

 Leading the human + AI enterprise

A new playbook for **strategic portfolio management**

Based on primary research with 308 senior enterprise leaders

FOREWORD LETTER

When good strategy meets reality

If you've led a strategy process recently, you know this feeling: The offsite went great. Everyone aligned on priorities. The plan made sense. Then three months later, you're in a QBR wondering why so little has moved forward.

Here's what we're learning: it's not a strategy problem. It's not even a leadership problem.

We analyzed responses from 308 senior leaders across strategy, portfolio, technology, and delivery functions. What we found was striking: 80% say their strategic priorities are clear or very clear. But fewer than 1 in 8 can show that three-quarters of their actual work connects back to those priorities in any system they use.

Think about that gap for a moment. They know what matters. They just can't consistently do it.

And AI? It's making this more obvious. The real returns on AI don't come from which models you choose. They come from whether your strategy, your work, and your execution live in one connected place or scattered across a dozen tools. AI amplifies whatever you already have. If your organization is fragmented, AI will show you exactly where and how.

“Our people are faster, but the organization doesn't feel faster.”

We heard this from leader after leader. It's the most honest description of where most enterprises are right now.

The organizations making progress aren't waiting for better AI or a better plan. They're asking a harder question: Why can't we move when we need to? What's actually in the way?

This report is built around that question. And the answer turns out to be more structural – and more actionable – than most of us have been willing to admit.



Tal Saraf

CIO and SVP of Engineering, Atlassian

Executive summary

The data is unambiguous on what is *not* holding organizations back: strategic clarity. 80% of senior leaders say their priorities are clear or very clear.

What the research points to instead is a structural constraint – the speed at which an organization can actually move humans, agents, compute, and capital when strategy demands a change. We call this mean time to pivot. In most enterprises, it is measured in months, not days.

That finding runs counter to where most organizations are investing. Better planning processes. Tighter OKRs. More frequent priority reviews. Data from 308 senior enterprise leaders shows none of it is moving the needle. The organizations rebalancing priorities quarterly are not outperforming those that do it annually. The bottleneck is not how often you re-rank. It is whether your organization can physically act on a change when one is needed.

The research points to **three structural barriers** that together determine mean time to pivot.



1. The Clarity Trap

The mistaken belief that strategic clarity is the binding constraint. It is not. 80% of leaders say their priorities are clear. Only 11% have three-quarters or more of their portfolio work explicitly linked to those priorities in any tracking system. The gap between knowing what to do and being able to do it is not a clarity gap. It is architectural.



2. Frozen capital


78% of organizations operate on funding models designed for a world where strategy changed once a year. Fixed annual budgets and project-based funding structures have no mechanism for the continuous reallocation that pivot speed requires. When a priority changes, three separate reallocation challenges unfold in three separate systems – headcount in HR, project budget in Finance, and compute and AI spend in IT – managed by three different teams, on three different timelines.



3. The AI paradox

Every enterprise in our study has invested in AI. Most leaders can point to individual productivity gains. What they cannot point to is organizational speed. The reason is structural: the single strongest predictor of AI effectiveness is not model sophistication. It is whether your strategy, your work, and your execution data exist in a connected system. AI amplifies whatever operating model it finds. In a fragmented organization, it produces more reports, more dashboards, and more AI initiatives untethered from the strategic bets they were meant to advance.

The finding that should most directly challenge your current SPM investment:
Organizations that have moved to quarterly or monthly priority reviews are not outperforming those that haven't.



The frequency of re-ranking has effectively zero predictive relationship with how fast the organization actually moves. If your current Strategic Portfolio Management investment thesis is built around planning cadence, this data challenges it directly. The organizations that move fastest are not running better planning cycles. They have built a different kind of operating model.

The organizations closing the gap share one structural shift: a living operating model where a change in direction actually changes what teams do, what capital funds, and what AI agents work on. That is what mean time to pivot measures – and why reducing it is now a strategic necessity, not a planning improvement.

In the pages that follow, we unpack the three structural barriers in detail, show how AI changes the stakes, and describe the architectural shift that distinguishes organizations that can pivot in days from those that take quarters.

The clarity trap

There is a pervasive assumption that the primary challenge of strategy is getting clear on direction. The research does not support this. Leaders are clear. The system that connects their clarity to action is not.

80%

have clear or very clear strategic priorities

54%

say priorities cascade effectively across teams

11%

have most execution work connected to strategy in a shared system

The cascade breakdown

Three factors consistently interrupt the translation from strategic intent to team-level action. In large enterprises, layers of management represent an opportunity for intent to be adjusted, softened, or lost. A third of organizations cascade strategy through functional silos using different tools, making cross-functional alignment structurally impossible. And qualitative data reveals a consistent pattern: objective, value-based prioritization is routinely overridden by whoever champions loudest or the HiPPO effect.

The consequence is measurable: 60% of leaders report that risks are identified sometimes late, mostly late, or after impact – the predictable result of a cascade process that distorts signal before it reaches the people who need to act on it.

Increasing how often priorities are reviewed does not resolve this. The data shows no relationship between review frequency and how quickly the organization actually moves – a finding the subsequent sections explain.

“

I think it is always the ones that championed the most... whoever gets to speak the loudest.

VP of IT

Mid-Market (1-5k employees),
Financial Services

The accountability gap

72% of leaders identify competing priorities as a chronic overstretch problem for key teams – the structural condition that makes meaningful accountability nearly impossible. Only 17% operate in a system where accountability is structurally embedded – where talent and funds shift based on transparent performance data. The remaining 83% run on manual status updates and personal judgment.

Accountability without visibility into strategic impact is not accountability. It is compliance for compliance sake. The question is not how to enforce it harder. It is what it actually takes for an organization to move when the signal is clear – and why most cannot.

The metric that matters: Mean time to pivot

We propose a single organizing metric defined by Atlassian’s Derek Huether: mean time to pivot – the time from a clear signal to an actual change in the allocation of humans, agents, compute, and spend. Not the time to update the strategy document. The time to physically move resources and encode that change into the operating model.

“How quickly can you actually react to a global, industry, or market event? How quickly can you figure out what to do, decide on a strategic pivot, and have that actually be meaningfully understood and executed throughout your organization?”



Derek Huether

Head of Solutions, Strategic Portfolio Management, Atlassian

The clarity trap established that strategic intent is not the constraint. 80% of leaders are clear on their priorities. The question this section answers is: what does it actually take to act on that clarity when conditions change?

The relationship between decision speed and organizational adaptability is stark. 70% of highly adaptable organizations rate their decision-making as effective. Among rigid organizations, the figure is 5%. Highly adaptable organizations are 14 times more likely to rate their decision-making as effective – and the structural factors that predict adaptability are the same ones that predict faster decisions: governance effectiveness, system connectivity, and funding model design.

This is not a mindset difference. It is an architectural difference. The organizations that decide fast do not have braver leaders or better offsites. They have built the structural conditions that allow routine reallocation to happen without requiring a governance meeting to initiate it.

Fewer than **one in three organizations** decide within a week of a performance signal. Nearly two-thirds take a month or more. 5% take a full quarter.



What the data shows drives organizational speed

To identify what actually predicts organizational speed, we ran a regression analysis across the full dataset. The findings point to four structural factors:

1. System integration

What the data shows about execution speed

When the tools teams use to execute work share data with the systems where strategy is set, blockers get resolved significantly faster. In our regression analysis, we found that system integration accounts for 17% of the variation in dependency resolution speed, and 21% of an organization's ability to automatically mitigate risks before they become blockers – the highest figures of any structural variable tested.

2. Governance effectiveness

The foundation that compounds

Effective governance and review frequency together explain 24% of the variation in organizational adaptability. When domain connectivity and AI use are layered onto strong governance, explained variance rises to 28%. The implication is important: AI and connected systems amplify the impact of good governance. They do not replace it. The organizations investing in AI without first addressing governance and connectivity will see diminishing returns.

17%

of dependency resolution speed explained by system integration

18%

of AI effectiveness explained by domain connectivity

14x

more likely: adaptable orgs rate decision-making as effective vs. rigid ones

3. Domain connectivity

The driver of AI effectiveness and decision confidence

The organizations getting the most from AI are not those with the most sophisticated models. They are the ones whose strategy, work, and execution data exist in a connected model. Domain connectivity is the single strongest predictor of AI effectiveness – accounting for 18% of the variation in AI effectiveness scores. It also accounts for 7% of the variation in data trust. Neither outcome is achievable by deploying more capable AI models without first building the connected model they require to reason over.

4. Funding model

The structural accelerant most organizations have ignored

Of all funding variables tested – review frequency, governance structure, budget ownership, rebalancing rhythm – only the shift toward dynamic, outcome-based funding shows a consistent relationship with how quickly organizations respond to strategic signals. Organizations on fixed annual or project-based budgets show no meaningful speed advantage regardless of how often they review priorities. The next section explains why.

What doesn't predict speed: how often you re-rank. Across regression models, the frequency of priority re-ranking based on performance data has effectively zero predictive relationship with decision velocity, dependency resolution, or AI effectiveness. This is the finding that should most directly challenge a common SPM investment assumption – and it runs throughout the rest of this report.

Three metrics organizations should be tracking but aren't

The current standard – portfolio health, initiative status, OKR completion rates – measures planning fidelity. It answers: are we doing what we said? It does not answer: can we do something different when the signal tells us we should? The three metrics that measure mean time to pivot directly are almost entirely absent from enterprise dashboards today.

Execution connectivity rate

What percentage of active execution work is natively linked to a current strategic bet in a shared system?

Benchmark: Only 11% of organizations achieve 75% or higher linkage. This is the most direct measure of whether the strategy-to-execution connection exists structurally, or is being managed manually.

Capital reallocation velocity

When a priority changes, how quickly do funding and capacity actually move to reflect that change?

Not the time to make the decision – the time for resources to follow it. An organization that tracks governance cadence but not reallocation velocity is measuring its intention to move, not its ability to.

Compute attribution coverage

What percentage of AI and cloud compute spend is attributed to a specific strategic bet with a defined outcome?

For most organizations, the answer is near zero. This gap is explored in the next section. It is increasingly the question that boards will ask when evaluating AI ROI – and most organizations have no mechanism to answer it.

3 questions every executive should be ready to answer

Mean time to pivot is the metric your board should be asking about alongside revenue growth and margin. It is a proxy for whether your leadership can actually execute a strategy change – or whether the operating model will absorb that change and continue on its original trajectory regardless. Pull three numbers before your next board conversation: what percentage of active projects are linked to a current strategic priority in a single system; what percentage of AI compute spend can you attribute to a specific outcome; and when you last changed a strategic priority, how long did it take for funding and headcount to follow. The gap between your governance decision date and your resource reallocation date is your mean time to pivot. It is almost certainly longer than you think.



Capital is frozen

Mean time to pivot is ultimately a capital problem. Even when leaders understand what needs to change and trust the data to act, they frequently cannot reallocate fast enough. The funding model is the most under appreciated structural barrier – and for most organizations, it has not changed in decades.

The governance tax on decision velocity

The mechanism by which slow decisions happen is governance design. 45% of organizations manage cross-functional dependencies during periodic planning cycles – meaning that when a dependency surfaces between cycles, the decision on how to resolve it can take weeks to finalize by design. This is not a failure of urgency. It is how the process was built.

The downstream effect is visible in decision latency data: when performance data indicates a problem, only 5% of organizations decide the same day, 27% within a week, 44% within a month, and 20% take a full quarter. Only 44% rate their governance cadence as effective. The gap is not a lack of governance – it is governance designed to minimize risk rather than accelerate decisions.

The funding model was built for a slower world

The governance tax on decision velocity exists because of how funding is structured underneath it. 42% of organizations operate on fixed annual budgeting. A further 36% use project-based funding tied to deliverables. Together, these describe nearly 80% of enterprises – and both were designed for a world where strategy changed once a year. That environment no longer exists.

Only 8% of organizations have moved to dynamic, outcome-based funding – the model the other 92% acknowledge they need. The reasons are structural, not motivational: CapEx/OpEx accounting rules constrain how quickly budget can be moved between categories; approval requirements add weeks to reallocation decisions; and the reputational cost of canceling something publicly committed creates powerful organizational inertia. The phrase that recurs across interviews: ‘We can’t be seen to cancel this.’ Strategic signals arrive continuously. The gap between signal cadence and reallocation cadence is where strategy dies.

“

Bureaucracy always slows things down... if you have to go through many levels of approvals before you can proceed on something, it’s going to take you some time.

VP Product Operations

Large Enterprise (10k-20k employees), Software

30%

find it easy to reallocate funds across initiatives

36%

identify budget overruns early enough to course-correct

8%

use dynamic, outcome-based funding models

The hidden capital problem: AI and compute

In a Human + AI enterprise, capital lock extends beyond headcount & budget. It extends to compute. Leaders cannot see which strategic bets are consuming AI compute and at what cost. They cannot identify which AI workflows are producing value and which are running on autopilot. They have no mechanism to ask: “What could we stop to free agents and compute for this new priority?”

The data exists – every major cloud platform produces granular cost data. The problem is that compute spend sits in IT infrastructure dashboards, entirely disconnected from the portfolio layer where bets are set. When an organization needs to pivot, three separate reallocation challenges are managed in three separate systems by three separate teams: headcount (HR), project budget (Finance), and compute/AI spend (IT or Engineering). These rarely move together – and no single leader has visibility across all three.

The ceiling on AI ROI is not model capability. It is capital governance. The organizations that solve this – that make compute spend as visible and attributable at the portfolio layer as headcount – will have a structural advantage in directing AI investment toward what actually matters.

What percentage of your AI and cloud compute spend can you attribute to a specific strategic bet with a defined outcome? For most organizations, the answer is near zero. This is the gap that will need an answer when the board asks how the AI transformation is paying off.



What this means if you're a CFO or COO

Your AI spend is probably not visible at the portfolio level. The question to ask your CIO or Head of Engineering: which AI workloads can we attribute to a specific strategic bet, and which are running without that attribution? The answer will tell you how much of your AI investment is strategically directed versus accumulated over time with no clear owner. That number matters more than the total AI spend figure, because it's the denominator against which you will eventually need to show ROI.

The AI paradox

AI does not create the structural barriers described in the preceding sections. It makes them more expensive.

Every enterprise in our study has made visible investments in AI. Most leaders can point to productivity gains. What they cannot point to is organizational speed – the ability to move faster as an enterprise when strategy demands a pivot. The reason is structural, and it follows directly from what capital lock and governance design already produce.

And without that structure, AI doesn't just fail to make you faster. It amplifies poor decisions and misaligned execution – helping parts of the organization move quickly in different directions, with more friction when they collide.

AI is a connectivity multiplier, not a strategy substitute

The single biggest determinant of whether AI produces strategic value is not which AI you deploy. It is whether the systems where strategy is set and the systems where work happens are connected.

This factor – domain connectivity – predicts AI effectiveness more strongly than any other variable we tested.

Regression analysis, N=308

Without a connected model, AI cannot answer the coordination questions leaders most need answered: “Which work is moving this bet?” “Who – including which agents – is doing it?” “Which systems and costs are at stake if we change direction?” Most organizations cannot answer these today, regardless of which AI tools they have deployed. Deploying more sophisticated AI on top of disconnected systems does not close this gap. It widens it.

The trust gap is wider than the usage gap

Only 34% of organizations regularly use AI in strategic contexts. But the more important finding is not the usage gap – it is the trust gap. AI trust does not improve meaningfully with AI usage alone. What improves trust is data connectivity: when leaders can see that AI is reasoning over an accurate, connected model of how their organization actually works, trust follows.

44%

agree AI improves
decision speed

41%

agree AI improves
decision quality

Only 26%

trust AI insights for
strategic decisions

The gap between perceived value and actual trust is not irrational. Leaders have seen AI-generated recommendations that could not be interrogated, dashboards that contradicted each other, and analysis built on data they could not verify. Trust is earned through transparency about what data AI is reasoning over – and that transparency is structurally unavailable in a fragmented operating model.

Where AI actually adds value

Today, only 2% of organizations use AI-driven insights to predict dependency risks – the baseline from which the opportunity needs to be measured. Leaders consistently identify three areas where AI could move beyond its current basic state: predictive risk detection – surfacing signals before they become crises; what-if scenario modeling before committing resources; and dependency and capacity visibility across teams and agent workloads. Each requires the same prerequisite: a connected, governed model of how work, people, systems, and priorities relate to one another. That model is what the preceding sections describe – and it is what most organizations have not yet built.

This observation recurs across the qualitative interviews in different forms. The bottleneck is not AI capability. It is organizational context – the structured, connected, continuously updated model of how the organization actually works. Building that model is the work that precedes AI delivering on its strategic promise.

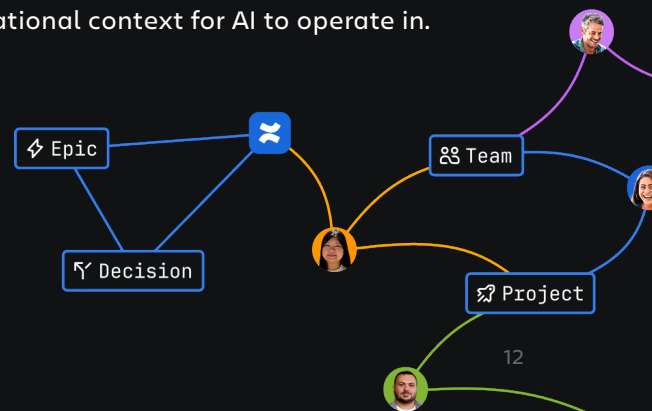
“To make AI really help in the strategic world, you need to basically deploy the model internally and give it that context, which is a very difficult job... The contextual data doesn’t exist. Contextual data is sitting in people’s heads.

VP Product Operations

Large Enterprise (10k-20k employees), Software

What this means if you’re a CTO or CIO

The question your organization is probably asking – which AI tools should we standardize on? – is the wrong first question. The right question is: what is the connected data model that AI will reason over, who owns it, and how confident are you that it accurately reflects how the organization actually works today? AI tool selection is downstream of that answer. The organizations getting the most from AI have not necessarily selected better models. They have built better organizational context for AI to operate in.



What winning looks like

Every enterprise is running an implicit market for capacity. Strategic priorities represent demand. Execution resources – teams, AI workflows, budget – represent supply. In most organizations, supply is locked in last year’s commitments and the mechanism to move it is slow, political, and fixed for rigid cycles. Winning organizations are better at running this market – sensing where demand has shifted and moving supply to meet it faster than their competitors can.

The structural practices that enable this are not new. SAFe, Lean Portfolio Management, and PMI standards have described them for over a decade. What this research adds is evidence of how wide the implementation gap remains – and two specific places where the standard playbook stops short of what the AI era requires.

What the established playbook gets right

Strategy in the same system as work

In most enterprises, strategy exists in documents while work lives in delivery tools – managed by different leaders, reviewed on different cadences. The result: strategic intent has no structural mechanism to direct daily work. Lean Portfolio Management has described this linkage since 2015. Our data shows 11% of organizations have achieved it. What remains rare is the connected organizational model that makes this linkage live rather than periodic.

Governance for decision velocity, not control

Effective governance and review frequency account for roughly a quarter of the variation in organizational adaptability. But the kind of governance matters. The most advanced organizations are moving beyond cadence optimization entirely – encoding decision rights into their operating model so routine reallocations happen without requiring a governance meeting to initiate. It is worth being precise: adaptability here is a perception variable, not an operational speed variable. The null finding in the next section shows that increasing how often you re-rank priorities has no effect on how fast the organization actually moves. These findings describe different things and are not in conflict.

Fund reallocation as an operational capability

Dynamic outcome-based funding – funding a hypothesis, validating it, then funding the next step – is standard doctrine in agile portfolio management. Our data shows 8% of organizations actually operate this way, constrained by CapEx/OpEx rules, approval structures, and the reputational cost of canceling something publicly committed.



Where the standard playbook stops short

The compute visibility gap

The standard playbook treats capacity as a human and budget question. That model no longer reflects reality. Capacity now comprises three resource types governed by three different leadership teams: headcount (HR), project budget (Finance), and compute and AI spend (IT or Engineering). Leaders cannot answer the question that matters most when strategy shifts: “What could we stop or slow to free agents and compute for this new priority?” The ceiling on AI ROI is not model capability. It is capital governance.

The null finding: planning cadence is not the answer

Here is the finding that should most directly challenge your current Strategic Portfolio Management investment thesis: making your planning process more frequent does not make your organization faster.

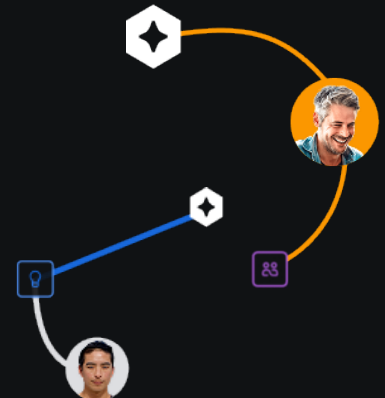
How frequently an organization re-ranks its priorities based on performance data has effectively zero predictive relationship with decision confidence, dependency resolution speed, or AI effectiveness.

Regression analysis, N=308

When a strategy changes, the decision propagates through governance and gets documented in planning systems. What it does not do – in the 89% of organizations without structural connectivity – is change anything in the execution layer: the tools where work lives, the funding models that govern capital, the systems determining what agents and compute are doing. The plan updates. The organization continues on its original trajectory. A system that helps an organization revise priorities more frequently, without connecting those revisions to the execution layer, is not improving mean time to pivot. It is improving mean time to update a document.

What this means if you're a CEO or COO

The organizations that close this gap are not commissioning another planning transformation. They are making two specific ownership decisions: who owns the organizational knowledge graph across HR, Finance, and Engineering; and who has authority to convene cross-functional reallocation when a strategic signal demands it. If neither question has a clear answer in your organization, that is the gap. Mean time to pivot will not improve until both do.



From fragmented to living: the operating model shift

Strategy as a Living Operating Model

The organizations that close the Strategy Execution Gap do not just plan better or deploy more AI. They run a living operating model:

- Strategy is treated as a **3-5-year hypothesis**, not a fixed plan.
- It is surrounded by a continuous loop: **Signals** → **Decisions** → **Reallocation** → **Updated Strategy**.
- It runs on a **shared, governed model** of priorities, people, agents, work, systems, and money.
- In this model, every enterprise is effectively running a market for capacity:
 - **Signals = demand** on the enterprise.
 - **Reallocation = moving supply** (humans, agents, compute) to the highest-value demand.
 - Humans set direction and guardrails; **AI augments sensing, modeling, and recommendation** – but only when it can operate over a reliable model of how the business actually works.

TODAY Manual & Periodic	TOMORROW Agentic & Continuous
Humans interpret and propagate decisions	Agents sense, recommend, and execute within guardrails
Manual strategic updates in documents	Strategy encoded in a living system that scales a connected operating model
Annual / quarterly plans; change seen as failure	Continuous hypothesis; plans adjust as signals change
Risk surfaced after impact	Risk surfaced and acted on in real time
Capacity = headcount + budget	Capacity = humans + agents + compute


Implications for Strategic Portfolio Management

The AI era is redefining what Strategic Portfolio Management needs to do – not because Strategic Portfolio Management has failed, but because the conditions it was designed for have changed faster than any framework could anticipate. The traditional mandate – govern the portfolio, track against plan, report upward – was built for a world where strategy changed annually and capacity meant headcount. That definition served organizations well for decades. It is now the bottleneck.

Highly adaptive organizations that can pivot quickly are not running a more sophisticated version of portfolio governance. They are maintaining a live, connected model of how capacity and strategy relate to one another, and using that model to make reallocation decisions in days rather than quarters. This demands a fundamentally different discipline than the one most organizations are practicing under the Strategic Portfolio Management banner today.

What Strategic Portfolio Management is becoming

The best description of where high-performing Strategic Portfolio Management disciplines are heading is not portfolio management. It is closer to managing an operating system – maintaining a live model of available resources, mediating access across competing processes, enforcing rules about priority, and surfacing information that allows the system to adapt and encode resilience. Portfolio governance asks: are we doing the right initiatives? Operating system management asks something harder: does the organization have a live, accurate model of its own capacity – across humans, agents, and compute – and can it reallocate quickly when signals demand it?



The question is no longer whether the portfolio is on track. It is whether the organization can see its own capacity clearly enough to move it.

The emerging Strategic Portfolio Management mandate

The traditional Strategic Portfolio Management mandate

Govern the portfolio. Track status and budget against plan. Report upward. Ensure delivery discipline. Identify and escalate risks.



The emerging Strategic Portfolio Management mandate

Maintain the organizational knowledge graph. Govern the blended capacity stack. Enable AI to reason over a live model of the organization. Shorten mean time to pivot.

Strategic Portfolio Management needs to own the organizational knowledge graph: a live, machine-readable model of how strategic priorities connect to work, how work connects to teams and agents, how teams and agents connect to systems and compute, and how all of it connects to the funding and outcomes the organization is accountable for. This is the infrastructure that makes both human judgment and AI reasoning possible at the organizational level.

Compute governance is the most urgent new frontier. AI and cloud compute spend is now a material line item for most enterprises – and almost entirely invisible at the strategy level. No one at the portfolio layer tracks which AI workloads are running against which strategic bets or whether they are producing value proportional to cost. This is not an IT problem. Capital allocation is exactly what Strategic Portfolio Management exists to govern.

No major Strategic Portfolio Management framework is currently addressing compute visibility as a portfolio governance problem. The organizations that move first will define what the discipline looks like for the next decade.



Governance itself needs to be redesigned – not around meeting rhythms, but around decision rights architecture: who can authorize what level of reallocation, under what conditions, without requiring a governance meeting to initiate. Routine reallocations happen continuously; genuine exceptions escalate to human judgment with the context required to decide quickly. AI operating over the organizational knowledge graph can continuously monitor the gap between strategic intent and current capacity commitments, surface conflicts before they reach governance, and recommend specific actions within defined guardrails.

What this means if you're a Chief Strategy Officer

The organizational knowledge graph is probably something you have been trying to build informally for years – pulling data from HR, Finance, and engineering to check whether execution is aligned with the strategy you set. The argument this report makes is that work should be structural, not heroic. In your organization, who owns the knowledge graph? If the answer is nobody, or five functions each owning a piece of it, that is the gap. The CSO often has both the strategic authority and the cross-functional reach to claim this mandate.

What this means if you're a CFO, COO, or Chief Transformation Officer

The Strategic Portfolio Management function described in this section is different from a project tracking PMO. It owns the organizational knowledge graph, governs AI and compute spend as a capital variable, and has authority to convene cross-functional reallocation when strategy shifts. If your organization wants to close the gap between strategic intent and organizational movement – and show measurable AI ROI to your board – this is the function that enables it.

Why existing Strategic Portfolio Management tools don't close the gap

Most Strategic Portfolio Management tools were built from the strategy layer down: strategic priorities as primary data objects, execution data pulled in through integrations. This architecture produces a portfolio view – an accurate picture of what was planned and how it is tracking. A portfolio view is not an organizational knowledge graph. The difference is not cosmetic: a portfolio view is assembled periodically; a knowledge graph is lived in. The entities in a portfolio view are planning constructs. The entities in a knowledge graph are operational realities – work items, dependencies, agent workloads, compute costs, team capacity.

Platforms built from the execution layer up have a different starting position. The knowledge graph is not something they need to build. It is something they need to surface, structure, and connect to the strategic layer – a materially more solvable problem.

Where Atlassian fits

Most Strategic Portfolio Management vendors built their products to serve the traditional mandate. These products are well-designed for the problem they were built to solve. That problem is no longer the binding constraint.



The binding constraint is the absence of the organizational knowledge graph. Without it, AI cannot reason about the organization accurately, the blended capacity stack cannot be governed, and mean time to pivot remains a function of how fast people can manually reconcile disconnected systems.

Atlassian's platform is already the native home of work. Atlassian did not start with a planning tool and add execution connectivity. It started with the execution layer – the tickets, the dependencies, the blockers, the velocity, the decisions being made at the team level every day – and has been accumulating the relationships, history, and context that constitute the organizational knowledge graph for twenty years across 300,000+ organizations running their execution inside Atlassian's ecosystem. Building the knowledge graph from the strategy layer down requires organizations to migrate their execution data into a new system. Building it from the execution layer up means it already exists.

Strategy

Strategy Collection models strategic bets, portfolios, and priorities as first-class objects in the same environment as execution work – the structural linkage that only 11% of organizations currently have.

Teamwork

The Teamwork Graph encodes the relationships between priorities, people, agents, work, systems, and history – the organizational knowledge graph made operational, enabling leaders to see headcount, project budget, and compute spend as a unified capacity stack. It's a structural advantage, encoding the real relationships, dependencies, and history required for true enterprise context.



Rovo operates over that connected model to surface drift, model reallocation options, highlight risks, and recommend specific actions within governance guardrails – AI doing the continuous monitoring work that currently consumes the majority of Strategic Portfolio Management practitioners' time.

CLOSING

The competitive advantage of structural adaptability

The organizations that build the connected operating model in the next two to three years will not simply be more efficient. They will operate in a structurally different competitive environment – and the advantage will compound.

When strategy, work, and capacity exist in a single connected model, strategy becomes a live steering mechanism rather than a planning artifact. Leadership can ask – and receive a data-driven answer to – whether the work the organization is doing today is still aligned with the bet they made six months ago. They can model the cost of a pivot before committing, and execute it in days rather than quarters. In markets where strategic windows open and close in months, the organization that reallocates in days rather than quarters captures advantages that are not easily replicated: more bets placed faster, resources concentrated behind what is working before competitors have finished their planning cycle.

The financial translation is direct: faster reallocation means less capital locked in weakened initiatives, AI investments attributable to outcomes rather than accumulated as unaccountable infrastructure spend, and the ability to show a board – with data, not narrative – that capacity is aligned with the highest-value strategic bets.


For investors and boards evaluating enterprise resilience, mean time to pivot is a proxy for something strategically important: whether leadership can actually execute on a strategy change, or whether the operating model will absorb that change and continue on its original trajectory regardless. The organizations that are most adaptable are not just better managed. They are structurally more valuable – faster to respond to disruption, more capable of attracting talent and capital, and more able to defend their AI investment thesis with evidence rather than aspiration.

Research Methodology

This report draws on a mixed-methods study conducted in Q1 2026 with 308 senior leaders across strategy, portfolio, technology, and delivery roles in enterprise-scale organizations. The quantitative survey was hosted and administered by GLG using a verified professional panel, fielded February 2–20, 2026, with an average completion time of 25–35 minutes. Qualitative depth was added through 60-minute interviews with 15 senior leaders, conducted in partnership with GLG. At the 95% confidence level, the margin of error for survey-level findings is approximately ± 5.6 percentage points (calculated from $N=308$ using standard statistical methodology).

Regression analysis was conducted to identify structural factors predicting organizational adaptability, decision velocity, AI effectiveness, and dependency management. Findings describe correlation and predictive association – not direct causation. Key results: governance effectiveness and review frequency account for 24% of adaptability variance (rising to 28% with AI use and domain connectivity added); domain connectivity accounts for 18% of AI effectiveness variance and 7% of decision confidence variance; system integration accounts for 17% of dependency resolution speed variance and 21% of automatic risk mitigation variance. The frequency of priority re-ranking based on performance data has effectively zero predictive relationship with any of these outcomes – a finding that held across regression models.

- **Sample:** N=308
- **Roles:** Strategy & Transformation (53%), IT/Technology (20%), Executive Leadership (20%), PMO (3%), Operations (3%)
- **Seniority:** Director (35%), C-Suite (27%), VP (23%), EVP/SVP (15%)
- **Regions:** North America (149), EMEA (104), APAC (55)
- **Industries:** Software/Technology (22%), Banking (20%), Manufacturing (17%), Healthcare (16%), Retail (12%), and others
- **Company size (FTE):**
 - Mid-Market - 1,000–5,000 (22%)
 - Enterprise - 5,000–10,000 (16%)
 - Large Enterprise - 10,000–25,000 (19%)
 - Very Large Enterprise - 25,000–50,000 (14%)
 - Mega Enterprise - 50,000+ (29%)
- **Fielding:** February 2–20, 2026; ~25–35 minute survey



About Atlassian

Atlassian unleashes the potential of every team. A recognized leader in software development, work management, and enterprise service management software, Atlassian enables enterprises to connect their business and technology teams with an AI-powered system of work that unlocks productivity at scale. Atlassian's collaboration software powers over 80% of the Fortune 500 and 350,000+ customers worldwide - including NASA, Rivian, Deutsche Bank, United Airlines, and Bosch - who rely on our solutions to drive work forward. Learn more about Atlassian at www.atlassian.com