elastic | The Search AI Company

The State of Search & AI 2026

Navigating the potential of search-powered AI: Insights from 330 IT leaders

Table of contents

Introduction

The intersection of search and artificial intelligence (AI), which brings together the defining features of advanced search retrieval and AI, is experiencing considerable growth and user enthusiasm. Organizations want to move beyond traditional keyword-based search and take advantage of generative and agentic workflows. However, enterprises seem to be stuck between exploration and execution. While many plan to implement the technology in the next 6 to 12 months, the exact plan forward is not as clear.

In this report, more than 330 IT leaders with decision-making responsibility for search solutions share the state of search and AI and its growth potential. The report examines how IT leaders are setting the vision for search in their organizations. It probes how they can address security and privacy needs, overcome skills deficits, and avoid technical debt to implement search and find a path to return on investment (ROI). This report concludes with recommendations for a successful transition. This vendor-neutral, third-party research was independently conducted by ViB Research and sponsored by Elastic.

Key findings

1

Search adoption is broad-based, with organizations embracing advanced use cases

- Eighty-two percent of respondents use keyword and text searching, while semantic search and hybrid search are used by 39% and 52% of respondents, respectively.
- Searching internal knowledge bases is the most common use case, cited by 74%.
- Vector and hybrid search bestow the most sought-after search goals, including more relevant and accurate search and increased productivity.

2

Organizations are interested in search and AI with a few successfully moving from ambition to adoption

- Search-powered AI is only integrated into business operations in 6% of respondent organizations, while over 70% of respondents are either evaluating or investigating.
- Visions for AI-powered search are ambitious, with the practical outweighing the theoretical: 79% said they wanted search-powered AI to “increase employee productivity and efficiency,” but only 15% want search to drive innovation and research.
- Sixty-five percent said that their AI apps are using less than 25% of their organizational data, though data use is expected to grow over the next 6–12 months.

3

Internal use cases are the pragmatic path to search and AI success

- Organizations that start with internal knowledge bases see higher implementation success rates (73% vs. 61% average), faster time to value (62% achieve ROI in under 12 months vs. 51% average), and fewer change management challenges (34% vs. 42% average).

4

Change is on the way as organizations seek gains in productivity

- Search requirements are undergoing dramatic changes. Over the past six months, 31% of respondents have seen an increased need for natural language search.
- Respondents are excited about agentic AI, with 70% saying it will improve productivity.

5

Search and AI projects face many obstacles to adoption, though expectations for ROI are optimistic

- Data privacy, security, and compliance are the most common concerns in search vendor selection, cited by 70% of organizations as a top criterion, second only to cost at 73%.
- Communication and expectation setting with leadership can hold search-powered AI projects back: 50% of respondents cited challenges in managing shifting goals, while 48% struggled with bridging solution complexity with leadership's simplified perceptions.
- Despite these obstacles, expectations for an ROI from search-powered AI feel optimistic, with 57% envisioning ROI within 12 months.

The current state of search and AI

Respondents currently employ a wide range of search technologies, going beyond basic forms of search like keyword search, which is in use at 82% of respondent organizations. Semantic search and hybrid search are in use at 39% and 52%, respectively. Nearly a quarter use visual or image search. These are sophisticated uses of search that suggest that search use has expanded past simple lexical-based inquiries in the enterprise toward implementing contextual and agentic search.

Demographic summary

This survey's 330 respondents are responsible in some capacity for their organization's search technologies, with 71% being technology decision-makers and 62% involved in vendor selection. They come from a variety of industries, the three most highly represented being healthcare (15%), manufacturing (12%), and technology/software (12%). Half work at companies with more than 5,000 employees. For detailed respondent demographics, see the Appendix.

Definitions

- **Full-text search** — Finds relevant information by analyzing complete documents or datasets
- **Hybrid search** — Merges traditional keyword matching with semantic search and vector search
- **Keyword search** — Matches words in the search query with the text of the web or dataset to return relevant information
- **Semantic search** — Uses AI and natural language processing (NLP) to understand the meaning, intent, and context behind a user's query
- **Vector search** — Represents data as numerical vectors, allowing the search engine to find conceptually related content versus only keyword matches
- **Visual or image search** — Enables the use of images in search queries
- **Voice search** — Allows voice commands to instruct the search engine

Types of search in use

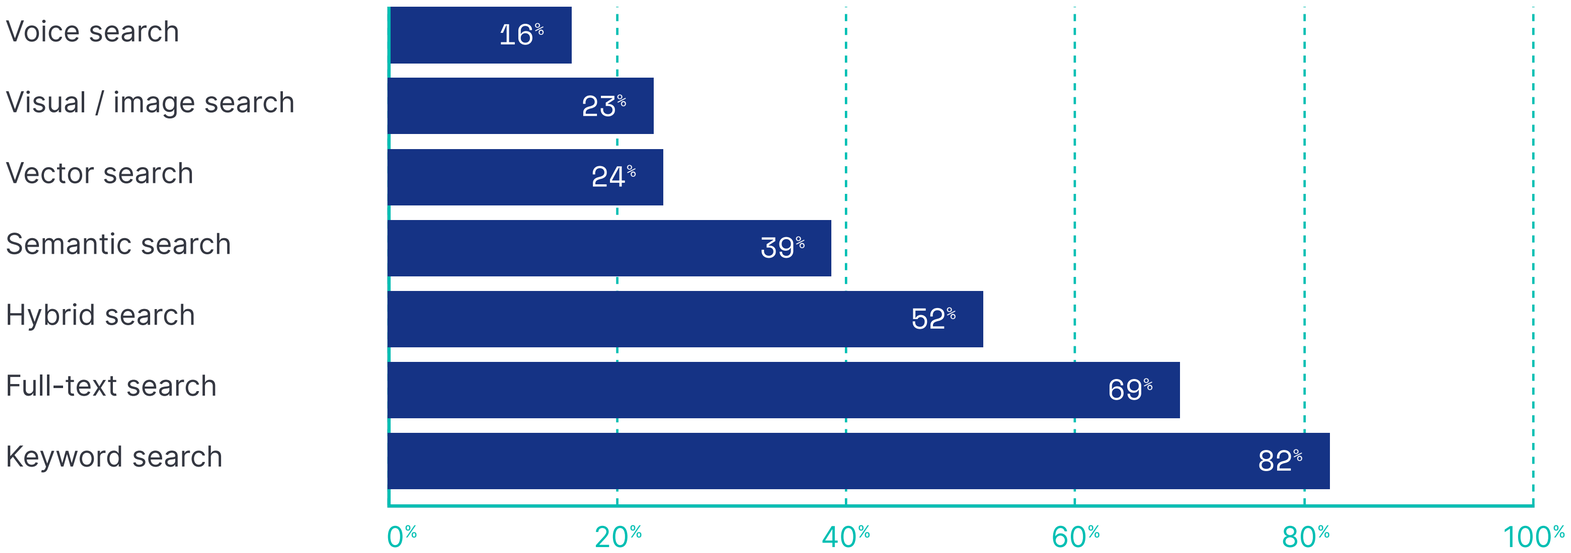

Figure 1. Answers to the question, "What types of search are currently in use at your organization?"

Most common search use cases

The most common search use case involves searching internal knowledge bases, cited by 74%, followed by website search (66%) and customer service (51%). These findings show that organizations prioritize internal efficiency.

Primary search use cases

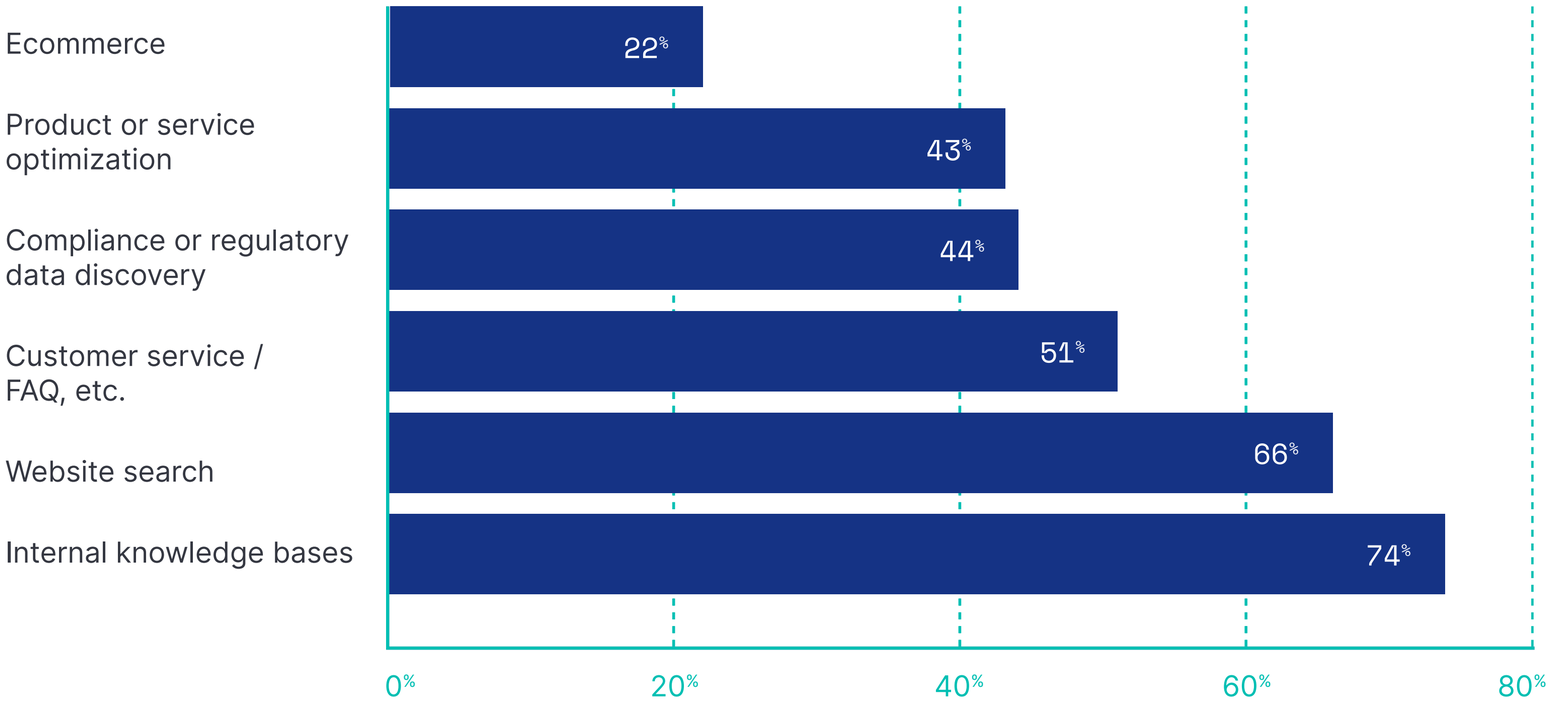

Figure 2. Answers to the question, “What are your primary search use cases within your organization?”

A revealing correlation emerges between use cases and successful search implementation. Organizations that start with internal knowledge bases experience:

- Higher implementation success rates (73% vs. 61% average)
- Faster time to value (62% achieve ROI in under 12 months vs. 51% average)
- Fewer change management challenges (34% vs. 42% average)

The pattern suggests a pragmatic path to success. By starting internally, search teams can demonstrate value to employees, build organizational confidence, and then expand to customer-facing applications. Organizations that began with ambitious customer-facing deployments show higher rates of stalled implementations and extended timelines.

Vector and hybrid search improve productivity and search accuracy

Vector search and hybrid search make search results more accurate and relevant, according to 63% of respondents. A similar number said that vector and hybrid search increased productivity, while 50% cited faster information retrieval across diverse data sources as a benefit of these two search modes. These benefits are the realization of respondents' most sought-after search goals.

Benefits of vector and hybrid search

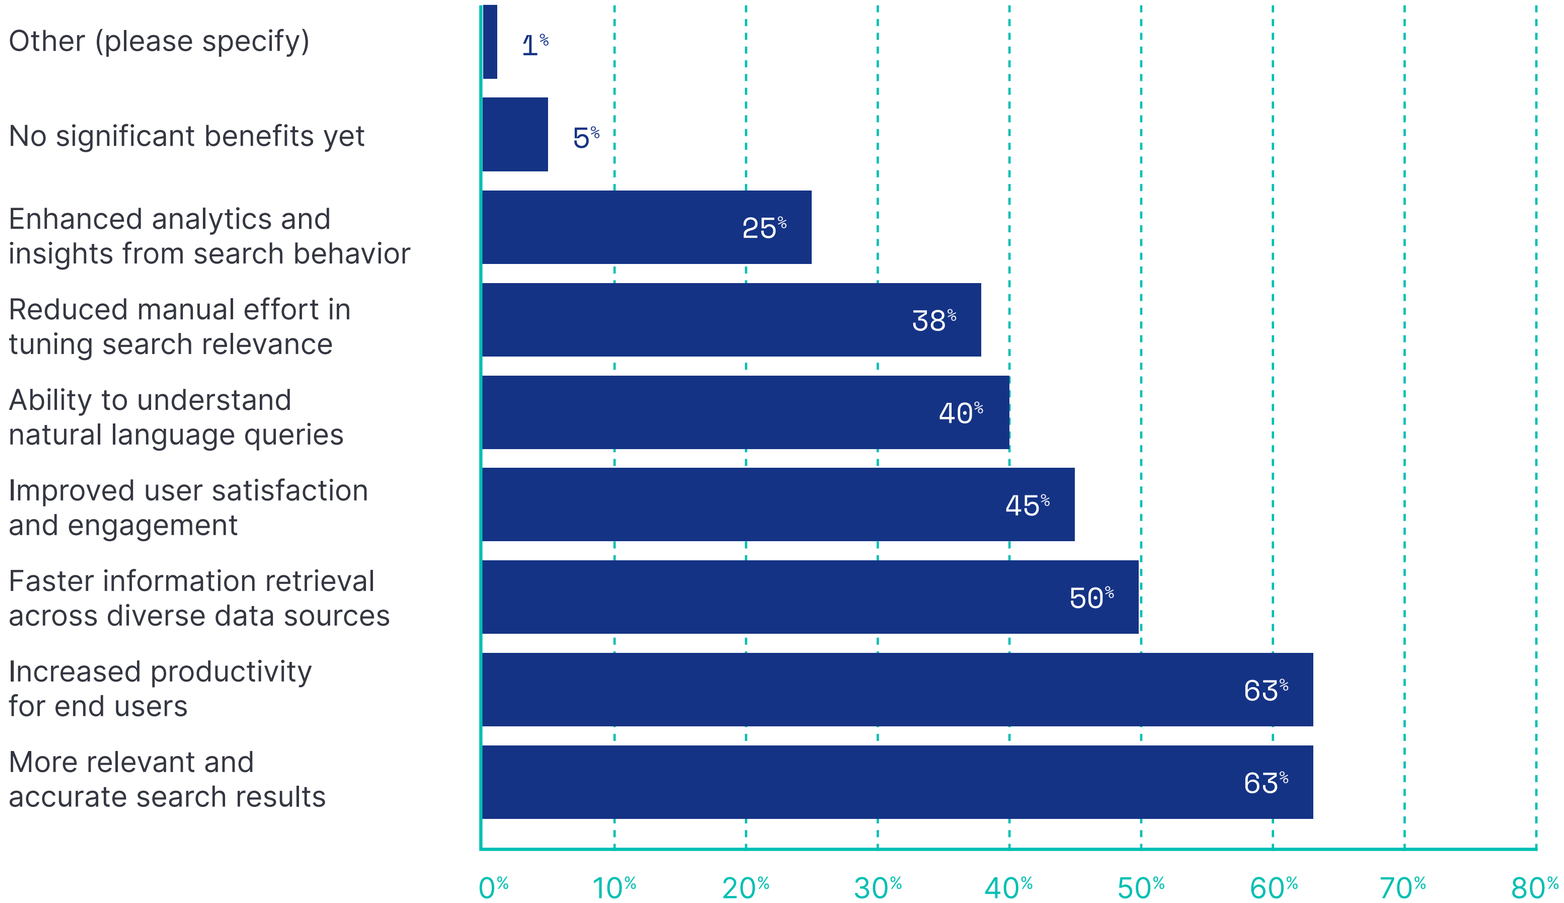

Figure 3. Answers to the question, “What benefits have you experienced or expect to experience from implementing vector and hybrid search at your organization?”

The uses for vector and hybrid search are varied. Use cases receiving priority run the gamut from AI-powered customer support and marketing content development to scientific research and skills-based talent matching. As a search administrator in the business services industry explained, “We are prioritizing leveraging vector and hybrid search primarily for applications involving knowledge management, incident analysis, and support automation.” A respondent who works in healthcare discussed their organization’s plans for retrieval augmented generation (RAG), saying they were building “complex recommendation engines that need to match user preferences with item attributes.”

The search-powered AI adoption paradox

Search-powered AI is on the horizon, but current deployments are still relatively rare. Today, search-powered AI is only integrated into business operations in 6% of respondent organizations. Twenty-two percent are currently implementing the technology. However, more than 7 out of 10 respondents are either evaluating or investigating search-powered AI solutions. These data points suggest a paradox at the heart of search-powered AI adoption: Companies understand the value of it, but they’re stuck between evaluation and action.

Current adoption of search-powered AI

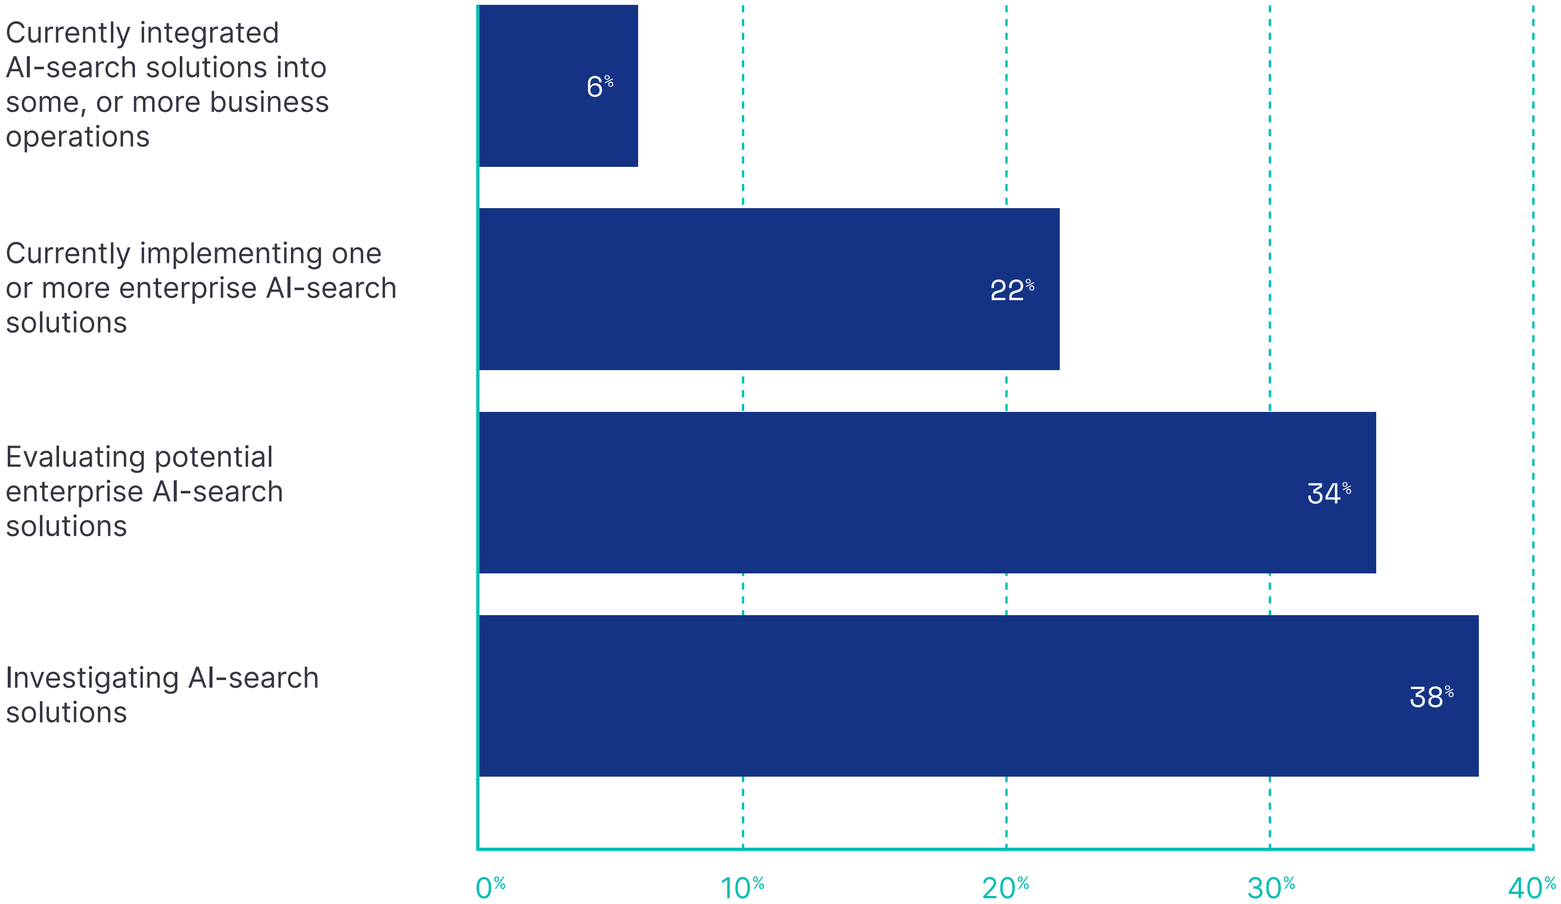

Figure 4. Answers to the question, “What is your organization’s current adoption status of AI search solutions?”

Organization size correlates with implementation confidence. Larger enterprises (10,000+ employees) show higher rates of active implementation (26%) compared to mid-sized organizations (500–2,499 employees) at 18%. This suggests that scale provides advantages, whether through larger budgets, deeper technical teams, or simply the capacity to fail forward and iterate.

One implication of the adoption paradox is that organizations in the investigation and evaluation phases risk widening the capability gap compared to their more aggressive competitors. As one senior technology leader noted, the challenge isn't identifying whether to adopt search-powered AI, but determining how quickly they can move from concept to production.

Search requirements are evolving as business needs grow more complex

Search requirements are undergoing dramatic changes. Over the past six months, 31% of respondents have seen an increased need for natural language search. A quarter said that their search requirements changed in the past six months to include a greater demand for real-time or near real-time results, while 17% that expanded data sources and types had become part of the requirements in this timeframe.

Looking ahead, 34% will require more advanced AI/machine learning (ML) capabilities in the next six months. Twenty-four percent said they will need to integrate with more apps, and 21% will need to support more datasets. Taking the past and future changes together, one sees a push toward more data sources, faster results, and more contextual relevance and semantic capabilities. Combined with an expectation of integrating with more business systems, these findings highlight how search-powered AI is expected to become more integral to business processes and day-to-day business workflows.

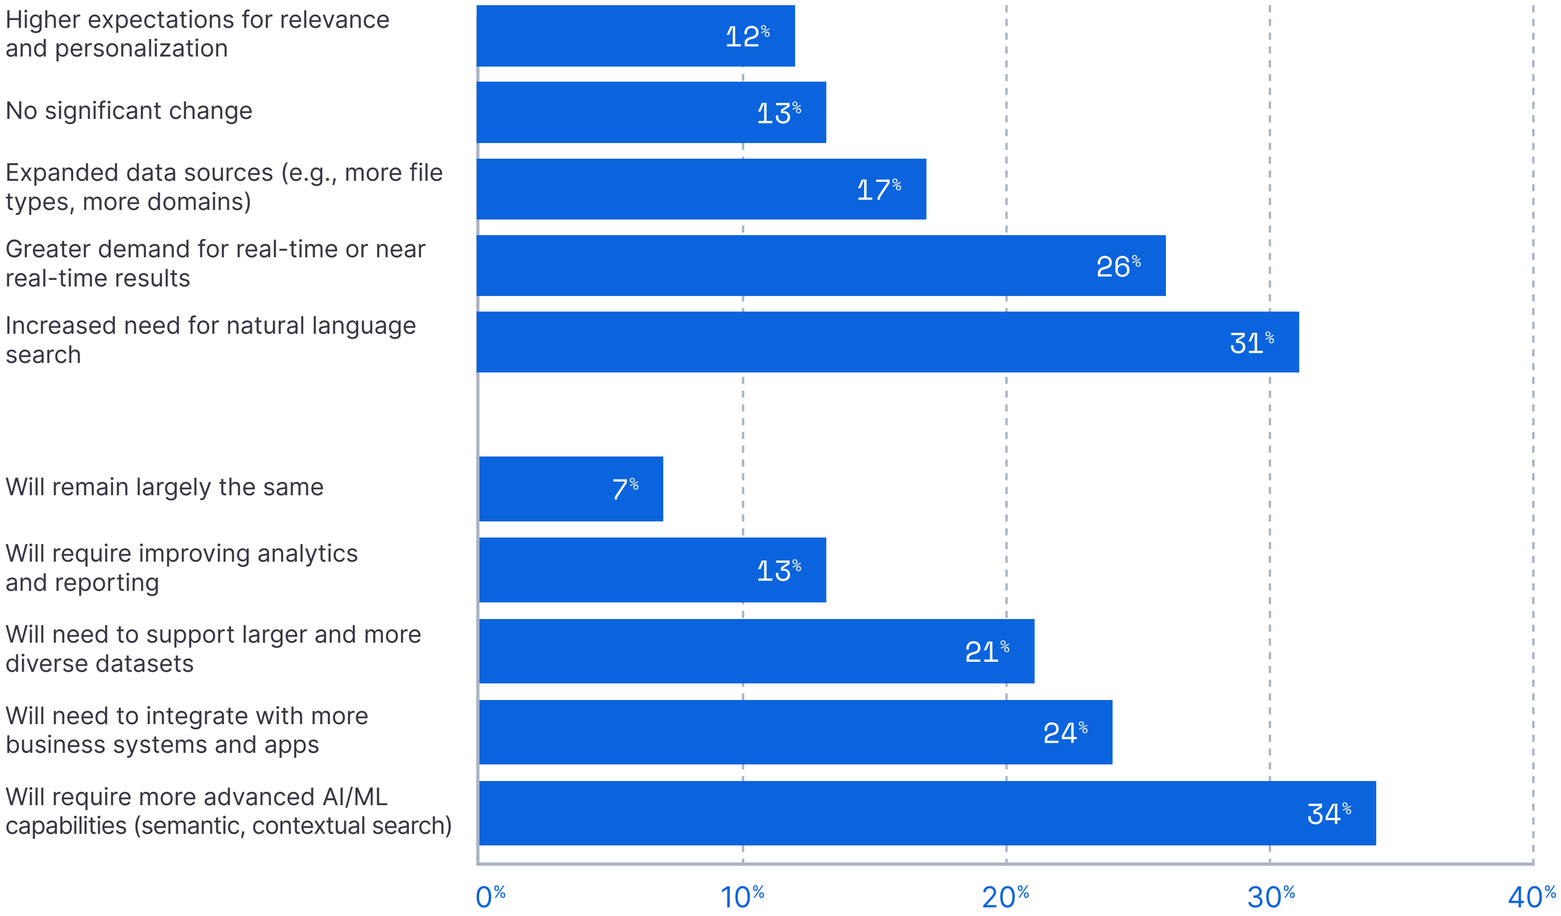

Changes in search requirements in past 6 months

Figure 5. Answers to the question, “How have your search requirements changed over the past six months?”

Predicted changes in requirements in the next 6 months

Figure 6. Answers to the question, “How do you predict your search requirements will change over the next six months?”

The need to support larger and more diverse datasets in the next six months makes sense in the context of how little data AI applications are used today.

- Sixty-five percent of respondents said their **AI apps use less than 25%** of their organization's data.
- Twenty-seven percent of respondents said they **use less than 10%.**

What's causing this low rate of data access? Possible reasons could include a lack of time and resources for data integration, as well as technical complexity and concerns about security. Whatever the issues, however, there's an expectation that AI will access more data in the future.

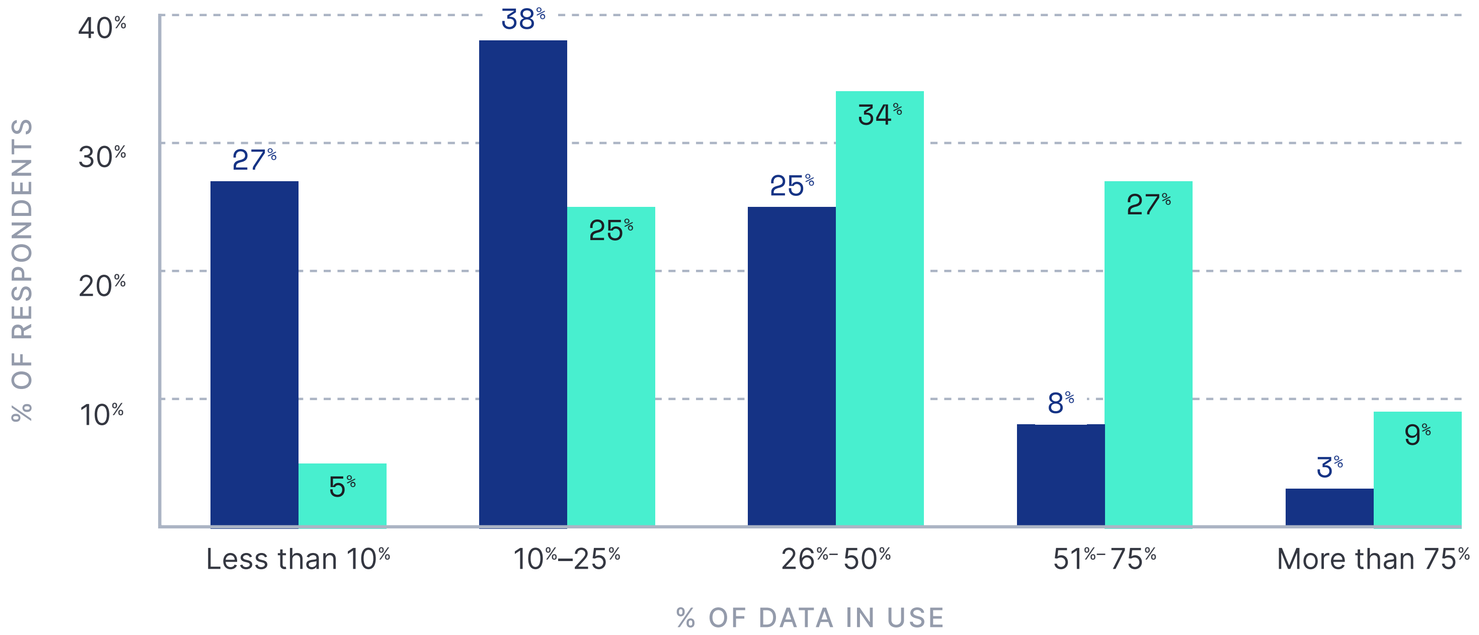

Percentage of organizations' data in use by AI apps

Figure 7. Answers to the question, "What percentage of your organization's data are you using in AI applications today?"

Anticipated future use of data by AI apps

Figure 8. Answers to the question, "What percentage of your organization's data do you anticipate using in AI applications in the next 6–12 months?"

Data use is expected to grow over the next 6–12 months. Seven in ten respondents plan to be using over 25% of their organization's data within 12 months. In contrast, 9% said they expect to use over 75% of their organization's data in AI apps.

Organizations have ambitious visions for search and AI

Search professionals have many ambitious goals, with the practical outweighing the theoretical. When asked, “What is your organization’s main goal in implementing search in your organization?,” 79% of respondents said “increase employee productivity and efficiency,” while 41% said “reduce operational costs.” Only 15% want search to drive innovation and research. Better knowledge management and retention was chosen by 26%.

Goals for search

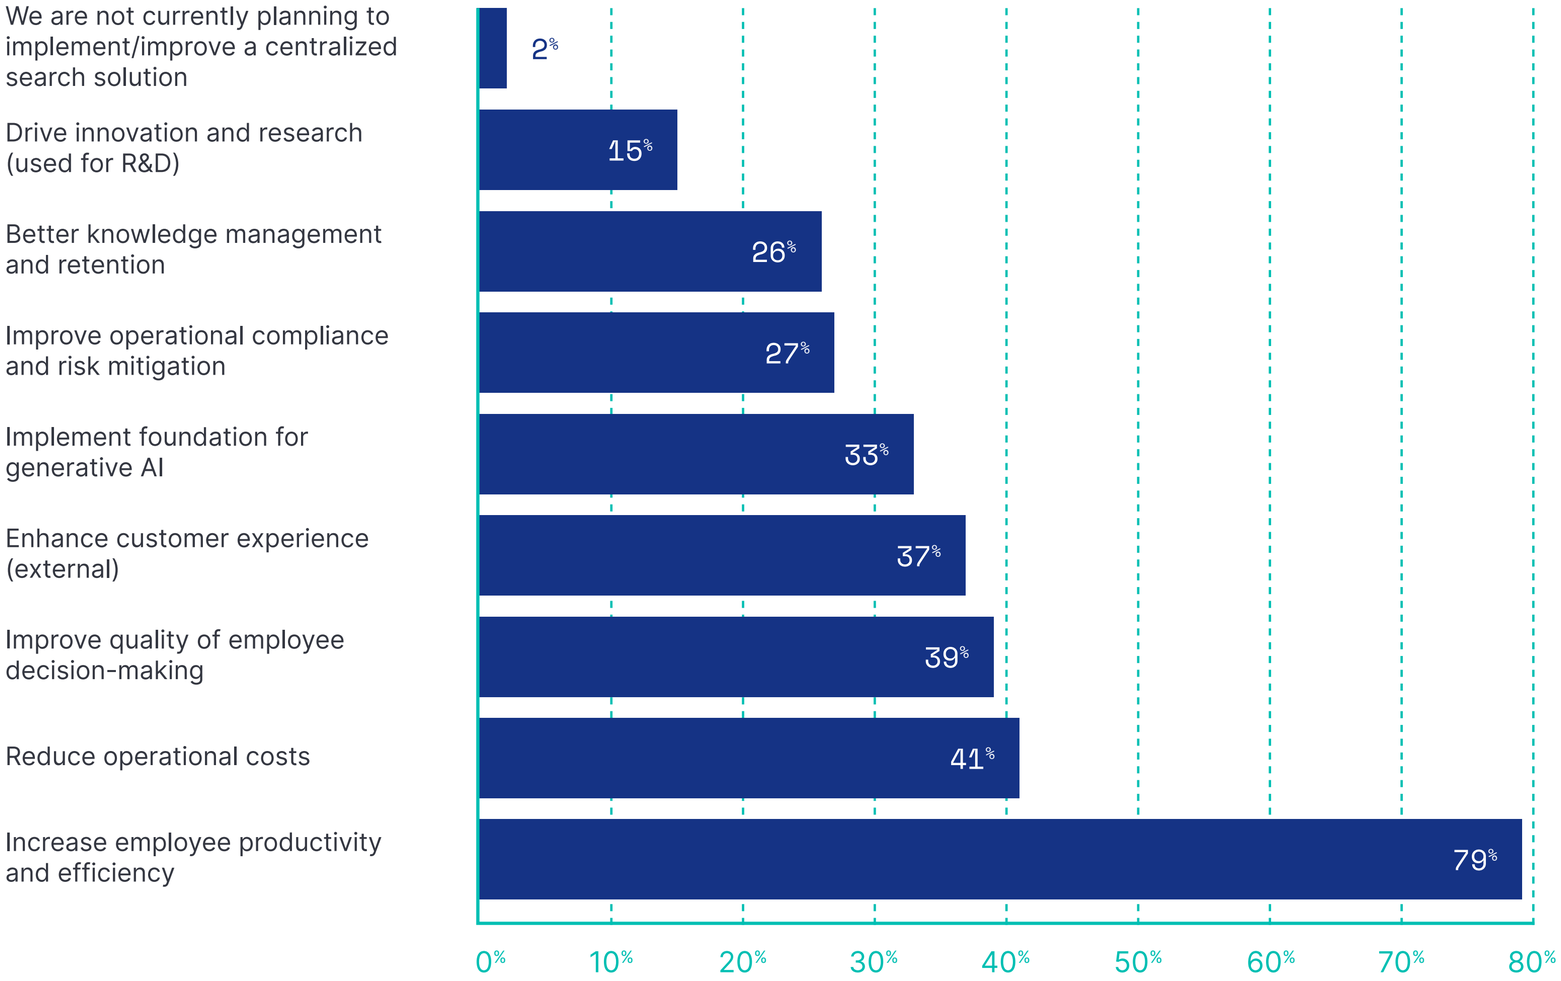

Figure 9. Answers to the question, “What is your organization’s main goal in implementing search in your organization?”

Near-term plans for search and AI are also ambitious: More than 80% either have or plan to adopt AI-powered search analytics, semantic search, personalized search results, question answering systems, and auto-complete.

Adoption plans for search and AI capabilities

	Adopted	Plan within 6 months	No plans	Total adopted and plan to adopt within 6 months
Retrieval augmented generation (RAG)	12%	42%	26%	54%
Content recommendation engines	18%	49%	18%	67%
Question answering systems	33%	41%	13%	74%
Semantic search/natural language understanding	26%	54%	11%	80%
Auto-complete and query suggestions	32%	51%	17%	83%
Personalized search results	33%	53%	15%	86%
AI-powered analytics (e.g., search trends, no-result queries)	25%	63%	12%	88%

Figure 10. Answers to the question "What AI search capabilities have you adopted or plan to adopt in the next 6 months?"

As shown in Figure 10, organizations are prioritizing context capabilities within their search experiences within the next six months. They plan to increase their adoption of conversational, context-aware technologies like AI analytics where planned adoption increases 38% over the next six months, RAG (up 30%), and semantic search/natural language understanding (up 28%). In Figure 11, the Elastic search maturity model shows where these changes fit in the broader evolution of search in the AI era.

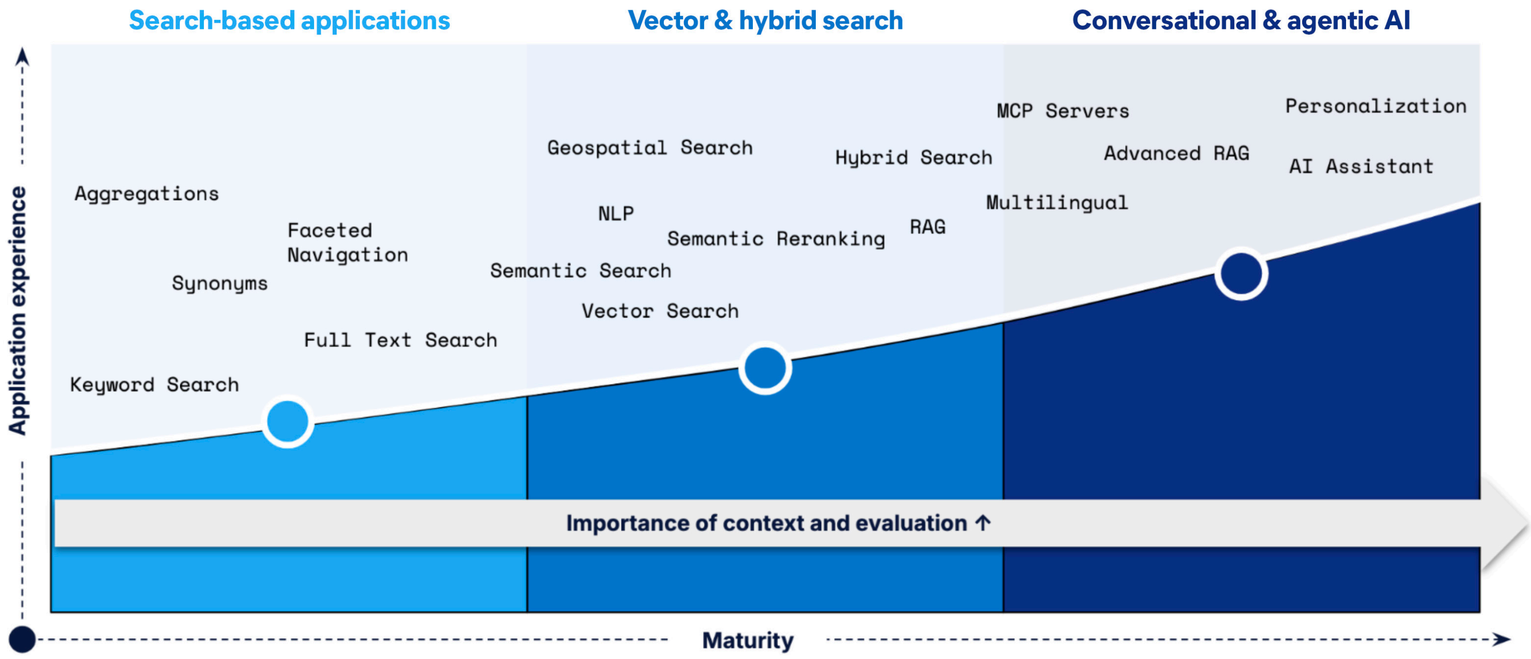

Figure 11. Elastic search maturity model

Defining ROI goals for AI investments

Expectations for ROI from search-powered AI feel optimistic, with 57% envisioning ROI within 12 months. This level of confidence may be due to the practical goals that IT leaders have for search-powered AI like improving productivity. Search has long demonstrated that it can deliver increased productivity, so it's not a big jump to imagine that new AI-powered search capabilities will lead to further productivity gains.

ROI expectations for search and AI

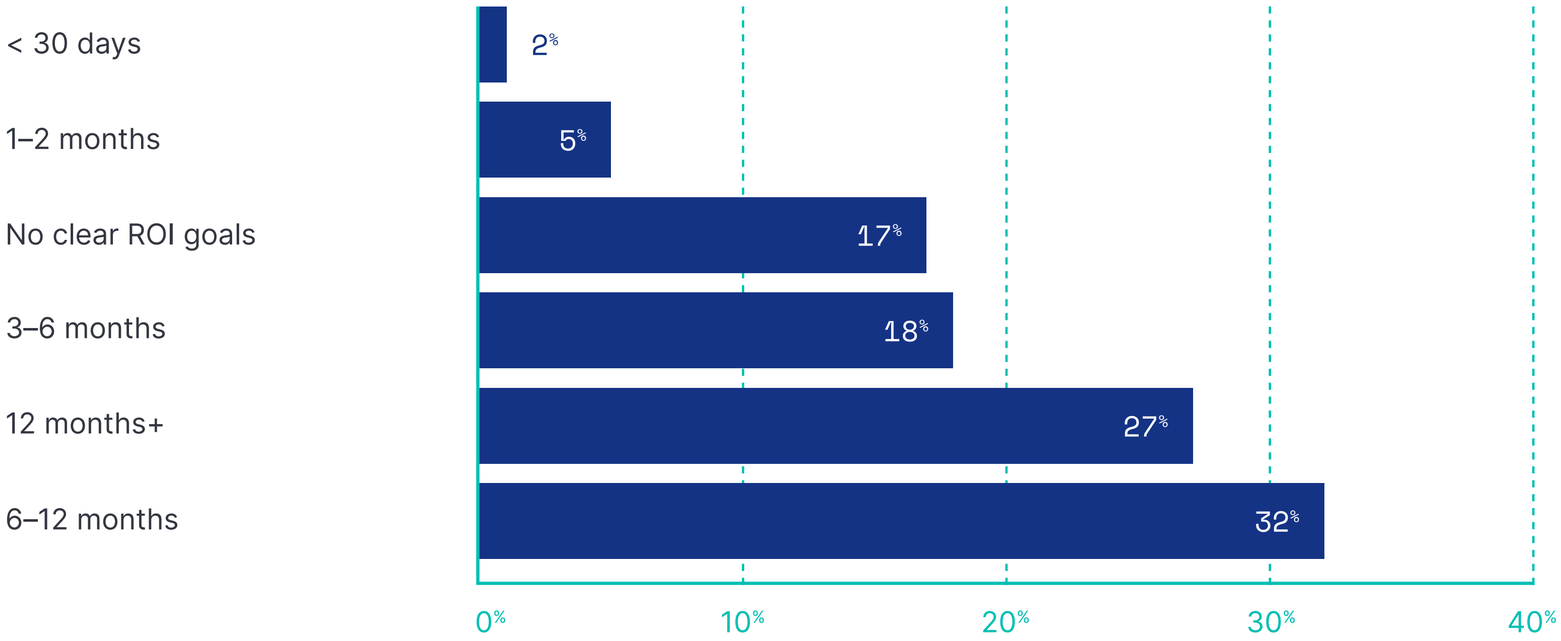

Figure 12. Answers to the question, "What are your organization's expectations for seeing returns on AI investments?"

About a sixth of respondents (17%) reported that they had no clear ROI goals for AI investments. The characteristics of this cohort are revealing:

- Forty-one percent had higher rates of stalled implementations versus 28% with clear ROI goals.
- Sixty-three percent had more frequent leadership challenges versus 48% with clear ROI goals.
- Forty-seven percent had lower budget allocation for projects versus 38% with clear ROI goals.

These findings suggest that a clear ROI framework is more than just an exercise in financial modeling. Rather, the ROI assessment process is essential for aligning stakeholders, justifying the expenditure of resources, and providing success metrics.

Underscoring this point, organizations with the most successful implementations define ROI across multiple dimensions:

- Productivity metrics (e.g., time saved, searches resolved)
- Cost reduction (e.g., support ticket deflection, manual search elimination)
- Quality improvements (e.g., decision accuracy, information discovery)
- Strategic enablement (e.g., foundation for generative AI, competitive positioning)

Respondents are excited about agentic AI

Respondents also shared big visions and ideas for agentic AI. A respondent who works in business services stated, “We see agentic AI becoming a core driver of productivity and innovation — both within our company and across the industry.” Specifically, their team “expects AI agents to take on more complex, multistep workflows, help teams make faster data-driven decisions, and automate routine operational tasks so people can focus on high-value work.”

Another perspective, offered from a respondent in the tech field, was that “agentic AI is rapidly transforming from assistive tools into autonomous digital collaborators capable of planning, reasoning, and executing complex tasks with minimal human input.” They cited the example of AI agents that can “orchestrate workflows across applications, manage infrastructure, and proactively secure environments.”

Present experience predicts future benefits. When asked about current benefits from agentic AI workflows, 70% said increased user productivity. Fifty-six percent said that agentic AI improved relevance and accuracy of results, while 55% said the technology reduced manual search and support workloads.

Benefits experienced or expected for agentic AI

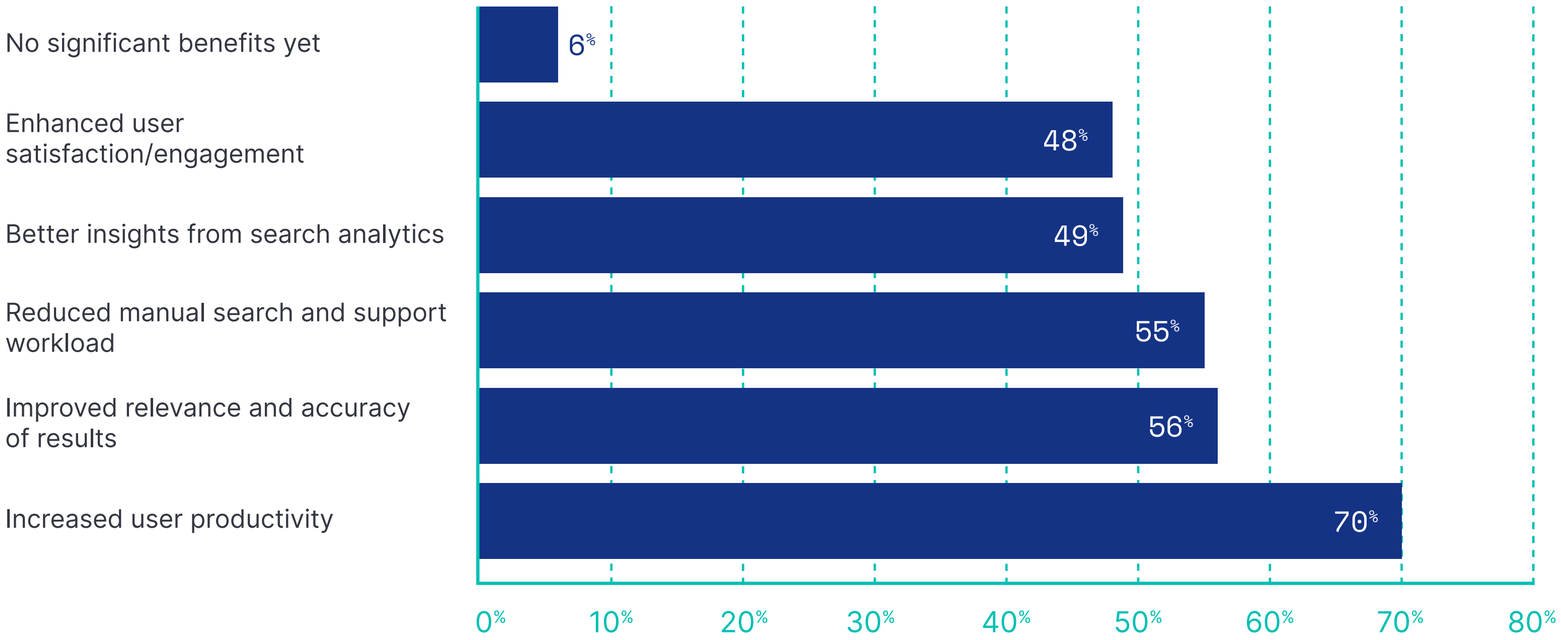

Figure 13. Answers to the question, “What benefits have you experienced or expect to experience from implementing agentic AI workflows at your organization?”

Obstacles stand in the way of search-powered AI implementation

Only a small percentage of search-powered AI apps are progressing from concept to production, with 68% of respondents reporting that fewer than 25% of their apps reach production. Nearly 4 out of 10 respondents say that fewer than 10% of their apps are progressing.

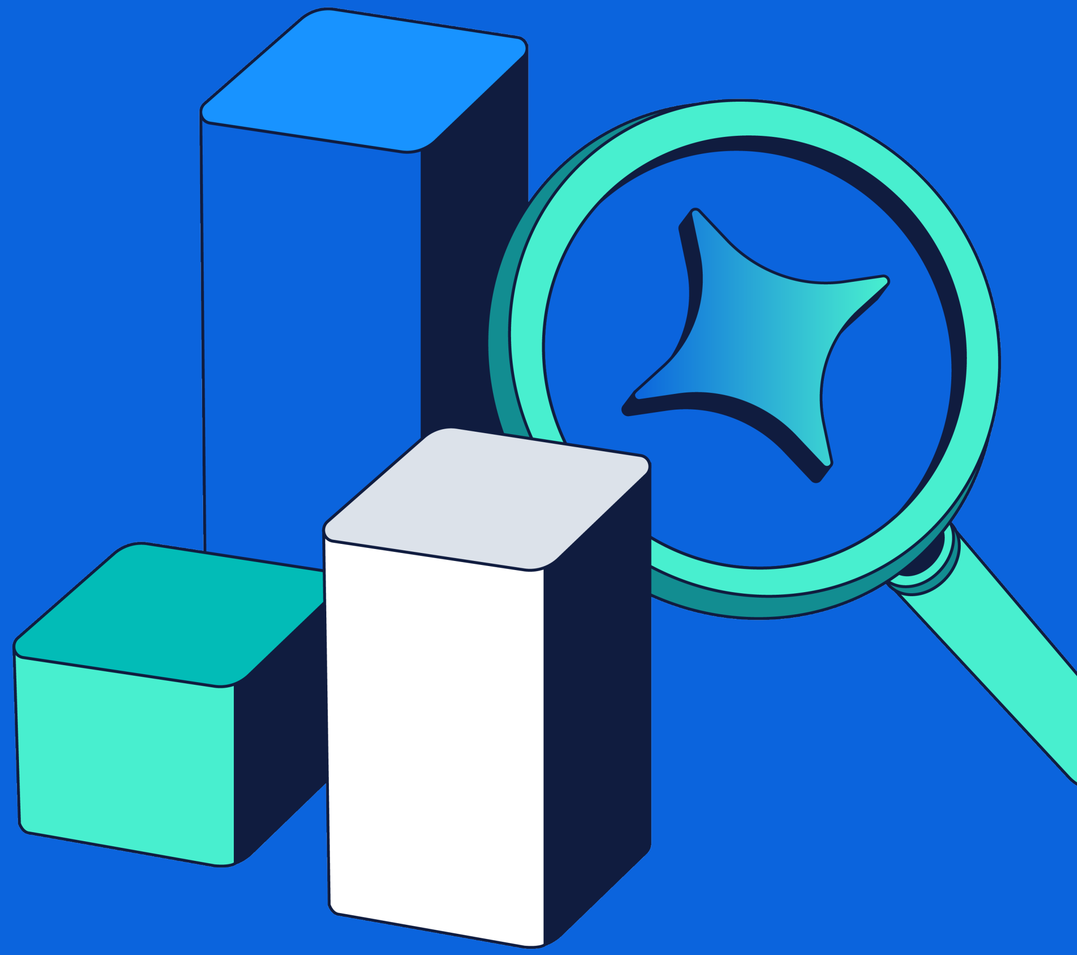

Search-powered AI apps progressing from concept to production

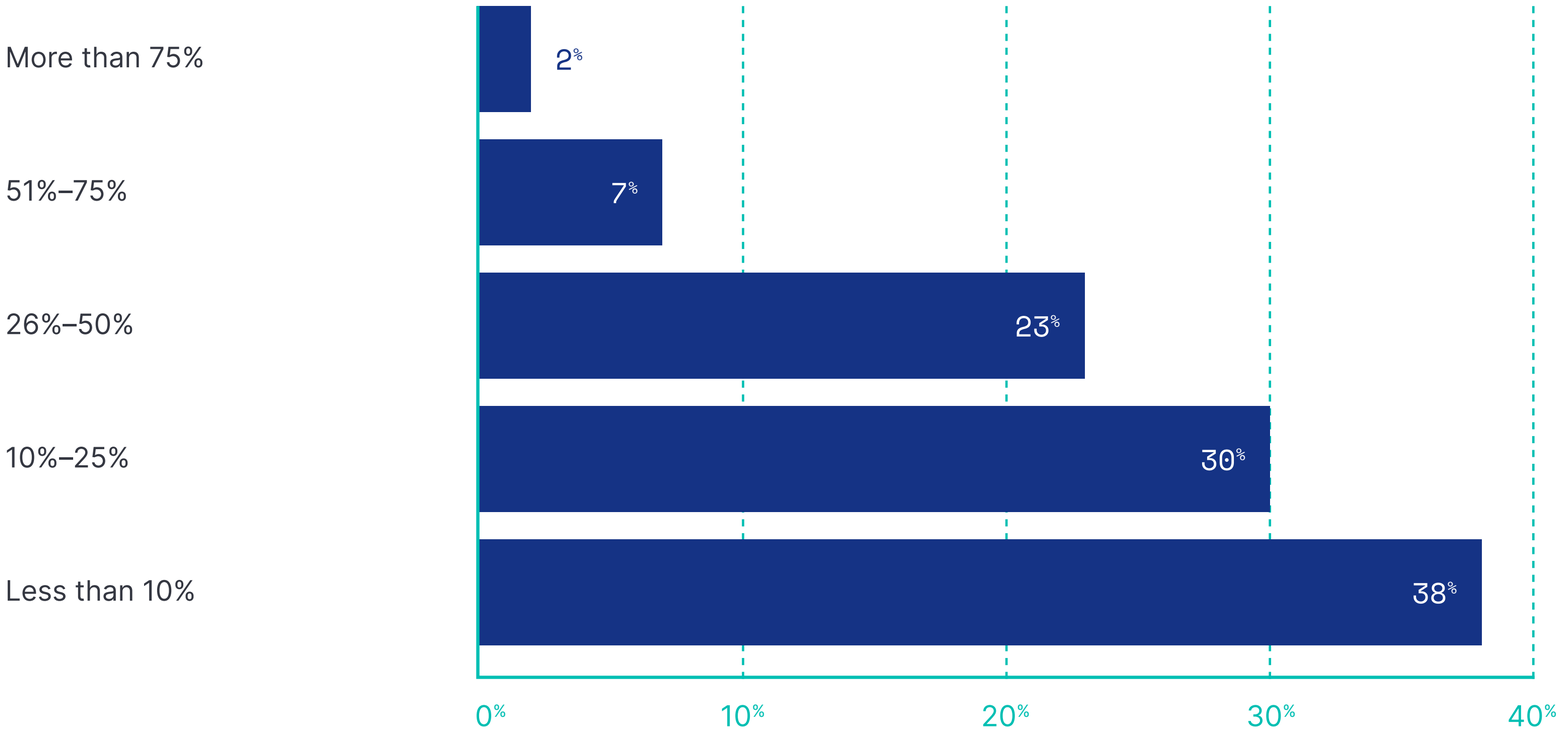

Figure 14. Answers to the question, “How many AI search apps are progressing from concept to production?”

What’s holding these apps back?

Budget and complexity are the two biggest challenges preventing search apps from progressing from concept to production, cited by 53% and 51%, respectively. Data quality and availability are also factors for 50% of respondents.

Key challenges in taking search solutions to production

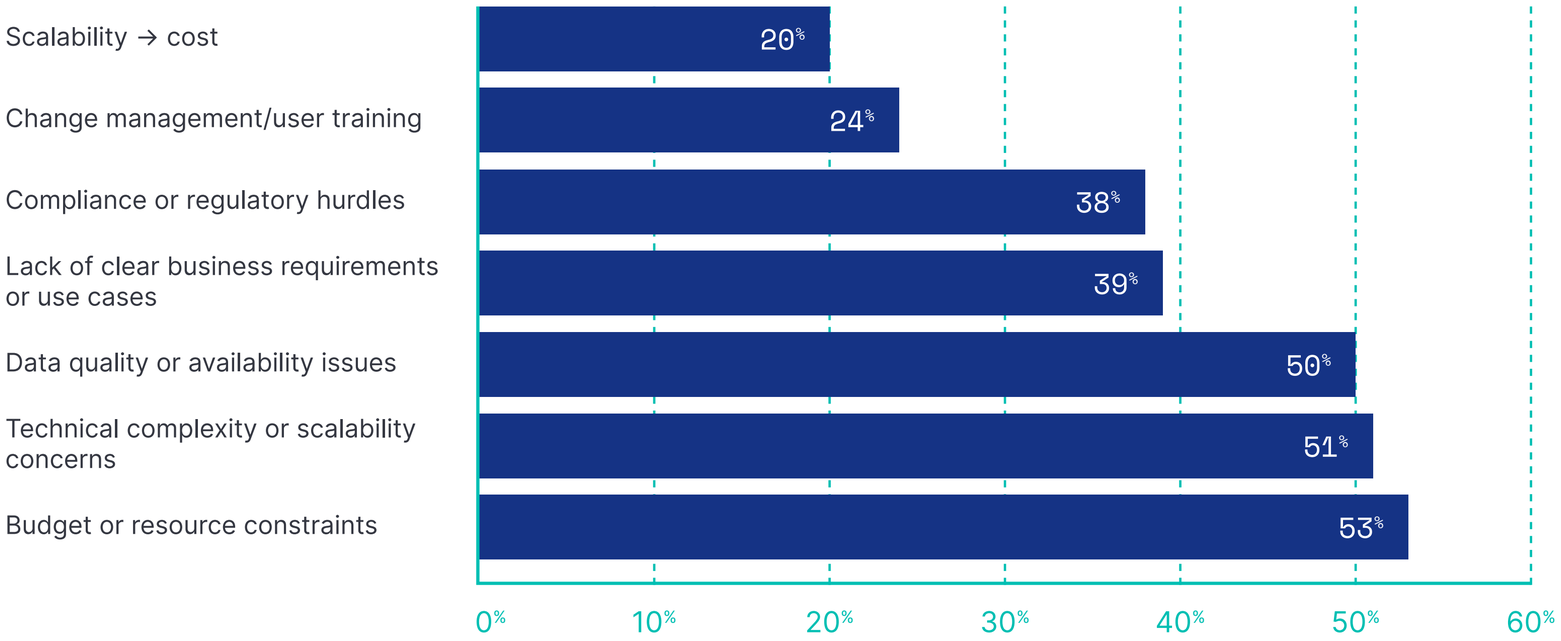

Figure 15. Answers to the question, “What are the key challenges in progressing search solutions from concept to production?”

Data privacy and security

Data privacy and security are concerns that potentially limit the use of data by AI apps. Indeed, the organizations planning the most aggressive data expansion also report the most significant concerns about data privacy. These concerns reflect the realities of operating amid increasingly tight regulations and powerful cyber threats. AI access to data represents an additional attack surface, so search managers are right to be cautious, even as they face pressure to increase data use. This caution shows up in search vendor selection, where 70% of organizations place security, privacy, and compliance capabilities as a top criterion, second only to cost at 73%.

The most common challenges facing AI projects

Search projects suffer from a range of challenges, most of which reflect a gap between vision and operational realities. For example, 51% of respondents experienced challenges in designs that are innovative yet impractical to maintain. The need to handle complex technical issues on tight timelines is a problem for 49%. Forty-eight percent feel that it’s challenging to balance technical quality with immediate business needs.

Common challenges in developing search projects

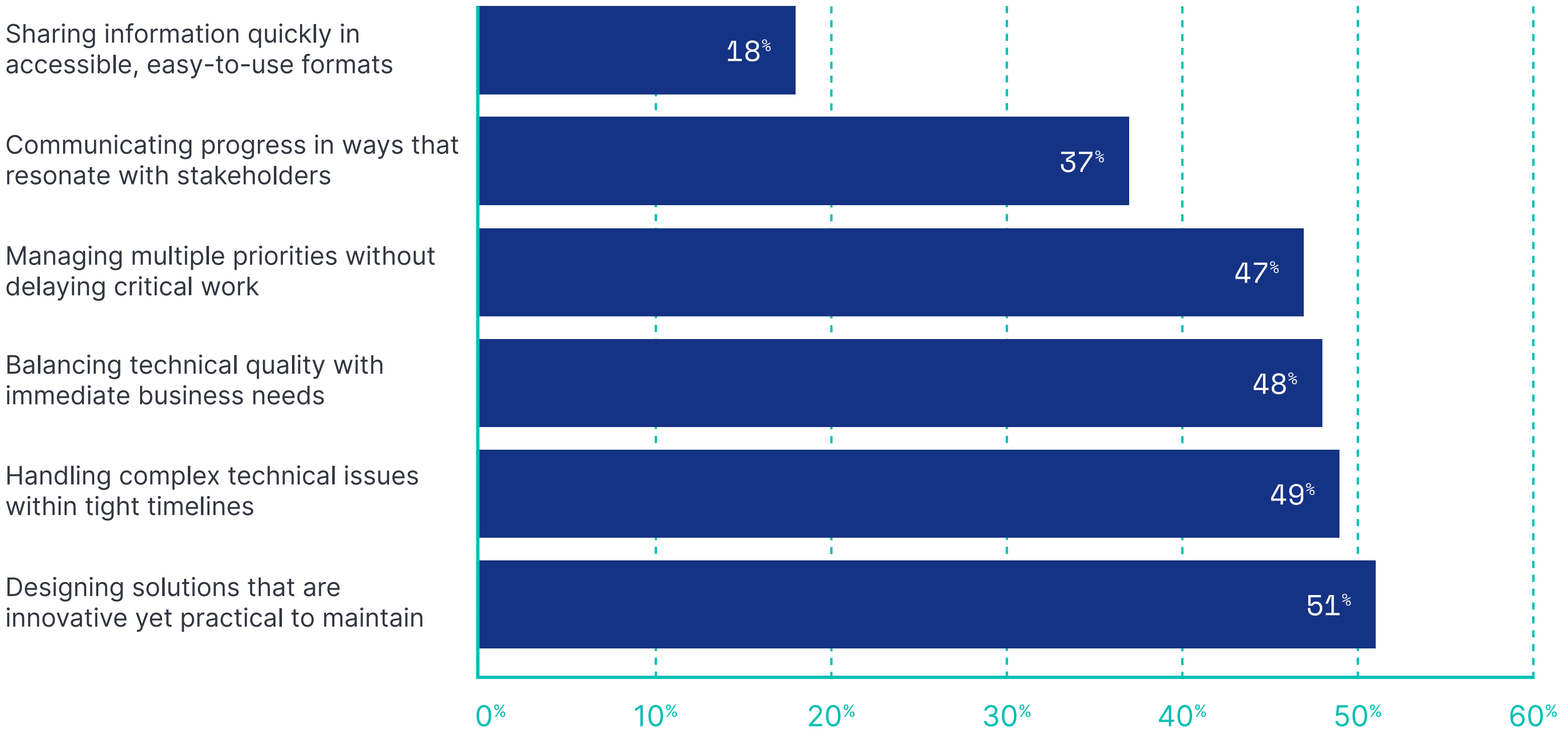

Figure 16. Answers to the question, “Which of the following are common challenges your team faces when developing search projects?”

Resource realities

Nearly half of the respondents cited a lack of internal expertise or resources as a primary challenge.
This manifests in three ways:

- **Expertise scarcity:** Only a small, sought-after talent pool has the specialized knowledge required for search-powered AI implementation, including understanding vector embeddings, RAG architectures, semantic search tuning, and NLP model selection. Organizations compete for this expertise, which drives up costs and lengthens project timelines.
- **Budget pressure:** 53% said that budget or resource constraints are barriers to moving from concept to production. This creates a catch-22 situation, where organizations need to demonstrate ROI to secure additional funding but need funding to reach the point of demonstrable ROI.
- **Vendor dependency:** The expertise gap drives 73% of organizations to prioritize cost and pricing models in vendor selection. This is not due to a desire to be economical, but rather because they're budgeting for extensive vendor support and services to compensate for internal capability gaps.

The correlation between expertise challenges and adoption stages is revealing:

- Organizations **investigating** solutions cited expertise gaps 52% of the time.
- Organizations **evaluating** solutions cited them 47% of the time.
- Organizations **implementing** solutions cited them 41% of the time.
- Organizations with **fully integrated** solutions cited them only 32% of the time.

This suggests that capability building follows commitment. Organizations develop expertise through implementation, not before it. Waiting for the expertise to be available in-house before starting means never starting. Successful organizations build capabilities in parallel with implementation, often heavily leveraging vendor partnerships and external expertise in early phases.

Challenges affecting vector and hybrid search

Three top challenges cited for implementing vector and hybrid search are: data privacy or security (57%), complexity in integration (48%), and lack of internal expertise (38%). Mitigating security and privacy risks involves a combination of purpose-built tools, skills, and policies. Finding expertise means recruiting highly paid employees or retaining consultants.

Challenges in implementing vector and hybrid search

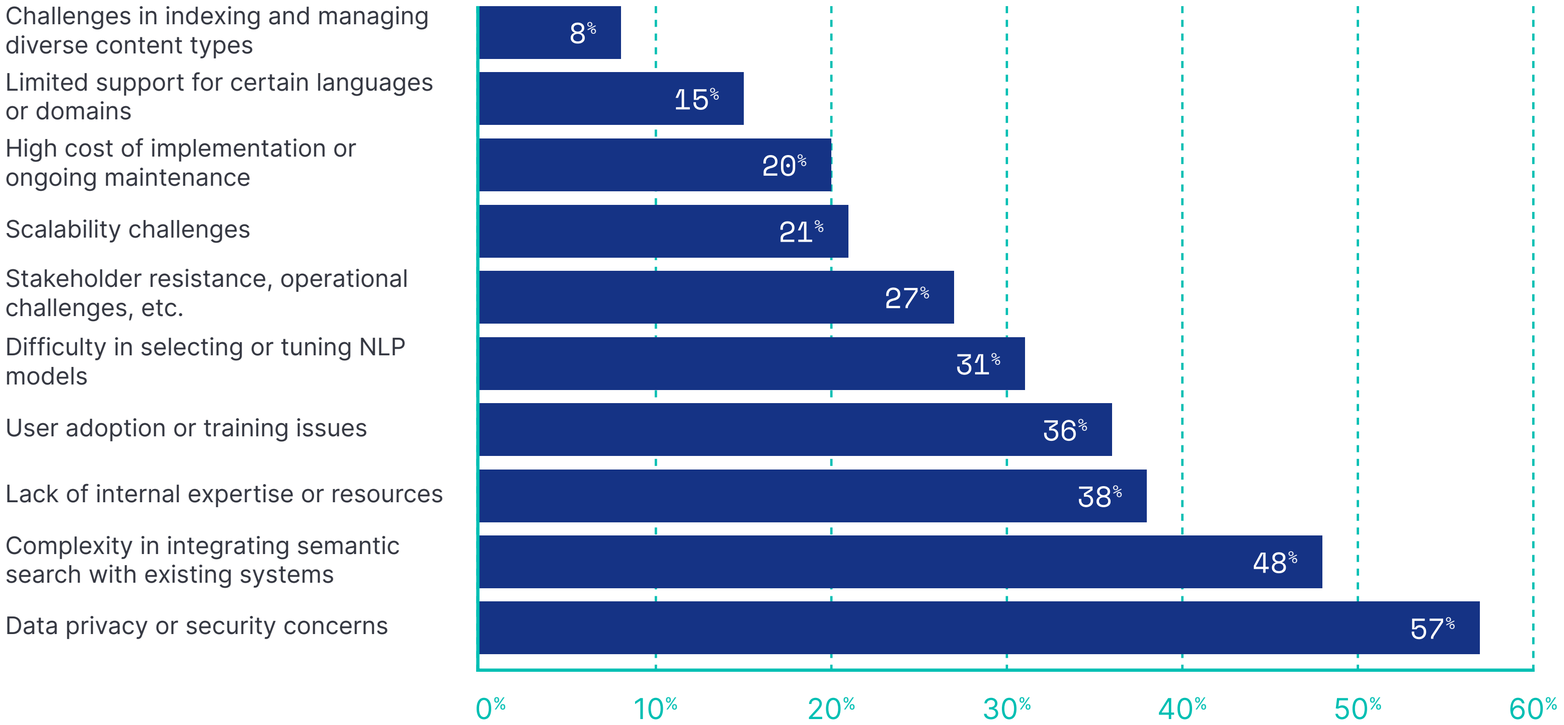

Figure 17. Answers to the question, "What challenges have you experienced or expect to experience from implementing vector and hybrid search at your organization?"

Challenges affecting agentic AI implementation

Agentic AI implementation faces similar challenges to those cited for vector and hybrid search. Nearly two-thirds of respondents saw data security and privacy as implementation challenges. Another 54% are worried about integration complexity. A lack of internal resources for this new technology was also cited (48%).

Concerns about the immaturity of the technology and its governance emerged. A respondent from the manufacturing industry remarked, “It’s got to get better, but will, fast.” Regarding security and governance, a search admin, also in manufacturing, said that their use of agentic AI will happen “after guardrails are in place.”

Challenges in implementing agentic AI

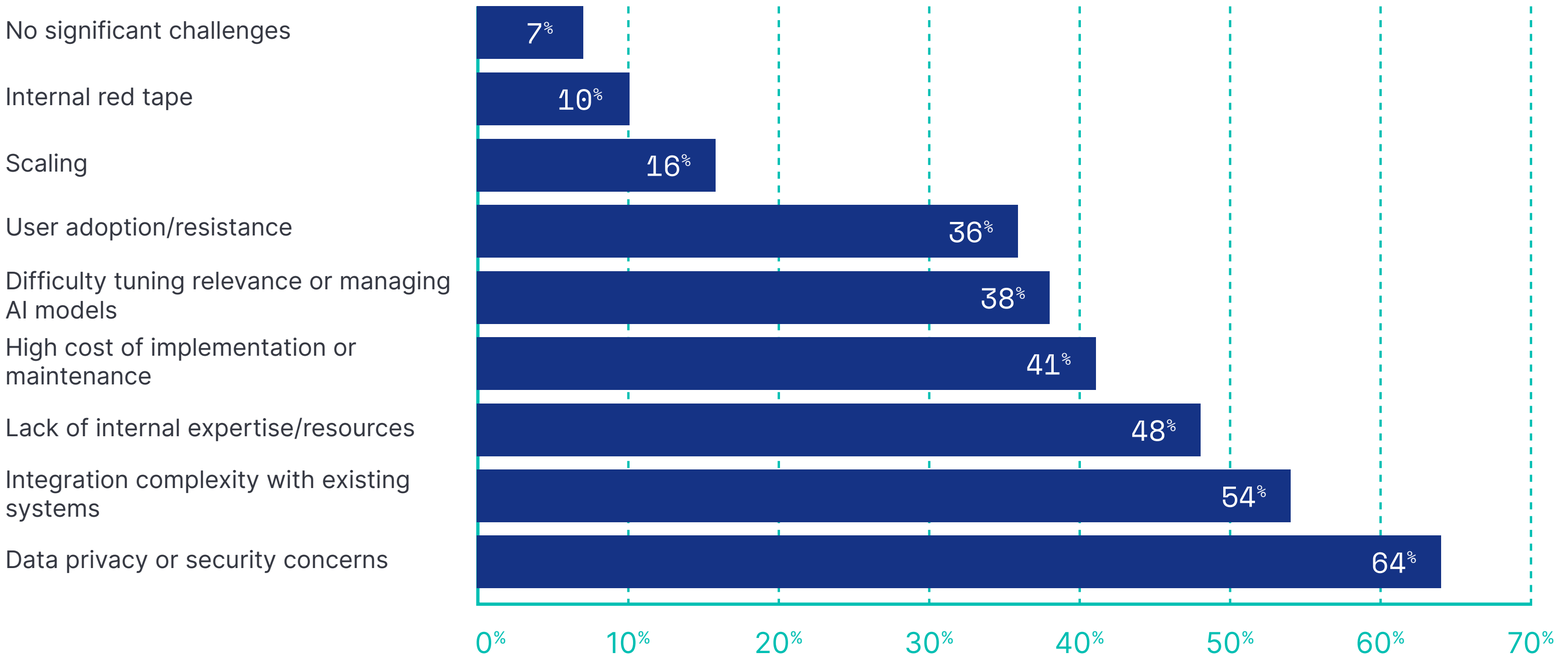

Figure 18. Answers to the question, “What challenges have you experienced or expect to experience from implementing agentic AI at your organization?”

Aligning leadership's expectations with search-powered AI implementation realities

Some of the resource and ROI challenges become easier to understand when taking leadership interactions into account. The problems mostly relate to managing expectations and perceptions. Fifty-six percent said aligning expectations with what is technically feasible is a common challenge when working with leadership. Half of respondents said the challenge was managing shifting goals, while 48% brought up the challenge of bridging solution complexity with leadership's simplified perceptions.

Common challenges in working with leadership

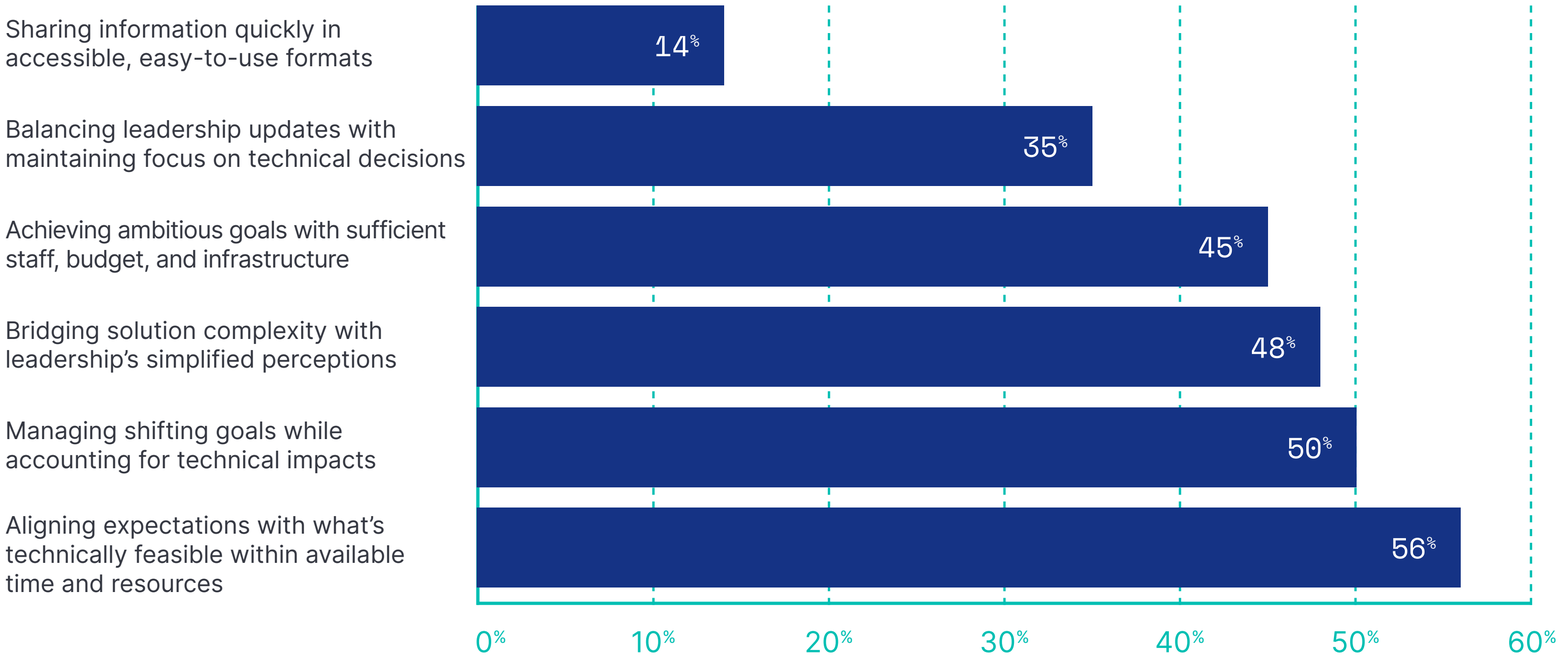

Figure 19. Answers to the question, "Which of the following are common challenges your team faces when working with leadership on search projects?"

Organizations that are struggling with leadership alignment can learn from IT leaders who have successfully implemented search and AI. They reported:

- More structured communication between technical teams and leadership (65% among respondents with successful search-powered AI implementations vs. 48% among those without)
- Clear ROI frameworks established *before* implementation (71% vs. 54%)
- Phased rollout strategies with defined success metrics at each stage
- Regular executive education on AI capabilities and limitations

The takeaway here is that successful search-powered AI implementation requires investments in internal, executive-level communication and expectation management, not solely investment in technical execution.

Top 3 criteria for selecting a search and AI vendor

What makes for a good search-powered AI vendor? According to the 330 IT leaders surveyed, three criteria are critical in the vendor selection review:

- Cost and pricing (73%)
- Security, privacy, and compliance (70%)
- Vendor market reputation (50%)

Technical criteria like speed of implementation (48%) and ease of integration (46%) rounded out the top five vendor selection priorities. The focus on security and privacy underscores the significant concerns about exposing organizations to risk when implementing search and AI.

Criteria for evaluating search-powered AI tools and vendors

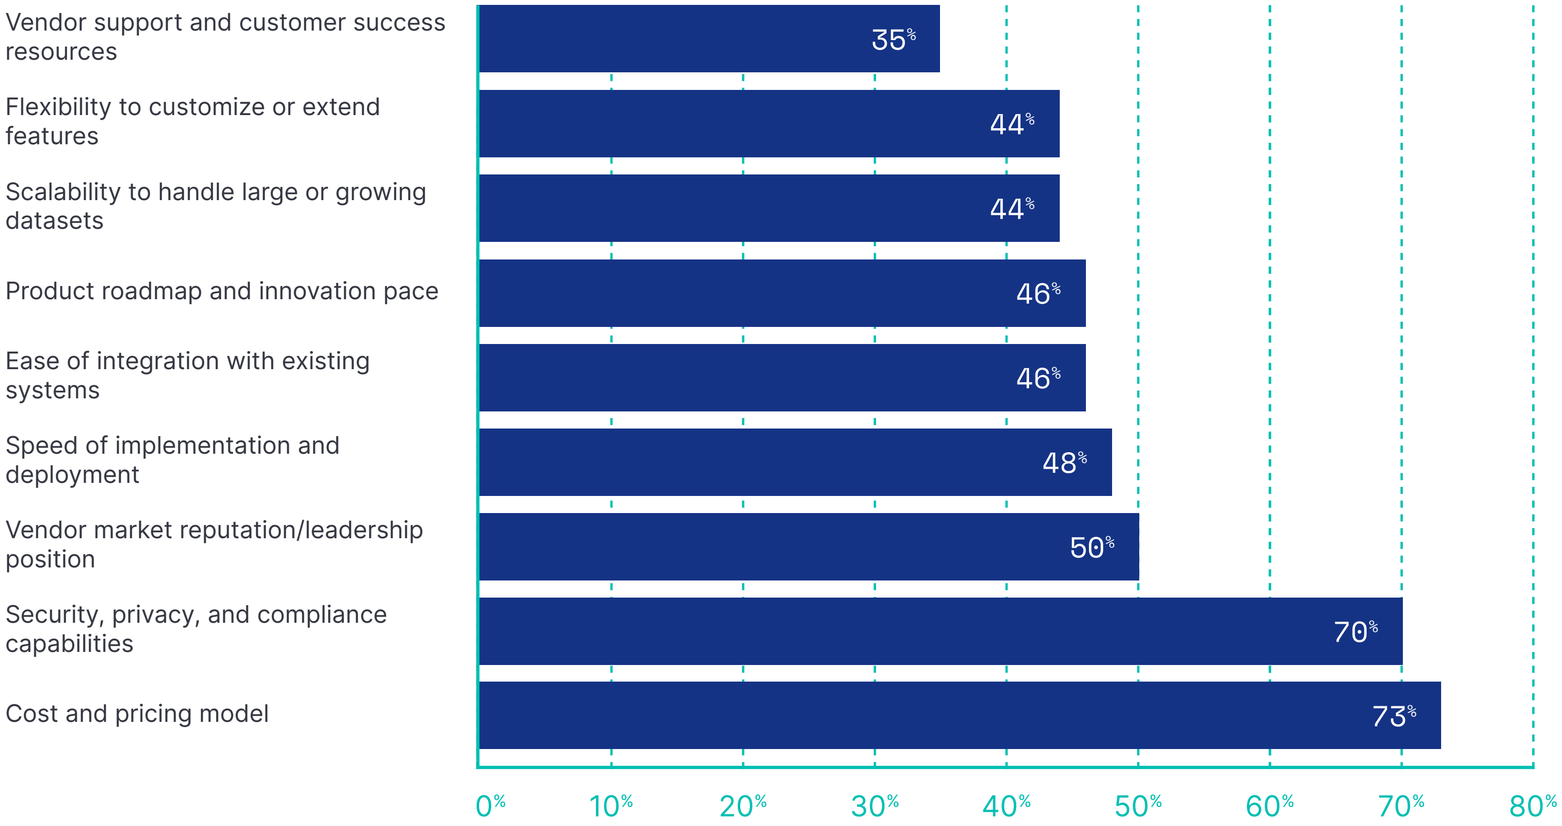

Figure 20. Answers to the question, “What are your criteria for evaluating AI search tools and/or vendors?”

Security concerns also vary by industry with highly regulated industries like healthcare (82%) and financial services (78%) citing security as their top priority for vendor selection.

Vendor research is evolving as LLMs enter the mix

Respondents mostly rely on traditional sources of information when deciding which search-powered AI vendor to select. Analyst firms are the most common, cited by 61%, followed by peer recommendations through professional networks (57%) and peer review sites (44%). Large language model (LLM) sources like ChatGPT garnered 37%, which is a big enough number to suggest that customer research methods are evolving.

Sources of information used in vendor selection

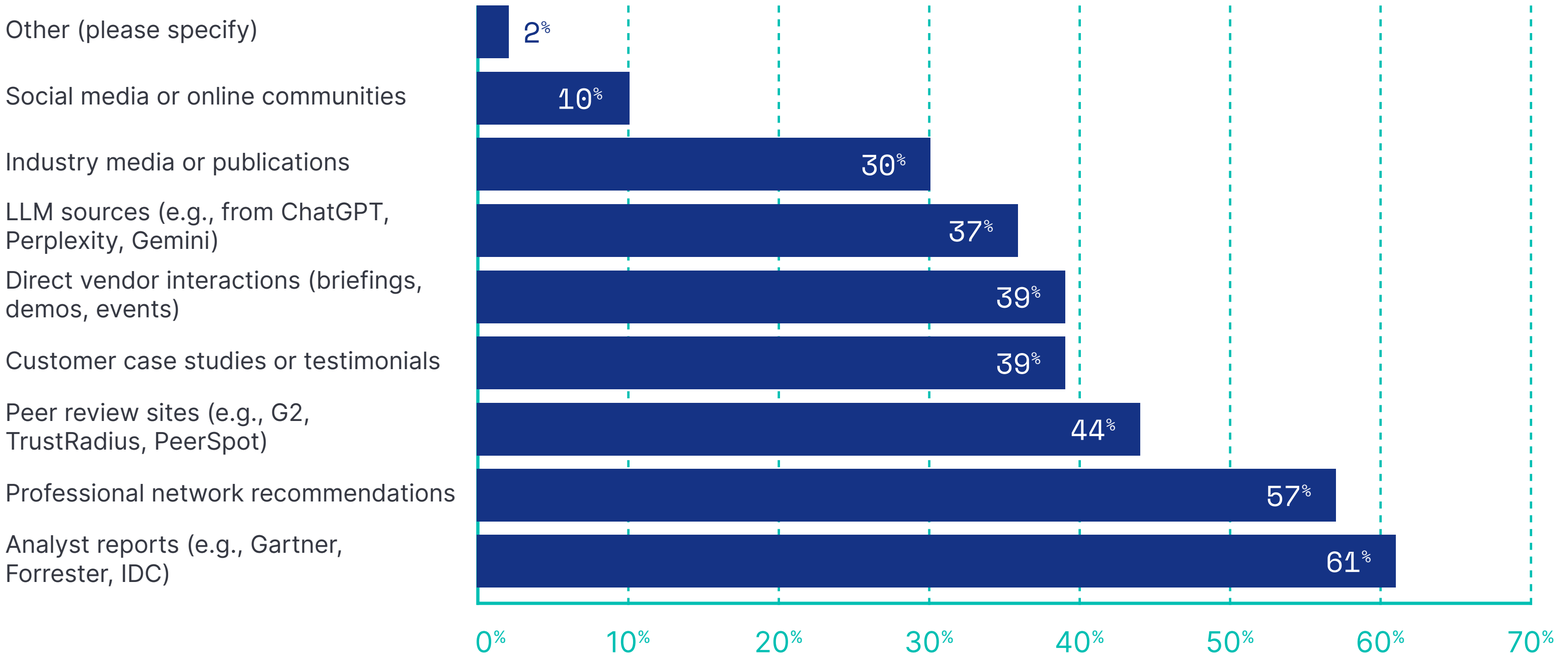

Figure 21. Answers to the question, "When assessing a vendor's market leadership, which sources do you rely on most?"

6 steps to implementing search and AI

The survey data suggests that there are ways to move forward and achieve success with search and AI, despite the obstacles. Here are six strategic recommendations:

1

Adopt a security-first architecture

The best practice is to deal with security up front and avoid treating it as a later-stage concern. Leading organizations establish security, privacy, and compliance frameworks before selecting a search and AI platform. This means defining policies for data classification, access controls, audit capabilities, and regulatory compliance. This approach should speed up deployment when the time comes and avoid the costly retrofitting that arises when security is an afterthought.

2

Plan for integration complexity

Organizations that successfully implement search-powered AI allocate 40%–50% of project resources to integration work. This includes API development, authentication/authorization setup, data pipeline construction, and workflow integration. In contrast, organizations that underestimate integration complexity show 3x higher rates of project delays.

3

Build skills and internal capabilities with vendor partnerships

It is not optimal to wait for internal talent to build specialized knowledge on vector embeddings, RAG architectures, semantic search tuning, and NLP model selection before starting a search-powered AI project. This will cause lengthy delays and perhaps even cause the project to be cancelled. Rather, as leading organizations do, it's best to combine vendor partnerships for initial implementation with aggressive internal skill-building programs. These companies hire strategically, upskill existing teams, and create centers of excellence that can eventually reduce vendor dependency.

4

Establish clear ROI frameworks early

Organizations with defined frameworks for ROI, including metrics and realistic timelines, show higher implementation success rates than those that lack such frameworks. The best practice is to establish a clear ROI framework early in the project lifecycle. It should include leading and lagging indicators across multiple value dimensions, such as productivity, cost reduction, quality improvement, and strategic enablement. It's also wise to assess the realism of the ROI timeline. The data show that aggressive ROI expectations (under three months) are associated with higher rates of project delays and scope changes.

5

Start with internal use cases, then expand

The data shows a higher rate of success for implementations that start with internal knowledge bases and then expand to customer-facing applications. This inside-out approach builds organizational confidence, develops use case expertise, and demonstrates ROI before tackling more complex deployments.

6

Invest in leadership expectation management and alignment from the start

Organizations that succeed with search and AI take the initiative to invest in executive communication. They establish regular executive education, create shared language around AI capabilities, and build realistic timeline expectations before projects begin.

How Elasticsearch operationalizes unused data

Agents are gaining traction driven by their potential to deliver efficiency gains and better customer experiences. But in practice, providing agents with the right context is difficult, especially when operating over messy, unstructured enterprise data. With deep, long-term investment in search relevance, retrieval augmented generation, and vector database, Elastic provides a complete set of capabilities to significantly simplify the approach to interacting with and building context-driven AI agents and applications that leverage the power of Elasticsearch.

Elasticsearch is a complete data platform for context engineering, purpose-built to deliver relevance at scale. It seamlessly combines the power of search and AI, making it easier than ever to unlock actionable insights. Elasticsearch is the most widely deployed open source vector database built to store and search both dense and sparse vectors and scale to billions of documents. With 5.5 billion downloads and a large open source community built around open source products, Elastic continues to empower developers and enterprises to build AI-native, search-powered applications.

Discover why Elastic was named a Leader in the Forrester Wave™: Cognitive Search Platforms, Q4 2025.

Read the report

Appendix: Respondents' demographic details

This vendor neutral, third-party research was independently conducted by ViB. ViB's best-in-class market research design and analysis methodology delivers the industry's most accurate insights. ViB leverages the industry's best practices and tools, incorporating extensive quality controls across the entire lifecycle from survey design to analysis, and presentation of findings.

Role

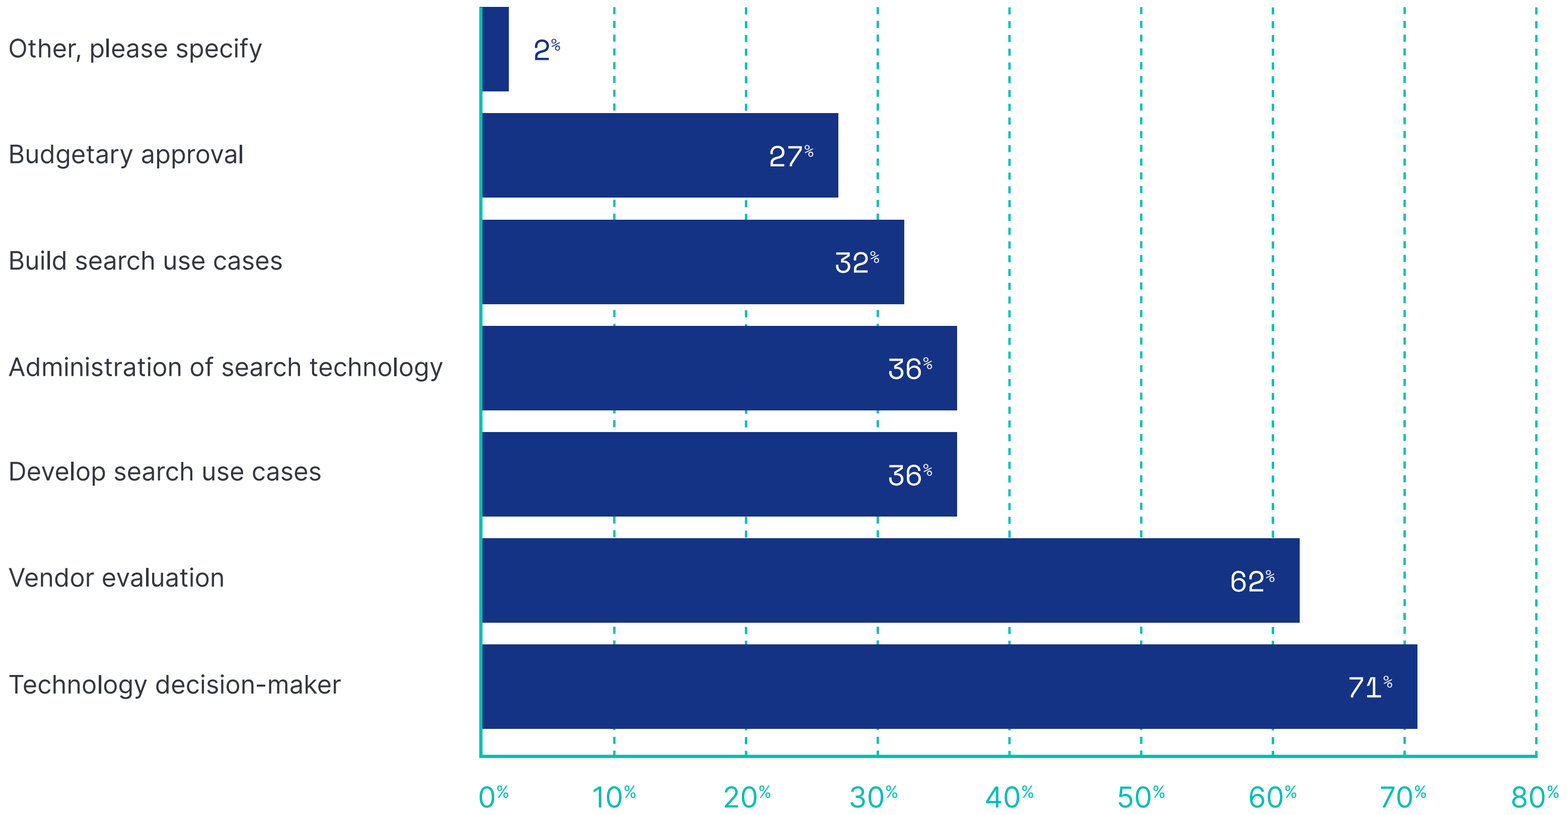

Size of organization (employees)

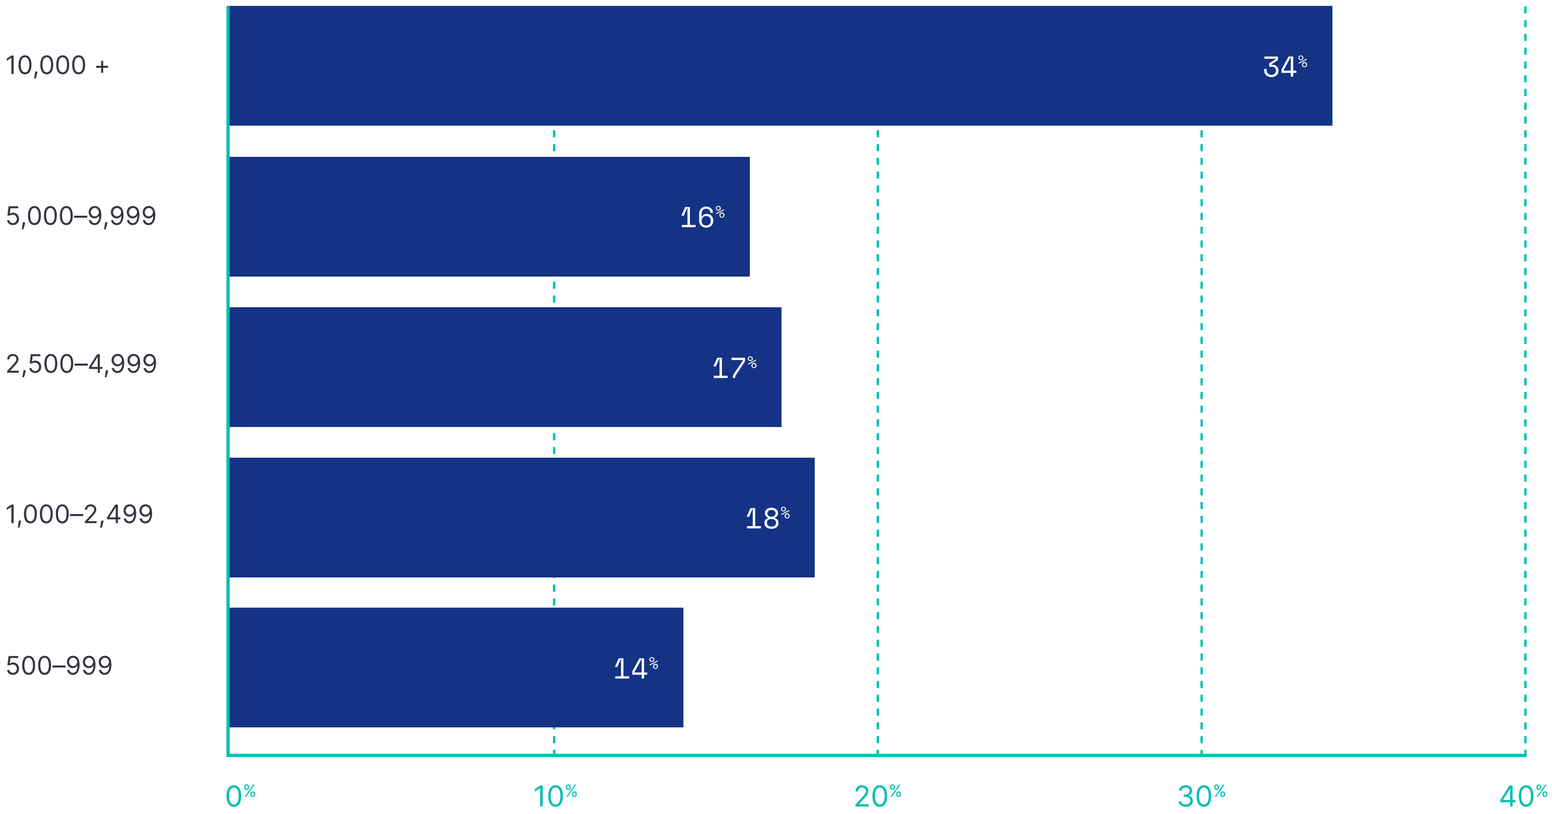

Industry