Report — AI Agents: builders, how they work, impact, and ties to the cryptocurrency industry

This report explains what **AI agents** are, how they're built, who's building them, practical use cases (including in crypto), the economic and safety impacts, and pragmatic recommendations for builders, investors and policy makers.

Executive summary

AI agents are software systems that **act autonomously to achieve goals** by planning, using tools, interacting with environments, and learning. Over the last 2–3 years "agentic" systems moved from research demos to production platforms: major AI vendors (OpenAI, Anthropic, Google/DeepMind, Microsoft) and many startups (Replit, Adept, Inflection, Mistral, and specialist toolmakers like LangChain) now offer agent frameworks, development platforms and orchestration tooling. At the same time, a distinct set of blockchain-native projects (e.g., **Fetch.ai**, **SingularityNET**, **Ocean Protocol**, **Alethea AI**) are designing agent-like services that run with token incentives, decentralized marketplaces, or data-privacy primitives — creating a real intersection between agents and crypto. (Reuters)

1. What is an "AI agent"?

Definition (practical): an AI agent is software that perceives inputs (text, APIs, sensors, on-chain state), **plans** a sequence of actions toward a goal, **executes** actions (calls APIs, submits transactions, changes state), and **adjusts** behavior using feedback or memory. Agents range from lightweight assistants that schedule calendar events to fully autonomous multi-step agents that research, plan, and execute complex tasks with little human supervision. Common properties include goal orientation, tool use, memory, planning, and sometimes multi-agent coordination.

Taxonomy (high level):

- **Assistant agents:** help humans (e.g., booking, summarizing).

- **Autonomous / goal-driven agents:** act to complete a defined objective with minimal supervision (e.g., AutoGPT variants).
- **Specialized agents:** domain-specific (code agents, trading agents, legal-tech agents).
- **Multi-agent systems:** many agents coordinate/compete to accomplish joint tasks.

(These categories help match architectures and safety controls to intended uses.) (Unity Communications)

2. Core technical architecture (how agents are built)

Agents are assembled from modular components:

1. **Base model / reasoning core:** an LLM or multimodal model provides language understanding, planning, and reasoning. Large vendors supply these models (OpenAI, Anthropic, Google; and many open models). (Reuters)
2. **Planner / chain-of-thought subsystem:** transforms goals into stepwise plans (task decomposition and subtask scheduling).
3. **Tooling layer (tool use):** connectors or "tools" agents call — web search, databases, task APIs, on-chain nodes, wallets, oracles, etc. Tool use enables grounding in trusted sources or real-world effect.
4. **Memory & state:** short/long term memory stores facts, context, and prior actions. Memory enables persistent agents that improve performance over time.
5. **Execution & orchestration:** runtime that runs the plan, dispatches tools, manages retries and error handling. Enterprise platforms add auditing and RBAC.
6. **Safety & monitoring:** authentication, action approvals, sandboxing, and human-in-the-loop gates.
7. **Observability & logging:** for auditability, debugging, and compliance.

Frameworks such as agent SDKs and orchestration platforms (LangChain ecosystem and many commercial equivalents) standardize these pieces; cloud vendors and enterprise software (GitHub Agent HQ, Salesforce Agentforce) embed agent orchestration into developer and business workflows. (The Verge)

3. Who's building AI agents? (leading companies & projects)

Major AI vendors (platform & model providers)

- **OpenAI** — large multimodal models and "agent" style features / operator affordances used to build autonomous workflows.
- **Anthropic** — Claude models and safety-focused agent products for enterprise use.
- **Google / DeepMind** — Gemini suite and internal agent frameworks for search, assistant and developer tooling.
- **Microsoft** — Copilot + integrations across Office, Azure, GitHub (supports agent workflows).

These companies provide the foundational models, platform APIs, and large-scale compute infrastructure that most agents depend on. (Reuters)

Enterprise & developer platform vendors

- **GitHub (Agent HQ)** — centralizes multi-agent coding workflows and integrates multiple model providers for devs. (The Verge)
- **Salesforce (Agentforce 360)** — enterprise agent platform integrated with business data and apps. (Reuters)
- **Replit, Adept, Inflection, Mistral, Hugging Face, LangChain (ecosystem)** — provide agent SDKs, orchestration, model hosting and tooling stacks for rapid agent development. (eWeek)

Blockchain-native / crypto projects building agentic systems

- **Fetch.ai** — builds autonomous economic agents and an agent framework designed to operate on decentralized infrastructure (agents negotiate, transact and coordinate services). This is explicitly agent-centric and pairs agents with blockchain incentives. (Platform to enable the agentic economy.)
- **SingularityNET (AGIX)** — a decentralized marketplace for AI services where agents and models are discovered, paid for, and composed on-chain. It aims to decentralize AI supply chains and let developers monetize services. (SingularityNET)
- **Ocean Protocol** — data marketplace and compute-to-data primitives that enable secure data provisioning for AI training and agent execution while preserving

privacy; Ocean's stack is often paired with agent frameworks to monetize data feeds. (Ocean Protocol)

- **Alethea AI** — creator of iNFTs and interactive crypto-native AI characters/agents (agentic NFTs that can interact, be owned, and evolve). (Alethea AI Labs)
- **Chainlink** — not an agent vendor per se, but Chainlink's oracle networks and integrations are critical tooling that lets agents access reliable off-chain data and model outputs inside smart contracts. Chainlink has explicitly described oracle–AI integrations. (Chainlink Blog)

(There are many startups also blending agent tech with Web3 primitives — the landscape is fast-moving.) (Marketer Milk)

4. Key use cases and examples where AI agents intersect with crypto

A. Autonomous market participants & DeFi automation

Agents can operate trading strategies, run liquidity provision, and perform cross-chain arbitrage. They monitor on-chain events, react to price movements, and execute transactions automatically. This accelerates market efficiency but also raises MEV, front-running and fairness concerns.

B. On-chain or hybrid agents for smart contract automation

Agents can watch contract state and call contract functions (e.g., execute liquidations, rebalance vaults, or manage positions). When paired with oracles and secure signing, these agents automate traditionally manual operations.

C. Decentralized AI marketplaces & data provisioning

Projects like **SingularityNET** and **Ocean Protocol** let agents discover, purchase and use AI models or datasets in a tokenized marketplace — enabling permissioned compute, paid model calls, and privacy-preserving compute-to-data patterns. This supports agent workflows that require proprietary data or paid model services. (SingularityNET)

D. Tokenized agent economies and iNFTs

Alethea and similar projects mint interactive, agentic NFTs (iNFTs) which can hold state, dialog, and simple agent policies. Owners can monetize agent behavior, license capabilities, or run them inside virtual worlds. This creates new digital-goods economics tied to agent capabilities. (Alethea AI Labs)

E. Oracles + agents: bridging model outputs to contracts

Agents that infer or compute off-chain values (market forecasts, risk signals, identity verification) can deliver signed outputs via oracle networks (e.g., Chainlink). This lets smart contracts act on AI insights in a verifiable way. (Chainlink Blog)

F. Privacy-preserving agent training & compute marketplaces

Ocean's compute-to-data architecture lets agents train or run inference on private datasets without exposing raw data, creating data monetization models for organizations that can't share raw datasets. This is valuable for regulated industries (healthcare, finance) where agents must learn but maintain privacy. (Ocean Protocol)

5. Economic, technical and social impacts

Economic & product impacts

- **Acceleration of automation:** agents reduce human cost for routine and semi-structured work (developer ops, customer support, trading).
- **New revenue models:** tokenized AI services, paid agent tasks, and iNFT economies create monetizable agent-driven micro-transactions. (CoinMarketCap)
- **Data monetization:** data marketplaces enable data owners to monetize training data while preserving privacy, shifting the data economy. (Ocean Protocol)

Technical & security impacts

- **Increased attack surface:** agents that control keys, execute transactions, or call external systems create new vectors for exploitation (compromised agent, poisoned data, tool abuse).
- **Oracle dependency:** agents that feed contracts through oracles depend on oracle integrity; attacks on oracle feeds or model supply chains can cascade on-chain. (Chainlink Blog)

Social & governance impacts

- **Workforce change:** agent automation will reshape jobs (higher-order supervision, prompt engineering, agent orchestration roles).

- **Decentralized governance challenges:** tokenized agent systems raise questions about ownership, liability, and how agent policies are updated — particularly if agents act autonomously in economic systems.

6. Risks and failure modes (special attention for crypto integration)

- **Misaligned goals & autonomy drift:** poorly specified goals can lead agents to harmful or unexpected behaviors (e.g., repeated on-chain spam or irrational trading that harms liquidity).
- **Oracle & data poisoning attacks:** bad inputs can make agents produce harmful actions; in DeFi this can cause mispriced liquidations or exploits. (Chainlink Blog)
- **Key/custody compromise:** agents often need signing keys to act on-chain — compromised keys allow theft.
- **Regulatory & compliance risk:** autonomous agents executing financial transactions cross borders and may run afoul of AML/KYC rules if not designed with compliance controls.
- **Sybil & economic attacks against decentralized agent markets:** token-based agent marketplaces can be manipulated by concentrated token holders or by synthetic data submissions.

Mitigations: layered safety (human-in-the-loop for high-risk actions), formal verification for critical contracts, robust oracle networks, key-management best practices, economic bonding/staking to penalize bad behavior, standardized audit trails and immutable action logs.

7. Governance, standards & policy considerations

- **Standards for agent provenance and audit trails:** machine-readable logs of agent plans, tool calls, and outputs should be standardized for auditability.
- **API & oracle standards:** common interfaces for agent→oracle→contract interactions reduce bespoke integration risk. Chainlink and other oracle providers are already publishing guidance on integrating AI outputs with smart contracts. (Chainlink Blog)

- **Regulatory sandboxes & testbeds:** governments and regulators should encourage sandboxes where agents can be stress-tested against AML, consumer protection and market-integrity rules.
- **Liability frameworks:** clarify where legal responsibility lies when an autonomous agent causes harm (developer, deployer, owner, or a tokenized governance collective).

8. Practical recommendations

For builders

- Start with **least-privilege** agent principals (narrow permissions, approval gates).
- Use **provable oracles** and signed attestations for data the agent relies on. (Chainlink Blog)
- Log every tool call and transaction in an auditable format; maintain immutable action hashes on-chain for post-hoc verification.
- Design human-in-the-loop checkpoints for high-value or irreversible actions (large transfers, governance votes).

For token projects / DAOs

- Build economic incentives and slashing mechanisms to deter malicious agent behavior.
- Vet agents before granting treasury access — use staged access and canary deployments.

For regulators & institutions

- Encourage interoperable reporting standards for agent-driven financial activity.
- Fund public testbeds that measure how agents interact with markets and critical infra (oracles, large DEXs).

AI agents are rapidly becoming a foundational abstraction in software: they combine reasoning, tool-use and autonomy to do work on behalf of humans. When paired with blockchain they enable **new economic models** (paid agent services, tokenized agent economies), **automation of market actions** (DeFi agentic operations), and **data**

monetization (compute-to-data, private training). That opportunity comes with substantial.

Systemic risks — from oracle attacks and key compromise to regulatory uncertainty — so the next 12–24 months will be shaped by how vendors, blockchain projects and regulators operationalize safety, auditability and economic governance.

Citation

1. **Chainlink Blog.** (2024). *The Intersection Between AI Models and Oracles.* Chainlink. (Chainlink Blog)
2. **Fetch.ai.** (n.d.). *Fetch.ai — Autonomous Economic Agents & Use Cases.* Fetch.ai website. (Platform to enable the agentic economy.)
3. **SingularityNET.** (n.d.). *SingularityNET — Decentralized AI Marketplace.* SingularityNET website. (SingularityNET)
4. **Ocean Protocol.** (n.d.). *Ocean Protocol — Tokenized AI & Data Marketplaces.* Ocean Protocol website. (Ocean Protocol)

5. **Alethea AI.** (n.d.). *About Alethea AI — intelligent NFTs and agentic characters*. Alethea.ai. (Alethea AI Labs)
6. **GitHub / The Verge.** (2025). *GitHub launches Agent HQ for multiple AI coding agents*. The Verge coverage of GitHub Agent HQ. (The Verge)
7. **Reuters.** (2025). *Salesforce deepens AI ties with OpenAI, Anthropic to power Agentforce platform*. Reuters reporting on enterprise agent platforms. (Reuters)
8. **Market & industry overviews (2025 lists).** eWeek / Forbes / curated lists of AI & agent platforms. (eWeek)