

Brought to you by:

WIZ[★]

Cloud Detection & Response

**for
dummies**[®]
A Wiley Brand



Visualize
threats

Use cloud-native
capabilities

Monitor cloud and
workload activity

Wiz Special Edition

Lisa A. Bucki

Hey readers,

Before you dive into this book, here's what you need to know about Wiz and how we're building a best-in-class Cloud Detection and Response (CDR) solution.

Wiz launched with a single objective: to revolutionize risk reduction in the cloud through a solution that security and development teams love. Wiz has changed the game for cloud security teams, by enabling them to efficiently identify and remediate risk with the power of context. But Security Operations teams have also been transformed by the shift to the cloud – as attackers develop new tactics, techniques, and procedures to target cloud environments, SecOps teams have to respond with new frameworks for Cloud Detection and Response.

In this book, we'll introduce the key aspects of an end-to-end CDR solution, show you how it differs from other cybersecurity efforts, and share practical best practices for real-time detection and response to cloud threats.

If you are interested in learning more about how Wiz can help your organization stop threats before they become breachers, check out wiz.io/demo or scan the QR code below.

Cheers,

Arie Zilberstein

VP of Threat Response, Wiz.io





Cloud Detection & Response

Wiz Special Edition

by Lisa A. Bucki

**for
dummies[®]**
A Wiley Brand

Cloud Detection & Response For Dummies®, Wiz Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2025 by John Wiley & Sons, Inc., Hoboken, New Jersey. All rights, including for text and data mining, AI training, and similar technologies, are reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Wiz and the Wiz logo are trademarks or registered trademarks of Wiz, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.dummies.com/custom-solutions. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-394-30413-4 (pbk); ISBN 978-1-394-30414-1 (ebk); ISBN 978-1-394-30415-8 (ebk)

Publisher's Acknowledgments

Project Editor: Jen Bingham
Acquisitions Editor: Traci Martin
Senior Managing Editor: Rev Mengle

Client Account Manager:
Jeremith Coward
Production Editor:
Tamilmani Varadharaj

Table of Contents

INTRODUCTION	1
About This Book	1
Icons Used in This Book.....	2
Beyond the Book.....	2
CHAPTER 1: Introducing Cloud Detection and Response	3
Understanding the Challenges of Threat Detection and Response in Cloud Environments.....	4
Who's Involved in Cloud Detection and Response?	6
Recognizing Differences in Cloud Attacker Tactics, Techniques, and Procedures	7
Building a Cloud Detection and Response Framework.....	8
Prepare.....	8
Detect	9
Investigate.....	9
Respond	9
Benefits of an Effective CDR Program	10
CHAPTER 2: Preparing for Cloud Attacks	11
Understanding Your Cloud Attack Surface.....	12
Constructing an Inventory of Cloud Assets.....	13
Assessing Which Assets Need Real-Time Monitoring	14
Using Logs and Runtime Data for Real-Time Visibility.....	16
CHAPTER 3: Detecting Attacks in Real Time	19
Approaches to Detection in Cloud Environments.....	20
Using Rules-Based Detection Tailored for Cloud Threats	20
Lowering Noise with Behavioral Analytics	22
Prioritizing Alerts Using Context.....	24
Aggregating and Monitoring Findings	25
Benchmarking Detection Coverage	25
CHAPTER 4: Investigating Incidents	27
Starting the Investigation	27
Understanding Who to Involve.....	28
Automating Context Gathering	29

Creating and Maintaining an Incident Timeline	30
Organizing a single source of truth	31
Expanding your view of an attack.....	31
Analyzing the Blast Radius	33
CHAPTER 5: Containing and Responding to Threats.....	35
Applying Containment Measures ASAP	36
What Cloud Forensics Entails and Why It's Important.....	37
Ensuring Forensic Data Availability	38
Best Practices for Dedicated Forensic Environments.....	39
Developing Automated Playbooks to Collect and Analyze Artifacts	40
Collaborating with Developers on Remediation.....	41
CHAPTER 6: Five Best Practices for Cloud Detection and Response.....	43

Introduction

Companies and individual technology users throughout the world continue to expand their reliance on cloud computing. Individual organizations and cloud service providers host an astonishing and growing amount of data in the cloud, including customer behavior data, healthcare information, financial records, social and gaming accounts, and beyond.

But although the cloud environment offers numerous advantages, it poses unique cybersecurity challenges for security operations teams. The cloud generates enormous amounts of security data, all of which is easier to access than ever before. But operationalizing that security data in an effective security operations program can be extremely challenging.

Implementing and automating end-to-end cloud detection and response (CDR) processes can reduce false positive alerts, help teams prioritize investigation and threat containment efforts, and enable proactive updates to make the cloud environment more secure.

About This Book

This book introduces key aspects of CDR, emphasizing how it contrasts with on-premises solutions. It begins with a broad overview of these capabilities, discussing how digital transformations and the increasing shift to cloud technologies have made security operations more complex.

Topics included in this book include:

- » Describing the unique challenges of threat detection and response in a cloud environment
- » Exploring the need for implementing a CDR framework in an organization
- » Understanding four key aspects of an end-to-end CDR solution: preparing for incidents, detecting malicious activity, investigating alerts, and responding to threats
- » Reviewing the opportunities that arise when implementing a CDR solution

Icons Used in This Book

Throughout the margins in this book, you see little icons that highlight different types of information. Here's what each means:



REMEMBER

The Remember icon reinforces must-know, fundamental concepts. After you're done reading this book (if you read the entire thing), you can always jump back to these items to refresh your memory.



TECHNICAL
STUFF

The Technical Stuff icon highlights details and facts beyond the basics. You might find this information interesting even though it's not essential.



TIP

The Tip icon points out ideas that can save time and frustration.



WARNING

The Warning icon identifies actions or situations to avoid.

Beyond the Book

This book helps business leaders and security operations professionals like you discover more about CDR techniques, tools, and processes. Given the breadth and complexity of the field, you may want to take a deeper dive into the subject on your own. To learn more, visit:

- » <https://threats.wiz.io>
- » <https://www.youtube.com/watch?v=K4wNQutuz5DI>
- » <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>
- » <https://attack.mitre.org/matrices/enterprise/cloud>

IN THIS CHAPTER

- » Discussing the differences between detection and response in traditional and cloud environments
- » Listing the cloud detection and response players
- » Understanding why cloud detection and response capabilities are necessary
- » Reviewing how cloud attackers operate
- » Outlining the key capabilities of a cloud detection and response system

Chapter 1

Introducing Cloud Detection and Response

We live data-driven lives now more than ever. Billions of devices connect to the cloud worldwide to perform a host of activities, from making online purchases to updating organizations' employee databases. Cloud adoption continues to grow rapidly to accommodate all these new services.

The expanding size and complexity of the cloud creates both challenges and opportunities for cyberdefenders. On the one hand, the cloud provides application programming interfaces (APIs) for everything, and you can get access to more information more quickly than ever before. But this information can be overwhelming, and operationalizing it for effective security operations is a challenge.

This chapter introduces you to cloud detection and response (CDR), including the unique threat detection challenges in cloud environments, who in your company is involved in CDR, the novel

tactics and techniques used by cloud attackers, the key components of an end-to-end CDR solution, and the opportunities CDR offers to an organization.

Understanding the Challenges of Threat Detection and Response in Cloud Environments

Cloud environments are often fundamentally different from traditional on-premises environments. In on-premises environments, organizations had complete control of their infrastructure, from physical deployment of machines to management of operating systems and the boundaries of the network perimeter.

In the cloud, organizations share these responsibilities with cloud providers, and may not be able to rely on the same types of monitoring to ensure security coverage. Plus, cloud environments are extremely complex, typically including dozens of services, widely distributed resources, numerous APIs, interconnected applications, different types of workloads, and short-lived cloud assets — access to all of which is generally controlled by a constantly changing identity layer.

This complexity leads to numerous unique challenges in securing a cloud environment, including:

- » **Greater risk of misconfiguration:** Errors in the security settings of cloud applications and systems present prominent attack vectors for data breaches. In the cloud, every service has an IP address by default, making errors in configuration more easily exploitable by attackers. Easily exploited misconfigurations include overprivileged identities, unsecured storage buckets, and using default credentials and passwords.
- » **Increased discoverability:** Security team members as well as would-be attackers can collect detailed information about cloud-based infrastructure, especially when APIs aren't well secured. Attackers can then use various assets and techniques to find new attack paths to compromise critical resources.

- » **Automated attack techniques:** Many environments use cloud management tools to perform tasks such as allocating resources without intervention by a team member, reducing manual workload. Attackers can use these automated processes to move much faster, expanding the blast radius of a breach.
- » **Lack of visibility:** When an organization mixes and matches cloud technologies from various providers on the fly, the resulting cloud environment becomes more complex, interconnected, and dynamic. Such an environment hampers the visibility needed to identify, contextualize, prioritize, and mitigate vulnerabilities.
- » **Greater blast radius:** If attackers can compromise the centralized control plane, they can gain access to stored data, network resources, and applications. The attackers can then not only damage the data, but even make changes to the environment itself.

Traditional security operations tools like endpoint detection and response (EDR), network detection and response (NDR), and security information and event management (SIEM) simply weren't designed to address the complexity of cloud environments.



REMEMBER

CDR involves identifying and mitigating security threats or incidents in the cloud. The process includes monitoring and investigation and may include automated or manual actions.

The new CDR standard focuses specifically on cloud environments. It includes automating tasks such as collecting forensic evidence and analyzing threat data to accelerate investigation and remediation of cloud attacks. CDR tools generate alerts in real time and automatically embed cloud-specific context within them, enabling a security operations (SecOps) team to prioritize remediation efforts.

CDR also can trigger automated responses such as quarantining workloads, creating asset and network isolation zones, or using approved images to rebuild workloads.



TIP

As with any new technology, terminology can vary and change over time. You may see CDR called cloud-native detection and response (CNDR) or cloud threat detection and response (CTDR).

Who's Involved in Cloud Detection and Response?

Cloud environments operate differently than traditional environments, and the lines between functions in the organization can get blurry when investigating and responding to cloud threats. It's important to understand the different stakeholders involved:

- » **Security operations (SecOps) teams** are primarily responsible for detecting and responding to real-time threats occurring in the environment, encompassing functions like security operations center (SOC) management, detection engineering, and incident response (IR).
- » **Cloud security teams** are responsible for the overall security posture of the organization's cloud environment, often including configuration management and overall security architecture.
- » **DevOps and developer teams** build and deploy the organization's applications and workloads. In cloud environments, where infrastructure is often managed with infrastructure-as-code (IaC) and there is a tight link between development and production environments, developers often have key context necessary to resolve threats and are often ultimately responsible for remediating root causes.
- » **Executives** like Chief Information Security officers are responsible for overall risk at the organization, and ultimately report to the board.

For most organizations, the SecOps teams lead cloud detection and response efforts. However, because so much activity in cloud environments is driven by fast-moving developers, close collaboration between SecOps and cloud security and developer teams is essential.

Recognizing Differences in Cloud Attacker Tactics, Techniques, and Procedures

Just as the shift to cloud has changed the game for defenders seeking to secure cloud environments, attackers are also adapting to the new reality. Although the underlying goals for threat actors operating in the cloud — such as resource hijacking or data exfiltration — remain the same, the techniques used to achieve these goals are often quite different. As a defender, examples of cloud-native attack techniques to be aware of include:

- » **Identity compromise:** The attacker might use a compromised identity to gain access to and move through the cloud control plane. Using credentials lifted from a compute instance, an attacker may be able to impersonate a user or nonhuman identity in other parts of the cloud environment.
- » **Cloud-native defense evasion:** Attackers seeking to remain undetected often tamper with audit logs, which are often stored in cloud-native storage solutions like object stores. By changing retention policies on those object stores through the control plane, attackers can erase evidence of their activities without even touching the actual log files.
- » **API attacks:** APIs may be improperly secured, such as having a lack of access authentication protocols, allowing unauthorized access to systems or data. APIs created within an organization may be improperly documented, unauthorized, or unknown to the organization. Such APIs might expose sensitive data or functionality, especially outdated APIs or ones that are misconfigured. API inputs and outputs also can be subject to attacks that inject malicious data, especially if the data is not properly validated.

In the cloud, attackers often don't even use malware. With a compromised identity, an attacker can move laterally, escalate privileges, and even exfiltrate data without even touching a machine workload.

The shift in attack techniques requires a shift in defensive strategy. For example, EDR has long been a common cybersecurity method. EDR uses agents — specialized software that performs a specific action, such as security scanning and reporting, installing software patches, runtime detection and response, or changing configurations. The agent must be installed and set up on every server or host in the system, from which it passes monitoring data to a centralized security service.

The agent-based approach works for the hosts that an organization controls. However, if an organization also uses cloud service resources such as an Amazon Simple Storage Service (Amazon S3) bucket where the underlying infrastructure is managed by the cloud service provider, the organization will be unable to gain agent-based visibility over these resources.

CDR systems provide automated, proactive tools and processes that identify and prioritize vulnerabilities and attacks. These include real-time monitoring and detection across the organization's entire cloud ecosystem, including in many instances the use of both agentless security and runtime sensors. Agentless security enables the centralized security service itself to scan hosts and collect and analyze data, whereas runtime sensors give full visibility into host-level activity.

Building a Cloud Detection and Response Framework

The complexity of detecting and responding to cloud threats can feel overwhelming. As environments change rapidly and attackers evolve with new techniques, it can be difficult to keep pace.

The best way to ensure you're ready for a cloud attack is to follow an established process. There are four main pillars of cloud detection and response, discussed in the following sections (and the rest of the book).

Prepare

The first step in stopping a cloud threat is ensuring that you have the data you need to see it in the first place. Preparing for cloud threats involves activities that may not necessarily have been

required in traditional environments. Higher level activities such as developing and implementing cloud security policies and an incident response plan can be a starting point. But from there, the SecOps organization should assess what telemetry is being collected and what compute assets have runtime visibility, evaluating their preparedness for a cloud incident.

Detect

The organization should implement continuous detection tools across cloud services to monitor all relevant attack surfaces in the cloud. Teams should be sure to monitor both runtime telemetry from workload sensors as well as audit and activity logs from their infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS), and software-as-a-service (SaaS) platforms. Using rules-based detection tools developed specifically for detecting cloud threats, coupled with cloud-native threat intelligence, enhances detection. Behavioral analytics augment the effort by revealing unusual activity.

Investigate

Investigating an attack reveals the best path toward mitigation and identifying any additional cloud security measures that should be implemented. The first step is to gather data that provides context about the incident to determine its scope and severity. Creating an incident timeline helps document how the attack unfolded for future review and learning. Analyzing the blast radius of the threat enables the SOC team to identify the appropriate containment measures as soon as possible.



TIP

An organization can deploy cloud investigation and response automation (CIRA) tools. These tools automate collecting and analyzing forensic data in cloud environments.

Respond

The security operations center (SOC) team needs to apply the identified containment and mitigation measures as soon as possible. Beyond that, the organization should focus on continuing to collect and analyze forensic data about an incident to enhance future planning and response.

If the investigation and response reveal significant vulnerabilities in need of remediation, the SOC team should collaborate with developers to close the attack paths in the cloud system. The forensic incident data collected should also guide future development efforts, to build a more secure cloud structure from the ground up.

Benefits of an Effective CDR Program

Although cloud security can seem like nothing but a huge problem, in truth, it's a problem that a security organization can transform into an opportunity. Implementing an end-to-end CDR solution can lead to benefits like the following:

- » Contextualizing and prioritizing threats reduces false positive alerts that waste SOC team time. By some estimates, SOC teams spend more than 30 percent of working time on false incident investigations.
- » Manual threat detection is inherently time consuming. The automations offered by CDR tools reduce the need for manual intervention, also reducing the need for additional hiring and allowing teams to focus on more strategic initiatives.
- » In the cloud, properly implemented CDR programs can help bring security operations, cloud security, and development teams together to collaborate more effectively. This collaboration helps makes the cloud environment more secure from the start.

- » Introducing you to the cloud attack surface
- » Inventorying your organization's cloud environment
- » Prioritizing assets for monitoring
- » Striking the right balance between visibility and cost

Chapter 2

Preparing for Cloud Attacks

Security operations center (SOC) professionals don't *think* that cloud attacks are coming, they *know* that cloud attacks are coming. To put your SOC team in a better position to respond to breaches, it is important to have a mechanism for gathering and recording threat activity.

This chapter covers the first major area of focus for developing a modern cloud detection and response (CDR) framework in an organization: prepare. The chapter begins with covering best practices for logging activity in cloud environments, moves on to discuss inventorying assets and selecting which assets need continuous monitoring, and concludes with a review of factors to consider when determining how much monitoring makes financial sense for your organization.

Understanding Your Cloud Attack Surface

There are lots of different types of assets in cloud environments — far more than the users and machines you might be used to from the on-premises world. Different asset classes in cloud environments include:

- » **Identity:** Both human and nonhuman identities and their associated roles and permissions, often managed through a combination of third-party identity providers (IdPs) alongside cloud-native identity layers.
- » **Data:** Both self-hosted as well as managed storage solutions like relational databases, object stores, and more.
- » **Network:** Virtualized network infrastructure like virtual private clouds (VPC), network security groups, and others.
- » **Compute:** Workloads running in your environment, including virtual machines (VMs), containers, and serverless.
- » **Control:** The management layer that enables creation, deletion, and modification of other resources in the cloud environment.

Responsibility over different types of assets will be shared between you and your provider, potentially limiting the information you have available for monitoring.

For example, if you self-host a database on a VM, you'll have full access to and visibility over the operating system underlying that database at the expense of additional complexity and potential risk. A managed database will often be easier to maintain and scale, but your visibility can be limited to audit logs.



REMEMBER

The control plane is important to emphasize, as it represents one of the key differences between cloud and traditional environments. In cloud environments, the control plane is crucial: through the control plane, an attacker can not only access an organization's environment, but they can even *change it*.

Constructing an Inventory of Cloud Assets

The first step in securing your cloud environment is understanding what's in it. Just as experts recommend that homeowners create an inventory of personal property to document its value for insurance purposes, cloud security experts advise creating an inventory or map of cloud resources. When constructing an inventory, you should ensure that each category discussed in the previous section is well-represented. You should also capture whether real-time visibility for activity related to each cloud asset is currently available.

One of the challenges with constructing a cloud inventory is that cloud environments grow and change extremely quickly. Services and features from cloud providers are added, moved, updated, and sometimes discontinued. IT staff members and vendors come and go, sometimes taking useful knowledge of the landscape with them. Plus, in the cloud, control of resource creation, destruction, and maintenance is often in the hands of developers and DevOps teams. Depending on the organization, these teams may or may not be in close communication with Security Operations, potentially leaving gaps in visibility as the environment changes.

All this change means that manually constructing an inventory is nearly impossible. Instead, automated tooling should be used to keep the inventory up-to-date.



Services for inventorying or mapping cloud resources use processes known as resource discovery, asset discovery, and/or service discovery. The service constantly communicates with cloud application programming interfaces (APIs), providing a continuously updated look at the status of the cloud's components.

Most cloud service providers offer automated inventory or mapping tools to their customers. These tools dynamically track the network and update as the status of various assets changes, giving real-time visibility into all the resources. AWS, Microsoft Azure, and Google Cloud all offer cloud inventory tools. Some examples of these inventory tools are below:

» AWS Config (<https://docs.aws.amazon.com/controltower/latest/userguide/config.html>)

- » Azure Resource Graph (<https://learn.microsoft.com/en-us/azure/governance/resource-graph>)
- » Google Cloud Asset Inventory (<https://cloud.google.com/asset-inventory/docs>)



TIP

Some third-party cloud security providers offer tools that can map interconnections between technologies in your organization's cloud, including multicloud scenarios, and identify pathways to a breach. These tools can be particularly useful when an organization uses multiple cloud service platforms across its cloud infrastructure.

Once developed, a cloud inventory or map provides a number of benefits, including:

- » **Improved visibility:** When an organization has a real-time view of all its cloud resources and their interconnections, the organization can be more proactive in managing and protecting the environment. The inventory can reveal previously unknown gaps in visibility.
- » **Enhanced resource optimization:** The inventory or map can help reveal how efficiently the organization uses various resources and whether resources may need to be reallocated.
- » **Better planning:** The cloud inventory or map can identify potential resource needs to help in planning future additions or reallocations.

Assessing Which Assets Need Real-Time Monitoring

The cloud inventory or map that your organization creates provides a starting place for prioritizing assets and processes for real-time monitoring. As you might expect, because a typical cloud consists of complex components housing significant amounts of data, the resulting volume of data gathered during monitoring can pose a challenge.

Practically speaking, an organization can't enable, collect, and analyze all the telemetry available from the cloud environment, because challenges arise in the form of both volume and cost.

The organization and SOC team should intentionally make tradeoffs between gaining greater system visibility and managing costs.

Further, the SOC may have to justify investing in logging and analysis with some objective return on investment (ROI) information. The SOC team should estimate the costs for specific forms of logging and the metrics delivered versus the value of the visibility afforded. Estimated savings could also include potential savings from faster incident remediation.



WARNING

Don't forget compliance and regulatory requirements when prioritizing logging costs versus benefits. For example, some fields may require a specific length of time to retain log data and could experience fines if discovered to be out of compliance with regulatory requirements.

The SOC team needs to determine which cloud service provider logging services to employ and understand how to configure these services while controlling costs. Collecting no more data than the SOC team requires while still optimizing coverage will help minimize costs. Depending on the provider, logging configuration choices such as these might be available:

- » Deciding which logs to enable
- » Specifying data retention and archiving times
- » Implementing data compression from certain sources



TIP

Some cloud service providers also enable the SOC team to configure alerts about high amounts of telemetry data or even daily caps on the amount of data collected. Some metrics might be updated and delivered as frequently as every minute.

Evaluating the data store the organization uses to store and analyze cloud data presents another opportunity for cost containment. Traditional security information and event management (SIEM) solutions are often far more expensive than modern data lakes per unit of data volume, and both SIEMs and data lakes require lots of upfront work to transform the data into actionable security insights for the SOC operation. Third-party CDR solutions offer a more out-of-the-box experience, but also leverage a scalable backend to manage costs.

Using Logs and Runtime Data for Real-Time Visibility

Data relevant to cloud security comes from two main sources: telemetry gathered via a cloud service provider's native software tooling that records data in audit or activity logs, and runtime data collected from sensors monitoring device operations and failures. The cloud platform typically provides some level of data aggregation and analysis, as well as creating various metrics about resource performance and health.

With audit logging, the types of log information available can vary from provider to provider. Amazon Web Services (AWS), for example, provides CloudTrail logs for account-related activities across services, VPC Flow logs for monitoring network traffic in the environment, and S3 access logs for recording operations on S3 buckets. Google Cloud Audit Logs offer a more centralized data source about who is performing operations on various assets. In Azure, the available logs include resource logs, activity logs, and Microsoft Entra logs for identity management.

The available tools also vary from provider to provider, and logging configuration is highly customizable. For example, AWS CloudTrail audit logs offer three types of logging at the account level — some of which are enabled by default, and others which need to be enabled manually:

- » **Management event (control plane operations) logging:** Tracks and records overall account changes, such as setting up the logging itself or creating or deleting an Amazon Simple Storage Service (Amazon S3) bucket.
- » **Data event (data plane operations) logging:** Tracks and records operations performed on or within resources in the account, often API activity.
- » **Insight event logging:** Tracks and records unusual activity and errors. The information captured might include items such as the event type, event time period, the impacted API, error codes, and other statistics.

Certain types of assets also by definition call for logging but present cost challenges. For example, VPC flow logs tracking network flow can accumulate a lot of data and cause teams to incur significant cost. Consider enabling these logs only in production accounts or in otherwise sensitive infrastructure areas. Likewise, object storage data event logs play a critical role in detecting data exfiltration (data theft). But there's generally no need to collect these logs for publicly facing buckets serving static content — all you'll get is lots of noise and lots of data you'll need to pay for.



WARNING

On some cloud service provider platforms, certain audit logs are disabled by default, such as data access audit logs in Google Cloud for most services. Your organization's SOC team should review each external provider's logs and enable any additional audit logs that the team identifies as essential.

Generally speaking, the log data collected from all the various sources must then be ingested and consolidated into a centralized log monitoring system. As a necessity, the system or tool centralizes the data into a repository, a necessary step because log files may be delivered in numerous formats, such as JavaScript Object Notation (JSON), common event format (CEF), and more. Cloud service providers often offer a centralized log monitoring tool to handle this function, and third party options are available, too. Security information and event management (SIEM) tools, for example, fall in the class of security solutions that perform this function.

Although the specific cloud activity to monitor and log will be determined by an organization's unique needs, experts recommend monitoring coverage over the following areas at a minimum:

- » **Identity:** Monitor the logs from any identity provider (IdP) such as Okta or Google Workspace used as part of the organization's cloud, and also events relevant to identity from the cloud service provider platform(s) in use. Compromised credentials, role assumption (using temporary security credentials to access resources), and other forms of access manipulation often open up cloud resources to attacks. Routine checks for unusual user access activity are essential.

- » **Data:** Track access logs for data resources that are important to your organization, such as PII databases and sensitive storage buckets, to look for unusual patterns of activity. For example, the AWS CloudTrail service logs S3 storage bucket data events. Cloud service providers also offer logging tools specific to the types of databases hosted on the service, including RDS Audit Logs and others.
- » **Network:** Track IP traffic flowing through a network using NSG and VPC flow logs. Proceed with caution, however. These logs tend to be verbose, holding extensive and detailed information. Selective logging in this instance reduces volume while still prioritizing critical information.
- » **Compute:** Monitoring runtime compute data provides the ability to identify resource hijacking and other suspicious activity. While services such as AWS GuardDuty can provide some compute coverage, third-party sensors deliver a higher level of protection in addition to that provided by logs.
- » **Control:** The cloud service provider control plane offers access and control over so many resources, keeping a close eye on control plane activity remains essential. The cloud service provider should provide central control logging, such as AWS CloudTrail, Azure Activity Logs, and Google Cloud Platform (GCP) Audit Logs.

- » Looking at benefits of rules-based threat detection
- » Applying behavioral analytics to filter out noise
- » Using cloud-native threat intelligence
- » Testing and benchmarking detection measures

Chapter 3

Detecting Attacks in Real Time

Detecting threats and intrusions on an organization's cloud infrastructure requires detailed logging and monitoring. Traditional techniques and tools for on-premises networks drill down on specific network features such as endpoints or network traffic. When used in the cloud, this approach often leads to siloed threat alerts that don't consider the impact of activity occurring on other planes.

This chapter covers the second major area of focus for developing a modern cloud detection and response (CDR) framework in an organization: detect. The chapter begins by covering three essential security approaches in CDR, and then discusses each of those approaches in more detail. The discussion then shifts to best practices for centralizing and monitoring detection findings and finishes with an overview of using attack simulation tools to test and evaluate the effectiveness of detection systems.

Approaches to Detection in Cloud Environments

A CDR framework relies on three complementary techniques (and the tools that perform them) to enhance cloud protection:

» **Rules-based detection tailored for cloud threats:**

Rules-based detection provides the ability to identify patterns of activity in logs that are indicative of malicious activity. A benefit of such rules is that they can be tailored to the needs of your organizations.

» **Behavioral analytics to lower noise in collected data:**

Behavioral analytics are techniques for combining multiple pieces of information about a given entity to develop a more holistic view of whether observed activity patterns are normal or abnormal in a given environment.

» **Context application to prioritize alerts:** CDR frameworks also need to assess alerts in the broader context of the cloud environment. For example, an alert involving assets containing sensitive data should be prioritized more highly than an equivalent alert about assets in a development environment.

By layering and integrating these approaches, a security operations center (SOC) team can achieve greater detection accuracy and develop better policies and automations for responding in real time. An integrated approach resulting in fewer false positives also frees up team capacity to investigate and contain dozens of critical issues per month, preserving operational capabilities and reducing the cost impacts of any breaches.

Using Rules-Based Detection Tailored for Cloud Threats

Detection rules do the work of identifying potentially malicious activity in a cloud-based system. Detection rules are based on activity that can be an indicator of tactics, techniques, and procedures (TTPs) that threat actors use to carry out attacks.

Rules-based protection provides clarity and explainability, because each rule corresponds to specific conditions for decision making. This explainability can begin to provide some context about the potential nature of a detected threat.

Rules-based approaches should be informed by threat intelligence: detailed data about the types of TTPs that attackers are using in the wild. Threat intelligence can be sourced through various means, but third-party providers provide feeds in standardized data formats to enable easy creation of detection rules using the data. When evaluating threat intelligence in a cloud context, it's important to ensure that cloud-specific indicators of compromise (IoCs) are emphasized in the intelligence feeds. For example, although malware is common in both traditional and cloud environments, identity-based attacker tactics are especially powerful in the cloud and cloud threat intelligence should provide information about attacker tactics in these areas.

The SOC team doesn't have to start from scratch in building rules-based detection for cloud assets. As would be expected, major cloud providers offer prebuilt rule sets as part of their detection tools and services offering. For example, the Amazon Web Services (AWS) GuardDuty service includes rules that a SOC team can leverage. Third-party vendors now offer even more tools for creating custom detections.

In addition, many regulatory frameworks and security standards require continuous monitoring and real-time threat detection. Implementing runtime detection for workload-level events helps organizations meet these compliance requirements and demonstrate a commitment to security.



TIP

Some runtime solutions offer the ability to leverage custom or preconfigured response policies to automate taking immediate action when a rules-based threat is detected. As infrastructure expands, automated systems can handle the increased volume of threats without the need for additional manual resources.



TECHNICAL
STUFF

The MITRE ATT&CK framework compiles knowledge of attacker TTPs for the cybersecurity industry. It graphs the TTPs used in various types of attacks by various types of attackers. Organizations should have a mix of rules that cover the entire matrix.

To ensure the best detection coverage, rules should be tailored for your organization's environment, and cover both runtime and cloud-level activity. Runtime detections focus on processes, files, and other workload-level elements. Rules are often created with boolean operations, string operators, and regular expressions. The SOC team can create rules that detect unusual activity based on runtime events such as the following:

- » **Process execution:** A process begins the execution of another.
- » **Network connection:** A process establishes an outbound connection.
- » **DNS query:** A process initiates DNS lookup activity.
- » **Network listen:** A process initiates network port listening activity.

Cloud-level detections often take the form of a query. Detection engineers can save queries over their log data, and alerts will trigger when results matching that query are returned. Cloud events used to detect malicious activity vary per cloud provider, but third-party CDR tools offer pre-built rules out-of-the-box.



WARNING

As with other aspects of the cloud, detection rules can become complex, with dependencies on other rules, functions, and various log sources. The SOC team should develop policies for maintaining and updating detection rules on a regular basis.

Maintaining rules also requires staying on top of emerging cloud threats. Major cloud service providers, researchers, threat intelligence providers, and other publications often publish lists of cloud-specific IOCs.

Lowering Noise with Behavioral Analytics

Because typical cloud monitoring in a CDR system involves collecting substantial volumes of log and runtime data, rules-based detection alone still might be noisy, generating many low-fidelity alerts that aren't valuable for SecOps teams. These false threat alerts can lead the SOC team to waste time chasing down non-issues.

Behavioral analytics refers to analyzing data to understand user and system behavior, in this case behavior in the cloud environment. Behavioral analytics tools typically include or can be augmented with artificial intelligence (AI) and machine learning (ML) models to perform a deep analysis of large amounts of behavior data captured in logs tracking items such as:

- » Control plane events
- » Network traffic flow
- » Database changes and activity
- » User sign ins and other user activity

The first part of the behavioral analysis identifies what constitutes normal, standard baseline pattern of network activity. From there, the behavioral analysis system can evaluate data in real time and identify and warn of usage behavior that differs from the baseline pattern, including potentially harmful behavior. Such systems also can usually specify the magnitude of how far the suspicious activity differs from the baseline pattern, enabling the SOC team to focus on the most severe potential threats.

A key fact about cloud environments is that users and systems generally perform many activities on a regular, predictable basis. Because cloud-native architectures often make use of microservices with highly specialized functions, individual cloud resources have much narrower responsibilities compared to on-premises. For example, a specific pod may only ever pull data from a specific database — and never perform any other activities. A deviation from that pattern would be easily spotted by behavioral analytics tools and identified for investigation and response. This makes behavioral analytics a strong detection tool to use in cloud environments.

All that said, behavioral analytics can still lead to false positives and false negatives in terms of identifying suspicious activity. Using it as part of a more comprehensive CDR strategy can help the SOC team compare alerts with other available information to determine whether a particular threat requires action.

Prioritizing Alerts Using Context

Most cybersecurity systems, including CDR platforms, display or send alerts (notifications) when suspicious activity has occurred. The SOC team can use the alerts to determine how serious a threat is and where to focus a response. Unfortunately, many solutions produce a high rate of false positives.



WARNING

When a system produces an overwhelming number of alerts, SOC team members can experience alert fatigue. Alert fatigue can reduce the team's ability to respond to the threats that really matter, as well as creating burnout among team members.

The SOC team can prioritize alerts using context from the cloud environment. In this case, *context* means setting up an alert to correlate multiple risk factors, so that the alert accurately prioritizes only the most critical incidents. For example, let's say your organization uses Amazon Simple Storage Service (S3) to store information. If your organization receives an alert rule about anomalous access to an S3 bucket, your team should prioritize that alert more highly if the S3 bucket stores sensitive customer data. If the bucket stores only static files for a public website, your team may not need to address that alert as quickly.

The SOC team should prioritize real-time alerts by correlating different sources of information about the cloud environment, including:

- » **Host-level information:** Does the alert involve a compute asset with a high-severity vulnerability?
- » **Identity information:** Does the alert involve a highly privileged user?
- » **Data inventory information:** Does the alert involve data assets containing PII, PHI, or other sensitive data?

Cloud service providers may not provide the capability to correlate information to spotlight the most critical risks in an organization's cloud infrastructure. However, other cloud security vendors do provide third-party solutions for contextualizing and prioritizing alerts.

Aggregating and Monitoring Findings

Having a SOC team member spend valuable time monitoring multiple individual dashboards wastes valuable time. Aggregating and monitoring findings in a centralized service or tool provides a best-practice solution for monitoring threats.

Centralizing your organization's threats into a single service or map enables the SOC team to monitor the environment more effectively. In a cloud environment with multiple sources of detections and alerting, centralizing provides concurrent threat visibility. Depending on your organization's cloud and SOC team setup, a centralization solution might be as simple as a basic Slack workflow, a lightweight dashboard, or a full-blown enterprise Security Information and Event Management (SIEM) solution.

Benchmarking Detection Coverage

Many industries apply some form of stress testing to evaluate the health or status of the organization and its systems or resources. The IT world uses various forms of stress testing, too, to evaluate system functions. For example, on a more traditional network, stress testing tools generate excessive network traffic to test the network's operating limits and quantify the level of traffic that might cause the network to fail.

In cybersecurity, stress testing often comes in the form of attack simulation. Attack simulation involves the SOC team and/or a third-party organization imitating a hack against its own infrastructure and assets. In an attack simulation, the security team aims to determine the security health of the network and uncover gaps in detection and response capabilities to target for improved security measures. These simulations usually involve a red team and a blue team:

- » The red team plays the role of the attacker, seeking to penetrate the network and access network resources.
- » The blue team plays the role of the defender, identifying the incoming attack(s) and initiating simulated attack responses.

Because the cloud environment can be subject to hundreds of unique threats and because the infrastructure and assets within any given cloud can vary widely, especially depending on which cloud service provider(s) an organization uses, an organization may need different tools to simulate attacks as part of its CDR process. A variety of open source red team tools and frameworks for simulating attacks exist, including Stratus Red Team, Microburst (which is a PowerShell toolkit for simulating Azure attacks), and Pacu (a framework for simulating AWS attacks).

The SOC should use any simulated attacks to benchmark detection coverage throughout the cloud infrastructure and take appropriate corrective action to secure any vulnerabilities revealed during the testing.

- » Applying automation to contextualize data
- » Constructing an incident timeline
- » Determining the scope of an incident

Chapter 4

Investigating Incidents

After verifying an attack on an organization's cloud, the security operations (SecOps) team should immediately move on to conducting an investigation that includes examining and recovering data needed to assess and document the attack.

This chapter examines the investigation aspect of cloud detection and response (CDR). The first section of the chapter provides guidance about automating the process of gathering forensic data to contextualize the incident. Next, the chapter covers how to create a timeline of the incident and document events and actions. The chapter concludes with a discussion of how to assess the scale and scope of the incident.

Starting the Investigation

The first part of any modern CDR framework — prepare — focuses heavily on gaining the right real-time visibility, using both cloud log telemetry and runtime sensors, throughout the organization's cloud infrastructure (see Chapter 2 for more.) Real-time threats can be detected using these data (see Chapter 3), but this information also should be captured and stored for a period of time for later use in investigation and response to any incidents that are detected.

When faced with a potentially malicious alert, the SecOps team needs to begin investigating and gathering evidence immediately. Sources of data relevant for cloud investigation include the following:

- » **Cloud service provider logs and records:** These include items covering your organization's cloud assets, such as audit logs, network flow logs, container orchestration logs, and more.
- » **Workload data from runtime sensors:** Sensors deployed on cloud workloads provide valuable runtime execution data to help identify malware or other malicious activity on the workload.
- » **CNAPP findings:** Some organizations use a cloud native application protection platform (CNAPP) solution to consolidate security coverage in multicloud infrastructures. These tools can provide cloud-specific context about the resources involved in the alert (for example, whether or not they contain sensitive data or have any exploitable vulnerabilities).

Understanding Who to Involve

Cloud investigations require close collaboration between Security Operations and DevOps and developer teams. Although Security Operations is primarily responsible for the investigation, DevOps and developer teams often have critical context about the environment that's required to understand the triggered alerts.

SecOps teams should be prepared to work with dev teams to investigate alerts, with a clear delineation of responsibility and pre-built workflows for communication and collaboration. SecOps teams should also ensure they are able to quickly and accurately identify relevant application owners for each part of their infrastructure to avoid any investigation delays.

Internal users of your organization's cloud resources can also provide evidence and context about unusual activity observed during the attack time frame. Reach out to these users and ask for their recollections, to confirm suspicious activity, or other evidence (screenshots or experiencing sign in difficulties or service interruptions, for example) they may have observed.

Automating Context Gathering

After relevant data have been gathered, the SecOps team should begin its analysis, focusing on the entities in the threat. They should try to build context about those entities and ask questions using their data. For example, teams should evaluate if the activity observed in the alert is relatively common or unusual for the involved entities. They should ask what other actions these entities were taking around the same time period of the alert. They should ask if any the suspicious entities had interacted with any other cloud resources — which may now in turn be compromised. And so on.

If after this process has concluded the SecOps team is able to confirm that malicious activity has occurred, the team should isolate or shut down the affected systems and capture forensic artifacts for later analysis.



WARNING

Without logs and other data covering all the actions and events involved in the breach, the later efforts to contain and analyze the breach will likely be stymied. Consider this need for logging in incident response when making proactive visibility versus cost tradeoff decisions when assessing which telemetry to collect.

As might seem obvious, gathering and correlating data from a variety of sources (multiple cloud service providers, multiple different cloud services, security tools for multiple different cloud storage and service objects, and so on) can be complex and time consuming.

Depending on the structure of the organization's cloud, the SOC team might be using a variety of different security tools from the cloud security provider and third parties in the cloud security stack. These tools comprise a whole alphabet soup of security tools: cloud security posture management (CSPM), Kubernetes security posture management (KSPM), cloud workload protection platforms (CWPP), vulnerability management, container and serverless security, cloud infrastructure entitlement management (CIEM), infrastructure-as-code (IaC) scanning, and more.

A cloud service provider may offer a tool on its platform for automating security incident investigation, such as Amazon Detective for AWS, but such tools cover only that platform, leaving the SOC team to invest time mining other security tools and raw telemetry for threat incident data.

Where possible, the SOC team should automate gathering event data from the various services, logs, and sources to minimize the time incident responders need to spend on repeated information gathering tasks. Some solutions provide for automating context and data gathering on an alert-by-alert basis. Each of these approaches would add automation to the data gathering process:

» **Security orchestration, automation, and response (SOAR)**

tools: SOAR tools automate and streamline incident response processes, including data gathering. These tools can consolidate and automate certain types of security activities, including threat response data gathering. Some cloud service providers offer SOAR-type tools, while others are available from third-party security providers. SOAR tools can enable teams to automate enrichment workflows, including use of third-party or open-source plugins to enrich alerts.

» **Cloud investigation and response automation (CIRA):**

CIRA enables and automates the collection and analysis of cloud forensic data. CIRA tools allow organizations to automate investigation and response to true-positive cloud incidents. By leveraging sophisticated data analysis techniques and machine learning (ML) algorithms, CIRA enables proactive threat detection, empowering organizations to stay ahead of cyberthreats. Many traditional security methods fall short in dynamic cloud environments facing novel threats and risks. CIRA is a crucial tool that effectively mitigates these risks. It can analyze log and event data automatically, helping to thwart security incidents before they escalate into full-fledged breaches across the cloud.

» **CNAPP tools:** As touched on earlier in the chapter, these third-party systems provide the most robust data gathering automation in complex cloud environments.

Creating and Maintaining an Incident Timeline

As the organization's SecOps team gathers evidence about a security incident, the team should organize all this information into a single incident timeline. The incident timeline gathers all the evidence into one document or chart, arranged in chronological

order, so that investigators can quickly grasp the full scope of the event.



REMEMBER

At a minimum, the timeline should identify the start of the attack, the end of the attack, how much and if or when data was stolen or released without authorization, and the cloud systems and resources that were compromised during the attack.

Organizing a single source of truth

The SOC team investigators should utilize the incident timeline as a single source of truth, or a central, definitive repository of the data.

Investigating a security incident often involves team members from various parts of the organization. The incident timeline helps keep the team members on the same page and can help members joining the team get up to speed quickly.

Decide where and how the team will collaborate to investigate and resolve an incident. The SOC team may want to establish a physical or virtual war room where investigators and subject matter experts (SMEs) can work together and communicate directly. The team also might communicate and collaborate on the timeline via a chat platform or Slack channel, for example.

While for large-scale incidents a dedicated timeline that the team constructs manually is warranted, for day-to-day use, some CDR products can help automate the creation of an incident timeline much more quickly than can be accomplished using manual approaches.



WARNING

Some sophisticated security breaches can take months to detect. When building the incident timeline, be sure to include a review of data in the period leading up to the alert. Otherwise, the timeline might lack important root cause information.

Expanding your view of an attack

As investigations progress, it becomes important to document the context of an incident with additional data. The types of data to seek out and use to annotate the incident timeline might include:

- » **Additional system behaviors:** What else were the users and resources involved in the incident doing around the time

of the incident? Did they communicate with other assets that might also be compromised? Did they touch any sensitive data?

- » **Unusual access behaviors:** Were there multiple failed login attempts? Did the activity come from a country that's infrequently active on your organization's cloud? Such behaviors can provide clues to the meaning of other data.
- » **Originating entities (identities):** In the cloud, identities often become obscured by complicated role assumption chains. Can you identify which actual users were ultimately behind each action you observed?
- » **Gaps in data and coverage:** Was the data needed for the timeline available from the expected sources? Did the team encounter poor record keeping, nonexistent logs, or other gaps in the record? Were the log files themselves compromised in any way? This may be indicative of evidence destruction on the part of the attacker.

CHAIN OF CUSTODY ISSUES

When a breach results in a crime such as data or identity theft, preserving the data (evidence) documenting the incident becomes even more critical. It is important to have procedures in place for managing the chain of custody. Chain of custody is a critical record keeping step to ensure that evidence collected in the course of the investigation can be relied upon in any future legal matters.

In a cloud environment, evidence of a breach may span multiple services, regions, and cloud providers. CDR tools can help maintain the chain of custody by collecting all relevant data and documenting key steps in the investigation process.

An organization should develop detailed forensics policies and procedures for evidence preservation, as well as understanding the rules and law in the jurisdictions where it purchases hosting or operates cloud assets.

Analyzing the Blast Radius

Cybersecurity professionals call the total impact of a security event the *blast radius*. Blast radius refers not only to the volume of cloud resources damaged or compromised, but also the overall seriousness of the attack. For example, having nonsensitive data of a few users compromised pales in comparison to having personally identifying information about a few hundred thousand users stolen and released online.

The SecOps team can examine the following factors to determine the blast radius:

- » **How many resources were impacted:** Identify resources that have interacted with the entities involved in the attack. Let's say the SOC team received an alert about anomalous access by an Amazon Elastic Compute Cloud (Amazon EC2) instance to a particular Amazon Simple Storage Service (Amazon S3) bucket. The team should assess what other buckets that particular EC2 also interacted with around the time of the alert.
- » **The attack path used:** How many cloud resources the attacker moved through provides an indication of the severity of the attack. An attack path that leads to sensitive data or a highly privileged cloud role also makes the breach more significant.



TIP

Some third-party security providers offer automated tools for blast radius analysis.

- » Containing the spread of a threat before business impact
- » Understanding cloud forensics
- » Using playbooks to automate collecting and analyzing evidence
- » Coordinating remediation efforts to fix root causes

Chapter 5

Containing and Responding to Threats

Investigating a security breach would be fruitless in the absence of taking action to address the breach. An organization's security operations center (SOC) team stands on the frontlines to put the fire out, so to speak. Containing the breach is a critical first step, but merely putting the fire out doesn't prevent the next fire. The SecOps team needs to perform detailed forensics to analyze the cause of each breach, developing important insights to guide data gathering and analysis automation choices, remediation steps, and security planning efforts.

This chapter covers the final major area of focus in modern cloud detection and response (CDR): respond. Discussion of the primary response, applying containment measures to contain the incident, comes first in the chapter. Moving on, the chapter discusses the importance of forensic investigation and remediation to prevent of future incidents.

Applying Containment Measures ASAP

Even while an organization's SOC team investigates a breach and builds an incident timeline, the team should implement measures to contain the security incursion and resulting damage. The team should first work to rapidly verify that the alert represents a true breach. If the SOC team is confident that the alert is a true positive (a real threat), it can move immediately to take actions such as:

- » Changing security group settings to isolate the involved cloud resources from the rest of the cloud environment to prevent the attack from spreading.
- » Rotating credentials for suspect identities at either the cloud service provider or identity provider (IdP) level.
- » Stopping or terminating compromised virtual machines.



WARNING

Before implementing any of the above actions, the SOC team should assess the potential impact each action would have on the organization's production environment. The team should also put measures in place to avoid downtime.

The SOC team should prebuild these containment actions so that they can be deployed quickly, as well as customize them for the organization's particular cloud environment. For example, depending on the cloud resources being used, the prebuilt actions might include:

- » Command line interface (CLI) or Terraform code snippets that the team can run manually.
- » Templated tickets for DevOps teams to apply changes at the cloud service provider level.
- » Fully automated playbooks that can apply changes immediately when certain criteria occur.



TIP

Many cybersecurity teams create a detailed incident response plan. Planning in advance can help reduce response time and clarify team member roles. If your organization doesn't have the bandwidth to develop its own plan, the company can hire experts in threat hunting, response, investigation, preparedness, and recovery to help develop a plan and respond to an incident at a time of crisis.

What Cloud Forensics Entails and Why It's Important

Forensics in cybersecurity is the process of extracting and recovering evidence after an attack, followed by a thorough analysis of the data gathered to develop a comprehensive understanding of the breach. Gathering information about a security incident enables the organization to identify the sequence of events or actions that occurred during the incident.

Forensic data in the cloud is composed of both cloud logs and traditional forensic artifacts (traces of malware and other software tools, system changes, timestamp changes, IP addresses, registry entries, and so on). The forensics team members complete a deep investigation of the data, artifacts, and other soft data from user interviews and other sources, finally documenting the information and analysis in a detailed report. The investigation provides the knowledge and insights needed to contain the incident, as well as performing other future mitigation actions and closing security gaps in the cloud infrastructure.

Every investigation begins with asking questions. In the case of investigating a cloud security incident, several questions can provide guidance about how the investigation and forensic analysis should proceed, including:

- »» What was the initial access point?
- »» Did the attacker spread to other resources, and if so, how?
- »» Did the attacker gain persistence in the environment, and if so, how?
- »» What is the impact of the breach?
- »» How can this be prevented next time?

In addition to answering the above questions, the analysis and reporting should cover the full scope of the incident timeline, including any reconnaissance or initial access that may have occurred well before the incident was finally detected. The report should not only document the impacts to various cloud resources and how the attacker moved through the complex cloud

environment, but also seek to understand the attacker's mindset to help connect the dots between seemingly disparate events in the chronology of the attack. The forensics team should also keep in mind that the evidence gathering and analysis could play a role in any charges or court cases that arise from the event and be ready to assist involved law enforcement and legal teams with evidence gathering and analysis.



TIP

At a minimum, the forensic report should include an executive summary, a brief overview of the incident, details about evidence collection, and a chronological timeline of how the attack unfolded. The report also may include findings (interpretation of the data), root cause information, lessons learned, and recommendations for future preventative measures.

Ensuring Forensic Data Availability

Preparation before a security breach serves the essential purpose of making sure that the necessary data for investigation and forensic analysis is readily available. Because so much forensic data normally exists only in the very same cloud environment that may be attacked and may not be retained by default, an organization's SecOps team needs to develop a strategy for obtaining and securing log data.

The organization leadership and SOC team should follow best practices for collecting and securing log data, including:

- » Develop and apply log retention policies appropriate to the organization's needs to ensure the availability of critical information. Often, forensic analysis requires looking backwards a over significant period of time before the incident to detect slow-moving threats. Some threats may take up to several months or more to unfold.
- » Ensure the organization's data gets pulled into a searchable, centralized location, such as a security information and event management (SIEM) tool. Some cloud service providers offer native SIEM tools, and third-party solutions exist, too.

- » Alternatively, the SOC team can consider establishing a data lake to store forensic data. A data lake is a modern, general-purpose data storage tool, and today's data lakes are often built with more scalable and cost efficient architectures than SIEM solutions. Third-party CDR tools can run on top of these data lakes to make searching the data and forensic analysis faster and easier.

Best Practices for Dedicated Forensic Environments

With log data preserved in a central place and appropriate retention policies applied, organizations can turn to management of forensic artifacts like machine images. Images of the cloud storage components deployed in an organization's cloud can yield valuable forensic information. When analyzing the images, the SOC forensics team needs a clean, standardized, preconfigured environment to investigate potentially malicious artifacts captured in an image. The images should be copied or transferred to a location or environment that's separate from the organization's production environment.

Establishing a dedicated forensics account with a cloud service provider can provide this separate environment. The forensic account is a new, separate account from the organization's main production account(s) with a cloud service provider.

A forensic account can help with:

- » Implementing access restrictions or legal holds to comply with retention or chain of custody requirements
- » Facilitating data transfer to law enforcement investigators, legal entities, or auditors
- » Isolating and sandboxing for malware analysis

At least one cloud service provider, Amazon, offers a prebuilt playbook for collecting forensic evidence. Called Automated Forensics Orchestrator for Amazon EC2, the tool enables the SOC

team to capture data from Elastic Compute Cloud (EC2) instances in the event of a potential security breach. It also can perform other actions, such as isolating an impacted EC2 instance and using automated investigation tools to analyze artifacts.



TIP

As it analyzes incident data, the digital forensics team should also compile a list capturing each indicator of compromise (IOC) left by the attacker in the environment. Each IOC identifies a specific malicious activity or artifact, or combinations of the two, that indicate a specific attack or attack methodology. Having a list of recognized IOCs enables the SOC to improve detections and alerting and speed future investigations.

Developing Automated Playbooks to Collect and Analyze Artifacts

A simple cybersecurity playbook outlines the steps the SOC team will take in the event of an incident. An automated playbook takes the concept a step further, defining preconfigured sets of remediation actions performed by scripts or code. Because security incidents often involve the same containment actions, such as isolating a compute resource or blocking an account, automating these activities speeds the threat response.



TECHNICAL
STUFF

Many playbooks fall under the umbrella of infrastructure-as-code (IaC), especially ones that function as provisioning tools.

The SecOps team can run playbooks manually, when needed. But cloud resources are ephemeral, making timely response and data collection critical. As such, some playbooks can be set up to run automatically when triggered by specific alerts or other trigger rules. In addition to performing containment measures, playbooks can assist with data gathering and analysis.

Consider these key points:

- » Ensure the SecOps team has deployed playbooks to automate taking image snapshots and copying them into your organization's dedicated forensics account.

- » Consider conducting analysis using open-source tools, such as Plaso, to create a comprehensive timeline of activity occurring on potentially compromised machines.
- » Consider using a third-party cloud investigation and response automation (CIRA) tool as a source of automated playbooks and other investigation and forensics tools. CIRA solutions typically boast a diverse library of prebuilt playbooks for common types of breaches. The SOC team can typically customize the available playbooks, or build new custom playbooks based on available templates. Some cloud service providers also offer services with playbook templates.

Security operations teams with more sophisticated threat models should make sure not to overlook gathering memory data for memory forensics. Memory forensics is the process of investigating memory contents to identify threats and forensic evidence. Memory forensics provides value in analyzing more sophisticated attacks, for example those leveraging fileless malware that executes and resides entirely within volatile memory. Important points in memory forensics include:

- » Volatile memory artifacts can be more challenging to collect, often requiring that an agent be deployed.
- » AWS provides guidance and best practices for capturing memory data using the AWS Systems Manager Agent (SSM). However, this also requires the SecOps team to configure the orchestration process using several AWS services.
- » Alternatively, consider using a third-party provider that can capture memory either through a pre-deployed sensor or that can deploy a collector on-demand to ensure capture of all data.

Collaborating with Developers on Remediation

Once the threat has been contained and the forensics team has gathered and analyzed data, identifying root causes as part of the forensics report, the remediation process can begin. Root cause remediation often requires changes to source code.

The SOC team must identify vulnerable compute instances, CSP misconfigurations, or other issues and trace these back to IaC, application code, or other source code and correct the issues that became the root cause for the attack. The SOC team should work with developers to remediate these issues and situations:

- » To encourage faster remediation, suggest process fixes using typical developer tools, such as ticketing systems, pull requests, and so on.
- » Ensure that remediation requests are actionable. For example, if the SOC team requests that a package be upgraded, include the minimum version required to resolve the vulnerability. If the SOC team requests a change to cloud configuration, include the relevant IaC snippet in your ticket.

Chapter 6

Five Best Practices for Cloud Detection and Response

The cloud threat landscape is shifting as businesses continue to adopt the cloud at increasing rates. To defend against these new threats, effective CDR is essential to safeguard your organization's cloud environment.



TIP

A third-party CDR platform can perform most of the best practices described in this chapter out of the box. An optimal CDR solution will use a single GUI that integrates breach readiness recommendations, detection and alerting, investigation, and containment and forensics.

As you build out a CDR framework at your organization, the best practices here can help you design your CDR program effectively from the ground up:

- » **Gain Real-Time Visibility:** Gain consistent real-time visibility over your organization's entire cloud environment, including both runtime data gathering and taking advantage of cloud service provider-native telemetry. You should make sure that you're collecting telemetry covering all critical elements of your infrastructure, and you should reassess your telemetry collection on an ongoing or at least periodic basis to make

sure no drifts have occurred, and no new blind spots have opened up.

- » **Benchmark Detection Coverage:** Benchmark your detection coverage using open source *red-teaming* (attack simulation) tools. Probing the vulnerabilities in and access points to your cloud environment ensures your security operations center (SOC) team can plug gaps and expand coverage for your organization's threat model. Commonly used open-source red teaming tools include Stratus Red Team, Pacu, and MicroBurst.
- » **Accelerate Context Gathering:** Accelerate context gathering using automated workflows for each alert. The SOC team should embed this context by correlating multiple conditions within each alert to reduce the need for manual querying. For example, at a basic level, teams should make sure that relevant cloud console pages for involved entities are linked to analysts in each alert. More advanced CDR platforms will surface context from the cloud environment within an alert — for example, showing whether an affected asset contained sensitive data or whether the behavior in question fits normal patterns without requiring manual querying.
- » **Build Automated Threat-Containment Playbooks:** Build automated threat-containment playbooks for each entity in your organization's cloud environment and ensure your SOC team can run each playbook immediately when the need arises.
- » **Automate Data Collection and Analysis:** Use tools to automate gathering and analyzing cloud forensic data to ensure critical information is preserved and each attack is fully eradicated.

Detect, Investigate, and Respond to Cloud Threats *Quickly* with Wiz



They say a demo is
worth a thousand words.
Watch Wiz in action
at wiz.io

WIZ★

Reduce false alerts, help teams prioritize investigation, and enable proactive updates!

This book introduces key aspects of cloud detection and response (CDR) and how it differs from other cybersecurity efforts, providing an overview of CDR for business leaders and security operations professionals. It outlines practical best practices for real-time detection and response to cloud threats.

Inside...

- Evaluating the threat landscape in a cloud environment
- Detecting threats targeting cloud-native infrastructure
- Investigating the scope of an incident
- Containment and response to cloud threats
- How to conduct cloud-native forensics
- Best practices for CDR

WIZ[★]

Author and content expert **Lisa A. Bucki** has educated others about computers and business topics since 1990. Bucki is cofounder of 1x1 Media, LLC (www.1x1media.com), which creates content on how-to topics for entrepreneurs, startup founders, makers, and other business professionals.

Go to **Dummies.com™**
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-394-30413-4

Not For Resale

for
dummies®
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.