orca security

EBOOK

How to Evolve Your Security Processes for a Cloud-Centric World

Introduction

The types of security risks that exist in the cloud are not fundamentally different from those that on-prem workloads face. Ransomware, software supply chain attacks, misconfigured access settings, and so on are examples of security challenges that you'll encounter in any type of IT environment.

In that sense, the migration of workloads to the cloud – which will be at the center of the IT strategies of 85 percent of businesses by 2025, according to Gartner[1] – hasn't drastically changed the nature of cybersecurity risks.

However, what has changed as businesses have moved to the cloud are the processes that security teams must embrace to protect cloud workloads. To put it bluntly, the approaches to discovering, assessing, and mitigating cybersecurity threats that are effective on-prem simply don't work in the cloud in many cases. Organizations need a fundamentally new approach to cyber risk management, even though the threats haven't fundamentally changed.

This eBook provides actionable guidance on how IT organizations broadly, and security teams in particular, can respond to this challenge. It begins by discussing the reasons why securing applications and data in the cloud requires a new set of approaches. It then details what those approaches entail and explains how to implement a security strategy that is just as effective for cloud workloads as for on-prem assets.

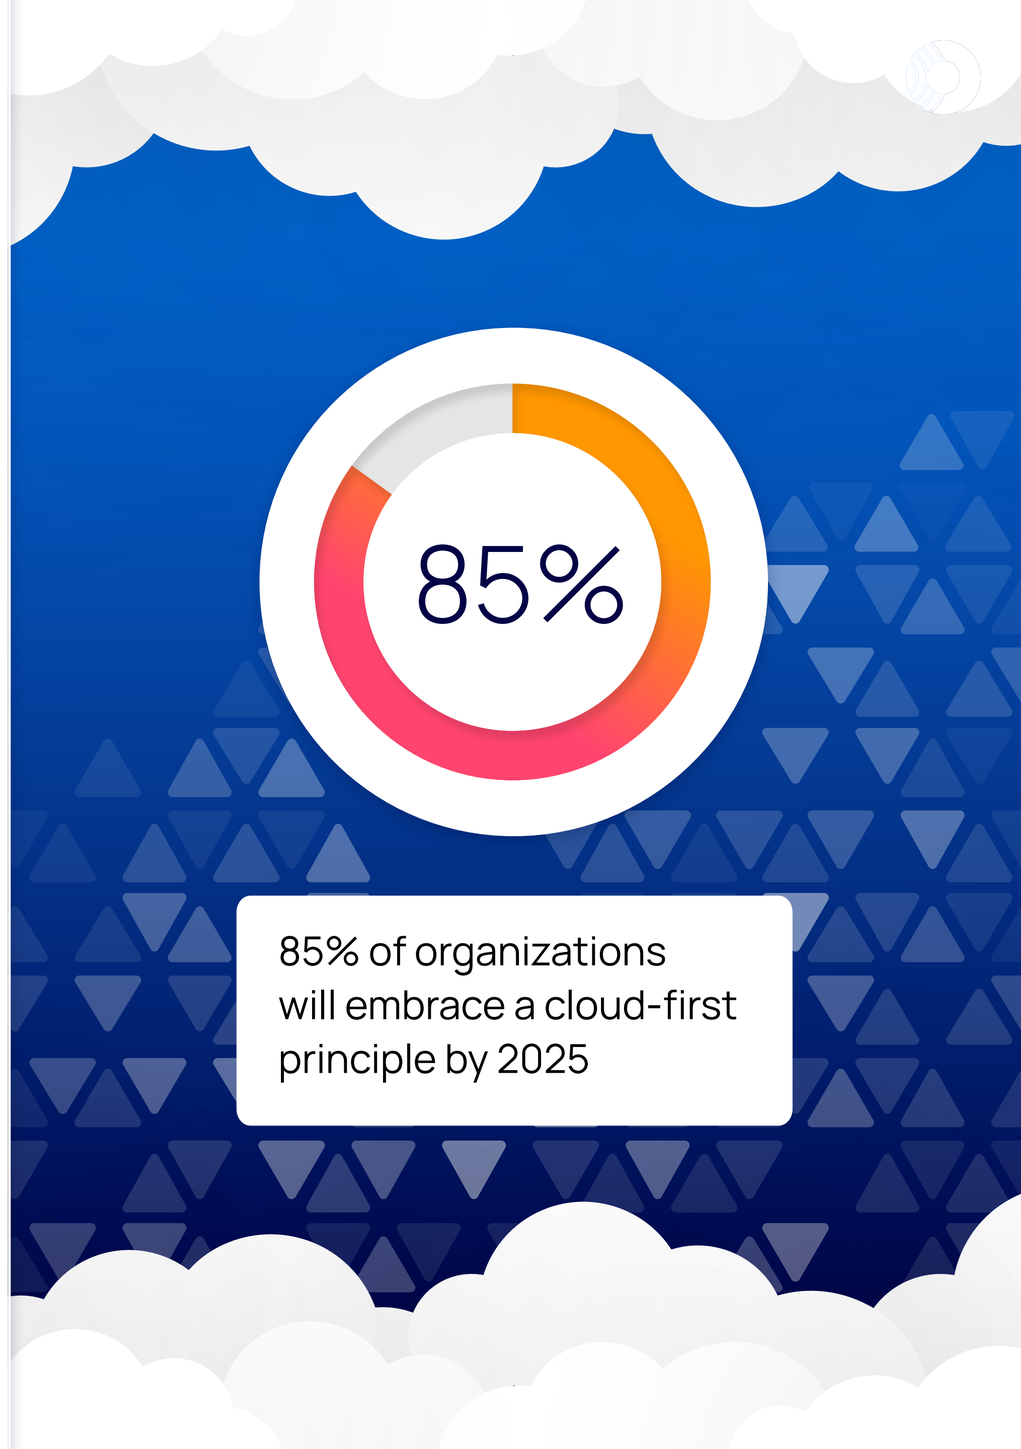

[1] Gartner, Inc., "Gartner Says Cloud Will Be the Centerpiece of New Digital Experiences," Laurence Goasduff, November 10, 2021.

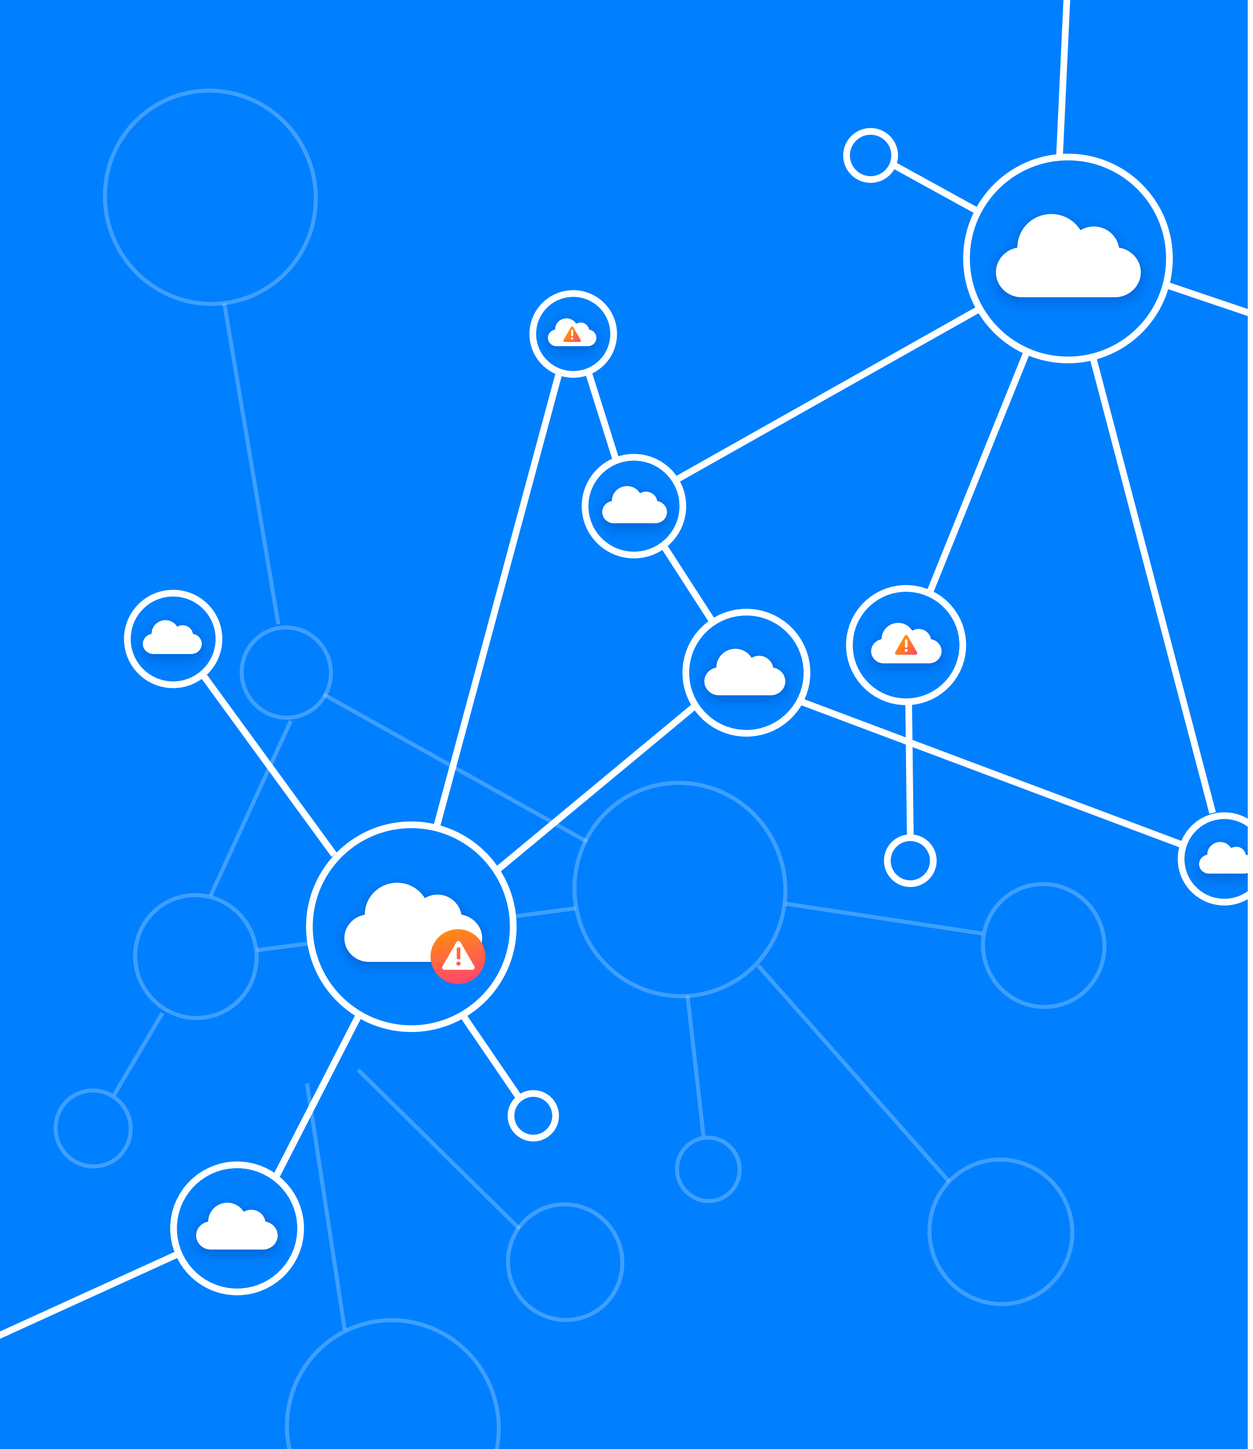

What Makes the Cloud Different From a Security Perspective?

At first glance, cloud environments don't always seem all that different from on-prem deployments, at least in some cases. Your typical cloud environment might consist of a set of virtual machines (VMs), which are similar to the collections of servers you'd run on-prem. And you deploy applications on top of those VMs in a fashion that is not drastically different from the deployment techniques that most teams use outside of the cloud.

But when you look beyond the surface, it becomes clear that the cloud is indeed different – and that the differences carry critical implications for security.

Cattle, Not Pets

One major difference between the cloud and conventional on-prem environments is that in the cloud, you can provision, launch, stop, and delete virtual servers almost instantaneously and without limit. This flexibility results from the massively scalable infrastructure that cloud providers offer and practices like Infrastructure-as-Code, which makes it easy to configure resources in a declarative and automated fashion. This agility leads teams to treat cloud-based servers as disposable assets – or as "cattle," to use a crude analogy that has gained popularity in the IT industry.

In contrast, on-prem servers demand careful setup and management practices. A failing or hacked server can't simply be torn out and replaced with a new one in a matter of minutes. On-prem servers are therefore akin to "pets," which IT teams must carefully manage and groom.

From a security perspective, the difference between cattle and pets means that in the cloud, addressing security issues requires the ability to modify the declarative and automated processes that teams use to deploy cloud infrastructure. Rather than addressing a problem interactively on an individual server, as you'd typically do in an on-prem environment, you need to be able to update the IaC template that you use to configure your VMs, then deploy a new, secure VM to replace the problematic one.

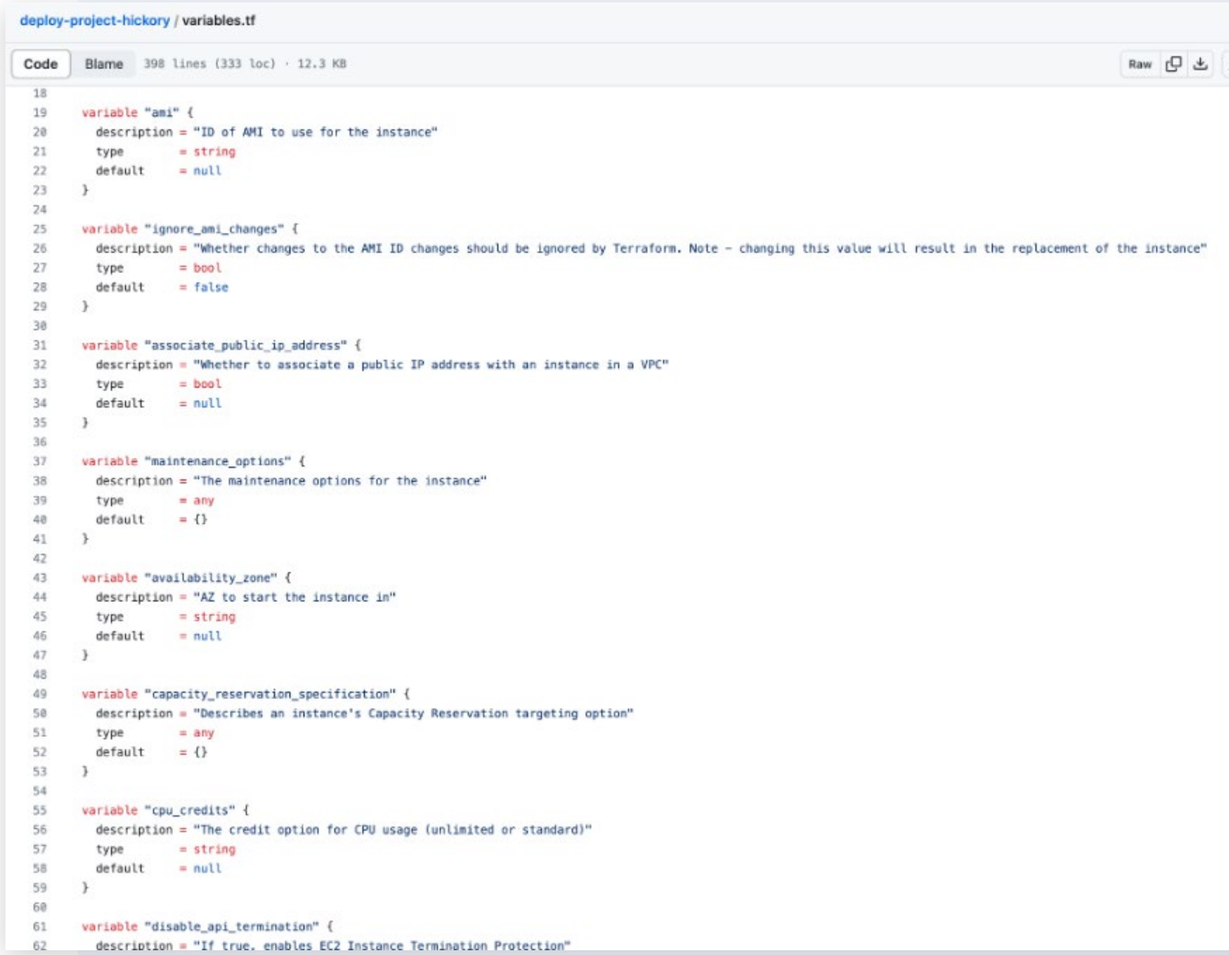
deploy-project-hickory / variables.tf

Code Blame 398 lines (333 loc) · 12.3 KB Raw

```
18
19  variable "ami" {
20    description = "ID of AMI to use for the instance"
21    type        = string
22    default     = null
23  }
24
25  variable "ignore_ami_changes" {
26    description = "Whether changes to the AMI ID changes should be ignored by Terraform. Note - changing this value will result in the replacement of the instance"
27    type        = bool
28    default     = false
29  }
30
31  variable "associate_public_ip_address" {
32    description = "Whether to associate a public IP address with an instance in a VPC"
33    type        = bool
34    default     = null
35  }
36
37  variable "maintenance_options" {
38    description = "The maintenance options for the instance"
39    type        = any
40    default     = {}
41  }
42
43  variable "availability_zone" {
44    description = "AZ to start the instance in"
45    type        = string
46    default     = null
47  }
48
49  variable "capacity_reservation_specification" {
50    description = "Describes an instance's Capacity Reservation targeting option"
51    type        = any
52    default     = {}
53  }
54
55  variable "cpu_credits" {
56    description = "The credit option for CPU usage (unlimited or standard)"
57    type        = string
58    default     = null
59  }
60
61  variable "disable_api_termination" {
62    description = "If true, enables EC2 Instance Termination Protection"
```

Microservices

Not all applications that move to the cloud are microservices (meaning they are broken into discrete services rather than operating as monoliths). Nor do you strictly need to use the cloud if you want to run a microservices app. Microservices can run on-prem, too.

That said, migration to the cloud has gone hand-in-hand with the adoption of microservices at many organizations. Refactoring existing applications so that they run as microservices is a great way to take full advantage of the distributed nature of cloud environments, and new applications that businesses deploy in the cloud are often built using microservices from the start.

Microservices and the deployment technologies that support them – such as containers and serverless functions – provide many benefits, including the ability to scale applications more efficiently and with greater reliability in the face of risks like failed servers.

However, microservices also significantly complicate security operations. They increase the attack surface of applications because they introduce more layers – such as an orchestrator like Kubernetes and an API gateway – into application hosting stacks. They also make it more difficult in many cases to determine exactly which microservice triggered a security incident. And the complex set of interactions and dependencies that exist between microservices increases the risk of configuration oversights that might expose microservices to attack.

In addition, remediation practices for addressing vulnerabilities in containers and serverless functions that host microservices are different from those of traditional environments. Most microservices apps are deployed using an immutable infrastructure strategy, which means that components that need to be patched are completely replaced with a new version instead of being updated in place. If you detect a vulnerability in a container or serverless function, you'd typically update the image for the container or function, then deploy an entirely new instance of it and delete the original one. That's different from patching a monolithic application, where you'd normally apply a patch without rebuilding the entire application.

In short, microservices translate to a much higher degree of complexity for security teams. Detecting and remediating security risks for microservices requires more complex processes than securing the monoliths that organizations typically run on-prem.

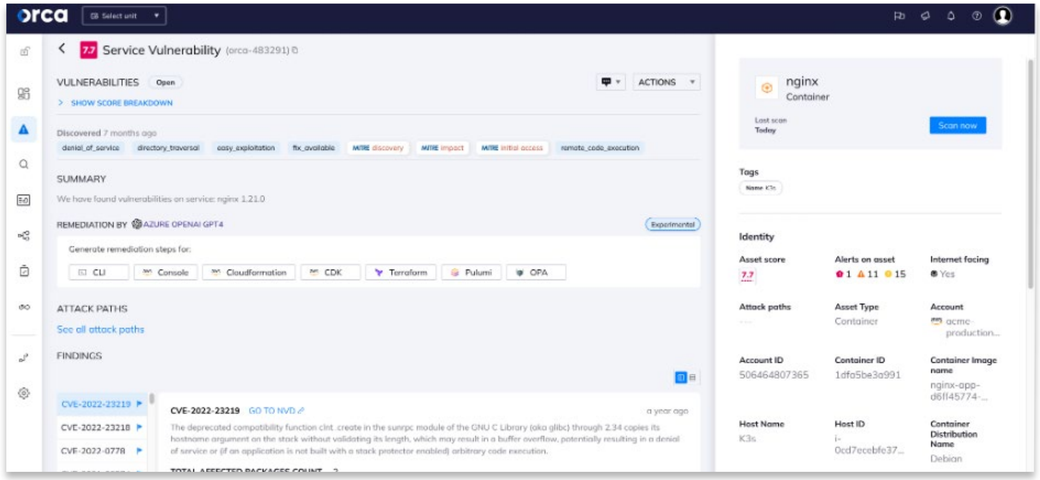

Multiple Cloud Platforms

There are many clouds out there. The most famous are the "Big Three" public clouds – Amazon Web Services, Microsoft Azure, and Google Cloud Platform – but organizations might also choose to deploy workloads on a variety of other public or private cloud platforms. Each of those clouds includes dozens of different types of services. In contrast, on-prem servers demand careful setup and management practices. A failing or hacked server can't simply be torn out and replaced with a new one in a matter of minutes. On-prem servers are therefore akin to "pets," which IT teams must carefully manage and groom.

Most clouds provide the same basic sets of security tools, such as Identity and Access Management (IAM) frameworks for restricting access to resources and monitoring services that can collect security-relevant data. However, each cloud provider implements these solutions differently.

As a result, security teams have more diverse sets of tools to work with, especially if their organizations use multiple clouds at the same time. In an on-prem environment, the only real platform difference you typically encountered was between Windows- and Linux-based servers, but in the cloud, there are hundreds of distinct cloud services on offer from among the various cloud providers.

Implementing processes for detecting and mitigating risks in this situation is, in a word, harder than when working with relatively homogenous on-prem environments.

Expansive User and Permission Settings

A single cloud environment might have hundreds or even thousands of users operating within it. Some are humans and others are machine users that run automated tasks. Each user receives a different set of permissions that define which resources they can access.

With such a vast number of users and permissions configurations to oversee, it can be difficult to adhere to security best practices related to access control. Ensuring that each user is assigned the least privileges necessary, for example, is difficult when there are so many users with unique (and constantly changing) access requirements. Enforcing zero trust can be difficult, too, because admins may be tempted to assign users permissions by default to simplify the permissions management process.

The problem becomes even worse if you use multiple clouds or cloud accounts with a different set of user and permissions configurations for each one.

These challenges are less acute in on-prem environments, where the total number of users tends to be smaller, and where permissions can be managed through a central system like Active Directory.

To manage these risks in the cloud, you need to be able to detect risks associated with weak permissions settings on a large-scale, highly dynamic basis. The practice that you use to manage users on-prem can't scale enough to work well in the cloud.

WHAT MAKES THE CLOUD DIFFERENT FROM A SECURITY PERSPECTIVE?

Limited Visibility

In an on-prem environment, you have complete control over your infrastructure and the workloads hosted on it. This translates to the ability to gain complete visibility into what's happening within your environment. You can collect every log and metric that your workloads and the underlying infrastructure produce.

Such is not the case in the cloud, where you can't access the underlying infrastructure. On top of that, it's often more difficult to monitor complex, microservices-based applications because they don't produce centralized logs and metrics in the way that on-prem monoliths do. For some types of cloud services – such as serverless functions – you can't even deploy security monitoring agents to help detect risks in your workloads because you can't even access the virtual servers that host your workloads.

Here again, this means that the security strategies that work on-prem can't keep up with the cloud. You must rethink your approaches to security monitoring and vulnerability management when you move to a cloud environment.

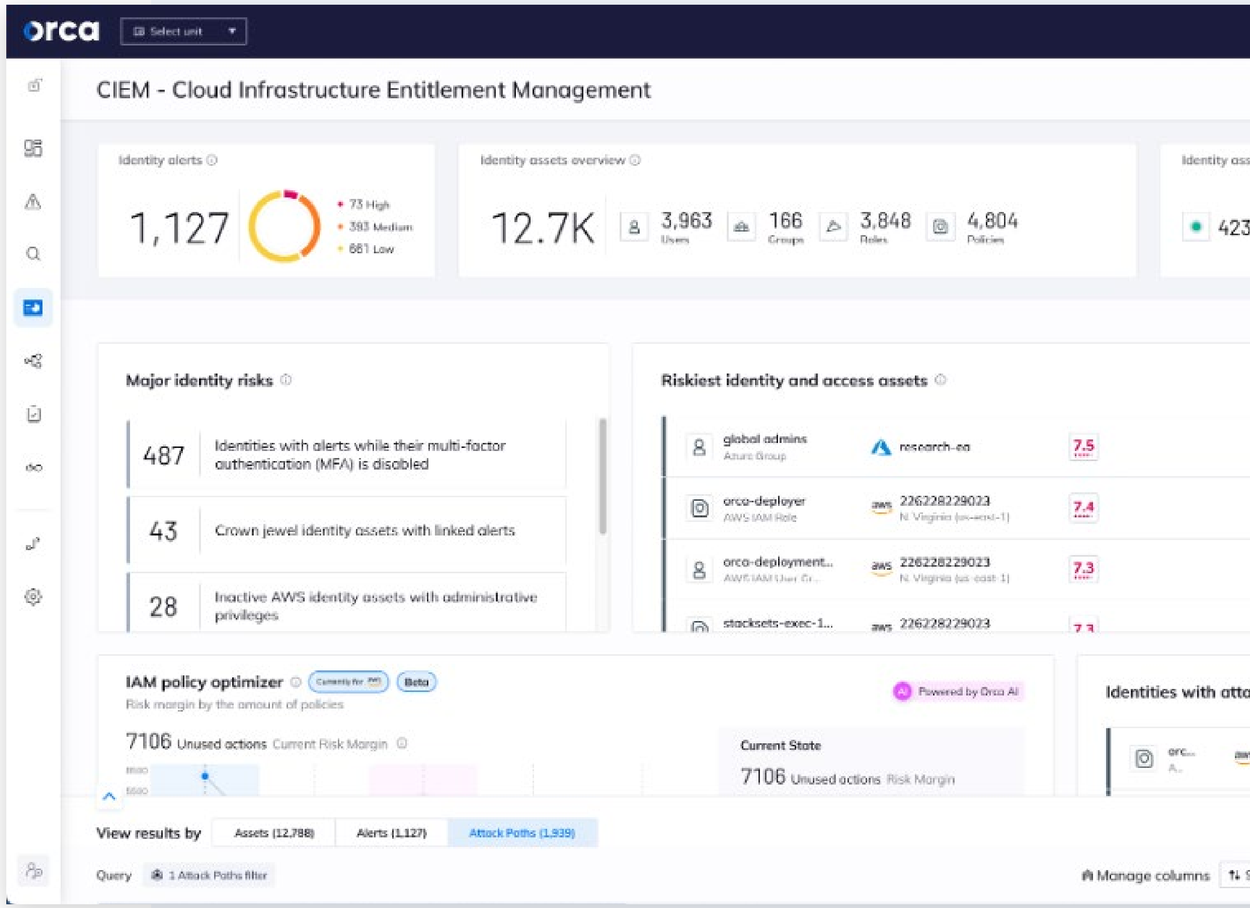

Cloud Security in Practice: A Real-Life Scenario

To drive home how different securing a cloud environment is from securing on-prem resources, imagine how your team would react in each setting to a critical vulnerability.

In the old world of on-prem servers, the vulnerability alert would cause some unlucky on-call operations engineer to have to apply a patch to the servers, reboot them if necessary, and then verify that the vulnerability has been resolved.

In the declarative world of cloud-based computing, however, that approach wouldn't work. You could patch your cloud servers or containers if you wanted, but doing so wouldn't resolve the vulnerability because it would reemerge as soon as you launched new server or container instances.

So, instead, you would need to go back to the IaC code, server images, and/or container images that your workload depends on and patch the vulnerability there. Then, you'd want to test the change to make sure it works as intended. Finally, you'd redeploy your workloads so that the fix takes effect in your production environment.

Each approach has its pros and cons – the on-prem remediation procedure is simpler but more onerous on engineers, while the security response in the cloud is more complex to execute but also easier to scale because a single change to infrastructure configuration can apply a patch across dozens or even hundreds of server or container instances.

But whether the cloud is better than on-prem is not what we're here to argue about. Our goal in this eBook is to emphasize that the cloud is simply different – and that it requires security detection and response practices designed specifically for cloud environments.

Conquering Cloud Security Challenges

Those are some of the security challenges that migrating to the cloud introduces. Now, let's talk about solutions.

Viewed from a very high-level perspective, the security practices that apply to the cloud aren't fundamentally different from those you'd use on-prem. Cloud security revolves around processes like scanning workloads for vulnerabilities and monitoring configurations for oversights that could introduce security risks, just as you would in an on-prem environment. However, the unique challenges of the cloud mean that these processes must be implemented in different ways.

Here's a look at how to adapt on-prem security strategies to the cloud.

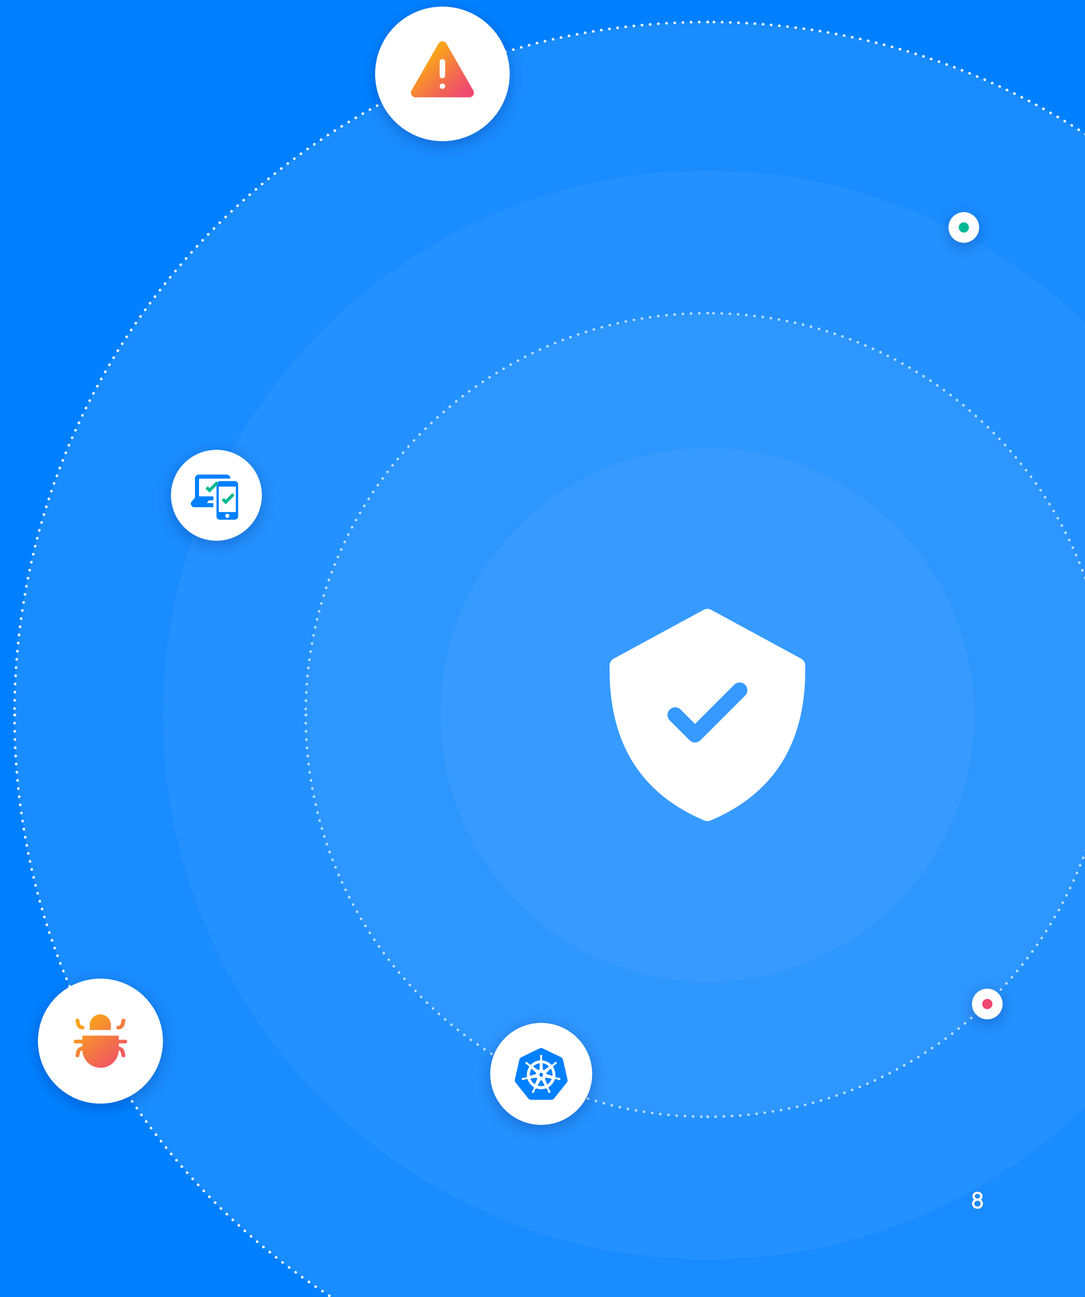

Bake Security Into Automation Pipelines

As we noted above, mitigating security risks on a case-by-case basis doesn't work in the cloud, where infrastructure management processes are typically automated and where individual resources may be destroyed and replaced by the time you can apply a patch.

Instead, effective cloud security includes integrating security into the automation pipelines that your business uses to manage its cloud infrastructure. You must be able to scan for vulnerabilities or risky configurations introduced by IaC templates, for example, and you should automate security scans as part of your continuous application delivery pipelines.

Practices like these – which shift security "left" by baking it into the software delivery process, rather than applying security only to post-deployment production environments – are the only way to manage cloud security threats at scale.

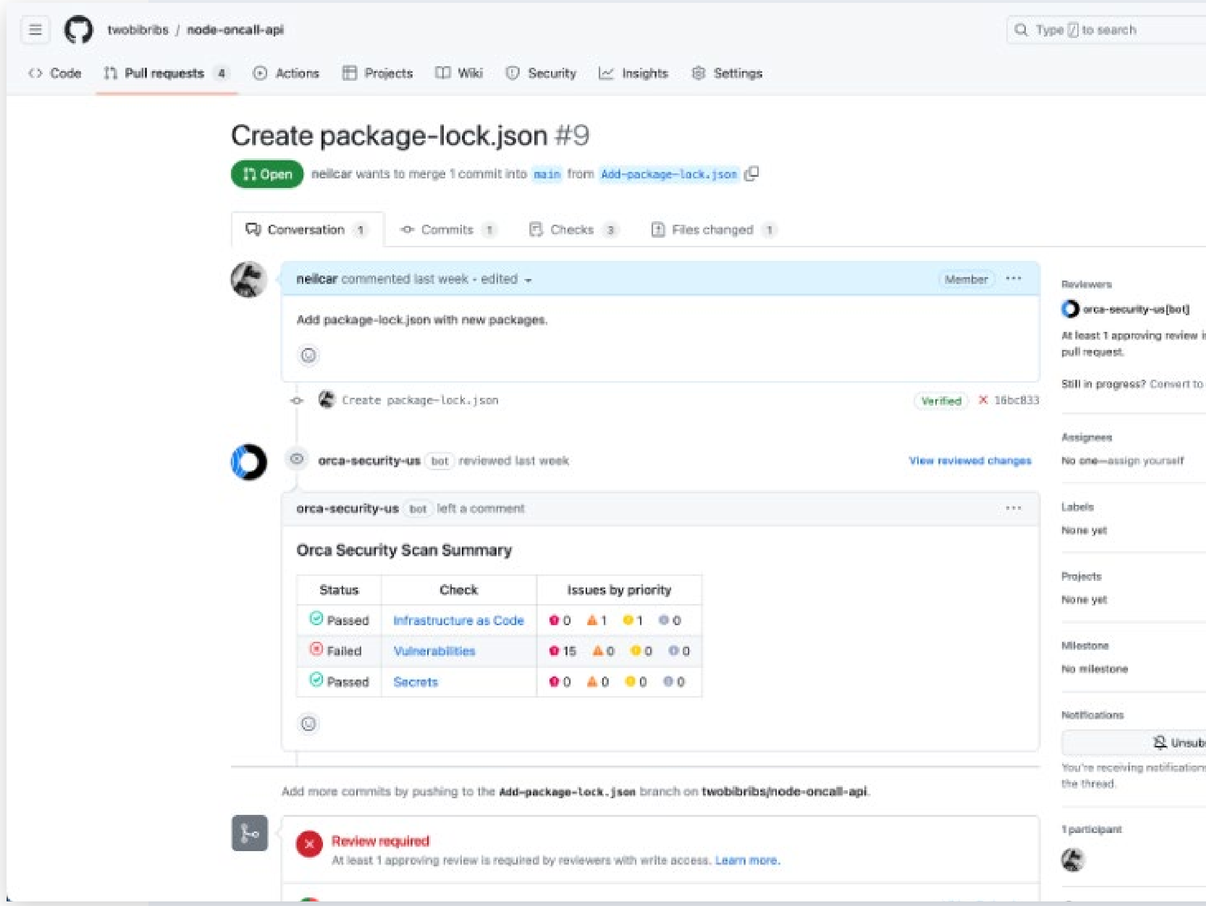

Centralize and Consolidate

Your cloud architecture might be complex and diverse because it includes multiple cloud providers, many accounts within each provider, and/or multiple types of cloud services. But that doesn't mean your security toolset should be complex and heterogeneous, too.

On the contrary, the best cloud security solutions are ones that can operate across all cloud environments and protect all types of cloud workloads. Instead of relying on security tools that work with just one cloud, or deploying different solutions for different categories of workload (like containers on the one hand and serverless functions on the other), choose security platforms that provide true end-to-end cloud security.

Centralized and consolidated security tools not only make it easier for security teams to do their jobs by saving them from having to juggle multiple tools or constantly switch between contexts. They also provide the richest possible context on security risks because they can compare different types of security data from across your cloud. That makes it easier to pinpoint the root cause of risks as well as to assess which risks pose the greatest threat.

Take an Agentless Approach

Agentless security means collecting security data without the use of traditional software agents that run directly on host infrastructure. An agentless approach solves many of the visibility challenges of the cloud because it makes it possible to gain deep insight into what is happening within cloud workloads even if you can't access the underlying infrastructure.

Agentless security solutions offer other benefits to boot, such as a reduced level of resource consumption (which translates to better performance for your applications and lower hosting costs). But the killer feature they bring to the cloud is their ability to monitor and help protect workloads in contexts where conventional security approaches just don't work.

Scan and Monitor Continuously

Centralized and consolidated security tools not only make it easier for security teams to do their jobs by saving them from having to juggle multiple tools or constantly switch between contexts. They also provide the richest possible context on security risks because they can compare different types of security data from across your cloud. That makes it easier to pinpoint the root cause of risks as well as to assess which risks pose the greatest threat.

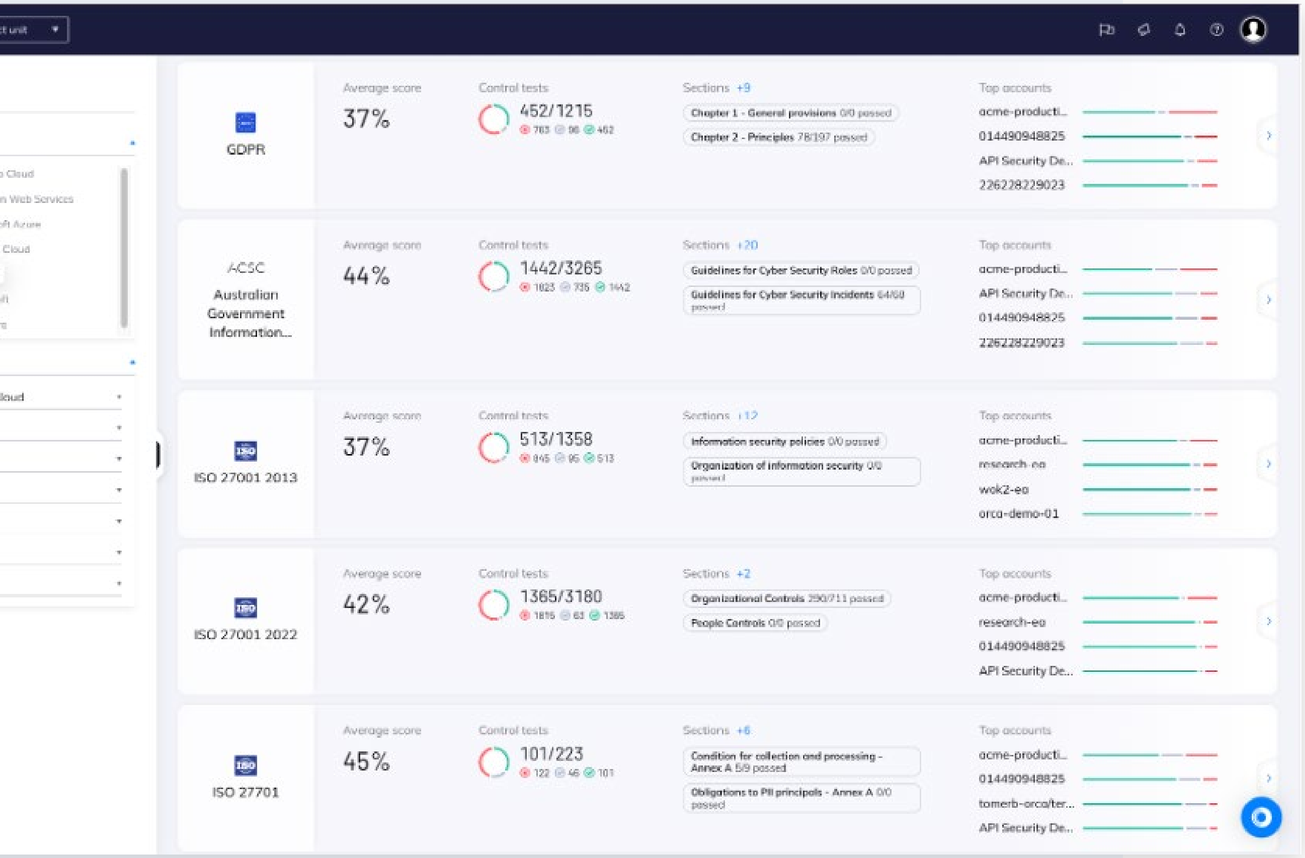

CONQUERING CLOUD SECURITY CHALLENGES

Leverage Compliance Benchmarks

Given the complex nature of cloud environments, it can be difficult to determine what a “good” security posture looks like. Even seasoned engineers can struggle to configure workloads in a highly secure way.

Rather than guessing or adopting default settings that may or may not make the most sense from a security perspective, leverage proven benchmarks – such as those provided by the Center for Internet Security (CIS) – to guide your cloud security strategy. Benchmarks help ensure that your cloud environment complies with recognized standards and best practices from a security perspective.

As for actually enforcing the benchmarks, you could audit your configurations manually, of course. But a better approach is to leverage security tools that are designed out-of-the-box to detect deviations from recognized security standards automatically. That way, you know your workloads adhere to best practices, and you don’t have to spend hours and hours configuring your security policies or assessing whether your cloud workloads comply with them.

Conclusion

In short, cloud security requires a new set of approaches to managing an old set of problems. Teams that excel in managing vulnerabilities, protecting against ransomware, and mitigating permissions configuration risks on-prem aren't prepared to succeed in the face of cloud security threats unless they revamp their security processes and strategies by:

- Baking security into cloud automation and application delivery pipelines.
- Centralizing security operations and tooling.
- Adopting agentless security tools to gain visibility in environments where conventional monitoring processes fail.
- Scanning and monitoring on a continuous, real-time basis.
- Using industry-standard benchmarks to guide security and compliance operations.

About the Orca Cloud Security Platform

Orca Security is the industry-leading agentless Cloud Security Platform that identifies, prioritizes, and remediates risks and compliance issues across your cloud estate spanning AWS, Azure, Google Cloud and Kubernetes. Instead of layering multiple siloed tools together or deploying cumbersome agents, Orca delivers complete cloud security in a single platform by combining two revolutionary approaches: SideScanning, which enables frictionless and complete coverage without the need to maintain agents, and a Unified Data Model, which allows for centralized contextual analysis of your entire cloud estate.

Orca's agentless platform connects to your environment in minutes and provides 100% visibility of all your assets, automatically including new assets as they are added. Orca detects and prioritizes cloud risks across every layer of your cloud estate, including vulnerabilities, malware, misconfigurations, lateral movement risk,API risk, weak and leaked passwords, and overly permissive identities.

Get a recorded or personalized demo of the Orca Cloud Security Platform.

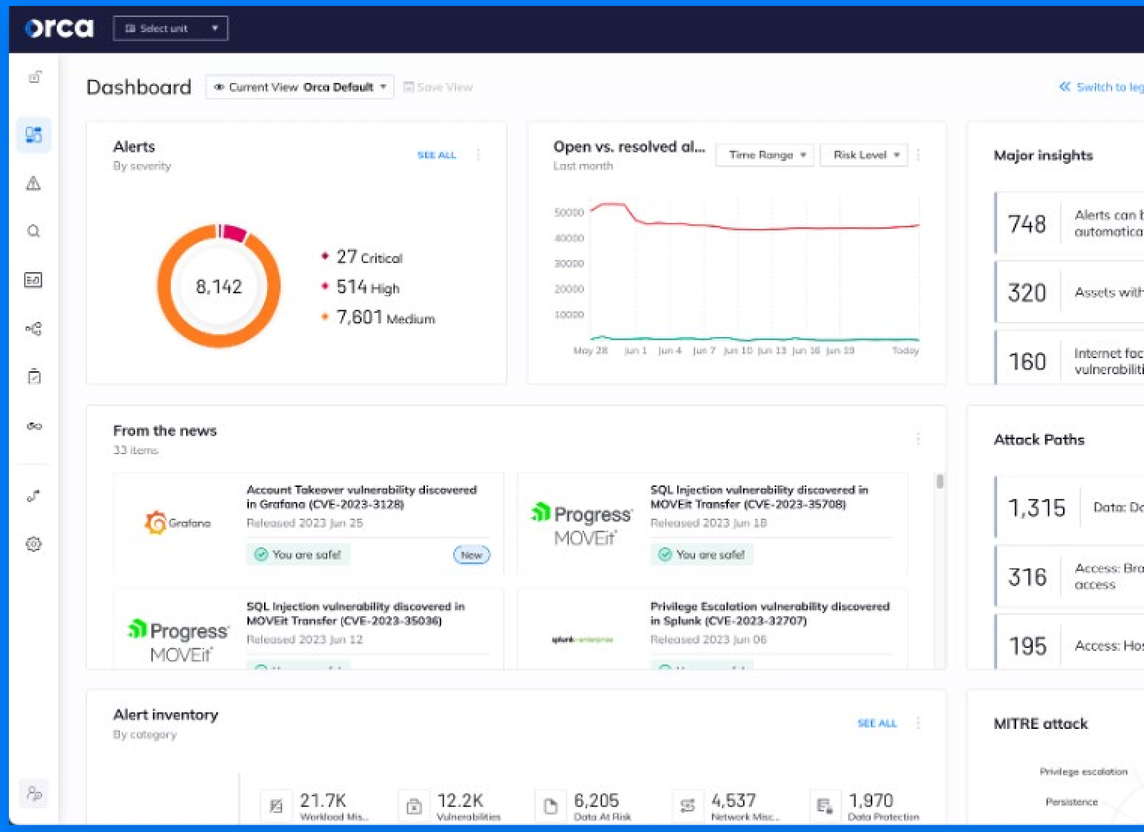


orca
security