

2020 State of AI Report



Our research into AI implementation unveils the characteristics that define success. Learn from those who have mastered their AI projects about the best approaches to stakeholder roles, technologies and tools, talent sources, budget responsibility, and more.

By Sue Troy

TABLE OF CONTENTS

3	Author's Page
4	Executive Summary
5	Research Synopsis
6	The Basics
7	Figure 1: Job Level or Type
8	Project Benefits
8	Figure 2: Organizational Attitude Toward AI
9	Concerns About AI
9	Figure 3: Realized Benefits
10	In-Use and Planned AI Technologies
10	Figure 4: Concerns
11	Use Cases
11	Figure 5: AI Tech to Avoid

12	Figure 6: IT Use Cases for Better-Results Group
13	Personnel Sourcing
13	Skill Gaps
13	Project Characteristics
13	Figure 7: Departmental Identification
14	Stakeholder Roles and Responsibilities
14	Figure 8: Implementation Time vs. Expectations
15	Tools and Systems Used
15	Figure 9: Specialized Applications
16	Figure 10: Virtual Digital Assistants
17	Use of Professional Services Providers
17	Figure 11: Professional Services

Appendix

18	Figure 12: Company Size
19	Figure 13: AI Maturity Level
20	Figure 14: AI Project Role of Respondent
21	Figure 15: AI Funding Sources
22	Figure 16: Expected Benefits
23	Figure 17: AI Tech Plans
24	Figure 18: AI Use Cases
25	Figure 19: Operations Use Cases for Better-Results Group
26	Figure 20: Skills Gaps
27	Figure 21: Closest Affiliated Project
28	Figure 22: Results Compared With Expectations
29	Figure 23: Implementation Time
30	Figure 24: Project Costs vs. Expectations
31	Figure 25: Stakeholder Responsibilities
32	Figure 26: Development Platforms
33	Figure 27: Workload Location
34	Figure 28: Chipsets
35	Figure 29: Professional Services Firms Used By Better-Results Group



Sue Troy

Sue Troy is senior content director for ITPro Today. She has spent more than 25 years writing and editing IT-focused content, with long stints at companies such as TechTarget and Ziff Davis Publishing. Before joining Informa in May 2017, she was editorial director of TechTarget's CIO, IT Strategy and Channel Media Group, where she was inspired by emerging technologies such as IoT, blockchain and artificial intelligence.

AUTHOR'S
PAGE

EXECUTIVE
SUMMARY

IT practitioners with AI project experience have learned a lot from those projects. And those who have found higher levels of success can teach us all about how the projects should be handled. For this research report, based on a survey conducted with InformationWeek and Interop, we set out to identify those learnings.

Some of the highlights from our research make logical sense:

- Those with higher levels of success tend to have more applications in production.
- They are more likely to see themselves on the bleeding edge of AI implementation.
- They are more likely to say that their projects exceeded expectations.
- They are more likely to say their projects went more quickly than they expected, and their projects were in fact completed more quickly than those with lesser results.
- They're more likely to work for a large organization.

Other highlights may be more instructive, offering hints for best practices that other

IT shops can follow.

- Stakeholder responsibility in the projects of those with better results tends to be more evenly distributed among team members – residing with those with the most specific knowledge for a particular function – rather than being concentrated among just a few people.
- Those with better results are generally more discerning as to which technologies they will pursue or avoid.
- Those with better results identify more areas where they lack skills.
- Those with better results are more likely to lean on professional services firms.
- Those with better results have more sources for AI talent within the company.
- Those with better results have different habits with regard to project sponsorship.

Now let's drill down into the details around those findings.

RESEARCH SYNOPSIS

Survey name: State of AI Survey


Survey date: April 2020

Region: North America

Respondent base: 154 IT and business professionals at companies with at least one AI-related project in general production. The margin of error for the total respondent base (N=154) is +/- 7.8 percentage points. Because these findings reflect those who are currently engaged in AI-related projects, data is not reflective of AI engagement across all enterprises everywhere.

Methodology: InformationWeek and ITPro Today surveyed technology decision-makers at mostly North American companies to uncover the ways companies are approaching and implementing emerging technologies – specifically artificial intelligence (AI) and the Internet of Things (IoT) – in order to grow and get ahead of the competition.

The survey was conducted online; respondents were recruited via emailed invitations containing an embedded link to the survey. The emails were sent to a select group of Informa Tech’s database; Informa is the parent company of InformationWeek, Interop and ITPro Today. Nearly 90% of respondents have an IT or technology-related job function, such as application development, security, Internet of Things, networking, cloud or engineering. Just over half of respondents work in a management capacity, with titles such as C-level executive, director, manager or vice president. Fifty-seven percent are from large companies with 1,000 or more employees, and 19% work at companies with 100 to 999 employees. Informa Tech was responsible for all aspects of survey administration, data collection and data analysis. These procedures were carried out in strict accordance with standard market research practices and existing U.S. privacy laws.



Artificial intelligence is transforming companies around the world, making operations more efficient and delivering greater understanding of customers. To stay competitive, businesses must at the very least keep pace in the worlds of machine learning, deep learning, natural language processing, robotic process automation, analytics and other related technologies. But to lead in their markets, organizations need to go beyond keeping pace to mastering these technologies.

Easier said than done. Once you've identified a particular AI technology to investigate as a potential implementation candidate, a wide range of choices remain. What are the best practices for AI implementation? What processes

should you follow? Who are the stakeholders and what are their responsibilities? Should you use an integrator? Where does the budget come from? What types of software platforms are required?

Our research helps answer those questions.

Our 2020 State of AI Survey, conducted in tandem with InformationWeek and Interop, sought answers to a wide range of questions about AI implementation. This report – which is an adjunct to an Information Week report that takes a broader view of the survey results – focuses on the respondents to the survey who had at least one AI application in general production. To understand the lessons those teams learned, we sliced the data according to respondents' results with their

projects; those with excellent and very good results were differentiated against those with only good or fair results – toward the goal of understanding the habits and choices of those with the best results. Here is the analysis of that data.

The Basics

Among those survey respondents with at least one AI application in general production, those with excellent and very good results comprised 63.5% of survey takers (excellent results accounted for 22.5% of respondents and very good results, 41%). In the other group are those with good and fair results, collectively accounting for 36.5% of respondents (good results at 28.5% of

respondents and fair at 8%). No one with AI projects at production scale characterized their results as poor.

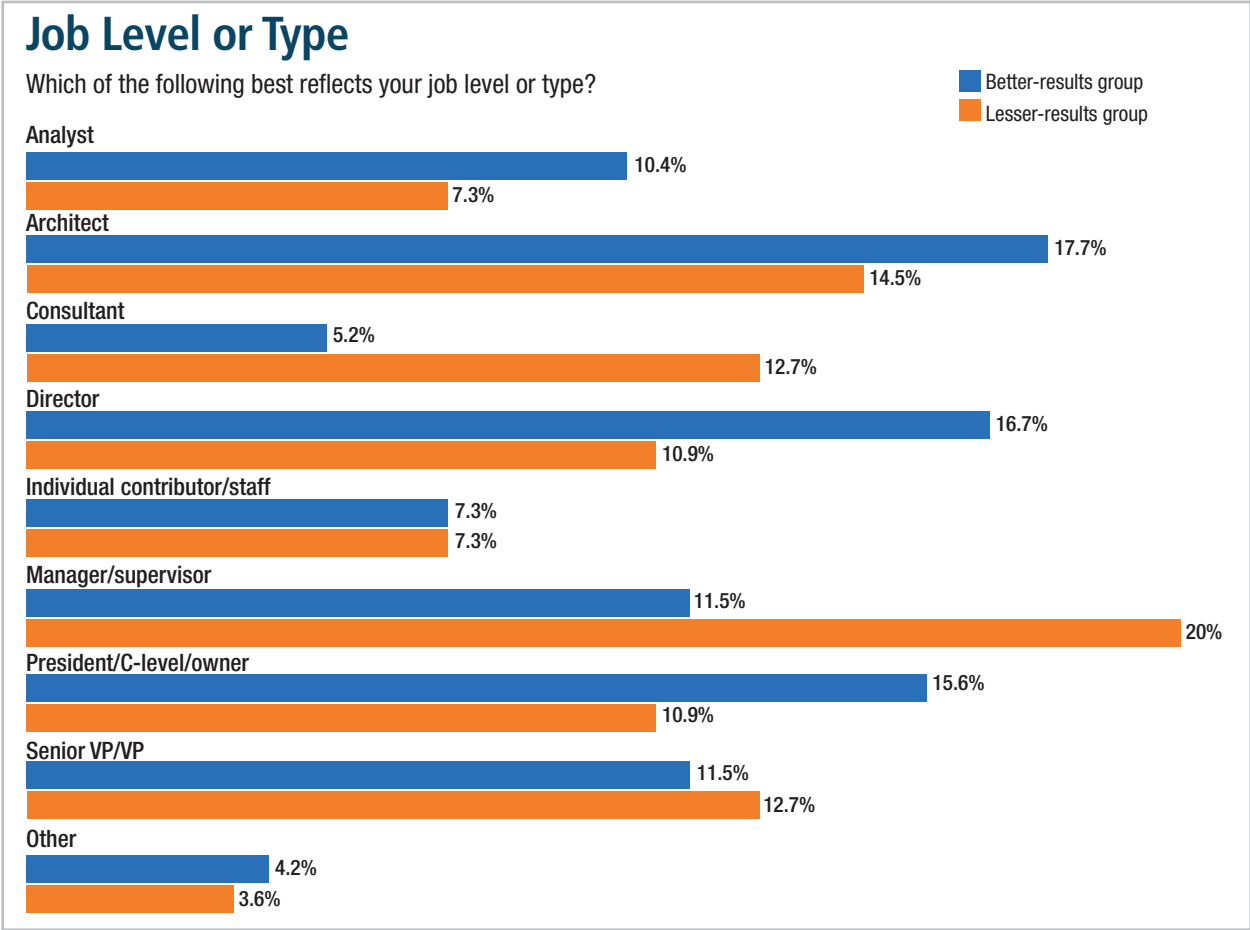
Among those with excellent or very good results, architects, directors and presidents/C-level/owners had outsized representation, accounting for 17.7%, 16.7% and 15.6% of that respondent base, respectively (see Figure 1). Among those with lesser results, those roles accounted for just 14.5%, 10.9% and 10.9%, respectively, of respondents — not huge deltas but deltas worth noting.

In considering company size of respondents, we found that those with better results were more likely to work for larger organizations (see Figure 12). Sixty percent of those with excellent or very good results hailed from companies with a thousand or more employees, versus 52.7% of those with lesser results. Large companies are generally better able to secure resources for emerging tech projects — and more likely to funnel those resources toward projects that, like AI, hold promise of competitive advantage.

Among those with excellent or very good results, almost two-thirds have multiple applications in general production, versus the remainder, with one application in general production and possibly others in development (see Figure 13). Among those with lesser results, the number of applications at production scale (multiple versus just one) is more evenly split. This distinction likely reflects the fact that experience with a higher volume of projects leads to better results. People learn from their mistakes and improve their approaches as they put more AI projects under their belts.

Among those with better results, respondents were more than twice as likely to say that they held the role in any of their organization’s AI projects of project leader

FIGURE 1



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

(22.9% vs. 9.1%) or senior IT exec (24% vs. 10.9%) compared with respondents with lesser results (see Figure 14). This is not necessarily the same as their role within the organization; rather, it strictly refers to

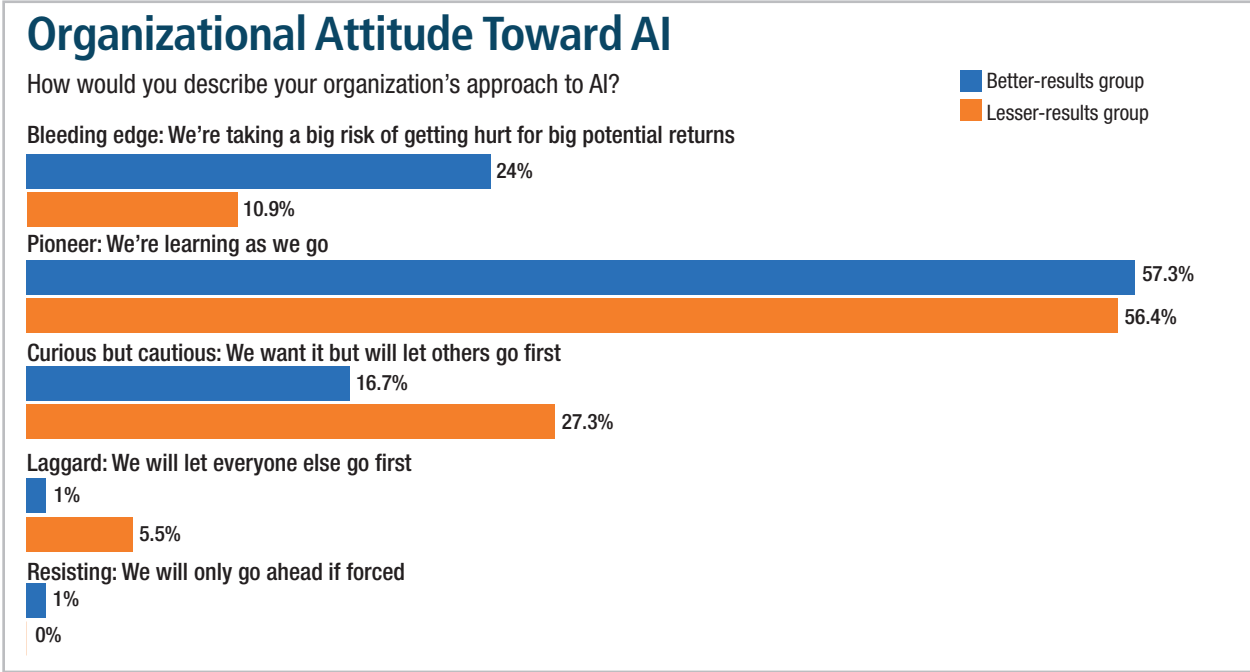
stakeholder roles within AI projects. Among those with lesser results, consultants had an outsized representation compared with the better-results group (20% vs. 9.4%); thus, internal ownership of AI projects correlates to more

favorable results. There are a few possibilities to explain those findings: Those companies with excellent or very good results may simply be more likely to have pushed past the need for consultants because of the volume of projects already under their belts. A more cynical take is that consultants might be able to help make AI projects happen but not necessarily happen well. No matter the reason for the differing influence of consultants, it seems clear that the roles of project leaders and senior IT execs can have a large positive impact on AI projects.

As you might expect, those with better results are more likely to consider themselves on the bleeding edge (“we’re taking a big risk of getting hurt for big potential returns”) than those with lesser results, cited by 24% of respondents in the better-results group vs. 10.9% in the lesser-results group (see Figure 2). This aligns with our survey finding that those with better results are more likely to have a greater number of AI projects at production scale. The largest chunk of both cohorts, however, described their organizations as pioneers (“we’re learning as we go”) — close to 57% in both groups. While it might seem incongruent that those with better results and more applications under their belt would describe themselves as pioneers at the same rate as those with lesser results and fewer applications under their belt, when you consider the fact that many survey respondents didn’t qualify to be included in this report, all who did qualify can be considered pioneers. To be included in this report, survey respondents had to be engaged in AI, and they needed to have at least one application in general production.

Those with lesser results were more likely than those with excellent or very good results to characterize their

FIGURE 2



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

organization as curious but cautious (“we want it but will let others go first”), at 27.3% vs. 16.7%.

One important question in any AI project relates to who will foot the bill. At businesses with excellent or very good results, 64.6% of respondents said that the funds come from the IT budget, while 57.3% cited business unit budgets and 47.9% cited allocations from corporate (see Figure 15). Among those with lesser results, a somewhat lower percentage (58.2%) said that IT paid for the projects, and a much lower percentage (40%) said that business units pay for them, while 49.1% cited allocations from corporate.

Project Benefits

When it comes to benefits of AI projects, the biggest difference between the better- and lesser-results groups was what they hoped for in new-product development (see Figure 16); those with lesser results were more likely to say they hoped for new products (36.4% vs. 24%). In terms of the top benefits, those with better results were most likely to cite greater efficiencies within IT operations (noted by 62.5% of respondents with excellent or very good results) and improved product support and customer experience (noted by 45.8% of that same group). The

former benefit was also the top choice among those with lesser results but to a lesser degree, cited by 56.4%. Improved product support and customer experience was cited by a higher percentage of the lesser-results group than the better-results group, so taking that plus the lesser-results group’s greater interest in new-product development suggests that they tend to think of AI benefits more in terms of external, customer facing projects than the better-results respondents.

In looking at the same factors from the perspective of benefits actually achieved (rather than simply expected), fewer respondents have experienced — than hoped for — greater efficiencies within IT operations, among both groups of respondents (see Figure 3).

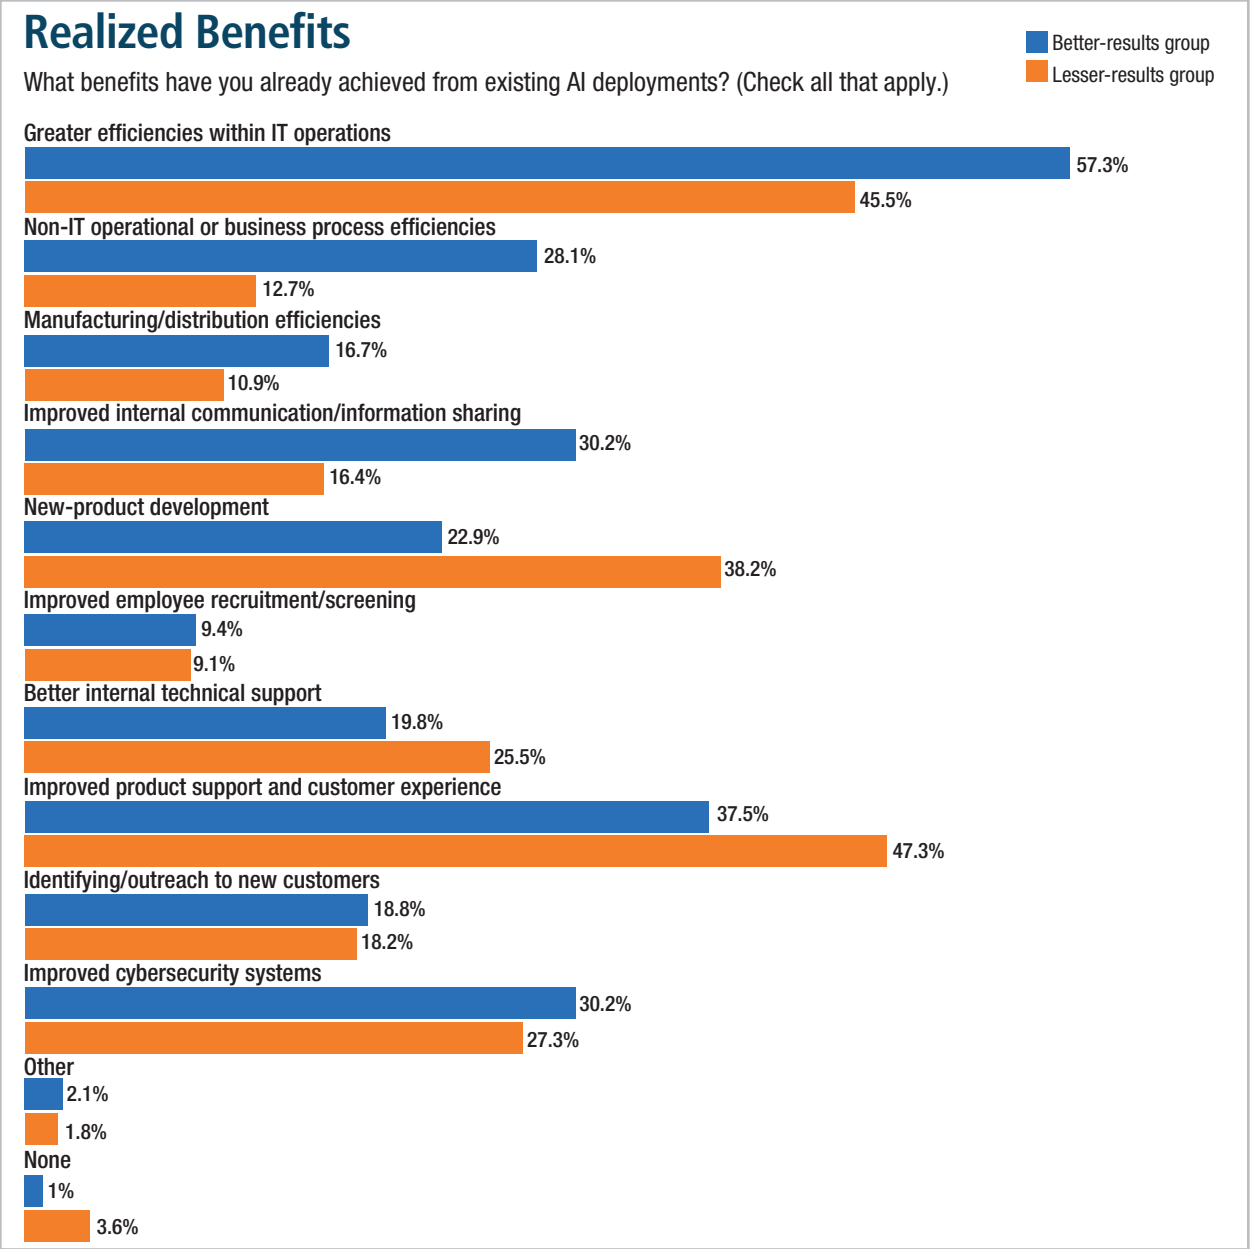
There are two areas where the opposite is true: where respondents across both groups have already experienced specific benefits but have a markedly lower level of hope for such benefits going forward. Those areas are improved employee recruitment/screening and better internal technical support. This suggests that those areas represent functions where AI is nearer to fulfilling its potential or nearer to meeting customer needs than other benefit areas. Or, perhaps a more negative take is that these benefits, once achieved, have turned out to be less of a win than originally anticipated.

In all, those with better results have achieved more benefits than they hope for, whereas among those with lesser results, the opposite is true: That group hopes for more benefits at a higher level than they’ve achieved them.

Concerns About AI

AI projects face obstacles in a number of areas, among them concerns the organization may hold about the technology. Although respondents in both the better-results and lesser-results groups cited security as their top organizational concern, a greater percentage of respondents who had lesser AI results stressed it (see [Figure 4](#)).

FIGURE 3



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

**LOOKING FOR MORE
RESEARCH ON AI?**
*Get AI insights from our sister
business Omdia, a research
firm addressing the entire tech
ecosystem.*

Forty-four percent of those with excellent or very good results cited security concerns, vs. 54.5% of those with lesser results.

What’s the reason for the lower concern about security among companies with better results? Perhaps those with lesser results have heard alarm bells around AI security but haven’t had enough experience with the tech (owing to the lower number of projects under their belt; see [Figure 13](#)) to either gauge their risks realistically or acquire sufficient knowledge and skill to offset their security concerns.

Conversely, a greater percentage of respondents in the higher-results group cited model transparency – or the degree to which the inner workings of AI algorithms are visible to users of the technology — as a top concern: 35.8% of that group vs. only 25.5% of the lesser-results group. Model transparency is an especially thorny issue. A high level of transparency can help mitigate bias and promote trust of the system, but it carries concerns that model explanations can be hacked, making the tech more vulnerable to attack.

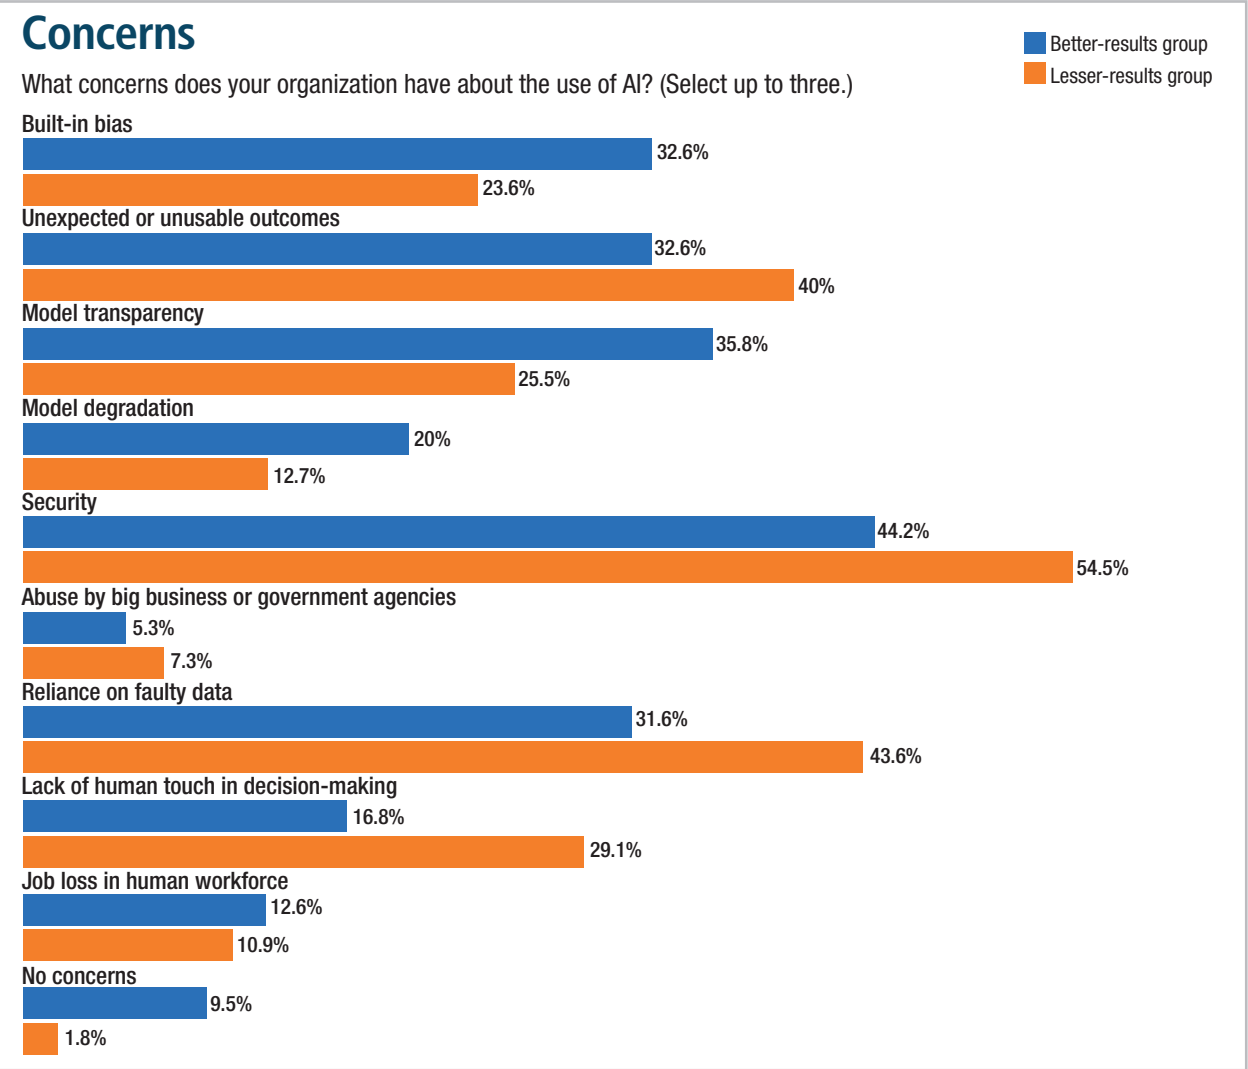
Those with better results are also significantly less likely to see organizational concerns around reliance on faulty data and lack of human touch in decision-making. (Reliance on faulty data was in the top five concerns for those with excellent or very good results, but a full 12 percentage points lower than those with lesser results.)

Interestingly, those with better results were more likely to say their organization had no concerns about AI tech (9.5% vs. 1.8%). Again, it seems that greater experience (in both quantitative and qualitative senses) has helped smooth out the wrinkles and quell fears in the AI project process.

In-Use and Planned AI Technologies

When we asked respondents about the AI-related technologies that their organizations will be incorporating over the next two years, the

FIGURE 4



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

“In general, those with excellent or very good results were more likely to avoid AI technologies than those with lesser results.”

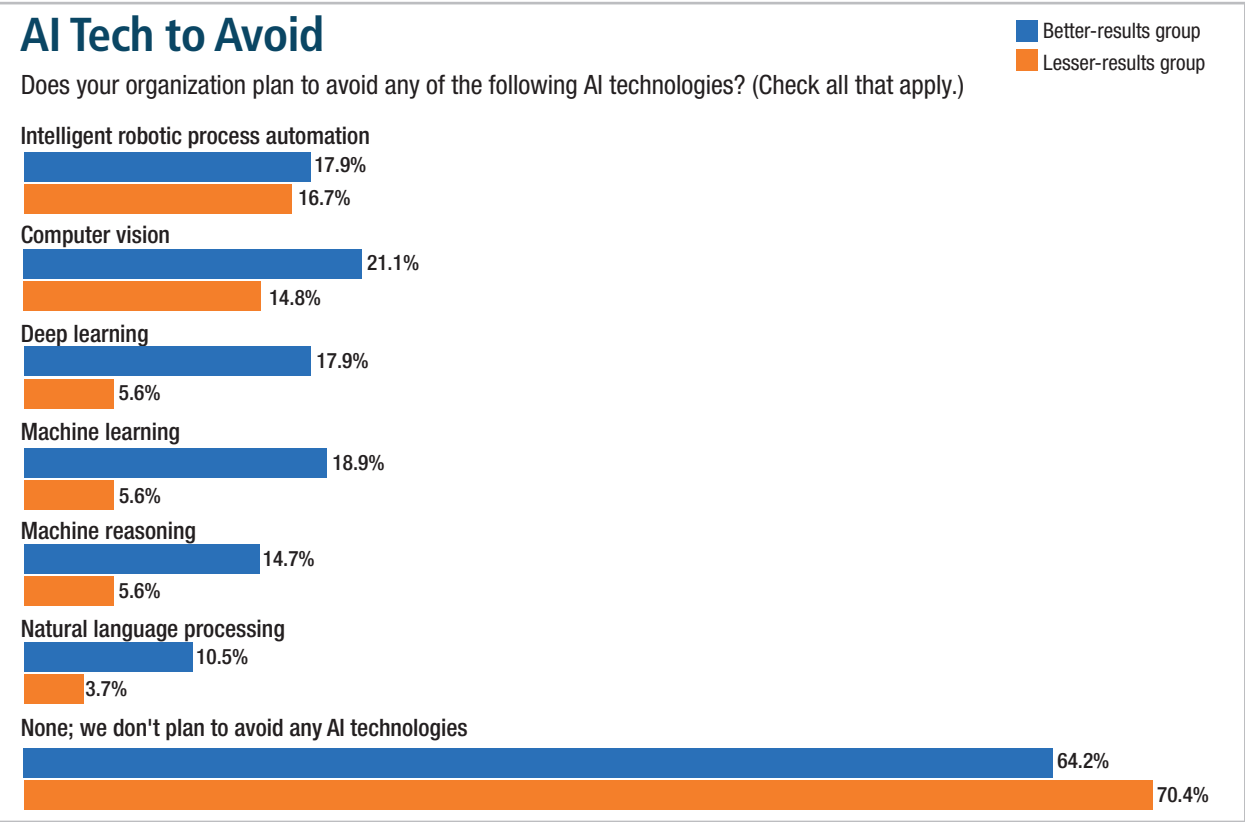
biggest delta between those with better results and those with lesser results was around intelligent robotic process automation (see [Figure 17](#)). Fifty-two percent of those with excellent or very good results said they would be implementing intelligent RPA vs. only 31% of those with lesser results – among the higher deltas between the two groups in our survey. Intelligent RPA is a bit like robotics process automation on steroids. Whereas RPA automates discrete tasks as a series of steps, intelligent RPA executes an entire business process with multiple, interrelated and successive tasks, incorporating interactions with the various applications a particular function relies on, manipulating unstructured data and communicating with users via chatbots along the way. Intelligent RPA could have a major impact on the labor market with its promise to significantly reduce costs and easily scale labor.

Intelligent RPA wasn’t the tech most frequently cited by those with better results (it was just a few percentage points behind machine learning and deep learning), but those other technologies had much smaller deltas with the lesser-results group.

AI Tech to Avoid

We asked respondents whether their companies were planning to *avoid* any particular AI technologies, which elicited some interesting results (see [Figure 5](#)). In general, those with excellent or very good results were more likely to avoid AI technologies than those with lesser results. For instance, 21.1% of those with better results were planning to avoid computer vision tech, vs. only 14.8% of those with lesser results. The deltas were more dramatic with deep learning (17.9% vs. 5.6%), machine learning (18.9% vs. 5.6%), machine

FIGURE 5



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

reasoning (14.7% vs. 5.6%) and natural language processing (10.5% vs. 3.7%). This further corroborates our findings that those with better results are more likely to have identified which AI technologies can most benefit their organization and have become more selective and focused on those technologies. They’re further along on the path to AI maturity.

Use Cases

We found that across both the higher-performing and lesser-performing groups, customer service, IT, operations and data management were all in the top four use cases (see [Figure 18](#)). The order shifted a bit between the two groups. Those with better results were

more likely to cite IT as the top use case (cited by 65.6% of those respondents, about nine percentage points above the lesser-results group) – which correlates with the finding that greater efficiencies within IT operations is the top benefit already achieved among those with better results. Those with lesser results, meanwhile, were more likely to cite operations as the top use case (cited by 61.8% of that group of respondents, about 10 percentage points higher than the better-results group).

We received enough responses to drill down into the use case results only for IT and operations cited by those with excellent or very good results.

Among those with better results, the most frequently cited IT-related use case (see Figure 6) was security analytics and predictive intelligence (at 71.4%), followed by help desk (55.6%), productivity (54%), knowledge management (52.4%) and DevOps (50.8%). According to VynZ Research, spending on AI-driven security is expected to reach \$30.5 billion by 2025, with a CAGR of 20.5%. While critics say that the [hype around AI-driven cybersecurity is overblown](#), clearly, IT departments are desperate to solve their cybersecurity problems, and, judging by this question in our survey, many of them are hoping AI will fill that need.


On the help desk, meanwhile, AI tools are using predictive analytics to improve decision-making around incident management and demand planning. And AI is being used for help desk chatbots and intelligent search recommendations.

Among those with excellent or very good results, the top operations-related use case was predictive

FIGURE 6

IT Use Cases for Better-Results Group

For which IT use cases does your company use AI-related technologies? (Check all that apply.)



Respondent count: Better results, 63
Base: Respondents with at least one AI application in general production, with excellent or very good results, and citing use of AI for IT use cases

maintenance, cited by 54% of those respondents, followed by inventory and supply chain optimization, and manufacturing analytics, each cited by 50% of those respondents (see [Figure 19](#)). Coherent Market Insights expects the [global predictive maintenance](#)

[market](#) to reach \$1.14 billion by 2027, with a CAGR of 22.4% between now and then, partly driven by acceleration of IoT technology. Using a cloud-based deployment model, it promises to reduce maintenance costs and downtime, primarily in manufacturing.

“Those with excellent or very good results were twice as likely as those with lesser results to say that the data science or AI personnel come from an external consulting group (29.5% vs. 15%).”

Personnel Sourcing


We asked respondents where the data science or AI personnel worked in their organization (see Figure 7). Their answers to this question were revealing in a few ways. First, the biggest delta in the two groups of respondents was in their citations of external consulting groups. Those with excellent or very good results were twice as likely as those with lesser results to say that the data science or AI personnel come from an external consulting group (29.5% vs. 15%).

But beyond that, those with better results were more likely than those with lesser results to report data scientists/AI personnel coming from any source, whether from central IT, within an independent group, within non-IT business units or from an external consulting group. In other words, those with better results have more sources for data science or AI expertise to work on their AI projects; there’s a 25% delta between the higher-performing and less-performing groups. This is something organizations should pay attention to. Though it might seem to be a no-brainer that additional data science expertise is helpful in an AI project, the data here confirms that a greater number of sources for data scientist or AI personnel is correlated with a beneficial impact on AI projects.

Skills Gaps

When it comes to identifying skills gaps within their organizations, those with better results were more likely in general to report gaps. Machine learning and data modeling was cited as a gap by 67.4% of those with better results, followed by data engineering (51.6%), compute

FIGURE 7



infrastructure (42.1%) and data governance (36.8%). This compares with 63.6%, 32.7%, 36.4% and 32.7%, respectively, among those with lesser results (see [Figure 20](#)). It stands to reason that those with more experience with a technology have a better sense of just where their weak spots are than those just getting started. The implication of that is that businesses with less experience with AI don’t yet understand what their skills needs will be and as a result they may be underestimating, which might perpetuate their lower results.

Project Characteristics

We asked respondents which use case was associated with the project they’re most familiar with. The top choice among both sets of respondents was IT. Almost 30% of those in the better-results group cited it, versus 23.4% of those with lesser results (see [Figure 21](#)). We then drilled down into characteristics of the project they were most familiar with. For instance, we asked respondents how the results of that project aligned with their expectations (see [Figure 22](#)). Here we saw a stark – though not surprising – contrast

Our sister site [AIBusiness.com](#) – the online companion to [The AI Summit series of events](#) – covers artificial intelligence across a range of vertical industries and for a variety of use cases. Get news, analysis, whitepapers, webinars and other sources there.

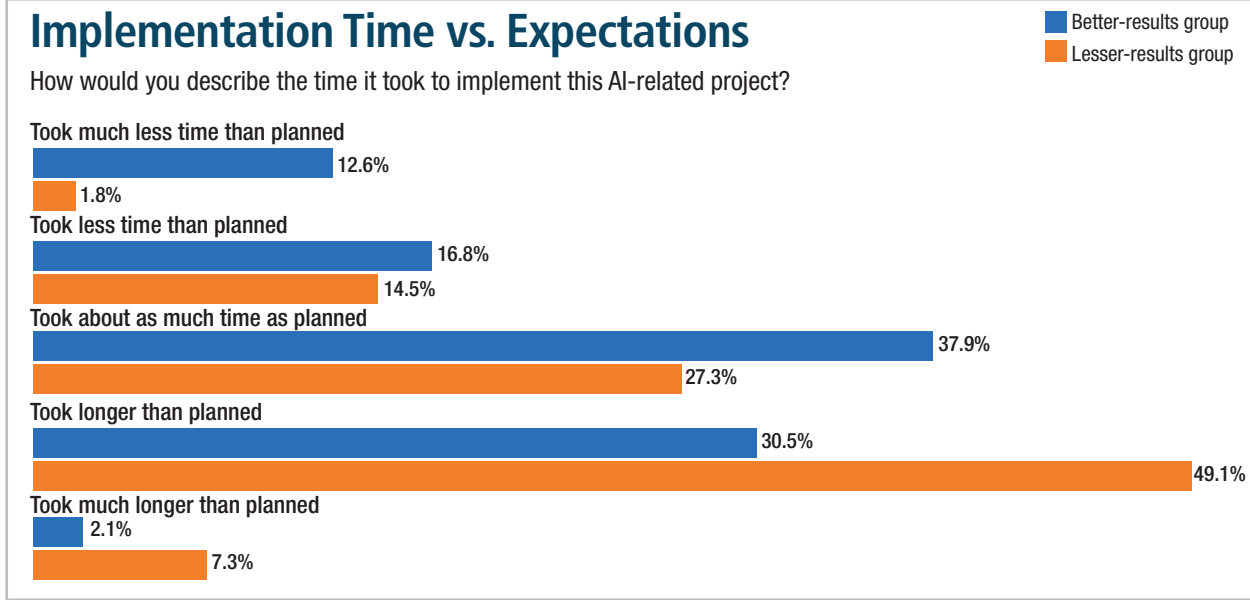
between the two groups. Those with better results were much more likely to say that the project far exceeded (25% of those respondents) or somewhat exceeded their expectations (33.3%). And those with lesser results were much more likely to say the projected simply met expectations (52.7% of those respondents) or even fell somewhat short of expectations (27.3%).

Likewise, in terms of how closely the project hewed to time expectations, there was a stark – but again, not surprising – contrast between the two groups (see Figure 8). Those with better results were much more likely to say that the project took much less time than planned (12.6% of those respondents), less time than planned (16.8%) or about as much time as planned (37.9%). With more projects under their belt, those with better results were able to more accurately predict how long a new project might take – and set project timelines accordingly. On the other hand, those with lesser results were much more likely to say that the project took longer (49.1% of those respondents) or much longer than planned (7.3%).

When we looked into the time for project completion, we again found a correlation between those with better results and a more favorable time frame (see Figure 23). Among those with better results, more than two-thirds (67.7%) said the project took them less than a year. Among those with lesser results, only 43.6% said the same. The upshot is those with lesser results are taking longer to get those lesser results. This is a symptom of their lack of experience with AI projects. With experience (and success) comes speed.

Again, when looking at the costs of AI projects and how

FIGURE 8



Respondent count: Better results, 95; lesser results, 55
Base: Respondents with at least one AI application in general production

closely they adhere to planning, we found a correlation between better results and costs that fell within planning (see Figure 24). Among those with better results, 73% said that the project either cost much less than planned, less than planned or about as much as planned. Among those with lesser results, only 43.6% said the same. The same theme as above holds true here: The more experienced IT practitioners are with AI, the better able they are to project costs and avoid going over budget.

Stakeholder Roles and Responsibilities

The most common division of stakeholder roles among those with excellent or very good results was line-of-

business specialist as primary sponsor; line-of-business leader as budgetary approver; IT leader as technology approver; data scientist as architect/specifier; and IT staff as implementer (see Figure 25). Among those with lesser results, the most common division of stakeholder roles was line-of-business *leader* as primary sponsor, budgetary approver *and* technology approver (so placing a larger burden of and responsibility for the project on the line-of-business leader) and IT staff as both architect/specifier and implementer — again, concentrating responsibilities within a single role.

The differences between these two sets of cohorts suggests that tasks and responsibilities should be widely

“Among those with better results, the architect/specifier role was most frequently handled by a data scientist, whereas those with lesser results put IT staff in the architect/specifier role more frequently than those with better results.”

distributed across a stakeholder group, and they’re best distributed to those with the most specific knowledge of an area. Data scientists have the most intimate knowledge of the data being used in an AI project, so they are well placed to specify the requirements of AI technology. The success or failure of an AI project is likely to have a direct (and likely day-to-day) impact on line-of-business specialists, which makes them naturally suited to the primary sponsorship role. And IT leaders, of course, should be entrusted with technology approval decisions. The data bears that out.


Line-of-business specialists, meanwhile, are the most likely stakeholder to be the idea originator among both those with excellent/very good results and those with lesser results — and were cited by the same percentage of respondents from both groups. That suggests that the role of the person most likely to bring forth an AI project idea does not have a bearing on the success or failure of the project.

However, those with better results were more likely to have ideas bubbling up from a wider group of people. And those with lesser results were more likely to say that a particular stakeholder was not involved in the AI project at all. There was a higher level of nonparticipation among every stakeholder role in the lesser-results group. Those with lesser results were also more likely to say that the implementer role was more heavily covered by a variety of stakeholders — in particular, data scientists, IT staff and consultants. This suggests that those with lesser results tend to overweight the implementation stage of the process. Among those with better results, the architect/specifier role was most frequently handled by a data

FIGURE 9

Specialized Applications

Does your company use software developed by AI solution (specialized application) providers?



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

scientist, whereas those with lesser results put IT staff in the architect/specifier role more frequently than those with better results.

This division of stakeholder responsibilities in AI projects is something businesses should pay close attention to. We recommend modeling project responsibilities along the lines of those used by the better-results group.

Tools and Systems Used

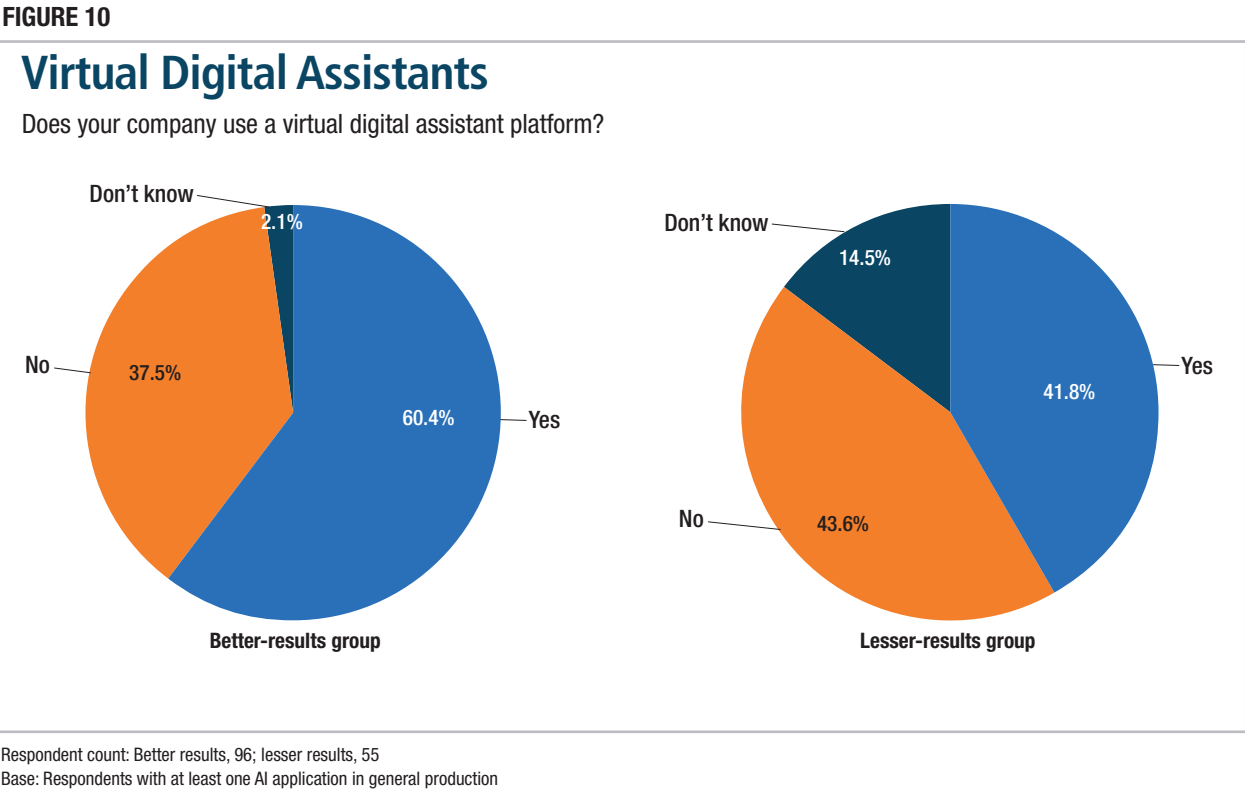
We looked at the types of tools and systems used

by survey respondents. With regard to specialized applications for AI, almost 60% of those with better results reported using them, vs. only 41.8% of those with lesser results (see Figure 9). Analyst firm Tractica (now under the Omdia umbrella) has tracked dozens (though many more exist) of AI application solution providers, including companies such as Conversable, Deep Genomics and Sana Labs. According to the firm’s [2018 AI market ecosystem report](#), the category includes AI solutions for enterprise end-user organizations, often

with a particular vertical or use case focus.

Virtual digital assistant platforms, meanwhile, focus on understanding natural language in spoken or written form. These tools can enable customer self-service access to a business’ knowledge base, for example, or provide speech-based interaction for customers. Vendors in this space include LogMeIn, Synthetix and Smart Action. In our survey, there was again a sharp contrast between the two sets of respondents (see Figure 10). Sixty percent of those with better results reported using a virtual digital assistant platform, versus only 41.8% of those with lesser results. The results for this question are a bit curious, given that only 37.5% of those with better results reported using natural language processing technology. There’s a pretty big gap between those using virtual digital assistant platforms and those using natural language processing. The disparity suggests that IT practitioners aren’t yet familiar enough with the underpinnings of virtual digital assistant to know that it’s fundamentally an NLP tech.

Another contrast exists with regard to use of an AI development platform. Almost 76% of those with better results report use of such a platform, versus 52.7% of those with lesser results (see [Figure 26](#)). According to Tractica, AI development platforms “typically take the ‘pain points’ out of the upfront AI processes by providing scalability, frameworks and automated model development.” Vendors in this space include Alteryx, Cognitive Scale and Petuum. The fact that there’s a significant delta in use of AI development tools between the better- and lesser-results respondents reflects the greater AI maturity level of those with better results. As



AI projects grow, businesses are more likely to take on development work internally.

When we looked into where AI infrastructure is running (see [Figure 27](#)) — whether in the cloud, on premises or both — we found that those with better results were less likely to say it was solely in the cloud (28.1% versus 43.6%, respectively). Those with better results were more likely to say they were running AI workloads on premises (32.3%) or both on premises and in the cloud (39.6%).

We also asked about specialized AI chipsets such as CPUs, GPUs, FPGAs, ASICs and neuromorphic chips (see [Figure 28](#)). Those with better results were much more likely to make use of such hardware than those with lesser results (53.1% versus 23.6%, respectively). Nvidia dominates the GPU market, but there are more than a dozen vendors in the overall AI chipset marketplace,

including big brands like Arm, Intel and Qualcomm as well as lesser-known companies like Graphcore and Mythic.

Use of Professional Services Providers


The use of a professional services provider appears to be associated with better results (see Figure 11). While the same percentage (47.9%) of respondents from that group said they used a professional services provider as those who said they did not use one, it was in contrast with those with only good or fair results. Among those with lesser results, only 34.5% said they use a professional services provider. Does a professional services provider contribute to positive results of a project, or is it more simply a characteristic of efforts that have seen success and been rewarded with money to offload some of the work to a specialized professional services provider? Or is it because businesses with higher levels of AI success are weighted more heavily toward larger organizations? All three factors could be in play here.

Among those with better results, the most frequently cited professional services firms were Accenture, PwC, Cognizant and Deloitte (see [Figure 29](#)). (Note that we had too little data from those with lesser results to report on their choice of specific professional services firms.)

FIGURE 11


Professional Services

Do you use a professional services provider for AI services?



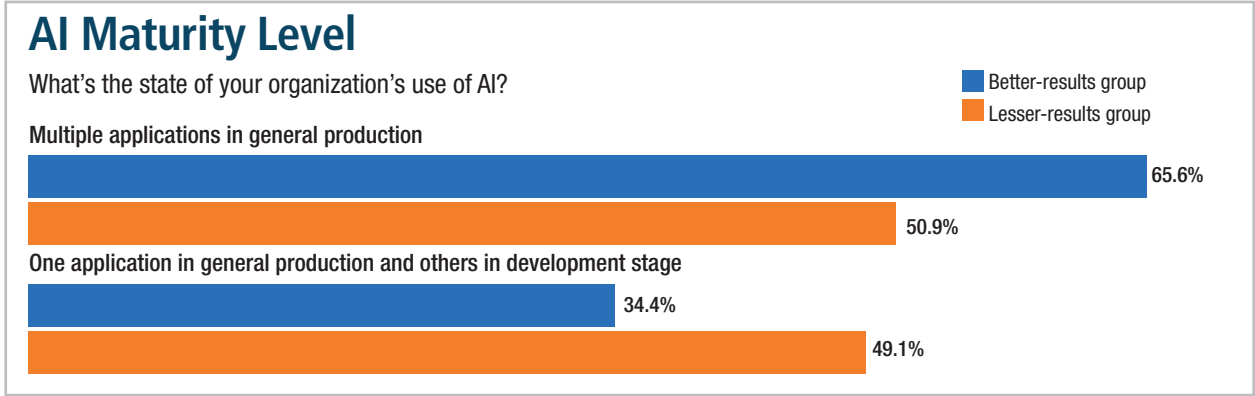
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 12



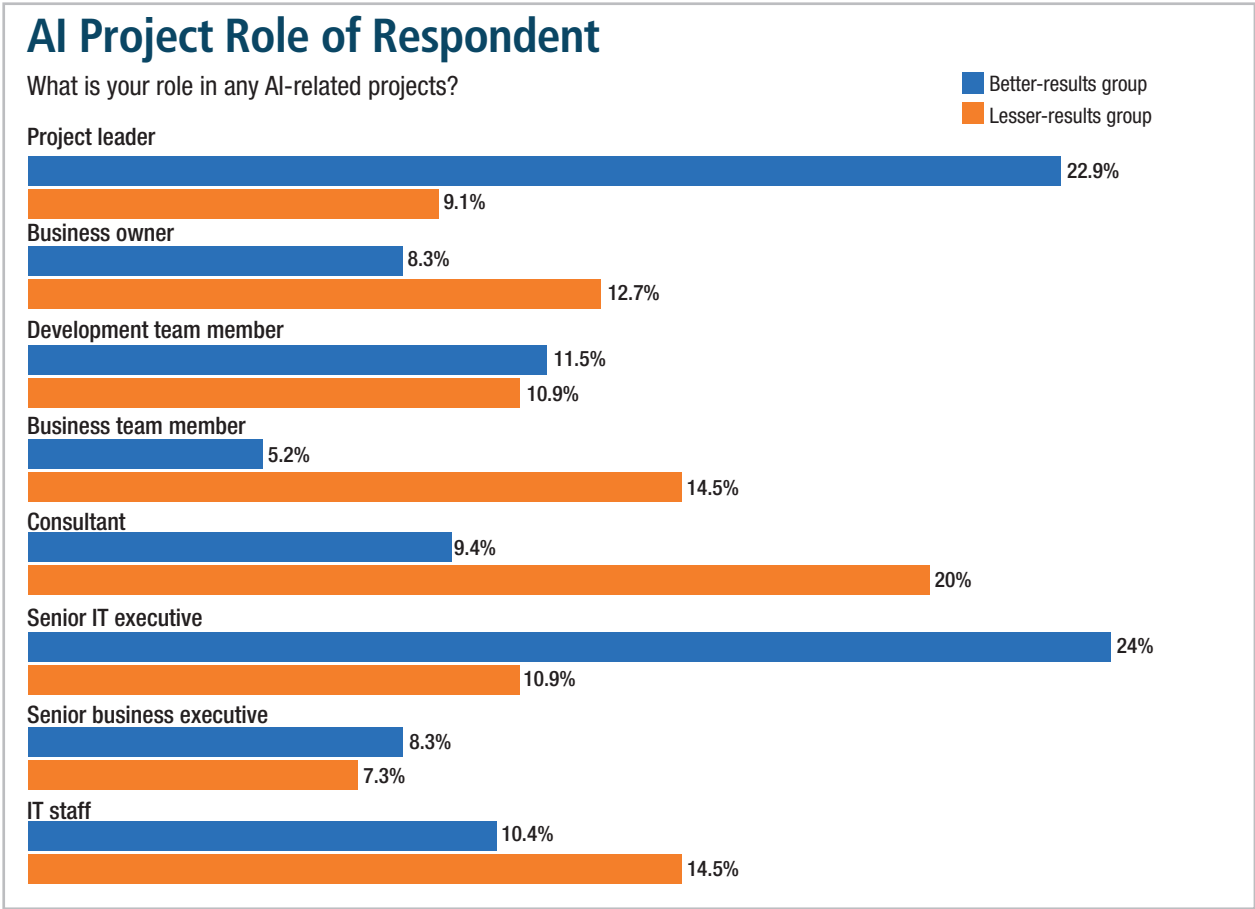
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 13



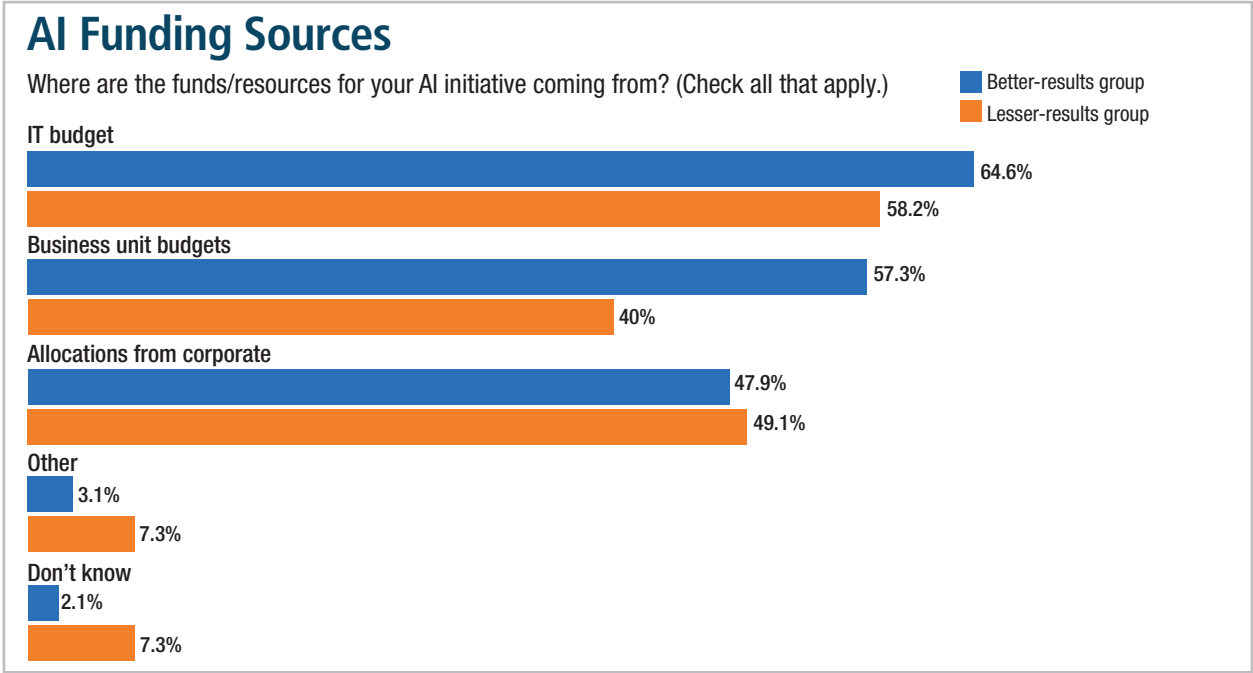
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 14



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 15




Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 16

Expected Benefits

What are the most important benefit(s) you hope to gain from an AI initiative? (Select up to three.)




Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 17

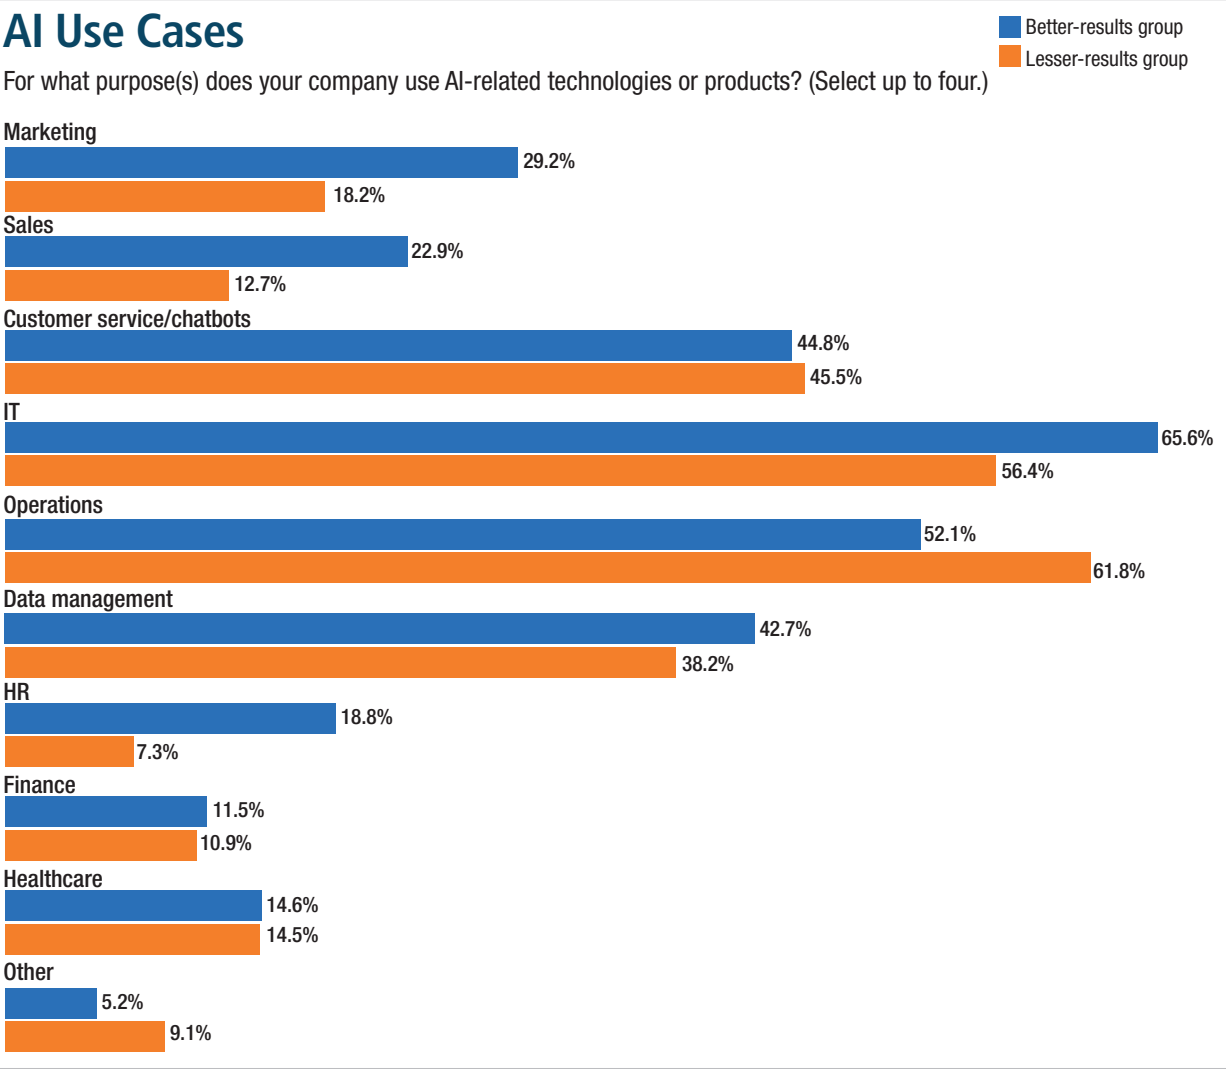
AI Tech Plans

What AI do you expect to incorporate into your workplace in the next six to 24 months? (Check all that apply.)



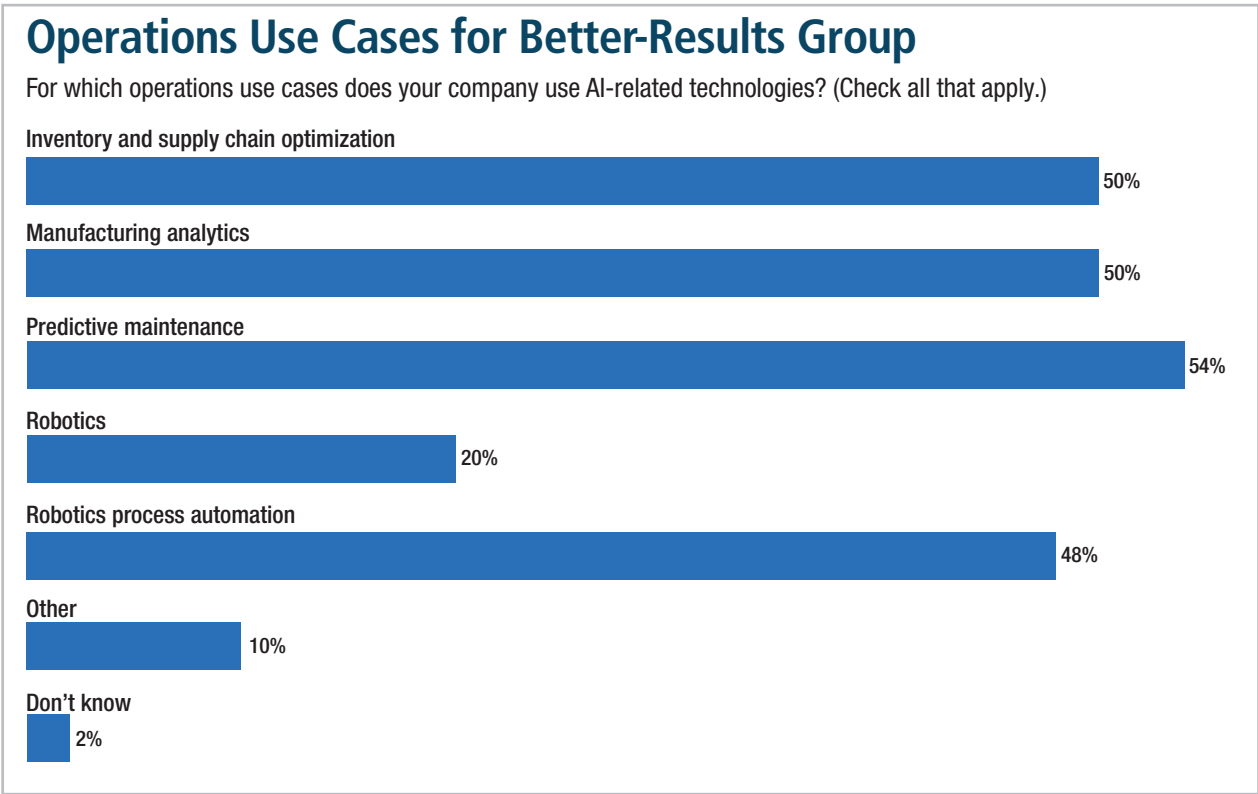
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 18




Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 19




Respondent count: Better results, 50
Base: Respondents with at least one AI application in general production, with excellent or very good results, and citing use of AI for operations use cases

FIGURE 20



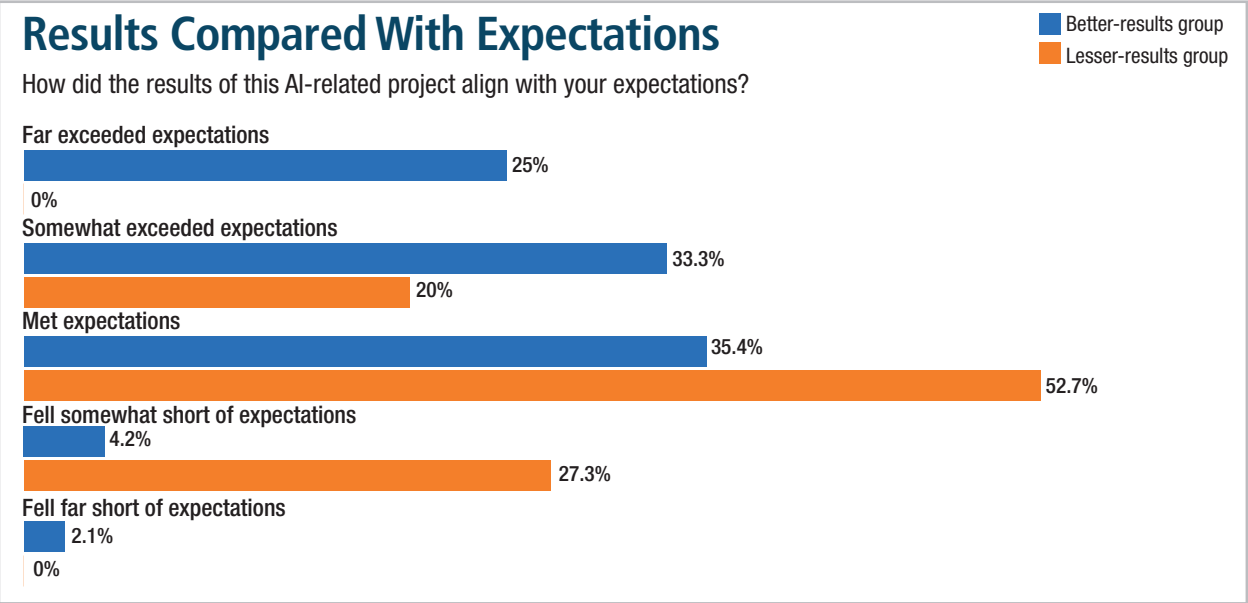
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 21



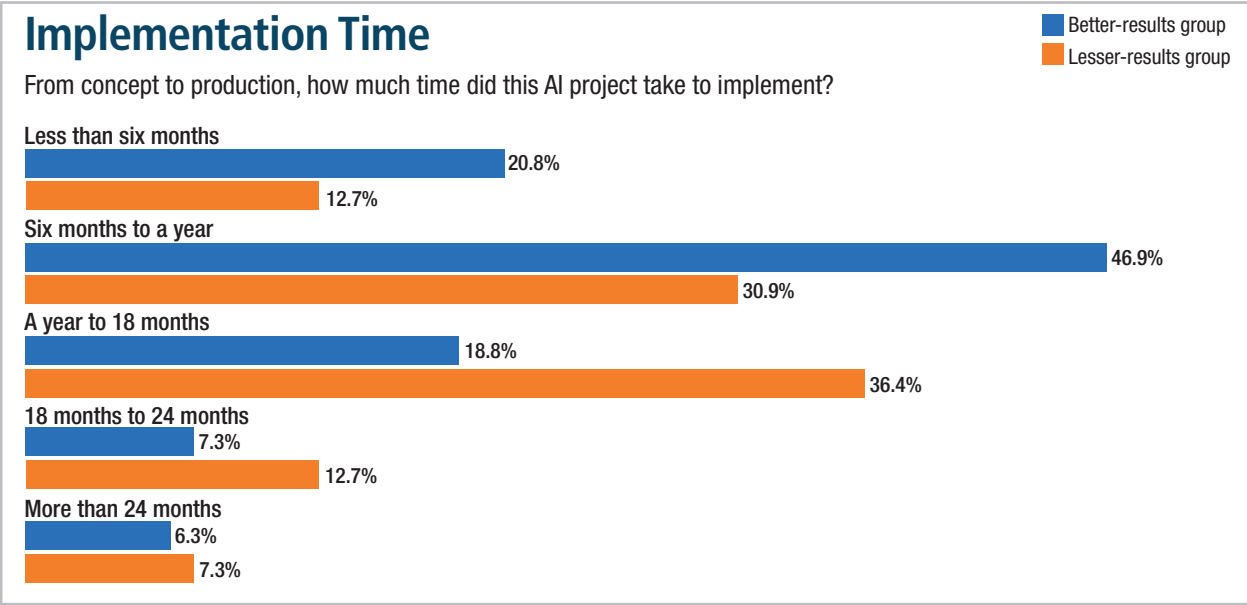
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 22



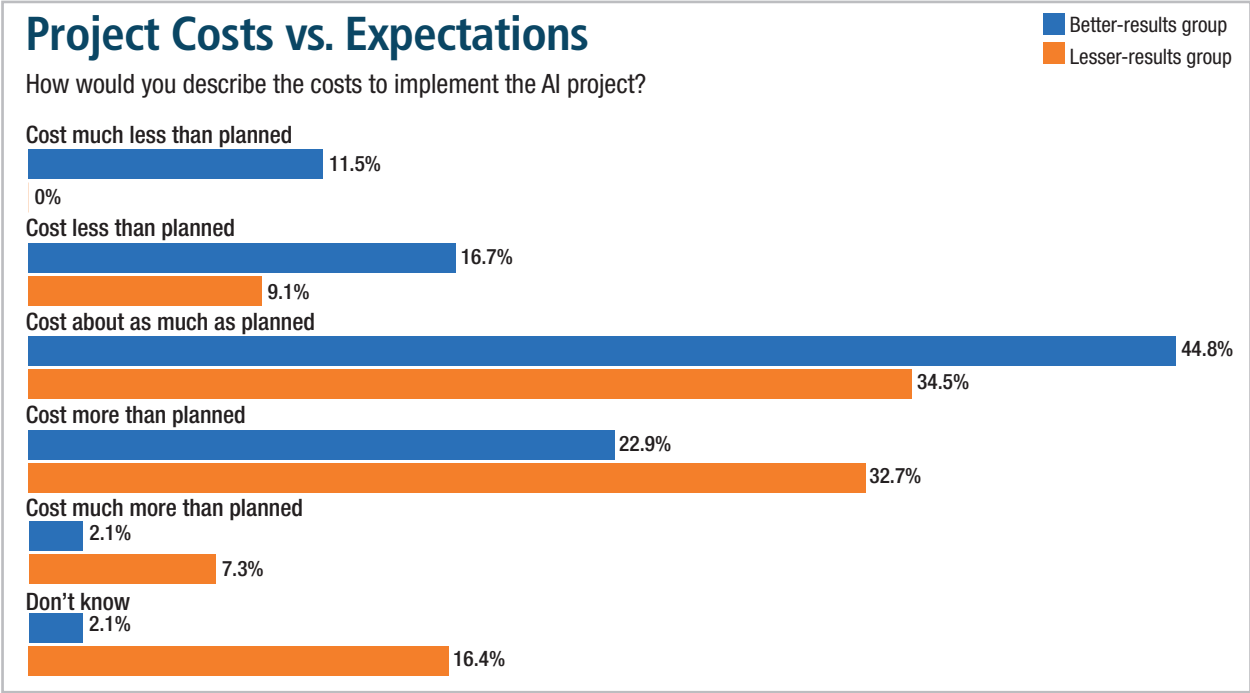
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 23



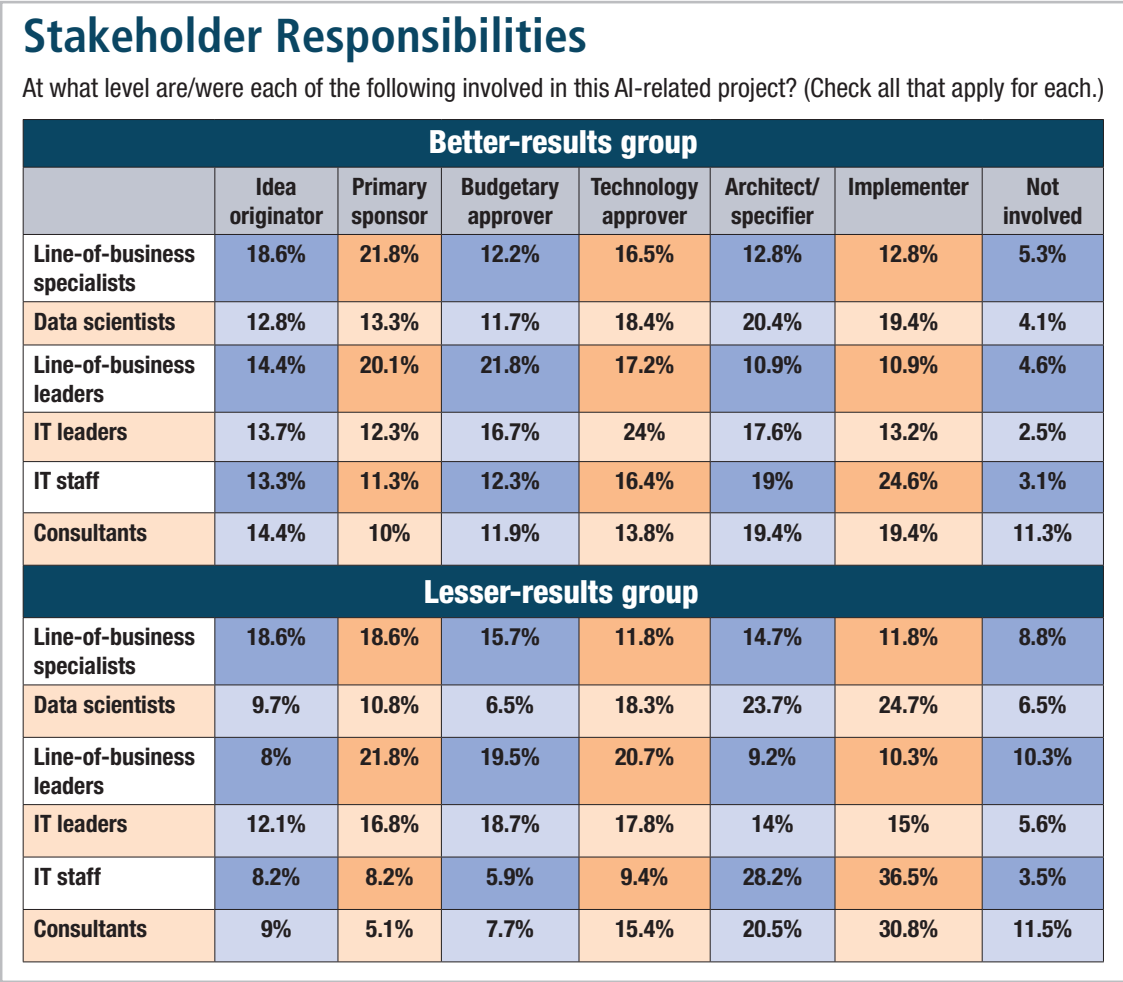
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 24



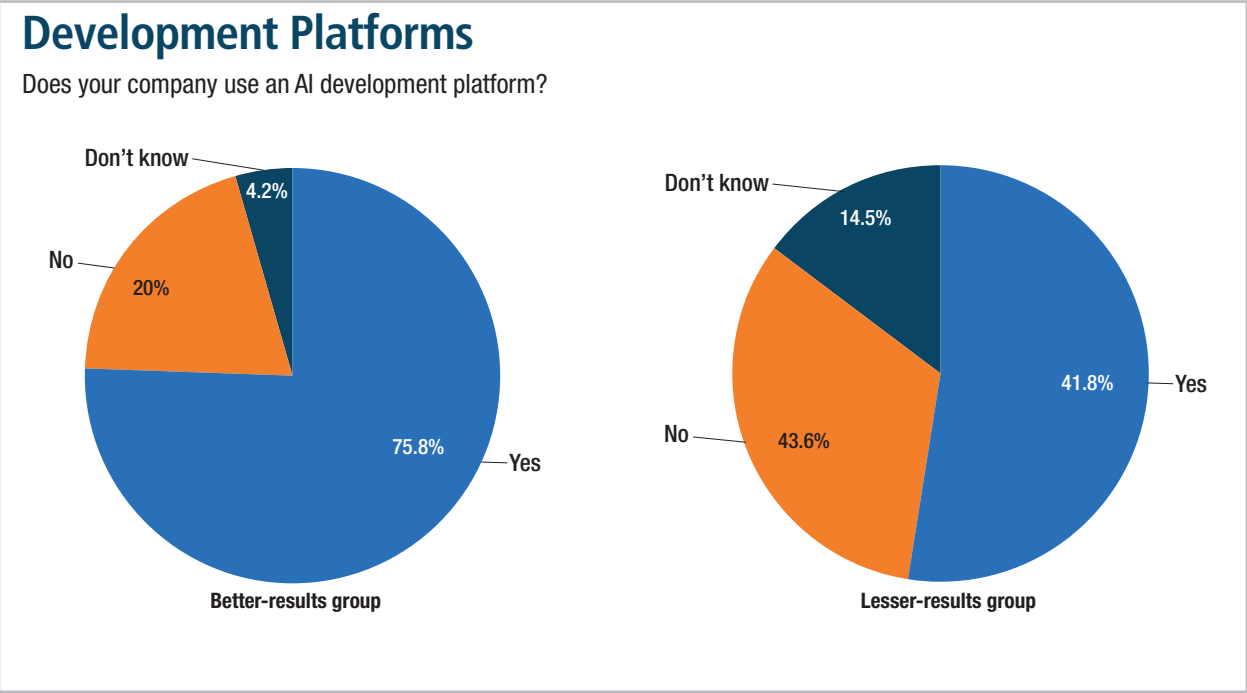
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 25




Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 26



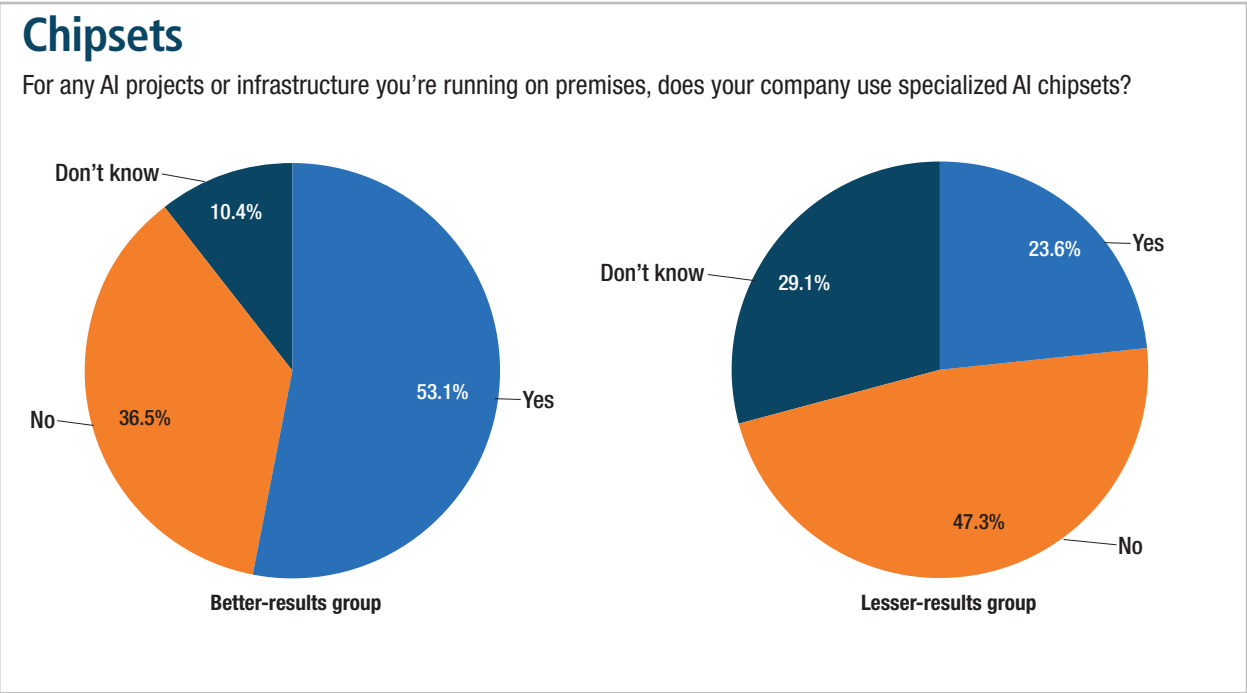
Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 27



Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 28




Respondent count: Better results, 96; lesser results, 55
Base: Respondents with at least one AI application in general production

FIGURE 29

Professional Services Firms Used by Better-Results Group

Which professional services providers does your company do business with? (Check all that apply.)



Respondent count: Better results, 46
Base: Respondents with at least one AI application in general production and citing use of a professional services firm