OPTIV

Building an Effective Vulnerability Management Program:

IS IT IN YOUR FUTURE?

White Paper

John Ventura
Practice Manager, Research
Optiv

Introduction

Penetration testing is much easier than it should be. It is not uncommon for even mediocre penetration testers to show up at a client on Monday morning and have working administrative domain credentials shortly after lunch time that same day. Easily detectable vulnerabilities with published exploits linger in production environments for years waiting to be the initial foothold for anyone trying to climb over whatever network security safeguards may otherwise be in place. Although this description may seem like an exaggeration, it is all too real today, and the implications for any situation where actual attackers are involved should give nightmares to even the most cynical IT professionals.

However, some rare organizations stand out as being resistant to common penetration testing methodologies. Those companies are the ones that deploy comprehensive centralized vulnerability management technology along with related policies and practices. Although no turn-key solutions really exist, administrators can better protect their data by enhancing their security programs with proven techniques that other organizations have found effective.

Although no turn-key solutions really exist, administrators can better protect their data by enhancing their security programs with proven techniques that other organizations have found effective

An effective vulnerability management program (VMP) can provide an organization with relevant data that can inform decision making across a broad spectrum of IT security administration. Analysis of this data can reveal large programmatic gaps in security as well as in IT best practices in general. For example, vulnerability scanners excel at identifying gaps in patch management practices that might otherwise be much easier to ignore or miss completely. In addition to identifying missing patches or dangerous configuration settings, these programs also identify resources like rogue networking equipment or forgotten servers whose presence within an environment may violate policy. Furthermore, policy initiatives, like the mandated removal of unsupported operating systems from production networks, become much easier to enforce.

Although the benefits are remarkable, even the best vulnerability management programs are not substitutes for regular diligence. Human beings must still provide analysis of the data and act appropriately. The effective security practices outlined here will stop some attackers and even a good portion of the professional penetration testers, but dedicated human beings are still more effective in spotting unique vulnerabilities that computers are not yet equipped for.

This paper divides vulnerability management programs into three components:

Data Acquisition

This component includes any automated means by which administrators learn about the vulnerabilities that exist within their environment. Data sources that report from vulnerability scanners or desktop agents, assuming that their data is collected, comprise the core of this acquisition technology.

Information Storage and Analysis

The vulnerability data must be stored in a centralized location and analyzed in order to present any value. Typically, these systems include a centralized database, like a Security Information and Event Management (SIEM) solution or other database solutions. Administrators utilize these systems to perform automated analysis for more informed decision making

Accountability Engine

Once the vulnerability management team understands the issues that are present, policies and procedures must be in place to encourage remediation. In larger organizations, responsibility for systems' management will be spread across multiple groups and departments, and an accountability engine is an organized system of communication that encourages administrators to exercise their control over systems to address identified issues.

Data Acquisition

To analyze data and meet any security goals, administrators must first devise a means of acquiring data. Typically, large organizations will deploy a fleet of vulnerability scanners across their environment along with desktop/server agent software. These systems will perform a number of automated scripted checks across the network with IP enabled hosts with the intent of identifying misconfigurations or patching deficiencies that could enable attacks against the IT environment.

What are vulnerability scanners?

Vulnerability scanners are network-enabled devices that run a series of scripted checks that probe other network-enabled nodes, like servers, desktops, or printers. Each check corresponds to a known security issue. They require the scanner to send data to another host

over known protocol and perform logical operations with any data potentially returned from that host. The pre-programmed tests of that data that the scanner performs will be the determining factor in whether or not the scanner reports a particular vulnerability as being present with the scanned host.

For example, security professionals would most likely consider the presence of an easily guessable password, like "password," attached to an administrative account on a production server as a vulnerability. This exact issue is common enough that attackers would observe it from time to time across otherwise diverse environments. Because this issue is predictable enough to be repeatedly present, many vulnerability scanners will include a check to verify whether or not this password is present. If a scanner detects an accessible secure shell (SSH) server, it should have scripted instructions to attempt to log into this server via SSH with the username "root" and the password "password." The scanner would most likely have some sort of internal logic to determine whether or not this login attempt was successful. If the SSH service responds with an error indicating that the credentials are invalid, the scanner would process this response and determine that the vulnerability is not present. However, if the login attempt was successful, the scanner would determine that the server is vulnerable and report this issue. Scanners typically repeat this testing process with a set of other pre-programmed checks to identify a wide variety of issues.

In addition to checks against misconfigurations, like the use of obvious passwords, vulnerability scanners will also attempt to determine significant patching deficiencies. Sometimes, the checks will rely on banners presented by network protocols, which often reveal software build or version data. Other times, these checks rely on differences in responses to network-based input. If legitimate administrators supply working credentials to these scanners, they can even log into servers and query systems for software versions and provide an authoritative source for server software version data and patch management.

Purchasing Concerns

The following factors are necessary for the consideration of the network architecture and budgeting concerns for individual scanning hosts:

- Large or even medium sized organizations will need multiple vulnerability scanners with regularly reoccurring subscription costs.
- These scanners must have unfiltered access to the hosts that they are scanning.

- Scanners must also have a network path to update servers or an alternate means of updating their internal database of checks to protect against newly discovered vulnerabilities.
- Scanners must also have a network path that allows them to report data to responsible parties.

Depending on the size and scope of the environment as well as the network architecture, administrators will have to determine the number of scanning devices they need as well as their location on the network. Larger or more complicated environments require the deployment of more scanning devices due to network filtering and other network complexity issues, as well as the capacity of a network scanner. Any vulnerability scanner will need unfiltered network access to the hosts or networks it scans in order to insure the acquisition of reliable data.

Depending on the size and scope of the environment as well as the network architecture, administrators will have to determine the number of scanning devices they need as well as their location on the network.

Although deploying additional scanners inside otherwise restricted networks complicates the VMP, it also promotes stability and leads to more accurate results. Firewalls, by their nature, will prevent hosts from outside a segmented environment from accessing resources. More intelligent firewalls may even alter network traffic making some vulnerability checks useless.

Furthermore, allowing vulnerability scanners to assess remote hosts through a firewall disrupts the network environment and leads to gaps in data acquisition. Firewalls that perform any kind of stateful packet inspection allocate computational resources for any newly established communication that traverses them. Scanners can perform hundreds of checks against entire networks of servers, which puts a tremendous strain on the firewall's internal resources. Scanning through a firewall will put the continued operation of the network in jeopardy due to the threat of resource exhaustion. Even if scans have completed without noticeable disruptions or isolations of network segments due to firewall crashes, unusually high loads occurring during scanning intervals or changes in network usage can lead to unforeseen outages or unacceptable network slowdowns. Administrators responsible for vulnerability management programs must take into account the means by which these devices update their sets of checks. For example, if a scanning host in an "air gapped" network accesses an update server via HTTPS, this connection may violate established network security policies. Similarly, the network filtering exceptions needed to make this update possible may also enable data leakage or weaken the security posture of the environment. Conversely, although some scanner software vendors allow more flexibility in how they update vulnerability checks, any procedures for the upkeep and maintenance of these devices will include some network security exceptions or, at the very least, administrative costs.

Reoccurring Costs

Commercial vulnerability scanner vendors rely on a business model similar to virus scanner vendors. Anti-virus companies respond to how the threat from viruses and other malware changes as attackers devise new software by introducing new checks. Similarly, vulnerability scanner vendors continually update the sets of checks by adding new ones to reveal newly discovered vulnerabilities. When security researchers or vendors release advisories describing new issues, developers working for scanner vendors respond by expanding coverage. The scanners have automated procedures for updating their internal database with these new checks, and customers typically pay a subscription fee to insure that they receive the most up-to-date coverage.

Assuming that the scanners are commercially licensed, each instance will cost the consumer a licensing fee. Failures to maintain these licenses will prevent the organization from accessing newly published vulnerability checks and may prevent future scans from running entirely. Without continued support from vendors, vulnerability management programs cannot contribute to an adequate level of protection.

Appliances, Virtual Machines and Servers

Minimum hardware requirements for vulnerability scanners do not tend to be terribly prohibitive. Although administrators may wish for more parallelism or more performance generally from their scanner solutions, hardware limitations are rarely the limiting factor. Many organizations have functional scanning solutions that utilize virtual machines or servers with limited hardware. Some vendors will also lease scanning appliances that have the scanning software pre-installed. Each solution comes with its own benefits and liabilities that administrators must account for before implementation.

Virtual Machines

Assuming that an organization has a robust virtualization environment, administrators can choose to deploy scanning from virtual servers. This approach has the advantage of being easily replicated and relatively cheap to deploy and maintain. As long as the staff maintains the scanning and operating system licenses, this solution is relatively easy.

However, virtual machines introduce another set of issues regarding data handling and backups. Vulnerability scan data often reveals critical vulnerabilities that are easily exploited. Insecure backup solutions can reveal a catalogue of easily exploitable vulnerabilities to attackers in a convenient and easily digestible way. Similarly,

administrators must be aware of threats introduced by the sharing of resources. Hypervisor vulnerabilities can allow attackers who compromise other virtual machines to gain access to vulnerability scan data. Shared physical network interfaces allow attackers to monitor scans and may also allow them to gain access to networks that must be open to the vulnerability scanner but restricted for other resources.

Dedicated Hardware

The use of stand-alone hardware will address some of the shortcomings of virtualized environments, including concerns regarding network architecture and the sharing of resources. However, administrators will have to devise solutions for various problems, including backups, remote administration, and software maintenance. However, the problems associated with "racking" any server will remain, including physical space, power consumption, and general maintenance.

Large real world VMPs will typically use a mixture of dedicated hardware and virtual machines. When designing these solutions, it is important to take the recommended hardware requirements into account regardless of what path or mixture of paths is required. Vulnerability scanners do not typically require the most state-of-the-art hardware or a massive investment in any particular instance.

Scanning Appliances

Some vendors will lease dedicated scanning hardware. These arrangements remove the burden of hardware maintenance issues, including backups and software maintenance. These devices also may come with an additional security benefit as they will be relatively independent of the rest of the environment. Attackers who compromise a Microsoft Windows domain will not likely have administrative access to these appliances, because they will not be part of any corporate domain.

Although the benefits of these turnkey solutions are many, they also require organizations to yield control of some very sensitive information to third-party vendors. Details about the internal servers, including exploitable vulnerabilities, will be sent to these vendors and stored within their cloud infrastructures. Security administrators should also be aware that if attackers ever compromise these vendors, the presence of these appliances will expose every portion of the IT environment. If attackers penetrate the vendor, they will have unfettered access to sensitive servers as well as an easily digestible database of vulnerabilities to choose from.

If attackers penetrate the vendor, they will have unfettered access to sensitive servers as well as an easily digestible database of vulnerabilities to choose from.

Such an apocalyptic scenario may seem unlikely, but it also might be beyond the risk appetite of some organizations.

Scheduling Vulnerability Scans

The vulnerability management team must work with asset owners to arrange regularly scheduled scans. Scans should not coincide with times when load times are high or when availability is heightened priority. Network architecture issues are also relevant. For example, if the scanners reach the servers, that they are scanning across a link that is shared with other applications, responsible parties need to make sure that scan times do not conflict with periods where availability is of a higher concern.

Additionally, servers that are accessible externally, typically with a real routable Internet enabled IP address, as well as internally with a non-routable IP address, present related problems. Scans that take place over the Internet should not occur concurrently with internal scans of the same host. Scanning the same hosts or networks two or more times simultaneously can lead to instability and network saturation.

Policy Concerns

Before purchasing or deploying any data acquisition technology, administrators must decide on policy issues regarding the deployment and operation of these devices. In particular, the following questions are especially pertinent, and administrators should address these issues during the initial planning phases:

- **Network Architecture** – What network architecture policies are already in place and how can vulnerability scanners update their database of checks and report back to a centralized reporting solution?
- **Scanner Access** – Will a vulnerability scanner's placement on a network cause it to miss important data by the imposition of a firewall or similar security technology?
- **System Upkeep and Administration** – How can the organization make sure that systems are maintained, including questions of power consumption, physical space, and secure backups?
- **Scanning Schedule** – What significant load times exist, and how can scans occur during times where the threat of service disruptions is minimized?

Information Storage and Analysis

Once vulnerability scanners begin generating data, they must send this data to a centralized source in order to enable analysis and make any of these efforts valuable. Decentralized solutions are possible, assuming that administrators are willing to access each scanner on a regular basis. But, centralized repositories deployed along with effective reporting, will reveal critical vulnerabilities in individual servers along with overall vulnerability trends and inform decision makingy.

Centralized repositories deployed along with effective reporting, will reveal critical vulnerabilities in individual servers along with overall vulnerability trends and inform decision makingy.

Choosing a Back End

Several SIEM vendors already have modules written to allow them to assimilate vulnerability scan data. Organizations that already have SIEM solutions on site should consider the integration of this data. Reporting and decision making can be enriched by incorporating data from other sources, the formats for which may already be supported by the SIEM.

However, several organizations have found success by importing vulnerability scan data into their own database back ends. Scanners from major vendors provide data in formats that can be easily imported into custom built solutions. Anyone managing these custom-built solutions should keep in mind that they will require ongoing support. Data management solutions that are dependent on the continued participation of their developers are doomed to fail, if and when, these developers' employment status changes. Documentation in these situations is essential for program continuity.

Just as with the vulnerability scanning infrastructure, vendors will provide a variety of solutions for the storage of this data. Some cloud-based vendors will store this data offsite and provide sophisticated database functionality for its reporting and management. These vendors will also include essential management features for backend server components, such as backups and power management. However, these solutions also require administrators to send their data offsite and create dependencies due to the fact that the third party has an ever increasing volume of data and possibly even a proprietary data management/reporting API over which they will retain complete control.

Data Management Concerns

For whichever data management solution works best for any organization, administrators must make sure that they take into account the specific needs of a vulnerability management program.

Simply locking data away into a secure database and running occasional queries on an ad-hoc basis is not a comprehensive enough solution for most environments.

In order to provide value and support comprehensive remediation strategies, the backend database needs to support functionality beyond simply logging vulnerabilities. Vendors typically have unique IDs for every vulnerability check that the servers perform, and each check will correspond to a single reported vulnerability. For example, a check, or test that the vulnerability scanner performs to look for a specific issue, might have a unique ID of 0x01234, and this check will only be involved in reporting a single issue, like a blank administrative password. If this issue exists on multiple servers, the same ID number will be reported along with the IP address or other unique identifiers for the servers. Instead of indexing vulnerabilities by titles, administrators should use these IDs to identify individual issues.

These database servers can also be a useful repository for remediation strategies. When administrators successfully address vulnerabilities, they should be encouraged to submit a brief description of how they addressed these problems. For example, if the presence of plaintext remote console services, like Telnet, is considered a vulnerability, the remediation for that issue would typically be to disable Telnet. Responsible administrators should inform the vulnerability management team that they resolved the issue by disabling the service and briefly describe how they did that (for example, editing inetd.conf and restarting inetd). Over time, this database will become more and more complete, and the administrators who need to remediate similar vulnerabilities in the future can use this resource to guide their efforts.

The backend vulnerability management database server must also track the occurrence of these vulnerabilities as they are detected on individual hosts. The first time an issue gets reported as well as the latest time it gets reported for specific servers are relevant data and are necessary for any comprehensive analysis. If vulnerabilities persist across the environment even after they are reported repeatedly, this pattern of deficiencies could indicate institutional or operational problems that merit review.

Additionally, vulnerability entries should contain a field to denote the severity of the vulnerability. Vulnerability scanners will already record a rating based on some sort of vendor defined severity rating system. Typically, these systems rate vulnerabilities either on a numerical scale or with escalating categories like “low,” “medium,” and “high.” Administrators should not take any of these mechanically assigned ratings as authoritative. However, CVSS scores or custom

severity ratings defined by vendors can act as a beginning to the process of prioritizing remediation efforts.

Data Management Policy Concerns

Due to the nature of the data being aggregated and stored, this portion of the VMP presents specific concerns. This portion of the VMP specifically deals with the handling of sensitive data, and administrators should consult existing policies, or draft new policies, if none exist, for issues that include the following:

- **Data Transmission and Storage** – Occasionally, vulnerability scanners or client side agent software will collect artifacts that represent sensitive data. Additionally, these solutions will result in the existence of a catalog of exploitable vulnerabilities, which is of obvious concern. The storage of this data and its transmission from the scanners or agents must be encrypted, and policies must reflect this need.
- **Cloud Security Concerns** – Pre-existing policies may already address what data may be stored offsite. Cloud data storage concerns are not already addressed explicitly by policy, administrators should consider these questions before going forward with a third party data management solution, including those offered by vulnerability scanner vendors.
- **Severity Grading** – What system will this organization use to grade vulnerabilities? Will it be a numerical system, with numbers like 1 through 10 assigned, or a more qualitative system with categories like "high," "medium," and "low"? What are the criteria for these grades?

Accountability Engine

Once a steady stream of vulnerability data has been established, the vulnerability management team can begin the task of guiding remediation efforts. Severity ratings can help administrators prioritize remediation efforts, and subsequent scans can allow the vulnerability management team to verify that vulnerabilities have been addressed successfully.

Severity Grading

Ideally, administrators will come up with a policy that defines a system of "modifiers" to customize vendor supplied severity ratings. Numerical ratings, like CVSS, or qualitative rating systems (high,

medium, low, etc.) should be weighted based on priorities assigned by a set of modifiers defined by policies. For example, vulnerabilities discovered in networks that organizations expect to be Payment Card Industry (PCI) compliant may be more serious than those discovered in other systems, so adding 20 or 30 percent to the numerical severity rating or raising the qualitative severity rating from "medium" to "high" may be in order. Pre-defined policies should include these modifiers to reflect real world security concerns and the priorities of the organization being protected. Administrators can then use these ratings more accurately to prioritize their remediation efforts.

Notification

Administrators responsible for the maintenance of the various servers being scanned must be notified when significant vulnerabilities are detected. Typically, the vulnerability management team will set up an automated process to email administrators vulnerability scan reports that are restricted to the relevant assets. Depending on the predefined severity thresholds, vulnerabilities that get reported more than a set number of times should be escalated. Sometimes, administrators are unable or simply unwilling to address serious security issues. In these cases, policies should include escalation mechanisms. After two or three notifications, many organizations choose to include higher layers of management in the alerts. In these organizations, administrators who ignore relevant alerts will find their supervisors notified.

Exception Management

In the real world, vulnerabilities do not always lend themselves to practical remediation. Design flaws, lack of vendor support, or other extenuating circumstances exist, and policies should allow administrators to make exceptions. Administrators who can justify the continued presence of significant vulnerabilities or their own non-compliance with security practices should submit their justifications to the vulnerability management team for review. If the vulnerability management team and the administrators can agree that exceptions are warranted, the escalation procedures should cease and exceptions should be entered into the centralized database. This exception process must also account for the fact that some reported vulnerabilities will be false positives. By design, these solutions will over-report vulnerabilities instead of taking chances of missing significant security issues. Although false positives are rare, the exception process must include mechanisms for their identification and handling.

Accountability Policy

This portion of the VMP is the most focused on actual human behavior. Consequently, policies must define how people are expected to behave regarding the handling of vulnerability data and the remediation or tolerance of vulnerabilities.

Identification – How will the vulnerability management team identify the responsible people behind the systems that are being examined? Before the accountability engine can notify anyone, policies must exist to establish responsibility for these systems.

Escalation – After non-compliance becomes apparent due to a repeated pattern of notifications, how will the matter be escalated? If managers get involved, how are they identified?

Data Handling – How will vulnerability data be transferred? Although email is very common, it is often completely unencrypted during transmission and storage. These messages may therefore violate policies, and administrators may wish to use an encrypted web-based system or devise a new solution that respects data handling concerns.

Weighting – What causes a vulnerability's severity to be upgraded or downgraded? Are some categories of vulnerabilities more severe due to business models or other concerns?

Exceptions – Under what circumstances is the organization willing to accept that a vulnerability cannot be remediated? Who in the organization ultimately makes this judgement?

Conclusion

An operational vulnerability management program will allow administrators to ensure the continued security of their organization's assets. Although some planning and maintenance are involved, **the benefits of increased visibility and real actionable data will guide decision making in ways that will most likely justify the effort**. Vendor supplied solutions exist for many of the needs expressed within this article. However, the purpose of this article was not to evaluate or compare these solutions. Administrators would be wise to compare what technologies are available to assist in the various aspects of their programs.

1125 17th Street, Suite 1700
Denver, CO 80202
800.574.0896
www.optiv.com

Optiv is the largest holistic pure-play cyber security solutions provider in North America. The company's diverse and talented employees are committed to helping businesses, governments and educational institutions plan, build and run successful security programs through the right combination of products, services and solutions related to security program strategy, enterprise risk and consulting, threat and vulnerability management, enterprise incident management, security architecture and implementation, training, identity and access management, and managed security. Created in 2015 as a result of the Accuvant and FishNet Security merger, Optiv is a Blackstone (NYSE: BX) portfolio company that has served more than 12,000 clients of various sizes across multiple industries, offers an extensive geographic footprint, and has premium partnerships with more than 300 of the leading security product manufacturers. For more information, please visit www.optiv.com.

11.16 | F1