



Research



OpenSearch Awareness and Perceptions Report

Stephen Hendrick, *The Linux Foundation*
Adrienn Lawson, *The Linux Foundation*

Foreword by
Jonah Kowall

October 2024

OpenSearch Awareness and Perceptions



Contents

- Foreword 04
- Introduction 05
- The Open Source Difference 07
 - The use and importance of open source software 07
 - Confirming the cost benefit of open source software 08
- Experience with OpenSearch 10
 - Experience with technology areas relevant to OpenSearch..... 10
 - Experience with OpenSearch 11
 - Experience across OpenSearch technology areas 12
 - OpenSearch usage patterns 13
 - Deployment patterns that OpenSearch users adopt..... 13
- User Perceptions of OpenSearch..... 14
 - User feedback on product capabilities and strategies 14
 - User feedback on product capabilities and strategies of concern..... 15
 - User feedback on selected OpenSearch product capabilities and strategies 16
 - User feedback on pivotal OpenSearch strategies 18
 - How OpenSearch should increase its adoption..... 19
 - How OpenSearch can improve its visibility 20
 - Clarifying the difference between OpenSearch software and the Amazon OpenSearch service..... 21
 - User opinions of OpenSearch..... 22

- Conclusions 23
 - Continued investment in product capabilities is necessary to retain product momentum and drive adoption..... 23
 - Product development, offering training & certification, better documentation, and how OpenSearch drives business value, are key to driving product visibility and adoption..... 24
 - OpenSearch should transition to oversight by a vendor-neutral foundation..... 24
- Methodology 25
 - About the survey 25
 - Data.World access..... 26
- Respondent demographics 27
- About the authors..... 28
- Acknowledgements 29
 - About OpenSearch 30
 - About Linux Foundation Research..... 30

Foreword

Open Source is essential to all, as this important research study confirms. The findings validate why we created the OpenSearch project and subsequently established the OpenSearch Foundation under the Linux Foundation. The importance of this project is now enshrined as an immutable software foundation, but it requires the contributions of people's time along with sponsoring organizations to make it happen.

When we created the OpenSearch project three years ago, we wanted to create a pillar for creators to build upon. The community will always be focused on user-centric development. The OpenSearch project has had a focus on delivering use cases around observability, security, search, and building the next generation of AI-assisted experiences. While these could each be different projects, the unified nature of OpenSearch is a strength and will continue to create a platform to solve multiple use cases with the same data. This unification is a trend we see across observability, specifically regarding logging and security. This is reinforced in the research findings as a key area of focus for most users.

The community has created and continued to build on these important features and capabilities. With the changes in the foundation, we will continue to deliver better integration and interoperability with key open source technologies along with broadening the contributors to this important foundational technology.

As we go forward, the key tenets of the project will remain in focus: keeping APIs consistent with a limited amount of breaking changes while advancing on scalability and efficiency. These core goals of the project are well reflected in this research as we listen and work with the community. OpenSearch will enable builders and creators to imagine what is possible. Doing this at scale with lower costs is a key finding of this research aligned with the future of OpenSearch.

Thank you to all the contributors and sponsors who have made this project possible. Here's to continued innovation and building.

JONAH KOWALL

OpenSearch Leadership Committee



Introduction

Modern systems rely on microservices architecture and containerization to package and run software applications. As Docker initially popularized it in 2013, open source containerization gained significant traction with Google's release of an open source container orchestration platform named Kubernetes in 2015. Industry adoption of Kubernetes quickly accelerated, leading to CNCF's acceptance of Kubernetes and a significant buildout of the Kubernetes ecosystem, resulting in Kubernetes becoming the de facto standard for container orchestration.

Some of the operational benefits of container-based systems include resource efficiency, portability, scalability, better security, and cost efficiency. The development of container-based infrastructure as open source projects drives additional benefits that include engagement across the sizable open source community, development in the open, and a combination of collaboration and knowledge sharing that results in faster innovation.

The use of containers and microservices represents a more granular approach to resource development and allocation and involves the use of many and often ephemeral containers to support scalability requirements. The complexity of these environments demands far more sophisticated approaches to resource management and the tools to monitor and manage system operations. The design of observability tools provides insight into the internal state and behavior of systems, typically within the context of software applications and infrastructure. They help developers, operators, and engineers understand how a system is performing, diagnose issues, and ensure reliability.

The use of container-based systems is not without its challenges. Systems are complex, resource management is complicated,

security is difficult, and observability is critical to ensuring both performance optimization as well as incident / event management. Real-time application monitoring, distributed search, security analytics, and visualization are core to monitoring, analyzing, and managing complex container-based systems and applications. Like most cloud service providers, Amazon Web Services (AWS) recognized the need to deliver a unified product with the following capabilities to manage container-based applications and systems:

- **Search:** to find relevant data across applications, websites, and data lakes to monitor and debug applications and infrastructure
- **Observability:** to detect, diagnose, and remedy issues that affect the performance, scalability, or availability of applications or infrastructure
- **Security analytics:** to address security information and event management (SIEM) and help detect, analyze, and respond to security threats

The solution was the development of OpenSearch. First announced in January 2021, the OpenSearch project is an open source fork of Elasticsearch and Kibana. Released for production under the Apache License, version 2.0 (ALv2), it is open to contributions from the OpenSearch community. Use cases for OpenSearch include search applications, log-an-event data analysis, business intelligence, security information and event management, and time series data analysis. For more information on OpenSearch, see <https://opensearch.org>.

The OpenSearch Project engaged Linux Foundation Research in 2023Q4 to develop and execute an empirical research study to understand the awareness, perception, and competitive characteristics of OpenSearch. The target audience included respondents who met the following criteria:

- Must be a Linux Foundation subscriber, member, or contributor
- Must be familiar with the Linux Foundation (to ensure understanding of or involvement with open source)
- Must work full time or part time in the information technology field

- Must be in an IT role aligned with development, deployment, operations, or support
- Must have familiarity with one or more areas or tools that address search, observability, or security analytics

Linux Foundation Research's survey development occurred in 2023Q4, and the survey happened in 2024Q1, yielding 689 completed surveys. For more information about the survey methodology and survey demographics, see the *About the survey* chapter at the end of this report.



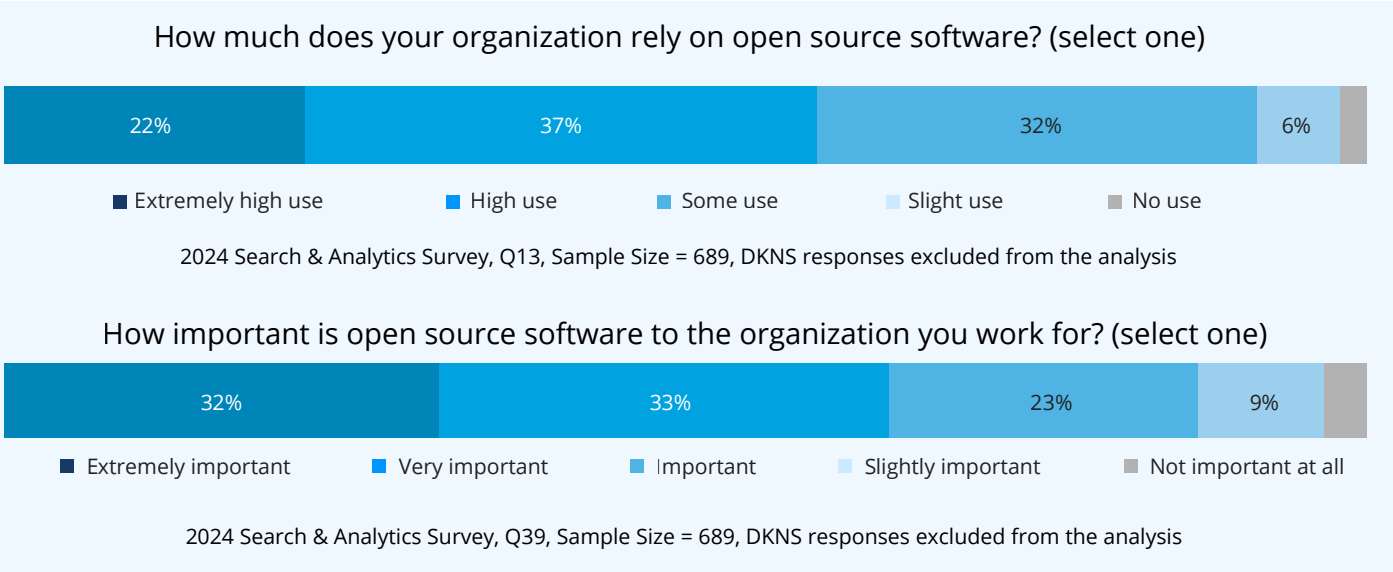
The Open Source Difference

The use and importance of open source software

Organizations' use of open source software is pervasive, and its importance is uniformly high. The top chart in Figure 1 shows that 92% of organizations use open source software at some, high, or extremely high levels. The addition of organizations that report a slight use of open source software drives this value to 98%. The use of open source software has remained consistently high over the last several years. Linux Foundation's research into SBOMs and cybersecurity readiness back in 2022 also shows open source use at 98%.¹

The bottom chart in Figure 1 also shows that 88% of organizations view open source software as either extremely important, very important, or important. Adding the 9% that identify open source software as slightly important to this total means that 97% of organizations view open source software as being important. The success and steadily increasing adoption of open source software over the last decade originated because of its cost-effectiveness, open standardization, and avoidance of vendor lock-in. More recently, open source has benefited from increasing developer support and collaboration, organization-sponsored contributions, Internet maturity, and the advent of cloud computing.

FIGURE 1: THE USE AND IMPORTANCE OF OPEN SOURCE SOFTWARE



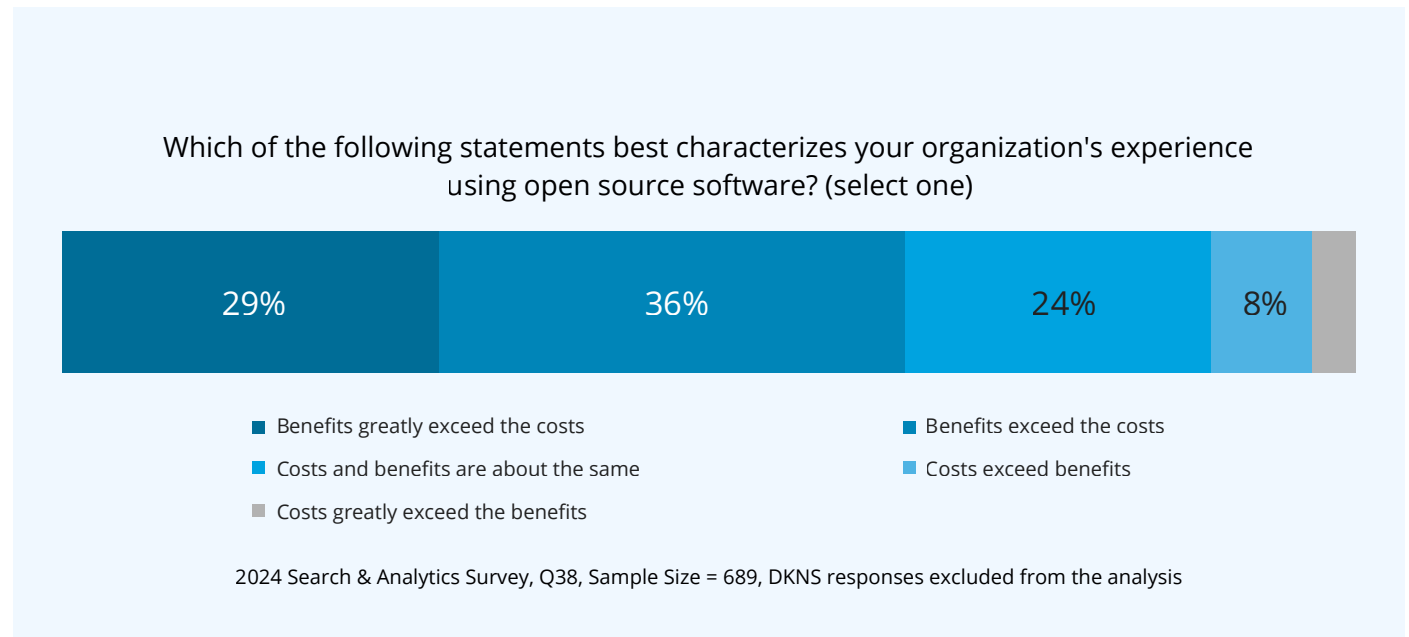
1. "Software Bill of Materials and Cybersecurity Readiness," Stephen Hendrick, January 2022.

Confirming the cost benefit of open source software

One of the most important attributes of open source software is that it is free to use. Cost benefits can also compound when factoring in free maintenance, upgrades, support, and added instances based on scalability requirements. Figure 2 examines the cost benefits of using open source software across organizations in our sample. It shows that 65% of organizations reported that open source benefits exceed costs, compared to 11% that reported that costs exceeded benefits. When decomposing Figure 2, we see that the 29% of organizations that reported that the benefits of open source greatly exceed the costs were 9.6x the 3% where the costs of open source greatly exceed the benefits. Similarly, the 36% of organizations where benefits exceeded the costs was 4.5x the 8% that experienced costs that exceeded benefits.

Some of the reasons why open source costs can exceed benefits originated from complex integration requirements with existing and proprietary software, lack of professional support, and insufficient documentation.

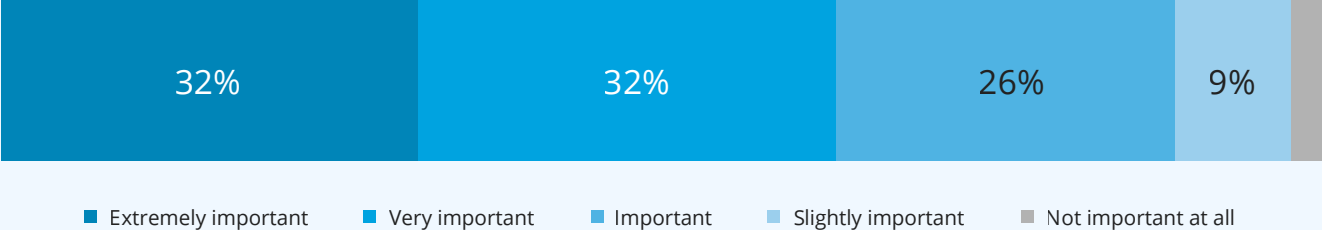
FIGURE 2: COST-BENEFIT CHARACTERISTICS OF OPEN SOURCE SOFTWARE




When we asked survey respondents more specifically about the importance of tools needed to support modern system operations, the results were nearly identical to the distribution of open source importance shown in Figure 1. Figure 3 shows that 89% of organizations view these tools and capabilities as either extremely important, very important, or important. The addition of those organizations that view these tools and capabilities as slightly important increases the total to 98%. The importance of these tools stems from the mission-critical nature of the systems they support combined with the added complexity of managing highly granular and ephemeral container-based systems.

FIGURE 3: THE IMPORTANCE OF TOOLS THAT SUPPORT MODERN SYSTEM OPERATIONS

How important are tools that address search, observability, and security analytics / SIEM to the organization you work for? (select one)



2024 Search & Analytics Survey, Q37, Sample Size = 689, DKNS responses excluded from the analysis



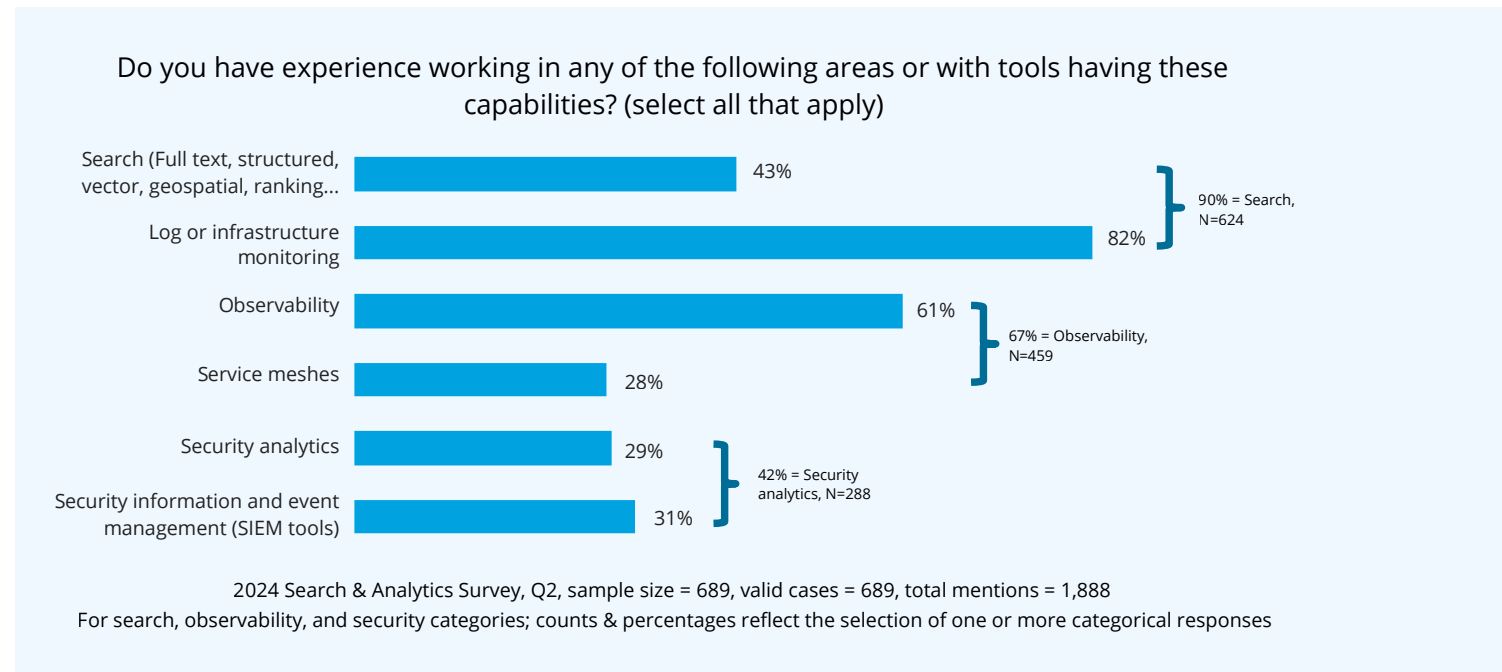
Experience with OpenSearch

To avoid bias, we made no references to OpenSearch during recruitment or fielding of the 2024 Search and Analytics Survey. During screening, we evaluated candidate experience with technology areas leveraged by OpenSearch to avoid naming the open source project.

Experience with technology areas relevant to OpenSearch

Respondents to the 2024 Search and Analytics Survey needed to have experience in one or more of the three technology areas currently addressed by OpenSearch and commercial products that are alternatives to OpenSearch. Figure 4 shows one of the initial questions in the survey, which identifies if a respondent had experience in one or more of these three areas (search, observability, and security analytics). This question also included a proxy for each of these areas in case respondents were unfamiliar with OpenSearch terminology. Figure 4 shows that out of a sample size of N = 689 completes, 90% had experience in search, 67% in observability, and 42% in security analytics.

FIGURE 4: SEGMENTING RESPONDENT EXPERIENCE IN SEARCH, OBSERVABILITY, AND SECURITY ANALYTIC TOOLS



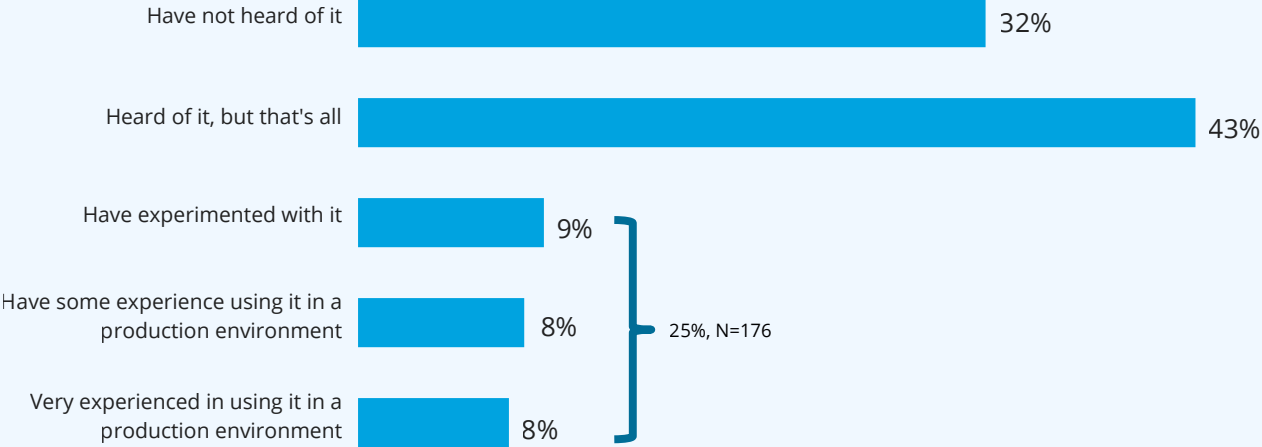
Experience with OpenSearch

Development of OpenSearch began in earnest at the beginning of 2022. This makes OpenSearch about 2.5 years old. Figure 5 shows respondent experience with the OpenSearch project, with 32% of respondents never having heard of OpenSearch and 43% having heard of OpenSearch but never having had any direct experience with it. This leaves just 25% (N = 176) who have experience or have experimented with OpenSearch. This is a compelling penetration rate for a young project entering into a variety of maturing markets.

The first mention of OpenSearch was in Q15 (Figure 5), where we asked respondents if they had experience with OpenSearch. Figure 5 shows that 75% of organizations have no experience with OpenSearch, and 25% have various degrees of experience. Of the 25% that had experience, 9% had only experimented with OpenSearch, 8% had some experience using OpenSearch, and 8% were very experienced using OpenSearch. For more information on the methodology behind identifying respondents with OpenSearch experience, see the methodology section in this report.

FIGURE 5: EXPERIENCE WITH OPENSEARCH

What is your experience with the open source OpenSearch software? (select one)

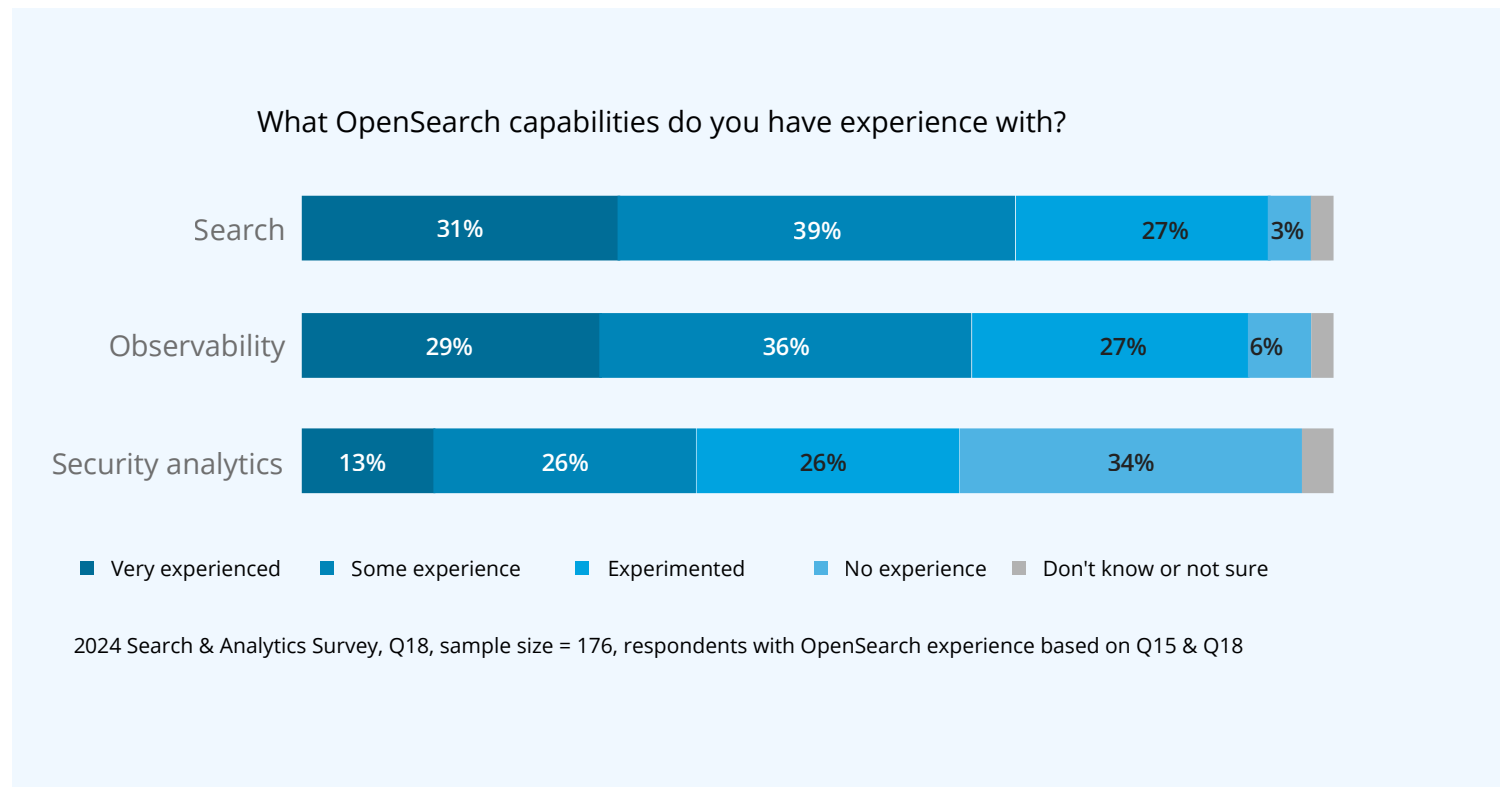


2024 Search & Analytics Survey, Q15, Sample Size = 689, respondents with OpenSearch experience based on Q15 & Q18

Experience across OpenSearch technology areas

For those respondents who reported experience or experimentation with OpenSearch, we then asked how much experience the respondent had in each of the three core areas currently supported by OpenSearch. Figure 6 shows that most OpenSearch users begin with search, presumably to support observability functions such as log monitoring and event data analysis. Figure 6 shows that 70% had experience with search, and 65% had experience with observability. Experience with security analytics is significantly less at 39% because most organizations already have identity and access management and /or SIEM tools in place. However, observability and security analytics are rapidly expanding, and evolving market segments can provide an important avenue for OpenSearch's future growth.

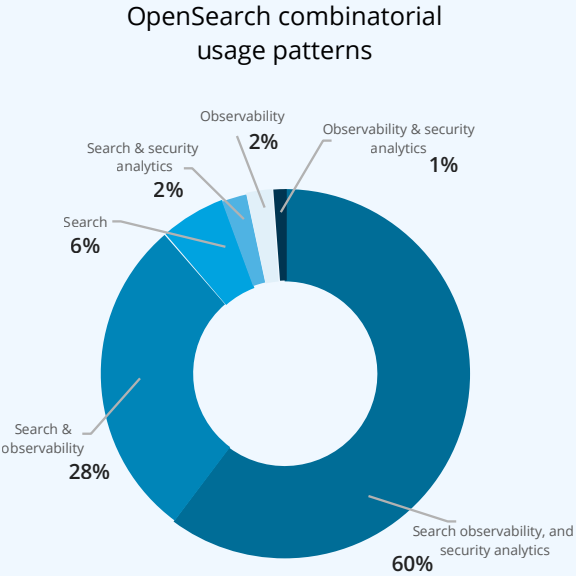
FIGURE 6: RESPONDENT EXPERIENCE WITH OPENSEARCH CAPABILITIES



OpenSearch usage patterns

OpenSearch offers seven different combinatorial ways to use its three core capabilities (see Figure 6). Figure 7 shows that nearly all OpenSearch users (94%) use two or more capabilities, with most users (68%) using all three capabilities. Figure 7 also shows that 96% of users rely on search, but just 6% use search alone. The reason for this is that search is a foundational capability for activities such as observability and security analytics.

FIGURE 7: OPENSEARCH COMBINATORIAL USAGE PATTERNS

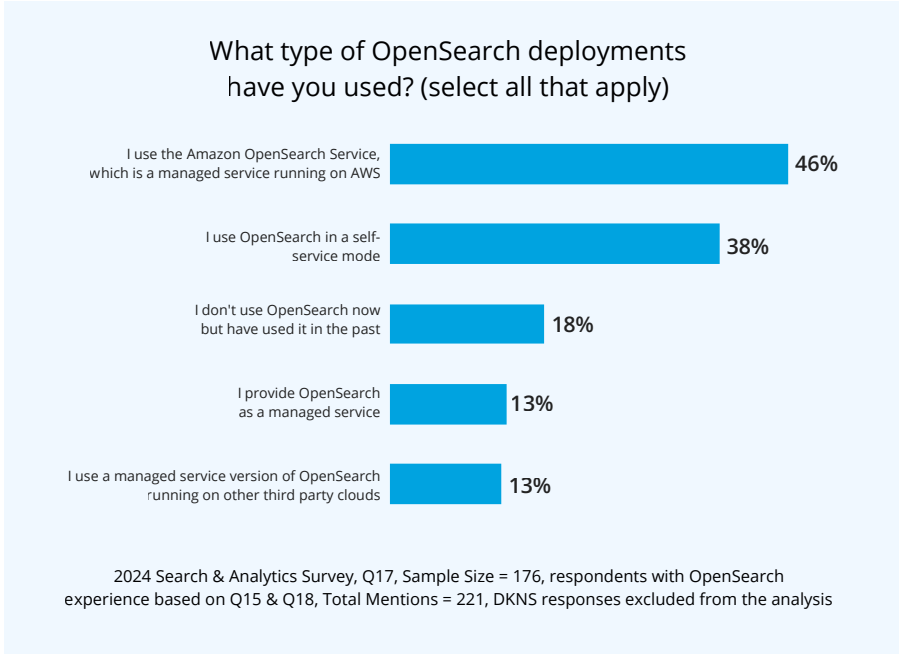


2024 Search & Analytics Survey, Q18, Sample Size = 176, respondents with OpenSearch experience based on Q15 & Q18

Deployment patterns that OpenSearch users adopt

The patterns organizations use to consume OpenSearch typically involve one of two methods: consuming it as software as a service (SaaS) (59%) or as an application running their data center (38%). Figure 8 shows that 46% of organizations use the Amazon OpenSearch service, a managed service running on AWS. An additional 13% of organizations consume OpenSearch running as a managed service on another cloud. Despite the multiple response structure of Q17 in Figure 8, the distribution of responses other than those focused on OpenSearch consumption implies that organizations consume OpenSearch in just one way.

FIGURE 8: OPENSEARCH DEPLOYMENT PATTERNS



2024 Search & Analytics Survey, Q17, Sample Size = 176, respondents with OpenSearch experience based on Q15 & Q18, Total Mentions = 221, DKNS responses excluded from the analysis

User Perceptions of OpenSearch

This section of the report focuses on user feedback regarding product characteristics and opportunities (Figures 9 to 12) as well as key questions regarding product positioning (Figures 13 to 16). Figures 9 to 12 were matrix questions, which scored attributes and strategies from not important or well (score = 1) to extremely important or well (score = 5). We can find the mean score for each attribute or strategy at the right of each chart. The data in the chart is in descending order based on the percentage of organizations responding extremely important / well and very important / well.

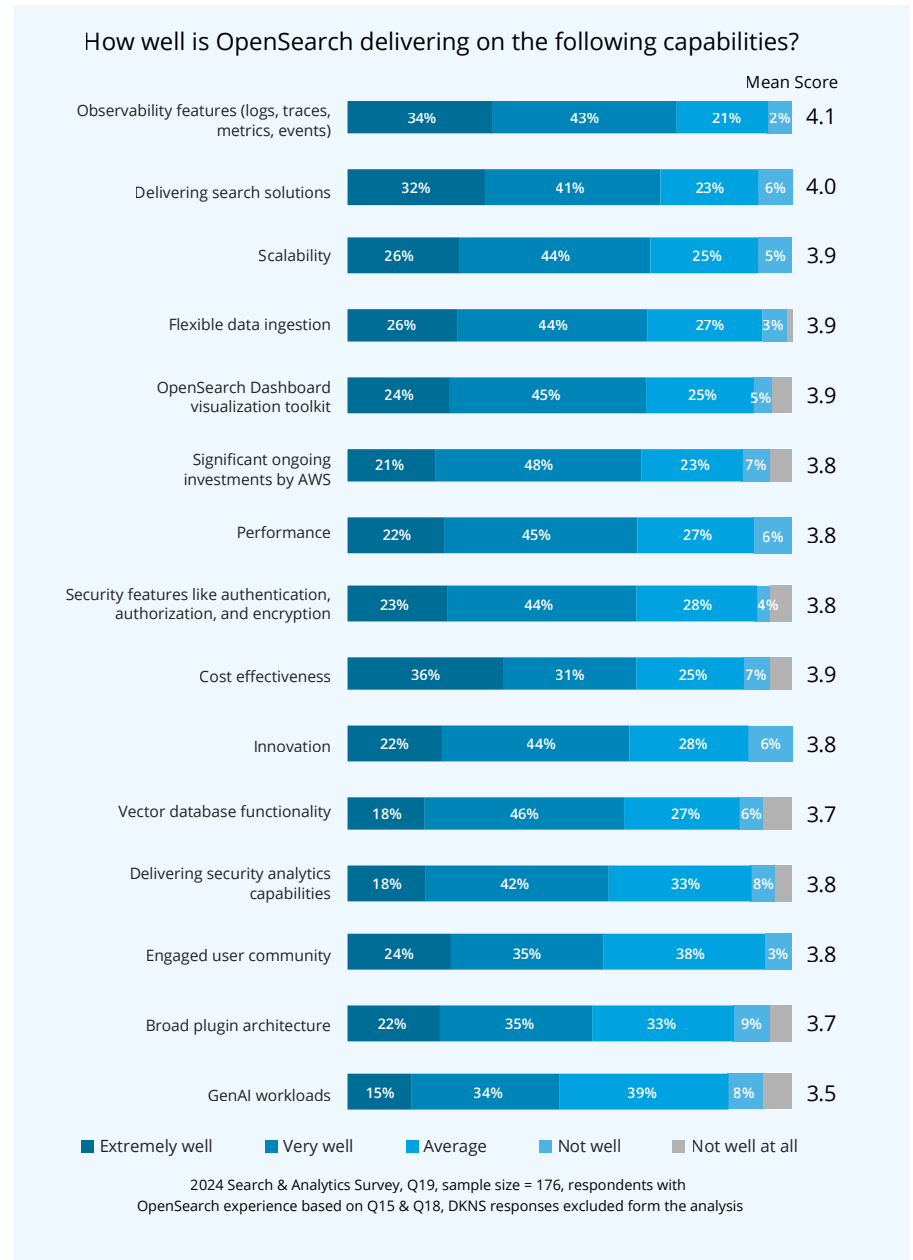
User feedback on product capabilities and strategies

Figure 9 shows the count of positive scores (extremely well and very well), which is far greater than the number of negative scores (not well and not well at all). Leading product attribute scores include high marks for observability features (76% positive and a mean score of 4.1) and the delivery of search solutions (73% positive and a mean score of 4.0). Additional product attributes, including scalability, flexible data ingestion, and visualization tools, also have a high rating of around 70% positive and a mean score of 3.9.

Even the lowest score for GenAI workloads included feedback that was 49% positive and only 11% negative, with a mean score of 3.5, which is above average. Especially interesting was the 36% of organizations that scored cost-effectiveness as extremely well. This was the highest score in the extremely well category, and while we have come to expect open source products as being cost-effective, the cost leadership exhibited here suggests capabilities that deliver benefits of high value.

Feedback from Figure 9 collectively shows that users are pleased with the capabilities offered by OpenSearch, although areas such as GenAI support, community engagement, and OpenSearch's plug-in architecture could benefit from more attention.

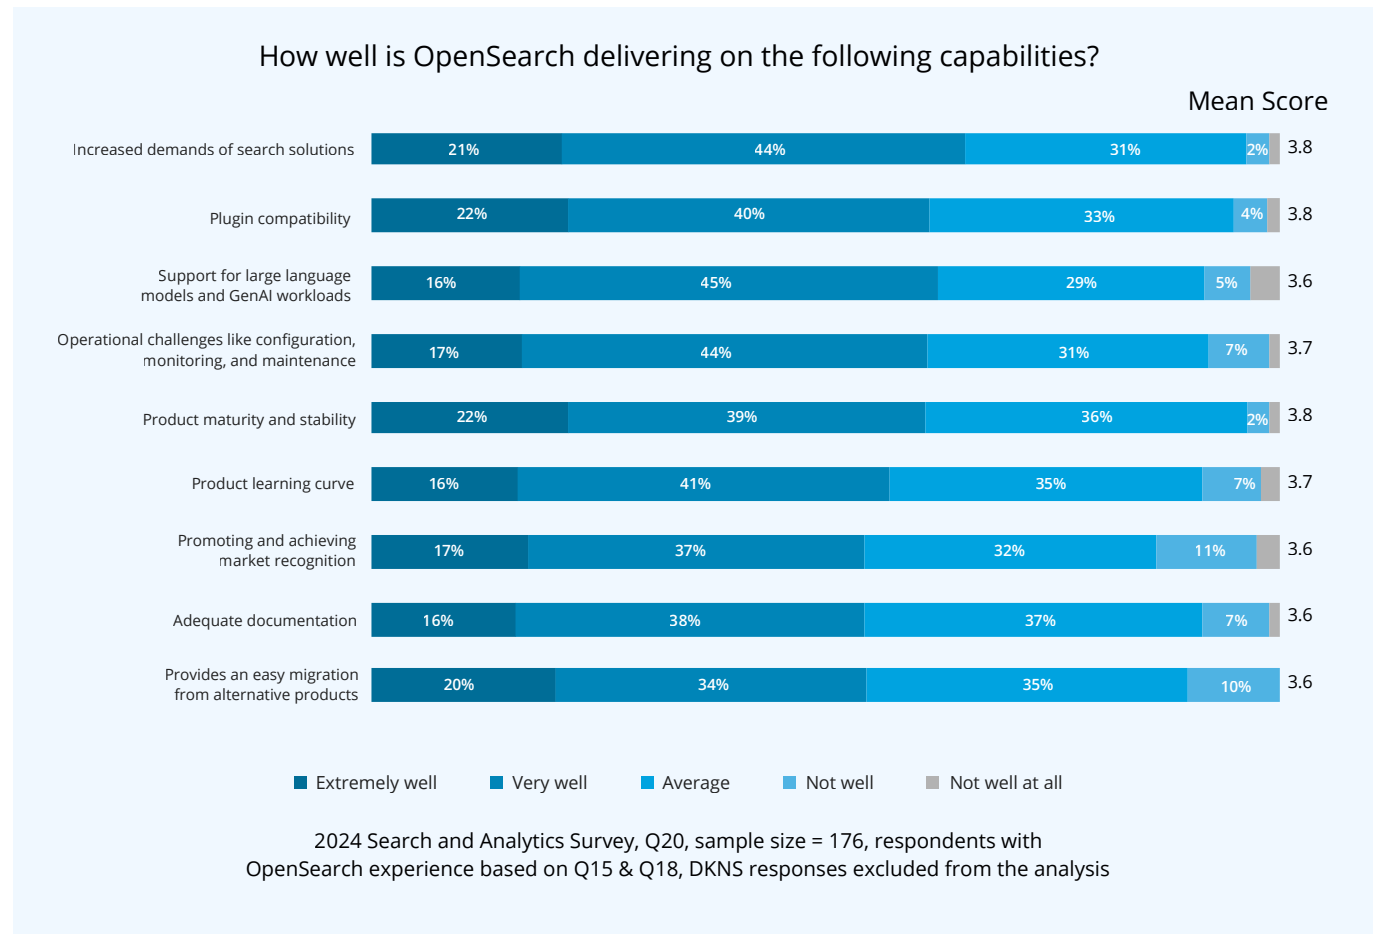
FIGURE 9: USER FEEDBACK ON SELECT PRODUCT CAPABILITIES AND STRATEGIES



User feedback on product capabilities and strategies of concern

The identification of product capabilities and strategies shown in Figure 10 occurred during discussions with the OpenSearch project team. This list includes those areas where the OpenSearch team believed users might express concern. While the positive scores, across Figure 10, ranged from 65% to 54%, our interpretation of this chart is that product capabilities and strategies that need additional attention are those with lower positive scores and higher negative scores. Using this criterion, the areas that require emphasis on the part of OpenSearch include reducing the product learning curve, better promotion of the product to achieve market recognition, improved documentation, and improved migration tools.

FIGURE 10: USER FEEDBACK ON PRODUCT CAPABILITIES AND STRATEGIES OF POTENTIAL CONCERN



User feedback on selected OpenSearch product capabilities and strategies

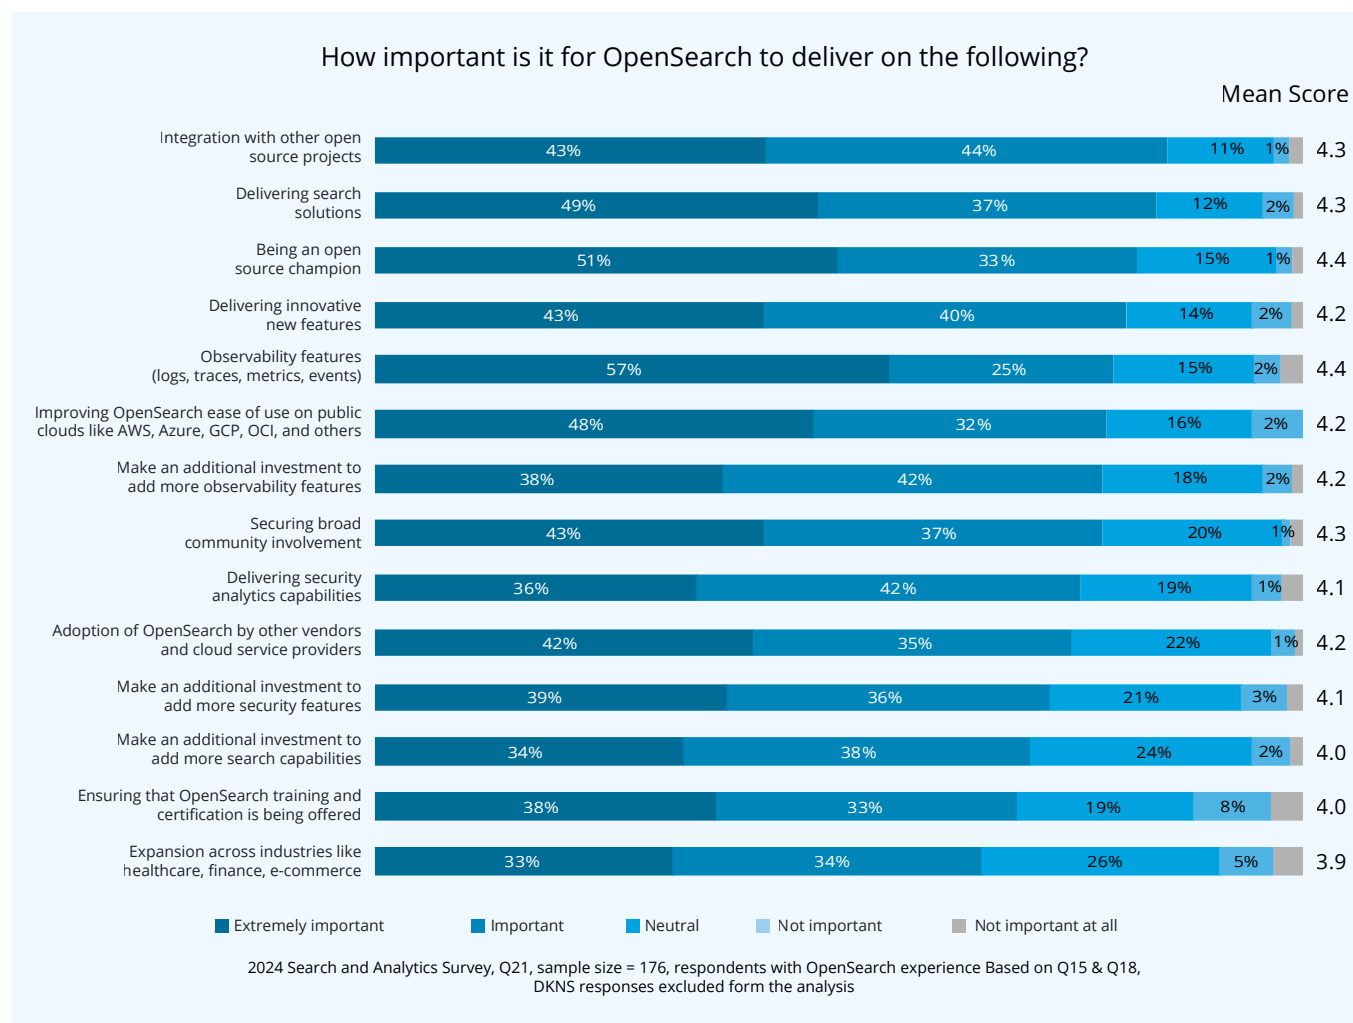
The OpenSearch project team was also interested in having users review and score a variety of product capabilities and strategies to determine which ones might take precedence. Figure 11 displays the importance that users attach to these actions. Positive scores range from 67% to 87%, and mean scores range from 3.9 to 4.4. Therefore, many of these potential actions have a high ranking, and the differences between the importance attached to them are marginal.

Leading strategies include integration with other open source projects (87% positive) and being an open source champion (84% positive). Integration with other open source projects represents a material opportunity because of the significant increase in scope that comes from bringing complementary projects together while avoiding the sizable investment in writing new code. Being an open source champion means having a deep commitment to the open source ethos and actively working to advance open source through advocacy, contribution, and community involvement.

Leading product capabilities include delivering search solutions (86% positive), delivering innovative new features (83% positive), and enhanced observability features (81% positive). The importance that users attach to nonspecific product-wide opportunities indicates a community that is not just deeply satisfied with OpenSearch but one seeking continued investment in core product areas to drive continued product success.



FIGURE 11: USER FEEDBACK ON THE IMPORTANCE OF PRODUCT CAPABILITIES AND STRATEGIES



User feedback on pivotal OpenSearch strategies

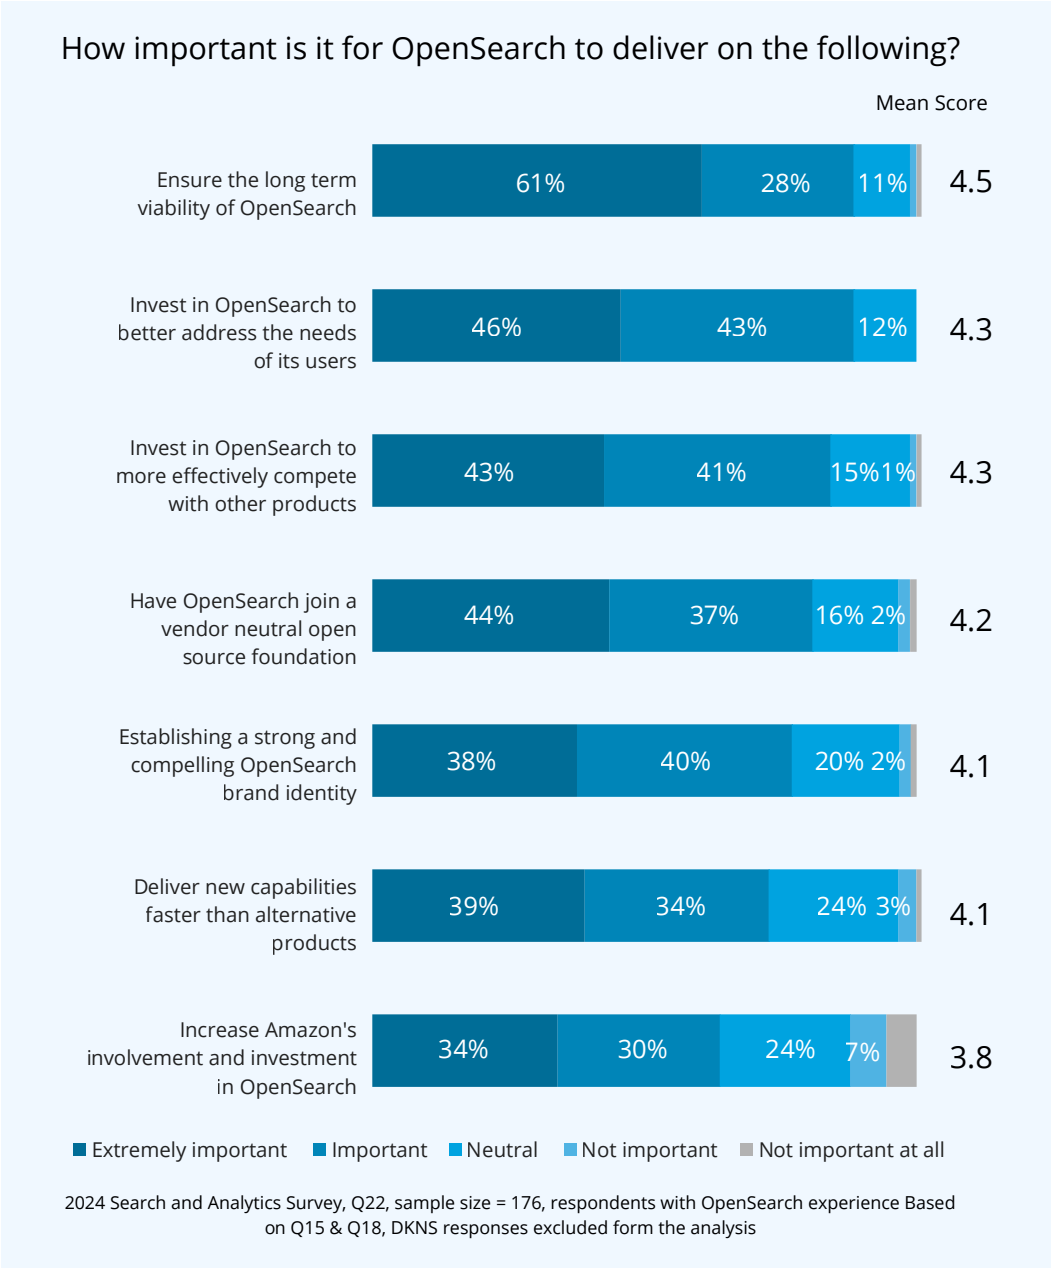
Over the last two years, AWS has made a significant and consistent investment in OpenSearch. Looking forward, AWS is exploring strategies that will define OpenSearch’s future road map. Figure 12 shows user feedback on some of these pivotal strategies.

Figure 12 shows extremely strong user sentiment in ensuring the long-term viability of OpenSearch, with 61% seeing this strategy as extremely important and 28% as important for a total of 88%. The strategy received the highest mean score (4.5) across Figures 9 to 12. The OpenSearch community clearly believes in the product and wants its strong open source ethos to persist.

The open source user community also wants to see continuing investment in OpenSearch to address user needs and more effectively compete with other products. Investing in OpenSearch to better address the needs of users was extremely important to 46% of users and important to 43% for a combined total of 88%. Similarly, investment to compete with other products more effectively was extremely important to 43% of users and important to 41% for a combined total of 84%.

One of the more fascinating findings in Figure 12 is that users are beginning to lean toward stronger community-based open source governance. While most users (64%) want to see increased involvement in investment by AWS, 7% do not see increased AWS involvement as important, and 5% report that this is not important at all. While increased AWS involvement by users is important, so is the objective to have OpenSearch join a vendor-neutral open source foundation, which was either extremely important or important to 81% of OpenSearch users.

FIGURE 12: USER FEEDBACK ON PIVOTAL OPENSEARCH STRATEGIES



How OpenSearch should increase its adoption

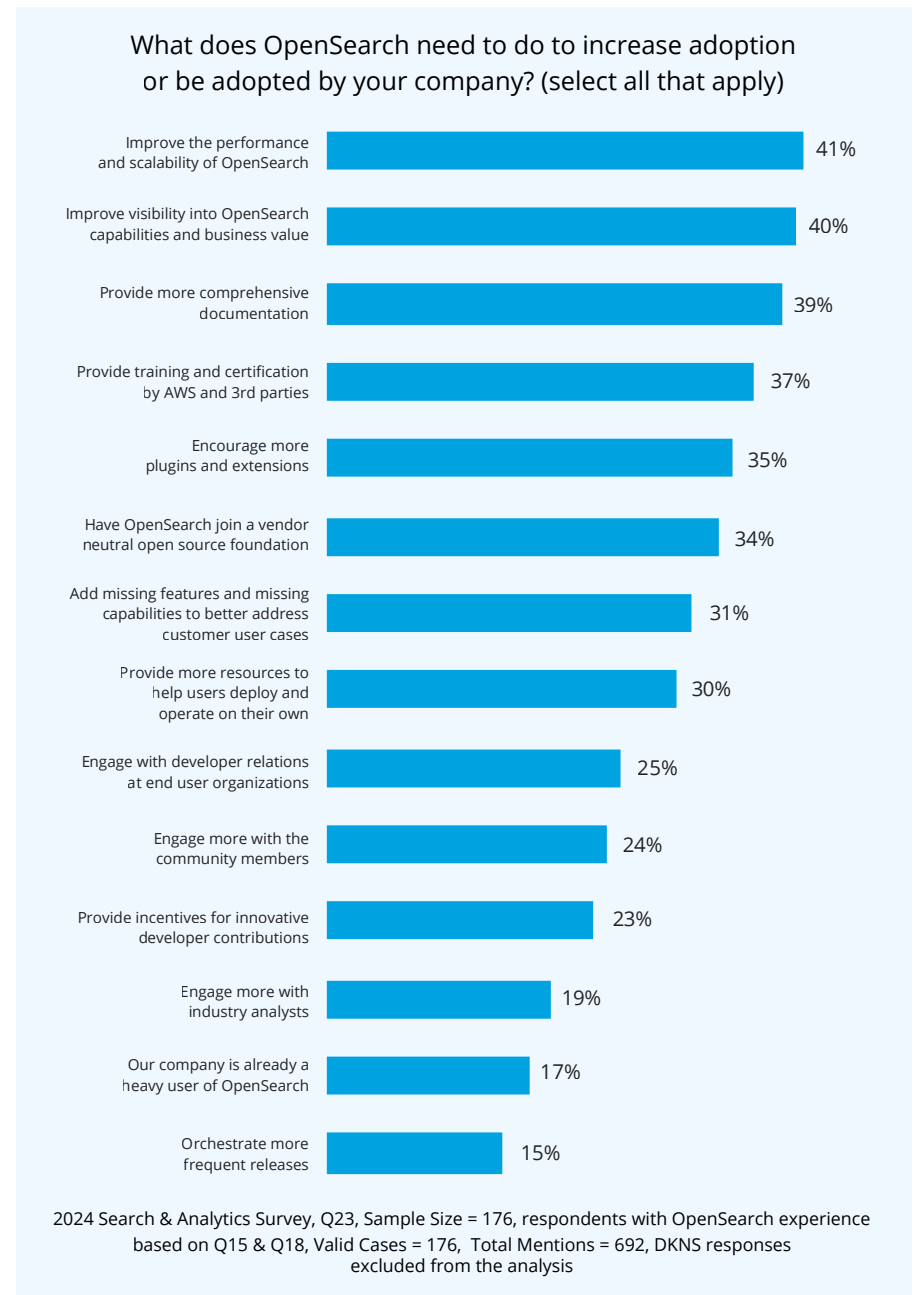
The OpenSearch community is small but strongly engaged. This begs a question. How should OpenSearch increase the size of its user community? Figure 13 addresses this question, yielding results that are sometimes enlightening and unusual.

The leading response in Figure 13, which 41% of users reported, was to improve the performance and scalability of OpenSearch. While Figure 9 shows scalability with a 70% positive score and performance with a 67% positive score, performance and scalability are complex issues. Data volumes have and will continue to increase at an accelerating rate—and this was before the advent of GenAI. Large language models are data intensive, with scalability key to handling large data volumes and distributed search and query optimization necessary for driving fast and accurate results. As organizations seek to harness the transformative capabilities of GenAI, optimizing scalability and performance will continue to be a key organizational objective.

The second leading response to the question of how to increase search adoption was to improve visibility into OpenSearch capabilities and business value, which 40% of users reported. One of the best ways to showcase product capabilities is by evaluating the business value that a product delivers to the organization. We can do this both qualitatively and quantitatively, and it is useful in explaining what a product can do for an organization and how effectively the product is performing this function.

Product documentation (39%) as well as training and certification (37%) were also top of mind for OpenSearch users. Improved documentation helps make complex products more approachable, and training and certification accelerate knowledge transfer and validate user competency. It is important to shortlist both of these actions, but training and certification should take priority because we can outsource them to third parties.

FIGURE 13: HOW OPENSEARCH SHOULD INCREASE ITS ADOPTION

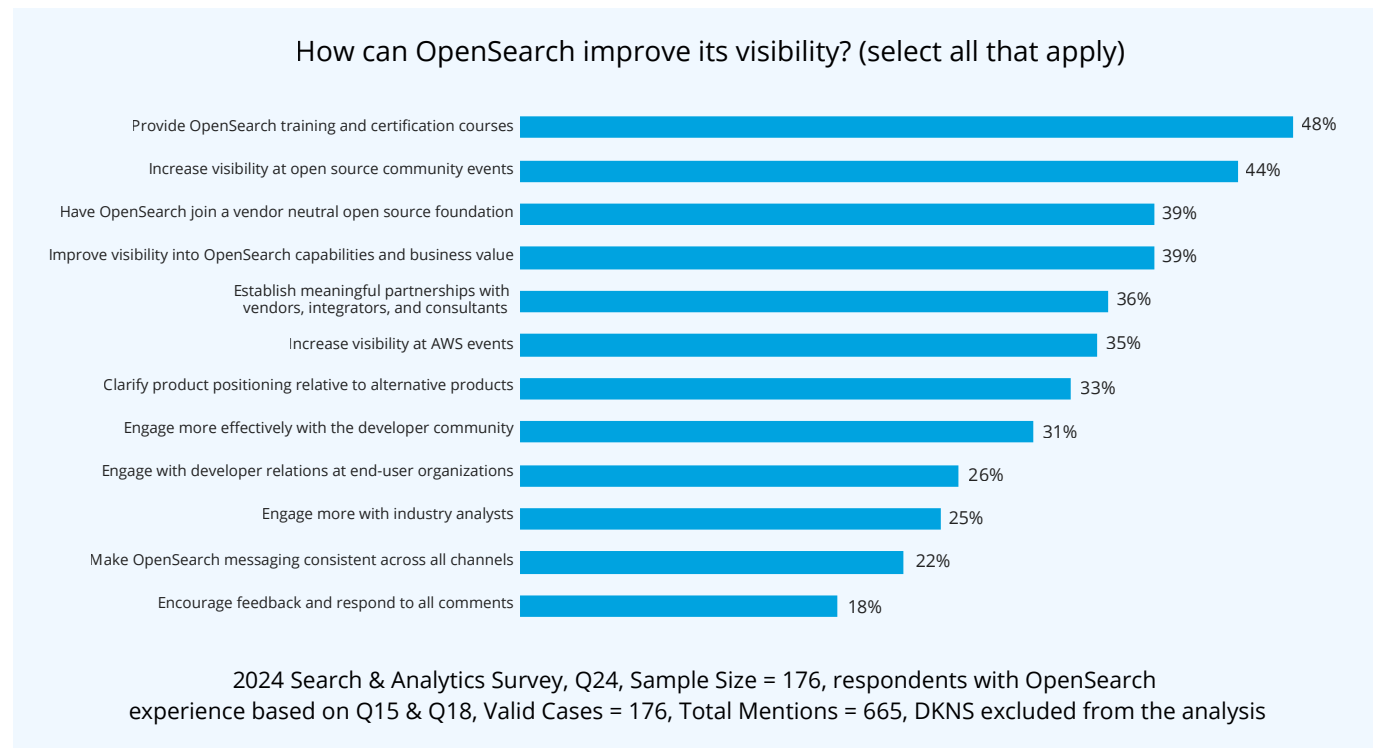


How OpenSearch can improve its visibility

Improving visibility and awareness are key prerequisites to driving product adoption. Figure 14 explicitly looks at a variety of techniques for improving OpenSearch visibility. User feedback confirmed that providing OpenSearch training and certification courses (48%), increased visibility at open source community events (44%), having OpenSearch join a vendor-neutral open source foundation (39%), and improving visibility into OpenSearch capabilities and business value (39%) were the leading ways to drive increased OpenSearch visibility.

Providing OpenSearch training and certification was the leading way in Figure 14 to improve visibility, and as we saw earlier, it helped drive adoption in Figure 13. Increased visibility and open source community events have time and budget implications for AWS but are otherwise easily achievable. Having OpenSearch join a vendor-neutral open source foundation is another theme we saw surface in Figure 13. We see this theme again in Figure 14. Finally, improving visibility into OpenSearch capabilities and business value is a recurring theme from Figure 13 but can be a cost-effective way to address visibility through marketing with resources that can reach a wide audience with a single investment.

FIGURE 14: HOW OPENSEARCH CAN IMPROVE ITS VISIBILITY



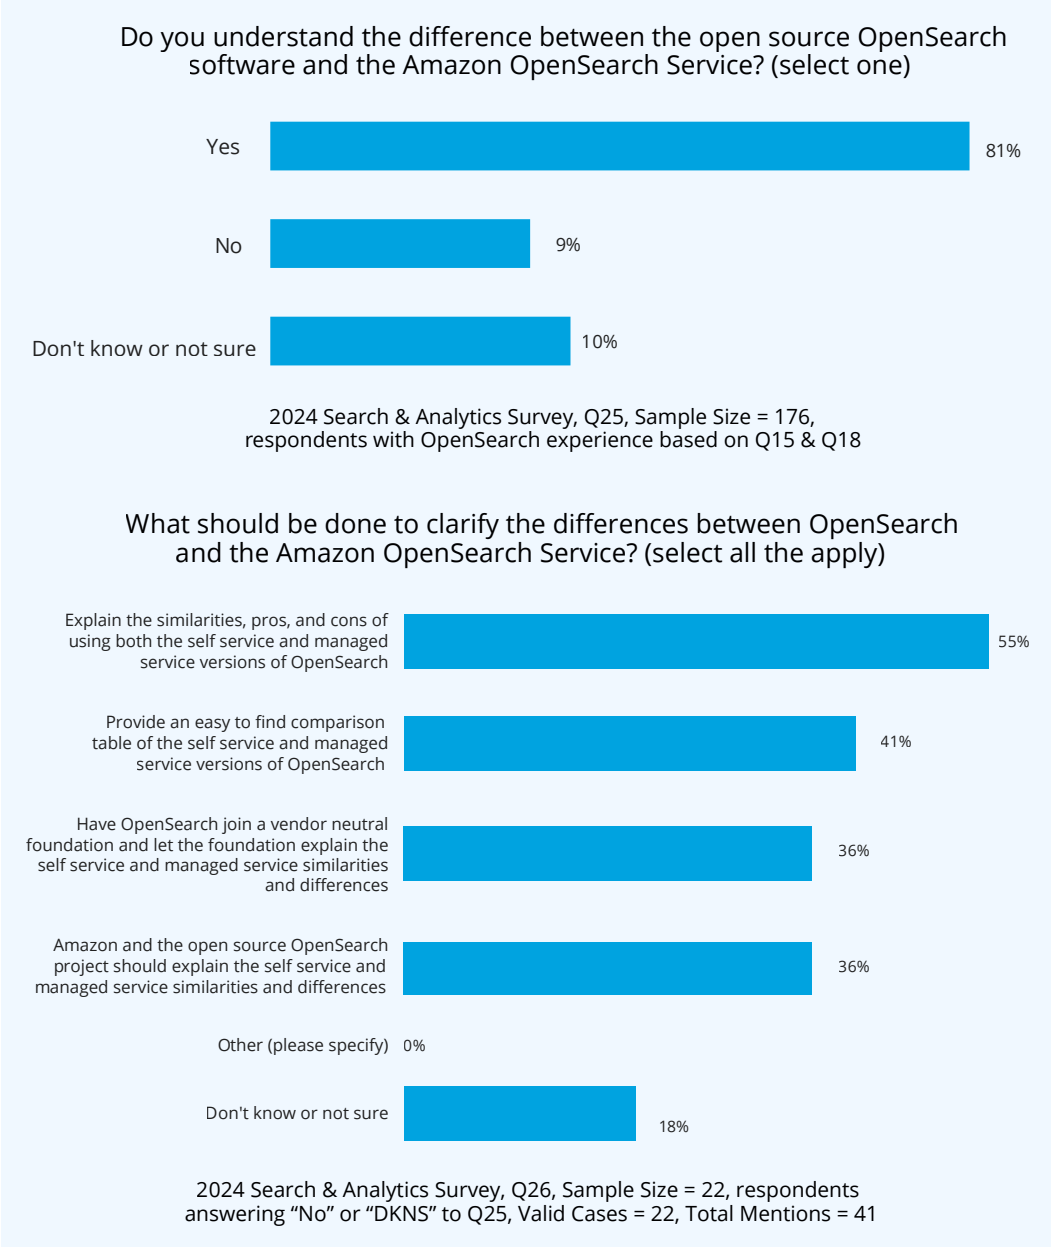
Clarifying the difference between OpenSearch software and the Amazon OpenSearch service

AWS has found that in some cases, there is confusion between OpenSearch the product and OpenSearch the service. This survey investigated to what extent there was this confusion and, when this confusion existed, what ways users were recommending to resolve it.

The top panel of Figure 15 queried users to determine if they understood the difference between OpenSearch software and the AWS OpenSearch service. The result was that 81% understood the difference while 19% did not. This confusion undoubtedly stems from users not appreciating that the OpenSearch software is also an open source project, meaning that the software is free for anyone to download. However, AWS has implemented the OpenSearch SaaS, which is an AWS SaaS product offering for those wanting to bypass OpenSearch management and maintenance.

The bottom panel of Figure 15 indicates that the primary solution to the issue, which 55% of users identified, is for Amazon to provide a more visible and clear explanation of OpenSearch in self-service mode or as a managed service. Secondary recommendations include providing an easy-to-find comparison table comparing self-service and managed service alternatives (41%), having OpenSearch join a vendor-neutral foundation and let them explain it (36%), and finally having Amazon and the OpenSearch project explain the differences (36%).

FIGURE 15: CLARIFYING THE DIFFERENCE BETWEEN OPENSEARCH AND OPENSEARCH SERVICES

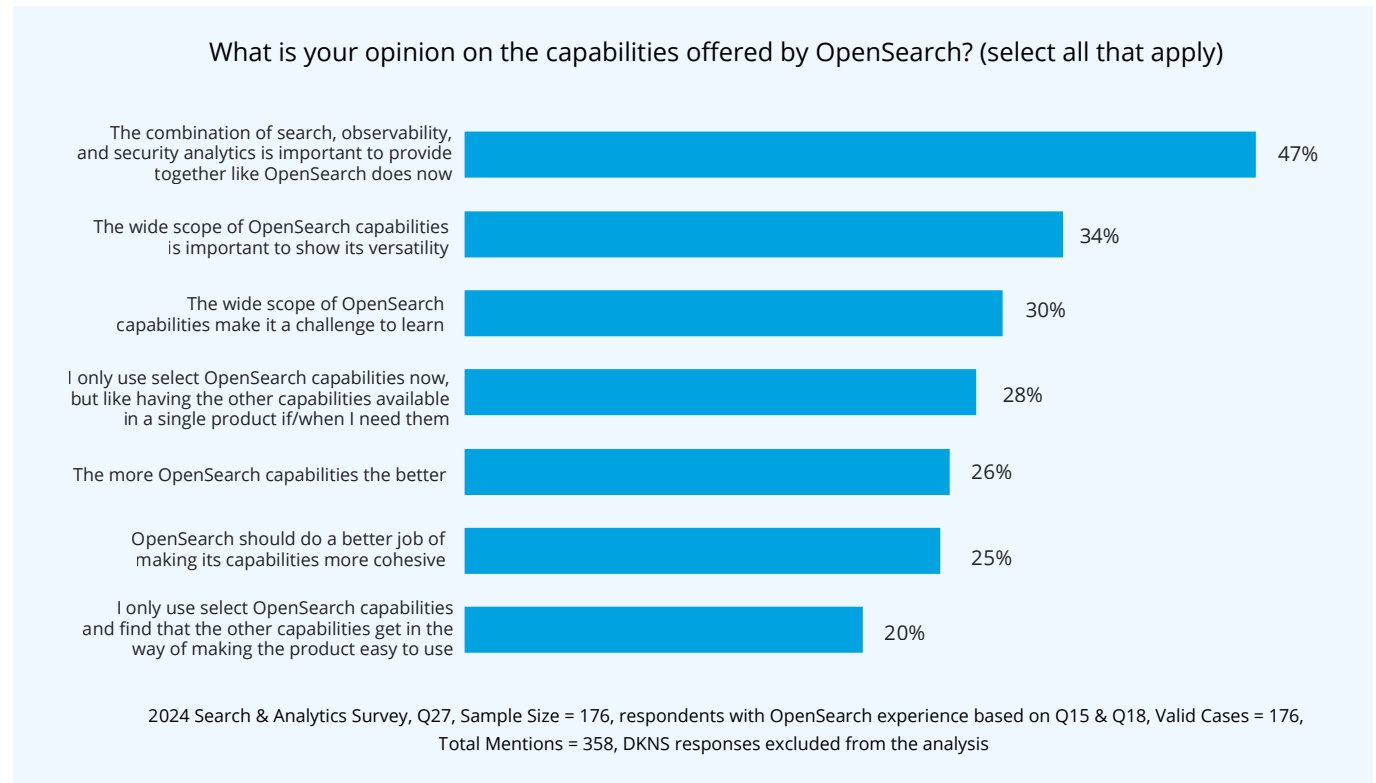


User opinions of OpenSearch

The design of the final question asked to OpenSearch users was to elicit recommendations on whether OpenSearch should be one expansive product or several products for use either independently or together.

The most compelling user feedback shown in Figure 16 was that 47% of users believe that the combination of search, observability, and security analytics is important to provide together, which is what OpenSearch does now. While this might cause a casual observer to infer that 53% of users were in favor of breaking the product apart, this is not the case. Other responses in Figure 16 include 34% who believe that the wide scope of OpenSearch capabilities is important to show its versatility, 28% who only use select OpenSearch capabilities but like having the capabilities available in a single product when needed, and 26% who said the more OpenSearch capabilities the better. Looking across all the above responses, we find that 81% of OpenSearch users in this sample selected one or more of these four responses, which all suggest keeping its capabilities together in one product.

FIGURE 16: USER OPINIONS OF OPENSEARCH



Conclusions

Continued investment in product capabilities is necessary to retain product momentum and drive adoption

From a product functionality standpoint, OpenSearch scored extremely well in both the user perceptions analysis and market analysis. This is significant given the brief tenure of OpenSearch in the market (~2.5 years). In the user perceptions analysis, Figure 9 showed high positive scores (the sum of extremely well and very well responses) for observability features (76%), delivery of search solutions (73%), scalability (70%), and flexible data ingestion (70%). In the market analysis, OpenSearch scored consistently well in each of its core domains—search, observability, and security analytics—despite the limited time in the market for OpenSearch compared to other products.

This suggests an excellence in product engineering that we should recognize and sustain. Achieving product leadership and penetration requires effective product development and marketing. According to our sample, OpenSearch has already solved what we consider to be the hardest problem: achieving product excellence.

Areas where users commented that OpenSearch could improve its functional capabilities, as discussed in Figure 10, include reducing the product learning curve, better promotion of the product to achieve market recognition, improved documentation, and improved migration tools.

When asked for feedback on a variety of OpenSearch products and market growth strategies, the leading user responses from Figure 11 were integration with other open source projects (87% positive), delivering search solutions (86% positive), being an open source champion (84% positive), delivering innovative new features (83% positive), and enhanced observability features (81% positive).

Users are also in favor of keeping all of OpenSearch's key capabilities within one product. Figure 16 shows that 47% of users believe that the combination of search, observability, and security analytics is important to provide together—which is what OpenSearch does now. Figure 16 also shows that 34% believe that the wide scope of OpenSearch capabilities is important to show its versatility, 28% only use select OpenSearch capabilities but like having the capabilities available in a single product when needed, and 26% said the more OpenSearch capabilities the better. Looking across all the above responses, we find that 81% of OpenSearch users in this sample selected one or more of these four responses, which all suggest keeping its capabilities together in one product.

Product development, offering training & certification, better documentation, and how OpenSearch drives business value, are key to driving product visibility and adoption

Improving the performance and scalability of OpenSearch was the leading response to how OpenSearch could increase product adoption, as shown in Figure 13. This is a continuation of the above theme of maintaining product excellence through continuing the current pace of R&D investment. Improving visibility into OpenSearch capabilities and business value, providing more comprehensive documentation, and providing training and certification were also all leading activities to drive OpenSearch adoption.

When asked how OpenSearch could increase its visibility, the leading response from users in Figure 14 was to provide OpenSearch training and certification courses, increase involvement at open source events, and improve visibility into OpenSearch capabilities and business value.

Maintaining the pace of product investment, providing training and certification, better documentation, and an increased focus on how OpenSearch drives business value are important recurring themes that are critical to ensuring OpenSearch's continuing success.

OpenSearch should transition to oversight by a vendor-neutral foundation

Another objective of this survey was to provide an opportunity for users to weigh in on the future of OpenSearch. One early indication came in Figure 11, where 84% of users reported that it was either extremely important or important that whoever had oversight on OpenSearch should be an open source champion. Being an open source champion requires a deep commitment to the open source ethos and actively working to advance open source through advocacy, contribution, and community involvement.

Figure 12 then asked how important it was for OpenSearch to join a vendor-neutral foundation. User response showed that 81% believed it was either extremely important or important to transition OpenSearch to a vendor-neutral foundation. Figure 13 and Figure 14 confirmed this finding by ranking having OpenSearch join a vendor-neutral foundation as a leading way to drive both visibility and adoption.

Methodology

About the survey

A web survey that Linux Foundation Research and its partners conducted from February to March 2024 provided the basis for this study. The survey's goal was to understand the awareness, perception, perspective, and product experience of respondents with search, observability, and security analytics experience. In this section, we present the study methodology and context regarding how we analyzed the data followed by the demographics of the respondents.

From a research perspective, it was important to eliminate any perception of sample bias and ensure high data quality. We handled the elimination of sample bias by sourcing our usable sample from Linux Foundation subscribers, members, partner communities, and social media. We addressed data quality through extensive prescreening, survey screening questions, and data quality checks to ensure that respondents had sufficient professional experience to answer questions accurately on behalf of the organization they worked for.

We collected survey data from industry-specific companies, IT vendors and service providers, and nonprofit, academic, and government organizations. Respondents spanned many vertical industries and companies of all sizes, and we collected data from several geographies, including the Americas, Europe, and Asia-Pacific.

The 2024 Search and Analytics Survey comprised 43 questions that addressed screening, respondent demographics, OpenSearch user perceptions, and respondent product experience. For information about access to the 2024 Search and Analytics Survey, its dataset, and survey frequencies, see the Data. World access information below.

The high-level design of the survey is as follows:

TABLE 1: SURVEY DESIGN

Pages	Questions	Question categories	Who answers the questions
P1	-	Introduction	All respondents
P2	Q1 - Q6	Tell us about yourself	All respondents (N=689)
P3 - P4	Q7 - Q8	Involvement in IT	Students only (N=27)
P4	Q9 - Q15	Tell us about the organization that you work for	All respondents except students (N=689)
P5 - P6	Q16 - Q22	Perceptions of OpenSearch	Respondents with OpenSearch experience (N=176)
P7 - P8	Q23 - Q27	Feedback to OpenSearch	Respondents with OpenSearch experience (N=176)
P9 - P10	Q28 - Q30	Perspectives on search tools	Respondents with search experience (N=555)
P11 - P12	Q31 - Q33	Perspectives on observability tools	Respondents with observability experience (N=450)
P13 - P14	Q34 - Q36	Perspectives on security analytics tools	Respondents with security analytics experience (N=241)
P15	Q37 - Q39	Closing questions	All respondents (N=689)
P16	Q40 - Q43	Interest in technology, training, and certification	Students only (N=2)
P17	Q44 - Q45	LFR Panel and reward information	All respondents

Survey screening involved the use of five variables to validate the respondent:

- Must be a Linux Foundation subscriber, member, or contributor
- Must be familiar with the Linux Foundation (to ensure understanding of or involvement with open source)
- Must work full time or part time in the information technology field
- Must be in an IT role aligned with development, deployment, operations, or support
- Must have familiarity with one or more areas or tools that address search, observability, or security analytics

A total of 689 respondents completed the survey. The margin of error for this sample size was $\pm 3.2\%$ at a 90% confidence level and $\pm 3.8\%$ at a 95% confidence level.

We stratified the data collection by company size, geographic region, and organization type. We segmented the data primarily by geographic region, company size, and type of organization.

Although there was a requirement for respondents to answer nearly all questions in the survey, we made a provision when a respondent was unable to answer a question by adding a “Don’t know or not sure” (DKNS) response to the list of responses for every question. However, this created a variety of analytical challenges.

One approach was to treat a DKNS just like any other response to know the percentage of respondents that answered DKNS. The advantage of this approach is that it shows the exact distribution of data collected. The challenge with this approach is that it can distort the

distribution of valid responses, i.e., responses where respondents could answer the question.

Some of the analyses in this report exclude DKNS responses. This is because we can classify the missing data as either missing at random or missing completely at random. Excluding DKNS data from a question does not change the distribution of data (counts) for the other responses, but it does change the size of the denominator used to calculate the percent of responses across the remaining responses. This has the effect of proportionally increasing the percentage values of the remaining responses. Where we have elected to exclude DKNS data, the footnote for the figure includes the phrase “DKNS responses excluded.”

The percentage values in this report may not total to exactly 100% due to rounding.

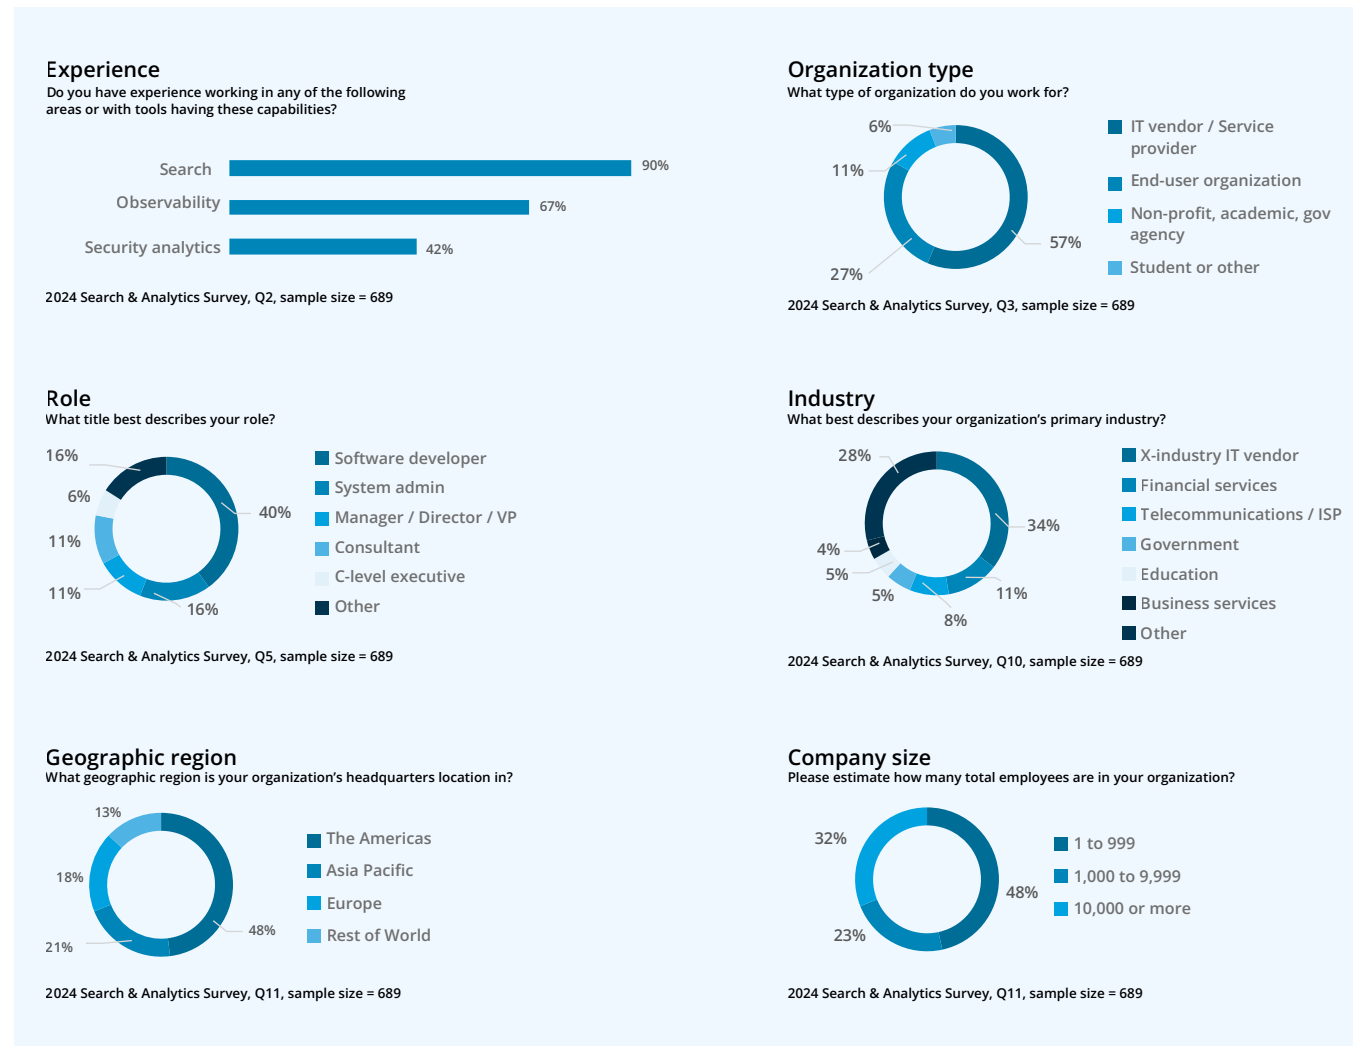
Data.World access

Linux Foundation Research makes each of its empirical project datasets available on Data.World. This dataset includes the survey instrument, raw survey data, screening and filtering criteria, and frequency charts for each question in the survey. You can find Linux Foundation Research datasets, including this project, at data.world/thelinuxfoundation. Access to Linux Foundation datasets is free but does require you to create a data.world account.

Respondent demographics

These demographics provide you with a profile of the 2024 Search and Analytics Survey respondents. We have regrouped all the demographics in Figure 17 to facilitate a more insightful analysis. For the original source data and study frequencies, please see the data.world access described above.

FIGURE 23: RESPONDENT DEMOGRAPHICS




About the authors

STEPHEN HENDRICK is vice president of research at the Linux Foundation, where he is the principal investigator on a variety of research projects core to the Linux Foundation's understanding of how open source software is an engine of innovation for producers and consumers of IT. Steve specializes in primary research techniques developed over 30 years as a software industry analyst. Steve is a subject-matter expert in application development and deployment topics, including DevOps, application management, and decision analytics. Steve brings experience in a variety of quantitative and qualitative research techniques that enable deep insight into market dynamics and has pioneered research across many application development and deployment domains. Steve has authored over 1,000 publications and provided market guidance through syndicated research and custom consulting to the world's leading software vendors and high-profile start-ups.

ADRIENN LAWSON is a data analyst at the Linux Foundation. Adrienn obtained a master's degree in social data science from the University of Oxford. She is responsible for survey development, analysis, and report writing. Adrienn has previously conducted research at the University of Oxford, the Budapest Institute for Policy Analysis, and the U.K.'s Office for National Statistics. She has a strong fascination with the collective power of open source collaboration within geographically dispersed communities. Additionally, she is most interested in researching trends and solutions for challenges related to open source software funding, sustainability, and supporting developers in their pursuit of responsible technological advancement.





Acknowledgements

We thank all the participants of the survey for kindly sharing their insights and experience in search and analytics. Special thanks to peer reviewers and Linux Foundation colleagues for their involvement in the various stages of the research process, including:

- Joshua Bright (AWS)
- Elizabeth Bushard
- Hilary Carter
- Mia Chaszeyka
- Eli Fish (AWS)
- Charlotte Henkle (AWS)
- Anna Hermansen
- Mukul Karnik (AWS)
- James McIntyre (AWS)
- Carl Meadows (AWS)
- Christina Oliviero
- Aparna Sundar (OpenSearch project)
- Daryll Swager (AWS)




About OpenSearch


OpenSearch is a community-driven, open source search and analytics suite used by developers to ingest, search, visualize, and analyze data. OpenSearch consists of a data store and search engine (OpenSearch), a visualization and user interface (OpenSearch Dashboards), and a server-side data collector (Data Prepper). Users can extend the functionality of OpenSearch with a selection of plugins that enhance search, analytics, observability, security, machine learning, and more.





About Linux Foundation Research

Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.

 x.com/linuxfoundation

 youtube.com/user/TheLinuxFoundation

 facebook.com/TheLinuxFoundation

 GITHUB.COM/LF-ENGINEERING

 linkedin.com/company/the-linux-foundation



Copyright © 2024 **The Linux Foundation**

This report is licensed under the [Creative Commons Attribution-NoDerivatives 4.0 International Public License](https://creativecommons.org/licenses/by-nc/4.0/).

To reference this work, please cite as follows: Stephen Hendrick and Adrienn Lawson, "OpenSearch Awareness and Perceptions Report," foreword by Jonah Kowall, The Linux Foundation, October 2024.

