


A photograph of two men in an office setting. The man in the foreground, wearing a red and black plaid shirt over a white t-shirt, is smiling and looking at a laptop. The second man, seen from the side and slightly out of focus, is also looking at the laptop. The laptop is open on a dark desk. A large blue diagonal graphic element is on the right side of the image.

AXIS

The Intersection of AI and Crypto

Risk and Insurance Considerations



This paper explores the convergence of artificial intelligence and blockchain technologies—an emerging domain reshaping digital infrastructure, value creation, and risk. Drawing on a16z's taxonomy of 11 use cases, we assess the risk implications, insurance gaps, and market maturity of each, with a particular focus on the insurability of AI-powered agents and decentralized compute systems. Intended for underwriters, brokers, and innovators operating at the frontier of tech risk, the paper provides a framework for underwriting strategy and highlights both coverage opportunities and structural limitations.

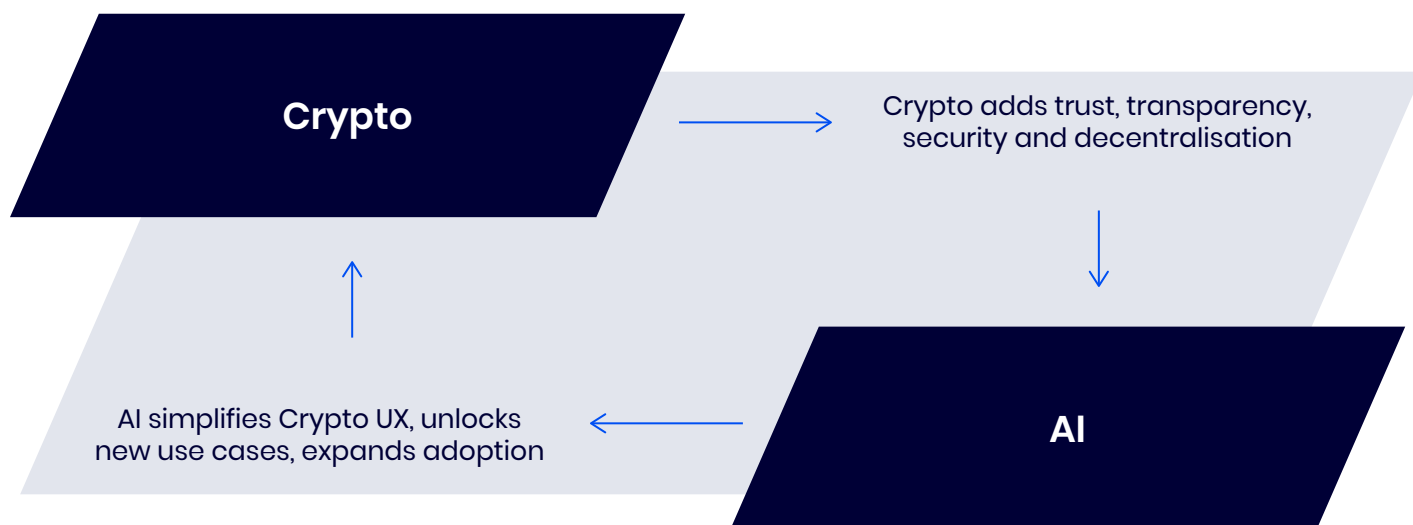
Introduction: The AI x Crypto Convergence

As artificial intelligence (AI) rapidly evolves to occupy core roles in business operations, infrastructure, and consumer applications, the convergence of AI and blockchain-based systems ("crypto") has emerged as a defining technological frontier. We believe this intersection is not only accelerating but will become a foundational layer of future digital infrastructure, creating a wave of opportunity for technologists, investors, and insurers alike.

While many of the underlying technologies are still maturing, there is strong momentum behind the development of AI agents that can autonomously interact with decentralized applications, data ecosystems, and cryptographic primitives. From distributed compute marketplaces to personalized, on-chain AI companions, this fusion is spawning entirely new models of value creation, coordination, and risk.

In this context, we are contextualizing and expanding on a list of 11 use cases for AI x Crypto convergence that were originally shared, without elaboration, by the **a16z** crypto team (Andreessen Horowitz). While a16z provided a concise taxonomy, this paper explores the real-world implications, risks, and insurability of each use case in detail.

Crypto & AI are symbiotic



Use Case Framework

To support analysis, each use case has been grouped into one of four categories, based on its functional focus and primary risk implications:

Identity



These use cases aim to establish verifiable, portable identities for users and AI agents across decentralized systems. They address how AI systems retain continuity of knowledge, distinguish between humans and bots, and participate in systems with provable reputation or credentials.

Decentralized Infrastructure



This category encompasses the physical and protocol-level systems that support AI processing, interaction, and collaboration. It includes distributed compute networks, smart contract-based guardrails, and blockchain-based code synchronization.

Incentive Models



Focused on new economic frameworks, these use cases explore how AI-generated value (e.g., content, insights, labor) is rewarded, how provenance is tracked, and how compensation flows through blockchain systems in fair and automated ways.

Owning AI



This final category envisions a future where individuals or organizations directly own, govern, and deploy their own AI agents. These agents may act semi-autonomously and are governed by decentralized architectures that respect user intent, data ownership, and goal alignment.

Structured Summary: Risk, Coverage, and Market Outlook

The following section presents a structured summary of each a16z use case, mapped across four key dimensions: investable trends, insurable risks, current availability of coverage, and market maturity. This format is intended to support decision-making across insurance coverage considerations, underwriting, product development, and strategic investment contexts.

	USE CASE (a16Z)	INVESTABLE TREND	INSURANCE RISKS	COVERAGE AVAILABILITY NOTES	MATURITY OUTLOOK
IDENTITY	Persistent data/context	Web3 identity platforms	Data privacy, cross-platform liability	Silent coverage possible under cyber/privacy liability	Early-stage
	Universal agent ID	Agent-authenticated smart contracts	Agent impersonation, fraud	Coverage emerging; some digital ID-focused MGAs	Pre-commercial
	Proof-of-personhood	Biometric credential platforms	Biometric misuse, liability	Manuscripted biometric coverage in niche markets	Mid-stage
DECENTRALIZED INFRASTRUCTURE	DePIN for AI compute	Decentralized GPU marketplaces	System downtime, model leakage	Available via ReIm, Nayms; performance-based triggers	Growth phase
	Agent interaction guardrails	On-chain agent governance	Misuse by AI agents, liability assignment	Conceptual only; bespoke clauses in development	Conceptual
	Live AI code sync	Decentralized Git/CodeOps	Code poisoning, IP disputes	Available with vetting; standard IP/media markets	Pre-seed
INCENTIVE MODELS	Micropayments/revenue sharing	Tokenized creator economies	Smart contract failures, payout disputes	Emerging; parametric smart contract cover available	Emerging
	IP provenance for AI	Blockchain IP registries	Attribution errors, content licensing	Standard in IP E&O; increasingly relevant	Early-stage
	Crawler payment systems	Consent-based AI crawlers	Retroactive LLM training claims	Reinsurer concern; not widely addressed yet	Very early
	Web Crawlers with Compensation Models	ZK Ads & privacy targeting	Targeting errors, privacy breach	ZK-specific risks not standardized; adtech E&O possible	Emerging
OWNING AI	User-owned AI companions	Autonomous personal AI agents	ZK system failure, targeting errors, privacy breach	Specialty E&O and cyber clauses under development	Conceptual, 2026+
			Malicious agent behavior, agent fidelity, AI liability	Coverage depends on privacy architecture; may require manuscript	

Insurable Risks and Current Coverage Landscape

In addition to the specific risks described below, there are broader structural and market-wide considerations that impact the availability and reliability of insurance coverage for AI x Crypto use cases:

Silent Coverage and Exclusionary Drift

Many existing technology E&O and cyber liability policies were not written with AI or autonomous agents in mind. As a result, coverage may exist implicitly (“silent coverage”) but is increasingly at risk of being narrowed through the addition of exclusionary language in renewals. This includes exclusions for autonomous systems, non-human actors, or large language model (LLM)-driven automation.

Crypto Market Aversion and Blanket Exclusions

A large portion of the mainstream insurance market (including global carriers) maintains explicit exclusions for cryptocurrency-related activities. This aversion can limit coverage even for AI use cases that are only partially crypto-integrated (e.g., blockchain-based agent coordination or tokenized incentives). Clients should be aware of broad crypto exclusions in their policies that may unintentionally restrict access to coverage.

AI Warranty Coverage (e.g., Armilla)

New warranty products are emerging to address the performance, safety, and bias risks of AI models. Armilla AI, for instance, provides a form of coverage that guarantees model behavior against predefined thresholds. While not a substitute for liability insurance, AI warranties can play a complementary role, especially for enterprises seeking third-party validation and contractual assurance.



Below is a detailed listing of the insurable risks associated with each major use case, stated explicitly and with commentary on whether coverage is currently available - including in specialty or niche markets:

Data privacy, cross-platform liability

Context: AI agents relying on persistent user data stored across decentralized systems.

Coverage: Available via cyber and privacy liability policies. Web3-specific endorsements may be required.

Agent impersonation, fraud

Context: Spoofed or hijacked AI agents acting under false credentials.

Coverage: Available through cyber liability or digital identity insurance. Some underwriters extend this to agent-based authentication layers.

Biometric misuse, liability

Context: Proof-of-personhood protocols involving iris scans or facial recognition.

Coverage: Available only in limited form; biometric-specific claims are often excluded unless manuscripted.

System downtime, model leakage

Context: DePINs experiencing service interruption or leaking proprietary model data.

Coverage: Specialty underwriters (e.g., specialty insurers) offer emerging forms of compute or performance-related coverage.

Misuse by AI agents, liability assignment

Context: Autonomous AI agents executing unauthorized transactions or actions.

Coverage: Not widely available but conceptually covered by AI-specific E&O or bespoke “agent fidelity” clauses under development.

Code poisoning, IP disputes

Context: AI models trained on tainted, misattributed, or unlicensed datasets.

Coverage: Available via IP liability and media E&O policies, often requiring robust vetting of training data.

Smart contract failures, payout disputes

Context: Automated revenue sharing or micropayments failing due to bugs or oracle failures.

Coverage: Smart contract performance insurance is emerging, typically with parametric triggers.

Attribution errors,
content licensing

Context: Failure to credit or compensate rightful content creators.
Coverage: Covered in IP/media E&O insurance. High relevance in creative, legal, and AI art sectors.

Retroactive LLM
training claims

Context: Legal challenges from content owners whose works were used to train AI models without consent.
Coverage: Growing concern for reinsurers; few policies address this directly but it is being reviewed under copyright/IP frameworks.

Targeting errors,
privacy breach

Context: ZK-ad tech failing to protect user identity during AI-driven targeting. ZK enables ad targeting to occur without the ad network ever seeing sensitive personal information. While privacy-preserving in theory, ZK systems also introduce unique risk profiles, including the possibility of inference attacks, failure of cryptographic implementation, or misuse of ZK credentials—all of which could carry both technical and legal liability
Coverage: Covered under privacy liability and ad tech E&O, but not yet standardized for ZK-specific failure modes.

Malicious agent
behavior,
AI liability

Context: User-owned or third-party AI agents acting harmfully or beyond scope.
Coverage: Specialty E&O or bespoke cyber extensions under exploration in Bermuda/London markets.



Non-Insurable Risks at the AI x Crypto Frontier

Not all risks emerging from the AI and crypto convergence can be effectively mitigated through traditional insurance instruments. Several structural and systemic risks fall outside the scope of current insurability frameworks:

Regulatory Arbitrage and Legal Ambiguity

The pace of innovation often outstrips regulatory clarity. Projects operating in legal gray zones (e.g., DAO-governed AI agents, decentralized model marketplaces) may face sudden legal action or jurisdictional bans that are not insurable.

Protocol Governance Failures

Decentralized infrastructure depends on token-based governance, which may be captured by malicious actors or become gridlocked, undermining security or roadmap execution without triggering insurable events.

AI Value Misalignment and Runaway Behavior

Autonomous agents, especially those capable of learning and evolving, may pursue goals that diverge from user intent or market expectations, posing existential platform risk without direct attribution.

Data Sovereignty and Ethical Conflicts

AI models trained on globally distributed data may face conflicting jurisdictional rules on copyright, consent, and data sovereignty, exposing projects to unenforceable or conflicting claims.

Social and Economic Externalities

Tokenized incentive systems may amplify inequalities, misinformation, or speculation-driven volatility, leading to broader systemic impacts that cannot be compartmentalized within conventional risk pools.

Conclusion

The fusion of AI and crypto is not a hypothetical horizon-it is materializing now across infrastructure, markets, and governance. While it unlocks compelling efficiencies and user empowerment, it also necessitates a complete rethinking of how risk is modeled, underwritten, and transferred. From the emergence of AI-native insurance risks to the expansion of decentralized compute economies, the underwriting lens must evolve in tandem. The 11 a16z use cases serve as a blueprint for where that future is being built-and where risk professionals must prepare to operate.



Next step?

Let's engage in dynamic discussion around the factors and variables that are unique to your business so we can obtain the best product, from reliable and stable insurers to meet your needs.

Axis Insurance Managers

#400 – 555 Burrard Street Box 275 – Bentall Two
Vancouver, British Columbia
Canada V7X 1M8

Phone: 604-731-5328

Toll-Free: 1 800-684-1911

www.axisinsurance.ca

Technology team



Clive Bird

Clive is an insurance risk specialist, investor, entrepreneur, and product developer for bespoke insurance risks. For over 15 years Axis Insurance enjoyed a reputation for quality, innovation, creativity and relationship building.

Clive began his Insurance career at Lloyd's of London, renowned for its technical underwriting expertise and a creative approach to risk, providing him with opportunities to push the boundaries of product innovation. Moving to the West Coast of Canada, he has expanded his broad Insurance knowledge and London market relationships to produce products for Canadian clients working across the globe.

Clive has worked extensively with public companies with a strong focus on mining and mineral exploration industry, addressing the broad range of risk exposures they face worldwide. Dynamic business enterprises are starved for the innovation and market relationships Clive can deliver. As an entrepreneurially-minded investor himself, he is embracing new technology and the shifting business landscape to stimulate new market capacity and technically efficient insurance products for the new business era.



Chris Jones

Hello, I'm Chris Jones, an Account Executive specializing in Technology here at Axis Insurance. With over 17 years in the insurance industry, I joined Axis Insurance in 2011, bringing a wealth of experience and knowledge to the table.

My expertise lies in managing technical risks, particularly in the ever-evolving technology industry. Of specific interest are software companies integrating emerging technologies such as blockchain and artificial intelligence into their offerings. Throughout my career, I have honed my skills to provide tailored insurance solutions that meet the unique needs of clients in these dynamic fields.

I am passionate about enabling clients to embrace risk as a catalyst for growth and success. By offering strategic insurance solutions that empower businesses to navigate uncertainties with confidence, I embody Axis Insurance's vision of shaping success through risk management.

Beyond the office, you can often find me unwinding on the golf course.

Glossary

AI x Crypto Use Cases

Persistent Data and Context in AI Interactions

A method for enabling AI agents to retain user-specific preferences, history, and learning across decentralized environments, ensuring continuity of service and personalization without centralized data silos.

Universal Identity for Agents

A framework that gives AI agents cryptographically verifiable, portable identities, allowing them to authenticate and operate autonomously across blockchain-based systems.

Forwards-Compatible Proof-of-Personhood

Blockchain-based tools that verify a user is human (not an AI), designed to remain valid and effective even as synthetic actors and bots become more sophisticated.

Decentralized Physical Infrastructure (DePIN) for AI

Community-powered networks that provide compute, storage, and bandwidth for AI systems, reducing reliance on centralized cloud providers and enabling distributed training and inference.

Infrastructure and Guardrails for Agent Interactions

Smart contract-based rules and permissions that govern how AI agents interact with each other, users, and services, reducing the risk of unintended or malicious behaviors.

Keeping AI/Live-Coding Apps in Sync

Using blockchain as a shared source of truth for real-time updates to AI-generated or AI-assisted code, ensuring version control and collaborative consistency across contributors.

Micropayments That Support Revenue Sharing

Systems that allow fine-grained, automated payments (e.g., per interaction or download) to compensate creators or developers for AI-generated outputs or participation.

Blockchains as a Registry for Intellectual Property and Provenance

Tools to log and verify the origin, ownership, and licensing of data or content used in AI training or output, supporting transparency and rights enforcement.

Webcrawlers That Help Compensate Content Creators

Crawling tools that pay or log compensation for the content they scrape and use to train AI models, aligning AI data collection with copyright and content creator incentives.

Privacy-Preserving Ads That Are Tailored, Not Creepy	Ad delivery systems using zero-knowledge proofs or other cryptographic techniques to personalize ads without exposing personal user data to the ad network or other parties.
AI Companions, Owned and Controlled by Humans	On-chain AI agents personally owned by users, trained on their data and governed by their goals, enabling autonomy without reliance on centralized control or platforms.

Technical / Acronym-Based Terms

LLM (Large Language Model)	A type of artificial intelligence trained on vast text datasets to generate human-like language and perform tasks like summarization, translation, and question answering.
E&O (Errors and Omissions Insurance)	A form of professional liability insurance covering legal costs and damages from mistakes or failures in professional services.
ZK (Zero-Knowledge Proof)	A cryptographic method that proves a statement is true without revealing the underlying information, often used in privacy-preserving systems.
MGA (Managing General Agent)	A specialized insurance agent or broker with authority to underwrite and bind policies on behalf of insurers.
DAO (Decentralized Autonomous Organization)	An organization governed by smart contracts and token-holder voting, rather than traditional corporate structures.
DePIN (Decentralized Physical Infrastructure Networks)	Blockchain-based systems that coordinate physical infrastructure like compute, storage, or bandwidth using crypto incentives.
IP (Intellectual Property)	Legal rights resulting from intellectual activity in the industrial, scientific, literary, or artistic fields.
Web3	A vision of the internet powered by blockchain and decentralized technologies, emphasizing user ownership and privacy.
Parametric Insurance	Insurance that pays out based on a predefined trigger event (e.g., smart contract failure or service downtime) rather than after a traditional claims adjustment.

Conceptual Risk and Insurance Terms

Silent Coverage	Unintended coverage that exists in a policy because the risk was not contemplated or explicitly excluded at the time of underwriting.
Exclusionary Drift	The gradual narrowing of policy coverage over time as insurers introduce new exclusions, often in response to emerging technologies or claim trends.
Agent Fidelity	The extent to which an autonomous AI agent acts in alignment with its user's intentions or constraints.
Malicious Agent Behavior	Actions taken by AI agents that are harmful, deceptive, or otherwise unintended by their human operator or system designer.
Regulatory Arbitrage	Exploiting differences in legal frameworks across jurisdictions to avoid regulation or oversight.
Inference Attack	A method by which an attacker deduces sensitive information by analyzing the outputs of an AI or cryptographic system.
Manuscripted Coverage	Customized insurance policy language negotiated between underwriters and insureds, often used for novel or complex risks.

Infrastructure & Governance Terms

Smart Contracts	Self-executing code stored on a blockchain that automates agreement enforcement without the need for intermediaries.
Token-Based Governance	Decision-making structures where voting power is allocated based on ownership of a blockchain token.
On-Chain	Activities, transactions, or data that are recorded directly on a blockchain ledger.
Off-Chain	Activities or data that are processed outside the blockchain and then optionally recorded on-chain.



axisinsurance.ca