Unified Security Monitoring Best Practices

June 8, 2011

(Revision 1)

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

Table of Contents

INTRODUCTION

Tenable's Unified Security Monitoring (USM) solution enables organizations to unify their system and event monitoring for centralized security and compliance reporting. This document outlines several best practices when deploying and optimizing a USM platform to perform security and compliance monitoring for enterprise networks. It also shows how various capabilities can be leveraged for greater maturity in how USM is used to monitor both systems and events for security and compliance.

USM COMPONENTS

Tenable's Unified Security Monitoring solution is comprised of several key components: Nessus, Passive Vulnerability Scanner (PVS), Log Correlation Engine (LCE) and SecurityCenter.

Nessus

Tenable Network Security's Nessus® vulnerability scanner is the world's leading active vulnerability scanner, featuring high-speed discovery, patch auditing, configuration testing, botnet monitoring, asset profiling and web application testing of your organization's systems. When managed by SecurityCenter, Nessus scanners can be placed behind firewalls, within enclaves, within discrete networks, inside a DMZ, outside of a DMZ and any other locations.

Nessus scanners are used by Tenable customers to audit thousands of hosts in a single scan and many of our customers scan more than 100,000 nodes on a routine basis.

Passive Vulnerability Scanner

Tenable Network Security's Passive Vulnerability Scanner (PVS) monitors network traffic in real-time to report on vulnerabilities and provide a forensic log of network activity. Each PVS sniffs vulnerabilities for both client and server side issues. It also logs a wide variety of network events such as each DNS lookup, email attachments, encrypted sessions, remote administration events, SQL database queries, Google searches and Windows network sharing file transfers.

A typical PVS deployment can monitor 500 Mb/s to 1 Gb/s of network traffic generated by 10,000 nodes.

Log Correlation Engine

Tenable's Log Correlation Engine (LCE) centralizes logs from a wide variety of platforms including firewalls, system logs, authentication logs, web applications and much more. Each LCE can be used to compress and normalize logs to enable reporting, visually intuitive dashboards, alerting and forensic investigations in support of compliance and security monitoring. The LCE provides a powerful array of tools to analyze and simplify events. The automation and analysis provided by LCE allows users of many different skill levels to identity compromised systems, help perform forensic investigations, identify and investigate anomalies, detect and track changes and identify malicious insiders.

A typical LCE can process in excess of 1,000 logs or events per second while maintaining a 20 to 1 compression ratio for any log received. Queries to the LCE for millions or billions of events are completed in just a few seconds, enabling users to analyze any type of log in context very quickly.

SecurityCenter

The SecurityCenter (SC) enables organizations to easily measure vulnerabilities and monitor security events, asset by asset. Multiple users and organizations can participate in role-based vulnerability scanning, patch auditing, passive network monitoring, log analysis, anomaly detection and more.

The SecurityCenter can also manage each attached Nessus vulnerability scanner, Passive Vulnerability Scanner or Log Correlation Engine. This enables deployments of multiple scanners for load-balanced scanning, multiple PVS sensors for distributed network monitoring and multiple LCEs for distributed log queries and centralized event analysis.

Typical USM deployments that leverage SecurityCenter and all of our technology include dozens of Nessus scanners, dozens of PVS sensors and one to three LCEs.

USM DEPLOYMENT AND CAPABILITIES MODEL

The following image outlines Tenable's Unified Security Monitoring Deployment and Capabilities model. It shows a roadmap that Tenable customers can use to ensure they are leveraging Tenable products appropriately.

The green boxes identify different levels of performance and monitoring that can be achieved by leveraging additional sensors or additional capabilities of SecurityCenter. These additions increase the capabilities of a USM deployment, but also increase the complexity.

For example, many organizations start out with one Nessus scanner managed by SecurityCenter. Deploying multiple scanners can load balance scanning for a large internal network and adding external scanners enables perimeter scanning. Two scanners are more complex than one and additional scanners deployed externally are more complex than deploying them internally.

The blue boxes signify a capability of a USM technology that can be leveraged, but requires a change in behavior in how your organization consumes the data. For example, patch auditing can be performed with Nessus. However if patch auditing is to be leveraged, some

form of host authentication is required which often requires interaction with external organizations. Further, since patch auditing is of great interest to IT system administrators, Tenable customers often transform their network monitoring from pure vulnerability scanning to scanning for rogue systems that aren't being managed by the IT department.

This paper describes these benefits and trade-offs for all aspects of the USM platform. For each additional maturity level, the general benefit and an example "quick win" is discussed.

Leveraging Security Center Users, Repositories and Organizations

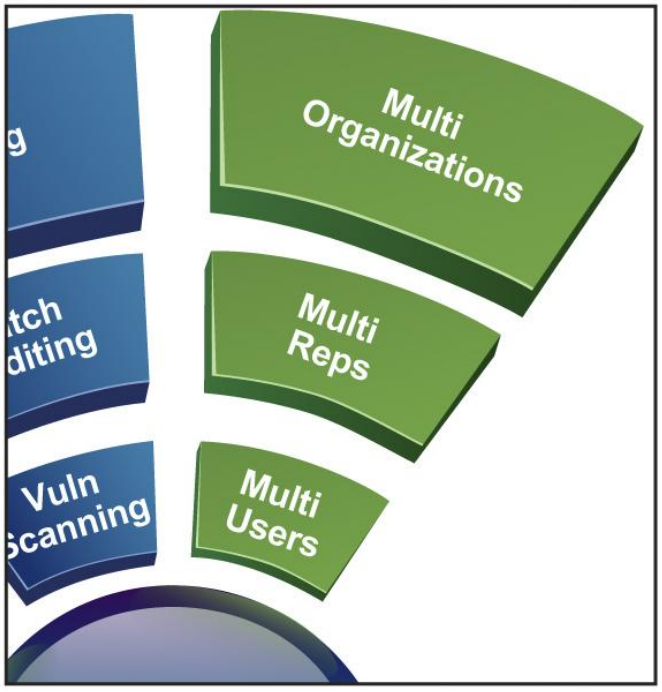

There are three levels of SecurityCenter access control that enable both small and large organizations to develop information security sharing programs.

Managing User Access

Benefit *– Share system and event security data with authorized users.*

Quick Win *– With simple asset lists, security audit work performed in bulk by one user can be easily parsed and shared securely with many different users.*

Every SecurityCenter deployment starts with adding at least one user. This user is part of an organization and has access to one repository of vulnerability data and/or log data.

Organizations often grow their user base by adding more user accounts that have the same access to all of the existing data with no restrictions. For example, we've seen many deployments where the CIO, CISO, Director of Incident Response and several key IT or compliance managers have access to all collected data.

However, when an organization wants to add more user accounts that are restricted to a limited set of systems, they leverage access control based on SecurityCenter asset lists. An asset list is set of IP addresses or ranges that define what a user can or cannot see. For example, a SecurityCenter may be used to manage scanning and log analysis for a 10.10.0.0/16 Class B address, yet there is a Payment Card Industry (PCI) requirement to perform scans and review logs of database servers in the 10.10.200.0/24 Class C address. In this case, an asset list could be generated for the Class C range and then users added that are restricted to only seeing data from there.

Managing Vulnerability Data with Repositories

Benefit *– Storing system audit data in discrete volumes creates many opportunities for sharing data, enhanced reporting and secure distribution.*

Quick Win *– Creating repositories based on audit type allows for transparent tracking of the overall audit process.*

Although SecurityCenter's ability to restrict access to vulnerability data based on an asset list is powerful, it cannot be used to further limit access to different types of data on a per-host basis.

For example, in our previous example of hosts in the 10.10.200.0/24 Class C network that are subject to PCI requirements, how could we perform a test on them for a different compliance standard such as a different government standard, a deeper network scan or even performing scans with Nessus that were not compliant with the guidelines outlined by PCI?

The answer is to leverage SecurityCenter repositories. A repository is a SecurityCenter component configured by an administrator. Each administrator defines one or more network ranges and one or more organizations to associate with the repository. Configured PVS sensors are associated with repositories by administrators and scan results are automatically imported into repositories by users who have scanning privileges.

Within an organization, users can be given rights to view any repositories associated with the organization. Users with access to multiple repositories can view all data together or partial data, as repositories can have overlapping ranges.

There are many use cases for leveraging SecurityCenter repositories. When Tenable designed this feature into SecurityCenter, we envisioned several practical uses that would allow complex organizations to collect a wide variety of security data, and share it securely with different consumers without "over reporting". Some use cases include:

- ***Perception management*** – Two repositories are managed for the same range. One is used to perform ongoing testing such as continuous PCI compliance scanning. The second is only used to perform scans or audits when we know that the asset is compliant. Auditors are only given access to the secondary repository.
- ***Active, Patch and Passive Data Separation*** – Although the SecurityCenter supports filtering vulnerabilities based on type of data delivered by Nessus for scanning, patch auditing and sniffing from the PVS, it can be very convenient to have these results available in a few selectable repositories. This simplifies dashboard, alert and report creation.
- ***Boundary auditing*** – Multiple repositories can be used to separate audit data based on its perspective to a target, such as an external scan, a scan through an Intranet VPN or

to test how a screening firewall or router is configured. This concept is a Tenable's USM Best Practice and is extensively covered in our paper "Firewall and Boundary Auditing" available at http://www.nessus.org/expert-resources/whitepapers.

- ***Scan type auditing and tracking*** – Nessus can be used to report on many different types of scans. Often, an organization will want to know the last time a host was scanned. Nessus plugin 19506 tracks this very well, but it does not keep a history of previous scans. Placing your scan results into different repositories allows easy reporting of plugin 19506 results. For example, consider two repositories, one for complete PCI scans and another for weekly light scans. Plugin 19506 from the PCI repository will accurately report the last time a host had a PCI audit performed against it.
- ***Separation of specific audit results*** – Certain types of audit data may need to be analyzed by an expert or may be extremely sensitive in nature and restricted to a few individuals. For example, leveraging Nessus for web application auditing typically requires an analyst to manually analyze the results. Similarly, performing audits across the network to look for sensitive information such as credit cards or social security numbers are also candidates for restricted access.
- ***Sharing results across organizations and SecurityCenters*** – Repositories can be shared between organizations on a SecurityCenter as well as shared between multiple SecurityCenters.
- ***Overlapping RFC 1918 networks*** – If your network is large enough that it encompasses overlapping RFC1918 IP address space, multiple repositories can be used to separate scans results as well as PVS passive scanning results.

Leveraging Multiple Organizations per SecurityCenter

Benefit *– Very large security organizations can implement complex reporting structures.*

Quick Wins *– Separation of duty is maintained by ensuring security and compliance groups only have access to the security data they are authorized to view.*

Each SecurityCenter deployment requires the administrator to configure an organization that all users initially belong to. Multiple organizations can be added later. Each organization can independently access their own scan zones for distributed Nessus scanners, Passive Vulnerability Scanner reports, repositories and one or more Log Correlation Engines.

Different organizations can be used to provide as much or as little transparency to the internal workings of how business units go about their security and compliance monitoring and reporting functions.

Many Tenable customers are content to build out their organizational chart within a single organization and leverage asset lists, repositories and access to credentials to provide the proper role-based access and oversight. However, at some point due to the sheer size of the organization, certain functions such as controlling scans, access to scan credentials and access to logs, may become cumbersome. Additionally, since SecurityCenter may have different consumers of data for compliance and security monitoring, there may be different levels of tolerance for over-sharing of data. For example, a web security team may want to analyze any type of potential vulnerability reported on a web server, but sharing this data with a compliance team could cause them to react to un-validated security issues.

Finally, the ultimate form of organizational maturity is to run them on different SecurityCenters. Some very large customers run multiple SecurityCenters and leverage independent organizations for various audit and reporting functions. In these cases, one SecurityCenter has credentials to log into another SecurityCenter and pull back the

repository as a snapshot or on a scheduled basis. This feature is very popular with our larger customers since the remote repository IP addresses do not count against the SecurityCenter license key.

MIGRATING TO CONTINUOUS CONFIGURATION AUDITING

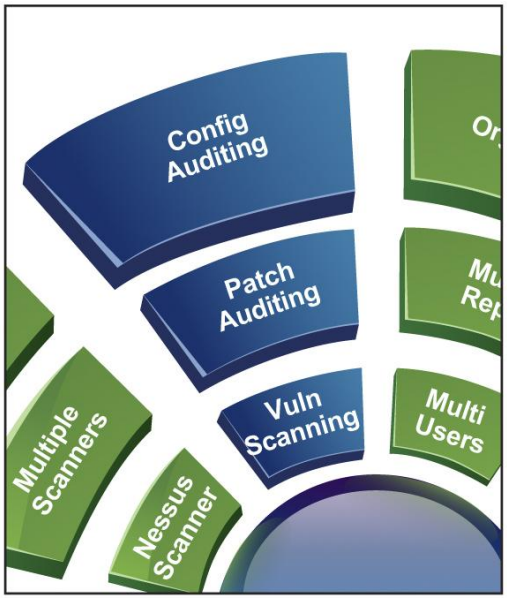

Nessus supports three types of audits: un-credentialed vulnerability scanning, credentialed patch auditing and credentialed configuration auditing. In this section, we will discuss how Tenable customers have transitioned from performing vulnerability scans to understanding how their configurations change.

Un-credentialed Vulnerability Scanning

Benefit *– Rapid identification of devices on the network and vulnerabilities in their services.*

Quick Win *– Immediate inventory of devices, applications and critical vulnerabilities.*

Nessus identifies vulnerabilities without logging into a target system by creating packets and network sessions that discover hosts and their services. For the discovered services, Nessus performs one or more non-destructive probes to enumerate information about the host and any potential vulnerability that may exist.

Nessus has over 40,000 unique plugins to identify a wide variety of systems and vulnerabilities. Many of these plugins focus on unique identification of operating systems and applications. Tenable has learned that no two networks are alike and often, assumptions about how certain techniques such as TCP/IP operating system fingerprinting work don't scale reliably in a complex enterprise.

However, since Nessus does such a good job identifying systems and server-side vulnerabilities, some users incorrectly assume that Nessus has nothing more to offer. As an independent auditing technology, Nessus can be leveraged for both patch auditing and configuration auditing, which are often key requirements of security and compliance monitoring programs.

Patch Auditing

Benefit *– Instantly see which systems are missing patches in a language your IT staff speaks and in less time than a typical vulnerability scan.*

Quick Win *– Identify all software as well as vulnerabilities in third-party applications without the use of an agent.*

When provided with credentials to log into the target Unix, Windows, router and other types of network devices, Nessus can enumerate all missing security patches. This type of audit is typically faster than a network scan since Nessus is making the same sort of API calls that an agent or operating system component would leverage. For example, to perform a full TCP port scan, any port scanner would have to send more than 100,000 packets per host. However, with credentials, Nessus sends less than 1,000 packets to enumerate all TCP and UDP ports, as well as the listening application. This type of data is much more useful than a simple list of ports.

Besides speed, patch audits have several major advantages over un-credentialed vulnerability scans:

- Reporting of specific patches in the nomenclature of the technology (i.e., specific Solaris patch IDs, Microsoft KBs, Red Hat RPMs, etc.).
- Identification of client-side vulnerabilities such as those used for email or web browsing.
- Identification of kernel level issues such as missing service packs.
- Identification of hand-compiled Unix and Windows services that may indicate unmanaged applications or malicious software.
- Identification of antivirus software that has out-of-date virus definitions.
- Identification of all software, BIOS information, running services, network interfaces, user data and much more.

To perform patch audits with SecurityCenter, users or organizations can manage credentials independently from scan policies. This allows for separation of duties so that one user could be responsible to ensure that scans are occurring with the right passwords, and another that the appropriate scans are being performed. Within SecurityCenter, a user can be given access to a credential without knowing what it is actually for. Credentials can make use of many different technologies to help your security staff comply with any type of password or credential policies your organization may have in place.

To leverage patch auditing, Tenable customers often create a variety of SecurityCenter organizations, users, scan policies, assets, credentials, scans and repositories to cover any combination of scan targets and reporting requirements.

An organization that is migrating from un-credentialed vulnerability scanning to patch auditing typically starts with patch auditing on a controlled number of systems. A server or server farm are typical targets to start performing patch audits. If a Windows domain is used to manage desktop computers, a patch audit that leverages those credentials is also possible.

When making the switch or adding additional types of patch auditing scans to SecurityCenter, Tenable recommends the following approaches:

- If a group within your organization has credentials but does not want to place them into SecurityCenter, they can perform a scan with Nessus and then upload the results manually.
- Consider creating a user role that only manages credentials leveraged by certain scans.
- Don't perform un-credentialed vulnerability scans and patch audits at the same time. Most Tenable customers operate with distinct scanning schedules for discovery of devices as compared to patch auditing.
- Create separate repositories to keep patch audit results distinct from un-credentialed vulnerability scans.
- Create dashboard elements, alerts and reports for each type of audit being performed. For example, you may have a "daily discovery scan" that attempts to ping hosts and enumerate some services, a "monthly full scan" that attempts to perform a more aggressive port scan and "weekly patch audits".
- When performing patch auditing in Windows environments, Nessus can be used to enable and disable the remote registry service, which is vital for enumeration of patches and software.
- Unix audits can leverage SSH keys as well as su and/or sudo.
- Nessus will report when an attempt to log in with credentials has failed. You can use these reports to identify when a system is present that isn't being managed or that the credentials have changed.

Many Tenable customers have added transparency to their patch management operations by monitoring and trending the actual age of patches for each organization. This can be done with any type of system data within SecurityCenter, but when reported on with patch audit data, it provides a very simple and easy to understand metric that can be compared between organizations.

Continuous Configuration Auditing

Benefit *– If you are already performing a patch audit of a system with Nessus, you can perform a configuration audit with a little extra effort.*

Quick Win *– Ensure that systems have logging enabled and strong password policies.*

Nessus also has the ability to audit the settings and configuration of Windows, Unix, router, database and other types of technologies. An example setting could be how often a system user is required to change their password. Tenable produces audit policies aligned with various industry and governing standards that can be downloaded and added to a scan policy.

When comparing patch auditing data to un-credentialed vulnerability scanning data, we saw that patch auditing added a lot of value missed with network scanning. Similarly, configuration auditing can identify many items that would greatly impact the security of any organization. For example, imagine a fully patched domain controller with no exploitable vulnerabilities yet:

- Logging is not enabled leaving no trace of user authentications or system errors
- Passwords may never expire and there are no complexity requirements
- It uses a DNS server that will be decommissioned in the next 90 days

The results of a configuration audit scan for each system include a list of settings that are compliant as well as a list of settings that are not. This is an important distinction from a Nessus vulnerability audit. To demonstrate compliance with a certain configuration standard, the actual list of settings is needed for transparency and auditing.

There are many different types of best practices used by Tenable customers to perform configuration audits.

Some Tenable customers have deployed Nessus and SecurityCenter solely to perform one type of audit. For example, many of our federal customers in the United States have deployed SecurityCenter and multiple Nessus scanners to perform credentialed configuration audits of thousands of Windows hosts for compliance with FDCC requirements. Other customers have taken policies written by Tenable that have been certified by the Center for Internet Security (CIS) to audit key systems such as their Microsoft Exchange email servers, Windows domain controllers or key Oracle database servers.

In other cases, customers may deploy different types of audits in an effort to identify new settings and recommendations that can be of use in their organization. Customers who want to harden their systems beyond one standard may run multiple policies for a variety of government, vendor and consensus type standards.

Tenable recommends the following best practices for performing configuration audits:

- Practice the same type of credential management as for patch auditing. For example, create SecurityCenter users who's only job is to manage passwords for credentialed scanning.
- Place configuration audit results into separate repositories.
- Do not perform un-credentialed vulnerability scans while you are performing credentialed configuration audits.
- Allow a sufficient time period to perform the audit. Depending on the policy, configuration audits can take longer than a typical patch audit. For example, some policies require searching of the *entire* hard drive to audit file permission issues.
- It is helpful to track how long scans take to make an informed decision on performing simultaneous patch and configuration audits or performing them separately.
- Some Tenable customers who have the Log Correlation Engine track systems that had a detected change on them, such as installed software, and then launch a configuration scan of that system.
- SecurityCenter supports filtering of results by audit policy and configuration scans can be stored in a specific repository. This provides many opportunities for reporting and creating views for people who need to see all data or auditors who need to see just one type of data.
- It is common for Tenable customers to leverage and customize the audit policies available to them. It is recommended that customized audit policies be renamed to indicate that they are modified and avoid confusion if an update is available, or if an auditor who is familiar with Tenable audits reviews your results.

Each of these types of audits will help build a reporting model that organizations can use to migrate from enumerating security and compliance data to that of exception management.

DISTRIBUTED NESSUS SCANNERS

Organizations usually start out leveraging multiple Nessus scanners to make their scans go faster or to scan "through" a firewall. In this section we will discuss how to leverage multiple scanners to perform true boundary monitoring of external and internal networks.

Tenable has published an in-depth paper titled "Firewall and Boundary Auditing" that describes several scenarios of distributed active scanners and passive sensors purposely deployed to monitor the perimeter firewall policy. This paper is available at http://www.nessus.org/expert-resources/whitepapers.

Single Nessus Scanner

Benefit *– Scans launched from one point centralizes ad hoc scans performed by different groups.*

Quick Win *– A single Nessus scanner can run on the same system as SecurityCenter to simplify initial deployments.*

Associating a single Nessus scanner to a SecurityCenter is the most basic form of vulnerability scanning supported. From this vantage point, users can launch any type of scan: vulnerability scan, patch audit or configuration audit. However, the scanner must have direct access to the scanned targets. If there is a firewall, network address translation boundary or proxy, Nessus may not be able to access the target directly.

Distributed Scanners

Benefit *– Scanners can be placed behind firewalls to scan the insides of various enclaves.*

Quick Win *– Decreased scan times due to SecurityCenter load balanced scans.*

Multiple Nessus scanners can be grouped together and deployed in logical components known as "scan zones". Each scan zone is treated by SecurityCenter like a single scanner. Scans can be launched from a scan zone and the work of performing the scan is distributed between each available scanner.

Deploying multiple scanners has many benefits including:

- Multiple scanners are more reliable than a single scanner. If a scanner fails within a scan zone, scans can still proceed.
- Scans have less impact on your infrastructure than when attempting to scan through boundaries. For example, your backbone routers don't have to carry port scans.
- Scans are faster since the scan work is divided up among the available scanners.
- With distributed scanners, scans can be launched from vantage points behind firewalls and on external networks.

Access to various scan zones can be limited to organizations or users just like many other SecurityCenter resources. This means you can allow power users to create scans that run from inside a target network or perform external scanning to see what sort of services are exposed. This also means you can force users to always use certain scanners minimizing the work they need to do when setting up a scan.

Each SecurityCenter allows more than 500 Nessus scanners to be managed at one time. This enables Tenable customers to think about their network in terms of the audit data they want to obtain and not try to maximize what any one scanner can or cannot scan. Common questions that shed light on best practices for distributed scanning include:

- Are enough scanners deployed to perform an adequate discovery scan as often and quickly as needed?
- Are there enclaves (behind firewalls or other boundaries) for which the organization has no scanning data?
- Are there secure external scanners positioned to provide Internet or Extranet views that hostile intruders may seek to exploit?
- Are there any NetFlow, sniffing, network based anomaly detection, network intrusion detection systems or other types of network security or performance monitoring solutions that are negatively affected by scanners?

In the next section, we will briefly discuss performing active scans to measure open ports and trust across network boundaries. Many organizations will jump to this type of monitoring for obvious boundaries, such as Internet facing services. However, Tenable recommends as much internal scanning and auditing be accomplished prior to deploying an infrastructure for boundary monitoring.

Boundary Monitoring

Benefit *– The actual number of devices and ports on any boundary can be audited.*

Quick Win *– Unauthorized backdoors and services can be found.*

Boundaries are particularly important to monitor since they are established to limit access. Firewalls, screening routers, intrusion prevention systems, spam gateways, network access control devices, VPNs, wireless access points and many other types of network components include filtering logic that prevents access to users based on their authentication, destinations and network locations.

Modern networks are often very complex in nature and have many different components that may be managed by different organizations. For example, Tenable has several customers where the perimeter firewalls are managed by a central "network group" and screening firewalls for various PCI services are managed by a separate e-commerce security team.

Beyond any one boundary, large organizations have a difficult access control problem to monitor. With an enterprise, trying to determine which groups have access to which servers can be very difficult. Most organizations focus on authentication and try to ensure that only certain groups and users have access to certain types of data. Unfortunately, this often means that the only thing preventing an internal user from gaining access to a restricted resource is a password.

The Tenable "Boundary and Firewall Monitoring" paper describes a variety of best practices to perform active scans with distributed Nessus scanners that can be used to:

- Enumerate boundary facing devices such as public web servers.
- Enumerate boundary facing services such as email and web servers.
- Alert and report against expected counts of open ports and system numbers.
- Leverage repositories and distributed scanners and compare the list of open ports and devices from various vantage points within the network.

The major point of the paper is that this sort of monitoring is measurable and can be automated. Trend lines of open ports, weekly counts of Internet facing systems, alerts for when new ports have been discovered are all examples of boundary monitoring that can be accomplished with SecurityCenter and distributed Nessus scanners.

Passive Monitoring Best Practices

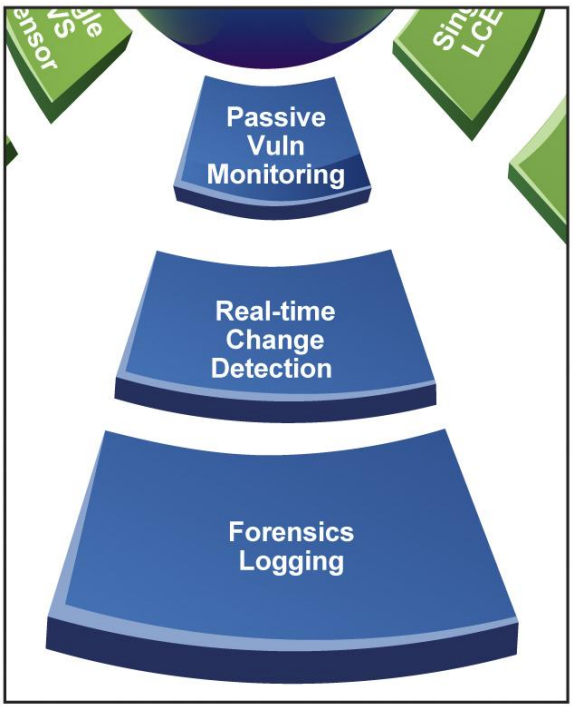

Each Passive Vulnerability Scanner can be leveraged for continuous vulnerability monitoring and real-time change detection logging of network data for forensic analysis. Once installed, they simply log data – no further configuration is required.

Passive Vulnerability Monitoring

Benefit *– Continuously identifies all systems and vulnerabilities associated with observed network traffic.*

Quick Win *– Identifies client-side vulnerabilities without the need for a credentialed network scan.*

The major feature of the Passive Vulnerability Scanner is the ability to continuously identify vulnerabilities through analysis of network traffic. As with Nessus, the data produced by PVS can be used to quickly enumerate and track all systems, servers, applications, network addresses, ports and many other forms of asset classification.

Passive traffic analysis identifies client-side applications and vulnerabilities that may not be readily obtainable with un-credentialed active scanning. Many Tenable customers who do not have the ability to leverage credentials to perform patch and configuration audits with Nessus deploy a PVS to specifically enumerate client-side software. PVS readily identifies web browsers, email clients and P2P software. It also tracks very useful client side information such as the version of Java in use or if a Windows computer is pulling patch updates directly from Microsoft.

PVS also has no impact on the network it is monitoring. Since it is passive, it will identify every open port, browsed port and system that it has seen on the network. This means that fragile networks or low bandwidth links can be monitored as well as very busy or near capacity high-speed links. In either case, the PVS does not add any additional traffic to perform its audit and the monitored systems have no impact on their resources.

Although many comparisons between active scans with Nessus and passive monitoring can be made, Tenable recommends a best practice approach that blends active and passive monitoring. Several best practices to consider include:

- Where credentials are not available, use the PVS to monitor client-side vulnerabilities.
- When scans are not allowed or are not allowed as often as desired, supplement PVS monitoring for additional coverage.

Tenable licenses the PVS uniquely from SecurityCenter and Nessus. Individual PVS deployments are available for a flat fee as are unlimited deployments of a PVS within a Class C or Class B network range. The Class B licenses are very popular with our larger customers who wish to deploy multiple PVS sensors throughout their infrastructure behind firewalls and inside various enclaves.

Real-time Change Detection

Benefit *–Identifies new applications, services and vulnerabilities without any operational impact.*

Quick Win *– Identify firewall rule changes that open vulnerabilities to your network in real-time.*

Since the PVS monitors traffic on a continuous basis, it can identify new systems, new ports and new vulnerabilities in real-time. The change identified by the PVS manifests in two ways.

First, every piece of information collected from the PVS by SecurityCenter is tagged with a timestamp. This allows SecurityCenter to track vulnerability and system data over time. By using a time filter, SecurityCenter users can select any type of data discovered by PVS. By selecting items that have been reported within the last few days, dashboards, alerts and reports can be produced that identify change.

Second, each PVS can be configured to send real-time vulnerability data to the Log Correlation Engine (LCE). The best example of this is the "New_Host" event. This is sent by PVS when an IP address is found to be active that was previously not active. Many other types of real-time logs are sent by PVS including discovery of new open ports, new trust relationships, new browser ports and new vulnerabilities.

Most Tenable customers send PVS data to a dedicated SecurityCenter repository. This allows any type of data to be treated uniquely, including working with passively discovered open ports. Using time filtering on repositories of passive data identifies changes in several items including:

- The number of open ports
- The number of active systems
- The number of client side vulnerabilities

Many Tenable customers take advantage of this data with the following best practices to identify change:

- Create trend lines or alerts based on repositories for the total number of open ports or systems per asset group.
- Leverage SecurityCenter dynamic asset lists to automatically create lists of IP addresses that are new or have had some sort of passively detected change. These lists can be scanned automatically or on a manually assisted basis.
- Leverage the Log Correlation Engine for real-time alerting. PVS logs sent to the LCE are normalized and are available for discrete real-time alerts. For example, when used in concert, SecurityCenter can schedule a policy alert to check if any "New_Host" alerts have been generated for a sensitive asset list such as a DMZ.
- Create reports or dashboards that compare the results between PVS and Nessus data to gain a unique insight into what is actually occurring on the network. Since the PVS and Nessus offer unique methods to enumerate hosts, the real-time changes found by the PVS are often compared to what Nessus scans have found.
- Leverage passively discovered boundary facing vulnerabilities, trust relationships, open ports and browsed port logs to identify changes in firewall or boundary policies.

Monitoring for change is a distinct best practice that is unique to passively monitoring for vulnerabilities. Organizations who wish to perform this type of monitoring need to identify what types of change need to be reported and who will investigate such changes.

Forensic Logging

Benefit *– All network file transfers, administration sessions and web searching are logged passively for analysis.*

Quick Win *– Analysis of malicious insider activity is easily understood when network logs are available.*

PVS utilizes specialized algorithms to provide deep analysis of vulnerabilities and identification of applications in complex protocols such as SQL, HTTP, DNS, FTP and Windows file sharing. A by-product of this analysis is the ability to produce a simple log message for many types of basic network events such as DNS lookups, emails with attachments and downloading of PDF files from a Windows network share. Web traffic searches and logins to social networking sites such as Facebook and Twitter are also logged. Tenable's Research team is constantly adding more and more protocols to PVS's arsenal of decoders.

There are several use cases for performing forensic logging of network activity that include:

- Use passive network analysis of employee activity. Logging DNS queries, web searches, social networking, network administration sessions and tracking files that have been downloaded enables employee profiling, network activity understanding and can help determine if an employee was malicious or not. The passive data also complements system logs, USB activity and other corroborating log sources.
- Perform malicious incident analysis from virus infections and successful penetrations. Antivirus defenses and intrusion prevention systems are not very helpful for incident analysis. An independent audit trail of all files downloaded, DNS queries made and web activity can greatly assist in identifying systems that have been compromised.
- Passive network traffic can substitute for a lack of network and system logs. For example, in the absence of a DNS server that logs IP address lookups, logs from the PVS can be used. Similarly, if database transaction logs are not available, all SQL queries, including insertions and deletions, can be logged. All VNC, RDP, SSH and SSL sessions are logged and can easily be used to drive correlation rules on the LCE or audit network activity. Finally, any type of unencrypted web query, either inbound or outbound, can be logged. This compensates for the lack of logging from internal web servers as well as outbound web queries logged from a web proxy.
- Any passively obtained log can be used for anomaly detection and correlation. When sent to the LCE, any PVS log can be used to look for first time events, continuous events and statistical anomalies. Considering the type of log, this is very useful to find malicious insiders, compromised systems and systems sending SPAM email.

Tenable views forensic logs from the PVS as a vital complement to any type of network or system monitoring effort. This functionality is available with any PVS deployment that is monitored with a Log Correlation Engine.

Distributed PVS Sensors

Organizations usually start out leveraging multiple PVS to monitor a perimeter choke point and progress to monitoring a variety of internal networks. Organizations that deploy passive sensors specifically to identify devices and ports allowed through a firewall have a near real-time method to monitor what is on their boundary.

Single PVS Sensor

Benefit *– A single PVS can identify a great deal of server and client vulnerabilities.*

Quick Win *– Continuous discovery of any type of network change and identification of new services, applications and vulnerabilities.*

Many Tenable customers enhance their security monitoring program with passive network monitoring. Customers who were solely performing un-credentialed or infrequent scans have provided feedback to Tenable that passive monitoring not only identified client-side security issues, it also identified ports and systems that were not targeted with active scanning.

A common shortcoming of any active scanning program is the tradeoff between how often to scan and how thorough the scans are. For example, consider simply picking a port range for your scan. It isn't optimal to try to probe every TCP or UDP port on every host. Nessus's default list of ports typically identifies a majority of the services and applications on the network. However, the few actively used services that Nessus might miss with a default network scan are easily found with the PVS.

Another shortcoming is when users make mistakes when configuring a set of scan targets. Nessus will perform scans of any targets assigned to it, but often organizations develop a scan list and never check it to see if it is accurate or updated. A typical optimization is to develop an actual scan list of targets instead of a network range. New hosts may be added to the list that have not been previously scanned. With passive monitoring, discovery of new hosts in these ranges is automatic.

Here are some recommended best-practices for working with data from the Passive Vulnerability Scanner:

- Send data from the PVS to its own repository. This simplifies reporting, provides a separate policy for expiring older data and facilitates comparisons with actively found system data.

- Use an LCE to log real-time forensic data from each PVS.
- Configure PVS with a network range in-line with the organization it is monitoring. Otherwise, it will over-report or under-report the actual amount of systems and vulnerabilities present.
- Automatically mark any vulnerability observed from a network session that originated outside of the network as being externally facing. This type of filter can quickly identify vulnerabilities that are exposed to the boundary.
- Record client-side vulnerabilities with a port of zero to clarify reports, dashboard elements and alerts.
- Use PVS to collect a list of all observed HTTP user-agent strings to identify many types of software, installed applications on smart phones and some types of malware.

Each PVS can also be configured to log client side port usage as well as trust relationships. Both have different levels of resources needed to support this type of tracking. For client side port usage, PVS will identify which ports a given host browses on. For trust relationships, PVS will actually track which port two given hosts communicate on. This is also useful because an IP address of 0.0.0.0 will be used for any set of outbound connections. Combinations of client side usage, as trust relationships and external connections are very useful for modeling and understanding how the various nodes on a network communicate with each other.

Distributed PVS Sensors

Benefit *– Multiple collection points can leverage internal and external network traffic.*

Quick Win *– Some boundary monitoring can be achieved, even with sensors placed "deep" within an enclave.*

Organizations that deploy multiple PVS sensors leverage one of two licensing models for PVS. The first is a single instance in which the PVS is licensed to run and monitor an unlimited set of IP addresses. The second license allows multiple PVS sensors to be deployed for monitoring a given Class B or Class C network range. PVS licenses limited to IP addresses defined by a network range can be deployed *anywhere* inside or outside a network.

Multiple PVS deployments can be leveraged to cover multiple network choke points. In some cases, it may make sense to unify multiple PVS sensor data into one repository and separate them in others. For example, consider an enterprise where three PVS sensors are all monitoring three discrete networks. If each PVS sensor is tuned to the network range of its discrete network, there will never be an overlapping set of vulnerability data. However, if each PVS sensor was tuned to look at the entire network range, then a given PVS might report information about hosts in multiple network segments. This may cause some reflection (over-reporting of the same vulnerability) or conflicting reports. Tenable recommends that each PVS be tuned to as discrete of a network as possible to avoid this.

In addition, for any given network range monitored by the PVS, if the discovered vulnerability was found from a network session that originated from an external connection, the report will mark the vulnerability as externally facing. With our previous example of three discrete internal enclaves and three PVS sensors, each PVS will consider any IP outside of its network to be an external source. If the network is the discrete enclave, then this will include IP addresses from the other two enclaves. If the network is the entire

network range of the organizations, then it will not report IPs from the organization as being "external".

Multiple RFC1918 networks can be monitored with distributed PVS sensors. Each deployed PVS must have network connectivity to a SecurityCenter, but it can sniff any set of IP addresses. As was discussed with active scanning and RFC1918 addressing, placing the results from these PVS sensors into their own repositories is a great method to manage overlapping RFC1918 ranges.

When deploying distributed PVS sensors it is important is consider your sniffing infrastructure. If your organization has deployed a dedicated sniffing infrastructure to monitor network traffic, a PVS may simply be able to plug into this. If not, you may need to deploy monitoring hardware or perform a configuration change to your network devices for the PVS to be able to monitor the desired network traffic.

Passive Boundary Monitoring

Benefit *– All inbound and outbound ports for each node can be tracked in near real-time.*

Quick Win *– Full external port scans won't be required anymore.*

Another strategy to leverage multiple PVS sensors is to deploy them to monitor your network boundaries. A boundary can be any type of perimeter defined by an access control policy. Examples include firewalls, screen routers, spam gateways, VPNs and much more.

Tenable has described how active and passive scanning can be leveraged for boundary monitoring in the "Firewalls and Boundary Auditing" whitepaper available at http://www.nessus.org/expert-resources/whitepapers.

When deployed in front of or behind any type of perimeter, each PVS can report on:

- All externally facing system addresses
- All externally facing services
- Any vulnerability data associated with these addresses and services
- Any exploitable vulnerability associated with these addresses and services
- All outbound client-side ports allowed
- All hosts that perform connections to the Internet

SecurityCenter can collect this data into various repositories such that any type of boundary can be analyzed and audited. This includes very simple boundaries such as an Internet facing firewall, more complex DMZs that have three boundaries and finally, very complex systems of trust between various enclaves within an organization.

Deploying a set of PVS sensors provides continuous verification of access control rules and independent proof of access control between various enclaves and assets within an organization.

DISTRIBUTED LOG CORRELATION

There is no best way to leverage log analysis and event monitoring for large scale enterprise reporting of security and compliance. In this section we outline several capabilities of the Log Correlation Engine, how the use of collection agents can enhance the types of logs and analysis and finally, how distributing multiple LCEs can not only enhance performance, but increase your storage and capacity to detect issues.

Deploying a Single LCE Instance

Benefit *– Automatically organize and analyze a wide variety of log sources for security and compliance monitoring.*

Quick Win *– All gathered logs can be searched by any authorized users.*

Adding a Log Correlation Engine to a SecurityCenter deployment allows organizations to combine event analysis with system analysis. Any type of log data sent to the LCE is available for search, can be normalized for analytic and reporting tools and will automatically perform the following correlations:

- Associating any log to a given user.
- Correlating intrusion detection logs with vulnerabilities from scanning, patch auditing or passive monitoring.
- Automatic identification of new types of logs that have never been seen before from a given server or desktop.
- Identification of events, such as DNS failure lookups or high CPU spikes, that occur continuously.
- Identification of statistical increases in any type of log including network activity, denied firewall events, login failures and much more.

Tenable has produced an extensive list of guidelines for gathering logs, interpreting the logs and generating reports and alerts in our "Log Correlation Engine Best Practices" paper. This 50 page technical document shows how LCE users can get the most benefit from their log data and is available from the Tenable web site at http://www.nessus.org/expert-resources/whitepapers.

.

A key component of system monitoring is to not only collect logs about the system and its activities, but to unify log data with vulnerability and configuration auditing. This ensures that any response to a potential security or compliance issue has all of the relevant data at hand.

SecurityCenter users can be segregated to have the ability to just perform log searches or to also work with normalized event data. When working with event data, a user can be limited to events that have occurred with IP addresses associated with their access. This enables organizations to re-leverage IDS, firewall, VPN and other types of IP-based security devices to share with security share holders in each enclave while maintaining separation of duties.

Adding LCE Agents for Enhanced Collection

Benefit *– LCE agents are available to collect logs from a variety of operating systems and technologies.*

Quick Win #1 *– Sending network data into an LCE effectively turns it into a network-based anomaly detection system.*

Quick Win #2 *– Sending Unix or Windows process accounting data helps identify software in use and exploits from viruses and other incursions.*

Organizations that deploy an LCE typically send logs via Syslog . These are the easiest forms of log gathering to configure. Simply add the IP address of your LCE and log away. However, there are many security issues with sending logs in the clear and since the messages are sent over UDP, there is no guarantee that a message will arrive.

Each LCE can be configured to receive logs over a secure and encrypted API from one or more "LCE Clients". LCE Clients are available that can collect a wide variety of logs including:

- Local Windows event logs
- Unix log files
- NetFlow data from one or more sources
- Remotely collect Windows event logs via WMI queries
- Checkpoint OPSEC firewall log data
- Cisco SDEE intrusion event log data

In addition to gathering the log file data noted above, LCE clients can also perform the following actions:

- Tail files, perform file integrity checking and gather CPU and disk information on Windows systems
- Gather process accounting data and perform file integrity checking on Unix systems
- Gather network session data by sniffing packets on one or more network interfaces

It is important to test before adding agents to production or supported desktop systems to make sure these solutions work as expected in your environment. Tenable typically offers more than one way to obtain data. For example, Windows event logs can be gathered from

a local Windows LCE agent, a Windows LCE agent can perform WMI queries to other remote Windows servers and there is also a dedicated Linux WMI client to obtain the same data.

The LCE Unix and Windows clients provide a great level of forensic monitoring. They gather process accounting events as well as perform file integrity checking. The LCE can summarize all commands run by every system and user and identify which commands were new for the system or the network. This is an excellent resource to identify changes in behavior as well as compromised systems. The combination of file integrity checking, command accounting and system integrity monitoring presents many possibilities for web server monitoring, identification of stealthy virus infections and searching for backdoors and rootkits.

LCE clients for network monitoring are also useful for tracking all network level sessions. The LCE has clients for NetFlow analysis as well as direct network session monitoring. Many of Tenable's customers that use both the PVS and the LCE place the network client on their deployed PVS platforms and monitor traffic there as well. This enables logging of all file transfers as well as network sessions into the LCE. The LCE logs short network sessions in one discrete log and specifically identifies network sessions that last a long time or transfer a large amount of data. When combined with other logs, such as web server or firewall logs, simple port filtering can paint a succinct picture of any type of network traffic.

In addition to the guidance outlined in the "Log Correlation Engine Best Practices" paper, Tenable highly recommends monitoring the configuration and status of the LCE clients. For example, many Tenable customers leverage the full Unified Security Monitoring architecture to:

- Alert if the expected number of LCE clients is less than it should be.
- Audit LCE client configurations with Nessus configuration audits.
- Ensure that Unix and Windows operating systems are configured correctly with logging enabled.
- Scan or monitor network traffic for evidence of hosts that don't have logging enabled.
- Leverage error, memory or CPU reporting within the LCE to monitor the status of SecurityCenter, Nessus or the Log Correlation Engine.

Deploying Multiple LCEs for Distributed Correlation

Benefit *– Multiple LCEs increases the amount of logs that can be stored, normalized and searched.*

Quick Win *– Enclaves that do not want to send their logs to a central location can operate their own LCE.*

SecurityCenter can manage multiple LCEs that can be used to increase the performance and reliability of your log collection process. Placing LCEs in various locations can also minimize the need to send logs to a central collection point. Separating LCEs can also create log storage silos that maintain sensitive data that can only be queried by authorized SecurityCenter users.

A typical LCE deployed on modern hardware (e.g., 3 Ghz CPU, 4 GB of memory and multiple terabytes of hard drive space) can sustain a rate of one thousand logs per second for long periods of time. Tenable has many customers who centralize all of their logs to one LCE.

These logs are sent via **`syslog`** or LCE clients. However, for larger customers and networks, there are many scalability issues to consider.

- What happens if your number of events per second rate is greater than 1,000?
- How many remote LCE clients can a single LCE support?
- What happens if you have a large number of users making many simultaneous queries to the LCE?

Each LCE can be queried simultaneously with other LCEs such that SecurityCenter users can make one query across all LCEs they have access to. For example, consider an effort to collect NetFlow records that generate close to 5,000 events per second from 20 different collection points. It would be easy to deploy five LCEs and give each LCE a NetFlow client that subscribed to four sources. Within the SecruityCenter, each LCE would be associated with an organization and any user of that organization could be given access as needed. Users could then analyze the NetFlow data across all LCEs, or they could select just one LCE to focus their queries.

Another common practice performed by our customers is to deploy LCEs for different types of technologies. We've seen customers deploy one or more LCEs dedicated to network security and network monitoring to collect NetFlow, intrusion detection and firewall logs while also deploying one or more LCEs to collect server logs from email, web, DNS and other devices.

If you have a sensitive enclave of logs, deploying an LCE to collect just the data from those systems may offer an alternative to centralizing all logs to ensure that sensitive data can't be viewed by unauthorized users. Logs containing potential customer data, credit card numbers, health care information, Social Security numbers and anything else considered "sensitive" must be logged in a manner to prevent access by unauthorized users.

SecurityCenter provides the ability to create roles for users who can and cannot search logs per LCE. However, SecurityCenter cannot prevent a user who is authorized to search logs on an LCE from finding a sensitive log. Having a separate LCE that contains the sensitive or potentially sensitive data enables this level of access control.

Finally, the ability to search through terabytes of logs by leveraging distributed event normalization is an excellent method to monitor your enterprise. Scaling a single LCE's ability to perform first seen event detection, identify anomalies and identify continuous activity with high-speed queries across multiple LCEs allows for very rich SecurityCenter dashboards, reports and alerts to be generated.

Event Correlation Maturity

Every organization that needs to monitor their security and compliance status can benefit from the ability to collect logs and events. However, there are many different strategies that can be leveraged depending on the environment. If the infrastructure to collect and search logs is not solid, then any type of business requirement to perform correlation or alerting could be called into question. On the other hand, simply performing correlation for the sake of creating alerts without any type of input from the organization is of limited use.

In this section we discuss best practices for deploying more and more sophisticated correlation levels. Each level requires slightly more effort than the previous method, but each provides very distinct alerting and monitoring for the business.

Correlating Attacks with Vulnerabilities

Benefit *– Create high quality security alerts regardless of what IDS/IPS is in use.*

Quick Win *– Unstructured intrusion events can be correlated with generic vulnerability scans with minimal configuration.*

The Log Correlation Engine performs standards-based vulnerability correlation with events from a variety of industry leading intrusion detection systems such as Snort/Sourcefire, TippingPoint/HP, IDP/Juniper, Proventia/IBM and Cisco IDS. As IDS events arrive at the LCE, it will mark them with a flag if the event correlates with a known vulnerability on the target system and port. Correlation is based on the vulnerability not the operating system technology.

When an intrusion detection event arrives for normalization by the LCE, it considers the event's metadata, such as the associated CVE vulnerability tag, the Symantec Bugtraq ID (BID) and in cases where supported, the Tenable Nessus plugin ID.

In many cases, Tenable's customers are already performing routine or continuous vulnerability assessments. This provides the LCE the absolute latest vulnerability data that it can have. For example, on a typical Microsoft Tuesday, most IDS and vulnerability scanning vendors ship new content to their customers on the same day. This means that an IDS sensor will likely have up to date detection rules in operation. However, a vulnerability

scanner may have up-to-date rules, but a scan needs to be run to generate results to share with the LCE.

To combat this, Tenable recommends several best practices:

- Leverage the real-time client and server capabilities of the PVS to report on the same type of network traffic an IDS device would monitor.
- Leverage credentials to scan more often. A system audited with credentials takes less time and generates less network traffic than an un-credentialed scan. Organizations that audit more often not only detect issues earlier, they have more relevant data for VA/IDS correlation.
- Schedule your weekly scans on Wednesdays to quickly detect vulnerabilities disclosed on Microsoft "Patch Tuesday".

When deploying this type of alerting within an organization, there are several issues to consider:

- Who should be receiving these alerts? Is it the actual system administrator owners, a security response team, a central security organization or some other group?
- Do these high-quality alerts require any different type of handling due to compliance regulations? For example, a correlated IDS/VA event from a system subject to PCI or one that holds customer data may be investigated differently than a less important network resource.
- Is there a feedback cycle that needs to be implemented between the scanning team and the IDS team? Often, Tenable's customers will report that some VA/IDS correlated events indicate poor signatures or "false positives" from various sensors. Tuning the sensors improves the VA/IDS correlation process.

A very basic best practice with any technology of this kind is to monitor the results prior to creating real alerts, tickets or pages. Since intrusion detection technology is based on stimulation from external sources, it is difficult to predict any type of alert rate or accuracy. The VA/IDS correlation cuts down a tremendous amount of alerts that need to be analyzed.

Enabling this functionality within the LCE is also trivial. All that is required is that the source IP addresses of your IDS sensors are specifically recognized by the LCE as being a source of potential IDS/VA correlation events.

Tracking Events by Users

Benefit *– Any normalized activity can be associated to a user on your network.*

Quick Win *– If a user is investigated, within seconds an accurate history of all network, system, security, authentication and other types of logs can be produced.*

The LCE normalizes events to a variety of well-known types. These events also are associated with IP addresses. The LCE can associate these IP addresses dynamically with user activity on the network.

LCE needs to receive data from one or more log sources that contain authentication information to perform user event correlation for an organization. For example, typical VPN authentication, Windows domain authentication or even authenticated IMAP email access all have logs that state a user authenticated from a certain IP address. Once configured to do

this, the LCE will associate any other normalized event whose IP address includes that of a known user. This enables dynamic learning of user IDs and performing analysis of user activity.

Deciding on which authentication sources to configure within an LCE could be the most difficult. Most organizations have multiple forms of authentication, and perhaps competing "single sign on" log sources. The good news is that as long as the LCE receives a log, it can pivot on multiple authentication sources.

Once user to event correlation is available, several new types of reports, dashboards and other capabilities can be achieved, including:

- Pivoting from any interesting event to the user or users involved with it.
- Sorting events by user such as top attacked user, top user with errors, top user with anomalies and more.
- The ability to quickly summarize all activity of users under investigation, including which programs they ran, which web site domains they visited, what their network traffic was like and much more.
- Leveraging the LCE's ability to learn active user names in order to demonstrate which users have accessed various systems.
- Determining which users have had access violations such as denied firewall events, login failures or anomalies within these types of logs.
- Tracking which users have accessed files such as PDF or other office documents.
- Learning if users are accessing corporate information with mobile devices from iPhones, Android devices and other types of platforms.
- Tracking which users leverage or make use of social networking sites such as Facebook.
- Keeping a history of which users were active on which IP addresses. This ability can be used to track a user in DHCP environments or roaming campus wireless networks.

The LCE is also very good about tracking which user accounts it has seen log into a server or system. This can help with transitive trust issues. For example, a user may be authorized to log into a Windows system as "Administrator" or a Unix system as "root". Differentiating which users logged into each system can be difficult. However, with user to IP address tracking, it can be observed which users logged into a system as a privileged user.

Anomaly Detection

Benefit *– All normalized logs are considered for first time seen, continuous or statistical alerting.*

Quick Win *– Leveraging the SecurityCenter dashboard to display anomalies for each asset or system is easy to accomplish and provides strong indicators of changes that impact compliance or compromise.*

For each log, the LCE will profile each host and look for the following four types of anomalies:

- Log entries that occur for the first time.
- Log entries that occur at a rate that is statistically significant from previous rates.
- Log entries that are continuous for periods of time in excess of 20 minutes.

- LCE anomalies that have been specifically programmed to automate detection of brute force password guessing, correlation with hostile IP addresses, identification of new programs, identification of change and much more.

The LCE enables this sort of tracking "out of the box". The complex part for most organizations is determining which of these alerts it makes sense to create dashboards, reports and alerts for.

Many of these concepts are discussed in depth in our "Log Correlation Engine Best Practices" paper. This 50 page technical document shows how LCE users can get the most benefit from working and collecting logs. This paper is available from the Tenable web site at http://www.nessus.org/expert-resources/whitepapers.

Let's consider an example asset class of five DNS servers for a small enterprise network. The LCE in place is collecting logs from the DNS servers, as well as NetFlow, firewall and intrusion detection logs. One of the DNS servers is compromised and a botnet is installed on it. There are many possible indicators that can be generated by the LCE.

- As the DNS servers are attacked, first time seen logs may be generated by the IDSs when observing the inbound attacks, as error messages from the DNS server and as error messages from the underlying OS. If the botnet starts attacking outbound, it's possible first time seen denied firewall events will occur as well as additional detection of unique outbound IDS events.
- If the attacks were sustained in nature, such as brute force password guessing, or remote exploits that require thousands of tries before succeeding, it's possible the LCE would have detected one or more types of continuous events.
- If the error rates (if any), IDS events, login failures, outbound denied firewall logs, NetFlow, etc. were statistically different before the attack or after infection, there could be dozens of potential alerts that indicate some sort of change in activity level has occurred.
- If a botnet ran any new programs in any new ways, if it installed software or users, if it caused any types of system crashes or kernel issues, these events would all be tracked and reported by the LCE.

Performing any type of analysis manually is very difficult. With SecurityCenter, dashboards that have various details of each type of these anomalies are available for quick installation. The only difficult part for most organizations is to determine what the names of their assets are and which ones they want to create dashboards for.

Many Tenable customers leverage multiple views. For example, there may be a high level dashboard that simply shows "DNS Servers" and if they've had any medium or high levels of anomalies in the last 48 hours. A second dashboard could have more detail, such as a display for each DNS sever, and separate columns for each type of detailed anomaly such as statistical increases in NetFlow, first time seen error events and continuous intrusion events.

For example, in the following example dashboard, "Never Before Seen" filters for application logs, events that indicate change, errors, logins and system events have been charted out for the past seven days. Multiple hosts, assets and entire networks can be used as filters to provide rows with "at a glance" views on the status of each column.

As previously noted, the "Log Correlation Engine Best Practices" paper documents many of the common strategies used by organizations to leverage reporting and alerting to monitor security and compliance. Depending on the types of logs collected, the types of technologies, the amount of logs and many other factors, what works for one situation may not work for others.

Tenable has designed the LCE to perform as much correlation at the "IP" level as possible. This means that finding anomalies on a per-IP basis occurs with the same logic for one web server on a small network as for hundreds of web servers on a busy large ISP. This also means that the any size network can leverage the LCE to look for signs of new events, continuous events or statistical increases in events across any system or log source.

Many pre-build dashboards that leverage the LCE's ability to detect and filter on events have been made available (http://blog.tenable.com/sc4dashboards/) to Tenable customers so they can mix and match them to solve a wide variety of monitoring issues. Customers who wish to share their dashboards can also do this on the Tenable Discussion Forums (https://discussions.nessus.org) in the LCE section.

CONCLUSION

This paper provides a roadmap for any Tenable customer leveraging any or all of our solutions. Many of these solutions are available to you right now and don't require any additional procurement of new licenses from Tenable. For each best practice of Tenable's Unified Security Monitoring platform, clear benefits and quick wins can be obtained by your organization.

This paper was also developed with feedback from many of our Tenable customers. If you have feedback you'd like to share with us, please contact us directly through your support or sales teams or post your suggestions and ideas to the Tenable Discussion Forums at https://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