ATTACKIQ

White Paper

The CISO's Guide to Using Attack Graphs and MITRE ATT&CK®

Advanced purple teaming using multistage adversary emulations and Attack Flow research.

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

Emulating Real-World Attacks

Historically, enterprise security teams have prepared their networks for attack by looking at singular tactics, techniques, and procedures (TTPs). Based on real-world adversary behaviors from the MITRE ATT&CK® framework, the "atomic" approach is a useful first step toward a threat-informed defense. Yet it fails to consider the interlinking elements of an entire attack sequence, limiting its effectiveness for defense teams. Atomic approaches often also fail to trigger responses from artificial-intelligence- and machine-learning-based security tools, as they do not resemble a real attack.

That is why the industry is moving toward the adoption of attack graphs (also known as attack flows). They emulate multi-stage attacks, and defenders can use them to evaluate their automated cybersecurity controls in a realistic and specific way. Attack graphs are particularly useful when applied to purple teaming exercises, as they enable security teams to move beyond a description of an attack sequence to outline methods for detecting threats and breaking the kill chain.

AttackIQ publishes attack graphs, along with recommendations for detection and mitigations, when new threats emerge. Given Vladimir Putin's war in Ukraine, this has recently focused specifically on responses to cyberattacks from Russia-based actors, but AttackIQ publishes attack graphs to help defenders prepare their defenses against a panoply of actors. We also work closely with the Center for Threat-Informed Defense (CTID) in its Attack Flow project, and we will incorporate any advances from the Center's research into our own attack graphs.

Attack graphs and the Attack Flow project provide defenders with an important tool to test their defenses in a realistic way, and will help CISOs answer when the CEO asks, "are we prepared?"

Moving on From Atomic Threat Analysis

Good boxers think like chess players; they don't just throw one punch and hope for the best. Rather, they plan complex combinations of techniques with the aim of unlocking their opponent's defenses. Conversely, a good defense relies on being able to anticipate these combinations and react appropriately and at speed. It is a skill that fighters can only learn through sparring—mock bouts in the ring. The closer it feels to a real fight, the better prepared the boxer will be when the match bell chimes.

The same is true of cybersecurity. Traditionally, CISOs and their defense (or "blue") teams have tracked adversary behaviors atomically, focusing on one specific action at a time. This is like a boxer concentrating on countering a single jab from an opponent, before then thinking about their uppercut.

That's not to say that the atomic method is redundant. It is a good first step towards threat-informed defense, focusing as it does on real-world attacker behavior. Yet it misses the bigger picture of the holistic attack sequence. As a result, atomic methods can make it difficult for blue teams to mount an adequate response to attacks.

There are two other key drawbacks with the atomic method:

1. It is neither realistic nor specific, and as a result **does not emulate real-world attacks**. Just as a boxer needs time sparring in the ring before a fight, defensive teams can only prepare for attacks through realistic simulations.

2. The lack of realism inherent in atomic testing also means that it is **unlikely to trigger a response from AI- or ML-based cyberdefense tools**, as most of these will not recognize an isolated action as being consistent with a real threat.

The requirement for CISOs and their blue teams is therefore clear: find a way to emulate realistic attack scenarios that takes account of attack sequences in their entirety.

Defending in the Round: MITRE ATT&CK and Multi-Stage Emulation

Security teams already have at their disposal all the information they require to build a comprehensive, multi-stage threat-informed defense. This is because attackers often use a range of known TTPs that are already in existence. These TTPs have been cataloged in the MITRE ATT&CK framework, a globally accessible knowledge base of adversary tactics and techniques based on real-world observations and threat intelligence.

Some of the most infamous cyberattacks of recent years have leveraged known TTPs found in the MITRE ATT&CK framework, such as the 2020 SolarWinds breach, which impacted more than 30,000 public and private organizations. While the Russian government first used a novel supply-chain attack to break past perimeter defenses, once inside the victims' networks, the attacker used known tactics that had long been in existence, including techniques that enabled them to move laterally.

The known behaviors described in MITRE ATT&CK therefore provide a good starting point for CISOs looking to evaluate their defensive shields. The aim should be to bring together MITRE ATT&CK behaviors in attack "graphs" or "flows" that illustrate a multi-stage attack in the round.

Attack Graphs in Action

Aligned to the MITRE ATT&CK framework, the AttackIQ Security Optimization Platform emulates a range of ransomware families, and strings together adversary TTPs in attack graphs that emulate the adversary with specificity and realism. Defenders can use these graphs to test their organization's security program continuously and automatically. Enumerating complete kill-chain sequences in this manner both better prepares organizations for real-world attacks and provides high-level efficacy when testing modern ML—and AI—based security controls.

Here is an example based on the WannaCry ransomware attack, which affected the operations of several high-profile organizations, including the U.K.'s National Health Service. Despite its relative age, WannaCry is still a threat to organizations all over the world — and the techniques within the malware strain continue to defy unprepared defensive teams.

Leveraging the MITRE ATT&CK framework, and building an attack graph in the AttackIQ Security Optimization Platform, we can see that:

- The WannaCry cryptoworm actively scans for systems (ATT&CK tactic T1595—Active Scanning) vulnerable to CVE-2017-0144, exploited by EternalBlue.
- Like the real threat, the AttackIQ scenario checks specified targets for an SMBv1 service being published. It can be configured to either check for this vulnerability only or exploit it if the check passes (T1190 —Exploit Public-Facing Application), opening the door for the worm to infiltrate the WannaCry payload onto vulnerable targeted assets.
- If the target asset is vulnerable, the WannaCry payload is saved to the file system, and, if not stopped by endpoint security controls, it executes peripheral device discovery and remote system discovery scripts.
- This emulates the behaviour of the worm to enumerate additional local filesystems and files for encryption as well as additional targets for malware propagation. It then establishes an encrypted Tor channel, enabling command and control (T1573), and, finally, the worm encrypts its first target (T1486), allowing the intruder to hold the data for ransom.

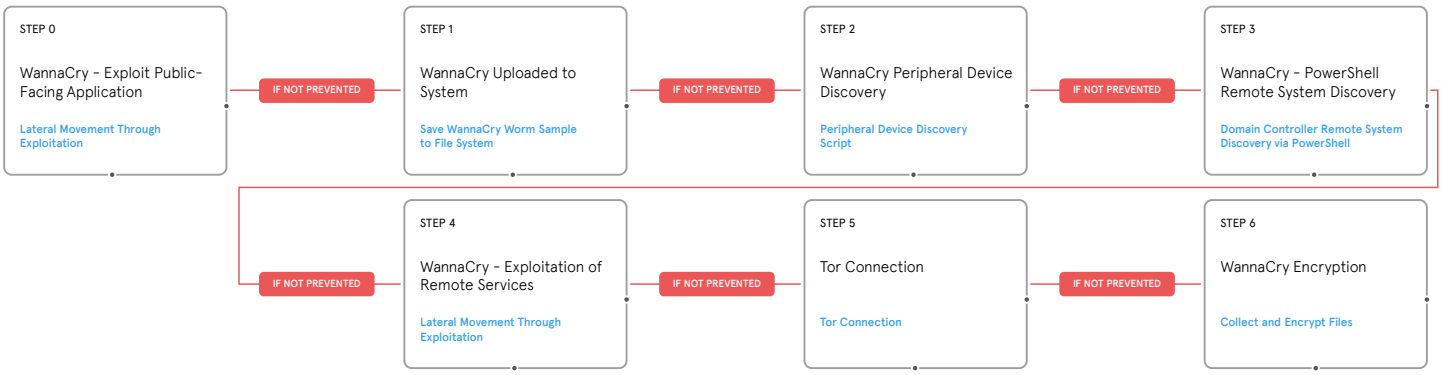

Attack graph for the Wannacry ransomware cryptoworm

WannaCry is a historic case, but organizations need to be prepared for the latest threats. As in boxing, speed and agility are crucial components of a strong defense. The better prepared with up-to-the-minute threat intelligence, the better able organizations are to respond. It is for this reason that AttackIQ continually monitors for alerts from the United States Computer Emergency Readiness Team (US-CERT) around current security issues, vulnerabilities, and exploits. We then use the alerts to generate attack graphs within days to help organizations prepare their defenses as soon as possible (see the example below in "Applying Attack Graphs in a Purple Team Environment").

Purple Teams and Attack Graphs

The relevance of attack graphs to blue teams is clear. It provides a schematic by which they can automatically test their cybersecurity defenses against real and relevant threats (i.e., those threats that history tells us are most likely to affect a given business in a given industry). But attack graphs are optimal in collaborative, purple team environments. This is because the demands of guarding against rapidly evolving attacks have outstripped the ability of blue or red teams to understand and adapt to the complete threat landscape alone.

Siloed security functions simply don't work as well as collaborative organizations aligned around known threats. Lines of communication must be open in every direction, and blue team members should work closely with red team colleagues to improve their knowledge of the anatomy of different types of attacks, starting with building consensus on which threat groups, tactics, and techniques pose the greatest risk to the organization.

Purple teams should focus first on reviewing attack variants and TTPs jointly, developing from there a prioritized list of attack graphs to investigate. Working together, red and blue teams can also better come up with mitigations to any flaws in their defenses identified by the attack graph, as well as the best approaches to break the kill chain in an attack.

What is a purple team?

In a purple team, the still-distinct red (attack) and blue (defense) teams develop highly communicative, supportive, and cooperative relationships across the functional boundary to better protect the enterprise from cyberattacks.

Applying Attack Graphs in a Purple Team Environment

On March 15, 2022, US CERT issued alert AA22-074A on Russia-based actors disabling multi-factor authentication (MFA). Just three days later, AttackIQ issued a new attack graph to emulate Russia-based threat actors as they exploit MFA protocols to disable MFA. The attack graph is worth describing here as it is a good example of how the tool can enable a purple team construct for cyberdefense operations, inclusive of detection and mitigation recommendations.

The alert followed the compromise of an NGO (non-governmental organization) by Russia-based actors, that played out as follows:

As early as May 2021, the threat actors leveraged a misconfigured account and enrolled a new device for MFA to gain initial access to the victim's network.

The PrintNightmare (CVE-2021-34527) vulnerability was then exploited to gain system privileges, which were used to modify the Windows hosts file to redirect the victim's MFA solution to a local host causing it to be unreachable and "fail open," disabling MFA for all accounts.

Next, the threat actors collected and used legitimate credentials to access the victim's virtual private network (VPN).

Once connected to the network, the intruder obtained additional domain account credentials, and established Remote Desktop (RDP) sessions to other Windows domain controllers.

With this additional access, the threat actors moved laterally to the victim's cloud and email accounts to access and exfiltrate data.

When any such alert is issued, organizations need to review the mitigation recommendations made by US-CERT and implement them immediately. They also need to validate their security program performance against the adversary to ensure that they are prepared.

It is here that the attack graph comes into its own. This attack graph emulates the attack covered in the US-CERT alert:

Attack graph in response to US CERT AA22-074A

The graph is just the start. To inform a purple team construct, it is also important to discern detection and mitigation recommendations to employ and improve your security posture. When US-CERT published the alert, we provided detection and mitigation recommendations to companies to help them get ahead of the techniques and procedures we know the adversary was likely to employ.

In total we outlined eight attack techniques using the graph, along with detection and mitigation recommendations. Here are two by way of example:

Scenario	Detection	Mitigation
Brute force password guessing (T1110.001). The attack is successful due to a victim account with a weak password. The AIQ scenario uses the top 25 RDP usernames and passwords identified in a leak of 1.3 million Windows RDP credentials published to an underground marketplace. The scenario needs to be configured with a host to target with the brute force attempts. **Remote System Discovery** (T1018). Once an initial foothold has been established, a threat actor is going to look for additional hosts they can access. This scenario uses a combination of Windows tools to find connected hosts using net view and then ping to validate they can be reached by the asset.	• If possible, use a SIEM product to correlate centralized logs and detect unusual login activity on your network. In the case of this US-CERT report, remote login attempts are of interest. Some parameters for your SIEM detections are as follows: • Source Type: WinEventLog:Security • Logon Type: 10 • Event Codes: 4624, 4625. • If possible, use SIEM, EDR or ProcMon tools to consume logs to monitor usage of interpreters using net.exe and ping commands in unusual ways indicating host discovery enumeration. • Cmd.exe or PowerShell.exe running commands such as "net view /all" from any unexpected user accounts • Nmap usage in the environment with the -sP flag indicating a ping sweep • Cmd.exe or PowerShell.exe running commands containing "for /L" AND "do ping," indicating a possible ping sweep script • Check Microsoft-Windows-DNS-Server-Service for Event ID: 6001 from any unusual servers or IPs for unauthorized DNS Zone Transfers.	• Ensure Account Lockouts are set and not waived • As you operationalize CTI, incorporate known password leaks if possible • Conduct Regular Active Directory security assessments • Ensure MFA policies and configurations are set to remain enabled for inactive accounts that have not yet been disabled • Regularly audit for inactive accounts and ensure inactive accounts are disabled; consider adding this to an offboarding process. • Ensure DNS Zone Transfers are restricted to a whitelist of known hosts • Restrict the availability of cmd.exe and PowerShell.exe to necessary users (admin groups) • Restrict usage of network diagnostic tools such as nmap.exe to necessary users (admin groups) using Application Control tools while considering adding detections for these types of tools in your SIEM.

AttackIQ and the Center for Threat-Informed Defense

As part of broader community engagement, AttackIQ is a founding research partner of the Center for Threat-Informed Defense's (CTID) and deeply involved in the Center's Attack Flow project. Through the project, CTID is developing a data format for describing sequences of adversary behavior to improve organizations defensive capability. Attack Flow will enable security teams to visualize, analyze, and share sequences of actions and the assets they affect, advancing the industry's understanding of adversary threats and how to manage them.

The Center has enlisted a range of industry leaders to engage in this project. As representatives from the Center state in a recent blog, the Center is working with participants including Anomali, AttackIQ, Citigroup Technology, Cybereason, Cyber Threat Alliance, The Global Cyber Alliance, Fortinet, Fujitsu, HCA, Microsoft Corporation, and Verizon Business Services. AttackIQ will reflect future updates from the Attack Flow project in the development of our own attack graphs. Having helped shaped the research behind the Attack Flow project, and given the broad industry support behind the project, we anticipate that the Center's work will prove an invaluable resource for CISOs and their purple teams.

Conclusion

In an environment where the threat environment is getting more challenging by the day, CISOs and their teams are under pressure to explain their defensive posture to the C-suite, understand and apply lessons learned from major incidents, and build realistic adversary emulation scenarios for purple teaming exercises. Attack graphs are the solution to all three of these needs. They provide practical schematics of an entire attack sequence, enabling defenders to map detection and mitigation responses so that they are prepared in the event that attackers turn their sites on their business. Champion boxers go into the ring with extensive sparring practice and an analytical understanding of the moves their opponents will likely take. Attack graphs enable security teams to do the same and land a knock-out blow on their adversaries.

ATTACKIQ

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

About AttackIQ

AttackIQ, the leading independent vendor of breach and attack simulation solutions, built the industry's first Security Optimization Platform for continuous security control validation and improving security program effectiveness and efficiency. AttackIQ is trusted by leading organizations worldwide to identify security improvements and verify that cyberdefenses work as expected, aligned with the MITRE ATT&CK framework. The Company is committed to giving back to the cybersecurity community through its free AttackIQ Academy, open Preactive Security Exchange, and partnership with MITRE Engenuity's Center for Threat Informed Defense.

For more information visit www.attackiq.com. Follow AttackIQ on Twitter, Facebook, LinkedIn, and YouTube.