

The AI Authority Ladder

A practical framework for deciding what AI should be allowed to do

ASSIST → RECOMMEND → DECIDE → ACT

AI capability and AI authority are not the same thing.

As AI moves from helping people work to recommending decisions, making decisions, and taking action, organizations need a clearer way to decide how much authority AI should actually have.

Jennifer Scarlato


The Visual Enterprise Architect

September 2026

AI is gaining authority, not just capability.

AI can help a person work, recommend what should happen, make a decision, or take action. Those roles are not interchangeable, and they should not be governed as though they are.

The AI Authority Ladder gives enterprise architects, business leaders, technology teams, risk teams, and product owners a simple language for deciding what AI should be allowed to do - and what needs to be true before it is allowed to do more.



The goal isn't to climb the Ladder. The goal is to choose the right level of authority for the work.

AI is not one thing.

People often describe AI as if it has a single setting: on or off, trusted or not, safe or risky, pilot or production. In reality, AI interacts with work along a spectrum of roles—and each role requires different expectations and controls.

Sometimes AI is simply helping someone work faster. It summarizes a case, drafts a response, finds a policy, or pulls key points from a document. A person is still clearly in charge.

Sometimes AI recommends something. It suggests a next step, a risk rating, a routing path, or a response. A person still decides, but AI is beginning to influence that decision.

Sometimes AI makes the decision. It classifies a request, prioritizes work, determines eligibility, or approves something within defined boundaries.

And sometimes AI acts. It creates a record, updates a system, sends a communication, triggers a workflow, or starts a downstream process.

Those are very different relationships between AI and the enterprise. Yet we often govern all of them under one broad heading: AI.

The question isn't only: What can AI do?

We also need to ask: What should AI be allowed to do?

Capability and authority are not the same thing.

Organizations spend a lot of time evaluating AI capability. Can the technology summarize this information? Determine the next best action? Make this decision? Execute this transaction?

Those are necessary questions, but technical capability shouldn't automatically become organizational permission.

An AI system may technically be capable of issuing a refund, modifying a customer record, approving an exception, changing infrastructure, communicating with another system, or initiating a purchase. That doesn't mean it should have the authority to do those things.

The architecture conversation should be asking: What authority are we willing to delegate to AI?

That question brings ownership, boundaries, human judgment, reversibility, accountability, and control into the conversation. Those are questions enterprise architects should be helping organizations answer.

Introducing The AI Authority Ladder


Over the past several years, I've watched the AI conversation shift quickly. We started by asking what AI could do. Now organizations are beginning to face a harder question: what should we allow it to do?

That distinction led me to create the AI Authority Ladder—a simple framework for describing the authority an organization delegates to AI.

Assist. Recommend. Decide. Act.

The framework is intentionally simple. The goal isn't to create another complicated AI taxonomy. It's to give business, technology, architecture, risk, and governance teams a common language for making better decisions about AI authority.

READ THE LADDER FROM THE BOTTOM UP. Each rung changes the role AI plays in the work. Higher authority does not automatically mean better AI.



01 ASSIST

AI helps a person do the work.

At the Assist level, AI helps someone work faster, more consistently, or with better context. It might summarize information, draft content, search approved knowledge, translate something, extract information, organize notes, or explain a complicated document.

The important distinction is that the person remains responsible for the judgment and the action.

Assist can create enormous value without giving AI much authority. It also gives organizations room to learn where data is strong, where knowledge is weak, where users trust AI, and where AI actually improves the work.

AT THIS LEVEL, ASK

Is the person clearly responsible for reviewing the output?

Is AI using trusted sources?

Can the user easily verify what AI produced?

Is sensitive information protected?

Authority: Help me.

02 RECOMMEND

AI suggests what should happen.

Recommend is where AI begins influencing a choice. It might recommend a case category, next best action, risk rating, routing path, remediation option, investment priority, or response.

The human still decides, but that distinction shouldn't make us underestimate the authority involved. People may accept an AI recommendation because it is fast, confident, or simply easier.

The human can technically remain in control while AI quietly becomes the dominant decision-maker. Recommendation therefore needs transparency, meaningful override, and measurement of whether the recommendation is actually improving outcomes.

AT THIS LEVEL, ASK

What is the recommendation based on?

Can the user see enough context to evaluate it?

Can it be rejected or overridden?

Are people evaluating the recommendation - or simply approving it?

Authority: Advise me.

03 DECIDE

AI chooses what should happen.

Decide represents a meaningful shift in authority. AI may classify incoming work, approve a low-risk request, determine eligibility, prioritize a queue, identify an exception, assign a risk level, or select a workflow path.

The system isn't suggesting the decision anymore. It's making it.

That isn't automatically unsafe, particularly when the decision is narrow, well understood, reversible, and supported by reliable data. But the decision boundary needs to be explicit, and the organization needs to know what happens when the data is incomplete, confidence is low, or the outcome is wrong.

AT THIS LEVEL, ASK

Is the decision boundary explicit?

Is the underlying data trusted enough?

What exceptions require human judgment?

Can we reconstruct why the decision was made?

Who owns the outcome?

Authority: Decide for me.

04 ACT

AI changes something.

Act is the highest level of authority on the core Ladder. AI is no longer only saying what should happen; it is making something happen.

It may create or update a record, send a communication, trigger a workflow, initiate fulfillment, close a case, modify a configuration, invoke another service, or start a downstream process.

Act contains degrees of autonomy. AI may act only after explicit approval, operate independently within predefined bounds, or orchestrate an outcome by determining steps and invoking tools or other agents. The core Ladder stays simple, while the governance inside Act can become much more sophisticated.

AT THIS LEVEL, ASK

What exactly can AI change?

Which systems and credentials can it use?

Which actions require approval?

What is logged and monitored?
Can the action be reversed?
When must AI stop and escalate?
Authority: Act for me.

The Ladder is not a maturity model.

The goal isn't to move every AI use case up the Ladder. More authority doesn't automatically mean a better or more mature solution.

Some use cases may always belong at Assist, and that may be exactly where they create the most value. Others may be ready to Recommend but not Decide. And for the right kind of work—narrow, low-risk, well understood, and easy to reverse—Act may make perfect sense.

There is no prize for reaching the top. The point is to choose the level of authority that makes sense for the work.

Authority changes the architecture.

As authority increases, several parts of the architecture conversation change with it.

	ASSIST	RECOMMEND	DECIDE	ACT
Primary concern	Output	Influence	Accountability	Impact
Data bar	Reviewable	Reliable enough to guide	Trusted enough to decide	Trusted enough to trigger
Human role	Review / edit	Choose / override	Own exceptions	Approve / monitor / intervene
Control emphasis	Sources / feedback	Transparency / override	Thresholds / audit	Permissions / logs / rollback

Risk changes as authority increases.

Not every Assist use case is low-risk, and not every Act use case is dangerous. Context matters. But the nature of the risk changes.

At Assist, the concern often centers on whether a person can verify the output. At Recommend, influence becomes important. At Decide, accountability becomes more significant. At Act, impact becomes central: what changes because AI took action, and what else might that action trigger?

| The controls should match the authority.

Data readiness changes too.

AI authority and data readiness belong in the same conversation. A use case may be ready for AI without being ready for the authority people want to give it.

AI may be ready to summarize supplier documentation but not approve supplier risk. It may be ready to recommend case routing but not close customer cases automatically. It may be ready to explain policy but not determine employee eligibility.

That doesn't make the use case unsuccessful. It means the authority needs to match the foundation.

| A use case can be ready for AI without being ready for the authority people want to give it.

Explainability matters more as authority increases.

People don't need to understand every technical detail behind an AI model, but they need enough information to trust, challenge, and govern what it is doing.

At Assist, that may mean seeing the source. At Recommend, it means understanding enough context to evaluate the suggestion. At Decide, the organization may need to reconstruct the information or criteria behind the choice. At Act, explainability and auditability come together: What happened? Why? What supported it? What else did it trigger?

As authority increases, "the AI did it" becomes a less acceptable answer.

Authority has to be designed into the work

Organizations frequently respond to AI risk by saying, "We'll keep a human in the loop." It sounds reassuring, but it doesn't tell us much about what that person is actually expected to do.

A human can technically be in the loop while providing almost no meaningful oversight. Someone clicking Approve hundreds of times a day without enough information or time to evaluate the recommendation isn't really exercising judgment. They've become a rubber stamp.

The more useful question is: Where does human judgment actually matter?

At Assist, that may mean reviewing and editing AI-generated work. At Recommend, the person needs to be able to evaluate the recommendation, make the final choice, and override it when necessary. Once AI begins to Decide, human involvement may shift toward exceptions or higher-impact decisions. At Act, people may approve sensitive actions, monitor outcomes, or intervene when AI reaches a boundary.

That's why "human in the loop" isn't really a control design. We have to design the loop.

The same is true for governance. AI controls shouldn't live only in a policy document or governance meeting; they need to show up in the way the work actually happens. Source visibility, thresholds, approvals, exceptions, permissions, logging, monitoring, and the ability to reverse an action all become part of the workflow as AI gains authority.

This also means AI governance can't belong to one team. Architecture, product, security, risk, data, and platform teams each see a different part of the picture. The work is bringing those perspectives together and turning them into practical boundaries for how AI operates.

AI governance cannot just be a meeting. It has to become design.

Start with the problem, not the AI

The Authority Ladder helps determine the right level of authority when AI is the right solution. It doesn't mean AI is always the right solution.

Sometimes a clearer form solves the problem. Sometimes it's better integration, a business rule, cleaner data, a simpler policy, or removing an unnecessary step altogether. If required information is consistently missing, validation may be more useful than AI. If approvals are unclear, the real problem may be decision rights. If people are entering the same information into three different systems, an integration is probably going to create more value than adding AI on top of the problem.

That's an important discipline for enterprise architects. Before deciding how much authority AI should have, we should first ask whether AI improves the work more than a simpler solution.

Consider a customer service organization that wants to "use AI to improve case resolution." At first, that sounds like a single AI use case. Through the Ladder, it becomes several very different possibilities.

AI could Assist by summarizing case history or drafting a response. It could Recommend the next best action or escalation path. It could Decide whether a case meets the criteria for a standard resolution. Or it could Act by updating the case, communicating with the customer, initiating fulfillment, or closing the work.

The business outcome hasn't changed. The technology may even be the same. What changes is the authority we've given it—and that creates four very different architecture, risk, and governance conversations.

That is where the Ladder becomes useful.

More authority should be earned

There will be cases where moving up the Ladder makes sense. But that decision should be based on what we've learned, not simply on what the technology is capable of doing next.

If AI is currently assisting, we should know whether the output is useful and reliable before allowing it to influence decisions. If it is recommending, we should understand whether those

recommendations are improving outcomes and where people are overriding them. Before allowing AI to decide or act, we should have confidence in the data, understand the exceptions, know that the controls are working, and be comfortable with what happens when the AI gets something wrong.

If those foundations aren't there yet, the answer isn't necessarily to abandon the use case. It may simply mean this is the right place on the Ladder for now.

There is nothing wrong with staying at Assist or Recommend when that is where the use case creates the most value with an acceptable level of risk. Moving up should happen because the evidence supports giving AI more authority—not because Act is somehow the destination.

Don't ask whether AI can move up the Ladder. Ask what would need to be true before it should.

Take It With You: The AI Authority Assessment

A framework is only useful if you can put it to work.

The AI Authority Assessment is designed to take the ideas in this paper and bring them into a real conversation. Use it when evaluating a new AI use case, moving a pilot into production, introducing AI into an existing workflow, or considering whether AI is ready for more authority.

This isn't meant to become another governance form that gets completed, filed away, and forgotten. Bring the worksheet into the room with the people who understand the work, the data, the technology, and the risk. Work through it together.


The questions are intentionally simple. What are we asking AI to do? What should remain human-owned? What could go wrong? What controls do we need? How will we know it's working? And what would need to be true before we gave AI more authority?

Pay particular attention to the questions that are difficult to answer. The blank spaces, disagreements, and “we don’t know yet” answers are often where the most important architecture

Good questions lead to better decisions.

AI AUTHORITY ASSESSMENT WORKSHEET

DEFINE. ASSESS. DECIDE. PUT IT TO WORK.



1 DEFINE THE USE CASE
What are we trying to achieve?

 Use case / name: _____
Business outcome: _____
What is AI being asked to do?

2 CHOOSE THE AUTHORITY LEVEL
What level fits this use case and why?

☐ ASSIST ☐ RECOMMEND ☐ DECIDE ☐ ACT

Why this level?

3 DEFINE THE BOUNDARIES
What stays human? What does AI need? What guardrails apply?

 **Human-owned**
What must remain with people?

 **Data / knowledge**
What data does AI need?

 **Controls**
What guardrails are required?

4 THINK THROUGH THE RISK
What could go wrong and what is the impact?

 **Risks**
What could go wrong?

 **Impact**
What's the impact if it does?

 **Reversibility**
Can the action be reversed?

5 WHAT WOULD EARN MORE AUTHORITY?
What would need to be true to move up a level?

 _____

6 DEFINE SUCCESS
How will we know this level of authority is working?

 _____

7 BEFORE YOU MOVE UP THE LADDER ...

If you can't check these boxes, you're probably not ready for more authority. →

☐ The current level is working.

☐ The data can support it.

☐ We understand what can go wrong.

☐ Human ownership is clear.

☐ Controls are in place.

☐ We can measure the outcome.

It's not about what AI can do. It's about what we should let it do.

THE VISUAL ENTERPRISE ARCHITECT

turn complexity into clarity.

JENNIFER SCARLATO

conversations begin.

Closing Thought

AI is going to keep getting more capable. The harder question for organizations will be deciding how much of that capability they are actually willing to turn into authority. That decision shouldn't be driven by what the technology can do or by a desire to move every use case toward greater autonomy. It should be based on the work, the risk, the data, the controls, and where human judgment still matters.

The AI Authority Ladder gives us a simple way to have that conversation. **Assist. Recommend. Decide. Act.** There is no prize for reaching the top. The value is in choosing the level that makes sense today, being clear about why, and understanding what would need to change before giving AI more authority tomorrow.

For enterprise architects, this is an opportunity to do what we do best: make the relationships visible, make the tradeoffs easier to understand, and help the organization make a better decision.

AI authority should be intentional. Not because AI can do more, but because we have decided when it should.