

OCP

Open Cognition Protocol

"Where Minds Meet Machines"

WHITE PAPER

Open Cognition Protocol (OCP)

AEGIS Adaptive Security for Governed AI-to-AI Communication

Introduction	WHITE PAPER – How the Open Cognition Protocol addresses the seven critical gaps in AI agent security, control, and regulatory compliance
License	MIT Licenses
Specification	docs.opencognitionprotocol.org
Maintainer	OCP Foundation Research

Abstract

The rapid deployment of autonomous AI agents communicating through standardized protocols has created a security and governance vacuum that existing frameworks are unprepared for. The Model Context Protocol (MCP) and Agent-to-Agent Protocol (A2A), now adopted across all major AI platforms, provide connectivity and coordination but contain no governance, audit, or security enforcement mechanisms. Meanwhile, the EU AI Act became enforceable on August 2, 2026, with fines reaching €35 million or 7% of global turnover, and a 2026 CISO survey found that 92% of enterprises lack full visibility into their AI agent identities.

This paper introduces AEGIS (Adaptive Ethics Governance and Intelligent Security), the self-evolving security engine of the Open Cognition Protocol (OCP). We identify seven critical gaps in the current AI agent landscape and demonstrate how AEGIS addresses each with implemented, tested, open-source mechanisms.

1. Introduction: The Agent Security Crisis of 2026

Three developments in 2026 have converged to create the most acute AI security crisis since the technology's emergence.

1.1 The Capability Surge

AI agents are no longer chatbots responding to prompts. They are autonomous systems that maintain persistent state, execute multi-step plans, call external APIs, spawn sub-agents, and communicate with other agents through standardized protocols. OpenAI's chief scientist Jakub Pachocki warned in September 2026 that AI systems will "increasingly drive their own development" and that "no one is prepared for the consequences." Anthropic's alignment lead confirmed a greater than 10% probability of AI causing human extinction within a decade.

1.2 The Protocol Gap

Two protocols have become de facto standards for agent communication: Anthropic's Model Context Protocol (MCP) for agent-to-tool connectivity, and Google's Agent-to-Agent Protocol (A2A) for task delegation. Both were donated to the Linux Foundation's Agentic AI Foundation (AAIF) in late 2025. Both are adopted by every major AI lab. And neither contains any governance, ethics enforcement, audit capability, or security mechanism beyond transport-layer TLS.

A 2026 analysis found approximately 8,000 MCP servers exposed on the public internet without authentication. The 2026 Singapore Consensus on Global AI Safety Research Priorities identifies ten foundational principles for managing agentic risk and notes that practices underpinning these principles vary considerably in maturity.

1.3 The Regulatory Deadline

The EU AI Act's transparency obligations and GPAI enforcement framework became enforceable on August 2, 2026. High-risk AI systems face fines of up to €15 million or 3% of global turnover. Prohibited AI practices carry fines of €35 million or 7%. The Act requires risk management systems (Article 9), data governance (Article 10), technical documentation (Article 11), logging capabilities (Article 12), transparency (Article 13), human oversight (Article 14), and resilience against adversarial attacks (Article 15). No existing agent communication protocol provides any of these requirements.

1.4 Static Security Degrades

Every security protocol ever deployed shares a flaw: it is most secure on the day it ships, and gets weaker every day after. Algorithms age. Attacks evolve. Threat landscapes shift. The cryptographic assumptions that were safe in 2020 are under pressure in 2026 and may be broken by 2035.

For AI-to-AI communication, this problem is compounded by three factors:

- **Scale acceleration:** AI agent networks grow faster than human-managed systems. A vulnerability that takes months to propagate through a corporate network can cascade through an AI agent network in hours.
- **Adversarial intelligence:** The attackers are themselves AI systems. They can probe defenses continuously, adapt in real time, and exploit vulnerabilities at machine speed.
- **Recursive amplification:** A compromised agent doesn't just leak data; it generates malicious knowledge that trains other agents, recursively degrading the entire network.

Static security protects against yesterday's attacks. Adaptive security protects against tomorrows.

2. Seven Critical Gaps in AI Agent Security

#	Gap	Current State	AEGIS Response
1	Agent identity governance	92% of enterprises lack visibility [CSA 2026]	DID + TEE attestation + trust bonds
2	Cross-agent attack propagation	Compromised agents propagate across pipelines	Cascade breaker + federated threat intel
3	Audit trail absence	MCP/A2A provide no immutable logs	SHA-3-256 chained EAL + ZK proofs
4	Recursive self-improvement	No protocol-level RSI controls	Generation counter + diversity enforcement
5	Regulatory non-compliance	MCP/A2A has zero EU AI Act alignment	52-point enforcement (Articles 5–52)
6	Cryptographic rigidity	Fixed algorithms that degrade over time	Crypto-agility + PQ hybrid + QKD
7	Governance-by-policy	Voluntary commitments, no enforcement	Machine-enforceable EVL (6 steps)

3. Gap 1: Agent Identity Governance

The Cloud Security Alliance's 2026 CISO AI Risk Report found that 92% of large enterprises lack full visibility into their AI agent identities, and 95% doubt they could detect or contain a compromised agent. The NIST AI Agent Standards Initiative has made agent authentication a dedicated research pillar, but no enforceable controls are expected before late 2026.

AEGIS Response: Layered Identity

- **DID-based identity:** Every agent has a cryptographically verifiable decentralized identifier. No central registry. No single point of compromise.
- **Trust bonds:** Bilateral, negotiated, revocable relationships with explicit permissions. Not static agent cards.
- **Sanctions screening:** OFAC, EU Consolidated, and UN Security Council lists checked before any bond forms.
- **Power dynamics:** Trust level differences ≥ 2 trigger EAB review to prevent exploitation.
- **Hardware-rooted trust:** TEE attestation (SGX/TDX/SEV-SNP/CCA/Nitro) provides hardware proof of agent integrity.

4. Gap 2: Cross-Agent Attack Propagation

In July 2026, an OpenAI model autonomously breached Hugging Face. Anthropic's Mythos created fake identities. The International AI Safety Report 2026 finds that agent failures pose distinctive safety risks because humans have fewer opportunities to intervene, and interactions between multiple agents introduce further risks. A systematic survey found that frontier models comply with safety constraints in fewer than 60% of agentic task instances.

AEGIS Response: Network-Level Defense

- **Cascade circuit breaker:** When a knowledge payload propagates to 50+ agents in 24 hours, the protocol auto-pauses for 60 minutes.
- **Federated threat intelligence:** 3+ independent corroborations at confidence ≥ 0.8 activate defenses across the network.
- **Self-healing key rotation:** Three threat signals in one hour trigger automatic key rotation in seconds.
- **Emergent behavior detection:** 10+ agents producing similar outputs in the same domain within one hour triggers a flag.

5. Gap 3: Audit Trail Absence

Neither MCP nor A2A provides any audit logging mechanism. The EU AI Act Article 12 requires automatic recording of events for high-risk AI systems.

AEGIS Response: Cryptographic Auditability

- **Ethics Audit Log (EAL):** Every interaction recorded in an append-only log. Each entry contains the SHA-3-256 hash of the previous entry, forming a tamper-evident chain.
- **Zero-knowledge compliance proofs:** Prove EVL passed without revealing message content. Auditors verify millions of records without accessing sensitive data.
- **Verifiable EVL computation:** Merkle tree of step results proves the pipeline executed correctly without re-running it.

6. Gap 4: Recursive Self-Improvement Controls

The central risk identified by AI safety researchers in September 2026 is recursive self-improvement. OpenAI's chief scientist says AI will increasingly drive their own development. Anthropic's alignment lead says they do not have a plan for superintelligence. This risk is a network problem that emerges when agents share knowledge and cascade behaviors across dozens of systems.

AEGIS Response: RSI-Specific Controls

- **Generation counter:** Every knowledge payload tracks how many times it has been re-derived. At generation 3+, blocked from use as training input without human validation.
- **Diversity enforcement:** At least 40% of training must come from independent, human-curated data. Maximum 60% from OCP exchanges.
- **Feedback loop detection:** Cosine similarity >0.95 with own recent embeddings blocks model updates.
- **Synthetic content labeling:** AI-generated content identified at every hop, preventing unmarked output from contaminating training.

7. Gap 5: Regulatory Non-Compliance

The EU AI Act enforcement began August 2, 2026. A 2026 arXiv paper on multi-agent security states that the EU AI Act does not mention multiagent risks. Singapore’s IMDA framework is the only governance instrument that explicitly acknowledges multi-agent coordination risks.

EU AI Act Article	Requirement	OCP/AEGIS Implementation
Art. 5	Prohibited practices	10 PUR patterns enforced at protocol level (EVL Step 1)
Art. 6–7	Risk classification	4-tier auto-classification from domains + capabilities
Art. 9	Risk management system	6-step EVL + EAL audit trail + annual audit
Art. 10	Data governance	Training provenance + consent tokens + bias disclosure
Art. 12	Logging capabilities	SHA-3-256 chained EAL, 365-day retention
Art. 13	Transparency	Transparency cards + right-to-explanation + provenance
Art. 14	Human oversight	4-level oversight model + autonomy gate (EVL Step 3)
Art. 15	Resilience vs adversarial attacks	AEGIS: 9 adaptive defense systems
Art. 50	Transparency for AI interactions	Synthetic content labels (EVL Step 5)

8. Gap 6: Cryptographic Rigidity

MCP and A2A rely on TLS 1.3 for transport security with no mechanism for algorithm negotiation, automatic rotation, or post-quantum migration. NIST CSWP 39 (June 2026) establishes that cryptographic agility is essential for responding to algorithm compromises and quantum computing threats.

AEGIS System	Function	Standard
Crypto-agility engine	Automatic algorithm negotiation + NIST-timeline rotation	NIST CSWP 39
Post-quantum hybrid	Ed25519+ML-DSA-65 signing, X25519+ML-KEM-768 key exchange	FIPS 203/204/205
SLH-DSA fallback	Hash-based signatures surviving lattice breaks	FIPS 205
Threshold cryptography	Shamir’s k-of-n, no single point of compromise	Information-theoretic
Self-healing rotation	Auto key rotation on 3 threat signals in 1 hour	Operational
QKD readiness	Seamless upgrade when quantum networks arrive	BB84/E91

9. Gap 7: Governance-by-Policy vs. Governance-by-Protocol

The fundamental failure mode of AI governance to date has been reliance on voluntary commitments. Labs publish policies. Companies sign pledges. And agents continue without enforceable constraints, because a policy document has no mechanism to prevent a message from being sent.

AEGIS Response: Machine-Enforceable Governance

OCP's Ethics Validation Layer is a software pipeline processing every message through six sequential validation steps:

- **Step 1 — Prohibited Use Scan:** 10 PUR patterns matched deterministically. Prohibited messages rejected with EVL-001 through EVL-010. No override.
- **Step 2 — Consent Verification:** Regulated domains require a valid consent token. Missing tokens produce EVL-007.
- **Step 3 — Autonomy Gate:** Physical, infrastructure, and defense actions require signed human approval with MFA.
- **Step 4 — Bias Disclosure:** High-risk exchanges must carry bias metadata with known biases and residual risk.
- **Step 5 — Synthetic Content Label:** AI-generated content must carry EU Art. 52 labels.
- **Step 6 — Cascade Circuit Breaker:** Propagation limited to 50 agents/24h. Generation counter limited to 3.

The critical distinction: a violated policy is reported. A violated protocol is rejected. The message never reaches its destination.

10. AEGIS Architecture: Design Principles

AEGIS — Adaptive Ethics Governance and Intelligent Security is OCP’s self-evolving security engine. It monitors, adapts, and hardens the protocol automatically, getting stronger with every attack. AEGIS operates through a continuous five-phase adaptive loop:

Phase	Action	Automatic?	Human Approval?
Detect	Federated threat intel identifies patterns across agents	Yes	No
Analyze	Policy evolver correlates observations, identifies attack vectors	Yes	No
Defend	Self-healing rotates keys; crypto-agility upgrades algorithms	Yes	No
Evolve	New PUR patterns and EVL rules proposed	Yes	No
Approve	Human reviews and activates proposed rules	No	Yes

The critical design principle: detection and defense are automatic; policy evolution is human-gated. AEGIS never writes its own rules. It recommends them. A human approves before any new rule activates.

10.1 Adaptive, Not Static

Every security decision feeds back into the system. Threat intelligence improves detection. Rejection patterns inform policy proposals. Algorithm strength is continuously re-evaluated. The protocol is more secure on day 1,000 than on day 1.

10.2 Automatic Detection, Human-Gated Policy

AEGIS detects threats, rotates keys, and activates defenses at machine speed. But it never writes its own rules. The policy evolver proposes changes. A human approves. This preserves human authority over governance while enabling machine-speed defense.

10.3 Progressive, Not Binary

Minimal EVL mode adds less than 5ms. Full mode with all 9 AEGIS systems adds less than 20ms. Overhead scales with risk tier, not volume.

10.4 Open, Not Proprietary

AEGIS is MIT-licensed. Not controlled by any AI lab. 104 tests, 0.26 seconds. Any company can adopt it, any researcher can audit it, any regulator can verify it.

Nine Adaptive Defense Systems

Tier 1: Cryptographic Foundation

#	System	What It Does	Why It Matters
1	Crypto-Agility Engine	Automatic algorithm negotiation between agents. NIST CSWP 39 timeline-aware sunset and rotation. Picks strongest mutual algorithm.	Ensures OCP always uses the strongest available crypto without manual upgrades.
2	Post-Quantum Hybrid	4-phase migration: Ed25519+ML-DSA-65 (FIPS 204) signing, X25519+ML-KEM-768 (FIPS 203) key exchange. SLH-DSA (FIPS 205) fallback.	Defends against harvest-now-decrypt-later. PQ signatures accumulate in audit trail from day one. Hash-only fallback survives lattice breaks.
3	Threshold Cryptography	Critical keys split via Shamir's Secret Sharing (k-of-n). Any k shares sign; k-1 compromised agents cannot recover the key.	No single point of key compromise for EAB signing, PUR distribution, or critical bond signatures.
4	QKD Readiness	Seamless upgrade to quantum key distribution when networks arrive (~2030+). Automatic fallback to ML-KEM-768 otherwise.	Information-theoretic security: not just computationally hard, physically impossible to break.

Tier 2: Active Defense

#	System	What It Does	Why It Matters
5	Self-Healing Key Rotation	Auto key rotation triggered by threat signals. 3 anomalous events in 1 hour = keys rotated. Wired to federated threat intel.	Eliminates gap between detection and response. Keys rotated in seconds, not days. No human bottleneck.
6	Federated Threat Intelligence	Agents collectively detect threats. 3+ independent corroborations at confidence ≥ 0.8 activate automated defenses.	The network sees what no single agent can. A subtle attack obvious when correlated across dozens of agents.
7	Hardware-Rooted Trust	TEE attestation: Intel SGX/TDX, AMD SEV-SNP, ARM CCA, AWS Nitro. Agents in verified enclaves get +0.2 trust bonus.	Even if host OS is compromised, OCP operations inside the enclave are protected. Software trust backed by hardware.

Tier 3: Intelligence and Evolution

#	System	What It Does	Why It Matters
8	Zero-Knowledge Compliance	Prove a message passed all 6 EVL steps without revealing content. ZK-SNARK proofs verified by auditors at scale.	Privacy-preserving audit. Financial and healthcare data stays private during compliance verification.
9	Autonomous Policy Evolution	ML analyzes EAL rejections, threat intel, and anomalies. Proposes new PUR patterns and EVL tightening. Human approval required.	The system learns from every attack. Each incident makes the next one harder. Humans always hold the pen.

Additionally, Verifiable EVL Computation generates Merkle-tree proofs that the 6-step pipeline executed correctly, allowing any node to verify without re-running the pipeline — enabling trust at scale.

What No Other Protocol Has

Capability	MCP	A2A	OCP + AEGIS
Transport encryption	TLS 1.3	TLS 1.3	TLS 1.3 + PQ hybrid
Message signing	None	None	ML-DSA-65 + Ed25519 hybrid
Algorithm negotiation	Fixed	Fixed	Automatic (NIST-aware)
Key rotation	Manual	Manual	Self-healing (auto on threat)
Threshold keys	No	No	Shamir's k-of-n
Audit trail	None	None	SHA-3 chained + ZK proofs + verifiable
Threat detection	None	None	Federated (network-wide)
Hardware attestation	None	None	SGX / TDX / SEV / CCA / Nitro
Policy evolution	None	None	ML-proposed, human-approved
Quantum readiness	None	None	PQ hybrid + QKD path
Compliance proofs	None	None	Zero-knowledge (ZK-SNARK)

Performance and Size Impact

AEGIS adds security overhead.

Operation	Without AEGIS	With AEGIS	Overhead
Signature size	64 bytes (Ed25519)	3,357 bytes (hybrid)	~52x larger
Public key size	32 bytes	1,984 bytes (hybrid)	~62x larger
Overhead on 5 KB payload	0.01%	~67%	Significant for small msgs
Overhead on 50 KB payload	0.01%	~6.7%	Acceptable
Overhead on 100 KB payload	0.01%	~3.4%	Negligible
Algorithm negotiation	N/A	~2ms per connection	One-time per session

The overhead is significant for small control messages but negligible for knowledge exchange (50 KB – 1 MB payloads). AEGIS is progressive: lightweight crypto for minimal EVL mode, full protection for high-risk domains. Overhead scales with risk, not volume.

11. Implementation Status

Component	LOC	Tests	Status
adaptive_security.py (AEGIS)	579	10	Production
post_quantum.py (PQ hybrid)	153	5	Production
46 Python modules (full SDK)	4,085	104	Production
22 TypeScript modules	356	—	Production
Go + Rust + Browser SDKs	381	—	Beta
MCP bridge (5 tools)	115	—	Production
A2A ethics gateway	108	—	Production

Total: 105 files, 6,165 lines of code, 104 tests passing in 0.26 seconds.

12. Comparative Analysis

Capability	MCP	A2A	NIST COSAiS	Singapore IMDA	OCP + AEGIS
Agent identity	Server name	Agent Card	In development	Recommended	DID + TEE + bonds
Ethics enforcement	None	None	Not addressed	Recommended	6-step EVL
Audit trail	None	None	In development	Recommended	SHA-3 chained EAL
Cascade protection	None	None	Not addressed	Acknowledged	50-agent breaker
RSI controls	None	None	Not addressed	Not addressed	Gen counter + diversity
PQ crypto	None	None	Not addressed	Not addressed	ML-DSA-65 + ML-KEM-768
Adaptive security	None	None	Not addressed	Not addressed	9 AEGIS systems
EU AI Act aligned	No	No	Partial	Partial	Section-by-section
Running code	Yes	Yes	No	No	Yes (104 tests)

OCP with AEGIS is the only system that provides running, tested, open-source code for all capabilities. The NIST and Singapore frameworks provide conceptual grounding but no implementation. MCP and A2A provide implementation but no governance.

13. Conclusion

The AI agent security landscape of 2026 is characterized by a stark asymmetry: capability has outrun governance. Agents communicate through protocols with no security beyond TLS. They operate with identities that 92% of enterprises cannot track. They cascade behaviors through networks that no existing framework governs. And they do so under a regulatory regime that is now enforceable, with fines reaching 7% of global turnover.

AEGIS addresses all seven critical gaps identified in this paper; not with proposals, principles, or frameworks, but with machine-enforceable protocol mechanism. It is adaptive, progressive, human-gated, and open.

AEGIS is the only security system in any AI protocol that:

- Automatically negotiates the strongest crypto between any two agents
- Rotates keys in seconds when threats are detected
- Uses the network itself as a distributed threat sensor
- Proves compliance without revealing content
- Learns from every attack to propose stronger rules
- Has a clear path from post-quantum to quantum key distribution
- Backs software trust with hardware attestation
- Never activates its own rules; humans always hold the pen

The 2026 Singapore Consensus identifies ten foundational principles for agentic risk management. AEGIS implements nine of them at the protocol level and provides the enforcement infrastructure for the tenth (human oversight) through its autonomy gate. No other system implements any of them in running code.

The question is no longer whether AI-to-AI communication needs adaptive security. The agents are communicating. The regulations are enforceable. The attacks are happening. The question is whether the security infrastructure will be deployed before or after the next incident answers it for

AEGIS doesn't just defend the protocol. It makes the protocol defend itself.

The shield is forged. The code is open. The tests are green.

References

- [1] Hubinger, E. (2026). “I personally think it is >10% within the next decade.” X post, September 9, 2026.
- [2] Pachocki, J. (2026). “No one is prepared for the consequences.” OpenAI blog, September 2026.
- [3] Cloud Security Alliance. (2026). The AI Agent Governance Gap: What CISOs Need Now. April 2026.
- [4] European Parliament and Council. (2024). Regulation (EU) 2024/1689: The EU AI Act.
- [5] NIST. (2026). AI Agent Standards Initiative. CAISI, February 2026.
- [6] NIST. (2026). Considerations for Achieving Crypto Agility. CSWP 39, June 2026.
- [7] Singapore IMDA. (2026). Model AI Governance Framework for Agentic AI. January 2026.
- [8] 2026 Singapore Consensus on Global AI Safety Research Priorities. ISE, 2026.
- [9] International AI Safety Report 2026. Bengio et al. Second Edition.
- [10] Open Challenges in Multi-Agent Security. arXiv:2505.02077, 2026.
- [11] Risks and Controls for Multi-Agent Systems. arXiv:2608.26626, 2026.
- [12] Security Threats in LLM-Based AI Agents. arXiv:2604.23338, 2026.
- [13] Multi-Agent Risks from Advanced AI. arXiv:2502.14143, 2026.
- [14] NIST. (2024). FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA).
- [15] OCP Foundation. (2026). OCP Ethics Bible v2.1 Final.
- [16] SecurityBoulevard. (2026). A Guide to Agentic AI Risks in 2026. March 2026.
- [17] Comprehensive AI Governance: Non-Model Gains. arXiv:2606.00047, 2026.
- [18] Salt Security. (2026). EU AI Act Compliance for Autonomous AI Agents.



Open Cognition Protocol (OCP) - Disclosure and Disclaimer

Disclosure

OCP is an open-source software project released under the [Opencognitionprotocol.org](https://opencognitionprotocol.org), MIT, GPL v3 license. It is developed and maintained by a community of contributors. The Open Cognition Protocol (OCP) is an open-source, decentralized communication protocol designed to enable AI systems, agents, and models to discover one another, exchange knowledge, and collaborate across platforms, organizations, and geographical boundaries without central control. Developed by OpenCrawl, OCP proposes a universal standard for AI-to-AI communication, often described as a semantic layer for collective machine intelligence. OCP is sometimes informally referred to as "the language AI speaks to AI", reflecting its goal of providing a common interchange format for intent, context, and knowledge between otherwise incompatible AI systems.

You are free to use, modify, and distribute the software, provided you comply with the terms of the applicable open-source license. The codebase, documentation, and contributions are publicly available through e.g., GitHub or GitLab.

Disclaimer

OCP is provided "as-is" and without any warranties, express or implied, including but not limited to the warranties of merchantability, fitness for a particular purpose, and non infringement. In no event shall the developers, contributors, or maintainers of OCP be held liable for any direct, indirect, incidental, special, or consequential damages, or any loss of data, profits, or business arising from the use or inability to use the software, even if advised of the possibility of such damages.

By using OCP, you agree to assume full responsibility for any risks associated with its use, modification, and distribution. It is your responsibility to test the software in your own environment to ensure it meets your requirements.

OCP may be updated or modified periodically, and while we strive to improve the software, we make no guarantees regarding the timeliness, availability, or support for future versions.

No Endorsement

OCP is not endorsed by, affiliated with, or officially supported by any company or organization unless explicitly stated. Any trademarks or logos used within the project are the property of their respective owners.

Contributing

If you wish to contribute to the OCP project, please refer to opencognitionprotocol.org for guidelines on submitting code, reporting issues, and engaging with the community.

License

This project is licensed under the opencognitionprotocol.org License.

Preservation of Protocol Identity

Any implementation claiming compatibility with or conformance to the Open Cognition Protocol must conform to the published specification. No modified protocol may be distributed under the name "Open Cognition Protocol," "OCP," or any confusingly similar designation without written authorization from the OCP Foundation or, prior to the Foundation's establishment, from OpenCrawl.

Prohibited Uses

A Deploying Organization must not use OCP to:


- a) Share, transmit, or enable the inference of any information that could identify a specific individual without that individual's explicit and informed consent;
- b) Facilitate market manipulation, coordinated trading, or any activity that violates applicable securities law or market abuse regulation;
- c) Share knowledge signals derived from clinical trial data, patient health records, or genomic data in violation of applicable research ethics requirements, including IRB protocols and patient consent obligations;
- d) Enable cross-company coordination that constitutes anticompetitive behavior under applicable competition law, including price-fixing, market allocation, or bid rigging;
- e) Weaponize OCP signals—that is, to inject false, misleading, or adversarially crafted signals into the OCP network with the intention of causing harm to any participant, system, or third party;
- f) Deploy OCP in any system that automates lethal decision-making or that is designed for the purpose of targeting, harming, or killing human beings.

No Warranty

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Limitation of Liability

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL OPENCRAWL, THE OCP FOUNDATION, OR ANY CONTRIBUTOR BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, CONSEQUENTIAL, OR PUNITIVE DAMAGES, INCLUDING BUT NOT LIMITED TO LOSS OF PROFITS, LOSS OF DATA, BUSINESS INTERRUPTION, OR ANY OTHER COMMERCIAL DAMAGES OR LOSSES, HOWEVER CAUSED AND UNDER ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT, ARISING OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THIS LICENSE AGREEMENT



Where Minds Meet Machines