



# The Tactical Guide to Securing Data on Cloud Platforms in 2021

How to Improve Security and Visibility of  
Amazon Web Services, Microsoft Azure,  
and Google Cloud Platform





# How can security teams maximize security and increase visibility on Amazon Web Services, Microsoft Azure, and Google Cloud Platform?

## ▲ Introduction

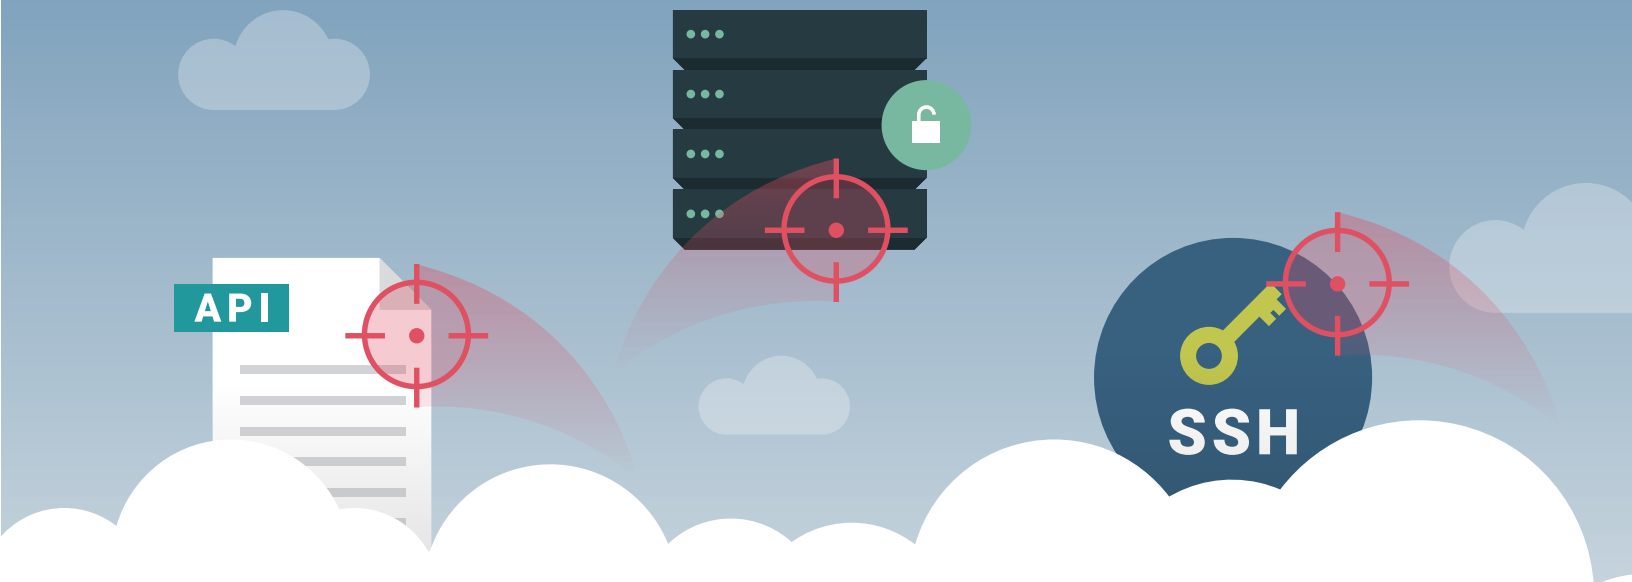
Enterprises have leveraged cloud platforms such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) to facilitate web applications for years, and the platforms have proven effective and reliable. That's why many enterprises are taking advantage of the scalability provided by these platforms to assist with traditional network needs such as directory services, user document storage, and internal operational applications.

While this approach can deliver accessibility and storage benefits, many organizations that move critical infrastructure to these cloud platforms mistakenly assume that the cloud provider also delivers sufficient visibility and monitoring of the cloud environment. They therefore fail to configure critical controls and secure architecture practices, leaving themselves vulnerable to attacks.

By properly determining where your most sensitive data lives, configuring the controls provided by popular cloud platforms, and monitoring performance and security information and event management (SIEM) integration, enterprise organizations can ensure consistent and convenient visibility and security of data stored on AWS, Microsoft Azure, and GCP.

## ▲ Common cloud platform risks

To begin securing cloud-based data, enterprise architecture and security teams must first understand that configuring cloud-based security options is not straightforward. Many of the risks facing a standard on-premise architecture – such as improper segmentation, overly-permissive firewall rules, or weak passwords for remote servers or VPNs – still exist in a cloud environment.



Systems must also log and monitor new risks, including exposure of application program interface (API) keys in source repositories or open web directories, overly permissive data storage buckets, and Secure Shell (SSH) keys available openly in source repositories. Attackers successfully use all these methods to breach systems.

Cloud environments are also not exempt from any compliance regulations required for on-premise environments. Business leaders often assume that cloud platforms provide these key controls, and only realize that's not happening when a violation occurs.

Default configurations for AWS, Azure, and GCP may not turn on event logging, encryption, data retention, multifactor authentication, or other preventative controls that ensure compliance. **They also don't automatically correlate information to SIEM and logging systems and will need to be set up properly to comply with General Data Protection Regulation (GDPR) rules and other compliance mandates.**

## ▲ In this paper, you'll learn:

**HOW** to configure and monitor your cloud platforms for improved visibility

**WHICH** cloud-native tools are needed to secure AWS, Microsoft Azure, and GCP

**WHY** integration is key to securing a multi-cloud environment



To begin securing cloud-based data, enterprise architecture and security teams must first understand that configuring these cloud-based security options is not straightforward. Many of the risks facing a standard on-premise architecture – such as improper segmentation, overly-permissive firewall rules, or weak passwords for remote servers or VPNs — still exist in a cloud environment.

## ▲ Step 1: Determine where your most sensitive data lives and prioritize integrations that increase visibility



By increasing visibility into cloud platforms, organizations will not only lower organizational risk, but also help cut costs, ease architecture design, and help keep a handle on sprawling deployments. Visibility also helps development teams with proper data organization and segmentation.

Increasing visibility begins with identifying the data worthy of protection and defining where it lives. The ease of cloud deployments means that sensitive data could live or be transferred anywhere, so properly understanding the location and usage of data is key. Without proper controls, it's much easier to transfer customer data from a private server to a public storage repository.

Visibility is essential because security teams must know how data moves throughout their environments in order to protect it. Usually, dataflows are best traced from the point at which an application is accessed back to where the company's developers eventually access the systems on which data is stored. Without a proper understanding of how data moves through the environment, security teams will waste resources securing potentially lower-priority infrastructure devices.

Two key benefits of cloud providers are out-of-the-box security tools and open APIs that make log ingestion easy. Once you know what you're looking for, proper tuning and integrations can make alerting and visibility into attacks as simple as an on-premise infrastructure. When you know where your sensitive data lives, you can properly configure security architecture up front to save months of headaches trying to secure data that has leaked into less-secure areas of the environment.



**Increasing visibility into cloud environments begins with knowing what the data worthy of protection is – and defining where this data lives.**



Amazon's combination of GuardDuty, CloudTrail, and CloudWatch offers a robust "out-of-the-box" security experience. CloudTrail and CloudWatch offer a consolidated approach to security of any account, granting a wide range of visibility through only needing two tool integrations.

- **GuardDuty**

GuardDuty's out-of-the-box security features provides comprehensive alerting. Think of this as a base-level defense for an AWS account.

- **Systems Manager**

Proper tagging and management of sprawling Elastic Compute Cloud (EC2) instances can be accomplished via this native AWS tool.

- **CloudTrail and CloudWatch**

CloudTrail and CloudWatch log everything that happens in the environment. While these logs are useless without proper alerting, they are indispensable when flowing through a SIEM. The amount of content in these logs can be massive, so appropriate filtering is recommended.

- **AWS Config**

Config offers an easy standardization of what deployments should look like.

- **Macie**

If you store sensitive information in the cloud environment, Macie helps look for sensitive data across S3 storage and enforces proper security controls.



Azure has gone through major revisions in how it handles logging and visibility, but recent updates have greatly improved the platform's visibility. Because Azure is generally used for more Windows-centric workloads, there are a few extra options for gaining visibility.

- **Security Center**

The Azure Security Center offers the main hub for security management in Azure. Security Center consolidates numerous controls into a security score that offers a useful starting place for remediations. Security Center can integrate with the Azure

Defender Line, including Azure Defender for Servers, Azure Defender for IoT, 365 Defender, and Azure Defender for SQL.

- **Azure Active Directory**

All user access flows through Active Directory; with proper monitoring, suspicious activities and changes can bring visibility into security team projects.

- **Azure Key Vault**

The Key Vault is the center for managing encryption in the Azure environment.

- **Azure Event Hub and Activity Logs**

This is the Azure parallel to AWS CloudTrail: activity log pipelines that offer critical functionality when ingested by a SIEM.



GCP has the major benefit of integrating well with G Suite, providing robust logging and integrations across an organization's infrastructure if G Suite is being utilized. While GCP integrates well with native Google tools, its log flow can be less intuitive than AWS and Azure.

- **Security Command Center**

Setting up Security Command Center tools gives greater visibility into each layer of the GCP environment. The Security Command Center offers many of the same features that are spread across multiple tools in other platforms.

- **Cloud Audit Logs**

Proper tuning and ingestion into a SIEM offers visibility into what's happening in your cloud environment.

- **Chronicle Detect**

Threat detection tool that includes an advanced rules engine built on Google Infrastructure.

**TIP:**

Leverage these additional tools for monitoring microservices in the cloud(s)

**• Datadog**

Datadog is a nice enhancement for performance monitoring, offering visibility across cloud platforms, Kubernetes clusters, and microservice performance. Datadog can serve as a single pane of glass to correlate logs across hybrid and multi-cloud infrastructures against specific performance metrics within your application. Datadog works well because it offers over 400 integrations across all the infrastructure that makes microservices possible.

**• Prisma Cloud**

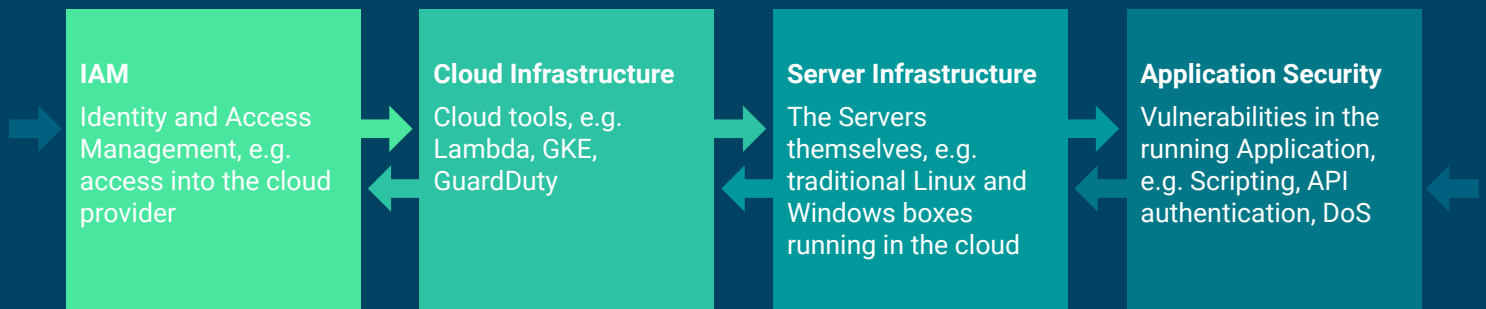
Palo Alto Network's microservice architecture vulnerability scanner, Prisma Cloud, can ingest CloudTrail logs at a lower cost than a SIEM and offers rule sets for vulnerabilities and compliance. At the microservice level, Prisma Cloud can be integrated at all stages of the pipeline to enforce image compliance, runtime protection, and vulnerabilities.

## ▲ Step 2: Configure cloud platforms to maximize the security of their architecture

Most cloud platforms include some built-in security features and applications to ensure proper data protection. Knowing which features the platform already provides is critical for determining the visibility and automation you can expect from your provider, and what internal security teams must perform.

Once the basic tools are enabled, the process of tuning and tightening controls can begin. The raw audit logs, in combination with identity and access controls, offer a powerful combination for security visibility. Set up alerts for unusual calls from accounts, repeated denials, policy changes, and several other pieces of content that easily identify attacker activity.

Look at the following as areas of vulnerability that an attacker would want to pivot through:



Vulnerabilities exist at each of these layers, with the most vulnerabilities existing at either extremity, as both are public-facing. Application security and Identity Access Management (IAM) should be the primary areas of concern.

IAM can be secured through the combination of access logs and regular auditing and tightening of permissions. AWS recently released a helpful tool called IAM Access Analyzer that attempts to automate some of this process.

Another way to ease the audit burden is to have users assume IAM roles based on their job roles through a single sign-on (SSO) provider. These centralized IAM roles can be tightened to different levels of access, with different levels of alerting on each. For example, consider segmenting the security team so that tier 3 access roles contain only a few authorized users that report on all activity, only to be used in emergencies. Regular duties are to be done with tier 1 roles.

Application security requires security by design, i.e., developers who care about the security of the application. There must be buy-in from developers on the importance of security; examples like encrypted communications are an initial challenge that development teams must prioritize. Once buy-in is established, several tools can assist in providing visibility into application security. Some key tools are static code scanners during the build pipeline, dynamic scanning of the running code, API security validation, web application firewall (WAF) controls, and library vulnerability scanning.



**Most cloud platforms include built-in security features and applications to ensure proper data protection. Knowing which features the platform already provides is critical for determining the visibility and automation you can expect from your provider, and what internal security teams must perform.**





AWS seems to have one of the more flexible automation and DevOps capabilities with its comprehensive API and scripting abilities, and includes many powerful networking options to secure the environment. Some common techniques used for securing the AWS environment include:

- **VIRTUAL PRIVATE CLOUDS (VPC)**

Provides network isolation that can take advantage of subnet segmentation with public and private subnets.

- **CLOUDTRAIL LOGGING**

Enables visibility into potentially unauthorized changes to the platform.

- **SECURITY GROUPS**

Locks down individual EC2 instances or categories of instances; functions much like a virtual firewall.

- **AWS RDS (RELATIONAL DATABASE SERVICES)**

Provides at-rest encryption for data.

- **IAM Analyzer**

Helps to identify external sharing and permissions boundaries.



Azure provides many built-in security mechanisms and currently has more application-specific protections than AWS or GCP. Here are the common items used to lock down the Azure environment:

- **OPERATIONS MANAGEMENT SUITE (OMS)**

Performs event viewer searches, runbook automation using PowerShell (similar to AWS Lambda functionality) and improves visibility into the environment with built-in dashboards.

- **AZURE APPLICATION GATEWAY**

Offers web application firewall (WAF) and load-balancing capabilities.

- **AZURE DEFENDER FOR SQL**

Monitors potential vulnerabilities, access, or injection attempts to the databases in Azure.

- **STORAGE ENCRYPTION**

Provides service to automatically encrypt when writing data; also can encrypt VM disk via BitLocker and perform SQL transparent data encryption (TDE).

- **AZURE VIRTUAL NETWORK (VNET)**

Similar to the AWS VPC, helps to isolate networks away from each other.

- **NETWORK SECURITY GROUP (NSG)**

Offers built-in firewall capabilities within VNets.

- **SECURITY CENTER**

Provides policy management, threat detection, and alerting and dashboard visibility.

- **AZURE ACTIVE DIRECTORY IDENTITY PROTECTION**

Provides features such as light user behavioral analytics (UBA) and SSO capabilities. An extremely powerful feature is the ability to extend your local Active Directory configuration to Azure, greatly reducing duplication of efforts around user management and permissions.



Google's cloud environment also provides a VPC environment like the AWS platform but with less overhead infrastructure than traditional VPCs. GCP also allows for the traditional segmentation, firewall, and access control features provided by most cloud platforms.

One of the major security-specific features is the Security Command Center, which provides capabilities for monitoring and securing the GCP platform, including:

- **ASSET DISCOVERY AND INVENTORY SCANNING**

Scans the App Engine, Compute Cloud, and Cloud Storage/Datastore platforms.

- **CLOUD DLP**

Discovers potentially sensitive data contained in data storage buckets.

- **APPLICATION VULNERABILITY SCANNING**

Uncovers potential injection and scripting vulnerabilities in hosted applications.

- **ACCESS CONTROL MONITORING**

Ensures the correct access control policies are in place.

- **ANOMALY DETECTION**

Identifies potential threats from botnets, malware communications, or other suspicious traffic in or out of the GCP environment.

- **INTEGRATION WITH TOP SECURITY TOOLS INCLUDING CLOUDFLARE, CROWDSTRIKE, PALO ALTO NETWORKS, AND QUALYS**

Detects potential distributed denial-of-service (DDoS) attacks, infected endpoints, policy violations, network attacks, and per-instance vulnerabilities.

- **INTEGRATION WITH THE OPERATIONS LOGGING PLATFORM AND CUSTOM APIS**

Collect metrics, logs, and traces across Google Cloud and your applications.

## ▲ Step 3: Monitor the cloud through integration

Most cloud environments generate an incredible number of logs that, without proper parsing, can quickly overrun security teams. Having the proper team, detections, and automations in place to constantly parse, tune, and respond to these alerts is the backbone of organizational security.

Beyond basic SIEM integrations, the cloud offers huge numbers of integrations through APIs. Businesses are faced with more and more tools that don't have native SIEM integrations, requiring security teams to pivot between multiple tools to triage a security alert. In order to efficiently and effectively respond to alerts, you must consolidate data into one place for monitoring.



To see what's happening in the cloud environment and export logs into an alerting engine such as a SIEM, enterprise organizations must plan a cloud-logging strategy, which begins by answering the following questions:

## **1** **WHERE WILL THE LOG AGGREGATION AND FILTERING TOOLS RESIDE?**

This often depends on the location of the data, but it's more efficient to have the collection tools doing the filtering and aggregating as close to the source as possible. Compliance may dictate where the raw log data must live. France, Germany, and many other countries have strict laws against exporting potentially sensitive personally identifiable information (PII) as well as the EU GDPR regulations, and upcoming U.S. versions in Nevada, New York, and California (California Consumer Protection Act).

## **2** **HOW BIG ARE THE INTERNET CONNECTIONS BETWEEN YOUR CLOUD ENVIRONMENT AND LOCAL DATA CENTERS?**

It may be more cost-effective to keep the raw data at both places and send the actionable events used in your alerting rules to the central SIEM or monitoring tool. Sometimes there may be a logical split in your organization to make this more feasible – for example, separate, distinct corporate divisions or business units.

## **3** **HOW WILL YOU COLLECT AND PARSE CLOUD INFRASTRUCTURE LOGS?**




In addition to the standard operating system or application logs from the servers, many essential “behind the scenes” cloud infrastructure logs should also be gathered and monitored for unauthorized/malicious activities.

**After answering those questions, consider the common methods of log collection and SIEM integration used by your cloud provider, detailed on the next page.**



Amazon includes a consolidated way to pull all AWS logs via CloudTrail, which logs all API calls to the environment from the console API or command-line interface (CLI). It uses a RESTful API to integrate into the other AWS applications and can aggregate logs from multiple instances and regions.

Some of the more actionable CloudTrail logs include starting or stopping of an instance and creation or deletion of users within AWS and logins to the platform. The CloudTrail log is formatted differently for many AWS resources but each event mostly contains the following information:

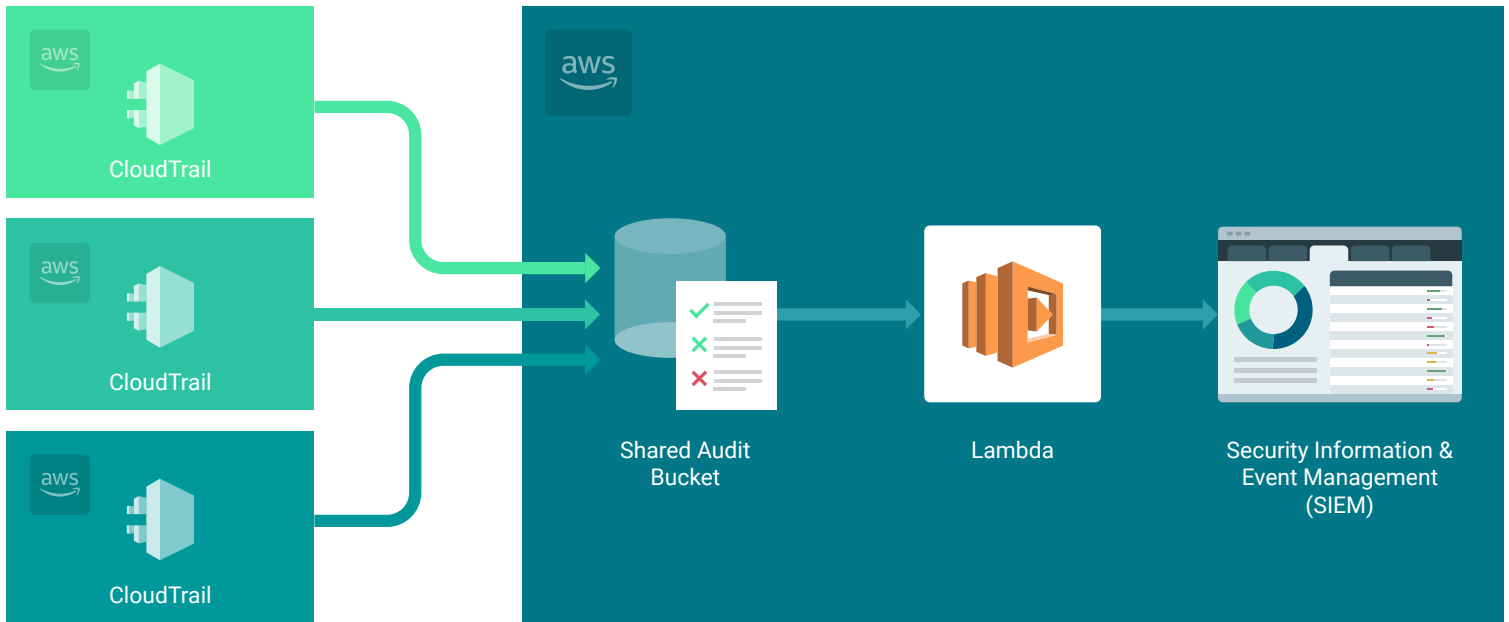
 <p>Time request occurred along with event ID and type</p>	 <p>User/service and source IP/AWS region that made the request</p>	 <p>Request parameters, error codes, and messages</p>
---	--	--

The CloudWatch application is often used in conjunction with CloudTrail to alert on real-time events generated by the CloudTrail logs. This is useful if certain types of log alerts need not stem from the SIEM, which may include items such as health or performance monitoring that go directly to the infrastructure team. CloudWatch has a mobile console or agent for viewing logs, the ability to create alarms for system health warnings, dashboards and reports with API integration for extended visibility, and options for effective monitoring of billing and auto-scaling.

The AWS built-in log viewing tools are good, but analysts typically don't want to log into a separate portal to import those logs into the SIEM for additional correlation. A simple example of exporting these logs would be to enable all the EC2 host instances logging to a S3 bucket, creating an AWS Lambda function (typically Python, but other languages are supported) to read from the S3 bucket and push the logs to an AWS syslog server, and finally forwarding the syslog messages to the SIEM log receiver or collector.

Leveraging a syslog architecture offers the flexibility of adding multiple destinations for sending log data, simplifies configuration, and permits most devices to send and receive syslog-formatted messages.

**Below is an illustration of the CloudTrail event flow into a SIEM:**



Source:

[https://medium.com/@marcusrosen\\_98470/real-time-log-streaming-with-cloudtrail-and-cloudwatch-logs-3389c4cc5ef4](https://medium.com/@marcusrosen_98470/real-time-log-streaming-with-cloudtrail-and-cloudwatch-logs-3389c4cc5ef4)

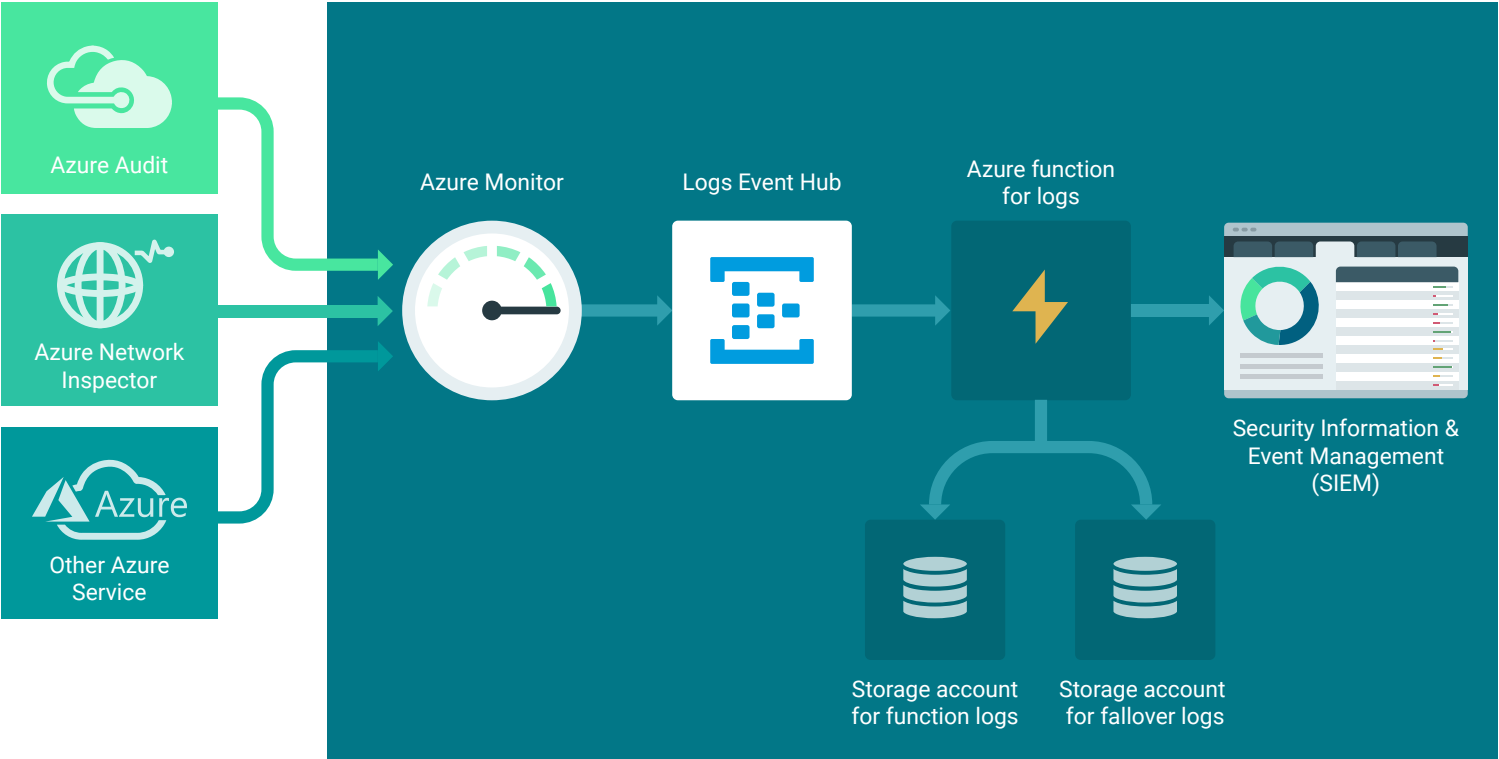


Azure has a few different ways to log events happening within the architecture. The easiest method uses the Azure Event Hub functionality to configure Azure applications to send all logs to Event Hubs – functioning like a typical event broker, such as Kafka or Redis, and providing queuing and streaming capabilities.

Applications that provide the most essential logs are from the Azure Resource Manager, which logs all the administrative activities and Security Center logs that collect from agents on all the Azure hosts and the built-in host-based Intrusion Detection Systems (IDS). Other useful applications include the Network Security Group Flows – which logs connections to and from virtual machines – and Azure Active Directory, which logs normal administrative Active Directory activities such as authentication success or failure, users and groups created or deleted, and account lockup information.

Azure’s Event Hubs service is the central location to process events. Event Hub is in many ways a hosted Kafka service that makes it easy to integrate with the SIEM tool of choice. The key logs to capture into the event hub are activity and diagnostic logs.

**Below is a sample event flow for sending logs to the SIEM, as described by Microsoft:**



Source: [https://medium.com/@marcusrosen\\_98470/real-time-log-streaming-with-cloudtrail-and-cloudwatch-logs-3389c4cc5ef4](https://medium.com/@marcusrosen_98470/real-time-log-streaming-with-cloudtrail-and-cloudwatch-logs-3389c4cc5ef4)

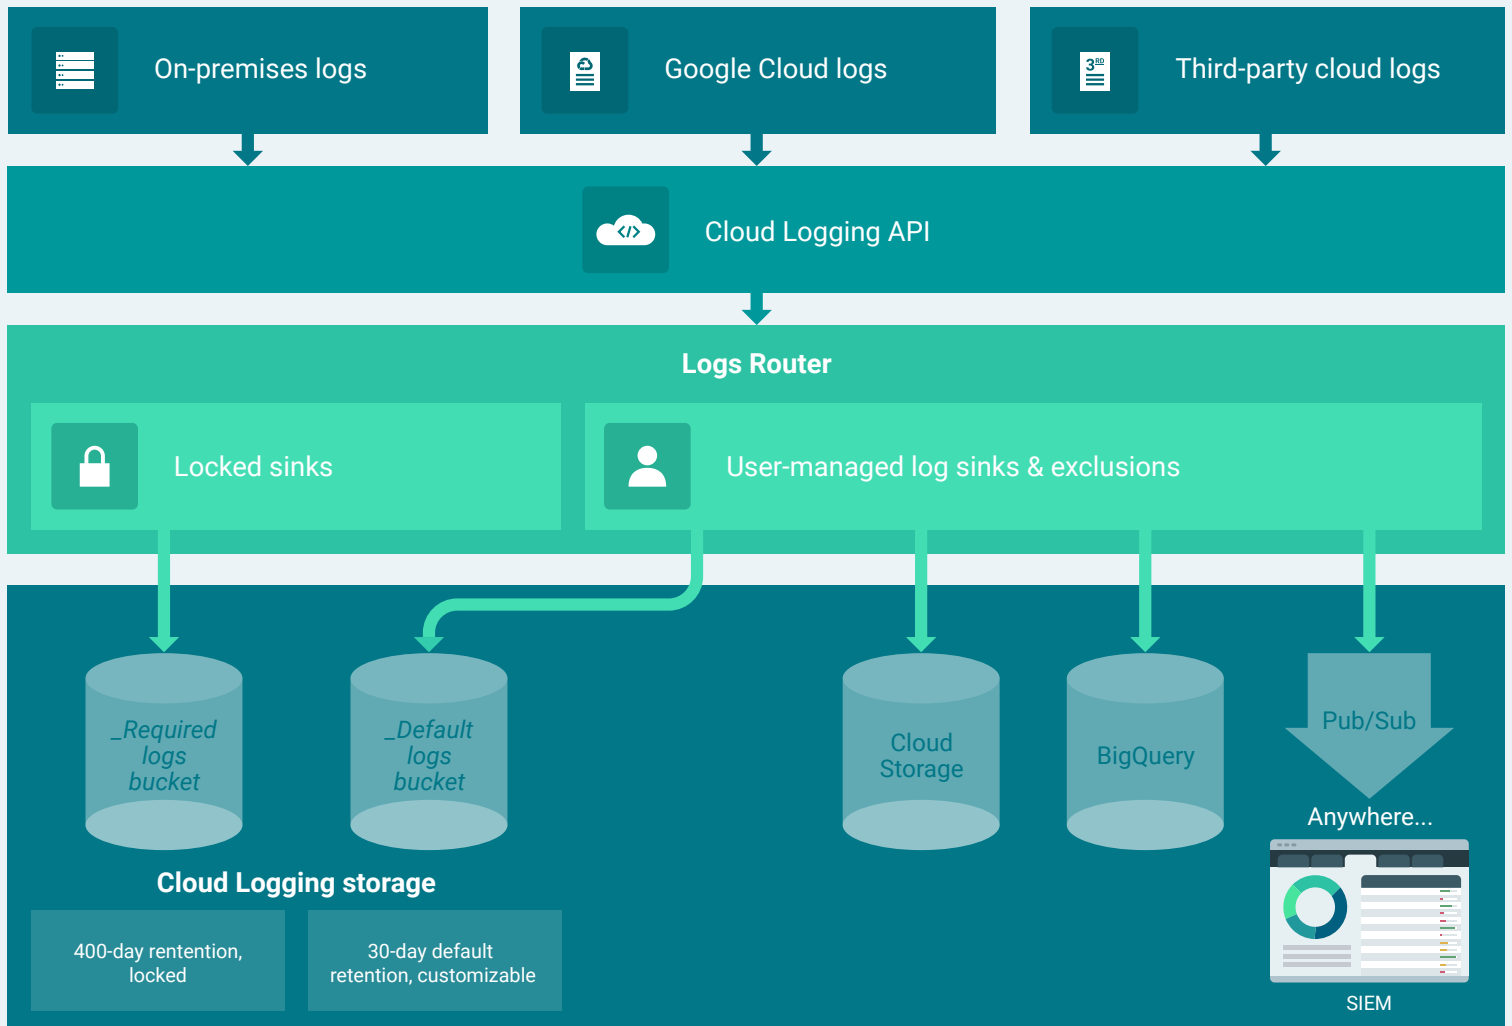
## Google Cloud

Google provides a convenient logging infrastructure with its operations platform (formally Stackdriver). Operations can feed into a publisher/subscriber (pub/sub) model that is easily integrated with a SIEM or open source logging tool. GCP Cloud Logging includes the ability to consolidate GCP and AWS logs if both services are being used, as well as BindPlane to capture logs from common app components.

Start by creating a project within GCP that contains the various Google resources that will be generating logs. Once the project is set up, create a pub/sub that contains all logs for the project, or isolate the logs for each resource by creating a pub/sub for each.

Different types of logs can be placed in different sinks to split between different types of storage and analytics tools. For example, you may want some logs to flow into BigQuery for long-term data analytics, but regular cloud storage for shorter terms.

Below is a sample logging pipeline to consume logs, as defined by Google:



Source:

[https://medium.com/@marcusrosen\\_98470/real-time-log-streaming-with-cloudtrail-and-cloudwatch-logs-3389c4cc5ef4](https://medium.com/@marcusrosen_98470/real-time-log-streaming-with-cloudtrail-and-cloudwatch-logs-3389c4cc5ef4)



Businesses are faced with more and more tools that don't have native SIEM integrations, requiring security teams to pivot between multiple tools to triage a security alert. To efficiently and effectively respond to alerts, you must consolidate data into one place for monitoring, detection, and response.




## INTEGRATION EXAMPLE: SECURITY AND VISIBILITY IN A MULTI-CLOUD ENVIRONMENT

At the end of the day, certain applications or frameworks do a better job at surfacing critical visibility for your teams. Operations teams may prefer the health monitoring of GCP Operations or AWS CloudTrail, security teams the monitoring features of Prisma Cloud, and developers the log searching offered by Splunk. Cloud platforms offer unparalleled log consolidation and ingestion abilities across all these different platforms, so it's essential to know where this data ultimately lives.

A use case for ingesting all this data might be using the following tool categories to gain visibility into different parts of your ecosystem:

- Application health data, EC2 instance monitoring, EKS health monitoring, Elasticsearch monitoring
- WAF logs, Kubernetes activity logs, alerts being forwarded from the EDR, CloudTrail logs for custom content
- Application troubleshooting logs, ad hoc data
- Endpoint data from the individual nodes or instances
- CloudTrail and CloudWatch logs for configuration management



Each of these tools offer different benefits and can provide visibility into what's happening in your cloud environment. A security professional can never know for certain where the first sign of an attack might come from. For example, the common attacker tool Mimikatz might one day be detected on a cloud worker node by your EDR. The EDR can provide some context into how it was installed and where it was pointed next; however, it won't tell you about anomalous API activity from a service account, attempted kubectl attacks, and how port conflicts are breaking a pod in the production environment.

In this scenario, an attacker compromised a vendor's service account, used that account to gain access to a Kubernetes pod, and is attempting to recon where to pivot next. Whereas an EDR might lead your security team to remove Mimikatz and reset local credentials on the pod, they would miss the larger picture of what happened. Only with full visibility across on-premise and cloud environments can security teams recognize modern attacks like these that can easily span the entire infrastructure.



**TIP:**

Expand security with Cloud Application Security Brokers (CASBs)

While not an all-encompassing enterprise platform, CASB solutions provide granular auditing capabilities at the application or site level, and require the same logging and monitoring considerations as the full-fledged cloud platforms. Most CASB solutions have similar capabilities or deployment considerations via proxy or API and are essential for incident response and forensic investigations since monitoring and alerting on unauthorized access to cloud services is often critical to detecting potential insider threats. Compromising Salesforce, Office 365, or Dropbox could be devastating to an organization, since many of the corporation's sensitive documents are stored in one or more of those applications. Data classification and/or DLP needs to also extend to these cloud services for consistent coverage and management.

While providing an additional layer of security is the primary function of CASBs, comprehensive IT security providers offer advanced services that turn security activities into valuable business insights.



**Only with full visibility across on-premise and multi-cloud environments can security teams detect, investigate, and mitigate modern attacks that can easily spread across the entire infrastructure.**

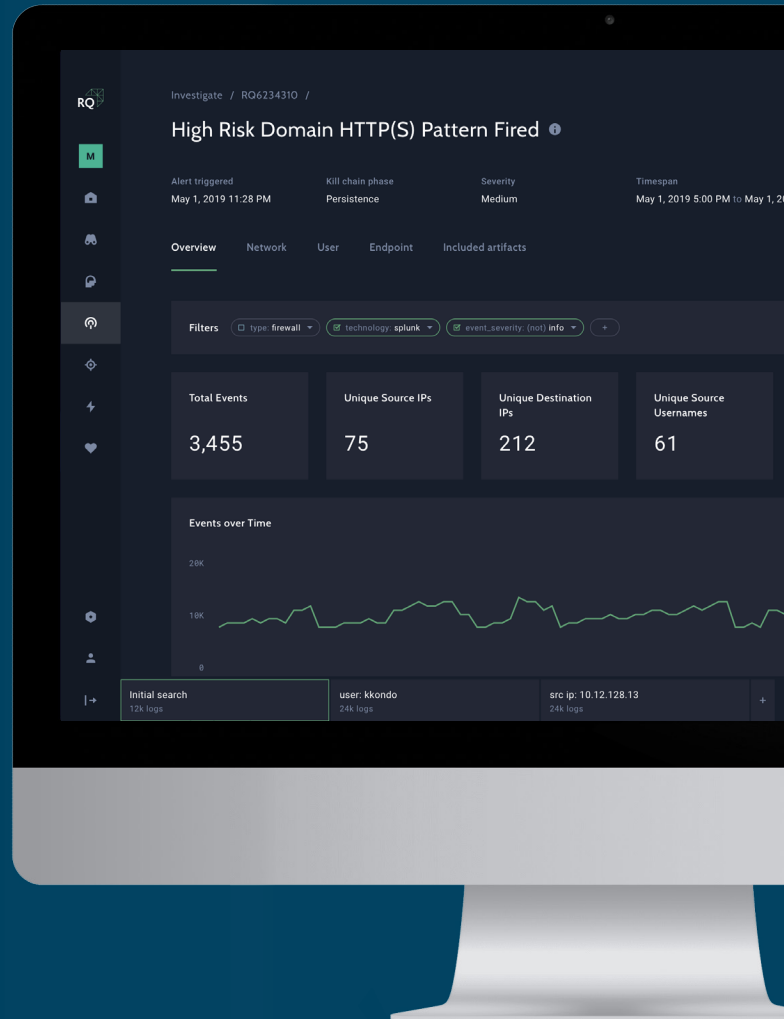
## How ReliaQuest Unifies and Secures Data in the Cloud

ReliaQuest, a global leader in cybersecurity, delivers industry-leading visibility and automation across complex environments with a platform purpose-built to protect organizations from security breaches. GreyMatter is the first cloud-native solution that unifies an enterprise's on premise and multi-cloud technologies, delivering a control plane for visibility, detection and response. By increasing visibility through the platform's proprietary universal translator and use of automation and artificial intelligence, GreyMatter saves security teams valuable time and increases effectiveness by enabling automatic and continuous threat detection, threat hunting, and remediation.

As organizations migrate their critical infrastructure and applications to cloud platforms, they are faced with alternatives to gain visibility from traditional methods that often leads to disparate pools of data. Without a centralized view or ability to recognize gaps, organizations are left vulnerable to attacks.

ReliaQuest GreyMatter provides holistic visibility to recognize threats across all of your tools and attack surfaces, including cloud, without the impracticality of creating a single data lake to perform analysis. Through the platform's universal translator, on-demand aggregation and normalization of targeted, fragmented data provided in a single, cohesive view, ReliaQuest GreyMatter enables faster, more comprehensive threat investigations as well as retrospective IOC searching and long-term behavioral analysis threat hunting.

“**ReliaQuest GreyMatter reduces complexity of data integration across disparate tools – from on-premise to cloud – so you can perform more efficient investigations and threat hunts across your technology stack.**”



**LEARN MORE ABOUT  
RELIAQUEST GREYMATTER**

**RELIAQUEST**

Make Security Possible™

(800) 925-2159

[www.reliaquest.com](http://www.reliaquest.com)

[info@reliaquest.com](mailto:info@reliaquest.com)

Copyright © 2018 ReliaQuest, LLC. All Rights Reserved. ReliaQuest, RQ, and the ReliaQuest logo are trademarks or registered trademarks of ReliaQuest, LLC or its affiliates. All other products names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies. All other information presented here is subject to change and intended for general information. Printed in the USA.