



Research



The 2026 Open Data Infrastructure Report

AI, Governance, and OpenSearch
Adoption Trends

September 2026

Csilla Zsigri, *Research Analyst & Strategist*
Adrienn Lawson, *The Linux Foundation*
Hilary Carter, *The Linux Foundation*

Foreword by Bianca Lewis,
OpenSearch Software Foundation



The 2026 Open Data Infrastructure Report

44% of organizations prefer **one core platform** with specialized tools, and 80% would consolidate to fewer platforms if it meaningfully reduced costs.



Gen AI and LLM-powered applications are the clear leader in AI use cases relevant to data infrastructure, cited by **82%** of organizations.



78% of organizations say **cross-vendor, cross-tool interoperability is critical** when evaluating data infrastructure.



Total cost of ownership (80%) and security and compliance (79%) top the list of what organizations look for in a data infrastructure platform.



Hybrid search (combining keyword and semantic search) is now the preferred approach for **68%** of organizations.




Top data infrastructure risk concerns are security vulnerabilities (58%), regulatory compliance and data exposure across multiple tools (33% each), and vendor lock-in (32%).



Running or planning to run **AI workloads** is standard practice for at least **83%** of organizations across every region, reaching over **90%** among large enterprises.



71% of organizations agree that **deploying infrastructure independently** of a specific vendor or cloud provider is a strategic priority.



Awareness of OpenSearch jumped from 68% in 2024 to 89% in 2026.



Top OpenSearch use cases are search and retrieval (91%), log analytics and observability (83%), and real-time analytics and operational intelligence (70%).



77% of organizations see OpenSearch as **relevant to AI**, with **44%** viewing it as **core AI infrastructure**. **62%** already use it in AI workloads.



OpenSearch **production adoption** nearly doubled between 2024 and 2026, from 19% to **36%**, with another 36% testing or planning to adopt.



Contents

Foreword	4
Executive summary	5
Data infrastructure priorities	6
AI momentum	12
OpenSearch perception	14
OpenSearch momentum	17
Community and foundation	22
Conclusions	24
Methodology	25
Appendix	29

Foreword

When I joined the OpenSearch Software Foundation as its first Executive Director just under a year ago, I came in from the vendor side of this ecosystem. Having spent years building businesses around search and analytics technology, I thought I understood what OpenSearch was. I did not, and could not, understand what it was becoming.

Five years ago, OpenSearch began as a fork. AWS built it in 2021 because the community needed an open alternative to Elasticsearch and Kibana when the license went proprietary, and needed one it could trust would stay open. Born out of necessity. In 2024, that necessity became something sturdier: the project moved under open governance at the Linux Foundation, and the newly formed OpenSearch Software Foundation became its steward, vendor neutral by charter and not just by good intentions. It's the moment this project stopped being a fork and started being its own data infrastructure platform for the entire community.

This fifth-anniversary research report is the first time I have seen that shift captured so clearly in hard data, and I will admit I did not expect the numbers to move this fast. Awareness of OpenSearch jumped from 68% to 89% of organizations in just two years. Production adoption has nearly doubled, from 19% to 36%. Another 36% are testing it or planning to adopt it right now, which means the number of organizations actively evaluating OpenSearch today is about as large as the number already running it in production. That's not a project quietly persisting, it's a project breaking through.

What makes me most excited, though, is where this is happening and why. This isn't adoption concentrated in one region, one industry, or one company size. We're seeing adoption across North America, Europe, and Asia-Pacific, across small businesses and global enterprises, across finance, retail, telecom, the public sector and more. And increasingly, the reason organizations are reaching for OpenSearch is AI. With vector search baked in, OpenSearch was part of the GenAI story from day one; today 83% of organizations are running or planning to run AI workloads, and among current OpenSearch teams, 62%

are already using it to power those AI pipelines. 44% now consider it core AI infrastructure, not a bolt-on. When I talk to practitioners at meetups and at OpenSearchCons, this is the part they light up about: OpenSearch sitting underneath the retrieval, the vector search, the observability that ensures their AI applications actually work in production, all without locking them into a single vendor's roadmap.

I want to be honest about what the data also tells us, because that is the kind of foreword I want to write. The data shows our biggest growth vector isn't convincing people to try OpenSearch, it's helping the 60% currently running experimental or replaceable setups turn them into mission-critical infrastructure. Most organizations who could join the Foundation have not yet. The barriers are not people rejecting OpenSearch; they are awareness and expertise gaps we can close. I read that as the clearest to-do list I have ever been handed, and I intend to spend the next year working through it alongside this community.

Five years from now, I want a member from our community to contribute a foreword to our research that says: A look back at a decade of dependable AI infrastructure, not just adoption but trust earned the hard way. Based on what is in this report, I believe we are exactly on that path. Whether you are running OpenSearch in production, evaluating it for your AI stack, or considering Foundation membership, I invite you to dig into these numbers and build the next five years with us.

Thank you to everyone who filled out this survey, sat for an interview, filed an issue, wrote documentation, or showed up to a community meeting this year. This report is really about you.

Bianca Lewis
Executive Director,
OpenSearch Software Foundation

Executive summary

OpenSearch was created in 2021, led by AWS, as a fork of Elasticsearch and Kibana. In 2024, OpenSearch transitioned to open governance under the Linux Foundation with the establishment of the OpenSearch Software Foundation, formalizing its status as a vendor-neutral, community-governed project. This report marks OpenSearch's fifth anniversary.

The 2026 Open Data Infrastructure Report presents findings from Linux Foundation Research's 2026 Data Infrastructure Survey, examining how organizations are approaching data infrastructure and where OpenSearch specifically stands within that landscape today. The findings are based on 294 survey respondents spanning IT vendors, end-user organizations, and independent consultants across multiple industries, company sizes, and geographies. In addition, 11 subject matter expert interviews were conducted in July and August, 2026, with volunteers sourced from the survey, and from an open call for participants by LF Research and OpenSearch leadership.

Perhaps the most significant finding is that OpenSearch has broken through. Awareness jumped from 68% in 2024 to 89% in 2026, and production adoption nearly doubled, from 19% to 36%. As another 36% of organizations are already testing OpenSearch or planning to adopt it, the number of organizations evaluating OpenSearch today is roughly as large as the number already running it in production.

This adoption curve is unfolding alongside a parallel shift toward AI. AI workloads have moved past early adoption, and are now running or planned, on average, at 83% or more of organizations, and depending on geography and company size, this figure could reach as high as 92%. OpenSearch is already delivering on core

use cases like search/retrieval and analytics/observability, and it's now being used for AI workloads at meaningful scale as well (62% of organizations). This coincides with the market's own buying criteria of cost, security, and vendor independence, which closely match OpenSearch's current capabilities. Perception data supports this alignment: OpenSearch is generally rated positively both on enterprise readiness and vendor neutrality, and 44% of organizations see it as core AI infrastructure.

Still, the clearest opportunity going forward lies in depth of use and not initial reach. 60% of current deployments remain experimental or production-but-replaceable rather than deeply embedded. This is a large, active population mid-trial with room to mature into committed, mission-critical usage. More maturity looks especially achievable given that the barriers holding it back are practical and solvable, rooted in awareness and expertise gaps rather than fundamental objections.

A similar pattern holds one level up, for the OpenSearch Software Foundation itself: membership is well positioned to expand, and the factors (ROI, enterprise readiness, and vendor neutrality) that would move prospective members, match the criteria driving platform adoption. Open source communities and developer forums already reach enterprise users at scale, representing existing channels through which that visibility could translate into membership growth.

Put simply, OpenSearch has established strong awareness and gained traction among the decision-maker and practitioner segments most relevant to its growth. The data points to a clear next phase: turning trial into production, production into mission-critical, broader use of both core and AI capabilities, and converting awareness into Foundation membership.

Data infrastructure priorities

Data infrastructure, meaning the search, analytics, observability, and related platforms and tools organizations rely on to collect, store, manage, process, and extract value from their data, has become a strategic priority. As data volumes grow, data types diversify, and use cases evolve and multiply, organizations are increasingly weighing not just what a platform can do, but how many tasks it can do at once, what it costs to run and maintain, and how much control they retain over their own data and vendor relationships, among other things.

Consolidated, multi-purpose platforms beat point solutions

Modern data infrastructure is needed for many tasks at once. As shown in Figure 1, use cases are broadly and evenly relevant across the survey sample which means that organizations are choosing platforms that can flex across multiple purposes rather than serve one task. AI-related use cases are among the fastest-growing demands on data infrastructure today.

FIGURE 1

Use cases relevant organizations

2026 Data Infrastructure Survey, Q10, N=294, 'All of the above' choices incorporated. For a complete dataset see Appendix A1.



Our research suggests that, overall, the platform story is more compelling than best-of-breed: 44% of survey respondents prefer one core platform with specialized tools, versus 26% who prefer picking the best individual tool per use case (see Figure 2).

“Organizations often favor standardized tools that have been thoroughly proven by the community. There is a clear preference for adopting a single core platform capable of addressing multiple use cases rather than managing highly specialized, individual tools.”

— ALFREDO SOTO NOLASCO, FULL STACK DEVELOPER, INDRA GROUP

That said, the best-of-breed segment is still sizable, with 26% of organizations selecting it as the preferred option. The split also varies by company size, with larger organizations favoring integrated platforms, while smaller ones show more varied tooling strategies (see complete dataset in Appendix A2.2).

Figure 3 demonstrates that the consolidated platform preference holds up under direct questioning: 80% of organizations said they would consolidate to fewer platforms if that meaningfully reduced costs, and 75% agreed that managing multiple specialized systems adds too much operational overhead.

But while cost is driving consolidation, interoperability remains non-negotiable. 78% of organizations say that interoperability across vendors and tools is critical when evaluating data infrastructure (see Figure 3).

“As network architectures increasingly disaggregate software from the underlying hardware, it is important that platforms—including technologies such as OpenSearch and SDN controllers—maintain open APIs and standards-based interoperability. This enables organizations to avoid unnecessary vendor dependence and preserves the flexibility to evolve infrastructure over time.”

— ANURADHA (ANU) UDUNUWARA, SENIOR ENGINEER, SRI LANKA TELECOM; TECH LEAD, XYNTAC ENTERPRISE

Data infrastructure spend is material but often invisible to practitioners

Average annual spend across search, analytics, observability, and related data infrastructure is around \$2.4 million. As shown in Figure 4, material six- and seven-figure annual budgets are concentrated in mid-market enterprises and above (\$250M+ annual revenue).

29% of respondents do not know or are not sure what their organization spends on data infrastructure, which is consistent

FIGURE 2

Organizations' tooling preferences

2026 Data Infrastructure Survey, Q11, N=294, Select one. For a complete dataset see Appendix A2.1-A2.2.




FIGURE 3

Agreement with statements on consolidation and interoperability

2026 Data Infrastructure Survey, Q15, N=294. For a complete dataset see Appendix A3.





FIGURE 4

Approximate annual spend across data infrastructure platforms

2026 Data Infrastructure Survey, Q22, N=294. For a complete dataset see Appendix A4.



with the pattern of decision-makers typically knowing the numbers while practitioners often do not.

“While organizations were previously less sensitive to running high cost compute jobs, leadership is increasingly clamping down on model costs and restricting access. As licensing and operational expenses accumulate, teams are becoming far more cost conscious when making technology decisions.”

— ACCORDING TO A MANAGER FROM A BIG FOUR CONSULTING FIRM

Cost, security, and vendor independence anchor buying decisions

Figure 5 illustrates that two criteria dominate platform selection: total cost of ownership (TCO) (very important to 80% of organizations) and security and compliance (very important to 79%). Vendor independence and deployment flexibility (69%), and enterprise readiness (67%) follow closely behind.


Consistent with the vendor independence trend, 71% of organizations agree that deploying infrastructure independently of a specific vendor or cloud provider is a strategic priority (Figure 3).

Security also shows up as the top data infrastructure risk concern. As depicted in Figure 6, security vulnerabilities are the strongest risk signal by a wide margin, ranking highest across all organization sizes and respondent roles. Regulatory compliance, data exposure, vendor lock-in, and operational complexity follow in a tight cluster. External threats and dependency risks dominate the risk register, while operational friction ranks lower.

FIGURE 5

Criteria rated “very important” when evaluating or selecting a data infrastructure platform

2026 Data Infrastructure Survey, Q12, N=294, Very important. For a complete dataset see Appendix A5.



“Data security is the absolute priority when evaluating data platforms, but it extends beyond underlying software bugs to how complex a tool is to configure properly. When a technology requires extensive effort to reach a secure state, maintainability becomes a major concern for corporate teams.”

— ACCORDING TO A MANAGER FROM A
BIG FOUR CONSULTING FIRM

Hybrid search is the default expectation

Search architectures have historically relied on keyword-based matching, ranking results by exact or approximate term overlap. As natural language interfaces and AI-powered applications have matured, semantic and vector search, retrieving results based on conceptual similarity rather than literal keyword matches, have emerged as a complementary approach particularly well-suited to use cases like retrieval-augmented generation. Hybrid search combines both methods, using keyword matching for precision and semantic/vector search for contextual relevance, and is increasingly viewed as the more complete solution rather than a tradeoff between the two.

Figure 7 shows that hybrid search is now the preferred approach for 68% of organizations, far outpacing keyword-only (13%) or semantic/vector-only (7%) preferences.

The 11% of do not know or not sure responses point to a continued education gap around what hybrid search actually means. Even so, hybrid search is a market-wide requirement, not a niche preference.

FIGURE 6

Biggest concerns around data infrastructure risk

2026 Data Infrastructure Survey, Q14, N=294. For a complete dataset see Appendix A6.





FIGURE 7

Preferred search approach

2026 Data Infrastructure Survey, Q13, N=294. For a complete dataset see Appendix A7.




“OpenSearch has emerged as a Swiss Army knife for solving many kinds of search problems using a well-known API, making it an appealing default choice for standard search needs. But teams facing complex query and ranking problems or certain high scale problems often require more bespoke or specialized solutions.”

— TREY GRAINGER, AI SEARCH CONSULTANT & FOUNDER, SEARCHKERNEL

FIGURE 8

Preferred vendor relationship and support model

2026 Data Infrastructure Survey, Q20, N=294; Q21, N=291. For a complete dataset see Appendix A8.1-A8.2



No single support model fits all

How organizations resource and support their data infrastructure is shaped by a mix of factors such as available in-house skills, risk tolerance, budget, and the maturity of the platforms themselves. Some organizations consolidate around a single trusted vendor to simplify accountability and reduce coordination overhead, while others deliberately work with multiple vendors or consultancies to avoid over-reliance on any one provider, or invest in building internal expertise to reduce external dependency altogether.

The same range of considerations extends to support models, where organizations weigh the assurances of commercial, vendor-backed agreements against the flexibility of community-driven support or the credibility of project-accredited providers.

That mix of considerations is reflected clearly in the survey data. On vendor relationships, 29% prefer working with multiple vendors or consultancies, 28% prefer consolidating around a single vendor, and 25% prefer building out in-house expertise, with the remainder expressing no clear preference or uncertainty (see Figure 8, first chart).

The same tradeoffs show up in support preferences. 37% prefer commercial, vendor-backed support; 23% prefer community-based support; and 21% prefer vendors accredited by the project (see Figure 8, second chart). Taken together, these results tell a consistent story, as illustrated in Figure 8: no single model fits all.

Looking more closely at who prefers what, three distinct support preference profiles emerge, as seen in Figure 9. VP/Director and large-enterprise segments lean toward commercial or multi-vendor support; mid-market organizations show a notable pull toward vendor-accredited support; and startups, alongside

practitioner teams, most often favor in-house expertise and community-based support.

Long-term support (LTS) commitments, the period during which a vendor or project guarantees security patches, bug fixes, and compatibility updates for a given release, have taken on new significance as regulatory frameworks increasingly hold organizations accountable for the software they run. The EU’s Cyber Resilience Act (CRA), in particular, requires manufacturers to provide security updates for a minimum of five years, or to match a product’s shorter expected lifetime. Additionally, manufacturers are required to clearly define, disclose, and document product support periods, placing pressure on organizations to understand and plan around the LTS commitments of the infrastructure they depend on.

When it comes to LTS windows, according to our 2026 survey, expectations vary widely, as shown in Figure 10. 24 and 36 months are the most common minimum durations expected, particularly among larger enterprises, while smaller organizations tend to favor shorter cycles. A full 20% say that they do not know or are not sure, suggesting that LTS policies and their implications are not yet well understood by the market.

“Managing software long term support requires balancing rapid turnaround cycles with legacy lifecycles. While the initial baseline is to design systems that can be updated within a three to five year window, operational realities mean engineering teams must ultimately remain prepared to maintain and support legacy setups that are decades old.”

— DARNELL OTTERSON, STAFF ENGINEER, DAIKIN APPLIED AMERICAS

FIGURE 9

Support preference profiles

2026 Data Infrastructure Survey, Q20, N=294; Q21, N=291. For a complete dataset see Appendix A9.1-A9.4.





FIGURE 10

Minimum LTS duration expected for production-critical platforms

2026 Data Infrastructure Survey, Q23, N=294. For a complete dataset see Appendix A10.




AI momentum

Artificial intelligence has moved from experimentation to infrastructure. As organizations build generative AI applications, retrieval-augmented generation pipelines, and AI-assisted analytics and observability, the underlying data platforms that search, index, retrieve, and monitor information have become just as central to AI success as the models and algorithms themselves. Search and analytics platforms play a foundational role in this shift, providing the retrieval layer that grounds AI applications in an organization's own data and the observability layer that keeps AI-powered systems running reliably in production.

FIGURE 11

Organizations running or planning to run AI workloads, by geography and company size

2026 Data Infrastructure Survey, Q16, N=294, Yes. For a complete dataset see Appendix A11.1-A11.2.



Understanding how widely AI workloads have spread, and what organizations need from their data infrastructure to support them, is essential context for evaluating OpenSearch's role and position in the market.

AI workloads are near-universal, and even more so at scale

Figure 11 shows that running or planning to run AI workloads is close to universal across regions, reaching 89% of organizations in the US/Canada, 86% in Europe, and 83% in Asia-Pacific. Propensity toward AI workloads scales with company size, too: 73% of small businesses are running or planning to run AI workloads, versus 89% of mid-size and mid-market enterprises and 92% of large and global enterprises. AI is no longer an early-adopter phenomenon. In any segment at the enterprise level it is standard practice.


“As enterprises adopt frontier AI models, control over data processing, storage, reuse, and governance is becoming a board level priority. Open source platforms such as OpenSearch enable organizations to build secure search, document storage, and vector-based RAG capabilities within enterprise boundaries. Features such as mature RBAC, audit logging, and snapshot management provide stronger safeguards for sensitive data while maintaining the flexibility needed to drive innovation.”

— DAN SITHAMBARAM, GLOBAL PRACTICE LEAD, REDINGTON GROUP

FIGURE 12

Leading AI use cases

2026 Data Infrastructure Survey, Q17, N=245. For a complete dataset see Appendix A12.



Gen AI/LLM-powered applications lead AI use cases

Gen AI and LLM-powered applications are the clear leader at 82%, ahead of AI-assisted operational analytics and observability (56%) and enterprise knowledge and business intelligence assistants or copilots (55%), as illustrated in Figure 12.

AI draws on a broad footprint of data, and single-platform access matters


Figure 13 shows that organizations power AI applications with a wide range of data: log, metrics, and operational data (64%), business and transactional data (59%), security and threat intelligence data (55%), internal knowledge documentation (55%), and real-time event data (53%) all see substantial current or planned use.

Against that breadth, 84% of organizations say that single-platform access to their data matters for AI (see the complete dataset in Appendix A13.2), reinforcing the same consolidation preference seen in the 'Data infrastructure priorities' section, now sharpened by AI delivery requirements.

FIGURE 13

Data used or planned to be used to power AI applications

2026 Data Infrastructure Survey, Q18, N=245. For a complete dataset see Appendix A13.1.



OpenSearch perception

OpenSearch is an open source search, analytics, and observability platform, originally forked from Elasticsearch (7.10.2) and Kibana (7.10.2) in April 2021, and now governed under the Linux Foundation via the OpenSearch Software Foundation (since September 2024). In 2024, Linux Foundation Research published a baseline study on OpenSearch awareness and perceptions. How have things changed since?

Awareness is breaking through

In 2024, 68% of organizations had at least heard of OpenSearch; by 2026, that figure reaches 89% (see comparisons in Figure 20). In this year's survey, 30% of respondents report being very familiar, 37% somewhat familiar, and 22% aware of OpenSearch but not very familiar (Figure 14).

Familiarity is particularly high among decision-makers, which may reflect strategic or market awareness more than hands-on technical experience. It's also higher among larger enterprises (mid-size and up) than smaller organizations (startups and small businesses). More detailed breakdowns can be found in Appendix A14.2-A14.3.


Practitioners describe OpenSearch as an open source search, analytics, and observability platform

As Figure 15 shows, when asked to describe OpenSearch in their own words, respondents most often reach for "search and retrieval" (43 mentions, the top theme) followed closely by "analytics, observability, and monitoring" (36 mentions). The two are frequently paired, reflecting how practitioner teams actually use the product day to day.

FIGURE 14

Familiarity with OpenSearch

2026 Data Infrastructure Survey, Q25, N=294. For a complete dataset see Appendix A14.1-A14.4.



“Do you need a search option that is also working for observability and MLOps? There is OpenSearch.”

— VP / DIRECTOR ENGINEERING, INFRASTRUCTURE, OR PLATFORM

Open source identity and community-driven governance are also prominently mentioned (31 mentions), closely coupled with both the Elasticsearch comparison and the vendor-independence angle. A small, but notable minority mention AI capabilities (10 mentions). These responses tend to be more technically sophisticated, describing OpenSearch's emerging identity as an AI data platform.

FIGURE 15

Top themes in how respondents describe OpenSearch

2026 Data Infrastructure Survey, Q26, N=105, Open-ended, Optional, Thematically categorized and ranked according to number of mentions. For a complete dataset see Appendix A15.

1. Search & retrieval (43 MENTIONS ▶ 47%)

Respondents describe OpenSearch primarily as a search engine — for data, logs, documents, or applications. Ranges from highly technical to intuitive descriptions.

2. Analytics, observability & monitoring (36 MENTIONS ▶ 39%)

Often paired with search. Reflects the core practitioner use case. This is how DevOps/SRE and platform engineers actually use the product day-to-day.

3. Open source & community (31 MENTIONS ▶ 34%)

Open source identity and community-driven governance are prominently mentioned. This theme is closely coupled with both the Elasticsearch comparison and the vendor independence angle.

4. Elasticsearch fork / alternative (26 MENTIONS ▶ 28%)

OpenSearch is commonly introduced by reference to Elasticsearch, as a fork, alternative, replacement, or ELK-stack successor. Tone is mostly neutral. A subset explicitly frames the fork positively and a minority is cynical about AWS's role.

5. AI, ML & semantic search (10 MENTIONS ▶ 11%)

A small but notable minority mention AI capabilities. These responses tend to be more technically sophisticated and describe OpenSearch's emerging identity as an AI data platform.

6. Scalability, flexibility & performance (10 MENTIONS ▶ 11%)

Scalability, flexibility, distributed architecture, and performance are mentioned as distinguishing attributes. Usually paired with another theme rather than standing alone.

7. Positive sentiment / community advocacy (8 MENTIONS ▶ 9%)

A cluster of responses that are enthusiastic but light on specifics. These represent genuine brand warmth and community affiliation rather than technical evaluation.

8. Security analytics (6 MENTIONS ▶ 7%)

The six respondents who mention security do so either as part of a broader multi-use-case description or as the primary lens. It appears most often in technically sophisticated responses from practitioners.

9. Data infrastructure / platform (6 MENTIONS ▶ 7%)

This theme captures descriptions at a higher level of abstraction. These tend to be shorter, more strategic framings, often from executives or less technical respondents.

10. Vendor neutrality & independence (4 MENTIONS ▶ 4%)

Few respondents explicitly mention vendor independence as part of their description, but when they do, the sentiment is strong.

Perceived as enterprise-ready and vendor-neutral, with room for improvement

As illustrated in Figure 16, on a 1-to-5 scale, respondents rate OpenSearch's readiness for mission-critical enterprise workloads at an average of 3.70, and its neutrality/vendor-independence today at 3.66 — both above the midpoint, indicating a generally positive but not unreserved baseline perception.

FIGURE 16

Average ratings of vendor neutrality and enterprise readiness

2026 Data Infrastructure Survey, Q27, Q28, N=250. For a complete dataset see Appendix A16.




Figure 17 shows that perception is closely aligned with commitment: 60% of current OpenSearch users describe their own deployment as either experimental or replaceable, versus 36% who call it deeply embedded or mission-critical.

FIGURE 17

How organization describe OpenSearch deployments

2026 Data Infrastructure Survey, Q38, N=128; Q41, N=73; Q32, N=185. For a complete dataset see Appendix A21, A26, A30.



On the non-adopter side, 24% cite vendor-neutrality concerns tied to OpenSearch's association with AWS as a reason for staying out, and 40% of prospective Foundation members cite greater vendor-neutrality confidence as a prerequisite for joining.

Relevant to AI, with room to grow

77% of organizations see OpenSearch as relevant to AI overall, with 44% seeing it as core AI infrastructure and 33% as a supporting component (Figure 18). Greater familiarity correlates strongly with viewing the project as core to AI, and active users are significantly more likely than non-users to see it that way (see data in Appendix A18.2). This suggests that perception of OpenSearch as an AI platform builds with hands-on exposure rather than preceding it.


“AI and machine learning teams coming out of data science programs usually prefer writing in Python and rarely reach for OpenSearch first. They often build initial solutions on newer tools, but once those pilots hit production scale and run into limitations, they take a second bite at the apple and rebuild their architecture on OpenSearch.”

— ERIC PUGH, CO-FOUNDER, OPENSOURCE CONNECTIONS & MEMBER, OPENSEARCH TECHNICAL STEERING COMMITTEE

FIGURE 18

Perceived relevance of OpenSearch to AI

2026 Data Infrastructure Survey, Q29, N=250. For a complete dataset see Appendix A18.1-A18.2.



view OpenSearch as core infrastructure for AI applications



see OpenSearch as a supporting component for AI applications

OpenSearch momentum

Awareness and perception of a technology do not necessarily translate into how it is actually used. Adoption depends on a separate set of factors, including how deployments mature over time, which use cases a platform delivers on in practice, how organizations operationally deploy and govern it, and what still holds non-adopters back. These findings trace OpenSearch’s momentum along each of these dimensions.

From awareness to production

Figure 19 shows that 36% of organizations are actively using OpenSearch in production today, and another 36% are in the pipeline (18% in test/dev and 18% planning to adopt), where the opportunity to convert is essentially as large as the current install base. Production adoption is highest among large organizations, especially in North America, while mid-sized enterprises in Europe and Asia Pacific show significant test/dev usage, pointing to active expansion momentum.

The awareness-to-adoption progression is possibly the most powerful narrative to emerge from comparing this year’s survey data to that collected two years ago (Figure 20). The OpenSearch Software Foundation was established in September 2024; the study establishing the 2024 baseline was conducted earlier that same year, ahead of the OpenSearch Software Foundation’s formation as a neutrally governed Linux Foundation project.

In 2024, 68% of organizations had at least heard of OpenSearch; by 2026, that figure reaches 89%. Production experience tells an even sharper story: it nearly doubled, from 19% to 36%, while trial/experimentation grew from 11% to 18% (Figure 20).

“Our use of data infrastructure technology dates back over ten years to early versions of our primary data platform. When licensing changes disrupted that ecosystem, we transitioned first to an intermediate distribution and eventually moved directly to OpenSearch. Today, we use the technology very heavily across multiple major areas, including global anti money laundering, public security, and anti-fraud solutions.”

— TERRY QUIGLEY, PRINCIPAL SOFTWARE DEVELOPER, SAS




FIGURE 20

Awareness and adoption 2026 vs. 2024

2026 Data Infrastructure Survey (N=294), 2024 Search and Analytics Survey (N=689). For a complete dataset see Appendix A20, and also A14.1, A14.2, A19.





FIGURE 21

OpenSearch deployment depth

2026 Data Infrastructure Survey, Q38, N=128. For a complete dataset see Appendix A21.1-A21.3.



Deployed broadly but not yet deeply entrenched

Product stickiness reflects how deeply integrated a product has become in users' daily routines, and by that measure, OpenSearch shows broad deployment but limited entrenchment to date. Among current users, 60% describe their OpenSearch deployment as experimental or production-but-replaceable, 40% specifically production-but-replaceable against 36% who call it deeply embedded or mission-critical (Figure 21).

Deployment depth varies by geography: North America is the most entrenched market, Europe has strong adoption but is largely stuck at the "replaceable" stage, and Asia-Pacific is active but still mostly exploratory (see Appendix A21.2).

"OpenSearch serves as a mission critical system at Uber because the friendly open source licensing and active community make it a reliable foundation."

— MINGMIN CHEN, DIRECTOR OF SOFTWARE ENGINEERING,
UBER & MEMBER, OPENSEARCH GOVERNING BOARD

Deployment is self-managed, cloud-first, and centrally governed

Operationally, 60% of organizations run OpenSearch self-managed in the cloud, 48% via the managed Amazon OpenSearch Service, 9% via a managed service from another vendor, and 43% self-managed on-premises (see Figure 22). Isolating organizations that rely on a single approach exclusively, 47% deploy in self-managed mode only, 38% mix managed services with self-managing, and 14% rely on managed services only.

Geography and company size sharpen the picture further (see Table 1 and Table 2). The EU market has the most balanced profile, reflecting data-sovereignty considerations; Asia-Pacific has a clear multi-modal story; and the US sits between the two on every dimension as it's more cloud-forward than the EU, less multi-modal than Asia-Pacific, and has AWS as the dominant managed option.

TABLE 1
OpenSearch deployment by geography

2026 Data Infrastructure Survey, Q37, N=128, Geographies

OPENSEARCH DEPLOYMENT	US	EU	ASIA-PACIFIC
Self-managed on-premises	35%	46%	81%
Self-managed in the cloud	69%	46%	68%
Managed service via Amazon OpenSearch Service	57%	29%	63%
Manager service from a vendor other than AWS	21%	11%	17%

By company size, self-managed on-premises deployment dominates for startups and small businesses, while the adoption inflection point for AWS managed services sits at mid-market.

TABLE 2
OpenSearch deployment by company size

2026 Data Infrastructure Survey, Q37, N=128, Company Size

OPENSEARCH DEPLOYMENT	START-UP	SMALL	MID-SIZE	MID-MARKET	LARGE	GLOBAL
Self-managed on-premises	67%	78%	45%	40%	36%	41%
Self-managed in the cloud	58%	44%	65%	62%	66%	85%
Managed service via Amazon OpenSearch Service	42%	22%	42%	67%	54%	67%
Manager service from a vendor other than AWS	0%	11%	32%	11%	23%	19%

FIGURE 22
Types of OpenSearch deployment in use

2026 Data Infrastructure Survey, Q37, N=128, DKNS excluded. For a complete dataset see Appendix A22.





FIGURE 23
OpenSearch deployment evolution (2026 vs. 2024)

2026 Data Infrastructure Survey, Q37, N=126; 2024 Search and Analytics Survey, Q17, N=174; DKNS responses excluded. For a complete dataset see Appendix A23.



Delivering across core and AI use cases

Figure 24 depicts that among active users, OpenSearch is used broadly: search and retrieval (91%) and log analytics and observability (83%) lead, reflecting its origins and core strengths. Real-time analytics and operational intelligence (70%), AI-related use cases (62%), and business intelligence and reporting (62%) follow. The elevated “do not know or not sure” responses for AI and BI suggest room to grow rather than a capability gap, therefore the barrier to deeper adoption looks like activation, not product quality (see Figure 25).

“We heavily use the core search and security features because we sell internationally to major banks and governments that require strict regulatory compliance. This need is what led us to directly contribute to the project, where we have contributed developer time helping to deliver FIPS compliance within the core and the security plugin.”

— TERRY QUIGLEY, PRINCIPAL SOFTWARE DEVELOPER, SAS

“Handling highly sensitive data under a consent framework requires extreme customizability to maintain strict privacy promises. While commercial solutions often violate those privacy bounds, OpenSearch provides a customizable framework that allows organizations to maintain aggregate data layers while upholding user trust.”

— ART ABAL, MANAGING DIRECTOR, VANA FOUNDATION

FIGURE 24

Where OpenSearch is delivering for active users

2026 Data Infrastructure Survey, Q36, N=128. For a complete dataset see Appendix A24.




FIGURE 25

How well is OpenSearch delivering on capabilities

2026 Data Infrastructure Survey, Q40, N=128. Capability ratings are based on active users only, non-adopters are not represented. For a complete dataset see Appendix A25.





FIGURE 26

Primary reasons for not adopting OpenSearch

2026 Data Infrastructure Survey, Q41, N=71, DKNS responses excluded. For a complete dataset see Appendix A26.



Adoption barriers are practical, and non-users are open to evaluating

Figure 26 shows that the leading reasons for not adopting OpenSearch are limited awareness or familiarity (30%), insufficient internal expertise or operational resources (27%), and integration or migration complexity (27%); followed by having no compelling business case to migrate (25%) and vendor-neutrality/AWS-association concerns (24%).


“There is a significant amount of noise surrounding search and analytics offerings, which creates a real question about what OpenSearch actually provides and how it compares to centralized market alternatives. The biggest challenge is really an education issue, as many corporate clients and engineering teams simply lack clarity on where OpenSearch fits within a modern tech stack.”

— ART ABAL, MANAGING DIRECTOR, VANA FOUNDATION

FIGURE 27

Non-user organizations evaluating OpenSearch

2026 Data Infrastructure Survey, Q42, N=73. For a complete dataset see Appendix A27.



Consistent with that, 81% of non-users are open to evaluating OpenSearch, though only 14% are actively planning to. The largest group (44%) is open to evaluating OpenSearch but not treating it as a priority (Figure 27).

Community and foundation

Beyond perception, usage, and product capability, open source projects rely on the strength of their surrounding ecosystem: the channels through which practitioners discover and learn about a technology, and the governance structures, such as foundations, that provide long-term stewardship, funding, and neutral oversight. How well a project performs on these dimensions shapes not only its reach today, but its resilience and independence over time.

OSS and developer communities dominate discovery

As illustrated in Figure 28, open source communities and project documentation (58%) and developer communities and technical forums (45%) are in a tier of their own as information channels, reaching practitioners and decision-makers alike.

Analyst reports matter most for the C-suite and large enterprises; social media drives awareness for executives and small businesses but has limited reach among practitioners; and conferences are a consistent touchpoint across roles (see complete dataset in Appendix A28.1-A28.3).


“To drive deeper community engagement, small face-to-face meetups offer tremendous value for engineering teams. Focused sessions where engineers from two or three companies can whiteboard, present, and directly exchange technical experience builds strong connections.”

— MINGMIN CHEN, DIRECTOR OF SOFTWARE ENGINEERING, UBER & MEMBER, OPENSEARCH GOVERNING BOARD

FIGURE 28

Popular information channels

2026 Data Infrastructure Survey, Q30, N=240, DKNS responses excluded. For a complete dataset see Appendix A28.1-A28.3.



Open source communities and project documentation



Developer communities and technical forums

FIGURE 29

OpenSearch Software Foundation / Linux Foundation membership

2026 Data Infrastructure Survey, Q31, N=250. For a complete dataset see Appendix A29.





FIGURE 30

Joining incentives for Foundation membership

2026 Data Infrastructure Survey, Q32, N=185. For a complete dataset see Appendix A30.



Against that reach, only 6% of organizations in our survey are members of both the OpenSearch Software Foundation and the Linux Foundation, 9% belong to the Linux Foundation only, 67% are not members of either, and another 18% do not know or are not sure. Foundation membership has significant room to grow as we can see in Figure 29.

Business case, enterprise readiness, and vendor neutrality are the clearest value points


Three key value propositions stand out for prospective Foundation members: a clear business case/ROI (53%), enterprise readiness proof (47%), and vendor-neutrality confidence (40%). These are the same themes, if not the same numbers, that anchor buying decisions more broadly in the data infrastructure priority section (Figure 30).

Once organizations are weighing membership, priority security patches and vulnerability notifications (79% very important) and an 18-month LTS window (70%) rank highest among prospective benefits. Early access to new features trails well behind (39% very important), and is the only benefit where a meaningful share (17%) say it's not important. This reinforces that the Foundation's growth case should lead with risk reduction, support commitments, and proof points, not roadmap access alone (see Figure 31).

FIGURE 31

Importance of potential Foundation membership benefits

2026 Data Infrastructure Survey, Q33, N=185. For a complete dataset see Appendix A31.



“OpenSearch’s open governance model and project management structure serve as an accelerator for commercialization, providing the executive management and marketing framework needed to compete on even ground with proprietary or venture-capital-backed alternatives.”

— TREY GRAINGER, AI SEARCH CONSULTANT & FOUNDER, SEARCHKERNEL

Conclusions

OpenSearch has broken through since 2024, moving from awareness to real production adoption, and it performs strongly against the market's own buying criteria, in particular, cost, security, and vendor independence.

By 2026, OpenSearch is already delivering in its core use cases (e.g., search/retrieval and analytics/observability) and is being deployed for AI workloads at meaningful scale, right as AI itself becomes a near-universal workload across the market.

The clearest opportunity going forward lies in depth of use rather than initial reach. A large base of organizations are already using or actively trialing OpenSearch, and for many of them, the natural next step is deepening that usage, extending into core AI data infrastructure along the way.

This shift looks likely to continue: self-managed deployment has grown faster than Amazon OpenSearch Service usage since 2024, indicating a broader diversification of deployment strategy. The remaining barriers to adoption are practical and solvable (led by understanding and expertise gaps, not fundamental objections); and no single support model fits every segment, pointing to demand for a flexible, tiered offering rather than one default.

A similar pattern holds for the Foundation. Membership has significant room to grow, and the factors (e.g., ROI, enterprise readiness, and vendor neutrality) that would move prospective members match the same buying criteria that drive platform selection more broadly. Open source communities and developer forums are already the dominant channel for reaching this audience, decision-makers and practitioners alike, providing an existing channel through which that awareness could translate into membership growth.

Vendor independence ranks among the top criteria organizations use to evaluate data infrastructure, and vendor neutrality is one of the more commonly cited reasons non-adopters give for staying out. OpenSearch's shift to a neutral, vendor-independent foundation directly addresses a concern the market continues to flag as significant.

Methodology

This study is based on a combination of a web survey and in-depth interviews conducted by LF Research in the second quarter of 2026. The survey and interviews were designed to examine the state of OpenSearch adoption within the broader data infrastructure landscape, as well as its likely growth paths. A previous study, conducted in 2024, examined awareness and perceptions of OpenSearch and provides the baseline for the year-over-year insights in this report.

We addressed data quality through extensive prescreening, in-survey screening questions, and post-collection data quality checks, to ensure that respondents had sufficient professional experience to answer accurately, either on their own behalf or on behalf of the organization they represented.

Survey data was collected from end-user organizations across a range of industries, as well as IT vendors and service providers, independent consultants, and nonprofit, academic, and government organizations. Respondents spanned multiple vertical industries, company sizes, and geographies.

The survey comprised 44 questions covering screening, respondent demographics, data infrastructure priorities, AI infrastructure priorities, OpenSearch perception and adoption, and Foundation membership. It was fielded in May 2026. A total of 294 respondents participated in the web survey, and 11 took part in the in-depth, qualitative interviews.

About the respondents


A PRACTITIONER-HEAVY MIX, WITH DECISION-MAKERS WELL REPRESENTED

The survey captures both a strategic, buying-decision perspective and a genuine practitioner's-eye view of the day-to-day experience.

FIGURE 32

Respondents' roles: decision-makers vs. practitioners


2026 Data Infrastructure Survey, Q4, N=294



IT COMPANY EMPLOYEES MAKE UP THE MAJORITY OF RESPONDENTS

65% of respondents work at IT companies, with another 18% employed at non-IT companies — together accounting for the large majority of the sample. The remainder is spread across independent consultants, non-profit organizations, academic/research institutions, and government/public sector organizations, each representing a small share of the total.


FIGURE 33
Respondents' primary occupation, by organization type
2026 Data Infrastructure Survey, Q3, N=294



PARTICIPATION IS CONCENTRATED IN WESTERN MARKETS

Respondent organizations are headquartered predominantly in the US/Canada (46%) and Europe (30%), together accounting for over three-quarters of the sample. Asia-Pacific (13%) and the rest of the world (11%) make up the remainder. Findings in this report should be read as most representative of North America and Europe, with lighter, though still meaningful, coverage elsewhere.

FIGURE 34
Respondent organizations, by headquarters region
2026 Data Infrastructure Survey, Q7, N=294



RESPONDENT ORGANIZATIONS SPAN ALL SIZES, WITH AN ENTERPRISE LEAN

Respondent organizations range from early-stage startups to global enterprises, with the largest single tier being large enterprises at \$1B–\$10B in annual revenue (24%). Grouped into broader size bands, large organizations (\$1B+ in annual revenue) make up 38% of the sample, mid-size organizations 29%, and small organizations (under \$50M in annual revenue) 24%, with public sector, non-profit, and other organizations accounting for the remaining 9%.

Data access

The raw survey data underlying this report is publicly available on Zenodo via the Linux Foundation Research community repository at <https://zenodo.org/communities/linux-foundation-research>. The dataset includes all anonymized survey responses and can be freely accessed, downloaded, and reused under the terms of the applicable open license. Researchers and practitioners who wish to conduct independent analysis or build upon this work are encouraged to consult the repository directly.

FIGURE 35

Respondent organizations, by revenue size

2026 Data Infrastructure Survey, Q8, N=294





FIGURE 36

Respondent organizations, by revenue size tier (grouped)

2026 Data Infrastructure Survey, Q8, N=294



Appendix

A1: Use cases relevant to organizations

2026 Data Infrastructure Survey, Q10, N=294, 'all of the above' choices incorporated

WHICH OF THE FOLLOWING USE CASES ARE RELEVANT TO YOUR ORGANIZATION? (SELECT ALL THAT APPLY)	
Real-time analytics / operational intelligence	79%
Log analytics and observability	78%
Security analytics	74%
Business intelligence and reporting	72%
AI-related use cases	71%
Search / retrieval	67%

A2.1: Organizations' tooling preferences

2026 Data Infrastructure Survey, Q11, N=294

HOW DOES YOUR ORGANIZATION PREFER TO ADDRESS THESE USE CASES FROM A TOOLING PERSPECTIVE? (SELECT ONE)	
One core platform with specialized tools where necessary	44%
Best individual tool for each use case	26%
One tool that supports multiple use cases	14%
No standardized approach	12%
Do not know or not sure	4%

A2.2: Organizations' tooling preferences by company size

2026 Data Infrastructure Survey, Q11, N=294

HOW DOES YOUR ORGANIZATION PREFER TO ADDRESS THESE USE CASES FROM A TOOLING PERSPECTIVE? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
One core platform with specialized tools where necessary	35%	35%	35%	45%	50%	43%
Best individual tool for each use case	29%	23%	28%	23%	27%	25%
One tool that supports multiple use cases	20%	19%	23%	23%	14%	16%
No standardized approach	12%	19%	12%	5%	4%	13%
Don't know or not sure	4%	4%	2%	5%	5%	3%

A3: Agreement with statements on consolidation and interoperability

2026 Data Infrastructure Survey, Q15, N=294

TO WHAT EXTENT DO YOU AGREE WITH THE FOLLOWING STATEMENTS: (SELECT ONE RESPONSE PER ROW)				
	Agree	Disagree	Neither	NA/NS
We would consolidate to fewer platforms if it meaningfully reduced costs.	80%	4%	9%	7%
Managing multiple specialized systems adds too much operational overhead.	75%	6%	15%	4%
Search quality and accuracy are more important than architectural simplicity.	60%	6%	27%	7%
Deploying infrastructure independently of a specific vendor/cloud provider is a strategic priority for us.	71%	8%	15%	5%
Interoperability across vendors and tools is critical when evaluating data infrastructure.	78%	3%	12%	7%

A4: Approximate annual spend across data infrastructure platforms

2026 Data Infrastructure Survey, Q22, N=294

APPROXIMATELY HOW MUCH DOES YOUR ORGANIZATION SPEND ANNUALLY ACROSS SEARCH, ANALYTICS, OBSERVABILITY, AND RELATED DATA INFRASTRUCTURE PLATFORMS? (SELECT ONE)	
Under \$100K	22%
\$100K-\$500K	10%
\$500K-\$1M	10%
\$1M-\$5M	16%
\$5M-\$10M	5%
Over \$10M	8%
Do not know or not sure	29%

A5: Criteria rated when evaluating or selecting a data infrastructure platform

2026 Data Infrastructure Survey, Q12, N=294

WHEN EVALUATING OR SELECTING A DATA INFRASTRUCTURE PLATFORM, HOW IMPORTANT ARE THE FOLLOWING FACTORS? (SELECT ONE RESPONSE PER ROW)				
	Very important	Somewhat important	Not important	NA/NS
Total cost of ownership	80%	18%	1%	1%
Vendor independence and deployment flexibility	69%	23%	5%	3%
Enterprise readiness (e.g., LTS versions, SLAs)	67%	27%	3%	3%
Community and ecosystem health	53%	40%	3%	4%
Security and compliance (certifications, regulatory alignment)	79%	16%	2%	2%
AI and vector search capabilities	49%	34%	12%	5%
Cloud native architecture and scalability	61%	28%	8%	3%
Availability of an open source alternative or community-driven version	53%	34%	6%	6%

A6: Biggest concerns around data infrastructure risk

2026 Data Infrastructure Survey, Q14, N=294

WHAT ARE YOUR ORGANIZATION'S BIGGEST CONCERNS AROUND DATA INFRASTRUCTURE RISK? (SELECT UP TO THREE RESPONSES)	
Security vulnerabilities	58%
Regulatory compliance	33%
Data exposure across multiple tools or environments	33%
Vendor lock-in	32%
Operational complexity across multiple platforms	31%
Loss of data control	27%
Vendor or project instability / discontinuation	22%
Migration complexity and switching costs	22%
Limited visibility into the technology stack	22%
Other	22%
Do not know or not sure	22%

A7: Preferred search approach

2026 Data Infrastructure Survey, Q13, N=294

WHICH SEARCH APPROACH DOES YOUR ORGANIZATION PREFER? (SELECT ONE)	
Hybrid search	68%
Keyword search	13%
Semantic / vector search	7%
Other	1%
Do not know or not sure	11%

A8.1: Preferred vendor relationship

2026 Data Infrastructure Survey, Q20, N=294

DOES YOUR ORGANIZATION PREFER TO WORK WITH A SINGLE VENDOR OR MULTIPLE VENDORS FOR SEARCH, ANALYTICS, OBSERVABILITY AND DATA INFRASTRUCTURE SUPPORT? (SELECT ONE)	
Multiple vendors / consultancies	29%
Single vendor	28%
In-house expertise	25%
No preference	13%
Do not know or not sure	5%

A8.2: Preferred support model

2026 Data Infrastructure Survey, Q21, N=291

WHAT TYPE OF SUPPORT MODEL DOES YOUR ORGANIZATION PREFER? (SELECT ONE)	
Commercial / vendor-backed support	37%
Community-based support	23%
Vendors accredited by the project	21%
No preference	13%
Do not know or not sure	6%

A9.1: Preferred vendor relationship by respondent role 2026 Data Infrastructure Survey, Q20, N=294

DOES YOUR ORGANIZATION PREFER TO WORK WITH A SINGLE VENDOR OR MULTIPLE VENDORS FOR SEARCH, ANALYTICS, OBSERVABILITY AND DATA INFRASTRUCTURE SUPPORT? (SELECT ONE)										
	CxO	VP Eng/ Infra	VP IT	VP Data	Platform / DevOps	SW Dev	Data Eng	AI/ML Eng	Security Eng	Observability
Multiple vendors / consultancies	44%	50%	76%	27%	26%	21%	30%	33%	28%	0%
Single vendor	32%	17%	18%	69%	23%	16%	10%	33%	22%	36%
In-house expertise	13%	22%	6%	0%	32%	39%	50%	7%	28%	36%
No preference	10%	11%	0%	4%	12%	13%	10%	20%	11%	18%
DKNS	1%	0%	0%	0%	7%	11%	0%	7%	11%	9%

A9.2: Preferred support model by respondent role 2026 Data Infrastructure Survey, Q21, N=291

WHAT TYPE OF SUPPORT MODEL DOES YOUR ORGANIZATION PREFER? (SELECT ONE)										
	CxO	VP Eng/ Infra	VP IT	VP Data	Platform / DevOps	SW Dev	Data Eng	AI/ML Eng	Security Eng	Observability
Commercial / vendor-backed support	35%	50%	48%	81%	25%	34%	10%	53%	28%	45%
Vendors accredited by project	35%	17%	36%	12%	19%	26%	40%	7%	28%	27%
Community-based support	19%	17%	6%	4%	32%	21%	20%	13%	22%	9%
No preference	12%	6%	9%	0%	14%	8%	20%	20%	6%	18%
DKNS	1%	6%	0%	0%	10%	8%	10%	7%	17%	0%

A9.3: Preferred vendor relationship by company size 2026 Data Infrastructure Survey, Q20, N=294

DOES YOUR ORGANIZATION PREFER TO WORK WITH A SINGLE VENDOR OR MULTIPLE VENDORS FOR SEARCH, ANALYTICS, OBSERVABILITY AND DATA INFRASTRUCTURE SUPPORT? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Multiple vendors / consultancies	15%	17%	58%	39%	44%	37%
Single vendor	30%	39%	19%	32%	31%	26%
In-house expertise	36%	22%	16%	14%	14%	23%
No preference	17%	17%	5%	9%	8%	8%
DKNS	2%	4%	2%	5%	3%	6%

A9.4: Preferred support model by company size 2026 Data Infrastructure Survey, Q21, N=291

DOES YOUR ORGANIZATION PREFER TO WORK WITH A SINGLE VENDOR OR MULTIPLE VENDORS FOR SEARCH, ANALYTICS, OBSERVABILITY AND DATA INFRASTRUCTURE SUPPORT? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Commercial / vendor-backed support	19%	26%	26%	43%	47%	47%
Vendors accredited by project	15%	22%	40%	34%	25%	19%
Community-based support	47%	30%	19%	8%	14%	15%
No preference	11%	13%	7%	12%	11%	8%
DKNS	9%	4%	7%	1%	3%	8%

A10: Minimum LTS duration expected for production-critical platforms

2026 Data Infrastructure Survey, Q23, N=294

WHAT IS THE MINIMUM LONG-TERM SUPPORT (LTS) DURATION YOUR ORGANIZATION EXPECTS FOR PRODUCTION-CRITICAL DATA INFRASTRUCTURE PLATFORMS? (SELECT ONE)	
12 months or less	12%
18 months	10%
24 months	21%
36 months	21%
48 months	6%
60 months or more	11%
Do not know or not sure	20%

A11.1: Organizations running or planning to run AI workloads by geography

2026 Data Infrastructure Survey, Q16, N=294

IS YOUR ORGANIZATION CURRENTLY RUNNING OR PLANNING TO RUN AI WORKLOADS/APPLICATIONS/PROJECTS? (SELECT ONE)										
	US	Canada	UK	EU	Other Europe	China	India	Japan	Oceania	Other APAC
Yes	90%	70%	91%	82%	85%	96%	56%	100%	60%	89%
No	5%	20%	5%	9%	8%	0%	19%	0%	20%	11%
DKNS	5%	10%	4%	9%	8%	4%	25%	0%	20%	0%

A11.2: Organizations running or planning to run AI workloads by company size

2026 Data Infrastructure Survey, Q16, N=294

IS YOUR ORGANIZATION CURRENTLY RUNNING OR PLANNING TO RUN AI WORKLOADS/ APPLICATIONS/PROJECTS? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Yes	78%	65%	82%	94%	90%	94%
No	12%	27%	12%	5%	6%	2%
DKNS	10%	8%	7%	1%	4%	4%

A12: Leading AI use cases

2026 Data Infrastructure Survey, Q17, N=245

WHAT KIND OF AI WORKLOADS IS YOUR ORGANIZATION RUNNING OR PLANNING TO RUN? (SELECT ALL THAT APPLY)	
Generative AI / LLM-powered applications	82%
AI-assisted operational analytics / observability	56%
Enterprise knowledge / BI assistants or copilots	55%
Retrieval-Augmented Generation (RAG)	51%
Document analysis / information extraction	47%
Predictive analytics / forecasting	47%
Threat detection / risk analysis	42%
Natural language interfaces over enterprise data	41%
Recommendations / personalization	29%
Other	1%
Do not know or not sure	1%

A13.1: Data used or planned to be used to power AI applications

2026 Data Infrastructure Survey, Q18, N=245

WHAT DATA DOES YOUR ORGANIZATION USE OR PLAN TO USE TO POWER AI APPLICATIONS? (SELECT ALL THAT APPLY)	
Log, metrics, and operational data	64%
Business and transactional data	59%
Security and threat intelligence data	55%
Internal knowledge documentation	55%
Real-time event data	54%
Customer and behavioral data	49%
Emails, chats logs, and collaboration content	41%
Product and inventory data	37%
External / third-party data sources	27%
Other	2%
Do not know or not sure	2%

A13.2: Access to AI data

2026 Data Infrastructure Survey, Q19, N=245

HOW IMPORTANT IS IT TO YOUR ORGANIZATION THAT ALL OF THIS DATA IS ACCESSIBLE FROM A SINGLE PLATFORM FOR AI APPLICATIONS? (SELECT ONE)	
Very important	51%
Somewhat important	33%
Not important	9%
Do not know or not sure	7%

A14.1: Familiarity with OpenSearch 2026

2026 Data Infrastructure Survey, Q25, N=294

HOW FAMILIAR ARE YOU WITH OPENSEARCH? (SELECT ONE)	
Very familiar	30%
Somewhat familiar	37%
Aware of it, but not very familiar	22%
Not familiar at all	11%

A14.2: Familiarity with OpenSearch 2026 by respondent role 2026 Data Infrastructure Survey, Q25, N=294

HOW FAMILIAR ARE YOU WITH OPENSEARCH? (SELECT ONE)										
	CxO	VP Eng/ Infra	VP IT	VP Data	Platform / DevOps	SW Dev	Data Eng	AI/ML Eng	Security Eng	Observability
Very familiar	61%	50%	52%	58%	27%	21%	40%	27%	22%	27%
Somewhat familiar	27%	28%	45%	38%	33%	26%	50%	53%	33%	45%
Aware of it, but not very familiar	10%	17%	3%	4%	25%	32%	10%	7%	11%	27%
Not familiar at all	2%	6%	0%	0%	15%	21%	0%	13%	33%	0%

A14.3: Familiarity with OpenSearch 2026 by company size 2026 Data Infrastructure Survey, Q25, N=294

HOW FAMILIAR ARE YOU WITH OPENSEARCH? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Very familiar	32%	26%	42%	42%	51%	48%
Somewhat familiar	32%	30%	33%	45%	31%	21%
Aware of it, but not very familiar	28%	13%	14%	12%	10%	21%
Not familiar at all	9%	30%	11%	1%	8%	10%

A14.4: Experience with OpenSearch 2024

2024 Search and Analytics Survey, Q15, N=689

WHAT IS YOUR EXPERIENCE WITH THE OPEN SOURCE OPENSEARCH SOFTWARE? (SELECT ONE)	
Very experienced in using it in a production environment	10%
Have some experience using it in a production environment	9%
Have experimented with it	11%
Heard of it, but that's all	38%
Have not heard of it	32%

A15: Describing OpenSearch

2026 Data Infrastructure Survey, Q26, N=105, Open-ended, Optional, Thematically categorized

HOW WOULD YOU DESCRIBE OPENSEARCH TO A COLLEAGUE UNFAMILIAR WITH IT? (OPTIONAL)		
	Mentions	Responses
Search & retrieval	43	47%
Analytics, observability & monitoring	36	39%
Open source & community	31	34%
Elasticsearch fork / alternative	26	28%
AI, ML & semantic search	10	11%
Scalability, flexibility & performance	10	11%
Positive sentiment / community advocacy	8	9%
Security analytics	6	7%
Data infrastructure / platform	6	7%
Vendor neutrality & independence	4	4%

A16: OpenSearch vendor neutrality and enterprise readiness

2026 Data Infrastructure Survey, Q27, Q28, N=250

	Average Score
ON A SCALE OF 1 TO 5, HOW NEUTRAL/VENDOR-INDEPENDENT DO YOU PERCEIVE OPENSEARCH TO BE TODAY?	3.66
ON A SCALE OF 1 TO 5, HOW WOULD YOU RATE OPENSEARCH'S READINESS FOR MISSION-CRITICAL ENTERPRISE WORKLOADS?	3.70

A18.1: Perceived relevance of OpenSearch to AI

2026 Data Infrastructure Survey, Q29, N=250

HOW DO YOU PERCEIVE THE ROLE OF OPENSEARCH IN AI ARCHITECTURES? (SELECT ONE)	
Core infrastructure for AI applications (e.g., retrieval layer for RAG)	44%
Supporting component (e.g., search or analytics tool)	33%
No relevance to AI use cases	6%
Other	1%
Do not know or not sure	16%

A18.2: Perceived relevance of OpenSearch to AI by familiarity

2026 Data Infrastructure Survey, Q29, N=250

HOW DO YOU PERCEIVE THE ROLE OF OPENSEARCH IN AI ARCHITECTURES? (SELECT ONE)			
	Very familiar	Somewhat familiar	Aware of it, but not very familiar
Core infrastructure for AI applications (e.g., retrieval layer for RAG)	61%	48%	21%
Supporting component (e.g., search or analytics tool)	33%	36%	34%
No relevance to AI use cases	4%	5%	10%
Do not know or not sure	2%	11%	34%

A19.1: Current status of OpenSearch use in the technology stack

2026 Data Infrastructure Survey, Q34, N=245

IS OPENSEARCH CURRENTLY USED ANYWHERE WITHIN YOUR ORGANIZATION'S TECHNOLOGY STACK? (SELECT ONE)	
Yes, actively in production environments	36%
Yes, but only in testing or development environments	18%
No, but we are planning to	18%
No and we have no plans to use it	12%
Do not know or not sure	17%

A19.2: Current status of OpenSearch use in the technology stack by geography

2026 Data Infrastructure Survey, Q34, N=245

IS OPENSEARCH CURRENTLY USED ANYWHERE WITHIN YOUR ORGANIZATION'S TECHNOLOGY STACK? (SELECT ONE)										
	US	Canada	UK	EU	Other Europe	China	India	Japan	Oceania	Other APAC
Production use	49%	22%	39%	50%	25%	67%	31%	63%	0%	25%
Test/dev use	19%	0%	32%	13%	33%	17%	15%	13%	50%	63%
Plans to use	9%	22%	17%	13%	33%	17%	23%	13%	0%	0%
No plans to use	9%	22%	10%	11%	0%	0%	8%	13%	50%	0%
DKNS	14%	33%	2%	13%	8%	0%	23%	0%	0%	13%

A19.3: Current status of OpenSearch use in the technology stack by company size

2026 Data Infrastructure Survey, Q34, N=245

IS OPENSEARCH CURRENTLY USED ANYWHERE WITHIN YOUR ORGANIZATION'S TECHNOLOGY STACK? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Production use	29%	38%	39%	52%	52%	45%
Test/dev use	5%	19%	32%	23%	26%	10%
Plans to use	37%	25%	11%	5%	14%	10%
No plans to use	16%	6%	9%	7%	5%	12%
DKNS	13%	13%	9%	13%	3%	22%

A20: Awareness and adoption 2026 vs. 2024

2026 Data Infrastructure Survey (N=294), 2024 Search and Analytics Survey (N=689)

HOW FAMILIAR ARE YOU WITH OPENSEARCH? (2026) IS OPENSEARCH CURRENTLY USED ANYWHERE WITHIN YOUR ORGANIZATION'S TECHNOLOGY STACK? (2024)		
	2026	2024
Production use	36%	19%
Experimentation	18%	11%
At least heard of OpenSearch	89%	68%

A21.1: OpenSearch deployment depth

2026 Data Infrastructure Survey, Q38, N=128

HOW EMBEDDED IS OPENSEARCH WITHIN YOUR ORGANIZATION'S CRITICAL INFRASTRUCTURE? (SELECT ONE)	
Used in production, but replaceable	40%
Deeply embedded into core systems	23%
Experimental or proof of concept	20%
Mission-critical	13%
Do not know or not sure	4%

A21.2: OpenSearch deployment depth by geography 2026 Data Infrastructure Survey, Q38, N=128

HOW EMBEDDED IS OPENSEARCH WITHIN YOUR ORGANIZATION'S CRITICAL INFRASTRUCTURE? (SELECT ONE)										
	US	Canada	UK	EU	Other Europe	China	India	Japan	Oceania	Other APAC
Production but replaceable	32%	50%	29%	64%	57%	25%	33%	17%	0%	29%
Deeply embedded	32%	0%	32%	21%	0%	20%	0%	17%	50%	14%
Experimental/PoC	23%	0%	29%	4%	43%	45%	50%	50%	0%	57%
Mission-critical	10%	50%	11%	11%	0%	10%	0%	17%	0%	0%
DKNS	3%	0%	0%	0%	0%	0%	17%	0%	50%	0%

A21.3: OpenSearch deployment depth by company size 2026 Data Infrastructure Survey, Q38, N=128

HOW EMBEDDED IS OPENSEARCH WITHIN YOUR ORGANIZATION'S CRITICAL INFRASTRUCTURE? (SELECT ONE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Production but replaceable	25%	56%	52%	36%	29%	26%
Deeply embedded	17%	0%	23%	24%	34%	22%
Experimental/PoC	42%	22%	23%	29%	26%	26%
Mission-critical	17%	11%	3%	11%	7%	22%
DKNS	0%	11%	0%	0%	4%	0%

A22: Types of OpenSearch deployment in use

2026 Data Infrastructure Survey, Q37, N=128, DKNS excluded

WHAT TYPES OF OPENSEARCH DEPLOYMENTS DOES YOUR ORGANIZATION CURRENTLY USE? (SELECT ALL THAT APPLY)	
Self-managed in the cloud	60%
Managed service via Amazon OpenSearch Service	48%
Self-managed on-premises	43%
Managed service from a vendor other than AWS	9%

A23: OpenSearch deployment evolution (2026 vs. 2024)

2026 Data Infrastructure Survey, Q37, N=126; 2024 Search and Analytics Survey, Q17, N=174;
DKNS responses excluded

WHAT TYPES OF OPENSEARCH DEPLOYMENTS DOES YOUR ORGANIZATION CURRENTLY USE? (SELECT ALL THAT APPLY)		
	2026	2024
Self-service mode		37%
Self-managed cloud	60%	
Self-managed on-premises	43%	
Amazon OpenSearch Service	48%	44%
Managed service from a vendor other than AWS	9%	12%

A24: Where OpenSearch is delivering for active users

2026 Data Infrastructure Survey, Q36, N=128

WHICH USE CASES DOES YOUR ORGANIZATION USE OPENSEARCH FOR? (SELECT ONE RESPONSE PER ROW)			
	Yes	No	NA/NS
Search / retrieval	91%	7%	2%
Log analytics and observability	83%	14%	3%
Real-time analytics / operational intelligence	70%	29%	2%
AI-related use cases (e.g., RAG, vector search)	62%	28%	10%
Business intelligence and reporting	62%	27%	11%
Security analytics (e.g., threat detection, SIEM)	52%	37%	11%

A25: How well is OpenSearch delivering on capabilities 2026 Data Infrastructure Survey, Q40, N=128

HOW WELL IS OPENSEARCH DELIVERING ON THE FOLLOWING CAPABILITIES? (SELECT ONE RESPONSE PER ROW)				
	Under-performing	Meeting expectations	Exceeding expectations	NA/NS
Search and retrieval capabilities	0%	49%	45%	6%
Observability capabilities	2%	67%	26%	5%
Security, privacy, and compliance	3%	75%	13%	9%
Scalability and resilience	4%	63%	27%	6%
Integration with other systems and tools	5%	64%	22%	9%
Operational experience	7%	64%	21%	8%
Dashboards and visualization	5%	66%	17%	12%
Contributors and community	3%	57%	27%	13%
AI and vector search capabilities	2%	38%	35%	25%

A26: Primary reasons for not adopting OpenSearch 2026 Data Infrastructure Survey, Q41, N=71, DKNS responses excluded

WHAT ARE THE PRIMARY REASONS YOUR ORGANIZATION HAS NOT ADOPTED OPENSEARCH? (SELECT UP TO THREE)	
Limited awareness or familiarity	30%
Insufficient internal expertise or operational resources	27%
Integration or migration complexity	27%
No compelling business case to migrate from current platform	25%
Concerns about vendor neutrality or association with AWS	24%
We are currently evaluating OpenSearch	24%
Concerns about enterprise readiness or support options	20%
We evaluated it and selected another platform	10%
Procurement, compliance, or regulatory constraints	9%
Other	7%

A27: Non-user organizations evaluating OpenSearch 2026 Data Infrastructure Survey, Q42, N=73

HOW OPEN IS YOUR ORGANIZATION TO EVALUATING OPENSEARCH IN THE NEXT 12-18 MONTHS? (SELECT ONE)	
Actively planning to evaluate or expand usage	14%
Open to evaluating if a strong business case is demonstrated	23%
Potentially open, but not currently a priority	44%
Not open – we are satisfied with our current platform	12%
Do not know or not sure	7%

A28.1: Popular information channels 2026 Data Infrastructure Survey, Q30, N=240

WHAT KIND OF MEDIA SOURCES DO YOU PRIMARILY USE TO GET INFORMATION ON PLATFORMS SUCH AS OPENSEARCH? (SELECT UP TO THREE)	
Open source communities or project documentation	58%
Developer communities and technical forums	45%
Conferences, events, and meetups	29%
Videos, webinars, and online technical content	28%
Social media and professional networks (e.g., LinkedIn, X, Reddit)	28%
Technology media and press coverage	25%
Vendor websites, blogs, and marketing communications	19%
Internal teams or peer recommendations	18%
Industry analyst firms and research reports	18%
Podcasts and industry interviews	6%
Other	1%

A28.2: Popular information channels by respondent role 2026 Data Infrastructure Survey, Q30, N=240

WHAT KIND OF MEDIA SOURCES DO YOU PRIMARILY USE TO GET INFORMATION ON PLATFORMS SUCH AS OPENSEARCH? (SELECT UP TO THREE)										
	CxO	VP Eng/ Infra	VP IT	VP Data	Platform / DevOps	SW Dev	Data Eng	AI/ML Eng	Security Eng	Observability
Open source communities or project documentation	39%	47%	36%	81%	56%	57%	70%	54%	42%	36%
Developer communities and technical forums	30%	47%	39%	65%	42%	61%	0%	62%	50%	18%
Conferences, events, and meetups	22%	29%	27%	15%	33%	25%	10%	23%	8%	45%
Videos, webinars, and online technical content	27%	35%	36%	35%	16%	36%	20%	38%	17%	55%
Social media and professional networks	37%	24%	21%	4%	15%	25%	50%	46%	50%	27%
Technology media and press coverage	28%	6%	42%	4%	33%	18%	40%	23%	25%	9%
Vendor websites, blogs, and marketing communications	34%	24%	36%	4%	20%	11%	0%	0%	25%	18%
Internal teams or peer recommendations	24%	24%	15%	19%	13%	18%	30%	15%	25%	27%
Industry analyst firms and research reports	31%	29%	21%	38%	9%	14%	30%	8%	25%	27%
Podcasts and industry interviews	10%	6%	9%	4%	9%	4%	0%	8%	8%	0%
DKNS	1%	6%	0%	0%	7%	4%	0%	0%	0%	0%

A28.3: Popular information channels by company size

2026 Data Infrastructure Survey, Q30, N=240

WHAT KIND OF MEDIA SOURCES DO YOU PRIMARILY USE TO GET INFORMATION ON PLATFORMS SUCH AS OPENSEARCH? (SELECT UP TO THREE)						
	Startups	Small	Mid-size	Mid-market	Large	Global
Open source communities or project documentation	53%	56%	38%	45%	46%	46%
Developer communities and technical forums	42%	50%	29%	42%	39%	36%
Conferences, events, and meetups	29%	38%	27%	19%	25%	29%
Videos, webinars, and online technical content	32%	13%	31%	30%	24%	38%
Social media and professional networks	37%	38%	21%	29%	27%	29%
Technology media and press coverage	26%	19%	38%	26%	22%	29%
Vendor websites, blogs, and marketing communications	21%	13%	15%	29%	31%	16%
Internal teams or peer recommendations	5%	6%	29%	16%	24%	21%
Industry analyst firms and research reports	11%	13%	25%	23%	32%	21%
Podcasts and industry interviews	3%	13%	4%	5%	14%	5%
DKNS	5%	0%	6%	3%	1%	2%

A29: OpenSearch Software Foundation / Linux Foundation membership

2026 Data Infrastructure Survey, Q31, N=250

IS YOUR ORGANIZATION A MEMBER OF THE OPENSEARCH SOFTWARE FOUNDATION OR THE LINUX FOUNDATION? (SELECT ONE)	
Yes, both	6%
Linux Foundation Only	9%
No	67%
Do not know or not sure	18%

A30: Triggers for Foundation membership

2026 Data Infrastructure Survey, Q32, N=185

WHAT WOULD NEED TO BE TRUE FOR YOUR ORGANIZATION TO JOIN THE OPENSEARCH SOFTWARE FOUNDATION? (SELECT ALL THAT APPLY)	
Clear business case – demonstrable ROI or cost advantage from formal membership	53%
Enterprise readiness proof – certified LTS versions, SLAs, and formal support options	47%
Vendor neutrality confidence – greater assurance that OpenSearch is truly independent	40%
Roadmap influence – ability to shape OpenSearch's direction and priorities as a member	31%
Exclusive member benefits – access to enterprise features, early releases, or capabilities not available to non-members	21%
Peer validation – evidence that similar organizations in our industry are already members	20%
Do not know or not sure	22%

A31: Importance of potential Foundation membership benefits

2026 Data Infrastructure Survey, Q33, N=185

IF YOU WERE CONSIDERING JOINING THE OPENSEARCH SOFTWARE FOUNDATION, HOW IMPORTANT WOULD THE FOLLOWING MEMBERSHIP BENEFITS BE? (SELECT ONE RESPONSE PER ROW)				
	Very important	Somewhat important	Not important	NA/NS
Priority security patches and vulnerability notifications	79%	14%	3%	4%
Long-term support (LTS) with an 18-month support window	70%	21%	3%	6%
Direct input into the product roadmap	52%	35%	5%	8%
Access to a network of certified vendors and consultancies	47%	39%	9%	5%
Early access to new features before general release	39%	37%	17%	6%

About the authors

CSILLA ZSIGRI brings a multicultural perspective to technology, innovation, and strategic insight, shaped by years of international work spanning market research, tech startups, and innovation consulting. With a background in economics, and management & business consulting, she spent over a decade as a technology industry analyst and advisor at 451 Research and S&P Global. She has also worked on EU-funded R&D&I projects, contributing market intelligence and strategic framing to help bring early-stage technologies to real-world impact, and served as Chief Strategy Officer for two technology startups, shaping market positioning and go-to-market strategy from the inside. Csilla is known for turning complex ideas into clear, actionable insight for clients ranging from emerging ventures to enterprises and institutions like the European Commission.

ADRIENN LAWSON serves as Director of Quantitative Research at the Linux Foundation, where she leads data-driven initiatives to understand open source ecosystems. With expertise in social data science from the University of Oxford and a background spanning academic and governmental research, she brings methodological rigor to analyzing distributed collaboration networks. At the Linux Foundation, Adrienn leads a team conducting cross-sectional research across industry verticals and geographic regions to provide comprehensive insights into open source dynamics. Her work encompasses empirical investigations into regulatory compliance, the implications of AI, and sustainable funding models. She produces evidence-based recommendations that inform strategic decision-making within the open source community.

HILARY CARTER joined the Linux Foundation in 2021 to launch and lead LF Research, established to deliver empirical insights into open source trends, opportunities, and challenges. Prior to joining the Linux Foundation, Hilary was the managing director of the Toronto-based Blockchain Research Institute, a global, syndicated think tank focused on blockchain technology. She has contributed to nearly 200 research projects focused on open source innovation and its adoption across industries. Hilary received a Master of Science in Management from the London School of Economics, and holds dual Canadian and Irish citizenship.

Acknowledgments

We thank all the people who participated in the survey. Special thanks to Linux Foundation colleagues for their involvement in the various stages of the research process: Lisa Briggs, Anna Hermansen, Bianca Lewis, James McIntyre, and Kristi Piechnik, as well as colleagues from the LF Creative Services and Hubspot teams.

LF Research uses generative AI tooling in varying degrees for research-related tasks including aggregating secondary sources, analyzing data, supporting report drafting, copyediting, and translation. All AI-generated content is reviewed and revised as needed.



 x.com/linuxfoundation

 facebook.com/TheLinuxFoundation

 linkedin.com/company/the-linux-foundation

 youtube.com/user/TheLinuxFoundation

 github.com/linuxfoundation

September 2026



Copyright © 2026 The Linux Foundation

This report is licensed under the Creative Commons Attribution-NonCommercial 4.0 International Public License.

To reference this work, please cite as follows:
Csilla Zsigri, Adrienn Lawson, and Hilary Carter,
“The 2026 Open Data Infrastructure Report: AI,
Governance, and OpenSearch Adoption Trends,”
foreword by Bianca Lewis, The Linux Foundation,
September 2026.



The [OpenSearch Software Foundation](https://opensearch.org/foundation) is a vendor-neutral community for search, analytics, observability, and vector database software. Hosted by the Linux Foundation and supported by premier members such as AWS, IBM, SAP and Uber, the OpenSearch Software Foundation works with community maintainers, developers, and member organizations to drive the continued growth of the OpenSearch project. With more than 2 billion software downloads since its inception and participation from thousands of contributors, the OpenSearch project and its community are transforming how information is managed and discovered. To learn more, please visit opensearch.org/foundation.



Founded in 2021, [Linux Foundation Research](https://linuxfoundation.org/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. By leveraging project databases and networks and committing 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 worldwide.