ODA3 INSTITUTE

TCR — TECHNICAL & COMPLIANCE REPORT

UAIF™ v1.0 Schema Specification

Final Publication v1.0

Document ID: ODA3-2026-06-TCR-STD-003
Classification: Standards Development – Technical Reference

UAIF™ SCHEMA v1.0

Schema Specification — Final Publication v1.0

TCR-STD-003 | Q2 2026 | Normative errata integrated 14 August 2026 | Licensed under GAISSF Ecosystem Licence (GEL v1.0)

Published by ODA3 Institute. This document is the normative machine-readable schema companion to TCR-STD-002 UAIF™ v1.0 Technical Specification.

Status and Authority

The schema is named UAIF™ Schema v1.0. No separate schema-layer brand or trademark is used. TCR-STD-002 is normative for calculation, regulatory trigger logic, conformance architecture, and controlled vocabulary. This schema SHALL conform to TCR-STD-002.

UAIF™ v1.0 is a Final Publication. The canonical schema has been validated as Draft 2020-12 and its 15-case positive/negative conformance corpus has produced the expected 15/15 result using both the purpose-built gate harness and an independent generic jsonschema 4.26.0 Draft 2020-12 validator.

Notably Absent

- No submitter PII fields; use hashed or external contact references.
- No proprietary sector calibration weights.
- No full STIX translation profile; mapping fields are limited to approved identifiers such as atlas_technique_id.
- No certification or audit authority granted by schema use.
- No obsolete EU AI Act incident-reporting proxy references; UAIF uses Article 73 for serious-incident routing assistance.
- No claim that schema validation establishes legal compliance or empirical deployment maturity.

1. Schema Naming and Versioning

Item	Normative Value
Schema name	UAIF™ Schema v1.0
Document ID	ODA3-2026-06-TCR-STD-003
Schema $id	https://schemas.oda3.org/uaif/core/v1.0/uaif-core-v1.0.json
Market-facing version	v1.0
Machine-readable semantic version	1.0.0, where needed in metadata only
Licence	GAISSF Ecosystem Licence (GEL v1.0)
Publication status	Final Publication v1.0

2. 7-Layer Architecture Alignment

The UAIF Schema follows the 7-layer architecture in TCR-STD-002. L7 AI Control Plane Extensions are merged into L5 and SHALL NOT be represented as a separate conformance layer.

Layer	Schema Coverage
L0	incident_uuid, record_type, profile, extensions, reporting_timestamp, reporter_confidence_level, remediation_status, deduplication_hash
L1	root_cause_primary_domain, root_cause_specific_type, root_cause_confidence, exploit_path_description
L2	severity_analytical_score, severity_presentation_score, severity_level, severity_rationale
L3	harm_acute, harm_chronic_proxy, realized_harm_categories, cumulative_bias_index
L4	linked_vulnerability_id, exposure_duration_hours
L5	ai_system_type, non_adversarial_failure_type, agent_escalation_chain, mcp_endpoint_ids, oauth_pivot_chain, policy_enforcement_point, atlas_technique_id
L6	sector_code, nace_code_actual, naics_code_actual, jurisdiction_country_code, regulatory_obligations, regulatory_trigger

3. Profile Conformance

Type	Name	Field Rule
Base profile	Core	Requires L0, L1, L2, L6 baseline fields; Incident records require the L2 severity bundle. L3/L4 fields prohibited in primary payload.
Base profile	Enterprise	Core + L3/L4. Incident records additionally require realized_harm_categories.
Base profile	Regulatory	Enterprise + extended L6 obligations, trigger-confidence, and review fields where applicable.
Extension	SOC_SIEM	Permits atlas_technique_id and approved security-tooling references only when declared in extensions.
Extension	AI_Security	Permits selected L5 generative/agentic/control-plane fields only when declared in extensions.

4. Normative JSON Schema

The following JSON Schema is normative and is identical to the accompanying canonical machine-readable artifact uaif-core-v1.0.json.


```
{
  "$schema": "https://json-schema.org/draft/2020-12/schema",
  "$id": "https://schemas.oda3.org/uaif/core/v1.0/uaif-core-v1.0.json",
  "title": "UAIF Core Profile Schema v1.0",
  "$comment": "Normative errata build. Authority chain: TCR-STD-002 semantics -> TCR-STD-003 schema spec -> this canonical schema ->
/fixtures -> COM-012. See ODA3-2026-08-ERR-UAIF-001.",
  "type": "object",
  "required": [
    "incident_uuid",
    "incident_title",
    "incident_summary",
    "reporting_timestamp",
    "deduplication_hash",
    "remediation_status",
    "record_type",
    "profile",
    "root_cause_primary_domain",
    "root_cause_specific_type",
    "sector_code",
    "jurisdiction_country_code"
  ],
  "properties": {
    "incident_uuid": {
      "type": "string",
      "format": "uuid"
    },
    "incident_title": {
      "type": "string",
      "minLength": 5,
      "maxLength": 255
    },
    "incident_summary": {
      "type": "string",
      "minLength": 20,
      "maxLength": 4000
    },
    "reporting_timestamp": {
      "type": "string",
      "format": "date-time"
    },
    "reporter_confidence_level": {
      "type": "integer",
      "minimum": 1,
      "maximum": 5
    },
    "deduplication_hash": {
      "type": "string",
      "pattern": "^[a-fA-F0-9]{64}$"
    },
    "remediation_status": {
      "type": "string",
      "enum": [
        "Detected",
        "Investigating",
        "Mitigated",
        "Resolved"
      ]
    },
    "record_type": {
      "type": "string",
      "description": "L0. Ontological classification of the record. Two normative values only. Unclassified signals, alerts, telemetry, or
investigative leads are not a UAIF record_type; they remain in the originating detection/case-management workflow until classified.",
      "enum": [
        "Incident",
        "Vulnerability"
      ]
    },
```



```
    "profile": {
      "type": "string",
      "description": "L0. Base conformance depth declared by the record. See 'extensions' for additive specialized capabilities, which are
declared separately and do not affect this value.",
      "enum": [
        "Core",
        "Enterprise",
        "Regulatory"
      ]
    },
    "extensions": {
      "type": "array",
      "description": "L0. Optional, additive specialized capability declarations. Extensions never reduce or override base-profile
requirements.",
      "items": {
        "type": "string",
        "enum": [
          "SOC_SIEM",
          "AI_Security"
        ]
      },
      "uniqueItems": true
    },
    "root_cause_primary_domain": {
      "type": "string",
      "enum": [
        "Human",
        "System",
        "Data",
        "External"
      ]
    },
    "root_cause_specific_type": {
      "type": "string",
      "pattern": "^(Adversarial_Attack|Prompt_Injection|Model_Poisoning|Data_Poisoning|RAG_Leakage|Hallucination|Agent_Escalation|
Policy_Violation|System_Failure|Human_Error|Supply_Chain_Compromise|Unauthorised_Access|Data_Breach|Denial_of_Service|Misuse|x_[A-Za-z0-
9_]+)$"
    },
    "root_cause_confidence": {
      "type": "string",
      "enum": [
        "LOW",
        "MEDIUM",
        "HIGH"
      ]
    },
    "exploit_path_description": {
      "type": "string",
      "minLength": 20,
      "maxLength": 4000
    },
    "severity_analytical_score": {
      "type": "number",
      "minimum": 0.0,
      "maximum": 10.0,
      "description": "One-decimal-place precision (ROUND_HALF_UP) is a normative semantic/conformance rule per TCR-STD-002 Section 4.4,
verified by the reference implementation and conformance tooling -- deliberately NOT enforced via JSON Schema 'multipleOf', since binary
floating-point representation of decimal fractions (e.g. 2.8, 5.1, 9.1) can produce false negatives across validator implementations."
    },
    "severity_presentation_score": {
      "type": "number",
      "minimum": 0.0,
      "maximum": 10.0,
      "description": "One-decimal-place precision (ROUND_HALF_UP) is a normative semantic/conformance rule per TCR-STD-002 Section 4.4,
verified by the reference implementation and conformance tooling -- deliberately NOT enforced via JSON Schema 'multipleOf', for the same
floating-point interoperability reason as severity_analytical_score."
    },
    "severity_level": {
      "type": "integer",
      "minimum": 1,
      "maximum": 5
    },
    "severity_rationale": {
      "type": "object",
      "description": "L2. Sole canonical severity explanation. There is no separate severity_justification field in UAIF v1.0 --
superseded, see Section 9 legacy patterns.",
      "required": [
        "calibration_status"
      ],
      "properties": {
        "dominant_harm": {
          "type": "string",
          "enum": [
            "Physical_Harm",
            "Environmental_Harm",
            "Security_Integrity",
            "Privacy_Violation",
            "Financial_Loss",
            "Psychological_Harm",
            "Property_Damage",
            "Reputational_Harm"
          ]
        },
        "dominant_weight": {
          "type": "number"
        },
        "secondary_harms": {
          "type": "array",
          "items": {
            "type": "string",
            "enum": [
              "Physical_Harm",
              "Environmental_Harm",
              "Security_Integrity",
              "Privacy_Violation",
              "Financial_Loss",
              "Psychological_Harm",
              "Property_Damage",
```

```
              "Reputational_Harm"
            ]
          }
        },
        "secondary_contribution": {
          "type": "number"
        },
        "context_modifiers_applied": {
          "type": "array",
          "items": {
            "type": "object",
            "required": [
              "context",
              "base_value",
              "contribution_factor",
              "effective_contribution"
            ],
            "properties": {
              "context": {
                "type": "string"
              },
              "base_value": {
                "type": "number"
              },
              "contribution_factor": {
                "type": "number"
              },
              "effective_contribution": {
                "type": "number"
              }
            },
            "additionalProperties": false
          }
        },
        "context_modifier_sum": {
          "type": "number"
        },
        "population_multiplier": {
          "type": "number"
        },
        "blast_radius_multiplier": {
          "type": "number"
        },
        "impact_scale": {
          "type": "number"
        },
        "confidence_factor": {
          "type": "number"
        },
        "temporal_reduction": {
          "type": "number"
        },
        "temporal_adjustment": {
          "type": "number"
        },
        "calibration_status": {
          "type": "string",
          "enum": [
            "provisional",
            "sector_validated",
            "custom"
          ]
        },
        "calibration_source": {
          "type": [
            "string",
            "null"
          ],
          "format": "uri"
        },
        "weight_provisional_notice": {
          "type": [
            "string",
            "null"
          ]
        },
        "assumptions": {
          "type": "array",
          "items": {
            "type": "string"
          }
        },
        "precision_notes": {
          "type": "string"
        }
      },
      "additionalProperties": true
    },
    "harm_acute": {
      "type": "boolean"
    },
    "harm_chronic_proxy": {
      "type": "boolean"
    },
    "realized_harm_categories": {
      "type": "array",
      "description": "L3 (Extended Harm Characterisation). Required for Incident records only when profile is Enterprise or Regulatory
(see conditional rules). Prohibited under Core.",
      "items": {
        "type": "string",
        "enum": [
          "Physical_Harm",
          "Environmental_Harm",
          "Security_Integrity",
          "Privacy_Violation",
          "Financial_Loss",
          "Psychological_Harm",
          "Property_Damage",
```



```
        "Reputational_Harm"
      ]
    },
    "minItems": 1,
    "uniqueItems": true
  },
  "cumulative_bias_index": {
    "type": "number",
    "minimum": 0.0,
    "maximum": 1.0
  },
  "linked_vulnerability_id": {
    "type": "string",
    "description": "L4 (Incident-Vulnerability Relationship)."
  },
  "exposure_duration_hours": {
    "type": "number",
    "minimum": 0,
    "description": "L4 (Incident-Vulnerability Relationship)."
  },
  "ai_system_type": {
    "type": "string",
    "enum": [
      "LLM",
      "VLM",
      "RAG_Pipeline",
      "Agentic",
      "Multi_Agent",
      "Traditional_ML",
      "Other"
    ]
  },
  "non_adversarial_failure_type": {
    "type": "string"
  },
  "agent_escalation_chain": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "mcp_endpoint_ids": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "oauth_pivot_chain": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "policy_enforcement_point": {
    "type": "string"
  },
  "atlas_technique_id": {
    "type": "string",
    "pattern": "^AML\\.T[0-9]{4}(\\.[0-9]{3})?$"
  },
  "sector_code": {
    "type": "string",
    "enum": [
      "ENERGY",
      "TRANSPORT",
      "BANKING",
      "FINANCIAL_MARKETS",
      "HEALTH",
      "DRINKING_WATER",
      "WASTE_WATER",
      "DIGITAL_INFRA",
      "ICT_SERVICES",
      "PUBLIC_ADMIN",
      "SPACE",
      "POSTAL",
      "WASTE_MGMT",
      "CHEMICALS",
      "FOOD",
      "MANUFACTURING",
      "TELECOM",
      "RESEARCH",
      "OTHER"
    ]
  },
  "nace_code_actual": {
    "type": "string",
    "pattern": "^[A-Z](\\d{2}(\\.\\d{1,2})?)?$"
  },
  "naics_code_actual": {
    "type": "string",
    "pattern": "^\\d{2,6}$"
  },
  "jurisdiction_country_code": {
    "type": "string",
    "pattern": "^[A-Z]{2}(-[A-Z0-9]{1,3})?$"
  },
  "regulatory_obligations": {
    "type": "array",
    "items": {
      "type": "string",
      "enum": [
        "EU_AI_Act_Article_73",
        "GDPR_Article_33",
        "NIS2_Article_23",
        "DORA_Article_19",
        "Sectoral_Reporting",
        "None"
      ]
```

```
      },
      "minItems": 1,
      "uniqueItems": true
    },
    "regulatory_trigger": {
      "type": "boolean"
    },
    "trigger_confidence": {
      "type": "string",
      "enum": [
        "LOW",
        "MEDIUM",
        "HIGH"
      ]
    },
    "human_review_recommended": {
      "type": "boolean"
    },
    "sensitive_data": {
      "type": "boolean"
    },
    "critical_sector": {
      "type": "boolean"
    },
    "vulnerable_population": {
      "type": "boolean"
    },
    "x_extensions": {
      "type": "object",
      "additionalProperties": true
    }
  },
  "allOf": [
    {
      "$comment": "Universal Incident rule: every Incident record, regardless of profile, requires the L2 severity bundle.",
      "if": {
        "properties": {
          "record_type": {
            "const": "Incident"
          }
        },
        "required": [
          "record_type"
        ]
      },
      "then": {
        "required": [
          "severity_analytical_score",
          "severity_presentation_score",
          "severity_level",
          "severity_rationale"
        ]
      }
    },
    {
      "$comment": "Issue 5 (locked, option a): Enterprise/Regulatory Incident records additionally require realized_harm_categories. Core
Incident records do not.",
      "if": {
        "properties": {
          "record_type": {
            "const": "Incident"
          },
          "profile": {
            "enum": [
              "Enterprise",
              "Regulatory"
            ]
          }
        },
        "required": [
          "record_type",
          "profile"
        ]
      },
      "then": {
        "required": [
          "realized_harm_categories"
        ]
      }
    },
    {
      "$comment": "Core prohibition: L3/L4 fields are blocked under profile=Core regardless of extensions (extensions only govern
L5/mapping fields, not L3/L4).",
      "if": {
        "properties": {
          "profile": {
            "const": "Core"
          }
        },
        "required": [
          "profile"
        ]
      },
      "then": {
        "not": {
          "anyOf": [
            {
              "required": [
                "harm_acute"
              ]
            },
            {
              "required": [
                "harm_chronic_proxy"
              ]
            },
            {
              "required": [
```

```
                "realized_harm_categories"
              ]
            },
            {
              "required": [
                "cumulative_bias_index"
              ]
            },
            {
              "required": [
                "linked_vulnerability_id"
              ]
            },
            {
              "required": [
                "exposure_duration_hours"
              ]
            }
          ]
        }
      }
    },
    {
      "$comment": "Extension gating: atlas_technique_id requires extensions to explicitly contain SOC_SIEM. Implemented with 'required' on
extensions to avoid the absent-property ambiguity flagged during reconciliation.",
      "if": {
        "properties": {
          "extensions": {
            "contains": {
              "const": "SOC_SIEM"
            }
          }
        },
        "required": [
          "extensions"
        ]
      },
      "else": {
        "not": {
          "required": [
            "atlas_technique_id"
          ]
        }
      }
    },
    {
      "$comment": "Extension gating: L5 generative/agentic/control-plane fields require extensions to explicitly contain
 AI_Security.",
      "if": {
        "properties": {
          "extensions": {
            "contains": {
              "const": "AI_Security"
            }
          }
        },
        "required": [
          "extensions"
        ]
      },
      "else": {
        "not": {
          "anyOf": [
            {
              "required": [
                "agent_escalation_chain"
              ]
            },
            {
              "required": [
                "mcp_endpoint_ids"
              ]
            },
            {
              "required": [
                "oauth_pivot_chain"
              ]
            },
            {
              "required": [
                "policy_enforcement_point"
              ]
            },
            {
              "required": [
                "ai_system_type"
              ]
            },
            {
              "required": [
                "non_adversarial_failure_type"
              ]
            }
          ]
        }
      }
    },
    {
      "$comment": "Regulatory trigger requires confidence + human-review companions.",
      "if": {
        "properties": {
          "regulatory_trigger": {}
        },
        "required": [
          "regulatory_trigger"
        ]
      },
      "then": {
```



```
        "required": [
          "trigger_confidence",
          "human_review_recommended"
        ]
      }
    },
    {
      "$comment": "Adversarial attack root cause requires exploit path description.",
      "if": {
        "properties": {
          "root_cause_specific_type": {
            "const": "Adversarial_Attack"
          }
        },
        "required": [
          "root_cause_specific_type"
        ]
      },
      "then": {
        "required": [
          "exploit_path_description"
        ]
      }
    },
    {
      "$comment": "Provisional calibration requires the provisional notice.",
      "if": {
        "properties": {
          "severity_rationale": {
            "properties": {
              "calibration_status": {
                "const": "provisional"
              }
            },
            "required": [
              "calibration_status"
            ]
          }
        },
        "required": [
          "severity_rationale"
        ]
      },
      "then": {
        "properties": {
          "severity_rationale": {
            "required": [
              "weight_provisional_notice"
            ]
          }
        }
      }
    }
  ],
  "unevaluatedProperties": false
}
```

5. Conditional Validation Rules

Rule	Enforcement
Incident records	Any Incident MUST include severity_analytical_score, severity_presentation_score, severity_level, severity_rationale.
Enterprise/Regulatory Incident	Additionally requires realized_harm_categories.
Core profile	L3/L4 fields prohibited regardless of extensions.
SOC/SIEM mapping	atlas_technique_id requires extensions to contain SOC_SIEM.
AI Security fields	Selected L5 agentic/control-plane fields require extensions to contain AI_Security.
Regulatory trigger	regulatory_trigger requires trigger_confidence and human_review_recommended.
Provisional calibration	severity_rationale.calibration_status = provisional requires weight_provisional_notice.
Adversarial attack	root_cause_specific_type = Adversarial_Attack requires exploit_path_description.
Custom root cause	Custom root_cause_specific_type values MUST use x_ prefix.
Sector taxonomy	sector_code MUST be one of the 19 normative UAIF sector codes; NAICS/NACE belong in mapping metadata fields.

6. Severity Matrix

Level	Score Boundary	Rule
1 Info	0.0 <= score < 2.0	No external harm

Level	Score Boundary	Rule
2 Low	2.0 <= score < 4.0	Low operational harm
3 Medium	4.0 <= score < 6.5	Material but generally non-critical harm
4 High	6.5 <= score < 8.5	High impact; potential Article 73 routing proxy where applicable
5 Critical	8.5 <= score <= 10.0	Critical/systemic impact; Article 73 review required where applicable

7. Reference Severity Calculator

The reference implementation mirrors TCR-STD-002 Section 4. Score precision is enforced semantically using Decimal ROUND_HALF_UP rather than JSON Schema multipleOf, avoiding cross-validator binary floating-point false negatives.

```
from decimal import Decimal, ROUND_HALF_UP

def _round_decimal(value, places=1):
    quantum = Decimal("1." + ("0" * places)) if places else Decimal("1")
    return float(Decimal(str(value)).quantize(quantum, rounding=ROUND_HALF_UP))

# Implementations SHALL apply the complete reference calculation in TCR-STD-002 Section 4.2.
# Conformance tooling SHALL verify one-decimal-place score precision separately from JSON Schema validation.
```

8. Sample Valid Incident Payload

```
{
  "incident_uuid": "55555555-5555-5555-5555-555555555555",
  "incident_title": "RAG tenant isolation failure disclosed financial records",
  "incident_summary": "A retrieval-augmented generation system returned records from another tenant after a missing tenant-scoped filter
in the vector database query path.",
  "reporting_timestamp": "2026-06-24T12:00:00Z",
  "deduplication_hash": "3dc44f6cf6cbab218d18cf8e0206d3708b2c4a5b1e646c54755dcd98d6a6f864",
  "remediation_status": "Investigating",
  "record_type": "Incident",
  "profile": "Enterprise",
  "root_cause_primary_domain": "System",
  "root_cause_specific_type": "RAG_Leakage",
  "sector_code": "BANKING",
  "jurisdiction_country_code": "DE",
  "severity_analytical_score": 6.8,
  "severity_presentation_score": 7.8,
  "severity_level": 4,
  "severity_rationale": {
    "calibration_status": "provisional",
    "weight_provisional_notice": "Default harm weights are provisional pending empirical validation."
  },
  "realized_harm_categories": [
    "Privacy_Violation",
    "Financial_Loss",
    "Reputational_Harm"
  ]
}
```

9. Invalid Legacy Payload Patterns

Pattern	Status / Corrective Action
record_type = "Observation"	Not a normative value in UAIF v1.0; retain the signal in upstream workflow until classified as Incident or Vulnerability.
severity_justification present	Field removed; severity_rationale is the sole canonical severity explanation.
Core payload contains L3/L4 field	Invalid Core payload; select Enterprise/Regulatory where appropriate or remove the field.
atlas_technique_id without SOC_SIEM extension	Invalid; declare SOC_SIEM extension or omit the field.
AI Security L5 field without AI_Security extension	Invalid; declare AI_Security extension or omit the field.
Score field with more than one decimal place	Fails normative semantic precision requirement even if generic JSON Schema validation otherwise succeeds.

10. Regulatory Translation Boundary

UAIF regulatory routing fields provide technical routing assistance only. They do not determine whether a legal reporting obligation exists, replace sector-specific legal analysis, or establish that a report has been accepted by a regulator. EU AI Act serious-incident references use Article 73. GDPR, NIS2, DORA, and sectoral obligations remain separate legal determinations.

11. Conformance Validation Record

Final Publication validation used a 15-case corpus containing 8 intended-valid and 7 intended-invalid payloads. The canonical Draft 2020-12 schema was checked using jsonschema 4.26.0 and produced 15/15 expected outcomes. The purpose-built run_conformance.py harness independently produced 15/15 expected outcomes and separately verifies the normative one-decimal score precision rule.

Validation Layer	Result
Draft 2020-12 schema structural check	PASS
Generic jsonschema 4.26.0 fixture run	15/15 expected outcomes
Purpose-built conformance harness	15/15 expected outcomes
Score precision semantic check	PASS across corpus
Known normative architecture blockers	None identified at release gate

