A SANS Survey

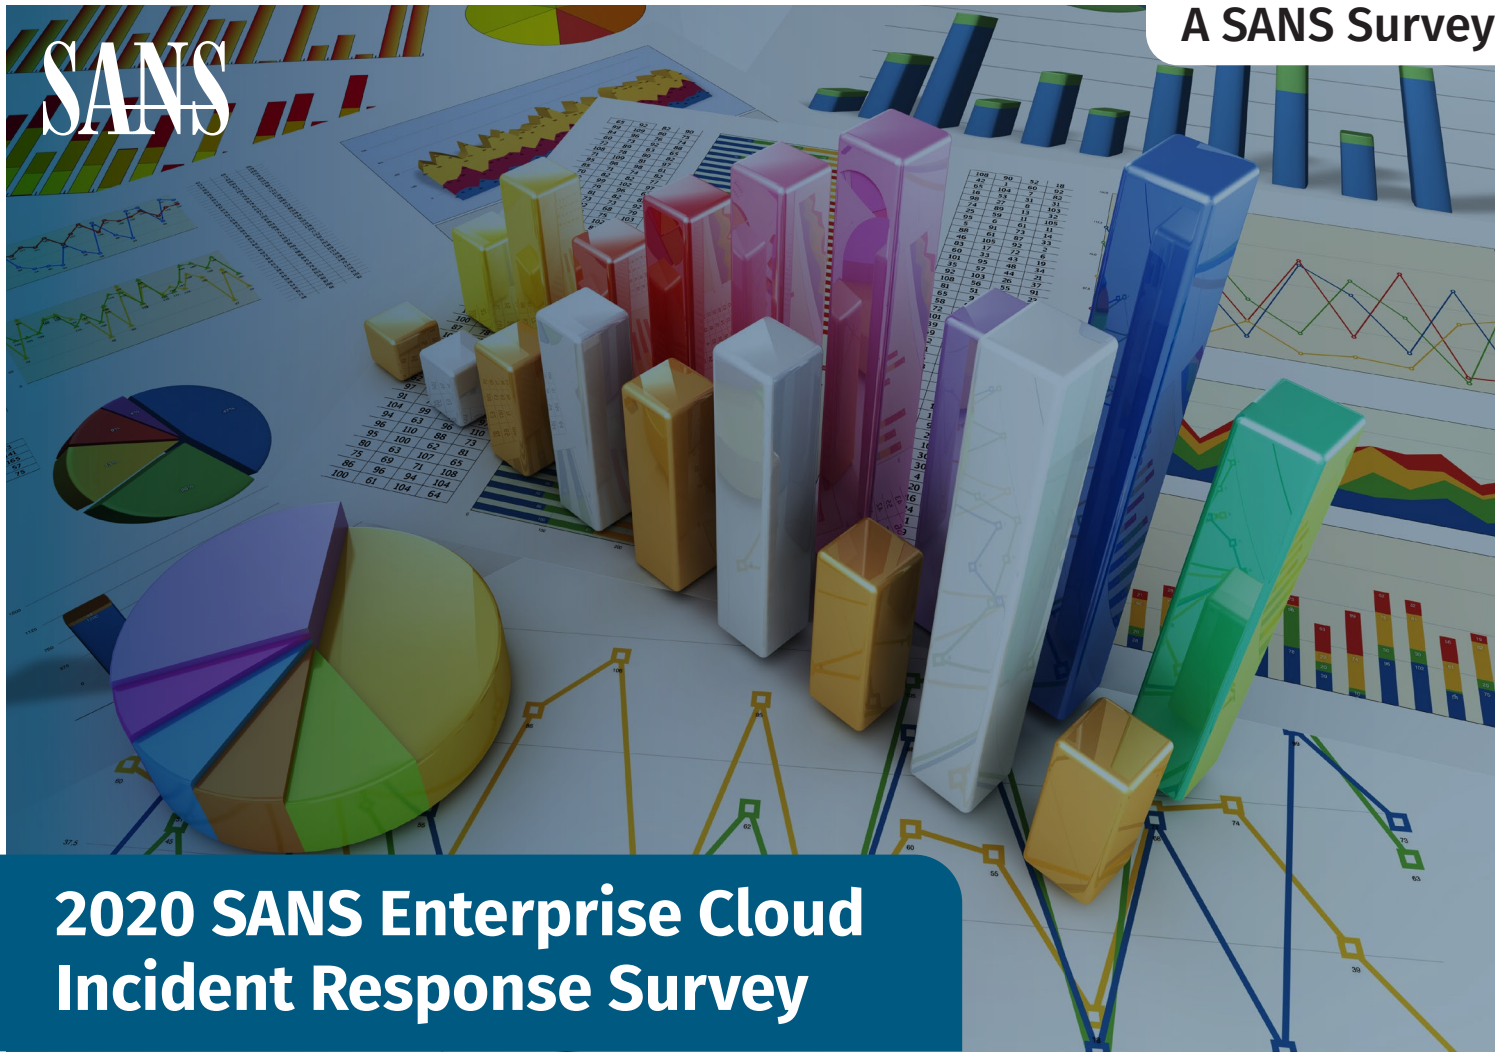

2020 SANS Enterprise Cloud Incident Response Survey

Written by **Chris Dale**
Advisor: **Matt Bromiley**

Sponsored by:
ExtraHop

September 2020

Analyst Program

Executive Summary

This whitepaper examines the results of the 2020 SANS Enterprise Cloud Incident Response (IR) Survey. The survey was promoted to the information security community during the first half of 2020 and garnered 218 respondents.

A solid 40% of respondents stated that they do not assess the effectiveness and maturity of cloud IR processes, indicating that the cloud shift hasn't fully grown and integrated into their businesses just yet. In cloud environments, most seek to solve before they understand, ensuring that the gap and visibility of effectiveness and maturity will likely increase in the future. The cloud is a great enabler, but without proper understanding of the different components deployed, challenges within IR will continue to develop.

The survey also explored which sources of data and tools incident responders consider most valuable. It is interesting that traditional sources of data (e.g., network data) are highly sought after, but often unavailable in many traditional cloud environments.

In cloud environments, incidents are plentiful and often related to data theft or compromised cloud components through leaked API keys, credentials and targeted attacks, for example.

Figure 1 provides a snapshot of the demographics for the respondents to this survey.

Top Takeaways

This year's survey shows a clear lack of staffing and skills necessary to respond in cloud environments—a top concern for many respondents, with more than 50% of respondents listing those shortcomings as the key impediment to effective cloud IR. Although hiring staff proportionally in tandem with increasing the number of cloud services in use is not always the best solution, it aligns with organizations' desire to increase automation. Automating functionality for IR is the top improvement that organizations are planning for cloud IR, and several respondents commented that they envision automation as a possible solution.

Top 4 Industries Represented

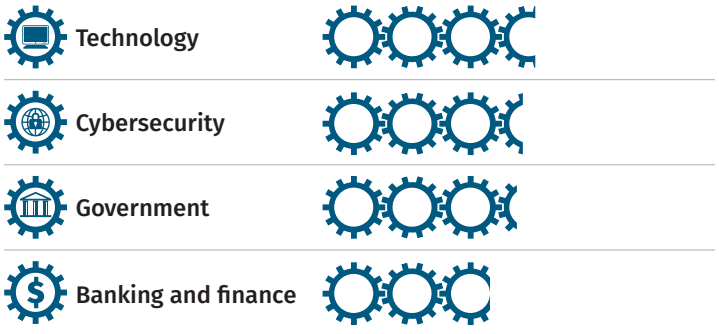

Each gear represents 10 respondents.

Organizational Size

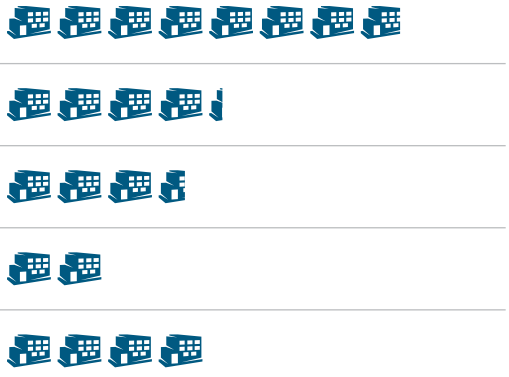

Each building represents 10 respondents.

Operations and Headquarters

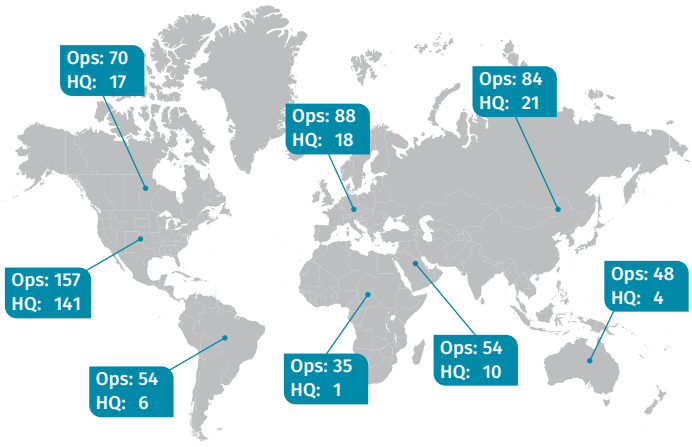

Top 4 Roles Represented

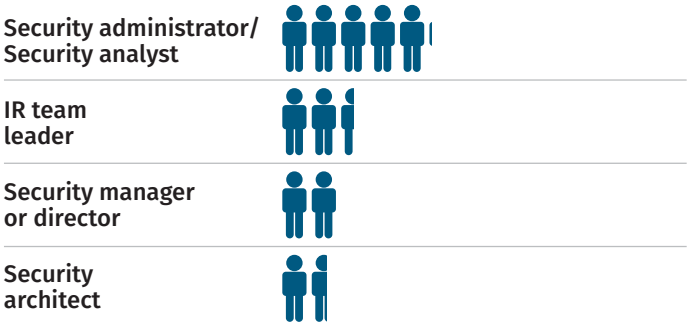

Each person represents 10 respondents.

Figure 1. Survey Demographics

Cloud Platforms

Cloud platforms seek to solve different challenges. Some providers, such as Google's G Suite and Microsoft Office 365, directly support companies with their business processes via their SaaS services, while others, such as Amazon Web Services (AWS) or Alibaba Cloud, are designed for provisioning their own custom services through platform-as-a-service (PaaS) and infrastructure-as-a-service (IaaS) services.

For cloud services such as compute, storage, networking and security, AWS and Microsoft Azure are the leading vendors. Seventy percent of respondents indicated that they are using Azure, while 64% are using AWS. We cannot neglect the 7% who selected "other," including popular cloud services such as DigitalOcean. See Figure 2.

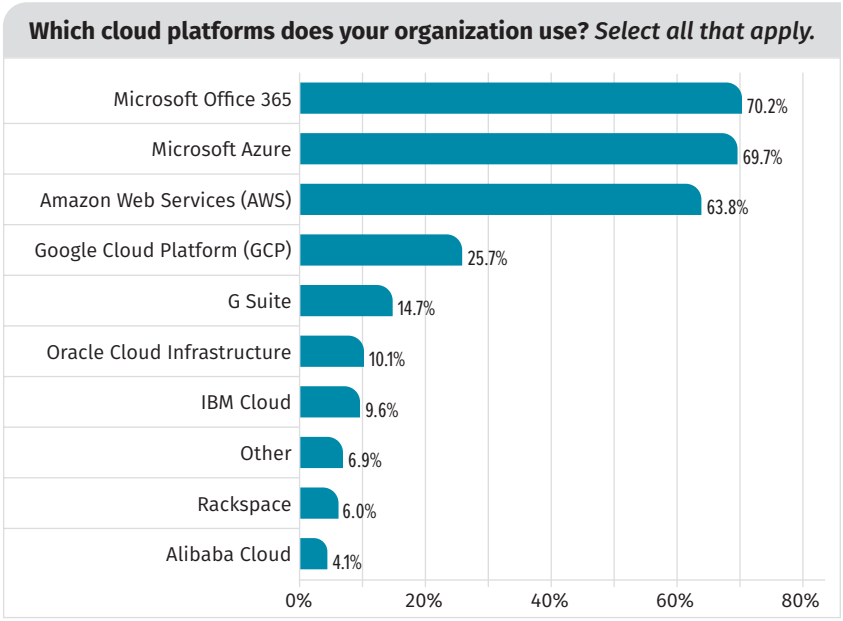

Figure 2. Cloud Platforms in Use

We can infer that organizations are using more than one cloud provider, possibly with competing services. Considering the survey's striking numbers on skill shortage, it's interesting to see that organizations are willing to invest in multiple cloud vendors, even if there are severe challenges on the existing choices made (see the "Impediments and Challenges" section later in this paper). Cloud platforms can also include productivity applications. Microsoft Office 365 (70%) dominates in this space, whereas the closest runner-up is G Suite, with a meager adoption rate of 15%. Office support platforms outside of Microsoft Office are still considered unusual and not very widespread. The upside of this is higher quantity in staff, skills and tools on this platform supporting IR efforts.

Cloud-Attacker and Cloud-Defender Maturity

A significant number of attacks on cloud environments do not appear to compromise the rest of the cloud infrastructure. There are multiple possible reasons for this:

- It's not yet a part of attackers' strategic motivations.
- Attackers have not yet matured to this point and are incapable of performing such attacks.
- Organizations fail to detect that attackers are already compromising their cloud infrastructure.
- Cloud services normally come with built-in segregation and least amount of privileges.

It is likely that attackers will explore such opportunities more in the future, perhaps when the attacks become more commonplace or tactics that support their efforts are present. One thing is certain: The potential to pivot to enterprises' cloud infrastructures is extremely fruitful. Attackers are likely to explore such opportunities and make the necessary investments in time and research. Technically speaking, the capabilities to pivot between cloud resources are present today—through compromising the necessary identities with the appropriate roles. However, our respondents are not seeing these compromises in their environments yet.

During their investigation of an incident or breach, respondents were unable to consistently and accurately discover impacted API keys (24%) and identity and access management (IAM) roles (36%), as compared with impacted users (74%), systems (65%) and data (41%). See Figure 3. This result is not surprising, because API and IAM roles are not compromised as often as users, systems and data. Data exfiltration (stealing sensitive data), at 41%, was the leading component in the breaches our respondents encountered.

When investigating these incidents and breaches, were you able to consistently and accurately discover the impacted API keys, IAM roles, users, systems, data, transactions and threat actors involved? ***Select all that apply.***

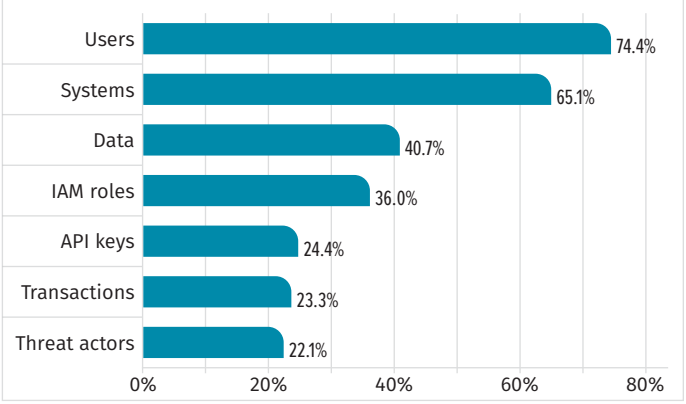

Figure 3. Consistency and Accuracy of Investigations

A given API key could lead to loss of many things throughout the application. However, if sensitive IAM roles are lost, much more damage and compromise could occur. (For example, if an owner of a resource were compromised, we could expect that resource would be fully compromised. In some cases, that resource could be used to stage other attacks against users of the platform.) It is likely that the built-in segmentation of many cloud services makes it substantially harder for attackers to pivot and perform lateral movement, which happens commonly in internal networks. There are many lessons to be learned on this account for hybrid and internal networks!

The defenders, however, reported very interesting numbers about how many incidents are detected externally rather than internally. Most respondents (32%) indicated that a third party detected more than 91% of their incidents, as shown in Figure 4.

What percentage of cloud incidents that you responded to were detected internally as opposed to being identified by an external party? ***Select the best answer.***

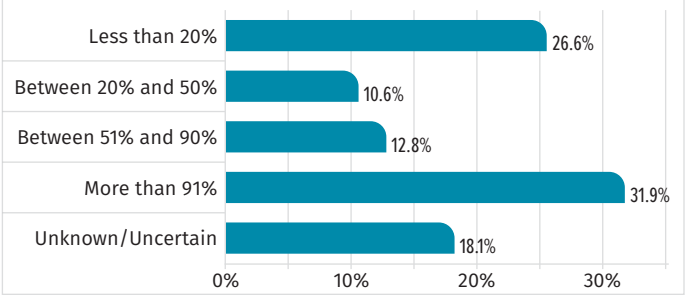

Figure 4. Internal vs. External Detection

This could be because of the built-in security centers of some cloud platforms, but perhaps the most interesting aspect of this is the mutual benefit of getting rid of security threats. As a customer, it's obvious why you wouldn't want incidents in your environment; but as a cloud provider, there are multiple reasons to assist your customers. An example—and the most obvious one—is attacks that might impact the performance of multiple tenants of the cloud environment (e.g., a DoS attack). Another aspect is keeping customers happy and ensuring that they're not billed for attacks that cause computing and networking fees. The inherent symbiosis of provider and customer is nothing new, but perhaps an aspect and developing factor of cloud environments.

Speed and Agility

The benefits of cloud environments are ripe for enabling and supporting organizations' speed and agility when dealing with incidents. The requirements in InfoSec must increase drastically if we are to deal with advanced attackers more effectively. Cloud environments support aspects such as elastic scaling, deployment pipelines and infrastructure-as-code, enabling organizations to take advantage of many easy wins in terms of containing and eradicating threats.

Reports of lower breakout times (the time attackers use from breaking in to pivoting) has been a concern for many. With built-in segmentation, the least amount of privileges and challenges in compromising more of the cloud environment, the cloud helps defend against the breakout times.

Unfortunately, the data from the survey is not in favor of the defenders just yet. The time from compromise to detection (also known as the *dwell time*) still tends to be weighted toward days and months, as opposed to less than 24 hours. See Table 1. The time for compromise until breaking out from the host, sometimes called *breakout time*, is normally within the matter of hours, but luckily not necessarily during cloud deployments.

The stats from detection to containment, and containment to remediation, shed a better light on response. We see that a majority of incidents are contained and remediated within one to five hours of identification. This is likely due to having an easier time in the cloud, because they often present better management utilities, less coupling between services and built-in segmentation. Automation will doubtlessly help in lowering the numbers even further, both from detection through containment and remediation.

Table 1. Compromise to Detection to Containment to Remediation

Duration	Time from Compromise to Detection	Time from Detection to Containment	Time from Containment to Remediation
Unknown	19.2%	7.7%	11.5%
Less than 1 hour	3.8%	19.2%	15.4%
1–5 hours	0.0%	30.8%	15.4%
6–24 hours	19.2%	15.4%	7.7%
2–7 days	30.8%	23.1%	23.1%
8–30 days	11.5%	3.8%	15.4%
1–3 months	15.4%	0.0%	3.8%
4–6 months	0.0%	0.0%	3.8%
>1 year	0.0%	0.0%	3.8%

Cloud environments normally provide better visibility into assets and the overall attack surface. You might argue that non-cloud environments are also more complex—in the sense that you might not have enough control and visibility of them. This complexity causes automation to be a dream further down the road; whereas, if services were moved into the cloud, assets could have an inherent strategy to build in capable IR and automation to support it.

Impediments and Challenges

According to respondents, a shortage of staffing and skills (55%) is the top impediment to effective cloud IR in their organizations. "There's not enough competency in existing staff, and there's not enough of us" is a frank reply from one respondent, but commonly not the best solution to a problem. Skills go a long way, but there are inhibitors to growing IT and security alongside the rapid and numerous new deployments that many organizations face. Tools and technology can go a long way toward combating this problem, but they aren't always feasible options for many companies. See Figure 5.

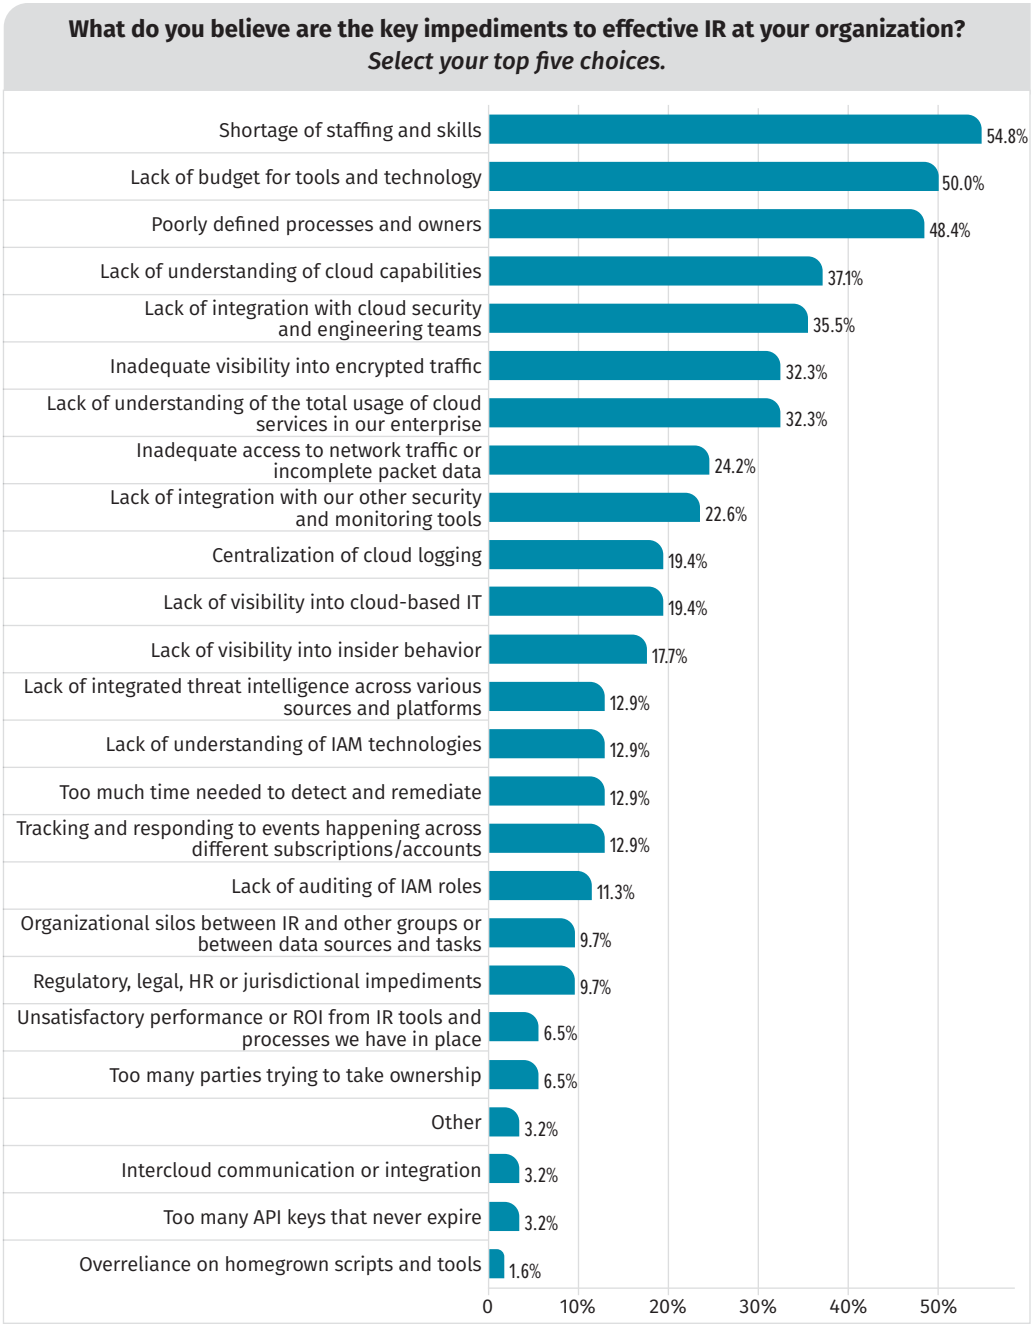

Figure 5. Impediments to Effective Cloud IR

The same applies to a lack of budget for tools and technology (50%) as an impediment. There is a continuous stream of new services offered in the cloud and new technology to understand, and one thing is for certain: We don't want security to be an inhibitor of innovation and development for the company.

To solve the overarching problem, governance is key. A cloud environment can help solve many issues and support the necessary level of development; however, perhaps instead of seeking only to solve, we must strive to understand and govern the platform in which we now reside. Governance should aid in doing more with less; ensuring security and IR are integral parts of existing efforts and regulations we're currently undergoing in the on-premises networks. Many organizations are leveraging the cloud, but based on survey results, frequently without the necessary governance and processes in place to ensure viable and secure operations. We see the opportunities and potential, but fail to see security as a natural and necessary part of the life cycle.

Assessing the effectiveness and maturity of one's own cloud IR capabilities is still a foreign concept to many (40%)—while 24% are using well-defined, public metrics (such as NIST) to help track, evaluate and update the plan. Assessing the effectiveness and maturity of cloud IR can help companies better understand the platforms they are operating in and help secure governance. See Figure 6.

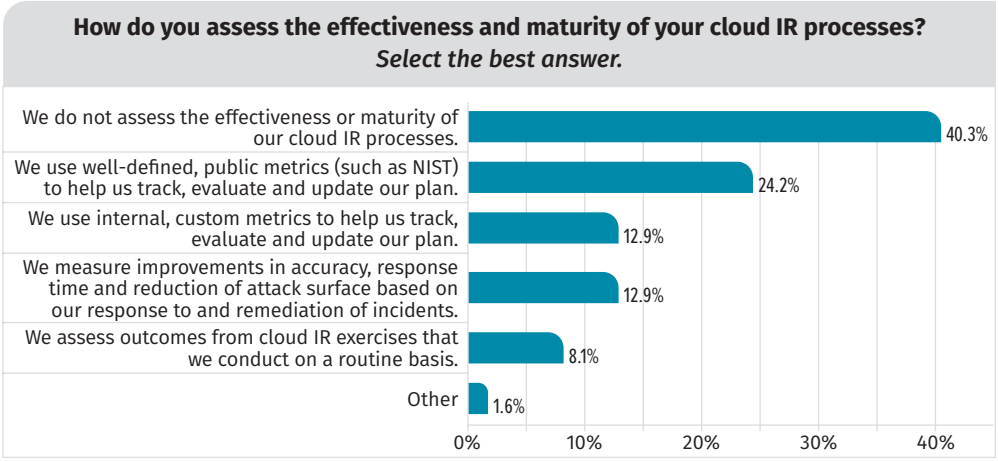

Figure 6. Cloud IR Effectiveness and Maturity

If we take the 20 CIS Controls,[1] the first and second controls are inventory and control of hardware and software assets, respectively. Can you properly define which resources your team will perform IR on? Inventory of assets and software has gone far beyond enumerating IP addresses and scanning them for service banners. The cloud can have all kinds of services deployed, some which can be hard to identify and find (even as a good guy or a bad guy). Consider, for example, message queuing functionality, serverless functions, PaaS exposing just about anything you can imagine, SaaS that connects to hybrid deployments—and the list goes on.

In terms of improvements, we're looking at a push for automation, perhaps in the light of Security Orchestration Automation and Response (SOAR) technology. Training can be used as a filler in the long run to bridge gaps in staff and competency, but to better understand the tactics, tools and procedures (TTPs) that adversaries are using against us, will training, bigger budgets for tech and more people in the equation be enough?

Breached Components and Impacted Systems

Cloud environments consist of many different components and systems, and a breach of any will impact the organization differently. Let's focus on how breaches affected our respondents' organizations and which of their systems were affected.

Data exfiltration is the component most commonly attributed to breaches, at 42%, and there should be additional measures to help organizations understand which data was lost. As discussed previously, while impacted users and systems are most readily identified during breach investigations, the affected data is identified less often, as reported by 41% of respondents. Data exfiltration can often have long-lasting effects on companies because the data might be sold and, in general, used against the company and the users it affect. Data might also have regulatory compliance requirements attached to it (e.g., when it regards personally identifiable information [PII] and healthcare). See Figure 7 on the next page.

[1] www.cisecurity.org/controls/

Analyzing the systems involved in breaches, we see file storage (Amazon S3, Dropbox) is heavily involved, and IaaS components such as EC2 are reported as being heavily involved too. The adoption rate of these traditional cloud capabilities is not surprising, because companies are likely much more mature and familiar with such technologies. A quick takeaway of this would be more and extensive logging and control of the data aspects involved with the services; after all, in most cases, data is what we're trying to protect. Businesses should also consider protecting the data itself, without relying solely on storage providers to implement protection mechanisms. An example would be file formats that support built-in encryption of data, such as Azure Information Protection. This prevents data from being compromised, even if the storage container hosting the data is compromised.

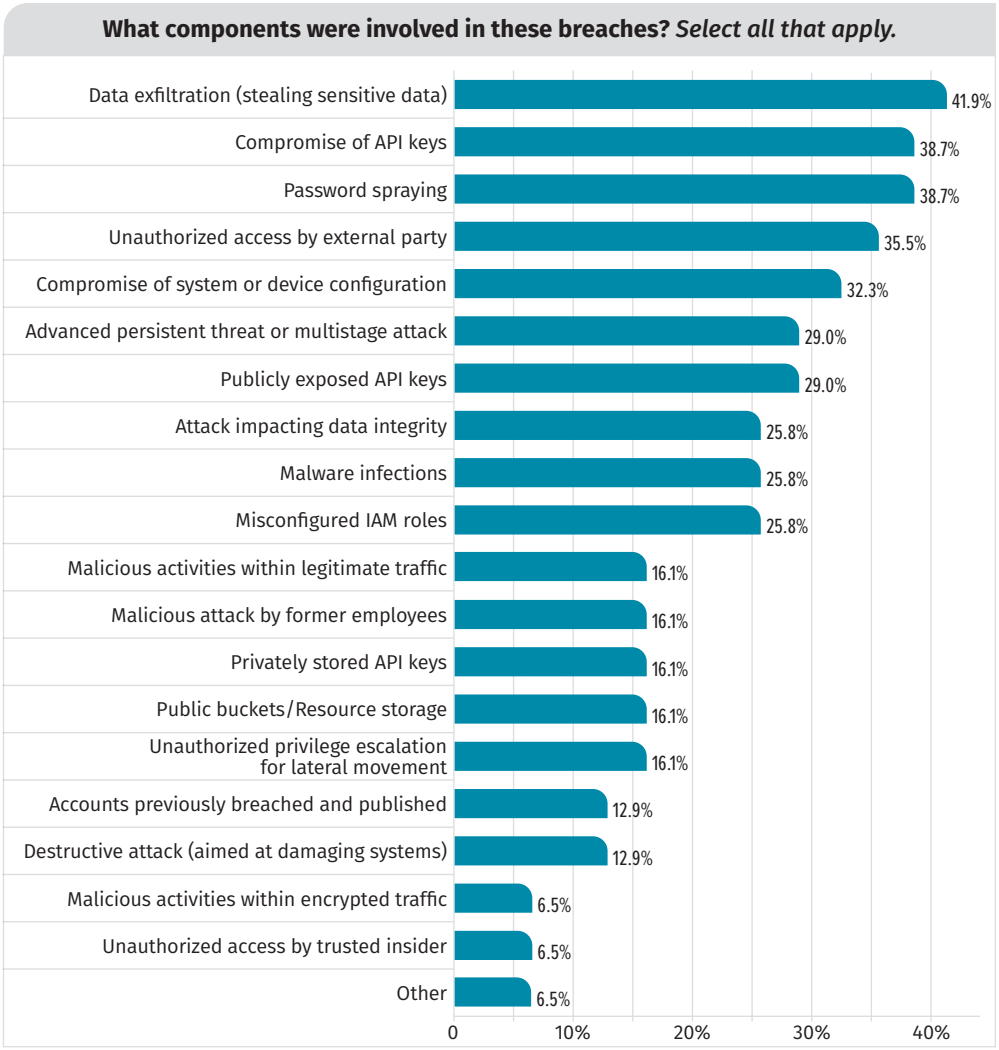

Figure 7. Components Involved in Breaches

Considering that data exfiltration rated as the most common component involved in breaches, we might be looking at misconfigured Amazon S3 buckets or things such as open service-buses in Azure. There's a myriad of problems that could relate to data exfiltration occurring in the cloud and when having to carefully control access across storage accounts. There are even search engines today that look for unsecured Amazon S3 buckets, attributing to attackers' maturity in attacking these kinds of services.

A whopping 68% of breaches involved API keys—39% compromise of API keys and 29% publicly exposed API keys. (See Figure 7.) APIs were a source of compromise for many organizations, publicly exposed and otherwise. This could be because of commits to version control systems, exposed online or to the compromising party, containing the keys by mistake. Mobile applications often embed API keys in configurations or simply expose them through JavaScript in web applications, enabling attackers to take advantage of APIs likely unrestricted if there's no built-in server-side access control.

A full 39% of respondents indicated that password spraying was a component of breaches. Password spraying (and likely credential stuffing) is the reason for many business email compromises (BECs). Also, when migrating a server to the cloud, sometimes organizations might be exposing more services than before (e.g., exposing Remote Desktop Protocol

[RDP] to the internet). Password spraying might now allow attackers a foothold, exposing the organization to attacks including ransomware and cryptocurrency mining. Services that were previously internal are now suddenly exposed to the internet. Zero trust models and proper governance would surely go a long way toward addressing this issue.

How is data collected and aggregated? Collecting and reviewing log data to help identify and conclude incidents has been a struggle for many organizations. In the cloud, ELK (according to 37% of respondents) is gaining popularity as a viable, free alternative to getting started in getting logs addressed, parsed and supported, or in support of a licensed solution. Splunk also seems to have quite a bit of traction, with an adoption rate of 54% by respondents.

A Crypto Paradox

The survey results show an interesting paradox: Incident responders want more insights into network traffic in the cloud environment for IR, and encrypted traffic is high on the list. But according to our respondents, it is also the hardest to acquire, as shown in Table 2.

Table 2. Cloud-Generated Data by Preference

Data Types	Need But Can't Acquire
Data from endpoints (virtual machines/containers)	8.5%
Host, domain and URL reputation data	5.6%
Indicator of compromise (IoC) threat intelligence data	9.9%
Short-term historical event data and logs (as much as seven days old) from SIEM	5.6%
Long-term historical event data and logs (older than seven days) from SIEM	12.7%
Virtual network TAPs	23.9%
Virtual network flows	18.3%
Transaction data from encrypted network traffic	38.0%
Cloud audit logs (Microsoft Azure Audit Logs, AWS CloudTrail, Microsoft Office 365 audit logs, etc.)	2.8%
ML/AI-assisted cloud provider detections (Amazon GuardDuty, Microsoft Azure Sentinel, Google Cloud Security Command Center)	11.3%
Related alarms from IPS, antivirus, network detection and SIEM	8.5%
Threat campaign data	21.1%
Vulnerability data	5.6%
Other	2.8%

This traffic is often hard to support because tenants in cloud environments share the network stack. Furthermore, decryption of network traffic is not ideal, because it is what protects us from other tenants spying on our communication and the cloud providers themselves. Best practice dictates end-to-end encryption, ensuring confidentiality and integrity of data, systems and users. Strategies indicate a push toward using metadata (e.g., Transport Layer Security [TLS] and Server Name Indication [SNI]) to take action on the data, but this becomes more problematic with TLS 1.3, which supports encryption of the SNI. Out-of-band decryption and analysis is one option for retaining the security and privacy benefits of encryption, while still gaining valuable insight into network traffic in the cloud.

Cloud Tools and Capabilities

SANS asked about the tools or capabilities organizations are using to identify cloud incidents as well as the level at which organizations are integrating these capabilities into their overall incident response. Already present in cloud platforms, endpoint detection and response (EDR) capabilities are the leading tools or capabilities that respondents are using to identify cloud incidents. Forty percent reported that these tools are highly integrated with their overall IR operations, and another 36% reported EDR being partially integrated (see Figure 8 on the next page). However, there is a caveat: EDR today only supports platforms in which the full operating system is available to the tool, whereas cloud platforms might be only a function-as-a-service (FaaS).

Network detection and response (NDR) capabilities and endpoint network filtering are almost equal in their respective integration rates. Response efforts can be applied directly to endpoints, but NDR will often grant better visibility in terms of scoping an incident and allows for containment actions for threats. Network analysis can inform investigations and decisions about what those responses should entail. Endpoints can be many different things in a modern cloud environment, and for this reason, IAM cloud controls will likely play a significant role in containing incidents.

Does your organization use any of the following tools or capabilities to identify cloud incidents?
Indicate how integrated each capability is with your overall IR and check N/A for those that don't apply.

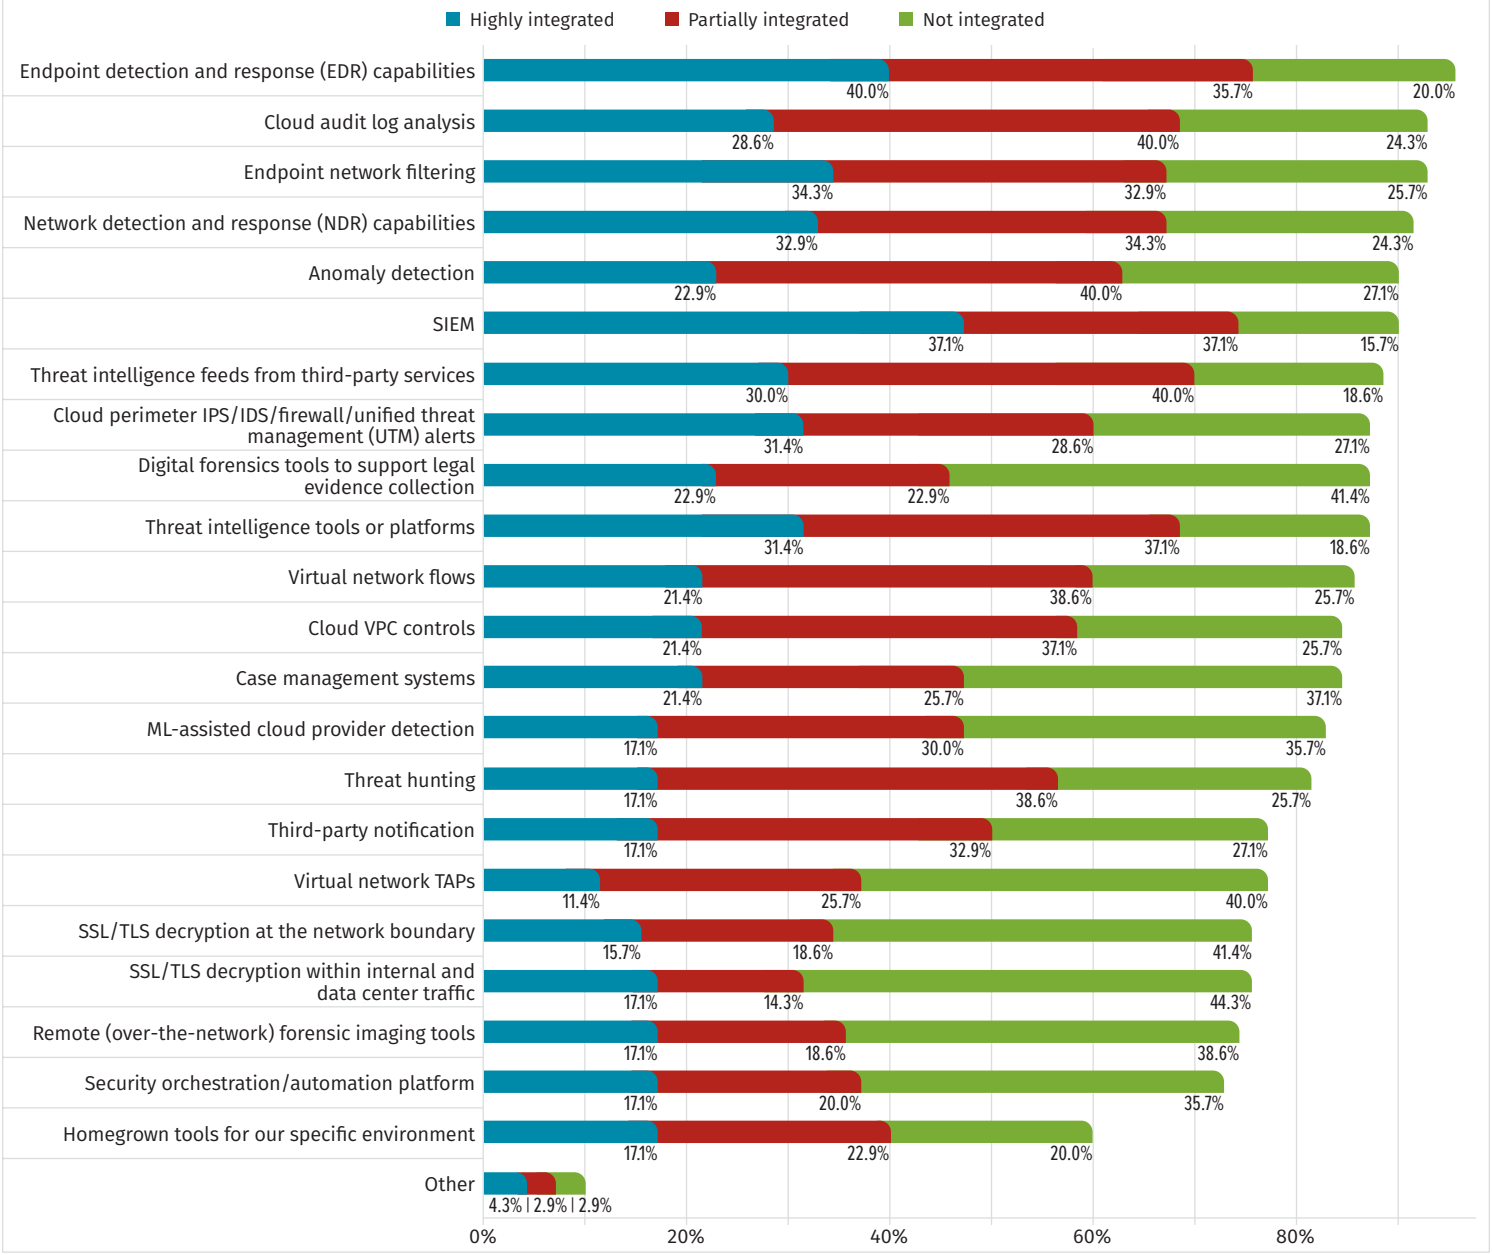

Figure 8. Tool and Capability Integration

Threat intelligence tools, platforms and feeds are highly integrated, per 31% of respondents. Threat intelligence could, for example, provide organizations with detailed information regarding the latest threat actors' TTPs in attacking cloud assets, or feeds could keep block lists up to date with the latest known malicious systems. The fact that 41% of respondents were breached multiple times by the same threat actor is concerning, and building threat intelligence capabilities would support IR efforts greatly in lowering this number. Organizations could also rely on threat intelligence to provide alerts when credentials of leaked users are exposed; however, we must assume that MFA will be implemented everywhere soon, because that helps mitigate this problem. Without MFA on public-facing assets today, we know we're doing something wrong.

Cloud audit log analysis is often a great place to start understanding what happened to the platform in the case of incident. Cloud audit log analysis is highly integrated for only 29% of respondents, although 40% claim partial integration. These logs, which are often already readily available for collection and analysis, could give companies more insights into how the platform is affected.

Threat hunting does stand out as something that is partially integrated (39%) in respondents' organizations. The concept of threat hunting has gained quite a bit of traction in the industry, but perhaps the slow adoption rate for cloud environments is lack of understanding and visibility. Cloud environments do not yet support a proper overview of the attack surface without using scripts and tinkering to understand it. If one were to generate good hypotheses for compromised cloud environments, threat hunting could become a great way to assist organizations in identifying gaps in important questions they must be able to answer.

Anomaly detection and ML-assisted cloud provider detection (with 17% highly integrated) are somewhat in the same alley. They require incident responders to investigate, perhaps following playbooks to support them in concluding alerts. This area of alerting often can help identify sophisticated threats in an on-premises environment, but it is not first in line for integration in cloud environments.

Many of the remediation actions are not currently automated, as shown in Table 3, whereas cloud environments often have great support for command-line tools, APIs and remote-management through automation. The lack of automation probably corresponds with the lack of skills and training within these environments; however, the same applies to many on-premises environments. Automation will be key to ensuring rapid responses for the future, so there is definitely a gap here that vendors and projects that support making automation easier can fill.

Table 3. Remediation Processes

Process	Manual	Automated	Both	Total
Create snapshots	43.3%	37.3%	16.4%	97.0%
Identify similar systems that are affected	67.2%	17.9%	11.9%	97.0%
Isolate infected virtual machines from the VPC while remediation is performed	68.7%	16.4%	11.9%	97.0%
ACL changes to block IPS	65.7%	17.9%	11.9%	95.5%
Identify and rescope IAM roles affected	61.2%	19.4%	14.9%	95.5%
Re-creating and redeploying API keys	67.2%	19.4%	9.0%	95.5%
Creating memory images	70.1%	17.9%	6.0%	94.0%
Automated VPC network security	59.7%	19.4%	13.4%	92.5%
Reimage or restore compromised machines from snapshots	58.2%	16.4%	17.9%	92.5%
Remotely manage virtual machine fleets to kill rogue processes	65.7%	14.9%	10.4%	91.0%
Update policies and rules based on IoC findings and lessons learned	71.6%	7.5%	11.9%	91.0%
Re-creating system configurations to determine what things were changed	53.7%	16.4%	19.4%	89.6%
Other	7.5%	3.0%	10.4%	20.9%

Team Composition and Budgets

It's clear that organizations are using IR services more widely to support internal IR teams than are using managed security service providers (MSSPs). In only 25% of the cases, one to four members of an MSSP support a case, whereas in 50% of cases, an IR service will use one to four members. These are the same figures that internal organizational functions, such as IT operations, development, security group and others, use to support IR cases in the cloud.

It's good to see that internal IR teams also consider internal functions as part of their core team, according to 52% of respondents. Including members such as DBAs, developers or system administrators in IR cases supports potential synergies, especially when organizations train to ensure that they're developing competency in-house.

IR is seeing an increase in planned budget allocations in the next 12 months, but not by significant percentages (see Figure 9). It's pretty even across the board from previous periods, but in general, more budget is being allocated to not only IR, but also to cloud IR and cloud IR using outside services.

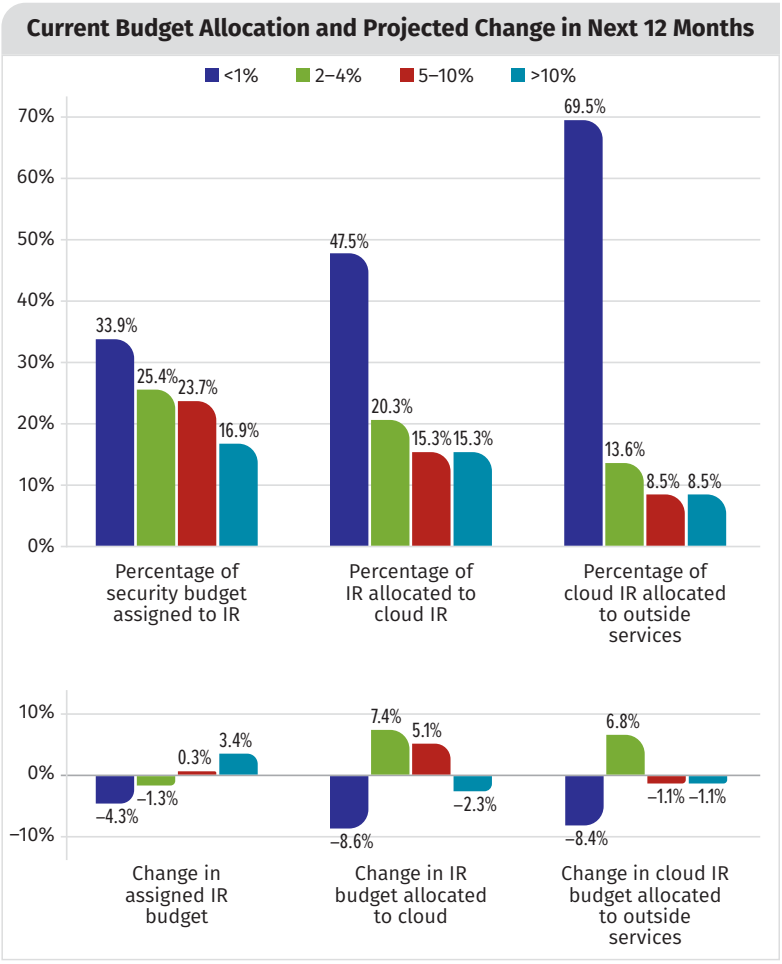

Figure 9. Current and Projected Budget Allocations

Conclusions

It's fair to say that the cloud offers great opportunities. Many of these opportunities come with risk, but the risk might be lower than the alternative. For example, the use and value of serverless functions have proven very useful. Serverless functions provide great value for developers, but because organizations might not fully understand the technology and the risk that comes with it, they might hold off on implementation.

What Is the Alternative?

One alternative to serverless deployments is to deploy a full server, which is typical of an on-premises deployment. At least now it's familiar and more comfortable to operations because it relates to what we're already doing on-premises. We don't want to stifle or stop innovation, at least not in 2020, so we'll go with the traditional route.

What this means is: We sought to solve before we understand. A full server is generally a more vulnerable environment than a serverless environment. The server is now filled with more attack surface than a potential serverless environment—more granting of privileges, a greater number of considerations in terms of patch management, and networks that now need to cater for segmentation.

With a serverless deployment, the attack surface is very small. Because of the tiny attack surface and least amount of privileges given to the serverless function, the harm caused by a mistake isn't usually dire. But if the same function were compromised on a traditional server, the consequences might be critical. Attackers can try all kinds of privilege escalation, install keyloggers, and pillage and pilfer just as they always have. Other aspects, such as configuration errors on the server, could pose vulnerabilities to the server itself—thus the overarching architecture.

So, with the marked lack of skills and people, should we still pounce on new opportunities that the cloud offers? The case is larger than a micro-perspective of deployment of services. Looking at the bigger picture, we now have a cloud provider supporting us with essential and core tools of any healthy IT operation, such as segmentation and the least amount of privileges. Even if we don't necessarily have a full understanding of *how* things work, our provider is likely to provide the good old IT-operations security measures we need to secure our services. This will potentially bring somewhat of an equilibrium between the continuous drive for innovation and the definitive need for security in our operations.

Stifling Innovation Is Not an Option

Information security cannot be a hinderance to continuous development of the organization and taking advantage of the potential ahead. The cloud typically requires different efforts in our responses than an on-premises network. It's important to gain visibility into how mature organizations are, measure that maturity, and understand how we deal with incidents in the cloud. We must ensure that we can identify gaps and start addressing them rather sooner than later. Later might be too late!

While the survey results reveal compromising factors in many organizations, the cloud providers tend to support us with capabilities that enable us to respond more easily and efficiently than before. If the choice is between hosting your own vs. sourcing parts of the responsibility to a cloud provider, for many, it's proven to be the latter.

Taking Back the Advantage

Considering that survey results strongly indicated a lack of skills, we are happy to see providers continuing to extend their training and offers in terms of their platforms. Budgets and adequate staff are hard to deal with in the short term, but training offers a potential way to gain an advantage. The trends in budgeting appear to be:

- An increase in IR investment
- A modest increase of 2% to 10% allocated to cloud IR
- A possible increase of 2% to 4% for cloud IR budget for outside services

Finally, more and more technology providers can provide their offerings in the cloud, providing services that are familiar and already integrated with on-premises solutions. Without a doubt, we need to start assessing cloud environments as we have with our on-premises solutions. And, at the very least, cloud often offers more advantages to defeat the adversaries before they're successful.

About the Author

SANS instructor **Chris Dale** teaches SANS SEC504: Hacker Tools, Techniques, Exploits, and Incident Handling. As founder and principal consultant at River Security, a specialist company operating out of Europe, Chris brings significant security expertise and a background in system development, IT operations and security management. This broad experience in IT is advantageous when managing penetration tests, incidents and while teaching. He has helped Fortune 500 understand their security challenges, contracted in government initiatives on securing democratic processes of the country, participated in expert groups in solving industry sector cyber related challenges country wide and helped build successful companies.

Sponsor

SANS would like to thank this survey's sponsor:

ExtraHop
Rise Above the Noise.