ESG Lab Review

Scanning One Million IP Addresses with McAfee Vulnerability Manager

Date: January 2014 **Author:** Tony Palmer, Senior Engineer and Analyst, ESG Lab

Abstract: *This report documents the results of ESG Lab's hands-on testing of the McAfee Vulnerability Manager MVM3100 Appliance. Over 1,000,000 public network IP addresses were scanned with a single appliance to demonstrate ease of setup and efficiency at scaling to large networks. A sampling of observed vulnerabilities, with lessons learned for corporate security teams, is featured in this review.*

The Challenges

Every network interface inside an organization or connected to it is vulnerable to a wide range of security exploits if it is not properly configured. Although compliance with corporate security policies is high, well-meaning employees may still unknowingly create openings for internal or external actors with malicious intent. With the advent of advanced persistent threats (APTs), even a brief window of vulnerability can have catastrophic long-term implications.

Figure 1. Frequency of Security Controls Effectiveness Testing

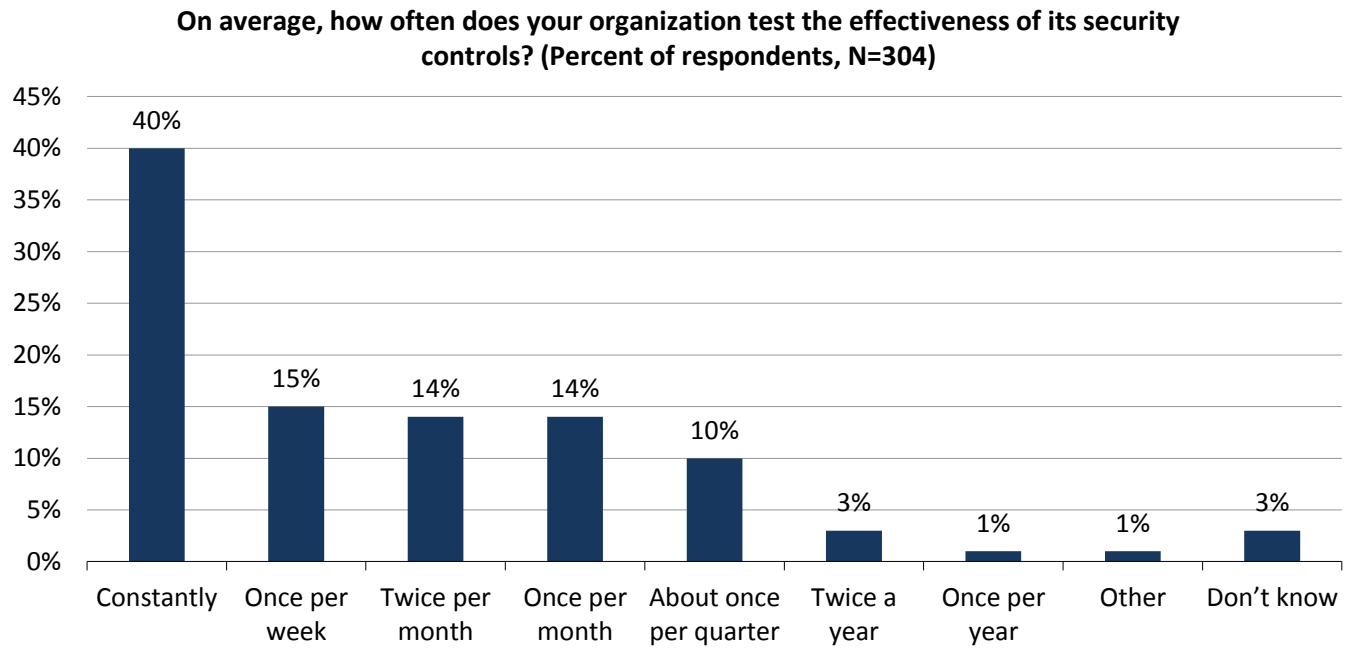

Source: Enterprise Strategy Group, 2013.

In an ESG research survey of large organizations (1,000+ employees), 58% of respondents reported that they use network/system scanning to test the effectiveness of their security controls and 48% reported scanning for rogue systems on the network, making those the two most popular methods reported by respondents. While some may consider this a high percentage, the bad news is that only 40% of the organizations reported testing the effectiveness of

their security controls constantly while 32% reported performing such work once per month or less often (see Figure 1).[1] Clearly there is a need for broader use of automated scanning solutions that are easy to deploy and capable of efficiently scaling to large enterprise networks.

The Solution: McAfee Vulnerability Manager

McAfee Vulnerability Manager (MVM) is designed to deliver scalable vulnerability scanning. MVM is available as software or bundled with a hardened OS in an appliance. MVM's configuration flexibility allows IT teams to distribute components of the solution inside and outside the corporate network, to scale horizontally, and to isolate scanning to particular network segments for traffic control. In its simplest configuration, all components can reside on a single rack-mounted appliance. Regardless of the configuration chosen, all findings are securely routed to a single reporting instance. Integration with other McAfee offerings is extensive, supporting a unified assessment of risks.

Figure 2. Deployment Options – Single Appliance or Multiple Distributed Systems

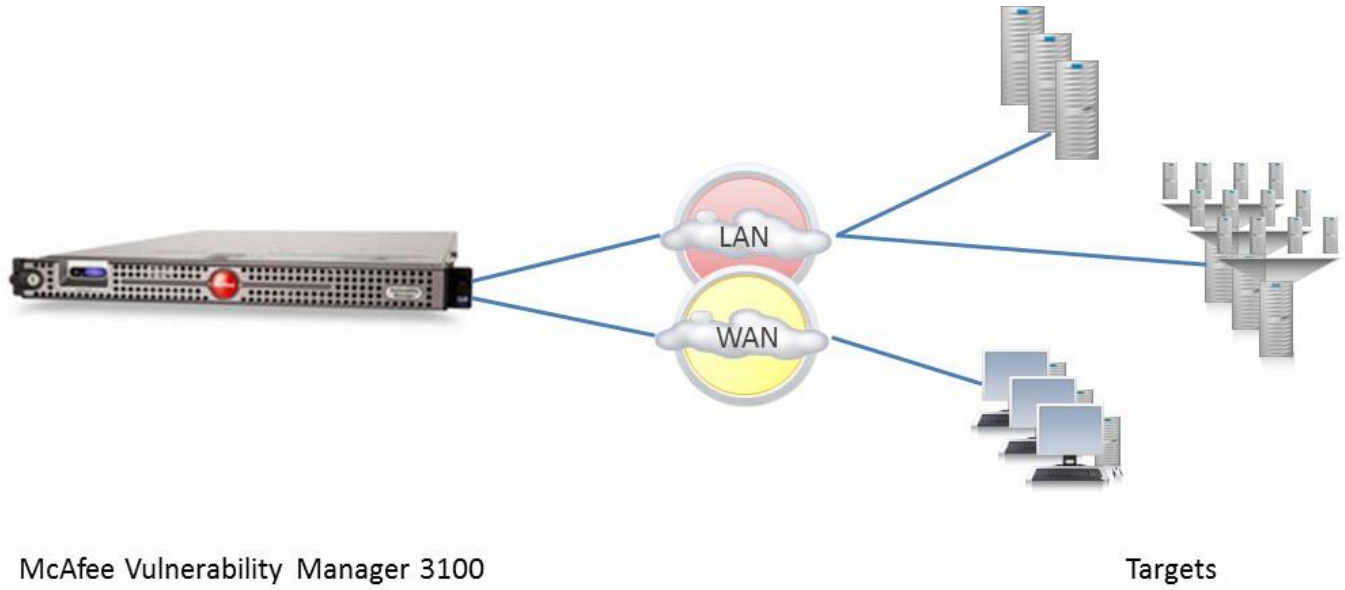

Frequent updates of content by McAfee Labs help ensure that the latest vulnerabilities are captured in assessments. Pre-configured templates and a best-practices guide address common assessment goals, such as quickly scanning for SANS Institute/FBI Top 20 Vulnerabilities, exhaustive "full" vulnerability scanning of all assets, or initial asset discovery. Regardless of the scan type chosen, intelligent defaults pre-populate most selections, speeding time to value.

[1] Source: ESG Research Report, *Security Management and Operations*, July 2012.

Figure 3. Pre-configured Scan Templates

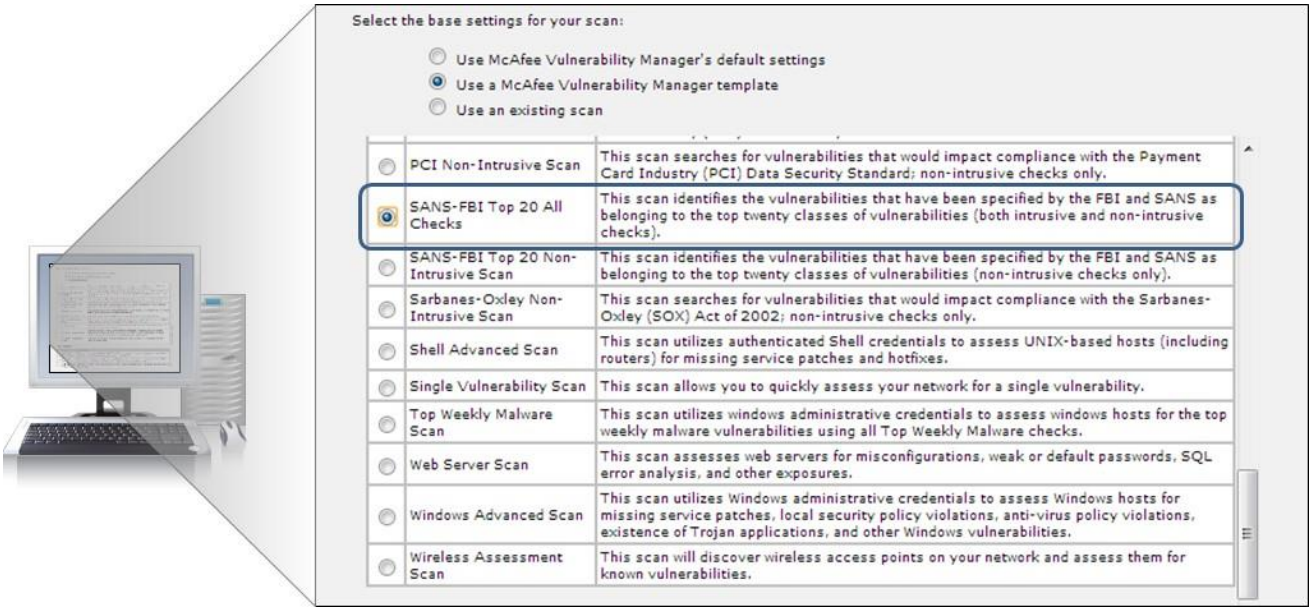

The MVM solution goes beyond being just a network scanner, incorporating features that directly address real-world business concerns, such as:

- **Deep Common Vulnerability and Exposure (CVE) coverage** – The Common Vulnerabilities and Exposures (CVE) system provides a reference-method for publicly known information-security vulnerabilities and exposures.
- **Rapid response time for critical events** – such as: Zero day vulnerabilities, "Patch Tuesday", and high impact, destructive vulnerabilities like Flame, PCAnywhere, NightDragon, etc.
- **Priority-based auditing and remediation** – Combines vulnerability, severity, and asset criticality information to quickly identify, rank, and address violations and vulnerabilities on networked systems and devices.
- **Active and passive network discovery** – Combines active and passive network discovery and monitoring to reveal virtualized, mobile, and hidden devices on your network.
- **Proof of "not vulnerable"** – A major requirement of auditors is to prove that you're not vulnerable to threats, which McAfee Vulnerability Manager provides as evidence directly in reports.
- **New threat identification and correlation** – Automatically ranks the risk potential of new threats by correlating events to your asset and vulnerability data.
- **Policy auditing and compliance assessments** – Defines values of policy checks and determines whether your organization complies with major regulations. Through an easy-to-use wizard, it gives you templates for SOX, FISMA, HIPAA, PCI, and more.
- **Flexible reporting** – Categorizes data by asset or network, and uses powerful filters to select and organize results in your reports. You can even create reports while scans are running.
- **Broad and deep reach** – Performs authenticated and unauthenticated checks, delving deep into operating systems and network devices to find vulnerabilities and policy violations.
- **Real-time Assessment** – New devices are scanned and assessed the moment they attach to the network.

ESG Lab Tested

ESG Lab performed hands-on evaluation and testing of a McAfee Vulnerability Manager appliance at a network colocation facility in Northern California. Testing was designed to demonstrate ease of setup, scalability, and effectiveness of the appliance. More than one million public IPs were scanned with a single MVM appliance, revealing some surprising findings about the state of unprotected network ports on the public Internet.

The configuration used in ESG Lab testing is shown in Figure 4. A single McAfee Vulnerability Manager MVM3100 appliance was placed on the public Internet. All components of the MVM solution were resident on this single rack-mountable system.

A Cisco Catalyst 3550 layer three switch was installed between the appliance and the Internet. Access control lists (ACLs) on the switch were employed to ensure that only ESG Lab and McAfee teams had access to the administrative interfaces on the appliance. These interfaces are already secured by authentication and encryption; the use of ACLs limiting access to IPs controlled by ESG Lab was an additional precaution given the direct exposure to the public Internet. These steps took less than ten minutes and required no knowledge specific to the MVM appliance. With the appliance racked and connected to the network, this completed the one-time setup steps.

Figure 4. The ESG Lab Test Bed

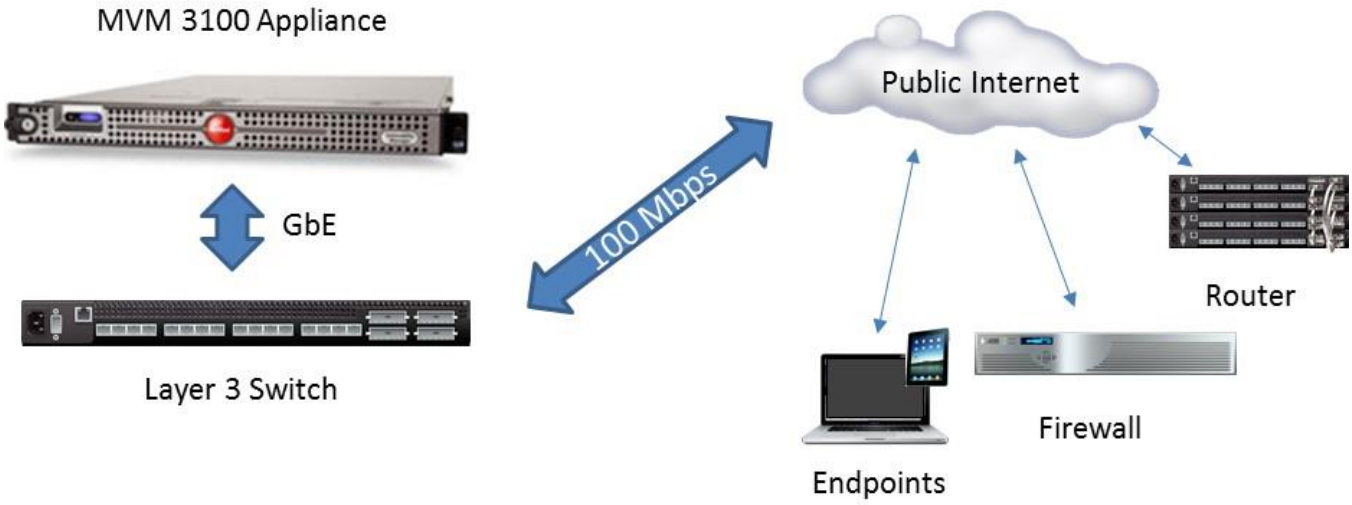

Ease of Setup and Management

ESG Lab explored the ease of setup of the MVM 3100 appliance. This included physical installation, basic network configuration, and setup of the scans to be run later in the lab. Because we were using an evaluation unit, the appliance arrived with both MVM and database software preinstalled.

ESG Lab reviewed the installation instructions and found them to be straightforward and well documented. It's worth noting that the MVM installer will run a "system check" for dependencies and will require all unmet dependencies to be addressed before allowing the install. This should not be an issue with McAfee-provided appliances, but reinforces the need for proper planning if an organization is opting to use existing hardware and OS.

ESG's focus in this testing was on the deployment steps, which occur after installing software on a lab bench: physical installation and network configuration in the targeted network location (one-time setup steps), and choice of scan templates and IP target ranges (ongoing setup steps performed each time scan policy is revisited).

ESG Lab Testing

ESG Lab next moved on to test the typical steps performed each time scan policy is revisited (choosing IP range targets and scan parameters). Administrators connect to the MVM Enterprise Manager Web interface via SSL to set up, monitor, and report on vulnerability scans. The home page of the Enterprise Manager is configurable and most tasks are initiated directly from drop-down menus on this page. ESG Lab chose "Scans," and then "New Scan" to get started.

The first step of configuring a scan is to define the scan targets. ESG Lab divided the world into five geographies and obtained IP ranges for the countries in each region from a commercial service.

MVM provides flexibility and leverage in defining the target IPs or systems to scan. The administrator can choose to enter Hosts, IPs, or URLs one at a time; provide a range of IP addresses by data entry or file import; or simply enter an entire CIDR block. ESG Lab used the Import option, browsed to the text file produced by the commercial IP range service, and clicked "Import." The imported lists contained as many as 50,000 IP range pairs per region.

The second step in configuring a scan is to choose the vulnerabilities to search for. If a template is chosen when creating the scan, this step is pre-populated but editable. Since the IP ranges were owned by neither ESG nor McAfee, ESG Lab and McAfee made edits to accommodate the public nature of this scan. The goal was to define and follow a clear set of ground rules:

- No use of credentials
- No delivery of payloads
- Assess for vulnerabilities, but do not exploit them

Vulnerabilities are grouped into categories and a "Hide Unchecked Values" option eliminates visual clutter when reviewing settings.

Figure 5. Editing Settings – Vulnerabilities to Scan for

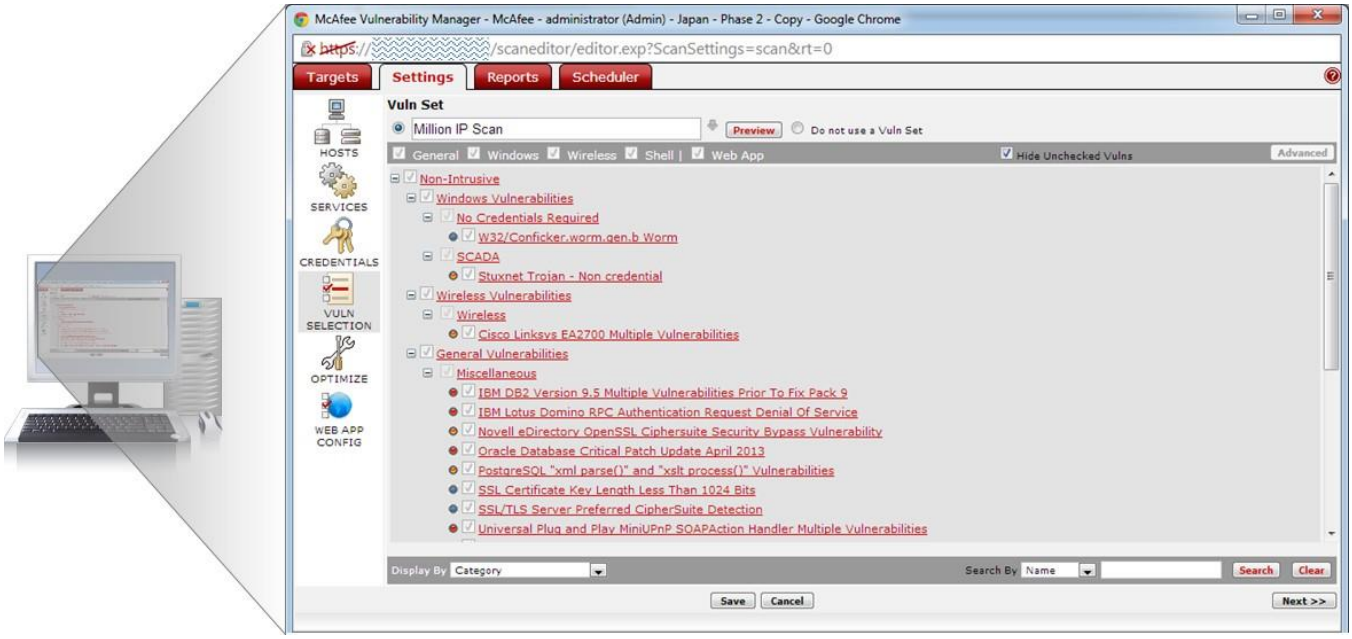

The third step in configuring a scan is to define the reports to be automatically generated as the scan completes. These reports can also be manually generated at any time, but defining them as part of the scan configuration supports a more proactive approach and a standardization of deliverables.

Figure 6. Defining Reports

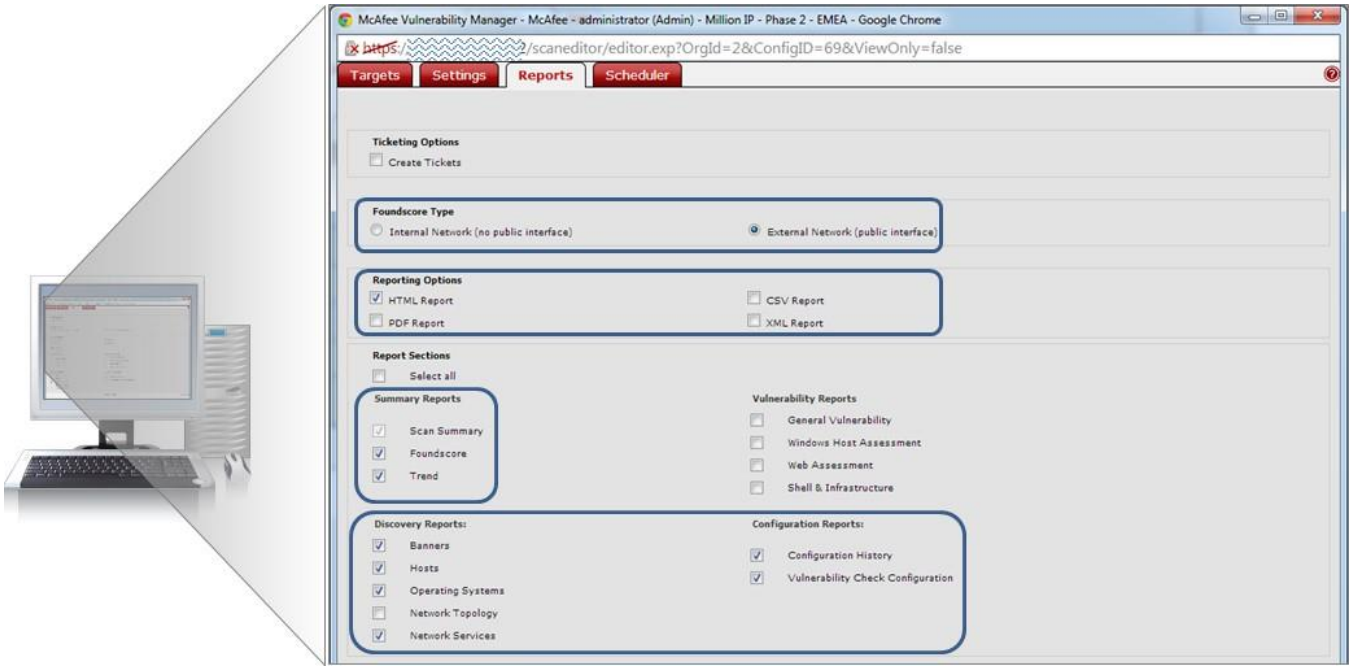

The final step in configuring a scan is to schedule the scan, whether it will be a one-time or recurring scan. ESG Lab launched the one-time scans for this testing immediately by making each scan "Active," choosing the Schedule Type of "Immediate," and clicking "Save and Scan Now."

Figure 7. Scheduling the Scan

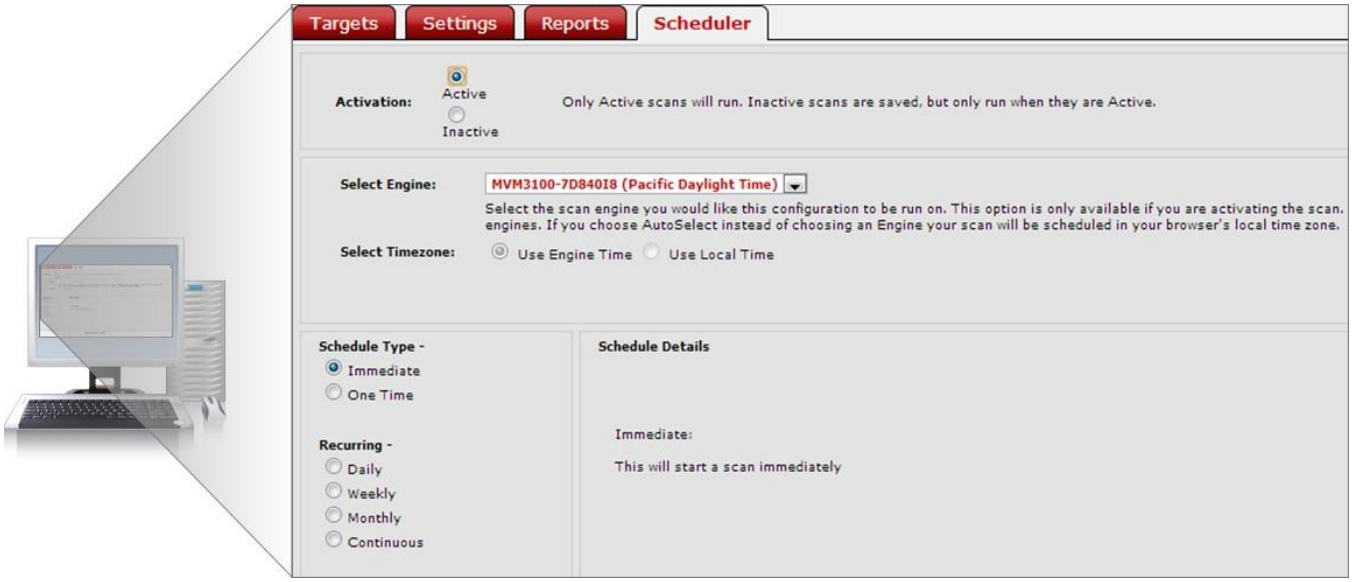

Why This Matters

While 55% of organizations surveyed by ESG stated that they were planning on increasing security headcount, finding information security help is becoming increasingly cumbersome: Nearly one-fifth of large organizations claim that it is extremely difficult to recruit/hire security professionals, while another 65% say it is somewhat difficult to recruit/hire information security professionals.[2] A risk-based scanning solution needs to minimize complexity and speed deployment time to enable organizations to quickly find and prioritize vulnerabilities without requiring specialized expertise.

ESG Lab was able to deploy and configure the MVM3100 appliance quickly and easily with no advance training. Ongoing management was simplified by the built-in templates and a well-structured scan setup workflow. Implementation was simple enough that it could reasonably be performed remotely with no expertise needed beyond setting the IP address in the OS. This enables IT to begin providing deep, targeted, and customizable vulnerability management services within minutes of powering on.

Efficiency at Scaling to Large Networks

The results of vulnerability scans are extremely perishable. For maximum value, the scanning infrastructure must be capable of frequent passes on *all* assets. Although a quarterly or monthly scan may meet the needs of some compliance-focused use cases, if the goal is securing the network from threats, then frequent or continuous scans are required. Efficiency in scaling up to large networks is a critical requirement.

McAfee's MVM architecture scales horizontally with multiple scan engines and scan controllers to distribute the workload and reduce network impact. But how well can each of those instances scale to inventory the network segment it is focused upon? By pointing a single scan engine at an extremely large target surface (the entire public Internet), ESG Lab sought to validate that McAfee's approach was inherently scalable.

ESG Lab Testing

ESG Lab launched all five of the geographic scans configured in the previous section within moments of one another.

The scans began with a status of "Pending" to indicate they were in queue to launch. Once scanning was underway, the status changed to "Running" and progress statistics were viewable for each scan.

The second through fifth scans remained "Pending" for nearly an hour. In a single-appliance configuration, this is by design; CPU priority is given to the first scan, which slows the IP address snapshots being made to the database as the subsequent scans prepare to launch.

ESG Lab monitored the progress of each scan on the Scan Status screen. Expanding the summarized scan status for each scan revealed the number of network devices discovered and assessed for vulnerabilities.

[2] Source: ESG Research Report, *Security Management and Operations*, July 2012.

Figure 8. Monitoring Scan Status

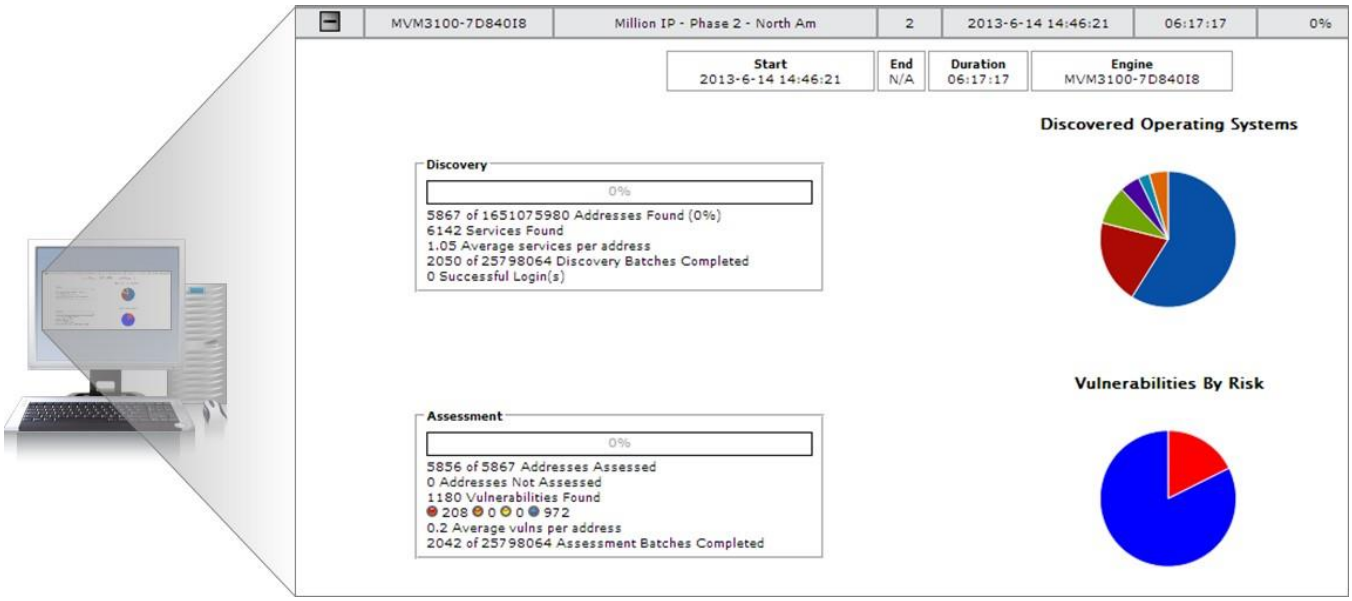

In order to ensure that each geographic region was equally represented in the lab findings, each of the five scans was paused after it reached approximately 200,000 discovered and assessed addresses. Some of the regions completed much sooner than others. For instance, Asia Pacific (not including Japan) and Latin America each reached approximately 207,000 and 213,000 assessments in just 96 hours. Europe, Middle East, and Africa (EMEA) crossed the 213,000 line in 119 hours. The two scans that ran the longest were Japan and North America, which reached 207,000 and 205,000 IPs around the 159-hour mark.

The overall goal of assessing 1,000,000 IPs was surpassed in approximately 150 hours. When the last two of the five scans were paused at the 159-hour mark, the MVM3100 had assessed an aggregate total of 1,046,553 IPs. In the process of discovering and assessing these IPs, ESG Lab estimates the MVM3100 attempted to connect to as many as six million IPs. (This estimate is based on an earlier random scan of 10,000 IPs where just 16% were reachable.)

Figure 9. Scan Progress by Hour

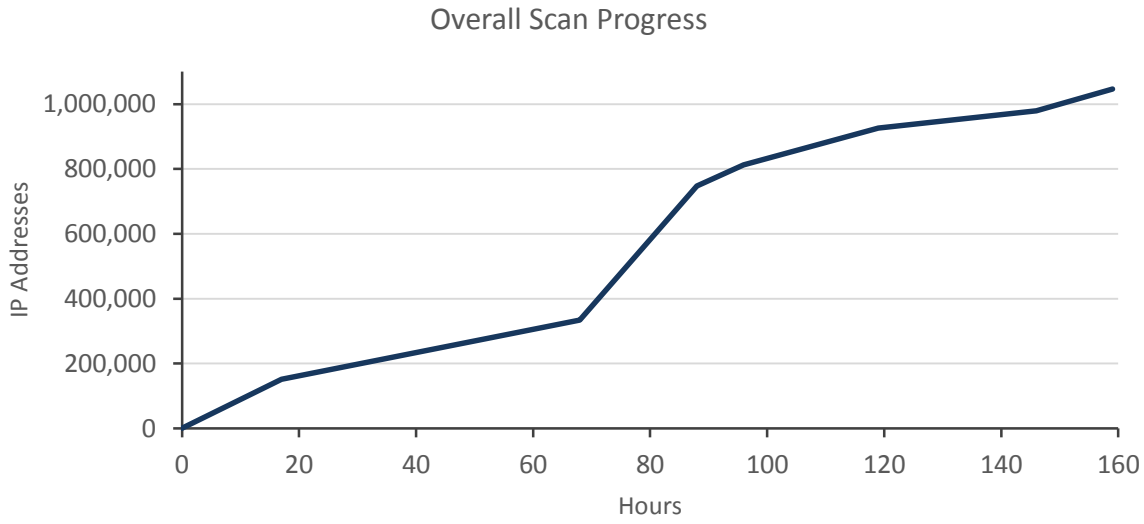

Why This Matters

According to ESG research, large organizations assess their security management capabilities using a number of metrics including the number of security events discovered, the number of security/IT audit violations or failures, and the number of vulnerable systems discovered.[3]

ESG Lab confirmed through hands-on testing that a single McAfee Vulnerability Manager appliance operated with the performance and efficiency required to deliver fresh result data (and the related security visibility), even at very large scale.

Uncovering Real World Exposures

ESG research found that protecting sensitive data/intellectual property was the most cited factor influencing information security strategy.[4] Quickly and accurately identifying vulnerabilities is a critical step in securing organizations' data. While the target of the scans for this testing was the public Internet, several surprising findings are relevant to the corporate security administrator. Three of the most notable vulnerabilities observed are summarized in Table 1.

Table 1. Observed Vulnerabilities

Security Exposure	Percentage of Hosts	Why It Matters
Outdated Secure Shell (SSH)	5.4%	Use of Secure Shell (SSH) is often considered a best practice for secure access to systems over the network. The problem is that outdated versions of SSH have numerous vulnerabilities, which introduce risk of denial of service (DOS) attacks, unauthorized disclosure of information, and other compromises of security.
Database listeners not filtered by firewall or access control list (ACL)	4.5%	An open TNS listener port allows remote attackers to execute arbitrary database commands by performing a remote registration of a database instance or service name that already exists, and then conducting a man-in-the-middle (MITM) attack to hijack database connections, also known as "TNS Poison."
Remote management cards not filtered by firewall or access control list (ACL)	2.3%	A remote management card allows users to monitor and manage servers and other network equipment by remote control regardless of whether the machine is powered on, or whether an operating system is installed or functional. In addition to being able to reboot, shut down, or power on a machine, users are able to broadcast video output to remote terminals and transmit input from a remote keyboard and mouse. They may also access local media like a DVD drive or disk images, and perform remote installations of software. The card vendors themselves warn that such cards are intended to be on a separate management network; they are not designed or intended to be placed on or connected to the Internet. Doing so could expose the connected system to security and other risks.

[3] Source: ESG Research Report, *Security Management and Operations*, July 2012.
[4] Source: ESG Research Report, *Security Management and Operations*, July 2012.

The Bigger Truth

While corporate security policies exist to protect organizations' networks and ESG has found compliance to be high, openings for internal or external actors with malicious intent can still be unintentionally created by well-meaning employees. Organizations' network interfaces–both internal and outward-facing–are potentially vulnerable to numerous security exploits if not properly configured. The ability to frequently and efficiently scan for these vulnerabilities is critical. With the advent of advanced persistent threats (APTs), even a brief window of vulnerability can have severe and lasting implications.

McAfee Vulnerability Manager (MVM) is engineered to enable IT teams to distribute components of the solution inside and outside the corporate network, to scale horizontally, and isolate scans to individual network segments for traffic control, providing enterprise-scale vulnerability scanning. Extensive templates regularly updated by McAfee Labs provide out-of-the-box support for broad and deep scans, with and without authentication credentials. Integration with McAfee ePolicy Orchestrator (McAfee ePO) software, other McAfee products, and McAfee Security Innovation Alliance partners is designed to reduce cost and optimize organizations' security posture, supporting a unified assessment of risks.

ESG Lab installed an MVM3100 appliance, which packaged all of the MVM components along with a database and hardened OS on a data-center-class system. One-time setup was minimal, requiring no special skills. Configuration of the scans was supported by a simplified workflow that prompted for targets (IPs to scan), vulnerabilities, reporting deliverables, and a schedule. Once scans were launched, the single appliance opened as many as 1,000 concurrent connections, leveraging significant parallelism to discover and assess over one million IPs in less than a week.

It should be noted that what ESG Lab tested is not the typical use case in a corporate environment. This was by design because we wanted to test the raw limits of the appliance–just how many addresses we could scan in a certain period of time. In a corporate setting, there would be additional considerations to contend with, such as isolating multiple separate subnets, and spanning both sides of a firewall. The delay we observed when simultaneously launching five massive scans on a single appliance would likely not occur at smaller scale or with horizontal scaling across multiple scan controller/scan engine pairs.

In closing, ESG Lab has validated that MVM can enable organizations of any size to simplify and accelerate scans of their networks no matter how large or how distributed. Analysis of vulnerabilities using the threat correlation module reduced the time required for impact assessment of new threats related to network access. In summary, McAfee offers a solution designed to enable enterprises to build a robust, cost efficient compliance ecosystem. Any organization that needs a clearer picture of the potential vulnerabilities in their networks would do well to take a close look at McAfee Vulnerability Manager.