4QUAYS CRYPTO AGILITY FOR AI VERIFICATION

Accountable AI you can prove.

AI agents now approve refunds, move data and talk to customers on your behalf. 4Quays gives every agent a verifiable identity and **proves what it said and did,** so you can scale AI without losing control, accountability or the trust of regulators and customers.

Less fraud exposure | AI Act-ready evidence | No agent redesign

AGENT ACCOUNTABILITY	TODAY	WITH 4QUAYS
Who acted	Guesswork	Agent-signed
Altered content	Goes unnoticed	Fails verification
Key approvals	Agent discretion	Human-approved
Proof for others	"Trust us"	Anyone can check
Rogue agent	Hunt for its keys	Revoked instantly

WHAT'S AT STAKE

Agents are making decisions. You need to verify them.

When an agent's output can't be attributed or checked, every approval, payment and customer commitment it touches can become a liability.

15%[1]

Of daily work decisions

Made autonomously by agentic AI by 2028, up from none in 2024.

92%[2]

Lacked AI access controls

Of organizations hit by an AI-related breach in 2026.

3%[3]

Of global turnover

EU AI Act maximum fine for high-risk obligations, which include automatic logging, or €15M if higher.

$25M lost to an impersonation.[4] A finance employee at Arup wired HK$200M after a video call with a deepfaked CFO. When anyone can be made to look and sound like someone else, only an irrevocable and immutable signature proves identity, human or agent.

WHAT THE BUSINESS GETS

Scale AI. Answer for every action.

4Quays makes agent accountability a property of your platform, not a promise each AI project has to keep on its own.

01

Reduce fraud

Impersonated or tampered instructions fail verification before anyone acts on them.

02

Humans in charge

High-value actions need a person to approve the exact content first.

03

Prove compliance

Automatic, tamper-evident records of what each agent did, ready for auditors.

04

Earn trust

Customers and partners verify your agents' work themselves, no account needed.

1

Identity per agent

Every agent answerable under one enterprise policy.

0

Agents redesigned

Works with the AI tools and vendors your teams already use.

100%

Of agent messages signed

Automatically, every turn, not only when the model remembers.

HOW IT WORKS

Enroll once. Prove every action.

Every agent becomes a cryptographic principal under one enterprise policy, across human-to-AI, AI-to-AI and AI-to-machine traffic.

Observe
Discover agents in Bedrock, OpenAI, Entra and frameworks in source.

Execute
Sign and verify through MCP, SDK and hooks. Keys never enter the chat.

Govern
ML-DSA level, grants, rotation and revocation set by central policy.

Prove
Public verification, every signature audited, evidence in your SIEM.

Agents as cryptographic principals
Each enrolled agent gets its own ML-DSA signing key, and optionally an ML-KEM encryption key, under policy.

Three layers of provenance
MCP gives capability, the Skill judgement, and hooks guarantee every prompt and response is signed.

Portable signed messages
Content travels inside the signature, so it verifies in email, tickets, chat, APIs and agent handoffs.

Public and local verification
Check signer and integrity on the public verify page, by API, or on-device against published keys.

One-time signing grants
A human approves exact content, such as a refund or a wire, before the agent can sign it. Nothing else qualifies.

Rotation and instant revocation
Rotate keys without invalidating past signatures. Revoke an agent and verification reports it immediately.

Common platform

INCLUDED IN EVERY 4QUAYS DEPLOYMENT

Discovery
Cloud KMS, code, network, endpoints and AI agents.

Prioritization
Risk scoring against NIST guidance and vendor readiness.

Monitoring
Latency, error rate, throughput and volume per policy.

Alerting
Threshold, drift and missing-activity rules.

Auditing
Every operation, with policy, algorithm and key.

SIEM & logging
Monitoring, audit and alert data to your SIEM.

TECHNICAL SPECIFICATIONS

Signatures	ML-DSA-44 / 65 / 87 (FIPS 204) · ECDSA (ES256) and RSA-PSS (PS256) for legacy verifiers
Message format	4Quays signed message (JWS, content embedded) · hosted verify links · public JWKS keys
Agent surfaces	Claude Code · Codex · Cursor · Gemini CLI (signing hooks) · any MCP client · TypeScript SDK
SIEM & logging	Splunk · Sentinel · Datadog · Elastic · QRadar · Sumo Logic · Falcon LogScale · webhook
Standards alignment	NIST FIPS 203 / 204 · IETF JOSE (JWS, JWK) · NIS2 · EU AI Act record-keeping (Art. 12)

1. Gartner, *Top Strategic Technology Trends for 2025: Agentic AI*, Oct 2024.
2. IBM Security & Ponemon Institute, *Cost of a Data Breach Report 2026*.
3. Regulation (EU) 2024/1689 (AI Act), Article 99(4); logging duty in Article 12.
4. CNN, "Arup revealed as $25 million deepfake scam victim", May 2024.

SEE IT LIVE

See 4Quays change cryptography in a live environment.

Book a demo

contact@4quays.com
4quays.com