



TENABLE
Network Security®

Firewall and Boundary Auditing

April 25, 2011

(Revision 3)

Copyright © 2011. Tenable Network Security, Inc. All rights reserved. Tenable Network Security and Nessus are registered trademarks of Tenable Network Security, Inc. The ProfessionalFeed is a trademark of Tenable Network Security, Inc. All other products or services are trademarks of their respective owners.

Table of Contents

Introduction	3
Tenable Building Blocks.....	3
<i>SecurityCenter</i>	<i>3</i>
<i>Active Scanning with Nessus.....</i>	<i>3</i>
<i>Passive Scanning with PVS</i>	<i>3</i>
<i>Vulnerability Repositories</i>	<i>5</i>
<i>Scanning Zones.....</i>	<i>5</i>
<i>Host Inventory.....</i>	<i>5</i>
Firewall Scanning.....	6
Pre-scanning Considerations	6
Firewall Scanning Example Scenarios	7
<i>Auditing One Firewall With Active Scanning</i>	<i>7</i>
<i>Auditing One Firewall with Passive Scanning</i>	<i>9</i>
<i>Auditing a DMZ with Active Scanning</i>	<i>10</i>
<i>Auditing a DMZ with Passive Scanning</i>	<i>12</i>
<i>Auditing Star-Topology Enclaves with Active Scanning</i>	<i>13</i>
<i>Auditing Star-Topology Enclaves with Passive Scanning.....</i>	<i>15</i>
Conclusion.....	15
About Tenable Network Security.....	17

INTRODUCTION

Does your organization audit firewall configurations, network boundaries or perimeter access control? Performing audits on complex networks with multiple control points can be manually intensive, difficult to model, prone to human error and too time-consuming to provide relevant security guidance.

Although Tenable's products are not specifically designed to audit network boundaries, they can be deployed in such a manner to perform near real-time analysis of the actual boundary as well as any trust relationships between different enclaves of large networks.

This document describes several simple strategies to leverage multiple Nessus scanners or multiple Passive Vulnerability Scanners with the SecurityCenter to automate boundary auditing. It also describes several tradeoffs to perform audits of multiple Internet and Intranet boundaries that are secured with firewalls, routers, intrusion prevention systems and other types of network policy enforcement devices.

Most importantly, these audits can be performed completely independent from analyzing actual firewall configurations, attempting to model complex firewall rule sets, communicating with the firewalls or communicating with the firewall management team. This is an ideal control for an independent audit and verification that the firewall policy in place is the one that should be in place.

TENABLE BUILDING BLOCKS

There are several key concepts to Tenable's Unified Security Monitoring architecture that enable auditing of multiple perimeter boundaries.

SecurityCenter

Tenable's SecurityCenter enables customers to easily measure vulnerabilities and discover security problems, asset by asset. SecurityCenter also helps manage asset discovery and correlates all of the information gathered from active and passive scanning with enterprise-wide log data to provide a comprehensive view of system and network activity across the enterprise.

Active Scanning with Nessus

Tenable Network Security's Nessus® vulnerability scanner is the world leading active vulnerability scanner, featuring high-speed discovery, asset profiling and vulnerability analysis of your organization's security posture. Nessus scanners can be placed behind firewalls, within enclaves, within discrete networks, inside a DMZ, outside of a DMZ and many other locations. Nessus does not care if the targets it is scanning are behind a firewall or it is in the same LAN subnet.

Each SecurityCenter can manage up to 512 Nessus scanners. This enables you to add more Nessus scanners to your topology without needing a larger license.

Passive Scanning with PVS

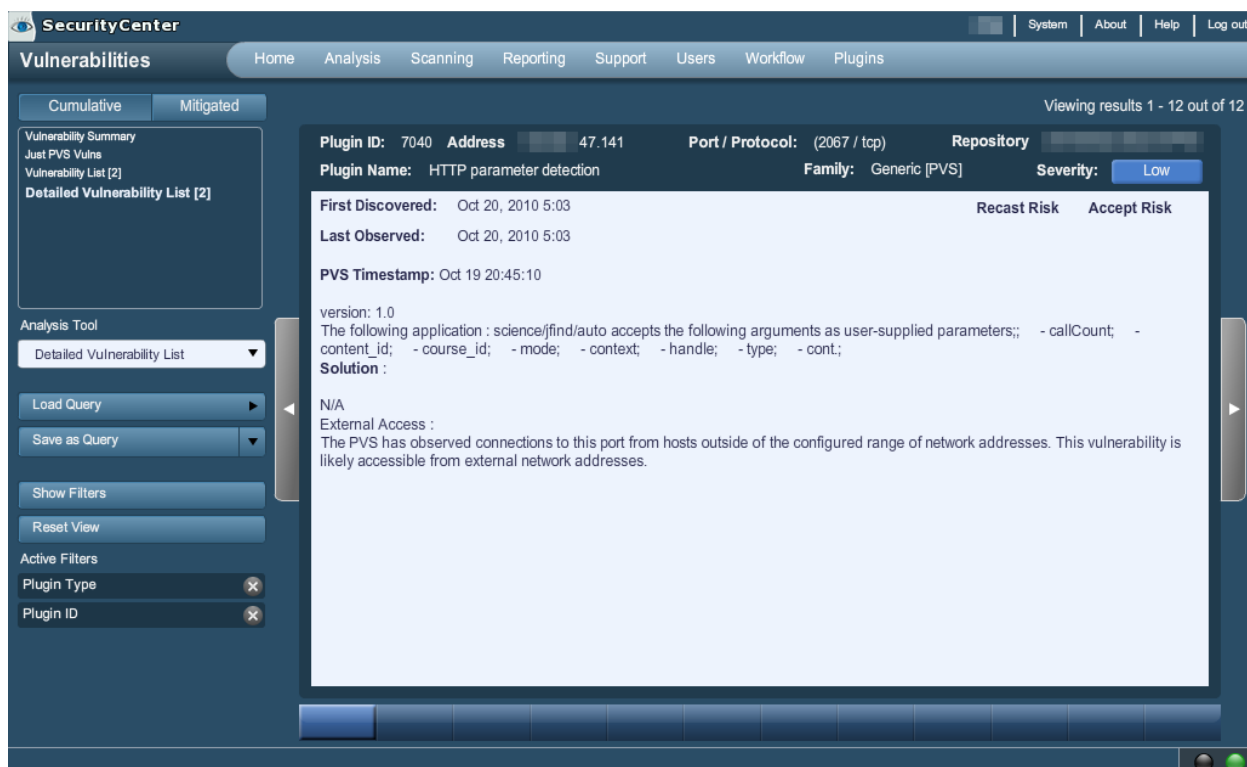
Tenable Network Security's Passive Vulnerability Scanner (PVS) is a network discovery and vulnerability analysis software solution that delivers continuous real-time network profiling and monitoring in a non-intrusive manner. PVS identifies every actively communicating host, enumerates all active services, itemizes all browsed ports and lists all client and server vulnerabilities. Each PVS is configured with the network addresses it is monitoring. When

reporting a vulnerability, if the source IP address of the observed network connection is outside local network, the vulnerability is automatically tagged as being externally facing.



In the image shown above, the PVS is configured to monitor traffic for an enclave. It observes PC1 log into SSH1 via secure shell across the boundary as well as PC2 log into SSH2 within the enclave. The PVS has been configured with the network ranges of the enclave and PC1 has an IP address that is not part of that range. The PVS would report on any SSH vulnerabilities found on SSH1 or SSH2 but it would also tag the vulnerabilities on SSH1 as being externally facing.

For example, in this screen capture of a PVS report shown within SecurityCenter, a sanitized IP address has been shown to be hosting a web form on port 2067 and the "External Access" reference has been added to the end of the report.



SecurityCenter | System | About | Help | Log out

Vulnerabilities | Home | Analysis | Scanning | Reporting | Support | Users | Workflow | Plugins

Viewing results 1 - 12 out of 12

Plugin ID: 7040	Address: 47.141	Port / Protocol: (2067 / tcp)	Repository: [Redacted]
Plugin Name: HTTP parameter detection	Family: Generic [PVS]	Severity: Low	
First Discovered: Oct 20, 2010 5:03		Recast Risk	Accept Risk
Last Observed: Oct 20, 2010 5:03			
PVS Timestamp: Oct 19 20:45:10			

version: 1.0
 The following application : science/ffind/auto accepts the following arguments as user-supplied parameters;; - callCount; - content_id; - course_id; - mode; - context; - handle; - type; - cont;
 Solution :
 N/A
 External Access :
 The PVS has observed connections to this port from hosts outside of the configured range of network addresses. This vulnerability is likely accessible from external network addresses.

The "External Access" reference can be used to quickly filter on any vulnerabilities found by a PVS that were observed to be accessed by clients from outside of the enclave.

The PVS is licensed by instance, Class C or Class B network. Within the Class C or Class B network license models, you can deploy as many PVS sensors you need to at no extra cost.

The “Firewall Scanning and Example Scenarios” section of this document compares the usage of deploying a PVS at the core of an enclave or on the perimeter. Some Tenable customers deploy multiple PVS sensors to have a clean capture of internal traffic or boundary traffic. Others only deploy one way. Depending on how your PVS is deployed affects how it can collect information on which ports, servers and applications are allowed through a boundary.

Vulnerability Repositories

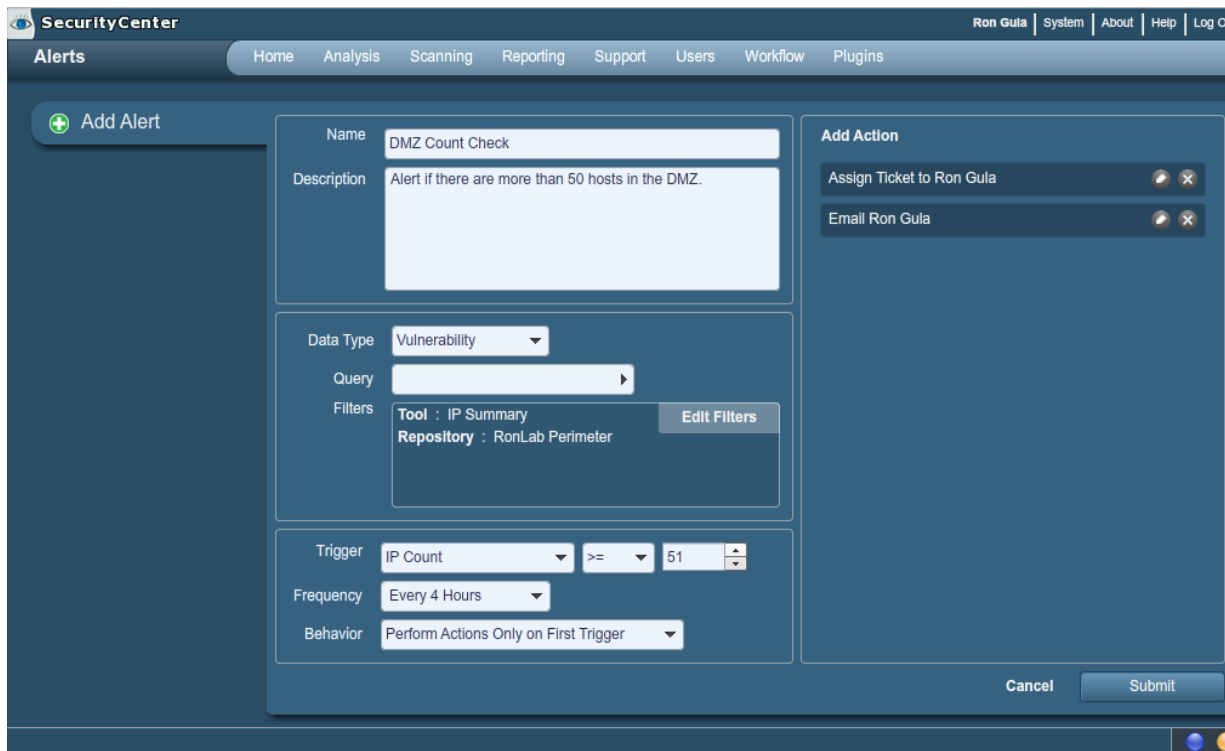
SecurityCenter can track scan results into one or more repositories of vulnerability data. Tenable customers typically leverage separate repositories for different organizations, different types of scans or audits of different types of compliance standards. For example, it is useful to have patch audits in one repository and unauthenticated vulnerability scan results in another. Repository data can be viewed as an aggregate, an ad hoc list or just one. Administrators can leverage this feature to create groups of repositories that represent boundary vulnerability scans distinctly from internal scans.

Scanning Zones

SecurityCenter can also group one or more Nessus scanners into a zone. These zones are logical groupings that can be independently assigned scan jobs. For example, you can tell SecurityCenter to perform a scan of some network addresses and have it use the “default” zone or you can override this by telling it to use the scanners from a different zone. Users can leverage this ability to create internal network scans and then use the same group of scanners to scan the boundary of another enclave.

Host Inventory

SecurityCenter can make use of any boundary scan (active or passive) results to count the total number of open ports or number of hosts. In the screen capture below, the SecurityCenter has been configured to alert if the number of IP addresses found in the DMZ from a particular scan is larger than 50.



Depending on how sophisticated you would like to be in your alerting, you can schedule alerts for many items such as:

- > Greater than or less than value of open ports or hosts
- > Not equal values of expected number of ports or hosts
- > Alerting on port counts with filters for specific port
- > Alerting on port counts with filters for ranges of ports such as greater than 1024

Note that for the purpose of this document, the term “firewall” is used to represent any type of boundary device that implements network filtering policies. This could be your router with some access control lists, an IP filtering intrusion prevention system, a virtual switch with packet filtering, etc.

FIREWALL SCANNING

Network boundary devices such as firewalls are designed to be sensitive to specific traffic patterns to block and log prohibited connection attempts. This aspect creates certain challenges to scanning such devices. It is important to configure scans in a manner that gathers accurate data without causing adverse reactions. This section describes pre-scanning considerations and provides some example scenarios for scanning through firewall devices.

PRE-SCANNING CONSIDERATIONS

If the firewalls you are scanning are busy, if they don't have enough available network bandwidth or if they are configured to log excessively, your scans may have an impact on performance and availability.

Before implementing any of the following techniques, start with less aggressive scans with the full knowledge of your firewall team so they can give feedback.

Tenable has numerous customers who perform full port scans across firewalled boundaries with no impact on availability or performance of the network. However, we also have some customers whose firewalling infrastructure is not as robust and have had to scale back the aggressiveness or thoroughness of their scans.

When selecting your policy to audit, consider the following issues:


- If you expect to perform a full 65,535 TCP port scan of each host, your scanners and firewalls will have to carry at least one connection attempt on each port. For 1000 potential hosts, this could lead to more than 65 million potential connections.
- It is possible to force Nessus to consider all hosts alive and thus perform a full port scan. By default though, it will attempt to only port scan hosts that are alive.
- Your firewalls may allow a connection on a port but the scanned host may not have that port open and show "closed" in your actual scan results.
- If you are interested in one type of port, such as port 22 for SSH access, it may be beneficial to configure your Nessus scanners to perform sweeps of all targets just on that or a handful of ports.

Although not covered in this document, if you have access to the Log Correlation Engine (LCE) and any type of NetFlow data, real-time Passive Vulnerability Scanner events, firewall connection events or even IDS events, this data can easily be saved as a list of IP addresses that can be targeted for an aggressive active scan.

FIREWALL SCANNING EXAMPLE SCENARIOS


Auditing One Firewall With Active Scanning

Our first scenario considers a very simple network that has one firewall. On one side is the Internet and on the other is a network of local Internet addressable nodes. There is no DMZ in this case. Following is an example topology:




The firewall implements a filtering policy. The policy may globally block certain types of ports or it may have different port filtering policies per destination address.

Let's add a Nessus scanner and a SecurityCenter to our network. The SecurityCenter and Nessus scanner will be deployed inside the network.



On the SecurityCenter we will create a repository named "InternalScans". This repository could contain vulnerability scan results, patch audits, web application scans – it doesn't really matter. All of our internal scans can go in here.

However, now we want to see what the "bad guys" from the Internet see when they scan our network. Add an external Nessus scanner to our topology as shown in the following diagram:



If we don't create a separate repository or set up a scan zone, we will have a big problem. Assume that host 1.1.1.1 is on the inside of our network but blocked by the firewall from anyone on the Internet. An internal scan with Nessus will identify 1.1.1.1 each time, but the external scans will consider 1.1.1.1 a dead host. To avoid this, we create a second repository named "ExternalScans".

With this new repository, we can leverage all of SecurityCenter's analytical tools such as port summaries, vulnerability summaries, counting hosts, etc. to create meaningful dashboards. There are many different types of reports you can do. For example, at Tenable we used our corporate SecurityCenter to count the number of Internet facing IP addresses and if the result was not an expected value, an email was generated.

Now that we have two repositories and two scanners in their own scan zones, we need to make a few decisions:

- How often do we want to launch scans from the outside to the inside?
- How many ports do we want to scan for?
- What do we care about?


How you answer these questions will be determined by your threat model, your firewall auditing requirements and maybe even external information such as port lists from NetFlow tools and sniffers.

SecurityCenter provides the ability to alert, trend or report on any type of data. You may want to create scheduled alerts that count the number of open ports and alert if it isn't an expected value. You may simply want to trend or report on the number of open ports over time. How you use the data is up to you. The point is that you can answer a wide variety of questions and create automated alerts and results when you need to.

Auditing One Firewall with Passive Scanning

Let's remove the Nessus scanners and add in a single PVS deployed just behind or just in front of the firewall so that it is monitoring all boundary traffic. The PVS needs to see a complete TCP session before logging a vulnerability, so it really doesn't make a difference if it is sniffing on the Internet side or internal side of the firewall.

Following is a screen capture of a PVS sniffing traffic. All vulnerabilities found by the PVS would be considered to be external. If there were two servers inside of the internal network, the PVS would neither see them nor report on them with this configuration.




Configuring SecurityCenter to send PVS data to its own repository creates a convenient opportunity to aggregate passive data for alerting or trending. In this example, we named the repository "BoundaryPassiveData" to avoid mixing this data with internal or external scan repositories.

With this repository in mind, we can create any type of alert, trend or report we want to. All vulnerabilities or open ports found in this manner are indeed Internet facing. We also have the advantage that we didn't have to scan through the firewall or conduct a scan on a limited group of ports.

But what if we had core switches and wanted to deploy the PVS internally in order to monitor all of our internal traffic? This type of deployment is still useful for auditing our

firewall. For this example, our internal PVS will have its data sent to a repository named "InternalPassiveData".




If the services and vulnerabilities found were stimulated by connections that originated from outside of the monitored network, the PVS will add the tag line of "External Access" to the report as shown in the introduction of this paper.

This means that for the internal PVS deployed, if you want to see what ports are open on through your firewall, you would create a filter that selected the "InternalPassiveData" repository and a vulnerability text string of "External Access". With this filter, any tools such as port summaries or listing of IP addresses would be for those services that were allowed through the firewall.

It is important that your internal PVS's configured network addressing is in line with reality. Tenable has worked with customers who leveraged PVS in this manner but the actual network addresses behind the firewall were different than those configured for the internal PVS.

Auditing a DMZ with Active Scanning


Let's add a DMZ (based on the term "de-militarized zone") to our network. DMZs sometimes behave as a layer between the internal and external connection and sometimes they are hung like a leaf off of the connections. It really doesn't matter from our point of view what the topology is because there are separate boundaries to be added at each layer. Here is an example leaf topology:



We still want to audit all of the connections:

- Internet to DMZ
- Internal to DMZ
- Internal to Internet
- DMZ to Internal
- Internet to Internal (egress filtering)
- DMZ to Internet (egress filtering)

Ignoring the egress filtering cases, consider adding Nessus scanners to our topology as shown below:



We could also envision creating repositories in SecurityCenter named and described as follows:

- Local DMZ
- Local Internal
- Boundary DMZ to Internal
- Boundary Internal to DMZ
- Boundary Internet to DMZ
- Boundary Internet to Internal

Configuring scan policies for targeted ports, scan schedules and deciding what to do with the results is again a function of how often and how thorough you want to audit the boundaries. Even though we have six different sets of data, we don't need to have six different cases for alerts, trending or reporting. SecurityCenter can leverage a filter with ad hoc repositories. For example, if you want to receive an alert for any Internet facing vulnerability with a severity of "high", you could create a query that selected both the "Boundary Internet to DMZ" and "Boundary Internet to Internal" repositories. If either had high severity vulnerabilities found, an alert would be generated.

Another important concept to consider is the exponential number of unique boundaries we may need to test. With one boundary we had two repositories, two scanners and perhaps

two scan policies. By simply adding the DMZ as one more enclave, we've gone to three scanners, six repositories and likely six different scan policies or schedules.

When performing Nessus scans, it may be difficult to test outbound egress filtering. Egress filtering is a good security practice because it can limit what a bad guy can do with a compromised system. They may not be able to communicate with a botnet or send spam email. However, Nessus needs a target to scan.

You may want to set up certain Internet based targets and not tell your firewall team about them to simulate this. In our case, having the Nessus scanners that are normally leveraged for local vulnerability scanning perform a full port scan on your sacrificial Internet host can expose flaws in outbound egress filtering.


If the outbound egress filtering policy is IP specific, such as having different outbound rules for the email server and another set of rules for the DNS server, the outbound scans from Nessus may not be relevant.

Auditing a DMZ with Passive Scanning

Passive scanning presents some unique benefits and limitations for monitoring multiple firewalled enclaves.


The largest benefit is tracking egress filtering. If the PVS is deployed to monitor boundary traffic, the "Client Side Port Usage" plugin (#00002) can be used to enumerate any port used in an outbound connection. This data is an excellent method to verify which ports are allowed to leave a firewalled enclave. The data can be used for alerting, trending and reporting. For example, selecting plugin #00002 and a repository of PVS data and the port summary tool would give you a list and port count of how many systems browsed on which ports.

A limitation of the PVS monitoring multiple enclaves is what it considers to be "outside" of the network. For example, consider the following PVSs deployed within our DMZ and in our internal network:




If we configured PVS #1 with the network address of the DMZ, then both connections from the Internet as well as from internal IPs would be considered "outside".

To avoid complexity, consider configuring both PVS #1 and PVS #2 with the entire network range. This would allow you to have each PVS focus on Internet facing security issues and you would be able to send data from both PVSs to a single repository, as shown below.



Even though there are two PVS sensors, we can pretend that there is really one network range or one giant enclave that consisted of the aggregation of the DMZ and the internal ranges. If each PVS was configured with those ranges, it would only report “externally” facing passively discovered vulnerabilities on either enclave that were accessible from the Internet.

Finally, keep in mind that your span ports or sniffing infrastructure may be able to differentiate packets from different boundaries. If so, then multiple PVSs may be able to monitor separate connections and have a better concept of inside vs. outside. Consider the following topology that shows multiple PVSs deployed on various Internet facing or internally facing span ports:



Depending on your sniffing topology and technology, you may be able to meet your needs to provide real-time and continuous boundary monitoring of your firewalls.

Such a topology may seem overly complex, but it serves to illustrate how configuring where PVS sniffs packets from as well as what it considers “inside” versus “outside” can be used for filtering.

Auditing Star-Topology Enclaves with Active Scanning

Consider the following large topology that has seven different enclaves and seven different scanners with seven different firewalls. These enclaves all have unique IP addresses and are not using Network Address Translation (NAT).



We can easily imagine using SecurityCenter to create a scan zone and a repository for each enclave to facilitate “internal” scans. However, for our boundary scans we could envision forty-two scans! For each enclave, there are six other potential enclaves to scan. Doing this seven times gives us forty-two potential scans. SecurityCenter supports generating this many scans policies, scan schedules and even repositories, but this is overly complex.

A lot of questions can be answered if we simply configured one scan of all the other enclaves for each enclave. For example, for enclave #1, we’d have a scan and repository for the internal scanning of that enclave, plus a second scan and repository for an external scan of enclaves #2 through #7. We may have a repository naming convention along these lines:

- Internal Enclave #1
- Boundary Enclave #1
- Internal Enclave #2
- Boundary Enclave #2
- Internal Enclave #3
- Boundary Enclave #3
- Internal Enclave #4
- Boundary Enclave #4
- Internal Enclave #5
- Boundary Enclave #5
- Internal Enclave #6
- Boundary Enclave #6
- Internal Enclave #7
- Boundary Enclave #7


That is 14 repositories, which is much more manageable than 42. From an auditing perspective, we can leverage a variety of filters to create queries for us.

To look for changes in expected values, alerts can be configured that count the number of hosts, number of open ports or number of vulnerabilities. These alerts can be applied to all of the boundary repositories at once for a single aggregate alert, or perhaps it would be more appropriate to have more discrete alerts that were tied directly to an enclave.

To audit what a specific enclave was showing to the rest of the organization, a filter that selected the other entire boundary enclaves along with an IP address filter for the IPs of the enclave in question would select all vulnerabilities found that were external. For example, if we wanted to see all of the externally facing vulnerabilities or open ports for enclave #5, we'd create a query that selected enclaves #1-#4 and #6 and #7 as well as the IP address ranges for enclave #5.

Auditing Star-Topology Enclaves with Passive Scanning

Consider the following topology that has a PVS deployed internally to each enclave.



If each PVS was deployed to watch internal traffic of the enclave, then the data reported would need to be filtered to discriminate "enclave facing" services versus "internally facing" services. Using the "External Access" vulnerability text filter is an easy way to perform this sort of analysis.

To create meaningful alerts or reports about firewall changes and open ports, each PVS could actually send its data to a single repository. Since each enclave has discrete IP addresses, there will not be any overlapping. If we wanted to see "enclave facing" vulnerabilities or services that an enclave's firewall was letting through, we'd set the IP filtering of the query to that enclave's ranges (or likely use a SecurityCenter asset list configured with this information before hand) and filter on vulnerability text that had the "External Access" tag. Because the PVS sensors were deployed internally, the data could only be used to look for services being offered through the firewall and not have the ability to audit egress filtering.


Alternatively, if the PVS sensors were actually deployed on the perimeter of each enclave then the data from them would be entirely based on the boundary traffic of the enclave. In this case, any repository or repositories that had this type of data would be focused entirely on traffic that was traversing the firewalled boundary.

CONCLUSION

Deploying combinations of active and passive vulnerability scanners at key locations in your network can help identify changes to your security boundaries. Nessus, the Passive Vulnerability Scanner and SecurityCenter can be leveraged to perform continuous and scheduled audits of your boundaries and have this data available for alerting, trending and reporting.

Tenable also offers the ability to audit firewall, intrusion prevention, NetFlow, network traffic and web proxy logs with the Log Correlation Engine. This tool enables Tenable customers to summarize traffic activity based on port and to also detect changes that have been made to security devices. These events can be used to further monitor your boundaries. The PVS also can perform real-time logging of network changes it finds such as new hosts, new open ports and new browsed ports. SecurityCenter users use these events from the PVS or from the firewalls to have alerts generated and Nessus scans scheduled.

Another benefit of deploying multiple Nessus scanners is that the topology generated by the Tenable 3D Tool is much more accurate. By leveraging `traceroute` information from multiple network vantage points, the 3D Tool can produce an accurate picture of the network such as the one shown below.



If you have an approach or feedback on performing boundary monitoring with Nessus, the Passive Vulnerability Scanner and SecurityCenter, please share it with Tenable and other customers on the Tenable Discussion Forums (<http://discussions.nessus.org>).

ABOUT TENABLE NETWORK SECURITY

Tenable Network Security, the leader in Unified Security Monitoring, is the source of the Nessus vulnerability scanner and the creator of enterprise-class, agentless solutions for the continuous monitoring of vulnerabilities, configuration weaknesses, data leakage, log management and compromise detection to help ensure network security and FDCC, FISMA, SANS CAG and PCI compliance. Tenable's award-winning products are utilized by many Global 2000 organizations and Government agencies to proactively minimize network risk. For more information, please visit <http://www.tenable.com>.

Tenable Network Security, Inc.
7063 Columbia Gateway Drive
Suite 100
Columbia, MD 21046
410.872.0555
www.tenable.com