KALWEIT ITS

RED TEAMING

The Ultimate Discipline for Identifying Vulnerabilities in Organizations

KALWEIT ITS

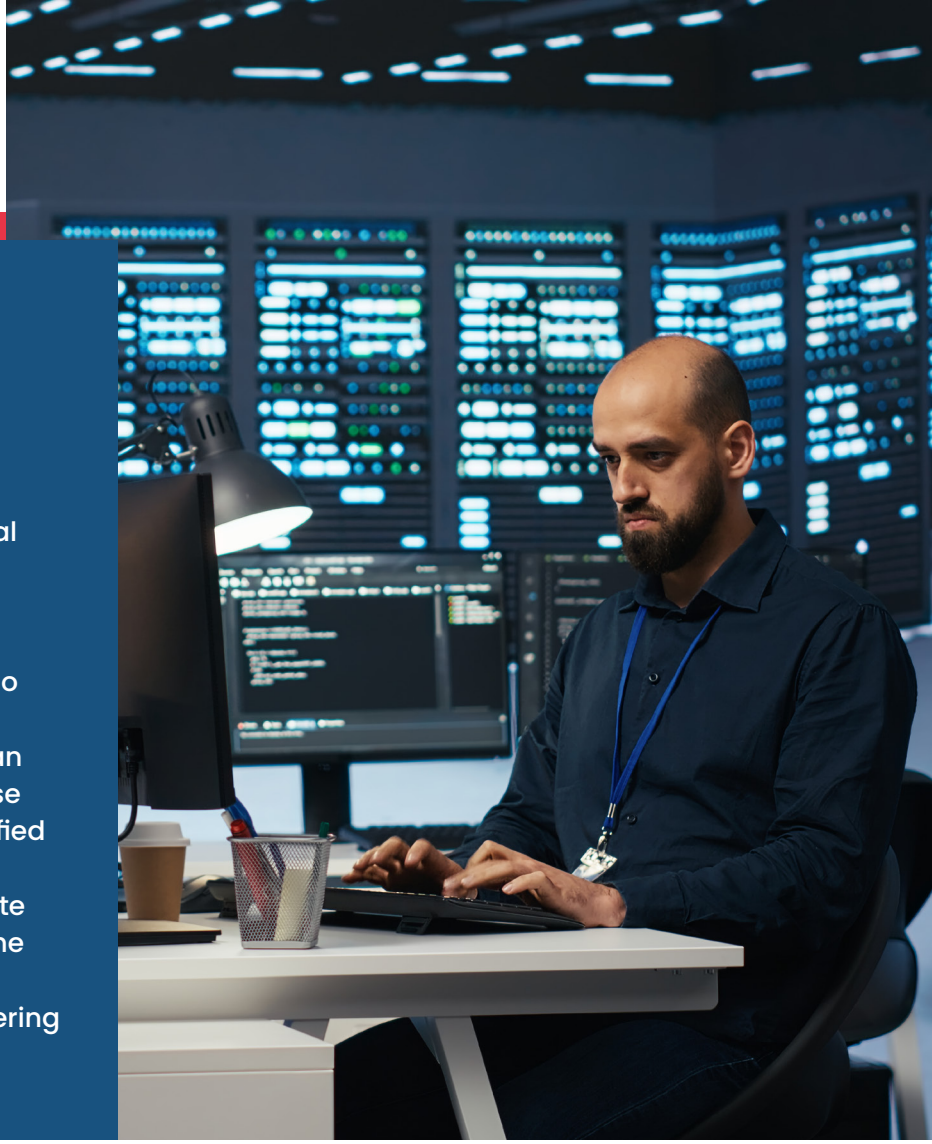

ABSTRACT

Red teaming, a realistic and holistic approach that emulates the tactics of actual attackers, is becoming essential in uncovering vulnerabilities in cyber defenses. In an era where traditional security measures such as audits and penetration testing (pentesting) are no longer sufficient, red teaming offers a more comprehensive assessment of an organization's security posture. The rise of artificial intelligence (AI) has amplified these challenges by introducing new, potent cyber threat vectors that elevate risk exposure. This paper delves into the mechanics of red teaming, providing insights into its effectiveness in countering the complexities of new cyber threats.

INTRODUCTION

The latest advances in artificial intelligence (AI) have increased cyber risk exposure for businesses and public sector organizations. Combined with increasingly sophisticated attack vectors, AI makes it more difficult than ever to identify and respond to serious cyber threats. Cybersecurity teams that rely on security audits and penetration testing (pentesting) may discover too late that these methods are not adequate for optimal risk mitigation.

Generative AI increases the attack surface

The increasing complexity and dynamics of AI systems make it difficult to identify and secure potential vulnerabilities. In addition, the secure integration and management of the technology requires specialized expertise, which is often lacking. As a result, GenAI has generally increased the attack surface at the companies surveyed.

GenAI changes IT environments and increases the **complexity** of defensive measures

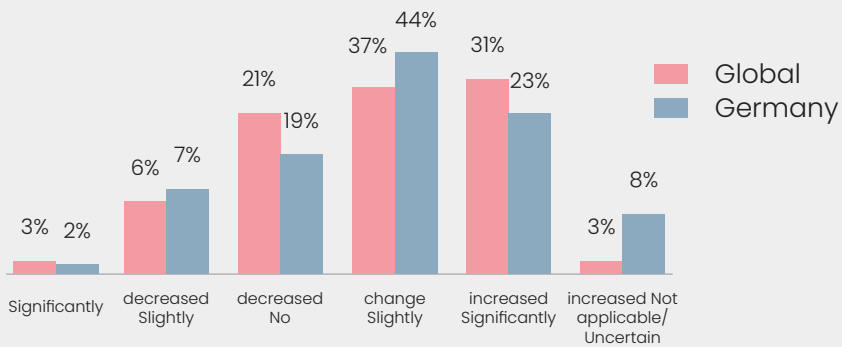

Red teaming, a realistic, holistic process that mimics how actual attackers operate, is far better suited to discovering weaknesses in cyber defenses. Red teaming is not new, but it may come as a surprise that this approach to testing controls and countermeasures has become more economical over time. It is no longer just for large enterprises. This paper explores how red teaming works, offering insights into its efficacy in this new era of AI-driven cyberthreats.

TODAY'S MAJOR CYBERSECURITY CONCERNS

It should be major news that enterprises have deep concerns about cyber security. The threat landscape never stops evolving, however, so specific worries change over time. A recent report from PWC highlights today's most pressing issues. Their research found, for example, that cloud-related threats were the most serious type of cyber incident facing businesses worldwide. Sixteen percent of the 1762 IT and security professionals surveyed by Pricewaterhouse Coopers (PWC), the global consultancy and audit firm, cited this form of threat as the most concerning. This was followed by "Hack and Leak Operations" (**15%**) and third-party breaches (**12%**). Ransomware, typically assumed to be the number one threat, actually garnered just **11%**.

Companies feel most threatened by these cyber incidents

The threat to digital infrastructures is a concern for most respondents worldwide. In Germany, the fear of hack-and-leak operations dominates, followed by cloud-related threats and attacks on networked products.

The secure integration and management of cloud technologies requires specialized expertise that is often lacking - and which, in view of the vast IT specialists are also hard to come by. Grant Waterfall, Cyber Security & Privacy Leader at PwC Germany and EMEA

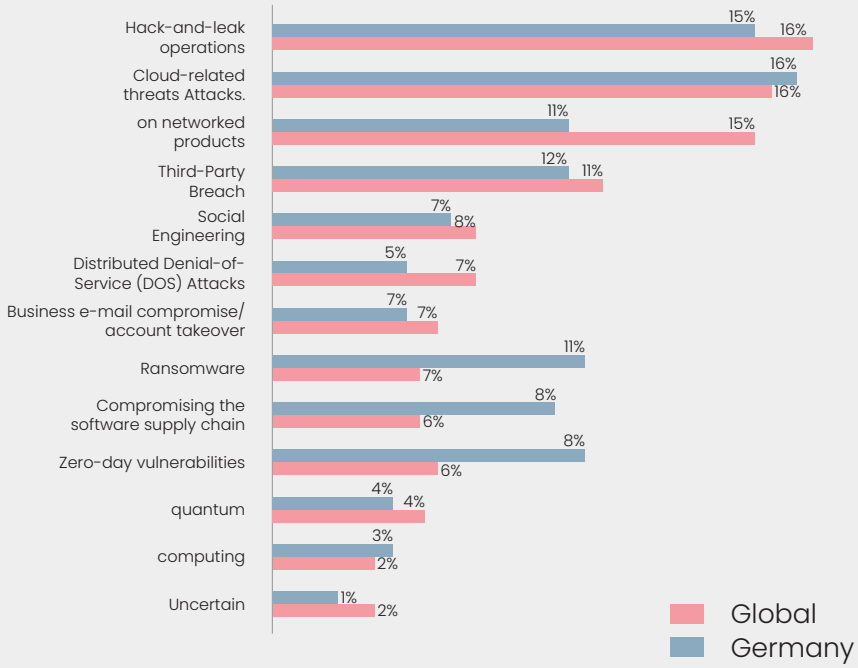

[1] Global Digital Trust Insights 2025, PWC

When it comes to areas where enterprises feel unprepared to defend their digital assets, quantum attacks came in first place, with **12%** of those surveyed saying they feel insufficiently prepared. Distributed Denial of Service (DDoS) attacks tied for the slot, also with **12%,** followed by attacks on networked products, at **11%**, and ransomware (**10%**).

Recent advances in AI are having an impact on cybersecurity, as well. According to PWC's research, **28%** of those surveyed shared that generative AI (GenAI) had either increased or "significantly increased" their attack surfaces over the last 12 months. Forty-four percent said GenAI had "slightly increased" their attack surface.

Why is GenAI such a problem, from a security perspective? There is more than one answer to this question. GenAI enables attackers to create thousands of synthetic identities, for instance, which can flood authentication tools and enable account takeovers. GenAI also enables hackers to create malware more efficiently than was previously possible.

WHAT IS RED TEAMING?

What is the aim of Red Teaming?

The aim of Red Teaming is to find vulnerabilities in IT security through a holistic approach, in which experts proceed in the same way as real attackers - i.e. testing human aspects, technical aspects and the ability to react or general emergency behavior.

- Human behavior
- Technical safety
- Behavior of your own IT in the event of an IT emergency

While today we mostly think of red teaming as a cybersecurity practice, its roots go back at least to the early 1960s. At that time, the U.S. military would engage in tabletop "war games," with a red team and a blue team facing off against each other in a theoretical military exercise. Whether you're doing red teaming to determine if you want to use nuclear weapons or detect a weakness in a computer network, the principles are the same, however.

The red team is the attacker. The blue team defends. In a cybersecurity red teaming exercise, the red team tries to break into the blue team's environment and successfully carry out an attack of some kind, e.g., exfiltrating data.

Typically, the blue team hires the red team and describes the digital assets they want to protect. For example, the blue team may say, "Our most valuable asset is our database. We need to protect it." The red team takes it from there, analyzing the blue team organization's perimeter and other cyber defenses—with an eye to penetrating them.

If the red team is successful, and they very often are, the blue team is eager to learn about the deficiencies in their countermeasures that allowed the attack to breach their defenses. For instance, the red team might use social engineering, e.g., a phishing attack, to get log in credentials to the blue team's network. This would suggest that the blue team (client) organization needs to do more to educate its users on social engineering attacks, and perhaps add an anti-phishing filter to its email, and so forth.

THE RED TEAMING PROCESS

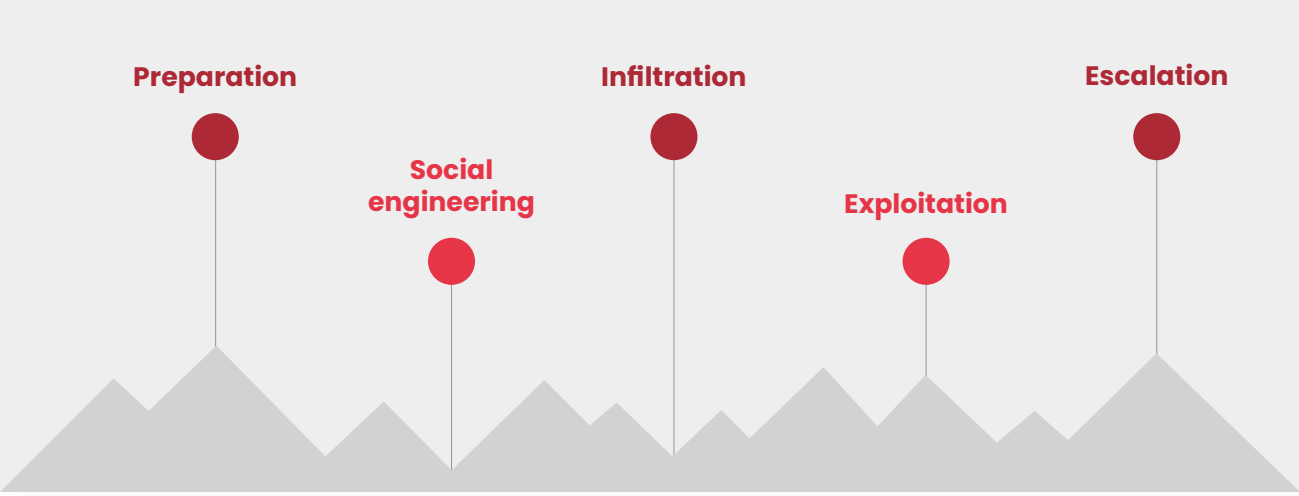

There is no single way to do red teaming. Rather, each red teaming exercise is unique, based on the client's IT estate, security controls, and so forth. That said, red teaming exercises follow a broadly similar pattern.

Most red teaming exercises begin with a conversation between the red team and the client to discuss their objectives and biggest security concerns. The client may want to understand its vulnerability to attacks that could affect its regulatory compliance, for example. Or, they might want to assess how well they can protect sensitive intellectual property, or retain the ability to function even if their network is compromised, among many examples.

The red team thinks and acts like an attacker. In particular, this means working through the following issues with the client in advance of conducting the exercise:

01

Attack path, i.e., how would an attacker approach the target?

02

Attack techniques, i.e., what methods will work best to breach the target's defenses?

03

Detection, i.e., how well will the client's security countermeasures perform in detecting the attack?

04

Reaction, i.e., will the client's defensive measures and incident response processes be effective in the event of an attack?

As the client's needs come into focus, the red team does whatever it can do to identify the client's digital assets and its cyber defenses. The best practice is to work "blind," and not ask the client to reveal how its countermeasures work. This is for realism. Real hackers don't have the advantage of knowing how a target defends itself, so neither should the red team.

The red team then goes on the attack. This process varies widely from case to case, but generally, a red team exercise involves mounting simulated attacks that probe human behavior along with technical security. Regarding the human side of red team attacks, while it's become increasingly common knowledge that people tend to the be weakest link in a cyber defense strategy, people remain a highly vulnerable attack surface.

People, from employees to contractors and employees of vendor firms, are vulnerable to spear phishing attacks, for instance. Even savvy and trained people can fall for a well-executed attack. Like, if your company's URL is acme.com, will your employees notice if they get an email from someone at acme.co? A good red team is skilled at social engineering attacks that coax network logins and other secrets, like maps of network segments, from people who think they're talking to friends and colleagues.

The red team then compiles a list of technical aspects affecting the client's security, such as:

- On the client side - security solutions and endpoint configurations
- Server side – security controls applied to servers, as well as their security configurations
- Active Directory, or equivalent – how the client manages identity, authentication, and authorization
- Network – including network security countermeasures and network topogrpaphy

This workload has grown significantly more complex as cloud computing has become the norm at most, if not all, enterprises. The cloud estate spans infrastructure-as-a-service (IaaS), software-as-a-service (SaaS), and platform-as-a-service (PaaS), along with private clouds, managed cloud hosting of enterprise applications and many hybrid modes of IT operations.

Such cloud deployments create a far larger attack surface for the red team to exploit. It is critical, however, for the client to include the cloud in their blue teaming efforts. Without taking the cloud into consideration, the client is leaving a major portion of its IT portfolio untested, from a security perspective.

At this point, it's time for the red team to get to work in earnest. This might involve simulating a DDoS attack on an endpoint, using social engineering to get login credentials, trying to get malware installed using an email attachment, and more.

The process assesses the following aspects of the client's security:

- Time to alert, i.e., measuring the time it takes the blue team to become aware of the attack (if they actually do)
- Time to react, i.e., the time that elapses between the alert notifying the blue team's security operations center (SOC) of the attack and SOC's response to the attack
- Capability to react, i.e., the scope and efficacy of the blue team's response

RED TEAMING VS. SECURITY AUDITS AND PENTESTING

Some organizations opt for security audits or pentesting as an alternative to red teaming. There are valid reasons for this, but it is not accurate to conflate red teaming and these other methods of testing the efficacy of security controls. The processes are different, as are the outcomes.

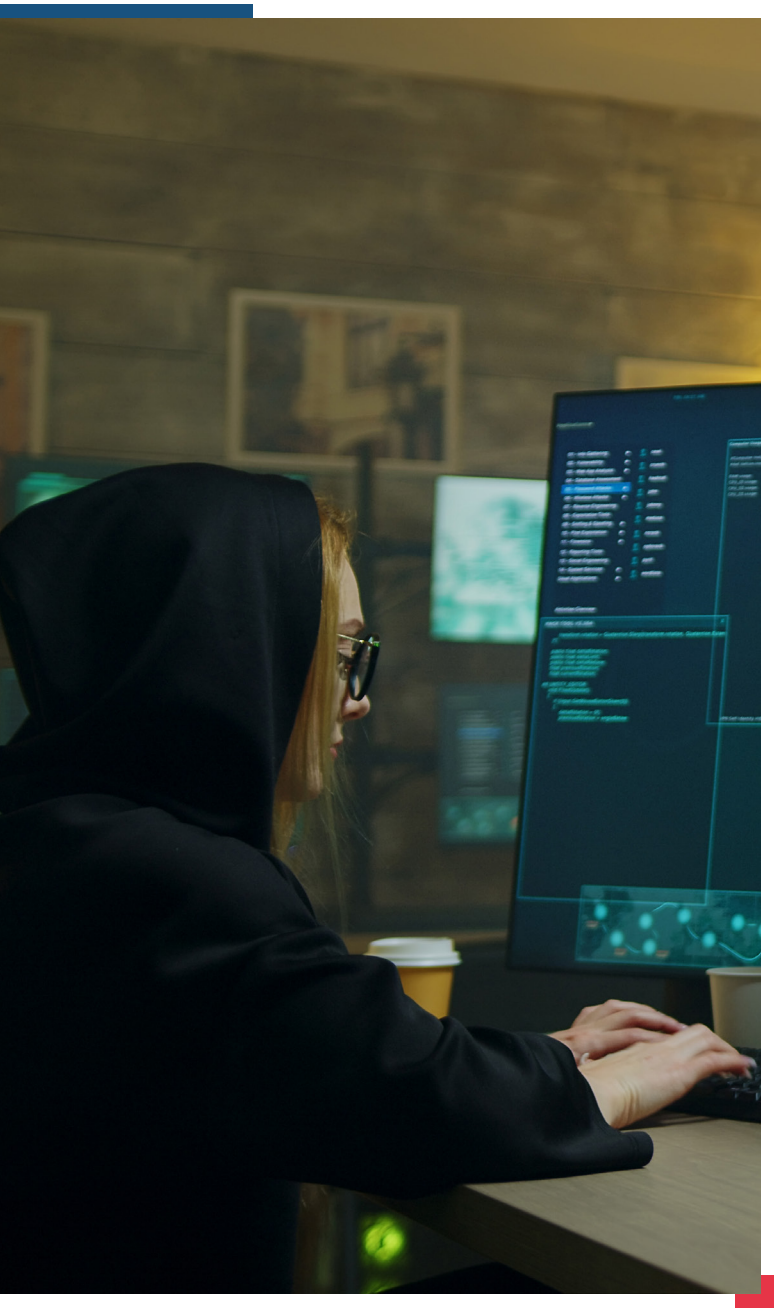

A security audit is a process of cataloging and inspecting security controls. The audit process and report typically map to a list of security policies. For example, if a company has a policy that requires all data to be encrypted at rest, there should be a corresponding security control, powered by suitable technology, to enforce the policy and encrypt data at rest. A security audit will check to see if that encryption tool is correctly deployed in order to validate that the control is in effect—and that the policy is being enforced as it is defined. This may involve a spot check of some kind, e.g., testing to see if data is encrypted in storage.

A pentest is comparable to red teaming, but much more limited in scope. A pentest will use attack simulation techniques, such as brute force password guessing processes, to see if a system can withstand a penetration attempt. Unlike a red team exercise, the process is transparent. The test target knows it's being pentested, and has usually disclosed its settings and functionality to the pentesters. It's not a holistic process like red teaming.

Security audits and pentesting are not substitutes for red teaming. Only a red team exercise can reveal security flaws that escaped the notice of the blue team. Security audits, and to some extent pentesting, are necessary. Without an audit, it's impossible to check whether security policies are being enforced. It would make little sense to do a red team exercise without this basic inspection as a prior step. Pentesting can be useful for checking whether a specific system, or upgrade, has been deployed securely. A red team exercise would represent an excessive effort for such a purpose.

A brief word about ethics in red teaming: An effective red teaming exercise will use real social engineering practices, along with other modes of attack. There is the potential for real damage to systems, along with invasion of an employee's privacy. The impacts are limited, but the client needs to be aware of potential risks when they agree to the exercise.

KALWEIT ITS

HOW ENTERPRISES CAN BENEFIT THE MOST FROM RED TEAMING

Enterprises stand to benefit from red teaming for a variety of reasons. For one thing, a well-conducted red team exercise will either reveal unknown weaknesses in security or affirm that security controls are doing their jobs. Both outcomes are desirable. Most blue teams eagerly await the findings of a red teaming exercise, even if they are a little bit embarrassing. They know it is far better to feel a twinge of embarrassment in a red team debriefing than suffer a data breach. It's much better to know about vulnerabilities than not know.

The full benefits of a red teaming exercise, however, come from what happens after they've concluded. The challenge for security managers is to take the findings from the exercise and operationalize them into new policies and programs. For example, if social engineering leads an employee to divulge a password, then more security training may be in order. Or, if the exercise reveals the existence of shared passwords, then the follow up should involve banning the practice or making it impossible for users to share passwords.

HOW TO IDENTIFY A GOOD RED TEAMING PROVIDER

How to Identify a Qualified Red Teaming Provider

A reliable Red Teaming provider can be recognized by these key attributes:

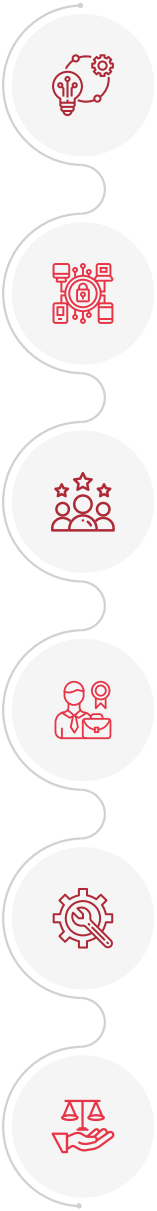

Use of Modern Tools and Techniques:
Deployment of advanced C2 frameworks (e.g., Brute Ratel, Cobalt Strike) and methods like EDR/XDR evasion, payload encryption, and social engineering.

Purpose-Built Red Teaming Infrastructure:
A mature, field-tested Red Teaming IT infrastructure is essential for simulating realistic and complex attack scenarios. This includes advanced tools, communication channels, and tailored evasion techniques that exceed the capabilities of standard penetration testing setups.

Experience and Specialized Knowledge:
Effective Red Teaming requires in-depth expertise in holistic attack strategies and social engineering. This specialized knowledge goes beyond the standard skill set of a penetration tester and is critical for identifying vulnerabilities across organizational and technical layers.

Dedicated Red Team Specialists:
Providers that assign generalist penetration testers to Red Team engagements may lack the depth required for this discipline. Red Teamers bring additional expertise in comprehensive attack emulation, which is not part of a typical penetration tester's toolkit.

Open Source Utilization:
Effective use of open-source tools to enhance attack techniques, particularly in evasion and communication strategies.

Legal and Ethical Compliance:
All actions are conducted within a pre-agreed "Permission to Attack" framework, ensuring adherence to legal and ethical standards.

A qualified provider combines advanced technical capabilities, a dedicated and specialized team, and a methodical approach to uncover and address critical vulnerabilities efficiently and comprehensively

CONCLUSION

Red teaming is one of the most effective ways to discover critical weaknesses in your cybersecurity architecture, policies, and operational practices. By taking a holistic approach and emulating the attack path used by sophisticated hackers, red teaming reveals attack surfaces that you may not have even thought about. Or, the process reveals problems you thought you had dealt with, like data encryption or network segmentation.

Advances in AI make the need for red teaming all the more compelling. Security audits and pentesting have their places in security strategy, but only red teaming will show you where you're truly deficient—and highlight how to remediate your security gaps. The process is no longer just for large enterprises. Smaller organizations can benefit from it, as well. If you're concerned about the impact of AI on your security posture, among other newly heightened risks, it's time to consider red teaming.

ABOUT US

As an independent consultancy specializing in manual penetration testing and red teaming, we deliver tailored solutions designed to meet the unique needs of our clients. By thinking like real attackers, we identify vulnerabilities that others often overlook. Whether for mid-sized businesses or large corporations, we bring fresh perspectives to IT security.

For over seven years, companies across various industries, including prominent DAX-listed firms and mid-sized enterprises, have placed their trust in us. With passion and extensive expertise, we support IT security teams in protecting their systems against one of worldwide greatest business risks: cyberattacks.

Our strength lies in our flexibility and competence. With a fully employed team of experts, we handle high-volume projects efficiently and accommodate even short-notice requests with ease.

Our dedication and expertise have earned us significant recognition, including being named to the Forbes 30 under 30 DACH list in 2019 and honored as "Hamburgers of the Month" by DIE ZEIT. These accolades reflect our commitment to advancing IT security at the highest level.

KALWEIT ITS

For more information, visit

https://kalweit-its.de/en/