SANS

A SANS Whitepaper

JumpStart Guide for SIEM in AWS

Written by **J. Michael Butler**

August 2019

Sponsored by:

AWS Marketplace
in conjunction with
Optiv

Introduction

Gone are the days of focused technicians in a darkened lab with a table full of terminals located somewhere deep below the data center. Thankfully, simple logging and manual reviews by a roomful of techs have morphed into more automated processes. With SIEM systems, logs are now normalized and collected in a central location for analysis. As SIEMs have matured, more automatic alerting, and even reactions to events, have moved us into the security orchestration and automated response (SOAR) world—or as it's also known in some circles, SIEM on steroids. Currently, according to Gartner, "Analytics are a core capability of all SIEM solutions."[1] Analytics and response are what SOAR is all about.

> *[SIEM] provides the ability to gather security data from information system components and present that data as actionable information via a single interface.*
> —National Institute of Standards and Technology

At its most basic level, the SIEM is defined by NIST as an "[a]pplication that provides the ability to gather security data from information system components and present that data as actionable information via a single interface."[2] Adding SOAR integrates additional data feeds, correlation, analysis and automated functions based on identified incidents, indicators, events and threats.

[1] "Critical Capabilities for Security Information and Event Management," www.gartner.com/doc/reprints?id=1-5VGLBIM&ct=181129&st=sb

[2] Computer Security Resource Center Glossary, https://csrc.nist.gov/glossary/term/Security-Information-and-Event-Management-Tool

In addition to SIEM log collection, some added data feeds for a SOAR system would likely include endpoint management system alerts, threat and vulnerability data from third parties (for example, STIX/TAXII feeds), and help desk and collected forensics data, all to be correlated with the SIEM data. Once that data is analyzed, remediation or other actions can automatically take place for those issues identified by the organization as reliably founded and actionable. The questionable issues can be referred to the SOC (security operations center) for further analysis as needed.

In this paper, we discuss needs, implementation options, capabilities, and various considerations for organizations seeking to implement SIEM/SOAR capabilities in Amazon Web Services (AWS). We discuss the integration of SIEM and SOAR in the cloud environment and how that compares to on-premises use. What does a cloud use case look like? What are the differences between cloud and on-premises deployments? Then we offer suggestions for planning integration of SIEM and SOAR into an AWS cloud environment in the way that is most beneficial to an organization. We hope to help organizations evaluate the options and make the best choice.

Understanding Your Needs

First, consider what technology your organization needs to adequately collect, analyze and react to SIEM data. If your organization can already determine the actionable events or incidents in the existing environment with current tools, the temptation may be to try to adapt those tools to the cloud or vice versa. In that case, be sure to review the security offerings available in the cloud that can improve on what the on-premises solutions offer. New features offering enhancements or alternatives for an on-premises system are being added regularly to the cloud.

After a careful determination of your organization's feature and function requirements, present those requirements to your vendors and start the discussions about what you need to make it all work. Look at the new technologies that may be needed.

Be certain to review existing gaps and what it would take to eliminate them. Be wary of the "gotchas" that will require (possibly significant) resource investments, such as additional subscription fees, personnel and training, and ongoing costs such as annual software maintenance fees. Also consider growth to scale and requirements to enable that growth, and, conversely, the ability to shrink to scale. Cloud environments make it easier to scale up and shrink down resources in response to users' needs. This is especially useful for organizations that experience seasonal change.

The organization should have a long-range plan to budget for implementation, ongoing operations, and hardware and software maintenance. No one needs one more software package to sit on the shelf without providing value. As the SIEM/SOAR project moves forward, revisit requirements regularly to make sure the organization's incident response needs are being met. Figure 1 illustrates the process.

Figure 1. Process for Understanding Your Needs

In the SANS 2019 Cloud Security Survey, 75% of the respondents reported using as many as 10 cloud providers for all operations, and 3% of the respondents said they use more than 100 providers![3] If your organization has multiple cloud providers, consider the need for SIEM/SOAR tools to be capable of accumulating and analyzing data from all of the cloud environments in use. This functionality is particularly needed if the organization has communication or network channels set up between multiple environments, causing incidents in one environment to have an undesirable impact on another.

Implementation Options in AWS

If your organization is thinking of leveraging current on-premises technologies for SIEM and SOAR as you move to the AWS cloud, be sure to take note of the new cloud-native solutions that were not previously available. As of this writing, AWS Security Hub, which provides compliance data, security alerts and security findings, is now generally available. Many desirable SIEM features are now native options in the AWS cloud. It is also important to note that third-party providers, including AWS partners Splunk and Sumologic, have already integrated with AWS Security Hub.

Cloud-Optimized

Consider the cost delta between using the cloud solutions versus the on-premises tools, as well as the costs for the significant storage requirements of SIEM/SOAR data in the cloud versus on premises. Also look at the license fees to be paid for the solution your organization needs versus any on-demand licensing available through AWS for access to its solution partners. At the very least, the cloud-native options can enhance other tools the organization uses, whether the SIEM data is stored in the cloud or on premises.

One advantage of working with off-premises options is the clearer pricing models when compared to running everything on premises. Many cost factors in the data center have to be included if the organization is to get a true picture of the total cost of ownership (TCO). For example, how much is being paid for CPU cycles, mass storage, power requirements, HVAC requirements, facility space, hardware, software, licensing, maintenance, upkeep, personnel and other hidden costs in on-premises environments? On the other hand, the pricing models will be much clearer from cloud providers, and TCO is more easily determined in the cloud.

Managed Services

Managed services are also an option, of course. If the organization does not have in-house expertise or resources, consider a third-party firm that can manage the SIEM/SOAR solution(s) of choice. It ultimately boils down to the requirements of the organization, the most efficient way(s) to meet those requirements and available budget. It may even be practical to start with managed services with a view to

[3] "SANS 2019 Cloud Security Survey," www.sans.org/reading-room/whitepapers/cloud/paper/38940 (registration required)

[4] This paper mentions product names to provide real-life examples. The use of these examples is not an endorsement of any product.

transitioning to an internal team over time. That way the organization can see a more immediate return on its investment in SIEM and SOAR while building out its systems and acquiring the needed resources and training to bring its program up to speed. Starting with managed services will mean more up-front cost but also much faster implementation and maturity.

Consulting Partner Private Offers

Customers can also engage through Consulting Partner Private Offers (CPPO) to work directly with trusted advisors to select and configure SIEM/SOAR solutions from AWS Marketplace. As organizations build out their cloud and cloud security strategy and plan, they may want to consider working with partners to accelerate their efforts or fill any gaps in knowledge or resources that are identified. All consulting partners may extend AWS Marketplace third-party solutions directly to customers through CPPO.[5] Not every organization will be able to find resources with deep cloud experience. Even experienced cloud technologists may have experience only in specific industries or with specific cloud vendors. A requirements document could be helpful when approaching prospective consultants.

Needs and Capabilities: The Business Case for SIEM and SOAR in the Cloud

Among the features that cloud architecture offers for SIEM and SOAR that an on-premises system cannot is visibility across multiple environments in different availability zones or regions. Such visibility could be even more important for global organizations. Consider also that the redundancy of the cloud practically guarantees reliable uptime, which is not available to an organization internally without great expense and multiple data centers.

Then factor in the ability of the cloud provider to offer pricing based on dynamic workloads and short life cycles, where entire environments can be spun up and shut down in a matter of minutes—again, not something a typical data center can provide to an organization. Even leveraging on-premises virtual hosts doesn't offer as much flexibility, especially compared to serverless implementations in the cloud.

Needs and Capabilities

Organizations require a lot of their SIEM/SOAR systems.

SIEM/SOAR

The need: Aggregating log events and security information from multiple systems, collecting data about threats and automatically responding to low-level security events without human intervention

[5] AWS Marketplace Channel Programs, https://aws.amazon.com/marketplace/partners/channel-programs

Capabilities

- Security threat and incident detection
- Bidirectional feeds with Amazon Security Hub
- Increased efficiencies
- Analytics and alerting
- Detailed drill-down compliance reporting
- Increased efficiencies for physical and digital security operations
- Event and threat intelligence correlation

For incident response functions, SOAR supplements SIEM and helps to:

- Define
- Prioritize
- Standardize
- Automate[6]

General Cloud SIEM and SOAR Considerations

Regardless of the SIEM/SOAR technology or cloud vendor selected, some general business, technical and operational considerations are associated with implementing security in the cloud. The following sections highlight many of these considerations.

Business Considerations

	Consideration	Details
	Policies and standards	Organizations will need to evaluate cloud capabilities to determine what changes are needed to ensure that compliance with policies and standards is achievable. Organizations should evaluate relevant retention policies for collected log data. They should determine what happens if a matter becomes litigious and a legal hold on certain data is necessary, as well as where and how data will be held in a secure state for the period of the legal hold.
	Governance model	Organizations need to decide whether to centralize or decentralize governance over cloud incident response and determine whether existing governance models used for traditional incident response can be extended to the cloud or if a cloud-specific model is required. Consider that cloud workloads can more easily span the globe and that data residency and visibility restrictions may apply in certain regions.
	Reporting and metrics	Providing the right metrics, key performance indicators (KPIs) and key risk indicators (KRIs) to the right stakeholders may require changes to account authorization for cloud architectures. Organizations will need to define reporting requirements specific to cloud workloads and evaluate features and products against these requirements.
	Funding and support	Funding and support for cloud SIEM and SOAR implementations may not currently be available. Management may not understand the shared responsibility model as it pertains to cloud usage and may assume that all needed features of SIEM and SOAR are included. Management will need to be educated to understand the implementation model and the related requirements as it determines the appropriate funding and support model.
	Risk classification	Acceptable risk vs. mitigated risk vs. transferred risk (NIST 800-30) is a consideration when determining what action(s) should or should not take place upon discovery of an incident or potential incident. The organization will need to determine the risk of automatically responding to SIEM alerts in an orchestrated manner as opposed to sending certain alerts to a manual queue or ignoring certain alerts altogether.

[6] Tech Target, https://searchsecurity.techtarget.com/definition/SOAR

Technical Considerations

	Consideration	Details
	SIEM capabilities	As organizations update policies and standards to address cloud workloads, they should also identify the technologies needed to comply with these new requirements. Some organizations may choose to be very prescriptive about which technologies should be used, while others may define the required capabilities and allow individual cloud operations teams to select their own technologies.
	Supported technology	Some technologies may not be supported for all cloud services or for all platforms running on cloud services. Organizations need to decide whether they will allow unsupported technologies, and if so, under what conditions.
	Agent-based technologies	No matter how lightweight, agent-based technologies decrease performance. In the cloud, they increase costs. Organizations may have a restriction on the number of agents that can be installed on each cloud resource. Determine how many security agents are already in place to decide whether a limit increase will be necessary. Any specific overhead allowance for agents should be evaluated during any proof of concept. Consider agentless technology options to preserve resources.
	Near-real-time logging and response	Logging is, or is near, real time. Organizations must determine their communication speeds and requirements. Organizations need to decide whether (near) real-time detection and response is required based on their cloud architecture. Consider data to be logged and storage requirements and location(s).
	Secure communication	As log data is collected by the SIEM and forwarded to SOAR, all communications must be secure, verifiable, immutable and forensically sound.

Operational Considerations

	Consideration	Details
	Operational responsibility and model	Operation of cloud resources is substantially different from the operation of traditional infrastructure, and that may affect who is responsible for implementing and configuring SIEM and SOAR capabilities. Organizations need to decide how best to implement and configure SIEM and SOAR technology, and which group(s) will be responsible for these tasks. Multiple teams may be involved, such as the identity management group, AWS architecture and administration group(s) and SIEM/SOAR admins. Determine whether operations should be centralized or decentralized, on premises or in the cloud.
	Monitoring and response	While implementation and configuration of SIEM and SOAR capabilities may be assigned to an existing cloud operations team, monitoring may be the responsibility of others, and response may be assigned separately. Organizations need to determine who will be responsible for monitoring and responding to endpoint security events. Will it be a centralized group, or does it make sense to separate out certain response functions to existing silos?
	Processes and procedures	Organizations may have specific processes and procedures for dealing with security events related to their traditional on-premises infrastructure. It is likely, however, that these processes and procedures will be different in the cloud. Organizations need to create new operational processes and procedures for SIEM and automated incident response in the cloud.

AWS Implementation Considerations

The general considerations discussed so far can help organizations lay the groundwork as well as secure funding and support for SIEM/SOAR functionality in the cloud. Now let's take a more detailed look at some specific considerations an organization will need to evaluate before implementing these solutions in AWS.

SIEM continues to mature, especially with the addition of analytics that allow for orchestration and automation (SOAR). Along with events and logs needed for SIEM and SOAR functionality normally being fed into Amazon-native tools, threat intelligence is also introduced to the AWS environment. Amazon GuardDuty provides additional monitoring and alerts for known threats. Such native AWS services help provide data for analytics. This analysis then leads to the needed detection of threats based on anomalous behavior known to be common to certain malicious activities.

In the considerations we have already enumerated, an organization can begin to determine budget and resource needs for implementing or enhancing SIEM and SOAR technologies. Let's take a look at considerations specifically related to SIEM and SOAR in the AWS environment.

	Consideration	Details
	Cloud context support	Due to the dynamic nature of the cloud, a resource that existed a few hours ago may not exist right now. Because SOAR technologies perform analysis of data or binaries external to the resource itself, there is a chance that when SOAR analysis is completed, the resource may no longer exist. **Evaluate:** • The flexibility for extension of log collections to include context • The additional cloud context (tags or image IDs, for example) that is captured, retained and used by SIEM and SOAR technology to allow correlation of findings and behavior with resources • The special concerns associated with studying resources that have potentially replaced the original resource from which data was gathered • The ability to ensure immutable accuracy with date/time stamps from all sources SIEM technologies typically send data and binaries to separate SOAR systems or to the vendor's cloud infrastructure to perform analysis. Depending on the cloud regions in use, the transfer of data and binaries to different systems could affect technology performance as well as cost.

	Consideration	Details
	Bandwidth and latency	SIEM technologies typically send data and binaries to separate SOAR systems or to the vendor's cloud infrastructure to perform analysis. Depending on the cloud regions in use, the transfer of data and binaries to different systems could affect technology performance as well as cost. **Evaluate:** • The architecture of the tools under consideration • The amount of data that will be transferred and where the data is being transferred from and to • Potential impacts on cost and performance due to bandwidth • Performance impact of latency between cloud regions and other relevant resources
	Logging sources—general	Centralized logging may include events from any or all of the following sources (these logging source lists should not be considered all-inclusive, given that requirements for events to log will vary in different organizations): • Host level • Operations • Security • Application • Firewall • DHCP • DNS **Evaluate:** • Which of the systems will be logged, and which events from those systems. This evaluation helps determine the space requirements for logs. • Storage; set up expandable elastic storage in case of a significant incident that fires off a large number of events. • Interfacing options with Amazon CloudWatch • Long-term storage; leverage Amazon S3 Glacier for long-term storage or overflow storage of logs, especially when review of particular logs may seldom be necessary.
	Logging sources—AWS	AWS CloudTrail offers logging of AWS-specific logging as well as logging common to any environment. • AWS CloudTrail – Security logs – Audit logs – VPC flow logs – API calls **Evaluate:** • Regulatory requirements • Retention requirements • Space requirements • Audit requirements • Amazon S3 Glacier for long-term storage or overflow
	Logging sources—endpoints	Endpoint tools and systems can feed logs to factor into the SIEM and SOAR, tying events together from servers and workstations with data collected from the host environment, network device, and other sources to provide a robust super-timeline related to incidents. Such timelines can paint a clear picture of the incident from birth to death and help with containment and eradication as well as lessons learned to avoid recurrence in the future. • Help desk tools • Asset management systems • Malware • Proxy data **Evaluate:** • Which events will be logged • The ability to manage date/time accuracy with the Network Time Protocol for the environment

	Consideration	Details
	Logging sources—security	Sophisticated security tools, especially those responsible for managing credentials, offer log entries to track such activities in detail. In addition, the origins of threat and vulnerability data, whether open source or commercial, should be factored into the SIEM for review and analysis. • Identity management tools • Credential secure storage • Vulnerability data • Threat data **Evaluate:** • Granularity of logging • Reputation of threat and vulnerability data feeds • Multifactor requirements for access to such powerful tools
	Incident response	Incident response (IR) will use the collected logs in the SIEM to determine when an event should be elevated to incident status. Once an incident is established, the IR team must determine an appropriate response. With the addition of SOAR, well-defined incidents can be contained automatically. The remaining incidents must be reviewed manually by some assigned security operations team for working through an established model, such as NIST SP 800-61. (See Figure 2.) • Automatic response • Manual response and intervention **Evaluate:** • How much manual response is needed? • What is the skill level needed to handle the manual response issues? • What alerts are based on events that are reliable indicators of incidents upon which action can immediately and automatically take place? • Can those incidents be separated from incidents that require further analysis before action can take place?
	Reporting	Reporting is one of the more important aspects of any SIEM/SOAR implementation. Reports will be used by technicians to help determine how to quickly identify and contain an incident as well as for determining the best strategy for eradication of the incident. Reports also document lessons learned to help eliminate or minimize recurrence. Reporting will have different audiences, all of which need the data communicated in the way most relevant for them. Those working in the areas of management, legal and compliance, for example, tend to have less technical backgrounds, so the approach and the language need to be different than a report intended for a database administrator or a web application programmer. • Analytics • Dashboards • Management • Compliance • Legal **Evaluate:** • What are the requirements from management, legal, compliance, security, operations and other teams for necessary reports to assist with evaluation of each area's gaps and to help them complete their tasks? • What report mechanisms and documentation will help pinpoint needed actions? • Are there reports that help with "lessons learned" meetings to reduce repeat occurrences?

Moving SIEM and SOAR to AWS requires the granular evaluation of impact on what needs to be logged. If the information, including context, is not complete enough to be actionable, it is of no use. Speed is also important. Ingestion of events, analysis of events, and alerting or automatic reactions to alerts all need to happen as close to real time as possible. Having all the pertinent data in one location with more-than-adequate CPU cycles, memory, storage space and bandwidth provides an advantage for response speed and resiliency. The other speed factor has to do with sourcing of the logged information. The sourcing will vary between organizations depending on how they utilize on-premises systems versus cloud systems and the connectivity between the two. AWS offers communication "pipes" through AWS Direct Connect that allow up to 10GB connectivity for getting the data from the organization to the cloud and back. Next, determine the sources providing log feeds to the SIEM. Finally, after analysis, determine what responses can be automated and what kind of alerting and reporting are necessary.

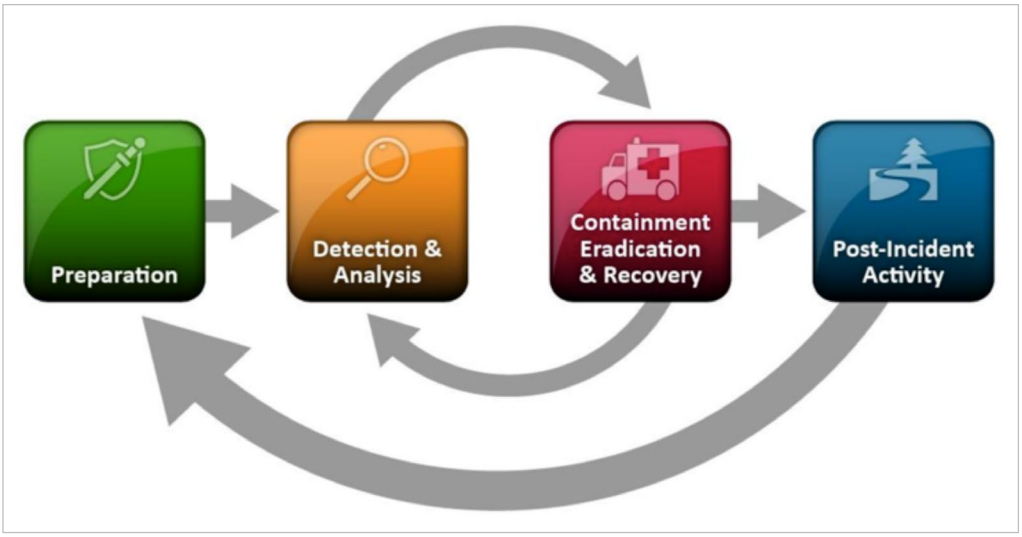

Figure 2. Continuous Integration Process[7]

Making the Choice

To summarize, the key considerations for implementing SIEM and SOAR in AWS include:

- Resources
- Cloud context
- Efficiency
- Ease of use
- Integration requirements
- Availability of built-in tools
- Time to alert and reaction

Have a Plan

Pull together resources from the appropriate teams; management, architecture, operations and information security are all important to the discussion. Determine the desired results from a SIEM system in the environment, then specify the requirements that will provide those results. Separate the "must haves" from the "nice to haves" and share that with the relevant vendors. Don't forget to discuss the requirements with every

[7] NIST, Computer Security Incident Handling Guide, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

relevant cloud vendor, such as any off-premises vendors used for HR, legal, change management, security threat and vulnerability management, or any other outsourced functions, in addition to the major cloud providers, such as AWS.

You must make decisions about what events from which systems must be included in the logs collected for analysis. How granular will the collections need to be in order to meet legal, regulatory, contractual and policy requirements? Don't forget to determine what events do not need to be collected, because every additional event collected will have an effect on data storage and a resulting cost.

Lastly, put together a team of subject-matter experts to decide what collection of events is a reliable positive indicator to trigger automatic response. Determine what the response(s) should be and put together a plan to refine and update those as needed on an ongoing basis.

Consider Partners

An organization should consider using CPPO partners who can accelerate integration of SIEM and SOAR into or with the cloud. As already mentioned, using a third-party vendor to manage the implementation provides the benefit of a quicker ROI and helps bring the organization up to speed operationally. Budgeting for adequate training is also crucial. SIEM/SOAR team members can gain some experience while working alongside partners. Consider the plethora of training videos and courses available from SANS and AWS and their partners that can lead to certification of the technical staff who will manage the cloud implementations. Make sure the partners you choose have a strong background in cloud use and/or consulting.

Don't overlook your cloud provider as a potential partner in achieving success as an infrastructure provider consultant. Speak to your chosen cloud provider to understand which SIEM providers work closely with them. Ask which have achieved security competency and thus are recommended by AWS for cloud environments, for example.

Conduct a Proof-of-Concept Test and Evaluate Options for Desired Features

Your choices must provide the results you expect, or get as close as is reasonably possible. The best way to see how close a vendor comes is to perform a proof-of-concept test. Fortunately, when working with the cloud, services and environments can be spun up temporarily for just such testing. Determine the services you need from the AWS Security Hub, for example, and test the capabilities online. Research which services and systems are available for free testing from AWS and take advantage of those options. Your organization needs to know what to expect from the options it chooses and determine whether those results will add value.

Conclusion

Back to our underground lab full of techies staring at multiple screens: With an adequately funded and implemented analytical SIEM system, supplemented by orchestration and automation (SOAR), security personnel will be spending less time hunting for evil and more time remediating the issues that cause the alerts. In an ideal world, many lower-level incidents will be handled automatically, freeing up personnel to address the more challenging issues that often present greater risk.

With SIEM and SOAR in the AWS cloud, the data center resource needs are handled by AWS. The hardware and everything needed to keep it running are no longer a concern for the organization, freeing up personnel and financial resources for other needs.

To get there, many decisions must be made. See Figure 3 for questions to address.

This paper provides talking points and direction for an organization that wants to move down a decision path. Hopefully, these choices will lead to a quicker implementation of the tools that fit best and provide the best return on investment.

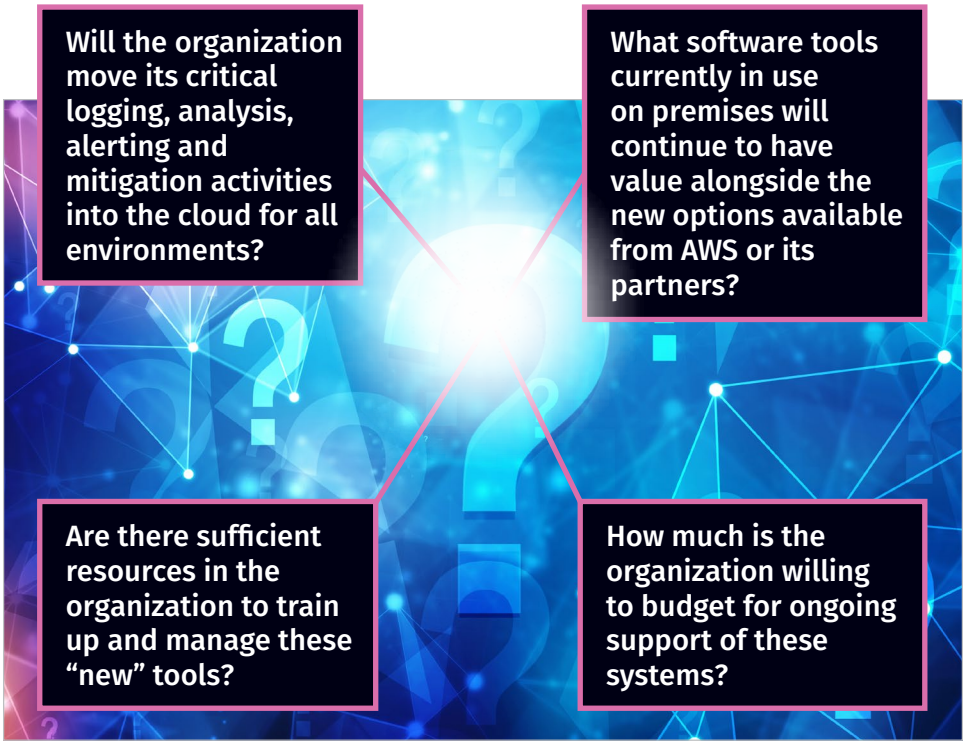

Figure 3. Questions for Cloud vs. On-Premises

Through this evaluation process, look at the features and functionality available from AWS. Many aspects of SIEM collection, analysis and SOAR implementation are already baked into the AWS environment. Careful consideration should be given to the cost delta between leveraging the features and functionality (including AWS partner options) in AWS, as compared to the local data center and its resources.

About the Author

J. Michael Butler is a SANS analyst who has also written SANS security training courseware and audited certification test questions; presents thought-provoking webcasts; and writes position papers, articles and blogs. He is an information security consultant with a leading provider of technical services for the mortgage industry, where he is involved in migration of assets to the cloud. Mike's responsibilities have included computer forensics, incident response, enterprise security incident management planning, internal auditing of information systems and infrastructure, information security policies, service delivery and distributed systems support. He holds the GCFA, GCIH, CISA, GSEC and EnCE certifications.

Sponsor

SANS would like to thank this paper's sponsor:

in conjunction with

About Optiv

Optiv is a market-leading provider of end-to-end cybersecurity solutions. Optiv helps clients plan, build and run successful cybersecurity programs that achieve business objectives through our depth and breadth of cybersecurity offerings, extensive capabilities and proven expertise in cybersecurity strategy, managed security services, incident response, risk and compliance, security consulting, training and support, integration and architecture services, and security technology. Optiv maintains premium partnerships with more than 350 of the leading security technology manufacturers.