Prevention of Information Security Breaches: Minimizing the Enterprise Attack Surface

A New Net Technologies Whitepaper

Mark Kedgley

CTO - New Net Technologies

www.nntws.com

NNT™
SECURITY THROUGH SYSTEM INTEGRITY
NOW PART OF netwrix

Abstract

When developing an information security strategy, prevention is still better than cure. Organizations get breached because they leave gaps in their defenses. Gaps that are subsequently exploited.

The aggregate of these vulnerabilities is referred to as the Enterprise Attack Surface.

This whitepaper looks at how you can assess and measure your Enterprise Attack Surface and what you can do to minimize your exposure to cyber attacks.

If there was a security product that gave 100% protection, wouldn't we all be using it?

Instead the breaches just keep coming –

- Target - 40 million payment cardholder details stolen and in total, personal information for more than 70 Million individuals has been stolen
- The Australian Information Commissioner reports a 20% increase in reported personal information breaches
- The discovery of the Windigo trojan and previously the Hand of Thief malware shows that there are now more occurrences of malware to non-Windows platforms

For many, security defenses are falling well short of the effectiveness needed.

The Enterprise Attack Surface

Organizations get breached because they leave gaps in their defenses. Gaps that are subsequently exploited.

Figure 1: The Enterprise Attack Surface

Information Systems comprise a range of database systems, operating system platforms, appliances and network devices. All of these components are vulnerable to attacks leading to a breach of data security.

The sum of these weak spots comprises the Enterprise Attack Surface.

Protection is provided by layered, overlapping defense measures and operational procedures. Any gap leaves the enterprise with a Cyber Security Achilles Heel, a weak spot or vulnerability that an attacker can exploit. We call this the Enterprise Attack Surface.

> “***What happened at Target?***
>
> In mid-December, we learned criminals forced their way into our system, gaining access to guest credit and debit card information. The investigation has recently determined that certain guest information was taken. That included names, mailing addresses, email addresses or phone numbers. We have partnered with a leading third-party forensics firm who is thoroughly investigating the breach
>
> ***How many guests were affected by the additional stolen information?***
>
> Up to 70 million individuals may be affected.
>
> ***How many credit or debit cards were impacted?***
>
> Approximately 40 million credit and debit card accounts may have been impacted between Nov. 27 and Dec. 15, 2013
>
> *source: target.com Jan 2014*”

How prone is your organization to an attack?

How do you measure or assess this and, working on the basis that what you can't measure, you can't ever manage, this should be a priority Information Management metric.

There are a variety of options available but the most commonly used is the Vulnerability Scanner. Tools like Qualys or Nessus provide an automated means to scan systems, effectively logging onto each and every device then running through a security audit checklist.

The report generated provides a list of vulnerabilities that are present for the system and an overall percentage score, sometimes with a weighting factor applied based on an attributed severity for each vulnerability.

The concept behind this kind of Vulnerability Scanner is that it is providing a transient measure of security at that point in time. In one respect, with new software vulnerabilities being discovered every day, there is merit in assuming that yesterdays' scan results are out of date and that you should always be starting with a fresh view.

However, this philosophy and approach results in a poor solution when considering the broader requirements for both vulnerability and configuration management.

These are explored in the following section, with solutions outlined to deal with the issues highlighted.

> “ 'The Traditional Vulnerability Scanner approach results in a poor solution when considering the broader requirements for both vulnerability and configuration management' ”

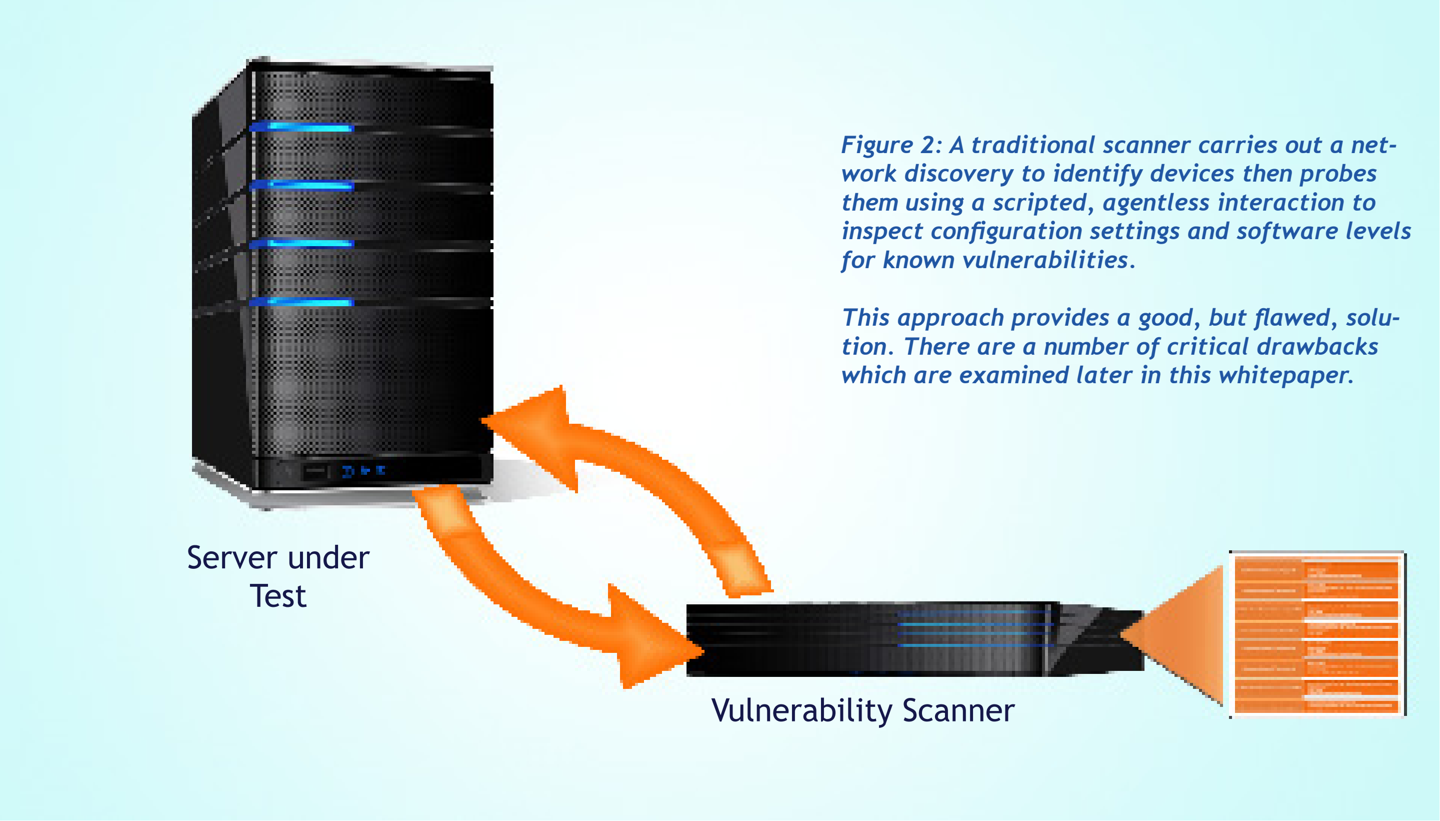

Figure 2: A traditional scanner carries out a network discovery to identify devices then probes them using a scripted, agentless interaction to inspect configuration settings and software levels for known vulnerabilities.

This approach provides a good, but flawed, solution. There are a number of critical drawbacks which are examined later in this whitepaper.

Software Flaw and Configuration-based Vulnerabilities

Organizations such as NIST (see http://csrc.nist.gov/publications/nistir/ir7502/nistir-7502_CCSS.pdf) define these in much detail in order to draw the distinction between '*software flaw vulnerabilities*' and what they term as '*security configuration issue vulnerabilities*'. Software Flaw vulnerabilities are defects in a software component that inadvertently provide a potential exploit. When a Software Flaw Vulnerability is discovered, a patch or update is required to replace the flawed software component with a modified version to eliminate the security exploit.

By contrast, Security Configuration Issue Vulnerabilities present a very different challenge. To give a simple example of a security configuration issue vulnerability, consider the use of a common setting on all platforms - maximum password age.

Who reading this welcomes being forced to change their password? It's always hassle and there is the counter argument which says forcing users to change passwords more frequently may *increase* risk as users are more likely to be writing down passwords in order to remember them!

However, recent data thefts at eBay (see http://www.scmagazine.com/ebay-hacked-all-users-asked-to-change-passwords/article/347967/) and, indirectly at Target, could have been prevented or their impact lessened through a more short-lived password age. Put simply, a more short-lived password has far less potential for damage than one that is valid for months.

Security versus Convenience, Ease of Use versus System Defenses

Therein lies the security vulnerability associated with password age. The example neatly encapsulates the core issue of security configuration vulnerabilities and why they represent a bigger problem than software vulnerabilities.

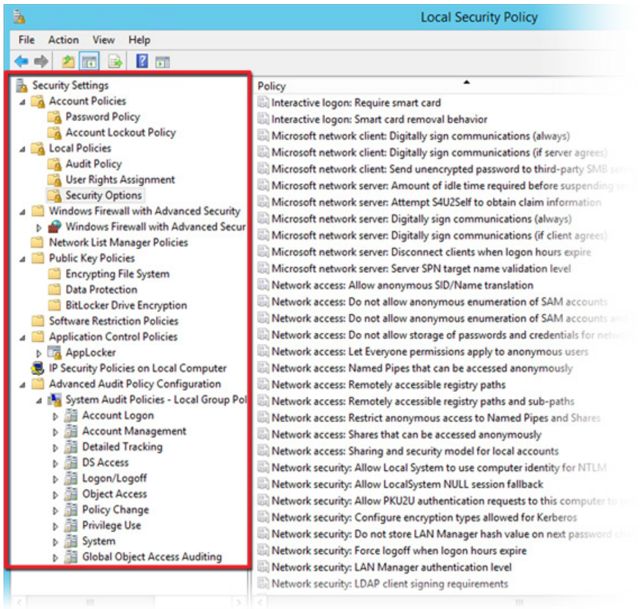

- Security Settings
  - Account Policies
    - Password Policy
    - Account Lockout Policy
  - Local Policies
    - Audit Policy
    - User Rights Assignment
    - Security Options
  - Windows Firewall with Advanced Security
    - Windows Firewall with Advanced Secur
  - Network List Manager Policies
  - Public Key Policies
    - Encrypting File System
    - Data Protection
    - BitLocker Drive Encryption
  - Software Restriction Policies
  - Application Control Policies
    - AppLocker
  - IP Security Policies on Local Computer
  - Advanced Audit Policy Configuration
    - System Audit Policies - Local Group Pol
      - Account Logon
      - Account Management
      - Detailed Tracking
      - DS Access
      - Logon/Logoff
      - Object Access
      - Policy Change
      - Privilege Use
      - System
      - Global Object Access Auditing

Interactive logon: Require smart card
Interactive logon: Smart card removal behavior
Microsoft network client: Digitally sign communications (always)
Microsoft network client: Digitally sign communications (if server agrees)
Microsoft network client: Send unencrypted password to third-party SMB ser
Microsoft network server: Amount of idle time required before suspending se
Microsoft network server: Attempt S4U2Self to obtain claim information
Microsoft network server: Digitally sign communications (always)
Microsoft network server: Digitally sign communications (if client agrees)
Microsoft network server: Disconnect clients when logon hours expire
Microsoft network server: Server SPN target name validation level
Network access: Allow anonymous SID/Name translation
Network access: Do not allow anonymous enumeration of SAM accounts
Network access: Do not allow anonymous enumeration of SAM accounts and
Network access: Do not allow storage of passwords and credentials for netwo
Network access: Let Everyone permissions apply to anonymous users
Network access: Named Pipes that can be accessed anonymously
Network access: Remotely accessible registry paths
Network access: Remotely accessible registry paths and sub-paths
Network access: Restrict anonymous access to Named Pipes and Shares
Network access: Shares that can be accessed anonymously
Network access: Sharing and security model for local accounts
Network security: Allow Local System to use computer identity for NTLM
Network security: Allow LocalSystem NULL session fallback
Network security: Allow PKU2U authentication requests to this computer to u
Network security: Configure encryption types allowed for Kerberos
Network security: Do not store LAN Manager hash value on next password ch
Network security: Force logoff when logon hours expire
Network security: LAN Manager authentication level
Network security: LDAP client signing requirements

Figure 3: Windows Security and Audit Policy comprises hundreds of settings to mitigate vulnerabilities - but they must be configured correctly to be effective

Any manufacturer of a software product aims to provide something that is easy to use, quick to deliver results and requiring as little user intervention as possible.

All of which, of course, are in direct conflict with the objectives of maximizing system security. As a consequence, default security configuration settings for any operating system, database system, or network device are typically weak.

The situation is compounded further in that, by definition, default settings are known to everyone including would-be hackers, which further lessens the effectiveness of default security settings. For example, an SNMP community string default of 'public' is common knowledge, with the net effect that it provides no protection whatsoever.

Now consider that a contemporary Windows platform has over 500 such settings spread across the Security Policy, the need for automation of the auditing function is essential in order to ensure vulnerabilities are kept at bay. A decent Vulnerability Scanner will highlight where vulnerabilities or deviation from compliance exist and usually provide advice for mitigating the vulnerability.

But if default settings typically leave a system prone to exploit, which settings should be used to provide the maximum level of security?

Configuration Hardening Standards - Does a Truly 'Authoritative' Source Exist?

Naturally, the manufacturer of the operating system, application, appliance or database system will typically provide security configuration best practice guidelines, for example the Microsoft Threat and Countermeasures Guide (see http://technet.microsoft.com/en-us/library/hh125921%28v=ws.10%29.aspx) is very comprehensive for 2008R2 and Win7, and similar resources can be found for other Windows versions and applications.

Move outside of the Microsoft arena, and most mainstream manufacturers provide their equivalent guides, for example RedHat's "Red Hat Enterprise Linux 6 Security Guide" (see https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/pdf/Security_Guide/Red_Hat_Enterprise_Linux-6-Security_Guide-en-US.pdf). The problem is that there is not - naturally enough - any consistency between the various manufacturers' content in terms of detail and presentation.

> *"CIS Benchmark material has been developed on a definitive, consensus basis to combine Manufacturer best practice guidance with input from security researchers and academic institutions, creating the most comprehensive source of authoritative guidance available"*

Alternatively, as mentioned in the previous section, NIST (National Institute of Standards and Technology, the measurement standards laboratory for the US) provides guidance in this area, with configuration checklists now provided via the National Vulnerability Database (see http://web.nvd.nist.gov/view/ncp/repository). However, while there is a good deal of content for a range of OS, applications, and database systems, most of this is presented in SCAP format only, requiring an SCAP compatible scanner to access the content.

Therefore for many security professionals, the preferred source of material will be the Center For Internet Security, or CIS (see http://benchmarks.cisecurity.org/downloads/benchmarks/).

The CIS Benchmarks repository provides consistently presented hardening checklists in both PDF and OVAL/SCAP formats.

The PDF Benchmarks provide detailed background on each hardened configuration setting recommended, the potential threat and the relevant commands required to both audit and remediate the vulnerability.

Best of all, the CIS Benchmark material has been developed on a definitive, consensus basis to combine Manufacturer best practice guidance with input from security researchers and academic institutions, creating the most comprehensive source of authoritative guidance available.

CIS SECURITY BENCHMARKS • MEMBER •

CIS SECURITY BENCHMARKS • CERTIFIED •

Note: NNT are one of a handful of CIS Certified Vendors, automating the auditing of systems for compliance with CIS Benchmark Checklist settings

Vulnerability Scanner - A Flawed Solution?

Earlier on it was suggested that a Vulnerability Scanner provided a good, but flawed, solution when seeking to minimize the Enterprise Attack Surface. Here we explore what the issues are that render the Vulnerability Scanner an incomplete solution.

1. How secure is the Enterprise right now, and will it still be as secure in 24 hours time?

Returning to the 'Security versus Convenience' conflict from earlier, the challenge in running secure IT operations is that configuration settings can 'drift' over time.

For example, an engineer needs to update the eCommerce Web Site to include some additional pages: *Get prompted for UAC every step, or speed things up by disabling this feature? Use sudo for every command, or just enable remote logon for root? Connect via the jump server, or configure a local account for access? Enable the Installer Service and disable the firewall temporarily while the work is undertaken?*

Whether it stems from time pressures to get a job done more quickly, from a process of elimination when troubleshooting, or through corner-cutting for convenience, security settings may become weakened during normal operational activities. Of course, the next time the scheduled scan runs this increased Attack Surface will be identified and addressed. However, up until that time, the Enterprise is left in a vulnerable state, more prone to exploits and security breaches.

The only way to address this issue is to schedule scans to be run more frequently.

Figure 4: 'Security versus Convenience' - the more hardened and secure a system, the more difficult it becomes for legitimate, desirable maintenance to be performed.

Engineers are human - and may just cut corners to make life easier for themselves, while at the same time, increasing the Enterprise Attack Surface...

2. The scan results show hardening measures are in place, but given that zero day malware, phishing attacks, and insider threats may yet result in a breach, how do I protect systems against these attack vectors?

Bad news! Even with systems hardened in line with CIS Checklist settings, the enterprise still can't be guaranteed to be 100% secure. New exploits are revealed all the time and zero day malware will evade AV defenses, while phishing attacks or 'inside man' threats can bypass security measures by exploiting employees' trusted-status.

This is where the Vulnerability Scanner begins to look very one-dimensional in its contribution to minimizing the Enterprise attack surface. Given that threats go beyond the exploitation of vulnerabilities, the need for a more comprehensive solution becomes apparent.

Vulnerability Scanner - A Flawed Solution? Continued...

The concept of providing a 'Host Intrusion Detection' function goes beyond the verification of compliance with a hardened build standard. Identification of malware requires a detailed inspection of files and the filesystem as a whole to provide visibility of new or modified file. This extended remit increases the work required of the scanner by an exponential factor.

For this malware-detection function, we now need the scanner to analyze tens of thousands of files rather than the few hundred needed to verify compliance. Furthermore, in order to detect Trojans masquerading as legitimate system files, the scanner needs a 'DNA fingerprint' of each file to be held. The security industry consensus is that the only way to achieve this is via a one-way cryptographic hash value being generated for each file.

3. How much does vulnerability and compliance management 'cost' in terms of Host Resources?

The previous two issues - the requirements for both a host intrusion detection function and for frequent (ideally continuous) verification of compliance with a hardened build standard - both have implications in terms of resource requirements.

There is always a price for monitoring - the resource requirements when scanning a host depends on the scope and breadth of the analysis being performed. A typical compliance report will take a few minutes to complete per host but in a large enterprise with hundreds or thousands of hosts, the aggregate resource usage in both host and network bandwidth terms quickly becomes significant, especially for geographically remote sites.

However, if a wide scale file integrity check is factored in to serve as a host intrusion detection safeguard, the resource issue becomes considerably more acute.

Three issues come into play - the first is the sheer numbers of files involved. A typical scan of all system files on a contemporary Windows server will amount to at least 15,000 files by the time all driver files, executables, and DLLs across the System 32, SysWOW64 and Program Files/Program Files (x86) folders. Just to inspect the files and their attributes, then compare to the previous baseline record of the filesystem immediately makes the scanner task an exponentially prolonged and more resource-intensive task.

The next problem is that there is a need to provide the filehashing capability, which means a 'dissolvable' agent has to be copied across the network to the host. The agent is called 'dissolvable' because it is a transient program/binary that is removed from the host once its work is done.

But the real issue arises once filehashes are being generated. Even a 'basic' SHA1 hash will require CPU muscle and, the more files there are and the bigger the file sizes involved will all go to magnify the hit on resources.

Now you are looking at scan workloads that require careful consideration. Run this type of scan on a busy, live system and you risk compromising performance and business services.

Vulnerability Scanner - A Flawed Solution? Continued...

3. How much does vulnerability and compliance management 'cost' in terms of Host Resources? Continued...

Therefore, scheduling scans to run out of hours is the only 'safe' option, and in a busy, 24/7 operation, even these scans should be spread as sparsely as possible.

Even if the agent is permanently deployed, such as in hybrid-SIEM systems that try and provide an element of file integrity monitoring, the periodic re-scan and re-hash of the filesystem will still need to be scheduled for 'dark hours' only.

For example, one enterprise grocery store chain is known to take a whole month to scan just 600 servers, before repeating the process again for the next month.

In other words, the resulting solution for compliance enforcement and host intrusion detection can only provide monthly alerts.

If you consider Target lost 40 million payment card numbers and personal information for 70 million customers in just over two weeks, monthly scans start to look like a toothless and ineffective measure.

But if real-time detection of breaches is now just as critical as ensuring preventative defenses are in place, how can the right balance be found between scan frequency and the required host/network resource costs?

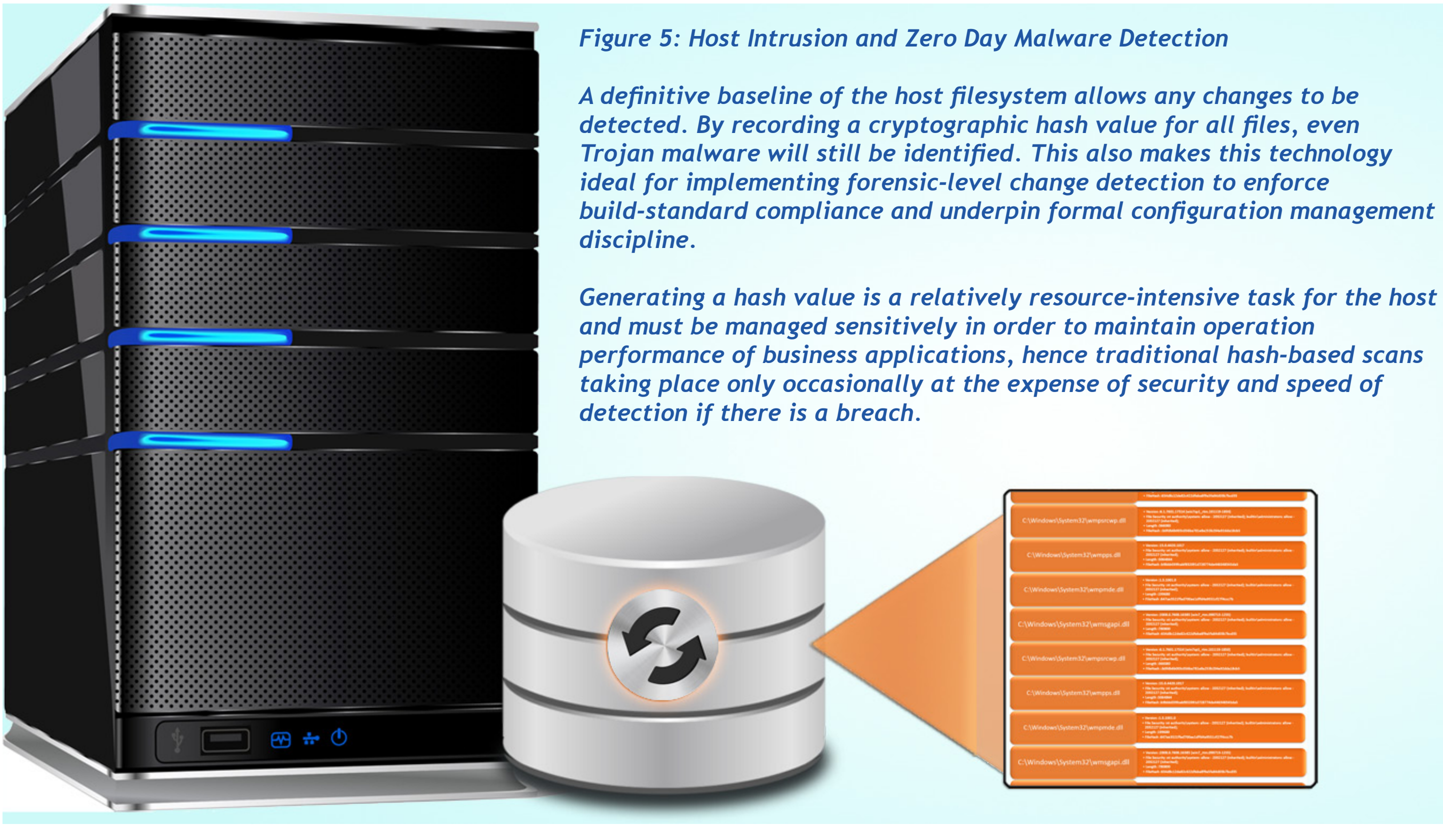
C:\Windows\System32\wmpsrcwp.dll
C:\Windows\System32\wmpps.dll
C:\Windows\System32\wmpmde.dll
C:\Windows\System32\wmsgapi.dll
C:\Windows\System32\wmpsrcwp.dll
C:\Windows\System32\wmpps.dll
C:\Windows\System32\wmpmde.dll
C:\Windows\System32\wmsgapi.dll

Figure 5: Host Intrusion and Zero Day Malware Detection

A definitive baseline of the host filesystem allows any changes to be detected. By recording a cryptographic hash value for all files, even Trojan malware will still be identified. This also makes this technology ideal for implementing forensic-level change detection to enforce build-standard compliance and underpin formal configuration management discipline.

Generating a hash value is a relatively resource-intensive task for the host and must be managed sensitively in order to maintain operation performance of business applications, hence traditional hash-based scans taking place only occasionally at the expense of security and speed of detection if there is a breach.

Real-Time Detection: Continuous, Real-Time Protection

To summarize, the three key problems with external vulnerability scanners are:

1. Results are ‘in the moment’ only - vulnerabilities will not be detected until the next scan

2. Host Intrusion/ Malware Detection not provided by standard scanner solutions as this requires a Hash-Value signature for each file to be captured

3. Generating the Hash-value signature is a Host-resource intensive task, so must be used sparingly (see problem 1 again - malware can cause damage every minute it is in place so ‘time to detection’ is critical)

The ideal solution will therefore provide a real-time, continuous detection mechanism, but without generating the repeated resource loads that a traditional Tripwire®-like agent or a SIEM/Scanner solution will.

Real-Time FIM Model

1. Run initial one-time baseline

There is no getting away from the need to use an agent if we are to use hash values which is, after all, the industry-standard for definitive file version identification.

However, in the Real-Time FIM model, the baseline only needs to be run once, providing a significant advantage over the traditional, repeated monthly re-scans.

2. Use real-time, triggered file change detection

Using a local agent on each host to provide the hash function isn’t new, whether it be a ‘dissolvable’ agent used by some scanners, or the permanently deployed variety like the traditional Tripwire® agent or SIEM system like LogRhythm®.

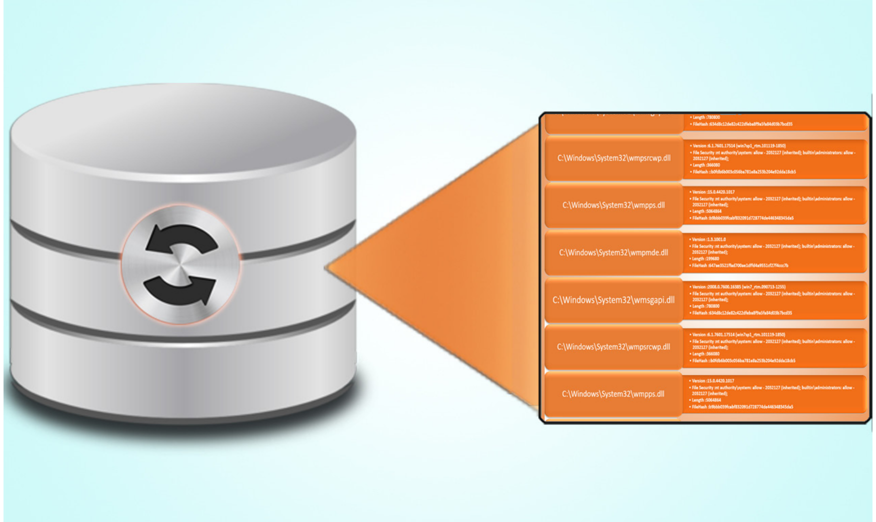
C:\Windows\System32\wmpsrcwp.dll
C:\Windows\System32\wmpps.dll
C:\Windows\System32\wmpmde.dll
C:\Windows\System32\wmsgapi.dll
C:\Windows\System32\wmpsrcwp.dll
C:\Windows\System32\wmpps.dll

***Figure 6**: Real-Time Continuous FIM provides the perfect solution for both enforcement of compliance and build-standards, and the fastest host intrusion and malware detection capability in the event of a breach.*

However, the real innovation is that, once the initial ‘fat’ baseline has been run, use a ‘file sniffer’ approach to ***ONLY*** subsequently generate a hash value when a qualifying filechange is observed.

No file changes this week? No hash generation required! No wasted load on the host.

But, if new files suddenly appear on the host, or existing files are changed as with a Trojan, the ***‘Continuous, Real-Time Protection’*** solution will pounce and analyze these new/ changed files ***BUT NO OTHER FILES***.

Instead of needing a complete re-baseline of the entire filesystem, we are now using laser-precision FIM techniques to minimize resources to the bare minimum AND at the same time providing real-time detection of threats. Talk about a perfect solution!

About NNT

NNT Change Tracker Gen provides continuous protection against known and emerging cyber security threats in an easy to use solution, offering true enterprise coverage through agent-based and agentless monitoring options.

- NNT analyzes every configurable component within your IT Estate and allows you to define a ‘Known, Good, Secure and Compliant State’ for all of your in scope systems.
- NNT-Change Tracker scans your devices and compares them to a standard policy, either user defined or based on an industry standard such as the Center for Internet Security (CIS).
- Policies can be automatically assigned based on the device type or priority via a centrally managed console.
- Gen7 is able to fully automate change approval for you, using the NNT FAST (File Approved-Safe technology) that combines unique intelligent change control knowledge base and whitelists.
- With NNT’s real-time capabilities, unlike traditional scanning or exclusively agentless technologies, potential breaches to systems or policies are spotted immediately.

NNT Change Tracker Gen 7 helps you to prevent security breaches of your systems by providing you with a powerful feature-rich, easy to use and affordable solution for validating, achieving and maintaining compliance with corporate governance or security standards.

www.nntws.com

©New Net Technologies

UK Office - Spectrum House,
Dunstable Road, Redbourn,
AL3 7PR
Tel: +44 8456 585 005

US Office - 9128 Strada Place,
Suite 10115, Naples, Florida
34108
Tel: +1-888-898-0674

Conclusion - The NNT View

You don’t have to search too far to realize that the discovery of new vulnerabilities is a daily event. Cyber security breaches are also becoming more professional and more effective, with Target proving how damaging even a 2 week breach can be.

There are some key questions to ask yourself if you are responsible for managing secure systems handling personal identification information or payment cards/financial data:

- *How well protected are your systems? Has your Enterprise Attack Surface been reduced to its absolute minimum?*
- *How comprehensive and up to date is your implementation of a Hardened Build Standard? Remember, eBay were breached by having a weak password-ageing policy, allowing much more damage to be done than might otherwise have been the case.*
- *If there was a drift from this hardened build standard, how long would you be at risk before discovering the vulnerabilities?*
- *And, if you did get breached, how long would it take you to realize this? As a result, how big would your losses be, both in direct financial terms and in terms of customer trust and competitiveness in your market?*

NNT Change Tracker Gen7 - Real-Time, continuous FIM...and Certified by the Center for Internet Security

- Change Tracker Gen7 has been certified by the CIS which means you can trust NNT to accurately deliver the most comprehensive, consensus-derived hardening checklists.
- NNT provide CIS Benchmark checklist coverage for all Windows, Unix and Linux Operating Systems, SQL Server and Oracle Database Systems, and for Network Devices and appliances such as Cisco ASA firewalls.
- Compliance is continuously enforced meaning vulnerabilities are highlighted more quickly than with traditional vulnerability scanners.
- Better still, NNT Change Tracker Gen7 provides continuous real-time FIM across all system, application, driver and configuration files providing peace of mind that system integrity is being maintained.
- And if the worst case scenario does happen and your systems are breached or infected with malware, this will be detected within seconds, minimizing damage and costs.

TO REQUEST A FREE TRIAL OR DISCUSS ANY AREA COVERED IN THIS WHITEPAPER, PLEASE CONTACT US AT info@nntws.com