SANS

A SANS Spotlight

Are Your Security Controls Yesterday's News?

Written by **Matt Bromiley**

July 2019

Sponsored by:

Cymulate

Introduction

Within information security (InfoSec), we find ourselves in a constant juxtaposition. On one hand, threat actors are constantly developing, honing and advancing their tactics, techniques and procedures (TTPs). On the other hand, defenders are often playing catch-up, simply waiting for patch updates or third-party reports to inform them of potential dangers. If you are reading about it in the news, it's usually too late to start thinking about your available defenses.

Furthermore, many organizations have resorted to testing their security controls using a handful of common, "modern" methods, such as:

- Vulnerability scans (whether patches are applied is another story!)
- External or internal penetration testing, with a goal of identifying security gaps
- Red teaming, either by an internal or external team

However, many organizations limit just how successful security control testing can be. Penetration tests, for example, are often given time or scope limits. Vulnerability scans are either largely ignored or pointed at small pieces of the entire enterprise. The success of a red team—and thus the takeaways for the organization—is completely dependent on the skill level of the red team itself. Here's the core problem with these methods:

Threats don't care about scope,
and they pay attention to *all* your vulnerabilities.

In this spotlight paper, one of a two-part series, we discuss just how successful an organization can expect to be if it's using old news, limited scope or "cookie-cutter" vulnerability scans as a way to assess its environment. Our belief is that security control testing needs to improve significantly to emulate *actual*—not hypothetical—threats to an organization.

With this paper, we have released a very brief poll[1] to gather results on these topics. We will also be publishing a follow-up Spotlight paper to examine those results and discuss how organizations are dealing with security control testing.

Right now, we challenge you to challenge your organization by asking the questions in Figure 1.

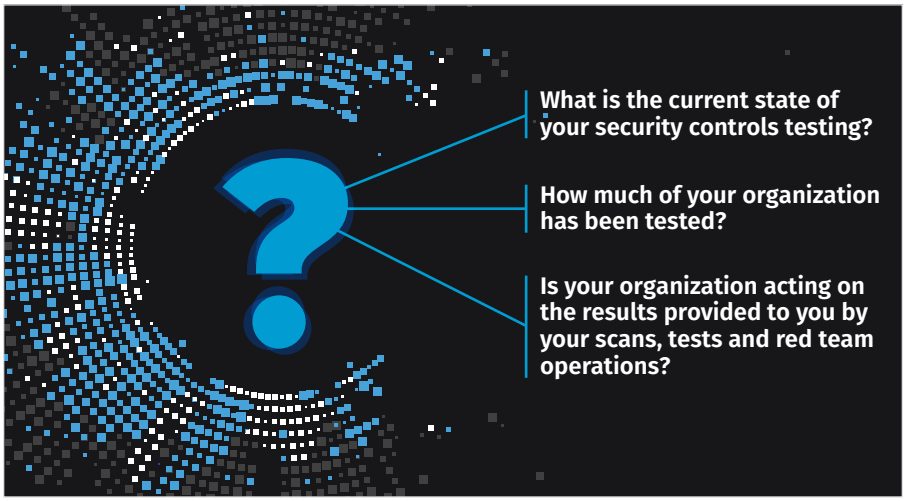

Figure 1. Questions to Ask Before Implementing Security Controls

Yesterday's News

When it comes to implementing security controls, many organizations rely on outdated or extremely limited data points. Let's examine some of the more common techniques to see whether they provide truly actionable output for an organization.

Bottom line: Could you actually detect and defend against a true actor—advanced or not—that is not bound by scope, time or permission?

Vulnerability Scans and Awareness

One common technique used by organizations to identify weaknesses in their environment is the use of vulnerability scanners. Vulnerability scanners allow for an organization to test certain areas of its environment (often the perimeter or dedicated systems or subnets) for well-known and documented vulnerabilities. And therein lie two issues: vulnerability scanners test for *known* vulnerabilities, against *predefined* sections of the enterprise network.

Aside from well-documented, media-reported or extremely critical, vendor-pushed vulnerabilities, by the time you've received a vulnerability scanning report, there's a chance the data is already outdated. Simply "knowing" that something is vulnerable cannot help you anticipate the steps that a *real* threat actor may take after a vulnerability is exploited.

If you are building your security controls in alignment with yesterday's threats, you are always playing catch-up.

Another issue is when vulnerabilities are so dangerous, they receive media attention and/or special, off-schedule updates. Despite the media coverage and publication of known vulnerabilities however, many organizations simply choose to ignore the provided advice! Look at the recent prolific spread of malware such as WannaCry and NotPetya, both of which reportedly accounted for billions of dollars in losses. The vulnerabilities exploited by this malware were well known, documented and patchable. Yet they were still wildly successful.

[1] Visit www.surveymonkey.com/r/SecurityTestingPoll to take the poll, which will close sometime in mid-August 2019.

Penetration Tests

Perhaps the most popular form of testing security controls is penetration testing, which can be done internally or by an external party. Penetration tests are a very common form of control testing in today's organizations, largely due to requirements from data compliance regulations such as PCI or HIPAA.

These regulations stipulate that organizations housing, processing or transmitting regulated data must undergo periodic testing to ensure they are protecting the data of interest. Some organizations adhere to the minimum compliance requirements, resulting in periodic testing (which is good) of limited areas exposed to only one type of data (which is bad).

Penetration tests required for compliance standards may still come with inherent scope, which means your environment is not being properly tested for true threat actor movements.

Unfortunately, this approach presents multiple challenges and should represent a minor footnote in a security program, at best. A compliance-driven penetration test may inherently limit the scope of the test, thus providing little value to the rest of the organization.

Even penetration tests without limited scope may not provide the most useful information for your security controls. Penetration tests are highly dependent on the skill level and speed of the penetration testers themselves. A very skilled penetration tester may use a particular technique, for example, to steal and harvest credentials. Defending against *that technique* may stop the penetration tester, but where do you stand in the grand scheme of advanced threat actors?

Another issue with penetration testing—and this extends to red teaming activities as well—is timing. These types of assessments are conducted at scheduled, periodic intervals. Often done monthly or quarterly (any longer time frame is simply too long), the information security team may have a leg up in preparing for a known, upcoming test.

Even worse, we have seen organizations delay their scheduled tests due to change freezes, holiday breaks or other concerns in the environment. As we've said, threat actors don't care about change freezes, corporate meetings or other blockers. Instead, they strike when it suits them.

True Threat Actor Simulation

Knowing that all commonly used techniques may not be providing the value they should, where do we go from here? Do we simply discard previous methods and hope our controls will work? Luckily, no. The answer to proper security control testing lies in *true* threat actor simulation.

Simulating a Real Attack

True threat actor simulation involves first mapping and detailing the steps that various actors take once they have gained a foothold in an environment. This is a crucial step, because a true simulation will help you recognize the *best controls* available to block certain threat actor activity.

One of the more common techniques used to detail threat actor activity has been aligning known activity with the MITRE ATT&CK™ Matrix. The ATT&CK Matrix, as shown in Figure 2, provides a way to utilize common language and terms to label the various techniques a *true* threat actor may employ during an intrusion.

The ATT&CK Matrix helps answer the most crucial security question that vulnerability scans and scope-limited penetration tests simply cannot answer: "What happens next?" If a penetration tester is successful in compromising an entire environment, the organization can react to the *specific route* the penetration tester took. By understanding how multiple threat actors approach their intrusions, organizations can begin to scale out and react to *all the routes* that true threat actors may take.

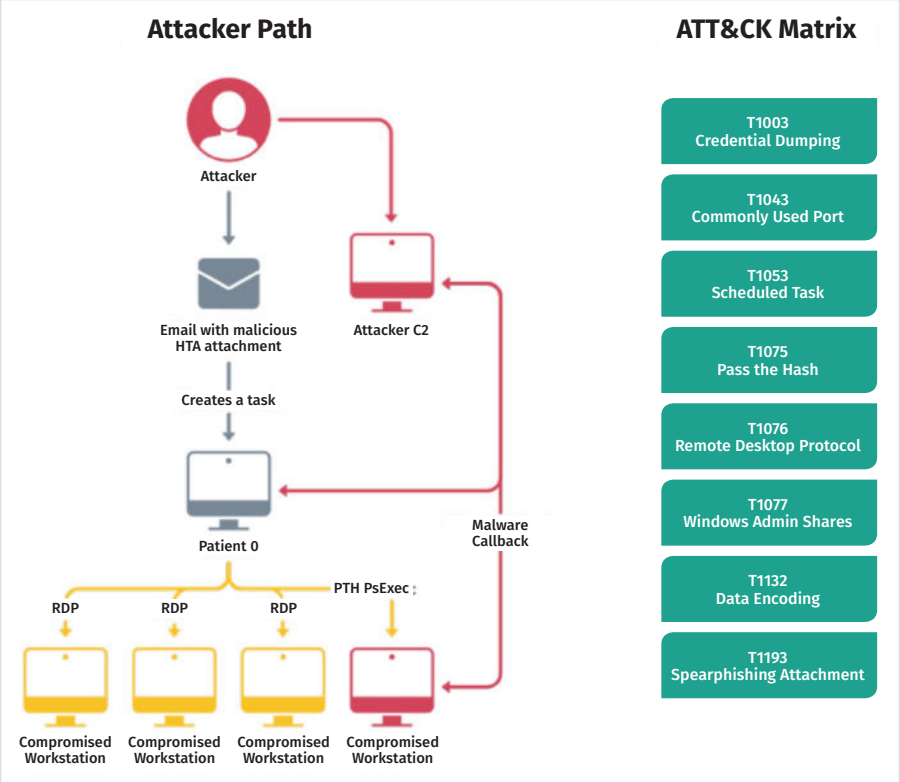

Figure 2. Threat Actor's Actions Mapped Against MITRE's ATT&CK Matrix

Time Is Your Advantage

Not only is *real* threat actor simulation the best way to determine whether your security spend can protect against threat actors, but it's the only way your organization can keep the advantage of time. By continually running simulations and tests, you are no longer bound by the restrictions of a penetration test schedule or the release of a vulnerability patch.

Very few threat actors use zero-day or unknown malware during their breaches. The reason why previously patched, well-known malware families succeed during intrusions is that organizations have lost the advantage of time. The concept of remaining up-to-date on software patches is not new, but it is still underpracticed by some organizations. By constantly testing, and escaping the limitations of scoped penetration tests, you can truly identify where security control spend will have the most impact in your environment.

To determine whether your organization can fall victim to an intrusion from an advanced threat actor, you need to test the actual steps the threat actor would take.

Closing Thoughts

In this paper, we set out to examine the techniques currently used by organizations to test their security controls. Unfortunately, years of extremely public data breaches and intrusions have proven that our current testing methods simply aren't working. Vulnerability scans and scoped-down penetration tests aren't allowing for organizations to implement security controls that stop *actual* threat actors—only hypothetical ones.

We've also set the stage for a poll, from which we hope to gain insight into exactly how organizations are testing their security controls. We will then provide a second spotlight paper, based on the poll results, to examine the collected data and to recommend additional steps organizations can take for better security.

If you find yourself or your organization living in the past, use our section on true threat actor simulation to move your security control testing activities toward achieving results your organization can actually act and improve upon.

About the Author

Matt Bromiley is a SANS Digital Forensics and Incident Response instructor, teaching Advanced Digital Forensics, Incident Response, and Threat Hunting (FOR508) and Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response (FOR572), and a GIAC Advisory Board member. He is also a principal incident response consultant at a major incident response and forensic analysis company, combining experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, memory and network forensics, as well as network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor: