Secureworks®

WHITE PAPER

DANGEROUS ASSUMPTIONS: SCIENCE, SECURITY, AND THE ADVERSARIAL TESTING IMPERATIVE

Untested assumptions are dangerous. They're also a primary cause of cybersecurity failure. Here's how you can avoid them.

Science is a vital human discipline. Without it, we would still be treating infections with leeches rather than antibiotics.

Contrary to popular mythology, however, science isn't just about some lone genius making a great discovery. It's also about others testing the claims about said discovery. Validation through testing is, in fact, central to the scientific method.

Despite this fact, organizations often fail to apply the scientific method to their own cybersecurity assumptions. As a result, they are horribly vulnerable to common threat actor tactics —especially those associated with ransomware — even though those vulnerabilities could easily be remediated if they were known.

These vulnerabilities exist almost without exception, simply because organizations don't put their cybersecurity assumptions to the test.

Fortunately, there's a proven way to apply the scientific method to your organization's cybersecurity: adversarial testing. Adversarial testing entails skilled, experienced cybersecurity professionals investigating, penetrating, and validating suspected vulnerabilities, stopping just short of actually compromising your environment. Testers do this all using the same advanced tools, tactics, and techniques as real-world cybercriminals do.

Adversarial testing thus uniquely enables you to discover exactly where your cyber defenses are working and where they fall short. Armed with that empirically grounded knowledge, you can then take concrete steps to strengthen your defenses precisely where and how they need strengthening—thereby dramatically reducing your exposure to ransomware and other cybersecurity risks.

So the real question is not whether your organization should engage in adversarial testing. That's a given. The real question is how you can most effectively leverage adversarial testing to achieve maximum risk mitigation return on your investment.

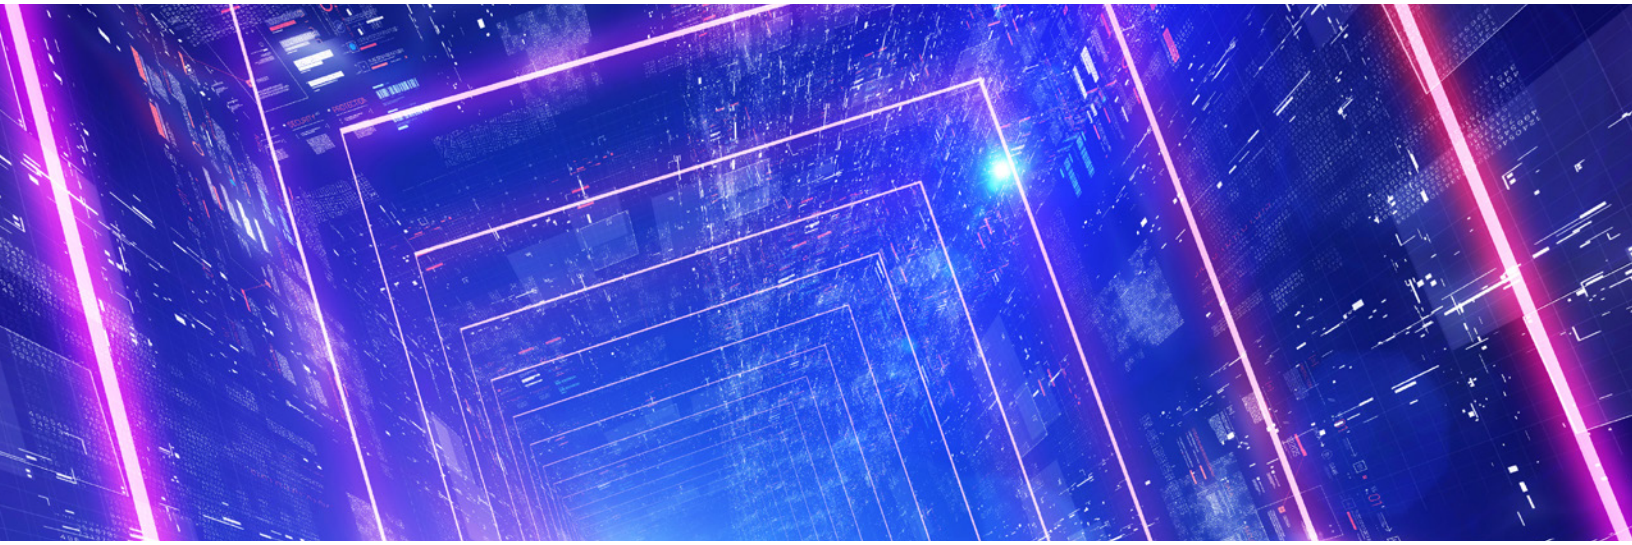

WHAT ARE "DANGEROUS ASSUMPTIONS?"

To understand the full value of adversarial testing, it is first necessary to understand the nature of the assumptions that adversarial testing evaluates —and why some of those assumptions can be dangerous.

Let's start by making it clear that these assumptions in no way suggest that SecOps teams are not doing their jobs. Just the opposite is true. SecOps teams are often so busy doing their jobs with extreme effort and diligence that they simply don't have the time to test every assumption they make. Indeed, all of us have to make assumptions every day in order to complete the most urgent items on our task lists.

When you wake up in the morning feeling fine, for example, you don't go to the emergency room at the hospital — because you naturally assume that there is nothing terribly wrong with you. That's why we get check-ups at the doctor: to test our assumptions about our health.

And that's a good thing, since doctors can tell us about a health problem that we would have otherwise failed to detect—with potentially disastrous consequences. Does this mean you were foolish not to realize your hematocrit was low? Of course not. That's not your job. Your job is to exercise, eat right, and get enough sleep. And to regularly get checkups to test your assumptions.

Also, assumptions do not have to be egregiously wrong to be potentially serious. A doctor may find something small that – upon investigation – could actually turn out to be very worrisome. One test doctors commonly use is a stress test.

In a stress test, doctors create conditions that challenge your body in order to test its response to that challenge. These principles hold true in cybersecurity as well. Assumptions don't have to be egregiously wrong to have potentially serious implications. After all, threat actors don't need a huge opening to achieve their malicious goals -- they just need an opening.

Here are some untested assumptions that SecOps teams commonly make regarding their organizations' environments:

- We patch all CVEs across all internet-facing systems.
- Our operational systems are fully segmented from our internet-facing systems.
- Our employees are all using strong passwords and multi-factor authentication.
- Our Data Loss Prevention solution will prevent insiders and outsiders from exfiltrating code.
- We can detect malicious PowerShell activity on any endpoint within ten minutes.
- No unauthorized outsider can get physical access to our physical Ethernet.
- Admin-level privileges are sufficiently segregated so that no single compromise will enable a threat actor to move laterally to our core application servers.

- Our cloud provider's security is sufficient—and a compromise in their cloud wouldn't enable a threat actor to breach our on-premise environment anyway.
- Our employees are sufficiently trained to minimize the possibility of social engineering.

It is well beyond the resources of most internal SecOps teams to test these assumptions while still doing their critical day-to-day work of shoring up defenses, securing adds and changes to the enterprise environment, responding to alerts, and proactively threat hunting.

This is why effective adversarial testing is so important.

ADVERSARIAL TESTING PROTIP #1:

Try to be specific about which assumptions most concern you. Or, to put it another way, which potential consequences of a falsifiable assumption are most likely to have disastrous consequences for your organization?

THE ROLE OF THE ADVERSARY

In an effective adversarial testing engagement, the offense actually has a three-fold mission:

Testing your defensive assumptions.
You perform an adversarial test of your cyber defenses to find out if what you believe to be true is actually true. Think you've done a sufficiently complete job of patching your internet-facing systems? Let a skilled hacker start probing for known weaknesses. If they don't find any, your assumption will be proven correct—and you will be able to move forward with confidence that your vulnerability management practices are working as they should.

Conversely and more commonly, your chosen adversary will find an exploitable vulnerability somewhere in your environment. That discovery serves a much greater purpose than a mere "Gotcha!" Instead, it will reveal exactly where your unpatched vulnerabilities are. You can then use that insight to improve your vulnerability management processes in appropriately targeted ways.

Testing their own offensive assumptions.
Threat actors adopt tactics that yield results with reasonable frequency within a reasonable amount of time and effort. Ransomware attackers in particular tend to be opportunists who will move on to other targets if their preferred methods of attack don't yield quick, easy results.

In their mimicry of criminal attackers, skilled adversarial testers will duplicate those same attack methods. That is, they will also start by making assumptions about your environment based on the common vulnerabilities they've learned about through a combination of up-to-date threat intelligence research and their own first-hand experiences.

These offensive assumptions can include:

- You haven't yet patched at least one of the CVEs that most recently came to their attention.
- They can get away with brute-force password cracking activity without being noticed.
- You cannot detect kerberoasting attacks with your current tooling.
- They can bypass your MFA with social engineering tactics.
- They can use admin-level credentials they stole from one system to gain admin-level access to another.

Testing these assumptions is important because they represent the most likely methods of an opportunistic threat actor. If your environment is well-defended against these most common tactics-of-choice, most threat actors will simply move on to another target.

Creative improvisation.
Skilled adversarial testers won't just test your existing assumption and theirs. They'll also employ that same creativity as an experienced malicious attacker — which means that they'll try to discover pathways into your environment that neither you nor they may have assumed were available before.

Examples of such creative improvisation abound. But here are a few taken from the experiences of Secureworks® Adversary Group (SwAG):

- Finding an ATM connected to the network with an Ethernet port that could be used surreptitiously by an intruder with a laptop.
- Compromising a user's Office365 account only to find that they had a file where they had all their other passwords stored in clear text.
- Discovering a tiny bit of not-yet-secure secured developer code negligently left active in a production environment, even though it was obviously intended to be there temporarily.

With the right skills and the right tooling, an expert adversarial team can get into almost any commercial environment on the planet. But, again, the objective of the exercise is not to simply pass or fail. It's to see exactly where an organization's defenses are strong — and exactly where they could use some improvement.

THE IMPORTANCE OF ITERATIVE COLLABORATION

Many enterprise cybersecurity leaders mistakenly think of adversarial testing as an isolated event followed by unilateral action. So they arrange for an adversarial test and then start acting on the results of that test as their schedules allow, given that they have a lot on their plates already.

This common approach is not entirely without value because you can certainly discover and remediate previously unknown problems in your cyber defense if you use adversarial testing this way.

But one-off, non-collaborative engagements are the least effective approach to adversarial testing. Musicians don't just practice once before they go on a concert tour. And sports teams don't just run one practice session before the season starts. In fact, musicians practice constantly. And coaches constantly run new drills in practice throughout the season to shore up the weaknesses they keep discovering in their team's performance.

The same principle holds true in cybersecurity. Adversarial testing is analogous to the drills coaches run when they have their players play against each other. These drills give everyone the opportunity to see how well certain offenses work against a given defense — and how well certain defenses work against a given offense.

That's why championship teams aren't always the ones with the biggest payrolls or the most raw talent. They're the teams led by coaches who excel at continuous improvement through ever-evolving practice.

Why iterate?
Here's the truth: No one does well on their first adversarial test. The SwAG team almost always finds several problems—and the severity of these problems generally demand further exploration for lesser issues.

When organizations follow the process model and best practices outlined here, they always do much better on the next test. And because they have far fewer problems, our adversarial testers can uncover other issues that—while certainly not as severe — could potentially result in a breach that leads to a serious compromise.

Does this mean that you should plan on endlessly iterating large-scale adversarial testing forever? Absolutely not. Just the opposite, in fact. If you apply adversarial testing best practices, you'll likely achieve a level of security maturity within 4-6 iterations that will allow you to apply adversarial testing in a much more focused and efficient manner going forward.

But no one gets there in one shot. And no one gets there by simply "fixing things." As in any scientific endeavor, you can't just implement a fix and assume that you've solved your problem. You have to test the assumption that you've solved the problem under valid experimental conditions.

If you have truly fixed the problem, great. You can move on to the next step. If not, you have to go back to the drawing board and get it right.

COLLABORATING ACROSS THE TESTING CYCLE

To get maximum business value from your investments in adversarial testing, it is essential to collaborate across the testing cycle. So make sure you:

Collaboratively define your experiment.
As noted earlier, you and your team should prepare for adversarial testing by first defining the set of cybersecurity assumptions you wish to test. You should also collaboratively determine what the parameters of the test will be. If you're evaluating your web application security, for example, it makes sense to collectively decide whether your web application firewall will be included in the scope of your test. It may also be necessary to let the attack team know that they must stop short of actually touching a database subject to regulatory reporting requirements — and that they instead only need to show that they gained a credential necessary for that access.

Also, to avoid unproductive debates after the fact, you should make an effort to get everybody on the same page about the data that will be used to evaluate the results of the experiment before you start it.

A good collaborative adversarial testing partner will provide you with a complete checklist of everything they believe you need to do in advance to make your experiment a productive one.

Collaboratively perform your experiment.
Once you've determined the size and shape of the experiment, run it. Don't get distracted as many organizations do by a lot of mid-experiment evaluation. Just complete the experiment and collect the required data.

Even though you're trying to reproduce real-world attacks and behaviors, communication between attacker and defenders may be critical at times. This is especially true if the attack team is concerned about taking a step that may adversely impact the business in real time. If the attack team compromises a senior manager's email account, for example, they should probably communicate with the appropriate manager to determine whether to continue along that vector — or spend their time probing elsewhere.

Collaboratively evaluate your experiment.
Once your experiment is complete, it's time to review the results. Everyone involved should openly share information with each other during this phase, despite the temptation not to show all their cards. Attack teams, for example, have a tendency to closely guard the methods they used to evade EDR — while defensive teams may want to avoid letting their adversary know exactly how they managed to detect a malicious implant.

While the desire to maintain an advantage is understandable, concealment at this stage is not productive. When people share everything that happened on their side during an experiment, it helps their opponent get better. And a better opponent eventually makes you better, as well.

Every chess player knows this. Every cybersecurity professional should know it too.

Also, during this step no one from either team should point fingers, make excuses, or boast about winning. Everyone is working together to learn how to make your environment more secure. They're doing so by assuming adversarial roles—but the process is ultimately a collaboration of peers pursuing a common goal.

Collaboratively define, prioritize, and pursue goals.
Once everyone agrees on what the experiment has revealed, you can chart a course of appropriate action. One set of actions will obviously be to shore up your cyber defenses wherever they were exposed to be lacking. However, it's best to define those goals as granularly as possible. For example, instead of just stating that your goal is to "***improve our MFA implementation,***" you're probably better off deciding that "we will enable number-matching to protect ourselves against users blindly accepting MFA push notifications."

Given that your staff has limited time and resources, you'll want to prioritize the pursuit of your goals based on their respective return on investment — or, as we like to call it at Secureworks, "bang reduction for the buck." If you address a potentially large risk with minimal effort, do that first. If it will take a lot of effort to fix a problem that doesn't create a major opportunity for threat actors, do it later.

And it's a good idea to do this prioritization collaboratively, because your attack team will often have the best insight into which defensive issues provided them with the most obvious pathway to a major business compromise.

Your attack team also needs to set goals based on the results of the experiment. For example, if a piece of malicious code they implanted on an endpoint got picked up by your XDR, they may want to work on finding a better way to hide on your endpoints the next time around.

After you work your way through this testing cycle, get ready to do it again. Your defensive assumptions for your next testing cycle should include the issues you worked on following the previous testing cycle. In other words, you're testing the assumption that "This solution we implemented achieves its goal." You'll also add any new assumptions that you need to validate.

Your offensive assumptions will change as well. Some of them will be along the lines of "Tactic X worked last time, so let's see if it works again." Others will be more like "We know the defense has implemented Measures A and B to protect themselves from Tactic Y, so let's see if we can use Tactic Z." Your offense should also try to take advantage of any brand-new published vulnerabilities to test the assumption that the lag-time in your organization's patching is still leaving you exposed to significant cyber risk.

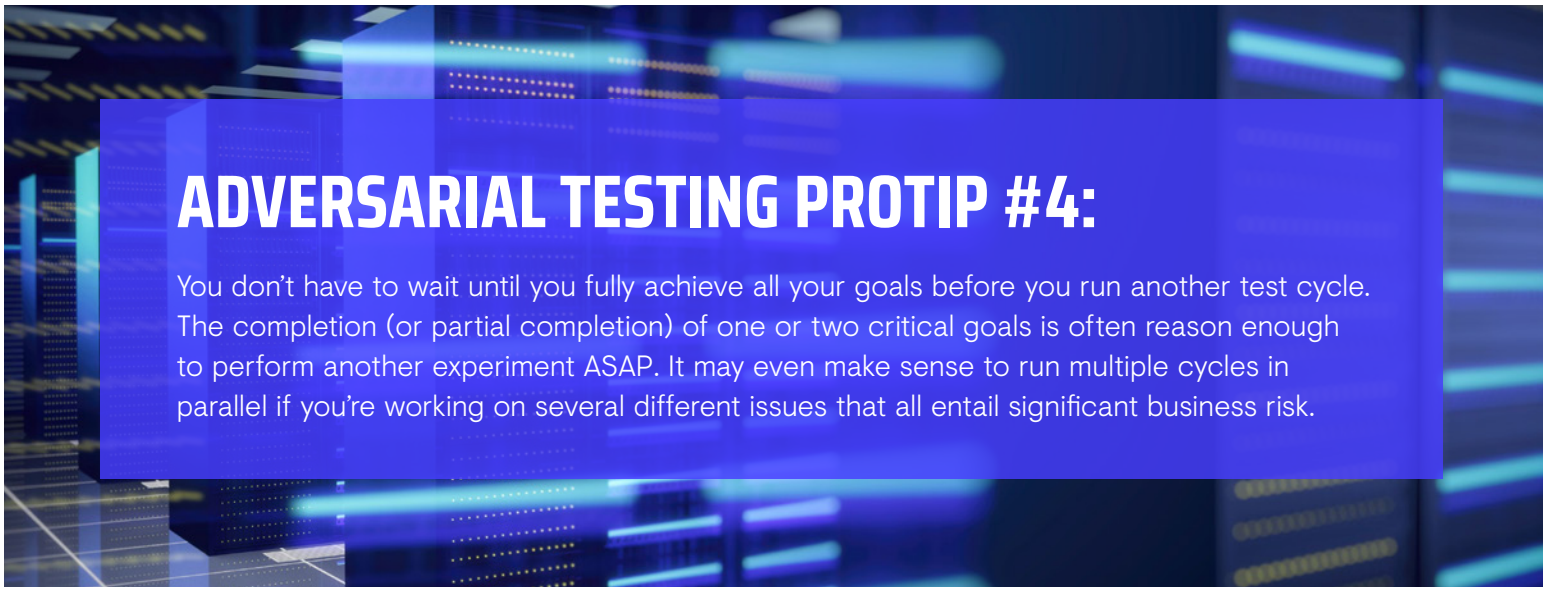

ADVERSARIAL TESTING PROTIP #4:

You don't have to wait until you fully achieve all your goals before you run another test cycle. The completion (or partial completion) of one or two critical goals is often reason enough to perform another experiment ASAP. It may even make sense to run multiple cycles in parallel if you're working on several different issues that all entail significant business risk.

THE BENEFITS OF BEST-PRACTICES ADVERSARIAL TESTING

Some organizations don't make sufficient use of adversarial testing as part of their cybersecurity operations. Others apply adversarial testing in a very narrow, non-collaborative way.

Both types of organizations miss the multiple benefits achieved by organizations that implement iterative, collaborative adversarial testing best practices. These benefits include:

- **Significantly reduced risk.** The more scientifically and aggressively you test your cybersecurity assumptions, the less likely you are to suffer a successful attack by a threat actor. The safest organizations are the ones that are constantly performing experiments to see just how safe they really are — or really aren't.
- **Smarter technology decisions.** A lot of security leaders think they're safe because of the technologies they've put in place. But no pre-purchase technology evaluation in a test bed can match a post-purchase technology evaluation in your production environment. With adversarial testing, you can discover if your vendors' claim hold water. You can also go to your vendors and show them your test results to see if they have any solid advice about how you can make better use of the technology they've sold you.
- **Improved SecOps skills.** Adversarial testing provides your team with unequaled hands-on training. They get to deal with real-life attacks in the real-life production environment they are charged with defending. And they get to do a postmortem with the attackers. No classroom or lab experience can match that.

- **Fact-based budget allocation.** Security budgets are growing, but they're not infinite. So you really want to apply every dollar where it will do the most good. By highlighting where your defenses are weakest, adversarial testing helps you pinpoint where to best spend your money — whether it's on new technologies, more staff, more training, or more user education.
- **Better cyberinsurance economics.** Organizations that can demonstrate better SecOps practices and greater cybersecurity maturity can negotiate better terms for their cyberinsurance coverage. In some cases, adversarial testing may even be a requirement for certain types of policies.
- **Mitigation of regulatory consequences.** If you operate in a regulated field such as healthcare or finance, you face all kinds of potentially punitive consequences in the event of a breach. Regulators take a dim view of victims of cyberattacks who cannot demonstrate an appropriate level of diligence around the care and control of their customers' data. Taking proactive measures, such as engaging in regular adversarial testing, shows a level of commitment to protecting that data.

ADVERSARIAL TESTING PROTIP #5:

If you're having trouble getting the budget you believe you genuinely need, think about showing a simplified version of your test results to the executives who hold the purse-strings. Those results can serve as empirical/impartial third-party validation that your organization needs to make a set of very specific additional investments in cybersecurity.

The bottom line: When it comes to your organization's security, you need to know what's wrong and what's right. And the best source of that knowledge is iterative adversarial testing against the most skilled adversary you can find. So don't put off adversarial testing any longer. Get started today!

ABOUT SECUREWORKS

Secureworks protects organizations with battle-tested, best-in-class cybersecurity solutions that reduce risk, optimize IT and security investments, and fill security talent gaps. We deliver solutions by security experts for security experts to prevent, detect, and respond to continuously evolving and diversifying threats.

Secureworks products are built on the Secureworks Taegis™ cloud-native security platform that continuously gathers and interprets telemetry from proprietary and third party sources. We use this telemetry to detect and prevent threats, automatically prioritizing the most serious ones. This enables faster, more confident response with time and cost-saving automation. Through active incidents, adversarial testing, and ongoing threat research, we continuously study, learn, and analyze our adversaries' behavior. These insights, coupled with advanced technologies, form the basis of Taegis. Learn more about Taegis at: secureworks.com/taegis

LEARN MORE

Learn how the **Secureworks Adversarial Group** can reduce your organization's exposure to ransomware and other risks through the science of cybersecurity testing,

For more information, call **1-877-838-7947** to speak to a Secureworks security specialist.

TYPES OF ADVERSARIAL TESTING

There are many different ways to perform adversarial testing — which can vary based on scope, attack types, and duration. Broadly speaking, however, adversarial testing can be split into two main categories:

Red Team

Under the Red Team testing model, your faux attackers will operate without advance knowledge of your defenders (aka the Blue Team). This model is intended for high levels of cyber maturity and tests your defenders' capabilities under real-world conditions — especially when it comes to their ability to detect and neutralize live threats in your environment as they happen.

Purple Team

Under the Purple Team testing model, attackers and defenders — the Red and Blue teams — work together continuously throughout the exercise to pinpoint issues and implement fixes. This model is used as an organization gains greater cybersecurity maturity, has completed a few penetration tests, and begins to zero in on more specific issues — such as tuning the parameters of its defenses against business email compromise (BEC) or testing how well a new cloud vendor's telemetry integrates into its XDR platform.

Secureworks®

Secureworks® (NASDAQ: SCWX) is a global cybersecurity leader that protects customer progress with Secureworks® Taegis™, a cloud-native security analytics platform built on 20+ years of real-world threat intelligence and research, improving customers' ability to detect advanced threats, streamline and collaborate on investigations, and automate the right actions.

CORPORATE HEADQUARTERS

United States
1 Concourse Pkwy NE #500
Atlanta, GA 30328
+1 877 838 7947
www.secureworks.com

EUROPE & MIDDLE EAST

France
8 avenue du Stade de France
93218 Saint Denis Cedex
+33 1 80 60 20 00

Germany
Main Airport Center,
Unterschweinstiege 10 60549
Frankfurt am Main Germany
069/9792-0

United Kingdom
One Creechurch Place,
1 Creechurch Ln
London EC3A 5AY
United Kingdom
+44(0)207 892 1000

1 Tanfield
Edinburgh EH3 5DA
United Kingdom
+44(0)131 260 3040

United Arab Emirates
Building 15, Dubai Internet City
Dubai, UAE PO Box 500111
00971 4 420 7000

ASIA PACIFIC

Australia
Building 3, 14 Aquatic Drive
Frenchs Forest, Sydney NSW
Australia 2086
1800 737 817

Japan
Otemachi One Tower 17F
2-1 Otemachi 1-chome,
Chiyoda-ku
Tokyo 100-8159, Japan
81-3-4400-9373
www.secureworks.jp