intruder

# The ultimate guide to vulnerability scanning

In today's hyper-connected world, reports of cyber-attacks and data breaches are commonplace. On any given week of the year, you can count on seeing news reports of the latest cyber incident.

With the average total cost of a data breach coming in at an astounding $3.92 million (according to the latest analysis by IBM), you can understand why cyber security is an increasing concern for businesses all over the world.

Vulnerability scanning is a fundamental component of all good cyber security strategies, but it can be complicated, and challenging to get right.

Whether your organization is just starting out on its journey to becoming more secure or looking to improve on existing security controls and learn more about vulnerability scanning best practices, this guide has something for you.

In this guide, you'll learn about: what is vulnerability scanning, the different types of vulnerability scanners available, vulnerability scanning frequency best practices, how to choose a vulnerability scanner, and how to get up and running with your chosen product. So let's get started.

# Table of contents

intruder

# An introduction to vulnerability scanning

## What is vulnerability scanning?

Vulnerability scanning is, at the simplest level, the use of software tools to identify and report on security issues (known as vulnerabilities) that affect your systems.

Vulnerability scanners often have many thousands of automated tests at their disposal, and by probing and gathering information about your systems, can identify security holes which could be used by hackers to steal sensitive information, gain unauthorized access to systems, or to cause general disruption to your business.

Armed with this knowledge, an organization looking to protect itself can then take action to remediate the security weaknesses discovered. This overall ongoing process of identifying and fixing your weaknesses is known as **Vulnerability Management**.

> Vulnerability scanners offer an excellent starting point, allowing an organization to identify their most serious and most exposed technical weaknesses so they can react before an attacker takes advantage.

## Who needs vulnerability scanning?

News headlines focus on the biggest security breaches, which usually affect large organizations. So you'd be forgiven for thinking that cyber security is a "big company" problem. However, when it comes to cyber security, unfortunately, small doesn't mean safe.

In the UK, a recent survey conducted by Ipsos Mori on behalf of the UK Government concluded that 32% of smaller businesses reported an attack or breach.

Whether it's delivering marketing or blog content via a website, operating internet-exposed applications or services, or simply the laptops your employees use for work; there's almost always a range of systems that could be attacked. All of these systems comprise an attack surface for hackers to target.

Successful breaches then lead to ransomware attacks, or even a compromise of less-sensitive data (such as names and addresses), which can result in an exodus of customers to a competitor, or a hefty GDPR fine.

Developing a comprehensive security strategy to mitigate these threats is a process that takes years to get right, and it's one which should constantly change and adapt as an organization grows and the threat landscape evolves. Vulnerability scanners offer an excellent starting point though, allowing an organization to identify their most serious and most exposed technical weaknesses so they can react before an attacker takes advantage.

In short, every business should understand where their cyber weaknesses are, and get them fixed.

---

## Vulnerability scanning vs penetration testing?

Vulnerability scanning isn't the only way to discover the vulnerabilities that affect your organization's systems – manual penetration testing is also a common way to check your systems. However, the two differ quite significantly in what they have to offer and how much they cost, so it's a reasonable question to ask which of these are appropriate for your organization and how often they should be performed.

Both have their pros and cons. Vulnerability scanning has the advantage that it can be performed automatically and continuously at a lower cost, so that new security issues can be identified soon after they are introduced.

Meanwhile penetration testing is usually performed on a consultancy basis, and it comes with time and cost overheads that can slow projects down, or be prohibitive in terms of how often testing is performed.

However, manual pen testing performed by skilled and qualified professionals can discover security issues which are more complex or specific to the business, which require a human level of understanding to discover.

Each has their place, and if budget allows, it's certainly best practice to employ both. However, for organizations that are looking to get started with protecting their business for the first time, we recommend first setting up a vulnerability scanner and regularly testing your publicly exposed attack surface. This makes most sense since penetration testers also use vulnerability scanners as part of their offering and there's no point in paying a professional to tell you something you could have found out for yourself. Even more importantly, if you can only run a penetration test once per year, you remain exposed to configuration mistakes and new vulnerabilities in between tests.

intruder

# Getting started with vulnerability management

## What should I scan?

In order to use a vulnerability scanner, you first need to know what you're going to point it at. It might sound simple, but if you're new to vulnerability scanning, you might find that there is no central record of the systems your organization is responsible for. If that's the case, it's time to start keeping track – if you don't know what you've got, you won't be able to protect it!

## Asset management

It's good practice for organizations to keep a centralized record of the systems they have under management (commonly referred to as Asset Management). Keeping up to speed with your organization as it grows or changes is essential. As new systems go live, or existing ones change their IP addresses or domains, keeping your documentation up to date will help make sure that systems don't fall through the gaps, and miss out on all the hard work your scanner is putting into identifying your security weaknesses.

If you're using modern cloud systems for some of your estate, then this may help somewhat, and modern vulnerability scanners will be able to hook into your cloud accounts to make this process seamless. However, you will undoubtedly have some systems that are outside this (employee devices and edge routers and firewalls at the very least), so it's still a good idea to keep an asset register.

## Scoping strategies

Now you know what you've got, how do you decide what to scan? It's common for organizations to deploy a range of systems, from laptops and workstations in the office or at home, to systems in cloud platforms like AWS, Azure, and Google Cloud. Deciding what to include in scanning can be hard work, but there's multiple ways to tackle it. Here we present three strategies, exposure based, sensitivity based and coverage based:

### Exposure-based

Any of your systems which are publicly accessible over the internet are effectively available for attack 24 hours a day. As a result, these systems are scanned for vulnerabilities by attackers on a constant basis. In fact, even unskilled attackers can automatically scan the entire internet for weaknesses using freely downloadable tools.

Say for example your company is a tech startup and offers services over the internet to its customers. This could be via a website, or web application, or anything else hosted online. While access to these applications may be secured under normal circumstances, just one weakness or mistake occurring in one of these systems could lead to an immediate data breach. As these systems are the ones which are being scanned day-in day-out by attackers, ideally you need to be doing the same to find those security issues before they do.

Serious vulnerabilities on publicly facing systems typically get exploited very quickly. Weaknesses can crop up overnight which didn't exist before, with farreaching consequences.

### Sensitivity-based

Your company may not have much on the internet that is sensitive. It could be just that your main website contains just marketing information, but all your sensitive customer information is stored in a central store that's firewalled off from the internet somewhere (whether that's an individual's laptop, or a network file share). If the reputation damage caused by a website defacement doesn't concern you, then in this case you may decide it makes more sense to perform vulnerability scans on the systems where the sensitive data is stored themselves, to make sure they are as hardened as possible against attacks.

### Coverage-based

After considering what's exposed to the internet, and where your most sensitive data is stored, it's worth considering that other company systems are still ripe for an attacker to compromise. Vulnerabilities can exist in any of your systems, and once an attacker breaches one system, their next move is often to use that position as a foothold to launch new attacks.

For example, hackers could breach an employee laptop by sending emails containing malicious files (or links to malicious websites) that exploit vulnerabilities on the system they are opened on. If a device is successfully compromised, it could be used to scan other systems on the same network, to exploit vulnerabilities on those too. Alternatively, information or access to other systems gained from the laptop may be used in further attacks.

For this reason, it often makes sense to attempt to cover as many systems as possible, especially where gaining access to one system could lead to breaching others. Which of these approaches is right for you will depend on your business resources, and where and how your most sensitive data is stored – often, the right answer will be a combination of all three.

intruder

# Types of vulnerability scanner

There are many types of vulnerability scanner which perform different security tasks, and cover off a range of different attack scenarios.

For example, an attacker could break into your internal network through a vulnerability on an exposed web server, or through unpatched software on an employee's workstation.

Identifying these different attack vectors require different use-cases of vulnerability scanners, and each are not always supported, so it's worth considering the risks to your business and finding a scanner that is suitable.

This section goes through the different use-cases in more detail.

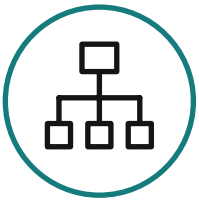

**Network-based vulnerability scanners**

**Agent-based vulnerability scanners**

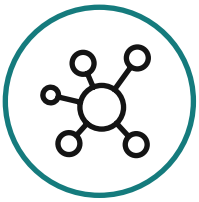

**Web-application vulnerability scanners**

## Network vulnerability scanners

Network vulnerability scanners are so called because they scan your systems across the network, by sending probes looking for open ports and services, and then probing each service further for more information, configuration weaknesses or known vulnerabilities. The way this works can differ, you might install a hardware appliance inside your network, or deploy a virtual appliance on a virtual machine, and then run scans from that machine against all others on the network.

One obvious benefit of network vulnerability scanners is that they can be quick to set up, simply install your scanner and get scanning. They can quickly become more complicated when it comes to maintenance though, keeping appliances up to date, and keeping them in-step with changes on your network. This is especially true the more complicated your networks become, and the number of scanners you need increases to cover each network segment.

### Internal vs External

Network scanners are often configured either to scan “internal” networks, or “external” networks.

As a general rule of thumb, the larger and more complex your private networks become, the more important it will be to also consider employing an internal network scanner that can check your internal systems for weaknesses that could lead to the compromise of a single system turning into a wider breach.

Smaller organizations with a much more modern digital footprint may not decide to put internal scanning in place early on, such as those with all their servers in the cloud and a simple private network for internet access only. However, these organizations would still benefit from agent-based internal scanning, so that vulnerabilities on employee’s devices can be identified for patching.

### External network scanning

An external vulnerability scan is simply one which scans your systems from the outside. What this really means is, the scanner’s probes come from an untrusted internet address which is outside of any of your organization’s private networks. They simulate the activities of a remote attacker looking at which systems and services are offered over the internet and start by looking for vulnerabilities on your internet exposed systems.

Although the amount of information that can be discovered by these scans can be limited compared to the other types described below, they can also be very revealing in understanding what attackers can see. That’s because if an attacker can see a weakness, it’s highly likely they are going to exploit it.

Some vulnerability scanners are well set up to deal with this situation, and have scanners ready in the cloud, so no deployment or management of your own systems are required. They are simply point-and-click.

Others may require you to set up your own scanning appliance, and manage this on an ongoing basis. While such services can be cheaper, the overheads here can sometimes be not worth the difference in cost.

### Internal network scanning

Internal network vulnerability scans are designed to find weaknesses on systems which do not expose ports or services to the internet.

This kind of vulnerability scanning helps to cover off a range of attack scenarios that couldn't be scanned for by external vulnerability scanners.

For example, if an outdated version of the Firefox browser is in use on a company laptop, the machine could be vulnerable to attacks if a user is convinced to visit a malicious website.

Similarly, there may be vulnerabilities in the ports or services a device exposes within a private network (such as weaknesses in the SMB service), which also could not be discovered by an external scanner.

Internal network-based scanners work broadly in the same way as external network scanners do, except the scanning device sits within an internal network, so services and devices which only expose themselves within a private network can be assessed.

Internal scans can be useful for identifying potentially vulnerable devices that were not known about in advance, as they can sweep a whole network range.

However, they can be highly ineffective at providing detailed information unless they are provided with credentials for logging into systems and querying for specific patch and configuration data. This is known as "authenticated scanning".

Authenticated scanning can provide much more detailed vulnerability information, but it can be tricky to configure and maintain. For that reason, a popular alternative is running "agent-based" scanners.

## Agent-based scanners

Agent-based scanning is performed by installing lightweight software scanners on each device to be covered, which can run local vulnerability scans and report back to a central server with the results.

In the same way as authenticated network scans, this type of scanner can pick up on a wide range of vulnerabilities, including weaknesses in software which doesn't expose ports or services for remote access at all (e.g. a vulnerable version of Firefox).

While it can be a little more time consuming installing agents across your digital estate, they have the benefit that once they are installed they can report back even if removed from the network (such as laptops being taken for home working).

Both types of internal scanner have their limitations and advantages.

For modern organizations with simple internal networks and the majority of their infrastructure in the cloud, an agent-based scanner would be the logical choice.

For organizations with complex internal networks, the choice of which type to go for is a little more difficult, and for the most mature organizations - and where budget allows - deploying a combination of both types of scanner should be considered.

For a full discussion on the advantages and disadvantages of agent-based and network-based scanners, we've written an article which goes into more depth.

## Web application scanners

Web application vulnerability scanners are a specialized type of vulnerability scanner which focus on finding weaknesses in web applications and websites. Traditionally, they work by 'crawling' through a site or application in a similar way as a search engine would, sending a range of probes to each page or form it finds to look for weaknesses.

Many vulnerability scanners include web application scanning as part of their offering, although it can be licensed separately. Others are dedicated purely to web application scanning, while some vendors include it along with a range of other checks.

One thing you might want to look out for is whether the scanner can perform authenticated web application scanning or not. Authenticated scanning is where the application is scanned past the login page, from the perspective of a malicious user, or attacker with credentials to log into the app.

Due to the amount of business logic and complexity that goes into making web applications, even the very best vulnerability scanners on the market today struggle to identify some application flaws effectively, and they still sadly don't come close to a human expert looking for flaws manually. It's important to understand what they are good at, and what they struggle with.

They are generally good at identifying straightforward weaknesses (such as simple SQL injections and cross-site scripting flaws). Weaknesses which are less straightforward to exploit cannot reliably be detected, in particular:

Access control weaknesses (such as unauthorized access to information which should require a higher privileged account)

- Exposure of sensitive information (scanners can often discover this information, but can't always tell it's sensitive!)
- Weaknesses in multi-step workflows (such as multi-page forms) - Weaknesses involving storing a payload which gets executed elsewhere (such as persistent cross-site scripting)
- Session-based weaknesses (weaknesses in the mechanisms used to manage user authentication)

### The OWASP Top Ten

The OWASP project has long since been the go-to resource for web application vulnerabilities, and they release a "Top Ten" summary of the most common web application flaws every few years. Of the ten issues OWASP lists in the latest version of the document, many are either poorly detected by web application scanners, or only certain types of the flaws can be reliably detected.

### Single Page Applications

Modern single-page apps are tough for automated scanners, as they fail to properly discover and generate legitimate application requests to perform their tests with.

### False Positives

Verifying whether vulnerabilities are false positives is a difficult task for an automated scanner to do effectively, and most do not attempt to do so.

As a result, you may end up trawling through long lists of non-issues which could quickly grow into a time-consuming process.

Automated scanners are certainly capable of discovering genuine web application security issues, but they are unlikely to catch everything and shouldn't be the only defense that's put in place.

Our recommendation is to ensure that web applications are regularly penetration tested, including where a web application scanner is in place. This strategy is a more robust way of discovering a wide range of weaknesses and, if budget allows, it is a best practice.

intruder

# Choosing a vulnerability scanner

Choosing the right vulnerability scanner for you can be difficult, and the range of quality in vulnerability scanning products is significant and not always transparent. You'll be relying on your chosen scanner to help prevent attacks daily, so how do you know a scanner is effective, and how are you supposed to compare one against another? We've included a few due diligence tips and best practices below.

## Take it for a spin

Most vulnerability scanners offer limited free trials, to give the user a chance to get used to how it works, and what the features are. This is a great way to get a feel for the product, its features and usability. A logical next step is to run a scan against a selection of your own systems and see what comes back. If you're comparing multiple scanners, it could be good to run them both against the same systems and see what is discovered. Unfortunately, a like-for-like comparison of two or more scanners doesn't always show a clear picture of how they compare.

For example, one of the scanners you're appraising may return more security issues which are false positives (issues which the scanner has mistakenly identified as a security issue). If you don't have the ability within your team to verify whether a security issue is valid or not, then this exercise may not be enough.

It's also worth noting that there might not be anything wrong with your systems right now, which reduces the value of doing this type of comparison of scanners. If the systems you're scanning do not have a wide range of security problems (that you already know about), it will be tough to gauge how good a scanner is.

Furthermore, a lot of vulnerability scanners stuff their results with 'Informational' issues which are not actually security problems. Watch out for these and remember that a simple comparison of the numbers of issues each scanner has discovered is missing the point. You'll likely be interested in which scanner can find the most genuine security problems, and the best way to do this is to scan systems which are known to be vulnerable.

## Find out what the scanner can check for

Most vulnerability scanners offer a list of security issues that the scanner checks for. This can be a good way to help you decide on which scanner is right for you. By reviewing the scanner's documentation, you can confirm that it is capable of checking for security issues in the range of software and applications which comprise your organization's digital estate. We've listed below some broad classes of vulnerability which a comprehensive vulnerability scanner should be able to check for, with some examples:

### Vulnerable Software

This class of vulnerabilities is the biggest category, as it includes checks for known weaknesses in all kinds of 3rd party software and hardware. These are weaknesses discovered by security researchers in certain versions a particular technology. Some examples of these would be a weak version of a Nginx web server, a vulnerable FTP service or weaknesses in a Cisco router or VPN.

### Web Application Vulnerabilities

These are weaknesses in your web applications. There's a wide range of weaknesses which could be used to gain unauthorized access to information, compromise the web server or attack web application users. Some examples of these would be weaknesses such as SQL injection, cross-site scripting and directory traversal weaknesses.

### Common Mistakes & Misconfigurations

These are a class of weaknesses which include identifying software which has been incorrectly configured, commonly made mistakes, and security best practices which aren't being followed. Some examples of these would be exposed SVN/git repositories, open email relays, and or a web server configured to reveal sensitive information.

### Encryption Weaknesses

A wide range of weaknesses in the encryption configurations used to protect data in transit between your users and servers can be identified by vulnerability scanners. These should include checks for weaknesses in SSL/TLS implementations, such as use of weak encryption ciphers, weak encryption protocols, SSL certificate misconfigurations, and use of unencrypted services such as FTP.

### Attack Surface Reduction

Some scanners can detect areas where you could reduce your attack surface. This is the principle of publicly exposing only the core services you absolutely need to. External vulnerability scanners can identify ports and services which could represent a security risk by leaving them exposed to the internet. Some examples of these are: exposed databases, administrative interfaces, and sensitive services such as SMB.

### Information Leakage

This class of checks report on areas where your systems are reporting information to end-users which should remain private.

Not all vulnerability scanners include checks for all of the above categories, and within each category the number and quality of checks vary too. Some scanners are focussed on one particular class of vulnerabilities - for example, web application vulnerabilities.

A web application focussed scanner wouldn't necessarily check for infrastructure level flaws, such as known vulnerabilities in the web server in use. If you're only employing one vulnerability scanner, it's certainly worth making sure that it can handle all of the above, so there aren't any gaps in your security coverage.

## Check for essential features

The range of vulnerability scanners on the market are greatly varied, and each offers a unique set of features which offer different core functionality and "nice-to-haves" which are non-essential features that are designed make your life easier.

Before choosing a scanner, it's worth asking yourself which features are essential for you, and which you don't need. Armed with this knowledge, you'll be able to more easily decide which product to go with. We've listed below some scanner features you may wish to consider:

- **Scheduling** – can you schedule scans to run out of peak hours, or during supported times?
- **Frequency** – how often can you run scans?
- **Reporting** – is the report easy to read and could you pass it on to a customer?
- **APIs** – can you programmatically trigger a scan after a development cycle?
- **Compliance** – is the scanner appropriate for your compliance requirements?
- **Cloud integrations** – does the scanner integrate with your cloud provider
- **Proactive scans** – can the scanner check your systems for the latest threats automatically?

## Reporting

Reporting is an important factor to consider on its own. There are two main uses for a security report from a vulnerability scanner: your developers and security engineers will use it to fix security weaknesses, and you may also need to use it to pass onto your partners and customers as proof of your secure practices.

It's important that the security issues detailed in the report give remediation advice in clear language that can easily be used to resolve the issue. Some vulnerability scanning reports are difficult to read and understand, whilst others present a clear, concise description of a security issue along with simple instructions on how to put a fix in place.

It's common for prospective customers or partners to ask for proof of security. When passing on a security report to a third party, you'll want to make sure you can easily pass on a well-formatted document that clearly details any remaining vulnerabilities and gives the reader a good insight into what's been tested for. Most vulnerability scanners will allow you to download a report during the trial period, so don't forget to take a look.

## Pricing

Price and available budget are always going to be a major consideration when choosing a vulnerability scanner. Cyber security budgets are often tight, and there are a wide range of security products and other costs which are competing for the same budget that will be spent on a vulnerability scanner.

Thankfully, most vulnerability scanners on the market are fairly priced in comparison with what they offer, so in general you do get what you pay for. That said, some vulnerability scanners are cheaper because they offer a cut-down set of features, which you might not require, so some shopping around to try out a few different scanners is time well spent.

Comparing price of vulnerability scanners is an area in which it's worth treading carefully too. The range of software which can be classed as a 'vulnerability scanner' is vast, and the quality and range of checks they offer varies greatly too.

If one of the options you are considering is significantly cheaper than the others, then some extra due diligence may be required to make sure that it's capable of performing the same range of security checks as the others you are comparing it to.

## Licensing and discovery scanning

Most modern vulnerability scanners have a pricing model which varies depending on the number of systems you'd like to cover. Pricing will vary depending on the type of scanner you're using, and which features you require, but broadly, you'll be charged depending on the size of your digital estate and how many systems are being scanned.

As we touched on in the 'Defining the scope' section above, some organisations may have difficulty answering the question "how many licenses do we need?" at this stage, as they may not know exactly how many live systems they are responsible for. This can be a challenge, especially since the answer to this question may have a significant effect on the cost of the scanner. Many scanners are able to help with this problem, using what's commonly known as 'discovery scanning'. Discovery scanning is a light-touch scan designed to discover which systems are live and which are not.

For example, you may have a range of public IP addresses, such as 1.2.3.4/24, which corresponds to 256 IP addresses. The chances are that not all of these are in use, and you may wish to save on costs by only paying for vulnerability scanning licenses for the systems which are active.

This is where discovery scanning can be useful. You can run a simple scan on your range of IPs to discover which respond – those that don't respond to probes on any port are either inactive or are not exposed to the scanner, so you won't need licenses for those.

Some modern scanners can save licenses for you automatically, by running discovery scans and only using licenses on live systems. This feature can save both time and money, as you can enter all of your known IPs, and the scanner will only charge you for those they are currently live and in use.

# Finally

We hope you have found some useful information on vulnerability scanning best practices in this introductory guide by Intruder.

Intruder offers a free trial of our continuous vulnerability monitoring product, so if you'd like to give our product a spin, you can get started here.

If you have any questions on vulnerability scanning best practices or would like to learn more about Intruder's comprehensive vulnerability scanner, please get in touch with the team at **contact@intruder.io**.