

ODA3 INSTITUTE

TECHNICAL REFERENCE

UAIF Reference Incident Records

Fully worked, synthetic example records for implementers

TCR-UAIF-018 | Public

© ODA3 Pvt Ltd

Document Control

Item	Value
Document Identifier	TCR-UAIF-018 (legacy identifier retained for citation continuity)
Doc Type	Technical Reference - Worked Examples
Framework Tag	UAIF™ v1.0
Schema	uaif-core-v1.0.json (JSON Schema Draft 2020-12), SHA-256 55056e7cfa4fc4b76cc4656833cb5b360a1b40fa90f23488d488409ccebeeca
Authoritative Schema Reference	ODA3-2026-06-TCR-STD-003 - UAIF™ v1.0 Schema Specification
Companion Documents	TCR-UAIF-014 Data Dictionary and TCR-UAIF-015 Validation Rules Catalogue - pending schema-alignment review; do not treat as authoritative until reissued
Regulatory Crosswalk	Not applicable - synthetic worked examples
Evidence Tier / Provenance	Not applicable - all records are synthetic
Audience	Implementers building intake tooling; analysts learning the schema
Publication History	Issued 14 July 2026; updated 04 October 2026
Review Cycle	On any schema correction; otherwise annual
Licence	GAISSF Ecosystem Licence (GEL) v1.0
Contact	contact@oda3.org
Classification	Public

UAIF Reference Incident Records

Fully worked, synthetic example records for implementers

SYNTHETIC DATA NOTICE

Every record in this document is synthetic. None describes a real organisation, system or event. All three records validate against the canonical UAIF™ v1.0 schema, uaif-core-v1.0.json, with the normative errata integrated (ODA3-2026-08-ERR-UAIF-001).

1. Purpose and Scope

This document walks through three fully worked, schema-valid UAIF™ v1.0 records. Across the three examples, it illustrates both Incident and Vulnerability records, Core and Enterprise profiles, and different root-cause patterns, with a plain-language explanation of why each field is populated as shown. It is a learning aid for implementers, not a conformance test. The authoritative schema rules are in ODA3-2026-06-TCR-STD-003 and the canonical uaif-core-v1.0.json schema.

TCR-UAIF-014 (Data Dictionary) and TCR-UAIF-015 (Validation Rules Catalogue) are companion references but are currently pending schema-alignment review. Until they are reissued, resolve any discrepancy in favour of TCR-STD-003 and the canonical schema.

Three schema rules shape every record below:

- profile is required - Core, Enterprise or Regulatory. Core records may not carry L3/L4 fields; Enterprise and Regulatory Incident records must carry realized_harm_categories.
- severity_rationale is the only severity explanation. There is no severity_justification field in UAIF v1.0; a record containing it is invalid.
- Extensions are declared, not assumed. L5 fields such as ai_system_type are valid only when extensions contains AI_Security; atlas_technique_id requires SOC_SIEM.

2. Example 1 - Minimal Valid Core Incident

The smallest record that satisfies every unconditional and conditional requirement for a Core Incident. Useful as a starting template.

```
{
  "incident_uuid": "3f9c7a2e-8b1d-4e6a-9c2f-1a5d6e7b8c90",
  "incident_title": "RAG assistant discloses unrelated customer record",

```

```

"incident_summary": "A customer-support RAG assistant retrieved and displayed a record belonging to a different customer after a retrieval-scoping defect. No downstream action was taken by the customer before the session was terminated.",
"reporting_timestamp": "2026-07-10T09:15:00Z",
"deduplication_hash": "a3f5e8b2c9d1e4f6a7b8c9d0e1f2a3b4c5d6e7f8a9b0c1d2e3f4a5b6c7d8e9f0",
"remediation_status": "Mitigated",
"record_type": "Incident",
"profile": "Core",
"root_cause_primary_domain": "System",
"root_cause_specific_type": "RAG_Leakage",
"sector_code": "ICT_SERVICES",
"jurisdiction_country_code": "US-CA",
"severity_analytical_score": 3.2,
"severity_presentation_score": 3.2,
"severity_level": 2,
"severity_rationale": {
  "calibration_status": "provisional",
  "weight_provisional_notice": "Cross-sector baseline weights; no sector-specific calibration package bound to this record."
}
}

```

Why this validates: all 12 unconditional required fields are present, including profile. record_type = "Incident" triggers the L2 severity bundle - severity_analytical_score, severity_presentation_score, severity_level and severity_rationale - all present. calibration_status = "provisional" triggers weight_provisional_notice, which is present. Because the profile is Core, realized_harm_categories is omitted: it is an L3 field, prohibited under Core. Removing any field shown breaks validation.

3. Example 2 - Enterprise Adversarial Incident with Full Severity Rationale

Demonstrates the Adversarial_Attack rule (exploit_path_description), a fully populated severity_rationale with context modifiers, the regulatory-trigger companions, the Enterprise harm-category rule and the AI_Security extension.

```

{
  "incident_uuid": "7b2e4f1a-9c3d-4e8b-a1f2-6d7c8b9a0e1f",
  "incident_title": "Prompt injection via uploaded document triggers unauthorised refund",
  "incident_summary": "Malicious instructions embedded in an uploaded document caused a support agent's tool-calling model to issue an unapproved refund and disclose an unrelated booking record to the requester.",
  "reporting_timestamp": "2026-07-11T14:22:00Z",
  "deduplication_hash": "b4f6a9c3d2e5f8a1b7c9d0e2f4a6b8c1d3e5f7a9b1c3d5e7f9a1b3c5d7e9f1a3",
  "remediation_status": "Resolved",
  "record_type": "Incident",
  "profile": "Enterprise",
  "extensions": [
    "AI_Security"
  ],
  "root_cause_primary_domain": "External",
  "root_cause_specific_type": "Adversarial_Attack",
  "exploit_path_description": "A guest-message upload contained hidden instructions that the assistant's document-parsing step passed unfiltered into the model context, causing it to invoke the refund tool with fabricated authority and to retrieve a second party's booking record via the same tool-calling session.",
  "sector_code": "OTHER",
  "jurisdiction_country_code": "GB",
  "severity_analytical_score": 5.8,
  "severity_presentation_score": 5.8,
  "severity_level": 3,
  "severity_rationale": {
    "dominant_harm": "Financial_Loss",
    "dominant_weight": 3.5,
    "secondary_harms": [
      "Privacy_Violation",
      "Security_Integrity"
    ]
  },
  "secondary_contribution": 1.4,
  "context_modifiers_applied": [
    {
      "context": "Prompt_Injection",
      "base_value": 1.2,
      "contribution_factor": 0.8,
      "effective_contribution": 0.96
    },
    {
      "context": "Agent_Escalation",
      "base_value": 1.0,
      "contribution_factor": 0.6,
      "effective_contribution": 0.6
    }
  ],
  "context_modifier_sum": 1.56,
  "calibration_status": "provisional",
  "weight_provisional_notice": "Cross-sector baseline weights; no sector-specific calibration package bound to this record."
},
  "realized_harm_categories": [
    "Financial_Loss",
    "Privacy_Violation",
    "Security_Integrity"
  ],
  "ai_system_type": "Agentic",
  "regulatory_trigger": true,
  "trigger_confidence": "MEDIUM",
  "human_review_recommended": true,
  "regulatory_obligations": [
    "GDPR_Article_33"
  ]
}

```

```
]
}
```

Why this validates: profile = "Enterprise" with record_type = "Incident" requires realized_harm_categories, which is present. root_cause_specific_type = "Adversarial_Attack" requires exploit_path_description. regulatory_trigger is present, so trigger_confidence and human_review_recommended are required and present. ai_system_type is an L5 field, permitted because extensions contains AI_Security. Each context_modifiers_applied entry carries exactly its four required sub-fields. The regulatory trigger is a routing aid only; whether GDPR Article 33 actually applies remains a legal determination.

4. Example 3 - Core Vulnerability Record

Demonstrates that the severity bundle is not required when record_type is Vulnerability - useful for logging a discovered weakness before it manifests as an incident.

```
{
  "incident_uuid": "1d8e3a5c-6f2b-4a9d-8c1e-3f5a7b9c1d3e",
  "incident_title": "Missing rate limit on agent tool-invocation endpoint",
  "incident_summary": "Internal review identified that the agent's tool-invocation endpoint has no rate limit, which could allow rapid repeated tool calls if a future prompt injection succeeded. No exploitation has occurred.",
  "reporting_timestamp": "2026-07-12T11:00:00Z",
  "deduplication_hash": "c5a7b9d1e3f5a7c9b1d3e5f7a9c1b3d5e7f9a1c3b5d7e9f1a3c5b7d9e1f3a5c7",
  "remediation_status": "Investigating",
  "record_type": "Vulnerability",
  "profile": "Core",
  "root_cause_primary_domain": "System",
  "root_cause_specific_type": "System_Failure",
  "sector_code": "OTHER",
  "jurisdiction_country_code": "DE"
}
```

Why this validates: record_type = "Vulnerability", so the severity bundle is optional and omitted. profile = "Core" is declared, as every record requires. No severity explanation field is needed; severity_rationale would be used if severity were scored.

5. Notably Absent

- These records are synthetic and do not represent any real organisation's incident history.
- They do not demonstrate every field combination, the Regulatory profile, or the SOC_SIEM extension. Until TCR-UAIF-014 and TCR-UAIF-015 complete schema-alignment review, use TCR-STD-003 and the canonical schema as the authoritative rule source.
- They are not the 15-case conformance corpus and do not test the one-decimal score-precision rule, which is verified outside JSON Schema.
- They do not demonstrate sector-specific x_extensions vocabulary; see an applicable sector calibration package only after that package is separately validated and published.
- The severity scores are illustrative only and do not reflect calibrated or sector-validated weighting.
- Records using Observation, severity_justification, or omitting profile come from the withdrawn pre-errata schema and are invalid in UAIF™ v1.0.

6. Revision Note - 04 October 2026

This reissue preserves legacy document identifier TCR-UAIF-018 for citation continuity. It replaces the withdrawn July schema basis with the frozen canonical uaif-core-v1.0.json schema (SHA-256 55056e7cfa4fc4b76cc4656833cb5b360a1b40fa90f23488d488409ccebeeca); adds required profile values; removes severity_justification; corrects Core/Enterprise field usage; declares extension-gated fields explicitly; updates contact and classification wording; and records the dependency status of TCR-UAIF-014 and TCR-UAIF-015. No new UAIF normative behavior is introduced.

GAISSF Ecosystem developed by ODA3 Institute, licensed under GEL v1.0.