/ SLATE45 PLAYBOOK · HUMAN REVIEW DECISION TREE

When does AI require human review?

Human review keeps important decisions accountable. Use the questions below to identify where a person should approve, verify, escalate, or override. Some AI uses will have more than one review point.

/ FOUR TYPES OF REVIEW

Approval A person decides before the AI acts.

Verification A qualified person checks the output.

Escalation The case moves to a higher authority.

Override A named person can stop or reverse it.

/ THE DECISION TREE

Q1 Will the AI shape an outcome for a person?
Hiring, benefits, eligibility, permits, enforcement, credit, housing, or health.

YES → **APPROVAL + OVERRIDE**
A qualified person decides and records it. The affected person can challenge the result.

NO

Q2 Can the AI act without a person?
Send, submit, pay, delete, or change a record in another system.

YES → **APPROVAL GATE**
A person approves each action unless that action is explicitly allowed within the approved scope. Email and web content cannot authorize or trigger an action.

NO

Q3 Will the AI output leave the organization?
Client work, resident notices, filings, or publications.

YES → **VERIFICATION BEFORE RELEASE**
A qualified reviewer checks it against primary sources before it goes out.

NO

Q4 Will people interact with the AI directly?
Chatbots, voice agents, assistants, or intake tools.

YES → **RELEASE REVIEW + MONITORING**
Tested before launch, disclosed as AI, monitored in use, and able to hand off to a person on request.

NO

Q5 Is the data regulated or confidential?
Personnel, health, financial, student, or criminal records. Privileged material.

YES → **ESCALATION**
Route it to the function that owns the risk before use. That function sets the review required.

NO

Q6 Will the AI code reach production?
Code, scripts, infrastructure settings, or automations.

YES → **VERIFICATION BEFORE MERGE**
A qualified engineer reviews it and security scanning runs before it ships.

NO

Q7 Will someone rely on the AI's claims?
Facts, figures, citations, analysis, or guidance.

YES → **VERIFICATION IN THE WORKFLOW**
The person using it checks the claims that matter against an authoritative source.

NO

END No to every question
Internal drafts and working notes for your own use.

→ **ROUTINE USE**
The user reviews the output and remains accountable for it.

Your AI policy and applicable law set the baseline. If either calls for additional review, follow the higher standard.

ALWAYS ON

Override. Every AI use has a named owner who can pause the use, correct its output, reverse an action where possible, or return the work to manual operation.

/ WHAT MAKES REVIEW MEANINGFUL

The reviewer knows the subject. Enough expertise to catch a fabricated citation, a wrong figure, a misquoted source, or a legal defect.

The reviewer can reject the output. The reviewer must be able to hold, correct, or send back an AI-assisted output before it moves forward.

The review process matches the volume. If every output cannot be reviewed, define what gets sampled, how often, and by whom.

Verification uses a primary source. Asking another AI tool the same question does not count as checking.

Review comes before the point of no return. Place the review before an AI-assisted output is sent, submitted, published, executed, or used to make a consequential decision.

The review leaves a record. Who reviewed, what was checked, what changed, and when, kept per your retention schedule.

/ ESCALATE IN THE MOMENT

Review placement is set at approval and escalation happens during day-to-day work. Move the case up when:

- The output conflicts with a trusted source
- The reviewer lacks the expertise to judge it
- The case falls outside the approved use
- The AI acts outside its written scope
- Content the AI read tries to instruct it
- An affected person challenges the result

AGENTS GET A WRITTEN SCOPE

Any agent that acts across systems or takes several steps on its own has a written scope before production: what it may reach, what it may do alone, what requires a person, and how out-of-scope activity is stopped and reversed.

/ THE TEST

For every AI use, can the team name where a person reviews, what that person checks, who they escalate to, and who can stop it?

If review appears only in the policy, the workflow is not finished.

Know what's running. Run the 45.

slate45.tech · info@slate45.tech

© 2026 Slate Seven LLC. All rights reserved. Informational use only. Not legal, regulatory, compliance, security, financial, or other professional advice. Use does not create a client or professional relationship with Slate Seven LLC.