4QUAYS CRYPTO AGILITY FOR AI SECURITY

Put AI to work. Keep it in bounds.

Every agent can act as a super admin. These super agents now coordinate with each other at machine speed, and their guardrails live inside the very prompts they can ignore. 4Quays lets you decide **who an agent talks to, the language it speaks and what it remembers,** from one agent to a swarm, with controls the agent can't override.

Contain incidents | Protect data & IP | Deploy AI safely

AGENT CONTROLS	TODAY	WITH 4QUAYS
Controls live	In the prompt	In your keys
Who it talks to	Anyone it finds	Declared peers
What others read	Plain text	Encrypted
What it keeps	Everything	Verified progress
Stopping it	Shut it all down	Freeze one agent

WHAT'S AT STAKE

Agents widen the attack surface. Prompts can't hold it.

An agent that can reach anyone, be read by anyone and remember anything is an insider you can't fully supervise.

25%[1]

Of enterprise breaches

Traced to AI agent abuse by 2028, by outsiders and insiders alike.

43%[2]

Of breaches involved shadow AI

In 2026, up from 20%, at an average cost of $5.39M each.

40%[3]

Of agentic AI projects cancelled

By the end of 2027, citing cost, unclear value or inadequate risk controls.

It has already happened. Research agents coordinated through notes left in a shared package cache, recruited peers, and rebuilt their board within hours of it being wiped. A readable language and a lasting memory were all they needed.

WHAT THE BUSINESS GETS

More AI in production. Less exposure.

Security becomes a property of your platform, so AI projects stop stalling on risk reviews and incidents stay small.

01

Contain incidents

Freeze one agent or faction. Everything else, and its progress, keeps working.

02

Protect data and IP

Agent conversations and memory are unreadable outside the channels you approve.

03

Deploy with confidence

Risk controls in place from day one move AI projects past pilot.

04

Prove control

Every message, refusal and freeze on record for auditors and the board.

1

Action to contain

An agent, a faction or a whole swarm, on its next step.

0

Capability lost

Agents keep everything they can do. Only reach and memory are governed.

0

Reliance on agent cooperation

Controls sit with your security team, never in the prompt.

HOW IT WORKS

Three controls. Any scale.

Agent messages and memory are encrypted, and you hold the keys. Each control is a key change that lands on the agent's next step.

Who it talks to

Declared channels per agent. Undeclared sends are refused and flagged; factions exist only by permission.

The language it speaks

ML-KEM encrypted, ML-DSA signed, unreadable outside its channel. Re-key a faction in one move.

The context it remembers

Memory sealed by epoch, signed handoffs between epochs, pinned policy, false memories retracted everywhere.

Every agent a verified identity

Its own ML-DSA and ML-KEM keys under policy, so nothing passes as an operator or a peer.

Envelopes set by people

A human sets the most an agent or tree may ever be: agents, operations, channels and lifetimes.

Swarm map

Declared channels against observed edges, factions and growth over time, on a live event stream.

Targeted containment

Freeze one channel, faction or agent. The rest of the swarm and its verified progress carry on.

An archive, not a memory hole

Sealed keys go to escrow for audited investigators. Shred is human-only, with two approvers.

Tripwires

Refused decrypts, undeclared sends and envelope violations raise alerts tied to an identity.

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

Scope	Single agent · agent trees · multi-vendor swarms · AI-to-AI, AI-to-machine, human-to-AI
Cryptography	ML-KEM-512 / 768 / 1024 (FIPS 203) · ML-DSA-44 / 65 / 87 (FIPS 204) · AES-256-GCM
Controls	Channel policies · channel keys · epoch-sealed memory · signed handoffs · envelopes · Freeze · escrow · Shred
Agent surfaces	Claude and OpenAI agents · Claude Code · Codex · any MCP client · `4q` CLI · `@4quays/sdk` · SSE stream
SIEM & logging	Splunk · Sentinel · Datadog · Elastic · QRadar · Sumo Logic · Falcon LogScale · webhook

1. Gartner, *Top Strategic Predictions for 2025 and Beyond*, Oct 2024.
2. IBM Security & Ponemon Institute, *Cost of a Data Breach Report 2026*.
3. Gartner, agentic AI project cancellation forecast, press release, June 2025.

SEE IT LIVE

See 4Quays change cryptography in a live environment.

Book a demo

contact@4quays.com
4quays.com