SPONSORED CONTENT | WHITE PAPER

The Shortcomings of Traditional Penetration Tests —and How Autonomous Pentesting Addresses Them

CSO

SPONSORED BY

HORIZON3.ai

~~TRUST BUT~~ VERIFY

With Research and Analysis by IDC

PENETRATION TESTING is an important part of any company's cybersecurity strategy, but most companies conduct them in a way that suffers from numerous flaws. That's the bad news. The good news is that a better, more automated, and faster approach to pentesting has arrived.

Among the flaws in the traditional approach to pentests is that they are conducted too infrequently and vary greatly in terms of quality, depending on the experience of the pentester. They often lack sufficient breadth and depth to provide real assurance that no security holes remain undetected, and they take a long time to complete.

Penetration testing

To ensure proper security, companies need a more automated and autonomous approach to pentesting, one in which they can run the tests as frequently as they choose, on their own. Such an approach would ensure that companies can run pentests that cover their entire attack surface, including internal, external, and cloud-based systems—just as a real attacker would. They would also get the results immediately, enabling them to address any issues instead of having to wait weeks as with traditional tests—leaving vulnerabilities exposed in the meantime.

Horizon3 is one company that is delivering automated, autonomous pentesting with NodeZero, its software-as-a-service (SaaS) pentesting platform. NodeZero democratizes pentesting, enabling any company to run pentests whenever and as often as they want.

Elements of effective security

Pentesting, of course, is just one element of an effective security strategy.

Proper security starts with tools that alert and log security events. A security information and event management (SIEM) tool that combines events from various tools is also helpful, serving as a central collection point for all alerts and

> **Given everything that goes into a traditional pentest**, it's little wonder that most companies don't conduct them with nearly enough frequency — **typically just once or twice per year.**

helping with alert correlation. Even so, a security professional is still required to assess the alerts and weed out the false positives from those that require attention.

Properly secured endpoints are also crucial. Endpoint security involves ensuring that all devices that connect to your network—including desktops, servers, and mobile devices such as laptops and phones—are properly secured. That means having antivirus and antimalware software in place, along with a firewall, encryption, virtual private network, identity and access management, and potentially more safeguards. Internet of Things (IoT) devices should also be protected, as they represent a growing attack surface.

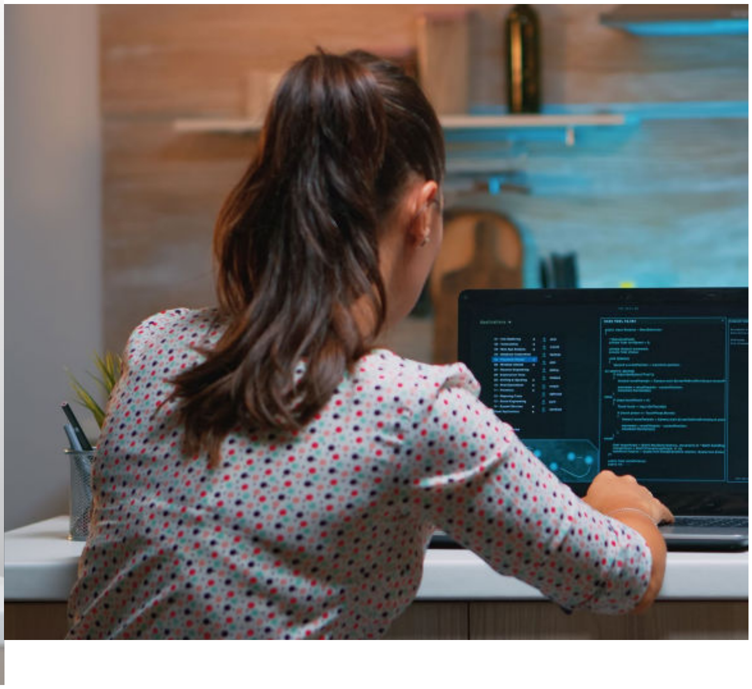

Having an effective password policy is likewise essential. Multifactor authentication is becoming table stakes in terms of identity protection. In-depth, frequent education about phishing is likewise a good idea, including in-house "white hat" attacks to identify employees who fall for phishing attacks.

Vulnerability scanning is not pentesting

Vulnerability scanning is another component of identifying potential vulnerabilities in your environment—but it is not at all equivalent to pentesting. A vulnerability scan is concerned with just one thing: finding potential vulnerabilities. The key word there is *potential*, because vulnerabilities are not all created equal. The crucial criterion is whether it's a vulnerability an attacker is likely to exploit.

Only a quality pentest can determine that. A good pentest, for example, checks for common misconfigurations and conducts credential attacks to see if a particular vulnerability may actually result in a breach. Credentials are one of the most-sought-after data types by hackers. They are involved in about 60% of all attacks, including ransomware attacks, according to

Verizon's "2021 Data Breach Investigations Report." In fact, 60% of ransomware cases involve the direct installation of desktop sharing apps, typically the result of either stolen credentials or brute force attacks, Verizon reports.

Additionally, a good pentest goes well beyond simply finding and reporting on vulnerabilities. Rather, the tester will attempt to use a vulnerability to gain access to a host, extract credentials housed on the host, and use them to gain access elsewhere in the network. The idea is to see whether a single vulnerability can lead to the compromise of the entire network, including sensitive business systems.

Routine pentesting of this type is essential to verify that all elements of your security strategy are working in concert—and to find the areas that need improvement.

Traditional pentesting

The "routine" part is key if you're going to keep up with all the latest threats. The problem is that traditional pentesting doesn't lend itself to routine testing, because it's a highly manual, time-consuming process.

After you've found a good pentester (which is no mean feat, as you'll see shortly), you've got to prepare your environment for the pentest. That means pulling together all relevant technical contacts the tester may need to communicate with, whether before, during, or after the test. You also need to inform key IT personnel about the test, so they don't think your company is facing a real attack when the test starts. Additionally, pentests do sometimes cause issues in the IT environment, so you need IT personnel to be on standby. You need to clearly define the scope of the test and ensure that the tester has appropriate permissions to conduct the tests.

Finally, you can onboard the pentester and let that person get to work. But don't hold your breath for the results, which typically take four to six weeks to be delivered.

At that point, you need to allot time to address any issues the test uncovered. This is also an opportunity to assess whether the tester directed you to the most critical issues, vs. low-level vulnerabilities that pose little to no actual threat.

Once you've addressed all the issues the test uncovered, you may need to conduct remediation testing if the penetration tester needs to validate fixes. This involves the same lengthy process as the initial test.

Problems with traditional pentests

Given everything that goes into a traditional pentest, it's little wonder that most companies don't conduct them with nearly enough frequency—typically just once or twice per year.

That's hardly enough to keep up with the ever-changing threat landscape and expanding attack surfaces. Consider that Google's Project Zero, which provides information on zero-day threats, identified more than 55 zero-day attacks that were exploited in 2021 by mid-November. By contrast, it identified 26 zero-days exploited in all of 2020, roughly the same as in each year since the project's inception, in 2014.

It's important to note that the actual number of vulnerabilities is not necessarily increasing; rather, it's that vulnerabilities are being exploited at a far greater rate. MIT Technology Review posits a reason for that:

"It's easier than ever to buy zero-days from the growing exploit industry. What was once prohibitively expensive and high-end is now more widely accessible. And cybercriminals, too, have used zero-day attacks to make money in recent years, finding flaws in software that allow them to run valuable ransomware schemes."

In order to ensure that companies stay safe in the face of cyberthreats, it's time to **fundamentally rethink the way we conduct penetration tests.**

Given all this activity, it would be easy to miss critical issues between pentests conducted only a couple of times per year.

What's more, traditional pentests are often driven by a need to stay in compliance with regulations and standards, such as SOC (Service Organization Controls) 2. In such a case, the pentest may reflect more of a "check the box" exercise rather than a real desire to ensure proper security.

Manual pentests also by nature restrict breadth of coverage. They are conducted in a constrained timeframe (also usually four to six weeks). In a large-scope environment, that limited engagement time inherently reduces the amount of attack surface a tester can cover.

No matter how experienced the pentesters may be, no tester can cover a large network in a short time. As a result, they may discover some attack paths but miss others simply due to time constraints.

On the other hand, four to six weeks is a long time to wait for results, when you consider that you likely have vulnerabilities that need to be addressed. The sooner you find out about and remediate them, the better.

Another issue with traditional pentests is that they are typically siloed into different categories of IT systems, such as internal, external, cloud-based, and IoT. But attackers don't think like that. Rather, they will employ any means possible to break into a network. If attackers see an opportunity to target an external system at an employee's home to steal credentials and then use those credentials to access an internal system, that's what they'll do.

Finally, quality can vary dramatically with traditional pentesting, depending entirely on the individual you hire. Some pentesting companies do little more than run a vulnerability scan and report the results. They don't try to think like a hacker and determine which vulnerabilities a hacker could reasonably exploit, vs. those that are relatively benign. And different testers may or may not show you how an attack was performed, which is important, because it proves that the vulnerability is real.

A better way: autonomous pentesting

In order to ensure that companies stay safe in the face of cyberthreats, it's time to fundamentally rethink the way we conduct penetration tests.

Today it's possible to capture the in-depth knowledge, insights, and practices of the best pentesters and model them in software. Similarly, we can model the behaviors of attackers in software, capturing the various paths they take to break into corporate networks. The result is a software-as-a-service (SaaS) platform that performs automated, autonomous penetration tests on demand.

Autonomous pentesting brings numerous benefits over traditional testing.

For one, you can run pentests as often as you want, even multiple times per week. It's a self-service platform, with little overhead for IT groups. Now there's no need to wonder if you're susceptible to whatever the latest threat is. Run a pentest and find out.

Automated SaaS-based pentesting also means that you can test orders of magnitude more assets than a human tester could in the same amount of time. It's perfectly feasible to test thousands of assets per week, vs. the four to six weeks it would take to conduct the same tests manually.

The ability to test more assets more quickly means that you can test across your entire attack surface: internal premises-based systems, external systems located in employee homes and regional offices, cloud-based infrastructure, and IoT systems.

Importantly, because they capture the knowledge of the best pentesters, SaaS-based tests look at the environment holistically, just as attackers do. The tests will find vulnerabilities that are actually exploitable, no matter where they are in your environment.

And there's no waiting four to six weeks for test results. Rather, you get results immediately, with clear reporting on all exploit chains. That includes proof of exploitation, meaning that you can see exactly how the attack was carried out, the exact devices that were compromised, and a determination of the resulting business impact. Reporting also includes an audit trail of the actions taken, so you can follow the attacker's footsteps.

This kind of SaaS-based testing democratizes pentesting, making the knowledge of the best pentesters available to anyone. Now network or systems administrators as well as blue team members can run tests on their own, whenever they see fit.

Horizon3 delivers with NodeZero

It's just this kind of advanced pentesting that Horizon3 offers with NodeZero, its SaaS-based offering that makes continuous autonomous pentesting available to any company.

With the ability to run pentests at any time, NodeZero enables your cybersecurity

> With the ability to run pentests at any time, **NodeZero enables your cybersecurity team to proactively find and fix attack vectors before attackers can exploit them.**

team to proactively find and fix attack vectors before attackers can exploit them. Frequent pentesting also means that you can verify the effectiveness of your security tools and address any holes.

NodeZero prioritizes the vulnerabilities it finds, based on risk level and how difficult each one is to address, so you can quickly zero in on the most serious alerts along with those that are relatively simple to address. Once you've remediated, you can test again right away to verify whether the steps you took were effective.

Detailed reporting gives you a picture of your current security posture and how it has improved over time. That's valuable information for corporate boards as well as regulators to prove the effectiveness of your cybersecurity program.

And NodeZero itself gets more effective over time, because as more vulnerabilities are identified, they are added to the system, creating an ever-increasing knowledge base of known vulnerabilities. ◆

Learn more about a better way to conduct penetration testing. Visit https://www.horizon3.ai.

Autonomous Penetration Testing as a Service: Offense Is the Best Defense

Autonomous penetration testing as a service is a new SaaS security capability. This IDC Analyst Brief intends to raise awareness of the solution and clarify how it differs from legacy pentesting services.

WRITTEN BY: PHILIP D. HARRIS, RESEARCH DIRECTOR,
WORLDWIDE CYBERSECURITY RISK MANAGEMENT SERVICES

APTaaS

Introduction

Businesses need reliable security controls more than ever to combat dynamic threats, protect their perimeter-less environments, stop malicious activity, and manage risk. Given all the variables in play, IT and security teams are hard pressed to keep up with relentless attackers. Autonomous penetration testing as a service (APTaaS) increases an organization's ability to protect itself by giving security teams new tools to help drive out threats immediately, boosting the effectiveness of cybersecurity teams by focusing on the right things.

APTaaS is a new solution that promises to usher in a future of continuous and proactive security assessment. As organizations invest in security solutions over time, many of these solutions may fail to work well together or to integrate easily into existing systems. APTaaS acts as a "sparring" partner against these solutions by automating penetration testing (pentesting) activities to expose critical vulnerabilities before an attacker can find them. Cybersecurity teams can swiftly spot and remediate vulnerabilities, resolve integration gaps, and avoid costly disruptions that may disaffect customers and tarnish brand reputation. For example, APTaaS was able to fully compromise a bank within seven minutes by leveraging vulnerabilities not identified by existing security solutions.

Unlike legacy point-in-time pentesting that covers only a small percentage of an attack surface, APTaaS offers comprehensive visibility, scope, and scale to match the threat environment, bringing context-awareness of and focus on the most urgent and impactful issues. Over time, APTaaS has the potential to displace a portion of legacy pentesting to:

- Provide offensive capabilities to organizations either without dedicated security teams or with defensive-only teams
- Be a force multiplier for organizations with red teams
- Validate service-level agreements with managed security service providers

What Is APTaaS?

APTaaS is a subscription service that runs proactively, discovering and exposing exploitable vulnerabilities. This empowers cybersecurity teams to focus on and fix the most urgent internal and external attack vectors before a real attacker exploits them.

APTaaS essentially becomes the "purple team" (see Table 1) within an organization. Blue teams and red teams focus on defense and attack separately and win only by defeating each other using point-in-time scenarios. The purple team approach unifies both blue and red teaming, producing more effective results by continuously scanning throughout the IT environment to expose high-impact vulnerabilities that must be remediated.

TABLE 1: Purple Teaming Offers a Distinct Advantage

What Matters	Traditional Approach	Purple Team Approach
Effort Required	**High** (multiteam, coordinated)	**Low** (self-service, on demand)
Test Frequency	**Annual or quarterly**	Agile and **continuous**
Total Cost	**High** for single pentest	**Low** for unlimited, ongoing pentests
Time to Value	**Weeks** to receive written report to act on	**Hours** to searchable, prioritized results
Coverage	**1%–2%** of environment	**99%+** of environment
Expertise Needed	**High** to execute	**Low** to execute
Resources	**External** professional services	**Internal** purple team partner
Ultimate Goal	**Pwn you** to demonstrate value	**Decreased risk** of being pwned at all

APTaaS functions are find (discover, leverage, and chain exploitable weaknesses to prove business impact), fix (mitigate or remediate weaknesses efficiently), verify (validate that mitigations or remediations were implemented and stay implemented), and assess continuously (compare results to find new weaknesses that have been added or fixed to address gaps). APTaaS completes these activities as a closed-loop service that allows cybersecurity teams to evaluate an environment for threat indicators and attack behaviors, unprotected assets, misconfigurations, human errors, log gaps, IT hygiene issues, and more. Armed with this information, cybersecurity teams can take the recommended actions to close gaps, fix misconfigurations, strengthen credential management, and more. In contrast, open-loop testing, such as blue team or red team exercises, involves people who find vulnerabilities; write them up; and hand them off for investigation, analysis, and remediation.

Legacy pentesting attack scenarios are mostly but not always based on playbooks designed to accomplish specific objectives, whether that is to bypass controls or uncover possible routes to critical assets at a point in time. APTaaS runs in the background, generally in production or test environments with attack patterns and scenarios that are predefined in algorithms and continuously updated by a team of experts. This approach offers the ability to take advantage of the latest attack scenarios immediately. The ability to function in a test environment can be significant when dealing with sensitive operational technology, industrial Internet of Things (IoT), or medical IoT devices that do not allow for even a hint of an invasive security test.

APTaaS testing options include on-demand, continuous, or set intervals. On-screen, real-time visualizations present a variety of findings, including vulnerabilities, severity ratings, and remediation steps. Standard and customizable reports present insights on diverse topics to audiences ranging from technical staff to executive leadership.

Why Is APTaaS Important?

Digital transformation initiatives, cloud migration, work from home, 5G, edge computing, and emerging technologies expand the attack surface, multiply the number of potential vulnerabilities, and increase the risk of breach. At the same time, adversaries use the latest technologies to launch zero-day threats and known threats using myriad tactics, techniques, and procedures (TTPs).

"Defense in depth" is a philosophy often espoused by organizations. However, it is getting the various layers of security solutions to work well together that ensures the integrity of the deployment. APTaaS helps organizations to achieve that goal and enables cybersecurity teams to focus on the key areas of remediation by thinking like an attacker.

For example, an organization's Local Security Authority Subsystem Service (LSASS) process was vulnerable, enabling exploits. The organization's firewall failed to detect this because it lacked the correct module, and the endpoint detection and response failed because it was not properly installed on certain machines. All this contributed to a failure in the defense-in-depth strategy. APTaaS uncovers these types of issues so organizations can resolve them quickly. In this instance, APTaaS was able to dump the organization's LSASS to block future exploits.

Complex security environments hinder swift and efficient detection, response,

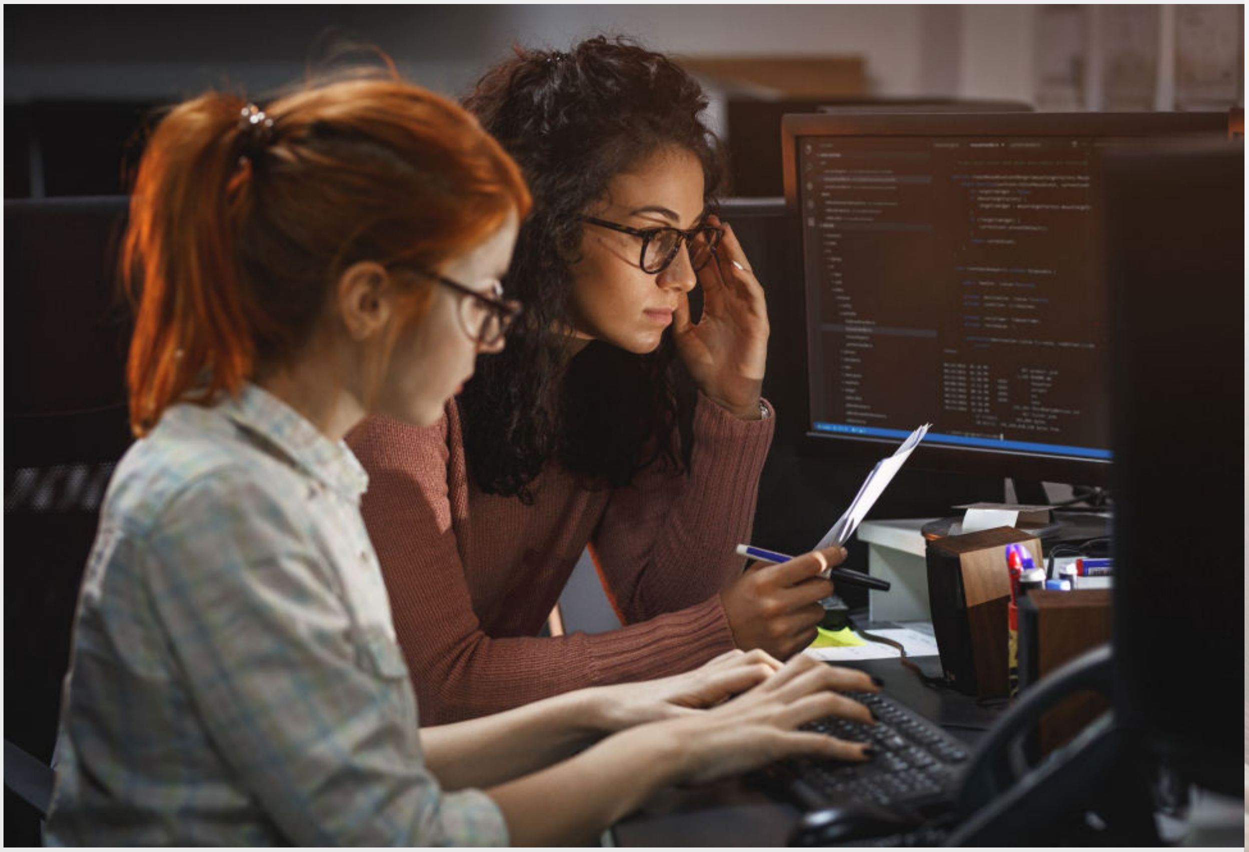

and remediation. To increase protection, some organizations purchase additional technology that may not be used fully due to time constraints, lack of skills, or scarcity of resources. Unfortunately, an excess of tools creates even more complexity, complications, and management difficulties. The challenge intensifies with an increase in tools, assets, remote workers, and legacy operating systems or hardware.

Patching offers an illustration of these issues. Cybersecurity teams may think that a system has been patched because the registry keys have been changed to reflect that status. However, the actual patching of the system can still fail, which may not be discovered by the organization until it is too late.

Without visibility into TTPs or organizational incident readiness, cybersecurity teams can lack clarity about mitigation steps, and vulnerabilities may go unmitigated.

APTaaS can strengthen organizational defenses by using artificial intelligence (AI) and machine learning (ML) or leveraging up-to-date attack frameworks and knowledge bases. APTaaS platforms deliver services to help cybersecurity teams answer questions such as:

- Are security controls working as intended?
- Are technologies leveraging default security settings, or are they poorly configured?
- Are there new attack patterns?

APTaaS offers an automated, scalable way to assess and improve an organization's security posture in an ongoing manner to thwart assailants searching for high-impact, exploitable vulnerabilities against critical assets.

By contrast, legacy pentests can be costly and labor intensive, reflecting a point in time within a tightly scoped section of the IT environment. Their effectiveness is reduced by their inability to be rapidly performed in a cost-effective and sustained way to reflect changes in cyberhygiene or changes to TTPs of cybermiscreants. Table 2 describes the differences between the various testing techniques and APTaaS.

TABLE 2: Differences Between APTaaS and Other Testing and Attack Simulations

Pentest Type	Core Functionality	Use Cases
Autonomous penetration testing as a service	Continuous pentesting service as a SaaS offering: - Requires no persistent or credentialed agents - Scopes specific IP ranges to scan or IP ranges to avoid - Intelligently identifies the scope for you - Enables or disables specific attacks	Enables understanding of the security posture across several dimensions so cybersecurity teams can: **- Find:** Proactively discovering, leveraging, and chaining exploitable weaknesses to: • Prove out the potential critical business impacts with proof-of-exploit in hand. • Understand the attack vectors leading to critical impacts to know exactly what to fix in order to disrupt the kill chain. **- Fix:** Focus on efficiently mitigating or remediating weaknesses that can actually be exploited instead of chasing down unexploitable vulnerabilities and false positives. **- Verify:** Validate that mitigations or remediations were implemented and remain implemented. **- Perpetual:** Continuously assess the security posture and quickly compare results to see what new weaknesses have been added or fixed.
Breach attack simulation	Automated testing of the existing security infrastructure; modeling attack chains to identify the most likely path an attacker would use to compromise an environment	This involves continues testing of security controls with gap remediation recommendations.
Traditional pentesting	Manual testing that is used to help test the effectiveness of an organization's vulnerability management program and associated controls within a defined scope	Test specific predefined networks, assets, platforms, hardware, or applications that are vulnerable to an attacker. Penetration tests are not focused on stealth, evasion, or the ability of the blue team to detect and respond, since the blue team is fully aware of the scope of the testing being conducted.
Red team	Designed to achieve specific goals, such as gaining access to a sensitive server or business-critical application	Emulate an advanced threat actor by using stealth, subverting established defensive controls, and identifying gaps in the organization's defensive strategy to better understand how an organization detects and responds to real-world attacks.
Blue team	The internal security team that defends against both real attackers and red teams; should be distinguished from standard security teams because of the mission to provide constant vigilance against an attack	An ongoing team of defenders may engage against known or unknown red team attack exercises. Defenders can also benefit from purple team exercises that integrate defensive tactics and controls from both the attacker and the defender teams.
Purple team	Manual, human-based exercise using real user behavior and exploits, with scenarios aligned to the organization's network to expose blind spots in security analyst response, tool efficacy, and gaps in security controls	Purple teaming aligns red and blue teams to provide an end-to-end and realistic autonomous-penetration-testing experience and prioritized vulnerabilities to the organization.

Several issues are associated with legacy penetration testing methods:

- **Scope limitation:** Typically, tests are set up on a project-by-project basis that is limited to a specific portion (range of IP addresses) of the IT environment desired by the customer.
- **Lack of context awareness:** Data exposure can be missed when testing only a portion of the environment, a select web application, or certain segments of the external network environment.
- **Existing tool sets:** Tools and services are noisy, unreliable, and a resource drain on organizations.
- **Cost:** Legacy testing methods become expensive as project scope and scale grow.
- **Point-in-time testing:** Due to the constant state of change, without ongoing testing, an organization can instantaneously fall victim to attack.
- **Expertise:** Without rigorous skills training and awareness of the latest TTPs, legacy pentesters can easily miss important test requirements.
- **Scalability:** Legacy methods do not scale well, due to time and expense needed to hire teams to conduct tests.

Benefits

APTaaS requires minimal configuration and can run in the background over the entire environment or a portion of it. Other benefits include:

- **Context:** Results are prioritized based upon criticality — a function of exploitability and impact.
- **New attack patterns:** APTaaS is continuously informed of new attacks and techniques by experts constantly updating algorithms.
- **Scalability:** No persistent agents are deployed, no additional scripting is required, humans are not involved in the day-by-day testing, and there is no limitation to environment size.
- **Efficient improvement of security posture:** Immediate resolution is ongoing through continued testing, identification, prioritization, proof, and remediation of exploitable high-impact risks and vulnerabilities.

- **Ease of manageability and maintenance:** Service manageability and maintenance are simplified and easily adjusted as needed.
- **Effectiveness and efficiency:** Cybersecurity operations teams can improve productivity and focus on critical issues rather than having to fix everything.

Considerations

Regulators need to be educated on the capabilities of APTaaS. Regulations also need to accommodate for APTaaS. Examples of such laws and regulations include medical device manufacturing rules, Health Insurance Portability and Accountability Act of 1996 (HIPAA), Payment Card Industry Data Security Standard (PCI DSS), Technology Services (SOC2), and Financial Industry Regulatory Authority (FINRA).

Conclusion

The goal of any organization is to be resistant to cyberattack. APTaaS provides an automated and scalable service that uncovers vulnerabilities, prioritizes them based upon context and impact, provides evidence of exploits, and remediates them immediately, protecting critical assets and reducing risk.

About the Analyst

Philip D. Harris, CISSP, CCSK,
Research Director, Cybersecurity Risk Management Services

Phil Harris is the Research Director for CRMS. He is responsible for developing and socializing IDC's point of view on Governance, Risk, and Compliance across people and process, focused on creating a foundation of privacy and trust with enterprises, IT suppliers, and service providers.