Vulnerability scanning tools and services

Advice on the choice, implementation and use of automated vulnerability scanning tools for organisations of all sizes.

Introduction

Vulnerability scanning is a broad term, used to describe the automated process of detecting defects in an organisation's security program. This covers areas such as the patch management process, hardening procedures and the software development lifecycle (SDLC). Services or products that offer vulnerability scanning are also commonly known as vulnerability assessment systems (VASs).

As part of an effective vulnerability management programme (VMP), vulnerability scanning solutions can be an affordable way to automatically detect security issues within an organisation's networks. However, the market for vulnerability scanning products and services covers many specialised areas and includes a broad range of options involving issues such as deployment models and licence costs. These complications can make it difficult to make the right choice when purchasing a vulnerability scanning solution for your own organisation.

This guidance seeks to provide a broad audience with the tools to select an appropriate vulnerability scanning solution.

Audience and structure

This guidance helps SMEs, large organisations and public sector bodies to:

- understand the basics of vulnerability scanning and how it integrates with a VMP
- decide on when and how to employ vulnerability scanning most effectively

- set the important criteria when purchasing a vulnerability scanning solution

The guidance is split into four steps, beginning with an assessment of your current vulnerability scanning set up, before moving on to consider the type of scanner you need. We then consider what to scan and when, before concluding with some general pointers.

Advantages of vulnerability scanning

There are a number of reasons why organisations should take advantage of vulnerability scanning:

- ***automation***: scanning can be run on a schedule, on-demand or in response to trigger events such as a new build of a software project or the deployment of a new server. This enables an up-to-date view of the vulnerability landscape to be maintained.
- ***speed***: scanners typically perform hundreds or even thousands of checks at a significantly faster pace than would be possible with manual testing.
- ***cost-effectiveness***: the benefits of speed and automation make it far more economical to perform vulnerability scanning against a target than testing it manually.
- ***scalability***: modern cloud-based architectures mean that services can increase or decrease their resources to enable small or large environments to be scanned within similar timeframes.
- ***compliance***: many vulnerability scanning solutions include bespoke checks to test compliance with common information security standards or an organisation's own baseline control set.
- ***accuracy***: by carrying out bespoke checks to confirm the presence of vulnerabilities, scanners can produce far more reliable results than simply referencing information held in software asset management solutions.

Most importantly, vulnerability scanning affords an organisation the ability to keep pace with individuals and groups intent on compromising systems, many of which use similar tools and techniques to discover security flaws.

Relationship to manual testing

It should be noted that automated vulnerability scanning cannot compare to manual processes such as penetration testing when it comes to the breadth and depth of test coverage.

Instead, automated scanning should be viewed as a cost-effective way of finding and managing common security issues, without needing to employ specialist security testers.

Similarly, by taking care of the ‘low hanging fruit’ through regular vulnerability scanning, penetration testing engagements can more efficiently focus on complicated security issues that are better suited to a human.

1. Evaluate your existing vulnerability management programme

Vulnerability scanning is only effective at reducing the risk to an organisation when used as part of a larger vulnerability management programme (VMP).

VMP programs typically include the following processes:

- ***System discovery***: Identifying assets owned by your organisation
- ***Asset classification***: Assigning assets into groups or categories based on common characteristics
- ***Vulnerability detection***: Finding and validating vulnerabilities in assets
- ***Vulnerability triage***: Prioritising vulnerabilities according to technical or business objectives
- ***Vulnerability remediation***: Advising on and verifying the fixing of identified issues
- ***Vulnerability disclosure***: Providing a mechanism for security researchers to disclose relevant vulnerabilities to you. Please refer to the NCSC’s

Vulnerability Disclosure Toolkit for information on creating your own vulnerability disclosure process.

Supporting your VMP

Vulnerability scanning solutions often contain features which support or integrate with a VMP, for example:

- Performing system discovery by regularly scanning for new hosts within your IP address range(s) or new web applications
- Validating systems discovered against existing asset management records
- Tailoring how vulnerability reports are presented, to align with the priorities of your business or organisation
- Supporting the remediation process by re-scanning specific issues and reporting when they are confirmed as fixed
- Integrating with other systems such as bug trackers or source code repositories to help co-ordinate and automate workflows
- Providing a secure authenticated portal for users to login and manage vulnerabilities collaboratively

What features do you need?

The extent to which the availability of these features will impact your choice of Vulnerability Scanning solution depends on your existing VMP and whether they enhance your situation or just provide unnecessary complexity.

For example, an organisation without a VMP already in place would likely benefit from a service that included a central portal to allow different administrators to view and manage the vulnerabilities relating to their own systems.

By contrast, organisations with a mature, established MVP may already have such capabilities and could therefore simply require a product that supports the exporting of results so that they may be easily integrated with the existing solution.

Other features that should be taken into consideration when purchasing a Vulnerability Scanning solution are documented at the end of this guidance.

2. Identify your assets

The word 'asset' is used in the context of vulnerability scanning to define an entity (either physical or virtual) that vulnerabilities are associated with. This can take several forms depending on the type of scanning being undertaken, such as:

- a network infrastructure component such as a router or switch
- a connected virtual or physical host such as a laptop, peripheral device or server
- an instance of a web platform or application
- cloud-based hosts or endpoints

It's common for organisations to own a variety of assets spanning some or all of these categories, though some may be more prevalent than others. It's important that these are identified and recorded (ideally in an asset register) so that the most appropriate type(s) of vulnerability scanner can be sought. Many vendors charge for their scanning services on a 'per asset' basis, so establishing an accurate idea of asset numbers is critical in estimating cost prior to procurement. This can be achieved with the help of (often freely available) port scanning tools to find active hosts on your network.

You may find that parts of your IT estate are highly distributed, for example due to users working remotely using their own mobile devices. In cases such as this, focus on any common services that are designed to be remotely accessed by these devices. For example, perhaps users are required to login to a single web portal or Virtual Private Network server that is externally accessible. Whilst the security of end user devices situated outside the perimeter of your internal network is still important, remote vulnerability scanning is unlikely to be of benefit in these circumstances. Instead, remote devices should aim to avoid common vulnerabilities by ensuring software is kept up to date.

Once you have identified all the relevant assets within your organisation, you should separate them into distinct logical groups. For example, you may wish to place any server hosts or web applications associated with your main web site into one category and your internal desktop estate into another. This helps to define separate, more manageable scopes for individual vulnerability scans.

As mentioned above in '*Evaluate your existing vulnerability management programme*', some solutions support this process by performing system discovery and classification automatically.

3. Choose an appropriate type of vulnerability scanner

Vulnerability scanners are typically categorised according to the type of target they are intended to assess. The broadest distinction being between 'infrastructure' and 'applications'.

Application scanners are further subdivided into those that target web applications and those that target native applications. Scanners also exist for a number of specialist subcategories such as cloud infrastructure, mobile applications or web applications built using a specific platform or technology.

Whilst specialised scanners can offer the most accurate and relevant results for the types of target they are designed to assess, an organisation's IT estate is likely to contain too much variation for such solutions to provide comprehensive coverage by themselves. You should therefore seek to first establish a foundational level of generalised scanning, to ensure a good level of coverage on the most common infrastructure issues.

If your organisation exposes assets in other specific categories, such as those mentioned above - and your budget allows for it - we recommend taking a layered approach to scanning, by supplementing your foundational scanning with more specialised scanners.

Infrastructure scanners

Infrastructure scanning solutions are typically focused on identifying and testing services that are accessible to the rest of the network or the Internet as a whole. For this, they often include host discovery and port scanning functionality.

Once an accessible network service is discovered, they typically probe it to discover as much information as possible. Using techniques such as ‘fingerprinting’ or ‘banner grabbing’, the scanner would collect details like the vendor and version number of the software. Many infrastructure scanners will also send safe test messages to some types of services to probe for more informative responses or directly test for the existence of a vulnerability. Once a service ‘fingerprint’ has been obtained, this is also referenced against a knowledge library of products known to contain security vulnerabilities.

Whilst some network vulnerability scanners also use more advanced methods than this and can even support checks that first require authentication, they typically aim for breadth instead of depth when it comes to coverage. For example, such scanners generally lack the ability to navigate web applications or detect vulnerabilities that require complicated interactions with specialised protocols. However, they may well be able to detect vulnerabilities stemming from the use of outdated software or weak encryption settings on those same ports.

Network vulnerability scanners are therefore an excellent choice for monitoring networks with large external footprints for new common vulnerabilities that could be exploited by attacks from the Internet or your internal corporate network. They are also more useful for IT estates that consist predominantly of ‘off-the-shelf’ solutions contain little or no custom developed software.

Web application scanners

Web application scanners are specifically designed to detect vulnerabilities in applications and web services exposed over HTTP/S.

They achieve this by interacting with applications in much the same way as a web browser would, albeit with the ability to send requests at a much faster rate, formulated to elicit responses from the web server that would indicate the presence of a vulnerability.

Web application scanners typically check for a wide variety of security issues that can affect both the web server itself and other users of the application. These are often aligned with publications such as the OWASP Top 10, which is a periodically updated lists of the most critical security risks to web applications.

Unlike network infrastructure scanners, web application scanners are designed to detect vulnerabilities in custom-built (and often complex) web applications.

Advanced web application scanners may also support more fine-tuned configuration. This may include the ability to specify a login page and credentials for the target application, or the ability to exclude specific types of scans or pages. Without these features, a scanner is not likely to achieve good test coverage against more complex web applications or may produce undesirable side-effects, such as large volumes of database entries from repeated form submissions. In general, the more tuned to the target web application the scanner is, the more relevant and useful the results will be.

Web application security scanners are an excellent choice when used in conjunction with network vulnerability scanners, or when custom web applications account for most of your external network footprint and therefore present most of the risk to your business or organisation. The NCSC's own Web Check service is an example of such a service, albeit one that can only be offered to the public sector. Web Check is specifically designed to be 'light touch' and aimed towards detecting the most common and widely applicable security issues.

Native software scanners

These scanning solutions are similar to their web application counterparts in that they are designed to identify common flaws in the construction and deployment of custom applications.

Unlike web application scanners, however, native software scanning solutions are designed to be run in an internal setting, often on the same host as the software product being evaluated, or somewhere with direct access to its source code. This enables checks to be performed that would not be possible from interacting with an external web application with limited network exposure.

Detecting and managing vulnerabilities within the software development process is beyond the scope of this guidance, however more information on this topic can be found here as one of our Secure Development Principles.

Comparing infrastructure and web application scanners

Type of vulnerability scanning	Associated assets	Examples of issues identified
Infrastructure	• Network infrastructure components • Physical hosts • Virtual hosts • End user devices • Cloud-based hosts or endpoints	• Missing Operating System or software application patches • Unsupported Operating Systems or software applications • Use of default or weak passwords • Use of weak cryptography or clear text services • Exposure of sensitive services or information • Lack of security hardening measures • Overly permissive access controls
Web application	• API endpoints • Web applications • Domains	• Injections from malicious user input • Broken authentication • Exposure of sensitive personal or system data • Broken access control • Use of vulnerable 3rd party components • Use of weak cryptography or unencrypted communications

4. Choose a deployment model

The marketplace for vulnerability scanning solutions and services includes both the traditional on-premises model and the increasingly popular vendor hosted model. You should choose a deployment model that best integrates with your infrastructure and satisfies your organisation’s security constraints

On-premises solutions

With an on-premises deployment, the customer is required to host the scanning product themselves on their own infrastructure. This could involve, for example, a virtual machine (VM) or physical appliance within your datacentre.

This type of deployment makes it far easier to scan areas of the network that do not have external network connectivity. Data is also stored locally in such deployments, therefore ensuring that you have full control over the location of any sensitive information relating to vulnerabilities in your systems.

However, the greater degree of administrative control comes at a price. Such deployments inevitably require some initial configuration and ongoing maintenance to ensure they stay up to date with the latest vulnerabilities.

In addition, on-premises solutions cannot scale easily to meet peaks in demand that may accompany scans of large portions of the IT estate at the same time. This can lead to the expense of maintaining excess capacity without the certainty that it may ever be needed. This problem is not specific to vulnerability scanners, but to on-site infrastructure hosting in general. As such, we recommend that on-premises solutions be used to scan systems that are not easily reachable from the Internet, or if your organisation already has on-site infrastructure hosting capability.

Vendor hosted solutions

Many solutions are also now offered as a service, whereby the scanning software remains hosted elsewhere, under the control and administration of the vendor.

This model is often referred to as Software as a Service, or SaaS. This can be a cost-effective way to overcome many of the shortcomings of on-premises solutions, however it also comes with its own disadvantages.

As an externally hosted service, it follows that SaaS scanners cannot easily access internal networks, situated behind firewalls and routers. This challenge can be overcome by installing agents on internal networks to form outbound connections to the vendor’s servers to receive instructions. If this is not possible, firewalls can be reconfigured to permit incoming connections from known scanners. This will inevitably involve a degree of initial configuration for network administrators, which may or may not be straightforward, depending on the structure of your IT estate.

Any changes made to your network to enable its scanning in this way will increase the level of risk to your organisation, due to the need to bestow a level of trust to the security scanning vendor. This should be clearly documented and factored into your security model.

Despite the considerations mentioned above, SaaS scanners also come with many advantages over on-premises solutions. The lack of an installed product or appliance removes the need to perform maintenance tasks, such as patching or updating of an internal vulnerability knowledgebase. In addition, SaaS solutions can typically scale to adjust to demand without the cost of permanently hosting unused capacity.

Finally, having your vulnerability scan results hosted by the vendor simplifies the task of applying your own protective measures to ensure such information remains confidential, whilst at the same time being accessible to those that need to see them (assuming you are content to trust the vendor's own controls).

We recommend using a hosted solution if the technical challenges and security concerns of allowing external access to your IT estate and having your organisation's vulnerability information retained by the vendor can be easily overcome. Such a solution would not be suitable for assessing an air-gapped network or one that holds highly sensitive information.

5. Decide which assets to scan, and when

Whilst it's true that greater coverage of your IT estate provides more complete visibility of your organisation's overall risk, it may not be practical or affordable to scan everything. In these cases, you should prioritise the assets that are Internet-accessible, host business-critical services, or contain the most sensitive data (e.g. database servers). It is important to maintain a record of the assets that are excluded from vulnerability scanning, to ensure the associated risks can be incorporated into your organisation's security model.

Extrapolating tests

Where multiple hosts have been built from a ‘golden image’ that guarantees the same configuration (as is common in standard desktop deployments) and no further changes have been made, it may be acceptable to scan a single host built from the image and extrapolate the findings across the other hosts.

Whilst vulnerability scanners are very unlikely to affect the availability of services or be otherwise disruptive, you may wish to consider first scanning non-production instances of servers that host business critical services. This will only produce results that are transferable to the live environment if the configuration of the two environments is consistent. Subsequent scanning of production systems should still be performed for cases where these configurations differ and where the scans were confirmed not to affect the availability of the non-production instances.

If you don’t have a representative non-production environment and have concerns about the fragility of certain business-critical hosts, it is permissible to temporarily omit these from potentially disruptive scans and record this in your risk register. This must be done with caution and for the shortest period possible, as doing so will create blind spots in your attack surface. The far better solution in such circumstances is to solve the underlying cause of the fragility so that the hosts can once more be scanned without fear of causing service disruption. Indeed, fragility should be seen as a vulnerability in itself and worthy of prompt remediation.

Scan regularly

You should aim to perform vulnerability scans of your infrastructure on a regular basis (at least once every month), or immediately after applying changes to remediate a critical issue.

Application scanners should be run any time there are changes made to the target application, such as when a new version is installed or when changes to the source code of your custom application have been committed. If possible, application scanning solutions should be incorporated into the software development process as part of secure build and deployment pipeline.

Additional considerations

There are many other things to consider when determining the suitability of vulnerability scanning services to your needs. Whilst it's often difficult to define an exact value for what 'good' or 'bad' looks like in each case, the following is a list of essential criteria that we recommend you ask prospective vendors to provide so that the answers may feed into your own evaluations:

- ***Responsiveness***: Can the solution detect a new vulnerability within a reasonable timeframe, once it has been publicly disclosed? This should be no more than a few days for critical issues.
- ***Coverage***: Does the scanner cover the categories of vulnerability that are relevant and important to you? For example, in the case of web application scanners, are all issues detected from the OWASP Top 10?
- ***Authentication support***: Does the scanner support authenticated checks? For example, is it able to log into Windows hosts to perform checks that are otherwise not available? Does it only support local authentication using an agent as well as remote authentication? Does it have safeguards in place to help prevent accounts getting locked out?
- ***Accuracy***: Does the scanner often produce false positives (where a vulnerability is reported to exist but doesn't) or false negatives (where a vulnerability exists but is not reported to)? For example, does it incorrectly identify old versions of software products or claim that patches have not been applied when you know they have?
- ***Reliability***: Is the scanner readily available to task both on an automatic schedule and manually on-demand?
- ***Scalability***: Does the scanner remain performant during periods of increased demand, with a cost model based on the required capacity at any one time?
- ***Reporting capabilities***: Can reports be tailored to suit your specific needs? Do the reporting capabilities provide the right information and metrics to support your security and administration teams?
- ***Support to other areas of a VMP***: Can the solution be easily integrated with your existing products or processes? Alternatively, does the solution provide

additional features above and beyond vulnerability detection that would supplement your existing VMP (for example, in-built issue tracking)?

- ***Integration with other Operating System components***: Can the solution provide additional value by leveraging existing installed software on target hosts? For example, does it integrate with Microsoft System Centre on Windows hosts to provide intelligent patch management capabilities?
- ***Support for different types of assets***: For example, does the solution support the scanning of virtual machines, containers or specialised database servers?
- ***Integration with cloud environments***: Can the solution interrogate common cloud providers to automatically discover and scan additional assets hosted in these environments?
- ***Safety***: Does the vendor guarantee that scanning activity will not disrupt the availability of the services being targeted? If not, is the solution configurable to exclude the more dangerous checks?

PUBLISHED

19 January 2021

REVIEWED

19 January 2021

VERSION

1.0

WRITTEN FOR

Small & medium sized organisations

Large organisations

Public sector

Cyber security professionals