

Survey

SANS 2023 Multicloud Survey: Navigating the Complexities of Multiple Clouds

Written by [Kenneth G. Hartman](#)

December 2023

Executive Summary

According to market data, more businesses than ever before are utilizing several cloud service providers. The first SANS Multicloud Survey, performed in 2022, indicated that the forces behind the tendency to adopt multiple cloud solutions were driven by a variety of factors, including mergers and acquisitions and concerns around ensuring business continuity. It is also clear that the major cloud service providers continue to innovate and differentiate their services in the face of intense competition.

This study takes a fresh look at multicloud adoption trends in the face of a looming recession, tech layoffs, and the excitement about the advancements of AI technologies.

Cloud adoption is on the rise, with more and more organizations using multiple cloud service providers (CSPs). This trend is being driven by several factors, including mergers and acquisitions, concerns around ensuring business continuity, and the desire to take advantage of the best-of-breed services offered by different CSPs.

As organizations adopt multicloud environments, they face a number of challenges, including managing security and compliance across multiple CSPs, ensuring data portability, and optimizing costs. However, there are also benefits to multicloud adoption, such as increased agility, flexibility, and choice.

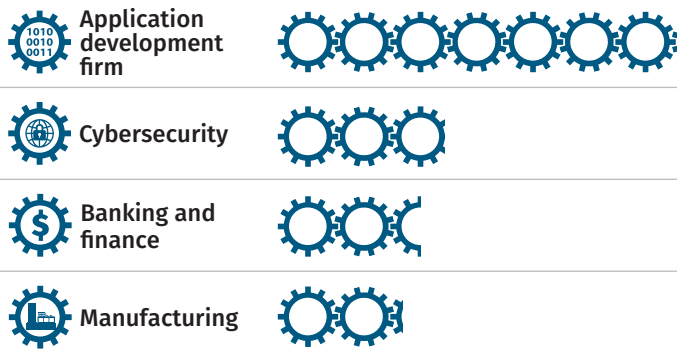
This report provides an overview of the multicloud landscape, including the key trends, challenges, and benefits. It also offers recommendations for organizations that are considering or already using a multicloud strategy.

Figure 1 on the next page provides survey respondent demographics.

Summary of Key Takeaways

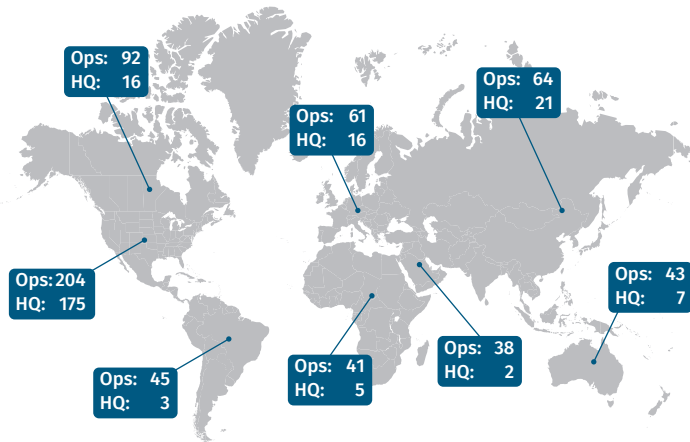
- **Security teams need to be prepared to secure cloud environments that are hosted on multiple cloud providers.**
- **Organizations should establish a centralized cloud security governance framework to manage cloud security across multiple cloud environments.**
- **Cost optimization is important, but it should not come at the expense of security.**
- **Organizations should carefully consider the security features and availability of a cloud service provider before selecting one.**
- **Security teams need to be able to deal with multiple SSO solutions.**
- **Organizations should consider using cloud access security broker (CASB) and secure access service edge (SASE) technologies to solve access control problems.**
- **Cloud security technologies can play an important role in increasing the security of a cloud environment, but each has a cost that needs to be considered.**
- **Organizations should use a SIEM solution to detect and respond to cloud security events.**
- **Organizations should have visibility into DNS traffic across all device categories.**
- **Organizations should use tools and practices to police cloud accounts as they grow.**
- **Organizations should ensure that developers know how to write secure cloud applications.**
- **Organizations should leverage third-party expertise to make sure their security team is aware of the evolving threats that face their cloud applications.**

Top 4 Industries Represented




Each gear represents 10 respondents.

Operations and Headquarters

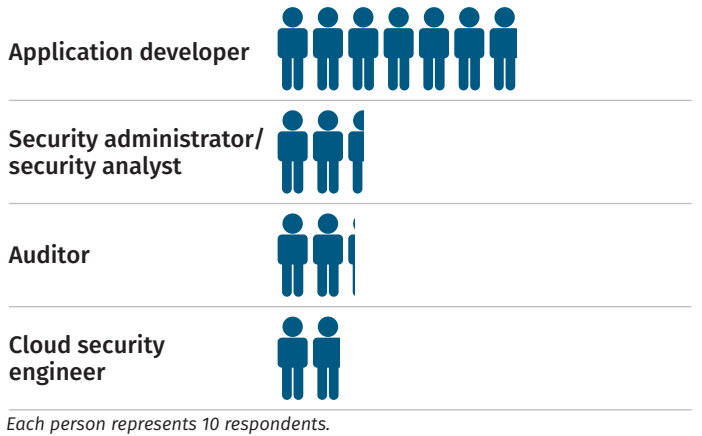


Organizational Size



Each building represents 10 respondents.

Top 4 Roles Represented



Each person represents 10 respondents.

Figure 1. Survey Participants, Demographic Data

Multicloud Operations

This survey targeted cloud users whose organization uses more than one cloud service provider. Of our 245 total respondents, 86% are using multiple cloud service providers, resulting in a survey population going forward of 210.

Amazon Web Services is still the leader based on the percentage of workloads running in AWS relative to the other cloud service providers used by each respondent's organization. (See Figure 2.)

What percentage of your workloads is running in the following clouds?




Figure 2. Percentage of Workloads Running in Each Cloud

Interestingly, 3% said that they are using AWS exclusively, while Azure and Google Cloud are used exclusively by only 1% each. Alibaba and Oracle are used exclusively by 2% each, while IBM Cloud was used exclusively by 1%. On the other end of the spectrum, 12% of respondents said they were not using AWS for compute workloads at all. It is also worth noting that 35% of respondents indicated they were not running any compute workloads in a private cloud, whereas more than 28% are using it for at least one fourth of their compute workloads.

KEY TAKEAWAY

More than 28% of respondents are using private cloud for at least one fourth of their compute workloads. Cloud security programs should include that focus as part of their security strategies.

According to Figure 3, Google Cloud Platform (GCP) tied with Azure as the cloud service provider used by the most respondents (118 or 63%), followed closely by AWS (54%). Alibaba is also very popular at 43%.

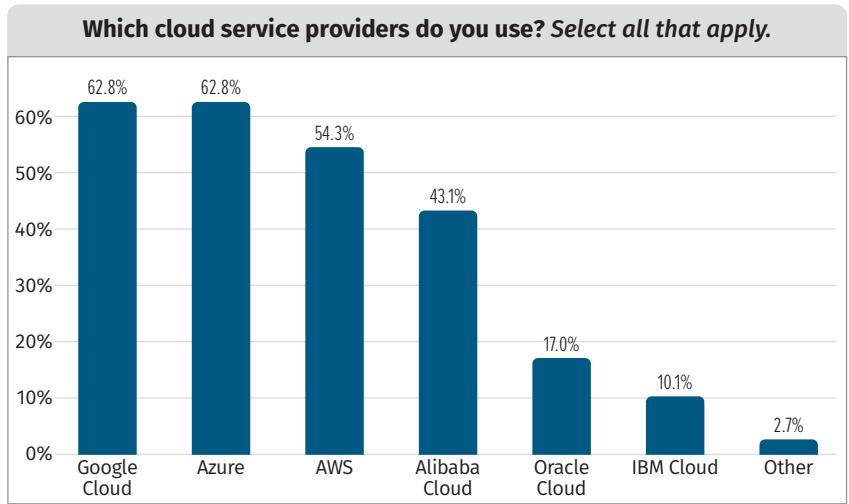


Figure 3. Cloud Service Providers in Use

KEY TAKEAWAY

In this survey, Google Cloud and Azure were the cloud service providers used by the most respondents, followed closely by AWS. Cloud security practitioners must be adept on all platforms that their organization uses.

GCP is the most likely CSP to be inherited as the result of a merger or acquisition, according to 47% of respondents. Considering that 63% of respondents said that they are currently using Google Cloud (see Figure 3), mergers and acquisitions (M&A) seems to be the most likely reason that most of these survey respondents started to use Google Cloud. This seems to be less the case for Azure, where 37% started using the platform because of a merger or acquisition, while 63% said they are currently using the platform. One possible explanation of this is that Microsoft has been known to leverage its Microsoft 365 solution to gain a foothold from which to encourage the adoption of Azure. This is represented graphically in Figure 4.




Figure 4. Organizations Using a CSP Because of M&A

Overall, based on the results displayed in Figure 4, it is clear that organizations must frequently adapt to a new cloud service provider because of either a merger or an acquisition. Considering this data, security organizations must be poised to promptly onboard new technologies in a secure manner. A single platform and unified product approach across all CSPs is highly recommended to help organization IT and security teams handle the complexity and variation they need to face with heterogeneous IT operations.

KEY TAKEAWAY

Security organizations that undergo an acquisition or a merger must be prepared to swiftly and securely integrate new cloud accounts, possibly with unfamiliar cloud service providers.

Of the 73% of respondents using multicloud (153/210) that answered our question, we see in Figure 5 that close to half have encountered improperly procured providers for various reasons, including shadow cloud accounts.

Shadow cloud accounts are a security risk because it is unlikely that they will be configured to meet the organization's security requirements. For the most part, the likelihood of identifying a shadow cloud account for a particular cloud service provider closely mirrors the likelihood of a respondent using that cloud service provider, with the notable exception of Azure and Google Cloud. This may be in part because an organization's Microsoft 365 accounts can be restricted from using cloud services that are not governed by the organization's policies, like email addresses associated with an organization's Google Workspace. This is not the case with a new AWS account because any email address can be used.




Figure 5. Organizations Encountering Improperly Procured Cloud Accounts

KEY TAKEAWAY

All organizations must be on the lookout for shadow cloud accounts and must take corrective action to either close these accounts or bring them under the organization's cloud security governance.

Cloud-Agnostic Workloads

Multicloud is a good option for product development organizations because it gives you more flexibility and choice. You can choose the best services from different providers, based on your specific needs and budget. This can help you avoid vendor lock-in and save money in the long run. Many companies have made their applications cloud-agnostic. Here is what we found:

- 88% of respondents said that it was at least somewhat important that their organization’s applications be cloud-agnostic; of that group, 13% indicated there was a company imperative.
- 82% of respondents said that their organization will run identical workloads in multiple clouds “often” or “sometimes.”
- 67% of survey respondents said that some or all of their applications are cloud-agnostic (see Figure 6). Given that 88% of the respondents said that cloud-agnosticism was at least somewhat important, we should expect this trend toward creating cloud-agnostic applications to continue.

A cloud-agnostic workload can run on any cloud service provider’s platform in a portable manner without requiring refactoring. A cloud-agnostic workload can be implemented by one or more virtual machines but is increasingly implemented using containers.

Does the pursuit of complete cloud-agnosticism make sense? Certainly, that is debatable. Although it may appear to be a good procurement strategy to consider cloud services as a generic commodity, the majority of cloud users do not merely utilize cloud infrastructure-as-a-service. Instead, they employ a combination of cloud services. As a workaround for this, they are integrating with cloud-managed services that have APIs and capabilities that are distinct from those of the competing providers.

Establishing a centralized cloud security governance framework is the single most critical security tip for a company that is contemplating becoming cloud-agnostic. In order to effectively manage cloud security across different cloud environments, this framework ought to include a detailed explanation of policies, procedures, and responsibilities. Additionally, it should include tools and technology that enable consistent monitoring, threat detection, and incident response activities across all cloud platforms.

KEY TAKEAWAY

The single most important security recommendation for a company considering becoming cloud-agnostic is to establish a centralized cloud security governance framework. This framework should outline clear policies, procedures, and responsibilities for managing cloud security across multiple cloud environments. It should also encompass tools and technologies that enable consistent monitoring, threat detection, and incident response across all cloud platforms.




Figure 6. Cloud-Agnostic Applications

Cost Optimization

When a workload is cloud-agnostic and portable, it is possible to migrate it to a different cloud service provider in an effective manner to take advantage of pricing changes. However, even if the transition to a new cloud provider requires a large amount of engineering effort, an organization may still be persuaded to make the switch if the potential financial gain is significant enough. We explored this by asking, “How likely are you to switch the cloud provider you are using for a particular service based on cost alone? For example, would you ever switch from Amazon EC2 to Azure VMs because of price alone?” Figure 7 shows that as the cost savings increase, an organization is more likely to switch service providers.

Despite the efforts of cloud service providers to differentiate their services, compute services are often envisioned as an interchangeable commodity. This can explain the desire for cloud-agnostic workloads. However, when cloud customers contemplate the effort that it would take to switch cloud service providers, they may balk.

Just over 18% of the respondents said that their organization makes real-time decisions regarding which cloud service provider to use based on current pricing, and another 18% said that their organization would never switch cloud service providers based on price alone. Some of the responses had commentary stating that they would factor in cost savings as a consideration but would also take into consideration the cost of making a change, how much time it would take to make the change, as well as other business factors.

Over 65% of the respondents indicated that they use a service to optimize costs based on where their virtual machines or containers would run. We asked them to provide the names of the services that they used, and were provided with the following list:

- Aptio Cloudability
- Cloudamize
- CloudBolt
- Google Cloud Billing
- IBM Cloud Transformation Advisor
- Kubecost
- Morpheus Cloud Management
- Opsani (which is now part of Cisco AppDynamics)
- ParkMyCloud
- RightScale Cloud Management
- SkyKick Cloud Manager
- VMware vRealize Business for Cloud

SANS is not making any recommendations regarding these companies, but generally speaking, the above-listed solutions offer visibility, governance, and security features in addition to cost optimization. Cost optimization efforts are to be applauded because there are generally security benefits related to eliminating compute and cloud storage that is no longer needed by the organization. Unmanaged systems are rarely secure.

How likely are you to switch the cloud provider you are using for a particular service based on cost alone? For example, would you ever switch from Amazon EC2 to Azure VMs because of price alone?




Figure 7. Likelihood to Switch Providers Based on Cost Savings

KEY TAKEAWAY

Cost optimization is commendable because removing unnecessary compute and cloud storage improves security. Cutting expenditures on security services is counterproductive.

Cloud service providers offer different pricing strategies for virtual machines based on different usage patterns. These include:

- **Reserved instances**—Customers that know they will need a virtual machine for a full year or more can commit to a specific term and receive a significant discount compared with an on-demand pricing plan.
- **Spot instances**—A spot instance is a type of cloud computing instance that is available at a discounted price. Spot instances are created when there is excess capacity in the cloud. Cloud providers offer spot instances to customers who are willing to have their instances terminated if capacity is needed for other customers. Spot instances can be a great way to save money on cloud computing costs, especially for batch jobs and workloads that are fault-tolerant and can be interrupted.
- **On-demand instances**—On-demand pricing is a cloud pricing model where customers pay for the resources they use, as they use them. There are no long-term commitments or upfront payments. Customers are billed by the hour or second, depending on the service they are using. On-demand pricing is the most flexible cloud pricing model, and it is a good option for workloads that have unpredictable or fluctuating demand.

Almost half (46%) of the respondents said they are using on-demand instances, but a large percentage (69%) were not sure which pricing plan their organization used. Reserved instances may be an opportunity for cost savings but could present the organization with a serious security risk if they are not maintained in a fully patched state. Reserved instances must be patched in place, whereas short-lived on-demand instances can be launched from a fully patched image that is built using continuous integration/continuous delivery (CI/CD) automation and rotated out frequently.

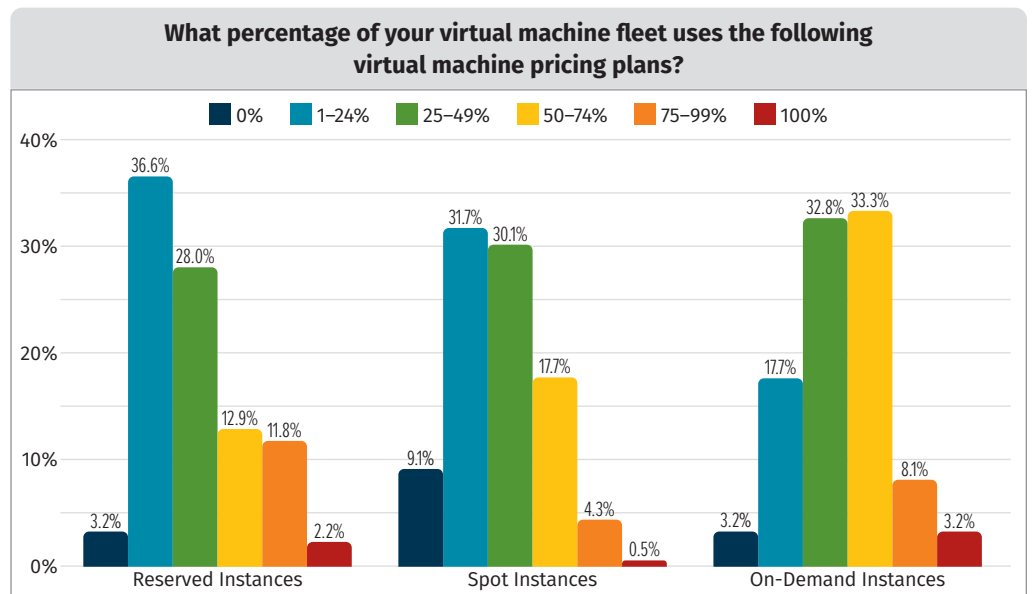


Figure 8. Virtual Machine Pricing Plans

KEY TAKEAWAY

Balance the costs of patching-in-place with the savings available with reserved instances pricing plans. There may be significant benefits to patching in the build process and using short-lived virtual machines.

The most respondents (43%) indicated that security features are very important, followed closely by availability (39% of respondents). Negotiating position was rated “important” by 52% of respondents and “somewhat important” by another 26% of respondents. Respondents were divided on how to rate “unique service offerings,” with 29% saying they were very important, 41% saying they were important, 24% saying they were somewhat important, and another 6% saying they were not important.

KEY TAKEAWAY

Security features and availability are the most important aspects of a cloud service provider’s service delivery. Negotiating position, developer preference, and unique service offerings are also important to consider when selecting a CSP.

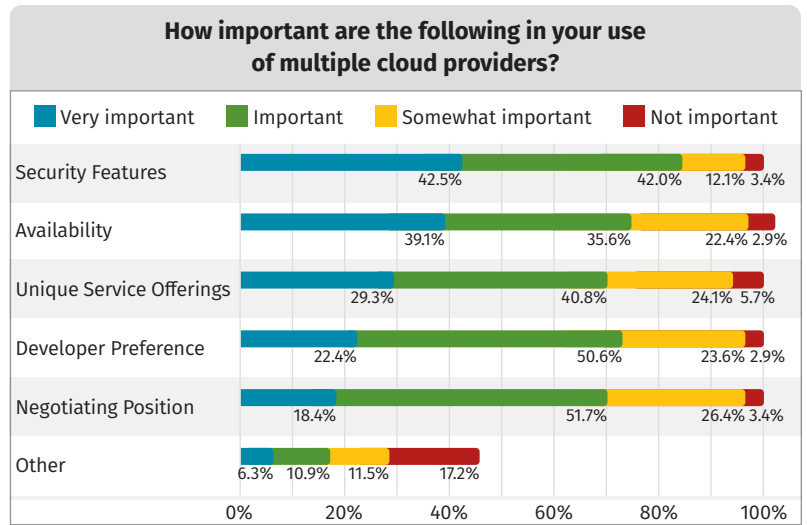


Figure 9. Important Features of Cloud Service Providers

Executive management has the most influence on the decision to use a specific cloud service. This audience may be particularly interested in both the availability and negotiating position factors in choosing to go multicloud. The information technology team and the security team, respectively, were deemed to have the next most influence on the decision to use specific cloud services.

KEY TAKEAWAY

Executive management has the most influence on the decision to use a specific cloud service, but other departments wield important influence as well.

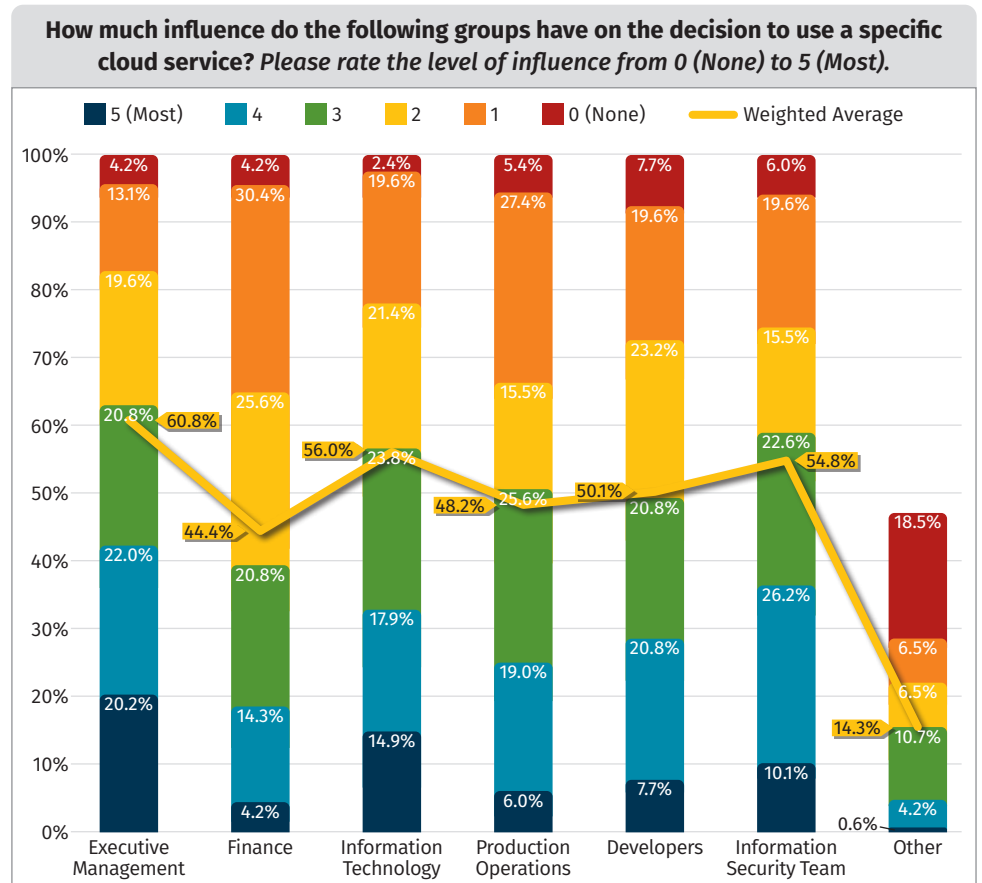


Figure 10. Influence of Specific Groups on Cloud Usage

We asked, “How long does it take for your organization to make decisions about which cloud service will be used?” The possible answers were days, weeks, months, quarters, 1 to 2 years, and more than 2 years. About one fourth (26%) stated that decisions are made in a matter of weeks or less, while 65% said that the decision would be made in a matter of months. Over 87% said that it was quarters, in other words, less than a year. Overall, it seems that decisions regarding cloud services are made in a collaborative manner in a relatively short period of time.

Implementation takes a little bit longer than making the decision, but still 81% said that the cloud services it decided to use are implemented within a matter of quarters or less. Approximately 67% of respondents said that the decision on which services to implement is completed within “months.”

KEY TAKEAWAY

Most cloud users make and act on decisions quickly. This places a huge strain on the information security team to be involved in cloud systems development to guarantee security requirements are considered during decision-making and deployment.

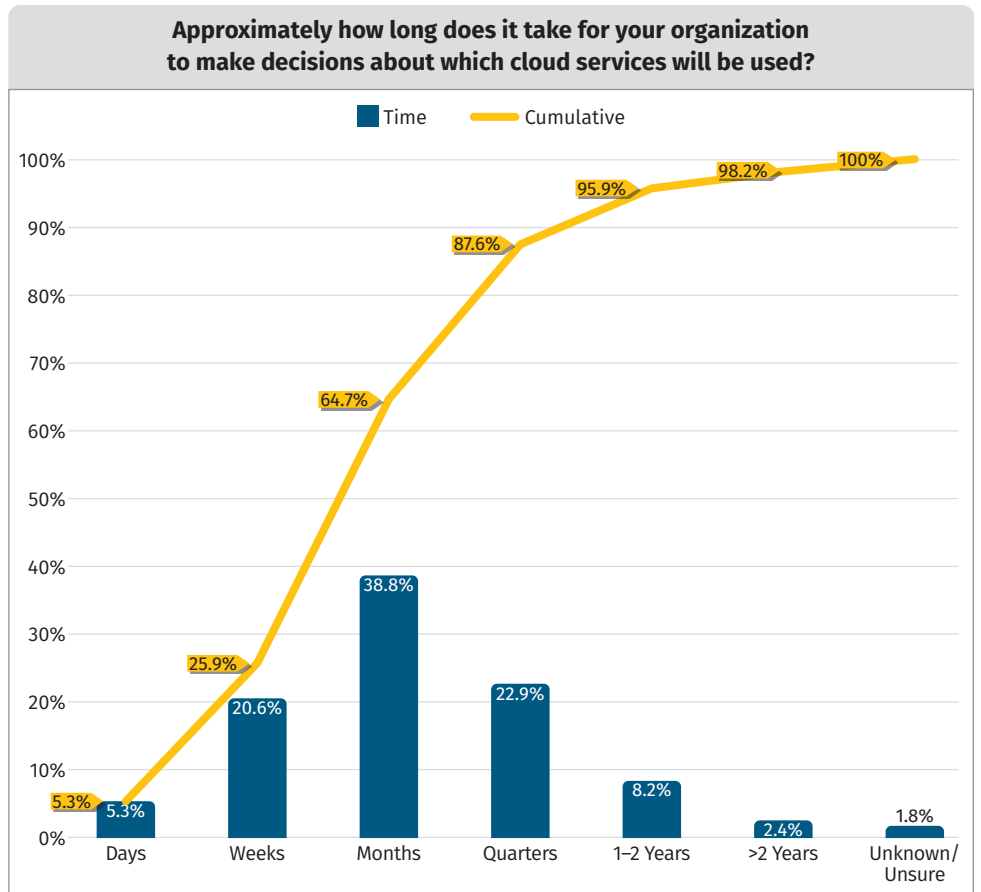


Figure 11. Time to Make Decisions

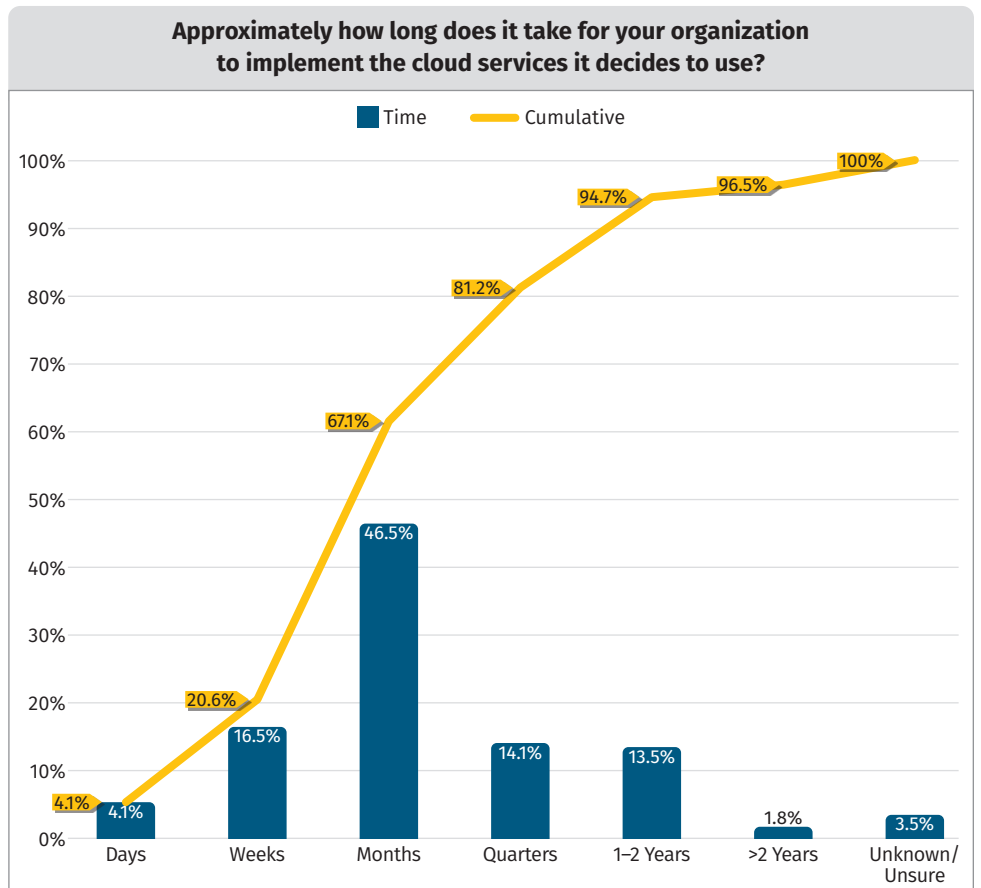


Figure 12. Time to Implement Decisions

Multicloud Access Control

During the first-ever SANS Multicloud Survey in 2022, we were surprised to discover that many of the respondents had replied to a question about which single sign-on (SSO) service provider their organization uses by writing in multiple SSO providers. This year, we decided to explore that with additional questions. Roughly half (55%) of respondents indicated that their organization uses multiple SSO providers, while another 11% do not, although they may use a single SSO provider. (See Figure 13.)

When we asked the participants why their organizations use multiple SSO services, allowing for more than one answer, we saw that the most frequent reason was because of mergers and acquisitions (55%). The next two reasons were: different teams support different services (47%) and a lack of a central authority mandating a single solution (48%). Perhaps the most interesting result is that only 14% said that the organization was working toward a single SSO solution. This suggests that there may be an acceptance that using more than one SSO provider is either good enough, seen as normal, or overshadowed by higher priorities.

KEY TAKEAWAY

Although a single SSO solution is preferred in an ideal world, security teams must be prepared to deal with multiple SSO solutions. This includes making sure there are no gaps and weaknesses in how users authenticate to corporate and cloud systems.

Table 1 lists the SSO providers in use. Microsoft Entra ID (formerly Azure Active Directory) is the most used SSO solution, followed by Okta.

Several businesses have begun utilizing cloud access security brokers (CASBs) in order to exercise control over the internet and cloud services that their users have access to. Whether they are on the company network or on a remote network, bring your own device (BYOD) and corporate-controlled devices can both have their connections handled by a CASB that has been well-architected. They are an effective defense mechanism against malicious software and shadow cloud accounts. Additional cloud network security services, such as zero trust, DNS filtering, and firewall-as-a-service, are layered on top to comprise a secure access service edge.

Does your organization use multiple SSO service providers?

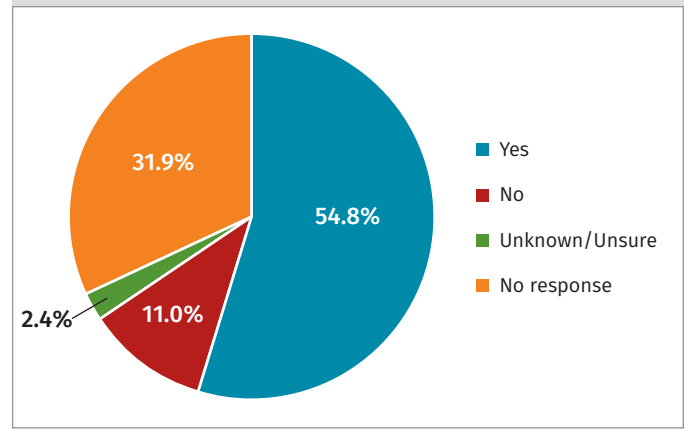


Figure 13. Percentage of Organizations Using Multiple SSO Providers

Table 1. SSO Providers in Use

SSO Providers In Use	%
Microsoft Entra ID	40.3%
Okta	35.1%
IBM Identity and Access Mgmt	29.9%
One Login	25.4%
Google Cloud Identity	23.9%
PingIdentity	20.1%
Ipsidy	15.7%
AWS Identity Center	10.4%
Ubisecure	8.2%
Oracle Identity Cloud Service	6.7%
Other	2.2%

Large organizations face an enormous challenge when it comes to protecting and securing their global-scale cloud footprint. One possible solution is to use either a CASB or a secure access service edge (SASE). Of the 79% of overall survey respondents that answered this question, only 7% said that their organization is using the technology to control all access; another 25% said their organization is using it to control some access. Refer to Figure 14. When we consider the large percentage of respondents that use multiple SSO solutions, it is not surprising to see these results for CASB and SASE.

KEY TAKEAWAY

Organizations are divided on their adoption of CASB and SASE. These technologies can solve some access control problems and deserve consideration.

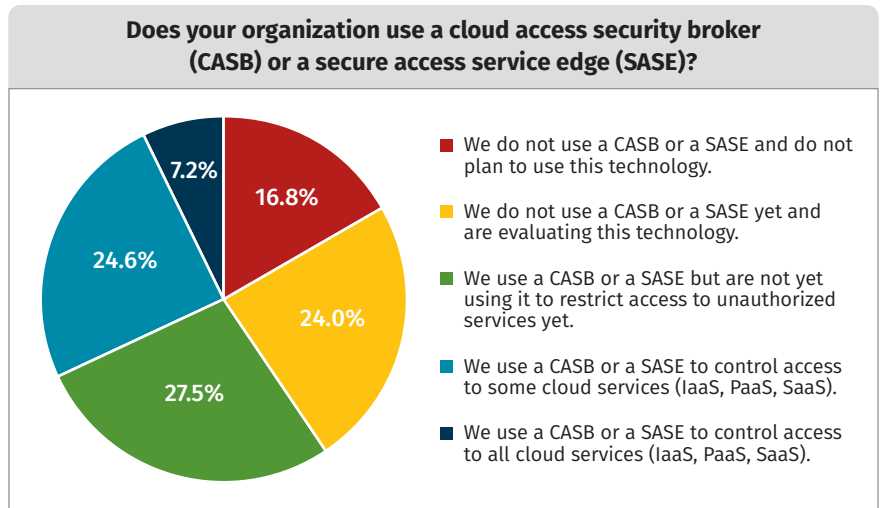


Figure 14. Percentage of Organizations Using CASB or SASE

Cloud Access Security Broker (CASB)

A CASB is a security solution that helps organizations manage and secure their access to cloud-based applications and services. CASBs sit between users and cloud applications, inspecting all traffic to and from the cloud to identify and mitigate security risks.

CASBs can provide a variety of security features, including:

- **Visibility**—CASBs provide organizations with visibility into their cloud usage, including which applications are being used, who is using them, and what data is being accessed.
- **Control**—CASBs allow organizations to control access to cloud applications and services, including who can access them, when they can access them, and what devices they can access them from.
- **Data security**—CASBs can help organizations protect their data in the cloud by encrypting data at rest and in transit, preventing data leakage, and monitoring for data loss and theft.

Secure Access Service Edge (SASE)

SASE is a cybersecurity architecture that combines network security and security services into a single cloud-delivered solution. SASE is designed to protect users, devices, and data regardless of location or application.

SASE converges the following capabilities into a single solution:

- **Secure web gateway (SWG)**—SWGs filter and block malicious web traffic, including malware, phishing attacks, and ransomware.
- **CASB**—CASBs provide visibility into and control over cloud usage, as well as data protection in the cloud.
- **Zero trust network access (ZTNA)**—ZTNA verifies the identity of users and devices before granting access to applications and resources.
- **Software-defined wide area network (SD-WAN)**—SD-WAN provides secure and reliable network connectivity between users, devices, and applications.

Cloud Technologies

Cloud service providers offer a variety of services to help cloud customers secure their environments. Some of these technologies are summarized in Table 2.

We were curious to what extent these technologies are being used. These results are displayed in Figure 15. The possible responses were never, sometimes, often, frequently, and extensively. As can be seen, the results are distributed within a range of 17% to 33% for cloud VPNs. VPNs are not appropriate for all situations and do incur a usage cost. This can explain why this category has the highest never-used percentage. It would behoove the cloud security practitioner to understand when, where, and why cloud VPN services are used.

On the other hand, private endpoints offer a significant security benefit that is just starting to be appreciated by cloud customers. Private endpoints allow for improved network access control rules and more restrictive identity and access management (IAM) policies by constraining authorized traffic to stay within the virtual private cloud (VPC). Although there may be incremental costs for using private endpoints, in many cases the improved security benefit is worth considering.

Flow logs allow cloud users to monitor traffic within a cloud network; collection of these logs for network forensics is recommended in situations where data security is paramount. Although flow logs capture just traffic metadata, they can still be used to identify data exfiltration, C&C traffic, network scanning, and other network attacks. Readers are cautioned to not rely on flow log monitoring for real-time detection because there can be significant delays when waiting for the measured traffic to show up in the log. Nonetheless, flow logs they have value from a network forensics perspective. Although flow logs are a summary of traffic, they can still accumulate rapidly, incurring a storage cost.

Because many of the cloud service providers' detective controls are provided on a "best effort basis" and may take time to reach a state of "eventual consistency," SANS recommends using preventive controls wherever possible. A good example of this is using organizational policies to technically enforce the entity's security policy.

Table 2. Common Cloud Security Technologies

Cloud Technology	Definition
Organizational policies	Organizational policies are technical controls applied at the top of the organization's cloud account to prevent configurations or actions taken in the cloud that are contrary to security policy. In AWS, these are <i>Service Control Policies</i> ; in Azure, they are known as <i>Azure Policies</i> ; and in Google Cloud, these are called <i>Organization Policy Constraints and Deny Policies</i> .
Flow logging	Flow logs summarize the connection between two systems in a cloud network within a capture window or aggregation interval. A flow log includes the source and destination IP address, source and destination port, protocol number, bytes transmitted, packets transmitted, timestamps, and additional optional metadata.
Private endpoints	A private service endpoint provides a private IP address from the range of private IP addresses allocated to a virtual network. This effectively brings the service endpoint into the virtual network so that it can be connected locally. Thus, traffic from systems on the virtual network to the service does not need to traverse the public internet and can be constrained to the virtual network.
Cloud VPN	Cloud service providers permit an IPSec VPN tunnel to be configured from their virtual networks to other remote networks, whether virtual or on-premises.
Cloud encryption keys	Customers of cloud services can manage their customer master keys (CMKs) in the cloud. Having complete control of the key life cycle means that a cloud customer can immediately disable all access to the data encrypted by that CMK.
Web application firewalls	A web application firewall (WAF) hosted in the cloud is priced on a pay-as-you-go model that factors in the number of rules in use along with the traffic that flows through the WAF. This makes it an attractive solution relative to an on-premises device because of its scalability.

KEY TAKEAWAY

Cloud security technologies like organizational policies, flow logging, private endpoints, cloud VPN, cloud encryption keys, and WAF can play an important role in increasing the security of a cloud environment; however, each has a cost that needs to be considered as well.

To what extent are you using the following security capabilities provided by your cloud providers?




Figure 15. Adoption of Various CSP Security Capabilities

It may make sense for a business with a single cloud service provider to utilize the specialized cloud security services provided by one of the main CSPs; however, this may not be the case for a business that utilizes multiple cloud providers. It may also be contingent upon the specific specialized cloud service in question. For instance, an organization may find it appropriate to utilize the security and compliance service offered by each of the cloud service providers it employs, while still relying on a third-party logging solution. Further illustrating this point, an organization that utilizes AWS, Azure, and Google Cloud may opt to implement all three of the following:

- AWS Security Hub
- Microsoft Defender for Cloud
- Google Cloud Security Command Center

To what extent are you using the following security services and providers for logging?




Figure 16. Adoption of Logging and Security Services

To aggregate its security logs, the same organization may utilize Splunk or another third-party solution as opposed to evaluating the unique cloud logging solutions provided by each cloud service provider. For those using cloud services to address logging, we see AWS CloudWatch is most popular (see Figure 16).

KEY TAKEAWAY

Regardless of which cloud logging services are used, it is imperative that adequate security logs are collected and monitored. The effort is compounded in a multicloud organization, but it is just as critical.

And for those using cloud services for security and compliance, we see that Microsoft Defender for Cloud appears more popular than AWS Security Hub (see Figure 17).

Intentionality is our recommendation regarding the utilization of cloud-specific security services. Assess the costs, benefits, assets, and limitations of each cloud-specific security service offered by each provider of cloud services you utilize.

Should a single service, such as Microsoft Defender for Cloud, be employed for all other cloud services (including AWS and Google Cloud) as well, or must the cloud-specific security service be utilized for each CSP utilized? Is there an alternative solution provided by a third party that addresses the constraints and overall cost of ownership associated with the utilization of numerous cloud-specific security services?

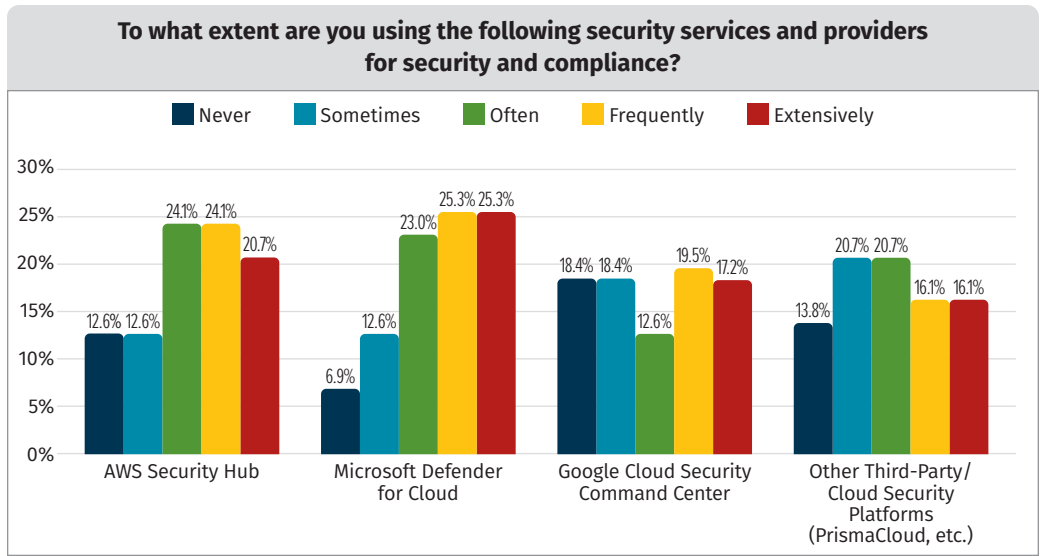


Figure 17. Adoption of Security and Compliance Services

KEY TAKEAWAY

Cloud-specific security services should be used intentionally. Compare the pricing, benefits, assets, and restrictions of each cloud service provider's cloud-specific security solution.

Network Protection

When the rudimentary packet screening of VPC firewalls is inadequate, the major cloud service providers provide managed firewall services that add an extensive array of firewall functionalities.

The survey responses depicted in Figure 18 indicate that for those using cloud services for network protection, the responses are remarkably similar for the AWS Network Firewall, Azure Firewall, and Google Cloud Armor. Most use traditional cloud service providers, but roughly half (45%) use third-party firewall/WAF services

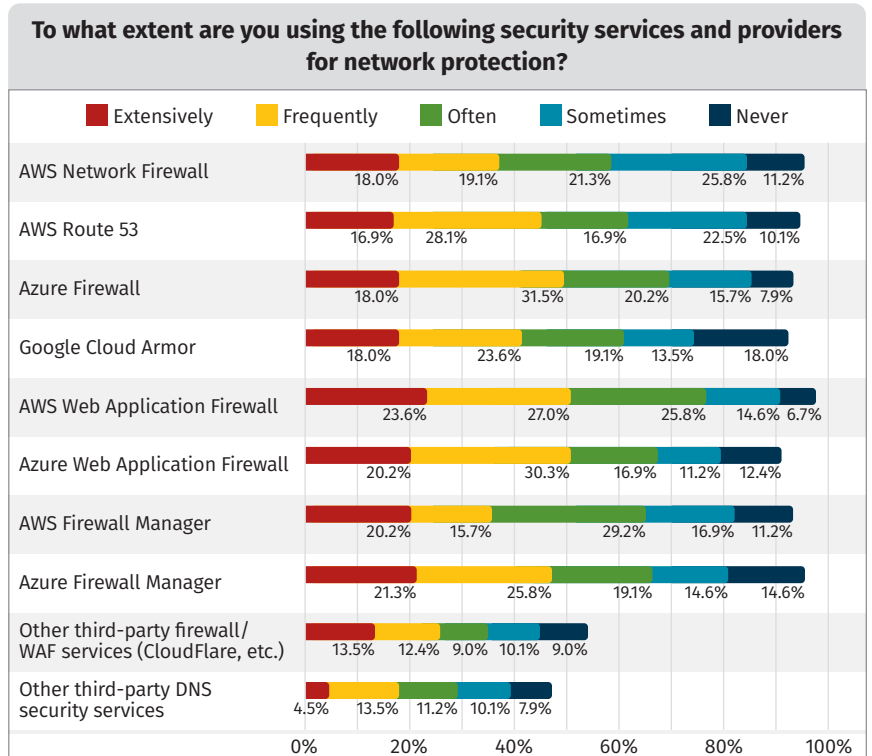


Figure 18. Adoption of Network Protection Cloud Services

A managed firewall serves a distinct purpose set than a WAF. One instance where a WAF is suitable for safeguarding a web application is when the virtual network permits access solely through HTTPS protocols. On the contrary, a managed firewall service can be utilized to safeguard one or more segments of a network against a wide variety of traffic types. Only 29% of participants said that they are using third-party DNS security services extensively, frequently, or often.

KEY TAKEAWAY

Leverage the network protection services as appropriate for your cloud environment and consider adding DNS security as part of your strategy.

The participants in the survey were provided with a rundown of the various data protection options made available by the Big Three cloud service providers (AWS, Azure, and Google Cloud). For those using cloud services for data protection, Figure 19 shows to what extent each service is being used.

KEY TAKEAWAY

Understand the features of the various data protection services that are offered by the cloud providers used by your organization and consider how they may be used to augment your security posture.

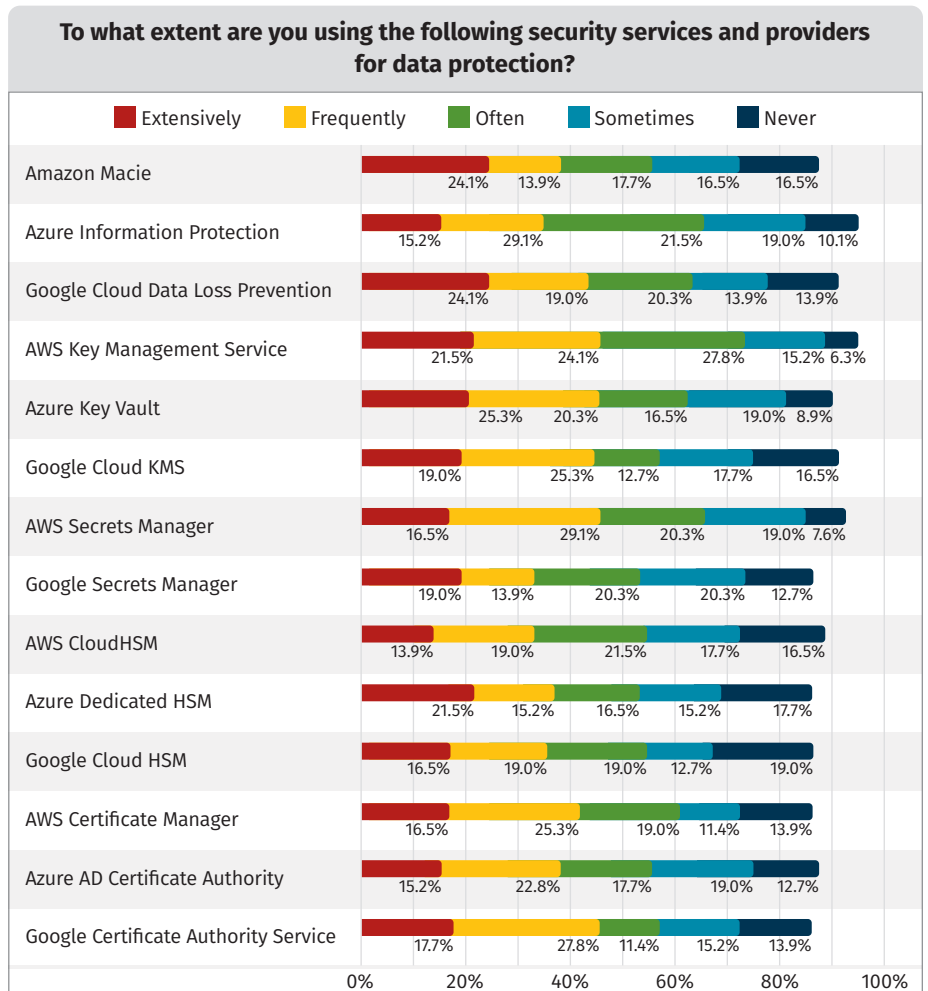


Figure 19. Adoption of Cloud Services Provider Data Protection Services

Incident Response

Each of the Big Three CSPs offers a “SIEM-as-a-service” to support security incident response and investigation. Azure Sentinel has the greatest overall adoption, followed closely by both Google Chronicle and AWS Detective. See Figure 20.

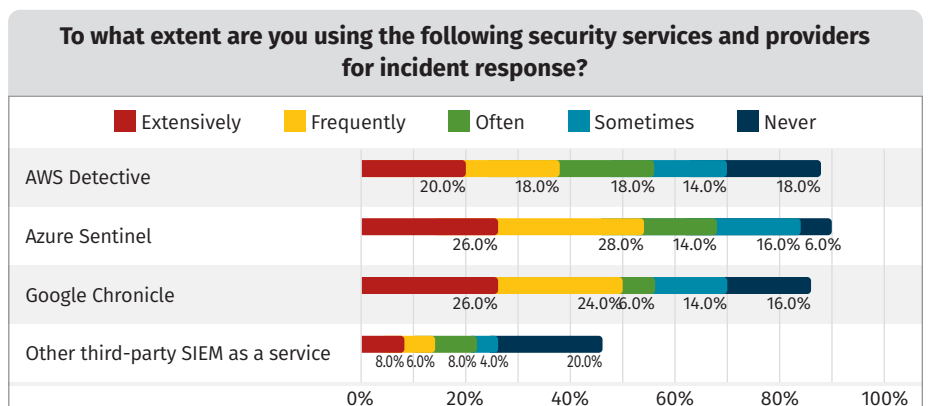


Figure 20. Cloud Services Provider Incident Response

Figure 21 shows the reported response times for the organizations that use a SIEM-as-a-service solution. Of those that responded, 81% said that they are responding at least monthly; of those, 39% said that they were responding daily or as often as the alerts occur.

Figure 22 illustrates how the participating organizations are split between those that use on-premises solutions and those that use cloud-based services. It's rather interesting to see that many of the participants use solutions from more than one SIEM provider. Of course, what is most important is that each company has the important, actionable information available immediately when it is needed.

KEY TAKEAWAY

All three of the largest CSP's offer a SIEM-as-a-service solution, and there are excellent third-party solutions as well. Organizations use a mix of on-premises and cloud-hosted solutions. Regardless of the solution, it is important that cloud customers develop the capability to detect and respond promptly to cloud security events.

DNS Visibility

In the context of cybersecurity, “DNS visibility” refers to the capacity to observe and comprehend all DNS communication on a network. This encompasses both incoming and outgoing DNS inquiries, in addition to the results that are obtained from those requests. The ability to detect and prevent malicious DNS behavior, such as DNS poisoning, DNS tunneling, and DNS beaconing, makes DNS visibility an essential component of multi-cloud security. Sixty-eight percent of respondents who answered this question said they use a DNS security solution for their cloud operations; of those respondents, 41% said it was a solution provided by a cloud service provider.

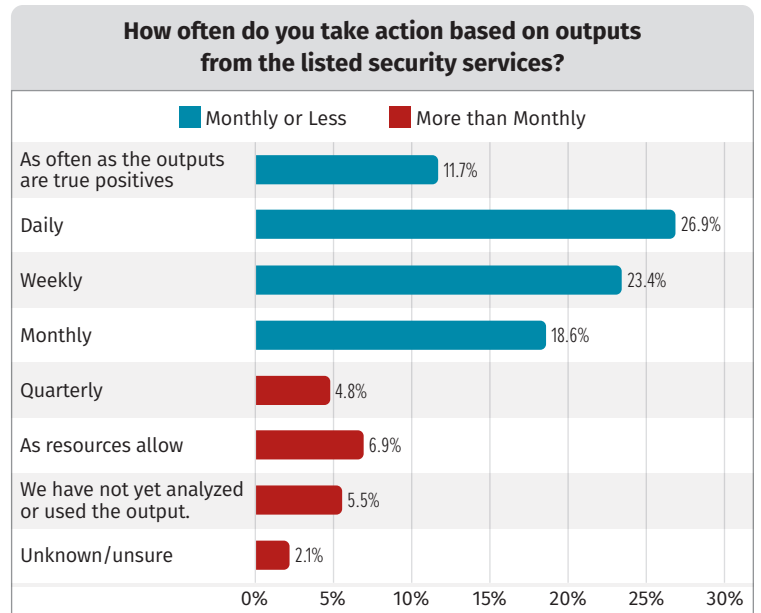


Figure 21. Response Time for Security Service Outputs

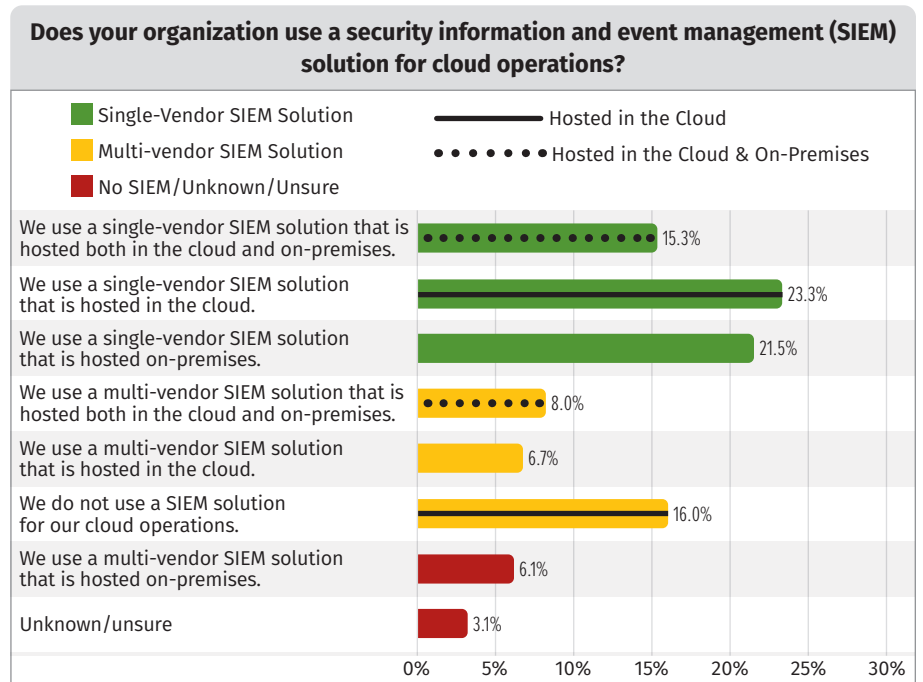


Figure 22. Deployment Models for Security Information and Event Management

The benefits of DNS visibility include:

- **Improved threat detection**—DNS visibility can help to detect a wide range of malicious DNS activity, including DNS poisoning, DNS tunneling, and DNS beaconing. This can help to stop attacks before they cause any damage.
- **Reduced risk of data breaches**—DNS visibility can help to reduce the risk of data breaches by blocking malicious domains that are known to be associated with malware or phishing attacks.
- **Improved compliance**—Many industry regulations require organizations to have visibility into their DNS traffic. DNS visibility can help organizations to comply with these regulations.

DNS poisoning is a type of attack where attackers spoof DNS responses in order to redirect users to malicious websites. DNS tunneling is a type of attack where attackers use DNS traffic to transport malicious data or commands onto or off of a network. With DNS beaconing, attackers use DNS traffic to communicate with a command-and-control server.

DNS security solutions can be used to detect and block these types of attacks by monitoring DNS traffic for suspicious activity. For example, if a DNS server suddenly starts receiving many requests for a particular domain name, this could be a sign of a DNS poisoning attack. Similarly, if a DNS server starts sending a large amount of traffic to a particular IP address, this could be a sign of a DNS tunneling attack.

DNS has some unique risks, which are highlighted here. But it is also a channel that can be used for any malicious activity that is normally monitored on HTTP(S) traffic. For example, it can be abused to exfiltrate (i.e., data theft) or infiltrate (i.e., introduce malware or send ransomware encryption keys) data. Because many SWG and other tools are limited to HTTP(S), and sometimes FTP, attackers have used DNS to easily bypass other defenses. This is one reason many are investing in DNS layer security instead of just add-ons to a next-generation firewall (NGFW) or the like.

DNS visibility can also be used to identify and block malicious domains. For example, a DNS security solution can be used to block domains that are known to be associated with malware or phishing attacks.

The survey indicates that 41% of respondents are using a DNS security defense solution provided

by their CSP, and 25% are using a DNS layer security tool. The balance of respondents (34%) are using a DNS add-on (see Figure 23). Figure 24 illustrates that less than 25% of respondents have implemented DNS visibility for the various devices listed, and between 30% and 41% have plans for the upcoming year.

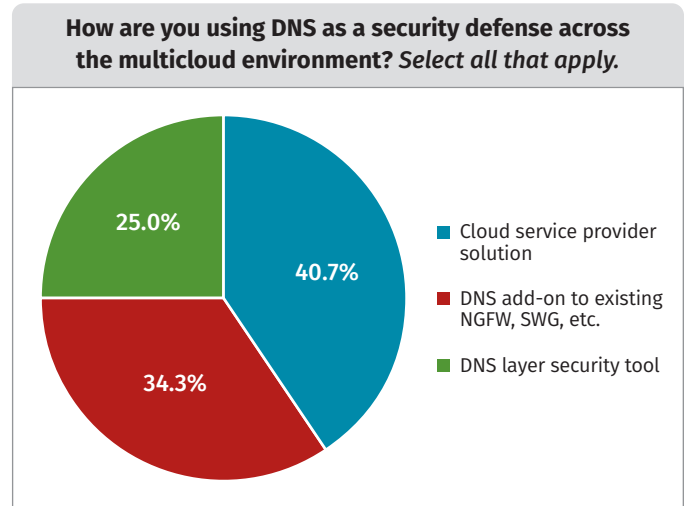


Figure 23. DNS as a Security Defense




Figure 24. DNS Visibility by Device Category

KEY TAKEAWAY

DNS visibility across all device categories is an important component of a comprehensive cloud security strategy.

Cloud Utilization

The extent to which organizations are increasing their adoption of the cloud is important for cloud security teams to monitor. Because cloud services are easy to purchase and convenient to use, costs can get out of control. Not only that, but it is also important to make sure that each new service instance is properly secured.

Also, scale changes everything. As companies move from tens of AWS Accounts, Azure Subscriptions, and Google Projects to hundreds or thousands, the tools to police these environments must also adapt.

Cloud Accounts

Sixteen percent of respondents said that their organization uses more than 100 AWS Accounts, 12% said that they use over 100 Azure Subscriptions, and 12% use more than 100 Google Cloud Projects (see Figures 25–27). This is up from 11%, 8%, and 5%, respectively, as of three years ago.

KEY TAKEAWAY

The tools and practices to police AWS Accounts, Azure Subscriptions, and Google Projects must keep pace as firms grow from tens to hundreds or thousands of cloud accounts.




Figure 25. AWS Accounts Over Time




Figure 26. Quantity of Azure Subscriptions Over Time




Figure 27. Google Cloud Projects Over Time

Storage Services

Another indicator of cloud scale is the number of storage containers a company is using. In this context, we are using “storage container” to collectively refer to AWS S3 Buckets, Azure Storage Containers, and Google Cloud Buckets. Close to 40% of all respondents stated that their organization uses more than 100 of each type of storage container (see Figure 28). When we consider that more than 15% of respondents said that the quantity of storage containers has increased by more than 20% in the past 3 years (see Figure 29), clearly cloud security teams need to put plans and practices in place to manage cloud storage scale and growth.

Cloud storage is relatively inexpensive and convenient to use. Over 14% of respondents said that their organization stores more than 100 terabytes in cloud storage across the Big Three providers (see Figure 30). This makes cloud storage services an attractive target for adversaries. Based on the author’s consulting experience, some companies are not sure what storage containers are needed or what type of data is stored in each. Effort should be made to identify and label each container, apply a retention policy, and eliminate any unneeded storage containers. It is quite possible that this exercise could pay for itself in terms of cost avoidance.

KEY TAKEAWAY

Identify the purpose of each storage container and label it, establish a retention strategy, and remove unnecessary storage containers.

How many storage containers does your organization use?




Figure 28. Quantity of Storage Containers

How has the quantity of storage containers changed in the past three years?

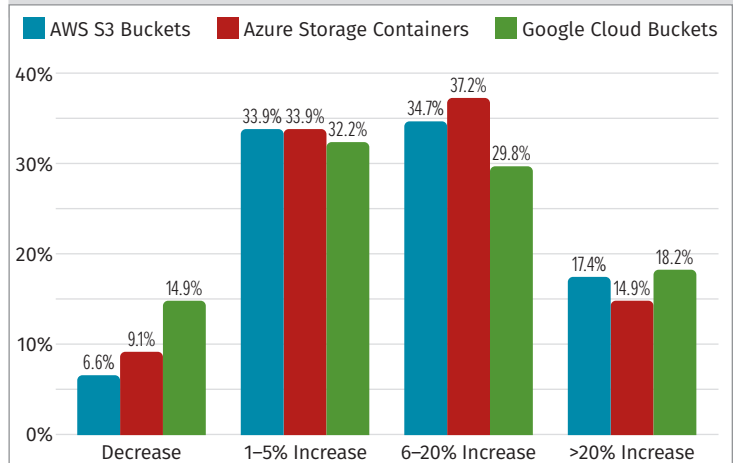


Figure 29. Change in the Quantity of Storage Containers Over the Past Three Years

How much data does your organization store across all AWS Accounts, Azure Subscriptions, and Google Projects?

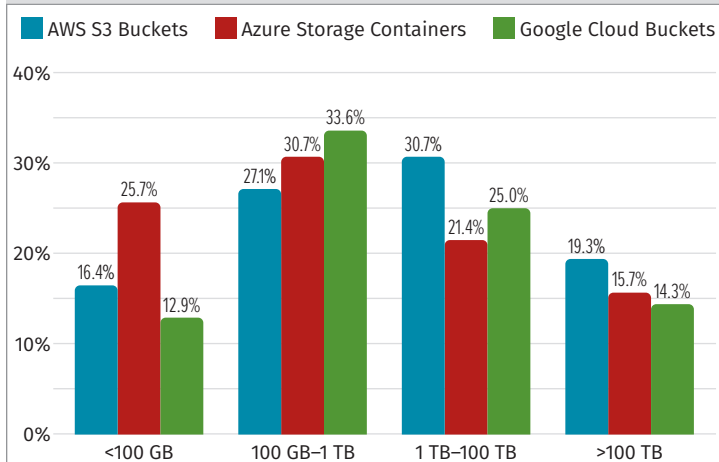


Figure 30. Data Stored Across All Cloud Accounts

How has the quantity of data stored changed over the past three years for your organization across all AWS Accounts, Azure Subscriptions, and Google Projects?




Figure 31. Change in the Quantity of Data Stored Across All Cloud Accounts Over the Past Three Years

Virtual Machine Age

Virtual machine age is an interesting indicator of a company's system delivery processes. Many companies start using the cloud after a "lift-and-shift" project to migrate systems from an on-premises data center to the cloud. The priority is to do whatever it takes to get the systems operational in the cloud. This typically means keeping the same configuration in the cloud that was used on-premises.

Companies that start with a cloud-first approach to launching a new software application will usually use a CI/CD pipeline to deliver system changes using automation. With this approach, a virtual machine image is "baked" with the latest application code and launched in production, fully patched, fully tested, and ready to go. In this scenario, there is no need to patch in place; just rotate out the VM with one that was just launched from a fresh image.

Table 3 shows the reported age distribution of the respondents' fleet of virtual machines. It is of critical importance to ensure that older machines are being patched in place. Also, ensure that these machines can be rebuilt. Sometimes a long-running machine is left alone because its administrators are afraid to touch it. Lastly, if it is necessary to have long-running virtual machines, these are good candidates for a reserved-instance pricing plan.

KEY TAKEAWAY

Identify the appropriate target age of each kind of virtual machine in your environment and manage to that age. Use short life spans to avoid patching in place, and put long-lived VMs on a reserved instance plan.

Table 3. Age Distribution of Virtual Machines

VM Age	Unknown	0%	1–5%	5–20%	20%+
0–1 Days	15.8%	25.3%	25.3%	22.6%	8.2%
1–7 Days	11.6%	19.2%	34.9%	24.0%	7.5%
7–30 Days	15.1%	15.1%	28.8%	30.1%	6.8%
1–6 Months	12.3%	14.4%	30.8%	28.8%	10.3%
6 Months–1 Year	13.7%	19.2%	29.5%	22.6%	11.6%
1–3 Years	17.1%	16.4%	26.7%	26.7%	11.0%
More than 3 Years	12.3%	24.7%	26.0%	22.6%	10.3%

Credential Rotation

Compromised credentials are a leading cause of security incidents in the cloud. Figure 32 shows various types of cloud credentials and the reported frequency at which those credentials are changed. Even more important than the rotation interval is to ensure that each credential has a designated owner and that the credential is being properly handled and managed by that owner throughout its lifecycle.




Figure 32. Rotation Interval for Various Cloud Credentials

KEY TAKEAWAY

Ensure that each cloud credential has an owner and is managed and rotated by that owner per policy.

Legacy Instance Metadata Service

One of the core services on a compute instance is the instance metadata service (IMDS), which internally provides credentials for the virtual machine, among other things.

These credentials can be used by the instance to authenticate its identity to other cloud services. An IMDS is used by virtual machines on AWS, Azure, and Google Cloud. However, a server-side request forgery (SSRF) attack could have exploited the early versions used by AWS and Google Cloud, allowing an attacker to steal the data from the instance. Google Cloud recently disabled all its vulnerable versions as a result. But by default, AWS's IMDS runs in a vulnerable manner, and the user must modify the setup to disable it. Sadly, many cloud security teams are ignorant of the significance of IMDS, one of the most important settings for a fully hardened EC2 instance.

We inquired of the respondents whether this had been addressed by their company yet. Figure 33 presents the findings. Less than a fourth (22%) of respondents said they have handled the vulnerability entirely or that it does not apply to them, which makes this response troubling. Nineteen percent stated that although they haven't started yet, they intend to address the problem. This can be because of worries about breaking older applications, a sign of inadequate testing environments, or other issues with the company.

The Capital One breach in July 2019 was the most well-known IMDS exploit. The Social Security information of 140,000 credit card holders, 80,000 bank account numbers connected to credit cards, and 106 million credit card applications were made public thanks to an exploit in this hack that took advantage of the outdated AWS IMDS. For this breach, Capital One settled a \$190-million class action lawsuit in 2022.

Sadly, as the prior data indicates, most firms continue to use the outdated IMDS that Capital One formerly employed. It won't be long until there is another breach of this magnitude, partly because of this service. Even worse, one may have already happened without anyone noticing.

If your organization uses AWS, has your organization turned off IMDSv1 in favor of IMDSv2 for your compute instances?




Figure 33. AWS IMDS Mitigation Status

KEY TAKEAWAY

Ensure all AWS EC2 users understand IMDS security and put preventive controls in place to prevent the deployment of EC2 instances that use IMDSv1.

Firebase

Google owns the Firebase cloud service suite, which has loose integrations with the Google Cloud infrastructure. The use of Firebase by development teams is quite common. The Firebase Realtime Database, the company's main product, is made to enable front-end clients to access data via the open internet.

Even though a real-time database can be configured securely, this is rarely done correctly. A 2018 report that was saved by a security research company that Symantec later purchased revealed that 3,000 Android apps made use of globally viewable and writeable Firebase databases. Hundreds of millions of records were stored in Firebase Realtime Databases by these applications, which have apparently been downloaded over 650 million times. These databases are easily copied and corrupted by anonymous users.¹

¹ Brandon Evans, "Firebase: Google Cloud's Evil Twin," October 8, 2020, <https://www.sans.org/white-papers/39885>

In July 2021, Avast Threat Labs discovered around 19,300 vulnerable Realtime Databases from a sample of approximately 180,300.² Similarly, in 2022, Check Point Research found that hundreds or more programs with unsafe Realtime Databases are posted each month to Virus Total.³

Security circles hardly ever mention Firebase. We were interested in finding out how much Firebase is used or known to responding organizations. Figure 34 presents the findings. We are most concerned about the 28% who have never heard of Firebase or are unsure if their organization is using it, because these organizations may not be secure.

KEY TAKEAWAY

Ensure that developers that use Firebase know how to write secure Firebase applications and that the Security Teams know how to validate Firebase security.

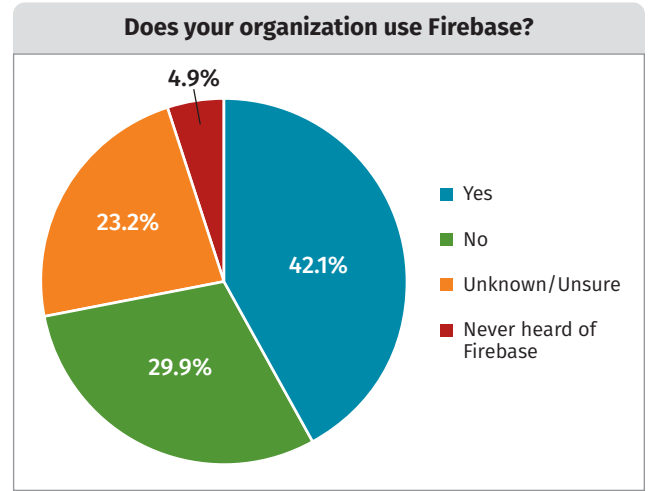


Figure 34. Firebase Usage

Professional Services

Figure 35 details the types of third-party professional services that are utilized. The most used professional service from a third party is penetration testing at about 71%, followed by application security assessments (51%) and security audits (51%). There is a benefit to outsourcing these services and changing up the vendors to get fresh eyes on the subject being evaluated.

KEY TAKEAWAY

Leverage third-party expertise to make sure your security team is aware of the evolving threats that face your cloud applications, especially with penetration testing and application security assessments.

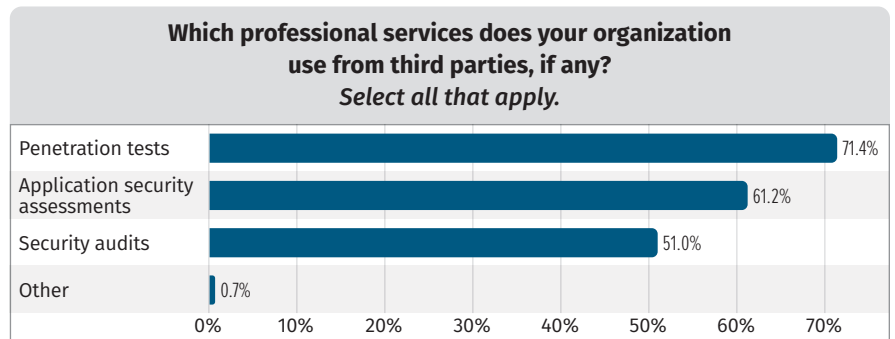


Figure 35. Usage of Professional Services

² Vladimir Martyanov, "Research Shows Over 10% of Sampled Firebase Instances Open," <https://decoded.avast.io/vladimirmartyanov/research-shows-over-10-of-sampled-firebase-instances-open>

³ Check Point Research Team, "Stop Neglecting Your Cloud Security Features," <https://blog.checkpoint.com/2022/03/15/stop-neglecting-your-cloud-security-features-check-point-research-found-thousands-of-open-cloud-databases-exposing-data-in-the-wild>

Conclusion

The use of multiple CSPs is a growing trend that is likely to continue. Organizations need to be aware of the security challenges that this trend introduces and take steps to mitigate those risks. By establishing a centralized cloud security governance framework and using cloud-specific security services, organizations can secure their cloud environments and protect their data.

Sponsor

SANS would like to thank this paper's sponsor:

