panther

REPORT

State of AWS Log Management

Insights from 250 Security Professionals who actively use AWS

panther.com

Introduction

AWS, as of April 2022, has ***226 services*** that generate a myriad of different log types and formats. In order to operationalize the vast amounts of data produced by these log entries, security teams must aggregate, normalize and automate based off of the log contents.

We wanted more insight into the current challenges, frustrations, and desires of teams using AWS. To answer these questions, we sought out security professionals who use AWS to better understand what they're seeing, what they're concerned about, and what they want to improve.

Key Findings

48.8% of respondents find it difficult to redirect or copy logs out of AWS into an external log management solution.

Increasing log retention and implementing a SIEM are the highest priorities for the most significant pluralities of respondents.

18.8% of respondents log data from more than 40 accounts, yet over 54.4% say their environments are “very complex,” and 64.8% have “only existed in the cloud.

The survey responses may indicate overconfidence in the capabilities of SIEMs.

Table of Contents

Methodology

On August 31 we surveyed 250 security professionals who use AWS. The survey was conducted online via PollFish using organic sampling through Random Device Engagement (RDE). Learn more about the Pollfish methodology ***here.***

Profile of Who We Surveyed

PART 1

For context about the respondents to our survey, we note all of them work in the security department for an organization that uses AWS. They are chiefly male, but only 12% more than females, with 70.8% between 25 and 44. Those younger or older than the core group are 17.2% and 12%, respectively.

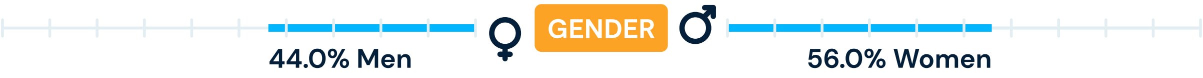

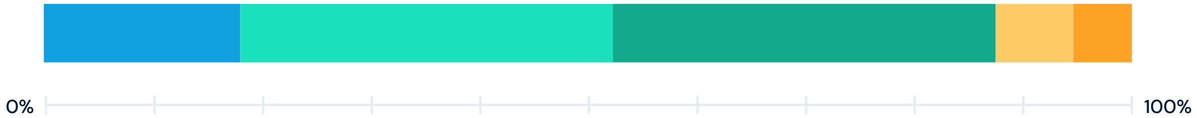

● Age: 18-24	17.2%
● Age: 25-34	35.2%
● Age: 35-44	35.6%
● Age: 45-54	6.8%
● Age: >54	5.2%

100.0% United States

Security Engineers and Analysts together make up 56.4%

What best describes your job title?

● Security engineer	38.0%
● Security analyst	18.4%
● Security consultant	6.4%
● Cybersecurity manager / admin	8.4%
● Network engineer / architect	6.8%
● Network engineer / architect	6.0%
● Software developer	10.0%
● Vulnerability Analyst / Penetration Tester	6.0%

What best describes the size of the company you work for?

Small business	Medium business	Large business	Major Corporation
1-100 Employees	100-999 Employees	1,000-10,000 Emp.	10,000+ Employees
20.4%	28.0%	35.6%	16.0%

How many people are on your security team?

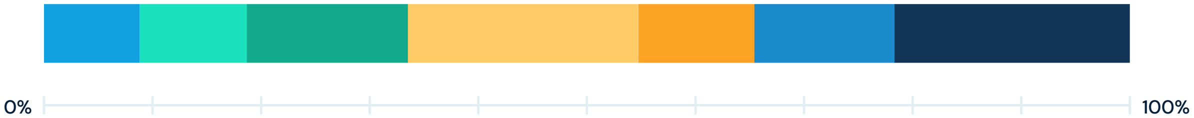

● 1 - 3	9.6%
● 4 - 10	10.0%
● 11 - 20	14.4%
● 21 - 30	21.6%
● 31 - 40	10.8%
● 41 - 50	12.4%
● 51+	21.2%

64.8% of the respondents work for companies that have only ever existed in the cloud

Would you describe the company you work for as a cloud-first organization? *[Cloud-first definition: the company has only ever existed within the cloud from its inception - not a company with a cloud-first strategy.]*

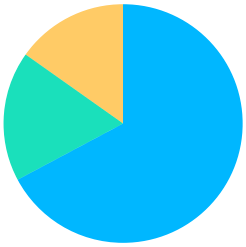

● Yes	64.8%
● No	19.2%
● I don't know	16.0%

State of AWS Log Management

PART 2

While AWS can quickly provide these logs to security practitioners, 60% (Q6) of survey respondents cite that collecting this data from AWS quickly is still a sticking point for their organization.

As the industry continues to grow in complexity of services and organizations continue to move faster with Infrastructure as Code deployments, logging even basic network logs can become a hurdle to overcome as 57% of respondents have found out (Q9). Panther can help to take security log aggregation from a complicated, compliance-driven cost drain to an operationally efficient and near real-time source of valuable information for responding to security events.

One of the top issues for security professionals is collecting and normalizing all the data produced by AWS and then making sense of that information. This survey captures and reports the current challenges, frustrations, and desires regarding solutions security professionals use daily. We shed light on the current state of AWS and logs and what the future may bring.

In this section, we present the findings relative to how the respondents are using AWS logs. We uncover the challenges they face and the strategies they use to overcome them.

54% say the AWS environment is very complex

How would you rate the complexity of your AWS environment?

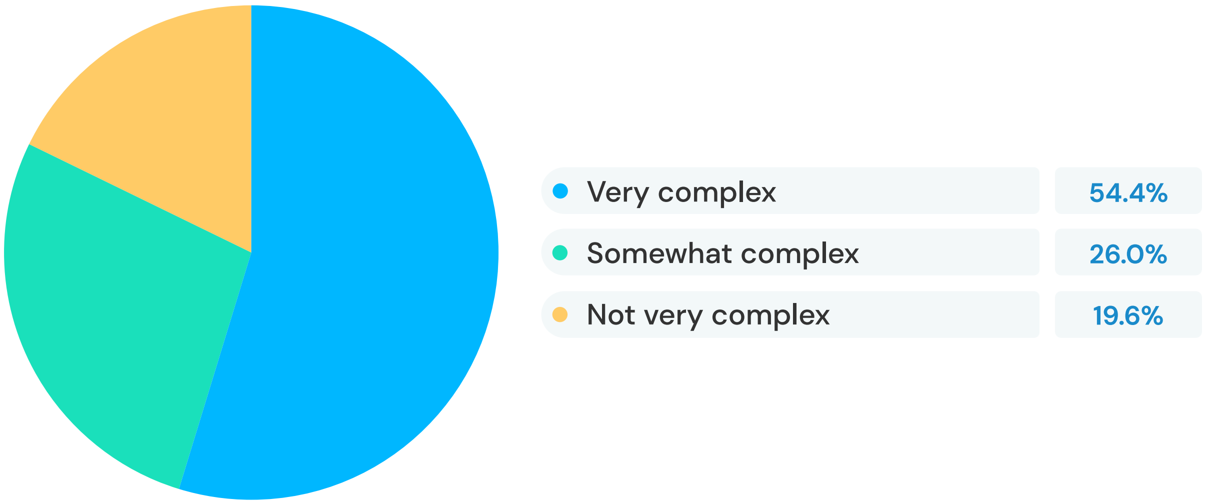

Top challenges

A large plurality (40.4%) of AWS security teams struggle with log correlation from different sources. This challenge underscores the need for SIEM solutions to alert based on events in multiple log streams automatically.

For “top challenges” overall, there is a slant towards just getting the logs into *something* vs. using the logs, and log ingestion must happen before anything else. Our later question about priorities shows that “increasing log retention” and “implementing a SIEM,” two essential but very preliminary steps in the SIEM journey, are the highest priorities for the most significant pluralities of respondents.

What are the top challenges you face with your current log management solution?

17.9%	Collecting large amounts of log data from multiple sources quickly
17.0%	Normalizing and transforming data received in different formats into a common schema
14.3%	Redirecting or copying logs out of AWS into an external log management solution
11.9%	Correlating related incident information from different log sources.
13.5%	Operationalizing the log data to identify security risks or threats
13.4%	Filtering log data or alerts to reduce false positives and benign incidents
11.7%	Scaling capacity for current log management solution to accommodate changing security demands

The largest group (24%) are logging from 21 - 30 AWS accounts

From our previous question about how complex teams feel their environments are, we learned that most consider their cloud infrastructure to be very complex. However, the answers to this question about the number of accounts they log indicate that these teams are logging a relatively small number of accounts. This fact begs the question of where the complexity comes into the environment.

How many different AWS accounts are you logging data from?

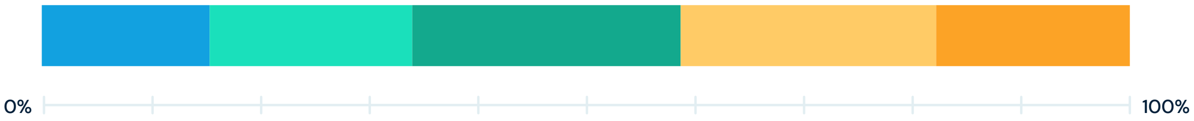

● 1-10	15.2%
● 11-20	19.2%
● 21-30	24.0%
● 31-40	22.8%
● 40+	18.8%

Volume ingested

What volume of logs are you currently ingesting per day?

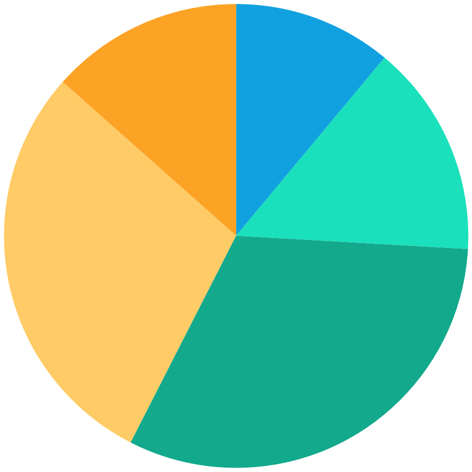

● < 10 GB /day	11.2%
● 10-100 GB/day	16.0%
● 100-300 GB /day	30.0%
● 300 - 500 GB/day	29.2%
● > 500GB /day	13.6%

Challenging log sources

What log sources do you find to be the most challenging to collect?

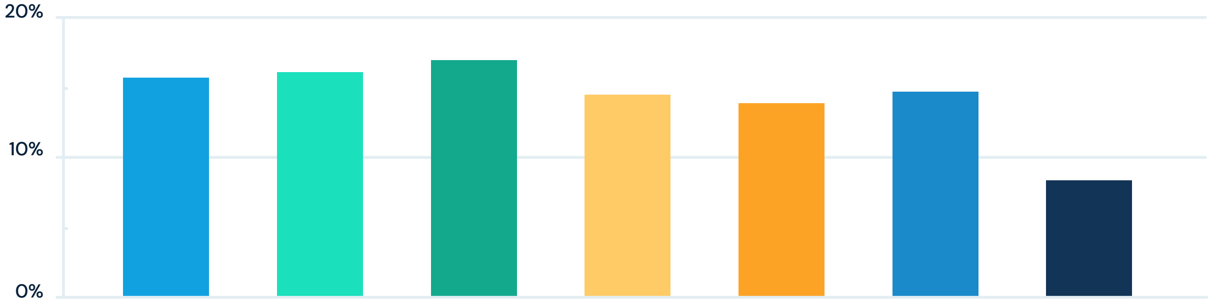

● Storage logs such as S3	15.5%
● Database logs (RDS, DynamoDB, etc.)	15.9%
● Networking logs (VPC, Cloudfront, Route53, etc.)	16.8%
● Container logs (EKS, ECS, Kubernetes, etc)	14.0%
● Host logs (EC2, etc)	13.6%
● Security, Identity, and Compliance logs (IAM, Guard Duty, etc..)	14.1%
● SaaS logs (Slack, Github, etc)	9.8%

Challenge copying logs out of AWS

Have you tried redirecting or copying logs out of AWS into another solution?

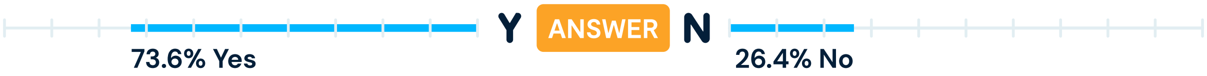

[If yes] were those efforts successful?

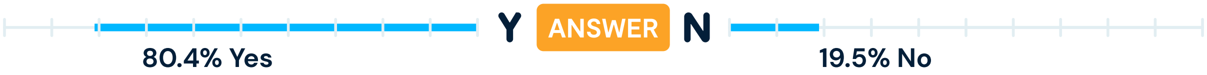

From the questions in this section, we learn that more than half of respondents say AWS is complex. They struggle with log correlation from different sources and getting logs into other tools. The largest group logs from 21 – 31 accounts at a rate of 100 – 300 GB/day.

Current State of Logs

PART 3

To better understand the current state of logs for cloud-first security teams, we asked a series of questions relating to visibility, alerts, and the volume change they are experiencing. We delve into the cost and manageability of log sources as well.

Visibility across accounts and services

Do you have visibility into the activity across all of your AWS services?

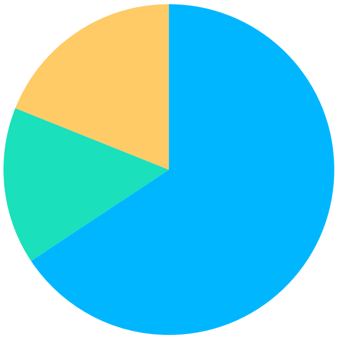

● Yes	65.2%
● No	15.6%
● I don't know	19.2%

Do you have visibility into the activity across all of your AWS accounts?

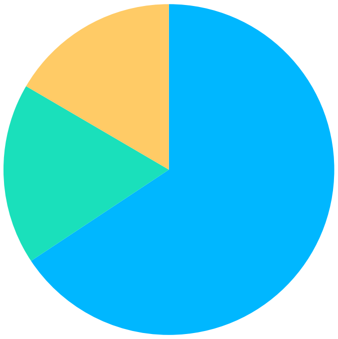

● Yes	64.0%
● No	18.8%
● I don't know	17.2%

Alerts

How many security alerts does your organization detect each day in AWS?

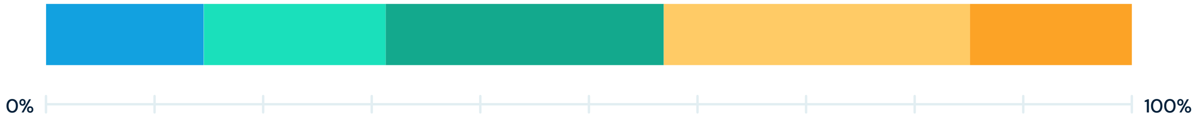

● <10	14.8%
● 10-50	16.4%
● 51-100	25.6%
● 101-250	29.2%
● 251+	14.0%

Incident Volume

How has the volume of security incidents changed over the past 12 months?

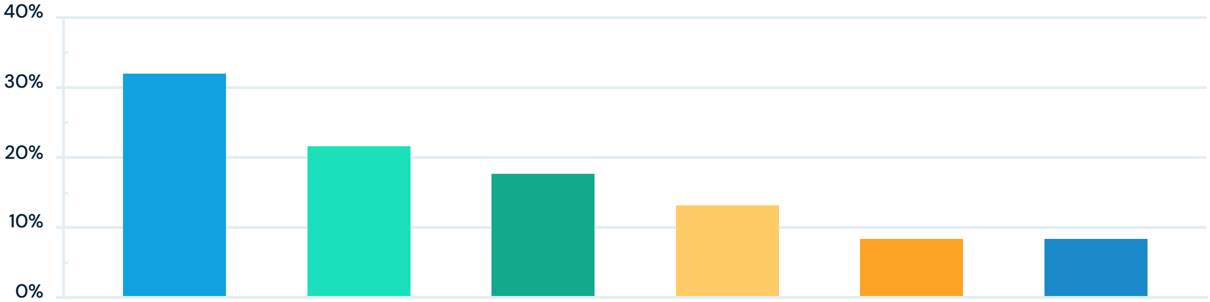

● Increased at a rate of 5x or more	30.4%
● Increases at a rate of 4x	20.4%
● Increases at a rate of 3x	17.5%
● Increases at a rate of 2x	12.8%
● No increase	9.6%
● Decreased	9.6%

Cost of log sources

Are there log sources that are too expensive to store and query for you to manage today?

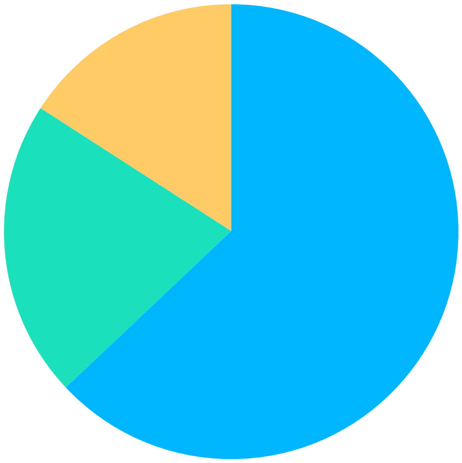

● Yes	63.6%
● No	20.8%
● I don't know	15.6%

Difficult to manage logs

Are there log sources that are too large for you to manage today?

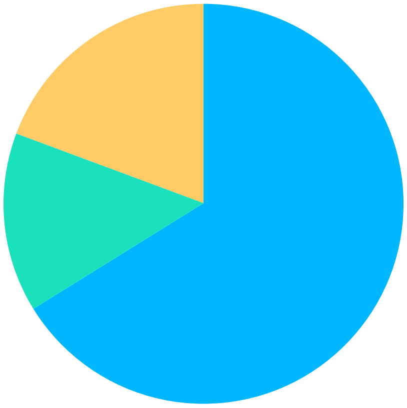

● Yes	65.2%
● No	14.8%
● I don't know	20.0%

[If yes] What log types? (by application)

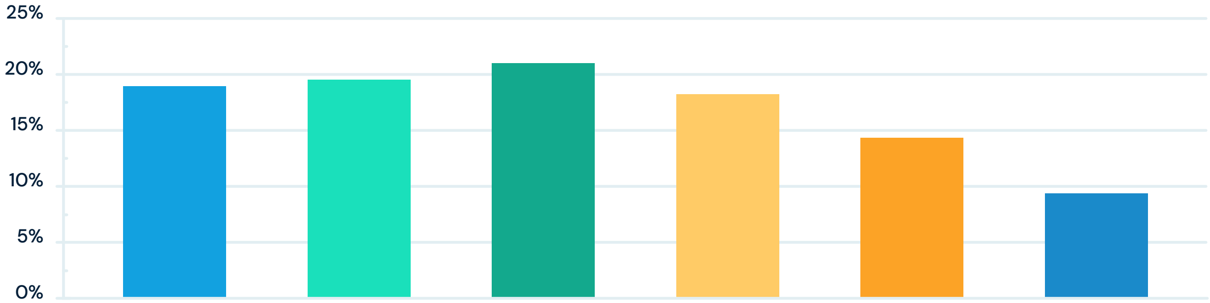

● Storage logs such as S3	18.6%
● Database logs (RDS, DynamoDB, etc.)	19.0%
● Networking logs (VPC, Cloudfront, Route53, etc.)	20.7%
● Container logs (EKS, ECS, Kubernetes, etc)	17.6%
● Host logs (EC2, etc)	14.4%
● Security, Identity, and Compliance logs (IAM, Guard Duty, etc..)	9.5%

This section exposes some reasons why visibility may, or may not, be an indicator of good log management, the number of daily alerts teams handle, and the growing increase in incident volume. Log source cost and size also contribute to the challenges practitioners face.

Technologies & Tools

PART 4

In this section, we turn our attention to monitoring configuration settings, the types of SIEMs our respondents use, and various aspects of how their SIEMs respond to queries.

Monitor configuration settings

Does your current log management solution monitor configuration settings to record and alert on changes to your environment?

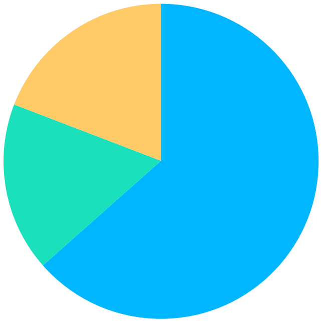

● Yes	64.0%
● No	16.4%
● I don't know	19.6%

Type of SIEM and query return time

What best describes the type of SIEM platform your SOC team uses?

● SaaS (e.g. Sumo Logic, Splunk Cloud)	32.8%
● Commercial On-Prem (e.g. Splunk On-prem)	22.0%
● Open-Source (e.g. Elastic)	11.2%
● Cloud provider solution (e.g. Azure Sentinel, Chronicle)	15.2%
● Custom Solution	11.2%
● We do not currently use a SIEM	7.6%

On average, how long does it take for a query to return with your current SIEM platform if the query is operating on the last 24h of data?

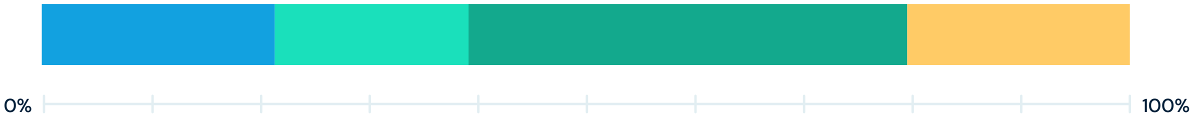

● Less than 1 hour	21.6%
● 1-12 hours	18.6%
● 13- 24 hours	39.3%
● 24 hours+	20.3%

On average, how long does it take for a query to return with your current SIEM platform if the query is operating on the last 7d of data?

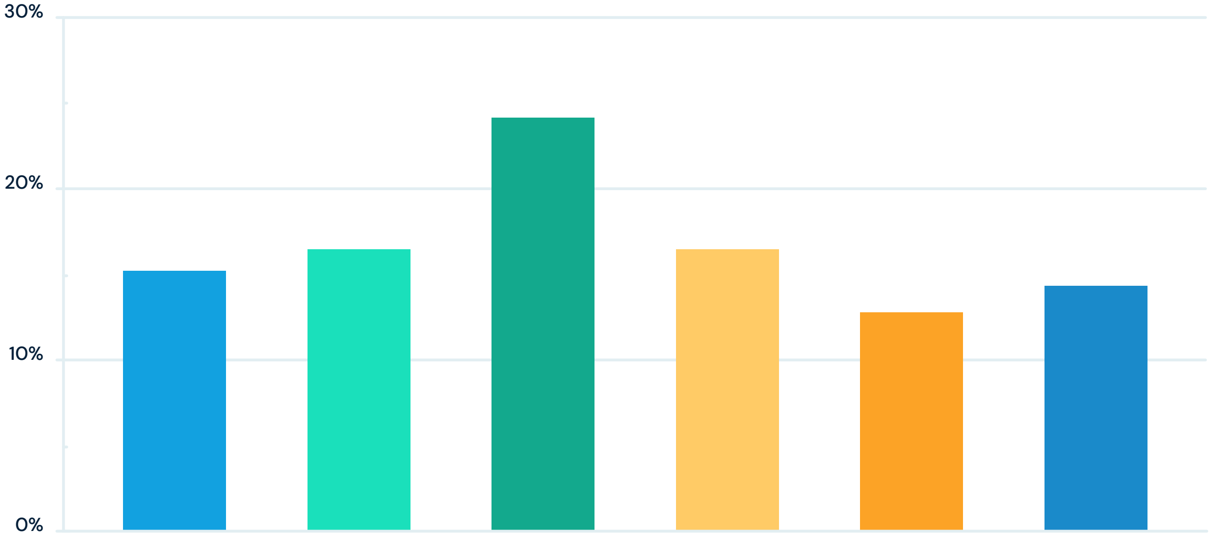

● Less than 1 hour	15.1%
● 1 - 12 hours	16.8%
● 13 - 24 hours	24.2%
● 24-48 hours	16.8%
● 2-4 days	12.5%
● 5 or more days	14.2%

On average, how long does it take for a query to return with your current SIEM platform if the query is operating on the last 30d of data?

● Less than 1 hour	12.5%
● 1 - 12 hours	15.5%
● 13 - 24 hours	20.3%
● 24-48 hours	16.0%
● 2-4 days	10.3%
● 5 days	11.2%
● 11 or more days	13.8%

On average, how long does it take for a query to return with your current SIEM platform if the query is operating on the last 6mo of data?

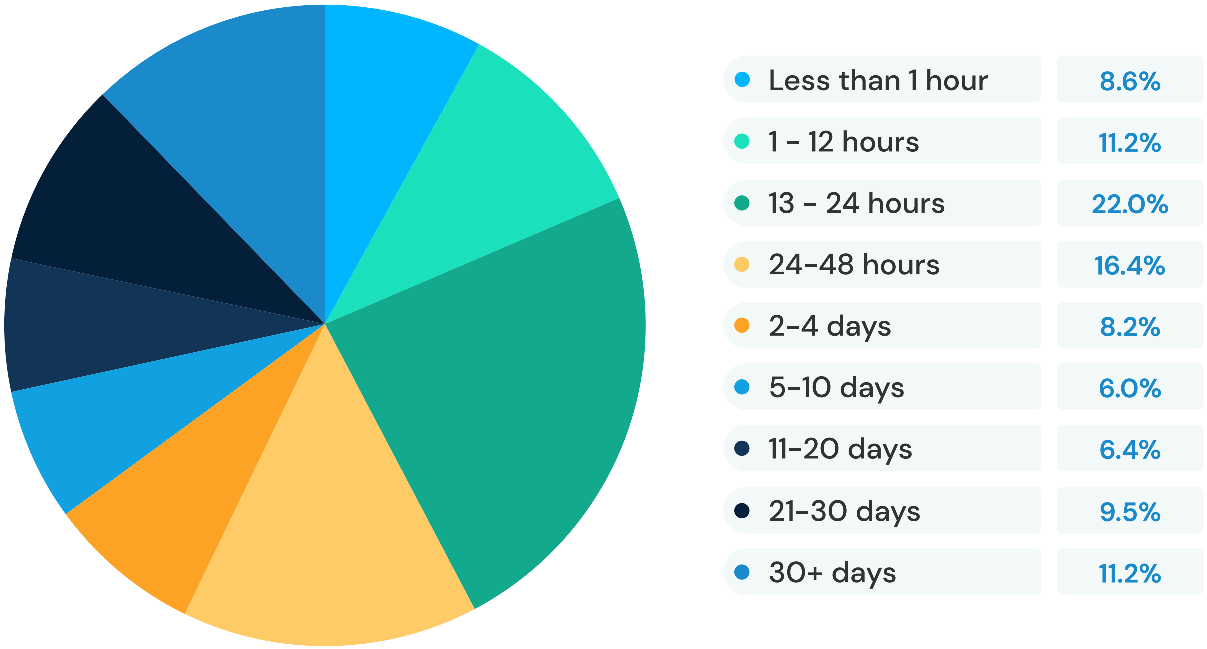

The results in this section can serve as a wake-up call. The SIEM query return times are extraordinary and should be cause for concern in these teams.

Priorities For the Future

PART 5

In this section, we'll take a look at what to expect in the future and what changes teams can make today to be well-positioned in the years to come.

Priorities

Using our respondent’s top priorities for the next 12 months and our guide, in the future, we can expect security teams to increase their log data retention (22.8%), implement a SIEM solution (18.4%), and add behavior detection and alerting (15.2%).

What are your top priorities for the next 12 months ?

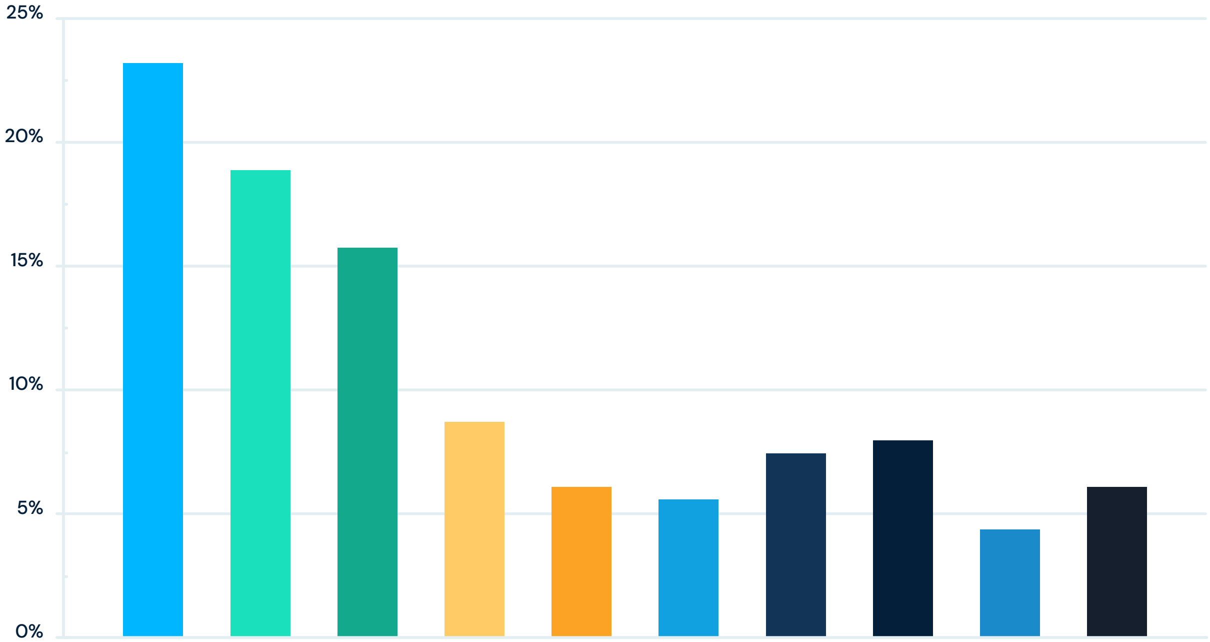

●	Increasing log data retention	22.8%
●	Implementing a SIEM solution	18.4%
●	Incorporating behavioral detection and alerting	15.2%
●	Increasing log retention	8.0%
●	Threat hunting	5.6%
●	Moving to real-time monitoring	5.2%
●	Increasing the number of log sources ingested	7.2%
●	Reducing mean time to detect and mean time to acknowledge	7.6%
●	Adopting a Detection-As-Code methodology	4.4%
●	Increasing log management scale/capacity	5.6%

Tactical changes to make today

Begin with the basics. It’s better to ingest and alert on a small set of logs efficiently and effectively than to ingest and alert a firehose of logs poorly. Define your top logging requirements and get to 80% of perfection with those log sources before moving to others.

Don’t think of your SIEM as a magical tool that does most of your thinking for you. Operating a SIEM is hard work, and it never ends. Changes in your environment and external threats mean that the only constant thing is change.

Segment functions into more AWS accounts. The AWS account logical boundary is an effective one. Affirmative steps have to be taken on both sides of an account boundary to allow any contact from one account to another. Indeed, struggling with creating such pathways is a significant troubleshooting issue.

Pay more attention to de-noising your security alerts. Noisy alerts are either masking important signals or not alerting on the correct data elements at all.

Queries shouldn’t take days! Make queries simpler, pre-process data to make data elements easier to parse out, or move to a database that can do faster distributed searches.

Powered by Panther, teams will leverage the ability to aggregate all AWS security log data into a centralized, normalized single view. Leverage out-of-the-box support for CloudTrail, GuardDuty Alerts, S3 Access logs, Application Load Balancers, VPC Network traffic flow, and other types.

Ingesting and aggregating log data is paramount for organizations with a presence in AWS. However, as products are developed using an increasing amount of services in AWS, each with their own unique log formats, and the amount of cloud accounts and complexity of an organization’s cloud footprint grows, current solutions can start to fall short of business needs. Legacy logging solutions that were not built with an AWS-first mindset continue to prove a pain point for many security teams. Panther’s out-of-the-box support for native AWS services helps to reduce development costs and implementation time for security teams, allowing them to focus on operationalizing their data for faster response times, less alert fatigue and more secure environments.

panther

REPORT

Detect Any Breach, Anywhere.

Detect suspicious activity in real-time, transform raw logs into a robust security data lake, and build a world-class security program with Panther.

Learn More

panther.com