ATTACKIQ

White Paper

The CISO's Guide to Cloud Security Using ATT&CK

How to identify, validate, and optimize native security controls within major public cloud platforms.

Notice

AttackIQ® publications are made available solely for general information purposes. The information contained in this publication is provided on an "as is" basis. Any additional developments or research since the date of publication will not be reflected in this report.

Table of Contents

Executive Summary

Threat-Informed Defense Missing in the Cloud

The public cloud is the Wild West of modern security systems, and many CISOs find themselves approaching the O.K Corral with one hand tied behind their back. They know they need to protect their company's cloud-based resources, but they aren't clear on what controls are available in their cloud platforms to provide the requisite protections. And they certainly don't know whether the controls in their Azure or AWS environment could successfully detect or respond to a threat.

Now is the time for the CISO to untie that hand. New research maps the security controls native to Azure and AWS to the adversary tactics, techniques, and procedures (TTPs) in the MITRE ATT&CK® framework. It also scores the effectiveness of each cloud platform in mitigating those TTPs. AttackIQ has been building on this research by creating templates within the Security Optimization Platform that streamline scenario testing of specific cloud-focused attacks, making it possible for security teams to validate cloud security controls at scale, continuously, and in an automated fashion to achieve cybersecurity readiness.

> *"Misconfiguration of cloud resources remains the most prevalent cloud vulnerability and can be exploited to access cloud data and services ... The rapid pace of CSP innovation creates new functionality, but also adds complexity to securely configuring an organization's cloud resources."*
>
> ***–National Security Agency (NSA)***

CISOs who leverage this solution can gain a comprehensive, evidence-based understanding of how well the security controls native to their AWS or Azure platform defend against real-world adversary TTPs. Establishing a regime of routine testing based on the AttackIQ templates enables a CISO to understand whether the cloud security controls, with the settings and configurations the internal security team has selected, are both appropriate and effective at preventing attacks.

Why Cloud Security Typically Lags Behind On-Prem

Cloud environments present a new security challenge. CISOs have traditionally been responsible for securing the clearly defined network edge; it may not have been easy, but it was understandable and discrete. By contrast, the idea of controlling risk in the cloud can appear difficult to fathom.

With a public cloud platform, the cloud service provider (CSP) — whether Microsoft, Amazon, Google, or another organization — is responsible for the physical security of the infrastructure and for keeping customers' data separated. At the same time, the companies using those cloud services are responsible for the security of their data and applications, operating systems, and identity and access management. The cloud platform might come with security solutions, but the customer organization is responsible for configuring those solutions and ensuring that they are working effectively.

The challenge of sharing security responsibility with an external organization is exacerbated by the facts that cloud offerings are constantly evolving and every cloud platform is at a different level of maturity. It's hard to stay on top of what protections are available from each CSP. For example, some cloud platforms have built-in firewalls, but many of their customers don't realize it.

Security operations center (SOC) teams frequently struggle to understand the capabilities and limitations of the security tools that are native to their cloud platforms. As a result, their companies may not be using tools that are available, or the SOC group may choose settings and configurations that are not effective at protecting their resources. Either way, that's a dangerous situation. A Help Net Security survey from 2020 revealed that 70 percent of companies had experienced a security incident in their public cloud environment during the past year.

According to the National Security Agency (NSA), "Misconfiguration of cloud resources remains the most prevalent cloud vulnerability and can be exploited to access cloud data and services ... The rapid pace of CSP innovation creates new functionality, but also adds complexity to securely configuring an organization's cloud resources."

Cloud resources are improving efficiency and information access across a wide range of corporate functions, which means they're not going away. As businesses in every industry sector pursue digital transformation, cloud security is a crucial component — perhaps the most crucial component — of the CISO's current job description.

A Key Resource for Threat-Informed Defense in the Cloud

The first step in building a threat-informed defense is to be well-informed about the threats that the organization faces. However, for SOCs defending cloud environments, information about the effectiveness of specific cloud controls in combating real-world threats is difficult to find. That's why MITRE Engenuity's Center for Threat-Informed Defense set out to develop a research-based solution to the perfect storm of threat-information challenges brewing in the public clouds.

The Center worked with its research partners (including AttackIQ), security vendors, major corporations, and CSPs to map known threats to the cloud-native security controls that were designed to thwart them. The project identified threats using the MITRE ATT&CK framework, a comprehensive knowledge base of adversary TTPs that have been used in cyberattacks around the world. Then it launched a process to map the appropriate threats to the security controls within public cloud security stacks. The project started with Azure, then moved on to Amazon Web Services (AWS), with more to come.

Mapping and Scoring Methodology

For every CSP, researchers start the process by exploring available documentation on the platform's security capabilities. Controls that fall within the scope of analysis must be native to the infrastructure-as-a-service (IaaS) platform. Next, for each selected control, researchers analyze documentation to identify the functionality that will indicate which ATT&CK techniques or sub-techniques it should mitigate, and they map the control to the relevant set(s) of ATT&CK techniques or sub-techniques.

At this point, researchers are ready to analyze the effectiveness of the control in protecting against the threats posed by those techniques or sub-techniques. The project's scoring rubric incorporates factors such as:

- the type of security function the control provides (i.e., is it a detective, response, or protective control, or some combination of the three);
- the control's success at mitigating most of the examples of this threat that are documented within the ATT&CK knowledge base; and
- the frequency with which the control enforces its applied security function.

Each control's final score is either "minimal," "partial," or "significant." Documentation of the project's results includes comments explaining scoring considerations for a particular control, which helps readers understand the degree to which the rating is applicable to their unique environments. For example, a rating of "minimal" detection of database threats might initially alarm a SOC team that uses database-as-a-service (DBaaS) solutions. However, if the comments specify that the rating reflects the cloud platform's inability to detect database threats on an IaaS platform, then that rating does not reflect security solutions' effectiveness in a DBaaS environment.

The Scoring Output

To ensure the results of this analysis are easy to use, the Center developed multiple means of accessing the scoring results. One is a command-line interface (CLI) tool through which users can query the database related to a particular CSP platform, to find the effectiveness scores for a selected control in detecting, responding, or preventing a specific attack.

Alternatively, users can write queries to understand which techniques map to a particular control. This approach can answer questions like: What security function is supposed to be providing detect capabilities for attacks related to privilege escalation, and how successful is it in thwarting those attacks?

Another result of the controls-to-threats mapping process is a visualization tool that shows the effectiveness of the CSP platform overall at detecting, responding to, and/or preventing specific TTPs. The result looks like Figure 1 for Azure and Figure 2 for AWS, as explained by the legend in Figure 3. These visuals represent all the controls within each cloud service provider's native security solutions.

Individually, each matrix reflects the cloud platform's ability to ward off attacker TTPs reflected in the ATT&CK framework. Collectively, they present a rich view of cloud security that has some interesting implications for CISOs and their SOC teams.

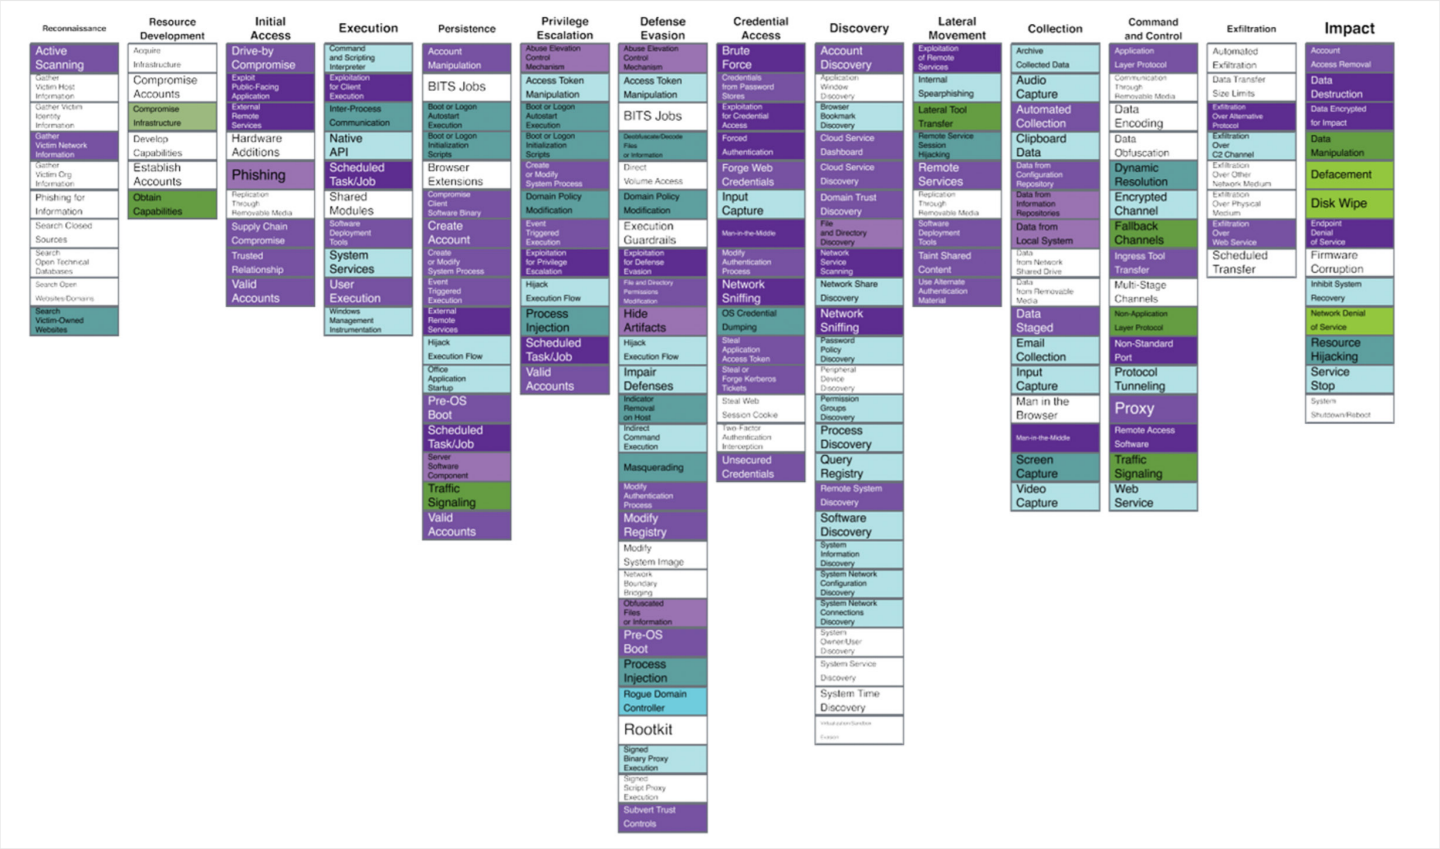

Figure 1:
The output of the Center's mapping and scoring project for Azure.

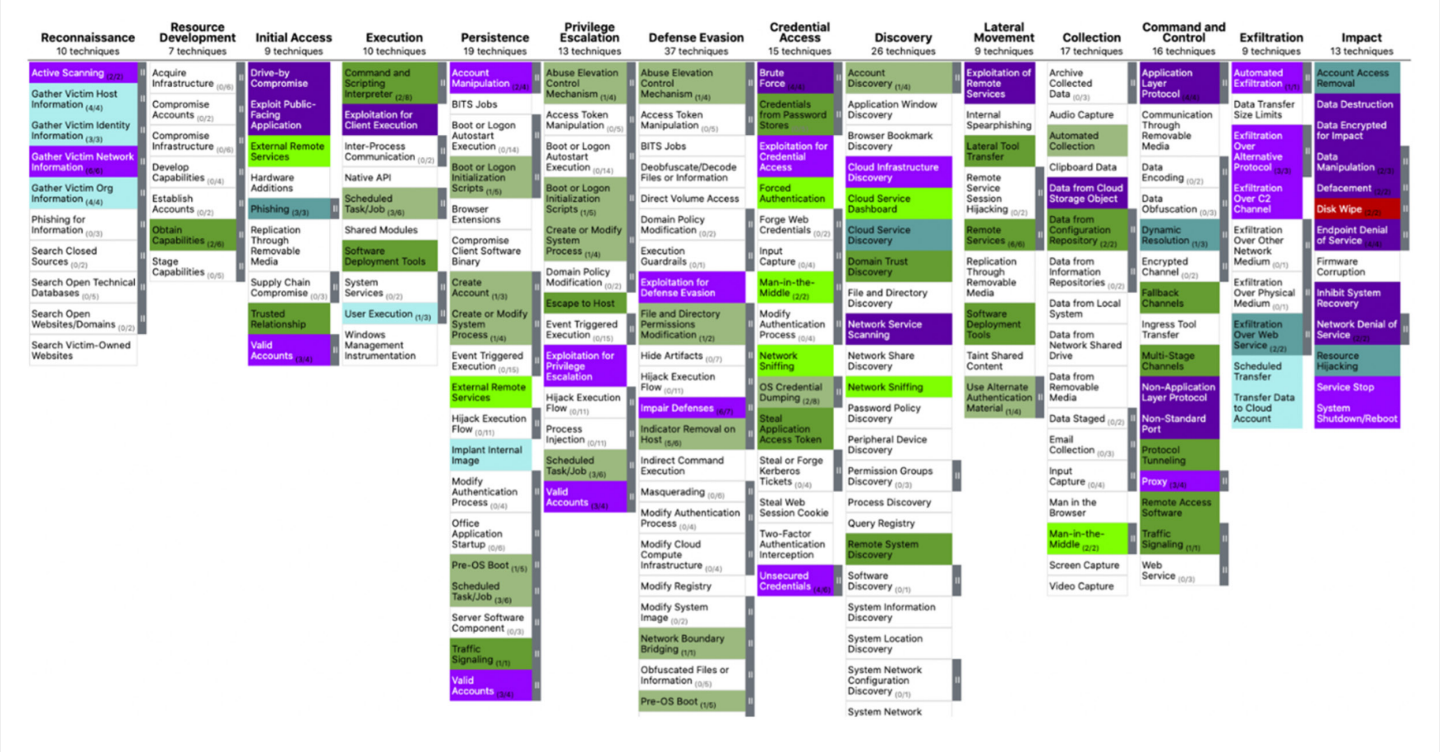

Figure 2:
The output of the Center's mapping and scoring project for AWS.

Using the Results

These independent assessments of which Azure and AWS controls are effective in mitigating threats empower security professionals to determine whether the platform(s) they use are, in principle, capable of protecting them against imminent threats to their organizations. CISOs should already have answers to these questions:

- **What are the business processes that we most need to protect?** The answer should pair the impact on the business if the process were to suffer data corruption, malware, or a distributed denial of service attack with the likelihood of an attack on that particular business process.

- **Which TTPs from the ATT&CK framework should we be most worried about?** The answer derives from determining which TTPs attackers are most likely to use to attack the company and which would have the greatest consequences for the business should the attack succeed.

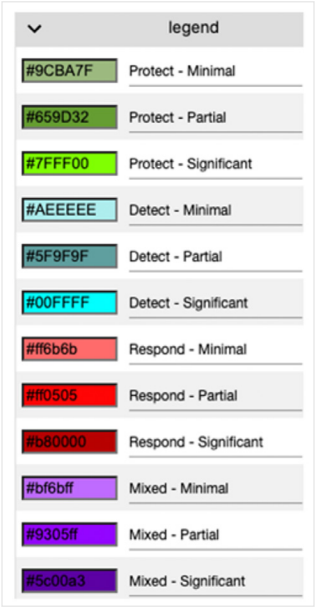

Figure 3:
Legend for both mapping/ scoring matrices.

Once CISOs have clear answers to these questions, they can compare their list of most-feared TTPs against the Center's research results. What are their cloud platform's capabilities in detecting each TTP on that list? What is the platform's response when it detects a threat, and can it successfully prevent most attacks of that type?

This evaluation brings to light any glaring security gaps, highlighting threats that the company's cloud platform is wholly incapable of defending against. In making these determinations, the CISO needs to be sure to look not only at the "minimal," "partial," and "significant" determinations, but also the context provided by the comments. Once the CISO and security team understand the gaps, they can determine next steps for bridging those gaps, perhaps by deploying a third-party security solution to protect the cloud platform.

Insights on Configuration

Meanwhile, delving into the Center's CLI tool can help a security team more effectively configure their cloud platforms' native security controls to protect the company's most important resources. Because some techniques map to multiple controls, the security team may find that they can use different combinations of controls to mitigate one technique.

By combining the mapping information with the scoring results, they may be able to gain new insights into possible approaches to minimize the number of controls they're utilizing while ensuring the effectiveness of the organization's overall security environment.

Developing an Ongoing Program of Threat-Informed Defense

Over the last years since cloud platforms expanded, SOC teams have struggled to manage cloud security, and the movement towards a threat-informed defense represents a big leap forward. Nevertheless, it is only one step in the journey. Just because a security platform is capable of providing a certain level of detection, response, or prevention does not mean all the platform's settings are set to reach that level of performance. Once CISOs understand how their cloud platform's controls map to the ATT&CK TTPs that are most problematic for their organization, they need to assess how well their platform is currently mitigating those threats.

Routine assessments of controls' performance against real-world attacks enables threat-informed defense and cybersecurity readiness. Such assessments gauge the effectiveness of both security staff and technology solutions, in their current configurations, at recognizing and thwarting attempted attacks. Success on these tests demonstrates that the security infrastructure is working as intended. Control failures highlight areas in which security needs to improve.

The most efficient way to incorporate such assessments into a company's day-to-day operations is a breach and attack simulation (BAS) tool, a system that emulates adversary attacks. Automation enables these assessments to run as frequently as needed, so with a BAS solution, a SOC can achieve much more regular and broad-based testing than human penetration testers could. That's why a BAS solution can provide the CISO with evidence about whether the organization's controls are prepared for likely attacks.

Continuous testing supports continuous improvement, which is exactly what's needed to effectively mitigate threats.

Optimizing Cloud Security with the Security Optimization Platform

The AttackIQ Security Optimization Platform is a BAS tool that automates and operationalizes simulations of attacks. Through tight integration with the MITRE ATT&CK framework, the Security Optimization Platform includes prebuilt templates that streamline scenario testing for specific ATT&CK TTPs. Its integration of up-to-date threat intelligence about adversary TTPs facilitates threat-informed defense companywide.

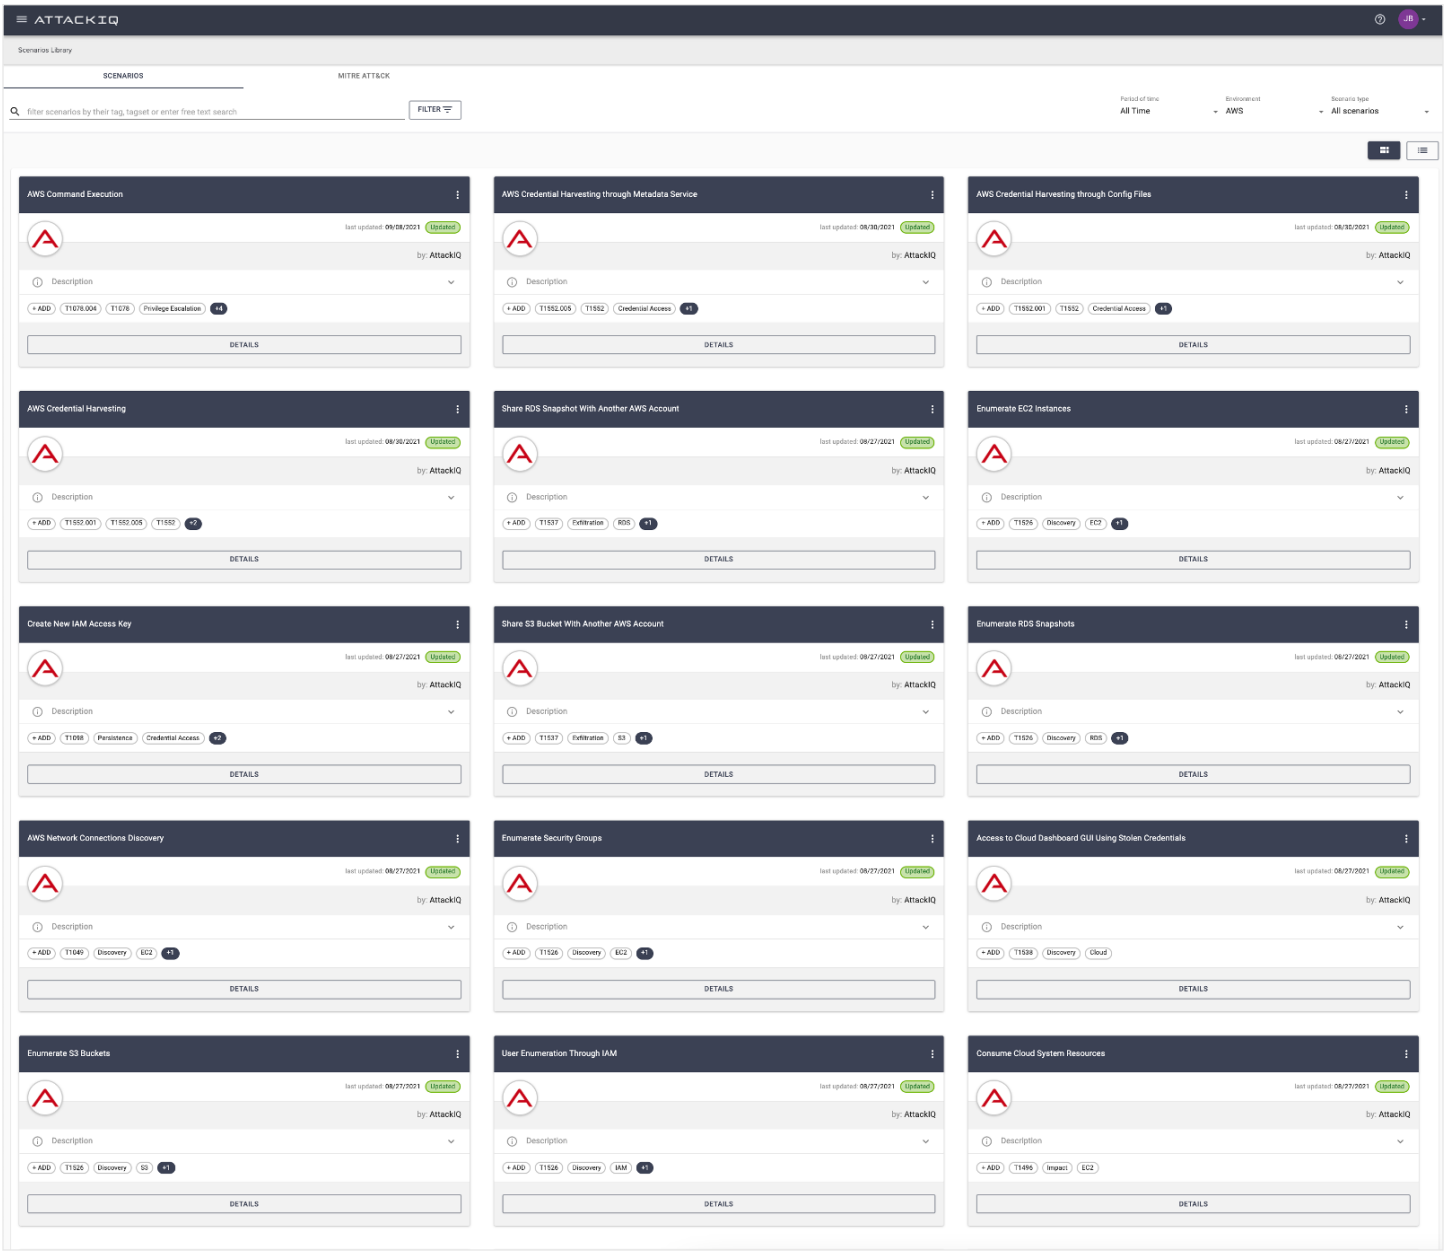

AttackIQ has created templates within the Security Optimization Platform that leverage the learnings of the Center's cloud mapping project. For Azure- or AWS-related TTPs on the company's "most wanted" list, the CISO and security team can run control assessments that answer three key questions:

- Do our defenses map to likely threats?
- Do the security measures we have in place work?
- Are we prepared to respond effectively if something goes wrong?

Chief information security officers should take this opportunity to build out a robust testing regime that simulates, on a regular basis, the attacks the company is most likely to face. Anytime a setting — or security functionality — within the cloud platform changes, the routine testing will identify whether that change introduced a critical security gap. If the security team takes steps to mitigate a discovered gap, they can re-test within the Security Optimization Platform to validate the effectiveness of their response.

Conclusion

Drawing a Map to Security in the Cloud

The transition to the cloud will continue to accelerate in businesses, and, just like in other areas of security, controls need to be testing continuously to ensure operational effectiveness. The good news is that this research from MITRE Engenuity's Center for Threat-Informed Defense, coupled with an automated testing platform, can help pave the way to cloud security readiness. A BAS solution enables the organization to roll scenario testing into day-to-day operations, and Azure- or AWS-oriented templates in the AttackIQ Security Optimization Platform simplify the process of testing native security controls in cloud platforms.

Chief information security officers looking to take tighter control over the security of their organization's cloud-based data and applications need visibility into the effectiveness of each public cloud platform's native security controls. Continuous testing that generates clear performance data is the best way to achieve it — and it gives CISOs the evidence-based analysis that they need to make informed decisions.

Learn more about how to achieve cybersecurity readiness and effectiveness by signing up for AttackIQ Academy's free courses on uniting threat and risk management, purple team operations, and putting MITRE ATT&CK into practice at www.academy.attackiq.com.

ATTACKIQ
U.S. Headquarters
171 Main Street Suite 656
Los Altos, CA 94022
+1 (888) 588-9116
info@attackiq.com