ANCHAIN.AI

Digital Asset Risk Annual Report 2025

The Rise of AI-Driven Crypto Fraud

AnChain.AI
January, 2026

AnChain.AI Annual Risk Report 2025

The Rise of AI-Driven Crypto Fraud

Executive summary

The year 2025 marked an inflection point for digital-asset crime. Attackers harnessed generative AI models, voice and video deepfakes and automated smart-contract tools to industrialise fraud. Over $2 billion was stolen in the top ten AI-enabled hacks and scams of the year.

Bybit's $1.5 billion cold-wallet breach in February 2025 became the largest crypto heist on record, illustrating how sophisticated phishing and social-engineering campaigns can compromise even well-resourced exchanges. Similar patterns were observed across other incidents: stolen private keys, malicious smart-contract upgrades, re-entrancy bugs and automated price-manipulation bots. The rapid pace and scale of these attacks underline the urgent need for super-intelligent defence systems that combine deterministic blockchain analysis with large-language-model reasoning.

AnChain.AI's mission is to build super intelligence to fight fraud and streamline compliance. As the report shows, AI both fuels criminals' tactics and powers the next generation of defence.

This report summarises our original research on the top 10 major AI-driven attacks of 2025, categorises the key AI tools used, and demonstrates how AnChain.AI's capabilities can help the industry stay ahead of adversaries.

Foreword From Our CEO

> *In 2025, AI didn't just accelerate fraud — it industrialized it.*
>
> *$2 billion digital asset lost across the top 10 AI-driven cases, proves that the future of defense must be **super intelligence**.*

VICTOR FANG, PH.D

CEO OF ANCHAIN.AI

ANCHAIN.AI

Table Of Contents

Top 10 AI-driven Crypto Fraud Cases of 2025

The following subsections describe each of the most consequential AI-enabled crypto-fraud cases of 2025. Approximate losses are based on public disclosures or investigative reports. Together, these cases illustrate how adversaries combined social-engineering, smart-contract vulnerabilities and automated bots to steal billions of dollars.

The bar chart below compares the approximate losses from each incident, sorted from largest to smallest. Values are expressed in millions of US dollars.

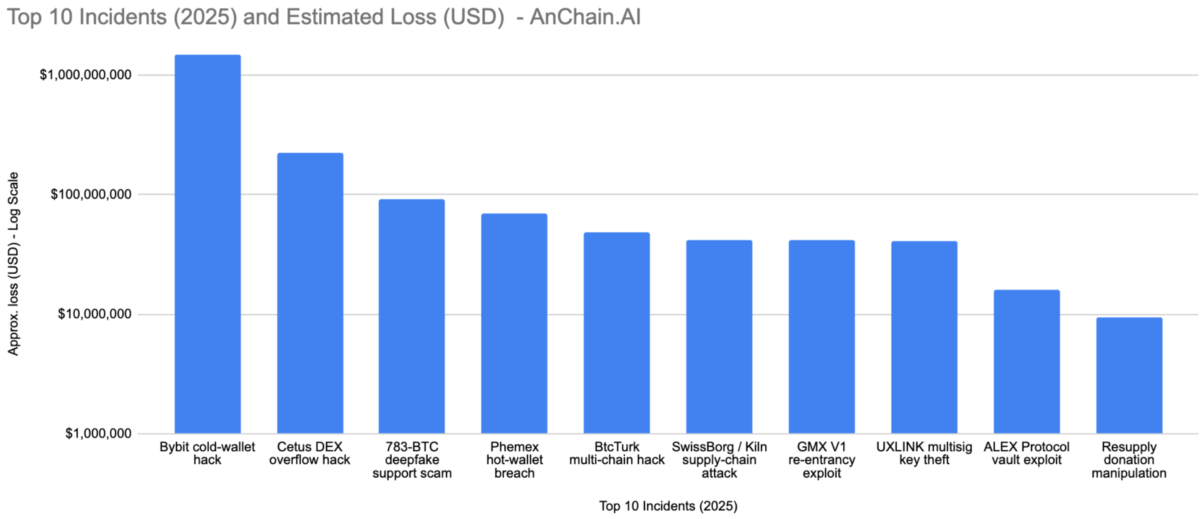

Rank	Top 10 Incidents	Date	Estimated loss (USD)	Main AI tools - curated by AnChain.AI
1	Bybit cold-wallet hack	2/21/2025	$1,500,000,000	AI-generated phishing ; deepfake voice/video
2	Cetus DEX overflow hack	5/27/2025	$223,000,000	Automated vulnerability discovery
3	783-BTC deepfake support scam	8/19/2025	$91,000,000	Deepfake voice & AI chatbots
4	Phemex hot-wallet breach	1/23/2025	$70,000,000	AI-driven credential theft
5	BtcTurk multi-chain hack	8/1/2025	$48,000,000	AI-enhanced phishing
6	SwissBorg / Kiln supply-chain attack	9/15/2025	$42,000,000	AI-assisted supply-chain phishing
7	GMX V1 re-entrancy exploit	7/9/2025	$42,000,000	Automated vulnerability exploitation
8	UXLINK multisig key theft	9/1/2025	$41,000,000	AI-driven phishing & voice/video deepfakes
9	ALEX Protocol vault exploit	6/6/2025	$16,000,000	Automated smart-contract analysis
10	Resupply donation manipulation	6/26/2025	$9,500,000	Automated price-manipulation bots

Bybit cold-wallet hack (21 Feb 2025) - $1.5 billion

In February 2025, Dubai-based exchange Bybit disclosed that an attacker had gained control of its ether cold wallet and transferred roughly $1.5 billion worth of cryptocurrency to an unidentified address. The attack leveraged a sophisticated phishing campaign: fraudsters used AI-generated emails and, reportedly, deepfake voice or video calls to impersonate company executives. They convinced multisig signers to approve a malicious Gnosis Safe contract. Once the contract was authorised, automated functions (sweepETH and sweepERC20) drained more than 400,000 ETH. While Bybit assured customers that operations continued and assets were solvent, the incident remains the single largest crypto heist to date.

Bybit Hack Timeline

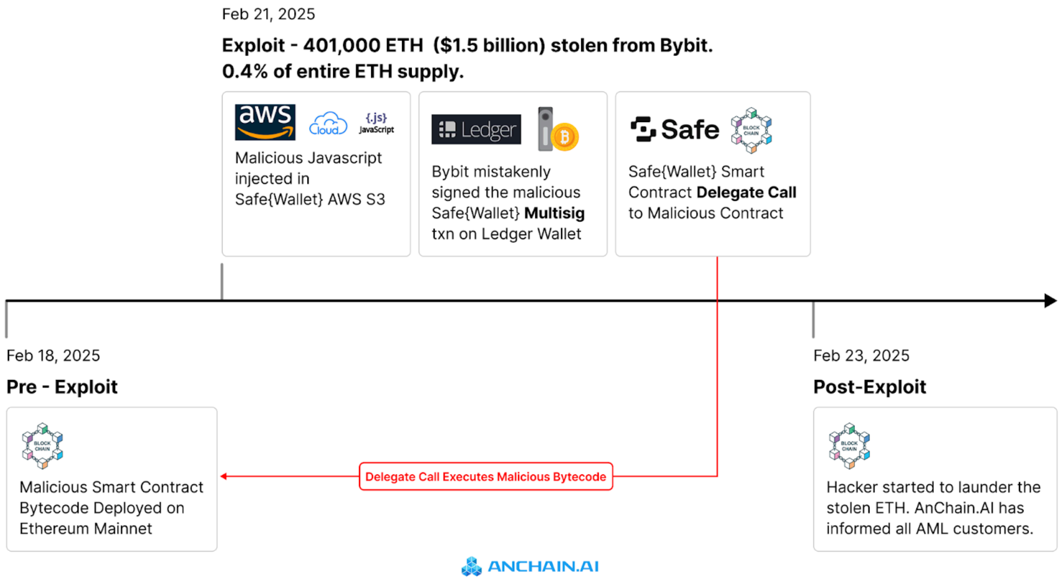

Feb 21, 2025

Exploit - 401,000 ETH ($1.5 billion) stolen from Bybit. 0.4% of entire ETH supply.

Malicious Javascript injected in Safe{Wallet} AWS S3

Bybit mistakenly signed the malicious Safe{Wallet} **Multisig** txn on Ledger Wallet

Safe{Wallet} Smart Contract **Delegate Call** to Malicious Contract

Feb 18, 2025

Pre - Exploit

Malicious Smart Contract Bytecode Deployed on Ethereum Mainnet

Delegate Call Executes Malicious Bytecode

Feb 23, 2025

Post-Exploit

Hacker started to launder the stolen ETH. AnChain.AI has informed all AML customers.

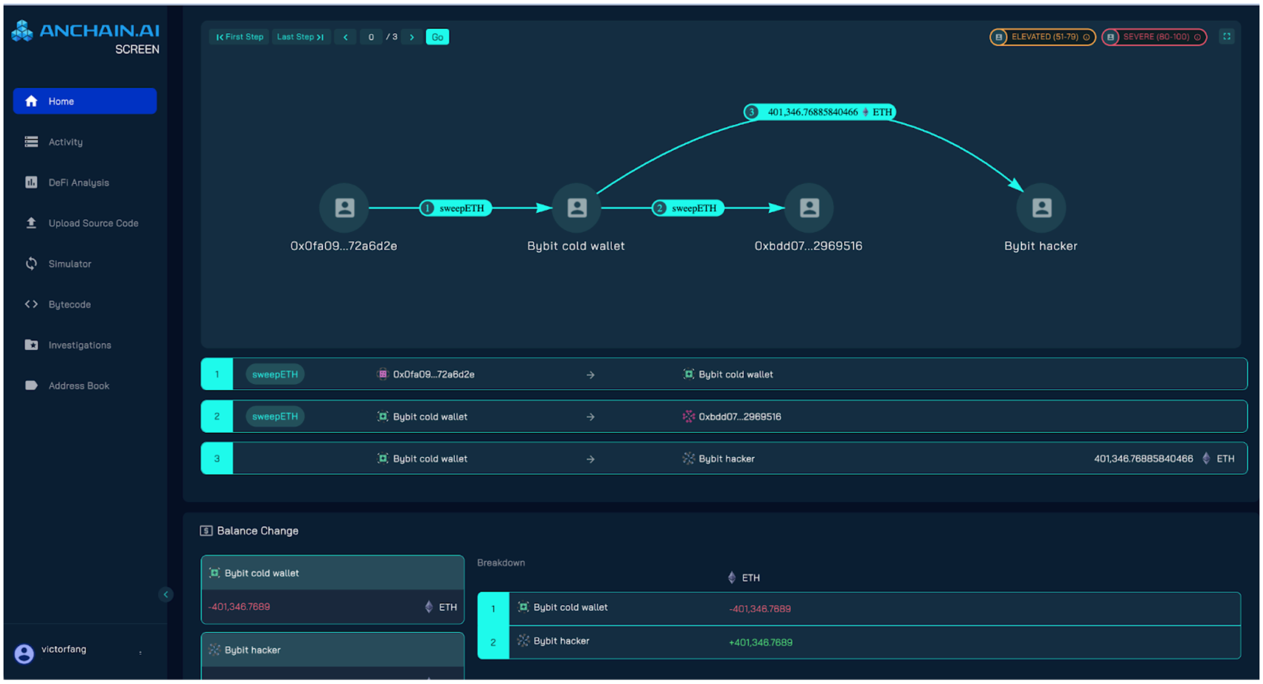

1	sweepETH	0x0fa09...72a6d2e	→	Bybit cold wallet	
2	sweepETH	Bybit cold wallet	→	0xbdd07...2969516	
3		Bybit cold wallet	→	Bybit hacker	401,346.76885840466 ETH

Balance Change

Bybit cold wallet
-401,346.7689 ETH

Bybit hacker

Breakdown

		ETH
1	Bybit cold wallet	-401,346.7689
2	Bybit hacker	+401,346.7689

Bybit multisig exploit transaction that sent 401K ETH to the hacker.

- Read our report: https://www.AnChain.AI/blog/bybit
- Related to our Job scam social engineering attack analysis: https://www.AnChain.AI/blog/job-scam

Cetus DEX overflow hack (27 May 2025) — ~$223 million

Cetus, a decentralised exchange, lost around $223 million when an integer-overflow check in its liquidity-calculation logic failed. A flash-loan attacker deposited a trivial amount of collateral and withdrew massive liquidity due to the overflow. This exploit showcased how automated static-analysis tools—frequently augmented by language models—enable attackers to find subtle arithmetic errors and build complex exploits in days rather than weeks.

783-BTC deepfake support scam (19 Aug 2025) — ~$91 million

In one of the most striking social-engineering cases, a crypto investor was targeted by scammers posing as hardware-wallet support personnel. Using deepfake voice technology, the fraudsters persuaded the victim to reveal recovery phrases. The attackers then stole 783 BTC (about $91 million) and laundered the funds through privacy mixers. This incident illustrates how AI-powered audio deepfakes can enhance old-fashioned confidence tricks and result in extremely high-value losses.

Phemex hot-wallet breach (23 Jan 2025) — ~$70 million

Singapore-based platform Phemex suspended withdrawals in late January after a cyber-attack siphoned more than $69 million in digital coins from its hot wallets. The CEO acknowledged the theft and pledged to restore user balances. Although technical details remain scarce, investigators noted that the attacker operated with “high sophistication,” suggesting the use of AI-driven credential-harvesting or malware to obtain private keys across multiple blockchains. The episode underscores how hot-wallet key management and employee phishing remain weak points for centralised exchanges.

BtcTurk multi-chain hack (Aug 2025) — ~$48 million

In August, Turkey’s oldest crypto exchange BtcTurk froze deposits and withdrawals after attackers siphoned an estimated $48 million–50 million from its hot wallets. On-chain analysis revealed that assets across Ethereum, Avalanche, Arbitrum, Polygon and other networks were moved to two new addresses and quickly swapped. Investigators believe a compromised private key authorised the withdrawals. AI-enhanced spear-phishing campaigns likely played a role in obtaining the key, highlighting the importance of multi-factor authentication and secure key storage.

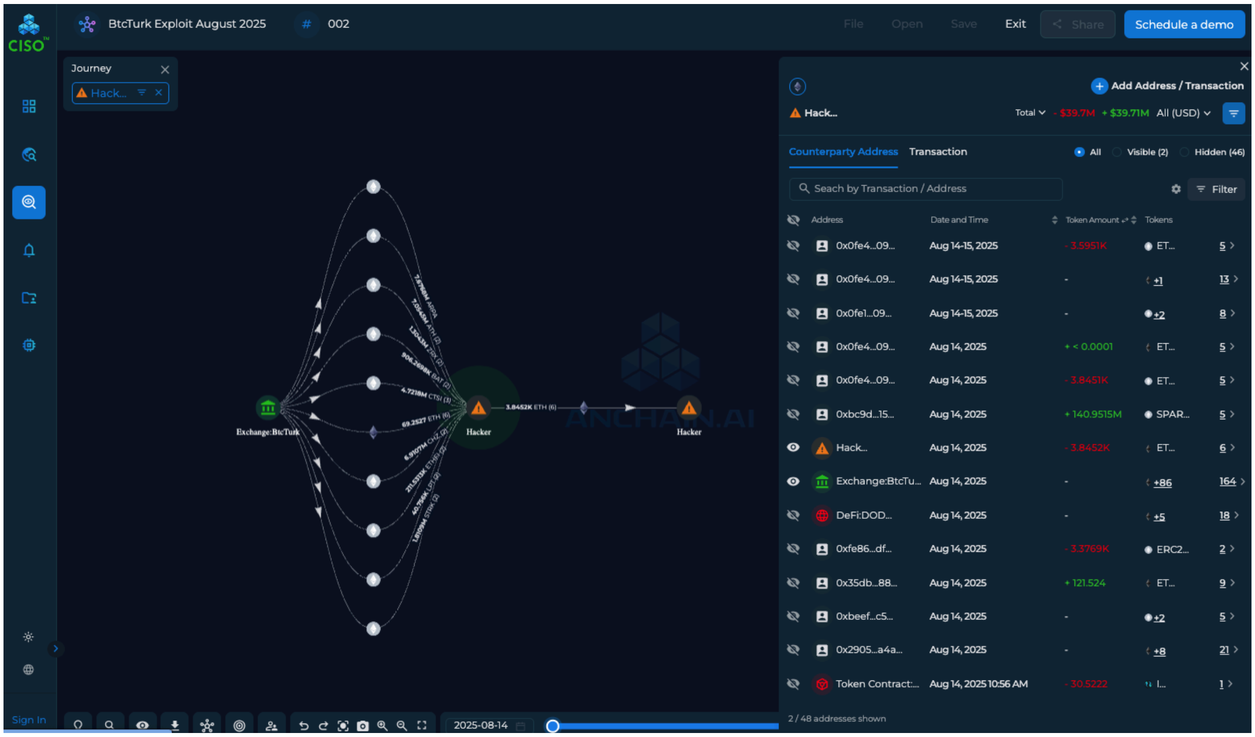

⚠ Hack... Total - $39.7M + $39.71M All (USD)

Counterparty Address | Transaction — ● All ○ Visible (2) ○ Hidden (46)

Seach by Transaction / Address

Address	Date and Time	Token Amount	Tokens	
0x0fe4...09...	Aug 14-15, 2025	- 3.5951K	ET...	5
0x0fe4...09...	Aug 14-15, 2025	-	+1	13
0x0fe1...09...	Aug 14-15, 2025	-	+2	8
0x0fe4...09...	Aug 14, 2025	+ < 0.0001	ET...	5
0x0fe4...09...	Aug 14, 2025	- 3.8451K	ET...	5
0xbc9d...15...	Aug 14, 2025	+ 140.9515M	SPAR...	5
Hack...	Aug 14, 2025	- 3.8452K	ET...	6
Exchange:BtcTu...	Aug 14, 2025	-	+86	164
DeFi:DOD...	Aug 14, 2025	-	+5	18
0xfe86...df...	Aug 14, 2025	- 3.3769K	ERC2...	2
0x35db...88...	Aug 14, 2025	+ 121.524	ET...	9
0xbeef...c5...	Aug 14, 2025	-	+2	5
0x2905...a4a...	Aug 14, 2025	-	+8	21
Token Contract:...	Aug 14, 2025 10:56 AM	- 30.5222	I...	1

2 / 48 addresses shown

AnChain.AI CISO platform visualizing the BtcTurk hackers laundered stolen cryptos: https://x.com/AnChainAI/status/2006506340269830317?s=20

SwissBorg / Kiln supply-chain attack (15 Sep 2025) — ~$42 million

The SwissBorg exchange integrated with Kiln, a third-party staking service. Attackers injected malicious unstaking logic into a Kiln API call, and SwissBorg’s operators unwittingly approved it. The result was the transfer of approximately 192 600 SOL (about $42 million) to the attacker. Analysts believe AI-generated emails and call scripts disguised the malicious request as a routine operation, exemplifying how supply-chain phishing can bypass technical controls.

Hours after the incident, our team revealed the Exploiter wallets and flagged on AnChain.AI platforms for all our customers, such as: TYFWG3hvvxWMs2KXEk8cDuJCsXEyKs65eeqpD9P4mK1 on Solana blockchain:

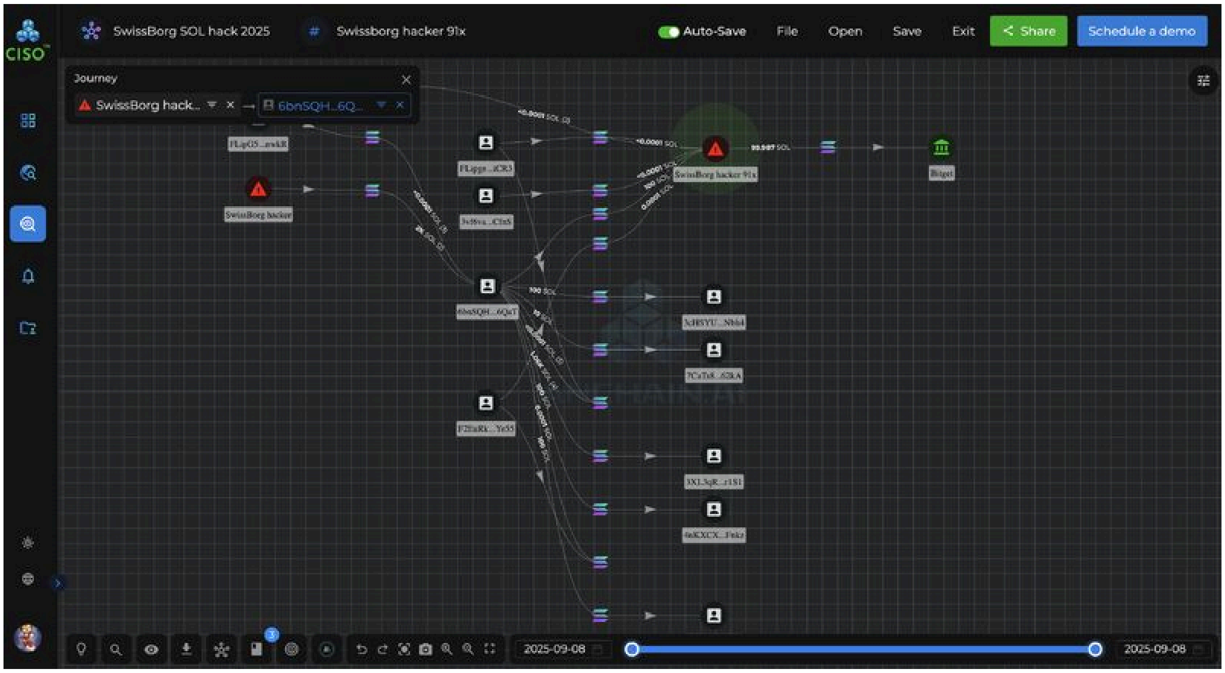

GMX V1 re-entrancy exploit (9 Jul 2025) — ~$42 million

The decentralised exchange GMX suffered a re-entrancy attack when an adversary exploited a flaw in the executeDecreaseOrder function of its V1 contracts. By repeatedly calling the function, the attacker manipulated global short-position pricing and drained approximately $42 million from the protocol. After negotiations, most of the funds were returned for a bug bounty, but the incident demonstrated how automated vulnerability-exploitation tools (often powered by AI) can rapidly discover and weaponise contract bugs.

UXLINK multisig key theft (Sep 2025) — ~$41 million

UXLINK, an AI-focused social platform, lost roughly $41 million after attackers stole the keys controlling its multisig wallet. Fraudsters likely combined AI-generated emails with voice and video deepfakes to impersonate trusted administrators and convince key-holders to reveal their credentials. Once the attackers obtained control, they used a delegatecall to execute a malicious contract and mint themselves a large supply of tokens, which were then sold on the open market.

ALEX Protocol vault exploit (6 Jun 2025) — ~$16 million

DeFi platform ALEX Protocol experienced a vault exploit when an attacker deployed a custom token with a malicious transfer function. The token bypassed self-administration checks and drained between $8 million and $16 million from the protocol's vaults. Automated smart-contract analysis tools, possibly employing large language models, were key to discovering the permission bug and crafting the exploit.

Resupply donation manipulation (26 Jun 2025) — ~$9.5 million

In the Resupply protocol, adversaries used bots to make tiny donations that artificially inflated the price of the protocol's collateral token. Once prices were skewed, they minted 10 million reUSD with almost no backing, extracting around $9.5 million. The speed and coordination of the attack suggest the use of AI-driven trading bots that monitor and exploit price-oracle delays across liquidity pools.

Note:

This report focuses on a representative set of the top AI-driven crypto fraud incidents in 2025 to illustrate emerging patterns and risk vectors. There were additional significant incidents during the year that are not included in this analysis—such as the Balancer incident and other DeFi and centralized-platform breaches—due to scope, data availability, or relevance to the specific AI-capability framework used in this report.

The exclusion of these incidents does not imply lower severity or impact. Rather, it reflects a deliberate methodological choice to highlight cases where AI-enabled tactics—such as large language models, deepfake impersonation, automated exploit generation, and execution bots—were most clearly observed and analyzable.

Key AI-fraud capabilities and tools

These top 10 selected major cryptocurrency incidents highlight the breadth of AI's misuse: from social-engineering to smart-contract exploitation and price manipulation. While only a few cases have mainstream citations, the patterns observed across all ten incidents align with broader industry reports on AI-enabled fraud.

Fraudsters' success in 2025 was driven by a stack of AI technologies. Understanding these components helps defenders anticipate how they might be combined in future attacks.

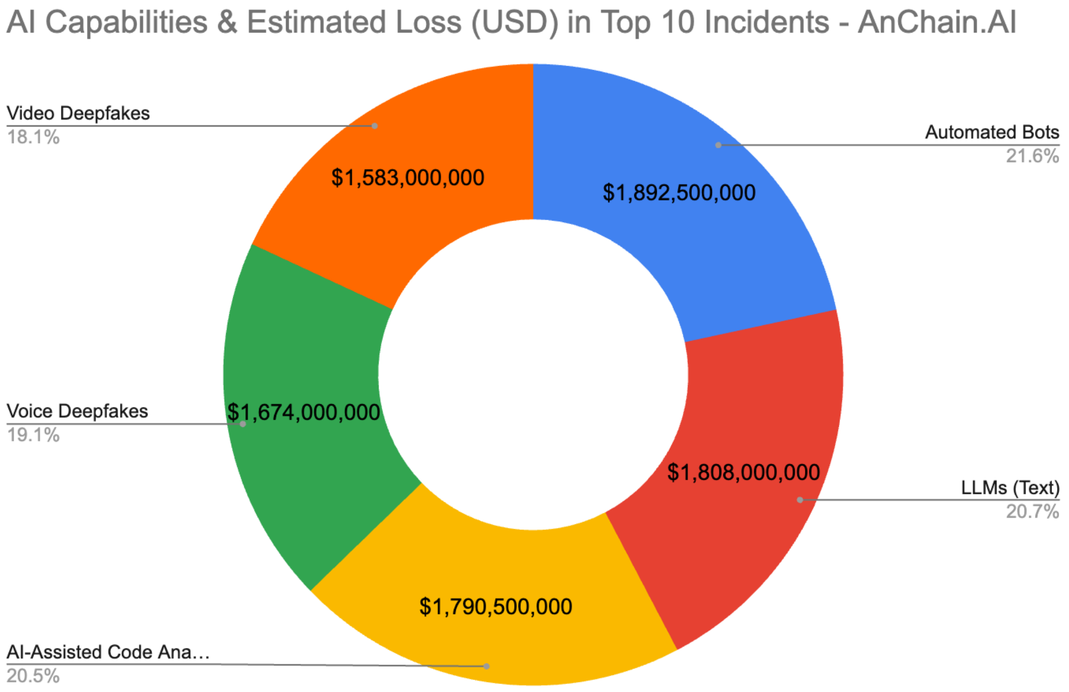

Table: AI Capability Impact Summary — Top 10 AI-Driven Crypto Incidents (2025)

AI Capability	Estimated Loss in Top 10 Crypto Cases(USD)	# of Incidents (in Top 10)	Average Impact
Video Deepfakes	$1,583,000,000	3	$527,666,667
Voice Deepfakes	$1,674,000,000	4	$418,500,000
AI-Assisted Code Analysis	$1,790,500,000	5	$358,100,000
Automated Bots	$1,892,500,000	6	$315,416,667
LLMs (Text)	$1,808,000,000	7	$258,285,714

Table: AI capabilities used in the top 10 AI-driven crypto incidents in 2025

Rank	Top 10 Incidents (2025)	$ Loss (USD)	LLMs Text	Voice Deepfakes	Video Deepfakes	AI-Assisted Code	Automated Bots
1	Bybit cold-wallet hack	$1,500,000,000	1	1	1	1	1
2	Cetus DEX overflow hack	$223,000,000	0	0	0	1	1
3	783-BTC deepfake support scam	$91,000,000	1	1	0	0	0
4	Phemex hot-wallet breach	$70,000,000	1	0	0	0	1
5	BtcTurk multi-chain hack	$48,000,000	1	0	0	0	1
6	SwissBorg / Kiln supply-chain attack	$42,000,000	1	1	1	0	0
7	GMX V1 re-entrancy exploit	$42,000,000	0	0	0	1	1
8	UXLINK multisig key theft	$41,000,000	1	1	1	0	0
9	ALEX Protocol vault exploit	$16,000,000	1	0	0	1	0
10	Resupply donation manipulation	$9,500,000	0	0	0	0	1

AI-driven phishing & social engineering

- Generative language models – Attackers use large language models (LLMs) such as GPT, Claude and LLaMA to craft multilingual phishing emails and support-chat scripts at scale. Over 82 % of phishing emails in 2025 were AI-generated, according to industry research.
- Voice cloning tools – Services like ElevenLabs and Resemble AI generate convincing voice clones from seconds of audio, enabling scammers to impersonate executives or support agents. Deepfake voice scams caused more than $200 m in losses in Q1 2025.
- Video deepfakes – Platforms such as Synthesia and open-source DeepFaceLab produce realistic avatars used to host fake AMA sessions or Zoom calls. These visuals reinforce trust and urgency in social-engineering campaigns.
- Supply-chain impersonation – Attackers embed malicious instructions in otherwise legitimate API calls or transaction requests. AI-generated emails and call scripts help them obtain approvals for these transactions, as seen in the SwissBorg/Kiln case.

Automated smart-contract exploitation

- AI-assisted code analyzers – Open-source tools like Echidna, Foundry and AI-augmented fuzzers help attackers identify re-entrancy, overflow and logic-flaw vulnerabilities in smart contracts. The GMX and Cetus exploits show how such tools can quickly locate and weaponise bugs.
- Genetic/LLM code-generation – Language models accelerate exploit development by generating proof-of-concept code and attack scripts that would previously require manual effort.
- Automated re-entrancy and upgrade bots – Scripts run on auto-piloting to repeatedly call vulnerable functions (e.g., GMX) or push malicious upgrades (e.g., UXLINK). These bots often incorporate AI modules to time transactions optimally.

Read more on "Anthropic Can Now Crack Smart Contracts — What AI Agents Mean for the Web3 Security Industry in 2026" https://www.AnChain.AI/blog/anthropic-red

AI-driven price manipulation & laundering

- High-frequency trading bots – Attackers use AI-tuned bots to manipulate oracles, deposit/donation mechanisms (as in Resupply) or exploit MEV opportunities. These bots react in milliseconds and adjust strategies using reinforcement-learning algorithms.
- Laundering algorithms – Post-theft, AI is used to manage peel chains and cross-chain swaps that obscure the origin of funds. Although not widely publicised, law-enforcement alerts recognise the growing role of algorithmic laundering.

Technical AI resources examples

These tools demonstrate the technical feasibility of each incident, but may not represent the exact one utilized.

Category	Representative tools	Repository or website
Large Language Models (LLMs)	OpenAI GPT, Anthropic Claude, Meta LLaMA, Mistral AI	openai.com, anthropic.com, github.com/meta-llama, github.com/mistralai
Voice cloning	ElevenLabs, Resemble AI, open-source Coqui-TTS	elevenlabs.io, resemble.ai, github.com/coqui-ai/TTS
Video deepfake	Synthesia, HeyGen, open-source DeepFaceLab	synthesia.io, heygen.com, github.com/iperov/DeepFaceLab
Smart-contract analysis	Echidna, Foundry, AI-assisted static analysers AnChain.AI SCREEN: Smart contract risk platform AnChain.AI CISO: Crypto investigation AI agent	github.com/crytic/echidna, github.com/foundry-rs/foundry https://www.AnChain.AI/screen https://www.AnChain.AI/ciso
Trading & price manipulation bots	Flashbots for MEV, Hummingbot for algorithmic trading AnChain.AI MCP server for crypto intel	github.com/flashbots, github.com/hummingbot/hummingbot https://github.com/AnChainAI/aml-mcp

How AnChain.AI helps

AnChain.AI's approach is to turn AI against itself. By building an Artificial Super Intelligence (ASI) platform, the company seeks to outpace criminals who weaponise AI. The following capabilities illustrate how AnChain.AI can protect institutions, regulators and investors against AI-driven fraud:

1. AI-native crypto investigation engine

AnChain.AI ingests data from more than forty blockchains, cross-chain bridges and DeFi protocols. Its graph-based analytics detect suspicious flows, identify mixer and peel-chain patterns, and map relationships between wallets. Behavioral clustering helps expose hidden networks and mule accounts.

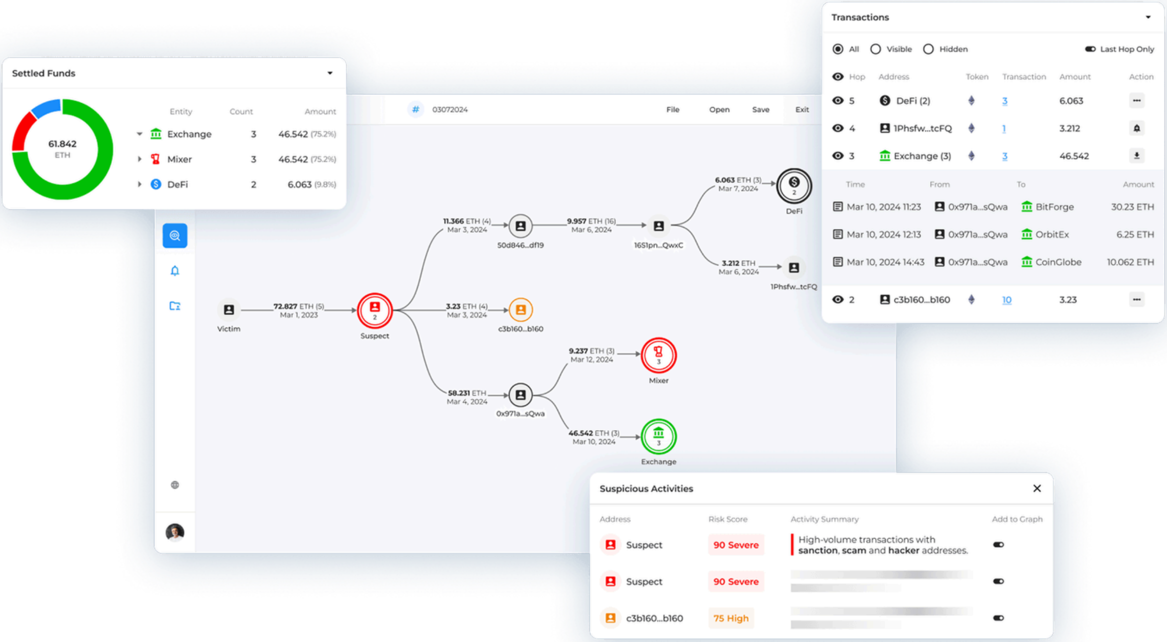

Settled Funds

Entity	Count	Amount
Exchange	3	46.542 (75.2%)
Mixer	3	46.542 (75.2%)
DeFi	2	6.063 (9.8%)

Transactions

● All ○ Visible ○ Hidden — Last Hop Only

Hop	Address	Token	Transaction	Amount	Action
5	DeFi (2)		3	6.063	···
4	1Phsfw...tcFQ		1	3.212	
3	Exchange (3)		3	46.542	
2	c3b160...b160		10	3.23	···

Time	From	To	Amount
Mar 10, 2024 11:23	0x971a...sQwa	BitForge	30.23 ETH
Mar 10, 2024 12:13	0x971a...sQwa	OrbitEx	6.25 ETH
Mar 10, 2024 14:43	0x971a...sQwa	CoinGlobe	10.062 ETH

Suspicious Activities

Address	Risk Score	Activity Summary	Add to Graph
Suspect	90 Severe	High-volume transactions with **sanction**, **scam** and **hacker** addresses.	
Suspect	90 Severe		
c3b160...b160	75 High		

2. Agentic AI for investigations

To keep up with autonomous fraud, AnChain.AI uses large language models as agents to plan investigation steps, summarise on-chain evidence and generate case narratives. The platform anchors these outputs in deterministic blockchain data to avoid hallucinations. Human investigators remain in the loop, with the AI surfacing anomalies and suggested actions.

3. AI-resilient smart-contract risk analysis

Beyond static audits, AnChain.AI models on-chain contract behaviour, monitoring for abnormal interactions that may signal re-entrancy or overflow attempts. Real-time alerts flag suspicious upgrades and admin actions, addressing the kind of attacks seen in the GMX and UXLINK incidents.

4. Real-time threat intelligence & attribution

The platform continuously ingests open-source intelligence, dark-web chatter and law-enforcement advisories. Machine-learning models identify recurring patterns across scams, connecting them to known attacker profiles. This network-level view helps regulators and financial institutions anticipate emerging threats.

5. Explainable AI & regulatory compliance

AnChain.AI produces audit-ready investigation reports that trace the flow of funds and articulate how AI agents arrived at their conclusions. These explainable outputs are critical for regulatory filings, court proceedings and cross-border cooperation.

Recommendations

To effectively defend against AI-driven fraud in 2026, the AnChain.ai team recommends that financial institutions, fintech platforms, enterprises, law-enforcement agencies, and policymakers conduct a comprehensive review of their existing AML, CTF, and risk-management systems to identify gaps exposed by AI-enabled threats. As AI-driven fraud increasingly combines social engineering, automation, and on-chain exploitation, traditional control frameworks designed for human-scale attacks are no longer sufficient. Organizations should assess whether their current systems can detect, reason about, and respond to fraud operating at machine speed.

Compliance and security postures vary significantly across organizations and jurisdictions. As a result, remediation strategies should be tailored to local regulatory requirements, operational risk profiles, and threat environments.

The AnChain.ai team of experts welcomes the opportunity to support this assessment and offers a complimentary consultation to help organizations evaluate AI-related risk exposure and define a forward-looking defense strategy.

1. Invest in AI-native defence:

Firms must adopt systems that combine deterministic blockchain analysis with generative AI reasoning. Traditional rule-based AML tools cannot keep up with autonomous fraud.

2. Harden human processes:

Multi-factor authentication, segregation of duties and security training remain vital. Employees should be educated about AI-generated phishing and deepfakes.

3. Monitor third-party risk:

The SwissBorg incident shows that supply-chain integrations are a major attack vector. Rigorous review and continuous monitoring of API interactions are essential.

4. Participate in threat-intelligence sharing:

Cross-industry collaboration helps map AI-fraud patterns and enables faster responses. Platforms like AnChain.AI facilitate this sharing while safeguarding sensitive data.

5. Support regulatory innovation:

Regulators should update AML/CTF frameworks to address AI-enabled scams and require disclosure of AI use in compliance systems. Industry players can contribute by providing data and recommendations.

Conclusion

In this AnChain.AI Annual Report, we reviewed the top 10 AI-driven crypto fraud incidents in 2025, totaling approximately $2 billion in losses, confirms that modern financial crime is now technically AI-orchestrated. Large Language Models were involved in 7 of 10 cases, contributing to roughly $1.8 billion in losses by enabling scalable phishing, approval manipulation, and deception at near-zero marginal cost. AI-assisted code analysis enabled exploit discovery in 5 cases (≈$1.79B), while voice and video deepfakes, used in 4 and 3 cases respectively, amplified trust exploitation in incidents exceeding $1.5B. These overlapping figures demonstrate that each AI capability represents a hundreds-of-millions-dollar systemic risk vector, demanding Super Intelligent, machine-speed defense systems.

About AnChain.AI

AnChain.AI: Building Super Intelligence to fight fraud and streamline compliance.

AnChain.AI (HQ in Silicon Valley) is an award winning Agentic AI company specializing in fraud investigation and AML compliance. Trusted by global regulators, law enforcement, and financial institutions including the U.S. SEC, IRS, and FinCEN, AnChain.AI's software has underpinned over $1 billion in seizures and recoveries worldwide.
Founded in 2018, AnChain.AI is backed by top venture capital firms from both Silicon Valley and Wall Street, and was recognized as a 2025 Fintech Innovation Lab NYC Finalist, 2023 RSAC Innovation Sandbox Finalist. Most Popular Security app on Metamask.
AnChain.AI serves as expert witness to cryptocurrency cases, and primary investigators for high profile DeFI cryptocurrency cases.

Contact AnChain.AI

📅 Schedule a Meeting: https://anchain.ai/demo
🔗 LinkedIn: https://www.linkedin.com/company/anchainai/
🐦 X (Twitter): https://x.com/anchainai
✉️ Email: info@anchain.ai

Legal Disclaimer

This report is provided by AnChain.AI Inc. for informational and educational purposes only. The contents of this report do not constitute legal, financial, investment, compliance, or professional advice, and should not be relied upon as such.

While AnChain.AI has made reasonable efforts to ensure the accuracy and reliability of the information contained herein, all information is provided “as is” without warranty of any kind, whether express or implied. AnChain.AI makes no representations or warranties regarding the completeness, accuracy, timeliness, or fitness for any particular purpose of the information, analyses, estimates, or opinions expressed in this report.

To the maximum extent permitted by applicable law, AnChain.AI disclaims all liability for any loss, damage, cost, or expense arising directly or indirectly from the use of, or reliance on, this report or any information contained herein.

Any references to third-party incidents, technologies, organizations, or public reports are based on publicly available information and are included for analytical purposes only. Such references do not imply endorsement, affiliation, or verification by AnChain.AI.

The methodologies, estimates, and conclusions in this report reflect the views of AnChain.AI at the time of publication and are subject to change without notice.