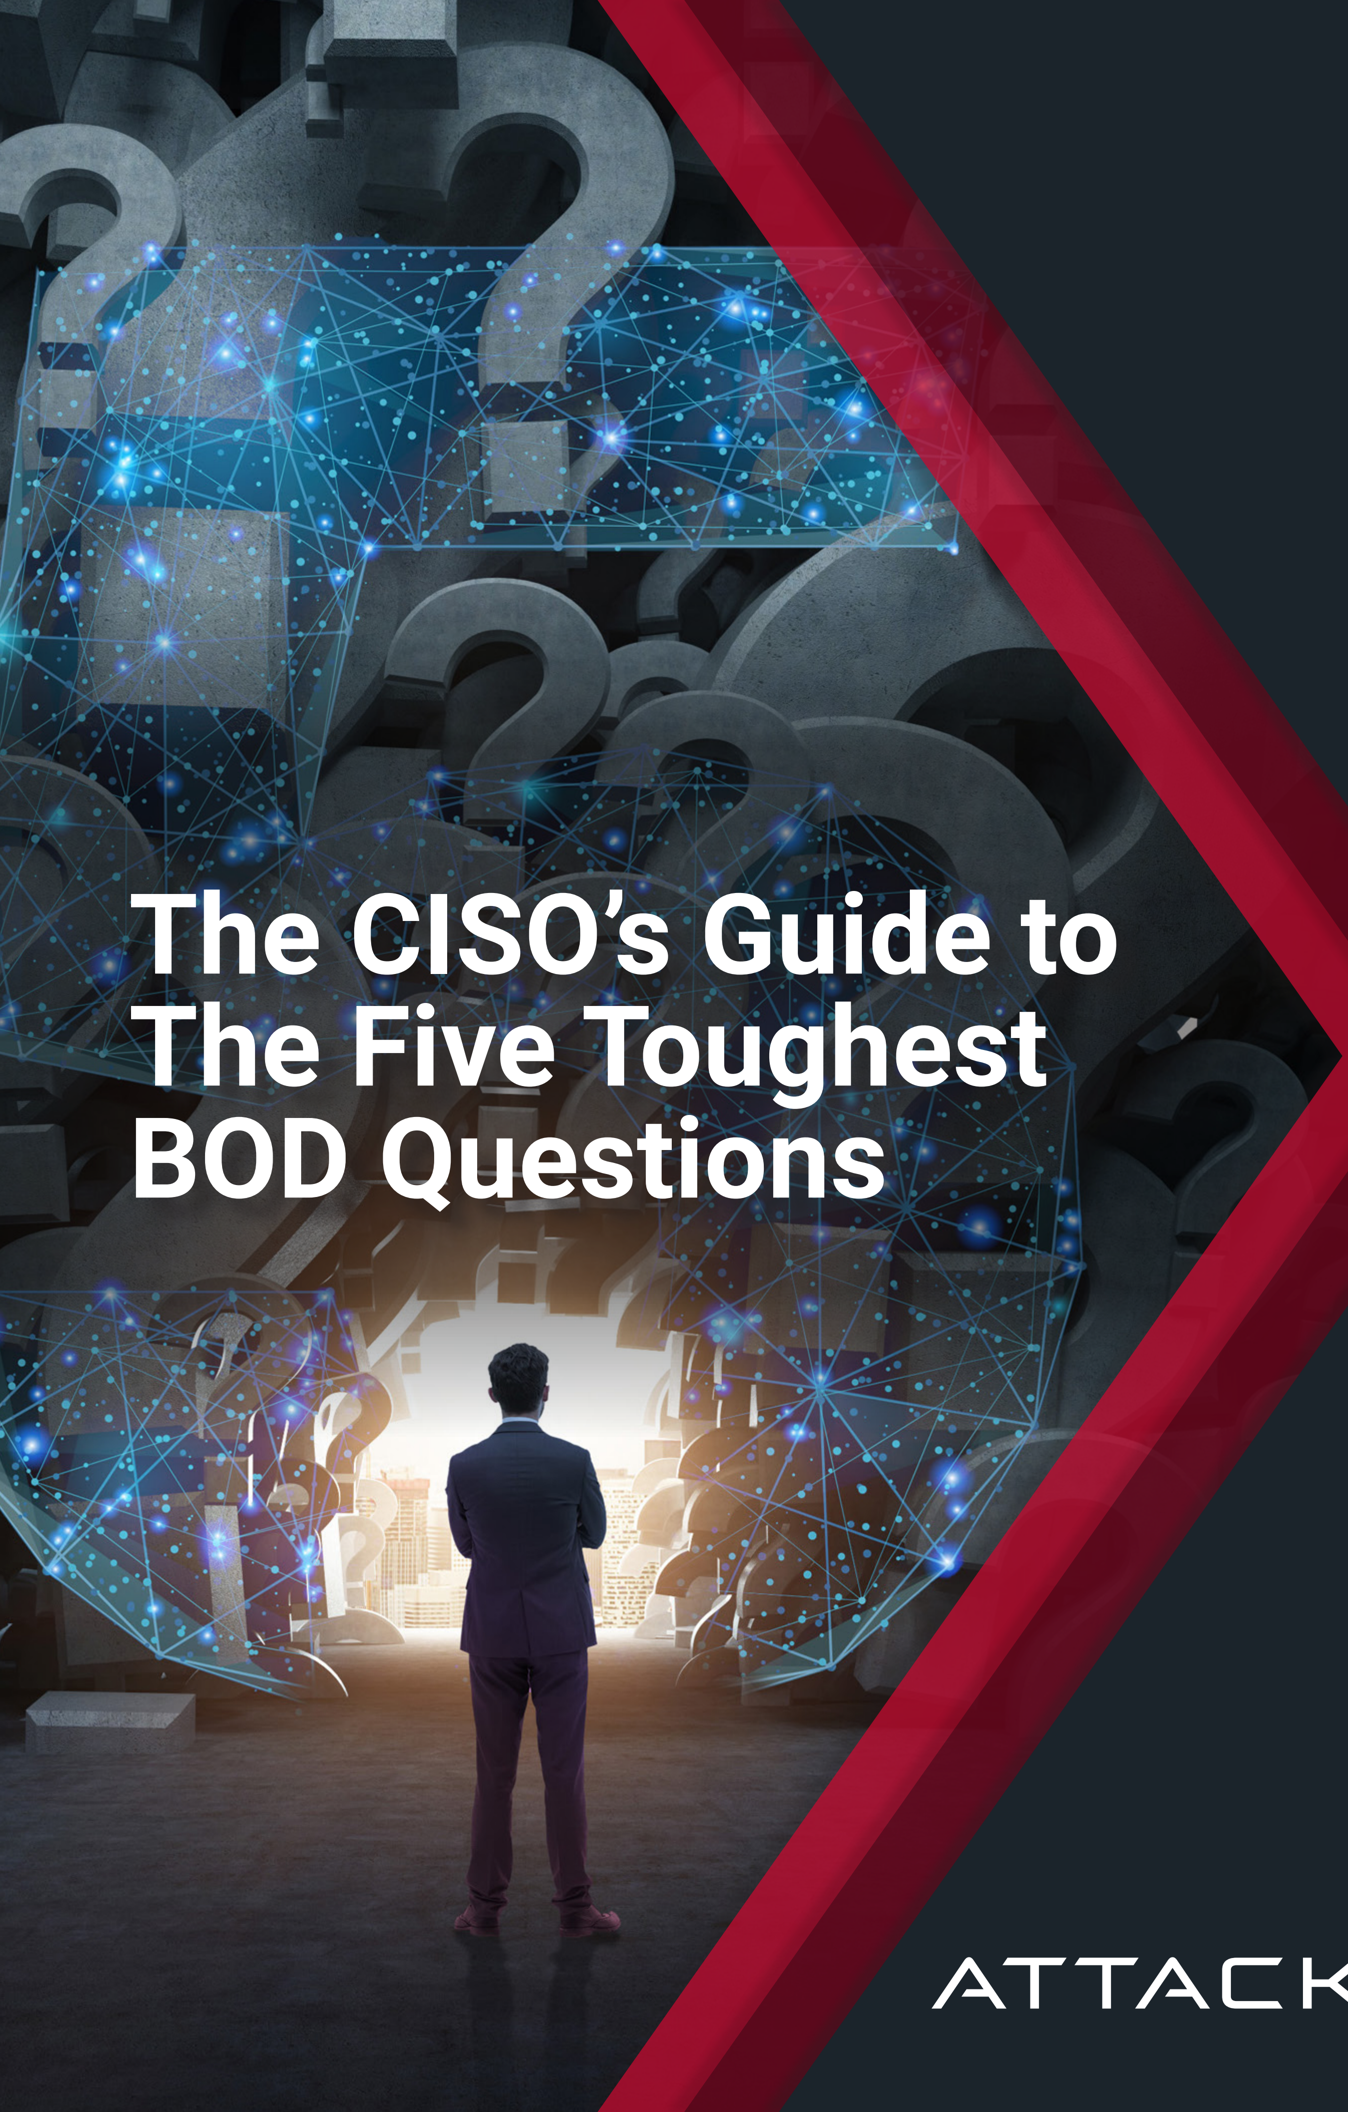

The CISO's Guide to The Five Toughest BOD Questions

ATTACKIQ®

Contents

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.

Executive Summary

This report, the CISO's Guide to the Five Toughest BOD Questions, shares our view of some of the most difficult challenges presented to CISOs today. Many of these questions are received regularly by CISOs from their chief executive officer, other executives in their enterprises, their board of directors (BOD), and their line of business managers, as well as from their partners in compliance and governance.

Why are these questions the toughest?
What makes them so difficult?

> "The board of directors is often concerned when faced with cyber defense budget requests that grow at a rate faster than any other department. The risk models and assessments which CISOs use are based upon many assumptions which are often neither validated nor supported by objective and measurable data. Too much of it is subjective. The cyber defenders truly need visibility to the key data and metrics that can enable them to most accurately identify their threats and best identify the security controls they need to fill the gaps.
>
> It is a fact that many of the largest enterprises in the world, some with massive cybersecurity budgets, still cannot determine objectively if their production security controls, procedures, and personnel are truly protecting their enterprise as they believe."
>
> **Carl M. Wright**
> **Chief Commercial Officer, AttackIQ**

All of this can change.

This report will summarize these five toughest questions, provide context for them, and identify key elements of the missing data that makes these questions so difficult. We will identify new best practice capabilities that bring the challenge full circle to closure by providing the objective data that CISOs and their management teams need to objectively and accurately answer these questions. We will overview new approaches that utilize important cybersecurity frameworks such as MITRE ATT&CK™ and introduce new technologies such as breach and attack simulation (BAS) as part of the new best practice path forward. Finally, we'll share our view of how to get started.

CISO Responsibilities in Context

The five toughest BOD questions relate directly to the position objectives of most CISOs. At the 1,000-foot view, successful CISOs must be able to manage and minimize cyber risk. Risk measurement requires the best working knowledge of the threats the enterprise will face, as well as the ability of the organization to precisely mitigate these threats. CISOs must be able to clearly communicate with top stakeholders about these risks, the perceived threats, and the planned mitigations. They have to make hard decisions about how to manage the residual risk that, in cybersecurity, is about the potential threats you cannot currently mitigate. In some cases, the residual risk can be considerable.

A major driver for commercial cybersecurity budgets is compliance and governance. The protection of confidential and sensitive data (data privacy) is a top-of-mind burning issue today that is being rapidly wrapped in growing layers of regulation and requirement. Most compliance is government-driven, but some is industry-driven, such as the Payment Card Industry Data Security Standard (PCI-DSS) in the financial community. CISOs must meet, manage, address, and document compliance and governance requirements and the solutions they have deployed to meet these requirements. These baseline operating levels and the written reports and compliance risk assessments that go with them are often critical to the well being of the enterprise and map directly to the overall top-level goals of the enterprise.

In government, military, and defense environments, the trade-offs which CISOs make will become necessarily more complex. Nation-state threats are highly sophisticated and often have malicious motives far beyond the theft of data. These motives may include the destruction of the network and information technology infrastructure as part of its ability to wage a future war. This raises the bar substantially on the CISO and other key stakeholders. As we have seen recently, these nation-state attackers may be turned towards our nation's commercial enterprise, with potentially disastrous results.

In the event of a cybersecurity incident, CISOs and cybersecurity defenders must have plans and resources in place to respond rapidly and effectively and to rapidly return the organization to normal operations. Most CISOs know they are always on a war footing. Attackers will penetrate their networks. They will get in - the hardest perimeter can no longer keep them out. So it becomes critical to have the security controls, procedures, and personnel in place to detect and defeat these attackers as early as possible in their Kill Chain® .

Finally, CISOs must not only prepare and request budgets but also rationalize the funding they are requesting. The more subjective the discussion becomes, the more difficult it is for the CISO to get the resources they need. It is frustrating for the CISOs and their teams as well as potentially dangerous for the enterprises they must protect.

The Five Toughest Board of Director (BOD) Questions

We'll take a close look at each of these questions and highlight the challenges. We'll then share how new best practices and new technologies to support them can give us the objective data we need to answer these questions accurately and completely.

1. What are the top threats to our organization and what is our strategy to deal with them?

2. What is our risk of breach?

3. Threat Du Jour! Are we protected against APTxyz?

4. How do we know if our security controls are working?

5. How do we rationalize new security controls?

What are the Top Cyber Threats to our Organization?

This question requires a consistent approach to describing the universe of threats that might impact your organization. It leads to explaining your strategy for dealing with them.

The strategy to deal with these threats is the sum of the mitigations such as security controls, personnel, and procedures. This also includes security architecture, your security framework, and the use of any strategies such as zero trust.

Ultimately, the alignment between the threats and the mitigations must be as precise as possible. The details absolutely count. One misconfigured security control or one over-looked vulnerability can make the entire enterprise vulnerable. You must also address your cloud, on-premise, internet of things (IoT), and mobile device security with one coherent and well-integrated plan. This is difficult work on a good day.

Threats intelligence can help. Threat intelligence is a critical part of identifying threats by geography, industry, enterprise size, and other parameters. You need to know which threats are most likely for your enterprise. When you are concerned about specific attack groups, you need to understand their goal and the specific techniques they use to penetrate your networks and information technology infrastructure.

Can you map your security controls with confidence to each of the techniques used by your most likely threats?

What is Our Risk of Breach?

This is always the million-dollar question. This question requires that the CISO and the cyber defense team be able to assess the risk of those attacks they view as most likely, weighed against their ability to successfully mitigate those attacks. In order to do this, they must understand the specific tactics, techniques, and procedures those attackers will use.

RISK = THREAT LIKELIHOOD VALUE * IMPACT VALUE

- In the most simple sense, risk can be calculated by multiplying the threat likelihood value by the impact value. Then risk can be categorized and compared on a relative basis.
- The threat likelihood value is the potential of a risk occurring as measured using qualitative values such as low, medium, or high.
- The impact value is the damage incurred by the event. If you put a dollar value on this risk you can compare one type of risk to another.
- The risk value can be offset by your ability to mitigate the threat successfully and confidently.

Once again, threat intelligence is a critical part of identifying threats by geography, industry, enterprise size, and other parameters. In order to assess the risk of specific attack groups, you again need to understand their goals and the specific techniques they might use to penetrate your networks and information technology infrastructure.

In the final analysis, the threat likelihood is both a function of your assessment of the subjective probability of a breach attempt, factored with your ability to defend against it. So if the probability of an APT29 attack in your industry is high, but you can identify that you have every mitigation in place to counter the techniques they use, then the threat risk, as mitigated, is reduced to a low value.

Can you objectively calculate your risk today?

Threat du Jour! Are we Protected Against APT29?

APT29 is just an example framed for this discussion. You have likely had to respond to this question in one of your board meetings already. You may have been asked point blank if you can defend against a certain type of malware or a specific recently publicized attacker. These tough questions come at you every quarter and will continue as successful breach attempts grow and continue to get well publicized by the media.

So let's look at our example.

APT29 is a Russian hacker group that is believed to be affiliated with the Russian intelligence service. The threat intelligence on this group is sometimes hard to track, as it has different names assigned by different threat intelligence sources. APT29 is also referred to as Office Monkeys, CozyCar, The Dukes, or CozyDuke, depending on the threat research organization source. In any case, APT29 uses advanced and heavily customized malware that helps identify its activity. The tactics, techniques, and procedures it uses are a clear fingerprint of its identity.

Your organization is a think tank in the United States. Think tanks are favorite targets of APT29 and you view an attack as likely. But is it high risk? Are you really protected against APT29?

In order to make this assessment, you must first successfully disambiguate the various attack groups, which go by different names but which are all in fact one group, to understand the threat. You need to see the malware tools that they use and understand the tactics and techniques that the group would deploy against you.

Now you need to map precisely how your security controls, processes, and personnel will address this expected Kill Chain of activity and mitigate enough of the techniques deployed by the APT29 attack group to successfully stop the attack. How will you get the data to do this? If you hire a red team, how will you know exactly how they decided to mimic APT29? How does this approach compare to industry data on APT29? How frequently can your red team test and will they retest everything again along with APT29? How will you know that your security controls actually work at this moment achieve this mitigation?

Once you know all of this, then and only then can you say that your defenses are well positioned to mitigate likely known APT29 attack techniques. The risk of a successful APT29 attack based upon currently known data is low. So how will you get the data to answer this question? Can you map your security controls with confidence to each of the known techniques used by APT29? Can you tell that your security controls are protecting you against APT29 now?

How do we Know if our Security Controls are Working?

Current events have shown us that the great majority of successful attacks are supported by failure of the cyber defense to perform as expected. Security controls were deployed, but for various reasons, whether misconfiguration, a missing software update, changes elsewhere in the networks, or any other number of reasons, security controls failed to work as expected.

Viewed differently, you previously selected security controls such as data loss prevention, based upon your expectation that this control would address key compliance and governance requirements. It was your belief, previously presented to your board of directors, that the installation of this important security control would allow you to protect your organization against a variety of expected attacker techniques designed to acquire and exfiltrate your organization's data. Yet tomorrow, you might have to explain to that same BOD how your data loss prevention failed to block an attack which successfully exfiltrated over 2 million customer financial records.

To make matters worse, you had your red team test your DLP the previous quarter. So why is the system not working the same today? Are our security controls working as we expect?

How do we Rationalize New Security Controls?

So far, we have carefully reviewed our threat intelligence and now understand the likely threats we will face by geography, industry, enterprise size, and other parameters. We've identified the specific attack groups we view as most likely to be these significant threats. We've considered the most common attacks and understand techniques they use to penetrate targeted networks and information technology infrastructure.

How do we know which attacker techniques can be mitigated by our security controls? Our new risk assessment has identified several new attacks plus all of the expected attacks being widely deployed against manufacturers like our firm.

Which security controls are best for us? Do we need these security controls? Do we have all of the security controls that we need? Exactly how do we know?

For example, firewalls are very time consuming and difficult to configure and set up. Which one gives us the features we need to protect against our expected threats?

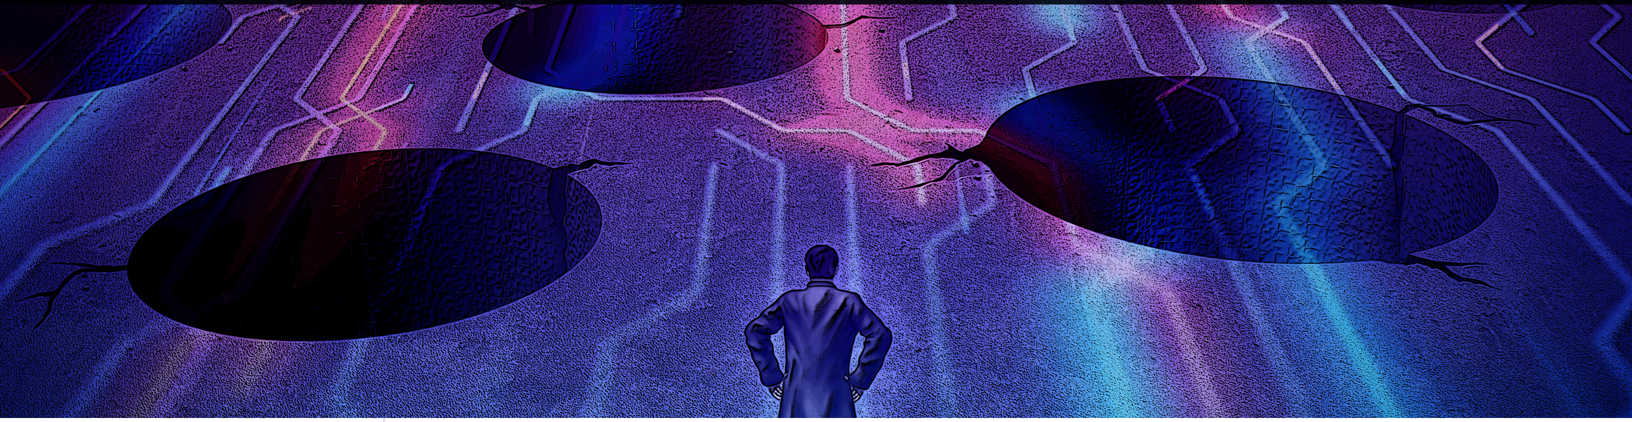

New Best Practice Solutions Fill the Gaps

There are new solutions that enable CISOs and their teams to best answer these difficult questions correctly and accurately.

Step 1 - MITRE ATT&CK Framework.

Deploy an attacker-oriented security framework such as MITRE ATT&CK. MITRE ATT&CK is enjoying a tremendous "15 minutes of fame" and there are good reasons for this. MITRE ATT&CK is, in both depth and breadth, the largest cyber attack knowledge base, providing suggested mitigation techniques, detection procedures, and other important technical information. MITRE has expanded the traditional Kill Chain to include the widest variety of tactics that are then supported by detailed techniques. This organized approach enables you to methodically select and analyze attacks and to compare them to the capabilities of your security controls so that you can understand the gaps. Once understood, you can then rationally expand your security controls and adjust your budgets. MITRE's stature in the cyber community and the independence of its intellectual property in the ATT&CK matrix make it the ideal platform from which your security operations management, executive staff, and board of directors can objectively evaluate and measure cybersecurity controls' performance, risk, and capability.

Step 2 - Breach and Attack Simulation Platform.

The next step is to automate security control validation and performance measurement using a breach and attack simulation system that operationalizes the MITRE ATT&CK framework. With automation, this new infrastructure will continuously validate your cybersecurity controls in your production environments.

Breach and attack simulation platforms allow enterprises to automatically simulate the full attack and expanded kill chain used by cyber attackers against enterprise infrastructure using software test points that allow testing across roaming laptops, user desktops, virtual machines, or cloud infrastructure. The result is detailed reports of the status and performance of your security controls and processes as well as the personnel that support them. Once BAS allows you to find the performance gaps, you can strengthen your security posture and improve your incident response capabilities. BAS can validate that your enterprise security systems are performing against known attacker behaviors.

Lets review how these work and how they help fill the gaps so you can accurately and completely answer the challenging questions in front of us.

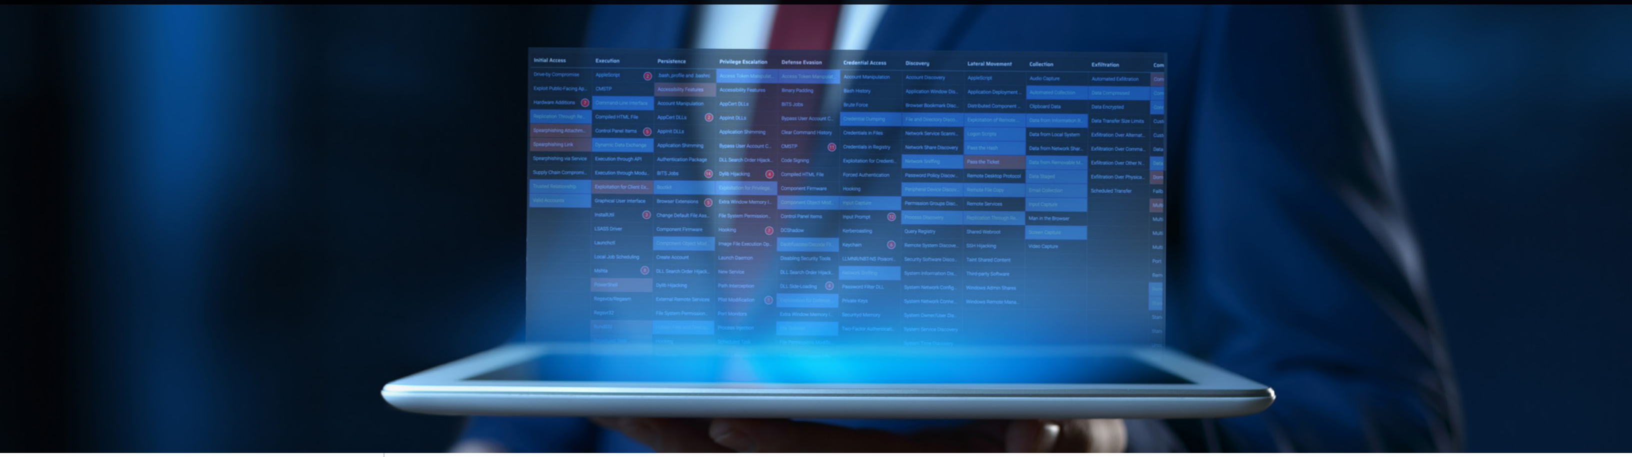

An Introduction to MITRE ATT&CK™

The MITRE ATT&CK enterprise matrix provides a tabular view of all attacker tactics and techniques that might leverage Windows, Mac, and Linux environments. Across the top are headings listing the 12 tactics defined by MITRE ATT&CK. Listed below each of those 12 tactics is a column that shows nine to 67 techniques that might be used to implement a particular tactic. It is often that case that several techniques are used in one or more tactics. A tactic clearly defines the goals of the attacker. A technique describes the different ways that a cyber attacker can achieve the end goals of the tactic.

The MITRE ATT&CK™ Matrix

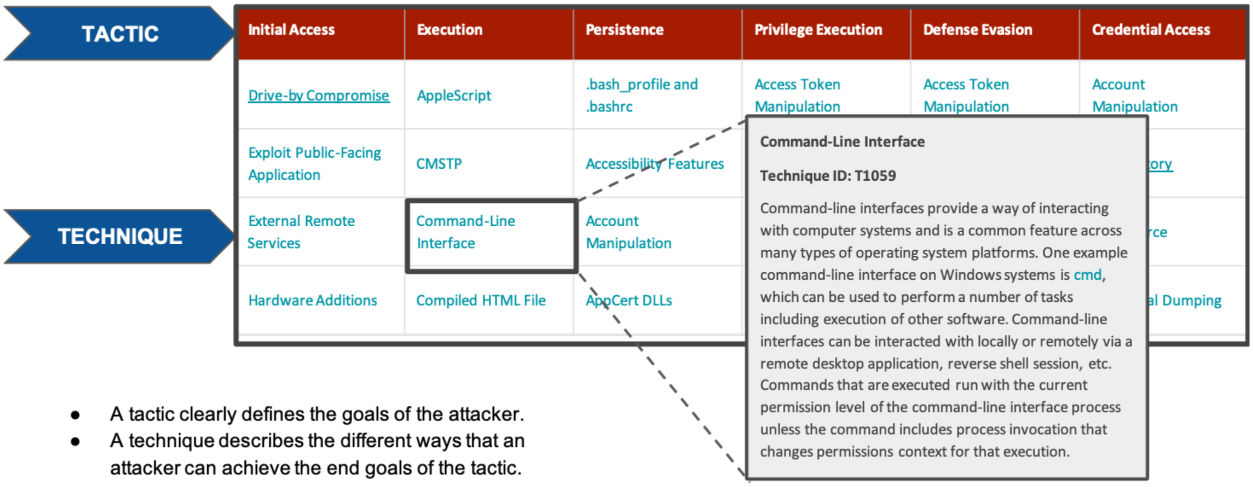

Initial Access	Execution	Persistence	Privilege Execution	Defense Evasion	Credential Access
Drive-by Compromise	AppleScript	.bash_profile and .bashrc	Access Token Manipulation	Access Token Manipulation	Account Manipulation
Exploit Public-Facing Application	CMSTP	Accessibility Features			ory
External Remote Services	Command-Line Interface	Account Manipulation			rce
Hardware Additions	Compiled HTML File	AppCert DLLs			al Dumping

Command-Line Interface

Technique ID: T1059

Command-line interfaces provide a way of interacting with computer systems and is a common feature across many types of operating system platforms. One example command-line interface on Windows systems is cmd, which can be used to perform a number of tasks including execution of other software. Command-line interfaces can be interacted with locally or remotely via a remote desktop application, reverse shell session, etc. Commands that are executed run with the current permission level of the command-line interface process unless the command includes process invocation that changes permissions context for that execution.

- A tactic clearly defines the goals of the attacker.
- A technique describes the different ways that an attacker can achieve the end goals of the tactic.

MITRE ATT&CK's benefits include:

A Common Lexicon

MITRE ATT&CK has compiled a well-matured taxonomy of the tactics and techniques that may be leveraged by any prospective attacker. This provides, for the first time, a common lexicon that enables business stakeholders, cyber defenders, and vendors to clearly communicate on the exact nature of a threat and the objective assessment of the cyber defense plan that can defeat it. This common lexicon brings a universal language that can be used to describe the procedures of an attacker or attack tools and exactly the techniques which they deploy. The precise lexicon of MITRE ATT&CK enables a more precise assessment of threats and a faster, better-targeted response.

The Largest Database of Documented Attacker TTPs

The largest depth and breadth of attack scenarios with suggested mitigation techniques, detection procedures and more. MITRE has expanded the kill chain to include the widest variety of tactics, which are then supported by detailed techniques. This organized approach enables you to methodically select the attack you need to validate your security controls and to understand the gaps so you can rationally expand your security controls set.

Use the Tactics, Techniques, and Procedures of a Real Attacker

MITRE ATT&CK lets you take on the mindset of the attacker. Imagine that one or more cyber attackers are working full time, with no other goal in mind than to break, enter, and compromise your intellectual property, damaging or destroying your information technology infrastructure.

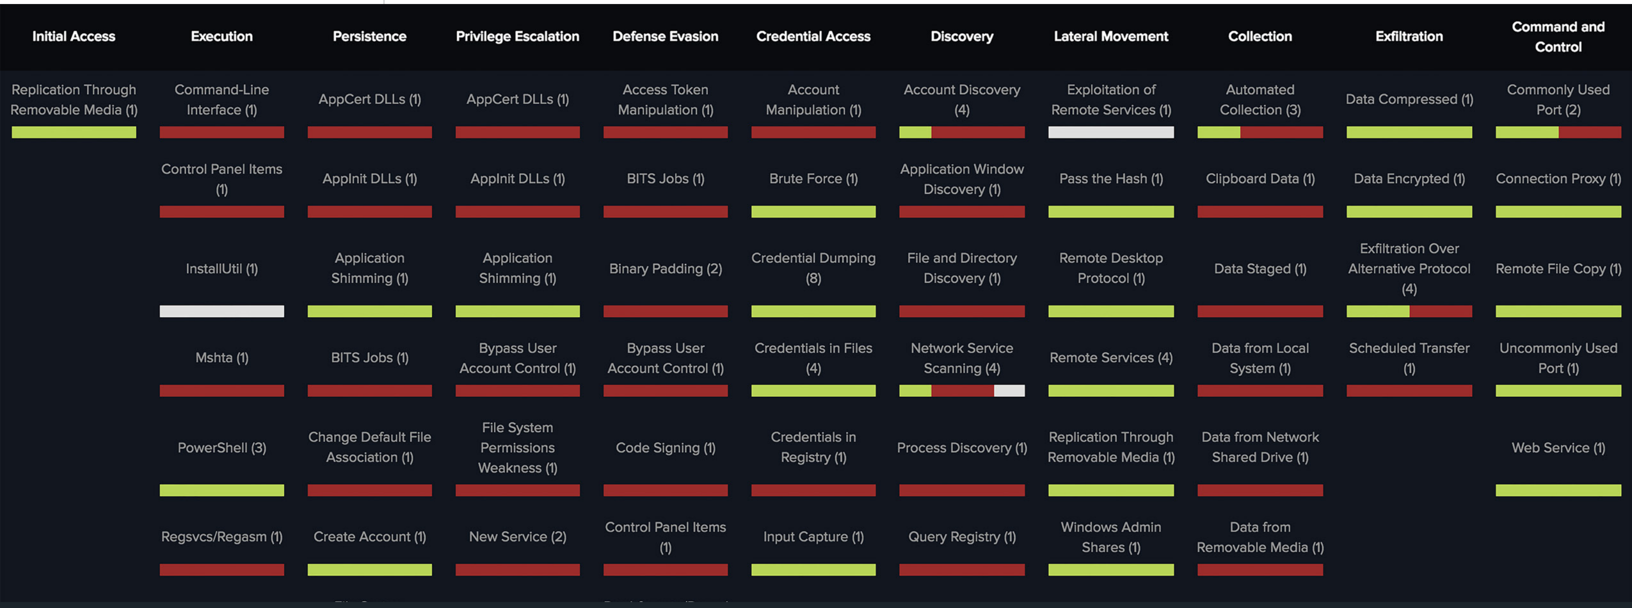

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Exfiltration	Command and Control
Replication Through Removable Media (1) 🟩	Command-Line Interface (1) 🟥	AppCert DLLs (1) 🟥	AppCert DLLs (1) 🟥	Access Token Manipulation (1) 🟥	Account Manipulation (1) 🟥	Account Discovery (4) 🟩🟥	Exploitation of Remote Services (1) ⬜	Automated Collection (3) 🟩🟥	Data Compressed (1) 🟩	Commonly Used Port (2) 🟩🟥
	Control Panel Items (1) 🟥	AppInit DLLs (1) 🟥	AppInit DLLs (1) 🟥	BITS Jobs (1) 🟥	Brute Force (1) 🟩	Application Window Discovery (1) 🟥	Pass the Hash (1) 🟩	Clipboard Data (1) 🟥	Data Encrypted (1) 🟩	Connection Proxy (1) 🟩
	InstallUtil (1) ⬜	Application Shimming (1) 🟩	Application Shimming (1) 🟩	Binary Padding (2) 🟥	Credential Dumping (8) 🟩	File and Directory Discovery (1) 🟥	Remote Desktop Protocol (1) 🟩	Data Staged (1) 🟥	Exfiltration Over Alternative Protocol (4) 🟩🟥	Remote File Copy (1) 🟩
	Mshta (1) 🟥	BITS Jobs (1) 🟥	Bypass User Account Control (1) 🟥	Bypass User Account Control (1) 🟥	Credentials in Files (4) 🟩	Network Service Scanning (4) 🟩🟥⬜	Remote Services (4) 🟩	Data from Local System (1) 🟥	Scheduled Transfer (1) 🟥	Uncommonly Used Port (1) 🟩
	PowerShell (3) 🟩	Change Default File Association (1) 🟥	File System Permissions Weakness (1) 🟥	Code Signing (1) 🟥	Credentials in Registry (1) 🟥	Process Discovery (1) 🟥	Replication Through Removable Media (1) 🟩	Data from Network Shared Drive (1) 🟥		Web Service (1) 🟩
	Regsvcs/Regasm (1) 🟥	Create Account (1) 🟩	New Service (2) 🟥	Control Panel Items (1) 🟥	Input Capture (1) 🟩	Query Registry (1) 🟥	Windows Admin Shares (1) 🟩	Data from Removable Media (1) 🟥		

Operationalizing MITRE ATT&CK With Breach and Attack Simulation Platforms

BAS platforms provide automation that enables the platforms to work autonomously and to scale to support the largest global enterprise. Support for live production environments enables you to see in real time how changes to configurations or administration can open new vulnerabilities in your cyber defense.

AttackIQ's BAS platform provides the setup of scenarios that are used to test your technology controls, validate your security posture, and instrument your environment. Scenarios will mimic malware and attack vectors so that you can confirm that your security controls are working as expected. The fast path to productivity is to test your existing security controls to validate they are performing as you expect.

Since our earlier example was APT 29, we'd like to share that early this year AttackIQ completed the **APT29 Assessment Template**: a group of scenarios that emulate many tactics, techniques, and procedures (TTPs) of the APT29 threat group. We built an assessment template to emulate the behavior of APT29, selecting and configuring the necessary scenarios to cover the entire post-exploitation attack chain of this threat group: from the first stage after compromising a machine to the later stages of communication with a command and control server and exfiltration of sensitive information.

In total, this assessment template contains **45 scenarios covering 56 MITRE ATT&CK techniques**. We based this selection of techniques on the one done by MITRE for their new round of evaluations, which will have APT29 as their simulated attacker.

We decided to group these scenarios in nine different tests according to their MITRE ATT&CK tactic, except for the last test where we decided to group the scenarios belonging to either Command and Control or Exfiltration tactics. In our case, the order of the tests corresponds to their position inside the MITRE ATT&CK matrix which, roughly speaking, corresponds to the depth of the intrusion after the initial breach.

This assessment is designed **to run almost out-of-the-box**, from the commands that certain scenarios will execute to the file types that the Collection scenarios will search, and from the credential dumping tools that will be used to the command and control servers (controlled by AttackIQ) to which some information will be exfiltrated. However, there are minor configuration options that the user will have to specify since they highly depend on the environment where the assessment is run.

The Five Toughest BOD Questions Objectively Answered

Let's take a look at these questions again and see how operationalizing the MITRE ATT&CK Framework with AttackIQ's breach and attack simulation platform will help you provide the data you need to answer these tough questions.

1. What are the top threats facing our organization and what is our strategy to deal with them?

Threat intelligence is a critical part of identifying threats by geography, industry, enterprise size, and other parameters. The strategy to deal with expected threats is the sum of the mitigations to include security controls, personnel, and procedures and a precise understanding of how they are performing.

MITRE ATT&CK provides detailed analysis of threat groups, helps disambiguate them, and shares data on the tactics, techniques, and procedures they use. Your threat intelligence can also supplement this data. You can notate in the MITRE ATT&CK enterprise matrix exactly which techniques are used by your top threats and then prioritize them based upon risk assessment.

You can then operationalize and test against these likely attacker techniques using AttackIQ's BAS platform to see how your current security controls mitigate the threat. If your security controls mitigate the threat, you can proceed to continuously validate your production systems going forward.

If your security controls do not mitigate this threat, you can see all of the holes in your defense and use this insight to prioritize decisions to acquire new security controls, check and reconfigure your existing security controls, change internal policies, and more to continue to mitigate and reduce the probability of a successful attack from these likely threats.

You can also perform complete one-to-one mapping between techniques that will be used in an attempt to breach your enterprise and the mitigations you have or are planning to acquire. Your strategy is supported by clear, crisp data with precise information on which attacker techniques you can prevent and which you cannot.

2. What is our risk of breach?

As before, the threats have been identified by threat likelihood value and impact value.

RISK = THREAT LIKELIHOOD VALUE * IMPACT VALUE

You have moved from a sea of unknowns to a much narrower set of prioritized target threats. Now you can articulate risk based upon successful mitigation of the specific techniques used by that threat. And you can clearly measure and report on that mitigation for all of the simulated likely threats using the BAS platform.

Mitigate it successfully and you've minimized the risk. If you're still missing the security controls, processes, and policies to address it, you now have an objective source to show a higher risk of breach quantified around these very specific parameters. And now you have an authoritative third-party reference, MITRE ATT&CK, to provide credibility to your assessment.

As before, threat intelligence is a critical part of identifying threats by geography, industry, enterprise size, and other parameters upfront.

In the final analysis, the threat likelihood is a function of both your assessment of the subjective probability of a breach attempt and of your ability to defend against it.

3. Threat Du Jour! Are we protected against APT29?

Board of directors meetings are almost always a forum for questions about the latest threats which are being publicized by the media. And for good reason! You may even want to raise the question yourself. A new threat has emerged in your industry and you are concerned - you may need additional funds for new security controls to better combat this problem. Do your existing security controls in your production protect against this new threat? One of your peers in another major bank just was hit by this APT … are you protected?

Once again, you can find complete data using the MITRE attack groups data, including the tactics, techniques, and tools that this new threat uses. Now you can test this scenarios we have developed for APT29 using the AttackIQ BAS to see how your existing security controls respond in your production systems.

"That's a great question. Based upon our threat intelligence sources, our organization is at significant risk for APT29. Recognizing that we might be attacked by this group, we ran multiple simulated scenarios for APT29 against our current security controls and we're concerned about a few specific areas. Here's exactly where we are unprotected against the specific techniques this attacker is known to use and here are the best possible solutions in terms of investment in security controls, processes, and personnel that we need to make. If we can make these investments we feel the risk of a successful AP29 attack on our organization will be substantially reduced."

4. How do we know if our security controls are working?

As previously mentioned, the great majority of successful attacks are supported by a failure of the cyber defense performing as expected. There are a plethora of variables that might, and often do, cause security controls to not perform as expected and leave thought-to-be-secure systems vulnerable.

Now you can continually measure the performance of your security controls through a feature called continuous security validation. Any configuration or other change in your software or network that might cause problems for your security controls will be immediately identified by AttackIQ's ongoing simulation.

A number of our partners also provide mapping of all of the techniques that their security control products mitigate successfully. You will be able to use their identified BAS testing scenarios to validate that their products are working to their specifications.

5. How do we rationalize new security controls?

Every department in the company has to justify budgets and expenditures. This has been traditionally hard for the cybersecurity team.

The use of a BAS platform allows you to map your expected threats to your existing security controls and to then identify precisely the weaknesses in your defense. You can quantify risk more precisely and present a compelling case to address the high risk and likely threats with new security controls, processes, and personnel.

This completely changes the tone of budget discussions. Every expenditure can be prioritized and presented in the context of likely threats and the objective measured ability of the security infrastructure to mitigate that threat.

Recommendations

Some of the toughest questions facing every CISO today can now be answered easily, accurately, and objectively by operationalizing a security framework such as MITRE ATT&CK with AttackIQ's breach and attack simulation platform. You can measure the performance of your cybersecurity controls, allowing you to identify the gaps and make improvements to meet the most likely threats in your environment.

AttackIQ BAS provides a powerful platform with which to implement and operationalize MITRE ATT&CK. The value provided by a BAS platform can be compelling for you and your organization and will enable you to continuously validate the performance of your production security controls, as configured, against these scenarios and many more.

To view a demonstration or participate in a free trial of our award winning BAS platform, please reach out to **info@attackiq.com** or visit us at **www.attackiq.com**.

ATTACKIQ
U.S. Headquarters
9276 Scranton Road, Suite 100
San Diego, CA 92121
+1 (888) 588-9116
info@attackiq.com

AttackIQ, a leader in the emerging market of breach and attack simulation, built the industry's first platform that enables red and blue teams to test and measure the effectiveness of their security controls and staff. An open platform, AttackIQ® supports the MITRE ATT&CK Matrix, a curated knowledge base and model for cyber adversary behavior used for planning security improvements and verifying that defenses work as expected. AttackIQ's platform is trusted by leading companies around the world. For more information visit **www.attackiq.com**. Follow AttackIQ on Twitter, Facebook, LinkedIn, and YouTube.