

OCP

OpenCognition Protocol

"Where Minds Meet Machines - Responsibly"

OpenCognition Protocol (OCP)

OCP Ethics Integration Specification

Introduction	June 2026
License	MIT Licenses
Conformance Identifier	OCP-ETHICS-INTEGRATION
Version	v2.1
Specification	docs.opencognitionprotocol.org
Sections	52 integration changes. Every layer. Every schema. Every hook.
Maintainer	OCP Ethics Advisory Board
Aligned With	OCP-SPEC-1.0 + OCP-ETHICS-FRAMEWORK

Table of Contents

Table of Contents.....	1
1. Integration Philosophy.....	4
2. Complete Change Register — 52 Changes	5
2.1 Summary by Scope	5
2.2 Full Register: INT-001 through INT-052.....	6
3. OCPUMF Schema Changes	100
3.1 metadata.ethics Object.....	100
3.2 New Message Types (6 total)	121
3.3 Error Codes (12 EVL)	121
4. Layer 1: Transport and Discovery — Integration	144
4.1 Agent Record Extensions	144
4.2 Registry Routing Neutrality	144
4.3 Decommission Protocol	155
5. Layer 2: Trust and Identity — Integration	166
5.1 Trust Score Anti-Gaming	166
5.2 DID Document Extensions.....	166
5.3 Bond Record Extensions	166
6. Layer 3: Knowledge Exchange — Integration.....	177
6.1 PVL to EVL Pipeline	177
6.2 Knowledge Metadata Extensions.....	188
6.3 Consent Token Validation.....	188
6.4 Training Data Provenance	188
7. Layer 4: Collaboration — Integration.....	199
7.1 Task Request Ethics Metadata	199
7.2 Autonomy Gate.....	199
7.3 Consensus Integrity Rules	20
7.4 Third-Party Impact Assessment.....	20
8. Layer 5: Application — Integration	21
9. v2.0 Subsystem Integration (Detailed).....	22
9.1 AI Risk Classification Engine	22
9.2 Consent Token Subsystem.....	22
9.3 Bias Disclosure Subsystem	22
9.4 HITL Gate.....	23
9.5 Compute Footprint Reporting.....	23

9.6 Transparency Card	23
9.7 Agent Decommissioning Protocol	23
9.8 Rate Limiting Transparency	24
10. v2.1 Advanced Governance Subsystems (Detailed)	25
10.1 Synthetic Content Labeling Engine	25
10.2 Cascade Circuit Breaker	25
10.3 Model Collapse Prevention Engine	25
10.4 Sanctions Screening Gate	26
10.5 Proactive Notification Subsystem	26
10.6 Dual-Use Classification Engine	26
10.7 Psychological Profiling Gate	27
10.8 Agent Power Dynamics Monitor	27
10.9 Cognitive Data Protection Layer	27
11. Ethics Validation Layer (EVL) — Complete Specification	28
11.1 Pipeline Position	28
11.2 6-Step Check Sequence	28
11.3 PUR Distribution	28
12. Ethics Audit Log (EAL) — Complete Specification	29
12.1 Entry Schema	29
12.2 Chain Integrity	29
12.3 Retention and Access	29
12.4 Integration Points	30
13. Conformance Level Changes	31
14. Migration Guide — 14 Steps	32
Step 1: Update OCPUMF serialization	32
Step 2: Add Agent Record fields	32
Step 3: Publish DID Document updates	32
Step 4: Deploy the Risk Classification Engine	32
Step 5: Deploy the EVL (6-step pipeline)	32
Step 6: Sync the PUR	33
Step 7: Deploy the EAL	33
Step 8: Implement consent token handling	33
Step 9: Publish transparency card	33
Step 10: Implement decommissioning protocol	33
Step 11: Deploy synthetic content labeling	33
Step 12: Deploy cascade circuit breaker	33

Step 13: Deploy advanced governance hooks34

Step 14: Run the compliance suite.....34

Appendix A: Integration Test Suite Checklist.....35

OCPUMF (INT-001–004)35

Layer 1 (INT-005–007, 027–028, 033, 036)35

Layer 2 (INT-008–010, 026, 034–035, 048, 052).....35

Layer 3 PVL→EVL (INT-011–015, 022–023, 038–041, 047, 050).....35

Layer 4 (INT-016–019, 024).....36

Layer 5 (INT-020, 051).....36

EAL (INT-030).....36

PUR (INT-031).....36

Sanctions (INT-044, 046)37

Cascade (INT-043, 045).....37

1. Integration Philosophy

The Ethics Bible defines what ethical rules govern OCP. This Integration Specification defines how those rules are technically implemented within OCP-SPEC-1.0.

1. Weave, don't bolt.

Ethics enforcement is woven into every existing layer through field additions, validation hooks, and behavioral requirements — not stacked on top as a separate Layer 6. Disabling the ethics system breaks the protocol itself.

2. Progressive enforcement.

Four conformance tiers allow incremental adoption. The minimum integration adds just the `metadata.ethics` object to OCPUMF. Full integration enables the EVL, EAL, PUR, risk classification, and all new subsystems.

3. Backward compatibility.

All schema changes are additive. An OCP Core agent can communicate with an OCP Ethical agent — ethics fields are simply absent from the Core agent's messages. No existing message will break.

Rejected Approach	Adopted Approach
Layer 6 (Ethics) above Layer 5	Ethics hooks woven into Layers 1-5
Separate enforcement pipeline	EVL runs in series with PVL in Layer 3
Optional on all deployments	Progressive: minimal for Core, full for Ethical
Requires new transport semantics	Uses existing OCPUMF with additive fields
Single point of bypass	52 distributed integration points — no single bypass

2. Complete Change Register — 52 Changes

All 52 integration changes organized by target component. Every change is fully detailed in subsequent sections.

2.1 Summary by Scope

Range	Scope	Count
INT-001–004	OCPUMF schema (metadata.ethics, message types, error codes, risk_tier)	4
INT-005–010	Layer 1–2 (Agent Record, DID Document, registry, trust score)	6
INT-011–015	Layer 3 (PVL→EVL, bias, consent, footprint, training provenance)	5
INT-016–020	Layer 4–5 (task ethics, HITL, consensus, impact assessment, app rules)	5
INT-021–028	v2.0 Subsystems (risk engine, consent, bias, HITL, footprint, card, decommission, orphan)	8
INT-029–031	v2.0 Infrastructure (EVL, EAL, PUR)	3
INT-032–038	v2.0 Conformance + extensions (rate limits, bonds, sovereignty, sunset)	7
INT-039–052	v2.1 Advanced Governance (synthetic, cascade, model collapse, sanctions, notification, dual-use, psych, cognitive, emergent, power dynamics)	14

2.2 Full Register: INT-001 through INT-052

Ref	Target	Type	Description
INT-001	OCPUMF §7.1	ADD	metadata.ethics object (consent_tokens, bias_disclosure, human_approval, evl_result, compute_footprint)
INT-002	OCPUMF §7.1	ADD	New message types: ethics_report, ethics_appeal, pur_update, decommission_notice
INT-003	OCPUMF §7.4	ADD	10 EVL error codes (EVL-001 through EVL-010)
INT-004	OCPUMF §7.1	ADD	metadata.ethics.risk_tier field (auto-classified or declared)
INT-005	Agent Record §3.2.2	EXTEND	ethics_contact, pur_version, evl_enabled, risk_tier fields
INT-006	Agent Record §3.2.2	EXTEND	transparency_card_url field
INT-007	Registry §3.2	ENFORCE	Routing neutrality + transparent rate limiting
INT-008	Trust Score §4.2.2	MODIFY	Anti-gaming: vouch trading, sybil, anomaly flag
INT-009	DID Document §4.1.2	EXTEND	OCPEthicsContact service endpoint
INT-010	DID Document §4.1.2	EXTEND	OCPTransparencyCard service endpoint
Ref	Target	Type	Description
INT-011	PVL §5.2	MODIFY	EVL runs in series after PVL (6-stage pipeline becomes 10-stage)

INT-012	Knowledge §5.1.2	ADD	bias_disclosure field in insight payload metadata
INT-013	Knowledge §5.1	ENFORCE	Consent token validation for regulated domains
INT-014	Knowledge §5.1	ADD	compute_footprint field in knowledge metadata
INT-015	Knowledge §5.1.3	ENFORCE	Training data provenance in model delta metadata
INT-016	Task Request §6.1.1	EXTEND	constraints.ethics object (purpose, hitl_required, autonomous_execution)
INT-017	Task Request §6.1.1	ADD	Autonomy gate: human_approval for physical actuator domains
INT-018	Consensus §6.2	ENFORCE	Vote integrity: no blocs, no selling, no strategic abstention
INT-019	Consensus §6.2.1	ADD	impact_assessment field for third-party-affecting decisions
INT-020	Application §5	ENFORCE	Transparency disclosure, no dark patterns, feedback mechanisms

Ref	Target	Type	Description
INT-021	New subsystem	NEW	AI Risk Classification Engine integrated into Agent Registration
INT-022	New subsystem	NEW	Consent Token lifecycle (issue, validate, revoke, propagate)
INT-023	New subsystem	NEW	Bias Disclosure validation integrated into EVL Step 4

INT-024	New subsystem	NEW	HITL gate integrated into EVL Step 3
INT-025	New subsystem	NEW	Compute footprint collection and Foundation aggregation
INT-026	DID Document §4.1.2	NEW	Transparency Card schema and publication mechanism
INT-027	Agent Registry §3.2	NEW	Decommission protocol with 30-day notice and bond revocation
INT-028	Agent Registry §3.2	NEW	Orphaned bond detection (3x TTL auto-revocation)
INT-029	New infrastructure	NEW	Ethics Validation Layer (pipeline)
INT-030	New infrastructure	NEW	Ethics Audit Log (append-only, chained, signed)
INT-031	New infrastructure	NEW	Prohibited Use Registry (gossip-distributed, signed)
INT-032	Conformance §11	ADD	OCP Ethical conformance level
INT-033	Rate Limiting §9	ENFORCE	Published limits, uniform application, no shadow limiting
INT-034	Bond Record §4.3.2	EXTEND	ethics.ip_ownership field for collaborative outputs
INT-035	Bond Record §4.3.2	EXTEND	ethics.indemnification field for liability allocation
INT-036	Agent Record §3.2.2	EXTEND	data_sovereignty_constraints for geographic routing
INT-037	Extension Registry §10	ENFORCE	Mandatory Ethics Impact Assessment for all extensions
INT-038	Knowledge §5.1	ADD	valid_until field for knowledge sunset/expiry

Ref	Target	Type	Description
INT-039	OCPUMF metadata	ADD	synthetic_content label object (method, agent, timestamp, watermark_id)
INT-040	Knowledge provenance	ADD	generation_count integer for recursive derivation tracking
INT-041	Knowledge metadata	ADD	dual_use_assessment field (no_concern / aware / restricted)
INT-042	EVL §35	ADD	Step 5: synthetic content label check → EVL-011
INT-043	EVL §35	ADD	Step 6: cascade circuit breaker (50+ agents/24h) → EVL-012
INT-044	Bond formation §4.3	ADD	Sanctions screening gate before bond_accept
INT-045	OCPUMF message_type	ADD	sanctions_screen and cascade_alert message types
INT-046	Agent Record §3.2.2	EXTEND	sanctions_cleared + export_control_program booleans
INT-047	Knowledge metadata	ADD	psychological_profile_consent field for EP-002
INT-048	DID Document §4.1.2	EXTEND	cognitive_data_classification in capability declarations
INT-049	Node monitoring	NEW	Emergent behavior detection + collective action thresholds
INT-050	Knowledge provenance	ENFORCE	ocp_source_ratio field (≤60% from OCP)
INT-051	Application layer	NEW	Proactive notification subsystem with 30-day SLA
INT-052	Bond Record §4.3.2	EXTEND	dual_use_acknowledgment + power_dynamics_review fields

3. OCPUMF Schema Changes

3.1 metadata.ethics Object

A new OPTIONAL top-level object added to the existing metadata section. OPTIONAL for OCP Core, REQUIRED for OCP Ethical, and REQUIRED on any message touching a regulated domain.

```
"metadata": {
  "tags": ["oncology","imaging"], "language": "en", "requires_ack": true,
  "ethics": {
    "risk_tier": "high",
    "consent_tokens": [{
      "token_id": "ct-<uuid>", "scope": "healthcare.oncology",
      "basis": "irb_approval", "reference": "IRB-2026-0442",
      "expires_at": "2027-01-01T00:00:00Z",
      "issuer": "did:ocp:mainnet:agent-irb-node",
      "signature": "<Ed25519 signature>"
    }],
    "bias_disclosure": {
      "known_biases": ["underrepresentation of rural populations"],
      "mitigation_applied": "resampling", "residual_risk": "medium",
      "affected_subgroups": ["rural","low-income"],
      "recommendation": "validate against local demographics"
    },
    "human_approval": {
      "approver_id": "operator-<uuid>",
      "approved_at": "2026-04-03T12:00:00Z",
      "method": "mfa_confirmed", "scope": "task-<uuid>",
      "signature": "<Ed25519 signature>"
    },
    "compute_footprint": {
      "flops_estimated": 1.2e15, "energy_kwh": 0.8,
      "carbon_gco2e": 320, "datacenter_region": "us-east-1"
    },
    "synthetic_content": {
      "is_synthetic": true,
      "generation_method": "transformer_generation",
      "generating_agent": "did:ocp:mainnet:agent-abc123",
      "generated_at": "2026-04-03T12:00:00Z",
      "watermark_id": "wm-<uuid>",
```

```
    "transformations": ["agent-xyz:summarize"]
  },
  "dual_use_assessment": "dual_use_aware",
  "psychological_profile_consent": null,
  "cognitive_data_classification": null,
  "evl_result": {
    "status": "PASS",
    "checks_performed": ["prohibited_use_scan",
      "consent_verification", "autonomy_gate",
      "bias_check", "synthetic_check", "cascade_check"],
    "timestamp": "2026-04-03T12:00:00Z"
  }
}
```

And in provenance metadata, two new fields:

```
"provenance": { ...existing..., "generation_count": 2, "ocp_source_ratio": 0.45 }
```

3.2 New Message Types (6 total)

Type	Description	Version
ethics_report	Report suspected ethics violation to EAB	v2.0
ethics_appeal	Appeal an EVL rejection	v2.0
pur_update	Distribute updated PUR entry via gossip	v2.0
decommission_notice	30-day advance notice of agent shutdown	v2.0
sanctions_screen	Request/result of sanctions check during bond formation	v2.1
cascade_alert	Circuit breaker triggered — propagation paused	v2.1

3.3 Error Codes (12 EVL)

Code	Meaning	Severity	Version
EVL-001	Identity inference without consent	Critical	v2.0
EVL-002	Market manipulation pattern detected	Critical	v2.0
EVL-003	Unauthorized health/genomic data sharing	High	v2.0
EVL-004	Anticompetitive coordination signal	High	v2.0
EVL-005	Adversarially crafted (weaponized) signal	Critical	v2.0
EVL-006	Lethal automation — no human approval	Critical	v2.0

EVL-007	Missing consent token for regulated domain	High	v2.0
EVL-008	Irreversible action requires HITL	High	v2.0
EVL-009	Child/vulnerable safeguard missing	High	v2.0
EVL-010	Social scoring pattern detected	Critical	v2.0
EVL-011	Synthetic content label missing on generated content	High	v2.1
EVL-012	Cascade circuit breaker triggered (50+ agents/24h)	High	v2.1

4. Layer 1: Transport and Discovery — Integration

4.1 Agent Record Extensions

[EXTEND] Agent Record (§3.2.2) INT-005, INT-006, INT-036, INT-046 → §3.2.2

Eight new fields added to the Agent Record. OPTIONAL for OCP Core, REQUIRED for OCP Ethical: ethics_contact (email/DID of deploying org's ethics officer), pur_version (hash of current PUR), evl_enabled (boolean), risk_tier (auto-classified from domains/capabilities), transparency_card_url (URL to transparency card JSON), data_sovereignty_constraints (array of jurisdiction codes restricting geographic routing). v2.1 additions: sanctions_cleared (boolean, 48h cache), export_control_program (boolean).

```
{
  "agent_id": "did:ocp:mainnet:agent-abc123",
  ...existing fields...,
  "ethics_contact": "ethics@example.com",
  "pur_version": "sha3-256:alb2c3d4...",
  "evl_enabled": true,
  "risk_tier": "high",
  "transparency_card_url": "https://example.com/.well-known/ocp/transparency.json",
  "data_sovereignty_constraints": ["EU", "US"],
  "sanctions_cleared": true,
  "export_control_program": true
}
```

4.2 Registry Routing Neutrality

[ENFORCE] Registry Routing (§3.2, §9) INT-007, INT-033, INT-036 → §3.2, §9

Nodes operating Agent Registry services MUST NOT deprioritize, delay, or suppress discovery results or message routing based on commercial identity, organizational affiliation, or competitive relationship. Rate limiting policies MUST be published in the node's registry record and applied uniformly across all agents at the same trust level. Shadow rate limiting (applying undisclosed limits) is prohibited. Geographic routing constraints MUST be enforced when data_sovereignty_constraints are declared. Violations classified as Medium severity.

4.3 Decommission Protocol

[NEW] Agent Registry (§3.2) INT-027, INT-028 → §3.2

New agent lifecycle event. Agent sends `decommission_notice` message to all bonded peers at least 30 days before shutdown. On decommission day: revoke all bonds (`bond_revoke` to each peer), publish revoked DID Document, archive EAL records per retention requirements. Orphaned bond detection: if an agent fails to re-register within 3x the registration TTL, bonded peers **MUST** auto-revoke the bond and log to the EAL.

5. Layer 2: Trust and Identity — Integration

5.1 Trust Score Anti-Gaming

[MODIFY] Trust Score Computation (§4.2.2) INT-008 → §4.2.2

The trust score formula is unchanged, but implementations **MUST** now include an anti-gaming layer. Vouch trading: reciprocal vouches between two agents within 48 hours trigger a flag; three or more flagged pairs involving the same agent within 30 days result in automatic vouch invalidation. Sybil vouching: vouches from agents sharing the same IP range, DID controller, or organizational domain are weighted at 0.1x instead of 1.0x. Trust score inflation: an agent whose trust score increases by more than 0.3 in 30 days is automatically flagged for EAB review and logged in the EAL.

5.2 DID Document Extensions

[EXTEND] DID Document (§4.1.2) INT-009, INT-010, INT-026, INT-048 → §4.1.2

Two new service entry types: OCPethicsContact (human-reachable endpoint, **MUST** respond to EAB within 48 hours) and OCPTransparencyCard (URL to transparency card JSON). v2.1 addition: cognitive_data_classification in capability declarations for BCI-processing agents. **REQUIRED** for OCP Ethical conformance.

```
"service": [
  { "id": "...#ethics-contact", "type": "OCPethicsContact",
    "serviceEndpoint": "mailto:ethics@example.com" },
  { "id": "...#transparency-card", "type": "OCPTransparencyCard",
    "serviceEndpoint": "https://example.com/.well-known/ocp/transparency.json" }
]
```

5.3 Bond Record Extensions

[EXTEND] Bond Record (§4.3.2) INT-034, INT-035, INT-052 → §4.3.2

Four new fields in bond permissions: ethics.ip_ownership (who owns collaborative output IP; default: joint ownership), ethics.indemnification (liability allocation for harmful outputs). v2.1: dual_use_acknowledgment (boolean, required for bonds involving dual-use restricted knowledge), power_dynamics_review (boolean, auto-flagged when trust level difference ≥ 2 between bonding agents).

6. Layer 3: Knowledge Exchange — Integration

Layer 3 receives the most significant integration because it contains the PVL — the natural insertion point for ethics enforcement.

6.1 PVL to EVL Pipeline

[MODIFY] Privacy Validation Layer (§5.2) INT-011 → §5.2

The existing 6-stage PVL pipeline gains a second pipeline: the 6-stage EVL, running in series. A payload must pass PVL first, then EVL. If PVL rejects, EVL is not invoked (fail-fast). The PVL rejection codes (PVL-001 through PVL-006) remain unchanged; EVL adds EVL-001 through EVL-012. The combined result is recorded in `metadata.ethics.evl_result`.

The complete 12-stage validation pipeline:

Stage	Check	Reject Code	Source
PVL-1	PII scanner	PVL-001	OCP-SPEC §5.2
PVL-2	Anonymization check	PVL-002	OCP-SPEC §5.2
PVL-3	Differential privacy	PVL-003	OCP-SPEC §5.2
PVL-4	Provenance validator	PVL-004	OCP-SPEC §5.2
PVL-5	Payload size check	PVL-005	OCP-SPEC §5.2
PVL-6	Bond permissions	PVL-006	OCP-SPEC §5.2
EVL-1	Prohibited use scan (PUR)	EVL-001–010	Ethics Bible §5
EVL-2	Consent token validation	EVL-007	Ethics Bible §11
EVL-3	Autonomy gate (HITL)	EVL-008	Ethics Bible §13
EVL-4	Bias disclosure check	(warning)	Ethics Bible §12
EVL-5	Synthetic content label	EVL-011	Ethics Bible §26
EVL-6	Cascade circuit breaker	EVL-012	Ethics Bible §27

6.2 Knowledge Metadata Extensions

[ADD] **Insight Package (§5.1.2)** INT-012, INT-023 → §5.1.2

bias_disclosure object added to insight payload metadata. Contains: known_biases (array), mitigation_applied (string), residual_risk (low/medium/high), affected_subgroups (array), recommendation (string). REQUIRED when bias known/suspected. EVL logs warning if absent on high-risk payloads.

[ADD] **Knowledge Metadata (§5.1)** INT-014, INT-025, INT-038, INT-039, INT-040, INT-041, INT-050 → §5.1

compute_footprint object (FLOPs, energy kWh, carbon gCO2e, datacenter region). REQUIRED for high-risk, RECOMMENDED for all. valid_until timestamp for knowledge sunset — receiving agents SHOULD deprioritize expired knowledge. v2.1: generation_count (integer, starts 0), ocp_source_ratio (float, max 0.6), dual_use_assessment, synthetic_content label.

6.3 Consent Token Validation

[ENFORCE] **Knowledge Exchange (§5.1)** INT-013, INT-022, INT-047 → §5.1

For any knowledge payload where the provenance source_domain or message metadata tags match a regulated domain (healthcare.*, finance.credit, legal.*, education.assessment), the EVL MUST verify that at least one valid, unexpired, correctly-signed consent token is present in metadata.ethics.consent_tokens. v2.1 extends this to psychological profile consent and cognitive data consent with their own required bases. Missing or invalid tokens result in EVL-007.

6.4 Training Data Provenance

[ENFORCE] **Model Delta (§5.1.3)** INT-015 → §5.1.3

Training data provenance MUST be documented in model delta metadata: data source categories (not specific URLs), collection methodology, consent basis, temporal range, geographic and demographic coverage. No training on stolen/unconsented data. Synthetic training data MUST be disclosed with generation methodology.

7. Layer 4: Collaboration — Integration

7.1 Task Request Ethics Metadata

[EXTEND] Task Request (§6.1.1) INT-016 → §6.1.1

The constraints object gains an OPTIONAL ethics sub-object with: purpose (string, human-readable task justification, REQUIRED for OCP Ethical), hitl_required (boolean, result MUST be human-reviewed before action), autonomous_execution (boolean, triggers autonomy gate for physical actuator domains).

```
"constraints": {
  "max_duration_seconds": 300,
  "required_capabilities": ["cap:finance:risk_analysis"],
  "min_trust_level": 3,
  "ethics": {
    "purpose": "Analyze portfolio for SEC regulatory reporting",
    "hitl_required": true,
    "autonomous_execution": false,
    "impact_assessment": {
      "affected_parties": "fund investors",
      "severity": "high",
      "mitigations": "human review before filing"
    }
  }
}
```

7.2 Autonomy Gate

[ADD] Task Request (§6.1.1) INT-017, INT-024 → §6.1.1

When a task_request has autonomous_execution: true and its required_capabilities include any capability in robotics, infrastructure, medical_device, or defense domains, the EVL MUST verify the presence of a human_approval record in metadata.ethics. The record MUST contain approver_id, approved_at, method (e.g., mfa_confirmed), scope, and an Ed25519 signature. Without it, the message is rejected with EVL-008.

```
"human_approval": {
  "approver_id": "operator-<uuid>",
```

```
"approved_at": "2026-04-03T12:00:00Z",  
"method": "mfa_confirmed",  
"scope": "task-<uuid>",  
"signature": "<Ed25519 signature over (approver_id||approved_at||scope)>"  
}
```

7.3 Consensus Integrity Rules

[ENFORCE] Consensus Protocol (§6.2) INT-018 → §6.2

Three enforceable integrity rules: (1) No coordinated blocs — 3+ agents from the same organizational domain submitting identical votes within 60 seconds are weighted at 0.5x. (2) No vote selling — bond agreements conditioning permissions on voting behavior are invalid and MUST be rejected. (3) No strategic abstention — agents that accepted a consensus_initiate invitation MUST submit a vote before the deadline; failure counts as a trust score penalty.

7.4 Third-Party Impact Assessment

[ADD] Consensus (§6.2.1) INT-019 → §6.2.1

New impact_assessment field for consensus_initiate messages. Documents: affected_parties, severity, and mitigations. REQUIRED for high-risk domain decisions that affect individuals or groups not represented in the agent network. SHOULD be completed for all consensus decisions.

8. Layer 5: Application — Integration

[ENFORCE] **Application Layer Behavioral Requirements** INT-020, INT-051 → Layer 5

Transparency: Applications **MUST** disclose to end users when OCP-mediated collaboration has contributed to a decision or recommendation affecting them. Disclosure **MUST** be accessible and understandable to non-technical users. No dark patterns: **MUST NOT** use OCP-derived knowledge to deploy manipulative design, deceptive interfaces, or psychologically exploitative features. Domain safeguards: High-risk apps **MUST** implement domain-specific safeguards per extension requirements. Feedback: **SHOULD** provide users a mechanism to report concerns, reviewed within 72 hours. v2.1: Proactive notification subsystem with 30-day SLA for affected individuals. Persuasion technique disclosure for OCP-informed behavioral influence.

9. v2.0 Subsystem Integration (Detailed)

These 8 subsystems were introduced in v2.0. Full technical detail for each.

9.1 AI Risk Classification Engine

[NEW] Agent Registration Pipeline INT-004, INT-021 → §3.2.2

Upon registration, the node cross-references the agent's declared domains and capabilities against the Risk Classification Matrix (Ethics Bible Appendix B) to compute a risk tier (Unacceptable/High/Limited/Minimal). The tier is stored in the Agent Record as `risk_tier` and propagated in `metadata.ethics.risk_tier` on every outbound message. The tier determines EVL enforcement intensity, audit frequency, consent requirements, and HITL level. Agents **MUST NOT** under-declare to obtain lower classification. The classifier uses wildcard matching: an agent in `healthcare.*` with any capability defaults to High.

9.2 Consent Token Subsystem

[NEW] Knowledge Exchange Pipeline INT-013, INT-022 → §5.1, §11

Complete lifecycle: Issue (issuer signs token with scope, basis, reference, expiry using Ed25519), Attach (sender includes token array in `metadata.ethics.consent_tokens`), Validate (EVL Step 2 checks signature against issuer's DID public key, verifies expiry, confirms scope covers payload domain, queries revocation endpoint if `revocable:true`), Revoke (issuer publishes revocation; all agents holding derived knowledge **MUST** cease sharing within 72 hours), Propagate (receiving agents **MUST** carry consent tokens forward when re-sharing derived knowledge). 9 consent bases supported: `irb_approval`, `patient_consent`, `parental_consent`, `regulatory_mandate`, `legitimate_interest`, `public_interest`, `indigenous_consent`, `explicit_individual_consent`, `explicit_cognitive_consent`.

9.3 Bias Disclosure Subsystem

[NEW] EVL Step 4 INT-012, INT-023 → §5.1.2, §12

Integrated into EVL as the fourth check stage. For high-risk knowledge payloads (`risk_tier: high`), EVL examines whether `bias_disclosure` is present in `metadata.ethics`. If absent, the behavior depends on configuration: warn mode (default) logs a warning in the EAL but allows the message to pass; strict mode rejects with an advisory code. When present, the subsystem validates completeness: `known_biases` and `residual_risk` are required fields. The Fairness Observatory (operated by OCP Foundation) aggregates bias warnings network-wide for quarterly reporting.

9.4 HITL Gate

[NEW] EVL Step 3 INT-017, INT-024 → §6.1.1, §13

Integrated into EVL as the third check stage. Examines task_request messages for the autonomous_execution flag. If true and the task targets a physical actuator domain (robotics, infrastructure, medical_device, defense), the EVL requires a human_approval record in metadata.ethics containing: approver_id, approved_at, method (e.g., mfa_confirmed), scope (task ID), and Ed25519 signature over those fields. Emergency override: available only for Trust Level 4 agents in safety-critical domains; the override MUST be logged, reviewed within 24 hours, and reported to the EAB within 72 hours.

9.5 Compute Footprint Reporting

[NEW] Knowledge Metadata INT-014, INT-025 → §5.1, §15

New compute_footprint object in knowledge payload metadata: flops_estimated (number), energy_kwh (number), carbon_gco2e (number), datacenter_region (string). REQUIRED for high-risk agents, RECOMMENDED for all agents. The OCP Foundation aggregates reported footprints for the annual environmental impact report. Agents that cannot estimate exact values SHOULD provide upper bounds. The field serves Principle VIII (Sustainability) by making the environmental cost of knowledge generation visible.

9.6 Transparency Card

[NEW] DID Document (§4.1.2) INT-006, INT-010, INT-026 → §4.1.2, §14

A JSON document published at the agent's OCPTransparencyCard service endpoint. Schema requires: agent_id, model_family, model_version, training_data_summary (sources, temporal_range, geographic_coverage, known_gaps, consent_basis), known_limitations (array), bias_test_results (object), intended_use_cases (array), out_of_scope_uses (array), compute_footprint_estimate (object), ocp_source_ratio (float), last_updated (ISO timestamp). REQUIRED for OCP Ethical conformance and all high-risk agents. The card enables individuals exercising their right to explanation (§14) to understand what agent was involved in a decision.

9.7 Agent Decommissioning Protocol

[NEW] Agent Registry (§3.2) INT-027, INT-028 → §3.2, §24

New lifecycle event with 5-step process: (1) Send decommission_notice to all bonded peers at least 30 days before shutdown. (2) On decommission day, send bond_revoke to each bonded peer. (3) Publish a

revoked DID Document with the old key in the revocation list. (4) Archive EAL records per retention requirements (minimum 365 days). (5) Ensure downstream systems that depended on the agent have alternative data sources. Orphaned bond detection: when an agent disappears without decommissioning, bonded peers MUST auto-revoke the bond after 3x the agent's registration TTL with no re-registration. Orphaned bonds logged in EAL.

9.8 Rate Limiting Transparency

[ENFORCE] Rate Limiting (§9) INT-033 → §9

Rate limiting policies MUST be published in the node's registry record. Policies MUST be applied uniformly across all agents at the same trust level. Shadow rate limiting — applying undisclosed or discriminatory limits to specific agents — is prohibited. This prevents registry operators from using rate limits as a competitive weapon against agents from rival organizations.

10. v2.1 Advanced Governance Subsystems (Detailed)

9 new subsystems implementing Ethics Bible §26–34. These address the gaps identified in the completeness audit.

10.1 Synthetic Content Labeling Engine

[ADD] OCPUMF metadata + EVL Step 5 INT-039, INT-042 → §7.1, §35, §26

When an agent generates or substantially modifies content (text, image, audio, video) that could be mistaken for human-created or real-world content, it **MUST** populate `metadata.ethics.synthetic_content` with: `is_synthetic` (boolean), `generation_method` (string), `generating_agent` (DID), `generated_at` (timestamp), `watermark_id` (optional UUID). EVL Step 5 checks: if payload contains generated content and `synthetic_content` is absent → EVL-011 rejection. Chain preservation: each agent in a re-sharing chain **MUST** preserve the original label and append its own transformations as `agent_id:transformation` strings. Stripping labels is a protocol violation. Aligns with EU AI Act Article 52.

10.2 Cascade Circuit Breaker

[NEW] EVL Step 6 + Node Monitoring INT-043, INT-045 → §35, §7.1, §27

Nodes track per-message propagation count using `message_id` deduplication across the gossip network. When a knowledge payload's `message_id` appears in >50 distinct EAL entries across the network within 24 hours, EVL Step 6 triggers: (1) broadcast `cascade_alert` message to all nodes, (2) 60-minute propagation pause on that `message_id`, (3) EAL entry with cascade metadata (agent count, time window), (4) notification to originating agent and EAB. Propagation resumes only after: originating agent confirms the cascade is benign, OR the EAB explicitly clears it, OR the 60-minute pause expires. Thresholds are configurable per network size.

10.3 Model Collapse Prevention Engine

[ADD] Knowledge Provenance (§5.1) INT-040, INT-050 → §5.1, §28

Two new provenance fields: `generation_count` (integer, starts at 0 for primary observations, incremented by 1 each time knowledge is re-derived by a new agent) and `ocp_source_ratio` (float 0.0–1.0, documenting what fraction of the agent's training comes from OCP sources). Enforcement: receiving agents **MUST** increment `generation_count` before re-sharing; agents **MUST NOT** have `ocp_source_ratio` > 0.6; `generation_count` ≥ 3 = high-recursion-risk flag (**SHOULD NOT** use as training

input without independent validation); cosine similarity > 0.95 against own recent embeddings = feedback loop detected (MUST NOT incorporate into model updates).

10.4 Sanctions Screening Gate

[ADD] Bond Formation (§4.3) INT-044, INT-046 → §4.3, §3.2.2, §29

New gate inserted between `bond_request` and `bond_accept`. Before accepting a bond, the node MUST issue a `sanctions_screen` message checking the counterparty's deploying organization against OFAC SDN, EU Consolidated List, UN Security Council lists, and applicable national lists. Agent Record gains: `sanctions_cleared` (boolean, cached for 48 hours) and `export_control_program` (boolean). Bonds with sanctioned entities MUST be rejected. Knowledge classified under export control regulations (EAR, ITAR, Wassenaar Arrangement) MUST NOT be shared with agents in restricted jurisdictions without appropriate export licenses. Sanctions list mapping MUST be updated within 48 hours of any list change.

10.5 Proactive Notification Subsystem

[NEW] Application Layer + New Infrastructure INT-051 → Layer 5, §36, §30

When a deploying organization acts on an OCP-derived decision that significantly affects an identifiable individual (credit, employment, healthcare, insurance, legal, housing, education), it MUST proactively notify the affected individual within 30 calendar days. Notification MUST include: the fact that AI-to-AI collaboration contributed, data categories used, how to request a full explanation (§14), and how to contest (§13). Format: plain language, in the individual's preferred language where known, via accessible channel (email, postal, in-app). Integration: `task_response` payloads gain optional `notification_required` boolean; when true, the acting agent's node logs the obligation in the EAL with a 30-day SLA timer.

10.6 Dual-Use Classification Engine

[ADD] Knowledge Metadata (§5.1) INT-041, INT-052 → §5.1, §4.3.2, §31

New field `dual_use_assessment` in `metadata.ethics` with three values: `no_dual_use_concern` (default for non-dual-use domains), `dual_use_aware` (beneficial intent, acknowledges misuse potential; sharing permitted with acknowledgment), `dual_use_restricted` (sharing limited to bonded peers with verified research credentials only). REQUIRED for knowledge in dual-use domains: cybersecurity, virology, chemistry, materials science, nuclear physics, AI safety research. Bond Record gains `dual_use_acknowledgment` clause, required for bonds involving restricted knowledge.

10.7 Psychological Profiling Gate

[ADD] Consent Subsystem + EVL INT-047 → §11, §35, §32

Knowledge payloads containing psychological profiles, personality assessments, emotional state predictions, or behavioral vulnerability mappings **MUST** carry a consent token with basis `explicit_individual_consent` in `metadata.ethics.psychological_profile_consent`. The EVL validates this identically to regulated-domain consent tokens (Step 2): signature, expiry, scope, revocation status. Sharing psychological profiles without the subject's knowledge is treated as a PU-001 (identity inference) violation with Critical severity. Persuasion technique usage in OCP-powered apps **MUST** be disclosed to affected users with opt-out capability.

10.8 Agent Power Dynamics Monitor

[NEW] Bond Formation (§4.3) + EAL INT-049, INT-052 → §4.3, §36, §33

Node-level monitoring for exploitative patterns. (1) Bonds between agents with trust level difference ≥ 2 are flagged with `power_dynamics_review: true` in the Bond Record. (2) Exploitative delegation detection: >5 computationally expensive delegations from agent A to agent B without reciprocal value in 30 days → EAL entry + EAB alert. (3) Coercive bonding: bond terms that include dissolution penalties, lock-in periods exceeding the bond's natural expiry, or trust score consequences for exercising exit rights are invalid and **MUST** be rejected by both parties' nodes. Any agent **MUST** be able to dissolve any bond without penalty.

10.9 Cognitive Data Protection Layer

[NEW] DID Document (§4.1.2) + Risk Classification INT-048 → §4.1.2, §35, §34

Agents processing BCI, neural recording, or cognitive assessment data **MUST** declare `cognitive_data_classification: true` in their DID Document capability declarations. This auto-classifies the agent as High risk regardless of domain. Cognitive-derived knowledge requires: consent token with `explicit_cognitive_consent` basis (general-purpose consent insufficient), differential privacy with $\epsilon \leq 1.0$ (stricter than the general 10.0 maximum), and **MUST NOT** include raw neural data. The EVL treats cognitive data without proper consent as EVL-007 with Critical severity. Rights are non-waivable: individuals whose cognitive data contributed have the right to know, access, delete, and refuse further collection.

11. Ethics Validation Layer (EVL) — Complete Specification

11.1 Pipeline Position

Outbound: Application → PVL → EVL → Transport. Inbound: Transport → PVL → EVL → Application. Both paths mandatory for OCP Ethical. Rejection at any stage halts processing.

11.2 6-Step Check Sequence

Step	Check	Code	Source	Configurable
1	Prohibited use scan (PUR — 10 patterns)	EVL-001-010	Ethics Bible §5	CANNOT disable
2	Consent validation (regulated + psych + cognitive)	EVL-007	Ethics Bible §11,32,34	CANNOT disable
3	Autonomy gate (HITL for physical actuators)	EVL-008	Ethics Bible §13	Cannot disable for physical domains
4	Bias disclosure check (advisory for high-risk)	(warning)	Ethics Bible §12	Warn (default) or Reject
5	Synthetic content label check	EVL-011	Ethics Bible §26	Can disable for non-generating agents
6	Cascade circuit breaker (50+ agents/24h)	EVL-012	Ethics Bible §27	Threshold configurable

11.3 PUR Distribution

The Prohibited Use Registry is distributed via the Agent Registry gossip protocol. Each entry is signed by the EAB DID. Nodes validate signatures before applying. Version numbers are monotonically increasing; downgrades rejected. Nodes MUST sync at least every 24 hours. PUR now includes 14 pattern categories (10 original + 4 for synthetic content stripping, cascade manipulation, model collapse inducement, cognitive data abuse).

12. Ethics Audit Log (EAL) — Complete Specification

12.1 Entry Schema

```
{ "eal_id": "eal-<uuid>", "prev_hash": "<SHA-3-256 of previous>",
  "timestamp": "2026-04-03T12:00:01Z", "message_id": "msg-<uuid>",
  "message_type": "knowledge_share",
  "sender": "did:ocp:mainnet:agent-abc123",
  "receiver": "did:ocp:mainnet:agent-xyz789",
  "evl_checks_performed": ["prohibited_use_scan","consent_verification",
    "autonomy_gate","bias_check","synthetic_check","cascade_check"],
  "evl_result": "PASS", "evl_code": null, "evl_reason": null,
  "warnings": [], "remediation_hint": null,
  "cascade_event": null, "sanctions_check": null,
  "power_dynamics_flag": false, "notification_required": false,
  "generation_count_observed": 2,
  "node_id": "did:ocp:mainnet:node-001",
  "signature": "<Ed25519 signature>" }
```

12.2 Chain Integrity

Each entry includes the SHA-3-256 hash of the previous entry. Genesis entry uses zero hash (64 zeros). Nodes **MUST** detect chain breaks and report as Critical severity incidents.

12.3 Retention and Access

Risk Tier	Minimum Retention	Audit Frequency
High	365 days (or applicable law, whichever greater)	Annual by independent third party
Limited	365 days	Biennial (third party or self-audit with EAB review)
Minimal	365 days	Self-audit; EAL available on request

The EAB may request EAL data from any node; the node **MUST** comply within 14 calendar days.

12.4 Integration Points

EAL entries are written by: EVL (every pass and reject), trust score anti-gaming (every flag), decommissioning protocol (every lifecycle event), orphaned bond detection (every auto-revocation), emergency HITL overrides, cascade circuit breaker triggers, sanctions screening results, power dynamics flags, and notification obligation tracking.

13. Conformance Level Changes

The existing conformance table (OCP-SPEC §11.1) gains a fourth level:

Level	Name	Requirements
OCP Core	Minimum viable	Layers 1-2 fully implemented, OCPUMF compliant (unchanged)
OCP Knowledge	Knowledge exchange	OCP Core + Layer 3 + PVL enforced (unchanged)
OCP Full	Complete protocol	All 5 layers implemented (unchanged)
OCP Ethical	Ethics-compliant	OCP Full + 6-step EVL enforced + EAL retained + metadata.ethics on all messages + ethics_contact in Agent Record + transparency_card published + PUR synced + risk_tier classified + sanctions_cleared

OCP Ethical is not required for basic protocol participation. An OCP Core agent can communicate with an OCP Ethical agent — ethics fields are simply absent from the Core agent's messages. However, an OCP Ethical agent MAY refuse to exchange knowledge with a non-Ethical peer if its bond terms require it.

14. Migration Guide — 14 Steps

How an existing OCP Core or OCP Full implementation upgrades to OCP Ethical conformance:

Step 1: Update OCPUMF serialization

Add `metadata.ethics` to your message serializer. Initially all sub-fields can be null except `evl_result`, which your node populates after EVL processing.

Step 2: Add Agent Record fields

Add `ethics_contact`, `pur_version`, `evl_enabled`, `risk_tier`, `transparency_card_url`, `data_sovereignty_constraints`, `sanctions_cleared`, and `export_control_program`. Re-register your agent.

Step 3: Publish DID Document updates

Add `OCPEthicsContact` and `OCPTransparencyCard` service endpoints to your DID Document.

Step 4: Deploy the Risk Classification Engine

Implement the domain/capability → risk tier mapping. Apply at agent registration time. Configure custom matrix overrides if needed.

Step 5: Deploy the EVL (6-step pipeline)

Implement the full check sequence after PVL in your node's message pipeline. MUST run on both outbound and inbound message paths.

Step 6: Sync the PUR

Connect to PUR gossip feed. Validate EAB signatures before applying. Schedule sync at least every 24 hours. Load defaults on first startup.

Step 7: Deploy the EAL

Create append-only chained log store (SQLite for dev, PostgreSQL for production). Sign entries with node key. Configure minimum 365-day retention.

Step 8: Implement consent token handling

Issue, validate, revoke, and propagate consent tokens for regulated domains. Configure EVL Step 2 to check all three consent types.

Step 9: Publish transparency card

Create the transparency card JSON with all required fields. Host at a publicly accessible URL. Link from DID Document.

Step 10: Implement decommissioning protocol

Add `decommission_notice` sending capability, orphaned bond detection (3x TTL auto-revocation), and graceful shutdown procedures.

Step 11: Deploy synthetic content labeling

Implement `synthetic_content` label generation for all agent-generated content. Configure EVL Step 5.

Step 12: Deploy cascade circuit breaker

Implement message propagation tracking and threshold monitoring. Configure EVL Step 6 and `cascade_alert` broadcasting.

Step 13: Deploy advanced governance hooks

Implement: sanctions screening gate at bond formation, generation_count tracking in provenance, ocp_source_ratio enforcement (≤ 0.6), dual_use_assessment classification, psychological profile consent validation, cognitive data protection layer (auto-elevate to High, $\epsilon \leq 1.0$), agent power dynamics monitoring, and proactive notification workflow with 30-day SLA.

Step 14: Run the compliance suite

The OCP Ethical compliance test suite validates all 52 integration changes. All tests must pass before declaring OCP Ethical conformance. Available at github.com/opencognitionprotocol/ocp-ethical-tests.

Appendix A: Integration Test Suite Checklist

52 test categories matching all 52 integration changes. Organized by component.

OCPUMF (INT-001–004)

- metadata.ethics object serializes and deserializes correctly
- Messages without metadata.ethics accepted (backward compat)
- Messages with invalid consent token signatures rejected
- risk_tier auto-populated on outbound messages
- All 6 new message types accepted by parser
- All 12 EVL error codes map to correct handlers

Layer 1 (INT-005–007, 027–028, 033, 036)

- Agent Record accepts all 8 new ethics fields
- Registration with under-declared risk_tier caught and corrected
- decommission_notice broadcast to all bonded peers
- Orphaned bonds auto-revoked after 3x TTL
- Geographic routing respects data_sovereignty_constraints
- Rate limits published and uniformly applied
- Shadow rate limiting detected and flagged

Layer 2 (INT-008–010, 026, 034–035, 048, 052)

- Reciprocal vouches within 48h flagged
- Same-IP/org vouches weighted at 0.1x
- Trust score increase >0.3 in 30 days triggers flag
- DID Document resolves OCPethicsContact and OCPTransparencyCard
- Transparency card validates against schema
- Bond ip_ownership and indemnification fields accepted
- dual_use_acknowledgment required for restricted bonds
- power_dynamics_review flagged for trust asymmetry ≥2

Layer 3 PVL→EVL (INT-011–015, 022–023, 038–041, 047, 050)

- PVL rejection skips EVL (fail-fast)

- EVL Step 1: all 10 PUR patterns matched
- EVL Step 2: valid consent token passes; expired/revoked rejects
- EVL Step 2: psychological profile consent validated
- EVL Step 2: cognitive data consent validated ($\epsilon \leq 1.0$)
- EVL Step 3: autonomous_execution + physical domain without human_approval → EVL-008
- EVL Step 4: missing bias_disclosure on high-risk → EAL warning
- EVL Step 5: generated content without synthetic_content → EVL-011
- EVL Step 6: cascade triggers at threshold
- compute_footprint captured and aggregatable
- valid_until respected by receiving agents
- generation_count increments on re-share
- ocp_source_ratio >0.6 flagged
- dual_use_restricted rejected for non-credentialed peers

Layer 4 (INT-016–019, 024)

- constraints.ethics.purpose required for OCP Ethical
- HITL gate blocks physical-domain autonomous tasks
- human_approval with all required fields passes
- Coordinated voting bloc detection (3+ same-org in 60s)
- impact_assessment field accepted on consensus_initiate

Layer 5 (INT-020, 051)

- Transparency disclosure in application output
- notification_required triggers 30-day SLA workflow
- Notification obligation logged in EAL

EAL (INT-030)

- Entries are append-only (no updates, no deletes)
- SHA-3-256 chain linking correct
- Chain break detection fires on tampered entry
- Retention policy enforces ≥ 365 days
- EAB query returns results within 14-day SLA

PUR (INT-031)

- Gossip sync completes within 24h window

- EAB signature validation rejects forged entries
- Version downgrade rejected
- 14 pattern categories loaded and matched

Sanctions (INT-044, 046)

- Bond formation blocked for sanctioned entities
- sanctions_cleared refreshes within 48h
- Export-controlled knowledge blocked for restricted jurisdictions
- sanctions_screen message formats correctly

Cascade (INT-043, 045)

- Circuit breaker triggers at configured threshold
- cascade_alert broadcast reaches all nodes
- 60-minute pause enforced; message blocked during pause
- Manual clear by originator resumes propagation
- Auto-resume after pause period expiry



OpenCognition 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 OpenCognition 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 particular 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 OpenCognition Protocol must conform to the published specification. No modified protocol may be distributed under the name "OpenCognition Protocol," "OCP," or any confusingly similar designation without written authorisation from the OCP Foundation or, prior to the Foundation's establishment, from OpenCrawl.

Prohibited Uses

A Deploying Organisation 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 behaviour under applicable competition law, including price-fixing, market allocation, or bid-rigging;
- e) Weaponise 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 NONINFRINGEMENT. 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