

Why AI Stalls: 4 Decisions Leaders Can't Afford to Delay

2026 Executive Research





Across hundreds of IT decision-makers, CapTech's 2026 Executive Research found a consistent pattern: organizations are investing in AI with confidence, but progress often hinges on operating decisions made too late in the AI lifecycle.

This friction shows up most clearly in four areas: success metrics, stakeholder alignment, legacy integration, and governance. Each decision shapes whether AI creates measurable value, gains organizational support, fits into the technology environment, and remains accountable once it's live.

These four decisions appear in every modernization effort. AI makes them continuous. Output falls within a range rather than landing on a fixed result, so "working" becomes a threshold the organization defines and tests. The model is licensed rather than owned, and even a pinned version carries a retirement date set by the provider. Each new capability can change what the system can do and what the organization must decide.

Agents raise the stakes further. As AI systems move from answering questions to taking action, each decision expands. Success becomes a completed task rather than a correct response. Alignment defines which actions the system may take without approval. Integration determines what it can access and change. Oversight creates the record of what it did, why it acted, and whether the result met the organization's standards.

The four decisions remain. AI makes the cost of leaving them open rise with every release.



86%

of IT decision-makers said their organizations could move faster with AI but are held back by internal decision-making, not the tools themselves.

AI Decision Discipline Connects Value, Authority, Architecture, and Oversight

The four decisions move together across the lifecycle. Each must be explicit, owned, and revisited as the system and capabilities change.

01 DEFINE SUCCESS

What does “working” mean?

- Own the eval set
- Set the release threshold
- Connect performance to business value

02 ALIGN AUTHORITY

Who may decide and act?

- Set the autonomy boundary
- Name output accountability
- Identify the quality-bar arbiter

Decision Discipline

every release

04 MAINTAIN OVERSIGHT

How will the system stay accountable?

- Log behavior and watch performance
- Keep a named owner on watch
- Define the trigger and fallback

03 PRESERVE OPTIONALITY

What can be changed later?

- Create a replaceable model boundary
- Keep critical steps deterministic
- Carry permissions through the system

CapTech’s research points to a practical challenge for AI leaders. Operational decisions are often treated as secondary to technical execution, then become more expensive to resolve once systems, teams, and expectations are already in motion.

The timing matters, but so does preserving room to adjust as AI capabilities change. Some decisions become more expensive when deferred. Others create more value when they can be revisited as requirements, models, and capabilities evolve. The challenge is deciding clearly enough to move forward while preserving the flexibility to adapt as technology evolves.

A decision made while a pilot is still flexible can clarify the work ahead. The same decision made later requires teams to rework architecture, renegotiate ownership, or rebuild measurement after the fact. The price of indecision climbs the longer it waits.

Every phase in the AI lifecycle offers an opportunity to course correct: identify the next decision that matters, make it explicit, and carry it forward as the initiative evolves. This decision discipline accelerates progress and improves AI outcomes.



Why AI Stalls in the Same 4 Places

The same four decisions appeared repeatedly in the research, and they live in the gap between technical capability and business value. These decisions define what success means, who has authority, how AI connects to existing systems, and how performance is monitored after deployment.

The priority rankings reveal why they're deprioritized: more visible technical concerns rise first. Security and compliance ranked as top-two priorities for 48% of IT decision-makers. Clear return on investment at 38%, cloud computing 32%, and data architecture at 30%. The operating decisions that most frequently determine execution sat lower, with governance, stakeholder alignment, and legacy compatibility clustered at 24% to 27%.

The risk rankings tell the other side of the story. Unaligned stakeholders were named as a top derailer by 28% of respondents. Legacy integration surfaced as a risk to progress for 26%. And governance approval was cited by 41% as a point where initiatives stall or reverse.

This data uncovered a real blind spot: the decisions that leaders ranked as lowest priorities are their highest risks to AI progress.



The Pattern Runs Deep

The priority-to-risk inversion extends further. Adoption rises from 10% priority to 22% risk. Tech debt ranks as a top-two priority for only 9% of respondents, yet 29% cite it as a risk to progress. When the pattern repeats across topics, it points to an operating challenge rather than an isolated oversight.

AI programs tend to focus early attention on visible technical work. Less visible decisions about ownership, incentives, measurement, and operating constraints receive less attention, until they affect progress.

The Cost of Late Decisions

Every AI lifecycle includes decision phases that either reduce uncertainty or carry it forward. The real cost of these decisions increases as an initiative moves from concept to pilot, from pilot to production, and from production to scale. Early decisions can still shape the build. Later decisions often become more expensive by requiring teams to adjust what has already been built, agreed to, or measured.

CapTech’s Research Found a Late Decision-Making Pattern Across All 4 Areas

- **Success metrics** are often decided after programs are already underway. Over a third of teams define at-scale key performance indicators in-program.
- **Stakeholder alignment** is settled later still, remaining unresolved until production for 38% of organizations.
- **Legacy integration** and tech-debt decisions are spread across every phase, and they’re rarely concentrated early enough to guide architecture.
- **Governance** waits the longest. A majority of IT decision-makers said their organizations delayed defining AI governance until systems were already live or close to production.




Rework Tax

Technical debt reflects the cost of choices embedded in technology. Rework tax is its operational counterpart: the time, budget, and attention teams spend revisiting work because a consequential decision remained unresolved. As an AI initiative advances, that tax appears in rebuilt measurement, renegotiated ownership, revised architecture, and governance added after deployment. Every delayed decision adds to the bill.

But AI creates a second cost. Some decisions are made too rigidly and have to be revisited when models, economics, or capabilities change.

The Cost of an Unresolved Decision Rises After Each Commitment

AI teams pay a rework tax when measurement, ownership, architecture, or oversight must be rebuilt after the work has advanced.



Rework tax creates a compounding effect across the AI lifecycle. Metrics added after launch are harder to tie to a clean baseline. Ownership clarified after handoff requires teams to reopen decisions. Architecture adjusted after integration can turn a design choice into rework. Governance added after deployment can protect what is already live, but the organization has less room to build oversight into the system from the start.

For organizations already in production, the cost curve isn't fixed. The next release, model update, handoff, or measurement cycle creates another opening to lower the cost of your next decision. It's never too late to course correct.

Decision discipline is the practice of defining the decision, assigning ownership, and revisiting it at the next phase before the cost of changing course rises again.



4 Decisions Leaders Can't Afford to Delay

01

SUCCESS METRICS

Success metrics determine what the organization is trying to prove before teams build toward it.

02

STAKEHOLDER ALIGNMENT

Stakeholder alignment determines what AI is expected to do, who has authority to make decisions, and how people will be rewarded under the new operating model.

03

LEGACY INTEGRATION

Legacy integration determines how AI connects to the systems, data, and architecture an organization already depends on.

04

ONGOING OVERSIGHT

Governance determines how AI is monitored, adjusted, and held accountable after deployment.

Success Metrics

Success metrics determine what the organization is trying to prove before teams build toward it. This decision sounds basic, but it’s one of the easiest to defer under the pressure to show progress. Without an early baseline, teams may know that a system is working but struggle to demonstrate how much value it created, where value stalled, or which release changed the result.

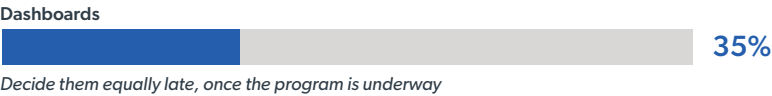
CapTech’s research indicates how often measurement trails execution. 85% of IT decision-makers said pressure to move quickly led their organizations to make key AI decisions before they had defined success criteria or guardrails. The metrics that should anchor a program frequently arrive after the work is underway. 36% of organizations define at-scale key performance indicators only once in-program, and 35% decide their dashboards equally late.

This decision becomes expensive when teams try to prove value without the evidence structure to support it. A vague sense of improvement leaves no clear threshold, no shared definition of progress, and no reliable way to connect the work to business outcomes. Organizations can still establish measurement after a system is live, but the baseline becomes harder to reconstruct and progress becomes harder to defend.

A live system creates a new starting point for measurement. Define success against the live baseline, select the few business metrics that matter most, and instrument them before the next release. The measurement may arrive later than ideal, but it can still improve the quality of the next decision.

Pressure to Move Quickly Leaves AI Measurement Decision Unresolved

WHERE MEASUREMENT LANDS LATE



WHAT RESPONDENTS FOUND HELPFUL



85%
Said pressure to move quickly led their organization to make key AI decisions before success criteria or guardrails were defined.

CAPTECH INTERPRETATION

The research indicates that measurement often trails execution, while respondents value tools that establish a starting point and a clear bar for progress.



What is an Eval Set

An eval set defines what “working” means for an AI system. It’s a fixed collection of representative tasks used to measure AI performance over time. Every release is scored against the same set, creating a consistent benchmark for determining whether the system is improving, holding steady, or declining. In practice, it’s the closest equivalent AI has to a regression test.

What to Decide Before the Next Release

Decide what today’s performance is, before it changes. Record the current state measure that the next return-on-investment claim will be compared against, such as cost per transaction or hours per case. 60% of IT decision-makers found establishing baseline metrics for later comparison helpful.

Decide the bar that the next release has to clear. Set the eval threshold required to proceed and name who approves the release once it is met. In the research, 62% of IT decision-makers found pre-set success thresholds or milestones helpful.

Decide who owns the eval set, and what belongs in it. Name the person accountable for the representative task set that defines “working,” including expected and failure cases. This turns the need for early success criteria into a release-level measurement decision for AI.

Decide the unit-economics bar and the two or three numbers leadership sees. For example, determine the cost per resolved task at production volume including human review, reported in business terms, with a named owner. Escalation rate is the line item that flips pilot ROI at scale. A majority of organizations found real-time tracking helpful.



SIGNALS TO WATCH: DIMINISHING RETURNS

IT decision-makers identified four warning signs that AI is no longer paying off at the expected rate.

- Business outcomes flatten (45%)
- Core metrics no longer improve between releases (44%)
- Integration complexity outweighs the benefit (39%)
- Opportunity costs to other projects rise (38%)

Each signal is easier to interpret when the team has already defined what progress should look like.

Stakeholder Alignment

Stakeholder alignment determines what AI is expected to do, who has authority to make decisions, and how people will be rewarded under the new operating model. When that alignment is unclear, execution can continue for a while, but unresolved questions tend to reappear at handoff points, budget discussions, adoption planning, and production support. This is where rework tax becomes most visible in the operating model.

CapTech’s executive research points to a meaningful gap between priority and risk. Stakeholder alignment ranks last as a strategic priority at 25% but first as a risk to progress at 28%. The broader operational effect extends further. **56% of IT decision-makers said decision-making and stakeholder misalignment slow AI modernization more than anything else.**

The human side of alignment often lands late. Incentive alignment, or how people are rewarded under the new system, is one of the latest calls organizations typically make. Only 22% of IT decision-makers settle it during the pilot, while 38% push it to production. Role and structure changes follow a similar pattern, with 48% landing at the transition. As a result, teams may build AI capabilities before the organization has resolved how work, accountability, and authority will shift around them.

A focused decision-rights session can restore alignment after work begins and prevent reopened discussions. To accelerate progress, clarify who owns each call, what has already been decided, and which questions need escalation before the next handoff.

Organizations Underestimate the Operational Impact of Alignment

Stakeholder alignment receives limited strategic attention, even as respondents identify misalignment as a major drag on AI modernization.

25%

Of IT decision-makers rank stakeholder alignment as their lowest strategic priority.

However...

56%

Of IT decision-makers say misalignment causes operational slowdown more than anything else.

The gap is not just priority versus risk. It shows up as slower decision-making and modernization.

What to Decide Before the Next Commitment

Decide who defines the autonomy boundary. Name the business, risk, legal, and technical stakeholders responsible for determining which actions the system may take independently, which require approval, and which are prohibited. The authority limits matter. Deciding who sets them matters first.

Decide who is accountable for the output, not the project. When the system is wrong in front of a customer, whose name is on the decision? Accountability for outcomes should be established before deployment, not assigned after a failure.

Decide who arbitrates the quality bar. The business may say the experience is not good enough while engineering points to passing eval scores. Both can be right. Name the tiebreaker before the disagreement. 48% of IT decision-makers found cross-functional decision structures helpful.

Decide what can change and who can change it. Some choices, like regulatory requirements, may be fixed. Others may need to evolve as the implementation matures. Establish ownership for revisiting decisions before a new capability, risk, or business need forces the conversation. Flexibility should reduce rework, not create recurring debate.



SIGNALS TO WATCH: MISALIGNMENT

- The same decision reopens in multiple meetings, which suggests ownership is unclear.
- Teams explain the project goal in different terms, pointing to competing definitions of success.
- Progress depends on one person being present because decision rights live in individual memory rather than in the operating model.
- Ownership questions routinely get deferred to a later phase.

When two or more of these signals appear, the issue is usually structural. Naming the decision rights before the next handoff reduces the amount of work that has to be renegotiated later.



Legacy Integration

Legacy integration determines how AI connects to the systems, data, and architecture an organization already depends on. The research suggests that legacy decisions create risk in two directions. Moving too slowly can constrain AI. Moving too quickly without a path to revisit key decisions can create a new form of lock-in.

Legacy integration ranked last as an AI priority at 24%, but it surfaced as the number-two risk to progress at 26%. The most opportune timing for this decision is nuanced. How you decide to modernize is just as important as when you decide. 72% of IT decision-makers said early modernization choices have limited their organization’s flexibility to change AI direction, and 73% said those decisions limited their ability to scale.

Our research supports a more precise modernization strategy: move early while preserving options. A vendor connection, integration pattern, or architecture decision can become a constraint if the next model, use case, or platform direction requires a rebuild. Anything you can’t change your mind about is functionally legacy, regardless of when it was built.

Reversibility is the principle that resolves this tension. In practice, the goal is to modernize early enough to support AI, and flexibly enough that a future change becomes a controlled switch instead of a full rebuild.

Organizations that already feel locked into a platform or architecture can still buy back flexibility. A translation layer, abstraction point, or phased retirement of hard dependencies can reduce the cost of the next change without requiring a large cutover.

Organizations Need Flexibility, but Many Design It Out Early

Early modernization choices can constrain both AI direction and scale, while respondents overwhelmingly associate reversibility with effective modernization.

CHANGE DIRECTION

72%

Said early modernization choices limited their organization’s flexibility to change AI direction.

SCALE

73%

Said early modernization choices limited their organization’s ability to scale.

REVERSIBILITY

92%

Agreed that preserving the flexibility to reverse a decision enables effective modernization.

Reversibility turns a future change into a controlled switch instead of a full rebuild.



Model Pinning Creates a Decision Deadline

An organization consuming a model from a vendor can pin a specific version, and should. But a pinned version is frozen in both directions: it will not degrade, and it will not improve.

Improvements do not arrive as occasional headline launches. Vendors ship updated versions steadily, and every gain in quality, speed, and cost comes bundled into one of them, so a pinned system takes none of them until it moves.

Providers also retire versions on their own schedules, which makes pinning temporary, rather than permanent. That leaves two clocks. One is the notice before retirement, which the provider sets. The other is the time it takes to qualify a replacement, which the organization controls. Exposure is the gap between them.

What to Decide Before the Next Architecture Commitment

Decide where the model boundary lives, and what it takes to leave. A single place in the architecture that calls models, with pinned versions, rather than direct Software Development Kit (SDK) integrations spread across services. The real test of lock-in is whether you can re-run evals against a candidate replacement model in a day or a quarter. Minimum viable architecture was the highest-ranked legacy integration tactic, selected by 35% of respondents.

Decide what stays deterministic. Determine which steps must never be probabilistic: calculations, entitlements, transactions, compliance checks. Deciding where AI must not be is an underrated architecture decision, and it’s what keeps blast radius small when something does go wrong.

Decide what the system may see, and how permissions travel with it. For AI, the legacy blocker is often entitlements, not integration. If retrieval can’t carry the user’s permissions through, you either have a leak, or a system no regulated function can use.

Decide which tech debt blocks scaling, and who signs off on deferring the rest. Name the known issues that must clear before scale and the ones deferred with intent. The decision should be explicit, documented, and owned. Tech debt rises from 9% as a priority to 29% as a risk.



SIGNALS TO WATCH: ONE-WAY DOOR DECISIONS

- Changing vendors would require updates across dozens of systems rather than one integration layer.
- The architecture assumes the current model, tool, or vendor will remain the long-term standard.
- Temporary standardization has become a permanent operating assumption.
- New use cases are being reshaped to fit the platform, instead of the platform flexing to support the use case.
- AI model replacement is handled as a transformation effort instead of a routine maintenance activity.

When leaving would cost more than arriving, modernization has created a new constraint. The next decision should reduce that dependency before it expands.

Technical debt is the future cost embedded in today’s technology choices.

Rework tax is the cost of operational decisions left unresolved while the work moves forward. AI programs often carry both, and each makes the next change more expensive.



LegacyLift® Keeps Systems Flexible

Most modernization efforts aim to move organizations away from legacy systems quickly. Convert-and-rehost tools can accelerate that work, but speed can come with a tradeoff: old system rigidity becomes new vendor dependence.

CapTech takes a different acceleration path. LegacyLift® reads a live legacy system and extracts the business rules, requirements, and dependencies first, then generates standards-based code the organization can own, staff, and extend. By making the logic explicit and the output portable, LegacyLift™ helps keep the next architecture or vendor decision open. Understanding the system is the mechanism. Optionality is the payoff.

Design for commoditization.

Organizations can reduce lock-in by treating AI models as commodities that can be replaced over time rather than permanent platform commitments.

Ongoing Oversight

Governance determines how AI is monitored, adjusted, and held accountable after deployment. In many programs, governance receives attention at approval, then becomes less visible once the system is live.

Governance approval is the most-cited point where initiatives stall or reverse, named by 41% of IT decision-makers, while governance itself ranks lower as a strategic priority, with only 27% naming it as a top-two AI priority. Timing is a major part of the challenge. 71% of IT decision-makers said their organizations delayed defining AI governance until systems were already live or close to production.

The data also shows that oversight often thins after approval. On average, departmental involvement drops 26 points from the approval phase into monitoring. Finance and product department involvement each drop 37 points, and sales department involvement drops 39 points. The result is a gap between the teams that approve AI and the teams that stay engaged while it runs.

Effective governance has to operate as part of the system. That means visible performance, named accountability, and pre-set intervention triggers that guide when to pause, adjust, or roll back a release. Oversight can still be retrofitted after systems go live by building these decisions into the next release more deliberately.

The Teams That Approve AI Often Don't Stay Involved Once It's Live

Governance is frequently defined late, and departmental involvement falls sharply between approval and ongoing monitoring.

7.8

Departments involved in AI approval



3.7

Departments involved in AI monitoring

CAPTECH INTERPRETATION

Our research indicates that this drop-off is significant enough to stall progress.

What to Decide Before the Next Deployment

Decide what you log, and what you watch. Capture full traces at launch, including inputs, retrieved context, tool calls, and outputs. Then define the measures that matter in production, including eval scores across releases, escalation rate, unsupported-claim rate, tool-call error rate, and cost per resolved task. Without a record of how the system behaves, teams can’t investigate incidents, explain outcomes, or improve performance over time.

Decide who stays on watch. Name the person accountable for the system in production, including who reviews performance, who investigates anomalies, and who has authority to escalate concerns. Oversight only works when responsibility continues after deployment.

Decide the trigger and the fallback. Define the conditions that force a pause, review, or rollback, and identify the pinned version the system will return to if those conditions are met. A trigger without a fallback is only an alert. Guardrails over perfection was the highest-ranked oversight tactic, selected by 40% of respondents.



SIGNALS TO WATCH: EMPTY WATCHTOWERS

- No one can name who is accountable for the system after the launch team has moved on.
- Monitoring only shows whether the system is running, not whether it’s still right.
- The last formal review happened before the system went live.
- Teams describe the system as working without a current measure that proves it.

If the organization can’t name what would trigger a pause, oversight is incomplete. The next decision is to define the trigger before the next issue forces the question.





Decision Discipline in Practice

Organizations that rate themselves highly on technology and AI maturity report stronger progress across all four decision areas. Their responses indicate **AI maturity comes from a repeatable operating habit: define the decision early, assign ownership, preserve flexibility, and maintain accountability as the initiative moves forward.**

The impact of this habit is significant: 91% of these high-tech, AI-mature respondents are confident they will hit their technology priorities over the next 18 to 24 months.

Another shared pattern is earlier clarity. Each decision resolved before the next commitment lowers the rework tax: less measurement to reconstruct, less ownership to renegotiate, fewer technical choices to unwind, and less oversight to retrofit after deployment.

Earlier, however, should be treated as a direction rather than a missed deadline. Every phase of the AI lifecycle creates another chance to improve the next decision.

AI-Mature Organizations Report Overwhelming Confidence in Near-Term Priorities

HIGH TECHNOLOGY + AI MATURITY

91% confident they will achieve their technology priorities over the next 18 to 24 months.

RESEARCH BOUNDARY

These findings show confidence and support for decision frameworks, not that frameworks caused the confidence result.



92%

of organizations agreed that modernization fails without a decision framework, even with the best technology in place. And 93% said the right framework made modernization more effective.

Better Decisions Lower the Cost of AI Progress

AI conversations have focused for years on models, platforms, and technical capability. Those factors still matter, but CapTech's 2026 Executive Research shows that many stalled initiatives share an operating pattern: the decisions that determine value, ownership, integration, and oversight are often resolved after the work has already advanced.

That sequence raises the cost of progress. A success metric defined late without an eval set is harder to prove. Alignment clarified late reopens decisions about system authority that should already guide execution. Legacy choices made without reversibility can constrain the next direction. Governance added after deployment can still protect the system, but it has less influence over how the system was designed.

The solution starts with the next decision. Name the decision that matters next, make it explicit while the initiative can still adapt, and carry it into the next phase with clear ownership and evidence.

AI initiatives deliver when operating decisions move with the technology: early enough to shape the work, reversible enough to survive the next capability shift, clear enough to assign ownership, and disciplined enough to avoid paying for the same work twice.



RESEARCH METHODOLOGY

This executive research was conducted by The Harris Poll on behalf of CapTech from May 8–22, 2026, through a survey of 300+ IT decision-makers at the director level and above. Participants work in IT functions accountable for technology outcomes at enterprise organizations of 250 or more employees, all running AI in production beyond the pilot phase, across 20 industries within the U.S.



Brian Bischoff

Chief Technology Officer

bbischoff@captechconsulting.com
804.301.0397



Bree Basham

Principal, Go-to-Market & Growth Strategy

bbasham@captechconsulting.com
704.604.1791



Let's do next together.®

CapTech®

captechconsulting.com

  [@captech_consulting](https://www.instagram.com/captech_consulting)