

L30_FRM_I02: Incident Review Protocol

L30_FRM Practical Sheet DocID: L30-FRM-I02 (v1.0.1)

Structured protocol for post-incident assessment of human refusal authority in AI systems, agents, or processes.

This form uses the L30-BAS structure and is not certification, legal compliance, deployment approval, safety approval, or proof of absence of risk.

Review Date	
Reviewer or Organization	
System or Case ID	
Purpose	Detailed post-incident review protocol
L30-BAS Status	Supporting / Detailed Form
Primary Starter Forms	L30_FRM_I01
Referenced Core Codes	L30-BX-01, L30-BX-02, L30-BX-03, L30-BX-04, L30-CI

Protocol Core — Mandatory

Purpose: To evaluate whether human refusal authority remained valid prior to irreversible external impact.

Primary Review Question

Was human refusal authority preserved in a way that allowed meaningful human intervention and the ability to stop the system before irreversible real-world impact occurred?

Interpretation guidance: if a human could not realistically intervene and stop the system in time, refusal authority was not preserved. This includes cases where intervention was technically possible but practically ineffective.

Evaluation Outcome

Effective

Not Effective

Minimum Requirement (G06): Human refusal must be exercisable and must allow a human to intervene and stop the system in time before irreversible external impact.

Absence Rule: If this condition is not demonstrably satisfied, including cases of uncertainty or insufficient evidence, the system is classified as Not Effective.

Evaluation Scope: This protocol applies exclusively to post-incident review. It must not be used for system design, optimization, or predictive control.

Decision Authority: Final classification must be explicitly assigned to a named reviewer or governing body.

Evidence Rule: Evaluation must be based only on documented, verifiable information. Assumptions must not be used to justify effectiveness.

Temporal Reference: Evaluation must be based on information available up to the point of irreversible external impact.

Applicability Rule: Any incident involving potential irreversible external impact must be evaluated under this protocol. Claims of non-applicability are considered invalid.

Exclusivity Rule: This evaluation must not be overridden or supplemented by external frameworks.

Scope Enforcement: Any use of this protocol outside post-incident review invalidates the evaluation.

Record Integrity Rule: This record must be append-only. Post-hoc modification of prior entries is not permitted.

Analytical Extension — Optional but Recommended

Evaluation Breakdown

A human decision-maker was identifiable

Refusal authority was accessible

Intervention could be executed in time

Intervention Attempted	Yes / No
Method	
Outcome	

Failure Point Identification

Not present

Not accessible

Too late

Overridden

Other

- System speed / autonomy level
- Human latency / response delay
- Organizational or procedural constraints

Conceptual Clarification

Human refusal authority refers to the ability of a human to actively intervene, halt, or prevent system actions. It must be operational in practice, not merely defined in principle.

Irreversibility Note

Irreversible external impact includes outcomes that cannot be undone, contained, or meaningfully mitigated once triggered.

Reviewer Discretion: This protocol defines minimum conditions. Reviewers may include additional observations or context where necessary.

Modification and Attribution Policy

- This protocol may be copied, extended, or modified.
- The name LUMINA-30 must be retained in all copies and derivatives.
- Only the unmodified version may be identified as LUMINA-30 Incident Review Protocol (Original).
- Any modified version must be clearly labeled as Modified version based on LUMINA-30 and must not be presented as the original protocol.
- Implementations may include a reference source or version identifier to clarify alignment with the original protocol.

Reviewer Notes