JUMPSEC

WHY ATTACKERS DON'T CARE ABOUT YOUR LISTS OF VULNERABILITIES

At the time it was first introduced, a penetration test accurately represented how an attacker was likely to target a network.

Today, that is no longer the case. As digital networks and business processes have evolved, so too have their security needs.

As we will explain in this article, this means that vulnerability exploitation is no longer as critical to an attacker as it once was. Despite these changes, penetration testing has continued to focus on the identification and remediation of vulnerabilities.

Our address
Unit 3E-3F, 33-34 Westpoint,
Warple Way, Acton, W3 0RG

Give us a call
0333 939 8080

Send us a message
hello@jumpsec.com

Find out more
www.jumpsec.com

Introduction

Penetration testing has been a staple of the cyber security industry for several years, relied upon by the vast majority of organisations for security assurance – demonstrating to internal and external stakeholders that they have taken appropriate steps to secure their network from cyber threats.

The first generation of penetration testing was designed to assess much simpler networks than those today. They lacked the complexity and scale of modern environments, with minimal traversal required for an attacker to move from the breach of a network's perimeter to the point of being able to perform a malicious action. With fewer assets to protect and shallower networks to contend with, safeguarding against the exploitation of vulnerabilities was fundamental to preventing an attacker from achieving their goal.

At the time it was first introduced, a penetration test accurately represented how an attacker was likely to target a network, simulating targeted actions against an environment designed to achieve attacker goals that could cause harm to the business. Today, that is no longer the case.

Networks today are much larger in size and scale, are much more diverse in terms of technologies implemented and the complexity of the assets, and are subject to more frequent development.

As we will explain in this article, this means that vulnerability exploitation is no longer as critical to an attacker as it once was. Despite these changes, penetration testing has continued to focus on the identification and remediation of vulnerabilities.

Over time, vulnerability-centric security audit procedures have become the norm, and the focus on vulnerabilities has increased further with the second generation of penetration testing. Vulnerabilities can be neatly quantified, categorised and scored, leading to misplaced belief in the equation that the more vulnerabilities that are identified and resolved, the more secure the network. Today, security teams are held to vulnerability-based compliance requirements by risk managers, whose standards may not have been updated since they were introduced.

The continued evolution of digital systems and technologies has created an infinite funnel of vulnerabilities to identify, manage, and remediate at an ever-increasing velocity. Over time this has created more work for less reward in terms of the security value of activities performed.

In reality, it is impossible to economically identify and remediate every vulnerability in a timely manner. Penetration testing's continued focus on vulnerability management means that it no longer meets its original goal – to prevent an attacker from performing actions that are likely to cause harm to the business.

We believe that with a shift in how security compliance is measured, organisations taking an alternative, threat-led approach can enhance the value of their cyber assurance activities. This article explores how such as an approach can deliver improved security outcomes and return-on-investment on penetration testing spend.

Analysing a typical 'List-based' security approach[1]

A typical security programme relies upon regular penetration testing to identify and prioritise vulnerabilities for remediation based on their risk score. Testing is performed individually against a list of known "critical" systems and applications to identify vulnerabilities.

This then produces a list of vulnerabilities to be prioritised and remediated according to perceived risk. However, employing a list-based approach to securing digital assets is not optimised in terms of building resilience to cyber risk.

The second generation "list-based" approach to penetration testing makes the following false assumptions:

Vulnerability exploitation is central to every cyber-attack

Attacks regularly see the abuse of legitimate functionality over vulnerability exploitation.

While the quantity of vulnerabilities as a metric is convenient to measure and benchmark, list-based approaches fail to identify the relationships between assets or consider how real-world attackers think and act when targeting an environment.

Not all attacks involve vulnerability exploitation, and not every vulnerability can reasonably be exploited as part of a cyber-attack. More to the point, if a vulnerability does not relate to the way that an attacker is likely to target a network, or cannot be reliably used as part of a probable attack to cause harm, can it even be characterised as a 'vulnerability'?

Attackers will look to leverage and exploit technically complex vulnerabilities

'Point and click' exploits pose a greater risk than complex, custom vulnerabilities. The most commonly observed vulnerabilities exploited in the wild are those affecting more prevalent third-party technologies, services, scripts, and tools, for which pre-weaponised exploit kits are likely to exist already.

These capabilities are not exclusive to the most advanced threat actors and can be trivially acquired over the internet for free or for hire. Very few attackers will spend their time conducting intensive vulnerability research. Even the most targeted attackers are highly opportunistic in their nature, capitalising on the latest disclosures and other actors' work to progress their offensive campaigns.

---

1 More information on the concept of Lists versus Graphs, introduced by John Lambert (Microsoft), can be found **here.**

Attackers will always look to identify the 'path of least resistance' to their objective, the most optimal path across the network. They are unlikely to expend valuable time and effort in producing complex technical exploits where low effort, low complexity paths to their objectives exist.

Attackers will target internally developed assets over third party applications and scripts

The majority of attacks leveraging vulnerabilities will target commonly used third-party technologies over internally developed custom applications. Third-party technologies with large-scale adoption are a highly valuable resource for attackers due to their widespread implementation.

Focusing on third-party products presents much greater return on investment for an attacker, presenting the opportunity to target numerous organisations instead of just one. They are therefore subject to higher levels of scrutiny from both the security community and the attackers themselves, naturally increasing the risk of 0-days being identified. When vulnerabilities are reported, they are typically much more accessible to the varying levels of threat actor, with exploit kits easily accessible over the internet.

Further, attackers targeting custom-built assets must necessarily resort to noisier discovery tactics; without access to an isolated version of the technology to research, attackers cannot develop custom exploits in the wild without targeting a live environment. This means they have to rely on more predictable discovery tactics, which are more likely to flag as obviously suspicious behaviour.

Attackers looking to compromise 'critical' systems must compromise them directly

An attacker does not see digital assets in the same way that an internal security team does. Not all are equally relevant to them, and what is most critical to the business is not necessarily the most valuable or pivotal to an attacker in terms of how they will look to traverse the network.

Instead of looking at a network as a list of critical applications against a list of potential exploits, skilled attackers see a graph – charting the nodes and relationships between systems and assets that can be followed to take the most direct course across the network to enable them to achieve their objectives.

The majority of vulnerabilities, in isolation, cannot be exploited to allow an attacker to reach their goal instantly – be it the theft of sensitive data, deployment of ransomware, or performing a fraudulent transaction. Equally, attackers will very rarely target an asset in isolation, nor do they have to directly target the asset or service that represents their technical goal. Attackers must chain actions to achieve their goals; however, not all stages of an attack chain necessarily involve vulnerability exploitation.

How List-based approaches fall short

List-based approaches fall short because they do not accurately represent how an attacker will target the network in reality.

In JUMPSEC's experience, most successful attacks will involve the following attack vectors:

- **Opportunistic exploitation of the latest public vulnerability disclosures** in third-party applications, scripts, or technologies running on external-facing infrastructure.
- **Use of novel or emerging malware types** in a way that evades existing detection controls before automated rulesets can be developed to counteract them.
- **Phishing attacks**. Phishing remains the most effective tool in an attacker's arsenal by manipulating human emotions and behaviours.

With these most prevalent attack vectors in mind, the bulk of vulnerability management activity performed under a conventional security assurance approach is of spurious value when considering the limited exploitability of the majority of vulnerabilities identified.

The issues most likely to be exploited are at a higher risk of being missed under a cyclical testing approach, as testing as infrequently as 6-9-month intervals often leaves organisations exposed to emerging exploits if they land in the window between tests.

According to a recent IBM report , 'scan and exploit' attacks – those leveraging recently released exploits in commonly used technologies, frameworks, and applications – was the top breach vector in 2020 at 35% of attacks recorded. This surpassed phishing for the first time, at 33%.[2]

During a conventional penetration test or vulnerability scan, assets are also examined in isolation. Evaluating vulnerabilities and misconfigurations without considering the context of an asset is likely to reduce the effectiveness and efficiency of any security hardening activities performed.

Notable characteristics that can inform testing include:

- Where the asset is located on the network
- Which systems it interfaces with
- What third-party technologies are running on the asset
- Whether it possesses functionality that an attacker could abuse

Vulnerabilities deemed objectively 'high risk' by general scoring frameworks (such as the Common Vulnerability Scoring System (CVSS)) may represent less of a risk when viewed in context.

2 https://www.ibm.com/security/data-breach/threat-intelligence

By revolving around the mitigation of vulnerabilities, list-based security programmes also focus exclusively on preventing malicious activity. It is widely recognised today that preventive controls will inevitably fail to a persistent and capable adversary.

For an organisation to effectively defend against cyber-attacks today, it must incorporate monitoring controls to detect malicious activity – both targeting the perimeter and on the internal network. Preventing an attacker from gaining a foothold on the network is no longer enough; secure organisations possess the capability to proactively locate and eject an attacker once they gain access to the network. Therefore, assurance activities should also extend to the assessment of detective controls that form part of a layered security model to thwart modern cyber threats.

By taking a vulnerability-centric, rather than risk-focused, view of security, list-based approaches can often lose sight of the goal; to reduce the risk of cyber attackers causing harm to the organisation. Testing is often performed more for compliance than security, with approaches designed around the need to meet often arbitrary internal and external requirements that do not reflect the organisations' real security needs, in the context of how a real-world attacker is most likely to target a network.

Thinking like an attacker by seeing the network as a graph

By considering the ways that an attacker is most likely to target their network, organisations can build a graph-based picture of their estate - identifying the systems and assets likely to be useful for an attacker to achieve their goals, and the associated business risk and impact were those goals to be achieved.

Thinking in graphs and considering the business context of a digital network ensures that security assurance activities accurately represent how an attacker is likely to target the network. Unexploitable vulnerabilities and low risk 'critical' assets that are not relevant to an attacker can be de-prioritised, enabling investment to be focused on the areas that matter most in terms of risk reduction.

One of the primary challenges with a typical security programme is that it is generic. It assumes that all organisations have similar security requirements and must perform the same types of activity. In reality, the most effective security approach for an organisation will be tailored to its unique requirements. Of course, this is something that organisations which rely on generic products and solutions resist; the more universal an approach is, the easier it is to deploy and scale across customers.

Organisations should consider the following questions when planning their security activities for the coming year; while few organisations will be able to confidently answer the majority, an effective approach should look to uncover the answers as part of the programme to contextualise activities and ensure they add value:

- Why would an attacker target you? What goals are they likely to have? What might the business impact be if they were to achieve those goals?
- Where on your network is an attacker most likely to land / interact with from the internet?
- What technologies do you have exposed over the internet?
- What are third-party technologies, products, and services present on your network?
- Which internal systems are most likely to be reached by an external attacker? Which systems offer attackers the most pivotal access to your internal systems?
- Have you considered the most probable routes an attacker will take to traverse your network?
- Do the assets you test feature on the most likely attack paths across your network?
- Do you offer authenticated access to your network to external / third-party users or systems?
- Which of your systems (internal and external) are most valuable to an attacker?
- Are you aware of any malicious activity outside your network, for example, where attackers are typosquatting or impersonating your brand?
- Have your employees have previously been targeted? For example, have your employees' credentials been exposed in a public breach, and have attackers attempted to use them your network?

Guidance on implementing an effective security assurance programme

Attackers today are continuously innovating; adapting their tradecraft to leverage emerging exploits and continuously probing target networks for weak points. They seek out low resistance paths into and through digital estates, seeking to circumvent traditional IT defences by using techniques which are not reliably discovered through traditional cyber-security approaches.

Testing performed for assurance purposes should represent how an attacker will target the network, rather than exhaustively identifying vulnerabilities that pose limited real-world risk. Offensive tactics are not limited to vulnerability exploitation, and include:

- Stealing credentials from trusted third-parties
- Harvesting of personal information to create sophisticated phishing campaigns
- Creating typosquatting websites to impersonate brands and defraud customers

Attackers are not bound by the internal testing schedules or corporate governance, and just because an asset is critical to the business does not mean it is of value to an attacker when charting a path through the network.

To stay ahead of this threat, defenders need to think like attackers. Third generation penetration testing should focus on understanding what, where, why, and how an attacker will look to target the network, in order to build appropriate and targeted defences.

In general, if planned cyber security investment in a test, control, product, or tool is not guided by an understanding of how it reduces an attacker's ability to cause actual harm, organisations should look to establish this context before making the purchase. Expecting investment to deliver risk reduction without first establishing the outcomes required inevitably means that it will fail to deliver the desired results.

Any organisation looking to purchase a product must always clarify why they need it, and how the product will meet this need, before investing. **The same is true of a penetration test.** Testing performed for assurance purposes should represent how an attacker will target the network, rather than exhaustively identifying vulnerabilities that pose limited real-world risk.

If the buyer doesn't know how or why an attacker will target their organisation – which assets they will target, what the most prominent attack paths are across their estate, and what controls they have deployed to reduce the risk – then they should look to increase the visibility of their security requirements before testing.

This will be the single most valuable cyber security exercise an organisation can conduct. Once an organisation knows how they are most likely to be attacked, they can align their defensive controls with the threat, in accordance with the risk posed.

Conclusion

The 'third generation' of penetration testing can deliver enhanced outcomes by ensuring investment is tied to actual cyber risk reduction in the context of how an attacker is likely to target the network – closing or building mitigating controls around high-risk attack paths to either prevent an attacker from performing a malicious action, or tune detection and response capability to control the risk.

Understanding how an attacker will target the organisation, based on its unique threat profile, technological composition, and business requirements, is the foundation of any effective security programme, and should form the basis for future penetration testing approaches if they are to deliver real security value.

By eliminating generic investment and shifting to a threat-led, graph-based approach, organisations can prevent waste and enable genuine security improvement to be realised over time, rather than playing catch-up chasing an infinite number of vulnerabilities that may not pose any real risk.

Unit 3E-3F, 33-34 Westpoint,
Warple Way, Acton, W3 0RG

0333 939 8080

hello@jumpsec.com

www.jumpsec.com