Action1

2024 AI Impact on Sysadmins: Survey Report

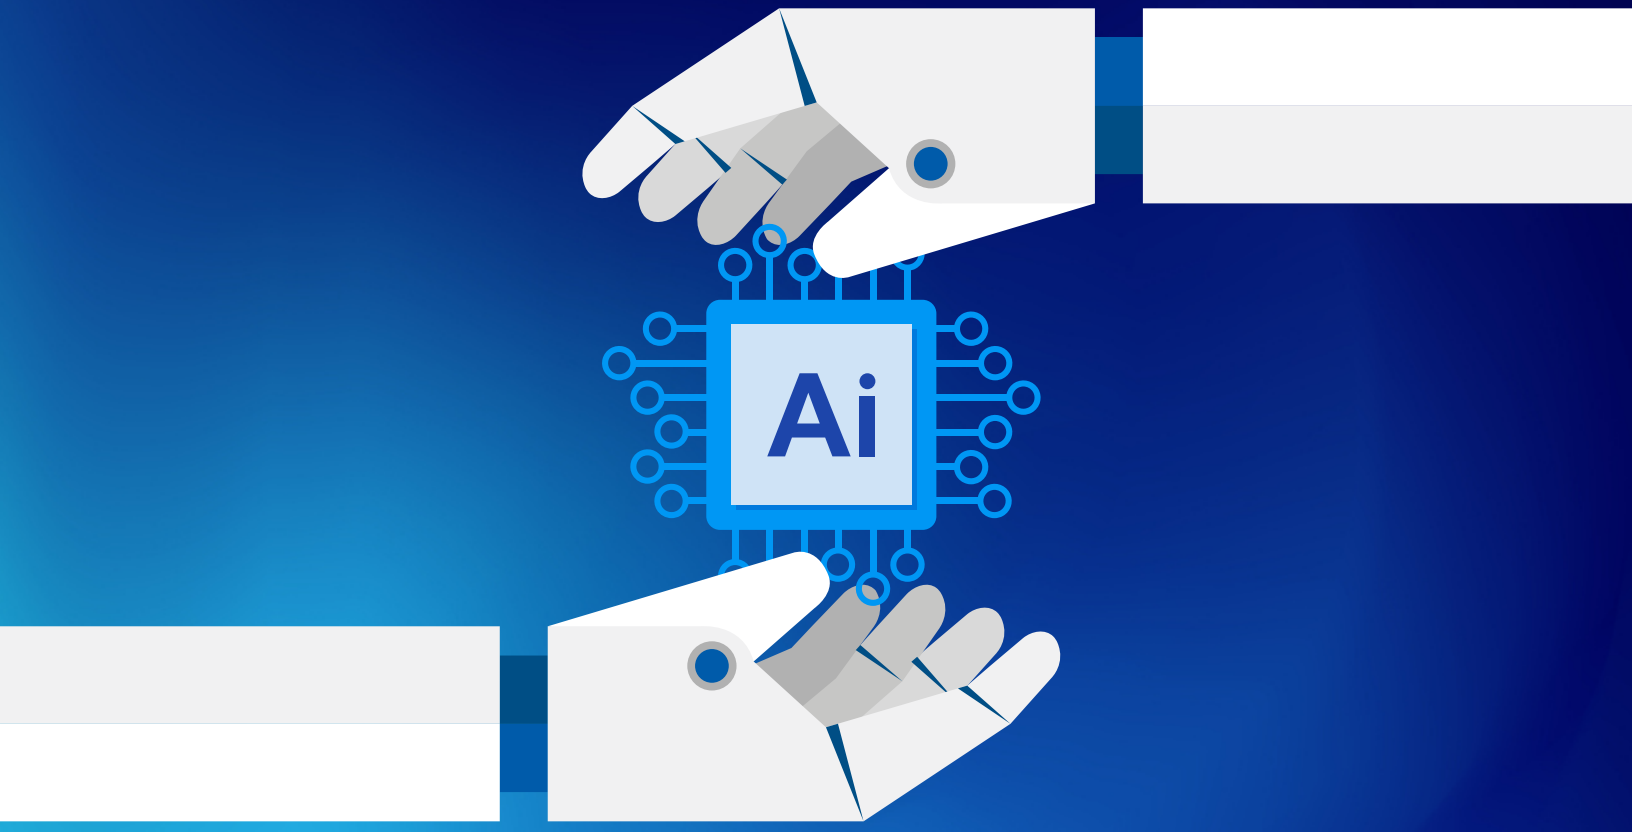

July 2024

Contents

Introduction 3

Executive Summary 4

Detailed Findings 7

- Sysadmin Functions and the Likelihood of AI Replacement 7
- Growing Demand for AI Training Among Sysadmins 8
- Many Sysadmins Still Struggle with AI Understanding 9
- Current Applications and Failures in AI Implementation 10

Key Recommendations 13

Appendix 14

- Methodology 14
- Detailed Responses in the “Other” Field 14
- Demography 15

Introduction

We are conducting this survey for the second consecutive year to explore the potential impact of generative AI on sysadmins' jobs. Our goal is to understand how sysadmins' perceptions, concerns, and expectations regarding the integration of AI into their job functions and its potential to contribute to organizational success have evolved over the past year. As the initial hype around AI transitions to a more pragmatic, mature, and practical approach, we aim to capture these changes.

In addition to exploring perceptions, this year's survey includes questions about experiences with trial and error in AI implementation.

The survey gathered insights directly from 450 sysadmins worldwide. By analyzing their responses, this report presents a comprehensive view of their thoughts and experiences on AI's potential.

This information will help IT professionals, software vendors, and organizations navigate the evolving landscape of AI, empowering sysadmins to adapt and excel in the changing technological environment.

Executive Summary

The survey's insights shed light on both the opportunities and challenges associated with AI adoption among system administrators (sysadmins). Here are the key takeaways:

Stability in AI Perceptions

Sysadmins have a clear and consistent understanding of the areas where AI can automate tasks. Their views have remained steady over the past year, identifying the top three areas for AI automation in the next two years: log analysis, server CPU and memory monitoring, and patch management (the latter has risen in expectations for AI automation from fourth to third place). However, areas requiring human judgment, such as user rights administration, single sign-on (SSO), and password management, are perceived as less likely to be automated by AI, as was the case last year.

Knowledge Gap and Training Need

Despite these clear perceptions, a significant 60% of sysadmins acknowledge a lack of understanding of how to leverage AI from a practical standpoint. This figure, while down from 73% last year, indicates a persistent gap in AI literacy. To address this, 72% of respondents expressed a need for additional training, up from 63% the previous year. Moreover, 45% of sysadmins are concerned about becoming obsolete in the job market due to their current level of AI literacy. This knowledge gap suggests that while there is interest and potential for AI, effective adoption will require substantial investment in education and training. Organizations need to address this gap to fully leverage AI technologies.

72%
of respondents expressed a need for additional AI training

Mixed Outcomes in Current AI Implementations

The varied success of AI implementations shows that while AI has potential, its performance is inconsistent. Specifically, AI is most commonly implemented in log analysis (26%), troubleshooting (25%), and incident detection & remediation (16%).

AI failures in troubleshooting outnumbered the successes

The highest failure rates occur in troubleshooting, with failures in more than half of the organizations using it. This underscores that current AI developments cannot cover the nuances of modern IT solutions, addressing only the most common systems and errors. AI may lack sufficient training data to cover all possible scenarios and may not understand the context in which problems occur, leading to incorrect diagnoses.

Failures in implementing AI for log analysis were reported in one out of every four organizations. This is due to the complex nature of logs, which generate massive amounts of data with varying structures. This makes it difficult for AI models to interpret meaningful data amid vast noise, overwhelming AI algorithms.

Risk of Disruptions

The survey found that AI has led to critical disruptions in 16% of organizations.

Indeed, an AI that misdiagnoses the root cause of a system problem can lead to incorrect remediation steps. This misdirection can prolong system downtime, impacting business operations and productivity. For example, misinterpreting a network problem as a hardware failure could delay resolution and exacerbate the issue. Incorrect log analysis can lead to false positives, misinterpreting normal activity as a security incident, which overwhelms security teams and diverts attention from real threats. Conversely, false negatives could allow breaches to go undetected, posing serious security risks. AI-driven automation errors could misallocate computing resources, such as improperly scaling virtual machines or misconfiguring load balancing. This would degrade performance and potentially bring down critical applications.

These conclusions underscore the importance of being realistic about AI's potential and the need to establish a careful and well-managed AI integration when it is deemed justified.

Limited Organizational Commitment To AI Implementation

The survey indicates that 80% of organizations do not require sysadmins to implement AI in their job roles, a figure only slightly down from 82% last year. This finding suggests that while there is interest in AI, there is still a significant gap between recognition of its potential and its mandated application. This gap highlights a cautious approach by organizations, potentially due to the risks and mixed effectiveness seen in AI applications so far.

Detailed Findings

Sysadmin Functions and the Likelihood of AI Replacement

In this section of the survey, we have listed the key areas of sysadmin functions and asked respondents to rate the likelihood of AI replacing each area in the upcoming two years.

- Log analysis
- Server CPU and memory monitoring
- Streamlining patch management

Interestingly, these results are not significantly different from last year, except that patch management has moved up from fourth to third place, while vulnerability prioritization has dropped to fifth.

Regarding areas least likely to be replaced by AI, the data also shows little change from last year. These areas include managing user rights and administration of user permissions, as well as single sign-on (SSO) and password management.

Indeed, AI can analyze system logs and network traffic patterns to identify anomalies and generate alerts, improving accuracy over time with machine learning. This reduces manual monitoring efforts, allowing administrators to focus on critical tasks. AI's ability to process large data sets swiftly enhances the efficiency and precision of log analysis.

For patch management, AI can prioritize updates based on risk, streamlining distribution and deployment of security updates. While AI can optimize patching plans by assessing vulnerability severity and system criticality, challenges remain for production servers, such as scheduling maintenance windows and getting approval from system owners.

In contrast, managing user rights, permissions, SSO, and password management requires human judgment to consider organizational dynamics and potential risks, making these tasks better suited for human oversight rather than AI.

TABLE 1. Which areas of your functionality are most likely to be fully automated through AI in the forthcoming 2 YRs? Please rate on a scale of 1 to 5, where 1 represents 'Extremely Likely' and 5 represents 'Extremely Unlikely.'

Category Name	Extremely Likely	Likely	Neutral	Unlikely	Extremely Unlikely	N/A
Log analysis	41%	42%	10%	4%	2%	1%
Monitoring of server CPU and memory utilization	30%	38%	17%	11%	4%	1%
Patch management processes optimization to improve prioritization, scheduling, and deployment of security updates	25%	45%	16%	8%	5%	1%
Detecting & remediating incidents	22%	45%	19%	9%	5%	1%
Vulnerability prioritization	20%	46%	18%	11%	4%	2%
Analyzing organization's security controls and comparing to compliance requirements	18%	45%	19%	13%	5%	1%
Troubleshooting	17%	32%	25%	20%	5%	1%
Providing IT staff with guidance and training	16%	36%	21%	17%	9%	1%
Performing post-incident reviews	15%	34%	24%	15%	9%	2%
Providing end-users with first-level IT support	14%	31%	19%	20%	15%	1%
Installing & maintaining software	14%	31%	21%	20%	13%	1%
Managing SSO & passwords	10%	18%	26%	25%	20%	1%
Administering user permissions & administration	9%	22%	27%	23%	18%	1%
Defining system usage policies & procedures	8%	26%	26%	24%	14%	1%
Managing files	7%	15%	34%	25%	16%	2%

Growing Demand for AI Training Among Sysadmins

Overall, 72% of the respondents indicated that they are looking to take additional training to enhance their understanding of AI integration into their job functions. This number is up from an already significant 63% who expressed this intention last year. These data show that, despite the initial hype around AI subsiding, the demand for skills to apply AI effectively is only beginning to grow, as more people consider it essential.

72%
of admins are seeking additional AI integration training.

CHART 1. Are you looking to take additional training to better understand how to integrate AI into your job?

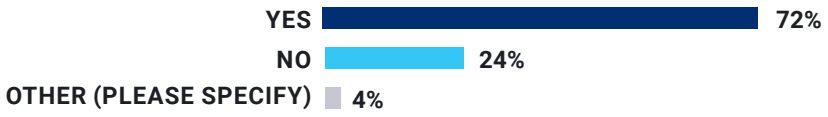

Nearly half of the respondents, 45%, expressed concerns about being left behind by other professionals who are more AI-literate. This figure has hardly changed from last year, when 47% of respondents reported fears of becoming obsolete in the job market due to insufficient AI skills.

CHART 2. Are you concerned that you might be left behind other professionals in your field who are more AI-literate?

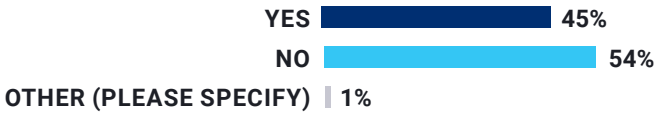

Many Sysadmins Still Struggle with AI Understanding

Despite the increasing awareness among sysadmins about the need to enhance their AI skills, companies are not yet requiring them to integrate AI into their work. This year's figure of 80% is not significantly different from last year's 82%. To some extent, this is logical, as AI technologies are still not fundamental in the IT stack.

At the same time, the question of missed opportunities for leveraging AI technologies remains open.

Ultimately, we see that the significant gap between the recognized importance of AI by the general public and the actual implementation within organizations has remained at the same level as last year.

80%
of organizations do not require admins to implement AI.

CHART 3. Does your company require you to start implementing AI into your job?

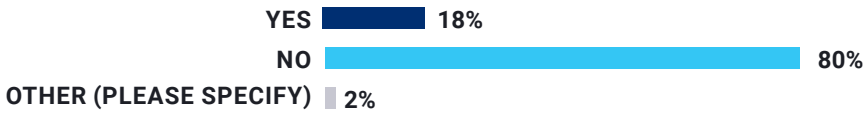

Although slightly fewer sysadmins than last year admit they don't understand how to implement AI—60% compared to 73% last year—the number is still substantial. This underscores that a significant knowledge gap among sysadmins when it comes to effectively incorporating AI into their job functions has persisted over the year. The existing disparity between strategic priorities and actual AI implementation within organizations is concerning, and it appears that many do not understand how to translate AI discussions into actionable steps. This disconnect between recognition and implementation underscores the need for greater organizational support, training, and clarity on how AI can be effectively integrated into sysadmin roles.

CHART 4. Do you have an explicit understanding of how to integrate AI into your job to improve the current processes and ensure that your company's IT stays on the cutting edge of the industry trends?

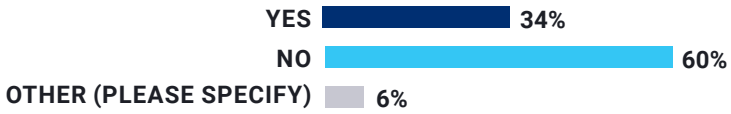

Current Applications and Failures in AI Implementation

Anticipating that companies should now actively start leveraging generative AI in their IT operations, we expanded this survey by adding a couple of questions about admins' real-life experience with AI.

Here are the top three areas where respondents have already implemented AI:

- Log analysis – 26%;
- Troubleshooting – 25%;
- Detecting & remediating incidents – 16%.

TABLE 2. In which of the listed areas have you already implemented AI?

Log analysis	26%
Troubleshooting	25%
Detecting & remediating incidents	16%
Providing IT staff with guidance and training	15%
Monitoring of server CPU and memory utilization	13%
Analyzing organization's security controls and comparing to compliance requirements	13%
Providing end-users with first-level IT support	12%
Vulnerability prioritization	11%
Installing & maintaining software	11%
Defining system usage policies & procedures	9%
Performing post-incident reviews	8%
Patch management processes optimization to improve prioritization, scheduling, and deployment of security updates	8%
Managing files	7%
Managing SSO & passwords	4%
Administering user permissions & administration	4%

Interestingly, 16% encountered errors due to AI that led to disruptions.

In most cases, failures occurred in the troubleshooting area as indicated by 16% of respondents. Log analysis was also named among the areas in which AI failed (6%). Another area for failures is providing end-users with first-level IT support (7%).

ANONYMOUS ANSWER:

> **The users want to speak with a person.**

Why does AI fail at troubleshooting?

Modern AI developments cannot cover the nuances of modern IT solutions; they can only fix problems related to the most common systems and errors.

System problems can be highly complex and unique. An AI may not have enough training data to cover all possible scenarios, leading to incorrect diagnoses. The variability of problems due to differences in infrastructure, software, and configurations also makes it difficult for AI to effectively generalize solutions.

Troubleshooting often requires a deep understanding of the context in which a problem occurs, including historical data, user actions, and system-specific nuances. AI is completely unarmed against new, unique problems that occur for the first time.

Why does AI fail at log analysis?

Logs generate massive amounts of data with varying structures. AI models may struggle to parse and interpret this data efficiently and accurately. Noise and irrelevant information within logs can overwhelm AI algorithms, especially if they are not designed to filter out extraneous data.

AI is highly dependent on proper training on log data; if the training is not effective, the errors in AI work will be much greater. Contextual relationships between log entries may be missed, leading to inaccurate conclusions.

Why does AI fail in first-level IT support?

First-level IT support involves interactions with end-users, who may describe problems in non-technical language or with varying degrees of detail and clarity. AI systems can struggle to accurately interpret these descriptions and provide relevant support. Even advanced AI chatbots may lack the conversational depth and flexibility needed to adapt to unexpected requests or follow complex dialog threads. The lack of empathy in AI can also negatively impact user satisfaction.

Most of the answers in the “Other” section indicate that either the organization hasn’t implemented AI, or it is still a work in progress, or it is not applicable. Some specific responses mention that AI failed in “writing PWS scripts” (2%). A detailed analysis of the answers in the “Other” section can be found in the Appendix.

CHART 5. Have you encountered errors due to the implementation of AI-based systems that caused a particular failure in the company's operations?

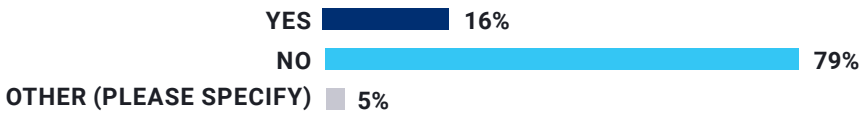

TABLE 3. In which of the listed areas did AI fail and why (please specify in the "Other" field, if applicable)?

Other	64%
Troubleshooting	13%
Providing end-users with first-level IT support	7%
Providing IT staff with guidance and training	6%
Log analysis	6%
Performing post-incident reviews	5%
Managing SSO & passwords	5%
Detecting & remediating incidents	5%
Managing files	4%
Installing & maintaining software	4%
Administering user permissions & administration	4%
Analyzing organization's security controls and comparing to compliance requirements	3%
Defining system usage policies & procedures	3%
Monitoring of server CPU and memory utilization	3%
Patch management processes optimization to improve prioritization, scheduling, and deployment of security updates	3%
Vulnerability prioritization	2%

Key Recommendations

The survey reveals that sysadmins are aware of AI's potential and have identified specific areas where AI can add value. However, the readiness for AI adoption is tempered by the need for further education and training, as well as the current limitations and risks associated with AI technologies. The fact that a substantial majority of organizations do not yet require AI implementation further emphasizes the tentative approach towards widespread AI adoption.

For organizations, the path to successful AI adoption involves:

1. **Investing in AI Literacy and Training Programs:** Equip sysadmins with the necessary knowledge and skills to effectively implement and manage AI solutions. Continuous education will help bridge the gap between AI potential and practical application.

2. **Implementing AI in Low-Risk Areas:** Start with well-defined, low-risk areas to build confidence and demonstrate the value of AI. This approach allows organizations to gain practical experience and refine AI applications without significant risks.

3. **Continuous Evaluation and Adjustment:** Regularly assess AI performance and make necessary adjustments to improve outcomes. This ongoing evaluation ensures that AI applications remain effective and aligned with organizational goals.

4. **Balancing AI and Human Expertise:** Maintain a balanced approach by leveraging human expertise in tasks where AI falls short. This strategy ensures that the strengths of both AI and human capabilities are maximized.

To effectively integrate AI, organizations should follow these detailed steps:

1. **Set Clear Objectives and Use Cases:** Identify specific problems and inefficiencies that AI can address, such as automated incident response, predictive maintenance, or intelligent log analysis. Ensure each use case has clearly defined success metrics.

2. **Engage System Administrators and End-Users:** Involve sysadmins and other end-users in the AI development process. Gather their feedback to design user-friendly interfaces and functionalities. Conduct training sessions and provide documentation to help users understand and effectively utilize AI tools.

3. **Ensure Performance Monitoring and Maintenance:** Implement robust monitoring systems to continuously track the performance of AI applications. Set up alerting mechanisms to detect and address any operational anomalies or failures promptly.

In summary, while there is significant potential for AI adoption among sysadmins, realizing this potential requires a strategic and informed approach. Addressing current limitations, equipping sysadmins with the necessary knowledge and tools, and aligning organizational requirements with the evolving capabilities of AI will pave the way for successful AI integration in IT.

Appendix

Methodology

To compile this report, we collected feedback from 450 sysadmins worldwide from Action1's customer base. Respondents were invited to participate in a giveaway for a chance to win a small monetary reward. The responses were collected in June 2024.

Detailed Responses in the "Other" Field

Below are the responses to the questions in the "Other (please specify)" field. In some questions that received many detailed answers, similar responses were averaged into the corresponding category for ease of information perception. Detailed answers to the question "Other (please specify)" are presented here only if the total number of responses is significant, namely, exceeding 5% of the total number of responses.

1. **Detailed answers to the question: "Do you have an explicit understanding of how to integrate AI into your job to improve the current processes and ensure that your company's IT stays on the cutting edge of the industry trends?" "Other (please specify)" field:**

Sort of, I use some of the available OpenAI Tools (Gemini, Bing, etc.) to help with writing scripts and giving me guidance on weird stuff but I could use more knowledge for sure.
We are aware but currently the tools are not ready
We aren't going to deploy AI
Not so precise, in process of further documentation and establishing most suitable features appropriate for my organization.
I have an understanding. Maybe not "explicit".
It's a WIP
Somewhat
It is in our 2-year plan, but we have not fully implemented yet
Partially. We are testing solutions

Work in progress
The so-called “AI” available is in no way ready for any of that.
Currently investigating different areas to make sure implementation is possible and effective.
Currently AI is the buzzword for marketing and a lot of tools already have automation built in. In many instances adding AI won’t make the department more efficient.
We are seeking ways to improve.
I use AI on an individual basis now. I guess that’s some kind of understanding.
Some understanding, but not explicit.
I could see how I would want to use it but not sure what tools it would take to get it done.
Currently working on this.
I don’t think anyone has.
Reasonable understanding in that it will requires a lot of work to build and implement.
WIP

Additionally, several respondents answered “N/A” to this question.

2. **Detailed answers to the question: “Have you encountered errors due to the implementation of AI-based systems that caused a particular failure in the company’s operations?” “Other (please specify)” field:**

I have had AI provide me with scripts and ‘information’ that were entirely wrong or would have caused significant problems if I didn’t review every single thing AI does for me right now.
Ai sometimes wrong
Some script development no failure in production but caught by manual review some bad code
No, but AI does still give incorrect answers in many instances
I’ve found AI to not be 100% reliable thus far.

Other responses to this question were mainly “N/A” or “Haven’t implemented yet.”

3. **Detailed answers to the question: "In which of the listed areas did AI fail and why?", "Other (please specify)" field:**

- Writing PowerShell scripts and coding – 2%
- Yet to implement/WIP – 7%
- Aren't going to implement/aren't using – 4%
- AI didn't fail – 4%
- None/nothing – 19%
- NA/non-applicable – 23%
- Other detailed responses on AI usage or failures – 5%, see below:

We only really use it for chatting and giving guidance to IT staff
We had issues ingesting data which caused false positives for a while
Verifiably inaccurate or nonsensical results, failure to comprehend want, and confusing intent.
Using ChatGPT to analyze means it often has accuracy issues.
Until AI companies can prove no data will leave my server, it isn't happening.
The users want to speak with a person.
We had issues ingesting data which caused false positives for a while.
Sometimes the AI makes wrong interpretation of the logs, and sometimes due to lack of knowledge of the latest software versions.
Some of the content generated did not apply at all to our org.
Organizational structure was too complex.
Not always right and sometimes faster to just google it.
LLM.
I've only used it for troubleshooting and it's done a remarkably good job.
It failed in troubleshooting due to inaccurate info but that can easily change in the future.
Incorrect code generation, unsafe GDPR practices.
If additional support software is needed in the install, the user will need to be involved in installing those first.
Have only used it for coding to this point.
Email filtering with AI is horrid .
ChatGPT to help learn and troubleshoot more efficiently. No full systems for AI implemented.
Analyzing images and detecting tumors on Pet-CT scans, the AI didn't manage to detect all tumors.
Ai in EndPoint protection, large number of false positives.
AI has never worked with 100% accuracy.
Accuracy was about 70%.

We thank all respondents for their detailed answers and suggestions. We will consider these suggestions in next year's survey and will add the missing answer options to these questions.

Demography

HEADCOUNT

TABLE 1. How many employees does your organization have?

1–100	47%
101–500	34%
501–1,000	9%
1,001–5,000	6%
5,001–10,000	4%

CHART 1. Location

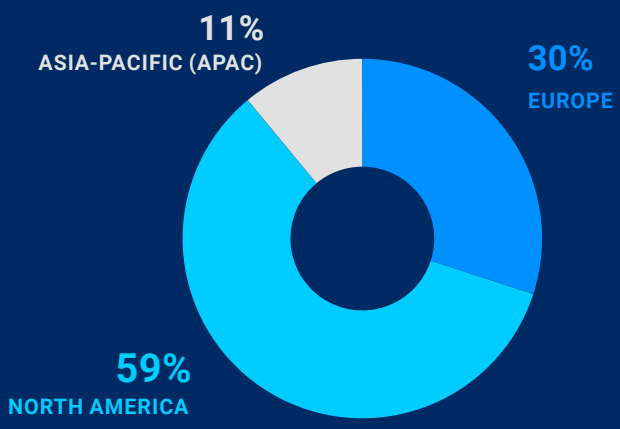

TABLE 2. Industry

Telecommunications, Technology, Internet & Electronics	25%
Education	10%
Manufacturing	9%
Nonprofit	8%
Prefer Not to Answer	7%
Finance & Financial Services	6%
Healthcare & Pharmaceuticals	5%
Construction, Machinery, and Homes	5%
Government	4%
Business Support & Logistics	3%
Retail & Consumer Durables	2%
Agriculture	2%
Airlines & Aerospace (including Defense)	2%
Automotive	2%
Entertainment & Leisure	2%
Food & Beverages	2%
Utilities, Energy, and Extraction	2%
Transportation & Delivery	2%
Real Estate	2%

About Action1 Research

The report is brought to you by Action1 Research, which conducts industry surveys among cybersecurity practitioners worldwide to discover trends in cybersecurity. For more information, please visit:

www.action1.com/resources/research/

About Action1

Action1 reinvents patch management with an infinitely scalable, highly secure, cloud-native platform configurable in 5 minutes—and it just works, with no VPN needed. Featuring unified OS and third-party patching with peer-to-peer patch distribution and integrated real-time vulnerability assessment, it enables autonomous patch compliance that preempts ransomware and security risks, all while eliminating costly routine labor. Trusted by thousands of enterprises managing millions of endpoints globally, Action1 is certified for SOC 2 and ISO 27001.

Action1 was founded by cybersecurity veterans Alex Vovk and Mike Walters, who previously founded Netwrix, which was acquired by TA Associates. Learn more at: www.action1.com.