

ODA3 INSTITUTE

IMPLEMENTATION AID

UAIF Incident Reporting Template

Human Intake Form for Populating a UAIF™ v1.0 Record

IMP-UAIF-043 | Public

© ODA3 Pvt Ltd

Document Control

Item	Value
Document Identifier	IMP-UAIF-043 (legacy identifier retained for citation continuity)
Doc Type	Implementation Aid - Human Intake Template
Framework Tag	UAIF™ v1.0
Canonical Schema	uaif-core-v1.0.json (JSON Schema Draft 2020-12)
Canonical SHA-256	55056e7cfa4fc4b76cc4656833cb5b360a1b40fa90f23488d488409ccebeeca
Authoritative References	ODA3-2026-06-TCR-STD-002; ODA3-2026-06-TCR-STD-003; TCR-UAIF-014; TCR-UAIF-015
Publication History	Issued 14 July 2026; updated 04 October 2026
Review Cycle	On every schema correction; otherwise annual
Audience	First-line analysts and reporters collecting information for UAIF records
Licence	GAISSF Ecosystem Licence (GEL) v1.0
Contact	contact@oda3.org
Classification	Public

HOW TO USE THIS TEMPLATE

This is a human-facing intake aid. It collects the information needed to build a UAIF™ v1.0 JSON record, but the completed form is not itself a valid submission. A downstream process must encode the record as JSON, generate or verify system-derived identifiers where applicable, and validate it against the canonical schema plus the separate semantic precision rule.

1. Section A - Always Required for Every UAIF Record

Complete or generate all twelve items below for every Incident or Vulnerability record. Items marked 'system-generated permitted' may be generated by the intake or case-management system rather than typed by the analyst.

Field	Guidance	Your entry
Incident UUID	Required UUID. May be system-generated. Preserve the same identifier across revisions of the same record.	
Incident title	5-255 characters. Short, factual description; avoid legal or severity conclusions unless already established.	
Incident summary	20-4000 characters. Describe what happened or what weakness was identified. Do not include credentials, secrets, unnecessary personal data, or unredacted third-party confidential material.	
Reporting timestamp	ISO 8601 date-time, e.g. 2026-10-04T09:15:00Z.	
Deduplication hash	Required 64-character hexadecimal value. May be system-generated according to the implementation's documented deduplication process.	
Remediation status	Detected / Investigating / Mitigated / Resolved.	
Record type	Incident / Vulnerability. Observation is not a normative UAIF™ v1.0 record type.	
Profile	Core / Enterprise / Regulatory. This is the base conformance depth and is always required.	
Root cause - primary domain	Human / System / Data / External.	
Root cause - specific type	Use one of the 15 canonical values in Section B, or an x_-prefixed custom value matching the schema pattern.	
Sector	Use one canonical UAIF sector_code. Use OTHER where appropriate; NACE/NAICS codes are separate optional mapping fields.	
Jurisdiction	ISO 3166-1 alpha-2 country code with optional subdivision, e.g. US-CA, GB, DE.	

2. Section B - Optional Baseline Metadata

Field	Guidance	Your entry
Reporter confidence level	Optional integer 1-5.	
Root cause confidence	Optional LOW / MEDIUM / HIGH.	
NACE code	Optional mapping metadata; must match the canonical nace_code_actual pattern.	
NAICS code	Optional mapping metadata; 2-6 digits.	
Sensitive data involved	Optional true/false.	
Critical sector	Optional true/false.	
Vulnerable population	Optional true/false.	

3. Section C - Root Cause Specific Type Reference

Value	Use when...
Adversarial_Attack	Deliberate attacker action caused or contributed to the event. Requires an exploit path description.
Prompt_Injection	Untrusted content altered model or agent behaviour through its input/context.
Model_Poisoning	Training, fine-tuning or feedback data was manipulated to corrupt model behaviour.
Data_Poisoning	Non-model data such as a reference dataset or retrieval corpus was manipulated.
RAG_Leakage	Retrieval-augmented generation exposed content it should not have.
Hallucination	The system generated fabricated output not grounded in its inputs.
Agent_Escalation	An autonomous or semi-autonomous agent exceeded intended authority or scope.
Policy_Violation	The system violated an internal policy or guardrail without an external attack.
System_Failure	Conventional technical fault, bug, outage or misconfiguration.
Human_Error	A person's mistake caused or contributed to the event.
Supply_Chain_Compromise	A vendor, dependency or third-party component was the source of compromise.
Unauthorised_Access	Someone accessed a system or data without authorisation.
Data_Breach	Confirmed or suspected unauthorised disclosure of data.
Denial_of_Service	Availability was degraded or denied.
Misuse	A legitimate user used the system in an unintended or prohibited way.
x_[custom]	Use only where none of the canonical values fit; value must begin with x_ and satisfy the schema pattern.

4. Section D - Required for record_type = Incident

Every Incident, regardless of profile, requires the four-field L2 severity bundle below. The scoring semantics are defined in ODA3-2026-06-TCR-STD-002.

Field	Guidance	Your entry
Analytical severity score	Required number 0.0-10.0. Final representation must use one decimal place under the UAIF semantic precision rule.	
Presentation severity score	Required number 0.0-10.0. Final representation must use one decimal place under the UAIF semantic precision rule.	
Severity level	Required integer 1-5. Canonical score bands: 1 = <2.0; 2 = 2.0-<4.0; 3 = 4.0-<6.5; 4 = 6.5-<8.5; 5 = >=8.5.	
Calibration status	Required inside severity_rationale: provisional / sector_validated / custom.	
Provisional notice	Required inside severity_rationale whenever calibration_status = provisional.	

Optional severity_rationale details may include dominant_harm, dominant_weight, secondary_harms, secondary_contribution, context modifiers, impact/confidence/temporal adjustments, calibration source, assumptions and precision notes. These are not all mandatory merely because the record is an Incident.

5. Section E - Profile-Specific Harm / Exposure Rules

Profile	Rule
Core	Do not populate harm_acute, harm_chronic_proxy, realized_harm_categories, cumulative_bias_index, linked_vulnerability_id, or exposure_duration_hours. These L3/L4 fields are prohibited under Core.
Enterprise	If record_type = Incident, realized_harm_categories is required. Other allowed L3/L4 fields remain optional unless another applicable requirement makes them necessary.
Regulatory	If record_type = Incident, realized_harm_categories is required. Other allowed L3/L4 fields remain optional unless another applicable requirement makes them necessary.

When required, realized_harm_categories must contain one or more of: Physical_Harm, Environmental_Harm, Security_Integrity, Privacy_Violation, Financial_Loss, Psychological_Harm, Property_Damage, Reputational_Harm.

6. Section F - Required for Adversarial_Attack

Exploit path description (20-4000 characters): describe the causal chain from entry point to effect. It is required whenever root_cause_specific_type = Adversarial_Attack.

7. Section G - Extensions and Extension-Gated Fields

Extensions are additive and must be declared explicitly. Populate the extensions array only when the record uses the corresponding specialized fields.

Extension		When required
SOC_SIEM		Declare SOC_SIEM before using atlas_technique_id.
AI_Security		Declare AI_Security before using ai_system_type, non_adversarial_failure_type, agent_escalation_chain, mcp_endpoint_ids, oauth_pivot_chain, or policy_enforcement_point.
Field	Guidance	Your entry
extensions	Optional array containing SOC_SIEM and/or AI_Security; values must be unique.	
AI system type	Optional: LLM / VLM / RAG_Pipeline / Agentic / Multi_Agent / Traditional_ML / Other. Requires AI_Security extension.	
ATLAS technique ID	Optional MITRE ATLAS technique identifier matching the schema pattern. Requires SOC_SIEM extension.	

8. Section H - Regulatory-Routing Fields

If regulatory_trigger is present at all - whether true or false - trigger_confidence and human_review_recommended are both required.

Field	Guidance	Your entry
Regulatory trigger	Optional true/false. If included, complete the next two fields.	
Trigger confidence	Required when regulatory_trigger is present: LOW / MEDIUM / HIGH.	
Human review recommended	Required when regulatory_trigger is present: true / false.	
Regulatory obligations	Optional array of routing tags: EU_AI_Act_Article_73, GDPR_Article_33, NIS2_Article_23, DORA_Article_19, Sectoral_Reporting, or None.	

LEGAL / REGULATORY BOUNDARY

These fields provide technical routing assistance only. They do not determine whether a reporting obligation legally applies, do not calculate a statutory deadline, and do not constitute a regulatory filing or legal notification.

9. Section I - x_extensions

x_extensions is an optional open object for separately governed sector- or organisation-specific data. Do not add undeclared custom fields at the top level: the UAIF top-level object is closed and unexpected top-level properties fail validation.

10. Before You Encode or Submit

- Confirm all twelve always-required fields are populated or generated.
- Confirm profile-specific prohibitions and requirements: Core prohibits L3/L4 fields; Enterprise/Regulatory Incidents require realized_harm_categories.
- Confirm every extension-gated field has the corresponding SOC_SIEM or AI_Security declaration.
- If regulatory_trigger is present, confirm trigger_confidence and human_review_recommended are also present.
- If root_cause_specific_type = Adversarial_Attack, confirm exploit_path_description is present.
- If calibration_status = provisional, confirm weight_provisional_notice is present.
- Do not use Observation or severity_justification; both belong to the withdrawn pre-errata model.
- Validate the encoded JSON against uaif-core-v1.0.json and separately verify one-decimal score precision.
- Do not include credentials, authentication secrets, or unnecessary personal/confidential information in narrative fields.

11. Notably Absent

- This form is not a machine-readable UAIF submission and does not replace JSON Schema validation.
- The form does not itself calculate the UAIF severity scores. The normative calculation semantics are published in ODA3-2026-06-TCR-STD-002.
- It does not treat realized_harm_categories as universally required for all Incidents.
- It does not permit Observation records or the removed severity_justification field.
- It does not assume optional AI/SIEM fields are valid without the matching extension declaration.
- It does not create regulatory, legal, certification, external-assurance or approval status.
- Certification is not currently offered. Any future UAIF certification programme would require separate approved and published governance.

12. Revision Note - 04 October 2026

This corrected reissue preserves legacy identifier IMP-UAIF-043 for citation continuity and aligns the template to the frozen canonical uaif-core-v1.0.json schema (SHA-256 55056e7cfa4fcf4b76cc4656833cb5b360a1b40fa90f23488d488409ccebceca). It removes Observation and severity_justification; adds incident_uuid, deduplication_hash and profile; separates the universal Incident severity bundle from Enterprise/Regulatory harm-category requirements; adds the Core L3/L4 prohibition; adds SOC_SIEM and AI_Security extension declarations/gating; corrects the score-precision and severity-level guidance; replaces the blanket 'unknown convention' instruction with schema-specific requirements; updates Public classification and contact details; and removes the incorrect statement that the published normative scoring logic is controlled. No new UAIF normative behavior is introduced.

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