EMA™
IT & DATA MANAGEMENT RESEARCH,
INDUSTRY ANALYSIS & CONSULTING

Prepared for:

beta systems

From Outcomes to Authority:
Defining the Enterprise Control Plane

July 2026 EMA Research Report
By **Dan Twing**, President and COO
Automation Information

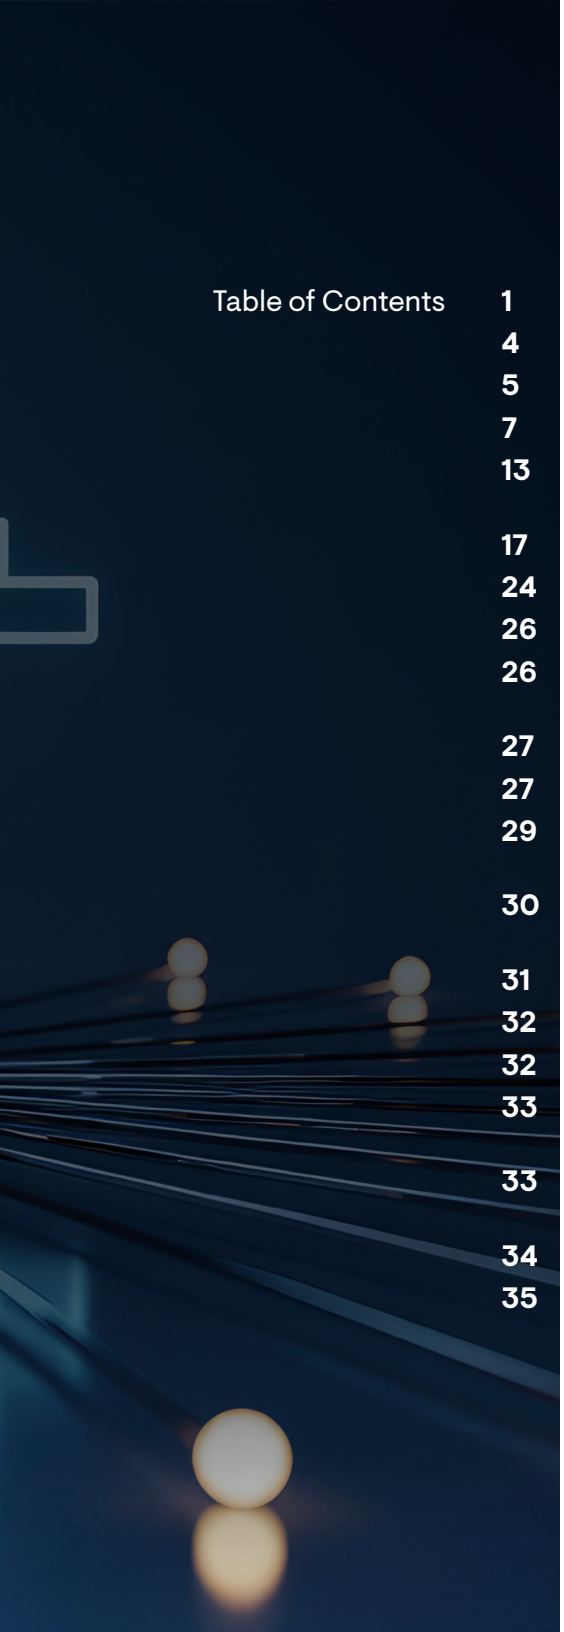

Table of Contents

Executive Summary

Enterprise automation is already federated. Organizations rely on multiple orchestration platforms operating simultaneously across different operational domains and they are not converging toward a single authority. The strategic challenge this creates is not a technology architecture problem. It is a governance problem: how do enterprises maintain accountability, visibility, and outcome confidence when operational authority is increasingly distributed across automation platforms, orchestration systems, observability environments, and AI-enabled technologies?

This report argues that the answer is the enterprise control plane – and that AI, rather than creating this requirement, has made it impossible to ignore.

EMA's survey findings confirm this directly: organizations view orchestration as a federated operational capability rather than a function any single platform category owns. Workload automation platforms, service management systems, cloud automation frameworks, ERP environments, CI/CD pipelines, observability platforms, business process automation tools, and emerging AI systems all participate in coordinating work. This is not a transitional state. It is the durable structure of enterprise operations – and the coordination challenge it creates grows more complex, not less, as each domain gains greater autonomy.

AI does not eliminate this federated reality. It intensifies it. As AI embeds throughout enterprise technology stacks, domain-specific platforms become increasingly autonomous while business outcomes increasingly depend on interactions that cross application, infrastructure, operational, and organizational boundaries. The strategic question is no longer which platform owns orchestration: it is how enterprises coordinate across the orchestration platforms they already possess.

This shift extends beyond technology architecture. For decades, enterprises governed operational authority through organizational structures, management hierarchies, policies, approval processes, accountability mechanisms, and organizational culture. These disciplines allowed people to coordinate actions toward common business outcomes. As operational action increasingly moves into automation platforms, orchestration systems, applications, observability environments, and AI-enabled technologies, enterprises must increasingly extend those same governance disciplines into the execution environment itself.

AI is accelerating the migration of operational authority and action into digital systems. The resulting challenge is not simply managing automation or AI. It is maintaining governance, accountability, visibility, and outcome confidence as authority and decision-making become increasingly distributed across human and digital actors.

The governance model for this environment remains unsettled. Organizations designed their existing governance approaches for deterministic automation operating within well-defined ownership structures. AI introduces adaptive decision-making, autonomous actions, and cross-domain coordination that those models were not built to handle. No single governance approach has emerged as dominant. What the research does find consistently is that organizations are treating autonomy as operational authority that must be earned – granted incrementally based on demonstrated reliability, compliance with governance policies, and outcome stability – not extended based on vendor assurances or broad organizational confidence.

Historically, organizations focused on validating automation before deployment and monitoring execution after deployment. AI changes this model by introducing reasoning into runtime operations. Successful execution no longer guarantees successful outcomes because interpretation now occurs during execution, increasing the need for operational assurance. Organizations increasingly require visibility not only into what actions occurred, but also why systems made those decisions, the information that influenced them, and whether intended business outcomes were achieved.

Practitioner voices throughout this research reinforce the survey findings. Across industries and platform types, practitioners consistently describe the same requirement: autonomy without context, governance,

observability, and human oversight remains unacceptable for critical operational environments. Organizations are not rejecting AI. They are calibrating it – granting operational authority incrementally, in proportion to demonstrated reliability, with governance and observability in place before expanding scope. When asked how they intend to expand AI authority, organizations cited accuracy and consistency of recommendations, demonstrated stability over time, and compliance with existing governance policies as the deciding criteria.

The research clarifies why this matters. Approximately 30% of respondents encounter incorrect or problematic AI outcomes frequently or very frequently. More than three-quarters of organizations that have deployed AI automation require human correction or rollback of AI-driven actions at some point. Successful process execution no longer guarantees correct outcomes. As reasoning becomes part of execution, assurance must become operational as well.

The research also identifies context degradation as a primary cause of AI outcome failures. As decisions cross application, infrastructure, operational, and organizational boundaries, critical context is often lost, reducing the reliability of recommendations and actions. Among organizations that have rejected AI recommendations, insufficient context ranks among the most frequently cited causes. The challenge is preserving the operational context and state awareness necessary to support reliable decisions as workflows move across increasingly complex and federated environments.

This finding points toward a broader reality emerging across enterprise operations. Intelligence alone is not sufficient. Intelligence without sufficient state awareness cannot be trusted to make consequential decisions. As AI assumes greater operational authority, the limiting factor increasingly becomes visibility into current conditions, dependencies, governance constraints, and business context. The challenge is no longer simply making systems smarter. The challenge is ensuring they possess sufficient state awareness to support trusted action.

Deterministic automation has always required state awareness. AI amplifies this dependency by expanding the scope of relevant state that must be understood, coordinated, and preserved, while increasing the consequences when that state is incomplete, fragmented, or unavailable.

These trends point toward the emergence of what EMA describes as the enterprise control plane. It is not another orchestration platform or a renamed version of an existing category. Instead, it represents a larger operational frame: the coordination layer that sits above specialized execution systems and makes federated autonomy governable.

EMA defines the enterprise control plane as the operational infrastructure through which enterprises coordinate governance, authority, accountability, visibility, decision transparency, context preservation, and outcome confidence across a federated ecosystem of human and digital actors.

Its purpose is not to eliminate specialization; it is to make specialized systems operate coherently together. Furthermore, its purpose is not to centralize decisions; it is to ensure that distributed decisions remain visible, governed, and accountable.

The strategic question for enterprise automation is no longer how to automate more. It is whether the coordination infrastructure exists to govern what has already been set in motion. As enterprise operations become increasingly federated, business outcomes depend upon decisions and actions occurring across multiple systems, domains, platforms, and actors. Maintaining visibility, governance, accountability, and confidence increasingly requires coordinating the state transitions those decisions create. Historically, enterprises governed operational authority through management hierarchies, policies, approval processes, and accountability structures. As operational authority moves into digital systems, those governance disciplines require an operational home. The enterprise control plane is that home – not a new orchestration platform, but the coordination layer that makes federated autonomy governable.

Introduction

The Road to the Enterprise Control Plane

Enterprise automation expanded across every operational domain – applications, infrastructure, data platforms, service management, observability, and increasingly AI-enabled technologies – without converging on a single platform. The result is a federated operational environment in which the challenge has shifted from automating individual activities to coordinating execution across otherwise disconnected platforms, teams, and processes.

As these platforms matured, execution responsibilities shifted in both directions: some activities moved into domain-specific platforms with the operational context to manage them directly, while other commonly required capabilities were absorbed into broader orchestration platforms. Regardless of where execution occurred, the need for technologies capable of coordinating activities across systems, teams, and operational domains continued to expand – and workload automation and orchestration platforms evolved to fill that role.

A note on terminology: the category of platforms that provides cross-domain automation coordination has carried several names over the years: workload automation, workload automation and orchestration, service orchestration and automation platforms (SOAP), and enterprise automation orchestration, among others. After the SOAP label was introduced, EMA tested it and competing names with practitioners; none achieved dominant recognition, but all describe the same functional category and largely the same vendor set. Throughout this report, EMA uses WLA as shorthand for this entire category, except when a specific label is directly relevant to the point being made.

Generative AI and agentic systems intensify this challenge. AI introduces new opportunities for adaptive decision-making and autonomous action, but it does so within operational ecosystems already characterized by complex interdependencies – ecosystems that already require significant coordination to function.

These developments raise important questions for enterprise leaders: how automation is coordinated across increasingly diverse operational environments, how authority is managed when multiple platforms participate in execution decisions, and how organizations are adapting their operating models to manage increasingly interconnected automation ecosystems.

To better understand these questions, EMA conducted this research study to examine the evolving structures organizations use to coordinate automation, govern autonomous actions, and manage operational outcomes. Drawing upon responses from IT leaders, architects, automation practitioners, operations professionals, and executives across multiple industries and regions, this study explores how organizations coordinate automation today, how governance practices are evolving, how trust is established and maintained, and what these developments reveal about the future of enterprise operations.

This study argues that the emergence of the enterprise control plane is not primarily a response to artificial intelligence. It is a response to a broader shift in enterprise operations. As authority and action become increasingly distributed across applications, automation platforms, orchestration systems, observability environments, and AI-enabled technologies, enterprises require new mechanisms for coordinating governance, authority, accountability, and outcomes across an increasingly federated operational environment.

Historically, these responsibilities were managed primarily through organizational structures. Management hierarchies, approval chains, policies, culture, and accountability mechanisms governed how people coordinated actions to achieve business outcomes. As more operational authority moves into digital systems, enterprises increasingly require operational infrastructure capable of extending those same disciplines into the execution environment itself.

Throughout this report, EMA uses the term “enterprise control plane” as shorthand for the emerging operational infrastructure required to coordinate governance, authority, accountability, visibility, and outcomes across increasingly federated operational environments. The concept will be developed progressively throughout the sections that follow.

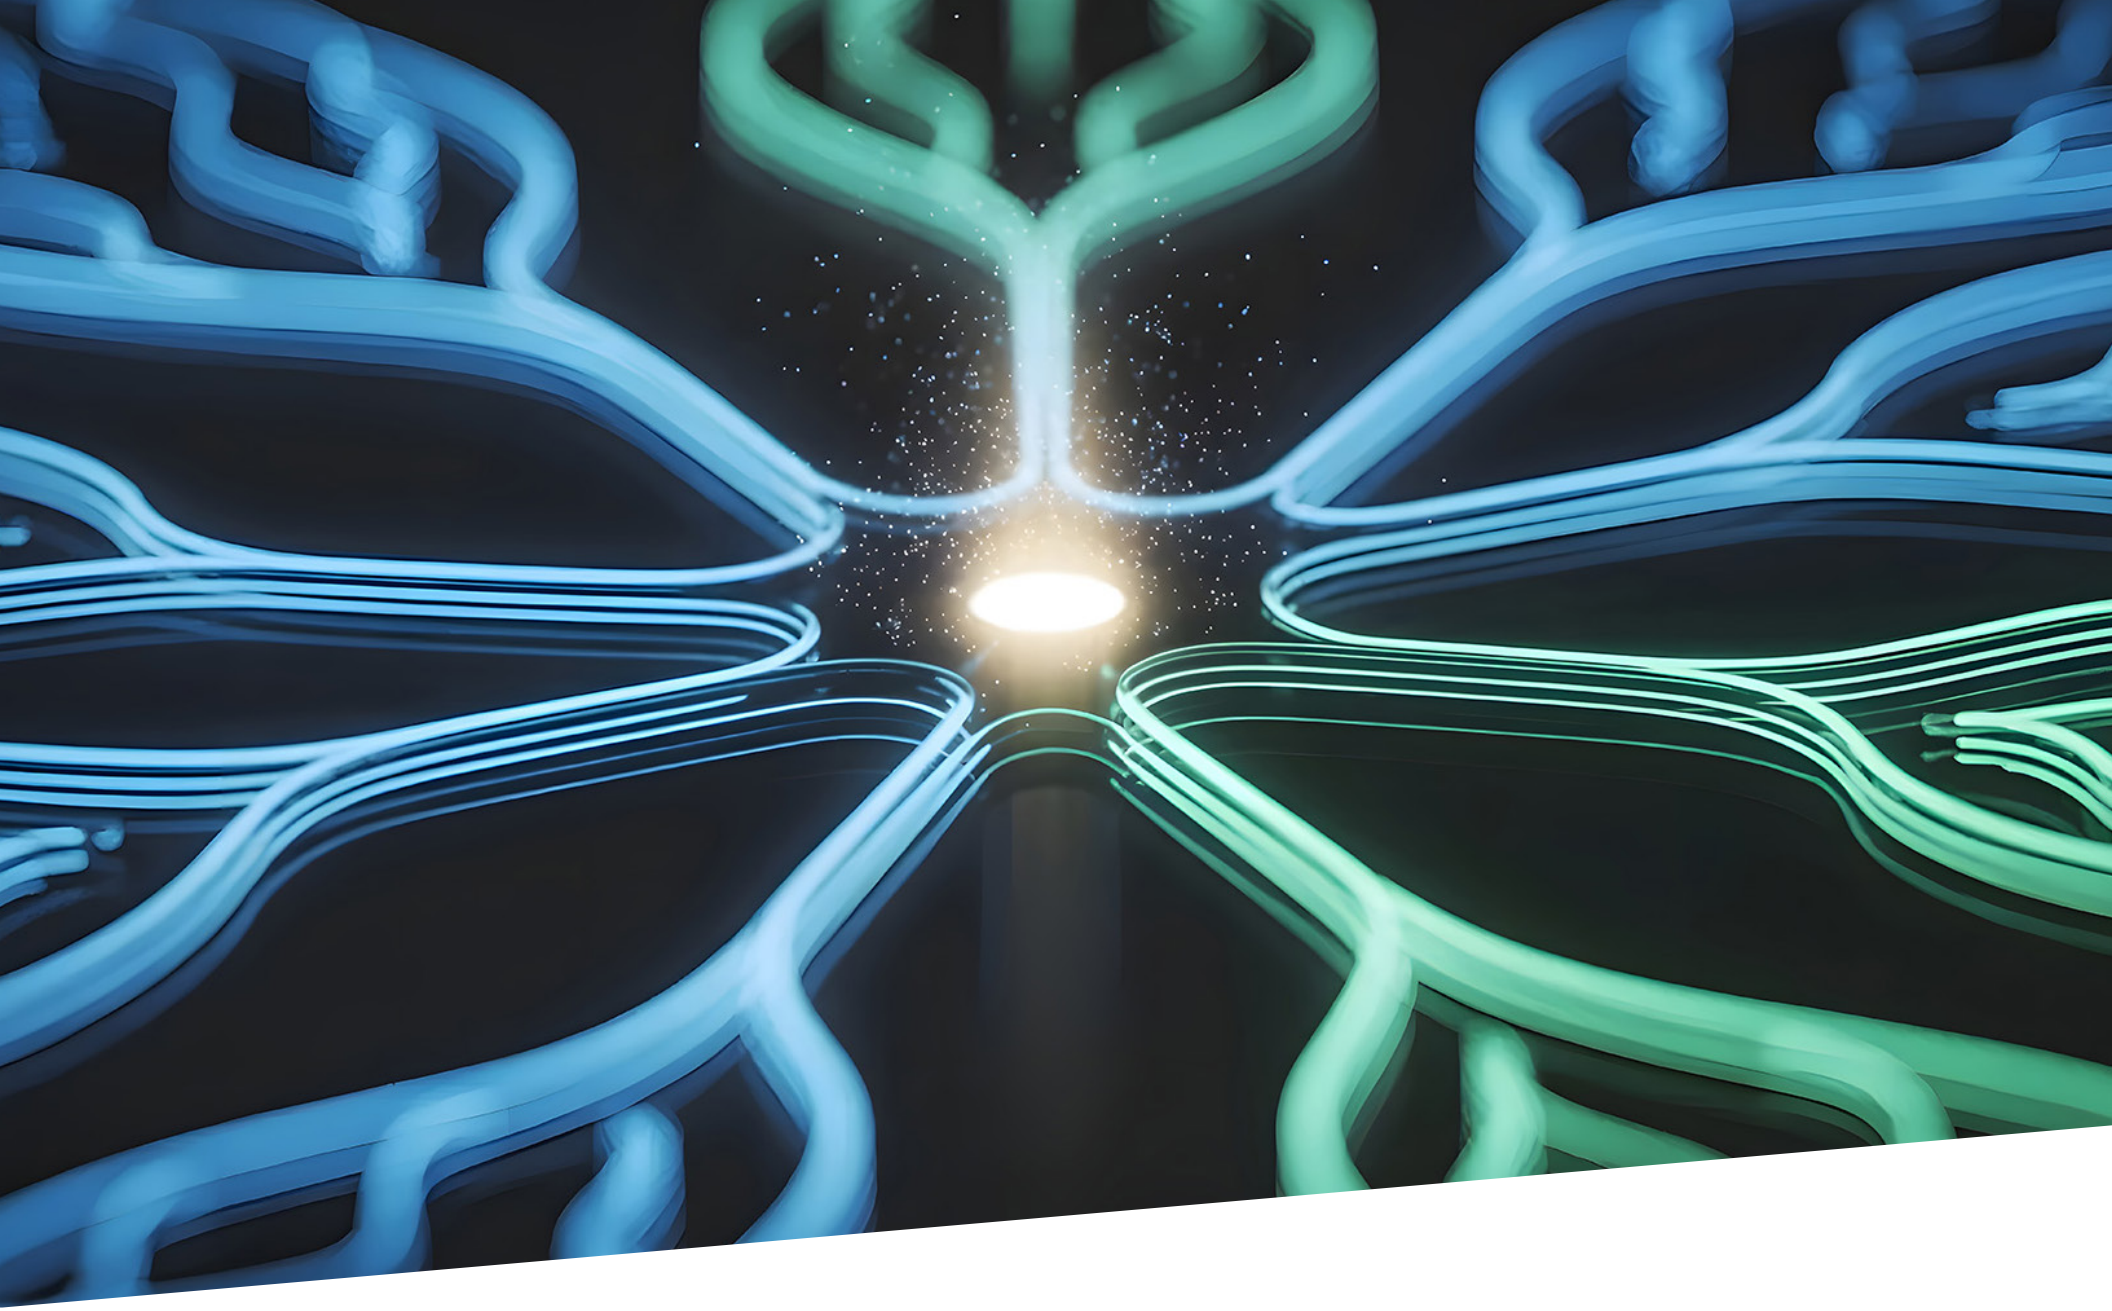

Enterprise Automation is Federating

The survey findings reveal what federated orchestration looks like in practice. The challenge is no longer simply automating work. It is coordinating automation that increasingly spans systems, teams, technologies, and organizational boundaries.

The first indication of this shift appears in the scope of modern automation itself. Nearly half of respondents reported that between 26% and 50% of their automation workflows span multiple domains, while more than one-third indicated that a majority of their automation crosses domain boundaries. Critical business outcomes increasingly depend upon activities coordinated across applications, infrastructure, cloud services, enterprise platforms, data environments, security systems, and AI-enabled technologies.

Approximately what percentage of your organization's automation workflows span multiple domains (e.g., IT operations, DevOps, data pipelines, SaaS applications, business processes)?

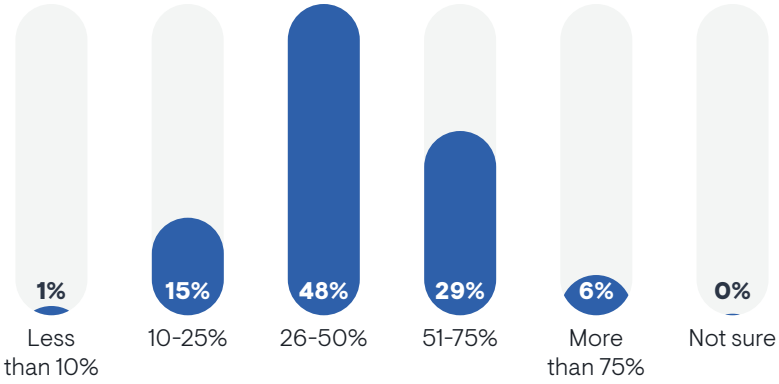

Figure 1.1: Percentage of Automation Workflows that Span Multiple Domains

Thinking about your organization's most critical cross-system workflow, which types of systems are coordinated within that workflow?

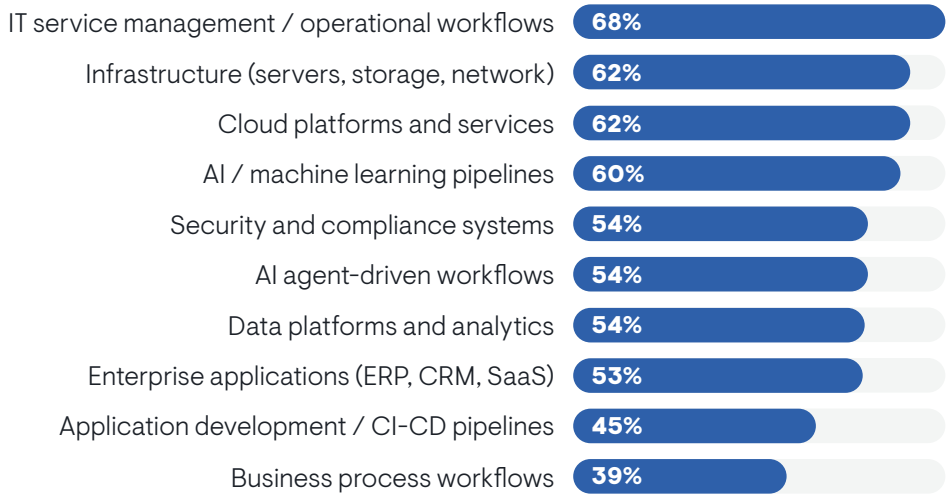

Figure 1.2: Systems Participating in Cross-Domain Workflows

As automation expands across domains, coordination authority becomes more distributed. Respondents reported responsibility spread across centralized orchestration platforms, shared governance models, AI-assisted environments, platform-specific control mechanisms, and other operational structures. No single coordination model emerged as dominant.

Up to this point, the findings describe a familiar story: automation spans multiple systems, coordination crosses organizational boundaries, and responsibility is distributed across teams and platforms. The more interesting finding emerges when organizations are asked a seemingly simple question:

Which platform actually serves as the primary orchestration system?

The answer depends on how the question is asked.

Where does primary control for coordinating workflows across systems reside in your organization? Include both automation and AI-driven workflows.

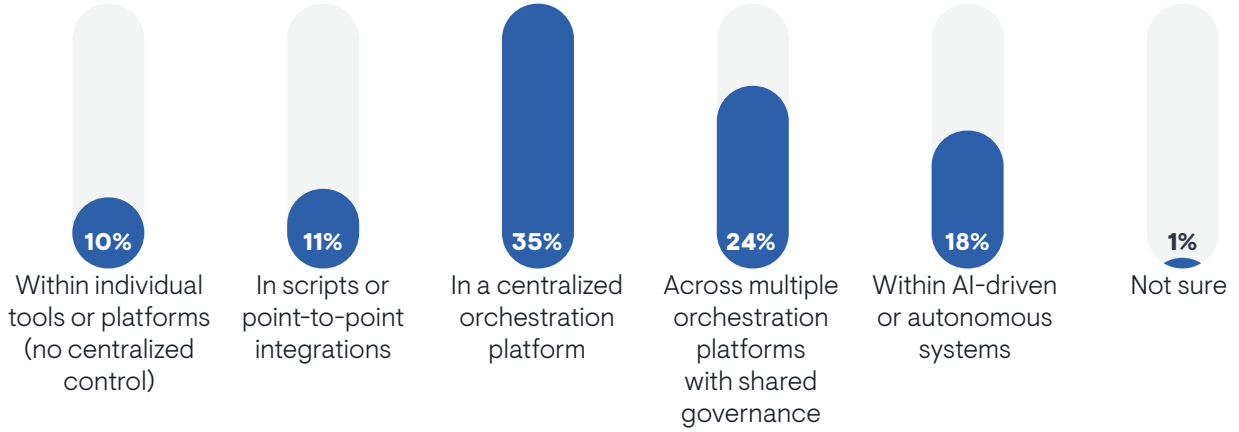

Figure 1.3: Where Workflow Coordination Authority Resides

Which type of system or platform currently plays the most significant role in coordinating and executing workflows across your environment?

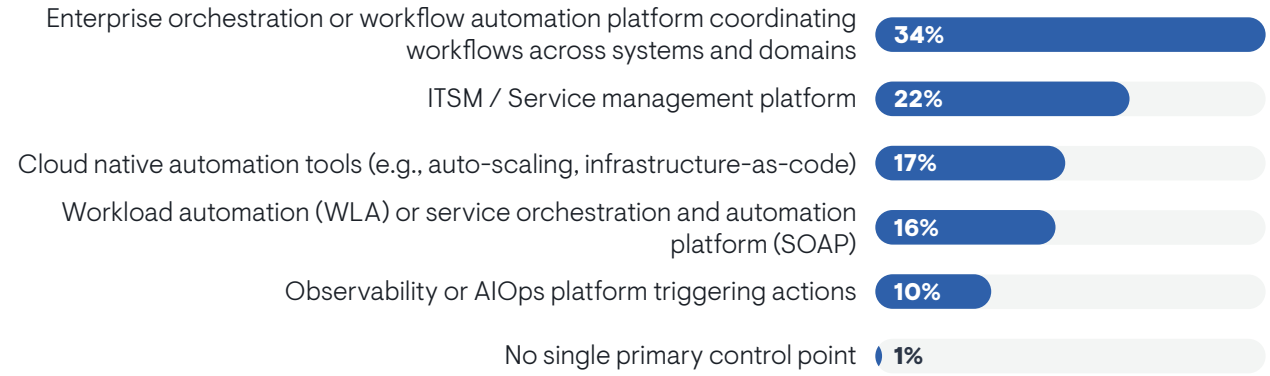

Figure 1.4: Platform Playing the Most Significant Coordination Role

When respondents selected from predefined orchestration categories, enterprise orchestration emerged as the largest response. At first glance, this appears to suggest a market converging around a common orchestration model.

However, that conclusion changes when respondents identify actual products.

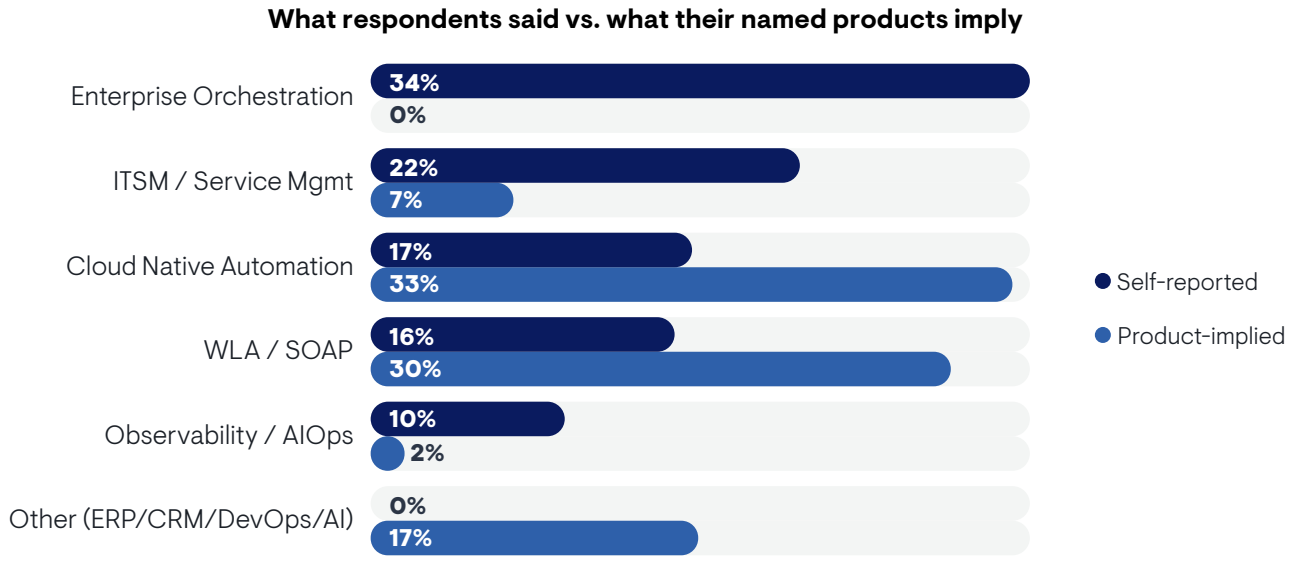

Figure 1.5: Location of Primary Workflow Coordination Control

The largest category in the self-classification exercise effectively disappears. Thirty-four percent of respondents selected enterprise orchestration when choosing from a category list. Yet, no corresponding product category emerges when respondents name the actual platforms coordinating work in their environments. Instead, respondents primarily identify cloud native platforms and workload automation systems.

This distinction is significant. Enterprise orchestration appears to function less as a product category than as a description of a role organizations believe must exist.

Open End – Platform Named for Coordinating & Executing Workflows By Category

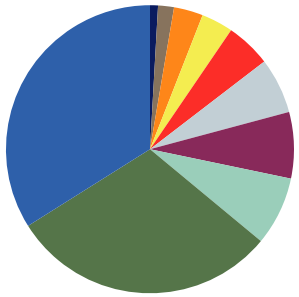

1%	Blank / Unrecognized
2%	Observability as Execution Authority (100% Dynatrace)
3%	Native AI / LLM Platform (OpenAI, deepseek, Zentra AI, Neuroother)
4%	DevOps / CI-CD (GitHub Actions, Atlassian, SoapUI, Puppet, Zapier, Red Hat, other)
5%	CRM / Sales Automation (100% Salesforce / Agentforce)
6%	Other / Unclassified (Symphony, Asana, Palo Alto, Crowdstrike, unknown vendors)
7%	ITSM / Service Management (ServiceNow 68%, TeamDynamix 12%, BMC Helix 4%, Resolve 4%, other 12%)
8%	ERP (Oracle 58%, SAP 42%)
30%	Workload Automation / WLA (Control-M/BMC 29%, Redwood 18%, IBM WLA/IWA 16%, JAMS 6%, AutoSys 5%, Stonebranch 4%, ANOW 3%, Automic 2%, ActiveBatch 1%, other 6%)
34%	Cloud Native / Hyperscaler (Microsoft 35%, AWS 19%, Google 14%, IBM Cloud/Watsonx 16%, Red Hat 4%, other 12%)

Figure 1.6: Product Categories Named for Workflow Coordination and Execution

The findings become even more revealing when operational dependence is examined.

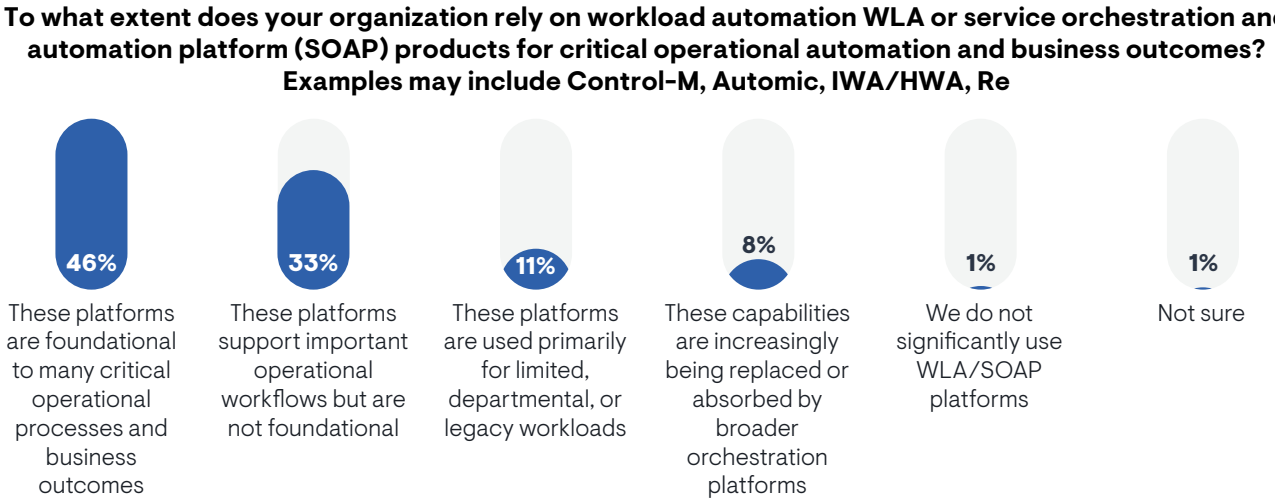

Figure 1.7: Workload Automation Platform Role and Persistence

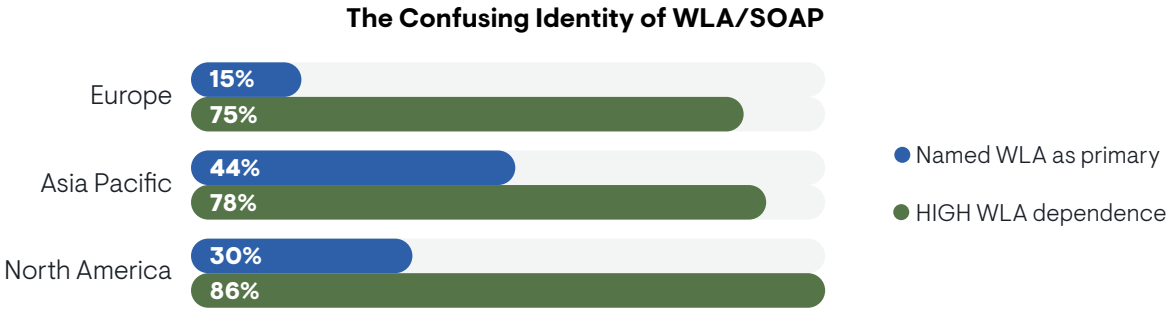

Figure 1.8: Regional Comparison – Categorized vs. Named Product

Across all three regions, organizations report substantially greater dependence on workload automation than primary-platform identification alone would suggest. Europe provides the clearest example. Only 15% of respondents identify workload automation as their primary orchestration platform, yet 75% report high or very high dependence on workload automation for critical operational outcomes.

These differences are not contradictions within the data. They are evidence of federation.

Different stakeholders view orchestration from different operational perspectives. Different platforms coordinate different domains. Different systems contribute to the same business outcomes.

The survey therefore reveals three realities operating simultaneously:

1. Organizations perceive orchestration through multiple operational lenses.
2. Organizations execute orchestration through multiple platform types.
3. Organizations depend upon multiple coordination systems simultaneously.

Taken together, the findings suggest that enterprise orchestration no longer operates through a single center of gravity. Coordination authority, execution responsibility, and operational dependence are already distributed across specialized domains.

Federation is not a future architecture.

Federation is the operating reality enterprises are managing today.

The Evolution of Enterprise Automation Coordination

EMA's research over the past decade traces this trajectory in detail: each wave of innovation delivered new value while simultaneously increasing the coordination requirements of the overall environment.

The findings in this report did not emerge in isolation. Over the past decade, EMA research repeatedly documented the expansion of automation from traditional scheduling into broader enterprise orchestration, cloud operations, data movement, AI-enabled decision-making, and increasingly distributed governance. Viewed collectively, these studies reveal a consistent pattern: as technology environments became more distributed, the need for coordination expanded with them. Table 2.1 summarizes that progression and its relationship to the emerging concept of the enterprise control plane.

The survey findings in Section 1 describe a federated orchestration environment. The broader historical question is how enterprise automation arrived at this point.

Period	Key EMA Research
2016 *Coordination Expands Beyond Scheduling*	**2016** Issues and Priorities in Modern Workload Automation *The Changing Landscape of IT Automation* **2018** The Shifting Role of Workload Automation *From Cost Center to Strategic Enabler* **2019** The Great Scheduler Migration *From Legacy to Modern Workload Automation* **2020** A Four-Stage Maturity Model for IT Automation *From Reactive to Autonomous IT Operations* **2021** Democratizing IT Automation in a Multi-Cloud World *The Rise of Self-Service Automation at Scale*
2022-2024 *Cross-Domain Coordination Becomes Central*	**2022** Global Workload Automation Market Size and Forecast 2022 to 2027 *Market Analysis and Industry Outlook* **2023** Workload Automation Transformation *Building the Foundation for Enterprise Agility* **2024** Data in Motion: Orchestrating File Transfers and Data Pipelines in the Cloud Era *Enabling Secure, Scalable, and Intelligent Data Movement*
2025 *Orchestration Becomes Strategic Infrastructure*	**2025** The Future of Workload Automation and Orchestration: Driving Digital Transformation with Orchestration and Generative AI *From Automation to Intelligent Orchestration*
Current Report *Governance Becomes Operational Infrastructure*	**Current Report** From Outcomes to Authority: Defining the Enterprise Control Plane *Establishing Governance as the Operational Backbone of the Autonomous Enterprise*

Figure 2.1 – EMA Research Evolution Toward the Enterprise Control Plane (2016–2026)

EMA's research suggests that the industry has not been solving a series of unrelated problems. Instead, it has repeatedly returned to the same fundamental challenge: coordinating work across increasingly complex environments. Each technology generation expanded the scope of coordination while preserving the underlying objective. The environment evolved. The coordination challenge remained.

Figure 2.2 illustrates this pattern as an orchestration helix.

The helix illustrates an important pattern visible throughout both EMA's historical research and the current survey findings. Enterprise automation repeatedly expanded into new domains while retaining responsibility for coordinating activity across those domains. New technologies changed what required coordination. They rarely reduced the need for coordination itself.

This pattern helps explain the findings presented in Section 1. Organizations increasingly rely on specialized orchestration platforms operating within cloud environments, service management systems, enterprise applications, observability platforms, development pipelines, and emerging AI ecosystems. Each domain continues to develop stronger orchestration capabilities tailored to its own requirements. Yet, business outcomes continue to span multiple domains simultaneously. The survey findings suggest that domain orchestration expands, but cross-domain coordination persists. As specialization increases, the need for cross-domain coordination increases with it.

EMA views this pattern as the persistent coordination problem. Every major technology wave expands enterprise capability while simultaneously increasing complexity, autonomy, and the amount of relevant state required to achieve trusted outcomes.

Viewed through this lens, the coordination challenge is fundamentally a state-awareness challenge. As workflows span domains, platforms, environments, organizations, and increasingly autonomous systems, the amount of relevant state expands while visibility into that state becomes increasingly fragmented. The coordination problem persists because the state landscape continues to expand.

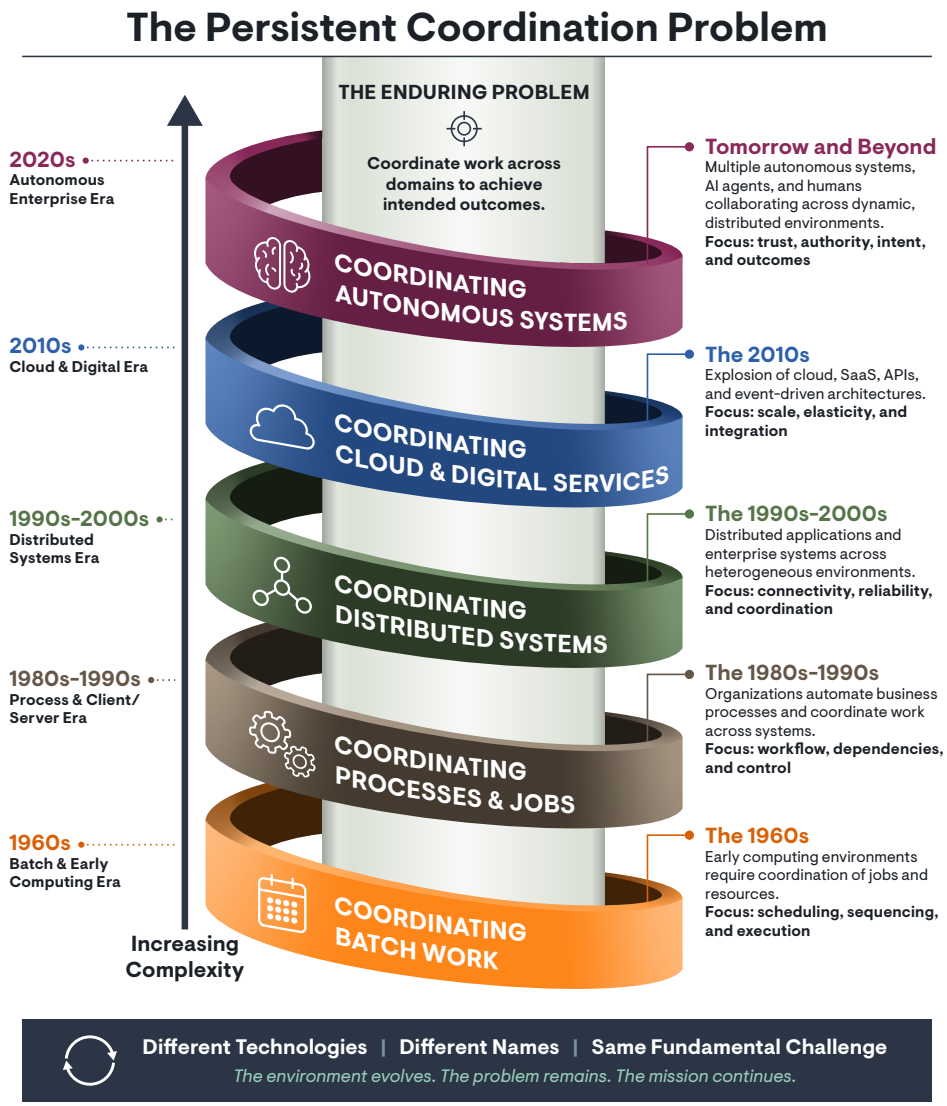

Figure 2.2

New technologies changed what required coordination. More importantly, they changed how much state had to be understood, preserved, interpreted, and acted upon. The environment evolved. The state expanded. The coordination challenge remained.

One of the more important findings emerging from this longitudinal view is that new technologies rarely reduced the need for coordination. Instead, they often increased it. Cloud computing simplified infrastructure provisioning but introduced new deployment models, services, and operational dependencies. Service-oriented architectures and APIs simplified integration while increasing the number of interactions among systems. Observability platforms improved visibility but generated additional information requiring interpretation and action. Automation repeatedly simplified individual activities while increasing the complexity of the overall environment.

A second shift involved velocity. Cloud native architectures and DevOps practices transformed release cycles from quarterly events to continuous deployment models. At precisely the moment enterprise automation was becoming more interconnected, organizations dramatically accelerated the rate at which those environments changed. Complexity increased while the time available to manage it decreased.

Artificial intelligence introduced yet another dimension of complexity. Organizations must now manage not only technical dependencies among systems, but also the behavior of systems capable of making their own decisions – a meaningful expansion of the coordination challenge beyond traditional automation models.

AI occupies a unique position within the evolution of enterprise automation because it appears on both sides of the complexity equation. On one hand, AI introduces new forms of complexity through probabilistic outcomes, reasoning processes, validation requirements, governance concerns, and autonomous decision-making. On the other hand, AI helps organizations manage complexity through natural-language interfaces, workflow generation, code generation, automation assistance, knowledge retrieval, operational guidance, and decision support. AI simultaneously increases complexity and provides new mechanisms for managing complexity.

This dual role becomes particularly important when considering agentic architectures. Much of the industry discussion treats agentic systems as a direct consequence of generative AI. The relationship is real, but it is not complete. The architectural principles underlying agentic systems – decomposition, specialization, delegation, abstraction, and coordination – long predate modern AI. They represent established engineering responses to complexity. Breaking large problems into smaller units, assigning specialized responsibilities, and coordinating outcomes across those units has been a recurring pattern throughout the history of computing and systems design.

Viewed through this lens, agentic architecture serves as both an AI-driven innovation and a complexity management strategy. AI accelerated its adoption and expanded its capabilities, but the underlying objective remains familiar: managing complexity by decomposing work into smaller, more manageable components while coordinating outcomes across the resulting ecosystem. In this sense, agentic systems continue a long tradition of architectural responses to increasing scale, interdependence, and operational complexity.

Rather than breaking this historical pattern, agentic systems represent another turn of the orchestration helix.

The evolution of enterprise automation coordination is fundamentally a story about complexity. New technologies repeatedly simplified individual tasks while increasing the scale, interdependence, velocity, and autonomy of the overall environment. AI continues this pattern, introducing new complexity while simultaneously providing new tools to manage it.

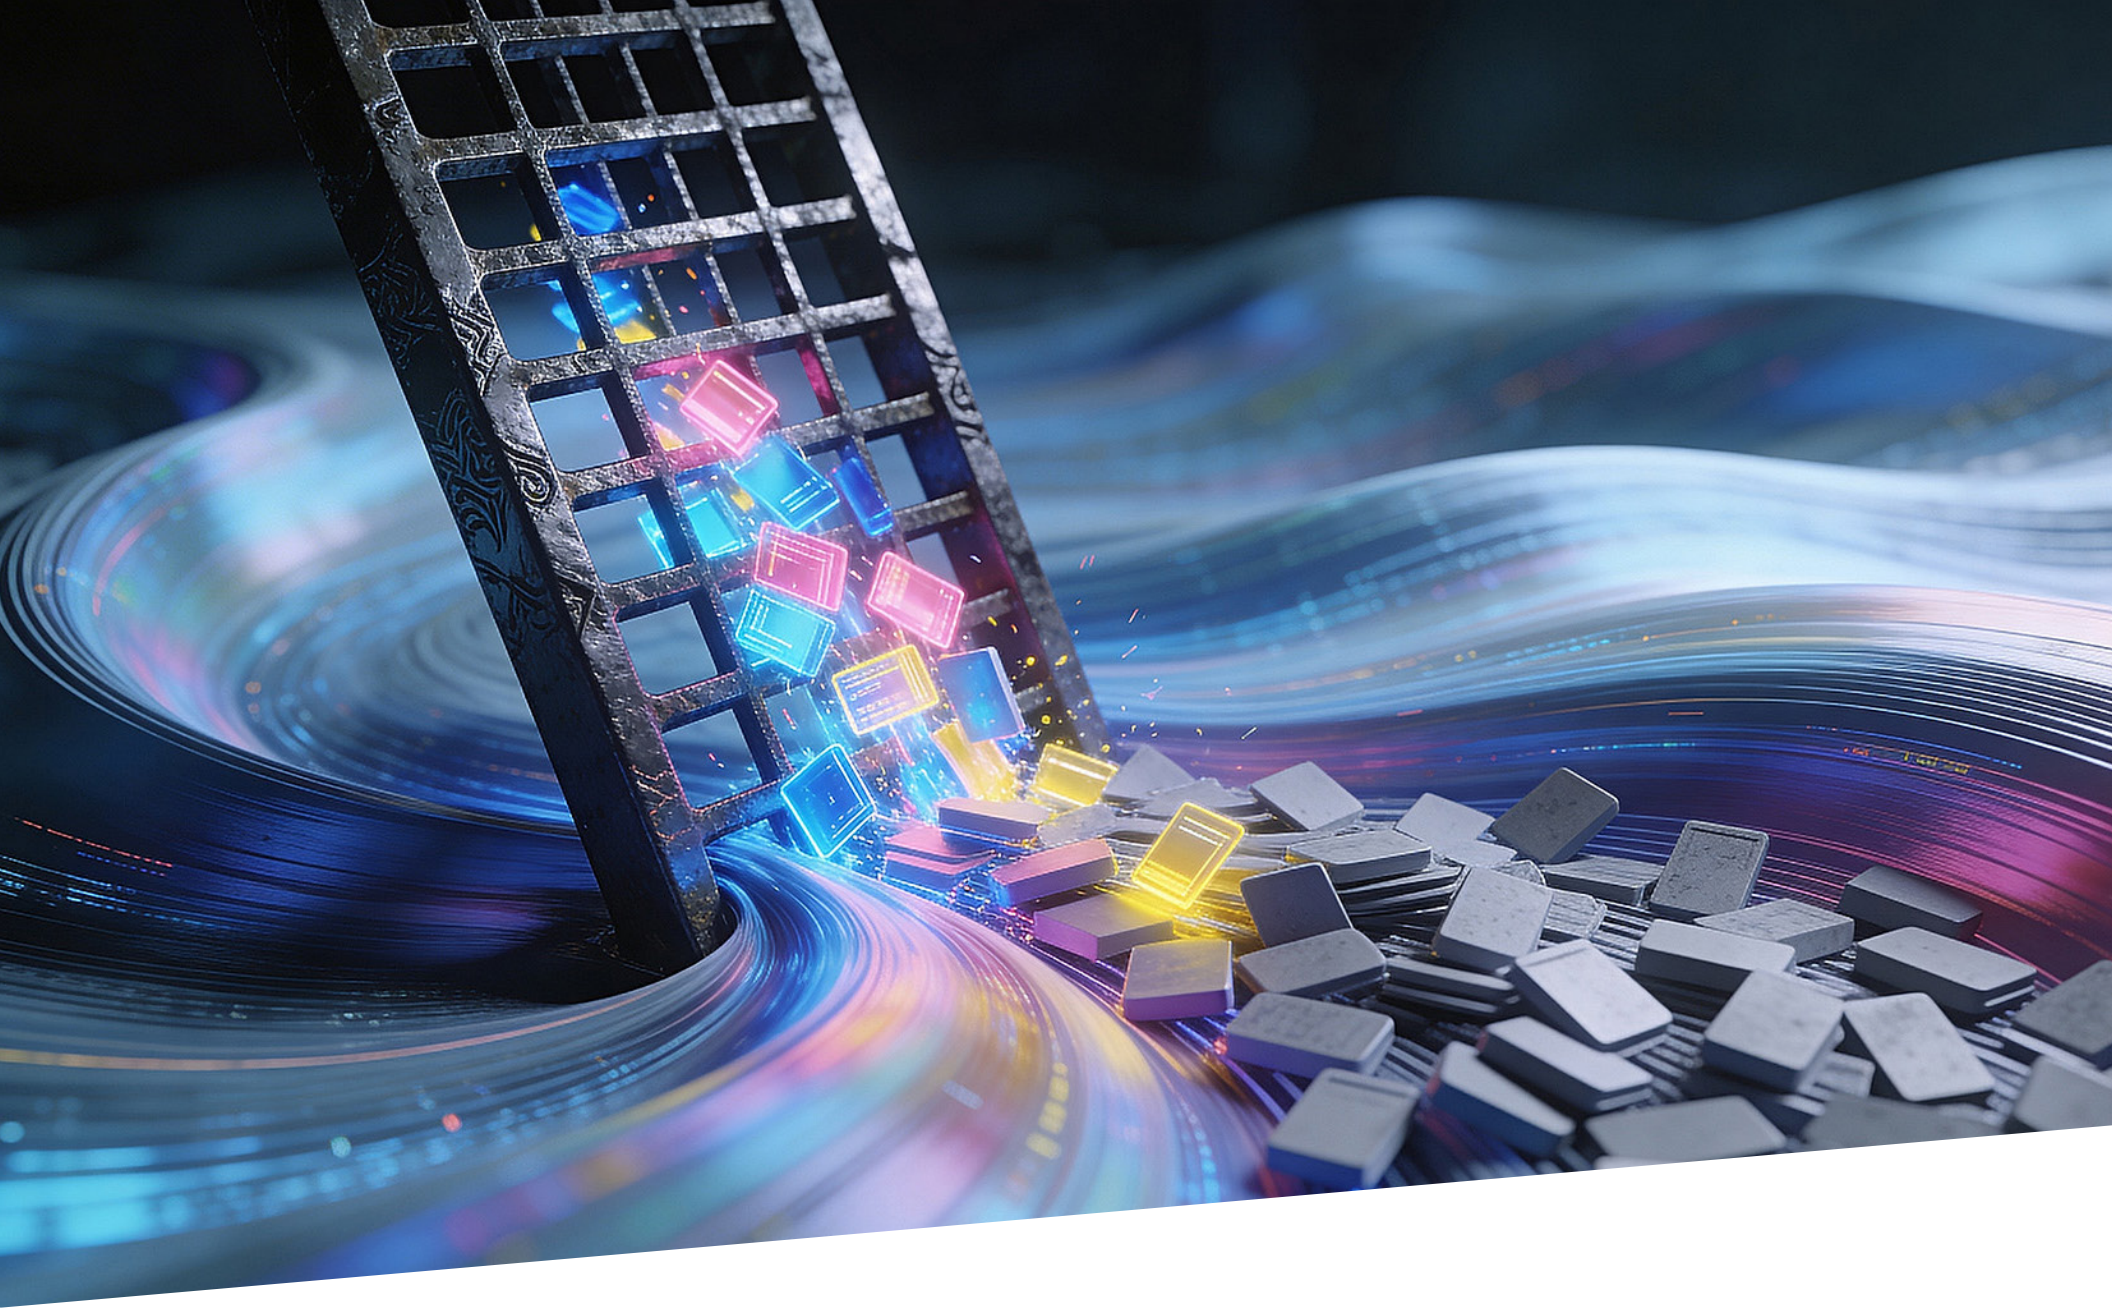

AI Changes Outcome Reliability

What AI changes is not the existence of automation failures – organizations have always encountered those. What AI changes is how organizations must think about reliability, validation, and operational confidence.

Historically, organizations concentrated most validation activities before automation entered production. They developed extensive processes for requirements analysis, testing, quality assurance, governance reviews, staging environments, and release controls, concentrating assurance efforts around relatively infrequent change events. Once teams tested and deployed a workflow, application, or automation process, organizations could generally expect it to execute repeatedly in a predictable and consistent manner until the next significant modification occurred. Outcome failures still happened, but most validation effort was focused on the points where change was introduced.

Artificial intelligence alters this model. Rather than simply executing predefined logic, AI-enabled systems increasingly interpret information, generate recommendations, select actions, construct responses, and participate in decision-making processes during execution. The challenge is no longer limited to validating automation before deployment. Organizations must increasingly consider how confidence is maintained while reasoning occurs in production environments.

In traditional automation:

- change events were relatively discrete
- validation was concentrated around those discrete changes
- operational confidence was largely inherited from prior validation

In AI-enabled systems:

- reasoning and adaptation can continue during production
- operational conditions evolve continuously
- confidence can no longer rely entirely on historical preproduction validation

This shift represents more than a technical difference. Traditional automation concentrated most assurance activities around relatively infrequent deployment events, workflow modifications, infrastructure updates, or configuration changes. Once validated, organizations generally expected systems to behave consistently until the next significant modification occurred. AI increasingly alters this assumption by introducing reasoning and decision-making directly into operational execution. As a result, organizations must increasingly determine which actions, decisions, and outcomes require runtime observability, governance, validation, or intervention within more dynamic operational environments.

In traditional automation environments, organizations rarely needed to evaluate every business outcome directly during execution because deterministic processing made successful execution a strong proxy for successful outcomes. If a workflow completed correctly, processed the expected data, and satisfied operational conditions, organizations could generally assume the intended result was achieved. As a result, operational monitoring focused heavily on execution status, availability, latency, throughput, completion codes, and SLA attainment rather than continuously evaluating the quality of business outcomes themselves. AI increasingly weakens this assumption. As reasoning, interpretation, and probabilistic decision-making become part of execution, successful process completion no longer guarantees correct outcomes, appropriate decisions, or acceptable business results.

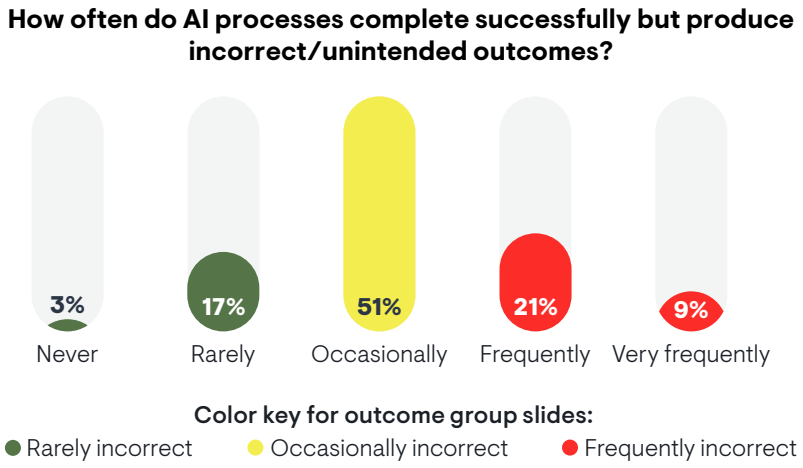

Figure 3.1: Outcome Gap Frequency

Early findings from EMA's research illustrate the significance of this challenge. Approximately 30% of respondents reported encountering incorrect, misleading, or otherwise problematic AI outcomes frequently or very frequently, while fewer than 3% reported never encountering such outcomes. These findings do not suggest that AI systems are inherently unreliable. Rather, they demonstrate that organizations are encountering outcome-quality challenges often enough to make assurance a meaningful operational concern.

The research also reveals the dominant reasons why AI recommendations require human intervention. The most frequently cited concerns involve governance conflicts, incorrect diagnosis, and insufficient context at the time recommendations are generated. Nearly half of respondents indicated that AI recommendations routinely require additional validation before execution. Among the most frequently cited causes: conflicts with governance or compliance policies, incorrect diagnosis or root-cause identification, and insufficient context or explanation at the time of the recommendation. This last factor is particularly notable. Outcome failures frequently stem not from errors in reasoning, but from gaps in the information available when decisions are made.

Much of the current market discussion frames AI reliability primarily as a data quality problem. High-quality data is unquestionably important, but enterprise outcome failures frequently involve issues extending beyond missing or inaccurate data itself. Operational context, business meaning, timing, policy interpretation, organizational priorities, and cross-domain dependencies may all influence whether a technically correct decision produces an acceptable outcome. The challenge increasingly involves preserving sufficient contextual understanding as decisions span multiple systems, operational domains, and reasoning environments.

Operational context, business meaning, timing, policy interpretation, organizational priorities, and cross-domain dependencies may all influence whether a technically correct decision produces an acceptable outcome.

The issue is not simply that errors occur. Enterprise automation has always experienced errors. The difference is that AI can introduce variability into processes that deterministic execution previously dominated. In traditional automation environments, a workflow might be tested once and then execute successfully thousands or millions of times with little variation in behavior. AI-enabled processes may encounter different contexts, different inputs, different interpretations, and different decision paths during normal operation. As reasoning becomes part of execution, the volume of decisions requiring confidence increases substantially.

As reasoning becomes part of execution, the assumptions underlying operational confidence change as well. Organizations that once concentrated assurance activities around relatively infrequent deployment events must now consider how confidence will be maintained as recommendations, adaptive behavior, and decision-making continue during production operations themselves. Operational conditions evolve continuously, introducing ongoing questions regarding where validation is necessary, how much validation is appropriate, and how frequently confidence must be reassessed.

The implications are both operational and economic. Instead of evaluating a workflow primarily when it changes, organizations may need confidence across potentially thousands or millions of reasoning events occurring during normal operation. Assurance itself increasingly becomes a workload consuming compute resources, operational effort, time, and money.

The scope across which AI operates appears to influence outcome reliability as well. Organizations reporting the highest frequency of incorrect outcomes were significantly more likely to deploy AI across multiple interconnected systems, while organizations reporting fewer incorrect outcomes were more likely to operate within isolated environments or highly governed enterprise-wide deployments. As AI decisions span applications, operational domains, and organizational boundaries, maintaining complete context becomes increasingly difficult. Information is abstracted, summarized, transferred, and transformed as workflows cross systems. The challenge is not simply improving AI reasoning.; it is preserving sufficient context as decisions move across increasingly complex operational environments.

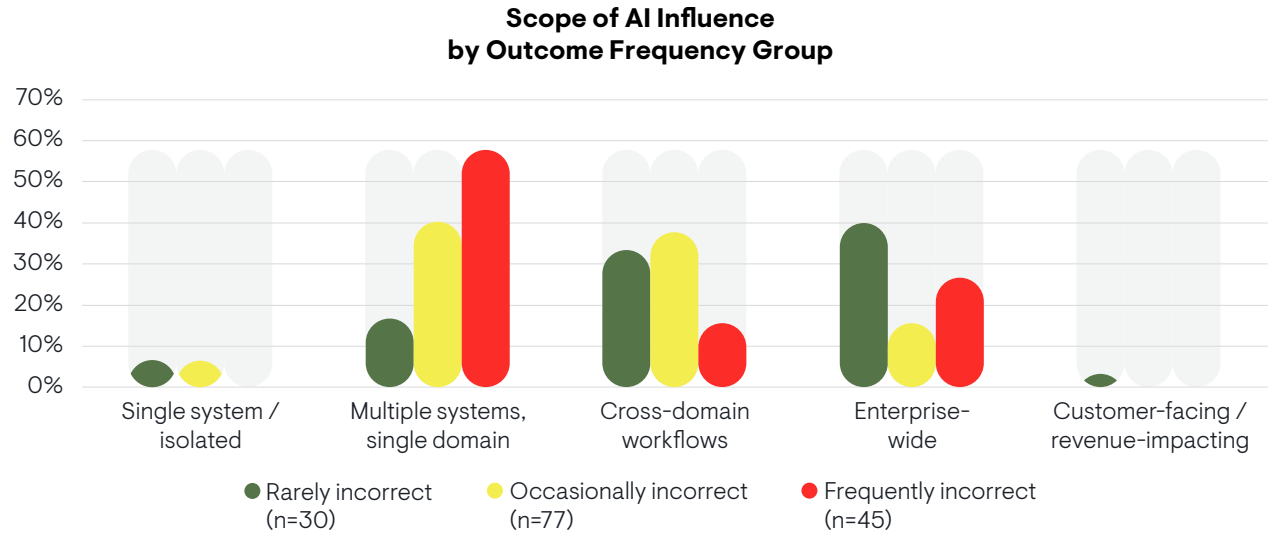

Figure 3.2: Scope of AI Influence by Outcome Frequency Group

Organizations are unlikely to address this challenge by validating every AI-driven decision indefinitely. At scale, exhaustive validation quickly becomes impractical. Instead, enterprises will seek mechanisms that reduce the amount of runtime assurance required while maintaining acceptable confidence in outcomes. Extensive simulation, production replay, testing, guardrails, confidence thresholds, trusted operating patterns, and deterministic constraints can significantly reduce operational risk before deployment. These approaches help establish confidence in common scenarios and create clear escalation paths for situations that fall outside expected boundaries.

However, pre-runtime assurance also has practical limits. No amount of testing can fully anticipate every future business condition, operational context, data variation, or interaction among systems. Eventually, AI-enabled processes encounter situations that were not explicitly exercised beforehand. Some degree of runtime assurance therefore remains necessary. The challenge is not choosing between pre-deployment validation and operational validation. The challenge is determining the appropriate balance between the two.

The operational reality of this challenge is visible in current practice. More than three-quarters of organizations that have deployed AI automation report requiring some level of human correction or rollback of AI-driven actions – ranging from occasional adjustments to frequent intervention. This is not evidence that AI systems are failing. It is evidence that runtime assurance is already an active operational discipline, not a future concern.

As organizations gain confidence in AI-enabled processes, assurance efforts will increasingly focus on exceptions, low-confidence situations, novel conditions, high-risk actions, and sampled reviews of trusted paths rather than exhaustive validation of every decision. The objective is not perfect certainty. The objective is sufficient confidence that outcomes remain within acceptable operational boundaries at a manageable operational cost.

Autonomy intensifies the assurance challenge. As AI systems assume greater operational independence, organizations inherently increase:

- exposure surface
- consequence scope
- independent decision volume
- assurance pressure

The relationship is structural, not incidental. Organizations reporting the highest rates of incorrect outcomes consistently operate with significantly greater levels of AI autonomy, while organizations experiencing fewer incorrect outcomes tend to maintain more constrained decision boundaries. This relationship should not be interpreted as evidence that autonomy itself is undesirable. Rather, it highlights the importance of matching governance, observability, validation practices, and operational controls to the level of authority granted to automated systems. As AI assumes greater responsibility for operational decisions, organizations increasingly require stronger mechanisms to maintain confidence in outcomes.

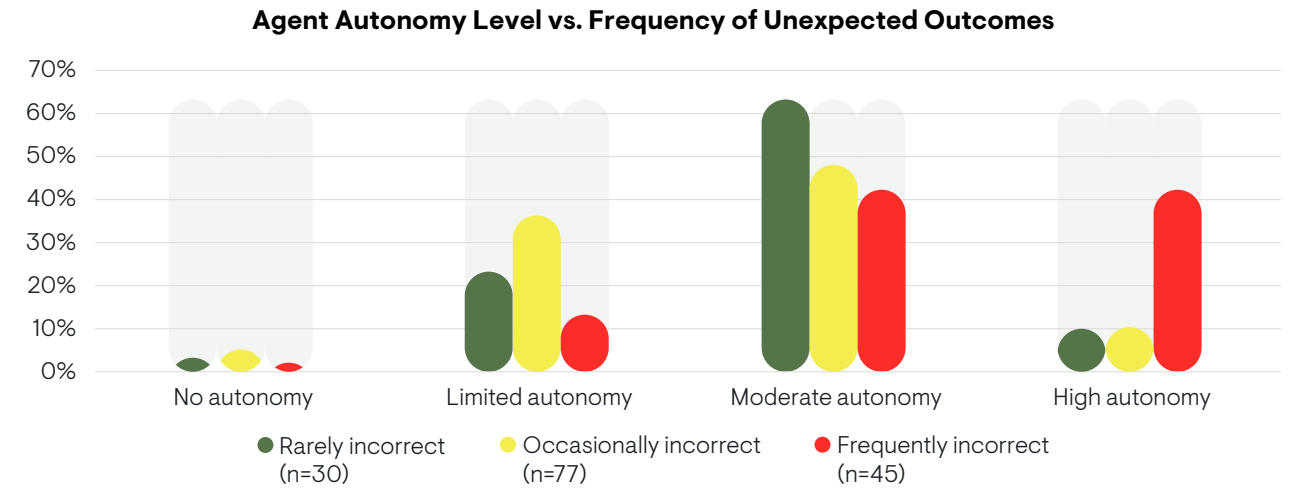

Figure 3.3: Agent Autonomy Level vs. Outcome Frequency

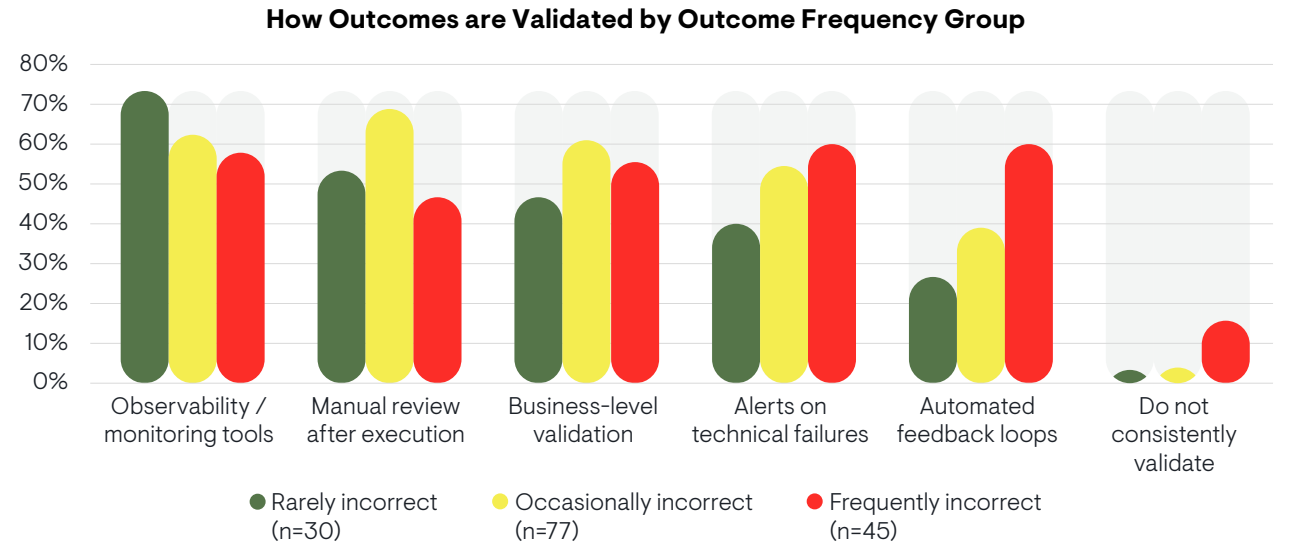

Figure 3.4: Outcome Validation Practices by Outcome Frequency Group

Validation practices themselves appear to influence outcome quality. Organizations reporting the fewest incorrect outcomes rely more heavily on observability, business-level validation, and consistent review processes than organizations experiencing frequent outcome failures. Perhaps most notably, business-level validation – the only validation approach that directly evaluates whether the intended business outcome was achieved rather than simply confirming successful execution – is used more extensively by organizations reporting the strongest results. The findings suggest that outcome assurance increasingly depends upon validating business outcomes rather than merely monitoring technical completion. Successful execution and successful outcomes are not always the same thing.

Successful execution and successful outcomes are not always the same thing.

The emergence of AI therefore changes outcome reliability in ways that extend beyond traditional measures of execution success. Completion status, availability, latency, and throughput remain important operational indicators, but they no longer provide a complete picture of success. Organizations increasingly need visibility into the quality of decisions, recommendations, actions, and outcomes produced during execution. As reasoning becomes operational, assurance must become operational as well.

The question is no longer simply whether a process executed successfully. Increasingly, organizations must determine whether the process produced the intended outcome, whether systems made decisions within acceptable boundaries, and whether organizations can maintain confidence at scale.

AI Changes Orchestration

Artificial intelligence is often discussed as a technology for improving individual decisions, accelerating software development, or increasing personal productivity. While those benefits are important, their impact on enterprise orchestration may prove even more significant. As AI becomes embedded within automation platforms, workflows, and operational processes, orchestration itself begins to change. The challenge is no longer simply coordinating predefined sequences of tasks. Increasingly, organizations must coordinate adaptive decisions, bounded autonomy, human oversight, governance controls, and continuously changing operational conditions.

The implications extend well beyond agents operating independently. Section 4 examines how AI changes automation definition, execution paths, the solution space, what must be coordinated, state awareness, and observability requirements – and what each change means for orchestration.

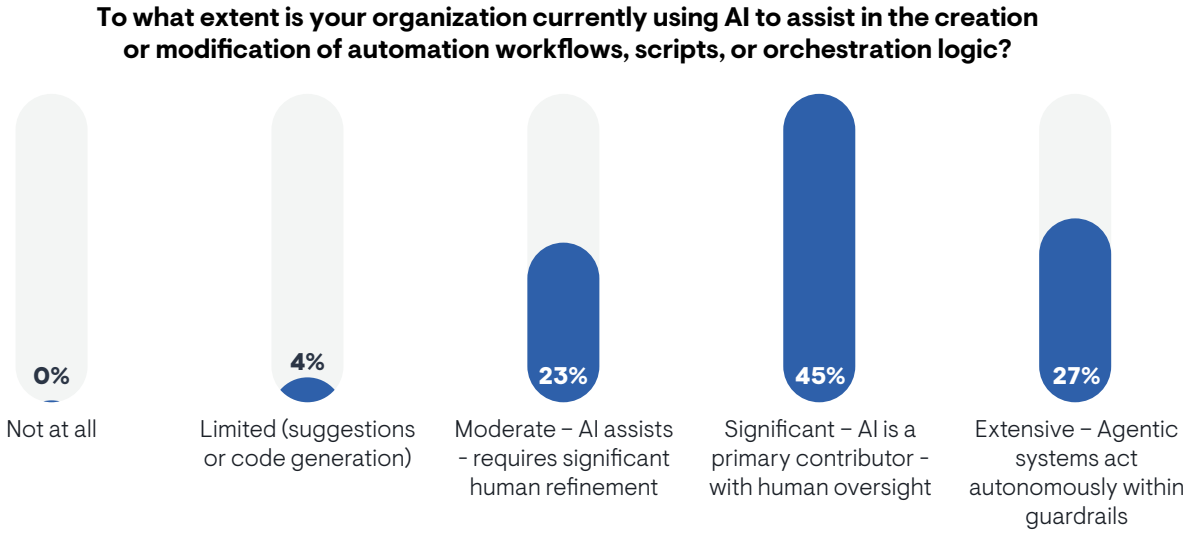

Figure 4.1: AI-Assisted Automation Development and Citizen Developer Findings

AI Democratizes Automation Definition

Historically, the design and coordination of enterprise-wide automation was largely the domain of specialists. Workload automation administrators, integration developers, workflow architects, and automation engineers designed and maintained cross-system automation logic using specialized tools and technical skills. At the same time, many applications and operational domains increasingly incorporated their own embedded automation capabilities, workflow engines, and process logic. While low-code platforms expanded participation over the past decade, designing, coordinating, and governing automation across complex enterprise environments still required significant technical expertise in most organizations.

AI dramatically lowers those barriers. Natural language interfaces, automation assistants, code-generation tools, and agentic development environments increasingly enable business analysts, process owners, operators, and subject matter experts to participate directly in automation creation. Users who may never have developed a workflow previously can now describe business objectives, desired outcomes, or operational requirements and receive working automation logic as a starting point.

This democratization represents one of the most important changes in enterprise orchestration. The number of people capable of defining automation expands dramatically. Automation becomes more accessible, development cycles accelerate, and organizations can automate processes that previously lacked sufficient technical resources. However, democratization also introduces new governance challenges. Organizations must determine who approves automation, how standards are enforced, how changes are managed, how automation is documented, and who provides support when automation behaves unexpectedly.

The challenge is no longer simply building automation. The challenge increasingly becomes governing, supporting, and coordinating a rapidly expanding population of automations and automation creators operating across the enterprise.

> The challenge is no longer simply building automation. The challenge increasingly becomes governing, supporting, and coordinating a rapidly expanding population of automations and automation creators operating across the enterprise.

AI Introduces Adaptive, Reasoned Execution Paths

Traditional orchestration systems primarily coordinated predefined execution frameworks rather than open-ended reasoning processes. Workflows often incorporated dynamic branching logic, exception handling, conditional processing, workload balancing, resource provisioning, retry mechanisms, SLA-driven prioritization, and throughput management based on changing operational conditions. If dependencies failed, workloads surged, or infrastructure constraints emerged, orchestration systems could reroute execution, allocate additional compute resources, delay processing, escalate events, or trigger alternative workflows. However, the available responses were generally bounded and explicitly defined in advance. The operational environment could be dynamic, but the decision framework governing the response remained largely predetermined.

AI introduces a fundamentally different capability. Instead of selecting among predefined branches, systems increasingly evaluate alternatives and choose among multiple possible execution strategies at runtime. A workflow that encounters an unavailable service may seek alternative sources, reorder processing activities, invoke different tools, request human guidance, or pursue another approved path based on current operating conditions.

Consider a simple currency conversion process. Historically, if the approved exchange rate provider became unavailable, the workflow might retry several times before escalating to an operator. An AI-enabled process might instead evaluate alternative data providers, determine whether unrelated processing can continue while waiting for rate information, or identify another approved method of completing the task. The workflow becomes more resilient and adaptive.

At the same time, adaptive behavior introduces new questions. Is the alternative source authorized? Is the information sufficiently accurate? Does reordering activities violate business policy? Could seemingly harmless adjustments create downstream consequences elsewhere in the process? The challenge is no longer whether alternative solutions exist. The challenge increasingly becomes determining which alternatives are acceptable and under what circumstances they may be used.

The important shift is not dynamic behavior itself. Enterprise automation has supported dynamic behavior for decades. The shift is reasoned selection among alternatives during execution.

AI Expands the Solution Space

The ability to reason fundamentally expands the number of possible actions available to an automated process. Traditional automation typically operated within a relatively constrained solution space. Developers anticipated potential conditions, defined acceptable responses, and implemented those responses explicitly within the workflow.

AI systems operate differently. Presented with a problem, they may identify numerous potential solutions that the workflow designer never explicitly anticipated. This flexibility represents one of AI's greatest strengths. It also creates one of the largest orchestration challenges.

Historically, automation platforms focused heavily on determining how work should proceed. Increasingly, orchestration platforms must focus on determining what work is permitted to proceed. Questions such as which systems may be accessed, which data sources are trusted, which tools may be invoked, which decisions require approval, and which actions exceed delegated authority become increasingly important.

The stronger the reasoning capabilities become, the more critical these boundaries become. AI's greatest strength is discovering alternatives. Orchestration increasingly becomes responsible for determining which alternatives are authorized, governed, and aligned with business intent.

AI Changes What Must be Coordinated

Traditional orchestration primarily coordinated operational objects such as jobs, applications, services, files, infrastructure resources, and business processes. While those responsibilities remain important, AI introduces entirely new classes of entities that must be coordinated within operational workflows.

Organizations increasingly need to coordinate AI agents, reasoning activities, confidence thresholds, human approvals, governance policies, escalation procedures, and decision outcomes alongside traditional execution activities. Future workflows are likely to combine deterministic execution, probabilistic reasoning, human judgment, and policy-driven controls within the same operational process.

This creates a significantly broader orchestration challenge. Success is no longer determined solely by whether tasks execute in the proper sequence. Success increasingly depends on coordinating decision-making, governance, oversight, and execution simultaneously. Human review may become part of the workflow. Confidence scores may determine whether processing continues automatically. Policy engines may influence which actions are permitted. Multiple agents may collaborate on portions of a process while humans remain responsible for final approval.

These changes also expand the operational state information associated with enterprise workflows. Traditional orchestration primarily tracked execution state, completion status, dependencies, availability, throughput, and infrastructure conditions. AI-enabled operations increasingly introduce additional forms of operational state, including reasoning state, confidence state, policy state, approval state, escalation state, validation state, and business-context state throughout execution. As orchestration expands beyond deterministic execution, managing these additional operational states becomes increasingly important.

As a result, orchestration increasingly coordinates decisions and governance alongside traditional execution activities.

In which scenarios does your organization require AI-driven actions to operate under orchestration or centralized control, rather than acting independently?

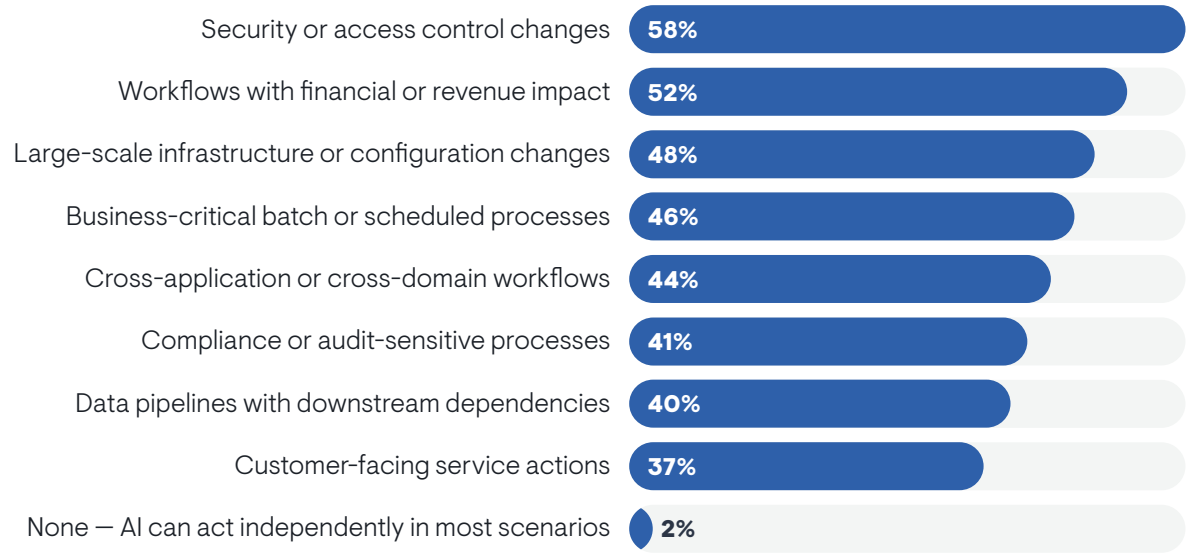

Figure 4.2: Human-in-the-Loop, Governance, and Agent Coordination Findings

AI Increases the Importance of State Awareness

Coordinating these increasingly complex operational entities requires significantly broader state awareness. Traditional orchestration systems primarily evaluated execution state, dependency status, resource availability, triggering conditions, and workflow completion status. AI-enabled orchestration increasingly depends on continuously changing operational context across the environment itself. Infrastructure health, application performance, network conditions, workflow progress, business process status, approval states, policy evaluations, reasoning confidence, and downstream operational readiness may all influence execution decisions in real time.

This is not a dependency AI introduces. Deterministic automation has always required state awareness. Before a workflow can determine what to execute, it must understand dependencies, prerequisites, schedules, resource availability, execution conditions, and operational status. State is not an enhancement to execution or intelligence. It is the foundation on which both operate.

AI does not create the dependency on state. It expands the scope of relevant state, elevates the consequences of incomplete or fragmented state, and makes state coordination increasingly central to reliable execution across every orchestration model.

The relationship between orchestration and observability therefore changes substantially. Historically, observability primarily provided visibility for operators and administrators. Increasingly, observability becomes a direct operational input into orchestration decisions themselves.

An AI-enabled workflow deciding how to proceed may need awareness of service availability, current latency conditions, API responsiveness, security status, approval state, workflow context, confidence levels, and downstream process readiness before selecting the next action. Expanded autonomy requires expanded awareness.

The orchestrator is no longer simply coordinating execution. It increasingly coordinates execution based on continuously changing operational state across the enterprise. As enterprise workflows span multiple domains, orchestration increasingly depends on maintaining awareness not only of local state, but also of state transitions occurring across the broader operational environment.

Which best describes how observability data is used to drive operational actions in your environment?

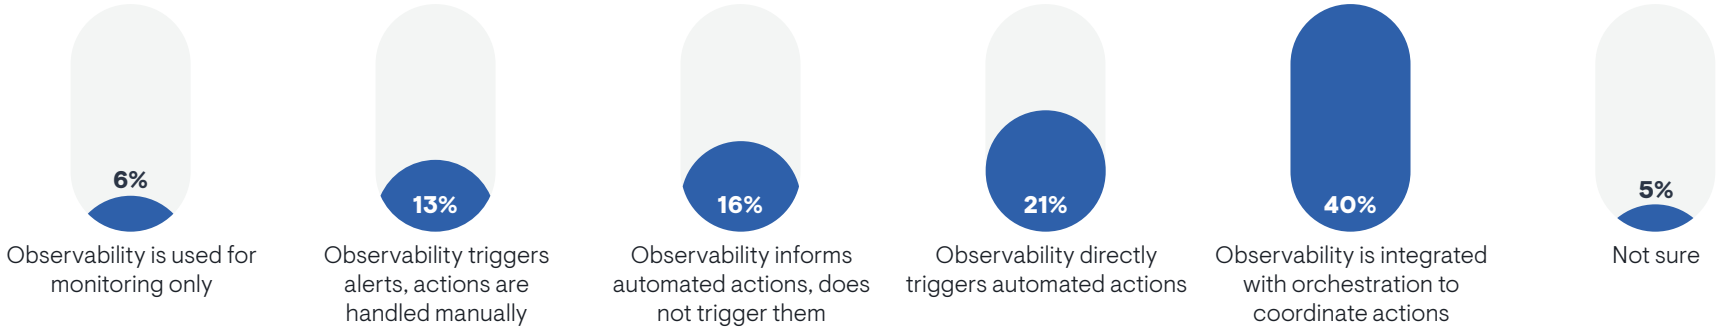

Figure 4.3: Observability-Driven Automation and Operational Intelligence Findings

AI Creates New Requirements for Automation Observability

Greater state awareness introduces a second challenge. Organizations increasingly need visibility not only into what happened, but also into why decisions were made.

Traditional observability focuses on logs, events, metrics, and traces. These telemetry streams help organizations understand execution behavior, identify failures, measure performance, and diagnose operational issues. They explain what occurred and where it occurred.

Adaptive orchestration introduces an additional requirement: decision observability.

Organizations increasingly need answers to questions such as: Why was this action selected? What alternatives were considered? What information influenced the decision? What confidence level existed at the time? What policies constrained the available options? Why was one path selected instead of another? Why was human review requested?

As reasoning becomes part of operational execution, reasoning itself becomes part of the operational record. Enterprises increasingly require visibility into automation state, decision state, reasoning state, confidence state, and policy state alongside traditional execution telemetry.

This requirement extends beyond troubleshooting. Auditability, governance, compliance, trust, and outcome assurance all depend upon understanding not only what actions occurred, but also why those actions occurred. A workflow that selected an alternate data source, reordered processing activities, escalated to a human reviewer, or invoked an additional agent may have executed correctly. Understanding why those decisions were made becomes equally important.

Traditional observability explains what happened.

Automation observability increasingly explains why the automation chose to make it happen. Organizations that have navigated the early stages of AI deployment identify this visibility as foundational: human oversight of AI-driven decisions, observability of automated actions, and enforceable guardrails on permitted behaviors consistently emerge as the most important operational disciplines for managing AI at scale.

Which conditions must be met before your organization is willing to grant AI greater operational autonomy?

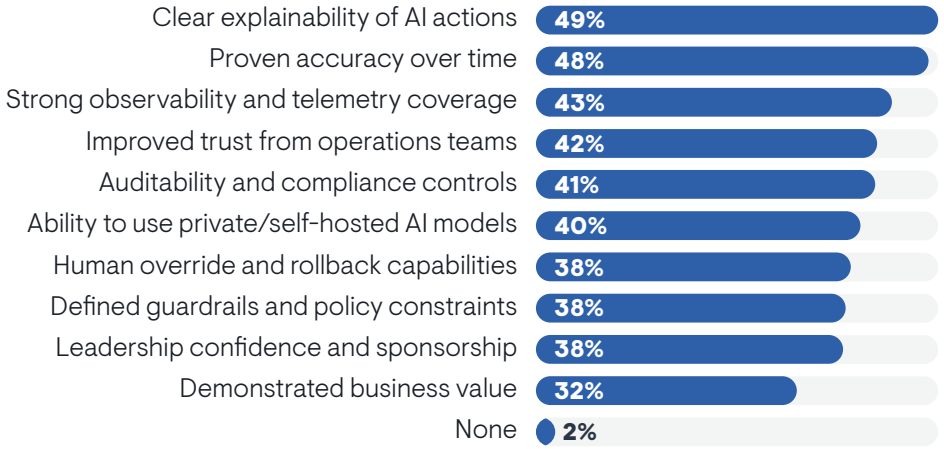

Figure 4.4: Trust, Explainability, Governance, and Outcome Assurance Findings

Taken together, these changes significantly expand the role of orchestration within modern enterprises. AI democratizes automation creation, introduces adaptive execution paths, expands the range of possible actions, broadens what must be coordinated, increases dependence on operational state, and creates entirely new observability requirements. The question is no longer simply how workflows are executed. The question increasingly becomes how autonomy, governance, trust, state awareness, and decision-making are coordinated across an increasingly complex automation landscape.

The Reality of Federated Orchestration

Section 1 established that enterprise orchestration is already federated. This section examines the structural reasons that federation emerged, why it is unlikely to reverse, and what it means for coordination as AI expands throughout every orchestration domain.

Enterprise Orchestration is Already Federated

The survey data reveals that no single platform dominates enterprise orchestration. Respondents identified a variety of technologies as playing the most significant role – spanning the full range of orchestration domains already operating across the enterprise. The distribution itself is more important than any individual percentage.

Most organizations already coordinate work through multiple orchestration centers, each responsible for different portions of the enterprise – business transactions through ERP, service operations through ITSM, application delivery through CI/CD, or infrastructure through cloud automation. Many operational workflows span several of these domains simultaneously.

The significance of the findings is therefore not which platform currently coordinates the largest share of activity. The significance is that enterprises already operate through a federated orchestration model. Federation is not a future state. It is the current operational reality.

> Federation is not a future state. It is the current operational reality.

Federation Emerged Through Specialization

Federation emerged because different operational domains required different forms of coordination, governance, expertise, and contextual awareness. As enterprises expanded digitally, specialized orchestration platforms evolved to optimize decision-making within specific operational environments.

Over time, orchestration capabilities evolved closer to the operational domains they supported. Different platforms optimized for different forms of coordination, operational priorities, governance requirements, ownership models, and decision-making patterns. Business systems emphasized transactional consistency and process governance. Service management platforms became optimized for operational workflows and support coordination. Cloud and infrastructure platforms prioritized elasticity, provisioning, and resource lifecycle management. Software delivery environments focused on deployment velocity and release automation. Each orchestration domain evolved around specialized operational context that could not easily be centralized elsewhere. Practical organizational needs drove the resulting architecture, not technological inefficiency.

This specialization delivered significant business value and remains an essential characteristic of modern enterprises. Organizations are unlikely to abandon specialized orchestration platforms simply because AI introduces new capabilities. If anything, AI increases the importance of domain expertise by making local decision-making more powerful within each orchestration environment.

The challenge therefore becomes coordinating specialized orchestration domains rather than replacing them.

Workload Automation Occupies a Distinct Cross-Domain Role

Many orchestration platforms primarily coordinate activity within a particular operational domain. Workload automation and service orchestration and automation platforms evolved to serve different purposes. If enterprise automation is viewed as the operational floor of the enterprise and individual orchestration domains represent the tiles, workload automation increasingly acts as the grout binding orchestration domains together into coordinated operational workflows.

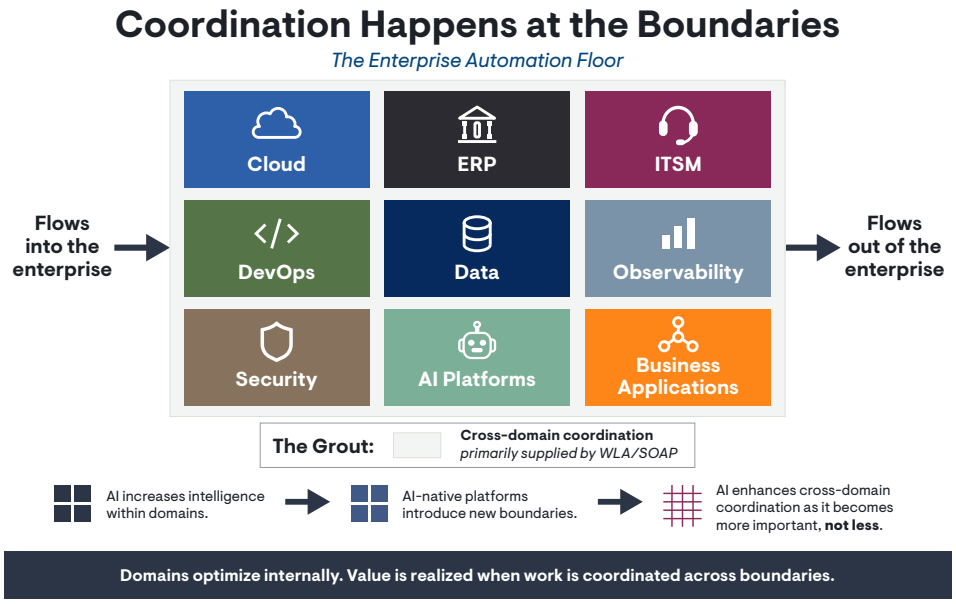

Figure 5.1 Coordination Happens at the Boundaries

A single workload automation workflow might coordinate ERP transactions, database processing, application services, cloud resources, external data providers, file transfers, business reporting systems, and third-party partners within the same end-to-end process. In many organizations, workload automation became responsible not for owning individual systems, but for coordinating activity between them.

This distinction becomes even more important in the AI era because enterprise outcomes increasingly depend upon interactions that cross organizational, technological, and operational boundaries. Financial closing processes, customer onboarding workflows, supply chain operations, business intelligence pipelines, and countless other processes require coordination among multiple systems rather than within a single domain.

The enduring value of workload automation therefore extends beyond executing jobs or managing schedules. Its role frequently involves coordinating dependencies, preserving operational continuity, and maintaining visibility across domain boundaries. As enterprise workflows become more distributed and adaptive, cross-domain coordination remains an important operational requirement regardless of which individual orchestration platforms participate.

AI Increases the Importance of Cross-Domain Coordination

AI capabilities are emerging simultaneously across every orchestration domain. Each domain is gaining more intelligent, autonomous capabilities tailored to its own operational context, from AI-assisted business processes in ERP to autonomous workflows in ITSM, to infrastructure optimization in cloud platforms, to AI-driven recommendations in observability. Agentic systems introduce entirely new forms of autonomous decision-making on top of these domain-specific advances.

As a result, enterprise workflows increasingly span multiple reasoning environments rather than merely multiple execution environments.

Orchestration primarily coordinates actions. Increasingly, orchestration must coordinate decisions made by multiple autonomous systems operating across different domains. Multiple systems may evaluate conditions, recommend actions, invoke agents, trigger workflows, or influence outcomes within the same operational process. The complexity of coordinating these activities grows significantly as AI becomes embedded throughout the enterprise.

The challenge is not simply that more decisions arise. The challenge is that those decisions are made within different domains operating under different objectives, constraints, policies, and context models – each with incomplete visibility into the others. This is the structural condition that makes context preservation the central operational problem of federated AI orchestration.

Federation Creates a Context Challenge

The effectiveness of AI reasoning depends heavily upon context. Decisions are only as good as the information available when those decisions are made. This creates a significant challenge within federated environments.

Historically, many decisions occurred relatively close to their source of expertise. Business decisions remained within business systems. Service decisions remained within service management environments. Infrastructure decisions remained within infrastructure management tools. The systems making decisions generally possessed rich contextual awareness regarding the activities they managed.

As workflows span multiple orchestration domains, maintaining that contextual richness becomes increasingly difficult. Systems abstract, summarize, translate, and transfer information as it crosses boundaries. State information may be incomplete, operational priorities may differ, and important assumptions may not accompany shared information. The further decisions move from their originating context, the greater the potential for information degradation.

> The further decisions move from their originating context, the greater the potential for information degradation.

This challenge directly influences outcome reliability. A workflow may execute successfully while still producing an undesirable outcome because critical context was unavailable at the moment a decision was made. AI systems amplify this concern because reasoning quality depends on the quality and completeness of available information. In many cases, execution errors may not cause outcome failures at all. Context failures may cause them instead.

The research supports this directly. Among organizations that have rejected AI recommendations, insufficient context or explanation ranks as one of the most frequently cited causes – named by nearly four in ten respondents. Context failures are not theoretical. They are among the most common reasons organizations override automated decisions in practice.

The challenge of federated orchestration therefore extends beyond coordinating execution. It increasingly involves preserving sufficient context and state awareness to support reliable decisions and coordinated outcomes across domain boundaries.

Coordination Becomes More Important Than Consolidation

Many orchestration categories increasingly seek broader operational influence. Workload automation platforms are expanding into agentic orchestration, governance, and broader automation coordination. Service management vendors increasingly position themselves as enterprise workflow hubs. ERP providers seek to coordinate AI-driven business execution. CRM platforms, such as Agentforce, aim to orchestrate customer-facing processes and autonomous agents. Observability vendors increasingly move from visibility toward operational action. Cloud providers continue extending automation capabilities deeper into enterprise operations.

Each category is pursuing a rational strategy. As AI expands automation opportunities, the value increasingly shifts toward coordinating decisions, workflows, and outcomes rather than simply executing individual tasks.

The survey findings nevertheless suggest that enterprises remain fundamentally federated. The more important challenge is enabling multiple orchestration platforms to work together effectively – preserving the governance, context, and accountability disciplines that federated environments make increasingly difficult.

Enterprise coordination increasingly becomes more important than orchestration consolidation.

As AI expands automation and increases cross-domain dependencies, the challenge shifts from orchestrating individual workflows toward coordinating decisions, policies, governance, and outcomes across the enterprise itself.

The Governance Challenge

The federated, AI-embedded environment described in Sections 3, 4, and 5 creates a direct governance challenge.

Historically, enterprises governed operational authority through organizational structures. Management hierarchies, approval processes, policies, accountability mechanisms, and organizational culture provided the framework through which authority was delegated and actions were controlled. These governance disciplines did not disappear as automation expanded. Instead, the actions being governed increasingly moved into applications, workflows, orchestration systems, automation platforms, and AI-enabled technologies.

Enterprise governance is therefore entering a new phase. The challenge is no longer simply approving predefined workflows before deployment. Increasingly, enterprises must coordinate distributed operational authority across human and digital actors operating throughout a federated environment.

The result is not simply a technology challenge. It is also a challenge of maintaining governance, accountability, authority, and operational control as decision-making becomes increasingly distributed across autonomous systems.

The survey findings suggest organizations are actively searching for governance models capable of balancing increasing autonomy with accountability, visibility, and operational control. While approaches vary, respondents consistently favor governance mechanisms that coordinate actions across domains rather than unrestricted autonomous operations.

Authority is Already Distributed

Decision authority is already distributed across the modern enterprise.

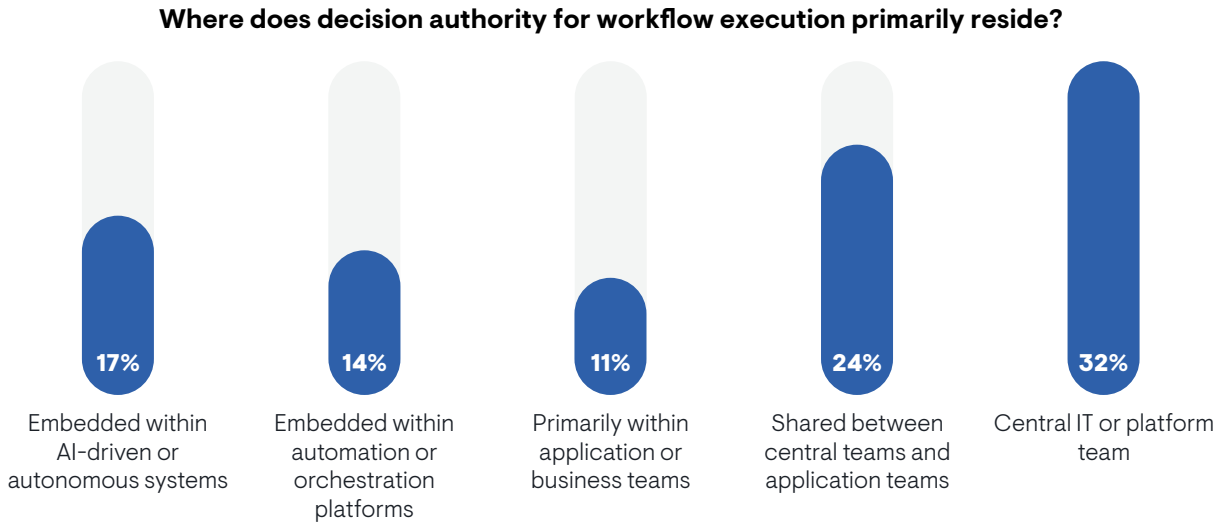

Figure 6.1: Decision Authority Distribution

While centralized IT organizations continue to play a significant role, respondents also identified shared governance structures, embedded automation platforms, business stakeholders, and AI-enabled systems as participants in operational decision-making. No single governance model dominates the results.

This finding reflects broader trends occurring throughout enterprise technology. Cloud adoption distributed infrastructure responsibility. DevOps distributed software delivery ownership. Platform engineering introduced new operational models. AI extends this progression by introducing software systems capable of influencing operational decisions directly.

As automation expands throughout the enterprise, authority increasingly becomes a shared responsibility. Decisions affecting business outcomes may involve application teams, operations teams, service management organizations, platform engineering groups, automation specialists, and AI-enabled systems operating within those environments. The challenge therefore becomes less about identifying a single owner and more about coordinating authority across multiple participants.

Autonomy Requires Governance

Greater autonomy increases both operational capability and operational consequence:

- greater autonomy expands decision authority
- greater autonomy increases operational consequence
- greater autonomy therefore increases governance requirements

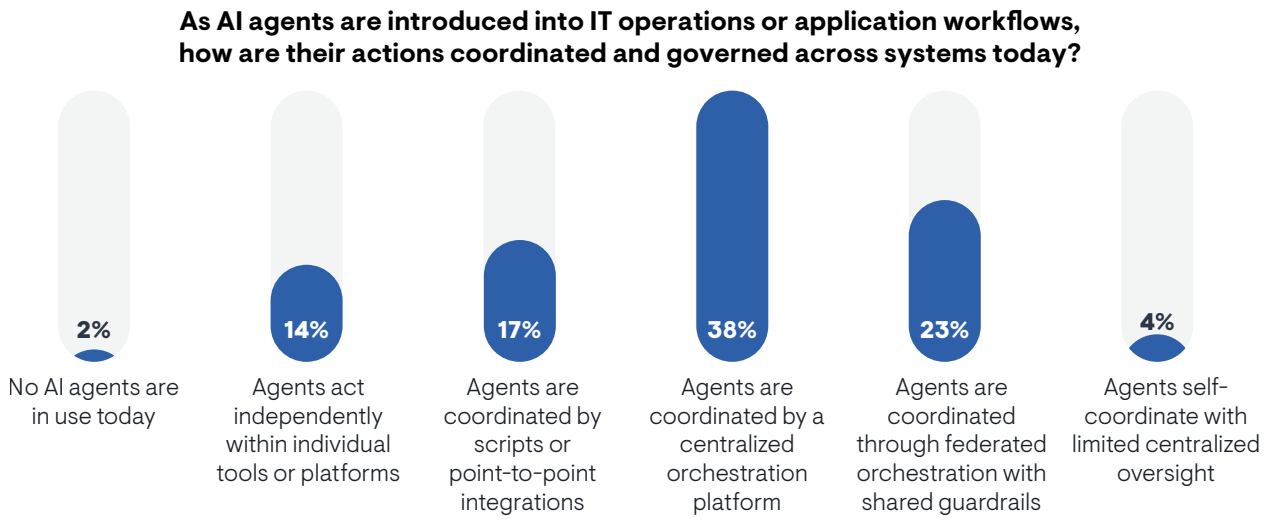

Figure 6.2: Preferred Agent Coordination Model

Respondents overwhelmingly favored approaches that combine automation with centralized coordination, federated governance, shared guardrails, and operational oversight. By contrast, relatively few organizations expressed a preference for completely independent or self-coordinating agents operating without broader governance structures.

Organizations clearly expect AI to play an increasingly important role in coordinating work, decision-making, and operational execution. At the same time, the findings demonstrate that enterprises do not view autonomy as something that should operate without governance boundaries. As operational authority expands, governance requirements expand with it.

Enterprises are treating autonomy as operational authority that must be earned incrementally through demonstrated reliability. Organizations increasingly recognize that AI systems require clearly defined boundaries regarding what actions may be taken, under what circumstances, using which information sources, and with what degree of human involvement. The challenge is not whether AI systems can make decisions. The challenge is determining which decisions remain subject to policy controls, review requirements, escalation paths, approval structures, and auditability standards.

> Enterprises are treating autonomy as operational authority that must be earned incrementally through demonstrated reliability.

When asked how they intend to calibrate the level of autonomy granted to AI systems, organizations identified concrete operational criteria rather than broad sentiment. Accuracy and consistency of recommendations ranked highest, followed closely by demonstrated stability and reliability over time and compliance with existing governance policies. Enterprises appear prepared to expand AI authority incrementally based on demonstrated operational performance rather than generalized confidence or vendor assurances.

Governance Must Span Operational Domains

Federated orchestration creates a cross-domain governance challenge.

Which best describes how observability data is used to drive operational actions in your environment?

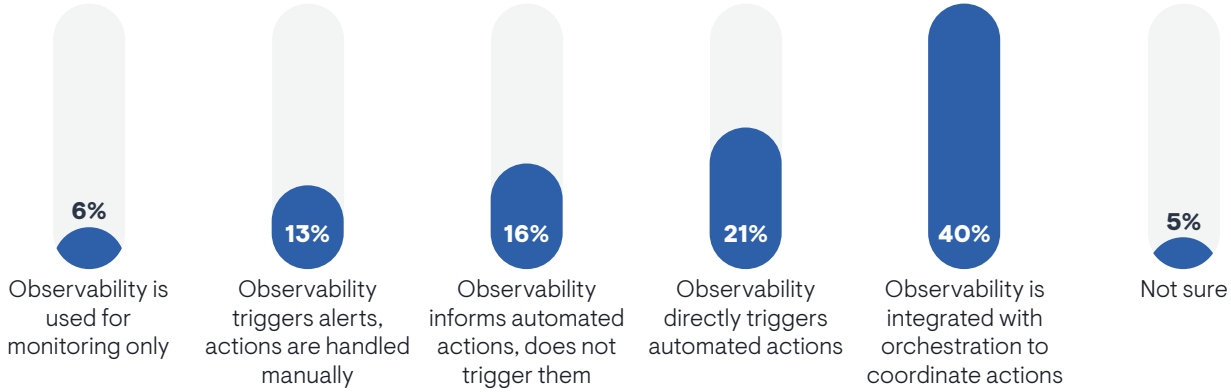

Figure 6.3: Observability as an Operational Input

Enterprise workflows increasingly span the full range of operational domains. AI expands this complexity further by introducing operational decisions that may originate in one domain while triggering actions, workflows, or policy consequences in another.

Governance models designed primarily around individual platforms increasingly struggle to manage these cross-domain interactions. Observability systems may influence operational actions. AI agents may coordinate work across platforms. Workflow decisions may affect systems with governance policies that originate elsewhere.

Cross-domain governance is not a new requirement that AI introduced. Workload automation and orchestration platforms have spent decades evolving governance mechanisms for coordinating operational activity across distributed enterprise environments. AI significantly expands the scale, variability, and runtime complexity of those governance requirements, but it does not eliminate the value of the existing coordination layer. Extending existing cross-domain coordination and governance capabilities may ultimately prove more practical than attempting to build entirely separate governance structures disconnected from operational execution itself.

Observability is Becoming Part of the Governance Model

Observability increasingly influences operational governance and automated decision-making.

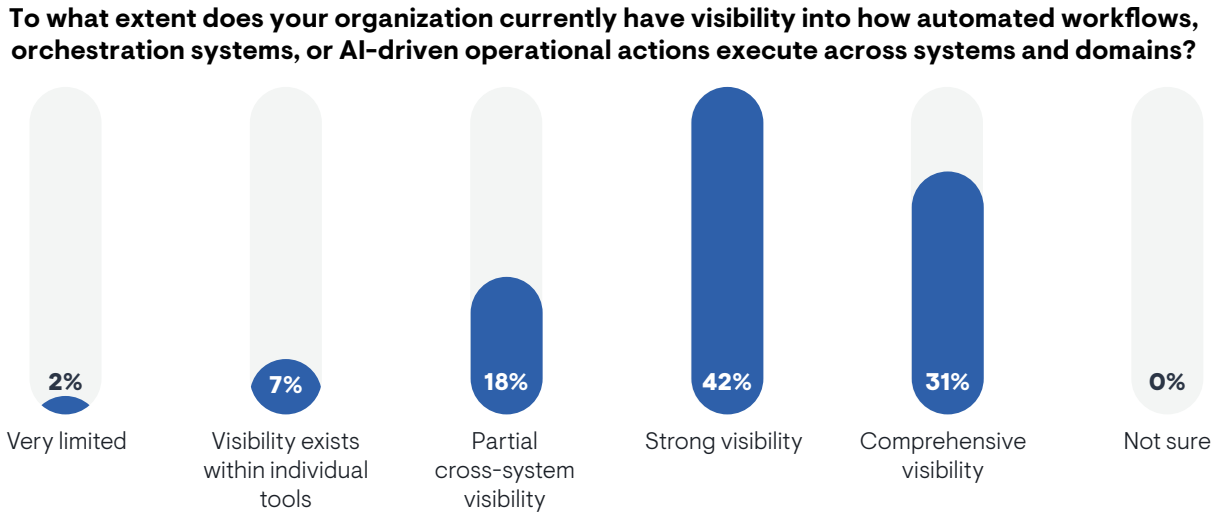

Figure 6.4: Observability Influence on Governance and Automated Actions

Metrics, logs, traces, telemetry, and operational state information increasingly function as direct inputs into orchestration decisions, automated actions, AI-assisted recommendations, escalation paths, and governance processes rather than serving solely diagnostic or monitoring purposes.

Organizations increasingly use observability data to influence automated actions, operational responses, orchestration decisions, and AI-assisted recommendations. Telemetry is becoming an operational input rather than merely an operational output.

This operational shift has significant governance implications. When operational decisions increasingly depend upon system state, observability becomes part of the decision-making process itself. Governance therefore requires visibility not only into system performance, but also into how telemetry influences actions taken by automation platforms, orchestration systems, and AI-enabled services.

The traditional distinction between observing operations and influencing operations continues to narrow. As organizations automate more operational decisions, observability increasingly becomes an active participant in governance rather than simply a monitoring capability.

Reliability Demands Accountability

The outcome reliability findings discussed earlier provide additional context for the governance challenge.

Organizations continue to experience incorrect outcomes, unexpected behavior, and operational failures even when workflows execute successfully. As autonomy increases, understanding how outcomes occur becomes increasingly important.

Trust is not created solely through successful execution. Trust depends upon visibility, accountability, and the ability to understand how outcomes were produced. Organizations increasingly require mechanisms capable of explaining what actions occurred, why those actions occurred, what information influenced decisions, and how business outcomes were affected.

This requirement extends beyond troubleshooting. It influences governance, compliance, operational risk management, auditability, and executive confidence in automated systems. As AI becomes more involved in operational decision-making, organizations increasingly need visibility into both actions and decisions.

The challenge is not simply enabling automation to act. The challenge is ensuring that automated actions remain understandable, governable, and accountable.

The Governance Model Remains Unsettled

Taken together, the findings suggest that organizations are entering a transitional period. Organizations largely designed existing governance approaches for deterministic automation operating within well-defined ownership structures. AI introduces adaptive decision-making, distributed authority, autonomous actions, and increasing coordination across operational domains.

Organizations are responding through a variety of approaches. Some emphasize centralized orchestration. Others adopt federated governance models with shared guardrails. Some increasingly embed decision-making authority within platforms, while others continue to rely on human review and oversight. No single governance model has yet emerged as dominant.

What does emerge consistently is that governance increasingly becomes a cross-domain coordination challenge rather than a platform-specific control problem. As orchestration becomes more federated and automation becomes more autonomous, the full set of governance and accountability disciplines must increasingly operate across distributed operational environments rather than within isolated systems.

> As orchestration becomes more federated and automation becomes more autonomous, the full set of governance and accountability disciplines must increasingly operate across distributed operational environments rather than within isolated systems.

Organizations are still determining how these governance models will evolve. What appears increasingly clear is that expanding autonomy does not reduce the need for any of the assurance disciplines defined in this report. It significantly increases the operational importance of all of them.

Practitioner Perspectives: Governing Autonomous Operations

Survey respondents consistently expressed confidence in AI's potential to improve operations. At the same time, their comments reveal a consistent set of concerns that cuts across platform types and industries: autonomy without context, governance, observability, and human oversight remain unacceptable for critical operational environments. The following voices illustrate the dominant themes that emerged across the survey dataset.

Human Oversight Remains Essential

"The primary concern is the lack of a 'human-in-the-loop' for critical decision-making. In IT operations, an autonomous AI could misinterpret a complex system anomaly and initiate a massive automated rollback or configuration change that leads to unexpected downtime. We are cautious because the risk of a catastrophic, automated error outweighs the efficiency gains at this stage."

Context Matters More Than Intelligence

"We still want to maintain human oversight over critical changes because AI sometimes does not fully understand the real operational context. If the core system is handled incorrectly, the impact can spread and directly affect customers and business operations."

The Risk of Errors is the Primary Barrier

"Our primary concern is the potential for unpredictable systemic impact and the loss of auditability. While AI can identify patterns faster than humans, granting it full autonomy risks cascading failures in which the AI's automated fixes create new, unforeseen issues that our engineers cannot quickly diagnose. Maintaining a human in the loop ensures that we preserve data integrity and operational stability for our users."

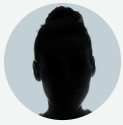

"The main reason my organization is cautious about granting AI more independence is a lack of ability to determine the decision-making reasons AI uses –more specifically, the inability to determine how or why something went wrong when AI executes an action on its own. There is also an inability to assign accountability for poor decisions or hallucinations."

Compliance and Security Create Hard Limits

"As a global financial institution, our primary constraint is maintaining strict regulatory compliance and absolute data integrity. Non-deterministic AI behavior poses an unacceptable risk to our auditability requirements. Uncontrolled changes to core transactional databases, cross-domain workflows, or security boundaries could result in severe financial exposure and regulatory non-compliance. Therefore, we mandate explicit centralized guardrails and human-in-the-loop oversight for all critical operational pipelines."

"Activities such as payment approvals, customer data processing, major operational changes, and legal and financial decision-making are considered too risky for AI to perform fully autonomously. We are concerned about accuracy, data security, and accountability when incidents occur."

Governance Creates Confidence

"AI works best in IT operations when kept within strict guardrails. A significant portion of automated actions are still rejected due to risk and lack of context, and fully agentic AI is still very limited in production. Human oversight and governance remain critical for safe adoption."

"Our primary concern is ensuring reliability, security, compliance, and human oversight before allowing AI to take autonomous operational actions. We need stronger guardrails, explainability, auditability, and rollback capabilities to prevent unintended cross-system impacts, service disruptions, or security risks."

Observability is a Prerequisite, Not an Add-On

"We learned that good observability is the critical foundation before applying AI to IT operations. When data between applications, infrastructure, and services is not fully linked, AI struggles to accurately identify root causes or the true extent of impact."

"AI cannot operate effectively in production without strict, centralized policy guardrails and a robust human-in-the-loop validation process. While AI excels at predictive anomaly detection and parsing massive volumes of telemetry data across hybrid cloud environments, its suggestions often lack the specific compliance and regulatory context required in tier-1 finance. Pre-execution human verification has proved essential for preventing non-deterministic model drift from causing unauthorized database or configuration changes, resulting in a production rollback rate of under 5%."

Data Quality Determines AI Reliability

"The biggest lesson we have learned is that AI is only as good as the data feeding it. Early on, we tried to automate incident response, but because our data was siloed and messy, the AI generated significant noise and false positives."

"AI is only truly effective when monitoring data is clean and consistent across systems. In the early stages, we encountered many false alerts due to data lacking context, but after standardizing logs, metrics, and tracing, incident detection improved significantly."

Trust Must be Earned

"We learned that AI should be deployed in small phases rather than scaling immediately across the entire system. Testing first on low-risk processes helps the operations team understand AI's limitations and build better trust over time."

EMA Perspective: From Cross-Domain Complexity to State Coordination

Cross-domain complexity is no longer an edge case. It is now the normal operating condition for enterprise automation. Nearly half of respondents report that between 26% and 50% of automation workflows span multiple domains, while more than one-third report that a majority of workflows cross domain boundaries. This means the coordination challenge is no longer limited to exceptional workflows or highly customized environments. It is structural.

Much of the technology industry remains focused on data. More recent advances in analytics, AI, semantic technologies, knowledge graphs, metadata management, and retrieval-augmented generation have shifted attention toward context. While important, context is not the endpoint. Enterprises ultimately operate on state.

Data describes what happened.

Context explains what matters.

State represents the current operational condition of a system, process, service, application, workflow, or business outcome.

It is state that determines what actions are appropriate next.

This distinction becomes increasingly important as enterprise workflows span multiple domains. Every operational domain maintains its own view of state, priorities, policies, dependencies, objectives, and operational conditions. As workflows move across applications, infrastructure, cloud services, observability platforms, data pipelines, business systems, service management environments, and AI-enabled technologies, state becomes increasingly fragmented across the enterprise.

The challenge is no longer simply coordinating tasks. It is coordinating state transitions across domains operating with incomplete visibility into one another.

AI intensifies this challenge. AI does not fundamentally change the need for orchestration. It changes how state is interpreted. Historically, enterprises relied on deterministic logic, business rules, policies, thresholds, workflow definitions, and human judgment to determine what should happen next. AI introduces a new mechanism for interpreting state, evaluating alternatives, and selecting actions.

As AI assumes greater influence over operational decisions, governance increasingly focuses on the state transitions those decisions create.

Observability reveals state

Intelligence interprets state.

Governance constrains allowable actions during state transitions.

Orchestration executes state transitions.

Observability, orchestration, governance, authority management, outcome assurance, federation, and AI governance are not separate problems that require separate solutions. They are all mechanisms for managing, interpreting, governing, coordinating, or validating state.

Viewed through this lens, the enterprise control plane is not fundamentally an AI architecture. It is a state coordination and governance architecture. Its purpose is to maintain awareness of state, coordinate state across domains, govern state transitions, constrain actions through authority and policy, and validate resulting state. AI increases the urgency of this requirement, but it did not create it.

The enterprise control plane coordinates and governs state transitions across systems, domains, workflows, and actors.

For EMA, this is the deeper significance of the survey findings. The enterprise control plane is not emerging because organizations need another orchestration platform, observability platform, or AI platform. It is emerging because enterprise operations increasingly consist of state transitions occurring across federated environments containing multiple systems of execution, multiple sources of intelligence, multiple governance domains, and multiple centers of authority.

Enterprises do not operate on data alone, or even context alone. They operate through a continuous series of state transitions. As AI expands autonomy and cross-domain complexity becomes the norm rather than the exception, maintaining awareness of state, governing interpretation, and coordinating state transitions becomes a strategic operational requirement. The enterprise control plane is the coordination layer that makes those transitions governable.

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Our team of analysts combines practical experience with a deep understanding of industry best practices and emerging vendor solutions to help clients achieve their strategic objectives. Learn more about EMA research, analysis, and consulting services at **www.enterprisemanagement.com** or follow EMA on **X** or **LinkedIn**.

2770 Arapahoe Road, Ste 132 - 544, Lafayette, CO 80026 +1 303.543.9500 **www.enterprisemanagement.com**

4753.070226