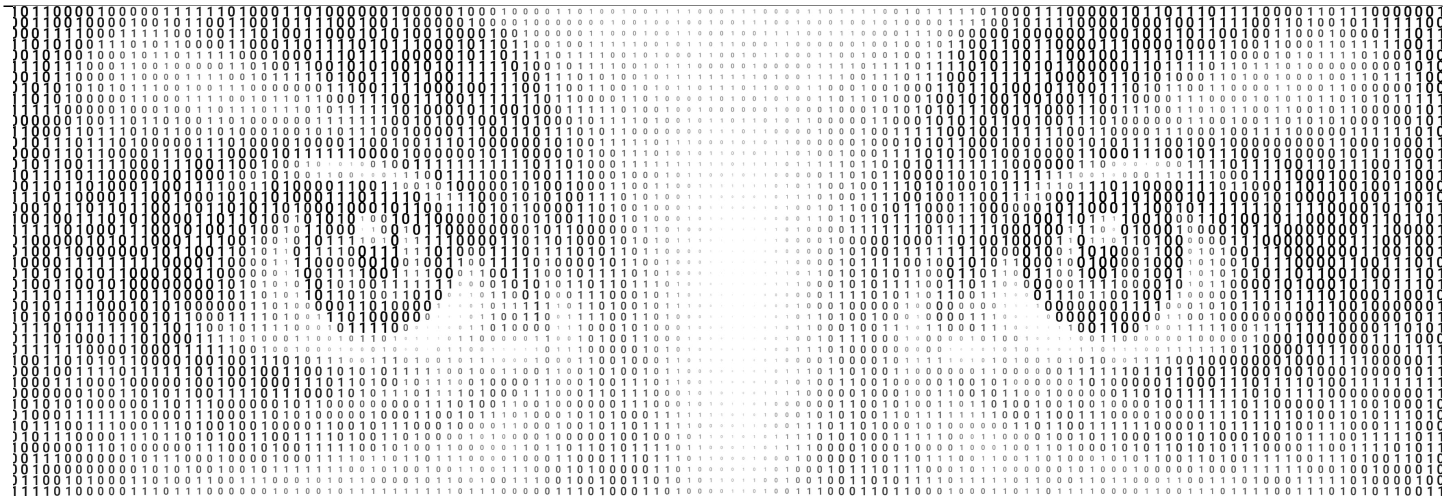


Does Your AI Actually Need a Human in the Loop?

Corporate, Securities, and Governance

Technology, Privacy, and eCommerce



Banner artwork by Kundra / *Shutterstock.com*

“Keep a human in the loop” has become standard advice in AI governance. It appears in policies, risk assessments and discussions with vendors, often without much explanation of what the human is supposed to do. For a general counsel reviewing a proposed use, the first point to establish is that AI does not generally require human review. There is no rule in either the European Union or the United States requiring someone to approve an output merely because AI produced it.

The question turns on what the system is actually authorized to do. Providing information is different from determining an outcome. Once the system begins making decisions about people, the nature of the decision, and the law governing it become relevant.

Even then, “human in the loop” does not necessarily mean human approval. What the human is required to do, and when that person becomes involved, depends on the rule that applies. Some regulated uses do not require human review at all.

When does the requirement arise?

What is the AI doing?	Human involvement?	What does the human do?	Putting it into practice
Providing ordinary assistance, such as drafting or summarizing	Usually not as a matter of AI-specific law	No legally required individual review merely because AI produced the output	Do not add human approval merely because AI is used
Recommending an important decision about a person, while a person genuinely decides	The human role may determine whether automated-decision rules apply	The person exercises genuine judgment rather than simply accepting the AI recommendation	Identify who makes the decision and confirm that person can disregard the AI
Making a legally or similarly significant decision about a person entirely automatically in	GDPR restrictions may apply	Where such automation is permitted on certain grounds, safeguards include human	Build the required intervention and contest process into the workflow

the EU		intervention and an opportunity to contest the decision	
Operating as an EU AI Act high-risk system	Human oversight is required	The person oversees the system during use and can intervene where appropriate	Assign responsibility and give the person enough authority to perform the oversight role
Producing an adverse consequential decision covered by Colorado law	Human review available on request	From Jan 1, 2027, meaningful human review and reconsideration must be available to the extent commercially reasonable	Create a route for requests and identify who will conduct the review
Making an automated US credit decision	Not merely because the decision is automated	Regulation B requires specific reasons for adverse action rather than human approval of every decision	Make sure the required reasons can actually be produced

The table addresses human involvement only. Other legal obligations may apply even when nobody is required to review individual outputs.

When the human makes the decision

The General Data Protection Regulation (GDPR) restricts certain decisions about individuals that are made entirely through automated processing and produce legal or similarly significant effects. The rule becomes relevant when the system itself determines the outcome rather than assisting the person who does.

In recruitment, an AI-generated assessment can form part of a decision that is still made by the recruiter. If the company relies on the recruiter’s involvement to treat the decision as human rather than solely automated, the recruiter must exercise real judgment and be able to depart from the AI’s recommendation. Simply adding an approval step to an otherwise automated process may not be enough.

Automatic rejection puts the system in a different position under the GDPR. The company would need to consider whether the decision has the kind of significant effect covered by the regulation and whether fully automated decision-making is permitted in those circumstances.

For legal, the practical consequence is fairly simple. If the proposed control is that a person makes the decision, the approval for the use case should identify that person or role and make clear that the AI output is not binding. The workflow should match what the company says happens.



When the human oversees the system

The EU AI Act separately requires human oversight for systems that qualify as high-risk. Certain uses in employment and credit fall within this category, although an AI tool is not automatically high-risk simply because it is used in one of those fields. For the listed high-risk uses most relevant to ordinary corporate deployments, the requirements apply from 2 December 2027.

Human oversight does not generally require approval of each output. The person responsible needs enough understanding of the system to recognize when intervention is necessary and enough authority to act when it is.

> The person responsible needs enough understanding of the system to recognize when intervention is necessary and enough authority to act when it is.

Legal should therefore avoid writing “human review required” into an approval without specifying what that means. If the requirement is oversight, the business needs to know who is responsible and whether the product gives that person the ability to perform the role.

This is also worth checking in the vendor contract. A promise that a product supports human oversight is of limited value if the customer does not have the information or controls needed to exercise it.

When the human becomes involved afterward

The GDPR also provides a right to obtain human intervention where certain solely automated

significant decisions are permitted. Certain solely automated significant decisions are permitted in defined circumstances. Where the relevant safeguards apply, the individual must be able to obtain human intervention and contest the decision. Human approval before the original decision takes effect is not necessarily required.

Colorado's Automated Decision-Making Technology Act takes this approach for certain automated outcomes. The law, as revised in 2026 and effective January 1, 2027, covers automated technology that materially influences consequential decisions. Following a covered adverse outcome, the affected person can request meaningful human review and reconsideration, to the extent commercially reasonable. The reviewer must have authority to change the decision and must consider the case rather than simply default to the automated result.

A company using a system covered by such a requirement needs somewhere for the request to go. Legal should identify the reviewer and make sure the process allows that person to reconsider the outcome. There is little benefit in requiring approval of every decision if the law instead requires a meaningful route to human review afterward.

When no human review is required

Under the Equal Credit Opportunity Act (ECOA) and Regulation B, a creditor taking adverse action must provide the applicant with the specific principal reasons for the decision. Those reasons must accurately describe the factors actually considered or scored, including where the creditor uses an automated credit-scoring system. ECOA does not require a credit officer to approve an automated decision before it is made.

Legal should treat the obligation it actually finds as the starting point. If the requirement is to provide specific reasons for an adverse decision, the control should ensure that those reasons can be provided. Adding manual approval does not answer that requirement.

> Legal should treat the obligation it actually finds as the starting point.

A company can still choose human review for a particular use. If it does, the approval should make clear that this is an internal control rather than a legal requirement. That becomes important if the business later wants to change the process.

What should legal do with this?

When a new AI use reaches legal, counsel needs a reasonably concrete description of what the system will do and what authority it will have. If it affects decisions about people, identify the law governing that use before deciding that a human reviewer is necessary.

Where the law does require human involvement, "human in the loop" is too vague for an approval record. The record should state what the person is expected to do and who is responsible. It should also be possible to tell from the actual workflow that the role exists.

The same question belongs in vendor diligence where the product will be used for regulated decisions. Legal does not need a generic representation that the product allows human oversight. It needs to know whether the company can perform whatever human role its own use of the product

requires.

The resulting approval does not need to be complicated. Depending on the use, the relevant entry might say:

Human involvement: Required. Recruiter makes the final decision and may disregard the AI recommendation. No automatic rejection permitted.

Or:

Human involvement: Required as operational oversight. Designated personnel can intervene in the system's operation. Individual outputs do not require approval.

Or simply:

Human involvement: Not legally required for individual decisions. Adverse-action explanation requirements apply.

That is more useful than making "human in the loop" a standard condition of AI approval. Legal can identify the requirement that actually applies and design the human role around it, without adding manual review where the law does not require it.

Join ACC for more AI insights!

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

Matías Cañibe



Associate General Counsel and Data Protection Officer based

Luxembourg

Matias Cañibe is Associate General Counsel, Assistant Corporate Secretary and Data Protection Officer at a NASDAQ-listed multinational company. He advises Boards, executive management and business teams on corporate governance, privacy, cybersecurity, artificial intelligence governance and cross-border legal matters. He holds law degrees from two civil-law jurisdictions and an MBA. He writes on the legal and governance implications of emerging technologies.