

The State of AI in ITSM

From Exploration to Enterprise Impact



Table of Contents

- 01 Executive Summary
- 02 **AI in ITSM Landscape:**
Where Are We Today
- 03 **What's Working:** The ROI
Case from Early Adopters
- 04 **AI Adoption Gap:**
Overcoming Barriers to AI
ITSM Adoption
- 05 **Strategic Priorities:** Where IT
Leaders are Investing Next
- 06 **From Insights to Action:** An
AI Readiness Framework
- 07 Conclusion

The Inflection Point is Already Here

AI in IT Service Management (ITSM) has moved past the question of whether it works. The question now facing IT leadership is how to move from cautious interest to confident execution, and how to do it before the window of competitive advantage closes.

This research surveyed IT professionals across industries, including professional services, financial services, healthcare, manufacturing, government, and higher education. The findings paint a picture that is at once encouraging and sobering: nearly every organization in the study believes AI will reshape IT Service Management, most are investing to make it happen, and yet only a small fraction have reached anything approaching scale.

The data reveals not a crisis of belief, but a crisis of execution. Understanding the gap between conviction and deployment is the central purpose of this report.

Key Findings IT Leaders Should Know

87%

Expect Production within 24 months

Among organizations with AI already in production, more than two-thirds report measurable reductions in average ticket resolution time.

99%

Resolution Time is Key

68% of organizations with AI in production have realized reduced ticket resolution time, with 99% citing this metric as a primary driver for adoption of AI in ITSM.

97%

AI is a Factor in ITSM Decision

While AI in ITSM adoption is still in early days, nearly all participants state that AI capabilities will influence their ITSM technology decisions with 59% citing that AI in ITSM in a major or critical factor.

55%

Still Researching or Piloting

The majority of organizations are engaged, but haven't reached widespread adoption, leaving room for increased ROI. The top two use cases are ticket resolution and improved customer satisfaction.

What This Means For IT Leaders

Organizations that believe in AI, but haven't yet built the data foundation, secured the executive mandate, or sequenced their use cases, are not behind. They are at precisely the right moment to learn from those who moved first. The question is whether they will act on it.



Data Readiness is Not a Precondition for Starting

Organizations waiting for clean data before beginning their AI journey are inverting the order of operations. Begin the pilot, identify the data gaps it reveals, and treat remediation as productive work rather than a prerequisite.



Executive Alignment Moves At the Speed of Evidence

Neutral leadership does not need a vision statement – it needs a proof point. Early pilots designed explicitly to generate executive – visible metrics are the fastest path from neutral to champion.



The Budget Conversation is the Wrong Conversation

Nearly every respondent in this study plans to increase AI investment. The real constraint is the business case, and CIOs who quantify the cost of inaction are consistently more successful at unlocking the resources they need.



The Platform Selection Decision is Already an AI Decision

With 59% of IT leaders treating AI capability as a major or critical factor in their next ITSM platform choice, organizations currently evaluating or renewing should treat native AI as a non-negotiable criterion.



Where AI in ITSM Stands Today

Before any organization can plan its next move, it needs an accurate picture of the current state of the market, not the version shaped by vendor roadmaps or conference keynotes, but the one reflected in how IT professionals are actually operating right now.

The landscape that emerges from this research is one of wide engagement and narrow deployment. AI in ITSM is not a niche conversation limited to technology-forward organizations with large IT budgets. It is active across professional services, financial services, healthcare, manufacturing, higher education, government, and beyond. What varies dramatically is not the level of interest, but the distance between where organizations are and where they need to be.

392

IT professionals surveyed
across 23 industry verticals

87%

Actively researching, piloting,
or deploying AI in ITSM today

13%

Have no current plans to
implement - a position the
data suggests will not hold

48%

Of executive leadership is
neutral on AI and awaiting
proof of value



Reading the Adoption Curve

The survey population maps cleanly onto a four-stage adoption arc. Understanding where the market sits collectively, and where your organization sits individually, is the starting point for any honest AI strategy.

Current Market Distribution Across Four Adoption Stages

17%

Exploring

No current plans; largely waiting for the market to mature or costs to fall

55%

Piloting

Researching options or running early pilots; largest cohort in the survey

28%

Deploying

Limited production deployment; AI is live in at least one use case

4%

Optimizing

Widespread deployment; AI is core to ITSM operations, not a supplement

The concentration at the Piloting stage is the most consequential feature of this data. More than half of all surveyed organizations are in a state of structured exploration, engaged enough to be testing AI, but not yet committed to a production deployment path.


This cohort will define the trajectory of the market over the next 18 to 24 months. How they move shapes everything that follows.

The 13% reporting no current plans deserve particular attention. In prior technology cycles, similar cohorts tended to be laggards by disposition. In this data, a significant portion appears to be organizations facing structural constraints, such as budget, compliance mandates, or legacy infrastructure, rather than fundamental skepticism.

With 59% of the broader survey population declaring AI a major or critical factor in their next platform decision, the "no plans" position carries meaningful long-term risk.

Where is AI Being Deployed?

Among the organizations that have widespread adoption, the capability picture is more developed. The headline finding that knowledge article suggestions and generation is the most widely adopted AI capability at 88%, reflects the importance of using AI to help remediate any data issues.




Knowledge article generation (88%) is particularly notable. It is less visible than a customer-facing chatbot (82%), but it consistently delivers faster ROI for IT by addressing a time-consuming task that often falls to the wayside for teams already stretched thin. This capability enables Knowledge Centered Service best practices, without creating more work for knowledge managers and technicians. Organizations currently piloting AI that have not yet evaluated knowledge management automation are leaving their highest-return use case on the table.

Intelligent ticket routing (71%) and AI-assisted agent responses (71%) are the capabilities most likely to be underreported. These are infrastructure-layer improvements. They make the system faster and the technician more effective, greatly improving the end users' experience.

Of note, many organizations that are operating in a "pilot" or "limited use" scenario are not actually deploying the full capabilities of AI. This is a critical factor when looking at return on investment, as these partial implementations will not have the same level of impact as a widespread deployment. Much of this is because when you use all the components together, the overall impact improves. Knowledge base article creation as a use case will improve the outcomes of Virtual Support Agents (VSAs) and ticket resolution suggestions, for example..

Why the Drivers Behind AI Investment Matter as Much as the Capabilities

When respondents identify their primary driver for AI adoption, the answer shapes not just what they build but how they measure success, how they frame the value to leadership, and which implementation challenges they are most likely to underestimate.



The dominance of end-user experience as a driver (47%) over cost reduction (16%) is a meaningful signal. It suggests that the most common internal AI narrative is not "how do we cut IT headcount" but "how do we provide service that employees actually want to use."

This framing is strategically important: it tends to generate broader organizational support, encounter less internal resistance from IT staff, and produce adoption metrics like satisfaction scores, self-service rates, and first-contact resolution that are both meaningful and executive-legible.

Organizations whose primary driver is cost reduction face a structurally different challenge: the ROI evidence comes slower, requires more baseline measurement discipline to demonstrate, and is more likely to be contested by finance teams who attribute savings to other factors. Neither driver is wrong. But understanding which one is actually operating inside your organization and building the measurement framework to match is essential groundwork before any pilot begins.

The organizations that will extract the most value from AI in ITSM **are the ones that move with the clearest sense of what success looks like and the ability to show metrics to prove it.**

The Platform Selection Signal

97% of respondents say AI capabilities will be a factor in their next ITSM platform decision. That finding, taken alongside the timeline data that 47% of organizations expect to be in production within 12 months, creates a clear and time-sensitive market dynamic: platform decisions being made today, and renewals coming due in the next one to two years, are already AI decisions, whether organizations frame them that way or not.

CIOs who approach their next ITSM evaluation without a structured AI capability assessment are setting their organizations back. The evaluation criteria that mattered most in the previous platform generation, things like service desk coverage, workflow configurability, and integration breadth, remain necessary, but AI capabilities should now be a critical part of the conversation.

Questions IT Leaders Should Answer Before Making An ITSM Platform Decision

✓ AI Architecture - Native, or Assembled From Parts You Manage?

- ☐ Which AI features work out of the box, and which require you to provision and manage your own LLM keys, vector database, or LLM subscription?
- ☐ When OpenAI or another model provider deprecates a model version, who is responsible for testing replacements and redeploying – the vendor or your team?
- ☐ Are AI suggestions for ticket category, priority, and routing generated from generic LLM inference, or from supervised ML trained on our organization's historical ticket data?
- ☐ Does the AI honor the platform's existing permission model when surfacing knowledge, related tickets, or report data through a virtual agent?

✓ Data Quality – Will It Work With What We Have Today?

- ☐ Does the AI require a curated, well-maintained knowledge base to be useful, or can it work with our actual mix of KB articles, resolved ticket history, uploaded documents, and SharePoint content?
- ☐ Does the platform help us identify and remediate data quality gaps as part of AI deployment, or is clean data a prerequisite we have to solve first?

✓ Automation – Where Are the Hard Limits on AI Driven Automations?

- ☐ When the AI identifies an action that needs to happen in another systems – HRIS, identity, CRM, monitoring, finance how can we configure that action with no code, or does it require coding or scripting?
- ☐ What happens if you don't have an out-of-the-box connector for a system you need?
- ☐ What are the hard limits on AI driven actions (automations)? Loops, concurrent runs, lookups, API calls per minutes, transactions per month – and what happens when we hit them?

✓ Compliance - How Will This Help Us Govern AI?

- ☐ How is AI activity governed, audited and partitioned – including credential management, role-based access, decision replay, and data retention controls?
- ☐ Can we configure data retention by data type (0-365 days), and does AI auditing with playback come standard or as an add-on?

✓ Platform Cost at Scale and Measuring ROI?

- ☐ What's the 3-year total cost of AI and automation at our expected volume – including consumption charges, Bring-Your-Own-AI infrastructure, and any add-on licensing?
- ☐ Do you currently track the metrics AI is designed to improve – resolution time, first contact resolution rate, ticket deflection, and cost per ticket?
- ☐ Does your current platform provide reporting granular enough to attribute outcome improvements specifically to AI?

The ROI Case: True Potential Realized with Widespread Adoption

Many AI capabilities are created to work together to drive continuous improvement and outcomes. **Most notable is the use of AI for KB content creation and gap finding.** When you look at organizations with widespread adoption of AI, the impact is profound. For those in pilot or limited use, ROI measured is only a portion of the true potential.

Measurable Outcome	Pilot Stage	Limited	Widespread
Reduced ticket resolution time	64%	77%	71%
Deflected ticket volume	52%	61%	82%
Improved satisfaction scores	59%	58%	76%
Better first-contact resolution	46%	60%	71%
Reduced IT support costs	40%	53%	65%

Percent of Companies Reporting Meaningful Impact By Stage

Skepticism about AI ROI is rational. Most enterprise technology investments that promise transformational outcomes deliver incremental ones, arrive late, and require far more integration work than the vendor's implementation estimate suggested.

AI in ITSM is not immune to those dynamics. But the respondents in this study who have deployed AI are not reporting aspirational outcomes. They are reporting measured results across five distinct areas, and the consistency of that pattern across different organizations, industries, and deployment models is the most credible signal in the data.

Two patterns in this data deserve particular attention from IT leadership.

First, reported outcomes improve as organizations move deeper into deployment. This is not simply a function of more AI in use. It reflects the compounding nature of ITSM improvements.

The reduction in resolution times frees technician capacity, which improves throughput, which in turn improves first-contact resolution rates and satisfaction scores. The organizations at widespread deployment are not just doing more with AI; they are operating within a fundamentally different service delivery model than the ones still piloting AI.

Second, cost reduction is the last metric to appear at scale. At the pilot stage, only 40% of organizations report meaningful cost impact, but by limited production, that rises to 53%, and by widespread deployment, 65%. For CIOs building the financial case for AI investment, this sequencing matters: the business case built on cost reduction alone will look weak early and strengthen over time. The smarter narrative leads with resolution time and satisfaction improvement, treating cost reduction as the lagging, but durable, confirmation of the investment thesis. **Even those in pilot stage are seeing real results.**

The Compounding Effect: Understanding How the Outcomes Connect

The five impact areas are not independent variables. They form a reinforcing system that accelerates as more of them come online simultaneously. For example, in a hypothetical mid-size environment processing 1,000 tickets per week, with an average resolution time of 4.2 hours and a first-contact resolution rate of 62%, you could see the following:

Sample AI in ITSM Outcomes: 1,000 tickets per week with an average resolution time of 4.2 hours and First Contact Rate of 62%

25%

Ticket Deflection

250 fewer tickets reach human technicians each week

~40%

Resolution Time Reduction

Technicians handle time can move from 4.2 to 2.5 hours.

+FCR

First-Contact lift

Better routing and suggested responses reduce repeat contacts

+CSAT

Satisfaction Improvement

Faster resolution and fewer repeat contacts drive satisfaction gains

↓ Cost

Support Cost Reduction


Capacity freed by deflection & efficiency becomes measurable cost avoidance

This cascade is why organizations that commit to full deployment consistently outperform those that treat AI as a supplementary layer on top of an unchanged service delivery model. The outcome is not the sum of individual capability improvements; it's the product of them.

Ticket Deflection: What Organizations Expect to Achieve


One of the most useful forward-looking data points in this study is the deflection expectation question: what percentage of tickets do respondents believe could realistically be deflected by AI-powered virtual agents in the next two years?

Expected Ticket Deflection Rate Via AI In The Next 2 Years



Nearly two-thirds of respondents expect between 10% and 30% deflection. This is a range that is achievable in the near term for most environments and represents a meaningful operational impact without requiring the organization to have solved every data quality and integration challenge first. The 24% who expect 30-60% deflection are likely organizations with high proportions of routine, repetitive ticket types (password resets, access requests, status inquiries) and can routinely achieve deflection rates in that range within 12 to 18 months of a well-configured virtual agent deployment.

The 9% expecting less than 10% deflection deserve particular attention. In some cases, this reflects environments with genuinely complex, non-routine ticket profiles where virtual agent deflection is structurally limited. In others, particularly in government and heavily regulated industries, it reflects compliance-driven conservatism about what AI can be trusted to handle autonomously. Both are legitimate positions. But for most ITSM environments, expecting less than 10% deflection suggests either an underestimation of what current AI capabilities can handle or an unresolved governance conversation about where autonomous resolution is permissible.



Building Blocks for Accelerated Adoption

Among organizations that have moved beyond the pilot stage, two challenges dominate post-deployment experience: integration with existing workflows (18%) and meeting accuracy and quality expectations (18%). These are essentially tied, and both require attention before deployment begins, not after.

18%

Integration with Existing Workflows

AI that operates outside native ITSM workflows creates parallel processes, adoption friction, and data inconsistency. Native platform integration is not a preference – it is an implementation requirement.

18%

Meeting Accuracy & Quality Expectations

Adopting AI will almost always show you gaps in data that need to be remedied in order to get the expected ROI. This should not be a roadblock to AI, but rather an opportunity. AI in ITSM should help repair the data.

12%

User Adoption & Acceptance

The most underestimated failure mode. Technically sound AI tools with low adoption deliver no ROI. Change management planning should begin at the same time as technical implementation, not after.

9%

Scaling Beyond Initial Use Case

The "pilot trap" – organizations that deliver a successful proof of concept but lack the architecture or governance to extend it. Use case sequencing and a scaling roadmap should be defined before the first pilot launches.

Evidence from Early Adopters is Clear

Organizations that **reach scale** aren't the ones that avoided these challenges. They are the ones who **anticipated and planned** for these challenges by adopting **ITSM with native AI**.

How Organizations Can Overcome Barriers to AI Adoption

The data reveals a striking paradox: there's a broad belief in AI's potential for IT Service Management, but friction stands between conviction and deployment. Understanding what's blocking progress is the first step to clearing the path.

96%

of organizations are missing the opportunity for exponential growth by not fully adopting AI ITSM.

62%

rate their organizational data as "fair" or "poor" for AI readiness

48%

report executive leadership is neutral and awaiting proof of value



Only 4% of surveyed organizations have reached widespread AI usage in ITSM. The majority of IT organizations are not disengaged from AI. They are actively engaged, actively piloting, and actively encountering the friction that separates proof-of-concept from production. The adoption gap is not a confidence problem. It is an execution problem, and it has identifiable causes.

The Budget Barrier is Real, But it's the Wrong Framing

Budget and resource constraints rank as the single largest barrier to AI adoption, cited by 26% of respondents. In government and higher education, that number climbs considerably, reflecting the structural spending constraints those environments face. But a closer look at the investment data complicates the budget narrative: every organization in this study is planning to increase AI spending over the next 12 to 24 months.

Not a single respondent reported plans to reduce investment.

This suggests that "lack of budget" is frequently a proxy for something else: the absence of a sufficiently compelling business case to unlock funds that already exist.

CIOs who have cleared this hurdle consistently describe the same approach, shifting the conversation from technology investment to operational ROI, and quantifying the current cost of doing nothing.

Security, Compliance, & the Accuracy Problem

Nearly 40% of respondents cite either security and compliance concerns (20%) or worries about AI accuracy and reliability (19%) as their primary barrier.

Security and compliance are governance and architecture challenges. They require clear data handling policies, a defined model for where AI operates within the service management environment, and alignment with legal and regulatory counsel. Most mature organizations are solving this through AI governance frameworks and vendor due diligence.

Accuracy and reliability concerns are a data quality challenge disguised as a model quality challenge. When AI produces poor outputs in an ITSM environment, the root cause is almost always insufficient or inconsistently structured training data, not a deficiency in the underlying model. **For this group, data readiness is the first use case for AI.**

Organizations that address data quality upstream consistently report faster time-to-value and fewer accuracy issues once deployed.

Barriers to Widespread Adoption: For those that have not met widespread adoption, where are the top challenges?

26%

Budget & Resources

Most frequently the symptom of a missing ROI narrative rather than absent funding

20%

Security & Compliance

A governance challenge – solvable by looking at native AI platforms with clear controls.

19%

Accuracy & Reliability

Waiting for perfect data is not the answer, using AI to help clean it will be the way forward

17%

Integration Complexity

Reflects the problems that come with using bolt-on vs. native AI in ITSM

Data Readiness: The biggest upside and your first use case

Of all the barriers surfaced in this research, data readiness may be the most consequential and the least discussed. Only 5% of organizations describe their data as excellent, meaning it's clean, structured, and accessible. Nearly half (47%) describe it as fair with significant quality issues to address, and 15% characterize it as poor, with major gaps in data collection and consistency.

Cross-referencing data readiness against adoption stage reveals a pattern worth noting: among organizations that have reached limited production deployment, 45% describe their data as good or excellent. This is roughly double the rate seen among organizations still in the research or pilot phase. Data readiness is not a prerequisite for AI success, but it is a leading indicator of it when organizations are starting AI pilots with virtual agents/chatbots, as 44% of respondents are.

Organizations with data quality issues will see better compounding success with AI if they instead start with knowledge article suggestions and intelligent ticket routing/classification to address any data gaps and build a solid foundation. If you treat your AI implementation as an exercise in remediating data quality, you'll see a greater return on investment faster.



The IT teams that will lead in AI-powered service desks over the next three years are not necessarily the ones with the largest AI budgets. They are the ones investing in the ITSM infrastructure with native AI and ecosystem integration.

The Executive **Alignment Gap**: Managing Upward Through Uncertainty

Perhaps the most politically complex barrier in this data is the one that doesn't appear on a technical roadmap. Nearly half of all respondents, 48%, describe their executive leadership as neutral, awaiting proof of value before committing meaningful support. Among organizations still in the pilot phase, that number reaches 45%, with a further 24% describing their executive environment as actively skeptical.

Compare that to organizations that have achieved widespread AI deployment: 73% of those respondents report a strong C-suite champion. The correlation is not subtle. Executive alignment is not a soft advantage — it is, statistically, the single most reliable differentiator between organizations that reach scale and those that remain in pilot indefinitely.

For CIOs navigating this dynamic, the implication is direct: the work of AI adoption is as much organizational as it is technical. Building a credible, evidence-based case for executive leadership grounded in peer data, early pilot results, and a realistic timeline is not a nice-to-have communication task. It is the critical path.

What Separates **Organizations That Scale** From Those That Stall

- A quantified business case tied to current operational costs, not aspirational efficiency language
- A plan to use AI in ITSM for data clean-up versus a barrier to adoption
- An executive sponsor with visibility into early pilot metrics – even before full deployment
- Use case sequencing that delivers IT-side value that can be measured with real metrics

Where IT Leaders are Investing Next

Survey data reveals a market on the cusp of an investment inflection, measured in its current pace but converging rapidly on a set of use cases, timelines, and capability priorities that will define the competitive landscape over the next twelve months.

Understanding where peers are directing investment, how quickly they expect to reach production, and which use cases they are prioritizing provides a benchmark against which every IT organization can assess its own roadmap and identify where it may be moving too slowly, or channeling resources toward lower-return use cases when higher-value opportunities are available.

46%

expect AI in production within 12 months - creating a near-term competitive divide

85%

plan modest investment increases - a market expanding steadily, not spectacularly

15%

plan significant investment increases - the vanguard that will set the performance benchmark



The Use Case Priority Stack: Top Investments

When respondents identify their highest-priority AI use case for the next 12 months, two capabilities pull well ahead of the field. Together, self-service virtual agents and intelligent ticket routing account for 61% of first-choice priority across the entire survey population. Understanding not just what these priorities are, but why they occupy the top two positions (and what the lower-ranked use cases signal about organizational maturity is essential context for any AI investment roadmap.

01

Self-service Virtual Agent/Chatbot

Top Priority

High visibility, end-user facing, executive-legible ROI through deflection metrics – the use case most likely to generate the proof of value needed to expand AI investment

Outcomes: Ticket Deflection & End User Satisfaction

02

Intelligent Ticket Routing and Prioritization

Lower visibility but faster time-to-value; reduces misrouted tickets, improves SLA compliance, requires less data preparation than virtual agents

Outcomes: Expedited Ticket Resolution, FCR, & End User Satisfaction

03

AI-assistance for Technicians

Suggested responses, next-best actions, contextual knowledge – delivers technician-side efficiency gains that compound over time as the model learns from resolved tickets

Outcomes: Faster Resolution, FCR, & End User Satisfaction

04

Knowledge Management and Article Generation

Consistently underrated relative to ROI delivered; addresses the highest time-cost activity for experienced technicians and feeds downstream AI capabilities

Outcomes: Faster Resolution & Accelerated AI Use Case ROI

05

Predictive Incident and Problem Management

The highest-maturity use case. Requires rich historical data and strong integration depth, but represents the greatest long-term strategic differentiation

Outcomes: Zero Touch Resolution Accelerates, Faster Resolution

The positioning of virtual agents at the top of the priority list is understandable. Virtual agents are the most visible AI capability to end users and senior leadership, which makes them an attractive choice for organizations seeking to demonstrate AI value quickly. But they are also the use case most sensitive to data quality and integration depth. Organizations that launch a virtual agent on a foundation of inconsistent ticket categorization or poorly configured knowledge bases frequently experience the accuracy and reliability challenges identified in Section 3, and those early failures have a disproportionate effect on organizational confidence and executive support.

For many organizations, intelligent ticket routing, ranked second overall, is the more strategically sound first deployment. It requires less end-user-facing infrastructure, delivers measurable SLA improvements quickly, and generates the clean routing and categorization data that a subsequent virtual agent deployment depends on. The sequence matters as much as the selection. Another key building block that is often ignored in initial pilots is the **knowledge management and article generation** which is the top actual outcome reported in section 3.

Getting Knowledge Management Right

Knowledge management and article generation, ranked fourth at 14%, is arguably the most under-invested category relative to its return. Article creation and maintenance is among the highest-burden, lowest-satisfaction tasks technicians and knowledge managers do. AI-assisted generation can reduce that burden substantially and improve knowledge base quality, with a caveat.

The AI used to do this needs to be more than just generative AI. Any tool you invest in should have a strong data integration and business logic layer that can be leveraged to create these articles within a templated structure you define for all knowledge base articles. Don't rely on a tool that offers only generative AI for knowledge creation – you won't get the results needed to make a meaningful impact.

A Note On Predictive Incident Management

At 10% of first-choice priority, predictive incident and problem management sits at the bottom of the near-term investment stack. This is rational as it requires the richest data history and the most mature integration architecture of any use case in the study. But the organizations currently investing in it are almost exclusively those already at limited or widespread production deployment. For IT leaders building a 24-to-36-month roadmap, predictive management should appear on the horizon now, even if it is not the immediate deployment priority. The data collection and categorization decisions made in the next 12 months will directly determine whether predictive capabilities are achievable in the following 6-12 months.

Investment Trajectory: Measured Acceleration, Not a Spending Surge

The investment picture that emerges from this data is one of deliberate, broad-based commitment rather than aggressive spending by a small vanguard. Eighty-four percent of organizations plan modest budget increases; 15% plan significant ones. No respondent plans to reduce AI investment.

The 15% planning significant increases deserve attention as a leading indicator. Cross-referencing against the primary driver reveals a clear pattern: organizations whose primary driver is resolution time improvement are significantly more likely to plan substantial increases (27% vs. the 15% average) than those primarily motivated by end user experience (9%).

This is likely explained by measurement clarity. Resolution time is precisely trackable, builds a compelling ROI case more rapidly, and gives finance stakeholders the confidence to approve expanded investment.

For CIOs whose budgets remain in the "modest increase" category, the implication is worth considering: the organizations now planning significant increases are likely to set the performance benchmarks against which the rest of the market will be measured within 18 months. Modest investment should be a deliberate strategic choice, not a default position.

The Production Timeline: A Market About to Accelerate

The timeline data may be the single most consequential forward-looking finding in this study. Taken together, 47% of organizations expect AI in production use within 12 months. Add the 40% expecting deployment in 12 to 24 months, and 87% of the survey population anticipates production AI within two years across the following time horizon:



The cross-tabulation between timeline expectations and current adoption stage adds important context. Among organizations currently in pilot or research mode, 47% still expect to reach production within 12 months — suggesting active, near-term deployment plans rather than open-ended exploration. Among those already at limited production, 49% expect to expand significantly within 12 months, indicating that partial deployment tends to accelerate rather than satisfy the appetite for broader AI integration.

For organizations that have not yet begun a structured pilot, these timelines define the competitive context. By the time a 12-month planning and procurement cycle concludes, a substantial majority of peers will be operating with AI in production. That is not an argument for rushed deployment. Poor implementations consistently outperform non-existent ones, but well-planned implementations outperform both. It is an argument for treating the planning process as already underway.

The Five Strategic Questions This Data Raises for IT Leadership

- 01** Is your highest-priority use case the one with the best risk-adjusted return, or the one with the highest executive visibility? Those are not always the same choice.
- 02** Does your current use case sequencing build toward predictive capabilities or does it risk locking you into a capability stack that cannot be extended without starting over?
- 03** If 47% of peers expect AI in production within 12 months, what is the service delivery gap between your organization and theirs – and is that gap acceptable to your end users and leadership?
- 04** Is knowledge management on your roadmap as a strategic investment, or only as a supporting task? The distinction determines whether your other AI capabilities compound or plateau.
- 05** What does "significant investment increase" look like for your organization – and what proof point from your current pilot would justify making that case to finance?

AI Readiness Framework for IT Leaders

The data tells a clear story about AI in ITSM adoption and the projected ROI. This section offers a practical framework for building your AI readiness and business case. Based on the patterns of organizations that have successfully navigated from pilot to widespread production, ROI is substantially increased when all the levers are working together.

Before introducing the framework, it is worth seeing what the practitioners have realized in terms of actual ROI by taking a look at some of the early adopters in the TeamDynamix customer community.

Early TeamDynamix Adopters of AI for ITSM are Seeing Real Value

“ We are seeing 20-25% more efficient triage through automation and intelligent routing with AI.

– United Federal Credit Union

“ AI has been a strong addition to our ITSM platform, accelerating the onboarding of new technicians.

– Vanderbilt University

“ We are deflecting 60% of our tickets with Virtual Support Agents tied to automation.

– Bowdoin

Remove Friction for Technicians

AI Service Assist takes the busywork out of IT by automatically suggesting field values, resolutions, related tickets and KB content, while even helping to create new articles on the fly. Technicians can move faster, stay in flow and spend less time wrestling with tickets.

Improve End User Satisfaction

Behind the scenes, AI agents accelerate every step of service, intelligently routing requests, surfacing the right answers and linking to relevant knowledge. Employees get accurate help in record time, with each interaction getting smarter as new agents come online.

Deflect Tickets with Virtual Support Agents

When you pair AI with automation and integration, you can leverage the AI Virtual Support Agents to actually take action. From password resets to onboarding to equipment replacement, you can easily resolve tickets without them touching the service desk.

The Four Levers of AI Readiness

Across the organizations in this study that have successfully moved from pilot to production, four distinct readiness factors consistently differentiate those that scale from those that stall. Importantly, none of them are purely technical. All four require deliberate organizational investment, and all four compound when they are addressed together.

01

Data Foundation

Clear, consistently categorized, historically complete ticket data gives AI the best foundation to build on. However, your data readiness shouldn't be a roadblock to innovation. **You can use your AI pilot to identify gaps in knowledge and leverage the AI for remediation.** Organizations that invest in categorization consistency and data cleanup during this implementation report fewer accuracy problems, faster time-to-value, and significantly higher confidence in AI outputs at the executive level. The 62% of respondents with fair or poor data readiness aren't blocked from starting. They are carrying a risk they need to actively manage from day one.

02

Native AI Architecture

AI that operates outside the native workflows of the ITSM platform – bolted on through middleware or accessed through a separate interface – consistently underperforms AI that is native at the platform level. The integration complexity barrier cited by 17% of respondents is frequently the consequence of evaluating AI capabilities in isolation from the platforms that house the data and processes they need to function. **Native AI is not a preference, it's a critical success factor.** For most ITSM AI use cases, native AI is the difference between a capability that works and one that is systemically flawed.

03

Use Case Sequencing

Not all use cases are created equal, and the order in which they are deployed has downstream consequences that most organizations do not adequately plan for. Virtual agents – the top priority for 33% of respondents – are the most executive-visible starting point but are also the most sensitive to data quality and integration depth. Organizations that begin with intelligent routing and knowledge management frequently discover that those deployments generate exactly the structured, well-categorized data that makes a subsequent virtual agent significantly more effective. Sequence is strategy. **When you combine AI capabilities for key building blocks such as KB content and rerouting with virtual agents, you accelerate the ROI.**

04

Change Management and Adoption Planning

Eleven percent of organizations that have deployed AI cite user adoption as their biggest implementation challenge – and that figure almost certainly undercounts the problem, since adoption failure is often interpreted as a technical quality issue rather than a change management failure. This concern surfaced directly in the open-ended responses: IT professionals worry about how their teams will respond to AI assistance. While there is a fear that the AI is perceived as a headcount reduction tactic, this is not playing out in reality. Instead, teams are able to offload mundane, repetitive tasks in favor of more impact driven work. The outcomes are faster resolution, reduced friction, and better service to end users. Addressing this narrative before deployment begins, not after adoption metrics disappoint, is the difference between a smooth rollout and an expensive proof of concept that never scales.

Use Case Sequencing Framework for AI Investment

With those four levers in place, the question becomes: where do you start, and what comes next? The following framework is built from the patterns in this study's data and reflects the deployment sequencing observed among organizations that have reached limited or widespread production deployment. It is a starting point, not a prescription. The right sequence for any individual organization will be shaped by its current data maturity, integration architecture, ticket volume profile, and executive priorities.

AI Deployment Sequencing Framework			Four stages: 0 to 36 month
Stage 1 Now Intelligent routing and classification Lowest data quality requirement. Immediate SLA and misrouting impact. Begins generating the clean categorization data downstream capabilities need. Why first: fast time-to-value, executive-legible SLA improvement, minimal end-user friction Outcomes: - Faster Resolution - Expedited Triage - Reduced Friction	Stage 2 3–6 Months Knowledge management and AI-assistance for Techs Addresses the highest technician time-cost activity. Builds knowledge base quality that virtual agent deployment depends on. Improves technician experience before end-user experience. Why second: technician-side ROI is measurable and credible; builds executive proof before public-facing deployment Outcomes: - Accelerates AI ROI - Faster Resolution - Improved FCR	Stage 3 6–18 Months Self-service virtual agent Highest executive visibility and deflection ROI. Requires the data quality and knowledge base completeness built in stages 1 and 2. End-user-facing deployment with 24/7 coverage potential. Why third: early-stage virtual agents on poor data foundations generate accuracy failures that erode confidence; the foundation must come first Outcomes: - Ticket Deflection - Faster Resolution - End User Satisfaction	Stage 4 18–36 Months Predictive incident and problem management Highest strategic differentiation. Requires rich historical data from stages 1–3. Shifts the IT organization from reactive to anticipatory service delivery. Why fourth: data collected in stages 1–3 is the prerequisite; decisions made now determine whether this is achievable in year three Outcomes: - Zero Touch Resolution - Ticket Deflection - End User Satisfaction

This sequence deliberately inverts the priority order that most organizations currently express. Virtual agents are the top stated priority for 33% of respondents; this framework places them third. That inversion is not a contradiction. It's a recognition that the organizations currently experiencing the accuracy and reliability problems cited in Section 3 are, in large part, organizations that deployed virtual agents before their data and knowledge foundations were ready to support them.

The organizations that will dominate the CX and efficiency leaderboards across the next 12 months are not necessarily the ones deploying virtualagents today. They are the ones doing the unglamorous foundational work, like routing improvements, knowledge base investments, and data cleanup, that make every subsequent capability exponentially more effective.

Building the Executive Case

A Three-part Structure

The 48% of organizations with neutral executive support are not waiting for a better AI product. They are waiting for a credible argument that the investment will produce measurable, attributable outcomes within a reasonable timeframe. The structure of that argument matters as much as its content. To build your case, you can utilize the following framework:

Part One	Part Two	Part Three
The cost of the status quo • Current weekly ticket volume and average resolution time • Estimated cost per ticket (technician time x average handle time) • First-contact resolution rate and cost of repeat contacts • After-hours coverage gap and associated risk or cost • Technician time spent on Level 1 vs. strategic work	Expected AI impact (peer-benchmarked) • 88% report ticket deflection as a primary outcome of implementing AI for the service desk • 71% of deployed orgs report faster ticket resolution • FCR improvement and associated satisfaction lift • 76% are experiencing improved customer satisfaction • 98% of orgs with ticket deflection are realizing rates of 10+% with 61% reaching 30%+	Continuous Optimization • 88% of orgs in deployment are leveraging AI for knowledge article creation and gap identification - this key step will allow for continuous optimization • Organizations are hitting challenges when it comes to enterprise integration - deploying AI along with automation and integration is a key success factor • Cost reduction is not necessarily measured by headcount reduction but rather by looking at speed of resolution - impacting overall staff productivity

The neutral executive does not need a vision statement. They need a proof point with a defined measurement date. Give them a number, a timeline, and a peer benchmark. Now, the conversation shifts from "should we?" to "how fast?"

The Gap Between Knowing and Doing is Closing

The organizations that will define the next generation of IT service delivery are making their foundational decisions right now. As a result, the window to learn from early adopters, rather than repeat their mistakes, is narrower than it's ever been.

87%

of organizations expect AI in production within 24 months - the market is not waiting

97%

say AI is an influence in their next ITSM platform decision, with native AI being a key differentiator

80%

of organizations at widespread deployment report meaningful ticket deflection

When this research was fielded, IT professionals across industries and organizational sizes described being at the inflection point of AI technology becoming operationally expected, of AI in ITSM moving from a differentiator to a baseline.

The data tells this story clearly. Ninety-seven percent of IT leaders say AI will be a major or critical factor in their next platform decision. Forty-six percent expect to be running AI in production within 12 months. Among those already there, 69% report reduced resolution times, 56% report measurable ticket deflection, and 80% of organizations at widespread deployment have demonstrated meaningful volume reduction through virtual agents.

The 96% of organizations that have not yet reached widespread deployment are not behind the curve by disposition. Many are managing real constraints, like budget cycles that move slowly, compliance environments that require careful governance, and technical debt that makes integration genuinely complex.

This research provides the evidence-based context needed to allow IT leaders to make genuinely informed decisions about the pace and sequencing of their AI investment, rather than defaulting to caution because the alternative felt too speculative to defend.

The neutral executive waiting for more evidence now has the evidence. The CIO without a pilot underway now knows that nearly half their peers will reach production before their procurement cycle concludes. The IT organization investing in virtual agents without a data or knowledge foundation now understands why so many early deployments encounter the accuracy and adoption problems they do.

AI is Reshaping the **Service Desk** at an Accelerated Pace

Based on the timeline data in this study, the next 6-12 months will see accelerated adoption and more discerning buyers as IT leaders move from early pilots to more widespread use of AI in ITSM. As that transition accelerates, three dynamics will reshape how AI is being evaluated for rapid time to value.

The service expectation gap will widen. End users in organizations with AI-powered ITSM will experience meaningfully faster resolution, 24/7 virtual support availability, and proactive communication on in-flight tickets. End users in organizations without it will increasingly notice the contrast — not as a technology observation, but as a service quality one. The pressure on IT leadership to close that gap will intensify as the use of AI ITSM grows.

AI Native ITSM with Integration Takes Center Stage

The platform decision will become the AI decision. With 59% of IT leaders already treating AI capabilities as a major or critical platform selection criterion, that proportion will only increase as early adopters demonstrate production outcomes. ITSM platform evaluations over the next 24 months will effectively be AI evaluations, and the organizations that have not yet developed clear criteria for what AI-native integration means in practice will be evaluating platforms without the most important variable on their scorecard.

AI is Being Used to Accelerate AI

Organizations that start building data quality now through the use of AI for categorization, routing, and knowledge management will have the foundation to scale quickly. By combining these capabilities with virtual agents, these IT teams will be consistently improving and accelerating ROI.

The compounding return on early, foundational investment is the most durable finding in this study.

A Final Observation

Every practitioner in this study who described the challenge they would most want AI to solve was describing a problem they already know how to solve — they simply do not have the capacity to solve it at the volume and speed their organizations require. AI in ITSM does not replace the judgment of experienced IT professionals. It removes the obstacles that prevent them from applying it.

The framing of AI as a capacity multiplier rather than a headcount substitute is not just the more accurate description of what the technology actually does in most ITSM environments. It is the framing that resolves the change management challenge, builds technician support for adoption, and reorients the executive conversation from risk to return.

Sequenced Use Case Planning Impacts Outcomes

The organizations in this study that are thriving with AI are not the ones with the largest AI budgets. They are the ones who defined what problem they were solving, built the foundation that the problem required, sequenced their investments in a rational order, and committed to measuring outcomes from the first day of deployment rather than retrospectively.

That approach is available to every organization in this study. The window to take it is open. What closes it is not the maturation of AI, but the increasing cost of starting later than your peers.



About the Study

We built this research to give IT leaders an honest picture of where the market stands, not a technology company's vision of where it should go. The findings here belong to the IT practitioners who shared their experiences, their challenges, and their aspirations for what AI can make possible in IT Service Management. Our role is to help organizations act on them with platform capabilities, implementation guidance, and a genuine commitment to the outcomes this data describes.

About TeamDynamix

TeamDynamix provides a modern, no-code platform that unifies ITSM/ESM with PPM and ITAM—so teams can manage services, projects, and assets in one place. Built on an AI-powered integration and automation layer, the platform helps organizations streamline service delivery and reduce manual effort. Customers report 30–60% ticket deflection, 40–90% faster resolution times, and the ability to eliminate two to three months of manual work for IT teams.

Designed to support real-world IT constraints, TeamDynamix delivers enterprise capability without overburdening teams or budgets. A no-code approach makes it easy to launch new workflows, forms, and automations quickly, while hundreds of prebuilt connectors simplify integration across the organization. And because the platform extends beyond IT, teams can support HR, marketing, legal, and other departments—without relying on heavy administration or specialized developers.

TeamDynamix is built for organizations that value agility, ease of ownership, and platform power—without complexity.

To explore how TeamDynamix approaches AI in ITSM, visit:

<https://www.teamdynamix.com/how-can-we-help/using-ai-itsm-for-faster-resolution>.



We are seeing 20-25% more efficient triage, and we are saving 5-10 minutes per ticket using AI in ITSM.

United Federal Credit Union



TeamDynamix