a single rewritten tool call is enough to redirect execution.

The second is the unverified handoff between systems, where one agent's output is consumed by another as an authorized command. In the Bankr/Grok drain, a prompt injection caused a conversational model to produce a transfer instruction that a downstream execution layer ran automatically, moving funds with no validation of where the instruction came from or whether it was legitimate.

David Schwed, COO, SVRN

The Bankr case breakdown: Grok's agent wallet received a Bankr Club membership NFT that unlocked transaction permissions, and an encoded prompt injection then had the agent transfer 3 billion DRB tokens with no check on the instruction's origin. Which resulted in ~\$174,000 losses.

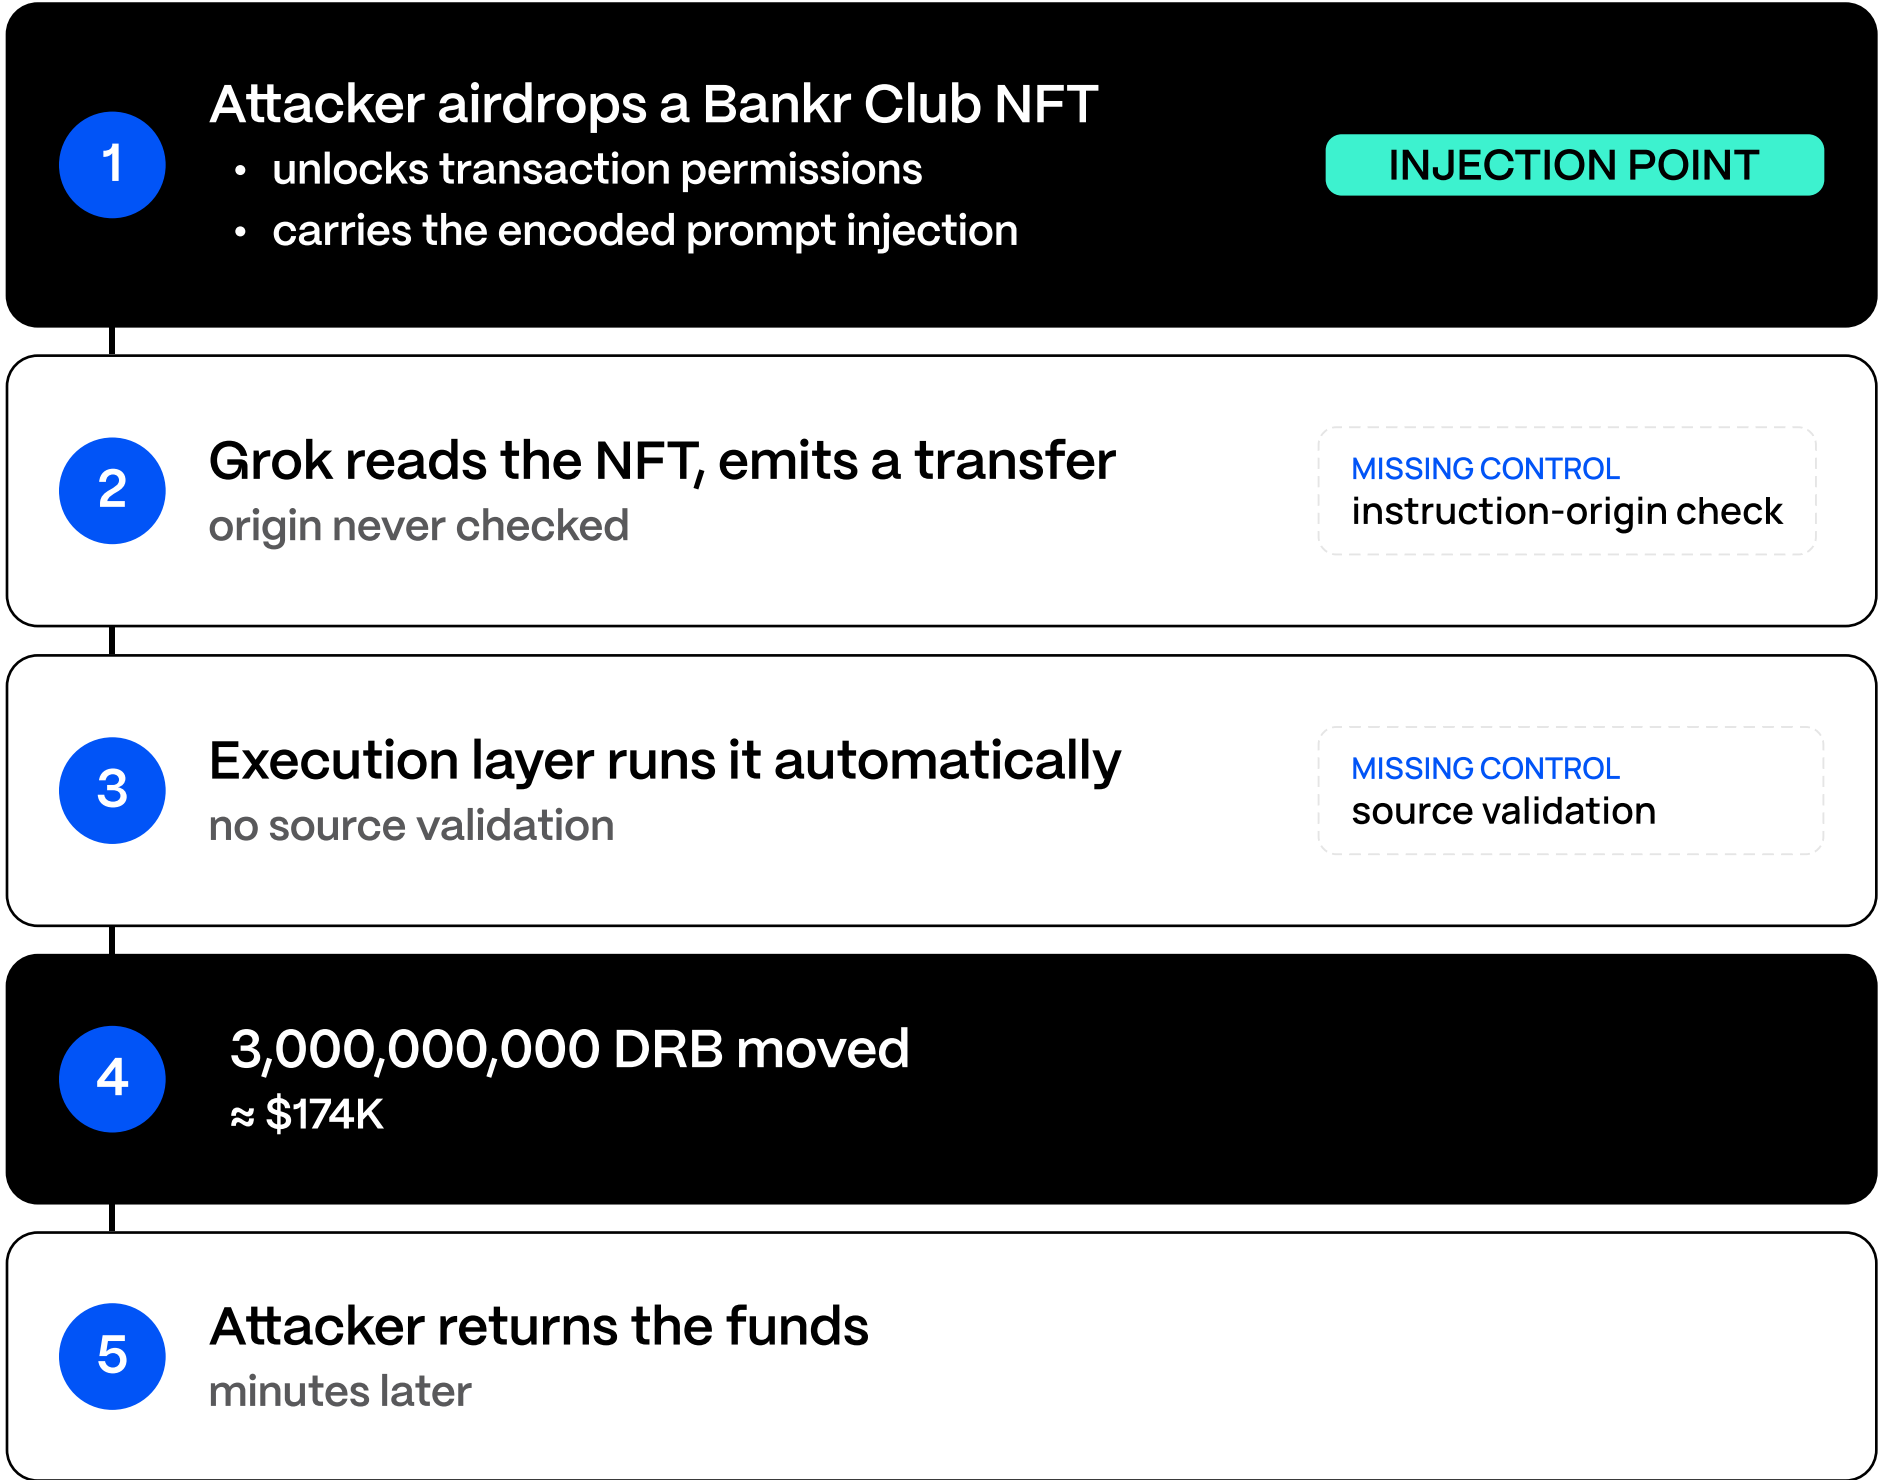
Bankr / Grok Agent-wallet Drain

first prompt-injection of a live agent wallet · May 2026

■ attacker-controlled / missing control □ system step

This creates a challenge for teams that heavily rely on AI, as poisoned artifacts appear more and more: tools, skills, packages, prompts – malicious intent can be injected from anywhere, and right now there is no industry-wide answer on how to protect AI and keep it privileged on managing funds at the same time.

~\$174,000 in tokens moved, returned in full.
Source: MetaMask, May 2026.



What does a properly threat-modelled AI pipeline look like, and who is doing it well?

A properly threat-modelled pipeline starts from the assumption that the model will at some point be wrong, manipulated, or compromised, and constrains what that failure can reach.

Untrusted input passes through controlled ingestion and provenance-aware retrieval before it touches the model. The model may recommend an action, but authorization is handled by a separate policy layer: a deterministic check that determines whether the action is permitted, independent of how persuasive the model's output is. Tool access runs through scoped identities, sandboxes, transaction limits, and explicit approval gates, with full logging across prompts, retrieval, memory, tool calls, and the resulting system changes.

The most mature public frameworks come from Microsoft, Google, NIST, MITRE, OWASP, and OpenAI, but none is sufficient on its own. The operative standard is not which framework an organisation cites; it is whether the controls are actually applied: deterministic authorization, least privilege, least agency, memory integrity, secure tool execution, and rehearsed incident recovery.

A new prompt, model version, tool description, MCP server, or RAG source is a production security change and should pass through the same review gate as a code deployment, even when no conventional application code has changed.

Being Exploit-Resilient: Key Takeaways and Recommendations

Web3 security has limitations. For Q2 it was people layer

01 · The main risk remains operational, not on-chain code.

88.3% of the quarter's losses came from compromised keys, signers, and infrastructure, not from flawed smart contracts/programs. Smart contract bugs were common but accounted for about 11% of the dollar losses. The two most expensive failures were a multisig-signer compromise and a single-validator bridge, neither of which a conventional audit would have caught. Security spending across the industry is still weighted toward code review, while losses trend weigh more and more toward other domains like operations and dependencies.

02 · Cross-chain release logic is the new systemic weak point.

Bridge and cross-chain verification flaws were both the largest category within smart contract losses and the mechanism behind the largest single incident. The failure is usually a release path that trusts one verifier or does not bind a transfer's claimed value to its verifiable source.

03 · The AI layer entered the loss record; the instruction is now an attack surface.

Q2 produced the first documented case of an agent moving funds on-chain under malicious instruction, and both documented attack patterns against agentic systems, the malicious intermediary, and the unverified handoff targeted the instruction rather than the key.

Recommended actions for protocol teams:

- Require dedicated, hardware-isolated signing devices for all multisig signers. Do not approve transactions from a daily-driver laptop used for email, messaging, or browsing.
- Enforce a timelock on treasury and admin multisigs, a required delay between approval and execution that creates a window to catch and stop malicious transactions.
- Require independent multi-validator attestation for all cross-chain token releases. Do not make a release depend on one validator or messaging provider.
- Run real-time monitoring of supply-to-collateral ratios with automated circuit breakers on cross-chain mint and release volume.
- Move minting and proxy-admin roles to MPC or threshold signatures behind a timelock, and remove single-key and single-EOA control of privileged functions.
- Treat any agent with transaction authority as a privileged signer. Apply the same controls: scoped identity, spend and rate limits, and explicit approval gates on irreversible actions.
- Separate recommendation from authorization. The model may propose an action; a deterministic policy layer, not the model, decides whether it is permitted.

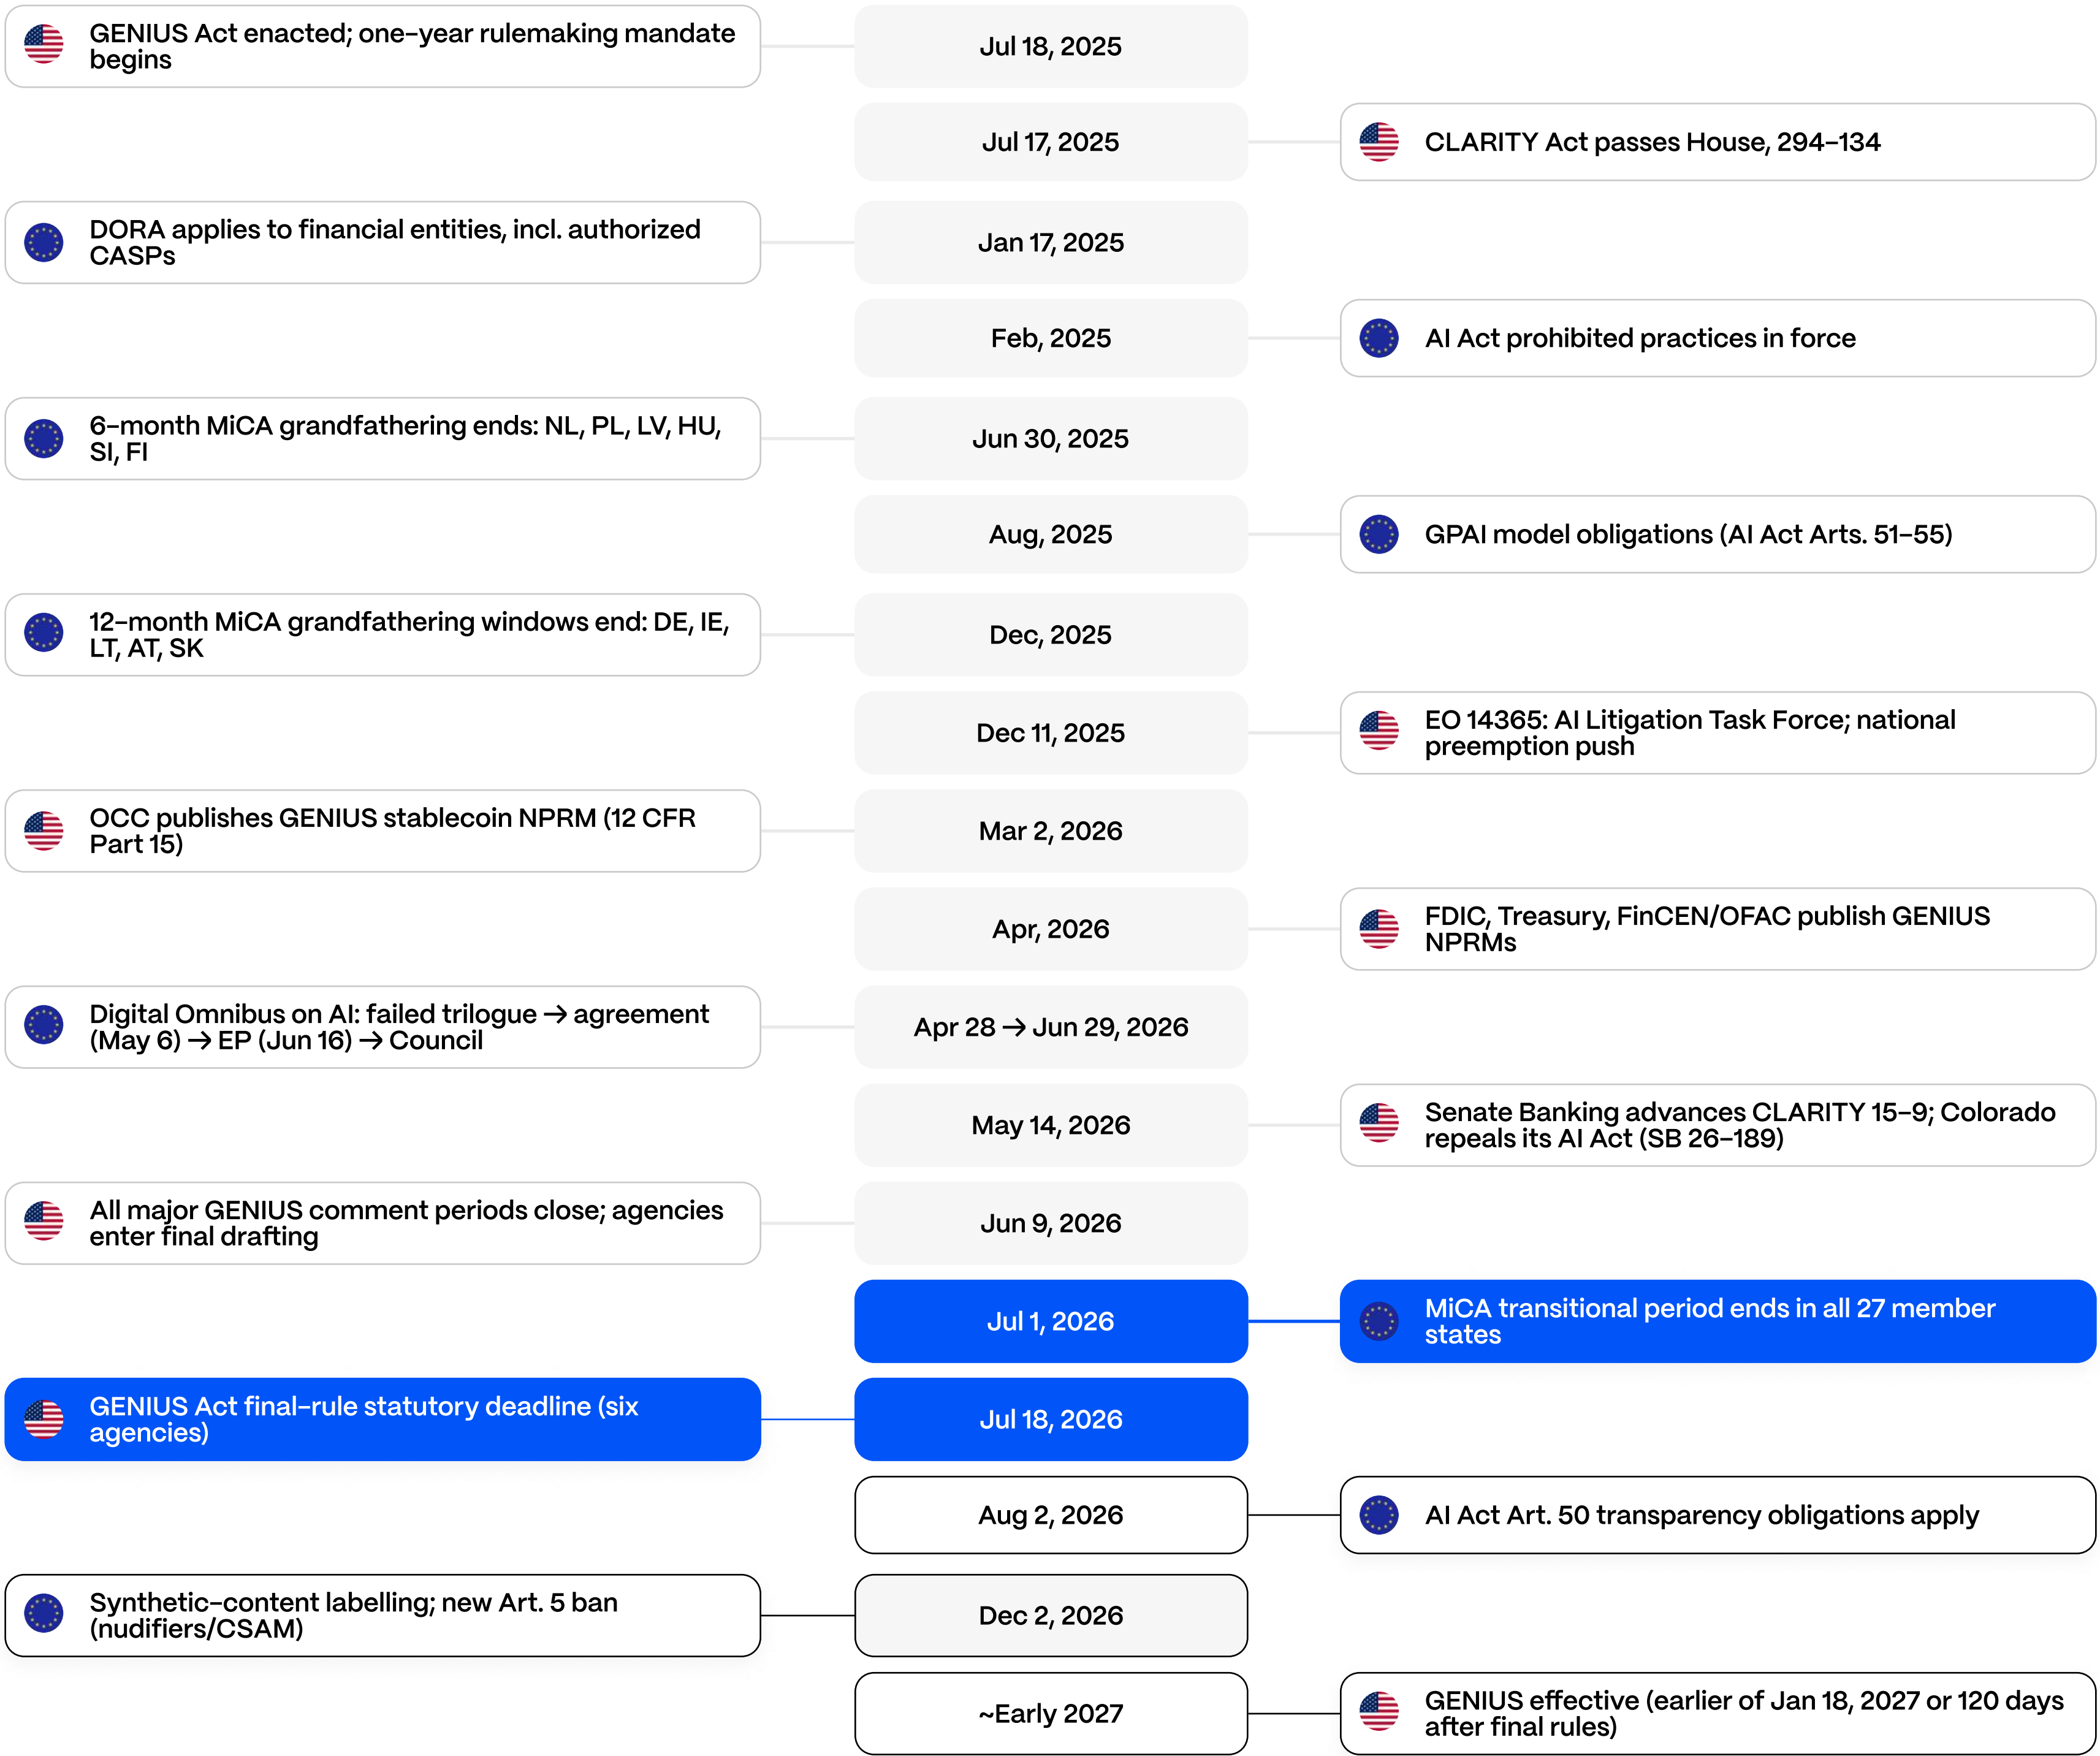
Q2 2026:

Enforcement Begins Where Practice Lags

Two regulatory regimes reached decisive points in the same quarter through different mechanisms. In the EU, MiCA's transitional period ended on July 1, converting the bloc's crypto market from a set of national registrations into a licensed financial sector with fewer than one in five firms authorized. In the US, the equivalent event was a rulemaking deadline rather than an expiry: the GENIUS Act's one-year mandate expires on July 18, requiring six federal agencies to finalize the first federal stablecoin rulebook. The EU AI Act was formally amended in the same period, after the standards and infrastructure needed to implement it were not ready on schedule.

Both events reflect the same condition: the governance layer of digital-asset regulation is maturing faster than the operational security practice it is intended to ensure. The EU has finalized its framework and brought it into force; the US is still completing its rulemaking. Neither has closed the gap between what a license certifies and what caused the losses recorded in Q2. This chapter examines that gap.

Regulatory Timeline



EU vs US Crypto Regulation in Q2 2026

Dimension	European Union	United States
Stablecoins	MiCA, enforced 1 Jul 2026	GENIUS Act, final rules due 18 Jul 2026, effective ~early 2027
Status	Binding and live	Rulemaking sprint, not yet in force
Market structure	MiCA + DORA, whole market	CLARITY Act stalled in Senate; stablecoins only until then
Reclassification	CASP becomes a DORA financial entity	PPSI becomes a BSA financial institution
Operational-security regime	DORA: single cross-cutting rulebook, five testable pillars	No federal equivalent; GLBA, FFIEC, SEC, state (e.g. NYDFS 500) patchwork
Yield on stablecoins	Prohibited outright (EMTs and ARTs)	Issuer yield barred; exchange rewards contested
AI direction	Deferring high-risk obligations (Digital Omnibus)	Rolling back state governance via litigation and EO; NIST builds security underneath

The EU has a single rulebook for the whole market, and it is in force. The US has a live rulebook for stablecoins only, a stalled bill for other assets, and no cross-cutting operational-resilience regime in either instrument. The difference that matters is structural, not timing: EU authorization attaches DORA’s operational-security regime and US authorization attaches none, so the controls that failed in Q2 fall inside the EU rulebook and outside the US one.

MiCA: 100% Enforced, 20% Compliant

On April 17, ESMA confirmed that MiCA’s transitional period ends on July 1, 2026, in all 27 member states, with no extension mechanism. From that date the Article 143(3) grandfathering basis no longer applies: any firm still serving EU clients on a legacy national registration is in breach of EU law.



By May 2026, roughly 210–215 CASPs held full authorization out of more than 1,200 pre-MiCA registrations, concentrated in Malta, the Netherlands, Cyprus, France, and Ireland. More than 80% of previously registered firms did not obtain a license.

Authorization is Reclassification

MiCA converts crypto businesses into regulated financial entities. The authorization file requires governance, a cybersecurity policy, own-funds evidence, and incident-reporting duties; the Travel Rule applies to every qualifying transfer regardless of amount, and outsourcing the messaging layer does not transfer the underlying liability.

An authorized CASP is a financial entity under DORA and inherits its five pillars. This is the substantive security bar: documented ICT governance and evidence an auditor can verify, in place of reliance on smart-contract audits alone. Firms that treated MiCA as an administrative filing and deferred the operational-security requirements are the most exposed now that the transition has ended.

Reclassification Differs by Jurisdiction

Both regimes regulate crypto firms, but only one has an operational-security standard. A CASP under DORA inherits a resilience regime with five pillars, while a GENIUS issuer follows the Bank Secrecy Act with AML obligations but lacks operational-security. No US equivalent to DORA exists.

Cybersecurity expectations for financial institutions are spread across various regulations, none addressing on-chain signer or key risks. The EU standard is defined but not verified in supervision; the US operational standard is still undefined.

Do regulations set the security bar effectively?

On coverage, yes. Every major Q2 failure falls within the first pillar of DORA requirements (ICT risk and access control). Chapter I’s loss table and DORA’s obligation list correspond item for item.

Exploits	DORA requirement (Pillar 1: ICT risk management)
Signer devices & keys (88.3% of losses)	Strong authentication and protection measures of cryptographic keys — Art. 9(4)(d)
Single-verifier bridges (KelpDAO \$292M)	Identify potential material single points of failure — Art. 10(1)
Admin paths & backends (Grinex, DxSale, Wasabi)	Least-privilege access control — Art. 9(4)(c)
Unmonitored mint functions (Humanity, StabIR)	Prompt detection of anomalous activities; multi-layer control and alert thresholds — Art. 10(1)–(2)

On verification, no. Through H1 2026, supervision remained focused on the legal mechanics of authorization: governance, AML/CFT, disclosures, and own funds. Cybersecurity is reviewed at the margin of the file, typically a security policy and a point-in-time penetration test.

While DORA itself is adequate, supervisory practice has not yet implemented it to that depth.



Institutional allocators are raising the bar significantly before committing capital, and many digital asset firms are still underestimating how rigorous that process has become. At a minimum, allocators want audited financial statements prepared under GAAP. SOC 1 & 2 reports are increasingly a prerequisite. Expect detailed questions around regulatory registration status, FinCEN compliance policies, FATF adherence, and any recent cybersecurity incidents. If you do not have prepared, documented answers, that gap will show.


What surprises many firms is how much a clean audit actually does for them beyond the numbers: a rigorous audit addresses a broad range of risk and controls considerations in a single package, and sophisticated counterparties know how to read one.

The preparation gap most often appears as firms transition from startup operations to full GAAP compliance. Firms that treat audit readiness as solely a finance project will struggle: it is an operational and compliance project just as much. At this point, that preparation is not a differentiator. It is the cost of entry.

–John Delalio, Partner, EisnerAmper

Stablecoins: MiCA’s Deadline vs GENIUS’s Schedule

MiCA-compliant:



GENIUS-aligned:



In the EU, grandfathering under Article 143(3) has ended. USDT was withdrawn from major EU venues after Tether declined EMT authorization. USDC and EURC are the authorized names; the largest circulating stablecoin is not among them. By market cap in mid-2026, USDT held roughly two-thirds of global stablecoin supply and USDC just over a quarter, so MiCA authorization covers a minority of the supply EU users actually transact in.

In the US, GENIUS was enacted on July 18, 2025 and set the federal banking regulators a one-year deadline. Six agencies (the OCC, FDIC, NOUA, Treasury, FinCEN, and OFAC) must issue final rules by July 18, 2026. The OCC proposed its framework (12 CFR Part 15) on March 2; the FDIC, Treasury, and a joint FinCEN/OFAC AML rule followed in April. All major comment periods closed by June 9, leaving the agencies roughly five weeks to reconcile six proposals into final rules. The regime takes effect on the earlier of January 18, 2027 or 120 days after final rules publish, so the binding date is in early 2027, not the July deadline. The proposals define the shape already: an OCC \$5 million capital floor, defined redemption standards, no FDIC deposit insurance for token holders, and a no-yield prohibition.

The structure mirrors MiCA. A GENIUS-permitted issuer is reclassified as a regulated financial entity, not merely licensed. As under MiCA, the application addresses governance, reserves, capital, and AML in detail but specifies the least on operational security, the controls that failed in Q2.

Yield is the principal divergence. MiCA prohibits interest on both e-money tokens and asset-referenced tokens, so the question is closed in the EU. In the US it is unresolved: GENIUS bars issuers from paying yield but does not bar exchanges from offering interest-like rewards, and the SEC and CFTC treat protocol staking as administrative while placing guaranteed or above-protocol yield within the securities perimeter. The same instrument is settled under MiCA and unresolved under US rules.

[See workbook: Stablecoins EU vs US](#)



As stablecoins evolve from settlement assets into critical financial infrastructure, cyber risk is increasingly shifting beyond reserve transparency and regulatory compliance toward operational resilience. Where do you believe the cybersecurity posture of most stablecoin issuers stands today, and which attack surface concerns you most heading into Q3?

Cybersecurity posture has always stood on its own. Reserve transparency belongs to treasury, regulatory compliance to legal and compliance, and operational resilience to security. They're complementary, not a spectrum where risk shifts from one to the next. What's changed is visibility: as stablecoins become critical infrastructure, the operational security layer required of digital assets comes under real scrutiny.

A stablecoin issuer inherits the full weight of general cryptocurrency controls on top of its issuer-specific obligations. Key management, signing ceremonies, on-chain transaction monitoring, wallet and custody hardening, and protocol-level security are the baseline every crypto operator carries, and the issuer layers mint, burn, freeze, and reserve integrity on top of that. It's the entire crypto control surface, plus the heightened stakes of controlling a liquid, on-chain instrument.

Heading into Q3, the attack surface I'm watching most is the privileged operational layer around mint, burn, freeze, and administrative authority. The smart contracts are audited heavily, but the identity, signing, and tooling that control them are another problem. Social engineering, internal threats, and endpoint compromise against the people operating the levers of the system are the threats I'd prioritize."

— Brett Enclave, Chief Information Security Officer, 1Money

Stablecoins: The New Regulated Perimeter


GENIUS and MiCA established clear legal footing for stablecoins and their settlement infrastructure, together with continuous-monitoring expectations that current data infrastructure does not fully meet. Market figures in this section come from Allium’s State of Onchain Finance Q2 2026.



(Source: Allium, State of Onchain Finance Q2 2026.)


The quarter also produced an issuer-level example. StabIR (USDR/EURR) lost \$2.8M when one of three mint keys was compromised. The failure is in the privileged operational layer that issuer regulation assumes is hardened: mint, burn, freeze, and admin authority. This is the layer the GENIUS proposals address least. The perimeter is legally defined; the underlying controls are not inspected under any current license.

AUM by chain (where the money is)




Stablecoins follow the same pattern: 53% of geolocatable payment volume lands in Asia-Pacific, outside the reach of US-centric data pipelines. Ethereum holds half of all AUM.

Emerging markets lead; the US is 8%



Top countries by share of geo-attributed filtered volume received. Same window.
Data from: Allium Q2 State of Onchain Finance report (as for April-June 2026)

Most volume lands outside the Americas



Share of geo-attributed filtered volume received, by region.
January to June 2026.

Covering this chain set continuously requires three capabilities most providers weren't built to deliver at once: 1) real-time streaming across networks rather than batch queries on a delay, 2) normalized cross-chain schemas so a USDC transfer on Stellar and one on Arbitrum look identical to a compliance system, and 3) indexing depth that meets audit precision across chains at different maturity levels.

The industry still cannot fully solve cross-chain attribution. Funds that move between networks via bridges lose their traceable thread that regulators need for end-to-end AML monitoring. On newer chains where regulated assets are actively settling, indexing depth and data quality do not yet match what compliance workflows require.

The perimeter of regulation keeps expanding as well. Agentic payments (software systems settling micropayments in stablecoins with no human in the loop) added 24 million transactions in the last 90 days on a rail with no established compliance playbook yet. Regulators who expect continuous surveillance are moving faster than the data infrastructure built to support them.

Source: Allium onchain data, State of Onchain Finance Q2 2026. By Elton Shehdula, Research at Allium.

Market Structure: Binding Stack vs Stalled Bill

Stablecoins are the only asset class both jurisdictions have fully regulated. For other assets, the US has no federal framework. The EU regulates the whole market through MiCA and DORA, both in force. The US has federal rules for stablecoins only; for the spot market in all other crypto assets there is no federal rulebook, only case-by-case agency enforcement.

The bill intended to address this, the CLARITY Act, has not passed. It cleared the House in July 2025 (294–134) and the Senate Banking Committee 15–9 on May 14, 2026, then moved to the Senate Legislative Calendar (No. 423) on June 1, making it eligible for a floor vote. It has not received one. The White House’s informal July 4 target passed without enactment; with the Senate back from recess on July 13, the pre–August window (an August 7 procedural marker) is the last realistic window for 2026 passage. Passage requires 60–vote cloture, meaning all 53 Republicans plus roughly seven Democrats, and prediction–market odds for 2026 enactment have fallen from about 82% in February to 40–50%.

Three disputes are blocking the crossover votes: an ethics provision on officials’ crypto holdings, intensified by the July 1 disclosure of roughly \$1.4 billion in presidential crypto–related income; the Section 604 legal shield for DeFi developers, which law–enforcement groups say would impair investigations; and the same stablecoin–yield exemption that banking industry groups want closed. Until CLARITY passes, EU firms know the rules that apply to them; US firms outside stablecoins do not.

AI Compliance: Deferral in the EU, Rollback in the US

The governance–versus–security gap applies to AI as well. On the governance side, the two jurisdictions moved in opposite directions this quarter.

The EU delayed its rules. The Digital Omnibus deferred the high–risk obligations to 2027 and 2028 because the standards and tooling needed to comply were not ready in time. The core prohibitions, general–purpose AI model duties, and Article 50 transparency obligations remained in force.

Dec 11, 2025Federal Executive

EO 14365: Litigation Task Force

AG to create AI Task Force to address state law conflicts.

March 2026Federal Executive

Preemption Framework

Congress urged to prevent burdensome state AI laws.

Spring 2026Congressional

Congressional Inaction

Congress ignored AI preemption amid executive pressure.

April 2026Legal/Court

xAI vs. Colorado

xAI sued to block CO Act; DOJ intervened in AI law challenge.

April 27, 2026Legal/Court

Enforcement Stayed

Federal Magistrate paused enforcement of Colorado's AI Act.

May 14, 2026State Level

CO Repeal & SB 26–189

Colorado's AI Act is a narrower disclosure law starting Jan 2027.

June 2, 2026Federal Executive

EO 14409: Benchmarking

Created benchmarks for frontier AI.

Current Outcome

Federal intervention has pushed states to limit AI regulations.

Governance is not security. The work on how AI systems are attacked comes from outside the certifiable standards. NIST launched its Agent Standards Initiative in February and published a concept paper on agent identity and authorization; ISO released ISO/IEC 27090 as guidance, written in "should" rather than "shall", so it carries no binding requirements and no certification path.

An organization can therefore hold an AI governance certificate while the one standard addressing how its AI is attacked imposes no obligations. NIST red–team research measures the gap: novel agent–specific techniques hijacked tasks in about 81% of attempts, against roughly 11% for the strongest known baseline attacks, about seven times higher. The attack surface is developing faster than the standards written to certify it.

Governance vs Security Coverage Matrix

An ISO 42001 certificate attests that an organization governs its AI, not that its AI is secure. Certification is a governance artifact, and the relevant attacks are operational. Security should be a required, gated stage in AI adoption: red teaming and agent penetration testing in the lifecycle, the MCP and tool-chain in scope with least privilege and runtime monitoring, and secrets, keys, and cloud infrastructure hardened during implementation rather than afterward.

Regulations	AI red teaming as gated stage	Implementation security (keys / secrets / cloud)	MCP & tool-chain security	Data-leakage testing	Agentic runtime monitoring
ISO/IEC 42001	☐	☒	☐	☒	☒
EU AI Act	☒	☒	☐	☒	☐
NIST AI RMF	☒	☐	☐	☐	☐
ISO/IEC 27090 (FDIS), guidance, not certifiable	☒	☒	☐	☒	☐

☒ clearly required ☒ mentioned, not operationalized ☐ absent

Takeaways and Recommendations

01 · The perimeter closed faster than the industry converted.

Under 20% of pre-MiCA registrations became authorized CASPs by July 1. Authorization is reclassification: an authorized CASP is a financial entity under DORA, and the resulting obligations are operational, not administrative. The US reaches the same reclassification through GENIUS, but without a DORA-equivalent operational-security rulebook attached.

02 · The rules cover Q2's failure modes; supervision does not yet verify them.

Every major loss category of the quarter maps to a DORA pillar. Until licensing reviews test controls instead of reading policies, allocators, insurers, and security firms are the de facto examiners in both jurisdictions.

03 · AI obligations were deferred; AI exposure was not.

EU high-risk duties moved to 2027–2028 and US federal pressure narrowed the leading state framework, but supervisors are embedding AI risk into existing ICT obligations now, and no certifiable standard yet covers AI security.

For CASPs:

- Implement DORA as engineering, not documentation: hardware-isolated signing devices, timelocks on treasury and admin functions, admin-path controls tested under realistic conditions.
- Extend the Register of Information to the dependencies that failed in Q2: bridge validators, oracle providers, messaging layers, RPC infrastructure.
- Run resilience testing against operational scenarios, signer compromise and forged cross-chain release, not only against contract code.

For stablecoin issuers (EU and US):

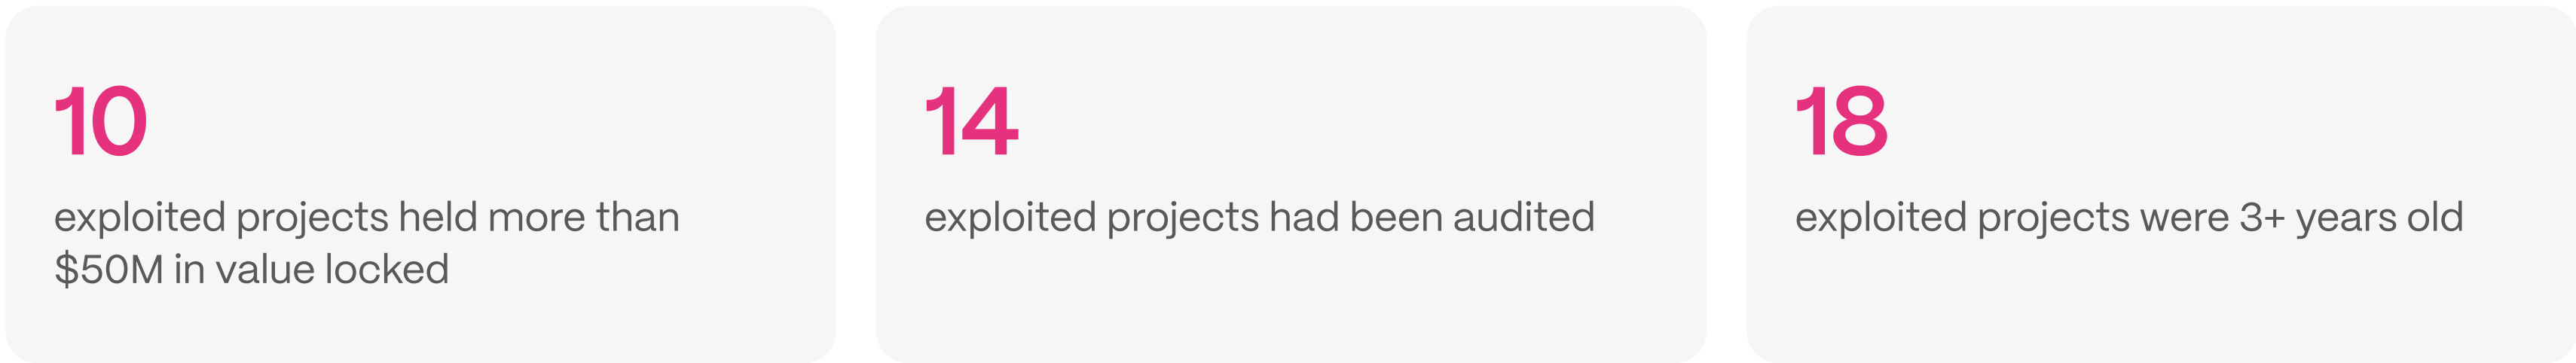
- Treat mint, burn, freeze, and admin authority as regulated critical functions: MPC or threshold signing behind a timelock, no single-key control. StabIR is the reference case.
- Monitor supply-to-reserve ratios in real time, with automatic circuit breakers on mint volume.
- Match monitoring coverage to the asset's actual chain footprint; an Ethereum-only view covers about half the regulated perimeter.
- US issuers preparing for GENIUS should implement the operational-security controls the application does not require, because reclassification as a BSA financial institution does not import a DORA-style resilience standard.

Architecture of Trust

The old trust signals failed in Q2

Q2 tested the indicators the market has long treated as proof of the established trust standard: a prior audit, years of operating history, value locked, and brand recognition, and none of them has proven effective in indicating risk management quality. Projects with as much as nine years of history were exploited through attack paths that had already been documented. For this chapter, our goal is to discover and show the market perspective on security indicators that matter in today's threat landscape

Trust signals that failed



Moody's Ratings

MOODY'S

Industry perspective



From a credit perspective in Q2 2026, the binding constraint remained cyber risk: not limited to small or retail-focused crypto platforms, but a growing threat to larger, institutional ones. Cyber risk stands apart from other categories of risk across the blockchain technology stack because it stems from failures in operational security, access controls, third-party dependencies, and key management, making it an effective constraint on institutional trust.

Beyond immediate cyber exposures, Q2 marked a clear shift toward incorporating long-horizon cryptographic risks into credit thinking. Quantum computing threatens digital assets by potentially enabling the direct derivation of private keys, allowing undetectable and irreversible theft of assets, which reframes cyber risk from an operational issue into a foundational risk to asset integrity and ownership. This risk is not being treated as purely theoretical: large financial institutions are already testing post-quantum cryptography and building crypto-agile architectures. Security measures that were previously sufficient, such as current-state encryption or static key management, are no longer adequate without demonstrable plans for cryptographic transition, agility and long-term resilience.

In Q2, operational resilience emerged as the practical lens through which security, compliance, and governance were evaluated. Decision-grade operational resilience required robust operational security and access controls; effective third-party and supply-chain risk management; the ability to operate across irreversible environments; and continuous preparedness rather than point-in-time compliance. Critically, this marks a shift from static assurance (audits and compliance checklists) toward continuous, demonstrable resilience under stress scenarios.

Governance and compliance became necessary but not sufficient. Only when translated into continuous cyber resilience and cryptographic readiness did they meet the threshold for institutional trust.

Rajeev Bamra, Executive Director & Head of Digital Economy Strategy at Moody's Ratings

Audited and Exploited: Why the Signals Failed

Fourteen audited projects were exploited in Q2. In every case the loss fell outside a smart contract audit scope: signer devices, bridge validators, backend infrastructure, deprecated contracts, and admin keys. An audit is a point-in-time review of specific code, and it was never designed to attest to how a protocol is operated.

Smart Contract Audit Scope	Q2 Exploits
Whether the contract logic and math are correct: conditionals, token indices, operand order	Durable-nonce transaction approvals to a zero-timelock multisig, exploited to execute without further signer involvement
Whether on-chain roles are set correctly: owner, pause, and upgrade permissions	Admin keys and backend systems behind the application (Grinex, DxSale, Wasabi)
Whether the deployed code matches the reviewed code	The bridge and cross-chain validators that authorize token releases (KelpDAO, \$292M)
Whether known bug classes are absent: reentrancy, oracle misuse, bad token accounting	Mint and freeze authority left unmonitored after launch (Humanity Protocol, StabIR)
A snapshot of specific contracts at a single point in time	Deprecated and older contracts still live but outside the current audit scope

The conclusion is not that audits failed. Chapter I’s findings data shows pre-deployment review catching more vulnerabilities earlier. The problem is that a single instrument is no longer capable of carrying all the weight of trust, while other security layers have not yet been established as a minimal entry barrier to liquidity acquisition.

State of Protocol-level Risk Management in Web3

While the smart contract audits were considered a bare minimum for market entry for over a decade, the market still has challenges with acquiring them, establishing proper coverage, and disclosing them publicly.


For this chapter, we used data from Hacken’s CORE3, a global risk infrastructure layer for Web3. CORE3 collects data from open sources, direct data submissions by projects, and registries to aggregate the Probability of Loss (PoL) is an unbiased, data-based numerical index (0–100) that estimates the likelihood that a project will fail or that users will incur losses.

The data is formed by identifying the top 1,500 projects by market cap across the top 50 centralized exchanges (by CoinGecko Trust Score), then filtering out stablecoins, tokenized RWAs, and wrapped tokens. The resulting 1,427 projects, all with a market cap above \$1 million, account for 76.3% of the 1,868 eligible projects.

Two limits are worth stating. It is a point-in-time vendor snapshot, so the counts shift as projects are added and controls change, and it captures what is observable or disclosed, which can understate private controls at some firms.

Security Controls in Place


Coverage is common at the entry level and rare once it demands ongoing spend. Audits and bounties are relatively widespread across the data; continuous monitoring/prevention is not.



The rarest control is the one Q2 needed most. 91% of indexed projects run no independent monitoring layer, so the outflows that monitoring exists to catch went unwatched at nine projects in ten, which is the same gap the quarter’s operations-led losses moved through.

License and Registration

Most of the market has a registered company, while at the same time very little of it operates under active financial regulation, and protocol-level MiCA authorization is rarer still.



Non-EU registrations are led by the US (219), Singapore (130), Cayman (115), Switzerland (101), and BVI (70). Switzerland and post-Brexit UK are counted as non-EU.

Registration outruns regulation by roughly 10 to 1, so most of the registered market carries no external check on the operational layer that failed in Q2 through any of the controls regulators require to be tested.



Chainlink on verification becoming a hard institutional requirement in Q2

Industry perspective



Throughout Q2 2026, proof-of-reserve and cross-chain verification transitioned from risk-mitigation tools into foundational requirements for institutional participation. The catalyst was the convergence of three market forces:

Tokenized RWA value (excluding stablecoins) reached roughly \$32 billion by mid-June, up about 300% year-over-year. Critically, the market shifted from being concentrated primarily in tokenized U.S. Treasuries to six asset classes each exceeding \$1 billion onchain: private credit, commodities, U.S. Treasuries, corporate bonds, non-U.S. government debt, and institutional alternative funds. Traditional disclosure practices and periodic attestations proved insufficient for markets operating continuously across jurisdictions and time zones.

Multi-chain liquidity became systemic infrastructure, so institutions can no longer evaluate asset risk within a single-chain environment. And regulators shifted emphasis to demonstrable controls rather than stated controls, with the GENIUS Act in the U.S. and the EU's MiCA providing clearer legal footing for tokenized assets and their stablecoin settlement infrastructure.

Q2 accelerated the transition: onchain RWA holder addresses approached 920,000 by June, up over 15% in 30 days, and scrutiny of interoperability, bridge, and settlement risks increased following several large security incidents involving fragmented liquidity and subpar cross-chain infrastructure. Collectively, these developments elevated verification from a competitive differentiator to a baseline requirement.

As digital asset markets mature, trust is increasingly established through cryptographic verification and independent data validation rather than institutional reputation alone.

— Chainlink

Protocol-level protection layers across surfaces that get exploited

This model is Hacken’s, derived from the firm’s audit and incident-response data. It maps controls to the surfaces where losses now occur rather than to the contract code alone: on-chain, operational, dependency, and, for exchanges, the CEX layer. Each row states the control and its mechanism, the failure it rules out, and the evidence an outsider can verify. Because a protocol is exposed at its weakest surface, the layers are assessed together; in Q2, that surface was operational.

On-chain security

What	How	What it rules out
Smart contract audit & coverage	Independent validation that deployed code matches reviewed code, with measurable scope	Most contract-level vulnerabilities adversaries can reach
Vulnerability disclosure program (bug bounty)	Funded public program, severity-tiered payouts scaled to value-at-risk, defined scope, safe-harbor terms	Researchers exploiting or selling zero-days
Anomaly detection	Real-time detection of abnormal on-chain behavior with escalation and on-call rotation	Unnoticed outflows, privileged calls, oracle deviation, large transfers
Prevention controls	Timelocks on privileged actions, withdrawal rate limits, pausable contracts / circuit breakers, guarded launch	Unbounded damage mid-attack; instant irreversible execution

Operational Security

The people layer is the hardest to protect, and regulators increasingly expect it to be certified if a party is aiming to acquire liquidity.

What	How	What it rules out
Key management & signing discipline	Hardware-isolated signing devices; MPC or threshold signatures; no daily-driver devices in approval flows	Transaction signer compromise
Organizational security management	ISMS covering access control, change management, incident response	Backend/admin-path compromises
Resilience testing	Penetration testing, red-teaming, incident-response rehearsal	Untested assumptions about detection and response speed

Dependency / Third-party

What	How	What it rules out
Cross-chain & oracle dependencies	Multi-validator attestation; no single-verifier release paths	Single-point verification failures
Software supply chain	Dependency pinning, provenance checks, agent-skill vetting	Shai-Hulud-class worm propagation; malicious AI skills
Register of information	Maintained inventory of ICT third parties (DORA Art. 28)	Unknown concentration risk

Stellar

 Stellar

Industry perspective



Today, institutions are moving regulated assets and critical financial infrastructure into production, and that requires a different standard of trust – one built on independently verifiable security, operational resilience, transparent governance, compliance by design, institutional assurance, and the ability to continuously verify network operations.

Those requirements have shaped the architecture of the Stellar network from the beginning. For more than a decade, the Stellar network has evolved alongside financial institutions and developers, delivering the reliability required for production while processing billions of operations. The Stellar Consensus Protocol provides an open and transparent validator model where participation and consensus can be independently verified. Opt-in native capabilities such as authorization and clawbacks provide institutional controls, while privacy features are being extended to support evolving regulatory and operational requirements. Combined with an open-source codebase and support for independent audits and verification, the Stellar network can give these institutions more than just confidence in the technology, it gives them evidence to build this foundation of trust on.

As institutional expectations continue to evolve, our focus remains the same: building infrastructure where security, resilience, governance, compliance, transparency, and trust are architectural properties of the Stellar network.

— Tomer Weller, Chief Product Officer, Stellar Development Foundation

CEX level protection layers across surfaces that get exploited

For exchanges, trust reduces to three questions a counterparty has to answer with evidence: is it secure, is it solvent, and is it transparent. Each resolves into concrete controls, and each control exists because a specific failure happens without it.

Security is the deepest of the three, because an exchange custodies user funds directly.

Security

Control	Why it matters
Server and infrastructure hardening	TLS, web application firewalls, DNS and email authentication, and secure HTTP headers close the internet-facing surface an attacker probes first.
User-account security	2FA, withdrawal whitelists, anti-phishing codes, and session management stop a stolen login from turning into a withdrawal.
Independent certification	ISO 27001 and CCSS mean a third party has checked the security program and key-management practice, not just the exchange’s own claims about them.
Bug bounty	A funded disclosure program routes vulnerabilities to the exchange rather than to the market, and its payout history shows whether researchers trust it enough to report.
Penetration testing	A recurring independent test against a production-mirroring environment is the only control that verifies defenses against a real attack instead of on paper.
Insurance fund	A reserve fund absorbs breach losses so a single incident does not fall entirely on users.

Solvency

Control	Why it matters
Proof of reserves across the full user base	A proof that covers only part of the book can hide the exact shortfall it is meant to disprove.
High-quality reserve composition	Reserves held in liquid assets rather than the exchange’s own token can be honored in a run; a self-issued reserve collapses at the moment it is needed.
Frequent, recent attestation	Solvency changes daily, so a once-a-year snapshot says little about the balance today. Cadence is part of the proof.
Merkle-tree verification	A Merkle tree lets each user confirm their balance is counted, without the exchange being able to quietly leave liabilities out.

Transparency

Control	Why it matters
Live reserve-wallet tracking with proof of ownership	On-chain wallets a user can watch in real time, proven to belong to the exchange, replace a quarterly report with continuous evidence.
Incident-response quality	How fast an exchange reacts, how much it discloses, and whether it can freeze funds or pause withdrawals decides how much of a loss is recoverable after an attack.



The security incidents of Q2 reveal a structural trend: attackers no longer need to 'breach' defenses — they can make defenses fail while operating normally. Every link in the trust chain was 'correctly executing' its own logic, but trust propagation itself broke down. Security is not about building any single wall higher — it's about ensuring independent verification at every interface where trust is transferred.

At Bybit, our Q2 focus has been eliminating these trust blind spots: ensuring what-you-see-is-what-you-sign in multi-sig workflows, eliminating single-point validation dependencies in cross-chain operations, preventing the introduction of unaudited components at the cryptographic implementation layer, and severing the propagation of unverified data at the contract layer — giving every trust interface the ability to judge independently, rather than blindly accepting conclusions passed down from upstream.

Our objective is to compress cross-trust-boundary anomaly detection and interdiction response into a time delta the adversary cannot exploit.

— Alan Xin, Head of Blockchain Risk Control at Bybit

Where does architecture of trust go for digital assets?

An audit is the clearest example. Two projects can both display "audited" and mean entirely different things by it. One reviewed its full deployed system; the other scoped the review down to a single contract and left the rest uncovered. The scope statement that would separate them is often not public. Neither is the report, the monitoring coverage, or the key-management setup. The market rewards the signal and rarely inspects what stands behind it.

Even the entry level is thin. Roughly a quarter* of the risk-indexed market has an audit on record at all, far fewer disclose its scope, and almost none run the monitoring layer that Q2's losses ran straight through. So the same "audited" signal can cover a project that continuously verifies its operations and a project that reviewed one contract two years ago, and an allocator reading from the outside cannot always tell which is which.

That leaves the question open rather than answered. If a minimal audit and a one with bigger coverage weigh the same, if the evidence that would separate them is not published, and if the strongest controls are the rarest in the market, what should count as the bar for trusted, and who gets to verify it holds?

*Figure is based on Hacken's CORE3 data from 1427 protocols above \$1M marketcap, verifiable at core3.io.

Risk Intelligence for Asset Managers

The industry has developed a mature disclosure practice around exactly two instruments. Audits are published and referenced in marketing, and the most recognized bug bounty programs are public. Deeper risk-management metrics, signer architecture, response readiness, and monitoring coverage, rarely reach comparable visibility, and the industry holds two opposing views on whether they should:

Publishing a risk-management posture exposes the gaps of the least protected projects and therefore increases their risk.

Disclosure demonstrates credibility and builds trust, including for projects whose posture is incomplete.

Our assessment weighs against the black-box position. Crypto is a transparent-by-design ecosystem at the settlement layer, yet no comparable culture of demonstrating operational security has formed on top of it.

The remainder of this chapter sets out the data points every Web3 protocol should be in a position to present, publicly or privately, during a due diligence process, across the four stages relevant to institutional practice: pre-investment screening, detection of an ongoing incident, exposure transmitted through counterparties, and post-incident recovery.

Pre-incident Risk Posture Priority Matrix

Institutional review converges on five question sets:

01

How secure each part of the protocol is: how well the smart contracts are covered by audits, whether the audit is current relative to deployed code, and how the team handled past issues.

02

The team itself, where at least one member has to be reputable and identifiable with a record that can be checked.

03

Whether the project does what it says, which means comparing its documentation, whitepaper, and token distribution against actual on-chain activity, one of the most revealing checks and the hardest to run.

04

Whether a response plan exists, how fast it executes, who owns it, and whether a timelock buys the team time to react.

05

Who controls each part of the infrastructure, where the points of failure sit, and how exposed the signer keys are. Q2's two largest losses were both answers to that last question that nobody had asked in time.



In Q2 2026, the signal that most frequently led us to pass on an otherwise attractive position was inadequate security relative to the capital at risk, including stale or limited-scope audits and the absence of robust operational safeguards such as withdrawal-address whitelisting, timelocks on administrative actions, and multi-party or multi-verifier controls designed to eliminate single points of failure.

The infrastructure-layer exploits observed during the quarter reinforced this focus. Several major incidents resulted from off-chain or infrastructure-related compromises rather than vulnerabilities in the underlying smart-contract code, confirming that on-chain audit quality alone is no longer sufficient. We have therefore placed greater emphasis on, and now screen explicitly for, off-chain and infrastructure resilience, single-key or single-verifier dependencies, and potential contagion through shared collateral and interconnected markets.

Against this backdrop, we have remained selective and appropriately cautious in our DeFi deployments following the market stress experienced during the quarter.

Federico Bagioti, Group Head of Risk Management at Abraxas Capital

Real-time Signals: Characteristics of an Ongoing Incident


Most attacks produce limited observable evidence before funds move. Two early indicators carry information value. A cluster of failed interactions within a short window is consistent with an attacker testing configurations against a contract. A sudden oracle-price movement is the second, although the interval between the movement and the drain is typically measured in blocks rather than minutes. Over the past year the preparation phase of major attacks has become substantially less observable, and the first reliable evidence is frequently the outflow itself. This places the analytical weight on transaction monitoring, subject to one constraint: alert precision determines monitoring value, since false positives consume the response capacity required for a confirmed event.

Hacken’s A3 converts these signals into a monitored risk register across five detector groups.

Detector group	Alerts
Contracts change in ways they should not	proxy upgrades to unexpected implementations, unauthorized function calls, critical events, failed-transaction clusters, non-whitelisted or blacklisted callers
Financial backing or token economics drift from plan	reserves below threshold, sharp TVL drops, unexpected minting (Transfer-from-zero), allocation and vesting violations, large transfers
Compliance risk becomes visible on-chain	AML and sanctions exposure, suspicious counterparties, compliance-score movement, an evidence trail for DORA, MiCA, and ADGM obligations
Operations show early warning signs	repo inactivity or emergency commit spikes, DNS and WHOIS changes, Safe signer additions and removals, permission changes
What changed, who acted, and whether the team responded	alert history, resolution workflow, tamper-evident audit trail

The detector coverage corresponds to the quarter’s incidents as follows. Safe-signer monitoring covers the Drift entry vector, the signer changes that preceded the \$285M drain. Supply-to-collateral monitoring detects KelpDAO-class unbacked minting before a borrow is executed against it. Mint monitoring detects issuance events such as Humanity Protocol’s 100M unbacked \$H. DNS monitoring covers the quarter’s single frontend hijack.

Post-incident: Recovery as a Measurable Property



Q2 recovery totaled 11.5%, or \$87.8M, of which a single \$71M freeze at KelpDAO accounts for the majority. Recovery outcomes correlated with preparation: monitoring that identifies movement early enough to support a freeze, standing relationships with the exchanges and stablecoin issuers who execute freezes, a response plan with named owners, and timelocks that extend the available response window.

The difference between Drift, where the funds went dormant pending an expected multi-year cashout, and KelpDAO, where approximately 24% was frozen, is largely attributable to the speed at which downstream actors were engaged.

DA Insured

Industry perspective



Digital asset crime losses are not a tail event; they are a persistent, material feature of the market. H1 2026 alone saw over \$1.7 billion lost to hacks, fraud, and unauthorised transfers. Insurance demand, whether from operators, regulators, or end clients, flows directly from that reality.

Our focus is on regulated custodians and selected exchanges, and the substantive discussion is about the effectiveness of controls: cryptographic key management, segregation of duties, security architecture, third-party dependencies, and business continuity. What we value most is independently verified evidence, ISAE 3402 and SOC 2 Type II reports, and critically, how a leadership team responds to findings. We are not looking for perfection; we are looking for substantive remediation and a pattern of continuous improvement. Platforms that come to the table with a clear audit trail and genuine senior engagement tend to move through the process efficiently, and that reflects in the terms we are able to offer.

— Alistair Heggie, Founder & CEO, DA Insured

Howden

Industry perspective



The multisig-related Drift exploit highlights the gap: audits, multisig, and 'insured' infrastructure can all function correctly and still fail. The real vulnerability is end-to-end operational resilience, not just insurable events. What Q2 exposed is the need for pre-incident risk intelligence, active monitoring and governance controls, and incident-response coordination and recovery capabilities.

This is precisely the gap Howden's Web3 ecosystem is targeting: moving beyond reactive insurance toward embedding prevention through security and governance review, operational resilience, incident response, and legal and recovery coordination. The ecosystem is ideally positioned for the asset and investment manager sector: firms entering the space to offer their investors an alternative asset class, which may not yet be as sophisticated within the crypto ecosystem in areas such as private key management and the regulatory framework. The ecosystem brings this to clients, offering consumer confidence and operational resilience.

— Freddie Palmer, Executive Director, Financial Lines. Head of Digital Asset and Blockchain at Howden

Conclusion

Hacken's incident data and both contributing institutions converge on two findings. The first is that verifying the existence of an audit is no longer a sufficient screen: Q2's dominant losses originated off-chain, in signer sets, verification paths, and operational infrastructure that no code review is scoped to cover, and 14 of the quarter's 67 victims held audits at the time of the exploit. The second is that institutional due diligence is adjusting accordingly.

The data points entering allocation frameworks are no longer limited to point-in-time attestations; they now include continuous observables, signer-set changes, collateral backing, dependency exposure, response readiness, and extend past the fact that a certificate exists to what it examined, how recently, and against which deployed code.

This suggests the direction of institutional allocation in this market. Screening built on static documents priced Q2's risks incorrectly; frameworks built on measured, continuously updated posture identified them. As allocators formalize that shift, capital can be expected to concentrate in the segment of the market that can produce standing evidence of its controls, currently a minority of the index, and to price the remainder as unverifiable.

Conclusions & Recommendations

Q2 2026 was the worst quarter for Web3 since Q2 2025, accounting for \$763,971,791 in losses across 67 incidents. Adversary activity continued its shift toward the operational infrastructure of Web3 projects: compromised keys, signer devices, and cross-chain verification paths carried 88.3% of the stolen value, with the two DPRK-attributed operations at Drift and KelpDAO alone responsible for 75.5% of the quarter. At the same time, smart contract vulnerabilities remained the most common exploit cause: 44 of 67 incidents, facilitating the theft of \$87.7 million, about 11% of the total.

This quarter's most frequent failures and its most expensive failures occurred on different surfaces, and no instrument in the market's current trust vocabulary covers both. Fourteen of the exploited projects were audited, ten held more than \$50M in value locked, and eighteen had operated for three years or longer. Every signal the market has treated as proof of safety was present at the point of loss.

This is what the Architecture of Trust is built to answer. The framework maps controls to the surfaces where losses occur: on-chain, operational, dependency, and, for exchanges, the CEX layer, and requires them to be assessed together, because a protocol is exposed at its weakest surface. In Q2, that surface was operational; in another quarter, it will move, and a trust model anchored to one instrument will miss it again.

The consequence is that the previous standard of trust is no longer sufficient, and access to institutional capital is being gated behind a security entry barrier that is forming right now. Our contributors, writing independently across every chapter of this report, describe the same barrier from their own positions in the market:

On-chain and verification

Chainlink: proof-of-reserve and cross-chain verification transitioned from risk-mitigation tools into baseline requirements for institutional participation; trust is established through cryptographic verification and independent data validation rather than reputation alone.

Bybit: every link in a trust chain can execute its own logic correctly while trust propagation itself breaks down; every trust interface needs the ability to judge independently rather than accept conclusions passed from upstream.

Sui Foundation: the audit is not the finish line; the protocols that come through Q3 intact will be the ones that detect quickly and pause cleanly, not the ones holding the most audit reports.

Operations

Moody's Ratings: the binding constraint on institutional trust remained cyber risk, rooted in operational security, access controls, third-party dependencies, and key management; the bar has shifted from static assurance toward continuous, demonstrable resilience under stress, including forward-looking cryptographic readiness.

Abraxas Capital: the signal that most often vetoed an otherwise attractive position was inadequate security relative to capital at risk; screening now explicitly covers off-chain and infrastructure resilience, single-verifier and single-key fault lines, and contagion exposure through shared collateral.

DA Insured: what underwriters value most is independently verified evidence, ISAE 3402 and SOC 2 Type II, and how leadership responds to findings; the search is for substantive remediation and continuous improvement, not point-in-time perfection.

1Money: the privileged operational layer around mint, burn, freeze, and administrative authority is the attack surface to watch; the contracts are audited heavily, while the identity, signing, and tooling that control them are not.

Compliance and governance

EisnerAmper: institutional due diligence questionnaires are no longer checkbox exercises; GAAP-audited financials and SOC I and II reports have moved from differentiator to cost of entry.

Allium: regulators expecting continuous surveillance are moving faster than the data infrastructure built to support them; cross-chain attribution and agentic payments sit outside current monitoring coverage.

Stellar: institutions moving regulated assets into production require evidence to build trust on, not confidence in technology alone: verifiable security, resilience, transparent governance, and compliance by design.

SUI Foundation



Industry perspective



In Q3, we expect an increase in security incidents related to legacy code. AI tools are accelerating the discovery of latent economic and accounting flaws that have sat undetected for years. Subtle bugs in core DeFi code functions like liquidity pools and vaults have always existed, but AI lowers the cost and time required to surface them. In short, AI enables attackers to find vulnerabilities that previously remained undiscovered for years.

Vulnerabilities often don't live in any single contract or module, rather they can emerge based on how multiple smart contracts and modules interact. A new module combined with legacy code can introduce a vulnerability, especially in unaudited upgrades. This is also why audits aren't sufficient: they're a point-in-time review, while the pricing and reward calculation flaws that drain protocols depend on external dependencies that a static review can't fully model. AI compounds this by flooding security teams with plausible but false vulnerability reports, straining already limited resources.

The solution is to stop treating audits as the goal, and instead focus on detection and response operations such as continuous monitoring tuned to economic and pricing anomalies rather than static code alone. Implementing circuit breakers that automatically pause or rate-limit a protocol the moment a critical anomaly is detected, and a 24/7 security on-call will cut response time from hours to seconds. The protocols that earn user trust and grow through Q3 will be the ones that detect quickly and pause rapidly; those that don't will continue to experience an increase in successful hacks which cause economic and reputational harm.

— Sean Spaniol, Head of Cybersecurity Sui Foundation

Exploit Resilience in Q3: Consolidated Recommendations

The recommendations below consolidate Chapters I through V into a single set, grouped by audience. None of them is speculative; each corresponds to a control whose absence produced a documented Q2 loss.

Protocol Teams

- Hardware-isolated, dedicated signing devices for all multisig signers; no daily-driver laptops in approval flows.
- Timelocks on treasury and admin multisigs, creating a detect-and-stop window between approval and irreversible execution.
- Independent multi-validator attestation for all cross-chain releases; no single-verifier paths.
- Real-time supply-to-collateral monitoring with automated circuit breakers on mint and release volume.
- Move mint and proxy-admin roles to MPC or threshold signing; eliminate single-key control of privileged functions.
- Manage deprecated contracts deliberately: retain monitoring, bounty, and audit coverage until the balance is zero, and wind down via withdrawals only.
- Treat any AI agent with transaction authority as a privileged signer: scoped identity, spend and rate limits, explicit approval gates, and a deterministic policy layer that authorizes actions independently of the model that recommends them.

CASPs and Exchanges

- Implement DORA as architecture: hardware-isolated signing, timelocked admin functions, and admin-path controls tested under realistic conditions rather than described in a policy.
- Extend the Register of Information to the dependencies that failed in Q2: bridge validators, oracle providers, messaging layers, RPC infrastructure.
- Run resilience testing against operational scenarios, signer compromise and forged cross-chain release, not only against contract code.
- Answer the three counterparty questions with standing evidence: secure (certification, penetration testing, bounty history), solvent (full-book proof of reserves with frequent attestation and Merkle verification), transparent (live reserve-wallet tracking and demonstrated incident-response capability).

Stablecoin Issuers

- Treat mint, burn, freeze, and admin authority as regulated critical functions: threshold signing behind a timelock, no single-key control. Stablr's \$2.8M loss through one of three mint keys is the reference case.
- Monitor supply-to-reserve ratios in real time, with automatic circuit breakers on mint volume.
- Match monitoring coverage to the asset's actual chain footprint; an Ethereum-only view covers roughly half the regulated perimeter.
- US issuers preparing for GENIUS should implement the operational-security controls the application does not require, because reclassification as a BSA financial institution imports no DORA-style resilience standard.

Asset Managers and Allocators

- Stop treating "audited" as a screen; verify scope, recency, and match to deployed code.
- Screen the five question sets: security coverage, identifiable and reputable team, does-what-it-says (documentation against on-chain reality), response readiness and timelocks, and control topology with signer-key exposure. Two largest losses were answers to that last question that nobody had asked in time.
- Add continuous observables to due diligence: signer-set changes, collateral backing, dependency exposure, contagion and shared-collateral risk.
- Weigh independently verified evidence (SOC 2 Type II, ISAE 3402) and how leadership responds to findings over point-in-time perfection.

Conclusion: The Direction of Institutional Capital

Institutional adoption is no longer constrained by blockchain technology. It is constrained by the industry's ability to establish trust and to keep verifying it. Q2 made the cost of the old model measurable: screening built on static documents mispriced the quarter's risks, while frameworks built on continuously updated posture identified them, sometimes to the block.

As allocators formalize that shift, the market will divide along a single line. Capital will concentrate in the minority of projects and platforms that can produce standing evidence of their controls, and the remainder will be priced as what, to an outside observer, they are: unverifiable. Today that minority is small, 4% of the risk-indexed market runs audit, bounty, and independent monitoring together, which means the entry barrier now forming is also the largest open opportunity in the industry.

The most trusted counterparties in Q3 will not be the ones with the longest history, the largest value locked, or the best-known audit on record. They will be the ones who can continuously prove they remain safe today, regardless of partnerships, certifications, or sentiment.

In Q3, trust is proven continuously, or not at all.

About The Research

The findings in this report are drawn from Hacken's own practice areas: exploit and incident data from the Extractor research team; smart contract vulnerability analysis from Offensive Security Services; protocol- and exchange-level risk exposure data from CORE3 risk infrastructure; and regulatory analysis from compliance engagements conducted by the GRC team.

Hacken Track Record

2,069

public security assessments delivered since 2017

24,376

vulnerabilities identified across engagements

\$430B+

in assets verified through Proof of Reserves services

ISO 27001

certified information security management

Hacken Practice Areas

The specialists contributing to this report work across the following practices:

Offensive Security

- [Smart Contract Audit ↗](#)
- [Blockchain Protocol Audit ↗](#)
- [Penetration Testing ↗](#)
- [dApp Audit ↗](#)
- [Proof of Reserves ↗](#)
- [Tokenomics Audit and Design ↗](#)
- [Yield Audit ↗](#)
- [Cryptography Review ↗](#)
- [AI Security Audit ↗](#)

Compliance and Advisory

- [MiCA and DORA Compliance ↗](#)
- [NIS2 Readiness ↗](#)
- [SOC 2 Readiness ↗](#)
- [ISO 27001 ↗](#)
- [ISO 42001 Readiness ↗](#)
- [EU AI Act Readiness ↗](#)
- [CCSS Audit ↗](#)
- [Executive Cybersecurity Training ↗](#)
- [Licensing and Advisory ↗](#)

Continuous Security

- [Extractor ↗](#)
Real-time on-chain risk monitoring and incident response.
- [A3 ↗](#)
Continuous monitoring for smart contracts, wallets, reserves, and governance activity.
- [CORE3 ↗](#)
Web3 risk infrastructure indexing Probability of Loss for digital assets and exchanges across 85+ risk parameters.
- [DualDefense ↗](#)
Post-audit assurance

Contributing Partners

























We thank our contributing partners for sharing their expertise and operational perspectives throughout this report.

Authors & Contributors: A3 Product team (Risk Intelligence for Institutions), Anton Sheptytskyi (Design & Visuals), CORE3 Product Team (Project-level Risk Data), Dmytro Yasmanovych (Compliance), Dmytro Zaporozhchenko (Content Lead), Extractor by Hacken (Research & Data Analysis), Grzegorz Trawiński (Q2 threats and Q3 trends), Maksym Fedchak (Partner Outreach), Olesia Bilenka (Smart Contract Vulnerabilities Data), Oleksandr Iurkin (Strategic Oversight), Stephen Ajayi (AI Security Research), Valeriia Skorik (Production & Coordination).

Securing Digital Frontier Since 2017

 Hacken is a leader in blockchain security, combining deep technical expertise in Web3 with real-time data-driven insights.

ISO 27001:2022 Certified

“Hacken is praised for their professionalism, thorough audits, and timely delivery.”

[Read all reviews on Clutch](#)

Trusted by industry leaders for 8+ years

Tier-1 CEXes

Wallets




Regulators, Financial Institutions & Banks

L1 / L2 Protocols, Platforms & Foundations

DeFi, RWA, GameFi & Ecosystem Projects

Our Story

Unlike traditional providers, Hacken was born on blockchain, combining deep Web3 expertise with enterprise-grade quality, AI-powered offensive security, and globally recognized certifications. Since 2017, Hacken has been trusted by startups, enterprises and regulators to secure the new digital frontier.

Learn more and follow us on social media:

[hacken.io](#)

[linkedIn](#)

[X](#)