Ultimate Guide to Incident Response **in GCP**

Introduction

Google Cloud Platform (GCP) presents its own landscape of challenges and opportunities when it comes to incident response. With a vast array of cloud services and tools at your disposal, navigating security incidents demands a nuanced approach. Unlike traditional on-premise environments, GCP's diverse range of over 100 services brings unique complexities to incident investigations and resolution.

Each GCP service generates its own set of logs and data sources, scattered across platforms like Cloud Logging, Cloud Monitoring, and Cloud Storage. Understanding the nuances of these data sources is crucial for effective incident response, but makes a day in the life of SOC analyst extremely complex. At the same time, the cloud unlocks innovative avenues for swift evidence collection, analysis, and resolution.

In this playbook, we aim to demystify the incident response process within GCP. We've compiled comprehensive guidance on addressing security incidents you're most likely to encounter across the most frequently used GCP services. Our focus is not just on remediation but on empowering you with the knowledge to efficiently leverage GCP resources for a streamlined incident response lifecycle.

Before the Incident

Being impacted by an incident is a ‘when’, not ‘if’ situation, so ensuring your organization is prepared to investigate and respond to potential cloud threats is vital to appropriately managing risk.

A proactive approach to cloud incident response enables security teams to understand whether they are prepared to quickly investigate and respond to threats before an incident occurs. This ensures that when an incident does occur, the security team will have the ability to quickly identify the root cause and remediate the threat. Being proactive also enables security teams to continuously improve their incident response program by preemptively identifying and rectifying any existing visibility gaps not to waste value time during the heat of an incident.

Know Your Data

Identify your crown jewels. Do you have particularly sensitive information, like Personally Identifiable Information (PII) or Payment Card Industry (PCI) data? If so, you need to know exactly where it lives and what systems process that data. This also includes any backups or logs that might shadow the original data.

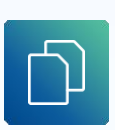

Have Backups and Test That They Work

A disaster recovery plan can mitigate not just security incidents like ransomware, but also other likely events such as data center hardware failure. Ransomware is a high risk due to both high impact and relatively high likelihood of occurrence.

Restrict Administrative Accounts

In general, follow the principle of least privilege. In particular, Google provides detailed advice on how to secure administrative accounts in IAM.

Require Multi-Factor Authentication for all User Accounts

This can be easily enabled by following this guide from Google.

Limit Network and Remote Access

Limit any connectivity to the internet from your machines as much as possible. A common security issue is machines with RDP accessible from the internet. This can put you at particular risk of brute-force ransomware attacks. Google has some great advice on how to set up RDP connections if they are required in your environment, such as Using IAP for TCP forwarding so that machines do not have to have external IP addresses.

Encryption

The general advice is to ensure data is always encrypted at rest and in transit. There are open discussions around how useful encrypting data at rest is with some cloud services. However, you may have particular requirements here if you are in a regulated industry such as finance or healthcare. Google offers advice on which of these services comply with different compliance standards.

Enable Logging

“Forensic readiness” will help you not only detect incidents earlier but also make investigations more thorough and efficient. As you can imagine, the more useful data you have, the more likely you will be able to find the root cause of an incident. Ensuring you have the right logs enabled can make all the difference.

GCP divides its logs in to a number of categories and subcategories:

- **Platform logs** serve as records generated by various Google Cloud services, aiding in the identification and resolution of issues while offering deeper insights into the functionalities of the employed Google Cloud services. Take VPC Flow Logs for instance, which capture network flow samples transmitted to and from VM instances.
- **Component logs**, akin to platform logs, stem from Google-provided software components operating within your systems. As an example, GKE deploys, by default, a per-node logging agent that reads container logs and adds helpful metadata These logs originate from the user's GKE instances, transferring to the associated Google Cloud project, and are instrumental in providing user support through log analysis and their metadata.

- **Security logs** play a crucial role in addressing the 'who, what, where, and when' of system activities:
    - **Cloud Audit Logs** provide insights into administrative actions and accesses within Google Cloud resources. Their activation fortifies security measures, enabling constant monitoring against potential vulnerabilities or external data misuse by security, auditing, and compliance entities. Refer to the list of Google Cloud supported services for available audit logs.
    - **Access Transparency Logs** detail actions performed by Google staff while accessing your Google Cloud content. These logs facilitate compliance tracking, aligning with your organization's legal and regulatory requirements. Review the supported Google Cloud services to access the corresponding Access Transparency logs.

- **User-written logs** comprise entries generated by custom applications and services. Typically directed to Cloud Logging, these logs utilize methods such as:
    - Ops Agent/Logging agent
    - Cloud Logging API
    - Cloud Logging client libraries for their integration

- **Multi-cloud logs** and **Hybrid-cloud logs** encompass entries from alternative cloud providers like Microsoft Azure, along with logs originating from on-premises infrastructure.

Both Data Dog and Secure Works have great tutorials on how to ensure full logging is enabled in your GCP environment.

Be Prepared

Periodically run tabletop exercises to simulate incidents and build muscle memory across both executive and operational teams.

Executives should be prepared to answer the following questions in advance of any incident:

- Under what circumstances do you notify law enforcement, regulatory authorities, auditors and the board?
- Will we pay a ransom? If so, how?
- If required, which outsourced incident response firm will you work with?
- If you lose access to core IT systems for an extended period of time? Do you have business continuity and disaster recovery plans in place?
- If the primary communication methods are either unavailable or compromised, do you have backup or out-of-band communications available?
- What working hours are incident responders expected to work in a high-severity incident?
- Do you have access to the data required to perform an investigation in all products and services?

Responding to the Incident

Gather the Incident Response Team

As part of incident response planning, organizations should craft a well-thought-out and rehearsed incident response and crisis communications plan with defined roles and responsibilities mapped out to limit the overall impact should an incident occur. This includes preparing internal teams and external incident response service providers in steps to take and actually exercising the plan end to end regularly. Ideally, this plan has been pre-approved during incident response planning so that incident response actions can kick off as soon as possible post-incident identification.

The roles in an incident response team will vary depending on both the size of your team and the scale of the incident. Most often, one person will take on a number of roles. A typical example of the roles in an incident response team is:

Leadership role - Commands the investigation and directs activities.

Investigator role - Identifies incident root cause and the full scope of compromised systems and data.

Responder role - Works with internal teams and 3rd parties to recover and restore systems and services and plan and coordinate remediation steps.

Documentation role - Enables the investigation, remediation and potentially legal representation. The legal representation may also be handled by inside or outside counsel (though only a small number of incidents end up bringing in a legal representative).

Understand the Environment

It is important to gain an understanding of the environment in which the incident occurred. If you are an internal SOC, you may already know the answers to these questions in advance of an incident:

- Where is sensitive data stored?
- How are users connected to what services and instances?
- Who are the administrators?
- Where are logs stored?
- What GCP Products and Services are in use?

Collect the Right Data

You can’t investigate what you don’t have access to, so it’s important to ensure you have the right access and you know which data sources will be most valuable in your investigation prior to an incident occurring. Once an incident is detected, it’s important to collect information relating to systems that may be compromised, including both meta-data (typically logs) and full content (disk images, volatile data, etc.). You will need to carefully scope this phase as it can prove difficult to find the right balance between collecting too much and collecting too little. A thorough investigation can require a lot of data - so collecting data in phases to gradually narrow the scope of your investigation is important. For example, it’s useful to perform an initial triage collection across a larger set of systems to determine which systems require a more in-depth full disk analysis to increase overall efficiency while getting the answers you need.

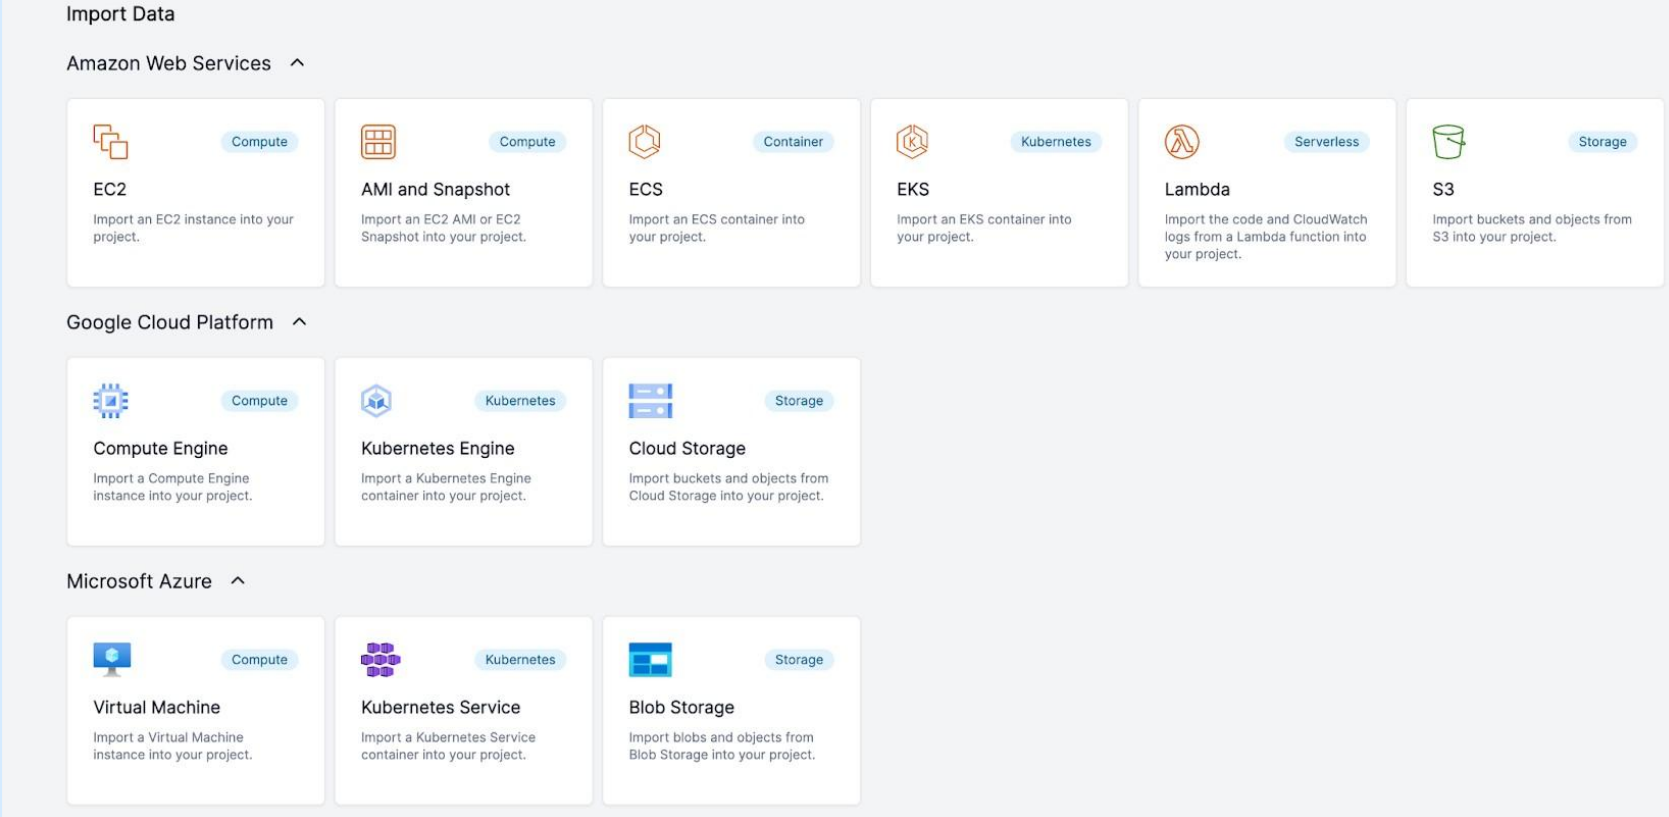

The Cado Platform can automatically collect and analyze full copies of data from GCP Compute Engine, Kubernetes Engine and Cloud storage as well as data sources from AWS and Azure in a single click, see the Cado in Google Cloud section for more information.

Perform the Investigation

First, identify the scope of the investigation by answering the following questions:

- Do you just need to recover services?
- Do you need to identify the root cause of the incident so it doesn't happen again?

Most investigations start with a suspicious event - such as a detection for malware on a system. And then the investigation progresses as you pivot based on timestamps or key findings and artifacts. For example:

- What other events happened just before or after the known bad event?
- Are there other suspect files in the same folder?
- Are other systems connected to known bad events or known compromised systems somehow?

Later we provide suggested investigative steps based on the GCP service involved, the type of incident, and recommendations for tools that may be useful.

Containment & Remediation

During the containment phase of an incident, some questions that will be important to answer include:

- Can you limit the damage before it gets worse?
- Do you need to isolate virtual machines or services?
- Can you permanently bring the environment back to a safe state?
- If you have identified the root cause, can you fix the original issue? If not, can you mitigate the risk with other preventative technology or additional monitoring to identify future use?
- Have you hunted for other potential compromises? For example, by importing key systems and scanning for malware.
- Have you reviewed the best practices above and confirmed if any need to be implemented?
- Have you enabled additional monitoring where gaps have been identified?
- Have you documented all findings and actions taken?
- Do you need to publish an incident report?
- Have you identified lessons learned and conducted a wrap-up meeting?

Post-Mortem

The postmortem process involves documenting incidents, their impact, resolution actions, the root cause, and subsequent preventive measures. These assessments aim to uncover systemic issues. Key practices when conducting a post mortem include collaborative documentation, real-time feedback, and formal review processes.

To integrate a postmortem culture effectively, organizations are encouraged to gradually introduce it into their workflow, celebrate successful implementations, and involve senior leadership to reinforce its value. Additionally, ongoing feedback mechanisms help refine and improve the postmortem process over time, contributing to a culture of continuous learning and improvement.

A blameless post-mortem can be a complicated thing to conduct as you are at the end of the day identifying the actions that directly lead to the incident, but it is crucial to foster an environment where people feel that they can escalate without fear. The culture surrounding postmortems should foster an environment where incidents are seen as opportunities for system strengthening rather than occasions for fault-finding.

A blame culture can lead to incidents being glossed over or swept under the rug and this can lead to worse outcomes in the future as root cause will not be identified and remediated until the problem is too big to ignore.

Service Specific Advice

Compute Engine

GCP Compute engine enables users to create, configure, and manage virtual machines in the cloud. VMs can be created from pre-configured images or from scratch and can be configured to run a variety of operating systems and applications.

GCP gives you the ability to export disk images of a Compute engine machine, this can be done via the gcloud command line using the compute images export command;

An example to Export a VMDK file named ‘image_1’ from a project called ‘project_1’ to a bucket named ‘bucket_1’:

```
gcloud compute images export --image=image_1 --destination-uri=gs://bucket_1/image_1.vmdk
--export-format=vmdk --project=project_1
```

The Cado platform can import these full disk images for processing and analysis. This includes the process of running threat intelligence against the contents and indexing file contents.

Kubernetes Engine

GCP Kubernetes Engine (GKE) is a managed Kubernetes service that lets you quickly deploy and manage containerized applications in the cloud. GKE reduces the complexity and operational overhead of managing Kubernetes by offloading much of that responsibility to Google Cloud. As a hosted Kubernetes service, GKE is quickly becoming a popular choice for developers and enterprises that want to deploy applications in containers.

Snapshotting a Running VM

To snapshot your persistent disk, first find the disks attached to your VM. Run the following command and look at the source field:

```
gcloud compute instances describe NODE_NAME --zone COMPUTE_ZONE \
    --format="flattened([disks])"
```

Look for the lines that contain disks[NUMBER].source. The output is similar to the following:

```
disks[0].source:
https://www.googleapis.com/compute/v1/projects/PROJECT_NAME/zones/COMPUTE_ZONE/disks/DISK_NAME
```

The disk name is the portion of the source name after the final slash. To complete the snapshot, run the following command:

```
gcloud compute disks snapshot DISK_NAME
```

Redeploying a Container

By redeploying your container, you start a fresh copy of the container and delete the compromised container. Redeploying makes sense when either you already know the cause of the vulnerability or you think it takes an attacker significant effort or time to compromise your container again.
You redeploy a container by deleting the Pod that hosts it. If the Pod is managed by a higher-level Kubernetes construct, deleting the Pod schedules a new Pod. This Pod runs new containers. To do this run the following command:

```
kubectl delete pods POD_NAME --grace-period=10
```

Deleting a Workload

Deleting a workload, such as a Deployment or DaemonSet, causes all of its member Pods to be deleted. All containers inside those Pods stop running. This makes sense to do when you want to stop an attack in progress if you are willing to take the workload offline. This can be the case if stopping the attack immediately is more important than application uptime or forensic analysis. For example to delete a deployment use the following command:

```
kubectl delete deployments DEPLOYMENT
```

Cloud Storage

GCP Cloud storage divides its logs into two categories, usage logs (these track requests made to a specified bucket and are generated hourly) and storage logs (these track the consumption of storage space by a bucket and are generated daily). Both are provided in a CSV file format that can be viewed inside the GCP console or downloaded.

GCP's Cloud Audit Logs also log usage activity with cloud storage raising the question around when should you use usage logs vs. when should you use Cloud audit logs.

Cloud Audit Logs are generally the recommended method of tracking operations performed in Cloud Storage:

- Cloud Audit Logs tracks access on a continuous basis, with delivery of events within seconds of their occurrence.
- Cloud Audit Logs produces logs that are easier to work with.
- Cloud Audit Logs can monitor many of your Google Cloud services, not just Cloud Storage.
- Cloud Audit Logs can, optionally, log detailed request and response information.

However, below are cases where using usage logs instead of or in addition to using Cloud Audit Logs can be beneficial. You most likely want to use usage logs if:

- You want to track access that occurs because a resource has allUsers or allAuthenticatedUsers in its access control settings, such as access to assets in a bucket that you've configured to be a static website.
- You want to track changes made by the Object Lifecycle Management or Autoclass features.
- You intend to use authenticated browser downloads to access objects in the bucket.
- You want your logs to include latency information, the request and response size of individual HTTP requests, or the full URL path and every query parameter.
- You want to track access to only certain buckets in your project and so do not want to enable Data Access audit logs, which tracks access to all buckets in your project.
- Note that usage logs are only generated hourly and can be delayed, particularly when reporting on buckets that experience high request rates.

Downloading logs

Storage logs are generated daily as mentioned above and should be crated before 10am PST. Usage logs are generated hourly and should be available 15 minutes after the end of the hour. They are both available in a CSV format through both the Google Cloud Console and the gcloud storage CLI

For the Google Cloud Console:

1. In the Google Cloud console, go to the Cloud Storage Buckets page.
2. Select the bucket in which your logs are stored.
3. Download or view your logs by clicking on the appropriate log object.

For gcloud storage CLI:

Use the following command:

```
gcloud storage cp
gs://BUCKET_NAME/LOG_OBJECT/DESTINATION
```

BUCKET_NAME is the Bucket where the logs are stored
LOG_OBJECT is the object you wish to download form the logs stored in the bucket
DESTINATION is where you want to deposit the logs

Open Source Tools

There are a number of community created tools that can be useful when responding to incidents in GCP here is a list of some of the most useful ones:

Cloud Forensic Utils - A Python based Cloud forensics toolkit supporting GCP, AWS and Azure useful for extracting evidence from virtual machines

Prowler - An Open source tool kit with support for GCP, AWS and Azure, can be used for best practices assessments, audits, incident response, continuous monitoring, hardening and forensics readiness.

SANS SIFT toolkit - Collection of free and open-source incident response and forensic tools designed to perform detailed digital forensic examinations in a variety of settings.

Security monkey - Monitors Cloud accounts for policy changes, works with both GCP and AWS

Cartography - Creates maps of infrastructure and the connections in between supports multiple platforms including, GCP, AWS and Azure

Automating Incident Response

In the cloud, things are fundamentally different than in an on-premises world. The scale is exponentially larger in terms of numbers of workloads. Applications are much more dynamic and elastic, spinning up and down continuously.

As such, the old, manual way of responding to threats simply doesn't work. By the time you've had a chance to kick off the necessary investigation and response steps from your legacy playbook, the attacker – and even the workload itself – is long gone and damage is already significant.

What's more, the risk equation is different. Traditionally, performing any sort of automated containment was a big deal, since it could mean interrupting critical business operations, especially if things don't go according to plan. However, in a resilient cloud-native environment, with thousands of load-balanced instances, the potential impact on the business when taking targeted, automated response measures, is much less.

Cloud security responders can really benefit by finding ways to ruthlessly automate manual steps in the incident response process. With automation, analysts save days and in some cases, even weeks during an investigation.

Cado in Google Cloud

You can deploy the Cado Platform into your Google Cloud environment to immediately gain the ability to collect and investigate resources from within Google Cloud. Note that Cado is now available in the Google Cloud Marketplace as well.

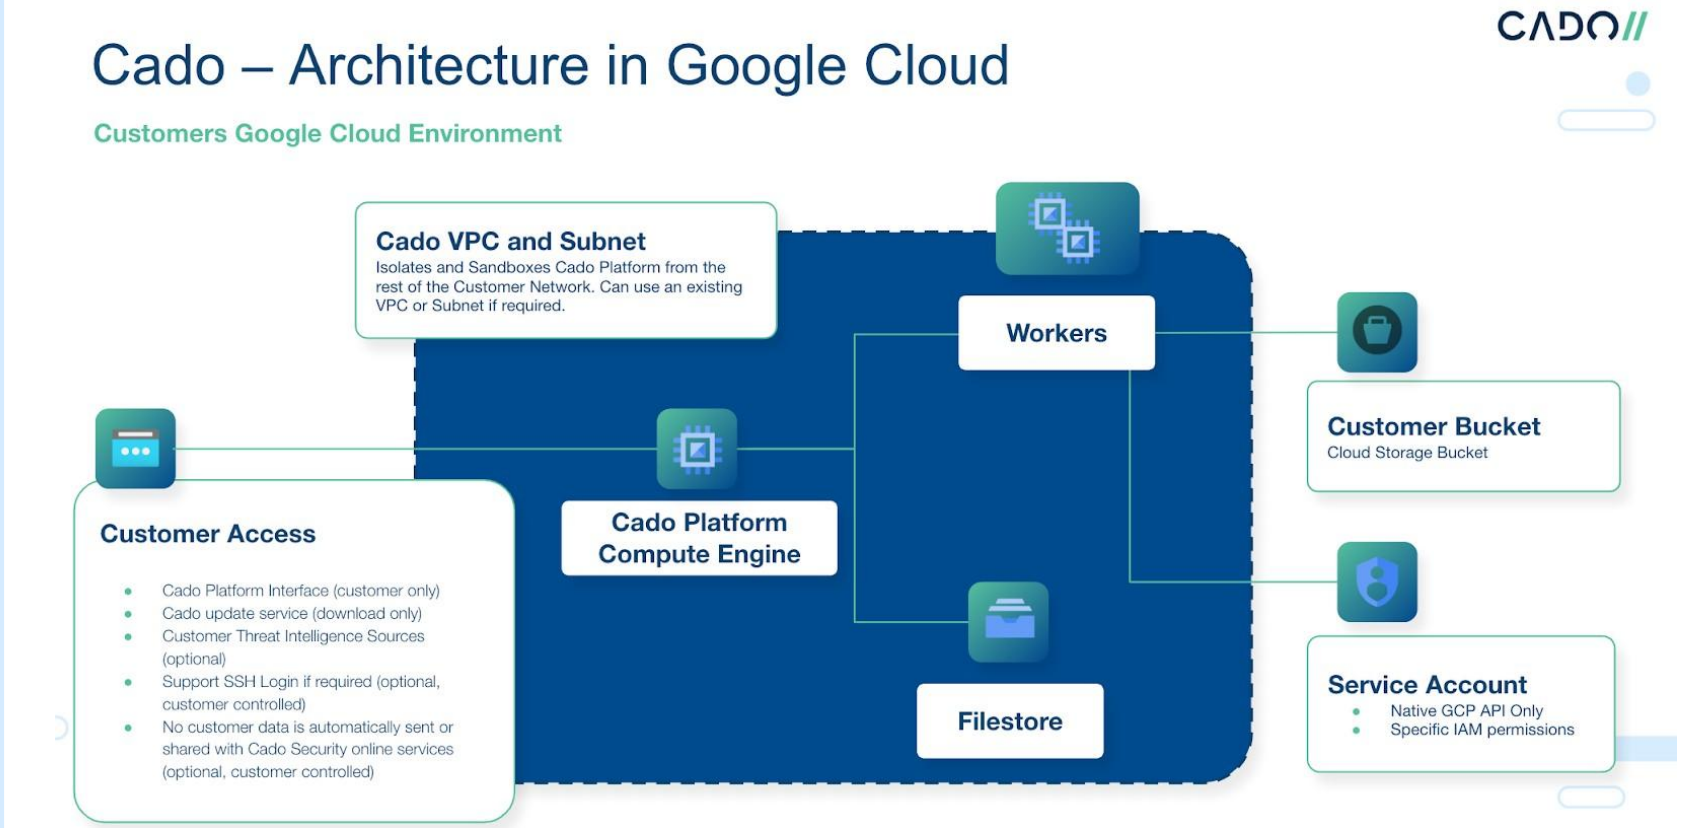

About Cado

Cado Security leads the charge in revolutionizing digital forensics away from conventional approaches, Cado pioneers a tailored platform explicitly crafted to navigate the complexities of cloud-based operations. As businesses rapidly transition to cloud-centric infrastructures, traditional tools struggle to keep pace, leaving security teams grappling with limited insights and sluggish response capabilities. The Cado redefines the norm by seamlessly automating the capture and analysis of forensic data across diverse cloud, container, and serverless environments. Cado empowers security teams to swiftly identify cloud risk, drastically reducing response times. Backed by esteemed investors including Eurazeo, Blossom Capital, and Ten Eleven Ventures, Cado maintains a formidable presence, serving as a beacon for cutting-edge security solutions both in the United States and the United Kingdom.