EXECUTIVE BRIEFING

The AI Control Gap

When AI Moves From Answering Questions to Taking Action

EXECUTIVE THESIS

AI governance is not about slowing adoption. It is about creating the visibility, boundaries, and proof required to trust AI with increasingly important business processes.

Prepared for mid-market executive teams

CEOs • CFOs • COOs • CIOs • CISOs • Finance, Risk and Compliance Leaders

entech

Strategy-led technology, cybersecurity, compliance and AI advisory

Governance That Makes Scale Possible

AI did not enter most organizations through a single decision. It arrived gradually, through employee tools, embedded SaaS capabilities, automation platforms, AI assistants, and increasingly, systems that can coordinate work across multiple applications.

That gradual adoption creates a new executive challenge. The question is no longer limited to whether a tool is approved. Leadership must understand how AI participates in business workflows, what information it can access, what authority it has, and how the organization proves that an AI-enabled outcome can be trusted.

> **ENTECH POINT OF VIEW**
>
> **The unit of AI governance is the workflow, not the AI tool.**

This briefing presents a practical approach for mid-market leaders. The focus is not on creating a large new bureaucracy. It is on building enough visibility, control, and validation to move from experimentation to governed execution.

Contents

The Control Problem Is Changing

An employee started using an AI assistant to summarize a document. A finance leader used AI to draft commentary for a management report. A sales team connected AI to the CRM. Operations began automating routine workflow decisions. IT introduced an assistant capable of retrieving information from internal systems. A business team began experimenting with an AI agent that could complete several steps of a process.

Each individual step may have appeared manageable. Together, they represent a much larger shift: AI is moving from helping people create content to becoming part of how work gets performed.

PROGRESSION

Governance requirements increase as AI takes on more of the workflow:

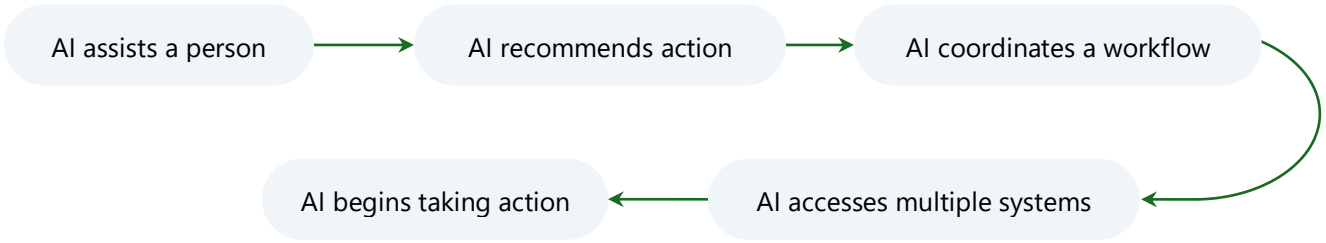

A Familiar Finance Example

Consider a straightforward financial planning and analysis scenario. A finance manager asks an AI-enabled system to prepare a variance analysis for the latest quarter. The system retrieves actual financial results, obtains budget data, compares performance, identifies material variances, drafts explanations, generates charts, and prepares an executive summary.

The finished report may look excellent. But that alone is not enough.

- Where did the data come from, and did the AI have permission to access all of it?
- Did it use the correct business units, entities, accounts, and reporting periods?
- Were any transformations applied to the data?
- Were required reviews completed?
- Where was the final report saved or sent?
- Could someone reconstruct the process six months later during an audit, investigation, or management review?

THE NEW EXECUTIVE QUESTION

How do we know what the AI did, what information it accessed, what controls were enforced, and whether the result can be trusted?

Organizations need a governance model designed for this new reality. Not governance that blocks experimentation. Governance that makes responsible scale possible.

From AI Tools to AI Actors

Traditional software generally follows predefined rules. A user selects an option. A system performs a known function. A transaction follows a defined path.

AI-enabled systems can operate differently. Depending on how they are designed and connected, they may interpret a request, develop a multistep plan, decide which tools to use, retrieve information from multiple sources, perform analysis, revise their approach, route exceptions, and execute permitted actions.

This creates an important difference for business leaders. The risk is not limited to whether the AI produces an incorrect sentence. The system may make several decisions before the final output is ever seen.

- It may interpret the scope of the request incorrectly.
- It may retrieve more information than the user needs, or pass data between systems.
- It may select a tool that introduces additional risk.
- It may take an action that traditionally required an employee to obtain approval.
- It may produce a correct result using a path the organization would never have approved.

WORKFLOW VIEW

Governance must follow the full path from business request to business outcome:

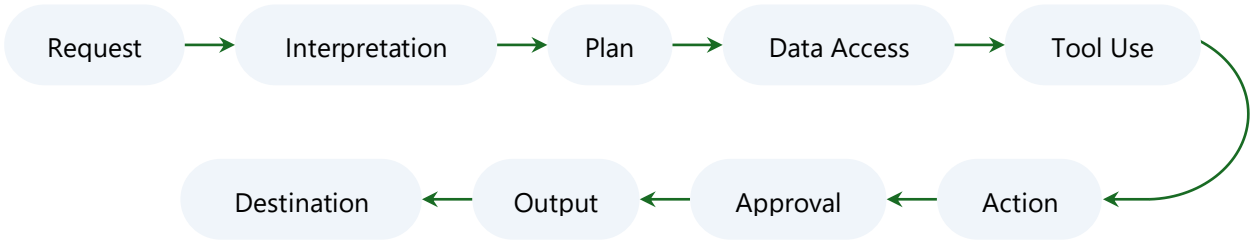

Every stage creates a governance question: Was the request appropriate? Did the AI interpret it correctly and stay within scope? Did it use approved data, and were permissions respected? Was approval required, and did approval actually occur? Was the result sent to an appropriate destination? Can the entire process be reconstructed?

EXECUTIVE IMPLICATION

An organization can have an approved AI platform and still have an ungoverned AI process.

Governance must follow the work from beginning to end. Three requirements become especially important:

Visibility & Oversight	Controlled Authority	Continuous Validation
See and reconstruct what happened.	Define what AI can access and do.	Test the outcome and the process.

Visibility and Oversight

The first requirement for governed AI is simple to describe: the organization needs to be able to see what happened. As AI becomes involved in business processes, leaders need more than a final answer. They need enough evidence to understand how the process operated: what request initiated it and who initiated it, how the request was interpreted, what steps the system performed, what information was retrieved, which applications and permissions were involved, what exceptions occurred, which approvals were completed, and where the result was sent or stored.

This is not only an audit issue. It is an operating issue. When a weekly operating report shows an unexpected decline in gross margin, the analysis may be wrong because the AI used a preliminary data set, a single division rather than the full company, inconsistent reporting periods, or an obsolete spreadsheet. The final number is only the visible symptom. The real management question is: *can we determine exactly how the number was created?*

> **WHY IT MATTERS**
>
> **A company cannot improve a process it cannot see.**

Practical Example: FP&A Reporting

Assume the CFO asks: Compare actual results against budget for the Southeast region, identify significant variances, and prepare a short executive explanation.

A well-governed process should preserve enough information to show how the system interpreted “Southeast region,” which entities and periods were included, which actual and budget sources were used, how materiality was defined, which calculations were performed, whether unusual results were escalated, who reviewed the analysis, and where the report was distributed.

That evidence allows finance to challenge, improve, and defend the process.

> **PRINCIPLE**
>
> **If the organization cannot explain how an AI-enabled process reached an outcome, it may not be ready to trust that process with a high-risk decision or action.**

Controlled Authority

AI creates value when it can do something useful. That usually means access: access to information, applications, tools, and sometimes the ability to change something.

The governance challenge is determining how much authority the system actually needs. The wrong approach is to ask what the AI system can do. The better question is what it should be allowed to do at this stage of maturity. A useful AI process does not need unlimited authority. In many cases, the safest and most effective approach is to separate capabilities.

AUTHORITY LADDER

Increase autonomy in stages as reliability is demonstrated:

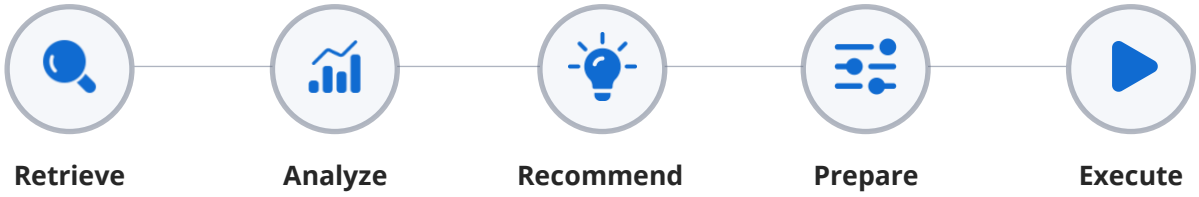

An AI system may be permitted to retrieve information, analyze it, draft a recommendation, prepare an action, and submit the action for approval. It may not initially be permitted to execute the action. This creates a practical path toward increasing autonomy without making autonomy an all-or-nothing decision.

Practical Example: Invoice Processing

An AI system is designed to review incoming invoices, compare them against purchase orders, identify discrepancies, and prepare invoices for payment. Over time, additional capabilities are added: updating vendor information, routing exceptions, preparing payment batches, interacting with the accounting system.

Individually, each capability may make sense. Collectively, the organization may have created a workflow in which one AI-enabled process can access vendor records, modify information, approve an exception, and initiate a transaction. The issue is not whether the AI is intelligent. The issue is whether the business has unintentionally weakened controls that existed for a reason.

WHAT CAN GO WRONG

Convenient connections can unintentionally weaken segregation of duties.

Define the Authority Envelope

Approved data sources	What information may the system retrieve?
Permitted tools	Which systems and functions may it use?
Allowed actions	What can it do without approval?
Restricted actions	What must it never do autonomously?
Approval gates	At what point must a person review or authorize the process?
Escalation rules	When should the system stop rather than continue?
Role separation	Should the same workflow be able to draft, approve, and execute?
Destination restrictions	Where can outputs be saved, shared, or transmitted?

> **PRINCIPLE**
>
> **AI authority should be earned through demonstrated reliability, not granted simply because the technology is capable of acting autonomously.**

Continuous Validation: Outcome and Process

AI-generated work can be persuasive. That is one of its strengths. It is also one reason governance cannot rely solely on whether a result looks reasonable. A professionally written explanation can contain a flawed assumption. A confident recommendation can rely on incomplete data. A correct total can conceal an incorrect process.

Executives are accustomed to reviewing outputs: the financial report is accurate, the contract summary is correct, the recommendation makes sense. With AI, this can create false confidence. There are two separate questions, and both matter.

OUTCOME QUALITY	PROCESS QUALITY
Was the answer correct?	**Was the answer produced correctly?**

MEMORABLE EXECUTIVE CONCEPT

A good answer does not prove a good process.

Two examples make the distinction concrete. An AI assistant produces the right answer but retrieves information from a sensitive data source the user was not supposed to access. An AI workflow prepares an accurate financial package and sends it to a distribution list containing an external address. In both cases the output passes review. The process does not.

Or consider a contract review workflow. An AI system summarizes a customer contract and identifies potential issues, and the summary looks accurate. The organization later discovers that the system reviewed an outdated contract version, omitted an exhibit, or failed to escalate a nonstandard liability clause. A human reviewer who checks only the summary may never see the process failure.

Organizations therefore need a repeatable way to test both AI outputs and AI behavior. This does not require every mid-market business to create an AI testing laboratory. Many organizations already have the foundations: reconciliations, tolerance levels, approval thresholds, escalation criteria, policy checks, and exception reviews.

ENTECH POINT OF VIEW

The governance opportunity is to translate existing business-control expectations into explicit tests for AI-enabled workflows.

What Should Continuous Validation Test?

Accuracy	Is the result factually and numerically correct?
Completeness	Were all required elements considered?
Source integrity	Did the system use approved and current information?
Scope	Did it use the correct entity, date, customer, employee, contract, or business unit?
Policy compliance	Did the process follow required business rules?
Permission compliance	Was all information accessed appropriately?
Consistency	Does the system perform reliably across similar cases?
Reproducibility	Can the organization recreate the result within acceptable tolerances?
Escalation behavior	Does the system recognize when it should stop and request review?

PRINCIPLE

AI governance must evaluate both the quality of the outcome and the quality of the process used to produce it.

Five Questions for Leadership

1. What AI systems and AI-enabled workflows are operating inside the organization?

WHY IT MATTERS

Many organizations have more AI adoption than leadership realizes. AI may exist in standalone tools, SaaS applications, Microsoft environments, departmental workflows, automation platforms, and unofficial employee practices.

A STRONG ANSWER

We maintain a current inventory of material AI use cases, including the business owner, purpose, users, connected systems, data accessed, actions performed, and risk classification.

2. What information can our AI systems access?

WHY IT MATTERS

AI risk often begins with access. A workflow may be helpful when connected to one approved source and risky when connected indiscriminately across shared drives, email, financial records, HR data, CRM information, and operational systems.

A STRONG ANSWER

For priority workflows, we know which data sources are approved, which identity and permissions the process uses, how sensitive information is restricted, and where results can be sent.

3. What decisions or actions is AI authorized to take?

WHY IT MATTERS

Recommendation and execution are not the same. An AI tool that drafts a journal entry presents a different risk than one that posts it. An assistant that recommends changing a customer account is different from one that can make the change automatically.

A STRONG ANSWER

We have defined the authority boundaries for material workflows, including what AI can retrieve, recommend, prepare, approve, execute, and distribute. High-impact actions have explicit approval and escalation requirements.

4. What happens when something goes wrong?

WHY IT MATTERS

Every important business process needs exception handling. The organization needs to know how the system behaves when information is missing, permissions fail, data conflicts, confidence is low, a threshold is exceeded, or the requested action falls outside policy.

A STRONG ANSWER

We have defined stop conditions, escalation paths, responsible owners, review expectations, and incident procedures for our higher-risk AI workflows.

5. Can we reconstruct and defend an AI-enabled decision?

WHY IT MATTERS

As AI becomes part of financial, operational, customer, employee, cybersecurity, and compliance processes, organizations will increasingly be asked to explain decisions by management, auditors, customers, insurers, regulators, legal counsel, and the board.

A STRONG ANSWER

For material AI-enabled processes, we can retrieve evidence of the request, scope, data sources, system interactions, approvals, actions, and final destination. We can explain both the result and how it was produced.

A Practical AI Governance Action Plan

Mid-market organizations do not need a large AI governance department to make meaningful progress. They do need clear ownership, focused priorities, and cross-functional involvement. A practical starting group may include executive leadership, finance or operations, IT, cybersecurity, compliance or risk where applicable, and leaders responsible for the business processes being evaluated.

> **OPERATING PRINCIPLE**
>
> **The first 90 days should focus on a small number of meaningful workflows rather than attempting to govern every theoretical AI use case at once.**

Days 1–30: Discover and Prioritize

- ✓ **Identify where AI is already being used**
  Look across enterprise AI platforms, embedded AI in SaaS applications, departmental tools, workflow automation, internal applications, and employee-created practices.
- ✓ **Inventory material AI tools and workflows**
  Capture business purpose, business and technology owners, users, information accessed, systems connected, outputs produced, actions performed, approval requirements, and destination of results.
- ✓ **Assign a business owner**
  The leader responsible for the business process should help define acceptable use, expected quality, required controls, and escalation requirements.
- ✓ **Classify workflows by risk**
  Distinguish lower-risk drafting and summarization from higher-risk sensitive data access, financial reporting, regulated processes, employment decisions, cybersecurity actions, transactions, and changes to systems of record.
- ✓ **Select three to five priority workflows**
  Choose workflows already in use, strategically important, connected to sensitive information, likely to scale, or capable of taking consequential actions.

Days 31–60: Define Guardrails

- ✓ **Define the Authority Envelope for each priority workflow**
  Work through the eight envelope dimensions: approved sources, permitted tools, allowed and restricted actions, approval gates, escalation rules, role separation, and destination restrictions.
- ✓ **Place approval gates where risk becomes significant**
  Including financial adjustments, external distribution, contractual commitments, employee actions, cybersecurity containment actions, material exceptions, and changes to systems of record.
- ✓ **Establish evidence requirements**
  Determine what needs to be retained to reconstruct the process.
- ✓ **Define what good looks like**
  Agree on standards for accuracy, completeness, scope, approved sources, permissions, policy compliance, approvals, and destination.

Days 61–90: Test and Operationalize

- ✓ **Create evaluation criteria and test failure scenarios**
  Turn business expectations into repeatable checks. Include incorrect, missing, and stale data, conflicting sources, unauthorized requests, permission failures, and requests outside defined scope.
- ✓ **Review workflow records, not just final answers**
  Review how selected workflows reached their conclusions.
- ✓ **Test permissions and escalation**
  Verify the process cannot retrieve or act beyond its approved authority, and that exceptions reach the correct person with enough context to act.
- ✓ **Establish governance metrics**
  Track ownership coverage, documented access, approval gates, exception rates, validation pass rates, recurring failures, and reconstruction time.
- ✓ **Determine whether autonomy should expand**
  Use evidence to decide whether a workflow remains in recommendation mode, progresses to preparation, or earns narrowly defined execution authority.

MID-MARKET GUIDANCE

Do not begin with 50 use cases. Learn on a small number of real processes, then apply what works more broadly.

AI Governance Maturity Checklist

Executives can use this self-assessment to identify the organization's current position. Most mid-market organizations will find themselves at different levels across different dimensions. That is normal. The objective is not to declare the organization "AI mature." It is to identify where increasing AI adoption is outpacing the control environment and decide where to strengthen governance first.

Dimension	Level 1: Experimenting	Level 2: Controlled Adoption	Level 3: Governed Scale
AI inventory	AI use is largely informal or unknown.	Material use cases are documented.	Inventory is maintained as AI adoption and integrations change.
Ownership	Responsibility is unclear or assumed to belong to IT.	Priority workflows have named business and technology owners.	Ownership, accountability, review, and escalation roles are embedded into operating processes.
Access, permissions & approvals	AI access depends largely on existing user access or convenient connections; human review is informal.	Approved sources, minimum access, and documented approval requirements are defined for priority workflows.	Access is actively reviewed and tested; authority is risk-based and autonomy expands only when reliability is demonstrated.
Workflow oversight	Teams mainly review final outputs.	Important processes capture key steps, sources, actions, and approvals.	Material workflows are reconstructable from request through final destination.
Validation & improvement	Users decide whether results look reasonable; problems are addressed when users complain.	Priority workflows have defined quality checks; exceptions are reviewed and recurring issues corrected.	Output and process validation occurs continuously; failures are systematically converted into controls, tests, and process improvements.

How to Use the Checklist

- Identify the highest-risk workflows first rather than averaging maturity across the entire company.
- Look for dimensions where AI use is expanding faster than the control environment.
- Define the next practical step for each priority gap. The target is improvement, not a perfect score.
- Reassess after major changes to AI tools, data sources, integrations, permissions, or workflow authority.

> **EXECUTIVE PERSPECTIVE**
>
> **The most important maturity question is not "How advanced is our AI?" It is "Are our controls keeping pace with how AI is actually being used?"**

The Competitive Advantage of Governed AI

The AI governance debate is often framed incorrectly. Innovation on one side, control on the other. Move quickly or move safely. That is a false choice.

Organizations will not capture the full value of AI by operating without boundaries; eventually, a lack of trust, visibility, accountability, or control will slow expansion. They also will not capture the value of AI by creating such a heavy governance structure that responsible experimentation becomes impossible.

> **ENTECH POINT OF VIEW**
>
> **The more useful objective is governed execution.**

A Practical Operating Model

- Know where AI is being used.
- Understand the complete workflow.
- Give systems only the authority they need.
- Maintain meaningful approval points.
- Define what a good result looks like.
- Test both the outcome and the process.
- Learn from failure.
- Increase autonomy when the evidence supports it.

This approach creates a different executive posture toward AI. Not blind enthusiasm. Not paralysis. Confidence. The organizations that understand what their AI systems are doing, control what those systems are allowed to do, and continuously validate their performance will be positioned to move faster. They will connect AI to more important workflows because leadership trusts the operating model around it, expand autonomy selectively because reliability has been demonstrated, and answer questions from boards, auditors, customers, and regulators because they understand how AI-enabled decisions are made.

> **CLOSING PERSPECTIVE**
>
> **The advantage will not go only to the organizations that adopt AI first. It will go to the organizations that learn how to scale it with control.**

Start With Clarity

For many executive teams, the first challenge is not selecting another AI tool. It is developing a clear view of what is already happening. An Entech AI Governance Assessment is designed to help leadership evaluate the organization's current position and establish practical priorities.

ENTECH AI GOVERNANCE ASSESSMENT

Turn AI Adoption Into a Clear, Governed Roadmap.

The Assessment Can Help Leadership

- Understand current AI adoption across the organization
- Identify significant governance and control gaps
- Clarify ownership and accountability
- Evaluate higher-risk AI-enabled workflows
- Assess how data, permissions, approvals, and oversight currently operate
- Prioritize immediate actions and create a practical roadmap for governed AI adoption

The objective is not to create a large new layer of bureaucracy. It is to help leadership build the visibility, control, and confidence required to trust AI with more of the business.

> THE QUESTION TO TAKE BACK TO YOUR EXECUTIVE TEAM
>
> **As AI becomes more capable, are we building the visibility, control, and confidence required to trust it with more of the business?**

About Entech

Entech helps organizations align technology, cybersecurity, compliance, and AI initiatives with business strategy. Our work combines practical advisory guidance with the operational expertise required to turn strategy into governed execution.

START A STRATEGY SESSION