

Autonomy Audit Framework

Reference Framework (v1.0)

1. Purpose

The Autonomy Audit Framework defines a structured method for evaluating whether an autonomous or intelligent system produces sufficient operational evidence to support claims of legitimacy, safety, and accountability.

2. Why Autonomy Requires Audit

As autonomy increases, intent and outcomes may diverge over time. Auditability ensures that delegated authority remains bounded and observable during operation.

3. Scope of an Autonomy Audit

Audits evaluate runtime behavior, boundary adherence, human override mechanisms, and evidence continuity rather than design claims alone.

4. Evidence Classes

Evidence may include operational logs, decision traces, override records, boundary violations, and remediation actions.

5. Audit Outcomes

Audit results may include pass, conditional acceptance, remediation required, or revocation of operational authority.

6. Relationship to Operational Legitimacy

The Autonomy Audit Framework provides the evidentiary backbone required to assess ongoing Operational Legitimacy.

7. Status and Stewardship

This document is published as a neutral, non-commercial reference (v1.0). Stewardship is maintained by the Operational Legitimacy Working Group.