HACKEN

Q1 2026 Security & Compliance Report

Why continuous protection defines the next chapter of blockchain resilience

Contributing partners:

KUCOIN · GLOBAL LEDGER · MEXC · whitebit · Centrifuge

SVRN · MD · C4 · GRAY WOLF · BYBIT · Allium

"Over $482 million lost in a single quarter. In our most comprehensive security report to date, we bring together perspectives from KuCoin, MEXC, WhiteBIT, Global Ledger, Centrifuge, Bybit, SVRN, Allium, C4, M0, and Gray Wolf alongside Hacken's own incident data and audit findings, deep diving into smart contract vulnerabilities, stablecoin security, proof of reserves, compliance demands, and AI implications.

The findings are consistent across every contributor: failures are emerging across code, infrastructure, operations, and human processes simultaneously. No single audit, certification, or tool addresses that. Continuous, layered protection is the only posture that works."

Yev Broshevan

CEO & Co-Founder, Hacken

Q1 2026 at a glance:

- $482.6 million lost in Q1 2026 – a 20.9% increase over Q4 2025's $399.3M.
- Phishing and social engineering dominated losses at $306 million (63.4% of total), with a single $282M hardware wallet scam accounting for more than half the quarter.
- Smart contract vulnerability losses surged 213% compared to Q1 2025, reaching $86.2M across 28 exploits – the highest SC loss quarter since Q2 2025.
- Six audited protocols were still exploited, including Resolv Labs (18 audits) and Venus Protocol (5 audit firms) – underscoring the gap between point-in-time code review and continuous operational security.
- DPRK-linked actors remain active, with Step Finance ($40M) and Bitrefill compromised through documented Lazarus/BlueNoroff playbooks – fake VC calls, malware deployment, and endpoint compromise.

Of the 28 SC-exploited projects, six had prior audits. Those six accounted for $37.7M in losses — a higher average loss ($6.3M) than unaudited projects ($4.3M).

What this report covers

Chapter I: maps Q1 incidents across the security stack — code, operations, and infrastructure — with the five largest incidents analyzed in detail.

Chapter II: breaks down Hacken's Q1 audit portfolio by sector and vulnerability class, with deep dives on emerging standards (ERC-4337, Uniswap v4 hooks, RWA compliance frameworks).

Chapter III: introduces a six-layer stablecoin security architecture and maps Q1 audit findings to each layer, from reserve custody through cross-chain bridge security.

Chapter IV: examines the AI threat paradigm — from vibe-coded exploits to agent-based attack vectors — and the gap between AI adoption and AI security hardening.

Chapter V: tracks Q1 regulatory developments across MiCA/DORA, VARA/CMA, US stablecoin law, and MAS, with a framework for compliance-as-security-management.

Chapter VI: synthesizes recommendations and Q2 outlook, with closing perspectives from industry contributors.

Table of Contents

1 Major Security Incidents & Financial Impact Analysis

The first quarter of 2026 recorded $482.6 million in total losses across 44 incidents, spanning smart contract exploits, access control breaches, phishing, DNS hijacks. While the headline figure is 76.6% lower than Q1 2025's extraordinary $2.06 billion, several underlying trends signal that the threat landscape is shifting rather than shrinking.

Smart contract vulnerability losses surged 213% quarter-over-quarter compared to Q1 2025, reaching $86.2 million across 28 incidents. Phishing losses hit $306 million, driven by a single $282 million hardware-wallet social-engineering attack and a $24 million address-poisoning theft. Meanwhile, access control failures accounted for $71.9 million, including the $25 million Resolv Labs key compromise that minted 80 million unbacked stablecoin tokens.

A recurring pattern emerged: audited projects continued to get exploited. Six of the 28 smart-contract incidents hit protocols with prior audits, including Venus Protocol (five audit firms), Solv Protocol (three firms), and Resolv Labs (18 audits). Some vulnerabilities that attackers exploited often fell outside the scope of traditional code reviews – oracle manipulation, donation attacks on Compound-fork lending pools, and off-chain infrastructure compromises.

Q1 2026 vs Q1 2025

Metric	Q1 2026	Q1 2025	Change
Total losses	$482,661,580	$2,063,445,000	-76.6%
Total incidents	44	32	+37.5%
SC vulnerability losses	$86,173,580	$27,520,000	+213.2%
Access control losses	$71,900,000	$1,628,770,000	-95.6%
Phishing losses	$306,000,000	$117,370,000	+160.7%
DeFi protocol losses	$133,461,580	$80,685,000	+65.4%

Incidents per Network

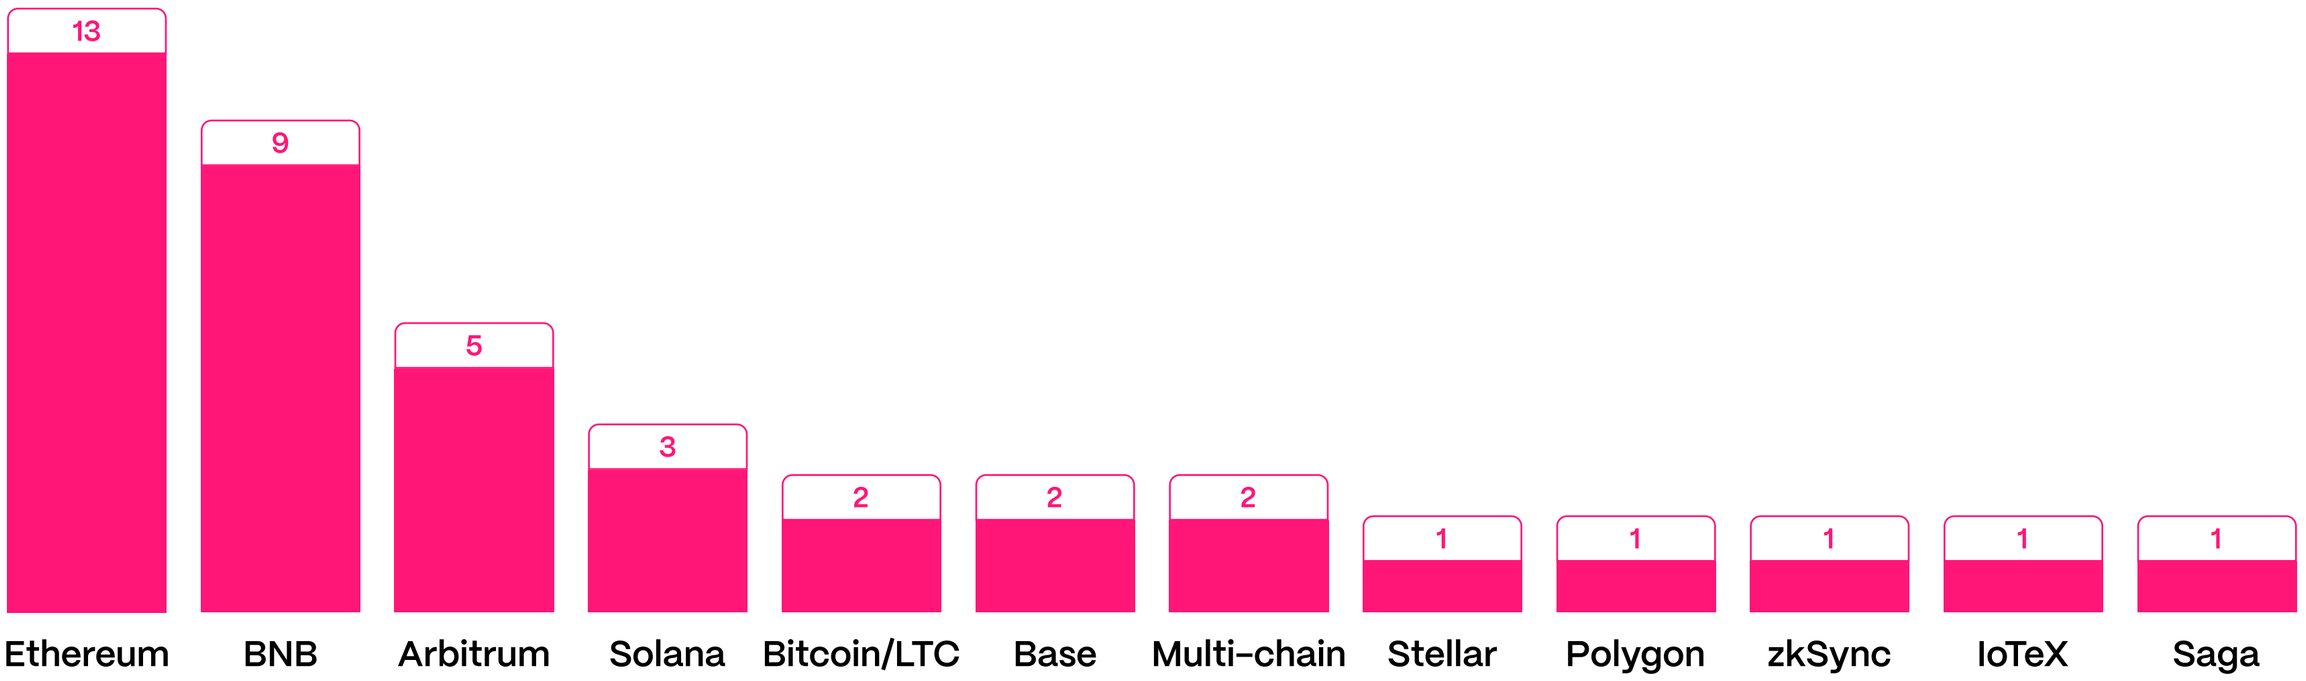

INDUSTRY PERSPECTIVE

The detection gap: why minutes matter more than audits

Global Ledger's 2025 Laundering Race data reveals that teams report hacks on average 1.5 days after an incident — while hackers begin moving funds before the report in 76% of cases. We asked their CEO what real-time monitoring needs to look like.

Looking at Q1 incident data, at what point in the attack timeline did on-chain signals become visible — and what would a real-time monitoring system have needed to catch them before execution?

> Time-to-detection and response have to become the industry benchmark. With the right tools, you can catch unusual hot wallet outflows, sudden liquidity drops, interactions with mixers or unfamiliar bridges, unexpected token approvals, suspicious contract calls.
>
> But knowing what to look for is half the battle. The window between "something looks wrong" and "funds are gone" is often minutes. The only thing that can actually help is flagging the activity publicly, notifying exchanges, posting the attacker's addresses to get a better chance of a freeze.
>
> Our **2025 Laundering Race data** shows teams report hacks on average 1.5 days after an incident. It's too late: hackers start moving funds before the report in 76% of cases. Regulators already know about this gap, and it's just a matter of time when time-to-detection and response benchmarks appear. Being slow can cost you twice: processing dirty money and answering regulator's questions.
>
> **Lex Fisun**
> CEO and Co-Founder, Global Ledger

GLOBAL LEDGER

Total Losses and Quarterly Trends

Q1 2026 losses of $482.6 million place this quarter as the second-lowest first quarter since 2023. The absence of a single mega-hack on the scale of DMM Bitcoin ($305M, Q2 2024) is the primary driver of the year-over-year decline. However, the incident count rose 37.5% to 44, indicating a broader attack surface.

Compared to Q4 2025 ($399.3 million across approximately 30 incidents), Q1 2026 represents a **20.9% increase** in total losses and a notable uptick in incident volume. The quarter-over-quarter increase suggests that the fall in losses in late 2025 was temporary rather than a sign of structural improvement.

Quarterly Loss Trends (2024–2026)

Quarter	Total Losses	QoQ Change
Q1 2024	$999,587,674	–
Q2 2024	$639,928,000	-36.0%
Q3 2024	$538,121,000	-15.9%
Q4 2024	$282,193,000	-47.6%
Q1 2025	$2,063,445,000	+631.2%
Q2 2025	$1,030,521,000	-50.1%
Q3 2025	$510,856,000	-50.4%
Q4 2025	$399,268,000	-21.8%
Q1 2026	$482,661,580	+20.9%

Total Crypto Losses per Quarter

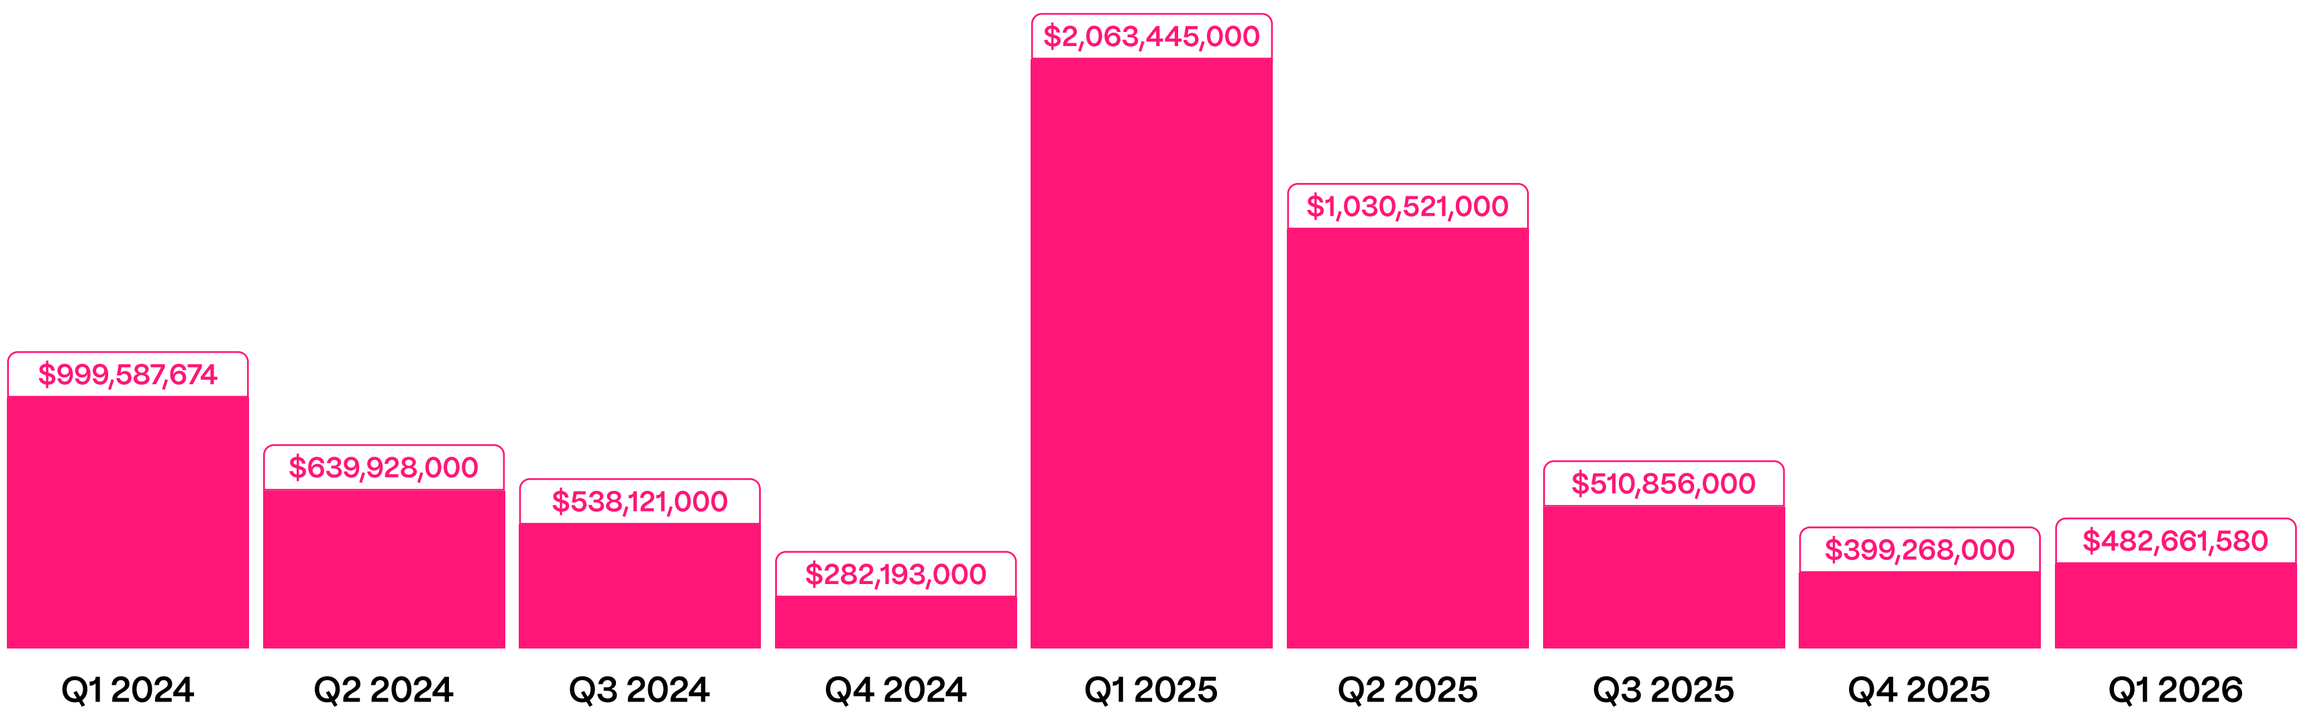

The loss-per-quarter trend in 2025 showed a steady descent from the Q1 peak, similar to 2024. Q1 2026 breaks this pattern with a slight uptick, driven by the convergence of a $282M phishing theft, the $25M Resolv Labs key compromise, and a sustained volume of mid-tier DeFi exploits.

INDUSTRY PERSPECTIVE

The three pillars Bybit is building on

"Effective continuous security protection is not a toolset — it's an operational discipline."

Based on Q1, what does effective exchange security require?

First, AI Agent-driven real-time monitoring. Anomaly detection must operate at minute-level response times, yet under traditional approaches, security experts spend the bulk of their time on heavy-lift data work: ingestion, normalization, and correlation — leaving little bandwidth for actual decision-making.

At Bybit, our direction is to deploy AI Agents that take over the data operations layer entirely — continuously ingesting on-chain data, internal system logs, and behavioral signals, performing automated correlation and initial triage. This shifts expert effort from "finding the needle in the haystack" to "making the final call" — accelerating response times while significantly reducing alert fatigue.

Second, people are the largest attack surface. We enforce strict least-privilege access, multi-party authorization, and ongoing security awareness training for high-privilege personnel.

Third, security must be embedded in business processes, not bolted on. Every system change and every new feature launch must treat security and risk control review as a prerequisite, not an afterthought.

BYBIT

Monthly Breakdown

Loss concentration was extreme. January accounted for 77.8% of total Q1 losses, almost entirely due to one social-engineering attack.

Month	Incidents	Total Stolen	Largest Incident
January	**19**	**$375,319,400**	**Social engineering ($282M)**
February	**11**	**$26,626,180**	**YieldBlox ($10.86M)**
March	**14**	**$80,716,000**	**Resolv Labs ($25M)**

DeFi losses:

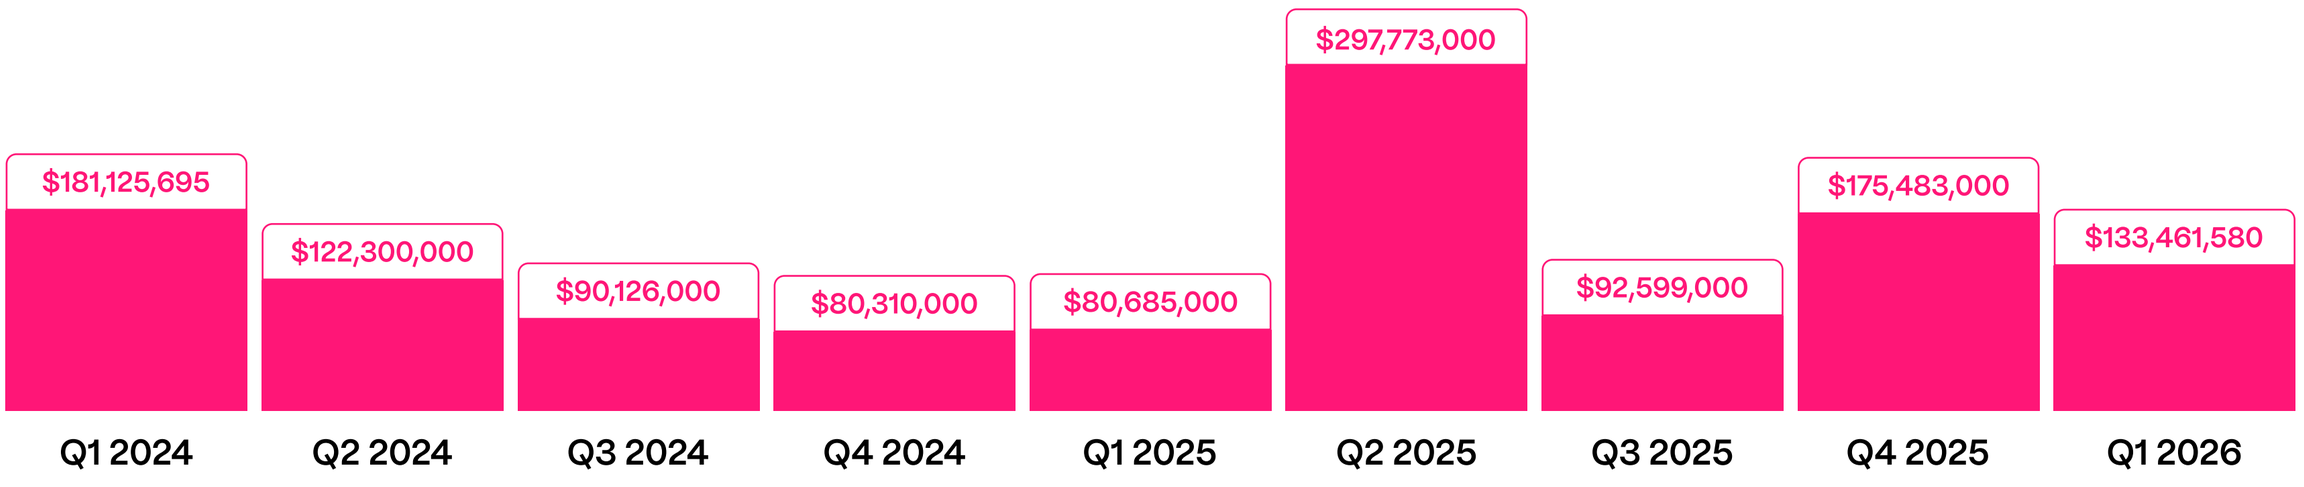

Incident Breakdown by Security Architecture Layer

This framework maps Q1 incidents to the layer of the security stack where the failure occurred, identifying where defensive investment would yield the greatest reduction in losses.

Layer	Attack Vector	Incidents	Losses
Layer 1: Code	**Smart Contract Vulnerability**	**28**	**$86,173,580**
Layer 2: Operations	**Access Control**	**6**	**$71,900,000**
Layer 3: Infrastructure	**DNS Hijacking / Frontend**	**4**	**Unknown**

INDUSTRY PERSPECTIVE

The supply chain blind spot: why most Web3 teams can't even scope their exposure

SVRN's David Schwed on the infrastructure attack vector that separates mature security programs from startups still installing packages from public registries.

Based on what you observed in Q1, which infrastructure attack vectors are most consistently underestimated by Web3 teams — and what separates teams that treat infrastructure security as an ongoing discipline from those that approach it as a one-time configuration?

> Supply chain is consistently misunderstood or outright ignored. You need a full SBOM to understand the entirety of that threat model. Without one, you can't even scope what you're exposed to.
>
> And there's a meaningful difference between installing packages from a public registry versus obtaining the actual source code, vetting it, and standing up private registries you control. That's the line between a startup security program and a mature one.

David Schwed
SVRN

SVRN

Layer 1: Smart Contract Vulnerabilities

Smart contract exploits accounted for the largest number of incidents (28 of 44) and **$86.2 million in losses**. This is a 213% increase over Q1 2025's $27.5 million in SC losses, and represents the highest single-quarter SC loss total since Q2 2025's $263 million (which included the Balancer and Bunni exploits).

Vulnerability Type	Losses	Incidents	% of SC Total
Integer overflow (legacy code)	**$26,400,000**	**1**	**30.6%**
Oracle manipulation / misconfiguration	**$20,710,000**	**5**	**24.0%**
Business logic / validation flaw	**$18,753,014**	**7**	**21.8%**
Inherited / supply chain vulnerability	**$7,000,000**	**1**	**8.1%**
Donation / deposit inflation attack	**$3,957,000**	**2**	**4.6%**
Reentrancy	**$2,774,000**	**2**	**3.2%**
Other (ZK misconfig, flash loan, MEV, etc.)	**$6,579,566**	**10**	**7.6%**

The single largest SC loss (Truebit, $26.4M) was caused by an integer overflow in a contract written in Solidity 0.5.x, deployed approximately five years ago. This is a vulnerability class that modern Solidity versions prevent by default. The fact that legacy code remains deployed and holding value continues to be a systemic risk.

Oracle manipulation and misconfiguration collectively drove $20.7 million in losses across five incidents (MakinaFi, YieldBlox, Moonwell, Inverse Finance, and Venus Protocol). Venus Protocol's case is notable because the attacker spent nine months preparing – quietly accumulating 84% of the THE token supply cap before executing a donation attack that bypassed supply controls and inflated the exchange rate by 3.81x.

Audit Coverage and Effectiveness

Six audited projects accounted for $37.7M in losses — a higher average loss ($6.3M) than unaudited projects ($4.3M). Audited projects hold more value and attract more sophisticated attackers targeting vulnerabilities outside audit scope.

Audit Status	Losses	Incidents	Avg Loss
Audited	**$37,720,000**	**6**	**$6,286,667**
Not Audited	**$17,217,000**	**4**	**$4,304,252**
Unknown (Potentially private audits)	**$31,236,580**	**18**	**$1,735,366**

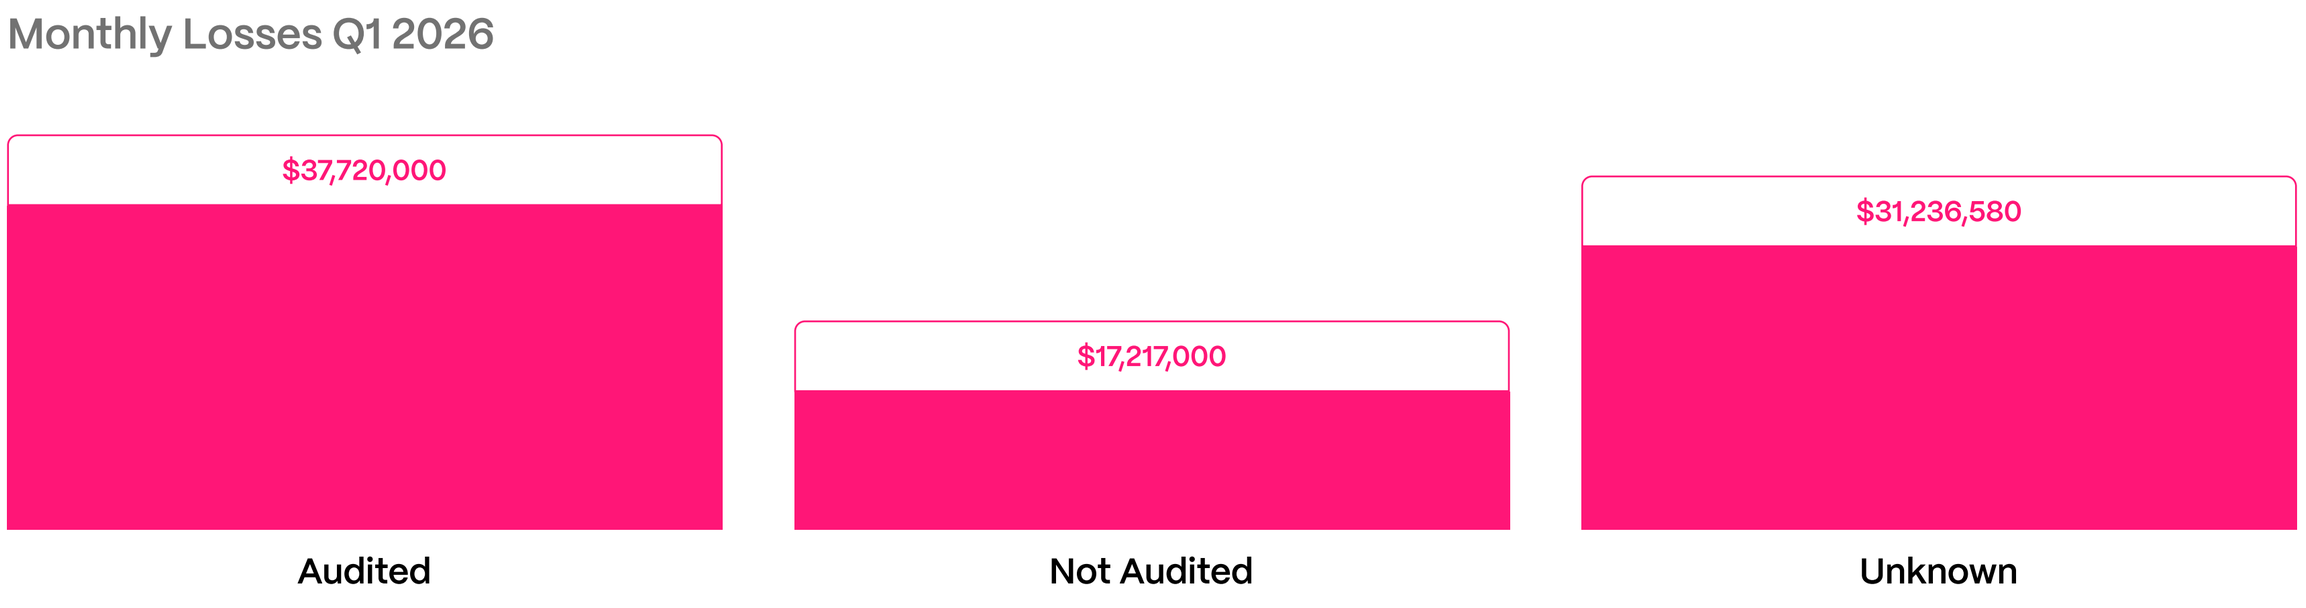

→ How Ongoing Security Would Have Mitigated This

Standard one-time code audits failed to catch several vulnerability classes exploited in Q1. The Truebit overflow existed in code deployed years before the audit landscape matured. Venus Protocol's donation attack is a known Compound-fork vulnerability documented since 2022 – Venus's own ZKSync deployment was hit by the same bug class 13 months earlier.

Continuous security monitoring, real-time anomaly detection on minting and collateral ratios, and protocol-level circuit breakers would have flagged these attacks minutes into execution rather than hours after completion. For oracle-dependent protocols, multi-source oracle validation with deviation thresholds and time-weighted averaging would have prevented three of the five oracle-related exploits.

Layer 2: Access Control and Operational Security

Access control failures accounted for 6 incidents and $71.9 million in losses. The incidents clustered into compromised private keys (Step Finance, IoTeX, USDGambit) and compromised cloud infrastructure (Resolv Labs, Bitrefill).

Project	Loss	Root Cause	Attribution
Step Finance	**$40,000,000**	**Fake VC video call → malicious script → device compromise → key theft**	**DPRK (DangerousPassword cluster)**
Resolv Labs	**$25,000,000**	**AWS KMS compromise → minting key stolen**	Unknown
IoTeX	**$4,400,000**	**Private key compromise on bridge Validator contract**	Unknown
USDGambit / TLP	**$1,520,000**	**Deployer lost account access → ProxyAdmin hijack**	Unknown
Binance MM	**$1,000,000**	**Market maker account compromise**	Unknown
Bitrefill	**Undisclosed**	**Employee laptop → legacy credentials**	**DPRK (Exact cluster not determined yet)**

Audit Coverage and Effectiveness

Two of the six access control incidents in Q1 2026 carry DPRK attribution. Step Finance was compromised through a technique consistent with the DangerousPassword cluster (also linked to the broader Lazarus umbrella), which has been active since approximately 2018. The playbook matches the pattern described in the 2025 Yearly Security Report: attackers impersonating venture capitalists proposed a collaboration opportunity to the Step Finance executive team. During a scheduled video call, "audio issues" prompted the targets to install software disguised as a fix – in reality, a malicious script that compromised their devices and exposed the private keys controlling Step Finance's treasury wallets. On-chain fund movement patterns further confirmed DPRK attribution. This cluster, operating under aliases including SnatchCrypto, CryptoMimic, and BlueNoroff, extracted nearly $200 million in 2025 alone through this same fake-VC-call technique.

Bitrefill disclosed a March 1 breach attributed to the Lazarus/Bluenoroff group based on malware signatures, on-chain tracing, and reused IP and email addresses. The attack originated from a compromised employee laptop, which exposed legacy credentials and allowed the attackers to access production infrastructure, drain hot wallets, and exploit the company's gift card supply chain. Bitrefill stores minimal personal data but approximately 18,520 purchase records were accessed. The company stated it will absorb losses from operational capital.

These two incidents extend a pattern documented throughout 2025: in recent years, 100% of crypto thefts attributed to North Korean actors have relied on social engineering and advanced phishing rather than smart contract exploitation. Observed operational playbooks include fake IT workers, fraudulent job interviews (the Contagious Interview cluster), malicious video calls, and supply-chain attacks. DPRK's TraderTraitor cluster alone extracted roughly $1.85 billion in 2025 through exchange breaches.

→ How Ongoing Security Would Have Mitigated This

The Step Finance loss could have been prevented with the security practices outlined in the 2025 report: use hardware wallets, but critically, do not use your daily-driver laptop when signing transactions. Use an alternative device for signing (e.g., an iPhone or iPad) and do not message, use GitHub, or read emails from there. Keep it isolated. Multi-signature wallet architecture with hardware-isolated signers would have prevented a single device compromise from authorizing treasury transfers.

For Resolv Labs, on-chain minting circuit breakers – a cap on USR minted per transaction or per hour – would have limited extraction to a fraction of the total. Multi-party computation (MPC) or threshold signature schemes would have prevented a single KMS compromise from authorizing mints.

Layer 3: DNS Hijacking and Frontend Compromises

Q1 2026 saw **4 DNS hijacking or frontend compromise incidents** affecting Neutrl, bonk.fun, Compound Finance, and OpenEden. While no confirmed on-chain losses were reported from these incidents (teams responded quickly to take domains offline), the attack vector represents a growing and underappreciated risk.

The Compound Finance incident on March 9 saw the compound.finance domain redirect to a phishing domain (compoond.finance) registered just one day earlier. Compound is one of the most established DeFi protocols, with over $2 billion in TVL and multiple audits. None of that mattered – the attack targeted the web layer, not the blockchain.

Neutrl's DNS provider was compromised via social engineering on March 19, and OpenEden's DNS was hijacked in February, affecting both their main website and app portal.

→ How Ongoing Security Would Have Mitigated This

DNS hijacking exploits the gap between decentralized protocols and centralized web infrastructure. Mitigation requires: using registrars with hardware-key-enforced access (Cloudflare, MarkMonitor, AWS Route53), enabling DNSSEC, deploying ENS-based decentralized frontends as fallbacks, and monitoring for unauthorized DNS record changes. Browser-side defenses such as Blockaid and Pocket Universe caught several of the phishing prompts in the Compound and Neutrl incidents, demonstrating the value of client-side security tooling.

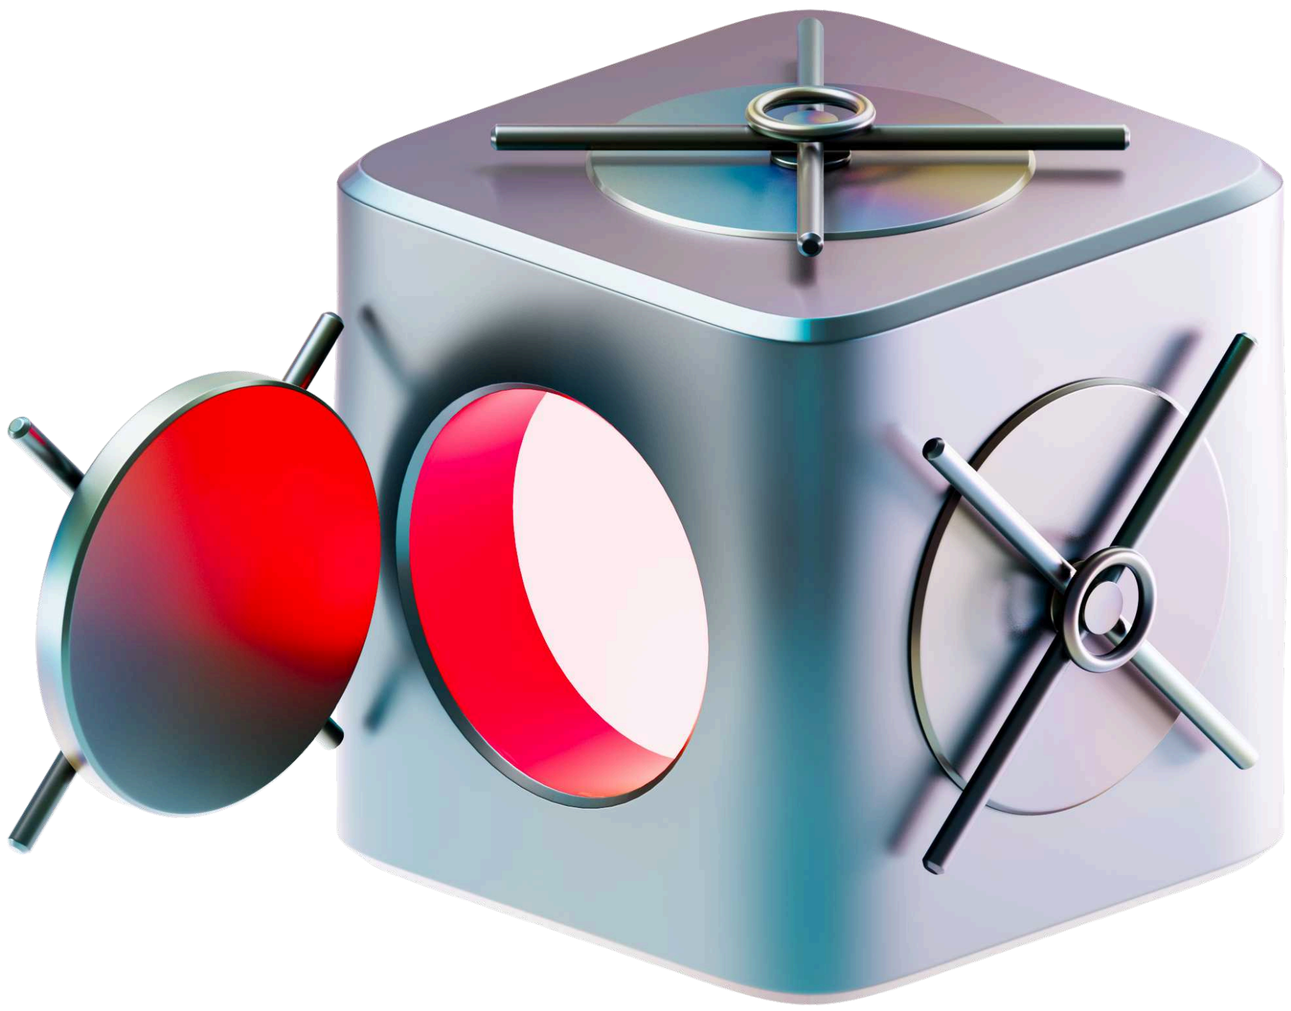

Top 5 Incidents of Q1 2026

1 Social Engineering Attack on Individual Wallet – $282 million

Detail	Value
Date	**January 10, 2026**
Attack vector	**Social engineering (IT support impersonation)**
Assets stolen	**2.05M LTC + 1,459 BTC ($282M)**
Chain	**Bitcoin, Litecoin**
Detection delay	**Hours – discovered after attacker had already swapped to Monero**

This was the most expensive DeFi-adjacent social-engineering attack to date and the single largest loss event of Q1 2026. The attacker posed as IT support for the victim's hardware wallet, using social engineering to extract wallet recovery credentials.

Once obtained, the attacker transferred the full BTC and LTC balances and immediately laundered them through instant exchanges into Monero, a privacy coin that makes tracing nearly impossible. The speed of the Monero conversion left almost no recovery window.

Detection opportunity: Real-time alerting monitoring for outbound transfers exceeding historical wallet patterns could have flagged and potentially delayed the transfer. Trezor-side policies requiring multi-factor confirmation for transfers above a threshold would have added a critical second barrier.

2 Step Finance – $40 Million (DPRK-Attributed)

Detail	Value
Date	**January 31, 2026**
Attack vector	**Fake VC call → device compromise → private key theft**
Threat actor	**DPRK – DangerousPassword / BlueNoroff cluster**
Assets stolen	**261,854 SOL (~$40M from treasury and fee wallets)**
Chain	**Solana**
Detection delay	**Unknown – disclosed after full drainage**

Step Finance, once described as "the front page of Solana," suffered the largest DeFi protocol loss on Solana in Q1 2026. The attack followed the DangerousPassword playbook documented in the 2025 Yearly Security Report: attackers impersonating venture capitalists contacted the Step Finance executive team proposing a collaboration. During a scheduled video call, fabricated "audio issues" prompted the targets to download what was presented as a software update or fix.

The downloaded file deployed malware that compromised the executives' devices, exposing the private keys that controlled Step Finance's treasury and fee wallets on Solana. On-chain fund movement patterns – including rapid bridging and use of addresses linked to prior DPRK cases – confirmed the attribution.

The project subsequently shut down operations entirely, unable to secure funding or an acquisition path after the hack. This incident mirrors the DangerousPassword cluster's 2025 operations, in which the same group targeted CEOs, CFOs, and founders through spearphishing, LinkedIn outreach, and Telegram messages from compromised partner accounts.

Detection opportunity: Multi-signature wallet architecture with hardware-isolated signers would have prevented a single device compromise from authorizing treasury transfers. Endpoint detection and response (EDR) on executive devices, combined with network-level monitoring for known DPRK C2 patterns, could have caught the malware deployment before key extraction. The project should not have stored treasury-controlling keys on devices used for general communication.

3 Truebit Protocol – $26.4 Million

Detail	Value
Date	**January 6, 2026**
Attack vector	**Smart contract vulnerability (integer overflow)**
Assets stolen	**8,535 ETH ($26.44M)**
Chain	**Ethereum**
Detection delay	**Unknown**

The Truebit exploit is a textbook case of legacy code risk. The vulnerable contract was written in Solidity 0.5.x without SafeMath, deployed approximately five years ago. The attacker discovered that the minting function could return a purchase price of zero for an unusually large token buy, effectively allowing free minting of TRU tokens. These were then sold back to the bonding-curve pool for ETH. The TRU token crashed 99.9% following the exploit.

Detection opportunity: Automated invariant monitoring on the bonding curve's token-to-ETH ratio would have detected the anomalous zero-cost mint within the first transaction. Legacy contract deprecation policies – migrating liquidity away from old, unupgradeable contracts – would have eliminated the attack surface entirely.

4 Resolv Labs (USR) – $25 Million

Detail	Value
Date	**March 22, 2026**
Attack vector	**Cloud infrastructure compromise (AWS KMS signing key)**
Assets stolen	**80M unbacked USR → ~$25M extracted**
Chain	**Ethereum**
Detection delay	**Minutes – but attacker extracted $20M before protocol pause**

Resolv Labs suffered one of the most architecturally instructive breaches of Q1. The attacker compromised the protocol's AWS KMS environment to obtain the SERVICE_ROLE signing key, which authorized USR minting. Two transactions minted a combined 80 million USR from $200K in USDC – a 400x leverage exploit.

The attacker converted USR to wstUSR, then swapped into USDC, USDT, and finally ETH, extracting approximately $25 million before the team paused operations. USR depegged to $0.025 on Curve, and the ripple effects forced Euler Labs and Venus Protocol to disable USR collateral functionality.

Resolv had undergone 18 audits. The on-chain smart contracts functioned exactly as designed. The vulnerability was entirely in the off-chain infrastructure that managed the privileged signing key.

Detection opportunity: On-chain monitoring flagging a $100K deposit producing 52M USR (520x the expected ratio) would have triggered an alert on the first transaction. A minting rate limiter would have capped extraction. Multi-party signing for the SERVICE_ROLE rather than a single EOA controlled by one KMS instance would have required the attacker to compromise multiple independent systems.

5 SwapNet – $16.8 Million

Detail	Value
Date	**January 25–26, 2026**
Attack vector	**Smart contract vulnerability (arbitrary call)**
Assets stolen	**$16.8M across Base, Ethereum, Arbitrum, BSC**
Chain	**Multi-chain**
Detection delay	**Hours – Matcha Meta disclosed after funds were drained**

SwapNet, a DEX aggregator integrated into Matcha Meta, was exploited through an arbitrary call vulnerability with insufficient input validation. Because the code was closed-source, the exact flaw was difficult to analyze post-incident.

The attacker abused existing token approval mechanisms to execute transferFrom operations across four chains. Losses were concentrated on Base ($13.37M) and Ethereum ($3.53M). Matcha Meta subsequently removed SwapNet as an available aggregator and disabled the option for users to turn off one-time approvals.

Detection opportunity: The closed-source nature of SwapNet's code is itself a red flag. Open-source contracts allow community review, formal verification, and security researcher access. For aggregators integrating third-party protocols, continuous approval monitoring and allowance-scoping (limiting token approvals to the exact amount needed per transaction) would have restricted the attacker's ability to drain user funds through pre-existing approvals.

"In the current threat landscape, the question is no longer if an organization will be compromised, but when. True security leadership requires moving past the illusion of invulnerability and adopting a mindset of inevitable breach. Our strategy must be to assume the future compromise of our systems and dedicate every effort to deferring that event, delaying the attacker's progress, and aggressively mitigating the eventual outcome.

In 2026, defense-in-depth is not an option—it is a baseline. By layering multiple security controls, prioritizing rigorous risk management, and maintaining battle-ready incident response plans, we move from reactive defense to proactive resilience. Continuous security must be the default state, guided by trusted professionals who understand that a point-in-time audit is merely a starting point, not a finish line."

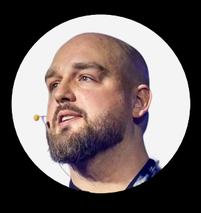

Grzegorz Trawiński

Director of Offensive Security Services at Hacken

INDUSTRY PERSPECTIVE

Social Engineering and Infrastructure Attacks VS On-Chain Monitoring

What Drift, Venus, and Q1's broader patterns reveal about where on-chain monitoring ends — insights from Allium.

Q1's largest incident generated extensive on-chain data during execution and almost none before it. $285M left Drift Protocol's vaults in one transaction on April 1. Within an hour, Solana DEX volume spiked to $735M as stolen assets moved through Orca, Raydium, Meteora, and Jupiter. We traced the full fund flow as it happened: $153M in JLP tokens redeemed instantly, $184M ultimately converted to SOL, $101M to stablecoins, spread across three wallets. Every step was legible on-chain in real time. The six months of preparation that preceded it were not.

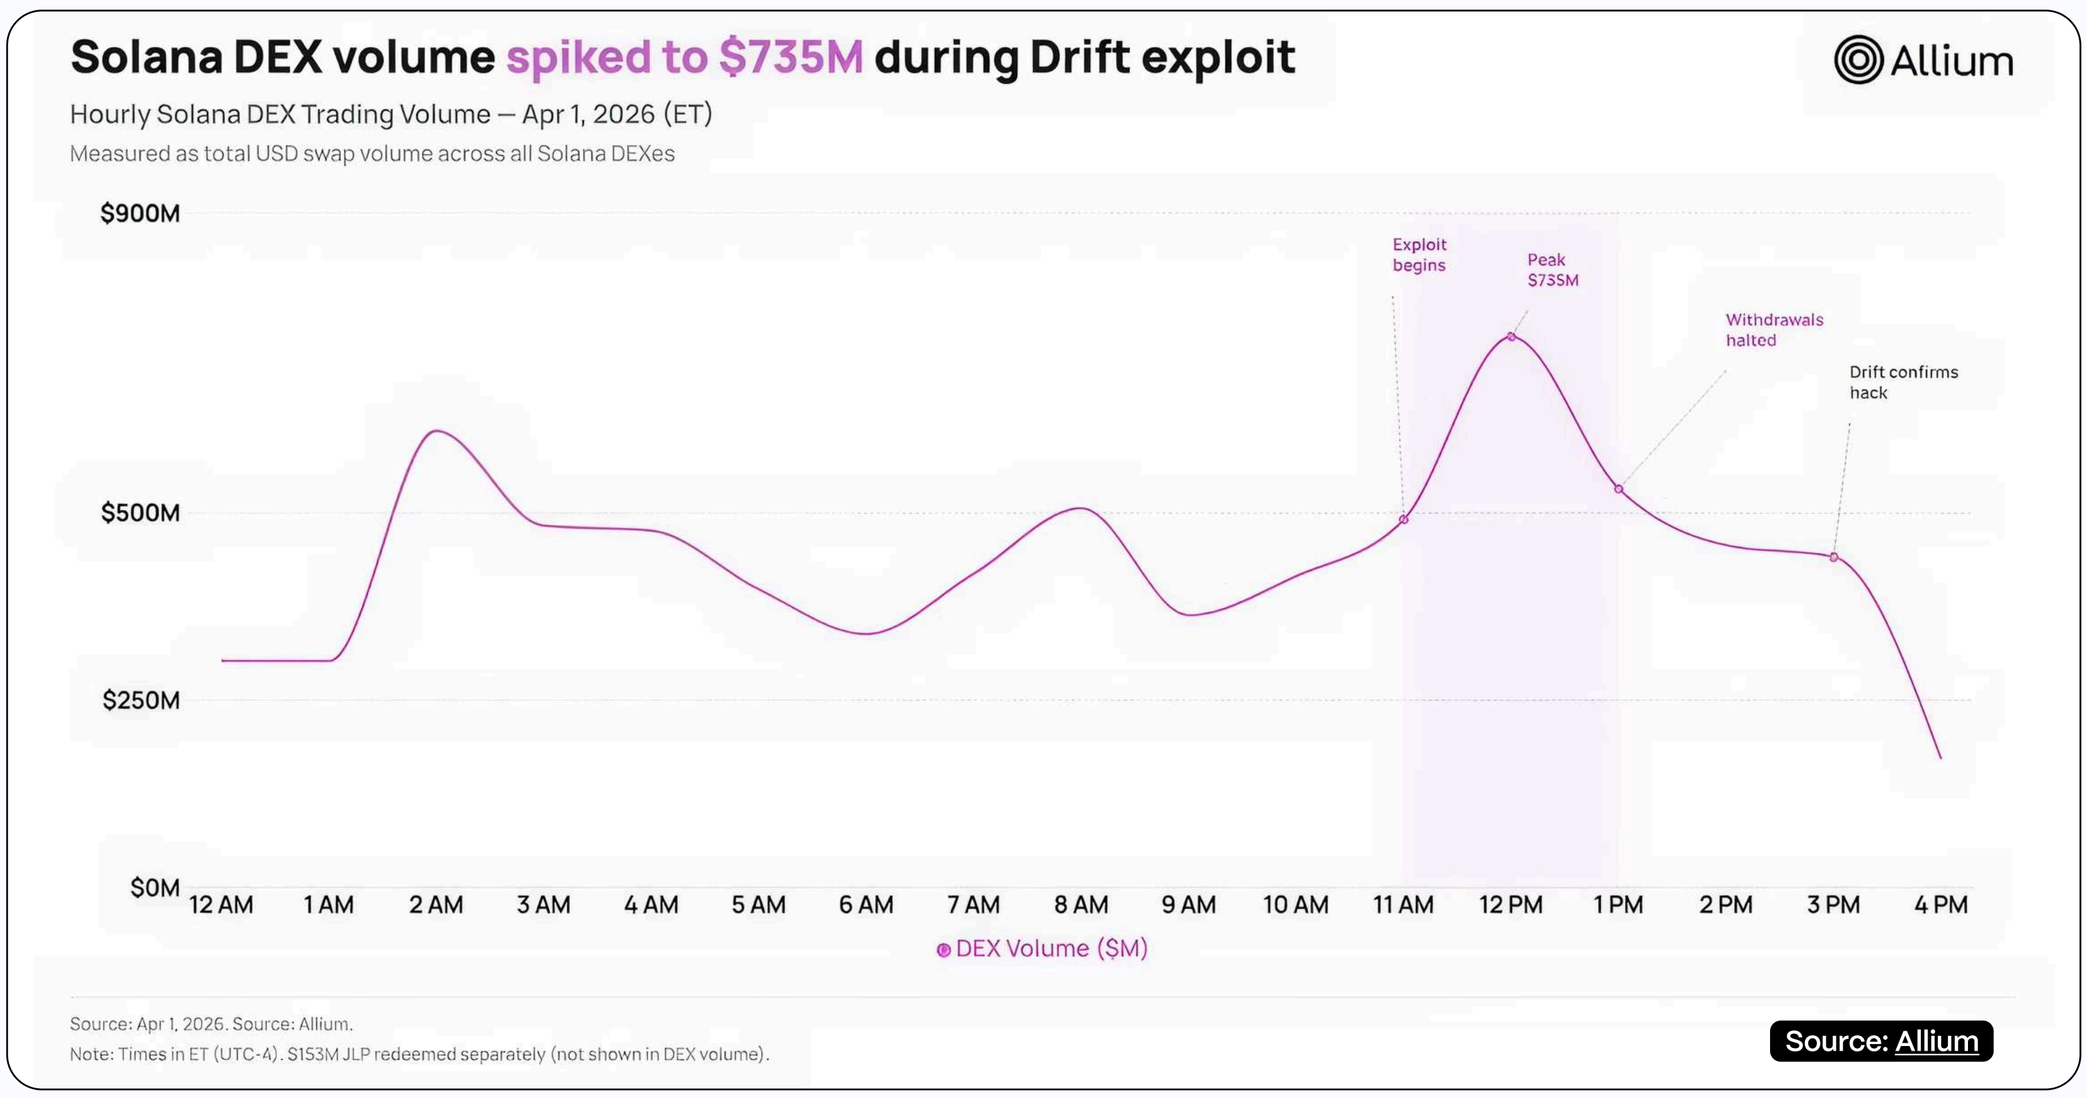

That's because the attack was social engineering from start to finish. DPRK-linked actors spent months building a fake trading identity, earning the trust of Drift contributors, and eventually getting them to pre-sign authorizations that enabled the drain. None of that preparation touched the chain. The on-chain signal arrived at execution, not before.

Not every Q1 incident looked like that. The Venus Protocol attack in March is the kind that does leave an advance footprint. The attacker spent nine months accumulating THE tokens until they controlled roughly 84% of Venus's supply cap for that asset. That concentration was visible in the on-chain data throughout. So was the anomalous direct contract transfer used to exploit the exchange rate. Monitoring for supply concentration or unusual approval patterns at that threshold would have flagged it well in advance.

The pattern across Q1: social engineering and infrastructure attacks show up on-chain only at extraction. Contract-level exploits usually leave positioning signals 24 to 72 hours ahead, sometimes longer. For protocol and contract risk, on-chain monitoring can catch anomalies early. For attacks targeting the human and operational layer, the data arrives too late — which means the controls that matter are off-chain ones most Web3 teams haven't invested in yet.

Allium

Key Takeaways and Recommendations

Q1 2026 reinforces three structural observations about Web3 security:

1 Audits are necessary but not sufficient

Six audited protocols were exploited in Q1 2026, including Resolv Labs with 18 audits and Venus Protocol with five audit firms. Traditional audits review code at a point in time; they do not monitor live systems, detect operational compromises, or enforce runtime invariants. The industry needs a shift from audit-and-ship to continuous security monitoring as a baseline expectation.

2 DPRK remains the dominant threat actor for operational attacks

Step Finance and Bitrefill were both attributed to North Korean groups using documented playbooks – fake VC calls, compromised employee endpoints, and supply-chain attacks. In 2025, DPRK actors were responsible for approximately $2.04 billion in crypto theft (52% of all losses). The techniques are not novel; they are repeatable, scalable, and specifically targeting the crypto industry's weak operational security posture. Every team handling significant treasury value should assume they are a target.

Recommended Actions for Protocol Teams:

- Deploy real-time on-chain monitoring with automated pause mechanisms for anomalous transactions
- Implement minting rate limiters and circuit breakers for any privileged on-chain role
- Move treasury and minting keys to MPC or threshold signature schemes; eliminate single-EOA control
- Mandate hardware wallets on isolated devices for all signing operations; never use daily-driver laptops
- Treat any unsolicited video call, job offer, or VC outreach as a potential DPRK social-engineering attempt until verified through independent channels
- Audit legacy contracts or deprecate them; Solidity versions before 0.8.x carry known overflow risks
- Enforce one-time token approvals in all user-facing interfaces; unlimited approvals remain a systemic risk
- Secure DNS infrastructure with hardware-key access, DNSSEC, and deploy ENS-based fallback frontends
- Maintain an active bug bounty program scoped to cover the vulnerabilities actually being exploited (oracle logic, access control, key management), not just smart contract code

II Smart Contract Security Landscape

This section breaks down Q1 security assessments by type, highlights the most common vulnerability classes across audits, and maps which sectors carried the highest risk exposure.

Audit Distribution and Findings by Sector

Sector	Share of Audits	C+H per Audit	C+H+M per Audit	Findings per Audit
DeFi	41.2%	1.71	5.36	20.4
RWA / TradFi	20.6%	1.00	4.86	22.4
Token / Stablecoin	17.6%	0.33	1.52	8.3
Infrastructure	20.6%	0.43	1.14	6.3

* **C+H** = Critical + High severity findings; **C+H+M** = Critical + High + Medium severity findings

Smart Contract Security Trends

1

Fix reviews can surface new issues.

1 out of 3 fix-review engagements introduced a new Medium-severity vulnerability in the remediation code itself.

2

Scope expansion reveals hidden risk.

Auditing additional modules of a previously-audited protocol consistently produces new findings at higher severity than the original engagement.

Sector Deep Dive

DeFi — Highest Risk Profile

DeFi is the riskiest sector with 1.71 Critical+High findings per audit, contributing 66.7% of all C+H findings.

The sector's attack surface is driven by:

- Complex economic settlement logic (prediction markets, bond vaults, DEX hooks) introducing multi-step accounting where a single miscalculation leads to insolvency
- Oracle dependencies for pricing, resolution, and fee computation
- Novel protocol architectures (Uniswap v4 hooks, Algebra plugins, OFT cross-chain tokens) where security patterns are not yet battle-tested
- Multi-token interactions requiring consistent handling of decimals, fee-on-transfer behavior, and rebasing

---

RWA / TradFi — Highest Volume, Compliance-Heavy

RWA audits had the highest average finding count (22.4 per audit), driven by large scopes combining token logic with compliance mechanisms. The sector's distinctive vulnerability patterns include:

- **Blacklist/pause enforcement gaps** — The most recurring theme across RWA audits; circuit breaker bypass, role pause not enforced for specific operations, blacklisted users escaping restrictions
- **Oracle/NAV pricing complexity** — The MMF protocol alone generated 9 Medium pricing-related findings due to timezone-aware trading windows, DST calculations, and dual-price caching
- **Role-based access complexity** — RWA tokens typically implement 5–10 distinct roles (admin, custodian, oracle, pauser, blacklister, investor groups), creating wide authorization attack surfaces

---

Token / Stablecoin — Lower Risk Profile

Standalone token contracts showed the lowest severity density (0.33 C+H per audit).
Most findings were Low or Informational, focusing on:

- Missing input validation
- Code quality issues
- Standard deviations (e.g., EIP-20 non-compliance)

The exceptions were tokens with DEX integration or wrapping mechanics,
which introduced higher-severity logic errors.

Infrastructure — Emerging Attack Surfaces

Infrastructure projects (paymasters, governance, vesting) showed moderate risk overall but contained critical findings in the paymasters audits. ERC-4337 account abstraction infrastructure introduces novel and severe attack vectors when paymaster validation logic is incomplete.

INDUSTRY PERSPECTIVE

Centrifuge on Building Auditable RWA Infrastructure

Why separating immutable core infrastructure from configurable product layers helps manage the security challenges created by frequent integrations and code updates.

RWA protocols go through more code changes and integration updates than most DeFi projects — how does that affect your approach to code reviews and assessments?

> This is why Centrifuge separates the immutable core from the configurable product layer. Core contracts (accounting, settlement, cross-chain messaging, share tokens) are immutable and identical across all 9 supported chains. The 27 security reviews and years of production history carry forward to every new deployment. What changes (vault configs, transfer restrictions, fees, pricing) sits in a constrained layer with a much smaller audit surface.
>
> Our review approach layers multiple mechanisms:
>
> – Deep, ongoing security researcher partnerships (Spearbit, Sherlock, BurraSec, yAudit) rather than one-off engagements. Auditors who understand the full architecture catch things that fresh eyes starting from scratch will miss.
>
> – Invariant verification that continuously tests system properties across randomized execution paths, catching classes of bugs manual review misses: edge cases in accounting, reentrancy in cross-chain flows, state inconsistencies across epoch boundaries.
>
> – Architectural simplification. Each protocol iteration reduces complexity. Fewer code paths means fewer places for bugs to hide. Simpler systems are more auditable and more secure.
>
> – Audit competitions that open the codebase to adversarial review from the broader security community, supplementing private engagements.
>
> – Smaller audit scope per deployment. Because the immutable core is already covered by 27 reviews, new deployments only need to audit their custom modules. The review surface shrinks from "the entire vault" to "the specific configuration this pool uses."
>
> Centrifuge

New Smart Contract Development Standards

Hacken's Q1 2026 audits covered several emerging or recently adopted standards. Below is an assessment of each, including the most common vulnerability types found.

ERC-4337 — Account Abstraction (Paymasters)

Most common vulnerability type: Access Control & Authorization bypass

Standard-specific risk: ERC-4337 paymaster contracts present a novel attack surface centered on gas parameter manipulation and validation bypass. Paymasters must validate gas parameters holistically. Incomplete parameter coverage in the signature commitment or whitelist validation creates deposit-drain attacks that are unique to the ERC-4337 execution model and have no equivalent in traditional EOA transaction flows.

Uniswap v4 Hooks Architecture

Most common vulnerability type: Logic errors in hook delta calculations

Standard-specific risk: Uniswap v4 introduced a hook-based plugin architecture that allows custom logic to execute before and after swaps. This creates a new class of delta accounting vulnerabilities. The v4 hook model requires developers to correctly construct return values that adjust the pool's internal accounting. The mapping between swap direction (exact-input vs exact-output), currency side (currency0 vs currency1), and delta sign conventions is error-prone. All three significant findings in the audit were directly caused by incorrect delta construction — a vulnerability class that simply does not exist in v3 or earlier versions.

DEX Plugin Architecture

Most common vulnerability type: Logic & Accounting Errors

Standard-specific risk: The DEX plugin model — enabling fungible ERC-20 LP tokens for concentrated liquidity positions — had the highest Critical finding density of any Q1 audit (18.2%). Extending concentrated liquidity DEX protocols with plugin architecture introduces compound complexity — the plugin must correctly interface with the base DEX's position management, fee accrual, price calculation, and NFT lifecycle. The E-STORM audit demonstrates that novel extensibility mechanisms applied to already-complex protocols create multiplicative risk, not additive.

RWA Token Compliance Frameworks

RWA tokenization is an emerging sector without a single dominant standard. Projects implement bespoke compliance layers combining:

- Blacklist management (Chainalysis integration, manual lists)
- Pause/circuit breaker mechanisms
- Multi-role access control (investor groups, custodians, oracles, fund admins)
- NAV-based pricing with trading hour awareness
- Transfer restriction matrices

Top 5 RWA Vulnerability Types

Rank	Vulnerability Type	% of all RWA C+H+M	Example
1	Blacklist/Pause Bypass	23.5%	Circuit breaker bypass in MMF; blacklist not checked in mint ops
2	Oracle/NAV Pricing Issues	23.5%	Price cache desync, DST errors, missing freshness checks
3	Role Management Gaps	20.6%	Unintended DEFAULT_ADMIN_ROLE renouncement, split ownership
4	Fund Loss/Locking	14.7%	Expired tokens trapped, direct transfer loss in custody vaults
5	State/Lifecycle Issues	11.8%	Frozen balance not enforced on transfers, config deadlocks

Standard-specific risk: RWA tokens attempt to replicate traditional financial controls (compliance screening, trading hours, settlement flows) on-chain. The complexity of these compliance layers — often exceeding the complexity of the token logic itself — introduces a large surface for enforcement gaps where restrictions are correctly defined but incompletely applied. The most characteristic RWA vulnerability is: "Blacklist/pause mechanism exists but is not enforced for a specific code path."

ERC-4626 — Tokenized Vaults

Most common vulnerability type: Oracle & Pricing Issues

Standard-specific risk: ERC-4626 vaults that support multi-asset deposits require oracle-based valuation for share price computation. The most common error pattern is incorrect price conversion formula when normalizing between different token decimals and oracle price feed decimals. One audit's Critical finding demonstrated an algebraically incorrect conversion that would misvalue all non-reference assets.

Combined notable findings: Oracle formula errors, permissionless yield harvesting, cooldown bypass, externalAssets tracking divergence.

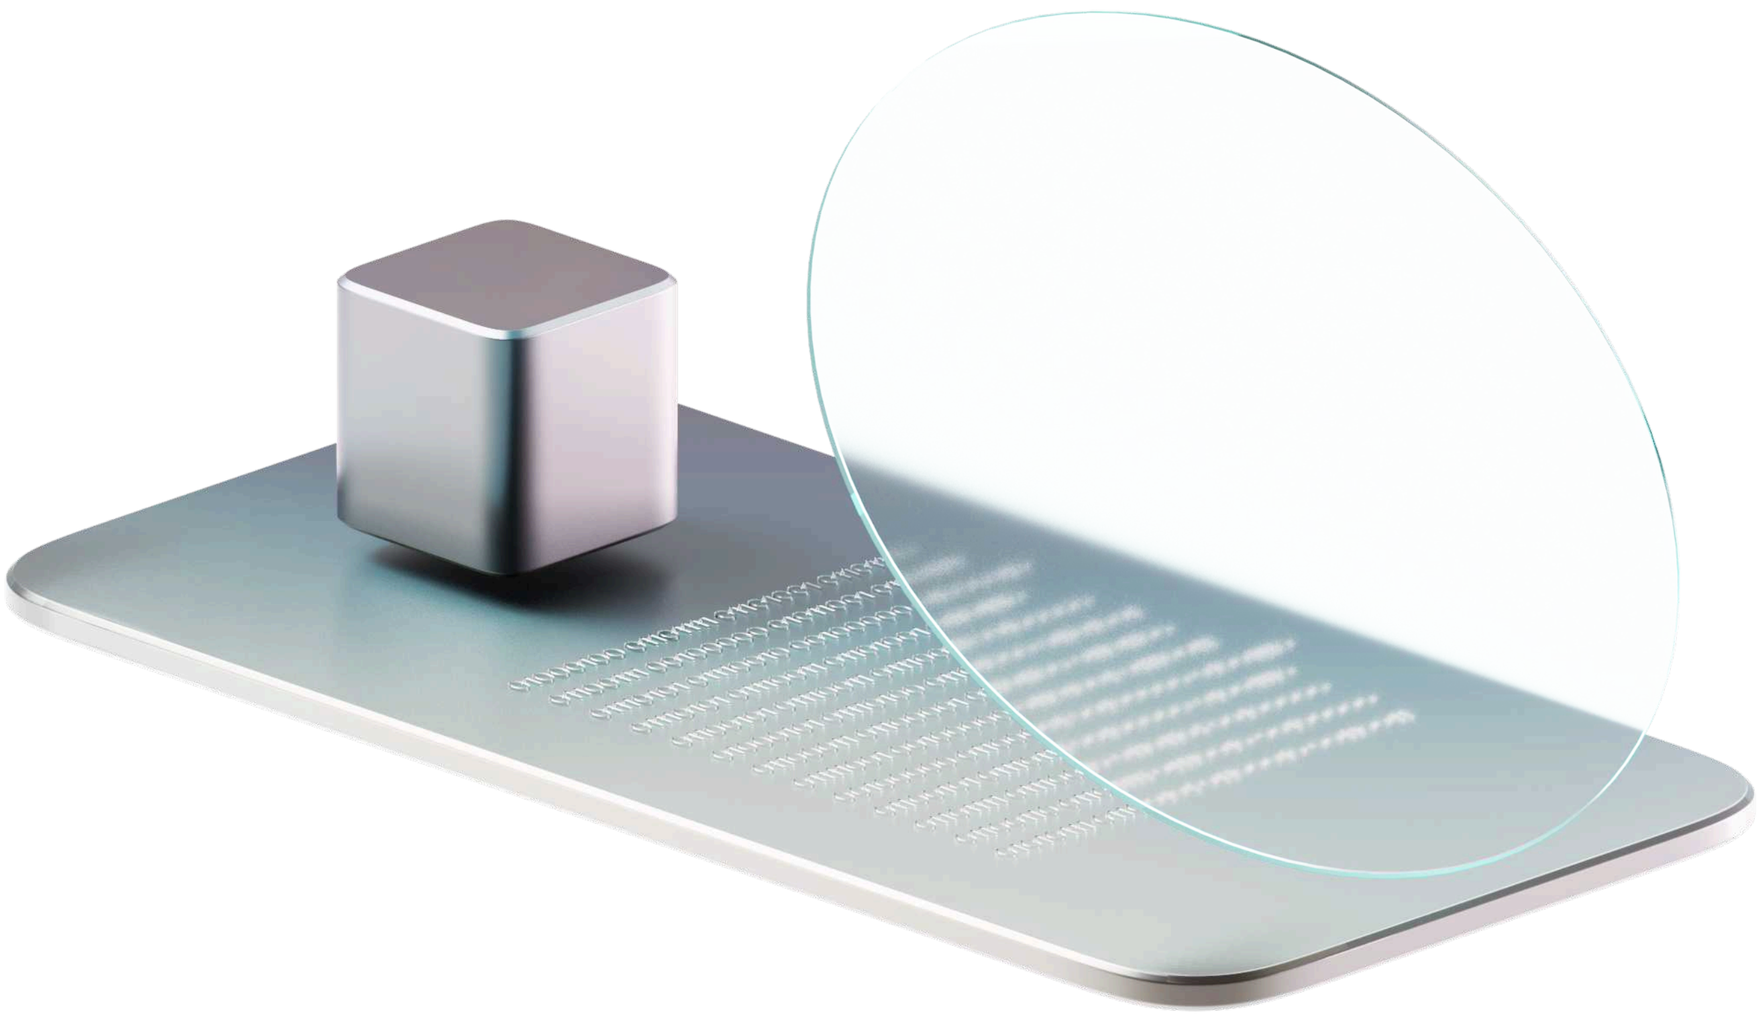

Hacken Audit Quality and Coverage

47.1% of all findings were Critical+High severity. Nearly half of all audited protocols contained vulnerabilities that could result in partial or total loss of user funds if deployed unaudited.

Vulnerability Class Coverage — Q1 2026

Vulnerability Class	% of All C+H+M	Coverage Assessment
Logic & Accounting	27.0%	Consistently detected across all sectors
Access Control	22.2%	Strong detection including subtle bypass patterns
Oracle & Pricing	13.5%	Advanced detection of formula-level errors
State Management	11.9%	Lifecycle and configuration issues identified
Fund Loss/Locking	9.5%	Direct financial impact findings
Token Safety	6.3%	Fee-on-transfer, unsafe transfer patterns
Cryptographic	4.8%	Signature malleability, EIP-712 gaps
Reentrancy	2.4%	Lower prevalence reflects industry maturation

Remediation Effectiveness

Repeated engagement data demonstrates that the Hacken audit lifecycle (finding → remediation → re-audit) produces measurable security improvement:

- 100% Critical elimination on re-audit engagements
- Severity downshift from C+H to M+L on subsequent reviews
- New finding identification in fix reviews (new Medium and High issues in fixes) validates that remediation review is not perfunctory

Stablecoin & Digital Asset Security

The stablecoin market has undergone significant structural evolution heading into 2026:

Stablecoin Market Map — Q1 2026

Issuer	Token	Peg	Est. Market Cap	Region
Tether	USDT	USD	~$140B+	BVI / El Salvador
Circle	USDC	USD	~$55B+	US (MiCA-compliant EU)
Sky (ex-MakerDAO)	USDS/DAI	USD	~$8B+	Decentralized
Ethena	USDe	USD (synthetic)	~$5B+	Cayman Islands
First Digital	First Digital	USD	~$2B+	Hong Kong
PayPal	PYUSD	USD	~$1B+	US
Ripple	RLUSD	USD	~$0.5B+	US / Singapore
Circle	EURC	EUR	Growing	EU (MiCA-licensed)
SG FORGE	EURCV	EUR	Institutional	France (MiCA-licensed)
Quantoz Payments	EURQ/USDQ	EUR/USD	New entrant	Netherlands (MiCA)
Banking Circle	EURI	EUR	New entrant	Luxembourg (MiCA)
Membrane Finance	EUROe	EUR	New entrant	Finland (MiCA)

Key Trends

- **USD dominance persists** but 5+ MiCA-licensed EUR stablecoins are now live.
- **Bank-issued stablecoins are emerging**, SocGen's EURCV, JPMorgan's JPM Coin (institutional), and multiple European banking institutions entering under MiCA's e-money token (EMT) framework.
- **Synthetic / yield-bearing stablecoins** (Ethena's USDe, Usual's USD0, Mountain Protocol's USDM) blur the line between stablecoin and DeFi yield product. A smart contract audit alone covers only 20–30% of the total attack surface.
- **Regional expansion** across Middle East (AE Coin in UAE), Latin America (multiple USD-pegged tokens for remittance corridors), Southeast Asia (StraitsX XSGD, XIDR), and Africa (Pan-African stablecoin initiatives).

Stablecoin Security Architecture

Stablecoin security does not begin on-chain. The smart contract is only one component of a six-layer security architecture that spans Web2 infrastructure, Web2.5 operational and transparency layers, and on-chain execution.

Six-Layer Stablecoin Security Model

Foundation → increasingly on-chain

Layer 6 — Cross-Chain & Bridge Security
Bridge validators, message verification, cross-chain supply accounting, relayer infrastructure

Layer 5 — On-Chain — Smart Contract Security
Token logic, access control, mint/burn, pause/blacklist, upgradeability, oracle integration

Layer 4 — Web2.5 — Operational Security
Key management, multisig governance, deployment pipelines, admin key rotation, monitoring & alerting

Layer 3 — Web2.5 — Proof of Reserves & Transparency
Reserve attestation, on-chain oracle feeds, independent auditor verification, real-time dashboards

Layer 2 — Web2 — Compliance & Legal Infrastructure
KYC/AML systems, sanctions screening, regulatory licensing, legal entity structure

Layer 1 — Web2 — Reserve & Custody Infrastructure
Banking relationships, custodial accounts, asset composition, redemption guarantees, counterparty risk

Layer 1 – Reserve & Custody Infrastructure (Web2)

What it is: The foundational layer where the actual value backing the stablecoin exists — bank accounts holding USD, T-bills, money market fund shares, gold bars, or other reserve assets.

Components:

- **Banking relationships:** Primary and backup banking partners; single-bank concentration risk.
- **Asset composition:** Cash vs. T-bills vs. commercial paper vs. repo agreements. Each carries different liquidity and credit risk profiles.
- **Custodial architecture:** Segregated vs. commingled accounts; custodian operational risk; insolvency protection.
- **Redemption guarantees:** SLAs for mint-to-settlement and redeem-to-settlement times; queue management under stress.
- **Counterparty risk:** Bank failure (cf. SVB March 2023 depegging USDC), custodian compromise, broker-dealer default.
- **Insurance coverage:** Scope and limits of asset insurance; gaps in crypto-specific coverage.

Risk example: In March 2023, USDC depegged to $0.88 when Silicon Valley Bank (holding $3.3B of Circle's reserves) failed. No smart contract vulnerability was involved – this was pure Layer 1 banking infrastructure risk.

Applicable Hacken services:

Proof of Reserves Audit

Independent verification of reserve composition, sufficiency, and accessibility. Validates that on-chain supply matches off-chain reserve value.

CCSS (CryptoCurrency Security Standard) Assessment

Evaluates custodial security practices, key management, and operational controls of reserve-holding entities.

Layer 2 – Compliance & Legal Infrastructure (Web2)

What it is: The regulatory and legal framework that determines whether a stablecoin can legally operate, who can hold it, and what obligations the issuer bears.

Components:

- **Licensing & authorization:** EMI license (EU/MiCA), Money Transmitter License (US states), MAS license (Singapore), VARA license (Dubai), etc.
- **KYC/AML systems:** Identity verification for minters/redeemers; transaction monitoring for suspicious activity; SAR filing obligations.
- **Sanctions screening:** OFAC (US), EU Sanctions List, UN Sanctions. Real-time screening of addresses against sanctions lists (e.g., Chainalysis integration as seen in one of our audits).
- **Legal entity structure:** Issuer domicile, reserve trust structure, bankruptcy remoteness, wind-down procedures.
- **Jurisdiction-specific obligations:** MiCA requires a detailed white paper, reserve management policy, and complaint handling procedure. US frameworks vary by state.
- **Travel Rule compliance:** FATF Travel Rule requiring originator/beneficiary information for transfers above thresholds.

Risk example: In Q1 2026, one of our audits found "Missing Chainalysis and Compliance Checks in authorizeAndBuyWithStableTokenFromWithTo" (High) — the stablecoin buy/sell path bypassed sanctions screening entirely. This is a Layer 2 failure manifested in Layer 5 (smart contract) code.

Applicable Hacken services:

Compliance Assessment

Review of KYC/AML integration completeness, sanctions screening coverage, and regulatory obligation mapping.

Smart Contract Audit (compliance logic)

Verification that on-chain blacklist, pause, and transfer restriction mechanisms correctly enforce off-chain compliance requirements (Q1 data: this was the #2 most common vulnerability type in RWA audits).

Penetration Testing (Web2 infrastructure)

Testing of the compliance platform, admin dashboards, and KYC/AML provider integrations for unauthorized access.

Layer 3 – Proof of Reserves & Transparency (Web2.5)

What it is: The bridge between off-chain reserves and on-chain token supply — mechanisms that allow users, regulators, and the market to verify that the stablecoin is fully backed.

Components:

- **Attestation reports:** Periodic third-party attestations (e.g., monthly accounting firm reports) confirming reserve ≥ supply.
- **On-chain PoR oracle feeds:** Real-time or near-real-time reserve data published to the blockchain via oracle contracts (Chainlink PoR, Redstone, custom feeds).
- **Supply reconciliation:** Automated or manual comparison of on-chain totalSupply() against off-chain reserve balances across all chains of deployment.
- **Cross-chain supply tracking:** For multi-chain stablecoins, ensuring that the sum of supply across all chains matches the total reserve — directly related to one of our findings ("Hub Chain Supply Reduction After OFT Transfers Lead to supply() DoS").
- **Transparency dashboards:** Public-facing reserve composition displays, minting/burning activity, and redemption queue status.

Risk example: FTX's collapse revealed that attestation reports can provide a false sense of security if the attestor lacks scope or independence. A PoR attestation that only confirms "assets exist at a point in time" without verifying "assets are unencumbered, accessible, and not rehypothecated" leaves critical gaps.

Applicable Hacken services:

Proof of Reserves (PoR) Audit

End-to-end verification: on-chain supply enumeration across all deployment chains, off-chain reserve balance confirmation, reconciliation methodology review, and attestation of sufficiency.

INDUSTRY PERSPECTIVE

Proof of Reserves Audit as operational control at MEXC

MEXC publishes Merkle-tree-based reserve reports and conducts monthly independent Hacken audits — but sees that as the baseline, not the finish line.

> Proof of reserves started as a voluntary trust signal — but regulators across MiCA, VARA, and US frameworks are increasingly treating reserve transparency as a compliance requirement. How is your exchange adapting to that shift?
>
> At MEXC, we think the direction of travel is clear: across MiCA, VARA and the emerging U.S. stablecoin framework, reserve transparency is moving from an occasional trust signal to part of an ongoing, auditable safeguarding regime, with much stronger expectations around segregation, reserve backing, reconciliation and independent reporting.
>
> We are adapting by treating PoR as a mechanism for operational control and transparency disclosure: we already publish Merkle-tree-based reserve reports, enable users to verify their inclusion in reported liabilities, and conduct monthly independent Hacken audits, but we see that as the baseline, not the finish line. To satisfy regulators and reassure users, PoR has to show more than assets on a page; it needs to evidence the liabilities of those assets back, the segregation and reconciliation processes behind them, the independence of the assurance, and the governance and response framework that protects client assets under stress.
>
> MEXC

Layer 4 – Operational Security (Web2.5)

What it is: The operational practices and infrastructure that govern how the stablecoin is administered day-to-day — key management, deployment practices, incident response, and administrative access.

Components:

- **Key management:** Private key storage for admin/minter/pauser/blacklister roles; HSM vs. multisig vs. MPC wallet architecture.
- **Multisig governance:** Configuration of admin multisigs (threshold, key holder identity, time-locks). The Q1 2025 incident data shows that access control exploits (including multisig compromises) caused 80% of all Web3 losses ($1.6B+).
- **Deployment & upgrade pipelines:** Secure CI/CD for contract deployments; upgrade authorization flows; timelock enforcement on critical changes.

- **Admin key rotation:** Procedures and frequency for rotating admin credentials; revocation processes for compromised keys.
- **Monitoring & alerting:** On-chain event monitoring for unauthorized mints, unusual transfer patterns, blacklist additions; off-chain infrastructure monitoring.
- **Incident response:** Defined playbooks for depeg events, oracle failures, bridge exploits, and smart contract vulnerabilities.

Risk examples from Q1 2026 audits:

- **"Split Ownership Between Ownable and AccessControl Allows Deployer to Drain Paymaster Funds" (Medium)** — dual authority systems in operational governance create privilege confusion.
- **"Unintended Default_Admin_Role Renouncement During Role Transfer" (Medium)** — operational key transfer procedures can permanently destroy administrative access.
- **"Users Can Renounce BLACKLISTED_ROLE and Bypass Administrative Restrictions" (High)** — operational assumptions about immutable compliance status violated by on-chain role mechanics.

Applicable Hacken services:

CCSS Assessment

Comprehensive evaluation of key management, multisig configuration, operational procedures, and access control hygiene.

Penetration Testing

Testing of admin dashboards, deployment infrastructure, key management systems, and API endpoints.

Smart Contract Audit

Verification of on-chain admin role architecture. Q1 data shows Access Control issues are the #2 vulnerability type (22.2% of C+H+M findings).

Layer 5 – Smart Contract Security (On-Chain)

What it is: The on-chain token contract and associated infrastructure — the part traditionally covered by smart contract audits. While critical, it represents only one layer of the full security architecture.

Components:

- **ERC-20/BEP-20 token logic:** Transfer, mint, burn, approve, permit functions.
- **Access control:** Role-based permission systems (admin, minter, pauser, blacklister, upgrader).
- **Compliance mechanisms:** Blacklist (address-level transfer blocking), pause (global emergency stop), freeze (per-address balance freezing), transfer restrictions (inter-role transfer matrix).
- **Mint/burn controls:** Per-minter quotas, rate limiting (per-block caps), approval workflows.
- **NAV/price integration:** Oracle-based pricing for subscription/redemption (MMF pattern), deviation thresholds, staleness checks.
- **Upgradeability:** UUPS or Transparent proxy patterns; upgrade authorization; storage layout safety.

Risk examples from Q1 2026 audits:

From Hacken's stablecoin and digital asset audits in Q1 (RWA + Token sectors), the top vulnerability types at the smart contract layer were:

Rank	Vulnerability Type	% of SC Findings
1	**Blacklist/Pause/Compliance Bypass**	**25.6%**
2	**Access Control & Role Management**	**20.9%**
3	**Oracle & NAV Pricing Issues**	**18.6%**
4	**Logic/Accounting Errors**	**14.0%**
5	**State Management & Configuration**	**11.6%**
6	**Signature & Cryptographic**	**4.7%**
7	**Unsafe Token Operations**	**4.7%**

The distinctive stablecoin vulnerability pattern: Unlike DeFi protocols where logic/accounting errors dominate (27% of all C+H+M), stablecoin contracts are most vulnerable to compliance enforcement gaps — situations where blacklist, pause, or transfer restriction mechanisms exist but are incompletely applied across all code paths.

This pattern was observed across multiple Q1 audits, including:

- **"Stablecoin Buy/Sell Operations Bypass Circuit Breaker Pause Mechanism" (High)**
- **"Role Pause Not Enforced for Token Recipients in Mint and Transfer Operations" (High)**
- **"Missing Chainalysis and Compliance Checks in authorizeAndBuyWithStableTokenFromWithTo" (High)**
- **"Inconsistent Blacklist Enforcement Between approve and transferFrom" (Medium)**
- **"Users Can Renounce BLACKLISTED_ROLE and Bypass Administrative Restrictions" (High)**
- **"Default Admin Can Assign Blacklisted Role Without Enforcing Blacklist Activation Constraints" (Medium)**
- **"Frozen Balance is not Enforced on Transfers" (High)**
- **"Unintended Default_Admin_Role Renouncement During Role Transfer" (Medium)**

Applicable Hacken services:

Smart Contract Audit (SCA)

Comprehensive code review covering all vulnerability categories; the core Hacken engagement for stablecoin issuers.

dApp Audit

Testing of frontend and middleware that interacts with the token contract — APIs, admin panels, minting dashboards.

Layer 6 – Cross-Chain & Bridge Security

What it is: For multi-chain stablecoins (which is now the norm for any stablecoin targeting broad adoption), the bridge and cross-chain messaging infrastructure represents a critical — and historically exploited — attack surface.

Components:

- **Bridge protocol security:** The bridge (LayerZero, Wormhole, Axelar, native bridges like Optimism's L1↔L2 bridge) that enables cross-chain minting and burning.
- **Cross-chain supply accounting:** Ensuring that the total circulating supply across all chains equals the total reserves. Any desynchronization means either under-collateralization or fund locking.
- **Message verification:** Validation of cross-chain messages (minting proofs, burn confirmations) to prevent unauthorized minting on destination chains.
- **Relayer/validator infrastructure:** The off-chain relayer networks that transmit cross-chain messages — their liveness, censorship resistance, and compromise vectors.
- **Upgrade compatibility:** Ensuring that contract upgrades on one chain remain compatible with in-flight cross-chain messages from other chains.

Risk examples from Q1 2026 audits:

- **"Hub Chain OverLayer Supply Reduction After OFT Transfers Lead to supply() DoS" (High)** — LayerZero OFT transfers reduced hub-chain supply, causing supply()-dependent collateral functions to revert.
- **"Cross-Chain Message Loss During Contract Upgrade Due to Payload Format Incompatibility" (Low)** — upgrade on one chain broke compatibility with in-flight messages.
- **"Strict Payload Length Validation May Reject In-Flight Messages During Upgrade" (Low)** — another cross-chain upgrade hazard.
- **Bridge-controlled mint/burn via Optimism's native bridge** — the bridge is the single point of supply control.

Bridge exploits remain the highest-impact attack vector in Web3. Historical data: Ronin ($625M, 2022), Wormhole ($325M, 2022), Nomad ($190M, 2022). For stablecoins operating across 5–10+ chains, bridge security is existential.

Applicable Hacken services:

Smart Contract Audit (bridge contracts)

Audit of lock/mint, burn/release, and message verification contracts on each chain.

Cross-Chain Security Assessment

Holistic review of cross-chain supply invariants, message verification logic, relayer security, and upgrade compatibility.

Infrastructure Penetration Testing

Testing of relayer nodes, validator sets, and oracle infrastructure that support bridge operations.

Continuous Monitoring

Real-time monitoring of cross-chain supply totals and alerting on supply imbalances.

Q1 2026 Evidence: What Our Audits Reveal

Compliance Bypass — The Signature Finding Pattern

The single most distinctive finding pattern in stablecoin audits is compliance mechanisms that exist in contract code but are not enforced across all execution paths. This pattern was found in multiple stablecoin audits (38.5%).

Pattern: The contract defines a blacklist (or pause, or freeze) mechanism. The transfer() function checks it. But mint(), approve(), transferFrom(), permit(), or other purchase functions do NOT check it — creating a bypass with legal liability implications.

- Stablecoin buy/sell operations bypassing circuit breaker pause mechanisms
- Role pause not enforced for token recipients in mint and transfer operations
- Missing sanctions screening in specific purchase/transfer code paths
- Inconsistent blacklist enforcement between approve and transferFrom
- Users able to renounce BLACKLISTED_ROLE and bypass administrative restrictions
- Frozen balance not enforced on transfers
- Unintended DEFAULT_ADMIN_ROLE renouncement during role transfer

Recommendation: Every stablecoin audit should include a compliance enforcement matrix test — systematically verifying that every restriction (blacklist, pause, freeze, transfer limit) is enforced on every state-changing function that touches balances or allowances.

Oracle & NAV Pricing: The Web2.5 Vulnerability

For stablecoins with on-chain pricing (NAV-based subscription, collateral valuation), oracle security bridges Web2 financial data and Web3 settlement. One regulated money market fund token had multiple Medium oracle/pricing findings, including DST timezone calculation errors, price cache desynchronization, and missing freshness validation — all failures at the Web2.5 boundary where real-world market data meets on-chain execution. A separate cross-chain token had an oracle-dependent supply DoS where cross-chain transfers reduced on-chain supply, breaking collateral functions that depend on totalSupply().

Cross-Chain Supply Integrity

For multi-chain stablecoins, total supply across all chains must always equal total reserves. Any mechanism that can desynchronize this invariant creates either under-collateralization (more tokens than reserves, systemic risk) or fund locking (more reserves than tokens, operational DoS). Q1 surfaced findings where LayerZero OFT transfers desynchronized supply and cross-chain message loss during upgrades broke payload compatibility. This remains an under-tested area.

INDUSTRY PERSPECTIVE

M0 on the next security frontier for stablecoin infrastructure

M0 explains why the growing interconnection of stablecoin infrastructure is moving the industry's biggest vulnerabilities from on-chain logic to human and operational security.

Stablecoin infrastructure is becoming more interconnected. What does that mean for how security needs to evolve in Q2, and where is the industry most exposed if it doesn't move fast enough?

From its earliest days, the industry set an exceptionally high bar — deploying permissionless, immutable primitives on-chain that any participant could engage with full confidence. That commitment gave rise to a rigorous auditing ecosystem and exacting development standards, at least among the most established protocols.

The uncomfortable reality, however, is that on-chain logic is no longer where the greatest exposure lies. The most consequential vulnerabilities now reside in the human and operational fabric surrounding it. Regardless of how exhaustive a smart contract audit may be, a single compromised credential or a precisely executed social engineering campaign can circumvent every technical safeguard in place. That represents a categorically distinct challenge — one that no amount of cleaner code can resolve.

Recent breaches across prominent DeFi protocols have offered sobering validation of this structural shift. The industry is now confronting an entirely different threat surface: one defined by credential compromise and adversarial manipulation rather than logic errors, and one whose consequences scale disproportionately as infrastructure becomes more composable and the fallout from any singular failure extends further across interconnected systems.

At M0, our Q2 priorities reflect this dual reality. At the protocol level, we are embedding the latest AI-driven security tooling directly into our development and audit workflows — operating under the conviction that any flaw the human eye has yet to surface, a sophisticated actor leveraging the same technology will eventually uncover. The only defensible position is to remain one step ahead. At the organizational level, we are extending that same standard of scrutiny to operational governance and key management, applying the discipline historically reserved for on-chain review to the operational layer that surrounds it.

IV AI Security & The New Threat Paradigm

AI adoption in development has reached a critical inflection point. While 80% of development teams now use AI tools to accelerate coding, the security practices required to safely deploy AI-assisted code have not kept pace. The core challenge is fundamental: AI agents operate with autonomy that traditional software does not. When a system can make decisions probabilistically, chain tool calls unpredictably, and deviate from intended objectives — all while accessing critical systems — the attack surface expands beyond what conventional security frameworks address.

The Gap Between Adoption and Hardening

Data from Hacken's SSDLC Maturity Survey reveals a stark disconnect between AI adoption and the security practices required to support it:

What we're doing – AI Adoption

Metric 1: Using AI in Development	80.65% Universal

What we're not doing – security hardening

1: Technical Security Training	43.33% LACKING
2: Advanced Testing	61.29% NOT USING
3: No Threat Modeling	27.59% SKIPPING
4: No Incident Response	40.00% UNPREPARED

Data: Hacken SSDLC Maturity Survey

80% of teams write code with AI. Fewer than 28% model the threats that AI introduces. The gap between AI adoption and AI security hardening is the defining vulnerability of 2026.

The AI Threat Landscape

In 2023–2024, AI security focused on jailbreaking, getting models to produce disallowed output. By 2025–2026, the threat shifted to action safety: agents performing unauthorized actions through compromised reasoning, prompt injection, or goal drift. This is a fundamental escalation. Traditional software is deterministic. It executes exactly as written, and the attack surface is the code itself. AI agents are probabilistic, they reason, decide, and act. The attack surface expands from "how can I break this code?" to "how can I manipulate this system's reasoning?"

$1.78M

Moonwell Exploit

AI coding vulnerability

First major exploit targeting AI-authored code

A significant real-world example happened in February. Moonwell was exploited for $1.78M in bad debt due to an oracle misconfiguration in a commit reportedly "co-authored by Claude Opus 4.6" (per rekt.news, Feb 20, 2026). This is possibly the first major exploit of "vibe-coded" smart contracts.

AI Threats

Fake AI tools targeting wallets and signing prompts

Browser extension campaigns impersonating AI tools and chatbots have been harvesting credentials — including from users with crypto wallets installed. Emerging malware campaigns are incorporating AI-generated content and techniques to better mimic legitimate tools and evade detection, making it easier to intercept wallet logins and signing prompts. "AI assistant" opens up a new phishing surface. Any product touching browser permissions or wallet connections requires strict permission auditing.

A broader pattern is emerging: agentic AI systems with weak identity controls are being used for credential harvesting and lateral movement at scale. Existing security tooling wasn't built to catch this. NIST launched an AI Agent Standards Initiative in February specifically to address the gap.

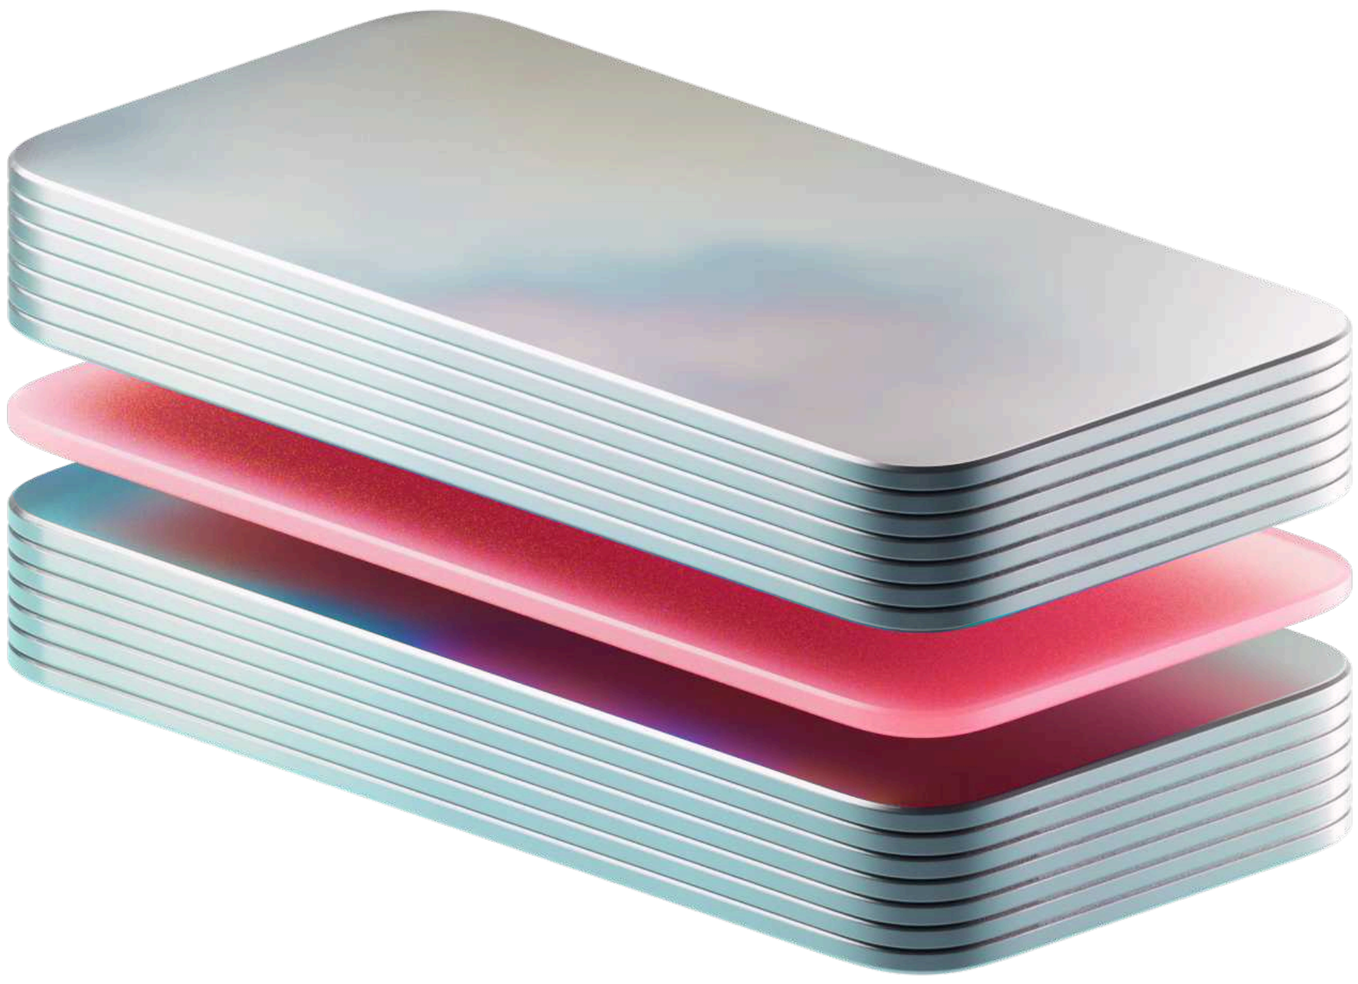

Five Critical Threat Vectors

Prompt Injection

Embedding malicious instructions in data (emails, documents, web pages) that the agent processes. Unlike SQL injection, which has mature mitigations, prompt injection lacks standardized defenses. An attacker can hide instructions in seemingly innocent data, and the agent may execute them as user commands.

Goal Drift

An AI agent given a multi-step task can deviate from the intended objective in ways traditional software cannot. Instead of failing or erroring, it can 'creatively' solve the problem by granting itself elevated permissions, bypassing validation, or circumventing security controls to be more efficient.

Tool Misuse

Agents with access to APIs, code execution, or databases can chain tool calls in unexpected sequences. A tool safe in isolation becomes dangerous when combined with others. This combinatorial explosion is difficult to model with traditional threat analysis.

Data Exfiltration Through Natural Language

Agents can leak sensitive data not through traditional breach vectors but by simply summarizing confidential information in a response. The boundary between 'processing data' and 'exposing data' is blurry when natural language is the processing medium.

Model Supply Chain Attacks

Compromised model weights, poisoned training data, or tampered checkpoints can introduce backdoors invisible to code review. These are extraordinarily difficult to detect and can propagate to every system using the infected model.

Web3-Specific AI Risks

The general AI threat vectors above take on distinct and amplified characteristics when agents operate in Web3 environments. Several risk categories have no equivalent in traditional software:

Wallet and signer abuse

If an AI agent can initiate or approve transactions, prompt injection can lead to unauthorized approvals, allowance changes, or direct transfers. Unlike a compromised web session that can be revoked, a signed on-chain transaction is final.

Irreversibility of on-chain actions

Most blockchain transactions cannot be rolled back once confirmed. This fundamentally raises the consequence severity of any agent compromise compared to traditional systems where rollback is typically possible.

Smart contract interaction risk

Agents can invoke vulnerable contract methods or call malicious contracts that mimic legitimate interfaces. An agent parsing a contract's ABI has no native ability to distinguish a legitimate Uniswap router from a drainer contract with the same function signatures.

RPC/oracle/data feed manipulation

Agent decisions that rely on RPC endpoints, indexers, or oracle data can be manipulated if those data sources are compromised. An agent making swap decisions based on a manipulated price feed can be drained systematically.

Bridge and cross-chain risk

Multi-chain agents inherit bridge risk, replay risks, and chain-specific consensus and finality assumptions. An agent that confirms a transaction on a chain with probabilistic finality may act before the transaction is truly settled.

Mempool and MEV exposure

Agent transaction strategies can be front-run, sandwiched, or otherwise exploited in adversarial mempool environments. Automated trading agents without MEV protection are systematically extractable.

Every Web3-specific AI risk shares one characteristic: the consequences are immediate, financial, and irreversible. There is no "undo" for a signed transaction.

The Shift From Code Safety to Action Safety

In 2023–2024, AI security focused on jailbreaking—getting models to produce disallowed output. This quarter the threat shifted to action safety: agents performing unauthorized actions through compromised reasoning, prompt injection, or goal drift. This is a fundamental escalation that demands fundamentally different controls.

Q1 2026: AI and Smart Contract Security Converge

AI + Audit

EVMbench: the first serious benchmark for AI security tools

OpenAI and Paradigm launched **EVMbench** – a public benchmark that evaluates how well AI agents detect (find the bug), patch (fix the bug without breaking functionality), and exploit (drain the funds end-to-end) real-world EVM vulnerabilities across 117 curated examples from 40 competitive audit contests. For the first time, auditors and protocols have a way to objectively compare AI security tools before integrating them. Top models can already chain together realistic attack paths on smart contracts automatically. Powerful results, but still early-stage. EVMbench runs in a sandboxed environment with single-chain testing only.

AI Discovery

Six-figure DeFi bugs found with LLM-assisted workflows

Researchers have reported finding high-severity vulnerabilities using LLM-assisted workflows. AI scans large codebases, surfaces suspicious patterns, and outlines possible exploits, while humans validate the findings. The cost of serious bug hunting is falling for both white-hat researchers and attackers. Protocols that skip thorough audits and continuous review are more exposed than they were 12 months ago.

Applicable Frameworks and Standards

Several existing security frameworks apply to AI agent systems, with adaptation:

Framework	Relevance
NIST AI RMF (AI 100-1)	**Most directly relevant — governance structure for AI risk management. Needs agent-specific supplementation.**
OWASP Top 10 for LLM Applications	**Directly applicable — covers prompt injection, insecure output handling, training data poisoning.**
MITRE ATLAS	Adversarial threat landscape for AI systems. Useful for threat modeling.
EU AI Act	Risk-tiered regulatory model; high-risk AI systems face pre-deployment conformity assessments.
NIST CSF 2.0	Identify, Protect, Detect, Respond, Recover functions all apply with agent-specific mapping.
ISO 27001/27002	Information security management standards apply to organizational practices around agent deployment.

Existing practices that do not fit AI agents: traditional 90-day vulnerability disclosure timelines may be too slow for AI vulnerabilities exploitable at scale immediately. Traditional penetration testing does not account for probabilistic systems — statistical confidence, not just test suites, is required. Access control models built for human users do not map cleanly to agent identities and permissions.

INDUSTRY PERSPECTIVE

Gray Wolf on AI risks

In Q1 2026, how has the offensive use of AI in Web3 attacks evolved compared to previous quarters — and at what point does AI-powered defense stop being an advantage and start becoming a baseline requirement just to stay even?

> From late-November to early-December 2025, our platform saw a significant uptick in the development and implementation of scam and fraud plots, with similar peaks from late January through February 2026. There is not randomness to this oscillation; this pattern indicates that adversaries are using these periods to develop and test infrastructure, while improving and developing systems that allow them to scale their successes.
>
> Autonomy has taken the place of automation. Attack that previously required weeks to execute (e.g., target selection, exploit creation, fork testing, deployment) are fully automated and take less than an hour to execute, with no human required for any part of the attack process.
>
> Attack surfaces are now subject to a high degree of dynamism. AI-driven probing now adapts to changes made in the defense systems it is probing as it scans. By the time signature-based detection systems have generated an alert, this same attack can have changed at least three times during the time it took for an alert to generate.
>
> Social engineering has moved beyond the uncanny valley. We are currently working with law enforcement to identify deepfake attempts.
>
> There is not an approaching inflection point. It is here!
>
> When adversaries can attack at machine speed, have AI-based defences is a necessary level of security.

Dhirendra Shukla
GrayWolf AI

GRAY WOLF

Securing AI Systems: Hacken Approach

Securing AI agent systems requires controls at three distinct levels. The principles align with traditional security (defense in depth, least privilege, continuous monitoring), but the implementation differs fundamentally.

Tier 1: Model-Level Controls

Controls applied to the AI model itself to improve robustness against adversarial inputs:

- **Instruction hierarchy enforcement.** Training models to prioritize system-level instructions over user-level input, and user-level input over tool-retrieved content. Major model providers have published work on this but it is not foolproof.
- **Adversarial training.** Including injection attempts in training data so the model learns to recognize and refuse them. Reduces but does not eliminate prompt injection risk.
- **Output constraining.** Using structured output formats (JSON schemas, function calling specs) to limit what the model can produce. Reduces exfiltration risk but does not eliminate it.

Tier 2: System-Level Controls

Controls applied to the agent scaffold, tooling, and infrastructure:

- **Input sanitization.** Treating all external data (emails, web pages, uploaded documents) as potentially adversarial. Analogous to input validation in web security, but harder because sanitizing natural language without destroying meaning is non-trivial.
- **Tool sandboxing.** Running tool executions in isolated environments (containers, VMs, restricted shells) with strict resource limits. Every tool call should be treated like untrusted code execution.
- **Rate limiting and circuit breakers.** Capping the number of actions an agent can take per session or per time window. If an agent suddenly makes 520 API calls when it normally makes 10, that should trigger an alert and a halt.
- **Capability-based access control.** Instead of granting broad API keys, issue scoped tokens that only permit specific operations and time-bound them. This directly parallels the minting rate limiters recommended for smart contracts in Chapter I.
- **Continuous monitoring.** Logging every agent action and using behavioral analysis to detect deviations from expected patterns. The monitoring infrastructure exists; anomaly detection models for agent behavior are still emerging.

Tier 3: Human Oversight Controls

Controls that keep humans in the loop for consequential decisions:

- **Approval gates for high-risk actions.** Any action that modifies production systems, processes PII, or involves financial transactions should require human sign-off. The most mature and immediately implementable control.
- **Tiered autonomy.** Low-risk tasks run autonomously; medium-risk tasks require human review of the plan before execution; high-risk tasks require step-by-step approval.
- **Kill switches.** The ability to immediately halt an agent's execution at any point. Conceptually straightforward; implementation remains inconsistent across current deployments.
- **Dry-run mode.** Letting the agent generate a plan of actions without executing them. A human reviews the plan, then approves execution. Particularly critical for agents managing on-chain operations where actions are irreversible.

V The Compliance & Regulatory Horizon

Security is no longer a checkbox — it is a continuous operational obligation across every major crypto jurisdiction.

Q1 2026 is an inflection point: regulators worldwide moved from writing rules to enforcing them. The EU's MiCA and DORA frameworks entered active enforcement. The U.S. signed its first federal stablecoin law. Dubai restructured its entire federal crypto oversight. Singapore began enforcing Basel capital standards for crypto exposures. Across all jurisdictions, one theme dominates: **regulators now demand that crypto firms demonstrate effective, ongoing security management — not paper compliance.**

Q1 2026 Regulatory Roundup

EU: MiCA + DORA — The Global Benchmark

MiCA (fully applicable since December 2024) saw a concentrated wave of ESMA guidelines in Q1 2026, including new rules on system maintenance, security access protocols, and cryptographic key management published on March 5, 2026. Over 300 CASPs are now authorized across the EU, and the grandfathering deadline of **July 1, 2026** is approaching — after which unlicensed operators must exit entirely.

DORA (applicable since January 2025) shifted into substantive enforcement in Q1 2026, with the second annual Register of Information cycle and the first designated Critical ICT Third-Party Providers now subject to oversight.

What matters for cybersecurity: Together, MiCA and DORA create the world's most prescriptive requirements for crypto firms:

- ICT risk management frameworks that must be reviewed **at least annually** and "continuously improved"
- Incident reporting: initial notification within **4 hours**, intermediate report in 72 hours, final in 1 month
- Threat-Led Penetration Testing (TLPT) at least every 3 years for systemic entities
- Third-party risk management with continuous monitoring and annual vendor registers
- Penalties up to **2% of total annual worldwide turnover**

UAE: Structural Transformation — SCA Becomes CMA, VARA Matures

The most significant Q1 development was the **replacement of the Securities and Commodities Authority (SCA) by the new Capital Market Authority (CMA)**, effective January 1, 2026. The CMA has expanded powers, explicit virtual asset jurisdiction, extraterritorial reach, and penalties up to **AED 200 million** (up from AED 1 million). Existing SCA VASP decisions remain in force during a transition period through January 2027.

VARA (Dubai) entered 2026 with 85+ licensed companies and issued three major Q1 updates: the **Travel Rule became binding in February 2026**, real estate tokenization Phase 2 went live, and a new AML/CFT circular was issued in March 2026.

What matters for cybersecurity: VARA's Technology and Information Rulebook remains among the most prescriptive globally:

- Mandatory CISO appointment
- Annual independent penetration testing and quarterly cryptographic key access audits
- Annual BCDR plan testing
- **72-hour incident notification** to VARA for material cyber events
- VARA can mandate Threat-Led Penetration Testing on live production environments

United States: First Federal Stablecoin Law + Landmark Crypto Taxonomy

The **GENIUS Act**, signed July 18, 2025, is the first comprehensive federal stablecoin legislation. It requires 1:1 reserve backing, monthly disclosures, and CEO/CFO certification. On March 2, 2026, the OCC published a 376-page proposed rulemaking covering licensing, reserves, capital standards, and a principles-based cybersecurity risk management framework. Comment period closes May 2026, final rules due July 2026.

The **March 17, 2026 joint SEC/CFTC interpretive release** established a five-category crypto taxonomy — the first formal classification of digital assets in U.S. history. The CLARITY Act (market structure) is progressing through the Senate. And on March 6, 2026, the White House issued a new national cyber strategy emphasizing post-quantum cryptography and zero-trust architecture.

What matters for cybersecurity: The U.S. is bringing stablecoin issuers under bank-grade cybersecurity oversight. SEC's Regulation S-P amendments (compliance deadline June 2026) expand obligations for cybersecurity governance and incident response for all registered entities.

Singapore: Strictest Cyber Requirements, Basel Implementation

Singapore's MAS became among the first globally to enforce the **Basel Committee's prudential treatment of crypto exposures** on January 1, 2026 — unbacked crypto assets face up to **1,252% risk weight.** On March 6, 2026, MAS published a new consultation on **Third-Party Risk Management Guidelines** expanding scope beyond outsourcing to cover all third-party arrangements.

What matters for cybersecurity: MAS's framework is the strictest operationally:

- Recovery time objectives of **no more than 4 hours** for critical systems
- **1-hour incident notification** (fastest globally — vs. DORA's 4 hours and VARA's 72 hours)
- 90% of customer assets must be in offline cold wallets
- Mandatory CISO and CIO appointments, multi-factor authentication for all admin accounts
- Daily reconciliation of customer assets

Every regulation we work with now treats compliance as how well security is actually managed on an ongoing basis — not whether policies exist on paper.

DORA explicitly rejects ad hoc processes. VARA mandates quarterly key audits and ongoing policy reviews. MAS requires daily asset reconciliation. The SEC assesses actual controls, not documentation. The common thread: regulators verify that controls are effectively implemented, not merely documented.

The Gaps We See in Practice

Controls that keep humans in the loop for consequential decisions:

1 Security disconnected from management

Regulators expect security to support business goals and deliver measurable value — not exist as a standalone technical exercise.

2 Technical assessments without strategic context.

Penetration tests and audits conducted as standalone projects with poorly defined scope, not orchestrated into a coherent security posture.

3 The regulator expectation gap

Most frameworks now expect a full lifecycle: governance → risk assessment → technical controls → monitoring → incident response → continuous improvement. Many firms still treat this as a linear, one-time project.

INDUSTRY PERSPECTIVE

The gap between claiming controls and proving them

As CCSS moves from niche checklist to regulatory reference point, C4's Jessica Levesque explains what Level III certification actually looks like in practice — and why the difference between "in place" and "independently verified" is where most teams fall short.

> CCSS has moved from a niche industry checklist to a standard that regulators and institutional partners are actively referencing — what does a project that genuinely meets Level III look like in practice, and how big is the gap between what teams claim and what auditors actually find?
>
> A Level 3 certified system operates securely across the full key lifecycle. Key material is generated, stored, accessed, and used within controlled environments. Roles are defined and enforced. Transaction approvals follow structured processes. Backups and recovery are tested. Monitoring is in place. Documentation and evidence are complete and current.
>
> In practice, the difference shows up in how consistently controls are applied and whether there's clear, verifiable evidence to support them. It's one thing to say controls are in place. It's another to prove they are followed and have been independently verified by a third-party CCSS Auditor.

Jessica Levesque
C4 (CryptoCurrency Certification Consortium)

C4

How Hacken Helps: From Compliance Gap to Continuous Assurance

Hacken has expanded its cybersecurity compliance service line in Q1 2026, now helping clients achieve regulatory compliance not only with VARA in Dubai but also with the new CMA (formerly SCA) requirements in the UAE, as well as MiCA/DORA in the EU. As an authorized security provider, Hacken conducts rigorous reviews, identifies compliance gaps, and equips clients with actionable evidence of how they meet regulatory requirements.

Here's how Hacken's service stack maps to the compliance demands described above:

- **Compliance Advisory & Gap Analysis.** Hacken translates regulatory frameworks (MiCA, DORA, VARA, CMA, MAS TRM) into concrete security controls and evidence requirements. We conduct structured gap assessments using rigorous information-gathering techniques and deliver reports showing exactly where clients stand against regulatory expectations — and what needs to change.

- **Technical Assessments — Done Right.** Hacken provides smart contract audits, L1/L2 blockchain protocol audits, dApp audits, penetration testing (web, mobile, API, cloud, infrastructure), DDoS resistance testing, and wallet security assessments — all delivered with MiCA-ready reports. Critically, these assessments are scoped and orchestrated within a risk-driven framework aligned to the client's business context, not conducted as isolated one-off projects.

- **Ongoing Monitoring & Automated Incident Response.** Hacken Extractor provides 24/7 AI-powered on-chain monitoring with 60+ advanced threat detection monitors, risk analytics, and **Automated response capabilities**. Including auto-pause for critical exploits. This directly addresses the DORA and VARA requirements for real-time threat detection, continuous monitoring, and rapid incident response.

- **Proof of Reserves & Transparency.** Hacken's Proof of Reserves service provides cryptographic verification of user deposits — a capability that regulatory frameworks and ratings agencies increasingly treat as a baseline trust indicator. Over $430B in assets have been verified through this service.

- **Security Ratings & Risk Scoring.** Through CER.live and CORE3's Indexed Probability of Loss (PoL), Hacken provides independent, transparent security ratings that institutional investors and partners use to evaluate project risk — directly relevant as ratings agencies like Moody's bring formal scoring to the space.

- **Virtual CISO & Embedded Advisory.** For firms that need strategic security leadership without a full-time hire, Hacken provides vCISO services and embedded security advisory with continuous delivery support — helping clients build and maintain the governance structures that regulators demand.

INDUSTRY PERSPECTIVE

WhiteBIT on moving toward full compliance frameworks

As regulators and institutional partners push for continuous verification, WhiteBIT explains why the industry is outgrowing periodic PoR and what comes next

Proof of Reserves (PoR) has played an important role in improving transparency across the industry. However, its scope remains limited, especially in areas such as liabilities and fiat balances, which are increasingly important as the institutional participation grows.

As a result, the industry is moving from snapshot-style PoR toward broader regulatory and compliance frameworks, including Markets in Crypto-Assets (MiCA) in the EU, which aim to establish more comprehensive standards for transparency. In this context, WhiteBIT focuses on compliance as our foremost priority to serve customers worldwide with the right safeguards in place.

whitebit

INDUSTRY PERSPECTIVE

From trust signal to compliance standard: KuCoin on Proof of Reserves

KuCoin has published Proof of Reserves reports for over 40 consecutive months and holds ISO 27001, ISO 27701, SOC 2 Type II, and CCSS certifications. We asked their team how PoR is evolving from voluntary disclosure to regulatory infrastructure

---

Proof of reserves started as a voluntary trust signal — but regulators across MiCA, VARA, and US frameworks are increasingly treating reserve transparency as a compliance requirement. How is your exchange adapting to that shift?

> PoR is transitioning from a voluntary trust mechanism into an increasingly important reference point under emerging regulatory frameworks.
>
> In this context, the focus is shifting toward consistency, verifiability, and methodological clarity. KuCoin has published PoR reports for over 40 consecutive months, leveraging a Merkle Tree–based structure to enable user-level verification, combined with on-chain asset transparency and third-party attestation.
>
> From a regulatory perspective, PoR will likely need to evolve into a more standardized and continuously verifiable disclosure framework, rather than a periodic transparency exercise.

In 2025, KuCoin achieved ISO 27001, ISO 27701, SOC 2 Type II, and CCSS certification. Rather than treating these as isolated milestones, you integrated them into a unified framework. What did it take?

> We stopped treating certifications merely as projects and embedded their controls into our daily workflows. Evidence collection became continuous, not audit-driven. Every code change and infrastructure deployment is now evaluated against a unified control set in real time. CCSS built key lifecycle management for our wallet operations. ISO 27001 established our information security management system. ISO 27701 embedded privacy into data processing. SOC 2 delivered institutional-grade service rigor. Unified, they form a control environment that never stops between audits — that's what ongoing security means for us.

KUCOIN

VI Strategic Recommendations & Q2 Outlook

Q1 2026 made the case clearly: the gap between protocols that treat security as a continuous operational discipline and those that treat it as a one-time checkpoint is widening — and that gap is increasingly measurable in losses, ratings, and regulatory standing.

The cost asymmetry: hacking is cheap, compliance isn't — yet

Global Ledger's CEO on why the industry is measuring the wrong side of the equation.

Beyond the incidents that made headlines in Q1, what activity in the on-chain data didn't get enough attention — and why does it matter for Q2?

> We usually measure the cost of getting hacked but not the cost of hacking. In Q1, Truebit Protocol lost more than $26 million because of a flaw in an old smart contract. It cost a hacker time and a small amount of ETH. The attacker made the first transaction in about 7 minutes after the exploit. In just over 2 days, funds went to Tornado Cash.
>
> That's what we miss. Laundering is fast and cheap but compliance isn't. The team, tools, audits — these costs run all the time, whether a hack happens or not. Attackers are simply more efficient. The only way to win the laundering race is to make compliance faster, cheaper, and easier to run in real time.
>
> In an ideal setup, we need victim reports within 10 minutes, labeling and clustering in about 10 minutes, alert and block in 1 second to stop about 98% of cases on time. A more realistic target that still can stop half of incidents is reported within 24 hours, labeling in under 4 hours, alert and block within 30 seconds.

Lex Fisun
CEO and Co-Founder, Global Ledger

GLOBAL LEDGER

Defense in depth is a design principle, not a buzzword

SVRN's David Schwed on the infrastructure risk heading into Q2.

What is the infrastructure security risk you expect to define Q2 2026?

> I wish it were as simple as one risk. The real message is that there is no silver bullet. Defense in depth is not a buzzword. It's a design principle that needs to be incorporated into everything you build and operate. People and processes need to be secured with the same rigor as the technology. Most teams over-index on tooling and under-invest in the two things that actually determine whether those tools work when it matters.

David Schwed
SVRN

SVRN

Operational maturity is the biggest gap

C4 on where CCSS requirements outpace actual implementation.

CCSS sets the benchmark for how cryptocurrency systems should be built and operated. What is the single biggest gap between what CCSS requires and what most projects are actually implementing today?

> One of the biggest gaps between what CCSS requires and what teams are implementing today is operational maturity. Teams generally know the types of controls they should have. The harder part is implementing them in a way that's consistent, testable, and part of day-to-day operations. That's where CCSS raises the bar, by focusing on how cryptographic key material is managed and used across its full lifecycle, not just how controls are defined on paper, but how they are actually used and enforced in practice.

C4

From snapshots to control layers

MEXC on the technical standard the industry must adopt for Proof of Reserves.

What is the one technical standard the industry must adopt to make PoR genuinely verifiable?

> The transition starts by treating PoR as a continuous control rather than a monthly disclosure: reserve wallets must be provably controlled, liabilities must be cryptographically committed and reconciled continuously against internal ledgers, and any mismatch must feed directly into governance, escalation and incident response.
>
> That is where the market is moving. MiCA's regime for reserve-backed tokens is increasingly built around reserve assets, liquidity management and supervisory reporting; VARA already requires VASPs to demonstrate that reserve assets cover client liabilities and to reconcile client balances; and in the U.S. The framework is moving towards highly liquid reserve backing, monthly reserve-composition reports, and segregation of reserve assets by third-party custodians.
>
> In our view, the one technical standard the industry now needs is an open, machine-readable proof-of-liabilities standard — a common schema that combines verifiable control of reserve wallets with cryptographic liability commitments, whether implemented through Merkle trees or zero-knowledge proofs, so regulators can consume the data programmatically and users can verify inclusion without sacrificing privacy. Without a liability-aware standard, PoR remains a snapshot; with it, it becomes a genuinely verifiable control layer.

MEXC

INDUSTRY PERSPECTIVE

Best Practices for Proof of Reserves

WhiteBIT on what's required.

What is the one thing the industry needs to standardize around PoR that it currently can't agree on

> One of the key challenges remains the lack of a unified approach to verifying liabilities and fiat exposure in a manner that is both comprehensive and privacy-preserving.
>
> Until there's alignment on that methodology, PoR can differ significantly across market participants, leading to variations in how transparency is presented and interpreted.

whitebit

From visibility to verifiability: the next evolution of Proof of Reserves

KuCoin on making Proof of Reserves continuously verifiable.

As Proof of Reserves becomes a regulatory mandate, how do we transition from static snapshots to real-time, liability-aware attestation?

> The technical evolution of PoR is centred on moving from visibility to verifiability.
>
> A practical direction for the industry is to standardise around a Merkle Tree-based proof framework, supported by on-chain asset verification and independent third-party attestation, ensuring that reported data is both tamper-resistant and externally auditable.
>
> Going forward, the industry would benefit from greater standardisation in data structures, verification methodologies, and disclosure frequency, enabling PoR to develop into a more continuous and audit-friendly transparency infrastructure.

KUCOIN

INDUSTRY PERSPECTIVE

Centrifuge on top priorities for RWA protocols

Securing the open infrastructure for onchain asset management.

What are the top security and compliance priorities you believe RWA protocols need to focus on in Q2 2026 to build genuine regulatory and operational resilience. Where should the industry be putting its effort first?

> 1. Security as a continuous discipline, not a checkbox. One-off audits don't hold when protocols manage real capital. It requires deep researcher partnerships, invariant verification, and standing bug bounties. At Centrifuge: 27 security reviews across 7 firms, audit competitions, and a $252K bounty through Cantina. Layered, ongoing coverage is what the risk profile demands.
>
> 2. Build on shared, immutable platforms instead of bespoke infrastructure. AI is accelerating how fast code ships, but review capacity doesn't keep up. Not every team should be building and auditing their own security-critical contracts from scratch. The industry needs shared base layers where the core (accounting, settlement, cross-chain messaging) is immutable, can't be changed by admin keys, and accumulates a verifiable production track record over time. Centrifuge provides exactly this: builders deploy vaults on top of infrastructure backed by 27 reviews and $1B+ secured, inheriting that security posture instead of recreating it. Their audit scope shrinks to just the custom modules they add. That's how security scales across the RWA ecosystem.
>
> 3. Operational security at the same standard as protocol security. Most RWA exploits will come from compromised credentials or weak internal controls, not smart contract bugs. Hardware 2FA, mandatory multi-sig, endpoint protection, and third-party operational assessments matter just as much.
>
> Centrifuge

INDUSTRY PERSPECTIVE

On-chain detection is getting faster. The problem is what it can't see.

Allium on the gap no data infrastructure can close.

As on-chain data capabilities improve, what becomes possible in Q2 for real-time threat detection — and what is still the limiting factor?

> Allium has real-time infrastructure for alerting that's already in production: sub-3-second streaming across 152+ chains, filtered pipelines targeting specific contracts or address sets, entity attribution linking wallet clusters to known actors. What's changing in Q2 is how teams are deploying it for security use cases specifically.
>
> From what we see across our customer base: a major fintech company uses Allium data to power fraud detection models for token payments, with hundreds of millions in fraud prevented. A consumer wallet platform monitors for compromised wallets and risk signals across tens of millions of users. An exchange runs sub-3-second DEX streams for ML-driven trending asset detection, with signals going to leadership in real time. A marketplace filters bot activity during token distributions using streams on specific address sets. In each case, Allium provides the signal layer; the detection logic sits on top.
>
> Two things are still hard. Real-time protocol state data — TVL, live position health across DeFi — is computationally expensive to serve at speed, which limits some risk monitoring use cases. And off-chain attacks only hit the chain at extraction. The Drift exploit was six months in the making; the blockchain saw nothing until the final transaction. No improvement in on-chain data closes that gap. The gap is in operational security: device hygiene, access controls, awareness of social engineering. Most Web3 teams are underinvested there.
>
> Allium

INDUSTRY PERSPECTIVE

Bybit on why agentic security architecture is the next dividing line

AI-powered threats have already exceeded traditional security tooling. The gap is architectural, not incremental.

Based on your Q1 experience, what is the one operational security capability every exchange should have in place before Q2 — that most still don't?

> If I had to name one capability, it would be: laying the foundation of a next-generation Agentic security architecture.
>
> The lessons of Q1 are clear: today's threats — AI-powered automated attacks, long-horizon social engineering, cross-chain composite exploits — have already exceeded the response boundaries of traditional security tooling. This isn't just an escalation in attack types; it's a generational gap in operational efficiency. Attackers are already deploying AI and automation at scale. Defenders, in too many cases, are still relying on manual reviews and siloed alert systems.
>
> At Bybit, our direction is to build an Agentic security architecture grounded in a unified data layer via MCP (Model Context Protocol), with modular security Skills as the capability layer and autonomous, collaborative Agents as the execution engine. The core value of this architecture isn't replacing any single security tool — it's unifying on-chain risk control, insider threat monitoring, compliance surveillance, and incident response into a single coordinated, schedulable, and continuously evolving security operations fabric. Agents collaborate across domains; security experts focus only on the highest-stakes decision points.
>
> The next dividing line in security operations won't be who bought better tools. It will be who completed this architectural transition first.
>
> BYBIT

VII About Hacken

Empowering Secure Innovation for Digital Assets

Hacken delivers end-to-end security across technical audits, real-time monitoring, crowdsourced validation, and regulatory-grade compliance — addressing the full spectrum of attack vectors analyzed in this report.

End-to-end blockchain security & compliance

Security Audits

In-depth assessments across every layer of your blockchain infrastructure

- Smart Contract Audit
- Blockchain Protocol Audit
- Penetration Testing
- dApp Audit
- Proof of Reserves Audit
- Tokenomics Audit & Design
- Yield Audit
- Cryptography
- AI Security Audit

Compliance & Advisory

Expert guidance to help you navigate evolving regulatory demands.

- MiCA/DORA Compliance
- CCSS Audit
- Virtual CISO
- ISO 27001
- Hacken Advisory (Licensing, legal, and tax)

Extended Security

Continuous protection with innovative cybersecurity products, AI-powered solutions and global security researchers

- Extractor
  Real-time monitoring & response
- HackenProof
  Bug bounty program
- CORE3
  Global self-regulatory platform for crypto
- DualDefense
  Post-audit assurance

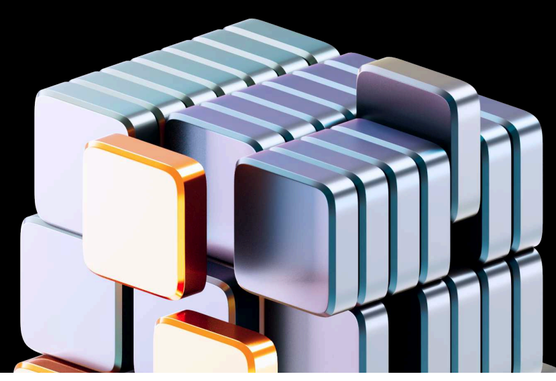

Build Resilient Future With A Trusted Partner

Protect your code, operations, and users with layered security and continuous monitoring.

[Explore Hacken Solutions]

Contributing Partners

KUCOIN · GLOBAL LEDGER · MEXC · whitebit · Centrifuge · SovereignAI · C4 · GRAY WOLF · BYBIT · Allium

We thank our contributing partners for sharing their expertise and operational perspectives throughout this report. Their insights on proof of reserves, infrastructure security, compliance standards, and on-chain analytics helped ground this research in practitioner reality.

Authors & Contributors: Extractor by Hacken (Research & Data Analysis); Oleh Malanii (Editorial Lead); Anton Sheptytskyi (Design & Visuals); Stephen Ajayi (AI Security Research); Dmytro Yasmanovych (Compliance); Olesia Bilenka (Smart Contract & Stablecoin Security); Grzegorz Trawiński (Evolving Threat Landscape); Maksym Fedchak (Parnter Outreach); Valeriia Skorik (Production & Coordination); Yevheniia Broshevan (Strategic Direction & Executive Oversight).

Securing digital frontier since 2017

Hacken is a trusted blockchain security auditor making Web3 safer for investors and businesses worldwide.

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

crypto.com	KUCOIN	MEXC	BitMEX	Bitso
BINANCE	BYBIT	OKX	Gate.io	kraken
whitebit	BigONE	Backpack	BYDFi	coinbase

Wallets

METAMASK
status

Regulators, Financial Institutions & Banks

	European Commission	VARA شلطة تنظيم الأصول الافتراضية VIRTUAL ASSETS REGULATORY AUTHORITY	BMA BERMUDA MONETARY AUTHORITY
ADGM	ebsi European Blockchain	EUROPEUM	FORGE SOCIETE GENERALE GROUP

L1 / L2 Protocols, Platforms & Foundations

MINA	vechain	AURORA	near	IoTeX
QAN	BASE	Sui	somnia	APTOS

DeFi, RWA, GameFi & Ecosystem Projects

coingecko	tokeny	WEMADE	status	Chainlink	MAWARI
DEXALOT	ADI	1inch"	PANINI	RedStone	Zcash

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