DARKReading black hat OMDIA

DECEMBER 2022

The Promise and Reality of Cloud Security

The head-spinning rush to the cloud in the wake of the COVID-19 pandemic laid bare a new category of security risks that has forced both enterprises and cloud providers to adapt their security practices.

Brought to you by
informa tech

This report – a compilation of Black Hat research, Dark Reading reporting, and Omdia analysis – explores how cloud security is rapidly evolving.

INSIDE:

DARK READING ANALYSIS

The Promise and Reality of Cloud Security

The head-spinning rush to the cloud in the wake of the COVID-19 pandemic laid bare a new category of security risks that has forced both enterprises and cloud providers to adapt their security practices.

By Kelly Jackson Higgins, Editor-in-Chief, Dark Reading

It wasn't long ago that many businesses shied away from relocating their applications and data to the cloud for fear of losing control of both their IT processes and data security. Cloud adoption boomed amid the pandemic as organizations fast-tracked their cloud plans, and, now, cloud security is getting its first real stress test.

During the course of the COVID-19 pandemic, cloud services matured as they gradually became the norm for many enterprise applications and services. However, the cloud's newfound popularity brought a fresh set of attack vectors for enterprises. In addition, security researchers who have been poking holes in the security of popular cloud services such as Amazon Web Services (AWS) and Microsoft Azure warn that cybercriminals and nation-state hacking teams are making their way to the cloud to target their victims. Cloud services promise stronger security — and deliver it in practice — but the reality is that cloud services have their own security vulnerabilities. There can be software flaws in cloud applications and servers, and, in some cases, weak default security controls can pose problems.

The good news: There is a whole a new generation of cloud security tools built specifically for cloud systems and accounts. Emerging technologies include cloud workload protection platforms (CWPP) for detection and response and cloud security posture management (CSPM) for proactive security. Cloud providers also have upped their security tool game to provide organizations with the tools they need to secure their own cloud environments. The bad news: Enterprise adoption of cloud security tools still lags in many corners.

The bad news: Enterprise adoption of cloud security tools still lags in many corners.

More than half of organizations now run more than 40% of their workloads in public cloud services, and organizations plan to increase their cloud workloads during the next year, according to a recent Cloud Security Alliance and Google Cloud survey. At the same time, the survey finds, the majority of organizations (85%) today don't run any cloud-asset discovery tools to inventory what they have running in their cloud services. Rather, they are sticking with manual methods, leaving them with perilously little insight and visibility into the devices, users, applications, and activity across their cloud-based operations.

Chaos and Confusion

The race to the cloud has exposed inherent security holes and vulnerabilities in popular infrastructure-as-a-service platforms. Security researchers have rooted out major flaws in AWS and Azure, including vulnerabilities in Azure's Open Management Interface (OMI) that could be used by an attacker for remote code execution and privilege escalation, as well as a flaw dubbed ChaosDB that gave Azure users full administrative access to other customers' Cosmos DB instances.

Ami Luttwak, and Shir Tamari, researchers from Wiz, discovered the Azure flaws, which Microsoft has since fixed.

The researchers initially realized they were onto something big when they found ways to break the isolation among different customers' AWS accounts. It was cloud customers' worst nightmare come true: An attacker could read and access data housed in their S3 cloud storage buckets — and even move the data to other storage buckets — for malicious purposes. The researchers, who first shared their groundbreaking research at Black Hat USA in 2021, were convinced the cross-account security flaws weren't isolated to AWS. Their theory was confirmed when they found similar issues in Azure, with the ChaosDB flaw in Cosmos.

But what struck Luttwak and Tamari most about their cloud vulnerability research was the opaque and inconsistent way cloud providers fix and alert their customers about security issues and updates. There's no standard, CVE-style process or repository for cloud providers and researchers to share vulnerability details, so cloud customers often aren't aware if or how a vulnerability affects their particular cloud service — if they are aware that the vulnerability even exists in the first place. Unlike in the pre-cloud era, when enterprises managed their own security, the cloud model comes with a shared responsibility relationship that is much less defined and could result in vulnerability patches and updates falling through the cracks.

In fact, patching and vulnerability management in the cloud is different than it is with client-server software. Cloud providers often handle the fixes, but, in some cases, customer response is required — for example, making a configuration change to reflect the new update.

Historically, cloud providers alerted users about fixes via email or a variety of other methods. In hopes of standardizing the process of cloud vulnerability reporting, the researchers helped launch a community-driven database. Called cloudvulndb.org, the database logs known security issues related to cloud service providers and provides information on how to mitigate or address the issues. Researchers from cloud security vendors, including Lightspin, Microsoft, Orca Security, and Wiz, have been populating the cloud vulnerability repository with details on the status of the latest cloud flaws, as well as the affected cloud services.

The cloud model comes with a shared responsibility relationship that is much less defined and could result in vulnerability patches and updates falling through the cracks.

IAM in Danger

At the core of cloud services is identity and access management (IAM). IAM encompasses not only user identities but also the requisite machine identities that manage everything from cloud API access to integrating various cloud applications and profiles via service accounts. Machine identities use digital certificates and cryptographic keys to communicate among one another, and they typically hold wide permissions, which makes them especially attractive and valuable to attackers.

In their Black Hat USA 2022 presentation in August, "IAM the One Who Knocks," researchers Igal Gofman and Noam Dahan of Ermetic warned that the complexity of cloud IAM makes it tough for organizations to manage their cloud identities, much less locate all of them. And AWS's, Azure's, and Google Cloud's services each handle and approach identity differently, so managing identities in a multicloud environment can be even more complicated.

Gofman and Dahan say the biggest risk for cloud identity security is attackers abusing the IAM "read" function and assigning their own permissions, which would allow them to steal data and escalate privileges. The researchers recommend utilizing the logging features offered by cloud providers to help map an organization's cloud identities, including the actions users are performing and the functions and resources to which they have access.

Key Cloud Security Tools

Omdia's 2022 Decision Makers Survey shows that more than 40% of respondents ranked the cost of cloud security tools as a top concern, followed by cloud services offering insufficient security functions (34%). Meanwhile, there are several CSPM commercial and open source tools available today, and Omdia notes that those organizations piloting or running CSPM in production are finding the technology useful for proactive security configuration and compliance issues. Organizations with the most cloud security experience that have CSPM in production also flagged concerns relating to securing data in the cloud, responding to cloud security incidents, and the lack of cloud technology chops within their security teams.

So, where does zero trust fit into the cloud? According to Omdia, cloud permissions management, or CPM, can provide zero-trust capabilities in cloud environments by tracking and right-sizing so-called identity "permission sprawl" in the cloud. That's basically a consequence of how cloud identities automatically inherit permissions that they don't actually need once they join another cloud group. CPM can stop or revoke overly broad or unneeded permissions, and monitors cloud instances for any additional expanded and risky permissions.

There's also emerging cloud detection and response (CDR) technology, which integrates cloud security into the security operations center's incident response process. Omdia's take: CDR shows promise for cloud security.

Read on for a deeper look at the promise and challenges surrounding cloud security — with news analysis from Dark Reading, research from Black Hat, and data and industry analysis from Omdia.

About the Author: *Kelly Jackson Higgins is the Editor-in-Chief of Dark Reading. She is an award-winning veteran technology and business journalist with more than two decades of experience in reporting and editing for various publications, including Network Computing, Secure Enterprise Magazine, Virginia Business magazine, and other major media properties. Jackson Higgins was recently selected as one of the Top 10 Cybersecurity Journalists in the US, and named as one of Folio's 2019 Top Women in Media.*

What Lurks in the Shadows of Cloud Security?

Organizations looking to get ahead in cloud security have gone down the path of deploying CSPM tooling with good results. Still, there's a clear picture that data security and security operations are next key areas of interest.

By Fernando Montenegro, Senior Principal Analyst, Cybersecurity, Omdia

As an industry, we're now at a point where we don't have to convince anyone that we have a massive digital dependence on cloud technologies and that securing cloud deployments is a key initiative for most organizations. There is widespread availability of cloud security posture management (CSPM) tooling — commercial- and community-driven alike — and CSPM itself is being incorporated into new tooling coming under the heading of cloud-native application protection platforms (CNAPP).

Many cloud security conversations focus primarily on making sure clouds are configured properly, but there is much more to cloud security than that. Much like traditional security is much more than patching, there is more to cloud security than configuration.

As we analyze cloud security trends, we recently collected survey data under our Omdia Decision Makers survey and found interesting results that highlight these conclusions.

The figure on the next page plots responses to the question "What are your top concerns in relation to cloud security?" The bars on the left show the aggregate view (n=186). We can clearly see a major concern with cost of security tooling, then other concerns aggregated together, including security tooling functionality, responding to events, data security, and others.

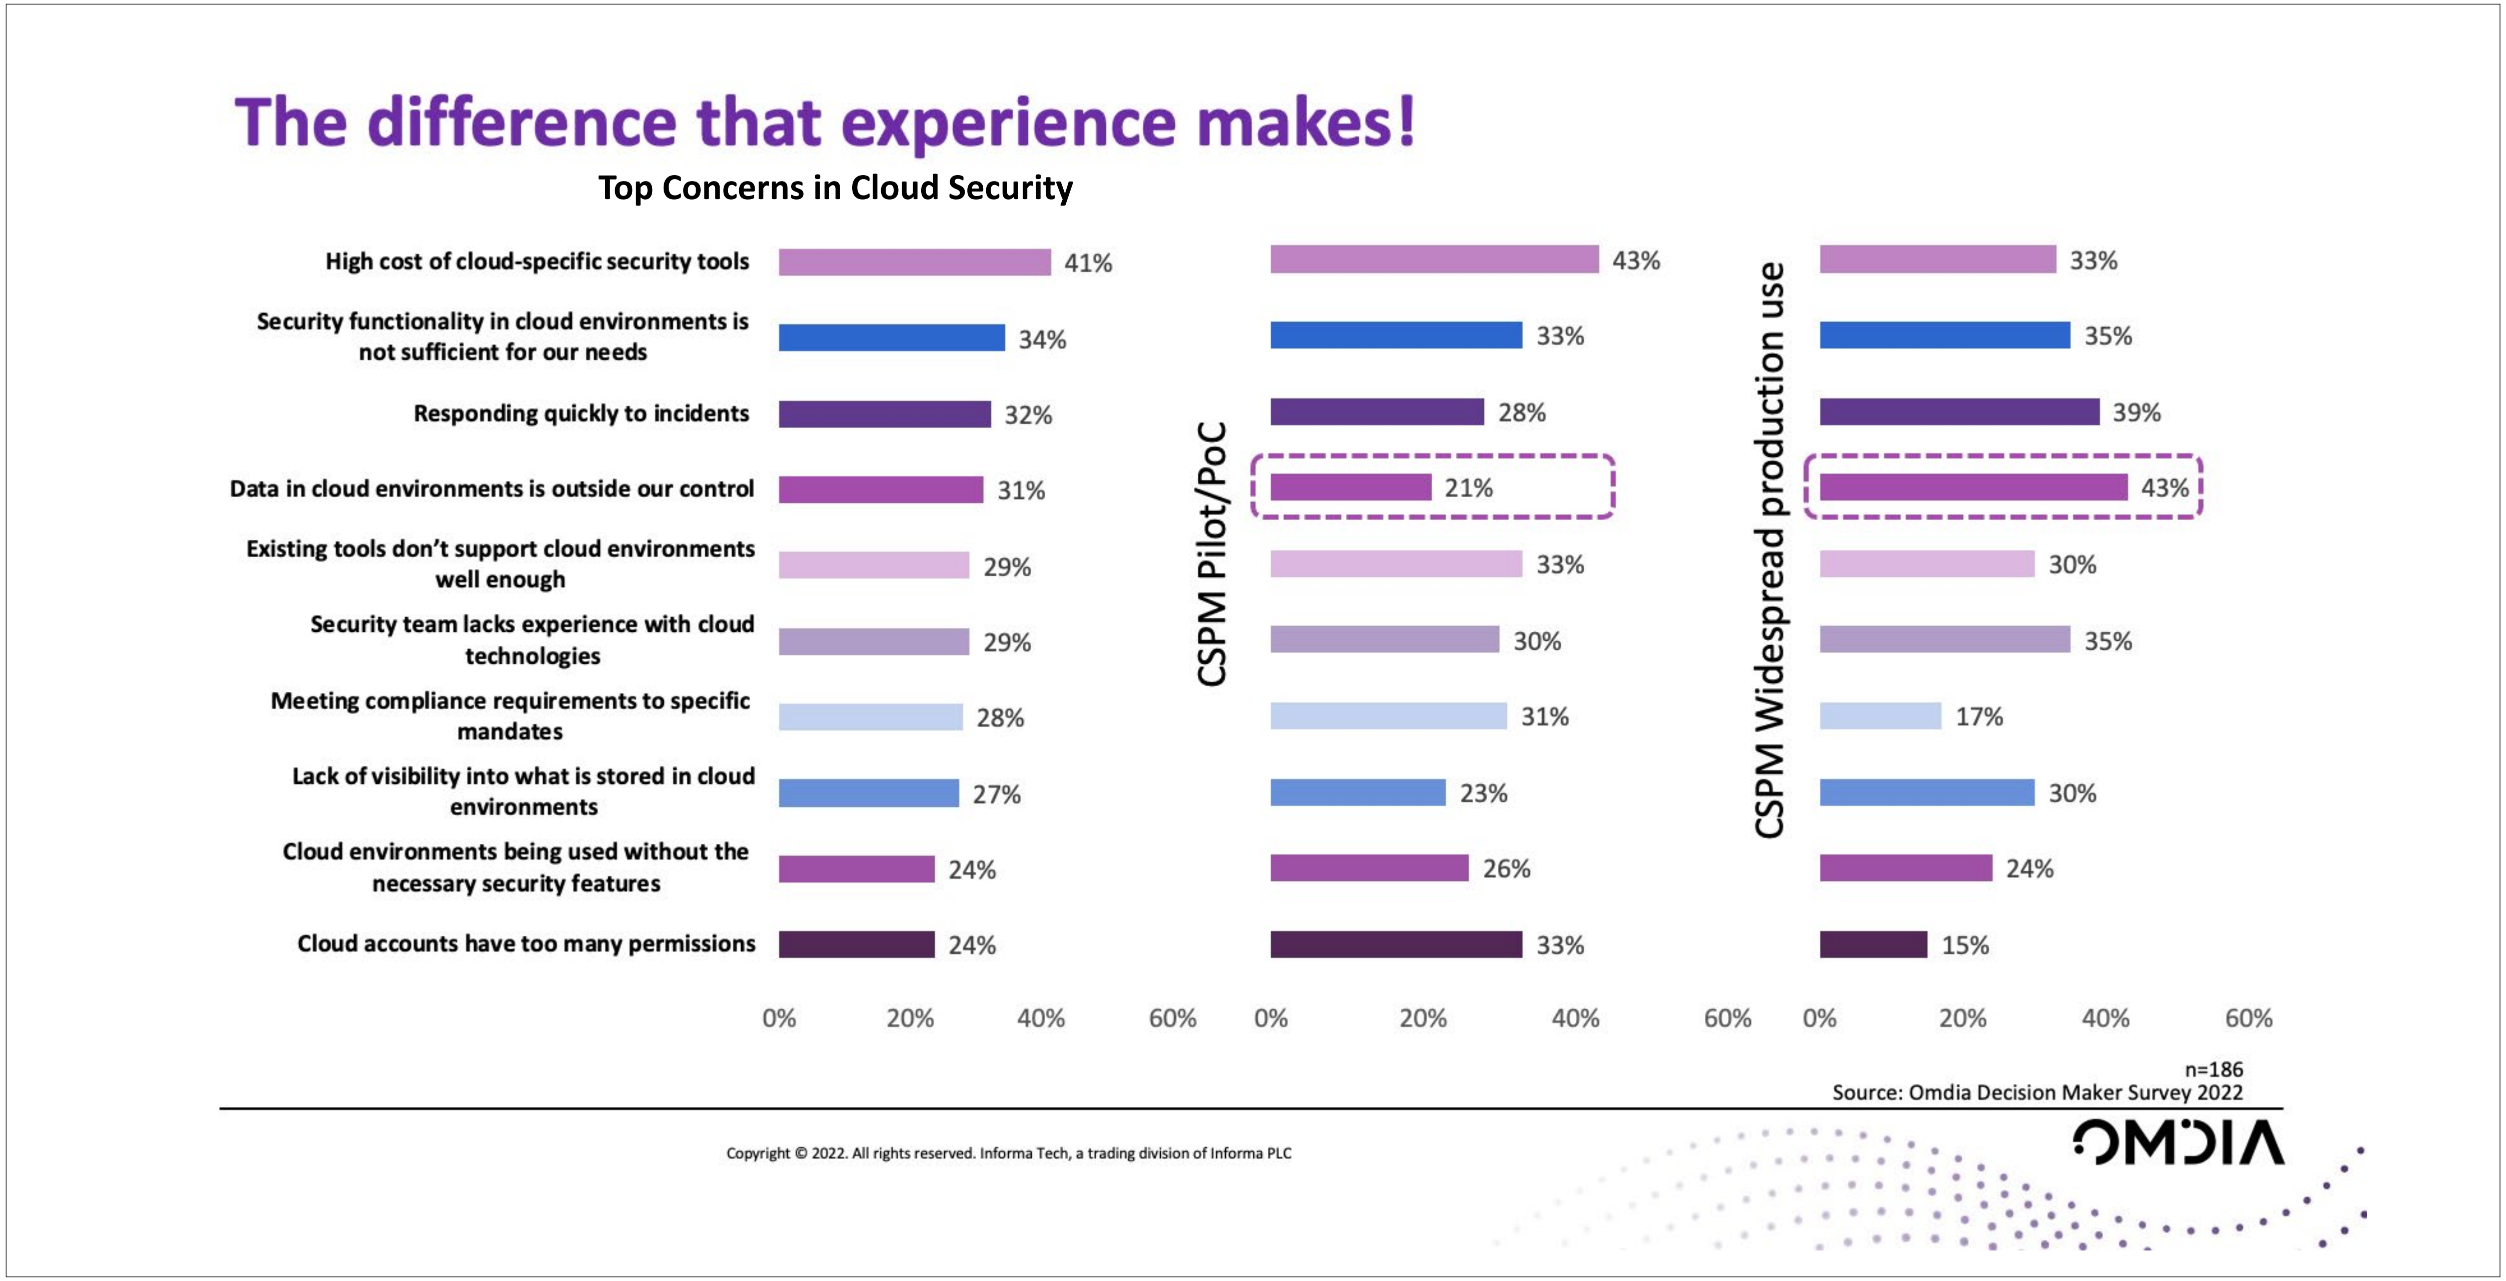

The difference that experience makes!

Top Concerns in Cloud Security

n=186
Source: Omdia Decision Maker Survey 2022

That said, we also segmented the population into two groups based on their response to a previous question about how advanced their deployments of CSPM tools were. Our ongoing industry interactions with multiple stakeholders point to CSPM tools as the type of tool most often associated with "cloud security" conversations, so we chose CSPM deployment experience as a proxy for cloud experience. Statistician George Box is famous for saying "all models are wrong, but some are useful," which we think is relevant here; there's no implication of causation, but some interesting variations show up in the response data.

For those that have what we consider "low" experience with cloud security (n=61) by virtue of having CSPM deployments in the pilot or proof-of-concept stage, concerns around cost are even more pronounced, as are concerns about cloud permissions and a slight bump for concerns about compliance.

For those that have more cloud security experience (n=54) — those that responded that they have deployed CSPM in widespread production use — responses shifted significantly. Now, concerns about data security are much more pronounced, as are concerns about how to respond quickly to incidents, with additional notable concerns regarding the ever-present skills gap in terms of cloud technologies.

Heightened Concerns

Our interpretation of this data is that customers are indeed seeing value from the CSPM tooling for configurations and compliance, but they now have heightened concerns on data security, security operations, and making sure their teams are skilled in cloud technologies. These concerns are already lurking in the shadows, and once CSPM clears the way of handling the more visible security configuration/compliance concerns, these issues come to the forefront.

The responses uncovered here point to interesting directions for future inquiry. It is increasingly clear that data security presents a key area of concern. What are the ways one gets to data? One way is via direct access to the data stores themselves. This is the provenance of cloud configuration (CSPM) and the increasingly popular DSPM (data security posture management) category. Another is getting access via the very APIs provided by the company; this then leads down a path of paying close attention to API security.

For security operations, the path forward appears to include more considerations about how to incorporate cloud security use cases in SOC response flows. Dubbed cloud detection and response (CDR), this is also a promising area of research that we're watching.

For end users, this means being ready to address these categories soon. For vendors, understand that there is much more to customer demand for cloud security than CSPM — or even CNAPP — alone.

About the Author: *Fernando Montenegro is a Senior Principal Analyst on Omdia's cybersecurity research team. He focuses on the Infrastructure Security Intelligence Service, which provides vendors, service providers, and enterprise clients with insights and data on network security, content security, and more. Fernando's experience in enterprise security environments includes network security, security architecture, cloud security, endpoint security, content security, and antifraud.*

Risk Disconnect in the Cloud

New Cloud Security Alliance (CSA) and Google Cloud study shows many enterprises struggle to measure and manage risk in their cloud workloads.

By Kelly Jackson Higgins, Editor-in-Chief, Dark Reading

Cloud adoption may be hopping, but many enterprises still wrestle with how to identify and manage their security risks with these services.

A new study conducted by the Cloud Security Alliance (CSA) and Google Cloud underscores that while the cloud ideally could help bolster security for organizations, many aren't adeptly handling their risk management in the cloud just yet. "Organizations are not taking advantage as aggressively of the capabilities to have a more secure environment" with cloud, says Jim Reavis, CEO of the CSA. "They're not being as proactive in monitoring and managing risk."

Interestingly, it appears many organizations may not know for sure the extent of their cloud adoption. Some 51% say that they now run 41% of their workloads in the public cloud, but it turns out most of them (85%) are not using cloud discovery tools to quantify that but, rather, estimating their use via manual methods. Those who use discovery tools including a cloud access security broker, or CASB (15%), to map their cloud workloads report 31% more cloud usage than those who performed manual assessments — a clue that most organizations

relying on manual tracking don't have a complete inventory of what's running in their cloud services, according to the study.

"You can't manage the risk of things you don't know about. The basic things lead to either breaches or data exposure, exfiltration, or a ransomware attack if you are not keeping your cloud assets updated and there are gaps in your usage of cloud," Reavis notes. But the cloud offers a better way to manage assets, he says, than traditional IT networks.

"There are tools there," and automated ways to detect and secure cloud assets, he says.

The study confirms a significant rise in cloud adoption. The average number of software-as-a-service (SaaS), platform-as-a-service (PaaS), and infrastructure-as-a-service (IaaS) services used by organizations was more than 147, up from 38 in 2020. Some 66% of organizations say they have 100 or fewer services; 32%, from 101 to 999; and 3%, 1,000 or more services.

The most commonly used infrastructure-as-a-service (IaaS) cloud platform is Azure (70%), followed closely by AWS (65%), and then Google Cloud at 24%, according to the study.

"Enterprises interviewed intend on increasing their workloads in the cloud over the next 12 months. With enterprises continuing to add production in the cloud and using more cloud services, managing cloud and digital assets will be critical in the management and measurement of risk in the cloud," according to the report.

The goal of the study was to gauge organizations' challenges of risk management in public cloud services, and Google and the CSA gathered survey data as well as interviews in 2021 with 600 IT and security professionals.

Cloud Escape

While the cloud is becoming more pervasive for IT operations, there has not been a correlation or increase in data breaches, Reavis notes.

To date, nearly all publicly disclosed breaches in the cloud have stemmed from misconfigurations, not cyberattacks, says Phil Venables, CISO at Google Cloud. "To prevent and address the risk of misconfigurations and compliance violations earlier in the development process, security leaders have started to embrace security as code to achieve the speed and agility of DevOps, reduce risk, and more securely create value in the cloud," Venables says.

For its part, Google offers a series of blueprints for its customers to help avoid misconfigurations and other cloud mistakes, such as its Risk and Compliance as Code (RCaC), Secure Foundations guide, and Cloud Architecture Center, for example.

"Blueprints help our customers rapidly configure cloud environments in a secure and compliant manner," notes Venables. "And ultimately, this level of secure hygiene helps prevent misconfigurations becoming a security risk or attacker entry point to cloud workloads."

According to the report, some 70% of organizations in the study say they don't have solid processes for mapping risk to their cloud assets. A tiny percentage — 4% — report that they have "highly effective" risk management in the cloud. Slightly more than 20% use cloud data-classification tools.

Meanwhile, the main security worries over applications in the cloud include loss of sensitive data (64%), improper configuration and security settings (51%), and unauthorized access (51%).

About the Author: ***Kelly Jackson Higgins is the Editor-in-Chief of Dark Reading. She is an award-winning veteran technology and business journalist with more than two decades of experience in reporting and editing for various publications, including Network Computing, Secure Enterprise Magazine, Virginia Business magazine, and other major media properties.***

The Role of Cloud Permissions Management in Cloud Security

Organizations looking to get ahead in cloud security have gone down the path of deploying Cloud Security Posture Management (CSPM) tooling with good results. Still, there's a clear picture that data security and security operations are next key areas of interest.

By Rik Turner, Senior Principal Analyst, Cybersecurity, Omdia

Cloud adoption continues apace around the world, driven by the digital transformation that has itself been turbocharged by the coronavirus pandemic. As application infrastructures move to the cloud, however, the need to secure corporate assets residing in cloud environments has become an ever-increasing requirement. While protection of workloads and data at runtime is an essential part of this process, there is growing interest in preemptive approaches that adopt an *a priori* stance to cloud security, reducing the attack surface before exploits can even take place.

Cloud Expands the Attack Surface

There can be little doubt that cloud is the direction of travel for most application infrastructures, whether they belong to enterprises or noncommercial entities. In adopting cloud, organizations frequently expand their attack surface. They have likely invested heavily over the years to protect their on-premises assets; however, cloud security is effectively a different ballgame, in which these traditional infrastructure security products frequently struggle to be relevant. New approaches to the problem are needed, not least because entirely new, cloud-native attack vectors have sprung up, requiring different types of technology to address them.

Cloud workload protection platforms (CWPP) provide reactive, runtime protection. The reactive — that is, "detect and respond" — approach in cloud security is embodied in so-called CWPP technology. This type of technology is of course essential, since attacks will almost certainly be launched against cloud assets, which makes detecting and blocking them a *sine qua non* of operating in such environments.

However, with the threat landscape continuing to flourish and skilled cybersecurity staff remaining hard to find and retain, interest is growing in more proactive approaches, which aim to reduce an organization's attack surface before any attacks happen. One example of this more proactive type of security is cloud security posture management (CSPM). This is technology that inspects assets once they are in production to detect signs of compliance or security drift.

A still more aggressively proactive approach to security is embodied in the trend now widely referred to as zero trust. This is a mindset or, if you prefer, a philosophical stance on security, which can be summed up as "never trust, always verify" — to which Omdia nowadays adds a third dimension, "… and continually monitor."

In cloud environments, an early example of zero trust was microsegmentation technology, which imposes workload isolation and enforces strict access policies, both for human and nonhuman identities. CPM — cloud permissions management — is a more recent development, and another way of achieving zero trust in the cloud. One of the problems in securing cloud environments is so-called sprawl. In essence, this issue arises because of the ease with which new cloud instances, whether they be workloads or data stores, can be spun up.

Similarly, cloud frequently creates the problem of permission sprawl. Identities can even gain new permissions simply by joining a particular group, even though they may have no need to access a specific asset. And, of course, IT operations don't necessarily keep track of all the permissions attached to every nonhuman identity in the environment.

CPM Addresses Permission Sprawl

CPM has arisen to address this situation. It aims to right-size a company's permissions estate; it enables it to curtail or even revoke permissions that are deemed excessive or simply unnecessary and, once that process is complete, to monitor the environment in an ongoing manner to detect any signs of renewed sprawl.

CPM begins by drawing up a full inventory of the extant permissions within an organization's cloud estate. Once that process is complete, a CPM platform carries out an analysis of all the permissions listed against the various identities to determine which ones are excessive or simply surplus to requirements. It then makes recommendations for how the permissions estate can be curtailed, with individual access rights being reined in or removed altogether. Some CPM platforms can also go further, actually performing the remedial action they have recommended in an automated fashion if the customer is happy to enable that feature. The figure below provides a diagrammatic overview of CPM.

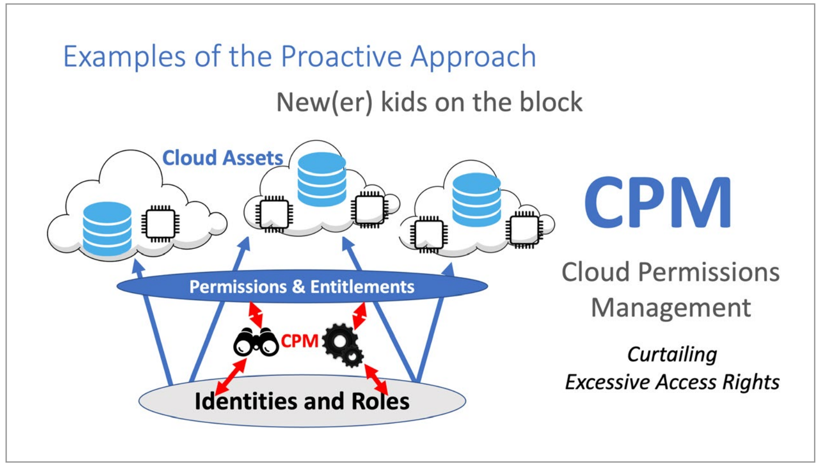
Examples of the Proactive Approach
New(er) kids on the block
Cloud Assets
Permissions & Entitlements
CPM
Identities and Roles
CPM
Cloud Permissions Management
Curtailing
Excessive Access Rights

Cloud permissions management. Source: Omdia

The Future of CPM

Omdia does not believe that CPM will remain a standalone capability, taken to market by vendors dedicated exclusively to its further development, for very long. It makes sense for CPM to be part of a broader portfolio of security capabilities, whether specifically for the cloud or for hybrid environments spanning both the cloud and on-premises infrastructure.

There are obvious synergies between CPM and well-established technologies, such as identity governance and administration (IGA) and privileged access management (PAM). The former seeks to manage access rights within an organization by setting identities up correctly to begin with and to guarantee their expungement when someone leaves, while the latter focuses on the rights of access to sensitive or confidential assets within an enterprise, aiming to impose the least privilege wherever possible. Equally, there are affinities between CPM and CSPM, in that both are proactive technologies that don't wait for a breach or an attack to happen before acting.

About the Author: *Rik Turner is a principal analyst in Omdia's IT security and technology team, specializing in cybersecurity technology trends, IT security, compliance, and call recording. He provides analysis and insight on market evolution and helps end users determine what type of technology and which vendor they should be pursuing.*

New Vulnerability Database Catalogs Cloud Security Issues

Researchers have created a new community website for reporting and tracking security issues in cloud platforms and services — plus fixes for them where available.

By Jai Vijayan, Contributing Writer, Dark Reading

Organizations traditionally have struggled to track vulnerabilities in public cloud platforms and services because of the lack of a common vulnerability enumeration (CVE) program like the one that MITRE maintains for publicly disclosed software security issues.

A new community-based database launched this spring seeks to address that issue by providing a central repository of information on known cloud service-provider security issues and the steps organizations can take to mitigate them.

The database — cloudvulndb.org — is the brainchild of security researchers at Wiz, who for some time have been advocating the need for a public catalog of known security flaws on platforms and services run by the likes of AWS, Microsoft, and Google. The database currently lists some 70 cloud security issues and vulnerabilities that security researcher Scott Piper had previously compiled in a document on GitHub titled "Cloud Service Provider Security Mistakes." Going forward, anyone is free to suggest new issues to add to the website or to suggest new fixes to existing issues. The goal is to list issues that a cloud service provider might have already addressed.

Centralized Vulnerability Repository

"The centralized database can help organizations review all past security issues in their [cloud service provider] at any time and check if they have not applied necessary remediation actions," says Alon Schindel, director of data and threat research at Wiz. "For example, organizations can check if they were using a certain service during a critical security issue's exploitability period and use the recommended detection methods — if available — to check if they were affected."

For now, the vulnerability database site does not have a system in place to automatically notify users when new security issues are added to it. But the goal is to add an RSS feed or mailing list for that purpose, says Schindel, one of the maintainers of the new database.

Schindel — like many other researchers — has noted how the lack of a formal and standardized system for publicly recording cloud security issues, and sharing information about them, is heightening risks for organizations. In a blog post last November, Schindel and another Wiz researcher pointed to vulnerabilities — such as one dubbed ChaosDB in Microsoft Azure and another called OMIGOD in Microsoft Azure — as specific reasons why a cloud vulnerability database has become a critical industry necessity. Both vulnerabilities were serious. And unlike many cloud vulnerabilities, the responsibility for mitigating risk with both vulnerabilities rested not just with the cloud provider but also with their customers.

ChaosDB impacted four Azure services and gave users overly permissive access to storage buckets belonging to other cloud tenants. OMIGOD was a set of four flaws in OMI, a Microsoft cloud middleware technology, that enabled remote code execution and privilege escalation. Though Azure and Microsoft addressed the vulnerabilities promptly, many organizations using the affected services had limited information on the changes they needed to make to address them, the Wiz researchers said.

“Typically, cloud service provider security issues do not have a patch in the traditional sense, as issues are fixed internally by the CSP without the need for any manual user action,” Schindel says. But no CVEs mean that there are no industry conventions for assessing severity, no proper notification channels, and no unified tracking mechanisms.

“This means that it’s difficult for a cloud customer to answer otherwise simple questions like, ‘Is my environment currently vulnerable to this?’ or, ‘Was it ever vulnerable to this?’” he adds.

Inconsistent Practices

Currently, all major CSPs accept responsibly disclosed vulnerabilities, and some have an official bug bounty or vulnerability reward program in place. Occasionally, a cloud service provider might even publish details of a fix they might have developed for a reported security vulnerability. However, there is little consistency among the various providers, Schindel says.

“Notification channels vary; vendors usually email affected customers only or send them a notification through a service health system,” he says.

Wiz has been unable to find any consistency in the publication cadence of security issues of the different CSPs, though Microsoft usually included fixes for Azure vulnerability in its monthly patch release cycle.

Wiz will maintain the new site, though anyone is free to contribute to it. The goal is to try and get major CSPs to engage with the effort or to use the site to provide more transparency around vulnerabilities discovered in their services. This can include information such as indicating the time periods during which a security issue might have been exploitable.

“We also hope that the value of such a database will help CSPs standardize their security issues publication processes,” he says.

About the Author: ***Jai Vijayan is a seasoned technology reporter with over 20 years of experience in IT trade journalism. He was most recently a Senior Editor at Computerworld, where he covered information security and data privacy issues for the publication. Over the course of his 20-year career at Computerworld, Jai also covered a variety of other technology topics, including big data, Hadoop, Internet of Things, e-voting, and data analytics.***

Why Some Cloud Services Vulnerabilities Are So Hard to Fix

Five months after AWS customers were alerted about three vulnerabilities, nearly none had plugged the holes. The reasons why underline a need for change.

By Karen Spiegelman, Features Editor, Dark Reading

It's a familiar story: A feature designed for convenience is used to sidestep security measures. In this presentation from Black Hat USA 2021, a pair of researchers show how they found three separate ways to hop between accounts on Amazon Web Services (AWS). Even though fixes for those vulnerabilities were released quickly, the holes reveal that cloud services do not offer the level of isolation expected. The long-term solution may mean changing how the cybersecurity sector handles CVEs.

For the first two isolation breaches, Wiz CTO Ami Luttwak and head of research Shir Tamari altered the path prefixes on AWS CloudTrail and Config to allow a user to write to another user's S3 bucket. The third method used the AWS command line to download files from another user's account via the serverless repository.

Sending the logs from several S3 buckets to one is intended for the convenience of an admin who runs several instances, Luttwak and Shir Tamari said in their presentation, "Breaking the Isolation: Cross-Account AWS Vulnerabilities."

Tamari added, "I learned from this behavior that CloudTrail can write to resources that are owned and managed in other accounts. And for me, a security researcher, there is a concern."

While Amazon did not have the power to fix the configurations for customers itself — because the fixes involved setting the source account you want, which only the users themselves can decide — it contacted all affected customers to explain the potential problem and how to fix it. Yet when Wiz went back after five months, it found that 90% of accounts had not applied the fixes.

Luttwak pointed out that the security team that AWS messaged often didn't get the warnings because of the sheer number of accounts they run. "How do you know this is an important fix to do?" he asked. "And the more we thought about it, the more we understood, this is a big, big problem."

As Luttwak said, "There's hundreds of services in AWS, and many of them are getting more and more cross account capabilities, because cross account is the main strategy today for organizations using AWS. So, the attack surface is just growing."

Watch the Black Hat video here.

About the Speakers: *Shir Tamari is an experienced security and technology researcher specializing in vulnerability research and practical hacking. Tamari is currently Head of Research of the cloud security company Wiz and also a member of the 5BC CTF team. Ami Luttwak is a serial entrepreneur, an experienced cybersecurity CTO, and a hacker at heart. Luttwak is mainly interested in cloud security and cloud exploits and understanding how the cloud is built to uncover its weaknesses. He is currently CTO of Wiz; before that, he led research as CTO of Microsoft cloud security and founded Adallom, a pioneering cloud security startup acquired by Microsoft in 2015.*

Black Hat USA 2022: Building Up IAM in a Multicloud World

In the cloud-first world, the security goal is to ensure only qualified users can access information across clouds.

By Karen Spiegelman, Features Editor, Dark Reading

The rise of multicloud environments brings with it the need to understand how to implement security policies across each cloud provider. The fact that each of the big three — Amazon's AWS, Microsoft Azure, and Google Cloud Platform — uses different nomenclature and configurations makes it that much more complicated to create a seamless and secure virtual network. A pair of researchers shared practical advice on how to secure one piece — identity and access management (IAM) — at Black Hat USA 2022.

Igal Gofman, Ermetic's head of security, and Noam Dahan, Ermetic's research lead, presented "IAM the One Who Knocks" at Black Hat USA 2022. "If there is one thing we would like you to take from this specific session, it is that IAM is the backbone service. It is the core service. It is the gateway that controls every access to your cloud resources, and it must be protected," Gofman emphasized.

Organizations have several reasons for using multiple clouds, as Gofman listed: adding in redundancy for better stability; reducing cost; taking advantage of multiple vendors' marquee features; or having conflicting platform requirements from different projects. But when you split resources among various clouds, he added, you need to be aware of and accommodate for the differences between how the platforms function.

"It's hard enough to be expert on one cloud platform," Gofman said. "But often we copy features and routines from one platform to another. And those may work differently from what we expect at the beginning."

Dahan then drilled down on the ins and outs of logging features from Azure, AWS, and GCP. Besides using logging for detection and incident response, he said logging is good for improving the permissions process. “In order to know whether you can take permissions away from someone, what you would usually do is try to examine the logs and see what they’re actually using, a sort of ‘use it or lose it’ philosophy,” he explained.

As Dahan put it, there are two main approaches to issuing permissions: sculpting from marble or from clay. Marble means starting with a full raft of permissions and then chipping away until you reach minimum necessary permissions; this can end up too permissive because you don’t want to remove too much. Clay means building up permissions until you have enough. Security staff likes this model, Dahan said, but developers hate it, because they don’t know what permissions they will need down the road. He recommended a hybrid approach of starting with a smaller hunk of permissions and then building up in places as needed.

The title of the talk comes from the TV series *Breaking Bad*, when science-teacher-turned-meth-kingpin Walter White reacts to a friend warning him that he’s in danger of someone coming to his door and killing him. White, incensed, asserts that he is the dangerous one by saying, “I am the one who knocks.” Perhaps IAM is the stand-in for White — it looks basic and unassuming, but underestimating its power is dangerous. Or maybe it’s just a turn of phrase.

Watch the Black Hat video here.

About the Speakers: *Igal Gofman is a Head of Security Research at Ermetic. Igal has a proven track record in cloud security, network security, research-oriented development, and threat intelligence. His research interests include cloud security, operating systems, and active directory. Noam Dahan is a Senior Security Researcher at Ermetic with several years of experience in embedded security. He is a graduate of the Talpiot program at the Israel Defense Forces and spent several years in the 8200 Intelligence Corps.*