

AI System Intake & Release Decision Record

A reusable record for registering an AI system, defining its authority and dependency boundary, capturing evaluation evidence, and recording a bounded release, pause, rollback, or retirement decision.

Use this record to	Track one reviewed AI system state from intake through release, material change, and retirement.
Decision options	Proposed · Discovery · Pilot · Conditional approval · Approved for bounded use · Paused · Retired
Public-safe boundary	Do not enter confidential, regulated, classified, privileged, source-selection-sensitive, production-secret, or credential information into a public copy.

1. Identity, ownership, and lifecycle state

System / use-case name	Record ID and version
Business owner	Technical owner
Risk / policy owner	Release authority
Current lifecycle state	Next review date or trigger
Record location / linked evidence	

2. Intended use and prohibited use

Business problem, expected benefit, and decision supported
Approved users, context, geography, environment, and operating conditions
Prohibited uses, unsupported decisions, and out-of-scope people or situations
Assumptions, unknowns, and conditions that invalidate the current approval

Completion signal

- A named owner can explain the intended use and its limits.
- The current state and next review trigger are explicit.
- A prohibited-use statement exists and can be enforced in practice.

3. System, data, and dependency boundary

Record the reviewed state. Link to versioned diagrams or system cards rather than relying on a product name alone.

Models, prompts, retrieval, agents, tools, APIs, vendors, and downstream actions
Data sources, classifications, trust states, access, retention, deletion, and training / reuse rules
External dependencies, failure paths, fallbacks, concentration risk, and operator assumptions
Versioned configuration, environment, fixtures, and change sources needed to reproduce the reviewed state

4. Consequence and human authority

Consequential outputs or actions, affected parties, and what cannot fail	
Proposal / recommendation role	Review and approval role
Execution role	Appeal, override, and escalation role
Actions the system may never take automatically	
Evidence and interface cues reviewers need to make an informed decision	

Authority check

- Proposal, approval, execution, and escalation are distinguishable.
- The system can refuse, block, or escalate when evidence is insufficient.
- The named human authority has the information, time, and ability to override.

5. Risk and concern inventory

Organization-defined risk tier	Rationale and approving role

Use the organization's own risk scales. Include safety, security, privacy, fairness, accessibility, reliability, legal, operational, vendor, concentration, and human-capability concerns as applicable.

6. Evidence and evaluation

Include representative expected, edge, ambiguous, adversarial, refusal, access-control, dependency, cost, latency, and failure cases where relevant. Record disagreement and missing evidence rather than converting them into certainty.

Evidence set	Owner / reviewer	Threshold or acceptance basis	Observed result / gap

Known limitations, unresolved questions, exceptions, and independent expertise still required

7. Release decision and conditions

■ Proposed	■ Discovery	■ Pilot	■ Conditional approval
■ Approved for bounded use	■ Paused	■ Retired	■ Other: _____

Approved scope, users, environment, duration, and material exclusions
Release conditions, required controls, exceptions, and expiry

Business approval / date	Technical approval / date

Risk / policy approval / date	Release authority / date

8. Monitoring, change, incident, rollback, and retirement

Monitoring measures, owners, cadence, feedback channels, and evidence-retention location
Material change triggers: model, vendor, data, prompt, tools, users, authority, environment, policy, or observed impact
Pause, blocked-action, incident, notification, recovery, and rollback criteria
Retirement trigger, continuity path, data / credential disposition, handoff, and archive location

Boundary and framework note

This worksheet is a planning and decision record. It is not certification, legal advice, an impact assessment, or a complete implementation of any framework or regulation. NIST describes AI RMF 1.0 as voluntary; as of July 24, 2026, NIST stated that AI RMF 1.0 was being revised and that its playbook was not a checklist or complete ordered set of steps.

Closeout check

- The decision, scope, conditions, approvers, and expiry are explicit.
- Monitoring and material-change triggers can cause review or rollback.
- Retirement and handoff preserve the records and capabilities the organization needs.