AI Automation / Task Orchestration

Design where and how humans should review, approve, or provide input within automated workflows — balances automation with necessary judgment.
Difficulty: Intermediate
Model: GPT-4 / Claude / Gemini
Use Case: Approval Workflows, Human Review
Updated: May 2026
Why This Prompt Exists
Not everything should be automated. Some decisions require human judgment — but putting humans in the wrong place creates bottlenecks; putting them in no place creates risk.

You get:

- automated decisions that should have human review (compliance risk)
- human review for routine decisions (bottleneck, slow, expensive)
- humans approving after the fact (too late to prevent errors)
- no timeout on human tasks (workflow waits forever)
- no handoff protocol (automation doesn’t know human approved)

But human-in-the-loop has patterns:

- pre-approval: human must approve before automation proceeds
- post-approval: human reviews after automation (audit, not control)
- exception handling: human only for edge cases
- escalation: human when automation confidence is low
- data enrichment: human provides missing information

Without design, human-in-the-loop creates friction, not value.

This prompt designs effective human-in-the-loop workflows.

The Prompt

```
Assume the role of a workflow designer who integrates human review
into automation.

Your task is to design where and how humans interact with an automated
workflow.

Generate:

1. HUMAN DECISION POINTS

| Decision Point | Trigger | Human Role | Timeout | Escalation |
|----------------|---------|------------|---------|------------|
| [e.g., high-value order] | [amount > $10,000] | Approve/Reject | 4
hours | Second reviewer |
| [e.g., refund request] | [user requests refund] | Review reason | 24
hours | Auto-approve |

2. HUMAN INTERFACE DESIGN
   - Where humans receive requests: [Slack / Email / Dashboard /
Webhook]
   - Information provided to human: [data snapshot, context, history]
   - Action options: [Approve, Reject, Request changes, Escalate]

3. TIMEOUT BEHAVIOR
   - If human doesn't respond in [X] minutes/hours:
```

* [Auto-approve / Auto-reject / Escalate / Retry notification]

4. AUTOMATION AFTER HUMAN DECISION
   - If approved: [what happens next]
   - If rejected: [what happens next, e.g., notify user, stop workflow]
   - If changes requested: [how to handle, e.g., return to user for edits]

5. AUDIT & LOGGING
   - What to log: [who approved, when, decision reason, data snapshot]
   - Retention period: [how long logs are kept]

6. OPTIMIZATION OPPORTUNITIES
   - Which human decisions could be automated with better rules?
   - Which automated decisions need more human review?

INPUTS:

Workflow description:
[E.G., "Customer refund process"]

Decision types requiring human judgment:
[E.G., "High-value refunds (>$500), unusual refund patterns, first-time customers"]

Human availability:
[E.G., "Support team 9-5 weekdays, on-call for emergencies"]

```
Compliance requirements:
[E.G., "SOX: all refunds >$1,000 require manager approval"]

RULES:
- Pre-approval for irreversible actions (payments, deletions, data exports)
- Post-approval for reversible or low-risk actions (audit trail sufficient)
- Always set a timeout on human tasks (prevent workflow stalls)
- Provide context, not just raw data (help humans decide quickly)
- Log every human decision for audit and training
- Review human decision patterns quarterly — automate what becomes routine
```

How To Use It

- Pre-approval for irreversible actions (payments, deletions, data exports).
- Post-approval for reversible or low-risk actions (audit trail sufficient).
- Always set a timeout on human tasks — prevent workflow stalls.
- Provide context, not just raw data — help humans decide quickly.
- Log every human decision for audit and training.
- Review human decision patterns quarterly — automate what becomes routine.

Example Input

Workflow description:
“Customer refund process — automatic for small amounts, manual review for larger ones”

Decision types requiring human judgment:
“Refunds over $500, multiple refunds from same customer, first-time customers”

Human availability:
“Customer support team, 8am-8pm ET, no weekend on-call”

Compliance requirements:
“All refunds >$1,000 require manager approval (SOX compliance)”

Why It Works
Most automation either avoids human involvement entirely (risky) or puts humans everywhere (inefficient). Neither is optimal.

This framework improves outcomes by forcing:

- human decision point identification (where are humans needed?)
- timeout specification (what if human doesn’t respond?)
- automation after human decision (what happens next?)
- audit design (tracking for compliance)
- optimization opportunities (which human tasks can be automated?)

Failure modes this prevents:

- Human bottleneck — workflow waits days for approval (add timeout + escalation)
- Human review for routine decisions (automate those)
- No audit trail — compliance failure (log everything)
- Stalled workflow — human on vacation (add fallback approver)

This improves on: Fully automated workflows that make mistakes, and fully manual workflows that are slow.

Related to: TO-04 (Recovery) for human escalation on failures; CRM-01 for approval workflows.

Build Better AI Systems

Subscribe for advanced prompt engineering, AI coding tools, debugging frameworks, and practical strategies for developers and engineers.

Carefully engineered prompts for people doing real work.

Share this:

- Share on Facebook (Opens in new window) Facebook
- Share on X (Opens in new window) X

See also Parallel Execution Designer