

PRACTICAL CONTROL GUIDE

Human-in-the-loop

A practical guide to placing one meaningful human checkpoint where automation carries the greatest commercial consequence.

A checkpoint should protect the outcome, not slow every step.

Let routine work run. Ask a person to review the point where money, commitments, customers or irreversible change enter.

01

Locate consequence.

Place control where being wrong would cost the business most.

02

Give the review evidence.

A person needs context, options and a clear decision to make.

03

Record the catches.

Visible value prevents useful controls being removed as overhead.

One well-placed checkpoint can protect an otherwise efficient flow.

[Use this guide →](#)

MISSION & SCOPE

Human review belongs at the point of consequence.

A checkpoint is not a person watching every automated step. It is a designed pause where evidence reaches a named owner before a high-impact action continues.

THIS GUIDE IS USEFUL IF

- Automation touches money or commitments.
- Exceptions are difficult to reverse.
- Existing reviews are routinely rubber-stamped.

THIS GUIDE WILL NOT

- Add approval to every low-risk action.
- Treat human review as a substitute for good automation.
- Remove accountability from the final owner.

How the checkpoint works

01

Detect

Identify the condition or threshold that requires review.

02

Decide

Give a person the evidence and options needed to act.

03

Record

Capture the decision, reason and next step for later review.

**A smoke detector does not put out the fire.
It makes sure the right person finds out in time.**

The checkpoint is valuable because it concentrates attention. It should interrupt only the cases where judgement or consequence justify the pause.

CHECKPOINT PLACEMENT

Six questions that find the highest-value human review.

Begin with consequence, then design the evidence and workload around it.

Protect the irreversible step.

A checkpoint placed after the commitment may create evidence, but it cannot prevent the outcome.

01

Consequence

Where could error affect money, safety, customers or commitments?

02

Reversibility

Which action is difficult or expensive to undo?

03

Evidence

What must the reviewer see to make a real decision?

04

Ownership

Which person or role is authorised to approve or stop?

05

Volume

How many cases can be reviewed without fatigue?

06

Learning

What record will show what the checkpoint caught?

RISK CHECKPOINT

Let routine work run. Pause before consequence.

The control belongs before the highest-impact action, not wherever it is easiest to build.

RUN UNATTENDED	REQUIRE REVIEW
01 Prepare and classify Draft, sort, calculate or assemble evidence inside known rules.	01 Commit or communicate Pause before money, promises, deletion or sensitive customer action.
02 Route normal cases Continue when inputs and conditions remain inside the safe path.	02 Handle exceptions Give a person context, options and authority to decide.
EFFICIENCY Routine cases avoid unnecessary interruption.	CONTROL High-consequence cases receive deliberate attention.

Need more flexibility? [Compare agents and workflows →](#)

COMMERCIAL EFFECT

A checkpoint creates value when it changes a decision in time.

Approval activity alone is not control. The review must occur before consequence, reach an authorised person and contain enough evidence to act.

Risk-based trigger

Useful attention

Only consequential or unusual cases interrupt the flow.

Easy-to-build trigger

Misplaced effort

Low-risk events create noise while important cases may pass untouched.

Decision evidence

Informed review

The reviewer sees context, source, options and consequence.

Bare notification

Rubber stamp

A generic alert asks for approval without supporting judgement.

Outcome recorded

Visible learning

Caught issues and decisions can improve rules and justify the checkpoint.

Approval disappears

Invisible control

Nobody can show what changed because the review occurred.

Ask what decision the human is expected to make.

The value appears through prevented errors and better exceptions. Record the cases and changes, not merely the number of approvals.

FIVE SIGNS THE HUMAN CHECKPOINT IS ONLY DECORATIVE

Problems to catch before review becomes a ritual.

A checkpoint must change what happens, not merely add a click.

Critical

Sign 1

No checkpoint protects customer money or commitments

The automation can commit the business before a person sees the case.

WHAT IT COSTS

High-consequence actions continue without deliberate human authority.

Major

Sign 2

The checkpoint exists but nobody monitors it

A queue or alert was created without assigning capacity and ownership.

WHAT IT COSTS

Cases wait, time out or proceed while responsibility remains ambiguous.

Major

Sign 3

Review sits at the easiest point

Implementation convenience replaced a consequence-based design.

WHAT IT COSTS

The control misses the action that is hardest or most costly to reverse.

Moderate

Sign 4

Too many low-value checks create fatigue

Review volume overwhelms attention and encourages rubber-stamping.

WHAT IT COSTS

Important cases receive the same quick click as routine ones.

Moderate

Sign 5

Nothing records what the checkpoint catches

Approvals are counted, but prevented errors and rule changes disappear.

WHAT IT COSTS

The control looks like overhead and cannot improve the automation.

A checkpoint without a real decision is only delay.

Give the reviewer evidence, authority and a meaningful reason to pause.

CORRECTION AND SELF-ASSESSMENT

Five questions before adding another approval step.

Common myth: human review defeats automation. Better framing: one targeted checkpoint protects the value of every unattended step around it.

01

Which consequence are we trying to prevent?

Name the money, commitment, customer impact or irreversible action.

02

Does the review happen before that consequence?

A late checkpoint can document failure but cannot prevent it.

03

What evidence reaches the reviewer?

Provide source, context, options and the effect of each choice.

04

Can the owner realistically monitor the volume?

Route only cases that justify attention and assign capacity.

05

What will be recorded and improved?

Capture decisions, caught issues and rule changes over time.

WORKING CHECKPOINT STANDARD

A reusable control for one high-consequence step.

Define the trigger, evidence, owner, decision and record together. A queue without these elements is not a functioning checkpoint.

01 Name the consequence

State what could happen if the automation is wrong.

02 Set the trigger

Route unusual or high-impact cases before action.

03 Supply evidence

Give source, context, options and consequence.

04 Authorise the owner

Name who may approve, reject, amend or escalate.

05 Protect timing

Set response expectations and a safe timeout state.

06 Record learning

Capture what changed and improve the rule.

INSTEAD OF

“Someone checks it.”

No trigger, evidence, authority or record.

Reusable checkpoint record

CONSEQUENCE:

TRIGGER:

EVIDENCE:

REVIEW OWNER:

DECISIONS ALLOWED:

TIMEOUT STATE:

OUTCOME RECORD:

RULE CHANGE:

DISCUSS THE CHECKPOINT

Protect the highest-consequence step.

Map one trigger, owner, decision and learning loop.

DUE DILIGENCE

Place attention where it can still change the outcome.

Three rules and two next steps for a useful human checkpoint.

01

Trigger by risk

Route the cases where being wrong matters commercially.

02

Provide evidence

Give the reviewer enough context to approve, change or stop.

03

Record change

Use caught cases to improve rules, training and ownership.

One meaningful checkpoint is stronger than many decorative approvals.

If a provider cannot show the trigger, evidence, owner and learning loop plainly, keep digging.

CONTINUE LEARNING

Agents vs Workflows

Choose the level of flexibility the task actually needs.

[Read next →](#)

AI Agents & Automation

Continue through the rest of this collection.

[Explore collection →](#)

Let routine work run. Pause before consequence.

A useful human checkpoint has a risk-based trigger, decision evidence, an authorised owner and a record that improves the system.

PRIMARY NEXT STEP

Agents vs Workflows

Choose between a fixed path and flexible action for the task.

[Continue →](#)

AI Agents & Automation

See every guide in this collection.

[Explore collection →](#)

All guides

Browse every topic and problem.

[View library →](#)

Need help placing the checkpoint?

Bring one automated process. Identify the consequential action, trigger, evidence, reviewer and safe timeout state.

[Discuss control →](#)